# ElectricAccelerator
# Electric Make User Guide

**Version 8.2**

**ElectricAccelerator Version 8.2**

# Contents

# Chapter 1: Electric Make Overview

Electric Make® (eMake), which is the main build application in ElectricAccelerator®, is a new Make version invoked interactively or through build scripts. It reads makefiles in several different formats, including GNU Make and Microsoft NMAKE. Electric Make distributes commands to the cluster for remote execution and services file requests.

**Topics:**

- Understanding Component Interactions
- ElectricAccelerator Virtualization
- Understanding Build Parts

# Understanding Component Interactions

To a user, ElectricAccelerator ("Accelerator") might appear identical to other Make versions—reading makefiles in several different formats and producing identical results. Using a cluster for builds is transparent to the Accelerator user.

Important differences in Accelerator build processing versus other distributed systems:

- Components work together to achieve faster, more efficient builds. Instead of running a sequential build on a single processor, Accelerator executes build steps in parallel on a cluster of hosts.

- For fault tolerance, job results are isolated until the job completes. If an Agent fails during a job, Accelerator discards any partial results it might have produced and reruns the job on a different Agent.

- Missing dependencies discovered at runtime are collected in a history file that updates each time a build is invoked. Accelerator uses this collected data to improve performance of subsequent builds.

## Electric Make and EFS

High concurrency levels in Accelerator are enabled by the Electric File System (EFS). When a job such as a compilation runs on a host, it accesses files such as source files and headers through EFS. EFS records detailed file access data for the build and returns that data to Electric Make.

Electric Make acts as a file server for Agents, reading the correct file version from file systems on its machine and passing that information back to the Agents. Agents retain different file version information and do not rely on Electric Make's file sequencing ability to provide the correct version for a job. The Agent receives file data, downloads it into the kernel, notifying EFS, which then completes the original request. At the end of a job, Electric Agent returns any file modifications to Electric Make so it can apply changes to its local file systems.

## Electric Make and Cluster Manager

When Electric Make is invoked on the build machine, it communicates with Cluster Manager to acquire a set of Agents it can use for the build. When Electric Make finishes, it sends Cluster Manager the build results, and tells Cluster Manager that Agents are free now to work on other builds. If more than one build is invoked, Cluster Manager allocates agents using a priority-based algorithm. Builds with the same priority share Agents evenly, while higher priority builds are allocated more Agents than lower priority builds. By default, agents running on the same host machine are allocated to the same build. In real time, Cluster Manager dynamically adjusts the number of agents assigned to running builds as each build's needs change, which allows Accelerator to make the best use of cluster resources.

# ElectricAccelerator Virtualization

ElectricAccelerator is designed to virtualize parts of the build setup so cluster hosts can be configured correctly for the specific build they are executing. Specifically, ElectricAccelerator dynamically mirrors the following host system properties on the Agent:

- Electric File System

- System registry (Windows only)

- User ID (UNIX only)

- Environment variables

## Electric File System (EFS)

Files are the most important resource to distribute to hosts. When Electric Make starts, it is given a directory (or list of directories) called `EMAKE_ROOT`. All files in eMake root(s) are automatically mirrored on cluster hosts for the build duration. This powerful feature means you do not need to configure file sharing between the build system and the cluster—simply specify the `EMAKE_ROOT` and the hosts can access files on the host build system.

Almost any file visible to the host build system can be mirrored to the cluster—regardless of how the host system accesses that file (local disk, network mount, and so on). Both sources and tools can be mirrored, but there are important flexibility/performance "trade-offs" to consider when you include tools in the eMake root. See Configuring Your Build.

When the build completes, EFS is unmounted and the files are no longer visible on the hosts—this ensures builds do not interfere with local host configuration.

Generally, EMAKE_ROOT can be set to include any existing directory on the host build machine. Files in these directories are automatically accessible by commands running on cluster hosts. The exceptions are detailed in Setting the Electric Make Root Directory.

## System Registry (Windows Only)

A side effect of running builds on Windows: Windows might execute tools that modify the system registry. For example, a build job step might install a program that includes registry modifications that are executed by a subsequent step. If job steps are run on different Agents without registry virtualization, the build might fail because registry modifications on one machine are not visible on another.

ElectricAccelerator solves this situation by mirroring the Windows registry in addition to the file system. You can specify a registry root using the command line --emake-reg-roots=<*path*>. Just as EMAKE_ROOT specifies a host file system subset to mirror, the registry root specifies which registry keys should be virtualized to Agents (for example, HKEY_LOCAL_MACHINE\Software\My Product). Access and modifications to keys and values under the registry root on Agents are then sent back to the host build system to ensure correct values are propagated to other Agents.

As with the file system, registry mirroring is active during the build only. After the build completes, the registry returns to its original configuration. Processes not running under an Agent on a host cannot see mirrored files and registry entries—they will see only the usual view of the machine.

## User Accounts

The same user ID is used by the Electric Agent [running on each host, executing processes] and the Electric Make process started on the host build system. Using the same user ID ensures processes have the same permission and obtain the same results from system calls (such as getuid) as they would running serially on the user machine. On Windows, processes are executed by the Electric Agent service user. The Agent user can be configured at installation time, see the *ElectricAccelerator Installation Guide* at http://docs.electric-cloud.com/accelerator_doc/AcceleratorIndex.html.

## Environment Variables

When Electric Make sends commands for execution to the Agent, it also sends environment variables. In most cases, this automatic environment virtualization is sufficient. Typically, variables that specify output locations, target architecture, debug-vs.-optimized, and so on, are propagated to tools on the host and their remote behavior correctly matches what they do locally. By default, environment variables sent by Electric Make override system-wide settings on the host.

For certain variables, override behavior is not desirable. For the build to operate correctly, tools running on the hosts must see the value from the local Agent, not the host build system. For example, the Windows SYSTEMROOT variable depends on the directory where Windows was installed on that machine—values might vary in different Window releases. The SYSTEMROOT value from the host build system is frequently different than the value from the local Agent.

To accommodate these instances, Electric Make allows you to exclude specific environment variables from virtualization. In addition to variables you might explicitly choose to exclude, Electric Make automatically excludes the TMP, TEMP, and TMPDIR variables on all platforms; and the COMSPEC and SYSTEMROOT variables on Windows. For a complete description of environment variables, see Electric Make Command-Line Options, Environment Variables, and Configuration File

# Understanding Build Parts

Software builds are often complex systems with many inputs that must be carefully set up and configured.

Correctly distributing system processes over a cluster of computers or hosts requires efficient replication of relevant parts of the build setup on each host.

ElectricAccelerator software is designed to help virtualize your build environment on each cluster host so distributed jobs display the same behavior as compared to how those jobs run serially on one computer.

Because input configurations can vary from one build to the next, ElectricAccelerator virtualization is active only while the build is running. When the build completes, the host returns to its original state. Consequently, different builds can share a cluster without danger of conflicting changes corrupting host configuration.

# Chapter 2: Setting Up ElectricAccelerator

After all components are installed, and before invoking Electric Make, you need to configure ElectricAccelerator to ensure accurate, reliable builds.

**Topics:**

- Defining Your Build

- Configuring Your Build

# Defining Your Build

Before using Accelerator, precisely define which components go into your build. Generally, a build has three input types:

**SOURCES + TOOLS + ENVIRONMENT = BUILD**

## Build Sources

Build sources include all compiled or packaged files to build your product. For example, these sources include:

- Intermediate files generated during the build (for example, headers or IDL files)

- Makefiles and scripts that control the build

- Third-party source files read during the build

- Any file read by the build in source control

## Build Tools

Tools used to create build output include make, compilers, linkers, and anything else operating on the sources during the build:

- Executable post-process tools (for example, *strip*)

- Static analyzers (for example, *lint*) if they run as part of the build

- Packaging tools (for example, *tar* or *zip*)

Sometimes the distinction between sources and tools is blurred. Consider these examples:

- A utility executable compiled from your sources during the build, run during later steps, but not part of final output

- Header source files that are part of the compiler (for example, `stdio.h)` or a third-party package, but not under source control

In this context, *sources* are those files that might change from one build to the next. Thus, a utility executable compiled from your sources as part of the build is considered a source.

By contrast, *tools* change infrequently—tools are often configured once and served from a central location (for example, an NFS share). A standard header such as `stdio.h` is usually considered part of the tool suite.

The distinction between inputs that can change between builds (sources) and inputs that can safely be assumed to be constant (tools) becomes important when configuring Accelerator virtualization.

## Build Environment

Your operating system environment is an essential part of your build. The operating system is easy to overlook because the environment usually is configured once per host and then ignored as long as builds function normally. But when a build is distributed across a shared cluster, it is important to identify which parts of the operating system could affect the build. Some inputs to consider include:

- User environment variables

- System registry (Windows only)

- Operating system version (including patches or service packs)

- User account and user permissions

- Host-specific attributes (for example, machine name, network configuration)

For each of these inputs, consider what impact (if any) they will have on the build.

Generally, some environment variables require correct settings for a build (for example, PATH or variables that specify output architecture or source/output file locations).

Another common example of how the operating system can affect the build occurs with the use of tools that require license management. If a tool license (for example, a compiler) is host-locked or it requires contacting a license server to operate, ensure the compiler on the host can acquire the license also (for example, cluster host names could be configured as valid license server clients).

# Configuring Your Build

After defining your build environment and identifying all of its inputs, configure the cluster and Electric Make so the system correctly virtualizes:

1. Set the Electric Make root directory.

2. Determine if you need to do additional configuration for tools.

3. Make additional configurations regarding the registry. (Windows only)

4. Configure environment variables if needed.

5. Set the Cluster Manager host and port.

6. Set Electric Make emulation.

7. Enable local agents (if applicable).

Some cases of virtualization and/or distribution of specific job steps are not desirable. For these cases you can configure Accelerator to:

- Run an individual command from the Agent back on the host system, using the "proxy command" function. See Using the Proxy Command.

- Prevent remote job execution by using the #pragma runlocal function. See Running a Local Job on the Make Machine.

# Chapter 3: Electric Make Basics

The following topics discuss basic Electric Make information to get you up and running.

**Topics:**

- Invoking eMake
- Setting the Electric Make Root Directory
- Configuring Tools
- Tools that Access or Modify the System Registry
- Configuring Environment Variables
- Setting the Cluster Manager Host and Port
- Setting Electric Make Emulation
- Electric Make Command-line Options and Environment Variables

# Invoking eMake

The Electric Make executable is called `emake`. The most important change to your build process is to ensure this executable is invoked in place of the existing Make.

For interactive command-line use, ensure the following:

- The ElectricAccelerator `bin` directory is in your `PATH` environment variable:

  - For Linux: `/opt/ecloud/i686_Linux/bin` or `/opt/ecloud/i686_Linux/64/bin`

  - For Solaris: `/opt/ecloud/sun4u_SunOS/bin`

  - For Windows: `C:\ecloud\i686_win32\bin` or `C:\ecloud\i686_win32\64\bin`

- You type `emake` in place of gmake or nmake.

**Note:** The 64-bit version of eMake is recommended for builds with very large memory requirements.

You can rename the eMake executable to either gmake or nmake because Electric Make checks the executable to determine which emulation to use. If the name of the submake is hard-coded in many places within your makefiles, a simple solution would be to rename gmake or nmake to *gmake.old* or *nmake.old*, and rename eMake to either *gmake* or *nmake* on cluster hosts only. In this way, you can maintain access to your existing make, but all submakes from an Accelerator build will correctly use Electric Make.

**Note:** Electric Cloud does *not* recommend running builds in `/tmp`.

To ensure eMake is called for recursive submake invocations in makefiles, use the `$(MAKE)` macro for specifying submakes instead of hard-coding references to the tool. For example, instead of using:

```
libs:
        make -C lib
```

use the following `$(MAKE)` macro:

```
libs:
        $(MAKE) -C lib
```

## Single Make Invocation

It is important to keep the build in a single Make invocation. At many sites, Make is not directly invoked to do a build. Instead, a wrapper script or harness is used to invoke Make, and users (or other scripts) invoke this wrapper. The wrapper script might take its own arguments and might perform both special set up or tear down (checking out sources, setting environment variables, post-processing errors, and so on). Because Electric Make behaves almost exactly like native Make tools, usually it can directly replace the existing makefile in wrapper scripts.

Sometimes, however, the script might invoke more than one Make instance. For example, the script could iterate over project subdirectories or build different product variants. In this case, each of these builds becomes a separate Accelerator build, with its own build ID, Cluster Manager entry, history file, and so on.

It is much more efficient for Make instances that are logically part of one build to be grouped under the control of a single parent Make invocation. In this way, Electric Make can track dependencies between submakes, ensure maximal parallelization and file caching, and manage the build as a single, cohesive unit.

If your build script invokes more than one submake, consider reorganizing makefile targets so a single Make is invoked that in turn calls Make recursively for submakes.

If a lot of the setup for each instance occurs within the build script, another possible solution is to use a simple top level makefile to wrap the build script; for example,

```
all:
        my-build-harness ...
```

In this instance, `my-build-harness` runs on the Agent much like any other command and sends commands discovered by submake stubs back to the host build machine. This approach works only if each submake's

output is not directly read by the script between Make invocations. Otherwise, it might be susceptible to submake stub output problems. See Invoking eMake  for more information.

# Setting the Electric Make Root Directory

The `--emake-root` option (or the `EMAKE_ROOT` environment variable) specifies the EFS root directory [or directories] location. All files under the Electric Make root directory are automatically mirrored on each Agent.

Electric Make uses the current directory as the default if no other root directory is specified. You must specify the correct root directory (or directories) or the build might fail because Electric Make cannot find the necessary files to complete the build or resolve dependencies.

For best results and performance, be specific when setting the Electric Make root location. Be sure to include:

- All files created or modified by the build.

- All source files.

- The location where build output files will go during the build, for example, object files, linker output, and so on.

- Other files read by the build such as third-party tools and system headers, or other files not modified if you need to. Be aware, however, that including these files can slow performance. See Configuring Your Build.

   If necessary, specify more than one directory or subdirectory. Separate each location using standard `PATH` variable syntax (a colon for UNIX, a semicolon for Windows).

   UNIX example:

   ```
   --emake-root=/src/foo:/src/baz
   ```

   In this example, you have streamlined the root path by excluding other `/src` subdirectories not necessary to the build.

   Windows example:

   ```
   --emake-root=C:\Build2;C:\Build4_test
   ```

**Note:** Any files used by the build, but not included under an Electric Make root directory, must be preloaded onto all hosts and identical to corresponding files on the system running Electric Make. If these files are not identical, Electric Make could find and use the wrong files for the build. This approach is appropriate for system compilers, libraries, header files, and other files that change infrequently and are used for all builds.

Generally, `EMAKE_ROOT` can be set to include any existing directory on the host build machine. Files in these directories are automatically accessible by commands running on cluster hosts. However, there are a few exceptions:

- `EMAKE_ROOT` cannot be set to the system root directory (for example, "/" on UNIX or `C:/` on Windows). It might be tempting to try this to specify "mirror everything," but in practice, this is not desirable because mirroring system directories such as `/var` or `/etc` on UNIX or `C:/Windows` on Windows can lead to unpredictable behavior. eMake will not allow you to specify the root directory as `EMAKE_ROOT`.

- `/tmp` and the Windows temp directory cannot be included in the eMake root.

- On Windows, another operating system restriction is imposed:
  `EMAKE_ROOT` is not a UNC path specification—it must be a drive letter specification or a path relative to a drive letter. It must also be a minimum of three characters.

# Configuring Tools

Cluster hosts must be able to execute all tools (compilers, linkers, and so on) required during the build. Three execution setup choices:

- **Tools are available on the cluster in the same locations as on the host build system**—this is the simplest setup and requires no special Electric Make or build system configuration.

  This is the most common setup if tools are provided from a network mount (for example, an NFS or SMB server). If tools are not supplied from a network share, but are installed on the local disk on each host, those tools must be updated whenever tools on the build system change.

- **Tools are available on the cluster, but not necessarily in the same locations as the host build system**—this configuration is common in environments where users are allowed to install tools on their host system in non-standard locations.

  In this case, care must be taken to ensure build steps executing on the hosts find the tools. Usually, this means making the PATH environment variable include standard locations used on the cluster hosts—any other variables referencing tool locations need to be modified similarly.

  As in the first case, if tools are installed on the cluster hosts' local disk, they must be updated when tools are updated on the host system.

- **Tools are not installed on the cluster**—in this case, when no tools are installed on the cluster, EMAKE_ ROOT is configured to include the tools directory so the compiler and linker are sent to the Agents by Electric Make [in the same manner as the source files].

  This is the most flexible setup because it eliminates the need to install or update tools on the hosts. If tools change frequently or vary between branches, this setup can dramatically reduce setup cost. However, to gain this flexibility, you typically trade performance because network overhead is increased.

  Another factor to consider is the possibility that configuration information for a virtualized tool might not work on the host. For example, it is not possible to virtualize the Microsoft Visual Studio IDE because of its tight integration with the local machine's registry.

In some special cases, it might be necessary to use tools outside of the virtualized file system. To support running such tools on the local build machine during a cluster build, see Using the Proxy Command.

> **IMPORTANT:** Electric Cloud does not recommend Proxy Command for frequent use; use it with caution.

# Tools that Access or Modify the System Registry

In addition to files, tools on Windows might access or modify the system registry. Use the --emake-reg-roots command-line option to specify a key to mirror. You can specify more than one key by separating multiple entries with semicolons:

```
--emake-reg-roots=HKEY_LOCAL_MACHINE\Software\Foo;
HKEY_LOCAL_MACHINE\Software\Bar
```

In addition, you can specify exception keys to not mirror the system registry by prefixing the key with a "–" character. For example:

```
--emake-reg-roots=HKEY_LOCAL_MACHINE\Software\Foo;
-HKEY_LOCAL_MACHINE\Software\Foo\Base
```

which means "mirror all keys and values under Foo except the keys in Foo\Base.

To ensure compatibility with Microsoft Visual Studio, the registry root specification automatically includes:

```
HKEY_CLASSES_ROOT;-HKEY_CLASSES_ROOT\Installer;
-HKEY_CLASSES_ROOT\Licenses
```

# Configuring Environment Variables

If your build environment is fairly constant, you might want to use *environment variables* to avoid respecifying values for each build. Environment variables function the same way as command-line options, but command-

line options take precedence over environment variables.

Because Electric Make environment variables are propagated automatically to Agents, most commands running on an Agent will run with the correct environment without modifications.

However, two important exceptions exist.

- As described in Configuring Tools, if tools are installed on the cluster host in different locations from the build system, the PATH variable (and other variables that reference tool locations) on the build system must be modified to include the locations for the cluster tools.

- Generally, differences between the build system and cluster hosts might indicate it is undesirable to override environment variables with eMake values. In this case, use the --emake-exclude-env command-line option or the EMAKE_EXCLUDE_ENV environment variable.

  For example, consider a build system environment variable called LICENSE_SERVER that normally contains the license server name that the system should contact to obtain a tools license. If this variable is machine-specific, eMake overrides the correct machine-specific value on the cluster hosts with the value from the build system. To ensure eMake does not override LICENSE_SERVER with the value from the build system, use the option, --emake-exclude-env, when running eMake:

  ```
  --emake-exclude-env=LICENSE_SERVER
  ```

  You can specify more than one value by separating them with commas:

  ```
  --emake-exclude-env=LICENSE_SERVER,TOOLS_SERVER
  ```

  Some variables are almost always used to describe the local machine state, so Electric Make always excludes them from mirroring. These variables are:

  ```
  TMP
  TEMP
  TMPDIR
  ```

  For Windows, this list also includes:

  ```
  COMSPEC
  SYSTEMROOT
  ```

### Configuring ccache

The only configuration required to use ccache with Accelerator is to set the CCACHE_NOSTATS environment variable. If you do not set this environment variable, the entire build becomes serialized because ccache continuously writes to a statistics file throughout the build. To learn more about ccache, refer to http://ccache.samba.org/.

# Setting the Cluster Manager Host and Port

The --emake-cm option (or the EMAKE_CM environment variable) sets the Cluster Manager host name (either IP address or machine name) and must be specified. If the host name is not specified by either of those methods, Electric Make operates in local mode and performs like a traditional make program by running jobs on the local machine in serial order—without using the cluster hosts.

### Cluster Build Example

The Cluster Manager host name is linuxbuilder.

- You can invoke a build against the cluster by running:

  ```
  % emake --emake-cm=linuxbuilder
  ```

- Alternatively, use the EMAKE_CM environment variable:

```
% setenv EMAKE_CM linuxbuilder
% emake
```

### *Local Machine Build Example*

- To force a local build (one *not* run on the cluster), leave the Cluster Manager host name undefined by running:

```
% emake --emake-cm=
```

- Or, leave the EMAKE_CM environment variable is undefined:

```
% emake
```

# Setting Electric Make Emulation

Electric Make can emulate different make variants: GNU Make (gmake), gmake in a Cygwin environment, and Microsoft NMAKE (nmake).

By default, the --emake-emulation=<*mode*> is set to gmake, which supports of a subset of GNU Make 3.81 (see Unsupported GNU Make Options and Features).

Available modes:

- gmake

- gmake3.82

- gmake3.81

- gmake3.80

- nmake

- nmake8

- nmake7

- cygwin

- symbian

**Note:** You can rename emake.exe to nmake.exe or gmake.exe to change the emulation type for all builds automatically. See the --emake-emulation-table option in Electric Make Command-Line Options, Environment Variables, and Configuration File.

### *Windows*

If you are using nmake, override the default setting from gmake to nmake:

```
emake --emake-emulation=nmake
```

This is a Windows-specific setting that must be set prior to invoking a build.

### *Cygwin*

If you are using Cygwin, override the default setting from gmake to cygwin:

```
emake --emake-emulation=cygwin.
```

# Electric Make Command-Line Options, Environment Variables, and Configuration File

You can configure Electric Make options from the command line for a specific build and can also use Electric Make environment variables or the emake.conf configuration file to make options persistent.

Following are caveats for using these methods for setting options:

- The environment variable `EMAKEFLAGS` can be used to set any command-line option. For example, this emake invocation:

```
% emake --emake-cm=mycm --emake-root=/home/joe
```

is equivalent to the following in `csh`:

```
% setenv EMAKEFLAGS "--emake-cm=mycm --emake-root=/home/joe"
% emake
```

The `bash` equivalent is:

```
$ export "EMAKEFLAGS=--emake-cm=mycm --emake-root=/home/joe"
$ emake
```

The Windows command shell equivalent is:

```
C:\> set EMAKEFLAGS=--emake-cm=mycm --emake-root=C:\home\joe
C:\> emake
```

- The hierarchy or precedence for setting an Electric Make option is:

    - Command-line options

    - Options set using the `EMAKEFLAGS` environment variable

    - Options set using other Electric Cloud environment variables

  Command-line options as well as options in the `EMAKEFLAGS` environment variable override options in the `emake.conf` configuration file.

## Editing the Electric Make Configuration File

The path to the Electric Make configuration file is `<ECloud install>/<arch>/conf/emake.conf` (on Linux and UNIX) or `<ECloud install>\<arch>\conf\emake.conf` (on Windows), where `<ECloud install>` is the top installation directory containing the Electric Make program, and `<arch>` is the architecture (such as `i686_Linux` or `sun4u_SunOS`). If you want to use the Electric Make configuration file but it does not exist, you must create it.

If you run an Electric Make executable that is not within an ElectricAccelerator installation, it will not use the installed configuration file (but will still use environment variables and options entered from the command line).

The configuration file has the following formatting rules:

- Only Electric Make-specific command line arguments are allowed (not those specific to GNU Make, NMake, and so on).

- One option (a switch-argument pair) is allowed per line.

- All allowed switches must be separated from their arguments by = (an equal sign).

- All text between = and the end of the line (including spaces and terminal spaces) is part of the argument.

- File order determines the option order (for many kinds of options, the final occurrence overrides all earlier ones).

- `NUL` characters are not allowed.

- Empty lines are ignored.

- Lines beginning with # are comment lines and are ignored.

Configuration file syntax errors (but not necessarily semantic ones) are warnings and cause the configuration file to be skipped (as if it was not present, except for the warnings).

# List of Command-Line Options

Command-line options are listed in alphabetical order except for platform-specific options that are listed *after* platform-independent options. Debug options are listed at the end of the table.

**Note:** The `--emake-volatile` command-line option is deprecated and no longer has any effect. If the option is specified, it is ignored.

