

BUILDING SOFTWARE WITH CLEARCASE

Release 4.0 and later

UNIX Edition

Rational
unifying software teams

800-012615-000

Building Software with ClearCase
Document Number 800-012615-000 December 1999
Rational Software Corporation 20 Maguire Road Lexington, Massachusetts 02421

IMPORTANT NOTICE

Copyright Notice

Copyright © 1992, 1999 Rational Software Corporation. All rights reserved.
Copyright 1989, 1991 The Regents of the University of California
Copyright 1984–1991 by Raima Corporation
Copyright 1992 Purdue Research Foundation, West Lafayette, Indiana 47907

Trademarks

Rational, the Rational logo, Atria, ClearCase, ClearCase MultiSite, ClearCase Attache, Clear DDTs, ClearQuest, ClearGuide, PureCoverage, Purify, Quantify, Rational Rose, and SoDA are trademarks or registered trademarks of Rational Software Corporation in the United States and in other countries. All other names are used for identification purposes only and are trademarks or registered trademarks of their respective companies.

Microsoft, MS, ActiveX, BackOffice, Developer Studio, Visual Basic, Visual C++, Visual InterDev, Visual J++, Visual Studio, Win32, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

Sun, Solaris, and Java are trademarks or registered trademarks of Sun Microsystems, Inc.

Oracle and Oracle7 are trademarks or registered trademarks of Oracle Corporation.

Sybase and SQL Anywhere are trademarks or registered trademarks of Sybase Corporation.

U.S. Government Rights

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the applicable Rational License Agreement and in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct 1988), FAR 12.212(a) 1995, FAR 52.227-19, or FAR 52.227-14, as applicable.

Patent

U.S. Patent Nos. 5,574,898 and 5,649,200 and 5,675,802. Additional patents pending.

Warranty Disclaimer

This document and its associated software may be used as stated in the underlying license agreement, and, except as explicitly stated otherwise in such license agreement, Rational Software Corporation expressly disclaims all other warranties, express or implied, with respect to the media and software product and its documentation, including without limitation, the warranties of merchantability or fitness for a particular purpose or arising from a course of dealing, usage or trade practice.

Technical Acknowledgments

This software and documentation is based in part on BSD Networking Software Release 2, licensed from the Regents of the University of California. We acknowledge the role of the Computer Systems Research Group and the Electrical Engineering and Computer Sciences Department of the University of California at Berkeley and the Other Contributors in its development.

This software and documentation is based in part on software written by Victor A. Abell while at Purdue University. We acknowledge his role in its development.

Contents

Preface	xv
About This Manual	xv
ClearCase Documentation Roadmap	xvi
Typographical Conventions	xvii
Online Documentation	xviii
Technical Support	xix
1. ClearCase Build Concepts	1
1.1 Overview of the ClearCase Build Scheme	2
1.2 Dependency Tracking of MVFS and Non-MVFS Files	3
Automatic Detection of MVFS Dependencies	4
Tracking Non-MVFS Files	4
1.3 Derived Objects and Configuration Records	4
1.4 Build Avoidance	5
Hierarchical Builds	6
Automatic Dependency Detection	6
1.5 Express Builds	7
1.6 Build Auditing with clearaudit	7
1.7 Compatibility with Other make Programs	8
1.8 Parallel Building	8
The Parallel Build Procedure	9
1.9 Building on a Non-ClearCase Host	10
2. Derived Objects and Configuration Records	11
2.1 Derived Objects Overview	11
Derived Object Naming	12
2.2 Configuration Records	13
Configuration Record Example	14

	Contents of a Configuration Record	15
	Header Section	15
	MVFS Objects Section	16
	Non-MVFS Objects Section	16
	Variables and Options Section.....	16
	Build Script Section	17
	Configuration Record Hierarchies	17
	Configuration Record Cache.....	20
2.3	Kinds of Derived Objects.....	20
	Shareable DOs.....	20
	Nonshareable DOs.....	21
	Storage of Derived Objects	21
	Promotion and Winkin	22
	DO Versions	25
2.4	Reuse of DO-IDs	25
2.5	Derived Object Reference Counts	26
	3. Pointers on Using ClearCase Build Tools	29
3.1	Invoking clearmake	29
	A Simple clearmake Build Scenario	30
3.2	Accommodating Build Avoidance.....	32
	Increasing the Verbosity Level of a Build	32
	Handling Temporary Changes in the Build Procedure	32
	Using a Build Options Specification (BOS) File	33
	Handling Targets Built in Multiple Ways.....	34
	Using a Recursive Invocation of clearmake.....	34
	Optimizing Winkin by Avoiding Pseudotargets	35
	Accommodating the Build Tool's Different Name.....	35
3.3	Declaring Source Dependencies in Makefiles	36
	Source Dependencies Declared Explicitly	36
	Explicit Dependencies on Searched-For Sources	37
3.4	Build-Order Dependencies.....	39
3.5	Problems with Forced Builds.....	39

3.6	How clearmake Interprets Double-Colon Rules	39
3.7	Continuing to Work During a Build.....	40
3.8	Using Config Spec Time Rules	41
	Inappropriate Use of Time Rules	42
3.9	Build Sessions, Subsessions, and Hierarchical Builds	43
	Subsessions.....	43
	Versions Created During a Build Session.....	43
	Coordinating Reference Times of Several Builds	44
	Objects Written at More than One Level	44
3.10	Build Auditing and Background Processes.....	45
3.11	Working with Incremental Update Tools.....	46
	Example: Building an Archive	46
	Makefile Restructuring for Incremental Archive Targets	47
	A Note on the Use of ar Keys	49
	Example: Incremental Linking	49
	Additional Incremental-Update Situations	49
3.12	Adding a Version String or Time Stamp to an Executable	50
	Creating a what String.....	50
	Implementing a -Ver Option.....	52
4.	Working with Derived Objects and Configuration Records	53
4.1	Setting Correct Permissions for Derived Objects	53
4.2	Listing and Describing Derived Objects	54
	Listing Derived Objects Created at a Certain Pathname.....	55
	Listing a Derived Object's Kind	55
	Displaying a DO's OID.....	56
	Displaying a Description of a DO Version	56
4.3	Identifying the Views that Reference a Derived Object	57
4.4	Specifying a Derived Object in Commands	57
4.5	Winking in a DO Manually	58
4.6	Preventing Winkin	59
	Preventing Winkin to Your View	59
	Preventing Winkin to Other Views	59
	Using Express Builds to Prevent Winkin to Other Views.....	59

Enabling Express Builds.....	60
Configuring an Existing View for Express Builds.....	60
Creating a New View that Uses Express Builds	60
Preventing Winkin to or from Other Architectures	61
4.7 Converting Derived Objects to View-Private Files.....	61
4.8 Working with DO Versions.....	62
Creating DO Versions	62
Checking In DOs During a Build	62
Accessing DO Versions.....	63
Displaying Configuration Records for DO Versions.....	64
Releasing DOs	65
4.9 Converting Nonshareable DOs to Shared DOs.....	66
Automatic Conversion of Nonshareable DOs to Shareable DOs	66
4.10 Creating Links to Derived Objects	67
4.11 Displaying VOB Disk Space Usage for Derived Objects	68
4.12 Deleting Derived Objects.....	68
Removing Data Containers for Derived Objects.....	68
Scrubbing Derived Objects and Data Containers	69
Degenerate Derived Objects.....	69
Data Container Deleted	69
DO Deleted from VOB Database.....	69
CR Unavailable	70
4.13 Displaying Contents of Configuration Records	70
4.14 Comparing Configuration Records.....	70
4.15 Attaching Labels or Attributes to Versions in a CR	71
4.16 Configuring a View to Select Versions Used to Build a DO	71
4.17 Including a Makefile Version in a Configuration Record.....	71
5. clearmake Makefiles and BOS Files	73
5.1 Makefile Overview	73
5.2 Build Options Specification Files	74
5.3 Format of Makefiles.....	75
Restrictions	75
Libraries	76

Command Echoing and Error Handling	76
Built-In Rules	77
Include Files	77
Macros.....	77
Order of Precedence of Make Macros and Environment Variables ..	77
Make Macros.....	78
Internal Macros.....	79
VPATH Macro	80
Special Targets	80
Special Targets for Use in Makefiles	81
Special Targets for Use in Makefiles or BOS Files.....	81
5.4 Using Makefiles in Snapshot Views	85
5.5 BOS File Entries	86
Standard Macro Definitions.....	86
Target-Dependent Macro Definitions	86
Shell Command Macro Definitions	86
Special Targets	87
Include Directives	87
Comments	87
6. Using clearmake Compatibility Modes	89
6.1 IRIX smake	90
Limitations	92
6.2 IRIX pmake.....	92
6.3 SunOS 5.x make.....	92
VPATH: Searches for Both Targets and Dependencies	93
VPATH Substitutions in Build Scripts	94
Limitations	95
6.4 IBM AIX make	95
6.5 Free Software Foundation Gnu make	96
VPATH Separator Character	96
Parallel Builds.....	96
Supported Gnu Make Command-Line Options.....	96

	Unsupported Gnu Make Command-Line Options	98
	Supported Gnu Make Features.....	99
	Unsupported Gnu Make Features.....	101
6.6	System V Release 3 make	102
	Unsupported Features and Options	102
7.	Using ClearCase to Build C++ Programs	103
7.1	Working with Templates.....	104
	Explicit Instantiation	104
	Alternative to Using the Procedures in This Chapter	105
	Precompiled Header Files	105
7.2	Working with Cfront-Based C++ Compilers.....	106
	Cfront Template Instantiation: Interaction with clearmake.....	106
	Link-Time Cfront Template Instantiation.....	107
	How Link-Time Instantiation Interferes with clearmake.....	108
	Models for Working With Cfront-Based Compilers	109
	The Simple Model.....	109
	How the Simple Model Works	110
	Sample Scenario Using the Simple Model.....	111
	Limitations of the Simple Model.....	111
	The Multiple Repositories Model.....	112
	Using a Recognized Compiler Macro.....	113
	Inserting Special Build Rules in Your Makefile	114
	Using an Alternate (CM-safe) Multiple Repository Model.....	115
	Example Makefile Using the Multiple Repository Model.....	116
	Testing the Makefile.....	116
	How the Multiple Repositories Model Works	117
	Limitations of the Multiple Repositories Model.....	118
	The Forced Instantiation Model	119
	Maintaining Dummy Source Files	120
	Setting Up the Makefile	121
	How the Forced Instantiation Model Works.....	122
	Limitations of the Forced Instantiation Model	123

7.3	Working with SPARCompiler C++	123
	SPARCompiler Template Instantiation: Interaction with clearmake	123
	Setting Up the Repository	124
	Cleaning the Repository	125
	Models for Working with SPARCompiler C++	125
	The Simple Model	126
	How the Simple Model Works.....	126
	Sample Scenario Using the Simple Model	126
	Limitations of the Simple Model	127
	Building Archives that Contain Template Code	128
	Managing Template References.....	128
	Building the Archive	129
	The Multiple Repositories Model	129
	Using a Recognized Compiler Macro	130
	Inserting Special Build Rules in Your Makefile.....	130
	Example Makefile Using the Multiple Repositories Model.....	132
	Testing the Makefile	132
	How the Multiple Repositories Model Works.....	133
	Limitations of the Multiple Repositories Model	133
	Building Archives that Contain Template Code	134
	Multiple Repositories Example.....	135
7.4	Working with the SGI Delta/C++ Compiler	136
	SGI Delta/C++ Compiler Template Instantiation:	
	Interaction with clearmake	137
	Automatic Instantiation	137
	Compile-Time Demand Instantiation	137
	Explicit Instantiation.....	138
7.5	Working with the IBM AIX XLC C++ Compiler	139
	XLC Compiler Template Instantiation: Interaction with clearmake	139
	Models for Working With IBM XLC	140

The Simple Model.....	140
Modifying the Source Files	141
Designing Your Makefile	141
Limitations of the Simple Model.....	141
The Compile-Time Demand Instantiation Model.....	142
Modifying the Source Files	143
Designing Your Makefile	143
Duplicate Symbol Warnings From the Linker	143
The Explicit Instantiation Model.....	144
Modifying the Source Files	144
Designing Your Makefile	145
7.6 Working with the HP aC++ Compiler.....	145
Automatic Instantiation.....	145
Command-Line Option Instantiation	146
Explicit Instantiation	146
8. Using ClearCase Build Tools with Java	147
8.1 ClearCase Build Problems with Java	147
Java Toolkits	148
Scope of the Problems	148
8.2 Benefits of Using make Tools with javac.....	148
Using javac Inside a Makefile	149
Using javac with clearmake Instead of make	149
8.3 Unnecessary Rebuilds and Prevention of Winkin.....	150
8.4 Building Java Applications Successfully	150
Writing Correct Makefiles	150
No Mutually Dependent Files	151
Mutually Dependent Files.....	151
Allowing Rebuilds.....	152
Configuring Makefiles to Behave Like make	153

9. Setting Up a Parallel Build	155
9.1 Overview of Parallel Building.....	155
9.2 Setting Up the Client Host.....	157
Creating Build Hosts Files	157
Setting Up Trust Relationships	158
9.3 Setting Up Server Hosts	159
Idleness Threshold	160
9.4 Starting a Parallel Build.....	160
Setting CCASE_HOST_TYPE in a Shell Startup Script	161
9.5 Preventing Parallel Builds of Targets.....	162
9.6 Preventing Exponential Invocations of abe.....	162
10. Building Software for Multiple Platforms	163
10.1 Issues in Multiple-Platform Development	163
10.2 Handling Source Code Differences	164
10.3 Handling Build Procedure Differences.....	165
Alternative Approach, Using imake.....	166
10.4 Segregating the Derived Objects of Different Variants	167
Approach 1: Use Architecture-Specific Subdirectories.....	167
Approach 2: Use Different Views	168
10.5 Multiple-Architecture Example	168
Scenario.....	169
Defining Architecture-Specific CPP Macros	169
Creating Makefiles in the Source and Build Directories	170
11. Setting Up a Build on a Non-ClearCase Host	173
11.1 Build Scenario	173
11.2 Setting Up an Export View	173
11.3 Mounting the VOB Through the Export View.....	174
11.4 Revising the Build Script.....	174
11.5 Performing an Audited Build in the Export View	176
Index	177

Figures

Figure 1	Building Software with ClearCase: Isolation and Sharing	3
Figure 2	Parallel Building.....	9
Figure 3	Non-ClearCase Access	10
Figure 4	Extended Pathname of a Derived Object	12
Figure 5	Kinds of Information in a Configuration Record	14
Figure 6	Configuration Record Hierarchy	18
Figure 7	Storage of a Shareable Derived Object.....	23
Figure 8	Storage of a Nonshareable Derived Object	24
Figure 9	clearmake Build Scenario.....	31
Figure 10	ClearCase Parallel Build Architecture	156

Preface

ClearCase is a comprehensive software version control and configuration management system.


About This Manual

This manual provides an overview of ClearCase build management features and describes how to use ClearCase's build tools. It is for new or experienced users of ClearCase who are familiar with software build concepts.

If you are not familiar with ClearCase build concepts and tools, read Chapter 1, *ClearCase Build Concepts*, Chapter 2, *Derived Objects and Configuration Records*, and Chapter 3, *Pointers on Using ClearCase Build Tools*.

For information about using ClearCase build tools with C++ programs or with Java tools, read Chapter 7, *Using ClearCase to Build C++ Programs* or Chapter 8, *Using ClearCase Build Tools with Java*.

ClearCase Documentation Roadmap



Typographical Conventions

This manual uses the following typographical conventions:

- *ccase-home-dir* represents the directory into which the ClearCase Product Family has been installed. By default, this directory is `/usr/atria`.
- *attache-home-dir* represents the directory into which ClearCase Attache has been installed. By default, this directory is `C:\Program Files\Rational\Attache`, except on Windows 3.x, where it is `C:\RATIONAL\ATTACHE`.
- **Bold** is used for names the user can enter; for example, all command names, file names, and branch names.
- *Italic* is used for variables, document titles, glossary terms, and emphasis.
- A monospaced font is used for examples. Where user input needs to be distinguished from program output, **bold** is used for user input.
- Nonprinting characters are in small caps and appear as follows: `<EOF>`, `<NL>`.
- Key names and key combinations are capitalized and appear as follows: `F1`, `SHIFT`, `CTRL+G`.
- [] Brackets enclose optional items in format and syntax descriptions.
- { } Braces enclose a list from which you must choose an item in format and syntax descriptions.
- | A vertical bar separates items in a list of choices.
- ... In a syntax description, an ellipsis indicates you can repeat the preceding item or line one or more times. Otherwise, it can indicate omitted information.

NOTE: In certain contexts, ClearCase recognizes “...” within a pathname as a wildcard, similar to “*” or “?”. See the **wildcards_ccase** reference page for more information.

- If a command or option name has a short form, a “medial dot” (·) character indicates the shortest legal abbreviation. For example:

lsc·heckout

This means that you can truncate the command name to **lsc** or any of its intermediate spellings (**lsch**, **lsche**, **lschec**, and so on).

Online Documentation

The ClearCase graphical interface includes a Microsoft Windows-like help system.

There are three basic ways to access the online help system: the **Help** menu, the **Help** button, or the F1 key. **Help**→**Contents** provides access to the complete set of ClearCase online documentation. For help on a particular context, press F1. Use the **Help** button on various dialog boxes to get information specific to that dialog box.

ClearCase also provides access to full “reference pages” (detailed descriptions of ClearCase commands, utilities, and data structures) with the **cleartool man** subcommand. Without any argument, **cleartool man** displays the **cleartool** overview reference page. Specifying a command name as an argument gives information about using the specified command. For example:

```
% cleartool man                                (display the cleartool overview page)
% cleartool man man                             (display the cleartool man reference page)
% cleartool man checkout                       (display the cleartool checkout reference page)
```

ClearCase’s **-help** command option or **help** command displays individual subcommand syntax. Without any argument, **cleartool help** displays the syntax for all **cleartool** commands. **help checkout** and **checkout -help** are equivalent.

```
% cleartool lsprivate -help
Usage: lsprivate [-tag view-tag] [-invob vob-selector] [-long | -short]
               [-size] [-age] [-co] [-do] [-other]
```

Additionally, the online *ClearCase Tutorial* provides important information on setting up a user’s environment, along with a step-by-step tour through ClearCase’s most important features.

Technical Support

If you have any problems with the software or documentation, please contact Rational Technical Support via telephone, fax, or electronic mail as described below. For information regarding support hours, languages spoken, or other support information, click the **Technical Support** link on the Rational Web site at **www.rational.com**.

Your Location	Telephone	Facsimile	Electronic Mail
North America	800-433-5444 toll free or 408-863-4000 Cupertino, CA	408-863-4194 Cupertino, CA 781-676-2460 Lexington, MA	support@rational.com
Europe, Middle East, and Africa	+31-(0)20-4546-200 Netherlands	+31-(0)20-4546-201 Netherlands	support@europe.rational.com
Asia Pacific	61-2-9419-0111 Australia	61-2-9419-0123 Australia	support@apac.rational.com

ClearCase Build Concepts

1

ClearCase® supports *makefile*-based building of software systems, and provides a software build environment closely resembling that of the **make** program. **make** was developed for UNIX systems, and has been ported to other operating systems. You can use ClearCase-controlled files to build software, and use native **make** programs, third-party build utilities, your company's own build programs, or the ClearCase build tools **clearmake** and **clearaudit**.

The **clearmake** build tool provides compatibility with other **make** variants, along with powerful enhancements:

- ▶ *Build auditing*, with automatic detection of source dependencies, including header file dependencies
- ▶ Automatic creation of permanent bill-of-materials documentation of the build process and its results
- ▶ Sophisticated build-avoidance algorithms to guarantee correct results when building in a parallel development environment
- ▶ Sharing of binaries among views, saving both time and disk storage
- ▶ Parallel building, applying the resources of multiple processors and/or multiple hosts to builds of large software systems

The **clearaudit** build tool provides build auditing and creation of bill-of-materials documentation.

clearmake and **clearaudit** are intended for use in *dynamic views*. You can use them in a *snapshot view*, but the features that distinguish them from ordinary **make** programs (build avoidance, build auditing, derived object sharing, and so on) are not enabled in snapshot views.

1.1 Overview of the ClearCase Build Scheme

Developers perform builds, along with all other ClearCase-related work, in *views*. Typically, developers work in separate, private views. Sometimes, a team shares a single view (for example, during a software integration period).

As described in *Developing Software with ClearCase*, each view provides a complete environment for building software that includes a particular configuration of source versions and a private work area in which you can modify source files, and use build tools to create object modules, executables, and so on.

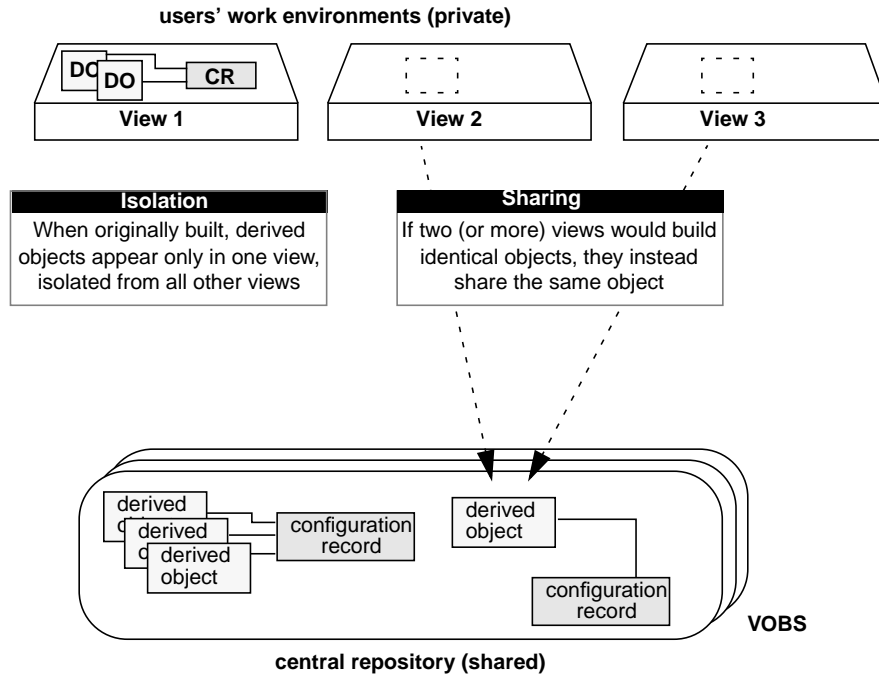
As a build environment, each view is partially isolated from other views. Building software in one view never disturbs the work in another view, even another build of the same program at the same time. However, when working in a dynamic view, you can examine and benefit from work done previously in another dynamic view. A new build shares files created by past builds, when appropriate. This sharing saves the time and disk space involved in building new objects that duplicate existing ones.

You can (but need not) determine what other builds have taken place in a directory, across all dynamic views. ClearCase includes tools for listing and comparing past builds.

The key to this scheme is that the project team's VOBs constitute a globally accessible repository for files created by builds, in the same way that they provide a repository for the source files that go into builds. A file produced by a software build is a *derived object (DO)*. Associated with each derived object is a *configuration record (CR)*, which **clearmake** uses during subsequent builds to decide whether the DO can be reused or shared.

Figure 1 illustrates the ClearCase software build scheme.

Figure 1 Building Software with ClearCase: Isolation and Sharing



Dependency Tracking of MVFS and Non-MVFS Files describes how ClearCase keeps track of the objects produced by software builds. *Build Avoidance* on page 5 describes the mechanism that enables such objects to be shared among views.

1.2 Dependency Tracking of MVFS and Non-MVFS Files

During build-script execution in a dynamic view, a host's *MVFS (multiversion file system)* audits low-level system calls performed on ClearCase data: **create**, **open**, **read**, and so on. Calls involving the following objects are monitored:

- Versions of elements used as build input
- View-private files used as build input (for example, the checked-out version of a file element)
- Files created within VOB directories during the build

Some of these objects are stored in the VOB, and others are view-private files. The view combines them into a *virtual work area*, where they appear to be located in VOB directories. They are called *MVFS files*, because they are accessed through the MVFS.

Automatic Detection of MVFS Dependencies

Because auditing of MVFS files is completely automated, you don't have to keep track of exactly which files are being used in builds. ClearCase does the tracking instead. For example, ClearCase determines which C-language source files referenced with `#include` directives are used in a build. Tracking eliminates the need both to declare such files in the makefile, and for dependency-detection tools, such as **makedepend**.

If you store your build tools (compilers, linkers, and so on) as ClearCase elements and run them from the VOB, they are recorded in the configuration record as implicit detected dependencies.

Tracking Non-MVFS Files

A build can also involve files that are not accessed through VOB directories. Such non-MVFS files are not audited automatically, but are tracked if you declare them as dependencies in a makefile. This tracking enables auditing of build tools that are not stored as ClearCase elements (for example, a C-language compiler), flag files in the user's home directory, and so on. Tracking information on a non-MVFS file includes its absolute path, time stamp, size, and checksum.

1.3 Derived Objects and Configuration Records

When it finishes executing a build script, **clearmake** records the results, including build audit information, in the form of derived objects and configuration records.

A derived object (DO) is a file created in a VOB during a build or build audit with **clearmake**. Each DO has an associated configuration record (CR), which is the bill of materials for the DO. The CR documents aspects of the build environment, the assembly procedure for a DO, and all the files involved in the creation of the DO.

NOTE: All derived objects created by executing a build script have equal status, even though some of them may be explicit build targets, and others may be created as side effects of the build script

(for example, compiler listing files). The term *siblings* describes a group of DOs created by the same script and associated with a single CR.

For more detailed information about DOs and CRs, see Chapter 2, *Derived Objects and Configuration Records*.

1.4 Build Avoidance

clearmake attempts to avoid rebuilding derived objects. If an appropriate derived object exists in the view, **clearmake** reuses that DO. If there is no appropriate DO in the view, **clearmake** looks for an existing DO built in another view that can be winked in to the current view. The search process is called *shopping*.

The process of qualifying a candidate DO is called *configuration lookup*. It involves matching information in the VOB from the candidate DO's config record against the user's current *build configuration*. This process guarantees correct results in a parallel development environment, which the standard time-stamp-based algorithm used by **make** cannot do. Even if an object module is newer than a particular version of its source file, the module may have been built using a different version. In fact, reusing object modules and executables built recently is likely to be incorrect when rebuilding a previous release of an application from old sources. The configuration lookup algorithm that **clearmake** uses guarantees that your builds will be both correct (inappropriate objects are not reused) and optimal (appropriate objects are always reused).

For a DO to be reused or winked in, the build configuration documented in its configuration record must match the current view's build configuration. The build configuration consists of:

Files	The versions of elements listed in the CR must match the versions selected by the view in which the build is performed. Any view-private files or non-MVFS files listed in the CR must also match.
Build procedure	The build options in the CR must match the build options specified on the command line, in the environment, in makefiles, or in build options specification files. The build script listed in the CR must match the script that will be executed if the target is rebuilt. The scripts are compared with all make macros expanded; thus, a match occurs only if the same build options apply (for example, "compile for debugging").

The search ends when **clearmake** finds a DO whose configuration matches the view's current build configuration exactly. In general, a configuration lookup can have three outcomes:

- **Reuse.** If the DO (and its siblings) in the view match the build configuration, **clearmake** keeps them.
- **Winkin.** If a DO built previously matches the build configuration, **clearmake** causes that DO and its siblings to appear in this view. This operation is termed *winkin*.

NOTE: **clearmake** does not contact all views to determine whether they contain DOs that can be winked in. Instead, it uses DO information in the VOB to eliminate inappropriate candidates. Only if it finds a candidate does it contact the containing view to retrieve the DO's config record.

- **Rebuild.** If configuration lookup fails to find a DO that matches the build configuration, **clearmake** executes the target's build script, which creates one or more new DOs, and a new CR.

Reuse and winkin take place only if **clearmake** determines that a newly built derived object would be identical to the existing one. Winkin takes place when two or more views select the same versions of source elements used in a build. For example, you can create a clone view, with the same configuration as an existing view. Initially, the new view sees all the sources, but contains no derived objects. Running **clearmake** winks in many derived objects from the existing view.

Hierarchical Builds

In a hierarchical build, some objects are built and then used to build others. **clearmake** performs configuration lookup separately for each target. To ensure a consistent result, **clearmake** also applies this principle: when a new object is created, all targets that depend on it are rebuilt. Note that winkin does not cause rebuilds of dependencies.

Automatic Dependency Detection

Configuration records enable automatic checking of source dependencies as part of build avoidance. All such dependencies (for example, on C-language header files) are logged in a build's configuration record, whether or not they are explicitly declared in a *makefile*.

1.5 Express Builds

During a regular audited build, **clearmake** writes to the VOB information about a newly built DO. Configuration lookup by future builds uses that information to determine whether or not the DO is a candidate for winkin.

There is a performance tradeoff when you create DOs. While the build is writing the DO information to the VOB database, other users cannot write to the VOB. This performance loss is offset when the DO is used by subsequent builds, which can make the builds faster. However, if the DO is never used by another view, the performance loss is not offset.

ClearCase express builds create derived objects, but do not write information to the VOB. Therefore, these DOs are nonshareable and are not considered for winkin by other views. They can be reused by the view in which they were built.

Express builds offer two advantages over regular builds:

- Scalability: During an express build, write access to the VOB is not blocked by time-consuming DO write operations. More users can build in a VOB without making VOB access slower.
- Performance: Express builds are faster than regular builds, because the build does not write DO information into the VOB.

Which kind of build occurs when you invoke **clearmake** depends on how your view is configured. To use express builds, you must use a dynamic view whose DO property is set to nonshareable. For information on enabling express builds, see *Using Express Builds to Prevent Winkin to Other Views* on page 59.

1.6 Build Auditing with clearaudit

Some organizations, or some developers, may want to use ClearCase build auditing without using the **clearmake** program. Others may want to audit development activities that do not involve *makefiles*. These users can do their work in an *audited shell*, a standard shell with build auditing enabled.

For example, a technical writer produces formatted manual page files by running a shell script that invokes **nroff(1)**. When the script is executed in a dynamic view in a shell created by **clearaudit**, ClearCase creates a single configuration record, recording all the source versions

used. All MVFS files read during execution of the audited shell are listed as inputs to the build. All MVFS files created become derived objects, associated with the single configuration record.

For more information, see the **clearaudit** reference page.

1.7 Compatibility with Other make Programs

Many **make** utilities are available in the multiple-architecture, multiple-vendor world of open systems. ClearCase's **clearmake** program shares features with many of them, and has some unique features.

You can adjust the level of compatibility that **clearmake** has with other **make** programs:

- ▶ Suppress special features of **clearmake**.

Use command options to turn off such features as `winkin`, comparison of build scripts, comparison of detected dependencies, and creation of DOs and CRs. You can turn off configuration lookup altogether, so that the standard time-stamp-based algorithm is used for build avoidance.

- ▶ Enable features of other make programs.

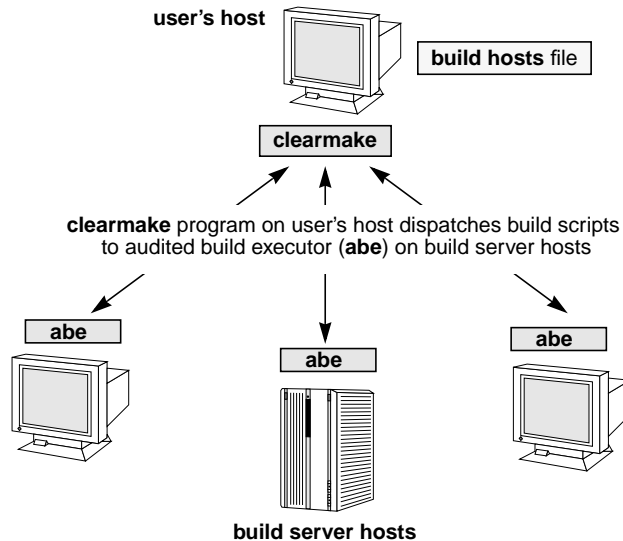
clearmake has several compatibility modes, which provide for partial emulations of popular **make** programs, such as Gnu Make and Sun Make. For more information, see Chapter 6, *Using clearmake Compatibility Modes*.

To achieve absolute compatibility with other **make** programs, you can actually use them to perform builds. However, builds with a standard **make** do not provide build auditing, configuration lookup, or sharing of DOs. The MVFS files that the build creates are view-private files, not derived objects. However, you can execute the **make** program in a **clearaudit** shell, which performs an audited build.

1.8 Parallel Building

clearmake includes support for *parallel* building (concurrent execution of a set of build scripts on one or most hosts). A command option specifies the number of hosts to use; host names to use are read from a *build hosts* file (Figure 2).

Figure 2 Parallel Building



For example, you can perform a three-way build, all of whose processes execute on a single multiprocessor server; an overnight build can be distributed across all the workstations in the local network.

The Parallel Build Procedure

Before starting a build, **clearmake** parses the makefile to analyze the hierarchy of intermediate targets needed to build the final target. In a serial build, **clearmake** constructs each target (or reuses an existing DO) before continuing the analysis of subsequent builds.

In a parallel build, when **clearmake** detects that a target is out of date, it dispatches the build script for that target to one of the hosts listed in the build hosts file. **clearmake** continues to analyze the build hierarchy to detect other targets that can be built at the same time. It dispatches the build script for subsequent targets to one of the hosts in the build hosts file. The total number of build scripts being executed at any particular time is equal to or less than the number you specify with **-J**. Each target is built as soon as system resources on the build host allow.

Execution of build scripts is managed on each remote host by the ClearCase *audited build executor* (**abe**), which is invoked through standard remote-shell facilities.

For more information about parallel builds, see Chapter 9, *Setting Up a Parallel Build*.

