MANAGING SOFTWARE PROJECTS WITH CLEARCASE

Release 4.1 and later

UNIX Edition

Rational®
the e-development company™

800-023557-000

ClearCase and MultiSite Release Notes
Document Number 800-023557-000 August 2000
Rational Software Corporation 20 Maguire Road Lexington, Massachusetts 02421

IMPORTANT NOTICE

Copyright Notice

Copyright © 1992, 2000 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, ClearDDTS, 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.

This product includes software developed by Greg Stein <gstein@lyra.org> for use in the mod_dav module for Apache (http://www.webdav.org/mod_dav/).

Rational®
the e-development company™

Contents

Preface xvii
About This Manual xvii
Organization xvii
ClearCase Documentation Roadmap xviii
Typographical Conventions xix
Online Documentation xx
Technical Support xx

1. Choosing Between UCM and Base ClearCase 1
1.1 Differences Between UCM and Base ClearCase 1
Branching 2
Creating and Using Baselines 3
Managing Activities 4
Enforcing Development Policies 5
1.2 Using Base ClearCase Tools with UCM 5

Part 1: Working in UCM

2. Understanding UCM 9
2.1 The Project Management Cycle 9
2.2 Creating the Project 12
Creating a PVOB 12
Organizing Directories and Files into Components 13
Shared and Private Work Areas 13
Starting from a Baseline 14
Setting Policies 14
Setting Up the UCM-ClearQuest Integration 15
2.3 Integrating Work into the Project (MultiSite) 16
2.4 Making a New Baseline 17

2.5 Promoting the Baseline 18
2.6 Overview of the UCM-ClearQuest Integration 20
Associating UCM and ClearQuest Objects 20
UCM-Enabled Schema 21
State Types 21
Queries in a UCM-Enabled ClearQuest Schema 22

3. Planning the Project 23
3.1 Using the System Architecture as the Starting Point 23
Mapping System Architecture to Components 24
Deciding What to Place Under Version Control 24
Mapping Components to Projects 25
Size of the System 25
Amount of Integration 25
Need for Parallel Releases 25
Example 26
Components and VOBs 26
3.2 Organizing Components 27
Considering VOB Capacity 27
Identifying Additional Components 28
Defining the Directory Structure 28
Identifying Read-Only Components 29
3.3 Specifying a Baseline Strategy 30
When to Create Baselines 30
Identifying the Initial Baseline 31
Ongoing Baselines 31
Defining a Naming Convention 32
Identifying Promotion Levels to Reflect State of Development 32
Planning How to Test Baselines 32
3.4 Planning PVOBs 33
Deciding How Many PVOBs to Use 33
Understanding the Role of the Administrative VOB 34
3.5 Identifying Special Element Types 35

Nonmerging Elements 35
Nonautomerging Elements 36
Defining the Scope of Element Types 36
3.6 Planning How to Use the UCM-ClearQuest Integration 36
Mapping PVOBs to ClearQuest User Databases 37
All Enabled Projects in a PVOB Must Link to the Same Database 37
Projects Linked to Same Database Must Have Unique Names 37
Use One Schema Repository for Linked Databases 38
Deciding Which Schema to Use 38
Overview of the UnifiedChangeManagement Schema 39
Enabling a Schema for UCM 40
3.7 Considering Which Development Policies to Enforce 41
Policies Available in UCM 41
Recommended Baselines 41
Modifiable Components 41
Default View Types 41
Rebase Before Deliver 42
Allow Deliveries from Stream with Pending Checkouts 42
Policies Available in UCM-ClearQuest Integration 43
Check Before Work On 43
Check Before ClearCase Delivery 43
Do ClearQuest Action After Delivery 43

4. Setting Up a ClearQuest User Database 45
4.1 Using the Predefined UCM-Enabled Schemas 45
4.2 Enabling a Schema to Work with UCM 46
Requirements for Enabling Custom Record Types 49
Setting State Types 49
State Transition Default Action Requirements for Record Types 50
4.3 Customizing ClearQuest Project Policies 52
4.4 Associating Child Activity Records with a Parent Activity Record 52
Using Parent/Child Controls 53
4.5 Creating Users 53

4.6 Setting the Environment 53

5. Setting Up the Project 55
5.1 Creating a Project from Scratch 56
Creating the Project VOB (PVOB) 56
Creating Components 57
Creating the Project 58
Defining Promotion Levels 59
Creating an Integration View 59
Creating and Setting an Activity 61
Creating the Directory Structure 61
Importing Directories and Files from Outside ClearCase 62
5.2 Creating a Project Based on an Existing ClearCase Configuration 63
Creating the PVOB 63
Making a VOB into a Component 63
Making a Baseline from a Label 64
Creating the Project 64
Creating an Integration View 64
5.3 Creating a Project Based on an Existing Project 65
Reusing Existing PVOB and Components 65
Creating the Project 65
Creating an Integration View 66
5.4 Enabling a Project to Use the UCM-ClearQuest Integration 66
Migrating Activities 67
Setting Project Policies 68
Assigning Activities 69
Disabling the Link Between a Project and a ClearQuest User Database 69
Fixing Projects That Contain Linked and Unlinked Activities 70
Detecting the Problem 70
Correcting the Problem 70
5.5 Creating a Development Stream for Testing Baselines 71

6. Managing the Project 73
6.1 Adding Components 73

Updating Snapshot View Load Rules 75
6.2 Integrating the Project 75
Finding WorkThat is Ready to Be Delivered 76
Completing Remote Deliver Operations 76
Undoing a Deliver Operation 77
6.3 Creating a New Baseline 77
Locking the Integration Stream 77
Verifying That the Code Base Is Stable 78
Making the New Baseline 78
Making a Baseline For a Set of Activities 80
Unlocking the Integration Stream 80
6.4 Testing the Baseline 80
Fixing Problems 81
6.5 Promoting or Demoting the Baseline 82
6.6 Tracking the Project 82
Comparing Baselines 83
Querying ClearQuest User Databases 84
6.7 Cleaning Up the Project 85
Removing Unused Objects 85
Projects 86
Streams 86
Components 86
Baselines 86
Activities 87
Locking and Making Obsolete the Project and Streams 87

7. Managing Parallel Releases of Multiple Projects 89
7.1 Managing a Current Project and a Follow-on Project Simultaneously 89
Example 90
Performing Interproject Rebase Operations 91
7.2 Incorporating a Patch Release into a New Version of the Project 92
Example 92
Merging Work to Another Project 94
7.3 Additional Merging Scenarios 95

Merging from a Project to a Non-UCM Branch 95
Merging to a System Project....95

Part 2: Working in Base ClearCase

8. Managing Projects in Base ClearCase 99
8.1 Setting Up the Project....100
Creating and Populating VOBs 100
Planning a Branching Strategy 100
Branch Names 101
Branches and ClearCase MultiSite....101
Creating Shared Views and Standard Config Specs 102
Recommendations for View Names 102
8.2 Implementing Development Policies 102
Using Labels 103
Using Attributes, Hyperlinks, Triggers, and Locks....103
Global Types....104
Generating Reports....104
8.3 Integrating Changes 105

9. Defining Project Views....107
9.1 How Config Specs Work 107
9.2 Default Config Spec....108
The Standard Configuration Rules 108
Omitting the Standard Configuration Rules 109
9.3 Config Spec Include Files 109
9.4 Project Environment for Sample Config Specs 110
9.5 Views for Project Development....111
View for New Development on a Branch 111
Variation That Uses a Time Rule....112
View to Modify an Old Configuration 112
Omitting the /main/LATEST Rule 113
Variation That Uses a Time Rule....114

View to Implement Multiple-Level Branching 114
View to Restrict Changes to a Single Directory 115
9.6 Views to Monitor Project Status 116
View That Uses Attributes to Select Versions 116
Pitfalls of Using This Configuration for Development 118
View That Shows Changes of One Developer 119
Historical View Defined by a Version Label 119
Historical View Defined by a Time Rule 120
9.7 Views for Project Builds 120
View That Uses Results of a Nightly Build 121
Variations That Select Versions of Project Libraries 122
View That Selects Versions of Application Subsystems 122
View That Selects Versions That Built a Particular Program 123
Configuring the Makefile 123
Fixing Bugs in the Program 124
Selecting Versions That Built a Set of Programs 124
9.8 Sharing Config Specs Between UNIX and Windows 125
Pathname Separators 125
Pathnames in Config Spec Element Rules 126
Config Spec Compilation 126
Example 126

10. Implementing Project Development Policies 129
10.1 Policy: Good Documentation of Changes Is Required 129
10.2 Policy: All Source Files Require a Progress Indicator 130
10.3 Policy: Label All Versions Used in Key Configurations 131
10.4 Policy: Isolate Work on Release Bugs to a Branch 132
10.5 Policy: Avoid Disrupting the Work of Other Developers 133
10.6 Policy: Deny Access to Project Data When Necessary 134
10.7 Policy: Notify Team Members of Relevant Changes 134
10.8 Policy: All Source Files Must Meet Project Standards 135
10.9 Policy: Associate Changes with Change Orders 136
10.10 Policy: Associate Project Requirements with Source Files 137
10.11 Policy: Prevent Use of Certain Commands 139

10.12 Policy: Certain Branches Are Shared Among MultiSite Sites 140
10.13 Sharing Triggers Between UNIX and Windows 141
Using Different Pathnames or Different Scripts 141
Using the Same Script 142
Notes 142

11. Integrating Changes 143
11.1 How Merging Works 143
Using the GUI to Merge Elements 145
Using the Command Line to Merge Elements 146
11.2 Common Merge Scenarios 146
Scenario: Selective Merge from a Subbranch 147
Scenario: Removing the Contributions of Some Versions 148
Scenario: Merging All Project Work 149
All Project Work Is Isolated on a Branch 149
All Project Work Isolated In a View 149
Scenario: Merging a New Release of an Entire Source Tree 150
Scenario: Merging Directory Versions 153
11.3 Using Your Own Merge Tools 154

12. Using Element Types to Customize Processing of File Elements 155
12.1 File Types in a Typical Project 155
12.2 How ClearCase Assigns Element Types 156
12.3 Element Types and Type Managers 157
Other Applications of Element Types 159
Using Element Types to Configure a View 159
Processing Files by Element Type 160
12.4 Predefined and User-Defined Element Types 160
12.5 Predefined and User-Defined Type Managers 161
12.6 Type Manager for Manual Page Source Files 161
Creating the Type Manager Directory 161
Inheriting Methods from Another Type Manager 162
The create_version Method 162

The construct_version Method 164
Implementing a New compare Method 166
Testing the Type Manager 167
Installing and Using the Type Manager 168
12.7 Icon Use by GUI Browsers 169

13. Using ClearCase Throughout the Development Cycle 173
13.1 Project Overview 173
13.2 Development Strategy 175
Project Manager and ClearCase Administrator 175
Use of Branches 175
Creating Project Views 178
13.3 Creating Branch Types 178
13.4 Creating Standard Config Specs 179
13.5 Creating, Configuring, and Registering Views 179
13.6 Development Begins 180
Techniques for Isolating Your Work 180
13.7 Creating Baseline 1 181
Merging Two Branches 181
Integration and Test 182
Labeling Sources 182
Removing the Integration View 183
13.8 Merging Ongoing Development Work 183
Preparing to Merge 184
Merging Work 186
13.9 Creating Baseline 2 187
Merging from the r1_fix Branch 188
Preparing to Merge from the major Branch 188
Merging from the major Branch 190
Decommissioning the major Branch 191
Integration and Test 191
13.10 Final Validation: Creating Release 2.0 191
Labeling Sources 192
Restricting Use of the main Branch 192

Setting Up the Test View 193
Setting Up the Trigger to Monitor Bugfixing 193
Fixing a Final Bug 194
Rebuilding from Labels 194
Wrapping Up 195

A. ClearCase-ClearQuest Integrations 197
A.1 Understanding the Two ClearCase-ClearQuest Integrations 197
Managing Coexisting Integrations 198
Schema 198
Presentation 198

Index 201

Figures

Figure 1 Branching Hierarchy in Base ClearCase 2
Figure 2 Branching Hierarchy Under UCM Streams 3
Figure 3 Project Management and Development Cycles in UCM 11
Figure 4 Baselines of Two Components 14
Figure 5 Rebase Operation 18
Figure 6 Promoting Baselines 19
Figure 7 Association of UCM and ClearQuest Objects in Integration 20
Figure 8 Components Used by Transaction Builder Project 26
Figure 9 Mapping Components to Projects 27
Figure 10 Using a Read-Only Component 30
Figure 11 Related Projects Sharing One PVOB 34
Figure 12 Projects in Multiple PVOBs Linked to the Same ClearQuest Database 37
Figure 13 Using the Same Schema Repository for Multiple ClearQuest Databases 38
Figure 14 UCM Tab of Record Form for a UCM-Enabled Record Type 39
Figure 15 Main Tab of Record Form for the BaseCMActivity Record Type 40
Figure 16 Associating a User Database with a UCM-Enabled Schema 46
Figure 17 Adding the UCMPolicyScripts Package to a Schema 47
Figure 18 Assigning State Types to a Record Type's States 48
Figure 19 Navigating to Record Type's State Transition Matrix 48
Figure 20 State Transition Diagram for UCM-enabled BaseCMActivity Record Type 51
Figure 21 Navigating to Integration Stream in Project Explorer 60
Figure 22 Step 2 of New Project Wizard 66
Figure 23 Enabling a Project to Work with a ClearQuest User Database 68
Figure 24 Navigating to the UCMProjects Query 69
Figure 25 Add Baseline Dialog Box 74
Figure 26 Find Posted Deliveries Dialog Box 76
Figure 27 Make Baseline Dialog Box 79
Figure 28 Comparing Baselines by Activity 83
Figure 29 Comparing Baselines by Version 84

Figure 30 Managing a Follow-on Release 90
Figure 31 Incorporating a Patch Release 93
Figure 32 Making a Change to an Old Version 113
Figure 33 Multiple-Level Auto-Make-Branch 115
Figure 34 Development Config Spec vs. QA Config Spec 117
Figure 35 Checking Out a Branch of an Element 118
Figure 36 Requirements Tracing....139
Figure 37 Versions Involved in a Typical Merge 144
Figure 38 ClearCase Merge Algorithm....145
Figure 39 Selective Merge from a Subbranch 147
Figure 40 Removing the Contributions of Some Versions 148
Figure 41 Merging a New Release of an Entire Source Tree 151
Figure 42 Data Handling: File Type, Element Type, Type Manager 158
Figure 43 User-Defined Icon Display 171
Figure 44 Project Plan for Release 2.0 Development 174
Figure 45 Development Milestones: Evolution of a Typical Element....177
Figure 46 Creating Baseline 1....181
Figure 47 Updating Major Enhancements Development 184
Figure 48 Merging Baseline 1 Changes into the major Branch 186
Figure 49 Baseline 2....188
Figure 50 Element Structure after the Pre-Baseline-2 Merge 190
Figure 51 Final Test and Release 191
Figure 52 Change Sets in ClearQuest GUI....199

Tables

Table 1 Recommended Directory Structure for Components 28
Table 2 State Types in UCM-Enabled Schema 50
Table 3 Environment Variables Required for Integration 54
Table 4 Queries in UCM-Enabled Schema 84
Table 5 Files Used in a Typical Project 156

Preface

ClearCase, a *configuration management* system, is designed to help software development teams track the objects used in software builds. You can use base ClearCase to create a customized configuration management environment, or you can adopt the Unified Change Management (UCM) process. UCM is an out-of-the-box process, layered on base ClearCase and ClearQuest functionality, for organizing software development teams and their work products.

About This Manual

This manual shows project managers how to set up and manage a configuration management environment for their development team using either UCM or the customizable features of base ClearCase.

Organization

The manual is divided into two parts:

- *Part 1: Working in UCM*. Read this part if you plan to use UCM to implement your team's development process.
- *Part 2: Working in Base ClearCase*. Read this part if you plan to use the base ClearCase features to implement a customized development process for your team.

ClearCase Documentation Roadmap

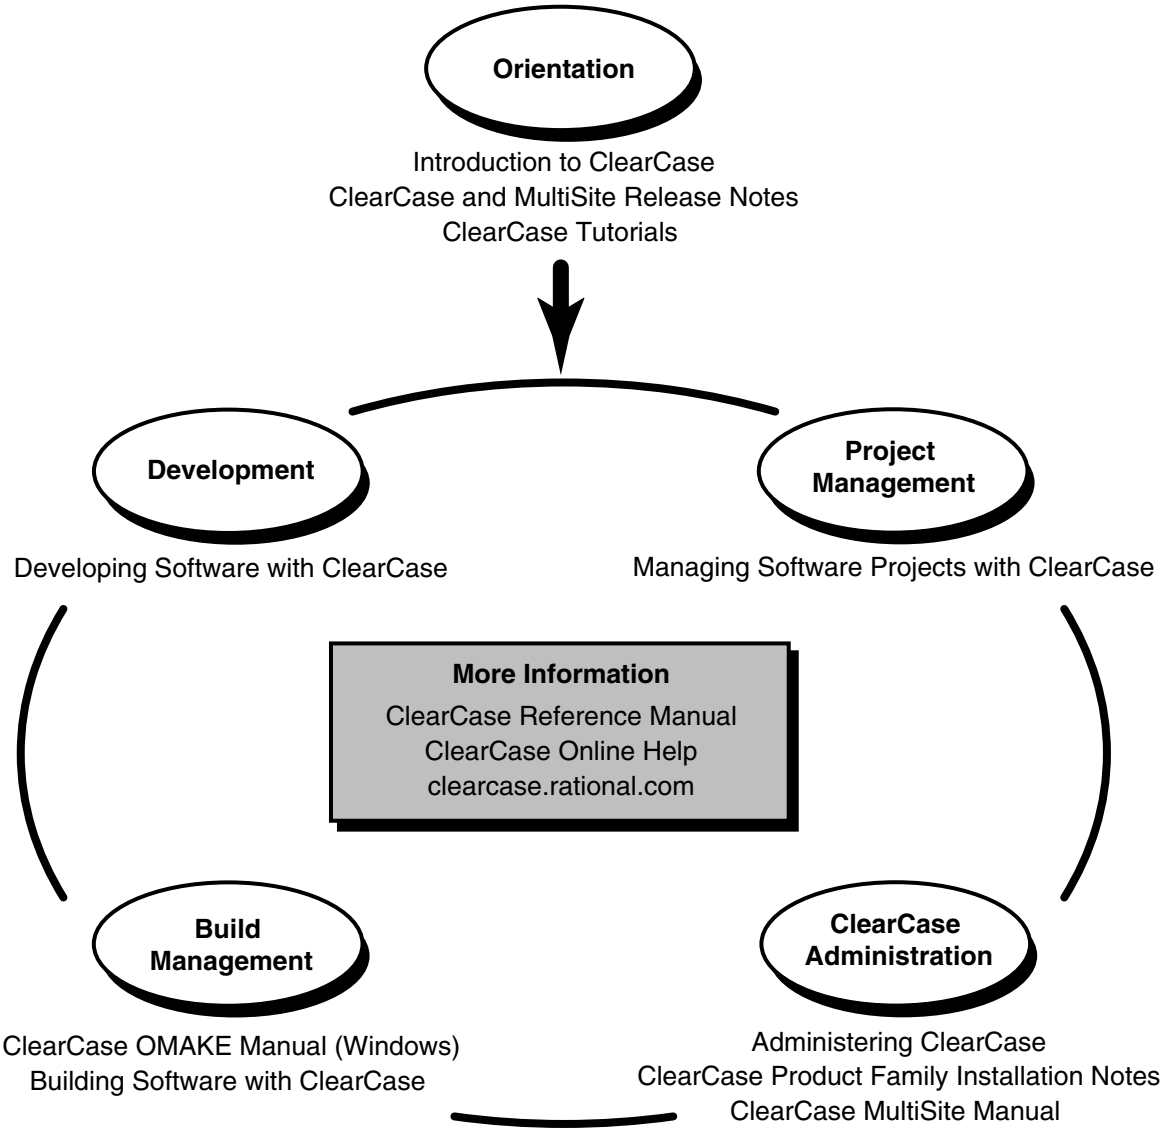

Typographical Conventions

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** on UNIX and **C:\Program Files\Rational\ClearCase** on Windows.
- *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: 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. To start the *ClearCase Tutorial* from the command line, type **hyperhelp cc_tut.hlp**.

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**

1 Choosing Between UCM and Base ClearCase

Before you can start to use ClearCase to manage the version control and configuration needs of your development project, you need to decide whether to use the out-of-the-box Unified Change Management (UCM) process or base ClearCase. This chapter describes the main differences between the two methods from the project management perspective.

The rest of this manual is organized into two parts. Part 1 describes how to manage a project using UCM. Part 2 describes how to manage a project using the various tools in base ClearCase.

1.1 Differences Between UCM and Base ClearCase

Base ClearCase consists of a set of powerful tools to establish an environment in which developers can work in parallel on a shared set of files, and project managers can define policies that govern how developers work together.

UCM is one prescribed method of using ClearCase for version control and configuration management. UCM is layered on base ClearCase. Therefore, it is possible to work efficiently in UCM without having to master the details of base ClearCase.

UCM offers the convenience of an out-of-the-box solution; base ClearCase offers the flexibility to implement virtually any configuration management solution that you deem appropriate for your environment.

Branching

Base ClearCase uses branches to enable parallel development. A *branch* is an object that specifies a linear sequence of versions of an element. Every element has one *main branch*, which represents the principal line of development, and may have multiple *subbranches*, each of which represents a separate line of development. For example, a project team may use the main branch for new development work while using a subbranch simultaneously for fixing a bug.

Subbranches can have subbranches. For example, a project team may designate a subbranch for porting a product to a different platform. The team may then decide to create a bug-fixing subbranch off that porting subbranch. Base ClearCase allows you to create complex branch hierarchies. Figure 1 illustrates a multilevel branch hierarchy. As a project manager in such an environment, you need to ensure that developers are working on the correct branches. To do that, you must tell them which rules to include in their *config specs* so that their views access the appropriate set of versions.

Figure 1 Branching Hierarchy in Base ClearCase

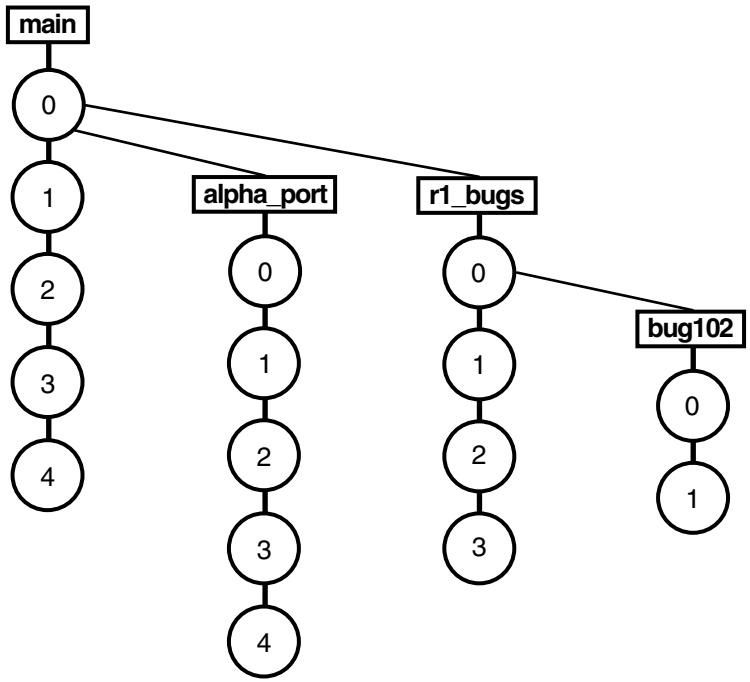

UCM uses branches also, but you do not have to manipulate them directly because it layers streams over the branches. A *stream* is a ClearCase object that maintains a list of activities and determines which versions of elements appear in a developer's view. In UCM, a project contains one *integration stream*, which records the project's shared set of elements, and multiple *development streams,* in which developers work on their parts of the project in isolation from the team. UCM does not allow for complex branch hierarchies. The project's integration stream uses one branch. Each development stream uses its own branch, which is a subbranch of the integration stream's branch. Development stream branches cannot have subbranches. Figure 2 illustrates the simple branching hierarchy that supports UCM streams.

As project manager of a UCM project, you need not write rules for config specs. Streams configure developers' views to access the appropriate versions on the appropriate branches.

Figure 2 Branching Hierarchy Under UCM Streams

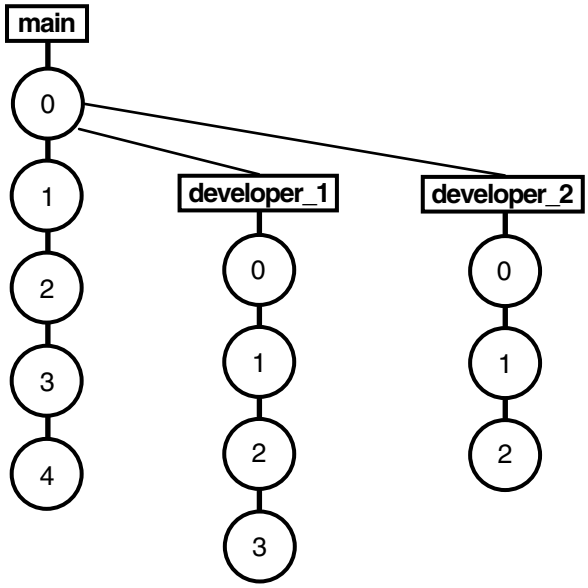

Creating and Using Baselines

Both base ClearCase and UCM allow you to create baselines. UCM automates the creation process and provides additional support for performing operations on baselines. A *baseline* identifies the set of versions of files that represent a project at a particular milestone. For example, you may create a baseline called **beta1** to identify an early snapshot of a project's source files.

Baselines provide two main benefits:

- The ability to reproduce an earlier release of a software project
- The ability to tie together the complete set of files related to a project, such as source files, a product requirements document, a documentation plan, functional and design specifications, and test plans

In base ClearCase, you can create a baseline by creating a version label and applying that label to a set of versions.

In UCM, baseline support appears throughout the user interface because UCM requires that you use baselines. When developers join a project, they must first populate their work areas with the contents of the project's recommended baseline. This method ensures that all team members start with the same set of shared files. In addition, UCM lets you set a property on the baseline to indicate the quality level of the versions that the baseline represents. Examples of quality levels include "project builds without errors," "passes initial testing," and "passes regression testing." By changing the quality-level property of a baseline to reflect a higher degree of stability, you can, in effect, promote the baseline.

Managing Activities

In base ClearCase, you work at the version and file level. UCM provides a higher level of abstraction: activities. An *activity* is a ClearCase object that you use to record the work required to complete a development task. For example, an activity may be to change a graphical user interface (GUI). You may need to edit several files to make the changes. UCM records the set of versions that you create to complete the activity in a *change set*. Because activities appear throughout the UCM user interface, you can perform operations on sets of related versions by identifying activities rather than having to identify numerous versions.

Because activities correspond to significant project tasks, you can track the progress of a project more easily. For example, you can determine which activities were completed in which baselines. If you use the UCM-ClearQuest integration, you gain additional project management control, such as the ability to assign states and state transitions to activities. You can then generate reports by issuing queries such as "show me all activities assigned to Pat that are in the Ready state."

Enforcing Development Policies

A key part of managing the configuration management aspect of a software project is establishing and enforcing development policies. In a parallel development environment, it is crucial to establish rules that govern how team members access and update shared sets of files. Such policies are helpful in two ways:

- They minimize project build problems by identifying conflicting changes made by multiple developers as early as possible.
- They establish greater communication among team members.

These are examples of common development policies:

- Developers must synchronize their private work areas with the project's recommended baseline before delivering their work to the project's shared work area.

- Developers must notify other team members by e-mail when they deliver work to the project's shared work area.

In base ClearCase, you can use tools such as triggers and attributes to create mechanisms to enforce development policies. UCM includes a set of common development policies, which you can set through the GUI or command-line interface (CLI).

1.2 Using Base ClearCase Tools with UCM

This manual is organized into two parts: Part 1 for UCM and Part 2 for base ClearCase. If you are managing a UCM project, you may occasionally want to extend UCM by using some of the tools in base ClearCase. In particular, you may want to use ClearCase attributes, triggers, and hyperlinks to customize development policies.

Part 1: Working in UCM

The following chapters describe how to plan, set up, and manage a UCM project to implement your team's development process.

.

Understanding UCM

2

This chapter provides an overview of Unified Change Management (UCM). Specifically, it introduces the main UCM objects and describes the tasks involved in managing a UCM project. Subsequent chapters describe the detailed steps required to perform these tasks.

2.1 The Project Management Cycle

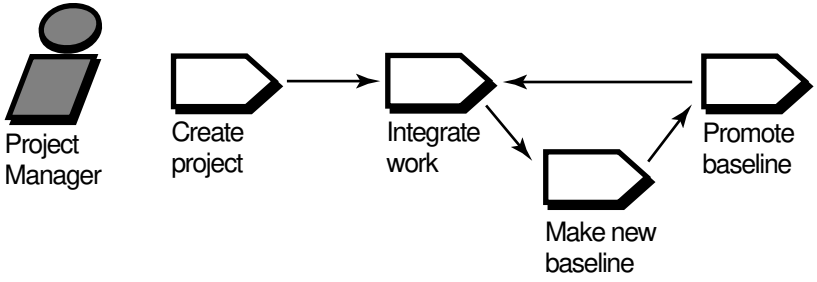

In UCM, your work follows a cycle that complements an iterative software development process. Members of a project team work in a UCM *project*. A project is the object that contains the configuration information needed to manage a significant development effort, such as a product release. A project contains one shared work area and typically multiple private work areas. Private work areas allow developers to work on activities in isolation. As project manager, you are responsible for maintaining the project's shared work area. Work within a project progresses as follows:

1. You create a project and identify an initial set of baselines of one or more components. A *component* is a group of related directory and file elements, which you develop, integrate, and release together. A *baseline* is a version of a component.

2. Developers join the project by creating their private work areas and populating them with the contents of the project's baselines.

3. Developers create activities and work on one activity at a time. An *activity* records the set of files that a developer creates or modifies to complete a development task, such as fixing a bug. This set of files associated with an activity is known as a *change set*.

4. When developers complete activities, and build and test their work in their private work areas, they share their work with the project team by performing *deliver* operations. A deliver operation merges work from the developer's private work area to the project's shared work area.

5. In the shared work area, you integrate the work delivered by developers.

6. Periodically, you create new baselines in the shared work area that incorporate the delivered work.

7. You perform quick validation tests to make sure that the new baselines build and appear to work correctly. A team of software quality engineers performs more extensive testing.

8. Periodically, as the quality and stability of baselines improve, you adjust the promotion level attribute of baselines to reflect appropriate milestones, such as Built, Tested, or Released. When the new baselines pass a sufficient level of testing, you designate them as the *recommended* set of baselines.

9. Developers perform *rebase* operations to update their private work areas to include the set of versions represented by the new recommended baselines.

10. Developers continue the cycle of working on activities, delivering completed activities, updating their private work areas with new baselines.

This list of UCM tasks can be seen as two cycles: project management and development. Figure 3 illustrates the connection between these cycles.

Figure 3 Project Management and Development Cycles in UCM

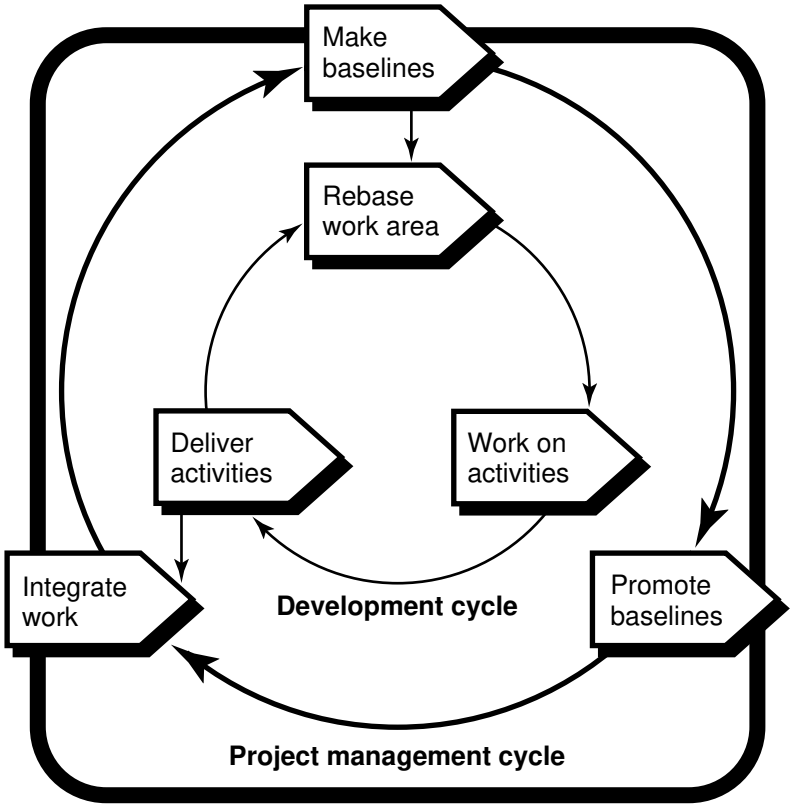

2.2 Creating the Project

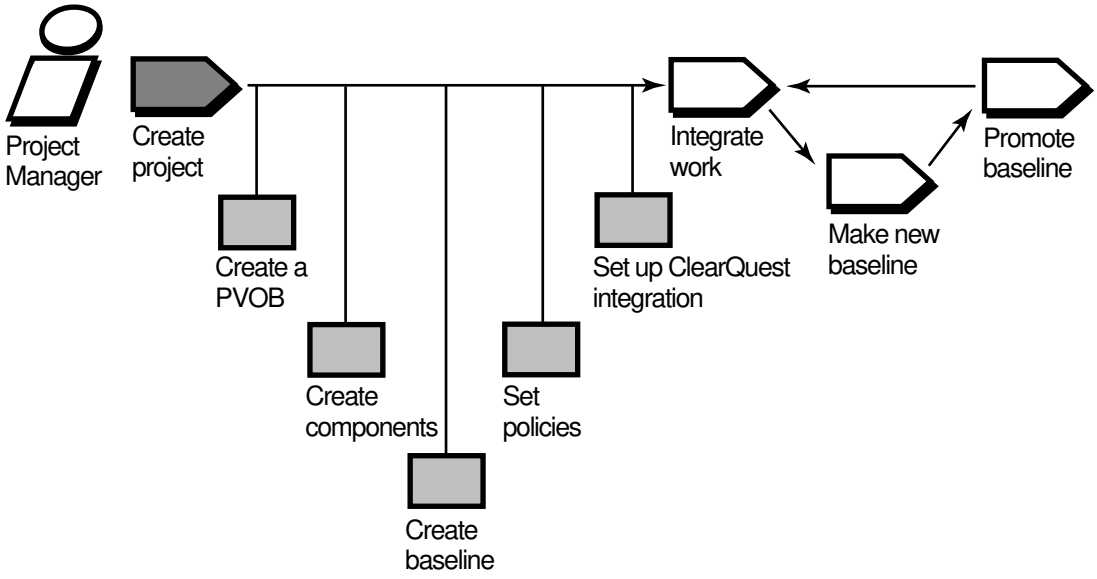

To create and set up a project, you must perform the following tasks:

- Create a repository for storing project information
- Create components that contain the set of files the developers work on
- Create baselines that identify the versions of files with which the developers start their work
- Select the development policies you want to enforce

Creating a PVOB

ClearCase stores file elements, directory elements, derived objects, and metadata in a repository called a versioned object base (*VOB*). In UCM, each project must have a project VOB (PVOB). A *PVOB* is a special kind of VOB that stores UCM objects, such as projects, activities, and change sets. A PVOB must exist before you can create a project. Check with your site's ClearCase administrator to see whether a PVOB has already been created. For details on creating a PVOB, see *Creating the Project VOB (PVOB)* on page 56.

Organizing Directories and Files into Components

As the number of files and directories in your system grows, you need a way to reduce the complexity of managing them. Components are the UCM mechanism for simplifying the organization of your files and directories. The elements that you group into a component typically implement a reusable piece of your system architecture. By organizing related files and directories into components, you can view your system as a small number of identifiable components, rather than one large set of directories and files.

Within a component, you organize directory and file elements into a directory tree. The component's root directory must be the root directory of a VOB. You can convert existing VOBs into components, or you can create a component from scratch. For details on creating a component from scratch, see *Creating Components* on page 57. For details on converting a VOB into a component, see *Making a VOB into a Component* on page 63.

Shared and Private Work Areas

A work area consists of a view and a stream. A *view* is a directory tree that shows a single version of each file in your project. A *stream* is a ClearCase object that maintains a list of activities and determines which versions of elements appear in your view.

A project contains one *integration stream*, which records the project's baselines and enables access to versions of the project's shared elements. The integration stream and a corresponding integration view represent the project's shared work area.

Each developer on the project has a private work area, which consists of a development stream and a corresponding development view. The *development stream* maintains a list of the developer's activities and determines which versions of elements appear in the developer's view.

When you create a project from the UCM GUI, ClearCase creates the integration stream for you. If you create a project from the command-line interface, you need to create the integration stream explicitly. Developers create their development streams and development views when they join the project. See *Developing Software with ClearCase* for information on joining a project.

Starting from a Baseline

After you create project components or select existing components, you must identify the baseline or baselines that serve as the starting point for the team's developers. A baseline identifies one version of every element visible in a component. Figure 4 shows baselines named BL1 and BL2 that identify versions in Component A and Component B, respectively.

The project's integration stream records the baselines. When developers join the project, they populate their work areas with the versions of directory and file elements represented by the baselines. This practice ensures that all members of the project team start with the same set of files.

Figure 4 Baselines of Two Components

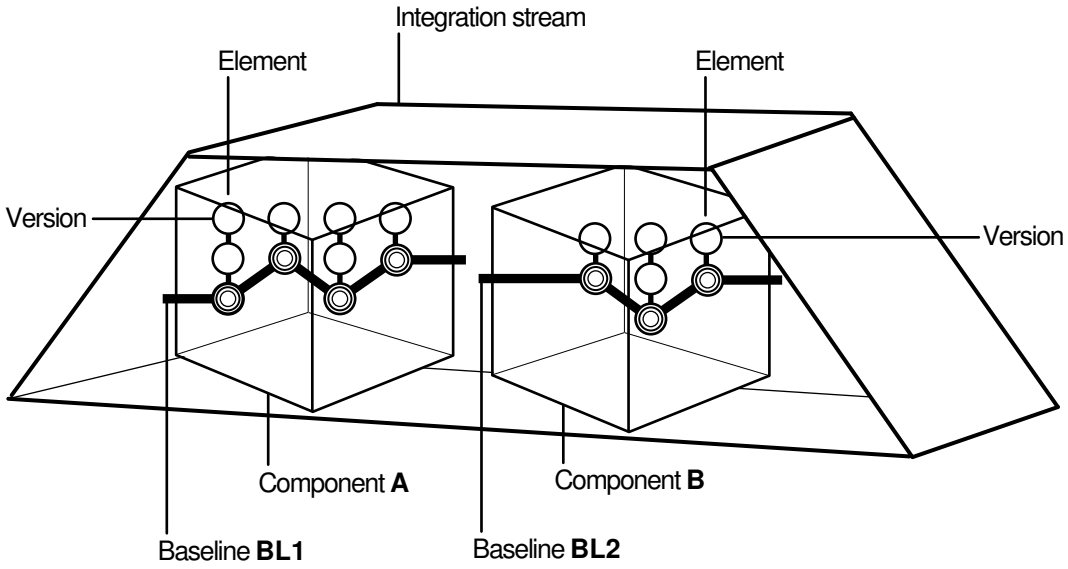

Setting Policies

UCM includes a set of policies that you can set to enforce development practices among members of the project team. By setting policies, you can improve communication among project team members and minimize the problems you may encounter when integrating their work. For example, you can set a policy that requires developers to update their work areas with the

project's latest recommended baseline before they deliver work to the integration stream. This practice reduces the likelihood that developers will need to work through complex merges when they deliver their work. For a description of all policies you can set in UCM, see *Considering Which Development Policies to Enforce* on page 41.

