

ElectricFlow 6.4 User Guide

Electric Cloud, Inc. www.electric-cloud.com

Copyright © 2002 – 2016 Electric Cloud, Inc. All rights reserved.

Published 5/24/2016

Electric Cloud® believes the information in this publication is accurate as of its publication date. The information is subject to change without notice and does not represent a commitment from the vendor.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." ELECTRIC CLOUD, INCORPORATED MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any ELECTRIC CLOUD software described in this publication requires an applicable software license.

Copyright protection includes all forms and matters of copyrightable material and information now allowed by statutory or judicial law or hereinafter granted, including without limitation, material generated from software programs displayed on the screen such as icons, screen display appearance, and so on.

The software and/or databases described in this document are furnished under a license agreement or nondisclosure agreement. The software and/or databases may be used or copied only in accordance with terms of the agreement. It is against the law to copy the software on any medium except as specifically allowed in the license or nondisclosure agreement.

Trademarks

Electric Cloud, ElectricAccelerator, ElectricCommander, ElectricInsight, and Electric Make are registered trademarks or trademarks of Electric Cloud, Incorporated.

Electric Cloud products—ElectricAccelerator, ElectricCommander, ElectricInsight, and Electric Make—are commonly referred to by their "short names"—Accelerator, Commander, Insight, and eMake—throughout various types of Electric Cloud product-specific documentation.

Other product names mentioned in this guide may be trademarks or registered trademarks of their respective owners and are hereby acknowledged.

Contents

Introduction to ElectricFlow	
Web-Based System	1
Automation Platform	34
What Makes ElectricFlow Unique?	34
ElectricFlow Architecture	35
Simple Architectural Overview	35
Expanded Remote Configuration	36
Other Configurations	37
Roadmap to ElectricFlow	37
Deployment Automation	39
Applications and Processes	39
Environments	42
Application Deployment Options	43
Master Components	46
Master Component Examples	47
Master Components List UI	68
Artifact Staging	70
At the Application Level	70
At the Pipeline and Release Levels	71
Rollback	71
Automatic Rollback	71
Manual Tasks and Steps	72
Process Branching	73
Using Process Branching	74
Process Branching States and Conditions	75
Custom Conditions in Process Branching	76
Snapshots	77
Example: Creating, Deploying, and Comparing Snapshots	
Configuration Drift	84
Deployment Examples	85
Guidelines for Modeling and Deploying Applications in ElectricFlow	85

When Modeling Applications	85
When Modeling Environments	86
When Deploying Applications	86
Getting the Real-Time Status of Application Runs and Troubleshooting	88
Example: Modeling and Deploying Applications	88
Creating a New Application	89
Authoring Component Processes	92
Authoring Application Processes	99
Creating Environments	108
Defining Tier Maps	112
Deploying and Troubleshooting Applications	113
Example: Manual Steps with Runtime Parameters	123
Process Flow Details	
Deployment Scenarios	
Setting Up Custom Parameters for Application and Component Processes	
Setting and Modifying the Parameter Label	
Looking Up Parameters in Application and Component Processes	
Attaching Credentials to Application and Component Processes	
Deploying Applications in Dynamic Environments	
About Resources	
Examples Modeling Dynamic Environments	
Usage Guidelines and Best Practices	144
Creating AMIs	147
AMIs with ElectricFlow Agents	147
AMIs with ElectricFlow Agents and Chef Configuration Management	147
Creating Resource Templates	148
Example: Resource Template for Infrastructure using Amazon EC2	152
Select Cloud Provider	152
Select account	158
Example: Resource Template for Infrastructure using OpenStack	159
Prerequisites and Configuration Guidelines	159
Select cloud provider	159
Select account	163
Example: Resource Template for Middleware Using Chef	164

Select configuration management	164
Select account	166
Viewing and Editing Resource Templates	166
Resource Template Details	167
Resource Template Properties	168
Access Control Settings	170
Change Tracking	171
Accessing the Resource Templates in the Automation Platform	171
Creating Environment Templates	173
Creating a New Environment Template	173
Creating an Environment Template Based on an Existing Template	182
Viewing and Editing Environment Templates	185
Deploying Applications With Provisioned Cloud Resources	188
Retiring Dynamic Environments	196
Dynamic Environment Example with Amazon and Chef	197
Resource Templates	197
Environment Templates	199
Deploying Applications to Dynamic Environments	200
Retiring Dynamic Environments	201
Developer Task: Creating Custom Plugins	202
Creating Custom Cloud Provider Plugins	203
How to Create a Custom Cloud Provider Plugin	203
Example: Property Structure for a Cloud Provider Plugin	206
Creating Custom Configuration Management Plugins	208
How to Create a Custom Configuration Management Plugin	208
Example: Property Structure for a Configuration Management Plugin	211
Inventory Tracking	212
Tracking at the Component Process Level	213
Environment Inventory	213
Pipelines	217
Pipeline Concepts	218
Tasks	218

Entry and Exit Gates	220
Pipeline UI	220
Pipeline Stage Summary	227
Creating User-Generated Data for the Stage Summary	227
Viewing the Stage Summary During a Pipeline Run	228
Viewing the Stage Summary for a Completed Pipeline Run	230
Credentials in Pipelines	231
Example: Authoring a Pipeline with Manual and Utility Tasks	232
Example: Authoring and Running Pipelines	236
Creating a Pipeline	239
Editing the Pipeline Details	244
Defining the Stage	246
Viewing and Editing the Pipeline Definition	251
Defining Gate Approvals	252
Running Pipelines	255
Examples: Viewing the Details of a Pipeline Run	257
Viewing Pipeline Runs	259
Opening the Pipeline Runs List	259
What's in the Pipeline Runs List	259
Searching the Pipeline Runs List	262
Troubleshooting Pipelines	266
Examples: Task List in a Stage	272
Change the Order in Which the Tasks Are Executed	272
Sample Task List for a Stage	272
Release Management	274
Release Concepts	
Modeling Releases	
Setting Up Releases	
Controlling the Releases	
Visibility and Status of Release Pipelines	
From the Release Dashboard	
From the Environment Inventory	
From the Change History	
Postp	
Example: Visibility in a Release Pipeline from the Release Dashboard	
Release Definition	
Example: Completed Release Definition	
Example: Traditional Multiple Application Release	

	Example: Continuous Delivery Style Releases for One Application	294
Re	elease Dashboard	300
	Accessing the Release Dashboard	300
	Release Dashboard	300
Pa	ath-to-Production View	305
Re	elease Summary	308
Rı	unning and Completing Releases	309
	Running the Release	309
	Aborting a Release Run	
	Completing the Release	312
Secu	rity	314
Ac	ccess Control	314
Pr	oject-Level Security	314
Cr	redentials and User Impersonation	
	What is a Credential?	
	Defining a Credential	
	Why Use Credentials?	
	Credential Access Control	
	Using Credentials for Impersonation at the Platform Level	317
	Setting the Impersonation Credential in the Automation Platform	317
	Attaching Impersonation Credentials to Steps, Procedures, and Projects	317
	Attaching Impersonation Credentials to Schedules	318
	Attaching Impersonation Credentials to Jobs	318
	Attaching a New Impersonation Credential	318
	Accessing Credentials From a Step at the Platform Level	319
	Attaching a Credential to a Procedure Step, Procedure, or Schedule	319
	Passing Credentials as Parameters	320
	Credential References	320
	Best Practices for Retrieving Credentials	321
	First Class	321
	Second Class	321
	Does the step's command use properties?	321
	Does the step execute code or commands that can be modified?	321
	Avoid passing passwords on the command line	322

Credentials and Impersonation in Application and Component Processes	322
Use Case: Attaching Credentials in Deployment Automation	324
Example	325
Credentials in Pipelines	326
Credentials in Releases	326
Change Tracking	327
Performance Consequences of Change Tracking	327
Use Case 1: Importing and Cloning Large Projects with Change Tracking Enabled	328
Use Case 2: Change Tracking of Non-Project-Owned Objects	329
Estimating Database Growth	330
Common Usage Patterns	330
Estimating Process	330
Alternative Estimating Methods	332
Best Practices for Change Tracking	332
When to Enable or Disable Change Tracking for a Project	332
Splitting Tracked Objects Into Separate Projects	333
Configuring Change Tracking	334
Enabling Change Tracking Globally	334
Enabling Change Tracking on a Per-Project Basis	334
Upgrading to ElectricFlow 6.x	
Customizing the Change History Page	336
Searching the Change History	
Viewing the Change History	
Viewing the Change History from the Applications List	
Viewing the Change History From an Application or Environment	341
Application Visual Editor	341
Component Process Visual Editor	344
Environments Visual Editor	346
Environment Tier	347
Viewing the Change History for Artifacts, Jobs, Projects, and Workflows	347
Artifacts	348
Jobs	349
Projects	350
Workflows	350
Change History Page	350

Time Line	353
Default Settings	353
Selecting the Time Range	354
Moving the Start and End Times	355
Specifying the Time Range	356
Number of Changes	
Paths to Objects	
Detailed Object Changes	
Filters	
Modifying What You See the Change History	
Change History Time Line	
Change History Filters	
Reverting Changes to a Tracked Object and Its Tracked Contents	370
Notifications	373
Application Processes and Process Steps	373
Jobs, Workflow States, and Procedures	373
Pipelines	374
Configuring Email Notifications in Application Processes	375
Selecting and Editing Email Messages for Application Processes	381
Automation Platform	387
Build-Test Automation	387
Getting Started	390
Overview	390
Automation Platform Terminology	390
Getting Started Scenarios Help Topic Series	391
Navigating the Automation Platform User Interface	391
In addition to the main set of tabs, the top bar includes:	391
Home Tab	392
Continuous Integration Subtab	393
Projects, Jobs, and Workflow Tabs	394
Cloud Tab	395
Resources Page	395
Pools Subtab	
Workspaces Subtab	
•	

	Zones and Gateways Subtabs	. 396
	Artifacts and Search Tabs	.396
	Artifact Versions Subtab	.396
	Repositories Subtab	.397
	Search Tab	.397
	Administration Tab	. 397
	Getting Started - Scenario 1 - Creating a Simple Procedure	.398
	Overview	. 399
	Begin Scenario 1	.399
	Scenario Extension - Adding Another Step	.405
	Finding an ElectricFlow Automation Platform Web Page	.406
	Summary	. 407
	Getting Started - Scenario 2 - Creating a Procedure That Uses an SCM	. 407
	Overview	. 409
	Begin Scenario 2	.409
	Scenario extension - Add a step to show workspace contents	.417
	Summary	.418
	Getting Started - Scenario 3 - Notification, Scheduling, and Reporting	.418
	Overview	. 420
	Begin Scenario 3	.420
	Scenario extension - Add a thumbnail report to your Home page	. 430
	Summary	. 432
	Getting Started - Scenario 4 - Multi-agent Build and Test	. 432
	For Linux	.433
	For Windows	.433
	Overview	. 433
	Begin Scenario 4	.434
	Scenario extension - Using postp	.439
	Summary	. 441
Au	tomation Platform Setup	.441

Automation Platform Home Page	441
Overview	442
Job Configurations	442
Create Job Configurations three ways	442
Shortcuts	443
Create Shortcuts two ways	443
Jobs Quick View	443
Create a job category	443
Reports	443
Customizing the ElectricFlow Platform UI	444
Customizing parameters	444
How do you customize parameters?	444
Custom parameter form contents	444
Example	445
Customizing the tab layout	446
Overview	446
View definition syntax	446
Storing views	449
Default views	449
Developing and troubleshooting	449
Home page configuration	450
User settings	450
Location	450
Installing the Home page	450
Reconfigure Contact Support link	450
Configuring ElectricFlow	451
Web Interface Online Help System	452
Setting Up Resources	452
What is a Resource?	452
What is a Pool?	453

ElectricFlow Supports Two Resource Categories	453
Standard Resources	453
Proxy Resources	454
Setting up SSH keys	454
To create a new resource	454
Setting Up Workspaces	455
To create a new workspace	455
Setting Up Email Configurations	455
Creating an Email Configuration	455
Setting Up a Source Control Configuration	456
Setting Up Directory Providers	457
To specify a directory provider	457
Enable/Disable Local ElectricFlow Users	457
wrapper.conf Properties	457
Using ElectricFlow in Your Environment	458
What's in a step?	458
Where's the script for a step?	459
Two ways to handle script execution	459
Two cases where it makes sense to store the script for a step in the ElectricFlow step	459
Environment variables	460
When do you use subprocedures?	460
How do you evolve procedures?	460
Porting existing scripts	461
ElectricFlow project version control	462
ElectricFlow Installed Tools	462
Usage	463
Commands	463
Server Communication Options	466
Global Options	466
Examples	466

	Example 1: Configure an agent to talk to any server (untrusted mode)	466
	Example 2: Configure an agent to accept connections only from a single remote ElectricFlow server	466
	Example 3: Configure an ElectricFlow server with additional host names in the certificate	
	Usage	
	ElectricFlow agent configuration options	
	Apache web server configuration options	
	ElectricFlow Server configuration options	
	Repository Server configuration options	
	General options	
Evo		
⊏ха	mples	
	Use Cases	
	Using ecdaemon	
	Command-line parsing	
	ec-perl considerations	
	roxy Algorithm	
есрі	roxy Operations	
	getDefaultWorkingDirectory	
	getDefaultTargetPort	478
	connect	478
	uploadFile	479
	generateWrapperScript	479
	uploadWrapperScript	480
	generateWrapperInvocationCommand	481
	runCommand	481
	cleanup	482
	ping	482
Ava	ilable Helper Functions	483
	mesg	483
	readFile	483
	writeFile	483

	484
setOperation	484
	484
loadFile	484
setSSHKeyFiles	485
setSSHUser	485
useMultipleSSHSessions	486
Examples	486
"Real World" Examples	487
ClusterExec	487
MySQL	487
Android	488
Setting up the process	490
Creating a Setup Step	491
Update the Postprocessor field for steps in your procedure	491
Add a Final Step to your procedure	491
Copying Other Files from the Workspace	492
Prerequisites	493
Command	493
Import Files	493
1	
Prerequisites	
	494
Prerequisites	494
Prerequisites	
Prerequisites Locations Command	
Prerequisites Locations Command Output	
Prerequisites Locations Command Output Automation Platform Tasks	
Prerequisites Locations Command Output Automation Platform Tasks Access Control	
Prerequisites Locations Command Output Automation Platform Tasks Access Control Reading and Using This Page	

Artifact Details	499
General Information section	499
The "tabbed" sections	500
Artifact Versions table	500
Properties table	500
Artifact - create new or edit existing artifact	500
To create a new artifact	501
To edit an artifact	501
Artifact Versions	501
Artifact Version Details	502
General Information section	503
The "tabbed" sections	503
Files	503
Retrievers table	503
Dependent Artifact Versions table	504
Properties table	504
Artifact Version - edit an existing artifact version	504
To edit an artifact version	504
Repositories	505
Repository - create new or edit existing repository	506
To create a new repository	506
To edit a repository	507
Database Configuration	
Defect Tracking Configurations	508
Link at the top of the table	508
Column descriptions	508
Defect Tracking - create new or edit existing configuration	509
To create a defect tracking configuration	509
Using the defect tracking integration	509
To edit an existing defect tracking configuration	510
Defect Tracking Reports	510

Email Configurations	510
Email Configuration - create new or edit existing email configuration	510
To create a new email configuration	511
To edit an existing email configuration	511
Email Notifier - create new or edit existing email notifier	511
To create a new email notifier	512
To edit an existing email notifier	514
Email notifier templates	514
Email Notifier Template - Html_JobStepTempl_AllSteps.txt	516
Email Notifier Template - Html_JobStepTempl_SingleStep.txt	517
Email Notifier Template - Html_JobTempl.txt	518
Email Notifier Template - Html_StateTemplate_ApproveWorkflow.txt	519
Email Notifier Template - Html_StateTemplate_FullWorkflow.txt	520
Email Notifier Template - JobStepTempl_FullProps.txt	521
Email Notifier Template - JobTempl_FullProps.txt	522
Email Notifier Template - StateTemplate_FullPropertyPaths.txt	523
Event Log	524
Column descriptions	525
Configuring the Event Log	525
Automation Platform Home Page	525
Overview	526
Job Configurations	526
Create Job Configurations three ways	526
Shortcuts	526
Create Shortcuts two ways	527
Jobs Quick View	527
Create a job category	527
Reports	527
Job Configuration	528
To create a new job configuration	528

To edit an existing job configuration	529
Shortcuts	529
Jobs Quick View	529
To create a new jobs quick view category	530
To edit an existing jobs quick view category	530
New Report	530
To populate the drop-down menu	531
Jobs	531
Links at the top of the table	531
Column descriptions	532
Tips	532
Job Step Details	533
Links and actions at the top of the table	533
Summary section (at the top of the page)	533
The "tabbed" Sections	534
Child Steps table	534
General table	535
Diagnostics table	535
Parameters table	535
Properties table	536
Notifiers table	536
Job Step Time and waitTime Properties Explained	536
Job Details	538
Links and actions at the top of the table	538
Summary section (at the top of the page)	538
The "tabbed" sections	540
Steps table	540
Diagnostics table	541
Parameters table	542
Properties table	542

Notifiers table	543
Published Artifact Versions table	543
Column descriptions	543
Retrieved Artifact Versions table	543
Column descriptions	543
Licenses	544
All Licenses	544
Column descriptions	544
Current Usage	545
License Types	546
Concurrent Resource License	546
Concurrent User License	548
Concurrent Step License	550
Registered Host License	550
View License	551
Import License	552
Plugin Manager	553
Tab and field descriptions above the table	553
Column descriptions	554
Using a plugin	555
Configuring the plugins directory	556
Procedure Details	556
Procedure Steps	556
Creating a new step	557
Parameters	557
Email Notifiers	558
Custom Procedure Properties	558
Procedure - create new or edit existing procedure	558
To create a new procedure	558
To edit an existing procedure	560

Step - create new or edit existing step	560
To create a new command or subprocedure step	560
To edit an existing command or subprocedure step	566
Publish Artifact Version Step	567
Retrieve Artifact Version Step	567
Artifact Retrieval Dependency Order Explained	569
Send Email Step	570
New Extract Preflight Sources Step	571
Parameter - create new or edit existing parameter	571
To create a new parameter	572
To edit an existing parameter	574
Projects	574
Projects have two purposes	575
ElectricFlow includes default projects	575
ElectricFlow Solutions	576
Project - Create New or Edit Existing Project	576
Creating a New Project	577
Editing an Existing Project	578
Project Details	578
The "tabbed" sections	579
Procedures tab	579
Workflow Definitions tab	579
Jobs tab	580
Workflows tab	581
Schedules tab	581
Credentials tab	582
Properties tab	582
Reports tab	582
Run Procedure	583
Using the Run Procedure page	583

Schedule - create new or edit existing schedule	584
To create a new standard schedule	584
To create a new continuous integration "schedule"	588
To edit an existing schedule	588
Credentials - create new or modify existing credential	588
To create a new credential	588
To edit a credential	589
Property - create new or edit existing property	589
Nested Property Sheet	590
Reports	590
Create a new report	591
Multiple Series Report	591
Category Report	592
Count Over Time Report	593
Procedure Usage Report	594
Resources	595
Supported Resource Categories	596
Resource Page Information and Functions	597
Links, Icons, and Buttons Above the Table	597
Table View — Column Descriptions	599
Grid View	602
Filters pane	603
New Resource Panel	604
New Proxy Resource Panel	608
Edit Resource Panel	
Edit Proxy Resource Panel	612
Resource Details panel	612
Switching a Non-Trusted Agent to Trusted	613
Manually switching a non-trusted agent to trusted status	613
Install or Upgrade Remote Agents	613

Prerequisites:	614
Using the Install or Upgrade Remote Agents Dialog Box	614
Parameters	616
Advanced	617
Resource Pools	621
Resource Pool – Create a New Pool or Edit Existing Pool	622
Creating a new resource pool	622
Editing an existing resource pool	624
Unusable resources	624
Resource Pool Details	625
The "tabbed" sections	625
Properties tab	626
Zones	627
Zone Details panel	. 628
Creating a new zone	628
Access Control notes	629
Gateways	629
Gateway Details panel	631
Create Gateway panel	631
Edit Gateway panel	632
Access Control note	632
Accessing the Resource Templates in the Automation Platform	632
Define Search	634
Search Results	635
Server	. 636
Settings - edit existing property settings	636
Source Control Configurations	637
Source Control Configurations - create new or edit existing configuration	. 638
AccuRev	. 638
ClearCase	639

File	639
Git	639
Perforce	639
Property	640
Subversion	640
Checking out code from a source control system	640
To edit an existing source control configuration	641
Active Users	641
Users and Groups	641
User - create new or edit existing local user	643
To create a new local user	643
To edit an existing local user	643
User Details	644
Edit User Settings	644
Groups	645
Group - create new or edit existing local group	645
To create a new local group	645
To edit a local group	645
Group Details	645
Directory Providers	646
Directory Providers - create new or edit existing directory providers	647
ElectricFlow supports Active Directory and LDAP directory providers	647
To create a new Active Directory provider	647
To create a new LDAP directory provider	650
Examples for directory provider field descriptions	653
To edit an existing directory provider	655
Test Directory Provider	655
Workflows	656
Workflow Definition - create new or edit existing workflow definition	656
To create a new workflow definition	656

	To edit an existing workflow definition	657
	Workflow Definition Details	. 657
	Graph view	659
	To create a new state	660
	To configure the new State Definition	661
	To create a Transition Definition	665
	Using the graph's "quick-access" features	. 667
	Show Legend	668
	List view	669
	Properties view	671
	Run Workflow	. 672
	Workflow Details	. 673
	Summary section - at the top of a running workflow page	. 674
	When the workflow is complete	675
	Graph tab	676
	Show Legend tab	676
	Using the graph's "quick-access" features	677
	Show History	677
	States tab	679
	State Details panel	. 680
	Parameters tab	. 683
	Properties tab	683
	Workflow Log	684
	Transition Workflow	685
	Workspaces	686
	Workspace - create new or modify existing workspace	. 687
	To create a new workspace	. 687
	To edit an existing workspace	689
	Workspace File	689
Ele	ctricFlow Platform Objects and Functionality	689

Access Control	690
Overview	690
Privileges	690
Users and Groups	691
Special Users and Groups	691
Access Control Lists - allow and deny	691
nheritance	692
Additional Information about "deny"	692
System Objects	693
Access Control and Jobs	694
Examples for Increased Security	695
Abbreviations used in the examples	695
Example 1 - Basic ACL Setup	695
Server ACLs	696
Custom Server Properties ACLs	696
System Object ACLs	696
Plugin Project ACLs	698
Project ACLs	698
Example 2 - Team ACL Setup	699
Server ACLs	700
Custom Server Properties ACLs	700
System Object ACLs	700
Plugin Project ACLs	702
ACLs for Existing Projects	702
ACLs for New Projects	703
Optional: Restricting Resources and Workspaces by Team	705
Resource ACLs	705
Workspace ACLs	706
Artifact Management	707
Overview	707

Artifact Objects	707
Artifact	708
Examples	708
Artifact Version	708
Version Strings	709
The artifactVersionState property	710
Repository	711
Using multiple repositories	713
Configuring the Artifact Repository Server for Amazon S3	713
The respositoryDisabled intrinsic property	715
Access Control	715
Publishing Artifact Versions	717
Enable Compression	718
nclude and Exclude Patterns	718
Retrieving Artifact Versions	719
Entering Filters During a Retrieve Operation	721
Dependent Artifact Versions	722
Artifact Cache	723
Cleaning Up Repositories and Caches	725
Authenticating Users for LDAP and Active Directory	725
Configuring LDAP	725
Sample LDAP Configuration	726
The following properties configure LDAP mapping:	726
Determining LDAP Mapping	727
Sample LDAP User Record	728
Sample LDAP Group Record	729
Sample Active Directory Configuration File	729
Change Tracking	730
Configuring Change Tracking	731
Enabling Change Tracking Globally	731

Enabling Change Tracking on a Per-Project Basis	731
In the ElectricFlow Platform UI	731
Using ectool	732
Using ec-perl:	732
Upgrading to ElectricFlow 6.x	733
Customizing the Change History Page	733
Viewing the Change History for Artifacts, Jobs, Projects, and Workflows	733
Artifacts	734
Jobs	734
Projects	735
Workflows	736
Modifying What You See the Change History	736
Change History Time Line	736
Default Settings	737
Number of Changes	737
Time Increment	738
Moving the Start and End Times Manually	739
Setting Custom Time Increments	739
Change History Filters	742
Searching the Change History	743
Reverting Changes to a Tracked Object and Its Tracked Contents	746
Change History Search Form	749
Change History Page	750
Paths to Objects	752
Detailed Object Changes	753
Credentials and User Impersonation	753
What is a Credential?	754
Defining a Credential	754
Why Use Credentials?	754
Credential Access Control	755

Using Credentials for Impersonation at the Platform Level	755
Setting the Impersonation Credential in the Automation Platform	756
Attaching Impersonation Credentials to Steps, Procedures, and Projects	756
Attaching Impersonation Credentials to Schedules	756
Attaching Impersonation Credentials to Jobs	756
Attaching a New Impersonation Credential	757
Accessing Credentials From a Step at the Platform Level	757
Attaching a Credential to a Procedure Step, Procedure, or Schedule	758
Passing Credentials as Parameters	758
Credential References	758
Best Practices for Retrieving Credentials	759
First Class	759
Second Class	759
Does the step's command use properties?	759
Does the step execute code or commands that can be modified?	760
Avoid passing passwords on the command line	760
Defect Tracking	760
Scenario example	761
Example: Enabling the JIRA integration in your procedure	761
ElectricSentry	762
How ElectricSentry works	763
Configuring ElectricSentry	763
Quiet time	763
Resource	764
Polling frequency	764
Time-of-day and day-of-week	764
Configuring a build for continuous integration	764
Source Control Management SCM) configurations	764
Create a procedure	765
Create a schedule	765

Optional - running ElectricSentry on multiple resources	765
Overview	765
Configuring ElectricSentry to use multiple resources	766
The Job Step Execution Environment	766
Terms and definitions	767
Preflight Builds	769
Why use Preflight Builds?	769
Preflight Build Solution	769
Preflights with ElectricFlow	770
Workflow	770
Components	770
Installation	770
Configuration	771
Server	771
Agent	772
Client	772
Samples	772
Default SCMs	777
Perforce	778
Subversion	779
AccuRev	779
ClearCase	779
Bazaar	780
Git	780
cvs	781
Mercurial	781
Repo	781
StarTeam	782
TFS	782
Vault	782

Troubleshooting	783
Client	783
Agent	783
Properties	783
Overview	783
Creating or modifying properties	785
Using property values	785
Property sheets and intrinsic properties	786
Property names and paths	786
Absolute property paths	787
Relative property paths	789
jobSteps Paths	789
Context-relative "shortcuts" to property paths	790
Property path shortcuts in ElectricFlow 5.0 and later	794
Property name substitutions	796
Expandable properties	796
Custom property names and values	797
The property hierarchy	798
Special property references	799
increment	799
timestamp	799
javascript	800
ElectricFlow Intrinsic Properties listed by object type	801
ElectricFlow Property Type definitions	802
ElectricFlow Object / Property tables	803
Property error codes	865
Postprocessors: Collecting Data for Reports	867
Overview	867
Postp	868
Extending postp: matchers	868

Postp functions	870
Integration with the ElectricFlow user interface	871
Custom property names and values	872
Diagnostic information	872
Postp integration with Java Tools	873
The process	873
From this output	874
Java Tool matcher examples	874
Two examples of postp Emma matchers:	874
An example of postp JUnit matchers:	875
An example of postp Clover matchers:	875
Artifacts directory	875
Reports	876
Run reports on a non-local resource	876
Real-time reports	877
Home page	877
Jobs	877
Job details / Step details	877
Resources	877
Build reports using ecreport	877
Example 1	879
Example 2	880
Default Batch Reports (Run Reports)	880
Report types	880
Default Batch Reports	881
Тір	882
Advanced reporting information - ecrptdata	883
How are default batch reports generated?	883
Filtering	884
Adding additional data columns to the report	885

Optional Batch Reports	
Creating a Report Job	886
Viewing the report	888
Optional Batch report examples	
Category sample report	
Procedure Usage Sample Report	890
Count Over Time Sample Report	
Multiple Series Report	894
Creating Custom reports	897
Data extraction	
Using ecextract.pl for data extraction	
Examples	901
BIRT Report Designer	904
Understanding additional report components	s905
Packaging reports for ElectricFlow	905
Helper functions provided in ElectricComma	nder::ReportUtils.pm907
Custom Report Examples	
	ding a "banner" heading909
Example 1: modifying an existing report - ad	
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report	ding a "banner" heading909
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report	909 ding a "banner" heading
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report To modify the copied report design Test your new report	909 ding a "banner" heading
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report To modify the copied report design Test your new report Example 2: complete end-to-end example .	909 ding a "banner" heading
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report To modify the copied report design Test your new report Example 2: complete end-to-end example Planning this report example	909
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report To modify the copied report design Test your new report Example 2: complete end-to-end example Planning this report example Creating a new BIRT rptdesign file	909
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report To modify the copied report design Test your new report Example 2: complete end-to-end example Planning this report example Creating a new BIRT rptdesign file Deploy your custom report	909 909 909 909 910 913 913 914
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report To modify the copied report design Test your new report Example 2: complete end-to-end example Planning this report example Creating a new BIRT rptdesign file Deploy your custom report Test your new report	909 909 909 909 910 910 913 913 914 924
Example 1: modifying an existing report - ad To copy an existing ElectricFlow report To modify the copied report design Test your new report Example 2: complete end-to-end example Planning this report example Creating a new BIRT rptdesign file Deploy your custom report Test your new report System Health Monitoring	ding a "banner" heading

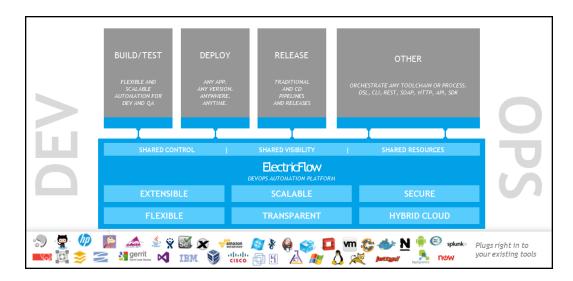
Basic workflow concepts	928
Workflow objects	. 929
Workflow Definition	.929
State Definition	. 929
Transition Definition	. 930
Workflow	.930
State	. 930
Transition	. 931
Access Control	. 931
Property Search Paths	. 931
Parameters	.932
Sending Notifications	. 933
Workflow Logs	933
Visualizing a Workflow	.933
Building a Workflow - a Tutorial	.933
Best Practice tip	934
A summary of the component sections to build our workflow:	. 934
To begin calling a job from within a workflow	.935
Collecting a parameter when a workflow is launched and passing its value to a job	.936
Retrying a state	.937
Automatically transitioning to different states based on the outcome of a job	.938
Invoking another workflow	941
Invoking another workflow	
	.942
Running jobs in parallel	.942 .943
Running jobs in parallel	.942 .943 .946
Running jobs in parallel Automatically transitioning to different states based on the outcome of jobs in another workflow Waiting for manual intervention	.942 .943 .946
Running jobs in parallel Automatically transitioning to different states based on the outcome of jobs in another workflow Waiting for manual intervention Restricting who can take a manual transition	.942 .943 .946 .950
Running jobs in parallel Automatically transitioning to different states based on the outcome of jobs in another workflow Waiting for manual intervention Restricting who can take a manual transition Sending email notifiers	.942 .943 .946 .950 .953

	Another workflow example	961
	Workspaces and Disk Space Management	. 961
	Overview	. 962
	Defining Workspaces	. 962
	Using Workspaces	. 963
	Workspace Directory Names	.963
	Working Directories	964
	Workspace Accessibility	. 964
	Local Workspaces (Disconnected Workspaces)	.964
	Access control	. 965
	Impersonation and workspaces	.965
	ElectricFlow managed files	. 966
	Disk space management	967
Tu	torials	. 967
	Adding a Link to a Job	. 967
	Implementation	. 968
	Related Information	969
	Calling a Subprocedure	969
	Implementation	. 969
	Related Information	970
	Checking the Outcome of Preceding Steps	970
	Implementation	. 970
	Related information	971
	Conditional Execution	971
	Implementation	. 971
	Related information	972
	Custom parameter layouts	.972
	Implementation	. 974
	Related information	974
	Email Notifications	. 974
	Implementation	. 975

Related information	976
Executing Tasks on All Resources in a Pool	976
Implementation	976
Related information	977
Factory Procedures	977
Implementation	978
Postp extension	979
Implementation	980
Related information	980
Publishing and Retrieving an Artifact	981
Implementation	982
Related information	982
Reserving a Resource for Job Step Duration	982
Implementation	983
Related information	983
Running Steps in Parallel	983
Implementation	984
Related Information	985
Step Timeouts and Steps that Always Run	985
Implementation	986
Storing and Retrieving Properties in a Job	986
Implementation	987
Related information	987
Working with Properties Stored in a Procedure	987
Implementation	988
Related Information	989
Glossary	989

Introduction to ElectricFlow

ElectricFlow™ is an enterprise-grade DevOps Release Automation platform that simplifies provisioning, build and release of multi-tiered applications. Our model-driven approach to managing environments and applications allows teams to coordinate multiple pipelines and releases across hybrid infrastructure in an efficient, predictable and auditable way.



Web-Based System

At its core, ElectricFlow automation platform is a web-based system for automating and managing the build, test, deployment, and release process. It provides a scalable solution, solving some of the biggest challenges of managing these "back end" software development tasks, including these challenges:

- Time wasted on script-intensive, manual, home-grown systems that
 - o Are error prone
 - Do not scale well
 - o Have little or no management visibility or reporting
- Multiple, disconnected build and test systems across locations, resulting in:
 - Redundant work
 - o Inability to share or reuse code files across teams
 - o Hard to manage build and test data
- Slow overall build and release cycles that directly impact:

- Release predictability
- Time-to-market

Automation Platform

The automation platform has a three-tier architecture, AJAX-powered web interface, and first-of-its-kind build and release analytic capabilities for reporting and compliance. With this solution, your developers, release engineers, build managers, QA teams, and managers gain:

- · Shared platform for disseminating best practices and reusing common procedures
- · Ability to support geographically distributed teams
- · Continuous integration and greater agility
- Faster throughput and more efficient hardware utilization
- · Visibility and reporting for better project predictability
- · Better software quality by integrating and validating against all target platforms and configurations

For examples of ElectricFlow architecture configurations, see ElectricFlow Architecture on page 35.

What Makes ElectricFlow Unique?

ElectricFlow provides enterprise-class speed and scalability for software build and release management. It is easy to install and use on a simple build, yet scales to support the largest and most complex build and test processes. ElectricFlow distributes jobs in parallel across multiple resources for faster overall cycle time.

ElectricFlow supports multiple teams, working in multiple locations, programming in multiple languages in an environment that can be centrally controlled and managed. Shared assets and reuse make individual teams more efficient by eliminating duplicate work, and gives organizations the power to deploy cross-company standards.

ElectricFlow's unique analytics provide visibility into one of the best indicators of project success: compiled, tested, working code. ElectricFlow's analytics database stores all build and test information for real-time and trend reporting giving your organization the power to collect pinpoint statistics and to gain visibility into important productivity metrics such as trends in error rates. Additionally, out-of-the-box reports provide information about cross-project status and build trends by project and resource utilization. ElectricFlow's integration with virtual lab automation (VLA) solutions also allows you to snapshot or reproduce a specific build for auditing or troubleshooting purposes.

ElectricFlow provides unified process automation across the entire build-test-deploy life cycle and across heterogeneous tools via integrations with leading ALM tools. Integrations with SCM tools enable continuous integration, triggering builds whenever code is checked into the specified repository/branch. When used with VMware Lab Manager, ElectricFlow can dynamically provision either physical or virtual resources without manual intervention. This feature delivers efficient, dynamic resource provisioning and reduces development and QA dependence on IT operations.

ElectricFlow Architecture

ElectricFlow was designed to support small, mid-range, or enterprise scale software production. Based on a three-tier architecture, ElectricFlow scales to handle complex environments. The ElectricFlow multi-threaded Java server provides efficient synchronization even under high job volume.

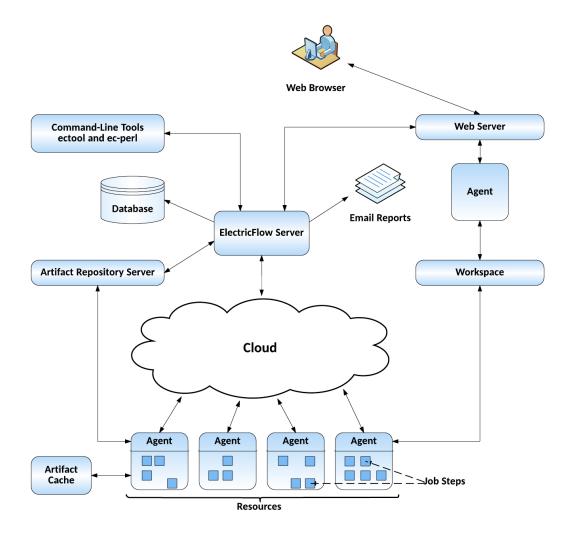
- The ElectricFlow server manages resources, issues commands, and generates reports.
- An underlying database stores commands, metadata, and log files.
- Agents execute commands, monitor status, and collect results in parallel across a cluster of servers for rapid throughput.

Simple Architectural Overview

This local configuration applies to all the use cases. The ElectricFlow server, web server, artifact cache, Artifact Repository server, workspace, command-line tools, resources, agents, and job steps are all in the automation platform.

In this local configuration:

- The ElectricFlow server manages resources, issues commands, and generates reports.
- · Resources, agents, and databases are managed in the automation platform.
- An underlying database stores commands, metadata, and log files.
- Procedures, which include job steps, are defined in the automation platform.
- Job steps are executed on resources in the defined environments.
- Applications (which include procedures), components, and environments are defined for deployment automation.
- · Pipelines, stages, and tasks are defined for pipeline management.



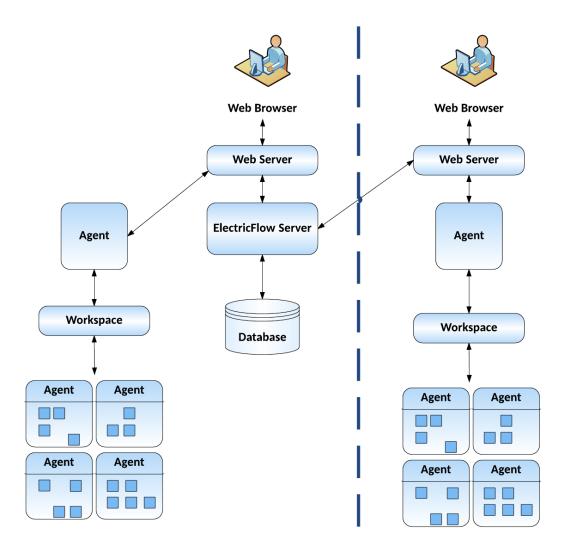
If you are only evaluating ElectricFlow, the ElectricFlow software, the database, the ElectricFlow server, the web server, and the repository server can reside on the same machine.

In a production environment, the database should reside on a separate machine from the ElectricFlow server to prevent performance issues. It is acceptable for the ElectricFlow server, web server, and repository server to reside on the same machine in a local configuration, but not required.

Expanded Remote Configuration

ElectricFlow is not limited by only the components shown in the previous configuration. This configuration applies to all the use cases.

The following shows a remote web server configuration and is an example for how you may set up a remote web server installation.



This type of remote web server configuration helps prevent network latency. If you have multiple sites, ElectricFlow can be configured to help you work more efficiently.

Other Configurations

Go to the ElectricFlow Installation Guide (http://docs.electric-cloud.com/eflow_doc/FlowIndex.html) for other architecture configurations:

- Proxy (universal) resources
- Remote database
- Multiple remote web servers
- · Multiple remote repository servers
- · Clustered configuration for horizontal scalability and high availability

Roadmap to ElectricFlow

Depending on what you want to do, go to the appropriate section.

Use Case	Definition	Go to
Build-test automation	Executing automated build and test processes, resulting in reduced costs, increase quality, reliability and traceability, and accelerated time to market. Build-test automation is a key enabler of continuous integration (CI).	Automation Platform on page 387
Deployment automation	Executing automated application deployment processes, resulting in reduced costs, increase quality, reliability and traceability, and accelerated time to market. Deployment automation is a key enabler of Application Release Automation (ARA) and continuous delivery (CD).	Deployment Automation on page 39 Deployment Examples on page 85
Pipelines	Executing automated end-to-end software development and delivery processes, resulting in reduced costs, increase quality, reliability and traceability, and accelerated time to market. Pipeline management is a key enabler of Application Release Automation (ARA) and Continuous Delivery (CD).	Pipelines on page 217
Release planning and management	The Release feature captures, executes, visualizes, and controls the life cycle of multi-application enterprise releases, where outputs from multiple teams are coordinated to produce a final release to be pushed to production. Releases allow you to execute and manage automated software release processes (including end-to-end software development and delivery processes), enabling Application Release Automation (ARA) and Continuous Delivery (CD). You can use the Release feature to deliver software releases through traditional methods and/or ARA and CD methodology. ElectricFlow manages all the releases, regardless of the release methodologies used by the release team.	Release Management on page 274
Automation platform	Where you create, configure, and manage the objects that make the build-test, deployment, and pipeline management possible.	Automation Platform on page 387

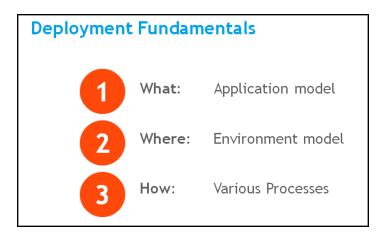
Deployment Automation

ElectricFlow™ DevOps automation accelerates development processes and simplifies operations. It accelerates the continuous delivery of software and makes software release processes more repeatable, visible, scalable, and efficient. It provides domain-specific capabilities to automate some or all phases of your software delivery process, including the build, test, integrate, deploy, and release processes. It gives distributed DevOps teams shared control and visibility into infrastructure, tool chains, and processes. It enables agility, availability, predictability, and security across many build-test, deployment, and release pipelines.

ElectricFlow uses a "model driven" approach to abstract the application architecture from the environment, infrastructure details so that the same application can be deployed to many environments using parameterized processes. ElectricFlow supports the creation, versioning, deployment, required middleware, and infrastructure details for a complete release package. You can model and version applications and environments to orchestrate the software release flow, and also to standardize and automate the deployment of applications between environments.

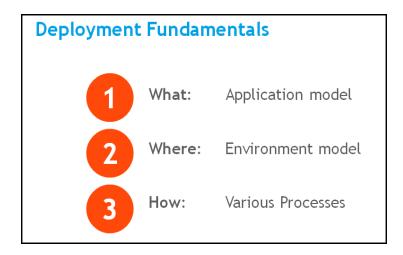
ElectricFlow Deploy makes deployments manageable, reproducible, and error-proof by modeling the application, the environments to which it will be deployed, and automating the workflow needed for the deployment. The model can be broken into these parts of your application deployment:

- What-You are deploying the application, which contains references to the files being deployed.
- Where-The application will be deployed to an ElectricFlow environment.
- How-You run processes that orchestrate deployment tasks to deploy the application to the environment.

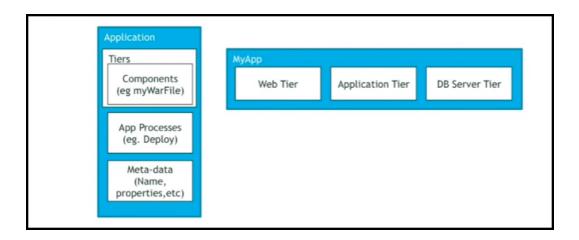


Applications and Processes

This section describes what you model (an application) and how you deploy it to an environment.

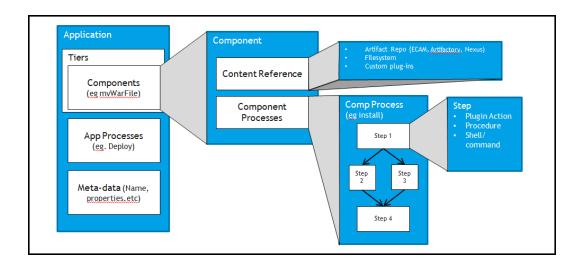


Applications are modeled by providing information about their components and defining the processes that orchestrate the deployment of those components. An application object in ElectricFlow consists of one or more application tiers. For example, MyApp is a three-tiered application with the Web, Application and database server (DB server) tiers.



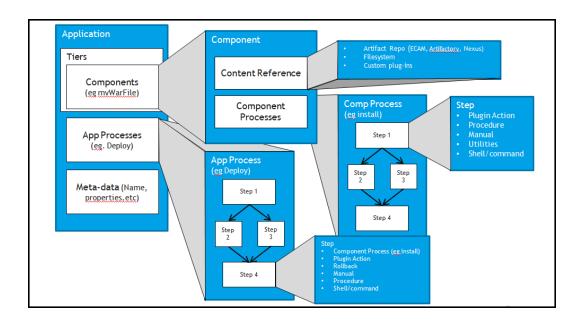
An application tier is a logical grouping of components. Each tier contains components. A component is based on an artifact and contains a reference to that artifact. This can be through the file system or an artifact repository. The application tier also contains component processes. A component process is a set of actions to be taken on that specific component when the application is deployed. The set of steps in a component process can be defined as

- · Direct commands
- Calls to ElectricFlow plugins, procedures, or utility functions
- Component operations based on the component definition
- Manual steps



Application processes are at the application (parent) level. These processes are invoked to orchestrate operations against the application. The steps in an application process can be defined as:

- Calls to component processes
- · Direct commands
- Calls to an ElectricFlow plugin, procedure, or utility function
- Manual steps
- A rollback step to roll back to a previous state or to a specific snapshot



Properties and parameters allow the same application to be deployed to many environments in a repeatable, reliable, and scalable way. They are used extensively throughout ElectricFlow and can be dynamically passed to all operations, including application deployments.

- Properties provide a flexible and powerful mechanism to manage data during all operations. They can
 be attached to any ElectricFlow object, such as applications, processes, tiers, components, and
 credentials. Go to Properties on page 783 for more information.
- Parameters provide a way to get input from users at runtime to deploy the application to a variety of
 environments. Parameters can be attached to processes, process steps, components in the application.
 Go to Deployment Examples on page 85 and Parameter create new or edit existing parameter on page
 571 for more information.

Go to Deployment Examples on page 85 for examples of how to author and deploy an application in the UI.

In an application, you can also:

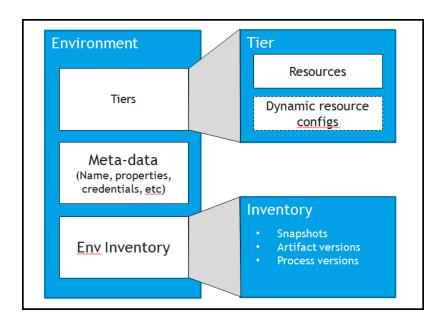
- Define process steps based on commands or scripts using the ElectricFlow DSL, ElectricFlow RESTful API, ElectricFlow Perl API, Groovy, or JRuby. See the ElectricFlow API Guide on the Electric Cloud Documentation page for the commands and code samples.
- Create a library of standardized reusable components to promote best practices across the enterprise
 and speed up application authoring with master components. Go to Master Components on page 46 for
 more information.
- Save snapshots of the application throughout the software release process. A snapshot is an immutable version of an application with specific artifact versions that can be used to optimize and troubleshoot the application. Go to Snapshots on page 77 for more information.
- Track changes for non-runtime objects such as projects, applications, environments, processes, components, artifacts, resources, and properties in the Change History. Go to Configuring Change Tracking on page 731 for more information.

Environments

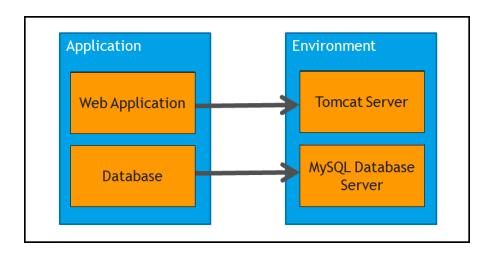
This section describes where applications are deployed.



Environments are modeled using resources that are statically defined on the system, dynamically spun up at deployment time, or a combination of both. An environment object in ElectricFlow contains one or more tiers. These tiers should map with the tiers in the application and are logical groupings of resources. Each tier defines the resources in that tier, which could be static resources in the ElectricFlow system or dynamic configurations that are spun up on deployment. An environment also contains an Environment Inventory that details the current versions of each component or snapshot deployed in that environment.



Once you have defined your application and environment, they must be connected with a tier map. This is a mapping of the application tiers to the corresponding environment tiers where the application will run. Tier maps provide a level of indirection, allowing you to freely add or remove resources from an environment without having to make any changes to the automation.



Go to Creating Environments on page 108 for more information about modeling static environments.

Go to Deploying Applications in Dynamic Environments on page 142for more information about modeling dynamic environments.

Application Deployment Options

ElectricFlow supports the following out-of-the-box options to deploy applications:

- · When authoring processes or procedures:
 - When defining a step for an application or component process, you must select the actions if a step fails:

On Error: Select Stop running to stop the process or Continue running to go to the next step in the process. The default is Continue running.

Run if: When there are conditions to be met before going to the next step, select one of the following options. The default is **all**.

- all: All conditions must be met before the step proceeds.
- any: At least one condition must be met for the step to proceed.

You can also attach parameters that affect how the step is executed at runtime. Depending on the type of step that you selected, different sets of parameters or fields are displayed.

Go to Example: Modeling and Deploying Applications on page 88 for examples of how to do define component and application process steps.

You can use process branching to specify the path through an application or component
process based on transition conditions other than out-of- the-box options. Decisions about
the next step in the flow are made while the process runs. This is similar to the transition
conditions for workflows in the automation platform (see Workflow Overview on page 927 and
Workflow Details on page 673).

Go to Process Branching on page 73 for more information about process branching states and conditions and how to set them in a process.

 When defining a step for a procedure, you must select the Error handling action if the step fails.

You can set the parameters for subprocedure steps.

You can also set preconditions and run conditions for command and subprocedure steps.

Go to Step - create new or edit existing step on page 560on how to set these options.

• When modeling environments:

Environment tiers in the environment model represent groups of targets defined by their purpose, such as web servers, databases, and artifact repositories. For example, an environment may have environment tiers for the different types of web servers called "Apache," "Tomcat," and "MS SQL", which contain all the files for your company's website in one or more locations. The resources in the Apache environment tier are the Apache servers in each location.

· When staging artifacts before the runtime:

Artifact staging retrieves artifacts that will be deployed before the deployment run starts. This ensures that all the artifacts are downloaded and available during runtime. It also reduces the application runtime. Artifact staging is enabled by default. Go to Artifact Staging on page 70 for more information.

• At runtime, select one of the following ways to deploy the application. Go to Example: Modeling and Deploying Applications on page 88 for deployment examples.

Smart deploy

This deployment method reduces risk and time by only deploying changed objects. ElectricFlow automatically detects the differences between what is about to be deployed and what has already deployed to an environment. It deploys only those objects not yet deployed to a specific resource or to resources where the object version is different. You can use this method throughout the software release life cycle.

• Full run

The system deploys the application with all the application processes, components, and artifacts.

· Partial run

The system deploys the application with only the selected application processes, components, and artifacts.

· Selecting artifacts with specific versions to deploy

The system runs the application with only the selected versions of the artifacts.

Snapshot

The system deploys a snapshot of the application. Also go to Example: Creating, Deploying, and Comparing Snapshots on page 78 for an example using snapshots.

• A combination of these ways:

	Smart Deploy	Full Run	Partial Run	Artifacts with Specific Versions	Snapshot
Smart Deploy		No	Yes	Yes	Yes
Full Run	No		No	Yes	Yes
Partial Run	Yes	No		Yes	Yes
Artifacts with Specific Versions	Yes	Yes	Yes		Yes
Snapshot	Yes	Yes	Yes	Yes	n/a

Artifact Staging

ElectricFlow retrieves the artifacts that will be deployed in an application run before the deployment starts. This ensures that all the artifacts are downloaded and available on the target. It also minimizes the time it takes to retrieve artifacts and reduces downtime during the deployment. Go to Artifact Staging on page 70 for more information.

• When an action fails in the process

When an application deployment fails, you can use *rollback* to restore the deployment to a previous state before the failure occurred. This operation is supported only in application processes in the deployment model. Go to Rollback on page 71 for more information.

Master Components

Master components promote component reuse. They provide an easy way to create new components by reusing existing component definitions and the underlying process definitions. Master components are components that can be created and managed independent of any application. A master component consists of a component definition and one or more component processes. Master components are then leveraged to create application components either by referencing or copying them.

You can use master components to create a library of standardized reusable components which can help to promote best practices across the enterprise (such as a standard WebSphere component, standard IIS component, and so on). When deployments have large numbers of similar components, this can be used to speed up application authoring by leveraging prebuilt component and component-process logic.

A component has two parts:

- The component details describe the artifact on which the component is based. It can include where the artifacts are stored (the source or repository server), artifact name, artifact version, the directory where the artifact can be retrieved, and so on. You can parameterize part of the artifact definition to account for variations in the artifact details every time you deploy the application. This allows you to reuse the master component in multiple deployment scenarios. For example, you can create a master component that has a standard process to deploy and undeploy, and the only thing that changes is which specific artifact to use based on parameters.
- Component processes are used to define the automation steps to perform on the artifact defined in the component details. These processes can be used to deploy, undeploy, or do any other action. Each master component can have one or more component processes, and each process can have specific automation steps. These steps can describe how to perform specific tasks such as publish or retrieve an artifact, stop a server process, copy files, restart the server, and so on. A master component must have at least one component processes. You can parameterize the master component processes so that certain behaviors can be achieved when master components are used in an application definition. For example, you can have port number as a component process parameter for a master component, and have different values for the port number in different application models where this master component is used.

After creating the master component, you can reference or copy it to create a new component in an application.

• You can reference a master component when authoring an application. The new application component will have the same properties as the master component. Any changes made to the master component will affect all the application components that reference it.

Use the *reference* method when you want to maintain lineage with the master component so that the application component continues to inherit definitions from the master component. Before making a change to a master component, check how many applications reference it and determine the impact of changing it.

• You can also copy a master component when authoring an application. The new application component will have the same properties as the master or existing application component when it is created.

Use the *copy* method when you want to use a master component just as a baseline for your application component. After an application component is created from a master component, any subsequent changes to the master component will not affect the application component. To get the latest changes in the application component, you delete and recreate the application component, and copy the latest master component to create a new application component.

You can also take an existing component that was originally copied from a master component, modify it, and then save it as a new master component that can be reused and shared by other users.

Master Component Examples

This example shows how to create a master component, create parameters for the artifact definition, define the component processes, create parameters for the component processes, and use the master component in an application, promoting the component reuse throughout your organization. It does not include how to author applications and environments or how to deploy the application. For information about these tasks, see the example in Example: Modeling and Deploying Applications on page 88.

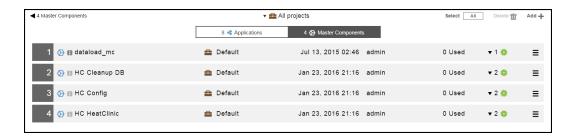
In the following steps, you will learn how to:

- 1. Create a master component and parameters for the artifact definition.
- 2. Author component processes and parameters.
- 3. Use the master component in an application.

Creating a Master Component and Parameters for the Artifact Definition

This section shows how to create a new master component and create parameters for the artifact definition.

Example:



By default, the master components for all projects are displayed.

If you want to see only the objects for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

If there are no matches, a message appears stating that there are no resource templates in the selected projects.

Starting from the Master Components List:

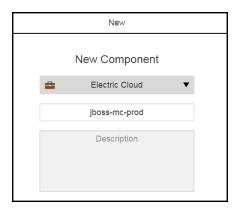
1. Click the Add + button.

The **New** dialog box opens.

- 2. Click Master Component > Create New to create a new master component.
- 3. Enter the name of the new master component and select a project to which the master component belongs.

Note: You can include hyperlinks as part of an object description for any ElectricFlow object.

Example:



4. Click Next.

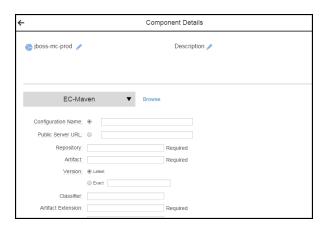
The Component Details page dialog box opens.

5. Click the down arrow in the Content Location field.

6. Select an artifact repository where artifacts are stored, which is represented as a plugin.

Example:

Select EC-Maven.



7. Enter the information about the artifact.

The information that you enter depends on the artifact repository (plugin) that you selected.

Example:

For **EC-Maven**, enter the following information.

• Server URL: http://repo-demo.jfrog.org/artifactory/

· Repository: jboss-releases

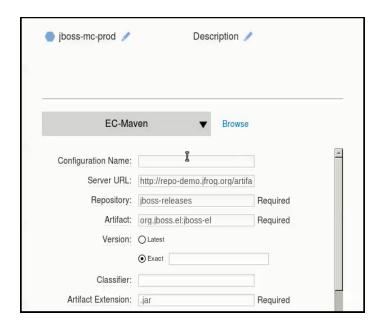
· Artifact: org.jboss.el:jboss-el

· Version: Exact

Artifact extension: jar

We will be parameterizing the artifact name and replacing **jboss-el** with a variable. For now, enter the information in the **Artifact** field above because it is required and the component cannot be created without this required information.

If you had selected another plugin, you may not have to do this.



8. Click OK.

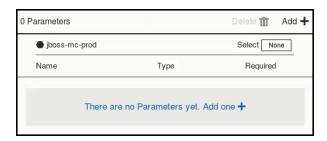
The component that you created now appears in the list.



- 9. Go to the row for the master component that you just created, and click the button.
- 10. Select Parameters.

11. In the Parameters dialog box, click There are no Parameters yet. Add one +, or click Add +.

Example:



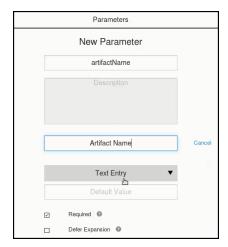
12. In the New Parameter dialog box, enter the parameter settings.

Example:

In the Name field, enter artifactName.

In the Label field, enter Artifact Name.

Keep the rest of the fields as is.

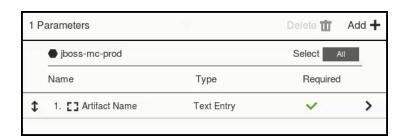


This parameter has default value and is required. The user must enter a value for this parameter when this master component is used in an application.

13. Click **OK**.

Example:

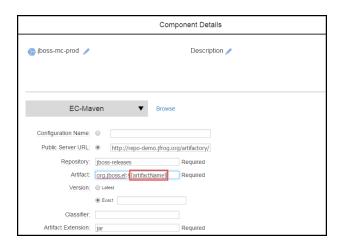
The Parameters dialog box now shows that there is one required parameter for the master component called Artifact Name and that the user has to enter the text (the artifact name) for it.





- 14. Go to the row for the master component that you created, and click the button.
- 15. Select **Details** to modify the artifact definition with the artifact-name parameter.
- 16. In the Component Details dialog box, go to the Artifact field and replace jboss-el with \$[artifactName].

Example:



17. Click **OK**.

You have now created a master component with one parameter for the artifact definition.

Note: The other ways to create a master component are:

- Create it from an existing master component by selecting an existing component from the "Create from existing component" list in the dialog box.
- Create from an existing application component by selecting an application and an application tier associated with that application and then an existing component in that application and application tier.

The next step is to add processes to the master components and create parameters for the processes.

Authoring Component Processes and Creating Parameters

This section shows how to author component processes for the master component created in the previous section and create parameters for those processes.

Starting in the Master Component list:



- 1. Go to the row for the master component that you created, and click the button.
- 2. Select Add process.

The New dialog box opens.

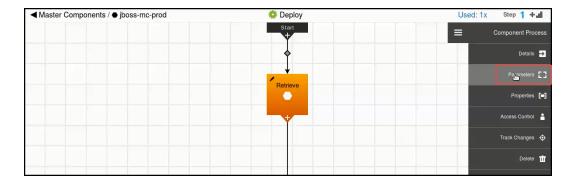
- 3. Enter the component process details, and click **OK**.
- 4. Define the component process in the Component Process Visual Editor.

For detailed instructions and examples, go to Authoring Component Processes on page 92.

5. Select the **Deploy** process.

The Component Process Visual Editor for it opens.

Example:





- 6. Click the button on the right side, and select **Parameters**.
- 7. In the Parameters dialog box for the Deploy process, click **There are no Parameters yet. Add one +**, or click **Add +**.

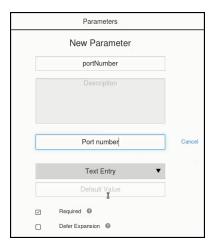
8. In the New Parameter dialog box, enter the parameter settings, and click **OK**.

Example:

In the Name field, enter portNumber.

In the Label field, enter Port Number.

Keep the rest of the fields as is.



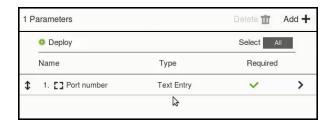
This parameter has no default value and is required. The user must enter a port number when this master component is used in an application.

9. Click OK.

Example:

The port number parameter is created for the Deploy process.

The Parameters dialog box shows that the port number parameter is required for the Deploy process.



10. Close the Parameters dialog box.

- 11. Reference the master component in an application, or copy it to an application.
 - To reference the master component in an application, go to Referencing a Master Component in an Application on page 56.
 - To copy the master component to an application, go to Copying a Master Component to an Application on page 61.

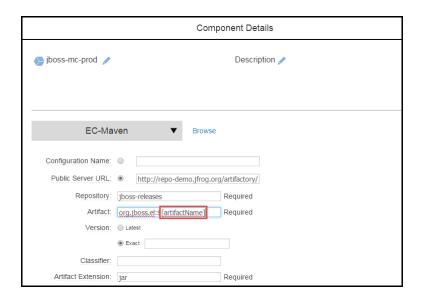
After you add the component to an application tier, a dialog box appears prompting you to enter values for the parameters for the master component process, starting with this step.

You have now authored component processes for the master component and have created a parameter for it.

Viewing the Master Component Details

You can ensure that the artifact definition and component process step defined by the master component are correct by viewing the master component details in the following places.

They first appear in the artifact definition:



They also appear in the component process step details. When you view the component process details for the Deploy process in the Component Process Editor, you can see that the Deploy process consists of one step called Retrieve, which is defined by the jboss-mc-prod master component.



Using a Master Component in an Application

After master components have been defined, they can be used in application models. There are two ways to leverage master components in an application:

- Reference the master component in an application.
- Copy the master component to an application.

To review the implications of these methods, go to Master Components.

Referencing a Master Component in an Application

When you reference a master component in an application, it is added to the application as an application component with the same definition, processes, parameters, and properties as the referenced master component. Any changes made to the master component will affect all the application components that reference it. Use the reference method when the component is used in multiple applications and you want to ensure that these application components continue to inherit all changes made to the master component.

This section shows how to author a new application that references the master component created in the previous sections.

Starting in the Applications List:

- Click the Add + button.
 The New dialog box opens.
- 2. Click **Application > Create New** to create an application.

3. Enter the pipeline name, select a project where the applications in the pipeline will be deployed, and enter a description of it.

Example:



4. Click OK.

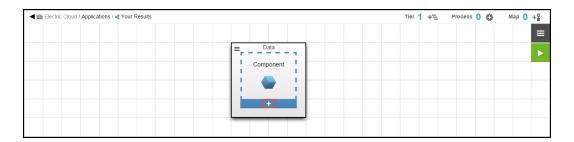
The Applications Visual Editor for your application opens. This is where you will author the application. Notice that there is an undefined application tier (Tier 1) with an undefined component in it.



- 5. Click the dialog box. button in the application tier and select **Details** to open the Applications Tier Details
- 6. Change the tier name to **Data** in the Name field, and click **OK** to rename the application tier.

The Applications Visual Editor now shows one environment tier called Data.

Example:



7. Click the + button in the component to add a component to the application tier.

8. Select Create from existing master component.

A list showing all available the master components opens.

You can also create a new application component using one of these options:

- Create a new application component

 —Clicking this creates a new component as described in Creating a New Application on page 89.
- Create from existing application component—Clicking this opens the Create from existing component dialog box. Click in the Select Application field and select an application. Then click in the Select Tier field and select an application tier. A list of components in the selected application and application tier appears. Select a component from this list.
- 9. In the Create a reference to a master list, select a master component.

In this example, select the master component called jboss-mc-prod.

10. Click Next.

The Create from existing component dialog box opens.

- 11. In the New dialog box, change the name to jboss-mc-results, and click Next.
- 12. Click Next.

If the master component has parameter for its artifact definition, you are prompted to enter values for those parameters in the next step.

Example:



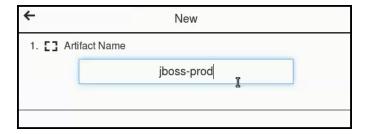
If the component does not have parameters for its artifact definition, click **OK**. Then go to the Component Process Visual Editor and add a parameter to the component, starting at Step 9 in Creating a Master Component and Parameters for the Artifact Definition on page 47Authoring Component Processes and Creating Parameters on page 53

13. Enter the values for the parameters, and click **OK**.

Example:

The jboss-mc-results component has one parameter called Artifact Name.

Enter jboss-prod in the Artifact Name field.



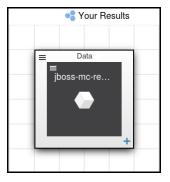
14. Click **OK**.

The Applications Visual Editor now shows that the application has an application tier with a referenced master component.

IMPORTANT: If you want to edit the master component, you need to do this in from the Master Component list. (There is no button in the lower right corner where you can access the component processes.) The UI also lets you know that this is a master component because the icon is different than that for a normal component.

Example:

The Applications Visual Editor now shows that the application has an application tier called Data with a referenced master component called jboss-mc-results.

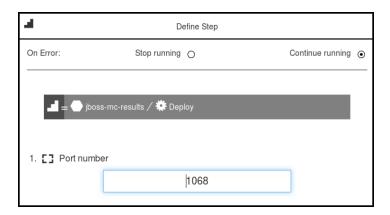


- 15. Define the application process to include a step that references the master component.
 - a. The first step is defined by the jboss-mc-results component and its Deploy process.

After you select the Deploy process, you are prompted enter to a value for the port number. Then click \mathbf{OK} .

Example:

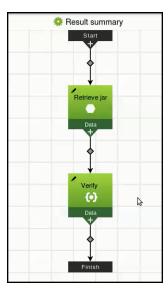
This is the parameter that is created for the master component process.



b. The second step prints the port number that you entered in the previous step.

Example:

This is the application process:



Now you can deploy your application as described in Deploying and Troubleshooting Applications on page 113.

Copying a Master Component to an Application

When you copy a master component to an application, it becomes an application component with the same definition, processes, parameters, and properties as the master component on which it is based. After an application component is created from a master component this way, any subsequent changes to the master component will not affect the application component. Use this method when you want to use a master component just as a baseline for your application component.

For details about how to create master components and parameters for them, see Creating a Master Component and Parameters for the Artifact Definition on page 47 and Authoring Component Processes and Creating Parameters on page 53.

This example describes how to copy a master component in an application.

Starting in the Applications list:

1. Click the Add + button.

The New dialog box opens.

- 2. Click Application > Create New.
- 3. Enter the name of the new master component and select a project to which the master component belongs.

In this example, we will create a new application called "Your Group Summary" that will include the master component called jboss-mcjc.

4. Click OK.

The Applications Visual Editor for your application opens. This is where you will author the application.

Notice that there is an undefined application tier (Tier 1) with an undefined component in it.



- 5. Click the dialog box. button in the application tier and select **Details** to open the Applications Tier Details
- 6. Change the tier name to **Data** in the Name field, and click **OK** to rename the application tier.
- 7. Click the + button in the application tier to add a component to it.
- 8. Select Create from existing master component.

A list showing all available the master components opens.

9. In the "Create a reference to a master" list, select a master component.

Select the master component called **jboss-mcjc**, select **Use a copy**, and select a project to which the master component belongs.

10. Click Next.

The Create from existing component dialog box opens.

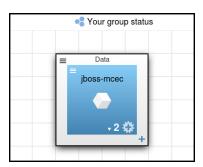
Example:



11. Change the name of the component to something that gives it more context.

Example:

After changing the component name to **jboss-mcec** and clicking **OK**, the Data tier now has a new application component called jboss-mcec, where "ec" refers to a software version:



- 12. If the master component has parameters for the artifact details, modify the artifact details of the new application component and remove the parameter reference.
 - a. Go to the Component Details dialog box.
 - b. Replace the parameter (\$[artifactName]) in the artifact details section with a specific value.

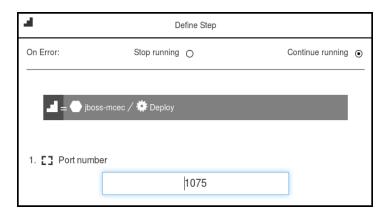
 In the example, replace \$[artifactName] with jboss-5-maintenance, where software versions starting with lower case letters are maintenance releases
 - c. Click OK.

- 13. Define the application process to include a step that references the new application component.
 - a. The first step is defined by the jboss-mcec component and its Deploy process.

After you select the Deploy process, you are prompted enter to a value for the port number. Then click \mathbf{OK} .

Example:

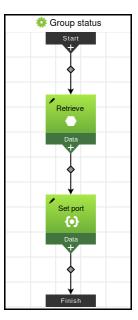
This is the parameter that is created for the component process.



b. The second step prints the port number that you entered in the previous step.

Example:

This is the application process:



14. Make other modifications to the application component and its component processes within application.

Now you can deploy your application as described in the Guidelines for Modeling and Deploying Applications in ElectricFlow on page 85.

Creating Master Components Based on Existing Application Components

Starting from the Master Components List:

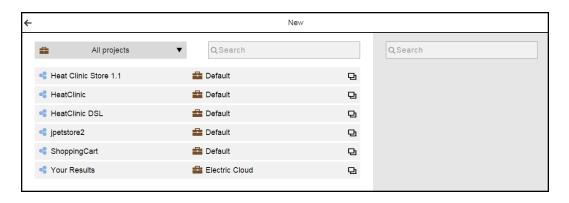
1. Click Add +.

The **New** dialog box opens.

2. Click Master component > Copy Application Component.

A list of existing applications opens. It shows the list of available applications. The default is to show the applications for all projects. If you want to see only the applications for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways: Click on the name of one or more projects. Enter the search criteria in the Search field. The projects that match the criteria appear in the list. You can also search for one or more applications by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no applications in the selected projects.

Example:



3. On the left side of the dialog box, select an existing application.

A list of existing components in that application appears on the right side.

Example:



4. Select an existing application component.

The **Create from existing component** dialog box opens with the name of the selected component highlighted.

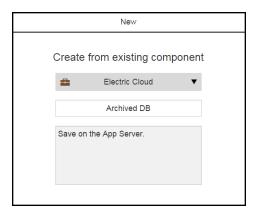
Example:



5. Change the name of the component, and enter an optional description.

You can also change the project to which the master component belongs.

6. Example:



7. Click OK.

The Master Component List now includes master component you just created. It shows the list of available master components. The default is to show the master components for all projects. If you want to see only the master components for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:Click on the name of one or more projects.Enter the search criteria in the Search field. The projects that match the criteria appear in the list.You can also search for one or more master components by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no master components in the selected projects.

Example:



Creating Master Components Based on Existing Master Components

Starting from the Master Components List:

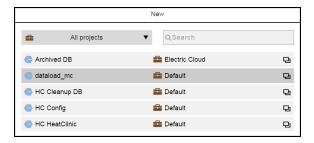
1. Click Add +.

The New dialog box opens.

2. Click Master component > Copy Master Component.

A list of existing master components opens. It shows the list of available master components. The default is to show the master components for all projects. If you want to see only the master components for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways: Click on the name of one or more projects. Enter the search criteria in the Search field. The projects that match the criteria appear in the list. You can also search for one or more master components by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no master components in the selected projects.

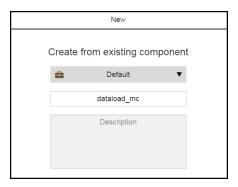
Example:



3. Select an existing application component.

The **Create from existing component** dialog box opens with the name of the selected component highlighted.

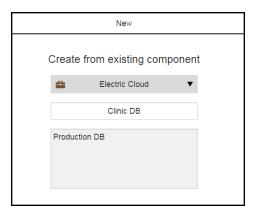
Example:



4. Change the name of the component, and enter an optional description.

You can also change the project to which the master component belongs.

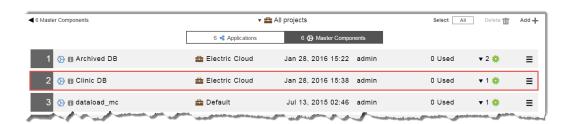
5. Example:



6. Click OK.

The Master Component List now includes master component you just created. It shows the list of available master components. The default is to show the master components for all projects. If you want to see only the master components for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:Click on the name of one or more projects.Enter the search criteria in the Search field. The projects that match the criteria appear in the list.You can also search for one or more master components by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no master components in the selected projects.

Example:



Master Components List UI

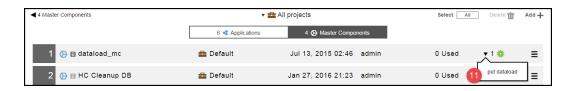
How to get here:

One of these ways:

- From the Home page:
 - 1. Click the Applications launch pad to open the Applications List.
 - 2. Click the Master components tab.
- Click the Main menu button, and then click Applications > Master Components.

This is a Master Components List:





1	Breadcrumb identifying the object you are viewing and the total number of master components.
2	Clicking the i button shows information about the element. This shows the description of the master component.
3	Name of the master component.
4	Name of the project to which the master component belongs.
	It shows the list of available master components. The default is to show the master components for all projects.
	If you want to see only the master components for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:
5	Click on the name of one or more projects.
3	Enter the search criteria in the Search field. The projects that match the criteria appear in the list.
	You can also search for one or more master components by entering the search criteria in the Search field next to the All projects field. If there are no matches, a message appears stating that there are no master components in the selected projects.
6	Date and time when the master component was created.
7	User who created the master component.
8	Number of applications dependent on this master component reference.
9	Number of component processes for the master component.

	Click the Menu button to open the context menu. You can get more details about the master component:
10	Details—The name and description of the object.
	Add Process—Add a component process. The Component Process Details dialog box opens.
	A master component can have one or more component processes.
	Parameters—The parameters defined for the master component.
	Properties—The properties in the object.
	 Access Control The access control configuration in the automation platform for the object.
	Track Changes—The change history of the object.
	Delete-Delete this object.
11	Clicking on the down arrow next to number of component processes (number 9 above) shows the list of processes for this master component.

Artifact Staging

You can use artifact staging to retrieve artifacts that will be deployed in an application run before the deployment starts. This ensures that all the artifacts are downloaded and available on the target, either the artifact cache (the default) or a directory specified in the component definition. Staging artifacts minimizes the time it takes to retrieve artifacts and reduces downtime during the deployment.

Artifact staging is enabled by default in the UI. However, when you use the API commands in the command line interface or in a script, make sure to enable artifact staging when modeling the application.

Go to Deploying and Troubleshooting Applications on page 113 for how to use artifact staging when deploying an application.

At the Application Level

When artifact staging is enabled in a deployment model, ElectricFlow dynamically creates and executes a job step to stage the artifacts. If the application process being deployed includes process branching, the staging job step retrieves all artifacts in the application process in parallel to reduce the staging time. In addition, the deployment processes, including the staging job step, are skipped on disabled resources.

You can deploy applications with the artifact staging enabled and following options:

- With smart deploy enabled, the staging job step retrieves artifacts only for those objects that have not been deployed to specific resources.
- With full deploy enabled, the staging job step retrieves artifacts only for all artifacts in the application.
- With partial deploy enabled, the staging job step retrieves artifacts only for artifacts in the application.
- With a selected snapshot, the staging job step retrieves artifacts only for artifacts in the selected snapshot.

Note:

For a component process, the process step to retrieve an artifact is optional because the component process is referenced in the application process. If you author a component process, the process type should be "Deploy" to ensure that the artifacts defining the component can be staged.

If an artifact staging for a component fails, it will then be retried when the process execution continues and when the actual component process execution happens.

You can view the progress of the staging job step in the Application View Run page.

- When the staging job step is successful, the artifact is retrieved and staged before the deployment run starts.
- When the staging job step fails, the application process automatically retrieves the artifact during the component retrieval step in the application run.

The environment inventory is updated as the job step runs. In addition to viewing the status of the resources in the environment, you can also view the state of other objects including the application processes, components and component processes, and artifacts being deployed.

At the Pipeline and Release Levels

Artifact staging is supported in pipelines and Releases.

- When you define the tasks in a pipeline stage, you can select application processes with artifact staging.
 The staging job step is executed as part of the application in the pipeline run as described in the previous section. Go to Tasks on page 218 for more information about stage tasks in a pipeline.
- The Release definition includes the pipeline and applications that you want to run in the Release. Go to Release Definition on page 283 for more information about the release definition.

Rollback

When an application process has a predefined rollback step and the application deployment fails, ElectricFlow provides several ways to restore the deployment model to a previous state before the failure occurred. ElectricFlow provides these ways to do this:

- Automatic rollback—ElectricFlow automatically reverses what has been deployed up to the point of failure in the process flow. Go to Rollback on page 71 and Example: Modeling and Deploying Applications on page 88 for more information.
- Snapshots
 –You manually create and save snapshots of previous successful deployments and run one
 of these the next time the application is deployed. Go to Snapshots on page 77 for more information.
- A component process defined as an Undeploy process—When defining a component process, select Undeploy as the process type.

Automatic Rollback

ElectricFlow supports a *rollback* step in application processes. This allows an automatic rollback to the environment before an erroneous deployment or to any previous environment snapshot. This reduces the time to author the rollback logic. For an example of showing how to add a rollback step to an application process, go to this step in Example: Modeling and Deploying Applications on page 88.

During an application deployment, the predefined rollback operation is invoked when the rollback step is executed. ElectricFlow reverses what it has executed in the process flow up to the application process step where the failure occurred. This operation is supported only in application processes. It puts the deployment in the same state as the last successful deployment, which is usually the same as the previous deployment state.

When defining an application process step as a rollback step, the application can roll back to the previous state up to where the deployment failed (the default) or to a specific snapshot. In addition, you can specify how the application will be rolled back.

- By default, smart deploy is enabled. If you want to ElectricFlow to execute a full or partial run the next time the application is deployed, disable this option.
- When the Undeploy Application before Rollback option and the application process that you want to
 undeploy (in the Select Process field) are selected, the environment inventory record is removed before
 the rollback operation is executed. This allows dirty files in the erroneous deployment to be
 automatically removed from the system.

When a specific snapshot is selected for rollback, ElectricFlow first tries to determine the last successfully deployed process and then tries to execute it from the snapshot.

While a native rollback step is available, ElectricFlow also allows you to have complete control over creating a custom rollback handling operation and using that as part of the error handling logic.

For an example of showing how to add a rollback step to an application process, go to this step in Example: Modeling and Deploying Applications on page 88.

Manual Tasks and Steps

You can define manual steps and tasks in application or component processes and in release delivery pipelines. They describe the actions that must be taken by a human to complete the process step and/or pipeline task. In application and component processes, these actions are referred to as *manual steps*. In pipelines, they are referred to as *manual tasks*. Any time a manual task or step is reached in the process flow, an email notification is sent to all users in the assignee list.

Manual steps and tasks consist of user-specified instructions and assignees.

- To set up a manual task or step:
 - Enter instructional text to help the approvers. This text can include URLs.
 - You can also define optional parameters to collect information during the approval step.
 - Specify who can approve. You can define the step or task to accept either user or group IDs or variables using the \${ [] notation to represent assignees as property references, but not both.

- To attach a parameter to a manual task or step:
 - Use the \${[] notation as a property reference to access other user-specified parameters at runtime, such as evidence and status.
 - For a pipeline manual task, use \$[/javascript myFlowRuntime.flowRuntimeStates [0].evidence] to collect the user's response in subsequent tasks. This command using (0) assumes that the manual task is the first task in the stage.
 - For an application process, use \$[/myJobStep/comment] or \$[/myJob/jobSteps/stepName/comment] to access the instructional text, which is usually in the notification emails sent to the users who perform the actions or who approve the completion of the manual step.
 - Put hyperlinks in the instructions and comments fields in the GUI. You can create a file on a server or other storage device and add the URL to that file in the GUI.
- To specify a default email template:
 - For a manual task in a pipeline, use /ec_deploy/ec_pipelineNotifierTemplates/ec_default_pipeline_manual_task_notification_template.
 - For a manual process step in an application, use /ec_deploy/ec_notifierTemplates/ec_default_manual_process_step_notification_template.

For examples showing how to define manual tasks and steps:

- Go to Example: Authoring a Pipeline with Manual and Utility Tasks on page 232 for how to define a manual task in pipeline stage.
- Go to this step in Example: Modeling and Deploying Applications on page 88 for an example showing how to define a manual step in a component process. The steps to define a manual step in an application process are the same.

Process Branching

You can use process branching to specify the path through an application or component process based on transition conditions other than out-of- the-box options. Decisions about the next step in the process are made while the process runs. This is similar to the transition conditions for workflows in the automation platform. If the application or component process applies to multiple use cases, you can design one process with two or more branches instead of designing multiple processes for each use case. You can also define steps that run in parallel.

For example, to install or upgrade software, you can define one process for multiple use cases and use the same steps except for the following:

- The source files can be in .zip or .tar format. The steps to extract the files depend on the format.
- The operating system can be Linux or Windows. The steps to download the files, install them on the server, and enter commands depend on the operating system.

ElectricFlow supports the following branching conditions. The default is **Always**.

- · Completion status of the previous process step
- A property set in another part of the system, not the in the previous step
- · Custom validation rules

Using Process Branching

You define process branching when authoring an application or component process.

Process Branching Example

This example shows an application process with process branching that deploys a .war file.

Clicking a connector in a link opens the branching conditions menu. Depending on the location of the connector, some or all of these options may be available (enabled):

- Always—Always go to the next step, referred to as the target.
- Successful—Go to the next step if the previous step, referred to as the source, is successful.
- Failure-Go to the next step if the previous step fails.
- Add Condition—Add a custom condition that must be met before going to the next step.
- Add Connector
 –Add a connector from the source of the link to a new target by selecting one of the
 highlighted eligible steps. You can only select an eligible step.
- **Change Source**—Change the source by selecting one of the highlighted eligible steps, which has a red outline. You can only select an eligible step.
- Change Target—Change the target by selecting one of the highlighted eligible steps, which has a red outline. You can only select an eligible step.
- Delete-Delete the selected connector and link.

When you select the connector between the Start step and first step after it, only some of conditions appear and only some are available. The condition between the Start and the next step is **Always**, the default branching condition

When you select the connector between subsequent steps, all of the conditions are available.

Usage Guidelines

Follow these guidelines when the process includes process branching:

- When you add a step, you must define it before adding another step.
- You can only configure branching conditions on a connector between two process steps.
- You cannot configure branching conditions between these objects:
 - The start of the process and the steps immediately after it.
 - The end of the process and the steps immediately before it.
- When defining a step in an application or component process, you configure what ElectricFlow does when an error occurs. *This setting overrides any job-step-level branching condition*.

You can select **Stop running** or **Continue running** in the **On Error** field in the Define Step dialog box.

If an error occurs in a job step and **Stop running** is selected, ElectricFlow aborts the process even if the branching condition is set to **Failure**.

Process Branching States and Conditions

State of the Branching Condition Connectors in the UI

In the ElectricFlow UI, the status of the link is based on the shape and color of the connector.

Shape	Color	Link Status
Diamond	Light gray	Always
Diamond	Dark gray	Disabled
Square	Green	Successful
Circle	Red	Failure

Examples of Branching Conditions

These are examples of branching conditions that you can apply in your processes.

· Based on the status of the previous step

Follow the branch based on the result of the previous step: Successful, Failure, or both (Always).

Example:

- Successful-If the file is downloaded successfully, the next step is to extract the files.
- Failure- If the file was not downloaded properly, the next step is to abort the process.
- Always-The next step is to always extract the files.
- · Based on a value of an operation during the step

Follow the branch that matches the result of an operation such as calculating a value or processing data during the step.

Example: The result of an operation is a file type.

- If the result is an XML zip file, the next step is open an XML text editor.
- If the result is a .htm file, open a web browser.
- If the result is a .mov, open an application to play the movie.
- Based on a property in another part of the system

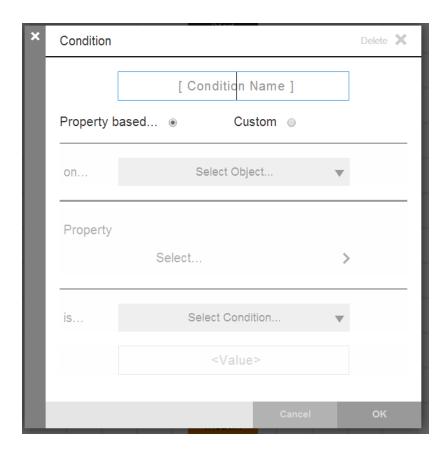
Follow the branch based on a property set in another part of the system, not in the previous process step.

Example:

- If the property os_type = linux is set on a resource, always follow the branch for Linux steps.
- If the property release_type is set to minor in the application, always follow the branch for minor releases when running the process.

Custom Conditions in Process Branching

In a component or application process with branching, when you click the connector on a link between two steps, the Condition dialog box opens.

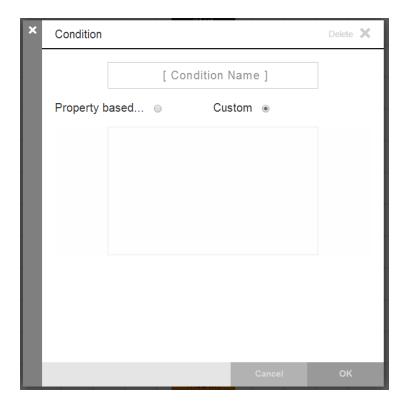


When you click the connector between the "Deploy War" and "Start Server" steps, these conditions appear:

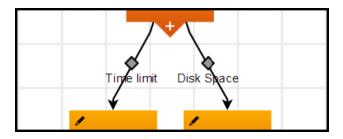
- Always—Always go to the next step, referred to as the target.
- Successful-Go to the next step if the previous step, referred to as the source, is successful.
- Failure-Go to the next step if the previous step fails.
- Add Condition—Add a custom condition that must be met before going to the next step.
- Add Connector—Add a connector from the source of the link to a new target by selecting one of the highlighted eligible steps. You can only select an eligible step.
- Change Source—Change the source by selecting one of the highlighted eligible steps, which has a red outline. You can only select an eligible step.
- **Change Target**—Change the target by selecting one of the highlighted eligible steps, which has a red outline. You can only select an eligible step.
- Delete-Delete the selected connector and link.

When you configure a **Property based** condition, the fields in the Condition dialog box remain the same.





After you configure your conditions, they appear near the affected connectors in the process.



Snapshots

In a software release process, the goal is to have a consistent deployment process that can be run often throughout the release life cycle. When the deployments are run in production, there are no unexpected results. Snapshots are used to implement these deployments.

A snapshot captures both the component versions and the process that were used for deployments. These snapshots make the deployments repeatable and stable.

ElectricFlow supports these types of snapshots:

- Environment snapshot: An environment snapshot is created from an environment where the application is deployed and tested. A snapshot will capture the exact artifact versions, process versions, configuration parameters, and so on used for the deployment to ensure repeatability. This is considered the reliable version of a snapshot for reliability and repeatability.
- Application snapshot—An application snapshot is created directly from the application model and does
 not depend on any environment deployment. When a snapshot is captured from the application model, it
 locks the application architecture, components, and process versions. If you specify the "Latest" artifact
 version as part of the component definition in the snapshot, the latest artifact versions will be used every
 time it is used for deployment. The main purpose of an application-based snapshot is to allow the
 creation of user-defined versions of the models.

For example, in a software delivery process, the developer checks in code, the components are built automatically by a continuous integration system, and the objects in the resulting build are published to a binary artifact management system. When the full complement of component artifacts is available, a release candidate pipeline is initiated. The first stage is to deploy to a development environment. From this deployment, a snapshot is created and used as the model for subsequent deployments to other environments. There will be some differences between the development, QA and production environments, such as database connection strings, credentials, and debug-vs-production artifacts. ElectricFlow can manage these differences in the application model using parameters and properties in templates, only modifying the deployment process flow in a controlled manner.

The process described in the example avoids some of the pitfalls encountered in traditional release scenarios. For example, configuration teams often perform their own builds. This means that the binaries supplied by the CM team may be different than what has already been tested. Release teams may also manually edit configuration files to adjust for differences in the deployment environments. Both of these practices are prone to error and can negatively impact the quality of the final software release.

In ElectricFlow, you can deploy snapshots for more reliable and repeatable software releases.

- You can model and save a version of your software release with specific artifact versions and rerun it later, even if the latest release version of it has changed.
- If you save snapshots during development and test phases, you can ensure that the components that
 were deployed and tested are the same as those in the release candidate. You can redeploy the
 snapshot any time.
- You can create and save more than one snapshot for different deployment scenarios.
- You can view the snapshots in the Snapshot List. From this list, you can manage all your snapshots, compare two snapshots, or get more information about them.
- Comparing snapshots helps you to deploy releases with reliable and repeatable results during ongoing software release cycles. You can build and test releases using snapshots, and do not have to design new applications for each release.
- You can use snapshots to refine and optimize applications that fit your deployment scenario and ensure that this version is properly developed, tested, and released.

Example: Creating, Deploying, and Comparing Snapshots

This example shows how to create application and environment snapshots, deploying snapshots, and comparing them during the software release life cycle.

Creating Application Snapshots



- 1. In the Applications List, choose an application and click the
- button.

2. Select Snapshot List.

The Snapshot List appears.

- 3. To add a new snapshot, click There are no Snapshots. Add one +.
 - The New Snapshot dialog box appears.
- 4. Enter the name for the snapshot that must be unique within the application and enter an optional description of it.
- 5. Click Next.

6. The New/Preview dialog box opens.

You can toggle between the **Component** and **App Process** views to see the components and application processes for the selected application.

Example:

In the Component view, change the version of the Config component to Latest.



This is the App Process view.



7. Click OK.

A messages appears stating that the snapshot has been created, and the Snapshot List is updated:



Deploying Application Snapshots



1. In the Applications List, choose an application and click the

button (the Run process button).

2. Select New Run.

The dialog box to select the settings for the application run opens.

3. Select these settings for the application run:

In the Select Process field, select Run.

In the Select Environment field, select heatclinic-qe.

In the Select a Snapshot field, select **HeatClinic v1.2**.





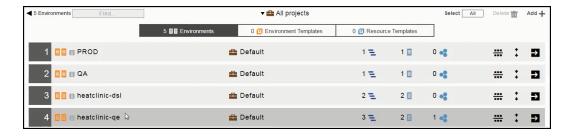
Note: Clicking next to **Compare** shows the comparison of the selected snapshot to last deployment of the **Run** process in the **heatclinic-qe** environment.

4. To deploy the snapshot, click **OK**.

Creating Environment Snapshots



1. In the Environments List. choose an environment and click the button.



2. Click the **Inventory** button to open the Environment Inventory.

Example:





- 3. Choose an application, and click the
- button.

4. Click New Snapshot.

The New Snapshot dialog box opens.

5. Enter the name and a description of the snapshot.

The name of this environment snapshot is *Heatclinic_QE2*.

6. Click Next.

The New/Preview dialog box opens. You can view the snapshot settings in the Component and App Process tabs but cannot change them because the application has been deployed to the heatclinic-que environment.

7. Click OK.

Comparing Snapshots

This example shows how to compare an application snapshot to an environment snapshot..

1. In the Snapshot List, select two snapshots that you want to compare.





2. Click the button (**Compare** button).

The Snapshot Comparison page opens.

Comparing Results

In the Snapshot Comparison page, the Component and App Process views show the differences between the snapshots, Heatclinic v1.2 (the application snapshot) and Heatclinic_QE2 (the environment snapshot). Heatclinic_QE2 is the more reliable, repeatable, and update-to-date version of the application.

The **Component** view shows that the Config component is different. Heatclinic v1.2 uses the "Latest" version of the artifact that may not be the same as the artifact version in Heatclinic_QE2, the more reliable and repeatable version of the snapshots.



The App Process shows one difference. Heatclinic v1.2 has one less application process than Heatclinic_QE2.



Note: Clicking the pull-down button next to the snapshot name allows you to select a different snapshot.

Configuration Drift

Throughout the software release process, ElectricFlow tracks changes in the release to ensure that the applications are repeatedly and reliably deployed and that the release is successfully completed. ElectricFlow uses snapshots and change tracking to track the changes and identify differences, and then can notify users about them.

Snapshots are a version of the application at a point in time, encapsulating the specific state of the application, its environment states, its components, and the artifacts and artifact versions that define the components. You can manually take snapshots and compare two to determine the differences between them. You can also automate a procedure that takes snapshots on a schedule, create custom reports, and send them to users, making it easier to identify differences between deployments and resolve issues in the release. Go to Snapshots on page 77 for more information about snapshots.

Change tracking is another way to track differences between deployments. Change tracking can track the changes between every state of non-runtime objects and allows you to revert to any previous state of the objects, such as applications, application processes, components, artifacts, procedures, workflows, resources, and workspaces. In the Change History, you can find what changed for a specific object, such as an artifact, and when the change occurred. Go to Change Tracking on page 730 for more information about this feature and how to use it.

Deployment Examples

This section has examples of application deployments and includes:

- How to identify ElectricFlow objects—These are easily identified by icons and colors in the Visual Editors (such as the Applications Visual Editor, Environments Visual Editor, and Application Process Visual Editor) and lists (such as the Applications List and the Environment List).
- · How to view details about the objects
 - In the Applications Visual Editor and Environments Visual Editor, you can view details and other information of the application tiers, components, environment tiers, and resources.
 - In the Application Process Visual Editor and Component Process Visual Editor, you can view details about the process steps.
- How a process can take different paths in the deployment based on parameter inputs.
- · What to expect with different deployment options.
- How to model dynamic environments that provision (commission) environments and spin them up at runtime.
- How to use master components.
- · How to use snapshots.

Guidelines for Modeling and Deploying Applications in ElectricFlow

When Modeling Applications

Keep the following in mind:

- By default, change tracking is enabled for all ElectricFlow Deploy projects.
- When authoring a component process, you specify the process type. The default is **Deploy**.

In the "Component Process Details" dialog box, select one of these values in the Type field:

- Deploy

 Enables Inventory Tracking. The ElectricFlow server tracks the artifacts deployed to
 environments.
- **Undeploy**—The next time that the process is run, the ElectricFlow server removes information about the artifacts deployed to environments.
- Other–Disables Inventory Tracking.
- You can also create a library of standardized reusable components to promote best practices across the
 enterprise and speed up application authoring with *master components*. This is an easy way to reuse
 existing definitions and their underlying process definitions. Go Master Components on page 46 for
 more information.

• After creating an application, you can version the entire application model by taking a snapshot of it. A snapshot is an immutable version of an application with specific artifact versions.

You can save more than one snapshot of an application during the software release life cycle. You can compare snapshots to optimize and troubleshoot the application. You can also deploy a snapshot even if the latest version of the application has been modified.

Go to Snapshots on page 77 for more information.

- You can include hyperlinks as part of an object description for any ElectricFlow object.
- Known issues:
 - If an application tier has two components that cannot be deployed to the same host, the components need to be in different application tiers.
 - If a formal parameter in a manual process step is renamed, the formal parameter properties are not updated.
- Go to Inventory Tracking on page 212 for more information.
- Go to Environment Inventory for more information.
- Go to Example: Modeling and Deploying Applications on page 88 for an example of the Application Inventory in the Application View Run page.

When Modeling Environments

You can optimize how resources are used in application deployments. You can model dynamic environments that are automatically spun up when an application is deployed. These environments can have cloud and static resources. Cloud resources are provisioned using resource templates, which are specified in environment templates.

Using dynamic environments allows you to:

- Provide ways to optimize how cloud resources are used.
- · Reuse provisioned resource pools.
- Track how provisioned cloud resources are used.
- · Provide the status of the provisioning process.
- · Verify the credentials of cloud resources before provisioning them.
- Configure the middleware of cloud resources on an on-demand basis.

When Deploying Applications

How an application is deployed depends on what you want to deploy. You can select the following deployment options.

- New run-Specify the runtime parameters and settings in the dialog box.
- Last run—Use the parameters from a previous run. You can modify one or more of these parameters as described in the dialog box where the runtime parameters and settings are set.

Note: The option does not appear the first time an application is deployed.

 Schedule

Set the application to run on a schedule as described in Running Applications with Schedules.

You can also deploy the application with these options and methods:

- Full deploy

 ElectricFlow deploys all the objects (including application processes, components, and
 artifacts) in the application. When you deploy an application for the first time, you must do a full deploy.

 Later, you can use this method to verify that the software is ready to be released.
- Smart deploy—This deployment method reduces risk and time by only deploying changed objects.
 ElectricFlow automatically detects the differences between what is about to be deployed and what has
 already deployed to an environment. It deploys only those objects not yet deployed to a specific
 resource or to resources where the object version is different. You can use this method throughout the
 software release life cycle.
- Partial deploy

 —The system deploys only objects that you select. You can use this method to verify incremental changes.
- Partial deploy with specific artifact versions—The system deploys only the artifacts with selected versions.
 You can use this method to verify incremental changes.
- Schedule–You can create schedules to deploy applications on a one-time, daily, weekly, or monthly
 basis. You can schedule nightly builds so that the developers and quality always have a new version to
 work on every day.
- Snapshot–You select a snapshot to deploy. A snapshot is an immutable version of the application with specific artifact versions. You can save and compare snapshots throughout a Release life cycle.

Note: If you select a snapshot and modify it, it is no longer a snapshot and becomes a different version of the application, which you can save as a new snapshot.

- Based on custom parameters defined in application processes—Set these parameters when you deploy
 the application or when you define an application process step. They determine how the application
 should be deployed.
- Artifact staging
 – ElectricFlow retrieves the artifacts that will be deployed in an application run before the
 deployment starts. This ensures that all the artifacts are downloaded and available on the target. It also
 minimizes the time it takes to retrieve artifacts and reduces downtime during the deployment. Go to
 Artifact Staging on page 70 for more information.
- Rollback—When an application deployment fails, you can use the rollback functionality to restore the
 deployment model to a previous state before the failure occurred. This operation is supported only in
 application processes. Go to Rollback on page 71 for more information.

You can deploy applications by combining methods, such as:

- · Smart deploy and partial deploy
- Smart deploy, artifact staging, and a partial deploy with specific artifact versions
- Full deploy and artifact staging with a snapshot
- · Smart deploy and artifact staging
- · Artifact staging and partial deploy

For example, this sequence is an example of how you may deploy your software over time:

- When you create an application in ElectricFlow and deploy it for the first time, you must do a full deploy. By default, smart deploy is disabled the first time that you run an application.
- Throughout these deployments, you can schedule the deployments to occur on a daily or weekly basis
 to provide builds that can be tested and installed on a regular basis.
- You can also save snapshots of the software at regular intervals or milestones. When you need to
 troubleshoot the software or want to do performance testing, you can use one of these snapshots for
 comparison.
- If the application does not deploy successfully, you can redeploy parts of it to troubleshoot the application or component processes that failed.
 - You can do a partial deploy and redeploy the application with only the objects that failed.
 - You can also do a partial deploy only with specific versions of artifacts to determine if one or more specific versions of artifacts are causing problems.
- Later, after successfully deploying the application, you can redeploy parts of the application when new versions of artifacts or new resources are available.
- When a new version of an artifact is released, you can deploy only the artifact by selecting the new version and doing a partial deploy.
- When you add artifacts and resources to the application, you deploy the new artifacts to resources and specific versions of selected artifacts to the new resources, a combination of smart deploy and partial deploy with specific artifact versions.

Getting the Real-Time Status of Application Runs and Troubleshooting

Use the Environment Inventory and Application Inventory to view the progress of the application as it runs and the results of previous runs. They show detailed results that can be used to troubleshoot the application. Go to Inventory Tracking on page 212for more information.

- In the Application Inventory (on the Applications Run View page), you can get information about the application, its application processes, components, and job steps and about the status of these objects.
- In the Environment Inventory, you can get more details about the environment, the applications mapped
 to it, number of deployed artifacts in the applications, where the artifacts are deployed, and the status of
 these objects.

Use Change Tracking to troubleshoot a failed job, look for more information about non-run object, and/or look for differences between objects. Go to Change Tracking on page 730 for more information.

Compare snapshots to refine and troubleshoot an application. Go to Snapshots on page 77 for more information.

Example: Modeling and Deploying Applications

Applications consist of components and the processes that orchestrate the deployment of these components. Each application object has one or more application tiers, where a tier is a logical group of components. A component is based on an artifact, and the component definition contains a reference to this artifact.

When defining a component, you specify the artifact repository (plugin) and the artifact version that you want to use. ElectricFlow supports artifacts from a variety of Artifact Repositories, such as the EC-Artifact repository, Maven based repositories like Nexus or Artifactory (using EC-Maven), or artifacts on any file system (using EC-

FileSysRepo). These artifacts could point to any build outputs, configuration scripts, database, scripts, and so on. The information that you enter in the component definition depends on the artifact repository that you select. For more information about artifacts, go to Artifact Management on page 707 and Artifacts on page 498.

You can use Change Tracking to track all the changes for non-runtime objects such as the application model, components, artifacts, resources, the environment model, and projects. Go to Configuring Change Tracking on page 731 for more information.

This example shows how to model and deploy an application in ElectricFlow.

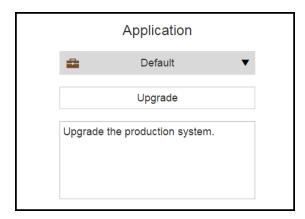
Creating a New Application

This example shows how to create an application consisting of multiple application tiers with multiple components.

In the Applications List:

- 1. Click **Add +** to add an application.
- 2. In the New dialog box, click Application > Create New.
- 3. Select a project to which the new application will belong in the **Select Project** field, enter a description of the application, and click **OK**.

Note: You can include hyperlinks as part of an object description for any ElectricFlow object. Example:



The Applications Visual Editor opens.

4. Define an application tier and the components in it.



In Tier 1, click and select **Details** to edit the details in the **Application Tier Details** dialog box.

- a. Rename the application tier to "App Server" and click **OK**.
- b. To define the component:
 - i. Click + under the Component.
 - ii. Click Create new application component.
 - iii. Enter "Config" as the component name, and click Next.
 - iv. In the Content Location field, select EC-FileSysRepo.
 - v. Enter the required Source and Artifact.
 - vi. Enter the artifact Version. In this example, select Exact, and enter 1.
 - vii. Click OK.
 - viii. Depending on the value that you select in the **Content Location** field, different fields appear.
 - ix. Enter the information to define the component.

Example:



c. In the "App Server" tier, click + in the lower right corner to add another component.Define another component similar to how to defined the first component in this step:Example:



#=

5. In the Applications Visual editor, click to add another tier to the application, and define the first component in it. The new tier is called "Tier 1." Repeat the procedure in this step to define application tier and the components in it. This tier will be called "DB" and will have two components.

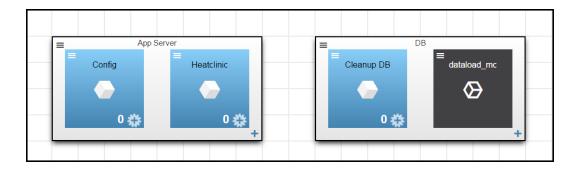
Example: Name it "Cleanup DB" with the following definition.



- 6. To a dd a master component to the DB tier:
 - a. Click + under the Component to define the component.
 - b. Click Create from existing master component.
 - c. Select dataload_mc from the list of master components, and click Next.

The Applications Visual Editor now shows the two application tiers with two components each.

Example:



Authoring Component Processes

In a component process, you define a set of actions that will be executed on a specific component during the deployment. A component process step can be defined as a call to an ElectricFlow plugin or procedure, a direct command or script in any scripting language, a manual task, or a utility function (a higher-order operation than a plugin).

This example shows how to author a component process with a manual step and step defined by a utility function.

In the Cleanup DB component of the "DB" tier:



- 1. Click to start authoring a component process.
- 2. In the Component Process Details dialog box, enter the following details:
 - Enter "Start server" in the Name field.
 - Enter an optional description in the Description field.
 - Keep the selection as **Deploy** in the **Type** field. This value determines track what is built, tested, and deployed during the application deployment. The default is **Deploy**.
 - **Deploy**–Enables Inventory Tracking. The ElectricFlow server tracks artifacts deployed to environments.
 - **Undeploy**—After the first successful job step in a component process with this setting, the automation platform removes the environment inventory record.
 - Other–Disables Inventory Tracking.



- Specify who can run the component process in the Credential field. Clicking displays a
 dialog box where you select a project and then a credential (user name and password) in
 that project for the process, which are created and managed at the platform level. You can
 only impersonate one credential. To set the credential type, see Adding Credentials.
- Select a workspace from the drop-down list of workspaces in the ElectricFlow platform in the Workspace field, where the files and results of the job step are stored (optional).
- Set a time limit in the **Time limit** field, the maximum amount of time that the step is allowed to run. After the time specified, the step is aborted. Enter the time, and select the unit of time: seconds, minutes, or hours. Go to the Job Step Details on page 533 for information how time limits work in procedure job steps.

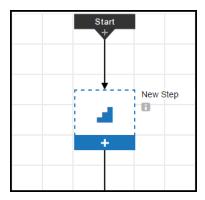
3. Click OK.

The Component Process Visual Editor opens.

The process model shows a process with **Start** and **Finish** steps with a **New Step** in between them.

 $\textbf{Note:} \ \textbf{Clicking i} \ \textbf{under the step name displays the details about the process step.}$

Example:



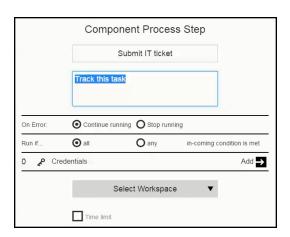
4. Click + in the first step.

- 5. Define the first step in the Component Process Step dialog box:
 - a. Enter "Retrieve" in the Name field.
 - b. Enter "Retrieve the new DB file" in the **Description** field.
 - c. Use the **Continue Running** setting in the **On Error** field. Go here for more information.
 - d. Use the all setting in the Run if field. Go here for more information.



- e. Specify who can run the component process in the **Credential** field. Clicking displays a dialog box where you select a project and then a credential (user name and password) in that project for the process step, which are created and managed at the platform level. You can only impersonate one credential. To set the credential type, see Adding Credentials.
- f. Click Next.

Example:



- 6. To define the first step in the Component Process Step dialog box as a manual step, select **Manual** as the **Step Type**, and enter the details about it:
 - a. Enter the users or groups who can perform and approve this task in the Assignees field.
 Email notifications are sent to assignees who perform the task and to those who approve the task before the process can proceed to the next step.
 - As you type the user or group name, ElectricFlow displays the possible matches.
 You can select one or more of the suggested users or groups or type the entire name to complete the entry.
 - You can also use the "\$[]" notation for parameter substitution to dynamically assign users or groups to perform and approve manual tasks at runtime.
 - b. Enter the instructions for the manual task.

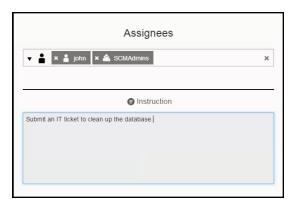
The instructions can include property references using the "\$[]" notation for file links, which is needed when referencing a task instruction in an email notification template. For example, a manual task could have this instruction:

Please check the Clover report, and make sure there is at least 80% code coverage before you approve.

http://myserver.com/reports/\$[/myPipelineRuntime/ec_summary/report_name]

c. Click OK.

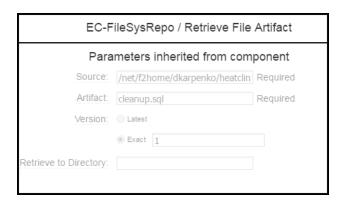
Example:



- 7. Click + on the bottom of the first step to add a new step below it.
- 8. Click + in the second step.
- 9. Define the second step in the Component Process Step dialog box:
 - a. Enter "Submit IT ticket" in the Name field.
 - b. Enter "Track this task" in the **Description** field.
 - c. Use the Continue Running setting in the On Error field. Go here for more information.
 - d. Use the **all** setting in the **Run if** field. Go here for more information.
 - e. Click Next.

10. To define the second step in the Component Process Step dialog box based on a component operation, select Component Operations as the Step Type. The parameters for this process step are automatically inherited from the "Cleanup DB" component.

Example:



- 11. Add a third step to the process as described in the previous steps. The step type is **Command**. You can enter a command or script when defining this step.
- 12. Add a fourth step called "Tear down environment" as described in the previous steps, and define the fourth step in the Component Process Step dialog box as a **Utility** step. This is also referred to as a utility function that is a higher-order operation than a third-party plugin.
 - a. Select Utilities as the Step Type

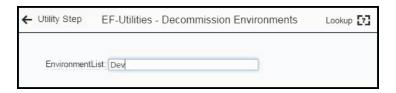


Select EF-Utilities - Decommission Environments and click
 Example:



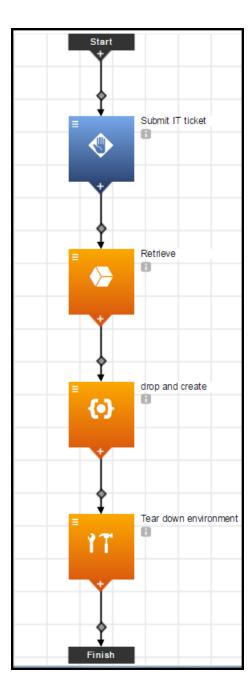
c. Enter the parameter values for the parameters for utility, and click \mathbf{OK} .

Example:



13. Repeat the previous steps to author component processes for other tiers in the application.

This is the component process:



- You can get details about each step by clicking i under the step name.
- It has one manual step, which is differentiated from the automated steps by color.

Manual steps are blue while automated steps are orange.

• Each step has an icon indicating the step type.

The first step is a manual step.

The second is based a component operation.

The third step is a command or script.

• These steps are deployed in series.

The second step can only start after the first step is completed.



- Clicking displays menu where you can select options:
 - Details-View step details
 - Properties-View step properties.
 - · Access Control-View access control settings.
 - Add Connector-Add another connector to the step.
 - Track Changes-View the Change History.

Authoring Application Processes

At the application (parent) level, you author application processes. When deploying an application, the application process that you select is executed to orchestrate operations against the application. An application process step can be defined as a call to an ElectricFlow plugin, procedure, or component process as well as a direct command, script, a manual task, or a utility function (a higher-order operation than a plugin).

This example shows how to author an application process with process branching and a rollback step.

In the Applications Visual Editor for the Upgrade application:

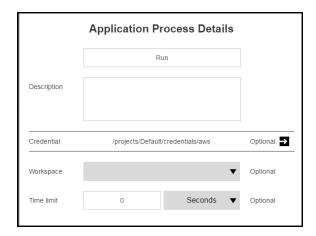


- 1. Click to start authoring an application component process.
- 2. In the Application Process Details dialog box, enter the following details:
 - Enter "Run" in the Name field.
 - Enter an optional description in the Description field.
 - Keep the selection as **Deploy** in the **Type** field. This value determines track what is built, tested, and deployed during the application deployment. The default is **Deploy**.
 - **Deploy**–Enables Inventory Tracking. The ElectricFlow server tracks artifacts deployed to environments.
 - **Undeploy**—After the first successful job step in a component process with this setting, the automation platform removes the environment inventory record.
 - Other—Disables Inventory Tracking.



- Specify who can run the application process in the Credential field. Clicking displays a
 dialog box where you select a project and then a credential (user name and password) in
 that project for the process, which are created and managed at the platform level. You can
 only impersonate one credential. To set the credential type, see Adding Credentials.
- Select a workspace from the drop-down list of workspaces in the ElectricFlow platform in the Workspace field, where the files and results of the job step are stored (optional).
- Set a time limit in the **Time limit** field, the maximum amount of time that the step is allowed to
 run. After the time specified, the step is aborted. Enter the time, and select the unit of time:
 seconds, minutes, or hours. Go to the Job Step Details on page 533or information how time
 limits work in procedure job steps.

Example:



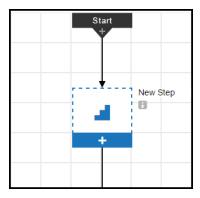
3. Click OK.

The Application Process Visual Editor opens.

The process model shows a process with **Start** and **Finish** steps with a **New Step** in between them.

 $\textbf{Note:} \ \textbf{Clicking i} \ \textbf{under the step name displays the details about the process step.}$

Example:



4. Click + in the first step.

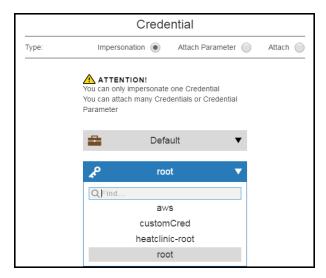
- 5. Define the first step in the Application Process Step dialog box :
 - a. Enter "Check" in the **Name** field.
 - b. Enter "Check the database version" in the **Description** field.
 - c. Use the ${f Continue}$ Running setting in the ${f On}$ ${f Error}$ field. Go here for more information.
 - d. Use the ${f all}$ setting in the ${f Run}$ if field. Go here for more information.

e. Add the root user in the Credential field. Only root user can run the application process.

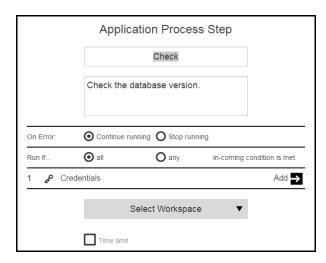


Clicking displays a dialog box where you select a project and then a credential (user name and password) in that project for the process step, which are created and managed at the platform level. You can only impersonate one credential. To set the credential type, see Adding Credentials.

Example:



This shows the process step has one credential.



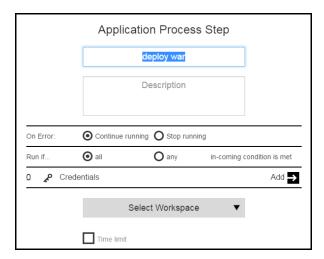
f. Click Next.

6. To define the first step as a command or script, select **Command** as the **Step Type**, and enter the command or script.

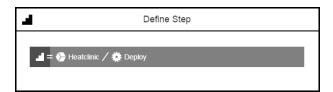
Example:



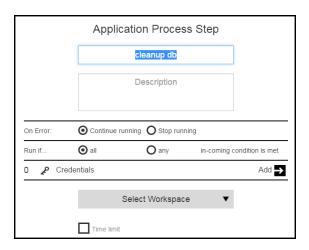
- 7. Click + on the bottom of the first step to add a new step below it.
- 8. Click + in the second step.
- 9. Define the second step in the Application Process Step dialog box as follows, and click Next:



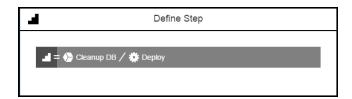
10. To define the second step in the Application Process Step dialog box based on a component, select Components > Heatclinic > Deploy. The step is defined by the component Deploy process for the Heatclinic component.



- 11. Click + on the bottom of the first step to add a third step below it. This step will run in parallel to the second step.
- 12. Define the third step as described in the previous steps.



The step type is Components. Select Components > Cleanup DB> Deploy.

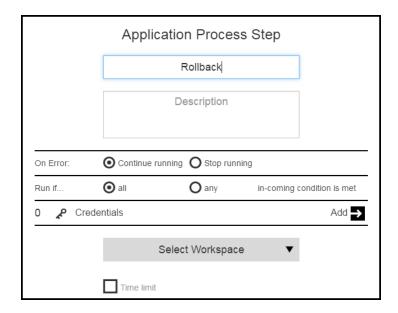


The step is defined by the component Deploy process for the "Cleanup DB" component.

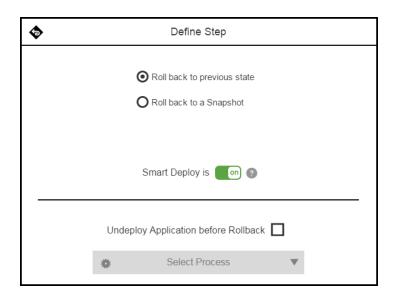
13. Repeat the previous steps to add more steps to the application process.

Note: Starting in ElectricFlow 6.3, you can also define an application or component process step as a utility function, a higher-order operation than a third-party plugin, for application modeling. Go here for example of how to define a utility step.

- 14. After the "start server" step, add a rollback step:
 - a. Click + at the bottom of the "start server" step.
 - b. Click + in the New Step to define it.
 - c. Click Next.
 - d. Select Rollback as the Step Type.
 - e. Define the rollback process.



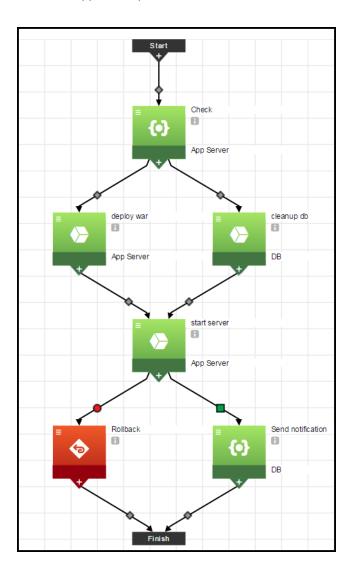
f. Click Next.



g. Click OK.

- 15. Complete the process by creating success and failure branches after the "start server" step.
 - Click on the connector between the "start server" step and the Rollback step, and change the condition from **Always** to **Failure**.
 - Click on the connector between the "start server" step and the "Send notification" step, and change the condition from **Always** to **Successful**.

This is the application process called Run:



- You can get details about each step by clicking i under the step name.
- It has one rollback step, which is differentiated from the automated steps by color.

The rollback step is orange-red while automated steps are green.

· Each step has an icon indicating the step type.

Two steps are defined by commands or scripts.

Three are defined by components.

One is for rollback.

• Some steps are deployed in parallel.

The second and third steps are both deployed after the first step.

• After the "start server" step is completed, the next step depends on the its outcome:

If the step completed successfully, the "Send notification" step is run where interested users and groups receive emails that the upgrade was completed successfully.

If the step failed, the rollback step is run and the step rolls back to its previous state.



- Clicking displays menu where you can select the options:
 - **Details**-View step details
 - Properties-View step properties.
 - Access Control-View access control settings.
 - Add Connector-Add another connector to the step.
 - Track Changes-View the Change History.

Creating Environments

Environments are logical grouping of machines to which various applications can be mapped and deployed. Using environments and templates, you can model the infrastructure and the middleware available for various applications. Environments consist of logical groups of resources with a similar function or role called environment tiers. An environment tier can be a group of machines with a similar function or role, such all the web servers or application servers or database servers for an environment. Resources are actual target end point machines (such as physical servers), virtual machines, or mobile devices.

Environments can be static, dynamic, or hybrid. A *static environment* has resources that are already provisioned and managed at the platform level. Each resource has its own logical name to identify it from the other resources in the system. It also can be assigned to one or more resource pools or to a zone (a collection of agents). Several resources can correspond to the same physical host or agent machine. Resources can also be configured as *standard* or *proxy*. Standard resources are machines running the ElectricFlow agent on a supported agent platform while proxy resources (agents and targets) are on remote platforms or hosts that exist in your environment and requires SSH keys for authentication. The ElectricFlow agent does not need to run on the remote platform or host. Go to Resources on page 595 for more information about to create, configure, and manage resources.

This example shows to how model a static environment to which the application will be deployed. For information about dynamic environments and how to model them, go to Deploying Applications in Dynamic Environments on page 142.

In the Environments List:

- 1. Click **Add +** (in the upper right corner) to add an environment.
- 2. In the **New** dialog box, select **Environment > Create New**.
- 3. Select a project to which the new environment will belong in the **Select Project** field, enter a description of the environment, and click **OK**.

Note: You can include hyperlinks as part of an object description for any ElectricFlow object.

Example:



The Environments Visual Editor opens.

4. Define an environment tier and assign resources it.



In Tier 1, click and select **Details** to edit the details in the **Environment Tier Details** dialog box.

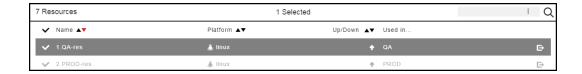
- a. Rename the application tier to "App" and click **OK**.
- b. To define the resource:

Click + under the Resource.

Click Add resources.

Select the resource that you want to assign to the environment tier, and click **OK**. A resource is available if it is enabled.

Example:



The Environments Visual Editor now has an environment tier called "App" with one resource.

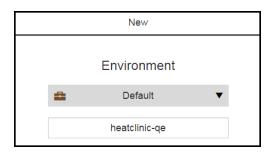


- 5. In the Environments List, click Add + (in the upper right corner) to add another environment.
- 6. In the New dialog box, select Environment > Create New.

7. Select a project to which the new environment will belong in the **Select Project** field, enter a description of the environment, and click **OK**.

Note: You can include hyperlinks as part of an object description for any ElectricFlow object.

Example:



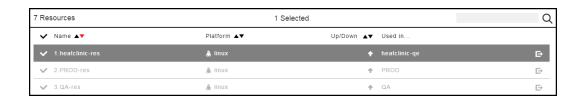
The Applications Visual Editor opens.

8. In the Environments Visual Editor, define the environment tier and assign resources to it.



- a. In Tier 1, click and select **Details** to edit the details in the **Environment Tier Details** dialog box.
- b. Rename the application tier to "mysql" and click **OK**.
- c. To define the resource:
 - i. Click + under the Resource.
 - ii. Click Add resources.
 - iii. Select the resource that you want to assign to the environment tier, and click **OK**. A resource is available if it is enabled.

Example:



The Environment Visual Editor now has an environment tier called "mysql".

+=

9. In the Environments Visual editor, click to add another tier to the environment, and assign resources to it. The new tier is called "Tier 1."

Repeat the procedure in this step to define the environment tier and the resources in it. This tier will be called "tomcat" with the resource called "heatclinic-res" assigned to it.

The Environments Visual Editor now has two environment tiers, each with one resource.

Example:



Defining Tier Maps

Before deploying applications, a tier map must be defined associating the application with an environment.

In the Applications Visual Editor for the application that you want to deploy:



1. Click

The Tier Map dialog box opens.

- 2. Select Environment.
- 3. From the list of environment that you have permission to view, select an environment. You can search for environments by project name or enter the search criteria in the **Search** field.

Example: Select QA.

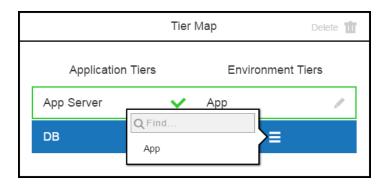


4. Map the application tiers to the environment tiers:



a. For each application tier, click in the corresponding **Environment Tiers** column, and select an environment tier. Select one from the list or enter the search criteria in the **Search** field.

Example:



b. Click **OK** when each application tier have been mapped to an environment tier.

These are examples of tier maps.

• For the QA environment:



• For heatclinic-qe environment:



Deploying and Troubleshooting Applications

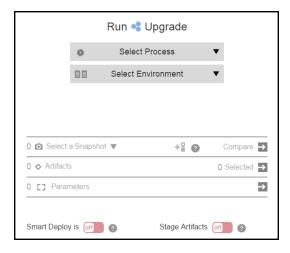
This example shows how to deploy the application multiple times with different runtime conditions.

First Run



In the Applications List, clicking runtime settings.

and selecting **New Run** opens the dialog box where you can set the



- In the Select Process field, select Run.
- In the Select Environment field, select QA.
- You can select a snapshot in the **Select a Snapshot** field if you have one or more.
- For the first run, all artifacts will be deployed.
- Smart Deploy and Stage Artifacts are enabled.



Click **OK** to start the application run.

The Applications View Run page shows the progress of the application run:

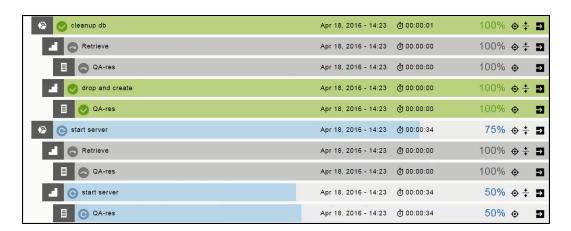
- The blue steps are running (in progress).
- The green steps were completed successfully.
- The red steps failed.



Clicking opens the Job Details page where you can begin the troubleshoot the step.

• The gray steps were skipped.





Subsequent Runs

After the first application run, you can deploy the **Run** process to the **QA** environment with other deployment options.

Based on the Last Run



Clicking and selecting **Last Run** opens the dialog box where you can set the runtime settings. It displays the runtime settings from the previous. You can deploy the application using the same settings as the previous run or you can modify one or more them for the current run.

Based on Schedules



Clicking and selecting **Schedule** opens the list of schedules for the application. If there are no schedules, the list looks like this:



Clicking There are no Run Schedules. Add one (or Add +) opens the Run Schedule dialog box.

- Select Create new to create a new schedule, and then enter the schedule details.
- Select from Previous Runs to create a schedule based on a previous application run.

A list of the last five application runs appears:



Select an application run, and then enter the schedule details.

In the Run Schedule Details dialog box, enter the schedule details:

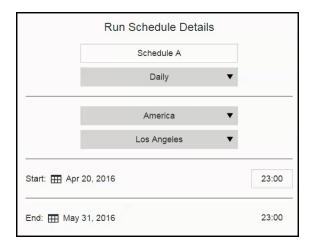
- In the **Schedule name** field, enter the schedule name.
- In the Frequency field field, click the down arrow and select Once, Daily, Weekly, or Monthly.
- In the Country field, click the down arrow and select a value in the drop-down list.
- In the City field, click the down arrow and select a value in the drop-down list.



• In the **Start** field, click if you want to change the start date, and click in the **hh:mm** box to set the start time using the 24-hour clock.



In the End field, click and select the end date from the pop-up calendar. The end time is the same as the start time.

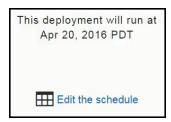


• Click Next.

The dialog box where you can set the runtime settings appears. You can also modify the runtime settings at this time.

• Click OK.

This message appears:



Click OK to close it.

When you view the list of schedules for the application again, the schedule that you created appears in the list.



Runtime Settings

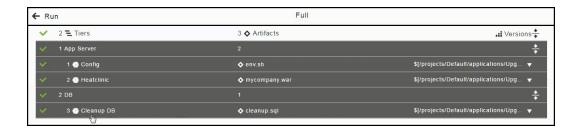
In the dialog box where you can set the runtime settings, the Artifacts field looks like this for a full run:



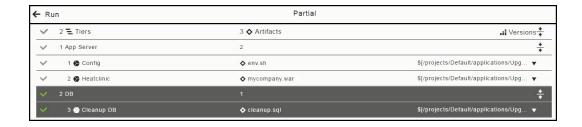
• Partial deploy with smart deploy and artifact staging are enabled:



Clicking opens the list of artifacts to be deployed in the application process.



If you want to only deploy the DB artifact (cleanup.sql), click in the "1 App Server" column to deselect (disable the selection of) the Config and Heatclinic components (based on the env.sh and mycompany.war artifacts), and click **OK**.

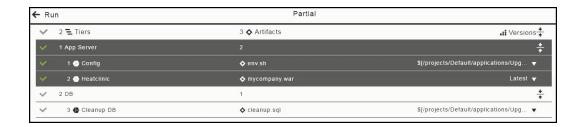


The Artifacts field now looks like this for a partial run:



• Partial deploy with specific artifact versions and with smart deploy and artifact staging are enabled:

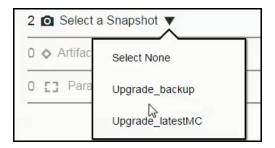
In the list of artifacts to be deployed in the application process, click in the "2 DB" row to deselect (disable the selection of) the "Cleanup DB" component (based on the cleanup.sql artifact), select the **Latest** version of the mycompany.war artifact, and click \mathbf{OK} .



The **Artifacts** field now looks like this for the partial run:



Snapshot with partial deploy and with smart deploy and artifact staging are enabled:
 In the Select a Snapshot field, click the down arrow and select the Upgrade_latestMC snapshot.



In the list of artifacts to be deployed in the application process, click in the "1 Config" row to deselect (disable the selection of) the Config component (based on the env. sh artifact) and click **OK**.



The **Select a Snapshot** and **Artifacts** fields now look like this for the partial run:

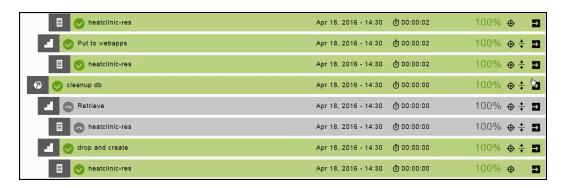


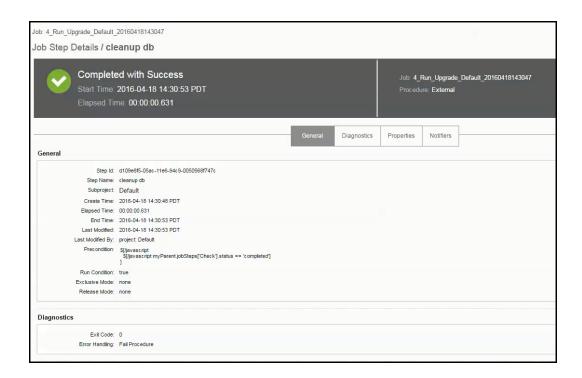
Viewing the Real-Time Progress of Deployments



This example shows the job details for a specific step. Clicking Details page:

in the "cleanup db" step opens the Job





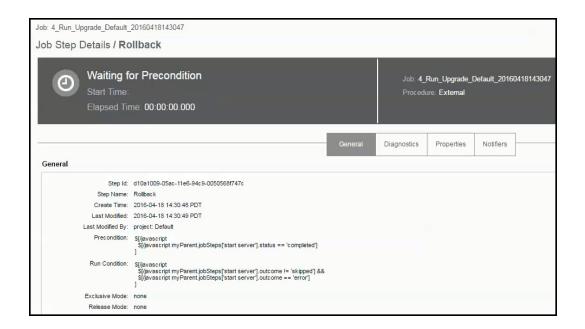
This example shows that the rollback step has not run yet:





Clicking in the rollback step opens the Job Details page, showing that the status is *Waiting for Precondition*. This step will not run until the precondition is met.





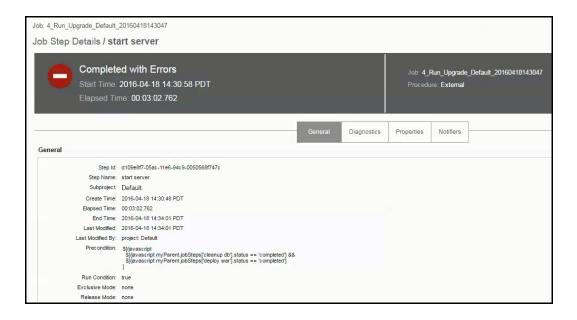
Troubleshooting Deployments

This example shows that the "start server" step was Completed with Errors.





Clicking in the "start server" step opens the Job Details page.



Clicking in the Diagnostics tab can provide more information about the errors.

Example: Manual Steps with Runtime Parameters

This topic describes manual steps with parameter inputs and how the parameter inputs influence the process flow to take different paths. You can define manual steps in an application process or component process.

In the following example, an application process called Deploy has the following steps. When deployed, this process updates the website for an online store.

- The first step is a command step to retrieve a .war file that will be deployed to the website.
- The second step is a manual step to log into a terminal where you can enter CLI commands to update the website.

If you have admin privileges, you can update the home (parent) page.

If you do not have privileges to update the home page, you can update only the product (child) pages.

Depending on your privileges, the parameter input to the manual step is either *pass* (you have the privileges to update the home page) or *fail* (you can only update the product pages).

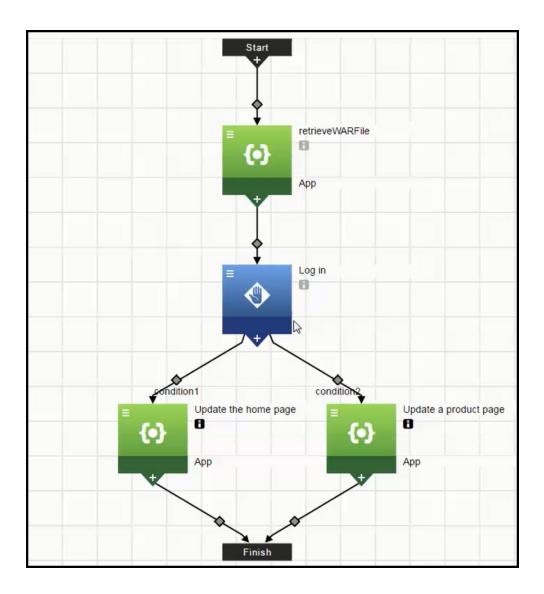
- The parameter inputs to the manual step are the preconditions to the third and fourth steps.
 - condition1 is fulfilled if the input to the manual step is pass.
 - condition 2 is fulfilled if the input to the manual step is fail.

These preconditions determine the step that is executed after the manual step.

• The next step is either the third or fourth step.

If the parameter input for the manual step is *pass*, the process continues to the "Update the home page" step.

If the input is fail, the process continues to the "Update a product page" step.



Process Flow Details

After starting the deployment, the .war file is automatically retrieved in the first step.

In the second step, email notifications are sent to one or more assignees. One of the required assignees performs the manual task, which is to log into a terminal and enter the appropriate CLI commands to set up the website update. After clicking on the link in the email, the assignee goes to the ElectricFlow UI and provides input to the manual process step using the pass fail value parameter.

- If the assignee has admin privileges, the value of the pass_fail_value parameter is pass. The process continues to the "Update the home page" step.
- If the assignee does not have admin privileges, the value is *fail*. The process continues to the "Update a product page" step.

In the third and fourth steps:

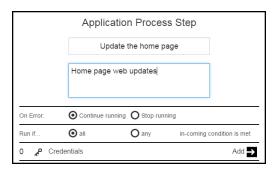
• The On Error setting is Continue running.

If the process encounters an error while the web page is being updated, it completes the step with errors and then continues to next step.

The default is **Stop running**, where the process will abort if the process encounters an error.

• The **Run** if setting is **all** (the default). All the preconditions to the step have to be met before the process goes to the next step.

If any is selected, only one of the conditions has to be met before the process continues running.



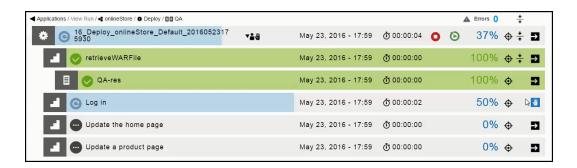
Deployment Scenarios

These are possible deployment scenarios.

Assignee has admin privileges and the deployment is completed successfully

The home page is updated in this scenario.

When the Deploy application process is deployed to the QA environment, you can view the progress of the deployment in the "Applications / View Run" page.



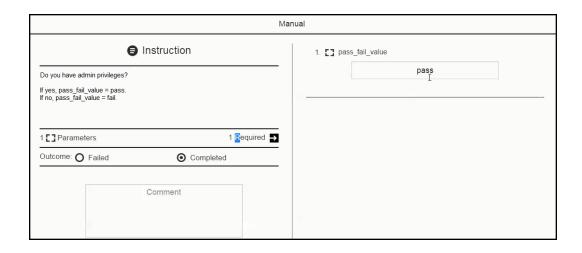




Clicking opens the Manual dialog box. Clicking parameters on the right side of the dialog box.

in the Parameters row displays the required

If you enter *pass* in the **pass_fail_value** field, select **Completed** as the **Outcome**, and click **OK**, the process continues to the "Update the home page" step.



The parameter value (*pass*) is used in the precondition to determine the next step in the process flow. It is also passed to the next step.

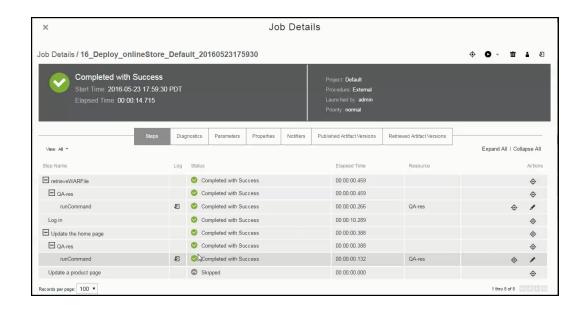
The step to update the home page is successfully completed. Notice that the "Update a product page" step is skipped.

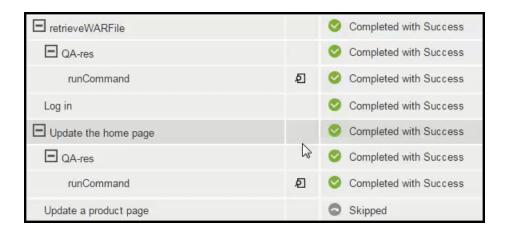




In the top row, clicking

takes you to the Job Details page.





Clicking in the runCommand row takes you to the log file for that step, which displays the parameter value from the manual step that was passed to the "Update the home page" step.

Job: 16_Deploy_onlineStore_Default_20160523175930

Workspace File / runCommand.cd53b36e-214a-11e6-b9d7-0050568f7ab6.log

pass

Assignee does not have admin privileges and the deployment is completed successfully

A product page is updated in this scenario.

When the Deploy application process is deployed to the QA environment again, you can view the progress of the deployment in the "Applications / View Run" page.



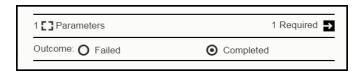
When the process reaches the second step, clicking

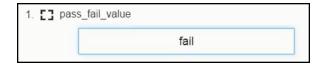
opens the Manual dialog box.



Clicking in the Parameters row displays the required parameters on the right side of the dialog box.

If you enter *fail* in the **pass_fail_value** field, select **Completed** as the **Outcome**, and click **OK**, the process continues to the "Update a product page" step.





The parameter value (*fail*) is used in the precondition to determine the next step in the process flow. It is also passed to the next step.

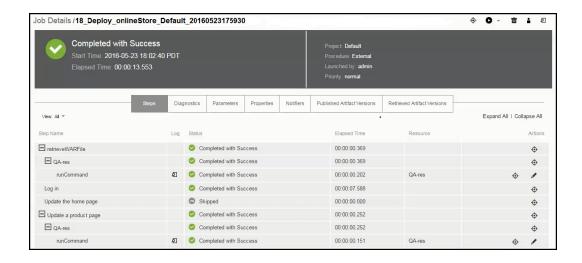
In this deployment, both the Deploy process and the step to update product page are completed successfully.





In the top row, clicking

takes you to the Job Details page.



Ð

Clicking in the runCommand row takes you to the log file for that step, which displays the parameter value from the manual step that was passed to the "Update the home page" step.

Job:18_Deploy_onlineStore_Default_20160523175930

Workspace File / runCommand.3e8ef99f-214b-11e6-a00e-0050568f7ab6.log

fail

Assignee does not have admin privileges and the deployment is completed with errors

A product page is updated in this scenario.

When the Deploy application process is deployed to the QA environment again, you can view the progress of the deployment in the "Applications / View Run" page.



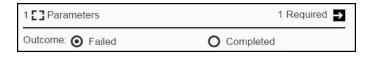
When the process reaches the second step, clicking

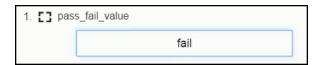
opens the Manual dialog box.



Clicking in the Parameters row displays the required parameters on the right side of the dialog box.

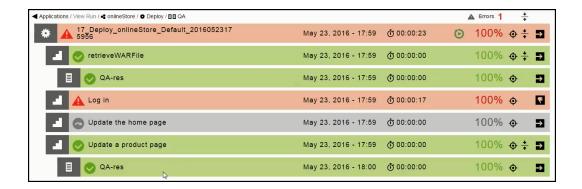
If you enter *fail* in the **pass_fail_value** field, select **Failed** as the **Outcome**, and click **OK**, the process continues to the "Update a product page" step.





The parameter value (*fail*) is used in the precondition to determine the next step in the process flow. It is also passed to the next step.

In this deployment, the step to update the product page is completed successfully even though the Deploy application was completed with errors. The error is in the second step where the outcome is **Failed**.





In the top row, clicking

takes you to the Job Details page.



Clicking in the runCommand row takes you to the log file for that step, which displays the parameter value from the manual step that was passed to the "Update the home page" step.

```
Job: 17_Deploy_onlineStore_Default_20160523175956

Workspace File / runCommand.3e8ef99f-214b-11e6-a00e-0050568f7ab6.log
fail
```

Setting Up Custom Parameters for Application and Component Processes

Parameters are used throughout the release planning and management process. They provide an easy way to model releases, pipelines, and applications. They can be inputs to procedures, workflows, applications, pipelines, and releases. You can use parameters in place of predefined credentials, property paths, user or group names for notifications, runtime conditions. Parameters are created and managed independent of any object at the platform level.

These examples show how to add custom parameters and setting the parameter labels in the GUI.

Adding Custom Parameters

In the Application Process or Component Process Visual Editor:



- 1. Click on the upper right corner of the Visual Editor.
- 2. Click Parameters.

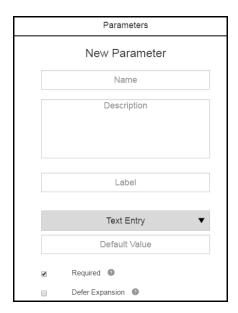
The Parameters dialog box opens.

Example:



3. Click **Add +** or if no parameters have not yet defined for the object, click **There are no Parameters yet. Add one +** to add a parameter.

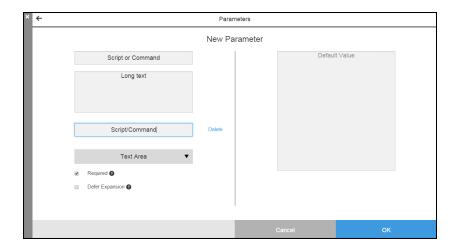
The New Parameter dialog box opens.



- 4. Enter the following information:
 - Parameter Name-Name of the parameter.
 - Description-This is optional.

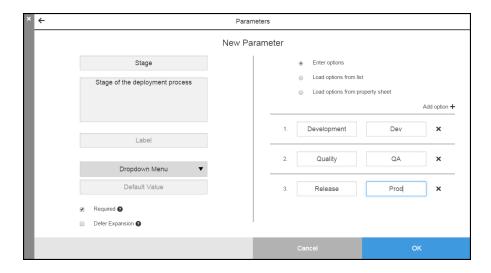
Note: You can include hyperlinks as part of an object description for any ElectricFlow object.

- Label-Label that appears on the GUI. This is optional.
- Text Entry—How the user selects the value for this parameter. Valid options are Text Entry, Test Area, Dropdown Menu, Radio Selector, Checkbox, or Credential.
- Default Value—Default value for the parameter. This is optional.
- Required—If the parameter is required, select the Required check box.
- **Defer Expansion**—If the parameter value contains \$[] and you want ElectricFlow to interpret it literally (not as a parameter reference), select the **Defer Expansion** check box.
- Depending on the parameter type that you select, other fields appear in the dialog box. Go to the following appropriate step to complete setting up the parameter.
- 5. If you select **Text Entry** as the parameter type, the **Default Value** field appears. You can enter a default value but is not required, and then click **OK**.
- 6. If you select **Text Area** as the parameter type, the **Default Value** field appears. You can enter a default value but is not required, and then click **OK**.

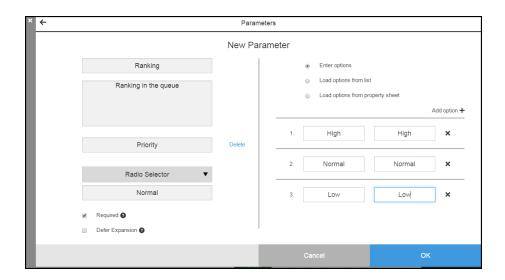


- 7. If you select **Dropdown Menu** as the parameter type, the **Default Value** field and ways to add the menu options appear. You can enter a default value but is not required.
 - a. On the right side of the dialog box, enter the menu options.
 - b. Click OK.

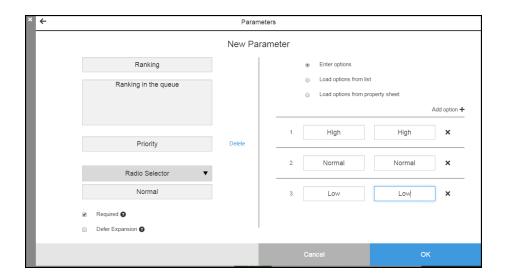
Example:



- 8. If you select **Radio Selector** as the parameter type, the **Default Value** field and ways to add the menu options appear. You can enter a default value but is not required.
 - a. On the right side of the dialog box, enter the menu options.
 - b. Click OK.



- 9. If you select **Checkbox** as the parameter type, the **Default Value** field and values for the check box appear. You can enter a default value but is not required.
 - a. On the right side of the dialog box, enter the values.
 - b. Click OK.



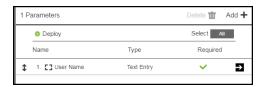
10. If you select **Credentials** as the parameter type, the **Default Value** field appears. You can enter a default value but is not required, and then click **OK**.

Example:



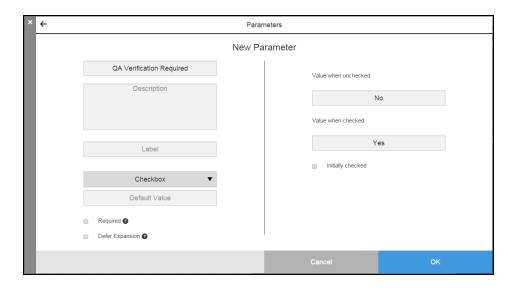
After completing these steps, the **Parameters** dialog box opens and shows the new parameter in the list.

Example:



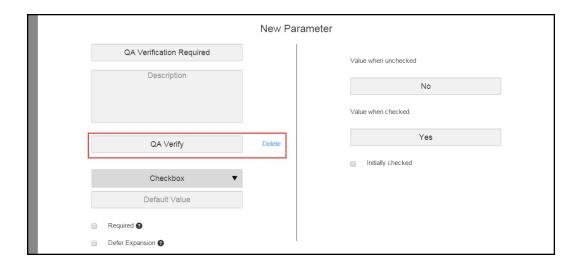
Setting and Modifying the Parameter Label

In the **New Parameter** dialog box:



1. Add a label in the **Label**. This is optional.

Example:

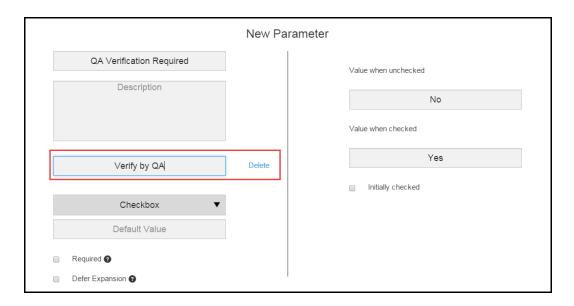


If you enter a label, it appears in the UI form when you deploy the application. If you do not enter a label, the parameter name appears in the UI form when you deploy the application.

2. Click OK.

- 3. To modify the label:
 - a. Open the **New Parameter** dialog box.
 - b. Clear the **Label** field.
 - c. Enter a new label.

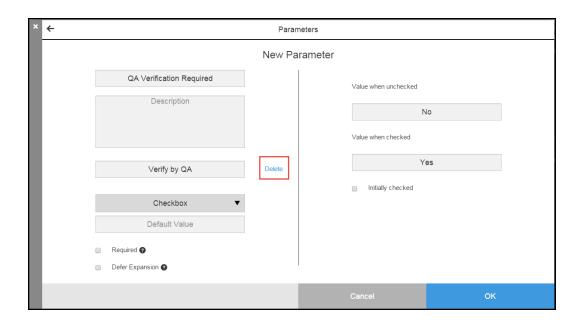
Example:



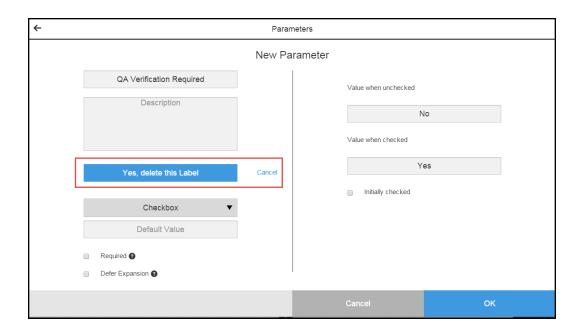
d. Click **OK** to save the change.

- 4. To delete the label:
 - a. Click Delete next to the Label field.

Example:



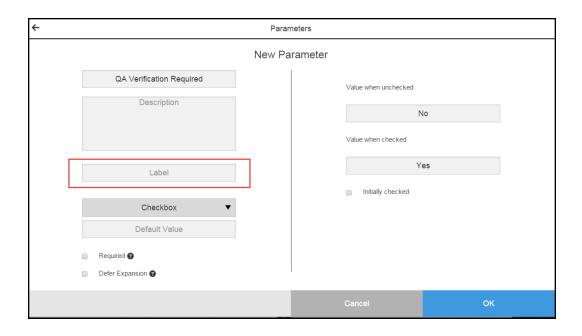
Notice that the text in the **Label** field changes.



b. Click in the Label field.

The label disappears.

Example:



c. Click OK.

The **Parameters** dialog box opens. The parameter name now appears in the name column.

Looking Up Parameters in Application and Component Processes

To apply parameters to an application or component process step, starting in the Application Process or a Component Process Visual Editor:

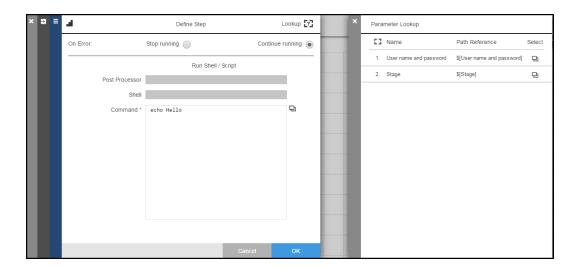
- 1. Add a new step to the process.
- 2. Define the process step in the **Define Step** dialog box.

3. When you define the process step with a plugin (**Plugins**), command or script (**Command**), or project (**Procedure**), click the **Lookup** button to open the **Parameter Lookup** dialog box.

Example:



The Parameter Lookup dialog box opens.



4. Choose a parameter and click the **Copy** button to copy the path reference.

A message appears in the row : < Parameter Name > has been copied.

The **Define Step** dialog box now has a Parameter field.

- 5. Click in the Parameter field and paste the path reference that you copied in it.
- 6. Repeat the previous two steps to apply another parameter to the process step.
- 7. Click OK.

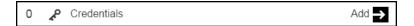
Attaching Credentials to Application and Component Processes

This topic describes how to add credentials and parameters to a component processes and apply custom parameters to them. These parameters are required at runtime when the application is deployed.

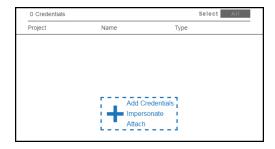
When defining an application or component process step:



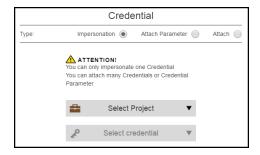
1. Click in the Credentials row in the dialog box where you can set the runtime settings.



The Credentials dialog box opens and shows the list of credentials for this process step. If there are no credentials, the dialog box looks like this:



2. Click + Add Credentials Impersonate Attach to add a credential.

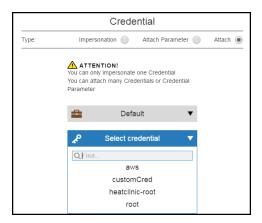


3. Select **Attach** to attach a credential to the process step. This credential is created and managed at the platform level.

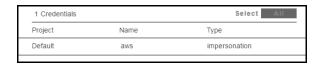
The other ways to add credentials to a process step are:

- If you want to impersonate one credential, select **Impersonation**, select the project to which the credential belongs, and the credential that you want to use for impersonation.
- If you want to attach a credential parameter to the process step, select **Attach Parameter**. Click in the **Select credential parameter** field to set the parameter.

4. Select the project to which the credential belongs in the **Select Project** field, the credential in the **Select Credential** field, and click **Next**.

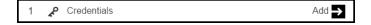


The list of credentials now shows the credential you added:



5. Click Next.

The dialog box where you can set the runtime settings now shows that one credential is attached to the process step.



6. Repeat the previous steps to attach another credential or use impersonation.

IMPORTANT:

When you impersonate a credential, make sure that the impersonated user has the absolute path to the bin directories in the \$PATH environment.

If you define a process step with a command, you must enter the absolute path in the **Post Processor** and **Shell** fields in the "Define Step" dialog box.

Deploying Applications in Dynamic Environments

This section describes dynamic environments and how to model them.

Environments can be static, dynamic, or hybrid. Any *dynamic environment* is automatically spun up when the application is deployed. Environments can have a combination of static and cloud resources that will be provisioned and configured dynamically. Go here for more information about static environments and how to model them.

Using dynamic environments allows you to:

- Provide ways to optimize how cloud resources are used.
- Reuse provisioned resource pools on an on-demand basis.
- Track how provisioned cloud resources are used.
- · Provide the status of the provisioning process.
- Verify the credentials of cloud resources before provisioning them.

You provision cloud resources by creating a resource template and adding it to an environment template. When the application deployment that includes a dynamic environment begins, the environment is spun up at runtime with the required cloud resources, and the application is deployed there.

Resource templates allow you to

- Define the cloud provider details to provision of machines at the appropriate time and the configuration management tools to install the appropriate middleware.
- Configure the resources for optimal performance.

About Resources

Each resource in an environment has its own logical name to identify it from the other resources in the system. It also can be assigned to one or more resource pools or to a zone (a collection of agents). Several resources can correspond to the same physical host or agent machine. Resources can also be configured as standard or proxy. Standard resources are machines running the ElectricFlow agent on a supported agent platform while proxy resources (agents and targets) are on remote platforms or hosts that exist in your environment and requires SSH keys for authentication. The ElectricFlow agent does not need to run on the remote platform or host. Go to Resources on page 595 for more information about to create, configure, and manage resources.

Examples

Infrastructure in ElectricFlow is modeled by configuring the cloud provider for dynamic environments. This configuration is versioned when you name and save the resource template. For an example of an Amazon EC2 infrastructure configuration, go to . For an example of an OpenStack infrastructure configuration, go here. Example: Resource Template for Infrastructure using Amazon EC2 on page 152 and Example: Resource Template for Infrastructure using OpenStack on page 159. For examples of infrastructure configurations, go to Example: Resource Template for Infrastructure using Amazon EC2 on page 152 and Example: Resource Template for Infrastructure using OpenStack on page 159.

Middleware is modeled by configuring the configuration management tool for dynamic environments. This configuration is versioned when you name and saved the resource template. For an example of a middleware configuration, go to Example: Resource Template for Middleware Using Chef on page 164.

Modeling Dynamic Environments

This is the high-level process to provision cloud resources that can be dynamically spun up during the application deployment.

1. Create and define one or more resource templates.

The resource template has the information required to spin up the resources on an on-demand basis and to provision dynamic resource pools. It has the following information:

- · Cloud provider and cloud instance details
- Configuration management settings
- Both cloud provider and configuration management settings
- 2. Create and define one or more environment templates.

When modeling an environment template, you define the environment tiers and then add resources to the tiers. You can either add static resources or a resource template to an environment tier.

When adding static resources, you can select one or more resources to add to the tier.

When adding resource templates, you can select only one to add to the tier and then enter the number of resources to provision.

Environment templates have this information:

- · Name and description
- · Environment tier details and properties
- · Resources assigned to the environment tiers
- Cloud resource details from the resource templates

Usage Guidelines and Best Practices

Review the following information before creating and using dynamic environments.

• This matrix shows the resource types that are allowed in environments and environment templates.

