



The TimeSafe[®] Configuration Management System

AccuRev Technical Notes

Version 4.6

December, 2007

AccuRev Technical Notes

Copyright © AccuRev, Inc. 1995–2007
ALL RIGHTS RESERVED

TimeSafe and **AccuRev** are registered trademarks of AccuRev, Inc.

AccuBridge, **AccuReplica**, **AccuWork**, and **StreamBrowser** are trademarks of AccuRev, Inc.

All other trade names, trademarks, and service marks used in this document are the property of their respective owners.

Table of Contents

Quick Evaluation of AccuRev (Windows)	1
AccuRev's Client-Server Architecture	1
Downloading the Installation Program	1
Running the Installation Program	2
Starting the Server as a User Other than the Administrator.....	6
Starting the AccuRev GUI	6
Running the Quick Setup Wizard	7
Registering Yourself as an AccuRev User, 7	
Creating and Populating an AccuRev Depot, 8	
Beginning Development Work	11
Seeing Other People's Changes	11
Adding a New File	11
What's Next?	11
 Quick Evaluation of AccuRev	
(Unix/Linux / Command Line Interface)	13
Downloading the Installation Program	13
Running the Installation Program	13
Starting the Server.....	17
Registering Yourself as an AccuRev User.....	17
Creating and Populating an AccuRev Depot	17
Removing a Depot, 18	
Beginning Work	18
Keeping Files, 18	
Promoting Files, 18	
Seeing Other People's Changes	19
Adding a New File	19
What's Next?	19
 Converting to AccuRev from Directories Containing Baselevels	21
Creating a Depot	21
Processing the First Baselevel.....	21
Recording the Baselevel with a Snapshot, 22	
Processing Subsequent Baselevels	22
Handling Additional Baselevel-to-Baselevel Differences, 23	
Cleaning Up	24
 Creating and Using a Maintenance Stream	25
Creating a Snapshot	25
Creating a Stream Based on the Snapshot	25
Performing Maintenance Work.....	25

Pathname Optimization:

ACCUREV_IGNORE_ELEMS and .acignore	27
Eligible “Whole-Workspace” Commands	28
GUI Counterparts to the “Whole-Workspace” Commands, 28	
Commands that Don’t Apply to the Whole Workspace	28
Values for ACCUREV_IGNORE_ELEMS.....	29
Specifying Directories and Their Contents, 30	
Setting ACCUREV_IGNORE_ELEMS on a Unix/Linux System.....	30
Setting ACCUREV_IGNORE_ELEMS on a Windows System	30
Per-Directory Pathname Optimization — the .acignore File	31

AccuRev Server Performance33

Overview	33
Key Factors Affecting AccuRev Performance.....	34
Resource Utilization and the Client-Server Architecture	34
AccuRev Server Machine Usage: Software.....	35
AccuRev Server Machine Networking	35
Secure Networking, 35	

Using Multiple AccuRev Servers37

Setting Up Multiple Server Machines.....	37
Setting Up Client Machines	37
Configuring a Client Machine to Use Multiple Servers, 37	
Workspaces and Servers, 38	
The ‘wspaces’ File, 39	
Format of the ‘wspaces’ File, 39	
Location of the ‘wspaces’ File, 40	
The –H Command-Line Option, 40	

Techniques for Sharing Workspaces41

Accessing a Windows Workspace From Multiple Windows Clients	41
Universal Access to a Workspace Located on a Share	42
The ‘share_map.txt’ File, 42	
Workspace Location Entries, 43	
Fixing Workspace Location Entries, 43	
Fixing Machine Name Entries, 44	
Example: Samba Share, 44	

What’s the Difference between Populate and Update?45

In a Nutshell	45
Example 1: Standard Update Scenario, 46	
Example 2: Restoring a Deleted File (“missing” by accident), 46	
Example 3: Adding Data to a Sparse Workspace (“missing” by design), 46	
Example 4: Handling Active Elements, 47	
Example 5: A Tale of Two Files, 47	
Data Structures Used by Populate and Update	47
How the Data Structures Get Their Data	49
Backing Stream, 49	
Workspace Stream, 50	

Workspace Tree, 52	
The Update Algorithm	53
Advancing the Scan Threshold, 55	
Incomplete Updates, 55	
Incomplete Update: Command Interrupted, 56	
Incomplete Update: Checksum Failure, 56	
Performing the “Fixup” Update, 56	
Using a Trigger to Maintain a Reference Tree	59
Notes for CVS Users	61
AccuRev Workspaces vs. CVS Sandboxes.....	61
Common Operations	61
Obtaining a copy of the source files, 61	
Placing files under version control, 61	
Bringing others’ changes into your workspace/sandbox, 62	
Saving your changes, 62	
Finding the history of files, 62	
Finding the status of files in your workspace/sandbox, 63	
Removing files, 63	
Reverting changes to files, 63	
Moving files, 64	
Checking out files to edit, 64	
Comparing versions of files, 64	
Version Control of Namespace-Related Changes	65
Twin Elements and Stranded Elements	65
Preventing Creation of Twins in Workspaces, 65	
Reporting of Twins in Dynamic Streams, 66	
Ability to Reuse an Element Name after a Rename Operation, 66	
When a Purge Operation Causes an Element to Disappear, 67	
Detection of All Stranded Elements, Including “Twins”, 67	
Ability to Operate on Stranded Elements Using Element-IDs, 68	
More Sophisticated Analysis of Namespace-Related Changes, 68	
Change to Merge Algorithm for Namespace-Related Changes, 68	
Handling Stranded Elements.....	70
Defunct element obscured by element with same name, 70	
Resolving the Situation, 71	
Elements under a defunct parent, 71	
Resolving the Situation, 71	
Elements under an excluded parent, 71	
Resolving the Situation, 72	
Dangling directory elements, 72	
Resolving the Situation, 72	
Elements under a non-existent (purged) parent directory, 72	
Resolving the Situation, 73	
Elements under a stranded parent directory, 73	
Notes on Cross-Links	75
Cross-Link Direction and Terminology, 75	
Cross-Links and Stream Namespaces, 75	
Source Stream: Workspace vs. Dynamic Stream, 77	

Multiple Cross-Links: Chaining, 78	
Double Vision: Seeing an Element Multiple Times in a Workspace, 79	
Double Vision and the ‘accurev name’ Command, 80	
Cross-Link Overlaps, 80	

Notes on Promote-by-Issue	83
Promote-by-Issue Basics.....	83
Promoting Issues to the Parent Stream	85
Cross-Promoting Issues to a Non-Parent Stream — Simple Case.....	86
Cross-Promoting Issues to a Non-Parent Stream — Patch Required	88
Changing the Way You Use the Change Palette, 89	
Working in the Workspace, 90	
Promotion / Creating a Tracking Issue, 91	
Working with the Tracking Issue, 92	

Quick Evaluation of AccuRev (Windows)

If you intend to use AccuRev for personal use or are just setting it up for evaluation, these steps will take you through a basic, single platform installation, initial depot setup, and basic AccuRev usage. Any installation parameters that you enter can be changed at a later time.

You do not need to log in as an administrator to evaluate the product. But a non-administrator won't be able to install AccuRev as a service. These instructions include a section on how to start the AccuRev server in console mode to get around this problem.

Should you need any help during the installation or afterwards, remember that your AccuRev download comes with 30 days of free support. Please feel free to contact us: email support@accurev.com or call 1-800-383-8170.

AccuRev's Client-Server Architecture

AccuRev operations use a simple client-server model:

- A user invokes a client program: either the command-line interface (CLI) program **accurev**, or the graphical user interface (GUI) program **acgui**. The computer on which the client program executes is termed the client machine.
- The client program bundles each user command into a transaction request, and sends the request to the server program, which is running on the server machine. The request is sent via TCP/IP, using a particular host/port-number combination. The server program executes the request by accessing the AccuRev repository on the server machine. If the request modifies the repository, the server program records the change as a transaction in one of the repository's storage depots.

The client machine and server machine can be the same computer. (That's the most likely setup for a pre-sales evaluation of AccuRev.) By changing the TCP/IP host/port-number combination, a client program can "switch its attention" back and forth among multiple server programs — and thus, multiple repositories.

Downloading the Installation Program

In a Web browser, go to our web site at <http://www.accurev.com>. Obtain an AccuRev license key file, by filling out the registration form on the Registration page. You'll need to specify the pathname of this file during the installation process. On the Downloads page, check the "Platform Support Notes" section to verify that the Java language will run correctly on your computer. You may need to install some operating system patches to establish this level of support. Both the installation program and AccuRev itself have a Java user interface.

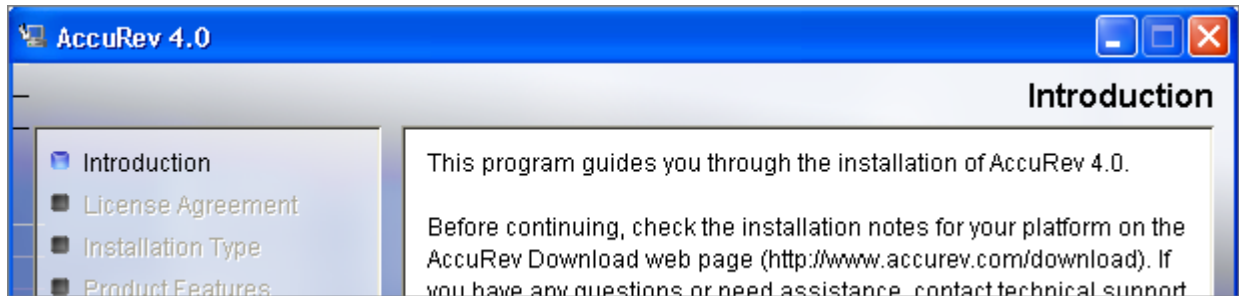
Then, download the installation program, **AccuRevInstall.exe**.

Running the Installation Program

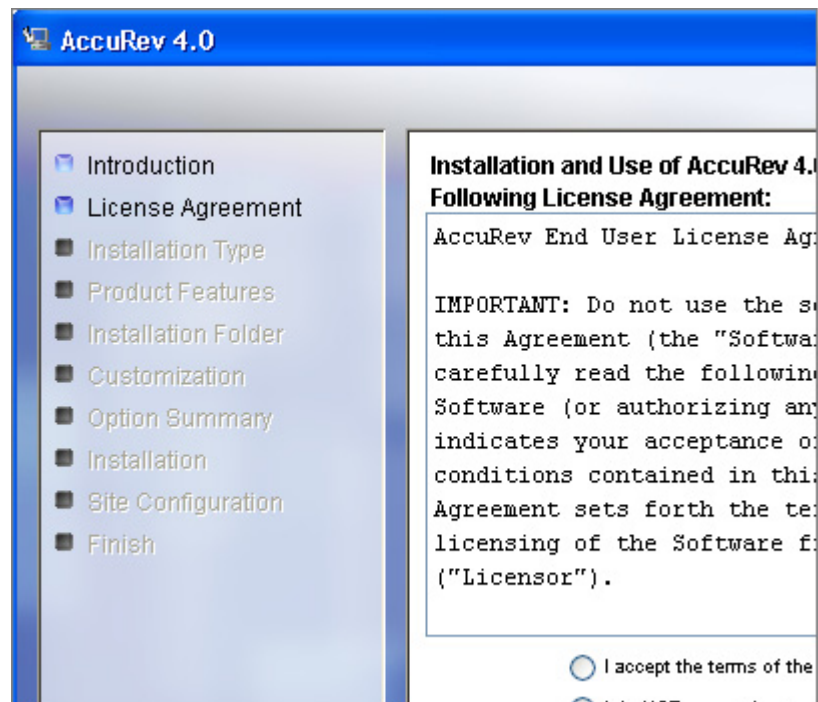
Note: after installation is complete, be sure to continue with the directions in this chapter. We'll guide you through initial depot setup and basic AccuRev usage.

Invoke the **AccuRevInstall** program. It's a "wizard": multiple screens that you navigate by clicking **Next** and **Previous** buttons. The wizard provides lots of helpful information, so you may not even need to consult the following screen-by-screen notes.

- **Introduction:** The screen explains how to use the wizard. Click **Next**.

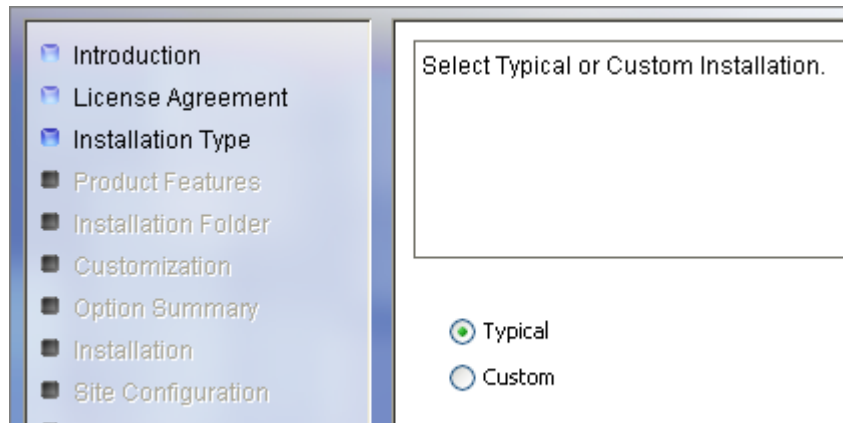


- **License Agreement:** Read the agreement, select "I accept", and click **Next**.

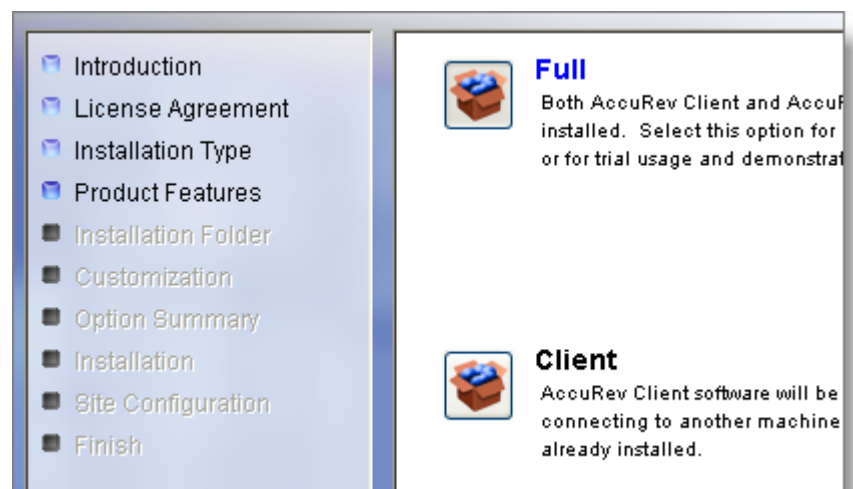


- **Choose Install Type:** If you select **Typical**, the installation program will make a few choices for you, skipping over a few of the subsequent screens.

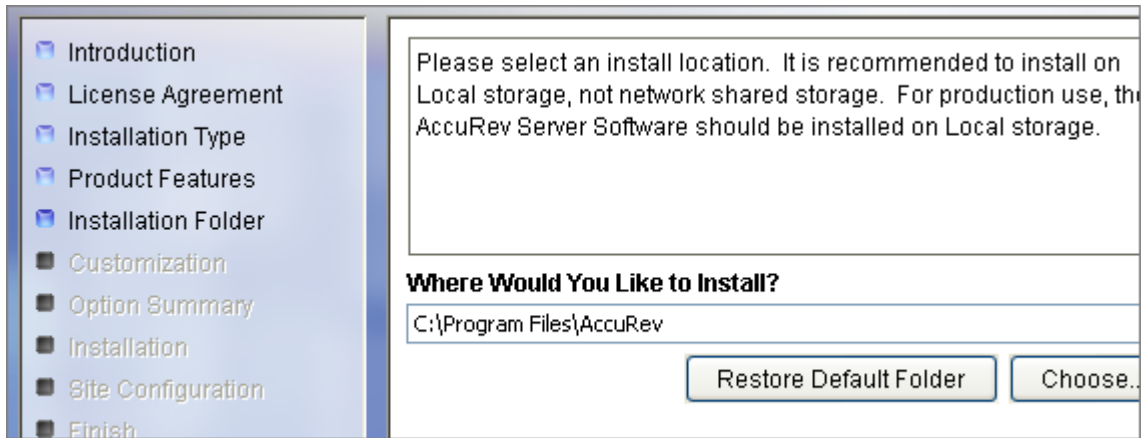
If you are not sure about the **Typical** vs. **Custom** choice, read ahead to the discussion of the various customizations before proceeding.



- **Choose Product Features:** For a product evaluation to be conducted on a single machine, choose **Full**, then click **Next**. If your evaluation involves running the AccuRev Server on a separate machine (e.g. to test network performance), choose **Full** when installing on the server machine; choose **Client** when installing on each client machine.



- **Choose Install Folder:** The wizard suggests a default location for installing AccuRev. This location will store the product itself (about 27MB of executables and configuration files), along with the repository (one or more “depots”) for your version-controlled files.



The installation program uses a heuristic (an intelligent guess) to determine whether you have Windows administrative privileges. If so, it suggests a public location — for example, within **C:\Program Files**. If not, it suggest a location within your home directory.

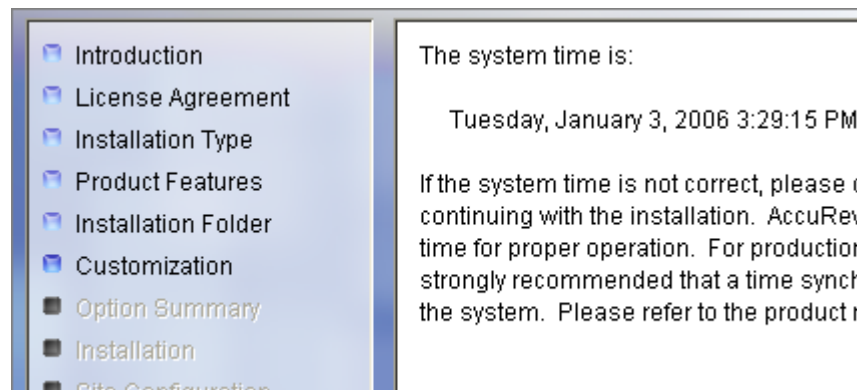
To choose an alternate location, type a pathname in the textbox or click **Choose** and navigate to an existing location. Before clicking **Next**, make sure the pathname in the textbox specifies a directory that does not yet exist. If you go the **Choose** route, select a location and then append “\AccuRev” (or maybe “\AccuRev3”) to the pathname before proceeding to the next screen.

- **Verify System Time:**

This appears for a **Full** installation only. It gives you a chance to change the Windows system clock before installing AccuRev. See *System Clock Synchronization* on page 21 of the *AccuRev*

Administrator’s Guide

for a discussion of the importance of an accurate system clock.



- **Optional Customizations:** These customization steps appear if you select a **Custom** install instead of a **Typical** install.

- **Configuration files:** A set of steps enables you to specify the contents of the client configuration file and, for a Full installation, the server configuration file. If you’re installing on top of an existing AccuRev installation, the wizard first asks if you wish to change the existing settings. The configuration file settings are:

- **Folder for Server Data Storage: (Full installation only)** Determines the location of the AccuRev repository, which stores your version-controlled files along with the databases that keep track of the files.

The default is to place the repository in a folder named **storage** under the main “install folder”.

You may want to place the repository on a separate disk — perhaps one that is faster and/or backed up more frequently than the disk where you’ve chosen to install the AccuRev executables. The repository’s disk must be located on the machine where the AccuRev server process executes.

- **License Key File:** (Full installation only) Names the license key file that you obtained after filling in the AccuRev registration form. If you’re installing on top of an existing AccuRev installation, the wizard first asks if you wish to replace your existing license key file.
- **Host and Port:** For a product evaluation to be conducted on a single machine, just click **Next** to accept the defaults: the AccuRev Server will run on your machine, using TCP/IP port 5050.

For a **Client** installation, enter the hostname or IP address of the server machine in the **Host** field. You may need to enter a fully-qualified hostname, such as **helios.newmoon.com**. You may also need to specify a different port number (see below).

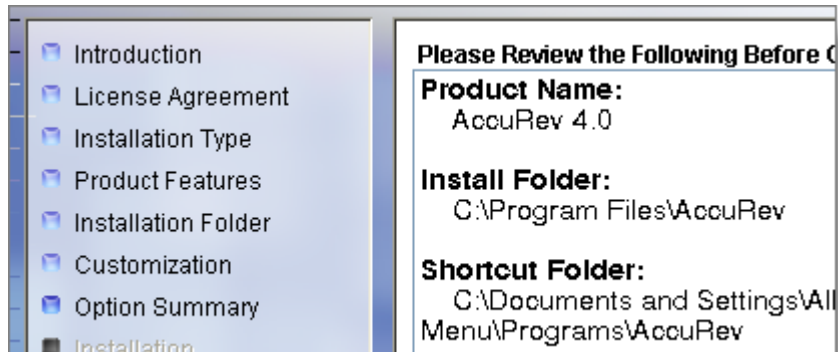
For a **Full** installation, make sure that no network applications running on this machine use port 5050. If this port is already being used, change the port number to one that is not used by any of the other applications. Then, make sure that all other users perform a Custom Client installation of AccuRev, specifying the same alternative port number.