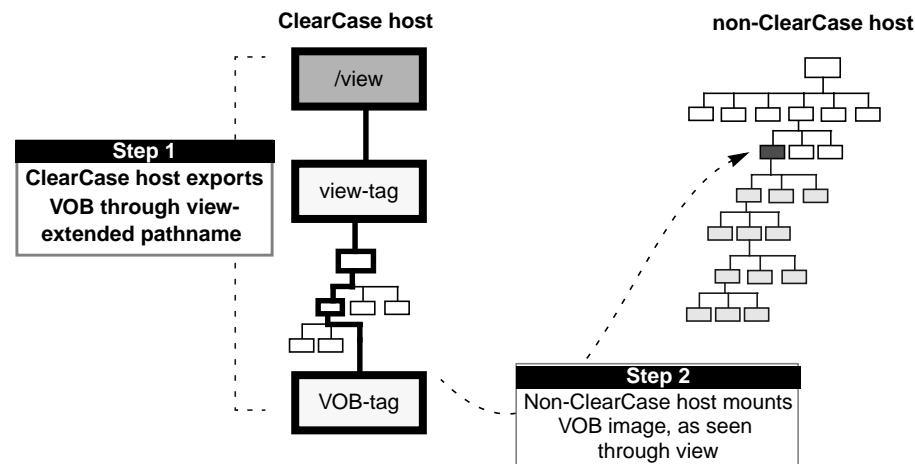
1.9 Building on a Non-ClearCase Host

Many organizations develop multiple variants of their products, targeted at different platforms. If ClearCase is not available for some of the platforms, you can still build all the required product variants by using either cross-development or *non-ClearCase access*:

- Cross-development allows you to perform a build on a supported host to produce executables for the unsupported host.
- Non-ClearCase access allows hosts on which ClearCase is not installed to access ClearCase data, using standard network file-sharing services (such as NFS).

Non-ClearCase access takes advantage of view *transparency*. Through automatic version-selection, a view makes any VOB appear to be a standard directory tree. Any such VOB image can be exported to a non-ClearCase host (Figure 3).

Figure 3 Non-ClearCase Access



Developers on the non-ClearCase host can perform builds within these VOB image directory trees, using native **make** utilities or other local build tools. The object modules and executables produced by such builds are stored in the view used to export the VOB.

For more information about building on a non-ClearCase host, see Chapter 11, *Setting Up a Build on a Non-ClearCase Host*.

Derived Objects and Configuration Records

2

This chapter describes derived objects and configuration records. ClearCase creates derived objects and configuration records only if you build in a dynamic view with one of the ClearCase build tools. For information on managing derived objects and configuration records, see Chapter 4.

2.1 Derived Objects Overview

As described in Chapter 1, *ClearCase Build Concepts*, derived objects are created during builds with ClearCase build tools. They are used for build avoidance and derived object sharing.

In a parallel-development environment, it is likely that many DOs with the same pathname will exist at the same time. For example, suppose that source file `msg.c` is being developed on three branches concurrently, in three different views. ClearCase builds performed in those three views produce object modules named `msg.o`. Each of these is a DO, and each has the same standard pathname, for example, `/vobs/proj/src/msg.o`.

In addition, each DO can be accessed with ClearCase extended names:

- Within each dynamic view, a standard UNIX pathname accesses the DO referenced by that view. This is another example of ClearCase's *transparency* feature.

`msg.o`

(the DO in the current view)

- You can use a view-extended pathname to access a DO in any view:

```
/view/drp/vobs/proj/src/msg.o           (the DO in view drp)
/view/2_integ/vobs/proj/src/msg.o       (the DO in view R2_integ)
```


Derived Object Naming

No name collisions occur among derived objects built at the same pathname, because each DO is cataloged in the VOB database with a unique identifier, its *DO-ID*. The DO-ID references a DO independently of views. The **lsdo** (list derived objects) command can list all DOs created at a specified pathname, regardless of which views (if any) can select them:

```
% cleartool lsdo hello.o
07-May.16:09 akp "hello.o@@07-May.16:09.623" on neptune
06-May.12:47 akp "hello.o@@06-May.12:47.539" on neptune
01-May.21:49 akp "hello.o@@01-May.21:49.282" on neptune
03-Apr.21:40 akp "hello.o@@01-May.21:40.226" on neptune
```

Together, a DO's standard name (**hello.o**) and its DO-ID (**07-May.16:09.623**) constitute a **VOB-extended pathname** to that particular derived object. (The extended naming symbol is host specific; most organizations use the default value, @@.)

Figure 4 Extended Pathname of a Derived Object



Standard software must access a DO through a dynamic view, using a standard pathname or view-extended pathname. You can use such names with debuggers, profilers, **rm**, **tar**, and so on. Only ClearCase programs can reference a DO using a VOB-extended pathname, and only the DO's metadata is accessible in this way:

ClearCase commands can use the extended pathname of a derived object

```
% cleartool describe hello@@07-Mar.11:40.217
created 07-Mar-96.11:40.217 by Allison K. Pak (akp.users@phobos)
references: 1 => cobalt:/usr1/tmp/akp/tut/old.vws
```



```

% cleartool catcr hello@@07-Mar.11:40.217
Target hello built by akp.user
Host "cobalt" running SunOS 4.1.3 (sun4m)
Reference Time 07-Mar-96.11:40:41, this audit started
  07-Mar-96.11:40:46
View was cobalt:/var/tmp/akp/tut/old.vws
Initial working directory was /vobs/akp_cobalt_hw/src
-----
MVFS objects:
-----
/vobs/akp_cobalt_hw/src/hello@@07-Mar.11:40.217
/vobs/akp_cobalt_hw/src/hello.o@@07-Mar.11:40.213
/vobs/akp_cobalt_hw/src/util.o@@07-Mar.11:40.215
-----
Variables and Options:
-----
MKTUT_CC=cc
-----
Build Script:
-----
  cc -o hello hello.o util.o
-----

```

Standard programs cannot use the extended pathname of a derived object

```

% ls hello.@@07-Mar.11:40.217
Cannot access hello@@07-Mar.11:40.217: No such file or directory.

```

EXCEPTION: You can use a view-extended pathname with the **winkin** command, to make the file system data of any DO available to your view. See *Winking in a DO Manually* on page 58.

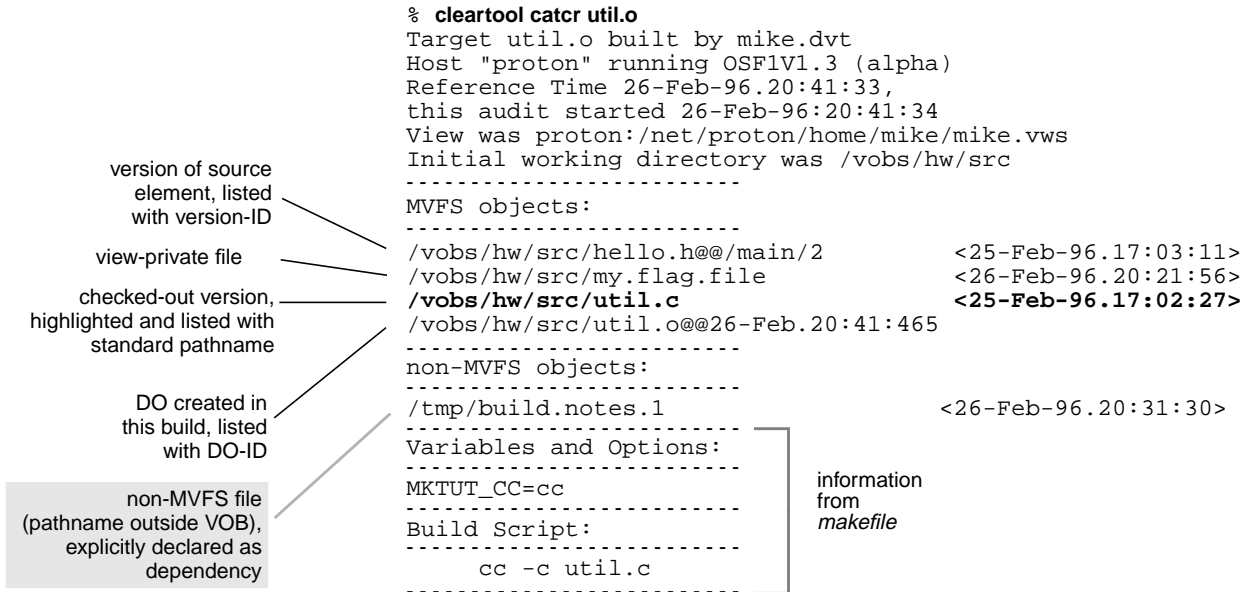
2.2 Configuration Records

A configuration record (CR) is the bill of materials for a derived object or set of DOs. The CR documents aspects of the build environment, the assembly procedure for a DO, and all the files involved in the creation of the DO.

Configuration Record Example

The **catcr** command displays the configuration record of a specified DO. Figure 5 shows a CR, with annotations to indicate the various kinds of information in the listing.

Figure 5 Kinds of Information in a Configuration Record



Some notes on Figure 5:

- Directory versions. By default, **catcr** does not list versions of the VOB directories involved in a build. To list this information, use the **-long** option:

cleartool catcr -long util.o

```
directory version /vobs/hw/.@@/main/1    <25-Feb-96.16:59:31>
directory version /vobs/hw/src@@/main/3  <26-Feb-96.20:53:07>
...
```

- Declared dependencies. One of ClearCase's principal features is the automatic detection of source dependencies on MVFS files: versions of elements and objects in view-private storage. In addition, a CR includes non-MVFS objects that are explicitly declared as dependencies in the *makefile*. Figure 5 shows one such declared dependency, on file **build.notes.1**, located in the non-VOB directory **/tmp**.

- ▶ Listing of checked-out versions. Checked-out versions of file elements are highlighted. Checked-out versions of directory elements are listed like this:

```
directory version /vobs/hw/src@@/main/CHECKEDOUT    <26-Feb-96.17:05:23>
```

When the elements are subsequently checked in, a listing of the *same* configuration record shows the updated information. For example,

```
/vobs/hw/src/util.c                                <25-Feb-96.17:02:27>
```

becomes

```
/vobs/hw/src/util.c@@/main/4                      <25-Feb-96.17:02:27>
```

The actual configuration record contains a ClearCase-internal identifier for each MVFS object. After the version is checked in, **catcr** lists that object differently.

NOTE: The time stamps in the configuration record are for informational purposes and are not used by ClearCase during rebuild or winkin decisions. ClearCase uses OIDs to track versions used in builds.

Contents of a Configuration Record

The following sections describe the contents of configuration records.

Header Section

As displayed by **catcr**, the header section of a CR includes the following lines:

- ▶ Makefile target associated with the build script and the user who started the build:

```
Target util.o built by akp.dvt
```

For a CR produced by **clearaudit**, the target is `ClearAudit_Shell`.

- ▶ Host on which the build script was executed, along with information from the **uname(2)** system call:

```
Host 'neon' running SunOS 5.5.1 (sun4m)
```

- ▶ Reference time of the build (the time **clearmake** or **clearaudit** began execution), and the time when the build script for this particular CR began execution:

```
Reference Time 15-Sep-93.08:18:56, this audit started 15-Sep-93.08:19:00
```

In a hierarchical build, involving execution of multiple build scripts, all the resulting CRs share the same reference time. (For more on reference time, see the **clearmake** reference page.)

- ▶ View storage directory of the view in which the build took place:

```
View was neptune:/home/akp/views/930825.vws
```

- ▶ Working directory at the time build script execution or **clearaudit** execution began:

```
Initial working directory was /vobs/hw/src
```

MVFS Objects Section

An *MVFS object* is a file or directory in a VOB. The MVFS Objects section of a CR includes this information:

- ▶ Each MVFS file or directory read during the build. This includes versions of elements and view-private files used as build input, checked-out versions of file elements, DOs read, and any tools or scripts used during the build that are under version control.
- ▶ Each derived object produced by the target rebuild.

Non-MVFS Objects Section

A *non-MVFS object* is an object not accessed through a VOB (compiler, system-supplied header file, temporary file, and so on). The Non-MVFS Objects section of a CR includes each non-MVFS file that appears as an explicit dependency in the makefile or is a dependency inferred from a suffix rule. See *Declaring Source Dependencies in Makefiles* on page 36.

This section is omitted if there are no such files or if the CR was produced by **clearaudit**.

Variables and Options Section

The Variables and Options section of a CR lists the values of make macros referenced by the build script and command-line options.

This section is omitted from a CR produced by **clearaudit**.

Build Script Section

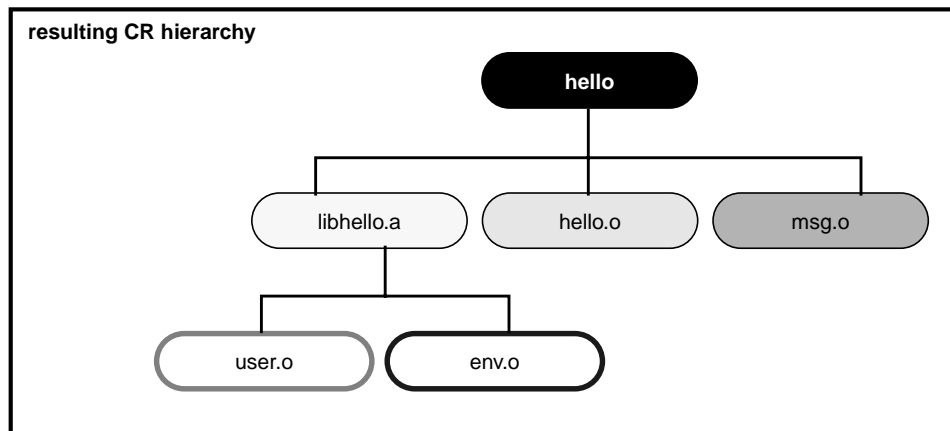
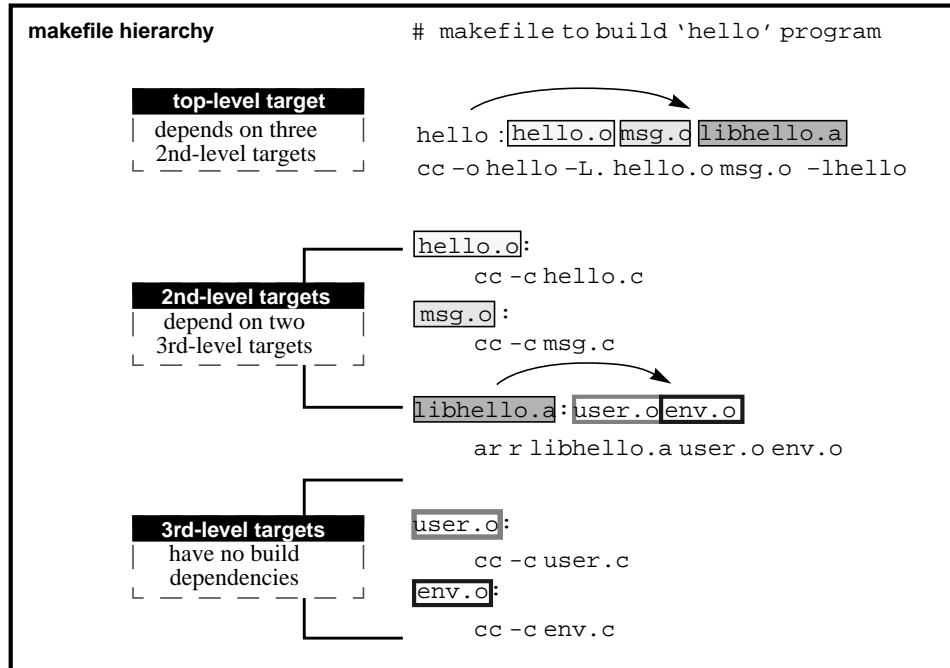
The Build Script section of a CR lists the script that was read from a makefile and executed by **clearmake**.

This section is omitted from a CR produced by **clearaudit**.

Configuration Record Hierarchies

A typical makefile has a hierarchical structure. Thus, running **clearmake** once to build a high-level target can cause multiple build scripts to be executed and, accordingly, multiple CRs to be created. Such a set of CRs can form a *configuration record hierarchy*, which reflects the structure of the makefile.

Figure 6 Configuration Record Hierarchy



An individual parent-child link in a CR hierarchy is established in either of the following ways:

- In a target/dependencies line

For example, the following target/dependencies line declares derived objects **hello.o**, **msg.o**, and **libhello.a** to be build dependencies of derived object **hello**:

```
hello: hello.o msg.o libhello.a
...
```

Accordingly, the CR for **hello** is the parent of the CRs for the **.o** files and the **.a** file.

- In a build script

For example, in the following build script, derived object **libhello.a** in another directory is referenced in the build script for derived object **hello**:

```
hello: $(OBJS)
    cd ../lib ; $(MAKE) libhello.a
    cc -o hello $(OBJS) ../lib/libhello.a
```

Accordingly, the CR for **hello** is the parent of the CR for **libhello.a**.

NOTE: The recursive invocation of **clearmake** in the first line of this build script produces a separate CR hierarchy, which is not necessarily linked to the CR for **hello**. The second line of the build script links the CR for **../lib/libhello.a** with that of **hello**.

The **catcr** and **diffcr** commands have options for handling CR hierarchies:

- By default, they process individual CRs.
- With the **-recurse** option, they process the entire CR hierarchy of each derived object specified, keeping the individual CRs separate.
- With the **-flat** option, they combine (or flatten) the CR hierarchy of each specified derived object.

Some ClearCase features process entire CR hierarchies automatically. For example, when the **mklablel** command attaches version labels to all versions used to build a particular derived object (**mklablel -config**), it uses the entire CR hierarchy of the specified DO. Similarly, ClearCase maintenance procedures do not scrub the CR associated with a deleted DO if it is a member of the CR hierarchy of a higher-level DO.

Configuration Record Cache

When a derived object is created in a view, both its data container and its associated configuration record are stored in the view's private storage area. The CR is stored in the view database, in compressed format. To speed configuration lookup during subsequent builds in this view, a compressed copy of the CR is also cached in a view-private file, `.cmake.state`, located in the directory that was current when the build started.

When a DO is winked in for the first time, the associated CR moves from the view's private storage area to the VOB database, as shown in Figure 7 and Figure 8.

2.3 Kinds of Derived Objects

The following sections describe the kinds of DOs and their lifecycles.

During a regular build, ClearCase build tools create shareable derived objects. During an express build, they create nonshareable derived objects. Both kinds of DOs have configuration records, but only shareable DOs can be winked in by other views.

Shareable DOs

When a ClearCase build tool creates a shareable DO, it creates a configuration record for the DO and writes information about the DO into the VOB. (At this point, the DO is shareable but unshared.) Builds in other views use this information during configuration lookup. If the build determines that it can wink in an existing DO, it contacts the view containing the DO and promotes the DO to the VOB. (The DO is now shareable and shared.)

As noted in *Express Builds* on page 7, you must consider whether the performance benefit of winking in DOs is worth the performance cost of making them available for winking.

NOTE: The process of looking for a DO to wink in does not slow performance, because the build tool uses an efficient algorithm to eliminate mismatches.

The configuration lookup process cannot guarantee that the DO is suitable for use. The process uses details in the config record to determine whether a DO is suitable for winking, but the config record does not record all parameters of a build. For example, a config record may list only a

compiler's name and the options used. If two builds use incompatible compilers with the same name, unwanted winking from one build to the other can occur.

NOTE: To minimize occurrences of incorrect winking, all developers must use the same set of tools. For example, put your build tools under version control and always run them from the VOB.

Nonshareable DOs

During an express build, the ClearCase build tool creates nonshareable DOs. The build tool creates a configuration record for the DO, but does not write information about the DO into the VOB. Because scanning the information in the VOB is the only method other builds use to find DOs, other builds cannot winkle in nonshareable DOs. However, a nonshareable DO can be reused by the view in which it was built.

A nonshareable DO can have shareable sub-DOs, but not shareable siblings. A nonshareable DO can be built using a winked-in shareable DO. (However, a shareable DO cannot have nonshareable sub-DOs or siblings.)

For information on enabling express builds, see *Using Express Builds to Prevent Winking to Other Views* on page 59.

You can use the same commands that you use with shareable DOs on nonshareable DOs, but some commands work differently on the two kinds of DOs. The reference pages for the commands describe the differences.

Storage of Derived Objects

When a DO is created, its data container is located in the view storage area. For a shareable DO, the ClearCase build tool creates the VOB database object for the DO, and also writes to the VOB information about the DO that can be used during configuration lookup. A nonshareable DO has no VOB database object, and the build tool does not write any configuration lookup information into the VOB.

A DO consists of the following parts:

- VOB database object (shareable DOs only)—Each DO is cataloged in the VOB database, where it is identified by an extended name that includes both its standard pathname (for example, `/vobs/hw/src/hello.c`) and a unique DO-ID (for example, `23-Feb.08:41.391`).

- Data container—The data portion of a derived object is stored in a standard file within a ClearCase storage area. This file is called a *data container*; it contains the DO's file system data.
- Configuration record—Actually, a CR is associated with a DO; it is not part of the DO itself. More precisely, a CR is associated with the entire set of sibling DOs created by a particular invocation of a particular build script. See *Configuration Records* on page 13.

When a shareable DO is first created, it is unshared:

- It appears only in that view.
- Its data container is a file in the view's private storage area.
- **clearmake** writes information about the DO into the VOB.

Promotion and Winkin

The first time a shareable derived object is winked in by another dynamic view, or when either kind of DO is promoted manually with a **winkin** or **view_scrubber -p** command, its status changes to shared:

- Its data container is *promoted* to a derived object storage pool in the VOB.
- (shareable DOs only) If the winkin was done by the build tool or the command was executed in another view, the DO now appears in two dynamic views.

When the winkin occurs during a **clearmake** build:

- The dynamic view to which the DO is winked in, and all other views to which the DO is subsequently winked in, use the data container in VOB storage.
- The original view continues to use the data container in view storage. (The **view_scrubber** utility removes this asymmetry, which causes all dynamic views to use the data container in VOB storage.)

When the winkin is done with the **winkin** or **view_scrubber -p** command, the data container in the view is removed after it is promoted to VOB storage. The original view and all other views to which the DO is subsequently winked in use the data container in VOB storage.

Figure 7 Storage of a Shareable Derived Object

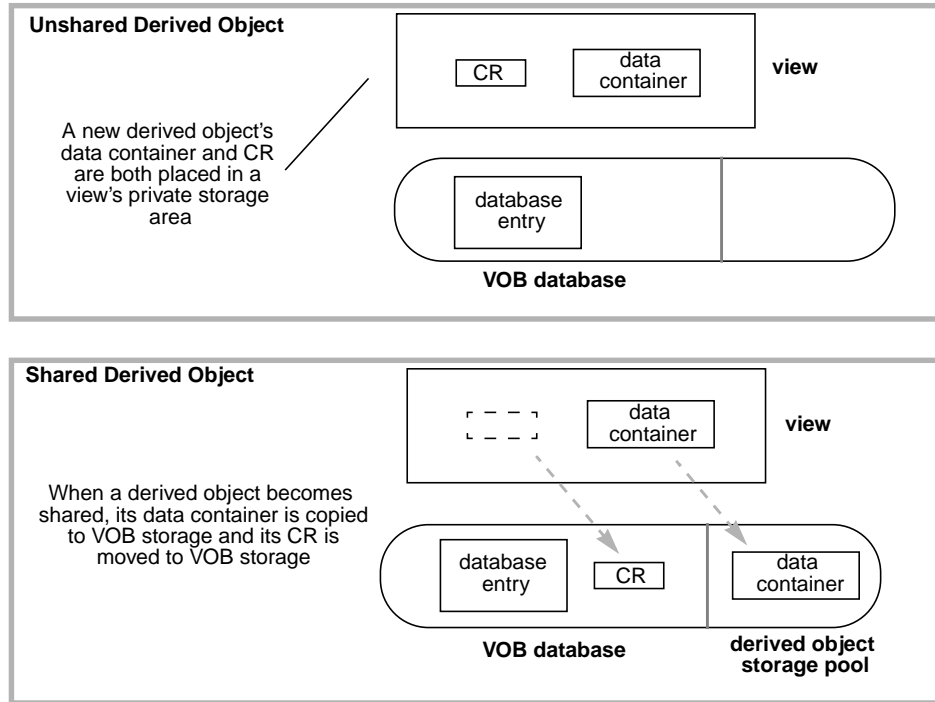



Figure 8 Storage of a Nonshareable Derived Object



After a derived object is winked in, it remains shared, no matter how many times it is winked in to additional dynamic views, and even if subsequent rebuilds or deletion commands cause it to appear in only one dynamic view (or zero views).

When a derived object's data container is in the VOB, any number of views can share the derived object without having to communicate with each other directly. For example, view **alpha** can be unaware of views **beta** and **gamma**, with which it shares a derived object. The hosts on which the view storage directories are located need not have network access to each other's disk storage.

For more information, see the **winkin** and **view_scrubber** reference pages.

DO Versions

You can check in a derived object as a version of an element, creating a *DO version*. Other versions of such an element can also be, but need not be, derived objects. A DO version behaves like both a version and a derived object:

- You can use its *version-ID* to reference it as both a VOB database object and a data file.
- You can apply a version label to it and reference it using that label.
- You can display its configuration record with **catcr** or compare the CR to another with **diffcr**.
- A **clearmake** build can wink it in if the element is located at the same pathname where the DO was originally built.
- You can wink it in with a **winkin** command.
- The **describe** command lists it as a `derived object version`. (The **lsdo** command does not list it at all.)

For more information on DO versions, see *Working with DO Versions* on page 62.

2.4 Reuse of DO-IDs

The DO-ID for a shareable derived object is guaranteed to be unique within the VOB, for all views. That is, if you delete a shareable DO, its numeric file-name extension is not reused (unless you reformat the VOB that contains it).

The DO-ID for a DO created by an express build (a nonshareable derived object) is unique only at a certain point in time. If you delete a nonshareable DO, the ClearCase build tools can reuse its numeric file-name extension. (Because ClearCase tracks derived objects using their VOB database identifiers, so no build confusion will occur if a file-name extension is reused.)

DO-IDs change when any of these events occur:

- The DO passes its first birthday. The time stamp changes to include the year the DO was created:

<code>util.o@@15-Jul.15:34.8896</code>	<i>(when first created)</i>
<code>util.o@@15-Jul-1998.8896</code>	<i>(after a year)</i>

- You convert a nonshareable DO to a shareable DO. (See *Converting Nonshareable DOs to Shared DOs* on page 66.)
- You process a VOB's database with **reformatvob**. All DO-IDs receive new numeric file-name extensions:

util.o@@15-Jul.15:34.8896	(before reformatvob)
util.o@@15-Jul.17:08.734	(after reformatvob)

The configuration record reflects these DO-ID changes.

2.5 Derived Object Reference Counts

A DO's reference count is the number of times the derived object appears in ClearCase dynamic views throughout the network. ClearCase also tracks the identifiers for the views that reference the DO. When a new derived object is created, **clearmake** sets its reference count to 1, indicating that it is visible in one view. Thereafter, each winkin of the DO to an additional view increments the reference count.

The **lsdo -long** command lists the reference count and referencing views for a DO. For example:

```
cleartool lsdo -long
01-Sep-99.18:56:45      Suzanne Lee (sgl.user@neon)
  create derived object "file.txt@@01-Sep.18:56.2147483683"
  size of derived object is: 10
  last access: 01-Sep-99.18:56:46
  references: 1 => neon:/home/sgl/views/sgl_test.vws
01-Sep-99.19:03:19      Suzanne Lee (sgl.user@neon)
  create derived object "util@@01-Sep.19:03.81"
  size of derived object is: 10
  last access: 01-Sep-99.19:03:33
  references: 2 (shared)
=> neon:/home/sgl/views/sgl_test.vws
=> neon:/home/sgl/views/point_of.vws
```

For a nonshareable DO, the reference count is always 1.

You can also create OS-level hard links to an existing shareable DO, each of which increments the reference count. Such additional hard links are sometimes subject to winkin:

- If the additional hard link was created in the same build script as the original DO, a winkin of the DO during a subsequent **clearmake** build causes a winkin of the additional hard link.

- Additional hard links that you create manually are not winked in during subsequent builds.

A reference count can also decrease. When a program running in any of the views that reference a shared derived object overwrites or deletes that object, the link is broken and the reference count is decremented. That is, the program deletes the view's reference to the DO, but the DO itself remains in VOB storage. This occurs most often when a compiler overwrites an old build target. You can also remove the derived object with a standard **rm** command, or if the makefile has a **clean** rule, a **clearmake clean** invocation.

A derived object's reference count can become zero. For example, suppose you build program **hello** and rebuild it a few minutes later. The second **hello** overwrites the first **hello**, decrementing its reference count. Because the reference count probably was 1 (no other view has winked it in), it now becomes 0. Similarly, the reference counts of old DOs, even of DOs that are widely shared, eventually decrease to zero as development proceeds and new DOs replace the old ones.

The **lsdo** command ignores such DOs by default, but you can use the **-zero** option to list them:

```
% cleartool lsdo -zero -long hello.o
.
.
08-Mar-96.12:47:54      Allison K. Pak (akp.user@cobalt)
  create derived object "hello.o@@08-Mar.12:47.259"
  references: 0
  ...
```

A derived object that is listed with a `references: 0` annotation does not currently appear in any view. However, some or all of its information may still be available:

- If the DO was ever promoted to VOB storage, its data container is still in the VOB storage pool (unless it has been scrubbed), and its CR is still in the VOB database. You can use **catcr** and **diffcr** to work with the CR. You can get to its file system data by performing a **clearmake** build in an appropriately configured view, or by using the **winkin** command.
- If the DO was never promoted, its CR may be gone forever. Until the scrubber runs and deletes the data container, the **catcr** command prints the message `Config record data no longer available for DO-pname`.

Pointers on Using ClearCase Build Tools

3

This chapter presents some pointers on making best use of **clearmake**.

3.1 Invoking **clearmake**

You can invoke **clearmake** from the command line or from the ClearCase File Browser (**xclearcase**). The command-line interface is designed to be as similar as possible to other **make** variants. Lowercase, single-letter command options have their familiar meanings. For example:

-n	No-execute mode
-f	Specify name of makefile
-u	Unconditional rebuild

clearmake recognizes additional options (also one letter, but uppercase) that control its enhanced functionality: configuration lookup, creation of configuration records and derived objects, parallel building, and so on. For a complete description, see the **clearmake** reference page.

You can run **clearmake** as a background process or invoke it from a shell script. (In **clearmake** output, some names are in bold, for clarity. On some architectures, running **clearmake** in the background suppresses the bold, but no characters are lost.)

A Simple `clearmake` Build Scenario

`clearmake` is designed to let developers in makefile-based build environments continue working in their accustomed manner. The following simple build scenario demonstrates how little adjustment is required to begin building with `clearmake`.

1. Set a view. Because working with ClearCase data requires a view context, it makes sense to set a view before starting a build.

(Strictly speaking, this is not required: if your process has a *working directory view* context, but not a *set view* context, `clearmake` sets the view to the working directory view by executing a `cleartool setview -exec clearmake` command. If your process has a working directory view context and a set view context, `clearmake` uses the working directory view. See *Setting a View* on page 78.)


2. Go to a development directory within any VOB.
3. Edit some source files. Typically, you need to edit some sources before performing a build; accordingly, you check out some file elements and revise the checked-out versions.
4. Start a build. You can use your existing makefiles, but invoke `clearmake` instead of your standard `make` program. For example:

<code>clearmake</code>	<i>(build the default target)</i>
<code>clearmake cwd.o libproj.lib</code>	<i>(build one or more particular targets)</i>
<code>clearmake -k monet CFLAGS=-g</code>	<i>(use standard options and make-macro overrides)</i>

(We recommend that you avoid specifying make-macro overrides on the command line. See *Using a Build Options Specification (BOS) File* on page 33.)

`clearmake` builds targets (or avoids building them) in a manner similar to, but more sophisticated than, other make variants. Figure 9 illustrates the results of a typical build.

Figure 9 clearmake Build Scenario



clearmake builds a new derived object for each checked-out source file, because no other view's build could possibly have used your checked-out version.

Note that **clearmake** does not attempt to verify that you have actually edited the file; the checkout makes a rebuild necessary. As you work, saving a file or invoking **clearmake** causes a rebuild of the updated file's dependents, in the standard make manner.

For source files that you have not checked out, **clearmake** may or may not build a new derived object:

- It may reuse a derived object that appears in your view, produced by a previous build.
- It may wink in an existing derived object built in another view. (It's even possible that a winked-in DO was originally created in your view, shared, then deleted from your view—for example, by a **make clean** rule.)
- Changes to other aspects of your build environment may trigger a **clearmake** rebuild: revision to a header file; change to the build script, use of a make-macro override; change to an environment variable used in the build script.

3.2 Accommodating Build Avoidance

When you first begin to build software systems with ClearCase, the fact that **clearmake** uses a different build-avoidance algorithm than other **make** variants may occasionally surprise you. This section describes several such situations and presents simple techniques for handling them.

Increasing the Verbosity Level of a Build

If you don't understand **clearmake**'s build-avoidance decisions, use the **-v** (somewhat verbose) or **-d** (extremely verbose) option. Equivalently, set environment variable **CCASE_VERBOSITY** to **1** or **2**, respectively.

Handling Temporary Changes in the Build Procedure

Typically, you do not edit a target's build script in the *makefile* very often. But you may often change the build script by specifying overrides for *make macros*, either on the command line or in the UNIX environment. For example, target **hello.o** is specified as follows in the makefile:

```
hello.o: hello.c hello.h
    rm hello.o
    cc -c $(CFLAGS) hello.c
```

When it executes this build script, **clearmake** enters the build script, after macro substitution, into the config record. The command

```
% clearmake hello.o CFLAGS="-g -O1"
```

produces this configuration record entry:

```
-----
Build script:
-----
    cc -c -g -O1 hello.c
```

So does this command:

```
env CFLAGS="-g -O1" clearmake -e hello
```

The **clearmake** build-avoidance algorithm compares effective build scripts. If you then use the command **clearmake hello.o** without specifying **CFLAGS="-g -O1"**, **clearmake** rejects the existing derived object, which was built with those flags. The same mismatch occurs if you create a **CFLAGS** environment variable with a different value, and then invoke **clearmake** with the **-e** option.

Using a Build Options Specification (BOS) File

To manage temporary overrides for make macros and EVs, place macro definitions in *build options specification* (BOS) files. **clearmake** provides several ways for using a BOS file. For example, if your *makefile* is named **project.mk**, macro definitions are read from **project.mk.options**. You can also keep a BOS file in your home directory, or specify one or more BOS files with **clearmake -A**. For details, see Chapter 5, *clearmake Makefiles and BOS Files*.