Setting Up the UCM-ClearQuest Integration

You can use UCM without Rational ClearQuest, the change request management tool. The integration with ClearQuest adds significant project management and activity management capabilities. When you set up a UCM project to work with ClearQuest, the integration links all project activities to ClearQuest records. You can then take advantage of UCM's state transition model and ClearQuest's query, reporting, and charting features. These features allow you to do the following:

- Assign activities to developers
- Assign states and state transition rules to activities
- Generate reports based on database queries
- Select additional development policies to be enforced

To set up the UCM-ClearQuest integration:

1. Enable a ClearQuest schema to work with UCM or use a predefined UCM-enabled schema.
2. Create or upgrade a ClearQuest user database to use the schema.
3. Enable your UCM project to work with ClearQuest.

See *Overview of the UCM-ClearQuest Integration* on page 20 for additional information about the integration.

2.3 Integrating Work into the Project (MultiSite)

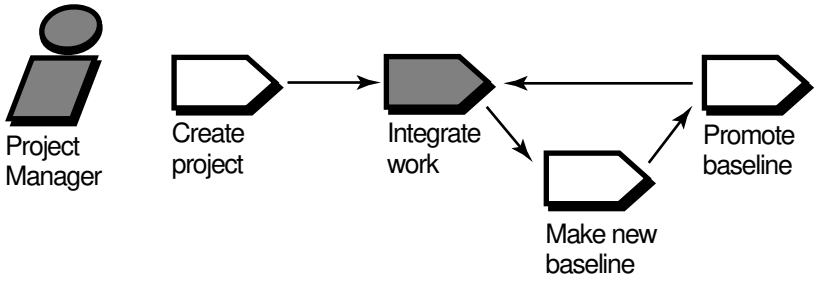

In most cases, developers complete the deliver operations that they start. If your project uses ClearCase MultiSite, you may need to complete some deliver operations. Many ClearCase customers use MultiSite, a product layered on ClearCase, to support parallel software development across geographically distributed project teams. MultiSite lets developers work on the same VOB concurrently at different locations. Each location works on its own copy of the VOB, known as a *replica*.

To avoid conflicts, MultiSite uses an exclusive-right-to-modify scheme, called *mastership*. VOB objects, such as streams and branches, are assigned a *master replica*. The master replica has the exclusive right to modify or delete these objects.

In a MultiSite configuration, a team of developers may work at a remote site, and the project's integration stream may be mastered at a different replica than the developers' development streams. In this situation, the developers cannot complete deliver operations to the integration stream. As project manager, you must complete these deliver operations. UCM provides a variation of the deliver operation called a *remote deliver*. When UCM detects a stream mastership situation, it makes the deliver operation a remote deliver, which starts the deliver operation but does not merge any versions. You then complete the deliver operation.

For details on completing remote deliver operations, see *Integrating the Project* on page 75.

2.4 Making a New Baseline

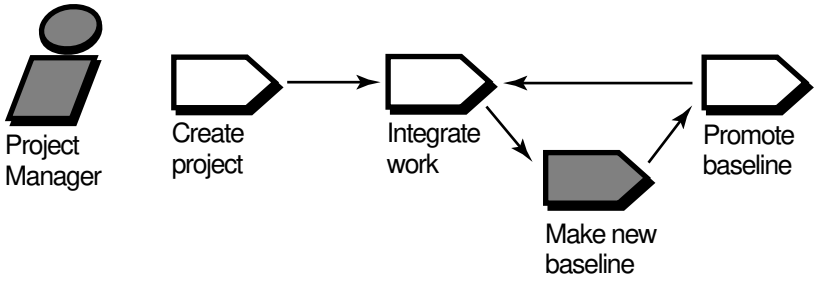

To ensure that developers stay in sync with each other's work, make new baselines regularly. A new baseline includes the work developers have delivered to the integration stream since the last baseline. To make a new baseline:

1. Lock the integration stream to prevent developers from delivering work while you create the baseline. Developers can continue to work on activities in their development streams.
2. Verify the stability of the project by building and testing its components.
3. Make the baseline.
4. Unlock the integration stream so that developers can resume delivering work.

After your team of software quality engineers test the new baseline more extensively and determine that it is stable, make the baseline the *recommended* baseline. Developers then update their work areas with the new baseline by performing a rebase operation, which merges files and directories from the integration stream to the development stream.

Figure 5 illustrates a rebase operation from baseline BL1 to BL2. For details on making baselines, see *Creating a New Baseline* on page 77.

Figure 5 Rebase Operation

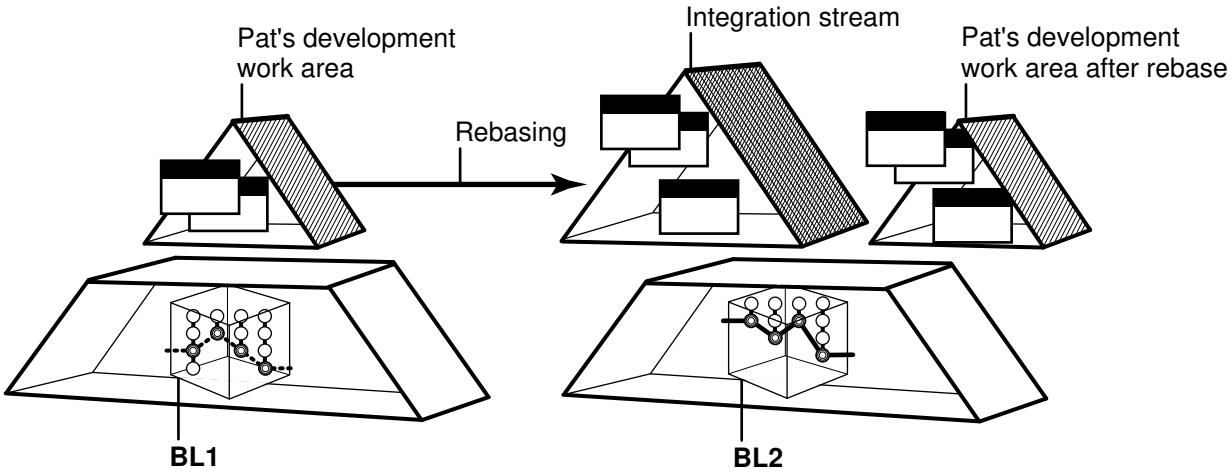

2.5 Promoting the Baseline

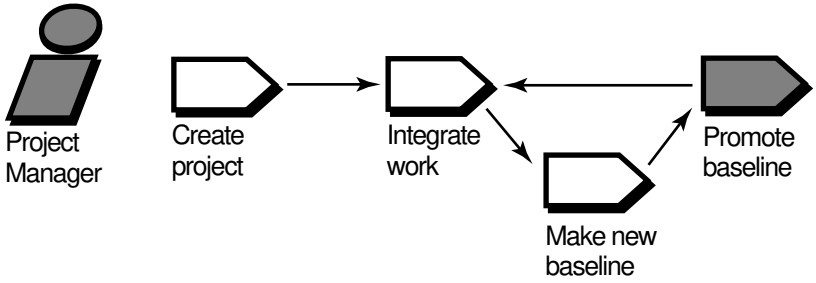

As work on your project progresses and the quality and stability of the components improve, change the baseline's promotion level attribute to reflect important milestones. The promotion level attribute typically indicates a level of testing. For example, Figure 6 shows the evolution of baselines through three levels of testing; the BL8 baseline is ready for production.

You can use promotion levels in development policies. For example, you can set a policy to make a baseline the recommended baseline when it reaches a particular promotion level, such as "tested." You can set another policy that requires developers to rebase their development streams to the set of recommended baselines before they deliver work. These policies help to

ensure that developers update their work areas whenever a baseline passes an acceptable level of testing.

For details on promoting baselines, see *Promoting or Demoting the Baseline* on page 82.

Figure 6 Promoting Baselines

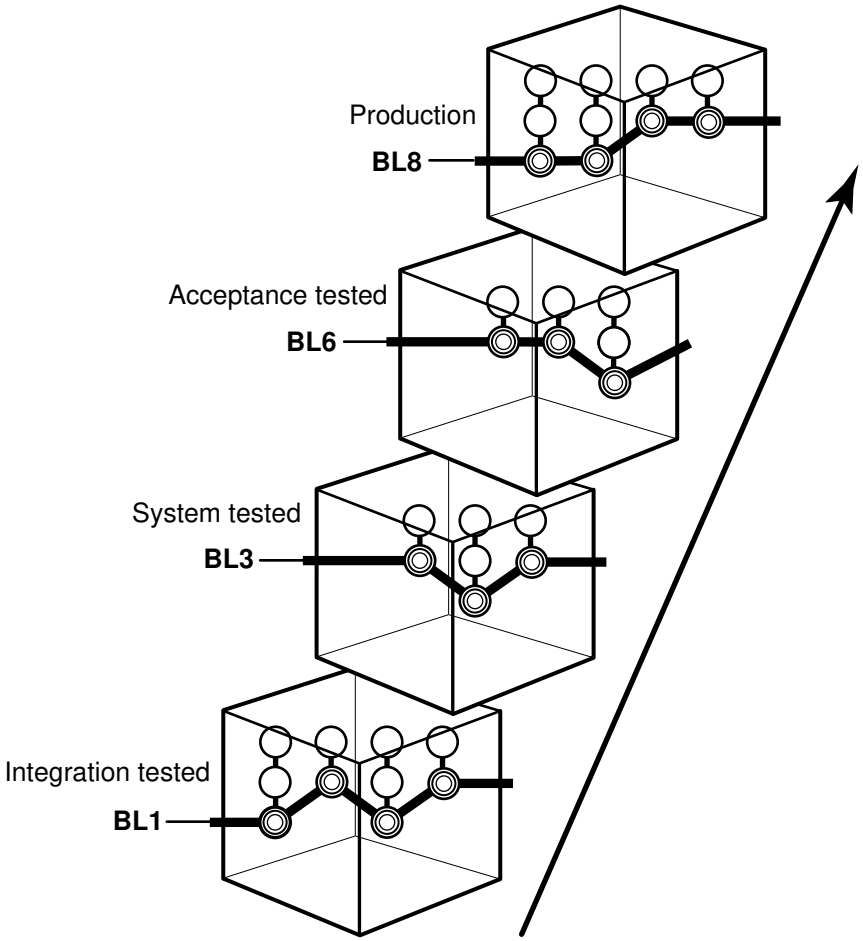

2.6 Overview of the UCM-ClearQuest Integration

This section describes the following UCM-ClearQuest integration concepts:

- Association of UCM and ClearQuest objects
- UCM-enabled schema
- Queries
- State types

Associating UCM and ClearQuest Objects

Setting up the integration links UCM and ClearQuest objects. Figure 7 shows the bidirectional linking of these objects.

Figure 7 Association of UCM and ClearQuest Objects in Integration

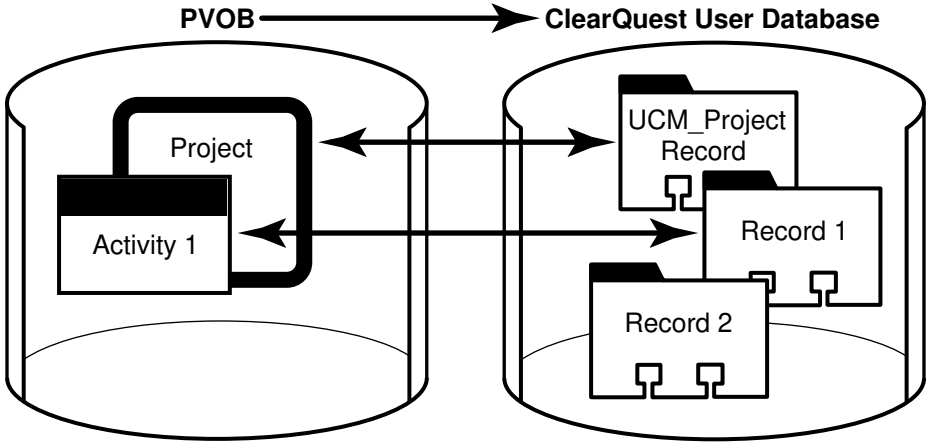

When you enable a project to link to a ClearQuest user database, the integration stores a reference to that database in the project's PVOB.

Every ClearQuest-enabled project is linked to a project record of record type **UCM_Project** in the ClearQuest user database.

Every activity in a ClearQuest-enabled project is linked to a record in the database. An activity's title is linked to the headline field in its corresponding ClearQuest record. If you change an

activity's title in ClearCase, the integration changes the headline in ClearQuest to match the new title, and the reverse is also true. Similarly, an activity's name is linked to the ID field in its ClearQuest record.

It is possible for a ClearQuest user database to contain some records that are linked to activities and some records that are not linked. Record 2 in Figure 7 is not linked to an activity. You may encounter this situation is if you have a ClearQuest user database in place before you adopt UCM. As you create activities, the integration creates corresponding ClearQuest records. However, any records that existed in that user database before you enabled it to work with UCM remain unlinked.

UCM-Enabled Schema

In ClearQuest, a schema is the definition of a database. To use the integration, you must create or upgrade a ClearQuest user database that is based on a UCM-enabled schema. A UCM-enabled schema contains certain fields, scripts, actions, and state types. ClearQuest includes two predefined UCM-enabled schemas, which you can use. You can also enable a custom schema or another predefined schema to work with UCM. For details on UCM-enabled schemas, see *Deciding Which Schema to Use* on page 38.

State Types

ClearQuest uses states to track the progress of change requests from submission to completion. A *state* represents a particular stage in this progression. Each movement from one state to another is a *state transition*. The integration uses a particular state transition model. To implement this model, the integration uses state types. A *state type* is a template that defines actions and other attributes of a state. You can define as many states as you want, but all states in a UCM-enabled record type must be based on one of the following state types:

- Waiting
- Ready
- Active
- Complete

For each state type, you can have multiple states. However, you must define at least one path of transitions between states of state types as follows: Waiting to Ready to Active to Complete. For details on state types, see *Setting State Types* on page 49.

Queries in a UCM-Enabled ClearQuest Schema

A UCM-enabled schema includes six queries. When you create or upgrade a ClearQuest user database to use a UCM-enabled schema, the integration installs these queries in two subfolders of the Public Queries folder in the user database's workspace. These queries make it easy for developers to see which activities are assigned to them and for project managers to see which activities are active in a particular project. For details on these queries, see *Querying ClearQuest User Databases* on page 84

3 Planning the Project

This chapter describes the issues you need to consider in planning to use one or more UCM projects as your configuration management environment. We strongly recommend that you write a configuration management plan before you begin creating projects and other UCM objects. After you create your plan, see Chapter 5, *Setting Up the Project* for information on how to implement it.

3.1 Using the System Architecture as the Starting Point

Essential to developing and maintaining high quality software is the definition of the system's architecture. The Rational Unified Process states that defining and using a system architecture is one of the six best practices to follow in developing software. A system architecture is the highest level concept of a system in its environment. The Rational Unified Process states that a system architecture encompasses the following:

- The significant decisions about the organization of a software system
- The selection of the structural elements and their interfaces of which the system is composed, together with their behavior as specified in the collaboration among those elements
- The composition of the structural and behavioral elements into progressively larger subsystems
- The architectural style that guides this organization, these elements, and their interfaces, their collaborations, and their composition

A well-documented system architecture improves the software development process. It is also the ideal starting point for defining the structure of your configuration management environment.

Mapping System Architecture to Components

Just as different types of blueprints represent different aspects of a building's architecture (floor plans, electrical wiring, plumbing, and so on), a good software system architecture contains different views to represent its different aspects. The Rational Unified Process defines an architectural view as a simplified description (an abstraction) of a system from a particular perspective or vantage point, covering particular concerns and omitting entities that are not relevant to this perspective.

The Rational Unified Process suggests using five architectural views. Of these, the implementation view is most important for configuration management. The implementation view identifies the physical files and directories that implement the system's logical packages, objects, or modules. For example, your system architecture may include a licensing module. The implementation view identifies the directories and files that make up the licensing module.

From the implementation view, you should be able to identify the set of UCM components you need for your system. Components are groups of related directory and file elements, which you develop, integrate, and release together. Large systems typically contain many components. A small system may contain one component.

Deciding What to Place Under Version Control

In deciding what to place under version control, do not limit yourself to source code files and directories. The power of configuration management is that you can record a history of your project as it evolves so that you can re-create the project quickly and easily at any point in time. To record a full picture of the project, include all files and directories connected with it. These include, but are not limited to the following:

- Source code files and directories
- Model files, such as Rational Rose files
- Libraries
- Executable files
- Interfaces
- Test scripts
- Project plans

- System and user documentation
- Requirements documents

Mapping Components to Projects

After mapping your system architecture to a set of components and identifying the full set of files and directories to place under version control, you need to determine whether to use one project or multiple projects. In general, think of a project as the configuration management environment for a project team. Team members work together to develop, integrate, test, and release a set of related components. For many systems, all work can be done in one project. For some systems, work must be separated into multiple projects. In deciding how many projects to use, consider the following factors:

- Size of the system
- Amount of integration required
- Whether you need to release multiple versions of the product concurrently

Size of the System

Consider the number of developers working on the system and the number of components. For a small system that consists of two or three components being developed by a dozen developers, one project probably makes sense. For a large system that consists of 20 components being developed by 100 developers, it may be wise to use several projects.

Amount of Integration

Determine the relationships between the various components. Related components that require a high degree of integration belong to the same project. By including related components in the same project, you can build and test them together frequently, thus avoiding the problems that can arise when you integrate components late in the development cycle.

Need for Parallel Releases

If you need to develop multiple versions of your system in parallel, consider using separate projects, one for each version. For example, your organization may need to work on a patch release and a new release at the same time. In this situation, both projects use mostly the same set of components. (Note that multiple projects can modify the same set of components.) When work on the patch release project is complete, you integrate it with the new release project.

Example

Figure 8 shows the initial set of components planned for the Transaction Builder system. A team of 30 developers work on the system. Because a high degree of integration between components is required, and most developers work on several components, the project manager included all components in one project.

Figure 8 Components Used by Transaction Builder Project

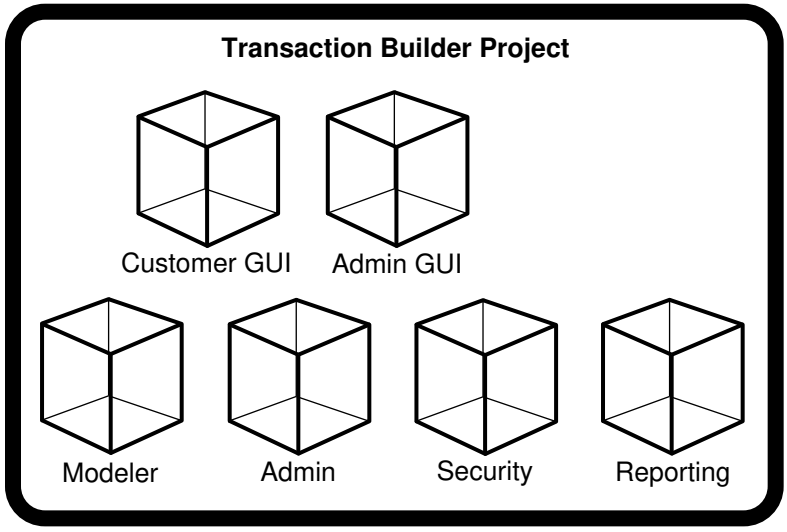

Components and VOBs

ClearCase implements components as *versioned object bases* (VOBs), the repositories for versions of file elements, directory elements, derived objects, and metadata associated with them. Figure 9 illustrates this implementation. In Figure 9, both Project1 and Project2 use Component2.

Figure 9 Mapping Components to Projects

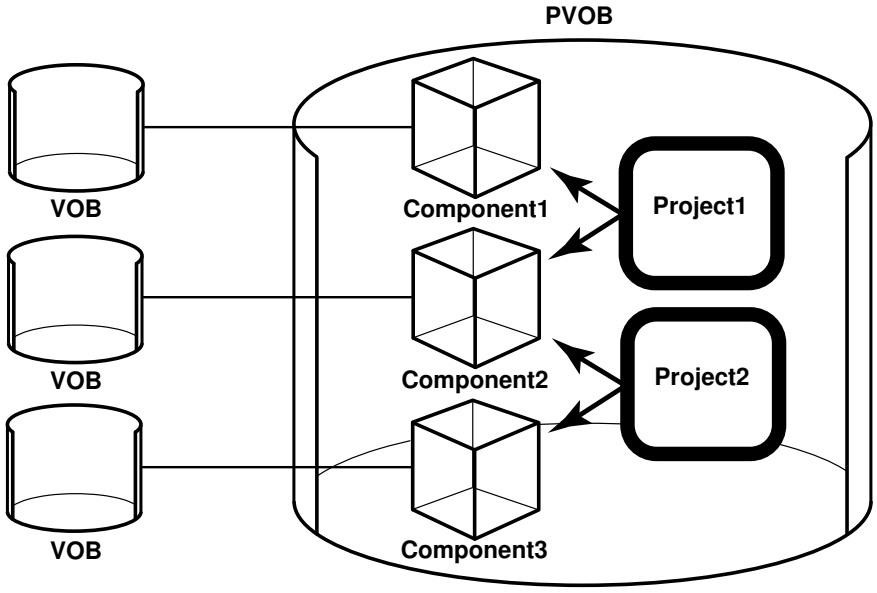

3.2 Organizing Components

After you map your system architecture to an initial set of components and determine which projects will access those components, refine your plan by performing the following tasks:

- Ensure that your components are a suitable size for VOBs
- Identify any additional components
- Define the component directory structures
- Identify read-only components
- Identify nonmerging elements

Considering VOB Capacity

Because ClearCase implements components as VOBs, you must ensure that the contents of each planned component do not exceed the capacity of a VOB. See *Administering ClearCase* for details on VOB capacity.

Identifying Additional Components

Although you should be able to identify nearly all necessary components by examining your system architecture, you may overlook a few. For example:

System component	It is a good idea to designate one component for storing system-level files. These items include project plans, requirements documents, and system model files and other architecture documents.
Testing component	Consider using a separate component for storing files related to testing the system. This component includes files such as test scripts, test results and logs, and test documentation.
Deployment component	At the end of a development cycle, you need a separate component to store the generated files that you plan to ship with the system or deploy inhouse. These files include executable files, libraries, interfaces, and user documentation.

Defining the Directory Structure

After you complete your list of components, you need to define the directory structures within those components. We recommend that you start with a directory structure similar to the one shown in Table 1; then modify the structure to suit your system's needs.

In Table 1, Component_1 through Component_n refers to the components that map to the set of logical packages in your system architecture.

Table 1 Recommended Directory Structure for Components

Component	Directories	Typical Contents
System	plans	Project plans, mission statement, and so on
	requirements	Requirements documents
	models	Rose files, other architecture documents
	documentation	System documentation

Table 1 Recommended Directory Structure for Components

Component	Directories	Typical Contents
Component_1 through Component_n	requirements	Component requirements
	models	Component model files
	source	Source files for this component
	interfaces	Component public interfaces
	binaries	Executable and other binary files for this component
	libraries	Libraries used by this component
	tests	Test scripts and related documents for this component
Test	scripts	Test scripts
	results	Test results and logs
	documentation	Test documentation
Deployment	binaries	Deployed executable files
	libraries	Deployed libraries
	interfaces	Deployed interfaces
	documentation	User documentation

Identifying Read-Only Components

When you create a project, you must indicate whether each component is modifiable in the context of that project. In most cases, you make them modifiable. However, in some cases you want to make a component read-only, which prevents project team members from changing its elements. Components can be used in multiple projects. Therefore, one project team may be responsible for maintaining a component, and another project team may use that component to build other components.

For example, in Figure 10, Project A team members maintain a set of library files. Project B team members reference some of those libraries when they build their components. In Project A, the

cm_libs component is modifiable. In Project B, the same component is read-only. With respect to the **cm_libs** component, Project A and Project B have a producer-consumer relationship.

Figure 10 Using a Read-Only Component

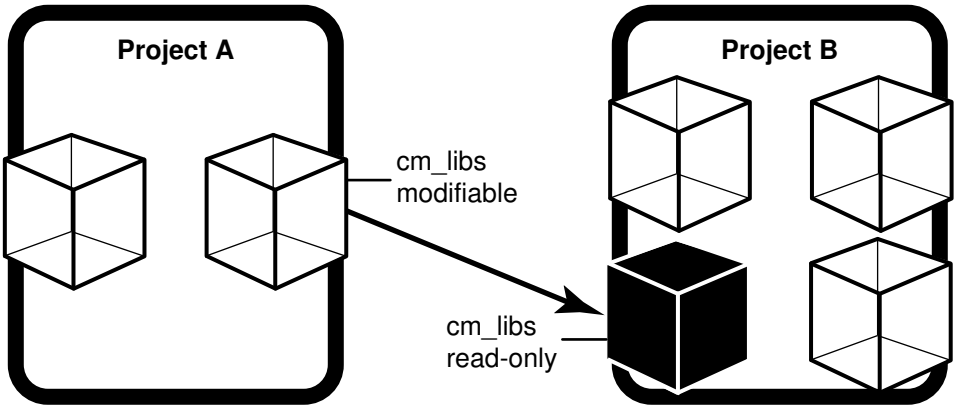

3.3 Specifying a Baseline Strategy

After you organize the project's components, determine your strategy for creating baselines of those components. The baseline strategy must define the following:

- When to create baselines
- How to name baselines
- The set of promotion levels
- How to test baselines

When to Create Baselines

At the beginning of a project, you must identify the baseline or baselines that represent the starting point for new development. As work on the project progresses, you need to create new baselines periodically.

Identifying the Initial Baseline

If your project represents a new version of an existing project, you probably want to start work from the latest baselines of the existing project's components. For example, if you are starting work on version 3.2 of the Transaction Builder project, identify the baselines that represent the released, or production, versions of its version 3.1 components.

If you are converting a base ClearCase configuration to a project, you can make baselines from existing labeled versions. Check whether the latest stable versions are labeled. If they are not, you need to create a label type and apply it to the versions that you plan to include in your project.

Ongoing Baselines

After developers start working on the new project and making changes, create baselines on a frequent (nightly or weekly) basis. This practice has several benefits:

- Developers stay in sync with each other's work.

  It is critical to good configuration management that developers have private work areas where they can work on a set of files in isolation. Yet extended periods of isolation cause problems. Developers are unaware of each other's work until you incorporate delivered changes into a new baseline, and they rebase their development streams.

- The amount of time required to merge versions is minimized.

  When developers rebase their development streams, they may need to resolve merge conflicts between files that the new baseline selects and the work in their private work areas. When you create baselines frequently, they contain fewer changes, and developers spend less time merging versions.

- Integration problems are identified early.

  When you create a baseline, you first build and test the project by incorporating the work delivered to the integration stream since the last baseline. By creating baselines frequently, you have more opportunities to discover any serious problems that a developer may introduce to the project inadvertently. By identifying a serious problem early, you can localize it and minimize the amount of work required to fix the problem.

Defining a Naming Convention

Because baselines are an important tool for managing a project, define a meaningful convention for naming them. A useful baseline name provides this information:

- Project name
- Milestone or phase of development schedule
- Date created

For example: **V4.0TRANS_BL2_990519**

Identifying Promotion Levels to Reflect State of Development

A promotion level is an attribute of a baseline that you can use to indicate the quality or stability of the baseline. ClearCase provides the following default promotion levels:

- Rejected
- Initial
- Built
- Tested
- Released

You can use some or all of the default promotion levels, and you can define your own. The levels are ordered to reflect a progression from lowest to highest quality. You can use promotion levels in development policies for the project. For example, you can set a policy that makes a baseline the *recommended baseline* when it has a certain promotion level or a higher one. By default, when developers join the project or rebase their development streams, they use the recommended baselines. Determine the set of promotion levels for your project and the criteria for setting each level.

Planning How to Test Baselines

Typically, software development teams perform several levels of testing. An initial test, known as a validation test, checks to see that the software builds without errors and appears to work as it should. A more comprehensive type of testing, such as regression testing, takes much longer and is usually performed by a team of software quality engineers.

When you make a new baseline, you need to lock the integration stream to prevent developers from delivering additional changes. This allows you to build and test a static set of files. Because validation tests are not exhaustive, you probably do not need to lock the integration stream for a long time. However, more extensive testing requires substantially more time.

Keeping the integration stream locked for a long time is not a good practice because it prevents developers from delivering completed work. One solution to this problem is to create a development stream to be used solely for extensive testing. After you create a new baseline that passes a validation test, your testing team can rebase the designated testing development stream to the new baseline. When the baseline passes the next level of testing, promote it. When you are confident that the baseline is stable, make it the recommended baseline so that developers can rebase their development streams to it.

For information on creating a testing development stream, see *Creating a Development Stream for Testing Baselines* on page 71.

3.4 Planning PVOBs

ClearCase stores UCM objects such as projects, streams, activities, and change sets in project VOBs (PVOBs). In addition to storing objects, PVOBs can function as administrative VOBs. You need to decide how many PVOBs to use for your system and whether to take advantage of the administrative capabilities of the PVOB.

Deciding How Many PVOBs to Use

Projects that use the same PVOB have access to the same set of components. If developers on different projects need to work on some of the same components, use one PVOB for those projects. For example, Figure 11 shows concurrent development of two versions of the **Webotrans** product. While most members of the team work on the 4.0 release in one project, a small group works on the 4.0.1 release in a separate project. Both projects use the same components, so they use one PVOB.

Figure 11 Related Projects Sharing One PVOB

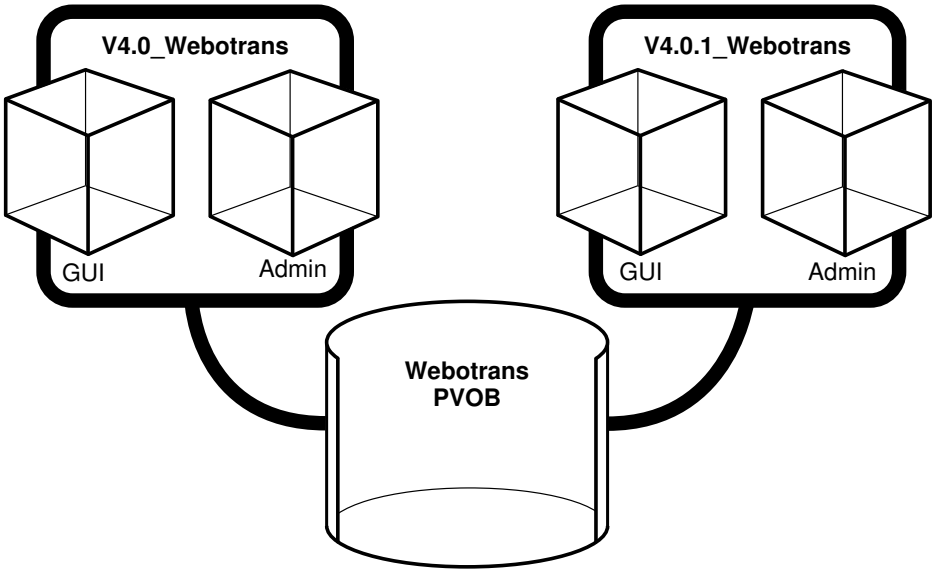

Consider using multiple PVOBs only when one or more of the following conditions applies:

- The projects do not share components and you anticipate that they will never need to share components.
- The projects are so large that PVOB capacity becomes an issue. For information on VOB capacity, see *Administering ClearCase*.
- You plan to use the UCM-ClearQuest integration, and you want to link projects to different ClearQuest user databases. See *Planning How to Use the UCM-ClearQuest Integration* on page 36 for information on the role of the PVOB in the integration.

Understanding the Role of the Administrative VOB

An *administrative VOB* stores global type definitions. VOBs that are joined to the administrative VOB with **AdminVOB** hyperlinks share the same type definitions without having to define them in each VOB. For example, you can define element types, attribute types, hyperlink types, and so on in an administrative VOB. Any VOB linked to that administrative VOB can then use those type definitions to make elements, attributes, and hyperlinks.

If you currently use an administrative VOB, you can associate it with your PVOB when you create the PVOB. ClearCase then creates an **AdminVOB** hyperlink between the PVOB and the administrative VOB. Thereafter, when you create components, ClearCase creates **AdminVOB** hyperlinks between the components and the administrative VOB so that the components can use the administrative VOB's global type definitions.

If you do not currently use an administrative VOB, do not create one. When you create components, ClearCase makes **AdminVOB** hyperlinks between the components and the PVOB, and the PVOB assumes the role of administrative VOB.

For details on administrative VOBs and global types, see *Administering ClearCase*.

3.5 Identifying Special Element Types

The concept of element types allows ClearCase to handle each class of elements differently. An *element type* is a class of file elements. ClearCase includes predefined element types, such as **file** and **text_file**, and lets you define your own. When you create an element type for use in UCM projects, you can specify a **mergetype** attribute, which determines how deliver and rebase operations handle merging of files of that element type.

When ClearCase encounters a merge situation during a deliver or rebase operation, it attempts to merge versions of the element. ClearCase requires user interaction only if it cannot reconcile differences between the versions. For certain types of files, you may want to impose different merging behavior.

Nonmerging Elements

Some types of files never need to be merged. For these files, you may want to ensure that no one attempts to merge them accidentally. For example, the deployment, or staging, component contains the executable files that you ship to customers or install in-house. These files are not under development; they are the product of the development phase of the project cycle. For these types of files, you can create an element type and specify **never** merge behavior.

NOTE: If you do not specify **never** merge behavior for these elements, developers could encounter problems when they attempt to deliver work to the project's integration stream. Developers create executable files when they build and test their work prior to delivering it. If these files are under version control as *derived objects*, they are included in the current activity's change set. During a deliver operation, ClearCase attempts to merge these executable files to the

integration stream unless the files are of an element type for which **never** merge behavior is specified.

Nonautomerging Elements

For some types of files, you may want to merge versions manually rather than let ClearCase merge them. One example is a Visual Basic form file, which is a generated text file. Visual Basic generates the form file based on the form that a developer creates in the Visual Basic GUI. Rather than let ClearCase change the form file during a merge operation, you want to regenerate the form file from the Visual Basic GUI.

For these types of files, you can create an element type and specify **user** merge behavior. For information on creating element types, see Chapter 12, *Using Element Types to Customize Processing of File Elements*, and the **mkeltype** reference page in *ClearCase Reference Manual*.

Defining the Scope of Element Types

When you define an element type, its scope can be ordinary or global. By default, the element type is ordinary; it is available only to the VOB in which you create it. If you create the element type in an administrative VOB and define its scope as global, other VOBs that have **AdminVOB** hyperlinks to that administrative VOB can use the element type. If you want to define an element type globally, and you do not currently use an administrative VOB, define the element type in the PVOB.

3.6 Planning How to Use the UCM-ClearQuest Integration

Before you can set up the UCM-ClearQuest integration, you need to make some decisions, which fall into two general categories:

- How to map PVOBs to ClearQuest user databases
- Which schema to use for the ClearQuest user databases

Mapping PVOBs to ClearQuest User Databases

This section describes three issues that you need to consider in deciding how many PVOBs to use for projects that link to ClearQuest user databases.

All Enabled Projects in a PVOB Must Link to the Same Database

When you enable a project to link to a ClearQuest user database, the integration stores a reference to that database in the project's PVOB. Any other projects in that PVOB that you enable must use the same database. Therefore, be careful when choosing projects to store in a PVOB. If you plan to link projects to different databases, use different PVOBs.

NOTE: If you use ClearCase MultiSite, all PVOB replicas must have access to the ClearQuest user database.

Projects Linked to Same Database Must Have Unique Names

Although UCM allows you to create projects with the same name in different PVOBs, you cannot link those projects to the same ClearQuest user database. Figure 12 illustrates this naming requirement.

Figure 12 Projects in Multiple PVOBs Linked to the Same ClearQuest Database

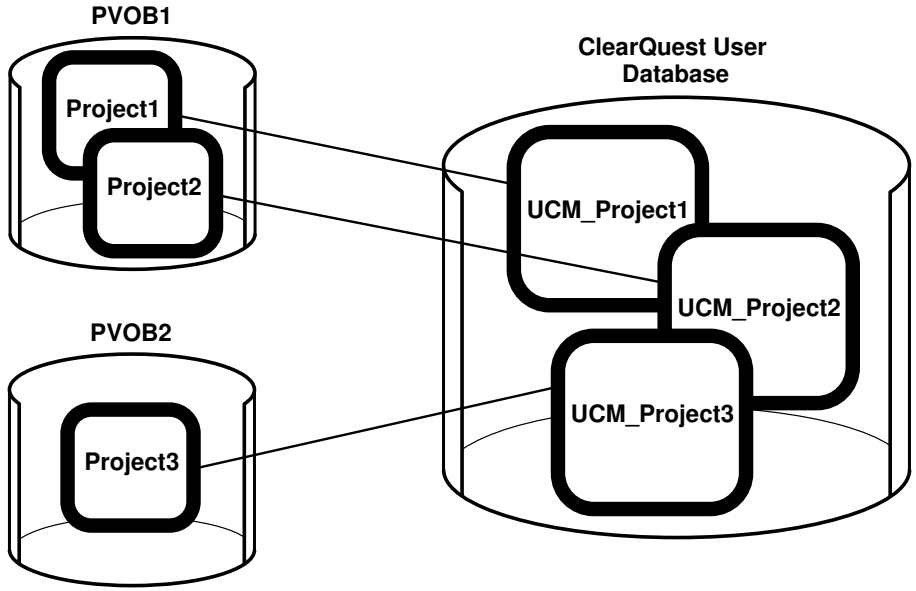

Use One Schema Repository for Linked Databases

If some developers on your team work on multiple projects, we recommend that you store the schemas for the ClearQuest user databases that are linked to those projects in one schema repository, as shown in Figure 13. This allows developers to switch between projects easily. If you store the schemas in different schema repositories, developers must use the ClearQuest Maintenance Tool to connect to a different schema repository whenever they switch projects.

Figure 13 Using the Same Schema Repository for Multiple ClearQuest Databases

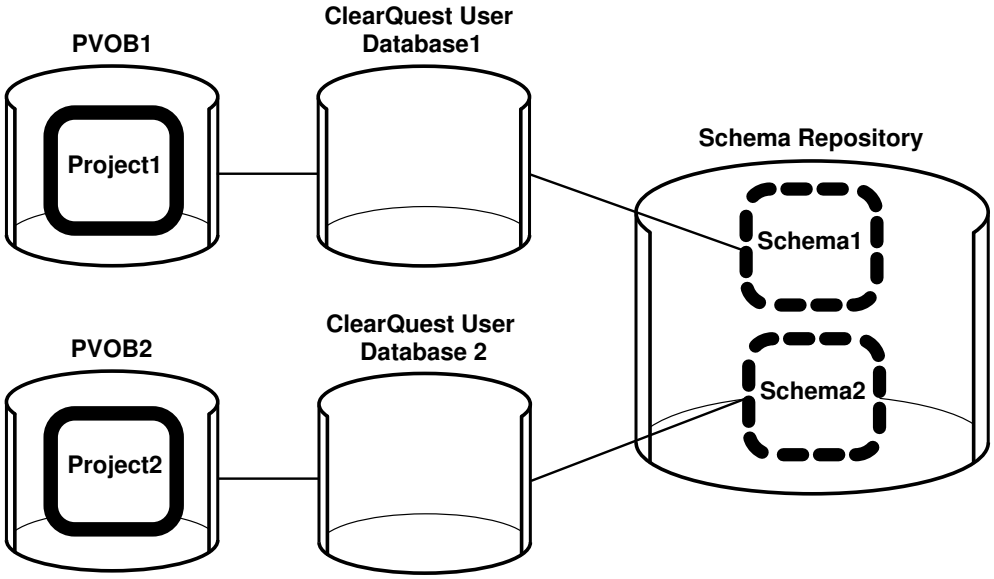

Deciding Which Schema to Use

To use the integration, you must create or upgrade a ClearQuest user database that is based on a UCM-enabled schema. A UCM-enabled schema meets the following requirements:

- The **UnifiedChangeManagement** package has been applied to the schema. A package contains metadata, such as records, fields, and states, that define specific functionality. Applying a package to a schema provides a way to add functionality quickly so that you do not have to build the functionality from scratch.