| Command-line Options | Environment Variables | Description |
|---|---|---|
| | `EMAKE_BUILD_MODE` | Always set to local.<br><br>Specifies that an individual emake invocation on the Agent does not enter stub mode, but instead behaves like a local (non-cluster) make. Electric Make automatically uses local mode when the `-n` switch is specified. |
| `--emake-annodetail= var1[,var2[,...]]` | | Specifies the level of detail to include in annotation output—a comma separated list for any of the following values:<br>`basic`: Basic annotation (enabled by default in cluster mode). If the JobCache feature is enabled, basic annotation includes information about cache hits and misses. JobCache is an add-on option and requires a separately-purchased license. The license can be used for Linux, Solaris, and Windows.<br><br>`env`: Enhanced environment variables<br>`file`: Files read or written<br>`history`: Serialization details<br>`lookup`: All file names accessed<br>`md5`: MD5 checksums for reads/writes<br>`registry`: Updates to registry<br>`waiting`: Jobs that waited<br>`none`: Disables all annotation. Note that this value disables basic annotation, even when the `--emake-annofile` option is used.<br><br>This option overrides the build class annotation settings set on the Cluster Manager. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-annofile=` `<file>` | | Specifies the name of the XML-formatted log output file. By default, the annotation file, `emake.xml`, is created in the directory where eMake is run. If specified, implies at least "basic" annotation details. <br><br> The following macros are available: <br><br> `@ECLOUD_BUILD_ID@` expands into the unique eMake build ID. <br><br> `@ECLOUD_BUILD_TAG@` expands into the build tag from the Cluster Manager (this is displayed in the Name column on the Build tab in the Cluster Manager UI). <br><br> `@ECLOUD_BUILD_DATE@` expands into an 8-digit code that represents the local system date where the build began, in the form YYYYMMDD. <br><br> `@ECLOUD_BUILD_TIME@` expands into a 6-digit code that represents the 24-hour local system time where the build began, in the form HHMMSS. <br><br> Example: <br> `--emake-annofile=annofile-@ECLOUD_BUILD_` `ID@-@ECLOUD_BUILD_TAG@-@ECLOUD_BUILD_DATE@-` `@ECLOUD_BUILD_TIME@.xml` <br><br> results in: <br> `annofile-4-default_4_20090220184128-` `20090220-184128.xml` |
| `--emake-annoupload=` `<0|1>` | | Enables (`1`) or disables (`0`=default) annotation file upload on the Cluster Manager. <br><br> This option overrides the build class annotation upload setting set on the Cluster Manager. |
| `--emake-autodepend=` `<0|1>` | | Enables (`1`) or disables (`0`=default) the eDepend feature. |
| `--emake-big-file-` `size=<N>` | | Sets the minimum file size (in bytes) to send through Agent-to-Agent transfers for direct file sharing between hosts. Default=`10KB`. |
| `--emake-build-label=` `<label>` | `ECLOUD_BUILD_` `LABEL` | Sets a customized build label. These labels are literal strings and do not use available tags when defining labels for build classes. |
| `--emake-class=` `<class>` | `ECLOUD_BUILD_` `CLASS` | Specifies the build class for the current build. The class must match an existing Cluster Manager class previously created by a user. If this option is not used, or if the class does not match, Electric Make assigns the default class to the build. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-clearcase=`<br>`var1[,var2[,...]]` | `EMAKE_`<br>`CLEARCASE` | Turns on support for specified ClearCase features—a comma separated list of any of the following values:<br>`symlink` : symbolic links<br>`vobs` : per-VOB caching (for speed)<br>`rofs` : read-only file system |
| `--emake-cluster-`<br>`timeout=<N>` | | If no Agents are available when the build starts, this option specifies the number of seconds to try to acquire Agents before giving up. The default is `-1` (infinite), so the build waits indefinitely for Agents. |
| `--emake-cm=<host>`<br>`[:port]` | `EMAKE_CM` | Sets the host name (either the IP address or machine name) and port for Cluster Manager. |
| `--emake-collapse=`<br>`<0\|1>` | | Turns history collapsing on or off. When collapsing is enabled, dependencies between a single or several jobs in another *make* instance are replaced with a serialization between the job and the other Make instance. This action typically results in significant history file size reduction, but might cause some over-serialization. In most builds, this has little or no impact on build time. In some builds, disabling collapsing improves performance at the cost of increased history file size. Default=1 (on) |
| `--emake-disable-`<br>`pragma=var1[,var2`<br>`[,...]]` | | Comma separated list of pragma directives to ignore—can be one or more of: `allserial`, `runlocal`, `noautodep`, or `all` to disable all pragmas. |
| `--emake-disable-`<br>`variable-pruning=`<br>`<0\|1>` | | Disables variable table pruning. Default=`0` (off) |
| `--emake-emulation=`<br>`<mode>` | `EMAKE_`<br>`EMULATION` | Sets Make-type emulation to mode. Default emulation type is `gmake`. You can rename `emake.exe` to `nmake.exe` or `gmake.exe` to change the emulation type for all builds automatically.<br>Available modes: `gmake`, `gmake3.80`, `gmake3.81`, `gmake3.82`, `symbian`, `nmake`, `nmake7`, `nmake8`, or `cygwin`.<br>If specifying `gmake3.82`, read the GNU Make 3.82 Note. |
| `--emake-emulation-`<br>`table=<table>` | | Configures default emulation modes for Make programs. TABLE is a comma separated list of NAME=MODE, where NAME is the name of a Make executable and MODE is the emulation mode to use if `emake` is invoked as NAME. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-exclude-env=`<br>`var1[,var2[,...]]` | `EMAKE_EXCLUDE_`<br>`ENV` | Specifies which environment variables must not be replicated to the hosts. |
| `--emake-hide-`<br>`warning=<list>` | `EMAKE_HIDE_`<br>`WARNING` | Hides one or more Accelerator-generated warning numbers. `list` is a comma-separated list of numbers that you want to hide. The `EC` prefix of a warning number is optional and can be omitted. |
| `--emake-history=`<br>`<read\|create\|merge>` | | Specifies the history mode creation model. Default=`merge`. |
| `--emake-history-`<br>`force=<0\|1>` | | Honors history mode even if the build fails. Default=`1` (on) |
| `--emake-historyfile=`<br>`<path/file>` | | Specifies which history file to use for a specific build. Allows you to change the default name and path for the history file `emake.data` set automatically by Electric Make.<br><br>The following macros are available:<br><br>`@ECLOUD_BUILD_ID@` expands into the unique eMake build ID.<br><br>`@ECLOUD_BUILD_TAG@` expands into the build tag from the Cluster Manager (this is displayed in the Name column on the Build tab in the Cluster Manager UI).<br><br>`@ECLOUD_BUILD_DATE@` expands into an 8-digit code that represents the local system date where the build began, in the form YYYYMMDD.<br><br>`@ECLOUD_BUILD_TIME@` expands into a 6-digit code that represents the 24-hour local system time where the build began, in the form HHMMSS.<br><br>Example:<br>`--emake-historyfile=historyfile-@ECLOUD_`<br>`BUILD_ID@-@ECLOUD_BUILD_TAG@-@ECLOUD_BUILD_`<br>`DATE@-@ECLOUD_BUILD_TIME@.xml`<br><br>results in:<br>`historyfile-4-default_4_20090220184128-`<br>`20090220-184128.xml` |
| `--emake-idle-time=`<br>`<N>` | | Sets the number of seconds (`N`) before idle Agents are released to the cluster. Default=`10` |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-ignore-all-intermediate=`<br>`<0\|1>` | | Causes eMake to not treat .SECONDARY targets having no prerequisites as meaning that all targets are intermediate. This option might increase performance on certain large builds with thousands of targets and reduce eMake runtime memory requirements (but breaks strict compatibility with GNU Make when emulating GNU Make 3.81 and later). |
| `--emake-impersonate-user=<name>` | `ECLOUD_ IMPERSONATE_ USER` | Run the build as this Cluster Manager user.<br>**Note:** Impersonate changes the user recorded by the Cluster Manager and that user's permissions. This option does not affect OS user permissions.<br>For additional information, see the online help topic, "Permissions." |
| `--emake-job-limit=`<br>`<N>` | | Limits the maximum number of uncommitted jobs to `N` where `0` means unlimited. Default=`0` |
| `--emake-ledger=`<br>`<valuelist>` | `EMAKE_LEDGER` | Enables the ledger capability. `Valuelist` is a comma-separated list that includes one or more of: `timestamp`, `size`, `command`, `nobackup`, `nonlocal`, and `unknown`. |
| `--emake-ledgerfile=`<br>`<path/file>` | `EMAKE_ LEDGERFILE` | The name of the ledger file. Default=`emake.ledger` |
| `--emake-maxagents=`<br>`<N>` | | Limits the number of Agents used by this build. `N=0`uses all available Agents. Default=`0`<br>This option overrides the build class Max Agents setting if `--emake-maxagents` is set lower than the build class setting. This option does *not* override the build class setting if `--emake-maxagents` is set higher than the build class setting.<br>This behavior prevents a user from overriding the upper limit for a class that was set up by the build administrator. |
| `--emake-mem-limit=`<br>`<N>` | | Controls the amount of memory Electric Make devotes to uncommitted jobs. When the limit is exceeded, Electric Make stops parsing new Make instances. Default=`1,000,000,000` (1 GB). |
| `--emake-mergestreams=<0/1>` | `EMAKE_MERGE_ STREAMS` | Indicates whether to merge the `stdout/stderr` output streams, yes (`1`) or no (`0`). The default is merge the streams (`1`). For most situations, this is the correct value. If you re-direct standard output and standard error separately, specify no (`0`) for this option. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-monitor <hostname/IP>:<port>` | | Sets the hostname/IP and port of the system from where you want to view the live monitor data in the Electric Cloud ElectricInsight® tool.<br><br>To monitor live build data, you must launch the ElectricInsight live monitor *before* you start the build. |
| `--emake-pedantic= <0|1>` | | Turns *pedantic mode* on (`1`) or off (`0`=default). When pedantic mode is on, warnings appear when invalid switches are used, or potential problems are identified (for example, rules with no targets or reading from a variable that was not written). When pedantic mode is off, Electric Make ignores irrelevant switches or exits without warning if it encounters unresolvable errors. |
| `--emake-priority= <low|normal>` | | Sets the priority for the current build's use of Agents in the cluster. When set to `normal`, the build uses at least the minimum number of Agents set in Cluster Manager. The default setting is determined by the Cluster Manager.<br><br>This option overrides the priority associated with a build class but can be used to lower the build priority only . |
| `--emake-read-only= <path>` | `EMAKE_READ_ ONLY` | All paths starting at the directories specified in `--emake-read-only` will be marked as read-only file systems when they are accessed on the agent. On UNIX, any attempt to create new files or write to existing files under those directories will fail with EROFS, "Read-only file system". On Windows, it will fail with ERROR_ACCESS_DENIED, "Access is denied". |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-readdir-conflicts=<0\|1>` | | Explicitly enables conflict detection on directory read operations (commonly called "glob conflicts," which is but one manifestation of the problem). Allowed values are `0` (disabled, the default value) and `1` (enabled). |
| | | If your build is susceptible to readdir conflict failures, you can enable these checks and get a correct build even if you do not conduct a single-agent build. The resulting history file is identical to a single-agent build result. Though the initial run with this feature might be over-serialized (a consequence of readdir conflicts), a good history file allows builds to go full speed, without conflicts, the next time. |
| | | You do *not* want to enable this option all the time. Correct usage: Enable it for one run if you suspect a globbing problem, and then disable it, but use the history file generated by the previous run. |
| | | Another possible strategy to use if you are not familiar with the build you are building, is to enable the option until you get a successful build, and then disable it after you have a complete, good history file. |
| `--emake-remake-limit=<N>` | | This option defaults to `10`. If set to `0`, makefiles are not added as goals at all and no remaking occurs. Setting the value to `1` is equivalent to the deprecated `--emake-multiemake=0`. |
| `--emake-resource=<resource>` | `EMAKE_RESOURCE` | The resource requirement for this build. |
| | | This option overrides the build class resource request setting set on the Cluster Manager. |
| `--emake-root=<path>` | `EMAKE_ROOT` | Specifies the Electric Make root directory(s) location. |
| | | Particularly on Windows, this parameter should *not* be used to virtualize your tool chain. Tools should be installed on each agent host for performance reasons and to avoid having to virtualize registry parts. |
| | | The semi-colon is the delimiter between drives. |
| | | Example: |
| | | `build@winbuild-cm$ emake --emake-cm=winbuild-cm --emake-emulation=cygwin --emake-root=/c/cygwin/tmp;/c/tmp`<br><br>`Starting build: 867`<br>`make: Nothing to be done for `foo'.`<br>`Finished build: 867 Duration: 0:00 (m:s)`<br>`Cluster availability: 100%` |
| | | In this example, the C: drive is mounted on /c. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-showinfo=` `<0\|1>` | | Turns build information reporting on (`1`=default) or off (`0`). Information includes build time to completion and average cluster utilization. |
| `--emake-tmpdir=` `<path>` | `EMAKE_TMPDIR` | Sets the Electric Make file temporary directory. |
| **Caching Commands (Dependency Optimization, JobCache, Parse Avoidance, and Schedule Optimization)** | | |
| `--emake-assetdir=` `<path>` | | Use the specified directory for cache locations for saved dependencies (for dependency optimization), JobCache, parse avoidance, and schedule optimization. The default name of this directory is .emake. By default, this directory is in the working directory in which eMake is invoked. JobCache is an add-on option and requires a separately-purchased license. |
| `--emake-jobcache=gcc` | | Enables the JobCache feature for all make invocations in a build. This option works for recursive and nonrecursive builds. JobCache is an add-on option and requires a separately-purchased license. |
| `--emake-optimize-` `deps=` `<0\|1>` | | Use the saved dependency information file for a makefile when dependencies are the same and save new dependency information when appropriate. |
| `--emake-optimize-` `schedule=` `<0\|1>` | | Turns the schedule optimization feature on (`1`=default) or off (`0`). This feature uses performance and dependency information from previous builds to optimize the runtime order of jobs in subsequent builds. |
| `--emake-parse-` `avoidance=` `<0\|1>` | | Avoid parsing makefiles when prior parse results remain up-to-date and cache new parse results when appropriate. |
| `--emake-parse-` `avoidance-ignore-` `env=<var>` | | Ignore the named environment variable when searching for applicable cached parse results. To ignore more than one variable, use this option multiple times. |
| `--emake-parse-` `avoidance-ignore-` `path=<path>` | | Ignore this file or directory when checking whether cached parse results are up-to-date. Append % for prefix matching. To ignore more than one path or prefix, use this option multiple times. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-suppress-include=<pattern>` | | Skip matching makefile includes (such as generated dependencies). Generally, you should not suppress makefile includes unless they are generated dependency files, and you have enabled automatic dependencies as an alternative way of handling dependencies. |
| | | **Note:** If the pattern does not have a directory separator, then the pattern is compared to the include's filename component only. If the pattern has a directory separator, then the pattern is taken relative to the same working directory that applies to the include directive and compared to the included file's entire path. |
| **UNIX-Specific Commands** | | |
| `--emake-crossmake=<linux\|solaris\|solarisx86>` | | Use this option to choose a specific operating system, so Cluster Manager will use Linux or Solaris hosts, specifically. |
| | `ECLOUD_ICONV_LOCALE` | Allows you to set the iconv locale. Use `iconv -l` to list the available locales. |
| | | **Usage Note:** If you receive an emake assertion failure that contains information similar to: |
| | | `emake: ../util/StringUtilities.h:406: std::string to_utf8(const std::string&): Assertion ` + "`c != iconv_t(-1)' failed."` |
| | | This could mean that your system is missing an internationalization package, or the locale package has a different name for ISO 8859-1. (This particular issue is due to the conversion between 8-bit strings and UTF-8.) |
| | | Example: |
| | | Your system recognizes "ISO8859-1", "ISO_8859-1", and "ISO-8859-1" only. Set `ECLOUD_ICONV_LOCALE` to one of those valid locale names. |
| **Windows-Specific Commands** | | |
| `--emake-case-sensitive=<0\|1>` | | Sets case sensitivity for target and pattern name matching. This is inherited by all submakes in the build. The option applies when using gmake/cygwin emulation modes only; nmake and symbian modes are not affected by this option. |
| | | Default for gmake and cygwin is on (`1`). |
| | | This option replaces `--case-sensitive`. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-cygwin=` `<Y\|N\|A>` | `EMAKE_CYGWIN` | `Y`=requires *cygwin1.dll* <br> `N`=ignore *cygwin1.dll* <br> `A`=use *cygwin1.dll* if available <br> Default=`A`, if launched from a Cygwin shell, if not, then `N`. <br> Default=`Y`, if eMake `emulation=cygwin` was set. |
| `--emake-ignore-` `cygwin-mounts=` `<mounts>` | `EMAKE_IGNORE_` `CYGWIN_MOUNTS` | Comma separated list of Cygwin mounts to ignore. Unless listed, Cygwin mount points are replicated on the Agent. |
| `--emake-reg-limit=` `<N>` | | Limits the number of registry keys sent automatically for each key. Default=`50`. |
| `--emake-reg-roots=` `<path>` | | Sets the registry virtualization path on Windows machines. The syntax is `--emake-reg-roots=path` `[;path]`. <br><br> Do not use this parameter to virtualize your tool chain. Tools must be installed on each agent host for performance reasons and to avoid figuring out which registry parts to virtualize. |
| **Debug Commands** | | |
| `--emake-debug=` `<value>` | `EMAKE_DEBUG` | Sets the local debug log level(s). For a list of possible values, see the `emake --help` message. |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-logfile=`<br>`<file>` | `EMAKE_LOGFILE` | Sets the debug log file name.<br>Default is `stderr`.<br>The following macros are available:<br>`@ECLOUD_BUILD_ID@` expands into the unique eMake build ID.<br>`@ECLOUD_BUILD_TAG@` expands into the build tag from the Cluster Manager (this is displayed in the Name column on the Build tab in the Cluster Manager UI).<br>`@ECLOUD_BUILD_DATE@` expands into an 8-digit code that represents the local system date where the build began, in the form YYYYMMDD.<br>`@ECLOUD_BUILD_TIME@` expands into a 6-digit code that represents the 24-hour local system time where the build began, in the form HHMMSS.<br>Example:<br>`--emake-logfile=logfile-@ECLOUD_BUILD_ID@-`<br>`@ECLOUD_BUILD_TAG@-@ECLOUD_BUILD_DATE@-`<br>`@ECLOUD_BUILD_TIME@.xml`<br>results in:<br>`logfile-4-default_4_20090220184128-`<br>`20090220-184128.xml` |
| `--emake-rdebug=`<br>`<value>` | `EMAKE_RDEBUG` | Sets the remote debug log level(s). For a list of possible values, see the `emake --help` message.<br>Enabling this option disables parse avoidance. |
| `--emake-rlogdir=`<br>`<dir>` | `EMAKE_RLOGDIR` | Sets the directory for remote debug logs. |
| **Local Agent Commands** | | |

| Command-line Options | Environment Variables | Description |
|---|---|---|
| `--emake-localagents=<y\|n>` | | Instructs eMake to use any available local agents (provided that a local-agent specific license is available on the eMake host). The default is:<br><br>• `y` if you specified a Cluster Manager *or* if agents appear to be running<br><br>• `n` if you did not request a Cluster Manager *and* it appears that local agents are unavailable<br><br>If you have installed local agents (and they are currently running), but you do not want to use them, then specify `--emake-localagents=n` (with or without the `--emake-cm` option). This lets you use cluster agents only (without needing to shut down running local agents). |
| `--emake-maxlocalagents=<N>` | | Limits the number of local agents used. N=0 uses all available agents (default 0). |
| `--emake-yield-localagents=<N,T>` | | If using more than *N* local agents, eMake releases the number of agents over *N* every *T* seconds so they can be used by another eMake that is looking for local agents. |

# ElectricAccelerator Sample Build

After all ElectricAccelerator components are installed and you are familiar with the concepts, try a test build. Using a text editor, create a makefile with the following content:

## UNIX

```
all: aa bb cc
aa:
        @echo building aa
        @sleep 10
bb:
        @echo building bb
        @sleep 10
cc:
        @echo building cc
        @sleep 10
```

### Windows

```
SLEEP=ping -n 10 -w 1000 localhost>NUL
all: aa bb cc
aa:
        @echo building aa
        -$(SLEEP)
bb:
        @echo building bb
        -$(SLEEP)
cc:
        @echo building cc
        -$(SLEEP)
```

**Note:** "ping" is used in the Windows example because Windows does not have a SLEEP utility.

If you were to run this file with GNU Make, you would expect it to finish in approximately 30 seconds—allowing for each 10-second command to run serially. Running against an ElectricAccelerator cluster with at least three Agents, the commands run in parallel allowing the build to complete much faster.

To start this sample build:

1. Specify the Cluster Manager by using the `--emake-cm=<host>` option. The Cluster Manager is responsible for assigning Agents to Electric Make for processing jobs. The example uses "linuxbuilder" as the Cluster Manager host.

2. Make sure the Electric Make root directory [or directories] specification includes all directories that contain source or input files required by the build. In the example, the only source file is the makefile, which is in the same directory where Electric Make is invoked. Because the default emake root is the current directory, `--emake-root=<path>` is not needed.

```
% emake --emake-cm=linuxbuilder

Starting build: 1
building aa
building bb
building cc
Finished build: 1  Duration: 0:11(m.s) Cluster availability: 100%
```

Cluster availability: 100% indicates the cluster was fully available for the build duration. For more information on cluster sharing and the cluster availability metric, see Annotation.

# Chapter 4: Additional Electric Make Settings and Features

The following topics discuss additional eMake settings and features.

**Topics:**

- Build Classes
- Using Priority Pools
- Using the Proxy Command
- Using Subbuilds
- Building Multiple Targets Simultaneously (GNU Make emulation only)
- Using Electric Make Variables
- Terminating a Build
- Shutting Down Cluster Hosts During Builds

# Build Classes

ElectricAccelerator provides features to organize builds in even the most complex environments. You do not need to create a build class—by default, a simple organizational structure is set up for you. But, if you have diverse product lines, or multiple product releases, Electric Cloud advises taking advantage of ElectricAccelerator build class functionality and setting up build classes.

A *build class* is a flexible, user-defined classification for a designated group of builds. Using build classes is optional, but if you do not assign a build class, Cluster Manager assigns the build to a default class. Electric Cloud recommends using build classes to organize the build management process.

Depending on your company requirements, you might use build classes to organize build groups by version/release, product type, development stage, or platform. You can decide how to use build classes to organize your builds into sets.

Classes have default priorities and boost values. Boost values range from -10 to +10 (default 0) where a higher boost value means builds in that class can use available agents ahead of builds with the same priority and less boost.

## *Tag Definitions*

Each build within a class is identified by a unique string called a *tag*. The build tag definition is a template that expands when a new build starts. The tag is user-defined and generally consists of a generic build name appended with build-specific data constructed from the following variables:

- `GC` - a globally unique number (Global Counter)

- `LC` - a number unique to the build class (Local Counter; the build serial number within the class)

- `BUILD_CLASS` - user-defined build class name

- `BUILD_CLASS_ID` - system-generated number Cluster Manager uses to identify each class

- `USER_NAME` - name of the user who invoked Electric Make

- `MACHINE_NAME` - name of the machine where Electric Make was invoked

- `USER_BUILD_LABEL` - label specified at the Electric Make command line
  for example, `--emake-build-label=my_build`

- `BUILD_OS_ID` - operating system ID under which the build was invoked (`0` = undefined, `1` = Windows, `2` = Solaris, and `3` = Linux)

- `DATE` - build start date and time using variables `Y, y, m d, H, M, S` (for example, `2005-01-18 10:14:32` is `20050118101432`)

- `Y` - the year at build start time (`YYYY`)

- `y` - the year at build start time (`YY`)

- `m` - sequential month number at build start time (`1-12`)

- `d` - sequential day of month at build start time (`1-31`)

- `H` - hour of the day at build start time (`0-23`)

- `M` - minutes at build start time (`0-59`)

- `S` - seconds at build start time (`0-60`)

- `a` - abbreviated day of week at build start time (`WED`)

- `A` - full name day of week at build start time (`Wednesday`)

- `b` - abbreviated month name at build start time (`AUG`)

- `B` - full month name at build start time (`August`)

- `c` - build start date and time using the variables `A, B, d, H, M, S, Y` (for example, `2005-01-18 10:14:32` is `18/01/05 10:14:32`)

Together, the build name and variables are referred to as the *tag definition*. Variable names are case-sensitive.

For example, the tag definition `%BUILD_CLASS%_%LC%_%DATE%` for a build class named `QA_BUILD` creates the following build tag:

    QA_BUILD_1234_20060123185958

When assigning build class tag definitions, choose from the list of tag variables above.

## Build Class Examples

Suppose your company has two major product lines: SuperSoftware and MegaSoftware. SuperSoftware runs on Windows and Solaris platforms. MegaSoftware runs on Windows only. You could begin by setting up three classes that include the product name, the platform, and the current version number for each product:

- You could name the first class Supersoftware_Win_v.2.1. The tag definition for this class would be:

    `%BUILD_CLASS%_%LC%_%a%_%b%_%d%_%H%_%M%_%S%`

  The result would be a series of builds each named, or *tagged*, with the product name, the platform, the version number, a serial number (unique to the class), and the date for each build, for example:

    Supersoftware_Win_v.2.1_12345_WED_AUG_22_14_37_12

- The second class could be named Supersoftware_Sol_v.1.7. The tag definition can be the same as in our first example because it would be distinguished by the second build class name. Build tags in the second class would look like:

    Supersoftware_Sol_v.1.7_12356_WED_AUG_23_11_14_39

- The third class could be named MegaSoftware_Win_v.1.3. For this product, the tag definition would be similar to the previous examples, but also could include the name of the user who started the build because the MegaSoftware team is spread over several different locations. For this class, the tag definition might look like:

    `%BUILD_CLASS%_%LC%_%USER_NAME%_%DATE%`

- As in the first two examples, the result would be a sequentially numbered series of builds with the product name, platform, version number, name of the user who ran the build, and the date of each build assigned through the build class:

    MegaSoftware_Win_v.1.3_12356_JSMITH_20050411100838

Additional classes could be created when the development of SuperSoftware or MegaSoftware entered a new phase, such as a new platform release or a new version release. In this way, the builds for each stage of development can be segmented into logical sets facilitating a more manageable and organized workflow.

## Creating a New Build Class Using the Cluster Manager

1. Type-in your Cluster Manager host name in the browser window location bar.

2. Click the Build > Build Classes tabs. The Build Classes page displays.

3. Click the **New Build Class** link. A blank class details screen appears.

4. Click the **Show Help** link on the right-side of the screen to see field descriptions and fill-in field information accordingly.

5. In the Tag Definition text box, enter the build class tag definition.

   **Note:** To avoid errors, follow standard naming conventions for tag definitions by using numbers, letters, and underscores only **without** leading or trailing white space. Use a percent sign on either side of any variables used. (For example, `%DATE%`.) Use underscores ( _ ) instead of spaces.

6. Continue filling-in the fields, accessing online help for additional information if needed.

7. Click **OK** when fields are filled-in.

**Note:** You can add comments to the build class also. Click the Build > Build Class tabs, then select the Build Class Name where you want to add a comment.

For more information on adding or editing build class comments, adding or editing class properties, or deleting a build class, see the web interface online help.

### Using an Existing Build Class

Assign the build's class name through the `--emake-class=<exact build class name>` eMake command-line option when the build is invoked. If you do not assign a build class, the Cluster Manager assigns the build to the default class. If the class name typed on the eMake command line does not match a class name already in Cluster Manager, eMake exits.

### Makefile Macros

Electric Make automatically creates makefile macros (`ECLOUD_BUILD_CLASS`, and `ECLOUD_BUILD_TAG`) from Cluster Manager build class data. These macros can be used to get generated values into your makefiles. For more information, see Using Electric Make Variables.

# Using Priority Pools

Priority Pools allow you to group resources into pools that can be prioritized differently among groups. Each pool's resources can potentially be used by any build, but builds originating from a pool's "owner" always have first priority to use that pool's resources. Using Priority Pools allows you to manage resource allocation for builds more efficiently.

### Using Priority Pools

Follow this procedure:

1. Enable Priority Pools through either the Cluster Manager or cmtool.

   Using the Cluster Manager, go to Administration > Server Settings and select Priority Pools.

   Using cmtool, run this command:

   ```
   cmtool modifyServer --resourceManagerType prioritypool
   ```

2. Define your pool using the Agents > Resources page. A pool resource uses this form: `__pool_xxx`, for example, `__pool_a`

**Note:** When you define a pool on the Resources page, you must include "`__pool_`". When you include a pool in the `--emake-resource` option, you can omit `__pool_` from the option.

3. Go to Agent > Agent Policies and set Agent allocation policy to shared.

4. Launch eMake and include the following option:

   ```
   --emake-resource="<pool resource>:<static resource>"
   ```

   Information about values for `--emake-resource=`

- When adding a pool resource to the option, you can omit `__pool_` from `__pool_xxx:` and use `xxx:` only.
- A pool name (before the : ) is not required. Not defining the pool name means the build will not use a pool resource.
- A resource name (after the : ) is not required. Not defining the resource name means the build will attempt to use any unused resources that it is allowed to use.

### Additional Information

When attempting to use a pool resource as a normal resource, there is a period (default is 60 seconds) during which, before your build starts, a build from the pool resource's owner can take back the resource.

## *Use Case 1 - High Performance Builds*

You have two pools of resources, one for high performance builds and one for low performance builds. You want to ensure that high performance builds can always use the more powerful 8-core machines and that low performance builds use the 4-core machines. You also want to allow high performance builds to use the 4-core machines when low performance builds are not running. And you want to allow low performance builds to use the 8-core machines when high performance builds are not running. You also have two special software packages, so you define a static resource for each.

Pool makeup:

- Pool a - High performance build resources, five 8-core machines, defined on the Resources page as *__pool_a*

  The machines are named h_1, h_2, h_3, s_1, and s_2.

- Pool b - Low performance build resources, five 4-core machines, defined on the Resources page as *__pool_b*

  The machines are named h_4, m_1, m_2, s_3, and s_4.

- Two static resources are also defined on the Resources page:

  - *s* - This resource includes these machines (which have a specific software package): s_1, s_2, s_3, and s_4 (from using host mask *s\**)

  - *m* - This resource includes these machines (which have a specific software package): m_1 and m_2 (from using host mask *m_1, m_2*)

| Launching eMake with–<br>`-emake-resource=` | means the build uses |
|---|---|
| `"a:"` | Pool "a" and any unused resources |
| `"a:s"` | Static resource "s" (four machines). If a low performance build is running, only machines "s_1" and "s_2" (two machines) are used. |
| `"b:m"` | Static resource "m" only |

| | |
|---|---|
| `"a:m"` | Static resource "m" only. If a low performance build is running, no machines are available and the build must wait. |
| `":s"` or `"s"` | Static resource "s" if low and high performance builds are not running. If a low performance build is running, then only machines "s_1" and "s_2" are used. If a high performance build is running, then only machines "s_3" and "s_4" are used. If low and high performance builds are running, no machines are available and the build must wait. |
| `":"` or `""` | Any unused resources (this is considered a common build) |

## Use Case 2 - Multiple Departments

There are two departments and each department has its own pool of machines. They want to contribute their machines toward a common pool so each department can use the other department's machines while still ensuring that their own machines are available for their department's builds. One department owns 20 machines and the other department owns 10 machines. Each department also owns a small number of 64-bit machines. IT contributes an additional four machines that any department can use.

Pool makeup:

- Pool Depta - Department A's resources, 20 machines, defined on the Resources page as __pool_Depta

  The machines are named a_01 through a_16 and a_17_64 through a_20_64 (these last four machines are 64-bit).

- Pool Deptb - Department B's resources, 10 machines, defined on the Resources page as __pool_Deptb

  The machines are named b_01 through b_08 and b_09_64 and b_10_64 (these last two hosts are 64-bit).

- Static resource *64bit* - This resource includes these machines: a_17_64 through a_20_64 and b_09_64 and b_10_64 (from using host mask *_64).

- General build machines - IT department-supplied general build machines, four machines

| Launching eMake with- -emake-resource= | means the build uses |
|---|---|
| `"Deptb:"` | Pool "Deptb" and any unused resources, using general build machines first and then Deptartment A's machines if that department is not running a build. |
| `"Depta:"` | Pool "Depta" and any unused resources, using general build machines first and then Deptartment B's machines if that department is not running a build. |
| `"Depta:64bit"` | Static resource "64bit" (six machines). If Department B is running a build, only "a_17_64" through "a_20_64" (four machines) are used. |
| `":"` or `""` | Any unused resource, using general build machines first and then Deptartment A's and Deptartment B's machines if those departments are not running builds (this is considered a common build). |

# Using the Proxy Command

Normally, Electric Make sends commands to the Agent for execution. In some cases, however, it might not be desirable or possible to execute a particular command on the Agent.

## Proxy Command Location

During installation, the `proxyCmd` binary is installed on every host:

- Linux: `/opt/ecloud/i686_Linux/bin/proxyCmd`

- Solaris: `/opt/ecloud/sun4u_SunOS/bin/proxyCmd`

- Windows: `C:\ECloud\i686_win32\bin\proxyCmd.exe`

When invoked by the Agent, it is: `proxyCmd <program> <arg1> ...`

ElectricAccelerator executes *<program>* on the host build system and proxies the result back to the Agent so it can continue remote execution.

## Determine If You Can Use the Proxy Command

You can use the "proxy command" feature to run a command locally on the Electric Make machine *if both of the following are true*:

- You have a command that cannot run on the Agent, either because it returns incorrect results or because it is not available.

- The command in question *does not read or write build sources or output*—it only makes external [outside of the build] read-only queries.

The second item is particularly important because the command runs on the host build machine, outside of the virtualized file system. Because ElectricAccelerator cannot track the activity of this process, the dependency and conflict resolution mechanisms that prevent build output corruption are circumvented. It is also important that the process is read-only [that the command make no changes to whatever system it is querying] because in parallel builds with conflicts, Electric Make could rerun the job producing unintended side effects.

**Note:** For these reasons, it is important to use the "proxy command" only when necessary. Where possible, try to ensure cluster hosts have the same tools installed as the build system.

## Proxy Command Example 1

The simplest *safe* use of the `proxyCmd` is a source control system query. For example, a particular build step queries the source control system for branch identification using a tool called `getbranch` that it embeds in a version string:

```
foo.o:
    gcc -c -DBRANCH=`getbranch` foo.c
```

It is preferable to avoid installing and configuring a full deployment of the source control system on the host when only this simple query command is needed.