Resource Type	Origin	Environment Tier	Environment Template Tier	Use Rules
Static resources		Yes	Yes	You may only add them to one environment. You cannot reuse them in more than one environment.
				You can add one or more static resources to an environment tier or an environment tier template.
Static resource pool	_	Yes	No	You may only add it to one environment. You cannot reuse it in more than one environment.
				You can add only one resource pool to an environment tier.

Resource Type	Origin	Environment Tier	Environment Template Tier	Use Rules
Resource pool provisioned in the resource template	Platform	Yes	No	You may only add it to one environment. You cannot reuse it in more than one environment. You can add only one resource pool to an
Resource pool provisioned in an environment template	ElectricFlow	No	No	environment tier. You cannot reuse the resource pools provisioned in an environment template. You can only use it in the dynamic environment created by the environment template.
Resource template	-	No	Yes	You may only add one resource template to an environment template tier.

In environments and environment templates, you can create a tier with static resources, a static resource pool, or a static resource template. When you are editing the tier, ElectricFlow maintains the initial resource-type association.

For example, if you initially add static resources to an environment tier, you can add only static resources to it. You are not allowed to add a resource this environment tier.

ElectricFlow validates the information you provide for configuring and provisioning the cloud provider.

For instance, when you configure Amazon EC2 as the cloud provider, ElectricFlow validates the Service URL using the Access ID and Access Key. The configuration information that you provide is also used to retrieve and display the available options from the cloud provider service for the provisioning parameters such as the Image or the Security group to use.

However, if your ElectricFlow server is unable to access external services such as the Amazon EC2 service, you may disable this validation and the dynamic options retrieval capability for the cloud provider plugin by setting the plugin property ec_disable_dynamic_operations to 1.

Creating AMIs

Before creating a resource template, you must create an Amazon Machine Image (AMI) with a pre-installed agent for the cloud provider. In ElectricFlow 5.4, Amazon EC2 and OpenStack are the supported cloud providers.

AMIs with ElectricFlow Agents

1. Install agents on ElectricFlow server node machines.

See the following sections in the ElectricFlow Installation Guide for detailed instructions:

- System Requirements and Supported Platforms > Agent Platforms
- System Requirements and Supported Platforms > Server and Agent Compatibility
- Installing ElectricFlow > Interactive Command-line Installation Method > Express Agent Command-Line
- Installing ElectricFlow > Non-Server Platform Agent Installation Method
- Configuration > Environment Proxy Server Configuration > Configuring Proxy Agents
- 2. Create the AMI on the Amazon EC2 or OpenStack platform.

Go to http://aws.amazon.com/amazon-linux-ami/ for Amazon EC2.

Go to http://docs.openstack.org/image-guide/content/index.html for OpenStack.

3. Create a resource template in ElectricFlow.

When you set the cloud provider, you can select **Amazon** or **OpenStack** and enter information in the fields in the dialog boxes. The fields that you see depend on the cloud provider that you select.

For detailed Amazon EC2 instructions, go to Configuring Amazon EC2 as the Cloud Provider.

For detailed OpenStack instructions, go to Configuring OpenStack as the Cloud Provider.

AMIs with ElectricFlow Agents and Chef Configuration Management

1. Install agents on ElectricFlow server node machines.

See the following sections in the ElectricFlow Installation Guide for detailed instructions:

- System Requirements and Supported Platforms > Agent Platforms
- System Requirements and Supported Platforms > Server and Agent Compatibility
- Installing ElectricFlow > Interactive Command-line Installation Method > Express Agent Command-Line
- Installing ElectricFlow > Non-Server Platform Agent Installation Method
- Configuration > Environment Proxy Server Configuration > Configuring Proxy Agents
- 2. Configure a Chef server with run-lists that will be applied to your cloud resources.

Go to http://docs.chef.io/.

3. Create the AMI on the Amazon EC2 or OpenStack platform.

Go to http://aws.amazon.com/amazon-linux-ami/ for Amazon EC2.

Go to http://docs.openstack.org/image-guide/content/index.html for OpenStack.

- 4. Create a resource template in ElectricFlow.
 - a. When you set the cloud provider, you can select **Amazon** or **OpenStack** and enter information in the fields in the dialog boxes. The fields that you see depend on the cloud provider that you select.

For detailed Amazon EC2 instructions, go to Configuring Amazon EC2 as the Cloud Provider.

For detailed OpenStack instructions, go to Configuring OpenStack as the Cloud Provider.

b. When you set the configuration management tool, select Chef.

For detailed instructions, go to Configuring Chef as the Configuration Management Tool.

Creating Resource Templates

1. Go to the Resource Templates List.

If there are no defined resource templates, the Resource Templates List is empty.



- 2. Add a resource template.
 - a. If the Resource Templates List is empty, click **There are no Resource Templates yet. Add one +** or click the **Add +** button.

In the New dialog box, click Resource Template.

Example:



The **New** dialog box opens.



b. If the Resource Templates List is not empty, click the **Add +** button.

The **New** dialog box opens.

Example:



3. Click **Create New** to create a new resource template.

The **New** dialog box opens.



4. Enter a name for the environment template, select a project to which it belongs, and enter an optional description.

Example:



- 5. Click Next.
- 6. Enter the cloud provider settings in the New Cloud Provider dialog box.
 - To use Amazon EC2 as the cloud provider, go to Example: Resource Template for Infrastructure using Amazon EC2 on page 152 for detailed instructions.

IMPORTANT: Before configuring the cloud provider, you must have an Amazon Machine Image (AMI) with a pre-installed agent available.

- To use OpenStack as the cloud provider, go to Example: Resource Template for Infrastructure using OpenStack on page 159for detailed instructions.
- 7. (Optional) Enter the configuration management settings in the New Configuration Management dialog box..

To use Chef as the configuration management tool, go to Example: Resource Template for Middleware Using Chef on page 164 for detailed instructions.

After completing these steps, the Resource Template List now shows the template that you created. The breadcrumb in the upper left corner shows the object hierarchy.



By default, the resource templates for all projects are displayed.

If you want to see only the objects for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

If there are no matches, a message appears stating that there are no resource templates in the selected projects.

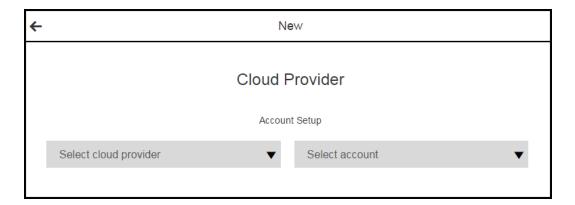
Example: Resource Template for Infrastructure using Amazon EC2

This section describes how to configure Amazon EC2 as the cloud provider for a resource template.

IMPORTANT: Before configuring the cloud provider, you must have an Amazon Machine Image (AMI) with a pre-installed agent available.

In the **New Cloud Provider** dialog box, "account" refers to the cloud provider setup. Select one of these options to enter the setup information:

- Click **Select cloud provider** to enter the cloud provider settings. Go to **Select Cloud Provider** on page 152 for the detailed instructions.
- Click **Select account** if you want to use an existing plugin configuration. Go to Select account on page 158 for the detailed instructions.

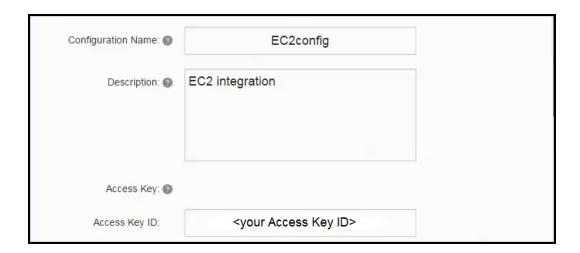


Select Cloud Provider

To enter the cloud provider settings:

1. Click **Select cloud provider > Amazon**.

- 2. Enter the following values:
 - Configuration Name—A unique name for the EC2 connection. (Required)
 - **Description**–A description for this configuration. The default is *EC2 integration*.
 - Access Key ID
 –The Access Key ID that is required for communicating with Amazon EC2.
 The configuration stores this Access Key ID in the User Name field of the credential.



- **Secret Access Key**—The Secret Access Key that is required for communicating with Amazon EC2. The configuration stores this value in the **Password** field of the credential.
- Retype Secret Access Key

 —The Secret Access Key in the Retype Password field of the
 credential.
- Service URL—The service URL for the EC2 service. The default is https://ec2.amazonaws.com. (Required)

When you provide the Access Key values and the Service URL, the system will validate that the provided values are valid for connecting to Amazon EC2.

- Resource—The resource or pool on which the integration steps can run. The default is local.
- Workspace—The workspace to use for resources dynamically created by this configuration.
 The default is default.
- Attempt Connection?—If the check box is selected, the system tries to connect to Amazon
 EC2 using the provided Access Key values and the Service URL to check credentials before
 creating the configuration.

• **Debug Level**—The debug level for the output. The default is 1.

The possible values are

- **0** Errors only.
- 1 Normal headers and responses.
- **2+** Debugging information included.



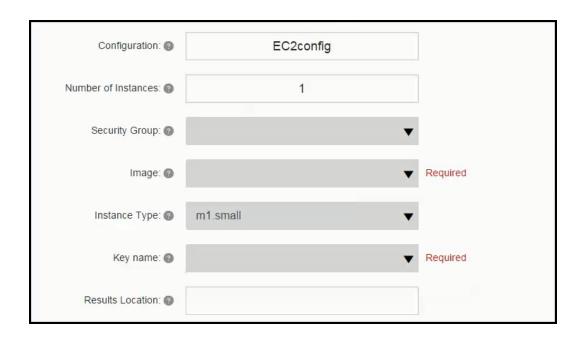
3. Click Next.

The **Provision Parameters** fields appear.

The Configuration field is prepopulated with the name you entered in the previous dialog box.

ElectricFlow tries to retrieve the provision parameters for this configuration:





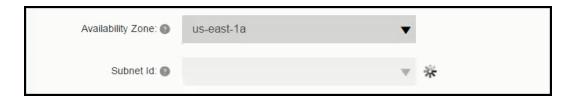
- 4. Enter the Provision Parameters settings:
 - Number of Instances—The number of instances to start. The default is 1.
 - Security Group-The default is default.
 - Image—Select the EC2 image to deploy. The drop-down menu displays the images that you own in your current region. (Required)



- Instance Type—You get the list of instance types at runtime. The default is Small(m1.small).
- Key name—Select the name of the key pair to use. (Required)



- Results Location—Location where the output properties will be saved. If not specified, it defaults to /myParent/parent.
- User Data-Extra user data to pass to the API runInstances procedure.
- Availability Zone
 –The zone where the instance will be created. (Required)
 Selecting a zone causes ElectricFlow to find a subnet in the selected zone.



Resource Pool

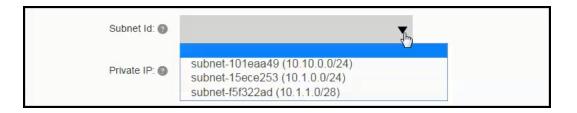
—The name of the resource pool. If you enter a name, a new resource pool is
created, and new resources are added to the pool for the instances created in Amazon EC2.



• Commander Workspace—If you specify a resource pool name in Resource Pool field, this is the workspace that will be used when creating the resource.

• **Subnet Id**—The ID of the subnet in which the instances are launched. This parameter is used with VPCs.

After selecting the zone where the instance will be created, you can select a subnet in that zone.



- Private IP—The private IP address to assign. If you specify a value, it must be from the IP
 address range of the subnet. If no value is specified, an available IP address from the IP
 address range of the subnet is selected. This parameter is used with VPCs
- Use Private IP for subnet?—If this check box is selected, the private IP address assigned to the instance will be used for the newly created resource. This is enabled by default.
- Resource Port-If you specify a resource pool name in 'Resource Pool' field, this is the port
 that will be used when creating the resource. If no value is specified, port 7800 will be used
 by default when creating the resource.
- Resource Zone Name

 —The zone to which the newly created resources will belong. The
 default zone is default.
- 5. Click Next to enter the configuration management settings.
- 6. Go to this step in Creating Resource Templates on page 148 to continue creating the resource template.
- 7. If you do not want to enter the configuration management settings, click **OK**.

The Resources Templates list now shows the resource template that you created.

Select account

Select the **Select account** option to use an existing plugin configuration. ElectricFlow automatically retrieves the cloud provider settings and puts them in the appropriate fields.

- 1. Click Select account
- 2. Select an existing plugin.
- 3. Click Next.

The dialog box shows for the settings for the selected the plugin configuration.

- 4. If you want to edit any of these settings, make your changes and then enter the password to save your changes.
- 5. Click **Next** to enter the configuration management settings.
- 6. Go to this step in Creating Resource Templates on page 148 to continue creating the resource template.
- 7. If you do not want to set the configuration management settings, click **OK**.

The Resources Templates list now shows the resource template that you created.

Example: Resource Template for Infrastructure using OpenStack

This section describes how to configure OpenStack as the cloud provider for a resource template.

Prerequisites and Configuration Guidelines

Review the following information before creating a resource template using OpenStack as the cloud provider:

• Prerequisite:

IMPORTANT: Before configuring the cloud provider, you must have an image with a pre-installed agent on OpenStack.

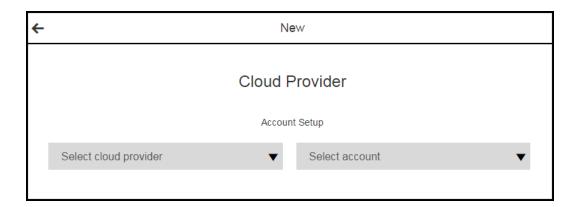
 When creating a resource template using OpenStack as the cloud provider, only one security group can be selected.

To select more than one security group, the workaround is to disable the <code>serveroptions</code> flag for the <code>security_groups</code> parameter in the EC-OpenStack plugin. In the plugin procedure called <code>_DeployDE</code>, edit the XML form as follows:

- 1. Go to the Custom Procedure Properties section.
- 2. Click ec_parameterForm.
- 3. In the Edit Property / ec_parameterForm dialog box, set the security_groups formElement called serverOptions to 0.

In the **New Cloud Provider** dialog box, "account" refers to the cloud provider setup. Select one of these options to enter the setup information:

- Click **Select cloud provider** to enter the cloud provider settings. Go here for the detailed instructions.
- Click Select account if you want to use an existing plugin configuration. Go here for the detailed instructions.

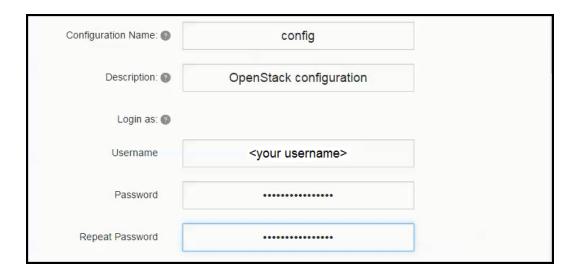


Select cloud provider

To enter the cloud provider settings:

1. Click Select cloud provider > OpenStack.

- 2. Enter the following values. As you enter them and click in an empty field, ElectricFlow tries to retrieve the cloud provider settings.
 - Configuration Name—A unique name for the OpenStack configuration. (Required)
 - **Description**–A description for this configuration. The default is *OpenStack configuration*.
 - **Login as**—The OpenStack account username and password. It must have enough privileges to perform API functions. (Required)



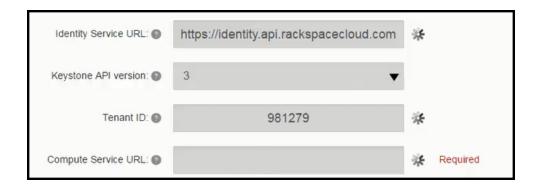
Identity Service URL

—The service end-point URL for the Identity Service for OpenStack. Do
not include the API version or the tenant ID in the URL. (Required)



• Keystone API version-The version of the Keystone Service API. The default is 3. (Required)

• Tenant ID-The ID of the OpenStack tenant to use. (Required)



When you click in the next empty field, ElectricFlow tries to retrieve the rest of the cloud provider settings for the selected configuration.

- Compute Service URL—The service end-point URL for the Compute Service for OpenStack.

 Do not include the API version or the tenant ID in the URL. (Required)
- Compute API version—The version of the Compute Service API. The default is 2.0. (Required)
- Block Storage URL

 —The service end-point URL for the Block Storage Service for OpenStack. Do not include the API version or the tenant ID in the URL. (Required)
- **Block Storage API version**—The version of the OpenStack block storage API. The default is 1. (Required)
- Image Service URL

 —The service end-point URL for the Image Service for OpenStack. Do not include the API version or the tenant ID in the URL. (Required)
- Image API version—The version of the Image Service API. (Required)
- Orchestration Service URL

 —The service end-point URL for the Orchestration Service for OpenStack. Do not include the API version or the tenant ID in the URL. (Required)
- **Resource**—The name of the resource or pool on which the integration steps can run. The default is *local*. (Required)
- Debug Level-The debug level for the output. The default is 1.

The possible values are

- **0** Errors only.
- 1 Normal headers and responses.
- 2+ Debugging information included.

3. Click Next.

The Provision Parameters fields appear.

The Configuration field is prepopulated with the name you entered in the previous dialog box.

ElectricFlow tries to retrieve the provision parameters for this configuration.

- Enter the Provision Parameters values.

 - Image-Select the image to use in OpenStack. (Required)
 - Flavor-Select the flavor to use in OpenStack. (Required)
 - Key Pair Name

 Select the key pair to use in OpenStack. (Required)
 - Security Group(s)—Select the security group to use in OpenStack.
 - Availability Zone—Select the availability zone in OpenStack in which to launch the server.
 - Customization script-Configuration information or scripts to execute when the server starts.
 - Results Location—Location where the output properties will be saved. If not specified, this
 defaults to /myJob/OpenStack/deployed.
 - Resource Pool

 Name of the resource pool. If you enter a name, a new resource pool is
 created, and new resources are added to the pool for the servers created in OpenStack.
 - Resource workspace—If you specify a resource pool name, this is the workspace that will be
 assigned to the created resources.
 - Resource Zone Name

 —The zone to which the newly created resources will belong. The
 default zone is default.
- 5. Click **Next** to set the configuration management settings.
- 6. Go to this step in Creating Resource Templates on page 148 to continue creating the resource template.
- 7. If you do not want to set the configuration management settings, click **OK**.

The Resources Templates list now shows the resource template that you created.

Select account

Select the **Select account** option to use an existing plugin configuration. ElectricFlow automatically retrieves the cloud provider settings and puts them in the appropriate fields.

- 1. Click Select account
- 2. Select an existing plugin.
- 3. Click Next.

The dialog box shows for the settings for the selected the plugin configuration.

- 4. If you want to edit any of these settings, make your changes and then enter the password to save your changes.
- 5. Click **Next** to enter the configuration management settings.
- 6. Go to this step in Creating Resource Templates on page 148 to continue creating the resource template.
- 7. If you do not want to set the configuration management settings, click OK.

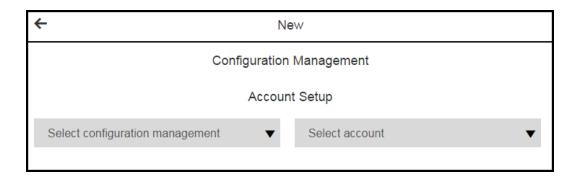
The Resources Templates list now shows the resource template that you created.

Example: Resource Template for Middleware Using Chef

This section shows how to configure Chef as the configuration management tool for a resource template after you have entered the cloud provider settings.

In the **New Configuration Management** dialog box, "account" refers to the configuration management setup. Select one of these options to enter the setup information:

- Click **Select configuration management** to enter the configuration management tool settings. Go here for the detailed instructions.
- Click **Select account** if you want to use an existing plugin configuration. Go here for the detailed instructions.

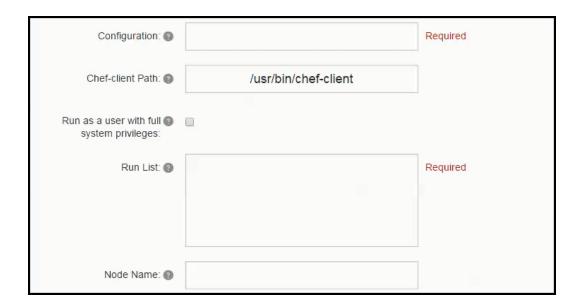


Select configuration management

To enter the configuration management settings:

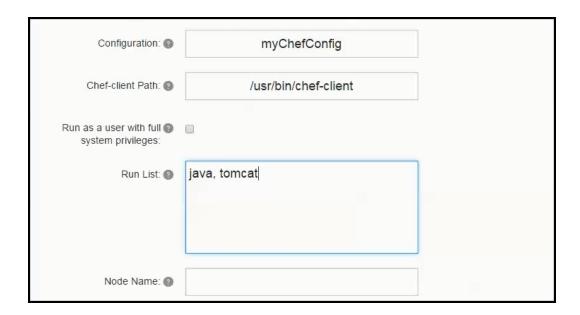
1. Click Select configuration management > Chef, and select Next.

The Converge Parameters fields appear.



2. Enter these settings:

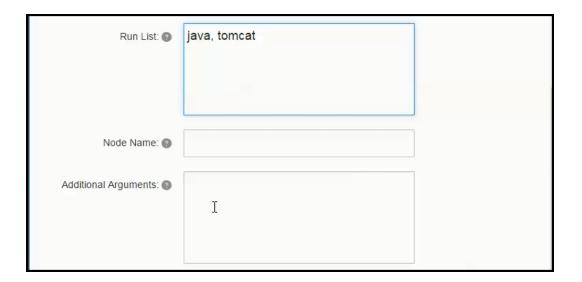
- Configuration—Name of the Chef configuration. This is required.
- Chef-client Path—The path to the chef-client executable. The default value is /usr/bin/chef-client. (Required)
- Run as user with full system privileges—If this check box is selected, chef-client is run as a user with elevated privileges on Unix systems (using sudo).
- Run List-The ordered list of Chef recipes to run. This is required.



Node Name

–Unique name for the node that will be maintained by the chef-client. If not
specified, Chef will use the Fully Qualified Domain Name of the node as the node name.
(Required)

• Additional Arguments—Any additional arguments that need to be passed to the chef-client.



3. Click OK.

The Resource Templates List now shows the resource template that you created.

Select account

Select the **Select account** option to use an existing plugin configuration. ElectricFlow automatically retrieves the configuration management settings and puts them in the appropriate fields.

- 1. Click Select account
- 2. Select an existing plugin.
- 3. Click Next.

The dialog box shows for the settings for the selected the plugin configuration.

- 4. If you want to edit any of these settings, make your changes and then enter the password to save your changes.
- 5. Click Next.

The Resource Templates List now shows the resource template that you created.

Viewing and Editing Resource Templates

You can view and edit this information about resource templates:

- Resource Template Details on page 167
- Resource Template Properties on page 168
- Access Control Settings on page 170
- Change Tracking on page 171

Resource Template Details

Starting in the Resource Template List, follow these steps to view and edit the resource template details:

1. Choose a resource template and click the **Menu** button.

A list of menu options appears.

Example:



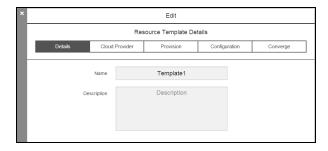
If you want to see only the objects for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:

- · Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

If there are no matches, a message appears stating that there are no resource templates in the selected projects.

2. Select Details.

The Edit Resource Template Details dialog box opens.



- 3. In the **Details** tab, edit the name and description of the resource template.
- 4. In the Cloud Provider tab, edit the cloud provider account information.
- 5. In the **Provision** tab, edit the information about the provision settings for the cloud provider.
- 6. In the **Configuration** tab, view and edit the configuration management details.

- 7. In the **Converge** tab, view and edit the details about how the virtual instances are converged in the defined configuration.
- 8. Click **Next** to go to the next tab.
- 9. Click **OK** to save your changes.

Resource Template Properties

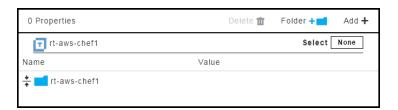
Starting in the Resource Template List, follow these steps to view and edit the resource template properties:

1. Choose a resource template and click the **Menu** button.

A list of menu options appears.

2. Select Properties.

The **Properties** dialog box opens.



3. To add a property:

a. Click the **Add +** button.

The **Property Details** dialog box opens.

If **Expandable** is selected, the property value can contain property references using the "\$[]" notation.



b. Enter the details and click **OK** to save them.

Example:



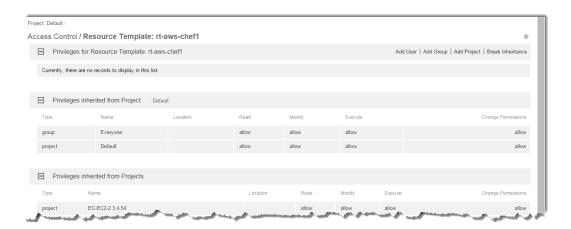
The **Property Details** dialog box now shows the new property that you added.

Example:



Access Control Settings

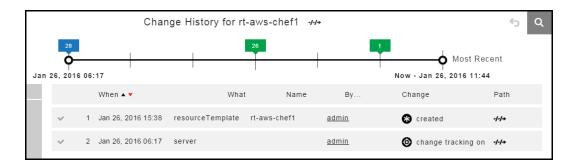
Clicking Access Control opens the Access Control page in the automation platform.



Change Tracking

Clicking Track Changes opens the Change History for the resource template.

Example:

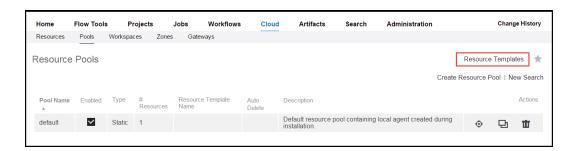


Accessing the Resource Templates in the Automation Platform

A resource template has the required information to provision and spin up cloud resources on an on-demand basis to model a dynamic environment. You set the cloud provider and configuration management details in the resource template. Then you define the environment tiers and add resource templates to these tiers in an environment template to complete the dynamic environment configuration.

You can generate new resource pools in ElectricFlow by accessing the resource templates from the automation platform UI.

1. Go to Cloud > Pools.



2. Click Resource Templates.

The Resource Templates list opens.

Example:





(Optional) To view details about a resource template, choose a template and click the Menu button (for it.

The Context menu opens.

- 4. Click one of these options:
 - Details-The Edit Resource Template Details dialog box opens.

The tabs that appear in the dialog box depend on the cloud provider specified in the resource template.

• Properties-The Properties dialog box opens.

You can view and edit the properties that apply to the resource template.

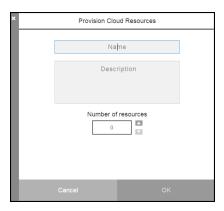
- Access Control-The Access Control page opens.
- Track Changes-The Change History for the resource template opens.



5. To provision the resource template, choose a resource template, click the **Menu** button (), and select **Details > Provision**.

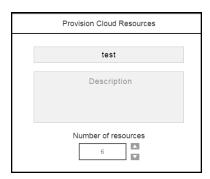
The Provision Cloud Resources dialog box opens.

Example:



6. Enter the resource pool name, an optional description, and the number of resources in the pool, and then click **OK**.

Example:



You go to the Job Details page in the automation platform.

Creating Environment Templates

You can create a new environment template from scratch or create one based on an existing template.

- Creating a New Environment Template on page 173
- Creating an Environment Template Based on an Existing Template on page 182

Creating a New Environment Template

Starting from the Environment Templates List, follow these steps to create an environment template from scratch:

1. Go to the Environment Templates List.

If there are no defined environment templates, the Environment Templates List is empty.



- 2. Add an environment template.
 - a. If the Environment Templates List is empty, click **There are no Environment Templates yet. Add one +** or click the **Add +** button.

In the **New** dialog box, click **Environment Template**.

Example:



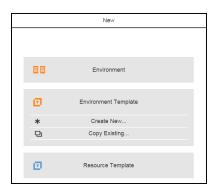
The **New** dialog box opens.



b. If the Environment Templates List is not empty, click the ${\bf Add} + {\bf button}.$

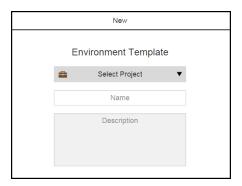
The **New** dialog box opens.

Example:



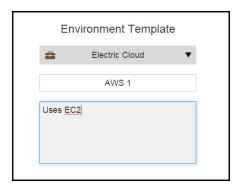
3. Click **Create new** to create a new environment template.

The **New** dialog box opens.



4. Enter a name for the environment template, select a project to which it belongs, and enter an optional description.

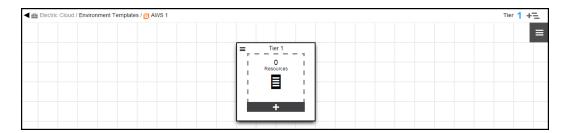
Example:



5. Click OK.

The Environment Templates Visual Editor opens. The environment template has one tier with no assigned resources. The breadcrumb in the upper left corner shows the object hierarchy.

Example:

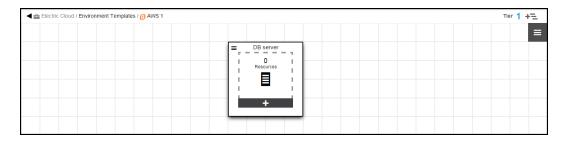




6. In Tier 1, click the button and select **Details**.

7. In the **Environment Tier Details** dialog box, enter a new name for the environment tier and an optional description, and then click **OK** to save the settings.

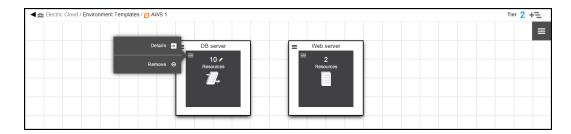
Example:



8. In the tier, click the + button to add a resource to it.

The **New** dialog box to add resources opens.

Example:



9. Click **Add resource template** to select a resource template.

The Resource Templates List opens.

10. Select a resource template, and click OK.

The dialog box to select a resource template opens.

It shows the list of available environment templates. The default is to show the environment templates for all projects.

If you want to see only the environment templates for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more environment templates by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no environment templates in the selected projects.

Example:



The **New** dialog box to select the number of resources to provision opens.



11. Enter the number of resources to provision, and click **OK**.

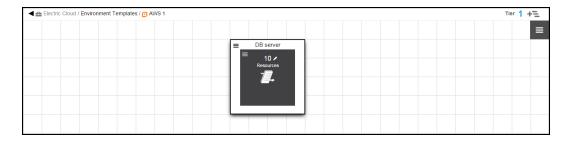
You must provision at least one resource (1).

Note: You cannot override the number of resources when deploying an application with an environment template.

Example:



The Environment Templates Visual Editor now shows that the tier (Tier 1) has 10 dynamic cloud resources set to be provisioned when the application is deployed.



- 12. To add static resources to an environment template:
 - a. Click the Add tier button to create a new tier.

A new environment tier appears in the Environment Templates Visual Editor.

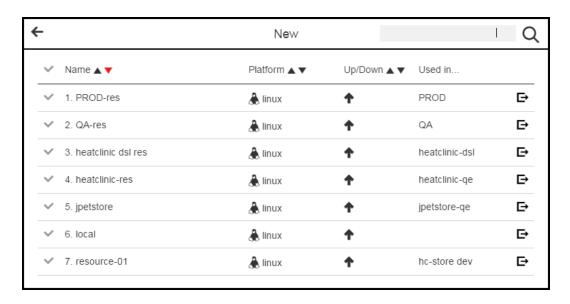
- b. Rename the new environment tier.
- c. Click the + button in the new tier to add resources to it.

The **New** dialog box opens.

d. Click Add resources to add static resources to the environment tier.

A list of static resources opens.

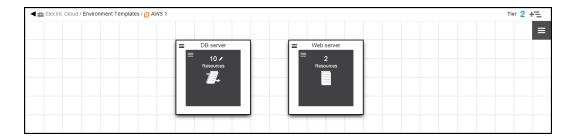
Note: You cannot override the setting for the number of static resources when deploying the application with an environment template



e. Select one or more resources in the list, and click \mathbf{OK} .

The Environment Templates Visual Editor now shows that the "Web server" tier has two static resources.

Example:



Creating an Environment Template Based on an Existing Template

Starting from the Environment Templates List, follow these steps to create an environment template that is based on an existing one:

- 1. Go to the Environment Templates List.
- 2. Add a resource template.

For detailed instructions, go to Add an environment template. on page 175

3. Click Copy existing to create a template based on an existing one.

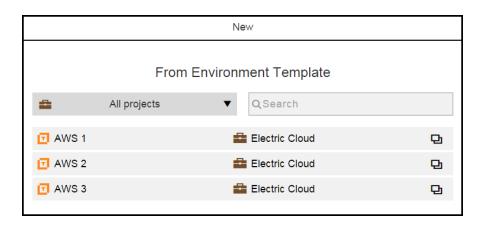
The list of existing templates opens.

It shows the list of available environment templates. The default is to show the environment templates for all projects.

If you want to see only the environment templates for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

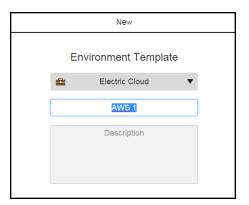
- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more environment templates by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no environment templates in the selected projects.



4. Select an environment template.

The **New** dialog box opens.

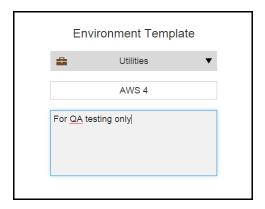


5. Rename the environment template, and enter an optional description of it.

You can also change the project to which the environment template belongs.

Click **OK** to save these settings.

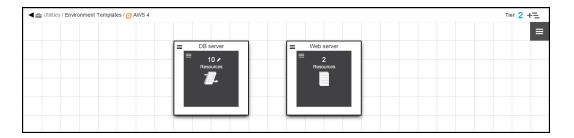
Example:



The Environment Templates Visual Editor now shows that the same environment tiers as the original version. The breadcrumb in the upper left corner shows the object hierarchy.

You can now modify the environment template and its objects (resources and resource templates) for your deployments

Example:

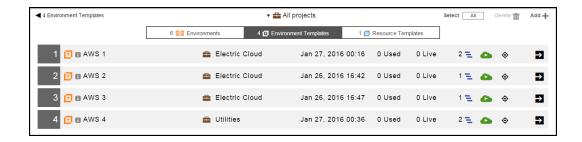


Viewing and Editing Environment Templates

Starting in the Environment Templates List, follow these steps to view and edit the environment template details:

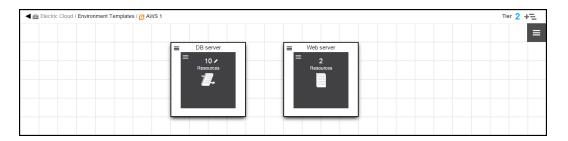


1. Choose an environment template and click the View Details button (



The Environment Templates for the selected template opens.

The breadcrumb in the upper left corner shows the object hierarchy.





- 2. Click to **Menu** button () to open the context menu for the environment template.
 - Click **Details** to view or edit the name and description of the environment template.
 - Click **Properties** to view or edit the environment template properties.
 - Click Access Control to go to the Access Control page in the automation platform.
 - Click Track Changes to open the Change History of the environment template.
 - Click Delete to delete this template.



- 3. To view and edit an environment tier with dynamic resources, click the **Menu** button () for the resources. The context menu opens with these options:
 - Click **Details** to view and edit the resource details for the environment tier.
 - Click **Remove** to remove the resources from the environment tier.

4. To change the resource template used by the environment tier, click **Details** and select a different resource template.

The list of available resource templates opens.

It shows the list of available resource templates. The default is to show the resource templates for all projects.

If you want to see only the resource templates for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- · Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more resource templates by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no resource templates in the selected projects.

The template being used is selected.



5. Select a different environment template.

A message appears. Select Yes if you want to continue.

Example:



The Edit dialog box number box opens and shows the number of provisioned dynamic resources.

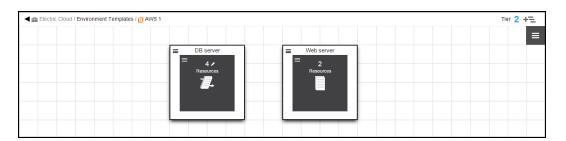
Example:



- 6. To change the number of provisioned dynamic resources, enter a new number or use the up/down arrows.
- 7. Click OK.

The Environment Templates Visual Editor now shows the changes that you made.

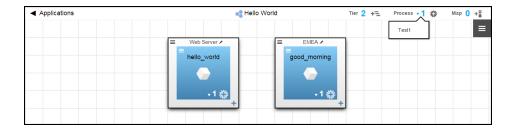
Example:



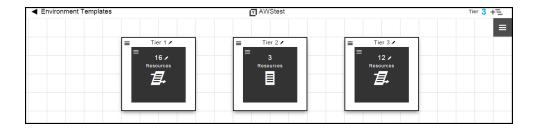
Deploying Applications With Provisioned Cloud Resources

About the example in this topic:

The example in this topic consists of an application called "Hello World" application with one application process called *Test*1, which has two steps.



Test1 will be deployed in a dynamic environment called *CloudEnv*. You model the CloudEnv environment using the *AWStest* environment template.



The AWStest environment template has three tiers:

- Tier 1 and Tier 3 have cloud resources defined in resource templates. These resource can be provisioned when you deploy the application.
- Tier 2 has static resources, which cannot be provisioned.

Starting in the Home page:

- 1. Go to the Applications List.
 - Starting from the Main menu, click the **Menu** button, and then select **Applications**.
 - Starting from the Home page, click Applications.

The Applications List opens.

Example:

This example uses the "Hello World" application.



2. Choose an application, and click the **Run process** button for that application.

Example:



3. Click New Run.

The **New** dialog box to deploy the application opens.

See the messages in this dialog box for hints about what you need to do to deploy the application.

Example:

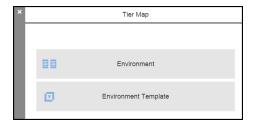


4. Select the application process to deploy.

5. Click **Map to Environment +** to create a tier map for the application.

The **Tier Map** dialog box opens.

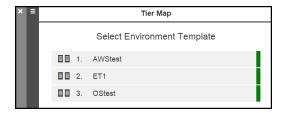
Example:



6. Select Environment Template.

The **Tier Map** dialog box to select an environment tier opens.

Example:

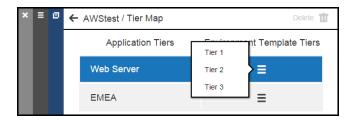


7. Select an environment template.

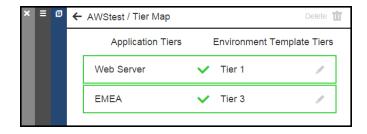
The <Environment template>/Tier Map dialog box opens.

8. For each application tier, click the **Menu** button to select an environment tier to which the application tier is mapped.

Example:



Example:



9. Click OK.

The **New** dialog box to deploy the application re-opens.



10. Click **Select Environment** to select an environment template.

A list of available environments and environment template based on the tier map opens.

Example:



11. Select 1 Templates.

A list of available environment templates appears.

Example:



12. Click the environment template that you want to use.

13. Enter a name for the dynamic environment that will be created from the selected environment template.

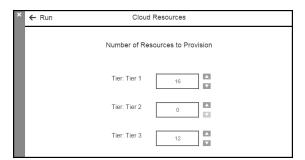
The **New** dialog box to deploy the application now shows the environment template name below the application process name. It also shows the number of cloud resources provisioned in the environment templates.

Example:



14. Click in the Cloud Resources row.

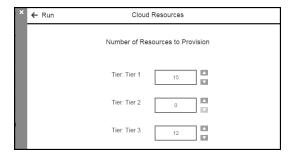
The Cloud Resources dialog box opens.



15. Change the number of cloud resources to provision

Example:

In this example, Tier 1 and Tier 3 have one or more resources to provision because they have cloud resources. You cannot provision resources in Tier 2 because it has only static resources.



16. Click **OK**.

The **New** dialog box now shows the new number of resources to provision.

The message "This application does not yet have an Environment to Run on" still appears.

When you click **OK**, ElectricFlow first attempts to create the dynamic environment. If this is successful, it deploys the application.

Example:



The Applications List opens.

17. To view the job details, click the View Details button to open to the Job Details page.

ElectricFlow first runs the job to create the dynamic environment. If this job is successful, it deploys the application.

Example:



You can also see the status of jobs as they progress in the Applications List.

Example:

This example shows the Application List when the job starts.



If you provision resources that are not available, such as 20 cloud resources in Tier 1 and 3 static resources in Tier 2, the job to create the dynamic environment fails and the application is not deployed.

Retiring Dynamic Environments

After you have deployed your application in a dynamic environment, you can retire it and destroy (decommission) the resources, which makes them unavailable for any future deployments.

Starting in the Environments List:

1. Choose a dynamic environment.



2. Click the Tear down button for that environment.

A message appears.

Example:



3. Click **OK** to verify that you want to retire (decommission) the environment.

The Job Details page opens.

Example:



4. View the results in the Job Details page.

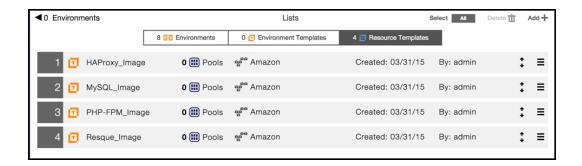
Dynamic Environment Example with Amazon and Chef

This example shows what you see in the ElectricFlow UI when your dynamic environment is configured with Amazon as the cloud provider and Chef for the configuration management.

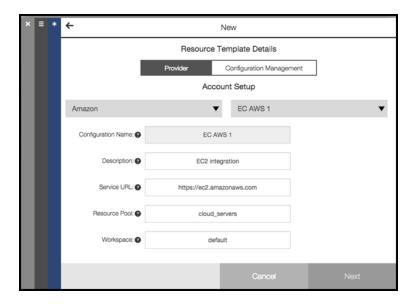
Resource Templates

In resource templates, you define cloud resources that will be provisioned for dynamic environments. Examples of cloud resources are a HAProxy server, MySQL server, PHP-FPM server, and Resque server.

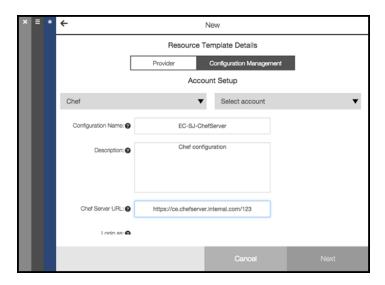
This is the Resource Template List.



These are the cloud provider settings for a resource template. It is defined by cloud provider account credentials and by an Amazon Machine Image (AMI).

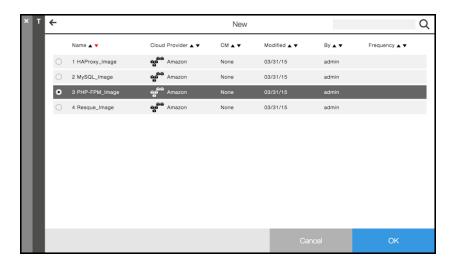


These are the configuration management settings for a resource template. It calls a Chef recipe to configure the cloud resources. You can configure the resources with other recipes or predefined configuration, such as a MySQL stack configuration or a HAProxy load balancer configuration.

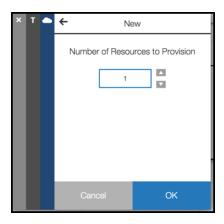


Environment Templates

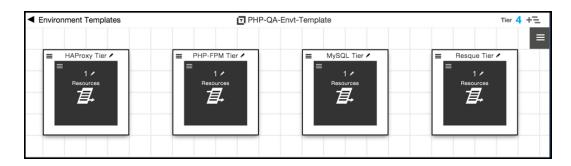
When you create a new environment template, you can add one resource template to an environment tier. These are the resource templates from which you can select.



After you select a resource template and click **OK**, you enter the number of resources to provision in an environment tier.

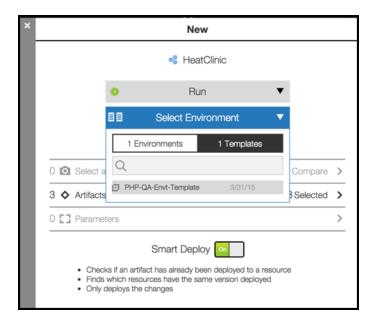


This is an environment template with four tiers. Each tier is configured with a resource template, and resource template has a provisioned resource.



Deploying Applications to Dynamic Environments

After you select an application to deploy, you select the application process to run and the environment in which to deploy the application.



To deploy the application in a dynamic environment, you can select an environment template. A dynamic environment is created with the provisioned cloud resources. After the resources are provisioned and the dynamic environment is created, the application is deployed.

You can view the status of the provisioning process.



Retiring Dynamic Environments

After the provisioning process completes, the new dynamic environment appears in the Environments List.



You can retire (decommission) dynamic environments on an on-demand basis to prevent excessive use of cloud resources and reduce costs.

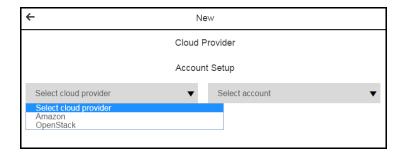
You click the **Tear down** button to retire the selected dynamic environment.



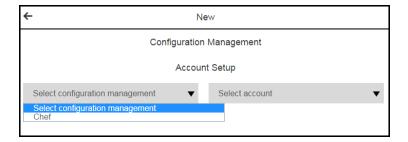
Developer Task: Creating Custom Plugins

ElectricFlow uses third-party plugins to provision cloud resources in resource templates for Dynamic Environments. Amazon EC2 (EC-EC2) and Openstack (EC-OpenStack) are supported as out-of-the-box cloud provider plugins, and Chef (EC-Chef) is supported as an out-of-the-box configuration management plugin.

To view the supported cloud provider plugins, open the **New Cloud Provider** dialog box and click **Select cloud provider** to view the drop-down menu of cloud providers.



To view the supported configuration management plugins, click **Select configuration management** to view the drop-down menu of configuration management tools.



If you have an existing procedure that you would prefer to use instead of an out-of-the-box third-party plugin to provision cloud resources, you can create a custom plugin based on your deployment scenario.

The process to create custom cloud provider and configuration management plugins uses metadata to loosely couple them to the application you want to deploy. The custom plugins define specific known properties that are automatically recognized by Dynamic Environments.

To create your own plugin for other cloud platforms, go to https://electric-cloud.github.io/index.html to get started. This website describes how to create open source plugins and has this information:

- For an introduction to the process of building and developing open source plugins, go to the Quick Start.
- To build and deploy your plugin, go to Build and Deploy.
- For more information, go to More information.
- Go to Creating Custom Cloud Provider Plugins on page 203 for the details to create custom cloud provider plugins.
- Go to Creating Custom Configuration Management Plugins on page 208 for the details to create custom configuration management plugins.

Creating Custom Cloud Provider Plugins

This section applies to developers who want to create custom cloud provider plugins for Dynamic Environments in ElectricFlow.

How to Create a Custom Cloud Provider Plugin

This procedure assumes that you already have a plugin and want to use it with ElectricFlow. It describes how to define properties in your plugin that ElectricFlow can read.

1. To convert your procedure to a plugin, define the name of the plugin and the ElectricFlow project to which the plugin belongs.

Later in this procedure, you will use these properties to create pre- and post-hooks to the plugin in ElectricFlow.

2. Define the following properties in your plugin under a top-level plugin property called ec_cloudprovisioning_plugin, which the system will read to Dynamic Environments.

Property Name	Description
ec_cloudprovisioning_ plugin/	Top-level plugin property directory
Properties defined by the custom plugin under ec_cloudprovisioning_plugin	
displayName	Name of the plugin that appears in the New Cloud Provider dialog box.
	Example: Amazon for the EC-EC2 plugin
	<boolean -="" 0 1 true false="" flag=""></boolean>
hasConfiguration	 Set this property is set to 0 if the plugin does not have any configuration procedures (CreateConfiguration and DeleteConfiguration).
	 Set this property is set to 1 if the plugin has configuration procedures.
	Name of the property sheet used by the plugin to store the saved configurations.
configurationLocation	This value is relative to the plugin's top-level properties.
Configurationsocation	For example, if this value is set as ec2_cfgs, the plugins configurations are read from /plugins/ <plugin_ key="">/project/ec2_cfgs.</plugin_>
operations/	Property sheet for the specific operations required by the Dynamic Environment system.
Properties under operations/ – These operations are mapped to the plugin operations through the following child properties.	
createConfiguration/	Properties for the procedure that creates the plugin configuration. It is usually called createConfiguration.
deleteConfiguration/	Properties for the procedure that deletes the plugin configuration. It is usually called deleteConfiguration.
provision/	Properties for the procedure that provisions virtual instances.

Property Name	Description	
retireResource/	Properties for the procedure that tears down (decommissions) a previously provisioned virtual instance backing the specified resource.	
retireResourcePool/	Properties for the procedure that tears down (decommissions) all previously provisioned virtual instances for the specified resource pool.	
Properties for the previously listed operations defined in the operations property sheet		
procedureName	Name of the procedure name in the plugin to which the operation is mapped.	
ui_formRefs/	Property sheet that references the ui_forms properties defined under the plugin's top-level properties.	
parameterRefs/	Property sheet for the input parameters to the procedure that can be used by the Dynamic Environment system. The parameter list is operation-specific.	

Go to Example: Property Structure for a Cloud Provider Plugin on page 206 for an example of the properties for the EC-EC2 (Amazon) plugin.

- 3. Save the plugin file in the appropriate location on your ElectricFlow server.
- 4. To import the plugin file to the automation platform:
 - a. In the automation platform, go to Administration > Plugins to open the Plugin Manager page.
 - b. Click the Install from File/URL tab.
 - c. In the File Install field, click Choose file to select the plugin file.
 - d. Click Upload to install it.

The plugin file appears in the Currently Installed tab.

- e. Find your plugin and click **Promote** in the Actions column to make it available for use by ElectricFlow.
- 5. In ElectricFlow, to create a resource template:
 - a. Go to Creating Resource Templates on page 148 for the details.
 - b. In the **Select cloud provider** field, select your plugin as the cloud provider.
 - c. Enter the plugin settings in the form.
 - d. (Optional) Set the configuration management tool.

Go to Creating Resource Templates on page 148 for the details.

- 6. In the ectool API, enter the following commands to set pre- and post-hooks linking the plugin to an application that you later deploy in a dynamic environment.
 - a. To create a pre-hook, enter

```
ectool createHook postConfigurationHook --hookType POST_CONFIGURATION --
procedureName <hookProcedureName> --procedureProjectName
<hookProjectName> --resourceTemplateName <resourceTemplateName> --
projectName <resourceTemplateProjectName>
```

where hookProcedureName is the plugin procedure name, hookProjectName is the name of the project to which the plugin procedure belongs, resourceTemplateName is the name of the resource template you created, and resourceTemplateProjectName is the name of the project to which the resource template belongs.

b. To create a post-hook, enter

```
ectool createHook preConfigurationHook --hookType PRE_CONFIGURATION --
procedureName <hookProcedureName> --procedureProjectName
<hookProjectName> --resourceTemplateName <resourceTemplateName> --
projectName <resourceTemplateProjectName>
```

where hookProcedureName is the plugin procedure name, hookProjectName is the name of the project to which the plugin procedure belongs, resourceTemplateName is the name of the resource template you created, and resourceTemplateProjectName is the name of the project to which the resource template belongs.

7. Create an environment template with the resource template that you created.

Example: Property Structure for a Cloud Provider Plugin

This is the EC-EC2 (Amazon) plugin property structure.

Property Name	Property Values
ec_cloudprovisioning_plugin/	
displayName	Amazon
hasConfiguration	1
configurationLocation	ec2_cfgs
operations/	createConfiguration/ deleteConfiguration/ provision/ retireResource/ retireResourcePool/
ec_cloudprovisioning_plugin/createConfiguration/	

Property Name	Property Values	
procedureName	CreateConfiguration	
ui_formRefs/		
ui_formRefs/parameterForm	ui_forms/EC2CreateConfigForm	
parameterRefs/		
parameterRefs/ configuration	config	
ec_cloudprovisioning_plugin/de	eleteConfiguration/	
procedureName	DeleteConfiguration	
ui_formRefs/	-	
parameterRefs/	configuration called config	
ec_cloudprovisioning_plugin/provision/		
procedureName	API_RunInstances	
ui_formRefs/	parameterForm called ec_parameterForm	
parameterRefs/	configuration called config resourcePool called res_poolName count called count	
ec_cloudprovisioning_plugin/retireResource/		
procedureName	API_TearDownResource	
ui_formRefs/	-	
parameterRefs/	resourcePool called res_poolName	
ec_cloudprovisioning_plugin/retireResourcePool/		
procedureName	API_TearDownResource	
ui_formRefs/	_	
parameterRefs/	resourcePool called res_poolName	

Creating Custom Configuration Management Plugins

This section applies to developers who want to create custom configuration management plugins for Dynamic Environments in ElectricFlow.

How to Create a Custom Configuration Management Plugin

This procedure assumes that you already have a plugin and want to use it with ElectricFlow. It describes how to define properties in your plugin that ElectricFlow can read.

1. To convert your procedure to a plugin, define the name of the plugin and the ElectricFlow project to which the plugin belongs.

Later in this procedure, you will use these properties to create pre- and post-hooks to the plugin in ElectricFlow.

2. Define the following properties in your plugin under a top-level plugin property called ec_ configurationmanagement_plugin, which the Dynamic Environment system can access.

Name	Description	
ec_ configurationmanagement_ plugin/	Top-level plugin property directory	
Properties defined by the custom plugin under ec_configurationmanagement_plugin		
displayName	Name of the plugin that appears in the Dynamic Environment UI Example: Chef for the EC-Chef plugin	
hasConfiguration	 <boolean -="" 0 1 true false="" flag=""></boolean> When this property is set to 0 or false, the plugin does not have any configuration procedures (CreateConfiguration and DeleteConfiguration). When this property is set to 1 or true, the plugin has configuration procedures. 	
configurationLocation	Name of the property sheet used by the plugin to store the saved configurations. This value is relative to the plugin's top-level properties. If this value is set as chef_cfgs, the configurations are in /plugins/ <plugin_key>/project/chef_cfgs.</plugin_key>	
operations/	Property sheet for the specific operations required by the Dynamic Environment system.	
Properties under operations/ – These operations are mapped to the plugin operations through the following child properties.		
createConfiguration/	Properties for the procedure that creates the plugin configuration. It is usually called createConfiguration.	
deleteConfiguration/	Properties for the procedure that deletes the plugin configuration. It is usually called deleteConfiguration.	

Name	Description
converge/	Properties for the procedure that converges the virtual instances to the defined configuration (including policies and roles).
teardown/	Properties for the procedure that deletes the configuration details on the specified dynamic resource or resource pool.
	For Chef, properties for the procedure that deletes the Chef API client and Chef node on the specified dynamic resource or resource pool.
Properties for the previously listed operations defined in the operations property sheet	Properties for the previously listed operations defined in the operations property sheet
procedureName	
ui_formRefs/	Property sheet that references the ui_forms properties defined under the plugin's top-level properties.
parameterRefs/	Property sheet for the input parameters to the procedure that can be used by the Dynamic Environment system. The parameter list is operation-specific.

Go to Example: Property Structure for a Configuration Management Plugin on page 211 for an example of the properties for the EC-Chef plugin.

- 3. Save the plugin file in the appropriate location on your ElectricFlow server.
- 4. To import the plugin file to the automation platform:
 - a. In the automation platform, go to **Administration > Plugins** to open the **Plugin Manager** page.
 - b. Click the Install from File/URL tab.
 - c. In the File Install field, click Choose file to select the plugin file.
 - d. Click Upload to install it.
 - The plugin file appears in the **Currently Installed** tab.
 - e. Find your plugin and click **Promote** in the Actions column to make it available for use by ElectricFlow.

- 5. In ElectricFlow, to create a resource template:
 - a. Go to Creating Resource Templates on page 148 for the details.
 - b. In the **Select cloud provider** field, select a cloud provider.
 - c. Enter the account settings for the selected cloud provider.
 - d. Enter the provision parameters on the next page.
 - e. In the **Select configuration management** field, select your plugin as the configuration management tool.

Go to Creating Resource Templates on page 148 for the details.

- 6. In the ectool API, enter the following commands to set pre- and post-hooks linking the plugin to an application that you later deploy in a dynamic environment.
 - a. To create a pre-hook, enter

```
ectool createHook postConfigurationHook --hookType POST_CONFIGURATION --
procedureName <hookProcedureName> --procedureProjectName
<hookProjectName> --resourceTemplateName <resourceTemplateName> --
projectName <resourceTemplateProjectName>
```

where hookProcedureName is the plugin procedure name, hookProjectName is the name of the project to which the plugin procedure belongs, resourceTemplateName is the name of the resource template you created, and resourceTemplateProjectName is the name of the project to which the resource template belongs.

b. To create a post-hook, enter

```
ectool createHook preConfigurationHook --hookType PRE_CONFIGURATION --
procedureName <hookProcedureName> --procedureProjectName
<hookProjectName> --resourceTemplateName <resourceTemplateName> --
projectName <resourceTemplateProjectName>
```

where hookProcedureName is the plugin procedure name, hookProjectName is the name of the project to which the plugin procedure belongs, resourceTemplateName is the name of the resource template you created, and resourceTemplateProjectName is the name of the project to which the resource template belongs.

7. Create an environment template with the resource template you created.

Example: Property Structure for a Configuration Management Plugin

This is the EC-Chef plugin property structure.

Name	Values
ec_configurationmanagement_plugin/	
displayName	Chef
hasConfiguration	1

Name	Values		
configurationLocation	chef_cfgs		
operations/	<pre>createConfiguration/ deleteConfiguration/ converge/ teardown/</pre>		
ec_configurationmanagement_	ec_configurationmanagement_plugin/createConfiguration/		
procedureName	CreateConfiguration		
ui_formRefs/	parameterForm called forms/CreateConfigForm		
parameterRefs/	configuration called config		
ec_configurationmanagement_plugin/deleteConfiguration/			
procedureName	DeleteConfiguration		
ui_formRefs/			
parameterRefs/	configuration called config		
ec_configurationmanagement_plugin/converge/			
procedureName	_RegisterAndConvergeNode		
ui_formRefs/	parameterForm called ec_parameterForm		
parameterRefs/	configuration called config		
ec_cloudprovisioning_plugin/teardown/			
procedureName	_DeleteNode		
ui_formRefs/	_		
parameterRefs/	resourceName called resource_name		

Inventory Tracking

Electric Flow uses Inventory Tracking to track what is built, tested, and deployed by the application, including artifacts, artifact versions, resources on which the applications are run, and environments to which the

resources are assigned. If there is an issue when an application is deployed you can find the details about what was deployed.

- Use the Environment Inventory to track and view details of the objects that were deployed and artifacts in
 the application. It shows the status of the application deployment at a point in time. Go to Example:
 Modeling and Deploying Applications on page 88 and Environment Inventory on page 213 for more
 information.
- Use the Application Inventory (on the Applications Run View page) to track and view the deployment results. It shows more details about the application at a point in time. Go to the Applications Run View page in the GUI to see these details. Go to Example: Modeling and Deploying Applications for an example.

Tracking at the Component Process Level

Inventory Tracking occurs at the component process level.

When defining a component process, you select one of these process types:

- **Deploy**–Enables Inventory Tracking. The ElectricFlow server tracks artifacts deployed to environments. This is the default.
- **Undeploy**–After the first successful job step in a component process with this setting, the automation platform removes the environment inventory record.
- Other–Disables Inventory Tracking.

Environment Inventory

The Environment Inventory is the state of the environment at a point in time.

- When an application is running, you can see its progress as it runs.
- After an application runs, you can see the details for the objects in application.

This example shows the Environment Inventory for the *heatclinic-qe* environment to which the HeatClinic application is deployed. The application has two application processes called Run and Stop.

In the Environments List, choose an environment and click the **Inventory** button to open the Environment Inventory.



Clicking in the row for an object in the Environment Inventory displays more details about that object, such as:

- Environment name
- · Application mapped to this environment
- · Number of deployed artifacts in the application
- · When the artifacts were deployed
- · Status of the deployment: success or failure



1. Go to the heatclinic-qe row, and click to open the Environment Inventory.



This is the Environment Inventory.





2. Click to view the list of components in the Run process of the HeatClinic application.



3. Click the **Process** button in any of the component rows to view the Run application process runs in the heatclinic-qe environment.

The results of running the Run application process, application process steps, and component process steps now appear in the inventory. An application process step can be based on a component (defined by component processes), a plugin, a command, or a procedure.

The Run application process consists of four process steps: Check, deploy war, cleanup db, and start server.

- The Check step is defined by a command.
- The remaining steps are defined by components in the HeatClinic application.

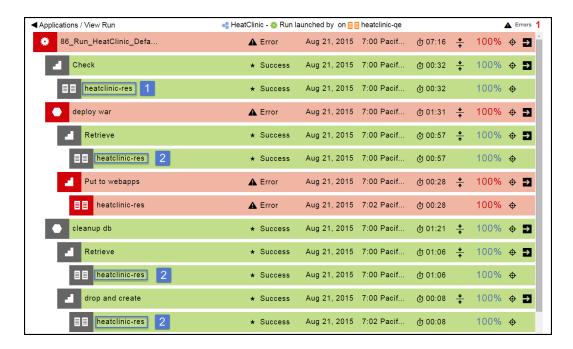


	Object	
1	Application process	
2	Application process step defined by a command	
3	Application process step defined by a component	
4	Component process step	

- 4. To open the Job Details page for an application process, an application process step, or component process step, do one of the following:
 - Click the object name to open the Job Details page for that object.
 - Click the View details button at the end of the row.



5. For a process step, click to view the results of deploying process steps.



	Object
1	Click here to view the results of running the application process step in the Job Details page.
2	Click here to view the results of running the component process step in the Job Details page.

Pipelines

A pipeline is a representation of your software delivery process on its journey to production, enabling Application Release Automation (ARA) and Continuous Delivery (CD). It takes the software through the hops and milestones (stages) of the software release life cycle, however you define it. You can model and execute one or more applications, procedures, or workflows in a pipeline.

Pipelines are a way to orchestrate the flow of the software release process. You can choose what part of the process in the pipeline to automate, based on your needs. Some examples are:

- From code check-in to production: After developers check in code, it is run through a pipeline that
 includes build, test, pre-production, and production stages, running various automation tasks in each of
 the stages like deploying the applications, running test automation suites, and so on, ending when the
 software is delivered to the production server.
- From system test to production: After the software passes the system tests, it is run through a pipeline to integrate the code from multiple teams, run user acceptance tests, and deploy the software to for beta testing and then all the way to the production.

Starting in ElectricFlow 6.2, multiple project support has been extended to ElectricFlow objects in application deployments, pipelines, and Releases. These objects as well as the objects belonging to them can be in any project within ElectricFlow. (In releases earlier than ElectricFlow 6.2, these objects were limited to the Default project.)

All objects in ElectricFlow have these capabilities for better management and maintenance of end-to-end software releases:

- Projects are top-level security containers that provide access control at the project level. All objects in a project inherit the access control settings from the project to which they belong.
- Projects are also ways to provide logical groupings of objects. If your deployment is large and has many
 production environments, you can organize production environments by project, with each project
 having its own access control settings. This makes it easier to manage and maintain multiple
 environments during the software delivery process.

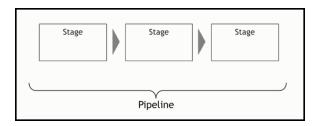
In this section, you will learn:

- About pipeline models–See Pipeline Concepts on page 218.
- How to author and run a pipeline-See Example: Authoring and Running Pipelines on page 236.
- About pipeline stage summaries—See Pipeline Stage Summary on page 227.

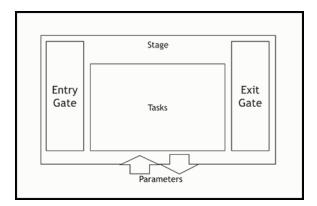
IMPORTANT: When you export a project while a pipeline is in progress, only the full export includes flow runtimes that have been completed. If you want to include in-progress pipeline runs in the path-to-production view and the visual indicators showing their percentage completed in the Release dashboard, set the excludeJobs argument to 0 or false in the export command. When the XML file is imported, the inprogress pipeline runs in the imported project are displayed in the path-to-production view and the Release dashboard.

Pipeline Concepts

A pipeline allows you to model the complete software delivery process. You can create multiple milestones, referred as to *stages*, in a pipeline. Examples of stages are development (Dev), QA, user acceptance test (UAT), pre-production, production, and so on. Most of the time, these stages will map closely to the environments where applications will be deployed, but do not have to be.



Each stage consists of an entry gate, an exit gate, and the automation tasks to complete the stage.



- An entry gate controls the approvals required before a pipeline can enter a stage. The entry gate ensures all the checks and balances are met before tasks in that stage can begin.
- The exit gate allows you to control when a pipeline can exit a stage and what approvals are needed.
- The tasks allow you to execute the automation logic. This is where you specify what to do.
 For example, if your stage is a test stage, it contains the tasks that must be executed to ensure that the test automation for your software delivery process is successfully completed.
- Parameters can be created for various objects and automation tasks in the stage. This makes the
 pipeline reusable, allowing orchestration of the inputs and outputs for the stage and of the properties
 passed between stages.

Tasks

Each task can specify automation logic defined by an application process, procedure, or workflow. You can also call integrations to third-party tools using plugins, or specify manual tasks performed by one or more users.

• Application process: Use this to deploy applications at the right time in a pipeline.

When artifact staging is enabled for an application, ElectricFlow tries to retrieve artifacts that will be deployed before the application run starts. This ensures that the artifacts are downloaded and available on the target (either the artifact cache [the default] or a directory specified in the component definition). This minimizes the time to retrieve artifacts and reduce downtime during the application deployment.

For an example showing how to deploy applications, go to Example: Modeling and Deploying Applications on page 88.

- Procedure: Use this to run a set of best practices, subroutines, modules, or functions that you can create and reuse at the platform level. Examples are scripts, command sets, or other automation logic objects.
- Manual task: Use this when a task should be performed manually. The appropriate notifications are sent
 to the user or group who performs that task and to the user or group who responds when the task is
 completed.
- Workflow: Use this to run a set of procedures in CI life cycles at the platform level. A workflow can combine all procedures for a build-test-deploy life cycle.

After the workflow is completed, it returns the workflow state to the pipeline as ERROR or SUCCESS:

- ERROR: The ERROR_LOG has entries, /myWorkflow/outcome or /myWorkflow/result is error, the workflow ended in a state that includes *Error* or *Fail* in its name, or it has any failed jobs.
- SUCCESS: /myWorkflow/outcome or /myWorkflow/result is exists and are not equal to error.
- Plugin: Use this to run third-party procedures that interface with external systems. Plug-in tasks are a
 perfect way to integrate and orchestrate your third-party tools. For example, if you can run Selenium
 automation at the right time using the Selenium plugin.
- Utility: Use this to run a utility function, a higher-order operation than a third-party plugin, as a stage task. You can also define an application or component process step as a utility function. Go here for more information about using utility functions in applications.

You may also select a Deployer task, a special type of task used when you want to deploy applications in a *Release* but will provide the details in the Release definition. This is explained in more detail in Release Management on page 274. When using Deployer tasks, the system will automatically know the applications to deploy, which are provided as part of a Release definition instead of creating separate Deploy tasks for each application.

When defining a task, you also set the error handling condition. In the **On Error** field in the UI, you can select **Continue running** or **Stop running** (the default).

- **Continue running**: If the task fails, the pipeline continues to the next gate or stage task with a known error. This behavior works best for pipelines in traditional software releases.
- **Stop running**: This is the default behavior. If the task fails, the pipeline aborts. This behavior usually occurs in Continuous Delivery pipelines.

For a manual task, the pipeline determines whether it stops running or continues running based on the rejection or completion of the task. If the task is rejected and **Continue running** is selected in the **On Error** field is selected, the pipeline continues to the next gate or stage task.

In pipelines created before ElectricFlow 6.4, error handling for tasks defaults to **Stop running** in the **On Error** setting in the task definition.

Entry and Exit Gates

Gates at the entries and exits of the stages control the pipeline's progress. Specific users or groups can be allowed to manually approve the criteria at a gate. During a pipeline run, you can see the progress of the pipeline at each stage using the pipeline stage summary. You can review the user-generated information in the pipeline stage summary to make approve or reject decisions at a gate.

- An entry gate controls if a stage may begin and its tasks can be executed.
- An exit gate controls if the stage may be completed and the pipeline can progress.
- You can set the approvals such that approval from only one of the approvers is sufficient for the pipeline
 to progress. In this case, you can have one approval rule and have all the approvers assigned to that
 rule.
- Instead, if you want all the key individuals to approve sequentially, you can specify multiple approval rules, each for a specific user.
- In addition to specific users, approval rules support groups. Approval or rejection from any user in the group will allow the pipeline to progress or stop.
- When defining a gate, you also set the error handling condition. In the On Rejection field in the UI, you
 can select Stop running (the default) or Continue running.
 - **Stop running**: If the user rejects the gate approval, then the pipeline will abort. This behavior usually occurs in Continuous Delivery pipelines.
 - Continue running: If the user rejects the gate approval, the pipeline still completes the task and then continues to the next gate or stage task with a known error. This behavior works best for pipelines in traditional software releases.
- If a gate has approvals, one or more of these actions occurs:
 - Email notifications are sent to the specified approvers.
 - If the approver approves the request, the pipeline's progress continues.
 - If the approver rejects it, the pipeline will end with appropriate error.
- Approval gates are optional and if a gate has no approvals, the pipeline's progress continues.

For example, if your pipeline has multiple stages for testing, you may want to set entry and exit criteria for each stage to ensure that the code is testing properly in the software delivery life cycle.

- In the system test stage, the QA engineer needs to approve the start of t the system testing workflow. The same user or someone else in the QA group can review the system test results before approving the exit criteria.
- Before anything can enter the production stage, you may want to have entry gate approvals by various individuals.

Pipeline UI

This section shows how pipelines are represented in the ElectricFlow UI.

All of the information described in the previous sections is summarized in the pipeline definition in the Pipelines List.

Pipelines List

The Pipelines List shows all the pipelines that you have permission to view.

To open it from the Home page, click **Pipelines > All Pipelines**.

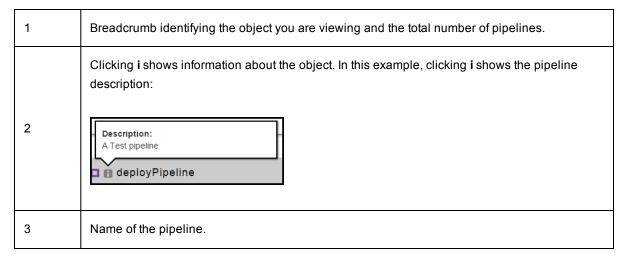
This is a Pipelines List:



Each row in the list is a *pipeline definition*, consisting of the pipeline name, the project to which the pipeline belongs, and the number of stages in the pipeline.

For each pipeline, you can perform these actions:

- Clicking the green **Run pipeline** button (number 7) starts a pipeline run.
- Clicking the **Expand** button (number 8) shows the previous pipeline runs.
- Clicking the View Details button (number 9) opens the Pipelines Visual Editor.
- Selecting All or None and then clicking the Delete button removes the selected pipelines from the list.
- Clicking the **Add+** button adds a pipeline.

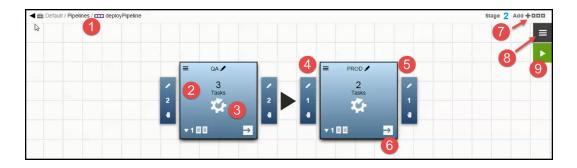


	The default is to show the applications for all projects.	
4	If you want to see only the objects for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:	
	Click on the name of one or more projects.	
	Enter the search criteria in the Search field. The projects that match the criteria appear in the list.	
	If there are no matches, a message appears stating that there are no resource templates in the selected projects.	
5	The project to which the pipeline belongs.	
6	Number of stages in the pipeline.	
7	Click the Run pipeline button to run the pipeline when the button is green.	
8	Click the Expand button to view the previous pipeline runs.	
9	Click the View details button to go to the Pipelines Visual Editor.	

Pipelines Visual Editor

In the Pipeline Visual Editor, you can author and edit a pipeline as well as view the underlying platform-level details.

This is a pipeline with two stages: QA and PROD.



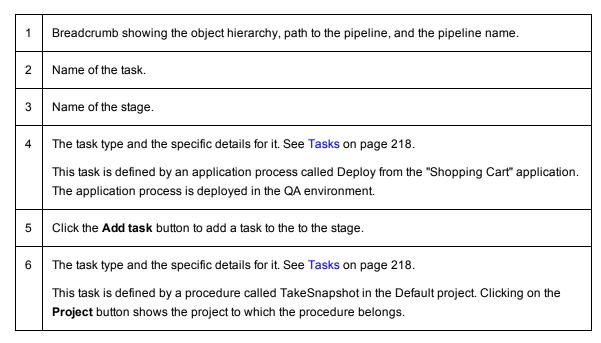
You can perform these actions:

- Clicking the **Add new stage** button (number 7) will add a new stage to the pipeline.
- Clicking the **Menu** button (number 8) will open a menu where you can get more details about the pipeline and take actions on the pipeline.
- Clicking the **Run pipeline** button (number 9) will start a pipeline run.
- In a stage, clicking the **Define stage tasks** button (number 6) will open a dialog box where you can define tasks for the stage.

1	Breadcrumb showing the object hierarchy, path to the pipeline, and the pipeline name.	
2	Stage in the pipeline, a milestone or checkpoint in the software delivery process.	
3	Tasks in a stage, the automation logic for a stage. See Tasks on page 218.	
4	Entry gate, consisting of the approvals to enter a stage. See Entry and Exit Gates on page 220.	
5	Exit gate, consisting of the approvals to exit a stage. See Entry and Exit Gates on page 220.	
6	Click the Define stage tasks button to open the task list for the stage.	
7	Click the Add new stage button to add a stage.	
8	Click the Menu button to open the menu where you get more details about the pipeline and take actions on the pipeline.	
	Clicking Details opens the Pipeline Details dialog box where you can edit the pipeline name and description.	
	Clicking Properties opens the Properties dialog box where you can view, add, and delete the properties associated with the pipeline.	
	Clicking Parameters opens the Parameters dialog box where you can view, add, and delete the parameters attached to the pipeline.	
	Clicking Access Control opens the Access Control page where you can view, add, and edit the privileges for the pipeline.	
	Clicking Track Changes shows the Change History for the pipeline and the objects in it. ElectricFlow tracks the changes to tracked objects including applications, procedures, workflows, workspaces, resources, and project-owned components such as library components and records a Change History of the historical states of the system and the changes between them.	
	Clicking Delete deletes the pipeline after you verify that you want to do this.	
9	Click the Run pipeline button to run the pipeline. The Pipeline Runs List shows the real-time progress of the pipeline run. Clicking in a stage will open the Stage Summary where you can view task details and progress of each task.	
	You can also view the results of previous pipeline runs.	

The task list for the stage shows all the tasks for the stage that are executed in sequential order.





Task Definitions

Tasks are defined in the Task Definitions dialog box described in this section.

There are two ways to get to this dialog box:

• When you want to add a task, click the Add task button.

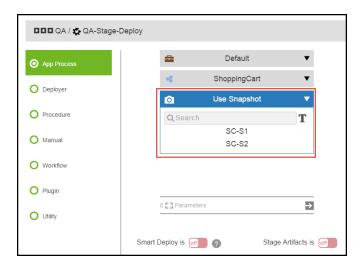


When you want to view or edit a task in the task list, select a task and click the button in its row.

Depending on the task type that you select, different fields appear on the right side oft the dialog box.

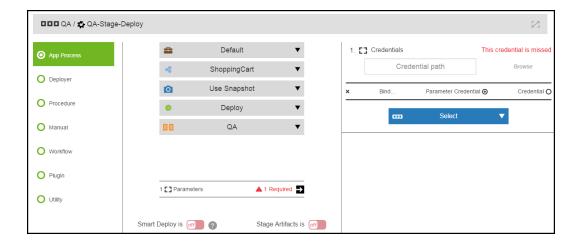
If you want to deploy an application, select **App Process**. You need to select an application, a snapshot, an application process, and the environment where the application process will run. If the parameters were created for application process, the Parameters field is enabled. You can also enable or disable **Smart Deploy**, where the system deploys only the artifacts that have not been deployed to a resource, specific versions of the artifacts, or artifacts that have not yet been deployed to new resources.

This example shows how to use a parameter to specify the snapshot. When you click in the "Select a Snapshot" field, these options appear:



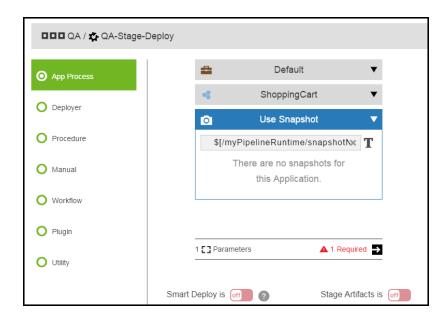
There are two ways to specify a snapshot:

• Select a specific snapshot for the application, such as **SC-S1**.



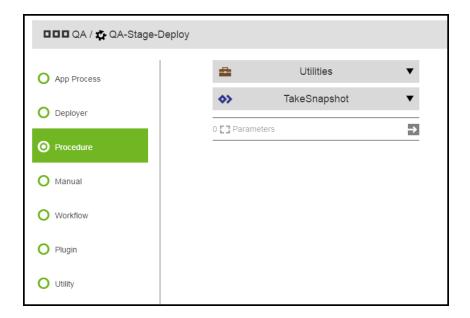
• Enter the snapshot as a parameter.

Enter this parameter: \$[/myPipelineRuntime/snapshotName].



If you want to run a set of best practices, subroutines, modules, or functions that you can create and reuse (such as a script), select **Procedure**. You need to select a project and a procedure in that project. If the parameters were created for procedure, the Parameters field is enabled.

This example shows that the task is a procedure called TakeSnapshot in the Utilities project.



Pipeline Stage Summary

A pipeline stage summary shows various status and metrics relevant to the pipeline when a pipeline stage is executed. For example, the summary could be the scan results from the code coverage tool or links to reports that are generated during the pipeline run. You can view the pipeline stage summaries during the pipeline run or for completed pipeline runs.

From the Pipeline Runs view, you can click on a specific stage to see that stage's summary. The stage summary shows two sets of stage-specific information.

- · Deployment summary consisting of:
 - A system-generated deployment summary.
 - Applications deployed through the pipeline and the environments where they are deployed.
- User-generated summary consisting of:
 - Links to web sites for reports and other information about the deployed artifacts.
 - Property information about the pipeline stage and the objects in it.

Creating User-Generated Data for the Stage Summary

If the task in a stage is defined by a procedure, a workflow, or an application process that includes a procedure as a process step, you can create user-generated data using ectool.

Stage Summary Information as Links to a Web Site

To create links to web site:

1. Create a procedure and the underlying steps with commands like the following:

```
ectool setProperty "/myPipelineStageRuntime/ec_summary/report" --value "<html><a hr
ef=\"http://www.myreportserver.com\" target=\" blank\">Global Report</a></html>"
```

2. Call this procedure from a task in the stage.

Generic Property Information

1. Create a procedure and the underlying steps with commands like the following:

To create a summary property:

```
ectool setProperty "/myPipelineStageRuntime/ec_summary/testCoverage "test value"
To create a property entry:
ectool setProperty "/myPipelineStageRuntime/ec_summary/testCoverage" --value "80%"
```

2. Call this procedure from a task in the stage.

Go to the following sections to see examples of user-generated data in the stage summary.

Viewing the Stage Summary During a Pipeline Run

To view the stage summary of a stage during its pipeline run, start the pipeline run by selecting a pipeline in the

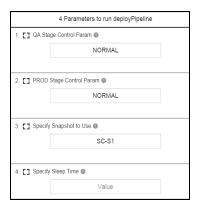


Pipelines list and clicking the

button:

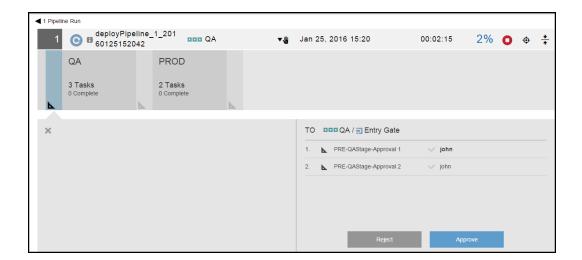


This pipeline requires input before the applications to be deployed through the pipeline. Enter the required parameters for pipeline, and click **OK**.



The pipeline runs. You can view its progress on the Pipeline Run page.

When manual approvals are required at the entry and exit gates to a stage, the system waits for the approvers to approve or reject before proceeding to the next step in the pipeline.



As the pipeline progresses through a stage, you can view its progress by clicking in the stage, which opens the pipeline stage summary.



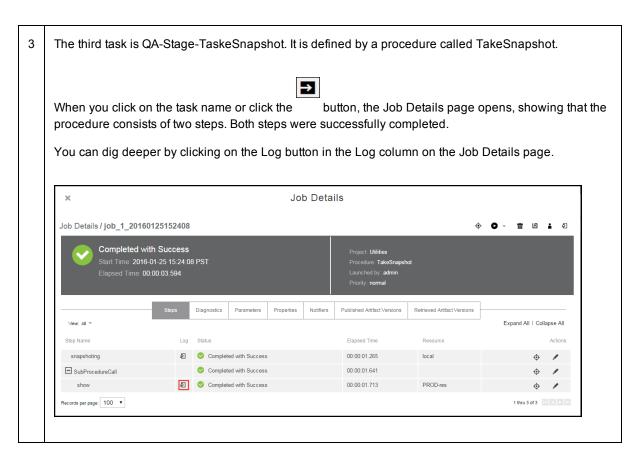
For the QA stage, the stage summary shows this information:

- Environment summary—The first two tasks (numbers 1 and 2 above) show what application processes were deployed and where they were deployed. The summary consists of the application, the application process being deployed, and the environment to which the application is deployed.
- User-generated summary–You can create two types of summary data, property information (consisting of text and/or numbers) and links to web sites (number 3 above).

In this example, the property information is testCoverage at 80%, and the link to the report is Global Report. When you click Global Report, the web site http://www.myreportserver.com opens.

The stage summary also shows information about the tasks:

- The first task is QA-Stage-Deploy. It is defined by the ShoppingCart application. The Deploy application process was deployed in the QA environment. The green check mark means that the task was successfully completed.
- The second task is QA-Stage-Deploy-with-Snapshot. It is also defined by the ShoppingCart application. The "Deploy" application process was deployed in the QA environment. The green check mark means that this task was also successfully completed.



Viewing the Stage Summary for a Completed Pipeline Run

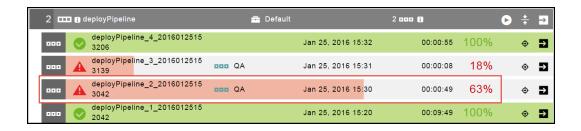
To view the stage summary of a stage in a completed pipeline run, start by selecting a pipeline in the Pipelines



List and clicking the button:

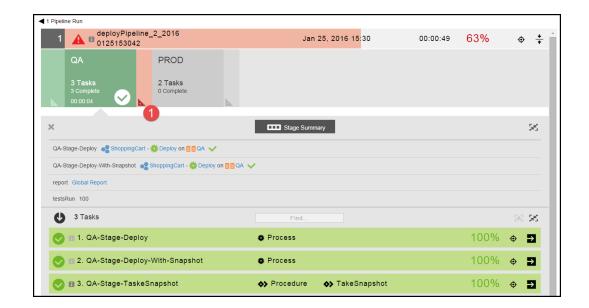


Clicking on the name of pipeline shows the list of current and completed pipeline runs.



To troubleshoot the pipeline, select the pipeline run that completed with errors. Click the name of the pipeline run to view a representation of the pipeline, and then click in a stage to view the stage summary for that stage.

This is an example of a stage summary:





To find out why the pipeline was completed with errors, click the button (number 1 above) in the stage summary.

The summary of the Exit Gate for the QA stage shows that one of the required approvals were rejected, so the pipeline stopped with errors.



Credentials in Pipelines

Credentials are passed to pipelines through applications in pipeline stages. If an application has required inputs (parameters) and any of the inputs is a credential parameter, you can enter the path to the credential, browse to it, select a credential parameter (**Parameter Credential**) or a credential binding to a pipeline task (**Credential binding**), or select a user-defined credential that is attached to the project associated with the pipeline.

A credential must be explicitly attached to the object using it so the server can perform an access control list (ACL) check at the definition time and limit the visibility of the password. To support accessing credentials at the pipeline task level, these tasks need to have the appropriate credentials attached. This is done implicitly by the UI when you are binding credentials to the pipeline tasks. If you are using ectool, use the attachCredential API command to perform the same operation.

Example: Authoring a Pipeline with Manual and Utility Tasks

This example shows how to tasks to a pipeline while defining a stage. The tasks include a manual task and a utility task, add a manual task to a pipeline while defining a stage.

- 1. Choose a stage in the pipeline.
- 2. Click the **Define stage tasks** button in the lower right corner of the stage.

The task list opens. When the task list has one or more items, it shows the tasks in sequential order for the stage.

- 3. Add a task to the stage.
- Enter the name of the new task in the Name field, and enter an optional description of it in the Description field.
- 5. Set the position of the task.
 - a. Select Before or After.
 - b. Click the down arrow to select the task that the new task will precede or follow.

Example:



- 6. If you want to add another task, click + Add another and repeat the previous two steps to add it.
- 7. If you do not, click Done.

The list shows the tasks in the order that they will be executed.

Example:



8. Click Requires Definition for the task you added.

The dialog box to define the task opens.

9. Enter the settings to define the task as a manual task, and click **OK** to save the task definition.

Example:

a. Select **Manual** to define the task as a manual task requiring human action in the **Instruction** field.

The instructions can include property references using the "\$[]" notation for file links, which is needed when referencing a task instruction in an email notification template. For example, a manual task could have this instruction:

Please check the Clover report, and make sure there is at least 80% code coverage before you approve.

http://myserver.com/reports/\$[/myPipelineRuntime/ec_summary/report_name]

- b. Click in the Assignees field.
- c. Enter the name of the users or groups assigned to the task, and select the one or more users or groups.

Example:



You can also use the \$[] notation to specify the parameter



d. Click **OK** to save the task definition.

The task list now shows the task as a manual approval requiring human intervention.

Example:





- 10. Click to add another task to the stage.
- 11. Enter the name of the new task in the Name field, and enter an optional description of it in the Description field.
- 12. Set the position of the task.

Example:



13. Click Done.

The task list now shows the new task.

Example:

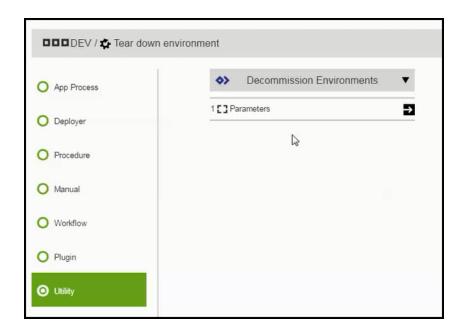


14. Click Requires Definition for the task you added.

- 15. Enter the settings to define the task as a utility task, and click **OK** to save the task definition. A utility task, also referred to as a utility function, is a higher-order operation than a third-party plugin, that you can use in pipeline modeling.
 - a. Select **Utility** to define the task as a utility task.
 - b. Click the down arrow in the "Select procedure" field and select **Decommission Environments**.

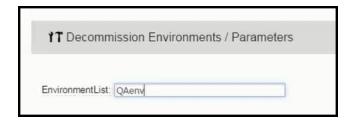


c. Click in the Parameters field to enter the parameters for the selected utility function.
 Example:



The dialog box to enter the parameter values opens.

d. Enter the parameter values and **OK**



The task list now shows a manual task as the first task and a utility task as the last task in the stage.

Example:



Example: Authoring and Running Pipelines

This example shows how to create a new pipeline, define it, and view the pipeline run results.

In the following steps, you will learn how to

- 1. Create a pipeline.
- 2. Edit the stage details.
- 3. Define the stage.
- 4. View and edit the pipeline definition.
- 5. Run the pipeline.
- 6. Approve the entry to or exiting from a stage
- 7. View pipeline runs.

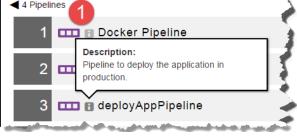
This is a Pipeline List:



Breadcrumb identifying the object you are viewing and the total number of pipelines.

Clicking i shows the description of the pipeline.

1



Clicking the down arrow next to **All Projects** opens a drop-down menu that shows the list of available pipelines. The default is to show the pipelines for all projects.

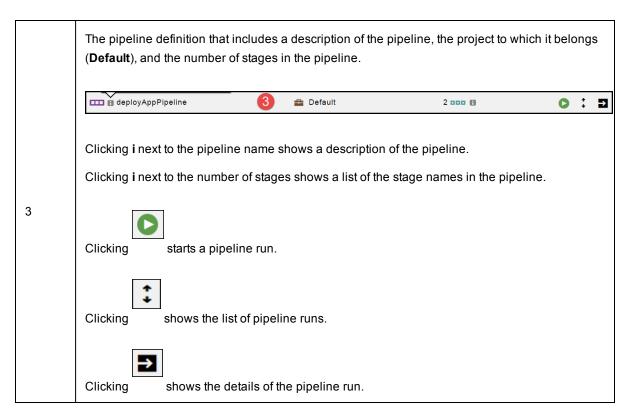


2

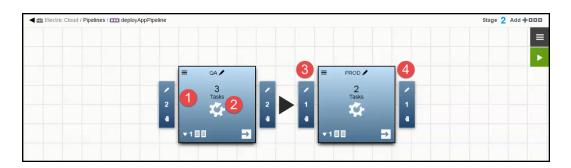
If you want to see only the objects for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

If there are no matches, a message appears stating that there are no resource templates in the selected projects.



This is the pipeline two stages: QA and PROD.



Element	Name
1	Stage
2	Tasks
3	Entry gate
4	Exit gate

IMPORTANT: When an application, pipeline, or release is cloned across different projects, you may have to fix the references after the objects are cloned. For example, if a pipeline with an application reference in

"default" project is cloned to a different project, the application reference needs to be fixed after the pipeline is cloned.

Creating a Pipeline

There are two ways to create a pipeline:

- Create a pipeline from scratch.
- Create a pipeline based on an existing pipeline.

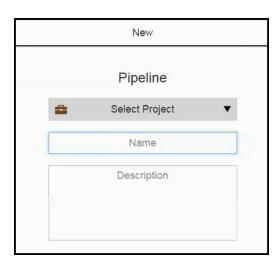
To create a pipeline component from scratch, start in the Home page and:

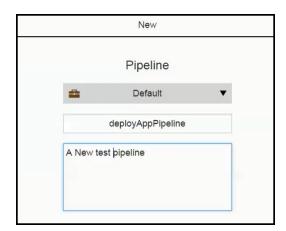
- 1. Go to the Pipelines List one of these ways:
 - Click on the Main menu button in the upper left corner, and then select Pipelines > All Pipelines.
 - Click Pipelines.

The Pipelines List opens.

- 2. Click Add + to add a pipeline.
 - a. To create a new pipeline:
 - i. Click Create new pipeline.
 - ii. Enter the pipeline name, select a project where the applications in the pipeline will be deployed, and enter a description of it.

Example:



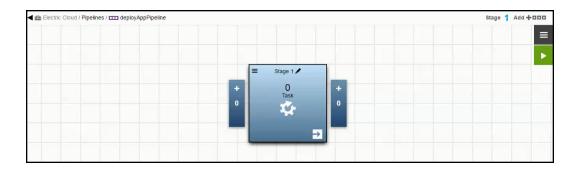


iii. Click OK.

The Pipelines Visual Editor re-appears.

It shows a pipeline with one stage.

Example:

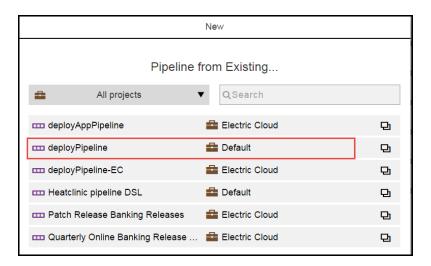


- b. To create a pipeline based on an existing pipeline:
 - i. Click Create from Existing.

A list of existing pipelines opens.

It shows the list of available pipelines. The default is to show the pipelines for all projects. If you want to see only the pipelines for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:Click on the name of one or more projects. Enter the search criteria in the Search field. The projects that match the criteria appear in the list. You can also search for one or more pipelines by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no pipelines in the selected projects.

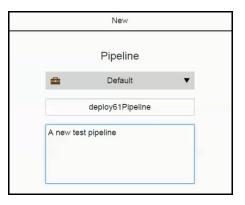
Example:



ii. Select a pipeline.

iii. Enter the pipeline name, select a project where the applications in the pipeline will be deployed, and enter a description of it.

Example:

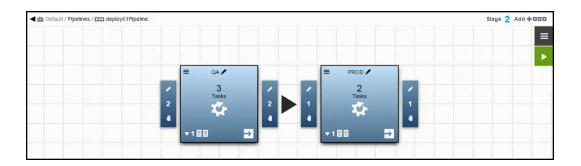


iv. Click OK.

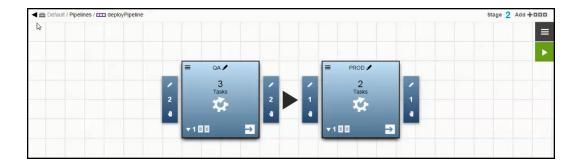
The Pipelines Visual Editor re-appears and shows a pipeline with the same stages, tasks, and gates as the existing pipeline.

Example:

The new pipeline that has the same settings as the deployPipeline pipeline:

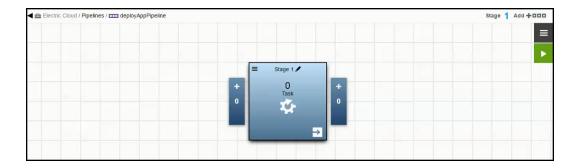


This is the deployPipeline pipeline:



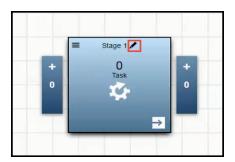
Editing the Pipeline Details

When a pipeline is created from scratch, it initially has one stage with the default stage name (Stage N) and no description of it. To provide meaningful information about it, change the stage name and description.



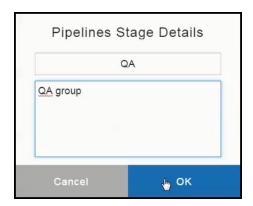
1. Click the **Edit** button next to the stage name:

Example:



2. Enter the stage name and a stage description.

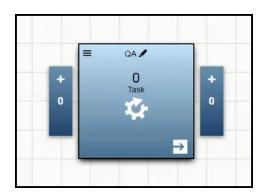
Example:



3. Click OK.

The pipeline now shows that the stage name is QA.

Example:



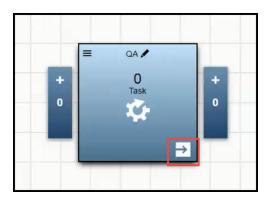
Defining the Stage

This section describes how to define a task in a stage.

Starting in the Pipelines Visual Editor:

- 1. Choose a stage in the pipeline.
- 2. Click the **Define stage tasks** button in the lower right corner of the stage.

Example:



The task list opens. When the task list has one or more items, it shows the tasks in sequential order for the stage.

- 3. If the stage has no tasks, do the following in a empty row:
 - a. Enter the name of the task in the Name field.
 - b. (Optional) Enter a description for the task in the Description field
 - c. Select the **On Error** condition. When the stage task fails, the pipeline either continues running and proceeds to the next task (**Continue running**) or stops running and aborts (**Stop running**). This gives you more control in the pipeline. The default is to **Stop running**.

Example:



Enter QA-Stage-Deploy in the Name field.



d. Click + Add another.

The task you just added is saved, and a new task is automatically added after it.

- 4. Enter the name of the new task in the Name field.
- 5. (Optional) Enter a description for the task in the Description field.
- 6. Set the position of the task.
 - a. Select Before or After.
 - b. Click the down arrow to select the task that the new task will precede or follow.
- Select the On Error condition. When the stage task fails, the pipeline either continues running and
 proceeds to the next task (Continue running) or stops running and aborts (Stop running). The default is
 to Stop running.

Example:



8. If you want to add another task, click + Add another and repeat the previous three steps to add it.

9. If you do not, click **Done**.

Example:



The list shows the tasks in the order that they will be executed.

Example:



10. Click Requires Definition for the first task in the list.

The dialog box to define the task opens.

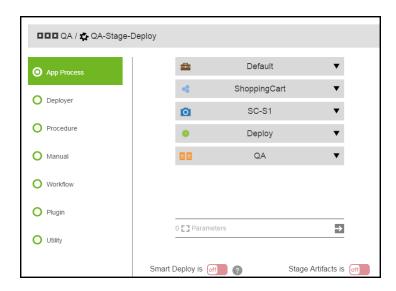
For more information about the task types, see Tasks on page 218.

11. Enter the settings to define the task, and click **OK** to save the task definition.

Example:

- a. Select App Process to define the task as an application process.
- b. In the Select Project field, select the Default project.
- c. Select the ShoppingCart application and Use Snapshot
- d. When you click in the **Select a Snapshot** field, a Search field opens where you can enter the criteria to search for a snapshot as well as the names of existing snapshots.
- e. Select SC-S1 as the snapshot to use.
- f. Select the **Deploy** process and the **QA** environment where application will be deployed.

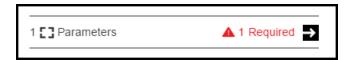
Example:



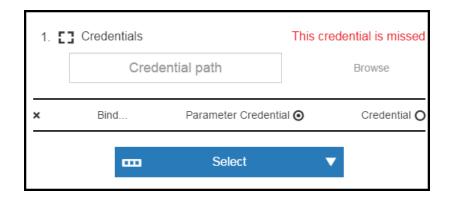
Go here for more information about the different ways to define a stage taskPipeline Concepts on page 218

g. The Parameter field is disabled because the **Deploy** process does not have any required parameters.

Example showing the Parameter field enabled:



If the process has a required parameter for the pipeline run, such as a credential parameter, a message with the number of required parameters (such as **1 Required**) appears in the Parameter field. You can enter the path to the credential, browse to it, select the **Parameter Credential** or **Credential** binding, or select a user-defined credential that is attached to the project associated with the pipeline.



12. Click **Requires Definition** for the second task in the list, define the task, and click **OK** to save the task definition.

The dialog box to define the task opens.

Example:

- a. Select App Process to define the task as an application process.
- b. In the Select Project field, select the Default project.
- c. Select the **ShoppingCart** application.
- d. Select the **Deploy** process and the **QA** environment where the **ShoppingCart** application will be deployed.
- e. If the **Deploy** process has one or more required parameter for the pipeline run, the Parameter field is enabled. Go here for more information.

13. Click **Requires Definition** for the third task in the list, define the task, and click **OK** to save the task definition.

Example:

- a. Select Procedure.
- b. When you click in the **Select Project** field, a drop-down list of available projects opens. Select the **Utilities** project.
- c. When you click in the **Select Procedure** field, a drop-down list of available procedures opens. Select the **TakeSnapshot** procedure.

Example:



If the **TakeSnapshot** procedure has required parameters for the pipeline run, the Parameter field is enabled. Go here for more information.

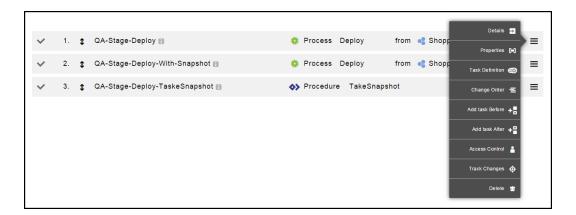
14. (Optional) Change the order in which the tasks will be executed. Go to Pipeline Concepts on page 218 for instructions.

Viewing and Editing the Pipeline Definition

To view the pipeline definition:

- 1. Go to the Pipelines List.
- 2. Choose a task and click the **Menu** button on the end of the row.

Example:



3. Click Task Definition.

The task definition appears.

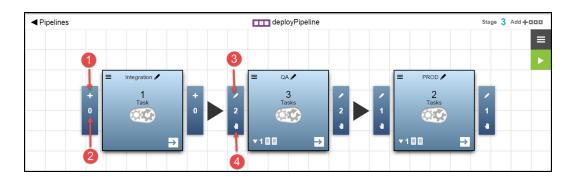
4. Change the settings that you want to modify, and click **OK** to save them.

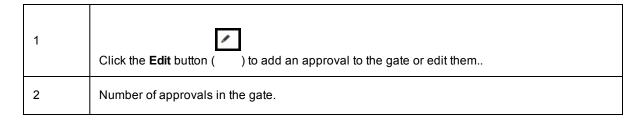
Defining Gate Approvals

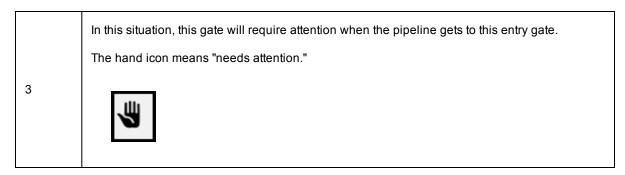
This section describes how to create approval rules on the entry and exit gates.

Gates for Pipelines

This example shows entry and exit gates in a pipeline in the ElectricFlow UI.







To add approvals to a gate that does not have any, click + in that gate:



Creating One Approval Rule in a Gate

This example shows how to create one approval rule with many approvers, where only one of them needs to take action.

Starting in the Pipelines Visual Editor:

- 1. Go to the gate where you want to add the approval rule.
- 2. Click + to add a gate approval rule.

The page to add or edit approval rules for the gate opens.

This is an undefined approval rule:



- 3. In the Name field, enter PRE-PRODStage-Approval.
- 4. In the **Approvers** field, add the users named **john** and **bob** and the group called **SCMAdmins**. Start typing the user or group name to find available users or groups.
- 5. Keep the **Default Pipeline Notification** as is (do not change it).
- 6. Click Done.

For the gate, there is one approval rule that one of the approvers (bob, john or the SCMAdmins) must approve for the pipeline to progress.



Creating Multiple Approvals in a Gate

This example shows how to create multiple sequential approval rules where sequential approvals from everyone is needed.

Starting in the Pipelines Visual Editor:

- 1. Go to the gate where you want to add the approval rules.
- 2. Click + to add a gate approval rule.

The page to add or edit approval rules for the gate opens.

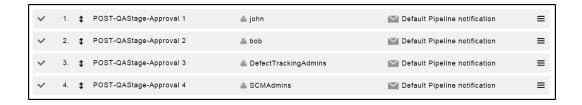
This is an undefined approval rule:



- 3. In the Name field, enter POST-QAStage-Approval 1.
- 4. In the **Approvers** field, add the user named **john**. Start typing the user or group name to find available users or groups.
- 5. Keep the Default Pipeline Notification as is (do not change it).
- 6. Click + Add Another.
- 7. In the Name field, enter POST-QAStage-Approval 2.
- 8. In the **Approvers** field, add the user named **bob**. Start typing the user or group name to find available users or groups.
- 9. Keep the **Default Pipeline Notification** as is (do not change it).
- 10. In the Name field, enter POST-QAStage-Approval 3.
- 11. In the **Approvers** field, add the group named **DefectTrackingAdmins**. Start typing the user or group name to find available users or groups.
- 12. Keep the **Default Pipeline Notification** as is (do not change it).
- 13. In the Name field, enter POST-QAStage-Approval 4.
- 14. In the **Approvers** field, add the group named **SCMAdmins**. Start typing the user or group name to find available users or groups.
- 15. Keep the **Default Pipeline Notification** as is (do not change it).
- 16. Click Done.

For the gate, there are multiple approval rules where each approvers needs to approve for the pipeline to progress. The approvers must approve in this order:

- 1. john
- 2. bob
- 3. DefectTrackingAdmins
- 4. SCMAdmins



Running Pipelines

This section shows how to run a pipeline. It also has examples showing different pipeline run scenarios.

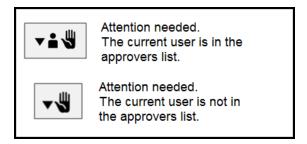
Starting from the Pipelines Visual Editor:



- 1. Click the **Run pipeline** button () to start a pipeline run.
- 2. If the pipeline has required parameters to run the pipeline, enter the parameter values.

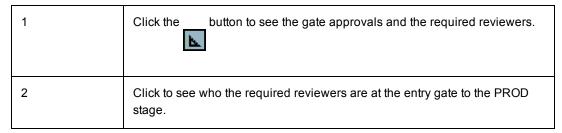
3. View the real-time progress of the pipeline run.

Note: These are the "Attention Needed" icons:

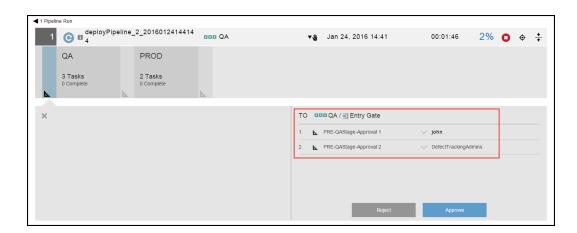


The following example shows that exit gate of the QA stage and/or the entry gate of the PROD stage has gate approvals that needs to be approved before the pipeline can enter the PROD stage.

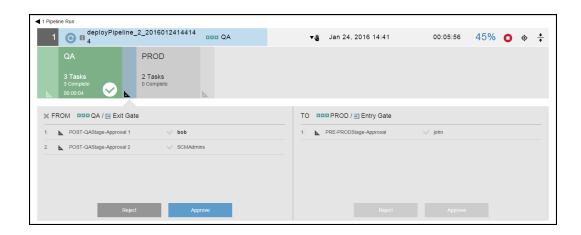




For the entry gate of the QA stage, the required approvers are the user named john and then the group named DefectTrackingAdmins.



For the exit gate of the QA stage, the required approvers are a user named bob (user) or a group named SCMAdmins and then a user named john.



4. Click in a stage to see the details of the pipeline stage summary.

For examples of pipeline stage summaries, see Pipeline Stage Summary on page 227.

5. Click View previous Pipeline Runs to see the previous runs of this pipeline.

For an example, see Example: Viewing Previous Pipe Runs on page 258.

Examples: Viewing the Details of a Pipeline Run

While a pipeline is running, you can click a stage to view its progress and open the pipeline stage summary.





For a task, click the () button to see the Job Details page:

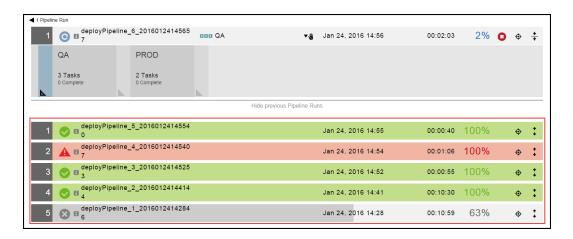


Example: Viewing Previous Pipe Runs

Click View previous Pipeline Runs to see the previous runs of this pipeline.



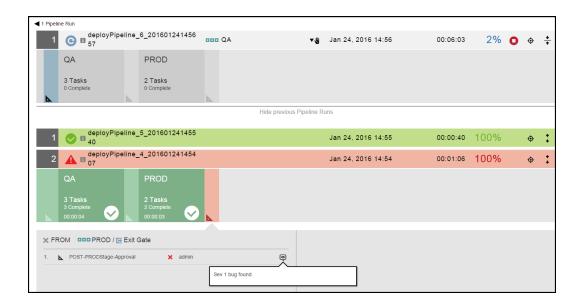
The previous pipeline runs appear.



*

Click the button for one or more previous pipeline runs.

In this example, the pipeline was not completed successfully. The approval at the exit gate of the PROD stage was rejected because a Sev 1 bug was found in the software release.



Viewing Pipeline Runs

The Pipeline Runs List summarizes the pipeline runs. Each row is an pipeline instance that has an icon and is color-coded to indicate the status. You can use filters and the Search field to find certain pipeline runs. See What's in the Pipeline Runs List on page 259 and Searching the Pipeline Runs List on page 262 for more information.

When you expand a row, you can view the details about the pipeline runs as well as the stage and task summaries for a pipeline instance. See Running Pipelines on page 255 for more details.

Opening the Pipeline Runs List

Click on the **Main menu** button on the Home page, and then click **Pipelines > Pipeline Runs** to open the Pipeline Runs page.

By default, the list shows all the running or completed pipeline runs that you have the permission to view.

What's in the Pipeline Runs List

Each row shows the status of the pipeline run, the pipeline ID and name, the user running the pipeline, when the pipeline was run, and how long the pipeline run lasts.

For each pipeline run, you can take the these actions:



- Clicking the Expand button () shows the pipeline run details and results of previous runs.
- Entering part or all of a pipeline name in the Find field (number 8 below) starts the search for pipelines
 matching the criteria. See Searching the Pipeline Runs List on page 262 for how to search in the
 Pipeline Runs List.
- Clicking on one of the filters (number 9 below) shows the pipeline runs whose status matches the filter.
 See Searching the Pipeline Runs List on page 262 for how to search in the Pipeline Runs List.

This is the Pipeline Runs List:



1	Breadcrumb identifying the object you are viewing and the total number of pipeline runs.
2	Find field. Enter the search criteria to find specific pipelines.

Filters presenting the status of the pipeline run.



Needs attention. Manual approval may be required to complete the current stage and go to the next stage.



Completed with errors



Running

3



Completed with success



Completed with warnings. The pipeline run could still be running, or it may have completed running.



Pipeline aborted. When you click the **Abort** button during a pipeline run, the pipeline will be aborted after the currently running tasks finish and no new tasks will be started.

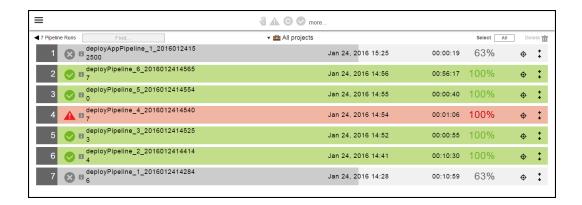
	Clicking the down arrow next to All Projects opens a drop-down menu.
4	Clicking the down arrow flext to Air Flojects opens a drop-down menu.
	It shows the pipeline runs for a specific pipeline. The default is to show the pipeline runs for all projects.
	If you want to see only the pipeline runs for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:
	Click on the name of one or more projects.
	 Enter the search criteria in the Search field. The projects that match the criteria appear in the list.
	You can also search for one or more pipeline runs by entering the search criteria in the Search field next to the All projects field. If there are no matches, a message appears stating that there are no pipeline runs in the selected projects.
5	Clicking the Abort button causes the current task to complete. No new tasks are started, and the current pipeline run stops.
6	Click the button to show more details about the pipeline run. Clicking it in this rows shows results in displaying more details about the pipeline run.
	1
7	This pipeline run is in progress. This shows that the pipeline is currently in the QA stage, an approver needs to do something to keep the pipeline running, when the pipeline is running and how long, and the percentage complete. You can abort the pipeline (number 5 above) and also see more details about this pipeline run (number 6 above).
8	The pipeline run was completed successfully.
9	The pipeline run was completed with errors.
10	The pipeline run was aborted.

Searching the Pipeline Runs List

To focus on specific pipeline runs, ElectricFlow provides several ways to quickly filter and search for pipeline runs, using a combination of filters and the **Find** field. You can also use these methods with the search operations on other ElectricFlow features.

- Filtering with Filter Icons on page 263
- · Searching With the Find Field on page 264
- Filtering by Projects on page 264
- Filtering and Searching With the All Methods on page 265

When no filter icons are selected, the **Find** field is empty, and **All projects** is selected, all of the pipeline runs appear in the list, such as:



Filtering with Filter Icons

When you select a filter, the Pipeline Runs List is updated and shows only the pipeline runs with a status matching the filter. These examples show how to filter pipeline runs using the filter icons:

Running



· Competed with success



· Completed with errors



• Pipeline aborted. When you click the **Abort** button during a pipeline run, the pipeline will be aborted after the currently running tasks finish and no new tasks will be started.



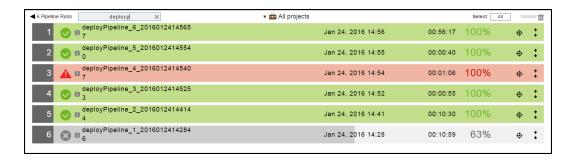
Searching With the Find Field

Start typing the name of the pipeline in the Find field.

For example, if you enter "my" as the search criteria, there are no pipelines that match the search criteria. The Pipeline Runs Lists returns with nothing in the list.

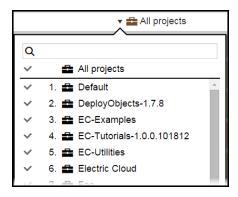


To find a pipeline starting with "deploypip" in the Pipeline Runs List, start to enter "deploypip" and notice that ElectricFlow updates the list and shows the pipelines with names that match the search criteria.



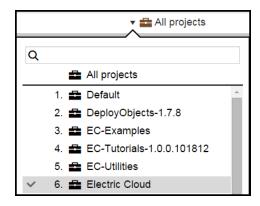
Filtering by Projects

Clicking the down arrow next to the **All projects** opens a drop-down list of all the projects you have permission to access.

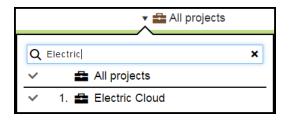


You can select the pipelines for specific projects one of these ways:

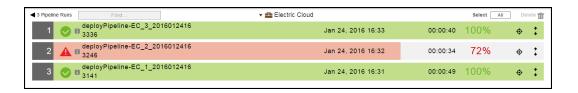
• Select one or more specific projects in the drop-down list, and then click outside the drop-down list.



• Enter search criteria to find a specific project, and then click outside the drop-down list.



The pipeline runs for the pipeline in the Electric Cloud project now appear in the Pipeline Runs List.



Filtering and Searching With the All Methods

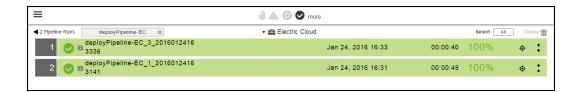
You can use all three methods to search for specific pipeline runs.

This example shows how to do this:



- Select (Completed with success).
- In the Find field, enter "deployPipeline-EC".
- Select the "Electric Cloud" project.

The pipeline runs that were completed successfully, start with "deployPipeline-EC", and are in the "Electric Cloud" project now appear in the Pipeline Runs List.



Troubleshooting Pipelines

This section shows how to troubleshoot your pipeline runs. If your pipeline result is not green, you may have a problem that you want to resolve. This topic describes some of the common steps that you may need to use when determining where there are issues in the pipeline process.

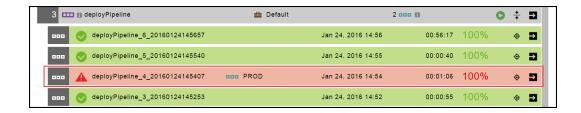
Example: Error Due to Gate Approvals

Starting in the Pipelines List:

1. Choose a pipeline to troubleshoot.

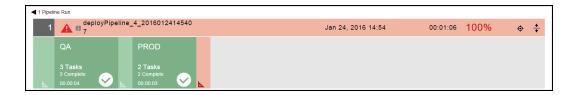
In the example in this topic, choose the deployPipeline.

2. To view details about a specific pipeline run, choose a run and click the View Details button in the row.



Clicking the **View Details** buttons on any task will show you the details of the underlying work that was done to complete the task.

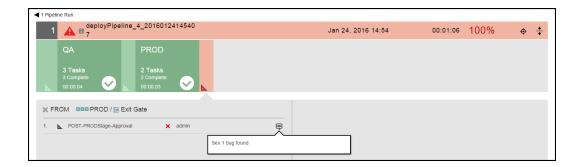
The pipeline run summary appears. There is an error in the exit gate of the PROD stage.



3. To get more information about the exit stage, click in the exit gate.

Clicking in the exit gate will show the approval status of gate.

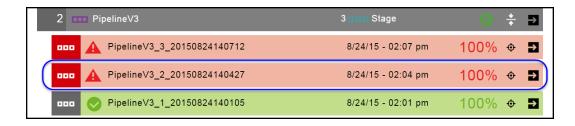
In this example, the pipeline was not completed successfully. The approval at the exit gate of the PROD stage was rejected because a Sev 1 bug was found in the software release.



Example: Error in a Pipeline Stage

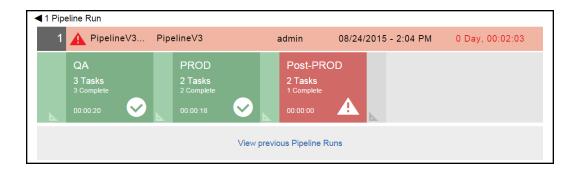
Starting in the Pipelines List:

- 1. Choose a pipeline to troubleshoot.
- 2. To view details about a specific pipeline run, choose a run and click the View Details button in the row.



Clicking the **View Details** buttons on any task will show you the details of the underlying work that was done to complete the task.

The pipeline run summary appears. There is an error in the Post-PROD stage.

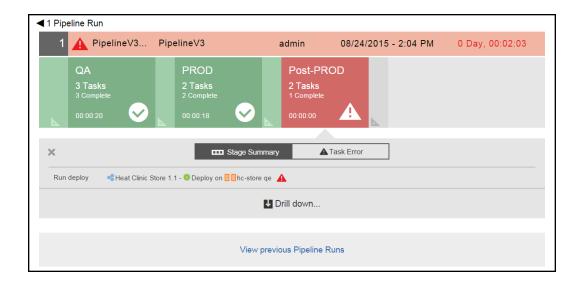


3. To get more information about the Post-PROD stage, click in the stage.

Clicking in the Post-PROD stage will show the Stage Summary and Task Error details.

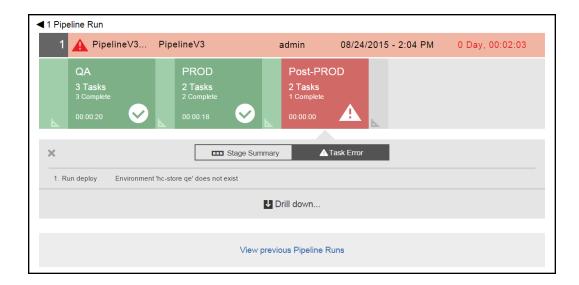
The Stage Summary shows the task definition.

In this pipeline, the task is an application process.

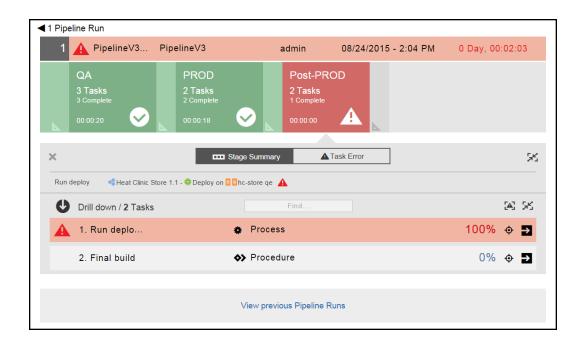


The Task Error information appears.

The error occurred in the "Run deploy" task. The problem is that the environment to which the application is mapped does not exist.



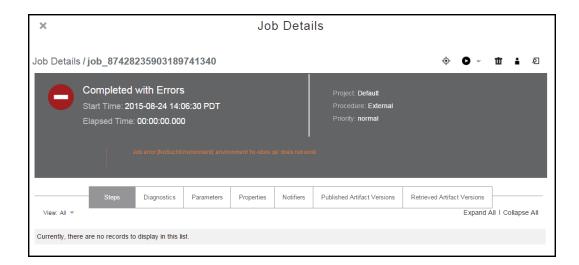
4. Click Drill down to see more stage details.



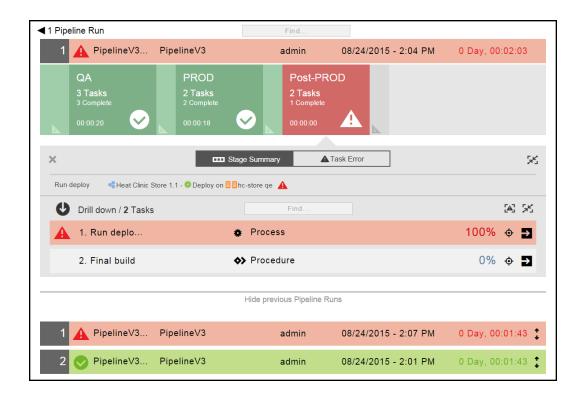
Notice that the "Run deploy" task failed and that the "Final build" task never ran because of the "Run deploy" task.

5. Click the View details button.

The Job Details page opens and shows that the "Run deploy: tasks was completed with errors. The problem is that the environment in the task definition does not exist.



6. Click View previous Pipeline Runs to see the previous runs of this pipeline.



7. Choose a pipeline in the Previous Pipeline Runs list, and click the **Expand** button.



The pipeline run summary for that pipeline appears.

8. Repeat Step 3 to Step 7 to do more troubleshooting.

Examples: Task List in a Stage

These examples show other actions you can perform in a task list.

Change the Order in Which the Tasks Are Executed

This section describes how to quickly move a task to a different position in the sequentially ordered task list and change the order in which the tasks will be executed.

Starting in the task list:

- 1. Choose the task that you want to move.
- 2. Click and hold the up/down arrow next to the task name to move the task row to its new position in the list.



The task is now in the first in the list.



Sample Task List for a Stage

This example describes the information in a sequentially ordered task list.



1	Breadcrumb showing the path to the pipeline.
2	Name of the stage. When you click the down arrow, a list of the stages in the pipeline opens. Select a stage to quickly to go to a different stage and view the task list for it.

3	Click the Select button to select All or None of the items in the list.
	Click the Delete button to delete the selected items.
4	Number of tasks in the stage. Click the Add task button to add a task to the stage.
5	Click the Menu button to open the menu where you get more details about the stage and take actions on the stage.
6	Task definition consisting of the task defined by an application process. It consists of a task named QA-Stage-Deploy and an application process called Deploy, which is part of the ShoppingCart application that is mapped to the QA environment. Clicking the Menu button opens the menu where you get more details about the task and take actions on the task.
7	Task definition consisting of the task defined by a procedure. It consists of a task named QA-Stage-TakeSnapshot and a procedure called TakeSnapshot in the Utilities project. Clicking the project button shows the project to which the procedure belongs.
	♦> Procedure TakeSnapshot
	Clicking the Menu button opens the menu where you get more details about the task and take actions on the task.

Release Management

The Release module captures, executes, visualizes, and controls the life cycle of multiple application enterprise releases. Using the Release capability, outputs from multiple teams can be coordinated to produce a final release to be pushed to production. You can create a Release model consisting of multiple applications deployed on different systems, such as traditional, cloud, mainframes, and remote servers. You can then run the Release pipeline reliably and repeatedly until the software release is complete.

You can use the Release feature to deliver software releases through traditional methods where changes from multiple applications are bulked up and are released together as a unit through a pipeline on a regular cadence (schedule) such as weekly, monthly, quarterly, or other frequency. In such traditional scenarios, the Release capability allows you to manage dependencies between multiple applications. In addition, you can even use the Release feature to cater to Continuous Delivery-style releases where code check-in from a developer can traverse all the way to production through various pipeline stages and approval gates. The Release feature allows ElectricFlow to manage all the releases in one platform, regardless of the release methodologies used by the release team. Anyone can quickly get a status of a release and then access the Release details. This information can also be used to troubleshoot and improve the software release process.

· Visibility and coordination

The *Release Dashboard*, also referred to as the *Release List*, allows the release team to get a bird's eye view of all the releases that are planned, active, or completed. For a specific release, you can easily see its status, what the milestones are, if the release is blocked for some human intervention, and the Release's progress.

Path to Production (or what is where)

For each stage in the Release pipeline, the Path-to-Production view shows the applications deployed in the release, the deployed versions, and the environments to which they are deployed. The details in this view include the application versions, the snapshots, and the environments that are not in compliance with the bill of materials throughout the Release.

Control

ElectricFlow keeps track of the data from the multiple teams involved in the release. The Release feature captures, validates, coordinates, and tracks all the details in one place, where everyone can access the *Release definition*, bill of materials (which applications to release), the pipeline controlling the release process, and environment to use across stages, approval conditions, release configurations, and so on.

When the Release is run, ElectricFlow automatically coordinates the pipeline information needed to deploy applications in each stage. If you need to add an application, change an application version, or deploy to a different environment, you need to change only the bill of materials. You do not need to revise the pipeline definition or pipeline tasks.

You can reuse the pipeline to deliver multiple releases of the same software.

The *Release manifest*, also referred to as the *Release definition*, specifies the release process flow for the application or applications to be released, which include:

- The pipeline controlling the process.
- The applications in the pipeline for multiple-tiered deployments.

- How to deploy tasks and processes using run-time parameters and other settings, such as smart deploy and artifact staging.
- Where versioned objects will be deployed (multiple-tiered environments).
- The approvers for manual process steps, manual tasks, and pipeline gates.

Release Concepts

Releases allow you to orchestrate multiple applications across multiple environments and to manage the software release life cycle using one platform. It brings together everything you need to take your software through its journey to production, regardless of your software release processes. You can manage traditional release processes (culminating in one big software release at the end of the software delivery life cycle), CI scenarios (build-test automation), CD scenarios (deployment automation), or a combination of these methods using a single platform. Dependencies can also be managed and controlled with this platform.

Modeling Releases

You can model a Release by defining the Release process that consists of

- The bill of materials consisting of the multiple applications that you want to deploy.
- One or more environments where the applications will be deployed or installed during the Release run.
- Pipelines and any inputs to the applications (parameters created for them) specifying how the
 applications will be deployed. Pipelines are the paths to production, and all the other objects in the
 Release are linked to the selected pipeline.
- The pipeline approval gates that determine when tasks in the stage can start or when the pipeline can enter or exit a stage.
- Approvers and task assignees perform specific activities as the Release runs. Individual users and teams are notified when an action needs to be taken.

The Release feature leverages the pipeline capability and a special type of pipeline task called *Deployer* to truly deliver an *auto-flexing* mechanism where the same pipeline can be used across multiple releases. For example, you have a standardized process for your major releases, and the only thing that normally changes are the applications and their released versions. In this case, you do not want to specify the exact applications to deploy as part the pipeline tasks, which do not allow for pipeline reusability. This is where a Deployer task in the pipeline and Release capability work together to provide reusability.

The Deployer task, previously mentioned in Pipelines on page 217, is a special type of task specified in the task list of a pipeline stage. It gets the actual list of applications (and the number of applications) to deploy from the bill of materials in the Release definition. This list is considered the payload of the Release. At runtime, the list of applications to deploy is passed to the pipeline through the Release object. The pipeline then deploys the applications in sequential order. So you can create one pipeline that works for different software releases.

In a pipeline stage, you can define a task using a out-of-the-box third-party integrations, also referred to as plugins, such as:

 Artifact repositories products such as Git, Perforce, Subversion, Microsoft Team Foundation Server (MS TFS), Accrue, ClearCase, Repo, and Vault.

- Issue tracking products that provide application developers with requirements tracking such as HP
 Application Lifecycle Management (HP ALM), Bugzilla, ClearQuest, Fortress, JIRA, HP Quality Center
 (HPQC), Rally, IBM Rational Team Concert, TeamForge, and TestTrack.
- IT service management (ITSM) products such as ServiceNow.
- Project management products such as Rally.

For a complete list of supported plugins, go to the Electric Cloud Plugins Directory.

For more information about tasks in pipeline stages, go to Tasks on page 218, Defining the Stage on page 246, and Viewing and Editing the Pipeline Definition on page 251.

ElectricFlow provides multiple project support for objects in application deployments, pipelines, and Releases. These objects as well as the objects belonging to them can be in any project within ElectricFlow.

All objects in ElectricFlow have these capabilities for better management and maintenance of end-to-end software releases:

- Projects are top-level security containers that provide access control at the project level. All objects in a
 project inherit the access control settings from the project to which they belong.
- Projects are also ways to provide logical groupings of objects. If your deployment is large and has many
 production environments, you can organize production environments by project, with each project
 having its own access control settings. This makes it easier to manage and maintain multiple
 environments during the software delivery process.

Setting Up Releases

This section describes only how to set up releases. To get started defining the Release process, you select a pipeline on which the process is based. This pipeline needs to have a Deployer task in its stages. Then you select the applications that you want to deploy (the bill of materials), the environment where the applications will be deployed, and any inputs to the applications. To see how to do this in the ElectricFlow UI, go to Release Definition on page 283.

To get a bird's eye view of all the releases that are planned, active, or completed, go to Release Dashboard on page 300.

To view the bill of materials, go to Path-to-Production View on page 305.

To view the pipeline's progress during the Release run, go to Pipeline Stage Summary on page 227.

For instructions on how to set up objects in a release:

- Go to Example: Modeling and Deploying Applications on page 88 to set up applications.
- Go to Master Component Examples on page 47 for examples of how to create and add parameters to your applications. The steps apply to applications with master and normal components.
- Go to Pipelines on page 217 to set up pipelines.

ElectricFlow supports resource exclusivity at the platform level. This allows a procedure to lock a resource for a job step and execute the tasks. The job step acquires and retains exclusive use of a resource, ensuring that the resource is always available for the step. Go to Reserving a Resource for Job Step Duration on page 982 for more information.

IMPORTANT: When an application, pipeline, or release is cloned across different projects, you may have to fix the references after the objects are cloned. For example, if a pipeline with an application reference in "default" project is cloned to a different project, the application reference needs to be fixed after the pipeline is cloned.

IMPORTANT: When you export a project while a pipeline is in progress, only the full export includes flow runtimes that have been completed. If you want to include in-progress pipeline runs in the path-to-production view and the visual indicators showing their percentage completed in the Release dashboard, set the excludeJobs argument to 0 or false in the export command. When the XML file is imported, the inprogress pipeline runs in the imported project are displayed in the path-to-production view and the Release dashboard.

Controlling the Releases

ElectricFlow provides these ways to ensure that the Release runs are successful:

· Manual gates:

For each stage in the Release pipeline, an entry gate controls when the tasks can start. If the conditions and approvals are met, the process flow can proceed and the tasks can be executed. The process flow can complete the stage and proceed to the next stage when the exit conditions and approvals are met.

ElectricFlow sends notifications to the appropriate approvers, users, and/or groups.

Go to Visibility and Status of Release Pipelines on page 277 for an example process flow.

· Preconditions:

You can set preconditions on procedure steps. Go to Step - create new or edit existing step on page 560 for how to set preconditions in procedure steps.

You can also set conditions in process branching when authoring application and component processes. These conditions determine the path that the Release takes through the process flow. Go to Process Branching on page 73 for how to set conditions on the transitions between process steps.

- ServiceNow integration: Use the ServiceNow plug to integrate IT service management (ITSM) with ElectricFlow.
- Plugins: You can use third-party plugins to execute the conditions and rules for approvals and quality gates.
- Postp: You can use postp to get data from the log files, make it appear on the Job Details page, and use it to take action on the Release run. Go to Postp on page 278 for more details.

Visibility and Status of Release Pipelines

From the Release Dashboard

The Release dashboard orchestrates all the information about the planned, active, or completed releases and displays the releases that you have permission to view in one place. At any time, you can quickly see the current status of the release, where it is in the process flow, any errors during the runtime, when human intervention is required to take some action or perform a task, approvals, and the start and end dates for the release. Go to Release Dashboard on page 300 for more information, and go to Example: Visibility in a Release

Pipeline from the Release Dashboard on page 279 for an example how to drill down from the top Pipeline Run view to the specific job details of the release tasks.

For a Release that is in progress or has been completed, clicking the Release status opens the Pipeline Run view. You can view the Stage Summary for the latest Release or view previous pipeline runs. The Stage Summary shows the pipeline's progress through each stage and evidence that the tasks were executed. Clicking in a stage task or a gate displays its Stage Summary. Go to Pipeline Stage Summary on page 227 for more information.

- When the Release is in progress, you can view the progress of the pipeline as it runs. Go to Viewing the Stage Summary During a Pipeline Run on page 228 for an example of this.
- You can also view a completed pipeline run. Go to Viewing the Stage Summary for a Completed Pipeline Run on page 230 for an example of how to do this.
- Clicking **View previous Pipeline Runs** displays the last five pipeline runs. After selecting one of these runs, you can view the Stage Summary to view more details about the pipeline run.
- If the pipeline has errors, you can also view what the error is and where it occurred in the pipeline in the Task Error view.



For example, you can view the progress and evidence of the Deployer task run. Clicking displays the payload of the Release. This is the actual list of applications (and the number of applications) to deploy that is specified in the bill of materials in the Release definition.

From the Release dashboard, you can go to the Path-to-Production view for a running or completed release. This view shows the bill of materials for the release. You can easily see the artifacts and snapshots that are in compliance with the Release manifest. Go to Path-to-Production View on page 305 for more information.

From the Environment Inventory

The environment inventory shows what is deployed in an environment at any time. Go to Environment Inventory on page 213 for more information.

From the Change History

The Change Tracking feature tracks the changes between every state of non-runtime objects, such as environments and resources. ElectricFlow records a *Change History* of the historical states of environment objects, including environment tiers and resources, and records changes between them. Go to Change Tracking on page 730 for more information.

Postp

Postp, a postprocessor included with ElectricFlow, is used for data collection and for report generation. For visibility in a Release run, you can get live streaming data from the log files using the standard output settings. Using postp commands in pipeline tasks (defined by application processes, procedures, or workflows), you can make this data appear in the Job Details page in ElectricFlow.



This example shows:

- During the Compile step, there were 955 compiles.
- In the "Command tests" step, 2982 tests were run.
- For the "Server tests (junit)" step, 14410 tests were run with 4 errors and 1 warning. Clicking on the links in this step provides details and evidence about how the job ran and what actions to take next.

Clicking on the log button for any of these steps will display the log file details.

For more information about using postp:

- Go to Defining the Stage on page 246 and Viewing and Editing the Pipeline Definition on page 251 to define, view, and edit pipeline tasks in the pipeline definition.
- Go to Postprocessors: Collecting Data for Reports on page 867 and Postp extension on page 979 to set up data collection, generate reports, and display data from the log files.

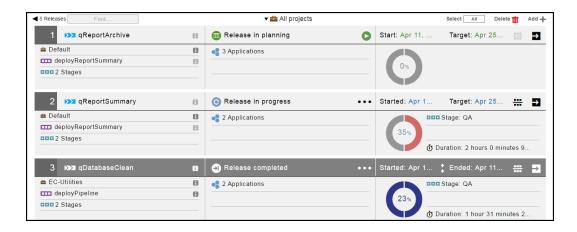
Example: Visibility in a Release Pipeline from the Release Dashboard

For an In-Progress Release

You can view the progress of the pipeline during the Release run starting in the Release Dashboard:

1. In the Release dashboard, clicking on the status of an in-progress or completed Release opens the Pipeline Run view for the Release pipeline.

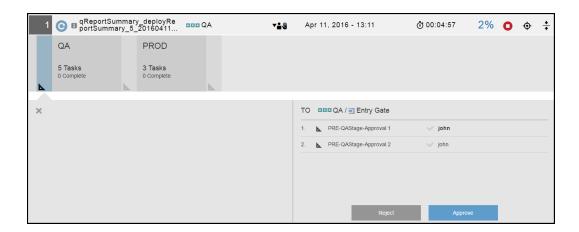
Example: Clicking **Release in progress** for the second pipeline shows the progress of a currently running Release.



2. When human intervention is needed at a gate (indicated by the hand icon), clicking **Approval waiting for john** displays the Stage Summary for an entry gate.

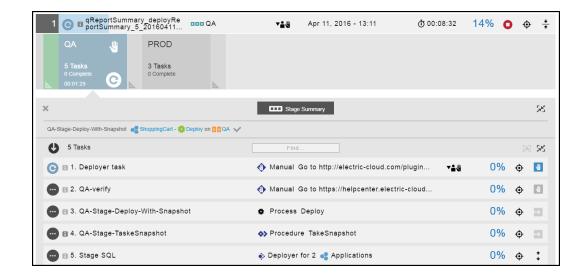


The Stage Summary for the enabled entry gate shows the required approvals.



After the approval rules or conditions are approved, the pipeline continues.

3. When a pipeline stage is running, it is enabled. Clicking in it displays the Stage Summary, which shows the tasks in the stage and their progress.



- a. In the Stage Summary, the top section shows information about the tasks deployed in the pipeline stage, which you can show or hide. You can also display user-generated information such as:
 - Links to web sites for reports and other information about the deployed artifacts
 - Property information about the pipeline stage and the objects in it



b. The next section shows the tasks in the pipeline stage. If the pipeline is in progress, click the down arrow in the row to view details for the manual task:



c. After completing the manual task, click **Complete** as the outcome, add a comment, and click **OK** in the Manual dialog box.



To view details about a completed manual task, click

to view the details. The Manual dialog



box opens. When the manual task has required input, click in the Parameters row to view the manual task details. The Manual dialog box now shows the **action** and **evidence** fields. The **action** field describes what was done to approve the manual task, and the **evidence** field describes what was done to approve it.

- 4. When human intervention is needed at the exit gate (indicated by the hand icon), an exit gate is enabled. Clicking in the exit gate displays the approvals required to complete the pipeline stage.
- 5. If the entry gate to the next stage has approval conditions before the tasks in it can start, this entry gate is enabled. Clicking in the entry gate displays the approvals required to start the tasks in a pipeline stage.
- 6. The sequence continues as the pipeline progresses.
- 7. Click View previous Pipeline Runs to view the previous five runs.



Clicking

displays the pipeline, where you can click in a gate or stage to view more details.



Clicking

hides this pipeline.



To view details about a manual task, click

to view the details of the manual task. The Manual dialog



box opens. When the manual task has required input, click in the Parameters row to view the manual task details. The Manual dialog box now shows the **action** and **evidence** fields. The **action** field describes the what was done to approve the manual task, and the **evidence** field describes what was done to approve it.



To view details about the other pipeline tasks, click

to view the Job Details

For a Completed Release

Go to Running and Completing Releases on page 309 for an example of running and completing a Release.

Release Definition

The Release process is defined in the Release definition, which consists of the underlying pipeline, the applications to deploy (the payload), where the applications will be deployed (the environments), and any inputs (parameters) to the application (referred to as configurations). The Release definition also includes notes that can be entered by team members, task assignees, and approvers as well as planned start and end dates for releases.

This section has examples of Release definitions and shows how to define a Release.

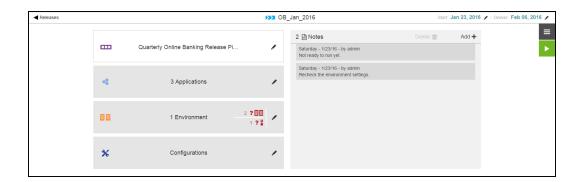
Example: Completed Release Definition

This release definition shows:

- The Release called "Q1 2016 Release OB" is based on the pipeline called "Quarterly Online Banking Release Pipeline."
- The **Applications** field shows that there are three applications to be deployed in this release. Clicking in the **Applications** field displays the list of applications to be deployed, which is passed to the Deployer task in the pipeline by the Release.

When artifact staging is enabled for an application, ElectricFlow tries to retrieve artifacts that will be deployed before the application run starts. This ensures that the artifacts are downloaded and available on the target (either the artifact cache [the default] or a directory specified in the component definition). This minimizes the time to retrieve artifacts and reduce downtime during the application deployment.

- The Environments field shows that the applications will be deployed in three environments. Clicking in the Environments field displays the environments where the applications in each stage will be deployed.
- Clicking in the **Configurations** field displays the inputs to the applications (parameters) to be entered at runtime.
- Two notes have already been added to the Release definition.



Example: Traditional Multiple Application Release

This example shows how to define a multiple application Release.

- 1. From the Home page, click Releases.
- 2. In the Release Dashboard, click Add+ to add a Release.

3. In the New dialog box, click Create new release.

Note: You can include hyperlinks as part of an object description for any ElectricFlow object.



4. Enter the name of the new Release and select a project to which the Release belongs.

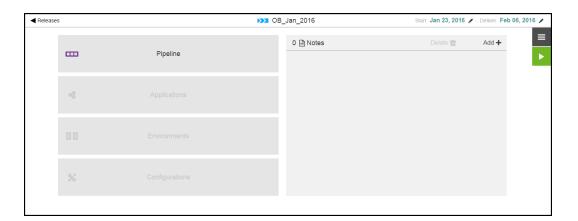


5. Click OK.

The Release Dashboard displays the new Release at the top.

6. In the Release Dashboard, click the new Release name.

The Release Definition page is displayed:



7. Click Pipeline.

A list of pipelines opens.



It shows the list of available pipelines. The default is to show the pipelines for all projects.

If you want to see only the pipelines for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more pipelines by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no pipelines in the selected projects.

8. Select a pipeline.

The pipeline called "Quarterly Online Banking Release Pipeline" is selected.

After the pipeline is selected, the Next button is enabled because the applications in the pipeline have

2[]

required inputs (parameters), as indicated by . This means that there are two required inputs to deploy the applications through the pipeline. If any input is a credential parameter, you can enter the path to the credential, browse to it, select the Parameter Credential or Credential binding, or select a user-defined credential that is attached to the project associated with the pipeline.

Enter the required inputs, and click **OK**.

If the applications did not have any required inputs, the **OK** button would be enabled instead.

More about this pipeline:

This pipeline has a Deployer task, as shown in the task list for the pipeline.

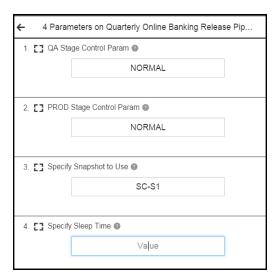
Remember that the Deployer task is specified in the task list of a pipeline stage but gets the actual list of applications to deploy from the Release payload (the bill of materials of the Release definition). At runtime, this list is passed to the pipeline through the Release object, and the pipeline deploys the applications in sequential order.



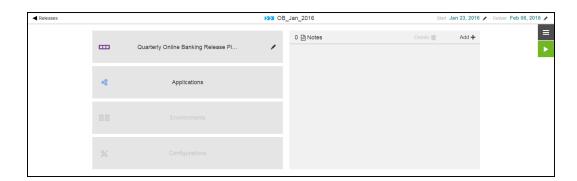
9. If the applications in the selected pipeline have inputs, a dialog box showing the inputs to the pipeline will be displayed.

If any input is a credential parameter, you can enter the path to the credential, browse to it, select the Parameter Credential or Credential binding, or select a user-defined credential that is attached to the project associated with the pipeline.

If any of required parameters do not have values, enter the appropriate values and click **OK**.



The Application field is now enabled in the Release Definition page.

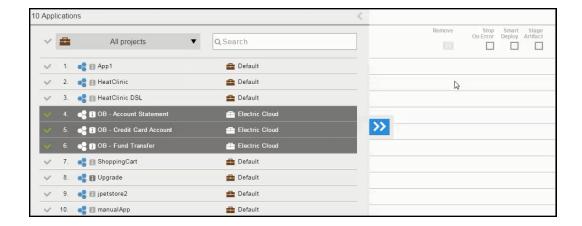


10. Click Applications.

The dialog box where you specify the applications to be deployed through the pipeline opens.



11. Select the applications that you want to deploy in the release, and then click the button to move them to the right side of the dialog box, which is the list of applications to be deployed through the pipeline.

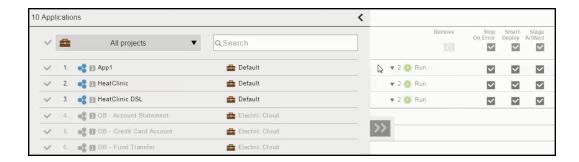


These applications are selected:





Clicking button moves the selected applications to the right side and displays the number of required parameters at run time and the application process to be deployed.





It shows the list of available applications. The default is to show the applications for all projects.

If you want to see only the applications for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more applications by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no applications in the selected projects.

12. Click < to set how you want the applications to be deployed during the Release run.

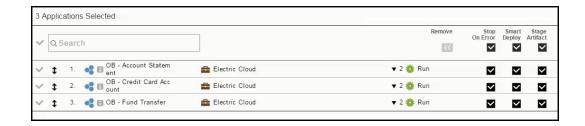
For each application, you can select these settings:

• **Stop on Error**: Select this error handling setting to abort the Release run when the Deployer task fails. The default is **Stop on Error**. If the application fails, the release pipeline aborts.

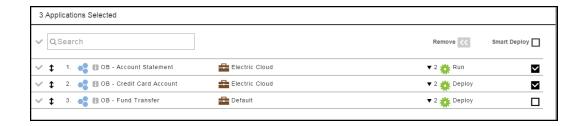
When this setting is not selected (continue on error, which is similar to the **Continue running** option in a pipeline stage or gate), the Deployer continues to the next application with a known error if the current application fails. This allows the release to continue running when an application run has errors and fails. For example, in a release with multiple applications, some applications can fail and not be deployed with other that were successfully deployed as delivered in the release.

- Smart Deploy: Select this setting if you want the Release to deploy the application only with
 artifacts that have not been deployed to a resource or with selected versions of the artifact
 have not been deployed to new resources since a previous run.
- Stage Artifacts: Select this setting if you want the Release to retrieve artifacts that will be deployed in an application run before the deployment starts.

In this example, all the settings are selected for the selected applications. The Release will stop if any of the applications fails.



Note: The default is that ElectricFlow uses smart deploy to deploy the applications in the pipeline. If you want to use smart deploy for all the applications, you do not need to change the settings.





13. If the applications have snapshots, click the down arrow next to each application, and click **OK**.

Using a snapshot as part of the Release definition is a best practice. Using snapshots, you can lock the exact versions of the artifacts to be deployed, which provides consistency and predictability. But you do not have to select a specific snapshot for an application. If a snapshot is not selected for an application here, the exact artifact versions to use would be based on how that application is modeled. That is, if the application model uses the *latest* artifact version for a component, the latest version would be used every time the Release runs when a snapshot is not selected.

When the Release runs, this list of applications (the payload) is passed to the Deployer task in the pipeline, and the applications are deployed through the pipeline sequentially. You can change the order in which the applications will be deployed by reordering the application input in the payload.

14. Click Environments to select the where the applications in each stage will be deployed.

The "Environment selection" dialog box opens.

Each cell in the table represents an environment where application in a stage will be deployed.

About the filters in the "Environment selection" dialog box:



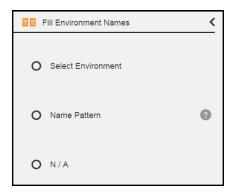
When you mouse over the filter icons, the tool tip text appears, from left to right:

- · View only cells with missing environments.
- · View only cells with missing tier maps.
- View only cells marked "N/A".
- · View only empty cells.
- 15. Select one or more cells in the table one of these ways:
 - For one cell (one application-stage combination)—You can select one cell by clicking in it. This allows you to select the environment for a specific application-stage combination.
 - For multiple cells (all the applications in a stage)—If the environment in a stage is the same for all the applications, you can select all the cells in a column (stage) by clicking in the header cell with the stage name.

Note: After one or more cells is selected, the Fill button in the upper right corner is enabled.

16. Click the > button in the upper left corner of the dialog box..

The "Fill Environment Names" section opens on the left side of the dialog box.

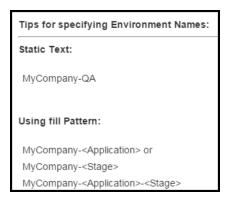


- 17. Select an environment using one of these options:
 - **Select Environment**–Click the down arrow to select an environment or environment template.

Clicking **Select Environment** opens the list of available environments in the **Environments** tab. These are static environments that are authored before deploying the application in the stage.

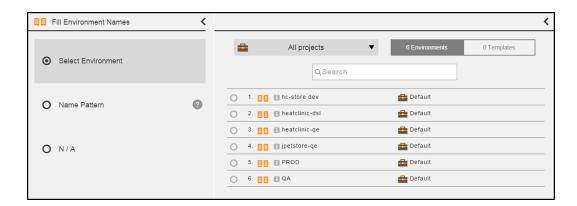
Clicking on the **Templates** tab opens a list of environment templates. These are dynamic environments with provisioned cloud resources that can be spun up during the application deployment. For more information about dynamic environments, see <u>Modeling Dynamic Environments</u> on page 143.

• Name Pattern–Click ? for tips on how to specify environment names. The tips describe how to specify environment names using static text or a fill pattern.



• N/A-Select this to not run the application in the stage.

In this example, the "hc-store dev" environment is selected from the Environments list.



It shows the list of available environments. The default is to show the environments for all projects.

If you want to see only the environments for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more environments by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no environments in the selected projects.

- 18. Click **OK** to close the "Fill Environment Names" section.
- 19. Click **OK** to save the settings.
- 20. Click **Configurations** to enter the inputs (parameters) for the applications in a stage.

If any of the required parameters have not been entered on one of the cells for an application in a stage,



you will see this:

You can enter the value for it now or at runtime.

21. Click OK.

The next step is to run the Release. Go to Running and Completing Releases on page 309 to learn more about running and completing the Release.

Example: Continuous Delivery Style Releases for One Application

This example shows that when you define a Release, you do not have to enter all the details. However, any configuration issues have to be resolved before the Release run.

After a Release is added, the Release definition opens:



Click in the **Pipeline** field. A dialog box opens where you can select a pipeline:



Note: When you click the button (number 1 above), a description of the deployPipeline appears. The icon on the right (number 2 above) means that two required parameters are attached to the pipeline.

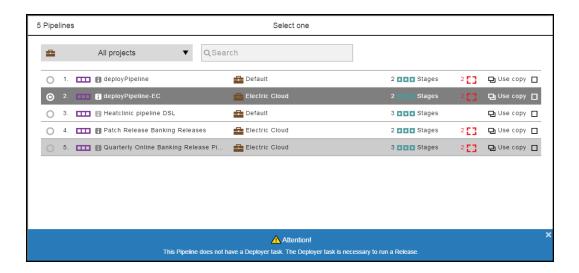
It shows the list of available pipelines. The default is to show the pipelines for all projects.

If you want to see only the pipelines for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more pipelines by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no pipelines in the selected projects.

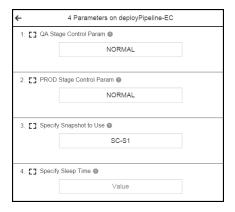
Select the deployPipeline application, and click Next.



If the pipeline does not have a Deployer task, a message appears reminding you that you need to define a Deployer task to run the Release. Close the message to continue, and click **Next**.

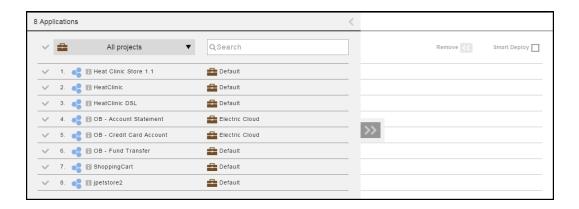
If the pipeline has a Deployer task, the message does not appear so you need to only click Next.

In the example, the pipeline has four required parameters. Some of them have default values that you can keep or change. Others do not, and you have to enter a value for them.



After you enter the values, click **OK**.

Now click in the **Applications** field. The dialog box to select one or more applications opens:



It shows the list of available applications. The default is to show the applications for all projects.

If you want to see only the applications for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more applications by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no applications in the selected projects.



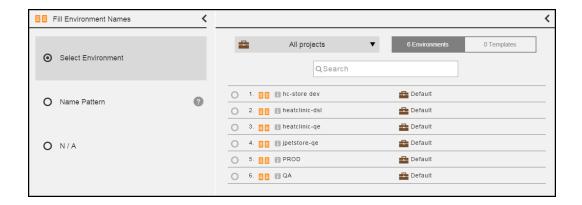
Select one or more applications, click the

button to add them to the pipeline, click **OK**, and click **<**:



Now click in the **Environments** field. The dialog box to select an environment where the applications in each stage will be deployed opens:





It shows the list of available environments. The default is to show the environments for all projects.

If you want to see only the environments for a specific project, click the down arrow in the **All projects** field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

You can also search for one or more environments by entering the search criteria in the **Search** field next to the **All projects** field. If there are no matches, a message appears stating that there are no environments in the selected projects.

About the filters in the "Environment selection" dialog box:



When you mouse over the filter icons, the tool tip text appears, from left to right:

- · View only cells with missing environments.
- · View only cells with missing tier maps.
- · View only cells marked "N/A".
- · View only empty cells.

For detailed instructions about how to select the environments in which the applications in each stage will be deployed, go here.

In this example, the environments are selected as follows:

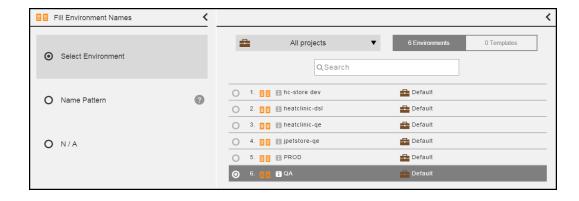
ShoppingCart application in the QA stage:

In the ShoppingCart row, select the cell in the QA column and double-click in it. Notice that the **Fill** button is now enabled.

A panel opens to the left where you can select the environment where the ShoppingCart application will be deployed.

Select the environment to which the ShoppingCart application is deployed: QA.

Click **OK** to save your selection.

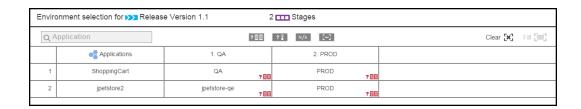


• jpetstore2 application in the QA stage:

In the jpetstore2 row, select the cell in the QA column and click in it. Notice that the **Fill** button is now enabled.

A panel opens to the left where you can select the environment where the jpetstore2 application will be deployed.

Select the environment to which the jpetstore2 application is deployed: jpetstore-qe.



Click OK.

• All the applications in the PROD stage:

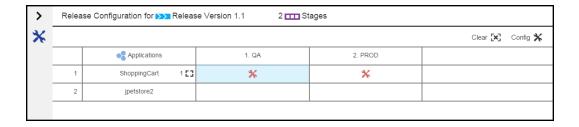
Click in the header cell for the PROD stage.

A panel opens to the left where you can select the environment where the applications will be deployed.

Select the environment to which the applications will be deployed: **PROD**.

Click OK.

Now click in the **Configurations** field. The dialog box to enter the required parameters for the pipeline to run opens:



If any of the required parameters have not been entered in a cell (an application to be deployed in a stage) with



in it, you can click in the cell to enter the inputs (parameters) for the application in a specific stage. You can enter the value for it now or at runtime. Then click **OK** to save your entries.

The Release definition now looks like this:



The next step is to run the Release. Go to Running and Completing Releases on page 309 to learn more about running and completing the Release.

Release Dashboard

The Release Dashboard, also referred to as the Release List, provides a way for the release team, task assignees, and approvers to quickly see the status of all the planned, active, or completed releases at any time. They can quickly see the current status, specific tasks and who is assigned to them, and the start and end dates for the stages in one place on one platform.

Accessing the Release Dashboard

To open the Release Dashboard from the Home page, do one of the following actions:

• Click Releases.



Click the Main Menu button () and select Releases.

Release Dashboard

The Release Dashboard shows all the Releases that you have permission to view.



By default, the releases for all projects are displayed.

If you want to see only the objects for a specific project, click the down arrow in the All projects field to select one or more projects one of these ways:

- Click on the name of one or more projects.
- Enter the search criteria in the Search field. The projects that match the criteria appear in the list.

If there are no matches, a message appears stating that there are no resource templates in the selected projects.

The Release Dashboard has this information:

Each row in the Release Dashboard contains this information:		
Release name	Example:	
	Q1 2016 Trading System	
Release status	The status can be: Release in planning Release in progress Release completed Examples:	
	Release in planning	
Start and target (end) dates for the Release	Example: Start: 10/30/15	
	The default time frame between the start and target (end) dates is two weeks.	
	The color of the dates indicates the Release status.	
	Green: The Release has not started yet, and the start and end dates are within the planned time frame.	
	 Medium blue: The Release is in progress, and the dates are the actual start and end dates. 	
	Dark blue: The Release is completed, and the dates are the actual start and end dates. Dad: The date is past the planned and date.	
	Red: The date is past the planned end date.	

Project to which the pipeline belongs, Example: name of the pipeline underlying the software release, and the number of Utilities 8 stages in this pipeline u Quarterly Online Banking Release Pipeine 8 ■■■ 3 Stages Number of applications in the pipeline Example: and any notifications for actions to be taken before the pipeline progresses 2 Applications Approval waiting for quincy This pipeline requires human intervention to progress. The user named quincy needs to approve something at an entry or exit gate of a stage before the pipeline progresses.