- The **UnifiedChangeManagement** package has been applied to at least one record type. This package adds fields and scripts to the record type, and adds the **Unified Change Management** tab to the record type's forms. Figure 14 shows the **Unified Change Management** tab.
- The **UCMPolicyScripts** package has been applied to the schema. This package contains the scripts for three ClearQuest development policies that you can enforce.

ClearQuest includes two predefined UCM-enabled schemas, named **UnifiedChangeManagement** and **Enterprise**. You can start using the integration right away by using one of these schemas, or you can use the ClearQuest Designer and the ClearQuest Package Wizard to enable a custom schema or another predefined schema to work with UCM. You can also use one of the predefined UCM-enabled schemas as a starting point and then modify it to suit your needs.

Figure 14 UCM Tab of Record Form for a UCM-Enabled Record Type

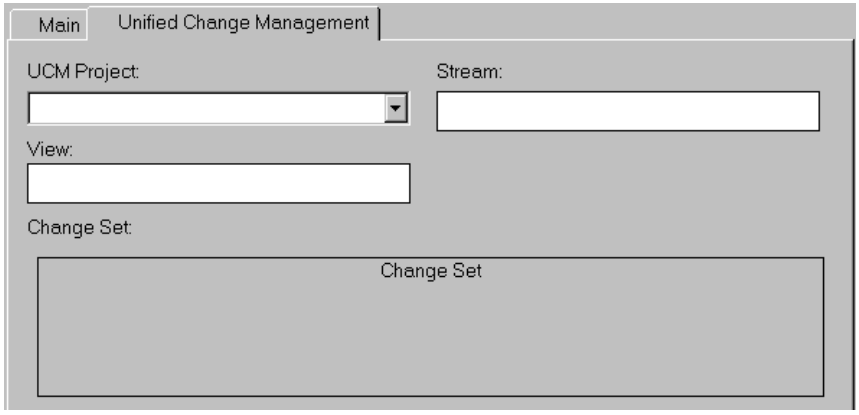

Overview of the UnifiedChangeManagement Schema

The UnifiedChangeManagement schema includes the following record types:

- **BaseCMActivity**
  This is a lightweight record type that you can use to store information about activities that do not require additional fields. Figure 15 shows the **Main** tab of the BaseCMActivity record form. You may want to use this record type as a starting point and then modify it to include additional fields and states.
- **Defect**
  This record type is identical to the record type of the same name that is included in

ClearQuest's other predefined schemas, with one exception: it is enabled to work with UCM. The Defect record type contains more fields and form tabs than the Activity record type to allow you to record detailed information.

- **UCMUtilityActivity**
  This record type is not intended for general use. The integration uses this record type when it needs to create records for itself, such as when you link a project that contains activities to a ClearQuest user database. You cannot modify this record type.

Figure 15 Main Tab of Record Form for the BaseCMActivity Record Type

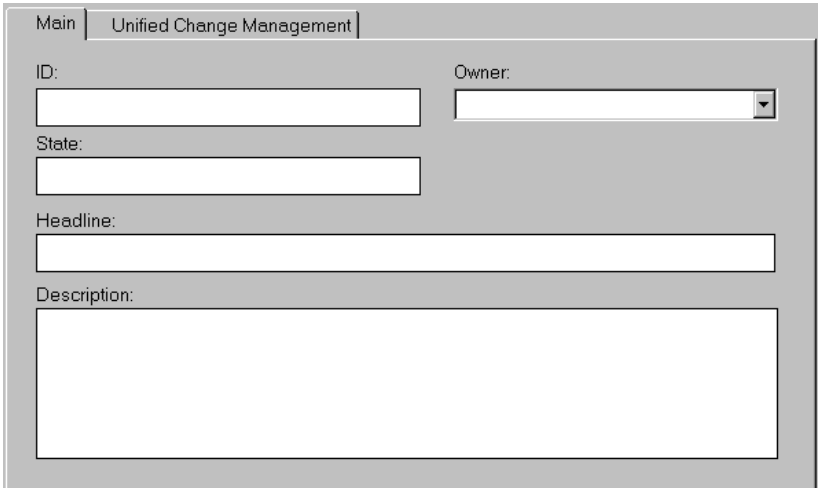

Enabling a Schema for UCM

If you decide not to use one of the predefined UCM-enabled schemas, you need to do some additional work to enable your schema to work with UCM. Before you can do this, you need to answer the following questions:

- Which record types are you enabling for UCM? You do not need to enable all record types in your schema, but you can link only records of UCM-enabled record types to activities.
- For each UCM-enabled record type:
  - Which state type does each state map to? You must map each state to one of the four UCM state types: Waiting, Ready, Active, Complete. See *Setting State Types* on page 49.
  - Which default actions are you using to transition records from one state to another? See *State Transition Default Action Requirements for Record Types* on page 50.

- Which policies do you want to enforce? The integration includes policies that you can set to enforce certain development practices. You can also edit the policy scripts to change the policies. See *Policies Available in UCM-ClearQuest Integration* on page 43 for details.

3.7 Considering Which Development Policies to Enforce

UCM includes policies that you can set to enforce certain development practices within a project. Some policies are available only if you enable the project to work with ClearQuest.

Policies Available in UCM

This section describes the policies that are available regardless of whether you enable the project to work with ClearQuest.

Recommended Baselines

Recommended baselines are the set of baselines that project team members use to rebase their development streams. In addition, when developers join the project, their development work areas are initialized with the recommended baselines. Select the promotion level at which baselines become recommended baselines.

NOTE: If a component does not contain a baseline whose promotion level is at or above the recommended baseline promotion level, ClearCase uses the component's foundation baseline for the project when developers attempt to rebase their development streams or join the project.

Modifiable Components

In most cases, you want components to be modifiable. For information on when to use read-only components, See *Identifying Read-Only Components* on page 29 .

Default View Types

When developers join a project, they use the Join Project Wizard to create their development views, integration views, and development streams. They use a development view and a development stream to work in isolation from the project team. They use an integration view to build and test their work against the latest work delivered to the integration stream by other

developers. ClearCase provides two kinds of views: dynamic and snapshot. Decide which type of view to use as the default for development and integration views.

Dynamic views use the ClearCase multiversion file system (MVFS) to provide immediate, transparent access to files and directories stored in VOBs. *Snapshot views* copy files and directories from VOBs to a directory on your computer.

We recommend that you use dynamic views as the default view type for integration views. Dynamic views ensure that when developers deliver work to the integration stream, they build and test their work against the latest work that other developers have delivered since the last baseline was created. Snapshot views require developers to copy the latest delivered files and directories to their computer (a *snapshot view update* operation), which they may forget to do.

Rebase Before Deliver

This policy requires developers to rebase their development streams to the project's current *recommended baselines* before they deliver work to the integration stream. The goal of this policy is to have developers build and test their work in their development work areas against the work included in the most recent stable baselines before they deliver to the integration stream. This practice minimizes the amount of merging that developers must do when they perform deliver operations.

Allow Deliveries from Stream with Pending Checkouts

This policy allows developers to deliver work to the integration stream even if some files remain checked out in the development stream. If you do not set this policy, developers must check in all files in their development streams before delivering work. You may want to require developers to check in files to avoid the following situation:

1. A developer completes work on an activity, but forgets to check in the files associated with that activity.

2. The developer works on other activities.

3. Having completed several activities, the developer delivers them to the integration stream. Because the files associated with the first activity are still checked out, they are not included in the deliver operation. Even though the developer may build and test the changes successfully in the development work area, the changes delivered to the integration may fail because they do not include the checked-out files.

Policies Available in UCM-ClearQuest Integration

This section describes the policies that are available only when you enable the project to work with ClearQuest. ClearQuest uses scripts to implement these policies. You can modify a policy's behavior by editing its script. See *Customizing ClearQuest Project Policies* on page 52.

Check Before Work On

ClearQuest invokes this policy when a developer attempts to work on an activity. The default policy script checks to see whether the developer's user name matches the name in the ClearQuest record's Owner field. If the names match, the developer can work on the activity. If the names do not match, the Work On action fails.

The intent of this policy is to ensure that all criteria are met before a developer can start working on an activity. You may want to modify the policy to check for additional criteria.

Check Before ClearCase Delivery

This default policy script is a placeholder: it does nothing. ClearCase invokes this policy when a developer attempts to deliver an activity in a UCM-enabled project. We recommend that you edit the script to implement an approval process to control deliver operations. For example, you may want to add an **Approved** check box to the activity's record type and require that the project manager select it before allowing developers to deliver activities.

Do ClearQuest Action After Delivery

ClearCase calls this policy at the end of a deliver operation for each activity included in the deliver operation. The default policy script uses the activity's default action to transition the activity to a Complete type state. If the default action requires entries in certain fields of the activity's record, and one of those fields is empty, the script returns an error and leaves the deliver operation in an uncompleted state. This state prevents the developer from performing another deliver operation, but it does not affect the current one. It does not roll back changes made during the merging of versions.

To recover from an error, the developer needs to fill in the required fields in the activity's record, and resume the deliver operation.

The integration runs this script for each activity in the deliver operation. The script may return success for any number of activities before returning an error on an activity. For the successful activities, the script may change their state when it invokes the default action. When you recover from an error and rerun the deliver operation, the script looks at all activities again. For those

that succeeded previously, the script does not attempt to change state. If you modify the script, be sure that it adheres to this behavior. ClearQuest returns an error if you attempt to change the state of a record to its current state.

4 Setting Up a ClearQuest User Database

This chapter describes how to set up a ClearQuest user database so that you can use the UCM-ClearQuest integration for your project. The steps in this chapter are typically completed by the ClearQuest database administrator. ClearQuest includes predefined schemas that are ready for use with UCM. You can also enable a custom schema, or another predefined schema, to work with UCM. See *Planning How to Use the UCM-ClearQuest Integration* on page 36 for information on the decisions you need to make before setting up the integration.

4.1 Using the Predefined UCM-Enabled Schemas

The predefined UCM schemas, named **UnifiedChangeManagement** and **Enterprise**, include the record type, field, form, state, and other definitions necessary to work with a UCM project. To set up a ClearQuest user database to work with UCM:

1. Create a user database that is associated with one of the predefined UCM-enabled schemas. In the ClearQuest Designer, click **Database>New Database** to start the New Database Wizard.
2. Complete the steps in the wizard. Step 4 prompts you to select a schema to associate with the new database. Scroll the list of schema names and select the new schema, as shown in Figure 16.
3. Click **Finish**.

Figure 16 Associating a User Database with a UCM-Enabled Schema

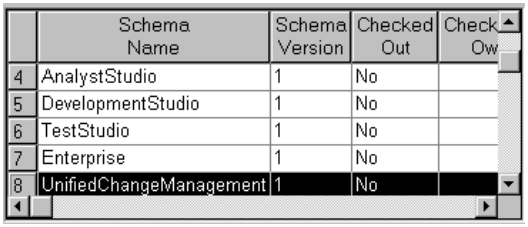

	Schema Name	Schema Version	Checked Out	Check Ow
4	AnalystStudio	1	No	
5	DevelopmentStudio	1	No	
6	TestStudio	1	No	
7	Enterprise	1	No	
8	UnifiedChangeManagement	1	No	

4.2 Enabling a Schema to Work with UCM

The predefined UCM schemas let you use the UCM-ClearQuest integration right away, but you may prefer to design a custom schema to track your project's activities and change requests, or you may prefer to use a different predefined schema. To enable a schema to work with UCM:

1. Ensure that the schema does not contain a record type named **UCM_Project**, which is a reserved name used by the UCM-ClearQuest integration.

2. In the ClearQuest Designer, click **Package>Package Wizard** to start the Package Wizard, as shown in Figure 17.

3. Add the **UCMPolicyScripts** package to your schema. If this package is not listed in the first page of the wizard, it has not been installed in your schema repository. To add the package to your schema repository, click **More Packages** to open the **Install Packages** dialog box; select the highest version of the package, and click **OK**. In the wizard, select the package, as shown in Figure 17. Click **Next**.

Figure 17 Adding the UCMPolicyScripts Package to a Schema

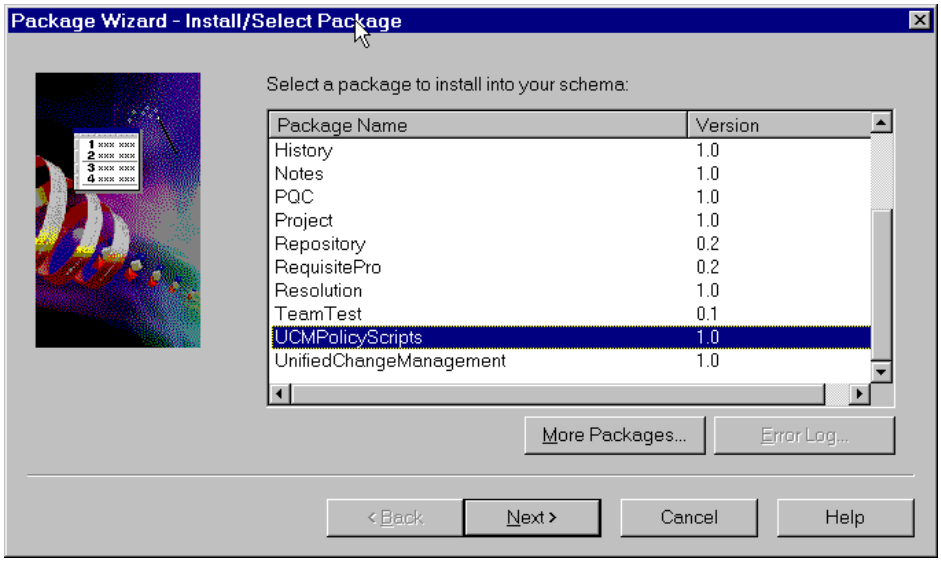

4. On the second page of the wizard, select your schema, and click **Finish**. To make the changes to the schema, ClearQuest checks out the schema for you. Check in the schema by clicking **File>Check In**. ClearQuest creates a new version of the schema.

5. Optionally, you can use the Package Wizard to apply the **BaseCMActivity** package to your schema. The **BaseCMActivity** package adds the BaseCMActivity record type to your schema. The BaseCMActivity record type is a lightweight activity record type. You may want to use the BaseCMActivity record type as a starting point and then modify it to include additional fields, states, and so on.

6. Apply the **UnifiedChangeManagement** package to the schema. Start the Package Wizard. Select **UnifiedChangeManagement**, and click **Next**.

7. In the second page of the wizard, select your schema. Click **Next**.

8. The third page of the wizard prompts you to specify the schema's record types. Select the check boxes of the record types that you want to enable. Click **Next**. All selected record types must meet the requirements listed in *Requirements for Enabling Custom Record Types* on page 49.

9. In the fourth page of the wizard, you must assign state types to the states for each record type that you choose to enable. For each state, click in the adjacent state type cell to display the list of available state types, as shown in Figure 18, and select one. To enable another record

type, click the arrow in the Record Type list to see the available record types. See *Setting State Types* on page 49 for a description of the four state types, and the rules for setting them.

When you are finished, click **Finish** to check out the schema.

Figure 18 Assigning State Types to a Record Type's States

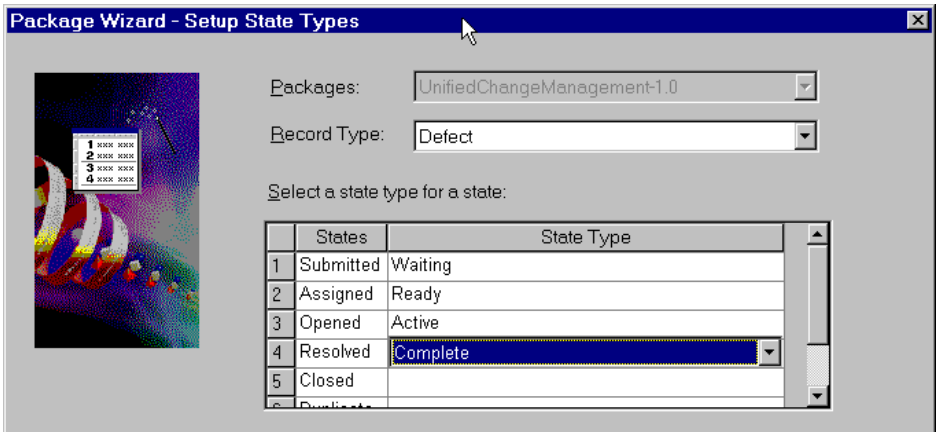

10. Before you can check in your schema, you must set default actions for the states of each enabled record type. Default actions are state transition actions that ClearQuest takes when a developer begins to work on an activity or delivers an activity. In the ClearQuest Designer workspace, navigate to the record type's state transition matrix, as shown in Figure 19.

Figure 19 Navigating to Record Type's State Transition Matrix

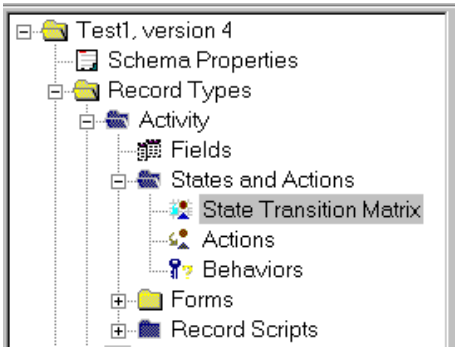

Double-click **State Transition Matrix** to display the matrix. Right-click the state column heading, and select **Properties** from its shortcut menu. Click the **Default Action** tab. Select

the default action. See *State Transition Default Action Requirements for Record Types* on page 50 for default action requirements. Before you can set default actions, you may need to add some actions to the record type. To do so, double-click **Actions** to display the Actions grid, and then click **Edit>Add Action**.

11. In the ClearQuest Designer workspace, navigate to the record type's **Behaviors**. Double-click **Behaviors** to display the Behaviors grid. Verify that the **Headline** field is set to **Mandatory** for all states. Verify that the **Owner** field is set to Mandatory for all Ready and Active state types.

12. Validate the schema changes by clicking **File>Validate**. Fix any errors that ClearQuest displays, and then check in the schema by clicking **File>Check In**.

13. Upgrade the user database so that it is associated with the UCM-enabled version of the schema by clicking **Database>Upgrade Database**.

Requirements for Enabling Custom Record Types

Before you can apply the **UnifiedChangeManagement** package to a custom record type, the record type must meet the following requirements:

- It contains a field named **Headline** defined as a SHORT_STRING, and a field named **Owner** defined as a REFERENCE to the ClearQuest-supplied **users** record type. The **Headline** field must be at least 120 characters long.
- It does not contain fields with these names:
  - **ucm_vob_object**
  - **ucm_stream**
  - **ucm_stream_object**
  - **ucm_view**
- It contains an action named **Modify** of type Modify.

Setting State Types

The integration uses a state transition model to help you monitor the progress of activities. To implement this model, the integration adds state types to UCM-enabled schemas. Table 2 lists and describes the four state types. You must assign each state to a state type. You must have at

least one state definition of state type Waiting, one of state type Ready, one of state type Active, and one of state type Complete.

Table 2 State Types in UCM-Enabled Schema

State Type	Description
Waiting	The activity is not ready to be worked on, either because it has not been assigned or it has not satisfied a dependency.
Ready	The activity is ready to be worked on. It has been assigned, and all dependencies have been satisfied.
Active	The developer has started work on the activity but has not completed it.
Complete	The developer has either worked on and completed the activity, or not worked on and abandoned the activity.

State Transition Default Action Requirements for Record Types

Record types can include numerous state definitions. However, UCM-enabled record types must have at least one path of transitions among state types as follows: Waiting to Ready to Active to Complete. The transition from one state to the next must be made by a default action.

For example, Figure 20 shows the actions and default actions between the states defined in the UCM-enabled BaseCMActivity record type included in the predefined UCM schema. The default actions are identified with an asterisk (*). The state types are in uppercase letters enclosed in brackets. The states appear immediately above their state types.

Figure 20 State Transition Diagram for UCM-enabled BaseCMActivity Record Type

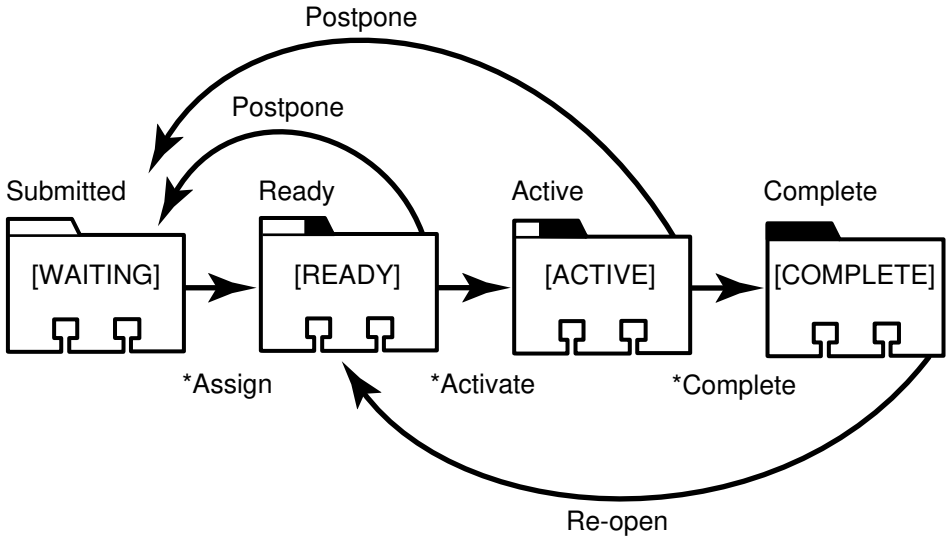

In addition to this single path requirement, states must adhere to the following rules:

- All Waiting type states must have a default action that transitions to another Waiting type state or to either a Ready or Active type state.
- If a Ready type state has an action that transitions directly to a Waiting type state, that Waiting type state must have a default action that transitions directly to that Ready type state.
- All Ready type states must have a default action that transitions to another Ready type state or to an Active type state.
- All Ready type states must have at least one action that transitions directly to a Waiting type state.
- For the BaseCMActivity record type, its initial state must be a Waiting type.

4.3 Customizing ClearQuest Project Policies

To implement the project policies, the integration adds the following pairs of scripts to a UCM-enabled schema:

- **UCM_ChkBeforeDeliver** and **UCM_ChkBeforeDeliver_Def**
- **UCM_ChkBeforeWorkOn** and **UCM_ChkBeforeWorkOn_Def**
- **UCM_CQActAfterDeliver** and **UCM_CQActAfterDeliver_Def**

Each policy has two scripts: a base script and a default script. The default scripts have **_Def** appended to their names and are installed by the **UnifiedChangeManagement** package. The integration invokes the base scripts, which are installed by the **UCMPolicyScripts** package. The base script calls the corresponding default script, which contains the logic for the default behavior. To modify the behavior of a policy, remove the call to the default script from the base script. Then add logic for the new behavior to the base script. Adhere to the rules stated in the base script.

Each script has a Visual Basic version and a Perl version. The Visual Basic scripts have a **UCM** prefix. The Perl scripts have a **UCU** prefix. For ClearQuest clients on Windows NT , the integration uses the Visual Basic scripts. For ClearQuest clients on UNIX, the integration uses the Perl scripts. If you modify a policy's behavior and your environment includes ClearQuest clients on both platforms, be sure to make the same changes in both the Visual Basic and Perl versions of the policy's script. Otherwise, the policy will behave differently for ClearQuest clients on UNIX and Windows NT.

For descriptions of these policies, see *Policies Available in UCM-ClearQuest Integration* on page 43.

4.4 Associating Child Activity Records with a Parent Activity Record

As project manager, you may assign activities for large tasks to developers. When the developers research their activities, they may determine that they need to perform several separate activities to complete one large activity.

For example, an "Add customer verification functionality" activity may require significant work in the product's GUI, the command-line interface, and a library. To more accurately track the progress of the activity, you can decompose it into three separate activities.

By using the parent/child controls in ClearQuest , you can accomplish this decomposition and tie the child activities back to the parent activity.

Using Parent/Child Controls

In ClearQuest, you use controls to display fields in record forms. A parent/child control, when used with a reference list field, lets you link related records. By adding a parent/child control to the record form of a UCM-enabled record type, you can provide the developers on your team with the ability to decompose a parent activity into several child activities.

To have ClearQuest change the state of the parent activity to Complete when all child activities have been completed, you need to write a hook. See *Administering Rational ClearQuest* for an example of such a hook.

4.5 Creating Users

Before you can assign activities to the developers on your project team, you must create user account profiles for each developer in ClearQuest. To do so:

1. In ClearQuest Designer, click **Tools>User Administration**.
2. Click **Add**.
3. Complete the **User Information** dialog box.

See *Administering Rational ClearQuest* and the ClearQuest Designer online help for details on creating user profiles.

4.6 Setting the Environment

Before you can enable a UCM project to work with a ClearQuest user database, you must define two environment variables as shown in Table 3. Developers who want to use the integration must also define these variables on their machines.

The ClearQuest installation directory includes a C shell script, **cq_setup.csh**, which you can execute to set the environment variables for you. For example:

% **source** *ClearQuest-install-directory***/cq_setup.csh**

Table 3 Environment Variables Required for Integration

Variable	**Setting**
$CQ_HOME	*ClearQuest-install-directory*/releases/ClearquestClient
$LD_LIBRARY_PATH ($SHLIB_PATH on HP-UX)	Must include: *ClearCase-install-directory*/shlib and *ClearQuest-install-directory*/releases/ClearquestClient/*architecture*/shlib

In addition, if you have multiple ClearQuest schema repositories, you must set the $SQUID_DBSET environment variable to the name of the schema repository you want to use.

5 Setting Up the Project

This chapter describes how to set up a project so that a team of developers can work in the Unified Change Management (UCM) environment. Before you set up a project, be sure to plan the project. See Chapter 3, *Planning the Project*, for information on what to include in a configuration management plan.

The chapter presents five scenarios:

- Creating a project from scratch
- Creating a project based on an existing base ClearCase configuration
- Creating a project based on an existing project
- Enabling a project to use the UCM-ClearQuest integration
- Creating a development stream reserved for testing new baselines

5.1 Creating a Project from Scratch

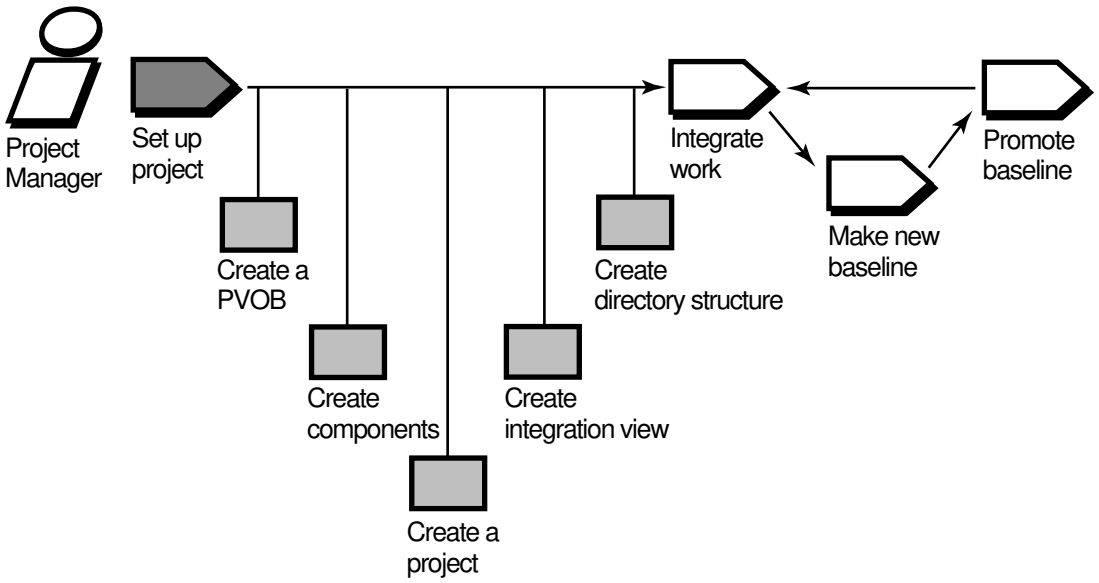

This section describes how to create and set up a new project that is not based on an existing project or on an existing set of ClearCase VOBs.

Creating the Project VOB (PVOB)

To create a PVOB:

1. Issue the **cleartool mkvob** command. For example:

   % **cleartool mkvob –tag /vobs/myproj_pvob –nc –ucmproject **
   /usr/vobstore/myproj_pvob.vbs

   The **–ucmproject** option indicates that you are creating a PVOB instead of a VOB. The **/usr/vobstore/myproj_pvob.vbs** path specifies the location of the PVOB's storage directory. A PVOB storage directory is a directory tree that serves as the repository for the PVOB's contents. A PVOB's storage directory contains the same subdirectories as a VOB's storage directory. For details about VOB storage directory structure, see *Administering ClearCase*.

2. Create the PVOB mount point to match the PVOB tag. For example:

   % **mkdir /vobs/myproj_pvob**

3. Mount the PVOB. For example:

   % **cleartool mount /vobs/myproj_pvob**

The PVOB assumes the role of administrative VOB. When you create components, ClearCase automatically makes **AdminVOB** hyperlinks between the components and the PVOB.

Creating Components

To create a component:

1. Make and set a view by using the **cleartool mkview** and **setview** commands. For example:

   % **cleartool mkview –tag myview /net/host2/view_store/myview.vws**

   % **cleartool setview myview**

2. Create a VOB by using the **cleartool mkvob** command. For example:

   % **cleartool mkvob –public –nc –tag /vobs/testvob1 /usr/vobstore/testvob1.vbs**

   A component's root directory must be the root directory of a VOB. Therefore, a VOB must exist before you can create a component.

3. Create the VOB mount point to match the VOB-tag. For example:

   % **mkdir /vobs/testvob1**

4. Mount the VOB. For example:

   % **cleartool mount /vobs/testvob1**

5. Issue the **cleartool mkcomp** command. For example:

   % **cleartool mkcomp –nc –root /vobs/testvob1 testcomp1@/vobs/myproj_pvob**

   In this example, the **mkcomp** command creates a component named **testcomp1** based on the VOB named **testvob1**. The VOB and PVOB must be mounted before you issue the command. All projects that use the **myproj_pvob** PVOB can access the **testcomp1** component.

As an alternative to using the **cleartool mkcomp** command, you can convert an existing VOB into a component by using the ClearCase Project Explorer. See *Making a VOB into a Component* on page 63 for details.

Creating the Project

To create a project:

1. On the command line, type **clearprojexp**. The Project Explorer appears. The Project Explorer is the graphical user interface (GUI) through which you create, manage, and view information about projects.

2. The left pane of the Project Explorer lists root folders for all PVOBs in the local ClearCase domain. Each PVOB has its own root folder. ClearCase creates the root folder using the name of the PVOB.

   ClearCase also creates a folder called **Components**, which contains entries for each component in the PVOB. Folders can contain projects and other folders. Select the root folder for the PVOB that you want to use for storing project information.

3. Click **File>New>Folder** to create a project folder. You do not need to create a project folder, but it is a good idea. As the number of projects grows, project folders are helpful in organizing related projects.

4. In the left pane, select the project folder or root folder. Click **File>New>Project**. The New Project Wizard appears.

5. In Step 1 of the New Project Wizard, enter a descriptive name for the project in the **Project Title** box. Enter a comment in the **Description** box to describe the purpose of this project.

6. Step 2 asks whether you want to create the project based on an existing project. Because you are creating a project from scratch, click **No**.

7. Step 3 asks you to choose the baselines that the project will use. These baselines are known as *foundation baselines* because they are the foundation upon which all work within the project is built.

   Click **Add** to open the **Add Baseline** dialog box. In the **Component** list, select one of the components that you previously created. The component's initial baseline appears in the **Baselines** list. Select the baseline. Be sure that the **Allow project to modify the component** check box is selected unless you want the component to be read-only. See *Identifying Read-Only Components* on page 29 for information on when you may want to use read-only

components. Click **OK**. The baseline now appears in the list in Step 3. Continue to use the **Add Baseline** dialog box until the project contains its full set of foundation baselines.

8. Step 4 prompts you to specify the development policies to enforce for this project. Select the check boxes for the policies you want to enforce. See *Considering Which Development Policies to Enforce* on page 41 for information about each policy.

9. Step 5 asks whether to configure the project to work with the ClearQuest integration. To enable the project to work with ClearQuest, click **Yes**, and select a ClearQuest user database from the list. See *Enabling a Project to Use the UCM-ClearQuest Integration* on page 66 for details about the integration.

Defining Promotion Levels

ClearCase provides five baseline promotion levels. You can keep some or all of them, and you can define your own promotion levels. To define the promotion levels that your project uses:

1. In the Project Explorer, select the PVOB root folder that contains your project, and then click **Tools>Define Promotion Level**. All projects that use that PVOB have access to the same set of promotion levels.

2. The **Define Promotion Levels** dialog box appears. To remove an existing promotion level, select it and click **Remove**. To change the order of promotion levels, select a promotion level and use the **Move Up** or **Move Down** buttons.

3. To add a new promotion level, click **Add**. The **Add Promotion Level** dialog box appears. Enter the name of the new promotion level and click **OK**. The new promotion level appears in the list of promotion levels in the **Define Promotion Levels** dialog box. Move it to the desired place in the order.

4. When you finalize the set and order of promotion levels, select one to be the initial promotion level for new baselines. The initial promotion level is the level assigned by default when you create a baseline.

Creating an Integration View

When you create a project, ClearCase creates the project's integration stream for you. To see and make changes to the project's shared elements, you need an *integration view*. To create an integration view:

1. In the Project Explorer, navigate to the integration stream by moving down the object hierarchy:

   a. Root Folder

   b. Project Folder

   c. Project

   d. Stream

   Figure 21 illustrates this hierarchy.

Figure 21 Navigating to Integration Stream in Project Explorer

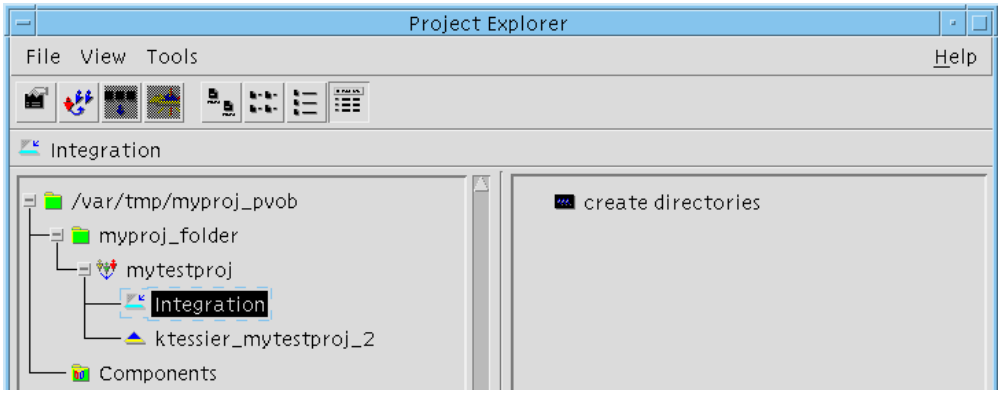

2. Select the integration stream and click **File>New>View**.

3. The **Create View** dialog box appears. Accept the default values to create an integration view associated with the integration stream. By default, the **Create View** dialog box uses this convention for the integration view name: *username*_**Integration**.

   ClearCase supports two kinds of views:

   - Dynamic views, which use the ClearCase multiversion file system (MVFS) to provide immediate, transparent access to files and directories stored in VOBs.

   - Snapshot views, which copy files and directories from VOBs to a directory on your computer.

   We recommend that you make the integration view a dynamic view to ensure that you always see the correct version of files and directories that developers deliver to the

integration stream. With a snapshot view, you have to perform an *update* operation to copy the latest delivered files and directories to your computer. For more information about dynamic and snapshot views, see *Developing Software with ClearCase*.

Creating and Setting an Activity

Before you can add elements to the integration stream, you need to create and set an activity.

1. Set your integration view. For example:

   % **cleartool setview kmt_Integration**

2. Issue the **cleartool mkactivity** command. For example:

   % **cleartool mkactivity –headline "Create Directories" create_directories**

   The ClearCase GUI tools use the name specified with **–headline** to identify the activity. The last argument, **create_directories**, is the activity-selector. Use the activity-selector when you issue **cleartool** commands.

3. By default, when you make an activity with the **cleartool mkactivity** command, ClearCase sets your view to that activity. ClearCase does not set your view to an activity if you create multiple activities in the same command line or if you specify a stream with the **–in** option. If you need to set your integration view to the activity, use the **cleartool setactivity** command. For example:

   % **cleartool setactivity create_directories**

Creating the Directory Structure

Because you are creating the project from scratch, you need to create the directory elements within the project's components to implement the directory structure that you define during the planning phase. See *Defining the Directory Structure* on page 28. To add a directory element to a component:

1. With your integration view set to an activity, navigate to the component. For example:

   % **cd /vobs/testcomp1**

2. Check out the component's root directory. For example:

```
% cleartool co –nc .
```

3. Issue the cleartool **mkelem** command. For example:

```
% cleartool mkelem –nc –eltype directory design
```

   This example creates a directory element called **design**. By default, the **mkelem** command leaves the element checked out. To add elements, such as subdirectories, to the directory element, you must leave the directory element checked out.

4. When you finish adding elements to the new directory, check it in. For example:

```
% cleartool ci –nc design
```

5. When you finish creating directory elements, check in the component's root directory. For example:

```
% cleartool ci –nc .
```

For additional information about creating directory and file elements, see *Developing Software with ClearCase* and the **mkelem** reference page.

Importing Directories and Files from Outside ClearCase

If you have a large number of files and directories that you want to place under ClearCase version control, you can speed the process by using the **clearexport** and **clearimport** command-line utilities. These two utilities allow you to migrate an existing set of directories and files from another version control software system, such as RCS or PVCS, to ClearCase. You can also use **clearexport** and **clearimport** to place directories and files that are not currently under any version control under ClearCase control.

To migrate source files into a component:

1. Create and set a non-UCM view by using the **cleartool mkview** and **setview** commands.
2. From within the view, run **clearexport** to generate a data file from your source files.
3. From within the view, run **clearimport** to populate the component with the files and directories from the data file.

4. In the component, create a baseline from a labeled set of versions. If the versions that you want to include in the baseline are not labeled, create a label type and apply it to the versions. See *Making a Baseline from a Label* on page 64 for details.

As an alternative, you can use **clearexport** and **clearimport** on VOBs, and then convert the VOBs to components. See *Creating a Project Based on an Existing ClearCase Configuration* on page 63 for details on converting VOBs into components.

For details on using **clearexport** and **clearimport**, see *Administering ClearCase* and the **clearexport** and **clearimport** reference pages.

5.2 Creating a Project Based on an Existing ClearCase Configuration

If you have existing VOBs, you may want to convert them into components so that you can include them in projects. This section describes how to set up a project based on existing VOBs.

Creating the PVOB

Use the **cleartool mkvob** command as described in *Creating the Project VOB (PVOB)* on page 56. If you currently use an administrative VOB, use the **cleartool mkhlink** command to create an **AdminVOB** hyperlink between the PVOB and the administrative VOB. When you create components, they then use the existing administrative VOB.

Making a VOB into a Component

To make a VOB into a component:

1. In the Project Explorer, select the PVOB. Click **Tools>Import VOB**. The **Import VOB** dialog box appears.

2. In the **Available VOBs** list, select the VOB that you want to make into a component. Click **Add** to move the VOB to the **VOBs to Import** list. You can add more VOBs to the **VOBs to Import** list. If you change your mind, you can select a VOB in the **VOBs to Import** list and click **Remove** to move it back to the **Available VOBs** list. When you are finished, click **Import**.

Making a Baseline from a Label

After you convert an existing VOB into a component, to access the directories and files in that component, you must create a baseline from the set of versions identified by a label type. To create the baseline:

1. If the set of versions that you want to use are not already labeled, use the **cleartool mklbtype** and **mklabel** commands to create and apply a label type. For example:

   ```
   % cleartool mklbtype –c "label for release 2" REL2
   Created label type "REL2".

   % cleartool mklabel -recurse REL2 .
   Created label "REL2" on "." version "/main/5".
   Created label "REL2" on "./src" version "/main/6".
   Created lable "REL2" on "./src/Makefile" version "/main/2".
   ```

   The **–recurse** option directs ClearCase to apply the label to all versions at or below the current working directory.