Using a BOS file to specify make macro overrides relieves you of having to remember which options you specified for the last build. If you have not modified the BOS file recently, derived objects in your view are not disqualified for reuse on the basis of build script discrepancies. Some of the sections that follow describe other applications of BOS files.

Handling Targets Built in Multiple Ways

Because **clearmake** compares build scripts, undesirable results may occur if your build environment includes more than one way to build a particular target. For example, suppose that the target **test_prog_3** appears in two makefiles in two directories. The first is in its source directory, **util_src**:

```
test_prog_3: ...
    cc -o test_prog_3 ...
```

The second is in another directory, **app_src**:

```
../util_src/test_prog_3: ...
    cd ../util_src ; cc -o test_prog_3
```

Derived objects built with these scripts may be equivalent, because they are built as the same file name (**test_prog_3**) in the same VOB directory (**util_src**). But by default, a build in the **app_src** directory never reuses or winks in a DO built in the **util_src** directory, because build-script comparison fails.

You can suppress build-script comparison for this target by using a **clearmake** special build target, **.NO_CMP_SCRIPT** in the makefile or in an associated BOS file:

```
.NO_CMP_SCRIPT: ../util_src/test_prog_3
```

To suspend build-script comparison once, you can use **clearmake -O**.

Using a Recursive Invocation of **clearmake**

You can eliminate the problem of different build scripts described in *Handling Targets Built in Multiple Ways* by adding a recursive invocation of **clearmake** to the makefile in **app_src**:

```
../util_src/test_prog_3: ...
    cd ../util_src ; $(MAKE) test_prog_3           ($(MAKE) invokes clearmake recursively)
```

Now, target **test_prog_3** is built the same way in both directories. You can turn on build-script comparison again, by removing the **.NO_CMP_SCRIPT** special target.

Optimizing Winkin by Avoiding Pseudotargets

Like other **make** variants, **clearmake** always executes the build script for a *pseudotarget*, a target that does not name a file system object built by the script. For example, in the section *Using a Recursive Invocation of clearmake*, you may be tempted to use a pseudotarget in the **app_src** directory's *makefile*:

```
test_prog_3: ...                                (shortened from ../util_src/test_prog_3)
    cd ../util_src ; $(MAKE) test_prog_3
```

A build of any higher-level target that has **test_prog_3** as a build dependency always builds a new **test_prog_3**, which in turn triggers a rebuild of the higher-level target. If the rebuild of **test_prog_3** was not necessary, the rebuild of the higher-level target may not have been necessary, either. Such unnecessary rebuilds decrease the extent to which you can take advantage of ClearCase's derived object sharing capability.

Accommodating the Build Tool's Different Name

The fact that the ClearCase build utility has a unique name, **clearmake**, may conflict with existing build procedures that implement recursive builds. Most make variants automatically define the make macro **\$(MAKE)** as the name of the build program, as it was typed on the command line:

```
% make hello.o                                (sets MAKE to "make")
% clearmake hello.o                          (sets MAKE to "clearmake")
% my_make hello.o                            (sets MAKE to "my_make")
```

This definition enables recursive builds to use **\$(MAKE)** to invoke the same build program at each level. The section *Optimizing Winkin by Avoiding Pseudotargets* includes one such example; here is another one:

```
SUBDIRS = lib util src

all:
    for DIR in $(SUBDIRS) ; do ( cd $$DIR ; $(MAKE) all ) ; done
```

Executing this build script with **clearmake** invokes **clearmake all** recursively in each subdirectory.

3.3 Declaring Source Dependencies in Makefiles

To implement build avoidance based on time stamps, standard **make** variants require you to declare all the source file *dependencies* of each build target. For example, object module **hello.o** depends on source files **hello.c** and **hello.h** in the same directory:

```
hello.o: hello.c hello.h
    rm -f hello.o
    cc -ccl /c hello.c
```

Typically, these source files depend on project-specific header files through `#include` directives, perhaps nested within one another. The standard UNIX files do not change very often, but programmers often lament that “it didn’t compile because someone changed the project’s header files without telling me.”

To alleviate this problem, some organizations include every header file dependency in their makefiles. They rely on utility programs (for example, **makedepend**) to read the source files and determine the dependencies.

clearmake does not require that source-file dependencies be declared in *makefiles* (but see *Source Dependencies Declared Explicitly on page 36*). The first time a derived object is built, its build script is always executed; thus, the dependency declarations are irrelevant for determining whether the target is out of date. After a derived object has been built, its configuration record provides a complete list of source-file dependencies used in the previous build, including those on all header files (nested and nonnested) read during the build.

You can leave source-file dependency declarations in your existing makefiles, but you need not update them as you revise the makefiles. And you need not place source-file dependencies in new makefiles to be used with **clearmake**.

NOTE: Although source-file dependency declarations are not required, you may want to include them in your makefiles, anyway. The principal reason for doing so is portability: you may need to provide your sources to another team (or another company) that is not using ClearCase.

Source Dependencies Declared Explicitly

The ClearCase *build auditing* facility tracks only the MVFS objects used to build a target. Sometimes, however, you may want to track other objects. For example:

- The version of a compiler that is not stored in a VOB
- The version of the operating system kernel, which is not referenced at all during the build
- The state of a *flag-file*, used to force rebuilds

You can force such objects to be recorded in the CR by declaring them as dependencies of the makefile target:

```
hello.o: hello.c hello.h /usr/5bin/cc my.flag
    rm -f hello.o
    cc -c hello.c
```

This example illustrates dependency declarations for these kinds of objects:

- ▶ **(hello.c, hello.h)** Dependencies on MVFS objects are optional. These are recorded by `clearmake` and MVFS anyway.
- ▶ **/usr/5bin/cc** — Dependencies on build tools are required to track the build tools that are not stored in VOBs. These dependencies are listed as non-MVFS objects in the configuration record.
- ▶ **my.flag** — Dependencies on view-private objects can implement a *flag-file* capability.

We suggest that you use view-private files as *flag files*, rather than using non-MVFS files (such as **/tmp/flag**). In a parallel build, a view-private flag file is guaranteed to be the same object on all hosts; there is no such guarantee for a non-MVFS file.

As an alternative to declaring your C compiler as a build *dependency*, you can place it (and other tools) in a tools VOB. The versions of such tools are recorded, eliminating the need for explicit dependency declarations. Additional issues in the auditing of build tools are discussed in the section *Explicit Dependencies on Searched-For Sources*.

Explicit Dependencies on Searched-For Sources

There are situations in which the configuration lookup algorithm that **clearmake** uses qualifies a derived object, even though rebuilding the target would produce a different result. Configuration lookup requires that for each object listed in an existing CR, the current view must select the same version of that object. However, in cases where you use search paths to find an object, a target rebuild may use a different object than the one listed in the CR. Configuration lookup does not take this possibility into account.

When files are accessed by explicit pathnames, configuration lookup qualifies derived objects correctly. Configuration lookup may qualify a derived object incorrectly if files are accessed at build time by a search through multiple directories, ; for example, when the **-I** option to a C or C++ compiler specifies a header file, or when the **-L** option to a linker specifies a library file. The following build script uses a search to locate a library file, **libprojutil.a**:

```
hello:
    cc -o hello -L /usr/project/lib -L /usr/local/lib \
        main.o util.o -lprojutil
```

The command **clearmake hello** may qualify an existing derived object built with **/usr/local/lib/libprojutil.a**, even though rebuilding the target would now use **/usr/project/lib/libprojutil.a** instead.

clearmake addresses this problem in the same way as some standard **make** implementations:

- ▶ You must declare the searched-for source object as an explicit dependency in the makefile:

```
hello: libprojutil.a
    ...
```

- ▶ You must use the **VPATH** macro to specify the set of directories to be searched:

```
VPATH = /usr/project/lib:/usr/local/lib
```

Given this makefile, **clearmake** uses the **VPATH** (if any) when it performs configuration lookup on **libprojutil.a**. If a candidate derived object was built with **/usr/local/lib/projutil.a**, but would be built with **/usr/project/lib/projutil.a** in the current view, the candidate is rejected.

NOTE: The **VPATH** macro is not used for all source dependencies listed in the config record. It is used only for explicitly declared dependencies of the target. Also, **clearmake** searches only in the current view.

Build Tool Dependencies. You can use this mechanism to implement dependencies on build tools. For example, you can track the version of the C compiler used in a build as follows:

```
msg.o: msg.c $(CC)
    $(CC) -c msg.c
```

With this makefile, either your **VPATH** must include the directories on your search path (if the **\$(CC)** value is **cc**), or you must use a full pathname as the **\$(CC)** value.

NOTE: If your C compiler is stored in a VOB and you invoke it from the VOB, ClearCase tracks its version and you do not have to include it as a dependency.

3.4 Build-Order Dependencies

In addition to source dependencies, *makefiles* also contain build-order dependencies. For example:

```
hello: hello.o libhello.a
    ...
libhello.a: hello_env.o hello_time.o
    ...
```

These dependencies are buildable objects, and are called *subtargets*. The executable **hello** must be built after its subtargets, object module **hello.o** and library **libhello.a**, and the library must be built after its subtargets, object modules **hello_env.o** and **hello_time.o**.

ClearCase does not detect build-order dependencies; you must include such dependencies in makefiles used with **clearmake**, as you do with other **make** variants.

3.5 Problems with Forced Builds

clearmake has a **-u** option (unconditional), which forces rebuilds. Using this option reduces the efficiency of derived object sharing, however. If you force **clearmake** to build a target in a situation where it would have winked in an existing DO, you create a new DO with the same configuration as an existing one. In such situations, a developer who expects a build to share a particular existing DO may get another, identically configured DO instead. This may confuse the team and waste disk space.

We suggest that you use a flag-file to force a rebuild, rather than using **clearmake -u**. (See *Source Dependencies Declared Explicitly* on page 36.)

3.6 How clearmake Interprets Double-Colon Rules

Double-colon rules are a special kind of makefile construct that allows several independent rules for one target, each with a possibly different build script. The semantics given to these rules by other **make** programs (such as Gnu make, Sun make, and for **clearmake** when CRs are not being generated) are that commands within each double-colon rule are executed if the target is older

than any dependencies of that particular rule. The result can be that none, any, or all of the double-colon rules are executed.

However, when **clearmake** creates CRs and associates them with the results of its builds, this interpretation runs the risk of generating incomplete CRs, which do not contain all the versions and build scripts used to build the targets. For this reason, **clearmake** interprets these rules in a more conservative way.

When building a target specified by a number of double-colon rules, **clearmake** concatenates all build scripts from all the double-colon rules for that target, and runs them in a single audited script.

To produce the correct results, any subtargets must already have been built, so **clearmake** builds any out-of-date subtargets before it executes the concatenated build script.

As a result, you may observe these differences in behavior between **clearmake** and other **make** programs concerning double-colon rules:

- **clearmake** runs more of the build scripts than other **make** programs.
- **clearmake** may run the build scripts in a different order than other **make** programs.

However, given the intended use and standard interpretation of double-colon rules, these differences still produce correct builds and complete correct CRs.

3.7 Continuing to Work During a Build

As your build progresses, other developers continue to work on their files, and may check in new versions of elements that your build uses. If your build takes an hour to complete, you do not want build scripts executed early in the build to use version 6 of a header file, and scripts executed later to use version 7 or 8.

To prevent such inconsistencies, any version whose selection is based on a **LATEST** config spec rule is locked out if it is checked in after the instant that **clearmake** was invoked. The moment that the **clearmake** build session begins is the *build reference time*.

The same reference time is reported in each configuration record produced during the build session, even if the session lasts hours (or days):

```
% cleartool catcr hello.o
Target hello.o built by drp.dvt
Host "fermi" running OSF1 V1.3 (alpha)
Reference Time 26-Feb-99.16:53:58, this audit started 26-Feb-99.16:54:10 ...
```

NOTE: The *reference time* is the build reference time, when the overall **clearmake** build session began. The *this audit started* time is when the execution of the individual build script began.

When determining whether an object was created before or after the build reference time, **clearmake** adjusts for *clock skew*, the inevitable small differences among the system clocks on different hosts. For more on build sessions, see *Build Sessions, Subsessions, and Hierarchical Builds on page 43*.

CAUTION: A build's coordinated reference time applies to elements only, providing protection from changes made after the build began. You are not protected from changes to view-private objects and non-MVFS objects. For example, if you begin a build and then change a checked-out file used in the build, a failure may result. Therefore, do not work on the same project in a view where a build is in progress.

3.8 Using Config Spec Time Rules

NOTE: If you use a UCM view, your config spec is generated by ClearCase. Do not add time rules to your config spec.

Using the reference time facility described in *Continuing to Work During a Build*, **clearmake** blocks out potentially incompatible source-level changes that take place after your build begins. But sometimes, the incompatible change has already taken place. ClearCase allows you to block out recently created versions.

A typical ClearCase team-development strategy is for each team member to work in a separate view, but to have all the views use the same config spec. In this way, the entire team works on the same branch. As long as a source file remains checked out, its changes are isolated to a single view; when a developer checks in a new version, the entire team sees the new version on the dedicated branch.

This incremental integration strategy is often very effective. But suppose that another user's recently checked-in version causes your builds to start failing. Through an exchange of e-mail, you trace the problem to header file **project_base.h**, checked in at 11:18 A.M. today. You, and other team members, can reconfigure your views to roll back that one element to a safe version:

```
element project_base.h .../onyx_port/LATEST -time 5-Mar.11:00
```

If many interdependent files have been revised, you can roll back the view for all checked-in elements:

```
element * .../onyx_port/LATEST -time 5-Mar.11:00
```

For a complete description of time rules, see the **config_spec** reference page.

Inappropriate Use of Time Rules

Your view interprets time rules with respect to the create version event record written by the **checkin** command. The checkin is read from the system clock on the VOB server host. If that clock is out of sync with the clock on the view server host, your attempt to roll back the clock may fail. Thus, don't strive for extreme precision with time rules: select a time that is well before the actual cutoff time (for example, a full hour before, or in the middle of the night).

Do not use time rules to freeze a view to the current time immediately before you start a build. Allow **clearmake**'s reference time facility to perform this service. Here's an inappropriate use scenario:

1. You check in version 12 of **util.c** at 7:05 P.M. on your host. You do not know that clock skew on the VOB host causes the time 7:23 P.M. to be entered in the create version event record.
2. To freeze your view, you change your config spec to include this rule:

```
element * /main/LATEST -time 19:05
```

3. You issue a **clearmake** command immediately (at 7:06 P.M.) to build a program that uses **util.c**. When selecting a version of this element to use in the build, your view consults the event history of **util.c** and rejects version 12, because the 7:23 P.M. time stamp is too late for the **-time** configuration rule.

3.9 Build Sessions, Subsessions, and Hierarchical Builds

The following terms are used to describe the details of ClearCase build auditing:

- Invoking **clearmake** or **clearaudit** starts a *build session*. The time at which the build session begins becomes the *build reference time* for the entire build session, as described on *Continuing to Work During a Build* on page 40.
- During a build session, one or more *target rebuilds* typically take place.
- Each target rebuild involves the execution of one or more *build scripts*. (A double-colon target can have multiple build scripts; see *How clearmake Interprets Double-Colon Rules* on page 39.)
- During each target rebuild, **clearmake** or **clearaudit** conducts a *build audit*.

Subsessions

A build session can have any number of *subsessions*, all of which inherit the reference time of the build session. A subsession corresponds to a nested build or recursive make, which is started when a **clearmake** or **clearaudit** process is invoked in the process family of a higher-level **clearmake** or **clearaudit**. For example:

- Including a **clearmake** or **clearaudit** command in a *makefile* build script executed by **clearmake** or **clearaudit**
- Entering a **clearmake** or **clearaudit** command in an interactive process started by **clearaudit**

A subsession begins while a higher-level session is still conducting build audits. The subsession conducts its own build audits, independent of the audits of the higher-level session; that is, the audits are not nested or related in any way, other than that they share the same build reference time.

Versions Created During a Build Session

Any version created during a build session and selected by a **LATEST** config spec rule is not visible in that build session. For example, a build checks in a derived object it has created; subsequent commands in the same build session do not select the checked-in version, unless it is selected by a config spec rule that does not use the version label **LATEST**.

An effect of this behavior is that you cannot check in and label a version during a single build session. Instead, you must check in the version during one build session, and label the version during another build session. Use the **mklabel -config** command to label versions associated with a specific derived object.

Coordinating Reference Times of Several Builds

Different build sessions have different reference times. The best way to assign a series of builds the same reference time is to structure them as a single, hierarchical build.

An alternative approach is to run all the builds within the same **clearaudit** session. For example, you can write a shell script, **multi_make**, that includes several invocations of **clearmake** or **clearaudit** (along with other commands). Running the script as follows ensures that all the builds are subsessions that share the same reference time:

```
% clearaudit -c multi_make
```

Objects Written at More than One Level

Problems occur when the same file is written at two or more session levels (for example, a top-level build session and a subsession): the *build audit* for the higher-level session does not contain complete information about the file system operations that affected the file. For example:

```
% clearaudit -c "clearmake shuffle > logfile"
```

The file **logfile** may be written twice:

- During the **clearaudit** build session, by the shell program invoked from **clearaudit**
- During the **clearmake** subsession, when the **clearaudit** build session is suspended

In this case, **clearaudit** issues this error message:

```
clearaudit: Error: Derived object modified; cannot be stored in VOB.  
Interference from another process?
```

To work around this limitation, postprocess the derived object at the higher level with a **cp** command:

```
% clearaudit -c "clearmake shuffle > log.tmp; cp log.tmp logfile; rm log.tmp"
```

3.10 Build Auditing and Background Processes

The ClearCase build programs—**clearmake**, **clearaudit**, and **abe**—use the same procedure to produce configuration records:

1. Send a request to the host's *multiversion file system* (MVFS), to initiate build auditing.
2. Invoke one or more child processes (typically, shell processes), in which *makefile* build scripts or other commands are executed.
3. Turn off MVFS build auditing.
4. If all the subprocesses have indicated success, and at least one MVFS file has been created, compute and store one or more configuration records.

Any subprocesses of the child processes invoked in Step #2 inherit the same MVFS build audit. (Recursive invocations of ClearCase build programs conduct their own, independent audits; see *Build Sessions, Subsessions, and Hierarchical Builds on page 43*.)

A problem can occur if a build script (or other audited command) invokes a background subprocess, and exits without waiting for it to complete. The build program has no knowledge of the background process and may proceed to Step #3 and Step #4 before the background process has finished its work. In such situations, ClearCase cannot guarantee what portion, if any, of the actions of background commands will be included in the resulting CR. The contents of the CR depend on system scheduling and timing behavior.

The ClearCase build programs audit background processes correctly only if both of the following conditions are true:

- The build script does not complete until all background processes are known to have finished.
- Each background process performs its first MVFS file access while it is still a descendant process of the **clearmake** or **clearaudit** process. (The ClearCase kernel component determines whether to audit a given process when that process first accesses the MVFS. If the process's ancestors include a process already being audited, the descendant process is similarly marked for auditing.)

If either or both of these conditions are not true, avoid using background processes in audited build scripts.

3.11 Working with Incremental Update Tools

The design of ClearCase's *build auditing* capability makes it ideal for use with tools that build derived objects from scratch. Because newly created objects have no history, ClearCase can learn everything it needs to know at build time. But this reliance on build-time file-system-level auditing can cause ClearCase to record incomplete information for objects that are updated incrementally, which do have a history.

In ClearCase, incremental updating means that an object is updated partially during the builds of multiple *makefile* targets, instead of generated completely by the build of one target. By default, **clearmake** does not update an existing CR incrementally when it builds a target. Instead, it does the following:

- Each time a build script incrementally updates an object's file system data, **clearmake** writes a completely new CR, which describes only the most recent update, not the entire build history.
- The new CR does not match the desired build configuration for any of the other targets that update the object incrementally.

This results in a situation that is both unstable and incorrect: all incremental-update targets are rebuilt each time that **clearmake** is invoked; when the build is finished, the DO has the correct file system data, but its CR may not describe the DO's configuration accurately.

clearmake provides a special makefile target **.INCREMENTAL_TARGET**, which can be used to guarantee correct CR information for incremental updates. The following sections give examples of how to use **.INCREMENTAL_TARGET**.

Example: Building an Archive

A common incremental-update scenario is the building of an archive by **ar(1)**. A traditional **make** program treats an archive as a compound object; it can examine the time stamps of the individual components (object modules) in the archive; and it can update the archive by replacing one or more individual object modules. Here is a simple *makefile* in which a special syntax enables multiple targets to update a single archive, **libvg.a**, incrementally:

```
libvg.a:: libvg.a(base.o)
libvg.a:: libvg.a(in.o)
libvg.a:: libvg.a(out.o)
```

If you edit one of the library's sources (for example, **out.c**), a traditional **make** program uses the special syntax and a **.c.a** built-in rule to update the library as follows:

1. It looks inside the archive **libvg.a**, and determines that it includes an **out.o** that is older than its source file.
2. It compiles a new **out.o** from **out.c**.
3. It uses **ar** to incrementally update **libvg.a**, replacing the old instance of object module **out.o** with the newly built instance.

clearmake does not implement this algorithm, and includes no support for treating an archive as a compound object. ClearCase build-avoidance is based solely on metadata (CRs), not on any analysis of the file-system data. **clearmake** interprets the above *makefile* as follows:

1. It considers all the **libvg.a(...)** dependencies to be multiple instances of the same double-colon build target.
2. Accordingly, whenever one of those double-colon targets requires rebuilding, **clearmake** rebuilds them all, using the standard **.c.a** built-in rule. The effect is to rebuild the entire archive **libvg.a** from scratch.

Thus, **clearmake** accepts the standard incremental-update syntax, but interprets it in a way that produces a nonincremental build procedure.

Makefile Restructuring for Incremental Archive Targets

The makefile in the previous example can be restructured to allow incremental updates to an archive:

```
.INCREMENTAL_TARGET: libvg.a
libvg.a: base.o in.o out.o
    ar rv libvg.a $?

base.o:
    cc -c base.c

in.o:
    cc -c in.c

out.o:
    cc -c out.c
```

Object modules built by this *makefile* are standard, shareable derived objects; typically, as library sources stabilize, most builds of target **libvg.a** reuse or wink in most of the object modules.

The **.INCREMENTAL_TARGET** directive tells **clearmake** to merge CRs incrementally for this target; that is, to merge the dependencies listed in the previous CR with those of the latest build. In this way, no information is lost from the CRs. Note that **.INCREMENTAL_TARGET** accepts patterns on its list of targets, as well as file names, so you can direct **clearmake** to merge all archive targets incrementally by including the directive **.INCREMENTAL_TARGET: %.a** in the makefile.

During a rebuild, the **\$?** macro expands to the list of dependencies that do not match the current configuration record for the target. **\$?** is useful in conjunction with **ar r** to replace, in the archive library, only those objects that have changed.

Avoid the following alternate restructuring; it causes a complete rebuild of the archive each time any object module is updated:

```
base.o: base.c
    cc -c base.c
    ar rv libvg.a base.o
.
. and so on
```

NOTE: When you use **.INCREMENTAL_TARGET** with an archive library, the full set of declared dependencies must be the same in all makefiles that update that library. Do not attempt to build up libraries incrementally from two different makefiles. For example:

```
../lib.a: base.o (makefile 1)
ar rv ../lib.a base.o
```

```
../lib.a: in.o out.o (makefile 2)
ar rv ../lib.a in.o out.o
```

The rules are executed in multiple steps, and **clearmake** doesn't combine them to verify that the target is up to date with respect to all of its dependencies. Using the construction given above can result in missing required rebuilds (with either **make** or **clearmake**).

A Note on the Use of ar Keys

Do not use the **u** key with **ar**; it is not reliable within a ClearCase environment.

As mentioned above, the **r** key may be used to direct **ar** to replace one or more object modules in an archive library without replacing the entire library. The **u** key directs **ar** to replace only those object modules with modification dates more recent than the archive library.

This creates problems within ClearCase. When determining whether a **.o** file needs to be rearchived, **ar** only looks at whether its time stamp is newer than that of the **.a** file. This check is not sufficient to determine whether a file inside a ClearCase VOB is out of date. For example, imagine a scenario in which a build winks in a **.o** file whose time stamp is older than the time stamp of the **.a** file. Because the file is different from the one used the last time you built the archive, you want the file to be rearchived. However, **ar** sees that the time stamp is older and does not rearchive the **.o** file.

Example: Incremental Linking

If your makefile is structured properly, configuration records are not likely to lose information during incremental links.

Incremental linkers typically work by determining which object files have changed since the last link, and relinking those objects only. Because the linker may read only some of the objects each time it links, a CR can, in theory, lose information as repeated links are made. However, in practice, because all dependencies of the link are listed in the build script, the build script does not change from one link invocation to the next. And, because you typically list the objects or predefined dependencies of the link, those dependencies are included in the CR.

Additional Incremental-Update Situations

You may encounter incremental updating in other situations, as well. For example, C++ compilers that support parameterized types (templates) often update type map files incrementally as different targets are built. ClearCase includes special *makefile* rules that store per-target type map files. For more information, see Chapter 7, *Using ClearCase to Build C++ Programs*.

Ada compilers often update certain common files in Ada libraries incrementally, as different compilation units are built. There are no current **clearmake** workarounds to implement per-target CRs for Ada libraries. To produce a CR for an Ada library, you can rebuild the library from its sources in a single **clearaudit** session.

3.12 Adding a Version String or Time Stamp to an Executable

This section describes simple techniques for incorporating a version string and/or time stamp into a C-language compiled executable. Including a version string or time stamp allows anyone (for example, a customer) to determine the exact version of a program by entering a shell command.

The techniques described below support use of the **what** command or the **-Ver** option. For example:

```
% what monet
monet R2.0 Baselevel 1
Thu Feb 11 17:33:23 EST 1996
```

```
% monet -Ver
monet R2.0 Baselevel 1 (Thu Feb 11 17:33:23 EST 1996)
```

After the particular version of the program is determined, you can use ClearCase commands to find a local copy, examine its config record, and if appropriate, reconstruct the source configuration with which it was built. (Presumably, the local copy is a derived object that has been checked in as a version of an element.)

You can identify the appropriate derived object by attaching a ClearCase attribute with the version string to the checked-in executable, or you can rely on the time stamp and your ability to run the **what** command on the checked-in executable to find it.

Creating a what String

The **what** program searches for a null-terminated string that starts with a special four-character sequence:

```
@(#)
```

To include this string in a C-language executable, define a global character-string variable. For example, these source statements produce the two-line **what** listing above:

```
char *version_string = "@(#)monet R2.0 Baselevel 1";
char *version_time   = "@(#)Thu Feb 11 17:33:23 EST 1996;
```

As an alternative, you can generate the time stamp dynamically when the **monet** program is linked, using this procedure:

1. Create a new source file that contains the statements that define the what strings. Instead of hard-coding a date string, use a **cpp(1)** macro in the source file. This example uses a source file named **version_info.h**, which contains a macro named **DATE**:

```
char *version_string = "@(#)monet R2.0 Baselevel 1";
char *version_time   = DATE;
```

2. Use shell command substitution to incorporate the current time dynamically into the value for the **DATE** macro:

```
SHELL = /bin/sh
OTHER_OBJS = main.o cmd_line.o (and so on)

monet: version_info.h $(OTHER_OBJS)
    cc -o monet -DDATE="\@(#) `date` \" " $(OTHER_OBJS)
```

A rebuild of **monet** is also triggered if the **version_string** variable is edited manually in **version_info.h**.

The *version_string* can be generated dynamically, too (for example, with environment variables). But it is more likely that the project manager edits this string's value before major builds.

NOTE: If you use **clearmake** to build **monet**, you need not declare **version_info.h** as an explicit dependency.

Implementing a `-Ver` Option

You need not depend on the `what` command to extract version information from your executable. Instead, you can have the program itself output the information stored in the `version_string` and `version_time` variables. Revise the source module that does command-line processing to support a `what` version option (for example, `-Ver`):

```
#include <stdio.h>

main(argc,argv)
    int argc;
    char **argv;
{
    /*
     * implement -Ver option
     */
    if (argc > 1 && strcmp(argv[1],"-Ver") == 0) {
        extern char *version_string;
        extern char *version_time;
        /*
         * Print version info, skipping the "@(#)" characters
         */
        printf ("%s (%s)\n",
                &version_string[4], &version_time[4]);
        exit(0);
    }
}
```


Working with Derived Objects and Configuration Records

4

This chapter describes the operations you can perform on derived objects and configuration records. For more information and examples, see the reference pages for the commands.

The information in this chapter applies only to dynamic views.

4.1 Setting Correct Permissions for Derived Objects

If you and other members of your team want to share derived objects (DOs), make sure that your views are configured to create shareable DOs, and that the DOs are created with a mode that grants both read and write access to team members. To accomplish this, use either of the following alternatives:

- Set your **umask** value to 2 in your shell startup file.
- Set the environment variable **CCASE_BLD_UMASK** to 2 and leave a more restrictive **umask** value in your **.login** file or its equivalent (the value 22 is commonly used, which denies write access to group members). When **clearmake** runs a build script, it performs the following steps:
 - a. Saves the current **umask** value.
 - b. Sets the **umask** value to the value of **CCASE_BLD_UMASK**.
 - c. Creates a shell (if necessary) to run the build script.

- d. Restores the original **umask** value when the build script (or shell) completes.

You can set `CCASE_BLD_UMASK` as a make macro, instead of as an environment variable.

NOTE: If you want to use `CCASE_BLD_UMASK`, do not set your **umask** value in your shell startup file. If you set the **umask** value in your startup file, the **umask** value is reset to its original value in Step #c when the startup file is read. Setting `CCASE_BLD_UMASK` in your startup file has no effect.

Other users cannot overwrite and destroy a DO that you are still using, even if you use a `CCASE_BLD_UMASK` value that grants write access to group members. If your DO has been winked in to another view, and the corresponding *makefile* target is rebuilt in that view, **clearmake** first breaks the link to your DO, and then creates a file in that view for the build script to overwrite.

Permissions on DOs affect the extent to which they are shareable:

- When you perform a build, the ClearCase build tool winks in a derived object to your view only if you have read permission on the DO.
- The ClearCase build tool can wink in DOs for which you do not have write permission. But `permission denied` errors may occur during a subsequent build, when a compiler (or other build script command) attempts to overwrite such a DO. To work around this problem, you can rewrite your makefile to remove the target before rebuilding it. You can also set a policy for how users must set their permissions.

For information on fixing the permissions of DO versions, see the **protect** reference page.

4.2 Listing and Describing Derived Objects

The following sections describe how to use the **lsdo**, **describe**, **ls**, and **lsprivate** commands to list derived objects.

Listing Derived Objects Created at a Certain Pathname

Use the **lsdo** command to list derived objects created at a specific pathname. For information on the kinds of DOs included in the listing, see the **lsdo** reference page.

- To list all DOs created at the pathname **adm.h**:

cleartool lsdo adm.h

```
01-Jul.13:49 "adm.h@@01-Jul.13:49.1286781"  
30-Jun.20:03 "adm.h@@30-Jun.20:03.1278990"  
30-Jun.18:14 "adm.h@@30-Jun.18:14.1277470"  
29-Jun.19:11 "adm.h@@29-Jun.19:11.1253509"  
29-Jun.18:13 "adm.h@@29-Jun.18:13.1252790"  
29-Jun.16:09 "adm.h@@29-Jun.16:09.1249897"
```

- To list all DOs created by you at the pathname **adm.h**:

cleartool lsdo -me adm.h

```
30-Jun.18:14 "adm.h@@30-Jun.18:14.1277470"
```

Listing a Derived Object's Kind

To display a derived object's kind, use the **cleartool** commands **ls -l**, **lsprivate -l -do**, or **describe -fmt "[%DO_kind]p"**. The kind can be one of the following values:

non-shareable	The DO was created during an express build and cannot be winked in by other views.
unshared	The DO was created during a regular build. Its data container is located in view storage, not in the VOB.
promoted	The DO's data container has been promoted to the VOB by a winkin or view_scrubber -p command. The DO is referenced by only one view.
shared	The DO's data container has been promoted to the VOB by a ClearCase build tool or a manual winkin or view_scrubber -p command.

- List, in long form, a particular DO.

cleartool ls -l util

```
derived object (non-shareable) util@@01-Sep.10:54.2147483681
```

- To list all DOs created in the current view in the `/vobs/dev` VOB, including the DO kind:

```
cleartool lsprivate -long -invob /vobs/dev -do
```

```
derived object (unshared) /vobs/dev/file2.txt@02-Jul.13:51.124
derived object (unshared) /vobs/dev/file2sub.txt@02-Jul.13:51.123
```

- To list the name and kind of all DOs created in the current view:

```
cleartool describe -fmt "%n\t%[DO_kind]p\n" `cleartool lsprivate -do`
```

```
/vobs/dev/file2.txt@02-Jul.13:51.124      shared
/vobs/dev/file2sub.txt@02-Jul.13:51.123   shared
/vobs/dev/dir1/foo.o@01-Jul.14:23.186     unshared
/vobs/dev/api/bin/adm.exe@04-Jul.04:01.776 unshared
...
```

Displaying a DO's OID

A derived object's OID is the permanent identifier recorded in the VOB database for the DO. It does not change over the life of the DO, unlike the DO-ID (see *Reuse of DO-IDs* on page 25). To display the OID, use the command **describe -fmt "%On"**. For example:

```
cleartool describe -fmt "%On\n" foo.o
```

```
b7afc83e.2f2311d3.a382.00:01:80:7b:09:69
```

Displaying a Description of a DO Version

The **describe** command displays descriptions of DO versions, as it does descriptions of regular versions. You can use the **-fmt** option to extract parts of the description. For example, the following command prints the name, predecessor version, and element type of a DO version:

```
cleartool describe -fmt "%n\t%[version_predecessor]p\t%[type]p\n" file1.o
```

```
file1.o@@/main/2      /main/1  text_file
```

For more information on the **-fmt** option, see the **fmt_ccase** reference page.