- **System PATH Variable:** Adjusting the search path facilitates usage of the AccuRev command-line interface: you’ll be able to invoke the CLI using the simple program name **accurev**, instead of having to type its full pathname. Adjusting the search path is not necessary if you will be using AccuRev’s graphical user interface only, not the command-line interface.
- **Source Control Integration:** Support for IDEs that use the Microsoft SCC interface is always installed. This setting controls whether AccuRev becomes the default provider of SCC services to such IDEs.
- **Java Virtual Machine:** The AccuRev graphical user interface is programmed in the Java™ language. This means that your computer must have a Java virtual machine (Java VM, also known as a “Java runtime environment”, or JRE) installed. The wizard will install a Java VM in the AccuRev installation area if you select “Install a Java VM”. We recommend this; it won’t affect any other Java installation or Java-based applications on your machine.

To enable AccuRev to use a Java virtual machine that is already installed on your machine, the wizard searches for an executable named **java** and displays its pathname

if it finds one. You can click **Search for Others** to have the wizard search for additional Java VMs on your computer. Or you can click **Choose Another**, then navigate to another Java VM.

- **Pre-Installation Summary:** Click **Install**. The wizard copies files to the installation area on your machine.




The installation process is now complete. An **AccuRev** shortcut is placed on the Windows desktop as part of the installation process. This shortcut starts the AccuRev graphical client user interface.

Starting the Server as a User Other than the Administrator

If you are unable to install the AccuRev service because you lack administrator privileges, just double-click the **AccuRev Demo Server** shortcut on the desktop. Alternatively, run the **server_start.bat** script in the AccuRev executables (**bin**) directory.

Starting the AccuRev GUI


Double-click the **AccuRev** shortcut on your desktop. The main screen appears:



Running the Quick Setup Wizard

The first time you start the GUI, the AccuRev Quick Setup wizard should appear. If it doesn't, run the command **Help > Quick Setup**. This wizard steps you through the process of:

- Registering yourself as an AccuRev user.
- Creating a new storage depot for permanent storage of a set of files.
- Placing an existing set of files under version control, by storing them in the new depot.



Before you begin, make sure that the set of files to be placed under version control are on a disk that is accessible from the machine where you've installed AccuRev.

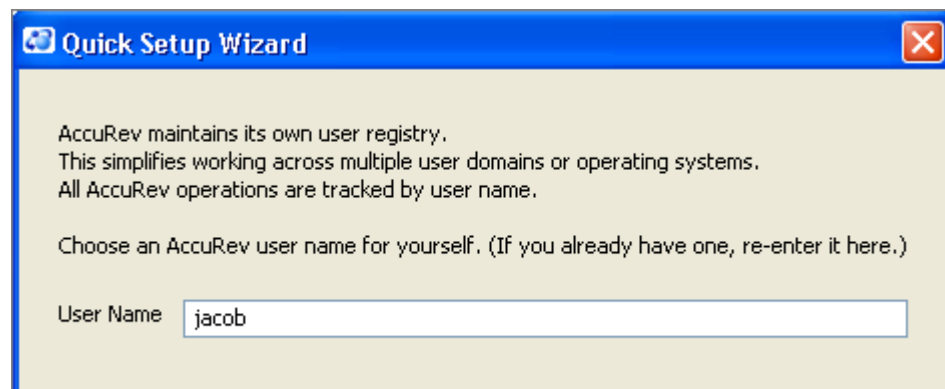
The following sections describe the wizard's operation in detail.

Registering Yourself as an AccuRev User

AccuRev maintains a username registry that is separate from the operating system's username registry. AccuRev usernames are called principal-names. Every AccuRev transaction log entry includes the principal-name of the person who initiated the transaction.

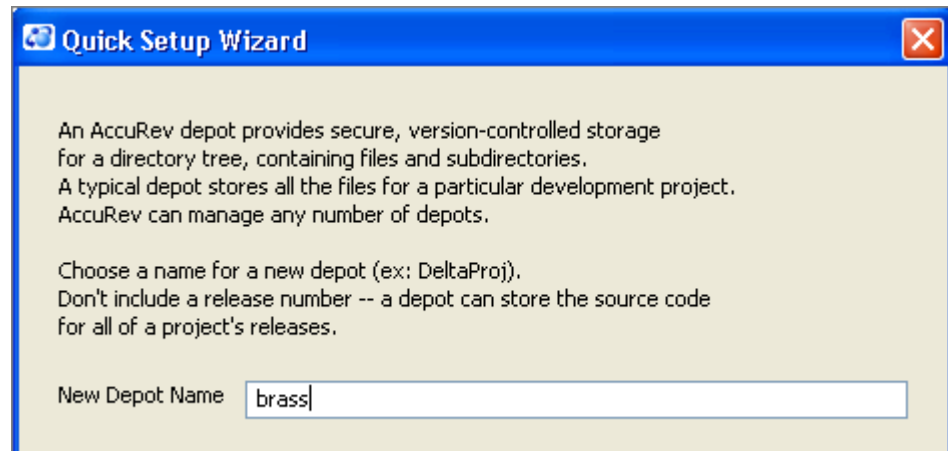
The wizard prompts you to create a principal-name for yourself. If you already have a principal-name, just enter it again in this dialog box.

The currently active principal-name is displayed in the status bar at the bottom of the GUI window. You can change it at any time with the **Tools > Change Active User** command.

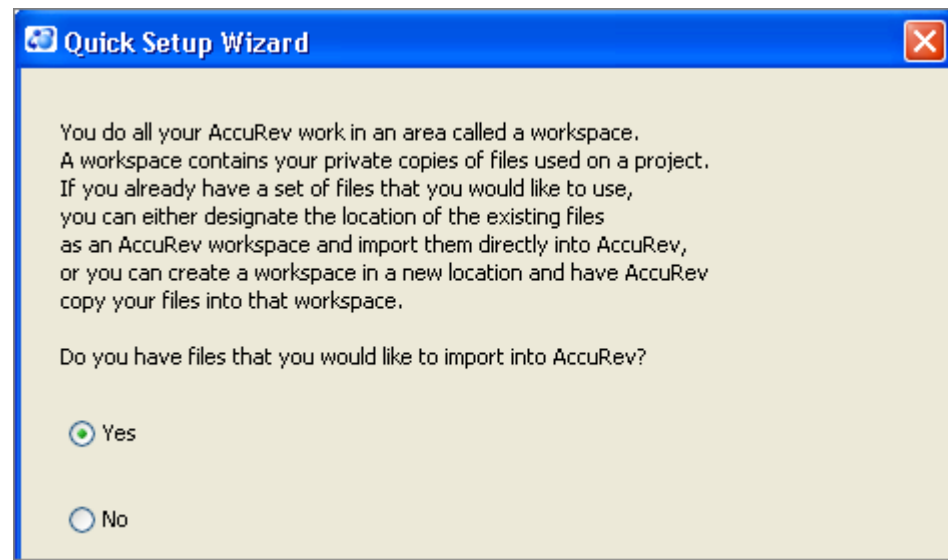


Creating and Populating an AccuRev Depot

Next, the wizard prompts you to specify the name for a new storage depot. Enter a name for the depot (e.g. **brass** for a project code-named “Brass”).



Now, you can indicate, by selecting **Yes**, that you have a directory tree whose files are to be placed in the new depot. The alternative (if you select **No**) is to create an empty depot.

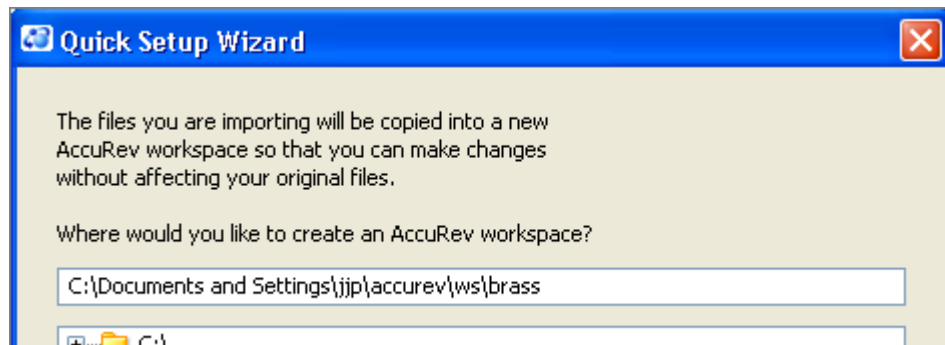
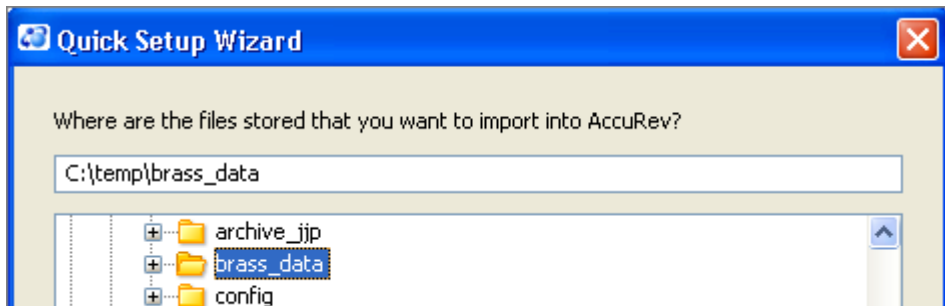


The next steps involve the creation of a new workspace, which enables you to work with the depot's files.

- **If you selected “Yes”, indicating that you have files to be placed in the new depot:**

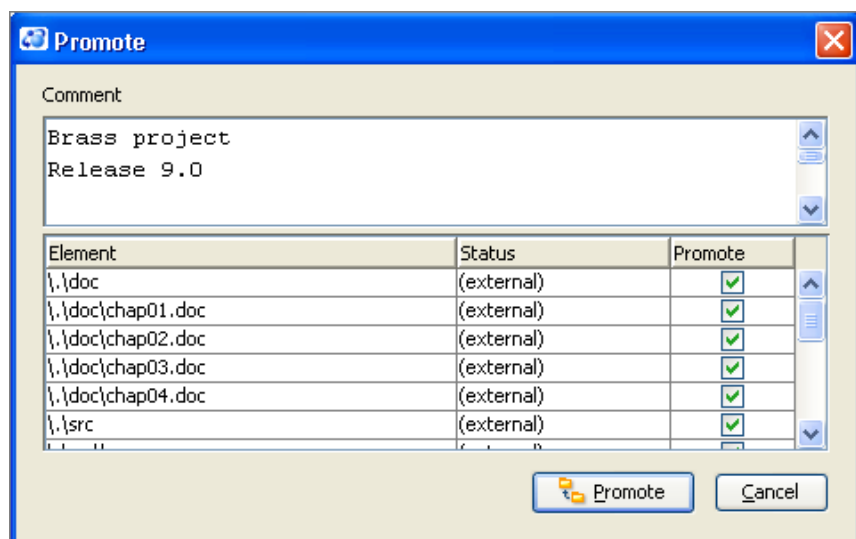
The wizard asks “Do you want to create a workspace in a new location ...”:

- Select **Yes** to create a workspace in new location, and copy your existing files to that new location. The wizard will (1) prompt you to specify the location of the existing files, and (2) prompt you to specify a location for the new workspace.



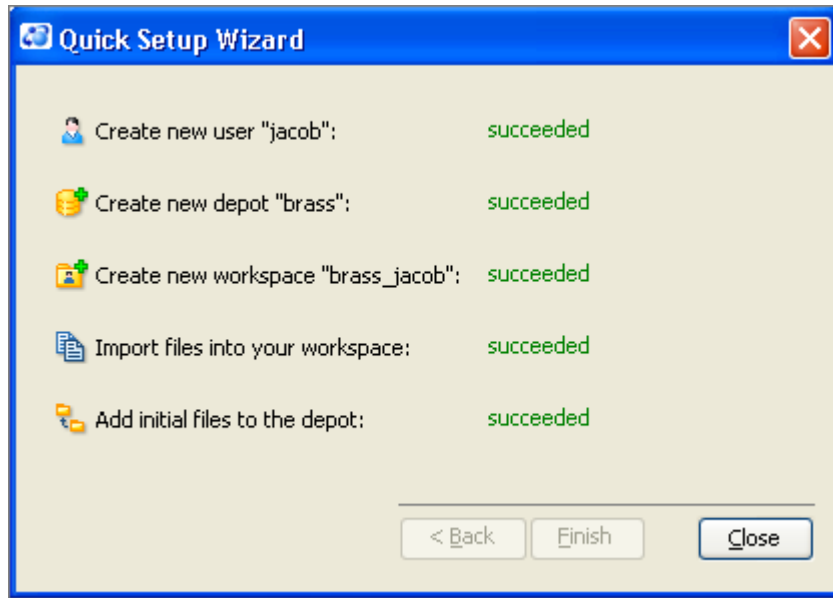
- Select **No** to convert an existing location into a workspace. For this alternative, the wizard will prompt you to specify the existing location.

In either case, your existing files are copied to depot in a series of operations that ends with a **Promote** command. click the **Ok** button to approve the promotion of the files. This makes the files available to other users' workspaces that you can create later:



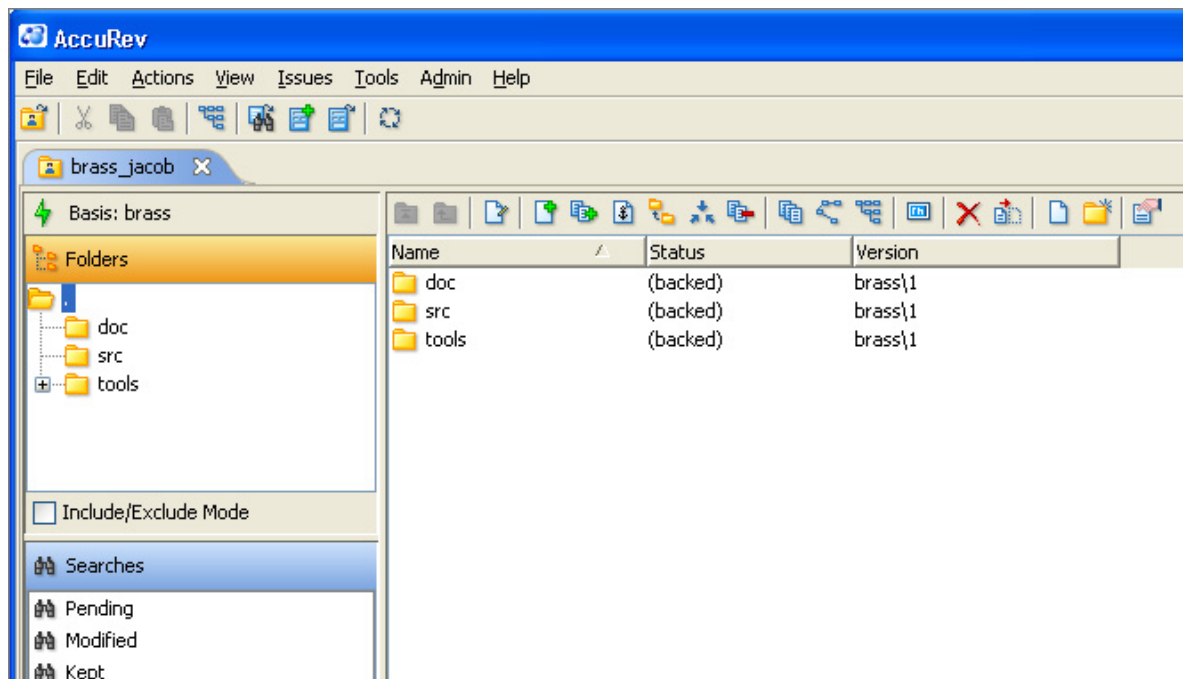
- If you selected “No”, indicating that you want to create an empty depot:
The wizard will prompt you to specify a location for the new workspace.

Finally, the wizard displays a message box, summarizing the work it has accomplished.



Note: To remove a depot from the repository, use the administrative command **maintain rmdepot**. This operation is irreversible! For details, see [Removing a Depot from the AccuRev Repository](#) on page 85 of the *AccuRev Administrator's Guide*.

After the Quick Setup Wizard ends, AccuRev displays the contents of the newly created workspace. This "Explorer-style" display includes folders (navigation) and details panes. It also includes a searches pane, with which you can locate files throughout the workspace, based on their AccuRev status.




Beginning Development Work

Now that you've created a depot and populated it with files, you're ready to begin making changes to those files. There's no need to “check out” files; simply edit them. To edit a file, double-click it. Alternatively, right-click it and select **Edit** from the context menu.

If you wish to keep some work that you have done, select the elements you want to keep, right-click on the selection and choose **Keep**. You are prompted for a comment, which is optional. This creates a new version in the depot, in your workspace's private stream. You can keep a file as often as you like.

When your changes are ready for other people to see, first keep the file, then promote it. To promote elements, select them, right-click on them, and select **Promote**. Promote places your most recently kept version of the file in the backing stream, making it available for others to use. The next time they update their workspace, they will get a local copy of your promoted file.

Seeing Other People's Changes

To update your workspace with other people's changes, click the  **Update** button above the folder pane's toolbar.

Adding a New File

Simply edit a file with your favorite editor. When you're ready for AccuRev to add the initial version of your new file, right-click on it and select **Add to Depot**. Enter a comment (optional, just as with **Keep**).

You can promote the new file immediately. Alternatively, you can continue editing, and then keep and promote.

What's Next?

This chapter has covered only the most basic installation options and features of AccuRev. See the *AccuRev User's Guide* to learn more about AccuRev's capabilities. There are both CLI and GUI editions of this manual.

Quick Evaluation of AccuRev

(Unix/Linux / Command Line Interface)

If you intend to use AccuRev for personal use or are just setting it up for evaluation, these steps will take you through a basic, single platform installation, initial depot setup, and basic AccuRev usage. Any installation parameters that you enter can be changed at a later time.

You do not need to log in as **root** to install and use AccuRev.

Should you need any help during the installation or afterwards, remember that your AccuRev download comes with 30 days of free support. Please feel free to contact us: email **support@accurev.com** or call **1-800-383-8170**.

Downloading the Installation Program

In a Web browser, go to our web site at <http://www.accurev.com>. Obtain an AccuRev license key file, by filling out the registration form on the Registration page. You'll need to specify the pathname of this file during the installation process. On the Downloads page, check the "Platform Support Notes" section to verify that the Java language will run correctly on your computer. You may need to install some operating system patches to establish this level of support. Both the installation program and AccuRev itself have a Java user interface.

Then, download the installation program, **AccuRevInstall.bin**, for your hardware/software platform.

Running the Installation Program

Note: after installation is complete, be sure to continue with the directions in this chapter. We'll guide you through initial depot setup and basic AccuRev usage.

Start the installation process with one of the following commands:

```
sh ./AccuRevInstall.bin                (graphical wizard)
sh ./AccuRevInstall.bin -i console      (non-graphical wizard)
```

Following is a transcript of a non-graphical installation, with some annotations. For help with a graphical installation, see the corresponding chapter for Windows systems, *Quick Evaluation of AccuRev (Windows)* on page 1.

- **Introduction and License Agreement:** The wizard starts with an explanation of how to use it, along with a license agreement. Press **Enter** several times to scroll through and read this material, then type **Y** and press **Enter** to accept the license agreement

```
This program will guide you through the installation of AccuRev 4.0.
Before continuing please check the installation notes specific to your
platform, which are found on the AccuRev Download web page
```

```
. . .
```

DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N) : **Y**

- **Non-Administrator User Identity Warning:** If you are not logged in as the **root** user, the wizard displays an appropriate message.
- **Choose Install Type:** The wizard has you choose between a typical installation and a custom installation:

Select Typical or Custom Installation.

->1- Typical
2- Custom

Typical or Custom?: **1**

- **Choose Product Features:** For a product evaluation to be conducted on a single machine, choose **Full**. If your evaluation involves running the AccuRev Server on a separate machine (e.g. to test network performance), choose **Full** when installing on the server machine; choose **Client** when installing on each client machine.

Please choose the Install Set to be installed by this installer.

->1- Full
2- Client

ENTER THE NUMBER FOR THE INSTALL SET, OR PRESS <ENTER> TO ACCEPT THE DEFAULT
: **1**

- **Choose Install Directory:** The wizard suggests a default location for installing AccuRev. This location will store the product itself, along with the repositories (“depots”) that hold version-controlled files. If you are the **root** user, it suggests a public location; otherwise, it suggests a location within your home directory.

Make sure the pathname specifies a directory that does not yet exist.

Where would you like to install?

Default Install Folder: /opt/accurev

ENTER AN ABSOLUTE PATH, OR PRESS <ENTER> TO ACCEPT THE DEFAULT
: **/usr/local/accurev**

INSTALL FOLDER IS: /usr/local/accurev
IS THIS CORRECT? (Y/N) : **Y**

- **Verify System Time:** This message appears for a **Full** installation only. It gives you a chance to change the Windows system clock before installing AccuRev. See *System Clock*

Synchronization on page 21 of the *AccuRev Administrator's Guide* for a discussion of the importance of an accurate system clock.

The system time is:

Wednesday, January 4, 2006 10:37:48 AM EST

. . .

PRESS <ENTER> TO CONTINUE:

- **Select License Key File:** This message appears for a **Full** installation only. Names the license key file that you obtained after filling in the AccuRev registration form. If you're installing on top of an existing AccuRev installation, the wizard first asks if you wish to replace your existing license key file. If you don't yet have a license key, just press **Enter**. When you obtain the license key file, you'll need to copy it to the **storage/site_slice** directory within the AccuRev installation area.

Choose a license key file [/home/jjp/keys.txt]:

- **Choose a Directory for AccuRev Server Data Storage:** This appears for a **Full** installation only. Press **Enter** to create the directory within the installation area.