In the following example, the `proxyCmd` provides an efficient solution. By replacing `getbranch` with `proxyCmd getbranch`, you avoid having to install the `getbranch` tool and its associated components on the host:

```
foo.o:
    gcc -c -DBRANCH=`proxyCmd getbranch` foo.c
```

## Proxy Command Example 2

A less invasive implementation that does not require makefile modifications and allows compatibility with non-ElectricAccelerator builds is to create a link on all agent machines using the name of the tool [for example, 'getbranch'] found on the eMake host to `proxyCmd`. For Windows operating systems that do not support `symlinks`, use the `copy` command. When invoked under a different name, `proxyCmd` knows to treat the linked name as the *<program>* to execute:

```
# ln -s /opt/ecloud/i686_Linux/bin/proxyCmd /usr/bin/getbranch
```

# Using Subbuilds

The Electric Make subbuild feature is designed to help speed up component builds through intelligent build avoidance. Currently, the subbuild feature's scope includes the following use case:

```
Makefile:
-------------
.PHONY: a b
all: a b
        @echo all
a b:
        $(MAKE) -C $@

a/Makefile
-------------
all: a.lib
        @echo a

a.lib:
        @touch $@

b/Makefile
-------------
all: ../a/a.lib
        @echo b
```

**Explanation**: If something from 'a' changes, and you are building from 'b', the only way to pick up the new a.lib is to build from the top level directory. With subbuilds, you know b's dependencies so you can build those dependencies directly without having to build *everything* from the top level directory.

The subbuild database must be built beforehand to make the dependency list available without having to parse any Makefiles that are not in the current directory.

The following sections describe how to use subbuilds. Refer to Subbuild Limitations for additional information about subbuild limitations.

## Subbuild Database Generation

The following command runs your build as normal and also generates a subbuild database with the name emake.subbuild.db.

```
emake --emake-gen-subbuild-db=1 --emake-root=<path> --emake-subbuild-db=emake.su
bbuild.db
```

Where <path> is the Electric Make root directory setting.

--emake-root is required for cluster builds.

--emake-subbuild-db is optional. If it is missing, the default emake.subbuild.db name is used.

## Run a Build Using Subbuild

The following command runs a build with subbuild information:

```
emake --emake-subbuild-db=emake.subbuild.db
```

Specify --emake-subbuild-db=<file> to run a build with subbuild information. When you invoke eMake with the --emake-subbuild-db option, it uses the dependencies extracted from the makefile and the subbuild database to determine which build components are prerequisites of the desired current make, then rebuilds those components before proceeding as normal.

When you specify `--emake-subbuild-db=<file>`, do *not* specify `--emake-gen-subbuild-db`, otherwise eMake regenerates the database.

## Subbuild Limitations

- There is no incremental building of the database. Each time you change something in a makefile in your build, you must rebuild the database by doing a full build.

- The database is not currently optimized for size. This might result in an extremely large database for very large builds.

- Subbuilds do not provide additional gains in non-recursive make builds.

- Because of the manner in which subbuilds are currently scheduled, there is interleaving output for the "Entering directory..." and "Leaving directory..." messages.

  For example: If a subbuild database was built for the following build:

```
Makefile:
-------------
.PHONY: a b
all: a b
a b:
        $(MAKE) -C $@

a/Makefile
----------------
all: a.lib
a.lib:
        echo aaa > $@

b/Makefile:
----------------
all: ../a/a.lib
        echo b
```

When you proceed to build just 'b' (maybe with "emake -C b") and a/a.lib is missing, you receive "entering directory a" after "entering directory b", even though 'a' is supposed to be built before 'b'.

```
make: Entering directory 'b'
make -C a
make[1]: Entering directory 'a'
echo aaa > a.lib
make[1]: Leaving directory 'a'
echo b
b
make: Leaving directory 'b'
```

### Information Applying to Local Builds Only

- Rules to build a sub-directory's output files must not overlap.
  For example: The rule to build `sub1/foo.o` must appear in `sub1/Makefile` only and not `sub2/Makefile`. Default suffix rules can cause eMake to find a way to build `sub1/foo.o` while trying to build `sub2`. In this situation, adding ".`SUFFIXES:`" to `sub2/Makefile` can resolve the issue.

- Subbuilds require that the build be componentized to some degree.

- Subbuilds require that you have practiced "good hygiene" in your build tree—there must be explicit dependencies mentioned in the component makefiles.

For example: If a build has two components, `foo` and `bar`, where `foo` produces a library `foo.dll` and `bar` uses that library, the rule might be written to produce `bar.exe` such as this in `bar/Makefile`:

```
bar.exe: $(BAR_OBJS)
        link $(BAR_OBJS) -l $(OUTDIR)/foo/foo.dll
```

For subbuilds to work (in local mode), it must be modified as in the following:

```
bar.exe: $(BAR_OBJS) $(OUTDIR)/foo/foo.dll
        link $(BAR_OBJS) -l $(OUTDIR)/foo/foo.dll
```

Note that it is explicitly stated that `bar.exe` requires `foo.dll`. Also note that it is NOT required to have a rule to build `foo.dll` in `bar/Makefile`.

There cannot be ANY rule at all to build `$(OUTDIR)/foo/foo.dll` in `bar/Makefile`, explicit or implicit, otherwise you will get the wrong information for building `foo/foo.dll` in the subbuilds database. The subbuilds database currently allows updates to existing entries while building the database.

# Extension for Building Multiple Targets Simultaneously (GNU Make Emulation Only)

Sometimes you might want one rule that creates multiple outputs simultaneously, and using a pattern rule might not be a suitable solution.

The `#pragma multi` directive enables you to create non-pattern rules that have multiple outputs.

`#pragma multi` causes the targets of the immediately following rule or dependency specification to be treated as updated together if/when they are updated. For example, the following produces one rule and one rule job only, rather than three of each:

```
#pragma multi
a b c: ; @echo building a b and c
```

## Important Notes

- If you apply `#pragma multi` to a target list, then you must apply it to all overlapping target lists, and those lists must specify the same set of targets (though they might do so in a different order).

- Target-specific and pattern-specific variable assignments for the targets of a `#pragma multi` rule must agree; otherwise it is unspecified whether eMake chooses the assignments for just one target, or somehow combines them all.

- You might not apply `#pragma multi` to static patterns, double-colon rules, or pattern targets that follow non-pattern targets on the same line.

- If you apply `#pragma multi` to a non-static pattern, a warning will issue.

- A `#pragma multi` rule with commands might not override or be overridden by other commands for the same targets.

- A `#pragma multi` dependency specification must correspond to a `#pragma multi` rule with commands having the same set of targets. Otherwise, eMake will error out. Implicit rules are not searched to find such missing commands.

- `$@` has the same meaning as it does in multiple-target patterns: the target that first caused the rule to be needed.

- Setting either of the following disables `#pragma multi`: `--emake-disable-pragma=multi` or `--emake-disable-pragma=all`.

# Using Electric Make Variables

Electric Make automatically defines several variables that can be used in makefiles to access Cluster Manager-specific values during a build. For example, you could insert the build tag into your compilation step by typing:

```
main.o:main.cpp
        gcc -DBUILD_TAG="$(ECLOUD_BUILD_TAG)" -o main.o main.cpp
```

The variables Electric Make automatically creates are described in the following table.

| Variable | Definition |
|---|---|
| ECLOUD_BUILD_CLASS | The build class as specified on the command line with --emake-class. For more information, see Build Classes. |
| ECLOUD_BUILD_ID | A globally unique build ID. This value is guaranteed unique across all builds in all classes for this Cluster Manager. |
| ECLOUD_BUILD_TAG | The tag as configured in Cluster Manager for this build class. For more information, see Build Classes. |
| ECLOUD_BUILD_COUNTER | A unique ID for this build within the current build class. This value is guaranteed to be unique across all builds in the same class, but not across builds in different classes. |
| ECLOUD_BUILD_TYPE | Allows you to discover the build type: *local* for a local build; or *remote* for a cluster build. |

# Stopping a Build

You can stop an in-progress build by one of these methods:

- Press **Ctrl-C** from the terminal where you invoked Electric Make, or

- Use ElectricAccelerator's web interface, or

- Use **cmtool**—see the *cmtool Reference Guide* at http://docs.electric-cloud.com/accelerator_doc/AcceleratorIndex.html.

Cluster Manager terminates builds that seem to be hung. If Cluster Manager does not receive a request from eMake for 60 seconds, it considers the build hung and the build will be terminated.

### Using the Cluster Manager to Stop a Build

1. Open the web interface by typing the Cluster Manager host name in the location bar of your browser window.

2. Click **Builds**.

3. Click **Stop Build** in the Action column on the row exhibiting your build ID and Name.

### Using cmtool to Stop a Build

**Note:** This is an "advanced user" option, which assumes you are already familiar with using command-line tools.

1. Request a list of running builds. The syntax is:

```
% cmtool --cm <clustermanager:port> getBuilds --filter <field name>=<value>
```

For example, if linuxbuilder is the Cluster Manager host name, type:

```
% cmtool --cm linuxbuilder getBuilds --filter "result !=-1"
```

A list of running builds will display accompanied by a number of attributes for each build—for example: Build ID, machine name, build class, owner, build start time, and so on. The Build ID is used to identify a build for termination. Also you can obtain additional information and/or include comments about the build. For example, to sort builds by start time and request only the first ten builds display, enter:

```
% cmtool --cm linuxbuilder getBuilds --order Id --filter="ID<11"
```

To get failed builds:

```
% cmtool --cm linuxbuilder getBuilds --order "start_time desc"--filter "resul
t !=0"
```

2. After you determine which build you need to terminate, use this syntax to end the build:

```
% cmtool --cm <clustermanager> stopBuild <buildId>
```

For example, to end build 4458, you would type:

```
% cmtool --cm linuxbuilder stopBuild 4458
```

# Shutting Down Cluster Hosts During Builds

You can remove agent machines from the cluster during a build without requiring downtime or disabling Agents in the Cluster Manager, with the following advisories:

- If possible, shut down Agents before rebooting agent machines. Not shutting down Agents before a reboot might cause commands to fail unexpectedly, which could manifest in spurious build failures.

- After rebooting, agent machines that automatically start Agents will connect to the Cluster Manager and will be assigned builds and so on.

# Chapter 5: Make Compatibility

ElectricAccelerator is designed to be completely compatible with existing Make variants it emulates. There are, however, some differences in behavior that might require changes to makefiles or scripts included in the build.

Almost all GNU Make and NMAKE options are valid for use with ElectricAccelerator. However, ElectricAccelerator does not support some GNU Make and NMAKE options.

The following topics documents those differences and what actions to take to ensure your build is compatible with ElectricAccelerator.

**Topics:**

- Unsupported GNU Make Options and Features
- Unsupported NMAKE Options
- Commands that Read from the Console
- Transactional Command Output
- Stubbed Submake Output
- Hidden Targets
- Wildcard Sort Order
- Delayed Existence Checks
- Multiple Remakes (GNU Make only)
- NMAKE Inline File Locations (Windows only)
- How eMake Processes MAKEFLAGS

# Unsupported GNU Make Options and Features

## Unsupported GNU Make Options

| Option | eMake response if specified |
|---|---|
| `-d (debug)` | error message |
| `-j (run parallel)` | ignored |
| `-l (load limit)` | ignored |
| `-o (old file)` | error message |
| `-p (print database)` | error message |
| `-q (question)` | error message |
| `-t (touch)` | error message |

## Unsupported GNU Make 3.81 Features

Electric Make does not support the following GNU Make 3.81 features:

- Though `$(eval)` is allowed in rule bodies, any variables created by the `$(eval)` exist only in the scope of the rule being processed, *not* at the global scope as in GNU Make. For example, the following is not supported:

```
all:
        $(eval foo: bar)
```

- Using `$*` for an archive member target

## GNU Make 3.82 Note

GNU Make 3.82 emulation supports the following GNU Make 3.82 features:

- Multi-line variable definition types

- Behavior in which pattern rule searches prefer the most-specific matching pattern, rather than the first matching pattern. For example, for a target {{foo.o}} and patterns {{%.o}} and {{f%.o}}, in that order, gmake 3.81 uses the {{%.o}} pattern, because it is the first match, but gmake 3.82 uses the {{f%.o}} pattern, because it is a more specific match

Other GNU Make 3.82 features are not supported.

# Unsupported NMAKE Options

```
/C (suppress output)
/Q (check timestamps, don't build)
/T (update timestamps, don't build)
/NOLOGO (suppresses NMAKE copyright message)
/ERRORREPORT (sends information to Microsoft)
```

# Commands that Read from the Console

When GNU Make or NMAKE invokes a makefile target command, that process inherits the standard I/O streams of the Make process. It is then possible to invoke commands that expect input during a build, either from the terminal or passed into the Make standard input stream. For example, a makefile such as:

```
all:
        @cat
```

could be invoked like this:

```
% echo hello | gmake
hello
```

More commonly, a command in a makefile might prompt the user for some input, particularly if the command encounters an error or warning condition:

```
all:
        @rm -i destroy

% gmake
rm: remove regular file `destroy'? y
```

Neither of these constructs is generally recommended because systems that require runtime user input are tedious to invoke and extremely difficult to automate.

Because Accelerator runs commands on cluster hosts where processes do not inherit the I/O streams and console of the parent eMake, makefiles (such as the examples above) with commands expecting interactive input are not supported.

In the majority of cases, tools that prompt for console input contain options to disable interactive prompting and proceed automatically. For example, invoking "`rm`" without the "`i`" enables this behavior. For those that do not, explicitly feeding expected input (either from a file or directly by shell redirection or piping) will suffice. For example:

```
all:
echo y | rm -i destroy
```

Finally, tools such as Expect (see http://expect.nist.gov) that automate an interactive session can be used for commands that insist on reading from a console.

# Transactional Command Output

Electric Make simulates a serial GNU Make or NMAKE build. Though Electric Make runs many commands in parallel, the command output (including text written to standard streams and changes to the file system, such as creating or updating files) appears serially in the original order without overlapping. This feature is called *transactional output* (or "serial order execution") and is unique to Accelerator among parallel build systems. This feature ensures standard output streams and underlying file systems always reflect a consistent build execution state, regardless of how many jobs Electric Make is actually running concurrently on the cluster.

Transactional output is achieved by buffering the results of every command until the output from all preceding commands is written. Buffering means that while the output *contents* on the standard streams matches GNU Make or NMAKE exactly, the *timing* of its appearance might be a little unexpected. For example:

- **"Bursty" output** – One of the first things you notice when running a build with Accelerator is that it appears to proceed in bursts, with many jobs finishing in quick succession followed by pauses. This type of output is normal during a highly parallel build because many later jobs might have completed and output is ready to be written as soon as longer, earlier jobs complete. The system remains busy, continuously running jobs throughout the build duration, even if the output appears to have paused momentarily.

- **Output follows job completion** – GNU Make and NMAKE print commands they are executing before they are invoked. Because Electric Make is running many commands in parallel and buffering results to ensure transactional output, command-line text appears with the output from the command *after* the command has completed. For example, the last command printed on standard output is the job that just completed, not the one currently running.

- **Batch output** – As a way to provide feedback to the user during a long-running execution, some commands might write to standard output continuously during their execution. Typically, these commands might print a series of ellipses or hash marks to indicate progress or might write status messages to standard error as they run. More commonly, a job might have several long-running commands separated with echo statements to report on progress during build execution:

For example, consider a rule that uses `rsync` to deploy output:

```
install:
        @echo "Copying output into destination"
        rsync -vr $(OUTPUT) $(DESTINATION)
        @echo "done"
```

With GNU Make, users first see the `Copying output` echo, then the state information from `rsync` as it builds the file list, copies files, and finally, they see the `done` echo as the job completes.

With Electric Make, all output from this job step appears instantaneously in one burst when the job completes. By the time any output from `echo` or `rsync` is visible, the entire job has completed.

## Stubbed Submake Output

Recursive Makes (also called *submakes* because they are invoked by a parent or *top-level* Make instance) are often used to partition a build into smaller modules. While submakes simplify the build system by allowing individual components to be built as autonomous units, they can introduce new problems because they fragment the dependency tree. Because the top-level Make does not coordinate with the submake—it is just another command—it is unable to track targets across Make instance boundaries. For a discussion of submake problems, see "Recursive Make Considered Harmful" by Peter Miller (http://miller.emu.id.au/pmiller/books/rmch/)

Submakes are particularly problematic for parallel builds because a build composed of separate Make instances is unable to control target serialization and concurrency between them. For example, consider a build divided into two phases: a `libs` submake that creates shared libraries followed by `apps` that builds executables that link against those libraries. A typical top-level makefile that controls this type of build might look like this:

```
all: makelibs makeapps
makelibs:
        $(MAKE) -C libs
makeapps:
        $(MAKE) -C apps
```

This type of makefile works fine for a serialized make, but running in parallel it can quickly become trouble:

- If `makelibs` and `makeapps` run concurrently (as the makefile "all" rule implies), link steps in the `apps` Make instance might fail if they prematurely attempt to read `libs` generated libraries. Worse, they might link against existing, out-of-date library copies, producing incorrect output without error. This is a failure to correctly serialize dependent targets.

- Alternatively, if `apps` is forced to wait until `libs` completes, even `apps` targets that do not depend on `libs` (for example, all the compilation steps, which are likely the bulk of the build) are serialized unnecessarily. This is a failure to maximize concurrency.

Also important to note: Submakes are often spawned indirectly from a script instead of by makefile commands,

```
makelibs:
        # 'do-libs' is a script that will invoke 'make'
        do-libs
```

which can make it difficult for a Make system to identify submake invocations, let alone attempt to ensure their correct, concurrent execution.

These problems are exacerbated with distributed (cluster) parallel builds because each make invocation is running ona remote Agent.

Correct, highly concurrent parallel builds require a single, global dependency tree. Short of re-architecting a build with submakes into a single Make instance, this is very difficult to achieve with existing Make tools.

An ideal solution to parallelizing submakes has the following properties:

- maximizes concurrency, even across make instances

- serializes jobs that depend on output from other jobs

- minimizes changes to the existing system (in particular, does not require eliminating submakes or prohibit their invocation from scripts)

## Submake Stubs

The parallel submake problem is solved by introducing submake stubs. Electric Make dispatches *all* commands, regardless of tool (compiler, packager, submake, script, and so on) to cluster Agents. After the Agent executes a command, it sends the results (output to standard streams and exit status) back to Electric Make, which then sends the next job command.

If the command run by the Agent invokes Electric Make (either directly by the expanded $(MAKE) variable or indirectly through a script that calls emake), a new top-level build is *not* started. Instead, an Electric Make process started on the Agent enters *stub* mode and it simply records details of its invocation (current working directory, command-line, environment, and so on) and immediately exits with status **0** (success) without writing output or reading any makefiles. The Agent then passes invocation details recorded by the stub back to the main Electric Make process on the host build machine, which starts a new Make instance and integrates its targets (which run in parallel on the cluster just like any other job) into the global dependency tree. Commands that follow a submake invocation are logically in a separate job serialized after the last job in the submake.

The following example illustrates this process:

```
...

Makelibs:
        @echo "start libs"
        @$(MAKE) -C libs
        @echo "finish libs"
```

This example is diagrammed in steps as shown in the following illustrations.



1. Electric Make determines that Makelibs target needs to be built.

2. Electric Make connects to Electric Agent.

3. Electric Make sends the first command, echo "start libs".

4. The Agent invokes the command, `echo "start libs"`, and captures the result.

5. The Agent returns the result, `"start libs"`, to Electric Make.

6. Electric Make sends the second command, `emake -f libs.mak`

7. The Agent runs the second command, `emake -f libs.mak`

8. `emake` enters stub mode and records the current working directory, command, and environment, then exits with no output and status code **0**

9. Agent returns the stub mode result and a special message stating a new Make was invoked with recorded properties (Electric Make).



10. Electric Make starts a new Make instance, reads the makefile, `libs.mak`, and integrates the file into the dependency tree.

11. New jobs are created and run against the cluster to build all targets in the `libs.mak` makefile.

12. Electric Make splits the last command in the `Makelibs` target, `echo "finish libs"`, into a continuation job that is defined to have a serial build order later than the last job in the submake, but there is no explicit dependency created that requires it to run later than any of the jobs in the submake. This means that it might run in parallel (or even before) the jobs in the submake, but if for some reason that is not safe, Electric Make will be able to detect a conflict and rerun the continuation job.



13. Electric Make sends the last command, `echo "finish libs"`.

14. Agent runs the command, `echo "finish libs"`, and captures the result.

15. Agent returns the result, `"finish libs"`, to Electric Make.

The Electric Make Stub solution addresses three basic parallel submake problems by:

- **Running parallel submakes in stub mode** – Stubs finish instantaneously and discover jobs as quickly as possible so the build is as concurrent as possible.

- **Creating a single logical Make instance** – New Make instances are started by the top-level Electric Make process only and their targets are integrated into the global dependency tree. Electric Make can then track dependencies across Make instances and use its conflict resolution technology to re-order jobs that might have run too soon.

- **Capturing submake invocations** – By capturing submake invocations as they occur, the submake stub works with your makefiles and builds scripts for the majority of cases "as-is." However, the stub introduces a behavior change that might require some makefile changes. See Submake Stub Compatibility.

# Submake Stub Compatibility

Submake stubs allow your existing recursive builds to behave like a single logical Make instance to ensure fast, correct parallel builds. Stubs do introduce new behavior, however, that might appear as build failures. This section describes what constructs are not supported and what must be changed to make stubs compatible with submake stubs.

At this point, it is useful to revisit the relationship between commands and rules:

```
all:
        echo "this is a command"
        echo "another command that includes a copy" ; \
        cp aa bb
        echo "so does this command" ; \
        cp bb cc
        cp cc dd
```

The rule above contains four commands: (1) echo, (2) echo-and-copy, (3) another echo-and-copy, and (4) a copy. Note how semicolons, backslashes, and new lines delimit (or continue) commands. The rule becomes a job when Make schedules it because it is needed to build the "all" target.

Submake stubs' most important features:

- A submake stub never has any output and always exits successfully.

- The Agent sends stub output back (if any) after each command.

- Commands that follow a stub are invoked after the submake in the serial order.

Because a submake stub is really just a way of marking the Make invocation and does not actually do anything, *you cannot rely on its output (stdout/stderr, exit status, or file system changes) in the same command*.

In the following three examples, incompatible Accelerator commands are described and examples for fixing the incompatibility are provided.

### Example 1: A command that reads submake file system changes

```
makelibs:
        $(MAKE) -C libs libcore.aa ; cp libs/libcore.aa /lib
```

In this example, a single command spawns a submake that builds `libcore.a` and then copies the new file into the `/lib` directory. If you run this rule as-is with Accelerator, the following error might appear:

```
cp: cannot stat 'libs/libcore.aa': No such file or directory
make: *** [all] Error 1
```

The submake stub exited immediately and the `cp` begins execution right after it in the same command. Electric Make was not notified of the new Make instance yet, so no jobs to build `libcore.a` even exist. The `cp` fails because the expected file is not present.

An easy fix for this example is to remove the semicolon and make the `cp` a separate command:

```
makelibs:
        $(MAKE) -C libs libcore.aa
        cp libs/libcore.aa /lib
```

Now the `cp` is in a command after the submake stub sends its results back to the build machine. Electric Make forces the `cp` to wait until the submake jobs have completed, thus allowing the copy to succeed because the file

is present. Note that this change has no effect on other Make variants so it will not prevent building with existing tools.

**Note:** In general, Accelerator build failures that manifest themselves as apparent re-ordering or missing executions are usually because of commands reading the output of submake stubs. In most cases, the fix is simply to split the rule into multiple commands so the submake results are not read until after the submake completes.

### Example 2: A command that reads submake *stdout*

```
makelibs:
        $(MAKE) -C libs mkinstall > installer.sh
```

The output is captured in a script that could be replayed later. Running this makefile with Accelerator always produces an empty `installer.sh` because submake stubs do not write output. When Electric Make does invoke this Make instance, the output goes to standard output, as though no redirection was specified.

Commands that read from a Make *stdout* are relatively unusual. Those that do often read from a Make that does very little actual execution either because it is invoked with `-n` or because it runs a target that writes to *stdout* only. In these cases, it is not necessary to use a submake stub. The Make instance being spawned is small and fast, and running it directly on the Agent in its entirety (instead of returning to the host build machine and distributing individual jobs back to the cluster) does not significantly impact performance.

You can specify that an individual `emake` invocation on the Agent does not enter stub mode, but instead behaves like a local (non-cluster) Make simply by setting the `EMAKE_BUILD_MODE` environment variable for that instance:

```
makelibs:
        EMAKE_BUILD_MODE=local \
$(MAKE) -C libs mkinstall >
        installer.sh
```

For Windows:

```
makelibs:
        set EMAKE_BUILD_MODE=local && $(MAKE) -C libs mkinstall >
                installer
```

Electric Make automatically uses local mode when the `-n` switch is specified.

### Example 3: A command that reads submake exit status

```
makelibs:
        $(MAKE) -C libs || echo "failure building libs"
```

This example is a common idiom for reporting errors. The `||` tells the shell to evaluate the second half of the expression *only* if Make exits with non-zero status. Again, because a submake stub always exits with **0**, this clause will never be invoked with Accelerator, even if it would be invoked with GNU Make. If you need this type of fail-over handling, consider post-processing the output log in the event of a build failure. Also see Annotation for more information.

Another common idiom in makefiles where exit status is read in loop constructs such as:

```
all:
        for i in dir1 dir2 dir3 ; do \
                $(MAKE) -C $$i || exit 1;\
        done
```

This is a single command: a "for" loop that spawns three submakes. The `|| exit 1` is present to prevent GNU Make from continuing to start the next submake if the current one fails. Without the `exit 1` clause, the command exit status is the exit status from the last submake, regardless of whether the preceding submakes succeeded or failed, or regardless of which error handling setting (for example, `-i`, `-k`) was used in the Make. The `|| exit 1` idiom is used to force the submakes to better approximate the behavior of other Make targets, which stops the build on failure.

On first inspection, this looks like an unsupported construct for submake stubs because exit status is read from a stub. Accelerator never evaluates the `|| exit 1` because the stub always exits with status code **0**. However, because the submakes really are reintegrated as targets in the top-level Make, a failure in one of them halts the build as intended. Explained another way, a submake loop is already treated as a series of independent targets, and the presence or absence of the GNU Make `|| exit 1` hint does not change this behavior. These constructs should be left as-is.

# Hidden Targets

Electric Make differs from other Make variants in the way it searches for files needed by pattern rules (also called suffix or implicit rules) in a build.

- At the beginning of each Make instance, Electric Make searches for matching files for all pattern rules *before it runs any commands.* After eMake has rules for every target that needs updating, it schedules the rules (creating jobs) and then runs those jobs in parallel for maximum concurrency.

- Microsoft NMAKE and GNU Make match pattern rules *as they run commands*, interleaving execution and pattern search.

Because of the difference in the way eMake and NMAKE match pattern rules, NMAKE and eMake can produce different makefile output with *hidden targets*. A hidden target (also known as a "hidden dependency") is a file that is:

- created as an undeclared side-effect of updating another target

- required by a pattern to build a rule

Consider the following makefile example:

```
all: bar.lib foo.obj

bar.lib:
        touch bar.lib foo.c

.c.obj:
        touch $@
```

Notice that `foo.c` is created as a side-effect of updating the `bar.lib` target. Until `bar.lib` is updated, no rule is available to update `foo.obj` because nothing matches the `.c.obj` suffix rule.

NMAKE accepts this construct because it checks for `foo.c` existence before it attempts to update `foo.obj`. NMAKE produces the following result for this makefile:

```
touch bar.lib foo.c
touch foo.obj
```

Electric Make, however, performs the search for files that match the suffix rule once so it can schedule all jobs immediately and maximize concurrency. Electric Make will not *notice* the existence of `foo.c` by the time it attempts to update `foo.obj`, even if `foo.c` was created. Electric Make fails with:

```
NMAKE : fatal error U1073: don't know how to make 'foo.obj'
Stop.
```

The fix is simply to identify `foo.c` as a product for updating the `bar.lib` target, so it is no longer a hidden target. For the example above, adding a line such as `foo.c: bar.lib` is sufficient for Electric Make to understand that `.c.obj` suffix rule matches the `foo.obj` target if `bar.lib` is built first. Adding this line is more accurate and has no effect on NMAKE.

GNU Make is similarly incompatible with Electric Make, but the incompatibility is sometimes masked by the GNU Make directory cache. GNU Make attempts to cache the directory contents on first access to improve performance. Unfortunately, because the time of first directory access can vary widely depending on which targets reference the directory and when they execute, GNU Make can appear to fail or succeed randomly in the presence of hidden targets.

For example, in this makefile, the file `$(DIR)/foo.yy` is a hidden target created as a side-effect of updating `aa` and needed by the pattern rule for `foo.xx`:

```
all: aa bb
aa:
        touch $(DIR)/foo.yy
bb: foo.xx

%.xx: $(DIR)/%.yy

        @echo $@
```

Depending on the value of `DIR`, this build might or might not work with GNU Make:

```
% mkdir sub; gmake DIR=sub
touch sub/foo.yy
foo.xx

% gmake DIR=.
touch ./foo.yy
gmake: *** No rule to make target 'foo.xx', needed by 'bb'.  Stop.
```

Electric Make does not attempt to emulate this behavior. Instead, it consistently refuses to schedule `foo.xx` because it depends on a hidden target (just as it did in the NMAKE emulation mode in the earlier example). In this case, adding a single line declaring the target: `$(DIR)/foo.yy: aa` is sufficient to ensure it always matches the `%.xx` pattern rule.