4.3 Identifying the Views that Reference a Derived Object

The VOB stores information about which views reference a derived object. To display this information, use the **lsdo** command:

cleartool lsdo -l hello.o

```
10-Mar-99.15:25:52 Allison K. Pak (akp.user@copper)
  create derived object "hello.o@@10-Mar.15:25.213"
  size of derived object is: 450
  last access: 15-Mar-99.14:22:17
  references: 2 (shared)
=> copper:/home/akp/tut/old.vws
=> copper:/home/akp/tut/fix.vws
```

4.4 Specifying a Derived Object in Commands

In general, you use standard pathnames to access DOs when you're working in a view that references them. To standard software (for example, linkers and debuggers), the standard pathname of a derived object (**util.o**) references the DO.

This is another example of ClearCase's transparency feature: a standard pathname accesses one of many different variants of a file system object. Note this distinction, however:

- ▶ A version of an element appears in a dynamic view because it is selected by a config spec rule.
- ▶ A particular derived object appears in a dynamic view as the result of a build or a winkin.

To access a DO in another dynamic view, use a view-extended pathname:

```
/view/drp/vobs/proj/src/msg.o           (the DO in view drp)
/view/R2_integ/vobs/proj/src/msg.o       (the DO in view R2_integ)
```

NOTE: You cannot use view-extended pathnames in makefiles.

To specify a certain DO in a ClearCase command, use the DO-ID. For example, you can use a DO-ID in the **catcr** command to view the contents of a specific DO's config record:

```
cleartool catcr foo.o@@29-Jun.14:40.88
```

NOTE: You cannot use a DO-ID in standard commands.

Because DO-IDs can change, avoid using them in files or scripts that operate on a DO. Instead, use a standard pathname or the derived object's object identifier (OID), which never changes. To determine a DO's object identifier, use **cleartool describe -fmt "%On\n"**. For example:

```
cleartool describe -fmt "%On\n" foo.o@@29-Jun.14:40.88  
2c5Fc68a.2e5311d3.a382.00:01:80:7b:09:69
```

Also, a derived object gets a permanent identifier when it is checked in as a version of an element. See *Working with DO Versions* on page 62.

4.5 Winking in a DO Manually

You can manually wink in any DO to your view, using the **winkin** command. For example:

```
cleartool lsdo hello  
08-Mar.12:48 akp "hello@@08-Mar.12:48.265"  
07-Mar.11:40 george "hello@@07-Mar.11:40.217"
```

```
cleartool winkin hello@@07-Mar.11:40.217  
Winked in derived object "hello"
```

```
% hello  
Greetings, sakai-san!  
Your home directory is /net/neptune/home/sakai.  
It is now Tue Mar 8 12:58:30 1994.
```

You can wink in a DO that doesn't match your build configuration for any of the following reasons: to run it, to perform a byte-by-byte comparison with another DO, or to perform any other operation that requires access to the DO's file system data.

The **winkin** command can also wink in the set of DOs in a hierarchy of CRs. You can use this recursive **winkin** to seed a new view with a set of derived objects. For example:

```
cleartool winkin -recurse foo@@20-Jul.14:32.146  
Winked in derived object "/vobs/smg_test/file2.txt"  
Winked in derived object "/vobs/smg_test/file2sub.txt"  
Promoting unshared derived object "/vobs/smg_test/foo".  
Winked in derived object "/vobs/smg_test/foo"
```

You can use the **winkin** command to convert a nonshareable DO to a shared DO. For more information, see *Converting Nonshareable DOs to Shared DOs* on page 66.

4.6 Preventing Winkin

The following sections describe how to prevent winkin to or from your view while using the ClearCase build tools. (You can prevent winkins altogether by building in a snapshot view or by not using the ClearCase build tools.)

Preventing Winkin to Your View

To direct **clearmake** to limit reuse to DOs created in the current view, use **clearmake -V**. For more information, see the **clearmake** reference page.

Preventing Winkin to Other Views

To prevent any derived objects you create from being winked in to other views, use one of the following techniques:

- Use express builds. See *Using Express Builds to Prevent Winkin to Other Views*.
- Use the **-T** or **-F** options to create view-private files with no config records. **clearmake** does not perform configuration lookup, but this does not matter if you are not changing other files.
- Use special targets that prevent winkin. For example, use **.NO_WINK_IN** with **clearmake**. For more information on special targets, see *Special Targets* on page 80.

Using Express Builds to Prevent Winkin to Other Views

During an express build, ClearCase creates DOs that are nonshareable and cannot be used by builds in other views. These nonshareable DOs have configuration records, but the ClearCase build tools do not write information into the VOB for these DOs. Therefore, the DOs are invisible to builds in other views.

NOTE: During an express build, the ClearCase build tools wink in DOs from other views. For information on avoiding winkins from other views, see *Preventing Winkin to Your View* on page 59.

Use express builds when DOs created by the build are not appropriate for use by other views. As a general rule:

- Use express builds for development builds that use relatively unstable, checked-out versions.
- Use regular builds for release or nightly builds that use stable, checked-in versions. DOs created by these builds are more likely to be winked in by other views.

Enabling Express Builds

When you invoke a ClearCase build tool, the kind of build that occurs depends on how your view is configured. To use express builds, configure an existing dynamic view with the nonshareable DOs property, or create a new dynamic view with the nonshareable DOs property. Then, run your ClearCase build tool (**clearmake** or **clearaudit**) in the view.

The following sections describe how to configure your view to use express builds.

Configuring an Existing View for Express Builds

Use the command **chview -nshareable_dos view-tag**. For more information, see the **chview** reference page.

Future builds in the view will create nonshareable derived objects. However, existing DOs in the view are shareable; they are not converted to nonshareable. These existing DOs can still be winked in by other views.

Creating a New View that Uses Express Builds

To create a new view, enter the **mkview** command and specify the **-nshareable_dos** option. For more information, see the **mkview** reference page.

NOTE: The ClearCase administrator can set a site-wide default for the DO property. If a site-wide default exists and you do not specify the DO property when you create the view, ClearCase uses the site-wide default. For more information, see the **setsite** reference page.

Preventing Winkin to or from Other Architectures

By default, **clearmake** winks in derived objects built on different architectures. For example, a build on a Sun machine may wink in a DO built on a Windows NT machine. If you want to prevent this behavior, use one of the following techniques:

- Differentiate the build script for different architectures:
 - > Add the architecture name to your build script so that **clearmake** differentiates among the build scripts. For example, include `echo $CPU`.
 - > Store the architecture name in a macro and pass it to **clearmake** on the command line.
 - > Use architecture-specific subdirectories to store DOs.
- Make your tools a build dependency by storing them in a VOB. Also, store your system header files in a VOB.

4.7 Converting Derived Objects to View-Private Files

Using a standard command or program to modify a derived object in any way converts it from a DO to a view-private file. For example, use **ls -long** to list a derived object:

```
% cleartool ls -long msg.o
derived object (shared)      msg.o@@@10-Mar.15:33.333
```

Modify the DO with a standard command:

```
% touch msg.o
```

The **ls -long** command now lists the file as a view-private object:

```
% cleartool ls -long msg.o
view private object         msg.o
```

4.8 Working with DO Versions

The following sections describe how to create and manipulate DO versions.

Creating DO Versions

You can convert DOs to elements, or check them in as versions of existing elements. The element-creation and version-creation processes are the same for shareable and nonshareable DOs, with this exception: when you check in a nonshareable DO, it is converted to a shared DO before being checked in.

For more information, see the **mkelem** and **checkin** reference pages.

Checking In DOs During a Build

You can write a build script that creates a derived object, and then checks it in or converts it to an element. However, the ClearCase build tool does not create a config record until the build script has completed (all commands after the *target-name:* have executed). Therefore, if the same build script that created the DO checks it in or converts it to an element, the resulting version is not a DO version.

For example, the version created by the following build script is not a DO version:

```
buildit : buildit.c
    cleartool co -unres -nc $@
    rm -fr $@
    cc -o $@ $*.c
    cleartool ci -nc buildit
```

You can work around this problem by building and checking in a derived object in two steps. For example, the makefile contains one build script that creates the DO, and another build script that checks it in, as shown here:

```

buildit : buildit.c
         cleartool co -unres -nc $@
         rm -fr $@
         cc -o $@ $*.c

stageit : buildit
         cleartool ci -nc buildit

```

The command **clearmake stageit** performs the following steps:

1. Brings the target **buildit** up to date. This creates a DO named **buildit** and an associated config record.
2. Brings the target **stageit** up to date. This step checks in the **buildit** derived object as a DO version.

Accessing DO Versions

When you check out a DO version, it is winked in to your dynamic view. You can use a standard pathname to access the DO's file system data. However, VOB-database access is handled in the following ways:

- A standard pathname to the DO references the version in the VOB database from which the checkout was made:

```
% cleartool checkout -nc hello (wink in derived object hello)
Checked out "hello" from version "/main/3".
```

```
% cleartool mklabel EXPER hello (use standard pathname to access version from which checkout was made)
Created label "EXPER" on "hello" version "/main/3".
```

- To access the checked-out placeholder version, you must use an extended pathname:

```
% cleartool mklabel -replace EXPER hello@@/main/CHECKEDOUT
Moved label "EXPER" on "hello" from version "/main/3" to
"/main/CHECKEDOUT".
```

If you process a checked-out DO version as described in *Converting Derived Objects to View-Private Files* on page 61, ClearCase reverts to its usual handling of checked-out versions. In this case, a standard pathname references the placeholder version in the VOB database.

Displaying Configuration Records for DO Versions

The **catcr** command displays the configuration record for a DO version. When you use **catcr -recurse** to display the CRs for a DO and all its subtargets, it does not display the CRs for DO versions unless you use the **-ci** option.

The following examples show how to display configuration records for DO versions.

- To display the configuration record for a single DO version:

cleartool catcr foo

```
Target foo built by smg.user
Host "radar" running SunOS 5.5.1 (sun4m)
Reference Time 20-Jul-99.14:32:32, this audit started 20-Jul-99.14:32:32
View was radar:/home/smg/views/smg_test.vws
Initial working directory was /vobs/smg_test
-----
MVFS objects:
-----
/vobs/smg_test/file1.txt@@/main/2                <30-Jun-99.15:35:12>
/vobs/smg_test/file2.txt@@30-Jun.15:33.115
/vobs/smg_test/foo@@/main/3
-----
Build Script:
-----
        cat file1.txt > foo
        cat file2.txt >> foo
-----
```

- To display the configuration record for a derived object, including the CRs of all subtargets except DO versions:

cleartool catcr -recurse foo

```
-----
-----
Target foo built by smg.user
...
Target file2.txt built by smg.user
...
Target file2sub.txt built by smg.user
...
-----
```

- To display the configuration record for a derived object, including the CRs of all subtargets:

```
cleartool catcr -recurse -ci foo
-----
-----
Target foo built by smg.user
...
Target file1.txt built by smg.user
...
Target file1sub.txt built by smg.user
...
Target file2.txt built by smg.user
...
Target file2sub.txt built by smg.user
...
```

For more information, see the **catcr** reference page.

Releasing DOs

A project team can use DO versions to make its product (for example, a library) available to other teams. Typically, the team establishes a release area in a separate VOB. For example:

- A library is built by its project team in one location—perhaps **/vobs/monet/lib/libmonet.a**.
- The team periodically releases the library by creating a new version of a publicly accessible element—perhaps **/vobs/publib/libmonet.a**.

You can generalize the idea of maintaining a development release area to maintaining a product release area. For example, a Release Engineering group maintains one or more release tree VOBs. The directory structure of the trees mirrors the hierarchy of files to be created on the release medium. (Because a release tree involves directory elements, it is easy to change its structure from release to release.) A release tree can be used to organize Release 2.4.3 as follows:

- When an executable or other file is ready to be released, a release engineer checks it in as a version of an element in the release tree.
- An appropriate version label (for example, **REL2.4.3**) is attached to that version, either manually by the engineer or automatically with a trigger.

- ▶ When all files to be shipped have been released in this way, a release engineer configures a view to select only versions with that version label. As seen through this view, the release tree contains exactly the set of files to be released.
- ▶ To cut a release tape, the engineer issues a command to copy the appropriately configured release tree.

4.9 Converting Nonshareable DOs to Shared DOs

NOTE: You cannot convert a shared or unshared DO to a nonshareable DO.

To convert a nonshareable DO to a shared DO, use the **winkin** command. **winkin** advertises the DO by making it shareable and writing information into the VOB, and then promotes it (makes it shared). The command also advertises the DO's sub-DOs and siblings, even if you did not specify the **-siblings** option. This process changes the DO-ID for each derived object.

The **view_scrubber -p** command performs the same operation. See the **winkin** and **view_scrubber** reference pages.

Automatic Conversion of Nonshareable DOs to Shareable DOs

Because you can change a view's DO property, and shareable DOs cannot have nonshareable sub-DOs or siblings, situations can occur in which **clearmake** must convert nonshareable DOs into shareable DOs.

For example, you set your view's DO property to nonshareable DOs, and then perform a build, creating nonshareable DOs. You then set your view's DO property to shareable DOs, and perform another build. The build tool determines that it can reuse some of the nonshareable DOs created in the first build to create shareable DOs in the second. It converts the nonshareable DOs to shareable DOs and reuses them.

4.10 Creating Links to Derived Objects

You cannot make a VOB hard link to a derived object. You can make one or more view-private hard links to a derived object, using the UNIX **In** command, with these restrictions:

- The derived object must be visible in the dynamic view where the view-private hard link is to be created; that is, it must appear in a standard UNIX **ls** listing. (You can use the **winkin** command to satisfy this requirement.)
- The pathname of the hard link must be within the same VOB as the original derived object.

All hard links to a derived object, including the name under which it was originally created, appear with the same *DO-ID* in a ClearCase **ls** listing; if there are multiple names for a derived object in the same directory, all are listed. For example:

```
% In hello hw
% cleartool ls
...
hello@@19-May.19:15.232
hw@@19-May.19:15.232
...
```

In a **catcr** or **describe** command, you can reference a derived object using any of its hard links; all the references are equivalent. But an **lsdo** command must reference a derived object by its original name, not by any of its subsequently created hard links. Likewise, a derived object can be winked in only at its original pathname.

SPECIAL CASE: If a hard link is created by the same build script as the derived object itself, the hard link becomes an additional name for the DO. **lsdo** lists the hard link, and **clearmake** can wink in using the hard link's pathname.

Each additional hard link increments a derived object's reference count. An **lsdo -l** listing includes the reference counts and the dynamic views in which the references exist. The (2) in this example shows that view **old.vws** has two references to **hello**:

```
% cleartool lsdo -long hello
08-Dec-92.12:06:19 Chuck Jackson (test user) (jackson.dvt@oxygen)
  create derived object "hello@@08-Dec.12:06.234"
  references: 2 => oxygen:/usr/vobstore/tut/old.vws (2)
```

4.11 Displaying VOB Disk Space Usage for Derived Objects

The **dospace** command reports VOB disk space usage for shared derived objects. For more information, see the **dospace** reference page and *Administering ClearCase*.

4.12 Deleting Derived Objects

The **rmdo** command removes the data container and the VOB database object for a derived object. See the **rmdo** reference page for more information.

Shareable derived objects and their data containers can be deleted independently. Deleting a nonshareable derived object deletes the DO.

Removing Data Containers for Derived Objects

The standard **rm(1)** command causes a shareable derived object to disappear from the dynamic view. The effect on physical data storage is as follows:

- If the DO's data container is in the view's private storage area, **rm** deletes that data container.
- If the DO's data container is in a VOB storage pool, the data container is not affected.

In both cases, the derived object in the VOB database is not deleted. The only change to the derived object is that its reference count is decremented.

When a build overwrites a nonshareable or unshared DO, the MVFS removes the old data container from the dynamic view's private storage area, and creates a new one there. It also creates a new CR. At the operating system level, the effect is that an existing file is overwritten.

Scrubbing Derived Objects and Data Containers

A reference count of zero means that the derived object has been deleted or overwritten in every view that ever used it. This situation calls for *scrubbing*: automatic deletion of DO-related information from the VOB. Scrubbing can remove the derived object from the VOB database, its data container from a VOB storage pool (if the DO had ever been shared), and in some cases its associated CR, as well.

The **scrubber** utility removes derived objects from a VOB database and data containers from VOB storage pools. The **view_scrubber** utility removes data containers from a dynamic view's private storage area. For more on scrubbing, see *Administering ClearCase*.

Degenerate Derived Objects

A derived object is complete if its VOB database object, data container, and configuration record (CR) are accessible. Because these entities exist independently, a derived object can become incomplete, or degenerate, if one entity is missing.

Data Container Deleted

When an unshared DO is removed with **rm** or by a target rebuild, its VOB database object continues to exist in the VOB database (with a zero reference count), but the data container no longer exists. Such DOs are usually ignored by **lsdo**, but can be listed with the **-zero** option. The **scrubber** utility deletes zero-referenced DOs.

The **checkvob** command can find and fix missing container problems.

DO Deleted from VOB Database

When an unshared DO is removed from its VOB database with **rmdo**, the data container continues to be visible:

```
% cleartool rmdo Vhelp.log (in general, avoid 'rmdo!')
Removed derived object "Vhelp.log@@14-Sep.72783".
```

```
% cleartool ls Vhelp.log
Vhelp.log [no config record]
```

CR Unavailable

A newly created CR is stored in the dynamic view where its associated DOs were built. If that view becomes unavailable (for example, it is inadvertently destroyed or its host is temporarily down), the DO continues to exist in the VOB database, but operations that must access the CR fail:

```
cleartool: Error: Unable to find view 'mars:/viewstore/pink.vws'  
from albd: error detected by ClearCase subsystem  
cleartool: Error: See albd_log on host mars  
cleartool: Error: Unable to contact View - error detected by ClearCase  
subsystem
```

4.13 Displaying Contents of Configuration Records

The `catcr` command displays the contents of a configuration record. See the `catcr` reference page for more information.

4.14 Comparing Configuration Records

Because config records provide complete records of how DOs are built, you can use them to determine how two builds differ. For example, you expected the build to reuse or wink in a DO that it rebuilt instead. You can compare the CRs for the two DOs to find out what aspect of the build environment was different.

To compare two existing CRs, use the `diffcr` command. For more information, see the `diffcr` reference page.

4.15 Attaching Labels or Attributes to Versions in a CR

You can attach a label or an attribute to the versions in the CR hierarchy of a derived object.

For example, to attach the `SMG_BUILD_5_99` label to the versions in the CR hierarchy of `file.o`:

```
cleartool mklabel -c "may 99 build" -config file.o SMG_BUILD_5_99
```

```
Created label "SMG_BUILD_5_99" on "/vobs/smg_test/" version  
"/main/CHECKEDOUT".
```

```
Created label "SMG_BUILD_5_99" on "/vobs/smg_test/acc.c" version "/main/2".  
Created label "SMG_BUILD_5_99" on "/vobs/smg_test/file.c" version "/main/1".
```

For more information, see the description of the `-config` option in the `mkattr` and `mklabel` reference pages.

4.16 Configuring a View to Select Versions Used to Build a DO

To select the versions in the CR hierarchy of a derived object, use the `-config` version selector in your view's config spec. For example, the following config spec selects the versions in the CR hierarchy for `hello.o`:

```
element * CHECKEDOUT  
element * -config /vobs/dev/lib/hello.o  
element * /main/v3.8/LATEST
```

For more information, see the `config_spec` reference page.

4.17 Including a Makefile Version in a Configuration Record

To record a makefile version in a CR, use one of the following methods:

- Declare it as an explicit dependency in the makefile. To do this, you can use the `$(MAKEFILE)` variable. You must explicitly list any included makefiles you want recorded.

The drawback to this method is that it causes targets that depend on the makefile to be rebuilt if there is any change to the makefile.

- Make it an implicit dependency by referring to it in a build script, and use the special target **.DEPENDENCY_IGNORED_FOR_REUSE** to ignore it in subsequent rebuild decisions. You must explicitly list any included makefiles you want recorded.

For example:

```
.DEPENDENCY_IGNORED_FOR_REUSE: $(MAKEFILE)
targ: dep1 dep2
    cat $(MAKEFILE) > /dev/null
    touch targ
```

The drawback to this method is that the makefile dependency is ignored for reuse, but it is not ignored for winkin.

This chapter describes *makefiles* processed by the ClearCase build program **clearmake**. This is a discussion of differences and ClearCase extensions, rather than a complete description of makefile syntax. This chapter also describes build option specification files (BOS files), which contain temporary macros and ClearCase special targets.

5.1 Makefile Overview

A makefile contains a sequence of entries, each of which specifies a build target, some dependencies, and the *build scripts* of commands to be executed. A makefile can also contain *make macro* definitions, target-dependent macro definitions, and build directives (*special targets*.)

- **Target/dependencies line.** The first line of an entry is a white-space-separated, nonnull list of *targets*, followed by a colon (:) or a double colon (::), and a (possibly empty) list of *dependencies*. Both targets and dependencies may contain ClearCase pathname patterns. (See the **wildcards_ccase** reference page.)

The list of dependencies may not need to include source objects, such as header files, because **clearmake** detects these dependencies. However, the list must include build-order dependencies, for example, object modules and libraries that must be built before executables. (See *Build-Order Dependencies* on page 39.)

- **Build script.** Text following a semicolon (;) on the same line, and all subsequent lines that begin with a <TAB> character, constitute a *build script*: a set of shell commands to be executed. A shell command can be continued onto the next text line with a \<NL> sequence. Any line beginning with a number sign (#) is a comment.

A build script ends at the first nonempty line that does not begin with a <TAB> or number sign (#); this begins a new target/dependencies line or a make macro definition.

Build scripts must use standard pathnames only. Do not include view-extended or version-extended pathnames in a build script.

Executing a build script updates the target, and is called a *target rebuild*. The shell commands in a build script are executed one at a time, each in its own instances of the subshell.

Note that **clearmake** always completely eliminates a \<NL> sequence, even in its compatibility modes. Some other **make** programs sometimes preserve such a sequence—for example, in a **sed(1)** insert command:

```
target: depdcy

sed -e '/xxx=0/i\

yyy=xxx;' depdcy > target
```

- ▶ **Make macro.** A *make macro* is an assignment of a character-string value to a simple name. By convention, all letters in the name are uppercase (for example, **CFLAGS**).
- ▶ **Target-dependent macro definitions.** A target-dependent macro definition takes the form *target-list := macro_name = string*

You can use macros in makefiles or in BOS files. For more information, see *Target-Dependent Macro Definitions* on page 86.

- ▶ **Special targets.** A line that begins with a dot (.) is a *special target*, which acts as a directive to **clearmake**.

5.2 Build Options Specification Files

A *build options specification* (BOS) file is a text file containing macro definitions and/or ClearCase special targets. We recommend that you place temporary macros (such as **CFLAGS=-g** and others not to be included in a makefile permanently) in a BOS file, rather than specifying them on the **clearmake** command line.

By default, **clearmake** reads BOS files in this order:

1. The default BOS files
 - a. The file **.clearmake.options** in your home directory (as indicated in the password database), which is the place for macros to be used every time you execute **clearmake**.
 - b. One or more local BOS files, each of which corresponds to one of the makefiles specified with a **-f** option, or read by **clearmake**. Each BOS file has a name in the form *makefile-name.options*. For example:

```
makefile.options
Makefile.options
project.mk.options
```
2. BOS files specified in the **CCASE_OPTS_SPECS** environment variable.
3. BOS files specified on the command line with **-A**.

If you specify **-N**, **clearmake** does not read default BOS files.

clearmake displays the names of the BOS files it reads if you specify the **-v** or **-d** option, or if **\$CCASE_VERBOSITY** \geq 1.

For information on the contents of BOS files, see *BOS File Entries* on page 86.

5.3 Format of Makefiles

The following sections describe the special considerations for using makefiles with **clearmake**.

Restrictions

clearmake does not support the use of standard input as a makefile.

Libraries

If a target or dependency name contains parentheses, it is assumed to be an archive (library) created by **ar(1)**. For example:

```
lib.a : lib.a(mod1.o) lib.a(mod2.o)
```

The string within parentheses refers to a member (object module) within the library. Use of function names within parentheses is not supported. Thus, **lib.a(mod1.o)** refers to an archive that contains object module **mod1.o**. The expression **lib.a(mod1.o mod2.o)** is not valid.

Inference rules for archive libraries have this form:

.sfx.a

where **sfx** is the file-name extension (suffix) from which the archive member is to be made.

The way in which **clearmake** handles incremental archive construction differs from other **make** variants. For more on this topic, see *Working with Incremental Update Tools* on page 46.

NOTE: The **u** key for **ar** is not reliable within a ClearCase environment. Do not use it. For more information, see *A Note on the Use of ar Keys* on page 49.

Command Echoing and Error Handling

You can control the echoing of commands and the handling of errors that occur during command execution on a line-by-line basis, or on a global basis.

You can prefix any command with one or two characters, as follows:

- Causes **clearmake** to ignore any errors during execution of the command. By default, an error causes **clearmake** to terminate.
The command-line option **-i** suppresses termination-on-error for all command lines.
- @ Suppresses display of the command line. By default, **clearmake** displays each command line just before executing it.
The command-line option **-s** suppresses display of all command lines. The **-n** option displays commands, but does not execute them.
- @ @– These two prefixes combine the effect of **-** and **@**.

The **-k** option provides for partial recovery from errors. If an error occurs, execution of the current target (that is, the set of commands for the current target) stops, but execution continues on other targets that do not depend on that target.

Built-In Rules

File-name extensions (suffixes) and their associated rules in the makefile override any identical file-name extensions in the built-in rules. **clearmake** reads built-in rules from the file *ccase-home-dir/etc/builtin.mk* when you run in standard compatibility mode. In other compatibility modes, other files are read. See Chapter 6, *Using clearmake Compatibility Modes*.

Include Files

If a line in a makefile starts with the string **include** or **sinclude** followed by white space (at least one <SPACE> or <TAB> character), the rest of the line is assumed to be a file name. (This name can contain macros.) The contents of the file are placed at the current location in the makefile.

For **include**, a fatal error occurs if the file is not readable. For **sinclude**, a nonreadable file is silently ignored.

Macros

The following sections describe the order of precedence of macros in a **clearmake** build, and the different types of macros.

Order of Precedence of Make Macros and Environment Variables

By default, the order of precedence of macros and environment variables is as follows:

1. Target-dependent macro definitions
2. Macros specified on the **clearmake** command line
3. Make macros set in a BOS file

4. Make macro definitions in a makefile
5. Environment variables

For example, target-dependent macro definitions override all other macro definitions, and macros specified on the **clearmake** command line override those set in a BOS file.

If you use the **-e** option to **clearmake**, environment variables override macro definitions in the makefile.

All BOS file macros (except those overridden on the command line) are placed in the build script's environment. If a build script recursively invokes **clearmake**:

- ▶ The higher-level BOS file setting (now transformed into an EV) is overridden by a make macro set in the lower-level makefile. However, if the recursive invocation uses **clearmake**'s **-e** option, the BOS file setting prevails.
- ▶ If another BOS file (associated with another makefile) is read at the lower level, its make macros override those from the higher-level BOS file.

See *Special Environment Variables* in the **clearmake** reference page for a list of build-related environment variables.

Make Macros

A *macro definition* takes this form:

```
macro_name = string
```

Macros can appear in the makefile, on the command line, or in a *build options specification* file. (See *Build Options Specification Files* on page 74.)

Macro definitions require no quotes or delimiters, except for the equal sign (=), which separates the macro name from the value. Leading and trailing white space characters are stripped. Lines can be continued using a `\<NL>` sequence; this sequence and all surrounding white space is effectively converted to a single `<SPACE>` character. *macro_name* cannot include white space, but *string* can; it includes all characters up to an unescaped `<NL>` character.

clearmake performs macro substitution whenever it encounters either of the following in the makefile:

```
$(macro_name)  
$(macro_name:subst1=subst2)
```

It substitutes *string* for the macro invocation. In the latter form, **clearmake** performs an additional substitution within *string*: all occurrences of *subst1* at the end of a word within *string* are replaced by *subst2*. If *subst1* is empty, *subst2* is appended to each word in the value of *macro_name*. If *subst2* is empty, *subst1* is removed from each word in the value of *macro_name*.

For example:

```
% cat Makefile
C_SOURCES = one.c two.c three.c four.c
test:
    echo "OBJECT FILES are: $(C_SOURCES:.c=.o)"
    echo "EXECUTABLES are: $(C_SOURCES:.c=)"

% clearmake test
OBJECT FILES are: one.o two.o three.o four.o
EXECUTABLES are: one two three four
```

Internal Macros

clearmake maintains these macros internally. They are useful in rules for building targets.

\$*	(Defined only for inference rules) The file name part of the inferred dependency, with the file-name extension deleted.
\$\$	The full target name of the current target.
\$<	(Defined only for inference rules) The file name of the implicit dependency.
 \$?	(Defined only when explicit rules from the makefile are evaluated) The list of dependencies that are out of date with respect to the target. When configuration lookup is enabled (default), it expands to the list of all dependencies, unless that behavior is modified with the .INCREMENTAL_TARGET special target. In that case, \$? expands to the list of all dependencies different from the previously recorded versions. When a dependency is an archive library member of the form <code>lib(file.o)</code> , the name of the member, file.o , appears in the list.
 \$%	(Defined only when the target is an archive library member) For a target of the form <code>lib(file.o)</code> , \$\$ evaluates to <code>lib</code> and \$% evaluates to the library member, file.o .
 MAKE	The name of the make processor (that is, clearmake). This macro is useful for recursive invocation of clearmake .

MAKEFILE During makefile parsing, this macro expands to the pathname of the current makefile. After makefile parsing is complete, it expands to the pathname of the last makefile that was parsed. This holds only for top-level makefiles, not for included makefiles or for built-in rules; in these cases, it echoes the name of the including makefile.

Use this macro as an explicit dependency to include the version of the makefile in the CR produced by a target rebuild. For example:

```
supersort: main.o sort.o cmd.o $(MAKEFILE)
cc -o supersort ...
```

For more information, see *Including a Makefile Version in a Configuration Record* on page 71.

VPATH Macro

The **VPATH** macro specifies a search path for targets and dependencies. **clearmake** searches directories in **VPATH** when it fails to find a target or dependency in the current working directory. **clearmake** searches only in the current view. The value of **VPATH** can be one directory pathname, or a colon-separated list of directory pathnames. (In Gnu compatibility mode, you can also use spaces as separators.)

Configuration lookup is **VPATH**-sensitive when qualifying makefile dependencies (explicit dependencies in the makefile). Thus, if a newer version of a dependent file appears in a directory on the search path before the pathname in the CR (the version used in the previous build), **clearmake** rejects the previous build and rebuilds the target with the new file.

The **VPATH** setting may affect the expansion of internal macros, such as **\$<**.

Special Targets

Like other build tools, **clearmake** interprets certain target names as declarations. Some of these special targets accept lists of patterns as their dependents, as noted in the description of the target. Pattern lists may contain the pattern character, **%**. When evaluating whether a name matches a pattern, the tail of the prefix of the name (subtracting directory names as appropriate) must match the part of the pattern before the **%**; the file-name extension of the name must match the part of the pattern after the **%**. For example:

Name	Matches	Does not match
/dir/subdir/x.o	%.o x.o subdir/%.o subdir/x.o	/dir/subdir/otherdir/x.o

The following targets accept lists of patterns:

- .DEPENDENCY_IGNORED_FOR_REUSE
- .INCREMENTAL_REPOSITORY_SIBLING
- .INCREMENTAL_TARGET
- .NO_CMP_NON_MF_DEPS
- .NO_CMP_SCRIPT
- .NO_CONFIG_REC
- .NO_DO_FOR_SIBLING
- .NO_WINK_IN
- .SIBLING_IGNORED_FOR_REUSE

Special Targets for Use in Makefiles

.DEFAULT :

If a file must be built, but there are no explicit commands or relevant built-in rules to build it, the commands associated with this target are used (if it exists).

.IGNORE :

Same effect as the `-i` option.

.PRECIOUS : *tgt ...*

The specified targets are not removed when a quit character (typically, `<CTRL-\>`) or an interrupt character (typically, `<CTRL-C>`) is typed.

.SILENT :

Same effect as the `-s` option.

Special Targets for Use in Makefiles or BOS Files

You can use the following special targets either in the makefile itself or in a *build options specification* file. See *Build Options Specification Files* on page 74.

.DEPENDENCY_IGNORED_FOR_REUSE: *file ...*

The dependencies you specify are ignored when **clearmake** determines whether a target object in a VOB is up to date and can be reused. By default, **clearmake** considers that a target cannot be reused if its dependencies have been modified or deleted since it was built. This target applies only to reuse, not to winkin. Also, this target applies only to *detected* dependencies, which are not declared explicitly in the makefile.

You can specify the list of files with a tail-matching pattern; for example, **Templates.DB/%.module**.

Unlike the files listed in most special targets, the files on this list refer to the names of dependencies and not the names of targets. As such, the special target may apply to the dependencies of many targets at once. This special target is most useful when identifying a class of dependencies found in a particular toolset for which common behavior is desired across all targets that have that dependency.

.INCREMENTAL_REPOSITORY_SIBLING: *file ...*

The sibling files listed are incremental repository files created as siblings of a primary target, may contain incomplete configuration information, and prevent **clearmake** from winking in the primary target. This special target is useful for situations where a toolset creates an incremental sibling object, and you want more control over how that object is used.

You can specify the list of files with a tail-matching pattern; for example, **%.pdb**.

Unlike the files listed in most special targets, the files on this list refer to the names of sibling objects and not the names of targets. As such, the special target may apply to the siblings of many targets at once. This special target is most useful when identifying a class of siblings found in a particular toolset for which common behavior is desired across all targets that have that sibling.

.INCREMENTAL_TARGET: *tgt ...*

Performs incremental configuration record merging for the listed targets; in other words, combines dependency information from instances of this target generated previously with the current build of this target. This special target is most useful when building library archives, because typically only some of the objects going into a library are read each time the library is updated.

You can specify the list of files with a tail-matching pattern; for example, **%.a**.

For information on restructuring a makefile to build incremental archive files, see *Working with Incremental Update Tools* on page 46.

NOTE: `.INCREMENTAL_TARGET` applies only to makefile targets built incrementally using a single make rule. Do not use it for the following kinds of files:

- Files built incrementally that are not makefile targets. For example, sibling objects like log files or template repositories.
- Files built incrementally from several different build scripts.