Choose a Directory for the AccuRev Server Software data repository. This location will store all data managed by AccuRev, so choose a location with enough space. The folder will be created if it does not already exist.

Choose a Directory [/usr/local/accurev/storage]:

- **Choose Hostname and Port Number:** For a **Client** installation, enter the hostname or IP address of the server machine in the **Host** field. You may need to enter a fully-qualified hostname, such as **helios.newmoon.com**. You may also need to specify a different port number (see below).

For a **Full** installation, make sure that no network applications running on this machine use port 5050. If this port is already being used, change the port number to one that is not used by any of the other applications. Then, make sure that all other users perform a Custom Client installation of AccuRev, specifying the same alternative port number.

Verify the host name and port number for the local AccuRev Server software. The default port does not need to be changed unless it conflicts with other services. The AccuRev Client will be configured to connect to the local server.

Select Host [wayback]:

Select Port [5050]:

- **Choose Link Location:** Control the creation of symbolic links to the executables.

Where would you like to create links? If you put links in your Path, you will be able to run AccuRev easily from anywhere on your system.

```
->1- Default: /usr/bin
  2- In your home folder
  3- Choose another location...
```

```
4- Don't create links
```

```
ENTER THE NUMBER OF AN OPTION ABOVE, OR PRESS <ENTER> TO ACCEPT THE DEFAULT
: 3
```

```
ENTER THE ABSOLUTE PATH TO THE SHORTCUT DIRECTORY
: /home/jjp/bin
```

```
SHORTCUT DIRECTORY IS: /home/jjp/bin
IS THIS CORRECT? (Y/N): y
```

- **Adjust System PATH Variable:** Accept the offer to add the AccuRev executables directory to your search path.

```
Add to Path? (Y/N): y
```

- **Choose Java Virtual Machine:** The AccuRev graphical user interface is programmed in the Java™ language. This means that your computer must have a Java virtual machine (Java VM, also known as a “Java runtime environment”, or JRE) installed. The wizard will install a Java VM in the AccuRev installation area if you select “1. Install a Java VM”.

To enable AccuRev to use a Java virtual machine that is already installed on your computer, the wizard searches for an executable named **java** and displays its pathname as choice #2 if it finds one. You can also specify a Java VM that has some other name (choice #3).

Please choose a Java VM for use by AccuRev. Installing a VM is recommended. AccuRev requires at least a v1.3 VM.

```
->1- Install a Java VM specifically for this application

  2- /workspaces/u2/java/jdk1.3.1_01/bin/java

  3- Choose a Java VM already installed on this system
```

```
ENTER THE NUMBER FOR THE JAVA VM, OR PRESS <ENTER> TO ACCEPT THE
CURRENT SELECTION: 2
```

- **Pre-Installation Summary:** After checking the summary, press **Enter**. The wizard copies files to the installation area on your machine.

Please Review the Following Before Continuing:

```
Product Name:
  AccuRev_3.x
```

. . .

PRESS <ENTER> TO CONTINUE:

- **Configure Installed Software — Set Server Host and Port:** For a product evaluation to be conducted on a single machine, click **Next**. If the client and server software will run on different machines, enter the hostname or IP address of the server machine in the **Host** field. You may need to enter a fully-qualified hostname, such as **helios.newmoon.com**.

To be truly careful, make sure that no network applications in your environment use port 5050. You can change the port number to any number that is not used by other applications. If you do, make sure to use the same port number for each AccuRev client and server installation.

- **Configure Installed Software — Start AccuRev Server:** On Windows 2000 machines, the AccuRev server can be configured as a service, so that the server starts automatically along with the operating system. Select **Yes** to make this happen.

If you select **No**, or if your machine has a version of Windows that does not support services, you must start the server manually.

Starting the Server

Start the AccuRev Server in the background:

```
accurev_server &
```

This sends a short message to **stdout**. All subsequent messages go to the server's log file.

Registering Yourself as an AccuRev User

Use **mkuser** to create a new AccuRev user. In most cases, your AccuRev username (also known as a principal-name) should match your Unix/Linux username.

```
accurev mkuser john_smith
```

Creating and Populating an AccuRev Depot

Follow these steps to place an existing source tree under AccuRev's control:

1. Create a depot with **mkdepot**.

```
accurev mkdepot -p gizmo
```

AccuRev will store the depot in a default location. (It can easily be moved at a later time.)

2. A workspace is a personal area used for editing and testing your work. Make a workspace out of your source tree with the **mkws** command. For example, if your sources are in **/home/jsmith/ws/my_gizmo**:

```
accurev mkws -w my_gizmo -b gizmo -l /home/jsmith/ws/my_gizmo
```

The argument to the **-b** option should be the same as the depot name.

3. **cd** to the top directory of your source tree containing the files you wish to put into the newly created depot:

```
cd /home/jsmith/ws/my_gizmo
```

4. To see all of the files in the workspace which need to be added, use the **stat -x** command.

```
accurev stat -x
```

5. Place all of the files in the workspace under AccuRev version control using the **add -x** command.

```
accurev add -x
```

(You can filter out files with certain extensions (e.g. **.bak**) if you don't want to place *all* the files under version control. See the **add** reference page for details.)

6. Make the newly added files available to other workspaces by promoting them:

```
accurev promote -k
```

Removing a Depot

A depot can be removed completely from the repository with the **maintain rmdepot** command. This operation is irreversible! For details, see *Removing a Depot from the AccuRev Repository* on page 85 of the *AccuRev Administrator's Guide*.

Beginning Work

Now that you've created a depot and have populated it with the files in your workspace, you're ready to begin making changes to the files you've added. There's no need to check out files; simply edit them.

Keeping Files

If you wish to keep some work that you have done on a file, use:

```
accurev keep -c "comment required" <filename>
```

This creates a new version in the repository in your private workspace stream. You can **keep** a file as often as you like.

Promoting Files

When your changes are ready for other people to see, first keep the file, then promote it:

```
accurev promote <filename>
```

The **promote** command places the most recently kept version of the file in the backing stream for others to use. The next time they update their workspace they will get a local copy of your promoted file.

Seeing Other People's Changes

To update your workspace with other people's changes, use:

```
accurev update
```

Adding a New File

Create a file in the workspace with an editor or by copying it from another location. When you're ready for AccuRev to add the initial version of your new file:

```
accurev add <filename>
```

Add is just like a keep, but it does not require a comment. The file must exist before you add it. You can **promote** now, or continue editing and then **keep** and **promote**.

What's Next?

This chapter has covered only the most basic installation options and features of AccuRev. See the *AccuRev User's Guide* to learn more about AccuRev's capabilities. There are both CLI and GUI editions of this manual.

Converting to AccuRev from Directories Containing Baselevels

Many development groups use source code provided by another group within the organization, or from another organization altogether. Typically, the code is imported on a periodic basis as a complete source tree, which we'll call a "baselevel". This note examines a scenario in which source-code baselevels are imported into AccuRev. Suppose each baselevel is stored as a directory tree:

```
D:\baselevels\gizmo1.0
D:\baselevels\gizmo2.0
D:\baselevels\gizmo2.5
```

It is easy to incorporate the multiple baselevels into AccuRev. Make sure you read these instructions all the way through before trying it out.

Creating a Depot

First, create an AccuRev depot, where AccuRev permanently stores all of the data for a programming project. For example:

```
accurev mkdepot -p gizmo
```

This creates a depot called **gizmo**. It has a single stream, also called **gizmo**.

Processing the First Baselevel

Next, populate **gizmo** with files from the first baselevel.

1. Go to the first baselevel directory:

```
cd \baselevels\gizmo1.0
```

2. Create a workspace to be used for importing files from the baselevel into AccuRev:

```
accurev mkws -w import -b gizmo -l .
```

Note that the command line ends with "dash-ell dot". This creates a workspace called **import**, which is based on stream **gizmo** (currently empty) in the current location.

3. Get a list of all of the files that AccuRev doesn't know about (which is all of them):

```
accurev stat -x > extfiles.list
```

The **-x** stands for external. You may see lots of files that you don't want to put under version control: object files, executables, text-editor backup files, etc.

4. You can have AccuRev ignore such files, by specifying patterns (wildcards) that match their names as the value of an environment variable. (Note that case is important in this pattern-matching). For example:

```
set ACCUREV_IGNORE_ELEMS=*.exe *.obj *.lnk *.err *.map (and so on)
```

(The syntax for setting environment variables varies among operating systems and command-line processors. For more on `ACCUREV_IGNORE_ELEMS`, see *Pathname Optimization: `ACCUREV_IGNORE_ELEMS` and `.acignore`* on page 27.)

5. Try the preceding two steps again, keep repeating this process until **stat -x** lists exactly the set of files that you want AccuRev to keep track of.
6. Create initial versions of these files in the depot:

```
accurev add -x
```

This creates versions in the **import** workspace, but has not yet made them available to others working on the **gizmo** project.

7. To make these new files public, promote them to the **gizmo** stream:

```
accurev promote -k
```

Recording the Baselevel with a Snapshot

At this point, you can create a snapshot of the **gizmo** stream. A snapshot is a special kind of stream, whose contents can never change. (Hence, a snapshot is also called a “static stream”, distinguishing it from a standard dynamic stream.) In this case, the snapshot will contain the versions in the first baselevel, because that’s exactly what the **gizmo** stream contains at the current time.

(The **gizmo** stream itself will change, as you incorporate additional baselevels. But any snapshot you create is guaranteed to be frozen forever!)

Use the **mksnap** command to create the snapshot:

```
accurev mksnap -s gizmo1.0 -b gizmo -t now
```

At any time in the future, you can use snapshot **gizmo1.0** to see the contents of the first baselevel. And if you need to fix a bug that existed at this baselevel, you can create a maintenance stream below the snapshot. See *Creating and Using a Maintenance Stream* on page 25.

Note: AccuRev does not implement snapshots with “version labels”, as do branch-and-label SCM systems. Since there’s no need to attach a label to each version in the baselevel, creating a snapshot is virtually instantaneous!

Processing Subsequent Baselevels

Now, you need to “layer” the files in the next baselevel on top of the files that you’ve already placed under AccuRev control.

1. Change the definition of the **import** workspace:

```
cd D:\baselevels\gizmo2.0
accurev chws -w import -l .      (again, “dash-ell dot”)
```

In effect, you’ve moved the workspace to where the files are, instead of moving the files into the workspace! The files fall into several categories.

2. Make sure that all files in this baselevel have timestamps that are later than the timestamps in the preceding baselevel:

```
accurev touch -R .
```

3. Process the files that changed from **gizmo1.0** to **gizmo2.0**.

To AccuRev, these files appear to be modified versions of the **gizmo1.0** files that you **add**'ed and **promote**'d in the preceding section. You can list all these “modified” files:

```
accurev stat -m
```

And you can keep the new versions of the files:

```
accurev keep -m -c "my comment"
```

4. Process the files that didn't change from **gizmo1.0** to **gizmo2.0**.

You don't need to do anything about these files. In particular, you don't need to **keep** new versions.

5. Process the files that didn't exist in **gizmo1.0**, but do exist in **gizmo2.0**.

These files are external, because AccuRev hasn't seen them before. (Just as *all* the files were external when you placed the first baselevel under version control.) Add the external files to the depot, just as you did in the preceding section:

```
accurev add -x
```

As above, you may want to use **stat -x** and the `ACCUREV_IGNORE_ELEMS` environment variable to filter out unwanted files before entering the **add** command.

6. Promote the new files and changed files:

```
accurev promote -k
```

You've now placed two baselevels under AccuRev control. Layering the third baselevel, **gizmo2.5**, on top of the second one is exactly the same as layering the second one on top of the first. Just repeat the steps in this section.

Handling Additional Baselevel-to-Baselevel Differences

In the discussion above, we broke a baselevel's “new layer” of files into three categories. This was a bit oversimplified — there are additional categories to consider.

- Files that existed in one baselevel, but were deleted in the next baselevel.

You can make such files disappear from the new baselevel by defuncting them:

```
accurev defunct <filenames>
```

- Files that were renamed from one baselevel to the next.

This will appear to be (1) a file that existed in one baselevel, but was deleted from the next baselevel, along with (2) a new file that didn't exist in the preceding baselevel. If you know

that file **oldname.c** in the preceding baselevel was renamed to **newname.c** in the next baselevel, use this series of commands to make the connection:

```
ren newname.c SAVEME (Unix/Linux: use the mv command)
accurev move oldname.c newname.c
ren SAVEME newname.c
```

Now, AccuRev knows that the element formerly known as **oldname.c** is henceforth to be known as **newname.c** (until the next name change, that is!).

Cleaning Up

Finally, deactivate the **import** workspace:

```
accurev rmws import
```

Creating and Using a Maintenance Stream

Many software development organizations have two main streams of development: work towards the next release, and maintenance of the previous release. Other SCM systems use a “branch based on a label” paradigm to accomplish this. AccuRev uses snapshots (static streams).

Creating a Snapshot

At the time of the release (say, “WidgetSoft Release 1.0”), create a snapshot:

```
accurev mksnap -s widget1.0 -b widget -t now
```

This creates a new snapshot called **widget1.0**. The snapshot contains whatever versions the **widget** stream contained at the time the **mksnap** command is executed. Subsequently, the **widget** stream can change as new versions are promoted to it, but the **widget1.0** snapshot never changes. Instead of **now**, you can specify any time in the past, such as **2005/05/18 10:10:24**.

Creating a Stream Based on the Snapshot

For maintenance work on this release, create a new dynamic stream based on the snapshot:

```
accurev mkstream -s widget_maint -b widget1.0
```

Initially, **widget_maint** will be identical to **widget1.0**, but it will change as people promote changes to it.

Performing Maintenance Work

Maintenance developers use workspaces based on the **widget_maint** stream. For instance, to make a maintenance fix, Mary might create a workspace like this:

```
accurev mkws -w widget_maint_mary -b widget_maint -l <wherever>
```

When Mary promotes her maintenance work, the changes will go to **widget_maint**.

All maintenance work is isolated from the main development stream, and vice-versa. Developers working on the next release create their workspaces off the development stream, not the maintenance stream. For example:

```
accurev mkws -w widget_justine -b widget -l <wherever>
```

Changes promoted from **widget_justine** will go to the main development stream, **widget**. The changes won’t appear in the **widget_maint** stream.

The Change Palette in the AccuRev GUI makes it easy to migrate changes back and forth between a main development stream (**widget**) and a maintenance stream (**widget_maint**).

Pathname Optimization: ACCUREV_IGNORE_ELEMS and .acignore

Like most command-line programs, AccuRev's CLI tool, **accurev**, accepts one or more filename/pathname specifications as command arguments:

```
accurev keep base.h  
accurev promote intro.doc chap*.doc
```

Such arguments cause the command to be invoked on a particular set of files.

But several **accurev** commands are capable of operating on *all* the files in the current workspace. For example, this command searches the entire workspace containing the current working directory, and lists the files that have not been placed under version control (external files):

```
accurev stat -x
```

Such “whole-workspace” commands are very powerful and useful, but they can be time-consuming. If your workspace contains many hundreds or thousands of files, you must wait while all the names are transmitted to the server machine, the AccuRev Server process determines the status of each file, and information on the matching files is returned to the client machine.

This large list of files to be processed may contain a significant number of “don't care” files. For example, a search for external files is probably intended to locate source files (with suffixes like **.c** or **.cc** or **.java** or **.bas**) that you've forgotten to place under version control. You probably don't care about program-generated files with suffixes like **.exe** (executables built in the source directory), **.bak** (editor backup files), **.msg** (copies of mail messages, and so on — because you don't intend to place them under version control).

You can use the environment variable **ACCUREV_IGNORE_ELEMS** to specify up to 50 patterns (or even individual filenames/pathnames). When it executes certain “whole-workspace” commands, the **accurev** tool ignores all external files that match this specification. For example, setting **ACCUREV_IGNORE_ELEMS** to the following value causes the **stat -x** command to ignore all **.exe** and **.bak** files:

```
*.exe *.bak
```

ACCUREV_IGNORE_ELEMS is also used — in a slightly different way — by certain commands that process a particular set of files, instead of the whole workspace.

AccuRev now supports specification of objects to be ignored on a per-directory basis, using configuration files, in addition to the global specification in environment variable **ACCUREV_IGNORE_ELEMS**. The following sections explain the details of using the environment variable and the configuration files.

Eligible “Whole-Workspace” Commands

The following **accurev** commands use the value of `ACCUREV_IGNORE_ELEMS` to filter names:

Command	Description
stat -x	List external (non-version-controlled) files
stat -m	List files that have been modified
stat -n	List files that have been modified, but are not active from AccuRev’s viewpoint (that is, are not in the workspace’s default group)
stat -p	List files that are pending promotion — files that have been modified, along with files with (kept) status
add -x	Place external files under version control
update	Update the workspace by incorporating versions from the backing stream (only if <code>USE_IGNORE_ELEMS_OPTIMIZATION</code> is set to <code>TRUE</code>)

These commands apply to the entire workspace if you *don’t* specify any filenames/pathnames or wildcard patterns on the command line:

```
accurev stat -n           (“whole-workspace” command)
accurev stat -n *.doc    (pattern specified; not a “whole-workspace” command)
```

When applying these commands to an entire workspace, the **accurev** tool, running on the client machine, uses `ACCUREV_IGNORE_ELEMS` to filter the list of filenames *before* sending the list to the AccuRev Server process. This can significantly reduce the amount of network traffic, and also reduce the amount of file-status computation the Server process must perform.

Note: by default, the **stat** command also uses a timestamp-based optimization to reduce the number of files it must process. For details, see *Optimized Search for Modified Files — the Scan Threshold* on page 218 in the *AccuRev CLI User’s Guide*.

GUI Counterparts to the “Whole-Workspace” Commands

The AccuRev GUI also uses the `ACCUREV_IGNORE_ELEMS` environment variable. Working in the searches pane of the File Browser corresponds to using the various “whole-workspace” forms of the **accurev stat** command. Thus, `ACCUREV_IGNORE_ELEMS` is used by these search criteria:

- External (**stat -x**)
- Modified (**stat -m**)
- Non-member (**stat -n**)
- Pending (**stat -p**)

Commands that Don’t Apply to the Whole Workspace

The **stat** and **add** commands above accept filename/pathname specifications, in several forms:

- Individual filename: **chap01.doc**
- Individual pathname: **doc/chap01.doc** or **./widgets/doc/chap01.doc**

- Wildcard pattern: ***.doc** or **docs/*.doc**
- list-file: **-l my_list_of_files**

Such specifications restrict the command to processing a certain set of files, not the whole workspace (even if you also specify **-x**, **-m**, **-n**, or **-p**). Similarly, the **files** command processes a certain set of files, not the whole workspace.

For these commands, the **accurev** tool still uses `ACCUREV_IGNORE_ELEMS` — but in a way that is less efficient than for whole-workspace commands:

1. Send the list of files that match the filename/pathname specification to the AccuRev Server.
2. Use `ACCUREV_IGNORE_ELEMS` to filter the data that the Server returns. Only names of external files are filtered out; files that are AccuRev elements remain in the listing, even if they match the value of `ACCUREV_IGNORE_ELEMS`.

This may produce similar, or even identical results as a whole-workspace command, but the fact that an unfiltered list is sent to, and processed by, the Server means that overall performance won't be as good.

Values for `ACCUREV_IGNORE_ELEMS`

The value of the `ACCUREV_IGNORE_ELEMS` environment variable must be a SPACE-separated list of filenames, pathnames, and wildcard patterns. Some examples:

```
*.exe
*.exe *.doc
manuals/*.doc
*.doc README.html
```

You can use any of these wildcards:

- **?** matches any one character
- ***** matches any sequence of characters (including a zero-length sequence — see note below)
- **[aekz]** matches **a**, **e**, **k**, or **z**
- **[a-e]** matches **a**, **b**, **c**, **d**, or **e**
- **{one,two,seven}** matches **one**, **two**, or **seven**
- **** parses the next character (for example, *****) literally, not as a wildcard

Note: the asterisk (*****) wildcard works a bit differently here than in standard Unix/Linux and Windows command processors. It matches any number of characters, including the directory separator (**/** or ****). With standard Unix/Linux and Windows command processors, the scope of ***** is restricted to a single pathname component.

So ***.doc** matches any of these names:

```
chap01.doc
manuals/chap01.doc
widgetproj/src/manuals/usergd/chap01.doc
```

The pattern **manuals/*.doc** matches any of these names:

```
manuals/chap01.doc
manuals/usergd/chap01.doc
```

... but not this name:

```
widgetproj/src/manuals/usergd/chap01.doc
```

The pattern ***/manuals/*.doc** (or ***manuals/*.doc**) *does* match the last name.

You can use the following pattern to ignore all items in the directory tree(s) named **usergd**:

```
*/usergd*
```

Specifying Directories and Their Contents

A typical application of `ACCUREV_IGNORE_ELEMS` is to have **stat -x** (“list all external files”) ignore temporary build directories. That is, you want the listing to exclude both the directories themselves and all the files within those directories. If the build directories are named **build_001**, **build_002**, etc., you might be tempted to use this pattern:

```
*/build_??*/*
```

But this pattern matches only the *contents* of the directories, not the directories themselves. Instead, use the following value for `ACCUREV_IGNORE_ELEMS`:

```
*/build_??? */build_??*/*
```

(The single pattern ***/build_??*** would match both directories and their contents. But it also might coincidentally match names of some source files, such as **lib/build_end.c**.)