2. In the Project Explorer, select the PVOB. Click **Tools>Import Label**. Step 1 of the Import Label Wizard appears.

3. In the **Available Components** list, select the component that contains the label from which you want to create a baseline. Click **Add** to move that component to the **Selected Components** list. If you change your mind, select a component in the **Selected Components** list and click **Remove** to move the component back to the **Available Components** list.

4. InStep #2, select the label type that you want to import, and enter the name of the baseline that you want to create for the versions identified by that label type. Then select the baseline's promotion level.

Creating the Project

Use the New Project Wizard to create the project as described in *Creating the Project* on page 58.

Creating an Integration View

Create an integration view as described in *Creating an Integration View* on page 59.

5.3 Creating a Project Based on an Existing Project

As you create new projects, you may need to create new versions of existing projects. For example, suppose you have released Version 3.0 of the Webotrans project and are planning for Version 3.1. You anticipate that Version 3.1 will use the same components as Version 3.0. Therefore, you want to use the latest baselines in the Version 3.0 components as the foundation baselines for Version 3.1 development.

Reusing Existing PVOB and Components

Because your project is a new version of an existing project and uses the same components as the existing project, do not create a new PVOB for this project. Continue to use the existing PVOB.

Creating the Project

Start the New Project Wizard, as described in *Creating the Project* on page 58, to create the project.

In Step 2 of the wizard, select **Yes** to indicate that the project begins from the baselines in an existing project. Then navigate to the project that contains those baselines. Figure 22 shows that the new project is based on the baselines in **mytestproj**.

Figure 22 Step 2 of New Project Wizard

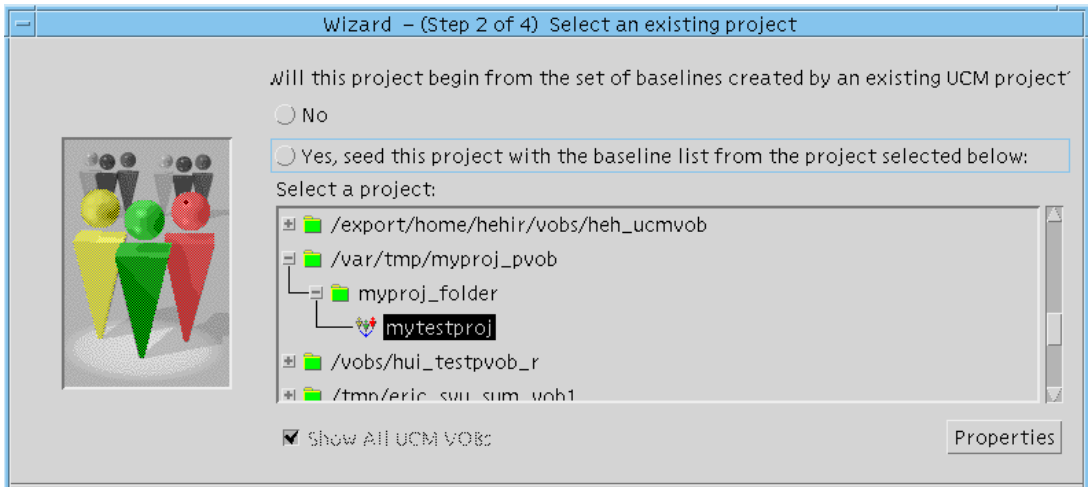

Step 3 lists the latest baselines in the project that you select in Step 2. You can add baselines from components that are not part of the existing project by clicking **Add** to open the **Add Baseline** dialog box. Similarly, you can remove a baseline by selecting it and clicking **Remove**.

Finish the remaining steps in the wizard as described in *Creating the Project* on page 58.

Creating an Integration View

When you create a new project, ClearCase creates a new integration stream for you. Therefore, you need to create a new integration view to access elements in the integration stream. Create an integration view as described in *Creating an Integration View* on page 59.

5.4 Enabling a Project to Use the UCM-ClearQuest Integration

Before you can connect a project to a ClearQuest user database, you must set up the database to use a UCM-enabled schema. See Chapter 4, *Setting Up a ClearQuest User Database*.

To enable a project to work with a ClearQuest user database:

1. In the left pane of the Project Explorer, right-click the project to display its shortcut menu. Click **Properties** to display its property sheet.

2. Click the **ClearQuest** tab and then select the **Project is ClearQuest-enabled** check box. Select the user database from the list, as shown in Figure 23. The first time that you enable a project, ClearQuest opens its **Login** dialog box. Enter your user name, password, and the name of the database to which you are linking the project.

3. Select the development policies that you want to enforce. See *Policies Available in UCM-ClearQuest Integration* on page 43 for a description of these policies. Click **OK** when you are finished.

If you are creating a new project, you can enable the project to work with ClearQuest by selecting **Project is ClearQuest-enabled** and selecting the user database in Step 5 of the New Project Wizard.

ClearCase does not require you to enable all projects in the PVOB to work with ClearQuest. However, all enabled projects in the same PVOB must use the same ClearQuest database. Therefore, give careful consideration to choosing the ClearQuest database. See *Mapping PVOBs to ClearQuest User Databases* on page 37 for details.

After you enable a UCM project to work with a ClearQuest user database, you may decide to link the project to a different user database. You can switch databases by selecting a different one on the ClearQuest tab of the project's property sheet if no other project in the same PVOB is ClearQuest-enabled, and no activities have been created.

Migrating Activities

If your project contains activities when you enable it to work with a ClearQuest database, the integration creates records for each of those activities by using the **UCMUtilityActivity** record type. If you want to store all of your project's activities in records of some other record type, enable the project when you create it, before team members create any activities. After the migration is complete, any new activities that you create can link to records of any UCM-enabled record type.

Figure 23 Enabling a Project to Work with a ClearQuest User Database

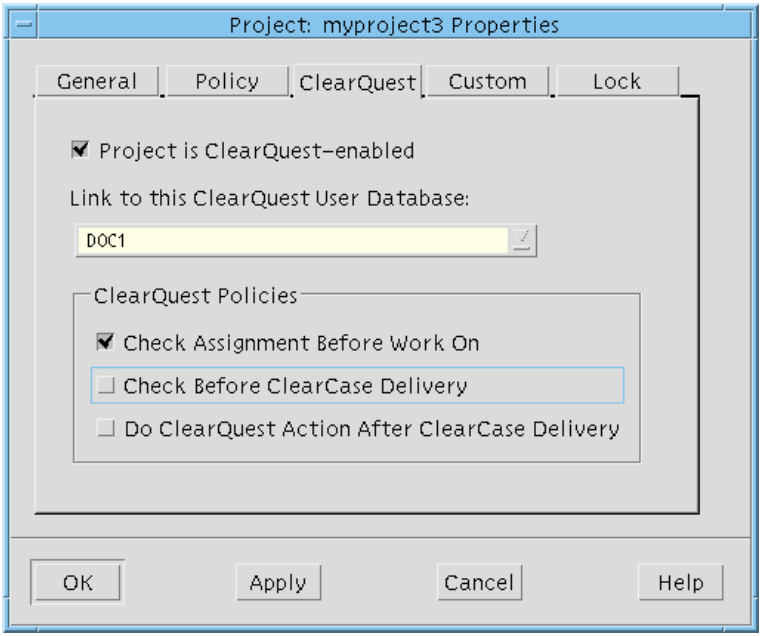

Setting Project Policies

A UCM-enabled schema includes three policies that you can set from either ClearCase or ClearQuest.

In ClearCase, set the policies by selecting check boxes on the ClearQuest tab of the project's property sheet, as shown in Figure 23.

To set policies from ClearQuest:

1. Start the ClearQuest client by entering **clearquest** at the shell prompt. In the ClearQuest client workspace, navigate to the **UCMProjects** query, as shown in Figure 24.
2. Double-click the query to display all UCM-enabled projects.
3. Select a project from the **Results** set. The project's form appears.
4. On the form, click **Actions** and select **Modify**. Select the check boxes for the policies you want to set.

See *Policies Available in UCM-ClearQuest Integration* on page 43 for descriptions of the policies.

Figure 24 Navigating to the UCMProjects Query

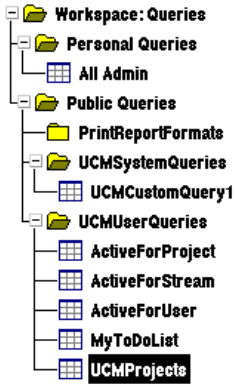

Assigning Activities

To create and assign activities in ClearQuest:

1. Start the ClearQuest client, and log in to the user database connected to the project.
2. Click **Actions>New**. The **Choose a record type** dialog box appears. Select a UCM-enabled record type, and click **OK**.
3. The **Submit** form appears. Fill in the boxes on each tab. On the **Main** tab, you must fill in at least the **Headline** and **Owner** boxes. On the **Unified Change Management** tab, select the project. When you finish filling in the boxes, click **OK**. ClearQuest creates the record.

User account profiles must exist in ClearQuest for the developers to whom you assign activities. See *Creating Users* on page 53 for details on creating user account profiles.

Disabling the Link Between a Project and a ClearQuest User Database

There may be times when you want to disable the link between a project and a ClearQuest user database. If another project in the same PVOB is ClearQuest-enabled or if activities have been

created, you must first disable the link between each ClearQuest-enabled project in the PVOB and the user database. To disable the links:

1. On the **ClearQuest** tab of the project's property sheet, clear the **Project is ClearQuest-enabled** check box.

2. Click **OK** on the **ClearQuest** tab. The integration disables the link between the project and the ClearQuest database. The integration also removes any existing links between activities and their corresponding ClearQuest records.

3. Display the project's property sheet again, select the **Project is ClearQuest-enabled** check box, and select another user database if you want to link the project to a different user database.

   NOTE: If you select the same user database that you just unlinked, the integration creates new ClearQuest records for the project's activities; it does not link the activities to the ClearQuest records with which they were previously linked.

Fixing Projects That Contain Linked and Unlinked Activities

It is possible that after you enable a project to work with ClearQuest, some of the project's activities remain unlinked to ClearQuest records. Similarly, when you disable the link between a project and ClearQuest, some activities may remain linked. Two scenarios can cause your project to be in this inconsistent state:

- A network failure or a general system crash occurs during the enabling or disabling operation and interrupts the activity migration.

- The project's PVOB is in a ClearCase MultiSite configuration, and unlinked activities were added by a MultiSite synchronization operation to the local PVOB's project, which is enabled to work with ClearQuest.

Detecting the Problem

If a developer attempts to take an action, such as modifying an unlinked activity in an enabled project, the integration displays an error and disallows the action.

Correcting the Problem

To restore the project to a consistent state:

1. In the Project Explorer, display the project's property sheet, and click the **ClearQuest** tab.
2. Click **Ensure all Activities are Linked**. The integration checks all the project's activities. If the project is enabled, the integration links any unlinked activities. The integration then displays the following summary information:
   - Number of activities that had to be linked.
   - Number of activities that were previously linked.
   - Number of activities that could not be linked because they are not mastered in the current PVOB replica. In this case, the integration also displays a list of replicas on which you must run the **Ensure all Activities are Linked** operation again to correct the problem.

5.5 Creating a Development Stream for Testing Baselines

When you make a new baseline, we recommend that you lock the integration stream so that you can build and test a static set of files. Otherwise, developers can inadvertently cause confusion by delivering changes while you are building and testing. Locking the integration stream for a short period of time is acceptable; locking the integration stream for several days can result in a backlog of completed but undelivered activities. To avoid locking out developers for a long period of time, you may want to create a development stream and use it for extensive testing of baselines.

To create a development stream:

1. In ClearCase Project Explorer, select the project, and click **File>New Stream**.

   The **Create a Development Stream** dialog box appears.

2. Enter a name and description for the new stream. Be sure that the **Prompt me to create a View for this stream** check box is selected. Click **OK**.

   The **Create View** dialog box appears.

3. Fill in the fields of the **Create View** dialog box to create a development view for the stream.

By default, ClearCase uses the set of *recommended baselines* when creating a development stream. Because the new baseline has not been tested extensively, you probably have not yet promoted

it to the level associated with recommended baselines. Therefore, you need to change the development stream's foundation baseline to be the one that you want to test:

1. In ClearCase Project Explorer, select the new development stream, and click **File>Properties** to display the stream's property sheet.

2. Click the **Configuration** tab. Select the component whose foundation baseline you want to change. Click **Change**.

   The **Change Baseline** dialog box appears.

3. Select the baseline that you want to test. Click **OK**.

4. If you want to change the foundation baseline of another component, select it from the **Configuration** tab, and repeat the process. When you are finished, click **OK** to dismiss the property sheet.

Now the development stream is configured so that you can build and test the new baselines, and developers can deliver changes to the integration stream without being concerned about interfering with the building and testing process.

6 Managing the Project

After you create and set up a project, developers join the project, work on activities, and deliver completed activities to the integration stream. As project manager, you need to maintain the project so that developers do not get out of sync with each other's work. This chapter describes the following maintenance tasks:

- Adding components
- Integrating work delivered by the remote deliver model
- Making new baselines
- Testing baselines
- Promoting and demoting baselines
- Tracking the progress of the project
- Cleaning up the project

6.1 Adding Components

Over time, the scope of your project typically broadens, and you may need to add components. To add a component to a project's integration stream:

1. Enter **clearprojexp** to start ClearCase Project Explorer.
2. In the left pane, select the project.
3. In the right pane, select the integration stream. Click **File>Properties** to open the integration stream's **Properties** dialog box.

4. Click the **Configuration** tab, and then click **Add**. The **Add Baseline** dialog box appears.

5. In the **Component** list, select the component that you want to add. The component's baselines appear in the Baselines list. Figure 25 shows the baselines available in the **testcomp3** component.

Figure 25 Add Baseline Dialog Box

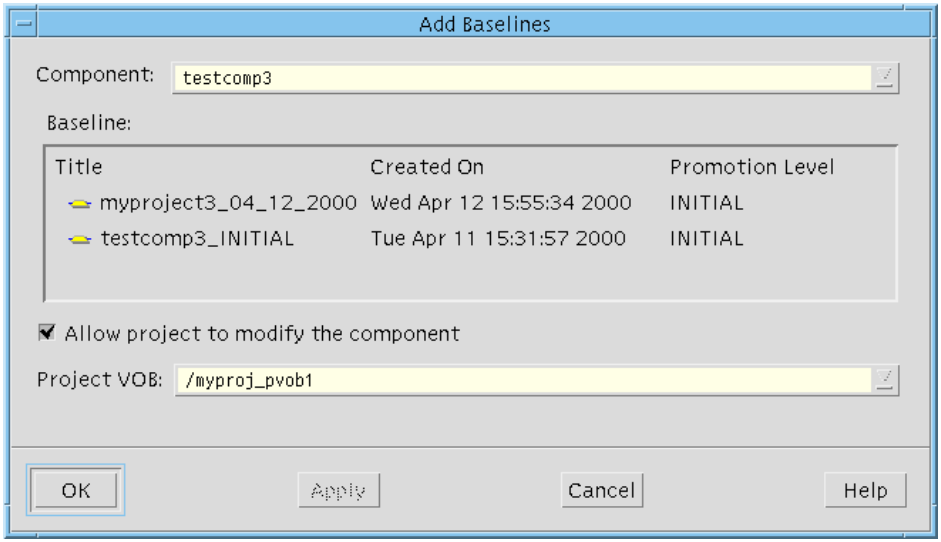

6. In the **Baselines** list, select the baseline that you want to add to the project.

7. Click **OK**. The **Add Baseline** dialog box closes, and the baseline that you chose appears on the **Configuration** tab.

8. Click **OK** to close the integration stream's **Properties** dialog box.

   The **Rebase Stream Preview** dialog box appears. To modify the integration stream's configuration to include the new foundation baseline, UCM needs to rebase the integration stream.

9. Click **OK** in the **Rebase Stream Preview** dialog box.

10. Click **Complete** to finish the rebase operation.

Updating Snapshot View Load Rules

If your integration view is a snapshot view, you need to edit the view's load rules to include the components that you add to the integration stream. A snapshot view's load rules specify which components ClearCase loads into the view. To edit the integration view's load rules:

1. In the Project Explorer, select the integration stream, and click **File>Properties** to display the integration stream's property sheet.
2. In the property sheet, click the **Load Rules** tab.
3. Select the component or components that you added to the integration stream.
4. Click **Add**. Click **OK** to dismiss the property sheet.

In addition, you need to know whether any developers working on the project use snapshot views for their development views. When a developer who uses a snapshot view rebases to a baseline that contains a new component, ClearCase updates the snapshot view's config spec, but it does not update the view's load rules. When you add a component, notify developers who use snapshot views that they need to update the load rules for their development views after they rebase their development streams to the new baseline.

6.2 Integrating the Project

In most cases, developers complete the deliver operations that they start. However, in a MultiSite configuration where the project's integration stream is mastered at a different replica than the developer's development stream, the developer cannot complete deliver operations. When ClearCase detects such a stream mastership situation, it makes the deliver operation a *remote deliver* operation.

In a remote deliver operation, ClearCase starts the deliver operation but leaves it in the posted state. It is up to you, as project manager, to find and complete deliver operations in the posted state. Developers who have deliver operations in the posted state cannot deliver from, or rebase, their development streams until you complete or cancel their deliver operations.

Finding WorkThat is Ready to Be Delivered

To find all deliver operations that are in the posted state:

1. In the Project Explorer, select the project.
2. Click **Tools>Find Posted Deliveries**. The **Find Posted Deliveries** dialog box appears, as shown in Figure 26, and lists all development streams within the project that contain deliver operations in the posted state.

Figure 26 Find Posted Deliveries Dialog Box

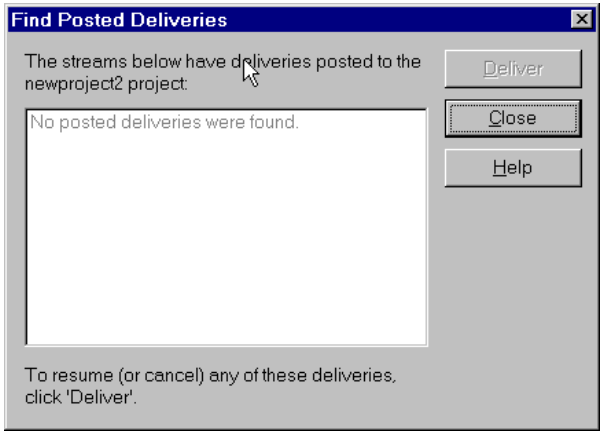

Completing Remote Deliver Operations

To complete remote deliver operations for a development stream:

1. Select the development stream from the list in the **Find Posted Deliveries** dialog box.
2. Click **Deliver**. The **Deliver** dialog box appears. Click **Resume** to resume the deliver operation. Click **Cancel** to cancel the deliver operation. See *Developing Software with ClearCase* for details on completing the deliver operation.

Undoing a Deliver Operation

In addition to the remote deliver scenario, there is another case where you may need to help developers with their deliver operations. At any time before completing the deliver operation, developers can back out of the deliver operation and undo any changes made during the operation. However, if developers check in their versions to the integration view, they cannot easily undo the changes. When this happens, you may need to remove the checked in versions by using the **cleartool rmver –xhlink** command.

NOTE: The **rmver** command erases part of your organization's development history, and it may have unintended consequences. Therefore, be very conservative in using this command, especially with the **–xhlink** option. See the **rmver** reference page in *ClearCase Reference Manual* for details.

Note that removing a version does not guarantee that the change is really gone. If a successor version was created or if the version was merged before you removed the version, the change still exists. You may need to check out the file, edit it to remove the change, and check the file back in.

6.3 Creating a New Baseline

As developers deliver work to the integration stream, it is important that you frequently make new baselines that record the changes. Developers can then rebase to the new baselines and stay current with each other's changes.

Locking the Integration Stream

Before you make a new baseline, lock the integration stream to prevent developers from delivering work. This ensures that you are dealing with a static set of files. To lock the integration stream:

1. In the Project Explorer, select the integration stream.
2. Click **File>Properties** to display the integration stream's property sheet.
3. Click the **Lock** tab.

4. Click **Locked** and then click **OK**.

Verifying That the Code Base Is Stable

After you lock the integration stream, we recommend that you build and test the project's executable files to make sure that the changes delivered by developers since the last baseline do not contain any bugs. For information on performing builds, see *Building Software with ClearCase*. Because you lock the integration stream when you build and test in it, we recommend that you use a separate development stream for extensive testing of new baselines. Perform only quick validation tests in the integration stream so that it is not locked for an extended period of time. See *Testing the Baseline* on page 80 for information about using a development stream for testing new baselines.

Making the New Baseline

To make a new baseline:

1. In the Project Explorer, select the project's integration stream.
2. Click **Tools>Make Baseline**. The **Make Baseline** dialog box appears, as shown in Figure 27.

Figure 27 Make Baseline Dialog Box

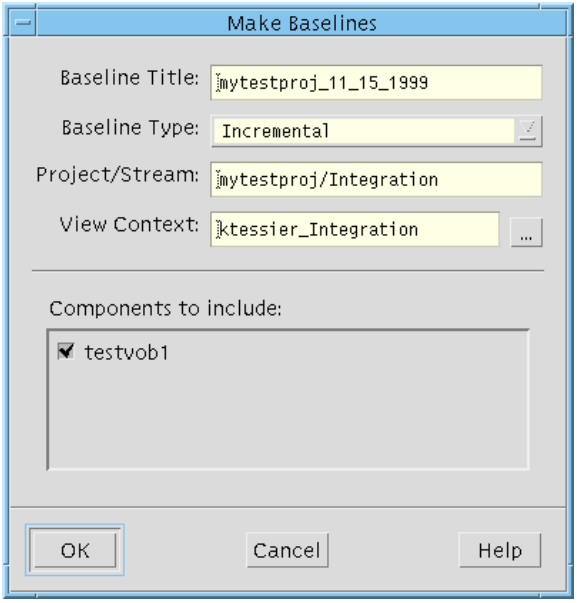

3. Enter a name in the **Baseline Title** box. By default, ClearCase names the baseline by appending the date to the project's name.

4. Choose the type of baseline to create.

   An *incremental baseline* is a baseline that ClearCase creates by recording the last full baseline and those versions that have changed since the last full baseline was created.

   A *full baseline* is a baseline that ClearCase creates by recording all versions below the component's root directory.

   Generally, incremental baselines are faster to create than full baselines; however, ClearCase can look up the contents of a full baseline faster than it can look up the contents of an incremental baseline.

5. Specify which components to include in the baseline. By default, ClearCase applies the baseline to all project components. If a component has not changed since the current baseline, ClearCase does not create a new baseline for it.

Making a Baseline For a Set of Activities

By default, all activities modified since the last baseline was made are included in the new baseline. There might be times when you want to create a baseline that includes only certain activities. To do so, use the **cleartool mkbl** command and specify the **activities** parameter. See the **mkbl** page in *ClearCase Reference Manual* for details.

Unlocking the Integration Stream

After you create a new baseline, unlock the integration stream so that developers can resume delivering work to the integration stream. To unlock the integration stream:

1. In the Project Explorer, select the integration stream.
2. Click **File>Properties** to display the integration stream's property sheet.
3. Click the **Lock** tab.
4. Click **Unlocked** and then click **OK**.

6.4 Testing the Baseline

To avoid locking the integration stream for an extended period of time, we recommend that you use a separate development stream for performing extensive testing, such as system, regression, and acceptance tests, on new baselines. See *Creating a Development Stream for Testing Baselines* on page 71 for information on creating a development stream.

After you create a new baseline and verify that it builds and passes an initial validation test in the integration stream, rebase the development stream:

1. In the Project Explorer, select the development stream and click **Tools>Rebase Stream**.

   The **Rebase Stream Preview** dialog box appears.

2. By default, ClearCase rebases your development stream to the *recommended baselines*. Because the new baseline has not been tested extensively, you probably have not yet promoted it to the level associated with recommended baselines. To rebase to the baseline, or baselines, you want to test, click **Change**.

The **Change Rebase Configuration** dialog box appears.

3. Select a component that contains a baseline you want to test. Click **Change**.

   The **Change Baseline** dialog box appears, listing all baselines for the component.

4. Select the baseline that you want to test, and click **OK**.

5. Select another component in the **Change Rebase Configuration** dialog box and repeat the process. When you finish selecting baselines, click **OK** to dismiss the **Change Rebase Configuration** dialog box.

6. Click **OK** in the **Rebase Stream Preview** dialog box to continue the rebase operation. See online help or *Developing Software with ClearCase* for details on rebasing a development stream. When you finish rebasing the development stream, you are ready to begin testing the new baselines.

Fixing Problems

If you discover a problem with a baseline while testing it, fix the affected files and deliver the changes to the integration stream as follows:

1. From the development view attached to the development stream, check out the files you need to fix. When you check out a file, you need to specify an activity.

2. Make the necessary changes to the files and check them in.

3. Build and test the changes in the development view.

4. When you are confident that the changes work, deliver the activity to the integration stream.

5. In the Project Explorer, make a new baseline that includes the fixes you delivered plus changes that other developers have delivered since you created the last baseline. See *Creating a New Baseline* on page 77.

6.5 Promoting or Demoting the Baseline

As work on your project progresses, and the quality and stability of the components improve, change the baseline's promotion level attribute to reflect a level of testing that the baseline has passed.

To promote a baseline's promotion level to the level specified for recommended baselines:

1. In the Project Explorer, select the integration stream.
2. Click **Tools>Recommend Baselines**.

To change a baseline's promotion level to something other than the level specified for recommended baselines:

1. In the Project Explorer, select the project's integration stream. Click **File>Properties** to open the integration stream's **Properties** dialog box.
2. Click the **Baselines** tab.
3. In the **Components** list, select the component that contains the baseline you want to promote. In the **Baselines** list, select the baseline. Click **Properties**. The baseline's **Properties** dialog box appears.
4. Click the arrow in the **Promotion Level** list to display all available promotion levels. Select the new promotion level.

On occasion, you may need to demote a baseline by changing its promotion level to one that is lower in the promotion level order. For example, suppose that after you create a new baseline, you discover that it contains a major bug. To prevent developers from introducing this bug to their development streams by rebasing, you can demote the baseline to a Rejected level.

6.6 Tracking the Project

ClearCase provides several tools to help you track the progress of your project. This section describes how to use those tools.

Comparing Baselines

ClearCase allows you to display the differences between two baselines graphically. To compare two baselines, use the **cleardiffbl** command. For example:

```
% cleardiffbl  newproject2_08_23_99.195 newproject2_latest.195
```

This command opens the **Compare Baselines** dialog box, as shown in Figure 28. You can also bring up the **Compare Baselines** dialog box from within the baseline's property sheet:

1. In ClearCase Project Explorer, select the integration stream, and click **File>Properties** to display the integration stream's property sheet.
2. Click the **Baselines** tab and then select the component that contains the baseline you wish to compare.
3. Select the baseline; then right-click it and select **Compare with Previous Baseline** or **Compare with Another Baseline**.

Figure 28 Comparing Baselines by Activity

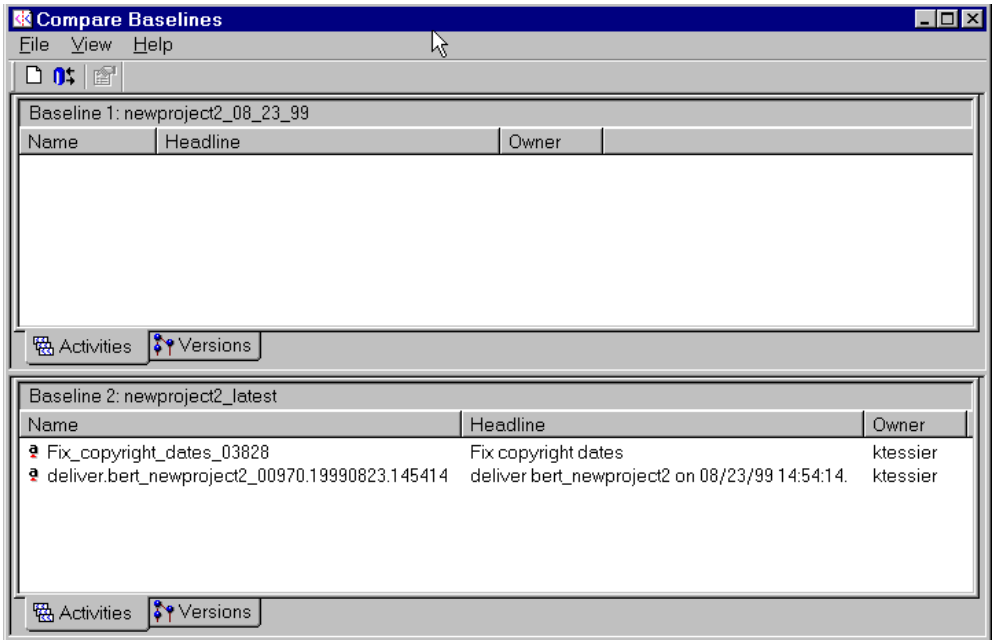

The Compare Baselines window in Figure 28 shows the results of a comparison of the **newproject2_08_23_99.195** and **newproject2_latest.195** baselines. The more recent baseline contains the **Fix copyright dates** activity. The Compare Baselines window also lists the integration activity that ClearCase created during the deliver operation.

To see the change sets associated with the activities, click **Versions**. Figure 29 shows the versions associated with the **Fix copyright dates** and integration activities.

Figure 29 Comparing Baselines by Version

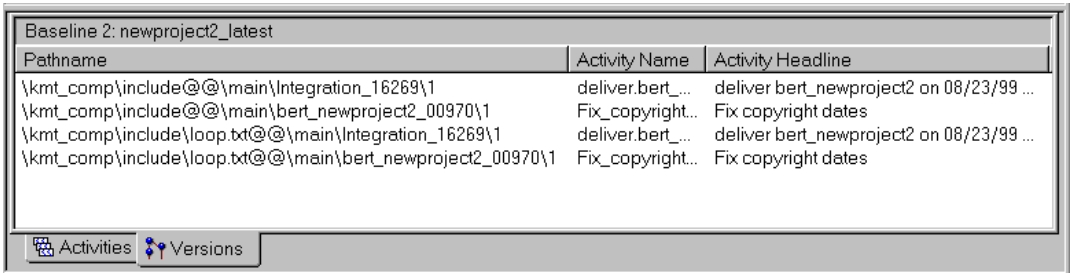

Querying ClearQuest User Databases

If you use the UCM-ClearQuest integration, you can use ClearQuest queries to retrieve information about the state of your project. When you create or upgrade a ClearQuest user database to use a UCM-enabled schema, the integration installs six queries in two subfolders of the Public Queries folder in the user database's workspace. These queries make it easy for developers to see which activities are assigned to them and for project managers to see which activities are active in a particular project. Table 4 lists and describes the queries.

Table 4 Queries in UCM-Enabled Schema

Query	Description
ActiveForProject	For one or more specified projects, selects all activities in an active state type.
ActiveForStream	For one or more specified streams, selects all activities in an active state type.
ActiveForUser	For one or more specified developers, selects all assigned activities in an active state type.

Table 4 Queries in UCM-Enabled Schema

Query	Description
MyToDoList	Selects all activities in an active or ready state type assigned to the developer running the query.
UCMProjects	Selects all projects linked to the ClearQuest user database.
UCMCustomQuery1	This query is not intended to be used by users; the integration uses it. When a developer checks out or checks in a file, or adds a file to source control and is prompted to select an activity, the integration calls this query to display the list of activities available in the stream associated with the developer's view.

You can also create your own queries by clicking **Query>New Query** within the ClearQuest client. In the **Choose a record type** dialog box that appears, select **All_UCM_Activities** if you want the query to search all UCM-enabled record types.

6.7 Cleaning Up the Project

When your team finishes work on a project and releases or deploys the new software, you should clean up the project environment before creating the next version of the project. Cleaning up involves removing any unused objects, and locking and hiding the project and its streams. This process reduces clutter and makes it easier to navigate in the Project Explorer.

Removing Unused Objects

During the life of the project, you or a developer might create an object and then decide not to use it. Perhaps you decide to use a different naming convention and you create a new object instead of renaming the existing one. To avoid confusion and reduce clutter, remove these unused objects.

To delete a project, stream, component, or activity, select the object in the Project Explorer, and click **File>Delete**. To delete a baseline, use the **cleartool rmbl** command.

Projects

You can delete a project only if it does not contain any streams. When you create a project with the Project Creation Wizard, the wizard also creates an integration stream. Therefore, you can delete a project only if you created it with the **cleartool mkproject** command, or if you first delete the integration stream. For more information on removing projects, see the **rmproject** reference page in *ClearCase Reference Manual*.

Streams

You can delete a development stream or an integration stream only if all of the following conditions are true:

- The stream contains no activities.
- No baselines have been created in the stream.
- No views are attached to the stream.

In addition, you cannot delete an integration stream if the project contains any development streams. For more information on removing streams, see the **rmstream** reference page in *ClearCase Reference Manual*.

Components

You can delete a component only if all of the following conditions are true:

- No baselines of the component other than its initial baseline exist.
- The component's initial baseline does not serve as a *foundation baseline* for another stream.

For more information on removing components, see the **rmcomp** reference page in *ClearCase Reference Manual*.

Baselines

You can delete a baseline only if all of the following conditions are true:

- The baseline does not serve as a foundation baseline.
- The baseline is not a component's initial baseline.
- A stream has not made changes to the baseline.
- The baseline is not used as the basis for an incremental baseline.

For more information on removing baselines, see the **rmbl** reference page in *ClearCase Reference Manual*.

Activities

You can delete an activity only if both of the following conditions are true:

- The activity has no versions in its change set.
- No view is currently set to the activity.

For more information on removing activities, see the **rmactivity** reference page in *ClearCase Reference Manual*.

Locking and Making Obsolete the Project and Streams

To prevent a project or a stream from appearing in the Project Explorer, lock the object and use the obsolete option. The obsolete option hides the object.

1. In the Project Explorer, select the stream or project that you want to hide, and click **File>Properties** to display its property sheet.
2. Click the **Lock** tab, and select **Obsolete**. Click **OK**.

To see objects that you have made obsolete, click **View>Show Obsolete Items** in the Project Explorer.

7 Managing Parallel Releases of Multiple Projects

The previous chapters describe how to manage a single project. However, you may need to manage multiple releases of a project simultaneously. To do so, you need to merge changes from one project to another. This chapter describes how to accomplish that merging in two common scenarios:

- Managing a current project and a follow-on project simultaneously
- Incorporating a patch release into a new release of the project

This chapter also describes other scenarios in which you can use these merging techniques between projects.

7.1 Managing a Current Project and a Follow-on Project Simultaneously

Given the tight software development schedules that most organizations operate within, it is common practice to begin development of the next release of a project before work on the current release is completed. The next release may add new features, or it may involve porting the current release to a different platform.

Example

Figure 30 illustrates the flow of a current project, Webotrans 4.0, and a follow-on project, Webotrans 4.1.

Figure 30 Managing a Follow-on Release

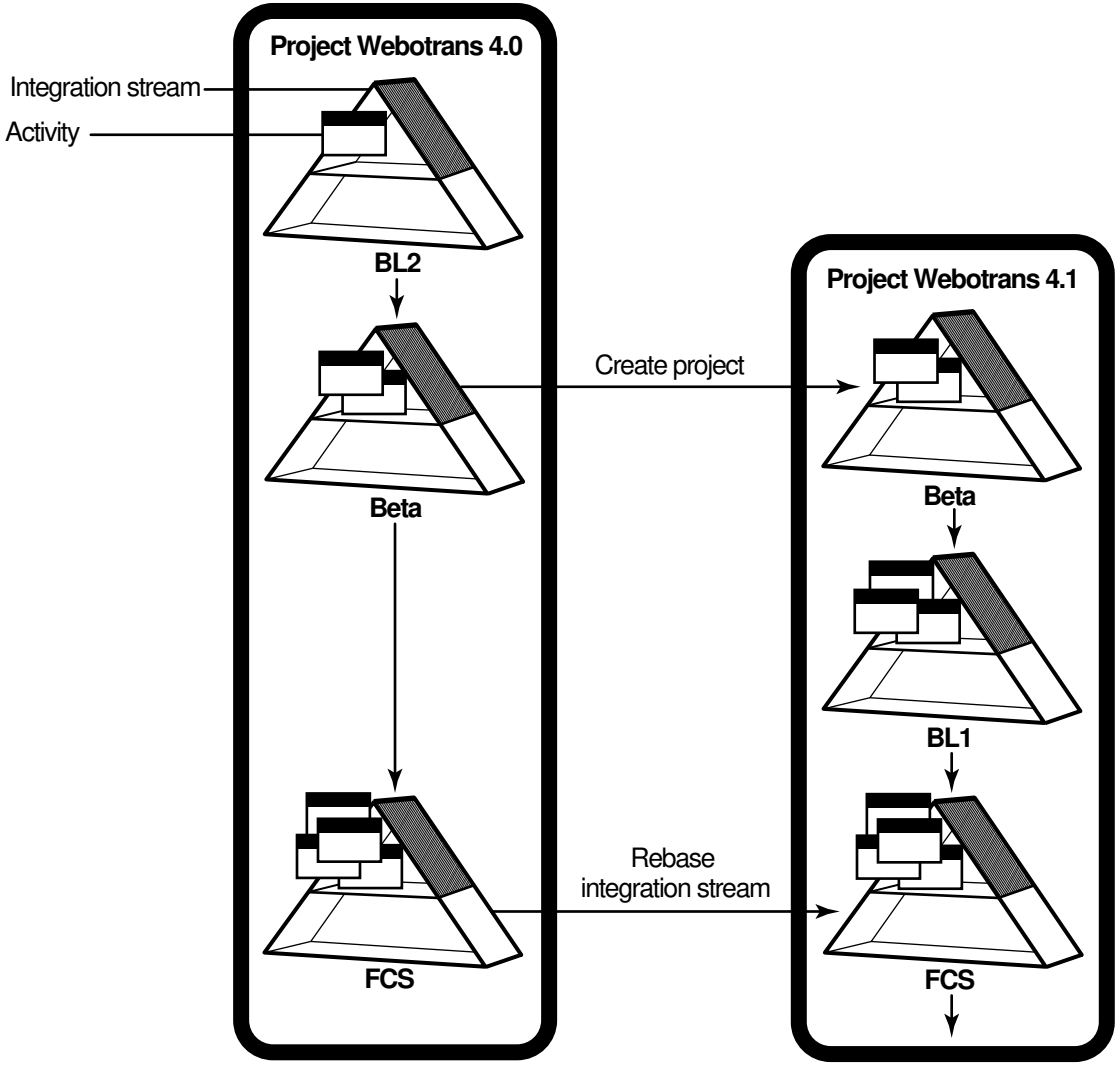

In this example:

- The project manager for the follow-on project created the Webotrans 4.1 project based on the Beta baselines of the components used in the Webotrans 4.0 project. Developers on both project teams then continued to make changes, and the 4.0 and 4.1 project managers continued to create new baselines that incorporate those changes.
- When the 4.0 team completed its work, the project manager created the final baselines, named FCS. The 4.1 project manager then rebased the 4.1 integration stream to the FCS baselines.

Performing Interproject Rebase Operations

To rebase an integration stream to a set of baselines in another project's integration stream:

1. Navigate to an integration view attached to the integration stream that you want to rebase.
2. For each component, issue the **cleartool rebase** command, specifying the component's baseline. For example:

   ```
   % cleartool rebase –baseline FCS.195 –gmerge
   Changed config spec for view "webotrans4.1_integration" to reflect its
   stream's new configuration.
   Build and test are necessary to ensure that the merges were completed
   correctly.

   When build and test are confirmed, run "cleartool rebase —resume —
   complete".
   ```

3. ClearCase merges nonconflicting changes automatically. You must resolve the changes that ClearCase cannot merge automatically. The **–gmerge** option directs ClearCase to start its Diff Merge graphical tool to help you resolve conflicting changes. For details on using Diff Merge, see the Diff Merge online help and *Developing Software with ClearCase*.
4. When you finish the merge, build and test the changes before completing the rebase operation.
5. Complete the rebase operation. For example:

   ```
   % cleartool rebase –resume –complete
   ```

In the example shown above, **FCS.195** is the full name of the baseline for one of the components in the Webotrans 4.0 integration stream. To determine a baseline's full name:

1. In the Project Explorer, select the integration stream.
2. Click **File>Properties**. The integration stream's property sheet appears.
3. Click the **Baselines** tab.
4. In the **Components** list, select the component that contains the desired baseline.
5. In the **Baselines** list, select the root name of the baseline. Click **Properties**. The baseline's property sheet appears. The **Name** box identifies the full baseline name.