**Note:** If a build completes successfully with Microsoft NMAKE or GNU Make, but fails with "don't know how to make <x>" with Electric Make, look for rules that create <x> as a side-effect of updating another target. If <x> is required by a suffix rule also, it is a hidden target and needs to be declared as explicit output to be compatible with Electric Make.

There are many other reasons why hidden targets are problematic for all Make-based systems and why eliminating them is good practice in general. For more information, see:

- "Paul's Rules of Makefiles" by Paul Smith at http://www.make.paulandlesley.org/rules.html. Among other useful guidelines for writing makefiles, the primary author of GNU Make writes, "Every non-`.PHONY` rule **must** update a file with the *exact* name of its target. [...] That way you and GNU Make always agree."

- "The Trouble with Hidden Targets" by John Graham-Cumming at http://www.cmcrossroads.com/content/view/6519/120/.

**Note:** In a limited number of cases, eMake might conclude that a matching pattern rule for an output target does not exist. This occurs because eMake's strict string equality matching for prerequisites determines that the prerequisites are different (even though the paths refer to the same file) and that there is no rule to build it.

# Wildcard Sort Order

A number of differences exist between GNU Make and Electric Make regarding the use of `$(wildcard)` and prerequisite wildcard sort order functions. When using the `$(wildcard)` function or using a wildcard in the rule prerequisite list, the resultant wildcard sort order might be different for GNU Make and Electric Make.

Different GNU Make versions are not consistent and exhibit permuted file lists. Even a GNU Make version using different system libraries versions will exhibit inconsistencies in the wildcard sort order.

No difference exists in the file list returned, other than the order. If the sort order is important, you might wrap `$(wildcard)` with `$(sort)`.

For example:

```
$(sort $(wildcard *.foo))
```

Do not rely on the order of rule prerequisites generated with a wildcard. For example, using `target: *.foo`.

Relying on the order of `*.foo` can be dangerous for both GNU Make and Electric Make. Neither GNU Make nor Electric Make guarantees the order in which those prerequisites are executed.

# Delayed Existence Checks

All Make variants process makefiles by looking for rules to build targets in the dependency tree. If no target rule is present, Make looks for a file on disk with the same name as the target. If this *existence check* fails, Make notes it has no rule to build the target.

Electric Make also exhibits this behavior, but for performance reasons it delays the existence check for a target without a makefile rule until just before that target is needed. The effect of this optimization is that Electric Make might run more commands to update targets (than GNU Make or NMAKE) *before* it discovers it has no rule to make a target.

For example, consider the following makefile:

```
all: nonexistent aa bb
aa:
        @echo $@
bb:
        @echo $@
```

GNU Make begins by looking for a rule for `nonexistent` and, when it does not find the rule, it does a file existence check. When that fails, GNU Make terminates immediately with:

```
make: *** No rule to make target 'nonexistent', needed by 'all'.  Stop.
```

Similarly, NMAKE fails with:

```
NMAKE : fatal error U1073: don't know how to make 'nonexistent' Stop.
```

Electric Make delays the existence check for `nonexistent` until it is ready to run the `all` target. First, Electric Make finishes running commands to update the `aa` and `bb` prerequisites. Electric Make fails in the same way, but executes more targets first:

```
aa
bb
make: *** No rule to make target 'nonexistent', needed by 'all'.  Stop.
```

Of course, when the existence check succeeds (as it does in any successful build), there is no behavioral difference between Electric Make and GNU Make or Microsoft NMAKE.

# Multiple Remakes (GNU Make only)

GNU Make has an advanced feature called Makefile Remaking, which is documented in the GNU Manual, "How Makefiles are Remade," and available at:
http://www.gnu.org/software/make/manual/make.html#Remaking-Makefiles

To quote from the GNU Make description:

"Sometimes makefiles can be remade from other files, such as RCS or SCCS files. If a makefile can be remade from other files, you probably want make to get an up-to-date version of the makefile to read in.

"To this end, after reading in all makefiles, make will consider each as a goal target and attempt to update it. If a makefile has a rule which says how to update it (found either in that very makefile or in another one) or if an implicit rule applies to it (see section Using Implicit Rules), it will be updated if necessary. After all makefiles have been checked, if any have actually been changed, make starts with a clean slate and reads all the makefiles over again. (It will also attempt to update each of them over again, but normally this will not change them again, since they are already up to date.)"

This feature can be very useful for writing makefiles that automatically generate and read dependency information with each build. However, this feature can cause GNU Make to loop infinitely if the rule to generate a makefile is always out-of-date:

```
all:
        @echo $@
makefile: force
        @echo "# last updated: 'date'" >> $@
force:
```

In practice, a well-written makefile will not have out-of-date rules that cause it to regenerate. The same problem, however, can occur when Make detects a clock skew—most commonly due to clock drift between the system running Make and the file server hosting the current directory. In this case, Make continues to loop until the rule to rebuild the makefile is no longer out-of-date.

In the example below, `DIR1` and `DIR2` are both part of the source tree:

```
-include $(DIR1)/foo.dd
all:
        @echo $@
$(DIR1)/foo.dd: $(DIR2)/bar.dd
%.d:
        touch $@
```

If two directories are served by different file servers and the clock on the system hosting `DIR2` is slightly faster than `DIR1`, then even though `foo.dd` is updated after `bar.dd`, it might appear to be older. On remaking, GNU Make will again see `foo.dd` as out-of-date and restart, continuing until the drift is unnoticeable.

This problem is particularly troublesome on Solaris, where GNU Make timestamp checking has nanosecond resolution:

```
% make
touch dir2/bar.dd
touch dir1/foo.dd
gmake: *** Warning: File 'bar.dd' has modification time in the future (2005-04-1
1 13:52:46.811724 > 2005-04-11 13:52:46.799573198)
touch dir1/foo.dd
touch dir1/foo.dd
all
```

Electric Make fully supports makefile remaking and can be configured to behave exactly as GNU Make. However, by default, to ensure builds do not loop unnecessarily while remaking, Electric Make limits the number of times it restarts a make instance to 10. If your build is looping unnecessarily, you might want to lower this value or disable remaking entirely by setting:

```
--emake-remake-limit=0
```

# NMAKE Inline File Locations (Windows only)

NMAKE contains a feature to create *inline files* with temporary file names. For example, the following makefile creates a temporary inline file containing the word "pass" and then uses "type" to output it.

```
all:
                type <<
        pass
        <<
```

With Microsoft NMAKE, the file is created in the directory where `%TMP%` points. Electric Make does not respect the `%TMP%` setting and creates the inline file in the rule working directory that needs the file.

# How eMake Processes MAKEFLAGS

eMake uses the following process:

1. Similar to GNU Make, eMake condenses no-value options into one block.

2. When eMake encounters an option with a value, it does what GNU Make does, it appends the value and starts the next option with its own –

3. Certain options are ignored/not created. This changes the layout of the options in MAKEFLAGS (for example $-j$, $-l$).

4. eMake-specific options are not added to MAKEFLAGS, but are handled through EMAKEFLAGS.

5. Passing down environment variables as TEST=test renders the same result as in GNU Make (an extra – at the end, followed by the variable=value).

6. On Windows, eMake prepends the --unix or --win32 flag explicitly.

# Chapter 6: Performance Optimization

The following topics discuss how to use ElectricAccelerator performance optimization features and performance tuning techniques.

**Topics:**

- Optimizing Android Build Performance

- Dependency Optimization

- Job Caching

- Parse Avoidance

- Javadoc Caching

- Schedule Optimization

- Running a Local Job on the Make Machine

- Serializing All Make Instance Jobs

- Managing Temporary Files

# Optimizing Android Build Performance

## Recommended Deployment/Sizing

For optimal Android build performance, make sure you have at least 48 ElectricAccelerator agents available.

## Recommended Steps

Electric Cloud recommends the following actions:

1. Upgrade your ElectricAccelerator deployment to v7.0 or later (procedure available in the *ElectricAccelerator Installation Guide*).

2. Configure a clean build environment.

3. To generate a correct history-file and record all the actual dependencies in the build, ElectricAccelerator needs to run a "learning" build.

   Sample command:

   ```
   emake \
     --emake-cm=$YOUR_CLUSTER_MANAGER \
     --emake-root=$YOUR_ROOTS \
     --emake-annofile=$YOUR_ANNOTATION_FILE \
     --emake-annodetail=basic \
     --emake-logfile=$YOUR_LOG_FILE_NAME \
     --emake-debug=g \
     --emake-optimize-deps=1 \
     --emake-parse-avoidance=1 \
     --emake-autodepend=1 \
     --emake-suppress-include=*.d \
     --emake-suppress-include=*.P \
     -f Makefile \
     -f noautodep.mk
   ```

   `noautodep.mk` contains "hints" used to tailor dependency management for better performance with Android builds. The contents are as follows:

   ```
   #pragma noautodep */.git/*
   $(local-intermediates-dir)/libbcc-stamp.c :

   #pragma noautodep */out/target/product/generic/system/bin/cat
   $(linked_module) :

   #pragma nolookup noproguard.classes-with-local.dex
   #pragma nolookup noproguard.classes.dex
   #pragma nolookup javalib.jar
   #pragma nolookup classes-full-debug.jar
   out/target/common/docs/doc-comment-check-timestamp :
   ```

4. Make "clean" and then re-run the build. ElectricAccelerator will use the previous learning build to optimize performance.

   Run the same sample command from step 3.

   **Note:** If you want to find out why a cached result was not used, add "P" debugging to `--emake-debug`.

# Dependency Optimization

ElectricAccelerator includes a dependency optimization feature to improve performance. By learning which dependencies are actually needed for a build, Accelerator can use that information to improve performance in subsequent builds.

When dependency optimization is enabled, eMake maintains a dependency information file for each makefile. If a build's dependencies have not changed from its previous build, eMake can use that stored dependency information file for subsequent builds.

## Enabling Dependency Optimization

You must first run a "learning" build with the dependency optimization feature enabled. To enable dependency optimization, set the following: `--emake-optimize-deps=1`

For the learning build, (because there is no stored dependency information file for that specific makefile) the argument only saves a new result. For subsequent builds, the argument enables the reuse of stored dependency information and saves a new file as appropriate. If you do not specify `--emake-optimize-deps=1`, then dependency information is not saved or accessed.

The following table describes dependency optimization-related options.

| eMake Option | Description |
| --- | --- |
| `--emake-assetdir=<path>` | Use the specified directory for assets such as saved dependency information. The default directory is named .emake. (This option also determines the cache locations for the parse avoidance feature and for the JobCache add-on option.) |
| `--emake-optimize-deps=<0/1>` | Use the saved dependency information file for a makefile when dependencies are the same and save new dependency information when appropriate. |

## Dependency Optimization File Location and Naming

Dependency information files are saved in the working directory where eMake was invoked under `<assetdir>/deps`. (The default asset directory is `.emake`.)

The file is named from the md5 hash of the root-relative path of the "main" makefile for the make instance that the dependency information file belongs to.

### Examples

If your eMake root is `/tmp/src`, and your makefile is `/tmp/src/Makefile`, then the dependency information file is named after the md5 hash of the string "Makefile".

If your root is `/tmp/src` and your makefile is `/tmp/src/foo/Makefile`, the file is named after the md5 hash of the string "foo/Makefile".

# Job Caching

JobCache is a feature that can substantially reduce compilation time. JobCache lets a build avoid recompiling object files that it previously built, if their inputs have not changed. JobCache works even after you clean the build output tree (for example, by using "make clean"). By caching and reusing object files, JobCache can significantly speed up full builds.

For ElectricAccelerator (enterprise), JobCache is an add-on option and requires a separately-purchased license. The license can be used for Linux, Solaris, and Windows.

For ElectricAccelerator Huddle, JobCache is available only on Linux platforms. JobCache is included with Huddle and uses the same license as Huddle.

JobCache uses cache "slots." When JobCache is enabled, eMake maintains a slot for each combination of command line options, relevant environment variable assignments, and current working directory. A slot can be empty or can hold a previously-cached result. If the appropriate slot holds an up-to-date result, a cache "hit" occurs, and compilation is avoided.

A cached result becomes obsolete if eMake detects file system changes that might have caused a different result (with the same command line options, environment variable assignments, and current working directory). Such file system changes include any files that are read during compilation, which means all source files, gcc precompiled headers, and compilation tools included in the eMake root.

Electric Cloud recommends that all components (Cluster Manager, Electric Agent/EFS, and Electric Make) on all machines in the cluster are upgraded to the latest version. However, you can still use JobCache with Agents on Electric Agent/EFS machines running versions as old as 7.2, as long as an appropriate backward compatibility package (BCP) is installed on those Electric Agent/EFS machines. For more information, see the *ElectricAccelerator Installation and Configuration Guide* at http://docs.electric-cloud.com/accelerator_doc/AcceleratorIndex.html.

## Benefits

- Speeds long, full builds (for example, when you do a "make clean" then a "make," or when you run a build in a new workspace)

- Builds faster than ccache

- Avoids certain false cache hits that might occur when you use ccache

## Limitations

- Supports gcc, g++, and cl

  - `-emake-jobcache=gcc`: turns on caching for all .o files.

    - Assumes use of only gcc and/or g++.

    - Should not be used to cache results from linking.

    - Is supported on Linux, Solaris, and Windows.

  - `-emake-jobcache=cl`: turns on caching for all .obj files.

    - Assumes use of only cl.exe.

    - Only supported debug option is /Z7

- Does not cache results from compilation jobs that invoke eMake

  - Date and time in an object file will still reflect the original compilation, not the current date or time, if you use a C preprocessor macro that expands to the date or time when the compiler runs, and ElectricAccelerator re-uses the resulting object file in a subsequent build

  - Source paths embedded in debugging information in object files will reflect the original compilation (even if you re-use the object file while building in a different workspace)

## Configuring JobCache

### *Licensing*

For ElectricAccelerator, JobCache is licensed based on the maximum number of builds that may use it simultaneously. This number is read from the `jobcacheMax` property in the Accelerator license file. Simultaneous builds that exceed this number occur without using this feature. For ElectricAccelerator Huddle, JobCache has no limit on the maximum number of builds that may use it simultaneously.

If the JobCache license entry is invalid, or if the number of simultaneous builds has exceeded the license limit, a `WARNING EC1181: Your license does not permit GCC object caching` message appears when a build tries to use JobCache. eMake will continue to work normally.

### Choosing a Disk for the Job Cache

To estimate the disk space required for the job cache, add the sizes of your object files together and multiply by 0.7. A specific example is that the Android KitKat Open Source Project (AOSP) requires about 4 GB. For best performance, choose a disk that is local to the eMake client host.

You can use the `--emake-assetdir=` eMake option to specify the directory for your job cache. The default name of this directory is .emake. By default, this directory is in the working directory in which eMake is invoked. (This option also determines the cache location for the parse avoidance feature and the location of the saved dependency information for the dependency optimization feature.)

### Building Multiple Branches

To maximize your cache hits, use the `--emake-assetdir=` option to specify a separate asset directory for each branch of code that you build.

### Setting the eMake Root

JobCache does not detect changes to compilation inputs that are not under your eMake root. You must ensure that your eMake root contains all sources and tools that might change.

## gcc Job Caching

### Enabling gcc JobCache for All Make Invocations in a Build

To enable JobCache for all make invocations in a build, use the `--emake-jobcache=gcc` eMake option. This option works for recursive and nonrecursive builds. Following is an example command that includes this option:

```
% emake --emake-cm=mycm --emake-root=/src/mysource --emake-jobcache=gcc --emake-ann
odetail=basic
```

If some of your ".o" targets are built by rules that do not invoke gcc or g++, and those rules do not behave enough like gcc or g++ for job caching to be suitable, then use the `jobcache` pragma to be more selective about which jobs are cached, either by enabling JobCache for fewer jobs or by selectively disabling it for particular jobs with `#pragma jobcache none`.

For example, the rules that build a particular ".o" file might use an environment variable that is not used by gcc or g++, and you might want to miss the cache if that environment variable changes its value.

The `--emake-annodetail=basic` option is not required to invoke eMake, but it is recommended for troubleshooting. For details, see the "Troubleshooting" section below.

### Using the Response File gcc Command Line Option

JobCache supports the response file (`@<file>`) gcc command line option, which reads gcc command-line options from a separate file specified by *<file>*.

### Enabling gcc JobCache for All Object File Targets in a Make Invocation

If you do not want to enable JobCache for all make invocations in a build, you can still enable JobCache for specific object file targets by applying the jobcache pragma inside the makefile. Following is an example makefile excerpt with the feature enabled for all targets ending in ".o", whether you compile them with gcc or g++:

```
...

#pragma jobcache gcc -exist *

%.o :

...
```

(You can apply the pragma to a pattern, but not to a suffix rule.) If cache misses occur because an insignificant file exists or does not exist, you can delete `-exist *`, and eMake will check only the existence of files whose names suggest that they are source files or gcc precompiled headers. You can augment the set of files that matters to eMake by adding `-exist` options that specify the desired glob patterns.

If you do not want to add the pragma to the main makefile, you can add it to an "addendum" makefile. This is a small makefile that you include in the eMake invocation from the command line. For example, you could create a makefile named `jobcache.mak` and then add `-f Makefile -f jobcache.mak` to your eMake invocation. An addendum makefile is useful when you do not control the content of your makefiles (such as when you include open source components that use build tools such as configure, CMake, or qmake to generate makefiles).

### *Enabling or Disabling JobCache for Specific Object File Targets*

To enable JobCache only for some object file targets, do one of the following:

- If an object file is built by an explicit rule (one without "%" in it), you can enable JobCache for that rule by inserting a pragma just above it in the makefile. You can also use an addendum makefile in this case. For example, to cache an explicit rule producing target.o, you can use this addendum:

        #pragma jobcache gcc -exist *

        target.o :

- You can use more narrow target patterns than "%.o" after a `jobcache` pragma. For example, "abc%.o" applies the pragma to "abcdef.o" but not to "xyzdef.o". (This is true even if the rule to build "abcdef.o" appears elsewhere in the makefile.)

Similarly, you can explicitly disable caching by using the `#pragma jobcache none` pragma.

If `jobcache` pragmas with different options apply to a target, then eMake selects one of them as follows:

- If the target is built by an explicit rule with a `jobcache` pragma, then eMake chooses that pragma.

- Otherwise, eMake chooses the pragma whose pattern has the most in common with the target name.

- If eMake still encounters ties, then it chooses the pragma that it encountered last during makefile parsing.

For example, the following makefile applies job caching to hello1.o and hello2.o, but not to hello3.o or date2.o:

        CC=gcc


        all : hello

                ./hello


        #pragma jobcache gcc -exist *

        hello1.o: hello1.c hello1.h

                $(CC) -c -o $@ $< $(CFLAGS)


        hello2.o: hello2.c hello2.h

                $(CC) -c -o $@ $< $(CFLAGS)

        # Uses __DATE__ to record when build occurred--do not cache.

        #pragma jobcache none

        date2.o: date2.c date2.h

```
        $(CC) -c -o $@ $< $(CFLAGS)


hello3.o: hello3.c hello3.h

        $(CC) -c -o $@ $< $(CFLAGS)


hello: hello1.o hello2.o hello3.o

        $(CC) -o $@ $^


clean:

        rm hello hello1.o hello2.o hello3.o


.PHONY: all clean


#pragma jobcache gcc -exist *

%2.o :
```

# cl Job Caching

### Enabling cl JobCache for All Make Invocations in a Build

On Windows, setting `-emake-jobcache=cl` will enable jobcache for all jobs that produce `.obj` files. These are normally created by the VC++ compiler (`cl.exe`). The option may be set on the command line, in an `emake.conf` file or in the `EMAKEFLAGS` environment variable.

When using the Visual Studio IDE add-in, add the `-emake-jobcache=cl` option to EMake Options in the command line settings.

For example, if you create a simple C++ Console Application using Visual Studio, it creates two compile jobs for `Example1.cpp` and `stdafx.cpp`



```
The Visual Studio IDE add-in converts this project into a makefile:
".\Example1\Debug\Example1.obj":: ".\Example1\Example1.cpp"
```

```
        @cl.exe /Od /Oy- /sdl /D WIN32 /D _DEBUG /D _CONSOLE /D _LIB /D _UNICODE /D UNICOD
E /EHsc /RTC1 /analyze- /MDd /GS /Zc:wchar_t /Zc:forScope /Yu"stdafx.h" /Fp".\Examp
le1\Debug\Example1.pch" /Fo".\Example1\Debug\Example1.obj" /W3 /WX- /nologo /c /Z7
/Gd /errorReport:none /fp:precise /TP ".\Example1\Example1.cpp"


".\Example1\Debug\Example1.pch" :: ".\Example1\Debug\stdafx.obj"


".\Example1\Debug\stdafx.obj" :: ".\Example1\stdafx.cpp"

        @cl.exe /Od /Oy- /sdl /D WIN32 /D _DEBUG /D _CONSOLE /D _LIB /D _UNICODE /D UNICOD
E /EHsc /RTC1 /analyze- /MDd /GS /Zc:wchar_t /Zc:forScope /Yc"stdafx.h" /Fp".\Examp
le1\Debug\Example1.pch" /Fo".\Example1\Debug\stdafx.obj" /W3 /WX- /nologo /c /Z7 /G
d /errorReport:none /fp:precise /TP ".\Example1\stdafx.cpp"
```

When `emake` runs with jobcache, it will store `Example1.obj` and `stdafx.obj` and reuse them on the next build if it gets a 'hit' (i.e., no relevant changes are detected since the last build).

> **IMPORTANT:** JobCache is not compatible with the compiler debug options `/Zi` or `/ZI` when updating same PDB file between different compilations. For example, each compilation must create a unique PDB file if debug is enabled. It is recommended that `/Z7` is used with `jobcache`. When using the Visual Studio IDE add-in, check 'Set Debug Information Format to C7 Compatible' to ensure `/Z7` is used:



## Running a "Learning" Build to Populate the Cache

You must first populate the cache by running a "learning" build with JobCache enabled. For the learning build (because the cache is empty), JobCache saves only a new result to the cache. For subsequent builds,

JobCache re-uses cached results and saves a new result to the cache as appropriate. If you do not enable JobCache, then the job cache is not accessed.

# Troubleshooting

## *Enabling Basic Annotation*

Basic annotation helps you to resolve your cache hit/miss problems by providing information about whether a cache miss occurred and the cause. Annotation that is related to JobCache is included in basic annotation. Basic annotation is not enabled by default, so you must do so by using the `--emake-annodetail=basic` eMake option. By default, the eMake annotation file is created in the working directory in which eMake is invoked; the file is named emake.xml by default.

## *Interpreting Job Cache Annotation Information*

Examine the relevant `job` XML element in the annotation file for a subelement named `jobcache`.

Following is an example of a `jobcache` subelement for a job with a job cache hit. The `status` attribute indicates whether eMake used a cached object file in that job (whether it had a cache hit), and if not, what else occurred:

```
...
<job id=...>
...
    <jobcache type="gcc" options=" -exist *" slot="6d00a0d9242610a075a98bc9400f8f1
1" duration="0.155831" status="hit">
...
    </jobcache>
...
</job>
...
```

The `duration` attribute is the duration (in seconds) of the job that populated the cache. If the current build is updating the cache, then the value of `duration` is the same as the duration of the job containing that `jobcache` subelement. If the current build is replaying from the cache, then it is the duration of the corresponding job from a previous build whose results were saved into the cache. eMake retrieves the figure from the cache.

Following is an example of a `jobcache` subelement for a job with a cache miss:

```
...
<job id=...>
...
    <jobcache type="gcc" options=" -exist *" slot="989324a736c583a306630930da80ba1
e" duration="0.266177" status="miss">
...        <differences>
                <diff name="/c/src/mysql-5.6.21/include/mysql.h" old="md5:d55ae58fd7
eec6055a50fcf6b83af99c" new="md5:071ed8e989b0d0e14d0b0bd96e94cd35"/>
            </differences>
        </jobcache>
...
</job>
...
```

Following is an example of a `jobcache` subelement for a job for which caching was disabled by using `#pragma jobcache none`:

```
...
<job id=...>
...
     <jobcache type="none" options="" status="uncacheable">
          <triggers>
               <trigger type="pragma" file="Makefile" line="11" options="none"/>
          </triggers>
     </jobcache>
...
</job>
...
```

Following is a complete list of the possible values of the `status` attribute:

- `hit`—eMake had a cache hit for that job.

- `miss`—eMake had a cache miss for that job. Each `diff` subelement shows a file system input to compilation that differs since the object file was cached and shows the difference that was observed.

- `newslot`—A new slot was created. This is because the previously-cached object files were from compilations that used different command-line arguments, environment variables, working directories, or any combination of these. To see the options used for that compilation, see the "key" file for the slot identified by the `slot` attribute. For example, the key file for slot 2b253a890c9745a0b500d888349ec2e2 has the following path name:

  `.emake/cache.16/i686_Linux/2b/25/3a/890c9745a0b500d888349ec2e2/key`

  The version number in the cache.*<number>* directory might vary with your ElectricAccelerator software version. The key file specifies the relevant environment variables, the working directory, and the command line. To allow cache hits when building in a new workspace, path names are specified relative to your eMake roots. If a target repeatedly has `newslot` status, get the `slot` identifiers from two consecutive builds for that same target and compare the key files.

- `uncacheable`—Caching was disabled by the `#pragma jobcache none` pragma, or an error occurred during the update of the relevant cache slot. Examine any `ERROR` and `WARNING` messages in the console output from eMake.

- `rootschanged`—There is no natural mapping from the old eMake roots to the new eMake roots.

- `unneeded`—JobCache was enabled for the job but not needed (because the target was already up to date according to ordinary GNU Make rules). The cache was not consulted, even though caching was requested for that target.

## *If Job Cache Annotation Information Is Missing*

If a particular job element in the annotation file has no XML `jobcache` subelement, this is because any combination of the following has occurred:

- The target name for that job does not match the pattern following any `jobcache` pragma, and if the target is built by an explicit rule, that rule does not follow a `jobcache` pragma.

- The intended `jobcache` pragma is misspelled.

- Either the `--emake-jobcache=` command line option was not used, or none of the target names ends in ".o".

- The appropriate licensing is not available to the eMake client.

### *If eMake Unexpectedly Used a Cached Object File*

If eMake should not have used a particular cached object file, then

- If eMake should have detected a change to a particular file, compare its path to your eMake roots.

- If a relevant environment variable changed, check whether the key file mentions it (see above), and if it does not, notify Electric Cloud.

### *Enabling Profiling Debug Logging*

To help troubleshoot performance issues, you should enable profiling debug logging. To do so, include a letter "g" in the argument of the `--emake-debug=<arguments>` option. Electric Cloud engineering and support staff use the eMake debug logs to help troubleshoot problems. For more information about debug logging and log levels, see the "Electric Make Debug Log Levels" section in Chapter 11,Troubleshooting.

## Viewing JobCache Metrics

The annotation file includes metrics about job cache activity. Following is an example that lists the metrics:

```
...
<metrics>
...
     <metric name="jobcache.na">923</metric>
     <metric name="jobcache.hit">679</metric>
     <metric name="jobcache.miss">3</metric>
     <metric name="jobcache.newslot">1</metric>
     <metric name="jobcache.rootschanged">0</metric>
     <metric name="jobcache.uncacheable">0</metric>
     <metric name="jobcache.unneeded">0</metric>
     <metric name="jobcache.workloadsaved">341.441121</metric>
...
</metrics>
...
```

For descriptions of these metrics, see the "Metrics in Annotation Files" section in Chapter 8, Annotation.

## Moving Your Workspace

If you want to move your workspace, make sure that the new eMake roots correspond to the old eMake roots. Also, because the asset directory defaults to .emake in the current working directory, you must either copy that directory to the new workspace or use `--emake-assetdir=` to specify an asset directory that you want the two workspaces to share. If you already use `--emake-assetdir=` to point to an asset directory within your old workspace and also want to move the asset directory, you must update its value to point to the new asset directory location.

## Deleting the Cache

In general, content in the cache is not deleted automatically (although it might be replaced by newer content). If the cache grows significantly beyond the size expected for a full build, you can delete the cache to save disk space.

For example, if you change the value of the `C_INCLUDE_PATH` environment variable, then the cache will grow to contain results for both the old and new values of that variable. In this case, you might want to clear the cache when you permanently change the value of this variable and therefore no longer need the old cache results.

To delete the cache, you delete the `<assetdir>/cache.*` directories. If you are using the default asset directory, enter

```
rm -r .emake/cache.*
```

# Parse Avoidance

ElectricAccelerator includes a parse avoidance feature that can almost eliminate makefile parse time. By caching and reusing parse result files, Accelerator can speed up both full builds and incremental builds.

Parse avoidance works when emulating GNU Make with clusters and/or local agents.

Parse avoidance uses the concept of cache "slots." When parse avoidance is enabled, eMake maintains a slot for each combination of non-ignored command line options, non-ignored environment variable assignments, and current working directory. (Ignored arguments and environment variables are listed here.) A slot can be empty or hold a previously cached result. If the appropriate slot holds an up-to-date result, parsing is avoided.

A cached result becomes obsolete if eMake detects file system changes that might have caused a different result (with the same command line options, environment variable assignments, and current working directory). Such file system changes include any file read by the parse job, which means all makefiles, all programs invoked by $(shell) during the parse (as opposed to during rule execution), the files they read, and so on.