Setting ACCUREV_IGNORE_ELEMS on a Unix/Linux System

When setting the `ACCUREV_IGNORE_ELEMS` environment variable on a Unix or Linux system, be sure to single-quote or double-quote the value, in order to protect any wildcard characters from being expanded by the shell:

```
export ACCUREV_IGNORE_ELEMS="*.exe *.doc"      (Bourne shell family)
setenv ACCUREV_IGNORE_ELEMS "/*.exe *.doc"     (C shell family)
```

To determine the current value of `ACCUREV_IGNORE_ELEMS`, use either of these commands:

```
env | grep ACCUREV_IGNORE_ELEMS                (or a shorter 'grep' pattern)
echo "$ACCUREV_IGNORE_ELEMS"                   (don't forget the quotes!)
```

Setting ACCUREV_IGNORE_ELEMS on a Windows System

On a Windows system, you can set the `ACCUREV_IGNORE_ELEMS` environment variable in the System applet (on the Control Panel). Alternatively, use the **set** command in a Command Prompt window:

```
set ACCUREV_IGNORE_ELEMS=*.exe *.doc           (no quotes!)
```

Don't use quote characters, even if the value includes SPACES.

To determine the current value of `ACCUREV_IGNORE_ELEMS` in a Command Prompt window, use either of these commands:

```
set  
echo %ACCUREV_IGNORE_ELEMS%
```

Per-Directory Pathname Optimization — the `.acignore` File

AccuRev's pathname optimization facility provides for faster performance by allowing certain objects to be ignored during various commands — notably **(external)**-status objects during whole-workspace searches. The value of environment variable `ACCUREV_IGNORE_ELEMS` is a pathname pattern — or several patterns, separated by SPACES; objects whose pathnames match a pattern are ignored by the various commands.

AccuRev now supports specification of objects to be ignored on a per-directory basis, in addition to the global specification in `ACCUREV_IGNORE_ELEMS`. One or more directories in a workspace can contain a text file named **.acignore**. An **.acignore** file works the same way as `ACCUREV_IGNORE_ELEMS`, with these exceptions:

- In an **.acignore** file, multiple patterns must appear on separate lines — not SPACE-separated on a single line.
- An **.acignore** file applies only to its own directory. In particular, it does *not* apply recursively to lower-level directories.

AccuRev Server Performance

This technical note is targeted to administrators planning a new enterprise deployment of AccuRev configuration management software, as well as to those planning the growth of an existing deployment. It will also be useful for administrators whose existing deployment is experiencing lower-than-expected performance.

When evaluating AccuRev performance, it is important to use the actual hardware that you will be using in production, or hardware of the same class.

Overview

AccuRev uses a standard client-server architecture:

- The **AccuRev Server** software runs on one machine, as a single multi-threaded operating system process, named **accurev_server**. The Server manages the AccuRev repository, where all data under AccuRev management is stored.
- **AccuRev Client** software runs on each AccuRev user's machine. The command-line interface is implemented as a single program, **accurev**, which executes a single command and then exits. The graphical user interface is implemented as two programs: the user-visible Java application **acgui** is long-lived; when the user invokes a command through the GUI, **acgui** runs the background application **accurevw**, which communicates with the Server and then exits.

User workspaces, where all development work is performed, are typically stored on users' machines. (But they can be stored on network file servers, or even on the AccuRev Server machine.)

- Each Client submits requests to the Server through a TCP/IP network. AccuRev Clients never communicate directly with each other.

As the above description indicates, the AccuRev Server is the heart of the AccuRev system. The sizing and configuration of the AccuRev Server machine significantly impacts the performance perceived by users.

We highly recommend that you first read the chapter *The AccuRev Data Repository* in the *AccuRev Concepts Manual*, along with the entire *AccuRev Administrator's Guide*. We also recommend that you attend at least the AccuRev QuickStart Training and the AccuRev Administrator's Training courses offered by AccuRev, Inc.

Key Factors Affecting AccuRev Performance

This diagram illustrates the AccuRev “performance stack”. The highlighted components are the ones that are under customer control — each component can influence the user’s experience of AccuRev either positively or negatively. In particular, a problem at any level can severely degrade the user’s experience of AccuRev performance. We strongly recommend that a top-to-bottom assessment of all of these components be performed prior to deploying AccuRev on the enterprise level.

The following sections discuss the aspects of an AccuRev deployment that can be “performance tuned” through knowledgeable choices.

User Knowledge
<i>AccuRev GUI</i>
<i>AccuRev Client Software</i>
AccuRev Setup/Usage
Client Hardware
Network Topology
<i>AccuRev Server Software</i>
AccuRev Server Configuration
AccuRev Server Hardware
Disk Storage

Resource Utilization and the Client-Server Architecture

The AccuRev Server is a multi-threaded application. Each individual AccuRev operation, such as executing a **History** command, uses only one thread, which runs on one CPU. When multiple client command requests arrive at the AccuRev Server at (approximately) the same time, these commands are executed through normal operating system multitasking. On a multiple-CPU machine, the threads are shared across all available CPUs.

AccuRev is architected to minimize client-server interaction and network bandwidth. Operations such as viewing files, editing files, and performing builds do not require any interaction with AccuRev. Thus, even during periods of intensive development, each user performs actions that require interaction with AccuRev relatively rarely. It is only when a user wishes to add a new file to a depot, rename a file, get the history of a file, get the latest version of a file, or similar operations that interaction with AccuRev will occur. Even then, these are short-running operations that put little load on the AccuRev Server software.

Experience indicates that, on average, a user will actually “hit” the AccuRev Server for less than 5 minutes per day. Typically, Server utilization is not uniform throughout the day, but tends to form peaks and valleys. For example, you might find that multiple users **Keep** their changes just before the nightly build is started, or that users tend to **Update** their workspaces at approximately the same time (e.g. when they arrive at work, or just after lunch).

Customer data has consistently shown that “peak load” rarely exceeds 15%. That is, you are unlikely to experience a situation in which more than 15% of the active AccuRev users have overlapping requests being processed by the AccuRev Server. (You can determine the number of active users by examining the Server log. Or you can simply assume that all the users listed by the

command **accurev show users** are active.) The most common command request during peak load periods is **Update**. This also is the AccuRev command that requires that most processor, disk, and network resources.

AccuRev Server Machine Usage: Software

AccuRev uses a single, central data repository, located on the AccuRev server machine. Any software running on the AccuRev Server machine that is not specifically supporting the AccuRev installation runs the risk of contending for the resources that AccuRev needs. For instance, if the machine is also acting as a file server for builds or other purposes, then users may find that AccuRev seems slow, even when there are no other AccuRev users currently making requests.

For optimal performance, the server machine should be completely dedicated to AccuRev usage. NFS, Samba, httpd, print queues, and other similar services should all be disabled. Only AccuRev administrators should use the system, and then only to service the AccuRev machine or software.

Reference trees and/or build workspaces should be located on client (user) machines or on dedicated build machines, not on the AccuRev Server machine. Software builds should not use the AccuRev Server machine's processor.

AccuRev Server Machine Networking

We strongly recommended that all build machines be located on the same subnet as the AccuRev Server machine, and that this subnet not include any other machines. The build subnet should use gigabit Ethernet. It is important that all of the equipment in the subnet be locked down to gigabit speeds and not be allowed to autonegotiate. Autonegotiation typically causes the subnet to operate, eventually, at the lowest supported speed.

AccuRev's network bandwidth usage is similar to that of **ftp**. For instance, a **keep** command on 100 files needs to send a list of 100 files to the server, in addition to the contents of those files. AccuRev is transaction-based; it often works on groups of files in a single transaction. Thus, the full list of file names and contents is sent to the server in one burst. This minimizes changes in the direction of data transfer, effectively removing the impact of any network latency.

Secure Networking

There are three ways to implement secure communication between two AccuRev sites:

- Most efficient is a dedicated connection between the two sites. Having a dedicated connection guarantees privacy and security. There is no performance penalty on a dedicated line.
- Next most efficient is to tunnel via a secure socket (SSH). The encryption has very little impact on bandwidth. Most SSH software offers compression, which can actually increase the effective bandwidth of the connection.
- Least efficient is VPN. Depending on the implementation, VPN overhead can consume up to 20% of the available bandwidth.

Using Multiple AccuRev Servers

The simplest way to use AccuRev is to have a single repository, managed by a single AccuRev Server. The server process runs on a particular machine, and listens for client requests on a particular IP port. But it's perfectly possible for an organization to have multiple servers running concurrently on different machines. Each server manages a separate repository, located on the server's own machine; the servers and their repositories are completely independent (and unaware) of each other.

Setting Up Multiple Server Machines

There are no special considerations, and no special procedures, for setting up multiple machines to run AccuRev Servers. You can install a server on as many hosts as desired within a local area (or wide area) network. The fact that AccuRev is installed on machines that are networked together is irrelevant; each installation is independent of all the others.

Setting Up Client Machines

Performing a client installation — or a combined client/server installation — on a machine creates a client configuration file, **acclient.cnf**, in the AccuRev executables directory. (That's subdirectory **bin** of the location designated as the “Install Folder” or “Install Directory” during installation.) The configuration file specifies the network location of the IP port on which an instance of the AccuRev Server is listening for client requests. For example:

```
SERVERS = jupiter:4079
```

Be sure to select a *custom* client installation, not a *typical* one. In a custom installation, you are prompted to supply the information to be stored in the **acclient.cnf** file — both the “host” (name of server machine) and “port” (IP port number). In a *typical* installation, the configuration file automatically gets a line specifying port 5050 on the local machine.

If more than one server has already been set up, you'll need to pick one of them during a custom client installation. Check with a system administrator, or look at file **acserver.cnf** in the AccuRev executables directory on the machine where the server is installed. For example, the server machine's **acserver.cnf** file might contain:

```
MASTER_SERVER = jupiter
PORT = 4079
...
```

Configuring a Client Machine to Use Multiple Servers

The client configuration file, **acclient.cnf**, is an editable text file. You can modify it at any time to identify additional AccuRev servers. For example, in a network where AccuRev servers are

running on machines **jupiter**, **venus**, and **pluto**, each client machine's **acclient.cnf** file might look like this:

```
SERVERS = jupiter:4079
SERVERS = venus:5050
SERVERS = pluto:6678
```

Each server's IP port must be listed on a separate line (even though the keyword is "SERVERS", not "SERVER"). If an existing AccuRev server isn't listed in a machine's **acclient.cnf** file, AccuRev client programs running on that machine won't be able to communicate with that server.

Note: the file must not contain any empty lines; be sure to check the end of the file.

In what order should the servers be listed? For users of the AccuRev GUI (**acgui**), it doesn't matter: users can switch among all the servers with the GUI's **Tools > Change Active Server** command. For users of the AccuRev CLI (**accurev**), the order is important: by default, each AccuRev CLI command is directed to the server listed on the *first* line of the client configuration file. This implies that if you want to change the "active server" in the CLI, you need to rearrange the lines in your machine's **acclient.cnf** file. This is quite doable, but cumbersome; the sections below describe a better way for CLI users to work with multiple AccuRev servers.

Note: the GUI's **Tools > Change Active Server** command revises the **acclient.cnf** file, moving the line for the new active server to the beginning of the file.

Workspaces and Servers

Each AccuRev workspace is associated with a particular AccuRev server: the workspace is attached to a particular stream, which belongs to a particular depot, which is managed by a particular server. When you execute a **mkws** command to create a new workspace, the command is directed to the server that is listed first in the machine's **acclient.cnf** file. (You must use the **-s** option to specify a backing stream for the new workspace; this stream must belong to one of the depots managed by that particular server.)

Note: It's more precise to describe the workspace's association as being with a particular AccuRev *repository*, rather than with a particular AccuRev *server*. This workspace-to-repository association is permanent: AccuRev has no concept of associating an existing workspace with a different repository (or even with a different depot in the same repository). On the other hand, many of the details can change: you can rename a workspace; you can move a workspace to a different location in the file system or to a different machine; you can move a repository to a different machine.

Suppose you want to create one workspace for use with a depot on server **venus**, and another workspace for use with a depot on server **pluto**. Here's how:

1. Make sure the **acclient.cnf** file lists **venus** first.
2. Create a workspace with **mkws**, specifying a stream in any of **venus**'s depots as the backing stream.
3. Edit the **acclient.cnf** file, so that the **pluto** entry is first.

4. Create a workspace with **mkws**, specifying a stream in any of **pluto**'s depots as the backing stream.

Many AccuRev CLI commands (examples: **keep**, **promote**, **merge**, **hist**) require that your current working directory be within an existing workspace. By default, such commands succeed only if the server associated with the workspace is the first one listed in the machine's **acclient.cnf** file. So, for example, an error would occur if we continued the example above as follows:

5. Go to the workspace created in Step 2.
6. Execute the command **accurev hist**.

An error would occur because the workspace is associated with the **venus** depot, but the first line in the **acclient.cnf** file still lists the server on **pluto** (changed in Step 3). Are we saying that every time you want to use the **venus** workspace, you need to make sure the **acclient.cnf** file lists the **venus** server first? Well, that's the way AccuRev works by default; but fortunately, there's a better solution to the "switching servers" problem.

The 'wspaces' File

If the workspaces you've created on your client machine are not all associated with the same AccuRev server, then you'll probably want to use a **wspaces** configuration file. This file records the workspace/server association for one or more workspaces. The AccuRev CLI client program, **accurev**, uses this file to decide which AccuRev server to use:

- If the current working directory is within one of the workspace-root-directory locations listed in the **wspaces** file, the CLI command is directed to the associated server listed on the same line of the file.
- Otherwise (including the case in which no **wspaces** file exists), the CLI command is directed to first server listed in the machine's **acclient.cnf** file. This is the default behavior that we've mentioned in the sections above.

If your site has a single AccuRev server (or if your workspaces are all associated with a single server), there is no need for a **wspaces** file: the default behavior directs each CLI command to the one and only server listed in your machine's **acclient.cnf** file.

Format of the 'wspaces' File

The **wspaces** file is a text file, each line containing four whitespace-separated fields:

- workspace name
- full pathname of workspace's root directory
- server machine name (or IP address)
- IP port number on which AccuRev server program listens

Thus, each line records the association between one workspace and its AccuRev server.

For example, here's a **wspaces** file recording workspaces associated with servers **venus** and **pluto**:

```
gizmo_dvt_jjp      /usr/jjp/gizmo_dvt      venus  5050
frammi_2.3maint_jjp /usr/jjp/frammi_2.3_maint pluto  6678
```

Some notes:

- The workspace pathname (2nd field) must match the pathname listed by the **accurev show wspaces** command. The pathname always uses forward slashes (/), not backslashes (\), even on Windows machines.
- The last two fields must exactly match an existing entry in the machine's **acclient.cnf** file — except for the fact that **acclient.cnf** uses a colon (:) to separate the machine name and the IP port number, whereas **wspaces** uses white space to separate them.
- Workspace names and pathnames that contain SPACE characters are not handled correctly by this facility.

Location of the 'wspaces' File

Each user has his own **wspaces** file, recording information about his own workspaces. The **wspaces** file must be located under your home directory, within subdirectory **.accurev**. On Windows systems, your home directory is determined by concatenating the values of environment variables HOMEDRIVE and HOMEPATH. On some versions of Windows, you must set these environment variables yourself; they are not set automatically by the operating system.

The -H Command-Line Option

For most commands in the **accurev** CLI program, you can specify the AccuRev server/repository to target on the command line, using the **-H** option:

```
accurev show -H pluto:6678 users
```

Note that the **-H** option follows the command name — in this example, **show** — not the program name, **accurev**. The hostname/port-number argument to this option has the same form as in the **acclient.cnf** file.

This mechanism bypasses the **acclient.cnf** file, though the file must still exist. It does *not* override a specification in the **wspaces** file.

Techniques for Sharing Workspaces

This note describes two techniques for accessing the same workspace from multiple machines.

Accessing a Windows Workspace From Multiple Windows Clients

Multiple AccuRev users, on Windows client machines, can share a workspace that is physically located on a Windows machine. (Or maybe there's just one user, who wants to access a workspace from multiple Windows client machines.)

1. Designate a directory that is *above* the top-level directory of the workspace tree as a Windows "shared directory". For example, if the workspace tree for workspace **widget_maint_derek** on machine **derekpc** is located at **C:\widget\workspaces\maintdrp**, you could set the shared directory as follows:

```
net share widgwork=C:\widget\workspaces
```

Note: the workspace tree's top-level directory itself (**maintdrp** in the example above) cannot be designated as the shared directory.

2. Determine how AccuRev records the workspace tree location, using the command **accurev show wspaces**. (The pathname will always use forward slashes, even if it's a Windows pathname.)

- If the workspace tree location incorporates the share name ...

```
widget_maint_derek    /widgwork/maintdrp    derekpc    ...
```

... skip to Step 3.

- But if the workspace tree location appears as an absolute pathname ...

```
widget_maint_derek    C:/widget/workspaces/maintdrp    derekpc    ...
```

... you must use the **chws** command to change the recorded location to a pathname that incorporates the share name. This involves mapping a network drive to the shared directory:

```
> net use K: \\derekpc\widgwork
The command completed successfully.
```

```
> K:
```

```
> cd \maintdrp
```

```
> accurev chws -w widget_maint_derek -l .      ("dash-ell dot")
Changed location.
Changed machine name.
```

```
> accurev show wspaces
```

```
...
widget_maint_derek    /widgwork/maintdrp    derekpc    ...
```

3. Now, all users on Windows client machines can access the workspace tree by mapping a network drive to the shared directory. Even the user on the machine where the workspace tree is located (**derekpc** in our example) must use a network drive to access the workspace tree.

```
> net use P: \\derekpc\widgwork
The command completed successfully.
```

```
> P:
```

```
> cd \maintdrp
```

```
> accurev info
```

```
...
Workspace/ref:  widget_maint_derek
Basis:          widget_maint
Top:            P:/maintdrp
```

Users on different machines can map the shared directory to different drive letters, and access the workspace as, for example, **Y:\maintdrp** or **R:\maintdrp**.

Universal Access to a Workspace Located on a Share

A workspace whose workspace tree is network-accessible through a share, can be accessed from any client machine — running Unix/Linux or Windows. The share can be configured through Samba/SMB or some other network file system. It must make the actual storage location available through a machine name and a simple “share name”: a name that looks like a single pathname component.

The ‘share_map.txt’ File

The technique in *Accessing a Windows Workspace From Multiple Windows Clients* on page 41 relies only on Windows operating system facilities. But the “universal workspace access” technique requires the maintaining of a pathname-mapping file for use by the AccuRev Server. If a “share” (that is, shared directory) has an entry in the pathname-mapping file, any workspace located on that share can be used on all AccuRev client machines capable of accessing the machine where the share resides.

The pathname-mapping file is a text file, **share_map.txt**, which must be located in the AccuRev **site_slice** directory on the AccuRev Server machine. It maps share names to absolute pathnames. Each line of **share_map.txt** consists of three TAB-separated fields, describing one share. For example:

```
jupiter    accwks    /public05/accurev_workspaces
```

This entry says, “a share named **accwks** is physically located on machine **jupiter**, at absolute pathname **/public05/accurev_workspaces**”. More generally:

- The first field (**jupiter**) names a machine where one or more workspace trees are (or will be) located.
- The second field (**accwks**) specifies the share name.
- The third field (**/public05/accurev_workspaces**) indicates the absolute pathname of the share on the machine. On a Windows machine, this includes the drive letter — for example, **C:/Public Directories/AccuRev Workspaces**.

Notes:

- All pathnames in **share_map.txt** must use forward-slash characters (/), even Windows pathnames.
- The field separator in the **share_map.txt** file must be single TAB character — don't use SPACES. If a specification (e.g. a share name) includes a SPACE character, do not enclose the specification in quotes.

Workspace Location Entries

The **accurev show wspaces** (or the GUI's **View Workspaces**) command displays the repository's workspaces table, showing the locations of existing workspaces. The pathnames always use forward slashes, even if they are Windows pathnames.

AccuRev can record the location as an absolute pathname on its machine:

```
C:/wks/light24/mnt_john      (Windows)
/bigdisk/home/john/widget_devel  (Unix/Linux)
```

Alternatively, it can record a location that incorporates a share name:

```
/accwks/wks_john      (Windows or Unix/Linux)
```

Universal workspace access requires that a workspace's location be recorded as an absolute pathname in the workspaces table. (Note that the sharing technique described in *Accessing a Windows Workspace From Multiple Windows Clients* on page 41 has the opposite requirement: workspace locations must incorporate the share name.) In addition, the "Host" name listed in this table for a workspace must exactly match the first field in some **share_map.txt** entry. Beware of domain name discrepancies — for example, **jupiter** vs. **jupiter.mycorp.com**.

Fixing Workspace Location Entries

If a workspace located on a share has the "wrong kind" of entry in the workspaces table, fix it to enable universal client access:

1. Make sure that **share_map.txt** has a valid entry for the share.
2. On any client machine that can "see" the workspace, go the top-level directory of the workspace tree.
3. Use the **chws** command to change the location recorded in the workspaces table to an absolute pathname:

```
> accurev chws -w <workspace-name> -l .      ("dash-ell dot")
```

You must fix each workspace location entry in this way individually.