The general guideline is that if you're not building a library with `ar` in a single makefile rule, and you're not building an executable using an incremental linker, you should not use `.INCREMENTAL_TARGET`.

`.NO_CMP_NON_MF_DEPS` : *tgt* ...

The specified targets are built as if the `-M` option were specified; if a dependency is not declared in the makefile, it is not used in configuration lookup.

You can specify the list of files with a tail-matching pattern; for example, `%.o`.

`.NO_CMP_SCRIPT` : *tgt* ...

The specified targets are built as if the `-O` option were specified; build scripts are not compared during configuration lookup. This is useful when different makefiles (and, hence, different build scripts) are regularly used to build the same target.

You can specify the list of files with a tail-matching pattern; for example, `%.o`.

`.NO_CONFIG_REC` : *tgt* ...

The specified targets are built as if the `-F` option were specified; modification time is used for build avoidance, and no CRs or derived objects are created.

You can specify the list of files with a tail-matching pattern; for example, `%.o`.

`.NO_DO_FOR_SIBLING`: *file* ...

Disables the creation of a derived object for any file listed if that file is created as a sibling derived object (an object created by the same build rule that created the target). These sibling derived objects are left as view-private files.

You can specify the list of files with a tail-matching pattern; for example, `ptrepository/_%`.

Unlike the files listed in most special targets, the files on this list refer to the names of sibling objects and not the names of targets. As such, the special target may apply to the siblings of many targets at once. This special target is most useful when identifying a class of siblings found in a particular toolset for which common behavior is desired across all targets that have that sibling.

.NO_WINK_IN : *tgt* ...

The specified targets are built as if the **-V** option were specified; configuration lookup is restricted to the current view.

You can specify the list of files with a tail-matching pattern; for example, **%o**.

.NOTPARALLEL : *tgt* ...

Without any *tgt* arguments, disables parallel building for the current makefile. **clearmake** builds the entire makefile serially, one target at a time. With a set of *tgt* arguments, prevents **clearmake** from building any of the targets in the set in parallel with each other. However, targets in a set can be built in parallel with targets in a different set or with any other targets. For example:

.NOTPARALLEL:%a

.NOTPARALLEL:foo bar

clearmake does not build any **.a** file in parallel with any other **.a** file, and **foo** is not built in parallel with **bar**. However, **clearmake** may build **.a** files in parallel with **foo** or **bar**.

.NOTPARALLEL does not affect lower-level builds in a recursive make, unless you specify it in the makefiles for those builds or include it in a BOS file.

You can specify the list of files with a tail-matching pattern; for example, **%a**.

See also Chapter 9, *Setting Up a Parallel Build*.

.SIBLING_IGNORED_FOR_REUSE: *file* ...

The *files* are ignored when **clearmake** determines whether a target object in a VOB is up to date and can be reused. This is the default behavior, but this special target can be useful in conjunction with the **.SIBLINGS_AFFECT_REUSE** special target or **-R** command-line option. This target applies only to reuse, not to winkin.

You can specify the list of files with a tail-matching pattern; for example, **Templates.DB/%.module**.

Unlike the files listed in most special targets, the files on this list refer to the names of sibling objects and not the names of targets. As such, the special target may apply to the siblings of many targets at once. This directive is most useful when identifying a class of siblings found in a particular toolset for which common behavior is desired across all targets that have that sibling.

.SIBLINGS_AFFECT_REUSE:

Build as if the **-R** command line option were specified; examine sibling derived objects when determining whether a target object in a VOB can be reused (is up to date). By default, when determining whether a target is up to date, **clearmake** ignores modifications to objects created by the same build rule that created the target (sibling derived objects). This directive tells **clearmake** to consider a target out of date if its siblings have been modified or deleted.

5.4 Using Makefiles in Snapshot Views

Because snapshot views do not make use of the MVFS, absolute VOB pathnames (for example, **/vobs/tools/foo.h**) are not supported in snapshot views on UNIX. For that reason, your makefiles must not include absolute VOB pathnames.

To eliminate absolute VOB pathnames from makefiles, use the **pwv -root** command to get the value of the current view-root directory. Use that value in one of the following methods:

- Pass the context-dependent view root to the makefile from an external definition (for example, with **Imake**).
- Define the context-dependent view root in the makefile. For example:

```
VWROOT=`cleartool pwv -root`  
TOOLS=$(VWROOT)/vobs/tools
```

The method shown in this example works for any **clearmake** compatibility mode. There are also methods specific to each compatibility mode. See your **make** documentation for more information.

5.5 BOS File Entries

The following sections describe the entries you can put in BOS files.

Standard Macro Definitions

A standard macro definition has the same form as a make macro defined in a makefile:

```
macro_name = string
```

For example:

```
CDEBUGFLAGS = -g
```

Target-Dependent Macro Definitions

A target-dependent macro definition takes this form:

```
target-pattern-list := macro_name = string
```

Any standard macro definition can follow the := operator; the definition takes effect only when targets matching patterns in *target-pattern-list* and their dependencies are processed. Patterns in the *target-pattern-list* must be separated by white space. For example:

```
foo.o bar.o := CDEBUGFLAGS=-g
```

Two or more higher-level targets can have a common dependency. If the targets have different target-dependent macro definitions, the dependency is built using the macros for the first higher-level target **clearmake** considered building (whether or not **clearmake** actually built it).

Shell Command Macro Definitions

A shell command macro definition replaces a macro name with the output of a shell command:

```
macro_name :sh = string
```

This defines the value of *macro_name* to be the output of *string*, any shell command. In command output, <NL> characters are replaced by <SPACE> characters. For example:

```
BUILD_DATE :sh = date
```

Special Targets

You can use some ClearCase special targets in a build options spec. See *Special Targets for Use in Makefiles or BOS Files* on page 81.

Include Directives

To include one BOS file in another, use the **include** or **sinclude** (silent include) directive. For example:

```
include /usr/local/lib/ux.options
sinclude $(OPTS_DIR)/pm_build.options
```

Comments

A BOS file can contain comment lines, which begin with a number sign (#).

Using **clearmake** Compatibility Modes

6

clearmake is designed for compatibility with existing **make** programs, which minimizes the changes you need to make to your makefiles. There are many variants of **make**, and each provides different sets of extended features. **clearmake** does not support all features of all variants, and we do not guarantee absolute compatibility.

If your makefiles use only the common extensions, they will probably work with **clearmake** without changes. If you must use features that **clearmake** does not support, consider using another **make** program in a **clearaudit** shell. This alternative provides build auditing (configuration records), but does not provide build avoidance (*winkin*).

NOTE: When building with configuration records, **clearmake** handles double-colon rules differently from other **make** programs. For details, see *How clearmake Interprets Double-Colon Rules* on page 39.

To specify a compatibility mode, do one of the following things:

- Use the environment variable `CCASE_MAKE_COMPAT` in a build options specification file or in your environment. For more information, see *Build Options Specification Files* on page 74.
- Use the `-C` option with **clearmake**. For more information, see the **clearmake** reference page.

In either case, you can use any of these compatibility modes:

sgismake	IRIX smake
sgipmake	IRIX pmake
sun	SunOS make
aix	IBM AIX make
gnu	Free Software Foundation Gnu make
std	Standard clearmake with no compatibility mode enabled (Use this option to nullify a setting of the environment variable <code>CCASE_MAKE_COMPAT</code>)

The following sections describe each compatibility mode, and **clearmake**'s compatibility with System V Release 3 **make(1)**.

6.1 IRIX smake

In **sgismake** compatibility mode, the following features are enabled:

- All extended macro-assignment operators:

<code>?=</code>	Assign if undefined
<code>:=</code>	Expand RHS immediately
<code>+=</code>	Append to macro
<code>!=</code>	Assign result of shell command

- All extended macro-expansion operators:

```
$(VAR:T)
$(VAR:S/pattern/replace/)
$(VAR:H)
$(VAR:R)
$(VAR:Mpattern)
$(VAR:E)
$(VAR:Npattern)
```

- Most makefile conditional directives:

```
#if (expressions may contain 'defined' operator and 'make' operator)
#ifdef, #ifndef
#ifdef, #ifndef
#ifdef, #ifndef
#else
#elif
```

```
#elifmake,  
#elifnmake  
#elifdef,  
#elifndef  
#endif
```

- Makefile inclusion with search rules similar to those of **cpp(1)**:

```
#include <file>
```

Look for file in **/usr/include/make**.

```
#include "file"
```

Look for file in current directory, then in directories specified with **-I** command-line options, and then in **/usr/include/make**.

- Command line option **-I**, for use with **#include** statements
- Aliases for internal make macros:

```
$(TARGET)   Alias for $@
```

```
$(PREFIX)   Alias for $*
```

```
$(OODATE)   Alias for $?
```

```
$(IMPSRC)   Alias for $<
```

```
$(ALLSRC)   Alias for $>
```

NOTE: \$> is not supported by standard **make (1)**.

- **smake**-specific built-ins file: **/usr/include/make/system.mk**
- Inference rules with nonexistent intermediates
- Search paths for dependencies (**.PATH** and **.PATH.suffix**)
- Deferring build script commands ("**...**" in build script)
- **.NULL** target: specifies file-name extension to use when target has no file-name extension
- **.NOTPARALLEL** target: disables parallel building
- **.MAKE** target: specifies that a target corresponds to a sub-make; that target's build script is to be invoked even when **-n** is used
- The **-q** command-line option (see the **clearmake** reference page)

Limitations

Using `-C -sgismake` on a non-IRIX system may cause errors because different systems have different names for their built-in makefiles. To disable use of built-in rules, use `clearmake -r`.

6.2 IRIX pmake

In `sgipmake` compatibility mode, all the SGI `smake` features listed in *IRIX smake* on page 90 are enabled, along with the following features:

- If no target description file is specified on the command line, search for **Makefile** before searching for **makefile**.
- Undefined macros in build scripts are not expanded.
- Undefined macros outside build scripts cause a fatal error.
- One shell per build script (with `-C sgismake`, each command in the build script is executed in a separate shell).

6.3 SunOS 5.x make

In `sun` compatibility mode, the following features are enabled:

- All extended macro-expansion operators:
 - `+=` Append to macro
 - `:sh=` Assign result of shell command
- Pattern-replacement macro expansions:
 - `$(macro:op%os=np%ns)`
- Shell-execution macro expansions:
 - `$(macro:sh)`

- Conditional (target-dependent) macro definitions:

```
tgt-list := macro = value  
tgt-list := macro += value
```

You can use target-dependent macro definitions in the makefile and in the build options specification (BOS) file.

- Special-purpose macros:

```
HOST_ARCH  
TARGET_ARCH  
HOST_MACH  
TARGET_MACH
```

- Built-ins file specific to Sun:

- `./make.rules` or `/usr/share/lib/make/make.rules` (SunOS 5.x)

- Pattern-matching rules specific to Sun:

```
tp%ts : dp%ds
```

- The `-q` command-line option (see the **clearmake** reference page)
- Delayed macro evaluation
- `MFLAGS` environment variable

VPATH: Searches for Both Targets and Dependencies

clearmake -C sun uses the **VPATH** search list (if it exists) to look in the current view for the target if both these conditions are true:

- The target's name is not an absolute pathname.
- There is no existing file corresponding to the target's name.

For each directory in the value of **VPATH**, the directory path is concatenated with the target's name, and if there is an existing file at the resulting path, that file is evaluated.

This feature works whether or not **clearmake** uses configuration lookup (that is, either with or without the **-T** or **-F** option). If it does use configuration lookup, **clearmake** prefers to use a DO in the current view:

1. As always, **clearmake** tries to reuse the candidate DO (if any) in the current view, built at the target's name.
2. If such a candidate does not exist or does not qualify for reuse, **clearmake** searches for a candidate in the current view that was built in directories on the **VPATH**.
3. If a candidate with an appropriate name exists in a **VPATH** directory but is rejected by the configuration lookup algorithm, **clearmake** looks in the VOB database for other candidates that were built in that same **VPATH** directory.
4. If no **VPATH** directory has any candidate with an appropriate name, **clearmake** proceeds to search the VOB database for other candidates in the directory corresponding to the target's name.

NOTE: In all these cases, all DOs on which **clearmake** performs configuration lookup were built in a single directory; it traverses multiple **VPATH** directories only in deciding where to begin performing configuration lookup.

VPATH Substitutions in Build Scripts

The names of targets and dependencies in build scripts are replaced by their **VPATH**-elaborated counterparts. If a file is found using the **VPATH**, all white-space-delimited occurrences of the file's name in a build script are replaced with the pathname at which the file was found. For example:

```
VPATH = tgtmdir:depdir

bar.o : bar.c
      cc -c bar.c -o bar.o
```

If **bar.c** is found in directory **depdir**, and **bar.o** is found in directory **tgtmdir**, and the target must be rebuilt, this build script is executed:

```
cc -c depdir/bar.c -o tgtmdir/bar.o
```

Limitations

Using `-C sun` on a non-SunOS system may cause errors because different systems have different names for their built-in makefiles. You can disable use of built-in rules with `clearmake -r`.

`clearmake -C sun` uses the SunOS `arch(1)` and `mach(1)` commands to set the values of special macros (for example, `HOST_ARCH` and `HOST_MACH`). This generates error messages on systems that do not support these commands. You can ignore such messages if your build scripts do not use the special macros. Some alternatives:

- Comment out the lines in `ccase-home-dir/etc/sunvars.mk` (the description file for Sun `make` compatibility) that define the `.CLEARMAKE_ARCH` and `.CLEARMAKE_MACH` macros.
- Write shell scripts to implement the `arch` and `mach` commands.

6.4 IBM AIX make

In `aix` compatibility mode, the following features are enabled:

- Execution of build script lines that begin with `+`, even if `clearmake` is invoked with the `-n` option.
- Execution of a shell command embedded in a dependency list:

```
myprog: `ls *.o`
```
- Specification of a built-ins file on the command line, using the macro `MAKERULES` (the default built-in file is `/usr/ccs/lib/make.cfg`).
- Environment variable `MAKESHELL` overrides the make macro `SHELL` to specify the shell in which build script commands are executed. (The environment variable `SHELL` is never used for this purpose.)
- `.POSIX` target prohibits execution of recursive makes during `make -n` invocations.
- The `-q` command-line option (see the `clearmake` reference page).

6.5 Free Software Foundation Gnu make

The following sections describe supported and unsupported Gnu make features.

VPATH Separator Character

You can use colons or spaces as separators in the `VPATH` macro. For more information, see *VPATH Macro* on page 80.

Parallel Builds

When you use the `-J` option with `clearmake -C gnu`, there is no limit to the number of parallel builds.

During a recursive invocation, the value `-j 1` is set in `MAKEFLAGS`. With `clearmake -C gnu`, the value of `-J` is not reset in `MAKEFLAGS` and the value of `CCASE_CONC` is not changed.

See also Chapter 9, *Setting Up a Parallel Build*.

Supported Gnu Make Command-Line Options

`clearmake` supports most of the single-character subset of Gnu make's command-line interface. However, `clearmake` does not accept any of the long-form spellings for Gnu make command options.

NOTE: If you need to use the long-form spellings, you can write a Perl wrapper that translates the long-form options into short-form and invokes `clearmake` with the short-form options.

`clearmake -C gnu` supports the following Gnu make command-line options:

- `-d` Prints debugging information in addition to normal processing messages.
- `-e` Gives variables taken from the environment precedence over variables from makefiles.

- f** *FILE*
Reads *FILE* as a makefile.
- i**
Ignores all errors in commands executed to remake files.
- k**
Continues as much as possible after an error.
- I** *DIR*
Specifies a directory *DIR* to search for included makefiles.
- n**
Prints the commands that would be executed, but does not execute them.
- P**
Prints the data base (rules and variable values) that results from reading the makefiles, and then executes as usual or as otherwise specified.
- r**
Eliminates use of the built-in implicit rules.
- s**
Silent operation. Does not print the commands as they are executed.
- v**
Prints the version of the make program.
- w**
Prints a message containing the working directory both before and after executing the makefile.
- q**
Question mode. Does not run any commands, or print anything. Returns an exit status of zero if the specified targets are already up to date, or one if any remaking is required.

Unsupported Gnu Make Command-Line Options

- `-C DIR`
`--directory=DIR`
- `--no-print-directory`
- `--warn-undefined-variables`
- `-h`
`--help`
- `-t`
`--touch`
- `-j [JOBS]`
`--jobs=[JOBS]`
- `-l [LOAD]`
`--load-average[=LOAD]`
`--max-load[=LOAD]`
- `-o FILE`
`--old-file=FILE`
`--assume-old=FILE`
- `-W FILE`
`--what-if=FILE`
`-new-file=FILE`
`--assume-new=FILE`
- `-S`
`--no-keep-going`
`--stop)`
- `-f -`

Supported Gnu Make Features

The following features are enabled with `-C gnu` (see the *Gnu Make* manual for details):

- Conditional makefile interpretation. For example:

```
ifeq ($(CC),gcc)
    $(CC) -o foo $(objects) $(libs_for_gcc)
else
    $(CC) -o foo $(objects) $(normal_libs)
endif
```

- Simply expanded variables

```
y := $(x) bar
```

in which the right side is expanded once when the assignment is first scanned

- The `+=` syntax to append to the value of a variable
- Stripping leading sequences of `./` from file names, so *.file* and *file* are considered to be the same target
- Variable references using pattern substitution:

```
${VAR:PATTERN_1=PATTERN_2}
```

- Text-manipulation functions, such as these:

```
$(subst FROM,TO,TEXT)
```

```
$(patsubst PATTERN,REPLACEMENT,TEXT)
```

```
$(strip STRING)
```

```
$(findstring FIND,IN)
```

```
$(filter PATTERN...,TEXT)
```

```
$(filter-out PATTERN...,TEXT)
```

```
$(sort LIST)
```

```
$(dir NAMES...)
```

`$(notdir NAMES...)`
`$(suffix NAMES...)`
`$(basename NAMES...)`
`$(addsuffix SUFFIX,NAMES...)`
`$(addprefix PREFIX,NAMES...)`
`$(join LIST1,LIST2)`
`$(word N,TEXT)`
`$(words TEXT)`
`$(wordlist START, END, TEXT)`
`$(firstword NAMES...)`
`$(wildcard PATTERN)`
`$(foreach VAR,LIST,TEXT)`
`$(origin VARIABLE)`
`$(shell COMMAND)`

- ▶ The **VPATH** variable for specifying a search path for every dependency
NOTE: **clearmake** searches only in the current view. See *VPATH Macro* on page 80.
- ▶ The **vpath** statement for specifying a search path for a specified class of names
- ▶ The **export** statement
- ▶ The **unexport** directive
- ▶ The **.PHONY** target declaration
- ▶ All of Gnu Make's built-in implicit rules

- Pattern rules. For example:

```
% .o : % .c
    COMMANDS
    ...
```

- Static pattern rules:

```
TARGETS ... : TARGET-PATTERN: DEP-PATTERNS ...
    COMMANDS
    ...
```

- The automatic variables

```
$$ $@ $* $< $% $? $^ $+
```

and their file-name and directory-name variants. For example:

```
$(@F) $(@D) ...
```

- Multiline variable definition:

```
define VAR
    TEXT
    ...
endef
```

- Error processing of rules

In standard make, a rule stops processing when any command in it returns an error. In Gnu make, the rule continues processing unless the shell was invoked with the `-e` flag. **clearmake -C gnu** conforms to the Gnu make behavior.

Unsupported Gnu Make Features

The following features are not currently supported:

- Automatic remaking of any makefiles that are declared as targets. You must rebuild them explicitly.
- Controlling sub-makes by manipulating the `MAKEFLAGS` variable.
- The declarations `.DELETE_ON_ERROR`, `.INTERMEDIATE`, `.SECONDARY`.

- Syntax and semantics for target-scoped variables introduced in 3.77 of Gnu Make.
- Automatic makefile regeneration and restart if the makefile and included makefile fragments are targets in the makefile itself.
- Automatic deletion of intermediate results of a chain of implicit rules.
- Special search method for library dependencies written in the form `-lname`. For each directory on the `VPATH/vpath` list, Gnu make searches in `DIR/lib`.
- When the `EV MAKEFILES` is defined, Gnu make considers its value as a list of names of additional makefiles to be read before the others, as though they were included implicitly.

6.6 System V Release 3 make

In its default mode, **clearmake** is designed to be compatible with System V Release 3 **make(1)**, standard **make**. Standard **make** represents a least common denominator for **make** functionality; most other variants of **make** are at least upward compatible with standard **make**.

We suggest that you limit yourself to standard **make** functionality, because it allows maximum portability among hardware/software platforms, and among different variants of **make** on any single platform. It is also the most complete and most fully tested of the **clearmake** modes.

In its default mode, **clearmake** supports most standard **make** description file syntax and command-line options.

Unsupported Features and Options

clearmake does not support standard **make** inference rules for SCCS files—special-case file-name extension rules using the tilde (`~`) character.

clearmake does not support the following standard **make** options:

- `-t` (touch option) This is not supported because **clearmake** configuration lookup differs significantly from standard **make**'s build-avoidance algorithm, which is based on file modification times.
- `-f -` Reading a description file from **stdin**.

Using ClearCase to Build C++ Programs

7

This chapter describes how to use **clearmake** effectively with C++ programs built with various compilers. If you use **clearmake** to build C++ programs, read the section *Working with Templates* and the appropriate section for the compiler you use.

Compiler	Section and Page
C++ compilers based on the Cfront translator from AT&T	<i>Working with Cfront-Based C++ Compilers</i> on page 106
SPARCompiler C++ Version 4.0 or later	<i>Working with SPARCompiler C++</i> on page 123
SGI Delta/C++ compiler	<i>Working with the SGI Delta/C++ Compiler</i> on page 136
IBM AIX XLC compiler	<i>Working with the IBM AIX XLC C++ Compiler</i> on page 139
HP aC++ compiler	<i>Working with the HP aC++ Compiler</i> on page 145

7.1 Working with Templates

Many C++ compilers allow you to work with templates. A *function template* is a pattern for a set of functions. Similarly, a *class template* is a pattern for a set of classes. For example, a template for stack classes describes the form of data and functions in terms of an unspecified element type. To create a stack class, you supply a parameter, the actual type of data in the stack. Using this mechanism, you can create classes for stacks of strings or stacks of integers. These stack classes may contain different types of items, but they have the same default behaviors.

At compile or link time, the C++ compiler generates code for each template class. The code generation process is called *template instantiation*. In some cases, the process of template instantiation can conflict with **clearmake** building. These are possible symptoms of the conflict:

- **clearmake** performs unnecessary rebuilds.
- **clearmake** rebuilds an object that it could have winked in.
- The C++ compiler does not recompile a template source that changed, resulting in a link error or a run-time error.
- When **clearmake** winks in incremental repository files, information is lost.
- The output of **cleartool** commands that display configuration records (for example, **catcr** or **diffcr**) is confusing.

The exact nature of the symptoms depends on the compiler you use. For many compilers, these problems do not arise. For other compilers, you can correct these problems by modifying your *makefile* or the program source.

Explicit Instantiation

If your compiler supports it, we recommend that you use the ANSI C++ explicit instantiation syntax as defined in the ISO *Standard for the C++ Programming Language*. This method requires you to request instantiation explicitly in the source code. Refer to your C++ compiler documentation for more information on using explicit instantiation.

The Explicit Instantiation method requires more effort to use; however, it allows you to control the placement of instantiated template code into object modules. In addition, as part of the ANSI C++ standard, it may be the most portable solution. Using explicit instantiation does not conflict with **clearmake** build avoidance.

Alternative to Using the Procedures in This Chapter

The Cfront compilers discussed in the next section, as well as some other platform-specific compilers, use repositories of files to manage template instantiation. We recommend that you maintain repository directories and files in a VOB. In fact, the procedures in this chapter assume that you do so.

For compilers that use repositories to instantiate templates, an alternative to following the procedures in this chapter is to keep the repository outside the MVFS. However, be aware of the shortcomings of this solution:

- Changes to template source code are not detected unless your makefile includes explicit dependencies on the template source code.
- Repositories stored outside the MVFS are not available for `winkin`, so you must handle sharing between views manually. This process is prone to error.
- Files that use the repository may be winked in without associated template information. The missing information may cause links to fail.
- For files stored outside the MVFS, **clearmake** must rely on time stamp information for build avoidance, and may not rebuild files that need to be rebuilt.
- No configuration records are maintained for files stored outside the MVFS.

Precompiled Header Files

Some compilers support options to generate precompiled header files. A precompiled header file is created as a sibling of one compilation and is referenced by later compilations. This may cause **clearmake** to track extra dependencies that can interfere with `winkin`. The resulting conflicts are similar to those caused by template instantiation. Check your compiler documentation to see whether your compiler supports such options.

7.2 Working with Cfront-Based C++ Compilers

This section incorporates work contributed by ClearCase user Steve Vinoski of The Hewlett-Packard Company. Permission to use this work has been granted to Rational Software Corporation.

This section describes how to use **clearmake** to build C++ programs when you use a C++ compiler that is based on the **Cfront** translator from AT&T. If you use one of the following compilers, read this section:

- CenterLine C++ (with ObjectCenter)
- HP C++
- SGI OCC
- Sun SPARCCompiler C++ 2.x (with Sun SPARCworks 3.x)

Cfront Template Instantiation: Interaction with clearmake

A compiler based on **Cfront** instantiates templates at link time. During the first step of the link phase, the compiler generates the template code the program needs. It later combines object files to create an executable.

The compiler manages template code in a repository, by default a subdirectory of the build directory. In the repository, the compiler maintains a database of information to refer to during template instantiation and a cache of instantiated template source and object files. By default, the name of the repository directory is **ptrepository**.

By maintaining a repository, the compiler can reuse template code that earlier builds have generated. The compiler can ensure that cached template code is current with its source.

If you use **clearmake** to build C++ programs that contain template code, and you are building in an MVFS directory, we recommend that you use one of the procedures in this chapter to avoid any incorrect behavior that may result. For example:

- **clearmake** rebuilds source files that use templates, regardless of whether the source changed. The result is a target that is never up to date.
- If you revert to an older version of a source file, **clearmake** invokes the linker to relink the program. However, when the linker invokes the compiler to reinstantiate template code, the compiler may fail to recompile templates that depend on the source file. The result can be a link error.

- **clearmake** rebuilds the executable rather than winking it in from another view.
- The configuration record for the final program contains confusing information about the dependency structure of the program. This problem becomes evident when you execute **cleartool** commands such as **catcr** and **differ** that display configuration records.

These behaviors occur because the compiler manipulates files in the repository in a way that conflicts with **clearmake**'s build avoidance and configuration records. The result is that when a program uses templates, **clearmake** often associates incorrect configuration information with the program file and the files in the repository.

The following sections describe these issues in more detail. For information about how to resolve the incorrect behaviors, see *Models for Working With Cfront-Based Compilers* on page 109.

Link-Time Cfront Template Instantiation

This section describes how compilers that use link-time template instantiation (Cfront compilers) work.

During compilation, when the C++ compiler encounters the use of a template, it creates a *type map* file (usually named **defmap**). It records in the type map file each C++ type that it encounters and the name of the file in which the type is declared.

When the compiler links object modules into an executable, it performs a prelink step, in which it examines each object module to locate unresolved template symbols. To resolve these symbols, the template instantiation system does the following:

1. Uses the information in the type map to create source files. The source files contain the template declarations, template definitions, and type declarations with which to instantiate the templates.
2. Compiles these source files in a special mode.
3. Removes from its list the template symbols that the resulting object modules resolve.
4. Adds to its list any unresolved symbols of templates to be instantiated, as found by the resulting object modules.
5. Repeats Step #1 through Step #4 until all template symbols are resolved.
6. Calls the real link editor, **ld**, to produce the executable.

The description above shows that the C++ compiler does the following:

- Maintains a repository of data between executions; it saves information about C++ types in the **defmap** file during each compilation.
- Generates source code that is compiled into template instantiation object files.
- Performs simple link processing by searching object modules for unresolved template symbols, and by attempting to resolve them.
- Performs build avoidance by detecting dependencies between template objects generated during an earlier link and both
 - > template sources (header and definition files)
 - > the sources that use the templates

It then reinstantiates the objects only when necessary.

The compiler performs build avoidance to minimize the costs of template instantiation. If the compiler reinstantiates the necessary templates whenever a program is linked, the use of many templates in a program requires very long link times. To minimize the overhead for subsequent builds, the compiler stores the type map and the template object modules in a directory (usually named **ptrepository**), so that they can be reused if possible. The compiler determines whether a template object module can be reused by analyzing the dependencies of the source files used to create it. If any of the sources have been updated since the template object module was created, the compiler must rebuild the object module.

This build-avoidance scheme is equivalent to the standard **make** scheme and simpler than the ClearCase scheme: it does not guarantee correct configurations that can be shared by multiple users in a parallel development environment.

How Link-Time Instantiation Interferes with clearmake

This section describes how typical C++ compilers that use repository-based template instantiation at link time can interfere with the incremental updating of ClearCase derived objects.

Whenever the C++ compiler encounters a template in a source file (for example, **foo.cc**), it updates the **defmap** file. When the **defmap** file is created within an MVFS directory, ClearCase detects the updating activity. When **clearmake** invokes the compiler, the **defmap** file becomes a derived object (DO), a sibling of the object file (**.o** file). Associated with both DOs is a configuration record (CR). The CR describes the build script that was executed and the object versions (for example, of **foo.cc**, **foo.h**, and so on) that were accessed.

If the compilation of another source file causes the **defmap** file to be updated again, a new DO and CR are created for the **defmap** file, thereby overwriting the previous **defmap** DO. Therefore, whenever the **defmap** file is updated, the derived objects for all targets that modified it previously are no longer candidates for winkin.

The C++ compiler updates the **defmap** file incrementally; that is, it does not destroy the data written there during the earlier compilation of **foo.cc**. However, ClearCase has no knowledge of the semantics of the update. It assumes that the **defmap** file was completely regenerated, and that its current configuration depends only on the most recent target rebuild that updated it. If another view's configuration matches that of the most recent target rebuild in the first view, **clearmake** can wink in the **defmap** file to the second view inappropriately.

Models for Working With Cfront-Based Compilers

To help **clearmake** and the C++ compiler work together effectively, we recommend that you design your *makefile* according to one of the models described in this chapter. The model you choose depends on your needs:

- The Simple model is the easiest to use. It works best for developers who work independently of others. It is not as effective for team projects in which sharing builds is important.
- The Multiple Repositories model is fairly easy to use and solves the problems of sharing builds in a team project. Its disadvantage is that compilations are slower.
- The Forced Instantiation model places a heavy burden on you, but it is the only model that allows you to build archives and shared libraries that contain template code.

CAUTION: If you use parallel builds to build applications in C++ that use templates, you must use the Multiple Repository Model. This avoids the problem of multiple processes trying to access the same template database simultaneously, thereby causing conflicts and build failure.

The remainder of this section describes each model in more detail.

The Simple Model

The Simple model solves some of the conflicts between ClearCase and **Cfront**-based C++ compilers. It allows you to build programs naturally. However, it does not allow you to rely on **clearmake** to wink in the results of previous builds.

Insert the following line into your makefile, at any point in the file that does not break a rule or definition that spans multiple lines:

```
include ccase-home-dir/config/clearmake/cfront_simple.mk
```

NOTE: Substitute your ClearCase installation directory for *ccase-home-dir* (default is **/usr/atria**).

How the Simple Model Works

The **cfront_simple.mk** makefile contains the following **clearmake** directives:

```
.SIBLINGS_AFFECT_REUSE:  
.INCREMENTAL_REPOSITORY_SIBLING: ptrepository/defmap  
.NO_DO_FOR_SIBLING: ptrepository/_lock \\  
    ptrepository/_instfile ptrepository/defmap.old  
.SIBLING_IGNORED_FOR_REUSE: ptrepository/defmap  
.DEPENDENCY_IGNORED_FOR_REUSE: ptrepository/defmap.old
```

The **.INCREMENTAL_REPOSITORY_SIBLING** directive in **cfront_simple.mk** prevents **clearmake** from winking in object files that contain template references. This is important because winking in an object file that contains template references can cause the instantiation step to fail.

The **.NO_DO_FOR_SIBLING** directive tells **clearmake** to treat temporary files created by the **cfront** compiler as noncritical. You can accomplish the same purpose (with less overhead) by removing these noncritical files in the build scripts that perform C++ compilations. For example, you can define a suffix rule as follows:

```
.C.o:  
    $(C++C) $(C++FLAGS) -c $<  
    rm -f ptrepository/_lock ptrepository/_instfile \  
        ptrepository/defmap.old
```

Removing these noncritical files in the build script also reduces confusing **catcr** and **diffcr** output.

If you also build your program with other **make** programs, the directives in **cfront_simple.mk** are ignored (because other **make** programs do not recognize these directives).

For more information about the directives in the **cfront_simple.mk** makefile, see *Special Targets* on page 80.

Sample Scenario Using the Simple Model

The Simple model prevents **clearmake** from winking in an object file that contains template references to avoid an inappropriate winkin. To take advantage of the results of a previous build (for example, a nightly build), perform manual winkin by using the **cleartool winkin** command. The **-recurse** option of the **winkin** command may be used to wink in a hierarchy of files that produced a build.

For example, a build in view **nightly** produces executable **hello** in VOB directory **/vobs/test**. To wink in the hierarchy of DOs for the version of **hello** from that build:

1. Set your config spec to match that of **nightly**. This selects the correct version of **hello**. (An alternative is to specify a VOB-extended pathname for the derived object in Step #2.)
2. Type the following command:

```
⌘ cleartool winkin -recurse /view/nightly/vobs/test/hello
```

Because the **winkin -recurse** command winks in a hierarchy of DOs without regard to the makefile or config spec selections in the current view, it is a good idea to issue a **clearmake** command immediately after issuing the **winkin -recurse** command. This ensures that all DOs are up to date. Invoking **clearmake** also ensures that files you have checked out are rebuilt.

See the **winkin** reference page for more information.

Limitations of the Simple Model

The following points describe limitations of the Simple model and some workarounds.

- The compiler may reuse out-of-date template code that ought to have been reinstantiated.

This can happen if you unreserve a file or change your config spec in such a way that you revert to a version of a template source file that is older than the version that the compiler used in a previous build. When **clearmake** detects the change, it recompiles the target. The compiler, however, relies on a time-stamp comparison to check cached template code against its source. As a result, it may inappropriately reuse code that is in the repository. (Note that regular **make** does not detect the change at all and does not issue a rebuild.)