### Console Output

When a cached parse result is reused, eMake replays the file system modifications made during the parse job and any standard output or standard error that the original parse job produced. For example, if the makefile includes

```
VARIABLE1 := $(shell echo Text > myfile)

VARIABLE2 := $(warning Warning: updated myfile)
```

then when using a cached parse result, eMake creates "myfile" and prints "Warning: updated myfile".

## Enabling Parse Avoidance

You must first run a "learning" build with the parse avoidance feature enabled. To enable parse avoidance, set the following: `--emake-parse-avoidance=1`

For the learning build, (because the parse cache is empty), the argument only saves a new result to the cache. For subsequent builds, the argument enables the reuse of cached parse results and saves a new result to the cache as appropriate. If you do not specify `--emake-parse-avoidance=1`, then the parse avoidance cache is not accessed at all.

**Note:** You should also disable generated dependencies (either by modifying your makefiles or by using `--emake-suppress-include`; see below)

The following table describes parse avoidance-related options.

| eMake Option | Description |
|---|---|
| `--emake-assetdir=<path>` | Use the specified directory for cached parse results. The default directory is named .emake. (This option also determines the location of the saved dependency information for the dependency optimization feature and the cache location for the JobCache add-on option.) |
| `--emake-parse-avoidance=<0/1>` | Avoid parsing makefiles when prior parse results remain up-to-date and cache new parse results when appropriate. |
| `--emake-parse-avoidance-ignore-env=<var>` | Ignore the named environment variable when searching for applicable cached parse results. To ignore more than one variable, use this option multiple times. |

| eMake Option | Description |
| --- | --- |
| `--emake-parse-avoidance-ignore-path=<path>` | Ignore this file or directory when checking whether cached parse results are up-to-date. **Append** % for prefix matching (the % must be the last character). To ignore more than one path or prefix, use this option multiple times. |
| | Incorrect placement of `%` will result in an error. |
| | Correct: `--emake-parse-avoidance-ignore-path=.foo%` |
| | Incorrect: `--emake-parse-avoidance-ignore-path=.%foo` |
| `--emake-suppress-include=<pattern>` | Skip matching makefile includes (such as generated dependencies). Generally, you should not suppress makefile includes unless they are generated dependency files, and you have enabled automatic dependencies as an alternative way of handling dependencies. |
| | **Note:** If the pattern does not have a directory separator, then the pattern is compared to the include's filename component only. If the pattern has a directory separator, then the pattern is taken relative to the same working directory that applies to the include directive and compared to the included file's entire path. |

If a file is read during a parse but changes before eMake attempts to reuse that parse's results, the cached parse result is normally considered to be obsolete. You can, however, temporarily override this decision with `--emake-parse-avoidance-ignore-path`.

eMake permanently ignores the effect on a given parse result of certain special files that might have existed when it was created, in all cases using the names they would have had at that time:

- Anything in ".emake" or whatever alternative directory you specified using `--emake-assetdir`
- Client-side eMake debug log (as specified by `--emake-logfile`)
- Annotation file (as specified by `--emake-annofile`)
- History file (as specified by `--emake-historyfile`)

**Notes:**

- It is recommend that your makefiles avoid any $(shell) expansion that uses these four special files and directories in a meaningful way, because parse avoidance will ignore changes to them.
- Enabling remote parse debugging (the `--emake-rdebug` option) disables parse avoidance.

You can use a pragma to instruct parse avoidance to detect the dependence of a parse result upon path wildcard match results. This pragma makes the parse cache sensitive to the existence (or nonexistence) of specific files in directories that it reads as well as in new subdirectories thereof. (eMake does not detect matching files in just any new subdirectory, but only those that are subdirectories of directories read in the original job and their subdirectories, recursively.)

You can specify more than one glob pattern, as in the following pragma:

```
#pragma jobcache parse -readdir *.h -readdir *.c
```

For Android-specific information about parse avoidance and other Android best practices, see KB article KBEA-00130 at https://electriccloud.zendesk.com/entries/23213988-KBEA-00130-Best-practices-for-Android-builds.

### *Parse Avoidance Example*

`noautodep.mk` is used for this example. The contents are:

```
#pragma noautodep */.git/*
$(local-intermediates-dir)/libbcc-stamp.c :

#pragma noautodep */out/target/product/generic/system/bin/cat
$(linked_module) :
```

The following parse avoidance example provides command line arguments for Android 4.1.1.

**Note:** You might need additional options such as `--emake-cm`

```
--emake-parse-avoidance=1 --emake-autodepend=1 --emake-suppress-include=*.d --emak
e-suppress-include=*.P --emake-debug=P -f Makefile -f noautodep.mk
```

## Deleting the Cache

To delete the cache, delete `<assetdir>/cache.*`. (The default asset directory is `.emake`.)

For example: `rm -r .emake/cache.*`

## Moving Your Workspace

If you want to move your workspace, make sure that the new eMake roots correspond to the old eMake roots. Also, because the asset directory defaults to .emake in the current working directory, you must either copy that directory to the new workspace or use `--emake-assetdir=` to specify an asset directory that you want the two workspaces to share. If you already use `--emake-assetdir=` to point to an asset directory within your old workspace and also want to move the asset directory, you must update its value to point to the new asset directory location.

eMake looks for eMake roots in the values of Make variables. When eMake replays a cached parse result in a new workspace, it replaces the old eMake roots in the values of those variables with the new eMake roots. This policy works well as long as no confusion exists between eMake roots and other text in those variable values. For example, if the value of a Make variable is `-I/we will look into this`, your old eMake root is /we, and your new eMake root is /wg, then the new value will be `-I/wg will look into this`. For best results, choose distinctive directory names for your workspaces.

## Limitations

- Windows is currently not supported.

- The parse avoidance feature does not detect the need to reparse in these cases:

  ○ Changes to input files and programs that are used during the parse (such as makefile includes) but are not virtualized (because they are not located under `--emake-root=...`).

  ○ Changes to non file system, non-registry aspects of the environment, for example, the current time of day

- If a parse job checks the existence or timestamp of a file without reading it, parse avoidance might not invalidate its cached result when the file is created, destroyed, or modified.

  ○ Using `--emake-autodepend=1` and `--emake-suppress-include=<pattern>` in conjunction with parse avoidance helps to avoid this limitation. If a generated dependency file did not exist when a parse result was saved into the cache, eMake might reuse that cached parse result even after that dependency file was created. Of course, you also benefit from eDepend's performance and reliability gains.

○ This limitation can be mitigated by the use of `#pragma jobcache parse -readdir ...`, which reruns the parse job after files are created or deleted that match the pattern and are located in a directory that was subject to a wildcard match or $(shell find ...) when the cache was populated.

# Troubleshooting

### Enabling Parse Avoidance Debug Logging

To help troubleshoot performance issues, you should enable parse avoidance debug logging. To do so, include a capital letter "P" in the argument of the `--emake-debug=<arguments>` option. Electric Cloud engineering and support staff use the eMake debug logs to help troubleshoot problems. For more information about debug logging and log levels, see the "Electric Make Debug Log Levels" section in Chapter 11, Troubleshooting.

### Debug Log Example

You can look in "P" debug logging for an explanation of why a cached result was found to be obsolete:

```
WK01: 0.062173 Input changed: job:J08345218 slot:b6189634a793fee2fe1929fbf47cc4e4 p
ath:/home/aeolus/t/Makefile thenSize:175 nowSize:176 ignore:0
WK01: 0.062228 Input changed: job:J08345218 slot:b6189634a793fee2fe1929fbf47cc4e4 p
ath:/home/aeolus/t/fog.mk thenSize:0 nowSize:1 ignore:0
...
WK01: 0.062284 Cache slot is obsolete: job:J08345218 slot:b6189634a793fee2fe1929fbf
47cc4e4
```

Notice that "`Makefile`" and "`fog.mk`" were both bigger; either of those changes would have triggered a reparse. The "ignore:0" comments indicate that `--emake-parse-avoidance-ignore-path` was not used to ignore the changes.

### Using Key Files for Debugging

To discover why a new cache slot was used, look in "P" debug logging for the name of the "key" file for that new cache slot; for example:

```
WK01: 11.853035 Saved slot definition: .emake/cache.11/i686_Linux/d9/0f/79/4fb975e8
ff94dfa2569a303e62.new.2NPR8J/key
```

**Note:** The ".new.2NPR8J" portion of the slot directory name should have already been removed automatically by eMake before you need to access the key file.

Then diff that key file against other key files in sibling directories to see what parameters changed. To determine which slot directories to compare, grep for the parse job IDs in "P" debug logging. You can get parse job IDs from annotation using ElectricInsight. Each key file contains an artificial "cd" command to express the working directory, all non-ignored environment variables, and the non-ignored command-line arguments. Ignored command-line arguments and environment variables are omitted from key files. The key file is stored in a directory whose pathname is derived from the md5sum of the file itself. Also, the command line might have been normalized to one that is equivalent to the original one. Currently, eMake quotes special characters according to UNIX/Linux/Cygwin "sh" shell rules.

# Ignored Arguments and Environment Variables

The following arguments do not affect which cache slot is chosen:

```
--emake-annodetail                    --emake-ledger
--emake-annofile                      --emake-ledgerfile
--emake-annoupload                    --emake-localagents
--emake-assetdir                      --emake-logfile
--emake-big-file-size                 --emake-logfile-mode
--emake-build-label                   --emake-maxagents
--emake-class                         --emake-maxlocalagents
--emake-cluster-timeout               --emake-mem-limit
--emake-cm                            --emake-parse-avoidance-ignore-path
--emake-debug                         --emake-pedantic
--emake-hide-warning                  --emake-priority
--emake-history                       --emake-readdir-conflicts
--emake-history-force                 --emake-resource
--emake-historyfile                   --emake-showinfo
--emake-idle-time                     --emake-tmpdir
--emake-job-limit                     --emake-yield-localagents
```

The following default set of environment variables do not affect which cache slot is chosen. You can specify additional environment variables with `--emake-parse-avoidance-ignore-env`. You must use `--emake-parse-avoidance-ignore-env` once for each variable to ignore.

_
BUILD_DISPLAY_NAME
BUILD_ID
BUILD_NUMBER
BUILD_TAG
CI_JENKINS_BUILD_ID
CI_JENKINS_BUILD_NUMBER
CI_JENKINS_BUILD_TAG
CI_JENKINS_HUDSON_SERVER_COOKIE
CI_JENKINS_JENKINS_SERVER_COOKIE
COLORFGBG
COLUMNS
COMMANDER_JOBID
COMMANDER_JOBSTEPID
COMMANDER_RESOURCENAME
COMMANDER_SESSIONID
COMMANDER_WORKSPACE
COMMANDER_WORKSPACE_UNIX
COMMANDER_WORKSPACE_WINDRIVE
COMMANDER_WORKSPACE_WINUNC
DBUS_SESSION_BUS_ADDRESS
DESKTOP_SESSION
DISPLAY
EMAKE_APP_VERSION
EMAKE_BUILD_MODE
EMAKE_LEDGER_CONTEXT
EMAKE_MAKE_IDEMAKE_PARSEFILE
EMAKE_PARSEFILE
EMAKE_RLOGFILE
ECLOUD_AGENT_NUM
ECLOUD_BUILD_CLASS
ECLOUD_BUILD_COUNTER
ECLOUD_BUILD_ID
ECLOUD_BUILD_TAG
ECLOUD_BUILD_TYPE
ECLOUD_RECURSIVE_COMMAND_FILE
GPG_AGENT_INFO
HOSTNAME
KDE_FULL_SESSION
KDE_MULTIHEAD
KDE_SESSION_UID
KDE_SESSION_VERSION
KONSOLE_DBUS_SERVICE

KONSOLE_DBUS_SESSION
LINES
LS_COLORS
OLDPWD
SESSION_MANAGER
SHELL_SESSION_ID
SHLVL
SSH_AGENT_PID
SSH_AUTH_SOCK
SSH_CLIENT
SSH_CONNECTION
SSH_TTY
TEMP
TERM
TERMCAP
TMP
TMPDIR
WINDOWID
WINDOWPATH
WRAPPER_ARCH
WRAPPER_BIN_DIR
WRAPPER_BITS
WRAPPER_CONF_DIR
WRAPPER_DESCRIPTION
WRAPPER_DISPLAYNAME
WRAPPER_EVENT_NAME
WRAPPER_EVENT_WRAPPER_PID
WRAPPER_FILE_SEPARATOR
WRAPPER_HOSTNAME
WRAPPER_INIT_DIR
WRAPPER_JAVA_HOME
WRAPPER_EVENT_JVM_ID
WRAPPER_EVENT_JVM_PID
WRAPPER_LANG
WRAPPER_NAME
WRAPPER_OS
WRAPPER_PATH_SEPARATOR
WRAPPER_PID
WRAPPER_SYSMEM_PP.P
WRAPPER_WORKING_DIR
XCURSOR_THEME
XDG_SESSION_COOKIE
XDM_MANAGED

```
WRAPPER_ARCH
WRAPPER_BIN_DIR
WRAPPER_BITS
WRAPPER_CONF_DIR
WRAPPER_DESCRIPTION
WRAPPER_DISPLAYNAME
WRAPPER_EVENT_NAME
WRAPPER_EVENT_WRAPPER_PID
WRAPPER_FILE_SEPARATOR
WRAPPER_HOSTNAME
WRAPPER_INIT_DIR
WRAPPER_JAVA_HOME
WRAPPER_EVENT_JVM_ID
WRAPPER_EVENT_JVM_PID
WRAPPER_LANG
WRAPPER_NAME
WRAPPER_OS
WRAPPER_PATH_SEPARATOR
WRAPPER_PID
WRAPPER_SYSMEM_PP.P
WRAPPER_WORKING_DIR
```

# Javadoc Caching

The ElectricAccelerator Javadoc caching feature can improve build performance by caching and reusing Javadoc files.

## Enabling Javadoc Caching

To enable Javadoc caching for a specific rule, add the following before that rule:

```
#pragma jobcache javadoc
```

Alternatively, you can mention the target again in a separate makefile:

```
#pragma jobcache javadoc
mytarget :
```

By default, the cache becomes obsolete when `*.java` files are added to or removed from directories that are read by the rule body. You can override this default wildcard pattern by invoking:

```
#pragma jobcache javadoc -readdir pattern
```

*pattern* is the pattern that you want to ignore. You can use multiple wildcard (glob) patterns.

eMake permanently ignores the effect on a given result of certain special files that might have existed when it was created, in all cases using the names they would have had at that time:

- Anything in ".emake" or whatever alternative directory you specified using `--emake-assetdir`

- Client-side eMake debug log (as specified by `--emake-logfile`)

- Annotation file (as specified by `--emake-annofile`)

- History file (as specified by `--emake-historyfile`)

**Note:** -hdf arguments are removed from consideration for which cache slot to use.

If you want to delete the cache, see Deleting the Cache.

For debugging information, see Troubleshooting.

## Limitations

- The following Javadoc invocations are not supported:

  ○ Javadoc is invoked indirectly, such as through a wrapper script

  ○ The rule body recursively invokes eMake

  ○ If Javadoc is named something other than "javadoc"

- Javadoc caching is sensitive to changes in `*.java` wildcard outcomes in directories that are actually read by the job, and any new subdirectories thereof.

- `#pragma jobcache ...` cannot be applied to patterns as such, only to explicitly specified targets (though the commands might be provided by a pattern).

- Windows is not supported.

- The Javadoc caching feature does not detect the need to rebuild Javadoc files in these cases:

  ○ Changes to input files and programs that are used by a rule body but that Accelerator does not virtualize (because they are not located under `--emake-root=...`).

  ○ Changes to non file system, non-registry aspects of the environment, for example, the current time of day

- If a rule job checks the existence or timestamp of a file without reading it, Javadoc caching might not invalidate its cached result when the file is created, destroyed, or modified. However, Javadoc caching will detect when files matching "*.java" (or the pattern specified by "-readdir") are added or removed from directories read by the rule body, or any new subdirectories thereof. In that case eMake will rerun Javadoc.

# Schedule Optimization

## How it Works

Schedule optimization uses performance and dependency information from previous builds to optimize the runtime order of jobs in subsequent builds.

This approach can improve performance on the same number of agents (compared to v7.0, or compared to v7.1 when the requisite information is not available), and enables eMake to realize the "best possible" build performance for a given build with fewer agents.

## Using Schedule Optimization

Schedule optimization is enabled by default, but it has no effect if scheduling data files are not available.

Emake automatically generates the scheduling data file in the asset directory (`.emake` by default) in the `sched` subdirectory under a platform-specific directory such as `Linux`, `SunOS`, or `win32`. For example, on Linux the scheduling data file is found in `.emake/sched/Linux/emake.sched`. The next eMake build that uses the scheduling data file will have improved performance.

You can use the `--emake-assetdir` eMake option to change the location of the asset directory.

## Disabling Schedule Optimization

To disable schedule optimization, add the following to the eMake command line:

```
--emake-optimize-schedule=0
```

# Running a Local Job on the Make Machine

Normally, ElectricAccelerator executes all build commands on the cluster hosts. Any file access by commands is served by in-memory EFS cache or fetched from Electric Make and monitored to ensure dependency information is captured, thus ensuring correct results.

Sometimes it is not desirable to run a command on an Agent (which could be in the cluster or local). For example, the command could exercise such heavy file I/O, creating overhead (running remotely) to significantly degrade performance. In these instances, you can use the `#pragma runlocal` directive. In a makefile, the `#pragma runlocal` directive allows you to specify a job to run locally on the host build machine, instead of remotely.

A run-local job executes on the host build machine where no Electric File System exists to virtualize and monitor file access. Therefore, special care must be taken for build success. The following are two important properties of run-local jobs:

- **Forced serial mode** – Electric Make executes a run-local job when all previous commands are complete, and does not allow other commands to run until the run-local job is complete. In other words, while the run-local job is running, no other build steps are executed.

- **Current directory changes only** – The run-local job must not make any changes outside of its current working directory. If changes occur, subsequent jobs will not see the output outside of the current working directory and could fail—precise interaction is dependent on the Electric Make file system state at the time the job is invoked.

> **IMPORTANT:** After a run-local job, eMake reloads the file system state for the current working directory of the Make, not the job.

**Note:** Marking a job that does not meet this requirement with `#pragma runlocal` is not a supported configuration.

Because of these restrictions, run-local jobs are intended only for jobs that:

- Perform heavy I/O (consequently, very inefficient to run remotely)

- Execute near the end of the build (for example, forcing the build into serial mode would have minimal performance impact)

- Logically produce a final output not consumed by steps later in the build

Typically, final link or packaging steps are the only commands that should be marked *runlocal*.

If a run-local job modifies the registry, these modifications will not be visible to later jobs in the build. Ensure you do not mark a job as *runlocal* if that job modifies the registry in a way that later parts of the build depend on.

To mark a job runlocal, precede its rule in the makefile with `#pragma runlocal`. For example:

```
#pragma runlocal
myexecutable:
        $(CC) -o $@ $(OBJS) ...
```

Also, you can specify `#pragma runlocal` for a pattern rule. In this case, all jobs instantiated from that rule are marked run-local:

```
#pragma runlocal
%.exe: %.obj
        $(CC) -o $@ $*
```

You can cause all jobs after a certain point in the build to run locally by using `#pragma runlocal sticky`. This variant means all jobs that occur later in the serial order than the job specifically marked as "runlocal sticky" are run locally, as if "runlocal" were specified for all of them. This is *not* the same as "all jobs declared after this

job in the makefile!" The order of job declaration in a makefile is not related to the serial order of those jobs, so the difference is significant.

# Serializing All Make Instance Jobs

Normally, Electric Make runs all jobs in a Make instance in parallel on multiple distinct Agents. For most builds, this process ensures the best possible performance.

In some cases, however, all jobs in a Make instance are interdependent and must be run serially, for example, a set of jobs updating a shared file. In particular, Microsoft Visual C/C++ compilers exhibit this behavior when they create a common program database (PDB) file to store symbol and debug information to all object files.

### Example

The following makefile invokes `cl` with the `/Zi` flag to specify the program database file be created with type and symbolic debugging information:

```
all:
        $(MAKE) my.exe
        $(MAKE) other.exe

my.exe: a.obj b.obj c.obj my.obj
                cl /nologo /Zi $^ /Fe'my.exe'

%.obj: %.c

                cl /nologo /Zi /c /Fo$@ $^
```

In this build, the `a.obj, b.obj, c.obj` and `my.obj` jobs are implicitly serialized because they all write to the shared PDB file (by default, `vc70.pdb`). In this case, jobs run in parallel, and running them on separate Agents only introduces unnecessary network overhead. This job type needs to run serially so it can correctly update the PDB file.

The Electric Make special directive, `#pragma allserial` used in the makefile, allows you to disable a parallel build in a Make instance and run the job serially on a single Agent. By inserting the `#pragma allserial` directive at the beginning of a line anywhere in the makefile, the directive specifies that all jobs in that make instance be serialized. This process maximizes network and file cache efficiency.

In the example above, by prefixing the `%.obj` pattern rule with the `#pragma allserial` directive:

```
#pragma allserial
%.obj: %.c
        cl /nologo /Zi /c /Fo$@ $^
```

Electric Make runs compiles and links for the `my.exe` Make instance in serial on the same Agent.

## Splitting PDBs Using hashstr.exe

The `hashstr.exe` utility creates a hash of the filename given a modulus (maximum number of PDBs that will be produced). A given file must always produce the same PDB or history would constantly change. The hash should only include the filename and not its full path. Precompiled headers (PCHs) must be turned off.

### Example

Usage: `hashstr "mystring" [modulus]`

Where mystring is the string from which to generate the hash value, and modulus is the number of hash bins you want to use.

You can add this to a pattern rule for builds that suffer from performance degradation due to PDB serialization, with something similar to the following:

```
%.o: %.c
$(CC) /c $(cflags) $(PCH_USE_FLAGS) $(cvars) $(cplus_flags) $(LOCAL_INCLUDE)
```

```
$(PCB_INCLUDE) $< /Fo$@ /Fd$(shell ${path-to-hashstr}/hashstr.exe "$@"
${hashstr-modulus}).pdb
```

# Managing Temporary Files

During the build process, temporary files are created, then deleted automatically when a build completes. To maintain optimum efficiency, it is important to identify where these files are created in relation to the final output files.

## Configuring the Electric Make Temporary Directory

Electric Make runs commands and collects output in parallel, but the results are written in order, serially. This feature ensures a build behaves exactly as if it were run locally. See Transactional Command Output.

The Electric Make *temporary directory* is used to store output files from commands (for example, object files) that finished running but are waiting their turn to be relocated to a final destination on disk. By default, Electric Make creates a temporary directory in the current working directory. The directory name will have the form:

```
ecloud_tmp_<pid>_<n>
```
where pid is the eMake process identifier.
where n is a counter assigned by eMake to differentiate multiple temporary directories created during a single build if you specify multiple directories.

Electric Make removes the temporary directory on exit (including **Ctrl-C** user interrupt). If Electric Make is terminated unexpectedly (for example, by the operating system), the temporary directory might persist and need to be removed manually.

Important temporary directory requirements include:

- Must be writable by the eMake process

- Must not be on an NFS share

- Must have enough space to contain the build output

  ○ Ideally, the temporary directory should have enough space to hold the complete build output; in practice however, it might need only enough space for output from the largest commands.

  ○ The exact temporary directory space requirement varies greatly with the build size, the number of Agents used, and build system speed.

You can use the --emake-tmpdir command-line option or the EMAKE_TMPDIR environment variable to change the temporary directory default location:

```
% emake --emake-tmpdir=/var/tmp ...
```

Files are relocated from the temporary directory to their final location. Keeping the temporary directory on the same file system as the EMAKE_ROOT helps performance because files can simply be renamed in place, instead of copied.

Electric Make and agent machines reset the values for TEMP, TMP, and TMPDIR. This is necessary to avoid possible conflicts with multiple build agents/jobs running on the same machine.

If you have more than one EMAKE_ROOT that spans multiple file systems, you can specify more than one temporary directory to ensure Electric Make can always rename files in place, instead of copying them.

For example, if your EMAKE_ROOT contains three directories:

```
% setenv EMAKE_ROOT /home/alice:/local/output:/local/libs
```

that reside on two different file systems

```
% df /home/alice /local
Filesystem     1K-blocks   Used   Available Use% Mounted on
filer:/vol/    76376484 47234436 25262352   66%  /home
```

```
vol0/space
/dev/hdb1      76896316 51957460 21032656    2%  /local
```

to ensure better performance, specify two temporary directories—write permission at both locations is required:

```
% emake --emake-tmpdir=/home/alice:/local ...
```

When specifying multiple temporary directories, note the following:

- If a temporary directory was specified for a particular file system, Electric Make automatically uses that directory for any files destined to reside on that file system.

- If a temporary directory was not specified for a particular file system, Electric Make uses the *first* directory specified for files destined to reside on that file system.

## Deleting Temporary Files

During a build, Electric Make creates a temporary directory inside the directory specified by the `EMAKE_TMPDIR` environment variable or in the directory specified by the `--emake-tmpdir` command-line option (or the current working directory if no Electric Make temporary directory is specified). If a specified directory does not exist, eMake creates it. All temporary directories created by eMake are automatically deleted when a build completes.

For example, if you set your temporary directory to `/foo/bar/baz` and only `/foo/bar` exists, eMake creates `/foo/bar/baz` and `/foo/bar/baz/ecloud_tmp_<pid>_<n>`, and then deletes both directories and all their contents when it exits. If `/foo/bar/baz` exists at the start of the build, only `/foo/bar/baz/ecloud_tmp_<pid>_<n>` is created and deleted.

However, if the build is aborted in-progress, eMake will not have the opportunity to remove the temporary directory.

Subsequent Electric Make invocations automatically delete temporary directories if they are more than 24 hours old. You can reclaim disk space more quickly by deleting temporary directories and their contents by hand. However, *do not delete the temporary directory while a build is in-progress*—this action causes the build to fail.

Temporary directory names are in this form:

```
ecloud_tmp_<pid>_<n>
```
where pid is the eMake process identifier.
where n is a counter assigned by eMake to differentiate multiple temporary directories created during a single build if you specify multiple directories.

# Chapter 7: Dependency Management

The following topics discuss ElectricAccelerator eDepend, the Ledger, and how ElectricAccelerator handles history data files.

**Topics:**

- ElectricAccelerator eDepend
- ElectricAccelerator Ledger File
- Managing the History Data File

# ElectricAccelerator eDepend

In its default configuration, Accelerator is designed to be a drop-in replacement for your existing Make tool—GNU Make or Microsoft NMAKE. Accelerator behaves exactly like your existing Make: it will rebuild (or declare up-to-date) the same targets following the same rules Make uses. The file dependency tracking technology in the Electric File System (EFS) and the Electric Make history feature is used to ensure the system reproduces exactly the same results in a parallel, distributed build as it would serially.

Because the system captures and records such detailed information about the relationships between build steps, it is uniquely capable of accomplishing much more than simply ensuring parallel builds are serially correct. In particular, by enabling ElectricAccelerator eDepend, you can wholly replace tools and techniques like **makedepend** or **gcc -M**—commonly used to generate makefiles too difficult to maintain by hand (for example, C-file header dependencies).

ElectricAccelerator eDepend is easier to configure, faster, more accurate, and applicable to a much wider range of dependencies than existing dependency generation solutions. If your current build does not employ dependency generation, you can enable eDepend and benefit from more accurate incremental builds without the overhead of configuring and integrating an external dependency generation tool.

The following sections describe the dependency generation challenge in more detail and how eDepend can improve your build speed and accuracy.

## Dependency Generation

Consider a build tree that looks like this:

```
src/
    Makefile            <--- top-level makefile: recurses into mylib and then into main to build the
                             program
    common/
        header1.h
        header2.h
    mylib/
        Makefile        <--- has rules to build mylib.o and create library mylib.a
        mylib.h
        mylib.c         <--- includes common/header1.h and mylib.h
    main/
        Makefile        <--- has rules to build main.o and then to link main using main.o and
                             mylib.a
        main.c          <--- includes common/header1.h, common/header2.h, and lib/mylib.h
```

## The Problem

Even in this simple example, the need for dependency generation is apparent: if you make a change to a header file, how do you ensure dependent objects are recompiled when you rebuild?

Makefiles could explicitly declare all header dependencies, but that quickly becomes too cumbersome: each change to an individual source file might or might not require an adjustment in the makefile. Worse, conditionally compiled code can create so many permutations that the problem becomes intractable.

Another possibility is to declare all headers as dependencies unilaterally, but then the build system becomes very inefficient: after a successful build, a modification to `header2.h` should trigger a rebuild only of the `main` module, not `mylib.a` as well.

Clearly, to get accurate, efficient builds, the system must have calculated dependencies automatically *before* it builds.

There are several ways to generate dependencies and update makefiles to reflect these dependencies (for example, **makedepend** or **gcc -M**), but they all have the drawbacks mentioned previously.

## eDepend Benefits

ElectricAccelerator eDepend is an Electric Make feature that directly addresses all problems with existing dependency generation solutions. Specifically:

- It is part of Electric Make and requires no external tool to configure, no extra processing time, and it is faster than other solutions.

- It is easily enabled by setting a command-line parameter to Electric Make. No tools or changes to makefiles are required.

- Like ElectricAccelerator itself, it is completely tool and language independent. eDepend automatically records any and all kinds of dependencies, including implicit relationships such as executables on libraries during a link step.