Visual indicator showing how far the Release has progressed with the percentage completed, the current stage, and how long the Release has been running.

The visual indicator is split into sections equal to the number of stages in the pipeline. If pipeline has three stages, it is split into three segments.

This example shows that the Release is running without errors, is at the INT stage, and has been running for over 11 days. The duration is the total time that the Release has been running before the Release is ended.



If the Release is still in the planning phase, the percentage completed is 0%, and the segments are all gray.



If the pipeline is running with errors, the part of the segment representing the current completed stage is red.



If the pipeline run is completed, the visual indicator is blue and the percentage completed is 100%.



Buttons and Icons in the Release Dashboard

Entering part or all of a pipeline name Find... in the Find field on the top of the dashboard starts the search for Releases matching the search criteria. Selecting All or None and then clicking Select All the **Delete** button removes the selected Delete 🎹 pipelines from the list. Clicking the Add+ button adds a Add + Release. Clicking this button opens to Path-to----Production view, which displays the applications being deployed in the Release and the environments where they are deployed. The details in this view include the application versions and the environments that are not in compliance with the bill of materials throughout the Release. Go to Path-to-Production View on page 305for more information.

Clicking this button opens the Release Definition dialog box. Go to Release Definition on page 283. Clicking this button starts a Release run. Select one of these actions: Clicking this button allows you to end, ••• re-run, or abort the Release pipeline. Note: To go to the Pipeline Run • End Release–Clicking this ends the Release. You cannot end the pipeline that is currently running. If you want to end it, you must wait for this run to complete or abort it in page, click the button in the the Pipeline Run page. When the pipeline is aborted, upper left corner of the Release tasks that are currently running will finish, and no new Dashboard, and select Pipelines > tasks will be started. Pipeline Runs. For more Re-run Pipeline

—Clicking this restarts the pipeline. information about the Pipelines • Abort Pipeline-Clicking this aborts the pipeline. If the Runs page, go to here. pipeline is aborted, tasks that are currently running will For more information about running finish and no new tasks will be started. and ending (completing) the Release, go to Running and Completing Releases on page 309.

Path-to-Production View

The *Path-to-Production view* shows the applications deployed in the Release, which versions are deployed, and the environments to which they are deployed. The details in this view include the application versions, snapshots, and the environments that are not in compliance with the *bill of materials* throughout the Release.

To go to the Path-to-Production view:

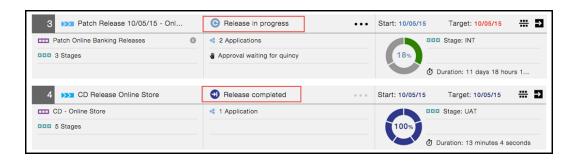
- 1. Go to the Release Dashboard.
- 2. Find a Release that is in progress or completed.



The button should be enabled.

Example:

Select one of these Releases:



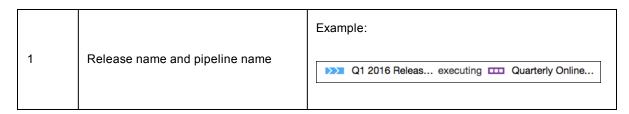
#

3. Click the button to open to Path-to-Production view.

This is a Path-to-Production view showing the bill of materials.



This information is in the bill of materials:



2 Stag	Stage	Each column shows the details for each stage in all the applications. Stage names are defined from the associated
		Pipeline. The stages in this example are <i>UAT</i> , <i>STG</i> , and <i>PROD</i> .
3	Summary	The bill of materials has three applications. In this example, the STG and PROD stages each
		have a compliance issue indicated by
4 Application		Each row shows the details for each application across all the environments.
		The application names are defined from the associated pipeline.
	Application	The applications in this example are:
		OB - Account Statements OB - Credit Card Statements
		OB - Fund Transfer
5	UAT stage	In this example, the UAT stage has no compliance issues.
		The versions of the applications deployed in the Release match what is in the bill of materials.
6	STG and PROD stages	In this example, the STG and PROD stages have compliance issues.
		The "OB - Credit Card Accounts" application version 2.0 is supposed to be deployed in the STG and PROD stages.
		However, version 1 of the application is deployed instead.

To get more details about compliance issues:



 Mousing over shows a message explaining the noncompliance. Use this information to troubleshoot and resolve the noncompliance.



• Clicking shows the environment to which the application is deployed, the deployment date, and information about the deployed components and artifacts.



You can take these actions in the bill of materials:



• Clicking the button shows less detail.



• Clicking the button shows all the details.



 Clicking the button starts opens to Search field, where you can enter the part of an application name to find specific applications, which is especially helpful when you have a long list.

Release Summary

You can add a Release summary using ectool. The summary can be found in the ec_releaseSummary property sheet that exists on the runtime associated with the Release. This property can be created using the following steps:

1. Create a procedure and the underlying steps with commands like the following:

```
ectool createProcedure 'Default' 'summaryproc'
ectool createStep 'Default' ''summaryproc'' 'step1' --command 'ectool setProperty /
myReleaseRuntime/ec_releaseSummary/summary1 --value "property value in plain tex
t>" '
ectool createStep 'Default' ''summaryproc'' 'step2' --command 'ectool setProperty /
myReleaseRuntime/ec_releaseSummary/summary2 --value "<html><a href=\"http://electric-c-cloud.com/\">Electric Cloud</a></html>"'
```

2. Call this procedure from a task in the stage.

During a Release run, after the task in a pipeline stage has completed, you can view the summary property in the Release Dashboard.

Running and Completing Releases

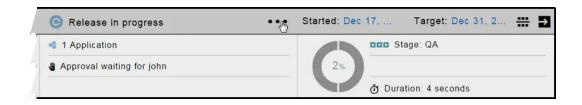
After a Release is defined, you can run it multiple times on the journey to production. At the end of the software release life cycle, when the Release process has been successfully completed and the software is ready for production, you can end the Release. Once the Release has ended, it cannot be run again.

Running the Release

This example shows how to run a Release called Q1_TimeBaseRelease.

You can start the Release run when the Release status is "Release in planning" and the **Run** button () is enabled.

The Release status changes to "Release in progress" and the **Run** button () has been replaced by the button. Clicking on it shows a pop-up window with more actions. The Start and Target dates are in blue because the Release is in progress.



Approval waiting for john

When is displayed, human intervention is required for the pipeline to progress to the first stage.

Approval waiting for john

Clicking on

causes the Pipeline Run view for the deployPipeline to be displayed.



The pipeline can progress after the user named john approves the entry to the QA stage of the pipeline.

As the pipeline progresses, you can click in the QA stage to view the pipeline stage summary:



The pipeline continues to progress to the PROD stage after the required approvers respond at the exit gate. Progress continues until this Release run is completed.

After the Release is completed, you can start a new Release run by clicking the button and selecting **Re-run Pipeline**. You can re-run the Release as many as times as you want during the software release life cycle.

Aborting a Release Run

There are two ways to abort a Release run.

• During the first Release run, you can abort it from the Release Dashboard.



Clicking the button and selecting **Abort Pipeline** will abort the pipeline. The tasks that are currently running will finish, and no new tasks will be started. This message appears:



This is the Release status after the Release run is aborted.



• During subsequent Release runs, you can abort the current Release run from the Pipeline Run view.



0

Clicking the Abort () button will abort the pipeline. The tasks that are currently running will finish, and no new tasks will be started. This message appears:



After you click **OK**, the Pipeline Run view now displays this status:

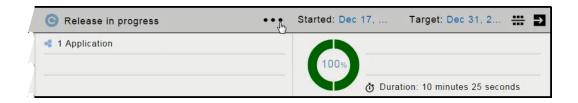


Completing the Release

When you complete the Release, no other actions can take place on it, and this action cannot be undone (reversed). The Release cannot be run again.

You cannot end the pipeline that is currently running. If you want to end it, you must wait for this run to complete or abort it in the Pipeline Run view. For more information, go to Aborting a Release Run on page 310.

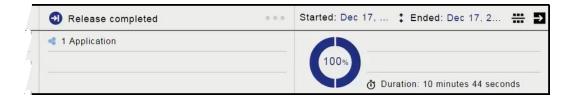
If the Release run is complete, you can complete the Release from the Release Dashboard.



Clicking the button and selecting **End Release** ends the Release. This message appears:



After you click **OK**, the Pipeline Run view now displays this status:



The Release status is now:

- The icon next to "Release completed" indicates that the Release has been completed (ended).
- The "percentage complete" icon on the right is split into two sections because the pipeline has two stages. It shows that the Release run is 100% complete.
- The Target date has been replaced by the Ended date.
- The Start and Ended dates are in dark blue because the Release is completed.
- The Duration is the total time for the Release.

Security

ElectricFlow provides security by assigning roles and privileges to specific users and groups on:

- System objects including applications, environments, projects, jobs, and schedules.
- · Actions performed on deployment models.

ElectricFlow uses access control, project-level security, and credentials and impersonation to enforce roles and privileges when executing deployment steps.

For an example of how to define roles and privileges when a specific user is allowed to deploy an application only to a specific environment, go to Use Case: Attaching Credentials in Deployment Automation on page 324.

Access Control

ElectricFlow uses access control to provide security for all system objects. This mechanism controls how users and groups use the system. Users must log in to view information or to perform operations (actions). After users log in, their system access is limited based on:

- The user name
- · The groups to which the user belongs
- The permissions specified for various ElectricFlow objects.

Go to Access Control for more information about how ElectricFlow enforces access control and for security examples using access control.

For instructions on how to set up access control, go to these topics:

- · Access Control on page 496 GUI page
- Access Control defining entries on page 497
- Privileges create new or edit existing privileges on page 498

Project-Level Security

Multiple Project support is available on applications, pipelines, Releases, environments, Master Components, resources, and environment templates as well as platform objects (for example, artifacts, procedures, jobs, schedules, and workflows) These objects as well as the objects belonging to them can be in any project within ElectricFlow.

This significantly improves object management at scale by allowing:

 ACL inheritance—All objects in a project inherit the access control settings from the project, providing better security for all the objects. They can be managed in their own projects and will inherit the ACLs setup at the project level. This significantly simplifies permissions management. Logical grouping—This allows users to better manage Deploy and Release objects under various
projects that are logically mapped by users, roles, geography, department, and so on, resulting in easier
maintenance.

For more information about specific objects, go to the appropriate sections in the *ElectricFlow User Guide*.

For an example showing how to select a project for an application in the GUI, go to Example: Modeling and Deploying Applications on page 88. You can also use API commands to do this:

- Use the createApplication API command to create a new application for a specific project.
- Use the createProcess command to create an application or component process for a specific project.

See the ElectricFlow API guide on the ElectricFlow documentation page for more information about using these API commands.

Also go to Example: Modeling and Deploying Applications on page 88 to see how to author and deploy an application in the GUI.

For an example showing how to select a project for a Release in the GUI, go to Release Definition on page 283. You can also use the createRelease API command to define a Release for a specific project.

Credentials and User Impersonation

What is a Credential? on page 315

Defining a Credential on page 316

Why Use Credentials? on page 316

Credential Access Control on page 316

Using Credentials for Impersonation at the Platform Level on page 317

Accessing Credentials From a Step at the Platform Level on page 319

Credential References on page 320

Best Practices for Retrieving Credentials on page 321

What is a Credential?

In ElectricFlow, a "credential" is an object that stores a user name and password for later use. Two credential types are available, *stored* or *dynamic*:

- stored credentials These credentials are given a name and stored in encrypted form in the database.
 Each Project has a list of stored credentials it owns. These credentials are managed from the Project Details page.
- **dynamic** credentials These credentials are captured when a job is created. Dynamic credentials are stored on the server temporarily until the job completes and then discarded.

Defining a Credential

After a credential is created, no one can view the password for the credential's account. This means one person can define a credential and enter the password, and other people can use the credential (and its account) without needing to know the password.

To create a new credential in the automation platform UI:

- 1. Click the Projects tab.
- 2. Select a project (first column) to see the Project Details page.
- 3. Select the Credentials tab and then click Create Credential.

Click the **Help** link on the New Credential page if you need more information about what to enter in the fields.

Why Use Credentials?

Use ElectricFlow credentials for:

- User impersonation
- · Saving passwords for use inside steps

When a step needs to run as a particular user, ElectricFlow can retrieve the user name and password from a stored credential. The credential is passed to the agent over an encrypted channel so the agent can authenticate itself to the operating system and set up a security context where the step runs with the user permissions in the credential.

In some cases, a step needs to enter credentials directly to an application.

- The credential may be a fixed application credential that needs to be presented every time the step runs.
 In this case, a stored credential provides a place to put this information without needing to embed the password in a script.
- Other times, the password is not known in advance. In this case, a dynamic credential created when a
 job is launched can prompt a user for information and then pass the credential to the step in a secure
 manner.

Credential Access Control

Because ElectricFlow stores passwords in a recoverable manner, it is important that credentials are accessible only under carefully controlled circumstances. To protect credentials, ElectricFlow uses two mechanisms: access control lists (ACL) and *attached* credentials.

To use a credential for any purpose, a user who attempts to reference the credential must have execute permission on the credential object—this allows you to explicitly control which users are allowed to use a credential anywhere in the system.

In addition to the user-based access control check, ElectricFlow requires a credential be explicitly *attached* to the object that is going to use it. After an object has a credential attached, object modifications are restricted to users who have both execute permission on the credential and modify permission on that object. This safeguard helps prevent credentials from accidentally being used by unauthorized users.

Attach a credential to an object in the automation platform UI one of these ways:

• Set the impersonation credential for a step, procedure, project, or schedule.

You can also set the impersonation credential for an application process, component process, or a process step.

Impersonation credentials are inherited, so attaching a credential to a container object like a procedure or a project implicitly attaches it to every object inside the container, such as all steps in a procedure.

Explicitly attach the credential to a step or schedule using the UI or the attachCredential API.

To access credentials from a step, the credential must be attached directly to the step, or the credential must be passed as a parameter to the containing procedure and have the parameter attached instead.

Using Credentials for Impersonation at the Platform Level

This information addresses the question, "When a job step is running on a resource, under which account is it running and how can I control that?" By default, a job step runs under the account used by the ElectricFlow agent that runs the step, which was chosen when the agent was installed. This approach works well in environments where it makes sense to run all jobs under a single user such as a "build" user. You install all agents to run as the desired user, and you do not have to worry about anything else—every step of every job runs as that user.

However, in other environments you may prefer to run different jobs, or even different job steps, under different accounts.

For example:

- If independent groups are sharing an ElectricFlow system, you may want each group to use a different account for its jobs.
- You may need to run jobs as particular users to access ClearCase views for those users.
- There may be certain steps that require special privileges.

For example, as part of your builds you may need to run a step that generates a certificate using privileged corporate information; that step must run under a special high-privileged account, but you want the remainder of the steps in the build to use a less-privileged account.

ElectricFlow allows you to select accounts on a per-job or per-job-step basis. This mechanism is called *impersonation*, where the ElectricFlow agent impersonates a particular user for the duration of a job step, and this *impersonation* is implemented using credentials. The use of credentials creates special security challenges.

Note: It is important to ensure privileged accounts can be used only for the purposes you intend and cannot be "hijacked" for other purposes. The ElectricFlow access control mechanism contains special facilities to ensure proper credential use.

Setting the Impersonation Credential in the Automation Platform

Attaching Impersonation Credentials to Steps, Procedures, and Projects

You can attach impersonation credentials to procedure steps, procedures, and projects before executing a job step.

ElectricFlow searches for a credential in the following order and uses the first credential it finds:

- 1. Procedure step
- 2. Procedure to which the step belongs
- 3. Project to which the project belongs

If the step is in a nested procedure, and no credential is found on the step, its procedure, or its project, ElectricFlow checks the calling step (and its procedure and project), its caller, and so on until it has worked up through the topmost procedure in the job. If no credential is found, the job step runs under the ElectricFlow agent account on its resource.

This approach makes it easy to manage your account usage. For example, if you want all jobs in one project to use a particular account, define a credential in that project for the account, and then attach the credential to the project. If you want all jobs in a project to use a particular account, except for one step that should use a different account, create two credentials in the project. Attach the first credential to the project and attach the second credential to the particular step.

Attaching Impersonation Credentials to Schedules

You can also attach a credential to a schedule. The credential in the schedule is used for steps where no other credential is available.

The schedule's credential receives the lowest priority.

Attaching Impersonation Credentials to Jobs

You can specify a credential when you launch a job manually one of these ways:

- Provide the name of an existing credential
- Enter an account and password.

If you choose the second approach, the account and password are stored only in ElectricFlow for the duration of the job.

For either way, the credential you specified when you launched the job is used as a last resort for steps only with no other credential.

Attaching a New Impersonation Credential

To attach a new impersonation credential for a procedure, step, or schedule:

- 1. Click the Projects tab.
- 2. Select a project.
- 3. Go to the Project Details page for the project you selected.

4. Do one of the following:

a. For a procedure:

- i. Click Create Procedure to go to the New Procedure page.
- ii. On this page, select the Impersonation Credential drop-down box to select a credential to use for impersonation.

b. For a step:

- i. Select a procedure to go to the Procedure Details page.
- ii. Select a step name to go to the Edit Step page.
- iii. On the Edit Step page, scroll down to the Impersonation section.

c. For a schedule:

- i. Select the Schedules subtab.
- ii. Select a schedule to go to the Edit Schedule page.
- iii. Scroll down to Impersonation Credential section.

For more information in the platform GUI, each of these web pages contains a **Help** link in the upper-right corner.

Accessing Credentials From a Step at the Platform Level

The ElectricFlow server supports an operation called <code>getFullCredential</code> that can return a password when called from inside a step. The ectool interface for this operation is:

```
getFullCredential <credential name>
    [--value <password|user>]
```

Without the --value option, the response is the same as that returned by getCredential, with the addition of a "password" element.

With --value, the simple text password or user name value is returned on *stdout* so it can be used directly without XML parsing.

The getFullCredential API is allowed only inside a running step, and is allowed to use credentials that were explicitly attached to the step only.

Attaching a Credential to a Procedure Step, Procedure, or Schedule

A credential must be explicitly attached to the object using it so the server can perform an access control (ACL) check at definition time and limit the visibility of the password. To support accessing credentials other than the one being used for impersonation, steps, procedures, and schedules contain a list of credential names.

In the platform GUI, to attach a credential (other than an impersonation credential) to a step, procedure, or schedule:

1. Go to an "Edit" page.

For example, to create a credential for a step, a step must already exist.

- 2. Select a step that needs a credential. The Edit Step page opens.
- 3. Scroll to the lower section of the page and click Attach Credential.

If you are an ectool command-line user, use the attachCredential API command. For more information on the command-line interface, go to the *ectool* Help topics within the online Help system.

Passing Credentials as Parameters

Sometimes you may not know which credential you need to use when you define a procedure step. In this situation, you can leave the credential choice up to the caller of the procedure.

The following example describes the procedure for passing credentials as parameters:

- 1. Create a project named InnerProj.
- 2. Within that project create a procedure named InnerProc with a parameter 'cred' of the type credential.
- 3. Within that procedure create a step named InnerStep with:
 - o Command: ectool getFullCredential cred
 - Attached Parameter Credential: cred
- 4. Create a project named OuterProj with two credentials: cred1 and cred2.
- 5. Within that project create a procedure named OuterProc with two parameters, cred1 and cred2, of the type credential.
- 6. Within that procedure create two subprocedure steps that both call InnerStep from the InnerProj project:
 - OuterStep1 with Attached Parameter Credential cred1, and cred1 in the Parameters section
 - OuterStep2 with Attached Parameter Credential cred2, and cred2 in the Parameters section

If you are an ectool command-line user, you can use the <code>createFormalParameter</code> and <code>attachParameter</code> API commands. For more information on the command-line interface, go to the <code>ectool</code> Help topics within the online Help system.

Credential References

For schedules, the value for any actual parameter passed to a credential-type formal parameter is interpreted as a reference to a previously attached credential on the schedule.

For procedure steps, the value for any actual parameter passed to a credential type formal parameter is interpreted as a reference to a previously attached credential or credential parameter on the calling step.

When a job step is created, all credential references are resolved and the content is copied into transient credentials on the job step. Credential parameter references are resolved by looking for a credential in the calling job step's transient credential map—looking for a credential with the name specified as the actual parameter value.

Credential references in the API can be specified as a relative or an absolute credential path.

- A relative path is similar to "rootCred" and is interpreted as referring to a credential in the current project.
- An absolute path includes a project name like "/projects/MyProject/credentials/rootCred" and can refer to a credential in any project.

Anywhere a credential reference is accepted (for example, setting an impersonation credential or attaching a credential to a step), either form can be used.

Best Practices for Retrieving Credentials

The ElectricFlow impersonation mechanism provides powerful mechanisms for using different accounts in job steps, and it can be used to handle a variety of challenges. However, if misused, impersonation can open up dangerous security loopholes. This section discusses these risks and how to avoid them.

Use of credentials and impersonation tend to fall into two classes.

First Class

In the first class, your goal is to open up access to accounts, you want to make some accounts generally available, and you trust large groups of users (such as everyone who can access a particular project) to use the accounts appropriately. This usage type is simplest and relatively safe.

- 1. Define an ElectricFlow group containing all trusted users.
- 2. Set permissions on the Project so only trusted users can access the project.
- 3. Define credentials for accounts in the Project and grant execute permission to everyone in the group.

After completing this task, everyone in the group can use the accounts for arbitrary purposes. No one outside the group will be able to access the project or the credentials.

Second Class

In the second class, you want to provide tightly-controlled access to a highly privileged function such as generating a certificate or accessing sensitive information. This kind of usage requires careful thought to avoid allowing unintended access to the credential. The best way to implement a solution for this scenario is to restrict access to the credential and use it on individual steps only after careful analysis of issues such as the following:

Does the step's command use properties?

If so, someone could potentially change the property value to change the step behavior.

For example, suppose the step's command is defined like this:

```
echo $[/myProject/x]
```

Anyone with write access to property x can hijack the step by placing the value "nothing; myCommand" in the property.

After property substitution, the final command will be:

```
echo nothing; myCommand
```

that will cause "myCommand" to execute in the step.

Either restrict access to the property sheet or avoid substitutions.

Does the step execute code or commands that can be modified?

In most jobs that execute build and test sequences, the job extracts various files from a source code control system and some of those files contain code executed during the build. In this case, anyone with access to the source code control system can affect the executed code.

For example, suppose a step runs a Make or Ant command using a configuration file extracted from the source code control system. Anyone with access to the source code control system can modify the configuration file to introduce new commands executed during the Make or Ant step. In this situation, it is virtually impossible to control what happens during the step, so you should never attach a credential for a privileged account to such a step.

The higher you attach a credential, the greater the risk of uncontrolled access to the credential's account.

For example, if you attach a credential to a project, there is a good chance anyone with write access to the project, or even the ability to execute the project's procedures, can hijack the credential using one of the loopholes described above. If you attach a credential to a step that invokes a subprocedure, you effectively give anyone with write access to the subprocedure complete access to the credential's account. If you care about access to an account, you should use its credential in the narrowest possible fashion, such as attaching it to a single job step.

Avoid passing passwords on the command line

When using credentials, avoid passing the password on the command line because [on most platforms] those passwords will be visible to all users logged into the system. The secure way to pass a credential to an application is to use a secured file in the file system or to pipe the value into the application via stdin.

Credentials and Impersonation in Application and Component Processes

You apply credentials and impersonation to control who can deploy applications and where the applications are deployed.

- You can attach one or more credentials to component or application process steps.
- You can attach only one impersonation credential to an application process, component process, or a process step.
- When you attach an impersonation credential in ElectricFlow, it specifies the user who can deploy the application and the environment in which the application is deployed.
- When you attach an impersonation credential to a platform-level object (such as a procedure), it
 specifies the account (user) that can run the job or job step. If you want to specify another condition, you
 have to attach another credential to the object.

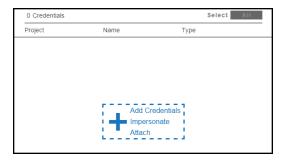
IMPORTANT:

When you impersonate a credential, make sure that the impersonated user has the absolute path to the bin directories in the \$PATH environment.

If you define a process step with a command, you must enter the absolute path in the **Post Processor** and **Shell** fields in the "Define Step" dialog box.

1. Click in the Add Credentials field.

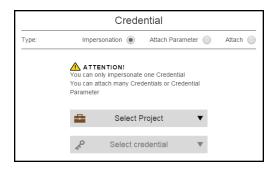
Example:



The Component Process Step / Add dialog box opens.

2. To impersonate one credential:

Example:

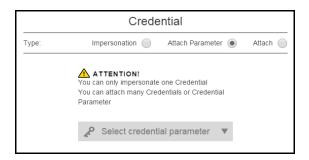


- a. Select Impersonate in the Type field.
- b. Click the **Select Project** field to select a project. You can select a credential from a project that is different than the one containing the application.
- c. Click the **Select Credential** field to open a drop-down list of credentials for the process step.
- d. Select a credential.
- e. Click OK.

The **Credentials** dialog box now shows the one credential for impersonation.

3. To attach one or more credential parameters to the process step:

Example:

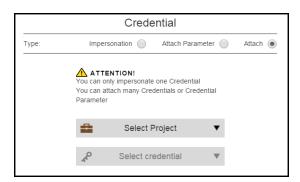


- a. Select Attach Parameter in the Type field.
- b. Click the **Select Credential Parameter** field to open a drop-down list of credentials for the process step.
- c. Select a credential.
- d. Click OK.

The Credentials dialog box now shows the attached credentials

4. To attach one or more credentials to the process step:

Example:



- a. Select Attach in the Type field.
- b. Click the **Select Project** field to select a project. You can select a credential from a project that is different than the one containing the application.
- c. Click the Select Credential field to open a drop-down list of credentials for the process step.
- d. Select a credential.
- e. Click OK.

The Credentials dialog box now shows the attached credentials.

Use Case: Attaching Credentials in Deployment Automation

When modeling an application, you can attach credentials these ways:

- For credentials, attach one or more credentials to component or application process steps.
- For credentials for impersonation, attach only one impersonation credential to these objects:
 - · Component process
 - · Component process step
 - · Application process
 - · Application process step

When you deploy the application, ElectricFlow applies the credentials based on the user permissions and deploys the application in one or more environments.

Example

This example describes how to define roles and privileges when an application is deployed to more than one environment and specific users are limited to specific environments.

You can attach impersonation credentials to an application in the GUI. The application has these credentials:

- Development (dev)
- Quality Engineering (qe)
- Production (prod)

Users have these privileges:

- User A is allowed to deploy the application to build a MyQSQL database in any environment and has admin privileges.
- User B is allowed to only deploy the application in the quality (qa) environments and is not trusted in the development (dev) and production(prod) environments.

The following user permissions determine what users are allowed to do in ElectricFlow. You can configure these credentials only from the GUI.

• You configure User A's profile in the automation platform and give User A higher-order privileges than other users. User A has the following credentials, including a credential for impersonation:

For each environment in ElectricFlow, set a property using a reference such as $property = 1000 \, \text{myEnvironment/dbConfigName}$ and define a unique value, which can be passed as a credential to a process or process step.

- In the development (dev) environment, set dbConfigName = dbUser_dev.
- In the quality engineering (qe) environment, set dbConfigName = dbUser qa.
- In the production (prod) environment, set dbConfigName = dbUser prod.
- When you configure User B's profile in the automation platform, User B is only given the credentials to
 deploy in the quality (qa) environment. You do not need to set properties to be passed as credentials
 when the application is deployed.

Attaching the same credentials from the automation platform is more complicated. Instead of setting only one credential for User A and one for User B, you need to create three unique credentials for the environments in

addition to credentials for the various user and environment combinations, such as User A and the development environment.

Credentials in Pipelines

Credentials are passed to pipelines through applications in pipeline stages. If an application has required inputs (parameters) and any of the inputs is a credential parameter, you can enter the path to the credential, browse to it, select a credential parameter (**Parameter Credential**) or a credential binding to a pipeline task (**Credential binding**), or select a user-defined credential that is attached to the project associated with the pipeline.

A credential must be explicitly attached to the object using it so the server can perform an access control list (ACL) check at the definition time and limit the visibility of the password. To support accessing credentials at the pipeline task level, these tasks need to have the appropriate credentials attached. This is done implicitly by the UI when you are binding credentials to the pipeline tasks. If you are using ectool, use the attachCredential API command to perform the same operation.

Credentials in Releases

In the Release definition, credentials appear in these objects:

- Pipelines: How the credential is applied to the Release depends on the selected pipeline.
 - If the pipeline has required inputs (parameters) and any of the inputs is a credential parameter, you can enter the path to the credential, browse to it, select a credential parameter (**Parameter Credential**) or a credential binding to a pipeline task (**Credential binding**), or select a user-defined credential that is attached to the project associated with the pipeline.
- Applications: How the credential is applied to the Release depends on the applications in the selected pipeline.

If any input to an application is a credential parameter, you can enter the path to the credential, browse to it, select a credential parameter (**Parameter Credential**) or a credential binding to a pipeline task (**Credential binding**), or select a user-defined credential that is attached to the project associated with the pipeline.

See the steps to select a pipeline and to select credential inputs for an application in the pipeline in the Release Definition on page 283for examples of these operations.

Change Tracking

Change Tracking is designed to track every change between every state of non-runtime ElectricFlow objects and to allow you to revert to any previous state of these objects. ElectricFlow tracks the changes to tracked objects including applications, procedures, workflows, workspaces, resources, and project-owned components such as library components and records a *Change History* of the historical states of the system and the changes between them. The change history has a copy of every state in which every non-runtime object has been, the object's current state, and indexing records for searching through the database. The tracked objects are not affected when ElectricFlow executes an object that is usually created or modified at execution time, such as a workflow or process.

Change tracking allows you to do the following:

- When debugging a failed job or looking for more information about a component in an ElectricFlow, see the Change History for the changes relevant to that object.
- When searching for specific change history records, filter the records by time frame, change type, entity type, or developer.
- Revert changes to an object or to an objects and its children.
- When you want to determine the differences between objects, export them at various levels in the object hierarchy.

In ElectricFlow, you can use Change Tracking in all aspects of the ElectricFlow end-to-end solution:

- In build-test automation, track the Change History of artifacts, resources, and ACLs.
- In deployment automation, track the Change History of components, environment tiers, and process steps (application or component). You also use Change Tracking with snapshots to make it easier to deploy reliable and repeatable software for Continuous Delivery.
- In Releases and Pipelines, track the Change History of stages and tasks.

Performance Consequences of Change Tracking

Change Tracking affects how ElectricFlow runs. The impact can be from low (minimal changes to the times to update the database) to high (significantly longer times to update the database).

ElectricFlow needs to store additional information in the database about the tracked objects in the Change History. In addition to storing normal information about ElectricFlow operations, ElectricFlow stores a separate archival copy of every state of every tracked object as well as the indexing information used to retrieve the Change History of tracked objects and the objects they own. This increase in database usage has performance consequences.

For most of ElectricFlow operations, the performance consequences of Change Tracking are insignificant, typically a few percent slower in one of these situations:

- None of the affected objects are tracked (job records and similar runtime-only objects are not tracked).
- Only a small number of tracked objects are changed by operations.

However, for operations that affect a very large number of tracked objects in a single operation (such as importing, cloning, or deleting a large project for which change tracking is enabled), the performance

consequences will be significant. The performance consequences are similar when enabling or disabling Change Tracking for a large project, because both operations require the server to write to the database the records for every tracked object owned by the project.

For these Change Tracking-intensive operations, the system performance will typically be significantly slower because several times as much data is being read and written between the ElectricFlow server and the database. In our testing, we have seen a wide range of degrees of performance decrease from 1.25x to 10x slower when Change Tracking is enabled on a large project, compared to a project on which Change Tracking is disabled. The decrease in performance is typically from 2x to 6x.

Use Case 1: Importing and Cloning Large Projects with Change Tracking Enabled

ElectricFlow does the following to increase the performance of Change Tracking-intensive operations in a large project with this configuration:

- The project contains over 20,000 audited objects.
- · Change Tracking is enabled.

To increase the performance when a large project is imported or cloned, ElectricFlow automatically decreases the amount of Change History indexing information that it saves, reducing the level of detail for these operations in the Change History. While this can improve the system performance, it can also make it harder when you want to revert an object to a specific state or to find information in the Change History while troubleshooting or debugging an issue. If you want suppress this behavior, use

--reduced DetailChangeHistory false in the import or clone $\ensuremath{\mathsf{API}}\xspace$ command.

By default, the objects listed in the table below and the tracked objects that they own are always tracked when Change Tracking is enabled (except when Change Tracking is globally disabled).

During an upgrade from ElectricFlow 5.0.x to ElectricFlow 6.0, when the ElectricFlow server starts, the server writes to the database a record for every tracked object and the objects that it owns. The objects listed in the table get new audit records. If any of the objects owns many tracked objects (such as properties), the startup process will be slow.

Objects (Entities)	Path
server	/server
system objects	/systemObjects/*
users	/users/*
groups	/group/*
plugin	/plugins/*
email configuration	/emailConfigs/*

Objects (Entities)	Path
workspaces	/workspaces/*
zones	/zones/*
gateways	/gateways/*
agents	/agents/*
resources	/resources/*
resource pools	/resourcesPools/*
repositories	/repositories/*
directories	/directories/*

Use Case 2: Change Tracking of Non-Project-Owned Objects

In this context, non-project-owned objects include all objects not contained within a project. Such objects include resources, artifacts, users, groups, zones, and other objects as well as the property sheets, properties, access control lists (ACLs), and ACL entries (ACEs) for them. In releases earlier than ElectricFlow 6.2, Change Tracking was enabled by default for non-project-owned objects, and there was no way to disable Change Tracking for them. When users upgraded from earlier versions of ElectricFlow to ElectricFlow 6.0.x, the upgrade process initialized and created Change Tracking-related data for them in the database. The performance impact was usually low, except when there was a very large number (in the several hundred thousands) of non-project-owned objects, such as many properties under the /server directory; in this situation, the performance of the first startup was significantly slower. The best solution to this situation is to relocate most of the projects under a project and disable Change Tracking for that project, which may not be a practical short-term solution.

Starting in ElectricFlow 6.2, Change Tracking can be disabled for all non-project-owned objects. (Change Tracking for these objects is enabled by default.) Using this feature is not recommended as a long-term solution; instead we recommend moving most of the objects under a project that has Change Tracking disabled. (Repeatedly disabling and enabling Change Tracking for non-project-owned project is expensive and time-consuming.)

In ElectricFlow 6.2, when there are many non-project-owned objects, the best way to disable Change Tracking for all non-project-owned objects is to change the database.propertiesconfiguration file to disable Change Tracking for these objects before starting the server running a pre-ElectricFlow 6.2 release. In the database.propertiesfile, make sure that it has this line:

COMMANDER DB NON PROJECT AUDITING ENABLED=false

Then install ElectricFlow 6.2 and start the server, running the system this way until most of the non-project-owned objects are relocated to a non-tracked project, and re-enable Change Tracking for non-project-owned objects, which requires restarting the server.

Estimating Database Growth

Here is one way to estimate the rate at which the database grows for Change Tracking:

Most of the information needed to estimate the database growth rate can also be parsed from an exported XML file of the project (such as a full import), except for the version count.

The likely rate of database growth is based on these factors:

- · How many objects of different types will be tracked
- · How big they are
- How often their state changes

The rate of database consumption is then the sum over all different types of tracked objects of the product of these three factors.

Common Usage Patterns

For most customer usage patterns, the most common objects (the ones that dominate on first factor) are ACLs and their ACEs, properties, or all of these objects. Typically:

- ACLs and ACEs: ACLs are consistently small, and usually have only a few ACEs.
- Properties: While most of them are small, some properties can be very large. The average property size is large.

For most users, a good quick estimate can be based on properties, especially properties that are large or are modified often.

Estimating Process

If you expect to enable Change Tracking for all projects, determine the values for these factors and calculate the database growth rate:

Database growth rate = Factor 1 x Factor 2 x Factor 3

· Factor 1: How many objects of different types will be tracked

The following procedure is only for properties:

- Join the ec_property table in the database joined to the ec_property_sheet table by ec_ property.parent_sheet_id = ec_property_sheet.id.
- Search for properties where the ec_property_sheet.entity_type value is one of the entity types corresponding to a database table that has a matching ec_*_aud table.

3. Determine the number of tracked objects per property type.

For most usage patterns, the great majority of properties will belong to entity types like job and jobStep that are not change tracked and so have no corresponding ec_*_aud table.

If you are already using ElectricFlow 5.3 or later, these properties and properties belonging to projects for which Change Tracking is disabled are also recognizable by having ec_property_tracked set to false.

• Factor 2: How big they are (size)

The following procedure is only for properties.

For small properties:

- 1. Search for objects that are less than 450 bytes in the ec property.string column.
- 2. Add an overhead value of half a kilobyte per property for the other columns.
- 3. Determine the average size of the tracked objects.

For larger properties:

- 1. Search for the length of the value in the ec_clob table linked to by the ec_property.clob column.
- 2. Add an overhead value of one kilobyte per property for the other columns.
- 3. Determine the average size of the tracked objects.
- Factor 3: How often new copies of the objects are generated

The following procedure is only for properties:

- 1. Use one of these methods:
 - Determine ratio of how long ago the ec_property.created and ec_ property.modified datetimes are, which gives an estimate concentrating on only the most recent changes.

and/or

- Determine the value of the ec_property.version counter divided how long ago the ec_property.created datetime is, which gives a long-term average rate of changes to the property since it was created.
- 2. Determine the average update frequency.

For properties that are frequently updated always to be a numerical value, it is possible to suppress Change Tracking of numerical-value updates by using one of these commands:

When the counter flag is set for a property, Change Tracking does not track changes to the property if the only change was a change in the property value from one numeric value to another. All other forms of changes to the properties that have this flag set are tracked normally.

CAUTION: Reverting an object that owns a counter property to a previous state will revert the value of the counter property to its value at the previous time that a change to this property was tracked—typically this is when its value is initialized. This may not be the desired behavior, so you may need to manually set a counter property if it is reverted,.

Alternative Estimating Methods

Depending on your usage patterns, one of these methods may result in a more accurate estimate than the previous method of determining the average factor values and multiplying them

Split the set of properties of change tracked objects into different populations

For each object, perform the previous method of determining the average factor values and multiplying them per tracked object.

Add the values for each tracked objects together.

Example:

- If you have one set of properties (such as those in one project, or belonging to one user) containing very large properties that are never changed, and another set containing very small properties that change often, multiplying the average property size by the average rate of property update would result in an overestimate of the total database consumption.
- If you put the properties in separate projects and calculate the database growth rate for each project, the database growth rates are more accurate.

Other columns that may be of interest include ec_property.owner (who created the property) and ec_property.last_modified_by usernames (who last modified the property).

Best Practices for Change Tracking

Change Tracking allows you to troubleshoot and debug issues, find information about specific objects, and revert an object to a specific state. However, Change Tracking affects how the system performance. In some situations, the performance consequences of Change Tracking can be significant.

This section describes how to improve the system performance.

When to Enable or Disable Change Tracking for a Project

In some situations, enabling or disabling Change Tracking on a project causes significant performance consequences. The ElectricFlow system performance becomes slower as the server writes to the database the records for all the tracked objects in the project.

To determine when to use Change Tracking on a project, do the following:

1. Export the project using this command:

2. Evaluate the results.

- a. The size of the results file gives you a good estimate of how big the project is:
 - If the file size is in the hundreds of MB, the project is large.
 - If the file size is in the tens of MB or less, the project is small.
- b. (Optional) Write a script to parse the results and determine these statistics:
 - Total number of objects that can be tracked in the project.

Use this value to determine the project size (for example, small, medium, or large), which is used to determine how long it would take to enable, disable, copy, import, or delete a project on which Change Tracking is enabled.

 The total size (in XML format) of the entities in the project that can be tracked by Change Tracking.

This gives a rough estimate of how much the database usage will immediately grow when Change Tracking is enabled for the project.

The rate at which tracked entities are modified in the project.

Look at the most recent last-modification dates of entities owned by the project for a pattern where certain entities are modified on a regular basis.

If there are objects that appear to be being modified frequently, the script can also parse out their path and the user name they were last modified by.

Next Steps:

- If the project is small, and its contents are not frequently changed by an automated process, you can safely enable Change Tracking for it.
- If the project is large, and if you frequently export, import, or copy it, and if the performance of these operations is very important to you, enabling Change Tracking for it may be problematic (see Estimating Process on page 330 for a possible workaround).
- If significant parts of the contents of the project are frequently updated by an automated process, enabling Change Tracking for the project will consume significantly more database space, and tracking the changes performed by the automated process is unlikely to be of much value, so enabling Change Tracking for the project is inadvisable.

Splitting Tracked Objects Into Separate Projects

If you export the project, evaluate the results, and determine that enabling or disabling Change Tracking on the project will significantly affect the performance of the ElectricFlow system, one way to improve the system performance is to split the project into separate projects based how the objects are modified.

For example, if you parse the results and find that there are two types of entities in your current project:

Objects that are manually modified over time by real users.

These objects represent real changes in project and are modified as procedures, jobs, and applications are run.

Objects that are frequently updated at regular intervals by an automated process using metaprogramming.

These objects are automatically generated and updated on a regular basis. The project file should have a large list of modifications to tracked objects.

Then you should split the objects into separate projects.

- Put the objects that are manually modified by real users in one project and enable Change Tracking on this project.
- Put the objects that are updated by an automated process in a different project and disable Change Tracking on it.

Configuring Change Tracking

Change Tracking must be enabled when ElectricFlow starts for your system to track changes and record the Change History.

By default, Change Tracking is enabled for projects created in ElectricFlow 5.3 or later. It is disabled for projects created in ElectricFlow versions before ElectricFlow 5.3, but can be enabled manually.

Enabling Change Tracking Globally

When installing ElectricFlow:

1. Add this line to the database.properties file:

```
COMMANDER_DB_AUDITING_ENABLED=true
```

2. Restart the ElectricFlow server.

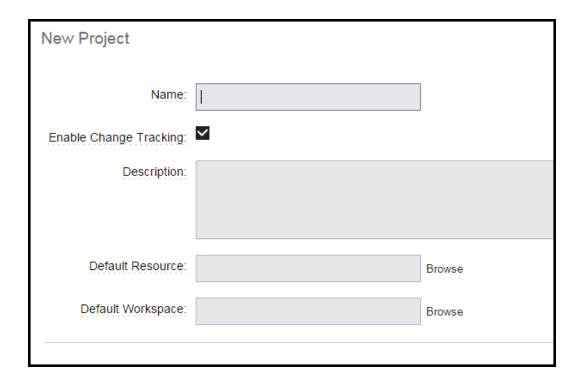
Enabling Change Tracking on a Per-Project Basis

IMPORTANT: Unless the allowEnablingDisablingChangeTrackingProjectLevel setting has been altered, only admin users are allowed to perform this task. The default setting is adminOnly. If the setting is changed to anyoneWithAccess, any user with access to the project can enable or disable Change Tracking for the project

You can enable Change Tracking one of the following ways:

In the ElectricFlow Platform UI

Change Tracking is enabled when the **Enable Change Tracking** check box is selected.



To disable Change Tracking, click the Enable Change Tracking check box to clear it. and click OK.

Using ectool

- Enter ectool modifyProject <projectName> --tracked false to disable Change Tracking.

Using ec-perl:

- Enter \$cmdr->modifyProject(<projectName>, {tracked => true}); to enable Change Tracking.
- Enter \$cmdr->modifyProject(<projectName>, {tracked => false}); to disable Change Tracking.

For properties that are frequently updated always to be a numerical value, it is possible to suppress Change Tracking of numerical-value updates by using one of these commands:

```
ectool:ectool modifyProperty cprojectName> -- path cpropertyPath> -- counter true

or
ec-perl:$cmdr-> modifyProperty (cprojectName>, {path => propertyPath>, counter => true});
```

When the counter flag is set for a property, Change Tracking does not track changes to the property if the only change was a change in the property value from one numeric value to another. All other forms of changes to the properties that have this flag set are tracked normally.

CAUTION: Reverting an object that owns a counter property to a previous state will revert the value of the counter property to its value at the previous time that a change to this property was tracked—typically this is

when its value is initialized. This may not be the desired behavior, so you may need to manually set a counter property if it is reverted,.

Upgrading to ElectricFlow 6.x

Change Tracking is enabled when you upgrade to ElectricFlow 6.x. This can significantly increase the time it takes to complete the upgrade.

If you want to upgrade with Change Tracking disabled, add this line to the database.properties file before starting the upgrade:

COMMANDER DB AUDITING ENABLED=false

Customizing the Change History Page

The performance of the Change Tracking feature is affected by number of records in the Change History as well as the number of entries being tracked. For example, a page showing 5000 entries may be slow to load and update and does not provide much useful information.

This ElectricFlow server uses the lowest of the following limits to determine the maximum amount of records to display in the Change History page:

· Maximum amount of records on the Change History page

To change the maximum number of records in the Change History page:

1. Set the CHANGE_TRACKING_HARD_MAX_RECORDS parameter in the wrapper.conf file to a new value.

The default value is 1000.

- 2. Restart the ElectricFlow server.
- · Maximum number of records retreived

Set the TrackingMax Records server setting to a new value not exceeding the CHANGE_TRACKING_HARD_MAX_RECORDS parameter in the wrapper.conf file.

To set TrackingMaxRecords, do one of the following:

Change the value in the UI. See the Server Settings page in the automation platform UI.

Use ectool to change the value. For example, enter the following command to limit the number of records retrieved to 100:

 $\verb| ectool setProperty / server/settings/changeTrackingMaxRecords --value 100| \\$

Searching the Change History

Follow these steps to start a search in the Change History.

You can start a Change History search from most pages in the ElectricFlow UI.

1. Click the **Search** button or click **Change History**.



The **Change History - Search** dialog box opens.

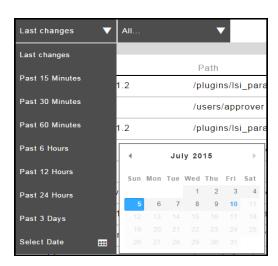


(F-	
	Time range field.
1	Click the down arrow to open the drop-down list of start times.
	The end time is the current time.
	Objects field.
2	Click the down arrow to open the drop-down list of objects to include in the search. You can select All or specific objects.
	By default, seven of the most commonly tracked objects are selected.
	Search criteria.
3	After you type, the system starts searching for objects based on the time range and objects that you selected.
	The search results are in the Change History.

- 2. To select a time range for the search:
 - a. Click the down arrow in the Time range field to open the drop-down list.
 - b. Select a time range.
 - c. If you want to use a time increment longer than three days, do the following:
 - a. Click Select Date.

The Date Picker opens.

Example:



b. Select a date.

The Date Picker closes and the date that you selected appears in the Time Increment field.

Example:



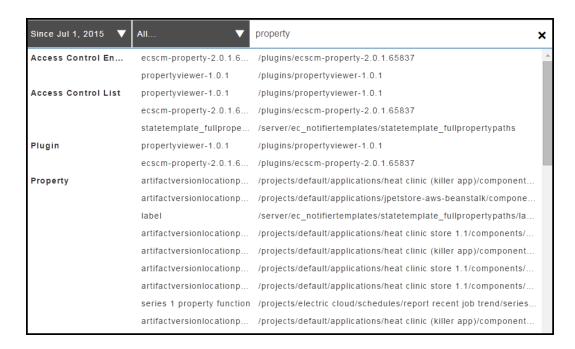
- 3. To select an objects for the search:
 - a. Click the down arrow in the **Object** field to open the drop-down list.
 - b. Select the objects for the search.
- 4. Enter the search criteria.

As you type, the system starts to search for objects that match your search criteria.

A list of objects matching your search criteria appears in the results section.

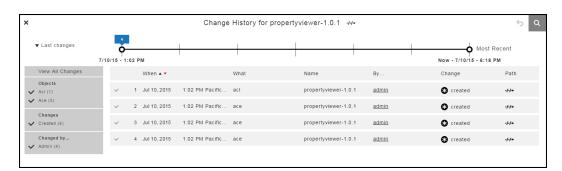
5. Select an object in the list.

Example:



The change history for the object that you selected appears.

Example:



Viewing the Change History

You can open and view the change history from the following objects in ElectricFlow:

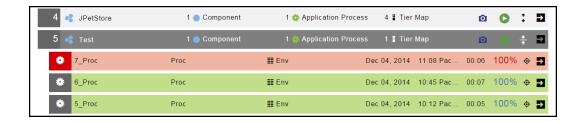
- · Applications List
- · Applications Visual Editor
- · Application Process Designer
- Artifacts
- Component

- Component Process Designer
- · Environments Designer
- Environment Tier
- Jobs
- Process Step (Application or Component)
- Projects
- Resources
- Workflows

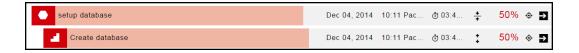
Viewing the Change History from the Applications List

You may want to view the Change History for these objects in an application:

· An application process that did not run successfully



· Component or component process in an application process that did not run successfully



• Resource that was not deployed successfully

Starting from the Home page:

- 1. Go to the Applications List.
- 2. Choose an application.
- 3. Click the View button.

Example:



A list of the application processes for the application appear.

- 4. To view the Change History of an application process:
 - a. Choose a process.
 - b. Click the Change History button.



The Change History for the application process opens.

- 5. To view the change history for an object in the application process,
 - a. Click the View Details button.

A list of objects in the application process (components and component processes) appears, and the breadcrumb changes to *Applications/View Run*.

- b. Choose an object.
- c. Click the Change History button to see the change history for the object.

Viewing the Change History From an Application or Environment

You may want to view the Change History for these objects:

- Applications
- Application processes
- · Component processes
- Environment
- Resources

Application Visual Editor

- 1. Go to the Applications List.
- 2. Select an application.

The Applications Visual Editor opens.

3. Click the Menu button.



4. Select Track Changes.

The Change History for the object opens.

The default time increment is Past 60 Minutes.

Example:



- 1. Go to the Applications List.
- 2. Select an application.

The Applications Visual Editor opens.

3. Select an application process.

The Application Process Visual Editor opens.

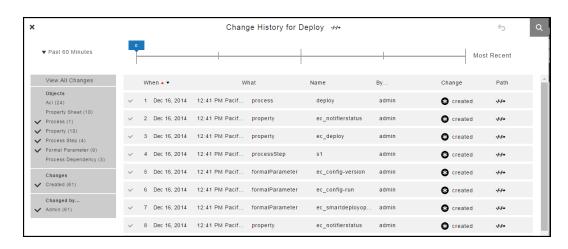
- 4. To view the Change History for the application process:
 - a. Click the Menu button.



b. Select Track Changes.

The Change History for the application process opens.

The default time increment is Past 60 Minutes.



- 5. To view the Change History for the application process step:
 - a. Choose a step.
 - b. Click the Edit button in the step to open the context menu.

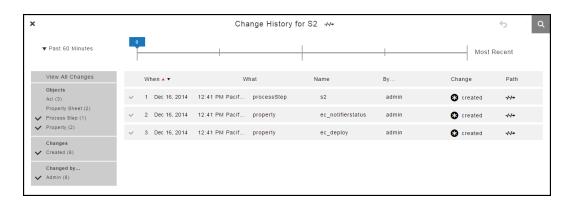


c. Select Track Changes.

The Change History for the application process step opens.

The default time increment is Past 60 Minutes.

Example:



Component Process Visual Editor

- 1. Go to the Applications List.
- 2. Select an application.

The Applications Visual Editor opens.

- 3. Choose a component.
- 4. Select a component process for the component.

The Component Process Visual Editor opens.

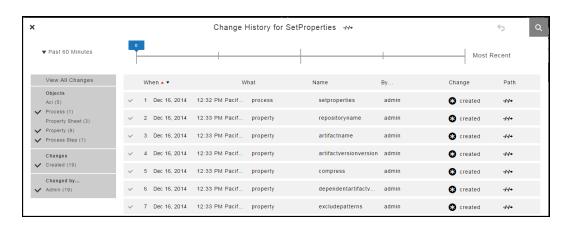
- 5. To view the Change History for the component process:
 - a. Click the Menu button.



b. Select Track Changes.

The Change History for the component process opens.

The default time increment is Past 60 Minutes.



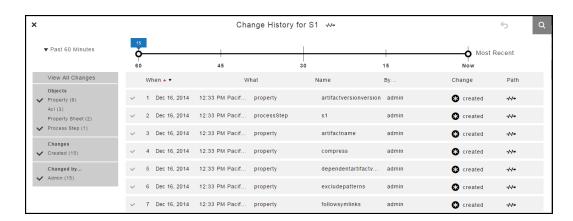
- 6. To view the Change History for the component process step:
 - a. Choose a step.
 - b. Click the Edit button in the step to open the context menu.



c. Select Track Changes.

The Change History for the component process step opens.

Example:



Environments Visual Editor

- 1. Go to the Environments List.
- 2. Select an environment.

The Environments Visual Editor opens.

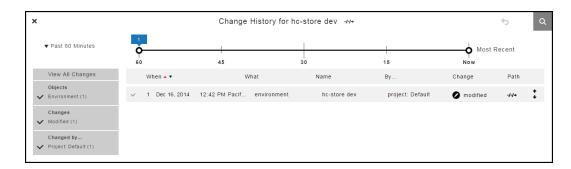
3. Click the Menu button.



4. Select Track Changes.

The Change History for the object opens.

Example:



Environment Tier

- 1. Go to the Environments Designer.
- 2. Choose an environment tier.
- 3. Click the Menu button.

Example:

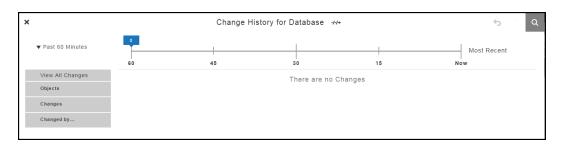


4. Select Track Changes.

The Change History for the object opens.

The default time increment is Past 60 Minutes.

Example:



Viewing the Change History for Artifacts, Jobs, Projects, and Workflows

When troubleshooting why a job failed, you can view the Change History for artifacts, jobs, projects, and workflows in the automation platform.

- Artifacts on page 348
- Jobs on page 349
- Projects on page 350
- Workflows on page 350

Artifacts

Starting from the Home page:

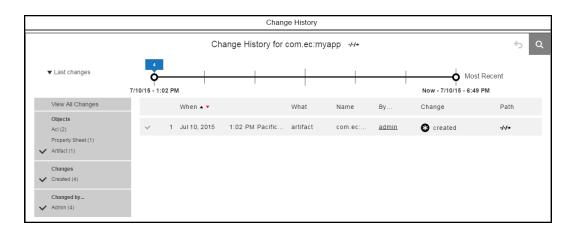
- 1. Go to the Artifacts tab.
- 2. Choose an artifact.
- 3. Click the Track Changes button.

Example:



The Change History for the selected artifact opens.

The default time increment is Past 60 Minutes.



Jobs

Starting from the Home page:

- 1. To go to the Job Details page, do one of the following:
 - Use the Jobs tab.
 - · Use the Jobs Quick View list.
- 2. If you use the Jobs tab, follow these steps:
 - a. Click the Jobs tab.

The Jobs page opens.

Example:



b. Click a job name to select a job.

The Job Details page opens.

Example:



3. If you use the Jobs Quick View list, click a job name to select a job.

The Job Details page opens.

- 4. Choose a job or job step.
- 5. Click Track Changes for the job or job step.

The Change History for the job or job step opens.

The default time increment is Past 60 Minutes.

Projects

Starting from the Home page:

- 1. Go to the Projects tab.
- 2. Choose a project.
- 3. Click the Track Changes button.

The Change History for the project opens.

Example:



Workflows

Starting from the Home page:

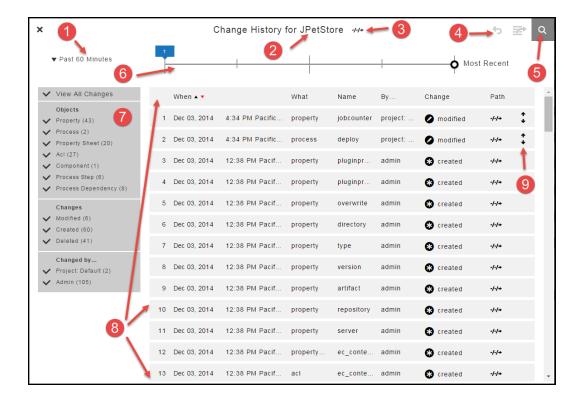
- 1. Go to the Workflows tab.
- 2. Choose a workflow.
- 3. Click Track Changes.

The Change History for the workflow opens.

Change History Page

How to get here: Click the **Search** button to open the "Change History - Search" form, and enter the search criteria.

The search results appear in the Change History page and include the following information:



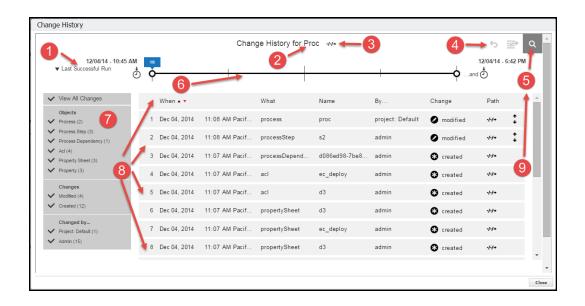
How to get here:

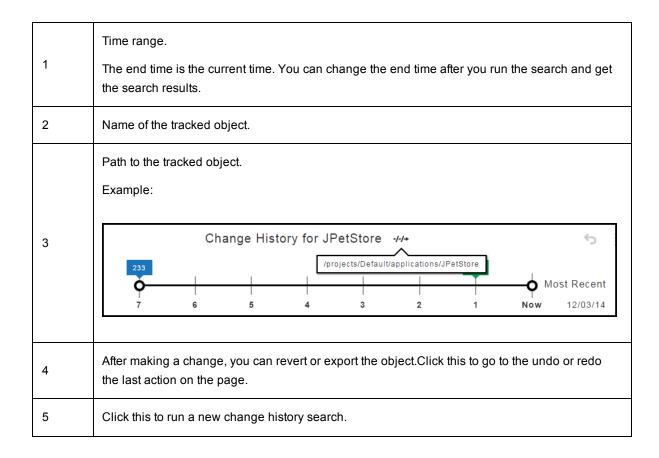
From the Home page in the automation platform, use one of the following methods to open and view the change history for jobs, projects, workflows, and artifacts:

From the Home page, use one of the following methods to open and view the change history for jobs, projects, workflows, and artifacts:

- Select a job that failed > Click the job name to go to the Job Details page > Click **Track Changes** in the upper right of the page.
- Click the Jobs tab > Select a job that failed to the Job Details page > Click **Track Changes** in the upper right of the page.

The search results appear in the Change History page for the job and include the following information:





	Time line.			
6	The start time is based on the time range that you selected.			
	The end time is the current time.			
	You can manually change the start and end times after you run the search and get the search results.			
	Filters for the change history.			
7	You can view all changes or view only selected changes.			
	The objects in the list are the objects in the change history search results.			
8	Change history for the selected object. When—the date and time that the object changed. What—The type of object. Name—The name of the object. By—The "user" that changed the object, which can be a project or a user. Change—The type of change. Path—Click the View Path button to see the path to the object.			
	When ▲ ▼ What Name By Change Path			
	1 Dec 03, 2014 4:34 PM Pacific property jobcounter project: omodified			
	2 Dec 03, 2014 4:34 PM Pacific process Before path: /projects/Default/applications/JPetStore/jobcounter After path: /projects/Default/applications/JPetStore/jobcounter			
	3 Dec 03, 2014 12:38 PM Pacif property pluginpr admin			
9	Click the View button to see more details about the object.			

Time Line

The time line for the change history is at the top of the Change History page.

The time line is automatically separated into divisions based on the number of minutes, hours, or days between the end time and the start time.

Default Settings

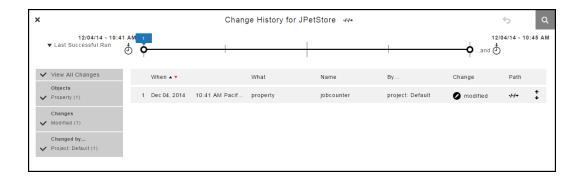
These are the default settings for the time line:

- The start time is the selected time range in the upper left of the page.
- The end time is **Most Recent** when the latest change to the object occurred.

- The default range is from the Last Successful Run to the Most Recent.
- The entire time line is displayed, and all the changes are in the list below the time line.

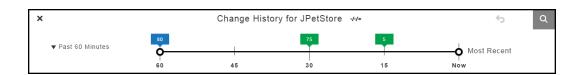
The time line on this page has the default range.

Note: The Last Successful Run range is available only after the first time that you run an application.

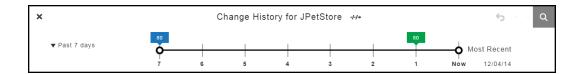


When the range is changed to **Past 60 Minutes**, the time line changes:

- The start time is 60 minutes from Now.
- The end time is when the **Most Recent** change occurred (**Now**).
- The time line has four divisions.



If the increment is Past 7 Days, the time line has seven one-day divisions.

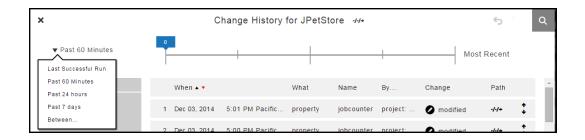


You can view the change history in different ways depending on the time range for the time line.

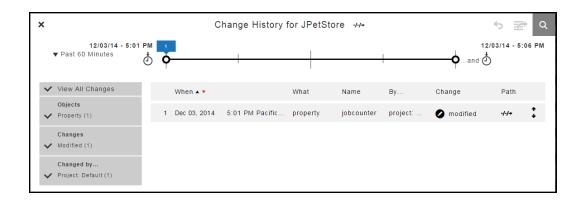
Selecting the Time Range

You can view the change history in different ways depending on the time range for the time line.

Click the down arrow next to the time range to select a different range for the change history.



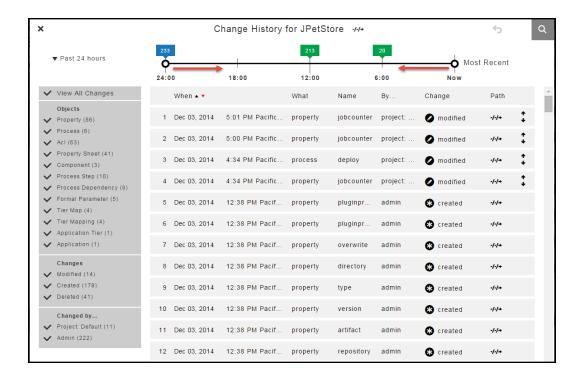
You can see the change history since the last successful run. Notice that the information on the left side shows a summary of the changes, which is a subset of the results that you got.



Moving the Start and End Times

You can manually move the start and end times on the time line.

When you move the start time to 18:00 and the end time to 6:00, notice that the list of objects in the change history changes.



When you move the start time to 18:00 and the end time to 6:00, notice that the list of objects in the change history changes.

Specifying the Time Range

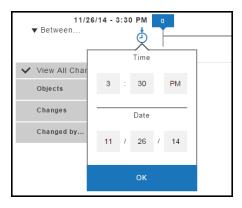
Example:



To manually set the times and dates for the start and end times:

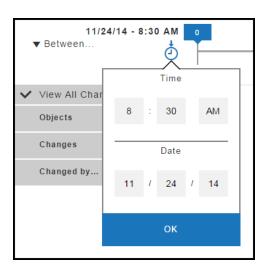
1. Select Between.

A drop down dialog box opens.



2. Select the time and date for the start of the time line.

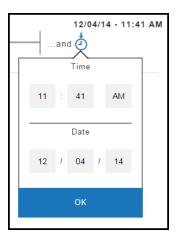
The default settings are 3:30 pm and eight days before the current date.



3. Click OK.

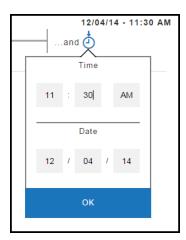
A drop down dialog box opens at the other end of the time line.

Example:



4. Select the time and date for the end of the time line.

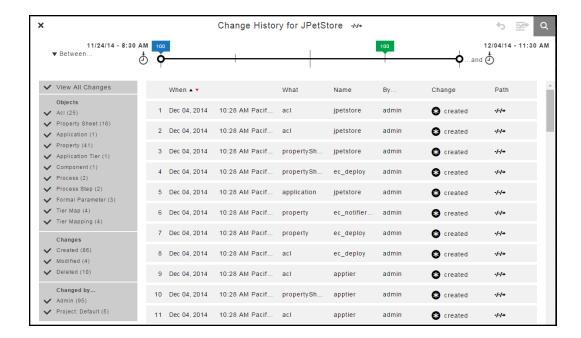
The defaults are 3:30 pm and the current date.



5. Click OK.

The time line changes to show only the changes from the start and end times and dates that you selected.

Example:



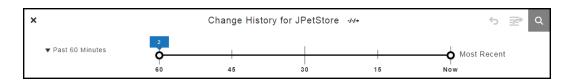
Number of Changes

The time range at the top of the change history shows the number of changes.

- There have been 233 changes in the last 24 hours.
- There have been 213 changes in the last 12 hours.
- There have been 20 changes in the last 6 hours.

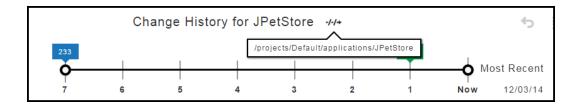


When you change the time range, the number of changes also changes. In the last 60 minutes, there have been only two changes.

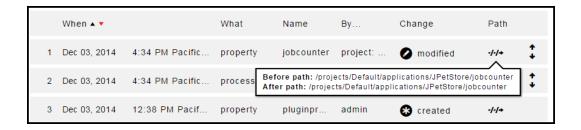


Paths to Objects

Click the View Path button next to the "Change History for JPetStore" title to see the path to the application.

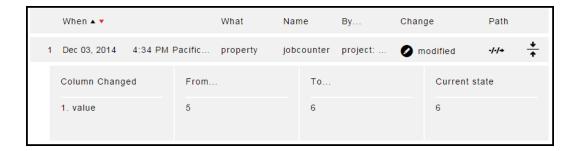


Click the View Path button to see the change in the path to the object before and after the change.



Detailed Object Changes

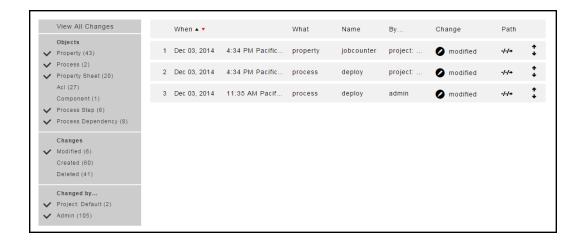
Click the View button to see the change in the property called jobcounter.



Filters

You can use filters to view changes to specific objects, the types of changes, and the users how made those changes.

Instead of selecting **View All Changes**, you can select specific objects, such as only properties, processes, property sheets, process steps, and process dependencies that have been modified by the Project:Default and Admin users.



When the list of filter criteria is long, not all of the criteria may appear in the filter list. To see all of the criteria, use the up or down arrows to see all the options.

This list does not show all of the users. Use the up and down arrows to see all four of the users.



Click the down arrow to see the other users.



Modifying What You See the Change History

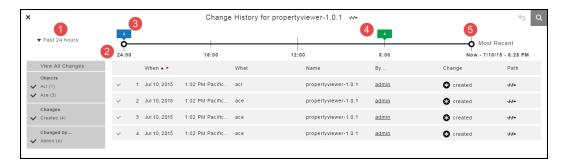
You can modify the information that appears in the Change History with these settings:

- Time line-Go to Change History Time Line on page 736.
- Filters-Go to Change History Filters on page 742.

Change History Time Line

The time line is at the top of the Change History page.

This example shows the Change History for a property called *propertyviewer-1.0.1*.



Time increment.

The default is Last Changes.

Click the down arrow to select another time increment.

2	The system automatically calculates the minutes, hours, and days since the last successful run. In the example, the last successful run occurred 24 hours ago. The time line is divided into four 6 hour subdivisions.
3	Total number of changes in the selected time increment.
4	The number of changes that occurred between 6 hours ago and now is 4. When you click the change number, the Change History is updated and shows only those changes.
5	Drag the start and end time markers to view specific changes.

Default Settings

The default time increment is **Last Change**.

The entire time line is displayed, and all the changes are in the list below the time line.

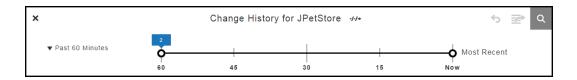
Number of Changes

The time line shows the number of changes throughout the time increment. In the following example:

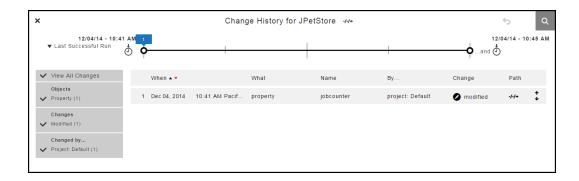
- There have been 233 changes in the last 24 hours.
- There have been 213 changes in the last 12 hours.
- There have been 20 changes in the last 6 hours.



When you change the time increment, there have been two changes in the previous 60 minutes.



Time Increment



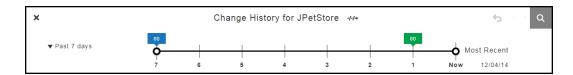
The system automatically determines how the time line is divided for the selected time increment.

When the range is changed to **Past 60 Minutes**, the time line changes:

- The start time is 60 minutes from Now.
- The end time is when the **Most Recent** change occurred (**Now**).
- The time line has four divisions.



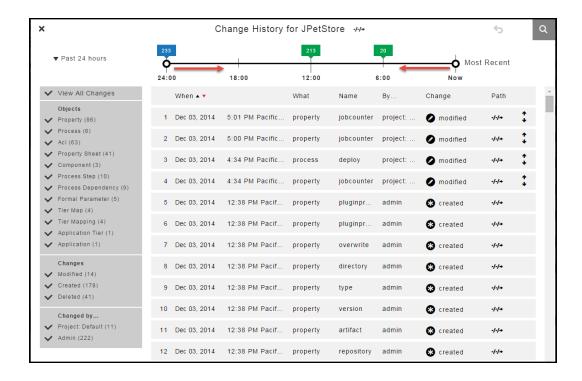
If the increment is $\bf Past~7~Days$, the time line has seven one-day divisions.



Moving the Start and End Times Manually

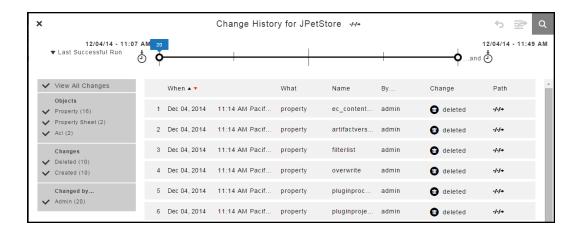
You can manually move the start and end times on the time line.

When you move the start time to 18:00 and the end time to 6:00, the list of objects in the change history changes.



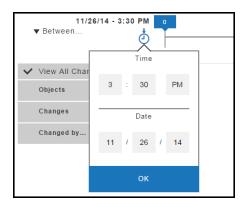
Setting Custom Time Increments

Example:



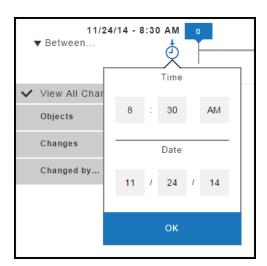
1. Select Between.

A drop down dialog box opens.



2. Select the time and date for the start of the time line.

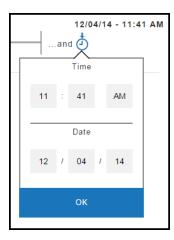
The default settings are 3:30 PM and eight days before the current date.



3. Click OK.

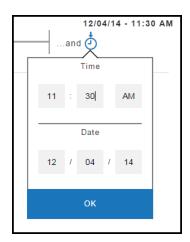
A drop down dialog box opens at the other end of the time line.

Example:



4. Select the time and date for the end of the time line.

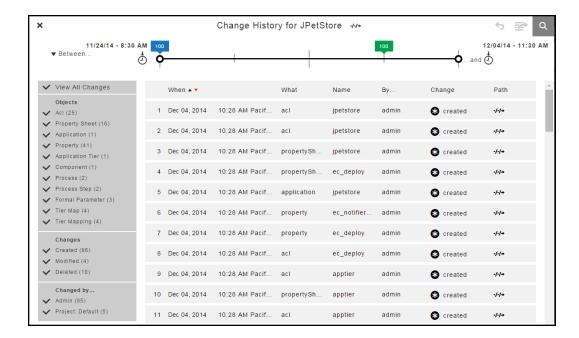
The defaults are 3:30 PM and the current date.



5. Click OK.

The time line changes to show only the changes from the start and end times and dates that you selected.

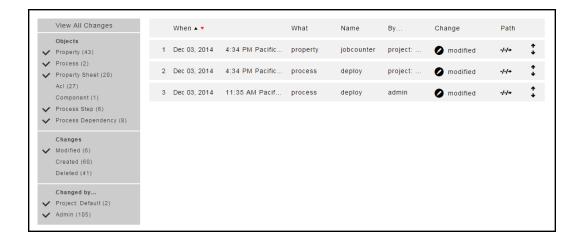
Example:



Change History Filters

The Change History has records only for an object and its tracked entities, not all entities belonging to the object. You can use filters to view changes to specific objects, the types of changes, and the users how made those changes if they are are tracked.

In this example, instead of selecting **View All Changes**, you can select specific objects, such as only properties, processes, property sheets, process steps, and process dependencies that have been modified by the Project:Default and Admin users.



When the list of filter criteria is long, not all of the criteria may appear in the filter list. To see all of the criteria, use the up or down arrows to see all the options.

This list does not show all of the users. Use the up and down arrows to see all four of the users.



Click the down arrow to see the other users.



Reverting Changes to a Tracked Object and Its Tracked Contents

The Change History displays a list of changes in reverse chronological order (from latest to oldest) for tracked entities in objects such as applications, procedures, workflows, workspaces, resources, and project-owned components such as library components. Each row in the Change History or a tracked object is a record of a change made either to the tracked object or to a tracked object owned by the object (directly or indirectly). For example, the Change History of a project would show both changes to the project itself and changes to any tracked object in the project, such as its procedures, the 'steps for each procedure, or properties belonging to the procedure steps.

Some types of objects (mostly run-time objects such as jobs, because tracking those would be resource-intensive and seldom useful) are not tracked by the Change Tracking feature. For a few types of tracked objects, some of their attributes are not tracked (for example, agents are tracked, but certain attributes of agents that vary frequently in real-time, such as their current status, are not). Changes to untracked objects, or to untracked attributes of tracked objects, are not shown in the Change History and cannot be reverted. In general, objects that are normally created or changed manually are tracked, and objects that are normally created or modified during automated processing are not, However, depending on your specific usage patterns, this behavior may not be recorded in the Change History, such as when you frequently run jobs that use ectool or API commands to automate the creation of new procedures or steps.

Follow these steps to select a tracked object and its tracked entities that you want to revert:

1. Go to the Change History for the object that you want to revert using one of these methods:



• Click the **Track Changes** button ()for the object.



- From the ElectricFlow Home page, click the Search button () to go to the "Change History Search" page and search for the object.
- 2. Locate the time period in which you are interested.

A list of changes for the tracked object appears.

- 3. If you decide to revert the changes for an object owned by the tracked object or for an object that owns the tracked object, use one of the methods to view the list of changes for this object:
 - Close the current Change History and navigate to the Change History for the object.



- Click the **View Path** button () on the top of the current Change History, copy the path, go to the "Change History Search" page, and search for a suitable addition to or truncation of that path.
- 4. Select a row (record) in the Change History corresponding to the earliest of the changes that you want to

You will be reverting the object and all its tracked contents to their state directly before the selected change was made.



a. If the **Revert** button () becomes enabled, ElectricFlow is capable of reverting the selected changes to the object. is tracking the changes to this object.

Go to Step 5.

b. If the **Revert** button is still disabled, ElectricFlow is not currently capable of reverting changes to the selected object. This could be due to tracking of this object being disabled and re-enabled, or the Change History feature being disabled and re-enabled. Also, not all types of tracked objects can currently be reverted (and due to mutual consistency requirements in the database format, certain types of objects cannot be reverted).

If you are unable to revert the selected object, you may still be able to achieve what you need to do one of these ways:

- Revert the parent object that owns the object you want to revert.
- Separately revert the individual objects owned by the selected object or even nested partial-reversions-of-reversions, possibly combined with manually reversion (undoing) certain changes.

Go to Step 3 to search for an for an owned or owning tracked object.

5. Before clicking the **Revert** button, make certain that you know all the changes that will be reverted, and have confirmed that this is what you want to do. (While it is possible to re-revert the changes caused by a mistaken reversion operation, doing so leaves a lot of activity in the Change History.)

IMPORTANT: All of the changes for tracked entities will be reverted, even the ones that you do not see in the Change History because of the timeline and filter settings.

- a. You may need to modify the timeline and filter settings to display all records in the Change History before you verify that all of the records are changes you want to revert.
- Make sure that the end of the timeline is set to Now and that the filters are set to View All Changes.
- c. Look through the entire list of changes from the most recent to the one you have selected, expanding them to see the details if needed, and confirm that you really want to revert all of them.

IMPORTANT: Tip on how to further refine what changes will be reverted

If you want to revert some of the changes shown but not others, you may be able to achieve what you want to do with some combination of the following methods:

- Reverting just certain objects owned by the object that you want to revert.
- Using nested or overlapping sets of reverts of owned objects.
- Undoing individual changes manually using the ElectricFlow UI.
- Using nested partial-reverts-of-reverts.

If you fully understand the XML format used for ElectricFlow exports, you can use the export API call to export the past and current states of a part of the object hierarchy, diff and merge them using a suitable tool for diffing and merging XML files, and then re-import the results. Caution is strongly suggested if you are using this approach to revert objects for which ElectricFlow does not currently support reversion, because:

- There are some places in the ElectricFlow XML data hierarchy where peer objects refer to each other in ways that need to be kept consistent (such as by index of current position in a list).
- For some types of objects, re-importing them can change their IDs and break references-by-id.
- 6. Click the **Revert** button. Depending on the size of the object and every tracked object it owns, this operation could take some time (as long as exporting a past state of them and then re-importing it).

If this operation is successful, the object and its tracked entities are reverted back to the state at the date and time of the record selected in Step 3. A message appears that the Revert operation was successful.

If this operation is not successful, the object and its tracked entities are not reverted, and a message appears about the failed operation.

Notifications

ElectricFlow supports email notifications, also referred to as email notifiers, for these ElectricFlow objects:

- Application processes
- Application process steps (type: process)
- · Jobs and job steps
- Workflow states
- · Procedure results
- Procedure step results
- · State definition results

Application Processes and Process Steps

ElectricFlow sends email notifications to users and groups at the application-process and application-processstep (type: process) levels. The notifications are triggered based on how the job finishes (the onCompletion event) and on the job outcome (success or failure).

When configuring email notifications:

- You specify users, groups, or email addresses as recipients of the notifications.
- You can also target notifications at specific environments.
- You can enable or disable notifications at the application-process and application-process-step levels.

Note: The user and group settings are defined and managed at the platform level. See the "Users and Groups" section for more information.

Go to Configuring Email Notifications in Application Processes about how to set up email notifications. When notifications are enabled and a default template is used, you do not have to perform additional steps to set up an email notification.

ElectricFlow provides default templates for success and failure outcomes. All the templates have a name, subject, body, and type content that are stored as properties in ElectricFlow.

You can edit existing templates or create new email templates in ElectricFlow. Go to Selecting and Editing Email Messages for Application Processes about how to edit HTML default notifications and to create new email templates.

ElectricFlow also sends notifications for assignees and approvers of manual process steps. Go to Manual Tasks and Steps on page 72for more information.

Jobs, Workflow States, and Procedures

For jobs and job steps, two types of email notifiers:

- On Start Notifier Sends an email when a job or job step starts.
- On Completion Notifier Sends an email when a job or job step completes.

For workflow states, three types of email notifiers:

- On Enter Notifier Sends an email when the state becomes the workflow's active state.
- On Start Notifier Sends an email after the state's subjob or subworkflow starts. If no subjob or subworkflow is defined for that state, the notifiers will not be sent.
- On Completion Notifier Sends an email after the state's subjob or subworkflow completes. If no subjob or subworkflow is defined for that state, the notifiers will not be sent.

Using email notifiers:

- Attaching an email notifier to a procedure results in a corresponding email notifier associated with the job that is created when the procedure is executed.
- Attaching an email notifier to a procedure step results in a corresponding email notifier associated with the job step created when the procedure referencing the procedure step is executed.
- Attaching an email notifier to a state definition results in a corresponding email notifier associated with the state that is created when the workflow is executed.

You must create an email configuration before you create and configure email notifiers at the platform level. Then you can send them to individuals and groups.

To create and edit email configurations, go to Email Configurations on page 510, and Email Configuration - create new or edit existing email configuration on page 510.

To create or edit email notifiers, go to Email Notifier - create new or edit existing email notifier on page 511. See the topics starting with "Email Notifier Template" for examples of HTML templates that you can use as-is or copy and rename to create a new template.

Pipelines

Notifications for pipelines are automatically sent in the following scenarios. You do not have to perform any additional steps to send email notifications.

- For a manual task in a stage, the appropriate notifications are sent to the user or group who performs that task and to the user or group who responds when the task is completed.
- When entry and exit gates have specified approvers, ElectricFlow sends approval notifications to them.
 - You can set the approvals such that approval from only one of the approvers is sufficient for the pipeline to progress. In this case, you can have one approval rule and have all the approvers assigned to that rule.
 - Instead, if you want all the key individuals to approve sequentially, you can specify multiple approval rules, each for a specific user.
 - In addition to specific users, approval rules support groups. Approval or rejection from any user in the group will allow the pipeline to progress or stop.

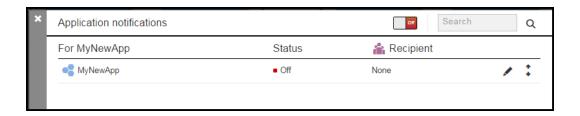
After approvers log into ElectricFlow, they can approve or reject the approval rules. The pipeline's progress will continue if the approval rules are approved. The pipeline will end with the appropriate error if the rules are rejected.

Go to Pipeline Concepts on page 218 for more information about notifications in pipelines.

Configuring Email Notifications in Application Processes

Review these guidelines before configuring email notifications:

- You can configure notifications for application processes and application process steps (type: process).
 When you open the Applications notifications dialog box and expand the application process details, only the application process steps (type: process) appear in the list.
- New email notifications are disabled in the application, its application processes, and the process steps before you configure them.



• You configure notifications in the "Application notifications / edit" dialog box.

IMPORTANT: The first time that you set notifications in this dialog box, the Notifications toggle changes to **On**. After you enter notification settings and click **OK**, email notifications are enabled at that level.

By default, the application expects that the user creates an email configuration called "default." The email configuration defaults to the server property /server/ec_deploy/ec_defaultEmailConfiguration, which is set to "default."

If you want to use a different name for the email configuration, change the value of /server/ec_deploy/ec_defaultEmailConfiguration to the new email configuration name.

Starting from the Home page, to set email notifications:

- 1. Go to the Applications List.
- 2. Select an application.

The Applications Visual Editor opens.

3. Click the Menu button.

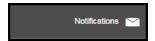
Example:



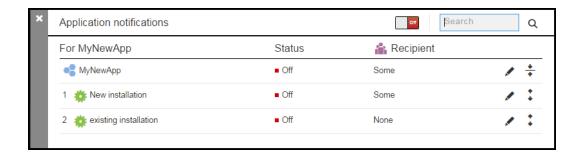
The Applications menu opens.

4. Click **Notifications** to add a new application.

Example:



The **Application notifications** dialog box opens.



5. Configure email notifications as follows:

You can configure one or more notifications in an application process or an application process step (type: process).

Configuring recipients

In the Who field, add users, groups, and/or email addresses.

When you start typing a user name, group name, or email address, a list of names or email addresses appear that match what you are typing.

Note: The user and group settings are defined and managed at the platform level. See the "Users and Groups" section for more information.

Example:



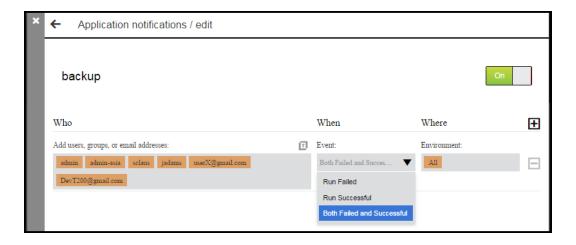
If one of the suggestions matches the name or email address, select it or continue typing. You can add more than one name or email address.



Configuring the event that triggers the notification

In the **When** field, select the event that triggers a notification to be sent to the recipients in the **Who** field. The default is **Both Failed and Successful**. Clicking in the **When** field shows a list of event triggers.

Example:



Configuring the environments where the notification applies

In the **Where** field, you select the environments to which the notifications apply. Clicking in the **Where** field shows the list of available environments to which the application is mapped in the tier map.



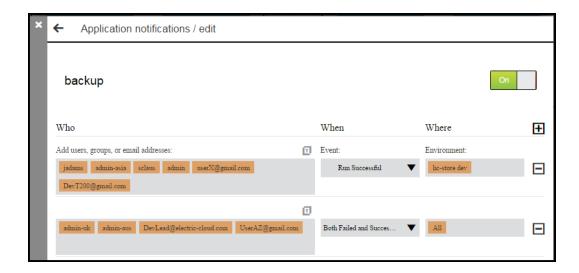
- 6. Select and edit the email notification message.
- 7. Add another notification for an application process or application process step.

Click the Add Notifications button to add a new notification.

Example:



After you have added your email notifications, click **OK** to save the settings and return to the **Application notifications / edit** dialog box.



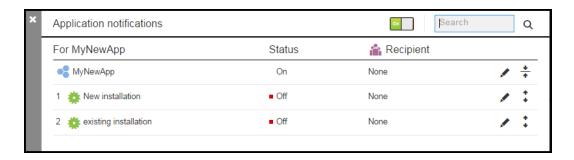
8. (Optional) Enable email notifications for the application processes or application process steps that are not already enabled.

To enable email notifications at the application level:

• Click the Notifications toggle and change it to On.

The status of the application changes to **On**.

Example:



• Click the Edit button to open the Application notifications / edit dialog box.

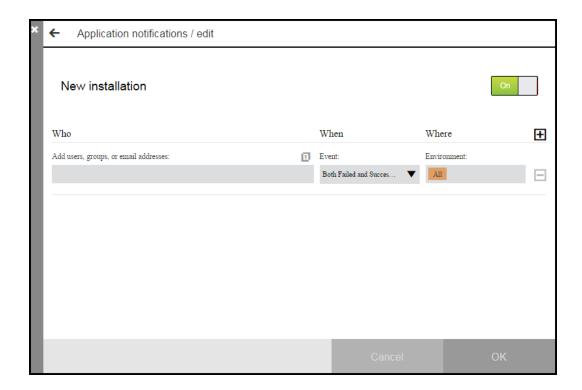
The Application notifications / edit dialog box appears. The Notification toggle changes to On.

Example:



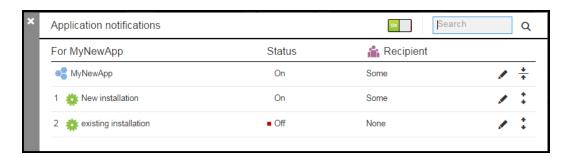
To enable notifications at the application process and process step levels, go to the **Application notifications** / **edit** dialog box for the specific process or process step.

The dialog box opens, and the Notifications toggle is now On.



When you enter notification settings in the dialog box and click **OK**, the settings are saved. The **Application notifications** dialog box appears and now shows that the application process status is **On**.

Example:



Selecting and Editing Email Messages for Application Processes

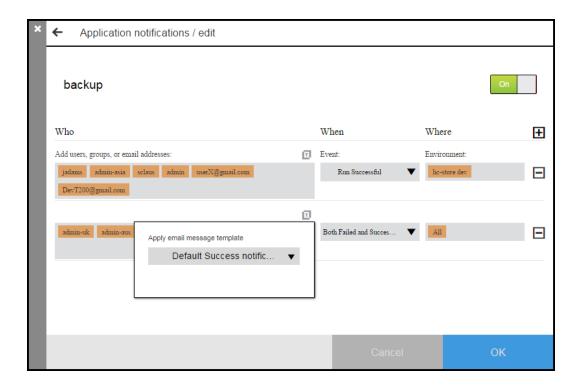
You can edit the email messages sent as notifications for application processes and application process steps (type: process).

Starting in the Application notifications / edit dialog box:

1. Click the Template button.

A drop-down box opens.

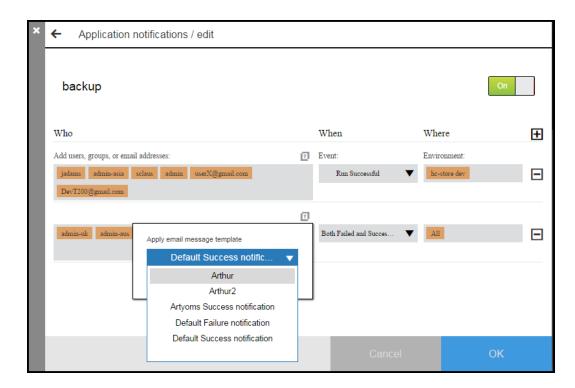
Example:



2. Click the down arrow to open the list of email message templates that can be applied to the application.

3. Select a template.

Example:



If the template is the current template applied to notification, the name of the template appears in the dialog box.

Example:

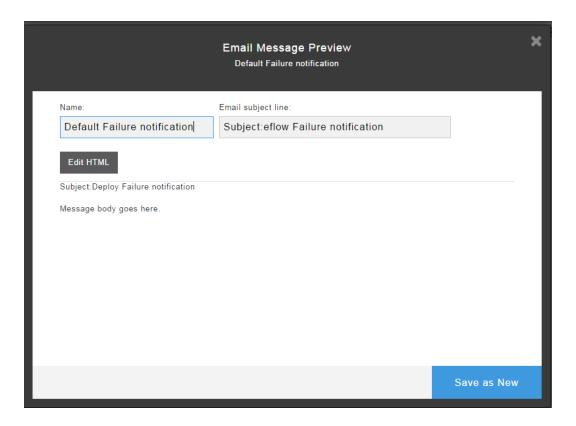


If the template is not the current template, the Apply and Edit buttons appear in the dialog box.



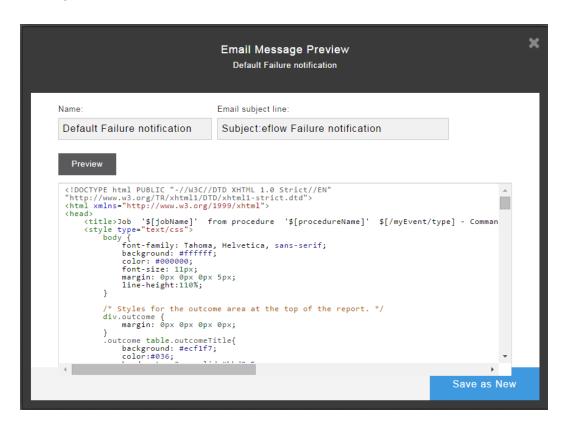
- 4. If you want to use the template that you selected instead of the current one and do not want to change selected one, skip the remaining steps in this task.
- 5. If you want to apply a different template or edit the template that you selected, do the remaining steps in this task.
 - If you click **Apply** to use the template as-is, skip the remaining steps.
 - If you click **Edit** to modify the template, go to the next step.

The **Email Message Preview** dialog box appears.



- 6. To edit the template:
 - a. Change the name of the template in the Name field.
 - b. Change the subject of the email in the Email subject line.
 - c. To modify the body of the email message, click **Edit HTML** and edit the HTML code.

Example:

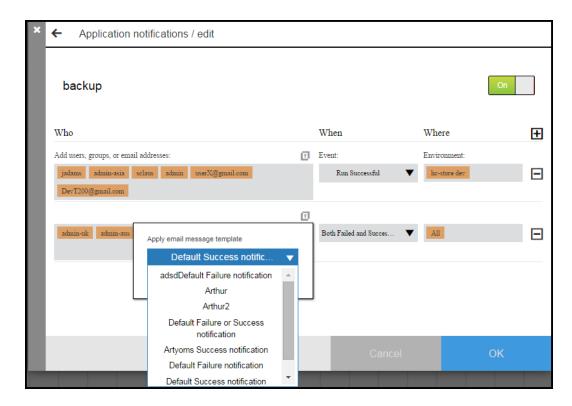


- 7. Click Preview.
- 8. To save your changes:
 - . Click Save Changes to save the changes in the existing template
 - Click Save as New to save the template as a new template.

The Application notifications / edit dialog box re-appears.

Click the down arrow to open the list of email message templates that be can applied to the application.
 The new email message template is in the list.

Example:



10. Click **OK** to save the settings.

Automation Platform

ElectricFlow is built on a powerful proven automation platform that natively integrates domain-specific capabilities for Enterprise-level continuous delivery. The automation platform gives distributed DevOps teams shared control and visibility into infrastructure, tool chains, and processes. It accelerates and automates the software delivery process and enables agility, availability, predictability, and security across many build-test, deployment, and release pipelines.

This section describes how to create, configure, and manage objects in the automation platform that make the automation of build-test processes, application deployments, and pipelines possible in ElectricFlow. Examples include:

- For build-test automation: Configuring resources and creating procedures to build and test your software.
- For deployment automation: Creating components based on artifacts, defining process steps with plugins, and assigning resources to environments.
- For pipelines: Defining tasks based on procedures, workflows, or plugins.
- For release management: Defining releases based on pipelines.

You use only the automation platform UI to configure ElectricFlow to perform these tasks. The objects and operations for build-test automation are the same as those for the automation platform.

You can also use Perl API commands through ec-perl and ectool, RESTful API commands, and DSL scripts. For information about using API commands and DSL scripts, see the ElectricFlow API Guide.

Build-Test Automation

Before you use ElectricFlow for build-test automation, deployment automation, pipelines, or release management, you must create, configure, and manage these objects in the automation platform:

For build-test automation, you must create, configure, and manage these objects in the automation platform:

Projects

A *project* is an object used in ElectricFlow to organize information. A project is a container object for procedures, steps, schedules, workflows, and properties. If you use ElectricFlow for different purposes, you can use a separate project for each purpose so different projects do not interfere with each other. When you work in one project, you do not normally see information in other projects. At the same time, a project can use information defined in other projects, which allows you to create shared library projects.

Resources

A *resource* is defined as an agent machine where steps can execute. A resource has a logical name and a host name. In some situations, it is convenient to have multiple logical resources associated with the same host. A resource can also be associated with one or more pools. Each resource has a *step limit* that determines the maximum number of steps that can execute simultaneously on the resource. Resources can be grouped into *resource pools*. Multiple resources can be defined on the same machine.

Procedures

Procedures and steps define tasks that you want ElectricFlow to execute. A procedure consists of one or more steps. A step includes a command or script executed on a single resource and is the smallest unit of work that ElectricFlow understands. Each step specifies a resource on which it should run (either the name of a specific machine or the name of a resource pool of equivalent machines, in which case ElectricFlow picks a machine from the pool). A step can be given a time limit, and if the step does not complete within the specified time, ElectricFlow automatically aborts it.

Steps are ordered within a procedure and normally execute sequentially. However, it is possible to mark a consecutive range of steps for parallel execution, so all steps in that range run concurrently.

You can define *parameters* for procedures. Parameter values are assigned when procedures are scheduled. Parameters can be required, optional, or have default values. Parameters are used for a variety of purposes such as specifying the branch to build or the set of platforms on which to run tests. Parameter values can be used in step commands and many other places.

Procedures can be nested. A step in one procedure can invoke another procedure instead of running a command. The invoking step provides parameters needed by the nested procedure, also referred to as a subprocedure.

Schedules

A schedule is used to execute procedures and determine when specific procedures run. A schedule can trigger at defined times, for example, every 2 hours from 10:00 pm to 6:00 am on Mondays, Wednesdays, and Fridays, or when modifications are checked into a particular branch of your source code control system. It is also possible to create a schedule that runs immediately and disappears after the job runs. When you create a schedule, you must provide the parameters required by the procedure that you want to invoke.

The Continuous Integration Dashboard works with your source code management (SCM) system and provides visibility into running builds, the ability to add a project to continuous integration quickly, and easily accessed configuration pages to setup or modify a continuous integration schedule.

Workflows

Managing a build-test-deploy product life cycle spanning multiple procedures and projects requires a significant amount of "meta-programming" and a heavy use of properties, and the *workflow* feature simplifies this process. Using the workflow object, you can create build-test-deploy life cycles by defining a set of states and transitions. Any ElectricFlow project can contain a workflow.

When a procedure is executed or run, a job is created. A job is an object that is created each time a
procedure begins to execute or run. The job keeps track of all data associated with the procedure's
execution, such as the running time of each step and any errors that may occur during the step.
ElectricFlow retains job information after the job completes so you can examine what occurred.

After setting resources, procedures, and schedules, ElectricFlow automatically runs the procedures that you created using these objects and facilities:

Zones and Gateways—A zone (or top-level network) that you create is a way to partition a collection of
agents to secure them from use by other groups. A gateway is a secured connection between two zones
when you want to share or transfer information between the zones. For example, you might want a
developers zone and a test zone. The ElectricFlow server is a member of the default zone, created
during ElectricFlow installation.

- Continuous Integration Builds and other schedules—Run jobs according to schedules that you define.
 Scheduled jobs can run at specific times or when source code changes are checked in to your source control system. ElectricFlow integrates with major source control systems. The Continuous Integration Dashboard allows you to add more projects easily and create build configurations quickly so you can visually see running builds, build status, and so on.
- Artifact Management functionality—Using artifacts can improve performance across builds, provide better
 reusability of components, and improve cross-team collaboration with greater tractability. For example,
 instead of developers repeatedly downloading third-party packages from external sources, these
 components can be published and versioned as an artifact. Developers then simply retrieves a specific
 artifact version from a local repository, guaranteeing a consistent package from build to build.
- Preflight build functionality—Used by developers to build and test code changes in isolation on their local machines before those changes are committed to a production build.
- Plugin capability-ElectricFlow is built with an extensible UI, enabling easy development of plugins that
 include integrations with other tools, custom dashboards, and unique user experiences based on roles.
 "Bundled" plugins, installed during ElectricFlow installation, provide easy integration with your SCM
 systems, defect tracking applications, and so on.
- Workflow functionality—Use a workflow to design and manage processes at a higher level than
 individual jobs. You can use workflows to combine procedures into processes to create build-test-deploy
 life cycles (for example). A workflow contains states and transitions that you define to provide complete
 control over your workflow process. The ElectricFlow Workflow feature allows you to define an unlimited
 range of large or small life cycle combinations to meet your needs.
- Resource management-If a resource is overcommitted, ElectricFlow delays some jobs until others are
 finished with the resource. You can define pools of equivalent resources, and ElectricFlow spreads
 usage across the pool.
- Recording a variety of information about each job, such as the running time and the success or failure of
 each step. A set of reports is available to provide even more information.
- Powerful and flexible reporting facilities—Various statistics such as number of compiles or test errors are
 collected after each step and recorded in the ElectricFlow database. A variety of reports can be
 generated from this information.
- Allowing you to observe jobs as they run and to cancel jobs.
- Credentials—Use a credential, consisting of an user name and password, can be attached directly to a
 step or schedule at the platform level. You can also attach impersonation credentials to procedure steps,
 procedures, and projects before executing the job step. It allows ElectricFlow to use a specific account
 with special privileges on a per-job or per-step basis.
- Workspace for each job, which is a disk area a job uses for storage

 —ElectricFlow also provides a facility for reclaiming space occupied by workspaces.
- Powerful data model based on properties—Properties are used to store job input data such as the source
 code branch to use for the build, to collect data during a job (such as number of errors or warnings), and
 to annotate the job after it completes (for example, a build has passed QA).
- Access control for users logged into the system–ElectricFlow uses this information to control their activities and integrates with Active Directory and LDAP repositories.
- Search, sort, and filter functions to minimize viewing or "wading" through information that is of no interest to you, allowing you to access to the information you need quickly.

- *Email notifications* to get important information or data to individuals or groups immediately and on a regular basis for a particular job or a specific job aspect.
- All ElectricFlow operations and features are available from a command-line application tool (Perl API), ectool, the RESTful API, DSL methods, and a user interface (UI).

For more information, go to ElectricFlow Platform Objects and Functionality on page 689.

Getting Started

Overview

Automation Platform Terminology

Getting Started Scenarios Help Topic Series

Navigating the ElectricFlow User Interface

Overview

ElectricFlow is an application for automating and managing your development life cycle, the software build, test, and deployment process. It helps development teams make these tasks repeatable, visible, and efficient.

ElectricFlow creates an environment where IT and Development organizations can work together to connect physical and virtual environments, processes, and tools already in use to create a private development cloud, also known as a *smart development cloud*. ElectricFlow is a scalable solution, solving some of the biggest challenges of managing "back-end" software development tasks.

Automation Platform Terminology

To get an overview of how the automation platform works and understand the ElectricFlow concepts and processes, review the following basic terms:

- Project
- Procedure
- Step
- Plugins
- Parameter
- Schedule
- · Continuous Integration
- Resource
- Job
- Reports
- Workspace
- Zones
- Gateways

See the glossary for a list of ElectricFlow terminology.

Getting Started Scenarios Help Topic Series

The series of scenario Help topics are designed to help you learn to use the ElectricFlow automation platform quickly. Follow the step-by-step scenario format to learn the automation platform basics:

- Creating a simple procedure–Scenario 1
- Creating a procedure that uses a Source Code Management (SCM) system—Scenario 2
- Setting up email notifications, reporting, and scheduling-Scenario 3
- Creating a multi-agent build procedure-Scenario 4

Each scenario builds on what you learned in the previous scenario, and each scenario ends with one or more scenario extensions to introduce you to other related ElectricFlow functionality.

Navigating the Automation Platform User Interface

To quickly link to task pages, view running or completed jobs, configure resources and so on, the automation platform web interface displays tabs across the top of the page. These tabs remain available at the top of all ElectricFlow automation-platform web pages.

When you select some of the main tabs, subtabs are provided. For example (see the next screen example), when the **Home** tab is selected, notice that the Overview tab is in bold font, which means you are looking at the Overview page. This page can be customized to show the information that you want to see quickly. You can use this page for shortcuts, quick links to jobs, "thumbnail" report views, and so on.

Selecting the Continuous Integration subtab takes you to the Continuous Integration Dashboard, which is place to configure your projects for continuous integration builds.

In addition to the main set of tabs, the top bar includes:

- Logged in as clearly shows the name of the logged-in user.
- Logout

 –Use this link to log out between ElectricFlow sessions if necessary.
- **Help**—This link provides a Help topic for the current page you are viewing, but if you select the down-arrow adjacent to the Help link, you have access to the following:
 - Tutorials—Use this link to display the Table of Contents for all currently available tutorial examples you may find useful as you become familiar with ElectricFlow.
 - Documentation

 —Use this link to access the entire ElectricFlow online Help system, which is fully searchable in the event you do not quickly see what you are looking for in the left-pane Table of Contents.

The Help Table of Contents contains feature overview/concept user-guide style Help topics, which are not linked to a particular web page. These topics are listed in the Table of Contents *above* the "Web Interface Help" section. The Web Interface Help section contains all of the page-specific Help topics.

Note: *To print a Help topic*— Right-click your mouse in the Help topic pane or go to your browser File menu.

- Visit the Community Site—This link takes you to the Electric Cloud Knowledge Base, all product documentation, and so on.
- Contact Support

 —Use this link to access the Electric Cloud Technical support phone number and
 email address. At a later time, you may choose to reconfigure this link to redirect it to your internal
 support time. See the "Reconfigure Contact Support" section in the "Customizing the ElectricFlow UI"
 Help topic.
- o About-Use this link to display the current version of ElectricFlow you are using.

Home Tab

You can customize the Home page to see the information you need at-a-glance or add Shortcuts to quickly return to the ElectricFlow pages you visit frequently. The Home page provides four categories:

• Job Configurations

ElectricFlow procedures can contain complex sets of parameters, making it difficult to remember the correct parameters for a particular situation and tedious to re-enter those parameters every time the procedure is invoked. Job Configurations provide a one-click solution to this problem.

Shortcuts

Use shortcuts to save frequently visited ElectricFlow web pages, so those pages are immediately accessible. You can also create a shortcut to any page on the web.

Jobs Quick View

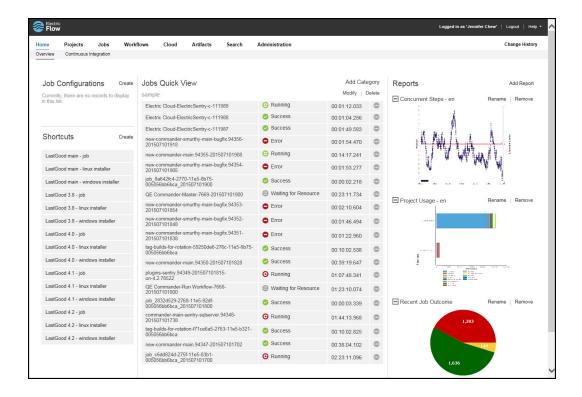
The Jobs Quick View allows you to define job categories that are interesting to you. You may be interested in a few jobs on this server. For example, you want only information about jobs that you launch manually and official builds for the products you work on. You do not want information about production builds for other products or personal jobs for other users.

Reports

You can configure reports you would like to see on a regular basis and display a thumbnail"report graphic in this section.

In the following Home page example, there are additional links. This page allows you to create, edit, and delete objects you see so that you can easily keep the page updated and current.

For more information about Home page functionality, see the Home Page Help topic.



Continuous Integration Subtab

The Continuous Integration Manager (CI Manager) provides a front-end user interface (the Continuous Integration Dashboard) for creating, managing, and monitoring continuous integration builds.

The Continuous Integration Dashboard provides:

- Visually seeing your running builds, build progress, build status, and accumulated build status.
- Easily accessed "Actions" to configure a continuous integration build.
- · Quick configuration of your preferred SCM system.
- A project that contains any number of continuous integration builds, depending on the work you have already setup for your procedures and steps to perform.

The following is an example of the dashboard. Your initial view will be blank until you add configurations (using your preferred SCM) to run your build procedures continuously. When this tab is opened to view the dashboard, click the **Help** link in the upper-right corner of the page for instructions to begin using the dashboard.



Projects, Jobs, and Workflow Tabs

Selecting the Projects, Jobs, or Workflows tab displays a table listing all previously configured projects, all completed jobs, or all previously configured workflows, respectively. Each table contains links in the table and links in the section above the table.

For example, selecting the Projects tab displays a table (see the next screen example), which has the projects you created and ElectricFlow default projects, which were added during ElectricFlow installation.

- Links at the top of the table include:
 - o A "star" icon-Click this icon to add this page to your Home page.
 - o A drop-down menu to choose which projects you want to see.
 - o Create Project link to go to the New Project page to define and add a new project.
 - New Search link to find a project whose name you may not remember or projects of the same type.
- Links within the table include:
 - Click on a Project name (first column) to go to the Projects Details page for that project.
 - Actions column–This column usually contains links to Track Changes, Edit, Copy, or Delete the object, which is a project in this example.

Each table, Projects, Jobs, or Workflows, has similar links to help you get more information, modify existing information, or create a new object. For more information on any of these ElectricFlow pages, see these Help topics: Projects, Jobs, or Workflows.



Cloud Tab

This tab opens to the Resources page and provides several other subtabs—Pools, Workspaces, Zones, Gateways, and the Cloud Manager plugin if it is installed.

Resources Page

The following example displays all resources that you configured ElectricFlow to use.

- In the information immediately after the **Resources** title at the top of the table, you can see how many licensed resources are in use.
- Use the row of buttons at the top of the page to perform resource management tasks.
- A resource name, in the Name column in the table, is a link to open the Resource Details panel for
 more information about the resource, and to access the Edit Resource panel to change or add access
 control privileges.

See the Resources Help topic for more information.



Pools Subtab

Click this subtab to display a list of all resource pools available to ElectricFlow. You may find it useful to group resources, creating one or more resource pools. For example, you can have resource pools for specific purposes, development teams, or other groups in your organization.

Similar to the Resources page, this page provides:

- A link to create a new resource pool.
- · A Search link to find an existing resource pool.
- Summary information for each resource pool in the list.
- Clicking the resource Pool Name takes you to the Resource Pool Details page for more information about that pool.
- And the Action column contains Copy and Delete links.

For more information about pools, see the Resource Pools or Resource Pool - create new or edit existing pool Help topics.

Workspaces Subtab

Similar to the Resources and Pools pages, the Workspaces page displays a table all available workspaces and summary information for each workspace.

Available links are:

- Create Workspace for adding another workspace.
- New Search to find an existing workspace.
- Clicking the Workspace Name takes you to the Edit Workspace page to make modifications.
- And the Action column contains Copy and Delete links.

For more information, see the Workspaces and Disk Space Management, Workspaces, or Workspaces - create new or edit existing workspace Help topics.

Zones and Gateways Subtabs

Similar to the other subtabs under the Cloud tab, the Zones and Gateways subtabs each have a table, listing all available zone or gateways, respectively.

- Selecting an object in the Name column takes you to the "Detail" page for more information.
- These pages are your management centers for zones or gateways, respectively.

For more information, see the Zones and Gateways Help topics.

Artifacts and Search Tabs

This page displays all artifacts available on this ElectricFlow server. It has Projects, Jobs, and Workflows pages, Create, Edit, and Delete links for the artifact. Clicking on an artifact Name (first column) takes you to the Artifact Details page. For more information, see the Artifact Management Help topic.



Artifact Versions Subtab

The Artifact Versions page displays all artifact versions available on the ElectricFlow server. Links in the table are the similar to other ElectricFlow pages. For more information, see the Artifact Versions Help topic.

Repositories Subtab

The Repositories page displays all artifact repositories available to the ElectricFlow server. For more information about repositories, see the Artifact Management and Repositories Help topics.

Search Tab

While many ElectricFlow web pages have a Search link, some do not. The Search page is always available. You can search any ElectricFlow supported object type on this page. For more information, see the Define Search Help topic.

Administration Tab

The next screen example shows all the current subtabs available when you select the Administration tab. Less-frequently used tasks are grouped for you as subtabs under the Administration tab. For example, after users and groups are set up, you do not need these pages unless a user or group configuration changes.



Administration subtab descriptions:

- Event Log-This page displays a log of events generated anywhere in the system, including jobs and workflows.
- Groups-Use this page to view a filtered list of ElectricFlow local users or groups.
- Users-Use this page to view a filtered list of ElectricFlow *local* users or groups.
- Licenses—This page displays all license information known to the ElectricFlow server. Typically, a single
 license is displayed, which describes the usage to which you are entitled. The license on the server may
 be based on concurrent resources, concurrent users, concurrent steps, registered hosts, or any
 combination of these licenses.
- Directory Providers—ElectricFlow uses account information from multiple sources. In most cases, the
 primary account information source is an external LDAP or Active Directory repository: both user and
 group information is retrieved from the repository. Local users and groups are defined in ElectricFlow.
- Email Configurations—This page displays all previously configured email configurations. You must create an email configuration so ElectricFlow can communicate with your mail server to send email notifications
- Database Configuration—If you do not use the ElectricFlow-supplied default database, use this page to
 configure another database to communicate with ElectricFlow. See the *ElectricFlow Installation Guide*for a list of approved databases for use with ElectricFlow.
- Plugins—This subtab displays the Plugin Manager page. A plugin is a collection of one or more features
 that can be added to ElectricFlow. Numerous plugins are bundled and installed with ElectricFlow,
 integrating third-party products seamlessly with ElectricFlow. For example, Source Code Management
 plugins are installed and so that you can configure your preferred SCM to communicate with
 ElectricFlow.

- Server–This page displays overall information about the ElectricFlow server. Use the Settings or Access Control links to make changes.
- Source Control–This page displays all SCM (source code management) configurations you have created to communicate with ElectricFlow.
- Defect Tracking This page displays previously configured Defect Tracking systems known to the ElectricFlow server. Use plugins to integrate ElectricFlow with numerous defect tracking systems.

Go to Scenario 1.

Getting Started - Scenario 1 - Creating a Simple Procedure

To begin, select the Projects tab, then the **Create Project** link at the top of the table. You need to create a project to contain the procedures you create.

On the New Project page, enter information in the fields as follows:

Field Name	Description
Name	Type a unique project name. You may want your project names to reflect the work groups or teams that will be using them. For example, you might set project names based on the products they support.
	For our scenario examples, we will use "FirstNewProject" for our project name.
Description	Enter a text description for your reference if you choose. ElectricFlow does not use or interpret this information.
Default Resource	(optional) For the purpose of "getting started", leave this field blank to use the Default resource already created during the ElectricFlow installation.
Default Workspace	(optional) For the purpose of "getting started", leave this field blank to use the Default workspace already created during the ElectricFlow installation.

Click **OK** to save the information you entered.

After clicking **OK**, you see the Project Details page and the name of the new project adjacent to the page title.

- For our scenarios, the project name is FirstNewProject.
- Your new project name will appear on the Projects page also.



Scenario requirement: No additional requirements for this scenario. This scenario uses the *local* agent (resource) that was included on your local machine when ElectricFlow was installed.

Overview

This scenario establishes that ElectricFlow and its components were installed successfully and guides you through creating a simple procedure with a parameter that will echo Hello World. After you create this procedure, you will run the procedure and review the results.

At the end of this scenario, you will be familiar with the following ElectricFlow concepts and features:

- Projects
- Procedures
- Steps
- Parameters
- Jobs
- · Navigating the Job Details web page

Begin Scenario 1

Step1. Select the Projects tab

On the Projects page, you can see the default projects installed during the ElectricFlow installation and the new project you just created.

Select your new project name to go to its Project Details page, then click the Create Procedure link.



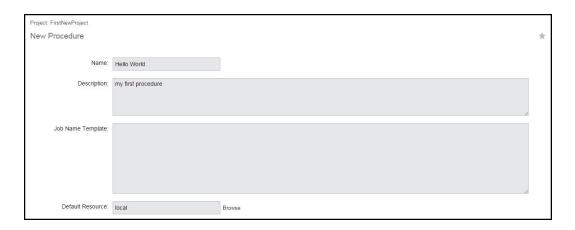
Step 2. Create a procedure

On the New Procedure page, you need to be concerned with the section of the page illustrated below only.

Enter the following information:

Field Name	Description
Name	For this scenario, Hello World is the procedure name. At a later time when you are creating your own procedures, you can choose any unique name that is most meaningful for procedures in your project.
Description	(optional) Enter a text description for your procedure if you choose to do so.
Default Resource	Use the Browse link and choose local. Using local accesses the agent installed during the ElectricFlow installation. At a later time after you have configured other resources, your Browse list could contain many choices.

Your New Procedure page will look similar to the following:



Before leaving this page, notice the "breadcrumb" in the top-left corner. Most ElectricFlow web pages contain a breadcrumb for easy return to the previous page.

Click **OK** to continue.

After clicking **OK**, ElectricFlow takes you to the Procedure Details page—this is the page you use to add important information to your procedure. When your procedure contains steps, parameters, or email notifiers, and so on, those objects will be displayed here for your reference.

Step 3. Create a parameter

An ElectricFlow procedure can define one or more parameters. These parameters are similar to the parameters of a subroutine in a programming language. Parameters allow you to define a "reusable" procedure by using symbolic variable names in place of a single, fixed value. When you run the procedure, you can type in a value to use for that particular job. You can refer to the parameter in a step for a procedure, using the standard

ElectricFlow "property reference syntax" - \$[...].

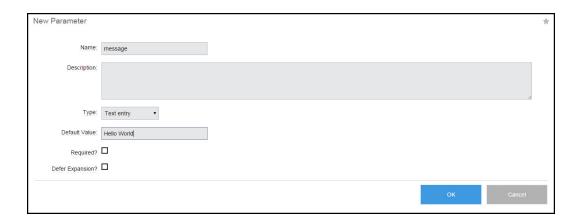
For example, see **Step 4**, specifically the text in the Command(s) text box.

On the Procedure Details page, click the **Create Parameter** link to go to the New Parameter page.

Enter the following information:

Field Name	Description
Name	Enter the parameter name message for this scenario.
Туре	Use the default Text entry for the parameter type.
Default Value	Type-in Hello World for the default parameter value.

Your New Parameter entries will look like the following:



Click **OK** after filling in the fields and ElectricFlow returns you to the Procedure Details page.

On the Procedure Details page:

Notice the Parameters table is populated with your parameter entry.



Also notice the name of the procedure we created (Hello World) is shown adjacent to the page title for your reference.

Step 4. Create a step

Now we can create a step on this procedure. On the Procedure Details page, to create a New Step for our current purpose, click the **Command** link to go to the New Step page.

On the New Step page:

The following screen displays the portion of the New Step page you need for this scenario.

Notice this version of the New Step page contains a Command text box to enter the information you need. At a later time, when you choose to create a different step type, the New Step page will be populated with the appropriate fields required for that step type.

Enter the following information:

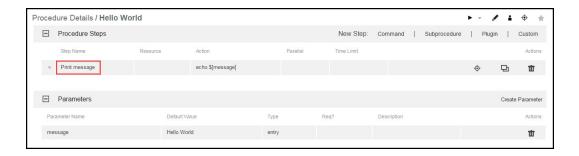
Field Name	Description
Step Name	Use Print message for the step name. For work outside this scenario, you can enter any unique name of your choice for the step name.
Description	(optional) Add a text description about this step if you need one.

Field Name	Description
	Type echo \$[message]
	In this command, message refers to the parameter we created.
Command(s) box	Note: The ElectricFlow Server does not actually interpret (or even understand) the contents of this field. ElectricFlow simply expands all property references and then passes the resulting text block to the ElectricFlow Agent. The Agent copies the text block to a temporary file and then invokes the Agent's native shell command, passing the name of the temporary file as a parameter to the shell command. The overall result is the same as if you had created this block of text as a BAT file (on Windows) or a shell script (on Linux). The operating system of the Agent machine takes over from that point to actually run your commands.
Resource(s)	For this scenario, leave this field blank to use the default resource you set on the procedure. If multiple resources are configured, you have the option to change the resource for each step if you prefer.

Your New Step page will look like the following example:



Click **OK** to create the step and ElectricFlow returns you to the Procedure Details page to see its entry as displayed in the next screen example.



Step 5. Run the procedure

Notice the Run-arrow icon among the links at the top of the tables. If you click the **Run** link, the procedure will run as-is immediately.

For this scenario, hover your mouse over the adjacent down-arrow and select the Run... option.

Selecting this option provides the Run Procedure page where you can add or modify the parameter values as necessary.



The Run Procedure page has the information we want for this scenario, so click the Run button.

When a procedure runs, it creates a "job".

After clicking **Run**, ElectricFlow takes you to the Job Details page to see the status of the "job" and watch its running progress.

This web page supplies a wealth of other information.



The Job Details page contains two sections. The section at the top of the page is the job's summary, displaying the job's status and other general information:

- The "green check" icon is displayed to show the job completed successfully.
- The "Completed with Success" section displays the result of the procedure that ran.
- The General Information section displays the project and procedure names and who launched the job.

The section below the summary area displays detailed job information, including tabs to expand information and links to access various other ElectricFlow web pages.

- Notice the "Success" entry in the Status column—this is your step result.
- In the Log column, click the icon in the Print message row to see the outcome/result of your Hello World Print message step. Your file will look similar to the following:

Job: job_6_20150711160208

Workspace File / **Print message.db7dd766-2820-11e5-8679-0050568f747c.log**Hello World

• Select the Parameters tab to see the parameter you created and its value.

For complete information about the Job Details page functions and available information, click the **Help** link in the top-right corner of this ElectricFlow web page.

Scenario Extension - Adding Another Step

Each procedure you create can contain as many steps as you need. We created only one step in this scenario. This example shows how to modify a procedure by adding additional steps. You can always add more steps or edit an existing step before you run a procedure again.

To add a step to an existing procedure:

- 1. On the Job Details page, select your procedure name (at the top of the page) to go to the Procedure Details page.
- 2. On the Procedure Details page, click the **Command** New Step link again.
- 3. On the New Step page, enter information in the fields as follows:
 - Name-Print Step Info
 - Command(s)–Enter the following text:

```
print "This step was launched by $[launchedByUser] on the resource
$[/myResource/resourceName].\n";
```

- Shell-ec-perl
- o Click **OK** to save the step and return to the Procedure Details page.

Note: You can move this new step to the top of the procedure by hovering your mouse over the icon next to the step name, then drag-and-drop the step to the top of the list. This is a useful function particularly if you find you have created steps in a different order in which you need them.

Click Run (from the Procedure Details page).

Finding an ElectricFlow Automation Platform Web Page

If you "lose" the Job Details page or any other ElectricFlow web page, several navigation methods are available:

- Use your browser's Back button if you were recently on a page you want to revisit.
- Use ElectricFlow "breadcrumbs." On some pages you see breadcrumbs on the left-side of the page that note your current position/level within the project.
- Use the tabs at the top of the page to drill-down to the page you need. For example:
 - To find the Job Details page:
 - Select the Jobs tab and notice that your procedure name is now listed in the Job column, accompanied by a date, job number, and other related job information on the same row. When you run a procedure, it becomes a job. Select the job you want to see and you will go to the Job Details page for that job.
 - To find the Procedure Details page:
 - Select the Projects tab, then select your project to go to the Project Details page to see a list of procedures in that project. Click on a procedure to go to the Procedure Details page for that procedure.

Note: A "project" is a container for all related procedures you create. You can create as many projects as you need, using the **Create Projects** link on the Projects page.

As you use ElectricFlow, you will quickly learn other navigation paths for your particular interests. In addition, most ElectricFlow web pages contain a "star" icon in the upper-right corner. Click this icon to add that page to your Home page for quick "one-click" access in the future.

Summary

This simple Hello World procedure demonstrated how to quickly create a procedure and how any command can be wrapped and parameterized using an ElectricFlow procedure. If you have an existing script you want to automate, you can call that script within a step command block. This scenario also introduced you to running a job, entering parameters for the job, and viewing results for the job you ran.

Go to Scenario 2

Getting Started - Scenario 2 - Creating a Procedure That Uses an SCM

Scenario requirements: Before you begin this scenario, ElectricFlow needs to be able to access your Source Control Management (SCM) system.

To set up a Source Control system (SCM) to use with ElectricFlow, you need:

- a running SCM system
- an active source depot in your SCM system
- an ElectricFlow agent that can communicate with your SCM server

Note: If the machine where you installed ElectricFlow can communicate with your Source Control (SCM) server, this requirement is met. *If not*, return to the installation executable file and run the ElectricFlow Agent install program on a machine that has access to your SCM system.

To install ElectricFlow agent software on another machine, copy the installation executable file to that machine, then double-click on the file to run the installation program. Select Express Agent when the install window opens.

When installation is complete, configure this new agent as an ElectricFlow Resource. See the beginning of Scenario 4 for more information.

- o Configure ElectricFlow to communicate with your SCM system.
- Select the Administration tab, then the Source Control subtab.
- The Source Control Configurations page displays a table listing Source Control configurations after you create them. Click the Create Configuration link and the following web page is displayed.



 On the New Source Control Configuration web page, use the drop-down menu to choose your SCM type.

For our example, we chose Perforce.

 Depending on which SCM you choose, the page expands to display appropriate fields for entering values to configure your SCM.

See the next screen example.

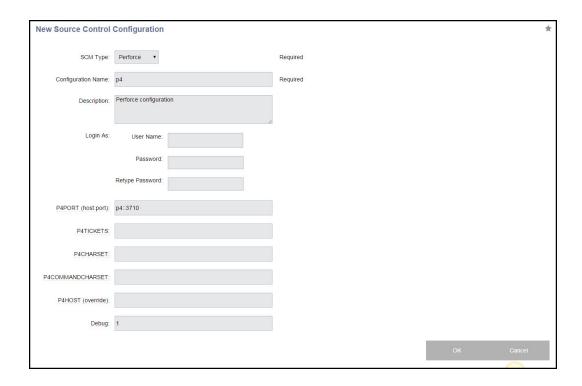
Enter information in the fields as follows:

• Configuration Name - For our example, we use the name, "p4". You can choose any name for your SCM Configuration.

Note: You will need this name later, so remember the name you choose for your SCM configuration. A configuration name is important because you may want to create more than one source control configuration.

 P4PORT - For our example, we use p4:3710 to designate the P4PORT. Your specification will be different.

Important: Other fields remain blank to use the defaults, but you may need to specify some or all information in the remaining fields depending on how your internal systems and access are set up for access to your SCM system.



Click **OK** after filling-in the fields.

For additional help with this ElectricFlow web page, click the **Help** link in the top-right corner of the web page.

Overview

This scenario guides you through creating a common build process integrated with your SCM system and build utility. And you will learn how to navigate and modify the procedure definition using the ElectricFlow Procedure Details web page.

At the end of this scenario, you will be familiar with the following ElectricFlow concepts and features:

- Source control configuration
- · Creating more complex steps
- Creating a build process
- Navigating the Procedure Details web page

Begin Scenario 2

Step 1. Select the Projects tab

On the Projects page, select the project name you created in Scenario 1 (in the first column) to go to the Project Details page, then select the **Create Procedure** link to begin.

Step 2. Create a procedure

On the New Procedure page, enter a Procedure Name - it is required.

You can use any unique name you choose—*do not* use the same procedure name you chose for Scenario 1. Procedure names within the same project must be unique.

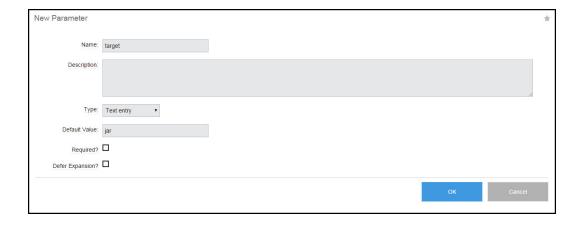
- Name For this scenario, we use Basic Build for the procedure name.
- Default Resource Use local for the resource name.

Click \mathbf{OK} to go to the Procedure Details page.



Step 3. Create a parameter

On the Procedure Details page (illustrated above), click the **Create Parameter** link to go to the New Parameter page.



- Name For our example, the parameter name is target. You can choose any unique name of your choice for your parameter name—parameter names must be unique within a procedure.
- Default Value We chose jar for this value.

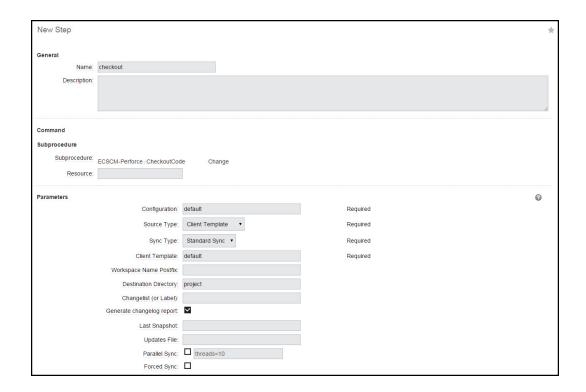
Click **OK** to continue and to return to the Procedure Details page to see the parameter you just created.

Step 4. Create a step to "checkout" source files from your SCM system

From the same Procedure Details page, click the Plugin link to create a New Step.

- In the Choose Step dialog, select Source Code Management from the left pane.
- Next, select ECSCM-Perforce from the list, then select "Perforce Checkout" from the right pane.
- Click Close to go to the New Step page, which will be populated with the fields you need to enter for the Perforce SCM.
 - Both the Subprocedure and the Parameters sections are populated with fields to enter information so ElectricFlow can communicate with your Perforce SCM system, using the ElectricFlow Perforce plugin.

Note: Depending on the SCM you chose to work with, including Perforce, the New Step screen you see may be different than the following example. If you need additional help configuring your SCM, see the corresponding SCM plugin on the Plugin Manager page. Each plugin has its own Help topic.



Enter information in the fields as follows:

Field Name	Description
Name	Enter a unique name for your subprocedure step. For this scenario, the step name we supplied is checkout.
Description	(Optional) Enter a text description for this step.

In the Subprocedure section:

This section displays the subprocedure name or a plugin name because the plugin is being called as subprocedure.

Note: The plugin name defined: "EC" is for Electric Cloud, "SCM" identifies the plugin as belonging to the Source Control Management category, and "Perforce" identifies the name of the source control system.

Our example illustrates the fields for the Perforce SCM plugin. Your fields will be different if you are using a different SCM. These parameter fields request some of the same values you already specified when you created the Source Control Configuration.

	Clicking this link takes you to the Project Details page for the plugin. Generally,
ECSCM-Perforce	and for this scenario, you do not need to access this page to make changes—it
	is for your reference only.

Field Name	Description	
CheckoutCode	Clicking this link takes you to the Procedure Details page for this procedure. Generally, and for this scenario, you do not need to access this page to make changes—it is for your reference only.	
Change	Clicking this link displays a dialog that allows you to change the project, plugin, or subprocedure. If you make a change in the Change Subprocedure dialog, the Subprocedure and Parameters sections on the New Step page will update automatically with corresponding fields for your new choice. Note: All SCM plugins that can checkout code contain the checkoutCode procedure.	
Resource	The local resource was specified earlier. If this field remains blank, the local resource is used.	
Note: The following two parameter fields are common among all SCM plugins:		
Configuration	The name of your Source Control Configuration.	
dest	The destination where you want to put your checked-out code—it is the directory within your job workspace. Each job adds a new directory to the workspace. At this point, you have not defined multiple workspaces, so the ElectricFlow default workspace is used.	

For more information about workspaces, see the "Workspaces and Disk Space Management" Help topic.

Click **OK** after entering values and to continue to the Procedure Details page, which should now look similar to the following example:



The Procedure Details page now displays the name of our new procedure, Basic Build, and its new checkout step using the Perforce SCM, and the target parameter with the jar value.

Step 5. Create a step to run the build

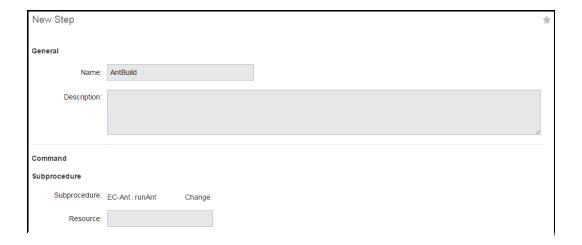
This step can be another subprocedure step, a command step, or a custom step. Any procedure can contain multiple types of steps, any of which may be either simple or complex.

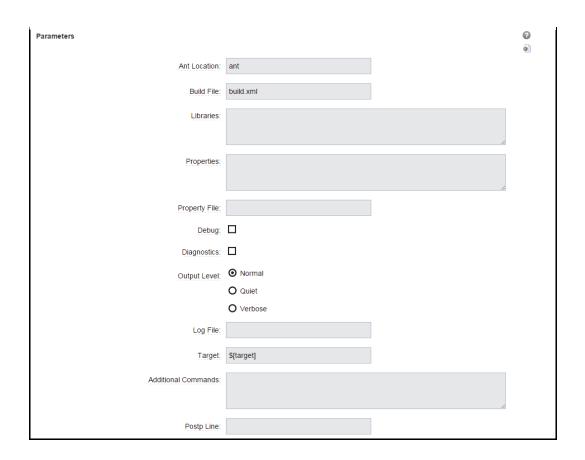
Steps can contain or use:

- · parameters and/or properties
- · many of your existing scripts, including shell scripts

For this step example (to run the build), we will create an Ant step. To invoke Ant, you could write a script that calls Ant directly, or you could call the ElectricFlow Ant plugin to invoke the runAnt procedure. For this example, we will use the ElectricFlow Ant plugin.

- From the Procedure Details page, click the **Plugin** link to see the Choose Step dialog and select the Build category.
- Select the EC-Ant plugin to see and select the "Ant Run Ant" step.
- When the New Step page is displayed, notice the Subprocedure section is populated with the chosen EC-Ant plugin, and the Parameters section provides the appropriate fields for this Ant step.





Enter information in the fields as follows:

Field Name	Description	
Name	For this example, we specified AntBuild for the step name. You can use any unique name you choose for this step.	
Subprocedure section:		
Note the EC-Ant plugin project name and the runAnt procedure name.		
Resource	Leave this field blank to continue using the local resource.	
Parameters section:		
Build File	Use build.xml.	
Target	For this example, we used \$[target].	
Working Directory	For this example, we used project.	

For this scenario, we leave the remaining fields blank, but you may need to enter additional information depending on your Ant invocation.

Click **OK** to continue and return to the Procedure Details page.

The following screen example illustrates the two steps we created for this scenario.



Step 6. Running the procedure

On the Procedure Details page, hover your mouse over the small down-arrow on the right-side of the **Run** link, choosing **Run...**

- Run Immediately This option runs the procedure immediately as set.
- **Run...** This option displays the Run Procedure web page where you may alter any existing parameters for this procedure, or set an existing credential.

Note: This scenario does not introduce you to the Advanced section of the Run Procedure page. For more information on these topics, click the **Help** link in the top-right corner of the Run Procedure page.

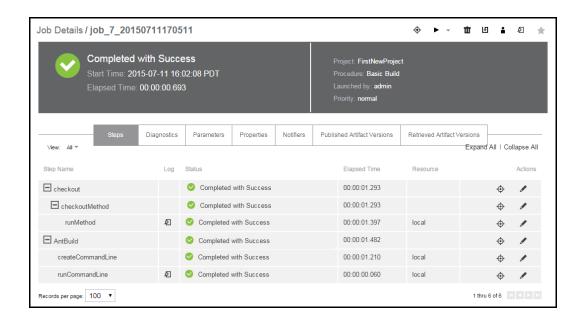


Click **Run** to run the procedure and continue to the Job Details page to see the status of your running job.

Step 7. On the Job Details page

The Job Details page now displays the status and results of running the Basic Build procedure.

Your Job Details page will look somewhat different if you supplied other values in this scenario.



More about the Job Details page

- The General Information section at the top of the table provides links back to the **Project > Procedure** that was launched to create this job.
- Click on a Step Name to go to the Job Step Details page.
- Click the Parameters tab to see parameters in this job.
- If there are errors, the Diagnostics tab has specific information.

More about the Procedure Details page

To go to the Procedure Details page at any time:

- · Select the Projects tab.
- Select the Project name (that contains the procedures you want to see) to go to the Project Details page.
- Select a Procedure name to go to the Procedure Details page.

You may want to review to the Procedure Details page frequently because it displays the steps that make up your procedure, the parameters contained in the procedure, email notifiers, and so on. Full edit capabilities are available by selecting any object (step, parameter, and so on).

Also, you can run this procedure again "on demand" by clicking the **Run** link at the top of the page.

For more information about the Procedure Details page functions, click the **Help** link in the top-right corner of the Procedure Details page.

Scenario extension - Add a step to show workspace contents

If you are curious about what your current default workspace contains at this point, you can create a step to see the workspace.

- Select the Projects tab, select FirstNewProject (or your project name if you used something different), select the Basic Build procedure name.
- Click the **Command** link to create a New Step—a new *command* step.
- The New Step page is displayed with a Command text box.
 - For Name, use seeWorkspace.
 - o In the Command(s) text box:
 - for Windows type dir /s
 - for Linux type ls -r
- Click **OK** to create the step and go to the Procedure Details page.
- Click Run to create the job and to go to the Job Details page.

When the job is completed, click the log icon in the Log column to see the contents of the current workspace for this procedure. Your Workspace File will be similar to the following example:

```
Job: job 7 20150711170511
Workspace File - seeWorkspace.1375.log
Volume in drive N has no label.
Volume Serial Number is EAB8-74CE
Directory of N:\job_7_20150711170511
07/11/15 09:43 AM <DIR>
07/11/15 09:43 AM <DIR>
07/11/15 09:43 AM <DIR>
07/11/15 09:43 AM
07/11/15 09:43 AM
                      <DIR>
                                      project
                               2,770 runMethod-checkoutCode-1377.log
                                 0 seeWorkspace.1375.log
             2 File(s)
                                2,770 bytes
Directory of N:\job_7_20150711170511\project
07/11/15 09:43 AM <DIR>
07/11/15 09:43 AM <DIR>
07/11/15 09:43 AM
                              1,141 build.xml
07/11/15 09:43 AM
                                338 Makefile
07/11/15 09:43 AM 07/11/15 09:43 AM <DIR>
                                 338 Makefile2
                                      src
           3 File(s) 1,817 bytes
Directory of N:\job_7_20150711170511\project/src
07/11/15 09:43 AM <DIR>
07/11/15 09:43 AM <DIR>
07/11/15 09:43 AM
1 File(s)
                                      . .
                                 144 HelloWorld.java
                                 144 bytes
    Total Files Listed:
                          4,731 bytes
              6 File(s)
              8 Dir(s) 244,679,413,760 bytes free
```

Summary

This scenario demonstrated how easily ElectricFlow can integrate with your SCM system and build utilities.

Go to Scenario 3

Getting Started - Scenario 3 - Notification, Scheduling, and Reporting

Scenario requirements: For this scenario, you need to set up an Email Configuration to use the email notification feature.

To set up an Email Configuration:

Select the Administration tab, then select the Email Configurations subtab.
 You will see a blank Email Configurations page. As you create one or more email configurations, they

will be listed in table format on this page.

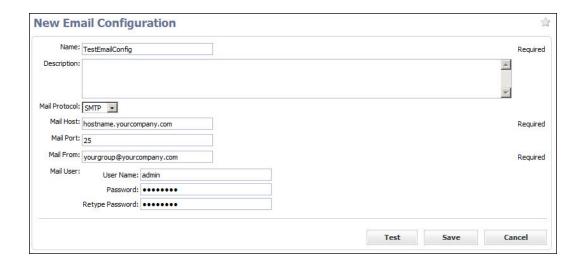
• Click the Add Configuration link.

On the New Email Configuration page, enter information in the fields to define your mail system.

Field Name	Description
Name	Enter a unique name for your email configuration.
Description	(Optional) Enter a text description for your reference.
Mail Protocol	Use the drop-down arrow to make a selection.
Mail Host	Enter the name of your mail host.

Enter information in the remaining fields according to the specifications you need.

Your New Email Configuration page should look similar to the following example:



Click **Save** after filling-in the fields or click the **Test** button first if you want to make sure your email configuration works.

Enter an email address in the pop-up box and click Send.

Note: For additional help with this ElectricFlow page, click the Help link in the top-right corner of the page.

Overview

This scenario guides you through creating an email notification, invoking your build process with a timed (standard) schedule, invoking your build process using a continuous build integration trigger, and introduces ElectricFlow reporting features.

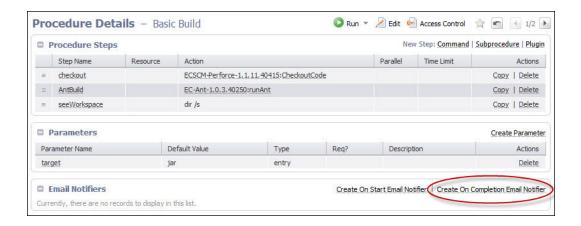
At the end of this scenario, you will be familiar with the following ElectricFlow concepts and features:

- Creating an email configuration
- · Creating an email notifier
- · Creating schedules
- Creating reports

Begin Scenario 3

Step 1. Creating an email notifier

On the Procedure Details page for the Basic Build procedure, click the **Create On Completion Email Notifier** link to begin. This link takes you to the New On Completion Email Notifier page.



Two email notifier types for jobs or job steps:

- On Start Notifier-Sends an email when a job or job step starts.
- On Completion Notifier–Sends an email when a job or job step completes.

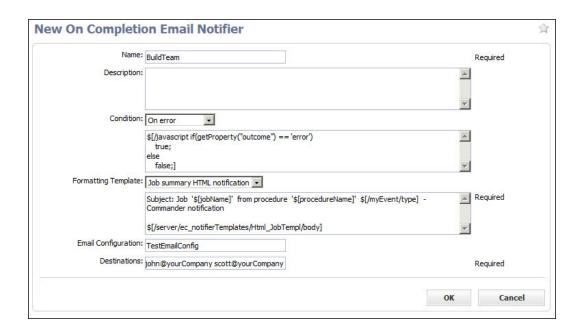
The fields for each of these notifiers are similar, but we chose the On Completion notifier because you may find you use this one more frequently.

On the New On Completion Email Notifier page:

Enter information in the fields as follows:

Field Name	Description
Name	For this example, we use BuildTeam. This name can be an arbitrary text string.
Description	(Optional) Enter a text description for your reference. ElectricFlow does not use or interpret this information.
Condition	Use the pull-down menu to select the type of condition you need for this email notifier. Edit the auto-supplied condition in the text box or add a completely new script for your purpose. The condition specifies whether the notifier should send a message depending on the result of a property expansion. If the result is empty, non-zero, or "true", the message is sent. If the result is "0" or "false", the message is not sent.
	For this example, choose the "Job summary HTML" notification template, or you can use the drop-down menu to select from a list of global, ready-to-use formatting templates. Depending on the type of email notifier you are creating, the available template choices in the drop-down menu will be different.
Formatting Template	To customize your template, edit the auto-supplied text in the Formatting Template text box, or you can add a completely new template for your purpose. Note: Any edits made in this text box will not be saved to the global template, and the template will appear as a Custom template when you return to this notifier definition. Additionally, make sure the content is formatted correctly, i.e., no illegal characters or spacing.
Email Configuration	Use TestEmailConfig, the email configuration we created.
Destinations	This is a space-separated list of valid email addresses, email aliases, or ElectricFlow user names, or a property reference that expands into such a list.

After filling-in the fields, your page may look similar to the following example:



Click **OK** to create your email notifier and to return to the Procedure Details page, where you can now see the Email Notifier you created.

Note: For more information, see the Email Notifiers Help topic.

Step 2. Creating a standard schedule

To create a standard "timed" schedule:

- · Select the Projects tab.
- Select your project name to go to the Project Details page.
- Select the Schedules subtab, then select the down-arrow to the right of the "Schedule" link, and choose "Standard Schedule" to go to the New Schedule page.

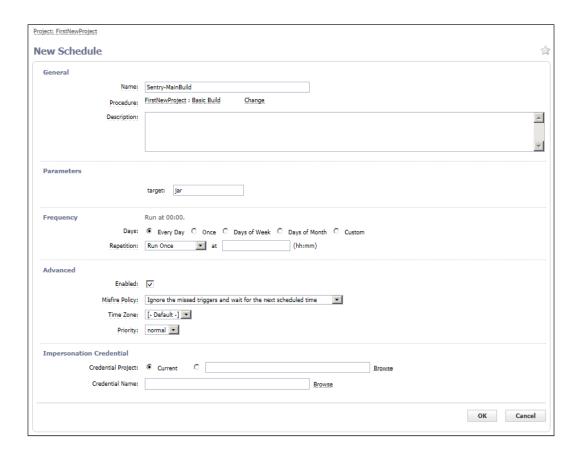


On the New Schedule page:

The following is an example of the New Schedule page.

• Notice the "breadcrumb" above the New Schedule page title—this entry displays the name of the project that will contain this schedule.

- After entering a name (unique within your project) for your schedule (we chose "Sentry-MainBuild"), choose the procedure your schedule will run.
 - o Click the Change link adjacent to the Procedure field.
 - o In the pop-up menu, leave the Current project selected.
 - o Place your cursor in the Procedure field to see a list of procedures in your project.
 - o Select the procedure you want this schedule to run. We chose "Basic Build".



Enter information in the remaining fields as follows:

Field Name	Description
Description	(optional) A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> <tt> <tt> </td></tr></tbody></table></style></pre>

Field Name	Description	
Parameters section		
If the procedure has parameters, enter their values in this section. The procedure cannot execute unless all required parameters are provided.		
Frequency section		
"Run at" is a summary section: As you make selections for the following Days and Repetition sections, a summary of what you selected appears here.		
Days	Select weekdays, month days, or create a custom frequency when you want the schedule to run.	
	Note: The "day" applies to the start time if the time range spans two days, that is, if the time range crosses the midnight boundary.	
Repetition	If Run Once - the schedule will run once per selected day, at the time you specify.	
	If Run Every - the procedure will execute repeatedly during the specified time range. For example, if you select a time range from 11:00-14:00 and an	
	interval of 40 minutes, the procedure will execute 5 times on each of the selected days: at 11:00 am, 11:40 am, 12:20 pm, 1:00 pm, and 1:40 pm.	
	If Run Continuously - as soon as one job ends, the scheduler will trigger immediately to run the job again.	
Advanced section		
Enabled	If this box is "checked," the schedule will run. You can disable this schedule whenever necessary.	

Note: For this scenario, you do not need to enter information for any remaining fields. Later, when creating an actual timed schedule for one of your procedures, refer to the Help topic for this page for more information.

Click **OK** to save your information and create this schedule.

At any time, you can return to the Project Details page for this project, select the Schedules subtab, and then select the Schedule Name from the table to go to the Edit Schedule page to modify this schedule.

Step 2a. Creating a continuous integration build trigger

At a later time when you have procedures and steps created to do the work you need, you may prefer implementing continuous integration build configurations for your software builds. Unlike a standard schedule

that runs on days and times you specify, a continuous integration build can be triggered to run each time code is checked into your source code management (SCM) system.

Adding a project to the Continuous Integration Dashboard is quick and easy. Each project can have one or more CI configurations, depending on the number of source branches you want to monitor. After the simple configuration process, you can visually see status and other information about your continuous integration jobs as builds are triggered to run.

To find the Continuous Integration Dashboard, select the Home tab, then select the Continuous Integration subtab. The dashboard will be empty until you start adding projects. For help using the Continuous Integration Manager and Dashboard, click the **Help** link in the upper-right corner of the dashboard web page.

Step 3. Creating a report

ElectricFlow provides multiple reports and custom report capabilities to help you manage your build environment.

Summary of available report types:

- Real-time reports filtered view of your ElectricFlow data in real-time
- Build reports summary reports produced at the end of a build and attached to the job
- Batch reports summaries of your build environment with trends over time, two types:
 - Default Batch reports automatically installed during ElectricFlow installation and scheduled to run daily (Cross Project Summary, Variant Trend, Daily Summary, Resource Summary, Resource Detail)
 - Optional Batch reports you can configure and schedule these reports to fit your requirements (Category, Procedure Usage, Count Over Time, or Multiple Series reports)
- Custom reports your choice to create and add at any time

Note: Batch and Custom reports must be run on the ElectricFlow server agent (local agent) only. These reports use the BIRT report engine, which is installed only on the ElectricFlow server.

Defining a "saved" search

Before you create a report, you need to define a "search" to focus the report on the information you want to see.

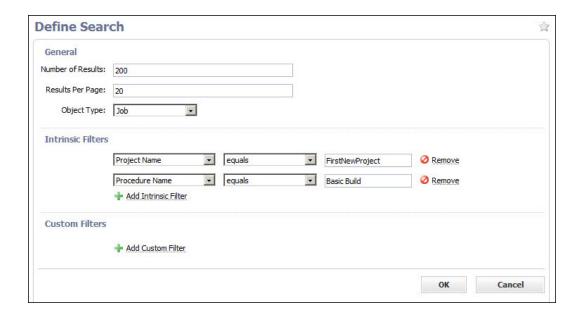
• Select the Search tab to go to the Define Search page.

Enter information in the fields as follows:

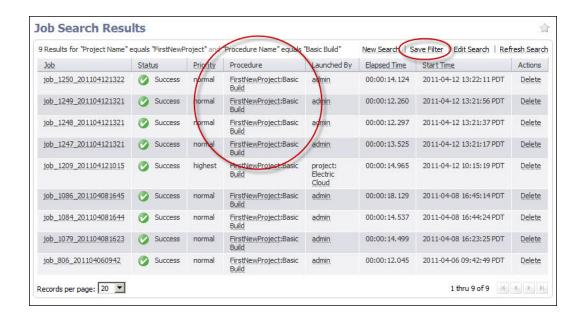
Field Name	Description
Number of results	For this example, do not alter the default settings in these fields as you see
Results Per Page	them. These fields do not affect the search you are creating.
Object Type	For this example, choose Job for this search.

Field Name	Description
Intrinsic Filters section	Click the Add Intrinsic Filter link one or more times as needed to further define the sort criteria for your search. We selected Project Name from the drop-down menu and typed-in our project name, then selected Procedure Name and typed-in our Basic Build procedure name.
Custom Filters section	Not needed for this scenario.

Your Define Search page will look similar to the following example:

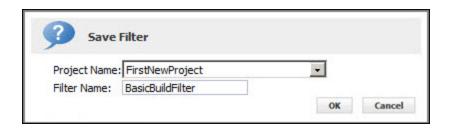


Click **OK** after filling-in the fields and to go to the Search Results page.



Click Save Filter to see the Save Filter dialog box.

In this dialog box, use the drop-down arrow to select the project name for your report, then enter a name of your choice for the filter. We used <code>BasicBuildFilter</code> for the filter name to the filter name to the Basic Build procedure for which we are creating a filter.

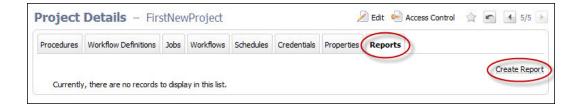


Click **OK** to return to the Job Search Results page—now we are ready to define the report we want to see.

Define the report

For this scenario, we will create an Optional Batch report. Select the Projects tab, select *your* project name to go to the Project Details page, then select the Reports subtab.

Click the Create Report link.



On the Optional Batch Reports page:

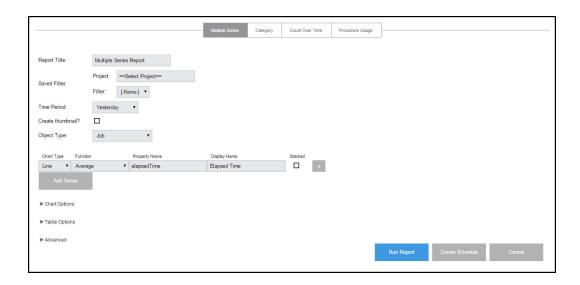
For this scenario, select the Multiple Series report tab.

Enter information in the fields or select the appropriate values as follows:

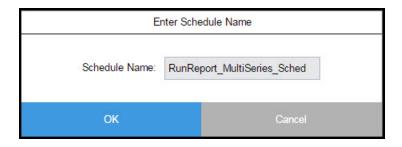
Field Name	Description	
Report Title	This is your report title. Type over the default report name, choosing any unique name for your report. We supplied Basic Build Report for our report name.	
Saved Filter	Project - Click your mouse inside this field to see a list of projects from which to make a selection. We select our FirstNewProject project name.	
	Filter - Use the BasicBuildFilter we created earlier.	
Time Period	Use the drop-down menu to select the time period.	
Create thumbnail?	Check this box so you can view this report on your Home page. (We will create a "thumbnail" report view on your Home page at the end of this scenario.)	
Object Type	Use the drop-down menu to select Job for this example.	
Table column choices		
Chart Type	Use the drop-down menu to choose the chart type.	
Function	Use the drop-down menu to choose the function you need.	
Property Name	Use the default property value or delete the text and click your mouse in the blank field to see a list of possible properties. We chose "outcome" for this example.	
Display Name	You can choose a different unique Display Name if you prefer to do so.	
Stacked	Select this check box to see your report results "stacked" versus overlaid.	
The "X" icon	Click this icon to delete any row you no longer need.	

Field Name	Description
Add Series button	Select this button if you would like to create additional table entries for additional report information.
Chart Options	Use the down-arrow to adjust the Time Grouping and see the defaults for the X and Y axis.

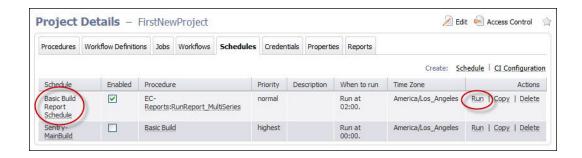
Your Multiple Series report page will be similar to the following example:



- Run Report button (information only for this scenario) this button takes you to the Job Details page to run the report immediately—one time only.
- Click the **Create Schedule** button this button displays a box with a default schedule name you can change.



- Enter a name for the schedule—again, we tied the schedule name to the procedure name for which we want the report.
- Click **OK** to go to the Project Details page.



Now you can see the report you just created listed in the Schedules table.

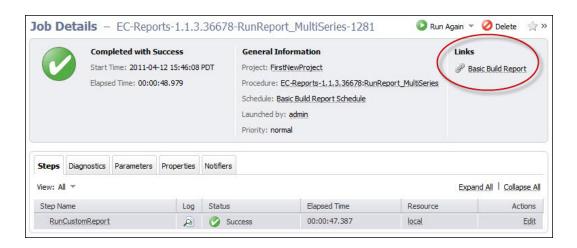
Click the Run link in the Actions column to run this report and then go to the Job Details page.

On the Job Details page, you will see your report running.

When the report is generated, you will see a Links section in the summary section at the top of the Job Details page.

This section contains a link to the report you generated.

Click this link to see your report.



For more information about the Job Details page, click here.

For more information about ElectricFlow reports, see the Reports, Creating Custom Reports, and other report Help topics.

Scenario extension - Add a thumbnail report to your Home page

When we created a report, we "checked" the Create thumbnail? checkbox, so we are ready to activate the thumbnail report view on the Home page.

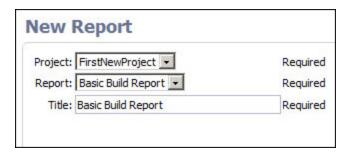
Note: Viewing an interesting report assumes there is ample data to report. While you can see a report after running one job, it is not very interesting unless you can see trends or comparisons after running numerous jobs.

To see a thumbnail report on the Home page:

- Select the Home tab.
- Click the Add Report link and a New Report page is displayed.



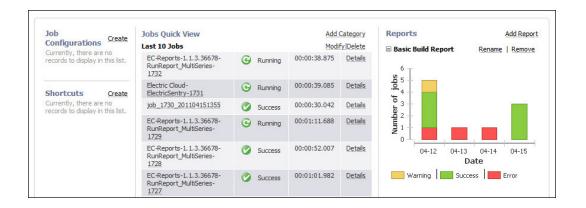
On the New Reports page:



- Project Use the down-arrow to select the project containing the report you want to view— FirstNewProject.
- Report This is the report name you entered in the Report Title field when you configured this report— Basic Build Report.
- Title This is the report title you want to see on your Home page to identify this report. We used the same Basic Build Report name.

Click **OK** to generate the thumbnail report view and to go to the Home page.

The following screen example illustrates how your Home page may look with the thumbnail report.



For more information about how to use the Home page, see The Home Page Help topic.

Summary

This scenario demonstrated how to create an Email Configuration so you could use an email notifier, how to create a continuous integration schedule, and how to set up a basic report with a "thumbnail" report view on your Home page.

Go to Scenario 4

Getting Started - Scenario 4 - Multi-agent Build and Test

Scenario requirements: For this scenario, you can use the configurations you created for Scenario 2 and 3. However, before you begin this scenario, you need to configure an additional agent machine. If you have a firewall running on the ElectricFlow server machine, disable it now or allow access to ports 8000, 8443, 61613, 80, and 443.

Installing ElectricFlow Agent software:

Install ElectricFlow Agent software on the machine you intend to use as an agent to run builds. For the simple purpose of getting started, install this agent on a different machine with the same operating system as the ElectricFlow server.

- Copy the ElectricFlow installation executable file to the machine you intend to use as an ElectricFlow agent.
- Double-click the executable file to begin the installation program.
- When the installation program opens, select the Express Agent installation.

Note: During the agent installation, you will have the opportunity to create a resource and provide a resource name. Enter any name you choose for your new resource. For this scenario, we will name our new resource, Resource2.

 After agent installation, verify that your new resource information is included in the ElectricFlow resource table. Select the Cloud tab to go to the Resources page.



Note: At a later time when you are more familiar with ElectricFlow, you can use the Resources page to install additional agents. See the Resources page Help topic, then specifically review the "Install or Upgrade Remotes Agents" section.

To configure a shared workspace:

For Linux

If your agents are installed on Linux, change the default workspace to point to a network location accessible to both machines:

- Select the Cloud tab, then select the Workspaces subtab, to see a list of workspaces.
- Select the workspace named "default" to edit the default workspace.
- Change the value of the UNIX Path to a network location accessible to the ElectricFlow server and agent machines.
- Make sure this location is readable and writable by the users chosen during the server and agent installations.

After fulfilling these scenario prerequisites, you are ready to proceed.

For Windows

If using a Windows environment, the shared workspace is created for you already.