To work around the problem, you must force the compiler to reinstantiate templates, for example, by removing **.o** files from the repository directory.

- In certain cases, **clearmake** does not wink in object files and executables.

The **.INCREMENTAL_REPOSITORY_SIBLING** directive in **cf_{front}_simple.mk** prevents **clearmake** from winking in object files and executables that contain template code. If you want to wink in such files produced by a previous build, you must do so manually, as described in *Sample Scenario Using the Simple Model* on page 111.

- In certain cases, **clearmake** does not detect changes to template implementation source files.

The **.DEPENDENCY_IGNORED_FOR_REUSE** directive in **cf_{front}_simple.mk** avoids unnecessary rebuilding by ignoring detected dependencies on certain objects. Thus, its behavior more closely resembles that of **make**. The disadvantage of ignoring these dependencies is that after performing repeated builds, **clearmake** may not detect that a change to a template implementation source file necessitates a rebuild. To force the rebuild, you must delete the instantiation object file from the repository.

- **cf_{front}_simple.mk** expects the name of the repository to be **ptrepository**. If you use the **-ptr** compiler option to change the name of the repository, you must edit **cf_{front}_simple.mk** to reflect the change. Alternatively, edit a copy of the file and specify the copy in the include line in your makefile.

NOTE: If you use **-ptr** to change the name of the directory containing the template repository but do not change the name of the repository, you do not have to edit **cf_{front}_simple.mk**. The patterns in **cf_{front}_simple.mk** match **ptrepository** in any directory.

- When you use the Simple model, the **cleartool** commands **diffcr** and **catcr** may produce unexpected output for programs that contain template code.

The configuration record for the final program contains confusing information about the dependency structure of the program. This problem becomes evident when you execute **cleartool** commands such as **catcr** and **diffcr** that display configuration records.

The Multiple Repositories Model

The Multiple Repositories model requires you to insert lines into your *makefile* and to follow certain naming conventions in the build rules. **clearmake** then invokes a set of scripts that perform extra repository management to avoid conflicts with **clearmake** build avoidance. The scripts maintain a separate repository for each object file or executable that contains template code.

We recommend that you use this model if you depend on sharing builds within a development team, or if the problems that the Simple model does not solve are unacceptable. The Multiple Repositories model is more difficult to use than the Simple model, and requires more development time to use.

The following sections describe how to perform these tasks:

1. Ensure that your makefile uses a recognized C++ compiler macro.
2. Insert special build rules in your makefile.
3. Verify that the special build rules take effect.

Using a Recognized Compiler Macro

Ensure that your makefile uses one of the following macro invocations to represent the C++ compiler in every place that the compiler is called:

- `$(C++C)`
- `$(C++)`
- `$(CXX)`
- `$(CCXX)`
- `$(CCC)`

Note that `$(CC)` is not in this list because it usually designates a C compiler, rather than a C++ compiler.

The macro invocation must appear in every build rule for C++ objects. For example:

```
.C.o:
    $(C++) $(C++FLAGS) -c $<
```

The macro invocation must also appear wherever you execute the C++ linker. For example:

```
myprog: $(MYOBS)
    $(C++) $(C++FLAGS) -o myprog $(MYOBS)
```

Inserting Special Build Rules in Your Makefile

Edit your makefile by following these steps:

1. Insert the following line:

```
$(CCASE_CXX_INC_LNK)
```

You must insert this line in a specific place:

- > If a C++ compiler macro is defined in your makefile, you must insert this line immediately after the line containing the compiler macro definition, because the special build rule redefines the C++ compiler macro. However, if you are using this makefile only with **clearmake**, you can remove the line containing the compiler macro definition.
- > Otherwise, insert the new line at the top of your makefile.

With the incremental rules in effect, **clearmake** can manage the template instantiation process efficiently. In most cases, these macros work well. In some cases, you may want to use the alternate (CM-safe) multiple repository model. Read *Using an Alternate (CM-safe) Multiple Repository Model* on page 115 to determine whether you need to use alternate rules.

NOTE: The **CCASE_CXX_INC_LNK** macro expands to an include directive:

```
include ccase-home-dir/config/clearmake/inc_cxx.mk
```

The macro is provided so that your *makefile* can be used transparently with other **make** programs. If that is not important to you, you can use the include directive instead to enable the special build rules.

2. Decide whether to define the pathname to the C++ compiler, as follows:
 - a. If the name of the compiler is **CC** and it does not require preceding pathname components, do nothing.
 - b. Otherwise, insert a line in the format

```
REAL_C++=your-CC-program
```

where *your-CC-program* is the name of the compiler as it appears in the definition of the C++ compiler macro.

If you use Purify with this model, insert the path to Purify before the path to the C++ compiler. For example:

```
REAL_C++='purify CC'
```

Note that you must use single quotes.

Similarly, if you use PureLink with this model, insert the path to PureLink before the path to the C++ linker, for example:

```
ATRIA_C++LINK=${CCASE_MAKE_CONFIG_DIR}/atria_cxxl \  
    @$ cfront timestamp_cache 'purelink $(REAL_C++)'
```

If you are using the CM-safe model discussed in the next section, substitute **CM-safe** for **timestamp-cache**:

```
ATRIA_C++LINK=${CCASE_MAKE_CONFIG_DIR}/atria_cxxl \  
    @$ cfront CM-safe 'purelink $(REAL_C++)'
```

Using an Alternate (CM-safe) Multiple Repository Model

In most cases, the incremental build rules, described in the previous section, work well. However, if unused template code accumulates in the repository, the build rules may prevent winking.

To use the CM-safe build rules, substitute the following line of code for the link-time template instantiation build rule:

```
$(CCASE_CXX_SAFE_LNK)
```

With the CM-safe rules in effect, **clearmake** forces the C++ compiler to instantiate all templates every time it links the program. The result is that there are no superfluous dependencies to prevent **clearmake** from winking in object files. The disadvantage of using CM-safe build rules is that for every link, extra time is spent on template instantiation.

NOTE: The **CCASE_CXX_SAFE_LNK** macro expands to an include directive:

```
include ccase-home-dir/config/clearmake/cmsafe_cxx.mk
```

The macro is provided so that your makefile can be used transparently with other **make** programs. If that is not important to you, you can use the include directive instead to enable the special build rules.

Example Makefile Using the Multiple Repository Model

Consider the following *makefile*:

```
MYOBS=main.o std.o istring.o site.o point.o

.SUFFIXES: .C .o
.C.o:
    $(C++) $(C++FLAGS) -c $<

C++=/net/elm/tools/bin/CC
C++FLAGS=-g

myprog: $(MYOBS)
    $(C++) $(C++FLAGS) -o myprog $(MYOBS)
```

After you make the modifications described in *Inserting Special Build Rules in Your Makefile* on page 114, the makefile looks like this:

```
MYOBS=main.o std.o istring.o site.o point.o

.SUFFIXES: .C .o
.C.o:
    $(C++) $(C++FLAGS) -c $<

C++=/net/elm/tools/bin/CC (this line can be removed if the makefile is being used only with clearmake)
C++FLAGS=-g
$(CCASE_CXX_INC_LNK)
REAL_C++=/net/elm/tools/bin/CC

myprog: $(MYOBS)
    $(C++) $(C++FLAGS) -o myprog $(MYOBS)
```

Testing the Makefile

To test your modifications:

1. Verify that your *makefile* is still valid for any **make** programs other than **clearmake** that you use to build your program. Your modifications must not affect how other **make** programs build your program.

2. When you build your program with **clearmake**, check whether there is a difference in the commands that **clearmake** executes:
 - If your build rules instantiate the template at link time, whenever **clearmake** executes a C++ translation, it echoes a command line that begins with the following text:

```
ccase-home-dir/config/clearmake/atria_cxx cfront timestamp-cache
```

- If you used CM-safe build rules, whenever **clearmake** executes a C++ translation, it echoes a command line that begins with the following text:

```
ccase-home-dir/config/clearmake/atria_cxx cfront cm-safe
```

If the result of either of these test steps is not what you expect, examine your makefile to verify that you modified it correctly.

How the Multiple Repositories Model Works

When **clearmake** executes the build rule, the C++ macro expands to code that includes the file *ccase-home-dir*/config/clearmake/inc_cxx.mk. If you use the CM-safe rules, **clearmake** includes the **cmsafe_cxx.mk** file instead. These files add special build rules for the **Cfront** compilers. Because the macros are ordinarily undefined, other **make** programs expand the macro lines to empty text, which is ignored.

The **inc_cxx.mk** and **cmsafe_cxx.mk** files provide a wrapper around the compiler and linker. The wrapper runs the **atria_cxx** script whenever the compiler or linker is invoked. The script performs actions that work with the compiler's template instantiation method, and also runs other scripts. To cause the **atria_cxx** script to echo the command lines it executes, invoke **clearmake** with the **-v** or **-d** command line option, or set the **CCASE_VERBOSITY** environment variable as described on the **env_ccase** reference page.

By default, the C++ compiler creates only one directory, **ptrepository**, for the template instantiation repository. The effect of modifying the build rules is as follows:

- The compiler creates an **xxx.ptrep** directory for each C++ target, regardless of whether the target is an object file or a linked executable. Maintaining separate repositories is an important part of the Multiple Repositories solution to the conflict between **clearmake** and **Cfront** template instantiation.
- **clearmake** runs the **atria_make_ptrep** script to combine the repositories before the link.

- The incremental macro calls the **atria_make_ts_cache** script to create a time-stamp cache. It uses this cache to determine when rebuilds are necessary.

The CM-safe rules do not use a time-stamp cache because they always delete and re-create the entire repository.

As a result of these actions, the new build rules facilitate better sharing (winkin) by avoiding incremental updates to repository files. The CM-safe rules also prevent the use of existing template information, which may create superfluous dependencies.

Limitations of the Multiple Repositories Model

These are the limitations of the Multiple Repositories model:

- You cannot use the **-ptr** compiler option.

To maintain separate template repositories, the build script inserts a **-ptr** option onto each C++ command line. Because the build rules assume control over the repository directories, you cannot use the **-ptr** option to specify another repository directory.

- An error during instantiation invalidates all template code.

The compilers discussed in this section instantiate templates at link time. If an error occurs during template instantiation, it is probably a compilation error in a template source file. When such an error occurs, the link fails and **clearmake** aborts the build.

When the build is interrupted in this way, the **clearmake** scripts do not record configuration information for the template code in the repository. On the next build, the scripts note the absence of configuration information and cause the C++ compiler to re-instantiate all the template code required by the program. This limitation can slow the development process.

- In certain cases, **clearmake** does not wink in repository files.

The template repository for the final executable accumulates files as the build progresses. The **Cfront** compiler and linker do not delete files from the repository. As a result, a repository may contain template code that is no longer used by the program. Normally, the contents of repositories can be shared (winked in) among views if the program and its dependencies match across the views. However, the presence of unused template code in a repository can interfere with sharing. To remove unused template code accumulated in a repository, delete the entire repository directory and rebuild.

- Multiple repositories can hold duplicates of instantiation objects.

Because a separate template repository is maintained for each object file and executable, the compiler is prevented from reusing template code that was instantiated during the compilation of another module, even if the instantiations are identical. Thus, multiple repositories can consume more disk space than a single repository. For an example of how to eliminate duplicate instantiation objects from an archive library, see Step #3 on page 122.

The Forced Instantiation Model

Compilers that are based on **Cfront** instantiate templates only during the link step of a program build. These compilers determine which templates to instantiate by examining external references in the object files to be linked; the programmer does not have direct control over which templates to instantiate. In situations where you need to control template instantiation, use the Forced Instantiation model.

The basic strategy of the Forced Instantiation model is to prevent the linker from instantiating templates automatically. To do so, you instantiate all templates that are needed by an executable or by an archive library, and then name them explicitly for inclusion in that executable at link time.

The Forced Instantiation model places the burden of managing template instantiation on you. However, it is the appropriate model to use in situations that require control over instantiation. When you use this model exactly as presented here, no conflict arises between the C++ compiler and **clearmake** building.

Use the Forced Instantiation model in the following situations:

- If you are developing code for multiple platforms, some compilers may not provide an automatic instantiation mechanism, and, therefore require a model similar to this one for instantiating templates. Using a common method of instantiation across multiple platforms may be easiest for you.
- If you build an archive or shared library containing instantiated template code, you must force the compiler to instantiate the code you need and extract the resulting object files from the repository.

To use this model:

- Maintain a set of “dummy” source files.
- Set up your makefile to build a dummy executable and build template code into the final executable or library.

Maintaining Dummy Source Files

The Forced Instantiation model requires you to maintain a set of dummy source files. Each dummy source file is a C++ source file that contains references to a set of template classes or functions. For each template class that you want to instantiate, add one variable declaration to the file. For each template function, add a call to that function.

Your project’s needs determine how to distribute the references among the dummy files. You may want to have one dummy file that contains all references, or you may want to group the references to reflect how the sources in your program are divided into libraries and executables.

For example, if **Array<char>** and **List<String>** are template classes your program requires, you can produce the following dummy file:

```
// dummy source file for forcing template instantiation.

#include "array.h"      // declares class template Array
#include "list.h"      // declares class template List
#include "string.h"    // declares class String

Array<char> dummy0;    // assumes void constructor
List<String> dummy1;  // assumes void constructor
```

Additional instructions for working with dummy source files:

- Include in the dummy source file the header files that contain requisite definitions.
- Define variables, as needed, to provide as arguments to constructors.
- If the template class you need to instantiate has no public constructors, you may substitute a line for the variable declaration of the form

```
((Class *)0)->MemberFunction ();
```

where *Class* is the name of the class and *MemberFunction* is a public member function of that class.

- If the template class has no public member functions, it must declare a friend class or friend function to be usable. In this case, write a dummy version of a friend function (or a member function of a friend class) in the file and put the variable declaration into the dummy friend function. Because this code is not linked into your program, you do not need to worry about multiple symbol definitions.

Setting Up the Makefile

To set up your makefile:

1. Assign separate repositories to each object file and executable.

Assign each object file and executable its own repository by specifying the `-ptr` compiler option on every C++ compiler command line. Add the following to the options you pass to the C++ compiler:

```
-ptr$@.ptrep
```

For example, if `$(C++FLAGS)` is expanded on every C++ compiler command line, add the option to the macro definition.

```
C++FLAGS=-ptr$@.ptrep    ...other compiler options...
```

2. Instantiate the templates by building the dummy executables.

The following example shows a **make** rule that compiles the **dummy.C** source file into the **dummy.out** executable, producing the **dummy.out.ptrep** repository. The `-pta` compiler option is required. This option directs the compiler to instantiate all template functions in a class rather than instantiating only the template functions that are referenced.

```
dummy.out: dummy.C main.o std.o
    rm -rf $@.ptrep
    $(C++) $(C++FLAGS) -pta $@ dummy.C main.o std.o
```

This **make** rule deletes the existing repository before rebuilding. The rule prevents reuse of existing template code, but ensures clean configuration information for the newly instantiated template code. It is important that you follow this model in your makefile, especially to ensure that the objects you build can be shared (winked in) by other views.

The example rule in this section depends on the additional object files, **main.o** and **std.o**. Add other objects to your makefile rule as necessary to link the program. In this section's example, assume that **main.C** contains an empty **main()** function, and **std.C** contains other definitions required to link the program.

3. Build the final executable or library.
 - a. Make the final executable or library target depend on the dummy executables.
 - b. In the build rules for the final target, invoke the **atria_list_obj** utility to collect all the names of the template object files produced by the compilations of the dummy sources.
 - c. Insert the names of the object files collected in Step #b onto the compiler or **ar** command line.

For example, suppose the target is the archive **lib.a**. The dummy source files **dummy0.C**, **dummy1.C**, and **dummy2.C** contain template references required by **lib.a**. Suppose also that **\$(OBJECTS)** expands to a list of the other objects required by **lib.a**.

In the following example, the utility program, **atria_list_obj** takes one or more repository directories as command line arguments and prints a list of the **.o** files that it finds in those directories or any of their subdirectories. This example uses **atria_list_obj** to collect the names of template object files and to put them on the **ar** command line.

```
lib.a: $(OBJECTS) dummy0.out dummy1.out dummy2.out
    rm -f $@
    ar qcv $@ $(OBJECTS) \
    `atria_list_obj dummy0.out.ptrep \
    dummy1.out.ptrep dummy2.out.ptrep`
```

For more information, examine the code for the utility program in *ccase-home-dir/config/clearmake/atria_list_obj*.

How the Forced Instantiation Model Works

By default, the C++ compiler creates only one directory, **ptrepository**, for the template instantiation repository. When you modify the build rules, the compiler creates an **xxx.ptrep** directory for each C++ target, regardless of whether the target is an object file or a linked executable. Maintaining separate repositories is an important part of the Forced Instantiation solution to the conflict between ClearCase and **Cfront** template instantiation.

When you specify the **-pta** option to the linker, the linker forces the instantiation of all template objects when it builds the dummy output executable. The template instantiation objects, which behave like a dictionary, are then available for placement on a command line to build a library or final executable. The **atria_list_obj** utility aids this process.

As a result of these actions, the new build rules facilitate better sharing (winkin) by avoiding incremental updates to repository files.

Limitations of the Forced Instantiation Model

- The dummy source files must be kept in sync with the actual uses of templates in the production library or executables.
- You must use the `-ptr` compiler option as described in Step #1 on page 121. You cannot use a single repository.
- Multiple repositories can hold duplicates of instantiation objects.

Because a separate template repository is maintained for each object file and executable, the compiler is prevented from reusing template code that was instantiated during the compilation of another module, even if the instantiations are identical. Thus, multiple repositories can consume more disk space than a single repository does. For an example of how to eliminate duplicate instantiation objects from an archive library, see Step #3 on page 122.

7.3 Working with SPARCompiler C++

This section describes how to use **clearmake** to build C++ programs using SPARCompiler C++ Version 4.x. This compiler is bundled with SPARCworks Version 3.0 or later. If you use this compiler, read this section. If you use an earlier version of SPARCompiler C++, see *Working with Cfront-Based C++ Compilers* on page 106.

NOTE: The SPARCompiler C++ Version 4.0 is still evolving. New patches to the compiler may allow different methods of managing repository files. The **clearmake** models outlined here are accurate as of the time that this document was written, but may become obsolete or incomplete as new versions of the compiler are released.

SPARCompiler Template Instantiation: Interaction with clearmake

When the SPARCompiler C++ compiles a source module into an object module, it instantiates all the template code to which the module refers. The compiler places the generated template code into a repository directory, **Templates.DB**. By default, the repository is a subdirectory of the build directory. At link time, the compiler links instantiated template code from the repository into the final executable as needed to resolve references in the program. The instantiated template code remains in the repository, so that the compiler can reuse template code generated

by earlier builds. The compiler employs its own scheme to check that cached template code is up to date with its source.

If you use **clearmake** to build C++ programs that contain template code, and you are building in an MVFS directory, we recommend that you use one of the procedures in this section to avoid any incorrect behavior that may result. For example:

- **clearmake** rebuilds components of the program, regardless of whether the source changed. The result is a target that is never up to date.
- If you revert to an older version of a source file, **clearmake** invokes the compiler to recompile the source files that depend on the template. However, the C++ compiler may fail to re-instantiate the template code.
- **clearmake** rebuilds program components rather than winking them in from another view.
- The configuration record for the program components contains confusing information about the dependency structure of the program. This problem becomes evident when you execute **cleartool** commands such as **catcr** and **differ** that display configuration records.

These behaviors occur because the compiler manipulates the files in the repository directory in a way that conflicts with **clearmake**'s build avoidance and configuration records. When a program uses templates, **clearmake** often associates incorrect configuration information with the program file and the files in the repository.

Setting Up the Repository

By default, the repository directory, **Templates.DB**, is a subdirectory of the compiler's working directory. (Although you can specify the **-ptr** compiler option to change the repository directory, we recommend that you use the default directory because the makefile models in this section assume that you are using the default directory.) For each repository directory in your build system, follow these steps to set up the repository:

1. Ensure that the **Templates.DB** subdirectory is a ClearCase directory element (under version control).
2. Ensure that the file **Templates.DB/Template.opt** exists and is a ClearCase element. If it does not exist, create an empty file and convert it to an element.

The SPARCompiler input file **Template.opt** tracks template specialization, that is, a specialized implementation of a template class member or of a template function. You must

always keep the **Template.opt** file under version control, even if you do not use template specializations, to help you avoid conflicts with the **clearmake** build avoidance mechanism.

NOTE: If you are using version 4.1 or later of the SPARCompiler, it is not necessary to make **Template.opt** a ClearCase element. However, we recommend that you do so if you are adding content to this file.

Cleaning the Repository

To force a complete rebuild or to reclaim disk space, you may want to clean the repository. A common practice is to add a **clean** target to your makefile to delete the entire **Templates.DB** repository. However, because we recommend that you maintain **Templates.DB** as a versioned directory, you may not delete the entire directory. To remove the derived objects from the repository while leaving the versioned files, remove only the files that match the following wildcard patterns:

```
Templates.DB/*.o
Templates.DB/*.state
Templates.DB/*.system
Templates.DB/*.module
Templates.DB/Module.DB/*.module
```

Models for Working with SPARCompiler C++

To help **clearmake** and the SPARCompiler C++ compiler work together effectively, we recommend that you design your makefile according to one of the models described in this section. The model you choose depends on your needs:

- ▶ The *Simple* model is easy to use. It works best for developers who work independently of others. It is not as effective for team projects in which sharing of builds is important.
- ▶ The *Multiple Repositories* model is fairly easy to use and solves the problems of sharing builds in a team project. Its disadvantage is that compilations are slower.

There is no *Forced Instantiation* model for this compiler, because this compiler does not provide a way to disable automatic template instantiations.

The remainder of this section describes each model.

The Simple Model

The Simple model solves some of the conflicts between ClearCase and SPARCompiler C++ compilers. It allows you to build programs naturally.

Insert the following line into your makefile, at any point in the file that does not break a rule or definition that spans multiple lines.

```
include ccase-home-dir/config/clearmake/sunpro_4_0_simple.mk
```

How the Simple Model Works

The `sunpro_4_0_simple.mk` makefile contains the following **clearmake** directives.

```
.SIBLINGS_AFFECT_REUSE:  
.SIBLING_IGNORED_FOR_REUSE: Templates.DB/%.system \  
    Templates.DB/%.state Templates.DB/%.module  
.DEPENDENCY_IGNORED_FOR_REUSE: Templates.DB/Dependency.state \  
    Templates.DB/%.module Templates.DB/Module.DB/%.module
```

If you also build your program with other **make** programs, the directives in `sunpro_4_0_simple.mk` are ignored (because other **make** programs do not recognize these directives).

For more information about each of these directives, see *Special Targets* on page 80.

Sample Scenario Using the Simple Model

To avoid an inappropriate winkin, the simple model prevents **clearmake** from winking in an object file that contains template references. To take advantage of the results of a previous build (for example, a nightly build), perform manual winkin by using the **cleartool winkin** command. The `-recurse` option of the **winkin** command may be used to wink in a hierarchy of files that produced a build.

For example, a build in view **nightly** produces executable **hello** in VOB directory `/vobs/test`. To wink in the hierarchy of DOs for the version of **hello** from that build:

1. Set your config spec to match that of **nightly**. This selects the correct version of **hello**. (An alternative is to specify a VOB-extended pathname for the derived object in Step #2.)

2. Type the following command:

```
% cleartool winkin -recurse /view/nightly/vobs/test/hello
```

Because the **winkin -recurse** command winks in a hierarchy of DOs without regard to the makefile or config spec selections in the current view, it is a good idea to issue a **clearmake** command immediately after issuing the **winkin -recurse** command. This ensures that all DOs are up to date. Invoking **clearmake** also ensures that files you have checked out are rebuilt.

See the **winkin** reference page for more information.

Limitations of the Simple Model

These are the limitations of the Simple model and the workarounds for each:

- The compiler may reuse template code that ought to have been reinstantiated.

This can happen if you unreserve a file or change your configuration specification in such a way that you revert to a version of a template source file that is older than the version that the compiler used in a previous build. When **clearmake** detects the change, it recompiles the program template. The compiler, however, relies on a file time-stamp comparison to check cached template code against its source. As a result, it may inappropriately reuse code that is in the repository.

To work around the problem, you must force the compiler to re instantiate templates, for example, by removing the **.o** files from the repository directory.

- In certain cases, **clearmake** does not detect changes to template implementation source files.

The **.DEPENDENCY_IGNORED_FOR_REUSE** directive in **sunpro_4_0_simple.mk** avoids unnecessary rebuilding by ignoring detected dependencies on certain objects. Thus, its behavior more closely resembles that of **make**. The disadvantage of ignoring these dependencies is that after performing repeated builds, **clearmake** may not detect that a change to a template implementation source file necessitates a rebuild. To force the rebuild, you must delete the instantiation object file from the repository.

- **sunpro_4_0_simple.mk** expects the name of the repository to be **Templates.DB**. If you use the **-ptr** compiler option to change the name of the repository, you must edit **sunpro_4_0_simple.mk** to reflect the change. Alternatively, edit a copy of the file and specify the copy in the include line in your makefile.

NOTE: If you use **-ptr** to change the name of the directory containing the template repository but do not change the name of the repository, you do not have to edit

sunpro_4_0_simple.mk. (The patterns in **sunpro_4_0_simple.mk** match **Templates.DB** in any directory.)

- The configuration record for the program components contains confusing information about the dependency structure of the program.

The **cleartool** commands **catcr** and **diffcr** often produce unexpected output for programs that contain template code.

Building Archives that Contain Template Code

If you use SPARCompiler C++ compiler with the Simple model, building an archive that contains instantiated template code is relatively easy. The next few sections describe how to perform the following steps to set up the makefile:

1. Identify a set of source files that refers to all the template code you require.
2. Compile the source files that you identified in Step #1.
3. Make the archive target depend on the compilation of the source files.
4. Pass the names of the object files in the repository to the archiver.

Managing Template References

The following points describe how to manage template references, depending on which symbols you want to define in your archive:

- You want to define in the archive every template symbol referenced by the files in your archive.

You may use the default repository **Templates.DB**.

- There are template symbols referenced by the files in the archive that you do not want to be defined in the archive.

You must specify separate repositories for the templates to be archived and for the templates that are not to be archived, and build your sources accordingly. You may use the **-ptr** compiler option to direct the compiler to look for a repository in a directory other than the current working directory.

- There are template symbols you want to define in your archive but they are not referenced in any of the object files in the archive.

You must force the compiler to instantiate the templates you need. Create one or more dummy source files that refer to the missing symbols.

Building the Archive

In the following example, the target is the archive, **lib.a**, and **\$(OBJECTS)** expands to a list of objects required by **lib.a**.

The source file **template_refs.C** refers to additional template code that you want to include in the archive. The template code to be archived resides in the default repository, **./Templates.DB**.

In the following example, the utility program, **atria_list_obj**, takes one or more repository directories as command line arguments and prints a list of the **.o** files that it finds in those directories or any of their subdirectories. The example uses **atria_list_obj** to collect the names of template object files and to put them on the **ar** command line.

```
lib.a: $(OBJECTS) template_refs.o
      rm -f $@
      ar qcv $@ $(OBJECTS) \
        `atria_list_obj Templates.DB`
```

For more information, you may examine the code for the utility program in *ccase-home-dir/config/clearmake/atria_list_obj*.

The Multiple Repositories Model

The Multiple Repositories model requires you to insert lines into your makefile and to follow certain naming conventions in the build rules. **clearmake** then invokes a set of scripts that perform extra repository management to avoid conflicts with **clearmake** build avoidance. The scripts maintain a separate repository for each object file or executable that contains template code.

This model provides the benefit of winking in the results of previous builds. However, it prevents the compiler from reusing template code cached previously, so compiling takes longer. If you build archives containing template code, you may find the Multiple Repositories model too complicated.

The following sections describe how to perform these tasks:

1. Ensure that your makefile uses a recognized C++ compiler macro.
2. Insert special build rules in your makefile.
3. Verify that the special build rules take effect.

Using a Recognized Compiler Macro

Ensure that your makefile uses one of the following macro invocations to represent the C++ compiler in every place that the compiler is called:

- `$(C++C)`
- `$(C++)`
- `$(CXX)`
- `$(CCXX)`
- `$(CCC)`

Note that `$(CC)` is not in this list because it usually designates a C compiler, rather than a C++ compiler.

The macro invocation must appear in build rules for C++ objects. For example:

```
.C.o:
    $(C++) $(C++FLAGS) -c $<
```

The macro invocation must also appear wherever you execute the C++ linker. For example:

```
myprog: $(MYOBS)
    $(C++) $(C++FLAGS) -o myprog $(MYOBS)
```

Inserting Special Build Rules in Your Makefile

Edit your makefile by following these steps:

1. Insert the following line:

```
$(CCASE_CXX_COMP)
```

- > If a C++ compiler macro is defined in your makefile, you must insert this line immediately after the line containing the compiler macro definition, because the special build rule redefines the C++ compiler macro. However, if you are using this makefile only with **clearmake**, you can remove the line containing the compiler macro definition.
- > Otherwise, insert the new line at the top of your makefile.

With these rules in effect, **clearmake** can manage the template instantiation process efficiently.

NOTE: The **CCASE_CXX_COMP** macro expands to an include directive:

```
include ccase-home-dir/config/clearmake/comp_cxx.mk
```

The macro is provided so that your makefile can be used transparently with other **make** programs. If that is not important to you, you can use the include directive instead to enable the special build rules.

2. Decide whether to define the pathname to your C++ compiler, as follows:

- a.** If the name of your compiler is **CC** and does not require preceding pathname components, do nothing.
- b.** Otherwise, insert a line in the format

```
REAL_C++=your-CC-program
```

where *your-CC-program* is the name of the compiler as it appears in the definition of the C++ compiler macro.

If you use Purify® with this model, insert the path to Purify before the path to the C++ compiler. For example:

```
REAL_C++='purify CC'
```

Similarly, if you use PureLink® with this model, insert the path to PureLink before the path to the C++ linker. For example:

```
ATRIA_C++LINK=${CCASE_MAKE_CONFIG_DIR}/atria_cxx1 \  
    @$ SunCC SunCC4.0 'purelink $(REAL_C++)'
```

Example Makefile Using the Multiple Repositories Model

Consider the following makefile:

```
MYOBS=main.o std.o istring.o site.o point.o

.SUFFIXES: .C .o

.C.o:
    $(C++) $(C++FLAGS) -c $<

C++=/opt/SUNWspro/bin/CC
C++FLAGS=-g

myprog: $(MYOBS)
    $(C++) $(C++FLAGS) -o myprog $(MYOBS)
```

After you make the modifications described in *Inserting Special Build Rules in Your Makefile* on page 130, the makefile looks like this:

```
MYOBS=main.o std.o istring.o site.o point.o

.SUFFIXES: .C .o

.C.o:
    $(C++) $(C++FLAGS) -c $<

C++=/opt/SUNWspro/bin/CC      (this line can be removed if makefile is only being used with clearmake)
C++FLAGS=-g
$(CCASE_CXX_COMP)
REAL_C++=/opt/SUNWspro/bin/CC

myprog: $(MYOBS)
    $(C++) $(C++FLAGS) -o myprog $(MYOBS)
```

Testing the Makefile

To test your modifications:

1. Verify that your makefile is still valid for any **make** programs other than **clearmake** that you use. Your modifications must not affect the way **make** programs other than **clearmake** build your program.

2. When you build your program with **clearmake**, check whether there is a difference in the commands that **clearmake** executes. Whenever **clearmake** executes a C++ translation, it echoes a command line that begins with the following text:

```
ccase-home-dir/config/clearmake/atria_cxx SunCC SunCC4.0
```

If the result of either of these test steps is not what you expect, examine your makefile to verify that you modified it correctly.

How the Multiple Repositories Model Works

When **clearmake** executes the build rule, the C++ macro expands to code that includes the file *ccase-home-dir/config/comp_cxx.mk*. This file adds the special build rules suited to the Sun 4.x compilers. Because the macro is ordinarily undefined, other **make** programs expand the macro lines to empty text, which is ignored.

The **comp_cxx.mk** file provides a wrapper around the compiler and linker. The wrapper runs the **atria_cxx** script whenever the compiler or linker is invoked. The script performs actions that work with the compiler's template instantiation method, and also runs other scripts. To cause the **atria_cxx** script to echo the command lines it executes, invoke **clearmake** with the **-v** or **-d** command line option, or set the **CCASE_VERBOSITY** environment variable as described in the **env_ccase** reference page.

By default, the C++ compiler creates only one directory, **Templates.DB**, for the template instantiation repository. When you modify the build rules, the compiler creates an *xxx.ptrep* directory for each C++ target, regardless of whether the target is an object file or a linked executable. The link process uses multiple **-ptr** options to search the repositories. The new build rules eliminate unnecessary rebuilds by preventing multiple targets from writing to the same file. Therefore, maintaining separate repositories is an important part of the solution to the conflict between **clearmake** building and Sun compiler template instantiation.

Limitations of the Multiple Repositories Model

These are the limitations of using the macro described in this section:

- ▶ You cannot use the **-ptr** compiler option.

To maintain separate template repositories, the build script inserts a **-ptr** option onto each C++ command line. Because the build rules assume control over the repository directories, you cannot use the **-ptr** option to specify another repository directory.

- The compiler does not reuse template code.

The build script for invocations of this compiler re-creates the template repository for the target object on every build. The rebuilds prevent the compiler from reusing template code that was instantiated during a previous build of the same module. They also prevent the compiler from sharing template code between modules. Therefore, compiling some files will take longer when you use **clearmake**.

In addition, multiple repositories may consume more disk space than a single repository does. For an example of how to eliminate duplicate instantiation objects from an archive library or executable, see *Multiple Repositories Example* on page 135.

Building Archives that Contain Template Code

This section describes how to use the Multiple Repositories model to build an archive or shared library that contains instantiated template code. Use the Multiple Repositories model for this purpose only when necessary. We suggest that where possible, you use the Simple model to build archives or shared libraries that contain instantiated template code.

To build a library that contains instantiated template code, you must pass to **ar** or to the linker the names of all repository files containing the code you want to include in the library or executable. In the Multiple Repositories model, the files you need may reside in different repositories. Also, duplicate files may exist among repositories.