- eDepend dependencies are recorded in Electric Make history files—transparently recorded and used without manifesting as makefile rules.

- eDepend is accurate because it uses file information discovered by the Electric File System at the kernel level as a job executes its commands.

## How Does eDepend Work?

Internally, eDepend is a simple application of sophisticated file usage information returned by the Electric File System.

1. As a job runs, the Electric File System records all filenames it accesses inside `EMAKE_ROOT`.

   This function has two very important implications:

   - eDepend can track dependencies within `EMAKE_ROOT` only.

   - eDepend can track dependencies for a job only after it has run—this is why you must start with a complete build rather than an incremental build.

2. After a job completes, Electric Make saves the following eDepend information to the eMake history file:

   - the working directory for the Make instance

   - the target name

   - any files actually read (not just checked for existence) and/or explicitly listed as prerequisites in the makefile

   **Note:** eDepend information is stored in the history file along with serialization history information. Commands operating on the history file (for example, those specifying file location or that erase it) apply to eDepend information as well.

3. In a subsequent build, whenever Electric Make schedules a target in a directory for which it has eDepend information, it evaluates file dependencies recorded in the earlier run as it checks to see if the target is up-to-date.

   In the example above, the rule to update `mylib.o` might look like this:

   ```
   mylib.o: mylib.c
           $(CC) ...
   ```

`mylib.c` includes `common/header1.h`, which is not explicitly listed as a prerequisite of `mylib.o`, so eDepend records this implicit dependency in the history file.

| Directory | Object | Dependency |
|-----------|--------|------------|
| src/mylib | mylib.o | common/header1.h |

If a change is then made to `common/header1.h`, `src/mylib/mylib.o` it will be rebuilt.

# Enabling eDepend

1. Start from a clean (unbuilt) source tree.

2. If your build system already has a dependency generation mechanism, turn it off if possible. If you cannot turn it off, you will still get eDepend's additional accuracy, but you will not be able to improve the performance or shortcomings of your existing system.

3. Build your whole source tree with eDepend enabled.

   Use the `--emake-autodepend=1` command-line switch:

   ```
   % emake --emake-autodepend=1 ...
   ```

   Alternatively, insert `--emake-autodepend=1` into the `EMAKEFLAGS` environment variable.

   ```
   % setenv EMAKEFLAGS --emake-autodepend=1
   % emake ...
   ```

4. Make a change to a file consumed by the build, but not listed explicitly in the makefiles.

   For example, touch a header file: `% touch someheader.h`

5. Rebuild, again making sure eDepend is enabled.

   ```
   % emake --emake-autodepend=1 ...
   ```

   Notice that without invoking a dependency generator, Electric Make detected the changed header and rebuilt accordingly.

## Important Notes

- The eDepend list is consulted only if all other prerequisites in the makefile indicate the target is up-to-date.

   Explained another way: If a target is out-of-date because it does not exist or because it is older than one of the prerequisites listed in the makefile, eDepend costs nothing and has no effect.

   If the target is newer than all its listed prerequisites, then eDepend is the "11th hour" check to ensure it really is up-to-date, and that there is not a newer implicit dependency. This is the only place eDepend interacts with your build: it forces a target that incorrectly appears to be up-to-date to be rebuilt.

- eDepend information, unlike traditional Makedepend rules, does not in any way imply anything about needing to build or update the implicit prerequisite.

   In the example above, if `header1.h` is renamed or moved, Electric Make just ignores the stale eDepend information. When eMake next updates the `mylib.o` target, it will prune stale dependencies from the eDepend list. This change to the history file occurs regardless of the setting of the `--emake-history-force` parameter to eMake.

   Unlike Make, Electric Make does not complain if it does not have a rule to make `header1.h` because eDepend dependencies are not used to schedule targets.

- eDepend's information scope is bound by a directory name and the target name. This means you can build cleanly from the top of a tree, then run accurate incremental builds from individual subdirectories and eDepend information will be used and updated correctly.

  However, it does imply if you have a build that

  ○ performs multiple passes or variants over the same directories

  ○ with exactly the same target names, but

  ○ runs significantly different commands

  For example, a build that produces objects with the same name for a different architecture or configuration, eDepend information might be over-used unnecessarily. In this case, Electric Make might rebuild more than is necessary, but with no incorrect build results. In this situation, you can achieve fast, correct builds by using separate history files, or ideally, by changing to unique target names across build variants.

## Using #pragma noautodep

Some build steps contain many implicit dependencies that might not make sense to check for up-to-dateness. Examples include symbol files consumed by a link step or archive packager input files (for example, *tar*). In both cases, any makefile explicit prerequisites are sufficient to determine if the target should be updated: eDepend information would just add overhead or cause unnecessary rebuilds.

You can selectively disable eDepend information for certain files from any step by supplying Electric Make with the makefile directive:

```
#pragma noautodep *.pdb
%.o: %.c
        $(CL) ...
```

The directive `#pragma noautodep` is applied to the next rule or pattern rule in the makefile. This directive specifies a class of files for eDepend to ignore. Note the following information about `#pragma autodep`:

1. Wildcards allowed for `#pragma autodep`:

   \* matches 0 or more characters

   ? matches 1 character

   [ ] matches a range of characters

2. The `noautodep` filters are matched against absolute pathnames. To limit a filter to files in the current directory for the job, use `./`:

   ```
   #pragma noautodep ./foo.h
   ```

   To specify "ignore foo.h" in any directory, use:

   ```
   #pragma noautodep */foo.h
   ```

3. If the supplied pattern has no wildcards and does not specify a path, it will never match.

   Electric Make ignores the directive and prints a warning as it parses the makefile:

   ```
   Makefile:2: ignoring autodep filter 'foo',
   does not match absolute path(s).
   ```

# ElectricAccelerator Ledger File

Traditional Make facilities rely exclusively on a comparison of file system timestamps to determine if the target is up-to-date. More specifically, an existing target is considered out-of-date if its inputs have a "last-modified"

timestamp later than the target output.

For typical interactive development, this scheme is adequate: As a developer makes changes to source files, their modification timestamps are updated, which signals Make that dependent targets must be rebuilt. There is, however, a class of workflow styles that cause file timestamps to move arbitrarily into the past or future, and therefore circumvent Make's ability to correctly rebuild targets.

Two common examples are:

- Using a version control system that preserves timestamps on checkout (also known as "sync" or "update").

  The default mode for most source control systems is to set the last-modified timestamp of every file updated in a checkout or sync operation to the current day and time. If you change this behavior to preserve timestamps (or if your tool's default mode is *preserve*), then updating your source files can result in modified contents but with a timestamp in the past (typically, it is the time of the checkin).

- Using file or directory sychronization tools (even simple recursive directory copies) to keep files updated against some other repository.

  Here again, while it is easy to modify source file content, the timestamp for modifications might be any of several possibilities: time of copy, last-modified time of source, last-modified time of destination, and so on.

## The Problem

In all modified source files cases, we would like the Make system to rebuild any dependent objects. However, because timestamps of modified files are not set reliably, Make might or might not force a target update. Here is an example Makefile:

```
foo.o: foo.c
        gcc -c foo.c

foo.o: foo.h
```

And a build is run without an existing `foo.o` object:

```
% make
gcc -c foo.c

% ls -lt
total 4
-rw-r--r-- 1 jdoe None 21 May 29 13:50 foo.o
-rw-r--r-- 1 jdoe None 21 May 29 13:50 foo.c
-rw-r--r-- 1 jdoe None 20 Apr 25 17:34 foo.h
-rw-r--r-- 1 jdoe None 41 Jan 19 09:27 Makefile
```

The `foo.o` target is updated. Next, suppose we ask our source control system to update the working directory, and it responds by giving us a newer copy of `foo.h`, one that is several weeks newer than what we have, and *that* timestamp is preserved:

```
% <sync>
% ls -lt
total 4
-rw-r--r-- 1 jdoe None 21 May 29 13:50 foo.o
-rw-r--r-- 1 jdoe None 21 May 29 13:50 foo.c
-rw-r--r-- 1 jdoe None 29 May 17 11:21 foo.h <-- notice timestamp change
-rw-r--r-- 1 jdoe None 41 Jan 19 09:27 Makefile
```

Traditional Make programs (here, GNU Make) will not notice the change because the timestamp is still in the past, and will incorrectly report that the target is up-to-date.

```
% make
make: `foo.o' is up to date.
```

Some Make facilities (notably, Rational 'clearmake' in conjunction with Rational ClearCase) have the ability to track timestamp information because they are integrated with the source control system.

## The Electric Make Solution

Electric Make solves this problem at the file level, completely independent of the source control system, by keeping a separate database of inputs and outputs called a ***ledger***. To use the Ledger, you specify which file aspects to check for changes when considering a rebuild. To do so, use the `--emake-ledger=<valuelist>` command-line switch (or the `EMAKE_LEDGER` environment variable). `<valuelist>` is a comma-separated list that includes one or more of: `timestamp`, `size`, `command`, `nobackup`, `nonlocal`, and `unknown`. For more information, see Ledger options in Electric Make Command-Line Options, Environment Variables, and Configuration File.

- `timestamp` – Any timestamp changes to either the target or the explicitly declared dependency, regardless of how it relates to the last modified time of the target input file, triggers a target rebuild.

- `size` – Any size change, regardless of the timestamp in the input file, triggers a target rebuild.

- `command` – Records the text of the command used to create the target. If makefile or its variables change, using `command` rebuilds the target. **Important caveat:** If you initialize a variable using the $(*shell*) function, be extremely careful to use the $(*shell*) function with a ':=' assignment to avoid re-evaluating it every time the variable is referenced. ':=' simply expanded variables are expanded immediately upon reading the line.

- `nobackup` – Suppresses the automatic backup of the ledger file before its use.

- `nonlocal` – Instructs eMake to operate on the ledger file in its current location, even if it is on a network volume. By default, if the file specified by `--emake-ledgerfile` (emake.ledger in the current working directory, by default) is not on a local disk, eMake copies that file (if it already exists) to the system temporary directory and opens the copy, then copies it back to the specified location when the build is complete.

  Using `nonlocal` removes a safety and might cause problems if the non-local file system has issues with memory-mapped I/O (IBM Rational ClearCase MVFS is known to have issues with memory-mapped I/O). If you are confident that you will get efficient and reliable memory-mapped I/O performance from the non-local file system, you can remove the safety for improved efficiency because eMake does not spend time at startup and shutdown copying ledger files. Electric Cloud strongly recommends against using `nonlocal` with ClearCase dynamic views. Electric Cloud does not support Ledger-related problems that occur when `nonlocal` is used in conjunction with the MVFS.

- `unknown` – Specifies that the Ledger feature consider a target to be out of date, if the Ledger database contains no entry for the target.

In the example above, the Ledger can detect if a rebuild is necessary as the timestamps change. If the original build was:

```
% emake --emake-ledger=timestamp
gcc -c foo.c
% <sync> <-- notice timestamp change
% emake --emake-ledger=timestamp
gcc -c foo.c
```

Electric Make consulted the Ledger and concluded the target needed to be rebuilt.

### Important Notes for the Ledger Feature

- The Ledger feature works by comparing an earlier input state with the current state: if the Ledger has no information about a particular input (for example, during the first build after it was added to a makefile), it will not contribute in the up-to-dateness check.

- Only one Ledger is used per build.

- The default ledger file is called `emake.ledger`

  It can be adjusted by the `--emake-ledgerfile=<path>` command-line option or `EMAKE_LEDGER_FILE=<path>` environment variable.

- If you specify `--emake-ledgerfile=<path>` but not `--emake-ledger=<valuelist>`, the Ledger still hashes the filenames, so the Ledger is triggered when the filename order changes or a file is added or removed.

- The Ledger automatically backs up the ledger file before using it. This ensures a non-corrupt file is available. If the ledger file is large, copying it could take some time on incremental builds. The ledger option, `nobackup`, suppresses the backup.

- Ledger works for local builds and those using a cluster, as well as local submakes in a runlocal job, see Managing Temporary Files.

- It is not possible, however, to share a Ledger between top-level make instances and local-mode submakes running on the cluster. See EMAKE_BUILD_MODE=local in Electric Make Command-Line Options, Environment Variables, and Configuration File.

- eMake consults Ledger information to trigger a rebuild only when a target would otherwise be considered up-to-date. Information in the Ledger never prevents a target from being rebuilt.

- In a GNU Make emulation, the Ledger feature changes the meaning of the '`$?`' automatic variable to be synonymous with '`$^`' (all prerequisites, regardless of up-to-dateness).

- You cannot change Ledger options for a particular ledger file—you must use the same combination of timestamp, size, and command that was used to create the ledger file.

- If you turn on `--emake-ledger` and `--emake-autodepend` at the same time, the Ledger keeps track of both implicit and explicit dependencies. This feature is comparable to using ClearMake under ClearCase, but is independent of ClearCase information records.

- Order-only prerequisites, in keeping with their semantic meaning, never affect Ledger behavior.

- Because the Ledger automatically rebuilds a target when there is no existing entry in the ledger file, a build that is using the Ledger for the first time might take longer than expected.

# Managing the History Data File

When Accelerator runs a build for the first time, it takes aggressive action to run all jobs as fast as possible in parallel. Jobs that run in the wrong order because of missing makefile dependencies are automatically re-run to ensure correct output. (These failed steps are referred to also as *conflicts*. See Conflicts and Conflict Detection for information about conflicts.)

To avoid the cost of re-running jobs on subsequent builds, Electric Make saves the missing dependency information in a *history data file*. The history data file evolves with each new run and enables Accelerator to continually run builds at peak efficiency.

You can choose the location of the history file and how it is updated.

# Setting the History File Location

By default, Electric Make creates the history file in the directory you use to invoke the build and names it `emake.data`.

The file location can be explicitly specified using the command-line option

```
--emake-historyfile=<pathname>
```

# Selecting History File Options

The history file (by default, `emake.data`) is used for two separate operations during an Electric Make cluster build:

- Input – Electric Make reads the history file as it starts a build to improve build performance.

- Output – Electric Make *writes* to the history file *as it completes a build* to improve performance of subsequent builds.

### Input Rules

If the history file (`emake.data` or whatever was specified with `--emake-historyfile`) exists, it is *always* read and used to improve build performance at the time of build execution.

### Output Rules

Data written to the history file when the build completes depends on the `--emake-history` option setting. Three options are available:

1. Merge – By default, Electric Make merges any new dependencies it discovers into the existing history file. In this way, the history file automatically evolves as your makefiles change, learning dependencies that accelerate the builds you run.

2. Create – If `--emake-history` is set to `create`, the old history file contents are completely overwritten by new dependencies discovered in the run that just completed. Use this setting to start a fresh history file, effectively eliminating any stale information from the file.

3. Read – If `--emake-history` is set to `read`, no data is written to the history file at build completion, and any new dependencies Electric Make discovered are discarded. Use this setting when you are sharing a single, static copy of the history file between developers.

By default, the history file is updated even if the build fails, regardless of the set value for `--emake-history`. You can override this behavior by setting the command-line option to `--emake-history-force=0`.

The history file directly impacts the number of conflicts the build can encounter. Ideally, an ElectricAcceleratorcluster build with good history should have close to "0" conflicts. If you see conflicts starting to rise, ensure you have a current history file for the build you are executing.

### Important Notes for Creating and Using the History File

- Use `--emake-readdir-conflicts=1` to guarantee correct history. Some parallel builds do not succeed without a good history file. In particular, builds that use wildcard or globbing operators to produce build-generated lists of files and operate on those lists might fail. For example, a makefile might use `ld *.o` as shorthand to avoid enumerating all the `*.o` files in a directory. Running the build with `--emake-readdir-conflicts=1` guarantees that the build succeeds and that a history file is created for use by subsequent parallel builds.

   **Note:** Do *not* enable this option all the time. Instead, enable it for one run if you suspect a globbing problem, and then disable it, but use the history file generated by the previous run.

- Relative `EMAKE_ROOT` locations must match. The history file records target file names relative to the `EMAKE_ROOT` specified during that run. For a subsequent build to use the history file correctly, target file

names must have the same path name relative to the eMake root.

For example:

If your eMake root is set to `/home/alice/builds`
and your build references a path name in that root: `/home/alice/builds/lib/foo.o`
then the history file records it as `lib/foo.o`
If a subsequent build sets the eMake root to `/home/bob/builds`
the history file will match correctly.

If, however, the eMake root is set to `/home/bob`
then the file that exists on the disk as `/home/bob/builds/lib/foo.o`
gets the root-relative name of `builds/lib/foo.o`
which does not match the name `lib/foo.o` in the history file generated above.

Because the history file does not match, performance can suffer.

**Note:** Be sure your `EMAKE_ROOT` matches the same location relative to sources as the `EMAKE_ROOT` used to create the history file that you are using.

- For builds with multiple roots, the roots must have the same alphabetical sorting order in each build in order for history to match.
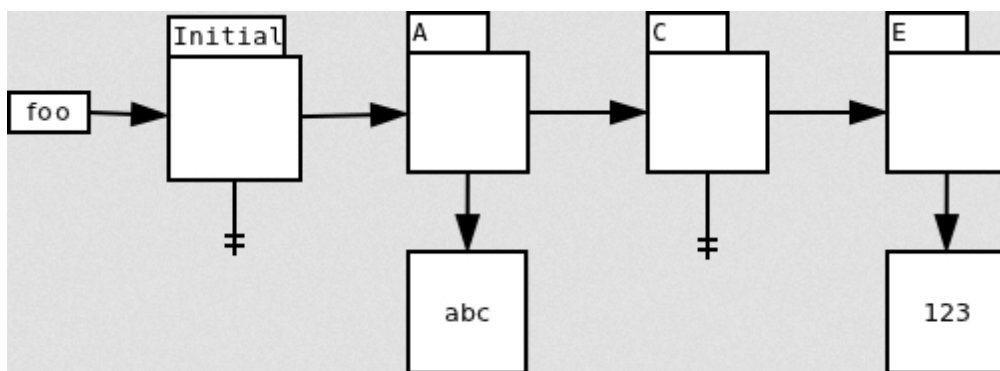
# Conflicts and Conflict Detection

This section discusses conflicts and conflict detection in Electric Make. Content was adapted from the "How ElectricMake guarantees reliable parallel builds" and "Exceptions to conflict detection in ElectricMake" blog posts on http://blog.melski.net/. This content represents the most recent information available when the articles were posted.

## How ElectricMake Guarantees Reliable Parallel Builds

The technology that enables emake to ensure reliable parallel builds is called conflict detection. Although there are many nuances to its implementation, the concept is simple. First, track every modification to every file accessed by the build as a distinct version of the file. Then, for each job run during the build, track the files used and verify that the job accessed the same versions it would have had the build run serially. Any mismatch is considered a conflict. The offending job is discarded along with any file system modifications it made, and the job is rerun to obtain the correct result.

## The Versioned File System

At the heart of the conflict detection system is a data structure known as the versioned file system, in which emake records every version of every file used over the lifetime of the build. A version is added to the data structure every time a file is modified, whether that be a change to the content of the file, a change in the attributes (like ownership or access permissions), or the deletion of the file. In addition to recording file state, a version records the job which created it. For example, here's what the version chain looks like for a file "foo" which initially does not exist, then is created by job A with contents "abc", deleted by job C, and recreated by job E with contents "123":

## Detecting Conflicts

With all the data emake collects — every version of every file, and the relationship between every job — the actual conflict check is simple: for each file accessed by a job, compare the actual version to the serial version. The actual version is the version that was actually used when the job ran; the serial version is the version that would have been used, if the build had been run serially. For example, consider a job B which attempts to access a file foo. At the time that B runs, the version chain for foo looks like this:



Given that state, B will use the initial version of foo — there is no other option. The initial version is therefore the actual version used by job B. Later, job A creates a new version of foo:



Since job A precedes job B in serial order, the version created by job A is the correct serial version for job B. Therefore, job B has a conflict.

If a job is determined to be free of conflicts, the job is committed, meaning any file system updates are at last applied to the real file system. Any job that has a conflict is reverted — all versions created by the job are marked invalid, so subsequent jobs will not use them. The conflict job is then rerun in order to generate the correct result. The rerun job is committed immediately upon completion.

Conflict checks are carried out by a dedicated thread which inspects each job in strict serial order. That guarantees that a job is not checked for conflicts until after every job that precedes it in serial order has been successfully verified free of conflicts — without this guarantee, we can't be sure that we know the correct serial

version for files accessed by the job. Similarly, this ensures that the rerun job, if any, will use the correct serial versions for all files — so the rerun job is sure to be conflict free.

# Exceptions to Conflict Detection in ElectricMake

### Non-Existence Conflicts

One obvious enhancement is to ignore conflicts when the two versions are technically different, but effectively the same. The simplest example is when there are two versions of a file which both indicate non-existence, such as the initial version and the version created by job C in this chain for file foo:



Suppose that job D, which falls between C and E in serial order, runs before any other jobs finish. At runtime, D sees the initial version, but strictly speaking, if it had run in serial order it would have seen the version created by job C. But the two versions are functionally identical — both indicate that the file does not exist. From the perspective of the commands run in job D, there is no detectable difference in behavior regardless of which of these two versions was used. Therefore emake can safely ignore this conflict.

### Directory Creation Conflicts

A common make idiom is `mkdir -p $(dir $@)` — that is, create the directory that will contain the output file, if it doesn't already exist. This idiom is often used like so:

```
$(OUTDIR)/foo.o: foo.cpp
    @mkdir -p $(dir $@)
    @g++ -o $@ $^
```

Suppose that the directory does not exist when the build starts, and several jobs that employ this idiom start at the same time. At runtime they will each see the same file system state — namely, that the output directory does not exist. Each job will therefore create the directory. But in reality, had these jobs run serially, only the first job would have created the directory; the others would have seen the version created by the first job, and done nothing with the directory themselves. According to the simple definition of a conflict, all but the first (serial order) job would be considered in conflict. For builds without a history file expressing the dependency between the later jobs and the first, the performance impact would be disastrous.

### Appending to Files

Another surprisingly common idiom is to append error messages to a log file as the build proceeds:

```
$(OUTDIR)/foo.o: foo.cpp
    @g++ -o $@ $^ 2>> err.log
```

Implicitly, each append operation is dependent on the previous appends to the file — after all, how will you know which offset the new content should be written to if you don't know how big the file was to begin with? In terms of file versions, you can imagine a naive implementation treating each append to the file as creating a complete new version of the file:

err.log → Initial [ ] → A [ abc ] → B [ abc / def ] → C [ abc / def / ghi ]

The problem of course is that you'll get conflicts if you try to run all of these jobs in parallel. Suppose all three jobs, A, B and C start at the same time. They will each see the initial version, an empty file, but if run serially, only A would have seen that version. B would have seen the version created by A; C would have seen the version created by B.

This example is particularly interesting because emake cannot sort this out on its own: as long as the usage reported for err.log is the very generic "this file was modified, here's the new content" message normally used for changes to the content of an existing file, emake has no choice but to declare conflicts and serialize these jobs. Fortunately, emake is not limited to that simple usage record. The EFS can detect that each modification is strictly appending to the file, with no regard to the prior contents, and include that detail in the usage report. Thus informed, emake can record fragments of the file, rather than the entire file content:

err.log → Initial [ ] → A [ abc ] → B [ def ] → C [ ghi ]

Because emake now knows that the jobs are not dependent on the prior content of the file, it need not declare conflicts between the jobs, even if they run in parallel. As emake commits the modifications from each job, it stitches the fragments together into a single file, with each fragment in the correct order relative to the other pieces.

## Directory Read Conflicts

Directory read operations are interesting from the perspective of conflict detection. Consider: what does it mean to read a directory? The directory has no content of its own, not in the way that a file does. Instead, the "content" of a directory is the list of files in that directory. To check for conflicts on a directory read, emake must check whether the list of files that the reader job actually saw matches the list that it would have seen had it run in serial order — in essence, doing a simple conflict check on each of the files in the directory.

That's conceptually easy to do, but the implications of doing so are significant: it means that emake will declare a conflict on the directory read anytime any other job creates or deletes any file in that directory. Compare that to reads on ordinary files: you only get a conflict if the read happens before a write operation on the same file. With directories, you can get a conflict for modifications to other files entirely.

This is particularly dangerous because many tools actually perform directory reads under-the-covers, and often those tools are not actually concerned with the complete directory contents. For example, a job that enumerates

files matching `*.obj` in a directory is only interested in files ending with `.obj`. The creation of a file named `foo.a` in that directory should not affect the job at all.

Another nasty example comes from utilities that implement their own version of the getcwd() system call. If you're going to roll your own version, the algorithm looks something like this:

1. Let `cwd` = ""

2. Let `current` = "."

3. Let `parent` = "./.."

4. stat `current` to get its inode number.

5. read `parent` until an entry matching that inode number is found.

6. add the name from that entry to `cwd`

7. Set `current` = `parent`.

8. Set `parent` = `parent` + "/.."

9. Repeat starting with step 4.

By following this algorithm the program can construct an absolute path for the current working directory. The problem is that the program has a read operation on every directory between the current directory and the root of the file system. If emake strictly adhered to conflict checking on directory reads, a job that used such a tool would be serialized against every job that created or deleted any file in any of those directories.

For this reason, emake deliberately ignores conflicts on directory read operations by default. Most of the time this is safe to do, surprisingly — often tools do not need a completely accurate list of the files in the directory. And in every case I've seen, even if the tool does require a perfectly correct list, the tool follows the directory read with reads of the files it finds. That means that you can ensure correct behavior by running the build one time with a single agent, to ensure the directory contents are correct when the job runs. That run will produce history based on the file reads, so subsequent builds can run with many agents and still produce correct results.

You can also use `–emake-readdir-conflicts=1` to force emake to honor directory read conflicts.

# Chapter 8: Annotation

As Electric Make runs a build, it discovers a large amount of information about the build structure. This information can be written to an "annotation" file for use after the build completes. Annotation information is represented as an XML document so it is easy to parse.

eMake collects many different types of information about the build depending on various command-line switches. The information eMake collects includes:

- makefile structure

- commands and command output

- list of file accesses by each job

- dependencies between jobs

- detailed timing measurements

- eMake invocation arguments and environment

- performance metrics

eMake can also be configured to upload annotation information to the Cluster Manager for centralized reporting purposes.

## Configuring eMake to Generate an Annotation File

By default, Electric Make collects configuration information and performance metrics only, which it sends to the Cluster Manager at the end of the build. This data is used to display reports on the Build Details page.

Optionally, you can configure eMake to collect additional information. This extra annotation information is written to an XML file in the build directory (`emake.xml` by default). The `--emake-annodetail` command-line switch controls the amount of information that eMake should collect. Default annotation detail is determined by the build class for that build.

Supported annotation detail flags:

- **basic** - The basic mode collects information about every command run by the build. Detailed information about each "job" in the build is recorded, including command arguments, output, exit code, timing, and source location. In addition, the build structure is represented as a tree where each recursive make level is represented in the XML output.

  If the JobCache feature is enabled, basic annotation includes annotation about cache hits and misses. For details, see the "Interpreting Job Cache Annotation Information" section in Chapter 6, Performance Optimization. JobCache is an add-on option and requires a separately-purchased license.

- **env** - The env mode adds information about environment variable modifications.

- **file** - The file mode adds information about files read or written by each job.

- **history** - The history mode adds information about missing serializations discovered by eMake. This includes information about which file caused two jobs to become serialized by the eMake history mechanism.

- **lookup** - The lookup mode adds information about files that were looked up by each job. **Note:***This mode can cause the annotation file to become quite large.*

- **md5** - The md5 mode computes MD5 checksums for files read and written by the build, and includes that information in annotation as an md5 attribute on appropriate `<op>` tags. The list of operation types that will include the checksum is read, create, and modify. No checksum is generated or emitted for operations on directories or symlinks, or for append operations. If a read file was appended to, and the read occurs before the appended update is committed, you will see a zero checksum on that read operation (by design because reading files that were appended to is a rare occurrence). The md5 mode implies "file" level annotation. This mode is configurable through the command-line only; it is not available on the web interface.

- **registry** - The registry mode adds information about registry operations.

- **waiting** - The waiting mode adds information about the complete dependency graph for the build.

All of the detail settings automatically enable "basic" annotation.

The `--emake-annoupload` command-line switch controls whether or not eMake sends a copy of the annotation file to the Cluster Manager as the build runs. By default, eMake sends minimal information to the Cluster Manager, even if more detailed annotation is enabled. eMake sends the full annotation file if annotation uploading is configured by the build class or by the eMake command line.

**Note:** You cannot disable `mergestreams` if you enable annotation. Enabling annotation automatically enables `mergestreams`, even if it was explicitly disabled on the command line.

# Annotation File Splitting

Because of limitations in the 32-bit version of the ElectricInsight tool, eMake as well as Electrify automatically partition annotation files into 1.6 GB "chunks." The first chunk is named using the file name that you specify with the `--emake-annofile` option or with "emake.xml," if `--emake-annofile` is not specified. The second chunk uses that name as the base but adds the suffix _1, the third chunk adds the suffix _2, and so on. For example, a four-part annotation file might consist of files named emake.xml, emake.xml_1, emake.xml_2, and emake.xml_3.

No special action is required to load a multipart annotation file into ElectricInsight. If all parts are present in the same directory, ElectricInsight automatically finds and loads the content of each file—simply specify the name of the first chunk when opening the file in ElectricInsight.

For loading large annotation files, Electric Cloud recommends the 64-bit version of ElectricInsight.