Overview

This scenario guides you through creating a multi-agent build and test process integrated with your SCM system, build utility, and unit/system tests. You can invoke this build using a schedule, Continuous Integration (CI), or "on demand". At the end of this scenario you will be familiar with the following Electric Flow concepts and features:

- · Agent-only installation
- · Copying a procedure to create a new procedure
- · Running steps in parallel
- Pools

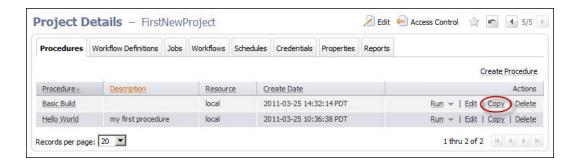
- · Post processor
- Search

Begin Scenario 4

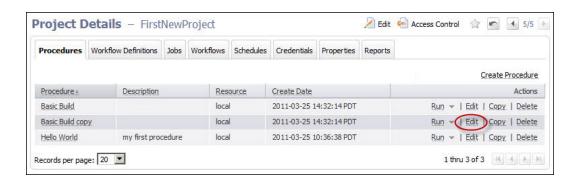
Step 1. Create a new procedure

This time, we are going to copy an existing procedure, rename it, and add another step.

- Select the Projects tab, then select the FirstNewProject project name.
- On the Project Details page, click the **Copy** link to copy the Basic Build procedure.



- When the Basic Build copy procedure is displayed, you can click its name to go to the Procedure Details page to verify it contains the same contents of the Basic Build procedure.
- Return to the Project Details page.



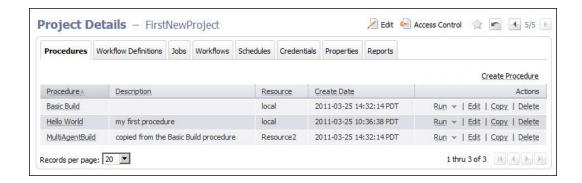
Select the Edit link for the Basic Build copy procedure.

On the Edit Procedure page:

- Name-Enter a new procedure name, MultiAgentBuild.
- Description—Add a description noting this procedure was copied from the Basic Build procedure.
- Default Resource–Select your new resource, Recource2, for this procedure.

Click **OK** after entering your information and ElectricFlow will take you back to the Project Details page.

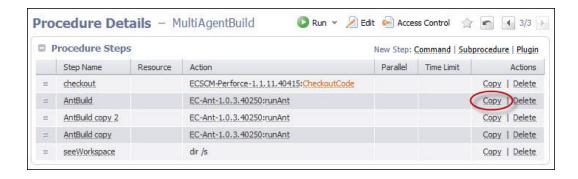
Your Project Details page should now look similar to the following example:



This page now displays the new MultiAgentBuild procedure with a description, and a new resource is designated for this procedure to use.

Step 2. Create two new Ant steps

- Beginning on the Project Details page, click the MultiAgentBuild procedure name to go to the Procedure Details page.
- Click the AntBuild step Copy link, twice.



- Now you see two additional steps names: AntBuild copy and AntBuild copy 2.
- Click the AntBuild copy step name to go to the Edit Step page.
- Notice the AntBuild copy step name is adjacent to the page title because this is the step we are going to edit.

Modifications for creating the new Ant steps:

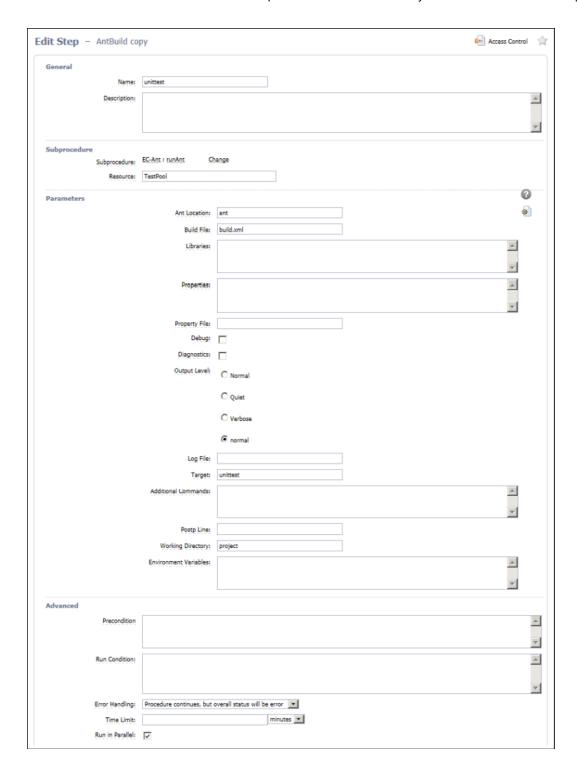
Because this is an "Edit" page, some of the fields already contain the information we originally supplied to create the first Ant step. The following list contains the fields we need to edit or add new information:

- Name—For this example, we chose unittest for the step name. Type over the existing text to edit the text.
- Resource–Enter the name, TestPool, because we will be creating a resource pool for testing. At this point, you need to enter a pool name only (no spaces in the name).
- Target-Edit this field to: unittest

In the Advanced section:

• Run in Parallel–Click the check box to select this option. You are designating that you want this step to run in parallel with one or more other steps.

Click **OK** to create the first of two new Ant steps and ElectricFlow returns you to the Procedure Details page.

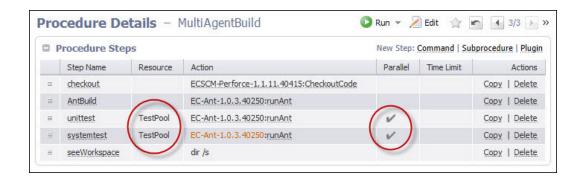


From the Procedure Details page, click the AntBuild copy two step name to go to the Edit Step page again.

To create the next Ant step, make the same edits (above) as you did to create the unittest step, except:

- Name-Change the step name to systemtest.
- target-Change this field to: systemtest

Click **OK** to create the second of two new Ant steps and to return to the Procedure Details page. You now have three Ant steps, including two new Ant steps, unittest and systemtest, with TestPool as the designated resource.



Also notice, the Parallel column now contains "checkmarks" for the 2 steps we want to run in parallel.

Step 3. Create a resource pool

Select the Cloud tab to see the following Resources page with a table listing the currently available resources.

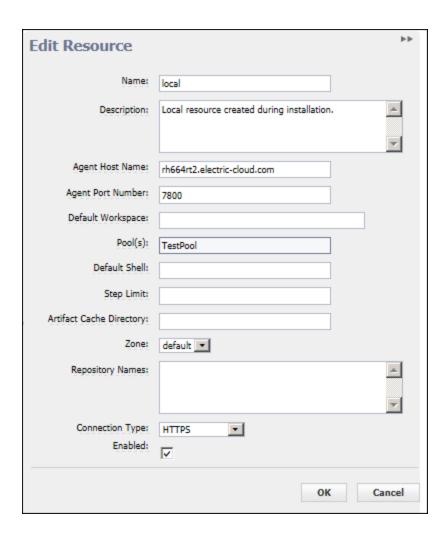


To put these two resources in a resource pool that we previously "named" when we created the new Ant steps, we need to edit each resource.

Click the local resource name to go to the Resource Details panel, then click the **Edit** link at the top of the panel.

On the Edit Resource panel:

The only change we need to make is to add the TestPool pool name to the Pool(s) field.

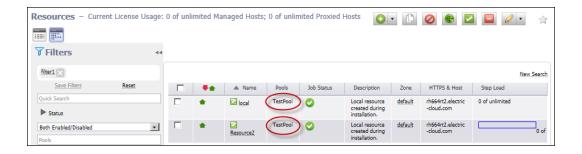


Click **OK** to save your change and return to the Resources page.

- Notice the Pool(s) column now has TestPool specified as a resource pool where the local resource is a member.
- Repeat this process to put Resource2 in the TestPool resource pool.

A resource can belong to one or more pools. If any resource belongs to multiple pools, all pools would be listed in this column. For more information about the Resources page, click here or click the **Help** link in the top-right corner of the web page.

The following screen example should be similar to your Resources page.



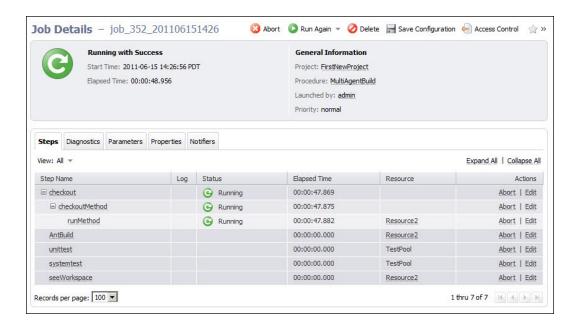
Step 4. Run the procedure

To run the new MultiAgentBuild procedure, select the Projects tab > FirstNewProject project name > MultiAgentBuild procedure.

On the Procedure Details page:

- Click the Run link at the top of the page.
- · Click Run on the Run Procedure page.

Your Job Details page should be similar to the following example.



Note: The screen example above illustrates a job that is still running—notice the green circular-arrow icon at the top of the page. This icon changes to a "checkmark" if this job completes successfully.

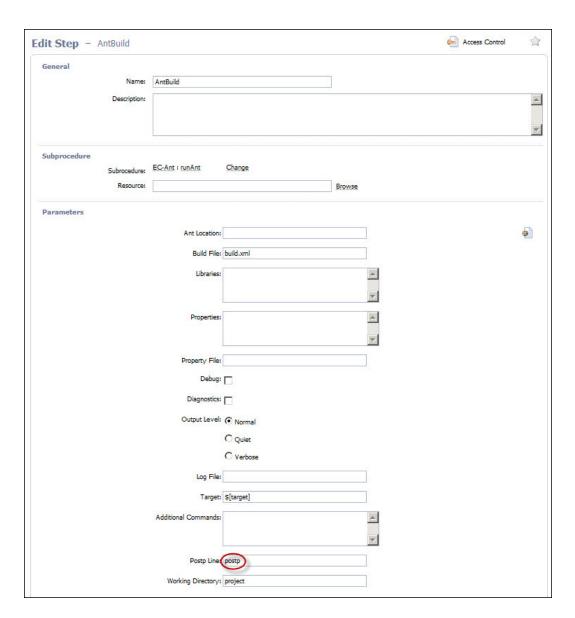
Scenario extension - Using postp

ElectricFlow implements data collection with a postprocessor. The postprocessor is a command associated with a particular procedure step. If the postprocessor is specified for a step, it executes concurrently with the main step command. The postprocessor runs on the same machine as the main command and in the same working directory, and it retrieves the log file from the step as its standard input.

The standard ElectricFlow postprocessor is called **postp**. **postp** scans the step's log file looking for interesting output such as error messages and then sets properties on the job step to describe what it found. For example, postp might create a property named "errors" whose value is the number of error messages in the log file, or a property named "tests" that counts the number of tests the step executes. Also, postp can extract portions of the step log that contain useful diagnostic information and save this information for reporting.

When you create a step, you can specify postp. If you already have a step and want to add postp reporting, you just have to edit the step to include postp.

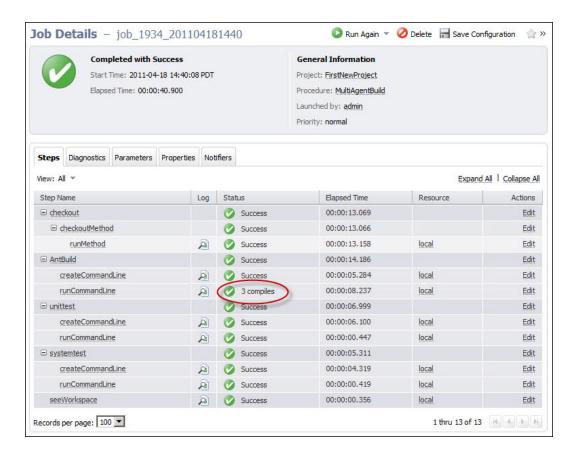
- From the previous Job Details page screen example, click the Edit link for the AntBuild step.
- On the Edit Step AntBuild page, in the Parameters section, add postp to Postp Line field.



Click **OK** to save your edit and return to the Job Details page.

Run this job again to see the postp result—click the ${\bf Run/Run\ Again\ link}.$

On the next screen example, notice the AntBuild step now shows "3 compiles" in the Status column. Using postp supplies additional information.



For more information about using postp, see the "Postprocessors" Help topic.

Summary

This scenario demonstrated how ElectricFlow can drive a complete build and test process—including building steps to run in parallel and creating a resource pool. This scenario also provided a brief look at how you can use ElectricFlow to execute rapid root cause analysis by reviewing postp diagnostics.

Automation Platform Setup

The topics in this section describe how to set up the automation platform UI.

After completing the tasks, you can create, configure and manage the objects needed to automate your build-test processes, application deployment, and pipelines.

Automation Platform Home Page

Job Configurations

Shortcuts

Jobs Quick View

Reports

Overview

This page provides a convenient console for running jobs and viewing results.

- This web page is your ElectricFlow Dashboard for tracking project health.
- You can customize the page to display only the jobs that interest you and to provide one-click access to the jobs you run frequently. Each of the Home page sections can be customized.
- The Reports section allows you to see thumbnail report images that are updated each time the report is generated.
- This page provides shortcuts for quick access to pages you use most frequently inthe automation platform or anywhere on the web.

IMPORTANT: If you are using or plan to use the EC-Homepage plugin to share your Home page sections, you will not see some links normally available to you. For example, if a particular Jobs Quick View category is shared, you **cannot** Delete or Modify that category. You can perform these functions only if the category belongs to you alone.

For information on sharing your Home page configuration, see the "Home page configurations" section in the "Customizing the ElectricFlow Platform UI on page 444" Help topic.

Job Configurations

Procedures in ElectricFlow can contain complex sets of parameters— making it difficult to remember the correct parameters for a particular situation and tedious to re-enter those parameters every time the procedure is invoked. Job Configurations provide a one-click solution to this problem. When you create a job configuration, you enter all the information needed to run a procedure, including parameters and/or a credential. All job configurations are displayed here on the Home page. Invoke a particular configuration by clicking its name in the Job Configurations section.

Create Job Configurations three ways

- On the Home page, create a job configuration from "scratch" by clicking the **Create** link in the Job Configurations section.
 - In the Create Configuration pop-up menu, select the project and procedure you want to use for creating this configuration.
- From the Job Details page for a previously invoked job, click the **Save Configuration** link at the top of the page.
 - Your saved job configuration is displayed on your Home page.
- From the Edit Schedule page, click the **Save Configuration** link at the top of the page. Your saved configuration is displayed on your Home page.

Shortcuts

Use shortcuts to save frequently visited ElectricFlow web pages, so those pages are immediately accessible. You can create a shortcut to any page on the web also.

Create Shortcuts two ways

- Mouse-over the "star" icon at the top of any automation platform page and click "Add current page" to
 add that page to the shortcut list. Mouse-over the star icon again and click "Remove current page" to
 remove the shortcut for that page. The star icon is yellow for pages saved as a shortcut and gray for
 those pages not saved as a shortcut.
- Click the Create link in the Shortcut section and provide a name and URL to create a shortcut.

To modify or update a shortcut, click the Edit link adjacent to the shortcut you want to change.

Shortcuts can be accessed conveniently from any automation platform web page. Mousing-over the star icon displays a list of shortcuts saved by the current user. Click on a shortcut name to view the page.

Note: The Shortcut section may contain static entries that cannot be deleted. And if you "share" your Home page, the Edit link will not be available for shared items.

Jobs Quick View

Perhaps only a few jobs on this server are of interest to you. For example, you may care about jobs you launch manually and official builds for the products you work on, but you may not care about production builds for other products or personal jobs for other users.

The Jobs Quick View allows you to define job categories that are interesting to you.

The Home page displays the most recent jobs in each category, and you can easily click-through to get more details about any of those jobs. For example, clicking on a job name takes you to that job's Job Details page.

Create a job category

- Click the Add Category link in the Jobs Quick View section.
 After creating a category, results are displayed on the Home page. Clicking the Details link displays a summary to the right of the category. In addition to job status and other diagnostic information, the summary displays running steps and failed steps (containing errors and warnings).
- Click the **Modify** link to edit the Jobs Quick View category or the **Delete** link to remove that job from the Jobs Quick View category.
- Click on the **Details** link to see job summary information—the summary remains visible, regardless of mouse location, until you click somewhere else on the page.

Note: If you "share" your Home page, the Modify and Delete links will *not* be available for shared items.

Reports

You can configure reports you would like to see on a regular basis and display a "thumbnail" report graphic in this section.

• Click the **Add Report** link to go to the New Reports page.

After filling-in the information on the New Reports page, you will see a thumbnail view of your report

(after it runs) on your Home page.

Note: If the drop-down menu on this page is empty, click the **Help** link on the New Report page for more information.

- Click the Collapse/Expand (+/-) box to see the full thumbnail report or just the report title.
- Click the Rename link if you need to rename your report.
- Click the **Remove** link if you need to remove this report from your Home page.

Note: If you "share" your Home page, the Rename and Remove links will not be available for shared items.

Customizing the ElectricFlow Platform UI

You can customize the ElectricFlow platform UI to get a more intuitive, task-specific user interface.

- Customizing parameters
- Customizing tab layouts
- Home page configurations

Customizing parameters

Two types of objects contain formal parameters – procedures and state definitions. For these objects, you can customize the way their formal parameters are presented. You can reorder the parameters and set the labels, form element types, default values, tooltip Help text, and whether or not the parameter is required. When defined, this customization shows up whenever the parameters are displayed—running a procedure, creating a step to call a procedure, creating a transition definition to a specific target state, and so on.

How do you customize parameters?

Create a property named ec_parameterForm on the object containing formal parameters (the procedure or state definition). The value of this property is an XML-style specification of form elements that map to the parameters. This property must be in sync with the formal parameters, which means all formal parameters must have a corresponding XML element, with no "extra" XML elements.

Custom parameter form contents

The value of the property in XML: Under a top-level element called <editor>, you need an element called <formElement> for each parameter. Under a <formElement>, you can specify the following tags:

- property* the name of the formal parameter
- $\bullet \quad type^{\star} \text{-} \text{ entry|textarea|select|radio|checkbox|project|savedSearch|credential}$
- label* displayed next to the element
- value the initial value of the element (does not apply to credentials, which must be specified at runtime)
- required if true, the user must enter a value for the element
- documentation text displayed in the tooltip when the form element is "moused" over
- numRows valid for textarea elements; represents the height of the element

- option a single option for select | radio elements; at least one is required
 - o name* the text displayed in the option
 - o value the value of the option
- checkedValue valid for checkbox elements; the value of the element when it is checked
- uncheckedValue valid for checkbox elements; the value of the element when it is unchecked
- initiallyChecked valid for checkbox elements; whether or not the element is "checked" by default

Note: An asterik (*) in the list above indicates a required tag.

Example

The following example is a parameter form that rearranges parameters and sets the labels, descriptions, default values, and form elements for each parameter.

```
<editor>
   <formElement>
       <label>One:</label>
       property>one
       <documentation>The first parameter.</documentation>
       <type>entry</type>
       <value>Test value
   </formElement>
   <formElement>
       <label>Two:</label>
       property>two
       <documentation>The second parameter.</documentation>
       <type>textarea</type>
       <required>1</required>
   </formElement>
   <formElement>
       <label>Three:</label>
       property>three
       <documentation>The third parameter.</documentation>
       <type>select</type>
       <option>
           <name>ABC</name>
           <value>abc</value>
       </option>
       <option>
           <name>XYZ</name>
           <value>xyz</value>
       </option>
       <value>xyz</value>
   </formElement>
   <formElement>
       <label>Four:</label>
       property>four
       <documentation>The fourth parameter.</documentation>
       <type>radio</type>
       <option>
           <name>First Option</name>
           <value>123</value>
       </option>
```

```
<option>
           <name>Second Option</name>
           <value>456</value>
       </option>
       <value>123</value>
   </formElement>
   <formElement>
       <label>Five:</label>
       property>five
       <documentation>The fifth parameter.</documentation>
       <type>checkbox</type>
       <checkedValue>true</checkedValue>
       <uncheckedValue>false</uncheckedValue>
       <initiallyChecked>1</initiallyChecked>
       <value>true</value>
   </formElement>
   <formElement>
       <label>Six:</label>
       property>six
       <documentation>The sixth parameter.</documentation>
       <type>project</type>
       <value>myProject</value>
   </formElement>
   <formElement>
       <label>Seven:</label>
       property>seven
       <documentation>The seventh parameter.</documentation>
       <type>savedSearch</type>
       <value>/projects/myProject/ec savedSearches/mySavedSearch</value>
   </formElement>
   <formElement>
       <label>Eight:</label>
       cproperty>eight
       <documentation>The eighth parameter.</documentation>
       <type>credential</type>
   </formElement>
</editor>
```

Customizing the tab layout

Overview

A "view" defines the layout of tabs in the ElectricFlow web UI. One or more tab views may be defined at the server, group, and user level. The default view the user sees when they first log in can be set at the server, group, and user level. The default defined on the user takes precedence over its groups and the groups take precedence over the server. Electric Cloud provides a system default view that will always show up in the user's list. Views can inherit from each other and add/modify/remove/reposition tabs and subtabs as needed.

If you define a custom view, you have to manually add tabs such as Continuous Integration so that they will appear in your custom view.

View definition syntax

A view is an XML property. The following is a list of elements that can be defined in a view:

- tab a top level tab
 - o label the text to display for the tab
 - o url the target URL of the tab
 - accesskey the keyboard shortcut to access this tab (a single letter)
 - o position the position of this tab relative to the other tabs (first tab is 1, second tab is 2, and so on)
 - show when inheriting from another base view, if 0, the tab is not visible, if 1 it is visible (default 1)
 - o tab a subtab displayed within this tab
 - label the text to display for the subtab
 - url the target URL of the subtab (relative to ElectricFlow base)
 - position the position of this subtab relative to the other subtabs (integer greater than 1)
 - show when inheriting from another base view, if "0", the tab is not visible, if "1" it is visible (default 1)
- base the name of a view from which to inherit tab and subtab definitions

The following is a basic definition of a two-tab view where one of the tabs has three subtabs.

```
<?xml version="1.0" encoding="utf-8"?>
  <view>
     <tab>
       <label>Home</label>
       <url>home.php</url>
       <accesskey>h</accesskey>
     </tab>
     <tab>
       <label>Administration</label>
       <url>workspaces.php</url>
       <accesskey>a</accesskey>
       <tab>
          <label>Workspaces
          <url>workspaces.php</url>
       </tab>
       <tab>
          <label>Directory Providers</label>
          <url>directoryProviders.php</url>
       </tab>
       <tab>
          <label>Licenses
          <url>licenses.php</url>
```

```
</tab>
</tab>
</view>
```

The following is a view that inherits from the above view and:

- 1. Changes the URL of the Home tab
- 2. Adds a new tab called Projects and places it at the beginning of the list
- 3. Changes the accesskey of the Administration tab
- 4. Moves the Workspaces subtab to the end of the list
- 5. Hides the Licenses subtab

```
<?xml version="1.0" encoding="utf-8"?>
<view>
   <base>firstView</pase>
   <tab>
     <label>Home</label>
     <url>customizedHome.php</url>
   </tab>
   <tab>
      <label>Projects</label>
      <url>projects.php</url>
     <accesskey>p</accesskey>
      <position>1</position>
   </tab>
   <tab>
      <label>Administration</label>
      <accesskey>z</accesskey>
     <tab>
        <label>Workspaces</label>
        <position>3</position>
      </tab>
      <tab>
         <label>Licenses
         <show>0</show>
      </tab>
   </tab>
</view>
```

Storing views

Tab views are stored in the server, group, and user property sheets. The view definitions are stored in a property sheet called ec_ui/availableViews. The name of the view is set to the property's description if defined, otherwise it is set to the property's name. The value of the view property is the XML definition document described above.

Default views

The property ec ui/defaultView, if set, determines the default view for users inheriting from that object.

- If this property is set on the server, it is the default for all users.
- If this property is set on a group, all users belonging to that group will see that view, overriding the server's default if it is set also.
- Finally, a user can explicitly set their view by defining the property on their own property sheet. If the user belongs to multiple groups that define defaults, then ElectricFlow chooses the first group alphabetically.

The user can set their view by clicking on their name in the navigation bar, then clicking on the **Edit Settings** link. The list of views shown here comprise all views defined in the ec_ui/availableViews property sheets for the server, groups to which the user belongs, and the user itself.

If views have the same name at different levels, then the user overrides the group which overrides the server. If a user selects a view and chooses to Save, then the ec_ui/defaultView property on their sheet is set accordingly.

Two special values are always available in this list: "EC Default" and "Inherit". The former is the default tab layout defined in the ElectricFlow web UI. The latter means to use the first view in the list. By default, all users inherit their tab view.

For more information on user settings, see the Edit User Settings Help topic.

Developing and troubleshooting

If you create an invalid view definition on a group or server property sheet, you will break the UI for all users who inherit that view. The best practice for developing tab views is to use a "test user". Store the view definition in the ec_ui/availableViews property sheet for that user, and set the ec_ui/defaultView for that user to the name of the new view. When the view is working properly and ready for other users, you can migrate the new view to a group or to the server.

If you continue to have difficulty and tab definitions are not working correctly and can no longer get to the user settings page, revert to the ElectricFlow default by running the following command:

```
ectool setProperty /users/$USERNAME/ec_ui/defaultView ec_default
```

When you log out and then log back in, you will have the default view again.

Note that changes to views are not visible until the next time a user logs in.

Home page configuration

The automation platform Home page is your configurable Dashboard. This page allows you to manage shared configurations and you can customize this page for your work preferences also. Refer to The Home page Help topic for details on creating Job Configurations, Shortcuts, Jobs Quick View, and Reports.

Shared Home page configurations can be set globally or for specific groups of users only. The general model for creating shared configurations is:

- Set up your personal configuration the way you want it to appear to others.
 To get to the configuration page, select the Administration tab > Plugins > and click the Configure link for EC-Homepage.
- 2. Click **Save** to make the configuration available to a specific group or globally to everyone on the server.
 - If you published the configuration to a specific group, no further action is needed. If you published the configuration globally, continue to the next step.
- 3. Copy and paste the Tab XML string into a specific user's or group's view.

This adds the public configuration to that user's or group's own home page and no further action is required for the user or group. If this step is not completed, additional action is required as described in Installing the Home page on page 450.

You can alter an existing shared configuration by loading from a previously saved location, make changes, then save the changes back to the original or an alternate location.

User settings

Use the Backup Settings action to save and restore your personal Home page configuration.

Select **Create** to back up and temporarily set aside your personal settings while you create or update shared settings.

Select **Restore** to retrieve your personal settings after updating shared settings.

Location

Select Global to share the new Home page with everyone on the server, or select Group to share with specific users only.

Installing the Home page

If a globally available Tab XML string (from the Configure EC-Homepage page) has **not** been pasted to the user's or group's view, it must be added to an existing view definition to replace the standard Home page with the one provided by this plugin. To do this, click **Load**. This replaces the current home page with the one that's publicly available.

Reconfigure Contact Support link

In the event you want to redirect the Electric Cloud Support URL (available from the Help link on each ElectricFlow web page) to your own support center, you can create a file to add to the ElectricFlow installed directory at <ElectricFlow Install Dir>/apache/htdocs/commander. Create a file named config_user.php with the following contents:

```
<?php
$config["contactSupportUrl"] = "your URL";
?>
```

Notes:

- 1. You must enter the value (URL) for your customer support site, replacing "your URL" in the example above. If no value is supplied, a blank window is displayed.
- 2. You must restart the Web Server for your change to take effect.

Configuring ElectricFlow

Use this page for quick access to configuring resources, workspaces, email configurations, your source control system, and defect tracking to communicate with ElectricFlow.

Note: If you are viewing this Help topic from the "Configuring ElectricFlow" web page, use these instructions:

Click the Add link for any object you want to configure and that objects configuration page is displayed.

- Each item you configure will be listed in this table for your easy reference.
- Configure objects listed in this table are linked to their respective "Edit..." page in the event you need to modify any values previously supplied. Select a configured object to edit its values.

The table below provides a brief description of ElectricFlow configuration tasks

Click on one of the **"Setting Up..."** links to see more information about that topic.

Setting Up Resources	Before resources can be set up, you need to know which agent machines are available to allocate to jobs. The ElectricFlow Resource web page displays all resources currently available to the ElectricFlow server.
Setting Up Workspaces	ElectricFlow provides each job step with an area on the disk it can use for "working files" and results. This disk area is called a job workspace.
	A job step can create whatever files it needs within its workspace and ElectricFlow automatically places files, such as step logs, in the workspace. The location of the job step workspace is displayed on the Job Details web page.
Setting Up Email Configurations	If you do not create an Email Configuration, you will be unable to send email notifications to individuals or groups.
	If you have multiple users or groups in remote locations who use a different mail server, you will want to create additional email configurations to accommodate those locations if they need to receive notifications.

Setting Up a Source Control Configuration	ElectricFlow needs to communicate with your Source Control system (SCM) if you intend to use the CI Dashboard for continuous integration schedules to start a new build based on new/modified source control checkins, or if you plan to configure Preflight builds (to build and test code changes before those changes are "committed").
Setting Up Directory Providers	ElectricFlow uses account information from multiple sources. In most cases, the primary account information source is an external LDAP or Active Directory repository: both user and group information is retrieved from the repository. Local users and groups can be defined within ElectricFlow.
Setting Up Defect Tracking	ElectricFlow uses plugins to integrate with numerous defect tracking systems. If the plugin you need for your defect tracking system is not automatically installed with ElectricFlow, see the Plugin Catalog to find the integration you need.

Web Interface Online Help System

Use the automation platform online Help system for more information. Click the **Help** link in the top-right corner of any product web page to see a specific Help topic for that page.

When the Help system opens, we recommend reviewing the Help table of contents. All Help folders above the Web Interface Help folder are user-guide style Help topics that provide more detailed information on each of their topics. If you generally prefer to use a command line rather than the ElectricFlow web interface, you will find complete ectool (the ElectricFlow command-line tool) commands and arguments here too.

Setting Up Resources

Select **Cloud > Resources** to view the Resources page.

Before resources can be configured, you need to know which agent machines are available to allocate to jobs. The Resources page displays all resources currently available to the ElectricFlow server. This page is your Resources management center where you can create, modify, or delete resources and see the resource status.

To view the Resource Pools page where you can manage the resource pools in ElectricFlow, go to **Cloud** > **Pools**.

To view the Zones page where you can manage the zones in ElectricFlow, go to **Cloud > Zones**.

What is a Resource?

A resource is defined as an agent machine where steps can execute. In addition:

- Each resource has a logical name.
- Each resource refers to an agent machine by its host name.

- Multiple resources can be defined on the same physical host.
- Each resource can be assigned to one or more pools, or one zone.
- Each resource has a step limit that determines the maximum number of steps that can execute simultaneously on the resource.

In Create/Edit Resource panels, you define the number of steps that can run on a resource. Remember that setting the Step Limit to "0" or leaving the field blank defaults to no step limit.

When ElectricFlow allocates resources for a job and steps are executed on the resource, it considers the step limit, not the CPU load from the steps. The step limit for a resource in your system depends on the agent machine performance. For example, using *unlimited* as the step limit for a heavy load can slow the agent and result in less consistent agent performance if this is not managed properly.

What is a Pool?

A pool is a group of interchangeable resources. For example, you might have a pool of Windows servers.

Naming a pool in a procedure step allows ElectricFlow to assign any resource from that pool to do the work for that step. By default, ElectricFlow will try to balance the requests against the set of resources in the pool, but you can also define your own rules for how resources will be selected.

To define a resource for a procedure step, you can do the following:

- Specify a resource or a pool name in a procedure step to execute that step on that resource/pool.
- Define custom properties for your resources or resource pools.
- Mark a resource as "exclusive" (the step acquires and retains its resource exclusively)...

For example, if configuration information varies from resource to resource, you can use custom properties to hold this information and then reference it from procedure steps where it is needed. Suppose you have a pool of test machines where the test hardware is in a different location on each machine. You could define a property named "testLocation" on each resource, then pass this property to procedure steps using a reference such as \$[/myResource/testLocation]. This approach allows you to define procedure steps to run on any resource and automatically handle configuration differences.

Note: When the ElectricFlow server and agents are installed on different hosts, make sure that the configuration for each agent specifies the Domain Name System (DNS) server.

ElectricFlow Supports Two Resource Categories

- Standard
- Proxy

Standard Resources

This resource category specifies a machine running an installed ElectricFlow agent on one of the supported agent platforms, as specified in the *ElectricFlow Installation Guide*, Chapter 2, "System Requirements and Supported Platforms."

Proxy Resources

This resource category requires SSH keys for authentication. You can create proxy resources (agents and targets) for ElectricFlow to use on numerous other remote platforms/hosts that exist in your environment. A proxy agent is an ElectricFlow agent, channeling to a proxy target.

- **Proxy agent** This is an agent on a supported Linux or Windows platform, used to proxy commands to an otherwise unsupported platform. A proxy agent is an ElectricFlow agent, channeling to a proxy target.
- Proxy target This is a machine on an unsupported platform that can run commands via an SSH server.
 Proxy targets have limitations, such as the inability to work with plugins or communicate with ectool commands.

When a step is running on a proxy resource, the proxy agent performs the following tasks:

- Uploads the commandFile to the Working Directory on the proxy target via SSH.
 - Working Directory the step runs in this directory on the proxy target, which defaults to the UNIX path to the workspace
- Creates a wrapper sh shell script to:
 - CD to the Working Directory
 - Set COMMANDER environment variables that exist in the proxy agent's environment
 - Run the commandFile that was uploaded earlier
- Uploads the wrapper script to the Working Directory on the proxy target
- Runs the wrapper script on the proxy target
- Cleans up script files on the agent and proxy target

Setting up SSH keys

If you are not familiar with setting up SSH keys, go to the knowledge-based article called "Setting up SSH Keys in preparation to use a Proxy Agent" in the Electric Cloud Support Web Site:

- 1. Go to https://electriccloud.zendesk.com/hc/en-us/categories/200176553-ElectricCommander.
- 2. In the ElectricFlow KB section, enter **KBEC-00049** in the Search field.
- 3. Click Search.
- 4. Click the link to the knowledge-based article called "KBEC-00049 Setting up SSH Keys in preparation to use a Proxy Agent" to open it.

On UNIX, the ElectricFlow agent looks in the ~/.ssh directory for the private and public key files to set up the SSH connection to the proxy target. Therefore, we recommend generating your key files in that directory.

On Windows, the ElectricFlow agent does not presume a default location. Therefore, in the proxy resource definition you must specify the key file location. Enter the following line in the resource "Proxy Customizations" field:

```
setSSHKeyFiles('c:\ foo\priv.key');
```

To create a new resource

Click the Cloud tab in the ElectricFlow web interface to go to the Resources page.

At the top right-side of the table, click the down arrow next to the plus sign and select **Create Resource** or **Create Proxy Resource** link to see the New Resource or New Proxy Resource panel.

Setting Up Workspaces

ElectricFlow provides each job step with an area on the disk it can use for "working files" and results. This disk area is called a job workspace.

A job step can create whatever files it needs within its workspace, and ElectricFlow automatically places files in the workspace, such as step logs. Normally, a single workspace is shared by all steps in a job, but it is possible for different steps within a job to use different workspaces. The location of the job step workspace is displayed on the Job Details page for the job under "Details" for the step.

To create a new workspace

ElectricFlow can handle any number of workspaces.

- To define a workspace from the web interface, click the Cloud > Workspaces tabs.
- On the Workspaces page, click the Create Workspace link at the top right-side of the table and the New Workspace web page is displayed.

Click the **Help** link in the top-right corner of the web page for assistance with filling in the fields to create a workspace.

To see the Workspaces Help topic now, click Workspace - create new or edit existing workspace.

After you create a workspace specification, you will see it listed on the Workspaces page.

Setting Up Email Configurations

This topic describes how to set up an email configuration. It must be created and configured before you set up email notifications, also referred to as *email notifiers* at the platform level. Then you can send them to individuals or groups.

If you have multiple users or groups in remote locations who use a different mail server, you need to create additional email configurations to accommodate those locations if they receive ElectricFlow notifications.

Creating an Email Configuration

To create an email configuration in the platform GUI:

- 1. Click Administration > Email Configurations to go to the Email Configurations page.
- 2. Click **Add Configuration** to go to the New Email Configuration page.

3. Enter information in the fields on the page.

There are two ways to get help in entering information in the fields:

- Click Email Configuration create new or edit existing email configuration to go to the Email Configuration help topic.
- If you are in the platform GUI, click the **Help** link in the upper-right corner of the page to see descriptions and other information.

After the email configuration is created, you will see it listed in the table on the Email Configurations page.

Setting Up a Source Control Configuration

You must configure your source control system to communicate with the ElectricFlow server if you plan to takes advantage of the ElectricFlow Continuous Integration Dashboard and associated functionality and Preflight functionality—both of these features are designed to work with a number of source control (SCM) systems.

- Continuous Integration Dashboard use this feature to create continuous integration schedules for your build environment. The Continuous Integration Manager is the front-end user interface for the ElectricSentry Continuous Integration engine.
- Preflight use this feature to build and test developer code before it is committed to the code base for your product.

ElectricFlow installs (bundles) and supports numerous source control types. After creating a source control configuration, your entry will appear in the table on the Source Control Configurations web page—to see this web page, select the Administration > Source Control tabs.

Note: If the SCM you prefer is not listed here, see the Plugin Manager page (Administration > Plugins) for a list of all currently available SCM plugin integrations available for ElectricFlow. To configure a different SCM, see the Help topic associated with that plugin.

Select one of the following links to go to the source control configuration page to configure your SCM system:

AccuRev

ClearCase

File

Git

Perforce

Property

Subversion

The Continuous Integration Dashboard has its own Help topic. For more information on ElectricSentry or Preflights builds, see their respective Help topics: ElectricSentry, Preflight Builds.

Setting Up Directory Providers

ElectricFlow uses account information from multiple sources. In most cases, the primary account information source is an external LDAP or Active Directory repository: both user and group information is retrieved from the repository. Local users and groups can be defined within ElectricFlow.

To specify a directory provider

ElectricFlow includes a web page to facilitate adding your existing LDAP or Active Directory users and groups to ElectricFlow.

- To access the Directory Providers web page, select the Administration > Directory Providers tabs.
- Click the Add Active Directory Provider or the Add LDAP Provider link to add a new provider.

Depending on which **Add... Provider** link you choose, you will see either the New Active Directory Provider or the New LDAP Provider web page. From either page, click the **Help** link in the upper-right corner of the page. The Help topic describes each field on either web page so you can easily enter information in the forms to have ElectricFlow use your existing user or group account information.

To see the Directory Provider Help topic now, click Directory Providers - create new or edit existing directory providers.

After adding a provider, you will see it listed in the table on the Directory Providers web page.

Enable/Disable Local ElectricFlow Users

Two server settings properties are supplied for local ElectricFlow users (both are "on" by default):

- enableAdminUser if set to '0', attempts to log in as the 'admin' user are denied as if the user did not
 exist
- enableLocalUsers if set to '0', attempts by a non-admin local user to log in are denied as if the user did not exist

These settings are stored in properties in the '/server/settings' property sheet and can be set from ectool by calling:

```
ectool setProperty /server/settings/enableAdminUser 0
```

Or, these settings can be set from the automation platform web interface:

- Select the Administration > Server tabs.
- At the top of the page, select the **Settings** link.

Note: Disabling these properties takes effect only for new sessions. Existing sessions continue to function "asis" until the user logs out and attempts to log in again.

wrapper.conf Properties

If you need to edit the wrapper.conf file, see http://wrapper.tanukisoftware.com/doc/english/properties.html for complete information about property formats.

Using ElectricFlow in Your Environment

This topic addresses some common issues you may encounter as you start using ElectricFlow and offers some tips about managing your ElectricFlow projects.

What's in a step?

Where's the script for the step?

Environment variables

When do you use subprocedures?

How do you evolve procedures?

Porting existing scripts

ElectricFlow project version control

What's in a step?

One of the first questions you may ask when you start to implement your build and test processes on ElectricFlow is how to divide processes into individual steps: "Should I use only a few steps, each of which does a lot, or a large number of fine-grain steps?" Here are some factors to help you decide how many steps to use:

- Reporting If you would like a separate report of success/failure for two activities, put those activities in separate steps. If you are happy to have a single report for both activities, a single step may make sense. For example, compilation and test phases should probably be in different steps because errors in the two phases will probably be handled differently. Unit tests for product components managed by different groups may make sense in separate steps; each group can watch for errors in its step.
- **Parallelism** If you want to use ElectricFlow to run two activities in parallel, put them in separate steps. If these activities are in the same step, ElectricFlow cannot run them in parallel.
 - ElectricFlow parallelism works best at a coarse grain, such as running different sets of unit tests in parallel or compiling for different platforms. Electric Cloud does not recommend trying to do fine-grain parallelism with ElectricFlow, such as compiling every individual source file in a separate step— this is likely to be complicated and brittle, resulting in an enormous number of job steps, which will make it difficult to view results. If you would like to use fine-grain parallelism for compilation, we recommend using our ElectricAccelerator product. In this case, you would make your compilation steps large, with as much work in each makefile as possible because ElectricAccelerator automatically subdivides the work and runs as many sub-steps as possible in parallel. The more work in a makefile, the more efficient the process becomes.
- **Resources** If two different activities need to run on different resources, they need to be in different steps. A single step runs entirely on a single resource.

- Conditional steps If you want to skip portions of your process during some jobs, but execute them during others, it probably makes sense to put those activities in separate steps and use the ElectricFlow "Run condition" mechanism to decide whether they run during each particular job. However, if you are invoking programs like *make* that already allow you to choose which actions to perform, it may make more sense to have a single large step and handle conditional behavior with those programs, rather than with ElectricFlow. For example, when running *make*, you can specify the targets to be built.
- **Setup** You might want to have a single step at the beginning of a procedure that processes the procedure's parameters and sets up the environment for the rest of the procedure. The setup step will create a snapshot of your source code or set up a ClearCase view also. After setup, the remaining steps should be easier to write because they just use information created during the setup step.

Where's the script for a step?

Typically, each step executes a script of some type, which is processed by a command language such as *cmd* on Windows, a *shell* on UNIX, or Perl.

Two ways to handle script execution

- The first approach Enter the script directly into the command field for the step. You can specify the language interpreter as the shell for the step.
- The second approach Store the script in a file that is part of the source code for your project, then specify a simple command for the step that invokes the command language to process the script file.

We recommend using the second approach in most cases. The main advantage of this approach is that it makes your ElectricFlow procedures more robust. For example, suppose the script needs to change as your product evolves. If you have kept the script outside ElectricFlow with the source code for the product, each product version can have its own copy of the script. When you extract the source code for the product at the beginning of a job, you also extract the scripts for its steps. You can change a step's script without worrying about its impact on other versions of the product.

However, if the script for a step is stored in the step, it is more difficult to evolve the script with new product versions. You probably still need to build older product versions, so you will worry whether the new script for the step will work with older product versions. If the script does not work, perhaps you can modify the step's script to test the version being built and take different actions for each version—this process becomes increasingly more complex as the number of versions increases. Or, you can make a separate copy of the build procedure for each product version (more on this below), but this process also gets complicated as you acquire more and more procedure versions. We find it is easier to store scripts for steps with the product code, so script changes are handled in the same way as changes to the product.

Two cases where it makes sense to store the script for a step in the ElectricFlow step

The first case is for the first step of a job.
 At this point you have not extracted the source code for the product, so you do not yet have access to any scripts stored with the source code.

The second case is for steps with one or two commands only, or steps already specified primarily by
information stored with your source code.
 For example, if a step is running a make or ant command, step behavior is already specified almost
entirely by a makefile or a build.xml file for Ant. In this case, the step's script invokes a single, relatively
simple command; no need to store this in a file. If you subsequently find that the script for a step is

changing frequently, consider moving it out of ElectricFlow and into a file stored with your source code.

The object is to make your ElectricFlow procedures as reusable as possible, so no change is needed with every small change to your product. Even better, organize your ElectricFlow procedures so a single procedure can be used for multiple products. To do this, store product-or version-specific information with the product, not with ElectricFlow.

Environment variables

Your build and test scripts probably depend on certain environment variables having certain values. In your existing system, you probably set those values at the beginning of your script. However, in ElectricFlow you need to set those values in each step that depends on them. The easiest way to set values is to create a short command file during the first step of the job and save it in the top-level directory of the job workspace. The command file should contain a sequence of commands to set all environment variables required by the job. Then, in each subsequent step, invoke that command file at the beginning of the step to set environment variables for that step. This approach simplifies management of your environment variables: if you need to change a variable, you change only the setup step at the beginning of the job, and the value is reflected in all of the following steps automatically.

When do you use subprocedures?

In ElectricFlow, the action for a step can invoke another procedure, passing parameters. This is a powerful tool for structuring your processes. Subprocedures tend to be used in two ways: for *encapsulating reusable processes* and for *managing concurrency*.

- encapsulating reusable processes, is the preferred use. If you need to perform certain activities
 repeatedly in different places, you can use a subprocedure for them. For example, when we build and
 test the ElectricFlow product, we do it on multiple platforms, but the mechanism is virtually the same on
 all platforms. We implement the basic build and test mechanism for one platform in a subprocedure
 named BuildAndTest. In our main procedure for production builds, we invoke BuildAndTest multiple
 times, once for each of the platforms. If the mechanism changes, we only need to change it once in
 BuildAndTest to fix all six platforms.
- managing concurrency suppose you have three steps A1, A2, and A3, that must run sequentially, and two other steps B1 and B2 that must also run sequentially, but the two groups can run in parallel. The way to implement this is to put A1, A2, and A3 in one subprocedure named "A", and B1 and B2 in another subprocedure named "B". Now you can create a top-level procedure with two steps: one invoking A and the other invoking B, and mark those steps for parallel execution. As a result, A and B will run in parallel but the steps inside each subprocedure will run serially.
 If you placed all five steps inside the top-level procedure, there would be no way to achieve this effect.

How do you evolve procedures?

After you start using ElectricFlow for managing your build and test processes, it will not be long before you encounter the following situation:

You have a set of ElectricFlow procedures that work fine on all existing versions of the product, but you are about to change the product in a way that affects ElectricFlow procedures. For example, you might be adding a new component, or perhaps you are going to change the way the product is installed for testing, or perhaps you have restructured the product to allow more steps to run concurrently. As a result, you need to change ElectricFlow procedures for the current version. At the same time, you need older product versions to continue building as well.

In most cases, the easiest way to handle these situations is to leave the current ElectricFlow procedures alone and keep using them for your older product version. This ensures you will not "break" older versions. Make copies of the procedures that need to change, then modify the copies and use them for your new software version.

You will have one version of procedures for each significant version of the product. For example, at Electric Cloud, we incorporate the version number into the procedure names: "Master-1.0" is the version of the Master procedure used with the 1.0 release, "Master-1.1" is used for 1.1, and "Master" with no version number is used for our current development.

Another approach would be to make a copy of the entire project for each release. This method may be easier if other information in the project, such as property values, needs to evolve also.

The "copying" approach gets more and more complicated as you add more and more versions of the procedure. To minimize this complexity, try to structure your procedures so they do not have to change frequently. Some techniques you can use:

- Use parameters for things that do change frequently, such as the branch of software to build.
- Where possible, store scripts and other information [used by the procedure] as part of the source code
 for the project as described above, so this information is versioned automatically, along with your source
 code.

Porting existing scripts

You probably already have scripts you have been using to build and test software, before switching to ElectricFlow.

Over time, you will probably want to do quite a bit of restructuring to take full advantage of ElectricFlow. However, you probably do not need to do major restructuring to get started with ElectricFlow. Here are some steps you might take to "get up and running" with ElectricFlow as quickly as possible and gradually convert to make the best use of the system.

To begin, see if you can take your existing script and run it monolithically as a single step in a procedure under ElectricFlow. This will get you running and start providing some ElectricFlow benefits for resource management, error analysis, and reporting.

Next, start dividing your previous script into separate steps for ElectricFlow. Begin with the steps easiest to separate from the rest of the script, such as the commands to compile your software. One of the challenges you face is how to pass data from one step to another. In a monolithic script, you can keep the data in variables that persist throughout the job. As you divide the script into steps, you need to figure out which variables are used only in a single step and which variables pass from step to step. For variables that pass between steps, use ElectricFlow properties to store their values. In many cases you can compute shared

data values in a single step at the beginning of the job, store their values in properties on the job, then access those values read-only in later job steps.

After dividing the steps, you can do finer-grain reporting. Start using the *postp* postprocessor to scan your log files and generate statistics and diagnostics. Over time you will probably discover you have a few error and warning messages peculiar to your site, and are not captured by postp's patterns. Learn how to extend postp with additional patterns to capture all the information that matters to you.

Now that reporting statistics are being recorded, you may wish to develop your own reports to summarize those statistics. You can easily use *ectool* to read statistics from the properties where they are stored.

Something else you can do after steps are divided: start tuning the performance of your procedures. For example, you can start marking steps to run in parallel, and you can choose resources on a step-by-step basis to finish your jobs faster and share resources more effectively.

ElectricFlow project version control

We recommend exporting your ElectricFlow project data at regular intervals and saving it in the same configuration management system you use for your source code (use the "export" ectool command to generate an XML file representing the contents of a project or procedure). This process allows you to track project changes and also allows you to "look back" at older versions if that should become desirable. Also, you should snapshot a version of ElectricFlow project data at the time of each major software release, so you can revert to the exact ElectricFlow configuration used to generate that release if necessary.

ElectricFlow Installed Tools

This reference topic can acquaint you with some tools you may not know about or provide quick access to information for tools you use already. Select from the links in the table below to go directly to information for that tool.

Tool Name	Description
eccert	A command-line tool used to manage the ElectricFlow Certificate Authority (CA) and the certificates configured in ElectricFlow Server and ElectricFlow Agent installations.
ecconfigure	A command-line tool that can change configuration values for any locally installed ElectricFlow server, web, agent, or repository service. ecconfigure is a more user-friendly mechanism for configuring aspects of ElectricFlow that would otherwise require manual configuration file updates. ecconfigure actually manipulates relevant service configuration files on your behalf.
	IMPORTANT: You must have root privileges to use thewebPluginsDirectory=path option.

Tool Name	Description
ecdaemon	A "wrapper" program that can be used to start another program from an ElectricFlow job step—the "started" program will run as a daemon process. The ElectricFlow agent uses the facilities of the underlying operating system to make sure the process runs in a separate process group on a UNIX-based system, or outside of the normal "Windows Job" grouping in a Windows system. In either case, the ElectricFlow agent does not treat the process as one it should wait for or one it should try to "kill" if ElectricFlow needs to abort the step.
ecproxy	A driver script with built-in support for SSH. Every major operation can be overridden by defining a Perl function in the Proxy Customization field on the New Proxy Resource panel, available from the Resources page,
ecremotefilecopy	Remote file copy: When ElectricFlow agents (on platforms other than Linux or Windows) run steps that create log files in a workspace the ElectricFlow web server cannot access (through Linux or Windows agents), use ecremotefilecopy to recreate job logs so they are visible on those ElectricFlow agents, which then enables the web server to retrieve and render those log files.
ZKConfigTool	A command-line tool that imports your ElectricFlow database configuration information into your ZooKeeper server.
ClusterInfoTool	A command-line tool that displays information on the running ElectricFlow server cluster from ZooKeeper.

eccert

A command-line tool used to manage the ElectricFlow Certificate Authority (CA) and the certificates configured in ElectricFlow Server and ElectricFlow Agent installations.

Usage

```
eccert [ options ] command [ arg ... ]
```

Commands

addTrustedServer crt	Add a server CA certificate to the agent's keystore.
getCRL	Retrieve the contents of the current certificate revocation list.
<pre>initAgent [local remote] [options]</pre>	Initialize the agent keystore with a new public/private key pair. Generates the agent certificate signing request. If run on the server host, the certificate will automatically be signed by the server CA, and the CA certificate and the signed agent certificate are installed in the agent's keystore. If run on a non-server host, the signing request is left in the agent directory. If CA Cert is provided, the CA certificate is installed in the agent's keystore.
	local Use the local server CA to sign the agent certificate.
	remote Connect to a remote ElectricFlow server to sign the agent certificate.
	force Replace any existing keystore.
	cname name Use the specified name as the common name (CN) in the agent certificate subject. This is normally the fully qualified domain name used by clients to connect to the agent.
	altNames entries Use the specified list of entries (comma or space separated) as the subjectAlternateNames list in the agent certificate. Simple names are interpreted as dns entries. Entries may begin with "dns:" or "ip:" to indicate the type (for example, "ip:192.168.0.1" or "dns:myHost"). If no entries are specified, then reverse DNS is used to look up the registered names of the host's IP addresses.
initCA	Initialize the server CA. Creates a new CA key and certificate.

<pre>initServer [options]</pre>	Initialize the server keystore. Creates and signs the server certificate. Installs the CA certificate and the signed server certificate into the server's keystore.
	force Replace any existing keystore.
	cname name Use the specified name as the common name (CN) in the server certificate subject. This is normally the fully qualified domain name used by clients to connect to the server.
	altNames entries Use the specified list of entries (comma or space separated) as the subjectAlternateNames list in the server certificate. Simple names are interpreted as dns entries. Entries may begin with "dns:" or "ip:" to indicate the type (for example, "ip:192.168.0.1" or "dns:myHost"). If no entries are specified, then reverse DNS is used to look up the registered names of the host's IP addresses.
list [agent server index [verbose]	Display certificate information for agent and/or server keystores or the CA certificate index. If no options are specified, both the agent and server keystores are listed.
	agent List the contents of the agent keystore.
	server List the contents of the server keystore.
	index List the contents of the CA issued certificates index.
	verbose Display additional details.
refreshCRL	Refresh the certificate revocation list from the ElectricFlow server.

revoke index	Revoke a previously issued certificate by index.
signCertificate csr crt	Sign the certificate signing request provided in file csr and write the signed result to the file crt. The request is rejected by the CA if there is a matching certificate already in the CA database.
updateAgentCertificate crt	Install a previously signed certificate crt into the agent's keystore.

Server Communication Options

server host	Address of the ElectricFlow server. Defaults to the value of the COMMANDER_ SERVER environment variable. If that does not exist,it defaults to localhost.
securePort port	HTTPS listener port on the server. Defaults to 8443.

Global Options

help	Print the Help message.
version	Print the version message.

Examples

Example 1: Configure an agent to talk to any server (untrusted mode)

This example generates a new self-signed certificate for the agent and recreates the keystore with no trusted authorities.

```
$ eccert initAgent -force
Generating keys
Generating certificate request
    cname=myAgent.company.com
    san=dns:myAgent.company.com
```

Example 2: Configure an agent to accept connections only from a single remote ElectricFlow server

This example generates a new certificate for the agent that is signed by the remove server's certificate authority and installs the signed certificate and its associated trust chain in the agent's keystore. After this point, the agent

will only accept requests from the specified server and will be used as a trusted resource by the server.

Example 3: Configure an ElectricFlow server with additional host names in the certificate

This example regenerates the ElectricFlow Server Certificate, the specified common name, and alternate subject names to allow trusted connections with multiple external dns names.

ecconfigure

A command-line tool that can change configuration values for any locally installed ElectricFlow server, web, agent, or repository service. ecconfigure is a more user-friendly mechanism for configuring aspects of ElectricFlow that would otherwise require manual configuration file updates. ecconfigure actually manipulates relevant service configuration files on your behalf.

Usage

```
ecconfigure [options]
```

ElectricFlow agent configuration options

agentInitMemory=percent	Initial java heap size as a percentage of the total system memory.
agentKeystorePassword=password	Password used to access the agent's keystore.
agentMaxMemory=percent	Maximum java heap size as a percentage of the total system memory.
agentInitMemoryMB=size	Initial java heap size in MB.
agentMaxMemoryMB=size	Maximum java heap size in MB.
agentLogFile=path	Path where the agent log file should be written.
agentLogMaxFiles=max	Maximum number of log files to accrue.
agentLogMaxSize=max	Maximum size of each log file. The value can be suffixed with a unit (MB, KB, B). Without a unit,the value is interpreted as bytes.

agentPluginsDirectory=path	The path used by the agent to get to the plugins directory on the ElectricFlow server where its resource definition lies.
agentLocalPort=port	Port used by the ElectricFlow agent for HTTP communication on the localhost network interface.
agentPort=port	Port used by the ElectricFlow agent for HTTPS communication on any network interface.
agentProtocol	Protocol used by the agent.
agentProxyHost=host	The IP address of the proxy server.

agentProxyPort=port	The port of the proxy server.
agentNoProxyHosts=hosts	Comma delimited list of hosts that should be reached directly, bypassing the proxy server.
agentEnableProxySettings=<0 1>	Enable or disable the proxy server configuration. If enabling for the first time, agentProxyHost and agentProxyPort must be specified.

Apache web server configuration options

webHostName=host	The host name of the current machine in the form that users will typically use in their browser to access the web server.
webHttpPort=port	The HTTP port of the web server.
webHttpsPort=port	The HTTPS port of the web server.
webTargetHostName=host	The host name of the ElectricFlow server to which the web server points.
webTargetHttpPort=port	The HTTP port of the ElectricFlow server to which the web server points.
webTargetHttpsPort=port	The HTTPS port of the ElectricFlow server to which the web server points.
webTimeZone=timezone	The Olson TimeZone format (example: America/Los Angeles) for the php web server.

webPluginsDirectory=path	The path used by the web server to get to the plugins directory on the ElectricFlow server to which it points. IMPORTANT: You must have root privileges to use thewebPluginsDirectory=path option.
webProxyUrl=url	The IP address and port of the proxy server in the following format: http:// <ip_address_proxy>:<proxy_port></proxy_port></ip_address_proxy>
webNoProxyHosts=hosts	Comma-delimited host list that should be reached directly, bypassing the proxy server.
webEnableProxySettings=<0 1>	Enable or disable the proxy server configuration. If enabling for the first time,webProxyUrl must be specified.
webDLC=url	The URL to use for downloadable content requests.

ElectricFlow Server configuration options

serverFileTransferPort=port	The file transfer port of the server.
serverHttpPort=port	The HTTP port of the server.
serverHttpsPort=port	The HTTPS port of the server.
serverInitMemory=percent	Initial java heap size as a percentage of the total system memory.
serverMaxMemory=percent	Maximum java heap size as a percentage of the total system memory.

serverInitMemoryMB=size	Initial java heap size in MB.
serverMaxMemoryMB=size	Maximum java heap size in MB.
serverPasskeyFile=path	Path to the server's passkey file.
serverProxyHost=host	The IP address for the proxy server.
serverProxyPort=port	The port for the proxy server.
serverNoProxyHosts=hosts	Comma-delimited host list that should be reached directly, bypassing the proxy server.
serverEnableProxySettings=<0 1>	Enable or disable the proxy server configuration.lf enabling for the first time, serverProxyHost and serverProxyPort must be specified.
serverZooKeeperConnection <zookeeper_ servers></zookeeper_ 	The location of the ZooKeeper servers.

Repository Server configuration options

repositoryPort	The repository server port.
repositoryInitMemory=percent	Initial java heap size as a percentage of the tot al system memory.
repositoryMaxMemory=percent	Maximum java heap size as a percentage of the tot al system memory.
repositoryInitMemoryMB=size	Initial java heap size in MB.
repositoryMaxMemoryMB=size	Maximum java heap size in MB.
repositoryStorageDirectory=path	Path to the repository backing store.The artifact repository will use this directory to store artifacts.
repositoryTargetHostName=host	The host name of the ElectricFlow server to which the repository server points.
repositoryTargetHttpPort=port	The HTTP port of the ElectricFlow server to which the repository server points.
repositoryTargetHttpsPort=port	The HTTPS port of the ElectricFlow server to which t he repository server points.

repositoryTargetProtocol= <http https></http https>	The protocol the repository server uses to talk to the ElectricFlow server.
repositoryProtocol	The protocol the repository server uses to talk to client applications.
repositoryProxyHost=host	The IP address for the proxy server.
repositoryProxyPort=port	The port for the proxy server.
repositoryNoProxyHosts=hosts	Comma-delimited host list that should be reached directly, bypassing the proxy server.
repositoryEnableProxySettings=<0 1>	Enable or disable the proxy server configuration. If enabling for the first time, repositoryProxyHost and repositoryProxyPort must be specified.
repositoryValidateFromDisk=<0 1>	Enable or disable disk validation.

General options

-v,version	Display version information.
-h,help	Display this information.

Examples

Setting initial and maximum memory settings

For example, to set the ElectricFlow Server initial memory percentage to 21% and the maximum memory percentage to 31%, specify:

```
ecconfigure --serverInitMemory 21 --serverMaxMemory 31
```

If your ElectricFlow server, web server, or repository server is deployed behind a proxy server that inhibits certain Internet access, you can use ecconfigure to set proxy settings for each server in your installation.

To use the following perl scripts, remove the brackets ("< >"), and replace the bracketed example text with the values you need.

• To set ElectricFlow Server proxy settings:

```
ecconfigure --serverProxyHost <IP_ADDRESS_PROXY> --serverProxyPort <PORT>
    --serverNoProxyHosts "<HOST1,HOST2>"
```

• To set Repository Server proxy settings:

```
ecconfigure --repositoryProxyHost <IP_ADDRESS_PROXY> --repositoryProxyPort
<PORT>
    --repositoryNoProxyHosts "<HOST1,HOST2>"
```

To set Web Server proxy settings:

```
ecconfigure --webProxyUrl <http://IP_ADDRESS:PORT> --webNoProxyHosts
<HOST1,HOST3>
```

Changing the Apache web server port

```
Run ecconfigure with --webHttpPort and --webHttpsPort.
```

You web server port setting will be changed appropriately in httpd.conf, ssl.conf, and config.php.

Configuring the backing store location for the Artifact Repository

```
ecconfigure --repositoryStorageDirectory <new-path>
```

ecdaemon

ecdaemon is a "wrapper" program that can be used to start another program from an ElectricFlow job step—the "started" program will run as a daemon process. The ElectricFlow agent uses the facilities of the underlying operating system to make sure the process runs in a separate process group on a UNIX-based system, or

outside of the normal "Windows Job" grouping in a Windows system. In either case, the ElectricFlow agent does not treat the process as one it should wait for or one it should try to "kill" if ElectricFlow needs to abort the step.

Use Cases

- ecdaemon is useful in the case where you are trying to deploy a "server-style" program in an
 ElectricFlow step. You do not want ElectricFlow to wait for that step to complete because it may run
 continuously, but you do want ElectricFlow to start the server program and then continue on to the next
 step.
- ecdaemon is useful if you want to "pre-load" some type of background process

ecdaemon launches the command and exits. Optionally, it sets a property in ElectricFlow with the pid of the program it spawned to make it possible for a later step to "kill" the daemonized program if desired.

For example:

```
ecdaemon c:/install.exe a b c
```

Using ecdaemon

Command-line parsing

ecdaemon supports the standard UNIX-style '--' flag to indicate there are no more ecdaemon options and all subsequent options should be treated as simple arguments to the command.

This is particularly important for commands that themselves take '--' arguments.

For example:

```
ecdaemon /usr/bin/myserver --config /etc/myserver.conf
```

will not run properly because ecdaemon will attempt to parse the --config option instead of passing it to the myserver program.

The correct way to invoke ecdaemon in this case is:

```
ecdaemon -- /usr/bin/myserver --config /etc/myserver.conf
```

If you want to store the daemonized process's pid in a property, do so as follows:

```
ecdaemon --pidProperty /myJob/serverPid -- /usr/bin/myserver --config
/etc/myserver.conf
```

As a daemon process, any output goes to /dev/null, therefore no output file is generated.

ec-perl considerations

Use the perl system () call to start ecdaemon.

System () returns an exit status, "backticks" capture and return output that waits for the daemonized command to complete on Windows, and <code>exec never</code> returns at all if it is successful.

ecproxy

A driver script with built-in support for SSH. Every major operation can be overridden by defining a Perl function in the Proxy Customization field on the New Proxy Resource panel, available from the Resources page (by specifying which operation this function re-implements. These operations must have certain "signatures" for the driver to invoke them properly—the operations are listed and described below. For more detail, see the SSH implementation in ecproxy.pl.

ecproxy Algorithm

ecproxy invokes the operations detailed below to perform the following actions:

- 1. Uploads the command-file to a "workingDirectory" on the proxy target, using the protocol specified in the proxy config. Currently, only SSH is supported.
- 2. Creates a wrapper sh shell script that CD's to workingDirectory, sets "COMMANDER_ environment variables" that exist in the proxy agent's environment, and runs the command-file previously uploaded.
- 3. Uploads the wrapper script to workingDirectory on the proxy target.
- 4. Runs the wrapper script on the proxy target and streams its output to the proxy agent's stdout.
- 5. Deletes the local wrapper shell script, the remote wrapper, and remote command-file.
- 6. Exits using the exit code of running the wrapper script.

ecproxy Operations

getDefaultWorkingDirectory

Description	Computes the default working directory where a command needs to run on the proxy target, if the step is not defined with a working directory.
Arguments	None
Returns	The path to a directory as it would be accessed on the proxy target.
SSH implementation function	Not SSH-specific, so the function is 'getDefaultWorkingDirectory'

Existing implementation	Return \$ENV{COMMANDER_WORKSPACE_UNIX};
Reason to override	If the "Working Directory" field is empty on a step that is going to run on a proxy target, the working directory for the step should be the workspace, just as it would be if the step were running on a non-proxy ElectricFlow agent. ecproxy is guaranteed to run in the workspace directory on the proxy agent, but it is not guaranteed that the proxy target has the same path to the workspace. For example, the workspace on a Windows proxy agent is not the same as the path a Unix proxy target uses to access the workspace—so the existing implementation of this operation simply returns the UNIX path to the workspace. However, if the proxy target has a different path for accessing the workspace, the existing implementation will give the wrong answer. Thus, a user can provide a different implementation that gives the right answer.
	I

Note: This operation is applicable only if "Working Directory" is empty, and is used as the working directory on the proxy target in that case for running the command.

getDefaultTargetPort

Description	Computes the default port where the proxy target is listening for this protocol.
Arguments	None
Returns	The default port.
SSH implementation function	ssh_getDefaultTargetPort
	,

Note: This operation is applicable only if the resource definition specifies no port value.

connect

Description	Opens a connection to the proxy target using the desired protocol.
Arguments	host, port (optional)

Returns	"connection-context hash-ref" on successful connection. This context can contain anything other functions can use to perform their tasks (for example, a connection handle).
Example 1	my \$context = connect('myhost', 22)
Example 2	<pre>my \$context = connect('myhost')</pre>
SSH implementation function	ssh_connect

Note: Because the port is optional, the implementation of connect can default to whatever is reasonable for the protocol. This means the 'Proxy Target Port' need not be specified in the ElectricFlow web UI for proxied agents reachable on the default port.

uploadFile

Description	Uploads the given srcFile to the proxy target as tgtFile.
Arguments	context, srcFile (typically simple file-name), tgtFile (typically workingDirectory/file-name)
Returns	Nothing; on failure it does a 'die' with an appropriate error message.
Example	<pre>uploadFile(\$context, 'agent123.tmp', '/opt/work/joe/agent123.tmp')</pre>
SSH implementation function	ssh_uploadFile

generate Wrapper Script

Description	Generates the script body that will run the command-file on the proxy-target in the workingDirectory.
Arguments	<pre>workingDirectory, cmd, cmdArg1,, cmdFileName (just the base-name, no directory)</pre>
Returns	A string containing the script to execute on the proxy target.
Example	<pre>generateWrapperScript('/opt/work/joe', 'perl', 'agent123.tmp'</pre>

SSH implementation function	Not SSH-specific, so the function is 'generateWrapperScript'
Existing implementation	cd workingDirectory; set COMMANDER_ environment variables; run command-file, properly quoting the command and args.
Reason to override	If the proxy target does not have sh, the wrapper script needs to be written in a language available on the target.

uploadWrapperScript

Description	Uploads the wrapper-script code to the proxy target.
Arguments	context workingDirectory.wrapperScriptBody
Returns	The path to the wrapper-script on the proxy target.
Example	<pre>my \$wrapperFile = uploadWrapperScript (\$context, '/opt/work/joe', 'cd /opt/work/joe; perl agent123.tmp')</pre>
SSH implementation function	3.1, 3.1.1: ssh_uploadWrapperScript
3.1.2 and later	Not SSH-specific anymore, so the function is 'uploadWrapperScript'

Notes:

- 1. This function must generate a uniquely named file that will not conflict with other ecproxy invocations that might be occurring in parallel steps. The recommended approach is to generate a file name containing the job-step-id.
- 2. Depending on the protocol and facilities available in the Perl implementation, you may or may not need to create a local tmp file to upload to the proxy target. If you do, record that fact in the context and clean up the local file in the cleanup operation. Setting the local wrapper file path in \$context->{wrapperFile} is recommended because the default cleanup operation implementation looks for that string.

generate Wrapper Invocation Command

Description	Generates the command-line for running the wrapper script on the proxy target.
Arguments	remoteWrapperFile (path on proxy target)
Returns	A string containing the command-line for running the wrapper script file on the proxy target.
Example	my \$wrapperCmdLine = generateWrapperInvocationCommand(\$wrapperFile)
SSH implementation function	Not SSH-specific, so the function is 'generateWrapperInvocationCommand'
Existing implementation	Return "sh \$remoteWrapperFile";
Reason to override	The default implementation of this function returns something like 'sh \$wrapperFile'. If the wrapper script is not an sh script, or if you want to pass different arguments to the shell, you must override this function.

runCommand

Description	Runs the given command-line on the proxy target.
Arguments:	context, cmdLine
Returns	exit-code from running the command on the proxy target, undef if the command could not be run for some reason.
Example	runCommand(\$context, \$wrapperCmdLine)
SSH implementation function	ssh_runCommand

cleanup

Description	Performs any cleanup task after the command has completed on the proxy target. Typically, it deletes any locally created temp files and uploaded files on the proxy target.
Arguments	<pre>context, cmdFile, wrapperFile (both are of the form workingDirectory/file-name)</pre>
Returns	1 on success, undef on failure.
Example	<pre>cleanupTarget(\$context, '/opt/work/joe/agent123.tmp', '/opt/work/joe/cmdwrapper.123.tmp')</pre>
SSH implementation function	3.1, 3.1.1: ssh_cleanup
3.1.2 and later	Not SSH-specific anymore, so the function is cleanup

Note: The default implementation deletes the locally created wrapper script file whose path is stored in \$context->{wrapperFile}, if it exists. Thus, if the uploadWrapperScript operation is overridden, it is recommended the overriding function set this attribute—that way, cleanup need not be overridden.

ping

Description	A test to see if the proxy target is usable.
Arguments	host, port (optional)
Returns	1 on success, undef on failure.
Example	ping('myhost', 22)
SSH implementation function	Not SSH-specific, so the function is ping.
Existing implementation	Opens a socket connection to the proxy target on the desired port.
Reason to override	The existing implementation may be deemed too simple for doing a ping; overriding ping to open a connection and do some protocol-specific handshaking might be more appropriate for some protocols / use cases.

Available Helper Functions

To make proxy customization easier, ecproxy provides the following helper functions.

mesg

Description	Debug logging function. Writes to the file referenced in the ECPROXY_DEBUGFILE environment variable (if it exists). No-op otherwise.
Arguments	message
Example	<pre>mesg("myCleanup: about to delete \$cmdFile on proxy target");</pre>

Note: This function automatically adds a newline to whatever it emits, so the caller does not have to incorporate a newline in message.

readFile

Description Reads a file.

Arguments fileName

Returns Contents of the file. If there is an error, it returns an empty string.

writeFile

Description	Creates a local file containing data.
Arguments	fileName,data
Returns	1 on success, undef on failure.
Example	<pre>writeFile("myWrapper.\$ENV{COMMANDER_ JOBSTEPID}.cmd", "perl foo.pl")</pre>

initDispatcher

Description	Initialize the operation dispatcher map to point to functions for the given protocol. For each operation, initDispatcher checks if a function named protocol_operation exists, and if so, assigns that function as the implementation for that operation.
Arguments	protocol
Example	<pre>initDispatcher("ssh") sets the "connect" operation to run "ssh_ connect", "uploadFile" => "ssh_ uploadFile", etc.</pre>

setOperation

Description	Sets the implementation of an operation to be a particular function.
Arguments	operation, function. The 'function' argument may be the name of a function or a reference to a function.
Example	<pre>setOperation("ping", "my_ping"); sets the "ping" operation to run the "my_ ping" function</pre>
Example	<pre>setOperation("ping", \&my_ping); same as above, but using a function ref</pre>

 $\textbf{Note:} \ \textbf{This function manipulates the $\tt gDispatcher hash, but provides a safe}$

loadFile

interface to it.

Description	Load proxy customizations from a file.
Arguments	fileName
Example	loadFile("custom.pl")

setSSHKeyFiles

Description	Set the paths to the public and private key files that ssh will use to authenticate with the proxy target.
Arguments	publicKeyFile, privateKeyFile
Example	<pre>setSSHKeyFiles('c:\foo\pub.key', 'c:\foo\priv.key')</pre>

Note: This is very useful on Windows proxies, where there is no reasonable default for ssh to use.

setSSHUser

Description	Set the name of the user to authenticate with the proxy target.
Arguments	userName
Example	setSSHUser('user1')

Note: By default, the user name the agent is "running as" is used to log into the proxy target. If key-based authentication is configured on the target system such that 'agentUser' can log into the 'user1' account on the proxy target, this function leverages that configuration.

useMultipleSSHSessions

Description	Normally, ecproxy uses one ssh session with a number of "channels" to perform tasks like uploading files, running the command, and running a cleanup command on the proxy target. Some SSH servers don't allow this. This method configures ecproxy to use a separate SSH session for each operation; this requires authenticating with the SSH daemon on the proxy target several times, and thus it may perform worse than the single-session-multi-channel mode.
Arguments	None
Example	useMultipleSSHSessions()

Examples

Specify public/private key files for SSH

- Set the proxyCustomization property on the resource like this: ?setSSHKeyFiles ('c:\foo\pub.key', 'c:\foo\priv.key');
- 2. Set the ECPROXY_SSH_PRIVKEYFILE and ECPROXY_SSH_PUBKEYFILE environment variables on the proxy agent as system environment variables.

Override one of the operations (for example, to enable SSH connection with username/password)

```
Set the proxyCustomization property on the resource like this: ?sub myConnect ($$) {...} setOperation ("connect", \&myConnect);
```

Load proxy customizations from a file rather than having all the logic in the proxyCustomization property on the resource

Set the proxyCustomization property on the resource like this: ?loadFile('c:\foo\custom.pl');

Implement a whole new protocol

Specify protocol as 'myproto' and have a proxy customization block like this:

```
sub myproto_getDefaultTargetPort() {
...
}
sub myproto_connect($;$) {
...
}
sub myproto uploadFile($$$) {
```

```
sub myproto_uploadWrapperScript($$$) {
# Note: As of 3.1.2, the default implementation is likely good enough, so
it may not be necessary to define this override
....
} sub myproto_runCommand($$) {
...
} sub myproto_cleanup($$$) {
# Note: As of 3.1.2, the default implementation is likely good enough, so
it may not be necessary to define this override
....
}
# Initialize the dispatcher to run these functionsinitDispatcher("myproto");
```

Override ping to do a connect operation (which does a full protocol handshake, authentication, and so on)

Write a specialized ping function for the proxy customization like this:

```
sub heavy_ping($$) {
   my ($host, $port) = @_;
   return ssh_connect($host, $port);
}
setOperation("ping", \&heavy ping);
```

"Real World" Examples

ClusterExec

A basic integration for using clusterupload and clusterexec to reach a proxy target is here. It has been tested on a Windows target with a Cygwin installation. It will not work "out-of-the-box" because it makes the following assumptions:

- The proxy target has sh and other UNIX tools (for example, rm).
- The locations of the clusterexec and clusterupload binaries are hard-coded at the top of the proxy customization.

To make this proxy customization work on a Windows machine that does not have Cygwin, the generateWrapperScript operation would need to be overridden with a function that generates a cmd batch script, and the generateWrapperInvocationCommand operation would have to be overridden to generate a "cmd /c ..." command rather than "sh ...".

MySQL

The idea here is that the proxy target need not be a host for running arbitrary commands. It could be a special entity like a db. This integration uses the <code>mysql clt</code> to run the step command (which should be SQL) on the db referenced by the proxy target host and port.

A bare-bones integration with MySQL:

```
# Set the path to the mysql binary; if the directory is in the proxy agent's # PATH, this variable can simply contain the name of the executable.
```

```
my $gMySQL = "c:/cygwin/usr/local/tools/i686_win32/bin/mysql.exe";
sub mysql getDefaultTargetPort() {
    return 3306;
sub mysql connect($;$) {
    # This "protocol" implementation is just going to use the mysql
    # command-line tool, so just save off host/port.
   my host = [0];
   my $port = $_[1] || mysql_getDefaultTargetPort();
    return {host => $host,
           port => $port};
sub mysql_uploadFile($$$) {
   my ($context, $cmdFile, $rmtCmdFile) = @ ;
    # We do not need to upload the command-file to the proxy target.
    \# We are going to run the mysql clt on the proxy agent to run
    # the query (contained in the local command-file),
    # so just save off the name of the command-file.
    $context->{cmdFile} = $cmdFile;
sub mysql uploadWrapperScript($$$) {
   my ($context, $workingDir, $wrapperScript) = @;
    # This has no meaning for this integration. No-op.
sub mysql runCommand($$) {
   my ($context, $cmdLine) = @ ;
    # cmdLine is a command-line for running the wrapper script, which
    # has no meaning for this integration. We just want to run
    # 'mysql' for the desired host/port with the command-file.
    system("$qMySQL -D commander -h $context->{host} -P $context->{port} " .
        "-u commander -pcommander -e \"source \context->\cmdFile\\"");
sub mysql cleanup($$$) {
    # We didn't create any temp files. No-op.
# Initialize the dispatcher to run these functions
initDispatcher("mysql");
```

Android

This example uses the adb tool to upload files to the device and run commands on it. Initial testing has been only against the android emulator, but it is implemented in such a way that it should work against a real android device attached using USB to the proxy agent, or a device on the network.

A first attempt at proxying to android devices:

```
# Set the path to the adb binary; if the directory is in the proxy agent's
# PATH, this variable can simply contain the name of the executable.
my $gADB = "c:/android-sdk-windows-1.6 r1/tools/adb.exe";
sub android getDefaultTargetPort() {
    # Not sure what a good meaningful value is here.
    return 0;
sub android connect($;$) {
    # This "protocol" implementation uses the adb
    # command-line tool. Depending on the value of
    # host, construct the appropriate adb command-line
    # argument.
   my $host = $ [0];
   my $context = {};
   if ($host eq "emulator") {
    if ($host eq "localhost") {
        # We want to talk to the emulator running on this host.
        $context->{targetArg} = "-e";
    } elsif ($host eq "usb") {
        # We want to talk to the single android device connected
        # to the computer via a USB.
        $context->{targetArg} = "-d";
    } else {
        # This must be the serial number of some device somewhere.
        $context->{targetArg} = "-s $host";
    return $context;
sub android uploadFile($$$) {
   my ($context, $srcFile, $tgtFile) = @ ;
   my($filename, $directories) = fileparse($tgtFile);
   my $result = `$gADB $context->{targetArg} push $srcFile
    "/data/tmp/$filename" 2>&1`;
    if ($? != 0) {
        die ("android uploadFile: Error uploading file $srcFile to
        /data/tmp/$filename: $result\n");
sub android runCommand($$) {
   my ($context, $cmdLine) = @_;
    # cmdLine is a command-line for running the wrapper script, which
    # has no meaning for this integration. We just want to run
    # 'adb' for the desired device with the command-file.
    system("$gADB $context->{targetArg} shell $cmdLine");
sub android cleanup($$$) {
   my ($context, $remoteCmdFile, $remoteWrapperFile) = @ ;
    # This was copied from ssh cleanup except that we do "rm",
    # not "rf -f".
```

```
mesg("cleaning up");
    # Delete the locally generate wrapper file.
   unlink($context->{"wrapperFile"});
    # Delete the cmd-file and wrapper script file on the proxy target.
    $gDispatcher{"runCommand"} ($context,
        "rm $remoteWrapperFile $remoteCmdFile");
sub android ping($;$) {
   my ($host, $port) = @;
   $port = $qDispatcher{"qetDefaultTarqetPort"}() unless isPortValid($port);
   my $socket = IO::Socket::INET->new(PeerAddr => $host,
                                       PeerPort => $port,
                                       Proto => "tcp",
                                       Type => SOCK STREAM)
                           or die "Couldn't connect to $host:$port : $@\n";
# Initialize the dispatcher to run these functions
initDispatcher("android");
1;
```

ecremotefilecopy

When ElectricFlow agents (on platforms other than Linux or Windows) run steps that create log files in a workspace the ElectricFlow web server cannot access (through Linux or Windows agents), use *ecremotefilecopy* to recreate job logs so they are visible on those ElectricFlow agents, which then enables the web server to retrieve and render those log files.

Using postp and *ecremotefilecopy*, the log file is populated and recreated in a workspace accessible to the ElectricFlow web server, allowing the Job Details page to display the log file. Although this functionality is supported, it is not a recommended method of operation. This method should be used only as a last resort when a shared file system (between alternate agents and primary platform agents [Linux and Windows]) is not possible.

The reasons ${\tt ecremotefilecopy}$ is not recommended are:

- You will not see logs in real time. Logs are not visible until the "recreate step" has completed running.
- There is a performance penalty, especially when running with large files.

Setting up the process

- Create a "Setup" step in your procedure
- Update the Postprocessor field for each step whose results you want to see on the server.
- Add a step (one or more times) to the procedure to recreate the ElectricFlow server log files.

Creating a Setup Step

In your procedure, create a step called "Setup". This step needs to be in your procedure **before** any step running on a remote workspace.

Note: This is your top-level procedure, not a subprocedure.

- · Navigate to your procedure.
- To create a new step, click the Command step link.
- Set the fields as follows:

Step Name: Setup

Command(s): ecremotefilecopy setup

Resource: local

Workspace: you have two choices:

- use default
- use:alternateWorkspaceForDisplay

There is a property on the Workspace called alternateWorkspaceForDisplay, which is a secondary location to look for workspace files. This secondary location is used when the workspace files are not accessible to the web server. If the Apache server cannot locate the file in the original workspace, it looks in the alternate one.

Update the Postprocessor field for steps in your procedure

This step defines a postprocessor that will run at the end of the steps you specify. Add the following information to every step running on a remote agent if you want to see its results in the web interface.

- · Navigate to a procedure and a step.
- In the Postprocessor field, enter:

```
postp --check none --loadProperty /myJob/jobSteps[Setup]/postpExtensions
```

• If you are using postp in this step to scan your step log for errors, warnings, and so on also, omit "-- check none" from the invocation line.

Add a Final Step to your procedure

This step adds a new step at the end of your existing procedure. This step finds all properties created by the postprocessor, then reads the properties and creates local log files based on the properties, then deletes the properties.

- Navigate to your procedure.
- To create a new step, click the Command step link.
- Set the fields as follows:

Step Name: Recreate the Log Files

Always run step: (Check the box)

Command(s): ecremotefilecopy recreateFiles

```
Resource: local
Workspace: default
```

After the final step runs, you should see links (icons) displayed in the Log column on the Jobs Details page.

Click the icon to display the log file.

Copying Other Files from the Workspace

By default, ecremotefilecopy copies only postp log and diag files, and step logs. You can also copy other files from the workspace using a function named postpEndHook2.

You must do the following in your step:

- 1. Make sure that the file you want to copy is in the step workspace. (It can be copied there, created there, etc.)
- 2. For your procedure, create a property (named postpEndHook2, for example).
- 3. Define a function named postpEndHook2 inside the property. For example:

```
sub postpEndHook2() {

    # Missing param does not cause an error
    $::gCommander->abortOnError(0);

    # Add filename to a "special" property such that it will be picked up by ecr
emotefilecopy
    my $fileName = 'paul.txt';
    copyFileToProperty ($fileName);

    # Restore default error handling
    $::gCommander->abortOnError(1);
}
```

4. Add the following line in the Postprocessor field of this step:

```
postp --check none --loadProperty /myJob/jobSteps[Setup]/postpExtensions --loadProperty /myProcedure/postpEndHook2
```

ZKConfigTool

You can ZooKeeper server with configuration information that all ElectricFlow server nodes will use in a clustered configuration. This command-line tool that imports your ElectricFlow database configuration information into your ZooKeeper server. The following minimum set of files is imported:

- database.properties
- keystore
- passkey
- commander.properties

Prerequisites

- The ElectricFlow Tools package must be installed on the system.
- The system must be running a version of Java supported by ElectricFlow. Java is automatically installed on a system with the ElectricFlow software as part of the Tools install.
- The ZooKeeper software must be installed on the network.

Command

\$ java -jar zk-config-tool-jar-with-dependencies.jar -<options>

Usage: ZKConfigTool

Option	Description
-commanderPropertiesFile <path_to_file></path_to_file>	Import the ElectricFlow commander.properti es file.
-databasePropertiesFile <path_to_file></path_to_file>	Import the ElectricFlow database.propertie s file.
-help	Show the command Help.
-keystoreFile <path_to_file></path_to_file>	Import the ElectricFlow keystore file.
-passkeyFile <path_to_file></path_to_file>	Import the ElectricFlow passkey file.
-wrapperConfFile <path_to_file></path_to_file>	Import thewrapper.conf file.
-writeFile <path_on_ zookeeper=""> <path_to_file></path_to_file></path_on_>	Write the specified file to the ZooKeeper server.

Import Files

Run the ${\tt ClusterTool}$ command to completely populate the ZooKeeper server with configuration information. The system must have the Tools install and can be able to communicate with ZooKeeper.

• If you run the command with defined file options, the tool immediately attempts to contact the specified ZooKeeper system and import the specified files.

Example command:

```
COMMANDER_ZK_CONNECTION=10.168.33.10:2181 java -jar zk-config-tool-jar-with-dependencies.jar --databasePropertiesFile database.properties --keystoreFile keystore --passkeyFile passkey --wrapperConfFile wrapper.conf --commanderPropertiesFile commander.properties
```

 If you run the command without defined file options, the tool displays a series of prompts for the required ElectricFlow information. Once the correct information is entered, the tool contacts the specified ZooKeeper system to import the specified files.

Example command:

```
$ COMMANDER_ZK_CONNECTION=10.168.33.10:2181 java -jar zk-config-tool-jar-with-dependencies.jar
```

ClusterInfoTool

Use ClusterInfoTool to get information on the running ElectricFlow server cluster from ZooKeeper.

Prerequisites

- The ElectricFlow server cluster must be installed and running on the network.
- Configuration files that all ElectricFlow server nodes will use in a clustered configuration must be uploaded to Apache ZooKeeper server using the ZKConfigTool.
- The ZooKeeper cluster must be running an odd number of Zookeeper nodes and there must be a leader node.
- The system must be running a version of Java supported by ElectricFlow. Java is automatically installed on a system with the ElectricFlow software as part of the Tools installation.

Locations

The ElectricFlow installer adds the ClusterInfoTool to the following default locations:

- Windows: C:\Program Files\Electric Cloud\ElectricCommander\server\bin\cluster-info-tool-jar-with-dependencies.jar
- Linux: /opt/electriccloud/electriccommander/server/bin/cluster-info-tool-jar-with-dependencies.jar

Command

```
$ java -DCOMMANDER_ZK_CONNECTION=<ZooKeeper_Server1_IP>:2181,<ZooKeeper_Server2_IP>:21
81,<ZooKeeper Server3 IP>:2181 -jar cluster-info-tool-jar-with-dependencies.jar
```

where DCOMMANDER_ZK_CONNECTION must point to the ZooKeeper system to which the ElectricFlow server cluster is connected.

Output

This is sample output generated by ClusterInfoTool:

```
Checking /commander/jgroups/hornetq: 582d3642-9736-e87f-4e19-c22d7776ccab ->
```

WIN-M3Q09A2PNFP-21066	6c80e9a9-2fed-6352-e437-fe76b65aa80d	10.0
.175.78:5446 F		
WIN-M3Q09A2PNFP-28728	0afa39df-d072-b05a-034a-aed74b7a39ee	10.0
.238.179:5446 F		
WIN-M3Q09A2PNFP-28295	c69f0430-df04-bf24-1294-b471a4a3f151	10.0
.2.207:5446 F		
WIN-M3Q09A2PNFP-15869	582d3642-9736-e87f-4e19-c22d7776ccab	10.0
.2.206:5446 Т		

Checking /commander/jgroups/commander:

0c32103a-d4e7-da91-f8c4-1fd8e125c156 ->

WIN-M3Q09A2PNFP-19713	0c32103a-d4e7-da91-f8c4-1fd8e125c156	10.0
.2.206:5447 Т		
WIN-M3Q09A2PNFP-14432	75afe3eb-7a04-d160-c8d0-d64f8ac3c796	10.0
.175.78:5447 F		
WIN-M3Q09A2PNFP-33530	1cf1795f-e658-3cae-c515-546082bffec9	10.0
.238.179:5447 F		
WIN-M3Q09A2PNFP-60743	cf8b4201-3fac-71c4-18da-0a885b3e1e61	10.0
.2.207:5447 F		

How to interpret ClusterInfoTool output:

- The nodes /commander/jgroups/hornetq and /commander/jgroups/commander contain information on these JGroups clusters:
 - commander for the ElectricFlow server cluster
 - hornetq for the HornetQ cluster
- The children nodes under each of the JGroups nodes represent the participating ElectricFlow servers in the cluster. Each child node entry is in this form:

```
<Logical Name> <UUID> <IP address>:<port> T|F
```

• The number of entries in both JGroups nodes should be the same, with matching IP addresses but with different port numbers and distinct logical names and UUIDs. The coordinator node in each JGroups cluster is identified with a 'T' against its entry.

Automation Platform Tasks

All the topics in this section are about a particular feature or function in the ElectricFlow automation platform. They are the same as the topics in the "Web Interface Help" section of the Help system on the automation platform UI.

From the automation platform UI, you can access the topics as follows:

All ElectricFlow Help topics above this directory are expanded, user-guide style Help topics about a
particular ElectricFlow feature or function. Direct access to these Help topics is from the Help system
Table of Contents. Although these topics are not directly linked within the ElectricFlow product, many
"page-specific Help topics contain links to one or more of these topics to provide more detailed
information quickly.

• The Help topics listed below this directory are "page-specific" Help topics, which means each of these topics contains specific information for the automation platform web page you are viewing. To access these Help topics while using ElectricFlow, click the **Help** link at the top-right corner on any automation platform web page. (The topics in this section of the AllHelp directory are listed singularly or in groups as is appropriate.)

Access Control

Use this GUI page to view and modify access privileges for a particular ElectricFlow object. Depending on the object where you want to set permissions, you will see that object's name as part of the page title (above the tables).

For example, if you clicked the Access Control link from a Project Details page, you will see the project name as part of the Access Control page title.

Reading and Using This Page

- This page displays one or more access control lists. The top list contains entries for the object itself
 (specified in the page title) and also identifies the object.
 For example, if the heading for the top list reads "Privileges for Procedure: buildAndTestAll," you are
 viewing access privileges for a procedure named "buildAndTestAll". Click the object name to view the
 main page for that object.
- Typically, you will see more than one list on the page. Each list below the first one contains privileges for
 an object that "contains" the objects above it.
 For example, a project contains all of its procedures and a procedure contains all of its steps. Privileges
 for the top-level object are determined by all the privileges in all of the displayed lists. The lists form an
 "inheritance chain" where it each object "inherits" permissions from the objects below it on the page.
- When a user attempts a particular operation on the object, ElectricFlow examines the lists on this page from top to bottom. If there is an entry specifying "deny" for the user (or a group containing the user) in the top list, access is denied. Otherwise, if an entry specifies "allow" for the user (or a group containing the user) in the top list, access is allowed. If access is neither allowed nor denied by the top list, ElectricFlow proceeds to the next list and processes it in the same way.
 - Note: If access is neither allowed nor denied by any list, ElectricFlow denies access.
- The inheritance mechanism makes it easy to control access for a large number of objects in a single place.
 - For example, project access control entries automatically apply to new objects created within the project. Each new object in the project has an empty access control list, but will inherit from the project.

Using the Links...

Use these links to add or increase Access Control for an object.

- Add User Use this link to add permissions for a specific user.
- Add Group Use this link to add permissions for a specific group, which means all users in that group would have the permissions allowed to the group.
- Add Project Use this link to set or redefine permissions for a project.

• Break Inheritance - Caution: Be very careful if you break inheritance! If you use the "Break Inheritance" action for any list, no additional inheritance occurs *below* that list and you no longer see other lists on this page. This action is useful if you want privileges for an object to be totally different than its containing object. If an object has no entries in its access control list and you break inheritance for that object, you make the object *completely inaccessible* - you will not even have the "Change Permissions" privilege, so you cannot restore inheritance. If this occurs, see your system administrator to restore inheritance.

· Actions:

- Edit Use this link to modify current permissions, but be careful if you modify permissions in an
 inherited access control list. Modifying inherited access control affects all other objects that inherit
 from the same list.
- o Delete Deletes the current privileges granted for that user, group, or project.

Privilege Definitions

The following four privilege types (for each ElectricFlow object) can be assigned *allow*, *deny*, or *inherit* permission.

Read - Allows object contents to be viewed.

Modify - Allows object contents (but not its permissions) to be changed.

Execute - If an object is a procedure or it contains procedures (for example, a project), this privilege allows object procedures to be invoked as part of a job. For resource objects, this privilege determines who can use this resource in job steps.

Change Permissions - Allows object permissions to be modified.

For more information, see the main Access Control Help topic. This Help topic also contains two examples that illustrate how you might use Access Control to increase ElectricFlow security.

Access Control - defining entries

Use this page to define access control entries selected from a filtered list of ElectricFlow users, groups, or projects, or by providing an explicit name.

Enter information in the fields as follows:

Note: The following instructions are generic. For example, if you chose Add User, the Filter field will be labeled "User Filter".

Field Name	Description
Radio selector	Choose between selecting "Search for existing" or select "Provide explicit name". Providing an explicit name bypasses the search and allows you to enter an arbitrary name.

Field Name	Description
Filter	Enter a partial name to retrieve a list of all users/groups/projects that include the partial name you supplied.
	[User/Group only] - Use this checkbox only if requesting external LDAP or Active Directory providers.
Include Inactive	If "checked," all available users or groups are returned.
	If not checked, only those users or groups known to the ElectricFlow server are returned, for example, users who have logged in or groups containing users who have logged in.

Click **OK** to retrieve your information.

If you entered filter criteria, the page will refresh with a list of matches. Select any row in this table to create an access control entry for that principal.

Column descriptions

Name - The name of the user/group/project.

Location - Only viewable for users or groups from non-local directory providers. The name combined with the location is the unique identifier for this user or group, local to the ElectricFlow server.

Privileges - create new or edit existing privileges

Use this page to create or modify an access control list entry. Each entry names a *principal*, which can be either an individual user or a group, and defines four privileges for the principal.

For each of the four privileges, **Read**, **Modify**, **Execute**, or **Change** Permissions:

- Creating a new entry or modifying an existing entry allows you to select either Inherit, Allow, or Deny access entries.
- on the server access control table, "Don't Care" replaces the "Inherit" option.

For more information about access control, see the Access Control Help topic.

Artifacts

This page displays all artifacts available on this ElectricFlow server.

Links and actions at the top of the table

- New Search Use this link to go to the Define Search page to filter the Artifacts page view.
- Create Artifact Use this link to go to the New Artifact page to create a new artifact.
- The "star" icon allows you to save this page to your Home page for quick access.

Column descriptions

Column Name	Description / Actions
Name	The name of the artifact—a system-generated name created by combining the Group Id and Artifact Key components. Note: Selecting an artifact name takes you to the Artifact Details page for that artifact.
Group Id	The user-generated group name for this artifact. This field may consist of alphanumeric characters, spaces, underscores, hyphens, and periods.
Artifact Key	The "key" component of the artifact name.
Description	The plain text or HTML description previously supplied for this artifact.
Actions	Edit - Use this link to go to the Edit Artifact page to modify the artifact on this row. Delete - Use this link to delete the artifact on this row.

Artifact Details

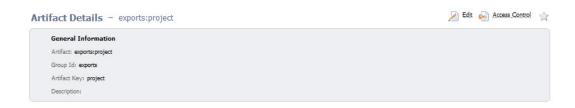
This page displays detailed information about an artifact.

Links and actions at the top of the table

- Edit Use this link to go to the Edit Artifact page to make any necessary changes for this artifact's description or artifact version name template.
- Access Control Use this link to set privileges for this artifact. For more information, see the Access Control Help topic.
- The "star" icon allows you to save this job information to your Home page.

General Information section

This section displays the artifact name (Artifact), Group Id, Artifact Key, and any Description previously supplied for this artifact.



The "tabbed" sections

The next section of tabs allows you to select the type of information you want to see.

Artifact Versions table

This table displays all artifact versions included in this artifact.



Column descriptions

Artifact Version - Click on an artifact version name to go to the Artifact Version Details page for that artifact version.

Description - The previously supplied description, if any, for this artifact version.

Actions

- o Click the **Edit** link to go to the Edit Artifact Version page.
- o Click the **Delete** link to delete this artifact version from the artifact.

Properties table

This table contains custom properties for this artifact.

- Click on a Property Name to edit that property.
- If a "folder icon" precedes a property name, it denotes a Nested Property Sheet.
- This section also provides Create Property, Create Nested Sheet, and Access Control links to add more custom properties or change or set privileges on this artifact.



Artifact - create new or edit existing artifact

To create a new artifact

Enter information in the fields as follows:

Field Name	Description
Group ld	Enter a group name of your choice for this artifact. A Group Id (groupId) acts as a namespace for grouping related artifacts. The idea is that artifact "sdk" for group "platform" can co-exist with artifact "sdk" in group "ui" without there being any name collisions. This field is limited to alphanumeric characters, spaces, underscores, hyphens, and periods.
Artifact Key	Enter an identifier of your choice for this artifact. This field is limited to alphanumeric characters, spaces, underscores, hyphens, and periods.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <dl> <i> <pre> <style> <tc> </td></tr><tr><td>Artifact Version Name Template</td><td>A template for the names of artifact versions published to this artifact. For example: \$[/myArtifactVersion/groupId]:\$[/myArtifactVersion/artifactKey]:\$ [/myArtifactVersion/version] Produces a name like: platform:sdk:1.0.0-36 Only absolute property paths and /myArtifactVersion-based paths are supported. If not specified, the "Artifact Version name template"server setting is used.</td></tr></tbody></table></style></pre></i></dl></div>

Click **OK** after filling-in the fields and to go to the Artifacts page to see your new artifact listed in the first column.

To edit an artifact

- Click the Access Control link at the top of the page to add permissions to this artifact.
- You can add a Description or modify your existing description.
- You can add an Artifact Version Name Template or modify your existing template.
- Click **OK** to save your changes and return to the Artifacts page.

Artifact Versions

This page displays all artifact versions available on this ElectricFlow server.

Links and actions at the top of the table

- **New Search** Use this link to go to the Define Search page to locate artifacts, repositories or other related information.
- The "star" icon allows you to save this page to your Home page for quick access.

Column descriptions

Column Name	Description / Actions
Name	The name of the artifact version. Select an artifact version name to go to the Artifact Version Details page for that artifact version.
Artifact	The name of the artifact. Select an artifact name to go to the Artifact Details page for that artifact.
Group Id	The name of the group where this artifact version is a member.
Artifact Key	The "key" identifier for this artifact version.
Version	The artifact version represented, by default, as major.minor.patch-qualifier-buildNumber. Note: You will not see all 5 artifact version components specified if only 3 components were defined.
State	The current state of this artifact version. Possible values are:available publishing unavailable.
Modify Time	The last time this artifact version was modified.
Actions	Edit - Use this link to go to the Edit Artifact Version page to modify the artifact version on this row.
	Delete - Use this link to delete the artifact version on this row.

Artifact Version Details

This page displays detailed information about an artifact version.

Links and actions at the top of the table

- Edit Use this link to go to the Edit Artifact Version page.
- Access Control Use this link to set privileges for this artifact version. For more information, see the Access Control Help topic.
- The "star" icon allows you to save this artifact version information to your Home page.

General Information section

This section lists current information about the artifact version. While most of the information is static, two items link to other pages:

Artifact - Click the artifact name to go to the Artifact Details page for the "owning" artifact.

Published By - Click this link to go to the Job Details page for the job that published this artifact version. If this artifact version was published by an end user outside of a job context, this field will display None.



The "tabbed" sections

The next section of tabs allows you to select the type of information you want to see.

Files

This section shows the collection of files included in this artifact version. The left-pane displays the artifact version directory structure and the right-pane displays files and subdirectories for the currently selected folder in the left-pane. This structure is similar to Windows Explorer.



Retrievers table

This table lists the 50 most recent jobs that retrieved this artifact version and includes links to the Job Details page for those jobs.



Dependent Artifact Versions table

This section displays dependent artifact version queries registered with the artifact version during a publish operation. These queries are evaluated during retrieval and the most current artifact version that matches each query is returned along with the primary artifact version.



Properties table

This table contains custom properties for this artifact version.

- Click on a Property Name to edit that property.
- If a "folder icon" precedes a property name, it denotes a Nested Property Sheet.
- This section also provides Create Property, Create Nested Sheet, and Access Control links to allow you to add more custom properties.



Artifact Version - edit an existing artifact version

This page displays an existing artifact version you can modify.

Links and actions at the top of the table

- Access Control Use this link to add or change permissions for this artifact version.
- The "star" icon allows you to save this page to your Home page for quick access.

To edit an artifact version

Field Name	Description / Action
Name	You can type-over the artifact version name in this field to change or modify the existing name.
Description	You can add a description or modify your existing description.

Field Name	Description / Action
Publisher	This is the job that completed the publish operation. Click this link to go to the Job Details page for this job. If this artifact version was published outside of a job context (for example, from ectool), the value for this field is "None".
State	An artifact version can be in one of three states: available, unavailable, and publishing. If the artifact version is:
	 in the Publishing state, the artifact version cannot be retrieved, but it will transition to the Available state when the publishing operation is complete. The web interface does not allow you to change the state if it is currently Publishing.
	 in the Available state, the artifact version can be retrieved. To confirm this state, the checkbox is "checked," and unchecking this box makes the artifact version unavailable.
	in the Unavailable state, the artifact version cannot be retrieved. "Checking" the box changes the artifact version state to Available. "The control of the control

_

_

_

Repositories

This page displays all artifact repositories available to this ElectricFlow server.

Links and actions at the top of the table

- Add Repository Use this link to go to the New Repository page to add another artifact repository.
- The "star" icon allows you to save this page to your Home page for quick access.

Column descriptions

Column Name	Description / Actions
Repository Name	The name of the artifact repository. Select a repository name to go to the Edit Repository page if you need to modify that repository. Note: If you have multiple repositories and want to change the search order for retrieving artifacts, drag the icon (in the far left column) up or down to reposition an artifact to the order you need.
Zone	The name of the zone where this repository resides.
URL	The server URL is in the form protocol://host:port/. Typically, the repository server is configured to listen on port 8200 for https requests, so a typical URL looks like https://host:8200/
Enabled	Select this checkbox to enable or disable this artifact repository. By disabling this repository, users will not be able to retrieve any artifacts stored in this repository.
Description	A plain text or HTML description previously supplied for this repository.
Actions	Delete - Use this link to delete the artifact repository on this row.

Repository - create new or edit existing repository

To create a new repository

Enter information in the fields as follows:

Field Name	Description
Name	Enter any name of your choice for this repository.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> <</td></tr></tbody></table></style></pre>

Field Name	Description
URL	The server URL is in the form protocol://host:port/. Typically, the repository server is configured to listen on port 8200 for https requests, so a typical URL looks like https://host:8200/.
Zone	Use the drop-down menu to select a zone for this repository. If no zone is selected, this repository will reside in the <i>default</i> zone.
Enabled	Select this checkbox to enable this repository. If this repository is disabled ("unchecked"), users will not be able to retrieve any artifacts stored in this repository.

Click **OK** after filling-in the fields and to return to the Repositories page to see your new repository listed in the first column.

To edit a repository

You can:

- Use the Access Control link at the top of the table to set permissions on this repository
- Change the repository name
- Add a description or modify the existing description
- · Modify the URL
- · Select a different zone.
- · Enable or disable the repository

Click **OK** after making any modifications and to return to the Repositories page.

Database Configuration

If you choose not to use the ElectricFlow built-in database, use this page to configure your MySQL, SQL Server, or Oracle database to communicate with ElectricFlow.

Note: The ElectricFlow built-in database is not recommended for production use.

Fill in the fields as follows:

Field Name	Description
Database Type	Select your database type from the drop-down menu. If "Built-in" is selected, no other information is required.

Field Name	Description
Database Name	Enter your database name.
Host Name	Enter the host name for your database server.
Port	Use the default port or enter a new port number.
Database Credentials	User Name–Accept the default "commander" user name or enter the appropriate user name for your database. Password–Enter the database password.

IMPORTANT: When you set the SQL server as the database, you can only use a non-named database.

Click Save and Restart Server after entering information in all fields.

Note: You may need to consult with your Database Administrator if you do not have all of the information required on this web page.

For more information, consult the ElectricFlow Installation Guide.

Defect Tracking Configurations

This page displays Defect Tracking configurations known to the ElectricFlow server—ElectricFlow integrates, using plugins, with numerous defect tracking systems.

Link at the top of the table

Use the Create Configuration link to create a new Defect Tracking configuration.

Column descriptions

Column Name	Description
Configuration Name	The name you provided to the configuration you created.
Description	This column describes the Defect Tracking configuration type.
Plugin	The name of the plugin where the Defect Tracking integration resides.
Actions	Edit - Click this link to modify an existing configuration.
	Delete - Click this link to remove the configuration.

Defect Tracking - create new or edit existing configuration

Use this page to define a new defect tracking configuration or modify an existing defect tracking configuration. A configuration is a collection of properties that define how ElectricFlow communicates with a particular defect tracking system.

After creating a defect tracking configuration, your entry will appear in the table on the Defect Tracking Configurations web page— to see this web page, select the Administration > Defect Tracking tabs.

To create a defect tracking configuration

From the Defect Tracking Type drop-down menu, select your defect tracking system—ElectricFlow integrates with numerous defect tracking systems. Each integration was created using plugin technology.

Click the **Submit** button and a set of fields is displayed to enter information to create your configuration.

Enter information in the fields as follows:

Field Name	Description
Configuration Name	Enter a unique name for this defect tracking configuration—any name you choose.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pr> <pr> <pr> <pr></pr></pr></pr></pr>
Login As	User Name - This is the name ElectricFlow needs to use to communicate with your defect tracking system. For example, you may be using a special "read-only" user name similar to "Build" for your user name.
	Password - This is the password for the specified User Name.
	Retype Password - Type the password again.
	URL - This is the URL to use to connect to a JIRA server. For example, you might use something similar to: http://10.10.10.10.8080 or http://yourJIRAserver

Click Submit to save your information.

Using the defect tracking integration

Go to Projects > select a project > select a procedure.

- To create a New Step for defect tracking, select the Plugin link.
- In the Choose Step panel, select Defect Tracking from the left pane, then select the defect tracking system you configured.
- The right-pane now shows the types of steps available for your configuration. Select the step you need and automatically go to the New Step page.
- On the New Step page, notice the Subprocedure section now contains the defect tracking integration you configured and the step you chose.
- Enter the remaining information to create your defect tracking step.

See the Defect Tracking Help topic for more information.

To edit an existing defect tracking configuration

Modify any of your defect tracking configuration information by "typing-over" any previously entered information.

Click **Submit** to save your modified defect tracking information.

Defect Tracking Reports

This is a report of URLs to JIRA defects associated with the ElectricFlow job.

The report was compiled by searching job properties for JIRA defect IDs then querying the configured JIRA server for the current defect state.

For more information about Defect Tracking, see the main Defect Tracking Help topic.

If you would like to create a defect tracking configuration, go to Administration > Defect Tracking and click the **Create Configuration** link to see the New Defect Tracking Configuration web page. Access the Help link on that page for more information.

Email Configurations

This GUI page displays all previously configured email configurations.

- Click on a Configuration Name to go to the Edit Email Configuration web page to modify an existing email configuration.
- Click on the Add Configuration link to create new email configuration.
- Action column This column contains two links to Copy or Remove an Email Configuration.

Notice this page also lists the email server name and who the email notification was from.

Email Configuration - create new or edit existing email configuration

Use this web page to specify an email server to ElectricFlow. If you do not create an Email Configuration, you will not be to send Email Notifications to individuals or groups.

If you have multiple users or groups in remote locations that use a different mail server, create additional Email Configurations to accommodate those locations if they need to receive ElectricFlow notifications.

To create a new email configuration

Enter information in the fields as follows:

Field Name	Description
Name	Enter a unique name to identify the Email Configuration.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </td></tr><tr><td>Mail Protocol</td><td>Either SSMTP or SMTP</td></tr><tr><td>Mail Host</td><td>The name of the mail server.</td></tr><tr><td>Mail Port</td><td>The port number for the mail server, but may not need to be specified. Protocol software determines the default value (25 for SMTP and 465 for SSMTP). Specify a value for this argument when a non -default port is used.</td></tr><tr><td>Mail From</td><td>An individual name or a generic name like "ElectricFlow" or "build."</td></tr><tr><td>Mail Username</td><td>An individual name or a generic name like "ElectricFlow."</td></tr><tr><td>Mail Password</td><td>The password for the mail user name account.</td></tr></tbody></table></style></pre>

Click Save after filling-in all fields.

Click the **Test** button to run a test to make sure your email configuration works. Enter an email address in the pop-up box and click **Send**.

To edit an existing email configuration

Use the Access Control link at the top of the page to add or change existing permissions.

You can modify any of the information already supplied or add new information.

Click **OK** after completing your edits.

Email Notifier - create new or edit existing email notifier

Use this page to set up email notifiers.

Note: You must have already set up an email configuration so ElectricFlow knows which email server to use when sending any notifications. If you have not configured an email server to communicate with ElectricFlow, click the **Administration** > **Email Configurations** tabs to do this configuration now.

For jobs and job steps, two types of email notifiers:

- On Start Notifier Sends an email when a job or job step starts.
- On Completion Notifier Sends an email when a job or job step completes.

For workflow states, three types of email notifiers:

- On Enter Notifier Sends an email when the state becomes the workflow's active state.
- On Start Notifier Sends an email after the state's subjob or subworkflow starts. If no subjob or subworkflow is defined for that state, the notifiers will not be sent.
- On Completion Notifier Sends an email after the state's subjob or subworkflow completes. If no subjob or subworkflow is defined for that state, the notifiers will not be sent.

Using email notifiers:

- Attaching an email notifier to a procedure results in a corresponding email notifier associated with the job that is created when the procedure is executed.
- Attaching an email notifier to a procedure step results in a corresponding email notifier associated with the job step created when the procedure referencing the procedure step is executed.
- Attaching an email notifier to a state definition results in a corresponding email notifier associated with the state that is created when the workflow is executed.

To create a new email notifier

Enter information in the fields as follows:

Field Name	Description
Name	This name can be an arbitrary text string.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with httml httml tags. The only HTML tags allowed in the text are: <a> <a> <a> httml tags. The only HTML tags allowed in the text are: <a> <a> httml httml https://www.channelle.com/html htt
Condition	Use the pull-down menu to select the type of condition you need for this email notifier. Edit the auto-supplied condition in the text box or add a completely new script for your purpose. The condition specifies whether the notifier should send a message depending on the result of a property expansion. If the result is empty or non-zero, the message is sent. If the result is "0", the message is not sent

Field Name	Description
	Use the pull-down menu to select from a list of global, ready-to-use formatting templates. Depending on the type of email notifier you are creating, the available template choices in the drop-down menu will be different. Make sure the content is formatted correctly, i.e., no illegal characters or spacing.
	To customize your template, edit the auto-supplied text in the Formatting Template text box or you can add a completely new script for your purpose. Any edits made in this text box will not be saved to the global template.
	Note: See the "Email notifier templates" section at the end of this Help topic for a list of available templates and instructions for making global modifications to these templates.
	To create a custom template, the basic structure is:
	zero or more email header lines
	blank line
	message body
Formatting Template	The template undergoes property expansion in a job, jobStep, or state context. In addition to normal property paths, an additional event object is accessible via /myEvent/.
	The following example is a sample formatting template for a job:
	Subject: Job '\$[jobName]' from procedure '\$[procedureName]' \$[/myEvent/type] - ElectricFlow notification cc: owner@example.org
	bcc: admin@example.org
	ElectricFlow Notification - from email notifier
	'\$[/myEvent/notifier]'
	Job '\$[jobName]' \$[/myEvent/type] at: \$[/myEvent/time]
	Project: \$[projectName]
	Procedure: \$[procedureName] Started: \$[start]
	Completed: \$[finish]
	Directory Name: \$[directoryName]
	Launched by User: \$[launchedByUser]
	Outcome: \$[outcome] From Codo: \$[orrorCodo]
	Error Code: \$[errorCode] Error Message: \$[errorMessage]
	Titor hessage. Y[errormessage]
	This example demonstrates how you can send this notification using CC or

Field Name	Description
	BCC fields in addition to specified Destinations. Because headers are interpreted by SMTP, normal email rules apply.
Email Configuration	Click inside this field or start typing to bring up a list of possible email configuration names. An email notifier that does not specify an email configuration will use the configuration named 'default' if it exists.
Destinations	A space-separated list of valid email addresses, email aliases, or ElectricFlow user or group names, or a property reference that expands into such a list, or you can enter an LDAP DL name (group name).

Click **OK** to save your email notifier configuration.

To edit an existing email notifier

Modify or add information to any of the fields and click **OK** to save your changes.

Email notifier templates

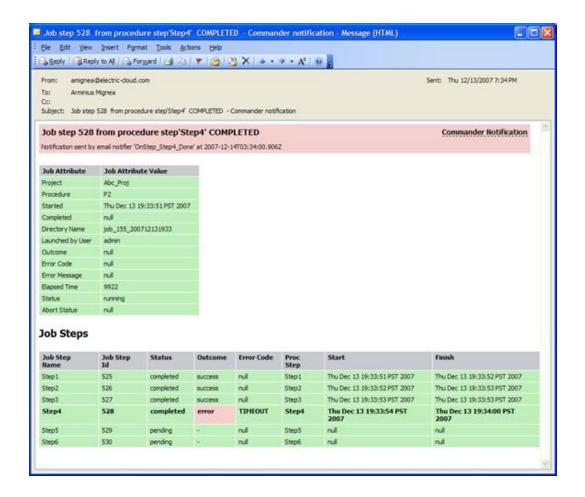
The table below lists global email notifier templates for your use. Each template name is a link to an example of the template output and the script.

Global templates are stored as property sheets under the "/server/ec_notifierTemplates" property sheet. You can modify template properties or add new global templates if you have modify privileges on the server property sheet. An email notifier will be updated if changes are made to its global template.

Template Name	Forma- t	Usage	Description
Job Summary text notification	Plain text	Procedur e	Formatting template for Procedure starting/completion notifications using the long form of property names (for example, \$[/myJob/jobName])
Step summary text notification	Plain text	Procedur e step	Formatting template for Procedure Step starting/completion notifications using the long form of property names (for example, \$[/myJobStep/jobStepId])
State summary text notification	Plain text	State	Formatting template for State notifications using the long form of property names (for example, \$[/myWorkflow/workflowNam e])
Job summary HTML notification	HTML	Procedur e	HTML formatting template for Procedure starting/completion notifications
Step summary HTML notification	HTML	Procedur e step	HTML formatting template for Procedure Step starting/completion notifications displaying information about the Job and the Job Step

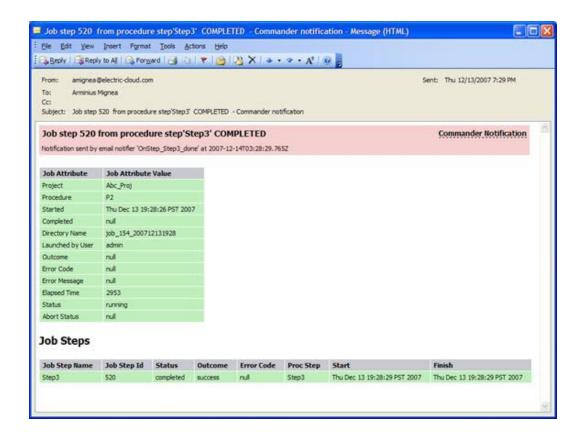
Template Name	Forma-	Usage	Description
All steps HTML notification	HTML	Procedur e step	HTML formatting template for Procedure Step starting/completion notifications displaying information about the Job and many Job Steps
Approval request HTML notification	HTML	State	HTML formatting template for State notifications displaying information about the workflow and all of its states
Workflow summary HTML notification	HTML	State	HTML formatting template for State notifications providing the recipient with links to view or take a manual transition

$Email\ Notifier\ Template-Html_JobStepTempl_AllSteps.txt$



- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

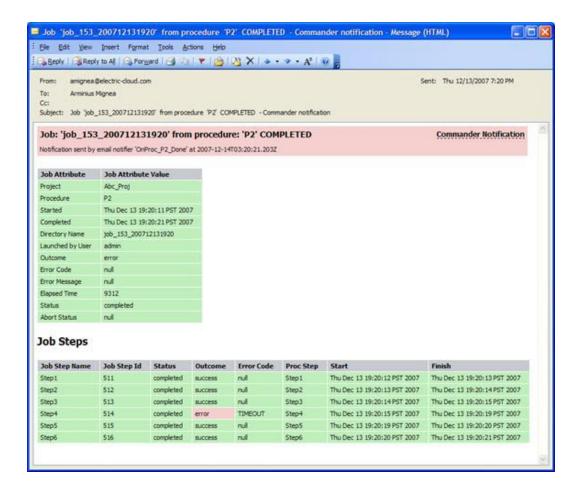
Email Notifier Template - Html JobStepTempl SingleStep.txt



Note: In the following script, myJob.steps and myJob.jobSteps return an array of step names.

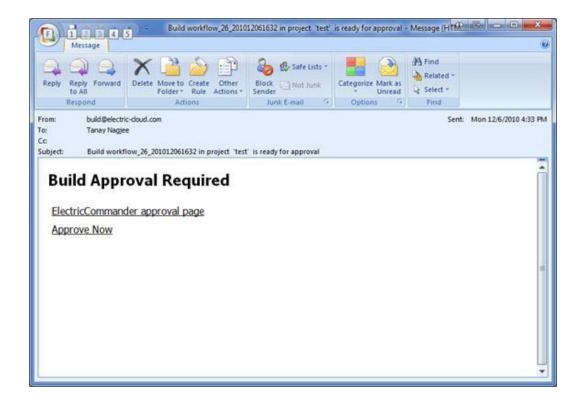
- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

Email Notifier Template - Html_JobTempl.txt



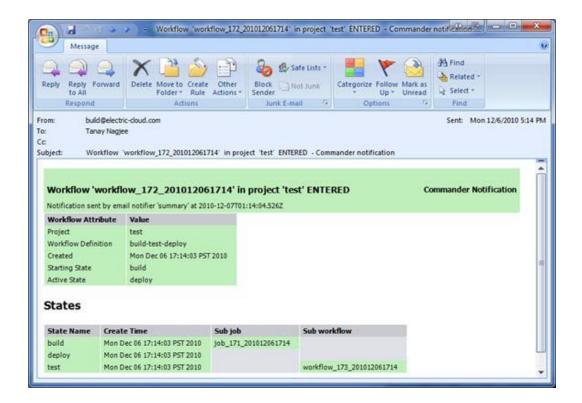
- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

Email Notifier Template - Html StateTemplate ApproveWorkflow.txt



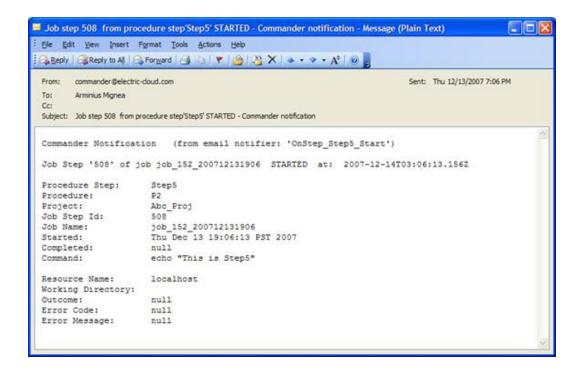
- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

Email Notifier Template - Html_StateTemplate_FullWorkflow.txt



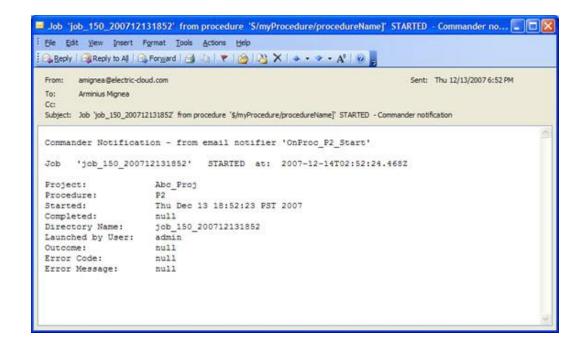
- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

Email Notifier Template - JobStepTempl_FullProps.txt



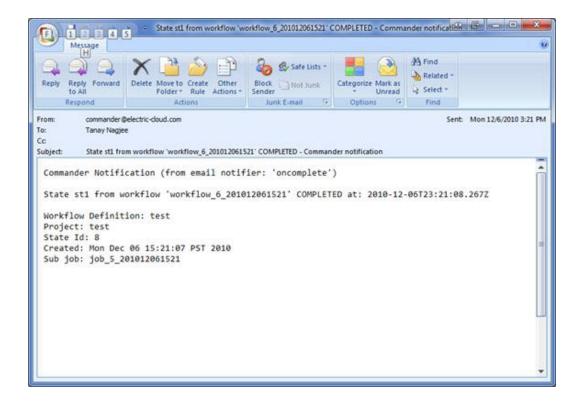
- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

Email Notifier Template - JobTempl_FullProps.txt



- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

Email Notifier Template - StateTemplate FullPropertyPaths.txt



- 1. Go to the corresponding template available in <installDir>\src\samples\notifier
- 2. Copy the script into a text editor to edit it for your purpose
- 3. "Copy and paste" your version into the Formatting Template field.

Event Log

This page displays a log of events generated anywhere in the system, including jobs and workflows.

- To see only events for a single workflow, select the Workflows tab, then a workflow Name to go to the Workflow Details page and click the **View Log** link at the top of the page.
- To see only events for a single job, select the Jobs tab, then the Job name to go to the Job Details page and click the **View Log** link at the top of the page.

Note: When running a job in a cluster, the name of the server managing the job is recorded in the joblog. If the management of the job moves to a different node (if the current node crashes for example) the change is logged into the job log as well.

To see only events for a specific object, select the Search tab to go to the Define Search page.

For example:

You might want to select the Object Type, "Log Entry", then click the **Add Intrinsic Filter** link. Select the down-arrow where you see "Container" auto-populated and select "Container Type". Use the "equals" operator, then select the next down-arrow to choose an object. Click **OK** to start the search.

Note:

From the Administration tab, the default view for this Event Log page is the warning (WARN) level.

For workflow and job event logs, the default view from their respective pages is the information (INFO) level.

Column descriptions

Time - This is the time the event was logged.

Severity - Severity can be either INFO, WARN, or ERROR. Use the pull-down menu at the top of the table to change between these views.

User - The user associated with the event—the user of the session at the time the event was generated. Clicking on a project link in this column takes you to the Project Details page for that project.

Container - Typically, this is the type and name of the workflow or job with a corresponding ID. Clicking on a link in this column takes you to the Job Details page for that job or workflow.

Subject - The object associated with the message. Clicking on a link in this column takes to you the source for this object.

Message - The message text will either be informational or display the warning or error message. The message may contain important information about a resource or workspace issue.

Configuring the Event Log

Because log entries can accumulate in large numbers, you can configure the Event Log for auto-deletion. By default, Event Log retention is 30 days. The ElectricFlow server automatically marks old log entries for deletion and the background deleter cleans out old log entries along with its other tasks.

To change the default settings, go to Administration > Server > Settings.

Modify the following two fields:

Event log retain time - Default is 30 days (in "minute" units).

Note: Setting the "retain" period to "0" disables the automatic deletion mechanism, allowing you to use an external cleanup script to implement a custom cleanup policy.

Maximum background delete delay - Default is 3600 seconds (or one hour).

Click **OK** to save your changes.

Automation Platform Home Page

Job Configurations

Shortcuts

Jobs Quick View

Reports

Overview

This page provides a convenient console for running jobs and viewing results.

- This web page is your ElectricFlow Dashboard for tracking project health.
- You can customize the page to display only the jobs that interest you and to provide one-click access to the jobs you run frequently. Each of the Home page sections can be customized.
- The Reports section allows you to see thumbnail report images that are updated each time the report is generated.
- This page provides shortcuts for quick access to pages you use most frequently inthe automation platform or anywhere on the web.

IMPORTANT: If you are using or plan to use the EC-Homepage plugin to share your Home page sections, you will not see some links normally available to you. For example, if a particular Jobs Quick View category is shared, you **cannot** Delete or Modify that category. You can perform these functions only if the category belongs to you alone.

For information on sharing your Home page configuration, see the "Home page configurations" section in the "Customizing the ElectricFlow Platform UI on page 444" Help topic.

Job Configurations

Procedures in ElectricFlow can contain complex sets of parameters— making it difficult to remember the correct parameters for a particular situation and tedious to re-enter those parameters every time the procedure is invoked. Job Configurations provide a one-click solution to this problem. When you create a job configuration, you enter all the information needed to run a procedure, including parameters and/or a credential. All job configurations are displayed here on the Home page. Invoke a particular configuration by clicking its name in the Job Configurations section.

Create Job Configurations three ways

- On the Home page, create a job configuration from "scratch" by clicking the **Create** link in the Job Configurations section.
 - In the Create Configuration pop-up menu, select the project and procedure you want to use for creating this configuration.
- From the Job Details page for a previously invoked job, click the **Save Configuration** link at the top of the page.
 - Your saved job configuration is displayed on your Home page.
- From the Edit Schedule page, click the **Save Configuration** link at the top of the page. Your saved configuration is displayed on your Home page.

Shortcuts

Use shortcuts to save frequently visited ElectricFlow web pages, so those pages are immediately accessible. You can create a shortcut to any page on the web also.

Create Shortcuts two ways

- Mouse-over the "star" icon at the top of any automation platform page and click "Add current page" to
 add that page to the shortcut list. Mouse-over the star icon again and click "Remove current page" to
 remove the shortcut for that page. The star icon is yellow for pages saved as a shortcut and gray for
 those pages not saved as a shortcut.
- Click the Create link in the Shortcut section and provide a name and URL to create a shortcut.

To modify or update a shortcut, click the **Edit** link adjacent to the shortcut you want to change.

Shortcuts can be accessed conveniently from any automation platform web page. Mousing-over the star icon displays a list of shortcuts saved by the current user. Click on a shortcut name to view the page.

Note: The Shortcut section may contain static entries that cannot be deleted. And if you "share" your Home page, the Edit link will not be available for shared items.

Jobs Quick View

Perhaps only a few jobs on this server are of interest to you. For example, you may care about jobs you launch manually and official builds for the products you work on, but you may not care about production builds for other products or personal jobs for other users.

The Jobs Quick View allows you to define job categories that are interesting to you.

The Home page displays the most recent jobs in each category, and you can easily click-through to get more details about any of those jobs. For example, clicking on a job name takes you to that job's Job Details page.

Create a job category

- Click the Add Category link in the Jobs Quick View section.
 After creating a category, results are displayed on the Home page. Clicking the Details link displays a summary to the right of the category. In addition to job status and other diagnostic information, the summary displays running steps and failed steps (containing errors and warnings).
- Click the **Modify** link to edit the Jobs Quick View category or the **Delete** link to remove that job from the Jobs Quick View category.
- Click on the **Details** link to see job summary information—the summary remains visible, regardless of mouse location, until you click somewhere else on the page.

Note: If you "share" your Home page, the Modify and Delete links will not be available for shared items.

Reports

You can configure reports you would like to see on a regular basis and display a "thumbnail" report graphic in this section.

- Click the Add Report link to go to the New Reports page.
 After filling-in the information on the New Reports page, you will see a thumbnail view of your report (after it runs) on your Home page.
 - **Note:** If the drop-down menu on this page is empty, click the **Help** link on the New Report page for more information.
- Click the Collapse/Expand (+/-) box to see the full thumbnail report or just the report title.

- Click the Rename link if you need to rename your report.
- Click the **Remove** link if you need to remove this report from your Home page.

Note: If you "share" your Home page, the Rename and Remove links will not be available for shared items.

Job Configuration

To create a new job configuration

Enter information in the fields as follows:

Name - Type a name you choose for the job configuration.

Procedure - Displays the current project and procedure. Click **Change** to use the pop-up menu to select a new project/procedure.

Parameters

Depending on the Project and Procedure you selected, a set of Parameters is displayed.

Enter all parameter values or just the required values if differentiated.

Advanced

Priority:

Use the drop-down menu to select the priority for this run procedure. Available priorities are: low, normal (default), high, or highest.

- You can select the job's priority here, on Run Procedure web page, or when you schedule the
 procedure to run at a regular time or interval. When a job is launched, its priority cannot be
 changed.
- Priorities take effect when two or more job steps in different jobs are waiting for the same
 resource. When the resource is available, it will be used by the job step that belongs to the job
 with the highest priority. If the priority level is the same, the resource will be used by the job step
 that belongs to the job with the lowest job ID number. If the job steps are in the same job, the
 resource will be used first by the step with the lowest job step ID number.
- If using ectool, the priority can be set by passing "-priority=<low|normal|high|highest>" to a runProcedure, createSchedule, or
 modifySchedule operation.

Note: To raise the job priority level above "normal," the user who launches the procedure must have execute permission on the system-level access control list for priorities.

Impersonation:

Click one of the three "radio" buttons to determine the default login user to use when the configuration is run.

Use pre-defined credential - If you select this option, ElectricFlow looks for a credential first in the
top-level procedure and then in its project. If it finds a credential, it uses the login from that
credential as the default for the job. If there is no credential in either place, by default, the job's

steps run as the user under which the agent for the step is running.

- Use specific credential If you select this option, you can choose a particular credential from those defined for the project containing the procedure the login user from that credential is used as the default for the job.
- Use a specific user If you select this option, you can enter a login name. When the configuration is run, you will be asked to enter the password for that account. The specified login will be used as the default for the job.

For more information on credentials and user impersonation, see the Credentials and User Impersonation Help topic.

Click **OK** when all selections are complete. Return to the Home page to see your new job configuration displayed.

To edit an existing job configuration

Change or modify any of your previously specified information.

Click **OK** when all selections are complete.

Return to the Home page to see your edited job configuration displayed.

Shortcuts

Use this page to create a New Shortcut or Edit an existing Shortcut.

Shortcuts appear on your Home page and provide quick access to other web pages, either in ElectricFlow or elsewhere.

- To create a new shortcut, type-in a shortcut name and the URL.
- To edit an existing shortcut, re-type the new shortcut name or URL (or both).

The easiest way to create a shortcut to an ElectricFlow automation platform page, such as the page for a particular project or procedure, is to navigate to that page and then click the "star" icon in the upper-right corner. This action automatically creates a shortcut and the star turns yellow to indicate a shortcut is in effect for the page.

If you click on the star icon again, the shortcut for that page is canceled.

Click **OK** to save your new or edited information and see your new/edited shortcut on the Home page.

Jobs Ouick View

A quick view category provides a convenient way to select a group of jobs that interest you, such as "all production builds for product xyz version 3.6" or "all jobs I invoke manually." ElectricFlow displays the most recent jobs for each category on your Home page so you can review them easily.

To create a new jobs quick view category

Enter information in the fields as follows:

Name-Type a name for your new job category.

Number of Jobs—Type the maximum number of jobs you want to see in your quick view summary. The summary displays the most recent applicable jobs.

Include Last Success–Select this check box if you want to see the last successful job in this category, even if it was not one of the most recent jobs.

Two Filter categories:

- Intrinsic filters—These filters provide a convenient way to access certain well-defined fields for jobs.
- Custom filters—These filters allow you to access a much broader range of values, including custom
 properties. Any values accessible through an intrinsic filter can be checked using a custom filter also
 (though not as conveniently).

Each filter specifies three values: the name of a particular property to check, a value it should be compared with, and relational operators such as "equals" or "less than." For example, if you select "Outcome" in an intrinsic filter with operator "not equals" and value "Success," the filter selects jobs that did not complete successfully because they had errors or warnings.

For an intrinsic filter, select a property by choosing one of several predefined properties from a list. For a custom filter, enter a property name. This name can be either the name of an intrinsic property for the job (such as "status" or "procedureName"), or the name of a property defined in the global property sheet for that job, or the name of a parameter for that job.

For the relational operator, select one of the values in the list. The behavior of the relational operators is determined by the underlying database. In most cases, comparisons are done using string comparison; operators such as "less than" may not be useful on numbers because "12" is considered less than "9". For some built-in properties, the database treats values according to a specific type such as number or date, so comparisons are implemented in the way appropriate for that type.

The relational operator "like" invokes an SQL wildcard comparison where "%" is the wildcard character. For example %test% matches "first test" and "tests failed" but not "Test 44".

Note: Job Quick View filters allow property path references. However, property path references are not allowed in search parameters on the Search page.

Click **OK** after making your selections to save your new job category and see it on the Home page.

To edit an existing jobs quick view category

Keep or change any of the previously entered information.

See the Filter information above to add or change existing custom or intrinsic filters.

Click **OK** after making your selections to save your modified job category and see it on the Home page.

New Report

Enter information in the following fields to define a new report to display on your Home page.

Field Name	Description
Project	The name of the project that contains the report you want to view.
Report	The report name you entered in the Report Title field when you configured this report. If you did not specify a "thumbnail" view when you defined your report, you will not see it here.
Title	This is the report title you want to see on your Home page to identify this report.

Click **OK** to continue.

After the report runs, the report you specified will be displayed on your Home page and updated whenever the report is generated again.

To populate the drop-down menu

- Go to the Project Details page for the project containing the reports you want to view.
- Select the Report tab.
- Click the Create Report link.
- Check the **Create thumbnail?** check box, which enters this report's name in the New Reports drop-down menu.

Note: All reports, *except* the Procedure Usage report, are available as "thumbnail" reports on your Home page.

Jobs

This page displays all jobs in the ElectricFlow system, both running and completed jobs. You can:

- Sort the Job, Status, Elapsed Time, and Start Time columns by clicking on the column name.
- Define a search to see only the jobs you choose.

Links at the top of the table

- **New Search** click this link to go to the Define Search page to select a specific set or range of jobs to view.
- Edit Search (available after a "search" is defined) click this link to redefine your search.
- Refresh Results click this link to see updated status for running jobs.

Column descriptions

Job - click on Job to sort the column alphabetically or click on a job name to go to that job's Job Details page.

Status - click on Status to sort this column by Running, Success, Warning, Error, or Aborted.

Priority - displays the job's priority set by a "run procedure" command or by the job's schedule.

Procedure - in this column you can click on a project name to go to the Project Details page or click on the procedure name to go to the Procedure Details page for that procedure.

Launched By - click on a name in this column to go to the page for the schedule (Edit Schedule page) or the user that ran the job.

Elapsed Time - click Elapsed Time to sort the time values from longest to shortest time, or the reverse.

Start Time - click Start Time to sort this column from the start time of the first job to the start time of the most recent job, or the reverse.

Actions - use this column to Abort a running job or Delete a completed job.

Aborting a job requires the execute privilege on the job—not just the modify privilege.

When you click Abort, the following dialog box appears:



Select either **OK** or **Force** to abort the job, or **Cancel** if you change your mind.

 Deleting a job on this page causes removal of job information from the ElectricFlow database, but information in the job's on-disk workspace area is not affected. You must delete workspace information manually.

Tips

If you would like to select a subset of the jobs, use the **Search** link at the top of the page.

From your Home page, you can add a specific "Search" for quick access on a regular basis.

Also notice: If you use RSS, an active RSS icon is provided in Windows Explorer on this page for your convenience. If you use Firefox, click **Bookmarks** > Subscribe to this page to display the feed. You can then add the feed URL to a viewer of your choice.

Job Step Details

This page displays detailed information about a job step.

Links and actions at the top of the table

- Access Control Use this link to set privileges for this job step. For more information, see the Access Control Help topic.
- "star" icon click this icon to add this page to your Home for quick one-click access in the future.

Summary section (at the top of the page)

Notice [in the screen example below] that the name of the job step you are viewing is adjacent to the page title, Job Step Details.

The Summary section provides:

- The job's Start, Wait (waiting for a resource), and Elapsed times
- You will see wait times specified if your job steps were restricted by resource, workspace, a
 precondition, or license availability.
 - For example, you will see the license wait time if a job step could not run because your license-allowed maximum concurrent step limit was met or exceeded, meaning no steps will run until the number of concurrent steps is reduced.
- A General Information section with links to get you to the job step's Job Details, Project Details, Procedure Details, or Edit Step pages.
- If the job created any custom reports, they appear in the summary area to the right of General Information. Reports on this page are for command steps only.

Reports

To create a special-purpose HTML report for a job: In the top-level of any of the job's workspaces, add a file that contains the word "report" or "Report" and ends in the .html file extension. The link is automatically created when the Job Details page is displayed. The job's workspace must be accessible to the web server. The link will be in the Reports section, within the Summary section.



Links

This page can include a Links section also, which will appear on the far right-side of the Summary section.

To add a link to this section:

• Create a link to any file:

To add a link, run a command in any job's job step using this format (this works for both local

[disconnected] and non-local workspaces):

```
ectool setProperty "/myJobStep/report-urls/<myReportName>"
"/commander/jobSteps/<jobStepId>/<artifactDirectoryName>/<file-path>"
```

Example command:

```
ectool setProperty "/myJobStep/report-urls/Test Report" "/commander/jobSteps/
$[/myJobStep/jobStepId]/testreport/index.html"
```

Note: If you are using this format (from ElectricCommander 4.2 and earlier):

```
ectool setProperty "/myJob/report-urls/<myReportName>" "/commander/jobs/<jobId>/
<workspaceName>/<artifactDirectoryName>]/<file-path>"
```

Example command:

```
ectool setProperty "/myJob/report-urls/Test Report" "/commander/jobs/$[/myJob/jobId]/
$[/myJobStep/workspaceName]/testreport/index.html"
```

It will continue to work onlyfor non-local workspaces. We recommend, however, that you change the command to use the newer format.

Create a link to any directory:

To add a link, run the following style command in any job's job step:

```
ectool setProperty "/myJobStep/report-urls/Main Job Workspace"
"file:///WinStor2/scratch/chron55build/$[/myJob/jobName]"
```

Note: Links to a directory automatically work in Internet Explorer, but if using Firefox, local links are disabled unless the default security policy is modified or an extension is used. Refer to http://kb.mozillazine.org/Links to local pages don%27t work for more details.

The "tabbed" Sections

The next section of tabs allows you to select the type of information you want to see. You may see six or fewer tabs, depending on the type of job step you are viewing.

Child Steps table

Note: The Child Steps table is available for subprocedure and broadcast steps only.

- Click on a Step Name to go to that step's Job Step Details page.
- Click the Log "icon" to see that job step's log information.
- Status This column provides the job's step status.
 - Aborted The step or job was aborted.
 - o Canceled The step did not run because the job was aborted or a previous step was aborted.
 - Error The command failed or postp detected an error in the output.
 - Running The step is currently running or a step inside the subprocedure call is running.
 - Skipped The step did not run because the run condition evaluated to "false".

- o Success The step completed with no errors or warnings.
- Waiting for Resource or Workspace The step is waiting to run as soon as a resource or workspace is available.
- Warning postp detected a warning message.

Note: The Status field can be over-ridden by postp to include more specific information. postp information may be linked directly to diagnostic information—also available by selecting the Diagnostics tab.

- Click on a Resource name to go to the Edit Resource page.
- Click the **Edit** link in the Action column to go to the Edit Step page.

General table

This section displays basic information including the step ID, step name, create and end times, and whether the step invoked a subprocedure or ran a command.



Diagnostics table

If the job step or child step (for subprocedure and broadcast steps) generated warning or error messages, you will see those messages here. This section contains one or more diagnostics, each of which is a portion of a step's log file that was identified as "interesting" by the step's postprocessor. Typically, each diagnostic relates to an error or warning detected by the postprocessor. All diagnostics for a particular job step appear together.



Parameters table

This section displays parameter values supplied to the top-level procedure when the job step was invoked and has these columns:

- Name
- Value

Note: This table is available for steps with child steps only.



Properties table

This section displays information computed and/or used by the job step while it is running. After the job step completes, these properties are normally read-only.

- · Click on a Property Name to edit that property.
- If a "folder icon" precedes a property name, it denotes a Nested Property Sheet.
- This section also provides Create Property, Create Nested Property, and Access Control links to
 update or enter additional information for the next time this job step is run. For more information on
 properties, see the Properties Help topic.

Notifiers table

If you have designated certain persons or groups to be notified of details for specific job steps, this section displays who receives an email notification, the type of notification, any condition, and a description [if supplied]. If a notifier is available, you can click on its name to view details or make modifications.

Job Step Time and waitTime Properties Explained

start - The time when this job step began executing.

When the ElectricFlow server assigns a resource to run the step, it issues a <runcommand> request containing the expanded command body to that resource's agent. The "start" time is when that <runcommand> is sent to the agent.

finish - The time when this job step completed.

When a command finishes executing, the ElectricFlow agent issues a <finishcommand> request containing the exit code and other metadata back to the ElectricFlow server. The "finish" time is when the <finishcommand> is received by the server.

elapsedTime - The number of milliseconds between the start and end times for the job.

This is exactly "finish" - "start".

• totalWaitTime - The sum of resource, workspace, and license wait times for this job step.

The sum of resourceWaitTime, licenseWaitTime, and workspaceWaitTime.

• waitTime - The number of milliseconds the step spent between runnable and running (for example, waiting for a resource).

Covers known reasons why a step couldn't run as soon as it became runnable (same as totalWaitTime), plus any other possible reason (slow transaction commits, system paging, etc).

resourceWaitTime - The length of time this job step stalled because it could not get a resource. This
could indicate that eligible resources for the step have reached their step limit, are in-use but the step
requires a resource exclusively, or resources are down.

For subprocedure steps, this is the sum of all child steps.

• licenseWaitTime - The length of time this job step had to wait to process because the license limit was reached or exceeded.

For subprocedure steps, this is the sum of all child steps.

• workspaceWaitTime - The time this job step had to wait because no workspace was found or available.

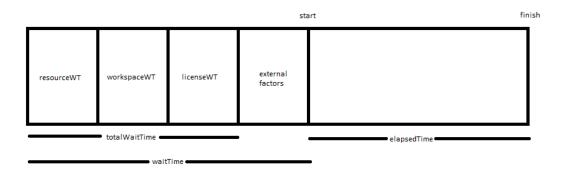
For subprocedure steps, this is the sum of all child steps.

When it is time to schedule the step, ElectricFlow checks the license and starts counting time against licenseWaitTime until the license is satisfied (e.g., running this step would not exceed the maximum number of licensed hosts). After the step is scheduled, ElectricFlow sends a runCommand message to the agent. If it fails because ElectricFlow cannot talk to the agent, ElectricFlow accrues time against resourceWaitTime. If ElectricFlow can talk to the agent, but it has trouble manipulating the workspace (for example, it cannot create the workspace dir, map the drive [on Windows], or create files), ElectricFlow accrues time against workspaceWaitTime.

It is possible for the following to occur:

- 1. A job step becomes stalls due to a license issue.
- 2. The license limit is no longer exceeded and ElectricFlow tries to run the step but hits a bad resource.
- 3. When the resource is good, the step stall because of a license again (adding more time to the licenseWaitTime),
- 4. When the license is satisfied, ElectricFlow reaches the resource, and everything runs.

This is why ElectricFlow continues cycling through these states until everything is satisfied.



The first 4 blocks (resourceWaitTime, workspaceWaitTime, licenseWaitTime, and external factors) can occur in any order.

External factors could include things such as system paging, slow transaction commits, etc.

Job Details

This page displays detailed information about a job. If the job is currently running, the page updates automatically as the job progresses. After the job is complete, the page will not continue to refresh. You can disable the automatic page refresh by clicking the **Stop Refresh** link at the top of the page.

Links and actions at the top of the table

- Abort Aborts a running job.
- Run Again Use this link to run the job again.
 - When you select the **Run Again** button, you are requesting to re-run the job with the same parameter values as the original invocation.
 - If you change the parameter values, you can select "Run..." from the dropdown menu (small downarrow) next to the Run Again button.
- Delete Deletes this job.
- Save Configuration Saves this configuration to the Quick View section on the Home page.
- Access Control Use this link to set privileges for this job. For more information, see the Access Control Help topic.
- **Stop Refresh** This link is available only if a job is currently running. Click this link if you do not want the page to refresh automatically while a job is running.
- View Log This link takes you to the Job Log page. If your job "completed with errors", this log contains
 error messages.
- The "star" icon allows you to save this job information to your Home page.
- The "curved arrow" icon returns you to the Jobs page.
- Pagination Use the "previous" and "next" arrow icons to view the previous or next job. The numbers
 between the arrow icons display the number of jobs you can view and the first number indicates which
 job [in the list] you are viewing.

Summary section (at the top of the page)

Along with the job's start and elapsed times, the summary section provides links to get you to the job's Project or Procedure Details pages or to the Edit Schedule page.

This section also displays:

- A job's priority and a "Calling State" category If the job was called by a workflow, you will see two links, one to the Workflow Details page and one to the State Details page.
- "Wait times" You will see *wait times* specified if your running job was restricted by resource, workspace, a precondition, or license availability.
 - **Note:** Total job wait time is more accurately thought of as *accumulated* wait time because it is the sum of all step wait times, including wait times for steps running in parallel.
- Two other useful sections: Reports and Links. If the job created any custom reports, they appear in the summary area to the right of **General Information**.

Reports

To create a special-purpose HTML report for a job: In the top-level of any of the job's workspaces, add a file that contains the word "report" or "Report" and ends in the .html file extension. The link is automatically created when the Job Details page is displayed. The job's workspace must be accessible to the web server. The link will be in the Reports section. The screen capture below displays an example of the HTML Reports section.



Click on a report file name to display the report.

Links

Notice the additional Links section in the screen example above. To add a link to this section:

· Create a link to any file:

To add a link, run a command in any job's job step using this format (this works for both local [disconnected] and non-local workspaces):

```
ectool setProperty "/myJobStep/report-urls/<myReportName>"
"/commander/jobSteps/<jobStepId>/<artifactDirectoryName>/<file-path>"
```

Example command:

ectool setProperty "/myJobStep/report-urls/Test Report" "/commander/jobSteps/
\$[/myJobStep/jobStepId]/testreport/index.html"

Note: If you are using this format (from ElectricCommander 4.2 and earlier):

```
ectool setProperty "/myJob/report-urls/<myReportName>" "/commander/jobs/<jobId>/
<workspaceName>/<artifactDirectoryName>]/<file-path>"
```

Example command:

```
ectool setProperty "/myJob/report-urls/Test Report" "/commander/jobs/$[/myJob/jobId]/
$[/myJobStep/workspaceName]/testreport/index.html"
```

It will continue to work onlyfor non-local workspaces. We recommend, however, that you change the command to use the newer format.

Create a link to any directory:

To add a link, run the following style command in any job's job step:

```
ectool setProperty "/myJobStep/report-urls/Main Job Workspace"
"file:///WinStor2/scratch/chron55build/$[/myJob/jobName]"
```

Note: Links to a directory automatically work in Internet Explorer, but if using Firefox, local links are disabled unless the default security policy is modified or an extension is used. Refer to http://kb.mozillazine.org/Links_to_local_pages_don%27t_work for more details.

The "tabbed" sections

The next section of tabs allows you to select the type of information you want to see.

Steps table

The Job Steps table displays all steps for this job.



- Mouse-over the View: All link [below the tabs] to see a filter list. You can choose a filter to apply to the
 Job Step table.
- The **Expand All | Collapse All** links are provided to see all subprocedures or to minimize the information presented.
- Click on a Step Name to go to the Job Step Details page for that step.
- Click the Log icon for any step to display the step's log file. This icon is not displayed if the logFileName property was removed from the step.
- Status column: This column provides the job's step status.
 - o Aborted The step or job was aborted.
 - o Canceled The step did not run because the job was aborted or a previous step was aborted.
 - o Error The command failed or postp detected an error in the output.
 - Running The step is currently running or a step inside the subprocedure call is running.
 - Skipped The step did not run because the run condition evaluated to "false".
 - Success The step completed with no errors or warnings.
 - Waiting for Resource or Workspace The step is waiting to run as soon as a resource or workspace is available.
 - Warning postp detected a warning message.

Note: The Status field can be over-ridden by postp to include more specific information. postp information may be linked directly to diagnostic information—also available by selecting the Diagnostics tab.

Action column:

- Click the Edit link to edit the procedure step definition.
- Aborting a step requires the execute privilege on the job—not just the modify privilege.

When you click **Abort**, the dialog box similar to the following appears:



Click **OK** to abort the step. If the step does not complete in approximately two minutes, you can "forcibly" abort the step to force the step to a completed status, enabling the rest of the job to continue.

Click Cancel if you change your mind.

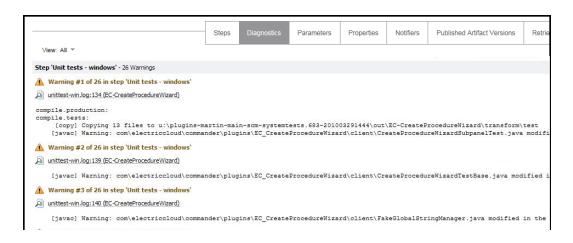
Diagnostics table

If the job generated warning or error messages, click this tab to see one or more diagnostics, each of which is a portion of a step's log file identified as "interesting" by the step's postprocessor. Typically, each diagnostic relates to an error or warning detected by the postprocessor. All diagnostics for a particular step appear together.

Each diagnostic contains a header line such as "Error #3 of 5" with an icon, followed by a line with additional identifying information and one or more lines extracted from the log file. The identifying line has a format similar to:

"systemTest-windows.log:16 (systemtest)"

- The first part of the line provides the name of the log file (systemTest-Windows.log) and the line number of the first line in the file that is part of this diagnostic (16).
- If you click on the log file name, ElectricFlow displays the full log file for you, with the portion in this diagnostic highlighted. If the postprocessor was able to provide additional information about the issue, it is displayed in parentheses after the log filename. The first value (systemtest-windows) gives the name of a particular module or package in which the problem occurred.
- The second value (systemtest) identifies a specific place where the problem occurred, such as the file name that did not compile or the name of a test that failed.



Parameters table

This section displays parameter values supplied to the top-level procedure when the job was invoked. This table does not allow you to edit the parameter's value.

These parameters are copied automatically to the top-level property sheet and you will see them in the properties table also.



Properties table

This table contains information computed and/or used by the job while it is running.

- All defined parameters are automatically copied to the top-level property sheet. Any custom properties defined for the state will be included in this table also.
- After the job completes, these properties are normally read-only.
- Click on a Property Name to edit that property.
- If a "folder icon" precedes a property name, it denotes a Nested Property Sheet.
- If you use square brackets in a property name, the property must be encapsulated in square brackets to avoid property sheet expansion. Note that this works for specific functions, such as creating a link on a job, but the drawback is that the property cannot be edited or deleted from the web interface. Example:

ectool setProperty "/myJob/report-urls[This property name has [some square brack
et] text in it]" "http://website.com"

• This section also provides **Create Property**, **Create Nested Sheet**, and **Access Control** links to update or enter additional information for the next time this job is run.



Notifiers table

If you designated certain persons or groups to be notified of details for specific jobs, this section displays who receives an email notification, the type of notification, any condition, and a description [if supplied]. If a notifier is available, you can click on its name to view details or make modifications.



Published Artifact Versions table

This table displays artifact versions published by steps in this job.



Column descriptions

- Name Select this name to go to the Artifact Version Details page for this artifact version.
- State Possible values for the state of the artifact version are: available | publishing | unavailable
- Description A previously supplied text or HTML description for this artifact version.
- · Actions column:
 - o Edit Select this link to go to the Edit Artifact Versions page.
 - **Delete** Select this link to delete the artifact version. You may be prohibited from deleting this artifact version if you do not have sufficient access control permissions.

Retrieved Artifact Versions table

This table displays artifact versions retrieved by steps in this job.



Column descriptions

- Name Select this name to go to the Artifact Version Details page for this artifact version.
- State Possible values for the state of the artifact version are: available | publishing | unavailable
- Description A previously supplied text or HTML description for this artifact version.

- · Actions column:
 - o Edit Select this link to go to the Edit Artifact Versions page.
 - Delete Select this link to delete the artifact version. You may be prohibited from deleting this artifact version if you do not have sufficient access control permissions.

Licenses

This page displays all license information known to the ElectricFlow server. Typically, a single license is displayed, which describes the usage to which you are entitled. The license on the server may be based on concurrent resources, concurrent users, concurrent steps, registered hosts, or any combination of these licenses.

All Licenses

Click **Import License** to import a new license file.

Column descriptions

Column Name	Description / Actions
Company	The name of the company or organization to which the license was granted. If this is not your company, please contact Electric Cloud Customer Support.
Feature	The type of license.
Expiration	The last day when this license is no longer valid. After your license expires (and after any grace period), running jobs will complete, but no new jobs can be launched.
Grace Period (days)	If this field is nonzero, you may continue using ElectricFlow for this many days after the expiration date. However, you will be warned frequently during the grace period that your license has expired.
Actions	Click the View License button to see additional details about this license.
	Click the Delete button to delete this license.

Current Usage

You see some or all of the following fields depending on the type of license or licenses you use for ElectricFlow.

Field Name	Description
Maximum Concurrent Steps	The maximum number of concurrent steps allowed by this license.
Concurrent Steps	The number of concurrent steps running now.
Maximum Concurrent Hosts In Use	The maximum number of unique host names defined in concurrent resources (non-proxy resources), where steps can run at any time.
Concurrent Hosts in Use	The number of unique host names defined in concurrent resources where steps are currently running.
Maximum Proxied Hosts In Use	The maximum number of unique host names defined in <i>proxied</i> resources, where steps can run at any time.
Maximum Concurrent Hosts	The maximum number of unique host names allowed to use ElectricFlow concurrently at any given time.
Concurrent Hosts in Use	The number of unique host names defined in concurrent resources where steps are currently running.
Maximum Proxied Host	The maximum number of unique host names define in proxied resources.
Proxied Hosts in Use	The number of unique host names defined in <i>proxied</i> resources where steps are currently running.
The following fields are sp	ecific to a concurrent user-based license.
Maximum Concurrent Users	The number of users allowed to use ElectricFlow concurrently at any given time.
Active Users	The number of users with an active license.
User IP	The "user@IP" combination that holds an active license.
Expiration	The time when the license will be released.
Last Use	The time when the license was last used.

Field Name	Description	
The following fields are specific to a registered host-based license.		
Maximum Registered Hosts	The maximum number of registered hosts allowed by this license.	
Registered Hosts in Use	The number of registered hosts where steps are currently running.	
Maximum Applications	The maximum number of applications that ElectricFlow can deploy at any given time.	
Applications in Use	The number of applications that are currently running.	

License Types

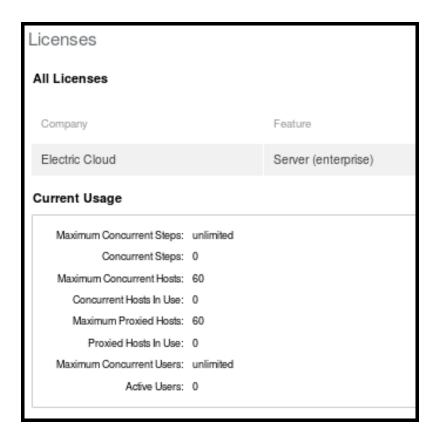
Concurrent Resource License

This license controls the number of unique host names on which steps can be scheduled at any time. Two properties, <code>maxHosts</code> and <code>maxProxiedHosts</code>, are defined on a license to control the maximum number of concurrent and proxied hosts that can be in use at the same time. The host named "localhost" and "127.0.0.1" refer to the local agent installed on the ElectricFlow server machine, and these host names are not part of the total count for hosts in use.