Fixing Machine Name Entries

If there's a discrepancy between a machine's name in a **share_map.txt** entry (say, **jupiter**) and its name in the workspaces table (say, **jupiter.mycorp.com**), change the workspace table entry:

```
accurev chws -w <workspace-name> -m jupiter
```

Example: Samba Share

Here's an example of how it can all work in a Samba environment, elaborating on the scenario above:

1. The organization decides that a directory, **/public05/accurev_workspaces**, on Unix host **jupiter** will be a location where users can create workspaces that can be shared across platforms.
2. The system administrator on **jupiter** turns that directory into a Samba share, named **accwks**. Here's the relevant excerpt from the Samba **smb.conf** file on host **jupiter**:

```
[accwks]
    comment = All users
    path = /public05/accurev_workspaces
    browseable = yes
    guest ok = yes
    writeable = yes
```

3. The AccuRev administrator makes this entry in the **share_map.txt** file, in the AccuRev Server's **site_slice** directory.

```
jupiter    accwks    /public05/accurev_workspaces
```

4. User **john**, working on a Windows machine, wants to create a workspace on the share. First, he makes the share accessible as a network drive:

```
net use T: \\jupiter\accwks
```

5. Then **john** creates his workspace on this network drive:

```
accurev mkws -w shrwks_john -b dvt_stream -l T:\wks_john
```

A **show wspaces** command indicates that the AccuRev Server uses an absolute pathname to record the new workspace's location on **jupiter**:

```
shrwks_john    /public05/accurev_workspaces/wks_john    jupiter    ...
```

6. **john** can now use this workspace from any client machine that can access the machine where the share resides.

What's the Difference between Populate and Update?

AccuRev users sometimes confuse the two commands **Populate** and **Update**. These commands seem similar because they both bring new data into your workspace. But they are quite different, both in their usage pattern — most people use **Update** far more often — and in what they accomplish. Understanding the difference between these two commands will enable you to choose the right command at the right time (always useful!), and will deepen your knowledge of how AccuRev really works.

Note: the AccuRev CLI command **accurev pop** corresponds to the GUI's **Populate** command.

This note starts with a brief statement of the difference between **Populate** and **Update**, along with a few examples. Then, we present a full discussion of the data structures and mechanisms involved in these commands.

In a Nutshell ...

The essential difference between **Populate** and **Update** concerns time. Roughly speaking, your workspace contains an informal “baseline” (the contents of the shared backing stream, at a particular moment) plus “changes” (the modifications that you make to some of the files). The **Update** command advances the workspace’s baseline from the time of the workspace’s last update to the present moment. This incorporates into the workspace data recently placed in the backing stream by other team members.

Note: AccuRev actually tracks the workspace’s baseline in terms of transactions, not timestamps.

The **Populate** command doesn’t advance a workspace’s baseline at all, but leaves it “stuck in the past”. Instead, **Populate** simply restores the appropriate “old” version of one or more elements that are currently missing from the workspace.

The two commands also differ in their scope: **Update** always processes the entire workspace; **Populate** processes just a selected set of elements or directory subtrees.

The capsule description above uses imprecise language, such as “advancing the workspace’s baseline” and “old version”. The following description is more precise, using AccuRev-specific terminology. The terms are explained fully in section *Data Structures Used by Populate and Update* on page 47 below.

The **Update** command changes both the workspace stream and the workspace tree:

- It advances the workspace stream’s update level to the depot’s most recent transaction — say, from current update level 32155 to new update level 34002. This allows a new set of versions — in this case, some or all the versions created by transactions 32156 through 34002 — to flow into the workspace stream from the backing stream.
- It copies the contents of the workspace stream’s new versions from the repository’s file-storage area to the workspace tree. When processing a sparse workspace, **Update**

bypasses elements with **(missing)** status. (It assumes that you have chosen not to load these elements into the workspace tree.)

By contrast, the **Populate** command changes the workspace tree only, not the workspace stream. In particular, it doesn't change the workspace stream's update level. **Populate** merely fixes a discrepancy between the workspace stream and the workspace tree: a certain version of a file is in the workspace stream, but there is no actual file in the workspace tree — that is, the file's status is **(missing)**. To fix this situation, you invoke **Populate**, which copies the version currently in the workspace stream to the workspace tree.

Note: it would be incorrect to conclude that **Update** *never* processes **(missing)** elements, and that **Populate** *only* processes **(missing)** elements. Examples 3 and 4 below show that exceptions exist for both these “rules”.

Example 1: Standard Update Scenario

You've just finished a coding project, so you're not actively working on any files in your workspace. Other team members create new versions of files **red**, **white**, and **blue** in their workspaces, then promote those versions to the team's backing stream. You invoke the **Update** command, which copies the most recent versions of **red**, **white**, and **blue** from the backing stream to your workspace.

Example 2: Restoring a Deleted File (“missing” by accident)

Since you have complete control over the files in the workspace tree, it's easy to accidentally delete a version-controlled file with an operating-system command or a third-party tool. If you do this, AccuRev knows that the file *should* be there, because a version of the element still exists in the workspace stream. Thus, the File Browser continues to list the element, but shows it as **(missing)** from the workspace tree. You select the element and invoke **Populate** to fix the accidental deletion.

Example 3: Adding Data to a Sparse Workspace (“missing” by design)

Note: this example has become less likely after Version 3.5, which no longer creates sparse workspaces. The include/exclude facility replaces the functionality of sparse workspaces.

A sparse workspace is one in which only some of the elements in the workspace stream are loaded into the workspace tree. That is, some of the files are missing by design — because you don't want them in this particular work area. But it's all the same to the File Browser: it shows the non-loaded elements as **(missing)**. As above, you select the elements and invoke **Populate**. In this case, instead of fixing an accidental deletion, it loads previously non-loaded elements into the workspace tree.

What would **Update** do in the situations of Examples 2 and 3?

- If you accidentally delete a file in a standard workspace, a subsequent **Update** *might* copy a new version of the file to your workspace tree; it does so only if someone else promoted a new version of the file to the backing stream since your last **Update**.
- But in a sparse workspace, **Update** *never* processes an element with **(missing)** status.

Example 4: Handling Active Elements

Update and **Populate** differ in how they handle elements that are active (are in the workspace's default group). The **Update** story is simple: it *never* overwrites the file in the workspace tree.

Populate usually doesn't overwrite the file, but there are a couple of cases to consider.

- It doesn't *need* to overwrite a file that you've kept and *not* subsequently edited, because the active version in the workspace stream is identical to the file in the workspace tree.
- But if you have subsequently edited the file in the workspace tree, so that the element status is **(modified)(member)**, then you can order **Populate** to overwrite the file and clobber those subsequent edits. This can also happen with a file that you've edited, but never kept, so that its status is **(modified)**.

Be careful — **(modified)** files will also be overwritten if you invoke **Populate** with both the **Recursive** and **Overwrite** options on a directory that directly or indirectly contains the active element.



Example 5: A Tale of Two Files

Let's see how **Update** and **Populate** differ in this situation:

You have a standard (non-sparse) workspace that is completely up to date. You delete two files, named **blue** and **green**. Someone creates a new version of **blue** in another workspace, and then promotes it to your workspace's backing stream.

If you select both **blue** and **green** in the File Browser, then invoke **Populate**, the two files that you deleted are restored to the workspace tree. This does *not* bring in the new backing-stream version of **blue**, because that version is not in the workspace stream — it's too new, having been created after your workspace's most recent update.


If you invoke **Update** instead of **Populate**, the workspace tree gets the new version of **blue**. No version of **green** is copied to the workspace tree, because **Update** only handles new versions — ones that enter your workspace stream as a result of advancing its update level.

Data Structures Used by Populate and Update

The simplicity of AccuRev's day-to-day usage model stems, in large part, from the fact that you don't need to worry about the "big picture" of your organization's development scheme and process. Instead, you only need to concern yourself with:

- the workspace in which you maintain your private copies of version-controlled files
- the workspace's backing stream, a "data switchboard" that organizes the sharing of files' changes with other members of your development team

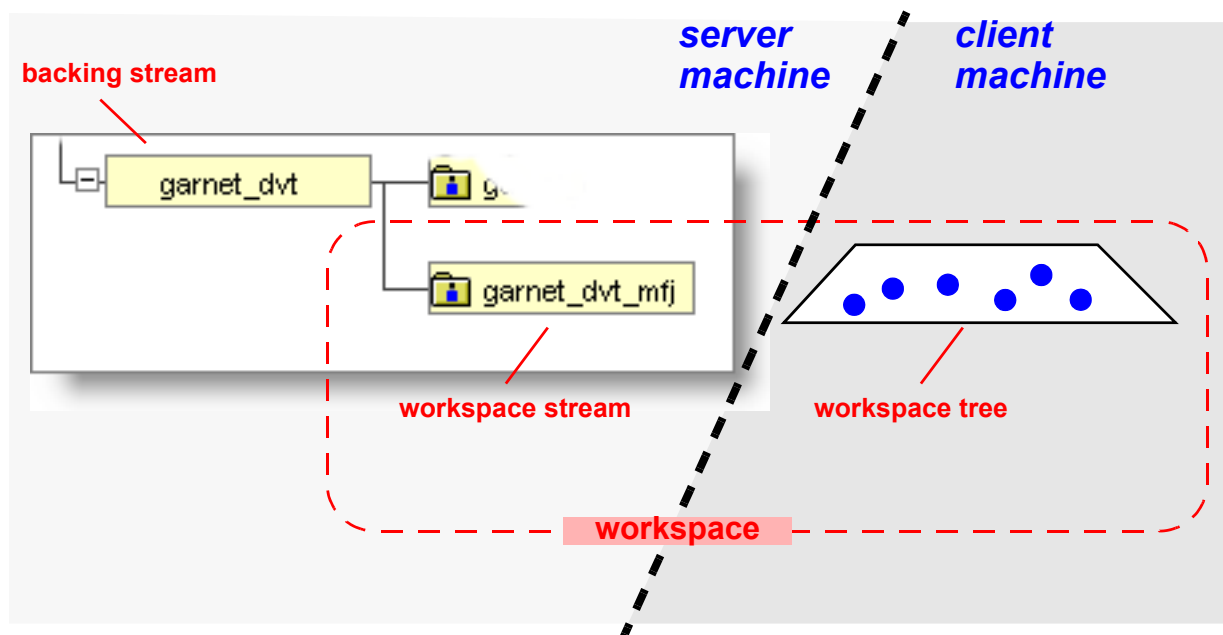
The illustration below shows how a workspace and its backing stream typically appear in the AccuRev StreamBrowser:



But a workspace actually consists of two parts:

- the workspace tree, an ordinary directory tree (“just a bunch of files”)
- the workspace stream, which contains all of the workspace’s configuration management information

So the picture looks more like this:



The above illustration shows one important difference between a workspace’s two parts: the workspace tree lives in “AccuRev client space”, while the workspace stream lives in “AccuRev Server space”. The following table summarizes all the important differences.

Workspace Stream	Workspace Tree
Resides in the database of an AccuRev depot, located on the AccuRev server machine	A standard directory tree, located on your client machine (or in other user-accessible storage)
Managed by the AccuRev Server process	Managed by you, the individual user

Workspace Stream	Workspace Tree
Contains all version control and configuration management information for the workspace, such as version-IDs	Contains no version control or configuration management information
Contains no actual files, just <u>version</u> objects that point to files in permanent storage	Contains <i>only</i> files and directories, which you can edit, compile, copy, etc.
Operating system commands and tools never change data here	Operating system commands and tools can change data here

The sections below expand on these differences.

How the Data Structures Get Their Data

Each of the data structures introduced above — backing stream, workspace stream, and workspace tree — is different in the way it gets changes (i.e. new data) from other parts of the development environment.

Backing Stream

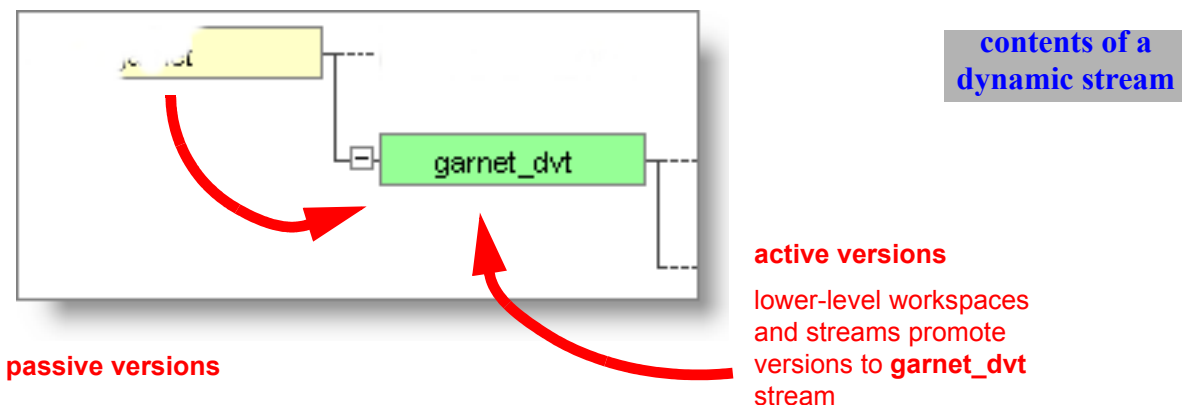
The backing stream is, in most cases, a dynamic stream. (It can also be a snapshot or a time-based stream.) A dynamic stream is a changing configuration of its depot. At any given moment, it (logically) contains a simple table that indicates particular versions of a set of elements. For example:

Element	Version-ID
doc	garnet_dvt\1
doc\chap01.doc	garnet_dvt\5
doc\chap02.doc	garnet\3
doc\chap03.doc	garnet_dvt\2
src	garnet\1
src\garnet.c	garnet_dvt\12
src\commands.c	garnet_dvt\7
tools	garnet\3
tools\start.sh	garnet_dvt\6
tools\end.sh	garnet\2

At any given moment, a dynamic stream's versions fall into two categories:

- **passive versions:** versions that the stream inherits from its parent stream. Inheritance is automatic and instantaneous: as soon as a new version enters the parent stream, it is inherited at once by the child stream.

- **active versions:** versions that have been **Promoted** to the stream, (usually) from lower-level workspaces and substreams. This set of versions constitutes the stream's default group.



In the configuration table above, all the **garnet_dvt\...** version-IDs indicate active versions in the **garnet_dvt** stream. All the **garnet\...** version-IDs indicate passive versions, inherited from the parent stream, named **garnet**.

Workspace Stream

The workspace stream is the “behind the scenes” part of your workspace. It lives inside the AccuRev repository, located on the AccuRev server machine. In many ways, it resembles a dynamic stream:

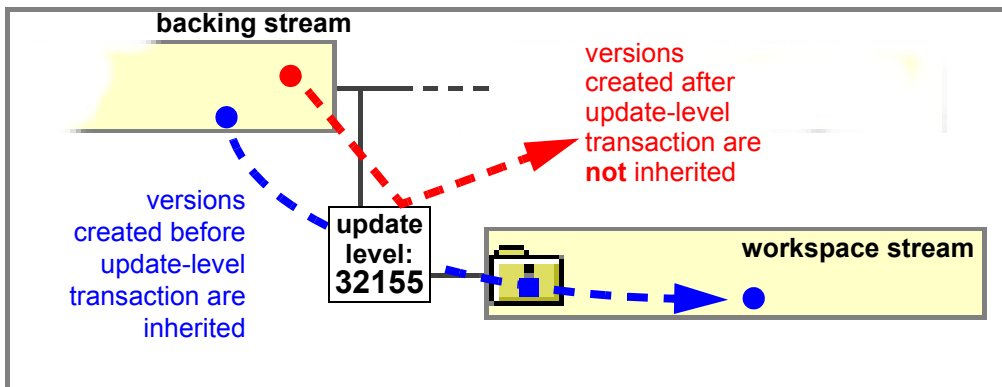
- It’s a changing configuration of one of the repository’s depots.
- It logically contains a particular version of some or all of the depot’s elements.
- It doesn’t contain actual files, but is logically just a table of elements and version-IDs.

Note: when creating a new version of a file, the **Keep** command copies the file to the repository’s file-storage area, not to the workspace stream itself. The workspace stream just gets a version-ID for the new version; the version-ID serves as a pointer to the file in the file-storage area.

- It contains active versions, created by explicit user commands: **Add**, **Keep**, **Rename**, **Defunct**, etc. The new versions preserve in the repository the changes that you’ve made to files in your workspace tree. (There’s only one way to create an active version in a dynamic stream: the **Promote** command.)
- It also contains passive versions, inherited from its parent stream, the workspace’s backing stream.

The last item is where the crucial difference between workspace streams and dynamic streams comes into play. The versions inherited by the workspace stream are not the ones *currently* in the backing stream, but the versions that *were* in the backing stream when the workspace was last updated. This is called the workspace’s update level. More precisely, AccuRev records the number of the depot’s most recent transaction — say, transaction #32155 — as the workspace stream’s update level. So we can rephrase the principle:

*The workspace stream inherits from the backing stream versions that were created in transactions up to and including the workspace's **update level**.*




Note: the update level of a workspace stream is very much like the optional basis time of a dynamic stream. Both mechanisms restrict the flow of versions to a child stream from its parent stream, based on a point in the depot's development history.

Thus, a workspace stream is not updated dynamically when changes occur to its parent stream. It gets new versions from the backing stream only when you issue an **Update** command. This is how AccuRev implements the workspace's user-controlled "priveness", isolating it from the changes regularly being recorded in the backing stream by other team members.

To summarize: at any given moment, a workspace stream contains:

- a set of passive, inherited versions, created in transactions that do not exceed the workspace's update level. (A file that you've **Promote**'d since the last update is an exception. The version is passive, but was created after the workspace's update.)
- a set of active versions, which you've created in that workspace with such AccuRev user commands as **Add**, **Keep**, **Rename**, and **Defunct**.



The Update Algorithm

This section describes the processing of the **Update** command in detail. It's not as simple as “get all the new versions”, because AccuRev takes care not to overwrite version-controlled files that you have changed, but whose changes have not yet been preserved in the repository.

1. AccuRev first searches the workspace tree for such “at-risk” files, by performing a CLI **stat** (file status) command on your workspace:
 - It uses the **-n** option to **stat**, which restricts the search to version-controlled files that you have modified but are not in the workspace's default group — the status of such “non-member” files includes the **(modified)** indicator but not the **(member)** indicator. It's safe to ignore default-group files, because these are never affected by an **Update**.
 - It uses timestamp optimization to speed the file search: it ignores files whose timestamps precede the workspace's scan threshold (the time that the workspace was most recently updated or otherwise searched for modified files). If you modify a file by overwriting it with a file with an old timestamp, the file will be ignored in this step. This can cause problems in Step 5 below.
 - If environment variable `USE_IGNORE_ELEMS_OPTIMIZATION` is set to `TRUE` (the value is case-insensitive), then it uses the value of environment variable `ACCUREV_IGNORE_ELEMS` to ignore certain pathnames in the file search. If an ignored file actually has **(modified)** status, an error will occur in Step 5 below.
 - To make sure that a file with a recent timestamp has actually been modified, it compares the file with the version currently in the workspace stream by performing a checksum. (This means that simply modifying a file's timestamp with a **touch** command won't prevent the file from being overwritten by **Update**. You have to make a real change to the file.)
2. If the preceding step found any “non-member” files — modified, but not in the default group — the **Update** command in AccuRev releases prior to Version 4.6 terminated immediately, without updating any file. Starting in Version 4.6, the **Update** command can sometimes proceed, even in the presence of such files:
 - A non-member file that is not due to be updated, because there is no newer version in the backing stream, does not prevent the update from proceeding.

- A non-member file that *is* due to be updated has **(modified)(overlap)** status. By default, the presence of one or more such files causes the update to terminate immediately, without updating any file. But you *might* be able to enable the update to proceed if you invoke an update option — **update -m** in the CLI, user preference **Update Resolves Trivial Merges** in the GUI. With this option invoked, the update proceeds only if a trivial merge can be performed for each file with **(modified)(overlap)** status. (The backing-stream version can be merged with the workspace-tree file automatically, because there are no conflicts that require manual intervention.)
3. AccuRev notes the number of the repository’s latest transaction, and sets that number as the workspace’s target update level.
 4. AccuRev decides which recently created versions in the repository should be delivered to the workspace tree. A version is a candidate for delivery if it became visible in the workspace’s backing stream by one of the transactions between the workspace’s *current* update level and the newly set *target* update level.
 5. AccuRev attempts to deliver all those versions to the workspace tree. Each time it is about to overwrite a file in the workspace tree, AccuRev first makes sure it won’t be “clobbering” unpreserved changes: it checksums the existing file to confirm that it matches the current version in the workspace stream (the version at the current update transaction level).

Note: it’s sometimes OK for the workspace tree to already contain the new version (the version at the target update level). See *Incomplete Updates* below for an explanation.

If a file about to be overwritten fails this checksum step, AccuRev reports a “crc mismatch” and the **Update** command terminates immediately. The most common cause of this error is your having overwritten the file in such a way that it gets an old timestamp. Such a file escapes detection in Step 1, but gets caught here — just in the nick of time to avoid being clobbered.

If the checksum succeeds, the file is safe to overwrite, so AccuRev updates it — finally! The update can be a replacement of the file’s contents, a change in its pathname, or both.

If AccuRev is processing files with **(modified)(overlap)** status (see Step 2 above), it automatically merges the backing-stream version with the file in the workspace tree, instead of simply overwriting the file. No checksum of the workspace-tree file is performed for such elements.

6. If *all* the recently created versions identified in Step 4 were successfully delivered to the workspace:
 - The target update level becomes the workspace’s current update level, indicating a successful, complete update.
 - The workspace’s scan threshold, used for the **stat -n** timestamp optimization (Step 1), is set to the time that this successful update began.