The makefile rules that build the library must specify all directories that contain template code. The rules must also eliminate duplicate code among the repositories.

To modify makefile rules to build a library:

1. Identify the source files that refer to all the template code you require.
 - Usually, you want to include in the library all template code referenced by the library source code. The set of source files that refers to all the template code you require (the *reference set*) is the set of library source files.
 - To exclude code, identify the source files that refer to it and eliminate them from the reference set.
 - To include in the archive template symbols that are not referenced by any of the source files in the reference set, you must force the compiler to instantiate these templates. Create one or more dummy source files that refer to the missing symbols. Add the dummy sources to the set of sources to compile.

The source files that you have identified will fall into the following categories:

- > Library source files that are members of the reference set
- > Library source files that are not members of the reference set
- > Dummy reference source files

The second and third categories are usually empty.

2. Build the library.

- a. Make the archive target depend on the compilation of the source file list.
- b. Collect the instantiated template code into a single directory.
- c. Pass the names of the object files in the directory to **ar** or to the linker.

The makefile rule that builds your library must depend on the compilation of all library sources and all dummy reference sources. When the rule invokes the linker or archiver, it must pass the names of all these files:

- > Library object files
- > Template object files derived from library source files that are members of the reference set
- > Template object files derived from dummy reference source files

Multiple Repositories Example

In the following example:

- The target is an archive, **lib.a**.
- **\$(OBJECTS)** expands to a list of objects required by **lib.a**.
- **\$(OBJECTS_EXCL_TMPL)** expands to a list of objects required by **lib.a** but containing references to templates you want to exclude from the library.
- **\$(DUMMYOBJ)** expands to a list of objects built from dummy reference source files.

The repositories are named **basename.o.ptrep**, where **basename.o** is the name of a compiled C++ object module. The makefile rule creates a temporary repository called **lib.a.ptrep** to collect the template object files before archiving them.

The utility program, **atria_list_obj**, takes one or more repository directories as command line arguments and prints a list of the **.o** files that it finds in these directories or any of their subdirectories. The example uses this program to collect the names of template object files and to put them on the **ar** command line.

```
lib.a: $(OBJECTS) $(OBJECTS_EXCL_TMPL) $(DUMMYOBJ)
    rm -r -f $@ $@.ptrep
    mkdir $@.ptrep
    for i in $(OBJECTS) $(DUMMYOBJ) ; do \
        cp /dev/null `atria_list_obj $$i.ptrep` $@.ptrep ; \
    done
    ar qcv $@ $(OBJECTS) $(OBJECTS_EXCL_TMPL) \
        `atria_list_obj $@.ptrep`
    ranlib lib.a
    rm -r $@.ptrep
```

This build script copies all the required template code into a single directory, and then archives all the code in the directory. The extra copy step eliminates duplicate template code from the archive.

For more information, examine the code for the utility program in *ccase-home-dir/config/clearmake/atria_list_obj*.

7.4 Working with the SGI Delta/C++ Compiler

This section describes how to use **clearmake** effectively to build C++ programs using the SGI Delta/C++ compiler. This compiler is available with SGI C++ Version 4.0 and later.

You may be using a version of the compiler that is earlier than Version 4.0. Both the Delta/C++ compiler and the earlier compiler, which is based on the Cfront translator from AT&T, are packaged with SGI C++ Version 4.0 and later.

You are using the earlier compiler if one of the following is true:

- You are using a version of SGI C++ that is earlier than Version 4.0.
- You are using Version 4.0 of SGI C++, and you run the **OCC** executable to invoke the compiler.
- You are using Version 4.0 of SGI C++, and you specify the **-use_cfront** command-line option to the compiler.

If you are using the earlier compiler, read *Working with Cfront-Based C++ Compilers* on page 106 instead of this section.

SGI Delta/C++ Compiler Template Instantiation: Interaction with `clearmake`

The Delta/C++ compiler offers several methods for building template code; the method you choose depends on your needs. The following sections describe these methods.

Automatic Instantiation

By default, the SGI Delta/C++ compiler instantiates templates automatically to build template code. During a prelink step, the compiler determines what template code the program requires and compiles the necessary code into some of the object files that make up the program. The compiler tracks how the program uses template code. It records this information in the `ii_files` subdirectory of the build directory.

If you use the Automatic Instantiation method, `clearmake` sometimes executes unnecessary rebuilds of program components. When the prelinker compiles template code into an existing object file, the dependency information that `clearmake` previously recorded for that object file is no longer up to date. The next time that `clearmake` is invoked, it will rebuild the object file. After this rebuild, the dependency information for the object file is again correct. At this point, `clearmake` no longer rebuilds that object file unnecessarily.

The Automatic Instantiation method is the easiest to use because it requires no programmer intervention. It is suitable for most applications. Aside from the unnecessary rebuilds described above, this method does not conflict with ClearCase configuration management.

Compile-Time Demand Instantiation

The Compile-Time Demand Instantiation method instantiates templates at compile time, rather than during a prelink step. To use this method, specify the `-ptused` and `-no_prelink` compiler options to the Delta/C++ compiler.

These options cause the compiler to compile all the template code that the source module refers to into the object module. If multiple source modules refer to the same template class or function, copies of the compiled template code appear in multiple object modules. When the program is linked, the linker removes duplicate template code originating from multiple object modules.

The Compile-Time Demand Instantiation method is easy to use, requiring only that you specify extra compiler options. It is suitable for most applications, especially for building archives and

shared libraries. Also, this method does not conflict with ClearCase configuration management. The disadvantage is that the compiler uses extra time and disk space to perform redundant template instantiation.

Explicit Instantiation

The Explicit Instantiation method is an alternative form of compile-time instantiation. The Delta/C++ compiler allows you to add directives to the source code to specify which template classes to instantiate.

When it compiles a source module, the compiler instantiates all the template classes specified by the directives in the source. The compiler instantiates each template classes completely, that is, it instantiates every member function and static data member of the class.

To use the Explicit Instantiation method:

1. For each template class to instantiate, add one **#pragma instantiate** directive to the source code. For example, if the program requires the **Array<String>** class, add the following directive:

```
#pragma instantiate Array<String>
```

2. In each source file that contains a **#pragma instantiate** directive, include the header files that contain definitions of the templates and classes used in the directives.
3. [Optional] Disable automatic instantiation by specifying the **-ptnone** and **-no_prelink** compiler options. Automatic instantiation does not interfere with explicit instantiation, but you can disable it, if you prefer.

The Explicit Instantiation method requires more effort to use. However, it allows you to control the placement of instantiated template code into object modules. This control is useful in some situations, especially when building archives of instantiated template code. Using explicit instantiation does not conflict with **clearmake** build avoidance.

7.5 Working with the IBM AIX XLC C++ Compiler

This section describes how to use **clearmake** effectively to build C++ programs using the IBM AIX XLC compiler.

XLC Compiler Template Instantiation: Interaction with clearmake

The XLC compiler instantiates a template when it encounters a **#pragma** directive in the source code. You can insert these directives into the source code manually, or you can direct the compiler to generate auxiliary source code that contains the necessary directives.

The XLC compiler's automatic template instantiation works as follows:

1. As the XLC compiler compiles a source module into an object module, it notes the templates that the module refers to.
2. When the compiler finishes the translation, it generates auxiliary source modules in a repository directory. By default, the repository directory is **tempinc**.
3. At link time, the compiler compiles the auxiliary sources and links the resulting object modules into the final executable. The template source and object modules remain in the repository, so that the compiler can reuse template code generated by earlier builds. The compiler verifies that cached template object modules are up to date with their sources.

If you use **clearmake** to build C++ programs that contain template code and are building in an MVFS directory, we recommend that you follow one of the procedures in this section to avoid any incorrect behavior that may result. For example:

- **clearmake** executes unnecessary recompiles of program components.
- **clearmake** fails to link in available derived objects.
- After a template source file changes, **clearmake** does not rebuild.
- **clearmake** rebuilds the executable, but the compiler fails to re-instantiate template code for which the source has changed.
- The output of **cleartool** commands such as **catcr** and **diffcr** that display configuration records is confusing.

This behavior occurs because the compiler manipulates files in the repository in a way that conflicts with **clearmake** build avoidance and configuration records. When a program uses templates, ClearCase often associates incorrect configuration information with the program file and the files in the repository.

Models for Working With IBM XLC

To help **clearmake** and the XLC C++ compiler work together effectively, we recommend that you set up your makefile and program sources according to one of the models described in this section. The model you choose depends on your needs:

- The *Simple* model is easy to use and solves some of the problems mentioned earlier. It is the only model that makes use of automatic template instantiation to build template code efficiently. It works best for developers who work independently of others. It is not as effective for team projects in which sharing builds is important.
- The *Compile-Time Demand Instantiation* model requires work to set up, but it is easy to use and avoids all the **clearmake** problems mentioned earlier. It also simplifies sharing builds in a team project. The penalty of using this model is that compile time can increase.
- The *Explicit Instantiation* model is difficult to use, because it requires you to track template references. This model avoids all the **clearmake** problems mentioned earlier, and simplifies sharing builds in a group project. It is also the best model to use for building archives and shared libraries that contain template code.

The remainder of this section describes each model in more detail.

The Simple Model

The Simple model relies on the automatic template instantiation of the XLC compiler. This model requires you to insert one line into your makefile.

The Simple model does not solve all the problems that result from conflicts between XLC automatic template instantiation and **clearmake**, but it allows you to make minimal changes to the project and to take advantage of the efficiency of automatic template instantiation. Configuration problems can result from version changes in the template sources. The Simple model is not appropriate for building libraries.

Modifying the Source Files

The XLC automatic template instantiation scheme requires that for every `.h` template header file, there is a corresponding `.c` source file with the same name that contains the implementation of the template. The `.c` extension is unusual in C++ source file names; you may want to set up your project source files so that each `.c` file refers to a corresponding source file with a more standard extension, such as `.cxx` or `.C`. You can set up a reference `.c` file in two ways:

- Create a new element with the `.c` extension that includes the corresponding file of the standard extension.
- Create a link with the `.c` extension that points to the corresponding file of the standard extension. See the **cleartool In** reference page for instructions on how to create a link in an MVFS directory.

NOTE: Do not set up template header files to include the corresponding source file. Doing so disables automatic template instantiation and defeats the Simple model.

Designing Your Makefile

To design your makefile:

1. Insert the following line into your makefile, at any point in the file that does not break a rule or definition that spans multiple lines.

```
.DEPENDENCY_IGNORED_FOR_REUSE: tempinc/%.C
```

This directive instructs **clearmake** not to rebuild program components that depend on a compiler-generated source file in the repository when that file changes or is deleted. Such rebuilds are unnecessary. This directive works if the leaf name of the repository directory is **tempinc**.

2. If you use the `-qtempinc=tempinc` compiler option to rename a repository directory, add to your makefile another directive of the form

```
.DEPENDENCY_IGNORED_FOR_REUSE: my_repository/%.C
```

where *my_repository* is the leaf name of the repository.

Limitations of the Simple Model

Some problems remain when you use the Simple model. The following list describes the problems and how to work around them.

- After you modify a template source file, **clearmake** sometimes fails to relink the executable.

To work around this problem, delete the executable before you build with a new version of a template source file in your program.

- When you revert to a version of a template source or header file that is older than the version the compiler used in a previous build, the compiler may reuse object modules in the repository that it should have recompiled. This can happen after you unreserve a template source or header file or after you modify your configuration specification. **clearmake** detects the change and issues a recompile. However, because the compiler relies on a file time-stamp comparison to check cached template code against its source, it incorrectly reuses the code already in the repository.

To work around the problem, remove all **.o** files from the repository when you revert to an older version of a template source or header file.

- **clearmake** rebuilds program components rather than winking in available builds from another view. This problem makes the Simple model more difficult to use for team projects in which sharing builds is important.

To manually wink in files produced by another build, use the **cleartool winkin** command. Use the **-recurse** option of the **winkin** command to wink in a hierarchy of files that produced a build. For more information, see the **winkin** reference page.

- The configuration record for the final program contains confusing information about the dependency structure of the program. This problem becomes evident when you execute **cleartool** commands such as **catcr** and **diffcr** that display configuration records.

The Compile-Time Demand Instantiation Model

The Compile-Time Demand Instantiation model does not use the XLC compiler's automatic template instantiation. Instead, it forces the compiler to instantiate all the template code used by a given module. At link time, there may be multiple definitions of template symbols, but the linker drops the duplicate symbols.

To use this model, you may have to edit some source files. The advantage of using this model is that avoiding automatic template instantiation solves the conflicts between the compiler and **clearmake**. The disadvantage is that compiling duplicate template code takes extra time and disk space.

Modifying the Source Files

The Compile-Time Demand Instantiation model requires that every template header file include its corresponding source file. Therefore, if you use this model, you must edit template header files if they do not include their corresponding source file.

We recommend that you use conditional compilation directives as shown in the following pattern.

```
// Header file my_tmpl.h

//

#ifdef _my_tmpl_h__ // header file exclusion -- use any
                    // unique name

#define _my_tmpl_h__

... Place template declarations here ...

#ifdef TEMPLATE_CODE_IN_HEADERS

#include "my_tmpl.cxx" // source file name suffix may vary

#endif

#endif // _my_tmpl_h
```

Designing Your Makefile

If you set up your header files as recommended in the previous section, you must also define the macro that enables source file inclusion on the compiler command line. For the sample above, enable source file inclusion by specifying the compiler option **-DTEMPLATE_CODE_IN_HEADERS**. The conditional directive allows you to build under a different model or on other platforms without changing the source code.

To disable automatic template instantiation, specify the **-qnotempinc** compiler option.

Duplicate Symbol Warnings From the Linker

When you set up template header files to include their source files, the implementation of each template is available to the compiler at compile time. The compiler instantiates the template

immediately and puts the compiled template code into the object module. The result is that every object module that refers to a certain template function or data contains the code for that template function or data. This leads to multiple definitions of template symbols at link time. The AIX linker drops the duplicate symbols without generating an error. As the linker drops each duplicate symbol, it issues a `Duplicate symbol` warning. You can ignore these warnings. To suppress them, specify the compiler option `-LOOK_IT_UP` in the final link step of the program.

The Explicit Instantiation Model

The Explicit Instantiation model is useful whenever it is necessary to assume full control over the placement of instantiated template code. For example, if you are building an archive or creating a shared library containing template code, you probably want to use this model. To use the Explicit Instantiation model, you must disable automatic template instantiation and on-demand compile-time instantiation, and explicitly identify the templates that your program needs to instantiate.

The Explicit Instantiation model requires that you keep track of the set of templates the program requires and maintain source directives to instantiate the templates. This maintenance burden is worthwhile only when you are required to specify the module in which instantiated template code appears, for example, when building code for an archive or shared library. Otherwise, the Compile-Time Demand Instantiation model is more suitable.

Modifying the Source Files

For each template class or function that your program requires, add exactly one `#pragma define` directive to the source code. For example, if your program requires the class

```
Array<String>
```

then add this `#pragma define` directive:

```
#pragma define(Array<String>)
```

In each source file that contains a `#pragma define` directive, include the source files that contain the implementations of the templates used. For example, if the source for the `Array` template is contained in the file `array.cxx`, add the following `#include` directive to the file that contains the `#pragma` directive:

```
#include "array.cxx"
```

Do not set up header files to include corresponding source files. If you do, duplicate template instantiation may result (see *The Compile-Time Demand Instantiation Model* on page 142).

Remove function definitions from source files that contain **#pragma define** directives.

Designing Your Makefile

In your makefile, disable automatic template instantiation by specifying the **-qnotempinc** compiler option.

7.6 Working with the HP aC++ Compiler

This section describes how to use **clearmake** to build C++ programs using the HP aC++ compiler (aCC).

If you are using the earlier HP Cfront-based compiler (CC), read *Working with Cfront-Based C++ Compilers* on page 106 instead of this section.

The HP aC++ compiler offers several methods of building template code; the method you choose depends on your needs. The following sections describe these methods.

Automatic Instantiation

By default, the HP aC++ compiler instantiates templates automatically to build template code. During a prelink step, the compiler determines what template code the program requires and compiles the necessary code into some of the object files that make up the program. The compiler tracks how the program uses template code. There is no repository; instead, the compiler records this information in **.o** and **.I** files in the build directory.

If you use the Automatic Instantiation method, **clearmake** sometimes rebuilds program components unnecessarily. When the prelinker compiles template code into an existing object file, the dependency information that **clearmake** previously recorded for that object file is no longer up to date. The next time **clearmake** is invoked, it rebuilds the object file; the dependency information for the object is again correct. At this point, **clearmake** no longer rebuilds that object file unnecessarily.

The Automatic Instantiation method is the easiest method to use because it requires no programmer intervention. It is suitable for most applications. However, it does cause occasional unnecessary rebuilds.

Command-Line Option Instantiation

The HP aC++ compiler provides various command line options for specifying what templates to instantiate for a given translation unit. For example, the **+inst_all** option requests instantiation of all templates and the **+inst_used** option requests instantiation of templates that are used. With these options, the compiler compiles all template code into the object module. If multiple source modules refer to the same template class or function, copies of the compiled template code appear in multiple object modules.

The command-line instantiation methods are easy to use; they require only that you specify extra compiler options. These methods are suitable for most applications and do not conflict with **clearmake** build avoidance. The disadvantage of using these methods indiscriminately is that the compiler requires extra time and disk space to perform redundant template instantiation. Indiscriminate use may also result in duplicate symbols.

Note that a special command-line option, **+inst_close**, is provided for building archives and shared libraries. With this option, the compiler uses the automatic instantiation method.

Explicit Instantiation

The HP aC++ compiler supports the ANSI C++ explicit instantiation syntax as defined in the ISO *Standard for the C++ Programming Language*. This explicit instantiation method requires you to request instantiation explicitly in the source code. For more information on using explicit instantiation, see the HP aC++ compiler documentation.

The Explicit Instantiation method requires more effort to use; however, it allows you to control the placement of instantiated template code into object modules. In addition, as part of the ANSI C++ standard, it may be the most portable solution. Using explicit instantiation does not conflict with **clearmake** build avoidance.

Using ClearCase Build Tools with Java

8

The building behavior of Java™ tools causes various problems for **clearmake**. This chapter presents these problems and some possible solutions.

8.1 ClearCase Build Problems with Java

Java source is kept in files with extension **.java**, where the file name must be constructed from the name of the class it defines. For example, a class **foo** referenced by other Java source files must be in a file named **foo.java**. Compiling **foo.java** with the Java compiler creates a file called **foo.class**. Subsequent compilations may read the **.class** files generated previously.

When the Java compiler encounters a reference to a class defined in another file, it rebuilds that class file if it is out of date or does not exist. This behavior puts extra information in the configuration record:

- Extra **.class** files as siblings of the target **.class** file
- Extra dependencies on **.java** sources from building the siblings

This information can cause unnecessary rebuilding, prevent winks, and create confusing **catcr** output. The major Java development toolkits exhibit this behavior.

Java Toolkits

The standard Java toolkit is the Java Development Kit from Sun, which includes the **javac** compiler and is available for many platforms.

The remainder of this chapter uses **javac** from Sun as the example.

Scope of the Problems

The build problems relate to conflicts between the dependency analyses of **clearmake** and **javac**. Because the Java compiler does some dependency analysis, some developers may not require **make** tools. Environments such as Visual J++[®] do not require that you use makefiles, nor do they generate and use makefiles themselves as Visual C++[®] does.

The need for **make** tools is reduced further because the Java language minimizes the need to recompile a dependent class file when the depended-on file changes. Java keeps interface and implementation in the same file, so changes to the implementation part of the file do not strictly require recompilation. This behavior contrasts with C/C++ applications, in which interface is separated from implementation by splitting the source into header (**.h**) and implementation (**.c**) files. Some developers may prefer to control when to rebuild Java sources. However, none of the current tool environments is based on such a language-sensitive recompilation; like **make**, **javac** uses time stamps to determine when to rebuild.

8.2 Benefits of Using make Tools with javac

Although **javac** does a good job taking care of dependency analysis, using **javac** by itself misses rebuilds that a **make** tool does not. There are additional benefits of using **clearmake** beyond those of using a non-ClearCase **make** tool, especially given the building behavior of **javac**.

Using javac Inside a Makefile

make detects modifications of indirect dependencies that **javac** does not. If **a.java** depends on **b.java** and **b.java** depends on **c.java**, when you change **c.java**, **javac a.java** does not rebuild **c.class**. Therefore, if you are using **javac** directly, you must recompile each file as you change it.

In addition, many Java applications have some components that are compiled natively or are written in another language. For at least those parts of their applications, developers need makefile-based building.

Using javac with clearmake Instead of make

clearmake provides additional benefits to those gained by using **make** or running **javac** directly. In addition to recording a build configuration, **clearmake** is better at rebuild determination than the Java tools, with or without **make**. For example, **clearmake** detects the following rebuild cases, but **javac** does not:

- Selection of an older version of a **.java** file. Because the rebuild decision is based on an older/newer comparison, **javac** does not detect that a rebuild is necessary.

NOTE: Clock skew between machines can cause similar time stamp problems outside of ClearCase.

- Change of the **javac** command line. If the command-line options used to build a **.class** file have changed since the last build, **clearmake** rebuilds the **.class** file. For example, if you add the **-g** switch to direct the compiler to rebuild with debugging information, you must invoke the compiler on all your **.java** files to ensure that they are rebuilt to contain the debugging information.
- Manual winkin of a **.class** file that is out of sync with, but newer than, the corresponding **.java** source selected by the view. Because the rebuild decision is based on an older/newer comparison, **javac** does not detect that a rebuild is necessary.

8.3 Unnecessary Rebuilds and Prevention of Winkin

`javac`'s build behavior causes `clearmake` to perform extra rebuilds and prevent winkins:

- Because a `.class` file is sometimes built as a sibling of another `.class` file, the build script for the sibling differs from what it would be if the `.java` file were compiled directly. `clearmake` rebuilds unnecessarily in this case because the build scripts do not match. Mutually dependent `.java` files are an extreme case of this behavior, because the build of one can change the build script of the other.
- Similarly, because the set of dependencies does not remain consistent from one build to the next, `clearmake` rebuilds because versions do not match.
- If a sibling derived object is overwritten, winkins are prevented.

8.4 Building Java Applications Successfully

The following alternatives allow you to successfully build Java applications with `clearmake`:

- Write the makefile correctly
- Allow `clearmake` to rebuild
- Configure `clearmake` makefiles to behave like `make`

The following sections describe each option in detail.

Writing Correct Makefiles

A correctly written makefile results in a correct set of configuration records, which gives you the full power of ClearCase configuration records and winkin without unnecessary rebuilding and without missing rebuilds. You can restructure a makefile to avoid `javac`'s automatic building behavior by enforcing that files on which other files depend are built before their dependents.

NOTE: `clearmake` detects implicit dependencies but cannot determine build order dependencies. If you want files to build in a certain order, you must declare that order in the makefile by adding additional dependencies.


You must take extra care when handling mutually dependent files, because there is not necessarily a correct order for building them. One possibility is to always generate all mutually dependent files as one unit, that is, in one configuration record. You can write the build script for a set of mutually dependent files to delete all class files corresponding to those files before building any of them. This ensures that they are not overwritten individually and makes them available as a unit for winkin.

The advantage of writing your makefile correctly is that it does not cause extra compilations or rebuilds. No special makefile directives are required, the configuration records have no unusual properties, and winkins will work fully. The disadvantage is that the makefile must always be synchronized with the structure and dependencies of the application.

The following sections are makefile examples for applications with particular dependency characteristics.

No Mutually Dependent Files

In this application, classes **x**, **y**, and **z** have a hierarchical dependency graph:




The makefile for such a dependency structure is very simple:

```
.SUFFIXES: .java .class

.java.class:
    javac $<
x.class: y.class
y.class: z.class
```

Mutually Dependent Files

This application consists of classes **top**, **a**, **b**, **c**, and **d**, which have a more complex dependency structure:



The makefile for this dependency structure is somewhat longer, but correct:

```
top.class: a.class b.class
    javac top.java

a.class: b.class

b.class:
    rm -f a.class b.class
    javac a.java b.java

b.class: c.class

c.class: d.class

d.class:
    rm -f c.class d.class
    javac c.java d.java
```

Allowing Rebuilds

If you continue to invoke **clearmake** until it determines that all files are up to date, the other ClearCase features work correctly. The configuration records record all class files as implicit dependencies rather than siblings, which allows **winkin** to work.

However, the number of rebuilds can become very large if the makefile is written incorrectly. It is possible to map out a correct set of dependencies as described in *Writing Correct Makefiles* on page 150; it is also possible to request (however inadvertently) that **clearmake** build the files in exactly the reverse, and most inefficient, order.

In addition, **clearmake**'s default behavior is to ignore modifications to siblings for the purposes of rebuilding. For **winkin** to work correctly, you must reenable that behavior by using a command-line option or special makefile directive.

Another drawback to this method is that the builds of mutually dependent source files do not fit well, because the files are never up to date. The makefile for these must be written carefully, as described in *Writing Correct Makefiles* on page 150.

Configuring Makefiles to Behave Like make

By using special targets, you can configure your **clearmake** makefile so that **clearmake** behaves as **make** does with regard to Java builds. The following targets eliminate the extra rebuilding described in *Allowing Rebuilds* on page 152:

```
.NOCMP_SCRIPT: %.class           (omake only)
.NO_CMP_SCRIPT: %.class         (clearmake only)
.DEPENDENCY_IGNORED_FOR_REUSE: %.class
```

.NO_CMP_SCRIPT disables build script checking. However, relevant build-script changes are ignored also. In addition, **.NO_CMP_SCRIPT** has no effect during **winkin**, so even when it is in use, **winkins** are prevented because of build script differences. Therefore, you must use manual **winkins** (see the **winkin** reference page) or forego them entirely.

.DEPENDENCY_IGNORED_FOR_REUSE disables the version checking of implicit dependencies when **clearmake** is looking for DOs to reuse. This can cause desired rebuilds to be missed, however. One benefit of using **clearmake** is automatic dependency detection (for example, of **.h** files in a C build), so it is not desirable to give this up.

To improve the missed implicit dependency checking caused by **.DEPENDENCY_IGNORED_FOR_REUSE**, you can add the missing dependencies as explicit dependencies in the makefile. However, this is a manual process, and you still lose build script checking and **winkin**. The remaining benefit of using **clearmake** is configuration records (though the **catcr** output for them may be confusing).

Setting Up a Parallel Build

9

This chapter describes the process of setting up and running builds that use several hosts in the local area network. It describes the control mechanisms for the client and server.

9.1 Overview of Parallel Building

ClearCase can perform builds in which multiple processes execute in parallel the build scripts associated with makefile targets. The processes executing the build scripts can run on a single machine or across a collection of machines around the local area network (parallel distributed build). By using more concurrent processes, parallel builds can reduce the overall build time significantly. Instead of one process running one build script at a time, you can have multiple processors working in parallel. For large software systems, this performance improvement can make a critical difference. For example, you can use parallel builds to enable a build of your entire software system to run overnight and finish before developers and testers arrive for work in the morning.

You start a parallel build the same way as a single-host build: by entering a **clearmake** command. A command-line option or environment variable setting causes the build to run in parallel mode.

A parallel build is controlled by specifications on all the hosts involved:

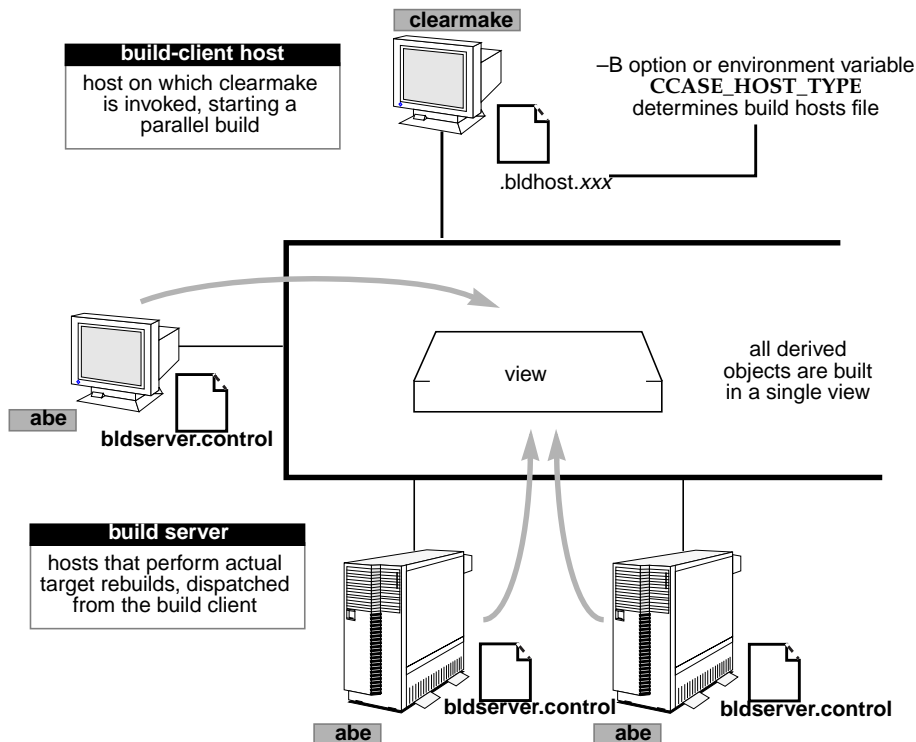
- ▶ Client controls. The host where you enter the **clearmake** command is the *build client* or *build controller*. On this host, you specify a limit to the number of build scripts to be executed concurrently. You can also specify a build hosts file, which lists the server hosts to be used for building.

- Server controls. Each build server host used in a parallel build can have an access-control file, named **bldserver.control**. To use the host as a build server, a build client must meet the access-control requirements.

When building in parallel, **clearmake** starts one or more Audited Build Executor (**abe**) processes. Each **abe** process sets to the same view and working directory as the **clearmake** process, then each process executes build scripts dispatched to it from the controlling **clearmake** process. When the **abe** process is being started on a remote machine, **clearmake** uses the standard UNIX remote shell facility, referenced through the *ccase-home-dir/etc/rsh* symbolic link. A build script runs under **abe** control as if it were being executed by **clearmake**, except that **abe** collects terminal output produced by the build script and sends it back to the build controller, where it appears in your **clearmake** window.

Figure 10 illustrates the ClearCase parallel build architecture.

Figure 10 ClearCase Parallel Build Architecture



9.2 Setting Up the Client Host

There are three issues to consider when you set up client-side processing for parallel builds:

- The number of parallel build processes to request. **clearmake** limits the number of concurrent target rebuilds to the value of the `CCASE_CONC` environment variable or the number you specify with the `-J` command-line option. **clearmake** does not start more **abe** processes than specified.
- The hosts to specify for builds. **clearmake** starts all **abe** processes on the local host unless you provide a *build hosts file*. You can specify a build hosts file with the `-B` command-line option or with the `CCASE_HOST_TYPE` environment variable. For information on the format of the build hosts file, see the **bldhost** reference page.
- The limitations and requirements for system loading on the build hosts. ClearCase tries to keep the build machines from becoming overloaded, by dispatching a build script to a host only if the host is sufficiently idle. The idleness threshold is specified in the build hosts file. If you do not specify an idleness threshold or do not use a build hosts file, **clearmake** submits as many build scripts as allowed by the `-J` or `CCASE_CONC` setting.

Creating Build Hosts Files

Build hosts files that you specify with the `-B` option can be located anywhere and do not have to have special names.

Build hosts files that you specify using the `CCASE_HOST_TYPE` environment variable must be located in your home directory, and each file must have a name that begins with **.bldhost**. Choose a file-name extension for each build hosts file that describes its intended use. For example:

.bldhost.sun5	List of hosts used to build SunOS 5 binaries
.bldhost.day	List of hosts used to perform parallel builds during the workday
.bldhost.night	List of hosts used to perform overnight parallel builds

Your build environment determines whether you need multiple build hosts files. In a heterogeneous network, for example, architecture-specific builds may or may not need to be performed on hosts of that architecture. (You may have cross-compilers, which eliminates this restriction.)

When you start a parallel build, you can specify a certain build hosts file or have **clearmake** select one. See *Starting a Parallel Build* on page 160.

For example, you can set up two build hosts files, for daytime and nighttime use, as follows:

1. Create a build hosts file for daytime use. For daytime builds, you can use the list of hosts that your system administrator has provided in `/usr/local/lib`, along with your own host. To minimize the disruption to other work, you can specify that each host is to be used only if it is not heavily loaded: that is, if it is at least 75% idle.

```
% cat > $HOME/.bldhost.day
-idle 75
neptune
#include /usr/local/lib/day_builds
<CTRL-D>
```

2. Create a build hosts file for overnight use. For overnight builds, you can use another list of hosts provided by the system administrator.

```
% cat > $HOME/.bldhost.night
#include /usr/local/lib/night_builds
<CTRL-D>
```

Because this file does not include an `-idle` specification, **clearmake** uses a host only if it is at least 50% idle.

For details on build hosts files, see the **bldhost** reference page.

If `CCASE_HOST_TYPE` is set, but **clearmake** cannot find or read the build hosts file, it does not perform the build. (**clearmake** assumes that if you set `CCASE_HOST_TYPE`, you want to perform a parallel build. Because the parallel build might use a host with a different architecture than the local host, performing the build on the local host might yield incorrect results. Therefore, **clearmake** does not default to building on the local host if `CCASE_HOST_TYPE` is set.)

clearmake will not use a host during a parallel build if your current view cannot be used on that host. (For example, the host might not be able to access the view's storage directory.)

Setting Up Trust Relationships

For parallel building to work correctly, the client host must be trusted by the build hosts (that is, remote login from the client host to the build hosts must work without a password being needed). You can set up this trust with **.rhosts** files on the hosts, or by having your system administrator set up general trust (for example, with `/etc/hosts.equiv`). To test remote login, execute the following command (substituting your ClearCase installation directory for `ccase-home-dir`):

```
ccase-home-dir/etc/rsh remote-hostname echo Good-to-use
```

If you are prompted for a password, the client host is not trusted.