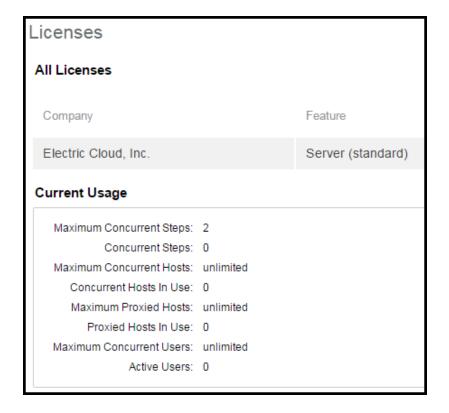
If two or more steps are running concurrently on the same host, the host name still counts once only towards the license limit. There is no limit to the number of resources pointing to a single host; from a licensing perspective, they all count as a single host. When the host limit is met, any new job step that attempts to run on a host [not already in use] will stay in the "runnable" state until its host is available. Current license usage statistics are displayed at the top of the Resources page in the web interface. If no license is available, jobs cannot be launched. When a license expires, running jobs are allowed to complete, but no new jobs can be launched.

A concurrent resource license can be a standard or enterprise license.

This is an example of an enterprise license:



This is an example of a standard license:



Concurrent User License

If the maxConcurrentUsers property is set on the license, its value defines the maximum number of users who can actively use the ElectricFlow server. A unique license is defined by the combination of the user's name and the IP address from which the user is making requests. A user can have multiple active sessions (through the web UI and ectool) from the same machine while only using a single license.

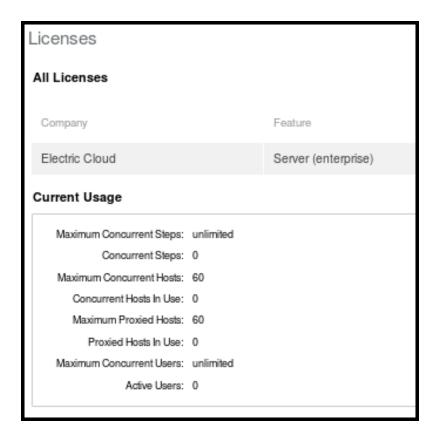
If a user makes a licensed API call without an active license, the server automatically attempts to assign a license. If a license is available (because the maximum has not been met), the user is assigned a license. If no licenses are available, the user will see a NoAvailableLicenses error and be redirected to the licenses page. Currently held licenses and their expiration dates can be viewed by calling the getLicenseUsage API or by viewing the Active Users table on the Licenses web page. When a license becomes available, the first user to make a licensed API call will be assigned that license. The only API calls not licensed are login, getLicenseUsage, getAdminLicense, getServerStatus, and getServerInfo.

When a license is first acquired, its expiration is set to 1 hour in the future. After the initial hour has passed, each time the user makes a licensed API call, the license is extended another 10 minutes (a configurable server setting). When a user's license passes the expiration date, it is released. When the user makes another licensed API call, a license is assigned automatically if one is available, otherwise the <code>NoAvailableLicenses</code> error is displayed.

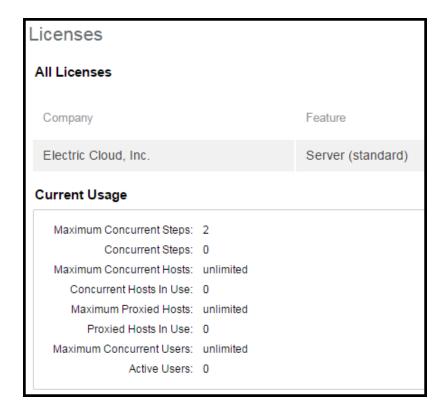
A single administrator license can be used in an emergency. All users who have *execut*e permission on the system-level licensing access control list can use the administrator license by calling the <code>getAdminLicense</code> API. This license has no expiration as long as all regular licenses are in use, but this license abides by the regular license expiration policy if regular licenses are available.

A concurrent user license can be a standard or enterprise license.

This is an example of an enterprise license:



This is an example of a standard license:



Concurrent Step License

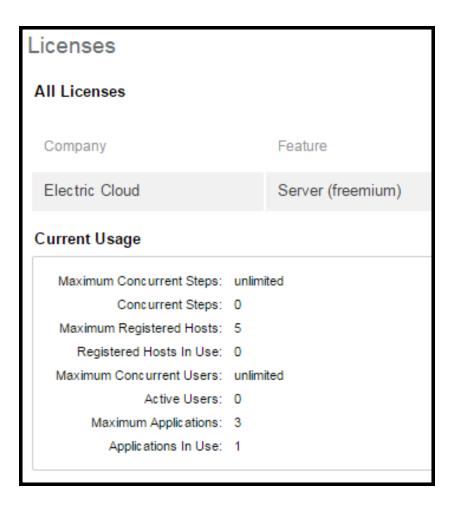
This ElectricFlow license controls the number of unique steps you are allowed to run in the system at the same time (in parallel). The maxConcurrentSteps property, defined on a license, controls the maximum number of steps that can be running at the same time.

Registered Host License

Registered hosts are dedicated hosts on which a procedure or step can always run. This license controls the number of unique registered host names. Two properties, <code>maxHosts</code> and <code>maxProxiedHosts</code>, are defined on a license to control the maximum number of registered and proxied hosts that can be in use at the same time. The host named "localhost" and "127.0.0.1" refer to the local agent installed on the ElectricFlow server machine, and these host names are not part of the total count for hosts in use. Concurrent hosts cannot be used with registered hosts.

If two or more steps are running concurrently on the same host, the host name still counts once only towards the license limit. There is no limit to the number of resources pointing to a single host; from a licensing perspective, they all count as a single host. When the host limit is met, any new job step that attempts to run on a host (not already in use) stays in the "runnable" state until its host is available. Current license usage statistics are displayed at the top of the Resources page in the web interface. If no license is available, jobs cannot be launched. When a license expires, running jobs are allowed to complete, but no new jobs can be launched.

This an example of a registered host license:



View License

This page displays all available license details for a single license.

Use the **Delete** link to remove this license.

Information

Name	Description
Company Name	The name of the company or organization to which the license was granted. If this is not your company, contact Electric Cloud Customer Support.
Product Name	Product to which this license applies (if it is not "ElectricFlow," it will not apply here).

Name	Description
Feature Name	Product feature to which this license applies: must be "Server" for ElectricFlow licenses.
Maximum Concurrent Steps	The maximum number of concurrent steps allowed by the license.
Maximum Concurrent Hosts In Use	The maximum number of unique host names defined in concurrent resources (non-proxy resources), where steps can run at any time.
Maximum Proxied Hosts In Use	This is the maximum number of unique host names defined in proxied resources, where steps can run at any time.
Maximum Registered Hosts	The number of registered hosts on which a procedure or step can always run. This appears when the license on the ElectricFlow server is based on a combination of concurrent and registered hosts or only registered hosts.
Maximum Concurrent Users	The number of users allowed to use ElectricFlow concurrently at any given time.
Expiration Date	The last day when this license is no longer valid. After your license expires (and after any grace period), running jobs will complete, but no new jobs can be launched.
Grace Period	If this field is non-zero, you may continue using ElectricFlow for this many days after the expiration date. However, you will be warned frequently during the grace period that your license has expired.

Import License

Use this page to provide a new license for ElectricFlow.

You should have received a license (as a block of text) from Electric Cloud already. If you have not received a license, contact Electric Cloud Customer Support or your sales representative.

- Use a text editor to open the license file.
- Copy the text into the Contents field.
- Click OK.

You will see your license information on the Licenses page.

Plugin Manager

A plugin is a collection of one or more features that can be added to ElectricFlow.

- Plugins are delivered as a JAR file containing the features implementation.
- When a plugin is installed, the ElectricFlow server extracts the JAR contents to disk into a configurable plugins directory.
- A plugin has an associated project that can contain procedures and properties required by the implementation.
- A plugin can provide one or more new pages for the web interface and may also provide a configuration page so you can provide any additional information that may be necessary to implement the plugin.

ElectricFlow provides default/bundled plugins, installed during the ElectricFlow installation. These plugins are "authored" and generally supported by Electric Cloud.

The information Plugin Manager displays across the top of the table:

Tab and field descriptions above the table

Currently Installed - The Plugin Manager default page view displays currently installed plugins. If you install additional ElectricFlow plugins from other sources, including those plugins you may create, they will also appear on this page.

Install from File/URL - Select this tab to see the following screen for plugin installation options:

File Install - Use this method to upload and install a local "jar" file from the machine running your web browser. Click **Upload** after specifying the path you need.

URL Install - install a plugin from a location specified by an URL. This location can either be an external web server (using http://), or a file on the ElectricFlow server host (using file://)

Install a plugin from the command-line:

ectool - the command-line tool, contains a full set of commands to perform plugin tasks. For more information on available ectool plugin commands, see the "API Commands - Plugin Management" topic in the ElectricFlow API Guide.

Note: The server has a configurable maximum file upload size that defaults to 50MB. Refer to the server "Maximum upload size" setting to adjust the default.

View Catalog - This page displays all plugins available for installation. If you do not see an Install option in the Actions column, you may not have the necessary access control permissions to install a plugin.

Catalog Install - click on the **Install** action (next to a plugin listed in the catalog) to download and install the plugin.

Category: a drop-down menu to filter the plugin list by a specific category.

Column descriptions

Column Name	Description / Actions
Plugin Label	The name of the plugin. Click on the label of an installed plugin to go to the Project Details page for that plugin.
Author	Identifies the person or organization that created the plugin and optionally provides a link to contact the author.
Version	The plugin version. Promoted plugin versions are also shown in this column. A promoted version number is labeled "(promoted)" immediately after the version number.
Category	The plugin type. Plugins that provide similar types of functionality are grouped into categories.
Description	A short text string that describes the plugin purpose or function.

Column Name	Description / Actions
	This column contains one or more of the following links to perform plugin actions or tasks. Not all plugins have the same task options. For example, some plugins may perform their own "setup" during installation and not require additional configuration information.
	Note: If you do not see the option you need in the Action column, it may be hidden from your view if you do not have sufficient access control permissions.
	Possible Actions: Configure - use this link to go to the configuration page for the plugin.
	Demote - use this link to make this plugin inactive and remove any tabs associated with the plugin.
Actions	Note: If you re-install a previously demoted plugin, previous values from the demoted version are not copied to the new version you install.
	Help - if provided, use this link to go to the plugin documentation.
	Install - Download and install a plugin from the plugins catalog site.
	Promote - use this link to upgrade the plugin to a new version. Old plugin values are copied to the promoted version.
	Re-install - use this link to download and reinstall a previously installed plugin.
	Uninstall - use this link to remove this plugin from your system.
	Note: the Demote, Promote, Install, Re-install, and Uninstall actions appear only if the user has modify permission to the "plugins" ACL.

Using a plugin

Multiple methods for adding plugin functionality to ElectricFlow are available.

- To configure ElectricFlow to recognize your SCM plugin, see the Configuring ElectricFlow Help topic and select the Setting Up a Source Control Configuration section.
 - o For additional plugin Help, see the Customizing the ElectricFlow UI Help topic.
 - For additional help using your SCM plugin with ElectricFlow Preflight, Default Tracking, and ElectricSentry, see those Help topics: Preflight Builds, Defect Tracking, ElectricSentry
- To add a subprocedure step using the Edit Step or New Step page in the ElectricFlow UI:
 - 1. Select the Projects tab and choose an existing project or click the **Create Project** link if you need to create a new project.

- 2. On the Project Details page, select an existing procedure or click the **Create Procedure** link if you need to create additional procedures.
- 3. On the Procedure Details page, to create a New Step using a plugin, select the **Plugins** link to go to the Choose Step panel.
 - In the left pane, select a plugin category and then choose a plugin.
 - In the right pane, select the step type you want, which will take you to the New Step page to add a new subprocedure step (using your selected plugin).
- 4. On the New Step page, notice that the plugin you selected is already displayed in the Subprocedure section.
- 5. Enter information in the Parameters section fields. These fields are specific to the plugin you chose. (These fields are frequently different for each plugin.)

Configuring the plugins directory

For instructions on configuring the plugins directory for your ElectricFlow server, remote agents, and/or remote web servers, see the *ElectricFlow Installation Guide*.

Procedure Details

This page displays procedure component tables, including a table for Procedure Steps, Parameters, Email Notifications, and Custom Procedure Properties. A procedure describes a process to execute that consists of one or more step.

Use this page to create steps.

Procedure Steps

This table displays all steps in the procedure [named at the top of the page]. The steps execute in order from top to bottom. Normally only one step executes at a time, but if a group of adjacent steps is marked as "parallel," all of those steps will execute simultaneously.

Important: If you want to change the order of the steps, click on the icon to the left of the step name and drag it up or down to the new step position.

Note: This table displays only a few of the fields for each step. To view (and modify) the complete details for a step, click on the step's name to go to the Edit Step page.

Links at the top of the page

- Use the **Run** link [hovering your mouse over the small down-arrow on the right-side of the link] to choose Run Immediately or Run...
 - Run Immediately Selecting this option runs the procedure immediately as set.
 - Run... Selecting this option displays the Run Procedure web page where you may alter any existing parameters for this procedure, or set an existing credential.
- Click the Access Control link to go to the Access Control page to add privileges for this procedure.
- Click the Edit link to go to the Edit Procedure page so you can make modifications if necessary.

- Use the Copy link in the Actions column to make an exact copy of this step.
- Choose one of the New Step links to create the type of step you need.

Creating a new step

The following list describes some of the more common steps you can create.

- **Command** this link takes you directly to the New Step page to write your own step. A Command(s) text box is available to add a script. This step invokes a bat, cmd, shell, perl script, or similar.
- **Subprocedure** the Choose Subprocedure Step dialog is displayed, allowing you to select a project or plugin from which you want to call an existing procedure to create a step. This step invokes another ElectricFlow procedure and will not complete until all subprocedure steps have finished.
- Plugin this link opens the Choose Step panel. The left pane displays plugin categories and the right pane displays available pre-configured steps for your left-pane selection. To quickly find your application/plugin, type your selection name in the Search box. Select the step you want to create and ElectricFlow takes you to the New Steps page.
 - On the New Step page, notice that the plugin step you selected is already displayed in the Subprocedure section.
 - Enter information in the Parameters section fields. These fields are specific to the plugin you chose. (These fields are different for each plugin.)

Enter information in any remaining fields to complete your step. See the New Step Help topic for more information.

A note about plugins

Plugins are the vehicle ElectricFlow uses to integrate source control systems (such as Subversion or Perforce), defect tracking applications (such as JIRA or MKS), reporting functionality, code analysis applications, build applications, and much more. To see the complete list of plugins bundled and installed with ElectricFlow, see the Plugin Manager page (go to Administration > Plugins).

Note: If you do not see the SCM system or defect tracking system you prefer, a Plugin Catalog is available from the Plugin Manager also. All plugins in the catalog are available for installation into ElectricFlow.

Parameters

Each procedure can define parameters, which are values provided to the procedure when it is invoked. Each parameter value is placed in a property associated with the job and the procedure can use the property to control its execution. For example, a parameter might specify a particular branch of a product to build; another parameter might indicate whether unit tests should be run during this build.

When you define a parameter, indicate whether the parameter is required (meaning the procedure will not run unless a value is provided for the parameter). Also, you can provide a default parameter value and a description to help callers understand how to use the parameter.

- Click the Create Parameter link to go to the New Parameter page.
- Click on the Parameter Name to edit its contents.

Email Notifiers

Use the **Create On Start Email Notifier** or **Create On Completion Email Notifier** links to go to one of the respective pages to create an email notifier for this procedure.

After creating an email notifier, click on the "bread crumb" link at the top of the page to return to the Procedure Details page.

Custom Procedure Properties

- To add new properties to this procedure, click the Create Property, Created Nested Sheet, or Access Control links.
- Click on the Property Name to edit its contents.
- Nested Property Sheets appear in this table preceded by a folder icon.

Tip

Want to rename this procedure or change its description? Click the Edit link at the top of the page.

Procedure - create new or edit existing procedure

To create a new procedure

Enter information in the fields as follows:

Field Name	Description
Name	Enter a name for the procedure. This name must be unique within the project.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <di> <i> <o> <pre> <pre> <style> <tc> <u>><u>></td></tr></tbody></table></style></pre></pre></o></i></di></div>

Field Name	Description
	This is the template used to determine the default name of jobs launched from the procedure.
	For example:
	<pre>\$[projectName]_\$[/increment /myproject/jobCounter]_ \$[/timestamp]</pre>
Job Name Template	Produces a name like:
l l l l l l l l l l l l l l l l l l l	projectFoo_1234_20140102130321
	Enter any combination of elements to create procedure names more meaningful to you. For example, you could choose to include the build number.
	Note: Electric Cloud does not recommend using the ElectricFlow-generated <code>jobId</code> because it is no longer a human-readable integer so it does not provide any indentifiable information and cannot be used as a counter.
Default Resource	Enter a resource name. This name must be unique among other resource names or click Browse to select a resource.
Default Workspace	Enter the name of the workspace or click Browse to select a workspace name.
Time Limit	If no time limit was specified on the calling step, time limits are copied to the calling step from the procedure. If the procedure is called from runProcedure (or a schedule), the time limit acts as a global job timeout. The "timer" for the procedure starts as soon as the calling step/job becomes runnable (all preconditions are satisfied). Specify a number and then use the pull-down menu to select the "time units" of
	the number you specified.
Impersonation Credential	
Credential Project	Default Credential Project is the current project, or select the other button and Browse for the project name you need.
Credential Name	The name of the credential for the project - you may need to Browse to find this name.

Click \mathbf{OK} to save the information to create a new procedure.

To edit an existing procedure

Use this page to modify certain overall information about a procedure. Additional information, such as procedure steps and properties, can be viewed and modified on the Procedure Details page. To access this page, click the **Project Name** link on the location path (breadcrumb), then select a Procedure name.

To learn more about a particular field on the form, mouse-over the field to pop-up descriptive text under the mouse or click the **Help** link at the top of the page.

Tip

Want to rename the procedure? Enter a new procedure name in the Name field, click **OK**.

Step - create new or edit existing step

Use this page to define a new step or to modify an existing step. The following information describes the common fields you need to enter to create any type of step.

Note: If you chose to create a step using the Plugin link, additional help is available in the Parameters section by clicking the "?" icon. The Parameter section is populated according to type you chose to create. Different step types require different information.

To create a new command or subprocedure step

Enter information in the fields as follows:

Field Name	Description	
General section	General section (for all step types)	
Name	Enter a name for the step that must be a unique name within the procedure.	
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <dl> <i> <pr> <style> <tc> </td></tr><tr><td colspan=2>Command section (for command steps only)</td></tr><tr><td>Command(s)</td><td>A script to execute the step functions that is passed to the step's shell for execution.</td></tr></tbody></table></style></pr></i></dl></div>	

Field Name	Description
Resource	Enter a unique resource or resource pool name or click Browse to select a resource. If you leave this field blank, the procedure or project provides a default resource.
	 The job step must have the execute privilege on the resource or resource pool. If it does not, the job step will be canceled with an access denied error.
	 The resource or resource pool must be enabled. If the resource or pool is not enabled, no error is generated, and the step remains in the runnable state until the pool is enabled.
Postprocessor	If this field remains blank, no postprocessor runs for the step. If this field is not blank, it specifies a command (passed to the step's shell for execution) that analyzes the log file for the step and collects diagnostic information for reporting. For more information, see the Postprocessors Help topic.
Subprocedure section (for subprocedure steps only)	
Procedure	Displays the current project and procedure. Click Change to use the pop-up menu to select a new project or procedure.
Resource	Enter a unique resource name or click Browse to select a resource. If you leave this field blank, the procedure or project provides a default resource.

Parameters section (for subprocedure steps only)

This section expands automatically, entering the fields you need specifically for the Subprocedure you chose in the previous section. For additional help with field descriptions, click the "?" icon if available.



To toggle between the standard view and property view for this section, click the icon. This icon appears in the standard form when one or more of the parameter types are not **Text entry** or **Text area**. It appears in the custom form for all the supported parameter types when the form has at least one parameter.

Advanced section (for both command and subprocedure steps)

Field Name	Description
	By default, if this field is blank, the step has no precondition and will run when scheduled. Set this property to make a step wait until one or more dependent conditions are met. When a job step is eligible to transition from pending to runnable, a precondition is evaluated. A precondition is a fixed text or text embedding property reference that is evaluated to TRUE or FALSE. An empty string, a \"0\" or \"false\" is interpreted as FALSE. Any other result string is interpreted as TRUE. The step will block until the precondition is TRUE.
	Precondition example: Assume we defined these 4 steps:
	Build object files and executables
	2. Build installer
Precondition	3. Run unit tests
	4. Install bits on test system
	Step 1 is an ordinary serial step. Steps 2 and 3 can run in parallel because they depend only on step 1's completion. Step 4 depends on step 2, but not step 3.
	You can achieve optimal step execution order with preconditions:
	Make steps 2-4 run in parallel.
	 Step 2 needs a job property set at the end of its step to indicate step 2 is completing (/myJob/buildInstallerCompleted=1). Set a precondition in step 4: \$[/myJob/buildInstallerCompleted]
Run Condition	By default, if this field is blank the step will always run. This field accepts "fixed text" or text embedding property references that are evaluated into a logical TRUE or FALSE. A "0" or "false" is interpreted as FALSE. An empty string or any other result is interpreted as TRUE.
	The property reference can be a JavaScript expression, for example, this expression could test whether the name of a step is equal to the value of a property called "restartStep". \$[/javascript (myStep.stepName == myJob.restartStep)]

Field Name	Description
Error handling	This field determines what happens to the procedure if the step fails.
	The first choice, "procedure continues, but overall status will be error" means an error in the step does not abort the procedure; subsequent steps will run as usual. However, an error will be reflected in the overall outcome of the procedure. If this is the top-level procedure in the job, the job will have an error outcome.
	If this is a subprocedure, the step invoking the subprocedure will have an error outcome, which still could cause the job to abort, depending on the error handling for that step.
	The second choice for this field is "abort procedure but allow running steps to continue." In this case, if an error occurs, no new steps are initiated in this procedure except those where "always run" was selected. Any steps currently running are allowed to complete, then the procedure will abort and its outcome will be set to error as described previously. Six error handling options are:
	• failProcedure - The current procedure continues, but the overall status is error (default).
	abortProcedure - Aborts the current procedure, but allows already- running steps in the current procedure to complete.
	abortProcedureNow - Aborts the current procedure and terminates running steps in the current procedure.
	abortJob - Aborts the entire job, terminates running steps, but allows alwaysRun steps to run.
	 abortJobNow - Aborts the entire job and terminates all running steps, including alwaysRun steps.
	ignore - Continues as if the step succeeded.
Time Limit	The maximum amount of time the step can execute. If the step exceeds this time, it will be aborted. Specify a number and then use the pull-down menu to select the "time units" of the number you specified.
Run in parallel	If this box is "checked," this step will run concurrently with any adjacent steps marked as parallel. If this box is not checked, the step will NOT run until all previous steps in the procedure have completed, and it will complete before any following steps execute.

Field Name	Description
Always run step	If a procedure is aborted (either because of an error in a step or because the job was aborted), under normal conditions no new steps will start in the procedure. However, if this box is "checked," the step will run even if the procedure is aborted. This facility is most commonly used for steps at the end of a job that generate reports or perform cleanup operations. Note: Remember that the Run Condition field is still evaluated and may cause the step to be skipped even if this box is checked.
	Select one of the following options:
	None - the "default", which does not retain a resource.
Retain Exclusive	 Job - keeps the resource for the duration of the job. No other job can use this resource, regardless of its step limit, until this job completes or "Release Exclusive" is used in a step. Future steps for this job will use this resource in preference to other resources—if this resource meets the needs of the steps and its step limit is not exceeded. Step - keeps the resource for the duration of the step.
	Call - keeps the resource for the duration of the procedure that called this step, which is equivalent to 'job' for top level steps.
Workspace	Use this field to define a workspace for the step to use. You can either type-in a workspace name or click Browse to select a workspace. If you leave this field blank, the workspace is determined by the procedure or project.
These "Advance	d" fields are used for command steps only
Broadcast	"Checking" this box means replicate the step to execute (in parallel) on each of the specified resources (for example, if a resource pool is specified, run the step on each resource in the pool).
	Select one of the following options:
Release Exclusive	 none - the "default" - no action if the resource was not previously marked as "retain". release - releases the resource at the end of this step. If the resource for the step was previously acquired with "Retain exclusive" (either by this step or some preceding step), the resource exclusivity is canceled at the end of this step. The resource is released in the normal way so it may be acquired by other jobs.
	releasetojob - allows a step to promote a Step exclusive resource to a Job exclusive resource.

Field Name	Description
	This shell will be used to execute the step's commands on a resource. For example, using sh or cmd /c, the agent saves the command block to a temporary file with a .cmd extension and runs it: sh foo.cmd or cmd /c foo.cmd
	If you do not specify a shell on a step, at step run-time the server looks at the resource shell. If a resource shell is not set, the shell line used by the agent is platform dependent: Windows: cmd /q /c "{0}.cmd" UNIX: sh -e "{0}.cmd"
	When you specify a shell (in the step or resource), and omit the cmd-file marker, the agent notices the omission and takes the correct action. For example: a user specifies sh -x. The agent converts this to sh -x "{0}.cmd"
	Two alternate forms of shell syntax where ElectricFlow uses a "marker," {0}, as a placeholder for the command file argument:
Shell	 <myshell> {0} <potential args="" extra="" shell=""></potential></myshell>
	 <myshell> {0}<.file extension> <potential args="" extra="" shell=""> In this example, the shell requires the command file to end in an extension other than .cmd. For example, powershell "& '{0}.psl'" This shell example runs Microsoft PowerShell against this step's command containing PowerShell commands.</potential></myshell>
	Notes:
	When the agent parses the shell, it will parse the extension as everything after {0}. until it sees a space or non-alphanumeric character.
	If your script uses International characters (non-ascii), add the following block to the top of your ec-perl command block: use utf8; ElectricCommander::initEncodings

Field Name	Description	
Working Directory	The directory where the commands for this step execute. A relative name is interpreted relative to the root directory for the job workspace. If you leave this field blank, the step's working directory will be the top-level directory in the job workspace. Note: If this step will run on a proxy resource, this directory must exist on the proxy target. The step will run in this directory on the proxy target.	
Log File	The step's log file name, specified relative to the root directory in the job workspace. If you leave this field blank, ElectricFlow picks a unique name for the log file based on the step name.	
Impersonation section (for both command and subprocedure steps)		
Credential Project	By default, the Current project is used, or click Browse and select a different project.	
Credential Name	If you select a credential in this field, the step runs under the account from that credential. This field is usually blank, which means a default credential is used. For more detailed information about credentials, see the Credentials and User Impersonation Help topic.	

Click **OK** to create the step.

To edit an existing command or subprocedure step

You can modify any of your previously entered step information and add information to one or all of the four additional sections.

Attached Credentials

Click the **Attach Credential** link to retrieve an existing credential for this step. If you need to create a new attached credential, return to the Project Details page for this step.

Attached Parameter Credentials

Click the Attach Parameter Credential link to retrieve an existing parameter credential for this step.

Email Notifiers

Choose **Create On Start Email Notifier** or **Create On Completion Email Notifier** to go the respective page to create the email notifier you need.

Custom Step Properties

Choose from the **Create Property**, **Create Nested Property**, or **Access Control** links to add a custom property to this step. Each Property pop-up box has its own Help topic or go to the main Properties Help

topic for more information. For more information about Access Control, go to the main Access Control Help topic.

Click \mathbf{OK} to save your modified step information.

Publish Artifact Version Step

Use the EC-Artifact plugin to create a step to publish a new artifact version.

The following parameters are available to create this step:

Field Name	Description
Artifact	Artifact name, in the form <pre><pre>cartifactKey</pre>. If the artifact does not exist, it will be created if this procedure's launching user or this project principal has the required permissions.</pre>
Version	A full version string takes the form: <major>.<minor>.<patch>-</patch>----- <qualifier>-<buildnumber>. The version specification must be unique across all of this artifact's versions.</buildnumber></qualifier></minor></major>
Repository	Name of the repository where this new artifact version will be published.
Enable Compression	Check this box to compress the artifact version before it is stored in the repository.
From Directory	Name of the directory in the job's workspace containing files that comprise the artifact version to be published. If not specified, the entire workspace is used.
Include Patterns	List file "include" patterns one pattern per line, to limit which files are published. If no patterns are specified, all files are included.
Exclude Patterns	List file "exclude" patterns one pattern per line, to limit which files are published. If no patterns are specified, no files are excluded.
Dependent Artifact Versions	List dependent artifact versions one per line, each in the form <groupid>:<artifactkey>:<versionrange>. All dependent artifact versions must exist for this artifact version to be retrievable. When this artifact version is successfully retrieved, its dependent artifact versions are retrieved also.</versionrange></artifactkey></groupid>

Retrieve Artifact Version Step

Use the EC-Artifact plugin to create a step to retrieve an artifact version.

The following parameters are available to create this step:

Field Name	Description
Artifact	Artifact name, in the form <pre><pre>cartifactKey>.</pre></pre>
Version	Select the latest version (Latest), the exact version of the artifact (Exact), or a version range (Range). Version is in the form: <major>.<minor>.<patch>-</patch>-</minor></major>
Retrieve to directory	If you want to retrieve this artifact version to a specific directory location, select the check box, then specify the full path to the directory where you want to retrieve this artifact version.
Overwrite	Use the drop-down menu to select one of the following options: • true - deletes previous content in the directory and replaces the content with your new version. • false - (existing behavior) if the directory does not exist, one will be created and filled with the artifact's content. If the directory exists, a new directory is created with a unique name and the artifact contents is supplied there. • update - this is similar to a merge operation—two artifact versions can be moved into the same directory, but individual files with the same name will be overwritten.

Field Name	Description
	Name or property sheet path used by the step to create a property sheet. This property sheet stores information about the retrieved artifact versions as XML in the artifactVersionXML child property, and the file system location of the retrieved artifact version in the cacheLocation child property.
	The initial value of this field includes \$[assignedResourceName] as one of the levels in the property path because different resources may have different cache directories, and you do not want a retrieval on one resource to clobber the data regarding a retrieval on a different resource.
Retrieved Artifact	Typically, a step that needs to use files from an artifact retrieved by this step will specify it in a Perl script similar to this:
Location Property	<pre>my \$cmdr = new ElectricCommander(); my \$xpath = \$cmdr->getProperty ("/myJob/retrieveArtifactVersions/\$[assignedResourceName]</pre>
	<pre>my \$retrieveDir = \$xpath->findvalue("//value")->value ();</pre>
	<pre>system("java", "-cp", "\$retrieveDir/lib/bar.jar", "Bar");</pre>
	This example shows how to retrieve the location of the retrieved MyGrp:MyKey artifact version on the resource, and add a file (bar.jar) in its lib directory to classpath, so the Bar class can be run.
Filters	Enter search filters, one per line, applicable to querying the ElectricFlow database for the artifact version to retrieve.

Artifact Retrieval Dependency Order Explained

The order of dependencies registered for an artifact version are significant.

Consider this scenario:

- A depends on B (any version) and C [1.0, 2.0)
- B depends on C (any version)
- C versions 1.0, 2.0, and 3.0 exist

When retrieving A, the dependency algorithm evaluates B first. The algorithm finds that the max version of B depends on any version of C, so the algorithm looks for max version C and finds C 3.0. Because this chain is satisfied, the algorithm returns to A and evaluates its next dependency "C [1.0, 2.0)". This results in matching C 1.0.

The returned artifacts are: A, B, C 1.0, and C 3.0.

Consider if the A dependency is changed to:

• A depends on C [1.0, 2.0), B (any version)

The algorithm will choose C 1.0 first. Then the algorithm evaluates B, determines that its "C (any version)" is satisfiable by the already chosen C 1.0.

The returned artifacts are A, B, and C 1.0.

Note: In the version range syntax [] indicates inclusive, () indicates exclusive.

Send Email Step

Use the EC-SendEmail plugin to create a step to send email notifications.

The following parameters are available to create this step:

Parameter Name	Description
Email Configuration	The name of the email configuration to use. If no configuration is specified, the configuration named "default" is used.
То	List the "To" recipients for the email message, one per line. A recipient can be a user or group name or a complete email address.
СС	List the "CC" recipients for the email message, one per line. A recipient can be a user or group name or a complete email address.
BCC	List the "BCC" recipients for the email message, one per line. A recipient can be a user or group name or a complete email address.
Subject	The subject line for the email message.
Message	The body of the email message or a workspace file containing the body. If both a plain text and an HTML message are provided, both values are sent as alternates in a multipart message. If a raw message is provided, both values are sent as alternates in a multipart message. If a raw message is provided, the value should be a properly formatted RFC822 message.
Advanced options	
Multipart Mode	The multipart mode of the email message. Defaults to "none" unless there are multiple parts, in which case it defaults to "mixedRelated".
Header(s)	One or more RFC822 email header lines (for example: "reply-to: user@host.com").

Parameter Name	Description
Attachment(s)	One or more files from the job's workspace to send as attachments. The filename extension is examined to determine the content-type.
In-line Attachment(s)	One or more inline attachments specified as a contentId and a workspace filename. The filename extension is examined to determine the content-type.

New Extract Preflight Sources Step

Choose one of the preflight plugins to create a step to retrieve an artifact version. There are numerous source code management (SCM) plugins. For example, Accurev, Bazaar, CVS, ClearCase, Git, Mercurial, Perforce, SVN, StarTeam, Team Foundation Server, and Vault.

Enter the following information:

Field Name	Description
Step Name	Enter a unique name for your preflight step. You can use any name of your choice.
SCM Configuration	Use the pull-down menu to select a source control configuration. If a configuration for your source control type does not appear in the drop-down menu, select the Administration > Source Control tabs to configure your source control system.
Destination Directory	This field appears after specifying your SCM from the pull-down menu. If this field does not specify "Required", you do not need to enter this information. The Destination Directory is a path relative to the job's workspace, where the source tree will be created.

Click Submit after entering your information.

After clicking Submit, the Edit Step page is displayed.

Use the Edit Step page to specify additional information for your preflight step. Also, you can attach a credential, create an email notifier, add custom properties, and so on to your step. For assistance using the Edit Step page, access the **Help** link in the upper right corner of the page.

For more information about preflight builds, see the Preflight Builds Help topic.

Parameter - create new or edit existing parameter

To create a new parameter

Enter information in the fields as follows:

Field Name	Description / Action
Name	Enter a unique name to specify the parameter when a procedure or workflow is invoked.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <o>> <pre> <pre> <style> <tc> <u>></td></tr><tr><td>Type - Select the parame</td><td>ter type from the drop-down menu. The following "types" are available.</td></tr><tr><td>Text entry</td><td>Allows a short text entry in the Default Value field.</td></tr><tr><td>Text area</td><td>Expands the Default Value field to allow adding a longer script-type entry.</td></tr><tr><td>Dropdown menu</td><td>Creates a drop-down menu from which to select a value when the parameter is presented. • Enter options - Click + Add Option to add a new row. Type-in the text and value for each option. The text is what will be displayed in the menu, and the value is the parameter value if that option is selected. • Load options from list - Enter a pipe-separated list of options (for example, foo bar baz). The text and value for the options will be the same. • Load options from property sheet - Enter the path to property sheet that contains options for the parameter. The property sheet must be created in a specific format: • An optionCount property must exist whose value is the number of options. • For each option, create a nested property sheet called optionN, where N is the option number, starting with 1. • In each nested sheet, create two properties - text and value. The value of the property text will be displayed in the menu. The value of the property value is the parameter value if that option is selected. • If optionCount is set to 3, you must create three nested sheets - option1, option2, and option3.</td></tr></tbody></table></style></pre></pre></o>

Field Name	Description / Action
	Creates "radio" buttons to select an entry when the parameter is presented.
	 Enter options - Click + Add Option to add a new row. Type-in the text and value for each option. The text is what will be displayed in the menu, and the value is the parameter value if that option is selected. Load options from list - Enter a pipe-separated list of options (for example, foo bar baz). The text and value for the options will be the same. Load options from property sheet - Enter the path to property sheet that
	contains options for the parameter. The property sheet must be created in a specific format:
Radio selector	 An optionCount property must exist whose value is the number of options.
	 For each option, create a nested property sheet called optionN, where N is the option number, starting with 1.
	 In each nested sheet, create two properties - text and value. The value of the property text will be displayed in the menu. The value of the property value is the parameter value if that option is selected.
	 If optionCount is set to 3, you must create three nested sheets - option1, option2, and option3.
	Creates a checkbox for a value to select (or not) when the parameter is presented.
	Value when unchecked - The value of the parameter when the checkbox is unchecked.
Checkbox	Value when checked - The value of the checkbox when the checkbox is selected.
	 Initially checked - Whether or not the checkbox should be checked initially. If true, set the "Default value" to match the "Value when checked". If false, set the "Default value" to match the "Value when unchecked".
Credential	Selecting this option requires the user to specify a user name and password to use this parameter at runtime. Note: Credential parameters are not supported on state definitions.
Project	Creates a project selector field on the Run Procedure page to choose a different project name where the parameter needs to find additional information.

Field Name	Description / Action
Saved Filter	Use this option to reference any saved filters, which is a passing the filter into the procedure at runtime (primarily used for reports). Saved Filters are stored as a property in your chosen project.
Default Value	You can specify a default value to assign to the parameter if no explicit value is provided.
Required?	Click the checkbox to select the parameter as "required." If the parameter is required, the procedure or workflow will not run without entering a value for that parameter. In this case, the default value is ignored. After the procedure or workflow begins execution, the parameter value is available in the job or workflow's property sheet with the same name as the parameter.
Defer Expansion?	Click the checkbox to set the parameter expansion to be "deferred." Deferred means the parameter value will not be expanded when the procedure call is expanded, but can be expanded from a command step instead.

Click **OK** after filling-in the fields.

To edit an existing parameter

You may change any information or add new information as necessary.

For example: Want to rename a parameter? Enter a new parameter name in the Name field, click **OK**.

For more information on property sheets, see the Properties Help topic.

Projects

This page displays all projects available on this ElectricFlow server—the projects you create, ElectricFlow-supplied default projects, and ElectricFlow solutions. From this page you can navigate into projects to view its procedures, steps, schedules, credentials, workflows, and other components.

A project is a top-level container within ElectricFlow. Most information about software production processes, such as procedures, schedules, jobs, and workflows are contained within a project.

• The drop-down menu (at the top of the table) displays all "tags" you previously defined to mark, label, or group related or similar type projects. When you select a tag, the project list changes to display only the projects with that tag.

To add a tag:

- 1. Click the Edit link for the project you want to tag.
- On the Edit Project page, enter a tag name in the Tags field.After creating a tag, the tag name will appear in the drop-down menu.

Note: For more information on "tagging" projects, see the Edit Project Help topic.

- To create a new project, click the Create Project link.
- To Edit, Copy, or Delete an existing project, click a corresponding link in the Actions column.

Note: If deleting a project, the "background deleter" will mark the project for deletion, but the deletion process may not start until the server completes other tasks already queued or in progress. Deleting a project could take considerable time depending on the size of your project. For example, a project containing hundreds of jobs, steps, and other objects will take longer to delete than a simple project with 3 procedures and 10 steps.

• To see project details, click on a project name in the Project column to go to the Project Details page.

Projects have two purposes

1. Projects allow you to create separate work areas for different purposes or groups of people so they do not interfere with each other.

For example, different projects can reuse the same names internally without conflict, and each project has its own access control that determines who can use and modify the project.

In a small organization, you might choose to keep all work within a single project, but in a large organization, you may want to use projects to organize information and simplify management.

2. Projects simplify sharing.

You can create library projects containing shared procedures and invoke these procedures from other projects.

After creating a library project, you can easily copy it to other ElectricFlow servers to create uniform processes across your organization.

ElectricFlow includes default projects

Currently, ElectricFlow includes the following projects:

• EC-Utilities—This project contains procedures to perform basic ElectricFlow tasks. Also, these procedures can be used as examples—copy a procedure to another project and then modify it for your purpose.

Note: By default, only the "admin" user has execute privileges on the EC-Utilities project. The admin user can enable privileges for additional users or groups.

- EC-Examples—This project contains templates for procedures that perform basic ElectricFlow tasks.
- Electric Cloud

 —This project contains procedures to manage ElectricSentry, specifically the Sentry

 Monitor for schedules and also contains ElectricFlow global reports.

Procedures within each of these projects are maintained by Electric Cloud.

IMPORTANT: When you create project names, do not use "Electric Cloud" or any other ElectricFlow-supplied project name. Duplicating a project name creates a conflict.

Also, the "EC-" prefix is reserved for ElectricFlow-supplied project names.

ElectricFlow Solutions

ElectricFlow solutions are developed in the field by Electric Cloud SEs—our field team has the best experience in applying ElectricFlow "best practices" to solve real problems. In addition, ElectricFlow solutions can be contributed by ElectricFlow users from any customer site.

Potentially, any available solution can be considered a new ElectricFlow feature "waiting" to be included in the product if that solution gains wide, general appeal within the ElectricFlow customer base. However, because solutions are not part of the ElectricFlow product, and not developed by the ElectricFlow Engineering team, they are not supported by the Electric Cloud Technical Support team. If and when a solution becomes an integral part of the product, that solution will then have full Technical and Engineering support like any other ElectricFlow feature.

Currently, ElectricFlow solutions are posted to the Electric Cloud Forum Site at http://forums.electric-cloud.com/ecforums/ - each solution, accompanied by its own installation instructions and Help file, can be downloaded from the forum site. If you have a solution you would like to contribute, the Electric Cloud / ElectricFlow Forum is where you need to post it.

When you install a solution on your ElectricFlow server, it will be displayed on the Projects web page as a project.

Installing an ElectricFlow Solution

Add new or upgrade existing ElectricFlow solutions by running the built-in setup script:

- 1. Download the ElectricFlow <solution name>.jar file from the Electric Cloud Community Forums web site.
 - For this example, the "solution name" is Test.
- 2. Save your download to a temporary directory on the system that hosts your ElectricFlow server.
- 3. Unpack the jar file in the temporary directory—the file creates a directory named Test.
 - If Java JDK is installed, use the included jar utility to unpack the .jar file: %jar xf Test.jar
 - For Linux, use any unzip utility, for example: unzip Test.jar
 - For Windows, use Windows Explorer to access the jar contents if you rename the file to Test.zip
- 4. Run the included setup.pl script in the jar through the ec-perl interpreter (installed in the bin directory). You must have "admin" or equivalent privileges to run the setup script.

```
ectool login admin
%cd <test>
%ec-perl setup.pl
```

Project - Create New or Edit Existing Project

Creating a New Project

Enter information in the fields as follows:

Field Name	Description
Name	Enter a unique project name. You may want your project names to describe the work groups or teams that will be using them. For example, you may set project names based on the products they support.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a>
Default Resource	Enter a resource name, or select Browse to find one. If you have not set up resource names or locations yet, click the Resources tab and then the Resources tab to see the Resources page.
	Note: This field is optional. If you specify a resource in this field, it will be used as the default for all jobs that run under the project. This is a convenient way to use a single resource for an entire project. You can change the resource by changing the value. There are several other places where you can specify resources, such as procedures and individual job steps. If you specify a resource in another location, it will take precedence over the one specified in this field.
Default Workspace	Enter a workspace name, or select Browse to find one. If you have not set up workspace names or locations yet, click the Resources tab and then the Workspaces tab to see the Workspaces page.
	Note: This field is optional. If you specify a workspace in this field, it will be used as the default for all jobs that run under the project. This is a convenient way to use a single workspace for an entire project. You can change the workspace by changing the value. There are several other places where you can specify workspaces, such as procedures and individual job steps. If you specify a workspace in another location, it will take precedence over the one specified in this field.

Click **OK** to save the information you entered.

Your new project name will appear on the Projects page.

Editing an Existing Project

- To rename the project, enter a new name in the project Name field and click **OK**.
- You may modify the description or enter a text description if this field is blank.
- You can modify any previously entered information.
- If you would like to use a particular credential (by default) for all jobs that run under this project, select it here. If you leave this field blank, credential information from other sources is used.

Tags - (Optional) Enter a tag if you want to categorize or "mark" this project to identify its relationship to one or more other projects or groups. For example, you may want to tag a group of projects as "production" or "workflow", or you may want to use your name so you can quickly sort the project list to see only those projects that are useful to you.

- The tag name you enter in this field is displayed in the drop-down menu on the Projects page. When you select a specific tag from the drop-down menu, you see only the projects with that tag.
- Assigning a tag creates an "ec_tag" property that can be seen (with its value the tag name you chose)
 after selecting the Properties subtab from the Project Details page,
- Do not use "spaces" in tag names. For example, if you want to mark a project with the tag name, "Eng Team A", you need to specify, "eng team a".
- Do not use symbols in a tag name—results are unpredictable.
- If a project is related to several project categories, you can enter a space-delineated list to add numerous tags at the same time.
 - Type a new name in the project Name field and click **OK**.
- To delete a tag, select the text and use your Delete key.

Impersonation Credential

Credential Project - The default Credential Project is the current project, or select the other button and Browse for the project name you need.

Credential Name - This is the name of the credential for the project. You can select Browse to find it.

Click **OK** after entering new or updated information.

For more information, see the Credentials and User Impersonation Help topic.

Project Details

This page displays project components, including procedures, workflow specifications, jobs, schedules, any credentials, properties, and reports.

Links and actions at the top of the page

- Edit use this link to go to the Edit Project page.
- . Access Control use this link to go the Access Control page for this project.
- The "star" icon allows you to save this job information to your Home page.
- The "curved arrow" icon returns you to the Projects page.

 Pagination - Use the "previous" and "next" arrow icons to view the previous or next project. The numbers between the arrow icons display the number of projects you can view and the first number indicates which project [in the list] you are viewing.

The "tabbed" sections

The tabs at the top of the table allow you to select the type of information you want to see. See the 8 tabs illustrated in the following screen example:



The Project Details page opens with the Procedures tab highlighted, so the first table you see is the Procedures table.

Procedures tab

The following links and actions are available in the Procedures table:

- Click the **Create Procedure** link (at the top of the table) to go to the New Procedure page to create another procedure.
- Click on a procedure name (first column) to go to the Procedure Details page for that procedure.
- Action column click on any of the available links (Run, Edit, Copy, or Delete) to perform that function for the procedure listed in that row.

Note: The Run link provides two choices when you hover the mouse over the down-arrow. Select **Run Immediately** or **Run...** .

- Run Immediately This option runs the procedure "as-is" immediately.
- Run... This option produces the Run Procedure web page where you can modify any existing parameters for this procedure, or set an existing credential.

Workflow Definitions tab

This tab provides a table listing all defined workflow definitions and includes the following functionality:

- Click the Create Workflow Definition link to go to the New Workflow Definition page to define additional definitions.
- Name column Click on a workflow definition name to go to the Workflow Definition Details page (for that workflow definition.
- Action column
 - Run use the down-arrow to select Run... or Run Immediately
 Selecting Run... takes you to the Run Workflow page where you can set the starting state for the workflow.
 - Selecting Run Immediately runs the workflow "as-is", immediately.
 - **Copy** use this link to make an exact copy of an existing workflow definition.

- Edit use this link to edit an existing workflow definition or just to change the name of the copy you created.
- **Delete** deletes an existing workflow definition on the same row.

Jobs tab

This tab provides a project-specific job table. Selecting the Jobs tab on the Project Details page displays only those jobs for the project you selected.

This table is a subset of the main Jobs tab that displays **all** jobs, regardless of the project of origin, and all Jobs page functionality is basically the same on either Jobs tab.

The following links and actions are available on this Jobs page:

- Job click on Job to sort the column alphabetically or click on a job name to go to that job's Job Details page.
- Status click on Status to sort this column by Running, Success, Warning, Error, or Aborted.
- Priority displays the job's priority set by a "run procedure" command or by the job's schedule.
- Procedure click on a Procedure name to go to the Procedure Details page for that procedure.
- Launched By click on a name in this column to go to the page for the schedule (Edit Schedule page) or the user that ran the job.
- Elapsed Time click Elapsed Time to sort the time values from longest to shortest time, or the reverse.
- Start Time click Start Time to sort this column from the start time of the first job to the start time of the most recent job, or the reverse.
- Actions use this column to Abort a running job or Delete a completed job. If your job has finished running, you will not see the Abort link.
 - Aborting a job requires the execute privilege on the job—not just the modify privilege.

When you click **Abort**, the following dialog box appears:



Select either **OK** or **Force** to abort the job, or **Cancel** if you change your mind.

Note: Deleting a job on this page causes removal of job information from the ElectricFlow database, but information in the job's on-disk workspace area is not affected. You must delete workspace information manually.

Workflows tab

The Workflows tab provides the following functionality:

- Name column Click on a workflow name to go to its Workflow Details page.
- State column Click on a state name to go to its State Details page.
- Modify Time This time value represents the most recent workflow activity.
- Workflow Definition column Click on a workflow definition name to go to the Workflow Definition Details page.
- Completed column a "checkmark" in this column identifies this workflow as completed.
- Actions column Click the Delete link to remove the workflow in that row.

For more information about workflows, see the Workflow Help topic.

Schedules tab

The Schedules tab provides the following functionality:

- To create a new schedule for this project, click either the Schedule link to create a new standard schedule to run at specific and/or days or the CI Configuration link to setup a new continuous integration schedule.
 - o Choosing "Schedule" displays the New Schedule page.
 - o Choosing "CI Configuration" displays the New CI Configuration page.

Table column descriptions

- To edit an existing schedule, click on the schedule name (first column).
- Select the check box in the Enabled column to enable or disable the schedule in that row.
- The Priority column displays the job's priority you selected while creating the schedule for this procedure.
- · Action column links -
 - Run runs the schedule with the permissions of the logged in user.
 For example, if this schedule is timed to run every day at 11:00 pm, you might decide to run the schedule at 6:00 pm on a particular Thursday, without changing the regular schedule settings.
 However, if you do not have the appropriate permissions to run this schedule, it will fail.
 - Copy makes a copy of the schedule on that row.
 - o Delete deletes the schedule in that row.

The following screen example illustrates the Schedules table:



For information on creating a standard schedule, see the Schedule - create new or edit existing schedule Help topic.

For information about CI Manager and adding a configuration, click the Home tab, the Continuous Integration subtab, then click the Help link in the upper-right corner of the page.

Credentials tab

The Credentials tab provides the following functionality:

- To create a new credential for this project, click the Create Credential link at the top of the table.
- To edit an existing credential, click on the credential name (first column).
- Action column the Delete link deletes the credential on that row.

For information about credentials and access control, see the Credentials and User Impersonation or Access Control Help topics.

Properties tab

The Properties tab provides the following functionality:

- To create a new property for this project, click the Create Property or Create Nested Sheet link.
- To view or change privileges on the property sheet, click the Access Control link.
- To edit an existing property, click on the property name (first column).
- Action column the Delete link deletes the property on that row.

For more information about properties, see the Properties Help topic.

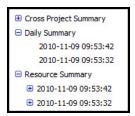
Reports tab

Selecting this subtab yields different results, depending on whether or not you have installed the additional ElectricFlow-supplied optional reports or have created custom reports for the specific information you need. Generally, if you have not yet configured additional reports, you will see the three installed, default ElectricFlow reports (Cross Project Summary, Daily Summary, and Resource Summary) as illustrated in the screen example below.

- When a report runs, the report "folder" is populated with the completed report.
 Reports are listed by date and time.
- Each individual report "date/time" entry is a link to see the full report.

• To create a new report, click the **Create Report** link to go to the web page to create a Multiple Series, Category, Count Over Time, or Procedure Usage report.

This is an example of the information on the Reports page :



For more information about reports, see the Reports Help topic series.

Tips

Want to rename this project?

Click the **Edit** link at the top of this page and enter a new name.

Want to rename a procedure in this project?

Click the **Edit** link for the procedure (on this page) and enter a new name.

Run Procedure

Using the Run Procedure page

- You may see a message informing you that no parameters were defined for this procedure. If this is the case, click **Run** to run the procedure.
- If parameters were defined for this procedure, you will see specific fields based on those parameters.
 You may make changes to any field you choose if you want to run this procedure under different
 parameters at this time. However, any changes you make are specific to this run procedure only. Your
 changes on this web page will not affect previously defined parameters for this procedure.
 Click Run to run the procedure.
- You may see a "Required" field. To run the procedure, you must enter the required parameter previously defined on the Procedure Details page.

Click Run to run the procedure.

Advanced section

Priority:

Use the drop-down menu to select the priority for this run procedure. Available priorities are: low, normal (default), high, or highest.

• You can select the job's priority here on the Run Procedure web page or when you schedule the procedure to run at a regular time or interval. When a job is launched, its priority cannot be changed.

- Priorities take effect when two or more job steps in different jobs are waiting for the same resource.
 When the resource is available, it will be used by the job step that belongs to the job with the highest priority. If the priority level is the same, the resource will be used by the job step that belongs to the job with the lowest job ID number. If the job steps are in the same job, the resource will be used first by the step with the lowest job step ID number.
- If using ectool, the priority can be set by passing "--priority=<low|normal|high|highest>" to a runProcedure, createSchedule, or modifySchedule operation.

Note: To raise the job priority level above "normal," the user who launches the procedure must have execute permission on the system-level access control list for priorities.

Impersonation:

Select one of the following:

- **Use pre-defined credential** (Default option) Use this option if you have not defined a credential for this procedure.
 - ElectricFlow looks for a credential first in the top-level procedure and then in its project. If it finds a credential, it uses the login from that credential as the default for the job. If there is no credential in either place, by default, the job's steps run as the user under which the agent for the step is running.
- Use specific credential If you select this option, more fields are displayed for you to choose a particular credential and enter its name from those defined for the project containing the procedure. The login user from that credential will be used as the default for the job.
- **Use a specific user** If you select this option, more fields are displayed to enter a User Name and Password. The specified login is used as the default for the job.

Click Run to run the procedure.

For more information on credentials and user impersonation, see the Credentials and User Impersonation Help topic.

Schedule - create new or edit existing schedule

Use this page to create or edit a standard schedule that specifies a set of times and days when a particular procedure should run. If you need to create a continuous build integration instead, see the Help information available from the Continuous Integration Dashboard web page.

For any schedule, click the "star" icon to add this schedule to the Shortcut section on your Home page.

To create a new standard schedule

Enter information in the fields as follows:

Field Name	Description
General section	

Field Name	Description
Name	Enter a Schedule name. This must be a unique name among other schedule names in this project.
Procedure	Displays the current project and procedure. Click Change to use the pop-up menu to select a new project or procedure.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with tags">https://ntml> tags tags">https://ntml> tags tags">https://ntml> tags tags">https://ntml> tags tags tags

Parameters section

If the procedure has parameters, enter information in their values in this section. The procedure cannot execute unless all required parameters are provided. If you move the mouse over a field for a particular parameter, inline Help displays the description of that parameter if one was provided in the procedure's definition.

- Add	Click this icon to displays fields to add one or more parameters.	
Frequency section	Frequency section	
Run at ???	This is a summary section: As you make selections for the following Days and Repetition sections, a summary of what you have selected appears here.	
Days	Select weekdays, month days, or create a custom frequency when you want the schedule to run. Note: The "day" applies to the start time if the time range spans two days, that is, if the time range crosses the midnight boundary.	
Repetition	If Run Once - the schedule will run once per selected day, at the time you specify. If Run Every - the procedure will execute repeatedly during the specified time range. For example, if you select a time range from 11:00-14:00 and an interval of 40 minutes, the procedure will execute 5 times on each of the selected days: at 11:00am, 11:40am, 12:20pm, 1:00pm, and 1:40pm. If Run Continuously - as soon as one job ends, the scheduler will trigger immediately to run the job again.	

Field Name	Description
Advanced section	
Enabled	If this box is "checked," the schedule will run. You can disable this schedule whenever necessary.
Misfire Policy	Choose between two options to manage how a schedule resumes in cases where the normal scheduled time is interrupted. The server may not be able to trigger a job at the scheduled time for several reasons: 1) a running job extends past the next scheduled runtime, or 2) the server cannot find enough available resources and a delay occurs, or 3) the server is down or disconnected from the network for a period that overlaps one or more scheduled times. If any of these situations occur, the misfire policy defines two possible responses: skip - all misfires are ignored and the job runs at the next scheduled time. run once - after one or more misfires, the job runs at the soonest time that occurs within an active region. For example, at server startup a scheduled job will not trigger immediately on startup.
Time Zone	Default is the time zone of the ElectricFlow server, or select a county from the first pull-down menu, then select an area from the second pull-down menu. Another option: Select "other" from the first pull-down menu, then select GMT [or another selection] from the second pull-down list.

Field Name	Description
Priority	Use the pull-down menu to select one of four priorities for this procedure. Also, you can select the job's priority on the Run Procedure web page. When a job is launched, its priority cannot be changed. Priorities take effect when two or more job steps in different jobs are waiting for the same resource. When the resource is available, it will be used by the job step that belongs to the job with the highest priority. If the priority level is the same, the resource will be used by the job step that belongs to the job with the lowest job ID number. If the job steps are in the same job, the resource will be used first by the step with the lowest job step ID number. If using ectool, the priority can be set by passing "priority= <low normal high highest>" to a runProcedure, createSchedule, or modifySchedule operation. Note: Schedules are launched by the project principal user. To raise the priority above "normal," this user must have execute permission on the system-level access control list for priorities.</low normal high highest>

Impersonation Credential section

If you select a credential for the schedule, it will be used as the default for jobs run from the schedule. For example, job steps will run under the login account specified in the credential. However, the schedule's credential is low in priority: it will be used only if there is no credential specified in the procedure or its project. Also, individual steps can override the credential with a specific credential for that step.

For more information, see the Credentials and User Impersonation Help topic.

Impersonation credentials are used to set the top-level impersonation credential for a job. If specified, the impersonation credential [on the job] is used as the default impersonation credential for all steps in the job.

Credential Project	The name of the credential for the project. By default, the current project is selected. Enter a different project name or select Browse to find this name.
Credential Name	The name of the credential for the project - you may need to Browse to find this name.

Click **OK** to save your schedule information.

To create a new continuous integration "schedule"

If you selected the **CI Configuration** link, the New CI Configuration page is displayed. *For information* on populating these fields, see the "Continuous Integration Manager" Help topic - click the **Help** link in the upper-right corner of the New CI Configuration page.

See the Standard Schedule information [above] to enter information in the fields common to both types of schedules.

To edit an existing schedule

Modify any of the information you previously specified.

The following additional links are available at the top of the page:

- Save Configuration if you click this link, the information in this schedule will be copied into a new Job Configuration and displayed on your Home page for easy access.
- Access Control click this link if you need to set permissions for this schedule.

Attached Credentials

Click the **Attached Credential** link if you need to attach an existing credential to this schedule. If you need to create a new credential for this schedule, return to the Project Details page.

Custom Schedule Properties

Select the **Create Properties**, **Create Nested Sheet**, or **Access Control** links if you want to assign properties to this schedule.

For more information on Properties or Access Control, see the following main Help topics: Properties, Access Control

Click **OK** to save your modifications for this schedule.

Credentials - create new or modify existing credential

To create a new credential

Enter information in the fields as follows:

Field Name	Description
Name	Enter a unique name of the credential object.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <div> <dl> <i> <o> <pre> <pre> <style> <tc> <u>><u>></td></tr></tbody></table></style></pre></pre></o></i></dl></div></div>

Field Name	Description
User Name	Enter the user name you choose to use for login.
Password	Enter the password that identifies the user account.

Click **OK** to save your new credential information.

To edit a credential

Modify any previously entered credential information and click **OK** to save your changes.

Use the Access Control link to add or change existing permissions.

You can add Custom Credential Properties for this credential also. Click the **Create Property**, **Create Nested Sheet**, or **Access Control** link in this section if you need to specify any of these items.

For more information, see the following Help topics:

Credentials and User Impersonation

Properties

Access Control

Property - create new or edit existing property

Use this pop-up window to create or modify a custom property attached to an object.

Enter information in the fields as follows:

Field Name	Description
Name	The property name can be an arbitrary text string, but unless you are an experienced ElectricFlow user avoid using slashes and brackets as a part of the property name. Note: The names "properties" and "project" are not valid property names.

Field Name	Description
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pr> <pr> <pr> <pr></pr></pr></pr></pr>
Value	This value can be an arbitrary text string.
Expandable	Do not select this check box if you do not want to allow property expansion.

Click **OK** to save your new or modified property information.

For more information on creating and using properties, see the Properties Help topic.

Nested Property Sheet

Use this pop-up window to define or modify a nested property sheet.

1. Enter information in the fields as follows:

Name-The property name can be an arbitrary text string.

2. Click **OK** to save your new or modified nested property sheet information.

For more information on creating and using properties, see the Properties Help topic.

Reports

If you have upgraded ElectricFlow this page displays historical reports (Cross Project Summary, Daily Summary, and Resource Summary) from your previous ElectricFlow version. These reports are no longer being updated.

Click the Historic/web server local reports link to see any previously created custom reports.

- Cross Project Summary this is a 30-day report summary
- Daily Summary this is a daily report showing all the procedures that ran
- Resource Summary this report shows resource usage for a specific day
- Historic/web server local reports

For more information, see the Reports Help topic.

Create a new report

Use this page to create a new report. Select one of the report-type tabs to begin.

Multiple Series Report

Field/Menu Name	Description
Report Title	This is your report title. Type over the default report name, choosing any unique name for your report.
	IMPORTANT: Do not use special characters in the report title. Some examples of special characters are : /? # [] @!, \$ & '() * + , ; =
Saved Filter	Project - Click your mouse inside this field to see a list of projects from which to make a selection. Filter - After selecting a project, this drop-down menu will contain any available filters for this project.
Time Period	Use the drop-down menu to select the time period.
Create thumbnail?	Check this box if you would like to make a thumbnail view of this report available on your Home page.
Object Type	Use the drop-down menu to select an object.
	Chart Type - Use the drop-down menu to choose the chart type.
Table column choices	Function - Use the drop-down menu to choose the function you need.
	Property Name - Use the default property value or delete the text and click your mouse in the blank field to see a list of possible properties.
	Display Name - You can choose a different unique Display Name.
	Stacked - Select this checkbox to see your report results "stacked" versus overlaid.
	The "X" icon - Click this icon to delete any row you no longer need.

Field/Menu Name	Description
Add Series button	Select the Add Series button if you would like to create additional table entries for additional report information.
Chart Options	Use the down-arrow to adjust the Time Grouping and see the defaults for the X and Y axis.
Table Options	Use the down-arrow to adjust the Time Grouping and see the default for the Time Column label.
Advanced	Use the down-arrow to select a Locale or add a Credential. Delete the Resource if you choose, then click your mouse in the blank field to see a list of other resources you might prefer to use.
Run Report button	Use this button to go to the Job Details page to run the report.
Create Schedule button	This button displays a box with a default schedule name that you can change if you choose. Click OK to go to the Job Details page.
Cancel button	Use this bottom to cancel the create report operation.

Category Report

Field/Menu Names	Description
	This is your report title. Type over the default report name, choosing any unique name for your report.
Report Title	IMPORTANT: Do not use special characters in the report title. Some examples of special characters are : / ? # [] @ !, \$ & '() * + , ; =
Saved Filter	Project - Click your mouse inside this field to see a list of projects from which to make a selection. Filter - After selecting a project, this drop-down menu will contain any available filters for this project.
Time Period	Use the drop-down menu to select the time period.

Field/Menu Names	Description
Create thumbnail?	Check this box if you would like to make a thumbnail view of this report available on your Home page.
Object Type	Use the drop-down menu to select an object.
Object Property	You can delete the default value and click your mouse inside the blank field to see a list of properties from which to choose.
Advanced	Use the down-arrow to select a Locale or add a Credential. Delete the Resource if you choose, then click your mouse in the blank field to see a list of other resources you might prefer to use.
Run Report button	Use this button to go to the Job Details page to run the report.
Create Schedule button	This button displays box with a default schedule name that you can change if you choose. Click OK to go to the Job Details page.
Cancel button	Use this bottom to cancel the create report operation.

Count Over Time Report

Field/Menu Name	Description
	This is your report title. Type over the default report name, choosing any unique name for your report.
Report Title	IMPORTANT: Do not use special characters in the report title. Some examples of special characters are : / ? # [] @ !, \$ & ' () * + , ; =
Saved Filter	Project - Click your mouse inside this field to see a list of projects from which to make a selection. Filter - After selecting a project, this drop-down menu will contain any available filters for this project.
Time Period	Use the drop-down menu to select the time period.

Field/Menu Name	Description
Create thumbnail?	Check this box if you would like to make a thumbnail view of this report available on your Home page.
Object Type	Use the drop-down menu to select an object.
Chart Options	Use the down-arrow to adjust the Time Grouping and see the defaults for the X and Y axis.
Table Options	Use the down-arrow to adjust the Time Grouping and see the default for the Time Column label.
Advanced	Use the down-arrow to select a Locale or add a Credential. Delete the Resource if you choose, then click your mouse in the blank field to see a list of other resources you might prefer to use.
Run Report button	Use this button to go to the Job Details page to run the report.
Create Schedule button	This button displays box with a default schedule name that you can change if you choose. Click OK to go to the Job Details page.
Cancel button	Use this bottom to cancel the create report operation.

Procedure Usage Report

Field/Menu Name	Description
	This is your report title. Type over the default report name, choosing any unique name for your report.
Report Title	IMPORTANT: Do not use special characters in the report title. Some examples of special characters are : /? # [] @ !, \$ & '() * + , ; =

Field/Menu Name	Description
	Click your mouse inside this blank field to see a list of projects. This does not have to be the current project.
Project	You can generate reports that show when one project's procedures call procedures in another project. You can also use the '%' wildcard, for example,
3	Project: %
	or
	Project: D%
Procedure pattern	Enter a SQL string match pattern for procedures to include in the report. The default % will use all procedures.
Advanced	Use the down-arrow to select a Locale or add a Credential. Delete the Resource if you choose, then click your mouse in the blank field to see a list of other resources you might prefer to use.
Run Report button	Use this button to go to the Job Details page to run the report.
Create Schedule button	This button displays box with a default schedule name that you can change if you choose. Click OK to go to the Job Details page.
Cancel button	Use this bottom to cancel the create report operation.

Resources

The **Cloud > Resources** page displays all resources known to this ElectricFlow server and provides easy access to resource configuration and all other resource management functionality.

- Each resource has a logical name—a unique name to distinguish this resource from other resource names.
- Each resource refers to an agent machine by its host name.
- Each resource can be assigned to one or more pools.

A pool is a group of interchangeable resources. For example, you may have a pool of Windows servers.

If you name a pool in a procedure step, ElectricFlow can assign any resource in the pool to that step, which allows ElectricFlow to choose a lightly loaded resource. You can change the resources in a pool without having to modify any procedures that use the pool.

• Each resource can be assigned to a zone. A default zone is created during ElectricFlow installation and all resources, by default, are members of that zone until they are assigned to a different zone.

A *zone* is a collection of agents. Every agent, and all resources defined on that agent, belong to only one zone. When ElectricFlow is installed, a *default* zone is created. If a zone is deleted, all agents in that zone are moved to the *default* zone. The ElectricFlow server resides in the default zone.

Note: The default zone cannot be deleted.

- Several resources can correspond to the same physical host or agent machine.
- When you specify a resource name in a procedure step, the step executes on that resource.

Supported Resource Categories

- **Standard**-This category specifies a machine running the ElectricFlow agent on one of the supported agent platforms, as specified in the *ElectricFlow Installation Guide*.
- **Proxy**–This category requires SSH keys for authentication. You can create proxy resources (agents and targets) for ElectricFlow to use on other remote platforms or hosts that exist in your environment.
 - Proxy agent—This is an agent on a supported Linux or Windows platform, used to proxy
 commands to an otherwise unsupported platform. A proxy agent is an ElectricFlow agent,
 channeling to a proxy target.
 - Proxy target This is a machine on an unsupported platform that can run commands through an SSH server. Proxy targets have limitations, such as the inability to work with plugins or communicate with ectool commands.

Help topic "Best Practice": We recommend reviewing this Help topic in its entirety to become familiar with new resource management features and expanded functionality. Later, when using the Resources page, use the following quick links to go to the Help section that you need to review.

Icons, Links, and Buttons Above the Table

Table view - column descriptions

Grid view

Filters pane

New Resource panel

New Proxy Resource panel

Edit Resource panel

Edit Proxy Resource panel

Resource Details panel

Switching a non-trusted agent to trusted

Install or Upgrade Remote Agents

Resource Page Information and Functions

Immediately after the *Resources* title above the table, you can see how many licensed resources are currently in use.

Links, Icons, and Buttons Above the Table

Hover your mouse over an icon or button to see the available actions. You can select one of these actions on a single resource or on multiple resources (selected using the check box in the first column) to perform the action on one or more resources at the same time.

• Select one or more resources from the table, and then select the button to perform the task or action.



• Exception: When you click the Action icon and select Create Resource or Create Proxy Resource, a panel opens where you can create one of these new objects. You can create only one object at a time.

Icon or Link Name	Description
+ *	Create Resource—The New Resource panel opens. For information about the fields, see the New Resource section. Create Proxy Resource—The New Proxy Resource panel opens. For information about the fields, see the New Proxy Resource section. Install Resource(s)—A dialog box opens where you enter information about the agents to install. For more information, see Install or Upgrade Remote Agents on page 613. Note: If you are upgrading one or more agents or resources, the label on this option is Upgrade Resource(s), which is available only for Linux agents. You cannot install or upgrade Windows agents using this process. IMPORTANT: If the EC-AgentManagement plugin is not installed, you will not see this option.
면	Copy Resource–Makes a duplicate copy of one or more resources that you selected.
TÍT	Delete Resource Deletes one or more resources that you selected. Note: If a resource is on a "gateway agent", the gateway will be deleted.

Icon or Link Name	Description
	Ping Resource ElectricFlow pings one or more resources to check if they are available.
©	When ElectricFlow "pings" a resource, it sends a message to the agent (for the resource) to make sure the agent is alive and running a version of software compatible with the ElectricFlow server. After the ping completes (or fails), the page refreshes to reflect the current resource state.
	Note: A gateway must exist before you ping a resource in a remote zone.
~	Enable Resource Enables one or more resources that you selected in the check boxes in the first column
_	Disable Resource –Disables one or more resources that you selected in the check boxes in the first column.
	Select one of these actions:
	When you click the button with no resources selected, the Select some resources to edit message appears.
	When you select resources and click this button, these options appear:
	Add to Pool Adds the selected resources to a pool.
	 Remove from Pool —This is available only if all the selected resources have at least one pool in common. If not, this link becomes Remove from Pool <name>.</name>
	Move to Zone–Moves the selected resources to a zone.
	Set Trusted—This is not available if the selected resources are already trusted. If not, all the selected resource will be changed to trusted.
	 Unset Trusted—This is available only if all the selected resources are trusted. If not, the selected resources will be changed to untrusted.
☆ ☆ or	Select this button to add this Resources page to your Home page for one-click access to return to this page. If the icon is yellow, the page is already accessible from your Home page.
New Search link	Use this link to go to the Define Search page if you need to search for a particular object or set of objects. For example, job, resource, artifact, workflow, and so on. Use the "back" button to return to the Resources page.

Table View -- Column Descriptions

Not all resource configuration information is displayed within the table, but you can see all other information about the resource when you edit the resource or access other slide-out panels. For example, the action to see or add access control privileges to a resource is available on the Edit Resource panel or from the **Access**Control link on that panel.

Note: Most column headings can be sorted. Click on a column heading to re-sort the information in that column.

Column name / icon	Description
	Select this check box to enable the corresponding resource in that row for a single or batch action. For example, deleting or enabling multiple resources at the same time is a batch action. Note: If you select this icon in the table header row, it will select or deselect all resources in the table.
**	Status of the agent. If an agent machine is not available, any resources associated with that host will not be available. Note: If an agent is down, the ElectricFlow server is not able to provide additional information about the cause.

Column name / icon	Description
Name	Name of the resource. The icon to the left of the resource name indicates
	whether this resource is <i>enabled</i> or <i>disabled</i> . These icons are: Enabled or Disabled
	If a resource is disabled, no job steps are assigned to it. If a step requests a particular resource and that resource is disabled, the step waits until the resource becomes enabled again. If the resource is in a pool, ElectricFlow uses any enabled resource in the pool to satisfy the request for the resource. Other actions:
	To change the enabled or disabled icon, select the check box (in first column) in the row for one or more resources, and then select the enabled or disabled icon above the table. Select multiple resources if this is a batch change.
	Note: You can change the resource enable/disable specification from the Edit Resource panel too, but this is an individual operation for one particular resource.
	Click on a resource name to see the Resource Details panel - this panel displays how your resources is configured currently. To make changes, click the Edit link to go to the Edit Resource panel.
	A space-separated list of pool names where this resource is a member.
Pools	Note: Spaces are not allowed within a pool name.
Туре	 A resource can be static or dynamic. A static resource is part of your system or network (not in the cloud), such as a server, database, or agent machine. A dynamic resource is a cloud resource that can be provisioned and later spun up on-demand when an application is deployed.
Resource Template Name	The name of the resource template to which the resource is applied. Note: Only dynamic resources can be applied to a resource template.
Time running	If the resource is a dynamic resource, how long it has been running.

Column name / icon	Description
	Indicates the current job status with the following icons:
	Running job - a job is still running, which means the resource is currently being used for one or more jobs. The names of all jobs using the resource are displayed next to the icon.
	Success - the job running on this resource completed successfully.
Job Status	Error - the job running on this resource encountered an error and may or may not have completed. One or more error messages will be listed in this column.
	Warning - the job running on this resource encountered a warning and may or may not have completed. One or more warning messages will be listed in this column.
	Busy - ElectricFlow is attempting to refresh the job status, but the resource is not responding or the UI cannot parse the response. Changing the value of 'Agent Host Name' to the IP address can sometimes resolve this issue.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <style> <tc> <tr</td></tr><tr><td>Zone</td><td>The name of the zone where this resource is a member.</td></tr><tr><td>HTTPS & Host</td><td>The "Connection Type" you specified when this resource was created.</td></tr></tbody></table></style>

Column name / icon	Description
Step Load	Resource <i>load</i> is indicated by one of the following:
	- this notation indicates 1 step is using this resource and an unlimited number of steps can use this resource at the same time. If 5 steps are using this resource, you would see "5 of unlimited".
	- this notation indicates that no steps are running on this resource and the step limit is set to 4. The bar indicates graphically how close you are to the step limit.
	Note: Use the Create/Edit Resource panels to define the number of steps that can run on a resource. Setting the Step Limit to "0" or leaving the field blank defaults to NO step limit.
	Actions you can perform on the resource, including:
Actions	Copy-Create a new resource by copying the this resource.
	Delete-Delete this resource.
	Tear Down –Remove (decommission) the dynamic resource from a dynamic environment.
	Track ChangesOpen the Change History for the resource.

Grid View



Two icons above the Filters pane:

- The first icon (on the left) is the default *Table view* icon.
- The second icon (on the right) displays the resource *Grid view*. Use the drop-down menus to filter the Grid view:

- Grouping Select the resource "grouping" you want to see—choose Pool, Job, Operating System,
 Platform, Proxy Agent, Version, or Zone.
 For example, if you have multiple resource pools defined and chose Pool, you will see groups
 labeled with your various pool names and a group (No Pool) for resources not included in a pool.
- Coloring Select what you want to view in each group—choose Platform, Step Count/Limit, Status,
 Held Exclusive, Version, or Operating System to sort your group.

Note: Colors are dynamically assigned per the view you select. Within a single view, the same color "square" implies similar information.

• Hovering your mouse over a grid "square" provides more detailed information about that resource.

Filters pane

Use the Filters pane to filter resources shown in either the Table or Grid view. The fields in the Filter pane correspond to columns in the Table view.

Enter the Filter criteria as follows:

Actions / Field	Description
Save Filters	Click this link to save your filter criteria. A dialog box is displayed so you can name your filter. For easy, repeat access, filters are saved in a "saved filters" box above the link. When you select a saved filter name, the Filter pane fields populate automatically.
Reset	Use this link to clear all entries in the Filter pane.
Quick Search field	Enter any text for your search. For example, enter a host name. You do not need to enter a full name or text string - this field filters on a partial text entry, sometimes 2-3 characters are all you might need.
Status	The status for the agent machine.
Drop-down menu	Default is both enabled and disabled resources. Use the drop-down arrow to choose enabled or disabled only.
Pools field	List one or more pool names, separated by a space. (No spaces are allowed within a pool name). Pool names must be entered exactly - a shorten version of the pool name will not yield a successful result.
Hosts field	List one or more full host names as shown in the Table view "HTTPS & Host" column (each name separated by a space). Entering a partial name will not yield a successful result.

Actions / Field	Description
Step Limit drop-down menu	Choose the step limit corresponding to your desired search.
Proxy Agent drop- down menu	Select No if not searching for a proxy agent or Yes if you are trying to locate a proxy agent.
Filter button	Select this button after you complete the fields/selections to define your filter criteria.

New Resource Panel

To define a new standard resource, enter information in the following fields:

Field	Description
Name	Enter a unique name for this resource. Do not use "spaces" in a resource name. This is the name used to select the resource for a job step—it need not be the same as the host name for the machine.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <dl> <i> <pr> <style> <tc> </td></tr><tr><td>Agent Host Name</td><td>Enter the domain name or IP address of the agent machine for this resource. This is the name all other machines in this agent/resource's zone will use to communicate with this agent host.</td></tr><tr><td>Agent Port Number</td><td>Enter the port number to use when connecting to the agent for the resource, Default port is 7800.</td></tr><tr><td>Default Workspace</td><td>Enter the workspace name or leave blank to use the local, default workspace. The workspace specification can be overridden at the project, procedure, or step level.</td></tr><tr><td>Pool(s)</td><td>A space-separated list of arbitrary names indicating the pools associated with this resource. Do not use "spaces" in a pool name.</td></tr></tbody></table></style></pr></i></dl></div>

Field	Description
Host Type	Use the drop-down menu to select the host type:
	Concurrent
	Registered
	Depending on the license installed on the ElectricFlow server, this field may appear.
	 If the license on the server is a concurrent resource license, the host type defaults to Concurrent and this field does not appear.
	 If the license on the server is a registered host license, the host type defaults to Registered and this field does not appear.
	 If the license on the server is a mixed-mode license (concurrent resources and registered hosts), you must specify the host type when adding a resource.

Field	Description
	The name of the shell program used to execute the step's commands on a resource. For example, using sh or cmd /c, the agent saves the command block to a temporary file with a .cmd extension and runs it: sh foo.cmd or cmd /c foo.cmd.
	If you do not specify a shell on a step, at step run-time the server looks at the resource shell. If a resource shell is not set, the shell line used by the agent is platform dependent: Windows: cmd /q /c "{0}.cmd" UNIX: sh -e "{0}.cmd"
	When you specify a shell (in the step or resource), and omit the <code>cmd-file</code> marker, the agent notices the omission and takes the correct action. For example: a user specifies <code>sh -x</code> . The agent converts this to <code>sh -x " {0}.cmd"</code> .
	Two alternate forms of shell syntax where ElectricFlow uses a "marker," {0}, as a placeholder for the command file argument:
Default Shell	 <myshell> {0} <potential args="" extra="" shell=""> In this example, the command file is not meant to be the last argument in the final command line. For example, mysql -e "source {0}" This shell example runs the mysql command against this step's command containing sql.</potential></myshell>
	• <myshell>{0}<.file extension> <potential args="" extra="" shell=""> In this example, the shell requires the command file to end in an extension other than .cmd. For example, powershell "& '{0}.psl'" This shell example runs Microsoft PowerShell against this step's command containing PowerShell commands.</potential></myshell>
	Notes:
	When the agent parses the shell, it will parse the extension as everything after {0}. until it sees a space or non-alphanumeric character.
	If your script uses International characters (non-ascii), add the following block to the top of your ec-perl command block: use utf8; ElectricCommander::initEncodings

Field	Description
	The maximum number of steps that can execute simultaneously on this resource—a step limit applies to a particular resource only, not to its underlying host.
Step Limit	For example, if you define two resources, resource1 with a step limit of 5 and resource2 with a step limit of 1, both specifying the same host machine, it is possible for a total of 6 steps to execute simultaneously on the underlying host.
	The Resource Details panel displays this information.
	Note: Setting the Step Limit to "0" or leaving the field blank defaults to NO step limit.
Artifact Cache Directory	The directory on this resource's agent host from which artifact versions are retrieved and made available to job steps. Enter an absolute path to the resource containing this cache directory.
Zone	The name of the zone where this resource is or will be a member—a zone is a top-level network containing one or more mutually accessible resources.
Repository Names	A "new-line" delimited list of repository names for this resource, if needed. When a step attempts to retrieve artifact versions from an ElectricFlow repository, it queries repositories specified here. If no repositories are specified, it will "fallback" to the default repository.
	Use the drop-down menu to select your Connection Type:
	HTTP - choose this option if you are not using SSL and the agent does not need to be "trusted".
Connection Type	HTTPS - choose this option if you are using SSL, but this agent does not need to be "trusted".
	Trusted HTTPS - choose this option if you are using SSL and this agent must be "trusted". A trusted agent is "certificate verified"—the agent verifies the server or "upstream" agent's certificate.
	Agents can be either trusted or untrusted:
	 trusted - the ElectricFlow server verifies the agent's identity using SSL certificate verification.
	untrusted - the ElectricFlow server does not verify agent identity. Potentially, an untrusted agent is a security risk.

Field	Description
Enabled	Select this box to enable this resource.
	When this box is checked, the resource is enabled, which means job steps can be assigned to it. If a job step requests a pool containing this resource, the step executes on a resource in that pool.
	If disabled, and this is the only resource that satisfies a particular job step, the step's execution is delayed until the resource is re-enabled. If a resource is disabled while job steps are running on it, the running steps continue to completion but no new steps are assigned to that resource.

Click **OK** to save your settings and see your new resource listed in the table.

New Proxy Resource Panel

These are the definitions of a proxy agent and proxy target:

- Proxy *agent*—This is an agent on a supported Linux or Windows platform, used to proxy commands to an otherwise unsupported platform.
- Proxy target This is a machine on an unsupported platform that can run commands via an SSH server.

To define a new proxy resource, enter information in the following fields:

Field	Description
General	
Name	Enter a unique name for this resource. Do not use spaces in a resource name. This is the name used to select the resource for a job step; it does not need to be the same as the host name for the machine.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <dl> <i> <pr> <style> <tc> </td></tr><tr><td>Proxy Agent Host Name</td><td>Enter the domain name or IP address of the proxy agent machine corresponding to this resource.</td></tr><tr><td>Proxy Agent Port Number</td><td>Enter the port number to use when connecting to the agent for the resource.</td></tr><tr><td>Default Workspace</td><td>Enter the workspace name for leave blank to use the local, default workspace.</td></tr></tbody></table></style></pr></i></dl></div>

Field	Description
Pool(s)	Enter a space-separated list of arbitrary names, indicating the pools associated with this resource. <i>Do not use</i> "spaces" in a pool name.
	Use the drop-down menu to select the host type:
	Concurrent
	Registered
	Depending on the license installed on the ElectricFlow server, this field may appear.
Host Type	If the license on the server is a concurrent resource license, the host type defaults to Concurrent and this field does not appear.
	 If the license on the server is a registered host license, the host type defaults to Registered and this field does not appear.
	 If the license on the server is a mixed-mode license (concurrent resources and registered hosts), you must specify the host type when adding a resource.
Step Limit	Specify the maximum number of steps that can execute simultaneously on this resource—a step limit applies to a particular resource only, not to its underlying host.
	For example, if you define two resources, resource1 with a step limit of 5 and resource2 with a step limit of 1, both specifying the same host machine, it is possible for a total of 6 steps to execute simultaneously on the underlying host.
Zone	The name of the zone where this resource will be a member.
	Use the drop-down menu to select your Connection Type:
Connection Type	HTTP - choose this option if you are not using SSL and the agent does not need to be "trusted".
	HTTPS - choose this option if you are using SSL, but this agent does not need to be "trusted".
	Trusted HTTPS - choose this option if you are using SSL and this agent must be "trusted". A trusted agent is "certificate verified"—the agent verifies the server or "upstream" agent's certificate.
	Agents can be either trusted or untrusted:
	trusted –The ElectricFlow server verifies the agent's identity using SSL certificate verification.
	untrusted –The ElectricFlow server does not verify agent identity. Potentially, an untrusted agent is a security risk.

Field	Description
Enabled	Select this box to enable this resource.
	When this box is checked, the resource is enabled, which means job steps will be assigned to it. If a job step requests a pool containing this resource, the step executes on a resource in that pool.
	If disabled, and this is the only resource that satisfies a particular job step, the step's execution is delayed until the resource is re-enabled. If a resource is disabled while job steps are running on it, the running steps continue to completion but no new steps are assigned to that resource.
Proxy Target	
Proxy Target Host Name	Enter the domain name or IP address of the proxy target machine corresponding to this resource.
Proxy Target Port Number	Enter the port number to use when connecting to the proxy target host or leave blank to use the default. The default port is 22, the SSH server port.

Field	Description
	The name of the shell program used to execute commands on this [the proxy target] resource—may be overridden at the step level.
	This shell will be used to execute the step's commands on a resource. For example, using sh or cmd /c, the agent saves the command block to a temporary file with a .cmd extension and runs it: sh foo.cmd or cmd /c foo.cmd.
	If you do not specify a shell on a step, at step run-time the server looks at the resource shell. If a resource shell is not set, the shell line used by the agent is platform dependent: Windows: cmd /q /c "{0}.cmd" UNIX: sh -e "{0}.cmd"
	When you specify a shell (in the step or resource), and omit the <code>cmd-file</code> marker, the agent notices the omission and takes the correct action. For example: a user specifies $sh - x$. The agent converts this to $sh - x$ " $\{0\}$.cmd".
Proxy Target Default	Two alternate forms of shell syntax where ElectricFlow uses a "marker," {0}, as a placeholder for the command file argument:
Sneii	 <myshell> {0} <potential args="" extra="" shell=""> In this example, the command file is not meant to be the last argument in the final command line. For example, mysql -e "source {0}" This shell example runs the mysql command against this step's command containing sql.</potential></myshell>
	 <myshell>{0}<.file extension> <potential args="" extra="" shell=""> In this example, the shell requires the command file to end in an extension other than .cmd. For example, powershell "& '{0}.ps1'"</potential></myshell>
	This shell example runs Microsoft PowerShell against this step's command containing PowerShell commands.
	 When the agent parses the shell, it will parse the extension as everything after {0}. until it sees a space or non-alphanumeric character.
	 If your script uses International characters (non-ascii), add the following block to the top of your ec-perl command block: use utf8; ElectricCommander::initEncodings
Proxy Customizations	This is proxy-specific customization data. Default is "none".

Click **OK** to save your settings.

Edit Resource Panel

All of the fields on this panel are the same as those on the New Resource panel:

- All information you previously provided to create this resource is filled-in for you.
- · You can change any information already supplied.
- You can add information in any blank fields.

Click **OK** to save your modifications.

IMPORTANT: If you want to change the Connection Type, using the UI by itself is insufficient. You must also use the ecconfigure utility to change the protocol. Run: ecconfigure --agentProtocol [http|https]

Edit Proxy Resource Panel

All of the fields on this panel are the same as those on the New Proxy Resource panel:

- All information you previously provided to create this resource is filled-in for you.
- · You can change any information already supplied.
- You can add information in any blank fields.

Click **OK** to save your modifications.

Resource Details panel

This panel displays how this resource is configured with the ElectricFlow server. Some of the information supplied is linked to its source. For example, the zone name for this resource is a link to its Zone Details panel.

- The name of the resource is shown below the panel title.
- The resource description is shown under the resource name.
- Links at the top of the panel:
 - Use the View Usage link to go to the Job Step Search Results page.
 - Use the **Edit** link to modify the resource specifications.
 - Use the Access Control link [at the top of the panel] to specify or view access privileges for this
 resource.
- The tabs:
 - o General by default, the panel opens to display the "general" resource configuration.

 Properties - select this tab to add properties to this resource or see existing resource properties if already configured.

Custom Resource Properties

You can define custom properties for any ElectricFlow object. For example, if configuration information varies from resource to resource, use custom properties to store this information and then reference it from procedure steps where it is needed.

For example, if you have a pool of test machines with test hardware in a different location on each machine, you could define a property named "testLocation" on each resource, then pass this property to procedure steps using a reference such as \$[/myResource/testLocation].

Switching a Non-Trusted Agent to Trusted

If you have upgraded to ElectricCommander 4.2 or later, none of your existing agents were converted to "trusted" during the upgrade process. By default, the *local* agent is not trusted during a new or upgrade installation.

After an upgrade, if you want to use a trusted configuration, any **pre-existing** agent/host machines must be manually switched to "trusted".

Manually switching a non-trusted agent to trusted status

Perform the following steps on the agent machine to change the agent status to trusted in your environment and in the customer's site.

- Enter ectool --server <COMMANDER_SERVER_IP> login admin to log into the server and save the session ID.
- 2. Enter eccert --server <COMMANDER_SERVER_IP> initAgent --remote --force to install the agent with a self-signed certificate that needs to be overwritten.
- 3. Restart the agent.

If you do not restart the agent, it tries to use a previously cached certificate if one exists. The old certificate is invalid because a new one was just created.

4. In the ElectricFlow UI, ping the agent.

Install or Upgrade Remote Agents



Select this button () to install or upgrade host machines you want to use as agent machines for your ElectricFlow resources. A dialog box opens where you enter information about the hosts to install or upgrade.

- If no hosts are selected in the Resources Table view, you will be installing host names you enter.
- If host names are listed already in the Resources Table view, selected one or more to perform an upgrade operation.

Notes:

- The dialog screens you see for an Install operation will be somewhat different from those for an Upgrade operation, depending on the options you choose.
- You can install or upgrade Linux, Solaris, Solaris x86, HP-UX, or MacOS hosts.
- When you install the agent software on a Solaris host, install it in a directory path that has less than 70 characters. .

Prerequisites:

- You must have an artifact repository server installed.
- The target machines must be running the SSH daemon.
- The user you select must be set up for passwordless *sudo* privileges.

 For example, in the /etc/sudoers file, you must add "<username> ALL=(ALL) NOPASSWD:ALL".

IMPORTANT: If the EC-AgentManagement plugin is **not** installed, the remote install/upgrade option will not be available and you cannot install/upgrade agents remotely from the automation platform user interface. During a new ElectricFlow installation, this plugin is installed automatically.

• For each remote agent machine you want to upgrade, make sure PasswordAuthentication=yes in /etc/ssh/sshd_config.

Example:

```
# To disable tunneled clear text passwords, change to no here!
PasswordAuthentication yes
```

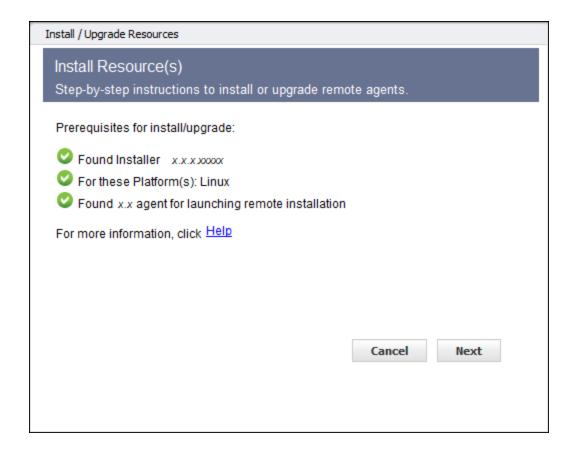
PasswordAuthentication=no is the default on SUSE platforms, so you must edit /etc/ssh/sshd_config. Restart the sshd service after changing the value.

- On Unix, SSH has the following additional prerequisite:
 - The user account on the target machine must have password-less sudo configured for running the installer with root privileges.

Using the Install or Upgrade Remote Agents Dialog Box

The ElectricFlow server runs several verification checks before you begin an Install/Upgrade operation, and when you see the following screen with all green "check mark" icons, you can proceed.

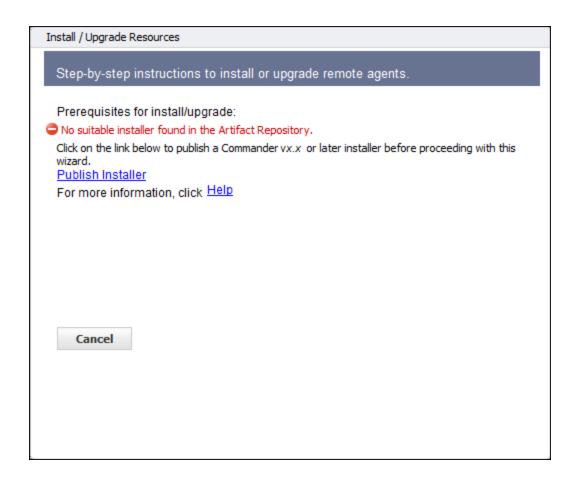
Note: Depending on your selections, you will see variations in some to the following screen examples.



Click **Next** to continue.

However, if the plugin is installed but an ElectricFlow 4.2 or later installer is not found in the artifact repository, you will see an error screen similar to the example below.

Click the **Publish Installer** link.



After clicking the **Publish Installer** link, the "Run Procedure – Publish Installer" page opens.

Enter information in the fields as follows:

Parameters

fromDir–Enter the path to the directory containing the installer on the ElectricFlow server machine. For example, /var/tmp for Linux.

platform—Use the drop-down menu to make a selection.

repository—Use default to use the repository installed during the ElectricFlow installation, or type-over default and enter a different repository name.

resource-Enter the name of resource that should be used by publishing procedure.

version–Enter the build version for your ElectricFlow install, for example, "5.0.0.56390". Your build version can be obtained from the name of the ElectricFlow installer you downloaded from the Electric Cloud FTP site.

Advanced

(Not required to complete the Publish Installer run procedure)

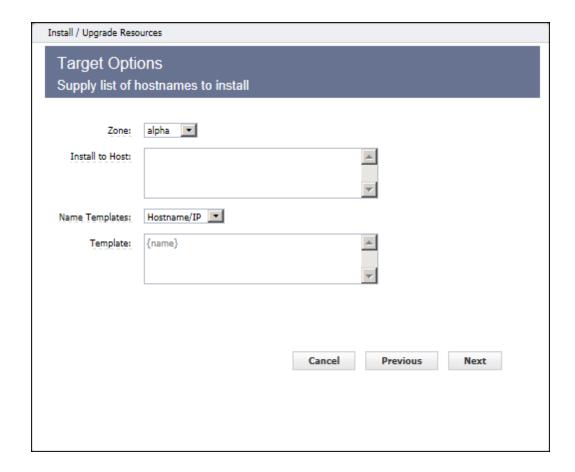
When the Publish Installer procedure runs, you can see the job status on the Job Details page.

For verification:

Go to the Artifacts page > select com.electriccloud:installer, which displays the Artifact Details page. This page displays all of your published installers.

Return to the Resources page and begin again

Assuming you now have all green "check mark" icons on the first dialog screen, and you clicked **Next**, the following dialog screen is displayed.



Enter information in the fields as follows:

• **Zone**—If you have not defined additional zones, the default zone is used and you will not see this field. If you have multiple zones defined already, use the drop-down menu to choose the zone where you are

installing or upgrading agents. You can install or upgrade agents to only one zone at a time. A functioning gateway must exist before an agent can be installed into the non-default zone. So installing the first agent into a zone must be done manually.

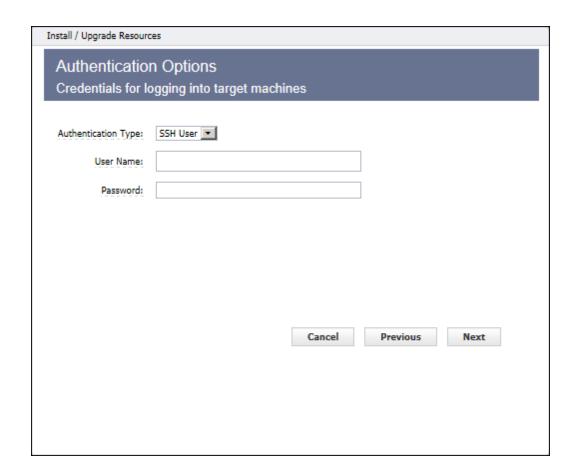
- Install to Host

 (Required) Enter one or more host names or IP addresses in a space-separated list.

 Individual host names cannot contain spaces. If this is an upgrade operation, host names or IP

 addresses are pre-populated in this field from the resources you selected for upgrade on the Resources page Table view, first column.
- Name Templates—Use the drop-down menu to choose a simple template for setting the names of newly
 created resources.
- **Template**—You can edit this field only if you selected Custom in the Name Template drop-down menu. This field undergoes property expansion in a global context and is scanned for special tokens {name} and {counter}.
- **Driving Resource**—This resource performs all actions for installing or upgrading a target host on behalf of the server.

Click **Next** to continue to the next dialog screen.



Enter information in the fields as follows:

Authentication Type-Use the drop-down menu to choose SSH User, SSH Key, or WMI:

• If SSH User:

- User Name-(Required) Enter a user name.
- Password–(Required) Enter a password for the user name.

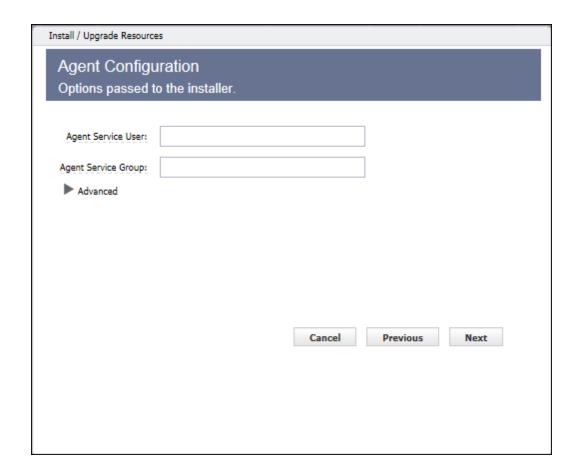
• If SSH Key:

- User Name-(Required) Enter a user name.
- Public Key Path-Required) The path to the SSH public key file.
- o Private Key Path-Required) The path to the SSH private key file.
- Passphrase–(Optional) The passphrase for unlocking the private key file.

• If WMI:

o No additional information is needed

Click **Next** to continue to the next dialog screen.



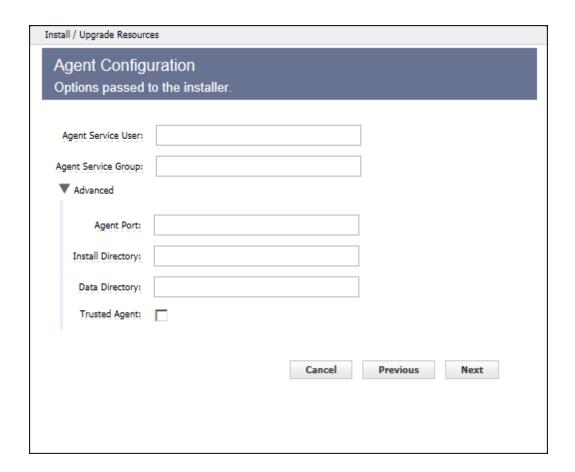
Note: You will not see this screen if you are upgrading hosts.

For installing hosts, enter information in the fields as follows:

- Agent Service User-(Required) the user the ElectricFlow agent "runs as".
- Agent Service Group-(Required) the group the ElectricFlow agent "runs as".

Selecting "Advanced" is optional—Select this option to define non-default values.

- If you did not select Advanced, click Next to continue (and skip the next dialog screen example).
- If you selected **Advanced**, see the next dialog screen example.



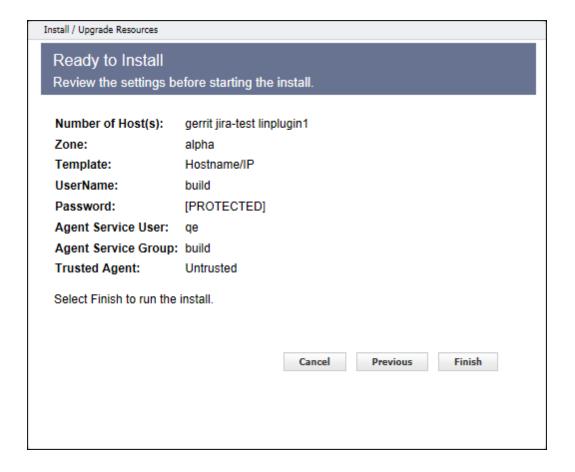
(Optional) Advanced options:

- Agent Port-Specify the listener port for the agent.
- Install Directory–Specify the path to the install directory.
- **Data Directory**—Specify the path to the directory where your modified ElectricFlow files are stored (configuration and log files).
- **Trusted Agent**–Select this check box and all agent machines being installed or upgraded will be "trusted". For more information, see Switching a Non-Trusted Agent to Trusted on page 613.

Click **Next** to continue to the summary dialog.

Your summary dialog will be similar to the following example, but it will contain the values you specified.

Note: If you are installing or upgrading more than 5 hosts, you will see a quantity number and NOT a space-separated. list of host names.



Click Finish to start the install or upgrade, and the Job Details page is displayed.

When Install/Upgrade is complete, you can return to the Resources page to see your newly installed or upgraded hosts listed in the Table or Grid view.

To verify a resource version, click a resource name (in the Resources page Table view) to open the Resource Details panel for that resource.

Resource Pools

This page displays all resource pools currently available to ElectricFlow.

- To create a new resource pool, click the Create Resource Pool link.
- To find an existing resource pool, use the **New Search** link.

Column descriptions

Column Name	Description
Pool Name	This is the name you supplied for this resource pool. Click on a Pool Name to go to the Resource Pool Details page.
Enabled	If this box is "checked", this resource pool is enabled.
Auto Delete	If this column is selected with a check mark, the pool will be deleted when the last resource is deleted or removed.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pr> <pr> <pr> <pr> <pr></pr></pr></pr></pr></pr>
Actions	Use the Copy link to make an exact duplicate of an existing resource pool, then select the copy to go to the Resource Pool Details page to change the resource pool name or add or remove resources from the pool.
	Use the Delete link to delete the resource pool on the same row. Note: Deleting a resource pool does not delete the resources referenced by the pool.

Resource Pool - Create a New Pool or Edit Existing Pool

Creating a new resource pool

Enter information in the following fields:

Field Name	Description
Pool Name	Enter a name for the resource pool. This name must be unique within a list of multiple resource pool names and cannot be the same as any resource name.
Enabled	Select this check box to enable the resource pool.
Туре	 A resource can be <i>static</i> or <i>dynamic</i>. A static resource is part of your system or network (not in the cloud), such as a server, database, or agent machine. A dynamic resource is a cloud resource that can be provisioned and later spun up on-demand when an application is deployed.

Field Name	Description
# Resources	Number of resources in the resource pool.
Resource Template Name	The name of the resource template to which the resource is applied. Note: Only dynamic resources can be applied to a resource template.
Auto Delete	If you select this check box, the pool is deleted when the last resource is deleted or removed. • The autoDelete property flag is automatically set to "true" if the pool was created on the Create Resource page, as a side-effect of the resource API calls. • By default, resource pools created explicitly on the Create Resource Pool page, have the autoDelete flag set to false—the checkbox is "unchecked".
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with tags">https://ntml> tags tags">https://ntml> tags allowed in the text are: <a> tags tags tags allowed in the text are: <a>
 tags tags tag
Actions	Actions you can perform on the resource pool, including: Copy—Creates a new resource by copying the this resource pool. Delete—Deletes this resource pool. Tear Down—Removes (decommissions) the dynamic resource pool from a dynamic environment. Track Changes—Opens the Change History for the resource pool.

Click **OK** to save the information to create a new resource pool. You will see your new resource pool listed on the Resource Pools page.

Note: Resource pools cannot be nested, which means if the name of a resource pool is added to a resource pool, it will be ignored.

A resource pool can override the default algorithm by providing a JavaScript fragment in the Ordering

 $\textbf{Filter} \ (\texttt{orderingFilter}) \ \textbf{intrinsic} \ \textbf{property}. \ \textbf{The JavaScript} \ \textbf{environment} \ \textbf{for this fragment} \ \textbf{is} :$

- A variable named resourceList an array of the resources in the pool, sorted as described above
- A variable named resourcePool the pool object

• The job step being scheduled is configured as the context object for property resolution

The evaluation result is expected to be an array of resources, in the order they should be considered by the scheduler.

Before evaluation, the JavaScript fragment is wrapped as a function, to make it possible to use the return statement in the script.

It is possible to return any resources in the result, not just those grouped in a pool, so this feature enables true dynamic resource selection. Resources are still subject to the "usable" criteria. See the "Unusable resources" section at the end of this Help topic.

Examples for the Ordering Filter field:

o Sort by resource name in descending order:

```
return resourceList.sort(function(a,b) {
return b.resourceName.localeCompare(a.resourceName);
});
```

o Ignore the default resources and return different resources:

```
var result = new Array();
result[0] = resources["resource1"];
result[1] = resources["resource2"];
return result;
```

Editing an existing resource pool

Use this page to modify any specifications for the resource pool.

Click **OK** to save the information and update the resource pool.

Unusable resources

Resources, whether in a pool or not, are not considered usable if any of the following are true:

- The resource does not exist.
- The resource is disabled.
- The resource does not have a host name.
- The agent where the resource is associated is not reachable.
- The number of steps running on the resource is equal to the resource step limit.
- The resource is exclusive to a different job using the job exclusiveMode.
- The resource is exclusive to a step in a different job using the step or call exclusive Mode.
- The resource is exclusive to a different step in the same job using the call <code>exclusiveMode</code>. and the other step is not an ancestor of the step.
- The resource is exclusive to a different step in the same job using the step exclusiveMode.
- The job step does not have execute privileges on the resource.
- The resource assignment exceeds the license limit (if any).

Resource Pool Details

This page displays information for the resource pool you selected, including a summary of the specifications supplied to create the resource pool, current resources assigned to this pool, and the ability to add or remove resources.

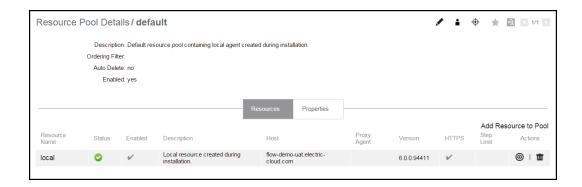
Links and actions at the top of the page

- Edit use this link to go to the Edit Resource Pool page.
- Access Control use this link to go the Access Control page for this resource pool.
- The "star" icon allows you to save this information to your Home page.
- The "curved arrow" icon returns you to the Resource Pools page.
- Pagination Use the "previous" and "next" arrow icons to view the previous or next project. The numbers between the arrow icons display the number of projects you can view and the first number indicates which project [in the list] you are viewing.

The "tabbed" sections

The tabs at the top of the table allow you to select the type of information you want to see. The Resource Pool Details page opens with the Resources tab highlighted, so the first table you see is the Resources table. See the following screen example.

This table displays the resources currently included in this pool. The name of the pool ("custom" in our example) follows the Resource Pool Details page title.



The summary section at the top of the table contains information previously defined when the pool was created:

Description - The text previously supplied for this object when it was created.

Ordering Filter - The Javascript ordering filter (see the New Resource Pool Help topic for more information) or "empty" if no filter was applied.

Auto Delete - When "yes" is specified, the pool will be deleted when the last remaining resource is removed or deleted from the pool.

Enabled - "yes" specifies this pool is enabled for use.

Links at the top of the table

Add Resources to Pool - Click this link to see a pop-up dialog to enter a resource name to add to this pool. You can type a name or select from a list of existing resource names. If other resources are already configured, you will see a list from which to select a resource. Click **OK** to add new resources to the pool.

Column descriptions

Column Name	Description / Actions
Resource Name	The name of the resource. Click on a resource name to go to the Edit Resource page.
Status	Indicates the current status of the resource.
Enabled	Indicates whether the resource is enabled. If a resource is disabled, no job steps will be assigned to it. ElectricFlow will use other resources within the pool to satisfy requests for the pool.
Description	The text previously supplied for this object when it was created.
Host	Steps assigned to run on this resource will use this host.
Proxy Agent	The name or IP address of the proxy agent machine.
Version	The version of the ElectricFlow agent installed on this resource.
HTTPS	A "checkmark" indicates HTTPS was selected when this resource was created.
Step Limit	The maximum number of steps that can execute simultaneously on this resource.
Actions	Ping - Use this link to check the status of the resource. Remove - Use this link to remove a resource from this pool.

Properties tab

This tab provides a table listing all properties for the resource pool and includes the following functionality:

Links at the top of the table

- To create a new property for this resource pool, click the Create Property or Create Nested Sheet link.
- To view or change privileges on the property sheet, click the Access Control link.

Column descriptions

Column Name	Description / Action
Property Name	The name of the property. Click on a property name to edit that property.
Value	Indicates the value assigned to this property
Description	A text description previously supplied for your reference only.
Actions	Delete - Use this link to delete this property.

Zones

This page displays zones currently available to ElectricFlow and provides other zone operations—create,edit, delete, and so on.

- A zone is a way to partition a collection of agents to secure them from use by other groups.
 For example, you might choose to create a developers zone, a production zone, and a test zone—agents in one zone cannot directly communicate with agents in another zone.
- A default zone is created during the ElectricFlow installation.
 The server implicitly belongs to the default zone, which means all agents in this zone can communicate with the server directly (without the use of a gateway).
- Every agent, and all resources defined on that agent, can belong to one zone only.
- Within a zone, agents can be either trusted or untrusted.
 - trusted the ElectricFlow server verifies the agent's identity using SSL certificate verification.
 - untrusted the ElectricFlow server does not verify agent identity. Potentially, an untrusted agent is a security risk.

Links and actions at the top of the table

- To create a new zone, click the Create Zone link.
- The "star" icon allows you to save this zone information to your Home page.

Column descriptions

Column Name	Description / Actions
Name	The name supplied when this zone was created. Click on a Zone Name to see the Zone Details panel for that zone.
Descriptions	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <div> <dl> <i> <pre> <pre> <style> <tc> </td></tr><tr><td>Actions</td><td>Delete - Use this link to delete the zone on the same row. Note: Deleting a zone does not delete resources in the zone. All agents (and their respective resources) move to the default zone if their current zone is deleted. The default zone cannot be deleted.</td></tr></tbody></table></style></pre></pre></i></dl></div></div>

Zone Details panel

This panel displays the zone name and description for the zone you selected.

Links and actions at the top of the table

- Resources-Use this link to go to the Resources page.
- Edit-Click this link to display the Edit Zone panel.
 - On the Edit Zone panel, you can change the zone name ,or you can add or change the zone Description.
 - Click **OK**after making any changes.
- Access Control Use this link to go to the Access Control page for this zone.
 - o Current access privileges are displayed, if any.
 - o Add the access control privileges you need for this zone.

Select the Properties tab to view any properties created for this zone, or use the "create" property links to create properties for this zone.

Creating a new zone

Enter information in the fields as follows:

Field Name	Description
Name	Enter a name of your choice for this zone. The name must be unique among other zone names.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <i> <pre> <pre> <style> <tc></td></tr></tbody></table></style></pre></pre></i>

Click **OK** to see your new zone displayed in the table.

To add resources to this zone, go to the Resources page.

Access Control notes

Resource and Resource Pool continue to inherit from Resources privileges. To create a resource, you must have modify privileges on Resources (this has not changed) and you must have modify privileges on the zone. In addition, to move a resource from one zone to another, you must have modify privileges on both zones and the resource you want to move.

Gateways

To communicate with a resource, workspace, or artifact repository server in another zone, a "gateway" must be created. A gateway object contains two resource (agent) machines, for example, GatewayResource1 and GatewayResource2—each configured to communicate with the other. One gateway resource resides in the *source* zone and the other in the *target* zone. A gateway is bidirectional and informs the ElectricFlow server that each gateway machine is configured to communicate with its other gateway machine (in another zone).

If your company requires the added security of a firewall between zones, gateway agents can be configured to communicate with/through the firewall.

This page displays all gateways currently defined in ElectricFlow and provides other gateway operations—create, edit, delete, and so on.

- A firewall between zones A gateway resource can be configured to communicate with an intermediary firewall in its path as a proxy to communicate with its peer on the other side of the gateway.
- Each gateway records the host/port combination each gateway agent/resource must use to communication with its peer on the other side of the gateway.

- Multiple gateways can be defined for a zone if required.
 For example, you may have multiple resources in zoneA that need to communicate with each other, but some of those resources also need to communicate with zoneB, while others need to communicate with zoneC only. In this scenario, zoneA would require two gateways—one to zoneB and one to zoneC.
- One resource can participate in multiple gateways.
 For example, assume we have 3 zones, zone1, zone2, and zone3, each created to contain agent/resource machines for a different, specific purpose (production, testing), but we want to share or pass data from a resource in zone1 to another resource in zone2 or zone3:
 - We need two gateways:
 - Gateway1 connects ResourceA in zone1 to ResourceC in zone3
 - Gateway2 connects ResourceA in zone1 to ResourceB in zone2
 - With this gateway-resource configuration, ResourceA can communicate directly with zone2 or zone3.

Links and actions at the top of the table

- To create a new gateway, click the Create Gateway link.
- The "star" icon allows you to save this gateway information to your Home page.

Column descriptions

Column Name	Description / Action
Name	The gateway name specified when this gateway was created. Click on a gateway name to see the Gateway Details panel for that gateway.
Enabled	A "checkmark" in the box indicates this gateway is enabled.
Resource 1	The first of two resources required to create a gateway.
Host 1	The name Resource 2 uses to communicate with Resource 1. If "blank", the Agent Host Name attribute in Resource 1's definition is used at runtime.
Resource 2	The second of two resources required to create a gateway.
Host 2	The name Resource 1 uses to communicate with Resource 2. If "blank", the Agent Host Name attribute in Resource 2's definition is used at runtime.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </td></tr><tr><td>Actions</td><td>Delete - Use this link to delete the gateway on the same row.</td></tr></tbody></table></style></pre>

Gateway Details panel

This panel displays properties and access control privileges assigned to this gateway.

Select the Properties tab to see any existing properties or to create properties for this gateway.

Links and actions at the top of the panel

- Edit Click this link to display the Edit Gateway panel.

 On the Edit Gateway panel, you can change the gateway name or add/change the gateway Description.
- Access Control Click this link to go to the Access Control page to set access privileges for this gateway.

Create Gateway panel

Enter the following information:

Field Name	Description / Action
Name	Enter a name of your choice for this gateway. The name must be unique among other gateway names.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <o>> <pre> <style> <tc> <u>> <u>> <u>> <u>> <u>> <u>> <u>> <u</td></tr><tr><td>Enabled</td><td>Select this box to enable the gateway.</td></tr><tr><td>Resource 1</td><td>The name of your choice for the first of two required gateway resources. Do not include "spaces" in a resource name.</td></tr><tr><td>Host 1</td><td>The agent host name where Resource1 resides. This host name is used by Resource2 to communicate with Resource1. Leave this field blank to use the host name from Resource1's definition.</td></tr><tr><td>Port 1</td><td>The port number used by Resource1 - default is to the port number used by the resource.</td></tr><tr><td>Resource 2</td><td>The name of your choice for the second of two required gateway resources. Do not include "spaces" in a resource name.</td></tr></tbody></table></style></pre></o>

Field Name	Description / Action
Host 2	The agent host name where Resource2 resides. This host name is used by Resource1 to communicate with Resource2. Leave this field blank to use the host name from Resource2's definition.
Port 2	The port number used by Resource2 - default is to the port number used by the resource.

Click **OK** to see your new gateway displayed in the table.

Edit Gateway panel

This panel is populated with previously supplied information to define the gateway.

You can change any existing specifications or add new information.

Click **OK** to save your changes.

Access Control note

A gateway inherits privileges from the ZonesAndGateways ACL. See the Access Control ("Server Objects" section) Help topic for more information.

Accessing the Resource Templates in the Automation Platform

A resource template has the required information to provision and spin up cloud resources on an on-demand basis to model a dynamic environment. You set the cloud provider and configuration management details in the resource template. Then you define the environment tiers and add resource templates to these tiers in an environment template to complete the dynamic environment configuration.

You can generate new resource pools in ElectricFlow by accessing the resource templates from the automation platform UI.

1. Go to Cloud > Pools.

Example:



2. Click Resource Templates.

The Resource Templates list opens.

Example:





(Optional) To view details about a resource template, choose a template and click the Menu button (for it.

The Context menu opens.

- 4. Click one of these options:
 - Details-The Edit Resource Template Details dialog box opens.

The tabs that appear in the dialog box depend on the cloud provider specified in the resource template.

• Properties—The Properties dialog box opens.

You can view and edit the properties that apply to the resource template.

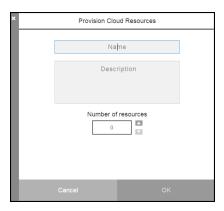
- Access Control-The Access Control page opens.
- Track Changes—The Change History for the resource template opens.



5. To provision the resource template, choose a resource template, click the **Menu** button (), and select **Details > Provision**.

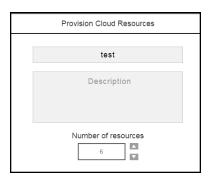
The Provision Cloud Resources dialog box opens.

Example:



6. Enter the resource pool name, an optional description, and the number of resources in the pool, and then click **OK**.

Example:



You go to the Job Details page in the automation platform.

Define Search

The page allows you to search for any supported object type.

Enter information in the fields as follows:

Field Name	Description
Number of Results	Default is 200. Change this number to see a greater or lesser search result set.

Field Name	Description
Results Per Page	Choose the number of search results you would like to see. Default is 20.
Object Type	Use the drop-down menu to select the object you wish to search.
Intrinsic Filters	Click the Add Intrinsic Filter link one or more times as needed to further define the sort criteria for your search. • Use the drop-down arrow to select an intrinsic filter argument. These arguments are properties. For a description of any of these properties, see the Properties Help topic. • Use the drop-down arrow to select an operator for the search. • To narrow your search, enter an actual object name. For example, if you selected Project Name to begin your search, you might want to enter the actual project name in this field or the first few letters if you have multiple projects named in a similar pattern.
Custom Filters	Click the Add Custom Filter link one or more times as needed to enter custom sort criteria to further define your object search. Type-in the object name to define your search criteria, then select your search operator.

Click **OK** after filling-in the fields and to go to the Search Results page.

Notes:

- 1. Only the first 450 characters of a property string are searchable. For example, if your property value happens to contain 1500 characters and the information you are trying to find is 800 characters into the property string, the search mechanism will not find it.
- 2. When defining a search on this page, property path references are *not* allowed in search parameters. However, property path references are allowed in the Home page Quick View filters.

Search Results

This page displays the results of the search you defined.

- If the results were not as expected, click the Edit Search link to modify your search criteria.
- Click the New Search link to define a new object you need to find.
- Click the Save Filter link if you want to save this filter for future use in a report. In the pop-up screen, enter a name for the filter and click OK. The Saved Filter will be stored as a property in your chosen project.
- To continually monitor the status of your found object, click the Refresh Search link.

Notes:

If you need to refer to this page on a regular basis, mouse-over the "star" icon to add this page to the Shortcut section on your Home page.

If you use RSS, an active RSS icon is provided in Windows Explorer on this page for your convenience. If you use Firefox, click **Bookmarks** > Subscribe to this page to display the feed. You can then add the feed URL to a viewer of your choice.

Server

This page displays overall information about the ElectricFlow server.

These links are at the top of the table:

- Click the **Settings** link to go to the Edit Server Settings page.
 Use this page to edit properties for the ElectricFlow server.
- Click the Access Control link to go to the Access Control page for the server.
 See the Access Control Help topic for more information.

The tables have this Information:

- The System Access Control table displays top-level properties.
 Click on an item in the Category column to see, add, or modify privileges for that object.
- In the Custom Server Properties table, click on a Property Name to view the property content.

 If using *ectool*, these properties can be accessed using the property name starting with "/server".

See the Properties Help topic for more information.

These links are provided within the table:

- Create Property Use this link to create a new server property.
- Create Nested Property Use this link to create a new nested server property.
- Access Control This access control link allows you to create or edit privileges for the property sheet.
- The Actions column allows you to Edit or Delete the property in that row.

Settings - edit existing property settings

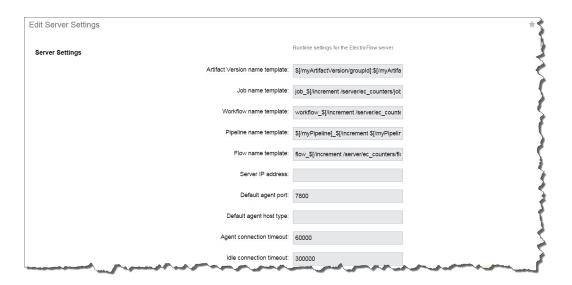
Use this page to edit property settings for the ElectricFlow server or any integrations you are using with ElectricFlow.

For example, if you access this page from the Server page, the page title will be Edit Server Settings and you will see fields for editing ElectricFlow server properties. From Perforce, the page will be titled Edit Perforce Settings, and the fields you can edit will be specific to Perforce functionality.

The following screen example is from the Edit Server Settings page.

Note: Depending on the ElectricFlow version you are using, the Edit Server Settings page may contain more or fewer fields than shown below.

For more information about any setting, hover your mouse over the setting name and click to see an inline Help description about the field value.



You may change any existing information or enter new information to empty fields. For Help information, hover your mouse over a field to see a Tool Tip.

Click **OK** when your changes are complete.

Also notice the "star" icon in the upper right corner. Click this icon to add this page to your Home page for quick access if you need to change these settings frequently.

Source Control Configurations

This web page displays all source control configurations you have created to communicate with ElectricFlow.

Link at the top of the table

• Use the **Create Configuration** link to create a new or additional source control configuration.

Table column definitions

Configuration Name - This is the name you chose for this source control configuration.

Description - This is the plain text or HTML description you provided for this source control configuration.

Plugin - This is the ElectricFlow plugin associated with this source control configuration.

Actions -

Edit - Use this link to modify information previously supplied to configure this source control system.

Delete - Use this link to remove this source control configuration.

Source Control Configurations - create new or edit existing configuration

Use this page to define a new source control configuration (SCM) or modify an existing source control configuration. A configuration is a collection of properties that define how to communicate with a particular source control system.

ElectricFlow bundles and supports numerous source control types.

- After creating a source control configuration, your entry will appear in the table on the Source Control Configurations web page— to see this web page, select the Administration > Source Control tabs.
- Other SCM systems are bundled with ElectricFlow and available in the Plugin Catalog. Each integration contains options specific to that SCM.
- The following SCM integrations are provided as examples.

AccuRev

ClearCase

File

Git

Perforce

Property

Subversion

Note: If you need a different SCM, not specified in this Help topic, go to the Plugin Manager page to locate your SCM (Administration > Plugins). See the Help topic associated with your plugin - located on the Plugins page. At a minimum, you need to enter a name for your configuration and you may want a minimal text description.

AccuRev

Enter information in the fields as follows:

Configuration Name - Type a name for the configuration— any name you choose, but the name must be unique. For example, you might use "myAccuRevServer."

Description - You can change or modify the default text description - ElectricFlow does not use this information.

Login As

User Name - The name ElectricFlow needs to use to communicate with your AccuRev system. For example, you may be using a special "read-only" user name similar to "Build" for your user name.

Password - The password for the specified User Name.

Retype Password - Type the password again.

Click **Submit** to save your information.

ClearCase

Enter information in the fields as follows:

Configuration Name - Type a name for the configuration— any name you choose, but the name must be unique. For example, you might use "myClearCaseServer."

Description - You can change or modify the default text description - ElectricFlow does not use this information.

Click **Submit** to save your information.

File

You may want to use this configuration with ElectricSentry to "watch" for changes in any files.

Enter information in the fields as follows:

Configuration Name - Type a name for the configuration—any name you choose, but the name must be unique.

Description - You can change or modify the default text description - ElectricFlow does not use this information.

Click **Submit** to save your information.

Git

Enter information in the fields as follows:

Configuration Name - Type a name for the configuration— any name you choose, but the name must be unique.

Description - You can change or modify the default text description - ElectricFlow does not use this information.

Click Submit to save your information.

Perforce

Enter information in the fields as follows:

Configuration Name - Type a name for the configuration— any name you choose, but the name must be unique. For example, you might use "myPerforceServer."

Description - You can change or modify the default text description - ElectricFlow does not use this information.

Login As

User Name (P4USER) - The name ElectricFlow needs to use to communicate with your Perforce system. For example, you may be using a special "read-only" user name similar to "Build" for your user name.

Password (P4PASSWORD) - The password for the specified User Name.

Retype Password - Type the password again.

P4PORT - Sets the hostname and port number ElectricSentry uses to contact the Perforce server (hostname: 1234).

P4HOST - The name of the host machine to impersonate.

P4TICKETS - The full path name of the ticket file used by the "p4 login." This value is set as an environment variable, not as a command-line option (fullPathAndFileName)

P4CHARSET - Sets the character set used for translation of Unicode files (characterSet)

P4COMMANDCHARSET - Used to support UTF-16 and UTF-32 character sets (commandCharSet).

Click **Submit** to save your information.

Property

You may want to use this configuration with ElectricSentry to "watch" for any property changes.

Enter information in the fields as follows:

Configuration Name - Type a name for the configuration—any name you choose, but the name must be unique.

Description - You can change or modify the default text description - ElectricFlow does not use this information.

Click **Submit** to save your information.

Subversion

Enter information in the fields as follows:

Configuration Name - Type a name for the configuration— any name you choose, but the name must be unique. For example, you might use "mySubversionServer."

Description - You can change or modify the default text description - ElectricFlow does not use this information.

Login As

User Name - This is the name ElectricFlow needs to use to communicate with your Subversion system. For example, you may be using a special "read-only" user name similar to "Build" for your user name.

Password - This is the password for the specified User Name.

Retype Password - Type the password again.

Click **Submit** to save your information.

Checking out code from a source control system

Go to Projects > select a project > select a procedure.

To create a New Step for source code management, select the Plugin link.

- In the Choose Step panel, select Source Code Management from the left pane, then select the SCM you configured.
- The right-pane now shows the types of steps available for your configuration. Select the step you need and automatically go to the New Step page.
- On the New Step page, notice the Subprocedure section now contains the SCM integration you configured and the step you chose.
- Enter the remaining information to create your SCM step.

To edit an existing source control configuration

Modify any of your source control configuration information by "typing-over" any previously entered information.

Click Submit to save your modified source control information.

Active Users

This page displays all users known to this ElectricFlow server, including users defined locally within the server and those defined in external repositories such as LDAP or ActiveDirectory.

Active Users are those persons who have logged into ElectricFlow, although if a another user views a selection of user accounts, those users [viewed] will appear on the Active Users page also.

Click the **Show Inactive Users** link to go to the All Users web page—this page lists all users: local, LDAP or Active Directory, active, or inactive. If you delete an LDAP or Active Directory user from this [Active Users] page, the user will not be deleted from the All Users page.

- To edit *local* user content, click on the user name you want to modify.
- To create a new local user, click the **Create Local User** link at the top, right-side of the table.
- To delete a local user, click the **Delete** link in the Action column for that user.
- Click the "star" icon at the top, right-side of the table if you want to add this page to your Shortcut section on the Home page.

While you can view **external** user content, you **must** go to the LDAP or ActiveDirectory repository to edit external user information. However, you may associate properties with external users and then use those properties in ElectricFlow.

To configure your existing LDAP and Active Directory account repositories to communicate with ElectricFlow, click the Administration > Directory Providers tabs.

Users and Groups

Use this page to view a filtered list of ElectricFlow local users or groups.

Enter information in the fields as follows:

Field Name	Description
Filter	Enter a partial name to retrieve the list of all users or groups that match that name.
Maximum Results	Enter the number of results you would like to see returned. Default is 100. Caution: Currently, ElectricFlow has no upward limit, but if you enter a number of 1000 or higher, you may overload your browser, which will result in performance degradation or worse.
Include Inactive [User/Group]	Use this check box only if requesting external LDAP or Active Directory providers. • If "checked," all available users or groups are returned. • If not checked, only those users or groups known to the ElectricFlow server are returned, for example, users who have logged in or groups containing users who have logged in.

Click **OK** to begin the search.

Links at the top of the page

Create Local [User/Group] - Click this link to create a new user or group - for users or groups *local* to the ElectricFlow server.

Column descriptions

Column Name	Description
Name	The name of the user or group. The name combined with the location is the unique identifier for this user or group, <i>local</i> to the ElectricFlow server. Click the name to see more information.
Repository	The repository where the user or group is located. "Local" means local to the ElectricFlow server. Other names in this column refer to defined directory providers.
Location	Only viewable for users or groups from non-local directory providers.
Real Name (users only)	The full name of the user.
Email (users only)	The email address for the user.

Column Name	Description
Actions	Edit - Only available for local users and groups. Click to edit the user or group definition.
	Delete - Click to delete the user from the ElectricFlow server. Note: This delete function does not delete non-local users from an LDAP or Active Directory repository.

User - create new or edit existing local user

Use this page to create or modify a local user only.

ElectricFlow supports two kinds of user accounts: those defined externally in an LDAP or Active Directory system, and *local* users defined in this ElectricFlow server.

Local users are not visible outside ElectricFlow.

We recommend using external accounts whenever available, but you may need to create local users if you do not have a shared directory service or if you need special accounts to use for ElectricFlow only.

To create a new local user

Enter information in the fields as follows:

Field Name	Description
User Name	Enter the name of the user account to be used for login
Real Name	Enter the user's real (given) name.
Email	Enter the user's email address.
Password	Enter the user's login password.

Click **OK** to save the new local user data.

To edit an existing local user

- You may highlight and type-over any previously entered user information.
- If you leave the password fields blank and click **OK**, the existing password remains unchanged.
- To change the user's password, enter the new password for the user, AND also enter the current admin password.

Click **OK** to save the modified user data.

Use this page to modify user properties or assign Custom User Properties. Select the **Create Property**, **Create Nested Sheet**, or **Access Control** links to set the properties you need.

Note: To configure your existing LDAP and Active Directory account repositories to communicate with ElectricFlow click the **Administration** > **Directory Providers** tabs.

User Details

This page displays external user information retrieved from a repository such as LDAP or Active Directory.

To edit external user information, you *must* connect directly to the repository to modify external user content. However, you can associate properties with an external user and then use those properties in ElectricFlow.

Use this page to modify user properties or assign Custom User Properties. Select the **Create Property**, **Create Nested Sheet**, or **Access Control** links to set the properties you need.

Definitions for summary information at the top of page

User Name: The name this user is known by in the system.

Repository: Usually LDAP or Active Directory.

Real Name: The real, given name for user.

Email: The user's email address as known within the system.

Groups: This field displays all groups where this user is a member. Clicking on a group name displays the Group Details page, which lists all members who belong to that group.

Note: To configure your existing LDAP and Active Directory account repositories to communicate with ElectricFlow, click the Administration > Directory Providers tabs.

Edit User Settings

Use this page to modify your User Settings for the product tab view. Depending on which view you choose, you can have limited access to ElectricFlow features and functions.

In the Tab View, use the pull-down menu to make a new selection:

Default - Selecting this view provides all regular/standard product tab views, including plugins, which by default are installed into this view.

Base - This view removes the ability to use a plugin. You can see the plugin list (accessed by selecting the Plugins tab), but you are denied access to configure, install, uninstall, or promote a plugin.

Inherit - Frequently this view is set for a group. The UI searches for the property, ec_ui/defaultView, first on the user, second in any groups the user belongs to, and last, on the server. The inherit setting clears the user's private defaultView setting and then picks up the first value from the rest of the search path.

Click **OK** after selecting a different tab view.

Groups

This page displays all groups known to this ElectricFlow server, including groups defined locally within the server and those groups defined in external repositories such as LDAP.

- To edit a *local* group, click on its name.
- To create a new local group, click the Create Local Group link.
- Click on an external group to view its contents, but you cannot edit the content through ElectricFlow.
 Instead, you must go to the group's repository directly. You can, however, associate properties with external groups, which can then be used in ElectricFlow.

To configure your existing LDAP and Active Directory account repositories to communicate with ElectricFlow, click the Administration>Directory Providers tabs.

Group - create new or edit existing local group

To create a new local group

Enter information in the fields as follows:

Name - Type a name for the group. Choose any name, but it must be unique among other local group names.

Users - Type a list of users who need to belong to this group— type one user per line.

Click **OK** after populating the group.

To edit a local group

- You can highlight the existing group name and type a new name—or
- You can scroll to the bottom of the list of users and type-in additional members for this group, one name per line—or
- You can highlight a user name and press the delete key (on your keyboard) to delete that user from the group.
- Click **OK** after completing your changes.

Also, you can associate custom properties with a group. Select either the **Create Property, Create Nested Sheet**, or the **Access Control** links to set up the properties you need.

Group Details

This page displays external group information retrieved from a repository such as LDAP or ActiveDirectory.

To edit external group information, you must connect directly to the repository to modify an external group. However, you can associate properties with an external group and then use those properties in ElectricFlow.

Click the **Access Control** link [at the top of the table] to see the privileges assigned to this group, inherited privileges, add a user or group, and so on.

Also, you can associate custom properties with a group. Select either the **Create Property**, **Create Nested Sheet**, or the **Access Control** links to set up the properties you need.

Note: To configure your existing LDAP and Active Directory account repositories to communicate with ElectricFlow, click the Administration > Directory Providers tabs.

Directory Providers

ElectricFlow uses account information from multiple sources. In most cases, the primary account information source is an external LDAP or Active Directory repository: both user and group information is retrieved from the repository.

Local users and groups are defined within ElectricFlow.

Links at the top of the table

Click the **Add Active Directory Provider** or the **Add LDAP Provider** link to go to a new web page to configure your Directory Provider.

After configuring your directory provider, that name will appear in the All Providers table, the Provider Name column.

Column descriptions

Column Name	Description / Actions
Provider Name	Click on a Provider Name to make changes to an existing directory provider.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with httml /ahtml> tags. The only HTML tags allowed in the text are: <a> <a> <a> <a>
Actions	Copy - Use this link to make a copy of the Provider Name configured on this row. Remove - Use this link to remove the Provider Name on this row.

Detailed Help for defining a directory provider in ElectricFlow is available from the New Active Directory Provider or the New LDAP Provider web page.

Click the Help link from either web page. The Help text also contains examples you may find useful.

Directory Providers - create new or edit existing directory providers

ElectricFlow supports Active Directory and LDAP directory providers

This Help topic contains instructions for creating an Active Directory or LDAP directory provider and editing those configurations if you need to make changes at a later time. Use the following links to go the section you need:

To create a new Active Directory provider

To create a new LDAP directory provider

Examples for directory provider field descriptions

To edit an existing directory provider

To create a new Active Directory provider

Enter information in the fields as follows to specify your existing Active Directory users and groups to communicate with ElectricFlow. For examples of information you enter in the these fields, see the table after the following description sections.

Field Name	Description	
General section		
Provider Name	This name identifies users and groups that come from this provider.	
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </td></tr><tr><td>URL Discovery</td><td colspan=2>Select the method to retrieve the URL for the Active Directory server. Auto- discovery using DNS automatically discovers Active Directory servers on the given domain. Alternatively, you can specify a custom URL to an Active Directory server.</td></tr><tr><td>Domain Name</td><td colspan=2>The domain where Active Directory servers are automatically discovered. For example, my-company.com</td></tr></tbody></table></style></pre>	

Field Name	Description	
Use SSL	Select this box if you want to use SSL when the ElectricFlow server contacts your Active Directory server. Note: Transport Layer Security (TLS) has replaced Secure Sockets Layer version 3.0 (SSLv3) on the ElectricFlow web server and the ElectricFlow server.	
URL	This is an explicit URL to the Active Directory server. The URL is in the form protocol://host:port/basedn. Protocol is either ldap or ldaps (for secure LDAP). The port is implied by the protocol, but you can override the port if you cannot use the default location (default is port 389 for ldap or 636 for ldaps). The basedn is the path to the top-level directory that contains users and groups at this site. Typically, this is the domain name where each part is listed with a dc= and separated by commas. Note: Spaces in the basedn must be URL encoded (use \$20 for spaces). If your site uses redundant directory servers for failover, you can enter multiple URLs separated by spaces.	
Query User Name	The name of a user who has read-only access to user and group directories in Active Directory. This is the user name to use for fetching user and group information. When you provide a domain name, you can provide the simple name, for example, myuser. When you provide an explicit URL, you need to provide a distinguished name, for example: cn=myuser, ou=Users, dc=electric-cloud, dc=com.	
Query User Password	The password for the query user.	
User Options section		
Note: When presting an Active Directory provider the Electric Flow convergutematically acts default		

Note: When creating an Active Directory provider, the ElectricFlow server automatically sets default values for any options (fields) that are empty. The default values match the most common Active Directory configurations. After the provider is created, you can view and modify defaults by modifying the provider.

User Base	This string is prepended to the basedn to construct the directory DN containing user records.	
User Search Filter	This LDAP query is performed in the context of the user directory to search for a user by account name. The string "{0}" is replaced with the user's login ID. Typically, the query compares a user record attribute to the substituted user login ID.	
User Name Attribute	This is the attribute in a user record that contains the user's account name.	

Field Name	Description	
Full User Name Attribute	(optional) This is the attribute in a user record that contains the user's full name (first and last) for display in the UI. If this attribute is not specified or the resulting value is empty, the user's account name is used instead.	
User Email Attribute	(optional) This is an attribute in a user record that contains the user's email address. If this attribute is not specified, the account name and domain name are concatenated to form an email address.	
Search User Subtree	Select this check box to search the specified directory by the user base and all directories below. If this check box is not selected, the search is limited to the specified directory only.	

Group Options section

Note: If you do not need group information at this time, do not enter information in the Group section, the group test fails without this information, but the test successfully authenticates users if both "user tests" are successful.

When creating an Active Directory provider, the ElectricFlow server automatically sets default values for the options/fields that remain empty. These default values match the most common Active Directory configurations. After the provider is created, you can view and modify the defaults by modifying the provider.

Enable Groups	Select this check box to enable external groups for this directory provider.	
Group Base	(optional) This string is prepended to the ${\tt basedn}$ to construct the directory DN containing group records.	
Group Member Filter	(optional) This LDAP query is performed in the groups directory context to identify groups that contain a specific user. Two common forms of group records in LDAP directories are: POSIX style groups where members are identified by account name, and groupOfNames or uniqueGroupOfNames records where members are identified by the full user DN. ElectricFlow supports both forms, so the query is passed two parameters: "{0}" is replaced with the full user record DN, and "{1}" is replaced with the user's account name.	
Group Member Attributes	(optional) This is a comma-separated attribute name list that identifies a group member. Most LDAP configurations only specify a single value, but if a mixture of POSIX and LDAP style groups exist in the directory, multiple attributes might be required.	

Field Name	Description	
Group Search Filter	(optional) This LDAP query is performed in the context of the groups directory to enumerate group records. You can choose from common templates that include either security or distribution groups (or both). These templates are based on the most common Active Directory settings.	
Unique Group Name Attribute	(optional) This is the group record attribute that contains the group name. To prevent group name overlap between multiple directory providers (or within the same provider in a multi-forested Active Directory server), we recommend setting this attribute to the distinguishedName.	
Common Group Name Attribute The Unique Group Name Attribute may not be searchable if using the distinguishedName. In this case, the Common Group Name Attribute used to search for groups, depending on the filter used.		

After filling-in all fields, click the Test button.

Three tests validate the information you supplied:

- · user authentication
- · user identified in Active Directory
- find all groups where the user is a member

If there is a test failure, correct the information you supplied and retest.

Click Save after successful test results.

New, defined directory providers will appear in the table on the Directory Provider web page.

To create a new LDAP directory provider

Enter information in the fields as follows to specify your existing LDAP users and groups to communicate with ElectricFlow. For examples of information you enter in the these fields, see the table after the following description sections.

Field Name	Description	
General section		
Provider Name	This name identifies users and groups that come from this provider.	

Field Name	Description		
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </td></tr><tr><td>URL</td><td colspan=2>The LDAP server URL is in the form <i>protocol://host:port/basedn</i>. Protocol is either <code>ldap</code> or <code>ldaps</code> (for secure LDAP). The port is implied by the protocol, but you can override the port if you cannot use the default (default port is 389 for <code>ldap</code> and 636 for <code>ldaps</code>). The <code>basedn</code> is the path to the top-level directory that contains users and groups at this site. Typically, this is the domain name where each part is listed with a <code>dc=</code> and separated by commas. Note: Spaces in the <code>basedn</code> must be URL encoded (use <code>%20</code> for spaces). If your site uses redundant directory servers for failover, you can enter multiple URLs separated by spaces.</td></tr><tr><td>Realm</td><td colspan=2>This is the realm of the LDAP directory provider, which is used to create unique user names when you have multiple providers. For example, if the realm is my-company.com, users are saved as myuser@my-company.com.</td></tr><tr><td>Query User Name</td><td colspan=2>This is the name of a user who has read-only access to the user and group directories in LDAP. This is the user name to use for fetching user and group information. When providing a domain name, you can provide the simple name, for example, myuser. When providing an explicit URL, you need to provide a distinguished name, for example: cn=myuser,ou=Users,dc=electric-cloud,dc=com.</td></tr><tr><td>Query User Password</td><td colspan=2>This is the password for the query user.</td></tr><tr><td>User Options sect</td><td colspan=2>User Options section</td></tr><tr><td>User Base</td><td colspan=2>This string is prepended to the basedn to construct the directory DN containing user records.</td></tr><tr><td>User Search Filter</td><td>This LDAP query is performed in the context of the user directory to search for a user by account name. The string "{0}" is replaced with the user's login ID. Typically, the query compares a user record attribute to the substituted user login ID.</td></tr><tr><td>User Name Attribute</td><td colspan=2>This is the attribute in a user record that contains the user's account name.</td></tr></tbody></table></style></pre>		

Field Name	Description	
Full User Name Attribute	(optional) This is the attribute in a user record that contains the user's full name (first and last) for display in the UI. If this attribute is not specified or the resulting value is empty, the user's account name is used.	
User Email Attribute	(optional) This is the attribute in a user record that contains the user's email address. If the attribute is not specified, the account name and domain name are concatenated to form an email address.	
Search User Subtree	Select this check box to search the specified directory by the user base and directories below. If this check box is not selected, the search is limited to the specified directory only.	

Groups Options section

Note: If you do not need group information at this time, do not enter information in the Group section, the group test fails without this information, but the test successfully authenticates users if both user tests are successful.

Enable Groups	Select this checkbox to enable groups.	
Group Base	(optional) This string is prepended to the basedn to construct the directory DN containing group records.	
Group Member Filter	(optional) This LDAP query is performed in the groups directory context to identify groups containing a specific user. Two common forms of group records in LDAP directories are: POSIX style groups where members are identified by account name, and groupOfNames or uniqueGroupOfNames records where members are identified by the full user DN. ElectricFlow supports both forms, so the query is passed with two parameters: "{0}" is replaced with the full user record DN, and "{1}" is replaced with the user's account name.	
Group Member Attributes	(optional) This is a comma-separated attribute name list identifying a group member. Most LDAP configurations only specify a single value, but if you have a mixture of POSIX and LDAP style groups in the directory, multiple attributes might be required.	
Group Search Filter	(optional) This LDAP query is performed in the context of the groups directory to enumerate group records.	
Unique Group Name Attribute	(optional) This is the group record attribute containing the group name.	

Field Name	Description	
Common Group Name Attribute	The Unique Group Name Attribute may not be searchable if using distinguishedName. In this case, the Common Group Name Attribute is used if searching for groups, depending on the filter.	

After filling-in all fields, click the Test button.

Three tests validate the information you supplied:

- user authentication
- user identified in LDAP
- find all groups where the user is a member

If there is a test failure, correct the information you supplied and retest.

Click Save after successful test results.

New, defined directory providers will appear in the table on the Directory Provider web page.

Examples for directory provider field descriptions

The following table provides examples for filling in the fields described above:

Field Name	LDAP example	ActiveDirectory example
Provider Type	LDAP	ActiveDirectory
Domain Name	N/A	my-company.com
Realm	my-company.com	N/A
URL	<pre>ldap://dir.company.com/dc=company,dc =com</pre>	ldaps://server/dc=company,dc=com
Query User Name	uid=JohnDoe,ou=People,dc=company,dc=com	cn=myuser,cn=Users,dc=company,dc =com

Field Name	LDAP example	ActiveDirectory example
Query User Passwor d	турж	турш
User Base	ou=People	cn=Users
User Search Filter	uid={0}	<pre>(&((userPrincipalName={0})) (sAMAccountName={0})) (objectClass=user))</pre>
User Name Attribute	uid	userPrincipalName
Full User Name Attribute	gecos	name
User Email Attribute	mail	mail
Group Base	ou=Group	cn=Groups
Group Member Filter	((member={0}) (memberUid={1}))	member={0}
Group Member Attribute	member, memberUid	member

Field Name	LDAP example	ActiveDirectory example
Group Search Filter	<pre>((objectClass=groupOfNames) (objectClass=posixGroup))</pre>	(objectClass=group)
Unique Group Name Attribute	distinguishedName	distinguishedName
Commo n Group Name Attribute	cn	cn

To edit an existing directory provider

You may change any previously supplied information in the fields.

After editing any fields, click the **Test** button.

The same three tests validate the information you supplied:

- · user authentication
- user identified in the directory provider
- find all groups where the user is a member

If there is a test failure, correct the information you supplied and retest.

Click Save after successful test results.

Edited, redefined directory providers will appear in the table on the Directory Provider web page.

Test Directory Provider

You must enter a valid User Name and Password on this page. Omitting or using an incorrect User Name and/or Password results in the test not completing or completing with an error. In either case, you will not obtain successful test results even if the Directory Provider information you supplied is valid and correctly entered into ElectricFlow.

Three tests validate the information you supplied:

- · user authentication
- user identified in a directory provider

• find all groups where the user is a member

If there is a test failure, correct the information you supplied and retest by clicking the **Test** button.

If you do not need group information at this time, and did not enter information in the Group section, the group test fails without this information, but the test successfully authenticates users if both user tests are successful.

Workflows

This page displays all workflows in the ElectricFlow system, both running and completed workflows.

Links and actions at the top of the table

- Sort the Name, Project, State, Modify Time, Workflow Definition, and Completed columns by clicking on the column name.
 - After sorting a column, the page changes to a Workflow Search Results page and more search options are available.
- Define a search to see only the workflows you choose. Click the **New Search** link.
- The "star" icon allows you to save this web page to your Home page.

Column descriptions

Name - This is the ElectricFlow-generated workflow name, which is defined by the Workflow Name Template setting. Click on any workflow name to go to that workflow's Workflow Details page.

Project - The name of the project containing this workflow. Click on any project name to go to the Project Details page for that project.

State - The current state of the workflow.

Modify Time - The time this workflow was modified.

Workflow Definition - The name of the workflow definition that ran to create this workflow. Click on any workflow definition name to go to that workflow definition's Workflow Definition Details page.

Completed - A "checkmark" in this column indicates the workflow has completed.

Actions - Delete - Use this link to delete the workflow on the same row.

Workflow Definition - create new or edit existing workflow definition

To create a new workflow definition

Enter information in the fields as follows:

Field Name	Description
Name	Enter a unique workflow definition name of your choice. You may want your workflow definition name to reflect the project where it belongs or its purpose. For example, you might set a name based on a product or team using this workflow.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </td></tr><tr><td>Workflow Name Template</td><td>This is the template used to determine the default name of jobs launched from the workflow definition. For example: \$[projectName]_\$[/increment /myproject/workflowCounter]_ \$[/timestamp] (substitute your values for the names above) Produces a workflow name like: projectFoo_123_20140102130321 Enter any combination of elements to create workflow names more meaningful</td></tr><tr><td></td><td>to you. For example, you could choose to include the build number. Note: Electric Cloud does not recommend using the ElectricFlow-generated workflowId because it is no longer a human-readable integer so it does not provide any indentifiable information and cannot be used as a counter.</td></tr></tbody></table></style></pre>

Click **OK** to save the information you entered.

Your new workflow definition will appear on the Project Details > Workflow Definitions page.

To edit an existing workflow definition

- 1. To rename the workflow definition, type a new name in the Name field and click OK.
- 2. Change or add to the description or modify the template.
- 3. Click **OK** when your edits are complete.

Workflow Definition Details

This Help topic provides a detailed overview for all views and functionality on this web page. You may want to read this topic in its entirety if this is your first experience with the ElectricFlow Workflow feature. If you are an experienced workflow user, the following quick links will take you to a specific topic section for quick review.

A note to existing workflow users: We recommend reviewing the Graph view section because workflow functionality previously found on other ElectricFlow web pages is now combined for ease-of-use on the Graph page.

Graph view

To create a New State

To create a State Definition

Action

Notifiers

Parameters

Properties

To create a Transition Definition

Using the graph's quick-access features

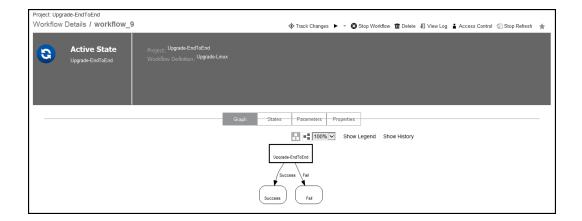
Show Legend

List view

Properties view

This page displays workflow definition components, including state and transition definition names, types, actions, descriptions, and various links to create additional transitions, state definitions, define access control, run the workflow, and so on. You can view any previously defined workflow objects and create new properties too.

Opening to the Graph view, your first glance at this page will look similar to the following:



Links and actions at the top of the page:

- Above the Workflow Definition Details page title you see (in this example) "Project: Upgrade-End to End". This breadcrumb information tells you which project you are viewing or working with and provides a link back to the project, itself.
- Immediately after the page title, you see "workflow_9", which (for our example) is the name of this
 workflow.



• Track Changes ()-Click this to open the Change History page.



- Run ()- Use this link to run the workflow definition. Hover your mouse over the drop-down arrow to see these choices:
 - Selecting Run... allows you to pick the "starting state" to run the workflow. After clicking OK from the Run Workflow pop-up dialog box, the Run Workflow page is displayed. If you previously created parameters for this workflow, they are displayed. You can accept the parameter values or change them before running the workflow.
 - Selecting Run Immediately this option uses the first (default) "starting state" to run the workflow.
 Note: Clicking Run, without using the down-arrow to make a selection, is the same as selecting Run Immediately.
- Access Control Use this link to set privileges for this workflow definition. For more information, see
 the Access Control Help topic.
- The "star" icon allows you to save this workflow definition to your Home page.

Navigation and view summary

Each button has its own expanded description section following this summary.

- Graph the page opens in this view. The main portion of this view contains the work area for visually creating your workflow.
- List this view provides a table displaying states, transitions, and so on in list form.
- Properties this view provides a table displaying previously created properties and provides links to create additional properties.
- Show Legend displays a legend for workflow objects.
- Create State this button produces the State Definition panel where you can create new states as you build or modify your workflow.
- The multi-square directional icons allow you to change your workflow graph from vertical to horizontal, depending on how you choose to see the graph.
- The drop-down menu for percentages allows you to diminish or expand the workflow graph size.

Graph view

To create a workflow, minimally you need to create state definitions and transition definitions. Optimally, those objects will have parameters and properties, and you will want to create or enforce access control and perhaps

create one or more email notifiers too. To begin, the following sections describing functionality in this view will familiarize you with how to create workflow objects and end with exploring the graph's quick-access features.

Click the **Create State** button to create a new state using the State Definition panel.

In the following screen example, you see a New State object added next to the Start state in an existing workflow graph.



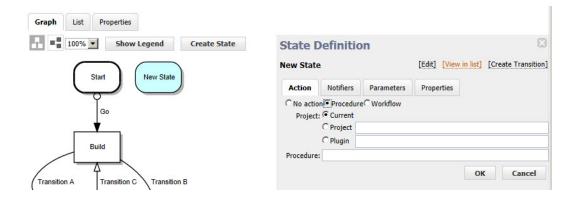
To create a new state

After clicking the Create State button, enter information in the State Definition panel as follows:

Field Name	Description
Name	(Required) Provide any name of your choice. The name must be unique to this workflow definition. If you do not provide a name, a system-generated name is created: New State, New State 2, and so on.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <div> <dl> <i> <1i> <pre> <pre> <style> <tc> </td></tr><tr><td>Startable</td><td>Select this checkbox if you want this state to be your workflow "starting" state. The first state in any workflow definition is saved as <i>Startable</i>, but you can change this later.</td></tr></tbody></table></style></pre></pre></i></dl></div>

Click **OK** and the State Definition panel provides options to configure your new State Definition.

To configure the new State Definition



In the State Definition Panel, the state name appears at the top. In this example in this topic, "New State" is the state name.

Configure the following settings.

Links at the top of the panel

- Access Control use this link to add or modify access control privileges for this state (link not illustrated in the State Definition panel example above).
- Edit use this link to modify this state's name, description, or whether or not the state is Startable.
- View in list displays the List view to see the state definition in the Workflow Definition Details table. This state will be highlighted in the List view.
- Create Transition displays the Transition Definition panel to create a transition for this state.

The "tabbed" or button sections

Action

As each state in a workflow becomes active, it performs an *action*. The state's action can be to create a job from a procedure or to start a workflow from a workflow definition. When the job or called workflow completes, On Completion transitions are evaluated to see if the workflow should change to a different state, possibly based on the outcome of the action.

Selecting No action, Procedure, or Workflow changes the fields that appear under these choices.

- No action if selected, no other information is required. Click OK.
- Procedure this selection displays the following fields:
 - o Current this selection refers to the project that contains the workflow definition.
 - Project if selected to call a different project, start typing in the text box to display a list of (non-plugin) projects from which to select. Only the first 10 matches are displayed. Type more characters to refine the list
 - Plugin if selected, start typing in the text box to display a list of plugin projects from which to choose.
 Only the first 10 matches are displayed. Type more characters to refine the list.

- Procedure start typing in this field to see a list of available procedures from which to select, depending on your previous choice of Current, Project, or Plugin. Only the first 10 matches are displayed. Type more characters to refine the list.
- Workflow this selection displays the following fields:
 - o Similar to selecting Procedure (above), select Current, Project, or Plugin.
 - Workflow if selected to call another workflow definition, start typing in this text box to display a list of workflow definitions from which to choose. Only the first 10 matches are displayed. Type more characters to refine the list.
 - Starting State start typing in this field to see a list of available starting states definitions in the workflow definition you selected. Only the first 10 matches are displayed. Type more characters to refine the list.

After selecting the procedure or starting state, press the <Tab> key to leave the field and see a list of parameters for the selected action. After entering any procedure values, click **OK** to save the changes.

Notifiers

Select this tab to see a list of email notifiers previously created for this state definition or to create an email notifier. Notifiers are commonly used to inform interested parties about transitions to the state being defined.

Note: Before you can set up an email notifier, you need an email configuration. If you have not already defined an email configuration, you can do it now. Go to Administration > Email Configurations and select the **Add Configuration** link.

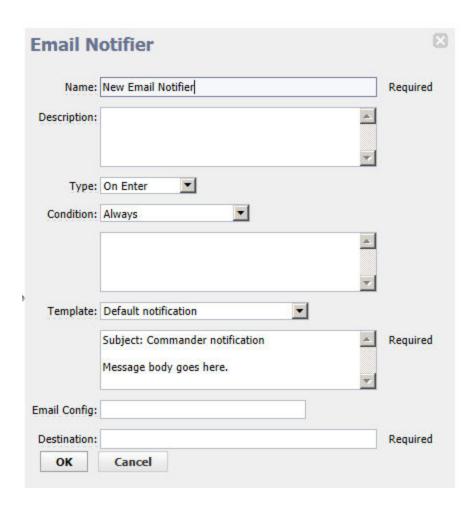
The Notifiers table contains a list of all previous created notifiers for this state definition.

Column descriptions

Column Name	Description / Actions
Name	The name of the email notifier. Select a name to go to the "edit" page for that notifier if you need to make changes.
Туре	The type of email notifier specified during creation.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with html html tags. The only HTML tags allowed in the text are: <a> <a> <a> <a> <tc> <a> <a> <a> <a> <a> <a> <a> <a> <a> <a< td=""></a<></tc></br>
Actions	Delete - Use this link to delete the notifier on that row.



To create an email notifier, click the Create Notifier link to access the Email Notifier panel.