Note that you can rebase your project's integration stream only if the baseline to which you are rebasing is a successor of your integration stream's current *foundation baseline*. In the above example, the FCS baseline is a successor to the Beta baseline, which is the current foundation baseline for the Webotrans 4.1 integration stream.

7.2 Incorporating a Patch Release into a New Version of the Project

Another common parallel development scenario involves working on a patch release and a new release of a project at the same time. This section describes this scenario.

Example

Figure 31 illustrates the flow of a patch release and a new release. In this example:

- Both the Webotrans 3.0 Patch and Webotrans 4.0 projects use the FCS baselines of the components in the Webotrans 3.0 project as their *foundation baselines*. The purpose of the patch release is to fix a problem detected after Webotrans 3.0 was released. Webotrans 4.0 represents the next major release of the Webotrans product.
- Development continues in both the 3.0 Patch and 4.0 projects, with the project managers creating baselines periodically.
- The developers working on the 3.0 Patch project finish their work, and the project manager incorporates the final changes in the BL2 baseline. The project manager then needs to merge those changes from the 3.0 Patch integration stream to the 4.0 integration stream so that the 4.0 project contains the fix.

Figure 31 Incorporating a Patch Release

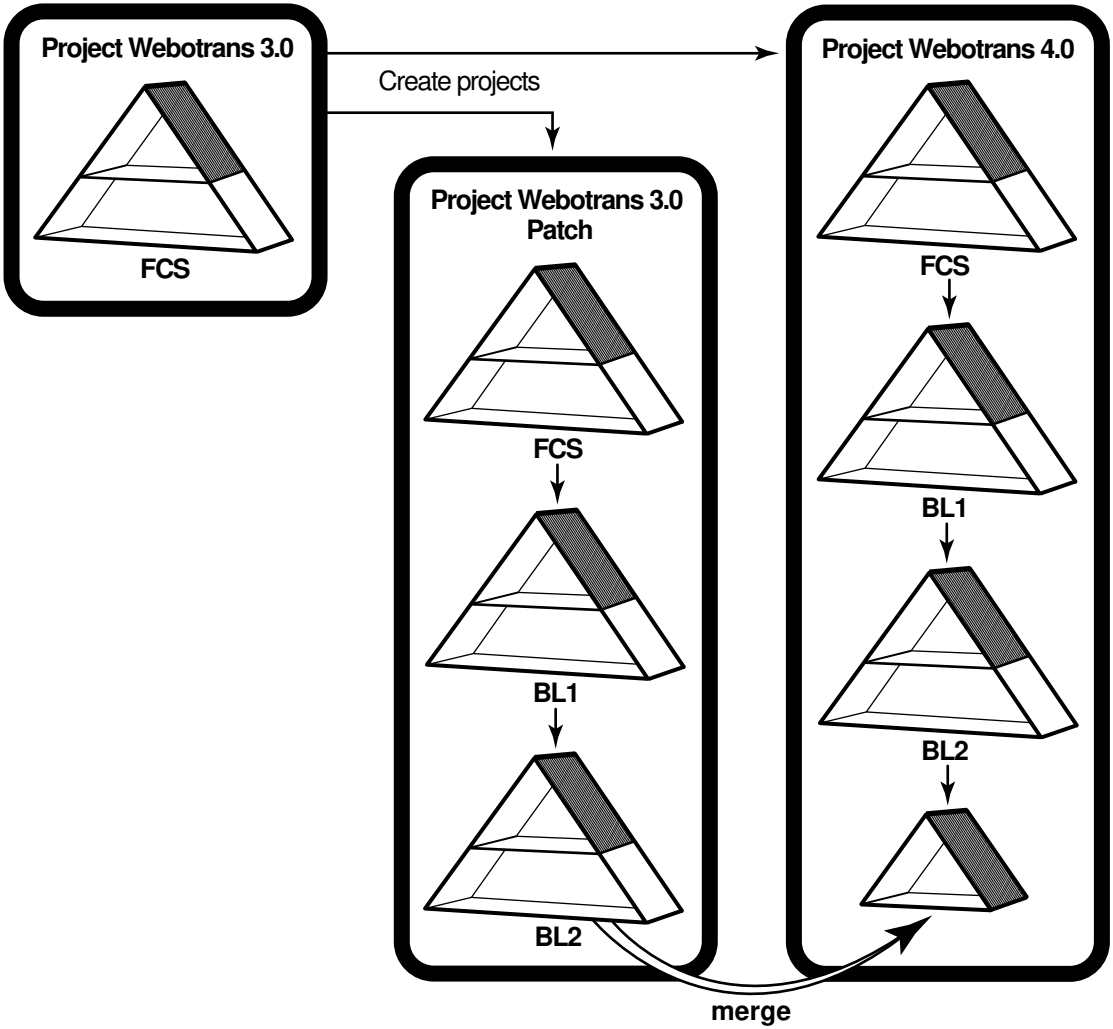

Merging Work to Another Project

UCM does not support interproject deliver operations. However, you can simulate a deliver operation by running a script such as the one shown here, which uses base ClearCase functionality to merge changes.

```
# Sample Perl script for delivering contents of one UCM project to another,
# or to a nonUCM project. Run this script while set to the integration
# view of the destination project.
#
# Usage: Perl <this-script> <project-name> <project-vob>

use strict;

my $mergeopts = '-print';
my $project = shift @ARGV;
my $pvob = shift @ARGV;
my $bl;

chdir ($pvob) or die("can't cd to project VOB '$pvob'");

print("######## Getting recommended baselines for project '$project'\n");

my @recbls = split(' ', `cleartool lsproject -fmt "%[rec_bls]p" $project`);

foreach $bl (@recbls) {

     my $comp = `cleartool lsbl -fmt "%[component]p" $bl`;
     my $vob = `cleartool lscomp -fmt "%[root_dir]p" $comp`;

     print("######## Merging changes from baseline '$bl' of $vob\n");

     my $st = system("cleartool findmerge $vob -fver $bl $mergeopts");
     $st == 0 or die("findmerge error");
}

exit 0;
```

The script finds the recommended baselines for the integration stream from which you are merging. It then uses the **cleartool findmerge** command to find differences between the versions represented by those recommended baselines and the latest versions in the target integration stream. For details on **findmerge**, see the **findmerge** reference page.

We recommend that you add error handling and other logic appropriate for your site to this script before using it.

7.3 Additional Merging Scenarios

This section describes two additional scenarios for which you may want to use a script similar to the one shown in *Merging Work to Another Project* on page 94.

Merging from a Project to a Non-UCM Branch

You may be in a situation in which part of the development team works in a project, and the rest of the team works in base ClearCase. If you are a longtime ClearCase user, you may decide to use UCM initially on a small part of your system. This approach would allow you to migrate from base ClearCase to UCM gradually, rather than all at once.

In this case, you need to merge work periodically from the project's integration stream to the branch that serves as the integration branch for the system. To do so, use a script similar to the one shown in *Merging Work to Another Project* on page 94.

Merging to a System Project

If you have a very large system, you may decide to use multiple projects with each project team working on a different part of the system. In this scenario, you would typically create an additional project to serve as the system project. The system project is the integration location for the system. When project managers for each of the projects create stable baselines, the system's project manager merges those baselines to the system project's integration stream.

To do this merging, use the script as shown in *Merging Work to Another Project* on page 94.

Part 2: Working in Base ClearCase

The following chapters describe how to use base ClearCase features to set up and manage a customized development environment for your project team.

8 Managing Projects in Base ClearCase

As a project manager, you are responsible for planning, staffing, and managing the technical aspects of a software development project. You decide what will be worked on, assign work to the project's team members, establish the work schedule, and perhaps the policies and procedures for doing the work.

When development is underway, you monitor progress and generate project status reports. You may also approve the specific work items included in a build and subsequently a baseline.

You may also be the project integrator, responsible for incorporating work that each developer completes into a deliverable and buildable system. You create the project's baselines and establish the quality level of those baselines.

Base ClearCase offers many features to make this work easier. Before development begins, you need to complete several planning and setup tasks:

- Setting up the project environment
- Implementing development policies
- Defining and implementing an integration policy

This chapter introduces these topics. The remaining chapters cover the implementation details. Chapter 13, *Using ClearCase Throughout the Development Cycle*, follows a project throughout the development cycle to show how you can use ClearCase.

Before reading this part of the manual, read *Developing Software with ClearCase* to become familiar with the concepts of VOBs, views, and config specs.

8.1 Setting Up the Project

This section describes the planning and setup work you need to do before development begins.

Creating and Populating VOBs

If your project is migrating to ClearCase from another version control product or is adopting a configuration and change management plan for the first time, you must populate the VOBs for your project with an initial collection of data (file and directory elements). If your site has a dedicated ClearCase administrator, he or she may be responsible for creating and maintaining VOBs, but not for importing data into them.

Administering ClearCase includes detailed information on these topics.

Planning a Branching Strategy

ClearCase uses branches to enable parallel development. A *branch* is an object that specifies a linear sequence of versions of an element. Every element has one *main branch*, which represents the principal line of development, and may have multiple *subbranches,* each of which represents a separate line of development. For example, a project team can use two branches concurrently: the **main** branch for new development work and a subbranch to fix a bug. The aggregated **main** branches of all elements constitutes the **main** branch of a code base.

Subbranches can have subbranches. For example, a project team designates a subbranch for porting a product to a different platform; the team then decides to create a bug-fixing subbranch off that porting subbranch. ClearCase allows you to create complex branch hierarchies. See Figure 1 for an illustration of a multilevel branching hierarchy. As a project manager in such an environment, you need to ensure that developers are working on the correct branches. To do that, you must tell them which rules to include in their *config specs* so that their views access the appropriate set of versions.

Chapter 9, *Defining Project Views*, describes config specs and branches in detail. Before you read it, a little background on branching strategies may be helpful.

Branching policy is influenced by the development objectives of the project and provides a mechanism to control the evolution of the code base. There are as many variations of branching policy as organizations that use ClearCase. But there are also similarities that reflect common adherence to best practices. Some of the more common branch types and uses are presented here.

Task branches are short-lived, typically involve a small percentage of files, and are merged into their parent branch after the task is completed. Task branches promote accountability by leaving a permanent audit trail that associates a set of changes with a particular task; they also make it easy to identify the task artifacts, such as views and derived objects, that can be removed when they are no longer needed. If individual tasks don't require changes to the same files, it is easy to merge a task branch to its parent.

Private development branches are useful when a group of developers need to make a more comprehensive set of changes on a common code base. By branching as much of the **main** branch as needed, developers can work in isolation as long as necessary. Merging back to the **main** branch can be simplified if, before merging, each developer merges the **main** branch to the private branch to resolve any differences there before checking in the changed files.

Integration branches provide a buffer between private development branches and the **main** branch and can be useful if you delegate the integration task to one person, rather than making developers responsible for integrating their own work.

Branch Names

It's a good idea to establish naming conventions that indicate the work the branch contains. For example, **rel2.1_main** is the branch on which all code for Release 2.1 ultimately resides, **rel2.1_feature_abc** contains changes specific to the ABC feature, and **rel2.1_bl2** is the second stable baseline of Release 2.1 code. (If necessary, branch names can be much longer and more descriptive, but long branch names can crowd a version tree display.)

NOTE: Make sure that you do not create a branch type with the same name as a label type. This can cause problems when config specs use labels in version selectors. For example, make all branch names lowercase, and make all label names uppercase.

Branches and ClearCase MultiSite

Branches are particularly important when your team works in VOBs that have been replicated to other sites with the ClearCase MultiSite product. Developers at different sites work on different branches of an element. This scheme prevents collisions—for example, developers at two sites creating version **/main/17** of the same element. In some cases, versions of files cannot or should not be merged, and developers at different sites must share branches. For more information, see *Policy: Certain Branches Are Shared Among MultiSite Sites* on page 140.

Creating Shared Views and Standard Config Specs

As a project manager, you want to control the config specs that determine how branches are created when developers check out files. There are several ways to handle this task:

- Create a config spec template that each developer must use. Developers can either paste the template into their individual config specs or use the ClearCase include file facility to get the config spec from a common source.
- Create a view that developers will share. This is usually a good way to provide an integration view for developers to use when they check in work that has evolved in isolation on a private branch.

  NOTE: Working in a single shared view is not recommended because doing so can degrade system performance.
- To ensure that all team members configure their views the same way, you can create files that contain standard config specs. For example:
  - **/public/config_specs/ABC** contains the ABC team's config spec
  - **/public/config_specs/XYZ** contains the XYZ team's config spec

Store these config spec files in a standard directory outside a VOB, to ensure that all developers get the same version.

Recommendations for View Names

You may want to establish naming conventions for views for the same reason that you do for branches: it is easier to associate a view with the task it is used for. For example, you can require all view names (called view-tags) to include the owner's name and the task (**bill_V4.0_bugfix**) or the name of the machine hosting the view (**platinum_V4.0_int**).

8.2 Implementing Development Policies

To enforce development policies, you can create ClearCase *metadata* to preserve information about the status of versions. To monitor the progress of the project, you can generate a variety of reports from this data and from the information that ClearCase captures in event records.

Using Labels

A label is a user-defined name that can be attached to a version. Labels are a powerful tool for project managers and system integrators. By applying labels to groups of elements, you can define and preserve the relationship of a set of file and directory versions to each other at a given point in the development lifecycle. For example, you can apply labels to these versions:

- All versions considered stable after integration and testing. Use this baseline label as the foundation for new work.
- All versions that are partially stable or contain some usable subset of functionality. Use this checkpoint label for intermediate testing or as a point to which development can be rolled back in the event that subsequent changes result in regressions or instability.
- All versions that contain changes to implement a particular feature or that are part of a patch release.

Using Attributes, Hyperlinks, Triggers, and Locks

Attributes are name/value pairs that allow you to capture information about the state of a version from various perspectives. For example, you can attach an attribute named **CommentDensity** to each version of a source file, to indicate how well the code is commented. Each such attribute can have the value **unacceptable**, **low**, **medium**, or **high**.

Hyperlinks allow you identify and preserve relationships between elements in one or more VOBs. This capability can be used to address process-control needs, such as requirements tracing, by allowing you to link a source file to a requirements document.

Triggers allow you to control the behavior of **cleartool** commands and ClearCase operations by arranging for a specific program or executable script to run before or after the command executes. Virtually any operation that modifies an element can fire a trigger. Special environment variables make the relevant information available to the script or program that implements the procedure.

Preoperation triggers fire before the designated ClearCase command is executed. A preoperation trigger on **checkin** can prompt the developer to add an appropriate comment. Postoperation triggers fire after a command has exited and can take advantage of the command's exit status. For example, a postoperation trigger on the **checkin** command can send an e-mail message to the QA department, indicating that a particular developer modified a particular element.

Triggers can also automate a variety of process management functions. For example:

- Applying attributes or attaching labels to objects when they are modified
- Logging information that is not included in the ClearCase event records
- Initiating a build and/or source code analysis whenever particular objects are modified

For more information on these mechanisms, see Chapter 10, *Implementing Project Development Policies*.

A lock on an element or directory prevents all developers (except those included on an exception list) from modifying it. Locks are useful for implementing temporary restrictions. For example, during an integration period, a lock on a single object—the **main** branch type—prevents all users who are not on the integration team from making any changes.

The effect of a lock can be small or large. A lock can prevent any new development on a particular branch of a particular element; another lock can apply to the entire VOB, preventing developers from creating any new element of type **compressed_file** or using the version label **RLS_1.3**.

Locks can also be used to retire names, views, and VOBs that are no longer used. For this purpose, the locked objects can be tagged as *obsolete*, effectively making them invisible to most commands.

Global Types

The ClearCase global type facility makes it easy for you to ensure that the branch, label, attribute, hyperlink, and element types they need are present in all VOBs your project uses. The manual *Administering ClearCase* has more information about creating and using global types.

Generating Reports

ClearCase creates and stores an event record each time an element is modified or merged. Many ClearCase commands include selection and filtering options that you can use to create reports based on these records. The scope of such reports can cover a single element, for a set of objects, or for entire VOBs.

Chapter 10, *Implementing Project Development Policies*, provides more detail on using event records and metadata to implement project policies. Event records and other metadata can also

be useful if you need to generate reports on activities managed by ClearCase (for example, the complete history of changes to an element). ClearCase provides a variety of report-generation tools. For more information on this topic, see the **fmt_ccase** reference page in the *ClearCase Reference Manual*.

8.3 Integrating Changes

During the lifetime of a project, the contents of individual elements diverge as they are branched and usually converge in a merge operation. Typically, the project manager periodically merges most branches back to the **main** branch to ensure that the code base maintains a high degree of integrity and to have a single "latest" version of each element from which new versions can safely branch. Without regular merges, the code base quickly develops a number of dangling branches, each with slightly different contents. In such situations, a change made to one version must be propagated by hand to other versions, a tedious process that is prone to error.

As a project manager, you must establish merge policies for your project. Typical policies include the following:

- Developers merge their changes to the **main** branch. This can work well when the number of developers and/or the number of changed files is small and the developers are familiar with the mechanics of merging. Developers must also understand the nature of other changes they may encounter when the merge target is not the immediate predecessor of the version being merged, which happens when several developers are working on the same file in parallel.

- Developers merge their changes to an integration branch. This provides a buffer between individual developers' merges and the **main** branch. The project manager or system integrator then merges the integration branch to the **main** branch.

- Developers must merge from the **main** branch to their development branch before merging to the **main** branch or integration branch. This type of merge promotes greater stability by forcing merge-related instability to the developers' private branches where problems can be resolved before they affect the rest of the team.

- The project manager designates "slots" for developer merges to the **main** branch. This is a variation on several of the mechanisms already described. It provides an additional level of control in situations where parallel development is going on.

There are other scenarios as well. Chapter 11, *Integrating Changes*, describes merging in detail.

9 Defining Project Views

This chapter explains how config specs work and provides sample config specs useful for project development work, for nondevelopment tasks such as monitoring progress and doing research, and for running project builds. It also explains how to share config specs between Windows and UNIX systems.

9.1 How Config Specs Work

When you create views for your project, you must prepare one or more *config specs* (configuration specifications). Config specs allow you to achieve the degree of control that you need to have over project work by controlling which versions developers see and what operations they can perform in specific views. You can narrow a view to a specific branch or open it to an entire VOB. You can also disallow checkouts of all selected versions or restrict checkouts to specific branches.

A config spec contains a series of rules that ClearCase uses to select the versions that appear in the view. When team members use a view, they see the versions that match at least one of the rules in the config spec. ClearCase searches the version tree of each element for the first version that matches the first rule in the config spec. If no versions match the first rule, ClearCase searches for a version that matches the second rule. If no versions of an element match any rule in the config spec, no versions of the element appear in the view.

The order in which rules appear in the config spec determine which version of a given element is selected. The various examples in this chapter examine this behavior in different contexts. For details about preparing config specs, see the **config_spec** reference page.

9.2 Default Config Spec

This config spec defines a dynamic configuration, which selects changes made on the **main** branch of every element throughout the entire source tree, by any developer:

```
(1)  element * CHECKEDOUT
(2)  element * /main/LATEST
```

This is the *default config spec*, to which each newly created view is initialized. When you create a view with the **mkview** command, the contents of file **default_config_spec** (located in *ccase-home-dir*) become the new view's config spec.

A view with this config spec provides a private work area that selects your checked-out versions (Rule 1). By default, when you check out a file, you check out from the latest version on the **main** branch (Rule 2). While an element is checked out to you, you can change it without affecting anyone else's work. As soon as you check in the new version, the changes are available to developers whose views select **/main/LATEST** versions.

The view also selects all other elements (that is, all elements that you have not checked out), on a read-only basis. If another user checks in a new version on the **main** branch of such an element, the new **LATEST** version appears in this *dynamic view* immediately.

By default, *snapshot views* also include the two *version-selection rules* shown above. In addition, snapshot view config specs include *load rules*, which specify which elements or subtrees to load into the snapshot view. See *Developing Software with ClearCase* for details on creating snapshot views.

The Standard Configuration Rules

The two configuration rules in the default config spec appear in many of this chapter's examples. The **CHECKEDOUT** rule allows you to modify existing elements. If you try to check out elements in a view that omits this rule, you can do so, but **cleartool** generates this warning:

```
% cleartool checkout –nc cmd.c
cleartool: Warning: Unable to rename "cmd.c" to "cmd.c.keep": Read-only
filesystem.
cleartool: Error: Checked out version, but could not copy to "cmd.c": File
exists.
Correct the condition, then uncheckout and re-checkout the element.
cleartool: Warning: Copied checked out version to "cmd.c.checkedout".
cleartool: Warning: Checked-out version is not selected by view.
Checked out "cmd.c" from version "/main/7".
```

In this example, the config spec continues to select version 7 of element **cmd.c**, which is read-only. A read-write copy of this version, **cmd.c.checkedout**, is created in view-private storage. (This is not a recommended way of working.)

The **/main/LATEST** rule selects the most recent version on the **main** branch to appear in the view.

In addition,a **/main/LATEST** rule is required to create new elements in a view. If you create a new element when this rule is omitted, your view cannot "see" that element. (Creating an element involves creating a **main** branch and an empty version, **/main/0**).

Omitting the Standard Configuration Rules

It makes sense to omit one or both of the standard configuration rules only if a view is not going to be used to modify data. For example, you can configure a *historical* view, to be used only for browsing old data. Similarly, you can configure a view in which to compile and test only or to verify that sources have been labeled properly.

9.3 Config Spec Include Files

ClearCase supports an include file facility that makes it easy to ensure that all team members are using the same config spec. For example, the configuration rules in this config spec can be placed in file **/public/c_specs/major.csp**. Each developer then needs a one-line config spec:

(1)
```
include /public/c_specs/major.csp
```

NOTE: If you are sharing config specs between UNIX and Windows NT computers where the VOB-tags are different, you must have two sources, or you must store the config spec in a UNIX directory that is accessible from both machines.

If you want to modify this config spec (for example, to adopt the no-directory-branching policy), only the contents of **major.csp** need to change. You can use this command to reconfigure your view with the modified config spec:

```
% cleartool setcs –current
```

9.4 Project Environment for Sample Config Specs

You can use different config specs for different kinds of development and management tasks. The three sections that follow present sample config specs useful for various aspects of project development, project management and research, and project builds. This section presents the development environment that these config specs are based on.

Developers use a VOB whose VOB-tag is **/vobs/monet**, which has this structure:

/vobs/monet	*(VOB-tag, VOB mount point)*
src/	*(C language source files)*
include/	*(C language header files)*
lib/	*(project's libraries)*

For the purposes of this chapter, suppose that the **lib** directory has this substructure:

lib/

libcalc.a	*(checked-in staged version of library)*
libcmd.a	*(checked-in staged version of library)*
libparse.a	*(checked-in staged version of library)*
libpub.a	*(checked-in staged version of library)*
libaux1.a	*(checked-in staged version of library)*
libaux2.a	*(checked-in staged version of library)*
libcalc/	*(sources for calc library)*
libcmd/	*(sources for cmd library)*
libparse/	*(sources for parse library)*
libpub/	*(sources for pub library)*
libaux1/	*(sources for aux1 library)*
libaux2/	*(sources for aux2 library)*

Sources for libraries are located in subdirectories of **lib.** After a library is built in its source directory, it can be staged to **/vobs/monet/lib**. The build scripts for the project's executable programs can instruct the link editor, **ld(1)**, to use the libraries in this directory (the library staging area) instead of a more standard location (for example, **/usr/local/lib**).

The following labels are assigned to versions of **monet** elements.

Version Labels	Description
R1.0	First customer release
R2_BL1	Baseline 1 prior to second customer release
R2_BL2	Baseline 2 prior to second customer release
R2.0	Second customer release

These version labels have been assigned to versions on the **main** branch of each element. Most project development work takes place on the **main** branch. For some special tasks, development takes places on a subbranch.

Subbranches	Description
major	Used for work on the application's graphical user interface, certain computational algorithms, and other major enhancements
r1_fix	Used for fixing bugs in Release 1.0

9.5 Views for Project Development

The config specs in this section are useful for project development because they enforce various branching policies.

View for New Development on a Branch

You can use this config spec for work to be isolated on branches named **major**:

```
(1)  element * CHECKEDOUT
(2)  element * .../major/LATEST
(3)  element * BASELINE_X -mkbranch major
(4)  element * /main/LATEST -mkbranch major
```

In this scheme, all checkouts occur on branches named **major** (Rule 2).

major branches are created at versions that constitute a consistent baseline: a major release, a minor release, or a set of versions that produces a working version of the application. In this config spec, the baseline is defined by the version label **BASELINE_X**.

Variation That Uses a Time Rule

Sometimes, other developers check in versions that become visible in your view, but are incompatible with your own work. In such cases, you can continue to work on sources as they existed before those changes were made. For example, Rule 2 in this config spec selects the latest version on the main branch as of 4:00 P.M. on November 12:

```
(1)  element * CHECKEDOUT
(2)  element * /major/LATEST -time 12-Nov.16:00
(3)  element * BASELINE_X -mkbranch major
(4)  element * /main/LATEST -mkbranch major
```

Note that this rule has no effect on your own checkouts.

View to Modify an Old Configuration

This config spec allows developers to modify a configuration defined with version labels:

```
(1)  element * CHECKEDOUT
(2)  element * .../r1_fix/LATEST
(3)  element * R1.0 -mkbranch r1_fix
```

Note the following:

- Elements can be checked out (Rule 1).
- The **checkout** command creates a branch named **r1_fix** at the initially selected version (the *auto-make-branch* clause in Rule 3).

A key aspect of this scheme is that the same branch name, **r1_fix**, is used in every modified element. The only administrative overhead is the creation of a single branch type, **r1_fix**, with the **mkbrtype** command.

This config spec is efficient: two rules (Rules 2 and 3) configure the appropriate versions of all elements:

- For elements that have been modified, this version is the most recent on the **r1_fix** subbranch (Rule 2).
- For elements that have not been modified, this version is the one labeled **R1.0** (Rule 3).

Figure 32 illustrates these elements. The **r1_fix** branch is a subbranch of the **main** branch. But Rule 2 handles the more general case, too: the **...** wildcard allows the **r1_fix** branch to occur anywhere in any element's version tree, and at different locations in the version trees of different elements.

Figure 32 Making a Change to an Old Version

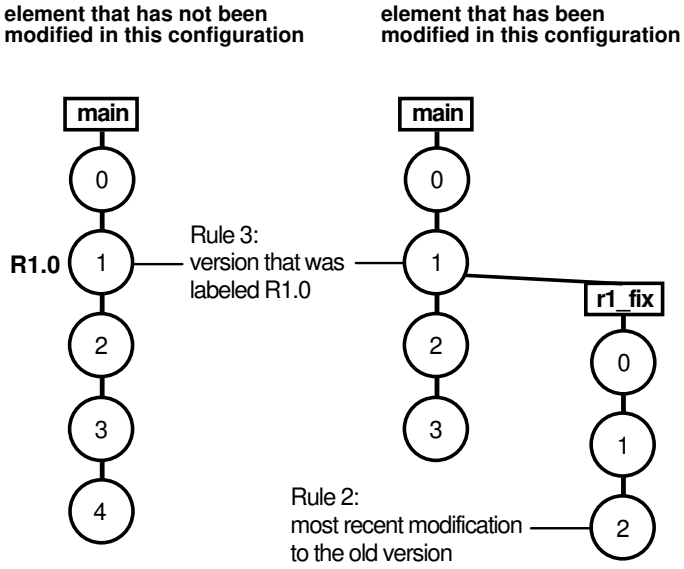

Omitting the /main/LATEST Rule

The config spec in *View to Modify an Old Configuration* on page 112 omits the standard **/main/LATEST** rule. This rule is not useful for work with VOBs in which the version label **R1.0** does not exist. In addition, it is not useful in situations where new elements are created. If your development policy is to not create new elements during maintenance of an old configuration, the absence of a **/main/LATEST** rule is appropriate.

To allow creation of new elements during the modification process, add a fourth configuration rule:

```
(1)  element * CHECKEDOUT
(2)  element * /main/r1_fix/LATEST
(3)  element * R1.0 -mkbranch r1_fix
(4)  element * /main/LATEST -mkbranch r1_fix
```

When a new element is created with **mkelem**, the **–mkbranch** clause in Rule 4 causes the new element to be checked out on the **r1_fix** branch (which is created automatically). This rule conforms to the scheme of localizing all changes to **r1_fix** branches.

Variation That Uses a Time Rule

This baseline configuration is defined with a **–time** rule.

```
(1)  element * CHECKEDOUT
(2)  element * /main/r1_fix/LATEST
(3)  element * /main/LATEST -time 4-Sep:02:00 -mkbranch r1_fix
```

View to Implement Multiple-Level Branching

This config spec implements and enforces consistent multiple-level branching.

```
(1)  element * CHECKEDOUT
(2)  element * .../major/autumn/LATEST
(3)  element * .../major/LATEST -mkbranch autumn
(4)  element * BASELINE_X -mkbranch major
(5)  element * /main/LATEST -mkbranch major
```

A view configured with this config spec is appropriate in the following situation:

- All changes from the baseline designated by the **BASELINE_X** version label must be made on a branch named **major**.
- Moreover, you are working on a special project, whose changes are to be made on a subbranch of **major**, named **autumn**.

Figure 33 shows what happens in such a view when you check out an element that has not been modified since the baseline.

Figure 33 Multiple-Level Auto-Make-Branch

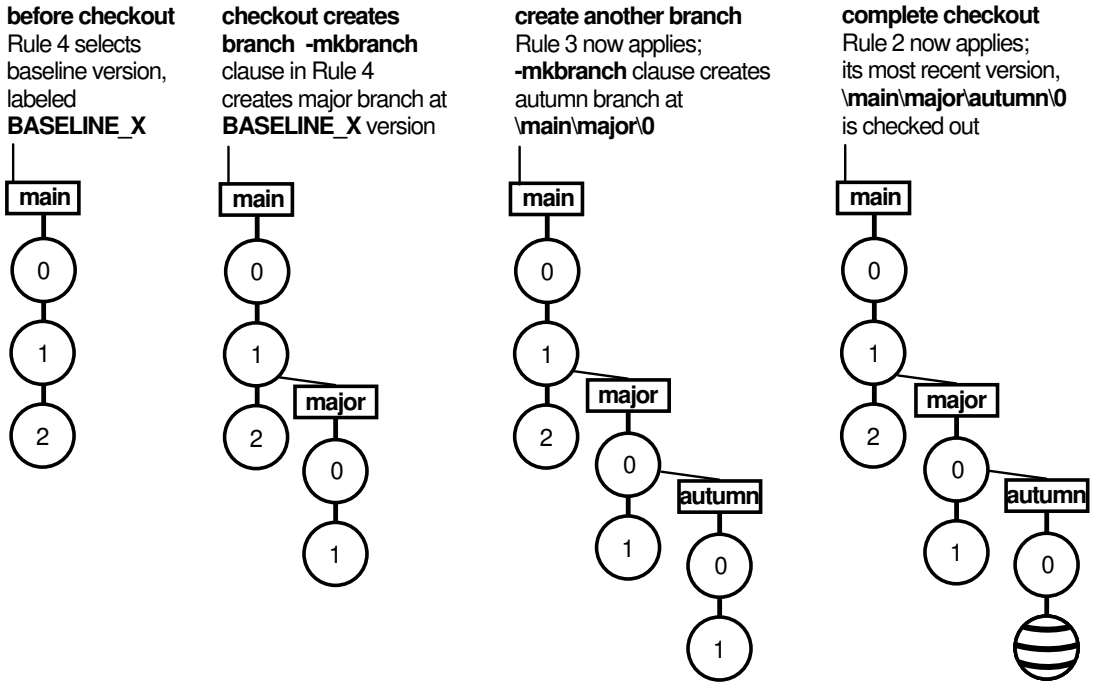

For more on multiple-level branching, see the **config_spec** and **checkout** reference pages.

View to Restrict Changes to a Single Directory

This config spec is appropriate for a developer who can make changes in one directory only, **/vobs/monet/src**:

```
(1)  element * CHECKEDOUT
(2)  element src/* /main/LATEST
(3)  element * /main/LATEST -nocheckout
```

The most recent version of each element is selected (Rules 2 and 3), but Rule 3 prevents checkouts to all elements except those in the directory specified.

Note that Rule 2 matches elements in any directory named **src**, in any VOB. The pattern **/vobs/monet/src/*** restricts matching to only one VOB.

This config spec can be extended easily with additional rules that allow additional areas of the source tree to be modified.

9.6 Views to Monitor Project Status

The config specs presented here are useful for views used for research and monitoring project status.

View That Uses Attributes to Select Versions

Suppose that the QA team also works on the **major** branch. Individual developers are responsible for making sure that their modules pass a QA check. The QA team builds and tests the application, using the most recent versions that have passed the check.

The QA team can work in a view that uses this config spec:

```
(1)  element -file src/* /main/major/{QAOK=="Yes"}
(2)  element * /main/LATEST
```

To make this scheme work, you must create an attribute type, **QAOK**. Whenever a new version that passes the QA check is checked in on the **major** branch, an instance of **QAOK** with the value `Yes` is attached to that version. (This can be done manually or with a ClearCase trigger.)

If an element in the **/src** directory has been edited on the **major** branch, this view selects the branch's most recent version that has been marked as passing the QA check (Rule 1). If no version has been so marked or if no **major** branch has been created, the most recent version on the **main** branch is used (Rule 2).

NOTE: Rule 1 on this config spec does not provide a match if an element has a **major** branch, but no version on the branch has a **QAOK** attribute. This command can locate the branches that do not have this attribute:

```
% cleartool find . –branch '{brtype(major) && \! attype_sub(QAOK)}' –print
```

The backslash (\ ) is required in the C shell only, to keep the exclamation point (**!**) from indicating a history substitution.

The **attype_sub** primitive searches for attributes on an element's versions and branches, as well as on the element itself.

This scheme allows the QA team to monitor the progress of the rest of the group. The development config spec always selects the most recent version on the **major** branch, but the QA config spec may select an intermediate version (Figure 34).

Figure 34 Development Config Spec vs. QA Config Spec

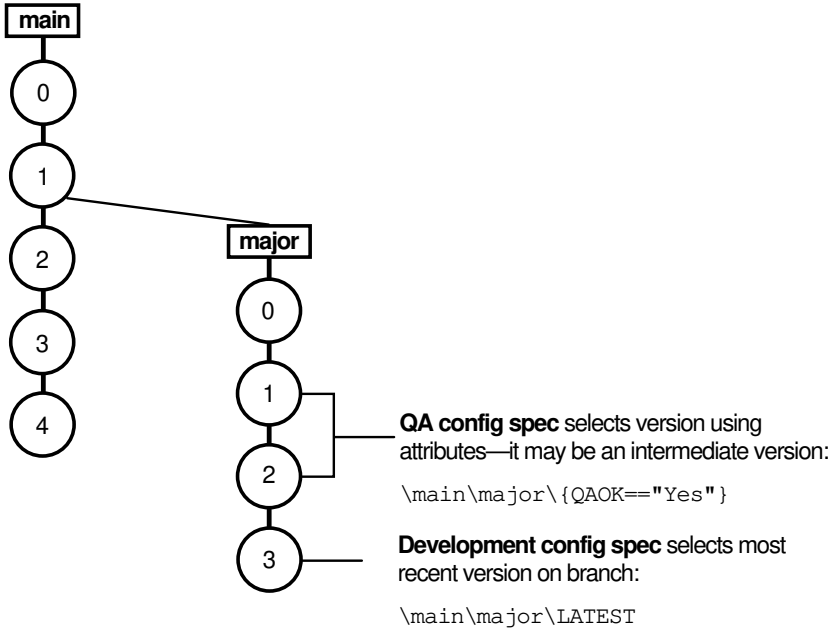

Pitfalls of Using This Configuration for Development

You may be tempted to add a **CHECKEDOUT** rule to the above config spec, turning the QA configuration into a development configuration:

```
(0)  element * CHECKEDOUT
(1)  element -file src/* /main/major/{QAOK=="Yes"}
(2)  element * /main/LATEST
```

It may seem desirable to use attributes, or other kinds of metadata, in addition to (or instead of) branches to control version selection in a development view. But such schemes introduce complications. Suppose that the config spec above selects version **/main/major/2** of element **.../src/cmd.c.** (See Figure 35.)

Figure 35 Checking Out a Branch of an Element

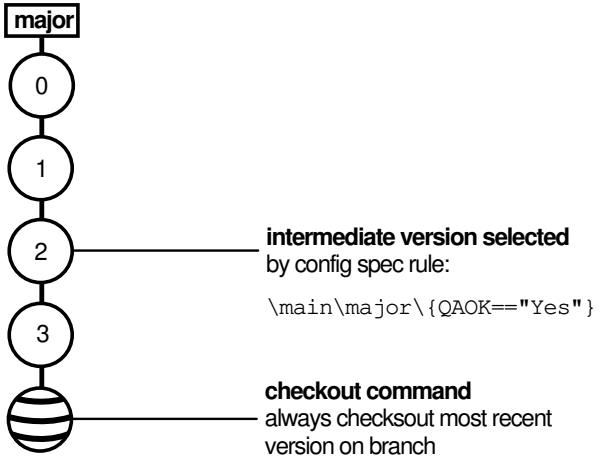

Performing a checkout in this view checks out version **/main/major/3**, not version **/main/major/2**:

```
cleartool: Warning: Version checked out is different from version previously
selected by view.
Checked out "cmd.c" from version "/main/major/3".
```

This behavior reflects the ClearCase restriction that new versions can be created only at the end of a branch. Although such operations are possible, they are potentially confusing to other team members. And in this situation, it is almost certainly not what the developer who checks out the element wants to happen.

You can avoid this problem by modifying the config spec and creating another branching level at the version that the attribute selects. This is the new config spec:

```
(0)     element * CHECKEDOUT
(0a)    element * /main/major/temp/LATEST
(1)     element -file src/* /main/major/{QAOK=="Yes"} -mkbranch temp
(2)     element * /main/LATEST
```

View That Shows Changes of One Developer

This config spec makes it easy to examine all changes a developer has made since a certain milestone:

```
(1)     element * '/main/{created_by(jackson) && created_since(25-Apr)}'
(2)     element * /main/LATEST -time 25-Apr
```

NOTE: Rule 1 must be contained on a single physical text line.

A particular date, April 25, is used as the milestone. The configuration is a snapshot of the main line of development at that date (Rule 2), overlaid with all changes that user **jackson** has made on the **main** branch since then (Rule 1).

The output of the **cleartool ls** command distinguishes **jackson**'s files from the others: each entry includes an annotation as to which configuration rule applies to the selected version.

This is a research view, not a development view. The selected set of files may not be consistent: some of **jackson**'s changes may rely on changes made by others, and those other changes are excluded from this view. Thus, this config spec omits the standard **CHECKEDOUT** and **/main/LATEST** rules.

Historical View Defined by a Version Label

This config spec defines a historical configuration:

```
(1)     element * R1.0 -nocheckout
```

This view always selects the set of versions labeled **R1.0**. In this scenario, all these versions are on the **main** branch of their elements. If the **R1.0** label type is *one-per-element*, not *one-per-branch*,

this config spec selects the **R1.0** version on a subbranch. (For more information, see the **mklbtype** reference page.)

The **–nocheckout** qualifier prevents any element from being checked out in this view. (It also prevents creation of new elements, because the parent directory element must be checked out.) Thus, there is no need for the **CHECKEDOUT** configuration rule.

NOTE: The set of versions selected by this view can change, because version labels can be moved and deleted. For example, using the command **mklabel –replace** to move **R1.0** from version 5 of an element to version 7 changes which version appears in the view. Similarly, using **rmlabel** suppresses the specified elements from the view. (The cleartool **ls** command lists them with a `[no version selected]` annotation.) If the label type is locked with the **lock** command, the configuration cannot change.

You can use this configuration to rebuild Release 1.0, verifying that all source elements have been labeled properly. You can also use it to browse the old release.