For example:

```
/usr/atria/etc/rsh puffin echo Good-to-use
```

```
Good-to-use
```

(current host is trusted by host puffin)

```
/usr/atria/etc/rsh auk echo Good-to-use
```

```
permission denied
```

(current host is not trusted by host auk)

9.3 Setting Up Server Hosts

Each build server host can have a **bldserver.control** file, which controls its use for parallel builds. This file can impose restrictions on who can use the host for parallel builds, and when the host can be used for parallel builds. If a build server host has no such file, it accepts all parallel build requests. The **bldserver.control** reference page describes this mechanism.

To set up a build server host that is used for your team's daytime builds and its overnight builds:

1. Create a **bldserver.control** file. Each line of the **bldserver.control** file defines a situation in which it accepts parallel build requests.

```
% cat > /var/adm/atria/config/bldserver.control
```

```
-time 08:30,19:30 -idle 60 (1)
```

```
-time 19:30,05:30 (2)
```

```
-user bldmeister (3)
```

```
<CTRL-D>
```

Line 1 specifies that during the interval between 8:30 A.M. and 7:30 P.M., this host accepts a parallel request when it is at least 60% idle. Line 2 specifies that during the interval between 7:30 P.M. and 5:30 A.M., this host accepts any parallel request, no matter how busy it is. Line 3 specifies that a parallel build request from a **clearmake** invoked by user **bldmeister** is always accepted.

2. Protect the **bldserver.control** file to make sure that your access-control settings cannot be deleted or altered:

```
% chmod 444 /var/adm/atria/config/bldserver.control
```

Idleness Threshold

You can specify the idleness threshold with **-idle** settings on both the client and server sides. If there is a conflict, the higher threshold controls whether the build host is used. For example:

- ▶ A **clearmake** process searches for hosts that are at least 50% idle (the default). A build server that appears to qualify because it is 70% idle is not used if its **bldserver.control** file includes an **-idle 75** specification.
- ▶ A **bldserver.control** file on a build server host permits access, because **-idle 60** is specified on a host that is currently 75% idle. But **clearmake** does not dispatch a build script to this host, because the build hosts file specifies an even higher threshold, **-idle 80**.

9.4 Starting a Parallel Build

To start a parallel build:

1. Set the **CCASE_HOST_TYPE** variable. The value of this variable determines which build hosts file **clearmake** looks for in your home directory:

Value	Build hosts file clearmake looks for
sun5	.bldhost.sun5
SUN5	.bldhost.SUN5
day	.bldhost.day
night	.bldhost.night

2. Invoke **clearmake**. To enable parallel building, use the **-J** command-line option or set the **CCASE_CONC** environment variable. To specify a build hosts file, use the **-B** option. To have **clearmake** choose a build hosts file, do not use **-B**.

For example, to specify a build hosts file and start a build that uses up to five hosts concurrently, use one of the following methods:

```
% clearmake -J 5 -B ~/.bldhost.day my_target           (command-line options)
% setenv CCASE_CONC 5                                 (environment variable and
% clearmake -B ~/.bldhost.day my_target                command-line option)
% setenv CCASE_HOST_TYPE day                           (environment variable and
% clearmake -J 5 my_target                             command-line option)
```

```
% setenv CCASE_CONC 5                                     (environment variables)
% setenv CCASE_HOST_TYPE day
% clearmake my_target
```

NOTE: If you specify `-J`, but do not set the `CCASE_HOST_TYPE` variable or specify a build hosts file with `-B`, `clearmake` builds run in parallel on the local host.

Setting `CCASE_HOST_TYPE` in a Shell Startup Script

In some parallel build environments, you may find it convenient to have your shell startup script set `CCASE_HOST_TYPE`. For example, your team may support an application on several architectures.

In this scenario, you build the application for a particular architecture like this:

1. Log in to a host of that architecture.
2. Set a view and go to the appropriate directory.
3. Enter a `clearmake -J` command to start a parallel build.

You can implement such a scheme as follows:

1. Use architecture-specific build hosts files. Give each build hosts file a file-name extension that names a target architecture: `.bldhost.hpux9`, `.bldhost.sunos5`, and so on. Typically, each of these files lists hosts of one architecture only. For example, all SunOS 5 hosts are listed in `.bldhost.sunos5`.
2. Set `CCASE_HOST_TYPE` according to the local host's architecture. Include a routine in your shell startup file that determines the hardware/software architecture of the local host, and sets `CCASE_HOST_TYPE` to one of the file-name extension strings: `hpux9`, `sunos5`, and so on. Here is a code fragment from a C shell startup script:

```
set ARCHSTRING = "`uname -s ; uname -r`"
switch ("${ARCHSTRING}")
  case "Solaris 5*":
    setenv CCASE_HOST_TYPE solaris5
    breaksw
  case "HP-UX 9*":
    setenv CCASE_HOST_TYPE hpux9
    breaksw
  ...
```

9.5 Preventing Parallel Builds of Targets

When **clearmake** builds a makefile target, there may be side effects you cannot address in a makefile. For example, one of your build tools may create temporary files that aren't guaranteed to have unique names, and then delete them at the end of its processing. When you use this tool serially, there are no problems. However, if you invoke this build tool in multiple parallel builds in **clearmake**, the tool may create identical files and cause the builds to interfere with each other.

You can solve this problem by using the **.NOTPARALLEL** special makefile target. To disable parallel building for a makefile, use this target without any arguments. For example:

.NOTPARALLEL:

To prevent specific targets from being built in parallel with each other, specify them as a set of arguments. Note that parallel builds are prevented only within the set of targets. For example:

```
.NOTPARALLEL: %.a
.NOTPARALLEL: foo.c bar.c
```

In this example, **clearmake** does not build any **.a** file in parallel with any other **.a** file, and **foo** is not built in parallel with **bar**. However, **clearmake** can build **.a** files in parallel with **foo**, **bar**, or any other file.

9.6 Preventing Exponential Invocations of **abe**

If **clearmake** is invoked recursively during a parallel build, the result may be more invocations of **abe** than you want or than the build servers can handle. To prevent too many **abe** invocations, use the **.NOTPARALLEL** special makefile target for high-level invocations of **clearmake**.

Building Software for Multiple Platforms

10

This chapter addresses the challenge of using a single source tree to develop an application for a variety of hardware/software platforms. We discuss various approaches, contrasting their advantages and disadvantages. An extended example incorporates some of the approaches.

10.1 Issues in Multiple-Platform Development

The following issues arise in an environment where developers are creating and maintaining several platform-specific variants of an application:

- Different source code is required for different variants. Different UNIX-based operating systems may use different functions to implement the same task (for example, `strchr(3)` vs. `index(3)`). Likewise, it may be necessary to include different header files for different variants (for example, `string.h` vs. `strings.h`).
- Differing variants and various platforms may each have differing requirements. The differences may involve such particulars as compiler locations, compiler options, and libraries.
- Builds for different variants must be kept separate. Because there is a single source tree, care must be taken to ensure that object modules and executables for one architecture are not confused with those for other architectures. For example, the link editor must not try to create an executable for one architecture using an object module that was built for another architecture.

Additional issues must be addressed if ClearCase does not run on one of the target platforms. See Chapter 11, *Setting Up a Build on a Non-ClearCase Host* for a discussion of one such issue.

10.2 Handling Source Code Differences

We recommend that you use the same files (that is, the same versions of file elements) in all builds, for all platforms. You can usually achieve this goal using the standard UNIX approach: conditional compilation using the C preprocessor, **cpp(1)**. If header file **string.h** is to be used for the architecture whose **cpp** symbol is **ARCH_A**, and header file **strings.h** is to be used for architecture **ARCH_B**, use this code:

```
#ifdef ARCH_A
#include <string.h>
#else
#ifdef ARCH_B
#include <strings.h>
#endif /* ARCH_B */
#endif /* ARCH_A */
```

If a file element cannot be compiled conditionally (for example, a bitmap image), the traditional solution is to put architecture-specific code in different elements (for example, **panel.image.sparc** vs. **panel.image.mc68k**). This approach requires that build scripts be made architecture specific, too.

With ClearCase, you also have the option of splitting the element into branches. The **ARCH_A** variant can be developed on the element's **/main/arch_a** branch; edits and builds for that variant are done in a view configured with this rule:

```
element * /main/arch_a/LATEST
```

Other variants are developed on similarly named branches, each using a different view, configured with a rule like the one above. In such a situation, the element's **main** branch may not be used at all.

We recommend that you use this branching strategy only when necessary, because of these disadvantages:

- Each time platform-independent code is changed on one of the branches, you may need to merge the change to all the other branches.

- Developers must create a view for each architecture. In each view, only one variant of the application can be built.

Instead of branching, you can organize your code into architecture-specific subdirectories or architecture-specific VOBs.

10.3 Handling Build Procedure Differences

Ideally, a single file (that is, a single version of a file element) drives all architecture-specific builds. One way to accomplish this is to revise *makefiles* as follows:

- Regularize build scripts
- Replace architecture-specific constructs (for example, `/bin/cc`) with make-macro invocations (for example, a `$(CC)`)
- Use `clearmake`'s include directive to incorporate architecture-specific settings of the make macros

For example, suppose that source file `main.c` is compiled differently for two different architectures:

```
main.o:
    /usr/ucb/cc -c -fsingle main.c
```

```
main.o:
    /usr/bin/cc -c main.c
```

To merge these build scripts, abstract the compiler pathname and options into make-macros, `CC` and `CFLAGS`, and place an architecture-specific include line at the beginning of the makefile:

```
include /usr/project/make_macros/$(BLD_ARCH)_macros
..
main.o:
    $(CC) -c $(CFLAGS) main.c
```

The files in the `make_macros` directory then have these contents:

```

CC      = /usr/5bin/cc                               /usr/project/make_macros/sun4_macros
CFLAGS= -fsingle

CC      = /usr/bin/cc                               /usr/project/make_macros/irix5_macros
CFLAGS=

```

The make-macro **BLD_ARCH** acts as a selector between these two files. The value of this macro can be placed in an environment variable by a shell startup script:

```
setenv BLD_ARCH `uname -s`
```

Alternatively, developers can specify the value at build time:

```
clearmake main BLD_ARCH="HP-UX10"
```

Alternative Approach, Using imake

NOTE: The **imake** utility is distributed with many UNIX variants and available for free from the MIT Consortium.

The **imake** utility provides an alternative to the scheme described in the preceding section. The **imake** methodology also involves architecture-specific make-macros, but in a different way. **imake** generates an architecture-specific *makefile* by running **cpp** on an architecture-independent template file, typically named **imakefile**.

A typical **imakefile** contains a series of **cpp** macros, each of which expands to a build target line and its corresponding multiline build script. Typically, the expansion itself is architecture independent:

```

MakeObjectFromSrc(main)                               (macro in 'imakefile')
main.o: $(SRC)/main.c                                 (expansion in actual makefile)
    $(CC) -c $(CFLAGS) $(SRC)/main.c

```

imake places architecture-specific make-macro settings at the top of the generated makefile. For example:

```

SRC = ..
CC  = /usr/5bin/cc
CFLAGS= -fsingle
RM  = rm -f

```

An idiosyncrasy of **imake** is that makefiles are derived objects, not source files. The architecture-independent template file (**imakefile**) is the source file, and must be maintained as a ClearCase element.

10.4 Segregating the Derived Objects of Different Variants

It is essential to keep derived objects (object modules, executables) built for different architectures separate. This section describes two approaches, though others are possible.

Approach 1: Use Architecture-Specific Subdirectories

Each variant of an application can be built in its own subdirectory of the source directory. For example, if executable **monet**'s source files are located in the directory **/usr/monet/src**, the variants can be built in subdirectories **/usr/monet/src/sun4**, **/usr/monet/src/irix5**, and so on. It is simplest to have the *makefile* create view-private subdirectories for this purpose. But if you want to use different derived object storage pools for the different variants, you must create the subdirectories as elements (**mkdir** command) and then adjust their storage pool assignments (**chpool** command).

Because the derived objects for the different variants are built at different pathnames (for example, **/usr/monet/src/sun4/main.o**), they are segregated by variant, and **clearmake** never winks in an object built for another architecture.

This approach has several advantages:

- All variants of the application can be built in a single view.
- You do not need to consider whether to suppress winkin for some or all targets.
- Because the derived objects for different variants have different pathnames, it is easier to organize multiple-architecture releases.

But this approach may require build script changes: the binaries for a build are no longer in the source directory, but in a subdirectory. Note that the build script in *Alternative Approach, Using imake* on page 166 is structured for this situation:

```
main.o: $(SRC)/main.c
.
```

Approach 2: Use Different Views

Perform builds for different platforms in different views (**sun4_bld_vu**, **irix_bld_vu**, and so on). A team of developers working on the same variant can share a view, or each can work in an architecture-specific view.

In most cases, the build script that creates a derived object varies from variant to variant, as described in *Handling Build Procedure Differences* on page 165. If so, **clearmake** prevents linking of derived objects built for another architecture. If this is not the case, force the build script to be architecture specific by including a well-chosen message or comment. For example, if **BLD_ARCH** is used as described in *Handling Build Procedure Differences* on page 165, you can include this message:

```
@echo "Building $@ for $(BLD_ARCH) "
```

The disadvantage of this approach is that when an element is checked out, you can build only one variant of the application. Because the checked-out version is visible only in one view, builds of other variants (which take place in other views) do not select the checked-out version. You must check in the element before building other variants.

Another disadvantage is the number of views that may be required. For instance, if seven developers want to maintain their own views in which to build four variants, 28 views are required.

10.5 Multiple-Architecture Example

This section presents an example of multiple-architecture development. This example uses **imake** to support building in architecture-specific subdirectories.

Scenario

This section shows how to set up multiple-platform development in the `/proj/monet/src` directory.

A developer can perform a build for a particular architecture as follows:

1. She logs in to a machine of the desired architecture, for example, a workstation running Solaris 2.5.
2. In her regular view, she goes to the source directory, `/proj/monet/src`, and enters a command to have **imake** generate a *makefile*.
3. She enters the command **clearmake Makefiles** to have **imake** create the appropriate **Makefile** in the architecture-specific subdirectory **sun5**. Note that the **Makefile** is a derived object, not a source file. Thus, there is no need to create an element from this file.
4. She goes to the **sun5** subdirectory, and builds software for that architecture, using **clearmake**.

The sections that follow describe how **imake** is involved in each of these steps.

Defining Architecture-Specific CPP Macros

Step #1 places you in an environment where the C preprocessor, **cpp**, defines one or more architecture-specific symbols. On a SunOS-4 host, **cpp** defines the symbols **sun** and **sparc**. This, in turn, causes **imake** to generate many architecture-specific (machine-dependent) **cpp** macros:

```
sun defined by C #ifdef sun
preprocessor    #undef sun
                #define SunArchitecture
                #ifdef mc68020
                .
                .
                .
                #endif
sparc defined by C #ifdef sparc
preprocessor    #undef SUN4
                #undef sun4
imake defines   #define MachineDep SUN4
longer symbols  #define machinedep sun4
                #endif
```

Additional **cpp** macros specific to Sun are read in from the auxiliary file **sun.cf**.

Creating Makefiles in the Source and Build Directories

A file named **Imakefile** in the source directory is the **imake** input file. This file controls the creation of *makefiles* both in the source directory itself, and in the architecture-specific subdirectories where software is built:

```
#ifndef InMachineDepSubdir
.  
<code to generate makefile in source directory>
.  
#else
.  
<code to generate makefile in an architecture-specific subdirectory>
.  
#endif
```

The **Imakefile** code used in the source directory defines a symbol to record the fact that builds do not take place in this directory:

```
#define IHaveMachineDepSubdirs
```

The **Makefile** generated by **imake** includes a Makefiles target that populates an architecture-specific subdirectory with its own *makefile*:

```
Makefiles::
    @echo "Making Makefiles in $(CURRENT_DIR)/$$CPU"
    -@if [ ! -d $$CPU ]; then \
        mkdir $$CPU; \
        chmod g+w $$CPU; \
    else exit 0; fi
    @$(IMAKE_CMD) -s $$CPU/Makefile \
    -DInMachineDepSubdir \
    -DTOPDIR=$(TOP) -DCURDIR=$(CURRENT_DIR)/$$CPU
```

CPU environment variable determines name of architecture-specific subdirectory

The command **clearmake Makefiles** invokes **imake** again, using the same **Imakefile** for input. This time, the symbol **InMachineDepSubdir** is defined, which causes the actual build code to be generated.

The **Imakefile** in **/proj/monet/src** contains these macros:

```
OBJS = cmd.o main.o opt.o prs.o
LOCAL_LIBRARIES = ../../lib/libpub/libpub.a
```

```
MakeObjectFromSrc(cmd)
MakeObjectFromSrc(main)
MakeObjectFromSrc(opt)
MakeObjectFromSrc(prs)
```

```
ComplexProgramTarget(monet)
```

The resulting **Makefile** generated in the build directory, **/proj/monet/src/sun4**, includes this build script:

```
$(AOUT): $(OBJS) $(LOCAL_LIBRARIES)
    @echo "linking $@"
    -@if [ ! -w $@ ]; then $(RM) $@; else exit 0; fi
    $(CC) -o $@ $(OBJS) $(LOCAL_LIBRARIES) \
        $(LDFLAGS) $(EXTRA_LOAD_FLAGS)
```


Setting Up a Build on a Non-ClearCase Host

11

This chapter describes a technique that creates configuration records for a build that involves ClearCase data, but is performed on a non-ClearCase host. *Non-ClearCase access* (exporting a VOB through a view) makes the data available to that host; a remote shell is invoked to perform the build on that host.

11.1 Build Scenario

Suppose you want to build library **libpub.a** for an architecture that ClearCase does not currently support, using a host of that architecture named **titan**. The VOB storage area for the library's sources is located at **/vobstore/libpub.vbs** on host **sol**. This VOB is also mounted on **sol**, at **/proj/libpub**.

11.2 Setting Up an Export View

A ClearCase *export view* allows limited access to one or more VOBs by using standard NFS export facilities. Each NFS export provides remote access to one VOB through a particular view. The following limitations apply when you use an export view to access a VOB:

- ▶ You cannot check out or check in versions in the VOB. If you need to check out a file, you must perform a remote login to a ClearCase host.

- ▶ Builds in the view do not perform build auditing, configuration lookup, or winkin. The builds do not create derived objects or config records. Any files created during the build are view-private objects.

NOTE: You can use remote-shell techniques to overcome this limitation, so that the files built on a non-ClearCase host become derived objects. See *Revising the Build Script*.

- ▶ You cannot reconfigure the view from the non-ClearCase host. If you need to reconfigure the view, you must perform a remote login to a ClearCase host.

NOTE: If you modify an export view's config spec, all users who may currently have the view mounted for non-ClearCase access must unmount and remount the view. Remounting the view ensures access to the correct set of files as specified in the updated config spec.

For information on setting up views and VOBs for export, see *Setting Up ClearCase Views* in *Administering ClearCase*.

NOTE: Export views are to be used only for non-ClearCase access to VOBs. To make a view accessible on a remote host, use the **startview** or **setview** command on that host. An export view can be mounted on a ClearCase host, but never try to mount it on the *dynamic-views* root directory, **/view**.

11.3 Mounting the VOB Through the Export View

On the non-ClearCase host, a standard NFS mount is performed on the exported pathname. For example, mount **/view/libpub_expvu/proj/libpub** at **/proj/libpub** (the same location at which the VOB is mounted on ClearCase hosts).

11.4 Revising the Build Script

To produce an audited build on a non-ClearCase host, you must revise the build script. Thus, it makes sense to build in an architecture-specific subdirectory, with a customized *Makefile*. (For more information, see Chapter 10, *Building Software for Multiple Platforms*.)

To create a CR that lists all of the build's input files and output files, the build script executed by **clearmake** must do the following:

- Declare all input files as explicit dependencies. Because the MVFS does not run on the non-ClearCase host, source dependencies are not detected.
- Invoke a remote shell to perform the build on the non-ClearCase host.
- If the build performed by the remote shell succeeds, perform a **touch(1)** of all output files from the ClearCase host. This converts the view-private files created by the remote shell command to derived objects.

A simple build script can be transformed as follows:

Native Build

```
OBJS = data.o errmsg.o getcwd.o lineseq.o
```

```
data.o:
```

(source dependencies need not be declared)

```
    cc -c data.c
```

```
    .
```

```
    .
```

```
    .
```

```
libpub.a: $(OBJS)
```

```
    ar -rc $@ $(OBJS)
```

(other object modules produced similarly)

Non-ClearCase Build

```
OBJS = data.o errmsg.o getcwd.o lineseq.o
```

```
data.o: data.c libpub.h
```

(must declare source dependencies)

```
    rm -f $@
```

```
    rsh titan 'cd /proj/libpub ; cc -c data.c'
```

```
    if [ -w $@ ]; then \
```

```
    touch $@ ; \
```

```
    fi
```

```
    .
```

```
    .
```

```
    .
```

```
libpub.a: $(OBJS)
```

```
    rm -f $@
```

```
    rsh titan 'cd /proj/libpub ; ar -rc $@ $(OBJS)'
```

```
    if [ -w $@ ]; then \
```

```
    touch $@ ; \
```

```
    fi
```

(other object modules produced similarly)

The remote shell command (**rsh** in the example above) varies from system to system.

The remote shell program typically exits with status zero, even if the compilation fails. Thus, you must use some other technique for checking the success of the build, after the remote shell returns. In this example, the build scripts assume that the remote build is successful if the target file exists and is writable.

11.5 Performing an Audited Build in the Export View

To perform the desired build:

1. Register and set the export view on your workstation, which is a ClearCase host:

```
% cleartool mktag -tag libpub_expvu /public/export.vws
% cleartool setview libpub_expvu
```

2. Build in the normal way, on your host:

```
% cd /proj/libpub
% clearmake
```

The script listed above specifies a particular non-ClearCase host, **titan**, on which remote shells are to be executed. If there is more than one non-ClearCase host on which builds are to be performed, you must generalize this script.

NOTE: Because the remote host name is part of the build script, `winkin` of derived objects built on the various hosts fails, unless you make further modifications (for example, using `clearmake -O` to disable build-script checking).

Index

A

- abe process, use for parallel builds** 156
- AIX make compatibility mode** 95
- ar, use of u key** 49
- archives**
 - format in makefile 76
 - incremental update example 46
- attache-home-dir directory** xvii
- attributes, attaching to versions in CR** 71
- auditing, *See* build auditing**

B

- bldserver.control file** 159
- BOS files**
 - about 74
 - clearmake read order 75
 - format of contents 86
 - recommended use 74
 - special targets for 81
- branch strategy for multiple platforms** 164
- build auditing**
 - about 3
 - effect of background processes 45
 - in export views 176
 - including non-MVFS files 36
 - incremental updates and 46
 - multiple levels, problems 44
 - without clearmake 7
- build avoidance**
 - about 5
 - algorithms for make and clearmake 47
 - Cfront-based compilers 108
 - differences in clearmake and make 32
 - multiple build scripts for target 34
 - scheme for in make 36
 - template instantiation method 104
- build environment**
 - for clearmake and make 30
 - views used 1

build hosts

- client setup for parallel builds 157
- idleness thresholds, setting 160
- non-ClearCase 10
- non-ClearCase, setting up 173
- parallel builds 158
- server setup 159

build options specification, *See* BOS files

build scheme in ClearCase 2

build scripts

- DO-IDs in 58
- format in makefile 73
- multiple for single target 34
- non-ClearCase hosts 174
- parallel builds 156
- temporary changes to 32
- when omitted from CRs 17

builds

- See also* express builds; hierarchical builds; parallel builds
- DOs and performance 7
- forced, problems with 39
- javac behavior 150
- labeling versions created in 43
- reference time 16
- starting 29
- subsessions 43
- verbosity levels, increasing 32
- working while in progress 40

builds for multiple platforms

- about 163
- branching strategy 164
- development issues 163
- DOs and subdirectories 167
- DOs and views 168
- example with imake 168
- handling source code differences 164
- makefiles for 165

built-in rules in makefiles 77

C

C++ templates

- about 104
- alternative instantiation procedures 105
- recommended method of instantiation 104

catcr command

- DO versions 64
- sample listing 14

CCASE_BLD_UMASK environment variable 53

CCASE_HOST_TYPE environment variable 160

CCASE_OPTS_SPECS environment variable 75

CCASE_VERBOSE environment variable 32

ccase-home-dir directory xvii

Centerline C++ compiler, *See* Cfront-based compilers

Cfront-based compilers

- about 106
- Forced Instantiation model 119
- how link-time template instantiation works 107
- interaction with clearmake 106
- Multiple Repositories model 112
- Simple instantiation model 109

clearaudit

- about 7
- contents omitted from CR 17
- coordinating multiple builds 44
- multiple log files, workarounds 44
- use with make programs 89

clearmake

- build scenario 30
- compatibility modes 8, 89
- declaring dependencies in makefiles 36
- double-colon rules 39
- format of makefiles 73
- increasing verbosity level for builds 32
- interaction with SGI Delta/C++ compiler 137
- interaction with XLC compiler 139
- interactions with Cfront-based compilers 106
- interactions with SPARCompiler C++ 4.x 123
- internal macros 79
- invoking 29
- Java behavior 147
- link-time template instantiation problems 108
- macro substitution 78
- recursive invocation 34
- standard input as makefile 75
- symptoms of template problems for C++ 104

clock skew

- about 41
- time rules and 42

commands, control of echoing during build 76

compatibility modes in clearmake

- about 89
- adjusting levels of 8

- AIX make 95
- Gnu make 96
- IRIX pmake 92
- IRIX smake 90
- SunOS 5.x 92
- System V Release 3 make 102

config specs, time rules in 41

configuration lookup

- about 5
- common outcomes 6
- in hierarchical builds 6
- problems with dependencies 37
- VPATH macro 80

conventions, typographical xvii

CRs (configuration records)

- about 4
- attaching labels and attributes to versions in 71
- comparing 70
- contents of 13
- contents of, effect of background processes 45
- displaying contents of 70
- displaying for DO versions 64
- double-colon rules and contents of 39
- effect on DOs when unavailable 70
- example 14
- hierarchy of 17
- hierarchy, and winkin 58
- hierarchy, processing by cleartool commands 19
- how created 45
- incremental updates and 46
- Java builds 147
- merging dependencies in 48
- recording makefile version 71
- storage of 20

D

.DEFAULT target 81

dependencies

- build order, in makefiles 39
- declaring in makefiles 36
- detected, log of 6
- format in makefiles 73
- merging in CRs 48
- problems when searching directories for 37
- tracking 4
- tracking non-MVFS files 4

.DEPENDENCY_IGNORED_FOR_REUSE target 82

describe command 56

diffcr command 70

DO versions

- about 25
- access to 63
- as release mechanism 65

- creating 62
- creating in builds 62
- displaying configuration records 64
- displaying description of 56
- documentation**
 - online help description xviii
- DO-IDs**
 - about 12
 - displaying 56
 - in build scripts 58
 - in cleartool commands 57
 - vs. OID 56
- DOs (derived objects)**
 - See also* nonshareable DOs; shareable DOs
 - about 4, 11
 - attaching labels and attributes to sources in CR 71
 - build avoidance role 5
 - converting to view-private files 61
 - costs of creating 7
 - criteria for reuse or winkin 5
 - degenerate 69
 - disk space usage, displaying 68
 - displaying kind of 55
 - effect of forced builds 39
 - incremental updating 46
 - incremental updating by Cfront-based compilers 108
 - incremental updating, example 46
 - incremental updating, links in 49
 - incremental updating, scenarios 49
 - kinds of 20
 - listing at specific pathnames 55
 - listing views that reference 57
 - multiple platform, separating 167
 - overwriting 68
 - removing 68
 - scrubbing 69
 - selecting versions for in view 71
 - siblings of 4
 - siblings of, types 21
 - specifying in commands 57
 - storage 21
 - when created 31
- dospace command** 68
- double-colon rules, how clearmake interprets** 39

E

- environment variables**
 - CCASE_BLD_UMASK 53
 - CCASE_HOST_TYPE 160
 - CCASE_OPTS_SPECS 75
 - CCASE_VERBOSITY 32
 - order of precedence in makefiles 77

- error handling, control of in makefile** 76
- export views**
 - about 173
 - auditing builds in 176
 - mounting VOBs 174
- express builds**
 - about 7
 - creating views for 60
 - reconfiguring views for 60
 - when to use 60
 - winkin to 59

F

- function names in makefiles** 76

G

- Gnu make compatibility mode** 96

H

- hard links**
 - VOB, creating to DOs 67
 - when winked in 26
- header files, precompiled** 105
- hierarchical builds**
 - configuration lookup in 6
 - reference time of 16
 - use of 44
- HP aC++ compiler**
 - about 145
 - template instantiation models 145
- HP C++ compiler**
 - See* Cfront-based compilers

I

- .IGNORE target** 81
- imake utility**
 - about 166
 - build example 168
- include files in makefiles** 77
- .INCREMENTAL_REPOSITORY_SIBLING target** 82
- .INCREMENTAL_TARGET target** 82
- IRIX pmake compatibility mode** 92
- IRIX smake compatibility mode** 90

J

Java compilers

- benefits of using make with 148
- build problems 147
- building in clearmake or omake 149
- configuring makefiles 153
- makefiles for 150
- rebuilding targets 152

L

labels, attaching to versions in CR 71

libraries, format in makefile 76

lsdo command 57

- examples 55

lsprivate command 55

M

macros

- internal clearmake 79
- order of precedence in makefiles 77
- substitution by clearmake 78
- target-dependent definitions 74

\$(MAKE) macro, defining for clearmake 35

make

- about 1
- build avoidance scheme 36
- use with Java 148

make macros

- format in makefile 74
- format of definition 78
- temporary overrides of 32

makefiles

- about 73
- builds for multiple platforms 165
- built-in rules 77
- contents for snapshot views 85
- controlling execution of 76
- declaring dependencies in 36
- design models for Cfront-based compilers 109
- double-colon rules and clearmake 39
- format for clearmake 73
- format of libraries 76
- function names in 76
- imakefiles 166
- include files in 77
- incremental updating and 47
- Java compilers 150
- javac, using with 149
- non-MVFS dependencies and 4
- order of precedence, macros and environment variables 77

- overriding build scripts in 32
- special targets 80
- standard input as, in clearmake 75
- version of in CR 71

mounting VOBs, in export views 174

MVFS

- template instantiation outside 105

MVFS files

- about 4
- in configuration records 16

N

.NO_CMP_NON_MF_DEPS target 83

.NO_CMP_SCRIPT target 83

.NO_CONFIG_REC target 83

.NO_DO_FOR_SIBLING target 83

.NO_WINK_IN target 84

non-ClearCase hosts

- about 10
- setup for build 173

non-MVFS files

- as dependencies, tracking 4
- in configuration records 16

nonshareable DOs

- about 7, 21
- automatic conversion to shareable 66
- converting to shareable 66
- promotion and winkin 22
- storage 21
- types of siblings 21
- unique DO-IDs for 25

.NOTPARALLEL target 84

- uses of 162

O

OIDs

- how used 15

online help, accessing xviii

order of precedence in makefiles 77

P

parallel builds

- about 8, 155
- client setup 157
- how clearmake works 9
- how controlled 155

- preventing parallel, of targets 162
- server setup 159
- starting 160
- pathnames**
 - accessing DOs 57
 - and DO-IDs 12
 - for DO versions 63
 - form in build scripts 74
 - MVFS, in makefiles for snapshot views 85
- .PRECIOUS target** 81
- pseudotargets, and winkin** 35

R

- rebuilding targets, See builds**
- reference count**
 - about 26
 - hard links and 26
 - when zero 69
- reference time**
 - about 16
 - effect on source control 40
 - for multiple builds 44
- release areas, structure and management** 65
- rmdo command** 68

S

- scrubbing DOs** 69
- SGI Delta/C++ compiler 4.x**
 - about 136
 - interactions with clearmake 137
 - template instantiation models 137
- SGI OCC compiler, See Cfront-based compilers**
- shareable DOs**
 - about 20
 - components of 21
 - converting to nonshareable 66
 - in views reconfigured for express builds 60
 - permissions to share 53
 - promotion and winkin 22
 - removing data containers 68
 - storage 21
 - types of siblings 21
 - unique DO-IDs for 25
- shell**
 - auditing build in 7
- shells**
 - setting CCASE_HOST_TYPE in 161
- .SIBLING_IGNORED_FOR_REUSE target** 84

- siblings of DOs**
 - about 4
 - shareable and nonshareable 21
- .SIBLINGS_AFFECT_REUSE target** 85
- .SILENT target** 81
- snapshot views, pathnames in makefiles** 85
- SPARCompiler C++ 2.x, See Cfront-based compilers**
- SPARCompiler C++ 4.x**
 - about 123
 - interaction with clearmake 123
 - makefile example 132
 - Multiple Repositories model 129
 - repository cleanup 125
 - repository setup 124
 - Simple model 126
- subsessions in builds** 43
- subtargets in makefiles** 39
- SunOS 5.x compatibility mode** 92
- System V Release 3 make compatibility mode** 102

T

- targets**
 - build rules and clearmake macros 79
 - format in makefiles 73
 - multiple build scripts for 34
 - preventing parallel builds 162
 - rebuilding by Java compilers 152
 - recursive invocation of clearmake 34
 - special 80
 - special, format in makefile 74
 - special, lists of 81
- technical support** xix
- time rules**
 - effect of clock skew 42
 - use in config specs 41
- time stamps, adding to C-language executables** 50
- typographical conventions** xvii

U

- umask setting to share DOs** 53

V

- version strings, adding to C-language executables** 50
- versions**
 - checked-out, how clearmake handles 31
 - created in builds, labeling 43
 - of DOs 25

view-extended pathnames for DOs 57

views

- configuring for express builds 60
- configuring to select versions for DO 71
- for builds 1
- preventing winkin to and from 59
- references to DOs, listing 57
- snapshot, contents of makefiles for 85
- time rules and 41
- using in builds for multiple platforms 168

VPATH macro 80

W

what string, creating 50

winkin

- about 6
- criteria for 5
- from other platforms, preventing 61
- hard links and 26
- javac behavior 150
- manual 58
- permissions for 54
- preventing 59
- pseudotargets and 35
- recursive, uses of 58
- reference count and 26

X

XLC compiler

- about 139
- Compile-Time Demand Instantiation model 142
- Explicit Instantiation model 144
- interaction with clearmake 139
- Simple model 140