Advancing the Scan Threshold

The timestamp optimization is most effective when it enables AccuRev to ignore as many files as possible. Accordingly, AccuRev takes advantage of opportunities to validly advance the scan threshold to a later time:

- If a workspace is completely up-to-date, the **Update** command advances the scan threshold anyway, setting it to the time the command began. (Thus, an unnecessary **Update** doesn't "do nothing".)
- The CLI command **stat -n** advances the scan threshold to the point in time just before the *earliest* timestamp among the files it finds (modified files that are not members of the workspace's default group). This preserves the validity of the *timestamp optimization principle*: for file elements that are not in the workspace's default group, the timestamp of a modified file is later than the workspace's scan threshold.

If **stat -n** does not find any non-member modified files, it advances the scan threshold to the time that the command began.

If **stat -nO** finds non-member modified files whose timestamps precede the scan threshold, it moves the scan threshold *backward*, in order to preserve the timestamp optimization principle described above.

Incomplete Updates

The **Update** command is not implemented as an atomic operation, and it is not recorded as an AccuRev transaction. Transactions are used to organize and serialize changes to the central repository, not to the workspace trees that implement user "sandboxes". An **Update** can take a significant amount of time, and is sometimes interrupted before it completes — by user request, by network failure, by loss of telephone connection, etc.

For purposes of discussion, assume that your workspace, **talon_dvt_mary**, has a current update level of 84, that the highest transaction in the repository is 109, and that 45 files in your workspace would be involved in a complete **Update**. You can monitor transaction levels using the CLI command **show wspaces**:

- Before the update the output of **show wspaces** might include:

```
talon_dvt_mary      c:\wks\talon_dvt      xlnr      13 84 84 1 0
```

The first **84** indicates the workspace's target update level; the second **84** indicates the current update level.

- The **Update** command sets the target update level to **109**, then proceeds. If **Update** processes all files successfully, it raises the current update level to the target:

```
talon_dvt_mary      c:\wks\talon_dvt      xlnr      13 109 109 1 0
```

But if the **Update** does not complete successfully, the target update level remains unchanged. A subsequent **show wspaces** reveals that the target update level differs from current update level:

```
talon_dvt_mary      c:\wks\talon_dvt      xlnr      13 109 84 1 0
```

The differing update levels — target vs. current — is the telltale sign that the most recent **Update** did not complete successfully.

Incomplete Update: Command Interrupted

Consider the case in which **Update** was interrupted — say, after it had processed 29 out of the 45 files to be updated. When a subsequent **Update** command is issued:

- The checksum of the 29 files will succeed, because those files are already at the target transaction level.
- The checksum of the 16 files will fail, because those files are still at the current update transaction level.

(See Step 5 above for a discussion of the checksum process.)

Incomplete Update: Checksum Failure

Now consider the case of an incomplete **Update**, due to one or more “crc mismatch” errors. Suppose that only 42 out of 45 files are updated, because 3 files fail the checksum match. You must fix the problem before issuing another **Update**.

If those three files had been overwritten by mistake, you can restore the proper versions using the **Revert to Backed Version** command (CLI: **purge**). Then, a second **Update** brings the new versions of those 3 files into the workspace.

Performing the “Fixup” Update

When it begins executing an **Update** command, AccuRev determines whether the preceding update of the workspace completed successfully or not:

- The **Update** completed successfully if the target and current update levels are the same.
- The **Update** was incomplete if the target and current update levels differ.

If the preceding update was incomplete, AccuRev performs two updates at once. First, it performs a “fixup” update that completes the preceding update; then it performs an additional update (if necessary), to process changes made to the backing stream after the incomplete update. During the “fixup” update, AccuRev avoids the unnecessary work: it does not retransfer files that were successfully delivered to the workspace during the incomplete update.

Example:

- Your workspace’s current update level is 84, and the highest transaction in the repository is 109.
- You issue an **Update**, but it fails to complete. At this point, the workspace’s target update level is 109, but its current update level is still 84.
- You wait until after lunch break to reissue the **Update** command. At this point, the highest transaction in the repository is 137.

- AccuRev performs a “fixup” update, which brings the current update level to the original target update level, 109. Then, it advances the target update level to 137 and performs another update. If this update succeeds, it advances the current update level to 137.

Using a Trigger to Maintain a Reference Tree

Reference trees allow you to have a physical copy of the most recent sources for a stream. They are available for reference, thus the name reference tree. Snapshots never change, so they only need to be updated once using **update -r** and then you can forget about them.

To create a reference tree, use the **mkref** command. To keep a reference tree up to date with its associated stream, you need to run **update** on the reference tree every time versions are promoted to the stream.

AccuRev supplies the following trigger scripts to automate this procedure:

server_post_promote.pl

A general-purpose script, which can be used to perform various tasks after completion of every **promote** command. In this case, we're going to have it call the **update_ref.pl** script.

update_ref.pl

A script that invokes the **update** command to update the files in a reference tree. On a Unix/Linux machine, this script must be setUID-root.

The indirection is necessary for security purposes.

To enable the automatic updating of one or more reference trees, follow these steps:

1. Make sure the following Perl scripts are installed in some directory on the search path of the AccuRev Server process's user identity:

```
server_post_promote.pl
update_ref.pl
```

See *Operating-System User Identity of the Server Process* on page 9 of the *AccuRev Administrator's Guide*.

2. Edit both the **server_post_promote.pl** and **update_ref.pl** scripts, and follow the step-by-step instructions contained within them.
3. Windows only: convert the Perl scripts to Windows batch files:

```
pl2bat server_post_promote.pl
pl2bat update_ref.pl
```

4. Tell AccuRev to run the **server_post_promote** script after every **promote** command:

```
accurev mktrig server-post-promote-trig server_post_promote.pl  (Unix/Linux)
accurev mktrig server-post-promote-trig server_post_promote     (Windows)
```

For more information, see the descriptions of **mkref**, **mktrig**, and **show triggers** commands.

Notes for CVS Users

This note contains information that will be helpful for CVS users who are migrating to AccuRev.

AccuRev Workspaces vs. CVS Sandboxes

Each directory in a CVS sandbox has a subdirectory named **CVS**. This subdirectory stores metadata: where the versions were checked out from and the version number of each file. Only these directories record the relationship between files in the sandbox and the repository. If you move a sandbox, CVS doesn't care because you are simultaneously moving the **CVS** subdirectories.

With AccuRev, the relationship between a workspace and the repository is tracked by the AccuRev Server. No metadata is stored in the workspace itself. AccuRev tracks the client machine where each workspace resides and the pathname of its top-level directory. If you move a workspace to a different location, you must inform AccuRev of the move using the **chws** command.

Common Operations

This section lists common version-control operations, and describes how to perform them with CVS, with the AccuRev CLI, and with the AccuRev GUI.

Obtaining a copy of the source files

CVS

```
cvs checkout <module>
```

AccuRev CLI

```
accurev mkws -w <workspace-name> -b <backing-streamname> -l <workspace-location>
```

AccuRev GUI

File > New > Workspace

Placing files under version control

CVS

```
cvs add <file(s)>
```

AccuRev CLI

```
accurev add <file(s)>
```

AccuRev GUI

Select files, right-click, *Add to Depot*

Bringing others' changes into your workspace/sandbox

CVS

1. `cvs update -dP`
2. Edit any merge conflicts.

AccuRev CLI

1. `accurev update`
2. `accurev merge -o` (edit any merge conflicts for each file)

AccuRev GUI

1. Click Update button.
2. Choose Overlap search.
3. Select files, right-click, *Merge*.

Saving your changes

CVS

`cvs commit`

AccuRev CLI

`accurev keep -m`
`accurev promote -k`

AccuRev GUI

1. Choose Modified search.
2. Select files, right-click, *Keep*.
3. Choose Kept search.
4. Select files, right-click, *Promote*.

Finding the history of files

CVS

`cvs history [<file(s)>]`

AccuRev CLI

`accurev hist [<file(s)>]`
... or ...
`accurev hist -a`

AccuRev GUI

1. Select file, right-click, *History* > *Show History*.
... or ...
1. *Admin* > *Depots*
2. Right-click depot, *History*.

Finding the status of files in your workspace/sandbox

CVS

cvs status <file(s)>

AccuRev CLI

accurev stat <file(s)>
... or ...
accurev files <file(s)>

AccuRev GUI

Automatically displayed in File Browser

Removing files

CVS

1. cvs remove <file(s)>
2. cvs commit

AccuRev CLI

1. accurev defunct <file(s)>
2. accurev promote <file(s)>

AccuRev GUI

1. Select files, right-click, *Defunct*.
2. Select files, right-click, *Promote*.

Reverting changes to files

CVS

cvs unedit <file(s)>

AccuRev CLI

accurev purge <file(s)>

AccuRev GUI

Select files, right-click, *Revert to > Backed Version*.

... or ...

Select files, right-click, *Revert to > Most Recent Version*.

Moving files

CVS

1. `cp <old-name> <new-name>`
2. `cvs remove <old-name>`
3. `cvs add <new-name>`

AccuRev CLI

1. `accurev move <old-name> <new-name>`
2. `accurev promote <new-name>`

AccuRev GUI

1. Select file, right-click, *Rename*.
2. Select file, right-click, *Promote*.

Checking out files to edit

CVS

`cvs edit <file(s)>`

AccuRev CLI

not necessary; just start editing the file

AccuRev GUI

not necessary; just start editing the file with right-click, *Edit*.

Comparing versions of files

CVS

`cvs diff -r <rev1> -r <rev2> <file>`

AccuRev CLI

`accurev diff -v <rev1> -V <rev2> <file>`

AccuRev GUI

1. right-click file, *History > Browse Versions*.
2. right-click version, *Diff Against > Other Version*, click other version

Version Control of Namespace-Related Changes

AccuRev SCM includes both management of changes to the *contents* of files and changes to the *pathnames* of files and directories (folders). During the course of development — and in particular, during periodic “refactoring” of the source code base — developers may make several kinds of namespace-related changes to the pathnames of version-controlled elements:

- Changing the names of files (for example, from **framework.java** to **gizmo_arch.java**)
- Changing the names of directories (for example, from **src** to **gizmo_src**)
- Moving files and directories to different locations in the source tree (for example, moving file **commands.java** from directory **gizmo_src** to a subdirectory named **gizmo_src/lib**)

AccuRev records each change to the pathname of a file or directory element as a new version of that element. As with content changes, all such namespace-related changes originate in workspaces, and are subsequently promoted up the stream hierarchy.

Version 3.8 introduced significant improvements to the handling of namespace-related changes. These improvements make AccuRev more flexible and intuitive, and they reduce the likelihood of creating “twins” unintentionally.

Twin Elements and Stranded Elements

The improvements affect these areas of namespace-related functionality:

- **Handling of “twin” elements** — Two or more distinct elements are described as twins if they have the same pathname within a depot. (Some SCM environments use the term “evil twins”; we won’t make that judgment.) AccuRev now tracks element names consistently, provides better information about the existence of twins, and provides tools for resolving unintended twin situations.
- **Detection of “stranded” elements** — an element is stranded in a particular workspace or stream if (1) it is active in that workspace or stream, but (2) does not have a pathname in that workspace or stream. AccuRev includes significant improvements to the detection and reporting of stranded elements.

These two areas are related. If a stream contains a set of twins at a particular pathname, only one of those elements is visible at that pathname, for most purposes. The other twin(s) are stranded.

The following sections describe the changes to namespace-related functionality in detail.

Preventing Creation of Twins in Workspaces

Two checks performed by the **Add to Depot** command (CLI command: **add**) help to prevent twins from being created:

- If the user’s workspace currently contains a defunct element at the same pathname, the **Add to Depot** command is cancelled.

- If an element with the same pathname currently appears in the workspace’s parent stream — and the element does *not* have **(defunct)** status in the parent stream — the **Add to Depot** command is cancelled. If the element is **(defunct)** in the parent stream, an **Add to Depot** command succeeds.

In previous releases, the first check was performed for renamed elements as well as for defuncted elements. This check no longer takes place — see *Ability to Reuse an Element Name after a Rename Operation* below.

Reporting of Twins in Dynamic Streams

The checks described in the preceding section apply to workspaces only, not to dynamic streams. There are various ways to create twins in dynamic streams — for example:

Defunct an existing element in a workspace, and promote the change to the parent stream. Then create a new element at the same pathname (in the same workspace or in a sibling workspace), and promote the new element to the parent stream. The parent stream now contains two elements at the same pathname — one is defunct, the other is “live”.

If a set of two or more twins exists in a dynamic stream (or snapshot), the File Browser shows only the one “live” element in the folders view and in the results of a Default Group or Defunct search. (If *every* element in a set of twins has **(defunct)** status, the most recently defuncted element is deemed to be “live”.)

In this situation, all the other twins in the set have **(stranded)** and **(twin)** status. These elements are displayed by the File Browser’s Stranded search, and by the AccuRev CLI command **stat -i**:

```
> accurev stat -s tin_dvt -i
\\.\doc\new.doc e:20 tin_dvt\2 (4\2) (defunct) (member) (stranded) (twin)
```

For more information, see *Detection of All Stranded Elements, Including “Twins”* below.

Ability to Reuse an Element Name after a Rename Operation

AccuRev now offers improved support for “refactoring” operations over previous releases. Now, all element renamings and directory hierarchy overhauls can be completed and tested in the developer’s workspace, before any changes are promoted to the parent stream.

This flexibility stems from the fact that after an element **Rename** operation (CLI command: **move**), the element’s former name becomes available for reuse immediately. (Previous releases “reserved” the former name, in case the change was purged from the workspace — see *When a Purge Operation Causes an Element to Disappear* below.)

Here’s a simple refactoring example:

```
> accurev move brass.h util.h
Moving \\.\src\brass.h to \\.\src\util.h

> copy c:\temp\temp_new_brass.h .\brass.h
1 file copied

> accurev add brass.h
```



```
Added and kept element \.\src\brass.h
```

```
> accurev promote util.h  
Validating elements.  
Promoting elements.  
Promoted element \.\src\util.h
```

```
> accurev promote brass.h  
Validating elements.  
Promoting elements.  
Promoted element \.\src\brass.h
```

Note that **util.h** (formerly named **brass.h**) must be promoted first, to “free” the name **brass.h** in the parent stream. Then the new element named **brass.h** can be promoted.

When a Purge Operation Causes an Element to Disappear

The preceding section describes new flexibility for refactoring. But it does introduce a complication: what happens if you rename an element, create a new element at the same pathname, then invoke **Revert to Backed** (CLI command: **purge**) on the renamed element?

The renamed element cannot revert to its old pathname, because there’s a new element at that pathname. Accordingly, the original element simply disappears from your workspace. This element does *not* assume (**stranded**) status — the purge operation makes the element inactive in the workspace, and (**stranded**) status applies only to active elements.

Note that at this point, your workspace contains a new element at the given pathname, and the parent stream contains the original element at that pathname. Attempting to promote the new element would produce a “name already exists in parent stream” error. These steps remove the original element from the parent stream: (1) defunct the original element in the workspace, using **defunct -e**; (2) promote this change to the parent stream.

Detection of All Stranded Elements, Including “Twins”

The File Browser’s Stranded search, and equivalently, the AccuRev CLI command **stat -i** now detect all known cases of stranded files:

- Defunct elements in same workspace/stream as a “live” element at the same pathname (“twins”)
- Elements located in a directory that is, itself, stranded
- Elements located in a directory that is defunct
- Elements located in a directory that is excluded from the workspace/stream
- Elements located in a directory that has been purged from the workspace/stream
- Elements with a “pathname cycle”

Stranded files are reported with the status flag (**stranded**). If a stranded file happens to be a twin of another element, it is also reported with the status flag (**twin**).

The final case, “pathname cycle”, occurs when two sibling workspaces make contradictory changes to the depot’s directory hierarchy, then promote the changes to the common parent stream. For example, one workspace might move directory **src** under directory **util**, while another workspace moves **util** under **src**. When both the changes are promoted to the parent stream, AccuRev won’t be able to determine the correct pathname for these directories and the elements under them. The two directories assume (**stranded**) status, and the elements under these directories become inaccessible.

Ability to Operate on Stranded Elements Using Element-IDs

The CLI now provides improved tools for relieving situations involving stranded elements. The **move**, **defunct**, and **undefunct** commands now support the **-e** option, which enables you to specify an element by its element-ID, rather than by its pathname. This is necessary for situations in which the desired element does not *have* a pathname in the workspace or stream.

See *Handling Stranded Elements* on page 70.

More Sophisticated Analysis of Namespace-Related Changes

AccuRev now distinguishes between these two changes to the pathname of an element:

- Renaming of the directory in which the element resides. This is not considered a change to the element itself; it is a change to the parent-directory element.
- Moving of the element from its current directory to another directory in the depot. This *is* considered a change to the element itself (and not a change to either of the parent-directory elements).

For example, suppose that file element **commands.java** resides in directory **cmd_interface**. A colleague changes the name of the directory to **cli** in her workspace, then promotes the change to the parent stream. When you update your workspace, the pathname of the file changes from **.../cmd_interface/commands.java** to **.../cli/commands.java**, as a “side effect” of the change to the parent directory. Note that this is not a change to the **commands.java** file element itself.

On the other hand, if the colleague moves file **commands.java** to another directory, say **.../cmd_interface/utis/commands.java**, this *is* a change to the file element. When you update your workspace, the pathname of the file changes accordingly (unless you have made a namespace-related change to the file, in which case a merge is required).

AccuRev implements this scheme by tracking an element’s parent directory by its element-ID (which never changes), rather than by its name (which can change, and can vary from stream to stream).

Change to Merge Algorithm for Namespace-Related Changes

The algorithm used by the **Merge** command (both in the GUI and the CLI) now uses the more sophisticated analysis described in the preceding section. **Merge** may perform two separate namespace-related steps:

- **Element name merge** — required when the simple name of the element being merged differs in the two contributor versions.

- **Path merge** — required when the parent directory of the element being merged differs in the two contributor versions.

Often, either or both of these steps will be performed automatically by **Merge**. If only *one* of the two contributors differs from the versions' closest common ancestor, then that contributor's change is applied automatically.

(Note that **Merge** may also need to perform a third step — a content merge — for a file element.)

The following example of the CLI **merge** command involves both kinds of namespace-related changes: (1) a file's simple name has been changed by two users, **john** and **mary**; (2) each user has moved the file to different sibling directory.

```
> accurev merge file01.mary
Current element: \.\dir02\sub03\file01.mary
most recent workspace version: 4/2, merging from: 5/5
common ancestor: 5/3
```

Both "path" and "element name" conflicts must be resolved manually, but the contributors' contents can be merged automatically.

```
Path merge will be required.
Element name merge will be required.
Automatic merge of contents successful. No merge conflicts in contents.
```

```
Path conflict for \.\dir02\sub03\file01.mary
```

*John moved the file from directory sub02 to directory sub01 (which has element-ID 68).
Mary moved the file from directory sub02 to directory sub03 (which has element-ID 82).*

```
Resolve path conflict by choosing path from:
```

```
(1) common ancestor: \.\dir02\sub02\ [eid=75]
(2) backing stream : \.\dir02\sub01\ [eid=68]
(3) your workspace : \.\dir02\sub03\ [eid=82]
```

```
Actions: (1-3) (s)kip (a)bort (h)elp
```

```
action ? [3] 2
```

*John renamed the file from file01.txt to file01.john.
Mary renamed the file from file01.txt to file01.mary.*

```
Resolve name conflict by choosing name from:
```

```
(1) common ancestor: \.\dir02\sub01\file01.txt
(2) backing stream : \.\dir02\sub01\file01.john
(3) your workspace : \.\dir02\sub01\file01.mary
```

```
Actions: (1-3) (s)kip (a)bort (h)elp
```

```
action ? [3] 3
```

The content merge is performed automatically.

```
Actions: keep, edit, merge, over, diff, diffb, skip, abort, help
```

```
action ? [keep]
```

```
Moving \.\dir02\sub03\file01.mary to \.\dir02\sub01\file01.mary  
Kept element \.\dir02\sub01\file01.mary
```

Handling Stranded Elements

As described above, an AccuRev workspace or stream can contain one or more elements that are stranded. An element is stranded in a particular workspace or stream if it is a member of the default group, but cannot be accessed because there is no pathname to the element in that workspace or stream. An element can be stranded in one stream but not be stranded in other streams.

In the AccuRev GUI, stranded elements are listed in the File Browser's **Stranded** filter. In the CLI, the command **stat -i** lists stranded elements. A stranded element is listed by its element-ID, along with a pathname that was once (but is not currently) valid in that stream.

The sections below describe the ways in which elements can become stranded, along with procedures for handling each situation.

Defunct element obscured by element with same name

Elements under a defunct parent

Elements under an excluded parent

Dangling directory elements

Elements under a non-existent (purged) parent directory

Elements under a stranded parent directory

Defunct element obscured by element with same name

This occurs in the parent stream of two workspaces if:

- The user in workspace #1 defuncts an element, then promotes this change to the parent stream.
- The user in workspace #2 updates the workspace (to incorporate the defuncting), creates a new element with the same name, then promotes the new element to the parent stream.

At this point, the defuncted element is stranded in the parent stream. It cannot be promoted to the “grandparent” stream by name, because it doesn’t have a name in the parent stream. The new element cannot be promoted to the grandparent stream at all, because the name in the grandparent stream belongs to the defuncted element.

Note: through repeated **add-promote-defunct-promote** cycles, it’s possible to have multiple elements with defunct status in the parent stream, all of which were created at the same pathname.