Historical View Defined by a Time Rule

This config spec defines a frozen configuration in a slightly different way than the previous one:

(1)
```
element * /main/LATEST –time 4–Sep.02:00 –nocheckout
```

This configuration selects the version that was the most recent on the **main** branch on September 4 at 2 A.M. Subsequent checkouts and checkins cannot change which versions satisfy this criterion; only deletion commands such as **rmver** or **rmelem** can change the configuration. The **–nocheckout** qualifier prevents elements from being checked out or created.

This configuration can be used to view a set of versions that existed at a particular point in time. If modifications must be made to this source base, you must modify the config spec to "unfreeze" the configuration.

9.7 Views for Project Builds

The config specs in this section are useful for running the various types of builds required for a project.

View That Uses Results of a Nightly Build

Many projects use scripts to run unattended software builds every night. The success or failure of these builds determine the impact of any checked-in changes on the application. In layered build environments, they can also provide up-to-date versions of lower-level software (libraries, utility programs, and so on).

Suppose that every night, a script does the following:

- Builds libraries in various subdirectories of **/vobs/monet/lib**
- Checks them in as DO versions in the library staging area, **/vobs/monet/lib**
- Labels the versions **LAST_NIGHT**

You can use this config spec if you want to use the libraries produced by the nightly builds:

```
(1)  element * CHECKEDOUT
(2)  element lib/*.a LAST_NIGHT
(3)  element lib/*.a R2_BL2
(4)  element * /main/LATEST
```

The **LAST_NIGHT** version of a library is selected whenever such a version exists (Rule 2). If a nightly build fails, the previous night's build still has the **LAST_NIGHT** label and is selected. If no **LAST_NIGHT** version exists (the library is not currently under development), the stable version labeled **R2_BL2** is used instead (Rule 3).

For each library, selecting versions with the **LAST_NIGHT** label rather than the most recent version in the staging area allows developers to stage new versions the next day, without affecting developers who use this config spec.

Variations That Select Versions of Project Libraries

The scheme described above uses version labels to select particular versions of libraries. For more flexibility, the **LAST_NIGHT** version of some libraries may be selected, the **R2_BL2** version of others, and the most recent version of still others:

```
(1)     element * CHECKEDOUT
(2a)    element lib/libcmd.a LAST_NIGHT
(2b)    element lib/libparse.a LAST_NIGHT
(3a)    element lib/libcalc.a R2_BL2
(3b)    element lib/*.a /main/LATEST
(4)     element * /main/LATEST
```

(Rule 3b is not required here, because Rule 4 handles all other libraries. It is included for clarity only.)

Other kinds of metadata can also be used to select library versions. For example, **lib_selector** attributes can take values such as `experimental`, `stable`, and `released`. A config spec can mix and match library versions like this:

```
(1)     element * CHECKEDOUT
(2)     element lib/libcmd.a {lib_selector=="experimental"}
(3)     element lib/libcalc.a {lib_selector=="experimental"}
(4)     element lib/libparse.a {lib_selector=="stable"}
(5)     element lib/*.a {lib_selector=="released"}
(6)     element * /main/LATEST
```

View That Selects Versions of Application Subsystems

This config spec selects specific versions of the application's subsystems:

```
(1)     element * CHECKEDOUT
(2)     element /vobs/monet/lib/... R2_BL1
(3)     element /vobs/monet/include/... R2_BL2
(4)     element /vobs/monet/src/... /main/LATEST
(5)     element * /main/LATEST
```

In this situation, a developer is making changes to the application's source files on the **main** branch (Rule 4). Builds of the application use the libraries in directory **/lib** that were used to build Baseline 1, and the header files in directory **/include** that were used to build Baseline 2.

View That Selects Versions That Built a Particular Program

This config spec defines a view that selects only enough files required to rebuild a particular program or examine its sources:

(1)
```
element * -config /vobs/monet/src/monet
```

All elements that were not involved in the build of **monet** appear in the output of ClearCase **ls** with a `[no version selected]` annotation.

This config spec selects the versions listed in the config record (CR) of a particular derived object (and in the config records of all its build dependencies). It can be a derived object that was built in the current view, or another view, or it can be a *DO version*.

In this config spec, **monet** is a derived object in the current view. You can reference a derived object in another view with an extended pathname that includes a *DO-ID*:

(1)
```
element * -config /vobs/monet/src/monet@@09-Feb.13:56.812
```

But typically, this kind of config spec is used to configure a view from a derived object that has been checked in as a DO version.

Configuring the Makefile

By default, a derived object's config record does not list the version of the makefile that was used to build it. Instead, the CR includes a copy of the build script itself. (Why? When a new version of the makefile is created with a revision to one target's build script, the configuration records of all other derived objects built with that makefile are not rendered out of date.)

But if the **monet** program is to be rebuilt in this view using **clearmake** (or even standard **make**), a version of the makefile must be selected somehow. You can have **clearmake** record the makefile version in the config record by including the special **clearmake** macro invocation **$(MAKEFILE)** in the target's dependency list:

```
monet: $(MAKEFILE) monet.o ...
     cc -o monet ...
```

clearmake always records the versions of explicit dependencies in the CR.

Alternatively, you can configure the makefile at the source level: attach a version label to the makefile at build time, and then use a config spec like the one in *Historical View Defined by a*

Version Label on page 119 or *View to Modify an Old Configuration* on page 112 to configure a view for building. You can also use the special target **.DEPENDENCY_IGNORED_FOR_REUSE**; for more information, see *Including a Makefile Version in a Configuration Record* in *Building Software with ClearCase*.

Fixing Bugs in the Program

If a bug is discovered in the **monet** program, as rebuilt in a view that selects only enough files required to rebuild a particular program, it is easy to convert the view from a build configuration to a development configuration. As usual, when making changes in old sources, follow this strategy:

- Create a branch at each version to be modified
- Use the same branch name (that is, create an instance of the same branch type) in every element

If the fix-up branch type is **r1_fix,** this modified config spec reconfigures the view for performing the fix:

```
(1)  element * CHECKEDOUT
(2)  element * .../r1_fix/LATEST
(3)  element * -config /vobs/monet/src/monet -mkbranch r1_fix
(4)  element * /main/LATEST -mkbranch r1_fix
```

Selecting Versions That Built a Set of Programs

It is easy to expand the config spec that selects only enough files required to rebuild a particular program to configure a view with the sources used to build a set of programs, rather than a single program:

```
(1)  element * -config /proj/monet/src/monet
(2)  element * -config /proj/monet/src/xmonet
(3)  element * -config /proj/monet/src/monet_conf
```

There can be version conflicts in such configurations, however. For example, different versions of file **params.h** may have been used in the builds of **monet** and **xmonet**. In this situation, the version used in **monet** is configured, because its configuration rule came first. Similarly, there can be conflicts when using a single **–config** rule: if the specified derived object was created by actually building some targets and using DO versions of other targets, multiple versions of some source files may be involved.

You can modify this config spec as described in *Fixing Bugs in the Program on page 124*, to change the build configuration to a development configuration.

9.8 Sharing Config Specs Between UNIX and Windows

You can, in principle, share config specs between UNIX and Windows systems. That is, users on both systems, using views whose storage directories reside on either platform, can set and edit the same set of config specs.

We recommend that you avoid sharing config specs across platforms. If possible, maintain separate config specs for each platform. However, if you must share config specs, adhere to the following requirements:

- Use slashes (/), not backslashes (\) in pathnames.
- Use relative, not full, pathnames whenever possible, and do not use VOB-tags in pathnames. You can ignore this restriction if your UNIX and Windows VOB-tags both use single, identical pathname components that differ only in their leading slash characters—**\src** and **/src**, for example.
- Always edit and set config specs on UNIX.

The following sections describe these requirements in detail.

Pathname Separators

When writing config specs to be shared by Windows and UNIX computers, you must use slash (/), not backslash (\), as the pathname separator. ClearCase on UNIX recognizes slashes only. (Note that **cleartool** recognizes both slashes and backslashes in pathnames; **clearmake** is less flexible. See *clearmake Makefiles and BOS Files* in *Building Software with ClearCase* for more information.)

Pathnames in Config Spec Element Rules

Windows and UNIX network regions often use different VOB-tags to register the same VOBs. Only single-component VOB-tag names, such as **\proj1**, are permitted on Windows computers; multiple-component VOB-tags, such as **/vobs/src**, are common on UNIX.

When VOB-tags differ between regions, any config spec element rules that use full pathnames (which include VOB-tags) can be resolved when the config spec is compiled (**cleartool edcs** and **setcs** commands) but only by computers in the applicable network region. This implies a general restriction regarding shared config specs: a given config spec must be compiled only on the operating system for which full pathnames in element rules make sense. That is, a config spec with full pathnames is shareable across network regions, even when VOB-tags disagree, but it must be compiled in the right place.

The restrictions do not apply if either of the following is true (see *Example* on page 126):

- The config spec's element rules use only relative pathnames, which do not include VOB-tags.
- Shared VOBs are registered with identical, single-component VOB-tags in both Windows and UNIX network regions. (The VOB-tags **\r3vob** and **/r3vob** are treated as if they were identical because they differ only in the leading slashes.)

Config Spec Compilation

A config spec that is in use exists in both text file and compiled formats. A config spec's *compiled* form is portable. The restriction is that full VOB pathnames in element rules must be resolvable at compile time. A config spec is compiled when you edit or set it (with the **cleartool edcs** or **cleartool setcs** command or a ClearCase GUI). If a user on the other operating system recompiles a config spec (by issuing the **edcs** or **setcs** command or causing the GUI to execute one of these commands) the config spec becomes unusable by *any* computer using that view. If this happens, recompile the config spec on the original operating system.

Example

This config spec element rule may cause problems:

```
element \vob_p2\abc_proj_src\*      \main\rel2\LATEST
```

If the VOB is registered with VOB-tag **\vob_p2** on a Windows network region, but with VOB-tag **/vobs/vob_p2** on a UNIX network region, only Windows computers can compile the config spec.

To address the problem, do one of the following:

- Use relative pathnames that do not include VOB-tags, for example:

```
element ...\abc_proj_src\*      \main\rel2\LATEST
```

- On UNIX, change the VOB-tag so that it has a single component, **/vob_p2**.

10 Implementing Project Development Policies

This chapter presents brief scenarios that show how you can implement and enforce common development policies with ClearCase. The scenarios use various combinations of these functions and metadata:

- Attributes
- Labels
- Branches
- Triggers
- Config specs
- Locks
- Hyperlinks

Sharing Triggers Between UNIX and Windows on page 141 describes how to define triggers for use on UNIX and Windows computers.

10.1 Policy: Good Documentation of Changes Is Required

Each ClearCase command that modifies a VOB creates one or more *event records*. Many such commands (for example, **checkin**) prompt for a comment. The event record includes the user name, date, comment, host, and description of what was changed.

To prevent developers from subverting the system by providing empty comments, you can create a preoperation trigger to monitor the **checkin** command. The trigger action script analyzes

the user's comment (passed in an environment variable), disallowing unacceptable ones (for example, those shorter than 10 words).

Trigger Definition:

```
% cleartool mktrtype –element –all –preop checkin –c "must enter descriptive comment" \
–exec /public/scripts/comment_policy.sh CommentPolicy
```

Trigger Action Script:

```
#!/bin/sh
#
#    comment_policy
#
ACCEPT=0
REJECT=1
WORDCOUNT=`echo $CLEARCASE_COMMENT | wc –w`

if [ $WORDCOUNT –ge 10 ] ; then
     exit $ACCEPT
else
     exit $REJECT
fi
```

10.2 Policy: All Source Files Require a Progress Indicator

You may want to monitor the progress of individual files or determine which or how many files are in a particular state. You can use attributes to preserve this information and triggers to collect it.

In this case, you can create a string-valued attribute type, **Status,** which accepts a specified set of values.

Attribute Definition:

```
% cleartool mkattype –c "standard file levels" \
–enum ' "inactive","under_devt","QA_approved" ' Status
Created attribute type "Status".
```

Developers apply the **Status** attribute to many different versions of an element. Its value in early versions on a branch is likely to be `inactive` and `under_devt`; on later versions, its value is

`QA_approved`. The same value can be used for several versions, or moved from an earlier version to a later version.

To enforce conscientious application of this attribute to versions of all source files, you can create a **CheckStatus** trigger whose action script prevents developers from checking in versions that do not have a **Status** attribute.

Trigger Definition:

```
% cleartool mktrtype –element –all –preop checkin \
–c "all versions must have Status attribute" \
–exec 'Perl /public/scripts/check_status.pl' CheckStatus
```

Trigger Action Script:

```
$pname = $ENV{'CLEARCASE_PN'};
$val = "";
$val = `cleartool describe -short -aattr Status $pname`;

if ($val eq "") {
exit (1);
} else {
     exit (0);
}
```

10.3 Policy: Label All Versions Used in Key Configurations

To identify which versions of which elements contributed to a particular baseline or release, you can attach labels to these versions. For example, after Release 2 is built and tested, you can create label type **REL2,** using the **mklbtype** command. You can then attach **REL2** as a version label to the appropriate source versions, using the **mklabel** command.

Which are the appropriate versions? If Release 2 was built from the bottom up in a particular view, you can label the versions selected by that view:

```
% cleartool mklbtype –c "Release 2.0 sources" REL2
```

```
% cleartool mklabel –recurse REL2 top-level-directory
```

Alternatively, you can use the configuration records of the release's derived objects to control the labeling process:

% **clearmake vega**

... sometime later, after QA approves the build:
% **cleartool mklabel –config vega@@17-Jun.18:05 REL2**

Using configuration records to attach version labels ensures accurate and complete labeling, even if developers have created new versions since the release build. Development work can continue while quality-assurance and release procedures are performed.

To prevent version label **REL2** from being used again, you must lock the label type:

% **cleartool lock –nusers vobadm lbtype:REL2**

The object is locked to all users except those specified with the **–nusers** option, in this case, **vobadm**.

10.4 Policy: Isolate Work on Release Bugs to a Branch

You may want to fix bugs found in the released system on a named bugfix branch, and to begin this work with the exact configuration of versions from that release.

This policy reflects the ClearCase baseline-plus-changes model. First, a label (**REL2,** for example) must be attached to the release configuration. Then, you or any team member can create a view with the following config spec to implement the policy:

```
element * CHECKEDOUT
element * .../rel2_bugfix/LATEST
element * REL2 –mbranch rel2_bugfix
```

If all fixes are made in one or more views with this configuration, the changes are isolated on branches of type **rel2_bugfix**. The **–mkbranch** option causes such branches to be created, as needed, when elements are checked out.

This config spec selects versions from **rel2_bugfix** branches, where branches of this type exist; it creates such a branch whenever a **REL2** version is checked out.

10.5 Policy: Avoid Disrupting the Work of Other Developers

To work productively, developers need to control when they see changes and which changes they see. The appropriate mechanism for this purpose is a view. Developers can modify an existing config spec or create a new one to specify exactly which changes to see and which to exclude.

To implement this policy, you can also require developers to write and distribute the config spec rule that filters out their checked-in changes. Some sample config specs:

- To select your own work, plus all the versions that went into the building of Release 2:

```
element * CHECKEDOUT
element * REL2
```

- To select your own work, plus the latest versions as of Sunday evening:

```
element * CHECKEDOUT
element * /main/LATEST -time Sunday.18:00
```

- To select your own work, new versions created in the **graphics** directory, and the versions that went into last night's build:

```
element * CHECKEDOUT
element graphics/* /main/LATEST
element * -config myprog@@12-Jul.00:30
```

- To select your own work, the versions either you (**jones**) or Mary has checked in today, and the most recent quality-assurance versions:

```
element * CHECKEDOUT
element * '/main/{ created_since(06:00) && ( created_by(jones) ||
created_by(mary) ) }'
element * /main/{QAed=="TRUE"}
```

- You can also use the config spec include facility to set up standard sets of configuration rules for developers to add to their own config specs:

```
element * CHECKEDOUT
element msg.c /main/18
include /usr/cspecs/rules_for_rel2_maintenance
```

10.6 Policy: Deny Access to Project Data When Necessary

Occasionally, you may need to deny access to all or most project team members. For example, you may want to prevent changes to public header files until further notice. The **lock** command is designed to enforce such temporary policies:

- Lock all header files in a certain directory:

  % **cleartool lock src/pub/*.h**

- Lock the header files for all users except Mary and Fred:

  % **cleartool lock –nusers mary,fred src/pub/*.h**

- Lock all header files in the VOB:

  % **cleartool lock eltype:c_header**

- Lock an entire VOB:

  % **cleartool lock vob:/vobs/myproj**

10.7 Policy: Notify Team Members of Relevant Changes

To help team members keep track of changes that affect their own work, you can use postoperation triggers to send notifications of various events. For example, when developers change the GUI, an e-mail message to the doc group ensures that these changes are documented.

To enforce this policy, create a trigger type that sends mail, and then attach it to the relevant elements.

Trigger Definition:

```
% cleartool mktrtype –nc -element –postop checkin \
      –exec /public/scripts/informwriters.sh InformWriters
Created trigger type "InformWriters".
```

Trigger Action Script:

```
#!/bin/sh
#
#Init
tmp=/tmp/checkin_mail

# construct mail message describing checkin

cat > $tmp <<EOF
Subject: Checkin $CLEARCASE_PNAME by $CLEARCASE_USER
$CLEARCASE_XPNAME
Checked in by $CLEARCASE_USER.

Comments:
$CLEARCASE_COMMENT
EOF

# send the message

mail docgrp <$tmp

# clean up

#rm -f $tmp
```

To attach triggers to existing elements:

1. Place the trigger on the *inheritance list* of all existing directory elements within the GUI source tree:

   % **cleartool find /vobs/gui_src –type d **
   –exec 'cleartool mktrigger –nattach InformWriters $CLEARCASE_PN'

2. Place the trigger on the *attached list* of all existing file elements within the GUI source tree:

   % **cleartool find /vobs/gui_src –type f **
   –exec 'cleartool mktrigger InformWriters $CLEARCASE_PN'

10.8 Policy: All Source Files Must Meet Project Standards

To ensure that developers are following coding guidelines or other standards, you can evaluate their source files. You can create preoperation triggers to run user-defined programs, and cancel the commands that trigger them.

For example, you may want to disallow checkin of C-language files that do not satisfy quality metrics. You may have your own metrics program, or you can run **lint(1)**. Suppose that you have defined an element type, **c_source**, for C language files (*.c).

Trigger Definition:

```
% cleartool mktrtype –element –all –eltype c_source \
–preop checkin –exec '/public/scripts/apply_metrics.sh $CLEARCASE_PN' ApplyMetrics
```

This trigger type **ApplyMetrics** applies to all elements; it fires when any element of type **c_source** is checked in. (When a new **c_source** element is created, it is monitored.) If a developer attempts to check in a **c_source** file that fails the **apply_metrics.sh** test, the checkin fails.

NOTE: The **apply_metrics.sh** script can read the value of **CLEARCASE_PN** from its environment. Having it accept a file-name argument provides flexibility because the script can be invoked as a trigger action, and developers can also use it manually.

10.9 Policy: Associate Changes with Change Orders

To keep track of work done in response to an engineering change order (ECO), you can use attributes and triggers. For example, to associate a version with an ECO, define **ECO** as an integer-valued attribute type:

```
cleartool mkattype –c "bug number associated with change" –vtype integer ECO
Created attribute type "ECO".
```

Then, define an all-element trigger type, **EcoTrigger**, which fires whenever a new version is created and runs a script to attach the **ECO** attribute:

Trigger Definition:

```
cleartool mktrtype –element –all –postop checkin –c "associate change with bug number" \
–execunix 'Perl /public/scripts/eco.pl' –execwin 'ccperl \\neon\scripts\eco.pl' EcoTrigger
Created trigger type "EcoTrigger".
```

Trigger Action Script:

```
$pname = $ENV{'CLEARCASE_XPN'};

print "Enter the bug number associated with this checkin: ";
$bugnum = <STDIN>;
chomp ($bugnum);
$command = "cleartool mkattr ECO $bugnum $pname";

@returnvalue = `$command`;
$rval = join "",@returnvalue;
print "$rval";

exit(0);
```

When a new version is created, the attribute is attached to the version. For example:

cleartool checkin –c "fixes for 4.0" src.c

```
Enter the bug number associated with this checkin: 2347
Created attribute "ECO" on "/vobs/dev/src.c@@/main/2".
Checked in "src.c" version "/main/2".
```

cleartool describe src.c@@/main/2

```
version "src.c@@/main/2"
...
  Attributes:
    ECO = 2347
```

10.10 Policy: Associate Project Requirements with Source Files

You can implement requirements tracing with *hyperlinks*, which associate pairs of VOB objects. The association should be at the version level (rather than the branch or element level): each version of a source code module must be associated with a particular version of a related design document.

For example, the project manager creates a hyperlink type named DesignDoc, which is used to associate source code with design documents:

**cleartool mkhltype –c "associate code with design docs" **
DesignDoc@/vobs/dev DesignDoc@/vobs/design

```
Created hyperlink type "DesignDoc".
Created hyperlink type "DesignDoc".
```

The *hyperlink inheritance* feature makes the implementation of requirements tracing easy:

➤ When the source module, **hello.c**, and the design document, **hello_dsn.doc**, are updated, the project manager creates a new hyperlink connecting the two updated versions:

cleartool mkhlink -c "source doc" DesignDoc hello.c /vobs/design/hello_dsn.doc
```
Created hyperlink "DesignDoc@90@/vobs/dev".
```

➤ When either the source module or the design document incorporates a minor update, no hyperlink-level change is required: the new version *inherits* the hyperlink connection of its predecessor.

cleartool checkin -c "fix bug" hello.c
```
Checked in "hello.c" version "/main/2".
```

To list the inherited hyperlink, use the **–ihlink** option to the **describe** command:

version that inherits hyperlink->

version to which -> hyperlink is explicitly attached

```
% cleartool describe –ihlink DesignDoc hello.c@@/main/2
hello.c@@/main/2
  Inherited hyperlinks: DesignDoc@90@/vobs/dev
/vobs/dev/hello.c@@/main/1 ->
/vobs/doc/hello_dsn.doc@@/main/1
```

➤ When either the source module or the design document incorporates a significant update, which renders the connection invalid, the project manager creates a *null-ended* hyperlink to sever the connection:

cleartool mkhlink -c "sever connection to design doc" DesignDoc hello.c
```
Created hyperlink "DesignDoc@94@/vobs/dev".
```

Figure 36 illustrates the hyperlinks that connect the source file to the design doc.

Figure 36 Requirements Tracing

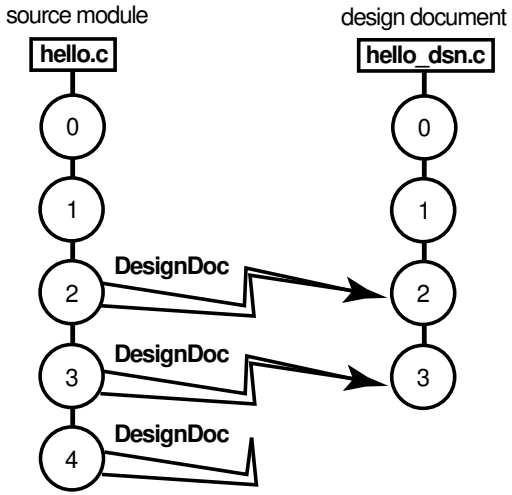

10.11 Policy: Prevent Use of Certain Commands

To control which users can execute certain commands on ClearCase objects, you can create a pair of trigger types. One of the types controls use of the command on element-related objects, and the other controls use of the command on type objects. Both trigger types use the **–nuser** flag to specify the users who are allowed to use the command.

NOTE: You cannot use triggers to prevent a command being used on an object that is not element-related or a type object. For example, you cannot create a trigger type to prevent operations on VOB objects or replica objects.

For a list of commands that can be triggered, see the **events_ccase** and **mktrtype** reference pages.

For example, the following commands create two trigger types that prevent all users except **stephen**, **hugh**, and **emma** from running the **chmaster** command on element-related objects and type objects in the current VOB:

**cleartool mktrtype –element –all –preop chmaster –nusers stephen,hugh,emma **
**–execunix 'Perl –e "exit –1;"' –execwin 'ccperl –e "exit (–1);"' **
–c "ACL for chmaster" elem_chmaster_ACL

**cleartool mktrtype –type –preop chmaster –nusers stephen,hugh,emma **
**–execunix 'Perl –e "exit –1;"' –execwin 'ccperl –e "exit (–1);"' **
**–attype –all –brtype –all –eltype –all –lbtype –all –hltype –all **
–c "ACL for chmaster" type_chmaster_ACL

When user **tony** tries to run the **chmaster** command on a restricted object, the command fails. For example:

**cleartool chmaster –c "give mastership to london" london@/vobs/dev **
/vobs/dev/acc.c@@/main/lex_dev

```
cleartool: Warning: Trigger "elem_chmaster_ACL" has refused to let chmaster
proceed.
cleartool: Error: Unable to perform operation "change master" in replica "lex"
of VOB "/vobs/dev".
```

10.12 Policy: Certain Branches Are Shared Among MultiSite Sites

If your company uses ClearCase MultiSite to support development at different sites, you must tailor your branching strategy to the needs of the different sites. The standard MultiSite development model is that a replica of the VOB is at each site. Each replica controls (masters) a site-specific branch type, and developers at one site cannot work on branches mastered at another site. (See *Introduction to MultiSite* in *ClearCase MultiSite Manual* for more information on MultiSite mastership.)

However, sometimes you cannot, or may not want to, branch and merge an element. For example, some file types cannot be merged, so development must occur on a single branch. In this scenario, all developers must work on a single branch (usually, the main branch). MultiSite allows only one replica to master a branch at any given time. Therefore, if a developer at another site needs to work on the element, mastership of the branch must be transferred to that site.

MultiSite provides two models for transferring mastership of a branch:

- The push model, in which the administrator at the replica that masters the branch uses the **chmaster** command to give mastership to another replica.

This model is not efficient in a branch-sharing situation, because it requires communication with an administrator at a remote site. For more information about this model, see *ClearCase MultiSite Manual*.

- The pull model, in which the developer who needs to work on the branch uses the **reqmaster** command to request mastership of the branch.

  This model requires the MultiSite administrators to enable requests for mastership in each replica, and to authorize individual developers to request mastership. If you decide to implement this model, you must provide the following information to your MultiSite administrator:

  - Replicated VOBs that should be enabled to handle mastership requests
  - Identities (domain names and user names) of developers who should be authorized to request mastership
  - Branch types and branches for which mastership requests should be denied; (for example, branch types that are site specific, or branches that must remain under the control of a single site)

  Implementing Requests for Branch Mastership in *ClearCase MultiSite Manual* describes the process of enabling the pull model and a scenario in which developers use the pull model. *Working On a Team* in *Developing Software with ClearCase* describes the procedure developers use to request mastership.

10.13 Sharing Triggers Between UNIX and Windows

You can define triggers that fire correctly on both UNIX and Windows computers. The following sections describe two techniques. With one, you use different pathnames or different scripts; with the other, you use the same script for both platforms.

Using Different Pathnames or Different Scripts

To define a trigger that fires on UNIX, Windows, or both, and that uses different path names to point to the trigger scripts, use the **–execunix** and **–execwin** options with the **mktrtype** command. These options behave the same as **–exec** when fired on the appropriate platform (UNIX or Windows, respectively). On the other platform, they do nothing. This technique allows

a single trigger type to use different paths for the same script or to use completely different scripts on UNIX and Windows computers. For example:

```
cleartool mktrtype –element –all –nc –preop checkin \
–execunix /public/scripts/precheckin.sh –execwin \\neon\scripts\precheckin.bat \
pre_ci_trig
```

On UNIX, only the script **precheckin.sh** runs. On Windows, only **precheckin.bat** runs.

To prevent users on a new platform from bypassing the trigger process, triggers that specify only **–execunix** always fail on Windows. Likewise, triggers that specify only **–execwin** fail on UNIX.

Using the Same Script

To use the same trigger script on both Windows and UNIX platforms, you must use a batch command interpreter that runs on both operating systems. For this purpose, ClearCase includes the **ccperl** program. On Windows, **ccperl** is a version of the Perl program available on UNIX.

The following **mktrtype** command creates sample trigger type **pre_ci_trig** and names **precheckin.pl** as the executable trigger script.

```
% cleartool mktrtype –element –all –nc –preop checkin   \
      –execunix 'Perl /public/scripts/precheckin.pl'   \
      –execwin 'ccperl \\neon\scripts\precheckin.pl'   \
      pre_ci_trig
```

Notes

- To conditionalize script execution based on operating system, use environment variables in Perl scripts.
- To collect or display information interactively, you can use the **clearprompt** command.
- For more information on using the **–execunix** and **–execwin** options, see the **mktrtype** reference page.

11 Integrating Changes

In a parallel development environment, the opposite of branching is merging. In the simplest scenario, merging incorporates changes on a subbranch into the **main** branch. However, you can merge work from any branch into any other branch. This chapter describes techniques and scenarios for merging versions of elements and branches. ClearCase includes automated merge facilities for handling almost any scenario.

11.1 How Merging Works

A merge combines the contents of two or more files or directories into a single new file/directory. The ClearCase merge algorithm uses the following files during a merge (see Figure 37):

- Contributors, which are typically one version from each branch you are merging. (You can merge up to 15 contributors.) You specify which versions are contributors.
- The base contributor, which is typically the closest common ancestor of the contributors. (For selective merges, subtractive merges, and merges in an environment with complex branch structures, the base contributor may not be the closest common ancestor.) ClearCase determines which contributor is the base contributor.
- The target contributor, which is typically the latest version on the branch that will contain the results of the merge. You determine which contributor is the target contributor.
- The merge output file, which contains the results of the merge and is usually checked in as a successor to the target contributor. By default, the merge output file is the checked-out version of the target contributor, but you can choose a different file to contain the merge output.

Figure 37 Versions Involved in a Typical Merge

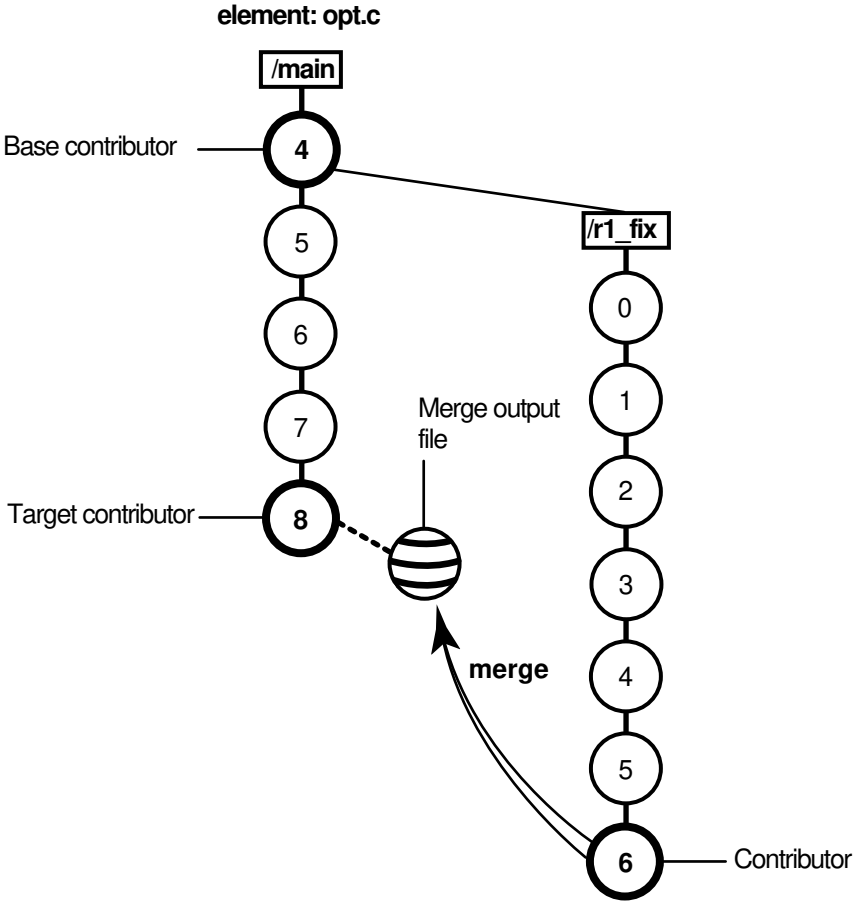

To merge files and directories, ClearCase takes the following steps:

1. It identifies the base contributor.
2. It compares each contributor against the base contributor. (See Figure 38.)
3. For any line that is unchanged between the base contributor and any other contributor, it copies the line to the merge output file.
4. For any line that has changed between the base contributor and one other contributor, it accepts the change in the contributor; depending on how you started the merge operation, ClearCase may copy the change to the merge output file. However, you can disable the

automated merge capability for any given merge operation. If you disable this capability, you must approve each change to the merge output file.

5. For any line that has changed between the base contributor and more than one other contributor, ClearCase requires that you resolve the conflicting difference.

Figure 38 ClearCase Merge Algorithm

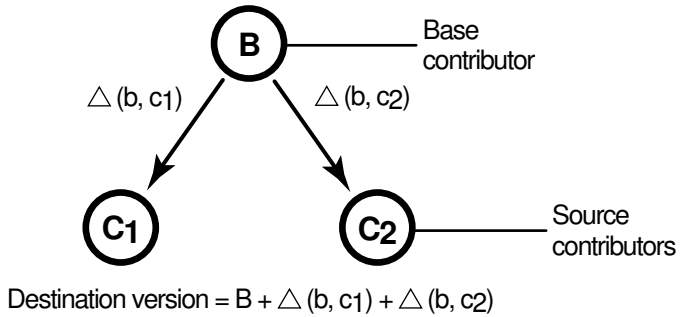

To merge versions, you can use the GUI tools, described briefly in the next section, or the command-line interface, described in *Using the Command Line to Merge Elements* on page 146.

Using the GUI to Merge Elements

ClearCase provides three graphical tools to help you merge elements:

- Merge Manager
- Diff Merge
- Version Tree Browser

The Merge Manager manages the process of merging one or more ClearCase elements. It automates the processes of gathering information for a merge, starting a merge, and tracking a merge. It can also save and retrieve the state of a merge for a set of elements.

You can use the Merge Manager to merge from many directions:

- From a branch to the **main** branch
- From the **main** branch to another branch
- From one branch to another branch

To start the Merge Manager, type **clearmrgman** at a command prompt.

The Diff Merge utility shows the differences between two or more versions of file or directory elements. Use this tool to compare up to 16 versions at a time, navigate through versions, merge versions, and resolve differences between versions.

To start the Diff Merge utility, type **xcleardiff** or use the **cleartool merge –graphical** command at a command prompt.

The Version Tree Browser displays the version tree for an element. The version tree is useful when merging to do the following::

- Locate versions or branches that have contributed to or resulted from a merge
- Start a merge by clicking on the appropriate symbol

The merge can be recorded with a merge arrow, which is implemented as a hyperlink of type **Merge**.

To start the Version Tree Browser, use one of these methods:

- At a command prompt, type **cleartool lsvtree –graphical**
- In the ClearCase File Browser, click an element and click **Versions>Show version tree**

Using the Command Line to Merge Elements

Use the following commands to perform merges from the command line:

- **cleartool merge**
- **cleartool findmerge**
- **cleardiff**

For more information on these commands, see the *ClearCase Reference Manual*.

11.2 Common Merge Scenarios

The following sections present a series of merge scenarios that require work on one branch of an element to be incorporated into another branch. Each scenario shows the version tree of an element that requires a merge and indicates the appropriate command to perform the merge.

Scenario: Selective Merge from a Subbranch

In this scenario, you want to incorporate the changes in version **/main/r1_fix/4** into new development. To perform the merge, you specify which versions on the **r1_fix** branch to include. See Figure 39.

Figure 39 Selective Merge from a Subbranch

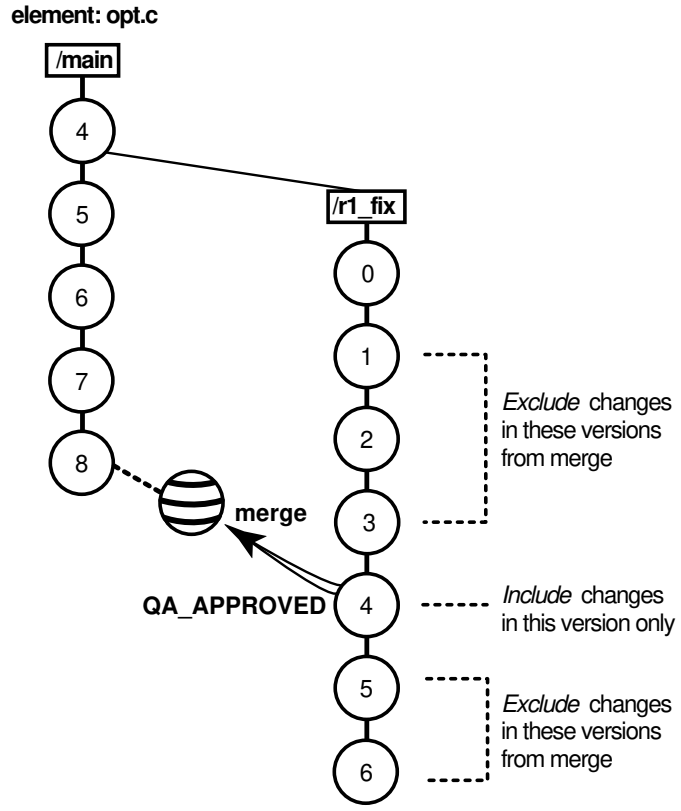

In a view configured with the default config spec, enter these commands to perform the selective merge:

```
% cleartool checkout opt.c
% cleartool merge –to opt.c –insert –version /main/r1_fix/4
```

You can also specify a range of consecutive versions to be merged. For example, this command merges only the changes in versions **/main/r1_fix/2** through **/main/r1_fix/4**:

% **cleartool merge –to opt.c –insert –version /main/r1_fix/2 /main/r1_fix/4**

No merge arrow is created for a selective merge.

Scenario: Removing the Contributions of Some Versions

A new feature, implemented in versions 14 through 16 on the **main** branch, will not be included in the product. You must remove the changes made in those versions. See Figure 40.

Figure 40 Removing the Contributions of Some Versions

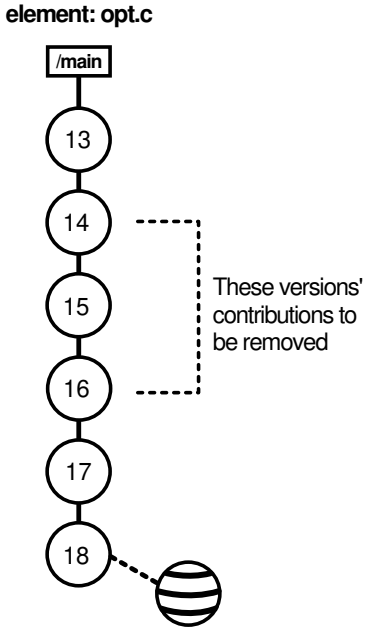

Enter these commands to perform this subtractive merge:

% **cleartool checkout opt.c**
% **cleartool merge –to opt.c –delete –version /main/14 /main/16**

No merge arrow is created for a subtractive merge.

Scenario: Merging All Project Work

Your team has been working on a branch. Now, your job is to merge all the changes into the **main** branch.

The **findmerge** command can handle most common cases easily. It can accommodate the following schemes for isolating the project's work.

All Project Work Is Isolated on a Branch

The standard approach to parallel development isolates all project work on the same branch. More precisely, all new versions of source files are created on like-named branches of their respective elements (that is, on branches that are instances of the same *branch type*). This makes it possible for a single **findmerge** command to locate and incorporate all the changes. Suppose the common branch is named **gopher**. You can enter these commands in a view configured with the default config spec:

```
% cd root-of-source-tree
% cleartool findmerge . –fversion .../gopher/LATEST –merge –graphical
```

The **–merge –graphical** syntax causes the merge to take place automatically whenever possible, and to start the graphical merge utility if an element's merge requires user interaction. If the project has made changes in several VOBs, you can perform all the merges at once by specifying several pathnames, or by using the **–avobs** option to **findmerge**.

All Project Work Isolated In a View

Some projects are organized so that all changes are made in a single view (typically, a shared view). For such projects, use the **–ftag** option to **findmerge**. Suppose the project's work has been done in a view whose view-tag is **goph_vu**. These commands perform the merge:

```
% cd root-of-source-tree
% cleartool findmerge . –ftag goph_vu –merge –graphical
```

NOTE: Working in a single shared view is not recommended because doing so can degrade system performance.

Scenario: Merging a New Release of an Entire Source Tree

Your team has been using an externally supplied source-code product, maintaining the sources in a VOB. The successive versions supplied by the vendor are checked in to the **main** branch and labeled **VEND_R1**, **VEND_R2**, and **VEND_R3.** Your team's fixes and enhancements are created on subbranch **enhance.** The views that your team works in are configured to branch from the **VEND_R3** baseline:

```
element * CHECKEDOUT
element * .../enhance/LATEST
element * VEND_R3 -mkbranch enhance
element * /main/LATEST -mkbranch enhance
```

The version trees in Figure 41 show various likely cases:

- An element that your team started changing at Release 1 (**enhance** branch created at the version labeled **VEND_R1**)
- An element that your team started changing at Release 3
- An element that your team has never changed

Figure 41 Merging a New Release of an Entire Source Tree

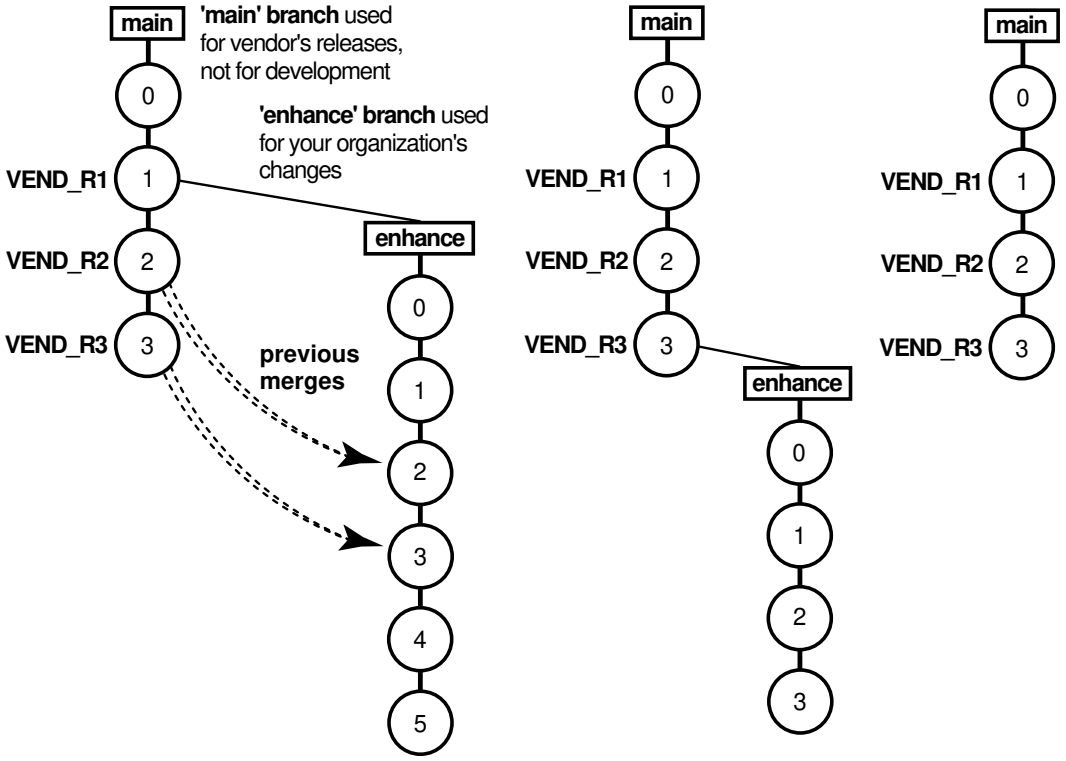

Release 4 has arrived, and you need to integrate this release with your team's changes.

To prepare for the merge, add the new release to the **main** branch and label the versions **VEND_R4**. Merging the source trees involves merging from the version labeled **VEND_R4** to the most recent version on the **enhance** branch; if an element has no **enhance** branch, nothing is merged.

This procedure accomplishes the integration:

1. Load the vendor's Release 4 media into a standard directory tree:

```
%cd /usr/tmp
% tar –xv
```

The directory tree created is **mathlib_4.0**.

2. As the VOB owner, run **clearexport_ffile**, to create a datafile containing descriptions of the new versions.

```
% cd ./mathlib_4.0
% clearexport_ffile
  .  (lots of output)
  .
```

3. In a view configured with the default config spec, start **clearimport** on the file **clearexport_ffile** created. This creates Release 4 versions on the **main** branches of elements (and creates new elements as needed).

```
% cleartool setview mainline

% cd /vobs/proj/mathlib
% clearimport /usr/tmp/mathlib_4.0/cvt_data
```

4. Label the new versions:

```
% cleartool mklbtype –c "Release 4 of MathLib sources" VEND_R4
Created label type "VEND_R4".
% cleartool mklabel –recurse VEND_R4  /vobs/proj/mathlib
  .  (lots of output)
  .
```

5. Set to a view that is configured with your team's config spec and selects the versions on the **enhance** branch:

```
% cleartool setview enh_vu
```

6. Merge from the **VEND_R4** configuration to your view:

```
% cleartool findmerge -nback /vobs/proj/mathlib –fver VEND_R4 –merge –graphical
```

   The **–merge –graphical** syntax instructs **findmerge** to merge automatically if possible, but if not, start the graphical merge tool.

7. Verify the merges, and check in the modified elements.

You have now established Release 4 as the new baseline. Developers on your team can update their view configurations as follows:

```
element * CHECKEDOUT
element * .../enhance/LATEST
```

```
element * VEND_R4 -mkbranch enhance          (change from VEND_R3 to VEND_R4)
element * /main/LATEST -mkbranch enhance
```

Elements that have been active continue to evolve on their **enhance** branches. When elements are revised for the first time, their **enhance** branches are created at the **VEND_R4** version.

Scenario: Merging Directory Versions

One of the most powerful features of ClearCase is versioning of directories. Each version of a directory element catalogs a set of file elements, directory elements, and VOB symbolic links. In a development project, directories change are as often as files do. Merging the changes to another branch is as easy merging files.

Take a closer look at the source tree scenario from the previous section. Suppose you find that the vendor has made several changes in directory **/vobs/proj/mathlib/src**:

- File elements **Makefile**, **getcwd.c,** and **fork3.c** have been revised.
- File elements **readln.c** and **get.c** have been deleted.
- A new file element, **newpaths.c**, has been created.

When you use **findmerge** to merge the changes made in the **VEND_R4** sources to the **enhance** branch, the changes to both the files and the directory are handled automatically. The following **findmerge** excerpt shows the directory merge activity:

```
*******************************
<<< directory 1: /vobs/proj/mathlib/src@@/main/3
>>> directory 2: .@@/main/enhance/1
>>> directory 3: .
*******************************
-------[ removed directory 1 ]-------|----------[ directory 2]------------
get.c  19-Dec-1991 drp                |-
*** Automatic: Applying REMOVE from directory 2
-----------[ directory 1 ]-----------|--------[ added directory 2]---------
                                    -| newpaths.c  08-Mar.21:49 drp
*** Automatic: Applying ADDITION from directory 2
-------[ removed directory 1 ]-------|-----------[ directory 2]------------
readln.c  19-Dec-1991 drp               |-
*** Automatic: Applying REMOVE from directory 2
Recorded merge of ".".
```

If you have changes to merge from both files and directories, it may be a good idea to run **findmerge** twice: first to merge directories, and then to merge files. Using the **–print** option to a **findmerge** command does not report everything that is merged, because **findmerge** does not see

new files or subdirectories in the merge-from version of a directory until after the directories are merged. To report every merge that takes place, use **findmerge** to merge the directories only, and then use **findmerge –print** to get information about the file merges needed. Afterward, you can cancel the directory merges by using the **uncheckout** command on the directories.

11.3 Using Your Own Merge Tools

You can create a merged version of an element manually or with any available analysis and editing tools. Check out the target version, revise it, and check it in. Immediately before (or after) the checkin, record your activity by using the **merge** command with the **–ndata** (no data) option:

```
% cleartool checkout nextwhat.c
Checkout comments for "nextwhat.c":
merge enhance branch
.
Checked out "nextwhat.c" from version "/main/1".

% <invoke your own tools to merge data into checked-out version>

% cleartool merge –to nextwhat.c –ndata –version .../enhance/LATEST
Recorded merge of "nextwhat.c".
```

This form of the **merge** command does not change any file system data; it merely attaches a merge arrow (a hyperlink of type **Merge**) to the specified versions. After you've made this annotation, your merge is indistinguishable from one performed with ClearCase tools.

12 Using Element Types to Customize Processing of File Elements

Most projects involve many different file types. For example, in a typical software release, developers may work on C-language source files, C-language header files, document files in binary format, and library files.

Every file that is stored in a ClearCase VOB is associated with an element type. ClearCase provides predefined element types for various kinds of file types, and every element type has an associated type manager, which handles the operations performed on versions of the element.

For some file types in your project, you may want to create your own element types so that you can customize the handling of the files. You can also create your own type managers.

This chapter describes how ClearCase uses element types and type managers to classify and manage files. It also describes how you can customize file classification and management.

12.1 File Types in a Typical Project

Table 5 lists the files used in a typical development project.

Table 5 Files Used in a Typical Project

Type of File	Identifying Characteristic
Source Files	
C-language source file	**.c** file-name extension
C-language header file	**.h** file-name extension
FrameMaker® binary file	**.doc** or **.mif** file-name extension, first line of file begins with `<Maker`
manual page source file	**.1 – .9** file-name extension
Derived Files	
ar(1) archive (library)	**.a** file-name extension
compiled executable	<*varies with system architecture*>file-name extension

12.2 How ClearCase Assigns Element Types

In various contexts, ClearCase determines one or more *file types* for an existing file-system object, or for a name to be used for a new object. When you create a new element and do not specify an element type, ClearCase determines the file type for the element.

The file-typing routines use predefined and user-defined *magic files*, as described in the **cc.magic** reference page. A magic file can use many different techniques to determine a file type, including file-name pattern-matching, **stat(2)** data, and standard UNIX magic numbers.

For example, the following magic file specifies several file types for each kind of file listed in Table 5.

Sample Magic File

```
(1)  c_src src_file text_file file:      -name "*.c" ;
(2)  hdr_file text_file file:            -name "*.h" ;
(3)  frm_doc binary_delta_file doc file: -magic 0, "<MakerFile" ;
(4)  manpage src_file text_file file:    -name "*.[1-9]" ;
(5)  archive derived_file file:          -magic 32, "archive" ;
(6)  sunexec derived_file file:          -magic 40,"SunBin" ;
```

12.3 Element Types and Type Managers

ClearCase can handle different classes of files differently because it uses *element types* to categorize elements. Each file element in a VOB must have an element type. An element gets its type when it is created; you can change an element's type subsequently, with the **chtype** command. (An element is an *instance* of its element type, in the same way that an attribute is an instance of an attribute type and a version label is an instance of a label type.)

Each element type has an associated *type manager*, a suite of programs that handle the storage and retrieval of versions from storage pools. (See the **type_manager** reference page for information on how type managers work.) Thus, the way in which a file element's data is handled depends on its element type.

NOTE: Each directory element also has an element type. But directory elements do not use type managers; the contents of a directory version are stored in the VOB database itself, not in storage pools.

Figure 42 shows how an element type is assigned to a newly created element.

Figure 42 Data Handling: File Type, Element Type, Type Manager

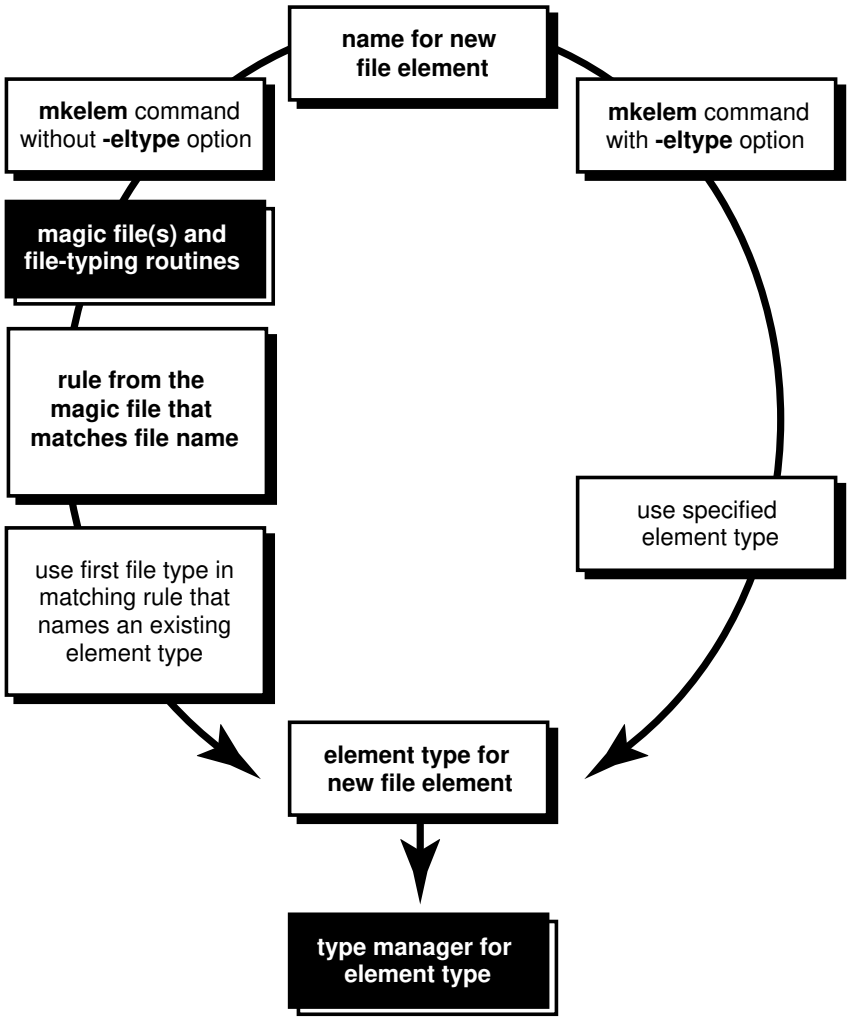

For example, a new element named **monet_adm.1** is assigned an element type as follows:

1. A developer creates an element:

   % **cleartool mkelem monet_adm.1**

2. Because the developer did not specify an element type (**–eltype** option), **mkelem** uses one or more magic files to determine the file types of the specified name.

   NOTE: ClearCase supports a search path facility, using the environment variable **MAGIC_PATH**. See the **cc.magic** reference page for details.

   Suppose that the magic file shown in *Sample Magic File* on page 157 is the first (or only) one to be used. In this case, rule (4) is the first to match the name **monet_adm.1**, yielding this list of file types:

   ```
   manpage src_file text_file file
   ```

3. This list is compared with the set of element types defined for the new element's VOB. Suppose that **text_file** is the first file type that names an existing element type; in this case, **monet_adm.1** is created as an element of type **text_file**.

4. Data storage and retrieval for versions of element **monet_adm.1** are handled by the type manager associated with the **text_file** element type; its name is **text_file_delta**:

   ```
   % cleartool describe eltype:text_file
   element type "text_file"
   ...
     type manager: text_file_delta
     supertype: file
     meta-type of element: file element
   ```

File-typing mechanisms are defined on a per-user or per-site basis; element types are defined on a per-VOB basis. (To ensure that element types are consistent across VOBs, the ClearCase administrator can use global types.) In this case, a new element, **monet_adm.1**, is created as a **text_file** element; in a VOB with a different set of element types, the same magic file may have created it as a **src_file** element.

Other Applications of Element Types

Element types allow differential and customized handling of files beyond the selection of type managers. Following are some examples.

Using Element Types to Configure a View

Creating all C-language header files as elements of type **hdr_file** allows flexibility in configuring views. Suppose that one developer has reorganized the project header files, working on a branch

named **header_reorg** to avoid disrupting the team's work. To compile with the new header files, another developer can use a view reconfigured with one additional rule:

```
element * CHECKEDOUT
element -eltype hdr_file * /main/header_reorg/LATEST
element * /main/LATEST
```

Processing Files by Element Type

Suppose that a coding-standards program named **check_var_names** is to be executed on each C-language source file. If all such files have element type **c_src**, a single **cleartool** command runs the program:

```
% cleartool find –avobs –visible –element 'eltype(c_src)' \
      –exec 'check_var_names $CLEARCASE_PN'
```

12.4 Predefined and User-Defined Element Types

Some of the element types described in this chapter (for example, **text_file**) are predefined. Others (for example, **c_src** and **hdr_file**) are not; the previous examples work only if user-defined element types with these names are created with the **mkeltype** command.

When a new VOB is created, it contains a full set of the predefined element types. Each is associated with one of the type managers provided with ClearCase. The **mkeltype** reference page describes the predefined element types and their type managers.

When you create a new element type with **mkeltype**, you must specify an existing element type as its *supertype*. By default, the new element type uses the same type manager as its supertype; in this case, the only distinction between the new and old types is for the purposes described in *Other Applications of Element Types* on page 159. For differential data handling, use the **–manager** option to create an element type that uses a different type manager from its supertype.

Directory *ccase-home-dir***/examples/clearcase/mkeltype** contains shell scripts that create a hierarchy of element types.

12.5 Predefined and User-Defined Type Managers

ClearCase provides predefined type managers. The type managers are described in the **type_manager** reference page. Each type manager is implemented as a suite of programs in a subdirectory of *ccase-home-dir***/lib/mgrs**; the name of the subdirectory is the name of the type manager.

The **mkeltype –manager** command creates an element type that uses an existing type manager. You can further customize ClearCase by creating new type managers and creating new element types that use them. Architecturally, type managers are mutually independent, but new type managers can use symbolic links to inherit some of the functions of existing ones.

The **type_manager** reference page describes the basic components of a type manager, and outlines the process of creating a new one. The file *ccase-home-dir***/lib/mgrs/mgr_info.h** provides comprehensive information on type managers. Before proceeding to the following sections, which present an extended example of creating and using a new type manager, read the **type_manager** reference page and the **mgr_info.h** file.

12.6 Type Manager for Manual Page Source Files

One kind of file listed in Table 5 is a manual page source file, a file coded in **nroff(1)** format. A type manager for this kind of file may have these characteristics:

- It stores all versions in compressed form in separate data containers, like the **z_whole_copy** type manager.
- It implements version-comparison (compare method) by running **diff** on formatted manual pages instead of the source versions.

The basic strategy is to use most of the **z_whole_copy** type manager's methods. The compare method uses **nroff(1)** to format the versions before displaying their differences.

Creating the Type Manager Directory

The name **mp_mgr** (manual page manager) is appropriate for this type manager. The first step is to create a subdirectory with this name in the *ccase-home-dir***/lib/mgrs** directory. For example:

```
# mkdir /usr/atria/lib/mgrs/mp_mgr
```

Inheriting Methods from Another Type Manager

Most of the **mp_mgr** methods are inherited from the **z_whole_copy** type manager, through symbolic links. You can enter the following commands as the *root* user in a Bourne shell:

```
# MP=$ATRIAHOME/lib/mgrs/mp_mgr
# for FILE in create_element create_version construct_version \
      create_branch delete_branches_versions \
      merge xmerge xcompare get_cont_info
> do
> ln –s ../z_whole_copy/$FILE $MP/$FILE
> done
#
```

Any methods that the new type manager does not support can be omitted from this list. The lack of a symbolic link causes ClearCase to generate an `Unknown Manager Request` error.

The following sections describe two of these inherited methods, **create_version** and **construct_version**, which can serve as models for user-defined methods. Both are actually implemented as scripts in the same file, *ccase-home-dir***/lib/mgrs/z_whole_copy/Zmgr**.

The create_version Method

The create_version method is invoked when a **checkin** command is entered. The create_version method of the **z_whole_copy** type manager does the following:

1. Compresses the data in the checked-out version
2. Stores the compressed data in a data container located in a source storage pool
3. Returns an exit status to the calling process, indicating what to do with the new data container

The file *ccase-home-dir***/lib/mgrs/mgr_info.h** lists the arguments passed to the method from the calling program (usually **cleartool** or **xclearcase**):

```
/*****************************************************************************
 * create_version
 * Store the data for a new version.
 * Store the version's data in the supplied new container, combining it
 * with the predecessor's data if desired (e.g for incremental deltas).
 *
 * Command line:
 *  create_version create_time new_branch_oid new_ver_oid new_ver_num
 *                 new_container_pname pred_branch_oid pred_ver_oid
 *                 pred_ver_num pred_container_pname data_pname
```

The only arguments that require special attention are `new_container_pname` (fifth argument), which specifies the pathname of the new data container, and `data_pname` (tenth argument), which specifies the pathname of the checked-out file.

The file *ccase-home-dir*/**lib/mgrs/mgr_info.sh** lists the appropriate exit statuses and provides a symbolic name for the create_version method:

```
# Any unexpected value is treated as failure
MGR_FAILED=1

# Return Values for store operations
MGR_STORE_KEEP_NEITHER=101
MGR_STORE_KEEP_JUST_OLD=102
MGR_STORE_KEEP_JUST_NEW=103
MGR_STORE_KEEP_BOTH=104
 .
 .
MGR_OP_CREATE_VERSION="create_version"
```

The example here is the code that implements the create_version method.

```
(1)     shift 1
(2)     if [ -s $4 ] ; then
(3)         echo '$0: error: new file is not of length 0!'
(4)         exit $MGR_FAILED
(5)     fi
(6)     if $gzip < $9 > $4 ; ret=$? ; then : ; fi
(7)     if [ "$ret" = "2" -o "$ret" = "0" ] ; then
(8)         exit $MGR_STORE_KEEP_BOTH
(9)     else
(10)        exit $MGR_FAILED
(11)    fi
```

The Bourne shell allows only nine command-line arguments. The `shift 1` in Line 1 discards the first argument (*create_time*), which is unneeded. Thus, the pathname of the checked-out version (*data_pname*), originally the tenth argument, becomes `$9`.

In Line 6, the contents of *data_pname* are compressed, then appended to the new, empty data container: *new_container_pname*, originally the fifth argument, but shifted to become `$4`. (Lines 2 through 5 verify that the new data container is, indeed, empty.)

Finally, the exit status of the **gzip** command is checked, and the appropriate value is returned (Lines 7 through 11). The exit status of the create_version method indicates that both the old data container (which contains the predecessor version) and the new data container (which contains the new version) are to be kept.

The construct_version Method

An element's construct_version *method* is invoked when standard UNIX software reads a particular version of the element (unless the contents are already cached in a *cleartext* storage pool). For example, the construct_version method of element **monet_admin.1** is invoked by the **view_server** when a user enters these commands:

```
% cp monet_admin.1 /usr/tmp          (read version selected by view)
% cat monet_admin.1@@/main/4         (read a specified version)
```

It is also invoked during a **checkout** command, which makes a view-private copy of the most recent version on a branch.

The construct_version method of the **z_whole_copy** type manager does the following:

1. Uncompresses the contents of the data container
2. Returns an exit status to the calling process, indicating what to do with the new data container

The file *ccase-home-dir*/**lib/mgrs/mgr_info.h** lists the arguments passed to the method.

```
/*****************************************************************************
 * construct_version
 * Fetch the data for a version.
 * Extract the data for the requested version into the supplied pathname, or
 * return a value indicating that the source container can be used as the
 * cleartext data for the version.
 *
 * Command line:
 *  construct_version source_container_pname data_pname version_oid
```

The file *ccase-home-dir*/**lib/mgrs/mgr_info.sh** lists the appropriate exit statuses and provides a symbolic name for the construct_version method:

```
# Any unexpected value is treated as failure
MGR_FAILED=1

# Return Values for construct operations
MGR_CONSTRUCT_USE_SRC_CONTAINER=101
MGR_CONSTRUCT_USE_NEW_FILE=102
 .
 .
MGR_OP_CONSTRUCT_VERSION="construct_version"
```

This example is the code that implements the construct_version method.

construct_version Method

```
(1)  if $gzip -d < $1 > $2 ; then
(2)      exit $MGR_CONSTRUCT_USE_NEW_FILE
(3)  else
(4)      exit $MGR_FAILED
(5)  fi
```

In Line 1, the contents of *source_container_pname* are uncompressed and stored in the cleartext container, *data_pname*. The remaining lines return the appropriate value to the calling process, depending on the success or failure of the **gzip** command.

Implementing a New compare Method

The compare method is invoked by a **cleartool diff** command. This method does the following:

1. Formats each version using **nroff(1)**, producing a pure-ASCII text file
2. Compares the formatted versions, using **cleardiff** or **xcleardiff**

The file *ccase-home-dir*/**lib/mgrs/mgr_info.h** lists the arguments passed to the method from **cleartool** or **xclearcase**.

```
/*****************************************************************************
 * compare
 * Compare the data for two or more versions.
 * For more information, see man page for cleartool diff.
 *
 * Command line:
 *  compare [-tiny | -window] [-serial | -diff | -parallel] [-columns n]
 *          [pass-through-options] pname pname ...
```

This listing shows that a user-supplied implementation of the compare method must accept all the command-line options that the ClearCase **diff** command supports. The strategy here is to pass the options to **cleardiff**, without attempting to interpret them. After all options are processed, the remaining arguments specify the files to be compared.

The file *ccase-home-dir*/**lib/mgrs/mgr_info.sh** lists the appropriate exit statuses and provides a symbolic name for the compare method.

```
# Return Values for COMPARE/MERGE Operations
MGR_COMPARE_NODIFFS=0
MGR_COMPARE_DIFF_OR_ERROR=1
 .
 .
MGR_OP_COMPARE="compare"
```

The Bourne shell script listed in *Script for compare Method* implements the compare method. (You can modify this script to implement the xcompare method as a slight variant of compare.)

Script for compare Method

```
#!/bin/sh -e
MGRDIR=${ATRIAHOME:-/usr/atria}/lib/mgrs

#### read file that defines methods and exit statuses
. $MGR_DIR/mgr_info.sh

#### process all options: pass them through to cleardiff
OPTS=""
while (expr $1 : '\-' > /dev/null) ; do
    OPTS="$OPTS $1"
    if [ "$1" = "$MGR_FLAG_COLUMNS" ] ; then
        shift 1
        OPTS="$OPTS $1"
    fi
    shift 1
done
#### all remaining arguments ($*) are files to be compared
#### first, format each file with NROFF
COUNT=1
TMP=/usr/tmp/compare.$$
for X in $* ; do
    nroff -man $X | col | ul -Tcrt > $TMP.$COUNT
    COUNT=`expr $COUNT + 1`
done

#### then, compare the files with cleardiff
cleardiff -quiet $OPTS $TMP.*

#### cleanup and return appropriate exit status
if [ $? -eq MGR_COMPARE_NODIFFS ] ; then
    rm -f $TMP.*
    exit MGR_COMPARE_NODIFFS
else
    rm -f $TMP.*
    exit MGR_COMPARE_DIFF_OR_ERROR
fi
```

Testing the Type Manager

You can test a new type manager only by using it on some ClearCase host. This process need not be obtrusive. Because the type manager has a new name, no existing element type—and therefore, no existing element—uses it automatically. To place the type manager in service, create a new element type, create some test elements of that type, and run some tests.

The following testing sequence continues the **mp_mgr** example.

Creating a Test Element Type. To make sure that an untested type manager is not used accidentally, associate it with a new element type, **manpage_test**, of which you are the only user.

```
% cleartool mkeltype –nc –supertype compressed_file \
      –manager mp_mgr manpage_test
% cleartool lock –nusers $USER eltype:manpage_test
```

Creating and Using a Test Element. These commands create a test element that uses the new type manager, and tests the various data-manipulation methods:

```
% cd directory-in-test-VOB
% cleartool checkout –nc .                              (tests create_element method)
% cleartool mkelem –eltype manpage_test –nc –nco test.1
% cleartool checkout –nc test.1                         (tests construct_version method)
% vi test.1                                             (edit checked-out version)
% cleartool checkin –c "first" test.1                   (tests create_ version method)
% cleartool checkout –nc test.1                         (tests construct_ version method)
% vi test.1                                             (edit checked-out version)
% cleartool checkin –c "second" test.1                  (tests create_ version method)
% cleartool diff test.1@@/main/1 test.1@@/main/2        (tests compare method)
```

Installing and Using the Type Manager

After a type manager has been fully tested, you can make it available to all users with the following procedure.

1. Install the type manager.

   A VOB is a networkwide resource; it can be mounted on any ClearCase host. But a type manager is a host resource: a separate copy must be installed on each host where ClearCase client programs run. If the copy is not installed, elements of the new type cannot be used. (It need not be installed on hosts that serve only as repositories for VOBs and/or views.)

   To install the type manager on a particular host, create a subdirectory in *ccase-home-dir***/lib/mgrs**, and populate it with the programs that implement the *method*s. You can create symbolic links across the network to a master copy on a server host.

2. Create element types.

   Create one or more element types that use the type manager, just as you did in *Testing the Type Manager* (do not include "test" in the name of the element type). For example, you can name the element type **manpage** or **nroff_src**.

3. Convert existing elements.

   You'll probably want to have at least a few existing elements use the new type manager. The **chtype** command changes an element's type:

   ```
   % cleartool chtype –force manpage pathname ...
   ```

   Permission to change an element's type is restricted to the element's owner, the VOB owner, and the *root* user.

4. Revise magic files.

   If you want the new element types to be used automatically for certain newly created elements, create (or update) a **local.magic** file in each host's *ccase-home-dir*/**config/magic** directory:

   ```
   manpage src_file text_file file: -name "*.[1-9]" ;
   ```

5. Inform the project team (and other teams, if appropriate).

   Advertise the new element types to all team members, describing the features and benefits of the new type manager. Be sure to provide directions on how to gain access to the new functionality automatically (through file names that match magic-file rules) and explicitly (with **mkelem –eltype**).

12.7 Icon Use by GUI Browsers

An **xclearcase** browser can display file-system objects either by name or graphically. In the latter case, **xclearcase** selects an icon for each file-system object as follows:

1. The object's name or its contents determines a list of file types, as described in *How ClearCase Assigns Element Types* on page 156.

2. One by one, the file types are compared to the rules in predefined and user-defined *icon* files, as described in the **cc.icon** reference page. For example, the file type **c_source** matches this icon file rule:

   ```
   c_source : -icon c ;
   ```

   When a match is found, the search ends. The token that follows **–icon** names the file that contains the icon to be displayed.

3. **xclearcase** searches for the file, which must be in **bitmap(1)** format, in directory **$HOME/.bitmaps**, or *ccase-home-dir*/**config/ui/bitmaps**, or the directories specified by the environment variable **BITMAP_PATH**.

4. If a valid bitmap file is found, **xclearcase** displays it; otherwise, the search for an icon continues with the next file type.

The name of an icon file must include a numeric extension, which need not be specified in the icon file rule. The extension specifies how much screen space **xclearcase** must allocate for the icon. Each bitmap supplied with ClearCase is stored in a file with a **.40** suffix (for example, **lib.40**), indicating a 40x40 icon.

This procedure causes **xclearcase** to display manual page source files with a customized icon. All manual pages have file type **manpage**.

1. Add a rule to your personal magic file (in directory **$HOME/.magic**) that includes **manpage** among the file types assigned to all manual page source files:

```
manpage src_file text_file file:  -name "*.[1-9]" ;
```

2. Add a rule to your personal icon file (in directory **$HOME/.icon**) that maps **manpage** to a user-defined bitmap file:

```
manpage : -icon manual_page_icon ;
```

3. Create a **manpage** icon in your personal bitmaps directory (**$HOME/.bitmaps**) by revising one of the standard icon bitmaps with the standard X **bitmap** utility:

```
% mkdir $HOME/.bitmaps
% cd $HOME/.bitmaps
% cp $ATRIAHOME/config/ui/bitmaps/c.40 manual_page_icon.40
% bitmap manual_page_icon.40
```

4. Test your work by having an **xclearcase** graphical directory browser display a manual page source file (Figure 43).

Figure 43 User-Defined Icon Display

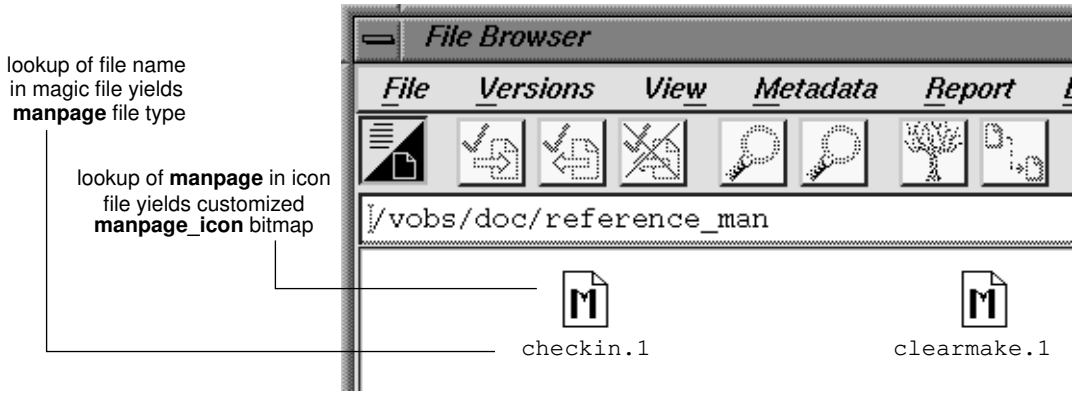

13 Using ClearCase Throughout the Development Cycle

The previous chapters describe various aspects of managing a project with ClearCase. This chapter presents one way in which you can use ClearCase to organize the work throughout a development project. During this cycle, developers create a new release and maintain the previous release.

This chapter describes concepts and methods to address typical organizational needs. There are many other approaches that ClearCase supports.

13.1 Project Overview

Release 2.0 development of the **monet** project includes the following kinds of work:

- Patches. Several high-priority bug fixes to Release 1.0 are needed.
- Minor enhancements. Some commands need new options; some option names need to be shortened (**–recursive** becomes **–r**); some algorithms need performance work.
- Major new features. A graphical user interface is required, as are many new commands and internationalization support.

These three development efforts can proceed largely in parallel (Figure 44), but critical dependencies and milestones must be considered:

- Several Release 1.0 patch releases will ship before Release 2.0 is complete.
- New features take longer to complete than minor enhancements.

➤ Some new features depend on the minor enhancements.

Figure 44 Project Plan for Release 2.0 Development

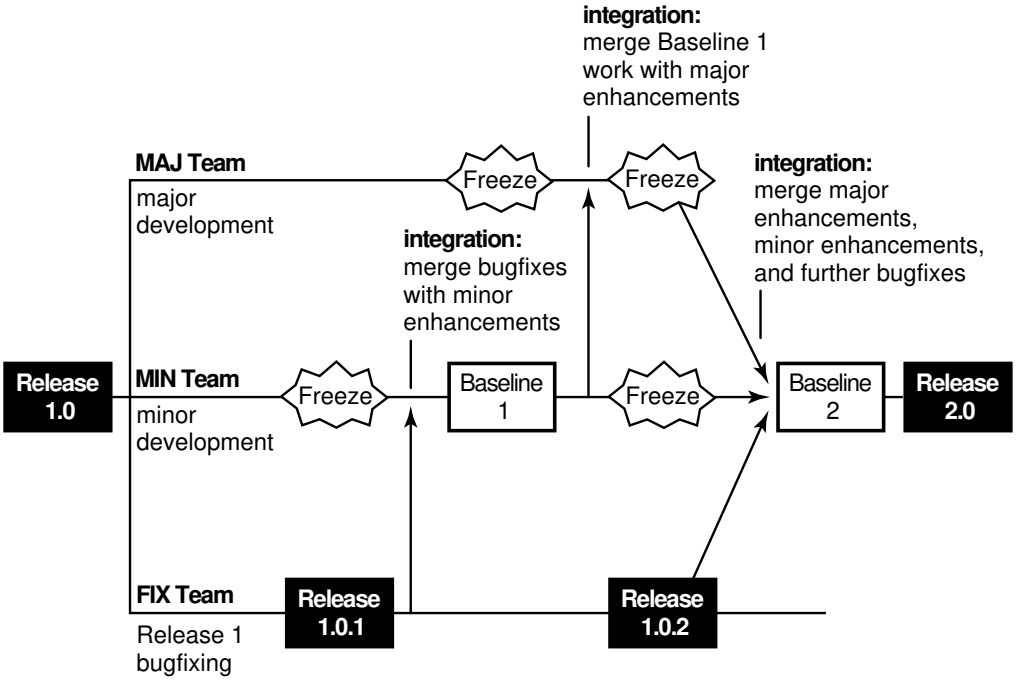

The plan uses a baseline-plus-changes approach. Periodically, developers stop writing new code, and spend some time integrating their work, building, and testing. The result is a *baseline*: a stable, working version of the application. ClearCase makes it easy to integrate product enhancements incrementally and frequently. The more frequent the baselines, the easier the tasks of merging work and testing the results.

After a baseline is produced, active development resumes; any new efforts begin with the set of source versions that went into the baseline build.

You define a baseline by assigning the same version label (for example, **R2_BL1** for Release 2.0, Baseline 1) to all the versions that go into, or are produced by, the baseline build.

The project team is divided into three smaller teams, each working on a different development effort: the MAJ team (new features), the MIN team (minor enhancements), and the FIX team (Release 1.0 bug fixes and patches).

NOTE: Some developers may belong to multiple teams. These developers work in multiple views, each configured for the respective team's tasks.

The development area for the **monet** project is shown here. At the beginning of Release 2.0 development, the most recent versions on the **main** branch are labeled **R1.0**.

```
/vobs/monet             (project top-level directory)
       src/             (sources)
       include/         (include files)
       lib/             (shared libraries)
```

13.2 Development Strategy

This section describes the ClearCase issues to be resolved before development begins.

Project Manager and ClearCase Administrator

In most development efforts, the project manager and the system administrator are different people. The user name of the project manager is **meister**. The administrator is the **vobadm** user, who creates and owns the **monet** and **libpub** VOBs.

Use of Branches

In general, different kinds of work is done on different branches. The Release 1.0 bug fixes, for example, are made on a separate branch to isolate this work from new development. The FIX team can then create patch releases that do not include any of the Release 2.0 enhancements or incompatibilities.

Because the MIN team will produce the first baseline release on its own, the project manager gives the **main** branch to this team. The MAJ team will develop new features on a subbranch, and will not be ready to integrate for a while; the FIX team will fix Release 1.0 bugs on another subbranch and can integrate its changes at any time.

Each new feature can be developed on its own subbranch, to better manage integration and testing work. For simplicity, this chapter assumes that work for new features is done on a single branch.

The project manager has created the first baseline from versions on the **main** branches of their elements. But this is not a requirement; you can create a release that uses versions on any branch, or combination of branches.

Figure 45 shows the evolution of a typical element during Release 2.0 development, and indicates correspondences to the overall project plan (Figure 44).