Enter information in the fields as follows:

Field Name	Description
Name	(Required) This name can be any text string you choose. The name must be unique among other notifier names in this project, on procedures or workflow definitions called from other projects.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pr> <pr> <pr></pr></pr></pr>
	Use the pull-down menu to choose a type:
	On Completion Notifier - sends an email after the state's job or workflow completes. If no job or workflow is defined for that state, these notifiers will not be sent.
Туре	On Enter Notifier - sends an email when the state becomes the workflow's active state.
	On Start Notifier - sends an email after the state's job or workflow starts. If no job or workflow is defined for that state, these notifiers will not be sent.
Condition	Use the pull-down menu to select the type of condition you need for this email notifier. Edit the auto-supplied condition in the text box or add a completely new script for your purpose. The condition specifies whether the notifier should send a message depending on the result of a property expansion. If the result is empty or non-zero, the message is sent. If the result is "0", the message is not sent.
	Use the pull-down menu to select from a list of global, ready-to-use formatting templates. Depending on the type of email notifier you are creating, the available template choices in the drop-down menu will be different.
Template	(Required) To customize your template, edit the auto-supplied text in the text box or you can add a completely new script for your purpose. Any edits made in this text box will not be saved to the global template.
	To create a custom template, the basic structure is:
	zero or more email header lines
	blank line manage body
	message body

Field Name	Description
Email Configuration	Click inside this field or start typing to bring up a list of possible email configuration names. An email notifier that does not specify an email configuration will use the configuration named 'default' if it exists.
Destinations	(Required) This is a space-separated list of valid email addresses, email aliases, or ElectricFlow user or group names, or a property reference that expands into such a list, or you can enter an LDAP DL name (group name).

Click **OK** to save your email notifier configuration. The next time you view the Notifier table, you will see this notifier in the list.

Parameters

Similar to Notifiers, a parameter table is displayed containing all formal parameters defined for this state definition. Parameters are presented when launching the workflow definition (if it is startable), or when taking a manual transition to the state. A transition may enter values for some or all of the formal parameters defined by the state definition, in which case only the unmapped parameters are presented to the user.

• **Create Parameter** - Use this link to go to the New Parameter page to create a parameter for this state definition.

If you need assistance with creating a parameter, click the **Help** link in the upper-right corner of the New Parameter page.

Properties

Similar to the Notifiers and Parameters table, a properties table is displayed containing available properties for this state definition. If no properties were defined for this state definition, the table does not exist.

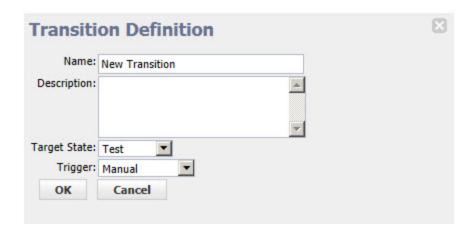
Click one of the available links to create one or more properties that will be displayed in the table when created:

- **Create Property** Click this link to go to the New Property pop-up box to create a new property for this state definition.
- Create Nested Sheet Click this link to go the New Nested Property Sheet pop-up box to create a
 nested property.
- . Access Control Click this link to go to the Access Control page for the property sheet.

The property pop-up boxes contain a Help link if you need assistance creating a property.

To create a Transition Definition

On the State Definitions panel, click the Create Transition link to see the Transition Definition panel:



Enter information in the fields as follows:

Field Name	Description
Name	(Required) Enter a unique name for the transition definition. It can be any name you choose, but the name must be unique within the state definition. If you do not provide a name, a system-generated name is created: <i>New Transition</i> , <i>New Transition</i> 2, and so on.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a>
Target State	This field displays the "best guess" for the target state. If not correct, or you intended to create a transition for a different state, use the down-arrow to choose an available target state from the list.

Field Name	Description
	Use the drop-down menu to choose the transition type.
	The four types of transitions are:
Trigger	On Completion - transition occurs when the action completes. These transitions are ignored if no action is specified for the source state. Note: On Completion transitions are taken only if the state is still active when the action completes, and are ignored if the workflow has transitioned to another state—this can occur if an On Start or Manual transition occurred before the action completed.
	 On Enter - transition occurs before sending notifiers or starting the action.
	 On Start - transition occurs immediately after starting the action. These transitions are ignored if no action is specified for the source state.
	 Manual - transition occurs when a user selects the transition in the UI and specifies parameters. The same action can occur using ectool or the Perl API by calling transitionWorkflow. Only users who have "execute" permission on the transition are allowed to use a Manual transition.
Condition	Use the drop-down menu to choose a condition. After selection, the text box is populated with a sample JavaScript string to edit for your purposes.
Condition	Note: If you select Custom, no sample is provided. Enter a script in the text box. We recommend using JavaScript.

Click **OK** to save your transition definition.

Notice the new links and options available now:

- Access Control—Use this link to add or modify access control privileges to this transition definition.
- Edit –Use this link to modify the transition definition.
- **View in list**—Use this link to see the new transition definition in the List view (table) on the Workflow Definition Details page. You can use the List view to move the transition to a different position within the state definition.
- Parameters—If parameters were defined on the State Definition Details page, you will not see them here.
 If parameters are defined on this page, you need to enter values. Manual transitions allow you to defer parameter assignment until transition time.
- Properties—Displays properties available for this object or you can create one or more properties at this time.

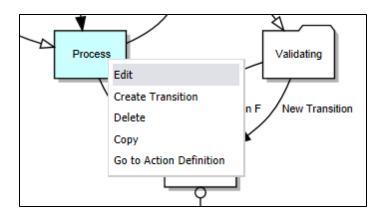
Using the graph's "quick-access" features

The following screen example is a portion of a workflow graph.

In the graph, states and transitions are links:

- Click on a state to open the State Definition panel to modify that state.
- Click on a transition to open the Transition Definition panel to modify that transition. (Hover your mouse over a transition or transition name and click when it changes color.)

However, if you *right-click* on a state or transition, a context-sensitive menu with other options is available. In the following example, right-clicking on the "Process" state provided the quick-link pop-up menu.



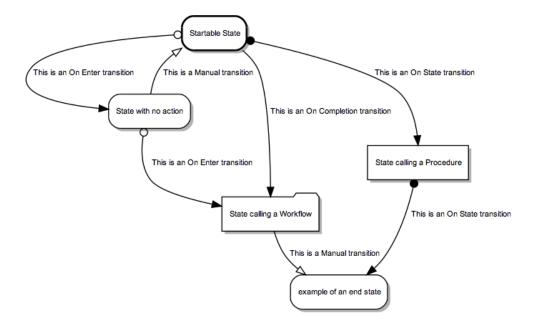
Available pop-up links include:

- For states calling a procedure or another workflow, menu choices can be:
 - Edit opens the State Definition panel to make modifications. This is the same as selecting the Edit link on the State Definition panel.
 - o Create Transition opens the Transition Definition panel.
 - o Delete deletes the state.
 - Copy makes a copy of the state with all associated transitions preserved.
 - Go to Action Definition for states calling a procedure, the Procedure Details page is opened. For states calling a workflow, the Workflow Definitions Details page for the "subworkflow" is opened.
- For states with "No action", the choices are the same, except the "Go to Action Definition" option is not available.
- For transitions, menu options include: Rename, Delete, and Copy.
- Right-clicking anywhere on the graph canvas provides two more choices:
 - Create State allows you to quickly create a new state. This is the same as clicking the Create State button at the top of the graph.
 - Rotate Graph acts as a "toggle", rotating the graph from a vertical to horizontal view and back again.

Show Legend

The following screen is another example of a workflow graph. Depending on your workflow, your graph could be very simple or much more complex than the example below.

Notice the transition line-endings and different shapes used for states.



The legend is available for reference anytime to help you become familiar with state shapes and line-ending definitions.



List view

This view will be empty until you build a workflow in the Graph view. You can view this page at any time to see a list of states and transitions in your workflow definition.

The following screen is an example of a populated Workflow Definition Details page. A **Create State Definition** link is available above the Action column, which has the same function as clicking the Create State button in the Graph view.



Column Name	Description / Actions
	This column displays state definition and transition definition names currently defined for this workflow definition.
	Transition definition names are indented under a state definition name.
Name	 Clicking a state definition name displays the State Definition panel (in the Graph view) for that state to see how it was defined or to make modifications.
	 Clicking a transition definition name displays the Transition Definition panel (in the Graph view) for that transition to see how it was defined or to make modifications.
	This is the type of action or trigger currently assigned to the state definition or transition definition.
Туре	 For states - the type may show calling a subjob, subworkflow, or no action.
	For transitions - the trigger type is provided.
	This is the action previously defined for the state or transition.
	For states - (where the Type is "subjob") the format for actions is
	<pre><pre><pre><pre><pre><pre><pre><pre></pre></pre></pre></pre></pre></pre></pre></pre>
	Clicking the second link displays the Procedure Details page for the procedure used by that state.
Action	 For states - (where the Type is subworkflow) the format for actions is <pre><pre><pre><pre><pre><pre><pre><pre></pre></pre></pre></pre></pre></pre></pre></pre>
	Clicking the first link displays the Project Details page for the workflow
	project. Clicking the second link displays the Workflow Definition Details page for the workflow used by that state.
	For transitions - clicking the Target action displays the State Definitions panel in the Graph view.
Startable	A "checkmark" in this column shows this state was defined as startable, which means this state is the beginning point for your workflow.

Column Name	Description / Actions
Description	A plain text or HTML description for this object.
	You can perform some or all of these actions: Create Transition - Click this link to go to the Transition Definition panel to create a new transition definition for this state.
	Access Control - Click this link to go to the Access Control page to define access privileges for this transition definition.
Actions	Copy - Use this link to make a copy of either the state definition or the transition definition. Note: Copying a state also copies its transitions, however, you must have modify privileges on the workflow project to make a copy.
	Move - Use this link to move a state to another position in the workflow. You will be prompted to choose a state name to move this state to a location before that state. If moving a transition, note that transition can be moved only to another location for the same state—transitions cannot be moved from one state to another state.
	Delete - Use this link to delete the state definition or transition definition on the same row.

Properties view

Select the Properties view button to see the Properties table. If no properties were defined for this workflow definition, no table appears.

The following screen is an example of the Properties table.



The following links and actions are available in the Properties table:

- Links and actions at the top of the table
 - **Create Property** Click this link to go to the New Property pop-up box to create a new property for this workflow definition.

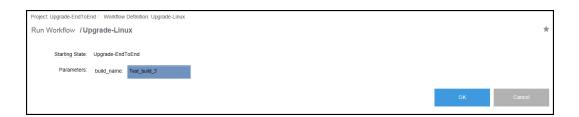
- Create Nested Sheet Click this link to go the New Nested Property Sheet pop-up box to create a nested property.
- · Access Control Click this link to change access control privileges on the property sheet.
- · Column descriptions

Column Name	Description / Actions
Property Name	Select a property name to edit that property. You can change its name, value, or add/change its description. If the property name is preceded by a folder icon, this is a property sheet. Click on the property name to open the "folder".
Value	The value currently assigned to the property.
Description	A plain text or HTML description for this object.
Actions	Delete - Click this link to delete the property on that row.

Run Workflow

Use this page to run a workflow.

- The "star" icon allows you to save this job information to your Home page.
- The "bread crumbs" Project: Upgrade-End to End / Workflow: Uprgrade-Linux provide links to a previous web page.
- The name after the Run Workflow page title is the name of the workflow you intend to run.



Starting State - This is the name of the current starting state for this workflow. If this workflow has multiple starting states and this is not the one you want to use, return to the Run Workflow pop-up dialog and select a different starting state.

Parameters -

- Any parameters previously specified for this starting start are displayed. If no parameters were defined, this area will be blank.
- If values are supplied, these are the default values specified when the parameter was created.
 If necessary, you can "type-over" these values to change them before running this workflow.
- $^{\circ}$ You must enter a value for any blank value field labeled "Required".

Click **OK** to run the workflow when your selections are complete.

Workflow Details

This page displays workflow details, including the workflow state, project name, workflow definition name, and any properties. The following links are provided for quick access to topic sections you may want to review again.

Graph

Show Legend

Using the graph's "quick-access" features

Show History

States

State Details panel

Parameters

Properties

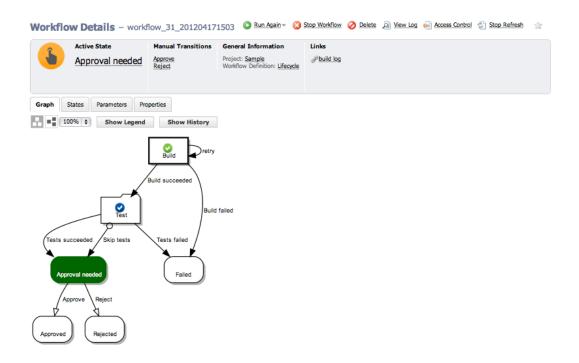
The name after the Workflow Details page title is the name of the workflow created after running a Workflow Definition. The number at the end of the workflow name is the workflow definition ID number. The ID number is auto-generated by ElectricFlow. Your workflow name will be different, depending on the information you supplied in the workflow definition.

Links and actions at the top of the table

Additional links are available if the workflow is running. When the workflow is complete, some links are no longer available.

- Run Again Use this link to run the workflow definition again.
 - When you select the **Run Again** button, you are requesting to re-run the workflow definition with the same parameter values as the original invocation.
 - If you want to change the parameter values, you can select "Run..." from the dropdown menu (small down-arrow) next to the Run Again button.
- **Stop Workflow** Select this option to stop the workflow. Any unfinished processes will complete, but no new transitions will occur.
- **Delete** Deletes this workflow.
- View Log Use this link to view the log created by this workflow.
- Access Control Use this link to set privileges for this workflow. For more information, see the Access Control Help topic.
- **Stop Refresh** This link is available only if a workflow is currently running. Click this link if you do not want the page to refresh the page automatically while the workflow is running.
- The "star" icon allows you to save this workflow information to your Home page.

While the workflow is running, your Workflow Details page will be similar to the following example:



Summary section - at the top of a running workflow page

- Next to the icon, notice that the state name, "Approval needed", is a link. Click this link to open the State Details panel for this state.
- The **Manual Transitions** section allows you select a manual transition if any were defined. However, if the manual transition contains any parameters whose values were not supplied at definition time, selecting that transition takes you to the Transition Workflow page to enter values.
- The General Information section provides links from the Project name to the corresponding Project
 Details page and from the Workflow Definition name to its corresponding Workflow Definition Details
 page.

Links

This page can include a Links section also, which will appear on the far right-side of the Summary section.

To add a link to this section:

• Create a link to any file:

To add a link, run a command in any job's job step using this format (this works for both local [disconnected] and non-local workspaces):

```
ectool setProperty "/myJobStep/report-urls/<myReportName>"
"/commander/jobSteps/<jobStepId>/<artifactDirectoryName>/<file-path>"
```

Example command:

ectool setProperty "/myJobStep/report-urls/Test Report" "/commander/jobSteps/
\$[/myJobStep/jobStepId]/testreport/index.html"

Note: If you are using this format (from ElectricCommander 4.2 and earlier):

ectool setProperty "/myJob/report-urls/<myReportName>" "/commander/jobs/<jobId>/
<workspaceName>/<artifactDirectoryName>]/<file-path>"

Example command:

ectool setProperty "/myJob/report-urls/Test Report" "/commander/jobs/\$[/myJob/jobId]/
\$[/myJobStep/workspaceName]/testreport/index.html"

It will continue to work onlyfor non-local workspaces. We recommend, however, that you change the command to use the newer format.

• Create a link to any directory:

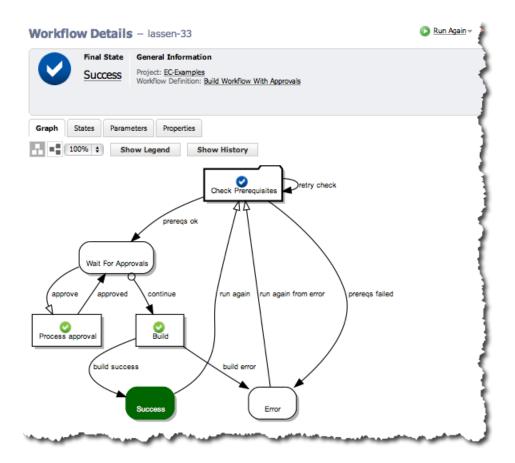
To add a link, run the following style command in any job's job step:

```
ectool setProperty "/myJobStep/report-urls/Main Job Workspace"
"file:///WinStor2/scratch/chron55build/$[/myJob/jobName]"
```

Note: Links to a directory automatically work in Internet Explorer, but if using Firefox, local links are disabled unless the default security policy is modified or an extension is used. Refer to http://kb.mozillazine.org/Links_to_local_pages_don%27t_work for more details.

When the workflow is complete...

Your Workflow Details page will be similar to the following example:



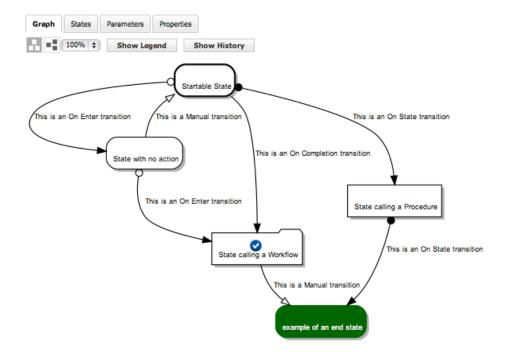
- When the workflow is complete, fewer links are available at the top of the table and the Manual Transitions section is no longer available.
- The "checkmark" icon indicates the workflow is complete. Click the state shown under Final State ("Success" in our example), to open that state's State Details panel.
- The **General Information** section is the same as when the workflow was running.

Graph tab

The Graph tab displays your workflow in a visual format. This format is the default view. The following screen is an example of a workflow graph. Depending on your workflow, your graph could be very simple or much more complex than the example below.

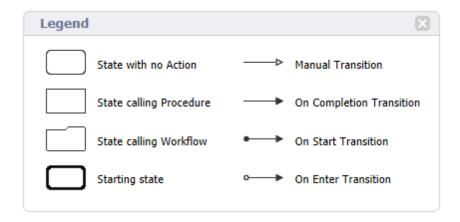
Note: This graph updates in real-time, which means you will see the "active" state highlighted with a color. An icon on the state represents an action.

Use the icons below the Graph tab to rotate the graph for better use of your screen viewing area, and you can re-size the graph using the "%" drop-down menu to choose another size.



Show Legend tab

The graph above illustrates all state and transition types. The state and transition names in this graph are labels for the type of state or transition represented. Use the Show Legend tab to see a quick reminder of state shape and transition line-ending definitions.



Using the graph's "quick-access" features

The following screen example is a portion of a workflow graph.

In the graph, states and transitions are links

- Clicking on a state takes you to the State Details panel to that state's transitions, parameters, and properties.
- Clicking on a transition takes you to the Transition Details panel to see details specifically about that transition, including its condition information.
 (Hover your mouse over the transition or transition name and click when the color changes.)

However, if you *right-click* on a state or transition, other link options are available. These options are different for a running or completed workflow versus the options on the Workflow Definition Details page that were available to you while you were defining your workflow.

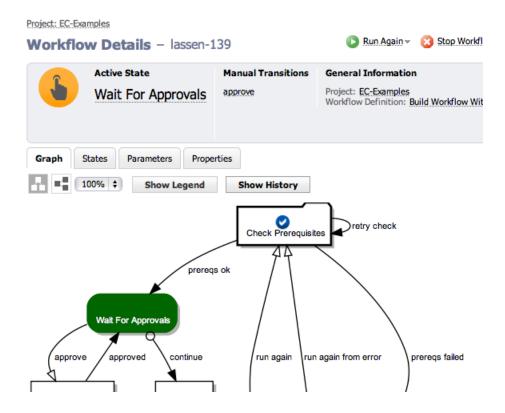
Available pop-up context-sensitive links include:

- For states, depending on which actions were defined for the state, you may have one or more of the following options:
 - o Go to Action opens the Workflow Details or Job Details page for that state.
 - Go to Action Definition for states calling a procedure, the Procedure Details page is opened. For states calling a workflow, the Workflow Definitions Details page for the "subworkflow" is opened.
 - Go to State Definition opens the State Definition panel for that state.
- For transitions, the only link available is: Go to Transition Definition, which opens the Transition Definition panel.

Show History

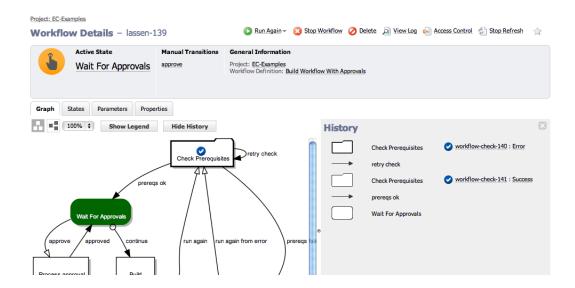
Selecting Show History allows you to see at-a-glance how your workflow progressed through its states and transitions, or for example, how many times the same state was visited and depending on outcome, which transitions were taken and then the outcome of those target states.

The following example shows a workflow started with the "Check Prerequisites" state and is now at the "Wait For Approvals" state.



So how did the workflow progress through these states?

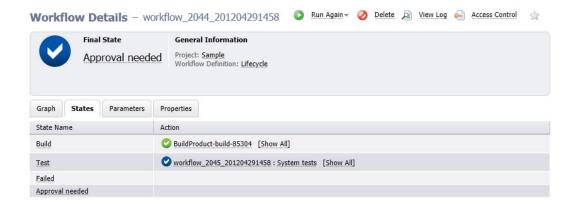
In the next screen example, with the History panel open, you can clearly see the workflow started and the check failed the first time, so the "retry check" transition was taken. The subworkflow was successful the second time, and the workflow has now transitioned to the "Wait For Approvals" state.



Notice the links displayed in the History panel. You can select links for completed jobs or workflows, which display the Job Details or Workflow Details page, respectively.

States tab

Selecting the States tab provides a table containing all state names and actions for this workflow.



The following links are available in the States table

- State Name Click any state name in this column to go to that state's State Details panel. (See the next section for more information about the State Details panel).
- Action Actions are available in this column only if you are using a subjob or a subworkflow. If so, links
 in this column connect you to the Job Details or Workflow Details page for the subjob or subworkflow,
 respectively.
- [Show All] link If the state executed multiple times during a workflow run, only the last subjob or subworkflow is displayed. However, clicking on Show All displays the results for all subjobs or subworkflows executed during a workflow run.

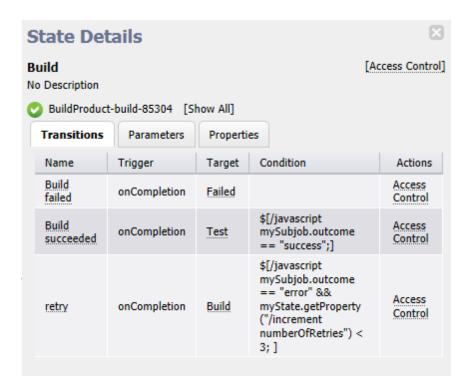
State Details panel

This panel displays state components, including transitions, parameters, and properties.

- Use the **Access Control** link to set privileges for this state. *For more information*, see the Access Control Help topic.
- The state name, "Build" (in our example), is shown at the top of the panel.
- Job or Workflow status links to the job or workflow created by the action.
- The Transitions tab is open to display the transition table listing all transitions for this state.

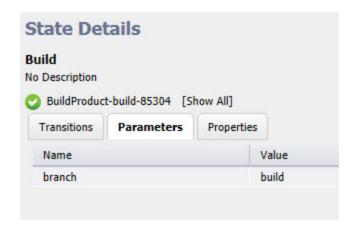
Column descriptions

Column Name	Description / Actions
Name	One or more names of transitions currently defined for this state. Click on a transition name to open the Transition Details panel.
Trigger	The type of trigger currently defined for this transition.
Target	The target state for this transition. Click on a target state name to open the State Details panel for that state.
Condition	This condition was specified when the transition was created. A condition specifies under which circumstances the transition is taken. You can edit this condition on the Transition Definition pane, but if you edit the condition during a running workflow, the change will not take effect until you run the workflow again.
Actions	Access Control - Use this link to set permissions on this transition.



The **Parameters tab** provides a table listing all defined parameters for this state. You will see any parameters passed to this state the last time this state was entered, but only if parameters were created on the State Definitions Details panel. This table does not allow you to edit the parameter's value.

These parameters are copied automatically to the top-level property sheet and you will see them in the properties table also.

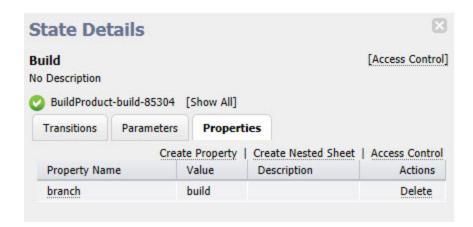


Column descriptions

Name - The name of the parameter defined for this state.

Value - The value currently defined for this parameter.

The **Properties tab** provides a table containing all defined properties for this state. If no properties were defined for this state definition, this table will be blank. All defined parameters are automatically copied to the top-level property sheet. Any custom properties defined for the state will be included in this table also.



Links at the top of the table

- **Create Property** Click this link to go to the New Property pop-up box to create a new property for this state.
- Create Nested Sheet Click this link to go the New Nested Property Sheet pop-up box to create a
 nested property sheet.
- Access Control Click this link to add or decrease access control on the property sheet.

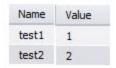
Column descriptions

Column Name	Description / Actions
Property Name	Select a property name to edit that property. You can change its name, value, or add/change its description. If the property name is preceded by a folder icon, this is a property sheet. Click on the property name to open the "folder".
Value	The value currently assigned to the property.
Description	A plain text or HTML description for this object.
Actions	Delete - Click this link to delete the property on that row.

Parameters tab

This tab provides a table listing all defined parameters for this workflow's starting state. You will see any parameters passed to the starting state when this workflow was launched, but only if parameters were created on the State Definitions Details page. This table does not allow you to edit the parameter's value.

These parameters are copied automatically to the top-level property sheet and you will see them in the properties table also.



Column descriptions

Name - The name of the parameter defined for this starting state.

Value - The value currently defined for this parameter.

Properties tab

Select the Properties tab to see the Properties table. If no properties were defined for this workflow definition, this table will be blank.

		Create Property Create N	ested Sheet Access Control
Property Name	Value	Description	Actions
property1	full build		Delete
property2	test only		Delete
test1	1		Delete
test2	2		Delete

The links at the top of the table

- Create Property Click this link to go to the New Property pop-up box to create a new property for this
- Create Nested Sheet Click this link to go the New Nested Property Sheet pop-up box to create a
 nested property sheet.
- Access Control Click this link to add or decrease access control on the property sheet.

Column descriptions

Column Name	Description / Actions
Property Name	Select a property name to edit that property. You can change its name, value, or add/change its description. If the property name is preceded by a folder icon, this is a property sheet. Click on the property name to open the "folder".
Value	The value currently assigned to the property.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </td></tr><tr><td>Actions</td><td>Delete - Click this link to delete the property on that row.</td></tr></tbody></table></style></pre>

Workflow Log

This page displays the workflow log.

Notice the name after the page title, "template-1257" — this is the name of the workflow, which includes the workflow name plus the ElectricFlow auto-generated ID number, and it is the object of the workflow log in our example.

Links and actions above the table

- Drop-down menu Use the down-arrow to select Error, Warn, or Info. The first time you choose a severity level, it will become your default level—the one you always see first when you view this page. Selecting another level changes your default view.
- The "star" icon allows you to save this workflow definition to your Home page.
- The "bread crumbs" *Project: SoftwareBuild / Workflow: template-1257* provide links to a previous web page.



Column descriptions

Time - Displays the time when an event occurred that caused the server to generate a message.

Severity - the three severity levels are:

- ERROR An unexpected failure was encountered while entering a state or launching a sub-action.
 Generally, this error indicates a critical problem with the workflow that requires fixing the workflow definition.
- WARN A non-critical issue was encountered while the workflow was running.
- INFO Provides workflow activity information including the state entered, transitions taken, and so on.

User - The name of the user or project principal that explicitly launched the job. This property is blank when the job is launched by a schedule.

Subject - Objects displayed in this column are the subject of the message. These objects are linked to either the Workflow Details or the State Details page.

Message - A text message generated by the ElectricFlow server while the workflow was running.

Transition Workflow

Use this page to transition your workflow to a different startable state.

Note the name after the page title, "template-99"—this is the name of the workflow in our example.

Links and actions above the table

- The "star" icon allows you to save the Transition Workflow page to your Home page.
- The "bread crumbs" Project: SoftwareBuild / Workflow: template-99 provide links to a previous web
 pages.



Field descriptions

Enter information in the fields as follows:

Field Name	Description
Transition	The name of this transition.
State	The name of the state that owns the transition— the active state of the workflow.
Target State	The name of the target state— the state where the workflow will transition.
Parameters	In the screen example above, "version" is the name of the parameter and you must enter a value for this parameter. This is a required field.

Click **OK** to transition the workflow to the target state.

Workspaces

This page displays all workspaces currently available to ElectricFlow. At a glance, you can see all paths to every workspace.

- To create a new workspace, click the Create Workspace link.
- To find a workspace, use the New Search link.

Column descriptions

Column Name	Description / Actions
Workspace Name	The name created for this workspace. Click on a Workspace Name in this column to edit that workspace.

Column Name	Description / Actions
Zone	The name of the zone where this workspace resides.
Enabled	This workspace is enabled if this box is "checked". Note: When this box is checked, the workspace is enabled, which means it can be accessed. In the case where a job or job step cannot access the workspace, the job will "queue", waiting for the workspace to become available.
Local	If the box is "checked", the workspace is "local", which means files in the local workspace are accessible only from the resource that originally created the file.
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <htm> <div> <dl> <i> <pre> <pre> <style> <tc> <ul</td></tr><tr><td>Drive Path / UNC Path / UNIX Path</td><td>The path to the workspace.</td></tr><tr><td>Actions</td><td>Copy - Use this link to make an exact duplicate of an existing workspace, then select the copy to go to the Edit Workspace page to change the workspace name.</td></tr><tr><td></td><td>Delete - Use this link to delete the workspace on the same row.</td></tr></tbody></table></style></pre></pre></i></dl></div></htm>

Note: Add access control privileges to a workspace on the Edit Workspace page.

Workspace - create new or modify existing workspace

To create a new workspace

Enter information in the fields as follows:

Field Name	Description
Name	Enter a unique name to distinguish this workspace from other workspace names.

Field Name	Description
Description	A plain text or HTML description for this object. If using HTML, you must surround your text with <html> </html> tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </td></tr><tr><td>Enter three paths to use or</td><td>resources to access the workspace root directory:</td></tr><tr><td>UNC Path</td><td>A path in UNC notation (such as //server/ec/a/b/c) to use on Windows machines. The path can use either slashes or backslashes for separators.</td></tr><tr><td>Drive Path</td><td>A second path for Windows machines, based on drive notation such as \mathbb{N}: or \mathbb{N}: /b/c. If the drive path contains more than just a drive letter, additional directory names must match the trailing directories of the UNC path. For example, if the UNC path is //server/ec/a/b/c, then \mathbb{N}: /b/c can be used as the drive path (agents automatically map \mathbb{N}: to //server/ec/a), but \mathbb{N}: /x/c is not allowed.</td></tr><tr><td>UNIX Path</td><td>The path to use for UNIX machines, typically based on an NFS mount. You must ensure appropriate mounts were created on all machines where the path will be used. ElectricFlow does NOT automatically mount filesystems for UNIX.</td></tr><tr><td>Conversely, if the workspa</td><td>be used only on UNIX machines, you can omit the UNC and drive paths. ce will be used only on Windows machines, you can omit the UNIX path. e the Workspaces and Disk Space Management Help topic.</td></tr><tr><td>Zone</td><td>Use the drop-down arrow to select a zone for this workspace.</td></tr><tr><td>Local</td><td>Select this checkbox if the workspace is "local", which means files in the local workspace are accessible only from the resource that originally created the file.</td></tr><tr><td>Enabled</td><td>"Check" this box to enable this workspace. When this box is checked, the workspace is enabled, which means it can be accessed. In the case where a job or job step cannot access the workspace, the job will "queue", waiting for the workspace to become available.</td></tr><tr><td>Credential Project and Credential Name</td><td>If you need to enter a user name and password to mount the workspace, you must enter the credential project name and then choose the appropriate credential name, which contains the required user name and password. Click the Browse link to select information for these fields.</td></tr></tbody></table></style></pre>

Click **OK** to save your workspace specification and see it listed on the Workspaces page.

If you need to add properties or privileges to this workspace:

Click **OK** to create the workspace and return to the Workspaces page, then click on the workspace name to go to the Edit Workspace page.

To edit an existing workspace

All previously supplied workspace information is available on this page and you can modify existing information or add new properties or privileges.

- Click the Access Control link [at the top of the page] to add access privileges to this workspace.
- To assign a new workspace name, highlight the existing name and type-in a new name.
- Re-assign paths if necessary.
- Use the drop-down menu to select a different zone if necessary.
- Add or modify the Credential Project name or the Credential Name if necessary.

Click **OK** to save your modified information.

If Custom Workspace Properties were already assigned, you may modify these properties or add new ones. If a property name is preceded by a "folder" icon, this denotes a Nested Sheet.

Click the "star" icon at the top of the page to add this page to your Shortcut section on the Home page.

For more workspace information, see the Workspace and Disk Space Management Help topic.

Workspace File

This page displays a file from the workspace for a job, typically the log file for a job step. For log files, this file contains both standard output and standard error from the step.

If you view a file while the job step generating the file is still running, the page refreshes automatically as information is added to the file.

Tip

If the log does not contain the output that you were expecting, check the job step commands. If the step redirects its output to a different location, this information will not appear in the log file.

ElectricFlow Platform Objects and Functionality

All the topics in this section are about a particular object or functionality in the ElectricFlow platform. They are the same as the topics in the Help system on the automation platform UI.

From the automation platform UI, you can access the topics as follows:

• Direct access to these Help topics is from the Help system Table of Contents. Although these topics are not directly linked within the ElectricFlow product, many "page-specific Help topics contain links to one or more of these topics to provide more detailed information guickly.

• To access these Help topics while using ElectricFlow, click the **Help** link at the top-right corner on any automation platform web page.

Access Control

Privileges

Users and Groups

Special users and groups

Access Control Lists - allow and deny

Inheritance

System Objects

Access Control and Jobs

Examples for Increased Security

Overview

ElectricFlow provides a comprehensive mechanism to control how individuals use the system.

- Users must log in to view information or perform operations.
- After you log in, system access is limited based on:
 - o user name
 - the groups to which that user belongs
 - o permissions specified for various ElectricFlow objects

Access Control is the ElectricFlow functionality that provides security for all system objects. *After* you are familiar with the following information describing the Access Control system, review the two examples, Basic ACL Setup and Team ACL Setup (at the end of this Help topic), which may provide guidance or insight for how you might setup enhanced ElectricFlow security at your site.

Privileges

ElectricFlow supports four privilege types for each object:

- Read Allows object contents to be viewed.
- Modify Allows object contents (but not its permissions) to be changed.
- Execute If an object is a procedure or it contains procedures (for example, a project), this privilege allows object procedures to be invoked as part of a job. For resource objects, this privilege determines who can use this resource in job steps.
- Change Permissions Allows object permissions to be modified.

Users and Groups

ElectricFlow uses account information from multiple sources. In most cases, the primary account information source is an external LDAP or Active Directory repository:

- Both user and group information is retrieved from the repository
- · Local users and groups can be defined within ElectricFlow

Note: To view user and group information, and to modify local user and group information, click the Administration tab and select either the Users or Groups subtab. External account information *cannot* be modified from ElectricFlow.

If the same user exists in multiple sources, only the highest priority name is used. A priority order is defined among external repositories, but local names have the highest priority. Thus, if you define a local user with the same name as an LDAP user, you will effectively mask the user in the LDAP account.

For local user accounts, only local groups are considered—group information from external repositories is not used. For accounts from a particular repository, groups from that repository are used along with local groups, but groups in other repositories are not considered.

In other words, you can define local groups in ElectricFlow to supplement groups defined in external repositories. When you view information about users in ElectricFlow, only relevant groups are shown. For example, when you view group information for a local user, only local groups are displayed.

Note: Groups are managed by name only, without regard to source. If a particular group name exists in different repositories, there is no way to distinguish between these groups inside ElectricFlow. Access given to one group is the same for any other group with the same name.

Special Users and Groups

The *admin* local user has special significance. If you are logged in as "admin," you automatically have all privileges on all objects, regardless of any other system settings. This privilege set is a fallback mechanism in the event too many privileges get removed for an object, leaving it unusable.

The admin account cannot be deleted.

If the admin account is missing, ElectricFlow will recreate the account the next time it starts up with password "changeme."

The *Everyone* group is predefined by ElectricFlow and cannot be deleted. Every user is automatically a member of the Everyone group.

For each project, ElectricFlow uses the project name and automatically defines a user associated with that project, called the project principal. For example, the *project principal* for a project named "nightly builds" is "project: nightly builds" (notice that there are two spaces in this name)— this principal is used for jobs running under the project, as described in the "Access control and jobs" section later in this Help topic.

Access Control Lists - allow and deny

Each ElectricFlow object, such as a project, procedure, job, property sheet, workspace, schedule, or resource contains an access control list (ACL). Individual properties do not have their own access control, but instead use

access control lists from their parent containers.

To view the access control list for an object, go to the web page displaying the object and click the **Access Control** link. An access control list can contain any number of entries— each entry names a particular user or group and indicates *allow*, *deny*, or is blank for each of the four privileges.

To determine whether a user can perform an operation on a particular object, ElectricFlow determines which of the four privileges is required for that operation, then searches all access control entries that refer to that user or groups containing that user. To be allowed access, at least one of the matching entries must specify "allow" and none of the entries can be "deny." A "deny" entry overrides an "allow" entry in the same ACL.

Inheritance

If an object's access control list does not indicate what to do for the user who is trying to access it, ElectricFlow looks in the access control list for the object containing the target object, then its parent, and so on up to a top-level object covering the entire server. This mechanism is called *inheritance*.

For example, a nested property sheet on a procedure inherits access control information from its parent property sheet, the procedure, the project containing the procedure, and the server. When you view the object access control list, you also see inherited access control lists so you can trace the object's inheritance chain.

When ElectricFlow performs an access check on an object, it begins with the access control list for that object.

- If a "deny" entry matches an entry in that list, access is denied.
- If there is no "deny" entry, but a matching "allow" entry is found, access is allowed.
- If there is no matching entry that specifies either "allow" or "deny," ElectricFlow moves to the next access control list in the inheritance chain and repeats the process.
- A matching entry in a lower-level access control list takes priority over entries in higher-level access control lists. If no matching entry in any list is found, access is denied.

Inheritance allows you to control access to a collection of objects because you can make changes in just one place. When new objects are created, their access control lists are empty, so all access control for the entire system is determined by the server-level list.

- Typically, when a new project is created, you define an access control list for the project and allow everything in the project to inherit from the project.
- If you do not define additional access control information in project components, all project access is determined by the project's access control list.

In some situations you may want tight control over object access, so changes in parent objects do not affect access to the object or its children. To achieve this, you can "break inheritance," which means a particular object does not inherit from its ancestors. The ElectricFlow web interface allows you to break and restore inheritance for any object on which you can change permissions. However, *if you break an object's inheritance, it affects all children of that object*.

Additional Information about "deny"

If a user matches a deny rule, however indirectly, the user is denied access. Even if a specific ACL entry granting (for example) read access by user name exists, a less specific deny ACL entry overrides that specific

rule. Therefore, avoid "deny" rules if at all possible (especially to groups), or you may experience some unexpected permissions issues.

You can, however, enforce an implicit "deny" by allowing an object to fall off the end of the inheritance chain. This practice allows denial behavior without the risks associated with an explicit ACL entry to deny access.

IMPORTANT: If you break inheritance on an object with an empty access control list, the object will become completely inaccessible. You will not be able to restore inheritance because you no longer have the right to change permissions on that object. If this happens, you must contact your system administrator who can login as admin and restore inheritance.

Consider ElectricFlow principals (actors) from two perspectives:

- A running job or workflow runs using the project as the identity, so if a job requires special permissions you must use the project as the principal in the ACL entry.
- Users and groups come into play when humans manipulate the GUI or execute ectool commands from the command line (or via a shell script) -- basically anywhere that a script or human needs to log in to ElectricFlow. Having logged in, a group and id exist, which can be used in ACL entries.

For example, consider a property that will store a build number. You can set an ACL that breaks inheritance (implicit deny), and then add an ACL entry granting permissions for the project to be able to increment the property value. This prevents humans from manipulating the build number but permits a running job to increment the build number at will.

Notes:

- When it is practical, a job will run with multiple identities—the project and the user that launched the job.
 You cannot, however rely on this behavior for security because jobs run from a schedule (such as CI jobs) and jobs started by a workflow do not have a user principal, just the project principals.
- A job can have multiple project principals because it "accumulates" new project principals when it calls
 into another project's subprocedure (and then releases the other project's identity when the
 subprocedure returns). This makes is possible to create trusted libraries.

System Objects

A few special *system objects* contain access control lists related to overall ElectricFlow system administration. These access control lists are available from the Administration page. The system objects are:

- **Server** an ElectricFlow system top-level object. Every other object in the system is contained in the server object and inherits access control information from the server object unless inheritance is broken.
- Administration Read permission allows access to the getStatus, getDatabaseConfiguration[s], getEmailConfig[s], and export(global) API functions.
 Modify permission allows access to the shutdown, setDatabaseConfiguration, createEmailConfig, deleteEmailConfig, modifyEmailConfig, and import(global) API functions.
- Artifacts Read permission allows access to the getArtifact API functions.

 Modify permissions allows access to createArtifact and deleteArtifact functions.

- **Directory** *Read* permission allows access to the getUsers, getGroups, and getDirectoryProviders API functions.
 - *Modify* permission allows access to the createUser, createGroup, deleteUser, deleteGroup, createDirectoryProvider, modifyDirectoryProvider, deleteDirectoryProvider, testDirectoryProvider, and moveDirectoryProvider API functions.
- Email Configurations *Modify* permission allows access to the createEmailConfig and deleteEmailConfig API functions.
- Force Abort Execute permission controls access to the --force flag on abort Job. By default, the
 ACL is created with Everyone: execute permission in addition to inheriting from the "Server". To force
 abort a job, the user must have execute permission on the job and execute permission on the
 forceAbort ACL.
- Logging Modify permission allows access to the logMessage API function.
- Licensing Read permission allows access to the <code>getLicense[s]</code> API functions.

 Modify permission allows access to the <code>importLicenseData</code> and <code>deleteLicense</code> API functions.

 Execute permission allows access to the <code>getAdminLicense</code> API function.
- Plugins Modify permission allows access to the createPlugin, deletePlugin, installPlugin, uninstallPlugin, promotePlugin API functions, and the modifyPlugin API function requires modify permission on the target plugin. For getPlugin, Read permission is required on the target plugin.
- **Priority** *Execute* permission allows the user who launches a procedure (using the runProcedure API function) to raise the priority of the job.
- Projects Modify permission allows access to the createProject and deleteProject API functions.
- Repositories Read permission allows access to the getRepository API function.
 Modify permission allows access to the createRepository, deleteRepository, modifyRepository, and moveRepository API functions.
- **Resources** *Modify* permission allows access to the createReource and deleteResource API functions.
- Session Execute permission allows access to the login API function.
- Workspaces Modify permission allows access to the createWorkspace and deleteWorkspace API functions.
- ZoneAndGateways Modify permission allows access to the createZone and deleteZone API functions. Modify permission also allows access to the deleteResource API function when the resource belongs to a Gateway.

Access Control and Jobs

When a job executes, it usually needs to access objects in ElectricFlow. For example, a job step command may refer to a parameter value, which is a property associated with the job object, or a step may invoke *ectool* to modify properties or any other ElectricFlow state. This process leads to questions:

• Under which user name does the job execute?

Procedures always run under the project principal user ID for the project that contains the procedure. If a procedure invokes a subprocedure in another project, that subprocedure will run under its own project's project principal and the project principal of its calling procedure. When a procedure is running under

multiple project principals, its steps will perform any operations for which any one of its project principals allow.

• How does ElectricFlow initialize job permissions when the job starts?

This question pertains to job object permissions. When a job starts, ElectricFlow sets full access control entries on the job for the project principal and the user who launched the job— assuming the job was launched by a user and not a schedule.

• What permissions are needed to abort a job?

Aborting a job requires *execute* permission on the job. If a job is launched by a user, that user is given all privileges on the job. If a job is launched by a schedule, the schedule's execute privileges are copied to the job.

The access control system determines whether jobs can be executed or not.

- For a user to run a job without creating a schedule, the user must have *execute* permission on the top-level procedure being executed.
- To create a schedule to run a procedure, a user must have *modify* permission for the project containing the schedule. After a schedule is created, no additional permissions are required to start jobs under the auspices of that schedule.

Examples for Increased Security

Initially, the *Everyone* group had Read and Execute permission on all ElectricFlow objects by default. The following examples illustrate how you might customize ElectricFlow default Access Control settings for your organization.

Caution: Make sure you are familiar with the Access Control system before making changes. Incorrectly setting ACLs can severely impact ElectricFlow behavior.

Abbreviations used in the examples

A = Allow

I = Inherit (some places referred to as "Don't Care")

D = Deny

Change Permission = Change

Example 1 - Basic ACL Setup

This example illustrates how you might change Access Control defaults to set up a basic level of security. The easiest way to manage ACLs is by using different ElectricFlow groups that allow users to perform specific roles.

For our example, we will create and use the groups called EC-administrator and EC-designer, and we will modify the default Everyone group. (When setting up your groups, you would choose any group names that suit your organization, but note that the "EC-" prefix is reserved for Electric Cloud use only.) This is how we want to define our groups:

- EC-administrator group

 This group has most of the power and abilities of the 'admin' user login, but allows for more flexibility.
- EC-designer group
 This role is more significant. Users in this group can create and modify most ElectricFlow objects.
- Everyone group
 This predefined group is used for people who only need to run ElectricFlow procedures, and have no need to create or modify them.

As you begin to modify the Access Control system, note that the system has some project principal ACEs that should remain in place for correct system operation.

Server ACLs

The Server is the foundation for almost all ACL checks. Generally, every other object in the system inherits these privileges. For this basic ACL configuration we want to set Server ACLs this way:

Туре	Name	Read	Modify	Execute	Change
group	EC-administrator	A	A	А	А
group	Everyone	I	I	I	I

Custom Server Properties ACLs

An often overlooked ACL setting is Server Property Sheet. Allow the Everyone group Read permission on this object.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	1	I

System Object ACLs

There are 14 categories of System objects with ACLs.

1 - Administration – allow the EC-designer group Read privileges.

Туре	Name	Read	Modify	Execute	Change
group	EC-designers	A	1	1	1

2 - Artifacts – allow the EC-designer group Read privileges.

Туре	Name	Read	Modify	Execute	Change
group	EC-designers	A	А	1	1
group	Everyone	А	I	I	I

3 - Directory – allow the Everyone group Read privileges.

Туре	Name	Read	Modify	Execute	Change
group	EC-designers	A	1	1	1

4 - Email Configurations – allow the Everyone group Read and Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	А	ı

- **5 Force Abort** no changes needed for this category.
- 6 Logging no changes needed for this category.
- 7 Licensing no changes needed for this category.
- **8 Plugins** allow the Everyone group Read privilege.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	1	1

- **9 Priority** no changes needed for this category.
- **10 Projects** allow the Everyone group Read and Execute privileges, and the EC-designer group everything except Change privilege.

Туре	Name	Read	Modify	Execute	Change
group	EC-designers	Α	А	Α	I
group	Everyone	А	I	А	I

11 - Repositories - allow the Everyone group Read and Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	1	1	1

12 - Resources - allow the Everyone group Read and Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	I	Α	I

13 - Session - allow the Everyone group Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	1	1	Α	ı

14 - Workspaces - allow the Everyone group to have Read and Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	А	I

Note: Sites that want a slightly more open system might consider granting the "EC-designers" group Modify access to Resources and Workspaces.

Plugin Project ACLs

Generally, you will want to allow the Everyone group Read privileges on every plugin you expect to use. (Each plugin has its own name, which is the project name also.)

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	I	I	I

Project ACLs

Next are special Electric Cloud projects, Default, EC-Utilities and EC-Examples.

Projects: Default, EC-Utilities, EC-Examples

Use caution when setting privileges on the EC-Utilities project. These are powerful procedures and some have their own ACL settings. Remember the 'admin' user will always be able to use these utilities. Allowing the EC-administrator group full privileges accomplishes the same goal for users in that group. We allow the Everyone group Read only privileges. This is a unique project because it does not inherit ACLs from anywhere else (for example, Server or Projects).

Туре	Name	Read	Modify	Execute	Change
group	EC-administrator	Α	Α	Α	А
group	Everyone	Α	I	I	I

Example 2 - Team ACL Setup

Some organizations may want to have two or more different and separate teams use ElectricFlow at the same time. In this example, people on one team have little knowledge about the other team who uses ElectricFlow and no direct access to any objects used by the other team. Our "Team" example describes how to modify default ElectricFlow Access Control settings to create a multi team security configuration using ACLs.

To divide your ElectricFlow system for use by separate teams, you will want to create one or more separate Projects for each team. You may decide to share some projects across teams.

Create your teams (groups). For our example, we will begin with two teams, T1 and T2. Each team has a designer group and a user group. Using groups allows the flexibility to have a designer from one group work as a user in another group.

- Create the EC-administrator group
 This group has most of the power and abilities of the 'admin' user login, but allows for more flexibility.

 Assigning people to the EC-administrator group allows better tracking of ElectricFlow administrative activity.
- Create two designer groups:T1-designer and T2-designer
 Members of the designer groups will have the ability to create and modify procedures and other
 ElectricFlow objects for their team. A member of the designer group from one team will not be able to view or use the objects of the other team.
- Create two user groups: T1-user and T2-user
 Members of the user groups will have the ability to run procedures that belong to their team. A member of the user group from one team will not be able to view or use objects that belong to the other team. who are not in any other group. This setup allows other users to access ElectricFlow without exposing the work of any of the teams. A person not in a group for either team will not be able to see objects of either of the teams.

As you begin to modify the Access Control system, note that the system has some project principal ACEs that should remain in place for correct system operation.

Server ACLs

The Server is the foundation for almost all ACL checks. Almost all objects in the system inherit these privileges. For this basic Team configuration, we will set the Server ACLs this way:

Туре	Name	Read	Modify	Execute	Change
group	EC-administrator	Α	Α	Α	А
group	Everyone	I	I	I	1

Custom Server Properties ACLs

An often overlooked ACL setting is Server Property Sheet. Allow the Everyone group Read permission on this object.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	1	I

System Object ACLs

There are 14 categories of System objects with ACLs.

1 - Administration – allow both designer groups Read privileges.

Туре	Name	Read	Modify	Execute	Change
group	T1-designer	Α	1	1	1
group	T2-designer	Α	I	I	1

2 - Artifacts – allow the Everyone group Read privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	I	I	I

3 - Directory – allow the Everyone group Read privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	I	I	1

4 - Email Configurations – allow the Everyone group Read and Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	А	1

- 5 Force Abort no changes needed for this category.
- **6 Logging** no changes needed for this category.
- 7 Licensing no changes needed for this category.
- **8 Plugins** allow the Everyone group Read privilege. Generally, privileges set here will be inherited by all plugins.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	1	1	1

- 9 Priority no changes needed for this category.
- **10 Projects** no changes needed for this category. Privileges for projects are now handled on a case by case basis.
- 11 Repositories allow the Everyone group Read privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	1	I

12 - Resources - allow the Everyone group Read and Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	I	А	-

13 - Session - allow the Everyone group Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	1	ı	Α	1

14 - Workspaces - allow the Everyone group Read and Execute privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	I	Α	1

Plugin Project ACLs

Generally, you will want to allow the Everyone group Read privileges on every plugin you expect to use. (Each plugin has its own name, which is also the project name.) Because you may want to use many plugins, you might want to use a script.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	I	I	I

ACLs for Existing Projects

Next are a few Electric Cloud projects, and two sample projects that can be seen and used by their respective team members only.

Project: EC-Examples

These are nice examples and we want to grant Everyone the ability to see and run them.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	1	Α	1

Project: Default

Inherited settings will allow the EC-administrator group full privileges.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	А	1

Project: EC-Utilities

Use caution when setting privileges on the EC-Utilities project. These are powerful procedures and some have their own ACL settings.

Remember the 'admin' user will always be able to use these utilities. Allowing the EC-administrator group full privileges accomplishes the same goal for users in that group. We allow the Everyone group Read only privileges. This is a unique project because it does not inherit ACLs from anywhere else (for example, Server or Projects).

Туре	Name	Read	Modify	Execute	Change
group	EC-administrator	Α	A	Α	А
group	Everyone	Α	I	I	I

Project: Electric Cloud

This project has one procedure and we want to grant everyone the ability to see it.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	A	1	1	1

ACLs for New Projects

Now we introduce the concept of ACL settings as they would be in a multiple team environment. For this example, we assume you have five projects being used by two teams in the following way.

- · Project-A used by Team1
- · Project-B used by Team1
- Project-C used by Team2
- Project-D used by Team2
- Project-E used by Team1 and Team2

Project-A

This project is visible and usable only by ElectricFlow users who are members of either the T1-user group or the T1-designer group. Notice that the designer group receives full permissions.

Туре	Name	Read	Modify	Execute	Change
group	T1-designer	Α	A	Α	Α
group	T1-user	А	I	А	I

Project-B

This project is the same as Project-A.

Туре	Name	Read	Modify	Execute	Change
group	T1-designer	Α	Α	Α	Α
group	T1-user	Α	I	А	1

Project-C

This project is visible and usable only by ElectricFlow users who are members of either the T2-user group or the T2-designer group. Notice that the designer group receives full permissions.

Туре	Name	Read	Modify	Execute	Change
group	T2-designer	Α	Α	Α	А
group	T2-user	Α	I	А	I

Project-D

This project is the same as Project-C.

Туре	Name	Read	Modify	Execute	Change
group	T2-designer	Α	А	А	А
group	T2-user	Α	I	Α	I

Project-E

This project is a superset that includes both teams.

Туре	Name	Read	Modify	Execute	Change
group	T1-designer	Α	А	А	А
group	T1-user	Α	I	Α	I

Туре	Name	Read	Modify	Execute	Change
group	T2-designer	Α	A	A	А
group	T2-user	А	I	А	I

Enabling Same Team Subprocedures

To enable procedures in one project to use procedures in another project, one more step is necessary.

Add the project principals for both T1 projects to the T1-user group.

In this case, these are called: "project: Project-A" and "project: Project-B".

Add the project principals for both T2 projects to the T2-user group.

In this case, these are called: "project: Project-C" and "project: Project-D".

Optional: Restricting Resources and Workspaces by Team

Resource ACLs

Resource: local

Because each team will have dedicated resources, it is useful to keep one resource available for everyone to use.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	I	Α	1

T1-resource

This resource is visible and usable only by ElectricFlow users who are members of either the T1-user group or the T1-designer group. This resource also needs Execute permission for the project itself to allow it to run from a schedule.

Туре	Name	Read	Modify	Execute	Change
group	T1-designer	Α	1	А	1
group	T1-user	Α	I	А	I

T2-resource

This resource is visible and usable only by ElectricFlow users who are members of either the T2-user group or the T2-designer group. This resource also needs Execute permission for the project itself to allow it to run from a schedule.

Туре	Name	Read	Modify	Execute	Change
group	EC-administrator	А	Α	Α	А
group	Everyone	I	I	I	I

Workspace ACLs

Workspace: default

Because there will be dedicated workspaces for each team, it is useful to keep one workspace available for Everyone to use.

Туре	Name	Read	Modify	Execute	Change
group	Everyone	Α	I	А	I

T1-workspace

This workspace is visible and usable only by ElectricFlow users who are members of either the T1-user group or the T1-designer group.

Туре	Name	Read	Modify	Execute	Change
group	T1-designer	A	I	А	_
group	T1-user	А	I	А	I

T2-workspace

This workspace is visible and usable only by ElectricFlow users who are members of either the T2-user group or the T2-designer group.

Туре	Name	Read	Modify	Execute	Change
group	T2-designer	Α	I	A	I
group	T2-user	А	I	Α	I

Artifact Management

Artifact Objects

Access Control

Publishing Artifact Versions

Enable Compression

Include and Exclude Patterns

Retrieving Artifact Versions

Entering Filters During a Retrieve Operation

Dependent Artifact Versions

Artifact Cache

Cleaning Up Repositories and Caches

Overview

During the course of the development process, project components or "outputs" are produced for consumption by other projects or the main project to create a whole entity. Ultimately, applications are assembled by combining these various outputs (dependencies) and packaging them together into a single deployable software application.

Outputs can be libraries, scripts, graphics, and so on. This output is an *artifact* that can be consumed by other projects.

Every stage in the application lifecycle produces and consumes artifacts. Artifacts are critical to the *build-test-deploy* process.

Using artifacts can:

- improve performance across builds
- provide better reusability of components
- · improve cross-team collaboration with greater traceability

For example, instead of each developer repeatedly downloading third-party packages from external sources, these components can be published and versioned as an artifact. A developer then simply retrieves a specific artifact version from a local repository, guaranteeing a consistent package from build to build.

Artifact Objects

ElectricFlow has three objects to support Artifact Management functionality: *artifact*, *artifact* version, and *repository*. Similar to other ElectricFlow objects, each of these objects supports custom properties and access control.

Artifact

An *artifact* is a top-level object containing artifact versions, a name template for *published* artifact versions, artifact specific properties, and access control entries to specify privileges.

If an artifact is deleted, all published artifact versions within that artifact will be deleted also.

You can create artifacts in several ways:

- Use the New Artifact web page
- Use the createArtifact API command with the command-line tool (ectool) or an ec-perl script
- An artifact is created when you publish an artifact version if the artifact did not already exist, and the user issuing the publish command has permission to create artifacts. See the "Access Control" and "Publishing Artifact Versions" sections below for more information.

To create an artifact, you enter a Group Id and Artifact Key and together, these create an artifact name in the form <code>groupId:artifactKey</code>. After ElectricFlow creates the artifact name from your supplied Group Id and Artifact Key specifications, you cannot modify the name.

Note: Plan your artifact Group IDs and Artifact Keys carefully to prevent the extra work of deleting a name and beginning your specifications again. The use of a groupId and artifactKey allow two "namespaces" for added flexibility in organizing artifacts within your company or organization.

Examples

- o If your development team uses GWT for web development and JUnit for unit testing, you may want to create artifacts "ThirdParty:GWT" and "ThirdParty:JUnit".
- Your organization may have two different development teams (OS and Apps) who want to use artifacts. In building each product, both development groups want to publish and use an "SDK" with their product. But each SDK is different and unique. One solution would be to create one artifact named "os:SDK" and another artifact named "apps:SDK". By using different group Id's, each development team is free to name their artifacts however they choose without worrying about name collisions with artifacts produced by the other team.

To reference an artifact, you can use the <code>groupId</code> and <code>artifactKey</code> combination or simply use the <code>artifactName</code>. The property path <code>/artifacts/<artifactName></code> allows you to access any properties of a given artifact.

Artifact Version

An *artifact version* is a collection of 0 to N files that were published to an artifact repository. Artifact version metadata is stored in an artifact Version object in the Electric Flow server.

Tips for working with artifact versions:

- To group a collection of files (typically from build output in an ElectricFlow workspace) into an artifact version, specify the "from directory" containing those files.
- If you want to include files from multiple directories in your artifact, you need to specify a common root directory to include these directories when publishing your artifact.

- To restrict which files and subdirectories in the "from directory" to include in the artifact version, you can specify "include" and "exclude" patterns.
- Artifact versions are stored on an artifact repository server in either uncompressed tar archives or compressed tar-gzip archives.
- Artifact versions can be created using the publishArtifactVersion API, which includes creating the object in the ElectricFlow server and publishing the artifact version to the artifact repository.

Note: Interactively creating an artifact version using the automation platform web UI is not supported at this time.

- When you retrieve an artifact version, the original directory structure of the directory where files were published "from" is preserved.
- When you retrieve/publish artifacts, directories containing symlinks are always resolved; they are not preserved.

Artifact versions are referenced by Group Id, Artifact Key (or artifact name) and a version string.

- The artifact version's name is set based on the name template on the "owning" artifact. By default, the name template is in the form <code>groupId:artifactKey:version</code>.
- Each artifact version name must be unique within an ElectricFlow server installation.

The property path /artifactVersions/<artifactVersionName> allows access to any properties for a particular artifact version. Because you may not know the artifact version name (because of the name template), this path supports the form <groupId>:<artifactKey>:<version> as a substitute for the name of the artifact version.

All ElectricFlow API's that take an artifactVersionName argument also accept this alternate form as a substitute, similar to how other API's that accept a jobId argument accept the job name as a substitute.

Version Strings

A version string must be provided when publishing an artifact version.

For example:

5.8.8-EN-55842 or 5.8.8.55842-EN - In either form, the version string is interpreted as:

Major version number: 5

Minor version number: 8

Patch level number: 8

Qualifier: EN

Build number: 55842

You must provide separator punctuation. When interpreting (parsing) the first version string form, ElectricFlow interprets the text after the first hyphen as the build number if this string contains numeric characters only. Otherwise, the string is interpreted as the qualifier.

Version string examples and how ElectricFlow interprets them:

	Input (version string)	Major	Minor	Patch	Qualifier	Build Number
1	2.0.3-44834	2	0	3	<empty nul<="" th=""><th>44834</th></empty>	44834
2	3.8.4-RC1	3	8	4	RC1	0
3	3.8.4-RC1-36	3	8	4	RC1	36
4	3.8.4.36	3	8	4	<empty nul<="" th=""><th>36</th></empty>	36
5	1-36	1	0	0	<pre><empty l="" nul=""></empty></pre>	36
6	3.8.4.36-RC1	3	8	4	RC1	36

In this table:

- Row 5 shows that you are not required to specify all three of <major, minor, patch> when
 publishing an artifact version if your organization's versioning conventions do not use all three fields.
 However, there is a caveat with this level of flexibility. You can create two artifact versions in the system
 with equivalent versions.
 - Version string "1-36" is equivalent to "1.0-36" but they are not the same so you can publish two artifact versions with these version strings successfully. When retrieving an artifact version, in some cases it could be unclear as to exactly which one you will get. For this reason, we recommend that you adopt a convention and abide by it for a particular artifact.

The artifact version object exposes components of the version string by using these intrinsic properties:

- version the version component is the combination of major.minor.patch-qualifier-buildNumber. A "version" component does not need to include all of the following intrinsic properties. See the table above.
- majorMinorPatch this is the major.minor.patch version component as specified in the version string. For example, if you specified a version string of "1-36, the majorMinorPatch component would be "1", not "1.0.0", which makes it possible for you to meet your naming conventions by reconstituting the version string a different way in the artifact version name template.
- qualifier the qualifier component of the version .
- buildNumber the build number component of the version.

The artifactVersionState property

Artifact versions have an artifact Version State property whose value can be one of the following:

publishing - The artifact version is currently being published and is not available for retrieval.

available - The artifact version is available for retrieval if needed.

unavailable - The artifact version is not available for retrieval while in this state.

You can manipulate the state between available and unavailable on the Edit Artifact Version web page. This action can be useful if an artifact version seems to be corrupt in some way, but you want to investigate before potentially deleting it from the system.

By making the artifact version unavailable, retrievers can acquire an older (potentially more stable) artifact version while you look into the problem. If the artifact version is good, you can simply make it available again.

Repository

The artifact repository is a machine where artifact versions are stored. The repository server is configured to store artifact versions in a directory referred to as the repository backing store.

By default, the backing store is the <datadir>/repository-data directory in the repository installation. This default setting can be changed by running ecconfigure --repositoryStorageDirectory on the repository server. The repository server listens on port 8200 for HTTPS requests to publish and retrieve artifact versions.

Connection information is stored in the repository object on the ElectricFlow server.

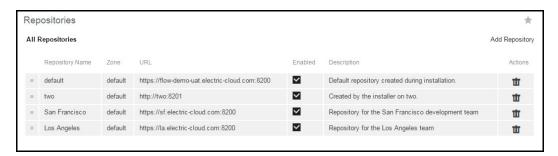
You can create a repository objects in several ways:

- Use the New Repository web page in the ElectricFlow UI
- Use createRepository API
- Repository objects can be created automatically during the ElectricFlow installation.

When you are installing a repository *server*, the installer creates a corresponding repository *object* on your ElectricFlow server. If you are installing a repository *server* on the same machine as your ElectricFlow server, a repository *object* named "default" is automatically created. If you are installing a repository *server* on a different machine than your ElectricFlow server, the installer will prompt you for information it uses to create the corresponding repository *object*.

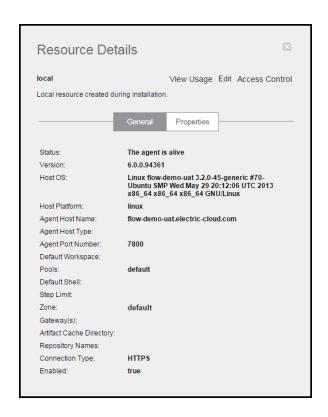
Note: When installing a repository server, we recommend installing an agent also to ease maintenance tasks such as clearing out stale artifact versions from the repository backing store.

Click the Artifacts tab, then the Repositories subtab to see your list of created repositories.



You can re-order the repository search order (for retrieving artifact versions) by using the mouse to grab and drag the icon in the left column to the row position you prefer, effectively re-ordering the repository list.

For distributed environments, where the preferred repository search order varies depending on which resource is performing the retrieval, you can specify a preference order on the Edit Resource panel. See the illustration below.

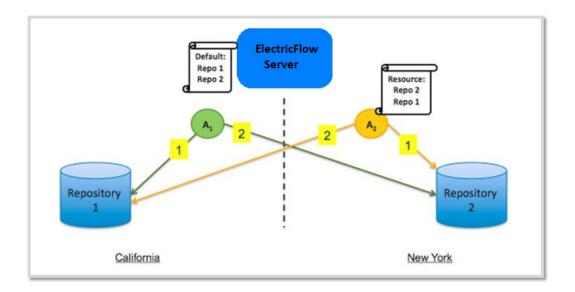


To go to the Edit Resource panel; from the Repositories page, select the Cloud tab, which opens to the Resources page, then select a resource name

Enter one repository name per line in the Repository Names field.

In the following example, a resource in California (A1) uses the "default" repository search order (that is, the order on the Repositories page). A1 searches for artifact versions by first requesting an artifact version from Repository 1. If the artifact version requested is not found, it will then request the artifact version from Repository 2

The resource in New York (A2) has a specific search order specified where it searches for artifacts from Repository 2, first, then searches Repository 1.



Using multiple repositories

ElectricFlow supports multiple repositories, however you cannot share a repository between multiple ElectricFlow servers. Multiple repository servers may be required in your organization because of business or organizational reasons, government and compliance requirements, or for performance and data protection.

Examples:

Business or Organizational Structure

Within a company, two different development teams want to have artifact repositories containing artifacts and artifact versions specific to their product. A repository sever is created for each team. The two teams publish artifact version to their respective repositories.

Performance

A company has multiple development sites. Instead of having a single repository server, which would create significant network performance issues from retrieving artifact versions during the build process, they create a repository at each development location and have artifact versions synchronized between the repository servers. In this way, when artifact versions are retrieved for use by a build process, they are retrieved from a local repository.

Data Protection and disaster recovery

A company has a single ElectricFlow server with two artifact repository server machines. The contents of the repository servers are synchronized and replicated between both machines. In the event repository 1 is unavailable during a retrieve, repository 2 is searched for the requested artifact retrieval.

Configuring the Artifact Repository Server for Amazon S3

You can configure the repository server to use Amazon S3 as the backing store and connect to Amazon S3 using IAM roles. This section describes the changes to make in the

/opt/electriccloud/electriccommander/conf/repository/server.properties file.

Before starting this procedure, make sure the following have been created and are available on the AWS Management Console (https://aws.amazon.com/console/):

- · An Amazon S3 bucket for the backing store.
- The IAM role with a policy that grants all the s3:* access to the Amazon S3 bucket you want to use with the respository server.

For example, the IAM role called s3repositoryFullAccess has a policy called s3fullaccessAndAccessToInstallers that grants all the s3:* access to the backing store bucket called ec-test-repository-backingstore.

Follow these steps to set up the artifact repository for Amazon S3:

REPOSITORY BACKING STORE=repository-data

 In the /opt/electriccloud/electriccommander/conf/repository/server.properties file, set REPOSITORY BACKING STORE to use Amazon S3 by changing

```
to
REPOSITORY BACKING STORE=s3://<name of the backing store bucket>
```

You can get the <name of the backing store bucket> by going to IAM > Roles > <name of the IAM role> page on the AWS Management Console. Then click Show Policy to open the Show Policy window to view the name of the backing store bucket.

Example:

In this example, <name of the backing store bucket is ec-test-repository-backingstore.

2. Restart the artifact repository server using a command like this:

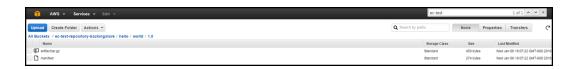
```
ubuntu@ip-10-101-214-28:~$ sudo /etc/init.d/commanderRepository restart
```

3. To test the REPOSITORY_BACKING_STORE, create two files in a /tmp/<your username> directory, publish the directory, and verify that it is published using commands like this:

```
ubuntu@ip-10-101-214-28:~$ cd /tmp
ubuntu@ip-10-101-214-28:/tmp$ mkdir <your username>
ubuntu@ip-10-101-214-28:/tmp$ cd <your username>
ubuntu@ip-10-101-214-28:/tmp/<your username>$ touch abc
ubuntu@ip-10-101-214-28:/tmp/<your username>$ ectool --server localhost login ad
min changeme
ubuntu@ip-10-101-214-28:/tmp/<your username>$ ectool publishArtifactVersion --ve
rsion 1.0 --artifactName hello:world
```

4. Verify that the /tmp/<your username> is published to the Amazon S3 backing store.

Example:



The respositoryDisabled intrinsic property

The repository object has a repositoryDisabled property that dictates whether or not artifacts can be published or retrieved from a particular repository. If you take a repository offline for maintenance, you can disable the repository in the ElectricFlow server rather than shutting down the repository machine or the service. Steps that attempt to publish to this repository will fail, and steps that attempt to retrieve artifact versions will skip this repository during the search for available repositories.

Access Control

You can apply access control permissions to artifacts and artifact versions to control who can create an artifact versus who can publish an artifact version versus who can read or retrieve an artifact version.

Action	Required permission	Object
Create an artifact	Modify	artifacts system object
Read artifact metadata	Read	the relevant artifact
Modify an artifact	Modify	the relevant artifact

Action	Required permission	Object
Delete an artifact, which deletes all artifact versions within the artifact	Modify	artifacts system object
Publish an artifact version	Modify	the artifact into which a new version is being published
Read artifact version metadata (required for retrieval)	Read	the relevant artifact version
Modify an artifact version	Modify	the relevant artifact version
Delete an artifact version	Modify	the artifact that owns the artifact version

Note: The Execute privilege does not take part in artifact or artifact version access control.

"Out-of-the-box," an <Everyone, modify> access control entry is available that allows anyone to publish an artifact version to any artifact, and allows anyone to create an artifact. You might consider removing this access control entry and setting up tighter control for individual artifacts. See the EC-Security plugin for help with this effort.

For more information about access control for artifact management, see the Access Control Help topic.

The following scenarios provide ideas for how you might set up access control for artifact management.

Control who can create new artifacts and publish artifact versions

ElectricFlow administrator, Adam, received communication from development manager, Max, requesting that only he should be able to create artifacts in ElectricFlow -- developers of ElectricFlow procedures under him should only be able to create steps to publish artifact versions under pre-defined artifacts.

His reasoning is that it is too easy for a developer to introduce a typo during publish (which does autocreate the artifact if the user has the appropriate privilege and the artifact doesn't exist).

Adam accommodates the request by giving Max the modify privilege on the artifacts system object and removing the <Everyone, modify> access control entry. Max then creates artifacts and assigns modify privileges to those artifacts for his development group, effectively allowing them to publish artifact versions.

At run-time, a step belonging to some project actually performs the publish, so Max adds project principal access control entries for the relevant projects as well, providing modify privileges.

Control which projects can retrieve artifact versions at run-time

Acme Corporation's router product has a software stack consisting of OS, PlatformLibraries, and Apps. Each of these has a ElectricFlow project with several procedures that work together to build the relevant component. PlatformLibraries builds against some of the C header files and libraries produced by the OS build. Similarly, Apps builds against libraries produced by PlatformLibraries.

The lead developer for the OS component, creates a Router:OSLib artifact and gives the PlatformLibraries project principal read privileges so that project can retrieve any OSLib artifact version it needs to build PlatformLibraries. Similarly, Paul from the PlatformLibraries team creates a Router:PlatformLib artifact and gives the Apps project principal read privileges.

Group 1 publishes an artifact version and wants only Group 2 to be able to retrieve it

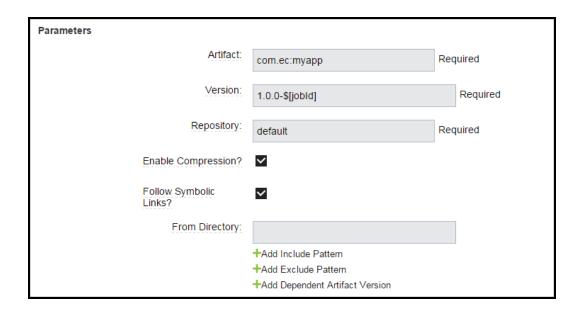
An organization wants to create artifacts for each of the third-party packages used by developers. Among the tools shared is the GWT package. The build manager chooses to create an artifact called <code>3rd-party-PKGS:GWT</code>. Because of licensing concerns, the build team must control or regulate who has access to the GWT and who can publish artifact versions used within the company's product.

In this case, the build team will publish a GWT artifact version and test it. During the test and approval cycles, only the build team has access to read and modify the artifact or publish new artifact versions. After a version has passed internal testing and approvals, the build team then grants read access to each of the developer groups.

Publishing Artifact Versions

An artifact version is published by using a procedure step or the *ectool / ec-perl* command-line interface. In most cases, using the Publish Artifact Version step is sufficient.

The following screen illustrates the parameters section for creating the Publish Artifact Version step:



Enable Compression

When publishing an artifact, consider whether or not to enable compression. Using compression reduces transfer time during publish. However, compression also adds overhead when computing the compressed data. If files included in the artifact version are primarily text files or are another highly compressible file format, the benefit of reduced transfer time outweighs the cost of computing compressed data.

Artifact versions that contain installers, jars, audio, and video are almost certainly **not** compressible because these file types are already compressed, so the cost of compressing outweighs the (near zero) benefit of reducing transfer time. Another consideration: Artifact versions are stored in the same format as they are transferred, so highly compressible artifact versions use less space in the repository backing store.

Include and Exclude Patterns

When publishing an artifact, you can choose which files to include by using include and exclude patterns. File patterns are expressed as relative paths under the From Directory. Pattern syntax and behavior is the same as Ant and uses the following wildcard characters:

- ? matches a single character
- * matches any number of characters, but only at a single directory level
- ** matches any number of directory levels

Examples:

```
Use *.jar to match any .jar file in the top-level directory.
```

Use */*.jar to match any .jar file in any child directory.

Use **/*.jar to match any .jar file at any level.

Retrieving Artifact Versions

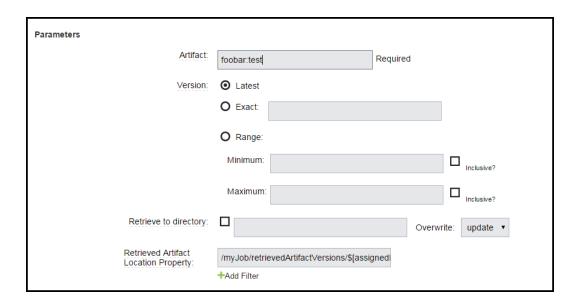
An artifact version is retrieved using a procedure step, or the retrieveArtifactVersion API call, used by the ectool/ec-perl command-line tools. This API call returns the most current artifact version that meets your criteria, per the algorithm described below.

Typically, you specify a version number range. Artifact version ranges are specified using interval notation. Brackets [] indicate versions to include and parentheses () indicate versions to exclude.

For example, the following artifact versions are published:

```
foobar:test:3.0.1-3645
foobar:test:1.2.3-2139
foobar:test:2.0.0-4000
foobar:test:2.0.0-DE-3395
foobar:test:2.0.0-DE-2445
foobar:test:2.0.0-DE
foobar:test:3.0.0-3539
foobar:test:4.0.0-5584
```

You want to retrieve the most current artifact version for the foobar: test artifact. To do this, you do not need to provide a version or version range, but instead enter retrieve parameters in the Retrieve Artifact Version step page as follows:



The equivalent command using ectool is:

```
ectool retrieveArtifactVersions --artifactName foobar:test
and with ec-perl, the syntax is:
    my $cmdr = newElectricCommander();
```

```
$cmdr->retrieveArtifactVersions({artifactName => "foobar:test"});
```

Version numbers for two artifact versions (Artifact Version 1 and Artifact Version 2) are compared in the server as follows:

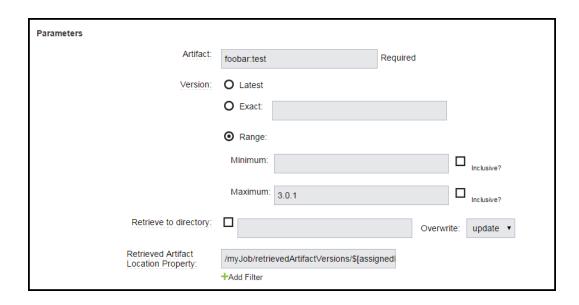
- 1. Compare major.minor.patch.
- 2. If they are equal, compare qualifiers.
- 3. If they are equal, compare buildNumbers.

Note: This comparison algorithm does not take the time of publish into account. "Current" means the artifact version with the greatest version per the above algorithm.

So the above list of artifact versions would be sorted (most current to least) as follows:

```
foobar:test:4.0.0-5584
foobar:test:3.0.1-3645
foobar:test:3.0.0-3539
foobar:test:2.0.0-DE-3395
foobar:test:2.0.0-DE-2445
foobar:test:2.0.0-DE
foobar:test:2.0.0-4000
foobar:test:1.2.3-2139
```

If you want to retrieve the latest 3.0.0 artifact version, set maximum version to 3.0.1 "exclusive". The parameters entered in the retrieve step would look like this:



Setting this constraint ensures the entire "version string" is properly evaluated. In other words, you want to ensure all artifact versions with 3.0.0, including qualifiers and build numbers, are included in the evaluation.

The equivalent command using ectool is:

```
ectool retrieveArtifactVersions --artifactName foobar:test
    --versionRange "(,3.0.1)"
```

And with ec-perl the syntax is:

Note the use of parentheses to exclude version "3.0.1".

Generally, while in development, you will want to specify a minimum artifact version and take the latest version published, then when an artifact has reached a "released" state, indicate a specific version.

Entering Filters During a Retrieve Operation

At times, a company might want to distinguish between artifact versions not just by version range but also by some other property of the artifact version.

Examples are:

• Specifying a qualifier in the version string when publishing.

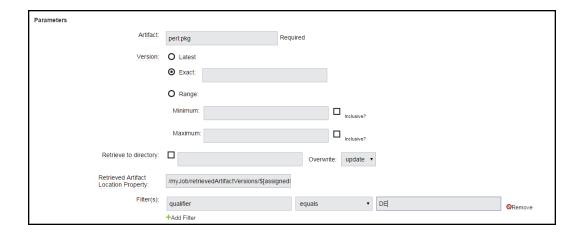
For example, you might want to publish an artifact version for your product's English or German variants. You could publish the English version as 2.0.4-EN-55497 and the German version as 2.0.4-DE-55497.

When the qualifier string is used this way to classify artifact versions, it should be used as a filter to retrieve the latest version. This is described in the following example.

Setting a custom property on the artifact version.

For example, after your QA team has tested an artifact version, they could show approval by setting a gaApproval custom property, including the name of the QA engineer.

For either case, filters can be applied when retrieving an artifact version. To retrieve the latest German version for the artifact, the parameters entered on the retrieve artifact custom step page look like this:



Because DE was used as the qualifier, ElectricFlow will retrieve the latest artifact version for the perl:pkg artifact where the qualifier equals DE.

Only the following filter operators can be applied to the intrinsic property version:

```
equals
greater than (greaterThan, for use in the API)
less than (lessThan, for use in the API)
greater than or equals (greaterOrEqual, for use in the API)
less than or equals (lessOrEqual, for use in the API)
```

"equals" is special because it does a string comparison on the version string of the artifact version, while the other operators compare the interpretation of the version string as described earlier.

All filter operators can be applied to version string components: majorMinorPatch, qualifier, or buildNumber. To find artifact versions with no qualifier, specify a filter on "qualifier" with operator "is not set" (or isNull in the API).

So for the QA approval case in a previous example, if a procedure is interested in an artifact version that was approved by QA, the retrieve step would specify a filter for "qaApproval is set".

ectool does not currently support the filters argument, but you can write a Perl script as follows:

```
my $cmdr = new ElectricCommander();
$cmdr->retrieveArtifactVersions({artifactName => "perl:pkg",
    filters => {propertyName => "qaApprover",
    operator => "isNotNull"}});
```

Dependent Artifact Versions

A published artifact version can be dependent on a list of artifact versions. These dependent artifact versions are retrieved when the primary artifact version is retrieved and they are specified with a query syntax when publishing the primary artifact version.

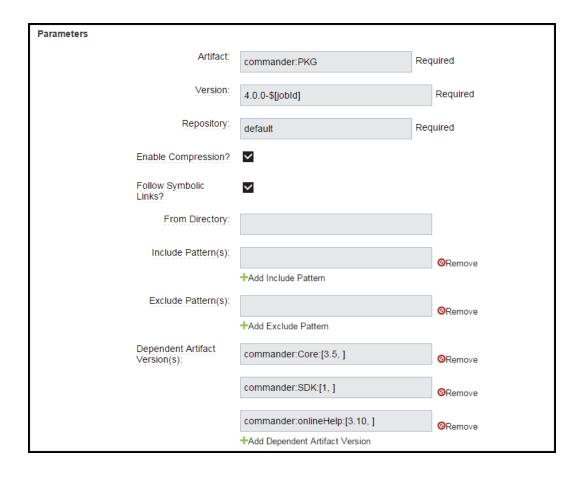
For example:

An artifact version for the ElectricFlow product could include dependent artifact versions for the core product, the SDK, and online Help:

- the latest commander: Core artifact version greater than or equal to version 3.5 (including 3.5)
- the most recent version of the commander: SDK artifacts
- the commander:onlineHelp artifact with a version greater than version 3.10 (excluding 3.10)

Because dependent artifact versions are evaluated during the retrieve process, specify them using the same syntax as you would for retrieving any artifact version, which is by including the group ld, the artifact key, and a specific artifact version or a version range.

Parameters for the "publish" custom step page would be similar to the following:



A published artifact version could have zero files, that is, no files are published as part of the artifact version, but it still could have a list of dependent artifact versions. An empty artifact version might be useful as modular convenience container for retrieving a group of artifact versions in downstream processes.

Artifact Cache

Artifact versions are retrieved to an artifact cache directory before being consumed by a build process. Properly configured, the artifact cache directory can significantly improve efficiency as you work with artifacts.

The artifact cache directory is set in one of these ways:

- Enter information in the Artifact Cache Directory field on the Edit Resource page.
- Use the modifyResource API to set the artifactCacheDirectory intrinsic property.
- Set the artifactCache field in the agent.conf file.

Thus, each resource can have its own artifact cache location (directory). Alternatively, multiple resources and an entire site can share an artifact cache if the cache directory resides on a shared file server.

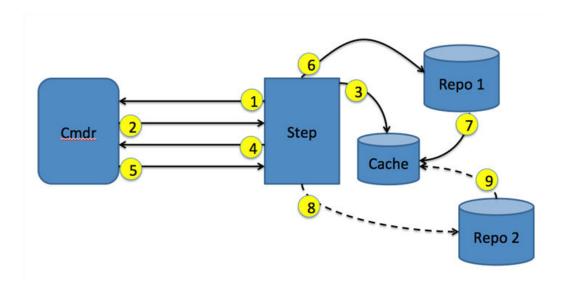
When an ElectricFlow step retrieves an artifact version, remember that often a version range or no version at all is specified. For example, to retrieve the latest ElectricFlow SDK artifact version, you could issue the following command:

ectool retrieveArtifactVersions --artifactName commander:SDK

No version is specified because we want the "latest" artifact version published.

In the following example, the step sends a request (1) to the ElectricFlow server.

- The ElectricFlow server replies (2) with a fully qualified artifact version name. In this case, it returns commander: SDK:1.2.0.43552.
- The step then looks in its artifact cache for the artifact version (3).
- If the artifact is not found in the artifact cache, a request is issued (4) to get a list of repositories to search to find the artifact version.
- The ElectricFlow server replies (5) with a list of repositories and the order in which to search them.
- The step requests the artifact version from the first (6) repository in the list.
- If the artifact version is found in the repository, the artifact version is retrieved to the cache (7).
- If the artifact version is not found, the step requests the artifact version from the next repository in the search order (8).
- If the artifact version is found in that repository, the artifact version is retrieved to the cache (9).
- Otherwise, if the artifact version is not found in any repository, an error is displayed.



In this example, if the requested artifact version was found in the artifact cache (step 3), the resource would not have communicated with either of the repository servers. Also, any potential network latencies would have been avoided from retrieving the artifact version from the repository server.

If you have a cache shared by multiple resources and steps that run on those resources need to use the same artifact version, the first step requesting an artifact version retrieves it from the repository to the cache. Subsequent job steps using the artifact version will have immediate access to the artifact version in the cache, eliminating the need for each step to retrieve its own copy of an artifact version.

Cleaning Up Repositories and Caches

Deleting an artifact version in the ElectricFlow server does not delete it in the repository backing store, and it does not remove the artifact version from an artifact cache.

ectool and the ElectricFlow Perl module include a function for each of these tasks: cleanupRepository and cleanupArtifactCache. Both of these API calls are intended to be run on the machine containing the cache/backingstore directory that needs to be cleaned up.

The cleanup algorithm is as follows:

- 1. Review the directory tree looking for directories three-levels deep.
- 2. For each such directory, interpret the three path components as <groupId, artifactKey, version>.
- 3. Issue a findObjects request to the ElectricFlow server to see if an artifact version with that specification exists.
- 4. Delete all directories identified by <groupId, artifactKey, version> and not recognized by the ElectricFlow server.

To operate correctly, using the cleanupRepository API requires read access to all artifact version objects stored in ElectricFlow and therefore must be run in an "admin" session. We also recommend running cleanupArtifactCache in an "admin" session—if not, artifact versions that are simply not readable by the user invoking the cleanup will be deleted from the cache.

In addition, both functions fail if the server does not recognize any of the chosen artifact versions. Thus, if the cleanup function is called on a directory that is not a backingstore directory or and artifact cache, it will not delete anything.

The recommended approach for removing stale artifact versions from a backing store is to use an ElectricFlow maintenance procedure that periodically runs on each repository server (which presumably has an agent install). Similarly, use a procedure with a broadcast step to run on each agent to clean up its cache. Note that the job step session does not necessarily have read privileges on all artifact versions, so it is important to login as a privileged user, then perform cleanup with those credentials.

Authenticating Users for LDAP and Active Directory

ElectricFlow uses account information from multiple sources. In most cases, the primary account information source is an external LDAP or Active Directory repository: both user and group information is retrieved from the repository. *Local* users and groups are defined within ElectricFlow.

Use the following LDAP information and examples to enter information in your LDAP template. If you are using the Active Directory template, you may still use the LDAP information as reference—the templates are similar with only a slight difference in the Active Directory template.

Configuring LDAP

A number of options need to be customized for any LDAP configuration. The following sample configuration shows how we have our own LDAP configuration set up. After the sample, see a list of properties with a description.

Sample LDAP Configuration

```
<bean id="LdapDirectoryProvider"</pre>
  class="com.electriccloud.security.ldap.
  LdapDirectoryProviderImpl">
  cproperty name="providerName" value="LDAP"/>
  property name="url" value="ldap://dir.
     electric-cloud.com/dc=electric-cloud,dc=com"/>
  property name="managerDn" value="uid=JohnDoe,ou=People,=
     electric-cloud, dc=com"/>
  cproperty name="managerPassword" value="****"/>
  cproperty name="userBase" value="ou=People"/>
  cproperty name="userSearchFilter" value="uid={0}"/>
  property name="userNameAttribute" value="uid"/>
  cproperty name="fullUserNameAttribute" value="gecos"/>
  cproperty name="emailAttribute" value="mail"/>
  cproperty name="groupBase" value="ou=Group"/>
  cproperty name="groupMemberFilter" value="(|(member={0}))
      (memberUid={1}))''/>
  property name="groupMemberAttributes" value="member,
     memberUid"/>
  cproperty name="groupSearchFilter" value="()(
     groupOfNames) (objectClass=posixGroup))"/>
  cproperty name="groupNameAttribute" value="cn"/>
</bean>
```

The following properties configure LDAP mapping:

emailAttribute - (optional) The attribute in a user record that contains the user's email address. If the attribute is not specified, the account name and domain name are concatenated to form an email address.

fullUserNameAttribute - (optional) The attribute in a user record that contains the user's full name (first and last) for display in the UI. If this attribute is not specified or the resulting value is empty, the user's account name is used instead.

managerDn - (optional) The DN of a user who has read-only access to LDAP user and group directories. If this property is not specified, the server attempts to connect as an unauthenticated user. Not all servers allow anonymous read-only access.

Note: This user does not need to be an admin user with modify privileges.

managerPassword - (optional) If the managerDn property is set, this password is used to authenticate the manager user.

providerName - This human readable name will be displayed in the user interface to identify users and groups that come from this provider.

userBase - This string is prepended to the basedn to construct the directory DN that contains user records.

userNameAttribute - The attribute in a user record that contains the user's account name.

userSearchFilter - This LDAP query is performed in the context of the user directory to search for a user by account name. The string "r; {0}" is replaced with the user's login ID. Typically, the query compares a user record attribute with the substituted user login ID.

url - The LDAP server URL is in the form <code>protocol://host:port/basedn</code>. Protocol is either <code>ldap</code> or <code>ldaps</code> for secure LDAP). The port is implied by the protocol, but can be overridden if it is not at the default location (389 for Idap, 636 for Idaps). The <code>basedn</code> is the path to the top level directory that contains users and groups at this site. This is typically the domain name where each part is listed with a <code>dc=</code> and separated by commas.

Note: Spaces in the basedn must be URL encoded (%20).

In addition to user information, the LDAP server can be queried for group information. This query is optional because the local group mechanism can refer to LDAP users and *local* users. However, the following elements can be used to tell the server how to map groups in LDAP

groupBase - (optional) This string is prepended to the basedn to construct the directory DN that contains group records.

groupMemberAttributes - (optional) A comma separated attribute names list that identifies a group member. Most LDAP configurations only specify a single value, but if there is a mixture of POSIX and LDAP style groups in the directory, multiple attributes might be required.

groupMemberFilter - (optional) This LDAP query is performed in the groups directory context to identify groups that contain a specific user as a member. There are two common forms of group record in LDAP directories: POSIX style groups where members are identified by account name, and groupOfNames or uniqueGroupOfNames records where members are identified by the full user DN. Both forms are supported, so the query is passed two parameters: "r; {0}" is replaced with the full user record DN, and "r; {1}" is replaced with the user's account name.

groupNameAttribute - (optional) The group record attribute that contains the name of the group.

groupSearchFilter - (optional) This LDAP query is performed in the context of the groups directory to enumerate group records.

Determining LDAP Mapping

A typical POSIX user record in LDAP looks similar the example below. To set up a mapping for this record, it is necessary to identify various record components. First, identify the path in the directory that contains user records. In this example, the build user has a distinguished name (dn) of

uid=build, ou=People, dc=mycompany, dc=com. This name uniquely identifies the build user account and this path splits into three parts:

```
base dn: dc=mycompany, dc=com
user base: ou=People
user element: uid=build
```

The baseDn is the parent of the directory that contains users. This value should be combined with the protocol and server to form the URL. In this case, the URL is ldaps://dir/dc=mycompany, dc=com.

Next, the userBase is the portion of the path that identifies the directory containing all user account records. This value is used directly as the userBase configuration element.

The remaining portion identifies the user without the People directory: uid=build. The user name is replaced in this value with the string "r; {0}" to form the userSearchFilter: uid={0}. This query allows the server to search for a user's account name by looking for a record with a matching uid attribute.

The final mapping step is to identify user record attributes that hold the account name, full user name, and (optionally) the user's email address. In this example, the account name is uid (identified earlier), the full user name attribute is gecos, and there is no email attribute.

At this point, the server is able to authenticate a user, look up a user by name, and determine the user's full name. For many installations this is sufficient.

Sample LDAP User Record

```
# build, People, electric-cloud.com
dn: uid=jdoe, ou=People, dc=mycompany,dc=com
loginShell: /bin/bash
uidNumber: 508
gidNumber: 508
objectClass: account
objectClass: posixAccount
objectClass: top
objectClass: shadowAccount
uid: jdoe
gecos: John Doe
cn: John
homeDirectory: /net/filer/homes/build
```

Also, you can configure the server to look for **LDAP** groups that refer to user accounts. A typical group record is shown below. Like a user account, an LDAP group has a distinguished name with a baseDn, a group base, and a group element. In this case, the basedn is still dc=mycompany, dc=com. The groupBase configuration element is ou=Group, and the group name is cn=build_users.

The server needs to identify records in the directory that correspond to groups—it does this by applying the <code>groupMemberFilter</code> to the records in the <code>groupBase</code> directory. In this case, group records are members of the <code>posixAccount</code> object class, so the filter can be set to <code>objectClass=posixGroup</code>. To display a group by its name, the server needs to know which attribute represents the group name. In this case, set the <code>groupNameAttribute</code> to <code>cn</code>.

Finally, the server needs a filter to determine which accounts belong to the group and the attribute name that represents a single group member. Group membership can be identified in one of two ways. Either the members are listed by account name, or by their LDAP distinguished name. In this case, POSIX group membership is determined by account name, so set the groupMemberAttributes property to memberUid, and set the groupMemberFilter to memberUid={1}.

Sample LDAP Group Record

```
# build_users, Group, mycompany.com
dn: cn=build_users,ou=Groups,dc=mycompany,dc=com
objectClass: posixGroup
objectClass: top
gidNumber: 100
memberUid: jdoe
memberUid: mary
cn: build users
```

Sample Active Directory Configuration File

The following XML file defines parameters needed to connect to an Active Directory server and the query to use for looking up user information.

```
<beans xmlns="http://www.springframework.org/schema/beans"</pre>
   xmlns:
  xmlns:tx
  xmlns:aop
  xsi:schemaLocation
  http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
  http://www.springframework.org/schema/tx
  http://www.springframework.org/schema/tx/spring-tx-2.0.xsd
  http://www.springframework.org/schema/aop
  http://www.springframework.org/schema/aop/spring-aop-2.0.xsd" default-lazy-
   init="true"
     <bean id="ADDirectoryProvider"</pre>
     class="com.electriccloud.security.ldap.ActiveDirectoryProviderImpl">
     property name="providerName" value="ActiveDirectory" />
     <!-- START OF CUSTOMIZATIONS -->
     <!-- URL to AD server of the form: "r;ldap://host:port/base dn" -->
     cproperty name="url" value="ldap://server:389/dc=company,dc=com" />
     <!-- Required credentials for an account that has read-only access to the server
     -->
     property name="managerDn"
     value="cn=myuser,cn=Users,dc=company,dc=com" />
     property name="managerPassword" value="mypw" />
```

```
cproperty name="userBase" value="cn=Users" />
     (objectClass=user))" />
     cproperty name="userNameAttribute" value="sAMAccountName" />
     cproperty name="fullUserNameAttribute" value="name" />
     <!-- Optional group configuration. This is only needed if you intend -->
     <!-- to manage groups externally via AD. -->
     <!-- <pre><!-- <pre>cycle="groupBase" value="value=""/>-->
     <!-- <pre><!-- <pre>cyroperty name="groupMemberFilter" value="member={0}"/> -->
     <!-- <pre><!-- <pre>cyroperty name="groupMemberAttributes" value="member"/>-->
     <!-- <pre><!-- <pre><!-- <pre>property name="groupSearchFilter" value="(objectClass=group)"/>-->
     <!-- <pre><!-- <pre>coupNameAttribute" value="cn"/>-->
     <!-- Max number of results to get back in one query. This must be -->
     <!-- no greater than the AD server's setting, which is typically 1000. -->
     cproperty name="pageSize" value="500" />
     <!-- END OF CUSTOMIZATIONS -->
     </bean>
</beans>
```

Change Tracking

Change Tracking is designed to track every change between every state of non-runtime ElectricFlow objects and to allow you to revert to any previous state of these objects. ElectricFlow tracks the changes to tracked objects including applications, procedures, workflows, workspaces, resources, and project-owned components such as library components and records a *Change History* of the historical states of the system and the changes between them. The change history has a copy of every state in which every non-runtime object has been, the object's current state, and indexing records for searching through the database. The tracked objects are not affected when ElectricFlow executes an object that is usually created or modified at execution time, such as a workflow or process.

Change tracking allows you to do the following:

- When debugging a failed job or looking for more information about a component in an ElectricFlow, see the Change History for the changes relevant to that object.
- When searching for specific change history records, filter the records by time frame, change type, entity type, or developer.
- Revert changes to an object or to an objects and its children.
- When you want to determine the differences between objects, export them at various levels in the object hierarchy.

In ElectricFlow, you can use Change Tracking in all aspects of the ElectricFlow end-to-end solution:

- In build-test automation, track the Change History of artifacts, resources, and ACLs.
- In deployment automation, track the Change History of components, environment tiers, and process steps (application or component). You also use Change Tracking with snapshots to make it easier to deploy reliable and repeatable software for Continuous Delivery.
- In Releases and Pipelines, track the Change History of stages and tasks.

Configuring Change Tracking

Change Tracking must be enabled when ElectricFlow starts for your system to track changes and record the Change History.

By default, Change Tracking is enabled for projects created in ElectricFlow 5.3 or later. It is disabled for projects created in ElectricFlow versions before ElectricFlow 5.3, but can be enabled manually.

Enabling Change Tracking Globally

When installing ElectricFlow:

1. Add this line to the database.properties file:

```
COMMANDER_DB_AUDITING_ENABLED=true
```

2. Restart the ElectricFlow server.

Enabling Change Tracking on a Per-Project Basis

IMPORTANT: Unless the allowEnablingDisablingChangeTrackingProjectLevel setting has been altered, only admin users are allowed to perform this task. The default setting is adminOnly. If the setting is changed to anyoneWithAccess, any user with access to the project can enable or disable Change Tracking for the project

You can enable Change Tracking one of the following ways:

In the ElectricFlow Platform UI

Change Tracking is enabled when the **Enable Change Tracking** check box is selected.

New Project		
Name:		
Enable Change Tracking:	✓	
Description:		
Default Resource:		Browse
Default Workspace:		Browse

To disable Change Tracking, click the Enable Change Tracking check box to clear it. and click OK.

Using ectool

- $\bullet \quad \textbf{Enter} \ \texttt{ectool} \ \ \texttt{modifyProject} \ \ \texttt{<projectName} > \ -\texttt{-tracked} \ \ \texttt{false} \ \textbf{to} \ \textbf{disable} \ \textbf{Change} \ \textbf{Tracking}.$

Using ec-perl:

- Enter \$cmdr->modifyProject(<projectName>, {tracked => true}); to enable Change Tracking.
- Enter \$cmdr->modifyProject(<projectName>, {tracked => false}); to disable Change Tracking.

For properties that are frequently updated always to be a numerical value, it is possible to suppress Change Tracking of numerical-value updates by using one of these commands:

```
ectool:ectool modifyProperty cprojectName> -- path cpropertyPath> -- counter true

or
ec-perl:$cmdr-> modifyProperty (cprojectName>, {path => propertyPath>, counter => true});
```

When the counter flag is set for a property, Change Tracking does not track changes to the property if the only change was a change in the property value from one numeric value to another. All other forms of changes to the properties that have this flag set are tracked normally.

CAUTION: Reverting an object that owns a counter property to a previous state will revert the value of the counter property to its value at the previous time that a change to this property was tracked—typically this is

when its value is initialized. This may not be the desired behavior, so you may need to manually set a counter property if it is reverted,.

Upgrading to ElectricFlow 6.x

Change Tracking is enabled when you upgrade to ElectricFlow 6.x. This can significantly increase the time it takes to complete the upgrade.

If you want to upgrade with Change Tracking disabled, add this line to the database.properties file before starting the upgrade:

COMMANDER_DB_AUDITING_ENABLED=false

Customizing the Change History Page

The performance of the Change Tracking feature is affected by number of records in the Change History as well as the number of entries being tracked. For example, a page showing 5000 entries may be slow to load and update and does not provide much useful information.

This ElectricFlow server uses the lowest of the following limits to determine the maximum amount of records to display in the Change History page:

· Maximum amount of records on the Change History page

To change the maximum number of records in the Change History page:

1. Set the CHANGE_TRACKING_HARD_MAX_RECORDS parameter in the wrapper.conf file to a new value.

The default value is 1000.

- 2. Restart the ElectricFlow server.
- · Maximum number of records retreived

Set the TrackingMax Records server setting to a new value not exceeding the CHANGE_TRACKING_HARD_MAX_RECORDS parameter in the wrapper.conf file.

To set TrackingMaxRecords, do one of the following:

Change the value in the UI. See the Server Settings page in the automation platform UI.

Use ectool to change the value. For example, enter the following command to limit the number of records retrieved to 100:

ectool setProperty /server/settings/changeTrackingMaxRecords --value 100

Viewing the Change History for Artifacts, Jobs, Projects, and Workflows

When troubleshooting why a job failed, you can view the Change History for artifacts, jobs, projects, and workflows in the automation platform.

- Artifacts on page 734
- Jobs on page 734

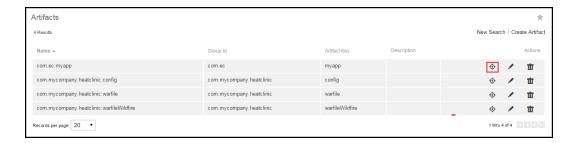
- Projects on page 735
- Workflows on page 736

Artifacts

Starting from the Home page:

- 1. Go to the Artifacts tab.
- 2. Choose an artifact.
- 3. Click the Track Changes button.

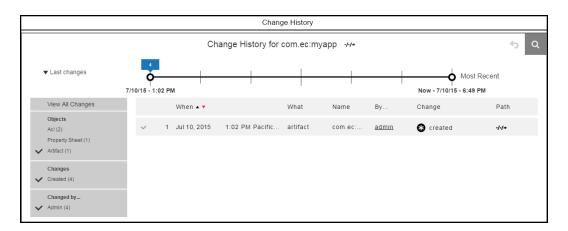
Example:



The Change History for the selected artifact opens.

The default time increment is Past 60 Minutes.

Example:



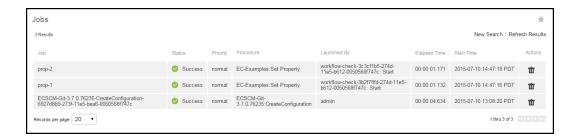
Jobs

Starting from the Home page:

- 1. To go to the Job Details page, do one of the following:
 - Use the Jobs tab.
 - · Use the Jobs Quick View list.
- 2. If you use the Jobs tab, follow these steps:
 - a. Click the Jobs tab.

The Jobs page opens.

Example:



b. Click a job name to select a job.

The Job Details page opens.

Example:



3. If you use the Jobs Quick View list, click a job name to select a job.

The Job Details page opens.

- 4. Choose a job or job step.
- 5. Click Track Changes for the job or job step.

The Change History for the job or job step opens.

The default time increment is Past 60 Minutes.

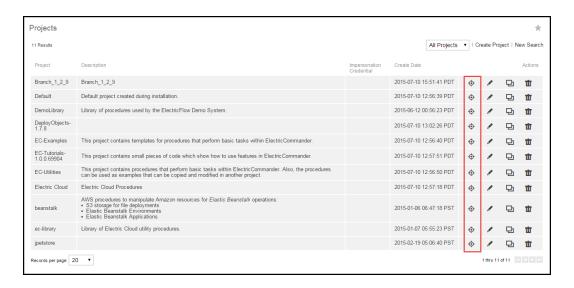
Projects

Starting from the Home page:

- 1. Go to the Projects tab.
- 2. Choose a project.
- 3. Click the Track Changes button.

The Change History for the project opens.

Example:



Workflows

Starting from the Home page:

- 1. Go to the Workflows tab.
- 2. Choose a workflow.
- 3. Click Track Changes.

The Change History for the workflow opens.

Modifying What You See the Change History

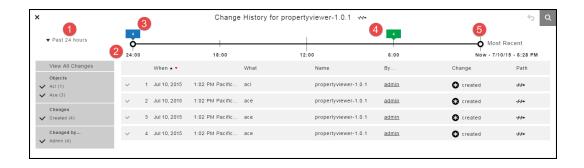
You can modify the information that appears in the Change History with these settings:

- Time line-Go to Change History Time Line on page 736.
- Filters-Go to Change History Filters on page 742.

Change History Time Line

The time line is at the top of the Change History page.

This example shows the Change History for a property called *propertyviewer-1.0.1*.



	Time increment.
1	The default is Last Changes.
	Click the down arrow to select another time increment.
	The system automatically calculates the minutes, hours, and days since the last successful run.
2	In the example, the last successful run occurred 24 hours ago. The time line is divided into four 6 hour subdivisions.
3	Total number of changes in the selected time increment.
	The number of changes that occurred between 6 hours ago and now is 4.
4	When you click the change number, the Change History is updated and shows only those changes.
5	Drag the start and end time markers to view specific changes.

Default Settings

The default time increment is **Last Change**.

The entire time line is displayed, and all the changes are in the list below the time line.

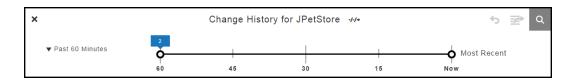
Number of Changes

The time line shows the number of changes throughout the time increment. In the following example:

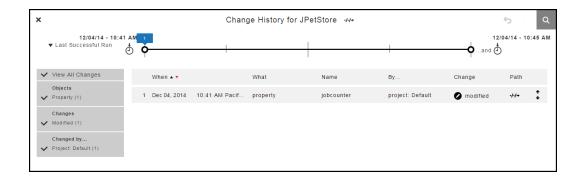
- There have been 233 changes in the last 24 hours.
- There have been 213 changes in the last 12 hours.
- There have been 20 changes in the last 6 hours.



When you change the time increment, there have been two changes in the previous 60 minutes.



Time Increment



The system automatically determines how the time line is divided for the selected time increment.

When the range is changed to **Past 60 Minutes**, the time line changes:

- The start time is 60 minutes from Now.
- The end time is when the **Most Recent** change occurred (**Now**).
- The time line has four divisions.



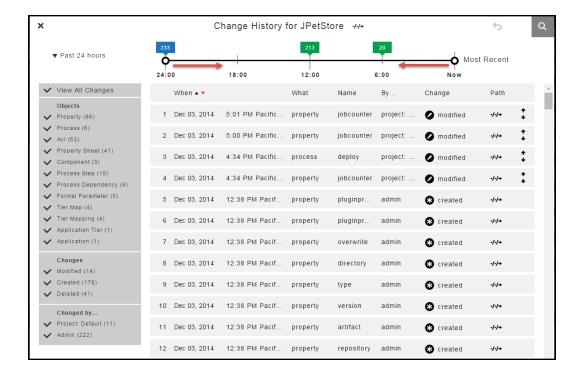
If the increment is Past 7 Days, the time line has seven one-day divisions.



Moving the Start and End Times Manually

You can manually move the start and end times on the time line.

When you move the start time to 18:00 and the end time to 6:00, the list of objects in the change history changes.



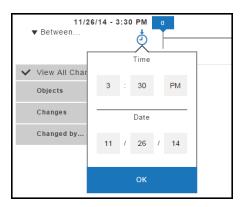
Setting Custom Time Increments



1. Select Between.

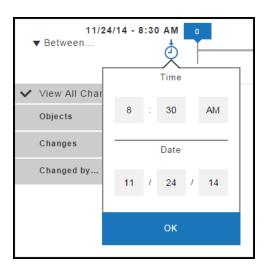
A drop down dialog box opens.

Example:



2. Select the time and date for the start of the time line.

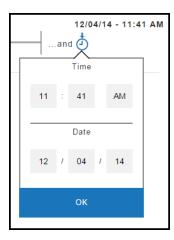
The default settings are **3:30 PM** and eight days before the current date.



3. Click OK.

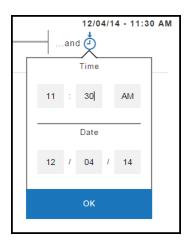
A drop down dialog box opens at the other end of the time line.

Example:



4. Select the time and date for the end of the time line.

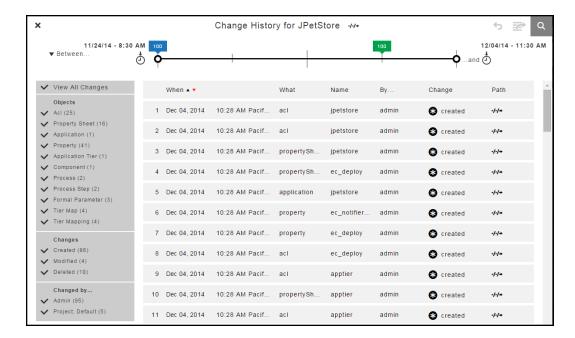
The defaults are 3:30 PM and the current date.



5. Click OK.

The time line changes to show only the changes from the start and end times and dates that you selected.

Example:



Change History Filters

The Change History has records only for an object and its tracked entities, not all entities belonging to the object. You can use filters to view changes to specific objects, the types of changes, and the users how made those changes if they are are tracked.

In this example, instead of selecting **View All Changes**, you can select specific objects, such as only properties, processes, property sheets, process steps, and process dependencies that have been modified by the Project:Default and Admin users.



When the list of filter criteria is long, not all of the criteria may appear in the filter list. To see all of the criteria, use the up or down arrows to see all the options.

This list does not show all of the users. Use the up and down arrows to see all four of the users.



Click the down arrow to see the other users.



Searching the Change History

Follow these steps to start a search in the Change History.

You can start a Change History search from most pages in the ElectricFlow UI.

1. Click the **Search** button or click **Change History**.

Example:



The **Change History - Search** dialog box opens.

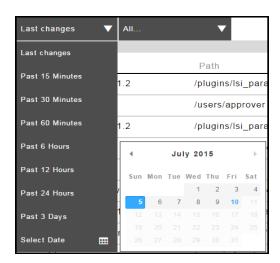


1	Time range field. Click the down arrow to open the drop-down list of start times. The end time is the current time.	
2	Objects field. Click the down arrow to open the drop-down list of objects to include in the search. You can select All or specific objects. By default, seven of the most commonly tracked objects are selected.	
3	Search criteria. After you type, the system starts searching for objects based on the time range and objects that you selected. The search results are in the Change History.	

- 2. To select a time range for the search:
 - a. Click the down arrow in the Time range field to open the drop-down list.
 - b. Select a time range.
 - c. If you want to use a time increment longer than three days, do the following:
 - a. Click Select Date.

The Date Picker opens.

Example:



b. Select a date.

The Date Picker closes and the date that you selected appears in the Time Increment field.

Example:



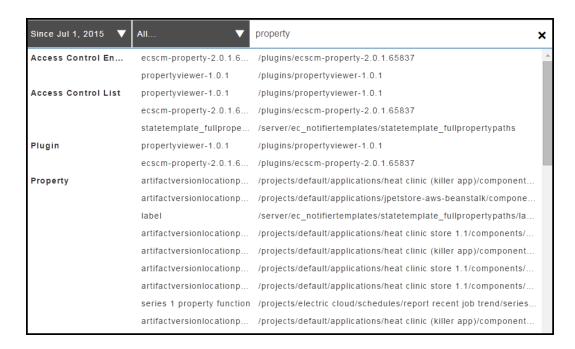
- 3. To select an objects for the search:
 - a. Click the down arrow in the Object field to open the drop-down list.
 - b. Select the objects for the search.
- 4. Enter the search criteria.

As you type, the system starts to search for objects that match your search criteria.

A list of objects matching your search criteria appears in the results section.

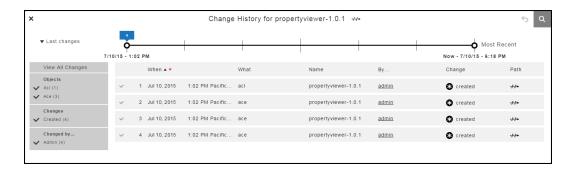
5. Select an object in the list.

Example:



The change history for the object that you selected appears.

Example:



Reverting Changes to a Tracked Object and Its Tracked Contents

The Change History displays a list of changes in reverse chronological order (from latest to oldest) for tracked entities in objects such as applications, procedures, workflows, workspaces, resources, and project-owned components such as library components. Each row in the Change History or a tracked object is a record of a change made either to the tracked object or to a tracked object owned by the object (directly or indirectly). For example, the Change History of a project would show both changes to the project itself and changes to any tracked object in the project, such as its procedures, the 'steps for each procedure, or properties belonging to the procedure steps.

Some types of objects (mostly run-time objects such as jobs, because tracking those would be resource-intensive and seldom useful) are not tracked by the Change Tracking feature. For a few types of tracked objects, some of their attributes are not tracked (for example, agents are tracked, but certain attributes of agents that vary frequently in real-time, such as their current status, are not). Changes to untracked objects, or to untracked attributes of tracked objects, are not shown in the Change History and cannot be reverted. In general, objects that are normally created or changed manually are tracked, and objects that are normally created or modified during automated processing are not, However, depending on your specific usage patterns, this behavior may not be recorded in the Change History, such as when you frequently run jobs that use ectool or API commands to automate the creation of new procedures or steps.

Follow these steps to select a tracked object and its tracked entities that you want to revert:

1. Go to the Change History for the object that you want to revert using one of these methods:



• Click the **Track Changes** button ()for the object.



- From the ElectricFlow Home page, click the Search button () to go to the "Change History Search" page and search for the object.
- 2. Locate the time period in which you are interested.

A list of changes for the tracked object appears.

- 3. If you decide to revert the changes for an object owned by the tracked object or for an object that owns the tracked object, use one of the methods to view the list of changes for this object:
 - Close the current Change History and navigate to the Change History for the object.



Click the **View Path** button () on the top of the current Change History, copy the path, go to the "Change History - Search" page, and search for a suitable addition to or truncation of that path.

4. Select a row (record) in the Change History corresponding to the earliest of the changes that you want to undo.

You will be reverting the object and all its tracked contents to their state directly before the selected change was made.



a. If the **Revert** button () becomes enabled, ElectricFlow is capable of reverting the selected changes to the object. is tracking the changes to this object.

Go to Step 5.

b. If the Revert button is still disabled, ElectricFlow is not currently capable of reverting changes to the selected object. This could be due to tracking of this object being disabled and re-enabled, or the Change History feature being disabled and re-enabled. Also, not all types of tracked objects can currently be reverted (and due to mutual consistency requirements in the database format, certain types of objects cannot be reverted).

If you are unable to revert the selected object, you may still be able to achieve what you need to do one of these ways:

- Revert the parent object that owns the object you want to revert.
- Separately revert the individual objects owned by the selected object or even nested partial-reversions-of-reversions, possibly combined with manually reversion (undoing) certain changes.

Go to Step 3 to search for an for an owned or owning tracked object.

5. Before clicking the **Revert** button, make certain that you know all the changes that will be reverted, and have confirmed that this is what you want to do. (While it is possible to re-revert the changes caused by a mistaken reversion operation, doing so leaves a lot of activity in the Change History.)

IMPORTANT: All of the changes for tracked entities will be reverted, even the ones that you do not see in the Change History because of the timeline and filter settings.

- a. You may need to modify the timeline and filter settings to display all records in the Change History before you verify that all of the records are changes you want to revert.
- b. Make sure that the end of the timeline is set to Now and that the filters are set to View All Changes.
- c. Look through the entire list of changes from the most recent to the one you have selected, expanding them to see the details if needed, and confirm that you really want to revert all of them.

IMPORTANT: Tip on how to further refine what changes will be reverted

If you want to revert some of the changes shown but not others, you may be able to achieve what you want to do with some combination of the following methods:

- Reverting just certain objects owned by the object that you want to revert.
- Using nested or overlapping sets of reverts of owned objects.
- Undoing individual changes manually using the ElectricFlow UI.
- Using nested partial-reverts-of-reverts.

If you fully understand the XML format used for ElectricFlow exports, you can use the <code>export</code> API call to export the past and current states of a part of the object hierarchy, diff and merge them using a suitable tool for diffing and merging XML files, and then re-import the results. Caution is strongly suggested if you are using this approach to revert objects for which ElectricFlow does not currently support reversion, because:

- There are some places in the ElectricFlow XML data hierarchy where peer objects refer to each other in ways that need to be kept consistent (such as by index of current position in a list).
- For some types of objects, re-importing them can change their IDs and break references-by-id.
- 6. Click the **Revert** button. Depending on the size of the object and every tracked object it owns, this operation could take some time (as long as exporting a past state of them and then re-importing it).

If this operation is successful, the object and its tracked entities are reverted back to the state at the date and time of the record selected in Step 3. A message appears that the Revert operation was successful.

If this operation is not successful, the object and its tracked entities are not reverted, and a message appears about the failed operation.

Change History Search Form

How to get here: From most pages, click the Search button to open the "Change History - Search" form.



The "Change History - Search" form has the following information:



1	Time range field. Click the down arrow to open the drop-down list of start times. The end time is the current time.
2	Objects field. Click the down arrow to open the drop-down list of objects to include in the search. You can select All or specific objects. By default, seven of the most commonly tracked objects are selected.
3	Search criteria. After you type, the system starts searching for objects based on the time range and objects that you selected. The search results are in the Change History.

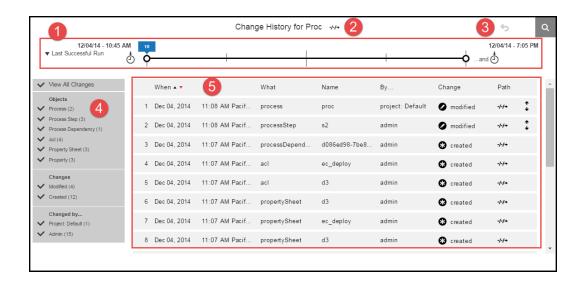
Change History Page

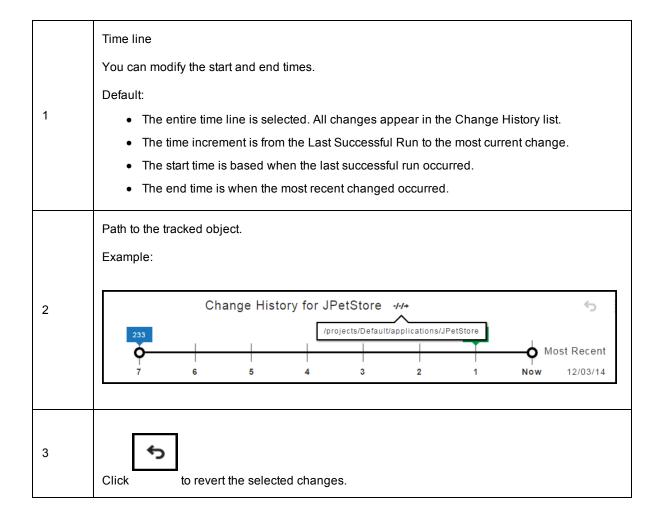
How to get here: Click the Change History button for a tracked object.

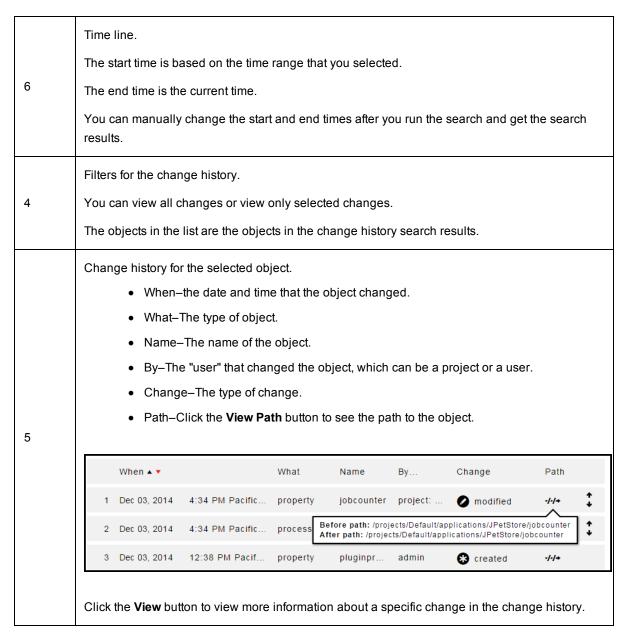
Example:



The Change History has this information about the object called *Proc*:

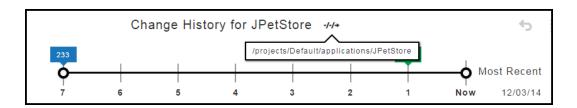






Paths to Objects

Click the View Path button next to the "Change History for JPetStore" title to see the path to the application.

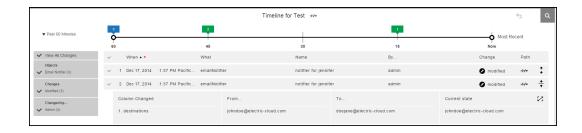


Click the View Path button to see the change in the path to the object before and after the change.



Detailed Object Changes

Click the View button to see the change in the property called emailNotifier.



Click the **Expand** button to all the changes to the property,

Example:



When you click the **Expand** button in a cell, you can see more details in the current change in the cell.

If you click the Select All button, all the changes about the object appear.



To select an object to revert or import the changes to an XML file, select the row of the object in the expanded view.



Credentials and User Impersonation

What is a Credential? on page 754

Defining a Credential on page 754

Why Use Credentials? on page 754

Credential Access Control on page 755

Using Credentials for Impersonation at the Platform Level on page 755

Accessing Credentials From a Step at the Platform Level on page 757

Credential References on page 758

Best Practices for Retrieving Credentials on page 759

What is a Credential?

In ElectricFlow, a "credential" is an object that stores a user name and password for later use. Two credential types are available, *stored* or *dynamic*:

- stored credentials These credentials are given a name and stored in encrypted form in the database.
 Each Project has a list of stored credentials it owns. These credentials are managed from the Project Details page.
- **dynamic** credentials These credentials are captured when a job is created. Dynamic credentials are stored on the server temporarily until the job completes and then discarded.

Defining a Credential

After a credential is created, no one can view the password for the credential's account. This means one person can define a credential and enter the password, and other people can use the credential (and its account) without needing to know the password.

To create a new credential in the automation platform UI:

- 1. Click the Projects tab.
- 2. Select a project (first column) to see the Project Details page.
- 3. Select the Credentials tab and then click Create Credential.

Click the **Help** link on the New Credential page if you need more information about what to enter in the fields.

Why Use Credentials?

Use ElectricFlow credentials for:

- User impersonation
- · Saving passwords for use inside steps

When a step needs to run as a particular user, ElectricFlow can retrieve the user name and password from a stored credential. The credential is passed to the agent over an encrypted channel so the agent can authenticate itself to the operating system and set up a security context where the step runs with the user permissions in the credential.

In some cases, a step needs to enter credentials directly to an application.

- The credential may be a fixed application credential that needs to be presented every time the step runs.
 In this case, a stored credential provides a place to put this information without needing to embed the password in a script.
- Other times, the password is not known in advance. In this case, a dynamic credential created when a
 job is launched can prompt a user for information and then pass the credential to the step in a secure
 manner.

Credential Access Control

Because ElectricFlow stores passwords in a recoverable manner, it is important that credentials are accessible only under carefully controlled circumstances. To protect credentials, ElectricFlow uses two mechanisms: access control lists (ACL) and *attached* credentials.

To use a credential for any purpose, a user who attempts to reference the credential must have execute permission on the credential object—this allows you to explicitly control which users are allowed to use a credential anywhere in the system.

In addition to the user-based access control check, ElectricFlow requires a credential be explicitly *attached* to the object that is going to use it. After an object has a credential attached, object modifications are restricted to users who have both execute permission on the credential and modify permission on that object. This safeguard helps prevent credentials from accidentally being used by unauthorized users.

Attach a credential to an object in the automation platform UI one of these ways:

- Set the impersonation credential for a step, procedure, project, or schedule.
 - You can also set the impersonation credential for an application process, component process, or a process step.
 - Impersonation credentials are inherited, so attaching a credential to a container object like a procedure or a project implicitly attaches it to every object inside the container, such as all steps in a procedure.
- Explicitly attach the credential to a step or schedule using the UI or the attachCredential API.
 - To access credentials from a step, the credential must be attached directly to the step, or the credential must be passed as a parameter to the containing procedure and have the parameter attached instead.

Using Credentials for Impersonation at the Platform Level

This information addresses the question, "When a job step is running on a resource, under which account is it running and how can I control that?" By default, a job step runs under the account used by the ElectricFlow agent that runs the step, which was chosen when the agent was installed. This approach works well in environments where it makes sense to run all jobs under a single user such as a "build" user. You install all agents to run as the desired user, and you do not have to worry about anything else—every step of every job runs as that user.

However, in other environments you may prefer to run different jobs, or even different job steps, under different accounts.

For example:

- If independent groups are sharing an ElectricFlow system, you may want each group to use a different account for its jobs.
- You may need to run jobs as particular users to access ClearCase views for those users.
- There may be certain steps that require special privileges.

For example, as part of your builds you may need to run a step that generates a certificate using privileged corporate information; that step must run under a special high-privileged account, but you want the remainder of the steps in the build to use a less-privileged account.

ElectricFlow allows you to select accounts on a per-job or per-job-step basis. This mechanism is called *impersonation*, where the ElectricFlow agent impersonates a particular user for the duration of a job step, and this *impersonation* is implemented using credentials. The use of credentials creates special security challenges.

Note: It is important to ensure privileged accounts can be used only for the purposes you intend and cannot be "hijacked" for other purposes. The ElectricFlow access control mechanism contains special facilities to ensure proper credential use.

Setting the Impersonation Credential in the Automation Platform

Attaching Impersonation Credentials to Steps, Procedures, and Projects

You can attach impersonation credentials to procedure steps, procedures, and projects before executing a job step.

ElectricFlow searches for a credential in the following order and uses the first credential it finds:

- 1. Procedure step
- 2. Procedure to which the step belongs
- 3. Project to which the project belongs

If the step is in a nested procedure, and no credential is found on the step, its procedure, or its project, ElectricFlow checks the calling step (and its procedure and project), its caller, and so on until it has worked up through the topmost procedure in the job. If no credential is found, the job step runs under the ElectricFlow agent account on its resource.

This approach makes it easy to manage your account usage. For example, if you want all jobs in one project to use a particular account, define a credential in that project for the account, and then attach the credential to the project. If you want all jobs in a project to use a particular account, except for one step that should use a different account, create two credentials in the project. Attach the first credential to the project and attach the second credential to the particular step.

Attaching Impersonation Credentials to Schedules

You can also attach a credential to a schedule. The credential in the schedule is used for steps where no other credential is available.

The schedule's credential receives the lowest priority.

Attaching Impersonation Credentials to Jobs

You can specify a credential when you launch a job manually one of these ways:

- · Provide the name of an existing credential
- · Enter an account and password.

If you choose the second approach, the account and password are stored only in ElectricFlow for the duration of the job.

For either way, the credential you specified when you launched the job is used as a last resort for steps only with no other credential.

Attaching a New Impersonation Credential

To attach a new impersonation credential for a procedure, step, or schedule:

- 1. Click the Projects tab.
- 2. Select a project.
- 3. Go to the Project Details page for the project you selected.
- 4. Do one of the following:
 - a. For a procedure:
 - i. Click Create Procedure to go to the New Procedure page.
 - ii. On this page, select the Impersonation Credential drop-down box to select a credential to use for impersonation.

b. For a step:

- i. Select a procedure to go to the Procedure Details page.
- ii. Select a step name to go to the Edit Step page.
- iii. On the Edit Step page, scroll down to the Impersonation section.

c. For a schedule:

- i. Select the Schedules subtab.
- ii. Select a schedule to go to the Edit Schedule page.
- iii. Scroll down to Impersonation Credential section.

For more information in the platform GUI, each of these web pages contains a **Help** link in the upper-right corner.

Accessing Credentials From a Step at the Platform Level

The ElectricFlow server supports an operation called <code>getFullCredential</code> that can return a password when called from inside a step. The ectool interface for this operation is:

```
getFullCredential <credential name>
   [--value <password|user>]
```

Without the --value option, the response is the same as that returned by getCredential, with the addition of a "password" element.

With --value, the simple text password or user name value is returned on *stdout* so it can be used directly without XML parsing.

The getFullCredential API is allowed only inside a running step, and is allowed to use credentials that were explicitly attached to the step only.

Attaching a Credential to a Procedure Step, Procedure, or Schedule

A credential must be explicitly attached to the object using it so the server can perform an access control (ACL) check at definition time and limit the visibility of the password. To support accessing credentials other than the one being used for impersonation, steps, procedures, and schedules contain a list of credential names.

In the platform GUI, to attach a credential (other than an impersonation credential) to a step, procedure, or schedule:

1. Go to an "Edit" page.

For example, to create a credential for a step, a step must already exist.

- 2. Select a step that needs a credential. The Edit Step page opens.
- 3. Scroll to the lower section of the page and click Attach Credential.

If you are an ectool command-line user, use the attachCredential API command. For more information on the command-line interface, go to the *ectool* Help topics within the online Help system.

Passing Credentials as Parameters

Sometimes you may not know which credential you need to use when you define a procedure step. In this situation, you can leave the credential choice up to the caller of the procedure.

The following example describes the procedure for passing credentials as parameters:

- 1. Create a project named InnerProj.
- 2. Within that project create a procedure named InnerProc with a parameter 'cred' of the type credential.
- 3. Within that procedure create a step named InnerStep with:
 - o Command: ectool getFullCredential cred
 - Attached Parameter Credential: cred
- 4. Create a project named OuterProj with two credentials: cred1 and cred2.
- 5. Within that project create a procedure named OuterProc with two parameters, cred1 and cred2, of the type credential.
- Within that procedure create two subprocedure steps that both call InnerStep from the InnerProj project:
 - OuterStep1 with Attached Parameter Credential cred1, and cred1 in the Parameters section
 - OuterStep2 with Attached Parameter Credential cred2, and cred2 in the Parameters section

If you are an ectool command-line user, you can use the <code>createFormalParameter</code> and <code>attachParameter</code> API commands. For more information on the command-line interface, go to the <code>ectool</code> Help topics within the online Help system.

Credential References

For schedules, the value for any actual parameter passed to a credential-type formal parameter is interpreted as a reference to a previously attached credential on the schedule.

For procedure steps, the value for any actual parameter passed to a credential type formal parameter is interpreted as a reference to a previously attached credential or credential parameter on the calling step.

When a job step is created, all credential references are resolved and the content is copied into transient credentials on the job step. Credential parameter references are resolved by looking for a credential in the calling job step's transient credential map—looking for a credential with the name specified as the actual parameter value.

Credential references in the API can be specified as a relative or an absolute credential path.

- A relative path is similar to "rootCred" and is interpreted as referring to a credential in the current project.
- An absolute path includes a project name like "/projects/MyProject/credentials/rootCred" and can refer to a credential in any project.

Anywhere a credential reference is accepted (for example, setting an impersonation credential or attaching a credential to a step), either form can be used.

Best Practices for Retrieving Credentials

The ElectricFlow impersonation mechanism provides powerful mechanisms for using different accounts in job steps, and it can be used to handle a variety of challenges. However, if misused, impersonation can open up dangerous security loopholes. This section discusses these risks and how to avoid them.

Use of credentials and impersonation tend to fall into two classes.

First Class

In the first class, your goal is to open up access to accounts, you want to make some accounts generally available, and you trust large groups of users (such as everyone who can access a particular project) to use the accounts appropriately. This usage type is simplest and relatively safe.

- 1. Define an ElectricFlow group containing all trusted users.
- 2. Set permissions on the Project so only trusted users can access the project.
- Define credentials for accounts in the Project and grant execute permission to everyone in the group.

After completing this task, everyone in the group can use the accounts for arbitrary purposes. No one outside the group will be able to access the project or the credentials.

Second Class

In the second class, you want to provide tightly-controlled access to a highly privileged function such as generating a certificate or accessing sensitive information. This kind of usage requires careful thought to avoid allowing unintended access to the credential. The best way to implement a solution for this scenario is to restrict access to the credential and use it on individual steps only after careful analysis of issues such as the following:

Does the step's command use properties?

If so, someone could potentially change the property value to change the step behavior. For example, suppose the step's command is defined like this:

```
echo $[/myProject/x]
```

Anyone with write access to property x can hijack the step by placing the value "nothing; myCommand" in the property.

After property substitution, the final command will be:

```
echo nothing; myCommand
```

that will cause "myCommand" to execute in the step.

Either restrict access to the property sheet or avoid substitutions.

Does the step execute code or commands that can be modified?

In most jobs that execute build and test sequences, the job extracts various files from a source code control system and some of those files contain code executed during the build. In this case, anyone with access to the source code control system can affect the executed code.

For example, suppose a step runs a Make or Ant command using a configuration file extracted from the source code control system. Anyone with access to the source code control system can modify the configuration file to introduce new commands executed during the Make or Ant step. In this situation, it is virtually impossible to control what happens during the step, so you should never attach a credential for a privileged account to such a step.

The higher you attach a credential, the greater the risk of uncontrolled access to the credential's account.

For example, if you attach a credential to a project, there is a good chance anyone with write access to the project, or even the ability to execute the project's procedures, can hijack the credential using one of the loopholes described above. If you attach a credential to a step that invokes a subprocedure, you effectively give anyone with write access to the subprocedure complete access to the credential's account. If you care about access to an account, you should use its credential in the narrowest possible fashion, such as attaching it to a single job step.

Avoid passing passwords on the command line

When using credentials, avoid passing the password on the command line because [on most platforms] those passwords will be visible to all users logged into the system. The secure way to pass a credential to an application is to use a secured file in the file system or to pipe the value into the application via stdin.

Defect Tracking

The Defect Tracking plugin enables linking existing defects to an ElectricFlow job. "Existing defects" are those defects previously created in the defect tracking system and already associated with the ElectricFlow job in some way.

For example, a defect is associated with a ElectricFlow job if the fix for the defect is part of the source code snapshot being built and tested in the job.

ElectricFlow pre-installs numerous defect tracking plugins including Bugzilla, ClearQuest, Fortress, JIRA, Quality Center, Ration Team Concert, Rally, Team Foundation Server, TeamForge, TestTrack, and more, using plugin integrations.

The following examples and instructions are for a JIRA integration. The steps for other defect tracking systems are similar to those for JIRA.

Scenario example

A developer fixes a defect in the "ABC" project and checks in the fix to the source control system, along with a comment that includes the fixed defect ID, for example:

```
ABC-123: fixed EFG bug
```

When the next ElectricFlow job is triggered for the ABC project, ElectricFlow checks out a source code snapshot from the source control system and queries the source control system for a log containing check-in details. This log, which should contain the above comment, will be stored in a property on the job.

The JIRA plugin will then do the following:

- Parse the property containing the source control system log to identify defect IDs.
- Query the configured JIRA server with identified defect IDs, including "ABC-123".
- Construct a descriptive URL to point to the JIRA defect on the JIRA server.
- Include the URL in the JIRA Report.
- Link to this JIRA Report from the Job Details page.

Example: Enabling the JIRA integration in your procedure

To ensure ElectricFlow links existing defects to a job, create a step to link the defects.

Go to Projects > select a Project > select a Procedure. To create a New Step, select the **Plugin** link.

- In the Choose Step panel, select Defect Tracking from the left pane, then select the defect tracking system you configured.
- The right-pane now shows the types of steps available for your configuration. Select the step you need and automatically go to the New Step page.
- On the New Step page, notice the Subprocedure section now contains the defect tracking integration you configured and the step you chose.

On the New Step page, enter information in the following fields:

General section:

Name-Enter a unique name for your subprocedure step (any name of your choice).

Description - (optional) A plain text or HTML description for this object. If using HTML, you must surround your text with <html> ... </html> tags. The only HTML tags allowed in the text are: <a>
 <div> <dl> <i> <style> <tc>

Subprocedure section:

Resource–(optional) You do not need to change the resource you previously chose to run the procedure.

Parameters section:

config-(required) This is the name of the Defect Tracking configuration you created on the New Defect Tracking Configuration page.

Note: If you did not previously create a Defect Tracking configuration, you must do that now before you proceed. Go to Administration > Defect Tracking > click the **Create Configuration** link to see the New Defect Tracking Configuration page. Access the **Help** link on that page for more information.

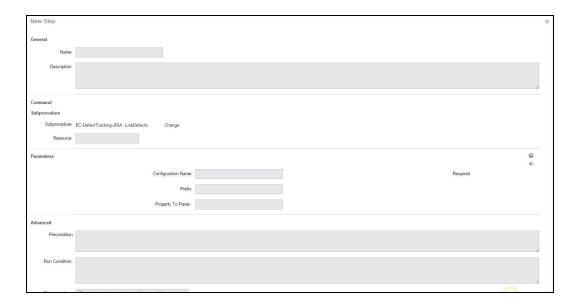
prefix–(optional) This is the key used by JIRA as the prefix for defects within a project. If the "prefix" field is blank, a regular expression is used to try and match defect IDs.

propertyToParse - (optional) This is the property or property sheet to search for defect IDs. If the "propertyToParse" field is blank, the default property /myJob/ecscm changeLogs is used.

Advanced section:

You do not need to enter any information in these fields.

The following screen is an example of the New Step page where you will create your Defect Tracking subprocedure step.



ElectricSentry

ElectricSentry is the ElectricFlow **engine** for continuous integration, integrating with Source Control Management (SCM) systems. ElectricSentry is installed automatically with ElectricFlow and is contained in an ElectricFlow plugin named ECSCM. When ElectricSentry detects a new or modified source file check-in, it

launches a user-defined procedure to build and test the latest version of source files. Using a few simple properties, the entire process is controlled by the user who determines when ElectricSentry is active and which projects and branches are monitored and built.

While you can interact with ElectricSentry directly, you may want to use the Continuous Integration Dashboard, the front-end user interface for ElectricSentry, which provides a visual display of your builds, and automates the configuration process for your branch, project, or procedure so you can start using continuous integration schedules quickly.

This Help topic describes using ElectricSentry directly, and NOT the Continuous Integration Dashboard, which has its own Help topic for dashboard use. To get started with the "dashboard", select the Home > Continuous Integration tabs.

How ElectricSentry works

ElectricSentry finds procedures the user wants to monitor by scanning for a special type of ElectricFlow schedule. From the schedule, ElectricSentry retrieves the name of the procedure to run, along with values to use as actual parameters when the procedure runs.

ElectricSentry can run at a regular interval you determine. When it runs, ElectricSentry scans all ElectricFlow projects, looking for schedules configured for use with ElectricSentry. When it finds one, ElectricSentry checks to see if there is already a job running based on that schedule and will not start a second job for the same schedule until the previous job completes.

Next, ElectricSentry queries the SCM system to see if new check-ins have occurred since the last attempted build. If the most recent check-in is different from the last attempted build, ElectricSentry then checks a "quiet period" by comparing the current time to the check-in time. If the user-configured time interval has elapsed since the most recent check-in, ElectricSentry runs the procedure to start a new job.

Finally, ElectricSentry deletes previous runs of itself, so only the most recent run appears in the jobs list. If using the Continuous Integration Dashboard, this is slightly different—you will see the previous 10 builds that ran.

Configuring ElectricSentry

ElectricSentry does not require any specific configuration, but it does allow fine tuning for some key settings. To change any setting, you must be logged into ElectricFlow and then navigate to the Electric Cloud project.

Quiet time

When setting up a continuous integration system, it is common to require an inactivity period before starting a build. This time period allows developers to make multiple, coordinated check-ins to ensure a build does not start with only some of the changes—assuming all changes are checked-in within the specified inactivity time period. This time period also gives developers an opportunity to "back-out" a change if they realize it is not correct.

With ElectricSentry, the inactivity time period can be configured globally for all projects or individually for a single project. The global setup is stored in the ECSCM-SentryMonitor schedule in the Electric Cloud project, using a property named ElectricSentrySettings/QuietTimeMinutes. You can set this property to the number of minutes you prefer for the quiet time prior to starting a build. If you remove this property, ElectricSentry will use a five-minute default time.

Resource

The ECSCM plugin contains a procedure named ElectricSentry. This procedure has a parameter named sentryResource that determines where ElectricSentry will run. The parameter defaults to a resource named *local*, which is created by the ElectricFlow installer. The local resource is configured to run on the machine named *localhost*, which refers to the machine where the ElectricFlow server is running.

You may need to change the resource where ElectricSentry runs, if for example, the server machine is not configured to run your SCM tools. You can change this configuration by modifying the ECSCM-SentryMonitor schedule to specify a different resource in the sentryResource parameter.

To do this:

- 1. Go to the Projects page.
- 2. Select the Electric Cloud project name to go to the Project Details page.
- 3. Click the Schedules subtab.
- Select the ECSCM-SentryMonitor schedule.
- 5. Change the sentryResource field to the name of the resource (or the resource pool name) where you want ElectricSentry to run.
- 6. Click **OK** to save your new information.

Polling frequency

ElectricSentry is set to look for new check-ins every five minutes. If you want ElectricSentry to check more or less often, open the ECSCM-SentryMonitor schedule in the Electric Cloud project. Change the number of minutes to speed up or slow down the ElectricSentry polling frequency entry.

Time-of-day and day-of-week

The initial setting allows ElectricSentry to run between 7 am and 11 pm every day of the week, but you can change this time period to monitor check-ins at different hours. The time period should cover the most likely hours developers would be checking-in changes, then frees the resources (after the last job started by ElectricSentry finishes) to run overnight builds. You can modify the ECSCM-SentryMonitor schedule to change the hours or days ElectricSentry monitors check-ins.

Configuring a build for continuous integration

You can control which projects and branches ElectricSentry monitors for continuous integration by creating a schedule that runs a procedure and by adding ElectricSentry settings. These tasks are made easier using the Continuous Integration Dashboard.

Source Control Management SCM) configurations

ElectricFlow bundles and supports a number of source control types. After creating a source control configuration, your entry will appear in the table on the Source Control Configurations web page—to see this web page, select the Administration > Source Control tabs.

IMPORTANT: Click the **Create Configuration** link to configure your source control system.

Note: Electric Cloud has tested numerous SCM versions and where possible, all SCM integrations were created in a generic manner to avoid SCM release-specific differences. In most cases, Electric Cloud supports any version of Perforce, ClearCase, Subversion, AccuRev, CM Synergy, Borland's StarTeam, and many others.

Select the Administration > Plugins tabs to go to the Plugin Manager page to the current list of available SCM plugins. If you do not see the SCM you need, it may be available in the Plugin Catalog—select the View Catalog tab.

Create a procedure

Navigate to your project and create a procedure that checks out sources from your SCM system and then runs your build process. we recommend having the procedure take a "branch name" parameter, then use this single procedure to build multiple branches of the same project.

Create a schedule

In your project, create a schedule that calls the procedure and specifies parameters to call the procedure. On the Project Details web page, select the Schedules tab and click the **CI Configuration** link to go to the New CI Configuration page. Enter a name for the new schedule and choose the SCM Configuration you want to use. For more information on this page, click the **Help** link in the upper-right corner of the screen. If you set up your procedure with a parameter for a branch, create one schedule for each branch you want to monitor.

Optional - running ElectricSentry on multiple resources

ElectricSentry normally runs as a single procedure executing queries against all SCM systems configured by ElectricSentry "trigger schedules." Because all queries are executed in a single step, these queries are executed by the same ElectricFlow resource. This process works very well for many companies, but for some large enterprise companies this process is not adequate.

You can benefit from using ElectricSentry on multiple resources if any of the following items apply to your organization:

- Connectivity No one single resource has access to all SCM systems in your organization.
- Security/Authentication Specific resources may contain Authentication tickets required for access to certain parts of your SCM depots, and no single resource contains all of these tickets.
- Scalability You may need too many trigger schedules to monitor in a single execution interval, which can cause unnecessarily long execution times for ElectricSentry jobs.
- Scheduling ElectricSentry is most useful during the work day when developers are adding new code
 check-ins. At night, it is common to redeploy resources to do nightly builds. If ElectricSentry runs on a
 central resource that serves development teams in multiple time zones, there is no way to schedule
 independent Sentry operation for different time zones.

Overview

The solution to all of these issues [above] is to run multiple ElectricSentry instances concurrently.

- Each Sentry instance can choose its own resource and its own operation schedule.
- Each instance can be restricted to monitoring schedules for one or more ElectricFlow projects.
- Also, you can create a global instance to monitor all projects not monitored by a restricted instance.

Each ElectricSentry instance is implemented by reusing the ElectricSentry procedure in the ECSCM plugin. This procedure has two parameters, <code>sentryResource</code> and <code>projectList</code>, that allow you to specify the configuration for different instances.

Each instance will have its own copy of the SentryMonitor schedule—one for each resource where you want ElectricSentry to run.

The default configuration has no projects specified in the "projectList" parameter—this is called the *global* instance. The global instance monitors all projects not explicitly monitored by any other instance. If a global instance is your only instance, it will monitor all projects.

Configuring ElectricSentry to use multiple resources

When ElectricFlow is installed, it creates a project called Electric Cloud that contains ElectricSentry's Sentry Monitor schedules. By default, all steps will run on the resource named local and all projects will be monitored. To use multiple resources, you need to create new schedules in the Electric Cloud project, creating new ElectricSentry instances.

- 1. Navigate to the Electric Cloud project and select it.
- 2. Select the Schedules tab.
- 3. Disable the schedule named ECSCM-SentryMonitor.
- 4. Click the **Copy** link next to the ECSCM-SentryMonitor schedule *repeat* for each new ElectricSentry instance you want to create.
- 5. Re-enable the schedule named ECSCM-SentryMonitor.
- Navigate to one of the new schedules you created and select it to edit it change the name from ECSCM-SentryMonitor copy n to something more meaningful, for example, "ECSCM-SentryMonitor for ABC team".
- 7. In the Parameters section:
 projectList box enter the project names you need to monitor one project name per line.
 sentryResource box enter the resource or pool name where you want this ElectricSentry instance to run.
- 8. In the Frequency section:

 Adjust the "time of day" and "days of the week" settings to schedule this ElectricSentry instance to run.

Note: If you are setting up ElectricSentry to run in a different time zone, make sure you set the Time Zone field for the location where ElectricSentry will run.

- a. In the Advanced section, click the Enabled check box.
- b. Click the Save button.

Note: For each new schedule you create, repeat steps 6-10 above.

The Job Step Execution Environment

Terms and definitions

Various factors influence job step execution. Both a job step and its postprocessor (if any) run in the same environment, except shell. The shell used to run a command comes from what is specified in the step. If a shell is not specified on the step, the shell specified on the resource is used—this is also the same shell used for running the postprocessor.

Machine

The machine where a job step is executed is determined by the resource specified in the corresponding procedure step. If a pool is specified in the procedure step, ElectricFlow picks a specific resource from that pool. A job step can determine its actual resource by querying the property /myResource/resourceName or assigned resourceName property on the job step. The host where a step is expected is /myResource/hostName. If the resource is a proxy agent, this property contains the name of the proxy target. The agent host name is in /myResource/proxyHostName.

OS-level access control

For its resource, a job step executes under the same account as the ElectricFlow agent. You may want to contact Electric Cloud technical support for help configuring ElectricFlow agents. Basically, the agent needs to know what account it will run as, and Windows agents require additional setup if impersonation is used. If using impersonation, the job step runs under the credential effectively attached to the step.

Environment variables

For its resource, a job step inherits environment variables from the ElectricFlow agent. The agent's environment variables can be configured as part of the agent configuration. In particular, the PATH environment variable typically includes the ElectricFlow installation directory for easy access to applications such as *ectool* and *postp* (the ElectricFlow postprocessor).

Working directory

The default job step working directory is the top-level directory in the job's workspace. However, you can configure the working directory with a property on the procedure step. When you run on a proxy agent, the step actually runs on the proxy target, in the working directory specified in the step. If the working directory is not specified, the step runs in the UNIX path to the workspace.

Standard I/O

Standard job step output and errors are both redirected to the step log file.

ElectricFlow access

For easy access to ElectricFlow, the following environment variables are set automatically for a job step:

COMMANDER_HOME

A variable whose value is the base installation directory for an agent. On Windows, this directory is typically C:\Program Files\Electric Cloud\ElectricCommander. On UNIX platform, this directory is typically /opt/electriccloud/electriccommander.

COMMANDER SERVER

IP address for the ElectricFlow server machine.

COMMANDER_PORT

Port number for normal communication with the ElectricFlow server.

COMMANDER PLUGINS

(configurable) The directory where installed plugins live (for example, C:\Documents and Settings\All Users\Application Data\Electric Cloud\ElectricCommander\Plugins)

COMMANDER_HTTPS_PORT

Port number for secure communication with the ElectricFlow server.

COMMANDER_JOBID

Unique identifier for the job containing the current job step.

COMMANDER_JOBSTEPID

Unique identifier for this job step.

COMMANDER_SESSIONID

ElectricFlow session identifier for the current job, which allows ectool to access ElectricFlow with job associated privileges, without an additional login.

COMMANDER_WORKSPACE_UNIX

Absolute path to the workspace directory for this job, in a form suitable for use on UNIX machines.

COMMANDER_WORKSPACE_WINDRIVE

Absolute path to the workspace directory for this job, in a form suitable for use on Windows, and starting with a drive letter.

COMMANDER_WORKSPACE_WINUNC

Absolute path to the workspace directory for this job, in a form suitable for use on Windows, specified using UNC notation.

COMMANDER WORKSPACE

Absolute path suitable for accessing the top-level workspace directory for this job on this machine. On a UNIX machine, this has the same value as COMMANDER_WORKSPACE_UNIX. For Windows, it is the same as COMMANDER_WORKSPACE_WINUNC.

COMMANDER_WORKSPACE_NAME

This is the name of the workspace on the ElectricFlow server.

These environment variables allow you to invoke ectool without specifying a --server argument. They also provide a context for accessing properties in ectool. For example, "ectool getProperty foo" looks for the property named "foo" on the current job step. Environment variables provide ectool with a session including all user privileges associated with the job.

Also, if the resource is a proxy agent, these environment variables are available in the step's environment on the proxy target.

Preflight Builds

Preflight Build Solution

Preflights with ElectricFlow

Samples

Default SCMs

Troubleshooting

Why use Preflight Builds?

When developers make changes, they generally build and test their code locally and *commit* their changes if the code ran successfully. Unfortunately, this process is not sufficient.

Code changes may break the production build because of environment differences, platform-specific issues, or incomplete commits. These issues leave the product in a broken state until the changes are backed-out of the product or the developer commits a fix, which often affects the entire development team. It is a time-consuming process to find the problem when multiple developers commit their changes simultaneously and the build breaks.

Preflight Build Solution

A **Preflight Build** is used to build and test a developer's changes before those changes are committed. A "post-commit" source tree is simulated by creating a clean source snapshot and overlaying the developer's changes on top of it. These sources are then passed through the production build procedure to validate the changes work successfully.

Developers are allowed to commit their changes only if the preflight build is successful. Because developer changes are built and tested in isolation, many common reasons for broken production builds are eliminated.

Preflights with ElectricFlow

Workflow

After a preflight build is configured, the general flow of interaction between the server, agent, and client is:

- A developer invokes the client preflight program either through an IDE or via the command-line. The following arguments are specified in the XML configuration file or on the command-line:
 - Define which ElectricFlow server to use (hostname, ports, and so on),
 - Define the project, procedure, and parameters to use to start the preflight build, and
 - SCM-specific information to use (ports, changelists, and so on).
- Next, the client preflight program connects to the ElectricFlow server and launches the specified procedure.
 - The job is started and the developer's changes are uploaded.
 - If auto-commit was turned on, the client program waits for the job to complete. Otherwise, the client program exits immediately.
- The job progresses until it gets to the special step added to the procedure to extract sources and overlay changes.
 - The agent-side driver creates a clean source snapshot based on configuration information passed to it from the client-side.
 - The step overlays changes [uploaded by the client] on top of the snapshot to simulate the developer's check-in.
- · The build continues as usual, using the modified sources.
- If the developer chose to auto-commit changes, the preflight program was waiting for the job to complete. If the job is successful, the changes are committed if:
 - o No files were added or removed from the change sets.
 - No files were modified since the preflight build started.
 - o All conditions set in the driver script were met.

Components

Three components are required for an ElectricFlow preflight build:

- Server This is the main ElectricFlow server that will run the preflight builds.
- Agent This is any ElectricFlow agent machine that is setup as a resource for the server. This machine
 must be able to communicate with the source control system because it is responsible for creating a
 clean source snapshot on which the developer's changes are overlayed.
- Client This is the machine where a developer has active code changes to submit for a preflight build. Like the agent, this machine must be able to communicate with the source control system because it is responsible for determining which files were modified so it can transmit these files to the agent.

Installation

Preflight is automatically installed with SCM plugins. ElectricFlow pre-installs numerous source code management (SCM) plugins including Accurev, Bazaar, CVS, ClearCase, Git, Mercurial, Perforce, SVN,

StarTeam, Team Foundation Server, and Vault. Before you can use your preferred SCM, you will need it to communicate with the ElectricFlow server.

Configuration

This section provides details for setting up and running Preflight builds in a typical build environment.

Server

Adding the preflight snapshot step

A typical production build procedure has a step that extracts a clean source code snapshot before the actual build starts. For preflight builds, a special step needs to run in place of this snapshot step. This special step is responsible for creating the base snapshot and then overlaying the developer's changes.

- On the Procedure Details page, to create a new step click the Plugin link.
- On the Choose Step panel, select Source Code Management.
- Choose the ECSCM plugin for your SCM or from the right-pane, select the "Preflight" step for your SCM.
- On the New Step page, notice that your SCM is displayed in the Subprocedure section.
 - Set the Resource name for the step to an agent that can communicate with the SCM system, so it can create the source snapshot.
- The Parameter section also displays appropriate fields to enter information for your SCM. The following two fields are common to most SCMs:
 - Configuration–Enter the name you created for your SCM configuration.
 - Destination Directory—This is a path relative to the job's workspace, where the source tree will be created.
- Enter all other information required for your new step.
- · Click OK.

Choosing between snapshot steps

Only one snapshot step should run depending on whether a production or preflight build is invoked. To enable the appropriate step, add a checkbox-style parameter to the build procedure to determine whether or not the build is a preflight. Now, use that parameter in the Run Condition field on the two snapshot steps. In the following example, the parameter is named "preflight", with an unchecked value of "false" and a checked value of "true".

- For the preflight snapshot step, set the run condition to <code>\$[preflight]</code>
- For the production snapshot step, set the run condition to
 \$[/javascript getProperty("preflight") == "false"]

Choosing how files are uploaded

By default, when developers run preflight builds, their changes are uploaded to the job workspace via the server. However, if the job workspace is on a network share, accessible to the developer's machine, a useful optimization is for files to be copied directly into the job workspace. To use the network share approach, set a flag on the top-level procedure that is invoked when running a preflight build. This property can be created on the Procedure Details web page as follows:

- 1. Create a nested property sheet on the procedure called ec preflight.
- 2. Create a property in that procedure called waitForStep with a value of "1". This property can be created from the command-line by calling:

Agent

The agent machine must be able to communicate with the SCM system to create the source snapshot.

Client

From a developer's machine, there are two ways to start preflight builds:

- From an IDE:
 - Eclipse
 - Open a Run dialogue, select Launch Commander Procedure, then select an existing preflight configuration or create a new one.
 - For more information, see the *ElectricCommander Eclipse Integration Tech Note* PDF file or the Eclipse plugin.
 - Visual Studio (See the ElectricCommander Visual Studio Tech Note PDF file or the Visual Studio plugin)
- From the command line:

The executable program, ecclientpreflight, is used to run a preflight build from the command-line. This program is included in the ElectricFlow "Tools" installation. When you run ecclientpreflight, configuration options can be passed in an XML configuration file or on the command-line.

If the same option is specified in both the configuration file and on the command-line, the value passed on the command-line takes precedence. The XML configuration file can be passed to the program via the --config option. If the XML configuration file is not passed in, current and ancestor directories are searched for a file named ".preflight", using the first one it finds as the configuration file.

Generally, a ".preflight" file is stored at the root of a developer's workspace, containing configuration information specific to that workspace.

Samples

In the following samples, the SCM was set to Perforce.

The following table is output from "ecclientpreflight --help", enumerating the command-line options. Some options may be specified using an abbreviated format (-c, with a single dash) or a more verbose format (-config, with two dashes).

Electric Cloud Client-Side Preflight Tool, Version: 4.x Copyright (C) 2012 Electric Cloud, Inc. All rights reserved. Usage: ecclientpreflight [options]

General Options

-c,config <file></file>	Load configuration options from the specified XML file. If <file> is '-', read standard input. If not supplied, the current and ancestor directories will be searched for a .preflight file. The first .preflight found will be used as the preflight configuration.</file>
load <file></file>	Evaluate <file> after option processing. May appear multiple times.</file>
-d,debug	Display debugging information.
-v,version	Display version information.
-h,help	Display this information.
Server Communication Opt	ions
server <name></name>	The host name (defaults to localhost).
port <port></port>	The HTTP port (defaults to 8000).
securePort <port></port>	The HTTPS port (defaults to 8443).
secure	If set, only secure connections are used. Off, by default.
timeout < <i>timeout</i> >	The response timeout in seconds (defaults to 180).
userName < <i>name</i> >	The name of the ElectricFlow user to login as.
password <password></password>	The password for the specified user. If blank and the user has no active session, a prompt to enter the password will be provided.
driverLocation	Property path to a property sheet containing the driver scripts that will be loaded and run after a connection to the server is established. Defaults to /server/ec_preflight.
mainDriver	The name of a property in the sheet specified by driverLocation. This is the main driver script that will be loaded and run after a connection to the server is established.
Logging Options	
-1,log	If specified, debug information is logged. Off, by default.

logDir	Where to store log and other files. Defaults to the user's home directory.
Procedure Invocation Optio	ns
projectName <name></name>	The name of the ElectricFlow project containing the procedure to invoke.
procedureName <name></name>	The name of the ElectricFlow procedure to invoke.
-p,param <name>=<value></value></name>	Enter additional parameters to the procedure. May appear multiple times. If any parameters were specified in the config file, those supplied on the command-line append to that list, overriding parameters with the same name.
priority	The priority of the job. Possible values are low, normal, high, highest. If left unspecified, defaults to normal.
jobTimeout <timeout></timeout>	The number of seconds to wait for the job to complete when auto-committing changes. Defaults to 3600 seconds (1 hour).
waitForJob	Wait for the job to complete and report its outcome. This action does not occur by default unless a set of SCM charges are being committed automatically.
runOnly	Run the procedure and exit immediately. The SCM driver is not downloaded in this case.
SCM Options	
scmType	The name of the SCM, required, unless the procedure is invoked in 'run only' mode. The driver is downloaded from \$driverLocation/clientDrivers/\$scmType.
	Valid values: ECSCM-Accurev, ECSCM-Bazaar, ECSCM-ClearCase, ECSCM-CVS, ECSCM-Git, ECSCM-Mercurial, ECSCM-Perforce, ECSCM-Repo, ECSCM-StarTeam, ECSCM-SVN, ECSCM-TFS, ECSCM-Vault.
autoCommit <1 0>	Whether or not the changes should be automatically committed if the job completes successfully. Off, by default.
commitComment <comment></comment>	A comment for the auto-commit.

Perforce Options	
p4port <port></port>	The value of P4PORT. May also be set in the environment. This is a required value.
p4user <user></user>	The value of P4USER. May also be set in the environment. This is a required value.
p4passwd <password></password>	The value of P4PASSWD. May also be set in the environment.
p4client < <i>client</i> >	The value of P4CLIENT. May also be set in the environment. This is a required value.
p4template <template></template>	The name of a Perforce client used to create a base snapshot before overlaying local changes. Defaults to the value of p4client if not specified.
p4changelist <change></change>	The changelist number (or default) whose changes are being tested. May be specified multiple times. If no changelists are specified, all client changelists will be used.

The following is a sample XML configuration file with all options specified:

```
<?xml version="1.0" encoding="utf-8"?>
 <data>
   <server>
     <userName>myUser</userName>
     <password>myPass</password>
     <hostName>ElectricFlowServer
     <port>1234</port>
     <securePort>2345</securePort>
     <stompPort>3456</stompPort>
     <secure>0</secure>
     <timeout>3456</timeout>
        <driverLocation>/server/ec_preflight/alternateDrivers</driverLocation>
        <mainDriver>myMainClientDriver
     </server>
     projectName>myProject
      oprocedureName>
```

```
<parameter>
       <name>branch</name>
        <value>main</value>
    </parameter>
    <parameter>
       <name>preflight</name>
       <value>true</value>
    </parameter>
    <priority>high</priority>
    <waitForJob>1</waitForJob>
    <jobTimeout>3600</jobTimeout>
  </procedure>
  <scm>
    <type>perforce</type>
    <port>perf:1234</port>
    <user>myUser</user>
    <password>myPass
    <client>myUser-main-client</client>
    <template>myUser-main-template</template>
    <changelist>default</changelist>
    <changelist>67382</changelist>
    <autoCommit>1</autoCommit>
    <commitComment>Fixing bug 38582.</commitComment>
  </scm>
</data>
```

In most cases, a developer does not need to override all default values, and will probably want to specify passwords and some SCM-specific options on the command-line.

The following is a more typical XML configuration:

```
ocedureName>myProcedure
   <parameter>
      <name>branch</name>
      <value>main</value>
   </parameter>
   <parameter>
      <name>preflight</name>
      <value>true</value>
   </parameter>
   <jobTimeout>3600</jobTimeout>
 </procedure>
  <scm>
      <type>perforce</type>
      <port>perf:1234</port>
      <user>myUser</user>
      <client>myUser-main-client</client>
 </scm>
</data>
```

If this file is stored in a .preflight file at the top-level of a source tree, the command to start the preflight would look something like this:

```
ecclient
preflight --p4changelist 56793 --autoCommit 1 --commit
Comment "Fixing bug 38582."
```

Default SCMs

ElectricFlow includes preflight support for some of the most common Source Control (SCM) systems. For an SCM plugin you installed, look for a Help link for that plugin (on the Plugins page). Command-line options are included in this Help topic for the following SCMs:

Acc uR ev	Git	Star Tea m
Baz aar	Me rcu rial	Sub vers ion

Cle arC ase	Pe rfo rce	TFS
CV	Re	Vau
S	po	It

Perforce

The command-line options for the Perforce driver for ${\tt ecclientpreflight}$ are:

Perforce Options / Descriptions	
p4port <port></port>	Required: The value of P4PORT may be set in the environment also.
p4user <user></user>	Required: The value of P4USER may be set in the environment also.
p4passwd <password></password>	The value of P4PASSWD may be set in the environment also.
p4client <client></client>	Required: The value of P4CLIENT may be set in the environment also.
p4template <template></template>	The name of a Perforce client used to create a base snapshot before overlaying local changes. Defaults to the value ofp4client if not specified.
p4changelist <change></change>	The changelist number (or default) whose changes are being tested and can be specified multiple times. If no changelists are specified, all client changelists will be used.
p4synctochange <change></change>	The changelist number the Preflight Job should use when sync'ing the source tree. Values are: head - The most recent changelist anywhere in the P4 depot (default). have - The changelist for the most recent file that was synced to 'p4client' changelist. This is a p4changelist number.

Subversion

The command-line options for the Subversion driver for $\operatorname{ecclientpreflight}$ are:

Subversion Options / Descriptions	
svnpath <path></path>	The path to the locally accessible source directory in which changes were made.
svnupdatetohead	Use this option to update the agent workspace created during preflight updated to HEAD. By default, the agent workspace is updated to the revision found in the client workspace.
svnignoreexternals	Causes the preflight process to ignore svn externals.

AccuRev

The command-line options for the AccuRev driver for ${\tt ecclientpreflight}$ are:

AccuRev Options / Descriptions	
accurevuser <user></user>	The value of ACCUREVUSER may be set in the environment also. If not specified, the default is to the ElectricFlow user.
accurevpasswd <password></password>	Required: The value of ACCUREVPASSWD may be set in the environment also. This is a required value.
accurevpending <pending></pending>	Use this option to scan the client workspace for all pending elements using the "stat -fn -p" command. By default, the workspace is scanned for all kept elements using the "stat -fn -k" command.
accurevpath <path></path>	Required: The value of the ACCUREVPATH. Also may be set in the environment.

ClearCase

Preflight support for ClearCase is available only for snapshot views. The command-line options for the ClearCase driver for ecclientpreflight are:

ClearCase Options / Descriptions	
ccpath <path></path>	The path to the locally accessible source directory where changes were made.
ccUnixCSpecPath <unixcspecpath></unixcspecpath>	The path to the appropriate config spec on a UNIX platform for the current view.
ccWinCSpecPath <wincspecpath></wincspecpath>	The path to the appropriate config spec on a Windows platform for the current view.
ccUnixRelativePath <unixrelativepath></unixrelativepath>	The relative path for the current view [on a UNIX platform] from the view root to the directory where changes were made.
ccWinRelativePath <winrelativepath></winrelativepath>	The relative path for the current view [on a Windows platform] from the view root to the directory where changes were made.

Bazaar

The command-line options for the Bazaar driver for ecclientpreflight are:

Bazaar Options / Descriptions	
workdir <path></path>	The developer's source directory.
branch <name></name>	Branch name for the preflight.
method= <local remote=""></local >	local - get tracked and untracked changes in the current workingdir remote - get changes between working tree and remote branch

Git

The command-line options for the Git driver for $\operatorname{ecclientpreflight}$ are:

Git Options / Descriptions	
gitdir=dir	The Git directory to process.
method= <local_all local_tracked="" remote="" =""></local_all>	local_all - get tracked and untracked changes between working tree and local repo
	local_tracked - get tracked changes between working tree and local repo
	• remote

CVS

The command-line options for the CVS driver for ecclientpreflight are:

CVS Options / Descriptions	
cvsroot <path></path>	The path to the repository.
module <module></module>	The module name for the preflight.
workdir <path></path>	The developer's source directory.

Mercurial

The command-line options for the Mercurial driver for ecclientpreflight are:

Mercurial Options / Descriptions		
hgpath <path></path>	The path to the locally accessible source directory in which changes were made. Generally, this path is to the root of the workspace.	

Repo

The command-line options for the Repo driver for ${\tt ecclientpreflight}$ are:

Repo Options / Descriptions	
repoworkdir <path></path>	The developer's source directory.
agentworkdir <path></path>	The path to the source directory, used by the agent.

StarTeam

The command-line options for the StarTeam driver for ecclientpreflight are:

StarTeam Options / Descriptions		
STProjectName <project></project>	The StarTeam project name.	
workingdir <path></path>	Working dir for the updated files.	

TFS

The command-line options for the Team Foundation Server driver for ecclientpreflight are:

TFS Options / Descriptions	
server <url></url>	The equivalent value to Collection in VS2008 and under. If you use this option, it is assumed you are using TF version VS2008 or under and it will use the option /server in the preflight. This field is required if you have VS2008.
collection <url></url>	The URL that points to /collection. This value is used for VS2010. If you specify this value, the command uses the option /collection when it executes the TF query for preflight. This field is required if you have VS2010.
localfolder	The path to the locally accessible source directory in which changes were made. Generally, this is the path to the root of the workspace.
workspace <workspace></workspace>	The workspace containing the data for the preflight.

Vault

The command-line options for the Vault driver for ${\tt ecclientpreflight}$ are:

Vault Options / Descriptions	
workingdir <path></path>	Working dir for the updated files.

Troubleshooting

Client

The --debug option to ecclientpreflight provides additional information about internal actions taken when starting the preflight build. By modifying the properties containing driver scripts, debug output can be added where needed. If you want to capture debug information to a log file, use the --log option. By default, logs are stored in the user's home directory. To change this location, use the --logDir option.

Agent

The step log for the custom preflight step provides information about the snapshot being created and the files being overlayed. Debug output can be added where needed by modifying the properties containing the driver scripts

Properties

Creating or modifying properties

Using property values

Property sheets and intrinsic properties

Property names and paths

Context-relative "shortcuts" to property paths

Property path shortcuts in ElectricFlow 5.0 and later on page 794

Property name substitutions

Expandable properties

The property hierarchy

Special property references

Intrinsic properties listed by object type

ElectricFlow Object / Property tables

Overview

ElectricFlow provides a powerful data model based around the notion of a property.

- A property is a string value with a name.
- Properties can have arbitrary names and values.
- You can attach properties to any object in the ElectricFlow system, such as a project, procedure, or job.
- After a property is created, it is stored in the ElectricFlow database.

- You can retrieve and modify the property value later.
- You can delete properties that you no longer need.

Properties are used extensively throughout the ElectricFlow system and provide a flexible and powerful mechanism to manage data about your builds.

Property Use Case Examples

- When a job starts, it computes a unique identifier for that build and saves it in a property. Later build steps can retrieve the property value to embed the build number in binaries generated during the job.
- When a build executes, it can set a property on the procedure to identify the source code version used
 for the build (for example, a tag or timestamp from your source code control system). The next time the
 build executes, it can retrieve the property value from the procedure and use it to extract information from
 your source code control system about the files that changed since the last run.
- Suppose you have a collection of machines for testing, and some machines have older test hardware
 and some machines have newer hardware with slightly different characteristics. You can set a property
 on each test machine resource to indicate which version of test hardware is available on which machine.
 Later, when a test executes, the test can retrieve the property for the machine where it ran and configure
 tests appropriately for that hardware. If you upgrade test hardware on a machine, all you need to do is
 change the property on that resource to reflect the new version.
- You can set job properties to indicate the build status produced by that job. For example, if your QA team
 finds fatal flaws in a build, it can mark the job accordingly. Builds that need to be preserved (release
 candidates or builds undergoing beta testing at customer sites) can be marked with properties so those
 builds are not deleted.
- When a job executes, properties are set for its job steps that hold metrics such as how many files
 compiled during a build step, how many tests executed during a test step, or how many errors were
 detected in the step. These property values are included in reports and can be examined later to
 compute trends over time. You can define additional properties for metrics useful to you.

ElectricFlow provides Intrinsic properties and allows you to create Custom properties. This Help topic enumerates properties available within ElectricFlow (Intrinsic properties), distinguishes between relative and absolute property paths, and describes property hierarchies.

Note: Intrinsic properties are case-sensitive. Custom properties, like all other object names in the ElectricFlow system, are "case-preserving," but *not case-sensitive*.

• Intrinsic properties

These properties represent attributes that describe the object to which they are attached, and are provided automatically by ElectricFlow for each similar type object. For example, every project has a Description property that can be referenced with a non-local property path such as /projects/Examples/description.

Custom properties

Custom properties are identical to intrinsic properties and when placed on the same object, are referenced in the same manner, and behave in every way like an intrinsic object-level property with one exception: they are *not* created automatically when the object is created. Instead, custom properties can be added to objects already in the database before a job is started, or created dynamically by procedure steps during step execution.

Creating or modifying properties

You can set a property value by using one of the following three methods or applications: the automation platform web interface, the *ectool* command-line application, or the Perl API.

- Using the web interface, you will see a table labeled Custom Properties on the pages that display details
 for projects, procedures, and so on.
 Click Create New Property to create a new property for that object, or click on an existing property to
 change its value.
- Using the ectool application, set a property value with a command like:

```
ectool setProperty /myJob/installStatus complete
```

This command sets the property named /myJob/installStatus to the value complete, and creates the property if it did not already exist (the meaning of property names like /myJob/installStatus is explained below). You can use ectool from within one job step to set properties accessed by later steps in the same job.

Using the Perl API, you can use an external script or a script in a step with ec-perl as the shell.
 A Perl code example:

```
use ElectricCommander ();
my $ec = new ElectricCommander;
$ec->setProperty ("/myJob/installStatus", "complete", {jobStepId =>
$ENV{COMMANDER JOBSTEPID}
```

Note: Using the Perl API may yield better performance if you are requesting multiple API calls in one step.

For more information, see the "Using the ElectricFlow API" topic in the ElectricFlow API Guide.

Using property values

A job step can access a property value by two methods. The first method is *substitution*, using the \$[] notation. Suppose you enter the following text as a command for a step:

```
make PLATFORM=$[platform]
```

Before running the step, ElectricFlow finds the property named platform and substitutes its value in the command string in place of \$[platform]. For example, if the property contains the value "windows", the actual command executed is:

```
make PLATFORM=windows
```

The substitution method can be used for a command in a step and for any step fields, such as the resource or working directory. This method allows you to compute configuration information in an early step of a job, then use that configuration information to control later steps in the job.

Another way for a step to access the property value is with the ectool getProperty command. For example, the following command returns the property value named platform:

```
ectool getProperty platform
```

Property sheets and intrinsic properties

A property value can be a simple string or a nested *property sheet* containing its own properties. Property sheets can be nested to any depth, which allows you to create hierarchical collections of information.

Most objects have an associated "property sheet" that contains "custom properties" created by user scripts. The property sheet is an intrinsic property of the containing object called "property sheet", so to reference a project's custom property "branchName", you could specify "/projects/aProject/propertySheet/branchName".

As a convenience, ElectricFlow allows the "property sheet" path element to be omitted and the path written as "/projects/aProject/branchName". If there is no intrinsic property with the same name, the path will find the property on the property sheet.

Custom properties in a property sheet can be one of two types: *string property* or a *property sheet* property. String properties hold simple text values. Property sheet properties hold *nested* properties. Nested properties are accessed via the property sheet property of their containing sheet.

For example:

```
/projects/aProject/propertySheet/topSheet/propertySheet/propB
- or -
/projects/aProject/topSheet/propB
```

All information managed by ElectricFlow exists in the form of properties and property sheets and your own custom-created properties. For example, each project, procedure, and step is represented internally as a property sheet—the command for a step is actually a property associated with the step, and so on. Every value in the ElectricFlow system can be accessed as a property, using the naming facilities described below. These properties are called *intrinsic* properties. *ElectricFlow enforces some restrictions on intrinsic property values, whereas custom properties can have any value you choose.*

To learn more about intrinsic properties defined for an object by ElectricFlow, see the "Using the ElectricFlow Perl API" topic in the ElectricFlow API Guide. For example, to learn about intrinsic properties associated with each project, find the documentation for the <code>getProject</code> ectool command. The result of running this command is an XML document whose field names and values represent the properties for the project.

Note: At the end of this Help topic, find the current list of ElectricFlow intrinsic properties sorted by each object and its associated properties [in table format].

Property names and paths

Properties are named using multi-level paths such as first/second/third, which refers to a property named "third" in a property sheet named "second" in a property sheet named "first." ElectricFlow also supports an equivalent notation using brackets instead of slashes. In this format, slashes are not considered separators when they appear between brackets.

For example, first[second]/third and first[second] [third] both refer to the same property as first/second/third. The bracket [] notation is based on matching brackets. For example, steps[build[1]] refers to a property named "build[1]" inside a property sheet named "steps" (it does not treat "build" as a property sheet containing a property named "1").

Property names/paths have two forms: absolute and relative.

Absolute property paths

Absolute property paths are referenced by a fully-qualified path syntax that begins with a slash character ("/") followed by a top-level name. This path syntax is similar to a file system path specification. The first component after the "/" must be one of several reserved words that select a starting location to look up the property name. For example, consider the property name /server/Electric Cloud/installDirectory. /server means "lookup" starts in the topmost property sheet associated with the ElectricFlow server. This property name refers to the "installDirectory" property inside the server property sheet.

The system defines the following top-level names (absolute property names):

```
/applications/...
```

Start in the property sheet containing all applications. For example, <code>/applications[deploy]</code> refers to an application named "deploy" and <code>/applications[deploy]/processes[install]</code> refers to a process named "install" in that application.

```
/artifacts/...
```

Start in the property sheet containing all artifacts. For example, /artifacts/myGrp:myKey/prop1 refers to the prop1 property on the artifact whose name is "myGrp:myKey".

```
/artifactVersions/...
```

Start in the property sheet containing all artifact versions. For example,

/artifactVersions/myGrp:myKey:1.0-36/prop1 refers to the prop1 property on the artifact version whose name is "myGrp:myKey:1.0-36".

Note: Throughout the API, you can substitute "groupId:artifactKey:version" wherever an artifactVersionName argument can be specified. This is the same if the artifactVersionNameTemplate specified in the artifact is in the form "groupId:artifactKey:version". When the template is different, it is convenient to be able to specify this tuple if you do not know the artifactVersionName.

```
/components/...
```

Start in the property sheet containing all components. For example, /components[warfile] refers to a component named "warfile" and /component[warfile] /processes[backup] refers to a component process named "backup" in that component.

```
/environments/...
```

Start in the property sheet containing all components. For example, /environments [qeserver] refers to an environment named "qeserver".

```
/groups/...
```

Start in the property sheet containing all groups. For example, /groups [dev] refers to the group named "dev".

```
/jobs/...
```

Start in the property sheet containing all jobs. For example, /jobs[ecloud.4096] refers to the job named "ecloud.4096." You can name a job using either its name (as in the preceding example) or using the unique identifier assigned to it by ElectricFlow.

```
/plugins/...
```

Start in the property sheet containing all plugins. For example, /plugins [EC-AgentManagement] refers to the currently promoted plugin named "EC-AgentManagement" and /plugins [EC-AgentManagement] /project/procedures [scpCopyFile] refers to a procedure named "scpCopyFile" in that plugin.

Note: There is a subtle difference between

```
ectool setProperty /plugins/EC-AgentManagement/project/foo 'bar'
and
ectool setProperty /plugins/EC-AgentManagement/foo 'bar'
```

The former places the property on the plugin's project and can be referenced by $f[\mbox{myProject/foo}]$ while the latter places the property on the plugin object and will not be found via $f[\mbox{myPlugin/foo}]$.

```
/projects/...
```

Start in the property sheet containing all projects. For example, <code>/projects[nightly]</code> refers to the project named "nightly," and <code>/projects[nightly]/procedures[main]</code> refers to a procedure named "main" in that project.

```
/repositories/...
```

Start in the property sheet containing all artifact repositories. For example,

/repositories/repo1/prop1 refers to the prop1 property on the repository whose name is "repo1".

```
/resourcePools/...
```

Start in the property sheet containing all resource pools. For example, /resourcePools[linuxA] refers to the resource pool named "linuxA".

```
/resources/...
```

Start in the property sheet containing all resources. For example, /resources [linux1] refers to the resource named "linux1".

```
/server/...
```

Start in the top-level server property sheet. For example, /server/a refers to the server property named "a".

```
/users/...
```

Start in the property sheet containing all users. For example, /users[Bob] refers to the user named "Bob".

```
/workspaces/...
```

Start in the property sheet containing all workspaces. For example, /workspaces[default] refers to the workspace named "default."

Relative property paths

Property names that do not begin with "/" are *relative* and looked up starting in the *current context*. For example, when executing a job step, the current context includes properties defined for the job step, parameters for the current procedure, and global properties on the current job.

Relative property paths are distinguished from absolute property paths because they do not begin with one of the top-level names. To avoid having to construct full property paths, ElectricFlow supports the concept of a relative property path.

In use, the relative property path value is resolved by its context and a defined search order, which results in accessing the value of an absolute fully-specified property value. Contexts and search orders are as follows:

- 1. In a job step context, ElectricFlow searches for relative property paths in the following order:
 - a property on the job step object
 - a property of the parent job step object, which includes parameters of the procedure on which the step is defined
 - · a property on the job object

Notes:

- The search can be enabled or disabled by using the --extendedContextSearch option on the ectool getProperty or setProperty commands. When searching for a property value, disable the search by setting the --extendedContextSearch switch to "false", requiring the property to exist on the job step object or return false. When writing a new value to a property, enable the search by setting the --extendedContextSearch switch to "true", allowing the search for a property within the search order before creating a new one. New properties are created on the job step object.
- o Parameters for subprocedure calls from job steps are searched for in a job step context.
- For procedure parameters for nested subprocedures, properties referenced by parameters are looked up as described [above] for job steps.
- 2. When expanding schedule parameters, ElectricFlow searches the relative property path in the following order:
 - · a property on the procedure being called
 - a property on the project on which the procedure being called is defined
 - · a property on the server
- 3. In a job name template context, ElectricFlow searches for the relative property path as a property on the job.
- 4. In any other context, ElectricFlow searches for the relative property path as a property on the context object.

jobSteps Paths

Some properties can be specified by using the property name and a globally unique identifier for the property, such as jobSteps. The format of the property path is similar to an absolute property path even though the property is not a top-level (absolute) property. For example, you can specify a job step as /jobSteps/<jobStepID>/.

You can also specify a property on a job step as /jobSteps/parent/propertyNamewhere:

- parent is the root parent of the job step that is not visible on the ElectricFlow UI.
- propertyName is the name of the property.

Context-relative "shortcuts" to property paths

A shortcut can be used to reference a property without knowing the exact name of the object that contains the property. You might think of a shortcut as another part of the property hierarchy. Shortcuts resolve to the correct property path even though its path elements may have changed because a project or procedure was renamed. Shortcuts are particularly useful if you do not know your exact location in the property hierarchical tree.

Shortcuts to property paths that provide convenient run-time access include:

```
/myApplication/...
```

Start in the property sheet for the application associated with the current job or job step.

```
/myApplicationTier/...
```

Start in the property sheet for the applicationTier associated with the current job step.

```
/myArtifactVersion/...
```

Start in the property sheet for the artifactVersion associated with the current context. The only context where this property has any value is in the artifactVersionNameTemplate field for the artifact.

```
/myComponent/...
```

Start in the property sheet for the component associated with the current context.

```
/myCredential/...
```

Start in the property sheet for the credential associated with the current job step. This form produces an error if the current job step does not have a credential.

```
/myEvent/...
```

This is a special property used only within an "email notifier." /myEvent/ allows you to refer to fields associated with the notifier itself. You can include these properties, for example, in the text of the email you send out for notification:

/myEvent/notifier - This property contains the name of the notifier that created the notifier event. You created this name when you created the notifier.

/myEvent/entity - This property contains the object where the notifier is attached. Two possible values are "job" or "jobStep," depending on where the notifier is attached.

/myEvent/source - This property contains the name of either the job or the jobStep where the notifier is attached.

/myEvent/type - This property defines whether the notifier is an "On Start" or an "On Completion" notifier. Two possible values are "STARTED" or "COMPLETED".

 $\label{eq:myEvent/time} $$\operatorname{Im}_{\tau} - This $$ property contains the time when the notifier occurred. The time is always specified in GMT and uses a formatted string, for example,$

```
2009-06-11T21:00:56.502Z
```

/myEvent/timeMillis - This property contains the "timestamp" in milliseconds when the notifier occurred. This value is the number of milliseconds since January 1, 1970 GMT. The timeMillis corresponding to the time in the previous example is: 1244754056502

```
/myEnvironment/...
```

Start in the property sheet for the environment associated with the current job or job step.

```
/myEnvironmentTier/...
```

Start in the property sheet for the environmentTier associated with the current job step.

```
/myGate/...
```

/myGate/ is currently limited to myGate.gateType (PRE or POST) or myGate.parentStage. It is useful to include information about the active gate and stage in email notifications templates.

```
/myJob/...
```

Start in the global property sheet for the current job. /myJob/ points directly to the job object so you can reference built-in job properties (jobName, createTime, outcome, and so on) and custom properties. This property sheet can also be used to hold working data that needs to be passed from one step to another within the job.

```
/myJobStep/...
```

Start in the property sheet for the current job step.

```
/myParent/...
```

Start in the global property sheet for the parent job step or job if this is a top-level step. /myParent/points directly to the job or job step object, so you can reference intrinsic or custom properties. For example, you can reference the parameters that were passed to this procedure.

In addition, you can reference a sibling step by calling /myParent/jobSteps/<sibling step name>.

Or, you can create additional properties on /myParent/ to pass information from one step in the procedure to another:

```
step1 calls setProperty /myParent/results 100
step2 can call getProperty /myParent/results
/myParent/.
/myPipeline/...
```

Start in the property sheet for the pipeline associated with the current job or job step. This can be used by a plugin, procedure, or process step executed from a stage task to retrieve property information defined in the pipeline.

```
/myPipelineRuntime/...
```

Start in the property sheet for the pipelineRuntime associated with the current job or job step. This can be used by a plugin, procedure, or process step executed from a stage task to retrieve property information stored on the pipeline runtime

```
/myPipelineStageRuntime/...
```

Start in the property sheet for the pipelineStageRuntime for a pipeline stage where the current job or job step was defined. This can be used by a procedure step executed from a stage task to retrieve information for the summary property of the stage runtime.

```
/myProcedure/...
```

Start in the property sheet for the procedure in which the current job step was defined. If the current job step is executing as part of a nested procedure, /myProcedure/ refers to the innermost nested procedure.

```
/myProcess/...
```

Start in the property sheet for the application or component process in which the current job or job step was defined.

```
/myProcessStep/...
```

Start in the property sheet for the application or component process step in which the current job step was defined.

```
/myProject/...
```

Start in the property sheet for the project in which the current job step was defined. If the job step has nested procedure invocations, this is the project associated with the innermost nested procedure, for example, the project associated with /myProcedure/.

```
/myResource/...
```

Start in the property sheet for the resource assigned to the current job step.

```
/myResourcePool/...
```

Start in the property sheet for the resource pool that provided the resource for the current job step. Returns null if the step did not specify a resource pool.

```
/myStage/...
```

Start in the property sheet for the pipeline stage in which the current job or job step was defined.

```
/myStep/...
```

Start in the property sheet for the current step. "Step" refers to the (static) definition of a step, which is part of a procedure—this is different from a job step, which represents a step when it executes (dynamically) in a job. Use /myJobStep/ to access the job step.

```
/myState/...
```

Start in the property sheet for the state object so you can reference built-in and custom state properties or find parameter values passed to that state. When accessed from a state, /myState/refers to that state. When accessed from a transition, /myState refers to the transition's owning

state. When accessed from a job or job step, /myState/ refers to the state that launched that job as a subjob.

```
/mySubjob/...
```

Start in the property sheet for the subjob so you can reference built-in and custom job properties, parameters passed to the job, and properties on steps within that job. When accessed from a transition, /mySubjob/ refers to the subjob started by the state that owns that transition. This property path is particularly useful in conditions for On Completion transitions because the outcome or other information for the subjob can influence which state the workflow transitions to next.

```
/mySubworkflow/...
```

Start in the property sheet for the subworkflow so you can reference built-in and custom workflow properties. When accessed from a transition, /mySubworkflow/ refers to the subworkflow started by the state that owns that transition. This property path is particularly useful in conditions for On Completion transitions because the active state and other information for the subworkflow can influence which state the workflow transitions to next. You can access information about states and transitions belonging to the workflow by using the path /mySubworkflow/states/someState or /mySubworkflow/states/someState/transitions/someTransition.

```
/myTask/...
```

Start in the property sheet for the pipeline task in which the current job or job step was defined.

```
/myTransition/...
```

Start in the property sheet for the transition object so you can reference intrinsic and custom transition properties. /myTransition/ is accessible only from a transition.

```
/myUser/...
```

This property can be used only if the current session is associated with: the predefined "admin" user, a user defined as "local", or a user defined by a Directory Provider (LDAP or ActiveDirectory). This property cannot be used if the user is a "project principal", which is normally the case when running inside an ElectricFlow step.

For example, with an interactive login you can use:

```
ectool getProperty /myUser/userName (to get the user name of the logged in user, or...) ectool getProperty /myUser/email (to get the email address)
```

```
/myWorkflow/...
```

Start in the property sheet for the workflow object so you can reference built-in and custom workflow properties. When accessed from a state or transition, /myWorkflow/ refers to the "owning" workflow. When accessed from a job or job step, /myWorkflow/ refers to the workflow whose state launched that job as a subjob. You can access information about states and transitions belonging to the workflow by using the path /myWorkflow/states/someState or

/myWorkflow/states/someState/transitions/someTransition.

```
/myWorkflowDefinition/...
```

Starts in the property sheet for the workflow definition object so you can reference built-in and custom workflow properties.

/myWorkspace/...

Start in the property sheet for the workspace associated with the current job step.

Property path shortcuts in ElectricFlow 5.0 and later

The following shortcuts to property paths are available in ElectricFlow 5.0 and later. For more information, go to Context-relative "shortcuts" to property paths on page 790.

Shortcuts	Descriptions	For jobs	For job steps
/myApplication/	Starts in the property sheet for the application associated with the current job or job step.	Yes	Yes
/myApplicationTier/	Starts in the property sheet for the application tier associated with the current job step.	No	Yes
/myComponent/	Starts with the property sheet for the component associated with the current job step.	No	Yes
/myCredential/	Starts with the property sheet for the credential associated with the current job step.	No	Yes
/myEnvironment/	Starts in the property sheet for the environment associated with the current job or job step.	Yes	Yes
/myEnvironmentTier/	Starts in the property sheet for the environment tier associated with the current job step.	No	Yes
/myGate/	/myGate/ is currently limited to myGate.gateType (PRE or POST) or myGate.parentStage. It is useful to include information about the active gate and stage in email notifications templates.	Yes	Yes
/myJob/	Starts in the property sheet for the job associated with the current job or job step.	Yes	Yes
/myJobStep/	Starts in the property sheet for the job step associated with the current job step.	No	Yes
/myParent/	Starts in the global property sheet for the parent job step.	No	Yes

Shortcuts	Descriptions	For jobs	For job steps
/myPipeline/	Start in the property sheet for the pipeline associated with the current job or job step.		Yes
/myPipelineRuntime/	Start in the property sheet for the pipelineRuntime associated with the current job or job step.		Yes
/myPipelineStageRuntime/	Start in the property sheet for the pipelineStageRuntime for a pipeline stage where the current job or job step was defined. This can be used by a procedure step executed from a stage task to retrieve information for the summary property of the stage runtime.		Yes
/myPlugin/	Starts in the property sheet for the plugin associated with the current job.	Yes	No
/myProcedure/	Starts in the property sheet for the procedure in which the current job or job step was defined.		Yes
/myProcess/	Starts in the property sheet for the process in which the current job or job step was defined.		Yes
/myProcessStep/	Starts in the property sheet for the process step in which the current job step was defined.		Yes
/myProject/	Starts in the property sheet for the project in which the current job was defined.	Yes	No
/myResource/	Starts in the property sheet for the resource associated with the current job step.	No	Yes
/myResourcePool/	Starts in the property sheet for the resource pool associated with the current job step.	No	Yes
/myStage/	Start in the property sheet for the pipeline stage in which the current job or job step was defined.	Yes	Yes

Shortcuts	Descriptions	For jobs	For job steps
/myState/	Starts in the property sheet for the state object so you can reference built-in and custom state properties or find parameter values passed to that state.	Yes	Yes
/myStep/	Starts in the property sheet for the step associated with the current job step.	No	Yes
/myTask/	Start in the property sheet for the pipeline task in which the current job or job step was defined.	Yes	Yes
/myWorkflow/	Starts in the property sheet for the workflow object so you can reference built-in and custom workflow properties.	Yes	Yes
/myWorkflowDefinition/	Starts in the property sheet for the workflow definition object so you can reference built-in and custom workflow properties.	Yes	Yes
/myWorkspace/	Starts in the property sheet for the workspace associated with the current job step.	No	Yes

Property name substitutions

Property names can contain references to other properties, which are then substituted into the property name before looking it up. For example, consider the following property name:

```
/myStep/$[/myProcedure/name]
```

If the value of /myProcedure/name is "xyz", the property above is equivalent to /myStep/xyz.

Expandable properties

Property values can contain property references using the "\$[]" notation.

For example:

- 1. Create a property named "foo" with a value of hello $\beta[bar]$.
- 2. Create a property named "bar" with a value of world.
- 3. Reference "foo" (either using foo] or ectool getProperty foo). The value "hello world" is returned.

If you want just the literal value of "foo" (useful in the UI, for example), you can use the expand option in ectool:

```
ectool getProperty foo --expand false. The value "hello $[bar]" will be returned.
```

Properties are expanded by default when you use getProperty or getProperties.

If the value of a property contains "s[]" but you do not want it to be interpreted as a property reference, you can use the expandable option:

```
ectool setProperty symbols '$[!@]#' --expandable false.
```

This option can be toggled in the web UI as well. Properties are expandable by default.

Because you cannot control where your expandable property might be referenced (and therefore which context is used during expansion), we recommend using absolute paths when referencing a property from the value of another property.

In the example above, if you define both "foo" and "bar" as properties on a project "proj1", you might assume there is no problem with the value of "foo". However, if you later reference "foo" from a job under "proj1" (for example, \$[/myProject/foo), foo will be referenced with the job step as its context. Therefore, when the value of "foo" is expanded, you will get a PROPERTY_REFERENCE_ERROR because "bar" is not defined in the context of the job step.

Custom property names and values

The following properties are used by the standard ElectricFlow UI, so you should use these property names whenever possible, and avoid using these names in ways that conflict with the definitions below.

- compiles the number of files compiled during the job step
- diagFile the filename in the top-level directory of the job's workspace, containing diagnostics extracted from the step's log file
- errors the number of errors (compilation failures, test failures, application crashes, and so on) that occurred during the job step. When property errors are set by postp, the step outcome is set to error also.
- tests the number of tests executed by the job step, including successes and failures
- testsSkipped the number of tests skipped during the job step
- warnings the number of warnings that occurred during the job step. When property warnings is set by postp, the step outcome is set to warning also.
- presummary if this property exists, its value is displayed in the "Status" field (on the Job Details page) for this step. This property appears *before* whatever would normally be displayed for status. If the property contains multiple lines separated by newline characters, each line is displayed on a separate line in the status field.
- postsummary if this property exists, its value is displayed in the "Status" field (on the Job Details page) for this step. This property appears *after* whatever would normally be displayed for status. If the property contains multiple lines separated by newline characters, each line is displayed on a separate line in the status field.
- summary if this property exists, its value is displayed in the "Status" field (in the job reports) for this step, replacing whatever would normally be displayed for status. If the property contains multiple lines separated by newline characters, each line is displayed on a separate line in the status field.

The property hierarchy

All ElectricFlow properties fall into a hierarchical structure, and you can reference any property using an absolute path from the root of the hierarchy. This includes properties you define and intrinsic properties defined by ElectricFlow.

For example, each step contains a property "resource" that provides the resource name to use for that step. All steps of a procedure exist as property sheets underneath the procedure. All procedures in a project exist as property sheets underneath the project, and so on.

The examples below illustrate some nesting relationships between objects.

For example, the notation "project/procedures[procedureName]" means each project object contains a property sheet named "procedures" holding all procedures defined within that project. Within the procedures property sheet, there is a nested property sheet for each procedure, named after the procedure. Thus, the name "/projects[a]/procedures[b]" refers to a procedure named "b" contained in a project named "a."

• Each project contains procedures, schedules, credential definitions, workflow definitions, and workflows:

project/procedures[procedureName]
project/schedules[scheduleName]
project/credentials[credentialName]
project/workflowDefinitions[workflowName]
project/workflows[workflowName]

- Each procedure contains steps and parameters. The parameters are "formal parameters," meaning they
 specify parameter names the procedure will accept, whether each parameter is required, and so on:
 procedure/steps[stepName]
 procedure/formalParameters[paramenterName]
- Steps that invoke subprocedures contain parameter values for the subprocedure. These are called "actual parameters" because they provide actual values that will be passed into the subprocedure: step/actualParameters[parameterName]
- Each job contains a collection of job step objects for steps in the outermost procedure, along with parameter values passed into the job when it was invoked: job/jobSteps[stepName] job/actualParameters[parameterName]
- If a job step invokes a nested subprocedure, its property sheet contains parameter values that were
 passed into the nested subprocedure, plus all job steps corresponding to that procedure:
 jobStep/actualParameters[parameterName]
 jobStep/jobSteps[stepName]
- Schedules contain parameter values for the procedures they invoke. These are called "actual parameters" because they provide actual values passed into the procedure: schedule/actualParameters[parameterName]
- Workflow definitions contain state definitions: workflowDefinition/stateDefinitions[stateDefinitionName]
- Transition definitions contain actual parameters: transitionDefinition/actualParameters[parameterName]

- State definitions contain transition definitions, actual parameters, and formal parameters: stateDefinition/actualParameters[parameterName] stateDefinition/formalParameters[parameterName] stateDefinition/transitionDefinitions[transitionDefinitionName]
- Workflows contains states: workflow/states[stateName]
- Transitions contain actual parameters: transition/actualParameters[parameterName]
- States contain transitions, actual parameters, and formal parameters: state/actualParameters[parameterName] state/formalParameters[parameterName] state/transitions[transitionName]

Special property references

ElectricFlow also supports several special property reference forms that are described in the following subsections.

increment

Use this form to increment the value of an integer property before returning its value. For example, suppose property xyz has the value 43. The property reference f[increment xyz] first increments the value of property xyz to 44, then returns 44.

timestamp

Use this form to generate a formatted timestamp value. For example, the property reference f[/timestamp yyyy-MMM-d hh:mm] returns the current time in a form such as "2007-Jun-19 04:36". The pattern following /timestamp specifies how to format the time and defaults to "yyyyMMddHHmm". The pattern follows conventions for the Java class SimpleDateFormat, where various letters are substituted with various current time elements. For more information, go to

http://docs.oracle.com/javase/7/docs/api/java/text/SimpleDateFormat.html for SimpleDateFormat class documentation.

Here are some of the supported substitutions:

```
y - Year, such as "2007" or "07"

M - Month, such as "April", Apr", or "04"

w - Week in year, such as "27"

W - Week in month, such as "3"

D - Day in year, such as "194"

d - Day in month, such as "18"

E - Day in the week, such as "Tuesday" or "Tue"

a - am/pm marker, such as "PM"

H - Hour in day (0-23), such as "7"
```

```
h - Hour in am/pm (1-12), such as "11"

m - Minute in hour, such as "30"

s - Second in minute, such as "15"

S - Millisecond, such as "908"

z - Time zone such as "Pacific Standard Time", "PST", or "GMT-08:00"

Z - Time zone such as "-0800"
```

Repeated letters cause longer forms to be substituted. For example, "yy" substitutes the last 2 digits of the current year, whereas "yyyy" substitutes the 4-digit year number. Single quotes can be used to substitute text directly without the above interpretations.

javascript

This form executes a Javascript code fragment inside the ElectricFlow server and returns the result computed by that code.

For example, \$[/javascript 4*2] returns "8.0". Javascript code can be arbitrarily long and include multiple statements. The value of the last statement is returned by the property reference.

Javascript code executes in an interpreter that provides access to ElectricFlow properties:

 The normal way to access property values is through Javascript objects. Global Javascript objects named "server," "projects," and so on, exist and correspond to all top-level objects in an absolute property path.

For example, there is a Javascript object named <code>server</code> that corresponds to the path <code>/server</code>, and a Javascript object named <code>myJob</code> that corresponds to the path <code>/myJob</code>. You can use either "." or "[]" notation to access properties within the object.

For a more complete list of top-level objects, see the "Absolute property paths" section.

Examples:

```
$[/javascript myJob.mightExist]
this is a safe way to refer to an optional parameter

$[/javascript (myJobStep.errors > 0) ? "step failed" : "no errors"]
test a value and return another string based on the result

$[/javascript server.settings.ipAddress]
refer to a property in a nested property sheet

$[/javascript server["Electric Cloud"].installDirectory]
if the property name has a space, use [""] notation

$[/javascript server["Electric Cloud"]["installDirectory"]]
do not use the "." in front of the "["
```

2. You can also call the function, "getProperty". It takes a property name as argument and returns the value of that property. The case where this is most useful is when you want to access another special property reference.

Examples:

```
$[/javascript getProperty ("/timestamp yyyy-MMM-d hh:mm")]
returns the current time

$[/javascript getProperty ("/increment /myJob/jsCount")]
this is the only way to change a property via javascript
```

For example, consider the following property reference:

```
$[/javascript (getProperty("/myJobStep/errors") > 0) ? "step failed" : "no
errors"]
```

This example returns "step failed" if the property "errors" on the current job step had a value greater than zero, and it returns "no errors" otherwise.

- 3. If calling the "setProperty" function, similar to getProperty, there are two variations on the function:
 - A global function variation that uses the current context object.
 The global function takes 2 or 3 arguments. The 3-argument version takes a context object, a path, and a value. The 2-argument version omits the context object and uses the current context object (for example, job step).

Example:

```
setProperty (myProject, "foo", "bar")
would set the value of the "foo" property on the current project to "bar".
```

An object function variation that can be called on objects, which uses that object as a context object.
 The object function takes two arguments: a "path" and a "value".

Example:

```
server.setproperty ("foo", "bar")
would set the "foo" property on the server object to the value "var".
```

ElectricFlow Intrinsic Properties listed by object type

This section describes the intrinsic properties for some (but not all) of the ElectricFlow objects. For each object type in the following list, there is a link to a table that includes a list of intrinsic properties applicable to that object and provides the property type and description.

```
acl repository artifact resource
```

artifactVersi resourcePool

on

credential resourceUsage

directoryProv schedule

ider

emailConfig server
emailNotifier state

formalParamet stateDefinition

er

gateways step

group transition

job transitionDefin

ition

jobStep

user

logEntry workflow

procedure workflowDefinit

ion

project workspace

property zones

propertySheet

ElectricFlow Property Type definitions

This table provides Property Type definitions for each property listed in the following tables. (A Property Type precedes each property description in the Description column.)

Туре	Definition
boolean	One of two possible values - either <i>true</i> or <i>false</i> . These values are more frequently represented by the numbers "0" and "1", where "0" equals false and "1" equals true.
date	A millisecond precision UTC date in ISO 8601 form: [YYYY] - [MM] - [DD] T [hh]: [mm] Z For example, 2007-06-19T04:36:22.000Z
id	Each time an object is created, ElectricFlow generates a unique ID number for that object.

Туре	Definition
name	This is a unicode string value with a maximum of 255 characters.
number	This is a simple integer numeric value.
reference	This property refers to another object.
string	This is a unicode string value with a maximum CLOB size of the database, but only the first 450 characters are indexed, which means a defined search will not "see" beyond the first 450 characters.

ElectricFlow Object / Property tables

In the following tables, the Description column displays the property type preceding the property description.

Object	Type:	acl
Object	iypc.	uo.

Description: An acl is an Access Control List.

Property Name	Description
aclId	id: The unique identifier for this acl object. Other objects can refer to this acl by its ID.
inheriting	boolean: If true, the ACL inherits ACEs from the ACL's parent.

Object Type: acl

Description: An acl is an Access Control List.

Property Name	Description	
ownerType	This is any type of object the ACL controls and can be any of the objects listed below.	
	ace acl admin artifact artifactVersion event formalParameter group job jobStep license notifier procedure procedureStep	project property propertySheet repository resource resourcePool server schedule systemObject user userSettings workspace plugin
parentId	id: The parent ACL.	

[back to top]

Object Type: artifact

Description: An *artifact* is an object that contains zero or more artifact versions.

An artifact has two purposes:

- 1. To group artifact versions and provide a template for naming the versions
- 2. To restrict who can publish artifact versions, based on groupId:artifactKey

Property Name	Description
acl	reference:acl
artifactId	id: The artifact's ID number.

Object Type: artifact

Description: An *artifact* is an object that contains zero or more artifact versions. An artifact has two purposes:

- 1. To group artifact versions and provide a template for naming the versions
- 2. To restrict who can publish artifact versions, based on <code>groupId:artifactKey</code>

Property Name	Description	
artifactKey	string: User-specified identifier for this artifact. This field may consist of alphanumeric characters, spaces, underscores, hyphens, and periods.	
artifactName	name: The name of this artifact.	
artifactVersionNameTempl ate	name: The template for artifact version names published to this artifact.	
createTime	date: The time when this object was created.	
description	string: A user-specified text description of the object.	
groupId	id: A user-generated group name for this artifact. This field may consist of alphanumeric characters, spaces, underscores, hyphens, and periods.	
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.	
modifyTime	date: The time when this object was last modified.	
owner	name: The person (user name) who created the object.	
propertySheet	reference: propertySheet	

[back to top]

Object Type: artifactVersion

Description: An artifact version is an object that represents a user-defined unit of related files typically produced by one job and consumed by one or more other jobs.

Property Name	Description	
acl	reference: acl	
artifactKey	string: User-specified identifier for this artifact. This field may consist of alphanumeric characters, spaces, underscores, hyphens, and periods.	
artifactName	name: The name of the artifact.	
artifactVersionId	id: The ElectricFlow-generated ID number for this artifact version.	
artifactVersionName	name: The name of the artifact version.	
artifactVersionState	<pre>string: Possible values are: available publishing unavailable</pre>	
buildNumber	number: User-defined build number component of the version attribute for the artifact version.	
createTime	date: The time when this object was created.	
description	string: A user-specified text description of the object.	
groupId	id: A user-generated group name for this artifact. This field may consist of alphanumeric characters, spaces, underscores, hyphens, and periods.	
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.	
majorMinorPatch	<pre>string: major.minor.patch component of the version attribute for the artifact.</pre>	
modifyTime	date: The time when this object was last modified.	
owner	name: The person (user name) who created the object.	

Object Type: artifactVersion

Description: An artifact version is an object that represents a user-defined unit of related files typically produced by one job and consumed by one or more other jobs.

Property Name	Description	
propertySheet	reference: propertySheet	
publisherJobId	id: The ElectricFlow-generated ID number for the job that published the artifact version.	
publisherJobName	string: The name of the job that published the artifact version.	
publisherJobStepId	id: The ElectricFlow-generated ID number for the job step that published the artifact version.	
qualifier	string: User-defined qualifier component of the version attribute for the artifact.	
repositoryName	name: The name of the artifact repository.	
version	string: An artifact version specification uses the following form: major.minor.patch.qualifier.buildNumber	

[back to top]

Object Type: credential

Description: In ElectricFlow, a *credential* is an object that stores a username and password for later use.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
credentialId	id: The credential's ID number.
credentialName	name: The name of this credential.

Object Type: credential

Description: In ElectricFlow, a *credential* is an object that stores a username and password for later

Property Name	Description
description	string: A user-specified text description of the object.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
password	string: The password corresponding to the user name and this credential.
projectName	name: The name of the project that contains this credential.
propertySheet	reference: propertySheet
userName	name: A saved string that represents the name portion of a credential, typically a user account name.

[back to top]

Object Type: directoryProvider

Description: A directoryProvider is an object containing information about an external directory service (LDAP or ActiveDirectory).

Property Name	Description
acl	reference: acl
commonGroupNameAttribute	string: The attribute in a group record that contains the common group name. If specified, this name is used only when searching for groups from an external provider.

Object Type: directoryProvider

Description: A directoryProvider is an object containing information about an external directory service (LDAP or ActiveDirectory).

Property Name	Description
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
directoryProviderId	id: The ID of the directory provider.
domainName	string: The domain name from which Active Directory servers are automatically discovered.
emailAttribute	string: The attribute in a user record that contains the user's email address. If the attribute was not specified, the account name and domain name are concatenated to form an email address.
enableGroups	boolean: Determines whether or not external groups are enabled for the directory provider. Defaults to "true".
fullUserNameAttribute	name: The attribute in a user record that contains the user's full name (first and last) for display in the UI.
groupBase	string: This string is prepended to the basedn to construct the directory DN that contains group records.
groupMemberAttributes	string: A comma-separated attribute name list that identifies a group member.
groupMemberFilter	string: Two common forms of group record in LDAP directories: POSIX style groups where members are identified by account name, and groupOfNames or uniqueGroupOfNames records where members are identified by the full user DN. Both forms are supported, so the query is passed to parameters: "{0}" is replaced with the full user record DN, and "{1}" is replaced with the user's account name.
groupNameAttribute	name: The group record that contains the name of the group.

${\tt Object\ Type:\ directory Provider}$

Description: A directoryProvider is an object containing information about an external directory service (LDAP or ActiveDirectory).

Property Name	Description
groupSearchFilter	string: This LDAP query is performed in the context of the groups directory to enumerate group records.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
managerDn	name: The DN of a user who has read-only access to LDAP user and group directories.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet
providerIndex	number: The index that specifies the search order across multiple directory providers. For example: 2 LDAP providers, one with index "0" and one with index "1" means the providers will be searched in that numerical order.
providerName	name: This human-readable name will be displayed in the user interface to identify users and groups that come from this provider.
providerType	string: <ldap activedirectory></ldap activedirectory>
realm	string: An identifier (string)used for LDAP directory providers so users and groups (within LDAP) can be uniquely identified in "same name" collisions across multiple directory providers. The realm is appended to the user or group name when stored in the ElectricFlow server. For example, <user>@dir (where the realm is set to "dir").</user>
url	<pre>string: The server URL is in the formprotocol://host:port/basedn. Protocol is either Idap or Idaps (for secure LDAP).</pre>

Object Type: directoryProvider

Description: A directoryProvider is an object containing information about an external directory service (LDAP or ActiveDirectory).

Property Name	Description
userBase	string: This string is prepended to the basedn to construct the directory DN that contains user records.
userNameAttribute	name: The attribute in a user record that contains the user's account name.
userSearchFilter	string: This LDAP query is performed in the context of the user directory to search for a user by account name. The string "{0}" is replaced with the user's login ID. Typically, the query compares a user record attribute with the substituted user login ID.
userSearchSubtree	boolean: If true, the subtree below the user base is searched recursively.
useSSL	boolean: This flag is used to specify SSL to communicate with your Active Directory servers.
	Note: Transport Layer Security (TLS) has replaced Secure Sockets Layer version 3.0 (SSLv3) on the ElectricFlow web server and the ElectricFlow server.

[back to top]

Object Type: emailConfig

Description: An *emailConfig* is an object that stores information created and used to communicate with the email server.

Property Name	Description
acl	reference: acl
configName	string: The name of the email configuration.

Object Type: emailConfig

Description: An *emailConfig* is an object that stores information created and used to communicate with the email server.

Property Name	Description
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
emailConfigId	id: The ElectricFlow-generated ID for the email configuration.
emailConfigName	string: The name of the email configuration.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
mailFrom	string: The email address used as the email sender address for notifications.
mailHost	string: The name of the email server host.
mailPort	number: The port number for the mail server. The protocol software determines the default value (25 for SMTP and 465 for SSMTP).
mailProtocol	string: This is either SSMTP or SMTP (not case sensitive). Default is SMTP.
mailUser	name: An individual or a generic name like "ElectricFlow" the name of the email user on whose behalf ElectricFlow sends email notifications.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet

[back to top]

Object Type: emailNotifier

Description: An *emailNotifier* is an object that stores information created and used to notify users about various types information, including status.

Property Name	Description
acl	reference: acl
condition	string: Only send mail if the condition evaluates to "true". The condition is a string subject to property expansion. The notification will NOT be sent if the expanded string is "false" or "0". If no condition is specified, the notification is ALWAYS sent.
configName	string: The name of the email configuration.
container	id: The object ID to which the email notifier is attached (for example: procedure-123).
containerType	string: The type of object that the notifier is attached to (procedure, step, job, jobstep).
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
destinations	string: A space-separated list of valid email addresses, email aliases, or ElectricFlow user names, or a string subject to property expansion that expands into such a list.
emailNotifierId	id: The ElectricFlow-generated ID for the email notifier.
eventType	string: <onenter onstart oncompletion> onStart triggers an email notification when the job or job step begins. onCompletion, default, triggers an email notification when the job finishes, no matter how it finishes.</onenter onstart oncompletion>
formattingTemplate	string: The email address used as the email sender address for notifications.

Object Type: emailNotifier

Description: An *emailNotifier* is an object that stores information created and used to notify users about various types information, including status.

Property Name	Description
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
notifierName	name: The name of the email notifier.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet

[back to top]

Object Type: formalParameter

Description: A *formalParameter* is a parameter expected by a procedure, including its name, a default value, and an indication of whether the parameter is required. Formal parameters are different from "actual parameters"--- formal parameters define the type of parameters a procedure is expecting, and actual parameters provide values to use at run-time.

Property Name	Description
container	string: An object ID for a "container" that contains formal parameters.
containerType	string: The type of object containing the formal parameter.
createTime	date: The time when this object was created.
defaultValue	string: The formalParameter's default value.

Object Type: formalParameter

Description: A *formalParameter* is a parameter expected by a procedure, including its name, a default value, and an indication of whether the parameter is required. Formal parameters are different from "actual parameters"--- formal parameters define the type of parameters a procedure is expecting, and actual parameters provide values to use at run-time.

Property Name	Description
description	string: A user-specified text description of the object.
expansionDeferred	boolean: Determines whether or not the property expansion is deferred.
formalParameterId	id: This is this formal parameter's ID.
formalParameterName	name: This is this formal parameter's name.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
required	boolean: If true, a value for this parameter must be supplied when the procedure is called.
type	name: The custom type of formalParameter.

[back to top]

Object Type: gateway

Description: To communicate with a resource, workspace, or artifact repository server in another zone, a "gateway" must be created. A gateway object contains two resource (agent) machines, for example, GatewayResource1 and GatewayResource2—each configured to communicate with the other. One gateway resource resides in the *source* zone and the other in the *target* zone. A gateway is bidirectional and informs the ElectricFlow server that each gateway machine is configured to communicate with its other gateway machine (in another zone).

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
gatewayDisabled	boolean: If set to 1 (true), the gateway is disabled.
gatewayId	id: The unique ElectricFlow-generated ID for this gateway.
gatewayName	string: The name of the gateway.
hostName1	string: The agent host name where Resource1 resides. This host name is used by Resource2 to communicate with Resource1. Do not specify this option is you want to use the host name from Resource1's definition.
hostName2	string: The agent host name where Resource2 resides. This host name is used by Resource1 to communicate with Resource2. Do not specify this option is you want to use the host name from Resource2's definition.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
port1	number: The port number used by Resource1 - defaults to the port number used by the resource.

Object Type: gateway

Description: To communicate with a resource, workspace, or artifact repository server in another zone, a "gateway" must be created. A gateway object contains two resource (agent) machines, for example, GatewayResource1 and GatewayResource2—each configured to communicate with the other. One gateway resource resides in the *source* zone and the other in the *target* zone. A gateway is bidirectional and informs the ElectricFlow server that each gateway machine is configured to communicate with its other gateway machine (in another zone).

Property Name	Description
port2	number: The port number used by Resource2 - defaults to the port number used by the resource.
propertySheet	reference: propertySheet
resourceName1	string: The name of your choice for the first of two required gateway resources. Do not include "spaces" in a resource name.
resourceName2	string: The name of your choice for the second of two required gateway resources. Do not include "spaces" in a resource name.

[back to top]

Object Type: **group**

Description: This is any *group* of users known to the ElectricFlow server, including groups defined locally within the server and those groups defined in external repositories such as LDAP or ActiveDirectory.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
groupId	id: The unique ElectricFlow-generated group ID.
groupName	name: This is this group's name.

Object Type: **group**

Description: This is any *group* of users known to the ElectricFlow server, including groups defined locally within the server and those groups defined in external repositories such as LDAP or ActiveDirectory.

Property Name	Description
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
mutable	boolean: If true, the member list of this group is editable within ElectricFlow via the web UI or the modifyGroup API.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet
providerName	name: The name of the directory provider that controls this group.

[back to top]

Object Type: **job**

Property Name	Description
abortStatus	string: If set, indicates the step was aborted. Values will be either ABORT or FORCE_ABORT - indicating how the step was aborted.
abortedBy	name: This is the user who issued the "abort."

Object Type: **jOb**

Property Name	Description
acl	reference: acl
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter, which is passed to a procedure when it is invoked. Actual parameters can be provided for jobs and nested subprocedures within a job. Actual parameters are different from "formal parameters"- formal parameters define the parameters a procedure is expecting, and actual parameters provide values to use at run-time.
callingState	string: The full property path to the "calling state", which can appear on subjobs and subworkflows of a workflow.
callingStateId	id: The ElectricFlow-generated ID number for the calling state.
combinedStatus	Provides more inclusive step status output - the resulting query output may contain up to three sub-elements: status message properties
createTime	date: The time when this object was created.
credentialName	name: The name of the credential being used for impersonation when the job runs commands on a resource.
deleted	The object was marked for background deletion. Possible values are "0" or "1". Default is "0" (not set).
directoryName	name: The name of this job's directory within each workspace for this job.

Object Type: **jOb**

Property Name	Description
elapsedTime	number: The number of milliseconds between the start and end times for the job.
errorCode	errorCode: When the outcome is error, this property displays the error code, identifying which error occurred.
errorMessage	string: When the outcome is error, this property displays the error description.
external	boolean: If 'true", the job is external.
finish	date: The time this job completed.
jobId	id: This is this job's ID number, which is a UUID.
jobName	name: This is this job's name.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
launchedByUser	name: The name of the user or project principal that explicitly launched the job. This property is blank when the job is launched by a schedule.
licenseWaitTime	The sum of time all job steps had to wait for a license.
liveProcedure	name: The current procedure name from which this job was created – if the procedure was renamed since the job launched, this will be the procedure's new name, and if the procedure was deleted, this will be null.

Object Type: **job**

Property Name	Description
liveSchedule	name: The current schedule name that launched this job – if the schedule was renamed since the job was launched, this will be the schedule's new name, and if the schedule was deleted, this will be null.
modifyTime	date: The time when this object was last modified.
outcome	The overall result of the job: success, warning, error, or skipped.
owner	name: The person (user name) who created the object.
priority	Values can be low, normal (default), high, or highest.
procedureName	name: The name of the procedure that defines the job's steps.
projectName	name: The name of the project that contains this job.
propertySheetId	reference: propertySheet
resourceWaitTime	The sum of time all job steps had to wait for a resource. This could indicate that eligible resources for the step have reached their step limit, are in-use but the step requires a resource exclusively, or resources are down.
runAsUser	name: The name of the user being impersonated in this job.
scheduleName	name: The schedule name that launched this job – this field differs from liveSchedule in that it is written at job creation time only, and not changed, even if the schedule is renamed or deleted.
start	date: The time when this job began executing.

Object Type: **jOb**

Description: A *job* is an ElectricFlow structure associated with invoking a procedure. A new job is created each time a procedure begins to execute. The job keeps track of all data associated with the procedure's execution, such as the running time of each step and any errors that may occur during the step. ElectricFlow retains job information after the job completes so you can examine what occurred.

Property Name	Description
status	Possible values are:
	pending - the job was created, but it is waiting for prerequisite steps to complete
	runnable - the job step is waiting for a resource
	scheduled - the job was assigned a resource, but the command has not started running
	running - the job/job step is running a command on the assigned resource
	completed - the job/job step has completed
totalWaitTime	The total sum of license, resource, a precondition, and workspace wait times job steps were restricted and had to wait to process.
workspaceWaitTime	The sum of time all job steps had to wait for a workspace.

[back to top]

Object Type: jobStep

Property Name	Description
abortStatus	string: If set, indicates the step was aborted. Values will be either ABORT or FORCE_ABORT - indicating how the step was aborted.
abortedBy	name: The user who issued the "abort."

Property Name	Description
acl	reference: acl
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter, which is passed to a procedure when it is invoked. Actual parameters can be provided for jobs and nested subprocedures within a job. Actual parameters are different from "formal parameters"- formal parameters define the parameters a procedure is expecting, and actual parameters provide values to use at run-time.
alwaysRun	boolean: If true, this step runs even if the job is aborted before the step completes. Note that a "force abort" will abort an alwaysRun step.
assignedResourceName	name: The name of the resource assigned to this step by the step scheduler.
broadcast	boolean: If true, this step runs on all resources in a pool.
combinedStatus	Provides more inclusive step status output - the resulting query output may contain up to three sub-elements: status message properties
command	string: This property specifies the command this step runs.
condition	string: If this element is not present, the event is ALWAYS triggered. If specified, the event is triggered only if the value of the condition argument is TRUE; if not true, a boolean value of "false" or "0" was used. Condition arguments can be a literal, a fixed value string, or a string subject to property expansion.
conditionExpanded	boolean: The result of the expansion on the step condition.
createTime	date: The time when this object was created.

Property Name	Description
delayUntil	For a step that was rescheduled due to a resource or workspace problem, this is the next time the step will be eligible to run.
elapsedTime	number: The number of milliseconds between the start and end times for the <code>jobStep</code> .
errorCode	errorCode: This property appears when the outcome is error and displays the error code, identifying which error occurred.
errorHandling	This is the error handling policy copied from the procedure step and indicates how the step responds to errors:
	failProcedure - The current procedure continues, but the overall status is error (default).
	abortProcedure - Aborts the current procedure, but allows already- running steps in the current procedure to complete.
	abortProcedureNow - Aborts the current procedure and terminates running steps in the current procedure.
	abortJob - Aborts the entire job, terminates running steps, but allows alwaysRun steps to run.
	abortJobNow - Aborts the entire job and terminates all running steps, including alwaysRun steps.
	failProcedure - The current procedure continues, but the overall status is error (default).
	ignore - Continues as if the step succeeded.
errorMessage	string: When the outcome is error, this property displays an error message description.
exclusive	boolean: If true, this step acquires and retains its resource exclusively.

Property Name	Description
exclusiveMode	string: Possible values are: none, job, step, call
exitCode	number: This is this step's exit code.
external	boolean: If "true", the job is external.
finish	date: The time when this job step completed.
hostName	name: The name of the host where this step was invoked (copied from the resource)
job	id: This is the ID number of the job that owns the job step. The string representation of the job is its name, so \$[job] evaluates to the job name, but \$[job/foo] is also legal and refers to the foo property of the job.
jobId	id: This is this job's ID number, which is a UUID.
jobName	name: This is the name of this step's job.
jobStepId	id: This is this step's ID number.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
licenseWaitTime	The length of time this job step had to wait to process because the license limit was reached or exceeded.
liveProcedure	string: The current procedure name for the procedure step from which the job or job step was created – if the procedure step was renamed since the job or job step was launched, this is the procedure step's new name, and if the procedure step was deleted, this will be null.

Property Name	Description
liveProcedureStep	name: This property shows the current procedure name for the procedure step from which this job step was created – if the procedure step was renamed since the job was launched, this will be the procedure step's new name, and if the procedure step was deleted, this will be null.
logFileName	name: The name of the log file produced by this step, relative to the job's workspace directory.
modifyTime	date: The time when this object was last modified.
outcome	<pre>string: The overall result of the job - possible values are: success, warning, error, or skipped</pre>
owner	name: The person (user name) who created the object.
parallel	boolean: If true,this step runs in parallel with other adjacent steps also marked to run in parallel.
postExitCode	number: This is this step's post processor exit code.
postLogFileName	name: The log file name produced by this step's post processor.
postProcessor	<pre>string: The post processor name used to gather information about this step.</pre>
precondition	string: By default, if a step has no precondition, it will run when scheduled. Set this property to make a step wait until one or more dependent conditions are met. When a job step is eligible to transition from pending to runnable, a precondition is evaluated. A precondition is a fixed text or text embedding property reference that is evaluated to TRUE or FALSE. An empty string, a \"0\" or \"false\" is interpreted as FALSE. Any other result string is interpreted as TRUE. The step will block until the precondition is TRUE.

Object Type: jobStep

Description: A *jobStep* is an instance of a step that ran - a single invocation of a command on a resource or a single call to a subprocedure within the context of a *job*.

Property Name	Description
procedureName	name: The name of the procedure that contains this jobStep.
projectName	name: The name of the project that contains this jobStep.
propertySheet	reference: propertySheet
releaseExclusive	boolean: A Boolean value indicating whether this step should release its resource upon completion.
releaseMode	string: Possible values are: none, release, releaseToJob
resourceName	name: The name of the resource or pool this step should use to run on.
resourceSource	This property indicates whether the resource for the job step is explicitly specified by the step or was defaulted from the procedure: procedure or procedureStep.
resourceWaitTime	The length of time this job step stalled because it could not get a resource. This could indicate that eligible resources for the step have reached their step limit, are in-use but the step requires a resource exclusively, or resources are down.
retries	number: This is a number - the number of retries before the request times out.
runAsUser	name: The name of the user being impersonated in this job step.
runnable	date: The time when the step became runnable.
runTime	number: The number of milliseconds the step command spent running on a resource.

Object Type: jobStep

Description: A *jobStep* is an instance of a step that ran - a single invocation of a command on a resource or a single call to a subprocedure within the context of a *job*.

Property Name	Description
shell	name: The shell used to execute the step's commands on a resource. The script name is inserted into the command at the position of a " {0}" marker in the command, or appended as a final argument if no marker is present. A filename suffix adjacent to the marker will be appended to the script name. For more information on shells, see the Shell definition in the Step Help topic.
start	date: The time when this job step began executing.
status	Possible values are: pending - the job step was created, but it is waiting for prerequisite steps to complete runnable - the job step is waiting for a resource scheduled - the job step was assigned a resource, but the command has not started running running - the job/job step is running a command on the assigned resource completed - the job/job step has completed
stepName	name: This is this step's name.
subprocedure	name: The nested procedure name to call when this step executes.
subproject	name: This property appears if the subprocedure element is present and specifies the project to which the subprocedure belongs (by default, the current project is used.)
timeLimit	number: The maximum length of time this step is allowed to run.
timeout	date: The time when this job step will be automatically aborted if it has not yet completed.

Object Type: jobStep

Description: A *jobStep* is an instance of a step that ran - a single invocation of a command on a resource or a single call to a subprocedure within the context of a *job*.

Property Name	Description
totalWaitTime	The sum of resource, workspace, and license wait times for this job step.
waitTime	number: The number of milliseconds the step spent between runnable and running (for example, waiting for a resource).
workingDirectory	name: The name of this step's working directory.
workspaceName	name: The workspace name used by the jobStep object.
workspaceWaitTime	The time this job step had to wait because no workspace was found or available.

Object Type: logEntry

Description: A *logEntry* is an object containing various types of information available in a log file generated from anywhere in the ElectricFlow system.

The *Event Log* can be configured for auto-deletion. By default, Event Log retention is 30 days. The ElectricFlow server automatically marks old log entries for deletion and the background deleter cleans out old log entries.

To change the default settings, go to Administration > Server > Settings and modify the Event log retain time and the Maximum background delete delay fields

Note: Setting the "retain" period to "0" disables the automatic deletion mechanism, allowing you to use an external cleanup script to implement a custom cleanup policy.

Property Name	Description
category	(currently not used)
container	string: Typically, this is the type and name of the workflow or job with a corresponding ID.
containerName	name: The name of the container.
containerType	string: The type of object the log entry pertains to (for example, procedure, job step, step).
deleted	byte: The object was marked for background deletion. Possible values are "0" or "1". Default is "0" (which means "deleted" is not set).
logEntryId	id: The log entry ElectricFlow-generated ID.
message	string: The message text will either be informational or display the warning or error message. The message may contain important information about a resource or workspace issue.
principal	string: The user or project principal from the session that was active when the event occurred.

Object Type: logEntry

Description: A *logEntry* is an object containing various types of information available in a log file generated from anywhere in the ElectricFlow system.

The *Event Log* can be configured for auto-deletion. By default, Event Log retention is 30 days. The ElectricFlow server automatically marks old log entries for deletion and the background deleter cleans out old log entries.

To change the default settings, go to Administration > Server > Settings and modify the Event log retain time and the Maximum background delete delay fields

Note: Setting the "retain" period to "0" disables the automatic deletion mechanism, allowing you to use an external cleanup script to implement a custom cleanup policy.

Property Name	Description
severity	string: Severity can be either TRACE, DEBUG, INFO, WARN, ERROR
subject	string: The object associated with the message.
subjectName	name: The name of the object associated with the message.
subjectType	string: (similar to container) Refers to the object the event concerns. This may be the same as the container, or it may be a different object that is related to the event in some manner.
time	string: The time the event was logged.

Object Type: procedure

Description: A *procedure* describes a series of actions to be performed on one or more resources. A procedure contains steps and properties and can define formal parameters.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
credentialName	name: The name of the credential assigned to this job step.
description	string: A user-specified text description of the object.
jobNameTemplate	name: The template used to name jobs created from this procedure.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
procedureId	id: This is this procedure's ID number.
procedureName	name: This is this procedure's name.
projectName	name: The name of the project that contains this procedure.
propertySheet	reference: propertySheet
resourceName	name: The name of the resource or pool this procedure should use to run on.
workspaceName	name: The workspace name used by the procedure object.

Object Type: **project**

Description: A *project* is an object used in ElectricFlow to organize information. A project contains procedures, schedules, credentials, and properties.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
credentialName	name: The name of the credential assigned to this project.
deleted	The object was marked for background deletion. Possible values are "0" or "1". Default is "0" (not set).
description	string: A user-specified text description of the object.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
pluginName	name: The name of the plugin associated with this project.
projectId	id: This is this project's ID.
projectName	name: The name of the project.
propertySheet	reference: propertySheet
resourceName	name: The name of the resource.
workspaceName	name: This is the workspace name used by the project.

Object Type: **property**

Description: A *property* is a string value with a name. Properties can have arbitrary names and values. You can attach properties to any object in the ElectricFlow system, such as a project, procedure, or job. After a property is created, it is stored in the ElectricFlow database. You can retrieve and/or modify the value later, and you can delete properties you no longer need. Properties provide a flexible and powerful mechanism to manage data about your builds.

IMPORTANT: Note: The names "properties" and "project" are not valid property names.

Property Name	Description
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
expandable	boolean: If set to true, the property value will undergo string expansion when it is retrieved.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
canonicalPath	string: The path that specifies the property object.
propertyId	id: This property's ID number.
propertyName	name: This property's name.
propertySheetId	reference: propertySheet - If the property is a nested property sheet, the propertySheet refers to the nested sheet.
value	string: The property or actual parameter's value - if this property is a string property.

Object Type: propertySheet

Description: A property value can be a simple string or a nested *propertySheet* containing its own properties. Property sheets can be nested to any depth, which allows you to create hierarchical collections of information.

Most objects have an associated property sheet that contains "custom properties" created by user scripts.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
propertySheetId	id: This property sheet's ID.

[back to top]

Object Type: repository

Description: A *repository* is an object that stores artifact versions. This object primarily contains information on how to connect to a particular artifact repository. Similar to steps in a procedure, repository objects are in a user-specified order. When retrieving artifact versions, repositories are queried in this order until one containing the desired artifact version is found.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
description	string: A user-specified text description for this object.

Object Type: repository

Description: A *repository* is an object that stores artifact versions. This object primarily contains information on how to connect to a particular artifact repository. Similar to steps in a procedure, repository objects are in a user-specified order. When retrieving artifact versions, repositories are queried in this order until one containing the desired artifact version is found.

Property Name	Description
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet
repositoryDisabled	boolean: Determines whether the repository is disabled. Default is "false".
repositoryId	id: The ElectricFlow-generated ID number of the repository.
repositoryIndex	number: The order of the repository, within a list of repositories.
repositoryName	name: The name of the repository.
url	string: The server URL is in the form protocol://host:port/. Typically, the repository server is configured to listen on port 8200 for https requests, so a typical URL looks like https://host:8200/.
zoneName	name: The name of the zone where this repository is or will reside.

Description: A *resource* is associated with a server machine available to ElectricFlow for running steps. A resource has a logical name in addition to a host name. A resource can be associated with one or more pools. The resource may be a "proxy target" machine that communicates with the ElectricFlow server though an ElectricFlow agent proxy machine.

Property Name	Description
acl	reference: acl
agentState	string: Specific information about an agent, including the state of the agent. Possible values are: unknown alive down
artifactCacheDirectory	string: The directory on the agent host where retrieved artifacts are stored.
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
exclusiveJobId	id: The ID number of the job that owns this resource, which occurs when one of the job's steps requests exclusive use of the resource for the duration of the job.
exclusiveJobName	name: The name of the job that owns this resource, which occurs when one of the job's steps requests exclusive use of the resource for the duration of the job.
exclusiveJobStepId	id: The ID number of the job step that owns this resource, which occurs when one of the steps request exclusive use of the resource for the duration of the job.
exclusiveJobStepName	name: The name of the job step that owns this resource, which occurs when one of the steps request exclusive use of the resource for the duration of the job.
gateways	The gateway names associated with this resource.

Description: A *resource* is associated with a server machine available to ElectricFlow for running steps. A resource has a logical name in addition to a host name. A resource can be associated with one or more pools. The resource may be a "proxy target" machine that communicates with the ElectricFlow server though an ElectricFlow agent proxy machine.

Property Name	Description
hostName	name: The computer name or IP address for the machine containing the ElectricFlow agent for this resource.
hostOS	string: The full name of the host operating system, plus its version. However, if this host is a proxy, "proxied" appears as the host description/name.
hostPlatform	string: Examples for "platform" are: Windows, Linux, HPUX, and so on. However, if this host is a proxy, "proxied" appears as the hostPlatform name.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
lastRunTime	date: The most recent time a job step ran on the resource.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
pools	string: A space-separated list of one or more pool names where this resource is a member. Steps defined to run on a resource pool will run on any available member (resource) in the pool.
port	number: The port number for this resource.
propertySheet	reference: propertySheet
proxyCustomization	string: This property displays the customization of the ElectricFlow proxy agent.

Description: A *resource* is associated with a server machine available to ElectricFlow for running steps. A resource has a logical name in addition to a host name. A resource can be associated with one or more pools. The resource may be a "proxy target" machine that communicates with the ElectricFlow server though an ElectricFlow agent proxy machine.

Property Name	Description
proxyHostName	name: The name of the ElectricFlow agent being used to proxy to another agent, the proxy target.
proxyPort	number: The port number of the ElectricFlow agent being used to proxy to another agent, the proxy target.
proxyProtocol	name: The name of the proxy agent protocol used for communication from the ElectricFlow proxy agent to the proxy target - default is SSH.
repositoryNames	A "new line" separated list of repository names.
resourceDisabled	boolean: A Boolean value indicating whether this resource was disabled.
resourceId	id: This is this resource's ID.
resourceName	name: The name of the resource or pool.
shell	name: The shell used to execute the step's commands on a resource. If no shell was defined on a step, the shell defined on the resource is used, or if no shell was defined on the resource, a default shell is used. For shells: The script name is inserted into the command at the position of a "{0}" marker in the command, or appended as a final argument if no marker is present. A filename suffix adjacent to the marker will be appended to the script name. For more information on shells, see the Shell definition in the Step Help topic.
stepCount	number: The current number of executing steps on this resource.
stepLimit	number: This property specifies the maximum number of steps that can run on this resource at one time.

Description: A *resource* is associated with a server machine available to ElectricFlow for running steps. A resource has a logical name in addition to a host name. A resource can be associated with one or more pools. The resource may be a "proxy target" machine that communicates with the ElectricFlow server though an ElectricFlow agent proxy machine.

Property Name	Description
trusted	boolean: If "true", this agent is trusted.
useSSL	boolean: A Boolean value indicating whether this resource uses SSL for communication. Transport Layer Security (TLS) has replaced Secure Sockets Layer
	version 3.0 (SSLv3) on the ElectricFlow web server and the ElectricFlow server.
workspaceName	name: The workspace name used by this resource.
zoneName	sting: The name of the zone where this resource resides.

[back to top]

Object Type: resourcePool

Description: A resource pool is a container for a group of resources.

Property Name	Description
acl	reference: acl
autoDelete	boolean: If "true", the resource pool is deleted when the last resource is removed or deleted.
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.

Description: A *resource pool* is a container for a group of resources.

Property Name	Description
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
lastResourceUsed	name: The name of the most recently used resource from the pool.
modifyTime	date: The time when this object was last modified.
orderingFilter	string: A Javascript block invoked when scheduling resources for a pool. Note: A Javascript block is not required unless you need to override the default resource ordering behavior.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet
resourcePoolDisabled	boolean: A Boolean value indicating whether this resource pool was disabled.
resourcePoolId	id: This resource pool's ID number.
resourcePoolName	name: The name of the resource pool.

[back to top]

Object Type: resourceUsage

 $\textbf{Description:} \ \ \text{Provides usage information for all resources in the system}.$

For any job step running on a resource, there is a resource usage record that contains the ID and name of the job, job step, and resource.

Property Name	Description
jobId	id: This is this job's ID number, which is a UUID.

Object Type: resourceUsage

Description: Provides usage information for all resources in the system.

For any job step running on a resource, there is a resource usage record that contains the ID and name of the job, job step, and resource.

Property Name	Description
jobName	name: The name of the job.
jobStepId	id: The ID number of the job step.
jobStepName	name: The name of the job step.
licenseWaitTime	The time this job step had to wait because no license was found or available.
resourceId	id: The ElectricFlow-generated resource ID number.
resourceName	name: The name (or names) of the resource.
resourcePoolId	id: The ElectricFlow-generated resource pool's ID number.
resourcePoolName	name: The name of the resource pool.
resourceUsageId	id: The unique ID of the resource usage record.
resourceWaitTime	The time this job step had to wait because no resource was found or available. This could indicate that eligible resources for the step have reached their step limit, are in-use but the step requires a resource exclusively, or resources are down.
waitReason	Possible values are: license, resource, or workspace.
workspaceWaitTime	The time this job step had to wait because no workspace was found or available.

Object Type: schedule

Description: Schedules are used to execute procedures and determine when specific procedures run.

Property Name	Description
acl	reference: acl
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter, which is passed to a procedure when it is invoked. Actual parameters can be provided for jobs and nested subprocedures within a job. Actual parameters are different from "formal parameters"- formal parameters define the parameters a procedure is expecting, and actual parameters provide values to use at run-time.
beginDate	string: The first day on which the schedule is active, in the form YYYY-MM-dd. For example, "2009-06-25"
createTime	date: The time when this object was created.
credentialName	name: The name of the credential being used for impersonation when the schedule launches a job.
description	string: A user-specified text description of the object.
endDate	string: The end of the range of dates the schedule is active, in the form: YYYY-MM-dd. The actual end date is not included in the range. For example, "2009-06-25"
interval	string: A floating point number that represents the time to wait between invocations of the schedule.
intervalUnits	string: These are the units to use when interpreting the <i>interval</i> value. The <i>units</i> can be hours, minutes, seconds, or continuous.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
lastRunTime	date: The last time a job was launched by a schedule.

Object Type: schedule

Description: Schedules are used to execute procedures and determine when specific procedures run.

Property Name	Description
misfirePolicy	string: The policy the schedule uses when it is unable to launch a job at the scheduled time: ignore - wait until the next scheduled time to run runOnce - run immediately, then resume normal scheduling
modifyTime	date: The time when this object was last modified.
monthDays	string: This property specifies which days of the month this schedule should run - shown as a space-separated list of numbers (1-31).
owner	name: The person (user name) who created the object.
priority	string: Values can be low, normal (default), high, or highest
procedureName	name: The name of the procedure that contains this schedule.
projectName	name: The name of the project that contains this schedule.
propertySheet	reference: propertySheet
scheduleDisabled	boolean: A Boolean value indicating whether this schedule was disabled.
scheduleId	id: This schedule's ID number.
scheduleName	name: This property displays this schedule's name and differs from the liveSchedule property found under a job in that the liveSchedule property is written at job creation time only, and not changed, even if the schedule is renamed or deleted.
startTime	string: The time of day when this schedule should start running, in the form: HH:mm:ss

Object Type: schedule

Description: Schedules are used to execute procedures and determine when specific procedures run.

Property Name	Description
stopTime	string: The time of day when this schedule should stop running, in the form: HH:mm:ss
timeZone	string: A Java-compatible time zone string.
weekDays	string: This property specifies which days of the week this schedule should run. This is a space-separated list of MONDAY TUESDAY WEDNESDAY THURSDAY FRIDAY SATURDAY SUNDAY. (Note: These are not localized strings.)

[back to top]

Object Type: Server

Description: This is the ElectricFlow server. There is only one such object in your database.

Read-only server properties

Two read-only server properties are available for use in a step, a script, or an email notifier, which can be used to identify the ElectricFlow server location.

/server/hostName - returns the ElectricFlow server name

/server/hostIP - returns the ElectricFlow server IP address

These properties are especially useful for building a URL link to an ElectricFlow web page. For example, a link in an email notifier that links back to a Job Details page.

A sample link: https://\$[/server/hostName]/commander/jobDetails.php?jobId=\$[jobId]

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
hostIP	string: The server's IP address.

Object Type: Server

Description: This is the ElectricFlow server. There is only one such object in your database.

Read-only server properties

Two read-only server properties are available for use in a step, a script, or an email notifier, which can be used to identify the ElectricFlow server location.

/server/hostName - returns the ElectricFlow server name

/server/hostIP - returns the ElectricFlow server IP address

These properties are especially useful for building a URL link to an ElectricFlow web page. For example, a link in an email notifier that links back to a Job Details page.

A sample link: https://\$[/server/hostName]/commander/jobDetails.php?jobId=\$[jobId]

Property Name	Description
hostName	name: Generally refers to the computer name or IP address for the machine containing the ElectricFlow server.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet
serverId	id: The server's ID number.
serverName	string: The ElectricFlow server's name - defaults to "server".

Object Type: **state**

Description: A *state* object belongs to a workflow and corresponds to the state definition, including a pointer to the workflow, formal parameters (cloned from definition), expanded actual parameters (form the last time a transition was taken to this state), and the "process" which includes the procedure or workflow (cloned from definition), unexpanded actual parameters (cloned from definition, expanded and passed to the process on entry), and the "last run" entity reference.

Property Name	Description
acl	reference: acl
active	boolean: If "true", the state of the workflow is active.
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter, which is passed to a procedure when it is invoked. Actual parameters can be provided for jobs and nested subprocedures within a job. Actual parameters are different from "formal parameters"- formal parameters define the parameters a procedure is expecting, and actual parameters provide values to use at run-time.
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
errorMessage	string: When the outcome is error, this property displays an error message description.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
projectName	name: The name of the project that contains this state.
propertySheet	reference: propertySheet
stateId	id: The ElectricFlow-generated ID number for the state.

Object Type: **state**

Description: A *state* object belongs to a workflow and corresponds to the state definition, including a pointer to the workflow, formal parameters (cloned from definition), expanded actual parameters (form the last time a transition was taken to this state), and the "process" which includes the procedure or workflow (cloned from definition), unexpanded actual parameters (cloned from definition, expanded and passed to the process on entry), and the "last run" entity reference.

Property Name	Description
stateName	name: The name of the state.
subjob	name: The name of the subjob.
subparameters	reference: propertySheet
subprocedure	name: The name of the nested procedure called when a step runs.
subproject	string: If a subprocedure argument was used, this is the name of the project where that subprocedure is found. By default, the current project is used.
substartingState	name: Name of the starting state for the workflow launched when the state is entered.
subworkflow	name: The name of the subworkflow - a workflow called by another workflow.
subworkflowDefinition	name: The name of the subworkflow definition.
workflow	name: The name of the workflow.

Object Type: **stateDefinition**

Description: A *stateDefinition* is a named object that belongs to a workflow definition, which includes specifications for whether or not the state is "startable", formal parameters, notifications, and the process. A process includes a procedure or workflow, the starting state (for workflows only), and actual parameters.

Property Name	Description
acl	reference: acl
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter, which is passed to a procedure when it is invoked. Actual parameters can be provided for jobs and nested subprocedures within a job. Actual parameters are different from "formal parameters"- formal parameters define the parameters a procedure is expecting, and actual parameters provide values to use at run-time.
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
projectName	name: The name of the project that contains this state definition.
propertySheet	reference: propertySheet
startable	boolean: "True" means this state definition can be the initial state of an instantiated workflow.
stateDefinitionId	id: The ElectricFlow-generated ID number for the state definition.
stateDefinitionName	name: The name of the state definition.

Object Type: stateDefinition

Description: A *stateDefinition* is a named object that belongs to a workflow definition, which includes specifications for whether or not the state is "startable", formal parameters, notifications, and the process. A process includes a procedure or workflow, the starting state (for workflows only), and actual parameters.

Property Name	Description
subprocedure	name: The name of the nested procedure called when a step runs.
subproject	string: If a subprocedure argument was used, this is the name of the project where that subprocedure is found. By default, the current project is used.
substartingState	name: Name of the starting state for the workflow launched when the state is entered.
subworkflowDefinition	name: The name of the subworkflow definition.
workflowDefinitionName	name: The name of the workflow definition.

Object Type: **Step**

Description: A *step* includes a command or script executed on a single resource and is the smallest unit of work ElectricFlow understands. Each step specifies a resource on which it should run (either the name of a specific machine or the name of a *pool* of equivalent machines, in which case ElectricFlow picks a machine from the pool). A step can be given a time limit and if the step does not complete within the specified time, ElectricFlow automatically aborts it.

Steps are ordered within a procedure and normally execute sequentially. However, it is possible to mark a consecutive range of steps for parallel execution, so all steps in that range run concurrently.

Property Name	Description
acl	reference: acl
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter, which is passed to a procedure when it is invoked. Actual parameters can be provided for jobs and nested subprocedures within a job. Actual parameters are different from "formal parameters"- formal parameters define the parameters a procedure is expecting, and actual parameters provide values to use at run-time.
alwaysRun	boolean: If true, this step runs even if the job is aborted before the step completes. Note that a force abort will abort an alwaysRun step.
broadcast	boolean: If true, this step will run on all resources in a pool.
command	string: This property specifies the command this step runs.
condition	string: If this element is not present, the event is ALWAYS triggered. If specified, the event is triggered only if the value of the condition argument is TRUE; if not true, a boolean value of "false" or "0" was used. Condition arguments can be a literal, a fixed value string, or a string subject to property expansion.
createTime	date: The time when this object was created.
credentialName	name: The credential name assigned to this step.

Object Type: **Step**

Description: A *step* includes a command or script executed on a single resource and is the smallest unit of work ElectricFlow understands. Each step specifies a resource on which it should run (either the name of a specific machine or the name of a *pool* of equivalent machines, in which case ElectricFlow picks a machine from the pool). A step can be given a time limit and if the step does not complete within the specified time, ElectricFlow automatically aborts it.

Steps are ordered within a procedure and normally execute sequentially. However, it is possible to mark a consecutive range of steps for parallel execution, so all steps in that range run concurrently.

Property Name	Description
description	string: A user-specified text description of the object.
errorHandling	This is the error handling policy copied from the procedure step and indicates how the step responds to errors:
	failProcedure - The current procedure continues, but the overall status is error (default).
	abortProcedure - Aborts the current procedure, but allows already- running steps in the current procedure to complete.
	abortProcedureNow - Aborts the current procedure and terminates running steps in the current procedure.
	abortJob - Aborts the entire job, terminates running steps, but allows alwaysRun steps to run.
	abortJobNow - Aborts the entire job and terminates all running steps, including alwaysRun steps.
	ignore - Continues as if the step succeeded.
exclusive	boolean: If true, this step acquires and retains its resource exclusively.
exclusiveMode	string: Possible values are: none, job, step, call
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.

Object Type: **Step**

Description: A *step* includes a command or script executed on a single resource and is the smallest unit of work ElectricFlow understands. Each step specifies a resource on which it should run (either the name of a specific machine or the name of a *pool* of equivalent machines, in which case ElectricFlow picks a machine from the pool). A step can be given a time limit and if the step does not complete within the specified time, ElectricFlow automatically aborts it.

Steps are ordered within a procedure and normally execute sequentially. However, it is possible to mark a consecutive range of steps for parallel execution, so all steps in that range run concurrently.

Property Name	Description
logFileName	name: The name of the log file produced by this step, relative to the job's workspace directory.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
parallel	boolean: If true, this step runs in parallel with other adjacent steps also marked to run in parallel.
postLogFileName	name: This property displays the log file name produced by this step's post processor.
postProcessor	string: This property displays the post processor name that is used to gather information about this step. Typically, this is either postp or postp <options>, or an appropriate command or path including options to a postprocessor available on the ElectricFlow server.</options>

Object Type: **step**

Description: A *step* includes a command or script executed on a single resource and is the smallest unit of work ElectricFlow understands. Each step specifies a resource on which it should run (either the name of a specific machine or the name of a *pool* of equivalent machines, in which case ElectricFlow picks a machine from the pool). A step can be given a time limit and if the step does not complete within the specified time, ElectricFlow automatically aborts it.

Steps are ordered within a procedure and normally execute sequentially. However, it is possible to mark a consecutive range of steps for parallel execution, so all steps in that range run concurrently.

Property Name	Description
precondition	string: By default, if a step has no precondition, it will run when scheduled. Set this property to make a step wait until one or more dependent conditions are met. When a job step is eligible to transition from pending to runnable, a precondition is evaluated. A precondition is a fixed text or text embedding property reference that is evaluated to TRUE or FALSE. An empty string, a \"0\" or \"false\" is interpreted as FALSE. Any other result string is interpreted as TRUE. The step will block until the precondition is TRUE.
procedureName	name: The name of the procedure that contains this step.
projectName	name: The name of the project that contains this step.
propertySheet	reference: propertySheet
releaseExclusive	boolean: A Boolean value indicating whether this step should release its resource upon completion.
releaseMode	string: Possible values are: none, release, releaseToJob
resourceName	name: The resource's name this step should use to run on.

Object Type: **step**

Description: A *step* includes a command or script executed on a single resource and is the smallest unit of work ElectricFlow understands. Each step specifies a resource on which it should run (either the name of a specific machine or the name of a *pool* of equivalent machines, in which case ElectricFlow picks a machine from the pool). A step can be given a time limit and if the step does not complete within the specified time, ElectricFlow automatically aborts it.

Steps are ordered within a procedure and normally execute sequentially. However, it is possible to mark a consecutive range of steps for parallel execution, so all steps in that range run concurrently.

A step in one procedure can invoke another procedure instead of running a command. The invoking step provides parameters needed by the nested procedure.

Property Name	Description
shell	name: The shell used to execute the step's commands on a resource. The script name is inserted into the command at the position of a "{0}" marker in the command, or appended as a final argument if no marker is present. A filename suffix adjacent to the marker will be appended to the script name. For more information on shells, see the Shell definition in the Step Help topic.
stepId	id: The step's ID number.
stepName	name: The step's name.
subprocedure	name: The nested procedure name to call when this step executes.
subproject	name: This property is displayed if the subprocedure element is present and specifies the project to which the subprocedure belongs (by default, the current project is used.)
timeLimit	number: A floating point number that specifies the maximum length of time this step is allowed to run.
timeLimitUnits	The units to use when interpreting the time limit. Units can be hours, minutes, or seconds.
workingDirectory	name: The name of the step's working directory.
workspaceName	name: The workspace name used by this step.

Object Type: transition

Description: A *transition* object belongs to a state and corresponds to the transition definition, and includes the target state (unexpanded actual parameters cloned from definition, expanded and passed to the target state on entry), the Javascript condition, and the trigger

(onEnter|onStart|onCompletion|manual).

Property Name	Description
acl	reference: acl
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter that is passed to the target state when the transition is taken.
condition	string: If empty or non-zero, the transition is allowed to trigger. If set to "0", the transition is ignored.
createTime	date: The time when this object was created.
description	string: A user-specified text description of the object.
index	The numeric index of a transition that indicates the transition order in the state definition's transition list.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
projectName	name: The name of the project that contains this transition.
propertySheet	reference: propertySheet
stateName	name: The name of the state.
targetState	string: The target state for the transition definition.
transitionId	id: The ElectricFlow-generated ID for the transition.

Object Type: transition

Description: A *transition* object belongs to a state and corresponds to the transition definition, and includes the target state (unexpanded actual parameters cloned from definition, expanded and passed to the target state on entry), the Javascript condition, and the trigger

(onEnter|onStart|onCompletion|manual).

Property Name	Description
transitionName	name: The name of the transition.
trigger	string: Possible values are: onEnter - before any actions onStart - after a subjob or workflow is created onCompletion - after a subjob or workflow completes manual - when a user manually requests a transition
workflowName	name: The name of the workflow.

[back to top]

Object Type: transitionDefinition

Description: A *transitionDefinition* is a named object that belongs to a state definition, which includes the target state definition (with actual parameters to the target state, the Javascript condition, and the trigger (onEnter|onStart|onCompletion|manual).

Property Name	Description
acl	reference: acl
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter passed to the target state.
condition	string: If empty or non-zero, the transition is allowed to trigger. If set to "0", the transition is ignored.
createTime	date: The time when this object was created.

Object Type: transitionDefinition

Description: A *transitionDefinition* is a named object that belongs to a state definition, which includes the target state definition (with actual parameters to the target state, the Javascript condition, and the trigger (onEnter|onStart|onCompletion|manual).

Property Name	Description
description	string: A user-specified text description of the object.
index	The numeric index of a transition that indicates the transition order in the state definition's transition list.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
projectName	name: The name of the project that contains this transition definition.
propertySheet	reference: propertySheet
stateDefinitionName	name: The name of the state definition.
targetState	string: The target state for the transition definition.
transitionDefinitionId	id: The ElectricFlow-generated ID for the transition definition.
transitionDefinitionName	name: The name of the transition definition.
trigger	string: Possible values are: onEnter - before any actions onStart - after a subjob or workflow is created onCompletion - after a subjob or workflow completes manual - when a user manually requests a transition
workflowDefinitionName	name: The name of the workflow definition.

Object Type: **USE**

Description: *User* refers to any user currently logged into the ElectricFlow system or any user who may use ElectricFlow, but may not be currently logged in.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.
email	name: Displays this user's email address.
fullUserName	name: The user's real name.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
mutable	boolean: If true, the user is editable within ElectricFlow via the web UI or the modifyUser API.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet
providerName	name: The name of the directory provider that controls this user.
userId	id: A unique ID number for a user object.
userName	name: Displays this user's name and also appears in a group object and displays the name of a user who belongs to this group.

Object Type: **workflow**

Description: A *workflow* object includes a pointer to the workflow definition, sets of states, the active state, the starting stare, parameters to the initial state, "complete" - a boolean value determining whether or not the workflow is complete, and a log containing the transactions taken and user events.

Property Name	Description
acl	reference: acl
activeState	string: The name of the active state on the workflow object.
actualParameter	reference: propertySheet An actualParameter is an object that provides the value for a parameter that is passed to the target state when the transition is taken.
callingState	string: The full property path to the state that created this workflow.
callingStateId	id: The ID number for the full property path to the state that created this workflow.
completed	boolean: If "true", the workflow is completed and no additional transactions will be evaluated.
createTime	date: The time when this object was created.
deleted	boolean: The object was marked for background deletion. Possible values are "0" or "1". Default is "0" (not set).
elapsedTime	The time this object ran from start to finish.
finish	date: The date/time when this object finished.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
launchedByUser	string: The name of the user or project principal that explicitly launched the workflow.
liveWorkflowDefinition	name: The current workflow definition name for the workflow definition from which the workflow was created.

Object Type: Workflow

Description: A *workflow* object includes a pointer to the workflow definition, sets of states, the active state, the starting stare, parameters to the initial state, "complete" - a boolean value determining whether or not the workflow is complete, and a log containing the transactions taken and user events.

Property Name	Description
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
projectName	name: The name of the project that contains this workflow.
propertySheet	reference: propertySheet
start	date: The date/time when this workflow began to run.
startingState	string: The initial state of the workflow.
workflowDefinitionName	name: The name of the workflow definition.
workflowId	id: The ElectricFlow-generated ID for the workflow.
workflowName	name: The name of the workflow.

[back to top]

Object Type: workflowDefinition

Description: A *workflowDefinition* object contains state and transition definitions, including the workflow name template (analogous to a job name template on a procedure), and ordered sets of state and transition definitions.

Property Name	Description
acl	reference: acl
createTime	date: The time when this object was created.

Object Type: workflowDefinition

Description: A *workflowDefinition* object contains state and transition definitions, including the workflow name template (analogous to a job name template on a procedure), and ordered sets of state and transition definitions.

Property Name	Description
description	string: A user-specified text description of the object.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
projectName	name: The name of the project that contains this workflow definition.
propertySheet	reference: propertySheet
workflowDefinitionId	id: The ElectricFlow-generated ID for the workflow definition.
workflowDefinitionName	name: The name of the workflow definition.
workflowNameTemplate	string: Template used to determine the default names for workflows launched from a workflow definition.

Object Type: workspace

Description: ElectricFlow provides each job step with an area on the disk it can use for "working files" and results. This disk area is called a job *workspace*. (The job workspace defaults to a directory under the workspace directory.) A job step can create whatever files it needs within its workspace, and ElectricFlow automatically places files such as step logs in the workspace.

Normally, a single workspace is shared by all steps in a job, but different steps within a job could use different workspaces. The location of the job step workspace is displayed on the Job Details page for the job under "Details" for the step.

Property Name	Description
acl	reference: acl
agentDrivePath	name: The drive-letter based mount point for this workspace on Windows agents.
agentUncPath	name: The UNC path for this workspace on Windows agents.
agentUnixPath	name: The UNIX path for this workspace on UNIX agents.
createTime	date: The time when this object was created.
credentialName	name: The name of the credential being used when a resource connects to a network share while setting up the workspace.
description	string: A user-specified text description of the object.
lastModifiedBy	name: This shows who (generally, a user name) last modified this object.
local	boolean: If "true", this workspace is local.
modifyTime	date: The time when this object was last modified.
owner	name: The person (user name) who created the object.
propertySheet	reference: propertySheet

Object Type: Workspace

Description: ElectricFlow provides each job step with an area on the disk it can use for "working files" and results. This disk area is called a job *workspace*. (The job workspace defaults to a directory under the workspace directory.) A job step can create whatever files it needs within its workspace, and ElectricFlow automatically places files such as step logs in the workspace.

Normally, a single workspace is shared by all steps in a job, but different steps within a job could use different workspaces. The location of the job step workspace is displayed on the Job Details page for the job under "Details" for the step.

Property Name	Description
workspaceDisabled	boolean: A Boolean value that indicates whether this workspace was disabled.
workspaceId	id: The workspace's ID number.
workspaceName	name: The workspace's name.
zoneName	name: The name of the zone where this workspace resides.

[back to top]

Object Type: **ZONe**

Description: ElectricFlow provides the ability to create zones—often used to enhance security.

- A zone is a way to partition a collection of agents to secure them from use by other groups.
 For example, you might choose to create a developers zone, a production zone, and a test zone—agents in one zone cannot directly communicate with agents in another zone.
- A default zone is created during ElectricFlow installation.
 The server implicitly belongs to the default zone, which means all agents in this zone can communicate with the server directly (without the use of a gateway).
- Each zone can have one or more "gateway agents", which you define. Gateway agents are used for communication from one zone to another zone. For more information, see the Gateways Help topic.

Property Name	Description	
acl	reference: acl	

Object Type: **ZONe**

Description: ElectricFlow provides the ability to create zones—often used to enhance security.

- A zone is a way to partition a collection of agents to secure them from use by other groups. For example, you might choose to create a developers zone, a production zone, and a test zone—agents in one zone cannot directly communicate with agents in another zone.
- A default zone is created during ElectricFlow installation.
 The server implicitly belongs to the default zone, which means all agents in this zone can communicate with the server directly (without the use of a gateway).
- Each zone can have one or more "gateway agents", which you define. Gateway agents are used
 for communication from one zone to another zone. For more information, see the Gateways Help
 topic.

Property Name	Description	
createTime	date: The time when this object was created.	
description	string: A user-specified text description for the object.	
lastModifiedBy	name: This person (generally, a user name) last modified this object.	
modifyTime	date: The time when this object was last modified.	
owner	name: The person (user name) who created the object.	
propertySheet	reference: propertySheet	
resources	string: A space-separated list of resources contained in this zone.	
zoneId	id: The ElectricFlow-generated ID for this zone.	
zoneName	string: The name of the zone.	

Property error codes

The following list provides error codes that may appear as a value within an errorCode property.

ABORTED

ACCESS_DENIED

AGENT_BAD_WORKINGDIR

AGENT_FAILED_CONNECT_BROKER

AGENT_FAILED_CREATE_FILE

AGENT_FAILED_CREATE_WORKSPACE

AGENT_FAILED_IMPERSONATION

AGENT_FAILED_MAP_DRIVE

AGENT_FAILED_PROXYTARGET_PING

AGENT_INCOMPATIBLE_VERSION

AGENT_INTERNAL_ERROR

AGENT_INVALID_CWD

AGENT_INVALID_MESSAGE

AGENT_INVALID_PING_TOKEN

AGENT_INVALID_WORKSPACE

AGENT_IO_CONNECTION_REFUSED

AGENT_IO_CONNECTION_RESET

AGENT_IO_ERROR

AGENT_IO_NO_ROUTE_TO_HOST

AGENT_IO_PORT_UNREACHABLE

AGENT_MALFORMED_XML

AGENT_NONABSOLUTE_PATH

AGENT_NONEXISTENT_DIR

AGENT_NONEXISTENT_FILE

AGENT_PING_FAILED

AGENT_STREAM_STOPPED

AGENT_TIMEOUT

AGENT_UNKNOWN_CMDID

AGENT_UNKNOWN_COMMAND

AGENT_WRONG_FILE_TYPE

BAD_AGENT_RESPONSE

BAD_LOGFILE_PATH

CANCELED

CIRCULAR_PROCEDURE_REFERENCE

CORRUPT_CREDENTIAL

DUPLICATE_JOB_NAME

EMPTY_SUBPROCEDURE

FAILED_JOB_RENAME

FORMAL_PARAMETER_ERROR

INTERNAL_ERROR

INVALID_EMAIL_HEADER

MY_EVENT_EXPANSION_ERROR

NO_EMAIL_HEADERS_SEPARATOR

NONEXISTENT_EMAIL_CONFIG

NONEXISTENT_PROCEDURE

NONEXISTENT_RESOURCE

NONEXISTENT_WORKSPACE

NOTIFIER_EXPANSION_ERROR

POST_PROCESSOR_ERROR

POST_PROCESSOR_OUTPUT_FAILURE

PROPERTY_REFERENCE_ERROR

RESOURCE_WITHOUT_HOSTNAME

SERVER_SHUTDOWN

TIMEOUT

UNKNOWN_HOST

Postprocessors: Collecting Data for Reports

Postp

Extending postp: matchers

Postp functions

Integration with the ElectricFlow user interface

Postp integration with Java Tools

Overview

In ElectricFlow, reporting is divided into two phases.

- The first phase is data collection: interesting information is extracted from job step logs and saved in the ElectricFlow database.
- The second phase is report generation: data collected previously is retrieved and organized into reports.

The report phases are separated so data is gathered once but then available to use in a variety of different reports either immediately or later. For example, one report might summarize errors within a particular job, and another report might display error trends from all jobs over the past month.

ElectricFlow implements data collection with a postprocessor.

- A postprocessor is a command associated with a particular procedure step. If a postprocessor is specified for a step, it executes concurrently with the main step command.
- The postprocessor runs on the same machine as the main command and in the same working directory, and it retrieves the log file from the step as its standard input.

ElectricFlow includes a standard postprocessor called postp that you can use and extend.

postp scans the step's log file looking for interesting output such as error messages and then sets properties on the job step to describe what it found. For example, postp might create a property named "errors" whose value is a count of the number of error messages in the log file, or a property named "tests" that counts the number of tests executed by the step. Also, postp can extract portions of the step log that contain useful diagnostic information and save this information for reporting.

Standard ElectricFlow reports, such as those on the Job Details page, display information collected by *postp* such as properties named "errors" and diagnostic log extracts. This information is available immediately, even before the step completes execution, so you can view it via ElectricFlow's web interface to monitor step execution. Also, you can create additional report generators of your own, which can use the same information displayed by ElectricFlow and/or any other additional information of your choice.

Postp

ElectricFlow's general-purpose postprocessor, *postp*, uses regular expression patterns to detect interesting lines in a step log.

- postp is already configured with patterns to handle many common cases such as error messages and
 warnings from gcc, gmake, cl, junit, and cppunit, or any error message containing the string "error." You
 can use the empty matcher group named "none" if just want to run a postpEndHook.
- postp is easy to use by simply setting the postprocessor for a step to "postp."
- postp also supports several useful command-line options.

To explore these options, invoke "postp --help" from your command-line.

Extending *postp*: matchers

If you find useful patterns in your log files undetected by postp, you can extend postp with additional patterns. This feature is easily implemented, but the extension interface currently involves writing simple Perl scripts and you need to look at the Perl source file for postp as you do this. The postp source code is installed during ElectricFlow installation and located in the src/postp.pl file in the distribution directory.

Postp is driven by a collection of *matchers*, which are regular expression patterns that select certain lines from step logs, and by a collection of Perl functions the matchers invoke to handle lines of interest. A matcher is a Perl hash with three values similar to the following example:

```
{
  id => "error",
  pattern => q{ERROR:|[Ee]rror:},
  action => q{
    incValue("errors"); diagnostic("", "error", -4)
  },
}
```

Explanation of the values in the Perl hash example above:

- id A unique name for the matcher—used to identify the matcher in command-line arguments and a few other places.
- pattern A regular expression tested against each line of the step log file. This particular pattern matches lines containing any of the strings "ERROR:", "Error:" or "error:"
- action A Perl script that executes whenever a log line matches the pattern for this matcher. The script in this example increments a variable named "errors" that is copied automatically to a job step property with the same name. The script also saves a portion of the step log beginning 4 lines prior and extending through the current line and associates those lines with this error so it can be displayed in the web interface. See below for more information on the incValue and diagnostic functions.

To extend postp with patterns of your own, write a Perl script to add new patterns to the @::gMatchers array. Here is a simple example:

```
my @newMatchers = (
    {
              => "coreDump",
      id
      pattern => "core dumped",
      action => q{
                      incValue("coreDumps")
    },
      id
              => "segFault",
      pattern => "segmentation fault",
      action => q{
                       incValue("segFaults")
                 }
    },
);
push @::gMatchers, @newMatchers;
```

These matchers detect lines containing the strings "core dumped" or "segmentation fault" and increment separate variables for each line type.

After writing extension code, you must ask postp to execute the code when it starts up. You can execute extension code in one of two ways:

- Place the code into a file and invoke postp with the --load option. For example, postp --load fileName
- Copy the code into a property in ElectricFlow and use the --loadProperty option to postp. For example, postp --loadProperty /myProject/extraMatchers

Postp extensions can contain arbitrary Perl code, which means you can use this mechanism to define additional functions to invoke in matcher actions if existing functions do not provide what you need.

Note: Additional postp matchers are available for your convenience. The matcher sample directory was installed during ElectricFlow installation. Go to src/samples/postp.

Postp functions

Postp contains several built-in functions to invoke in your matchers. The most useful functions are summarized below. Also, you can scan postp code for additional functions.

debugLog(format, arg, arg, ...)

Outputs information to the debug log, if the --debugLog command-line switch is set. Format provides a format string similar to printf, and each argument provides a value to substitute into the format string.

backTo(pattern, start)

This function searches backward to find the first line in the step log matching "pattern" (a regular expression) and returns the offset of that line relative to the current line. The result is normally used as an argument to the "diagnostic" function. "Start" is optional; it specifies the first line to check and is specified as an offset relative to the current line (it defaults to -1).

backWhile(pattern, start)

This function searches backward to find the first line in the step log that does not match "pattern" (a regular expression) and returns the offset relative to the current line of the line just after the first non-matching line. The result is normally used as an argument to the "diagnostic" function. "Start" is optional; it specifies the first line to check and is specified as an offset relative to the current line (it defaults to -1).

currentModule()

Returns the name of the current module (as determined by previous calls to pushModule and popModule), or an empty string if there is no current module.

diagnostic(name, type, first, last)

This function extracts a group of contiguous lines from the log file and saves them along with additional information for reports such as the log extracts, which appear at the bottom of the Job Details web page. The range of lines to extract is indicated by "first" and "last," each of which is an offset relative to the current line. For example, if "first" is -3 and "last" is 2, then 6 lines will be recorded: 3 lines before the current line, the current line, and 2 lines after the current line. In many cases, the values for "first" and "last" are computed by calling functions such as "forwardTo" or "backWhile." "Name" provides an identifier for this particular diagnostic, such as the name of a test that failed or a file that did not compile. "Type" specifies which kind of information this is, and must be "error," "warning," or "info." In addition to log lines, this function records "name," "type," the current module, if any, and the name of the current matcher.

forwardTo(pattern, start)

This function searches forward to find the first line in the step log matching "pattern" (a regular expression) and returns the offset of that line relative to the current line. The result is normally used as an argument to the "diagnostic" function. "Start" is optional; it specifies the first line to check and is specified as an offset relative to the current line (it defaults to 1).

forwardWhile(pattern, start)

This function searches forward to find the first line in the step log that does not match "pattern" (a regular expression) and returns the offset relative to the current line of the line just before the first non-matching line. The result is normally used as an argument to the "diagnostic" function. "Start" is optional; it specifies the first line to check and is specified as an offset relative to the current line (it defaults to 1).

incValue(name, increment)

Adds "increment" to a value named "name" and arranges for that value to be written eventually to a property by the same name on the current job step. If this is the first call for "name," its value is initialized to 0. "Increment" is optional and defaults to 1.

Note: postp does not check the job step for a pre-existing property with the same name; it simply overwrites it.

logLine(lineNumber)

Returns the line from the step log given by lineNumber. 1 corresponds to the first line in the step log and the Perl variable \$::gCurrentLine holds the number of the current line. This function caches a sliding window of lines in the file, allowing you to go back to retrieve lines preceding the current line (as long as they do not precede it by too many lines). If the requested line is off the end of the file then undef is returned. If the requested line is before the beginning window of cached lines, an empty string is returned.

popModule()

Cancels the effect of the most recent call to pushModule, resetting the current module name to whatever it was before the corresponding call to pushModule.

postpEndHook()

If you define a function with this name, it invokes after postp has finished processing the log file, but before it makes its final properties update on the job step. Use this function to perform your own operations such as generating an error if the log file did not contain a particular line you were expecting.

pushModule(name)

In some situations it is possible to divide the log file into parts corresponding to different modules. For example, with recursive make invocations, there are typically notifications in the log output before and after each recursive make. This function is invoked to indicate a new module is being entered, where "name" is the name of the module. After this function is called, the "diagnostic" function will include "name" with error or warning messages to provide additional information in job reports. The previous module name, if any, is saved; you can return to it by calling popModule.

setProperty

This function sets a property in the ElectricFlow server. If the named parameter is a relative path, like <code>moduleCount</code>, the property is set or created on the current job step. You can use an absolute path, like <code>/myJob/fileLocation</code> also. Calling <code>setProperty</code> does not result in an immediate call to the ElectricFlow server. The property is added to an update list for updating at the next "update interval", typically every 30 seconds.

Integration with the ElectricFlow user interface

postp interacts with the ElectricFlowUI using two methods. The first method: Create custom properties with special names that are recognized by the UI itself. The second method: Create a file that contains "diagnostics",

which are used to display errors or warnings, and link them to specific sections in the step's log file.

Custom property names and values

postp can be used to create properties in ElectricFlow, on the job step or anywhere else in ElectricFlow. However, the following properties are used by the standard ElectricFlow UI, so you should use these property names whenever possible, and avoid using these names in ways that conflict with the definitions below.

- compiles the number of files compiled during the job step
- diagFile the filename in the top-level directory of the job's workspace, containing diagnostics extracted from the step's log file
- errors the number of errors (compilation failures, test failures, application crashes, and so on) that occurred during the job step. When property errors are set by postp, the step outcome is set to error also.
- tests the number of tests executed by the job step, including successes and failures
- testsSkipped the number of tests skipped during the job step
- warnings the number of warnings that occurred during the job step. When property warnings is set by postp, the step outcome is set to warning also.
- presummary if this property exists, its value is displayed in the "Status" field (on the Job Details page) for this step. This property appears *before* whatever would normally be displayed for status. If the property contains multiple lines separated by newline characters, each line is displayed on a separate line in the status field.
- postSummary if this property exists, its value is displayed in the "Status" field (on the Job Details page) for this step. This property appears *after* whatever would normally be displayed for status. If the property contains multiple lines separated by newline characters, each line is displayed on a separate line in the status field.
- summary if this property exists, its value is displayed in the "Status" field (in the job reports) for this step, replacing whatever would normally be displayed for status. If the property contains multiple lines separated by newline characters, each line is displayed on a separate line in the status field.

Diagnostic information

The second form of postprocessor generated output contains diagnostic extracts from the step's log file, typically providing additional information about problems. The postprocessor stores diagnostic information in an XML file in the top-level directory of the job workspace, then sets the step's **diagFile** property with the name of the file. Diagnostic files must have a format like the following example:

XML elements in the diagnostic file are as follows:

- diagnostics Overall container; its children consist of all of the diagnostic elements.
- diagnostic Describes one diagnostic extract; the elements described below are all children of this element.
- matcher (optional) Identifier for the matcher that triggered this diagnostic; used primarily for debugging.
- o name (optional) Identifier that indicates the problem or situation that resulted in this diagnostic, such as the name of a failed test or the name of a file that did not compile.
- o type Type of message: must be "error," "warning," or "info."
- module (optional) Name of the module in which the issue occurred, such as the name of a source code module being compiled at the time of a compile error.
- firstLine Line number in the log file for the first line of the diagnostic extract (1 means the first line of the file). This element is used to provide a link from the diagnostic extract to the full log file.
- o numLines Total number of lines included in the diagnostic extract.
- o message The actual lines from the log file.

Postp integration with Java Tools

ElectricFlow integrates with three standard Java tools through postp matchers. These Java tools are:

- EMMA an open source toolkit for measuring and reporting Java code coverage (Emma v2.0)
- JUnit a framework for writing and running automated tests (tested with Ant v1.7)
- Clover Atlassian's Java code coverage (generated by system, functional, or unit tests) analysis tool (Clover v2.4)

Note: If one of these Java tools is invoked in an ElectricFlow job step, ElectricFlow automatically detects the invocation [if you are using Postp] and adds any reports generated by these tools to the list of links at the top of the Job Details page.

The process...

Postp parses output from the invoked JAVA tool to match paths to reports it has already created. Then, postp copies all files that comprise the report to a unique location in the "Artifacts" Directory. Next, postp generates a

link to the location in the Artifacts Directory to make the report available in the Links section on the Job Details page.

The following example of generated output from Emma illustrates this process:

```
emma:
init:
[mkdir] Created dir: C:\Documents and Settings\ptharani\Desktop\emma\emma-
2.0.5312\examples\out
compile:
[javac] Compiling 4 source files to C:\Documents and
Settings\ptharani\Desktop\emma-2.0.5312\examples\out
[emmajava] EMMA: processing classpath ...
[emmajava] EMMA: [3 class(es) processed in 47 ms]
[emmajava] main(): running doSearch()...
[emmajava] main(): done
[emmajava] EMMA: writing [txt] report to [C:\Documents and
Settings\ptharani\Desktop\emma-2.0.5312\examples\coverage\coverage.txt] ...
[emmajava] EMMA: writing [html] report to [C:\Documents and
Settings\ptharani\Desktop\emma-2.0.5312\examples\coverage\coverage.html] ...
all:
BUILD SUCCESSFUL
Total time: 1 second
```

From this output

The actual link may reside here:

```
/home/commanderWorkspace/job_112_200901081732/artifacts/javaTools/238/emmaCoverage/2/coverage.html
```

While the link name may be:

```
Step Id 238 ant-on-the-fly - emma report# 2
```

And the value of the link may be:

```
jobSteps/238//javaTools/238/emmaCoverage/2/coverage.html
```

Java Tool matcher examples

Two examples of postp Emma matchers:

```
id => "emmaReport1",
  pattern => q{EMMA: writing},
  action => q{emmaExtractReport()},

},

id => "emmaReport2",
  pattern => q{\[report\] writing},
  action => q{emmaValidateOutput()},
},
```

An example of postp JUnit matchers:

```
id => "junitReportCapture",
  pattern => q{\[junitreport\] Processing},
  action => q{junitExtractReport ()},
},
```

These matchers correspond to the following output:

```
junit.report:
[junitreport] Processing
/net/WinStor2home/ptharani/junitDemo1/sample/testreport/TESTS-TestSuites.xml to
/tmp/null1214791178
[junitreport] Loading stylesheet jar:[file:/usr/local/tools/common/apache-ant-
1.7.0/lib/ant-junit.jar\!/org/apache/tools/ant/taskdefs/optional/junit/xsl/junit-
frames.xsl
[junitreport] [file:/usr/local/tools/common/apache-ant-1.7.0/lib/ant-
junit.jar!/org/apache/tools/ant/taskdefs/optional/junit/xsl/junit-frames.xsl
[junitreport]] Transform time: 622ms
[junitreport] Deleting: /tmp/null1214791178
```

An example of postp Clover matchers:

```
id => "cloverHtmlReportAntTask",
  pattern => q{\[clover-html-report\] Writing HTML report to},
  action => q{cloverExtractReport()},
},
```

These matchers correspond to the following output:

```
clover.report:
[clover-html-report] Clover Version 2.4.0, built on November 05 2008 (build-747)
[clover-html-report] Loaded from: /home/cloverDemo/clover-ant-2.4.0/lib/clover.jar
[clover-html-report] Clover: Evaluation License registered to electric cloud.
[clover-html-report] You have 27 day(s) before your license expires.
[clover-html-report] Loading coverage database from: '/home/cloverDemo/clover-ant-2.4.0/tutorial/.clover/clover2_4_0.db'
[clover-html-report] Writing HTML report to '/home/cloverDemo/clover-ant-2.4.0/tutorial/clover_html'
[clover-html-report] Done. Processed 1 packages in 2559ms (2559ms per package).
```

Artifacts directory

The value of the artifacts directory determines the scope of what is visible in the ElectricFlow UI from a job's workspace. By default, the Artifacts Directory is set to "artifacts", so the artifacts directory would have the form:

```
<path to job workspace>/artifacts
```

The value of the artifacts directory can be set in any of the following properties:

```
/myJob/artifactsDirectory
/myProject/artifactsDirectory
/server/settings/artifactsDirectory
```

Postp queries these properties [in the order listed] to determine the location of the Artifacts Directory. If no value is found (because the property was never set), the default "artifacts" is used.

Postp has a feature where it recognizes that data belongs to some standard tool (such as Junit), copies the logs produced by that tool to the artifacts directory, and then creates the report-url properties. These actions are done with the junitReportCapture matcher.

If you want to disable this matcher in your postp invocation in the step that runs junit tests, do the following:

```
postp --dontCheck junitReportCapture
```

It is possible that the artifacts directory will still be created even if nothing is put in it. If that occurs, set the artifactsDirectory property on your job (or the owning project) to empty-string.

Reports

ElectricFlow provides multiple reports and custom report capabilities to help you manage your build environment.

- · Real-time reports filtered view of your workload in real-time
- . Build reports summary reports produced at the end of a build and attached to the job
- Batch reports summaries of your build environment with trends over time, two types:
 - Default Batch reports automatically installed during ElectricFlow installation and scheduled to run daily
 - (Cross Project Summary, Variant Trend, Daily Summary, Resource Summary, Resource Detail)
 - Optional Batch reports you can configure, rename, and schedule these reports to fit your requirements
 - (Category Report, Procedure Usage Report, Count Over Time Report, Multiple Series Reports)
- · Custom reports your choice to create and add at any time

You can use also StatsD and Graphite to generate custom reports. For more information, go to System Health Monitoring on page 927.

Note: Batch and Custom reports must be run on the ElectricFlow server's agent (local agent) only. These reports need the BIRT report engine, which is installed on the ElectricFlow server.

Run reports on a non-local resource

If you want to configure reports to run on a non-local resource, do the following:

- In the web interface, go to Projects > Electric Cloud > runReports > runEcrptdata.
- 2. Change the Parameter Resource from the default value "local" to the name of a different resource. the resource must have access to the plugins directory.

Real-time reports

Home page

The Home page is a user-customizable dashboard that can be configured to show you just the ElectricFlow jobs you care about.

The Jobs Quick View section (on the right-side of the Home page) shows jobs you define using filters on job properties. You can create multiple categories filtered on specific project names, procedure names, users, or any other job property. Summary details for the job "pop-up" if you hover your mouse over one of the jobs.

Jobs

Unlike the Home page, which allows you to see specific subsets of work, the Jobs page shows all jobs that are running and all jobs that have completed. This is a good report to view when you need to see a list of all jobs run by any user or schedule. System administrators and managers can use this page to monitor overall system throughput.

Job details / Step details

After a job starts, you can get execution details from the Job Details page. To view job details, click on the job from the Home page or the Jobs tab. From within the Job Details page you can drill down to see details for an individual step. This is the finest grain of reporting, showing the detailed results of a specific job.

Resources

The Resources page reports the current state of your defined resources. This page also shows the current state of the agent (connected or not) and any steps currently running on that resource. Use this report to monitor your resources in near real time.

Build reports using ecreport

ElectricFlow includes a utility, *ecreport*, that summarizes build results and sends that summary through email. This customizable utility allows you to create a simple matrix based on results stored in properties and log files. Details such as SCM change logs and errors can be included.

In addition, you can attach your own reports to any job and have them show up as attachments to the Job Details Report. To attach custom reports, put them in the top-level workspace directory in a file that ends with ".html" and contains the string "report" or "Report." Names of all matching files are displayed in the Summary section in the upper right corner of the page. Click on a report name to view the report.

The following is an example of ecreport output:

thunder-main.4315-200706	261344 is good!		View Online
	-		view Oillile
- 1	Windows	Linux	
Compile Build installer	1183 compiles	1121 compiles	
	ok	ok	
Agent unit tests	748 tests	642 tests	
Server unit tests Web unit tests	1875 tests	1860 tests	
	1556 tests	1557 tests	
Tool unit tests	410 tests	338 tests	
Install	ok	ok	
System tests	22 tests	22 tests	
General Information			
Start Time	6/26/2007 13:44:19		
Elapsed Time	1 hour, 27 minutes, 6.0 s	econds	
Workspace	//f2/scratch/buildserver3	build/thunder-main.4315-2007	06261344
Recent Updates			
chris			
Change 24249 by chris@chris=hundar on 2007/06/26 13:24:31			
Change 24248 by <u>chris@chris</u> -thunder on 2007/06/26 13:24:31			
Adding "index-redirect.php" only type in the host name directory, and renamed "ind	for the web server.		
This will resolve the issue	of https redirection	on not working when us	ing the http
Jobs fixed			
NMB-2761 on 2007/06/26 by t	com *closed*		
https://buildserver	2:443 redirects to h	nttp://www.thunder.com/	-
Affected files			
//depot/thunder/main/web/dpgroot/index-redirect.php#1 add			
•			

ecreport extracts information about a particular job from the ElectricFlow server and writes an HTML report to standard output The report is organized as a 2-dimensional table summarizing the results of steps in the job. The table layout can be customized by defining its format using the ectool (or ec-perl) --load option. Invoke the program in the top-level directory of the job's workspace so it can access log file extracts in that workspace (or, use the --dir option).

The following options are available for use with the *ecreport* utility:

Option	Description
culprit	If specified and the job failed, the report will be emailed to each user mentioned in the updates file. The value of this option provides the subject line for the message (%s will be replaced with the job name).
data	Name of file containing job data for the report. If specified, this data is used in place of querying the ElectricFlow server. Used primarily for testing.
dir	Change to this working directory before doing anything else.
help	Print this message and exit without doing anything.

Option	Description
jobId	Identifier for the ElectricFlow job to report on (required unlessdata is supplied). This is a UUID.
load	Name of a file containing Perl code to evaluate after option processing. Used primarily for debugging. There can be multipleload options.
emailConfig	The name of the Email Configuration to use when emailing the report. Defaults to "default".
mailto	Email the report to this addresses.
replyto	Value for the "Reply-To" field in emailed reports.
server	Location of the ElectricFlow server (hostname or hostname:port), to retrieve job data. Defaults to the COMMANDER_SERVER environment variable.
subject	Value for the "Subject" field in emailed reports. If the value contains a %s, a string summarizing the job replaces the %s.
updates	Name of a file containing information about recent updates reflected in the job. If specified, the contents of this file are included in the report.
version	Print ecreport version number.
webRoot	Base URL for the ElectricFlow Web server where the job was run, such as http://myServer. URLs in the report will refer to ElectricFlow Web pages relative to this URL.

Example 1

```
ecreport --jobId $[/myjob/jobId] --load ./rptformat.pl --updates update.log --
webRoot
```

http://myServer/commander > ecreport.html

This example writes a report to the file ecreport.html. Because this file has an .html extension, it will be displayed in the Reports section at the top of the Job Details page. The format of the table for this report is specified in rptformat.pl. An example of the contents from this file would be:

```
@::gTableRows = (
    ["",
         "Windows",
          "Linux",
         "Solaris"],
    ["Extract",
```

```
"span",
            "span",
            "update.log"],
    ["Compile",
            "windows-compile.log",
            "linux-compile.log",
            "solaris-compile.log"],
    ["Unit test",
            "windows-unittest.log",
            "linux-unittest.log",
            "solaris-unittest.log"],
    ["System test",
            "windows-systemtest.log",
            "linux-systemtest.log",
            "solaris-systemtest.log"]
);
```

The "Recent Updates" report section is specified in update.log. The --webRoot argument defines the root for any URL that appears in the report.

Example 2

This report will be sent as the body of an email message to user1@company.com with subject "Build Report".

Default Batch Reports (Run Reports)

These reports are installed during the ElectricFlow installation and automatically scheduled to run daily.

To access these reports: Select the Projects tab, then select the Electric Cloud project, then click the Reports subtab. This report collection graphs trends over time, including success/failure, build times, and resource consumption. These reports allow you to customize two levels of grouping into "variants." These variants may represent products, branches, locations, architectures, or any other descriptors you would like to use to group your workload.

Note: After a Default Batch Report runs, a link for each report appears on the Job Details page in the Links section. When you select a report from the Links section, it will be displayed in full-screen view.

Report types

Cross Project Summary - This report shows the status of multiple variants over the past 30 days. This high-level summary of each of the 30 days shows green, yellow, or red depending on the result of the best build of the day. The summary also shows average times, total invocations, and other data.

Variant Trend - This report shows more details about a particular variant such as build outcome over time, build quantity over time, elapsed build times over time, and a list of actual jobs that ran during the period.

Daily Summary - This report shows the most recent results of each variant, including the last successful build.

Resource Summary - This report shows resource utilization over the last 24 hours. You can see which resources are idle and which resources are busy. ElectricFlow has advanced features that allow you to select resources for your job at runtime, using late binding of property names. This feature provides unparalleled flexibility and control over your environment, but could make it difficult to find out which resources were used. This report provides the visibility you need into resource loading.

Resource Detail - This report shows the utilization of a particular resource over the course of one day.

Wait Time - This report shows cumulative wait time due to license, resource, and workspace waits. It also shows the number of steps that were delayed for any of those reasons.

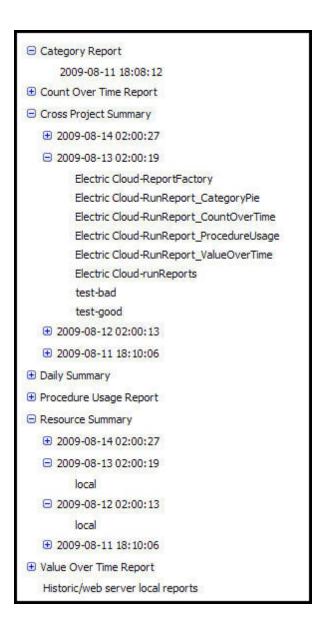
Default Batch Reports

- Cross Project Summary
 - o Variant Trend
- Daily Summary
- Resource Summary
 - Resource Detail



The navigation tree on the left-side of the screen allows you choose the report and date of interest. These reports, which show long term trends over all projects and resources, are run from the runReports procedure in the Electric Cloud project.

The following example shows the Default Batch Reports, the historical reports link, and Optional Batch Reports (described below). Notice that you can expand a selected report to three levels. The third level represents drill-down links available within that report. In addition, if you click on a report name, the most current report is displayed.



Tip

You may change the number of reports displayed under each report type. Locate the propertySheet named ec_reportConfiguration on the project or server (project overrides server) and find the "groupings" section. Next, you can create nested property sheets under ec_reportConfiguration with the properties groupingName, sortDirection, and limit. If you set this information using ectool, you will also see it on ElectricFlow's web interface.

If the $ec_reportConfiguration$ propertySheet does not exist, you will need to create it, then copy the example below into the propertySheet.

Example:

ectool setProperty ec reportConfiguration/reportDaysLimit --value 90 --projectName foo

```
ectool setProperty ec reportConfiguration/groupings/1/groupingName --value reportType
--projectName foo
ectool setProperty ec reportConfiguration/groupings/1/sortDirection --value ascending
--projectName foo
ectool setProperty ec reportConfiguration/groupings/1/limit --value 100 --projectName
foo
ectool setProperty ec reportConfiguration/groupings/2/groupingName --value reportDate
--projectName foo
ectool setProperty ec reportConfiguration/groupings/2/sortDirection --value descending
--projectName foo
ectool setProperty ec reportConfiguration/groupings/2/limit --value 5 --projectName
ectool setProperty ec reportConfiguration/groupings/3/groupingName --value
reportResource --projectName foo
ectool setProperty ec reportConfiguration/groupings/3/sortDirection --value ascending
--projectName foo
ectool setProperty ec reportConfiguration/groupings/3/limit --value 100 --projectName
foo
```

This example explicitly defines the default behavior for project foo's reports.

The 1, 2 and 3 propertySheets define the order of the groupings. The actual names 1, 2, 3, could be anything as long as they are listed alphanumerically in the order you need.

Additionally, any groupings that do not exist will be skipped during output. For example, if you added a reportRunByUser grouping #4 and added it to some reports that define reportType="User Builds", reportDate, and reportRunByUser properties (but not reportResource), you would get something similar to a resource usage report structure with users rather than resources at the deepest level.

Note: If you choose to display 500 or more reports at one time, you may experience some performance degradation.

Advanced reporting information - ecrptdata

How are default batch reports generated?

Reports are scheduled as part of a special procedure, runReports, in the Electric Cloud project. This procedure has one step, runEcrptdata, that runs the ecrptdata program. This program reads data from ElectricFlow, summarizes it for reports, and generates HTML reports using the BIRT reporting engine. The program takes input parameters. During installation, reporting is configured as follows:

```
ecrptdata --end yesterday --grain 300 --debug 2 --reports all
```

By default, all reports are run every day at 2am. The Cross Project Summary and Variant Trend shows data from the previous 30 days. The Daily Summary, Resource Summary, and Resource Detail reports show data from the last full day (yesterday). You can adjust the schedule to run more often or at different times by altering the "Report Schedule" schedule in the Electric Cloud project.

Changing the report time period

If you would like to alter the report time period, you must change the input parameters to the ecrptdata program within the procedure step "Electric Cloud:runReports:runEcrptdata". Select the Ecrptdata step name to go to the Edit Step page. Make the changes yo need in the Parameters section. If the --end option

is omitted, the end of the reporting period is the time when the report runs (now). If the keyword "yesterday" is provided, reports will include full data for yesterday. Otherwise, a specific date and time can be given in the format "YYYY-MM-DDTHH:MM:SS".

Changing the report grouping

You can tell ElectricFlow how to group the Cross Project Summary and Variant Trend reports. By default, the outer grouping is "Project Name" and the inner grouping is "Procedure Name." To change these groupings, use the following ecrptdata options:

group1Name xxxx	Set the column header for the first column of the Cross Project Summary report to "xxxx".
group1Property yyyy	Use property "yyyy" as the name to group on in the first column of the Cross Project Summary report.

All lookups are done with respect to an individual job and will be converted to absolute property paths. The values projectName, procedureName, and jobId are set in context for the job being examined.

Examples:

To show the value of the procedure name as the outer grouping and the value of the parameter "branch" as the inner grouping,

use the following ecrptdata options:

```
--group1Name Procedure --group1Property "/myProcedure/procedureName"
--group2Name Branch --group2Property /myParameter/branch
```

Note: The <code>/myParameter</code> "shortcut" to the property path was created specifically for <code>ecrptdata</code> only—it cannot be used as other "<code>/my</code>" property shortcuts are used. For more information on property shortcuts, see the Properties Help topic. In addition, when working with <code>ecrptdata</code>, any <code>/my</code> value is case-insensitive. For example, you can use <code>/myParameter</code> or <code>/myparameter</code>.

Filtering

ecrptdata supports "findObjects" type filters when searching for jobs and job steps.

This feature is implemented with two new command line switches: "--jobFilter <property, operator, operand>" and "--stepFilter <property, operator, operand>". The jobFilter option filters the CrossProjectSummay, DailySummary, and VariantTrend reports. The stepFilter option filters the ResourceSummary and ResourceDetail reports.

Examples:

```
--jobFilter "branch,isNotNull,1"
--jobFilter "jobName,like,my-procedure%"
--stepFilter "outcome,notEquals,aborted"
```

Adding additional data columns to the report

You can add custom data columns to each report. These columns are specified by providing a comma separated list of property references in the following ecrptdata options:

crossProjectProperties	sets extra columns for the Cross Project Summary Report
variantTrendProperties	sets extra columns for the Variant Trend Report
dailySummaryProperties	sets extra columns for the Daily Summary Report
resourceSummaryProperties	sets extra columns for the Resource Summary Report
resourceDetailProperties	sets extra columns for the Resource Detail Report

You must choose properties appropriate for the type of report and grouping you have chosen. For example, if you choose Project Name and Procedure Name for your groupings in the Cross Project Summary report, each line of the report will be in the context of a project or procedure. You can add extra columns to show properties on the project or procedure, but it would not make sense to show properties on a resource, job, or job step. For resource reports, the properties must be in the context of a resource.

The column header is the name of the property.

Examples:

```
--crossProjectProperties "/myproject/QA_State,/myprocedure/PartNumber" --resourceSummaryProperties "/myresource/OS,/myresource/Architecture"
```

Secure login

ecrptdata provides options for secure login to the ElectricFlow server.

Currently, there are three ways to login:

1. Use both the "--user <user>" and "--pass <password>" options.

This method is not secure... the password will be in plain text and visible to others.

2. If you omit the "--pass <password>" option, the password will be read from stdin. You can then read the password from a secure file or from a stored credential.

For example:

```
cat mySafePasswordFile.txt | ecrptdata --user <user> ...
ectool getFullCredential myCred --value password | ecrptdata --user <user> ...
```

3. Another option, "--credential <credentialName>", allows you to specify a credential name. If you use this option, "--user <user>" and "--pass <password>" are ignored. Note that the credential must be attached to the runReports job step. Both the user and password are retrieved from the credential.

Logic:

- 1. If --credential is used, attempt to lookup credential and extract user and password
- 2. If user is not set, return error
- 3. If password is not set, read from stdin
- 4. Attempt to login to server specified in --server <server> option using user and password

Optional Batch Reports

Additional reports are available, but these reports are not scheduled to run by default. You can schedule these reports to suit your purpose.

Note: When you run an Optional Batch Report, a link for each report appears on the Job Details page in the Links section. When you select a report from the Links section, it is displayed in full-screen view. The following optional reports are available:

- Category Report
- Procedure Usage Report
- Count Over Time Report
- Multiple Series Report (formerly known as the "Value Over Time" report)

These reports are available from the Reports tab on the Project Details page for your Project.

Note: For the Category, Count Over Time, and Multiple Series reports, a **View Search Results** link is available at the bottom of the report if a Saved Filter was used to generate the report.

Creating a Report Job

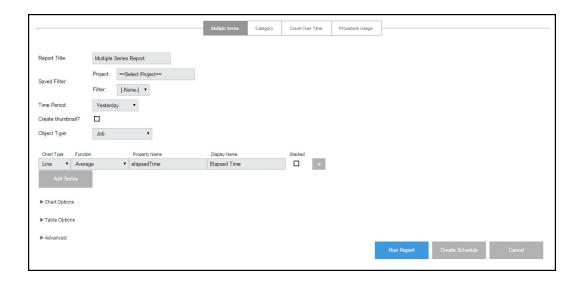
ElectricFlow runs reports as a job, so reports can be managed like any other job.

To create a report for your project:

- 1. Go to the Project Details page for your selected project.
- 2. Select the Reports subtab.
- 3. Click the Create Report link.



After clicking **Create Report**, the next screen is displayed with the default values for the selected tab (Multiple Series) in this example.



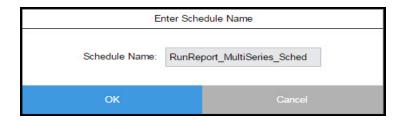
Enter information in the fields or select the appropriate values as follows:

- Report Title Your report title. Type over the default report name, choosing any unique name for your report. We supplied Basic Build Report for our report name.
- Saved Filter:
 - Project Click your mouse inside this field to see a list of projects from which to make a selection. We select our FirstNewProject project name.
 - o Filter Use the BasicBuildFilter we created earlier.
- Time Period Use the drop-down menu to select the time period.
- Create thumbnail? Check this box so you can view this report on your Home page. (We will create a "thumbnail" report view on your Home page at the end of this scenario.)
- Object Type Use the drop-down menu to select Job for this example.
- Table column choices:
 - Chart Type Use the drop-down menu to choose the chart type.
 - Function Use the drop-down menu to choose the function you need.
 - Property Name Use the default property value or delete the text and click your mouse in the blank field to see a list of possible properties. We chose "outcome" for this example.
 - o Display Name You can choose a different unique Display Name if you prefer to do so.
 - o Stacked Select this check box to see your report results "stacked" versus overlaid.

- The "X" icon Click this icon to delete any row you no longer need.
- Select the Add Series button if you would like to create additional table entries for additional report information.
- Chart Options Use the down-arrow to adjust the Time Grouping and see the defaults for the X and Y axis.

Buttons at the bottom of the page:

- Run Report button this button takes you to the Job Details page to run the report immediately—one time only.
- Click the Create Schedule button this button displays a box with a default schedule name you can change.



 Enter a name for the schedule, Basic Build Report Schedule—again, we tied the schedule name to the procedure name for which we want the report.

Viewing the report

After the report job runs, view the report by selecting the Reports subtab within the project. **Note:** All reports are viewed from within a project context. If you have reports that span multiple projects, run them from the Electric Cloud project or create a project to store summary reports.

After selecting the Reports subtab, see the navigation tree on the left-side of the page. This tree allows you to select the report you want to view. The report tree contains three levels.

- Report Title Click the "+" icon next to the report type to see available report versions.
- Report Date Click this link to see report details.

Note: Click the report type to display the most recent report version.

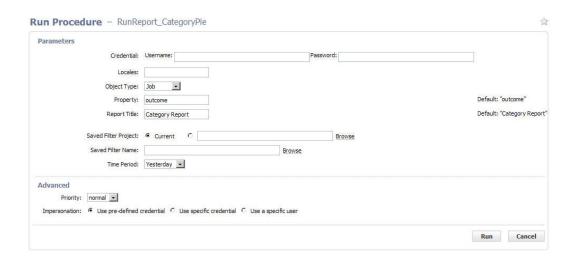
Optional Batch report examples

Category sample report

This report shows the distribution of property values in a pie chart.



Creating this report on the Run Procedure page:



Field descriptions:

Credential - these are credentials to be used when collecting data. If no credentials are specified
(Username/Password remain blank), the report runs with the privileges of the project from where it is run.
If credentials are supplied, the data is collected using the ElectricFlow user security access in the
credential. Use credentials when you want to access data outside of your project.

- Locales -a comma separated list of locales where you would like to render the report. Locale is specified as one or more of the following: en, zh, ja, or ko, which translate respectively to English, Chinese, Japanese, or Korean. If no locale is specified, the default is to use the locale of the ElectricFlow server. For example, a server already set for Chinese will see reports in Chinese with no other changes required. Note: Currently, ElectricFlow supports four locales only.
- Object Type the type of object to search (default job)
- Property the property of interest (default outcome) and must be a valid property for the object type selected
- Report Title your report title (default Category Report)
- Saved Filter Project select "Current" or Browse to find the project where you want to save this filter.
- Saved Filter Name the name of a saved filter to use when filtering collected data.

Use the Search tab to create the filter. After entering your search criteria, click **OK**. If the preview displays the data you want, click **Save Filter** and provide a Project Name and Filter Name to specify where the filter is saved.

On the Run Report parameters page, select the filter. If you leave the Filter Name blank, no filtering is performed.

Note: "Saved filters" can filter on start and end times. Because start and end times are specified in the Time Period parameter, there is a possibility you could select a combination of filter date and time periods that exclude all data. For example, you could select a time period of "Yesterday" and a filter of Start=less than Jan 1, 1900. To avoid this problem, we recommend you do NOT include dates or times in your filter.

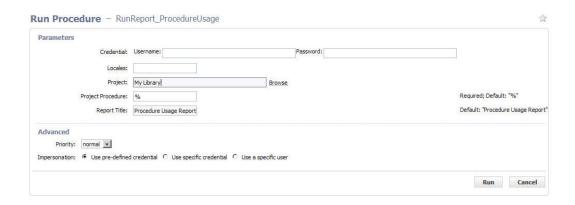
- Time Period these are keywords to define the report start and end times. The range is used to filter collected data. All values are in the local time zone.
 - Today midnight until now
 - · Yesterday previous full day
 - . This week from the start of the previous Sunday until now
 - Last week from Sunday to Sunday of the last complete week
 - This month from the start of the first of this month until now
 - Last month from the last complete calendar month, which means on January 3rd the last full month is all of December
 - This year from January 1st until now
 - Last year all of last year (Jan 1 to December 31)
 - All all dates until now

Procedure Usage Sample Report

For a specific procedure, this report shows procedures called by the procedure and which procedures it called.



Creating this report on the Run Procedure page:



Field descriptions:

- Credential these are the credentials to use for collecting data. If no credentials are specified
 (Username/Password remain blank), the report runs with the privileges of the project from where it is run.
 If credentials are supplied, data is collected using the ElectricFlow user security access in the credential.
 Use credentials when you want to access data outside of your project.
- Locales a comma separated list of locales where you would like to render the report. Locale is specified as one or more of the following: en , zh, ja , or ko, which translate respectively to English, Chinese, Japanese, or Korean. If no locale is specified, the default is to use the locale of the ElectricFlow server. For example, a server already set for Chinese will see reports in Chinese with no other changes required. Note: Currently, ElectricFlow supports four locales only.
- Project this is the name of the project where you want to install the report.
- Project Procedure enter a project procedure. "called" and "called by" data is shown for that
 project/procedure combination. You can use "wildcards." For example, to see data for ALL projects and
 procedures, enter the % character in each field.
- · Report Title this is the title of the report

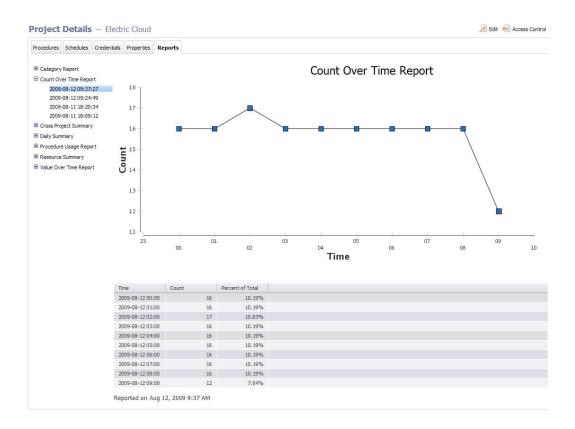
Count Over Time Sample Report

This report shows search results from counting multiple job, job step, or resource objects within the selected time range.

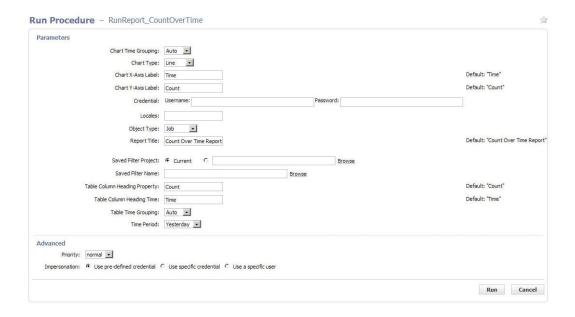
For example, this report counts the number of occurrences for the specified object-type over the selected time range--if a filter exists, it will be applied. This report does not aggregate the value (use ValueOverTime if you want aggregation), each occurrence counts as one.

Use this report when the property contains string values. The ValueOverTime report can be used only for numeric properties.

Note: Generating large numbers of datapoints can exceed BIRT's charting capability.



Creating this report on the Run Procedure page.



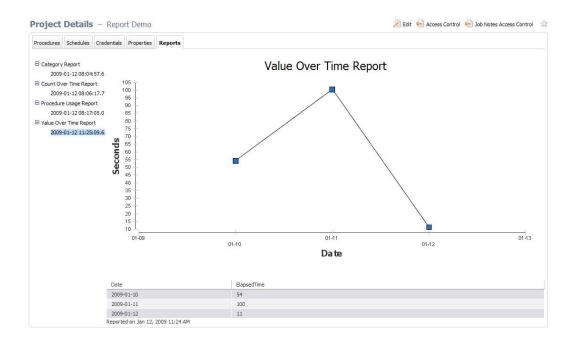
Field descriptions:

- Chart Time Grouping these are: Auto (default), Minute, Hour, Day, Week, Month, or Year
- Chart Type these are: Line (default), Bar, Cone, Tube, Triangle, or Area
- Chart X-Axis Label the text to display on chart x-axis (default Date)
- Chart Y-Axis Label the text to display on chart y-axis (default Count)
- Credential credentials to be used when collecting data. If no credentials are specified
 (Username/Password remain blank), the report runs with the privileges of the project from where it is run.
 If credentials are supplied, the data is collected using the ElectricFlow user security access in the
 credential. Use credentials when you want to access data outside of your project.
- Locales a comma separated list of locales where you would like to render the report. Locale is specified as one or more of the following: en, zh, ja, or ko, which translate respectively to English, Chinese, Japanese, or Korean. If no locale is specified, the default is to use the locale of the ElectricFlow server. For example, a server already set for Chinese will see reports in Chinese with no other changes required. Currently, ElectricFlow supports four locales only.
- Object Type the type of object to search (default job)
- Report Title your report title (default Count Over Time Report)
- Saved Filter Project select "Current" or Browse to find the project where you want to save this filter.
- Saved Filter Name the name of a saved filter to use when filtering collected data.//Use the Search tab to create the filter. After entering your search criteria, click OK.//If the preview displays the data you want, click Save Filter and provide a Project Name and Filter Name to specify where the filter is saved.// On the Run Report parameters page, select the filter. If you leave the Filter Name blank, no filtering is performed.//Note: "Saved filters" can filter on start and end times. Because start and end times are specified in the Time Period parameter, there is a possibility you could select a combination of filter date and time periods that exclude all data. For example, you could select Time Period of "Yesterday" and a filter of Start=less than Jan 1, 1900. To avoid this problem, Electric Cloud recommends you do NOT include dates or times in your filter.

- Table Column Heading Property the table heading for the right column (default Count)
- Table Column Heading Time the table heading for the left column (default Date)
- Table Time Grouping these are: Auto (default), Minute, Hour, Day, Week, Month, Year
- Time Period these are keywords to define the report start and end times. The range is used to filter collected data. All values are in the local time zone.
 - Today midnight until now
 - · Yesterday previous full day
 - This week from the start of the previous Sunday until now
 - Last week from Sunday to Sunday of the last complete week
 - . This month from the start of the first of this month until now
 - Last month from the last complete calendar month, which means on January 3rd the last full month is all of December
 - This year from January 1st until now
 - Last year all of last year (Jan 1 to December 31)
 - All all dates until now

Multiple Series Report

This reports shows the cumulative value of a property over time. This report should be run against properties that contain numbers. (Use the Count Over Time report if your property contains string values.) You can specify the function to use for aggregating values (sum, count, average). For example: Assume a property that contains the number of failed unit tests is stored on each job. Using the Multiple Series report, you can see the failure trend over time.



Run Cancel

Run Procedure - RunReport_ValueOverTime Parameters Chart Time Grouping: Auto Chart Type: Line Chart X-Axis Label: Date Default: "Date" Chart Y-Axis Label: Seconds Default: "Seconds" Credential: Username: Locales: Object Type: Job Property: elapsedTime Default: "elapsedTime" Property Expression: Property Function: Average Report Title: Value Over Time Report Default: "Value Over Time Report Browse Saved Filter Name: Table Column Heading Property: Default: "Date"

Creating this report on the Run Procedure page:

Table Column Heading Time: Date

Table Time Grouping: Auto

Time Period: Yesterday

Impersonation: ● Use pre-defined credential ○ Use specific credential ○ Use a specific user

Field descriptions:

Advanced

- Chart Time Grouping these are: Auto (default), Minute, Hour, Day, Week, Month, or Year
- Chart Type these are: Line (default), Bar, Cone, Tube, Triangle, or Area
- Chart X-Axis Label the text to display on chart x-axis (default Date)
- Chart Y-Axis Label the text to display on chart y-axis (default Count)
- Credential these are credentials to be used when collecting data. If no credentials are specified
 (Username/Password remain blank), the report runs with the privileges of the project from where it is run.
 If credentials are supplied, the data is collected using the ElectricFlow user security access in the
 credential. Use credentials when you want to access data outside of your project.
- Locales a comma separated list of locales where you would like to render the report. Locale is specified as one or more of the following: en, zh, ja, or ko, which translate respectively to English, Chinese, Japanese, or Korean. If no locale is specified, the default is to use the locale of the ElectricFlow server. For example, a server already set for Chinese will see reports in Chinese with no other changes required. Note: Currently, ElectricFlow supports four locales only.
- Object Type the type of object to search (default job)
- Property the property of interest (default outcome) and must be a valid property for the object type selected
- Property Expression an optional mathematical function to apply to your data. For example, elapsed time is recorded in milliseconds. To make a useful report with elapsed times, you may want to show seconds, so this field would be entered as "/1000".

• Property Function - the following list contains the aggregation functions to use for data in the chart and table.

Count	Median
Count Distinct	Mode
First	Std Deviation
Last	Sum
Min	

- Report Title your report title (default Value Over Time Report)
- Saved Filter Project select "Current" or Browse to find the project where you want to save this filter.
- Saved Filter Name the name of a saved filter to use when filtering collected data.

Use the Search tab to create the filter. After entering your search criteria, click **OK**. If the preview displays the data you want, click **Save Filter** and provide a Project Name and Filter Name to specify where the filter is saved.

On the Run Report parameters page, select the filter. If you leave the Filter Name no filtering is performed.

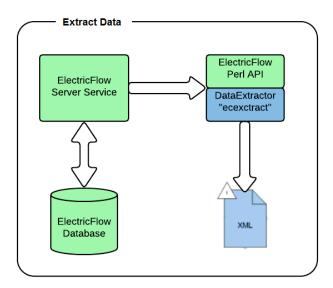
Note: "Saved filters" can filter on start and end times. Because start and end times are specified in the Time Period parameter, there is a possibility you could select a combination of filter date and time periods that exclude all data. For example, you could select Time Period of "Yesterday" and a filter of Start=less than Jan 1, 1900. To avoid this problem, Electric Cloud recommends you do NOT include dates or times in your filter.

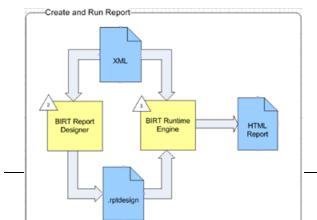
- Table Column Heading Property the table heading for the right column (default Count)
- Table Column Heading Time the table heading for the left column (default Date)
- Table Time Grouping these are: Auto (default), Minute, Hour, Day, Week, Month, Year

- Time Period these are keywords to define the report start and end times. The range is used to filter collected data. All values are in the local time zone.
 - · Today midnight until now
 - · Yesterday previous full day
 - This week from the start of the previous Sunday until now
 - Last week from Sunday to Sunday of the last complete week
 - This month from the start of the first of this month until now
 - Last month from the last complete calendar month, which means on January 3rd the last full month is all of December
 - This year from January 1st until now
 - Last year all of last year (Jan 1 to December 31)
 - All all dates until now

Creating Custom reports

The ElectricFlow batch mode reports which, by default, run once each night and are accessed through the Reports tab [in the Electric Cloud project] in the web UI. These reports are built using a combination of the ElectricFlow API and the BIRT Report Designer tool. BIRT is an open source Eclipse-based reporting system you can use to create your own custom reports, or modify existing ElectricFlow reports. This section describes how to write your own custom report and integrate it into ElectricFlow.





There are several steps to using BIRT to create reports:

- · Create an XML file with data using the ElectricFlow Perl API
- Design a report using BIRT RCP Report Designer
- Run the report with the BIRT Report Engine
- · View the report from within ElectricFlow

Data extraction

To begin creating a custom report, you must first extract report data from ElectricFlow and move it into a standalone XML file. The recommended way to create the XML file is to use the ElectricFlow ecextract utility program, which is derived from the Perl API. The benefits of using ecextract to create XML files are:

- Security ecextract enforces all ElectricFlow security and ACLs. (Direct DB access does not honor ACLs.)
- Maintainability ecextract is supported by Electric Cloud, including backward compatibility and insulates you from database schema changes.
- Simplicity ecextract provides easy-to-use search and filtering options. Because of the complex nature of our schema, particularly for properties, it may be difficult to create custom SQL statements.
- Reuse You may find that XML files created by built-in reports already contain all the data you need. If
 you simply want to present data differently or "slice" the data in a different way, you can reuse the XML
 from the existing built-in reports.

Note: If you already use ectool or Perl for data extraction, and have already created the scripts you need for custom reports, you may continue to use those scripts. However, we recommend using ecextract if you are new to creating custom reports.

Using ecextract.pl for data extraction

A Perl program, ecextract.pl, is provided so you do not have to write Perl code—or you can use it as a starter for your own Perl API calls. This script allows you to specify the most common query types on a command line and generates an XML file. The following table lists currently available ecextract arguments:

Usage: ecextract [options]

Data Extract Options:	
type job jobStep resource	The object type to query.
property [label=]property	Add a property to the output. If label is specified, it is used as the XML tag name instead of the property name.
filter property, operator, operand1 [,operand2]	Add a property filter to the query.

savedfilter filterName	Use a saved filter (ignores filter)		
period period	The time period. One of: today, yesterday, thisweek, lastweek, thismonth, lastmonth, thisyear, lastyear, all, window, <number days="" of="">, or range, <start-time>, <end-time> Range start, end time is date in ISO 8601 format: YYYY-MM-DDTHH:MM:SS[Z]If the date ends with 'Z', it is considered to be GMT time, otherwise the local timezone is used.</end-time></start-time></number>		
sort property, order	Add a sort to the query. Order is either 'ascending' or 'descending'.		
outputPath	Path and name of output file. Default is out.xml.		
max	Maximum number of objects to return.		
Server Communication Options:			
server server	Address for the ElectricFlow server. Defaults to the value of the COMMANDER_SERVER environment variable. If it does not exist, defaults to localhost.		
port port	HTTP listener port on the server. Defaults to 8000.		
securePort port	HTTPS listener port on the server. Defaults to 8443.		
secure 0 1	If set to true, HTTPS is used to communicate with the server. Defaults to true.		

timeout seconds	Set the number of seconds for the timeout. Defaults to 280.		
Login Options:			
user user	ElectricFlow login user.		
pass password	Password for login user.		
credentialName	Credential name (valid only when running inside a job step).		
Global Options:			
help	Print this message and exit without doing anything.		
version	Debug level. 0 - errors only 1 - 0 + System progress 2 - 1 + Property Warnings 3 and higher give detailed diagnostics		
Plugin File Options:			
load loadFile	Plugin perl program to execute		
	Stop processing ecextract options. Allows plugins to process their own options following this option.		

ecextract.pl is part of the EC-Reports plugin—this script is located in the "agent/bin" directory of the EC-Reports plugin.

Notes:

--property

Properties can be specified as high-level values (--property propertyName), properties in a sheet (--property sheet1/sheet2/propertyName), or as relative paths (--property /myProject/propertyName).

--filter

Each entry consists of a property name to test and an operator to use for comparison. The property can be an intrinsic property defined by ElectricFlow or a custom property added by the user. Each operator takes zero or one operand to compare against the desired property.

The operators are:

```
contains, equals, greaterOrEqual, greaterThan, in, lessOrEqual, lessThan, like, notEqual, notLike, isNotNull, isNull

Example: --filter jobName, equals, Hello
```

Examples

The following 3 examples contain the script input and the output.

Example 1: Find properties on job steps for one day in April

```
ec-perl ecextract.pl ^
  --user report ^
  --pass report ^
  --server localhost ^
  --outputPath jobStepOut.xml ^
  --type jobStep ^
  --max 1000 ^
  --property jobId ^
  --property finish ^
  --property jobStepId ^
  --property stepName ^
  --property status ^
  --property outcome ^
  --property runTime ^
  --property waitTime ^
  --filter jobName, like, commander-2.2%% ^
  --filter createTime, greaterOrEqual, 2008-04-01T00:00:00.000Z ^
  --filter createTime, lessOrEqual, 2008-04-02T00:00:00.000Z ^
  --sort jobId, descending
Output from example 1
<?xml version="1.0" encoding="utf-8"?>
<response>
    <object objectType="jobStep" objectId="jobStep-3511785">
        <jobId name="jobId">4fa765dd-73f1-11e3-b67e-b0a420524153</jobId>
        <finish name="finish">2009-04-01T23:54:45.424Z</finish>
        <jobStepId name="jobStepId">4fa765dd-73f1-11e3-b67e-b0a420524153</jobStepId>
        <stepName name="stepName">Setup</stepName>
        <status name="status">completed</status>
        <outcome name="outcome">success</outcome>
```

Example 2: Find all jobs that were created in April 2008

```
ec-perl ecextract.pl ^
--user report ^
--pass report ^
--server localhost ^
--outputPath jobOut.xml ^
--type job ^
--max 100 ^
--property jobName ^
--property jobId ^
--property projectName ^
--property procedureName ^
--property branch ^
--filter createTime, greaterOrEqual, 2009-04-01T00:00:00.000Z ^
--filter createTime,lessOrEqual,2009-04-30T00:00:00.000Z ^
--sort projectName, ascending ^
--sort procedureName, ascending ^
--sort createTime, descending
```

Output from example 2

```
<?xml version="1.0" encoding="utf-8"?>
<response>
   <object objectType="job" objectId="job-158461">
       <jobName name="jobName">commander-3.2.23916-200904011650</jobName>
       <jobId name="jobId">4fa765dd-73f1-11e3-b67e-b0a420524153</jobId>
       projectName name="projectName">Commander
       cedureName name="procedureName">Master
       <branch name="branch">3.2</branch>
   </object>
   <object objectType="job" objectId="job-158455">
       <jobName name="jobName">commander-anders-main-ad.23915-200904011627</jobName>
       <jobId name="jobId">4fa765dd-73f1-11e3-b67e-b0a420524153</jobId>
       projectName name="projectName">Commander
       cedureName name="procedureName">Master
       <branch name="branch">main
   </object>
</response>
```

Example 3: Find all resources

```
ec-perl ecextract.pl ^
 --user report ^
 --pass report ^
 --server localhost^
 --outputPath resourceOut.xml ^
 --type resource ^
 --max 100 ^
 --property resourceName ^
 --property stepCount ^
 --property stepLimit ^
 --property hostName ^
 --filter stepCount, greaterThan, 0 ^
 --sort resourceName, ascending
Output from example 3
<?xml version="1.0" encoding="utf-8"?>
<response>
```

BIRT Report Designer

Now that you have a dataset, you can create a BIRT report. You need an XML file to use while designing the report. If you did not create an XML data file yet, now is the time to do so. After you have sample data, you are ready to start designing a report.

Report definitions are created with the BIRT Report Designer, currently version 2.5.2. These definitions are then stored in .rptdesign files. You have several options for installing the designer tool:

- Use a BIRT Report Designer bundled and already installed with Eclipse, or
- download BIRT Report Designer to use with your existing Eclipse installation, or
- because Eclipse is not required, you can download only the BIRT Report Designer (this is the easiest and preferred method).

We recommend downloading the BIRT RCP Report Designer only, which is the standalone design tool. The url for downloading the Report Designer is:

http://www.eclipse.org/downloads/download.php?file=/birt/downloads/drops/R-R1-2_5_2-201002221500/birt-rcp-report-designer-2_5_2.zip

The general download page for BIRT 2.5.2 with the other options discussed is http://download.eclipse.org/birt/downloads/build.php?build=R-R1-2 5 2-201002221500

Using the BIRT Report Designer is beyond the scope of this Help topic. Documentation and tutorials can be found on the BIRT website: http://www.eclipse.org/birt. You may find it helpful to view the report designs for built-in ElectricFlow. These designs can be found in the agent/reports/<report name> directory in the EC-Reports plugin.

Understanding additional report components

In addition to the XML dataset and report design, there are optional report components you may wish to include in your report.

- External style sheet
 BIRT supports external style sheets for reports. If you use these style sheets, they need to be located with your report design. For example, built-in ElectricFlow reports use an external style sheet file "styles birt.css".
- Externalized string properties

 String literals (used in your report design) can be externalized into property files to support localization.

 These property files need to be located with your report design also. The properties file for the default locale will have the same name as your report and a ".properties" file extension, for example,

 CategoryPie.properties. Text for additional locales are stored in files named <report name>_

 <locale>.properties, for example, CategoryPie_ja.properties contains text localized for the Japanese language.

You may review examples of these additional components for built-in ElectricFlow reports. These files are located in the EC-Reports plugin, in the agent/reports/<report name> directory.

Packaging reports for ElectricFlow

After you have an XML dataset (using ecextract.pl), a BIRT report design (using BIRT Report Designer), and any additional optional report components completed, you are ready to integrate your report into ElectricFlow. To do this:

- Store all the report components in a location accessible to your ElectricFlow agent—you can use the ElectricFlow plugins directory. For example, for your report files, you could create a directory called "MyReport" in the ElectricFlow plugins directory.
- Create an ElectricFlow procedure to perform the necessary steps to generate a report, including:
 extracting data, running the report design, and registering report output to appear on the Report tab on
 the Project Details page. You may want to copy one of the procedures from the built-in ElectricFlow
 reports to use as a template. The following command creates a copy of the RunReport_CategoryPie
 procedure called RunReport_MyReport in the project of your choice. Substitute your project name for
 my-project-name in the following command.

Example: customizing the command body

The following is an example of customizing the procedure contents for the RunReport_MyReport procedure, cloned from the RunReport_CategoryPie procedure above.

One customization is required and two are optional. The remainder of the command body does not need to be changed.

Customizing the location of the BIRT report design file
 Required: update the location of the BIRT report design file.

```
    Original value: our $reportDesignFile =
        "$pluginDir/agent/reports/CategoryPie/CategoryPie.rptdesign";
    New value: our $reportDesignFile = "$ENV{'COMMANDER_PLUGINS'}
        /MyReport/MyReport.rptdesign";
```

· Customizing data extraction

To customize data extraction using ecextract.pl, modify the following section of the command body:

```
#-----
       # extractReportData
              Extract report data from Commander using ecextract utility
       #-----
       sub extractReportData()
                  # Time property varies based on object type
                  my $timeprop = "modifyTime";
                  if (
                           $timeprop = "finish";
                  }
                   my $command = "ec-perl \"$pluginDir/agent/bin/ecextract.pl\" "
                         . " --debug 3"
                         . " --outputPath \"$ ecextractXmlFile\""
                         . " --credential \"$[Credential]\""
                         . " --type \"$[Object Type]\""
                         . " --property time=$timeprop"
                         . " --property series1=\"$[Property]\"";
                  if ("$[Saved Filter]" ne "") {
                            command = $command . " --savedFilter \"$[Saved Filter]
\""
                  $command = $command . " --period \"$[Time Period]\"";
                   print "data extract [$command]\n";
                  my $result = `$command`;
                  print "extract result:\n$result";
                  if ($?) {
                       exit (1);
                  }
```

Customizing report parameters

To customize the parameters passed to your report design, modify the following section of the command body:

```
# getReportArgs
      Get report arguments
  -----
 sub getReportArgs() {
           # args for creating report
           my $categoryHeading = ucfirst($[Property]);
           my @args = (
                "-f", "$reportFormat",
                "-o", "$reportName",
                "-i", "images/",
                "-p", " xmlfile=../$ecextractXmlFile",
                "-p", "ReportTitle=$reportTitle",
                "-p", "TableColHeadingTime=$categoryHeading",
                "-p", "TableColHeadingSeries1=Count",
                "-p", " SearchURL=" . getSearchUrl(),
                "$reportDesignFile",
           );
            # add locale args
           foreach my $locale (@locales) {
                     if ($locale ne "") {
                          unshift @args, "-1 $locale";
                     }
           return @args;
```

Helper functions provided in ElectricCommander::ReportUtils.pm

ecgenReport - runs the BIRT report, using the BIRT Report Designer

- installationDirectory this is the directory where ElectricFlow is installed and where you will find JAVA and BIRT
- artifactDirectory this is the directory where reports, images, and other artifacts should be stored.
 To display using the Reports subtab, these items must be in a directory named "artifacts" at the top of the job's workspace
- args an array of arguments to the BIRT run-time engine, which specifies parameters, design templates, and so on.

registerReport - marks the report so it will be displayed when you select the Reports subtab

- commander an ElectricCommander() object used to make ElectricFlow calls. If you need to use a specific user context, call the login method before calling registerReport
- jobId this is the report job ID, which is used to set properties on the job
- url this is the URL to register and it takes the form workspace name/html file name
- title this is the text string for the report navigation tree
- treeGroupings a hash grouped name/values set as properties to define the report tree groupings—report tree groupings are defined on a propertySheet named ec_reportConfiguration on the project or on the server
 - hash key this is a property name that will be created on job's ec_reportData/\$title/ property sheet and used for report tree grouping
 - o hash value property value

registerArtifactsDirectory - creates the property required to set the artifacts directory

- commander an ElectricCommander() object used to make ElectricFlow calls
- jobId this is the report job ID, which is used to set properties on the job
- artifactsDir this is a directory path relative to the workspace where job artifacts will be stored

extractFile - this function pulls content from the report definition (stored in properties) and creates files from them

- commander an ElectricCommander() object used to make ElectricFlow calls. If you need to use a specific user context, make sure to call the login method before calling registerReport
- jobId this is the report job ID, which is used to set properties on the job
- reportType this is used to lookup report attributes
- files an array of hashes that specify which properties should be loaded into which files
 - o prop the property name from the installed report type
 - dest the destination file name (relative to the current working directory)
 - o expand this is "1" if you want ElectricFlow to expand property references in the property, "0" if not

getReportTimestamp - return hires timestamp formatted in a specified time zone

- TimeZone (optional) a time zone specifier, default is "local"
- FormatString (optional) a format specifier for the timestamp string

Custom Report Examples

This Help topic contains two examples illustrating how you might use the BIRT Report Designer to modify a built-in ElectricFlow report to create a custom report more meaningful for your purposes.

Example 1: modifying an existing report - adding a "banner" heading

The easiest way to customize a report is to take an existing report and modify it. ElectricFlow provides a set of built-in reports you can modify.

Process overview for modifying an ElectricFlow built-in report

- Make a copy of an existing report
- · Make changes to the report
- · Test your new custom report

To copy an existing ElectricFlow report

For this example, we will make a copy of the built-in Value Over Time report. The process to copy a report is:

• Make a copy of the file system report components of the Value Over Time report located in the ElectricFlow plugins directory.

Rename the copied directory to MyReport.

The default location of the ElectricFlow plugins directory is:

```
On Linux: /opt/electriccloud/electriccommander/plugins
```

 $\begin{tabular}{ll} On Windows: C:\Documents and Settings\All Users\Application Data\Electric Cloud\ElectricCommander\plugins \end{tabular}$

Linux example - The following commands will make a copy of the Value Over Time report file components and then rename the files:

```
$ cd /opt/electriccloud/electriccommander/plugins
$ cp -r EC-Reports-1.0.0.*/agent/reports/ValueOverTime/ MyReport
$ cd MyReport
$ rename ValueOverTime MyReport ValueOverTime*
```

The contents of the MyReport directory should be:

```
$ ls -1
MyReport_en.properties
MyReport_ja.properties
MyReport_ko.properties
MyReport.properties
MyReport.rptdesign
MyReport_zh.properties
```

```
styles_birt.css
```

Make a copy of the RunReport_ValueOverTime procedure from the EC-Reports plugin that performs the
steps to generate a report, including extracting data, running the report design, and registering the report
output to appear on the Report tab of the Project Details page. In this example, we will create a Project
called "MyReportProject" and the procedure we copy will be called "RunReport MyReport".

- Edit the "RunReport_MyReport" procedure in "MyReportProject" to reference the report design file copied earlier.
- Navigate to project MyReportProject, procedure RunReport_MyReport and select the RunCustomReport step to edit the Command block:

Find the following line:

```
our $reportDesignFile =
"$pluginDir/agent/reports/ValueOverTime/ValueOverTime.rptdesign";
```

and change this line to:

our \$reportDesignFile = "\$ENV{'COMMANDER_PLUGINS'}/ MyReport/MyReport.rptdesign";



Click **OK** to save the change.

To modify the copied report design

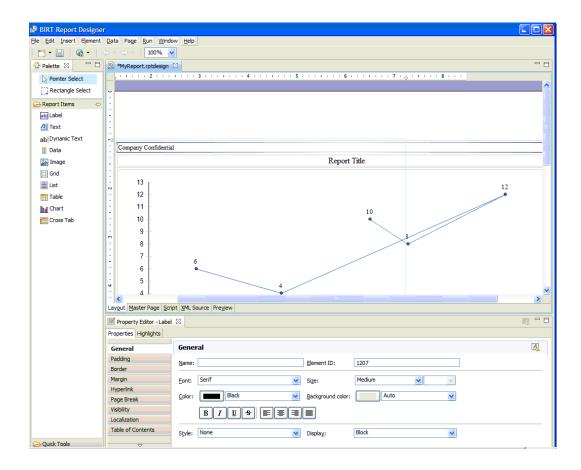
Start the BIRT Report Designer and open the report design file (MyReport.rptdesign) located in the MyReport directory previously created in the ElectricFlow plugins directory.

At this point You are now ready to begin editing an existing report to create a custom report.

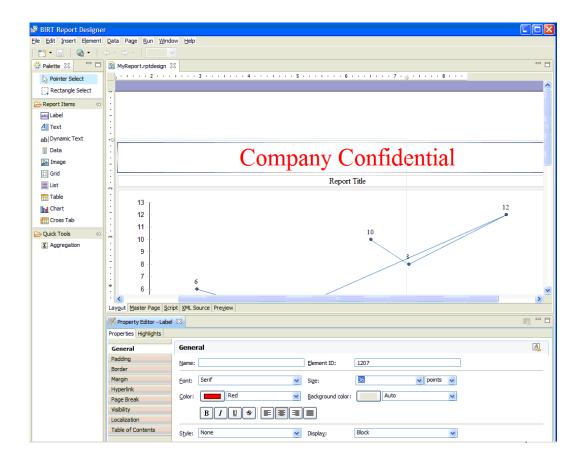
For our example, we are adding a new heading, "Company Confidential," (in large red letters) to the top of the existing report.

- 1. On the BIRT screen, select the Layout tab.
- 2. Drag a "Label" from the Palette window into the top of the report and type "Company Confidential" (or whatever text you would like to use for your test report) into the Label box.

See the next BIRT screen example.



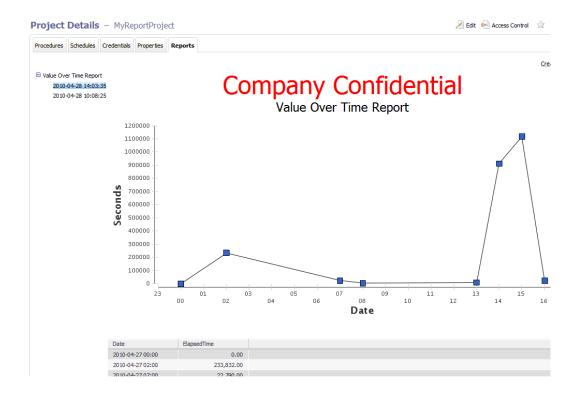
- 3. Now you can change the font size and color to create a banner appearance. Select the Property Editor window to customize your text see the next screen example.
- 4. After customizing your text, Save your report.



Congratulations! You have just created a custom report.

Test your new report

Run the report by running the RunReport_MyReport procedure in the MyReportProject project. When the job completes, your new report will be available in the Links section on the Job Details page and from the Reports tab if the MyReportProject project. You will see the new header in the report output.



Example 2: complete end-to-end example

Using BIRT and some scripting expertise, you can create almost any kind of report you can imagine. This example creates a new report that shows the run-time trend of all job steps. Tasks for creating this report can be grouped into the following steps:

- · Plan your report
- · Create a new BIRT rptdesign file
- Deploy your custom report
- · Test your new report

Planning this report example

We decided our report will be a single-series type report that will show job steps runtime. The report will include a line chart at the top of the page whose X and Y axis represent the number of seconds and the job step ID respectively. We also want a table with two columns—column names will be "Job Step Name" and "Run Time".

Next, we need to gather all materials:

Generate sample data to test in the report.

You can use data from a previous run of a Value Over Time report—this file is located in the workspace of a previous reporting job and has the name ecextract.xml.

Also from the copy of an ElectricFlow Value Over Time report, we will "copy and paste" a few values from it to save some typing.

Note: This example briefly describes items we are working with, but for more information about BIRT, you may need to consult BIRT documentation.

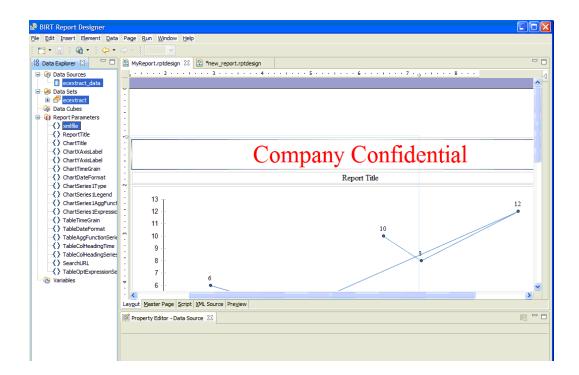
Creating a new BIRT rptdesign file

In BIRT Report Designer, create a new rptdesign file. Open the report file you created in "Example 1: modifying an existing report - adding a banner heading" so you can use the previous example report to copy some items into your new report.

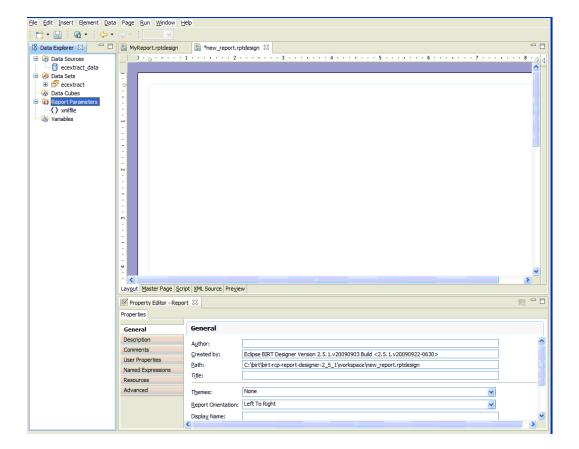
Items to copy into the new report:

- From Data Sources, copy ecextract_data
- From Data Sets, copy ecextract
- From Report Parameters, copy xmlfile

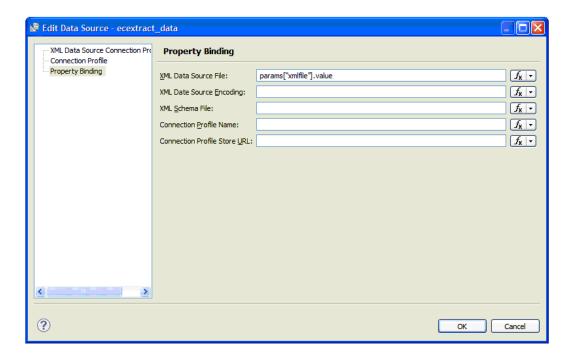
See the next screen example to help find these items.



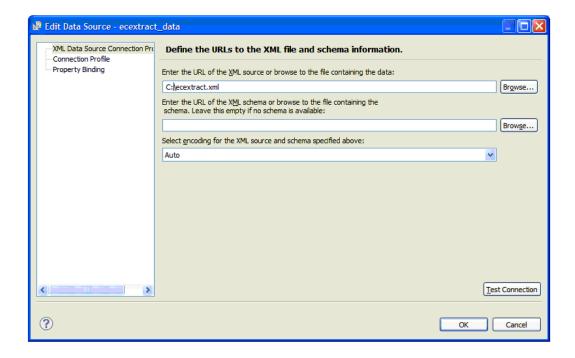
After copying these items, your new report design will look like the following example:



Double-click the Data Sources > ecextract_data file to ensure the XML file is referenced as a parameter. If it is not, type in the value as shown below.

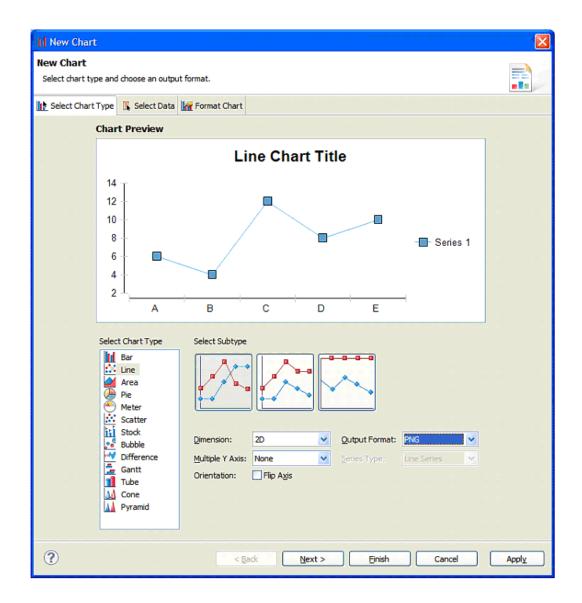


Double-click Data Sources > $ecextract_data$ to ensure the XML Data Source Connection refers to the data file you downloaded earlier. In this example the file is C: ecextract.xml.



From the Palette window, drag a Chart object into your report. Your view will be similar to the following:

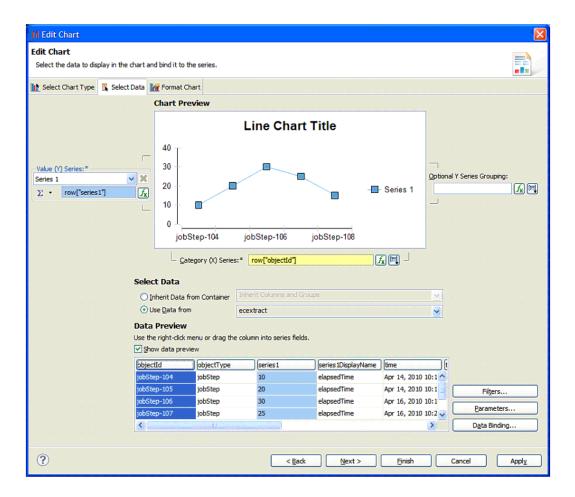
In the New Chart dialog, select Line under Select Chart Type and browse to select PNG for the Output Format.



Click Next.

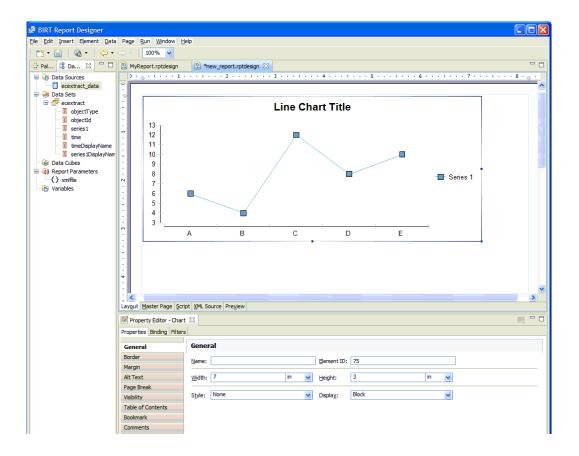
On the next screen:

- Under Select Data, browse to select ecextract.
- Make sure the Use Data From radio button is selected.
- Drag the "series1" header to the Value (Y) Series field.
- Drag the "objectId" header to the Category (X) Series field.



Click **Next** to apply any other formatting changes you choose, then click **Finish**.

On the next screen, after clicking Finish, you have to option to resize the chart as you wish.

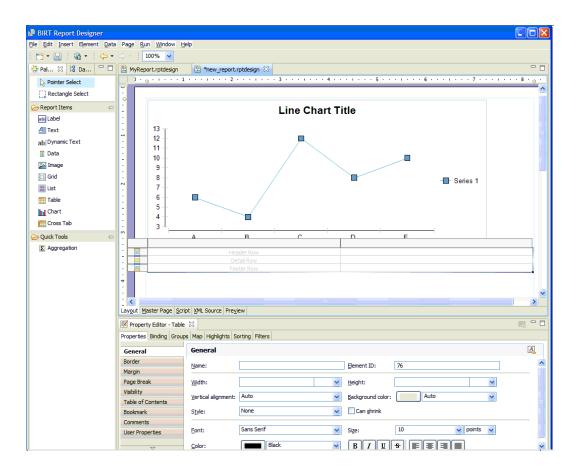


You are now ready to drag a Table item from the Palette to the report layout below the chart.

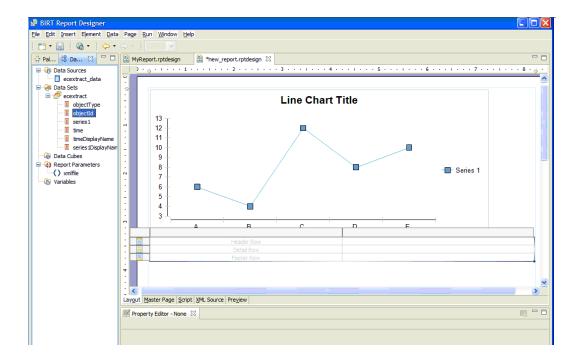
Select "2" columns and the ecextract dataset.

№ Insert Table	X	
Table size		
Number of <u>c</u> olumns:	2 🗘	
Number of <u>d</u> etails:	1 🗘	
Data <u>S</u> et:	ecextract 💉	
□ <u>U</u> se these dimensions for new tables.		
? ок	Cancel	

Click OK.



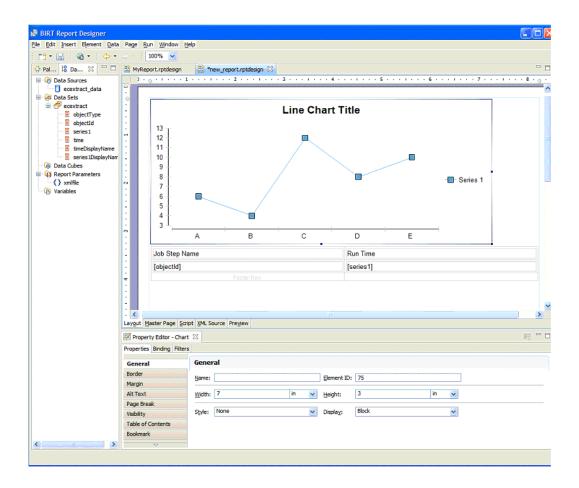
Select the Data Explorer tab and expand Data Sets > ecextract. See the next screen.



Now you are ready to:

- Drag "objectId" to the left Detail Row.
- Drag "series1" to the right Detail Row.
- Replace the Header Row for "objectId" and "series1" with Job Step Name and Run Time, respectively.

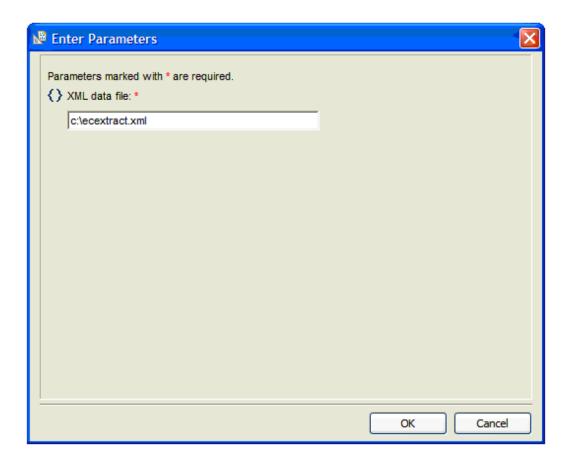
See the next screen example.

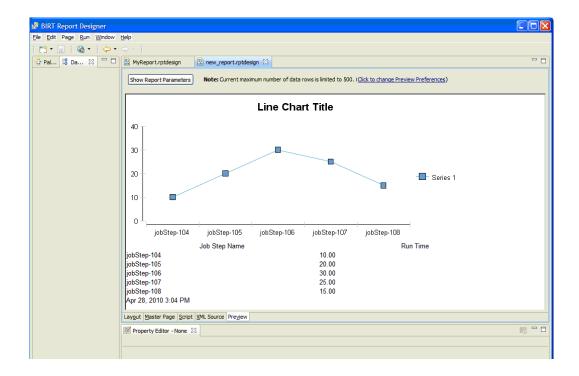


Select Preview to see your report.

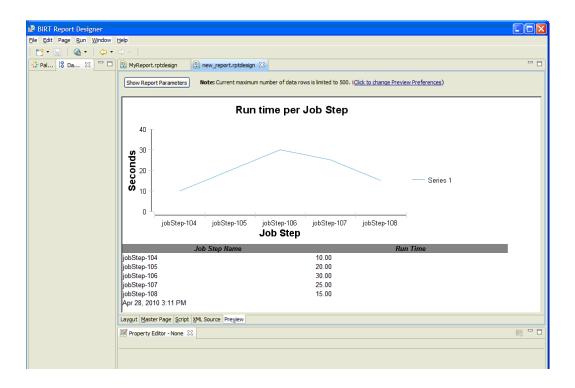
In the "Enter Parameters" dialog, type the location of your ecextract.xml data file.

Click OK.





At this point, you can continue to make formatting changes to make your custom report look exactly how you want it to look. Our example report, when finished, looked like the following example::



Deploy your custom report

If you have not already done so, save your new report design now. We have named our report design "JobStepRunTime.rptdesign".

The new report design needs to be stored in a location accessible to the ElectricFlow Agent that will run the report. As in the previous example, we will use the ElectricFlow plugins directory for this purpose.

 Save the JobStepRunTime.rptdesign file to a new JobStepRunTimeReport directory in the ElectricFlow plugins directory.

The default location of the ElectricFlow plugins directory is:

On Linux: /opt/electriccloud/electriccommander/plugins

On Windows: C:\Documents and Settings\All Users\Application Data\Electric Cloud\ElectricCommander\plugins

Linux example - The following commands will make a copy of the Value Over Time report file components and then rename the files:

- \$ cd /opt/electriccloud/electriccommander/plugins
- \$ mkdir JobStepRunTimeReport
- Now copy the previously saved JobStepRunTime.rptdesign file to the JobStepRunTimeReport directory.
- Make a copy of the RunReport_ValueOverTime procedure from the EC-Reports plugin that performs the
 steps to generate a report, including extracting data, running the report design, and registering the report
 output to appear on the Report tab of the Project Details page. In this example, we will create a Project
 called "JobStepRunTime" and the procedure we copy will be called "RunReport_JobStepRunTime".
- Edit the "RunReport_JobStepRunTime" procedure in project "JobStepRunTime" to reference the report design file copied earlier.
- Navigate to project JobStepRunTime, procedure RunReport_JobStepRunTime and select the RunCustomReport step to edit the Command block:



Find the following line:

```
our $reportDesignFile =
   "$pluginDir/agent/reports/ValueOverTime/ValueOverTime.rptdesign";
and change it to:
```

```
our $reportDesignFile = "$ENV{'COMMANDER_PLUGINS'}/
JobStepRunTime/JobStepRunTime.rptdesign";
```

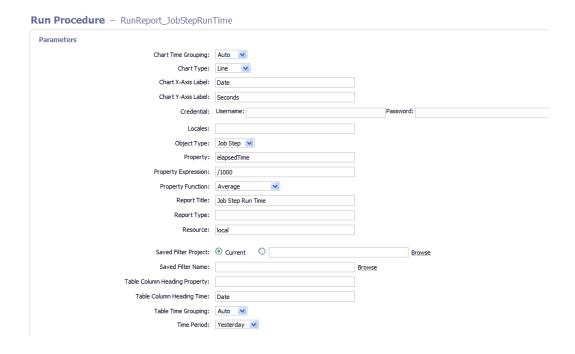
Edit Step — RunCustomReport				
General				
	Name:	RunCustomReport		
	Description:			
Command				
	Command(s):	100000000000000000000000000000000000000		
		<pre>if (%reportType eq "") { %reportType = %reportTitle; } our %reportFormat = "html"; #html,pdf,doc,xls</pre>		
		our freportName = safeFileName ("freportTitle freportEormat");		
		our fartifactsDir = "artifacts";		
		<pre>our \$agentInstallDir = getAgentInstallDir(); our \$pluginDir = "\$ENY{'COMMANDER PLUGINS'}/\$[/plugins/EC-Reports/pluginName]";</pre>		
		our *reportDesignFile = "%ENY("COMMANDER PLUGINS")/JOStepRunTime/JoStepRunTime.xptdesign";		
		our \$birtInstallDir = "\$ENY('COMMANDER_PLUGINS')/\$[/plugins/EC-ReportEngine]/agent";		
		our secentractMnlFile = safeFileName("ecentract-sreportType-s[/increment /myJob/xmlCounter].xml");		
		<pre>our @locales = getLocales();</pre>		

Click **OK** to save the change.

Test your new report

Run the report by running the RunReport_JobStepRunTime procedure in the JobStepRunTime project. On the Run Procedure page, change the values of the following parameters:

- Object Type => Job Step
- Property Expression => /1000
- Report Title => Job Step Run Time



Click Run.

When the job completes, your new report will be available in the Links section of the Job Details page and from the Reports tab in the JobStepRunTime project.



System Health Monitoring

You can also use StatsD and Graphite to generate custom reports for system health monitoring in ElectricFlow. After StatsD and Graphite are enabled, reports about a predefined list of items are generated.

This Help topic does not include instructions to set up StatsD and Graphite. Go to the following sites for documentation and tutorials:

- https://github.com/etsy/statsd/
- http://graphite.readthedocs.org/en/latest/index.html#
- https://www.digitalocean.com/community/tutorials/an-introduction-to-tracking-statistics-with-graphite-statsd-and-collectd
- https://www.digitalocean.com/community/tutorials/installing-and-configuring-graphite-and-statsd-on-an-ubuntu-12-04-vps

Workflow Overview

Workflow objects

Access Control

Property Search Paths

Parameters

Sending Notifications

Workflow logs

A Workflow Tutorial

Use a Workflow to design and manage processes at a higher level than individual jobs. Workflows allow you to combine procedures into processes to create build-test-deploy lifecycles (for example). A workflow contains states and transitions you define to provide complete control over your workflow process. The ElectricFlow Workflow feature allows you to define an unlimited range of large or small lifecycle combinations to meet your needs.

Some benefits of defining a workflow include

- Manual intervention workflows make it easy to introduce manual decision points into your process.
 With workflows, you can enable specific users or groups to choose how the workflow should proceed.
 Using an email notification, you can notify one or more users about the workflow status. If the workflow is waiting for someone to manually intervene, an email (including a screen link) can provide immediate attention and access to choices for how the workflow should proceed.
- **Aggregation** workflows allow you to collect information from a set of related jobs into a single location. A workflow can act as a central place to monitor the full lifecycle of your build, test, and deploy process.

- **Looping** workflows provide the native ability to run a job more than once. You have full control over the loop condition—you can re-run the job any number of times, only re-run the job if it did not succeed, or make this decision based on a custom job property.
- **Branching** workflows allow you to run multiple jobs in parallel and wait for them to complete before moving to the next state in the workflow. If you combine branching with looping, you can re-run these jobs based on a condition you defined (for example, only re-run jobs that failed).
- **Multipathed** design a workflow to take multiple paths automatically, depending on the outcome of each state's action, and the pre- and post-triggers you defined for each state.

Basic workflow concepts

Workflows contain states—one or more states can be designated as "starting" states to provide multiple entry points into the workflow. When a workflow is launched, a starting state is specified. You may prefer to design a complex or multi-purpose workflow so you can start the workflow at a different state to complete only a portion of the workflow for a particular outcome. For example, you might choose to run only the test suite workflow section.

As each state in a workflow becomes active, it performs an *action*. The state's action can be to create a job from a procedure or to start a workflow from a workflow definition. When the job or called workflow completes, "On Completion" transitions are evaluated to see if the workflow should change to a different state, possibly based on the outcome of the action.

You can create as many state definitions as you need. For example, you might have a small workflow definition with only 5 states, calling 5 procedures, or you might create 50-100 or more states to process a like-number of procedures from multiple projects.

A state can send email notifications before or after an action to notify interested users about workflow progress or to request manual user intervention to move the workflow to the next state.

Transitions are used to move workflow progress from one state to another state. Four types of transitions are available to move a workflow to the next state:

On Enter – evaluates before sending email notifiers or starting the action

On Start – evaluates immediately after starting the action. These transitions are ignored if no action is specified for the source state.

On Completion – evaluates when the action completes. These transitions are ignored if no action is specified for the source state.

Note: On Completion transitions are taken only if the state is still active when the action completes. These transitions are ignored if the workflow has transitioned to another state.

Manual – evaluates when a user selects the transition in the UI and specifies parameters. The same action can occur using ectool or the Perl API by calling transitionWorkflow. Only users who have "execute" permission on the transition are allowed to use the manual transition.

Similar to other objects in the ElectricFlow system, workflow objects can contain access control privileges, properties, and parameters.

Workflow objects

Workflow objects are split into two types: **Definition** objects and **Instance** objects. Definition objects provide the template for a running workflow instance. You create a new workflow by defining a Workflow Definition along with its State Definition and Transition Definition objects.

When you run the workflow definition, the system creates a new Workflow object with an equivalent set of State and Transition objects that represent the run-time instances of the workflow definition.

Note: We omit the "Instance" qualifier for brevity in the API and the UI.

Workflow Definition

This is the top-level workflow object, which is a container for states, and transitions, and other information defining your workflow. In the same way a procedure defines the behavior of a job, the workflow definition defines the behavior of a workflow.

Create a project to get started:

- Select the Projects tab, then click the **Create Project** link. After creating a project, its name will appear in the table on the Projects page.
- Select your project name to go to its Project Details page, select the Workflow Definitions subtab, then select the **Create Workflow Definition** link.
- On the New Workflow Definition page, name your workflow definition. Access the **Help** link on that page if you need help creating the Workflow Name Template (similar to the Job Name Template). Click **Save**.
- Now, on the Workflow Definition Details page (Graph view), you are ready to create the first state for your workflow.

Click here to see the Help topic for the Workflow Definition Details page that contains an example of how this page might look as you add states and transitions to your workflow definition. Notice also that you can add properties to the workflow definition at this point.

State Definition

Each workflow can contain one or more states. Defining states for a workflow is analogous to defining steps for a procedure.

You might choose to name your states for their intended purpose (for example: waiting for user, building, testing, or test1, test2, and so on) just as you might name steps for your procedure. You can define multiple states as "startable" so you can choose different starting points when you begin using the workflow. By default, the first defined state is always "startable".

As you create states, they will be displayed in the graph on the Workflow Definition Details page.

Selecting a state in the Graph view opens the State Definition panel to define:

- the action you want this state to perform, along with any parameters required by the action,
- transitions for this state,
- · parameters for this state,

- · email notifiers for this state,
- and any properties you want to assign to this state.

Transition Definition

Each state can contain one or more transitions. The transition definition requires a name for the transition. This transition name will appear on the Workflow Definition Details graph and the Workflow Definition Details List view.

You can define one or more transitions for each state.

When defining a transition, you can:

- Add a text description for your reference.
- Select the "target" state,
- Select the type of transition trigger you want to use (On Enter, On Start, On Completion, or Manual)
 See the Workflow Definition Details Help topic for trigger type definitions and other information for defining a transition,
- Select (or enter) a Condition to specify when the transition is allowed to trigger,
- and assign parameters to a transition. These parameters are passed to the target state.

Workflow

While we generically refer to or think of a "workflow" as an automated process to perform multiple tasks, it is actually the end result of running a Workflow Definition. Running a Workflow Definition produces a Workflow, which is analogous to running a Procedure, which produces a Job.

After you start running Workflow Definitions, select the main Workflow tab to see all workflows that ran (or are still running) listed in a table, each linked to its own Workflow Details page. The Workflow Details page main view is a graphical representation of the workflow, and also provides links to stop the workflow process, run the workflow again, see the workflow log, and so on.

A Workflow contains a pointer to:

- its workflow definition,
- sets of states and their transitions.
- · the current and initial state,
- · parameters to the initial state,
- a "complete" value to specify whether or not the workflow is complete,
- and a log containing transitions taken and user events.

State

The run-time instance of a State Definition is a State. The state is part of a workflow and contains information about actual parameter values specified by the user when the workflow was run or from transitions taken to enter the state.

If the state has an *action*, then the state contains information about the most recently created job or workflow. Because a state may be entered multiple times, the state contains a copy of the state definition properties needed to launch the action again.

Note: After the state is created, it no longer depends on the state definition, so any changes made to the definition will not affect workflows already running. Changes do not take effect until the next time the workflow runs.

Transition

The run-time instance of a Transition Definition is a Transition. The transition is part of a state and contains a copy of information from the transition definition. Each transition object corresponds to a transition definition, but once defined, each transition is recognized by its unique name. A transition contains unexpanded actual parameters cloned from definition, expanded and passed to the target state on entry, the Javascript condition, and the trigger type instructing the transition when to occur.

Note: After the transition is created, it no longer depends on the transition definition, so any changes made to the definition will not affect workflows already running. Changes do not take effect until the next time the workflow runs.

A transition moves the workflow from its state to a target state.

Access Control

Setting access control privileges for those persons who can or cannot run a workflow or execute a manual transition may relate to security policies in your organization.

Specific workflow access control notes:

- All workflow, state, and transition ACLs are copied from the definition object to their corresponding instance when the workflow is run initially.
- To run a workflow, you must have execute permission on the state definition you are using to start the workflow.
- To take a manual transition, you need execute permission on that transition.
- To strictly control who can take a manual transition:
 Grant read and modify privileges to anyone who can view or edit the workflow, but only give execute permission to those users or groups who can take the manual transition.
- You can control who can start each of the startable states by setting ACLs on state definitions accordingly.

For more information about access control, including examples you might use, see the Access Control Help topic.

Property Search Paths

A workflow serves as a hub of information that can be shared between states, transitions, jobs, and other workflows.

Depending on the context, an implicit search path is used to reference a property by name. Using a simple property name like <code>\$[someCustomProperty]</code> "walks" down the following paths until it finds a property by that name. If the search fails to find a property, an error is produced.

The search path followed to resolve a simple property name depends on where the reference occurs:

Transition context:

- 1. Transition property sheet
- State property sheet includes parameter values passed to that state and any custom properties.
 Because a state can be entered multiple times, parameter values are available for the most recent entry only.
- Starting state property sheet used to get to parameter values for the initial entry into the workflow.
- 4. Workflow property sheet used to access information shared across the workflow.

State context:

- 1. State property sheet includes parameter values passed to that state and any custom properties. Because a state can be entered multiple times, parameter values are available only for the most recent entry.
- Starting state property sheet used to get to parameter values for the initial entry into the workflow.
- Workflow property sheet used to access information shared across the workflow.

For information on context-relative property path shortcuts that allow passing information from one object to another or accessing a property from different parts of the workflow, review the Properties Help topic. The Properties Help topic also includes all intrinsic workflow-related properties.

Parameters

States may use formal parameters to determine what information to provide when the workflow enters that state. Parameter values can come from the initial run workflow call or from a subsequent transition into the state.

An automatic transition must specify all values required by the target state. A manual transition may omit some values to be passed to its target state. If a manual transition does not specify all values needed by the target state, then any missing values are collected from the user when the transition is taken.

Parameters to the starting state for the workflow are collected when the workflow is launched. These parameters are accessible throughout the workflow and act like global workflow parameters. See the Property Search Paths section above.

It is important to distinguish between formal parameters on a state and formal parameters on the action the state calls (procedure or workflow starting state). The procedure (or workflow) may have its own parameters, so the state is responsible for entering required actual parameter values. In some cases, values may be static values, in other cases the values may come from the state's parameters or those of the starting state (for example, global workflow parameters)

Parameters on the calling state do not need to be the same as the parameters on the action:

• A state parameter value can be passed as the value for an action's parameter by specifying a property reference in the value field on the Action tab for the state definition.

- All parameters on workflow current and starting states are accessible with a simple property reference (for example, '\$[param1]')
- Parameters on a workflow starting state act like global workflow parameters because they are available everywhere within the workflow.

Sending Notifications

A state can be defined to send email notifications when it is entered, when it starts its action, and/or when the action completes. Notifiers are created by visiting the Notifications tab on the State Definition panel. See the Email Notifier Help topic for details on defining email notifiers. Sample notifier templates are provided also.

Workflow Logs

Each workflow keeps a log of events that relate to the workflow. The log contains information about the evaluation of transition conditions, transitions taken, actions started, and other information that might prove useful when debugging a workflow. Access the log by clicking the **View Log** link on the Workflow Details page. **For more information**, click here to go to the Workflow Log Help topic, which contains a screen example of the Workflow Log page.

Visualizing a Workflow

Two Workflow web pages, Workflow Definition Details and Workflow Details, open to the Graph view.

- Workflow Definition Details page As you build your workflow, you can see a visual representation of state and transition definitions you define.
- Workflow Details page displays a graph, updated in real-time, for your running or completed workflow.

Your workflow Graph (on either page) contains context-sensitive links to various options. Options are different depending on whether or not you are building or running your workflow. Workflow Definition Details and Workflow Details Help topics include information about the context-sensitive link options provided within the workflow graph.

The next section, "Building a Workflow - a Tutorial," begins with a "sketched" representation of the workflow the tutorial will create. At the end of the tutorial, see the actual workflow graph created by the ElectricFlow workflow process.

Building a Workflow - a Tutorial

This overview tutorial is an example of how to create a "build-test-approve" workflow, building each state and transition we want the workflow to include, and building the workflow in the order we need to use to achieve our purpose. The diagram below outlines our workflow process.

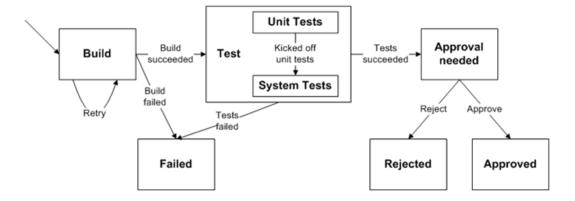
You can read through the tutorial or actually recreate this example workflow on your machine. If you want to use this tutorial as a real, working example, you must do some initial setup first:

- Create a project named "Sample"
- · Create a procedure name "BuildProduct"
- On the "BuildProduct" procedure, add a parameter named "branch"

The following diagram is a visual illustration of the workflow we are going to build.

Reading the diagram:

- the starting state is Build, which could transition to the Test state or the Failed state
- two test states Unit Tests and System Tests, which could transition to a Failed state or an Approval needed state
- the Approval needed state could transition to one of two states Rejected or Approved
- all transitions are represented by the lines between the states, which include "arrows" to show the
 possible workflow directions (depending on the outcome of each state's action



Best Practice tip

Designing your workflow on paper or whiteboard, creating your own state and transition diagram, is a great way to organize your thoughts and map your workflow process.

A summary of the component sections to build our workflow:

- Calling a job from within a workflow
- Collecting a parameter when a workflow is launched and passing its value to a job
- Retrying a state
- · Automatically transitioning to different states based on the outcome of a job
- · Invoking another subworkflow
- Running jobs in parallel
- · Automatically transitioning to different states based on the outcome of jobs in a workflow
- · Waiting for manual intervention
- Restricting who can take a manual transition
- · Sending email notifiers
- · Adding a global parameter to use later in the workflow

- Setting the name of your workflow
- Workflow list

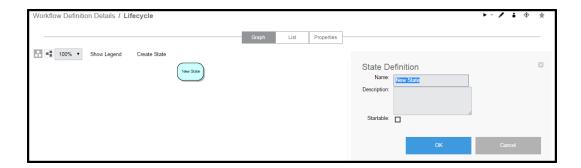
To begin calling a job from within a workflow

Or, running a job by calling its procedure from a workflow state.

- Select the Projects tab, select the Sample project to go to the Project Details page, then select the Workflow Definitions subtab.
- Click the Create Workflow Definition link to go to the New Workflow Definition page..
- Create a new Workflow Definition named Lifecycle.



Because the workflow definition, *Lifecycle*, is a new definition and does not have any state or transition definitions, the State Definition panel opens, ready for you to create a new state definition for your workflow.



- Change the system-generated state Name to Build and click **OK**.
- The State Definition panel is displayed.



- Specify the Action for this state, which is Procedure.

 In this example, the procedure *BuildProduct* already exists with a single parameter named *branch*.
- Set the value of this parameter to "3.0".
- The Build state now has an Action associated with it.
- Click OK.
 Note the shape of this state in the graph. The state shape will be different depending on the state Action selected.
- Click Run to launch the workflow.
- Running the workflow shows a Job was launched by the *Build* state. The workflow job completed successfully, indicated by the "check mark" on the state.
- You can get more details about the job by clicking on the *Build* state, which opens the State Details panel, and clicking on the Job link displays the Job Details page.
- Alternately, clicking on the Show History button opens the History panel to show what ran and the workflow outcome.

Collecting a parameter when a workflow is launched and passing its value to a job

Or, "prompting" for a parameter value when a workflow is launched and passing its value to a procedure.

- Return to the Workflow Definition page for the **Lifecycle** workflow.
- Click the Build state (from the graph) to open the State Definition panel, then select the Parameters tab.
- Click the Create Parameter link to display the New Parameter page.
- Create a parameter called "branch".
- Add a description and make sure the Required? box is "checked" and click OK.
- On the Workflow Definition Details page again, change the value passed to the job from 3.0 to \$[branch].

Now, when the workflow is launched, you will be prompted for the value of *branch* and you can specify a different value if necessary.

For example, a "3.5" value is now specified in the screen below.

The value is passed to the job when the workflow is launched.

Retrying a state

If a State fails (technically, the job started by the workflow state), you may choose to retry it a few times before giving up.

- Create a new transition from the Build state back to itself.
- Right-click on the Build state and select the Create Transition link.
- Create a new On Completion transition named "retry" from the *Build* state that returns to the same state when the job fails.
- Add a limit of 3 retries by using the following JavaScript condition:

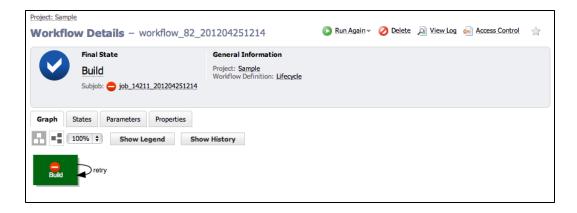
```
$[/javascript mySubjob.outcome == "error" && myState.getProperty('/increment
numberOfRetries')< 3; ]</pre>
```



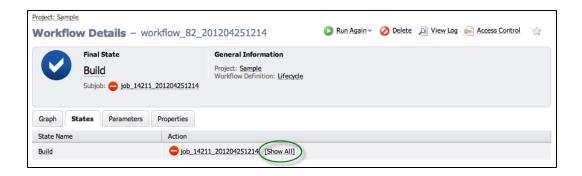
• The new transition is displayed on the Workflow Definition Details page.



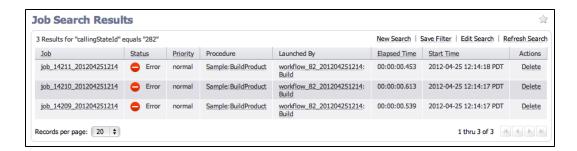
• When we run the workflow, we see only the last job that was run.



Clicking the States tab and then the Show All link displays all jobs that were invoked by the Build state.



On this screen, you can see the build was attempted 3 times because that was the limit we specified in the JavaScript condition for the "retry" transition.

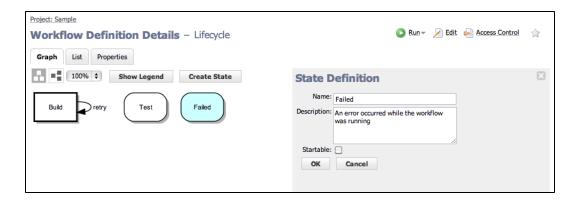


Also, you can see the number of times a state was visited and which transitions were taken by clicking the **Show History** button.

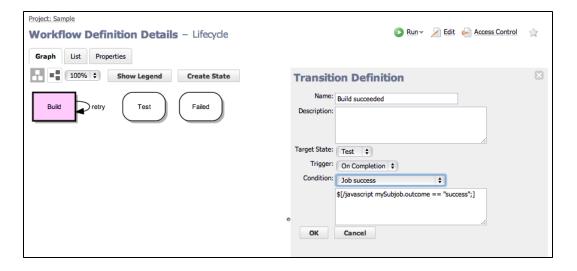
Automatically transitioning to different states based on the outcome of a job

If the *Build* job fails, we might want to end the workflow. If the job succeeds, we might want to continue to the next stages of the workflow.

- Add new states to transition to, depending on the outcome of the job:
 - Create the Test state where the workflow will transition to when the build state succeeds.
 - o Create a Failed state.



- Create an On Completion transition from Build to Test when the job succeeds.
- Notice that the Condition dropdown menu provides a template condition called *Job success* that can be used here.

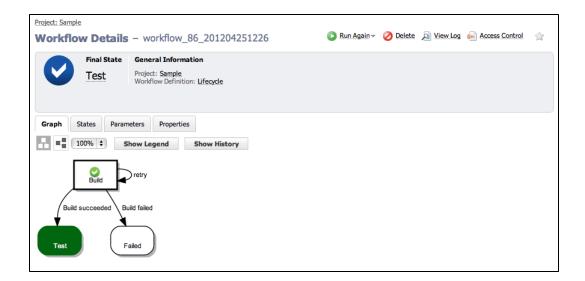


- Create an On Completion transition from Build to Failed when the job fails.
- Notice that the Condition field remains blank.
 Because transitions are evaluated in order, if the workflow gets to this transition, we will know that the "success" case will be false.

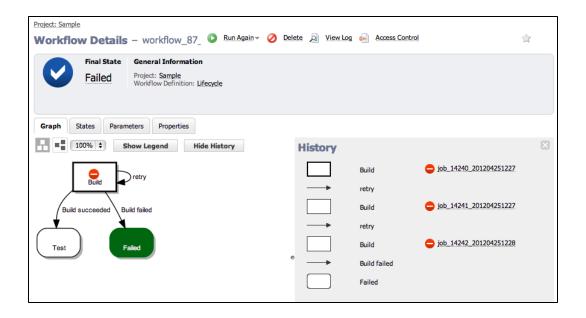
Leaving the Condition field "Always" allows us to handle all "non-successful" outcomes (error or warning).



• When you run the workflow, if the build succeeds the workflow will end up in the *Test* state.



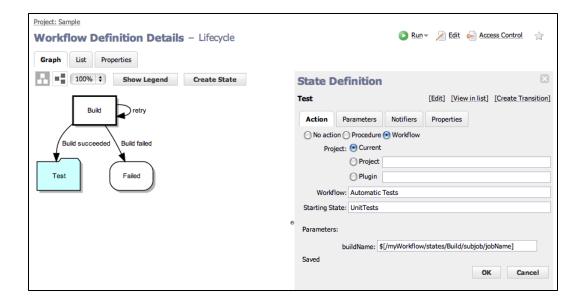
• The workflow will end up in the *Failed* state if the build fails after trying three times. (In this example, we show the Workflow History panel open.)



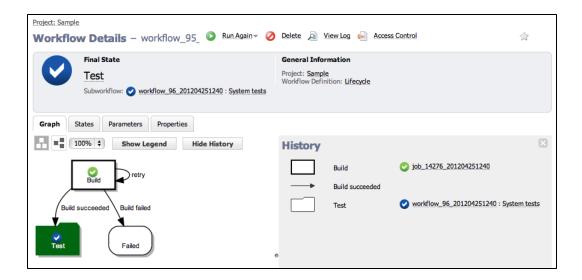
Invoking another workflow

Just as a state can invoke a job as its "Action", it can invoke another workflow, which then becomes its subworkflow.

- From the Workflow Definition Details page for *Lifecycle*, click on the *Test* state in the graph.
- Specify the "Action" for the Test state, which is a Workflow.
- In this example, the workflow *AutomaticTests* already exists and so does a state called *UnitTests*, which has a single required parameter called *buildName*.
- We can get the name of the build from the *Build* state by using the property path \$[/myWorkflow/states/Build/subjob/jobName].



- The Test state now has an Action associated with it.
- Because the Test state Action is Workflow, the object shape in the graph has changed.

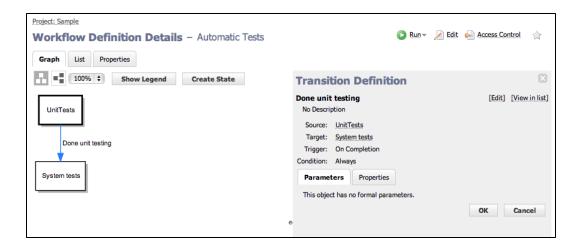


- Running the workflow shows another workflow was launched by the Test state.
- Click on the workflow link in the History panel to see more details about that workflow.

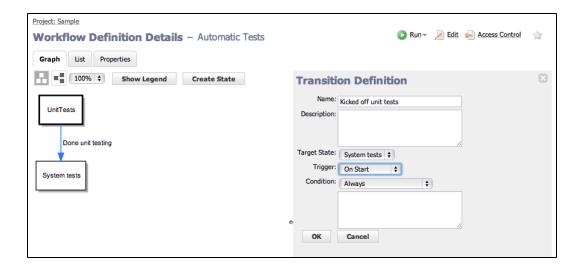
Running jobs in parallel

Run multiple jobs at the same time.

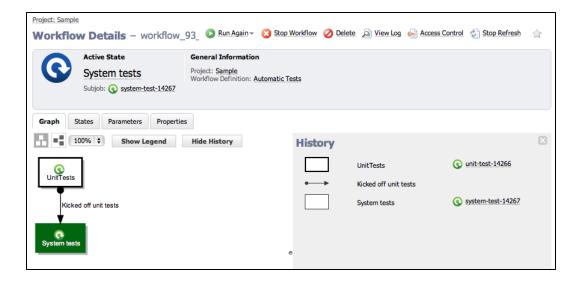
• Look at the *AutomaticTests* workflow that was invoked (in the previous example) by the *Test* state in our main *Lifecycle* workflow. The first state in this workflow runs unit tests, and when the tests complete, the state automatically transitions to the second state that runs system tests.



- Now we will change the behavior of the transition to move from the *Unit tests* state to the *System tests* state as soon as the first job is launched.
- Edit the existing transition by renaming it to *Kicked off unit tests* and changing it to an On Start trigger type.



- Now, when the AutomaticTests workflow is launched, both jobs will run in parallel.
- The workflow completes automatically when both jobs complete.



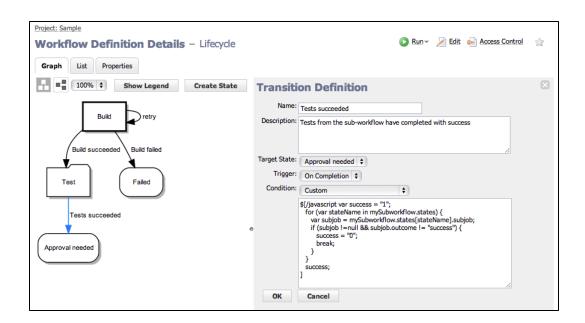
Automatically transitioning to different states based on the outcome of jobs in another workflow

If the jobs in the *Test* workflow fail, we will want the workflow to end. If the jobs succeed, we will want to continue to the next stages of a calling workflow.

- · Click the Create State link.
- Create the Approval needed state that the workflow will transition to when tests succeed.
- Create a transition named "Tests succeeded" from Test to Approval needed when all jobs in the workflow succeed.
 - (Right-click on the *Test* state to create the transition.)
- The following JavaScript condition checks the outcome of all subjobs for a subworkflow and evaluates to "true"

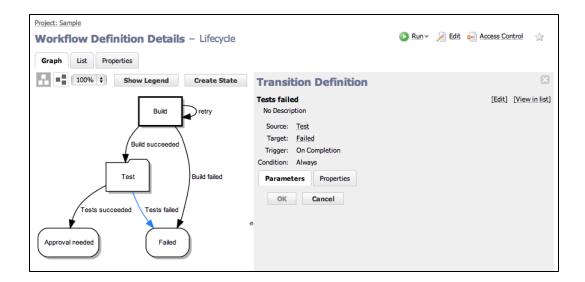
only if they all succeeded:

```
$[/javascript
  var success = "1";
  for (var stateName in mySubworkflow.states) {
    var subjob = mySubworkflow.states[stateName].subjob;
    if (subjob !=null && subjob.outcome != "success") {
        success = "0";
        break;
    }
  }
  success;
}
```

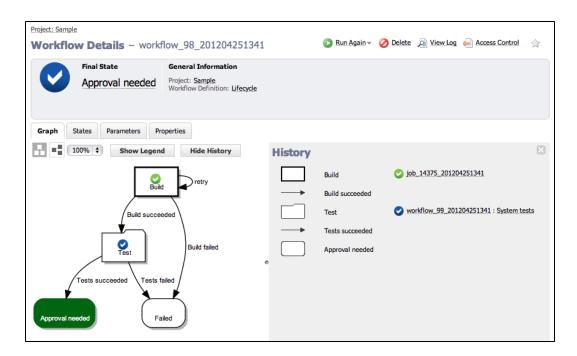


Note: We created the *Approval needed* state first, then we created the *Tests succeeded* transition to connect the two states.

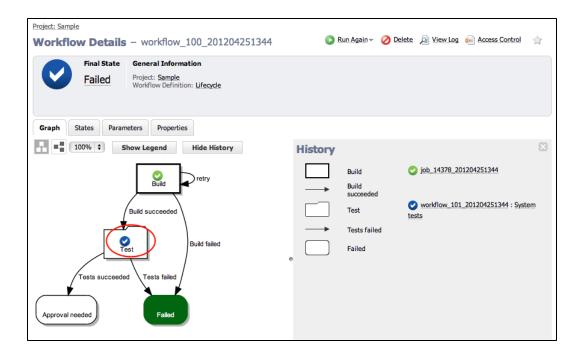
- Create a transition from the Test to the Failed state when either of the jobs in the workflow fails.
- · Notice the Condition field remains blank.
- Because transitions are evaluated in order, if the workflow gets to this transition we know the success case was "false".



• When you run the workflow, if the test jobs succeed, the workflow will end up in the *Approval needed* state.

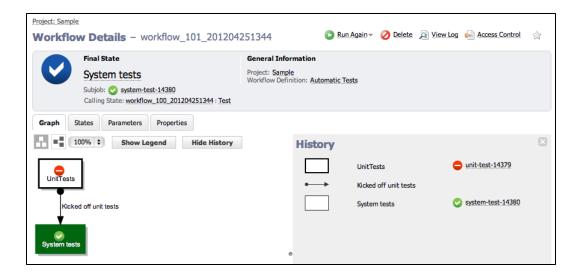


• In the following example, one of the tests in the *Test* workflow encountered an error, so the workflow will end up in the *Failed* state if either of the test jobs fail.



Note: The "blue checkmark" indicates the workflow (or subworkflow) has completed, which means it is no longer running. This icon does NOT indicate the workflow outcome (completed successfully, or with warnings or errors).

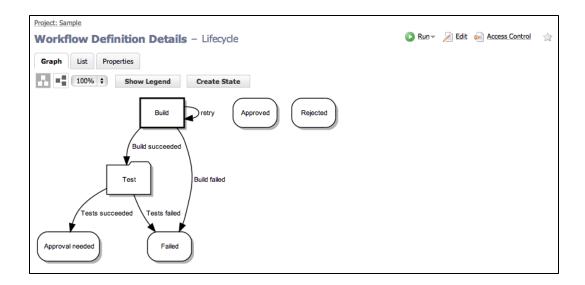
To see exactly what happened in the workflow, you can "drill-down" by clicking on the workflow link shown in the History panel.



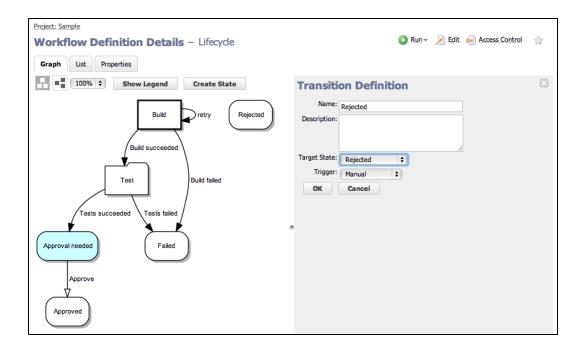
Waiting for manual intervention

Allow users to choose how to progress the workflow.

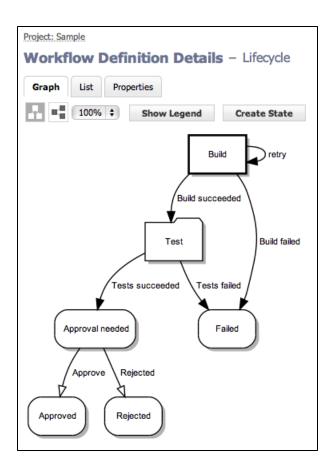
- From the Workflow Definition Details page for *Lifecyle*:
 - Create an Approved state and a Rejected state for the workflow to transition to when a user approves or rejects the build.



• For the *Approval needed* state, we create a manual transition called *Approve* that the user can select to progress the workflow to the *Approved* state.

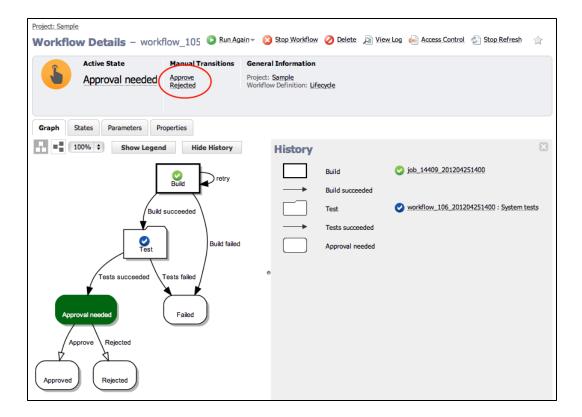


- Next, we create a manual transition called *Reject* that the user can select to progress the workflow to the *Rejected* state.
- At this point, our workflow now looks like the following example.

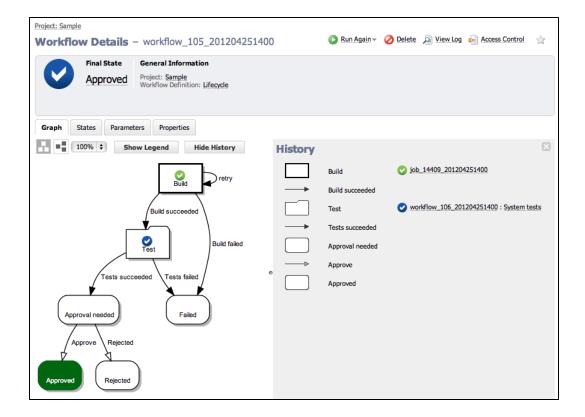


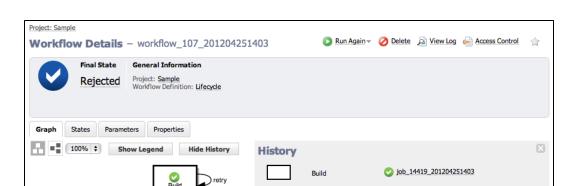
• When you run the workflow and the build and test phases succeed, the new manual transitions are made available to users on the Workflow Details page.

Note: If a user does not have execute permission for manual transitions, the Manual Transitions section will not be included in the summary section at the top of the Workflow Details page.



• If you choose the Approve transition, the workflow will end up in the Approved state.





Tests succeeded

Rejected

workflow_108_201204251403 : System tests

• If you choose the Reject transition, the workflow will end up in the Rejected state.

Build failed

Restricting who can take a manual transition

Approval needed

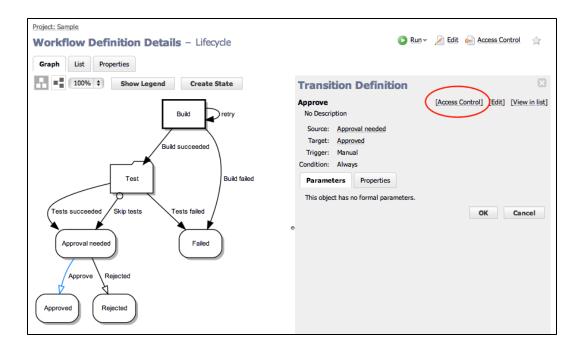
Approved

Only allow certain users to take manual transitions.

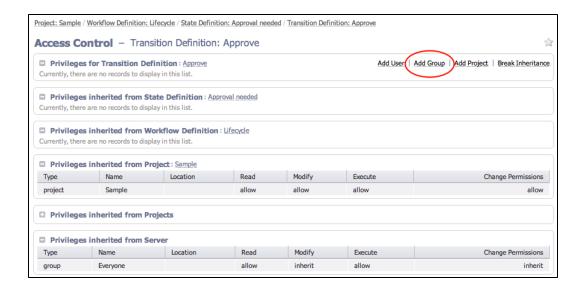
• From the Workflow Definition Details page for Lifecycle:

Failed

- When the workflow is created, access control entries from the Transition Definition are copied to the corresponding transition.
 - These are the access control entries for the Approve transition before we edit them (see below).
- o To add access control to a transition, select the transition or transition name.
- When the Transition Definition panel opens, click the **Access Control** link,



• When the Access Control page is displayed, click the **Add Group** link to add privileges for a group that can take this transition.



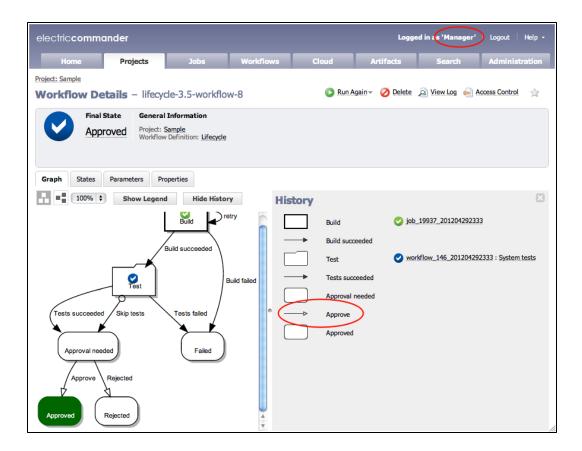
Note: If you already have groups created, you can use them here or create two groups now: *Lifecycle Approvers* and *Lifecycle Modifiers*.

- Allow read and execute access for a group named *Lifecycle Approvers*. These users can see and take a manual transition in a workflow.
- Allow read, modify, and change permission access for a group named *Lifecycle Modifiers*. These users can edit the Transition Definition, but will not have access to take a manual transition in a workflow.

- "Break the inheritance" so only specified groups have access to this transition.
- Duplicate this access control configuration for the Reject transition.



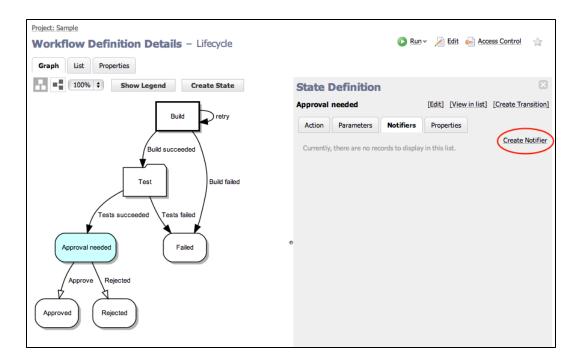
- When a workflow is launched and reaches the *Approval needed* state, a user who is not a member of the *Lifecycle Approvers* group will not be allowed to execute any manual transitions.
- A user who is a member of the Lifecycle Approvers group can execute manual transitions.
- The Workflow Details page updates in real-time, which means you will see changes at the top of the
 page pertaining to the current state in the running workflow, and current changes in the graph also. For
 example, if the previous state had a manual transition, but the current state now processing does not
 have a manual transition, the Manual Transition section will not be displayed.



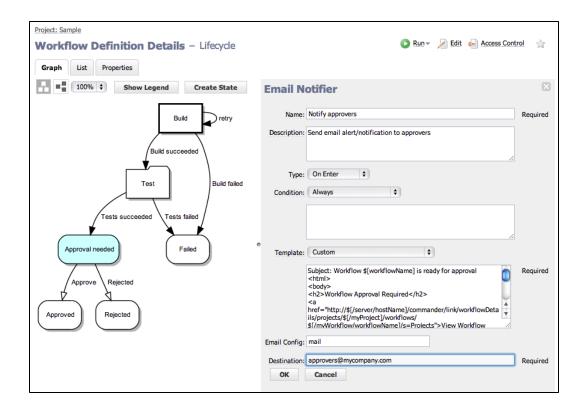
Sending email notifiers

Send email notifications so a user or group knows they need to take an action on a workflow.

- From the Notifications tab for the *Approval needed* state definition, click the Notifiers tab.
- Click the Create Notifier link.



- We want to create an On Enter email notifier.
- When a workflow enters the Approval needed state, this email notifier will be sent out.



- In this example, we want to send an email notification to the "approvers".
 The email will contain a link to the Workflow Details page and links to all available transitions.
- Use this formatting template for the email notifier: http://docs.electric-cloud.com/commander_doc/5_ 1/PDF/WorkflowTutorialScript.txt.

It includes a link to the ElectricFlow workflow, embeds a block of JavaScript that cycles through all the manual transitions from the active state, and has a link to go directly to those transitions. To ensure that the script works properly, use the script as is and make sure there are no white spaces before <!DOCTYPE...>, <html>, and </html>.

 This script is a sample of the type of script you can insert into the Formatting Template field illustrated in the following screen example. • When a workflow is run and it reaches the *Approval needed* state, an email similar to the following example can be sent to specified destinations.

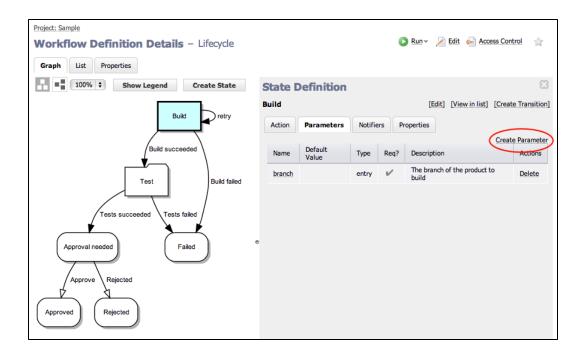
Only users who have execute permission on these transitions will be able to approve them by clicking the link.

Workflow Approval Required View Workflow workflow_7_20141212073210

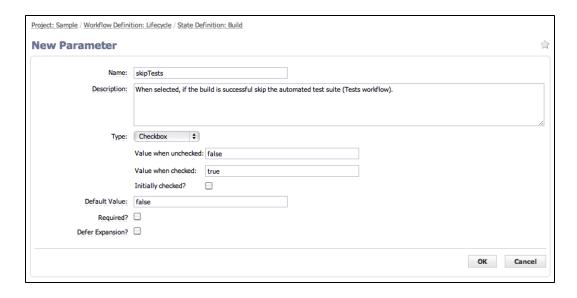
Adding a global parameter to use later in the workflow

Prompt for a parameter when the workflow is launched, and use its value later in the workflow.

- Currently, there is one parameter for the starting state Build.
- Click on the Build state in the graph and select the Parameters tab.
- Click the Create Parameter link to create a second parameter.

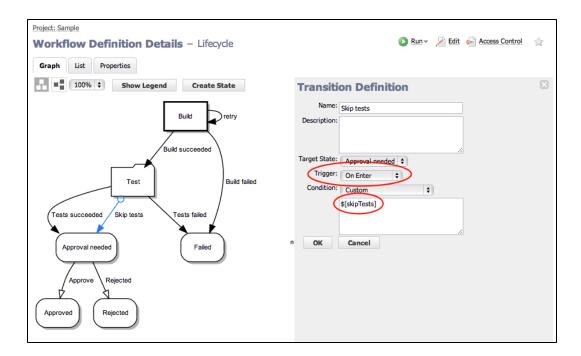


• Add a checkbox-type parameter named *skipTests* which, if selected, should skip the test suite after the build is completed.



• Create an On Enter transition from *Test* to *Approval needed* if the *skipTests* parameter evaluates to "true"

Because this is an On Enter transition, it will be taken before the automated tests workflow is started.

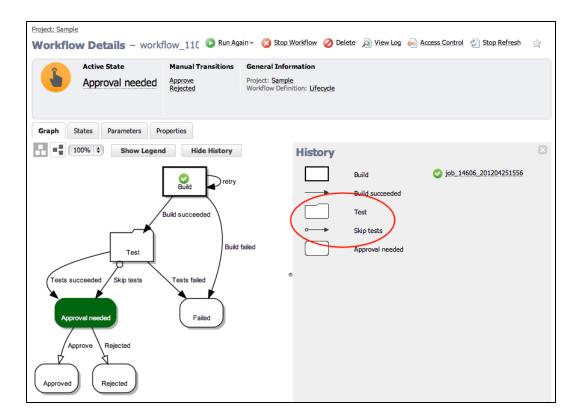


• When the workflow is launched, you will see the skipTests checkbox.



• If the parameter is checked, and after the *Build* state is completed, the workflow will immediately transition from *Test* to *Approval needed*, without starting the subworkflow.

Note: Looking at the workflow History, we see the *Test* state was entered, but the subworkflow (and subjobs) was not executed.

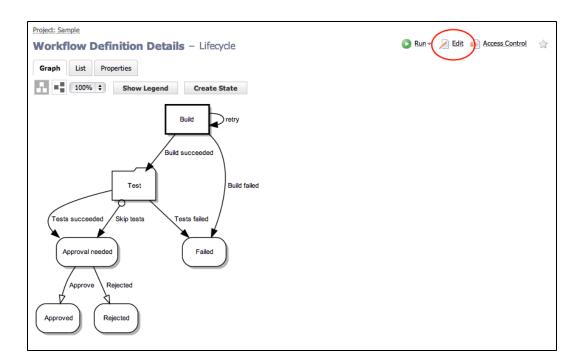


Setting the name of your workflow

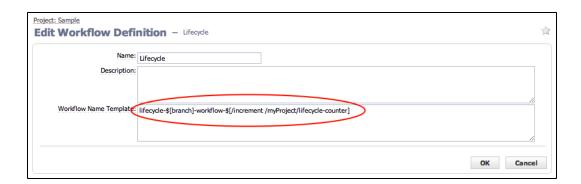
Choose appropriate names for workflows so they are easy to identify.

- If the Workflow Name Template was not set, workflows are created with a default name (workflow <jobId> <timestamp> or workflow_239_201012091208, for example).

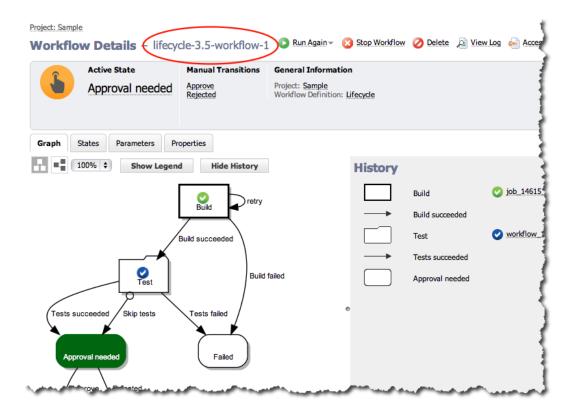
 However, because you may have many different workflow definitions, you might want to create a distinguishing name for workflows launched from each definition.
- From the Workflow Definition page for *Lifecycle*, click the **Edit** link at the top of the page.



- Setting the Workflow Name Template (as shown in the sample below) achieves this goal.
- lifecycle-\$[branch]-workflow-\$[/increment /myProject/lifecycle-counter]
- \$[branch] refers to the value of the branch parameter as specified when the workflow is launched.
- \$[/increment /myProject/lifecycle-counter] creates a property called lifecycle-counter on the project and increments the property every time a new workflow is launched so the names are unique.



• Now, when the workflow is launched, its name is much more descriptive than the default workflow name.

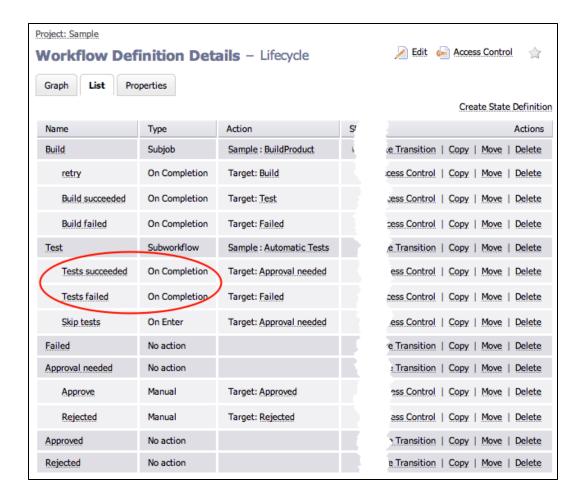


Workflow List

Most of the time, you will want to interact with your workflow from the default Graph view. However, if you want to see your workflow in a list or table view, click the List button. See the Workflow Definition Details Help topic for detailed information about this page.



The transition order is important because transitions are evaluated in order and the first one with a condition that evaluates to 'true" is taken. In the next screen example, the *Test* state has two On Completion transitions. When the state completes, it evaluates the *Tests succeeded* transition first. If a third transition was added for the *Test* state, you could use this List view and Move the transition to the correct order for evaluation.



Another workflow example

An advanced sample Workflow Definition is pre-defined for you in the EC-Examples project. This is a "working" example, also modeling the workflow process, with which you can interact. To navigate to this example: Select the Projects tab > EC-Examples > Workflow Definitions subtab > then select Build Workflow With Approvals.

Workspaces and Disk Space Management

Defining workspaces

Using workspaces

Workspace directory names

Working directories

Workspace accessibility

Local workspaces

Access control

Impersonation and workspaces

ElectricFlow managed files

Disk space management

ecremotefilecopy

Overview

ElectricFlow provides each job step with an area on the disk it can use for "working files" and results. This disk area is called a **job workspace**. (The job workspace defaults to a directory under the workspace directory.)

- A job step can create whatever files it needs within its workspace, and ElectricFlow automatically places
 files in the workspace, such as step logs. Normally, a single workspace is shared by all steps in a job,
 but it is possible for different steps within a job to use different workspaces.
- The location of the job step workspace is displayed on the Job Details page for the job under "Details" for the step.
- Workspaces are grouped together, with workspaces for different jobs allocated as children of the same workspace root. When you define procedures, you can specify which workspace roots to use, then ElectricFlow creates a child directory within that workspace root for the job's steps to use.
- The term "workspace" is used at different times to refer to either a job workspace or a workspace root.
 Depending on the context, it should be easy to discern which "workspace" is being discussed or referenced.

Defining Workspaces

ElectricFlow can handle any number of workspace roots, which are defined by selecting the Cloud > Workspaces tabs on the web interface. To create a workspace root, provide a workspace name and three file paths for jobs to access the workspace:

Workspace name

The workspace name must not have trailing spaces.

When two workspaces have similar names, such as "Server1" and "Server1" (with an extra space at the end), the database records them as two entries, one without trailing spaces and one with trailing spaces.

If you run a job against one of these workspaces, it may hang and not proceed because ElectricFlow is waiting for a resource, an available workspace, or a step to complete.

UNC path

A UNC path that Windows machines can use to access the workspace via the network, such as //server/share/x/y/z.

You must ensure this path is accessible on any resource where the workspace is used.

Drive path

A Windows path, using a drive letter that refers to the same directory as the UNC path. Before running a job step using the workspace, the ElectricFlow agent creates a drive mapping for this path to match the UNC path.

For example, if the UNC path is //server/share/x/y/z and the drive path is N:, the agent will map N: to //server/share/x/y/z. If the drive path is N:/y/z then the agent will map N: to //server/share/x.

The directories in the drive path after the drive letter (y/z in the preceding example) must match the last directories in the UNC path.

UNIX path

A path used on UNIX machines to access the workspace via the network, typically using an NFS mount. You must ensure appropriate mounts exist on all resources where the workspace is used.

If your environment contains Windows machines only, you can omit the UNIX path, and if your environment contains UNIX machines only, you can omit the UNC and Drive paths.

During the ElectricFlow installation, you had the option of defining a workspace. If you defined a workspace, it was named "default" and will be the default workspace for all jobs, as described below.

By default, remote workspaces are accessed with the agent user's credentials. For Windows resources, you can override that user's credential by attaching a credential to the workspace. Using a credential attached to a workspace is useful if the agent machine is not a member of the domain, but the workspace is located on a machine in the domain. An attached credential for a domain user allows the agent to manipulate the workspace as that domain user.

Using Workspaces

The simplest way to use workspaces is to define a single workspace named "default" that is readable and writable on all resources, and never specify any other workspace information. With this approach the default workspace will be used for all steps in all jobs.

However, you can also arrange for different jobs to use different workspaces, or even for different steps within a single job to use different workspaces.

You can specify a workspace name in any of the following ElectricFlow objects:

- If you specify a workspace in the definition of a procedure step, it applies to that step.
- If you specify a workspace in the definition of a procedure, it applies to all steps in the procedure.
- If you specify a workspace in the definition of a project, it applies to all steps in all procedures defined in that project.
- If you specify a workspace in the definition of a resource, it applies to all steps assigned to that resource.

If workspaces are defined in several of these places, the workspace for a particular step is chosen in priority order corresponding to the list above: a workspace specified in a step takes priority over all others, followed by a workspace in the procedure, project, and finally resource. If no workspace is specified in any of these locations, the workspace named "default" is used.

Workspace Directory Names

A job workspace directory name, within its workspace root, is derived from the name of the job.

For example, if a job named "job_63" uses a workspace whose root directory is //ec/workspaces, the job workspace directory will be //ec/workspaces/job_63. Changing the name of a job has no impact on workspace names: the initial name of the job is always used for workspace directory names.

A job is allocated a single job workspace only within a particular workspace root, which will be shared by all steps that specify that root.

For example, suppose a job has 5 steps: 3 steps specify workspace A and 2 steps specify workspace B. The 3 steps that specify A will all share a single directory under A, and the other 2 steps will share a single directory under B.

Working Directories

By default, each job step begins execution with its current working directory set to the top-level directory in the job's workspace. A job step can override this location by specifying an explicit working directory. If a job step specifies a relative path for its working directory, the path is relative to the top-level workspace directory.

Workspace Accessibility

The simplest way to set up ElectricFlow is to make every workspace root accessible on every resource. This accessibility mode provides the most flexibility and simplifies ElectricFlow system management.

Unfortunately, some environments cannot allow such universal access.

For example, some sites do not provide file sharing between UNIX and Windows machines. In large enterprise environments where ElectricFlow is shared among multiple groups, you may want to partition resources among the groups and limit access by each group to the other groups' resources. In the most extreme case, you may not have network file sharing at all. ElectricFlow can handle all of these cases.

If a workspace root is not accessible on a particular resource, you should not use that workspace in any step that runs on the resource. If you do, the resource will not be able to create the job workspace directory and the step will stall until the issue is resolved. A workspace root (on a resource) must be both readable and writable to be used for job steps.

Local Workspaces (Disconnected Workspaces)

Normally, workspaces are shared, which means if the same workspace is used on different resources, they see the same set of files. However, it is possible to define a *local* workspace (also referred to as a disconnected workspace), which means it will refer to a different local directory on each resource where it is used.

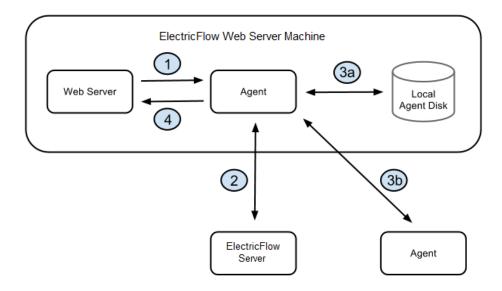
To define a local workspace on Windows, specify a local path such as C:/Windows/Temp/ecworkspace for the drive path.

For UNIX, specify a local UNIX path such as /usr/tmp/ecworkspace. Local workspaces are advantageous because they do not require remote network access and access may be faster in some cases.

However, you cannot share workspace information among job steps running on different resources.

Note: Web servers "ask" the relevant agent for access to log files in workspaces. The agent streams the file back to the web server, which then streams the log file to your browser. However, this mechanism works only for Windows and Linux 4.2 (or later) agents. Solaris, MacOS, and HPUX agents must transfer log files to a

workspace proactively so a Linux or Windows agent can access them for the web server—this task can be made easier, using the <code>ecremotefilecopy</code> tool. See the ecremotefilecopy section in ElectricFlow Installed Tools on page 462 for details.



The diagram illustrates the architecture when the web server requests a log file and you have defined a local workspace.

- 1. The web server asks the local agent to retrieve the file from a workspace.
- 2. The agent (local to the web server) asks the ElectricFlow server how to find the workspace, and the ElectricFlow server replies with a route to an agent and validates the session ID.
- 3. The route to the agent that has the file could be: (a) the local agent itself or (b) a different agent that ran the step that produced the file.
 - a. The agent (local to the web server) gets the file from its local disk.

or

- b. The agent gets the file from some other agent.
- 4. The agent (local to the web server) sends the file to the web server.

Access control

You can restrict access to workspaces using the ElectricFlow access control mechanism. Before a step can use a workspace, it must have the "execute" privilege on the workspace, where "it" means the ElectricFlow user ID under which the step executes (the project principal). You can set permissions on a workspace to enable or disable access as you choose.

Impersonation and workspaces

Before a job step can use a particular workspace, three access limitations must be satisfied:

The job step must have execute access to the workspace object.

If not, ElectricFlow aborts the step with an access control violation. See the Access Control Help topic for details.

The workspace must be accessible on the resource where the step runs.

If not, the ElectricFlow agent is unable to create the step's log file and the step's execution will fail. See Workspace Accessibility (above) for more information.

The Windows or UNIX agent account and/or job step must have read/write access to workspace files.

If not, the step's execution will fail in one of several ways. This issue is the topic of this section.

Potentially, two system accounts can impact each job step execution:

The first account is the one used by the ElectricFlow agent—you selected this account when ElectricFlow was installed. The agent account must have the ability to write the workspace root directory (to create the job workspace directory), and it must have the ability to write the job workspace directory to create log files.

The second account is the one under which the step executes. If the step does not use user impersonation (in other words, no credential has been specified for it), the step runs under the same account as the agent so there are no additional issues. However, if the step is running with credentials, the credentialed account must also have read/write access to the workspace area.

Because of these requirements, you may end up with a configuration where any job step running in a workspace can read or write any file under that workspace root, including files from other jobs that share the same workspace root. In many environments this solution works, but in some situations you may want to prevent one job from accessing files from a different job.

To create job privacy, run all of the job's steps with credentials for a separate account. In addition, run a job step at the beginning of the job that creates a subdirectory within the job workspace and change the protection on that directory to permit access to the credentialed account only. Next, place all private files for the build in the protected subdirectory. These files will not be accessible to other jobs or the agent. The top-level directory in the job workspace still needs to be accessible to the agent so it can read and write log files, but everything in the lower-level directory will be private.

ElectricFlow managed files

ElectricFlow automatically places certain files in the top-level job step workspace directory:

Step logs

When a step's command executes, its standard output is redirected to a log file unique to that step. The log filename is derived from the step name and its unique identifier within ElectricFlow.

For example, a step named "step1" will have a log file with a name like "step1.2770.log", where 2770 is a unique identifier assigned by ElectricFlow. A unique identifier is needed in situations where the same step name occurs multiple times in a job such as multiple invocations of a single nested procedure.

Postprocessor logs

If a step specifies a postprocessor, standard postprocessor output is redirected to a file in the workspace. The postprocessor output file will have a name similar to the step log. For example, "step1.2770-postp.log".

Diagnostic extracts

If a postprocessor extracts diagnostic information from a step's log file, it usually places that information in an XML file in the top-level job workspace directory, and then it sets a property that contains the filename. The ElectricFlow *postp* postprocessor uses filenames like diag-2770.xml, where 2770 is the unique identifier for the step. Other postprocessors may use different filenames.

Disk space management

The current ElectricFlow version does not provide a facility for automatically deleting old workspaces. You are responsible for deleting obsolete job workspaces yourself. If you delete a job, ElectricFlow does not automatically delete its workspace. (A future ElectricFlow version may provide automated facilities for identifying and deleting obsolete workspaces—until then, you may wish to create a procedure that runs on a regular schedule and deletes old job workspaces.)

Tutorials

The ElectricFlow Help topics in this directory help you learn more about ElectricFlow and specific ElectricFlow tasks.

You can directly access these Help topics from the Help system Table of Contents.

You can also access these Help topics while using the automation platform. Click **Help** at the top-right corner on any web page in the automation platform.

Adding a Link to a Job

ElectricFlow provides an easy method for adding links to the Job Details page to view files and reports for any job where you need to access additional information quickly. Adding direct links eliminates logging into where these files may be stored.

Note: The Workflow Details page and several other ElectricFlow web pages also allow adding links .

This tutorial shows how to add links to a job and have them immediately available on the Job Details page.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

• The Procedure Details page for the "Adding a link to a job" procedure contains two steps: Create File and Create Job Link

The Action column displays a code fragment of the script used to create the readme.txt file and a
reference for the Create Job Link step. For this example, we want to link to a "readme" file.
 (You can click on the Create File step name to go to the Edit Step page to see the full ec-perl script
shown partially in the Action column.)

Click **Run** to run this sample procedure and see the resulting job status on the Job Details page, where the new link to the readme.txt file is displayed under the Links heading in the summary section at the top of the page.

Implementation

Use the following instructions if you would like to practice creating a link on the Job Details page now.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow. These instructions only coach you through the process you would use to create links on your Job Details pages.

- In the automation platform UI, click Projects > EC-Tutorials-<version> > Adding a link to a job.
 The Procedure Details page opens.
- In the Procedure Steps section, click Subprocedure to create a new step based on a subprocedure.
 The Choose Subprocedure Step dialog box opens.
- 3. To select a project:
 - a. Select **Project**. In the text box, start typing **EC-Utilities**.
 As you type, possible matches appear in a drop-down list.
 - b. When you see EC-Utilities, select it.
- 4. To select a procedure:
 - a. Select Project. In the text box, start typing Create Job Link.
 As you type, possible matches appear in a drop-down list.
 - b. When you see Create Job Link, select it.
- 5. Click OK.

The New Step page opens.

- 6. In the General section, Google search in the Name field.
- 7. In the Parameters section:
 - a. In the Link Label field, enter Google Search.(This is any name you choose for the name of your link.)
 - In the Link Location field, enter http://www.google.com
 (This is the URL to your link location, or it can be a relative path to where you want to create a link.)
 - c. Click **OK** to save your entries and return to the Procedure Details page.
- 8. On the Procedure Details page, click Run.
- 9. On the Job Details page, as soon as the links are processed you will see the new "Google Search" link in addition to the "readme.txt" link.

Clicking on the new Google Search link will take to you to the Google website.

Related Information

Job Details - See this Help topic for more detailed information about creating links and complete information on other aspects for using the Job Details page.

Calling a Subprocedure

The ability to create *reusable components* is a major ElectricFlow strength. Procedures in any project or plugin can be called as a subprocedure in a step for another procedure. Creating a subprocedure is one way of reusing existing functionality already implemented in ElectricFlow.

This tutorial shows how to use the step configuration process to call an existing procedure to use as a subprocedure.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

- The Procedure Details page for the Calling a subprocedure procedure contains a step named Execute a
 procedure from this project.
 (Clicking on this step takes you to the Edit Step page.)
- The Action column displays a plugin (or project) name, EC-Tutorials, and a procedure name, Working with properties in a stored procedure.
 (Clicking on the plugin/project name takes you to the Project Details page for that project, and clicking on the procedure name takes you to the Procedure Details page for that procedure.)

Click **Run** to run this sample procedure and see the resulting job status on the Job Details page.

Implementation

Use the following instructions if you would like to practice creating another subprocedure for this "Calling a subprocedure" procedure.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow and cannot be transferred to your procedures. These instructions only provide a clearer explanation for how you can call subprocedures in your projects.

To create a step that executes a subprocedure:

- 1. On the Procedure Details page, click the **Create Step** link to see the Choose Step Type dialog.
 - · Select the Subprocedure tab.
 - Select Plugin and place your cursor in the text field and start typing EC-Tutorials.
 - For Procedure, click inside the text box and select Reserving a Resource for job step duration procedure or whichever procedure you would like to choose.
 - Click **Create** to go to the New Step page to create the step to call an existing procedure to use as a subprocedure.
- On the New Step page, enter any new step name of your choice. Notice that no parameters are required for this step.

- 3. Click **OK** to return to the Procedure Details page to see your new step in the table, then click **Run**.
- 4. When the Job Details page is displayed, you can see your new step and the procedures it is calling.

Related Information

Procedure - create new or edit existing procedure - Help topic Step - create new or edit existing step - Help topic

Checking the Outcome of Preceding Steps

ElectricFlow steps can have a *success*, *warning*, *error*, or *skipped* outcome. Sometimes it is useful to execute or skip subsequent steps based on the outcome of a previous step. Used with error handling behavior available at the step level, sophisticated flow control can be achieved easily.

This tutorial demonstrates how to see the outcome of a preceding step and use run conditions to execute steps based on that outcome.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

The Procedure Details page for the procedure, "Checking outcome of preceding step", contains 3 steps:

- step 1 force an error This step forces an error by attempting to run on a resource that does not exist.
 - This action was achieved by hard-coding the step named "step 1 force an error" to run on a resource called "this-does-not-exist".
- step 2 execute if step 1 failed This step uses JavaScript to check the outcome of the first step and executes if the first step failed.
 - You can view the JavaScript for the run condition by clicking on the step name to take you to the Edit Step page—see the Run Condition field in the Advanced section.
 - This statement checks the result of "step 1" for an error. If there is an error, "step 2" will execute.
- step 3 execute if step 1 succeeded This step also uses embedded JavaScript to check the outcome of the first step and executes if the first step succeeded.
 - This is a slight variation of "step 2". Instead of looking at "step 1" for an error, this step checks for success and executes only if "step 1" succeeds.
 - You can view the JavaScript for the run condition by clicking on the step name to take you to the Edit Step page—see the Run Condition field in the Advanced section.

Click **Run** to run this sample procedure and see the resulting job status on the Job Details page.

Implementation

Use the following instructions if you would like to practice creating steps to check the outcome of a preceding step.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow and you cannot transfer any steps from this tutorial to your projects. These instruction are provided only to give you a clearer definition for how to recreate this process in your own projects.

- Click the Create Step link.
 - o In the Choose Step Type dialog, choose **Command** step to go to the New Step page.
 - Enter a Name for your step.
 - In the Command(s) text box, you can use the information in the above section. "Copy and paste" the
 provided JavaScript, changing the step name (between the brackets) to the new step name you
 supplied.
- Repeat this process to create 3 steps, clicking **OK** to save each step.
- When your 3 new steps are displayed on the Procedure Details page, click **Run** to see your steps running on the Job Details page.

Related information

Step - create new or edit existing step - Help topic

Conditional Execution

Process automation frequently contains actions to be executed only if certain conditions are met. Conditional execution can be applied directly to your procedures rather than having to maintain larger scripts to handle these actions. For example, you may need to control whether or not to build on different operating system types.

This tutorial shows you how to implement simple conditional execution in a step.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

This procedure defines one parameter and then uses a run condition to check the parameter value to decide whether or not to execute a step.

- The Procedure Details page for the Conditional execution procedure contains a step named step 1. (Clicking on this step takes you to the Edit Step page.)
- The Action column displays the command to be executed.
- The Parameters section displays a parameter with values.

 The parameter named "Execute step 1" was created as a checkbox.

If checked, the parameter equates to "true".

In the run condition for "step 1", the parameter is read using the following property reference notation: $\frac{1}{2} \left(\frac{1}{2} \right) = \frac{1}{2} \left(\frac{1}{2} \right) \left$

```
$[Execute step 1].
```

(Clicking on the Parameter Name takes you to the Edit Parameter page to make modifications you may need.)

Click **Run** (at the top of the page) to run this sample procedure, which then takes you to the Run Procedure page where you could cause step 1 to execute by selecting the Execute step 1 checkbox.

Click Run again to go to the Job Details page to see the job created by running this procedure.

Implementation

Review the following instructions if you would like to see how to reproduce a conditional execution step.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow. These step-by-step instructions are provided only for more detailed information so you can easily see how to create your own "conditional execution" solutions within your procedures.

- Click the Create Parameter link to go to the New Parameter page.
 - Enter a new parameter name.
 - Click the Type down-arrow and select Checkbox.
 - Select Initially checked?
 - For Default Value, type "true".
 - o Select Required?
 - Click **OK** to save your parameter and return to the Procedure Details page.
- Click the step name to go to the Edit Step page.
 - Change the Run Condition text to check the new parameter name you supplied above.
 - Click **OK** to return to the Procedure Details page.
 - Click Run to go to the Run Procedure page to see 2 parameters where you can now decide if both are required or not.
 - Click Run to go to the Job Details page.
 - If only 1 parameter was "checked" as required, you will see the "skipped" notation in the Status column.

Related information

More complex run conditions can be implemented using inline JavaScript.

See the tutorial, "Checking outcome of preceding step" for a more complex run condition example.

Custom parameter layouts

By default, parameters for procedures are displayed in alphabetical order based on the parameter name. Using the ec_parameterForm property allows parameters to be displayed in a specified order, different from the default, alphabetized order by name.

This tutorial shows how to create a custom parameter layout for procedures.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

- This procedure implements two parameters.
- The Action column contains the command to execute. In this case, the action shows the ec_ parameterForm property which was attached to the procedure.
- Notice the checkmarks in the Req? column, which indicates both parameters are "required" each time this procedure runs.

- In the Custom Procedure Properties table, notice the "formElement" values supplied to determine the new parameter order.
 - If the procedure is run *without* the ec_parameterForm property, the parameter input form displays parameters in default, alphabetical order:

200	-
first parameter:	
second parameter:	

Click Run.

On the Run Procedure page, notice you are presented with a Parameter section that shows the parameters in reverse order because values supplied for the ec_parameterForm property were called to create this new order.

Parameters		
	this is the label for "second parameter"	
	this is the label for first parameter	

The following XML creates the new order for these parameters:

Click **Run** to run this sample procedure and see the resulting job status on the Job Details page.

Implementation

Note: This tutorial is intended for viewing, but not modifications within its procedure, and any changes you might make will not be saved when you upgrade ElectricFlow.

However, in your own projects, you can use the sample XML above to change the parameter order in any of your procedures.

If you copy and paste the code samples we provided, using the ec_parameterForm property in one or more of your procedures, make sure you replace our example values with your parameter names, labels, and so on. If you have more than two parameters to reorder, create any number of additional <formElement> sections you need.

Related information

Customizing the ElectricFlow UI/Customizing Parameters - Help topic
Step - create new or edit existing step - Help topic
Parameter - create new or edit existing parameter - Help topic

Email Notifications

Use email notifiers to send information to people who need or want ElectricFlow notifications, whether or not they actually use ElectricFlow. For example, you might want to set up an email notifier to send log file error excerpts matched by *postp* to a team or individual responsible for investigating the error.

This tutorial illustrates how email notifiers can be attached to steps to automatically send an email when jobs, steps, or workflows execute.

Prerequisite: If you do not have an Email Configuration, you must create one before you can set up an email notification. An Email Configuration tells the ElectricFlow server which host to use for mail delivery. Remember the email configuration name you choose because you will need it later.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

- A step named "do something" that puts a short message in the log file
- A parameter that defines the email address for where you want to send the notification. This parameter is
 used as the destination value for the notifier named "start notifier". This notifier sends an email when the
 procedure runs.
- An email notifier named "start notifier" with its Type specified as "onStart".
 - Click on the notifier name to view the notifier details.
 - Make sure the value in the Email Configuration field matches the email configuration name from the "Prerequisite" section above.
 - Click **OK** to save any changes and return to the Procedure Details page.

Click Run.

On the Run Procedure page, in the Parameters section, enter an email address, an address list, or email group name.

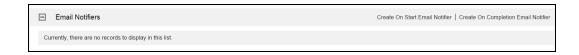
Click Run again to run this sample procedure and see the resulting job status on the Job Details page.

Implementation

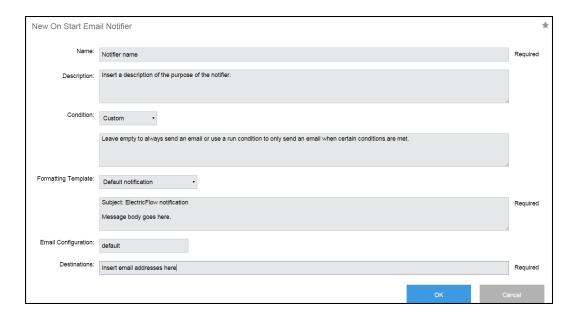
Note: Use the following information to create an email notifier in your own project, assuming you have already setup an email configuration. Any changes you might make within this tutorial will not be saved when you upgrade ElectricFlow.

Using the ElectricFlow automation platform UI to add an email notifier:

 On the Procedure Details page, go to the Email Notifiers section, and click Create On Start Email Notifier.



2. Enter information in the fields on the New On Start Email Notifier page. See the example below.



Note: If the Email Configuration field is left blank, the system uses the email configuration with the name "default".

Click **OK** to save the notifier.

When the procedure is run, an email notification is sent to the email addresses specified in the Destinations field. In this tutorial example, the Destinations field refers to the parameter value passed in to the procedure when it is run.

The value is referenced using the following notation:

\$[/myJob/Send notification email to]

Related information

Email Configuration - create new or edit existing email configuration - Help topic Email Notifier - create new or edit existing email notifier - Help topic

Executing Tasks on All Resources in a Pool

Creating resource pools (grouping resources) means you can reference many resources using one name, the pool name. When configuring a step, you can "broadcast" the step to all resources in a pool. For example, you might want to start an application deployment at the same time on all resources in a pool.

This tutorial shows how to group resources into a resource "pool".

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

This procedure implements three steps:

1. The first step, Setup environment, creates a pool for resources.

Clicking on the step name takes you to the Edit Step page to see the full command text supplied to create this step—a partial section is displayed in the Action column.

Note: Functionality to create this step is beyond the scope of this tutorial. For more information, see the link to "Using ectool and the ElectricFlow API" below.

The following resources were created and added to a pool named EC-Tutorial-pool:

EC-Tutorials-resource1

EC-Tutorials-resource2

EC-Tutorials-resource3

- 2. The second step, Execute command on all resources in pool, executes an echo statement on the resource EC-Tutorial-pool and has the "broadcast" option selected, which results in the step executing on all resources in the pool simultaneously.
- 3. The third step, Cleanup, deletes the resources created in the Setup environment step.

Click **Run** to run this sample procedure and see the resulting job status on the Job Details page.

Implementation

Note: This tutorial is intended for viewing only and/or clicking on the available links to see more information. Any modifications you might make within this tutorial cannot be transferred to your projects and will not be saved when ElectricFlow is upgraded.

Before you can create the steps illustrated in this tutorial procedure for your own purpose, you must have already defined your resources and created resource pools.

See the following list of related links for more help with creating and using resource pools.

Related information

```
Resources - Help topic

Resource Pool - create new or edit existing pool - Help topic

Resource Pool Details - Help topic

Step - create new or edit existing step - Help topic
```

Factory Procedures

Factory procedures allow dynamic behavior in process automation. To implement a factory procedure, you must first create a procedure that other procedures can use as a template. A procedure uses the template to construct new procedures at runtime. Instead of creating steps or complex scripts to account for all possible scenarios, procedures can be created "on the fly" based on user input, resulting in a very dynamic system.

This tutorial shows how to create factory (dynamic) procedures, which are procedures that modify themselves at runtime based on parameters passed into them.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-

<version>, and click the procedure name.

An explanation for what you see and what is actually happening:

This tutorial contains 2 procedures and creates a third procedure "on the fly":

- "Factory procedures" this procedure uses the procedure template (Utility Procedure Factory Template).
 - This procedure has a parameter called "number of steps". This parameter takes an integer and is used to specify how many new steps should be generated.
 - **Note:** For this tutorial example, only one procedure is constructed to use the template procedure, but you may have multiple procedures that use the factory template procedure.
- "Utility Procedure Factory Template" you do not see this procedure on this Procedure Details page, but it is included in the EC-Tutorials project. This is the factory procedure template used when generating new steps in the main procedure. This procedure implements a single step that writes some text to the log file.

This procedure is designed to do a repetitive task, repeating the task a variable number of times. For example, you might need to run a test framework "n" number of times.

Note: You can view this procedure in the EC-Tutorials project.

The main procedure, "Factory procedures", implements these steps:

 The first step: "create dynamic procedure" is a Perl script that uses the Perl API to create a new procedure.

The following script creates a new procedure containing the number of steps specified by the parameter:

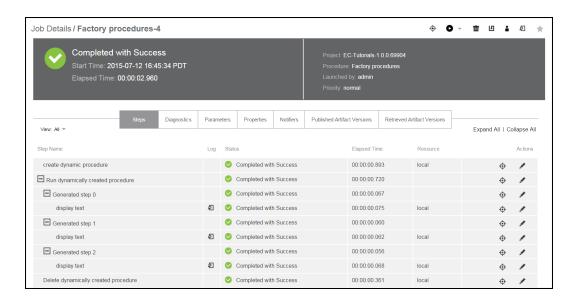
```
my $procedureName = "Dynamically Created Procedure $[jobId]";
$ec->createProcedure("$[/myProject/projectName]", $procedureName);
for(my $i=0;$i<$[number of steps];$i++)
{</pre>
```

```
$ec->createStep("$[/myProject/projectName]", $procedureName, "Generated step
$i", {subprocedure => "Utility Procedure - Factory Template",
    actualParameter => {actualParameterName => 'text', value => "Executing
generated step $i"}});
}
$ec->setProperty("/myJob/dynamicProcedure", $procedureName);
```

- 1. First, the procedure generates a unique name for the new procedure to be created. The name is created by appending the value of the \$[jobId] property string.
- 2. A new procedure is created in the current project.
- Based on the number of steps to be created that were passed in as a parameter, a number of subprocedure steps are created in the procedure. The subprocedure steps call the utility procedure associated with this tutorial.
- 4. The name of the newly created procedure is stored in /myJob/dynamicProcedure
- The second step: "Run dynamically created procedure" calls the procedure that was created in Step 1. Step 2 appears to call nothing, but in fact it is calling a procedure referred by \$[/myJob/dynamicProcedure].

Click Run to run this sample procedure and go to the Job Details page.

Running the procedure and specifying the number of steps to be generated as 3, results in the following job being executed:



Notice the steps that execute in parallel in the dynamically created procedure: Generated step 0, Generated step 1, and Generated step 2

Implementation

To create a "factory procedure" process in your project:

- 1. Create a utility procedure this procedure performs a task.
 - Create the steps you need for your task.
- 2. Create a second procedure the "factory procedure," providing a name of your choice for this procedure.
 - Create a minimum of 3 steps in this procedure.
 - step 1 create a new procedure that contains steps, each calling the utility template procedure.
 - o step 2 this step executes the procedure created in step 1.
 - o step 3 this step deletes the procedure created in step 1.

Each time you run the "factory procedure," it creates a temporary procedure to run the utility procedure.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow.

Postp extension

postp is a powerful ElectricFlow feature (postprocessor) you can use to monitor step output in real-time and take action. Frequently, the *action* is to extract information from step output, store it in the ElectricFlow database, and set properties for reporting. In addition, postp can be used to execute functionality depending on whether or not a pattern is found. Postp, an essential part of the ElectricFlow toolbox, provides fast feedback to users about job step status.

Note: Postp matchers are written in Perl. Postp extensions are either loaded from a file located on the ElectricFlow server or from a stored property.

This tutorial demonstrates how to construct a simple custom *postp* matcher to extend postp capabilities, and loads the postp extension from a property.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-

<version>, and click the procedure name.

An explanation of what you see:

- This procedure implements a single step, "get directory listing", that retrieves a directory listing using a simple Perl script. Perl is used for the example to ensure the same output no matter which operating system is used for this tutorial. Understanding the Perl code used to generate the tutorial listing is not required for this tutorial.
- A property was created. This procedure has a postp extension attached in a property named matcher and contains the following Perl code:

```
use ElectricCommander;
push (@::gMatchers,
{
  id => "fileCountMatcher",
  pattern => q{(\d+ File\(s\))},
  action => q{
  setProperty("summary", "Matcher $matcher->{id} found the following output\n\n$1");
  }
});
```

The postp matcher has an identifier ("id"), a regular expression ("pattern") that matches text and an "action" to be performed when text matching the regular expression is found. In this case, the regular expression "id" is fileCountMatcher, and an example of the output being matched is "2 File(s)". When the pattern is matched, the action to perform is to set the summary property to the text matched by the regular expression.

The custom postp matcher is loaded using the following command in the Postprocessor field when creating the step:

```
postp -loadProperty /myProcedure/matcher
```

Click on the step name to go to the Edit Step page to see how this step was created—note the text in the Command box and the Postprocessor field.

Click Run to run this sample procedure and see the resulting job on the Job Details page.

See the Status column. The number of files in the file system root where the workspace is located is displayed.



Implementation

To try using postp functionality in your own project:

- Edit an existing procedure or create a new one with a "matcher" as a property.

 For your practice example, you can "copy and paste" the matcher example provided in this tutorial (above).
- Create a new Command step for this procedure filling in the Postprocessor field on the New Steps page.

Again, for practice, you can use the postp string provide above in this tutorial.

When you run your new procedure, you will see results similar to those you saw after running the tutorial example procedure.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow.

Related information

Postprocessors - Help topic

Step - create new or edit existing step - Help topic

Publishing and Retrieving an Artifact

To ensure traceability in processes, output needs to be versioned and tracked as it is handed from one process stage to another. With ElectricFlow Artifact Management, implementing your processes is auditable and secure.

This tutorial shows how to produce and consume artifacts as part of an ElectricFlow process.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see - the following concepts are illustrated in this procedure:

The first five steps are simple ec-perl scripts that create files and directories as use as an artifact.

The first step, "Create myartifiact.txt file", creates a text file in a workspace that can be added as an artifact. Click this step to go to the Edit Step page to see the full "dummy" script used to set up the artifact in the Command(s) field. This script creates a file named myartifact.txt.

Note: There is no code specific to ElectricFlow in this step, so understanding the Perl code used to generate the text file is not necessary.

The next four steps create a directory, a subdirectory and the second text file in it.

• The "Publish artifact" step calls the EC-Artifact: Publish procedure. This step is configured to place the file myartifact.txt, created in the Setup step, into an artifact.

To do this, the EC-Artifact: Publish procedure requires the following information:

- The artifact where you want to publish. This is in the format Artifact Group: Artifact Key.
 - In this case, the name of the artifact (where you want to publish) is EC-Tutorials:MyArtifact.
- The artifact version being published. The version is in the form of major.minor.patch-qualifier-buildnumber or major.minor.patch.buildnumber-qualifier. For this tutorial, the version is set to 1.0.\$[jobId]-tutorial.

 Using the \$[jobId] property expansion results in a new artifact being created every time the procedure is run.
- The repository server from which you need to retrieve the artifact. This tutorial uses a repository named default that was created when ElectricFlow was installed.

Note: To see how this information is specified in this step, click the "Publish artifact" step name to go to the Edit Step page.

- The "Retrieve artifact" step calls the EC-Artifact: Retrieve procedure. This step takes two required
 parameters and several optional parameters (see the Edit Step page/Parameter section for this step).
 The required parameters are:
 - Artifact EC-Tutorials:MyArtifact (this is the name of the artifact to retrieve).
 - **Version** 1.0.\$[jobId] -tutorial (this is the artifact version to be retrieved either an exact version or a version range). This tutorial uses an exact version.

The optional parameters are:

- **Retrieve to directory** the location where the artifact files are downloaded when they are retrieved. This is where the retrieved artifact versions are stored.
- Overwrite the default is update where the artifact files are updated when they are retrieved.
 The other values are true (the files for the retrieved artifact versions are overwritten) or false (the files for the the retrieved artifact versions are not overwritten).

The **Retrieved Artifact Location Property** parameter is the name or property sheet path that the step uses to create a property sheet. The default value is

/myJob/retrievedArtifactVersions/\$[assignedResourceName].

This property sheet has information about the retrieved artifact versions, including their location in the file system. It displays the location from where the artifact is to be retrieved and retrieves the property value in the <code>/myJob/retrievedArtifactVersions/\$[assignedResourceName]</code>.

For the **Filter(s)** parameter, enter search filters, one per line, applicable to querying the ElectricFlow database for the artifact version to retrieve.

For more information about these parameters, go to Retrieve Artifact Version Step.

- The "Output location that artifact was retrieved to" step prints the message "Artifact extracted to: <directory where the retrieved artifact will be located>".
- Click Run to run this sample procedure and see the resulting job on the Job Details page.

Implementation

To publish and retrieve artifacts, apply these concepts to your artifact projects, creating the steps as required to your procedure.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow.

Related information

```
Artifact Management - Help topic

Artifact Details - Help topic

Job Details - Help topic

Publish Artifact Version step - Help topic

Retrieve Artifact Version step - Help topic
```

Reserving a Resource for Job Step Duration

Sometimes you may want a resource to perform work for only one job at a time. Because a ElectricFlow agent is capable of executing steps for many jobs on a resource, this can produce undesirable performance results. For example, you might want a dedicated resource to stress test an application after you have ElectricFlow initiate a load testing task.

This tutorial shows how to reserve a resource for one job step only while that job step is running.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

- This procedure locks the resource for the step and executes the tasks.
- To demonstrate this functionality, two steps are implemented.
 - Notice that both steps are set to run in parallel on the local resource.
 - Both steps request a lock of the resource for the duration of the step execution.

The functionality implemented in each step:

- A simple Perl sleep statement to force a lock over a period of time when the procedure is run
- Select "Step" in the Retain Exclusive field in the Advanced section on the New/Edit Step page

Click a name in the Step Name column (to go to the Edit Step page for that step). You will see the Perl sleep statement added to the Command text box, and "Step" selected on the Retain Exclusive field.

Instead of running the steps in parallel as you may have previously requested, the steps will now execute serially. The second step must wait for the resource to become available because it was reserved by the first step.

Click **Run** to run this sample procedure and see the resulting job on the Job Details page.

Implementation

Create or edit Command Steps for your procedures if you need this functionality.

- Go to the Procedure Details page for the procedure where you want a step to reserve a resource.
- Add a Command Step to your procedure (use the step creation link at the top of the table), or edit an existing step.
- In the Command text box, enter a script if you want to specify any time constraints.
- On the New/edit Steps page, in the Advanced section, choose "Step" in the Retain Exclusive field.

The next time you run this procedure, your job step will have exclusive use of the resource you specified.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow.

Related information

Step - create new or edit existing step - Help topic

Running Steps in Parallel

The main advantage to running steps in parallel is increased performance. You can set a step to run serially as a synchronization point and then execute more steps in parallel. For example, running steps to execute tests on different platforms concurrently would mean numerous tests could complete in the same time period.

This tutorial demonstrates how to run steps in parallel.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

This procedure contains 5 steps:

- Step 1 and 2 run in parallel.
- The third step (sync) is a sync point where the procedure waits until steps 1 and 2 finish executing.
- Steps 4 and 5 execute in parallel after the sync point.
- Steps 1, 2, 4, and 5 execute simple Perl scripts that print a message and "sleep" for a period of time.

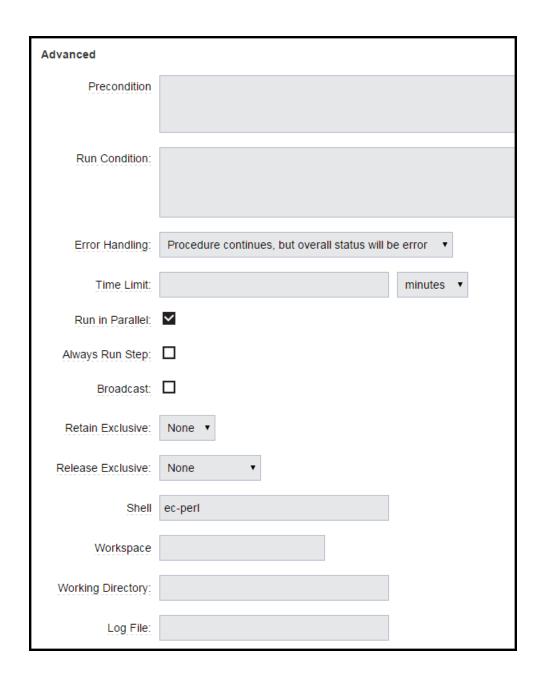
Note: Using "sleep" is not required, but used here to make sure each parallel step runs long enough to be visible while running this tutorial.

Click **Run** to run this sample procedure and see the resulting job status on the Job Details page.

Implementation

Note: Use the following information to create steps to run in parallel in your procedures. Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow.

Steps will run in parallel if you select the Run in Parallel option in the Advanced section on the Step - create new or edit existing step page.



Related Information

Step - create new or edit existing step - Help topic

Step Timeouts and Steps that Always Run

Complex processes can be brittle so ElectricFlow provides a way to account for different errors in your procedures. You can specify timeouts, error handling methods, and steps that always run no matter what errors have occurred in preceding steps in the procedure. These facilities can make sure your build does not "hang" in a situation where you would prefer the build to continue.

This tutorial shows how to construct a procedure so processes are terminated after a set time period and how to implement a step that always runs in the event of an error (to perform cleanup).

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-

*version
, and click the procedure name.

An explanation of what you see:

This procedure implements 2 steps:

(Clicking on the step name takes you to the Edit Step page to see how the step was created.)

• The "force timeout" step has a 10-second timeout and is set to abort the procedure and terminate running steps in the event of a timeout or error.

This step executes a sleep command using ec-perl, which makes the step execute for 15 seconds. Step error handling is set to "Abort procedure and terminate running steps". This step will always fail because the 15-second sleep time is longer than the 10-second timeout.

• The "always run" step set to run no matter what the error handling behavior is for any preceding steps.

This step uses an echo statement to output some text. The step property, Always Run Step, is set to "true" so this step runs regardless of the result for the preceding step.

Click **Run** to run this sample procedure and see the resulting job on the Job Details page.

On the Job Details page: Notice that the "always run" step completed with Success while the "force timeout" step completed with a TIMEOUT error.

Implementation

Note: This tutorial is intended for viewing only. Any changes you might make within this tutorial cannot be transferred to your projects will not be saved when you upgrade ElectricFlow.

As you create steps for your projects, you can duplicate these tutorial concepts. On the New Steps page, review the Advanced section. Notice the Error Handling and Always Run Steps options.

- For Error Handling, choose the option that best suits your purpose.
- For the Always Run Step option, select the checkbox to ensure this step will always run.

Storing and Retrieving Properties in a Job

When a job is running, it is often necessary to share data between steps. Because the data being shared is specific to the running procedure instance, the job, the data is more appropriately stored at the job level rather than the procedure level. Use the techniques illustrated in this tutorial to safely pass data between steps in a job.

This tutorial shows you how to create, set, and retrieve properties stored on a job.

To view an example procedure for this tutorial , go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see:

This procedure implements two steps.

- Write data to a property attached to the running job:
 "Step 1 Create a property attached to the job" calls ectool to create a property name property1.
 - The path to the property being set is "/jobs/\$[/myJob]/property1.
 - \$ [/myJob] is a shortcut to the running job name and can be used while the job is running. If the job number was 1000, the path "/jobs/\$[/myJob]/property1" might expand to "/jobs/job-1000/property1".
- Read data from the property created in Step 1:
 - "Step 2 Read the property attached to the job by Step 1" this step could read step values stored on a job, running as part of the job, or could be executed using a simple property expansion.
 - In this case, property1 was set in Step 1 and is referenced in Step 2 using the \$[property1] notation.
 - Two other ways to reference a property (shown by example in the Action column):
 - Using \$[/myJob/property1]
 - Using inline JavaScript with the notation \$[/javascript myJob.property1]

Click **Run** to run this sample procedure and see the resulting job on the Job Details page.

• After the job completes, click the icon in the Log column for Step 2 to view three lines each showing the data (the property value) that was stored in the property that was retrieved.

Implementation

If you need to retrieve properties, adapt these step concepts to your procedures.

- For help getting started, click on a step name in the tutorial procedure to go to the corresponding Edit Step page.
- You can reuse the text supplied in the Command(s) box, changing the "property1" name and other values to those more meaningful to your project/procedure.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow.

Related information

Procedure - create new or edit existing procedure - Help topic

Step - create new or edit existing step - Help topic

Working with Properties Stored in a Procedure

Using ElectricFlow properties is one of the keys to robust yet flexible process automation. Configuration and historical data persistence are common tasks in process automation.

This tutorial shows common ways to read and write properties on a procedure.

Note: For this tutorial, familiarity with creating procedures and steps is assumed. See the "Related Information" links below if you would like to review these topics.

To view an example procedure for this tutorial, go to the automation platform UI > **Projects** > **EC-Tutorials**-<*version*>, and click the procedure name.

An explanation of what you see—this procedure implements the following functionality:

- This procedure implements five steps. Each step name is a description of its action.
 - Click on a step name to go to the corresponding Edit Step page to see how the step was created.
 This page also contains text in the Description text box to describe what you see in the Action column on the Procedure Details page. The full text of what appears in the Action column is provided in the Command text box.
 - In the Custom Procedure Properties table, notice two properties were created: property1 and property2
- There are numerous ways to work with stored properties. Each of the 5 steps described below illustrates a different method.
 - Read data from property1 using the property substitution notation (syntax):
 The property value is read in this step using the property substitution reference
 \$[/myProcedure/property1].
 - Read data from property1 using inline JavaScript:
 The property name property1 is read from the procedure level using this JavaScript notation:
 \$[/javascript myProcedure.property1].
 - Perform a calculation using data from property 2 and inline JavaScript:
 In this example, the value of property2 is multiplied by 5 using this JavaScript notation:
 \$[/javascript myProcedure.property2 * 5].
 - Read data from a property1 using the ElectricFlow Perl API:
 Perl code is used to call the ElectricFlow API to retrieve the value of property1.

 The getProperty API returns an object so the property value must be retrieved after the getProperty call, using the findnodes method for the returned object.

```
#read the property
my $property1node = $ec->getProperty("/myProcedure/property1");
#retrieve the property value from the object returned to $property1node
my $property1 = $property1node->findnodes("//value")->string value();
```

Read data from property1 using ectool:
 ectool code is used to retrieve the value of property1:

```
ectool getProperty "/projects/$[/myProject/projectName]/procedures/Working
with properties stored in a procedure/property1"
```

Click **Run** to run this sample procedure and see the results on the Job Details page.

On the Job Details page, click the icon in the Log column to see the value read from the property.

Implementation

If you need to retrieve properties stored in a procedure, adapt any one of these step concepts to your procedures.

You can reuse the text supplied in the Command(s) box (from any Edit Step page in this tutorial), changing names and other values to those more meaningful to your project/procedure.

Note: Any changes you make within this tutorial will not be saved when you upgrade ElectricFlow.

Related Information

Procedure - create new or edit existing procedure - Help topic

Step - create new or edit existing step - Help topic

Properties - Help topic

Glossary

This glossary is a reference topic containing short descriptions for ElectricFlow objects, terms, and concepts.

Term	Description
access control	An access control list (ACL) determines if a particular user can perform a particular operation on a specified object. The list contains access control entries (ACE), each of which specifies a user or group and indicates whether certain operations are allowed or denied for that user or group. Using access control provides security for ElectricFlow system use. See the Access Control topic for more information.
ACE (Access Control Entry)	An ACL determines if a particular user can perform a particular operation on a specified object. The list contains access control entries (ACE), each of which specifies a user or group and indicates whether certain operations are allowed or denied for that user or group. Using access control provides security for ElectricFlow system use. See the Access Control topic for more information.
ACL (Access Control List)	An ACL determines if a particular user can perform a particular operation on a specified object. The list contains <i>access control entries</i> (ACE), each of which specifies a user or group and indicates whether certain operations are allowed or denied for that user or group. Using access control provides security for ElectricFlow system use. See the Access Control topic for more information.
actual parameter	An actual parameter is an object that provides the value for a parameter, which is passed to a procedure when it is invoked. Actual parameters can be provided for jobs and nested subprocedures within a job. Actual parameters are different from "formal parameters": formal parameters define parameters a procedure is expecting, and actual parameters provide values to use at run-time.
admin	"admin" is a special built-in user that has universal ElectricFlow access. If you log in as admin, you can perform any operation in the system, regardless of access control limitations.

Term	Description
agent	An agent is an ElectricFlow component that runs on each machine where job steps can execute. The agent works under the ElectricFlow server's control to execute job steps, monitor their progress, and record information about their completion. A single agent process can manage multiple job steps executing concurrently on a single machine. See the Web Interface Help > Resources topic for more information.
artifact	An artifact is a top-level object containing artifact versions, a name template for published artifact versions, artifact specific properties, and access control entries to specify privileges.
artifact key	An artifact key is an identifier for an artifact and the "key" component of the artifact name.
artifact repository	See the Artifact Management > Artifact objects > Repository topic for more information.
artifact version	An artifact version is a collection of 0 to ${\tt N}$ files that were published to an artifact repository.
backing store	The backing store is the directory on the repository server where artifact versions are stored. By default, the backing store is the <datadir>/repository-data directory in the repository installation—this default setting can be changed.</datadir>
compression	Compression reduces transfer time when publishing an artifact. However, compression also adds overhead when computing the compressed data. If files included in the artifact version are primarily text files or are another highly compressible file format, the benefit of reduced transfer time outweighs the cost of computing compressed data.
continuous integration	Using continuous integration means a build is launched every time code changes are checked into a Source Control Management (SCM) system.
	The ElectricFlow ElectricSentry component is the engine for continuous integration, while the CI Continuous Integration Dashboard is the front-end user interface for ElectricSentry.
credential	A credential is an object that stores a user name and password for later use. You can use credentials for user impersonation and saving passwords for use inside steps. Two credential types are available: stored or dynamic.

Term	Description
custom property	Custom properties are identical to intrinsic properties and when placed on the same object, are referenced in the same manner and behave in every way like an intrinsic object-level property with one exception: they are not created automatically when the object is created. Instead, custom properties can be added to objects already in the database before a job is started, or created dynamically by procedure steps during step execution.
	Custom properties in a property sheet can be one of two types: <i>string</i> property or a <i>property sheet</i> property. String properties hold simple text values. Property sheet properties hold nested properties. Nested properties are accessed by way of the property sheet property of their containing sheet.
description	A description is an optional plain text or HTML description for an object. Description text is for your use, ElectricFlow does not use this information. If using HTML, you must surround your text with httml tags. The only HTML tags allowed in the text are: <a> <pre> <style> <tc> </h></td></tr><tr><td>diagnostic extract</td><td>A diagnostic extract is a log file portion from a job step, typically describing an error or interesting condition, extracted by a postprocessor and saved for reporting. The postprocessor usually places this information in an XML file in the top-level job workspace directory, and then sets a property that contains the filename.</td></tr><tr><td></td><td>The ElectricFlow <i>postp</i> postprocessor uses filenames like diag-2770.xml, where "2770" is the unique identifier for the step. Other postprocessors you may use can have a different filename configuration.</td></tr><tr><td>dynamic credential</td><td>Dynamic credentials are captured when a job is created. Dynamic credentials are stored on the server temporarily until the job completes and then discarded.</td></tr><tr><td>ec-perl</td><td>ec-perl is a small wrapper program installed as part of ElectricFlow. When the ecperl wrapper runs, it sets up the environment, finds, and calls ElectricFlow's copy of Perl, passing all of its parameters to Perl.</td></tr><tr><td>ectool</td><td>ectool is the ElectricFlow command-line application that provides control over the ElectricFlow system if you prefer using a command-line interface rather than the ElectricFlow web interface. Most functions that can be invoked through the ElectricFlow web interface can be invoked using ectool.</td></tr></tbody></table></style></pre>

Term	Description
ElectricAccelerator	ElectricAccelerator is a software build accelerator that dramatically reduces software build times by distributing the build over a large cluster of inexpensive servers. Using a patented dependency management system, ElectricAccelerator identifies and fixes problems in real time that would break traditional parallel builds. ElectricAccelerator plugs into existing Make-based infrastructures seamlessly and includes web-based management and reporting tools.
ElectricSentry	ElectricSentry is the ElectricFlow engine for continuous integration, integrating with numerous Source Control Management (SCM) systems. ElectricSentry is installed automatically with ElectricFlow and is contained in a plugin named ECSCM and in the Electric Cloud project. Note: The CI Continuous Integration Dashboard is the front-end user interface for ElectricSentry.
email configuration	Before you can send an email notifier, you must set up and email configuration, which establishes communication between the ElectricFlow server and your mail server.
email notifier	After setting up the ElectricFlow server and your mail server to communicate, you can send email notifications (notifiers). You can attach email notifiers to procedures, procedure steps, and state definitions.
Event log	See log.
Everyone	A special intrinsic access control on page 989group that includes all users.
filter	Two filter categories: Intrinsic filters - these filters provide a convenient way to access certain well-defined fields for jobs. Custom filters - these filters allow you to access a much broader range of values, including custom properties. Any values accessible through an intrinsic filter can be checked using a custom filter also (though not as conveniently).
formal parameter	A formal parameter is an object that defines a parameter expected by a procedure, including its name, a default value, and an indication of whether the parameter is required. Formal parameters are different from "actual parameters": formal parameters define the kinds of parameters a procedure is expecting, and actual parameters provide values to use at run-time.

Term	Description
gateway	To communicate with a resource, workspace, or artifact repository server in another zone, a "gateway" must be created.
	A gateway object contains two resource (agent) machines. For example, GatewayResource1 and GatewayResource2 are each configured to communicate with the other. One gateway resource resides in the <i>source</i> zone and the other in the <i>target</i> zone.
	A gateway is bidirectional and informs the ElectricFlow server that each gateway machine is configured to communicate with its other gateway machine (in another zone).
group	A group defines a collection of users for access control purposes. A group can be defined externally in an LDAP or Active Directory repository, or <i>locally</i> in the ElectricFlow server.
	See the ElectricFlow Help > Web Interface Help > Users and Groups > Groups topic for more information.
impersonation	Impersonation is a mechanism that allows a job step to execute under a particular login account (the ElectricFlow agent "impersonates" a particular user during the execution of that step). Impersonation is implemented using credentials.Impersonation is implemented using credentials. See the ElectricFlow User Guide for more information.
inheritance	A feature of the ElectricFlow access control mechanism where access to a particular object is determined by the access control list for that object, and also by the access control lists of the object's parent and other ancestors. Each object can be configured to enable or disable inheritance from its ancestors.
intrinsic property	Intrinsic properties represent attributes that describe the object to which they are attached. ElectricFlow automatically provides intrinsic properties for each similar type object within ElectricFlow. For example: Every project has a description property that can be referenced with a non-
job	local property path such as /projects/Examples/description. A job is the output associated with invoking a ElectricFlow procedure. A new job
	is created each time you run (execute) a procedure.

Term	Description
job configuration	A job configuration is an object containing all parameter and credential information needed to run a procedure. A Job Configuration section is provided as part of the ElectricFlow Home page to make it easy for you to invoke your favorite configurations with a single mouse click. You can create job configurations in three ways:
	 From the Job Details page for a previously invoked job, click the Save Configuration link at the top of the page. Your saved job configuration will be displayed on your Home page.
	 Create a job configuration from "scratch" by clicking the Create link in the Job Configurations section (on the Home page). In the Create Configuration pop-up menu, select the project and procedure you want to use for creating this configuration.
	 On the page for editing a schedule, click the Save Configuration link at the top of the page. Your saved configuration will be displayed on your Home page.
job name template	This is the template used to determine the default name for jobs launched from the procedure. You can create a Job Name Template when you create a procedure. For example:
	<pre>In the Job Name Template field, you might enter: \$[projectName]_\$[/increment /myproject/jobCounter]_\$[timestamp]</pre>
	which produces a name like: projectFoo_1234_20140102130321
	You can enter any combination of elements to create procedure names more meaningful to you. For example, you could choose to include the build number and procedure name.
jobs quick view	A Jobs Quick View section is one of the facilities provided on the ElectricFlow Home page. This section allows you to define a category of jobs interesting to you (such as all running jobs or all jobs for a particular product version). Your Home page can display the last several jobs in each category you define.
job step	After a procedure is executed, the resulting job contains one job step for each step in the original procedure. The job step records information about the procedure step execution, such as the command executed, the resource where it executed, execution time, and error information.

Term	Description
job workspace	A directory (containing all files and subdirectories) allocated by ElectricFlow for a particular job. Each job workspace is allocated as the child of a workspace root directory. See workspace.
local group	A group defined <i>inside</i> ElectricFlow, as opposed to a group defined in an external repository. A local group can refer to both local and remote users, whereas a group in an external repository refers to users in that repository only. See group.
local user	A user defined <i>inside</i> ElectricFlow, as opposed to a user defined in an external repository. If a user defined in an external repository has the same name as a local user, the external user is not accessible. Local users are not visible outside ElectricFlow. We recommend using external accounts whenever available, but you may need to create local users if you do not have a shared directory service or if you need special accounts to use for ElectricFlow only. See user.
log	 ElectricFlow provides a log for events generated anywhere in the system, including jobs and workflows. Note: From the Administration tab, the default view for the Event Log page is the warning (WARN) level. For workflow and job event logs, the default view from their respective pages is the information (INFO) level. To see only events for a single workflow, select the Workflows tab, then a workflow Name to go to the Workflow Details page and click the View Log link at the top of the page. To see only events for a single job, select the Jobs tab, then the Job name to go to the Job Details page and click the View Log link at the top of the page. To see only events for a specific object, select the Search tab to go to the Define Search page. For example: You can select the Object Type, "Log Entry", then click the Add Intrinsic Filter link. Select the down-arrow where you see "Container" autopopulated and select "Container Type. Use the "equals" operator, then select the next down-arrow to choose an object. Click OK to start the search. See the Event Log topic for more information. See the ElectricFlow Help > Web Interface Help > Event Log topic for more information.

Term	Description
matcher	A <i>matcher</i> controls the postp postprocessor. Use matchers to extend postp with additional patterns if you find useful patterns in your log files undetected by postp. A matcher contains a pattern that matches lines in a step's log and actions to carry out if/when the pattern matches.
misfire policy	A misfire policy allows you to manage how a schedule resumes in cases where the normal scheduled time is interrupted. Available options are: skip (all misfires are ignored and the job runs at the next scheduled time) and run once (after one or more misfires, the job runs at the soonest time that occurs within an active region). See schedule.
parameter	A property value passed into a procedure when it is invoked (at <i>run</i> time), and used by the procedure to change its behavior. Two types of parameters: actual and formal.
plugin	A plugin is a collection of one or more features, or a third-party integration or tool that can be added to ElectricFlow. Plugins are delivered as a JAR file containing the functional implementation. When a plugin is installed, the ElectricFlow server extracts the JAR contents to disk into a configurable plugins directory.
	A plugin has an associated project that can contain procedures and properties required by the implementation. A plugin can provide one or more new pages for the web interface and may also provide a configuration page so you can provide additional information that may be necessary to implement the plugin.
polling frequency	The <i>polling frequency</i> is how often the ElectricSentry continuous integration engine is set to look for new code check-ins. The default is set to every 5 minutes, but this number can be adjusted.
pool	Also known as "resource pool". A pool is a collection of resources. If a step specifies a pool name as its resource, ElectricFlow can choose any available resource within that pool.

Term	Description
postp	postp is a postprocessor included with ElectricFlow. postp uses regular expression patterns to detect interesting lines in a step log. postp is already configured with patterns to handle many common cases such as error messages and warnings from gcc, gmake, cl, junit, and cppunit, or any error message containing the string "error." postp also supports several useful command-line options, and it can be extended using "matchers" to handle environment-specific errors. See matcher.
postprocessor	A postprocessor is a command associated with a particular procedure step. After a step executes, the postprocessor runs to analyze its results. Typically, a postprocessor scans the step log file to check for errors and warnings. Also, it records useful metrics such as the number of errors in properties on the job step, and extracts step log portions that provide useful information for reporting. ElectricFlow includes a standard postprocessor called <i>postp</i> for your use and you can "extend" postp. See matcher.
preflight build	A preflight build provides a way to build and test a developer's changes before those changes are committed. A "post-commit" source tree is simulated by creating a clean source snapshot and overlaying the developer's changes on top of it. These sources are then passed through the production build procedure to validate the changes work successfully. Developers are allowed to commit their changes only if the preflight build is successful. Because developer changes are built and tested in isolation, many common reasons for broken production builds are eliminated.
privileges	 ElectricFlow supports four privilege types (for access control and security) (for access control and security) for each object: Read - Allows object contents to be viewed. Modify - Allows object contents (but not its permissions) to be changed. Execute - If an object is a procedure or it contains procedures (for example, a project), this privilege allows object procedures to be invoked as part of a job. For resource objects, this privilege determines who can use this resource in job steps. Change Permissions - Allows object permissions to be modified.

Term	Description
procedure	A procedure defines a process to automate one or more steps. A procedure is the ElectricFlow unit you execute (run) to carry out a process. A step in one procedure can call another procedure (in the same or different project), and this procedure then becomes known as a "subprocedure" (also known as a "nested" procedure). The step can pass arguments to the subprocedure.
project	A project is a top-level container for related procedures, workflows, schedules, jobs, and properties, which is used to isolate different user groups or functions, and also encapsulate shared facilities. Projects have two purposes:
	 Projects allow you to create separate work areas for different purposes or groups of people so they do not interfere with each other. In a small organization, you might choose to keep all work within a single project, but in a large organization, you may want to use projects to organize information and simplify management.
	 Projects simplify sharing. You can create library projects containing shared procedures and invoke these procedures from other projects. After creating a library project, you can easily copy it to other ElectricFlow servers to create uniform processes across your organization.
project principal	Project principal is a special user ID associated with each project. If a project name is "xyz," the project principal for that project is "project: xyz" (with an embedded space). This principal is used when procedures within the project are run, so you can create access control entries for this principal to control runtime behavior.

Term	Description
property	A property is a <code>name-value</code> pair associated with ElectricFlow objects to provide additional information beyond what is already built into the system. Built-in data is accessible through the property mechanism also. Two types of properties: <code>intrinsic</code> and <code>custom</code> .
	ElectricFlow provides Intrinsic properties and allows you to create Custom properties. Note: Intrinsic properties are case-sensitive. Custom properties, like all other object names in the ElectricFlow system, are case-preserving, but not case-sensitive.
	• Intrinsic properties These properties represent attributes that describe the object to which they are attached, and are automatically by ElectricFlow for each similar type object. For example, every project has aDescription property that can be referenced with a non-local property path such as /projects/Examples/description.
	Custom properties Custom properties are identical to intrinsic properties and when placed on the same object, are referenced in the same manner, and behave in every way like an intrinsic object-level property with one exception: they are not created automatically when the object is created. Instead, custom properties can be added to objects already in the database before a job is started, or created dynamically by procedure steps during step execution.
property sheet	A property sheet is a collection of properties that can be nested to any depth. The property value can be a string or a nested property sheet. Most objects have an associated "property sheet" that contains custom properties created by user scripts.
proxy agent	A proxy agent is an agent on a supported Linux or Windows platform, used to proxy commands to an otherwise unsupported agent platform. Proxy agents have limitations, such as the inability to work with plugins or communicate with ectool commands.
proxy resource	This resource type requires SSH keys for authentication. You can create proxy resources (agents and targets) for ElectricFlow to use on numerous other remote platforms/hosts that exist in your environment.
proxy target	A proxy target is an agent machine on an unsupported platform that can run commands via an SSH server.
publisher	A publisher is the job that completes the <i>publish</i> operation for an artifact version.

Term	Description
quiet time	An inactivity period before starting a build within a continuous integration system. This time period allows developers to make multiple, coordinated check-ins to ensure a build does not start with some of the changes only—assuming all changes are checked-in within the specified inactivity time period. This time period also gives developers an opportunity to "back-out" a change if they realize it is not correct. Using ElectricSentry, the inactivity time period can be configured globally for all projects or individually for a single project.
reports	ElectricFlow provides multiple reports and custom report capabilities to help you manage your build environment.
	Real-time reports - filtered view of your workload in real-time
	Build reports - summary reports produced at the end of a build and attached to the job
	Batch reports - summaries of your build environment with trends over time, two types:
	Default Batch reports - automatically installed during ElectricFlow installation and scheduled to run daily (Cross Project Summary, Variant Trend, Daily Summary, Resource Summary, Resource Detail)
	 Optional Batch reports - you can configure, rename, and schedule these reports to fit your requirements (Category Report, Procedure Usage Report, Count Over Time Report, Multiple Series Reports)
	Custom reports - your choice to create and add at any time

Term	Description
repository or repository server	The artifact repository is a machine where artifact versions are stored in either uncompressed tar archives or compressed tar-gzip archives. The repository server is configured to store artifact versions in a directory referred to as the repository backing store.
	By default, the backing store is the <datadir>/repository-data directory in the repository installation—this default setting can be changed.</datadir>
	A <i>repository</i> is an object that stores artifact versions. This object primarily contains information about how to connect to a particular artifact repository. Similar to steps in a procedure, repository objects are in a user-specified order. When retrieving artifact versions, repositories are queried in this order until one containing the desired artifact version is found.
	Connection information is stored in the repository object on the ElectricFlow server.
resource	A resource specifies an agent machine where job steps can be executed. Resources can be grouped into a "pool", also know as a "resource pool." ElectricFlow supports two types of resources:
	 Standard - specifies a machine running the ElectricFlow agent on one of the supported agent platforms Proxy - requires SSH keys for authentication. You can create proxy resources (agents and targets) for ElectricFlow to use on numerous other remote platforms/hosts that exist in your environment.
schedule	A schedule is an object used to execute procedures automatically in response to system events. For example, a schedule can specify executing a procedure at specific times on specific days. Three types of schedules are available: Standard, Continuous Integration, and Custom (custom schedules are typically continuous integration schedules that do not use the ECSCM plugin).
Sentry schedule	A continuous integration schedule created using the ElectricSentry engine for continuous integration or the CI Continuous Integration Dashboard, which is an easy- to-use front-end user interface for the ElectricSentry engine.

Term	Description
shortcut	One type of shortcut is part of the ElectricFlow Home page facility and records the location of a page you visit frequently (either inside or outside of ElectricFlow), so you can return to that page with a single click from the Home page.
	Another type of shortcut is a context-relative shortcut to property paths. This shortcut can be used to reference a property without knowing the exact name of the object that contains the property. You might think of a shortcut as another part of the property hierarchy. These shortcuts resolve to the correct property path even though its path elements may have changed because a project or procedure was renamed. Shortcuts are particularly useful if you do not know your exact location in the property hierarchical tree.
state	Workflows always have a single active <i>state</i> . Each state in a workflow, when it becomes active, can perform an action. A state can run a procedure to create a subjob or run a workflow definition to create a subworkflow—in the same way that procedures can call other procedures. One or more states can be designated as "starting" states to provide multiple entry points into the workflow. See state definition.
state definition	Workflow objects are split into two types: <i>Definition</i> objects and <i>Instance</i> objects. Definition objects provide the template for a running workflow instance. You create a new workflow by defining a Workflow Definition along with its State Definition and Transition Definition objects.
	When you run the workflow definition, the system creates a new Workflow object with an equivalent set of State and Transition objects that represent the run-time instances of the workflow definition. Note: We omit the "Instance" qualifier for brevity in the API and the UI.
	Each workflow can contain one or more state objects. Defining states for a workflow is analogous to defining steps for a procedure.

Term	Description
step	A step is a procedure component. Each step specifies a command to execute on a particular resource or a subprocedure (nested procedure) to invoke. Commonly created steps include:
	Command - This step invokes a bat, cmd, shell, perl script, or similar.
	Subprocedure - This step invokes another ElectricFlow procedure.
	 Plugin step - These include task-specific steps. Depending on which step-type you choose, the information you need to enter is somewhat different. Some of the step types bundled with ElectricFlow include: Publish or retrieve artifact version
	Send Email
	Various SCM step types
	Build tools
stored credential	Stored credentials are given a name and stored in encrypted form in the database. Each project contains a list of stored credentials it owns. These credentials are managed from the Project Details page.
subprocedure	Creating subprocedures is a way of "nesting" procedures. A step (from any procedure) can call a procedure from another project or the same project. The procedure called by the step then becomes a subprocedure.
substitution	A mechanism used to include property values in step commands and elsewhere. For example, if a step command is specified as "echo \$[status]", and when the step executes there is a property named "status" with value "success", the actual command executed will be "echo success".
system object	This is a special object whose access control lists are used to control access to some ElectricFlow internals. System objects are: admin, artifactVersions, directory, emailConfigs, forceAbort, licensing, log, plugins, priority, projects, repositories, resources, server, session, and workspaces.
tag	A way to categorize a project to identify its relationship to one or more other projects or groups. You can edit a project to add a tag. Enter a tag if you want to categorize or "mark" a project to identify its relationship to one or more other projects or groups.
	For example, you might want to tag a group of projects as "production" or "workflow", or you might want to use your name so you can quickly sort the project list to see only those projects that are useful to you.

Term	Description
transition	Transitions are used to move workflow progress from one state to another state. Four types of transitions are available to move a workflow to the next state:
	On Enter - transitions before sending notifiers or starting the sub-action
	On Start - transitions immediately after starting the sub-action. These transitions are ignored if no sub-action is specified for the source state.
	On Completion - transitions when the sub-action completes. These transitions are ignored if no sub-action is specified for the source state. Note: On Completion transitions are taken only if the state is still active when the sub-action completes, and are ignored if the workflow has transitioned to another state—this can occur if an On Start or Manual transition occurred before the sub-action completed.
	 Manual - transitions when a user selects the transition in the UI and specifies parameters. The same action can occur using ectool or the Perl API by calling transitionWorkflow. Only users who have "execute" permission on the transition are allowed to use the Manual transition. See transition definition.
transition definition	Workflow objects are split into two types: <i>Definition</i> objects and <i>Instance</i> objects. Definition objects provide the template for a running workflow instance. You create a new workflow by defining a Workflow Definition along with its State Definition and Transition Definition objects.
	When you run the workflow definition, the system creates a new Workflow object with an equivalent set of State and Transition objects that represent the run-time instances of the workflow definition.
	Note: We omit the "Instance" qualifier for brevity in the API and the UI.
	Each state can contain one or more transition objects. The transition definition object requires a name for the transition. This transition name will appear on the Workflow Definition Details page for quick reference and also on the State Definition Details page when you select the Transition Definitions tab.
	You can define one or more transitions for each state, depending on which transition options you want to apply to a particular state.
user	A user defines an account used to log into the system and control access to ElectricFlow objects. A user can be defined externally in an LDAP or Active Directory repository, or locally in ElectricFlow. See local user.

Term	Description
workflow	You can use a workflow to design and manage processes at a higher level than individual jobs. For example, workflows allow you to combine procedures into processes to create build-test-deploy lifecycles.
	A workflow contains <i>states</i> and <i>transitions</i> you define to provide complete control over your workflow process. The ElectricFlow Workflow feature allows you to define an unlimited range of large or small lifecycle combinations to meet your needs. See workflow definition.
workflow definition	Workflow objects are split into two types: <i>Definition</i> objects and <i>Instance</i> objects. Definition objects provide the template for a running workflow instance. You create a new workflow by defining a Workflow Definition along with its State Definition and Transition Definition objects.
	When you run the workflow definition, the system creates a new Workflow object with an equivalent set of State and Transition objects that represent the run-time instances of the workflow definition. Note: We omit the "Instance" qualifier for brevity in the API and the UI.
workflow name template	This is the template used to determine the default name of jobs launched from the workflow definition. For example:
	<pre>\$[projectName]_\$[/increment /myproject/workflowCounter]_ \$[timestamp]</pre>
	(substitute your values for the names above)
	Produces a workflow name like: projectName_123_20140102130321

Term	Description
workspace	A workspace is a subtree of files and directories where job file data is stored. The term "workspace" typically refers to the top-level directory in this subtree.
workspace root	A workspace root is a directory in which ElectricFlow allocates job workspace directories. Each workspace root has a logical name used to refer to it in steps and procedures.
zone	A zone is a way to partition a collection of agents to secure them from use by other groups—similar to creating multiple top-level networks. For example, you might choose to create a developers zone, a production zone, and a test zone—agents in one zone cannot directly communicate with agents in another zone. A default zone is created during ElectricFlow installation. The server implicitly belongs to the default zone, which means all agents in this zone can communicate with the server directly (without the use of a gateway).