Resolving the Situation

To get the defuncted element out of the way, promote it by element-ID to the grandparent stream:
promote -e <eid> -s <parent_stream>.

To recover the defuncted element in workspace #1, use **undefunct -e <eid>** on the defuncted element. This has the side effect of making the new element inaccessible in workspace #1. Depending on your needs, use **defunct -e** or **move -e** on the new element.

Elements under a defunct parent

This occurs in the parent stream of two workspaces if:

- The user in workspace #1 defuncts a directory element, then promotes this change to the parent stream.
- The user in workspace #2 modifies a file element in that directory, then keeps and promotes it to the parent stream.

At this point, the file element is stranded in the parent stream. In addition, the user in workspace #1 cannot access the file element by name.

Resolving the Situation

To propagate the file element’s change to the grandparent stream, promote it by element-ID:
promote -e <eid> -s <parent_stream>.

To access the file’s contents, use its element-ID: **cat -e <eid> -v <version-ID>**.

The only way to work with the file element in workspace #1 is to first **undefunct** the directory, which makes the file visible again.

Elements under an excluded parent

This occurs in the parent stream of two workspaces if:

- The user in workspace #1 modifies a file element in that directory, then keeps and promotes it to the parent stream.
- The user in workspace #2 sets a rule that excludes a directory element from the parent stream (**excl -s <parent-stream> <directory-name>**).

At this point, the file element is stranded in the parent stream. In addition, the user in either workspace cannot access the file element by name (after updating the workspace).

Resolving the Situation

To propagate the file element's change to the grandparent stream, promote it by element-ID: **promote -e <eid> -s <parent_stream>**.

To access the file's contents, use its element-ID: **cat -e <eid> -v <version-ID>**.

The only way to work with the file element in either workspace is to first remove the exclude rule (**clear** command) from the parent stream, and then update the workspace. This makes the file visible again.

Dangling directory elements

This contradictory situation — a particular directory seems to be both above and below another directory — occurs in the parent stream of two workspaces if:

- The user in workspace #1 moves directory A under directory B, then promotes directory A.
- The user in workspace #2 moves directory B under directory A, then promotes directory B.

At this point, both directories are stranded in the parent stream. An update of workspace #1 causes directory B to be removed; an update of workspace #2 causes directory A to be removed.

Resolving the Situation

The only way to untangle this knot of inconsistency is to checkout (**co** command) a previous version of each directory that has the “correct” (that is, consistent with the other directory) pathname, then promote these old versions to the parent stream.

The simplest way to do this is to specify the transaction that created the directory at its correct pathname: **co -t <add-transaction-number>**. But this method can be “messy” if the **add** transaction also created other elements, such as the files within the directory.

Another method is to use a workspace under a time-based stream to see the relevant directories with their correct pathnames. Checkout the “old” directory versions, promote these versions from the workspace to the time-based stream, then use **promote -s <time-based-stream> -S <parent-stream>** to promote to the parent stream.

Note: with either method, you'll probably need to use the **-O** option to the **promote** command, in order to avoid the need to merge the “old” directory versions.

Elements under a non-existent (purged) parent directory

This occurs in the parent stream of a workspace if:

- The user creates a new directory and file within the new directory, and promotes both new elements to the parent stream.
- The user purges (GUI: **Revert to Backed**) the new directory from the parent stream.

At this point, the new file is stranded in the parent stream.

Resolving the Situation

You cannot propagate the file element's change to the grandparent stream, because the new directory never existed in that stream.

To access the file's contents, use its element-ID: **cat -e <eid> -v <version-ID>**.

The only way to work with the file element is to first checkout (**co** command) the version of the directory that was originally created in the workspace. The simplest way to do this is to specify the transaction that created the directory: **co -t <add-transaction-number>**. But this method can be “messy” if the **add** transaction also created other elements.

Another method is to use a workspace under a time-based stream to see the directory before it was purged from the parent stream. Checkout the directory, promote it from the workspace to the time-based stream, then use **promote -s <time-based-stream> -S <parent-stream>** to promote to the parent stream.

Note: with either method, you'll probably need to use the **-O** option to the **promote** command, in order to avoid the need to merge the “old” directory version.

Elements under a stranded parent directory

To access an element under a stranded parent directory, restore the accessibility of the parent, as described in the sections above. This restores the accessibility of the element in question.

Notes on Cross-Links

This note clarifies and supplements the basic documentation of AccuRev's cross-link feature, introduced in Version 4.5.

Cross-Link Direction and Terminology

A cross-link is created in a workspace by the **Include from Stream** command (CLI: **incl -b**). The command name implies that a connection is being established *from* a specified backing stream *to* the workspace. But an existing cross-link is listed by the CLI command **lsrules** like this:

```
xlink <pathname> from <workspace> to <backing-stream>
```

That is, the direction of the cross-link “arrow” is the opposite of the direction implied by the “include from” command name. When describing a cross-link, we use this terminology:

- The workspace (or stream) where the cross-link has been created is the cross-link's source stream.
- The designated backing stream is the cross-link's target stream.

In CLI messages, “cross-link” is abbreviated to “xlink”.

Cross-Links and Stream Namespaces

Each AccuRev stream (including snapshot streams and workspace streams) provides a namespace: a set of pathnames to some or all of the depot's elements. For example:

```
\\.doc
\\.src
\\.tools
\\.doc\chap01.doc
\\.doc\chap02.doc
\\.src\commands.c
\\.src\topaz.c
\\.src\topaz.h
\\.tools\cmdshell
\\.tools\perl
\\.tools\python
\\.tools\tools.readme
\\.tools\cmdshell\bash
\\.tools\cmdshell\csh
\\.tools\cmdshell\bash\end.sh
\\.tools\cmdshell\bash\start.sh
\\.tools\cmdshell\csh\end.csh
\\.tools\cmdshell\csh\start.csh
\\.tools\perl\add_cr.pl
\\.tools\perl\remove_cr.pl
```

```
\\.tools.python.setup.py
\\.tools.python.vars.py
```

Since this set of depot-relative pathnames defines a hierarchy, it's often clearer to list the pathnames component-by-component, like this:

```
\\.
  doc
    chap01.doc
    chap02.doc
  src
    commands.c
    topaz.c
    topaz.h
  tools
    tools.readme
    ...
```

To locate an element, AccuRev interprets its specified pathname component-by-component (just like the operating system does). The cross-links facility provides a way to make AccuRev switch namespaces in the middle of the pathname-interpretation process.


Note:

For example, consider this pathname:

```
\\.tools.cmdshell.csh.end.csh
```

And suppose you've created a cross-link at subdirectory **cmdshell**, with workspace **W** as the source stream and stream **S** as the target stream. AccuRev will process the pathname, component-by-component, as illustrated here:

- Pathname components up to *and including* the cross-linked component, are interpreted in the original (source stream) namespace.
- Additional pathname components, if any, are interpreted in the new (target stream) namespace.



Note that in workspace **W**, you continue to access the cross-linked element, subdirectory **cmdshell**, through its "local" name in the workspace's namespace. It's quite possible (but you don't need to know) that this element has a different name — even a different pathname — in the target stream:

```
\\.tools.shell_scripts
\\.tools.common/scripts\
\\.scripting
... etc.
```

Pathname components below “cmdshell” are interpreted in the namespace of stream **S**, the target stream. For example, if script **end.csh** has been renamed in stream **S** to **topaz_exit.csh**, then that’s the name you must use in workspace **W**, as well:

```
\\.tools\cmdshell\csh\topaz_exit.csh
```

The File Browser and the CLI commands **stat** and **files** make this namespace-switching transparent: AccuRev shows you the element names and pathnames that will enable you to access the data from your current workspace or stream context.

Source Stream: Workspace vs. Dynamic Stream

The example in the preceding section uses a workspace as the “source stream”. The same pathname-interpretation principles apply if the source stream is a dynamic stream.

But the basic difference between workspace streams and dynamic streams affects the way cross-links work in them:

- In a dynamic stream, the **Include from Stream** command incorporates all changes from the target stream immediately. This reflects the fact that a dynamic stream inherits versions from its backing stream automatically and instantly.
- In a workspace, the **Include from Stream** command respects the workspace’s update level. That is, it incorporates only those changes that occurred in the target stream before the workspace’s most recent update. A subsequent **Update** command will bring in the more recent changes from the target stream.

Example: to see how cross-links work with a workspace’s update level, suppose that the following changes have been made in stream **topaz_mnt**:

- directory element **\\.tools\cmdshell\cmd** has been **Defunct**’ed
- directory element **\\.tools\cmdshell\csh** has been renamed to **\\.tools\cmdshell\c_shell**
- file element **\\.tools\cmdshell\c_shell\start.csh** has been edited

You use the **Include from Stream** command to create a cross-link from your workspace to stream **topaz_mnt**, at pathname **\\.tools\cmdshell**. The immediate change to your workspace depends on its update level:

- If the changes in stream **topaz_mnt** occurred after your workspace’s most recent update, you won’t see the changes immediately in your workspace: directory **cmd** will still exist, directory **csh** won’t be renamed to **c_shell**, and you won’t see the edits to file **start.csh**. But the status of these elements includes the **(stale)** indicator, showing that the changes are in the backing stream, waiting to be incorporated:

```
\\.tools\cmdshell\cmd topaz\1 (9\1) (backed) (xlinked) (stale)
\\.tools\cmdshell\csh topaz\1 (9\1) (backed) (xlinked) (stale)
\\.tools\cmdshell\csh\start.csh topaz\2 (9\4) (backed) (xlinked) (stale)
```

At this point, performing an **Update** will bring the changes into the workspace.

- If the changes in stream **topaz_mnt** occurred before your workspace’s most recent update, all those changes will be brought into the workspace immediately.

The procedure in the first bulleted paragraph can be described as “Include then Update”; the second bulleted paragraph’s case can be described as “Update then Include”. The final result is the same in both cases: the changes to the cross-linked elements in their new backing stream are incorporated into your workspace. We consider the second case to be an AccuRev best practice:

Best Practice:

Update your workspace before performing an **Include from Stream** command

If you **Update** first, other backing-stream changes won’t be “mixed in” with the **Include from Stream** changes during the next workspace update. Moreover, fully establishing the link from your workspace to the target stream will involve a single step (Include), rather than two steps (Include then Update).

Note: because it respects — but does not change — your workspace’s update level, **Include from Stream** more closely resembles the **Populate** command than the **Update** command.

Multiple Cross-Links: Chaining

AccuRev can traverse two or more cross-links in the same pathname. For example, you might use this pathname in workspace **W**:


```
./aaa/bbb/ccd/DDD/eee/fff/GGG/hhh/foo.java
```

And suppose there are two cross-links:

- At subdirectory **DDD**, a cross-link from workspace **W** to stream **S1**
- At subdirectory **GGG**, a cross-link from stream **S1** to stream **S2**

As AccuRev traverses the pathname component-by-component, it interprets the components as illustrated here. As it progresses down the pathname, AccuRev also traverses a “chain” of cross-links:

- start in workspace **W**, then ...
- cross-link to stream **S1**, then ...
- cross-link to stream **S2**



“Chaining” of cross-links can continue to any number of levels. The same principle applies repeatedly: a cross-linked pathname component is interpreted in the *source* stream’s namespace; subsequent non-cross-linked components are interpreted in the target stream’s namespace.

But you must take care when “chaining” cross-links in this way. It is possible to create ambiguous configurations, which AccuRev handles by removing the affected elements. See [Cross-Link Overlaps](#) on page 80.

A special case of cross-link chaining occurs when you create a configuration in which two or more cross-links occur at the *same* pathname component. For example, consider again this pathname:

```
\\.tools\cmdshell\csh\end.csh
```

And suppose there is a chain of two cross-links at the same pathname component:

- At subdirectory **cmdshell**, a cross-link from workspace **W** to stream **S1**
- At subdirectory **cmdshell**, a cross-link from stream **S1** to stream **S2**

In workspace **W**, the subdirectory will continue to have its “original” name, **cmdshell**. But the subtree under the subdirectory will come from the stream **S2** namespace. By extension, you could chain any number of cross-links at the **cmdshell** component: **W** > **S1** > **S2** > **S3** > **S4** ... As above, the directory retains its “original” name in the workspace, and the workspace sees the directory’s subtree as it exists in the final target stream.

Double Vision: Seeing an Element Multiple Times in a Workspace


One consequence of AccuRev’s cross-link facility is that two (or more) different versions of the same element can appear at different pathnames in the same workspace or stream. We call this phenomenon double vision. This is not an error — at least, not from AccuRev’s perspective. Seeing the same element twice might be exactly what you intended, or it might signify that you’ve left some refactoring work unfinished.

Here’s an example: suppose you are tasked with doing some cleanup on the Topaz project’s development tree:

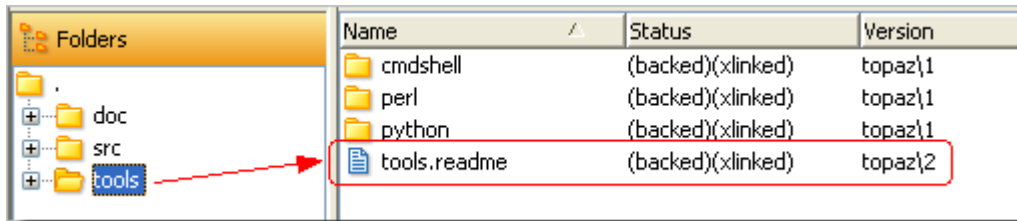
- Flatten out the subdirectories under **tools**.
- Move file **tools.readme** to the depot’s root directory, and rename it to **README-tools.txt**.
- Improve the source file comments in the **src** directory.

You perform this work in your workspace, named **topaz_refact**. But when the dust settles, you find that the programs in the **tools** subdirectory no longer work. You are not sure whether the problem is in the **tools** directory or the **src** directory. So you decide to “back out” your refactoring of the **tools** directory, by cross-linking to the known-to-work version of the **tools** directory in snapshot stream **topaz_2.3.9**.

Now, you have two different versions of the “README” element in your workspace! In your refactoring, you created a new version in your workspace, at pathname **\\.README-tools.txt**:



But your workspace now cross-links to the Release 2.3.9 version of the **tools** subdirectory, which contains the Release 2.3.9 version of the same element, at pathname `\\.tools\\tools.readme`:



Name	Status	Version
cmdshell	(backed)(xlinked)	topaz\1
perl	(backed)(xlinked)	topaz\1
python	(backed)(xlinked)	topaz\1
tools.readme	(backed)(xlinked)	topaz\2

This case of double-vision is clearly an error, reflecting the fact that your refactoring work is still ongoing. In other cases, you might want two (or more) versions of a commonly used source file, say **topaz.h**, to appear in a workspace. Perhaps several different versions of the file are required, in order to build different executables using that file. Version skew is the executables' other dependencies might mandate the different versions of **topaz.h**.

Double Vision and the 'accurev name' Command

The **accurev name** command lists the pathname for a given element (specified by element-ID) in your workspace. It can also list the pathname for a specific version of an element, or the version in a specific stream:

```
accurev name -e 28
accurev name -v topaz_mnt -e 116
```

In a double vision situation, the **name** command can list all of an element's pathnames in a workspace or stream:

```
> accurev name -e 28 -v topaz_refact
\\.tools\\tools.readme
\\.README-tools.txt
```

Cross-Link Overlaps

Section *Multiple Cross-Links: Chaining* on page 78 describes how a set of cross-links can define a "chain" of backing streams to be used at different components in a pathname:

```
/. /aaa/bbb/ccccc/DDD/eee/fff/GGG/hhh/foo.java
```

start interpreting
pathname in
workspace W

switch from
workspace W
to **stream S1**


switch from
stream S1
to **stream S2**

Chaining works correctly if each switch to the next link in the chain occurs at the same pathname component or at a lower component. But here's a situation that violates this rule:

`/. /aaa/bbb/ccccc/DDD/eee/fff/GGG/hhh/foo.java`
 |
 start interpreting
 pathname in
workspace W
 |
 switch from
stream S1
 to **stream S2**
 |
 switch from
workspace W
 to **stream S1**

In this case, the second link in the cross-link chain (**S1** > **S2**) occurs at a *higher* pathname component, **DDD**, than the first link (**W** > **S1**, at component **GGG**). AccuRev recognizes this situation as a cross-link overlap.

When a workspace that has a cross-link overlap gets updated, AccuRev removes the subtree below the component where the first link was created.



Notes on Promote-by-Issue

This technical note describes AccuRev's **promote-by-issue** feature, both reviewing the basic concepts and presenting several usage scenarios. This information supplements and updates the basic documentation of promote-by-issue in the AccuRev GUI's help screens. The most important points are:

- In some cases, promote-by-issue requires you to create an additional version of one or more elements, using the **Patch** or **Merge** command. You must create a new AccuWork issue, called a tracking issue, to enable AccuRev to account for the additional version(s).
- For certain usage scenarios, you must use the Change Palette in a way that might differ from your previous usage procedure.


Promote-by-Issue Basics

(This section paraphrases the information in section “Viewing a Stream's Current Development Activity” of the “Stream Browser” help topic for the AccuRev GUI.)

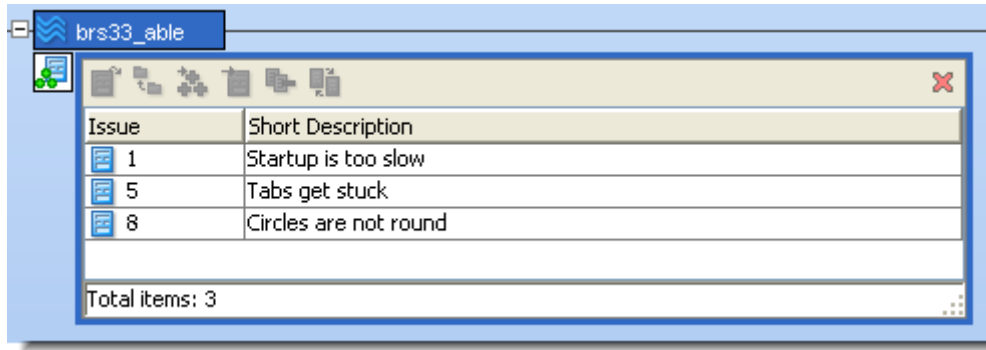
The Stream Browser can show the development activity currently taking place in each stream or workspace. A control below the stream or workspace opens or closes a subwindow that displays the details of the development activity. The activity details can appear in several ways — by element, by transaction, or by issue record.



Use the development-activity mode controls at the right side of the Stream Browser toolbar to determine how the activity details will be displayed. (The icons on the controls below the streams and workspace change accordingly.) You can change modes either before or after opening a development-activity subwindow.



When displaying a stream's activity by issue record, the subwindow displays the issue records that are in a particular stream:



(This display is the same as the results of a **Show Active Issues** or **Stream Diff by Issues** command. In this note, all the examples use the Stream Browser's development activity subwindow, but you can perform promote-by-issue operations in these other contexts, also.)

Promote-by-issue involves selecting one or more issue records and performing any of the following operations:

- Invoking the **Promote** command from the subwindow's toolbar or the context menu of the issue(s).
- Drag-and-drop'ing the issue record(s) onto another stream or workspace in the Stream Browser display. The destination can be the parent (backing) stream, another dynamic stream, or one of your workspaces. This invokes a **Promote** or **Send to Workspace** operation on all the elements in the selected issue record(s).

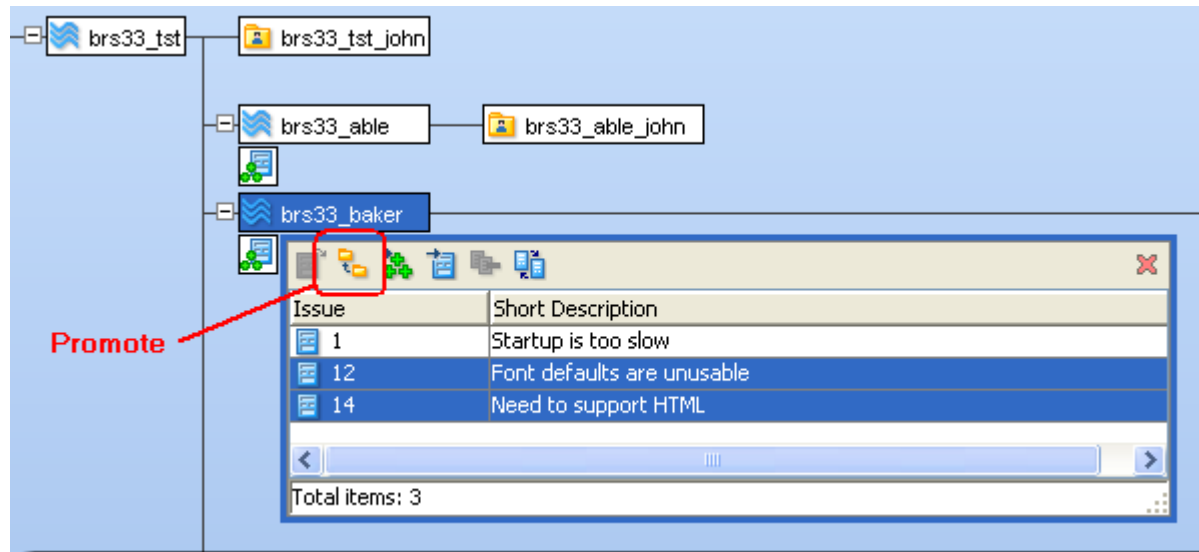
You can also drag-and-drop the subwindow control to another stream. This performs the **Promote** or **Send to Workspace** operation on all the issue records in the subwindow (even if only some of them are currently selected).