Figure 45 Development Milestones: Evolution of a Typical Element

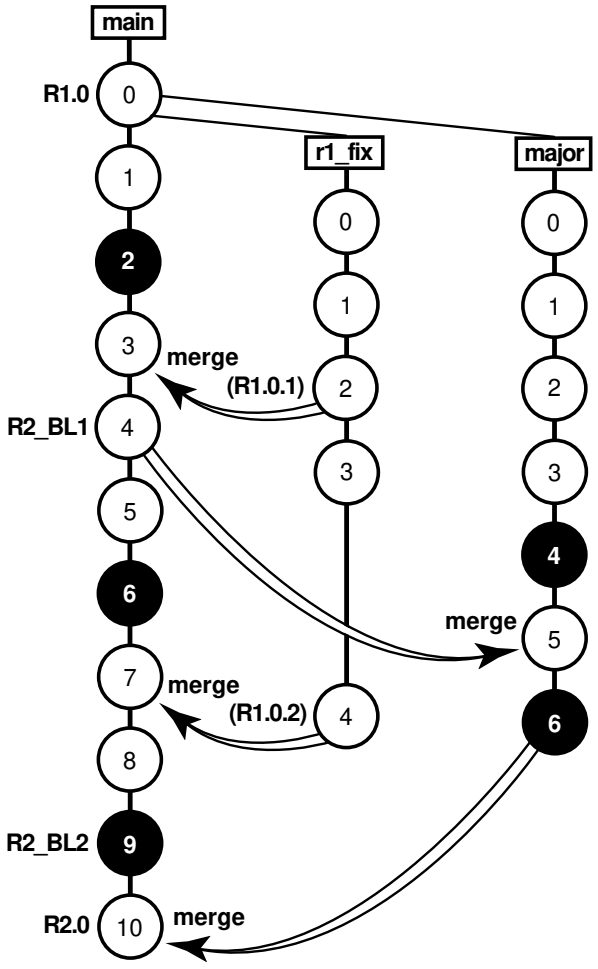

1. (All branches) Start minor and major enhancements, along with **R1.0** bug fixing
2. (**main**) Freeze minor enhancements work
3. (**main**) Merge bug fixes from Release 1.0.1 into minor enhancements
4. (**main**) Baseline 1 release
5. (**major**) Freeze major enhancements work
6. (**major**) Merge Baseline 1 changes into major enhancements
7. (**main**) Freeze minor enhancements work
8. (**main**) Merge additional bugfixes into minor enhancements
9. (**major**) Freeze major enhancements work
10. (**main**) Merge major enhancements work with minor enhancements work
11. (**main**) Baseline 2 release
12. (**main**) Final testing period
13. (**main**) Release 2.0

The MAJ team works on a branch named **major** and uses this config spec:

```
(1)  element * CHECKEDOUT
(2)  element * .../major/LATEST
(3)  element * R1.0 -mkbranch major
(4)  element * /main/LATEST -mkbranch major
```

The MIN team works on the **main** branch and uses the default config spec:

```
(1)  element * CHECKEDOUT
(2)  element * .../main/LATEST
```

The FIX team works on a branch named **r1_fix** and uses this config spec:

```
(1)  element * CHECKEDOUT
(2)  element * .../r1_fix/LATEST
(3)  element * R1.0 -mkbranch r1_fix
(4)  element * /main/LATEST -mkbranch r1_fix
```

For the MAJ and FIX teams, use of the *auto-make-branch* facility in Rule (3) and Rule (4) enforces consistent use of subbranches. It also relieves developers of the task of creating branches explicitly and ensures that all branches are created at the version labeled **R1.0**.

13.3 Creating Branch Types

The project manager creates the **major** and **r1_fix** branch types required for the config specs in *Creating Project Views* on page 178:

**cleartool mkbrtype –c "monet R2 major enhancements" **
major@/vobs/libpub major@/vobs/monet

```
Created branch type "major".
Created branch type "major".
```

cleartool mkbrtype –c "monet R1 bugfixes" r1_fix@/vobs/libpub r1_fix@/vobs/monet

```
Created branch type "r1_fix".
Created branch type "r1_fix".
```

NOTE: Because each VOB has its own set of branch types, the branch types must be created separately in the **monet** VOB and the **libpub** VOB.

13.4 Creating Standard Config Specs

To ensure that all developers in a team configure their views the same way, the project manager creates files containing standard config specs:

- **/public/config_specs/MAJ** contains the MAJ team's config spec.
- **/public/config_specs/FIX** contains the FIX team's config spec.

These config spec files are stored in a standard directory outside a VOB, to ensure that all developers get the same version.

13.5 Creating, Configuring, and Registering Views

Each developer creates a view under his or her home directory. For example, developer **arb** enters these commands:

```
% mkdir $HOME/view_store
% cleartool mkview –tag arb_major $HOME/view_store/arb_major.vws
Created view.
Host-local path: phobos:export/home/arb/view_store/arb_major.vws
Global path:     /net/phobos/export/home/arb/view_store/arb_major.vws
It has the following rights:
User : arb      : rwx
Group: user     : rwx
Other:          : r-x
```

A new view has the default config spec. Thus, developers on the MAJ and FIX teams must reconfigure their views, using the standard file for their team. **arb** edits her config spec with the **cleartool edcs** command, deletes the existing lines, and adds the following line:

```
/public/config_specs/MAJ
```

If the project manager changes the standard file, **arb** must enter the command **cleartool setcs –current** to pick up the changes.

13.6 Development Begins

To begin the project, a developer sets a properly configured view, checks out one or more elements, and starts work. For example, developer **david** on the MAJ team enters these commands:

```
% cleartool setview david_major
% cd /vobs/monet/src
% cleartool checkout –nc opt.c prs.c
Created branch "major" from "opt.c" version "/main/6".
Checked out "opt.c" from version "/main/major/0".
Created branch "major" from "prs.c" version "/main/7".
Checked out "prs.c" from version "/main/major/0".
```

The auto-make-branch facility causes each element to be checked out on the **major** branch (see Rule (4) in the MAJ team's config spec in *Creating Project Views* on page 178). If a developer on the MIN team enters this command, the elements are checked out on the **main** branch, with no conflict.

ClearCase is fully compatible with standard development tools and practices. Thus, developers use the editing, compilation, and debugging tools they prefer (including personal scripts and aliases) while working in their views.

Developers check in work periodically to make their work available to other team members (that is, those whose views select the most recent version on the team's branch). This allows intrateam integration and testing to proceed throughout the development period.

Techniques for Isolating Your Work

Individual developers may need or prefer to isolate their work from the changes made by other team members. To do so, they can use these techniques to configure their views:

- Time rules. When someone checks in an incompatible change, a developer can reconfigure the view to select the versions at a point before those changes were made.
- Private subbranches. A developer can create a private subbranch in one or more elements (for example, **/main/major/anne_wk**). The config spec must be changed to select versions on the **/main/major/anne_wk** branch instead of versions on the **/main/major** branch.
- Viewing only their own revisions. Developers can use a ClearCase query to configure a view that sees only their own revisions to the source tree.

13.7 Creating Baseline 1

The MIN team has implemented and tested the first group of minor enhancements, and the FIX team has produced a patch release, whose versions are labeled **R1.0.1**. It is time to combine these efforts, to produce Baseline 1 of Release 2.0 (Figure 46).

Figure 46 Creating Baseline 1

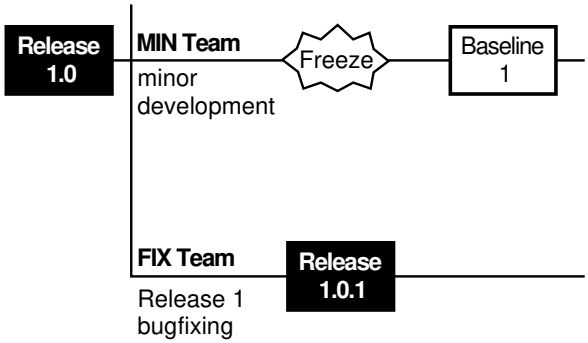

Merging Two Branches

The project manager asks the MIN developers to merge the **R1.0.1** changes from the **r1_fix** branch to their own branch (**main**). All the changes can be merged by using the **findmerge** command once. For example:

```
% cleartool findmerge /vobs/libpub /vobs/monet/src \
      –fversion .../r1_fix/LATEST –merge –graphical
.
. <lots of output>
.
```

Integration and Test

After the merges are complete, the **/main/LATEST** versions of certain elements represent the efforts of the MIN and FIX teams. Members of the MIN team now compile and test the **monet** application to find and fix incompatibilities in the work of both teams.

The developers on the MIN team integrate their changes in a single, shared view. The project manager creates the view storage area in a location that is accessible from all developer hosts:

```
% umask 2
% mkdir /netwide/public
% cleartool mkview –tag base1_vu /netwide/public/base1_vu.vws
Created view.
Host-local path: infinity:/netwide/public/base1_vu.vws
Global path:     /net/infinity/netwide/public/base1_vu.vws.
It has the following rights:
User : meister  : rwx
Group: mon      : rwx
Other:          : r-x
```

The **umask** value of 2 allows all members of the MIN team to use the view. Any member of the **mon** group can use the following command to begin working in this view:

```
% cleartool setview base1_vu
```

Because all integration work takes place on the **main** branch, there is no need to change the configuration of the new view from the ClearCase default. MIN developers set this view (**cleartool setview base1_vu**) and coordinate builds and tests of the **monet** application. Because they are sharing a single view, the developers are careful not to overwrite each other's view-private files. Any new versions created to fix inconsistencies (and other bugs) go onto the **main** branch.

Labeling Sources

The **monet** application's minor enhancements and bug fixes are now integrated, and a clean build has been performed in view **base1_vu**. To create the baseline, the project manager assigns the same version label, **R2_BL1**, to the **/main/LATEST** versions of all source elements. He begins by creating an appropriate label type:

```
% cleartool mklbtype –c "Release, Baseline 1" R2_BL1@/vobs/monet R2_BL1@/vobs/libpub
Created label type "R2_BL1".
Created label type "R2_BL1".
```

He then locks the label type, preventing all developers (except himself) from using it:

```
% cleartool lock –nusers meister lbtype:R2_BL1@/vobs/monet lbtype:R2_BL1@/vobs/libpub
Locked label type "R2_BL1".
Locked label type "R2_BL1".
```

Before applying labels, he verifies that all elements are checked in on the **main** branch (checkouts on other branches are still permitted):

```
% cleartool lscheckout –all /vobs/monet /vobs/libpub
```

No output from this command indicates that all elements for the **monet** project are checked in. Now, the project manager attaches the **R2_BL1** label to the currently selected version (**/main/LATEST**) of every element in the two VOBs:

```
% cleartool mklabel –recurse R2_BL1 /vobs/monet /vobs/libpub
Created label "R2_BL1" on "/vobs/monet" version "/main/1".
Created label "R2_BL1" on "/vobs/monet/src" version "/main/3".
 <many more label messages>
```

Removing the Integration View

The view registered as **base1_vu** is no longer needed, so the project manager removes it:

```
% cleartool rmview –force –tag base1_vu
```

13.8 Merging Ongoing Development Work

After Baseline 1 is created, the MAJ team merges the Baseline 1 changes into its work (Figure 47). The team now has access to the minor enhancements it needs for further development. Team members also have an early opportunity to determine whether any of their changes are incompatible.

Figure 47 Updating Major Enhancements Development

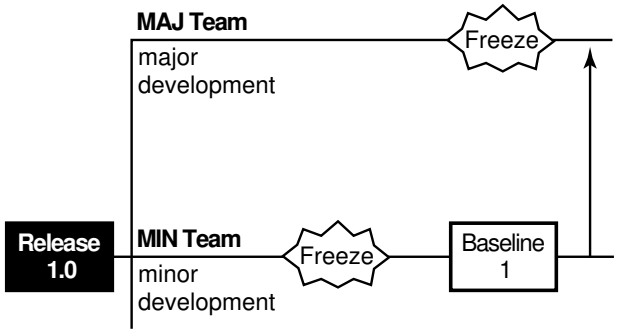

Accordingly, the project manager declares a freeze of major enhancements development. MAJ team members check in all elements and verify that the **monet** application builds and runs, making small source changes as necessary. When all such changes have been checked in, the team has a consistent set of **/main/major/LATEST** versions.

NOTE: Developers working on other major enhancements branches can merge at other times, using the same merge procedures described here.

Preparing to Merge

1. The project manager makes sure that no element is checked out on the **major** branch:

   ```
   % cleartool lscheckout –all /vobs/monet /vobs/libpub
   ```

   NOTE: Any MAJ team members who want to continue with nonmerge work can create a subbranch at the "frozen" version (or work with a version that is checked out as unreserved).

2. The project manager performs any required directory merges:

```
% cleartool setview major_vu                    (use any MAJ team view)

% cleartool findmerge /vobs/monet /vobs/libpub –type d \
–fversion /main/LATEST –merge
Needs merge /vobs/monet/src [automatic to /main/major/3 from /main/LATEST]
   .
   . <lots of output>
   .
Log has been written to "findmerge.log.04-Feb-99.09:58:25".
```

3. After checking in the files, the project manager determines which elements need to be merged:

```
% cleartool findmerge /vobs/monet /vobs/libpub –fversion /main/LATEST –print
   .
   . <lots of output>
   .
A 'findmerge' log has been written to
"findmerge.log.04-Feb-99.10:01:23"
```

This last **findmerge** log file is in the form of a shell script: it contains a series of **cleartool findmerge** commands, each of which performs the required merge for one element:

```
% cat findmerge.log.04-Feb-99.10:01:23
cleartool findmerge /vobs/monet/src/opt.c@@/main/major/1 –fver /main/LATEST –merge
cleartool findmerge /vobs/monet/src/prs.c@@/main/major/3 –fver /main/LATEST –merge
   .
   .
cleartool findmerge /vobs/libpub/src/dcanon.c@@/main/major/3 –fver /main/LATEST –merge
cleartool findmerge /vobs/libpub/src/getcwd.c@@/main/major/2 –fver /main/LATEST –merge
cleartool findmerge /vobs/libpub/src/lineseq.c@@/main/major/10 –fver /main/LATEST –merge
```

4. The project manager locks the **major** branch, allowing it to be used only by the developers who are performing the merges:

```
cleartool lock –nusers meister,arb,david,sakai brtype:major@/vobs/monet \
brtype:major@/vobs/libpub
Locked branch type "major".
Locked branch type "major".
```

Merging Work

Because the MAJ team is not contributing to a baseline soon, it is not essential to merge work (and test the results) in a shared view. MAJ developers can continue working in their own views.

Periodically, the project manager sends an excerpt from the **findmerge** log to an individual developer, who executes the commands and monitors the results. (The developer can send the resulting log files back to the project manager, as confirmation of the merge activity.)

A merged version of an element includes changes from three development efforts: Release 1.0 bug fixing, minor enhancements, and new features (Figure 48).

Figure 48 Merging Baseline 1 Changes into the major Branch

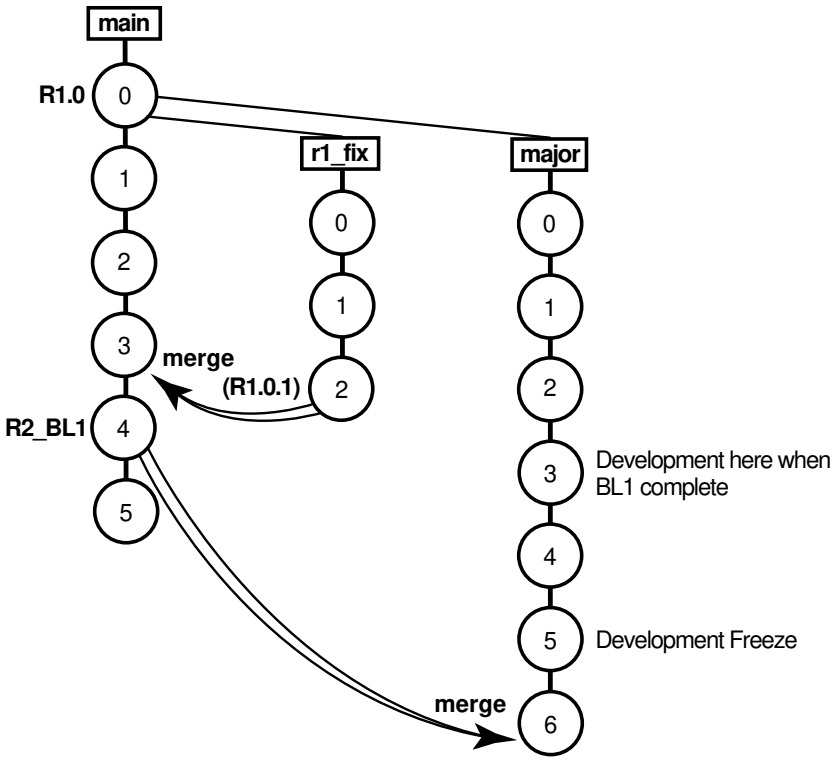

The project manager verifies that no more merges are needed, by entering a **findmerge** command with the **–whynot** option:

```
% cleartool findmerge /vobs/monet /vobs/libpub –fversion /main/LATEST –whynot –print
  .
  .
No merge "/vobs/monet/src" [/main/major/4 already merged from /main/3]
No merge "/vobs/monet/src/opt.c" [/main/major/2 already merged from /main/12]
  .
  .
```

The merge period ends when the project manager removes the lock on the **major** branch:

```
% cleartool unlock brtype:major@/vobs/monet brtype:major@/vobs/libpub
Unlocked branch type "major".
Unlocked branch type "major".
```

13.9 Creating Baseline 2

The MIN team is ready to freeze for Baseline 2, and the MAJ team will be soon (Figure 49). Baseline 2 will integrate all three development efforts, thus requiring two sets of merges:

- Bug fix changes from the most recent patch release (versions labeled **R1.0.2**) must be merged to the **main** branch.
- New features must be merged from the **major** branch to the **main** branch. (This is the opposite direction from the merges described in *Merging Ongoing Development Work* on page 183.)

Figure 49 Baseline 2

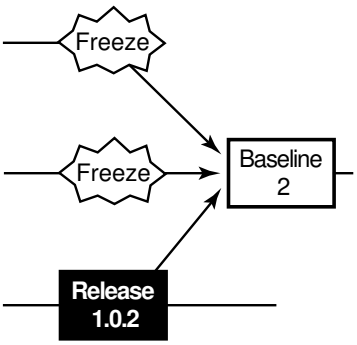

ClearCase supports merges from more than two directions, so both the bug fixes and the new features can be merged to the **main** branch at the same time. In general, though, it is easier to verify the results of two-way merges.

Merging from the r1_fix Branch

The first set of merges is almost identical to those described in *Merging Two Branches* on page 181.

Preparing to Merge from the major Branch

After the integration of the **r1_fix** branch is completed, the project manager prepares to manage the merges from the **major** branch. These merges are performed in a tightly controlled environment, because the Baseline 2 milestone is approaching and the **major** branch is to be abandoned.

NOTE: It is probably more realistic to build and verify the application, and then apply version labels before proceeding to the next merge.

The project manager verifies that everything is checked in on both the **main** branch and **major** branches:

```
% cleartool lscheckout –brtype main –recurse /vobs/monet /vobs/libpub
% cleartool lscheckout –brtype major –recurse /vobs/monet /vobs/libpub
%
```

No output from these commands indicates that no element is checked out on either its **main** branch or its **major** branch.

Next, the project manager determines which elements require merges:

```
% cleartool setview minor_vu                    (use any MIN team view)

% cleartool findmerge /vobs/monet /vobs/libpub –fversion .../major/LATEST –print
   .
   .  <lots of output>
   .
A 'findmerge' log has been written to
"findmerge.log.26-Feb-99.19:18:14"
```

All development on the **major** branch will stop after this baseline. Thus, the project manager locks the **major** branch to all users, except those who are performing the merges; locking allows ClearCase to record the merges with a hyperlink of type **Merge**:

```
% cleartool lock –nusers arb,david brtype:major@/vobs/monet brtype:major@/vobs/libpub
Locked branch type "major".
Locked branch type "major".
```

Because the **main** branch will be used for Baseline 2 integration by a small group of developers, the project manager asked **vobadm** to lock the **main** branch to everyone else:

```
% cleartool lock –nusers meister,arb,david,sakai \
brtype:main@/vobs/monet  brtype:main@/vobs/libpub
Locked branch type "main".
Locked branch type "main".
```

(To lock the branch, you must be the branch creator, element owner, VOB owner, or *root* user. See the **lock** reference page.)

Merging from the major Branch

Because the **main** branch is the destination of the merges, developers work in a view with the default config spec. The situation is similar to the one described in *Preparing to Merge* on page 184. This time, the merges take place in the opposite direction, from the **major** branch to the **main** branch. Accordingly, the **findmerge** command is very similar:

```
% cleartool findmerge /vobs/monet /vobs/libpub –fversion /main/major/LATEST \
–merge –graphical
   .
   .  <lots of output>
   .
A 'findmerge' log has been written to
"findmerge.log.23-Mar-99.14:11:53"
```

After checkin, the version tree of a typical merged element appears as in Figure 50.

Figure 50 Element Structure after the Pre-Baseline-2 Merge

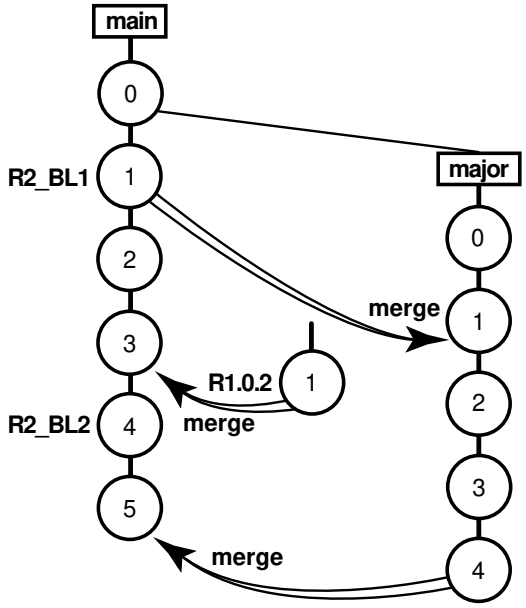

Decommissioning the major Branch

After all data has been merged to the **main** branch, development on the **major** branch will stop. The project manager enforces this policy by making the **major** branch obsolete:

```
% cleartool lock –replace –obsolete brtype:major@/vobs/monet  brtype:major@/vobs/libpub
Locked branch type "major".
Locked branch type "major".
```

Integration and Test

Structurally, the Baseline 2 integration-and-test phase is identical to the one for Baseline 1 (see *Integration and Test* on page 182). At the end of the integration period, the project manager attaches version label **R2_BL2** to the **/main/LATEST** version of each element in the **monet** and **libpub** VOBs. (The Baseline 1 version label was **R2_BL1**.)

13.10 Final Validation: Creating Release 2.0

Baseline 2 has been released internally, and further testing has found only minor bugs. These bugs have been fixed by creating new versions on the **main** branch (Figure 51).

Figure 51 Final Test and Release

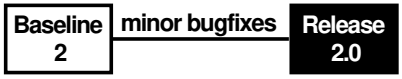

Before it is shipped to customers, the **monet** application goes through a validation phase:

- All editing, building, and testing is restricted to a single, shared view.
- All builds are performed from sources with a particular version label (**R2.0**).
- Only the project manager has permission to make changes involving that label.

- All labels must be moved by hand.
- Only high-priority bugs are fixed, using this procedure:

  a. The project manager authorizes a particular developer to fix the bug, by granting her permission to create new versions (on the **main** branch).

  b. The developer's checkin activity is tracked by a ClearCase trigger.

  c. After the bug is fixed, the project manager moves the **R2.0** version label to the fixed version and revokes the developer's permission to create new versions.

Labeling Sources

In a view with the default config spec, the project manager creates the **R2.0** label type and locks it:

cleartool mklbtype –c "Release 2.0" R2.0@/vobs/monet R2.0@/vobs/libpub
```
Created label type "R2.0".
Created label type "R2.0".
```

cleartool lock –nusers meister lbtype:R2.0@/vobs/monet lbtype:R2.0@/vobs/libpub
```
Locked label type "R2.0".
Locked label type "R2.0".
```

The project manager labels the **/main/LATEST** versions throughout the entire **monet** and **libpub** development trees:

cleartool mklabel –recurse R2.0 /vobs/monet /vobs/libpub
<*many label messages*>

During the final test phase, the project manager moves the label forward, using **mklabel –replace**, if any new versions are created.

Restricting Use of the main Branch

At this point, use of the **main** branch is restricted to a few users: those who performed the merges and integration leading up to Baseline 2 (see *Merging from the major Branch* on page 190). Now, the project manager asks **vobadm** to close down the **main** branch to everyone except himself, **meister**:

```
% cleartool lock –replace –nusers meister brtype:main
Locked branch type "main".
```

The **main** branch is opened only for last-minute bug fixes (see *Fixing a Final Bug* on page 194.)

Setting Up the Test View

The project manager creates a new shared view, **r2_vu**, that is configured with a one-rule config spec:

```
% umask 2
% cleartool mkview –tag r2_vu /public/integrate_r2.vws
% cleartool edcs –tag r2_vu
```

This is the config spec:

```
element * R2.0
```

This config spec guarantees that only properly labeled versions are included in final validation builds.

Setting Up the Trigger to Monitor Bugfixing

The project manager places a trigger on all elements in the **monet** and **libpub** VOBs; the trigger fires whenever a new version of any element is checked in. First, he creates a script that sends mail (for an example script, see *Policy: Notify Team Members of Relevant Changes* on page 134).

Then, he asks **vobadm** to create an all-element trigger type in the **monet** and **libpub** VOBs, specifying the script as the trigger action:

```
% cleartool mktrtype –nc -element –all –postop checkin –brtype main \
–exec /public/scripts/notify_manager.sh \
r2_checkin@/vobs/monet r2_checkin@/vobs/libpub
Created trigger type "r2_checkin".
Created trigger type "r2_checkin".
```

Only the VOB owner or *root* user can create trigger types.

Fixing a Final Bug

This section demonstrates the final validation environment in action. Developer **arb** discovers a serious bug and requests permission to fix it. The project manager grants her permission to create new versions on the **main** branch, by having **vobadm** enter this command.

```
% cleartool lock –replace –nusers arb,meister brtype:main
Locked branch type "main".
```

arb fixes the bug in a view with the default config spec and tests the fix there. This involves creating two new versions of element **prs.c** and one new version of element **opt.c**. Each time **arb** uses the **checkin** command, the **r2_checkin** trigger sends mail to the project manager. For example:

```
Subject: Checkin /vobs/monet/src/opt.c by arb
/vobs/monet/src/opt.c@@/main/9
Checked in by arb.

Comments:
fixed bug #459: made buffer larger
```

When regression tests verify that the bug has been fixed, the project manager revokes **arb**'s permission to create new versions. Once again, the command is executed by **vobadm**:

```
% cleartool lock –replace –nusers meister brtype:main
Locked branch type "main".
```

The project manager then moves the version labels to the new versions of **prs.c** and **opt.c**, as indicated in the mail messages. For example:

```
% cleartool mklabel –replace R2.0 /vobs/monet/src/opt.c@@/main/9
Moved label "R2.0" on "prs.c" from version "/main/8" to "/main/9".
```

Rebuilding from Labels

After the labels have been moved, developers rebuild the **monet** application again, to verify that a good build can be performed using only those versions labeled **R2.0.**

Wrapping Up

When the final build in the **r2_vu** passes the final test, Release 2.0 of **monet** is ready to ship. After the distribution medium has been created from derived objects in the **r2_vu**, the project manager asks the ClearCase administrator to clean up and prepare for the next release:

- The ClearCase administrator removes the checkin triggers from all elements by deleting the all-element trigger type:

  cleartool rmtype trtype:r2_checkin@/vobs/monet trtype:r2_checkin@/vobs/libpub
  ```
  Removed trigger type "r2_checkin".
  Removed trigger type "r2_checkin".
  ```

- The ClearCase administrator reopens the **main** branch:

  cleartool unlock brtype:main
  ```
  Unlocked branch type "main".
  ```

ClearCase-ClearQuest Integrations

A

A.1 Understanding the Two ClearCase-ClearQuest Integrations

The integration of ClearQuest and ClearCase associates one or more ClearQuest records with one or more ClearCase versions allowing you to use features of each product. ClearCase supports two separate integrations with ClearQuest:

- The base ClearCase-ClearQuest integration
- The UCM-ClearQuest integration

Online help that describes the base ClearCase-ClearQuest integration is available from within the ClearQuest Integration Configuration GUI on Windows. To open the GUI, click **Start>Programs>ClearCase Administration>Integrations>ClearQuest Integration Configuration**. Note that this integration cannot be used with UCM projects.

For further information on the UCM-ClearCase integration, see Part 1 of this book.

In general, we recommend that you use the base ClearCase and UCM integrations separately, and avoid using a common ClearQuest user database. However, it is possible for both integrations to use the same ClearQuest user database. This can be useful if you are moving a project to UCM and have a substantial amount of information in a ClearQuest user database that was created with the base ClearCase-ClearQuest integration. You may want the new work in UCM to be reflected in new ClearQuest records in the same ClearQuest user database.

The remainder of this appendix discusses considerations in managing the coexistence of the base ClearCase-ClearQuest integration and the UCM-ClearQuest integration.

Managing Coexisting Integrations

When a ClearQuest user database that had been integrated with ClearCase previously is configured for integration with UCM, the existing change sets are preserved intact in the ClearQuest user database, but cannot be migrated to the UCM integration.

Change sets of existing records in the ClearQuest user database are preserved, and you can access them from ClearQuest. To continue work on a task in a project that has been migrated to UCM, create a new, corresponding, UCM activity and continue work there.

See *Planning How to Use the UCM-ClearQuest Integration* on page 36 for related information.

Schema

A ClearQuest schema can contain modifications from both the base ClearCase-ClearQuest integration and the UCM-ClearQuest integration. A record type in such a schema would include both the ClearCase package and the Unified Change Management package.

An individual record of that record type can store either ClearCase or UCM change set information, but not both.

Presentation

The form for a record type that uses both integrations includes two tabs to show the change set information associated with each integration, as shown in Figure 52. The **Unified Change Management** tab lists the change set for a UCM activity. The **ClearCase** tab shows the change set associated with a ClearQuest record.

Figure 52 Change Sets in ClearQuest GUI

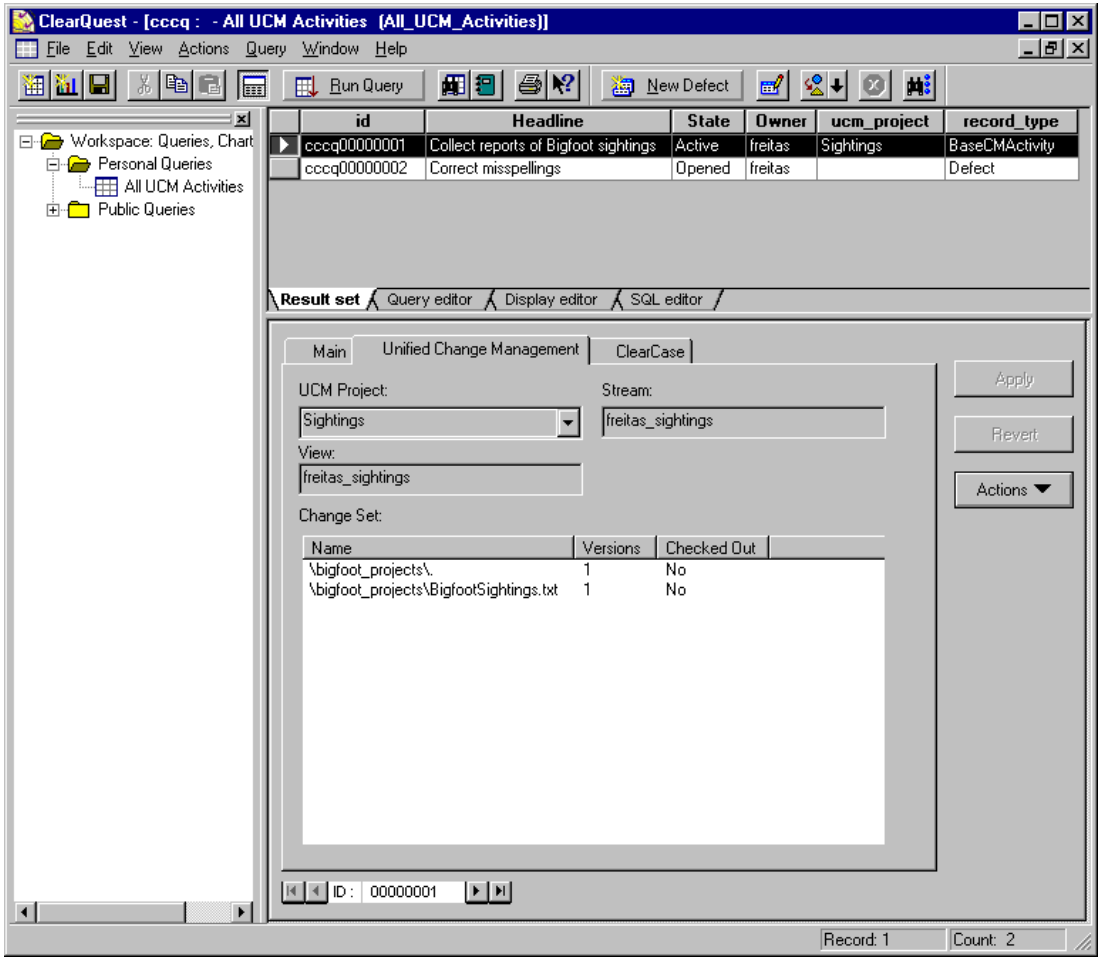

Index

A

activities
- about 10
- creating and assigning in ClearQuest (procedure) 69
- creating and setting in new project (procedure) 61
- decomposing in ClearQuest 52
- fixing ClearQuest links 70
- migrating to ClearQuest integration 67
- state transition after delivery 43
- verifying owner of 43

administrative VOBs and PVOBs 34

assignments, verifying 22

***attache-home-dir* directory** xix

attributes
- about 103
- change request policy 136
- use in config specs 116
- use in monitoring project status 130

B

base ClearCase and UCM, compared 1

baselines in base ClearCase
- creating, extended example 181, 187
- labeling policy 131

baselines in UCM
- about 14
- benefits of frequent 31
- comparing (procedure) 83
- creating 17
- creating for imported files (procedure) 64
- creating new (procedures) 77
- creating streams for testing (procedure) 71
- fixing problems (procedure) 81
- foundation 58
- naming convention 32
- promoting and demoting (procedure) 82
- promotion levels 18
- recommended, promotion policy 41
- strategy for 30
- test planning 32
- when to delete 86

branch types, example 178

branches
- about 100
- bug-fix policy 132
- config spec rules for 111–112, 114
- controlling creation of 102
- example of project strategy 175
- in MultiSite 101
- mastership transfer models 140
- merge policies 105
- merging elements from UCM projects 95
- merging to main 149
- multiple levels, config specs for 114
- naming conventions 101
- stopping development on 191

building software, view configurations 120

C

***ccase-home-dir* directory** xix

change requests
- tracking in base ClearCase 136
- tracking states 21

change sets 10

ClearQuest integration
- about 15, 20
- customizing policies 52
- database, setting up 45
- decomposing activities 52
- disabling links to project 69
- enabling custom schema (procedure) 46
- enabling projects to use (procedure) 66
- environment variables 53
- planning issues 36
- policies available 43
- querying database 84
- recommended use of 197
- setting up 15
- setting up UCM schemas (procedure) 45

components
- about 13
- adding to integration stream (procedure) 73
- ancillary 28

candidates for read-only 29
conversion of VOBs (procedure) 63
creating new (procedure) 57
design considerations 24
importing files for (procedure) 62
mapping to projects 25
organizing for project 27
recommended directory structure 28
when to delete 86
config specs
about 102, 107
default, standard rules in 108
examples for builds 120
examples for development tasks 111
examples for one project 178
examples of time rules 112, 114, 119–120
examples to monitor project 116
include file facility 109
project environment for samples 110
restricting changes to one directory 115
selecting library versions 122
sharing across platforms 125
use of element types in 159
conventions, typographical xix

D

deliver operations
between projects 94
element types and merging 35
finding posted work (procedure) 76
MultiSite and 16, 75
pending checkouts policy 42
rebase policy 42
remote deliver 75
remote, completing (procedure) 76
state transition policy 43
development policies
See policies in base ClearCase; policies in UCM
development streams 13
creating for testing (procedure) 71
rebasing (procedure) 80
when to delete 86
directories, merging 153
directory structure
creating new (procedure) 61
recommended, for UCM components 28
documentation
online help description xx

E

element types
customizing 155
how assigned 156
predefined and user-defined 160
element types in UCM 35
environment variables for ClearQuest 53
event records 104

F

foundation baselines 58

G

global types 34, 104

H

hyperlinks
about 103
requirements tracking mechanism 137

I

importing files and directories 62
include file facility 109
integration streams
about 13
adding components (procedure) 73
locking 71
locking (procedure) 77
locking considerations 33
merging multiple 95
merging to base ClearCase branch 95
rebasing between projects (procedure) 91
unlocking (procedure) 80
updating development view load rules 75
when to delete 86
integration views
creating for UCM project (procedure) 59
recommended view type 42

J

Join Project Wizard 41

L

labels
about 103
baselines in base ClearCase 131
use in config specs 119, 121
load rules, updating for integration stream 75
locks
about 104
examples 134

M

main branch 100
makefiles and config specs 123
mastership
about 16
models of transfer 140
merging files
how it works 143
merging in base ClearCase
about 105
commands for 146
directory versions 153
entire source tree 150
extended example 183, 188
GUI tools for 145
how it works 143
non-ClearCase tools 154
removing merged changes 148
selective merge 147
to main branch 149
merging in UCM
See deliver operations; rebase operations
MultiSite
branches and 101
ClearQuest links in PVOBs 70
mastership transfer models 140
remote deliver 75
use in UCM 16

N

naming conventions
branches 101
ClearQuest schema 37
UCM baselines 32
views in base ClearCase 102

O

online help, accessing xx

P

parallel development
base ClearCase mechanisms 100
extended example in base ClearCase 173
UCM scenarios 89
parent/child controls in ClearQuest 52
patch release in UCM project 92
policies in base ClearCase
access to project files 134
bug-fixing on branches 132
change requests 136
coding standards 135
documenting changes 129
enforcement mechanisms 102, 129
labeling baselines 131
monitoring state of sources 130
notification of new work 134
on merging 105
requirements tracking 137
restricting changes visible 133
restricting use of commands 139
transfer of branch mastership 140
policies in UCM
about 14
approval before delivery 43
customizing ClearQuest 52
default view types 41
delivery transition state 43
delivery with pending checkouts 42
modifiable components 41
promotion levels 18
rebase before deliver 42
recommended baselines 41
setting ClearQuest (procedure) 68
verify activity owner before checkout 43
Project Explorer 58
projects in base ClearCase
branching strategy 100
config specs 102
development policies 102
extended example of lifecycle 173
generating reports 104
merging policies 105
planning and setup 100
views to monitor progress 116
projects in UCM
about 9
cleanup tasks 85
concurrent, managing 89

creating 12
creating from existing configuration 63
creating from existing projects 65
creating new (procedure) 58
disabling links to ClearQuest database 69
factors in gauging scope 25
fixing ClearQuest activity links 70
importing components 62
incorporating patch release 92
maintenance tasks 73
mapping components to 25
merging multiple 95
merging to base ClearCase branches 95
planning issues 23
setting up new 56
tools to monitor progress 82
promotion levels
about 18
changing (procedure) 82
default 32
defining in new project (procedure) 59
policy for recommended baselines 41
PVOBs
about 12
as administrative VOBs 34
ClearQuest links and MultiSite 70
creating from existing configuration 63
creating new (procedure) 56
mapping to ClearQuest database 37
number needed 33
requirements of ClearQuest database 67

Q

querying ClearQuest database 22, 84

R

Rational Unified Process 23–24
rebase operations
between projects (procedure) 91
element types and merging 35
policy for deliver operations 42
updating development view load rules 75
recommended baselines 41
record types for schemas, custom 49
remote deliver operations 75–76
reports
ClearQuest queries 84
for base ClearCase projects 104

S

schemas (ClearQuest)
about 21
enabling custom for UCM 40
enabling custom for UCM (procedure) 46
predefined, using 45
queries 22
requirements for UCM 38
storage issues 38
selective merge 147
smoke tests 32
state types
about 21
default transition requirements 50
setting for custom schemas 49
streams 13
subtractive merge 148
system architecture 23

T

technical support xx
time rules in config specs 112, 114, 119–120
triggers
about 103
checkin command example 129
example script for 134
sharing in mixed environments 141
to disallow checkins 135
to notify team of new work 134
to restrict use of commands 139
type managers
about 157
creating directory for 161
how they work 161
implementing compare method 166
inheriting methods 162
predefined 161
testing 167
user defined 161
typographical conventions xix

U

UCM and base ClearCase, compared 1
UCMPolicyScripts package 38
UnifiedChangeManagement package 38–39
user accounts
creating ClearQuest profiles (procedure) 53

V

version control, candidates for 24
views
config specs 107
configuring for builds 120
configuring for development tasks 111
configuring historical 119–120
configuring to monitor project 116
naming conventions in base ClearCase 102
policy for default types in UCM 41
restricting changes visible in 133
sharing for merges 149
VOBs
converting to UCM components (procedure) 63
creating and populating in base ClearCase 100

W

work areas 13