# Working with Annotation Files

The simplest way to use an eMake annotation file is to load it into ElectricInsight. This program allows the user to see a graphical representation of the build, search the annotation file for interesting patterns, and perform sophisticated build analysis using its built-in reporting tools.

Also, you can write your own tools to perform simple tasks using annotation output. For example, reporting on failures in the build can be accomplished by looking for "failed" elements inside job elements and then reporting various details about the failed job such as the commands, their output, and the line of the makefile that contains the rule for the command. Refer to the DTD for the annotation file format below.

## Annotation XML DTD

```
<!-- build.dtd -->
<!-- The DTD for Emake's annotated output. -->
<!-- -->
<!-- Copyright (c) 2004-2008 Electric Cloud, Inc. All rights reserved. -->

<!ENTITY % hexnum "CDATA">
<!ENTITY % job "(message*, job)">
<!ENTITY % valueName "name NMTOKEN #REQUIRED">
<!-- Can't use NMTOKEN because Windows has environment variables like
     "=D:". -->
<!ENTITY % envValueName "name CDATA #REQUIRED">

<!ELEMENT build
    (properties?, environment?, (message* | make)+, fs?, metrics? )
>
<!ATTLIST build
    id    CDATA #REQUIRED
    cm    CDATA #IMPLIED
    start CDATA #REQUIRED
>
<!-- Out of band build messages -->

<!ELEMENT message (#PCDATA) >
<!ATTLIST message
    thread    %hexnum;             #REQUIRED
    time      CDATA                #REQUIRED
    code      CDATA                #REQUIRED
    severity  ( warning | error ) #REQUIRED
>
<!-- Properties list -->

<!ELEMENT properties (property*) >
<!ELEMENT property   (#PCDATA) >
<!ATTLIST property
    %valueName;
>
<!-- Environment list -->

<!ELEMENT environment (var*) >
<!ELEMENT var    (#PCDATA) >
<!ATTLIST var
    %envValueName;
    op ( add | modify | delete ) "add"
>
<!-- File system dump -->

<!ELEMENT fs       (roots, symRoots, (content|name)*) >
<!ELEMENT roots    (root+) >
<!ELEMENT root     (#PCDATA) >
<!ATTLIST root
    nameid CDATA #REQUIRED
>
<!ELEMENT symRoots (symRoot*) >
<!ELEMENT symRoot (#PCDATA) >
<!ATTLIST symRoot
    symLinkPath CDATA #REQUIRED
>
```

```
<!ELEMENT content (contentver+)>
<!ATTLIST content
    contentid CDATA #REQUIRED
>
<!ELEMENT contentver EMPTY>
<!ATTLIST contentver
    job CDATA #REQUIRED
>
<!ELEMENT name (namever*) >
<!ATTLIST name
    nameid CDATA #REQUIRED
    dir    CDATA #REQUIRED
    name   CDATA #REQUIRED
>
<!ELEMENT namever EMPTY>
<!ATTLIST namever
    job       CDATA #REQUIRED
    contentid CDATA #REQUIRED
>
<!-- Metrics list -->

<!ELEMENT metrics (metric*) >
<!ELEMENT metric  (#PCDATA) >
<!ATTLIST metric
    %valueName;
>
<!-- Make subtree -->

<!ELEMENT make
    (environment?, ( message | job | make )*)
>
<!ATTLIST make
    level CDATA #REQUIRED
    cmd   CDATA #REQUIRED
    cwd   CDATA #REQUIRED
    mode  ( gmake | nmake | symbian ) #REQUIRED
>
<!-- Job -->

<!ELEMENT job
    (environment?,(output | command | conflict)*,depList?,opList?,
registryOpList?,timing+,failed?,waitingJobs?)
>
<!ATTLIST job
    thread    %hexnum;  #REQUIRED
    id        ID    #REQUIRED
    status    ( normal | rerun | conflict | reverted | skipped ) "normal"
    type      ( continuation | end | exist |
               follow | parse | remake | rule ) #REQUIRED
    name      CDATA #IMPLIED
    file      CDATA #IMPLIED
    line      CDATA #IMPLIED
    neededby  IDREF #IMPLIED
    partof    IDREF #IMPLIED
    node      CDATA #IMPLIED
>
<!-- Command and related output, output blocks can contain nested -->
<!-- make subtrees in local mode. -->
```

```
<!ELEMENT command
    (argv,inline*,(output | make)*)
>
<!ATTLIST command
    line  CDATA #IMPLIED
>
<!ELEMENT argv (#PCDATA)>
<!ELEMENT inline (#PCDATA)>
<!ATTLIST inline
    file CDATA #REQUIRED
>
'<!ELEMENT output (#PCDATA)>
<!ATTLIST output
    src   ( prog | make ) "make"
>
<!-- Conflict description -->

<!ELEMENT conflict EMPTY>
<!ATTLIST conflict
    type        ( file | cascade | name | key | value ) "cascade"
    writejob   IDREF #IMPLIED
    file        CDATA #IMPLIED
    rerunby    IDREF #IMPLIED
    hkey        CDATA #IMPLIED
    path        CDATA #IMPLIED
    value       CDATA #IMPLIED
>
<!-- Job failure code -->

<!ELEMENT failed EMPTY>
<!ATTLIST failed
    code CDATA #REQUIRED
>
<!-- List of jobs waiting for this job, local mode only -->

<!ELEMENT waitingJobs EMPTY>
<!ATTLIST waitingJobs
    idList IDREFS #IMPLIED
>
<!-- Start and stop times of this job -->

<!ELEMENT timing EMPTY>
<!ATTLIST timing
    invoked    CDATA #REQUIRED
    completed CDATA #REQUIRED
    node       CDATA #IMPLIED
>
<!-- Dependency list, only used when annoDetail includes 'history' -->

<!ELEMENT depList (dep*)>
<!ELEMENT dep     EMPTY>
<!ATTLIST dep
    writejob IDREF #REQUIRED
    file     CDATA #REQUIRED
>
<!-- Operation list, only present when annoDetail includes -->
<!-- 'file' or 'lookup' -->
```

```
<!ELEMENT opList (op*)>
<!ELEMENT op     EMPTY>
<!ATTLIST op
    type     ( lookup | read | create | modify | unlink | rename |
               link | modifyAttrs | append | blindcreate ) #REQUIRED
    file     CDATA #REQUIRED
    other    CDATA #IMPLIED
    found    ( 1 | 0 ) "1"
    isdir    ( 1 | 0 ) "0"
    filetype ( file | symlink | dir ) "file"
    atts     CDATA #IMPLIED
>
<!-- Registry operation list, only present when annoDetail includes -->
<!-- 'registry' -->

<!ELEMENT registryOpList (regop*)>
<!ELEMENT regop     (#PCDATA)>
<!ATTLIST regop
    type     ( createkey | deletekey | setvalue | deletevalue |
               lookupkey | readkey ) #REQUIRED
    hkey     CDATA #REQUIRED
    path     CDATA #REQUIRED
    name     CDATA #IMPLIED
    datatype ( none | sz | expandsz | binary | dword | dwordbe |
               link | multisz | resourcelist | resourcedesc |
               resourcereqs | qword ) "none"
>
```

# Metrics in Annotation Files

The following values are available when you select the Metrics option from the dropdown menu on the Build Details page in the web interface.  Some performance metrics are available only with --emake-debug=g.

| Metric | Description |
|---|---|
| chainLatestReads | Number of times the latest version of an FSChain was requested. |
| chainSerialReads | Number of times an FSChain was requested to match a particular spot in the serial order. |
| chainWrites | Number of new versions created, which occurs any time a new name is created, or the content of a file changes. |
| clusterAvailability | Cluster availability percentage. |
| compressBytesIn | The number of bytes passed in for compression. |
| compressBytesOut | The number of bytes returned from compression. |
| compressTime | Time spent compressing data. |
| conflicts | Number of jobs that ran into conflicts and had to be re-run. |
| decompressBytesIn | The number of bytes passed in for decompression. |

| Metric | Description |
|--------|-------------|
| decompressBytesOut | The number of bytes returned from decompression. |
| decompressTime | Time spent decompressing data. |
| diskReadBytes | The number of bytes read from the local disk. |
| diskReadTime | Time spent reading data from the local disk. |
| diskReadWaitTime | Time spent waiting for data to be read from the local disk. |
| diskWriteBytes | The number of bytes written to the local disk. |
| diskWriteTime | Time spent writing data to the local disk. |
| diskWriteWaitTime | Time spent waiting for data to be written to the local disk. |
| duration | Duration of build (in seconds). |
| elapsed | The total elapsed time for the build. |
| emptyNameVersions | Total number of Name versions with no associated Content. Available with `--emake-debug=g`. |
| fitness | Indicates the "fitness" of the history, which is how closely the history used for a given build "matched" that build. The range of values is 0 (history did not match at all) to 1 (history already had information about all implicit dependencies). Available with `--emake-debug=g`. |
| freezeTime | Time the job queue was frozen, which means only high priority items are taken of the queue. |
| hiddenWarningCount | The number of warning messages hidden by the eMake client and all remote parse jobs, with one count for every message number for which at least one message was hidden. The count does not include messages hidden by eMake stubs or rlocal-mode eMakes. |
| jobcache.hit | Number of jobs for which eMake obtained job cache hits. A hit means that the job was replayed from the cache. JobCache is an add-on option and requires a separately-purchased license. |
| jobcache.miss | Number of jobs for which a file system input changed since the last time the job was saved into the cache. JobCache is an add-on option and requires a separately-purchased license. |

| Metric | Description |
|--------|-------------|
| jobcache.na | Number of jobs for which job caching was not applicable. For example, jobs for which the output files are not object files. JobCache is an add-on option and requires a separately-purchased license. |
| jobcache.newslot | Number of jobs for which eMake created a new cache slot. This occurs when the previously-cached jobs differed from these jobs in their command-line arguments, relevant environment variables, or working directories. Together, these parameters of a job choose the slot for that job, which is where eMake stores the results of that job. JobCache is an add-on option and requires a separately-purchased license. |
| jobcache.rootschanged | Number of jobs for which there is no natural mapping from the old eMake roots to the new eMake roots. JobCache is an add-on option and requires a separately-purchased license. |
| jobcache.uncacheable | Number of jobs for which eMake encountered an issue updating the job's relevant cache slot. Look for ERROR and WARNING messages in the console output from eMake. Jobs that failed, jobs for which caching was disabled by the #pragma jobcache none pragma, and jobs that did not access files inside the eMake root are included in this metric. JobCache is an add-on option and requires a separately-purchased license. |
| jobcache.unneeded | Number of jobs for which JobCache was enabled but not needed (because the target was already up to date according to ordinary GNU Make rules). The cache was not consulted, even though caching was requested for that target. JobCache is an add-on option and requires a separately-purchased license. |
| jobcache.workloadsaved | Difference between agent usage (in seconds) with and without the JobCache feature. This represents an estimate of the agent usage saved by the JobCache feature. JobCache is an add-on option and requires a separately-purchased license. |
| localAgentReadBytes | Bytes copied from local agents. |
| localAgentReadTime | Time spent copying data from local agents. |
| localAgentWriteBytes | Bytes copied to local agents. |
| localAgentWriteTime | Time spent copying data to local agents. |

| Metric | Description |
| --- | --- |
| maxArenaCount | Peak active arenas during the build. |
| maxMakeCount | Peak active make instances during the build. |
| noAgentsWaitTime | Amount of time spent waiting with no agents allocated to the build. |
| nodeReadBytes | The number of bytes read from agents. |
| nodeReadTime | Time spent reading data from agents. |
| nodeReadWaitTime | Time spent waiting for data to be read from agents. |
| nodeWriteBytes | The number of bytes written to agents. |
| nodeWriteTime | Time spent writing data to agents. |
| nodeWriteWaitTime | Time spent waiting for data to be written to agents. |
| runjobs | Number of jobs that did work. |
| symlinkReads | The number of times the application read symlinks from the local disk. |
| termDiskCopiedBytes | The number of bytes committed by copying. This should be a small number if possible. If you are using ClearCase and this number is nonzero, look into the `--emake-clearcase=vobs` option. |
| termDiskCopiedFiles | The number of files committed by copying. This number should be small if possible. |
| termDiskMovedBytes | The number of bytes committed by moving. You want this to be big number if possible. |
| termDiskMovedFiles | The number of files committed by moving. You want this to be a big number if possible. |
| termDiskRemovedBytes | The number of bytes committed by remove (original file). |
| termDiskRemovedFiles | The number of files committed by remove (original file). |
| termDiskRemoved TmpBytes | The number of bytes committed by remove (temporary file). |
| termDiskRemoved TmpFiles | The number of files committed by remove (temporary file). |
| terminated | Number of jobs that ran to completion. These jobs might not necessarily have done anything. |

| Metric | Description |
|--------|-------------|
| totalAgentCount | Total number of agents in cluster. |
| totalChains | Total number of FSChains in the build—anything in the file system tracked by eMake for versioning is an FSChain, for example, file names, file contents. Because both names and contents are tracked, this should be at least twice the number of files accessed in the build. |
| totalNameVersions | Total number of Name versions. Available with `--emake-debug=g`. |
| totalVersions | Total number of FSChain versions. Available with `--emake-debug=g`. |
| usageBytes | The number of bytes received for usage data. |
| workload | Sum of agent usage over all agents (in seconds). This value is also reported to the Cluster Manager. |
| writeThrottleWaitTime | Amount of time spent waiting for the write throttle (which is in place to avoid slowing the build by seeking the disk head back and forth too often). |

### Timers

Note that most timers will not be available unless you run with `--emake-debug=g` (for profiling). These timers correspond to the amount of time that eMake spent in certain areas of the code or in a certain state.

| Metric | Description |
|--------|-------------|
| timer:agentManager.should RequestAgents | Time spent checking to see if eMake should be requesting agents, which involves talking to the Cluster Manager. |
| timer:agentManager.startup | The time spent for the Agent Manager to start up. |
| timer:agentManager.stop | Time spent for the Agent Manager to shut down. |
| timer:agentManager.work | Time spent with the Agent Manager actively doing work outside shouldRequestAgents. |
| timer:annoUpload.startup | Time spent starting up the thread to upload annotation. |
| timer:annoUpload.work | Time spent uploading annotation. |
| timer:bench | Benchmark showing the cost of 100 invocations of the timer code (start/stop). |
| timer:directory.populate | Time spent making sure that eMake's model of the directory contents is fully populated. |

| Metric | Description |
|---|---|
| timer:history.parsePrune | Time spent in parse jobs signaling stale history entries to prune stale events. |
| timer:history.pruneFollowers | Time spent signaling stale submakes (for jobs that have followers) to be pruned. |
| timer:history.pruneNo Follower | Time spent signaling stale submakes (for jobs that have no followers) to be pruned in jobs. |
| timer:idle.agentManager | Time spent in the Agent Manager sleeping between shouldRequestAgents checks. |
| timer:idle.agent ManagerRequest | Time spent in the Agent Manager waiting for the Cluster Manager to respond to a request for agents. |
| timer:idle.agentRun | Time spent by worker threads waiting for requests from the agent. |
| timer:idle.annoUpload | Time spent by the annotation upload thread waiting for data. |
| timer:idle.noJobs | Time spent by worker threads waiting for a new runnable job to enter the job queue. |
| timer:idle.untilCompleted | Time spent by the Terminator thread waiting for jobs to be completed. |
| timer:idle.waitForAgent | Time spent by worker threads waiting for an agent to become available. |
| timer:jobqueue | Time spent within the lock guarding the job queue. |
| timer:Ledger.close | Time spent closing the Ledger and flushing data to disk. |
| timer:Ledger.commit | Time spent committing Ledger data to the database. |
| timer:Ledger.isUpToDate | Time spent querying the Ledger to find if a file is up to date. |
| timer:Ledger.staleAttributes | Time spent by the Ledger code statting files to ensure recorded attributes match the actual attributes on disk. |
| timer:Ledger.update | Time spent updating the Ledger. |
| timer:main.commit | Time spent by the Terminator thread committing jobs, less time spent flushing deferred writes to disk. |
| timer:main.history | Time spent reading and writing history files. |
| timer:main. lockedWriteToDisk | Time spent by the Terminator thread flushing deferred writes to disk. |

| Metric | Description |
|---|---|
| timer:main.terminate | Time spent by the Terminator thread terminating jobs. |
| timer:main.writeToDisk | Time spent by the Terminator thread writing operations to disk, less the time covered by any of the other main.writeToDisk timers. |
| timer:main.writeToDisk. append | Time spent appending to existing files on disk. |
| timer:main.writeToDisk. createdir | Time spent creating directories on disk. |
| timer:main.writeToDisk. createdir.attrs | Time spent changing directory attributes on disk for newly created directories. |
| timer:main.writeToDisk. createdir.chown | Time spent changing file ownership on disk for newly created files. |
| timer:main.writeToDisk. createdir.chown | Time spent changing directory ownership on disk for newly created directories. |
| timer:main.writeToDisk. createdir.times | Time spent changing directory times on disk for newly created directories. |
| timer:main.writeToDisk. createfile | Time spent writing file data to disk for newly created files. |
| timer:main.writeToDisk. createfile.attrs | Time spent changing file attributes on disk for newly created files. |
| timer:main.writeToDisk. createfile.times | Time spent changing file times on disk for newly created files. |
| timer:main.writeToDisk.link | Time spent creating links on disk. |
| timer:main.writeToDisk. modify | Time spent modifying existing files on disk. |
| timer:main.writeToDisk. modifyAttrs | Other time spent in writing attribute changes to disk (mostly notifying the file system that attributes have gone "stale"). |
| timer:main.writeToDisk. modifyAttrs.attrs | Time spent modifying attributes of existing files on disk. |
| timer:main.writeToDisk. modifyAttrs.chown | Time spent modifying ownership of existing files on disk. |
| timer:main.writeToDisk. modifyAttrs.times | Time spent modifying times of existing files on disk. |

| Metric | Description |
| --- | --- |
| timer:main.writeToDisk. unlink | Time spent unlinking existing files on disk. |
| timer:main.writeToDisk. unlink.data | Time spent recording the fact that a file was removed. |
| timer:main.writeToDisk. unlink.tree | Time spent removing entire trees on disk. |
| timer:mergeArchiveRefs | Time spent modifying word lists for multi-word archive references like `lib(member1 member2 ...)`. |
| timer:mutex.DirCache | Time spent waiting for a directory cache. |
| timer:mutex.filedata.nodelist | Time spent within the lock guarding the list of agents allocated to the build. |
| timer:mutex.jobcreate | Time spent within the lock used to synchronize the terminator and worker threads when creating new jobs. |
| timer:mutex.joblist | Time spent within the lock used to protect the job list. |
| timer:mutex.jobrunstate | Time spent within the lock used to coordinate starting and canceling jobs. |
| timer:mutex.nodeinit | Time spent within the lock used to protect the list of hosts while initializing agents. |
| timer:mutex.target | Time spent within the lock protecting failure tracking on a target. |
| timer:node.putAllVersions. getShortName | Time spent getting file short names on Windows when doing an E2A_PUT_ALL_VERSIONS. |
| timer:node.setup | Time spent connecting to hosts and initializing them. |
| timer:node.svc.getData | Time spent handling A2E_GET_FILE_DATA and A2E_RESEND_FILE_DATA, not including the time spent sending data in response. |
| timer:node.svc.getData. acquireLock | Time spent waiting for the ChainLock when handling A2E_GET_FILE_DATA and A2E_RESEND_FILE_DATA. |
| timer:node.svc.getData.copy | Time spent copying file data for E2A_LOAD_LOCAL_FILE. This is the local agent version of timer:node.svc.getData.send. |
| timer:node.svc.getData. insideLock | Time spent holding the ChainLock when handling A2E_GET_FILE_DATA and A2E_RESEND_FILE_DATA. |

| Metric | Description |
|--------|-------------|
| timer:node.svc.getData.send | Time spent sending file data during E2A_PUT_FILE_DATA and E2A_PUT_BIG_FILE_DATA. |
| timer:node.svc.getDir | Time spent sending directory entry data to the agent. |
| timer:node.svc.getVersions | Time spent handling A2E_GET_ALL_VERSIONS requests. |
| timer:node.svc.getVersions.acquireLock | Time spent waiting for the ChainLock when handling A2E_GET_ALL_VERSIONS. |
| timer:node.svc.getVersions.insideLock | Time spent holding the ChainLock when handling A2E_GET_ALL_VERSIONS. |
| timer:node.svc.runCommand | Time spent handling A2E_RUN_COMMAND requests. |
| timer:usage.io | Time spent reading and responding to usage data. |
| timer:usage.io.makedata | Time spent creating file data from usage. |
| timer:usage.io.makedata.local | Time spent copying file data from usage reported by local agents; a subset of timer:usage.io.makedata. |
| timer:usage.latency | Time spent dispatching incoming usage data and saving output files. |
| timer:usage.record | Time spent recording usage data, including resolving new name IDs, removing duplicate lookup records, and so on. |
| timer:worker.continuationjob | Time spent by worker threads running continuationJobs. |
| timer:worker.endjob | Time spent by worker threads running endJobs. |
| timer:worker.existencejob | Time spent by worker threads running existenceJobs. |
| timer:worker.followjob | Time spent by worker threads running followJobs. |
| timer:worker.invoke | Time spent invoking remote jobs. |
| timer:worker.invokelocal | Time spent invoking local jobs. |
| timer:worker.other | Otherwise unaccounted time spent by worker threads. |
| timer:worker.parsejob | Time spent by worker threads running parseJobs. |
| timer:worker.remakejob | Time spent by worker threads running remakeJobs. |
| timer:worker.rulejob | Time spent by worker threads running ruleJobs, less the time spent actually running commands and figuring out if it needs to run. |

| Metric | Description |
| --- | --- |
| timer:worker.rulejob. needtorun | Time spent by ruleJobs figuring out if they need to run. |
| timer:worker. runcommands | Time spent by worker threads running commands in commandJobs. |
| timer:worker. shouldRequestAgents | Time spent by worker threads checking to see if they should request agents. |
| timer:worker.startup | Time spent by worker threads initializing. |
| timer:worker.stop | Time spent shutting down worker threads. |
| timerThreadCount | Number of threads in eMake. |

# Chapter 9: Third-party Integrations

The following topics provide information about how ElectricAccelerator integrates with the following environments:

- ClearCase
- Cygwin
- Eclipse

# Using ClearCase with ElectricAccelerator

If your build environment relies on the ClearCase source control system, there are some special considerations for running eMake. ClearCase views can be either "snapshot" or "dynamic."

- ClearCase *snapshot views* behave like a normal file system, so no special support is required.

- ClearCase *dynamic views* have non-standard file system behavior that requires explicit handling by eMake.

**Note:** ElectricAccelerator does not currently support ClearCase integration on Solaris x86. If you have need of this support, please contact your Electric Cloud sales representative. You can make use of ecclearcase_fake.so to provide information about your ClearCase setup through an `ini` file. Refer to eMake's "fake" Interface for ClearCase.

## Configuring ElectricAccelerator for ClearCase

### ecclearcase Executable

Set the EMAKE_CLEARCASE_SERVER environment variable to the path of the ecclearcase executable. This ensures eMake can locate the correct version of ecclearcase if, for example, you are using 64-bit eMake.

### LD_LIBRARY_PATH

Using ElectricAccelerator in a ClearCase environment requires your LD_LIBRARY_PATH (on UNIX) or PATH (on Windows) to contain a directory that includes libraries required to run "cleartool." Library filenames for Windows or UNIX begin with "libatria" (Windows - *libatria.dll*, UNIX - *libatria.so).*

If you plan to use ClearCase with eMake, you must add the ClearCase shared libraries to the LD_LIBRARY_ PATH on your system.

For sh:

```
LD_LIBRARY_PATH=/usr/atria/linux_x86/shlib:$LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

For csh:

```
setenv LD_LIBRARY_PATH /usr/atria/linux_x86/shlib:${LD_LIBRARY_PATH}
```

(`/usr/atria/linux_x86/shlib` is an example and might differ on your system depending on what OS you use and where ClearCase is installed.)

To ensure ElectricAccelerator knows where ClearCase is installed, edit `/etc/ld.so.conf` to include the ClearCase installation location. As a second option, you can include the ClearCase installation location in LD_ LIBRARY_PATH.

### ClearCase Views on agents

When ElectricAccelerator replicates a ClearCase view on an agent, it appears as a generic file system—ClearCase commands that run as part of a build will not work on the host, even if ClearCase is installed on that machine. The Electric File System masks ClearCase's VOB mounts. If your build runs ClearCase commands, these commands must be *runlocal* steps. For additional information, see Managing Temporary Files.

When replicates a ClearCase view on an agent, it appears as a generic file system—ClearCase commands that run as part of a build will not work on the host, even if ClearCase is installed on that machine. The Electric File System masks ClearCase's VOB mounts. If your build runs ClearCase commands, these commands must be *runlocal* steps. For additional information, see Managing Temporary Files.

**Note:** Because of the potential adverse interaction between two different file systems (ClearCase and ElectricAccelerator), Electric Cloud recommends that you do *not* install ClearCase on ElectricAccelerator Agent machines. If you must run ClearCase on an ElectricAccelerator Agent machine, ensure that whichever one you need to start and stop frequently is configured to start "second" at system startup time.

### --emake-clearcase

The Electric Make command-line option `--emake-clearcase` controls which ClearCase features are supported for a build. By default, ClearCase integration for `rofs`, `symlink`, and `vobs` is disabled. To turn on support for these specific ClearCase features, if your build relies on these options, use `--emake-clearcase=LIST`, where `LIST` is a comma-separated list of one or more of the following values:

- `rofs` : detect read-only file systems

  eMake queries ClearCase for each file it accesses to determine whether the file should be considered 'read-only'.

- `symlink` : detect symbolic links (Windows only)

  eMake queries ClearCase for each file it accesses to determine whether the file is a ClearCase symbolic link.

- `vobs` : configure separate temporary directories for each vob

  Normally, eMake uses the `--emake-tmpdir` setting to determine where to place temporary directories for each device. With the 'vobs' option enabled, eMake automatically configures one directory per VOB. On Windows, eMake also communicates with ClearCase to determine which VOB a file belongs to so it can select the correct temporary directory.

**Note:** If `--emake-clearcase` is not specified on the command line and the environment variable `EMAKE_CLEARCASE` is present, Electric Make takes the options from the environment.

### eMake's "fake" Interface for ClearCase

In addition to a direct interface to ClearCase, eMake also provides a "fake" interface that allows the end user to pass information manually to eMake about the ClearCase environment. Normally, you invoke ClearCase functionality by specifying `--emake-clearcase=LIST` to eMake, at which point eMake attempts to load ecclearcase6.so and ecclearcase7.so (.dll on Windows). Whichever library successfully initializes in the ClearCase environment is used to talk to ClearCase through a provided API that is no longer maintained or supported. You can specify the precise library to load by setting the environment variable `EMAKE_CLEARCASE_LIBRARY` to the path to the desired library.

Under some conditions, the ClearCase API does not function properly. For this circumstance, eMake provides ecclearcase_fake.so (.dll on Windows). If you point EMAKE_CLEARCASE_LIBRARY to the fake interface, eMake loads that instead. The fake interface then loads the file specified in the environment by ECCLEARCASE_FAKE_INI, defaulting to ecclearcase_fake.ini. The `ini` file has two sections: `[vobs]` and `[attrs]`.

The `[vobs]` section maps a VOB path to a comma-separated set of attributes. Currently, `public` should be present for a public VOB and `ro` for read-only.

The `[attrs]` section maps a filename to *symlink\*type*, where *symlink* might be empty if the file is not a symbolic link and *type* can be `null`, `version`, `directory_version`, `symbolic_link`, `view_private`, `view_derived`, `derived_object`, `checked_out_file`, `checked_out_dir`. If *symlink* is not empty, `symbolic_link` is assumed. If *type* is `version` or `directory_version` and `--emake-clearcase=rofs` is active, the EFS returns EROFS (or STATUS_ACCESS_DENIED on Windows) when an attempt is made to write the file.

If CLEARCASE_ROOT is set in the environment (as by `cleartool setview`), all `[attrs]` entries are tracked under their exact path as well as one with the CLEARCASE_ROOT prepended. If CLEARCASE_ROOT is set to `/view/testview`, setting `/vobs/test/symlink2` in `[attrs]` is the same as setting both `/vobs/test/symlink2` and `/view/testview/vobs/test/symlink2`.

Sample `ini` files for UNIX:

```
[vobs]
/vobs/test=public
/vobs/readonly=public,ro
```

```
[attrs]
/vobs/test/symlink2=symlink
/vobs/test/symlink/alpha=*directory_version
/vobs/test/symlink/beta=alpha
```

and Windows:

```
[vobs]
\test=public
\readonly=public,ro

[attrs]
S:/test/symlink2=symlink
S:/test/symlink/alpha=*directory_version
S:/test/symlink/beta=alpha
```

## Where ClearCase Dynamic Views Affect eMake Behavior

### Read-only mounts

ClearCase can mount files in a read-only mode, which means they appear to be writable, but any attempts to modify these files fail with a "read-only file system" (UNIX) or "access denied" (Windows) error message. Because eMake cannot tell whether a file is modifiable using normal file system interfaces, it does not know to disallow modifications performed by commands running on the agents. This activity leads to failures when eMake attempts to commit changes (incorrectly) allowed on the agent.
A simple test case:

```
unix% cleartool ls
Makefile
clock@@/main/2  Rule: /main/LATEST

unix% cat Makefile
all:
mv clock clock.old
```

The file "clock" is checked in to ClearCase. Makefile attempts to rename it. If you just run "make", it fails immediately, but can be instructed to ignore the error:

```
unix% make -i
mv clock clock.old
mv: cannot move `clock' to `clock.old': Read-only file system
make: [all] Error 1 (ignored)
```

Note that this file system is *not* mounted *read-only,* so a Makefile can be created. Because "clock" is checked-in, it cannot be renamed without checking it out first, and ROFS is the error ClearCase gives.

Now try this with eMake:

```
unix% emake --emake-root=/vobs -i
Starting build: 114626
mv clock clock.old
ERROR EC1124: Unable to rename file
/vobs/test/drivel/clock to /vobs/test/drivel/clock.old: Read-only
file system (error code 0x1e): Read-only file system
Interrupted build: 114626   Duration: 0:00 (m:s) Cluster
availability: 100%
```

Without activating ClearCase support, eMake does not know "clock" cannot be moved, so the operation succeeds on the agent, then fails when eMake attempts to commit it to disk. Specifying the "-i" flag to ignore errors will not work here.

```
unix% /home/user/Projects/4.2/i686_Linux/ecloud/emake
/emake --emake-clearcase=rofs --emake-root=/vobs -i
```

```
Starting build: 114630
mv clock clock.old
mv: cannot move `clock' to `clock.old': Read-only file system
make: [all] Error 1 (ignored)
Finished build: 114630   Duration: 0:00 (m:s)
Cluster availability: 100%
```

When eMake knows to replicate ClearCase's behavior, the error occurs on the host and can be handled normally.

### Multiple VOBs

eMake writes uncommitted files into temporary directories, and moves them into their correct location after resolving any conflicts. eMake automatically places a temporary directory in the current working directory where it is invoked, and also creates a temporary directory in each location specified by the `--emake-tmpdir` option or the `EMAKE_TMPDIR` environment variable. When possible, eMake writes uncommitted files to the same physical device where the file will be saved when it is committed, which makes the commit operation a lightweight "rename" instead of a heavyweight "copy" operation.

Under ClearCase, each VOB functions as a separate physical disk, so to achieve optimal performance, a temporary directory must be specified for each VOB where the build writes files. `--emake-clearcase=vobs` sets up this directory for you automatically.

- On UNIX, each VOB has a distinct physical device ID, and this option is nothing more than a "shorthand" for specifying `EMAKE_TMPDIR=/vobs/foo:/vobs/bar:...`. in the environment.

- On Windows, you must interface with ClearCase directly to make this distinction, so using `--emake-clearcase=vobs` is important to get the most speed for a build that writes to multiple VOBs.

### Symbolic links

On Windows, ClearCase conceals the nature of its symbolic links from other programs, so what is actually a single file appears to be two different files to other programs. This situation creates an issue for eMake's versioning mechanism as it tracks two separate chains of revisions for one underlying entity; a job's view of the file can get out of sync and lead to build failures. `--emake-clearcase=symlink` interfaces directly with ClearCase to determine whether a particular ClearCase file is a symbolic link and represents it on the agent as a reparse point, which is the native Windows equivalent of a symbolic link. All file operations are redirected to the target of the symbolic link, avoiding synchronization problems.

This issue does not occur on UNIX platforms because ClearCase uses native file system support for symbolic links.

#### A simple test case:

Beginning with a directory, "alpha", and a symlink to that directory, "beta":

```
windows% cleartool ls -d alpha beta
alpha@@/main/1   Rule: \main\LATEST
beta --> alpha
```

And a makefile:

```
all:
        @echo "Furshlugginer" > alpha/foo
        @echo "Potrzebie" > beta/foo
        @cat alpha/foo

windows% emake --emake-root=. -f symlink.mk
Starting build: 50070 Furshlugginer
Finished build: 50070   Duration: 0:02 (m:s)   Cluster availability: 100%

windows% emake --emake-root=. -f symlink.mk --emake-clearcase=symlink
Starting build: 50071 Potrzebie
Finished build: 50071   Duration: 0:01 (m:s)   Cluster availability: 100%
```

*Explanation:* ClearCase has no way to tell the Windows file system that the symlink is a symlink, so `alpha/foo` and `beta/foo` appear to be distinct files. (On UNIX, this is not an issue because symlinks are a standard operating system feature, which means ClearCase can show them as such.) If a build does not contain any ClearCase symbolic links, there is no reason to turn on the integration; if it does, though, you can run into the problem where eMake assumes there are two different files when there is actually just one underlying file, in which case you need to turn on the "symlink" part of our ClearCase integration.

## Performance Considerations

Running builds from ClearCase dynamic views can impose a considerable performance cost depending on the ClearCase configuration and your build. The best performance is achieved by using ClearCase snapshot views. If using snapshots is not possible, there are a few things to consider when setting up an eMake build.

Enabling the "symlink" or "rofs" options incurs a performance cost because of the need to communicate with the ClearCase server when accessing a file for the first time. Many builds do not need these features, even if they are running inside a ClearCase dynamic view, so consider leaving them turned off unless you encounter unexpected build failures.

Enabling the "vobs" option should have minimal performance cost, and might significantly speed up your build if build output is written back to your dynamic view.

Because of improved caching, eMake might provide a significant performance boost beyond that provided by running build steps in parallel. eMake caches much of the file system state, reducing the total number of requests to the ClearCase server during the build. Depending on how heavily loaded your ClearCase server is, this can significantly improve build performance. If you notice build speedups higher than you would expect given the number of agents in your cluster, improved caching might be the reason.

Using the "fake" interface for ClearCase (see eMake's "fake" Interface for ClearCase), which lets you specify the details of VOBs and files in a static file, is much faster than communicating with ClearCase. This might suffice for many users.

# Using Cygwin with ElectricAccelerator (Windows Only)

Cygwin is a Linux-like environment for Windows that consists of two parts:

- A DLL (`cygwin1.dll`) that acts as a Linux API emulation layer, providing substantial Linux API functionality.

- A collection of tools that provide a Linux look and feel.

If your builds used gmake in a Cygwin environment, you might need to use Electric Make's `--emake-emulation=cygwin` option.

For more information about other Cygwin-specific Electric Make command-line options and corresponding environment variables, see Windows-Specific Commands. Specifically, the following command-line options:

```
--emake-cygwin=<Y|N|A>
```

and

```
--emake-ignore-cygwin-mounts=<mounts>
```

# Using Eclipse with ElectricAccelerator

To configure Eclipse to run eMake, follow this procedure:

1. Open your C++ project.

2. Go to the project's Properties > Builders and click **New**.

3. Select **Program** and click **OK**.

4. Fill-in the following information for the new builder under the Main tab:

- Name

- Location (the full path to `emake`, which is OS dependent)

- Working Directory

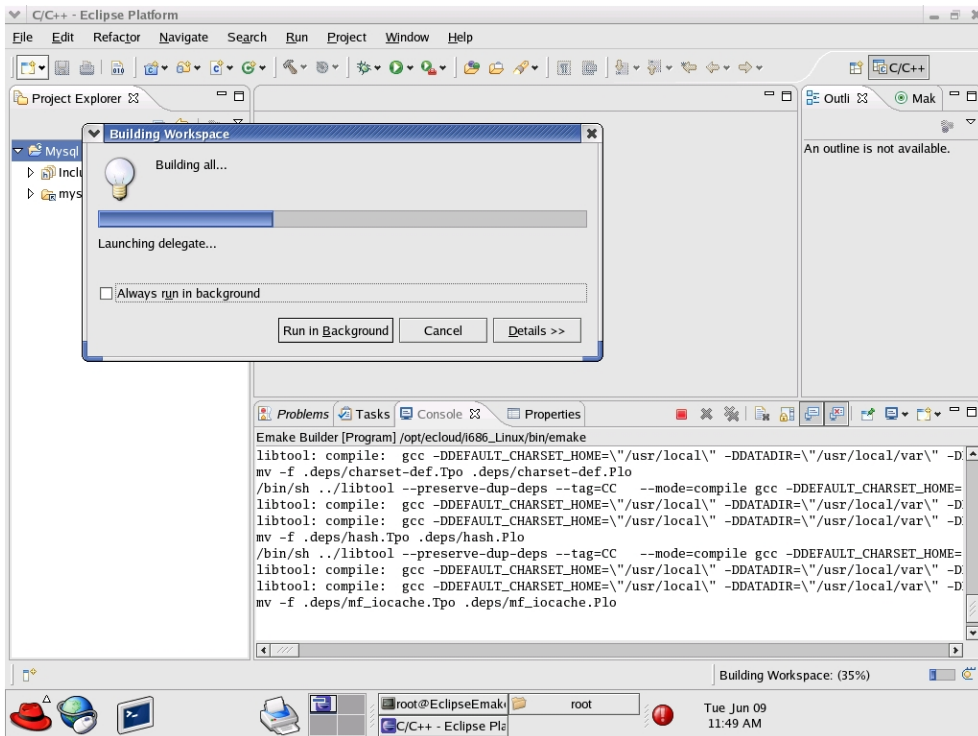- Arguments (arguments are specific to your configuration)

The following screenshot illustrates the Edit Configuration dialog.
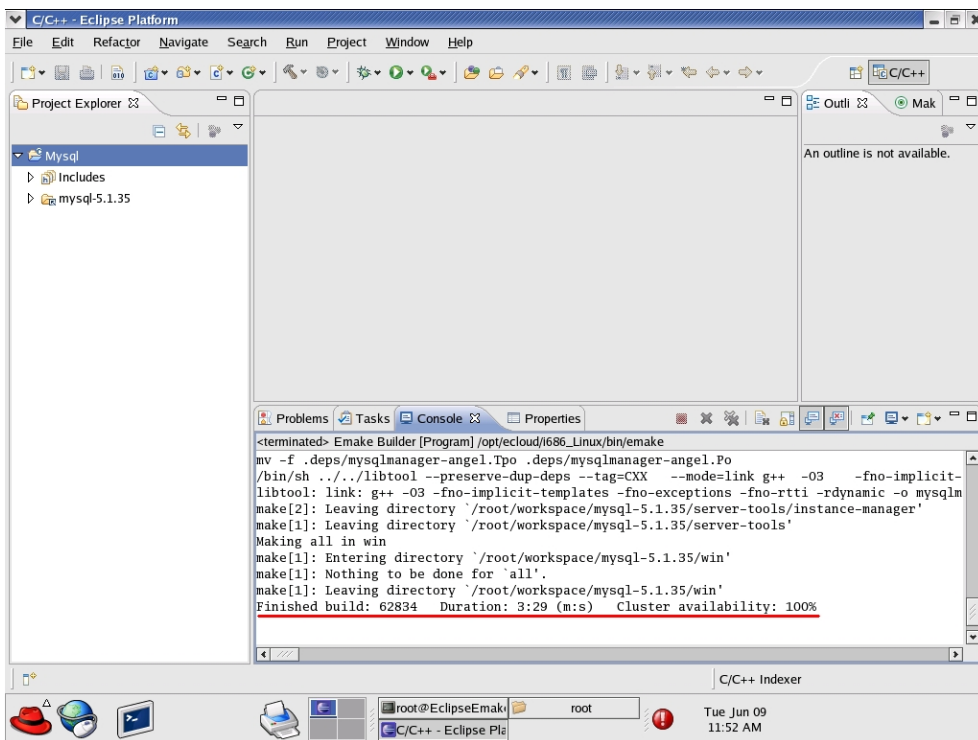


5.  Click the Build Options tab. Enable Run the builder for the following *only*:

    - After a "Clean"

    - During manual builds

    - During auto builds

6.  Click **OK**. Your new builder is displayed in the Builders pane.

7.  Create another builder for "cleans" only. On its Main tab, ensure `clean` is included for Arguments. On its Build Options tab, enable Run the builder for the following *only*:

    - During a "Clean"

8.  Click **OK**. Your second builder is displayed in the Builders pane.

9.  Deselect CDT Builder in the Builders pane and then click **OK**.

Now you can build your project. Click **Project > Build all**.

The following screenshot illustrates a build in progress.

The following screenshot illustrates a successfully completed build.

# Chapter 10: Electrify

Electrify accelerates builds by parallelizing the build process and distributing build steps across clustered resources.

### Known Limitations

- The tool you want to monitor must provide parallel support, such as SCons.

- Electrify does not provide any of Electric Make's dependency detection or correction features. The build tools you use with Electrify must be capable of accurate parallel execution on their own.

- The information written into annotation is more limited with Electrify than what is provided by Electric Make. Electrify annotation provides information only on the commands executed on the cluster, including command lines, file usage, and raw command output. Electrify does not provide information about dependencies, job relationships, targets, or other logical build structure data.

### Recommendations

Electric Cloud has evaluated the following build tools for use with Electrify.

- SCons - Electric Cloud recommends using SCons with Electrify. Using SCons does not have any known limitations.

- Ant - Electric Cloud does not currently recommend using Ant with Electrify.

# Electrify as Part of the Build Process

The following sections describe how to run builds using Electrify.

## Running Electrify on Windows

### Example

```
electrify [args] other tools' command line
```

## Running Electrify on Linux

### Example

```
electrify [args] other tools' command line
```

### SCons Example

```
electrify --emake-cm=<CM name> --electrify-remote=g++:gcc:ranlib:ar scons -j 4
```

### GNU Make C++ Example

```
electrify --emake-cm=<CM name> --electrify-remote=g++:gcc:<any other tools used
in the build> make -j 4 -f makefile
```

# Important Reminders About Electrify

- Ensure `cl.exe`, `link.exe`, and so on, are those of Microsoft Visual Studio. The wrapper application might have changed them to its version.

- Though all of the usual eMake arguments are available, Electrify uses only a subset of them.

- On 64-bit Windows platforms, if you did not install ElectricAccelerator in its default install location, you must specify the complete location of `electrifymon.exe` (including the executable name) for the `EMAKE_ELECTRIFYMON` environment variable.
  For example, if your custom install location is `C:\programs\ECloud`, then set the `EMAKE_ELECTRIFYMON` environment variable using:
  `set EMAKE_ELECTRIFYMON = C:\programs\ECloud \i686_win32\64\bin\electrifymon.exe.`

- On UNIX platforms, electrifymon must locate `electrifymon.so` so it can tell monitored programs to load the monitoring library that reports back to electrifymon. By default, electrifymon looks in the following locations:

| Platform | 32-bit | 64-bit |
|---|---|---|
| Linux | /opt/ecloud/i686_Linux/lib | /opt/ecloud/i686_Linux/64/lib |
| Solaris (SPARC) | /opt/ecloud/sun4u_SunOS/lib | /opt/ecloud/sun4u_SunOS/64/lib |
| Solaris (x86) | /opt/ecloud/i686_SunOS.5.10/lib | /opt/ecloud/i686_SunOS.5.10/64/lib |

  You can override these locations using the following environment variables: `ELECTRIFYMON32DIR` and `ELECTRIFYMON64DIR`.

- On Linux, you can change the mode that Electrify uses for monitoring. Use `--emake-electrify=<mode>`, where mode can be: `preload` for LD_PRELOAD intercept, or `trace` for ptrace.

# Electrify Arguments

All arguments are optional.

| | |
|---|---|
| `--electrify-remote=<x;y>` | x and y are commands that are distributed to the cluster. Use the command's full name, such as `cl.exe`, `link.exe`, `gcc.exe`, without the path. The name is case insensitive. In a Cygwin environment, you can use ':' (colon) instead of ';' (semicolon). |
| | Limited to 2048 characters |
| | Environment variable: `ELECTRIFY_REMOTE` |

| | |
|---|---|
| `--electrify-not-remote=<`*`x;y;`*`>` | x and y are commands that *are **not*** distributed to the cluster. Use the command's full name, such as `cl.exe`, `link.exe`, `gcc.exe`, without the path. The name is case insensitive. In a Cygwin environment, you can use ':' (colon) instead of ';' (semicolon).<br><br>`--electrify-not-remote` and `--electrify-remote` are mutually exclusive.<br><br>If you use `--electrify-not-remote`, all other tools' command lines are executed remotely, by default. Generally, this is not what you want, so to do this, you must add a command to this list. |
| `--electrify-not-intercept=<`*`x;y;`*`>` | xxx and yyy are commands that you do ***not*** want to be monitored, meaning the monitor process does not inject a dll to them and their child processes, so they will ***not*** be distributed.<br><br>Environment variable: `ELECTRIFY_NOT_INTERCEPT` |
| `--electrify-log=<`*`fullpath`*`>` | fullpath is the path of the file you want to log. This logs all process creation and interception information.<br><br>Environment variable: `ELECTRIFY_LOG` |
| `--electrify-localfile=<`*`x`*`>` | (Windows only) Integrates local file access (create, rename, and so on) by locally running tools with the remote file system.<br><br>You can set x to two different flags: `NT` or `y`.<br><br>Set `nt` if you want to monitor undocumented low-level file access Nt functions. This monitors the following functions: NtCreateFile, NtDeleteFile, NtClose, NtWriteFile, and NtSetInformationFile. Though this includes only five functions, their functionality is rich, so this selection includes nearly all scenarios where the local file system changes.<br><br>Set `y` to monitor documented Win32 APIs for file access. This monitors the following Win32 APIs: CreateFileW, CreateDirectoryA, CreateDirectoryExA, CreateDirectoryExW, CreateDirectoryW, DeleteFileA, DeleteFileW, MoveFileA, MoveFileExA, MoveFileExW, MoveFileW, RemoveDirectoryA, RemoveDirectoryW, SetFileAttributesA, and SetFileAttributesW. Though there are many functions, their functionality is less than Nt functions, particularly because some tools such as Cygwin cp.exe use NtCreateFile and so on. In general, use the `y` flag for testing purposes only. |
| `--electrify-allow-regexp=`<br>`<`*`perl-regular-expression`*`>` | (Linux only) Sends all processes to the cluster whose full command-line matches the regular expression. |
| `--electrify-deny-regexp=`<br>`<`*`perl-regular-expression`*`>` | (Linux only) The processes whose command-line match the expression will be executed locally. You can use this after the `--electrify-allow-regexp` option to tune the selection of processes that are sent to the cluster more precisely. |

# Using Whole Command-Line Matching and efpredict

## Whole Command-Line Matching

On Linux, you can use a process's entire command-line to determine if it should be sent to the cluster for execution. Prior to Accelerator v7.0, the only way to discriminate was by the process name. Unfortunately, this did not work well for scripting languages or for languages that run in a VM, such as Java, because the process name is always 'java' no matter which particular program is being executed. This means that in previous versions you could send all Java programs to the cluster or none, and that was the limit of your discretion.

Additional Electrify command-line options:

```
--electrify-allow-regexp=<perl-regular-expression>

--electrify-deny-regexp=<perl-regular-expression>
```

These options allow you to specify which sub-processes should/should not be executed in the cluster. You use a perl-style regular expression that is matched against both the process name and all of its arguments, such as the name of the script or JAR file that is being executed. When Electrify detects that a process is started, it constructs the command-line for that process by joining all of the components of its "argv" array together with spaces and then applying the list of "allow" and "deny" regular expressions in the sequence that they were supplied on the command-line.

### *Whole Command-Line Matching Example:*

You have three processes:

```
java -jar runstep.jar -x86

java -jar otherjar.jar

java -jar runstep.jar -armv7
```

The following options send the first process to the cluster but not the second or third:

```
--electrify-regexp-allow="[^ ]+java\s.*runstep.jar.*"  --electrify-regexp-deny
=".*\-armv7.*"
```

Prior to Accelerator v7.0, you were only able to send all or none.

Initially, a process is considered to be for local execution only, but successive 'allow-regexp' options can change this state if any of them match. Any "deny-regexp" in the sequence will, if it matches, short-circuit the decision immediately and cause that process to be executed locally.

## efpredict

efpredict helps you verify that the expressions you entered actually select the correct processes. Without efpredict, you would need to perform a full build and then examine the annotation file to see if the correct decisions were made. You would have to repeat this process each time there were any mistakes, and it could result in a long process. Instead, you can test settings with efpredict. Provide the same options as you would for Electrify and then enter command-lines into efpredict's standard input to see if it selects them for local or cluster execution. One easy way to do this is to pipe an old build log into efpredict. Check the resulting output visually to see if the desired processes were executed remotely.

### *efpredict Example:*

```
cat oldlog | efpredict --electrify-regexp-allow="[^ ]+java\s.*runstep.jar.*"  --ele
ctrify-regexp-deny=".*\-armv7.*"
```

gives this output:

```
remote_allow: java -jar runstep.jar -x86

remote_deny: java -jar otherjar.jar

remote_deny: java -jar runstep.jar -armv7
```

## Important Notes

- Whole command-line matching is Linux only

- efpredict is Linux only

- If a process was executed by a shell, variables will be expanded, quotes will be removed, and white-space between tokens will be replaced with single spaces before Electrify matches the process. This means that if you look at a process invocation in a shell script or makefile, that might not be the exact text that Electrify sees when it attempts to intercept the invocation of that process.

  For example, in a script you might see:

  ```
  'gcc        "$SOURCE/myfile.c"       -o "$OUTPUT/myfile.o" -c    '
  ```

  but when Electrify intercepts this and tries to reconstruct the command-line, it will see:

  ```
  "gcc src/myfile.c -o out/myfile.o -c".
  ```

  Regular expressions must be written to match what Electrify will be able to see.

- One regexp can match many commands. For example, to send both the gcc and ld commands to the cluster you can use:

  ```
  --electrify-regexp-allow='[^ ]*((gcc)|(ld))(\s.*)?'
  ```

# Additional Electrify Information

### Selecting Commands to Parallelize

A large portion of the build process acceleration will be achieved through parallelizing a small number of specific commands, such as compiling and linking. Expending additional effort to select and parallelize many different additional commands might not result in a significant amount of further acceleration.

**Note:** If a file is created or modified by one or more parallelized commands, then you should parallelize all commands that use that file.

### Using Electrify with GNU Make

If you intend to use Electrify with GNU Make, Electric Cloud recommends using eMake instead. eMake provides superior performance and correctness and full annotation information.

### How an Electrify Build Differs from an eMake Build

An important difference between an eMake build and an Electrify build is what portion of the build activity occurs remotely versus locally. In an eMake build, effectively all build activity (except "#pragma runlocal" jobs) takes place on the cluster, where the EFS is used to monitor file system accesses and propagate changes made by one job to other jobs in the build.

In an Electrify build, a greater portion of the build activity takes place on the local system—at the very least, the build process itself (such as SCons) runs locally. Typically, file system modifications made by processes running locally are "invisible" to Electrify, and therefore to processes running on the cluster—just as "#pragma

runlocal" jobs might make changes that are invisible to eMake. Electrifymon provides a means to update the virtual file system state in Electrify in response to file system modifications made by local processes.

If you know that a build will not run any processes locally that modify the file system, you need not use electrifymon when invoking Electrify. However, some build tools will themselves make changes to the file system (for example, when SCons employs its build-avoidance mechanism by copying a previously built object in lieu of invoking the compiler), so the safest choice is to use electrifymon to start.

In addition to file system monitoring, electrifymon on Windows provides a sophisticated mechanism for intercepting processes invocations and determining which processes to distribute to the cluster. On Linux, process interception is handled by the explicit use of proxy commands.

# Chapter 11: Troubleshooting

The following topics discuss information to assist you with troubleshooting.

**Topics:**

# Agent Errors Establishing the Virtual File System

Agent errors regarding the establishment of the virtual file system for a particular build will be displayed if there are at least three errors. These errors would occur during the initial setup of the agent's build-specific environment but before any particular build step is run on that agent. The most common type of error involves eMake roots or Cygwin mounts, where the virtual file system setup is specific to the build but not to any particular build step.

# Agents Do Not Recognize Changes on Agent Machines

If you manually mount a file system or change automounter settings on agent machines *after* they are started, you must restart the Agents for them to recognize your changes.

# Electric Make Debug Log Levels

This section discusses Electric Make debug log levels. Content was adapted from the "Electric Make debug log levels" blog post on http://blog.melski.net/ and was the most recent information available when the article was posted.

Disclaimer: eMake debug logs are intended for use by Electric Cloud engineering and support staff. Debug logging contents and availability are subject to change in any release, for any or no reason.

Often when analyzing builds executed with Electric Make, all of the information you need is in the annotation file—an easily digested XML file containing data such as the relationships between the jobs, the commands run, and the timing of each job. But sometimes you need more detail, and that is where the eMake debug log is useful.

## Enabling eMake Debug Logging

To enable eMake debug logging, specify this pair of command-line arguments:

`--emake-debug=<value>` specifies the types of debug logging to enable. Provide a set of single-letter values, such as "`jng`".

`--emake-logfile=<path>` specifies the location of the debug log.

## eMake Debug Log Level Descriptions

Available log levels:

| | |
|---|---|
| a: agent allocation | l: ledger |
| c: cache | m: memory |
| e: environment | n: node |
| f: file system | o: parse output |
| g: profiling | p: parse |
| h: history | P: parse avoidance |
| j: job | r: parse relocation |
| L: nmake lexer | s: subbuild |
| | Y: security |

### a: agent allocation

Agent allocation logging provides detailed information about eMake's attempts to procure agents from the Cluster Manager during the build. If you think eMake might be stalled trying to acquire agents, allocation logging will help to understand what is happening.

### c: cache

Cache logging records details about the file system cache used by eMake to accelerate parse jobs in cluster builds. For example, it logs when a directory's contents are added to the cache, and the result of lookups in the cache. Because it is only used during remote parse jobs, you must use it with the `--emake-rdebug=value` option. Use cache logging if you suspect a problem with the *cached local file system*.

### e: environment

Environment logging augments node logging with a dump of the entire environment block for every job as it is sent to an agent. Normally this is omitted because it is quite verbose (could be as much as 32 KB per job). Generally, it is better to use env-level annotation, which is more compact and easier to parse.

### f: file system

File system logging records numerous details about eMake's interaction with its versioned file system data structure. In particular, it logs every time that eMake looks up a file (when doing up-to-date checks, for example), and it logs every update to the versioned file system caused by file usage during the build's execution. This level of logging is very verbose, so it is not usually enabled. It is most often used when diagnosing issues related to the versioned file system and conflicts.

### g: profiling

Profiling debug logging is one of the easiest-to-interpret and most useful types of debug logging. When enabled, eMake emits hundreds of performance metrics at the end of the build. This is a very lightweight logging level and is safe (and advisable) to enable for all builds.

### h: history

History logging prints messages related to the data tracked in the eMake history file—both file system dependencies and autodep information. When history logging is enabled, eMake will print a message every time a dependency is added to the history file, and it will print information about the files checked during up-to-date checks based on autodep data. Enable history logging if you suspect a problem with autodep behavior.

### j: job

Job logging prints minimal messages related to the creation and execution of jobs. For each job you will see a message when it starts running, when it finishes running, and when eMake checks the job for conflicts. If there is a conflict in the job, you will see a message about that, too. If you just want a general overview of how the build is progressing, j-level logging is a good choice.

### L: nmake lexer

eMake uses a generated parser to process portions of NMAKE makefiles. Lexer debug logging enables the debug logging in that generated code. This is generally not useful to end-users because it is too low-level.

### l: ledger

Ledger debug logging prints information about build decisions based on data in the ledger file, as well as updates made to the ledger file. Enable it if you believe the ledger is not functioning correctly.

### m: memory

When memory logging is enabled, eMake prints memory usage metrics to the debug log once per second. This includes the total process memory usage as well as current and peak memory usage grouped into several "buckets" that correspond to various types of data in eMake. For example, the "Operation" bucket indicates the amount of memory used to store file operations; the "Variable" bucket is the amount of memory used for makefile variables. This is most useful when you are experiencing an out-of-memory failure in eMake because it can provide guidance about how memory is being utilized during the build, and how quickly it is growing.

### n:node

Node logging prints detailed information about all messages between eMake and the agents, including file system data and commands executed. Together with job logging, this can give a very comprehensive picture of the behavior of a build. However, node logging is extremely verbose, so enable it only when you are chasing a specific problem.

### o: parse output

Parse output logging instructs eMake to preserve the raw result of parsing a makefile. The result is a binary file containing information about all targets, rules, dependencies, and variables extracted from makefiles read during a parse job. This can be useful when investigating parser incompatibility issues and scheduling issues (for example, if a rule is not being scheduled for execution when you expect). Note that this debug level only makes sense when parsing, which means you must specify it in the `--emake-rdebug` option. The parse results will be saved in the `--emake-rlogdir` directory, named as `parse_jobid.out`. Note that the directory might be on the local disk of the remote nodes, depending on the value you specify.

### p: parse

Parse debug logging prints extremely detailed information about the reading and interpretation of makefiles during a parse job. This is most useful when investigating parser compatibility issues. This output is very verbose, so enable it only when you are pursuing a specific problem. Like parse output logging, this debug level only makes sense during parsing, which means you must specify it in the `--emake-rdebug` option. The parse log files will be saved in the `--emake-rlogdir` directory, named as `parse_jobid.dlog`. Note that the directory might be on the local disk of the remote nodes, depending on the value you specify.

### P: parse avoidance

Parse avoidance logging indicates when a new parse is required, and if so, why it was required.

### r: parse relocation

Parse relocation logging prints low-level information about the process of transmitting parse result data to eMake at the end of a parse job. It is used only internally when we the parse result format is being extended, so is unlikely to be of interest to end-users.

### s: subbuild

Subbuild logging prints details about decisions made while using the eMake subbuild feature. Enable it if you believe that the subbuild feature is not working correctly.

### Y: authentication

Authentication logging is a subset of node logging that prints only those messages related to authenticating eMake to agents and vice-versa. Enable this debug level, if you are having problems using the authentication feature.

# Index