The following sections explore several promote-by-issue scenarios.


Note: in the remainder of this technical note, we shorten the term "issue record" to "issue".

Promoting Issues to the Parent Stream

In many cases, promote-by-issue is simple and easy. You promote one or more issues to the parent stream ...




... and the issues simply move into the parent stream:



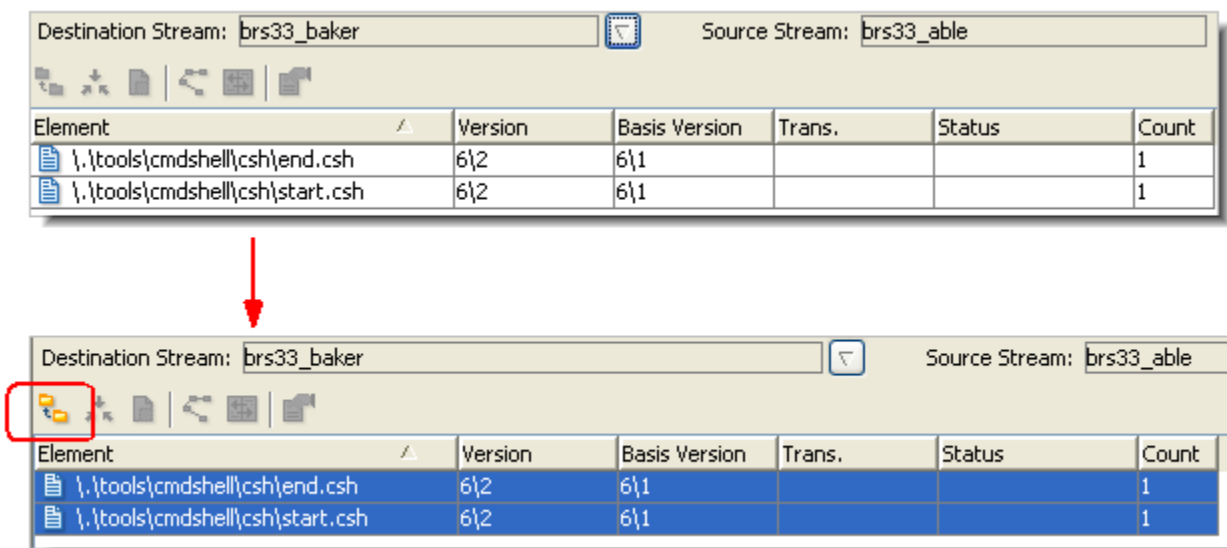
Note: if the **Promote** command fails with a “merge required” error, you must perform a merge for one or more elements in the source workspace (or a workspace below the source stream). When promoting the merged version(s), assign them to the same issues as the original version(s).

Cross-Promoting Issues to a Non-Parent Stream — Simple Case

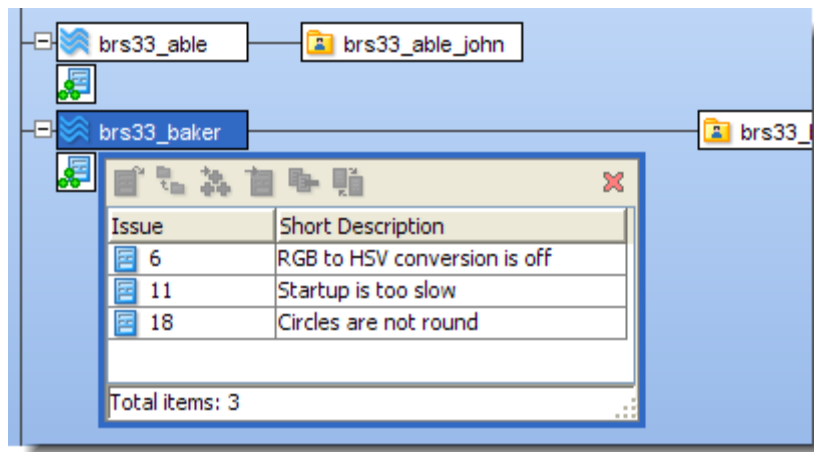
Similarly, in many cases cross-promoting from one dynamic stream to another dynamic stream proceeds without complication. First, you perform a drag-and-drop operation:



This populates the Change Palette with the versions in the selected issues. If the **Promote** command is enabled when you select all the elements, you can proceed to propagate all the issues (and their versions) to the destination stream:




(In this case, issue #11 consists of element **start.csh**, and issue #18 consists of element **end.sh**.)
Back in the Stream Browser, the promoted issues now appear in the destination stream:



Cross-Promoting Issues to a Non-Parent Stream — Patch Required

In this scenario, your intention is the same as in the preceding one — to cross-promote the versions in one or more issues from one dynamic stream to another. You start the same way, with a drag-and-drop operation from stream **brs34_able** to stream **brs34_baker**:




As before the elements in the issue(s)' change packages are loaded into the Change Palette. But in this scenario, one or more of the elements is listed with **(patch)** status:

Destination Stream:

Source Stream:

Element	Version	Basis Version	Trans.	Status	Count
\.src\brs34.c	6\8	6\1			1
\.tools\cmdshell\bash\start.sh	6\9	6\6		(patch)	1
\.tools\tools.readme	6\5	6\1			1

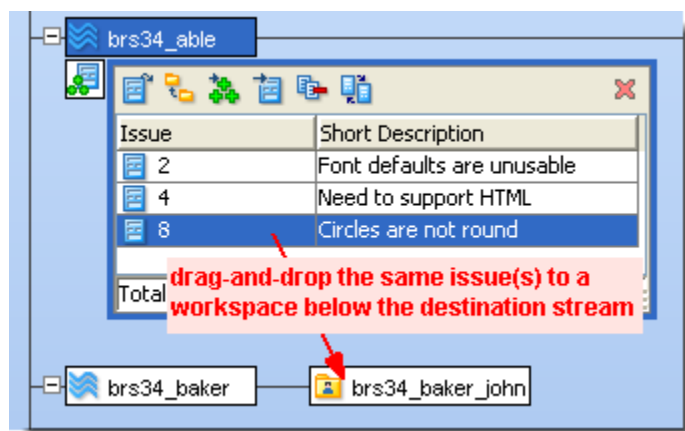
This occurs when the element's change-package entry does not contain the complete set of changes to the element that (1) are in the source stream and (2) have not yet been promoted to the destination stream:



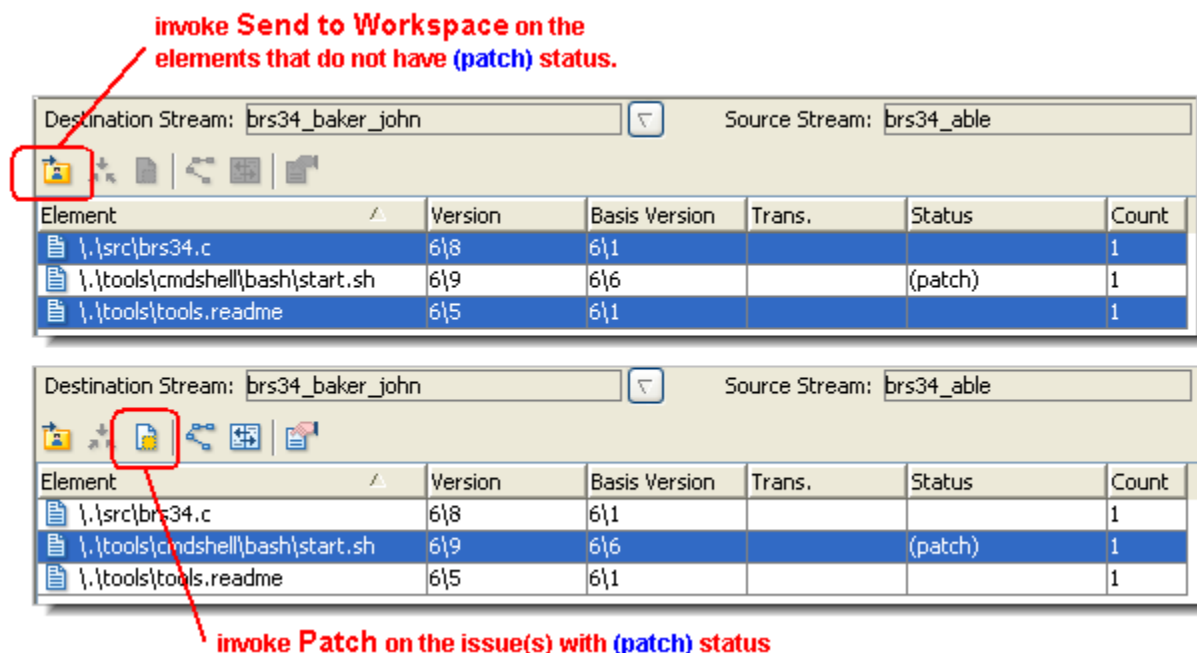
Changing the Way You Use the Change Palette

In the situation described in the preceding section, you might be accustomed to dealing with the **(patch)**-status element separately from the other elements. But do not proceed in this manner! Instead, close the Change Palette tab without propagating any of the changes to the destination stream.

Now, populate another instance of the Change Palette with a drag-and-drop operation of the same issue(s) from the source stream to a workspace below the destination stream:




The new Change Palette display is similar to that of the previous instance. Process the elements as indicated below:

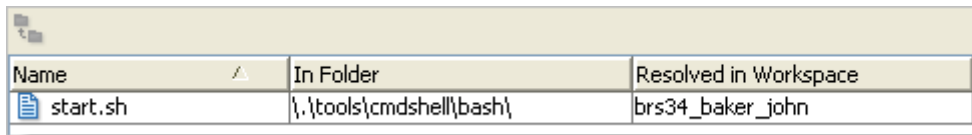


What about (overlap) status? If you have made changes to one or more of these elements in this workspace, some of the elements will have **(overlap)** status. For such elements, you can use **Merge** instead of **Send to Workspace**.

If an element has both **(patch)** and **(overlap)** status, use the **Patch** command, not the **Merge** command. AccuRev warns you if you attempt to merge. In this case, click **Cancel** and invoke the **Patch** command instead.

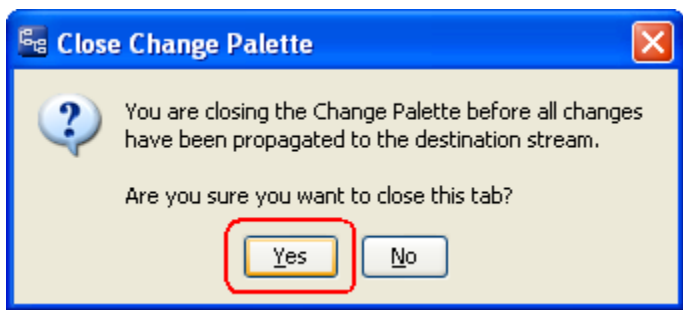


After you have processed all the elements in the issue(s), each Patch'ed element is listed at the bottom of the Change Palette:

A screenshot of a table within a software interface, likely the Change Palette. The table has three columns: "Name", "In Folder", and "Resolved in Workspace".

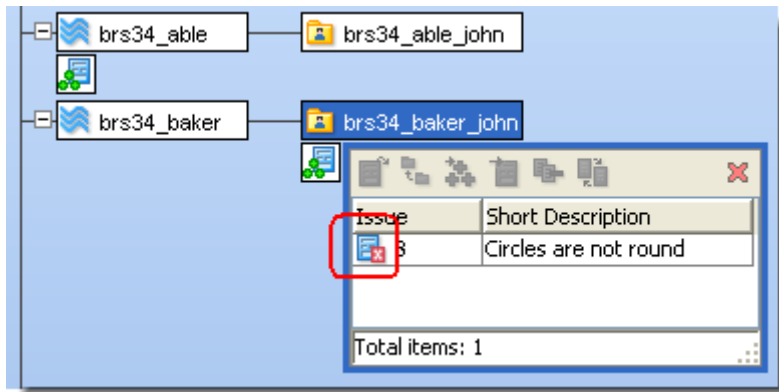
Name	In Folder	Resolved in Workspace
start.sh	\\.tools\cmdshell\bash\	brs34_baker_john

Although you might be accustomed to promoting Patch'ed and Merge'd elements from the Change Palette, do not proceed in this manner. Instead, close the Change Palette tab without invoking **Promote**.



Working in the Workspace

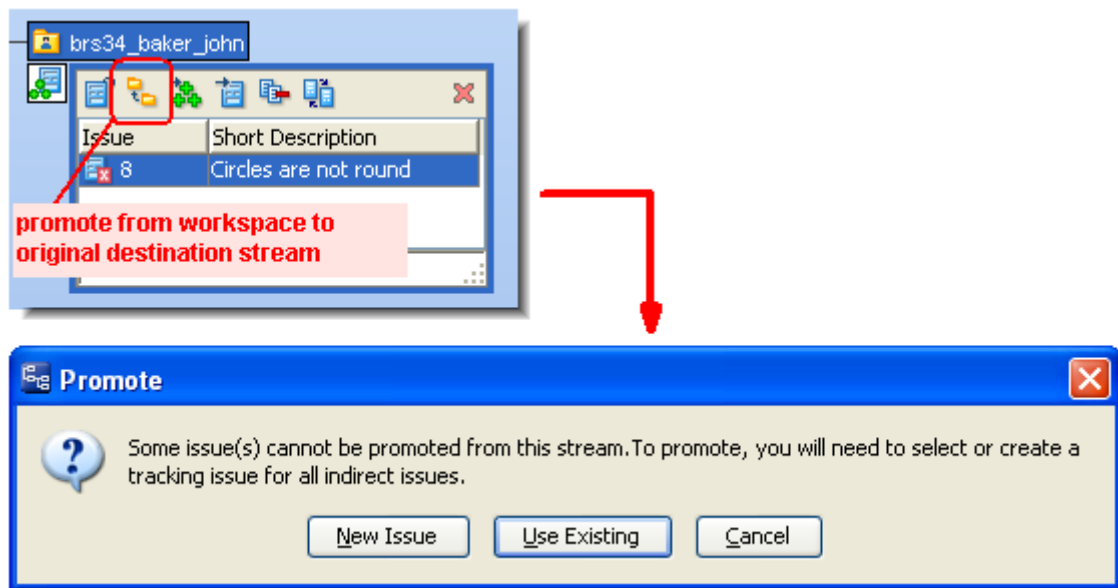
At this point, the selected issue(s) have been propagated to the workspace below the destination stream. Each issue is listed in the workspace's development activity subwindow with a special variant of the issue icon:



This icon indicates that AccuRev will track changes made in this workspace to the issue(s)' elements. For each Patch'ed element, the new version created by the **Patch** command is one such change. You can create any number of additional versions of both the Patch'ed elements and the elements you processed with **Send to Workspace**.

Promotion / Creating a Tracking Issue

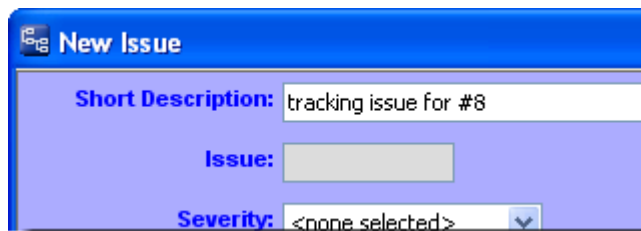
When your work on the issue is finished, you must use promote-by-issue to send your changes to the workspace's backing stream (the original destination stream of the cross-promote). Don't use promote-by-element or promote-by-transaction, which would defeat AccuRev's ability to track changes related to the selected issue(s). When you invoke the **Promote** command, AccuRev prompts you to specify an additional issue, called a tracking issue. This is the mechanism that AccuRev uses to record the changes made in this workspace to the original issue(s)' elements.



A single tracking issue can keep track of the additional changes for any number of original issues. Thus, it makes sense to select **New Issue** the first time you need to specify a tracking issue, and select **Use Existing** on subsequent uses of promote-by-issue in the same workspace. You might

also need to select **Use Existing** if you're not allowed to create new issues — for example, if you've integrated a third-party issue-tracking system with AccuWork.

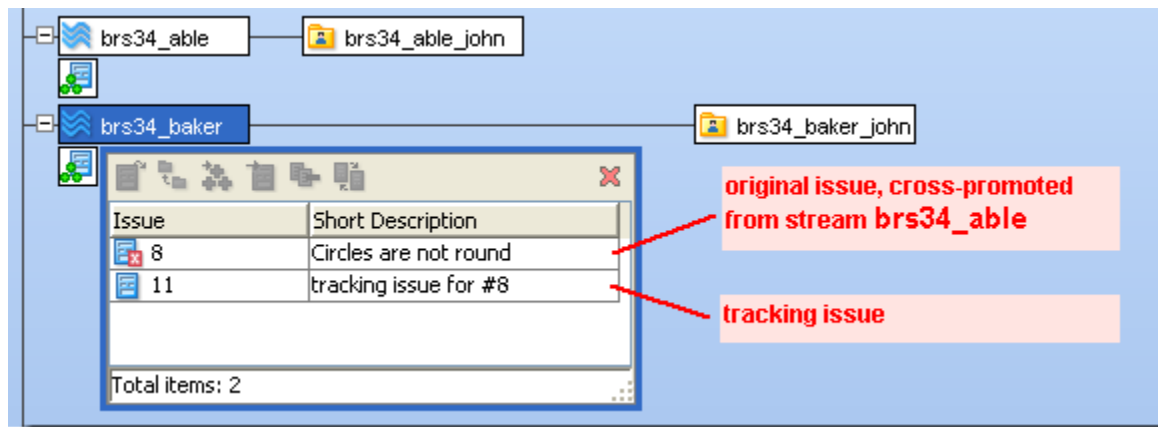
Either way, an AccuWork edit form appears, in which fill in the fields of the new or existing tracking issue. As usual, click the edit form's **Save** button when you are finished filling in the fields.



AccuRev automatically proceeds with your original command, **Promote**: its prompts you for a comment, then sends the versions from the workspace to the original destination stream. The cross-promotion of the original issue(s) is now complete — after just a bit of a detour!

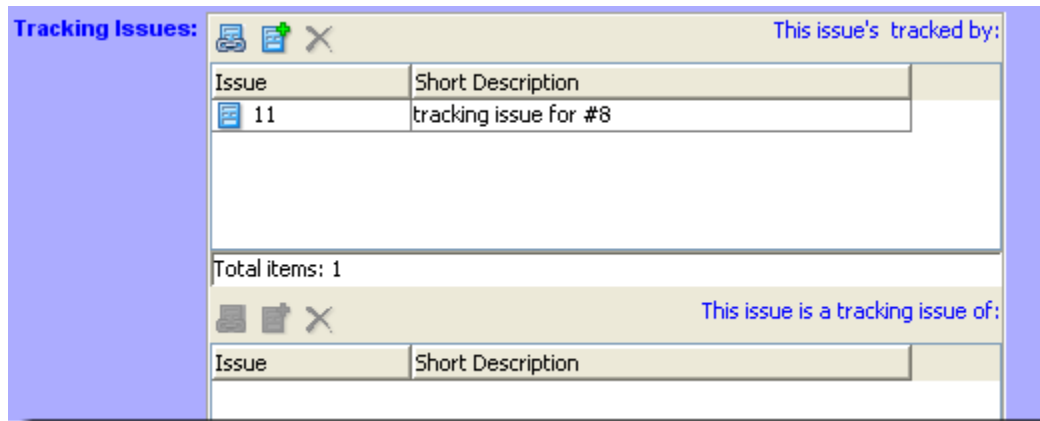
Working with the Tracking Issue

The Stream Browser shows that the original issue (in our example, #8) has been promoted from the source stream (**brs34_able**) to the destination stream (**brs34_baker**). The tracking issue (#11) records the additional changes to issue #8's elements that have been propagated to the destination stream.



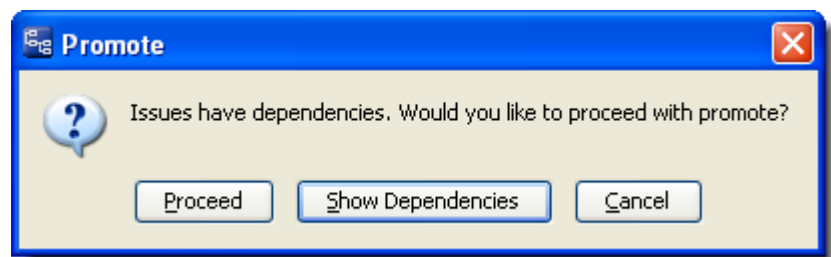
AccuRev Version 4.6 introduced change package dependency tracking, whose goal is to ensure that a **Promote** operation sends a self-contained, consistent set of changes to the destination stream. A tracking issue and the corresponding original issue(s) are connected by a change package dependency: the tracking issue depends on the original issue(s).

The best way to monitor such connections is with a Relationship field whose type is **Track**. For example, when issue #8 is viewed in an edit form, its connection to issue #11 might be displayed like this:



As with all change package dependencies, AccuRev warns you if you attempt to promote a tracking issue without its dependencies, the original issue(s). In this situation, don't click **Proceed**! It's important always to promote both the

original issue(s) and the tracking issue at the same time. So the correct procedure is to **Cancel**, select both the original and tracking issues, then invoke **Promote** again.



As you propagate issues up (or across) the stream hierarchy, you must continue to obey this same rule:

Always promote original issues and the corresponding tracking issues at the same time.

If you attempt to **Promote** an original issue alone, without including its tracking issue, AccuRev refuses:

