

AccuRev[®]

Administrator's Guide

**Version 5.4.1
February 2012**

Revised 23-February-2012

Copyright

Copyright © AccuRev, Inc. 1995–2012

ALL RIGHTS RESERVED

This product incorporates technology that may be covered by one or more of the following patents:
U.S. Patent Numbers: 7,437,722; 7,614,038.

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

AccuBridge, **AccuReplica**, **AccuWork**, **Kando**, 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

The AccuRev Repository	1
Repository Access Permissions	1
READ ME NOW: Assuring the Integrity of the AccuRev Repository	1
Backing Up the Repository	3
Restoring the Repository	4
Archiving Portions of the Repository	5
Moving a Workspace or Reference Tree	5
Moving a Depot	5
Removing a Depot	6
A Word of Caution on Windows Zip Utilities	6
Storage Layout	6
The AccuRev Server	9
Operating-System User Identity of the Server Processes	9
UNIX/Linux Identity	9
Windows Identity	9
AccuRev User Identity of the Server Process	10
Starting the AccuRev Server	11
Running the Server Automatically at Operating System Startup	11
Starting the Server Manually	12
Server Configuration File	12
UNIX/Linux: Controlling the Server's Operating-System User Identity	13
Controlling Login Session Longevity	13
Controlling Multithreading of the AccuRev Server	14
Configuring Access to the Database	14
Server Logging	14
Time zone offset	15
Logging Levels	15
UNIX/Linux: Log File Rotation	15
Controlling Server Log Verbosity	15
Verbose Server Logging	15
Server Watchdog	17
Watchdog Logging	18
Controlling Server Operation	18
UNIX/Linux: 'acserverctl' Utility	18
Windows: 'Services' Control Panel	20
Server-Control Files	20
Open Filehandle Limits and the AccuRev Server	21
Changing the Per-Process Open File Descriptor Limit	22
AccuRev Server Performance	24

AccuRev Architecture Overview	24
Key Factors Affecting AccuRev Performance	25
Resource Utilization and the Client-Server Architecture	25
AccuRev Server Machine Usage: Software.....	26
AccuRev Server Machine Networking.....	26
Using Multiple AccuRev Servers	26
Setting Up Client Machines	27
Workspaces and Servers	28
Specifying a Server By Using the –H Option.....	28
Configuring the Web User Interface (Web UI).....	28
System Clock Synchronization	31
Detecting System Clock Discrepancies -- Timewarp	31
AccuRev’s Timewarp Detection Scheme	31
Fixing System Clock Discrepancies	32
Automatic, Gradual Convergence of System Clocks	32
Manual Synchronization Tools.....	33
Archiving of Version Container Files	35
The ‘archive’ Command	36
Determining Which Versions to Archive	36
Dry Run Capability.....	36
Archiving the Versions	36
The ‘reclaim’ Command.....	37
Attempts to Access Archived Versions.....	38
Using ‘hist’ to Research Previous ‘archive’ Commands	38
Restoring Archived Versions -- The unarchive Command.....	38
Replication of the AccuRev Repository	41
Master and Replica	41
AccuRev Licensing in a Replication Environment.....	43
Replica Licensing Using the Traditional (Named) keys.txt Manager	43
Replica Licensing Using the Standard/Flexible accurev.lic Manager	43
Installing and Configuring AccuRev in a Replication Environment.....	43
Configure an AccuRev Server as a master server.....	43
Configure a Replica server	44
Establish an AccuRev User Identity for the AccuRev Server Process	44
Synchronize the Site Slice	45
Indicate the Depots to be Replicated	45
Enabling Compression.....	46
Setting Up a Client Machine to Use a Replica Server	46
Using a Replica Server	47
The Update Command	47
Command Interaction in a Replicated Environment.....	47

Removing Storage Containers on a Replica Server	48
Removing a Replica Server	48
Improving Replica Performance	48
Triggers and Replication	49
Creating New Depots	50
Adding and Removing Depots from a Replica Repository	50
Synchronizing a Replica Manually	50
On-Demand Downloading of a Version's Storage File	51
Automating Replica Synchronization	51
Synchronization Security	52
The replica_site.xml File	52
Moving the AccuRev Server and Repository to Another Machine	53
Procedure for Moving the Repository	53
On the Source Machine	53
On the Destination Machine	54
AccuRev Security Overview	55
Users and Groups	55
User Authentication	55
The "accurev_login" User-Authentication Method	56
The "custom" User-Authentication Method	56
Authenticating a Replica User on the Master	56
The 'server_auth_trig' Script	57
Selecting the User-Authentication Method	58
How AccuRev Records the User-Authentication Method	58
Restriction on Access to the "Add User" Command	58
Locks on Streams	58
Access Control List (ACL) Permissions	59
Element-Level Security (EACLs)	59
Features	60
Basic Terms	60
Important Concepts	60
Permissions	61
Inheritance	62
Access Denied	62
Installing AccuRev with Element Security for the First Time	63
Auditing	63
Setting Permissions for a Replica Server	63
EACL Usage Scenarios	64
Restricting Access to Commands using Triggers	69
Which Security Feature Should I Use?	69

AccuRev Triggers	71
Pre-Operation Triggers	71
Client-Side Triggers.....	71
Server-Side Triggers.....	71
Post-Operation Triggers.....	73
Triggers in a Replication Environment.....	73
Triggers and Security	73
Transaction-Level Integration Trigger.....	73
Preparing to Use an AccuRev-Provided Trigger Script.....	74
Enabling a Trigger	74
pre-create-trig, pre-keep-trig, pre-promote-trig, server-post-promote-trig.....	75
server_admin_trig	75
server_preop_trig	75
server_dispatch_post.....	75
Notes on Triggers in Multiple-Platform Environments	76
The Trigger Parameters File	76
Format of the “pre-create-trig” Trigger Parameters File	77
Format of the “pre-keep-trig” Trigger Parameters File	79
Format of the “pre-promote-trig” Trigger Parameters File.....	81
Format of the “server-post-promote-trig” Trigger Parameters File.....	82
Format of the “server_preop_trig” Trigger Parameters File.....	83
Format of the “server_admin_trig” Trigger Parameters File.....	86
Format of the “server_dispatch_post” Trigger Parameters File	87
Encoding of Element Lists.....	88
Encoding of Command Comments.....	89
Trigger Script Contents.....	89
Trigger Script Exit Status	90
Trigger Script Execution and User Identities.....	91
‘Administrative Users’ in Trigger Scripts.....	91
The Trigger Log File.....	91
Integrations Between AccuRev and AccuWork	92
Change-Package-Level Integration.....	92
Transaction-Level Integration.....	95
Implementation and Customization of the Transaction-Level Integration.....	96
If Both Integrations are Enabled.....	97
Disabling Triggers.....	97
The ‘maintain’ Utility	99
Specifying a Database Admin Username and Password	99
‘maintain’ Command Reference	99
Backup/Restore of the AccuRev Repository	103
Removing a Depot from the AccuRev Repository	103
Before You Begin	103
Depot Removal Procedure	103

Reusing a Depot's Name	104
License Management	105
License Types.....	105
To Install and Enable Licenses	105
Types of Standard/Flexible Licenses	106
Replication Server Licenses.....	107
Standard Licenses vs. Flexible Licenses.....	107
How Flexible Licenses Work.....	108
License Files, Options Files, and Configuration File	109
ISV options and accurev.opt.....	110
Multiple AccuRev Servers.....	110
Configuration Example.....	110
License Manager Administration Page.....	111
Additional Documentation.....	111

The AccuRev Repository

The AccuRev Server program manages a data repository, which provides long-term storage for your organization's development data — for example, all versions of all source files. The repository consists of:

- a database called **accurev**, which contains:
 - a **site** schema, which contains the user registry, list of depots, list of workspaces, and other repository-wide information.
 - a schema for each depot, each of which contains depot-specific metadata and AccuWork issue data.
- the **site_slice** directory, which contains repository-wide AccuWork data, workflow configuration data, server preferences, and triggers
- the **depots** directory, which contains a set of subdirectories, each storing an individual depot. A depot subdirectory stores one or both of:
 - a version-controlled directory tree: all the versions of a set of files and directories.
 - AccuWork schema, query, and change package configuration data.

When it starts, the AccuRev Server program determines the location of the **site_slice** and **depots** directories by looking at the settings in configuration file **acservr.cnf**. This file must reside in the same directory as the Server program (**accurev_server**) itself. See *Server Configuration File* on page 12 for more information.

Repository Access Permissions

The operating-system user identity of the AccuRev Server process must have full access to all the files and directories within the **storage** directory. For maximum security, this should be the *only* user identity with permission to access the repository.

If you create an **acadmin** AccuRev administrator account, as suggested in *Operating-System User Identity of the Server Processes* on page 9, this user identity must also have access to the **bin** directory where the AccuRev executables are stored.

READ ME NOW: Assuring the Integrity of the AccuRev Repository

The integrity of the AccuRev repository is critically important. If information in the repository is lost or corrupted, your organization's ability to do business may be severely compromised. The integrity of the repository relies on the integrity of underlying software (the database software and the file system, including the device drivers for data storage devices) and underlying hardware (the data storage devices themselves). Certain practices will enhance the safety and reliability of these underlying facilities. We strongly recommend that you:

- Use high-quality disk drives and disk controllers.
- Reduce the impact of a hard-disk failure by using disk mirroring (for example, using a RAID system) or other fault-tolerant disk subsystems.
- Power the AccuRev server machine with an uninterruptible power supply (UPS), with automatic shutdown of the server machine if the UPS is running out of power. This reduces the likelihood of interrupted data transfers to disk.
- Establish a good data-backup regimen, and make sure your backups are reliable by doing test restores on a regular basis. (See *Restoring the Repository* on page 4.)

This section focuses on one aspect of data integrity: guaranteeing “write” operations to the repository. The AccuRev Server process does not, itself, perform the act of writing data on the disk. Like all application programs, it makes a “write” request to the operating system (UNIX/Linux, Windows). In turn, the operating system performs a “write” operation to the disk itself. (On some larger systems, there may be additional links in this chain of write operations.)

Operating systems and disk subsystems often use special techniques that boost the performance of write operations, but can compromise data integrity. For example, when an application program makes a write request, the operating system might:

- Acknowledge the request immediately — good, because the application program can then proceed to its next operation.
- Delay actually sending the data to the disk (“write-behind”) — bad, because a system failure at this point might result in the data never being stored on the disk.

It is essential that such techniques *not* be used when the AccuRev Server process sends information to the disk containing the AccuRev repository. The Server always follows each write request with a “synchronize the disk” request. Sometimes, this ensures that data is safely on disk before the Server proceeds to its next task. For example, this is typically the case if the repository is stored on a disk that is local to the machine on which the Server is executing.

But in some situations delayed-write techniques may be used even when the AccuRev Server makes “synchronize the disk” requests. This is typically the case if the repository is located on a network shared file system. In such situations, the Server's “synchronize the disk” requests are effectively ignored, so that successful completion of write operations to the AccuRev repository cannot be guaranteed. (Some disk subsystems implement such a guarantee by having their own battery backup; buffered data is flushed to disk when the power fails.)

In an attempt to avoid such unsafe situations, the AccuRev Server process attempts to determine whether the file system where the repository is stored guarantees the successful completion of write operations. If it decides “no”, the Server refuses to use the repository. This determination is not foolproof — both “false positives” and “false negatives” are possible.

If you have any question about the safety of your data-storage system, please contact AccuRev Support Services.

Backing Up the Repository

The **backup** command makes a copy of the AccuRev metadata. During **backup** command execution, clients can continue to work. Only transactions that are complete at the time the **backup** command is invoked are included in the metadata backup. The backup may take a few seconds or a few minutes, depending on the amount of metadata on your system. After running the backup, you can make a complete copy of the file storage area (the **storage** directory tree), without worrying about synchronization or time-skew.

CAUTION: Do not execute the **backup** command while you are copying the file storage area. This can place incorrect data into the backup copy of the repository.

With the metadata backup and the copy of the file storage area, you can restore the repository to its state at the time you executed the **backup** command.

Before You Begin

- Read the section titled *A Word of Caution on Windows Zip Utilities*, below.
- Note that the live-backup scheme relies on the ability to copy files that are currently in use by the AccuRev Server process. To support this scheme, your backup/restore tool must be able to copy files that are currently “open” at the operating system level.
- Your backup/restore tool should also have these features:
 - Ability to preserve files’ timestamps.
 - Ability to preserve files’ ownership and execute permissions.
 - Ability to back up zero-length files. (See *Server-Control Files* on page 19.)

If you have any questions, contact AccuRev Support Services.

Backup Procedure

1. Back up the AccuRev metadata:

```
accurev backup <backup-file-spec>
```

By default, the backup is stored in the **storage\site_slice\backup** directory.

2. If your backup utility cannot copy files that are currently open at the operating system level, stop the AccuRev Server. (See *Controlling Server Operation* on page 17.)
3. Use a backup/restore tool to create a backup copy of the entire directory tree below the **storage** directory, except for the **db** directory. This backup can be all-at-once or piecemeal; for example, you can back up the **site_slice** directory and the subdirectories within the **depots** directory separately.

Note: if your site slice is in a non-standard location (as specified by the `SITE_SLICE_LOC` setting — see *Server Configuration File* on page 12), or if some depots are in non-standard locations (perhaps moved with the **chslice** command), then your job in backing up the entire repository is more complicated than simply to copy the **storage** directory.

4. If you have scripts or triggers that are crucial to your usage of AccuRev, back those up as well.
5. If you stopped the AccuRev Server in Step 2, start it again. (See *Controlling Server Operation* on page 17.)

Backing up a Replica

The procedure described above can also be used to back up a replica. If you are backing up a replica, you will also need to make a copy of the **acserver.cnf** file from the **bin** directory and save it with the backed-up data.

Out of Shared Memory Error

If you have a large number of depots, you may encounter error messages similar to the following:

```
pg_dump: WARNING: out of shared memory
pg_dump: attempt to lock table <table name> failed: ERROR: out of
shared memory
HINT: You may need to increase max_locks_per_transaction.
```

You can address this issue in the following ways:

- perform a more selective backup using "-p" to back up only *some* depots
- increase "max_shared_locks_per_transaction" as suggested by the error message

To increase "max_shared_locks_per_transaction", edit the value of this variable in postgresql.conf using the following formula:

```
20 x (number_of_depots + 1) / max_connections
```

NOTE: Supporting the backup of more than 1,000 depots may also require that you increase the value of SHMMAX to more than 32MB.

Restoring the Repository

If you have backed up the repository according to the directions above, you can easily restore the repository to the time at which you executed the **backup** command:

1. Stop the AccuRev Server. (See *Controlling Server Operation* on page 17.)
2. Restore the backup copy of the **storage** directory, using the backup/restore tool that you used to create it.

Note: if your site slice is in a non-standard location (as specified by the SITE_SLICE_LOC setting in the **acserver.cnf** file — see *Server Configuration File* on page 12), or if some depots are in non-standard locations (perhaps moved with the **chslic** command), then your job in restoring the backup of the entire repository is more complicated than simply to restore the **storage** directory.

3. Restore any backups of scripts or triggers, using the backup/restore tool that you used to create them.

Steps 2 and 3 will restore the backed-up data to standard locations on an existing AccuRev system. For more complex restore operations, see *Moving the AccuRev Server and Repository to Another Machine* on page 53.

4. Restore the backed-up metadata:

```
maintain restore <backup-file-spec> <db-admin>
```

The *maintain* command will prompt you for the database admin password.

Note: the AccuRev Database Server must be running for this command to succeed.

5. Restart the AccuRev Server. (See *Controlling Server Operation* on page 17.)

Note: suppose a particular depot's files were not backed up for several hours after the *backup* command was executed. Even if several new versions of file **gizmo.c** were created with the *keep* command during that interval, the backed-up AccuRev metadata will have no record of those transactions. But you can still retrieve a copy of any or all of those versions from the backup medium: it's in a container file in the **data** subdirectory of the depot directory.

Restoring a Replica

The procedure described above can also be used to restore a replica. After restoring the metadata, refer to the backed-up **acservr.cnf** file to ensure that the MASTER_SERVER, PORT, LOCAL_SERVER, and LOCAL_PORT settings are correct for your replication environment.

Archiving Portions of the Repository

The container files that store the contents of individual file versions can be moved to offline storage, in order to save online storage space for the repository. For details, see *Archiving of Version Container Files* on page 35.

Moving a Workspace or Reference Tree

Note: before you start, consult *A Word of Caution on Windows Zip Utilities* below.

First, make sure that no user or script process is currently using the workspace or reference tree. Move the physical contents of the workspace tree or reference tree with a backup/restore tool (e.g. *tar*, *zip*, *xcopy* /s). Then, let AccuRev know about the move:

```
accurev chws -w <workspace-name> -l <new-location>
```

```
accurev chref -r <reftree-name> -l <new-location>
```

Moving a Depot

Note: before you start, consult *A Word of Caution on Windows Zip Utilities* below.

First, make sure that no user or script process is currently using the depot. (To guarantee this, you may wish to stop the AccuRev Server.) Move the physical contents of the depot with a backup/restore tool (e.g. *tar*, *zip*, *xcopy* /s). Then, let AccuRev know about the move:

```
accurev chslice -s <slice-number> -l <new-location>
```

(Use *accurev show depots* to determine the slice number of the depot.)

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 103.

A Word of Caution on Windows Zip Utilities

Be careful when using WinZip[®] or PKZIP[®] on a Windows machine to move a workspace, reference tree, or depot. You may want to use *tar* on a UNIX/Linux machine to “pack up” a directory tree, and then use the Zip utility on a Windows machine to “unpack” it.

- When moving the entire repository or an individual depot, be sure to disable conversion of line-terminators during the “unpack” step:
 - In WinZip, make sure the option “TAR file smart CR/LF conversion” is not selected (*Options > Configuration > Miscellaneous*).
 - In PKZIP, make sure the “CR/LF conversion” setting is “None -- No conversion” (*Options > Extract*).

Enabling conversion of line-terminators during the “unpack” step will corrupt the text files in a depot’s file storage area (see [File Storage Area](#) below). The AccuRev Server always expects lines in these text files to be terminated with a single LF character, no matter what kind of machine the server is running on.

- Conversely, when moving a workspace or reference tree, you may wish to enable “TAR file smart CR/LF conversion”. The files in a workspace or reference tree are handled directly by text-editors, compilers, testing tools, etc. Many Windows text-editors are incapable of handling text files whose lines are terminated with a single LF character.
- Zip utilities typically refuse to copy files that are open at the operating system level. Typically, you can work around this limitation by stopping the *accurev_server* program, but this defeats AccuRev’s “live backup” feature.

Storage Layout

Each AccuRev depot is stored in a separate directory tree under the installation area’s **storage** directory. The **storage** directory is a sibling of the executables (“bin”) directory. An example directory layout would be:

```
<ac-install>  
  bin
```

```
storage
  db
  depots
    project1
    project2
    project3
  site_slice
```

A depot consists of three parts:

Configuration Files

The *mktrig* command creates a one-line configuration file that names the script to be executed when the trigger fires for transactions involving this particular depot. For example, making a trigger of type “pre-keep-trig” creates a configuration file in the depot named **pre-keep-trig**. (This file might contain the pathname `/usr/local/bin/accurev_prekeep.pl`.)

Database

The database stores all the versions of a set of files and directories, as well as AccuWork schema, query, and change package configuration data.

File Storage Area

Whenever a user creates a new real version of a file with the *keep* command, the AccuRev Server copies the file from the user’s workspace to the depot’s file storage area. The newly created storage file is permanently associated with the real version-ID in the workspace stream (e.g. 25/13), and also with subsequently created virtual version-IDs in higher-level streams (7/4, 3/9, 1/4).

Storage files are located in subdirectory tree **data** within the depot directory. The files may be in compressed or uncompressed form. Compressed files may correspond to more than one real version. Conceptually, storage files are numbered sequentially starting with 1. Within the **data** directory, they’re arranged in a hierarchy for faster access. For example, storage file #123456 would be stored as **data/12/34/56.sto**.

You can relocate a depot’s file storage area onto other disk partitions or even onto remote disks. Exercise extreme caution when relocating storage in this area. Make sure you have first done a full backup and have shut down the *accurev_server* program.

The AccuRev Server

The AccuRev repository is managed by a single program, the AccuRev Server (*accurev_server*). This program communicates with the AccuRev Database Server. Both of these must be started prior to running any AccuRev client commands. No user should attempt to work directly with the repository, unless it is an emergency. In this case, please contact AccuRev Support Services.

Operating-System User Identity of the Server Processes

Like all processes, the AccuRev Server and AccuRev Database Server processes have an operating-system user identity. It should be a unique user identity, not used by any other program. This helps to ensure that no other user or process has access to the repository.

UNIX/Linux Identity

CAUTION: The AccuRev Database Server cannot be run as the **root** user. In addition, do not attempt to run the AccuRev Server as **root**. Some user-supplied trigger scripts run under the operating-system identity of the AccuRev Server, which poses a significant security risk. (See *Trigger Script Execution and User Identities* on page 92.)

We suggest that you create an operating-system user named **acserver**, belonging to a group named **acgroup**. (Any similar names will do.) Only the AccuRev Server should run as **acserver**.

For emergency “manual” access to the repository, you can create another user identity — say, **acadmin** — and place that user in the same group, **acgroup**. You can configure UNIX/Linux-level auditing and place other appropriate controls on this account; this leaves the **acserver** account (and thus, the AccuRev Server process) unencumbered by such controls.

Configure the AccuRev Server to run with the **acserver/acgroup** identity by placing these names in the server configuration file, **acserver.cnf**. See *UNIX/Linux: Controlling the Server’s Operating-System User Identity* on page 13.

Windows Identity

The AccuRev Server and AccuRev Database Server run as Windows services. By default, these services run as the built-in local user named **System**. This user identity must have access to the AccuRev executables (**bin**) directory and to the repository. See *Repository Access Permissions* on page 1.

You can use the **Services** control panel to configure the services to run under another identity (“account”).



AccuRev User Identity of the Server Process

In addition to its user identity at the operating system level, the AccuRev Server process sometimes needs an AccuRev username (principal-name) identity:

- When it executes a server-side trigger script that invokes AccuRev client commands, such as *annotate* or *promote*.
- When it performs a synchronization with the master repository — explicit or implicit *replica sync* command. This applies only if the AccuRev Server is managing a replica repository.

If either of these situations applies to the AccuRev Server that you are administering, you must take steps to establish a valid AccuRev username for the AccuRev Server. The AccuRev username need not be special or reserved. Just make sure that any security controls — ACL permissions and/or *server_admin_trig* script — are configured to allow that particular AccuRev username to perform the required operations. See [AccuRev Security Overview](#) on page 55.

Note: for security reasons, we recommend that the operating-system identity of the AccuRev Server process (for example, **acserver**) should *not* also be an AccuRev username.

With the “AccuRev login” user-authentication scheme, a session file establishes the AccuRev user identity of a process. Create a long-lived session file for the AccuRev Server’s operating-system identity as follows:

1. Access the command line.
2. Set environment variable `ACCUREV_HOME` to the home directory of the operating system user that the AccuRev Server runs as (for example, **acserver**).

If the AccuRev Server is running as the local **System** account on a Windows machine, the home directory is `C:\`.

3. Create a long-lived session file for the AccuRev username that the AccuRev Server will use:

```
accurev login -n john
Password: *****
```

Note: if you are creating a session file on a replica server machine, to be used for communicating with the AccuRev Server process on the master server machine, direct the *login* command to the master server machine. For example:

```
accurev login -n -H bingo_master:5050 john
Password: *****
```

This session file will be valid indefinitely, thanks to the *-n* option.

Starting the AccuRev Server

The following sections describe how to start the AccuRev Server, either automatically at operating system bootstrap time, or manually at a command prompt. (You can also perform a “manual” startup with a UNIX/Linux shell script or a Windows batch file.)

Running the Server Automatically at Operating System Startup

Typically, the Server program is started automatically when the operating system boots on the server machine. On UNIX/Linux systems, an “rc” or “init.d” startup script starts the *accurev_server* program. The AccuRev installation program does not install the startup script automatically. You must customize and install the sample startup script, named *accurev*, located in the **extras/unix** subdirectory of the AccuRev installation directory. See the **README** file in that subdirectory for complete instructions.

Note: Red Hat Linux will not run an rc.d script at shutdown unless it finds a lock file named [/var/lock/subsys/accurev](#). Without this lock file, the “acserverctl stop” command will not run, and the log file will not record a clean shutdown.

On Windows systems, the AccuRev installation program automatically configures the *accurev_server.exe* program as a Windows service. Use the *Services* control panel to stop, start, and configure the Server program.

Starting the Server Manually

The AccuRev Server may also be started manually:

- **Windows systems:** If you've changed the startup type of the **AccuRev** service to "Manual", you can start the service from the **Services** control panel. Alternatively, run the **server_start.bat** script, located in the AccuRev executables (**bin**) directory.
- **UNIX/Linux systems:** Start the Server with the **acserverctl** utility:

```
<AccuRev-executables-dir>/acserverctl start
```

The Server will run with your operating-system user identity. (Make sure that the server configuration file's USER and GROUP settings are commented out. See [UNIX/Linux: Controlling the Server's Operating-System User Identity](#) on page 13.)

Server Configuration File

When it starts, the AccuRev Server program reads configuration file **acserver.cnf**, located in the AccuRev executables directory. This configuration file is generated during installation, but can be edited manually thereafter.

*Important! After editing this file, you must restart the AccuRev server. If you have edited the values for SITE_SLICE_LOC, MASTER_SERVER, or PORT, you **must** also run the **maintain server_properties update** command before the new values will take effect. See [The 'maintain' Utility](#) on page 99 for more information.*

Here are some sample **acserver.cnf** settings:

```
MASTER_SERVER = accurev_server_machine.company.com
PORT = 5050
SITE_SLICE_LOC = /partition0/site_slice
DEPOTS_DEFAULT = C:\Program Files\AccuRev\storage\depots
```

Important! The white space surrounding the equals sign (=) in configuration files is mandatory.

The MASTER_SERVER name should be the fully-qualified name of the server machine, including a domain name and Internet extension. Using just the server name may work in most situations, but fully-qualified is preferred. Alternatively, you can use the IP address of the server machine.

The PORT setting contains the port that the AccuRev Server is available on.

The SITE_SLICE_LOC setting points to the directory that the Server uses for storing repository-wide AccuWork data, workflow configuration data, server preferences, and triggers. This directory:

- Must be owned by the operating-system account that the Server runs as (for example, **acserver**).

- Must be physically located on the server machine. The `SITE_SLICE_LOC` location must not be within a remotely mounted file system (UNIX/Linux) or within a shared directory (Windows) on a remote machine.

The `DEPOTS_DEFAULT` setting points to the directory that the Server uses for new depots, for storing the files created when a permanent copy of a version is created with the *keep* command. Note that if you move an existing depot, this variable is not used — you must tell AccuRev about the new location with *chslice*.

UNIX/Linux: Controlling the Server’s Operating-System User Identity

The following specifications determine the user identity and group membership of the operating system process in which the AccuRev Server runs:

```
USER = <user-name>
GROUP = <group-name>
```

When the AccuRev Server is started automatically (usually at system boot), it changes its user identity and group membership according to the `USER` and `GROUP` settings in the `acserver.cnf` file. (See *Controlling Server Operation* on page 17.)

With these settings commented out, the AccuRev Server runs under the identity of the user who started it.

UNIX/Linux: Setting the Server’s Home Directory

In addition to having a user identity and group membership, the AccuRev Server has an AccuRev home directory. This directory is used for a variety of purposes — for example, to store a login session file created by a *server_admin_trig* trigger script.

By default, the AccuRev home directory is the same as the operating-system home directory, as indicated by the environment variable `HOME`. It’s a best practice to override the `HOME` value by setting the value of environment variable `ACCUREV_HOME`. If the AccuRev Server is started automatically at system startup time by a script in the “rc” or “init.d” directory, the most logical place to set the AccuRev home directory is in this startup script:

```
export ACCUREV_HOME=/users/acserver
```

Controlling Login Session Longevity

A successful user login creates a session that by default expires 4 hours (240 minutes) after the last AccuRev command is executed. You can change this behavior by creating or modifying this `acserver.cnf` setting:

```
SESSION_TIMEOUT = <number-of-minutes>
```

On UNIX/Linux systems, a user can control the timeout for an individual session by setting environment variable `SESSION_TIMEOUT` before logging in. For example, to set a 15-minute timeout interval for a single session:

```
export SESSION_TIMEOUT=15
accurev login derek
```

Non-expiring login sessions

The following setting in **acservr.cnf** causes user sessions never to expire:

```
SESSION_TIMEOUT = 0
```

No matter what the setting, users can create non-expiring sessions with *login -n*.

Controlling Multithreading of the AccuRev Server

The AccuRev Server is a multi-threaded program, architected to support a maximum of 256 concurrent threads. To conserve system resources, you can specify a lower maximum in the **acservr.cnf** file:

```
MAX_THREADS = 25
```

As it's running, the AccuRev Server may reduce the maximum even further than the specified `MAX_THREADS` level, depending on the available computing resources.

Configuring Access to the Database

The AccuRev Server (as of Version 5.0) uses the third-party database PostgreSQL. The settings listed below are configured at installation.

```
DB_DRIVER = Postgres
DB_CONNECT = localhost:5075@accurev
DB_USER = acserver
DB_PASS = e8c5ed8f07bfaf44d2f2eebc215e3cf3
DB_HOME = C:\Program Files\AccuRev\postgresql
```

- `DB_DRIVER` – Identifies the third-party database used by AccuRev. This setting should not be changed.
- `DB_CONNECT` – Identifies the database server, port, and name used by AccuRev. This setting should only be changed if you change the port that the AccuRev database server uses.
- `DB_USER` – Identifies the database user AccuRev uses to connect to the third-party database. This setting can be changed using the *maintain setcnf* command.
- `DB_PASS` – Identifies the (encrypted) password for `DB_USER`. This setting can be changed using the *maintain setcnf* command.
- `DB_HOME` – Identifies the location of the binaries and libraries used by the third-party database. This setting is used by the *accurev backup* command, and should not be changed.

Server Logging

The AccuRev Server maintains a log file, **acservr.log**, in subdirectory **logs** of the **site_slice** directory. Each log message includes a timestamp, the AccuRev username that invoked a client command, and the IP address of the client machine.

Time zone offset

As of AccuRev V5.4, the format of **acsrvr.log** entries has been expanded to include a time zone offset. For example:

```
2011/12/20 14:33:41.329-05:00 Server locale is C
```

where "-05:00" is the ISO-8601 TZ offset with respect to GMT.

This enables you or AccuRev Support to correlate logs from masters, replicas, and clients spread across different timezones.

Logging Levels

Logging information can be preserved at various levels of detail, as specified in **acsrvr.cnf**:

```
# log level 2 or 3 is recommended by AccuRev support team
LOG_LEVEL = 2
```

UNIX/Linux: Log File Rotation

On UNIX/Linux server machines, log file rotation keeps the log file from growing too large. Periodically, the AccuRev Server timestamps the current log file and moves it to subdirectory **logs** of the **site_slice** directory. For example, the log file might be renamed **acsrvr-2002-01-23-04-47-29.log**. The Server then creates a new **acsrvr.log** file. The log file is rotated weekly; it is also rotated whenever the AccuRev Server is restarted.

Controlling Server Log Verbosity

The verbosity of the server logs is controlled by the LOG_LEVEL entry in the **acsrvr.cnf** file:

```
LOG_LEVEL = 3  #enable the highest (most verbose) level of logging
```

At log level 1, each execution thread produces one line in the log. At log level 2, each execution thread can produce multiple log lines. Log level 3 essentially combines levels 1 and 2. At higher log levels, some of the messages detail the work of server subtasks.

Verbose Server Logging

The Server is a multi-threaded program, so it can handle several client commands concurrently. A typical AccuRev client command causes the AccuRev Server to execute a set of server subtasks. For each client command, the Server's "master thread" creates a new "worker thread" to perform the set of subtasks for that particular command. When the worker thread has performed all the subtasks, it exits. When the LOG_LEVEL is 2 or 3, the log messages indicate many of the details of server subtask execution.

For example, a single **update** command can generate a set of log messages like this:

```
2008/02/09 16:24 connection 1076 on 00000E98 cache 0 started
2008/02/09 16:24 1076 mary *update 00000E98 1.2.3.101
2008/02/09 16:24 1076 mary cur_wspace 00000E98 1.2.3.101
2008/02/09 16:24 1076 mary ws_type 00000E98 1.2.3.101
```

```

2008/02/09 16:24 1076 mary      stream_top    00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      check_time   00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      update       00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      end          00000E98 1.2.3.101
2008/02/09 16:24 connection 1076 on 00000E98 success 0.563 0 0 0 1.2.3.101 mary

```

These messages may or may not appear on consecutive lines of the log file. If multiple client commands are being executed concurrently by different worker threads, the log messages that the threads generate will be interleaved in the log file.

Let's examine each message in the above example:

```
2008/02/09 16:24 connection 1076 on 00000E98 cache 0 started
```

The first message is generated at the time (2008/02/09 16:24:20) a client request is accepted by the Server's master thread. This is `connection 1076` between the client and the server). The master thread creates a new worker thread (worker thread-ID `00000E98`) and hands the request off to it.

```
2008/02/09 16:24:20 1076 mary      *update      00000E98 1.2.3.101
```

This message indicates the user who invoked the command (`mary`), the name of the command, marked with an asterisk (`*update`), and the IP address of the client machine (`1.2.3.101`).

```

2008/02/09 16:24 1076 mary      cur_wspace   00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      ws_type      00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      stream_top   00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      check_time   00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      update       00000E98 1.2.3.101
2008/02/09 16:24 1076 mary      end          00000E98 1.2.3.101

```

Each time the worker thread begins a particular subtask, it sends a message to the log. In the lines shown above, the client command is implemented through server subtasks `cur_wspace`, `ws_type`, `stream_top`, `check_time`, `update`, and `end`. (The last subtask is always named `end`.) Each message also includes the username, thread-ID, and client IP address.

```
... connection 1076 on 00000E98 success 0.563 0 0 0 1.2.3.101 mary
```

The last message is generated by the worker thread after it has completed all subtasks and is about to exit. (If the `LOG_LEVEL` is 1, this is the *only* message generated for each client command.) In addition to the username, thread-ID, and client IP address data also included in the preceding messages, this message reports summary measures, listed in boldface above:

- **success / failure** (`success` in the example above): The overall result of the attempt to execute the client command.
- **run time** (`0.563`): the total time, in seconds, that the worker thread took to process the entire client command.
- **last check** (`0`): the time, in seconds, elapsed since last progress update from worker thread. In a success message, this value is 0. In a failure message, this value is non-zero.
- **time delta** (`0`): the time difference between the clocks on the client and server machines.

- *exit status* (0): the exit code for thread: 0 = success, non-zero = error.

Server Watchdog

The AccuRev Server is designed for high reliability, to ensure the integrity of the repository and its availability to AccuRev users. But even the most robust software systems are occasionally compromised; the AccuRev Server can be brought down by a bad disk sector or an administrator's mistaken command.

The reliability of the AccuRev Server is further enhanced by a companion program, called the "Watchdog", which runs on the same machine. The sole function of the Watchdog is to monitor the Server and restart it in the event of a failure. The effect of the Watchdog on Server performance is insignificant.

Note: both the Server and Watchdog show up in the operating system's process table with the same name: *accurev_server*.

Every 10 seconds, the Watchdog sends a simple command to the Server. If the Watchdog detects that the Server is not responding or is not functioning properly, the Watchdog restarts the Server. If the Watchdog detects five such failures within a three-minute timespan, it doesn't restart the Server; such a situation indicates the need for server reconfiguration or investigation by the AccuRev support team. (If `ACCUREV_WATCHDOG_FAST_FAIL_DISABLE` is set in the Watchdog's environment, it keeps trying to restart the Server indefinitely.)

For the most part, the functioning of the Watchdog process is transparent, making administration simple:

- The Watchdog process starts automatically when the Server process is started (typically, at operating system bootstrap time).
- The administrative commands for stopping the Server process cause both the Watchdog and Server to stop. These commands have been reworked to terminate the Watchdog directly; before it exits, the Watchdog terminates the Server.

Tools that control the execution of the Server and Watchdog are in described in section *Controlling Server Operation* on page 17.

Watchdog Logging

The Watchdog maintains a simple log file, **acwatchdog.log**, in subdirectory **logs** of the **site_slice** directory. On UNIX/Linux server machines, the Watchdog log file is rotated similarly to the Server log file.

Controlling Server Operation

AccuRev includes facilities for controlling the operation of the AccuRev Server, the Watchdog, and the AccuRev Database Server. The user interface varies by platform:

- UNIX/Linux: the *acservectl* command-line utility
- Windows: the *Services* control panel

UNIX/Linux: ‘*acservctl*’ Utility

If the AccuRev Server and the AccuRev Database Server are running on a UNIX/Linux machine, you can control its operation with the *acservctl* script. This is a Bourne-shell script, located in the AccuRev **bin** directory. (It is based on the control script for the Apache Web server.)

The *acservctl* script also controls the AccuRev Database Server. The AccuRev Server process is dependent upon the Database Server process; if it is not running the AccuRev Server will not start.

Note: by default, *acservctl* assumes that AccuRev is installed at **/opt/accurev**. If this is not the case, you must run *acservctl* in an environment where `ACCUREV_BIN` is set to the pathname of the AccuRev **bin** directory. For example:

```
env ACCUREV_BIN=/var/accurev/bin acservctl ...
```

acservctl provides a set of non-interactive commands. The format of each command is:

```
acservctl <command-name>
```

(Omitting <command-name> is equivalent to executing *acservctl help*.) The commands are:

dbstart

Start the Database Server process.

dbstop

Tell the Server, the Watchdog, and the Database Server processes to stop gracefully.

start

Start the Database Server (if it is not running), then start the Server and Watchdog processes.

stop

Tell the Server and Watchdog processes to stop gracefully.

dbstatus

Report whether the Database Server is running or not.

status

Report whether the Server is running or not.

pause

Tell the Server to stop accepting new requests from AccuRev clients.

resume

Tell the Server to start accepting new requests from AccuRev clients again.

dbrestart

Stop the Server and Watchdog processes, then restart the Database Server, the Server, and the Watchdog processes.

restart

Tell the Server process to stop gracefully; this allows the Watchdog to restart it. If the Watchdog is not running, a *start* or *hardrestart* is performed.

kill

Forcibly stop the Server and Watchdog processes. This is accomplished by sending a TERM signal to each process. The script gets the process-IDs from files **acserver.pid** and **acwatchdog.pid**, located in the **site_slice** directory. These files are written automatically when the processes are started.

hardrestart

Perform a *kill*, followed by a *start*.

help

Display an *acserverctl* command summary.

The various “tell a process” capabilities are implemented through server-control files. (See *Server-Control Files* below.)

Windows: ‘Services’ Control Panel

If the AccuRev Server and AccuRev Database Server are running as services on a Windows machine, you can control their operation from the *Services* control panel. In some versions of Windows, this control panel is in a subdirectory called *Administrative Tools*.

The context (right-click) menu available for the services includes these commands:

start
stop
pause
resume
restart

For descriptions of these commands, see *UNIX/Linux: ‘acserverctl’ Utility* above. The *restart* command brings down both the Server and the Watchdog, by performing a *stop* followed by a *start*. For the AccuRev Database Server, the *start*, *stop*, and *restart* commands perform the equivalent of the *dbstart*, *dbstop*, and *dbrestart* commands in *acserverctl*.

Server-Control Files

On all platforms, the AccuRev Server and Watchdog processes check, once per second, for the existence of several server-control files in the **site_slice** directory. The existence of the server-control file causes the process to perform a particular action. In most cases, the contents of the file are irrelevant; a zero-length file will do.

acserver.command.pause

(used by the *pause* server-control command) Tells the Server to stop accepting new requests from AccuRev clients. The Server completes transactions that are already in progress and logs

its “paused” status to the log file. Then, it continues to run, but the only thing it does is monitor the **acserver.command.pause** file. When this server-control file is removed, the Server resumes normal operation.

This server-control file is *not* removed when a new Server process starts up. If the file exists, the Server starts up in the paused state.

acserver.command.shutdown

(used by the **stop** and **restart** server-control commands) Tells the Server to “finish up quickly” and exit. The Server immediately stops accepting new requests from AccuRev clients. It continues to work on transactions that are already in progress, but it aborts any transactions that are not completed within 10 seconds. Then, the Server exits.

If the Watchdog is running, it detects the Server’s shutdown and starts up a new Server immediately. Thus, this server-control file typically causes a Server restart. In any event, this file is automatically removed whenever a new Server process starts up.

If 10 seconds is not the appropriate interval for “finishing up quickly”, place another integer (such as 120) in the **acserver.command.shutdown** file. The Server exits when there are no more transactions to work on, or when the timeout interval has passed, whichever comes first.

acwatchdog.command.shutdown

(used by the **stop** server-control command) Tells the Watchdog to exit cleanly. When the Server detects that the Watchdog has exited, it exits also, just as if it had found an **acserver.command.shutdown** file (see above). In this case, however, there is no longer a Watchdog process, so no restart of the Server takes place.

This server-control file is automatically removed when a new Server process starts up.

acserver.command.taskkill

Use this capability only as a last resort. Typically, it is preferable to restart the entire AccuRev Server (which allows in-progress tasks to complete), rather than terminating just one of its worker threads.

To terminate a particular worker thread:

- Go to the **site_slice** directory.
- Determine the ID number of the worker thread by examining the thread-status table (**Tools > Server Tasks** command).
- Place the thread’s ID number (e.g. **42**) in the flag file:

```
echo 42 > tempfile
mv tempfile acserver.command.taskkill
```

Using the **mv** (or **move**) command instead of the **echo** command to create the **taskkill** file prevents a race condition that might cause the server to see **taskkill** as an empty file.

Important! After terminating a thread, restart the AccuRev Server as soon as possible. This minimizes the likelihood that terminating the thread will cause a memory resource leak in the Server process, impairing overall system performance.

Open Filehandle Limits and the AccuRev Server

The AccuRev Server is designed to handle multiple client commands concurrently: any number of requests that “read” data, along with one command that “writes” data. Accomplishing such concurrency typically requires that the AccuRev Server have many files open at the same time. Each operating system imposes limits on how many files can be open simultaneously. There may be an “open file descriptor” limit for each user process, or an overall limit for all user processes, or both. If the AccuRev Server hits the open file descriptor limit, additional client requests will be queued until file descriptors become available. (No client command is cancelled, and no data is lost. Hitting the open file descriptor limit just slows AccuRev Server performance.)

To check the current open-file-handle count:

- (UNIX/Linux) Create a file (**anyname.xml**) containing the following line:

```
<serverInfo/>
```

Then run this AccuRev command:

```
accurev xml -l anyname.xml
```

- (Windows) Use Microsoft’s **Handle** utility on the AccuRev Server machine to monitor the open file handles. (See the “Process Utilities” link at <http://www.sysinternals.com>.)

The table below indicates the default open file descriptor limits for the various AccuRev-supported operating systems. Following the table are instructions for increasing these limits.

Operating System	Default Limit on Open Filehandles	Command to Change Limit on Open Filehandles
Windows 2003, 2008 Windows XP, Vista	2048 per process	none
Solaris 9	255 per process (soft) 65536 per process (hard)	/usr/bin/ulimit -n <number> /usr/bin/ulimit -Hn <number>
Linux (Red Hat 4)	about 3500 per system	/sbin/sysctl -w fs.file-max=10000
Linux (Power PC 2.6)	1024 per process	
AIX 5.3	2000 per process	
HP-UX 11i V1	360 per process	no command; must rebuild UNIX kernel (see below)

Note: if you are performing a pre-purchase evaluation of AccuRev in an environment with a limited number of users and a limited amount of data, there is no need to make any changes. The default limits will be more than adequate.

Changing the Per-Process Open File Descriptor Limit

The procedure for increasing a process's maximum number of open files varies from operating system to operating system.

Note: in all cases, be sure to remove file **acserver.handle.limit**, located in the AccuRev **site_slice** directory, before restarting the AccuRev Server or rebooting the operating system. This file caches the current value of the open-files limit.

Linux

You must be the **root** user to perform the following procedure.

1. Change the overall limit on the number of open file descriptors each process can have (e.g. to 10,000):

```
> /sbin/sysctl -w fs.file-max=10000
```

The number you specify is stored in file **/proc/sys/fs/file-max**.

2. Add this line to file **/etc/pam.d/login**:

```
session    required pam_limits.so
```

3. Change the capabilities of the Bash shell command **ulimit**, by creating or editing the “nofile” (number of open files) lines in file **/etc/security/limits.conf**. Example:

```
*    soft    nofile    1024
*    hard    nofile    10000
```

These lines specify that:

- By default, a Bash shell (and its subprocesses) can have as many as 1024 open file descriptors.
 - A Bash shell can execute the command **ulimit -n number**, with 65535 as the maximum value of **number**. This enables that particular shell (and its subprocesses) to have up to **number** open files. (The person executing the **ulimit** command doesn't need to know what the “hard limit” specified in **/etc/security/limits.conf** is — he can just enter the command as **ulimit -n unlimited** to get the maximum value.)
4. Test your change by entering a **ulimit** command in a Bash shell, setting the limit somewhere in between the **soft** and **hard** specifications you made in Step 3. For example:

```
> /bin/bash

$ ulimit -n
1024

$ ulimit -n 5000

$ ulimit -n
5000
```

5. Restart the AccuRev Server from a Bash shell:

```
> <AccuRev-bin-directory>/acserverctl stop (stop the AccuRev Server)
```

```
> /bin/bash(start a Bash shell)
```

```
$ <AccuRev-bin-directory>/acserverctl start (restart the AccuRev Server)
```

HP-UX

You must be the **root** user to perform the following procedure.

1. Enter this command to start the System Administration Manager utility:

```
> sam
```

2. Select *Kernel Configuration*, then *Configurable Parameters*.

3. Select the **maxusers** parameter.

4. Increase the value of this parameter, for example to 128.

5. Invoke the *Actions > Process New Kernel* command, to create a new HP-UX kernel.
6. Exit the System Administration Manager utility.
7. Reboot the operating system.

AccuRev Server Performance

This section 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.

AccuRev Architecture Overview

AccuRev uses a standard client-server architecture:

- The **AccuRev Server** software runs on one machine, as a 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 AccuRev Data Repository* of the *AccuRev Concepts Manual*, along with the rest of this manual. We also recommend that you attend 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.

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.

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

AccuRev Server Machine Usage: Software

AccuRev uses a 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 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.

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 (subject to license restrictions) 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

Installing AccuRev 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
```

If more than one server has already been set up, you'll need to specify one of them during the 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  
...
```

You can modify the client configuration file 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"). The file must not contain any empty lines; be sure to check the end of the file.

The AccuRev GUI allows you to configure the contents of the **acclient.cnf** file via the **Tools > Login** dialog.

In what order should the servers be listed? 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. For users of the AccuRev GUI, it doesn't matter. The GUI's **Tools > Login** 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.) To access multiple workspaces on multiple servers, create a workspace configuration file as described in *Working with Multiple Repositories* on page 4 of the *CLI User's Guide*.

Note: It is 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.

Specifying a Server By Using the `-H` 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
```

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.

Configuring the Web User Interface (Web UI)

The optional Web UI can be used to access AccuRev functionality through a Web UI server. However, you can also access the Web UI from the AccuRev Java client to provide access to AccuWork features. Configuring the Web UI on the AccuRev server is also required if you use the AccuRev integration with Eclipse.

To configure the AccuRev Web UI for use with:

- AccuWork
- AccuRev integrations

use the following procedure:

1. Install the Web UI server as described in the *AccuRev Web User Interface Installation and Release Notes*, and make note of the host and port where it is installed.

2. On the AccuRev server, create a file named `settings.xml` in the directory `ac_install/storage/site_slice/dispatch/config`
3. This file should contain the following entry:

```
<settings>
  <webui url="http://yourwebhost:yourwebport/accurev"/>
</settings>
```

where you must substitute valid values for your site for `yourwebhost:yourwebport`.

This entry allows the AccuRev server to kick off a web browser and have it point to the correct web server.

4. Optionally, you can also add the following line within the `<settings>` tags if you wish to enforce that all of your users use the AccuRev Web UI client instead of the traditional Java client when working with AccuWork issues:

```
<usewebForIssues value="true"/>
```

Note that this setting is case-sensitive: make sure that `usewebForIssues` is capitalized exactly as shown here.

This setting will override the `preferences.xml` files for individual users, and the “Enable Issue Preview” and “Open Issues In Web Only” options in their AccuRev Preferences dialog will be disabled.

System Clock Synchronization

Time plays a fundamental role in AccuRev's architecture and in its day-to-day operations. Some examples: each transaction is recorded in the database at a particular time; a snapshot reconstructs the state of a stream at an arbitrary time; the *stat* command and the AccuRev GUI use timestamps to optimize the lookup of modified files within a workspace.

AccuRev is a networked product: programs execute on one or more server machines and multiple client machines. In a perfect world, the system clocks on all these machines would always be perfectly synchronized. The data items on the (master) server machine (say, versions created by *keep* commands) and corresponding data items on a client machine (the files that were kept) would have timestamps that are consistent with each other.

Software systems do exist that synchronize all the machines in a network to within milliseconds. If your organization uses such a system, then you don't need to read any further in this chapter!

Most software development organizations don't have — and don't need — synchronization at the millisecond level. AccuRev defines a 5-second tolerance as “good enough for software configuration management”. This chapter describes AccuRev's own facilities for detecting system-clock discrepancies, along with other related facilities commonly available on Windows and UNIX/Linux systems.

Detecting System Clock Discrepancies -- Timewarp

A timewarp (clock skew) occurs when the discrepancy between the system clocks on a client machine and a server machine exceeds the allowable tolerance of 5 seconds. A timewarp can also occur in a replication environment, when the clocks on the master server and a replica server differ by more than 5 seconds. (See *Replication of the AccuRev Repository* on page 41.)

Timewarp problems typically occur during initial system setup and during time zone adjustments. For example, the change from Eastern Standard Time to Eastern Daylight Time can cause a timewarp on a machine that is not configured correctly to handle the time zone adjustment.

For most AccuRev operations, a timewarp check is performed when the client contacts the server, or when a replica server forwards an AccuRev command to the master server.

AccuRev's Timewarp Detection Scheme

Note: In previous releases, AccuRev sometimes performed small system-clock adjustments automatically. Now, AccuRev never adjusts a machine's system clock; it only reports discrepancies.

Each time a client program contacts the server program — or a replica AccuRev Server contacts the master AccuRev Server — AccuRev compares the system clocks on the two machines. If the discrepancy is less than or equal to 5 seconds, AccuRev proceeds to execute the user's command.

If there's a timewarp exceeding 5 seconds between a client machine and a server machine *and* the user's command specifies the `-t` option with a time specification (not a transaction number), AccuRev uses the (case-insensitive) value of variable `AC_SYNC` in the user's environment to determine how to proceed:

- If `AC_SYNC` is unset or has the value **ERROR**, an error occurs and a message like this appears:

```
client_time:    2008/07/14 10:54:54 Eastern Standard Time (1216047294)
server_time:    2008/07/14 10:49:56 EDT (1216046996)
timewarp:       298 seconds
```

```
The time on this machine is more than 5 seconds different than the
time on the server. Please fix this and try again.
```

```
You may have a problem with your system clock.
```

```
You can force the time on your system to match the server time
using the accurev synctime command.
```

- If `AC_SYNC` has the value **WARN**, no error occurs, but a warning that includes the timewarp details is displayed. Command execution proceeds.
- If `AC_SYNC` has the value **IGNORE**, no error occurs and no warning is displayed. Command execution proceeds.

Notes on timewarps when a client interacting with a replica server:

- A client command that “reads” information from the repository is handled by the replica AccuRev Server; only a timewarp between the client and replica server machines is relevant.
- A client command that “writes” information to the repository is handled by the master AccuRev Server; only a timewarp between the client and master server machines is relevant.
- If there's a timewarp exceeding 5 seconds between a replica server machine and the master server machine *and* a client of the replica server invokes the command *accurev replica sync*, a warning message is recorded in the AccuRev Server log on the replica server machine.

Fixing System Clock Discrepancies

The sections below describe schemes for dealing with discrepancies between system clocks in the AccuRev client-server environment. We begin with the more desirable scheme: automatic, smooth clock adjustment. Then we describe the less desirable scheme: manual, sudden clock adjustment.

Automatic, Gradual Convergence of System Clocks

An optimal scheme for synchronizing machines' system clocks has these attributes:

- All machines in the network participate in the scheme, so the entire network is kept synchronized.
- Each machine's system clock is adjusted automatically (perhaps requiring some initial installation or configuration task).

- System clock adjustments can be made smoothly: for example, a discrepancy of 10 seconds can be gradually eliminated over the span of a few minutes by a minor speed-up or slow-down of a machine's clock. Presumably, such adjustments are imperceptible to human users and won't cause any "surprises" in time-sensitive applications.

Synchronization systems fitting this description are typically based on the standard Network Time Protocol (NTP) or its variant, the Simple Network Time Protocol (SNTP). For example, the Windows Time Service is available on recent versions of Windows. It provides a complete solution if all machines in your network are running Windows. For a more general, multi-platform solution, see the Network Time Protocol (<http://www.ntp.org>).

AccuRev-Related Guidelines

Here are guidelines for using a gradual-convergence system in an AccuRev network:

- Configure the system so that a single machine in the network acts as the time source that other machines synchronize with.
- Ideally, have all AccuRev machines participate in the synchronization system.
- If this isn't possible, make sure that the AccuRev server machine participates in the synchronization system. (AccuRev itself will take care of synchronizing its client machines to the server machine; see the next section.)

The purpose of these guidelines is to ensure that no AccuRev client machine gets into a situation of synchronizing itself with two different, and possibly conflicting, machines: the AccuRev server machine and a separate time-source machine.

Manual Synchronization Tools

A less desirable scheme for keeping system clocks synchronized is to occasionally type clock-adjustment commands manually on one or more of the machines. This method can be improved a bit by using scripts and scheduling tools such as *cron* (UNIX/Linux) and *at* (Windows). Only the **root** user (UNIX/Linux) or a user with administrator privileges (Windows) can set the system clock manually.

Setting the System Clock on the AccuRev Server Machine

On a UNIX/Linux machine, the *date* command changes the system clock. What time should you set the clock to? In many cases, you can use *rsh* or *telnet* to determine the time on another "time source" machine.

On a Windows machine, use the *net time* command to synchronize with a specified "time source" machine, or with the domain controller machine. To set the clock to a particular time, use the *date* command in a Command Prompt window, or double-click the digital clock in the Windows taskbar (lower-right corner of the screen).

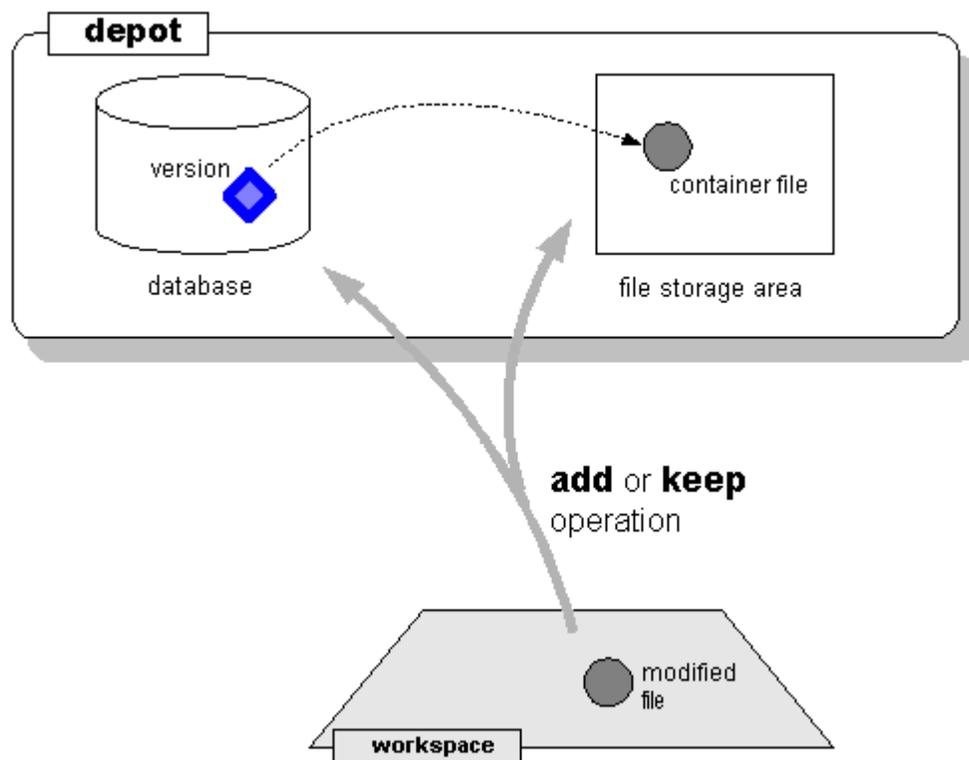
Setting the System Clock on AccuRev Client Machines

The *accurev sync* command changes a client machine's system clock to match the clock on the AccuRev server machine. The GUI command is *Tools > Synchronize Time*.

Archiving of Version Container Files

Users execute a **keep** command to preserve the current contents of a version-controlled file (“file element”) in an AccuRev depot. Similarly, users execute an **add** command to place a file under version control. The **add** and **keep** commands:

- copy the current contents of the file to a container file, located in the depot’s file storage area.
- create an associated version object in the database.



The version object can never be deleted from the database or modified in any way. The corresponding container file is always accounted for, and can be in either of these states:

- **normal** — the container file is located in the depot’s file storage area (the **data** subdirectory of the depot directory). AccuRev commands, such as **update**, **cat**, and **diff**, can access the contents of the version.
- **archived** — the container file has been moved to a gateway area outside the depot’s file storage area. AccuRev commands cannot access the contents of an archived version. After container files have been moved to the gateway area, an administrator can use standard operating system or third-party tools to transfer the container files to off-line storage: tape, CD-ROM, removable disk drive, Web-accessible storage, etc.

The AccuRev CLI commands **archive** and **unarchive** shift container files back and forth between the normal and archived states. Before using **unarchive**, the administrator must transfer the appropriate container files from off-line storage back to the gateway area. Then, invoking **unarchive** moves the container files back into the depot’s **data** directory.

The ‘archive’ Command

The command *accurev archive* processes one or more versions of file elements, shifting the versions’ container files from **normal** status to **archived** status. The command has this format:

```
accurev archive [ -E <element-type(s)> ] [ -i ] [ -p <depot> ]  
  [ -a | -I <stream-category(s)> ] [ -s <stream> ]  
  [ -t <transaction-range> ] [ -c <comment> ] [ -R ]  
  [ -Fx ] { -l <list-file> | <element-list> }
```

Determining Which Versions to Archive

archive determines the set of versions to archive as follows:

- Start with a particular set of file elements, which you specify as command-line arguments in the *<element-list>*, or in a list-file (plain-text or XML format). You can include directories in this list; in this case, use the *-R* option to include the recursive contents of those directories.
- Optionally, take the subset of versions whose element type matches the specification made with *-E*. (Note that different versions of an element can have different element types.)
- Optionally, take the subset of versions that were created in a particular streams (*-s*, for example, your current workspace stream). You can also archive versions from all streams in the depot (*-a*).
- Optionally, take the subset of versions created in a specific transaction, or range of transactions:

```
-t <number>                single transaction  
-t <number>-<number>      range of transactions
```

- In addition to the multiple subsettings of versions described above, you can use the *-I* option to *include* versions in the set, based on where in the stream hierarchy they are referenced. For example, you can include versions that were not originally included in the set because they are referenced by one or more snapshots.

If you don’t use the *-a* or *-I* option, *archive* refuses to archive any version that is currently visible in any stream or snapshot. Specified versions that are already archived are silently ignored.

Dry Run Capability

Using the *-i* option (in addition to the other options described above) generates an XML-format listing of the desired versions, but does not perform any actual archiving work. It is highly recommended that you do this “dry run” before actually archiving any versions, to avoid any surprises.

Archiving the Versions

After determining which versions to process, the *archive* command moves a version’s container file from a “normal” location under the **data** directory:

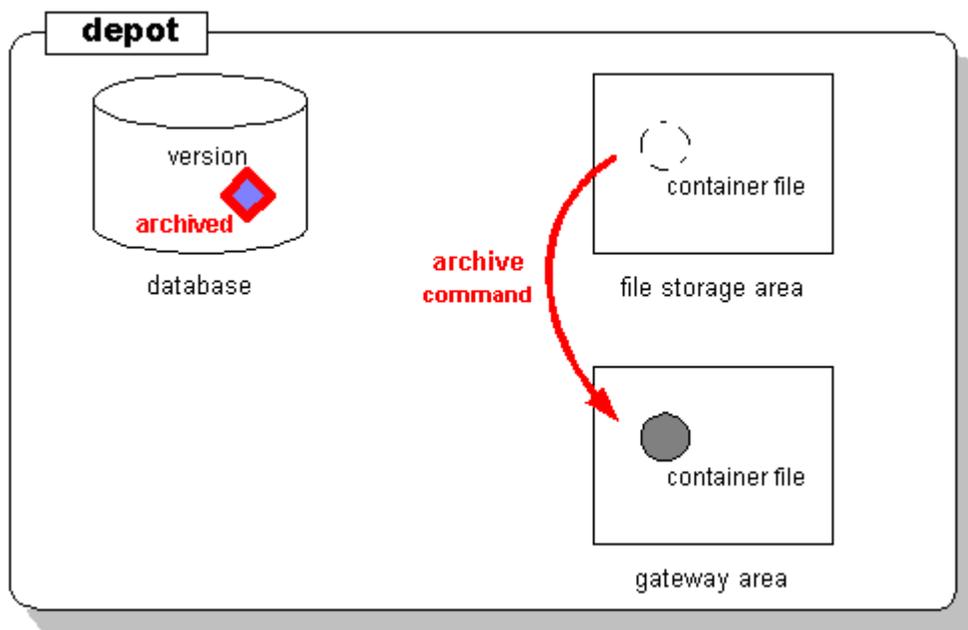
```
.../storage/depots/gizmo/data/25/07.sto
```

... to a corresponding “archived” location in the **archive_gateway/out** area:

```
.../storage/depots/gizmo/archive_gateway/out/data/25/07.sto
```

archive also marks the version as “archived” in the database.

Subsequent attempts by AccuRev commands to retrieve the contents of the archived version will fail.



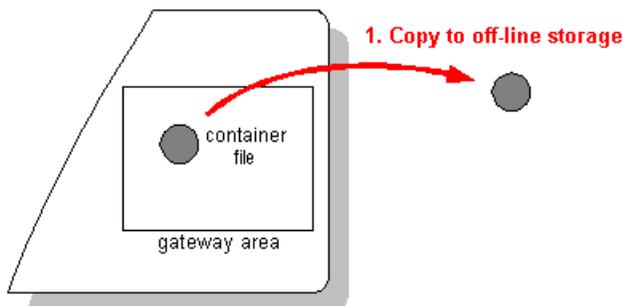
The changes made by this command are recorded in the database as a transaction of type **archive**. You can use the **-c** option to specify a comment string to be stored in this transaction. You can search for particular comment strings when using the **hist** command to locate previous **archive** transactions. See *Using ‘hist’ to Research Previous ‘archive’ Commands* on page 38.

The ‘reclaim’ Command

The **archive** command merely moves container files from one location (the depot’s **data** area) to another location (the depot’s **archive_gateway** area). To reduce the amount of disk space consumed by the archived versions, you must:

1. Copy the files from the **archive_gateway** directory tree to off-line storage. You can use operating system commands (**copy**, **xcopy**, **cp**, **tar**) and/or third-party data-backup utilities to accomplish this.

Be sure to use a tool that preserves the source data’s directory hierarchy in the copied data.

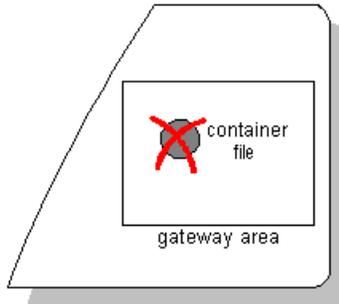


WARNING! AccuRev has no way of tracking which tool you use for this purpose, or what off-line storage medium you copy the files to. It's up to you to maintain good records of these activities!

2. Delete the files from the **archive_gateway** directory tree, using the **reclaim** command:

```
accurev reclaim [ -p <depot> ]  
                -t <archive-transaction>
```

You must specify a single transaction of type **archive**, created by previous **archive** command(s).



2. Reclaim disk space

Attempts to Access Archived Versions

The **archive** command affects depot storage only. It has no immediate effect on any workspace. But you might subsequently enter an AccuRev command that attempts to access a version that has been archived. For example, if version **gizmo_int/8** of file **floor_layout.gif** has been archived, then this command fails:

```
accurev cat -v gizmo_int/8 floor_layout.gif > old_layout.gif
```

In such cases, a message is sent to **stderr** and the command's exit status is 1.

Using 'hist' to Research Previous 'archive' Commands

Within the database, there is a complete record of all version-archiving activity for each depot. Execution of the **archive** command is recorded as a transaction of kind **archive**. You can use the **hist** command to locate all such transactions:

```
accurev hist -a -k archive
```

You can also select just those **archive** transactions that were created with a particular comment string:

```
accurev hist -a -k archive -c "stadium images"
```

In a **reclaim** command, you must indicate the storage space to be reclaimed by specifying the number of an **archive** transaction.

Restoring Archived Versions -- The unarchive Command

After you have **archive**'d some versions and **reclaim**'ed the disk space:

- The versions' container files are no longer in the depot's **data** area.
- Copies of the container files are no longer in the depot's **archive_gateway/out** area (since you've transferred them to off-line storage).

If you decide you need to restore some or all of the archived versions, you must first copy the container files from off-line storage back to the **archive_gateway** area. You must place the files under **archive_gateway/in**, at the same relative pathname as they were originally placed under **archive_gateway/out**. For example, if the *archive* command places a container file at:

```
.../storage/depots/gizmo/archive_gateway/out/data/25/07.sto
```

... you must restore the file from off-line storage to this pathname:

```
.../storage/depots/gizmo/archive_gateway/in/data/25/07.sto
```

After placing all the container files in the **archive_gateway/in** area, you can execute the *unarchive* command. This command has the same format as *archive*. That is, you specify the versions to be restored in exactly the same way as you originally archived them (with one exception — see below).

For example, to archive all non-active versions of GIF image files in stream **gizmo_maint_4.3**:

```
accurev archive -s gizmo_maint_4.3 *.gif
```

Later, you can restore all those versions:

```
accurev unarchive -s gizmo_maint_4.3 *.gif
```

Exception: you can use the *-a* option to unarchive *all* the versions currently in the **archive_gateway/in** area. In this case, the *unarchive* command syntax doesn't mimic the *archive* command syntax exactly:

```
accurev unarchive -s gizmo_maint_4.3 -a
```

This set of restored versions might have been archived in a single step or in multiple steps.

Replication of the AccuRev Repository

This chapter describes how to set up and use AccuRev's repository replication feature. One server machine stores the "master" copy of the AccuRev data repository; any number of additional server machines can store "replicas" of the repository. Each replica contains some or all of the repository's depots. Users can send commands to the AccuRev Server software running on any of these machines.

Master and Replica

One host in the network acts as the AccuRev server machine: it runs the AccuRev Server process and contains the AccuRev repository. In a replication scenario, this original host (or more precisely, this instance of the AccuRev Server process) is termed the master server.

One or more additional hosts in the network can act as replica servers. Each such host runs its own instance of the AccuRev Server process; likewise, each such host has its own copy of the AccuRev repository. The diagram below shows the servers in a replication scenario, along with various client machines.

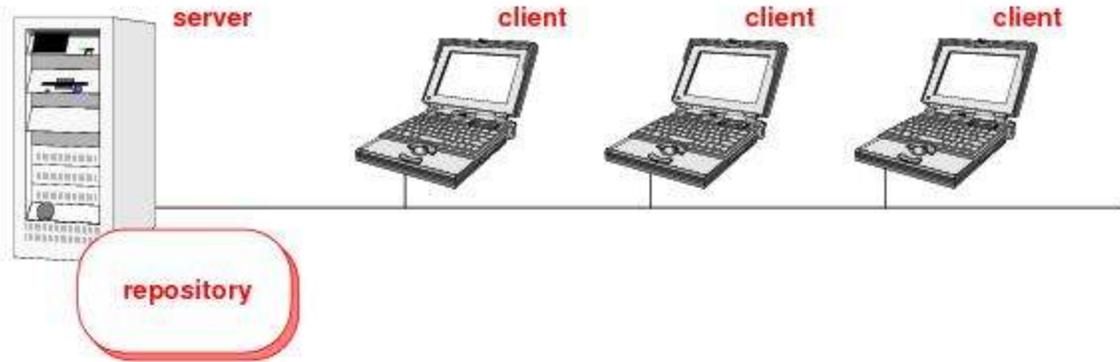
We use the terms master repository and replica repository to distinguish the multiple repositories in a replication scenario. The master repository is always complete and up-to-date; all transactions (operations that change the repository) are handled by the master server and are logged in the master repository.

By contrast, a replica repository can become out of date during day-to-day usage: it can be missing recent transactions initiated by clients using other replica servers or the master server. You can issue a simple synchronization command to download such missing transactions from the master repository to the replica repository. This makes the replica repository database into an exact copy (temporarily, at least) of the master repository database. Synchronization also occurs automatically whenever a transaction is initiated by a client using that replica server.

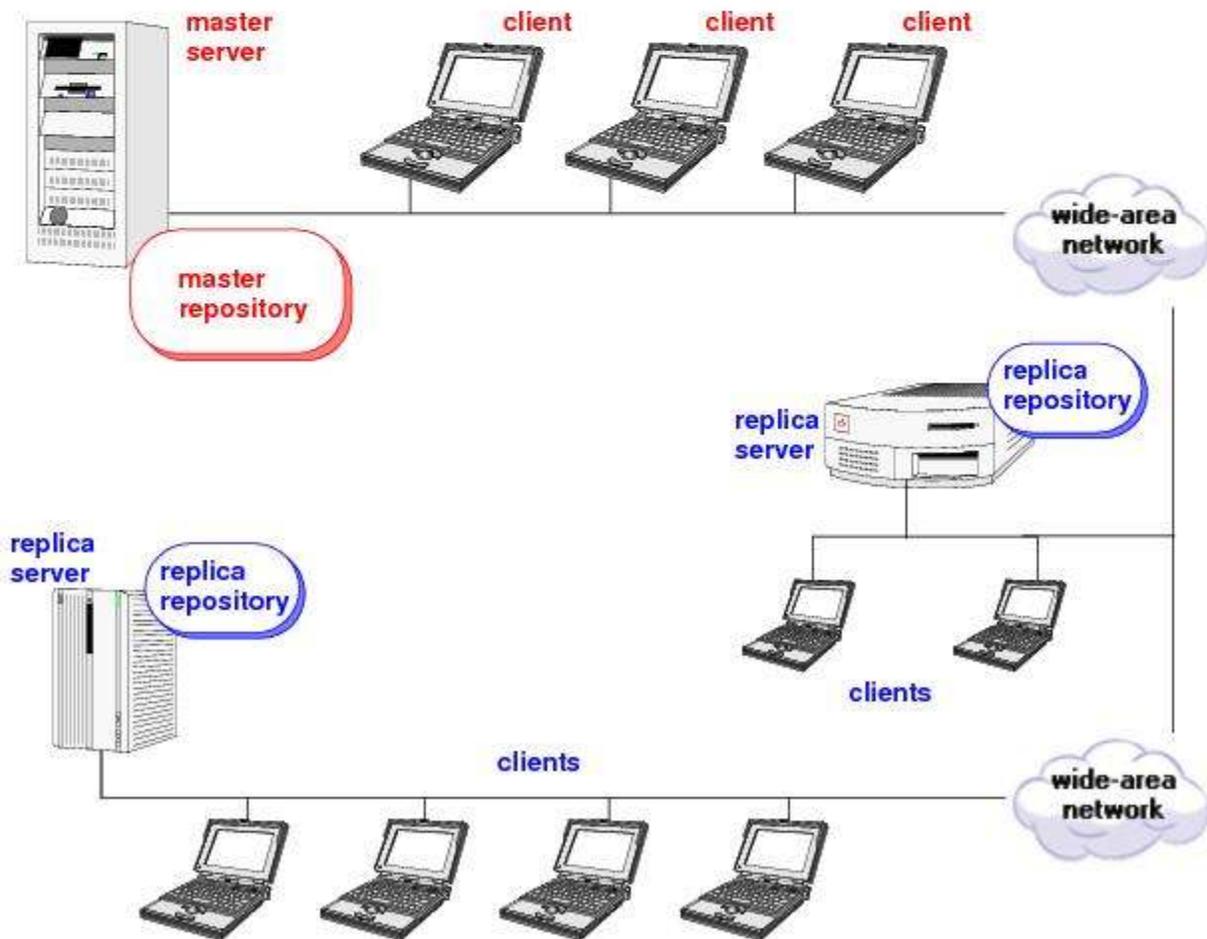
A replica repository can contain a selected subset of the depots in the master repository. If the master repository contains 10 depots, one replica repository might be configured to contain 4 of the depots, another replica repository might be configured to contain 7 of them, and a third replica repository might be configured to contain all 10 depots.

For more details on day-to-day operations involving master and replica repositories, see the sections starting with *Using a Replica Server* on page 47. First, we address licensing issues and describe the replica setup process.

Before Replication



After Replication



AccuRev Licensing in a Replication Environment

Replication server license requirements differ depending on whether you use the older proprietary named “keys.txt” license manager or the newer standard/flexible “accurev.lic” license manager (see the *License Management* chapter for more details).

Note: The named keys.txt license manager is no longer supported as of AccuRev 5.4.

Replica Licensing Using the Traditional (Named) keys.txt Manager

With named licenses, an AccuRev license enables your organization to maintain a single data repository, managed by a single AccuRev Server process running on a particular machine. This license also enables a certain number of users to access the repository with AccuRev client software.

The most obvious (and easiest) way to add replication to this environment is to have the existing server machine become the master server. Let’s say that you want to replicate the repository at two remote sites — one with 10 client users, the other with 25 client users. In this case, your organization might purchase:

- ten individual Remote User Licenses
- one Remote Site License (for the 25 users at the second remote site)

Replica Licensing Using the Standard/Flexible accurev.lic Manager

With standard/flexible licensing, AccuRev replica servers require no special license administration: license checking for the replica servers takes place on the master server automatically.

See *Replication Server Licenses* on page 107 of the *License Management* chapter for information about calculating numbers of license required by replica servers.

Installing and Configuring AccuRev in a Replication Environment

To install AccuRev software for replication, see the *AccuRev Installation and Release Notes*.

Once AccuRev replication software has been installed, use the procedures below to correctly configure your replica and master server(s).

Caution: Enabling replication poses a potential security risk. Before proceeding, be sure to read *Synchronization Security* on page 52.

Configure an AccuRev Server as a master server

1. Log in to the master and stop the AccuRev Server as described in *Starting and stopping AccuRev* in the *AccuRev Installation and Release Notes*.
2. Edit the **acservr.cnf** file, located in the AccuRev **bin** directory. Add the following line:

```
REPLICATION_ENABLED = true
```

3. If you would like the master server to authenticate logins on the replica, ensure that the authentication methods on the master server and the replica server are the same, then add the following line to the **acserver.cnf** file:

```
MASTER_AUTHENTICATES_LOGIN = true
```

4. Note the MASTER_SERVER and PORT settings in the **acserver.cnf** file. You'll need these settings when configuring the replica server below.
5. Restart the AccuRev Server process.
6. Create a new AccuRev username (*replica-user*) that will be used as the user identity for requests from a replica server.

Configure a Replica server

1. On the replica server, stop the AccuRev Server process.
2. Edit the **acserver.cnf** file, which is located in the AccuRev **bin** directory.
 - Change the keyword MASTER_SERVER to LOCAL_SERVER, and change the keyword PORT to LOCAL_PORT. But don't change the value of either setting.
 - Add new MASTER_SERVER and PORT settings, using the values of these settings on the master server. (These are the settings you noted in Step 4 above..)

After these edits, the four lines might look like this:

```
MASTER_SERVER = masthost
PORT = 5050
...
LOCAL_SERVER = replhost
LOCAL_PORT = 5050
```

Note: there is no relationship between the LOCAL_PORT and PORT numbers. They can be the same or different.

*Important! After editing acserver.cnf, you must restart the AccuRev server. If you have edited the values for SITE_SLICE_LOC, MASTER_SERVER, or PORT, you **must** also run the **maintain server_properties update** command before the new values will take effect. See [The 'maintain' Utility](#) on page 99 for more information.*

Establish an AccuRev User Identity for the AccuRev Server Process

These steps ensure that the replica server process has an AccuRev username (*replica-user*) with enough rights to access all the files in the master repository. AccuRev ACL permissions control access to depots and streams for specified AccuRev users and groups.

Caution: As of AccuRev 5.2, your site may have also implemented the optional element-level security (EACLs) feature. If you are using element security (EACLs) to control access to your data, you *must* ensure that permissions are correctly set for *replica-user*. If they are not correctly set, the replica may not be able to fetch any data from the master server, or it may be

possible for a user to log into the replica and access elements that the replica is not supposed to have access to. See *Setting Permissions for a Replica Server* on page 63 for more information.

Be sure to use the same *replica-user* username on both the replica and the master.

1. Create an operating-system user *replica-user*, and then log in as that user.
2. Configure the AccuRev server:
 - **UNIX/Linux:** Edit the server configuration file so that *replica-user* is the identity of the replica AccuRev Server. See *Operating-System User Identity of the Server Processes* on page 9.
 - **Windows:** Reconfigure the AccuRev Server service to run as *replica-user*, instead of as **LocalSystem**. In the Control Panel's Services program: open the Properties window for the AccuRev service, go to the Log On tab, select "This account", and enter the actual value of *replica-user* and its Windows password.
3. Next, establish *replica-user*'s AccuRev-level credentials:
 - Open a command-shell or C-prompt window.
 - Set environment variable ACCUREV_HOME to *replica-user*'s home directory. (**Note:** Make sure that the path to the home directory does not include any spaces.)
 - Create a "permanent" session file for user *replica-user*, for accessing the AccuRev Server on the master server from the client machine replica:

```
accurev login -n -H <mast-server-host>:<mast-server-port> replica-user
Password: *****
```
4. Continue to set up replication as described in the following sections.

The `-n` option makes this session file valid indefinitely.

Synchronize the Site Slice

Perform these steps on the replica server:

1. Start the AccuRev Server and AccuRev Database Server.
2. Run the *accurev synctime* command to ensure that the replica server time is synchronized with the master server time.
3. Run the *accurev replica sync* command to copy site-specific data from the master server to the replica server. In particular, this command makes the AccuRev Server on the replica aware of all the depots on the master server.

Indicate the Depots to be Replicated

The AccuRev repository on the replica server now has an up-to-date site slice, but the repository doesn't yet contain detailed data on any depots.

1. Log in to the replica server. See *Authenticating a Replica User on the Master* on page 56 if you would like to authenticate all replica server logins on the master server.

```
accurev login repl_user
Password: *****
```

Note: If *repl_user* on the replica server and the replica user defined on the master server bear the same name, there are now two session files for **repl_user**, one for accessing the master server and the other for accessing the replica server.

2. List all the depots in the master repository, by executing this command on the replica server:

```
accurev show -fix depots
```

In the XML-format output, the depots that exist in the master repository, but are not replicated on the replica server, are listed with this attribute:

```
ReplStatus = "missing"
```

3. For each depot that is to be replicated on the replica server, execute a *mkreplica* command. For example, if depots named **widget**, **gadget**, and **cust_support** are to be replicated:

```
> accurev mkreplica -p widget
Created replica of depot 'widget'.
Synchronizing ...
Done.
> accurev mkreplica -p gadget
...
> accurev mkreplica -p cust_support
...
```

Enabling Compression

To enable compression between the master and the replicas, add the following line to *acserver.cnf* on the master server:

```
COMPRESSION_ENABLED = true
```

Note: If your installation makes use of hardware compression, enabling this setting might decrease performance.

Setting Up a Client Machine to Use a Replica Server

A machine on which the AccuRev client software is installed can use any server — either a replica server or the master server. As always, the *SERVERS* setting in the client configuration file — **acclient.cnf** in the AccuRev **bin** directory — specifies which AccuRev Server process is to be sent client command requests. Examples:

```
SERVERS = replhost:5050          (use replica server)
SERVERS = masthost:5050         (use master server)
```

You can switch a client back and forth among multiple replica servers (and possibly the master server, too). It's as simple as editing the client's **acclient.cnf** file and then running **accurev sync** to synchronize the server times.

Using a Replica Server

When your client machine is set up to use a replica server, you can issue all AccuRev commands in the usual way. In general:

- Configuration management commands that *read* data from the repository — such as **files**, **diff**, and **cat** — use the replica repository.
- Configuration management commands that *write* data to the repository — such as **keep**, **promote**, and **merge** — use the master repository. After the master repository has been modified, the local replica repository is automatically brought up to date. For details, see [Synchronizing a Replica Manually](#) on page 50, which describes how you can bring the local replica repository up to date when you are *not* writing data to the repository.
- All AccuWork issue management operations are handled by the master server. Thus, replication does not improve AccuWork performance.

The Update Command

The **update** operation works as follows when you execute it on a client that uses a replica server:

1. An implicit **replica sync** command is performed, copying database transactions from the master repository to the replica repository. This brings the replica database completely up to date.
2. A **stat** operation is performed on the replica server, to determine the state of the workspace stream and its backing stream.
3. Data files representing new versions of elements are copied from the file storage area in the master repository to the file storage area in the replica repository.
4. Data files are copied from the replica repository to your workspace tree. In addition to the files retrieved from the master repository in the preceding step, this can include files that have already been “downloaded” to the replica repository through other users' commands.
5. On both the replica server and the master server, the transaction level of the workspace is set to the most recent transaction (or to the transaction specified with **update -t**).

Command Interaction in a Replicated Environment

If a particular replica server is executing a **mkreplica** command, a request to perform a **replica sync** on that server might fail. Be careful to run **mkreplica** at a time when it is unlikely that any user or script is invoking **replica sync**.

If a **replica sync** command does fail for this reason, first make sure that no **mkreplica** command is executing, then invoke **replica sync** again.

Similarly, if a *replica sync* command is already in progress, in rare situations a client command sent to the same replica server might not complete correctly, causing an error message to be displayed. The change is made correctly on the master server in such situations. The user who issued the command can perform an *update* command to have the change reflected in the workspace.

Removing Storage Containers on a Replica Server

Use the *replica archive* command to remove unused container files from the depot's **storage** directory on a replica server.

The *replica archive* command activates the **server_preop_trig** trigger with a **replica_archive** command. The command lists each element for which the storage container was removed. For example:

```
Removing storage for element /file1.txt (2/5)
Removing storage for element /file1.txt (2/6)
```

For more information on this command, see the *replica* command in the *AccuRev CLI User's Guide*, or enter *accurev help replica* in the CLI.

Removing a Replica Server

Since the master repository is always complete and up-to-date, removing a replica server is very simple:

1. Inform any users on the replica server of the change. Every user must log out and edit the **acclient.cnf** file (located in the AccuRev **bin** directory) to point to another server.
2. On the replica server machine:
 - Stop the AccuRev Server and AccuRev Database Server.
 - Uninstall AccuRev by running the uninstall program.
 - Access the AccuRev installation directory and remove any remaining files.

Improving Replica Performance

The two techniques described below can greatly improve the user experience with AccuRev replicas.

Performance of a Newly Created Replica

Use the following procedure to download version files to a replica when it is first created, instead of when the replica's first user performs an *Update*.

Perform these steps on the master server:

1. Stop the AccuRev Server and AccuRev Database Server.
2. For each depot to be replicated, use a ZIP utility (or *tar*, or some other tool) to copy the **data** subdirectory of the depot's slice. Example:

```
cd /opt/accurev/storage/depots
zip -r /tmp/oscar_data.zip oscar/data
zip -r /tmp/felix_data.zip felix/data
```

3. Restart the AccuRev Server.

Perform these steps on the replica server:

4. Replicate the depots with the command **accurev mkreplica**. Example:

```
accurev mkreplica -p oscar
accurev mkreplica -p felix
```

5. Stop the AccuRev Server and AccuRev Database Server.

6. For each replicated depot, move the slice's **data** subdirectory aside. Example:

```
cd /var/applications/accurev/storage/depots
mv oscar/data oscar/data.ORIG
mv felix/data felix/data.ORIG
```

7. For each replicated depot, recreate the slice's **data** subdirectory, using the copy created in Step 2. Example:

```
unzip oscar_data.zip
unzip felix_data.zip
```

8. Restart the AccuRev Server.

Performance of Heavily Used Replicas

In a distributed computing environment, many new versions might be created in a replica at Site A, while no one is working at Site B. The first person to update his workspace at Site B pays the penalty of waiting for all those new versions to be transferred from the master server to Site B.

To avoid this situation, create a “dummy” workspace at Site B for each heavily used stream, and create a script that updates all these workspaces. Then, set up a **cron** job (UNIX/Linux) or a Scheduled Task (Windows) to run the script on a regular basis. Ideally, the script would run at least once a day — before the start of Site B's work day, but after the end of the work day of Site A (and any other sites). For example, this **crontab** line runs script **update_replica_workspaces** at 4:30AM every day:

```
30 4 * * * /usr/local/bin/update_replica_workspaces
```

Running the script at additional times during the work day can further improve replication performance from the users' viewpoint — if there is sufficient file-transfer bandwidth.

Triggers and Replication

A user's invocation of an AccuRev command on a client machine may cause a client-side trigger and/or a server-side trigger to fire. A client-side trigger fires on the user's client machine — no ambiguity there. But in a replication environment, there are two AccuRev Servers to consider: the

one on the replica server machine and the one on the master server machine. Where does a server-side trigger fire?

- A *rmreplica* command fires a *server_admin_trig* trigger on the replica server. The AccuRev user identity (“principal-name”) of the user who invoked the command is passed to the trigger script.
- A *replica sync* command fires a *server_admin_trig* trigger on the master server. The AccuRev user identity of the AccuRev Server process on the replica server is passed to the trigger script — not the identity of the user who invoked the command. The command-name is passed to the trigger script as **replica_sync**.
- For all other commands that write to the repository, a server-side trigger fires on the master server. The AccuRev user identity (“principal-name”) of the user who invoked the command is passed to the trigger script.
- For commands that do not write to the repository, a server-side trigger fires on the replica server.

Creating New Depots

New depots can be created only in the master repository, not in a replica repository. If a client using a replica repository issues a *mkdepot* command, an error occurs:

```
Cannot create a new depot on the replica server
```

After creating a new depot in the master repository, you can include it in a replica repository with this sequence of commands, issued on a client that uses the replica server:

```
accurev replica sync
accurev mkreplica -p <depot-name>
```

Adding and Removing Depots from a Replica Repository

After you have set up a replica repository, you can use the commands *mkreplica* and *rmreplica* to change which depots are included in the replica repository. These commands are described in the *AccuRev CLI User’s Guide*, at *mkreplica* on page 167, and *rmreplica* on page 218.

Synchronizing a Replica Manually

During the course of development, your local replica repository typically becomes out-of-date with respect to the master repository. This occurs when other users send commands to other replica servers or directly to the master server. In both such cases, new transactions are entered in the master repository, but are not entered in the your local replica repository.

At any time, you can enter this CLI command to bring your local replica repository up to date:

```
accurev replica sync
```

This transfers data from the master repository site slice to the replica repository site slice. It also transfers database transactions from the master repository to the replica repository — but only for

the depots that are included in the local replica. It does not transfer the corresponding storage files for *keep* transactions. See *On-Demand Downloading of a Version's Storage File* on page 51.

A *replica sync* command is performed automatically on the local replica after each operation that is initiated by a client of the local replica, and that makes a change to the repository. See *Using a Replica Server* on page 47.

Note: you never need to synchronize directly with other replicas; synchronizing with the master is sufficient to bring your replica up to date.

On-Demand Downloading of a Version's Storage File

As a performance optimization, AccuRev copies database transactions only — not storage files that hold the contents of *keep* versions — when it synchronizes the master repository with a replica repository during:

- a *replica sync* command
- the automatic replica synchronization that follows an operation, invoked by a client using a replica server, that modifies the repository

Storage files for versions are downloaded from the master repository to the local replica repository during an *update*. (See *The Update Command* on page 47.) The storage file for an individual version is downloaded when a client using a replica server explicitly references that version. Examples:

```
accurev cat -v talon_dvt/12 mywork.c
accurev diff -v talon_dvt/12 mywork.c
```

Both these commands cause the storage file for version **talon_dvt/12** of file **mywork.c** to be downloaded to the local replica repository before the command itself is executed.

Automating Replica Synchronization

If a workgroup is much less active than other workgroups, its local replica repository can “fall behind” the master repository significantly. This can also occur if the workgroup uses the local replica repository mostly as a reference — for frequent read operations, but infrequent write operations. Falling behind in this way does no harm, but it can be bothersome. When some user finally does perform a write operation — keeping a new version of a file, or changing the location of a workspace — the local replica repository automatically “catches up”, which might involve downloading tens or hundreds of transactions.

To prevent the local replica repository from falling too far behind, we recommend that you use operating system tools to perform an *accurev replica sync* command automatically, at regular intervals — say, every 15 minutes. On a Windows machine, use the Scheduled Tasks program in the Control Panel. On a UNIX/Linux host, set up a *cron* job to execute this command.

Synchronization Security

Note: this section describes a security risk that exists only for organizations using the *AccuRev Replication Server* product. This risk does not apply to organizations that use the standard AccuRev software, without the replication option.

The repository synchronization scheme poses a potential security risk: the **acservr.cnf** server configuration file on an AccuRev server machine can name *any* master server machine in a MASTER_SERVER setting. And by default, the targeted master server will comply with *any* synchronization request — even an *accurev replica sync* command executed on a completely unrelated client machine.

We strongly recommend using the *server_admin_trig* trigger on the master server machine to implement an authentication scheme, so that the master server will send repository data over the wire only to valid requestors. The following Perl code might be added to the sample *server_admin_trig* script included in the **examples** subdirectory of the AccuRev distribution:

```
if ($command eq "replica_sync") {
    if ($principal ne "rep01_acadmin" and $principal ne "rep02_acadmin") {
        print TIO "Repository synchronization disallowed:\n";
        print TIO "Authentication by the server_admin_trig script failed.\n";
        close TIO;
        exit(1);
    }
}
```

This code allows users **rep01_acadmin** and **rep02_acadmin** to perform repository synchronization, rejecting requests from all other user identities.

Note: a *server_admin_trig* script identifies the command as *replica_sync*, even though the actual CLI command is *replica sync*.

The replica_site.xml File

Each replica repository's site slice directory contains an XML-format file, **replica_site.xml**. This file contains information about the depots that are replicated in that repository. The *mkreplica* and *rmreplica* commands maintain the contents of this file.

Moving the AccuRev Server and Repository to Another Machine

The AccuRev data repository should be physically located on the machine that runs the AccuRev Server process. (This is not an absolute restriction — see *READ ME NOW: Assuring the Integrity of the AccuRev Repository* on page 1.) The repository consists of multiple slices: the site slice contains information that pertains to the entire repository, and each depot has its own slice.

From time to time, you may want (or need) to have the AccuRev server process run on a different machine. To accomplish this, you must:

- Install AccuRev on the new machine.
- Move your data repository to the new machine.

The procedure described below is safe to use even if the new machine has a different byte order than the old machine.

Procedure for Moving the Repository

Make sure you perform each of the following steps on the appropriate server machine. We call them:

- The *source* machine — where the AccuRev server is currently running and the data repository is currently located.
- The *destination* machine — the machine to which you want to move the data repository.

Note: the steps below always show UNIX/Linux pathname separators (/). When you're executing commands on a Windows machine (either source or destination), be sure to use Windows pathname separators (\).

The procedure calls for multiple stops and starts of the AccuRev Server process. For details on how to accomplish this, see *Controlling Server Operation* on page 17.

On the Source Machine

1. Execute the *accurev show slices* and *accurev show depots* commands, and save the output for reference in the following steps.
2. Perform a backup of the AccuRev repository, as described in *Backing Up the Repository* on page 3.
3. Make the backed-up files accessible to the destination machine.

On the Destination Machine

1. Request a new license key for the destination machine by filling out the form at <http://www.accurev.com/license-request.shtml>. Save this key file to a location on the destination machine.
2. Perform a full AccuRev Server installation on this machine as described in “To install an AccuRev 5.0 server” in the *AccuRev 5.0 Installation and Release Notes*. The installation directory pathname need not be the same as on the source machine.
3. Stop the AccuRev Server and AccuRev Database Server, as described in *Controlling Server Operation* on page 17.
4. Perform a restore of the AccuRev repository you backed up on the source machine, as described in *Restoring the Repository* on page 4. Check that the file ownership and permissions are correct; they may need to be altered if the username used to run the AccuRev server is different on the source and destination machines, or if your backup/restore tool does not correctly preserve these settings.
5. Make sure that the AccuRev Database Server (but not the AccuRev Server) is running, which should be the case after performing the “Restoring the Repository” procedure in the previous step.
6. Update the server properties:

```
maintain server_properties update
```

See *The ‘maintain’ Utility* on page 99 for more information.
7. Change the location(s) of the slices to their locations on this machine using the output from the *accurev show slices* and *accurev show depots* commands on the source machine. To do so, run the following command for each slice:

```
maintain chslice <slice-number> <new-location>
```
8. Start the AccuRev Server.

AccuRev Security Overview

This chapter presents an overview of AccuRev’s security-related features. We discuss and compare the following topics, and provide pointers to more detailed documentation.

- *Users and Groups*
- *User Authentication*
- *Locks on Streams*
- *Access Control List (ACL) Permissions*
- *Element-Level Security (EACLs)*
- *Restricting Access to Commands using Triggers*

Users and Groups

Each AccuRev user must have an AccuRev username (also called a “principal-name”). AccuRev maintains its own user registry, which is separate from the registry maintained by your machine’s operating system (or the network).

Optionally, you can create AccuRev user groups, and add users to the groups as “members”. Starting in Version 4.5, groups can be nested within one another; that is, a group can be a member of another group. AccuRev groups are independent of operating system groups.

Groups are used by the ACL facility to grant or deny access to a particular resource for an entire set of users. They are used by the Locks facility to specify a set of users to which a stream lock does or does not apply. (See *Locks on Streams* on page 58 and *Access Control List (ACL) Permissions* on page 59.)

Usernames and Groupnames

Each AccuRev group has a user-defined name. Usernames and group names share a “namespace” in AccuRev. This means that a user and a group cannot have the same name.

User Authentication

With a few exceptions, a user cannot execute AccuRev commands unless he is authenticated as a registered AccuRev user. For authentication purposes, each registered AccuRev user has a username/password pair recorded in the database. A user’s password can be empty.

AccuRev automatically distinguishes two categories of users — those who have empty passwords, and those who have non-empty passwords. The keywords **authuser** and **anyuser**, respectively, are used by the ACL facility to designate these categories. See *secinfo* on page 221 of the *AccuRev CLI User’s Manual* for more information.

The “accurev_login” User-Authentication Method

AccuRev supports authenticating a user through an explicit login to the AccuRev Server. Using the GUI or the CLI, the user invokes the **Login** command and enters his password (possibly empty). This launches a user session, which is typically of limited duration (a few hours). When the session expires, the user must login again to continue using AccuRev.

The “custom” User-Authentication Method

AccuRev also supports script-based authentication of AccuRev users. This **custom** method works as follows:

1. The user invokes the **Login** command through the AccuRev GUI or CLI client, and types a password.
2. The client encrypts the password and transmits the username/password combination to the AccuRev Server.
3. The AccuRev Server verifies that the username exists in AccuRev’s user registry. It does not check the password, however.
4. The AccuRev Server invokes a script named *server_auth_trig*, passing it the username and password combination. The exit status of this script determines the success/failure of the **Login** command. See *The ‘server_auth_trig’ Script* on page 57 for details.

This feature provides tremendous flexibility. For example, an authentication script could use an external user database, accessed through an LDAP interface, to perform AccuRev user authentication.

Authenticating a Replica User on the Master

Prior to V4.7.3, any authentication scheme on the master was only as trustworthy as the security of all the replica servers. The changes to replica authentication in AccuRev 4.7.3 allow sites using replication to deploy replicas to sites they may not trust to do user authentication, by allowing user authentication to be performed on the master server only.

Use the MASTER_AUTHENTICATES_LOGIN parameter in the **acservr.cnf** file on the master server:

- If this parameter is set to **true**, user authentication is performed on the master server only. The master and replica servers must have the same authentication method. Only the *accurev_login* or *custom* authentication methods are allowed. The replica server must be able to contact the master server to log in; if the connection fails, the login fails.

To change the authentication method, change and restart the master server first, then change and restart all replicas.

- If this parameter is set to **false**, user authentication can be performed on the replica, as before.

After changing the value of the parameter, restart the master server.

As a result of these changes, the XML input to the *server_auth_trig* script has changed to add **server** and **port** elements, which should contain the hostname and port of the master server.

The 'server_auth_trig' Script

The *server_auth_trig* user-authentication script, used by the **custom** user-authentication method, is similar to the *server_admin_trig* administrative trigger script. A sample script, implemented in Perl, is included in the AccuRev software distribution, in the **examples** subdirectory.

To deploy a user-authentication script, place an executable file in subdirectory **triggers** of the **site_slice** directory:

- UNIX/Linux: the file must be named **server_auth_trig** or **server_auth_trig.pl**
- Windows: the file must be named **server_auth_trig.bat** or **server_auth_trig.exe**

The script is called when a user invokes the AccuRev *Login* command. It executes in a subprocess of the AccuRev Server. The script's standard input is a simple XML document, with the structure shown in this example:

```
<triggerInput>
  <hook>server_auth_trig</hook>
  <command>login</command>
  <ip>192.168.6.186</ip>
  <username>derek</username>
  <password>MyP@ssw0rd</password>
  <server>myserver</server>
  <port>5050</port>
</triggerInput>
```

(No password encryption is necessary, because this XML document is passed to the subprocess through an operating system pipe, not through a file.) The script's standard output is appended to file **triggers.log**, located in the **logs** subdirectory of the **site_slice** directory. No output is required, however; only the script's exit status is significant:

- If exit status is zero, the user's Login command succeeds.
- If exit status is nonzero, the user's Login command fails.

The user's **Login** command also fails if the script does not execute properly — for example, if a runtime error occurs or the script file is not executable.

Selecting the User-Authentication Method

When you install AccuRev on the machine where the AccuRev Server process will run, the AccuRev Installation Wizard automatically sets the authentication method. Thereafter, you can switch methods with the **authmethod** command. Example:

```
accurev authmethod accurev_login
```

The switch takes place immediately. This might cause user confusion; it is not a command to run casually.

How AccuRev Records the User-Authentication Method

The current user-authentication method is stored in file **preferences.xml** in the **site_slice** directory. (Don't confuse this with individual users' **preferences.xml** files.) For example:

```
<preferences>
  <authmethod>accurev_login</authmethod>
</preferences>
```

As an alternative to the **authmethod** command, you can change the user-authentication method by editing this file, then stopping and restarting the AccuRev Server process.

Restriction on Access to the “Add User” Command

User authentication is performed by the add-new-user command (GUI: **Add User**, CLI: **mkuser**). To prevent a chicken-and-egg problem, user authentication is bypassed if the AccuRev user registry is currently empty. After the first AccuRev user has been created, only an authenticated user can add users to the AccuRev user registry.

Locks on Streams

Each stream can have one or more locks on it. A lock prevents a certain set of users from using the stream to create new versions of elements. That is, it prevents execution of the **Promote** command — either promoting from the designated stream, or promoting to the designated stream, or promoting in either direction.

A lock can be absolute, applying to all users. Alternatively, you can specify that a lock applies to a particular AccuRev user, or to a particular AccuRev group. Conversely, you can specify that a lock applies to everyone *except* a particular AccuRev user, or to everyone *except* a particular AccuRev group.

Locks can also prevent reconfiguration of the contents of a stream with the include/exclude facility.

For more on locks, see the *lock* reference page in the *AccuRev CLI User's Manual*.

Access Control List (ACL) Permissions

In addition to (or instead of) locks, each stream can have one or more permissions on it. Whereas a lock controls the ability to create new versions (through the *Promote* command), a permission is more general: in addition to controlling *Promote*, it controls the ability to read data from the stream, using such commands as *Annotate*, *Diff*, and *Open*. A permission also controls workspace-specific commands, such as *Update* and *Populate*.

Unlike locks, which always apply to individual streams, ACL permissions can be configured to apply to entire stream subhierarchies.

You can create an ACL permission that applies to an entire depot. This provides a way of controlling access to all of a depot's file system data, in all streams. It also provides a way to control access to a depot's AccuWork issues.

For more on ACL permissions, see the *setacl* reference page in the *AccuRev CLI User's Manual*.

Element-Level Security (EACLs)

Historically, AccuRev has supported a security mechanism for depots and streams using Access Control Lists (ACLs). To address customer needs for more extensive security, AccuRev 5.2 introduces greatly enhanced ACL security control at the element level.

A depot ACL protects all the elements, streams, and issues in the depot. A user who has no access to the depot can still see the stream hierarchy, but none of the files in the streams. They can also not edit the streams in anyway. A stream ACL prohibits editing streams and viewing content in a stream. It can also be inheritable down the hierarchy. The table below shows what stream, depot and element ACLs restrict:

	Viewing Content	Editing Stream	View and Editing Issues
Stream ACL	Yes	Yes	No
Depot ACL	Yes	Yes	Yes
Element ACL (new)	Yes	No	No

Features

There are three main features of the new element security mechanism:

- The ability to allow or deny access to all versions of an element, no matter what stream they are in, for a specified user or group. If a user is denied access they cannot see or view the element.
- The ability to specify who can modify the access to an element. This is a higher privilege than just being allowed access.
- The ability to view the history of access changes for auditing purposes. This means being able to trace who assigned access to an element, when they did it and what they did.

Basic Terms

To control access to elements (files and directories), AccuRev 5.x implements a new *eacl* command (see *eacl* on page 104 of the AccuRev *CLI User's Guide*) to set and modify *Access Control Lists (ACLs)* and *Access Control Entries (ACEs)*. An ACL is a list of security protections that applies to an element. An ACE is an entry in an ACL that defines a *principal* and a *privilege*.

A principal can be:

- a user
- a group
- *all*

A privilege can be:

- *Full* – the ability to see and view the element and to modify its ACL
- *Allow* – the ability to see and view the element, but not modify its ACL
- *Deny* – the inability to see and view the element or modify its ACL

An ACL contains zero or more ACEs. An element can have only one ACL assigned to it at any point in time. You specify whether to set, add, or remove an ACE (principal and privilege) to an element, and AccuRev takes care of the ACLs automatically. Note that ACLs cannot be created or modified without an element.

Important Concepts

Before jumping into implementation details, it is important to keep the following concepts in mind to avoid confusion and incorrect assumptions:

Element-based security

AccuRev security is based on AccuRev *elements* rather than pathnames, and this makes AccuRev security both more powerful and more complex than typical filesystem security. Don't forget: an AccuRev element can be either a file or a directory, and can also be an element link (elink) or symbolic link (slink).

An element is defined by its element id (“EID”), which never changes. It also has a pathname (defined by its “parent EID”) and a name, both of which can change.

A symbolic link (slink) is an element whose contents is a pathname which can point to either AccuRev or non-AccuRev data.

An element link (elink) is an element whose contents is a pointer to another element, which must be in the same depot. The target element can be a directory element, a file element, another element link, or a symbolic link.

A key concept to master is how ACLs apply to the element versus the element’s contents. See *Inheritance* on page 62 to understand how access to *contents* and *namespace* differs,

Likewise, if you use element links (elinks), you must understand how AccuRev element security does and does not work apply to the link and to the destination file, otherwise a file that you think is denied to a principal may still be accessible.

In AccuRev, the same element can have multiple pathnames, depending on what streams it is in at different times. While pathnames and filenames can change, the element id (EID) never does, and this is what AccuRev security is based upon. This allows you to secure the contents of an element no matter what its current pathname is.

Remember:

- A simple pathname may refer to more than one element (such as both ends of an elink), and likewise a single element may have more than one pathname (such as one pathname to the actual file, and a different pathname for that file as a target of an elink).
- An element may have different paths in different streams, and may have different paths at different points in time.

ACLs are not TimeSafe[®]

The third concept to keep in mind when using AccuRev element security is that ACLs are always current: they are NOT TimeSafe[®], and they are not affected by what stream they happen to appear in. This means that you cannot recreate ACLs as they appeared at a certain point the same way as you can recreate file versions or stream configurations. (You can, however, view ACL changes through the *hist* command.) It also means that the appearance of a snapshot stream can change for a particular user or group, The snapshot doesn’t change—that is TimeSafe. But a user’s *access* to the contents of that snapshot can change. If your access changes, the next time you perform a *pop* command on your workspace, you may see files appear or disappear.

Permissions

The default behavior for an element whose ACL doesn’t contain an ACE for a user is to deny access to that user. The *deny* privilege overrides *allow* and *full*. This makes it easy to give access to everyone (all), but then deny a few. If a user has both *allow* and *full*, then *full* privilege is granted.

The *full* privilege means the ability to change which ACL is associated with the element and to modify the ACL as well. A user only needs to have *allow* privilege to show ACLs on an element

or to view the ACL. When setting or modifying an ACL for an element, it is the current ACL associated with the element that controls the ability to set or modify the ACL. The table below describes what a user can or cannot do based on privileges:

	View ACL	Modify ACL	View Elements	Modify Elements
Full	Yes	Yes	Yes	Yes
Allow	Yes	No	Yes	Yes
Deny	No	No	No	No

Inheritance

Element security inheritance differs for content and namespace.

Content is based on a *static inheritance* model. This means that an ACL is set for an element when assigned by the user or when it is created and inherits the ACL from its parent.

Namespace access is based on *dynamic inheritance*. This means that it is computed by the element's pathname when it is accessed.

This means that if you are denied access to a particular element, you cannot see the contents or even the name of the element. However, if you are denied access to directory, but not denied access to a file in that directory, you can still perform certain operations on that file by using its EID. For example, if you know its EID, you can *cat* the file. You will not be able to see its name, but you can see its contents.

The user also has the ability to set an ACL recursively down the directory tree, making it easy to assign an ACL to multiple elements. Multiple elements can be specified on the command line as well.

Renaming, moving, or defuncting an element does not change its ACL or access. Only setting or modifying an ACL on the element can change access to the element.

Access Denied

This section describes what it means when access is denied to the content or the name space of an element:

- A user is denied access to the content and namespace of an element if they are denied access to the element itself.
- A user is also denied access to the namespace of an element if they are denied access to any of the directories in the element path.
- If a user is denied access to the namespace, but not the element, the user will be denied access to the content if they try to specify the element name.
- If a user is denied access to the namespace, but not the element, the user will NOT be denied access to the content if they try to specify the element EID.

Installing AccuRev with Element Security for the First Time

When you install an AccuRev release that contains element security for the first time, all elements are assigned a single default ACL created during upgrade that allows "all access". This provides new and existing customers with the exact same behavior provided by previous versions of AccuRev.

Since `allow` is the only privilege set during the upgrade process, you need the ability to assign ACLs or modify them afterwards. To allow this, AccuRev introduces a new *superuser* role. A superuser has special security privileges to modify any ACL and assign an ACL to any element no matter what ACLs are already there. A superuser always has full access on all elements.

The superuser is needed to begin the process of setting up element security, and to fix problems that users may cause, such as denying everyone access to an element.

During the installation process you (the AccuRev Administrator) are prompted to specify an existing user or a new user to be the superuser. An administrator can also set or unset a user as a superuser through the **maintain** program.:

```
maintain su -a <accountname>
```

The `-r` option allows you to remove superuser status from a user,

```
maintain su -r <accountname>
```

Auditing

To provide auditing (or history) capability for element security, AccuRev keeps track of all ACL changes to elements. This doesn't mean that element ACLs are TimeSafe® (and it is important to realize that element ACLs are **not** TimeSafe). The most recent version of an ACL and the most recent ACL assigned to an element is always used for authorization of access. Historical information on changes to ACLs is only used for audit reporting. The *hist* command is used to display the history of ACL changes on an elements.

Setting Permissions for a Replica Server

If you are using a replica server in an EACL environment, it is critical that you set EACL permissions correctly not only for the users on the replica, but for the special replica user account (*replica-user*) that is used to communicate between the replica and the master server.

For example, assume that you want the users of a replica server to be able to access all of the files underneath folder `offshore_files`, but that you do not wish them to be able to see anything under `hq_files`. If you fail to assign **allow** permission to *replica-user* for `offshore_files`, the replica may not be able to fetch any files from the master.

More importantly, if you fail to assign **deny** permission to *replica-user* for `hq_files`, it would be possible for a privileged user to log into the replica and bring the denied files down to the replica. Continuing this example, say that a privileged user from corporate headquarters visits the offshore site where the replica is being used. If the *replica-user* is not correctly configured, the privileged user could log into the replica server and inadvertently bring down `hq_files` elements by doing an

update. If the *replica-user* is configured with **deny** access to these files, the server will not send them to the replica.

EACL Usage Scenarios

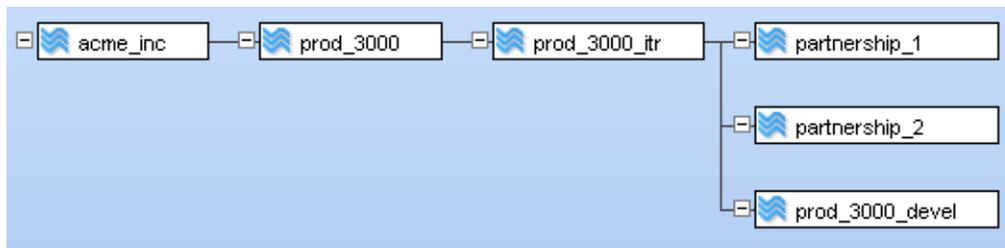
The following examples illustrate two basic approaches to setting up element ACLs:

- **less restrictive**—By default, all users have **allow** access to all elements, unless explicitly denied.
- **more restrictive**—By default, all users are denied access to all elements, unless explicitly allowed.

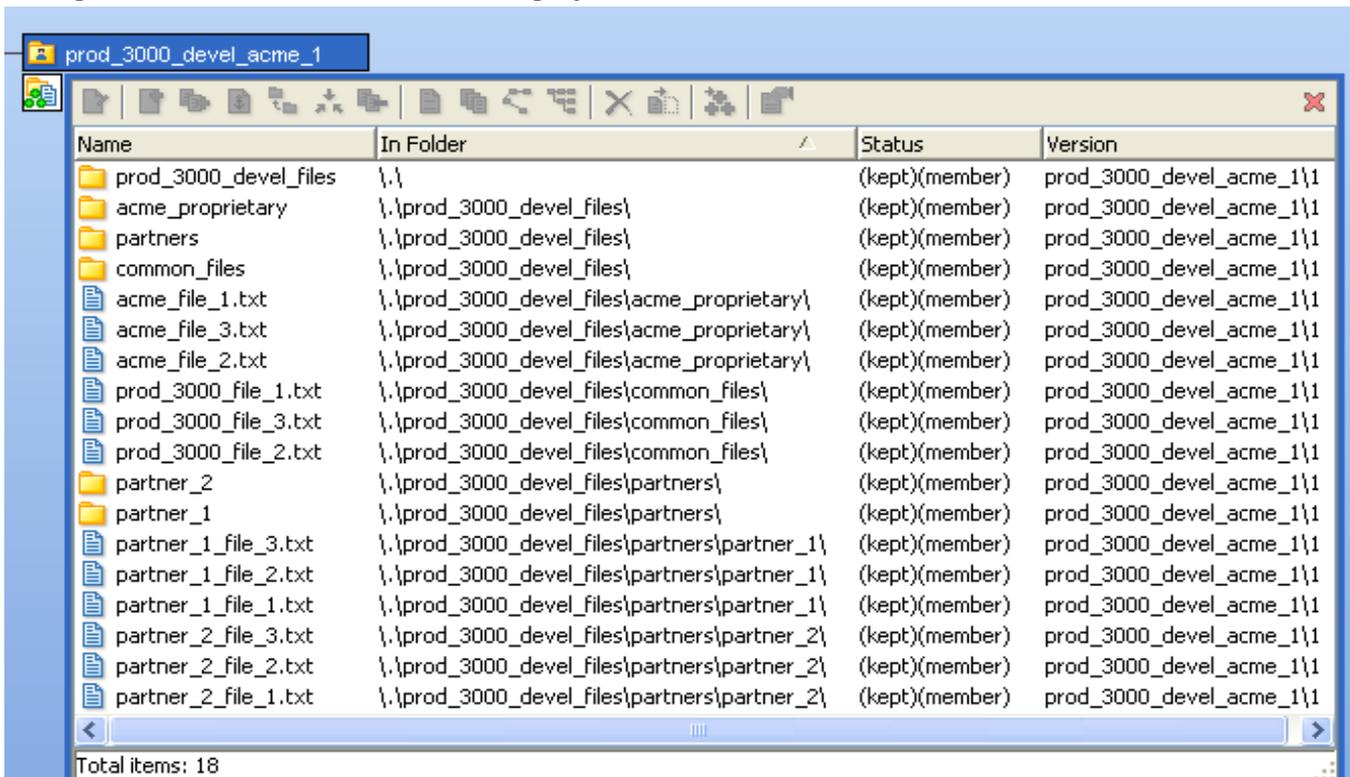
Both examples make use of the same environment: Assume that your company (“Acme Agile, Inc.”) has a new product (“Product3000”), and you will be collaborating with two partners to develop the code (“Partnership_1” and “Partnership_2”). For simplicity, these examples illustrate settings ACLs for specific users. In a real-life scenario, you would typically assign many users to various groups, and then set the element ACLs by group.

Set Up

The AccuRev administrator initially set up a basic stream structure:



You, Acme employee user “acme_1”, create a workspace off of stream `prod_3000_devel` and set up the initial files and folders for the project:



Since you are a trusted user, the AccuRev administrator grants you “superuser” status so that you will be able to assign element ACLs as needed.

To make you a superuser, the AccuRev administrator logs in to the AccuRev server machine and uses the AccuRev *maintain* command. **Note:** The administrator must first shut down the AccuRev server before using *maintain*.

```
> maintain su -a acme_1
AccuRev 5.2 (2011/04/15) Enterprise Edition
Copyright (c) 1995-2011 AccuRev Inc. All rights reserved
Changed user 'acme_1' to a super user
>
```

The administrator restarts the AccuRev server, and you log in to the AccuRev CLI as `acme_1` and start assigning EACLs as shown below.

Less Restrictive Scenario

Your starting point is the system default for first time AccuRev installations: everybody has *allow* access to all files and namespaces, and you will explicitly change privileges to certain elements for certain users to either **full** or **deny**.

Remember:

- **deny** means that the user cannot even see the element.
- **allow** means that the user can see and work with an element, but not change its ACL

- **full** means that the user not only can see and work with an element, but can also set the ACL on that element.
- **deny** overrides **allow** and **full**, and if there is no explicit **allow**, a user is denied.

In this case, you want to grant the user from Partnership_1 (“**part_1**”) **allow** privileges on the **partner_1** files, but **deny** that user any privileges on the **partner_2** files or the **acme_proprietary** files. Likewise, the user from Partnership_2 (“**part_2**”) should have **allow** privileges to the **partner_2** files, but **deny** privilege to the **partner_1** and **acme_proprietary** files.

Note: It is conceivable that partners might want **full** access to their own files so that they could set specific ACLs for their own users. However, the **full** privilege is very powerful and should not be granted unless absolutely necessary. A partner with **full** access could deny you access to files on your own system, or inadvertently **deny** access to everybody, both of which would require superuser intervention.

Both partners should retain **allow** access to the Acme **common_files** folder and its contents.

Finally, you (**acme_1**) want to have **full** privileges on all elements in the project. To set up these ACLs:

1. Log in to the AccuRev CLI with the **acme_1** superuser account.

```
> accurev login acme_1
```

2. Change to the workspace directory for the **prod_3000_devel** project and grant yourself (**acme_1**) full access to all project files.

```
>accurev eacl -a acme_1:full -R prod_3000_devel_files
```

Note: Since we want to leave the default **all:allow** ACL in place for all the files in the project, and explicitly adjust the ACLs of specific files and specific users, it is critical that we use the **-a** (add) option here, and not **-n** (new). If we had specified **-n**, user **acme_1** would have been granted **full** access, while *removing* access for everybody else. Removing access is the same as specifying **deny**.

3. Explicitly deny both partners any access to the Acme proprietary elements:

```
>accurev eacl -a part_1:deny,part_2:deny -R prod_3000_devel_files\acme_proprietary
```

Again, note the use of **-a**, which retains the default **allow** access to these elements for other users (such as other Acme employees). Using **-n** here to set **deny** access to the partners would have been the same as setting **all:deny**, since it would have removed the default **all:allow** setting.

4. Deny Partnership_1 any access to Partnership_2 files, while granting **allow** access to them to the Partnership_2 user (and retaining the default **all:allow** access to all other users).

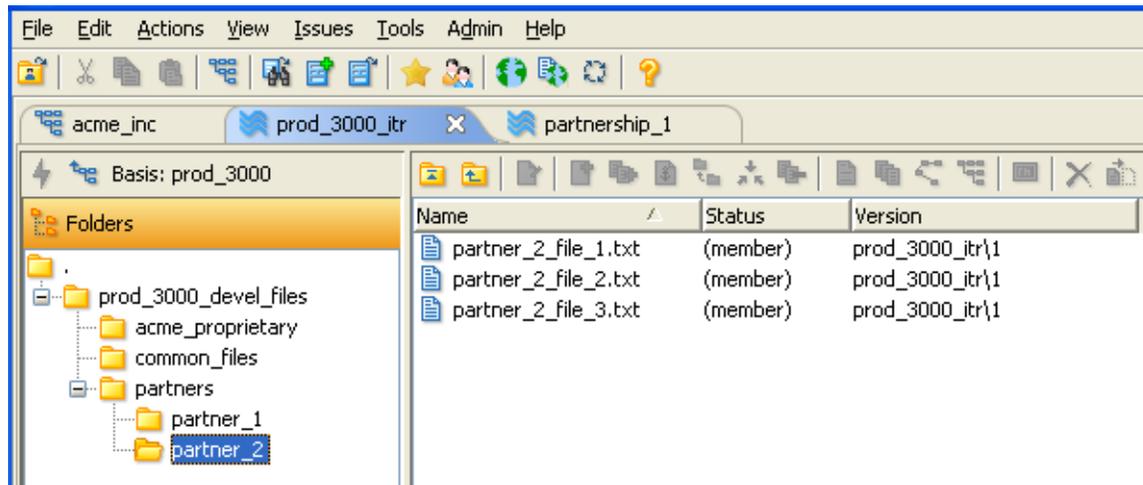
```
>accurev eacl -a part_2:allow,part_1:deny -R
prod_3000_devel_files\partners\partner_2
```

5. Deny Partnership_2 any access to Partnership_1 files, while granting **allow** access to them to the Partnership_1 user (and retaining the default **all:allow** access to all other users).

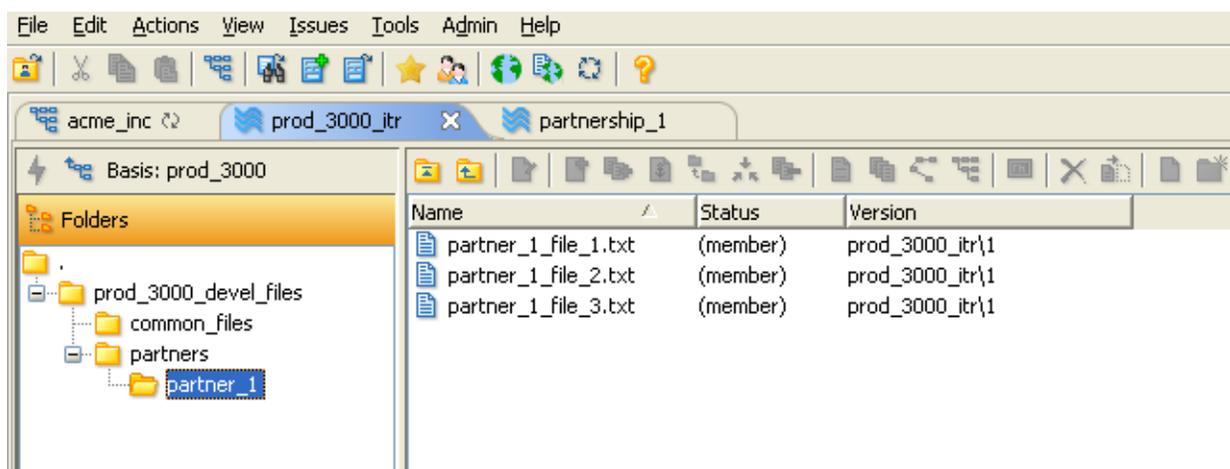
```
>accurev eacl -a part_1:allow,part_2:deny -R
prod_3000_devel_files\partners\partner_1
```

Once you ([acme_1](#)) promote all these elements up the stream, the effect of the ACLs becomes obvious. **Note:** The *promote* has nothing to do with setting EACLs; it is just that the files were being set up for the first time in [acme_1](#)'s workspace and do not even exist in the stream hierarchy until they are promoted.

When you ([acme_1](#), who has *full* access to all elements) view the contents of the `prod_3000_itr` stream, all folders and elements are visible. In the following example, the elements in `partner_2` are displayed



However, if user `part_1` logs in and views the same stream, he or she sees a very different display. Only the `partner_1` files and the common Acme files are visible. From user `part_1`'s perspective, the `partner_2` and Acme proprietary files and folders do not even exist.



User `part_2` would see a similar display, except the `partner_2` files and folder would be visible, while those of `partner_1` could not be seen.

Note: Although the ACL setting is not visible from the GUI, an ACL list from the CLI will reveal that while user `acme_1` has *full* access to the `partner_1` elements and user `part_2` has no access to the `partner_1` elements, user `part_1` has *allow* access to the Acme common files through the default *all* ACL:

```
>accurev eacl common_files
```

```

/prod_3000_devel_files/common_files
  all:allow
  acme_1:full

>accurev eacl partners\partner_1
/prod_3000_devel_files/partners/partner_1
  all:allow
  acme_1:full
  part_1:allow
  part_2:deny

```

More Restrictive Scenario

This scenario is similar to the less restrictive scenario above, except that the default privilege for all users to all files will be *deny*. We will then adjust ACLS to increase privileges as needed.

Similar to the previous scenario, we will start by giving user `acme_1` full access to all project files, but this time we will use the *-n* option instead of *-a*:

1. Log in to the AccuRev CLI with the `acme_1` superuser account.

```
> accurev login acme_1
```

2. Change to the workspace directory for the `prod_3000_devel` project and use the *-n* option to grant yourself (`acme_1`) full access to all project files, replacing the default *all:allow* ACL.

```
>accurev eacl -n acme_1:full -R prod_3000_devel_files
```

Note: Unlike the previous scenario which used *-a*, all users except `acme_1` are now denied access to all elements in the project by default.

3. Explicitly grant the user from each partnership *allow* access to the top of the project tree (but do not recurse through the tree structure, since we want to be selective about what they can and cannot access):

```
>accurev eacl -a part_1:allow,part_2:allow prod_3000_devel_files
```

4. Explicitly grant the Partnership_1 user (`part_1`) *allow* access to the Partnership_1 files and to the Acme common files. Note that unlike the previous example, it is necessary now to explicitly grant allow access to the common files, since the *all:allow* default ACL was implicitly removed in Step 2.

```
>accurev eacl -a part_1:allow -R
  prod_3000_devel_files\partners\partner_1
>accurev eacl -a part_1:allow -R prod_3000_devel_files\common_files
```

5. Likewise, explicitly grant the Partnership_2 user (`part_2`) *allow* access to the Partnership_2 files and to the Acme common files.

```
>accurev eacl -a part_2:allow -R
  prod_3000_devel_files\partners\partner_2
>accurev eacl -a part_2:allow -R prod_3000_devel_files\common_files
```

Because the default setting in this scenario is to deny access to users who have not been explicitly granted access, the above steps only address users `acme_1`, `part_1`, and `part_2`. You would need to continue these procedures to grant access to other Acme users as needed. In a real-life scenario, you would typically assign users to groups and assign ACLs to those groups rather than to individual users.

Restricting Access to Commands using Triggers

By default, any registered AccuRev user can execute any AccuRev command. Many organizations wish to restrict users' access to certain commands, such as the ability to maintain users, groups, and passwords, the ability to lock streams and create ACL permissions, and so on. Triggers provide a flexible mechanism for implementing command-based security.

Many AccuRev commands can be configured to “fire a trigger”. This causes a user-defined script to execute either:

- before the command executes (pre-operation trigger), or afterward (post-operation trigger)
- on the client machine, or on the server machine

A pre-operation trigger can affect the execution of the command or cancel it altogether. Typically, a security-related trigger checks the identity of the user invoking the command, then decides whether or not to allow the command to proceed.

Some triggers are configured on a per-depot basis, using the *mktrig* command. These triggers monitor individual commands (*add*, *keep*, and *promote*). Three are pre-operation triggers that fire on the client machine; one is a post-operation trigger that fires on the server machine.

Other triggers are configured, on a per-depot or whole-repository basis, by placing a script in a well-known location on the server machine. These triggers monitor groups of commands, rather than individual commands.

For more on triggers, see the chapter *AccuRev Triggers* on page 71.

Which Security Feature Should I Use?

AccuRev's security features overlap to a considerable extent. For example, when a user invokes the *promote* command, any of these mechanisms can prevent the command from proceeding:

- a lock on the source or destination stream
- an ACL permission on the destination stream, on a higher-level stream, or on the entire depot
- a *pre-promote-trig* trigger script, which runs on the client machine
- a *server_preop_trig* trigger script, which runs on the server machine

How do you decide which feature to use in a given situation? There are no absolute rules, but here are some guidelines:

To script or not to script?

The trigger mechanism depends on execution of user-supplied scripts, written in Perl, Python, or some other scripting language. There's a trade-off: scripting required development time and significant expertise, but is infinitely flexible.

In many situations, you may be able to use the AccuRev software distribution's sample Perl scripts, which are designed for fill-in-the-blanks customization.

It makes sense to implement as much security as possible with locks and ACL permissions (and perhaps the sample trigger scripts), before delving into original trigger scripting.

Locks vs. ACL permissions

Roughly speaking, a lock controls the placing of data *into* a stream, whereas an ACL permission controls the reading of data *from* a stream. (There are plenty of exceptions to this general rule.)

Both locks and permissions can be targeted at specific users or groups. The ACL is more flexible: you can create any number of permissions for the same stream, but only limited number of locks.

Running trigger scripts: client machine vs. server machine

Running trigger scripts on the client machine conserves networking and server resources. But keep in mind that all client machines must have copies of the scripts (or must have network access to a central location where scripts are kept).

Running trigger scripts on the server machine provides administrative simplicity and centralized logging.

AccuRev Triggers

A trigger is a callback built into certain AccuRev commands. When a user enters the command, the corresponding trigger causes a user-defined or built-in procedure to be performed just before or after the command executes. Typically, a user-defined procedure is implemented as a script in the Perl scripting language. Sample Perl scripts are available in the **examples** subdirectory of the AccuRev installation directory.

Note: in this chapter, “trigger script” refers to any executable program, written in any language, that is executed when a trigger fires.

AccuRev supports pre-operation triggers, post-operation triggers, and triggers that integrate issue management into AccuRev’s configuration management. The latter category of triggers have pre- and post-operation components.

Some triggers are set with the *mktrig* command; others are set by placing the script at a special location; yet others are set through the Schema Editor in the GUI.

Pre-Operation Triggers

The following triggers execute a procedure before the user-requested command executes. Each of these triggers has the ability to cancel execution of the user’s command. (See *Trigger Script Exit Status* on page 91.) Some of the triggers fire on the client machine, and others on the server machine. It’s possible for a single command (e.g. *keep*) to cause triggers to fire both on the client and on the server.

Client-Side Triggers

The following pre-operation triggers fire on the client machine:

- *pre-create-trig*: fires on the client machine prior to execution of an *add* command (GUI: *Add to Depot*). It does not fire for an *In* command (GUI: *Paste Link*), which creates a link element.)

The trigger script must specify the element type (directory, text, binary, or ptext) of each element to be created by the command. This overrides the element type specified with the *add -E* option.

- *pre-keep-trig*: fires on the client machine prior to execution of a *keep* command.
- *pre-promote-trig*: fires on the client machine prior to execution of a *promote* command or a *purge* command (GUI: *Revert to Backed Version* or *Revert to Most Recent Version*).

Server-Side Triggers

The following pre-operation triggers fire on the server machine.

- *server_admin_trig*: fires on the server machine prior to execution of certain commands. This is a repository-wide trigger — it fires no matter what depot, if any, the user’s command applies to. The following commands cause *server_admin_trig* to fire:

addmember	chstream	mkdepot	mkws, mkref	rmws
authmethod	chuser, chgroup	mkgroup	reactivate	setacl
chdepot <i>(see note)</i>	chws <i>(see note)</i>	mkreplica	remove	unlock
chpasswd	ismember	mkstream	rmmember	<u>defcomp</u>
chref	lock	mktrig	rmreplica	replica_sync
chslice	lsacl	mkuser	rmtrig	write_schema

Note: when the *chdepot* command renames a depot, *server_admin_trig* fires twice: once for the renaming of the depot object (chdepot), then again for the renaming of the depot's root-stream object (chstream). Similarly, when the *chws* command renames a workspace, *server_admin_trig* fires twice: once for the renaming of the workspace object (chws), then again for the renaming of the workspace stream object (chstream).

The last three commands are not standard AccuRev CLI commands:

- The *defcomp* command is not visible to the user; it is used in the implementation of the include/exclude facility CLI commands *incl*, *excl*, *includo*, and *clear*.
- The *replica_sync* command recognized by the *server_admin_trig* trigger corresponds to the *replica_sync* command in the CLI.
- The *write_schema* command is generated by the AccuRev GUI when the *Save* button is clicked in the Schema Editor.
- *server_preop_trig*: fires on the server machine prior to execution of certain commands. This is a depot-specific trigger — it fires only for commands that operate on the depot(s) where the trigger has been activated. The following commands cause *server_preop_trig* to fire:

add	defunct	promote	replica archive
archive	keep	purge <i>(see note)</i>	revert
co	ln	reclaim	unarchive

Note: *purge* is equivalent to the GUI command *Revert to Backed*.

For *add* or *keep*, the *server_preop_trig* script can specify the exclusive file locking state (parallel or serial) of the element(s) processed by the command. This overrides any specification made with the *-E* command-line option.

The *server_admin_trig* and *server_preop_trig* triggers are independent of each other and are fired by different sets of commands — for a given command, only one of these triggers will fire. But these triggers can fire in addition to the triggers enabled with the *mktrig* command (*pre-create-trig*, *pre-keep-trig*, *pre-promote-trig*, and *server-post-promote-trig*) and the *server_dispatch_post* trigger.

Post-Operation Triggers

The following triggers execute a procedure after the user-requested command executes successfully. If the user's command fails, the post-operation trigger does not fire. A post-operation trigger always fires on the server machine.

- ***server-post-promote-trig***: fires on the server machine subsequent to execution of a ***promote*** command.
- ***server_dispatch_post***: fires on the server machine each time an AccuWork issue record is created or modified. This trigger is intended to enable email notification to interested users. A sample Perl script is available in the **examples/dispatch** subdirectory of the AccuRev installation directory.

Triggers in a Replication Environment

See *Triggers and Replication* on page 49.

Triggers and Security

See *The 'server_auth_trig' Script* on page 57.

Transaction-Level Integration Trigger

You can achieve tight coordination of your organization's configuration management and issue management capabilities by enabling one or both of the integrations between AccuRev's configuration management and issue management facilities. The transaction-level integration is enabled by a trigger on a depot-by-depot basis:

```
accurev mktrig -p WidgetDepot pre-promote-trig client_dispatch_promote
```

The "client_dispatch_promote" integration routine is built into the AccuRev software — no scripts are required — and includes both pre-operation and post-operation components:

1. On the client machine, a user invokes the AccuRev ***promote*** command.
2. The pre-operation part of the trigger fires on the client machine, prompting the user to specify one or more AccuWork issue records (SPACE-separated). If this part of the trigger fails (e.g. the user specifies a non-existent issue record), the ***promote*** command itself is cancelled.
3. The ***promote*** command completes, and is recorded in the AccuRev repository as a transaction.
4. The post-operation part of the trigger fires on the server machine, updating the issue record that the user specified by adding the number of the ***promote*** transaction to the **affectedFiles** field.

If you use the built-in "client_dispatch_promote" integration routine as the **pre-promote-trig** trigger, you must not also set a **server-post-promote-trig** trigger. Doing so would suppress the post-operation component of the "client_dispatch_promote" routine.

For information on handling this situation and other aspects of customizing the transaction-level integration, see *Implementation and Customization of the Transaction-Level Integration* on page 98. (This section also describes another integration between configuration management and issue management, which works at the change-package level instead of the transaction level.)

Note: the *purge* command (GUI command *Revert to Backed*) also triggers this integration, because it uses the depot's *pre-promote-trig* trigger capability.

Preparing to Use an AccuRev-Provided Trigger Script

Sample trigger scripts are installed with AccuRev, in the **examples** subdirectory. These sample scripts are implemented in the platform-neutral Perl scripting language. Use the following procedure to install and use one of these scripts:

1. **Install Perl.** There are many sources on the Web for Perl. We recommend the ActivePerl distribution from <http://www.activestate.com>. This distribution includes a conversion utility, *pl2bat*, which makes a Perl script executable under Windows, by embedding the Perl code in a Windows batch file (**.bat**).

Be sure to install Perl on all appropriate machines. Note that some pre-operation triggers run on the client machine, while others run on the server machine. All post-operation triggers run on the AccuRev server machine.

For faster processing on large XML files, try installing the **XML::SAX::ExpatXS** Perl module and using that module in your Perl scripts.

2. **Get a copy of the sample script.** Copy the sample script from the *examples* subdirectory of the AccuRev installation directory to an AccuRev workspace. Then use the *add* command to place the script under version control.
3. **Prepare the script.** Open the script in a text editor, and customize the script according to the instructions included as comment lines. Before embarking on complex script customization, be sure to read *The Trigger Parameters File* on page 76.

Note: the path to the AccuRev executable in the script must not contain unquoted spaces. Either properly quote the spaces for Perl, or use short filenames on Windows platforms (for example, **progra~1** if **accurev.exe** is located in the **Program Files** directory).

4. **Enable the trigger.** Enable the trigger, either with the *mktrig* command or by placing the script in the proper location. See *Enabling a Trigger* below for details.

Enabling a Trigger

Depending on its type, an AccuRev trigger is enabled in one of these ways:

- Executing an *accurev mktrig* command, specifying the location of the script. AccuRev simply records the location you specify in the repository; it doesn't make a copy of the script. Make sure that no one moves it!
- Placing the executable script file in the location prescribed for that type of trigger.

For details, consult the appropriate subsection below:

pre-create-trig, pre-keep-trig, pre-promote-trig, server-post-promote-trig

Use the *mktrig* command to enable use of the script in a particular depot. For example:

```
accurev mktrig -p WidgetDepot pre-keep-trig /usr/ac_scripts/addheader
```

The *-p* option isn't necessary if your current directory is in a workspace associated with that depot. When the trigger fires, AccuRev will search for the script at the specified pathname (in the example above, */usr/ac_scripts/addheader*).

We strongly suggest specifying an absolute pathname. Otherwise, when the trigger fires, AccuRev will use the user's search path (for **pre-create-trig**, **pre-keep-trig**, or **pre-promote-trig**) or the search path of the AccuRev Server's user identity (for **server-post-promote-trig**) to find the specified script file.

server_admin_trig

Place an executable file in subdirectory **triggers** of the **site_slice** directory:

- UNIX/Linux: the file must be named **server_admin_trig** or **server_admin_trig.pl**
- Windows: the file must be named **server_admin_trig.bat**

Example:

```
C:\Program Files\AccuRev\storage\site_slice\triggers\server_admin_trig.bat
```

server_preop_trig

Place an executable file in subdirectory **triggers** of the slice directory of one or more depots (*accurev show slices* displays slice directory locations):

- UNIX/Linux: the file must be named **server_preop_trig** or **server_preop_trig.pl**
- Windows: the file must be named **server_preop_trig.bat**

Example:

```
/opt/accurev/storage/depots/talon_tests/triggers/server_preop_trig
```

server_dispatch_post

Place an executable file in the AccuRev executables (**bin**) directory on the AccuRev Server machine:

- UNIX/Linux: the file must be named **server_dispatch_post** or **server_dispatch_post.pl**
- Windows: the file must be named **server_dispatch_post.bat**

Note: for compatibility with previous AccuRev releases, the script can also be named **dispatch_email**, with the appropriate suffix.

Example:

```
C:\Program Files\AccuRev\bin\server_dispatch_post.bat
```

Notes on Triggers in Multiple-Platform Environments

If you have a mixed environment where you have both Windows and UNIX/Linux clients accessing the same depot, you can set up triggers that will operate for both.

- Create trigger scripts that have the exact same name for both platforms, but ensure that the Windows version has the appropriate extension (e.g., **check_for_comments.bat**), and that the UNIX/Linux script has *no* extension (**check_for_comments**).
- Test them and ensure that the scripts run correctly on their respective platforms.
- Place both versions of the script in the same directory. (You can place the Windows and UNIX/Linux scripts in separate directories if you need to, but we recommend keeping them together for simplicity and ease of maintenance.) This directory needs to be visible to all clients--it needs to be exported or shared from the server, and mapped or mounted on the clients.
- Adjust the PATH on your clients to point to the proper directory for their platform. For example, the PATH setting “s:\triggers” on Windows clients might be “/mnt1/ac/triggers” on UNIX/Linux machines.
- Execute the **mktrig** command, specifying the script name without a suffix, and without a qualifying path:

```
accurev mktrig -p WidgetDepot pre-keep-trig check_for_comments
```

When called by a Windows client, the trigger script with an extension will get executed. When called by a UNIX/Linux client, the trigger script without the extension will get executed.

Remember to revise *all* versions of a script when you revise any one of them.

The Trigger Parameters File

When a trigger fires and executes a user-supplied script, AccuRev passes two arguments to the script:

- The first argument is the pathname of a flat-text file containing information about the transaction that is about to be performed (or was just completed).
- The second argument is the pathname of an XML-format file containing the same information. (In some cases, detailed below, the XML-format file contains a small amount of additional information that is not contained in the flat-text file.)

Exceptions: only one argument, the pathname of an XML-format file, is passed to a **server_preop_trig** script or a **server_admin_trig** script.

These files are called trigger parameters files. The flat-text file contains a series of values — usually one value per line — in a prescribed order. The XML-format file contains a set of elements below the top-level **<triggerInput>** element. Each element contains the information for

one parameter: the parameter name is the element tag, the parameter value is the element contents (sometimes encoded as a set of subelements). For example, here are two trigger parameters files generated by the same user command:

Flat-text trigger parameters file	XML-format trigger parameters file
pre-create	<triggerInput>
talon	<hook>pre-create</hook>
talon_dvt_john	<depot>talon</depot>
4	<stream1>talon_dvt_john</stream1>
adding some files	<changePackages></changePackages>
this multi-line	<comment>adding some files
comment has	this multi-line
four lines	comment has
C:/wks/talon/dvt_john	four lines</comment>
john	<topDir>C:/wks/talon/dvt_john</topDir>
/tools/cont.sh	<principal>john</principal>
/tools/end.sh	<elemList>
/tools/start.sh	<elem>/tools/cont.sh</elem>
	<elem>/tools/end.sh</elem>
	<elem>/tools/start.sh</elem>
	</elemList>
	</triggerInput>

The information contained in the trigger parameters file varies among the trigger types, as described in the following sections.

Format of the “pre-create-trig” Trigger Parameters File

The following table presents the information in the trigger parameters file sent to a **pre-create-trig** script. This information describes the creation of one or more new elements to a depot (CLI: *add*, GUI: *Add to Depot*).

Note: the trigger fires on creation of a new file or directory element, but not on creation of a new link element (CLI: *ln*, GUI: *Paste Link*).

The order of the parameters in this table reflects the order in which they appear in the flat-text trigger parameters file.

Parameter	Description
hook	Type of trigger: pre-create .
principal	AccuRev username of person invoking the command.
depot	Name of depot targeted by the command.
stream1	Name of the workspace stream in which the new elements are to be created.
topDir	Pathname to the top-level directory of the user’s workspace tree, as it is listed by the <i>show wspaces</i> command.

Parameter	Description
comment	Zero or more comment lines specified by the user (see <i>Encoding of Command Comments</i> on page 90 below).
elemType	(XML-format parameters file only) Integer indicating the element type specified by the user: 0 =none specified, 2 =text, 3 =binary, 4 =ptext.
elemList elements	One or more files/directories to be added to the depot. For general notes, see <i>Encoding of Element Lists</i> on page 90 below.

In the flat-text trigger parameters file, the elements to be created are listed, one per line, at the end of the file:

```
/tools/cont.sh
/tools/end.sh
/tools/start.sh
<-- end-of-file of trigger parameters file
```

There is no need to supply an element count, since an end-of-file condition signals the end of the element list.

In the XML-format trigger parameters file, the element paths are encoded as **<elem>** sub-elements of **<elemList>**:

```
<elemList>
  <elem>/tools/cont.sh</elem>
  <elem>/tools/end.sh</elem>
  <elem>/tools/start.sh</elem>
</elemList>
```

(There is also an **<elements>** element with the same **<elem>** sub-element data.)

Overwriting the ‘pre-create-trig’ Trigger Parameters File

A *pre-create-trig* script must overwrite its flat-text parameters file with data that indicates the type of each element to be created. Each line must describe one new element:

```
<element-pathname> <element-type>
```

... where *<element-pathname>* is a pathname from the input “elemList”, and *<element-type>* is a numeric code:

- 1 directory
- 2 text file
- 3 binary file
- 4 ptext file

Make sure that the element-type value is 1 for every directory in the original list. You can’t change the element-type of a directory! You can however, change among the text-file, binary-file, and ptext-file types. For example, you might override AccuRev’s default classification of file **ReadMe.html** as 2 (text-file), setting the element-type to 3 (binary-file) instead.

See *Controlling the Element Type and Exclusive File Locking State* on page 40 of the *AccuRev CLI User's Guide* for a discussion of element types.

Example: to have an **add** command create two text-file elements, two binary-file elements, and a directory element, a **pre-create-trig** script would replace its flat-text parameters file with this data:

```
/tools/end.sh 2
/tools/icons 1
/tools/icons/end.png 3
/tools/icons/start.png 3
/tools/start.sh 2
```

Note: there is currently no provision for the script to overwrite the XML-format trigger parameters file. The data to be passed to the AccuRev Server must be in flat-text format.

Format of the “pre-keep-trig” Trigger Parameters File

The following table presents the information in the trigger parameters file sent to a **pre-keep-trig** script. This information describes the creation of a new versions of one or more existing elements in a depot (CLI: **keep**, GUI: **Keep**).

The order of the parameters in this table reflects the order in which they appear in the flat-text trigger parameters file.

Parameter	Description
hook	Type of trigger: pre-keep .
principal	AccuRev username of person invoking the command.
depot	Name of depot targeted by the command.
stream1	Name of the workspace stream in which the new versions are to be created.
topDir	Pathname to the top-level directory of the user's workspace tree, as it is listed by the show wspaces command.
comment	Zero or more comment lines specified by the user (see <i>Encoding of Command Comments</i> on page 90 below).
elemType	(XML-format parameters file only) Integer indicating the element type specified by the user: 0 =none specified, 2 =text, 3 =binary, 4 =ptext.
elemList elements	A specification for each new element version to be created. For general notes, see <i>Encoding of Element Lists</i> on page 90 below.

In the flat-text trigger parameters file, the versions to be created are listed, one per line, at the end of the file. Each line contains three specifications:

```
<element-pathname> <version-ID> <element-type>
```

There is no need to supply an element count, since an end-of-file condition signals the end of the element list. For example:

```
/tools/icons/end.png talon_dvt_john/5 3
/tools/icons/end.sh talon_dvt_john/9 2
/tools/icons/start.png talon_dvt_john/2 3
/tools/icons/start.sh talon_dvt_john/13 2
```

In the XML-format trigger parameters file, each version to be created is encoded as an `<elem>` sub-element of `<elemList>`. The element's attributes specify the version-ID (**stream** and **version** attributes) and the element-type (**elemType** attribute). The element pathname is encoded as the contents of `<elem>`.

The following example contains the same data as the flat-text example above:

```
<elemList>
  <elem
    stream="talon_dvt_john"
    version="5"
    elemType="3"/>/tools/icons/end.png</elem>
  <elem
    stream="talon_dvt_john"
    version="9"
    elemType="2"/>/tools/icons/end.sh</elem>
  <elem
    stream="talon_dvt_john"
    version="2"
    elemType="3"/>/tools/icons/start.png</elem>
  <elem
    stream="talon_dvt_john"
    version="13"
    elemType="2"/>/tools/icons/start.sh</elem>
</elemList>
```

In either format, the element-type value can be either **2** (text file), **3** (binary file), or **4** (ptext file). Note that different versions of an element can have different types.

In addition to the `<elemList>` element, the parameters file includes an `<elements>` element, with additional information on each file: its element-ID and the real version-ID of the workspace's *current* version (not the one about to be created):

```
...
<elem
  eid="58"
  ver="3/4"/>/tools/icons/end.png</elem>
...
```

Format of the “pre-promote-trig” Trigger Parameters File

The following table presents the information in the trigger parameters file sent to a **pre-promote-trig** script. This information describes the creation of a new versions of one or more existing elements in a depot (CLI: *promote*, GUI: *Promote*).

Note: the **pre-promote-trig** trigger also fires on execution of a CLI *purge* command (GUI: *Revert to Backed*) — but only when the version is being purged from a dynamic stream, not a workspace.

The order of the parameters in this table reflects the order in which they appear in the flat-text trigger parameters file.

Parameter	Description
hook	Type of trigger: pre-promote .
principal	AccuRev username of person invoking the command.
depot	Name of depot targeted by the command.
stream1	Name of the workspace or stream that the versions are to be promoted from.
action	<i>(XML-format parameters file only)</i> (purge) The character string purge .
option_I	<i>(XML-format parameters file only)</i> The issue number (or numbers, SPACE-separated) specified by the user with the <i>promote -I</i> command-line option.
stream2	<i>(XML-format parameters file only)</i> Name of the stream that the versions are to be promoted to.
topDir	Pathname to the top-level directory of the user’s workspace tree, as it is listed by the <i>show wspaces</i> command.
comment	Zero or more comment lines specified by the user (see <i>Encoding of Command Comments</i> on page 90 below).
elemType	<i>(XML-format parameters file only)</i> Integer indicating the element type specified by the user: 0 =none specified, 2 =text, 3 =binary, 4 =ptext.
elemList elements	One or more files/directories to be promoted. For general notes, see <i>Encoding of Element Lists</i> on page 90 below.

In the flat-text trigger parameters file, the elements to be created are listed, one per line, at the end of the file:

```
/tools/cont.sh
/tools/end.sh
/tools/start.sh
<-- end-of-file of trigger parameters file
```

There is no need to supply an element count, since an end-of-file condition signals the end of the element list.

In the XML-format trigger parameters file, the element paths are encoded as `<elem>` sub-elements of `<elemList>`:

```
<elemList>
  <elem>/tools/cont.sh</elem>
  <elem>/tools/end.sh</elem>
  <elem>/tools/start.sh</elem>
</elemList>
```

In addition to the `<elemList>` element, the parameters file includes an `<elements>` element, with additional information on each element: its element-ID and the real version-ID of the version to be promoted):

```
<elements>
<elem
  eid="51"
  ver="8/13">/tools/cont.sh</elem>
</elements>
```

Overwriting the ‘pre-promote-trig’ Trigger Parameters File

A *pre-promote-trig* script can work in tandem with a *server-post-promote-trig* script, providing customized “before and after” processing around the execution of *Promote* commands:

- The *pre-promote-trig* script overwrites its flat-text triggers parameters file.
- The *first line* of the overwritten parameters file becomes the value of the `<fromClientPromote>` parameter passed to the *server-post-promote-trig* script.

Note: there is currently no provision for a *pre-promote-trig* script to pass data to a *server-post-promote-trig* script by overwriting the XML-format trigger parameters file.

Format of the “server-post-promote-trig” Trigger Parameters File

The following table presents the information in the trigger parameters file sent to a *server-post-promote-trig* script. This information is generated by AccuRev, and describes the *Promote* command that has just executed. The first line of this file provides a mechanism for passing user-specified data from a *pre-promote-trig* script to a *server-post-promote-trig* script. See [Overwriting the ‘pre-promote-trig’ Trigger Parameters File](#) on page 82.

The order of the parameters in this table reflects the order in which they appear in the flat-text trigger parameters file.

Parameter	Description
hook	Type of trigger: server-post-promote .
principal	AccuRev username of person invoking the command.
depot	Name of depot targeted by the command.

Parameter	Description
stream1	Name of the stream that the versions were promoted to.
source_stream	Name of the stream the versions were promoted from.
dest_stream	Name of the stream the versions were promoted to.
comment	Zero or more comment lines specified by the user in the Promote command (see <i>Encoding of Command Comments</i> on page 90 below).
fromClientPromote	A single text line containing the issue number (or numbers, SPACE-separated) specified by the user during the Promote command. If no issue number was specified, the flat-text parameters file contains a blank line and the XML-format file contains an empty element.
transNum	The transaction number of the Promote transaction that just completed.
transTime	The time of the Promote transaction that just completed.
changePackages	(XML-format parameters file only) A set of <changePackageID> subelements, specifying the same information as <fromClientPromote>.
elemList elements	A specification for each version that was promoted. For general notes, see <i>Encoding of Element Lists</i> on page 90 below.

The following example shows the data encoded in <elemList> and <elements>:

```
<elemList>
  <elem
    stream="9"
    version="7"
    elemType="2">/dir00/sub00/file04.txt</elem>
</elemList>
<elements>
  <elem
    eid="8"
    ver="9/7">/dir00/sub00/file04.txt</elem>
</elements>
```

Format of the “server_preop_trig” Trigger Parameters File

The parameters file passed to a *server_preop_trig* script is in XML format:

```
<triggerInput>
  <hook> ... </hook>
  <command> ... </command>
  <principal> ... </principal>
  <ip> ... </ip>
```

```
...
</triggerInput>
```

The set of subelements under the **<triggerInput>** element depends on the user’s command—the **issues** parameter is generated only for the *promote* command, for example. The following table provides a summary. For full details, see the sample **server_preop_trig** script in the **examples** directory in the AccuRev installation area.

Parameter	Description
hook	Type of trigger: server_preop_trig .
command	The user command. See <i>Server-Side Triggers</i> on page 71 for a list of commands that can fire this trigger.
principal	AccuRev username of person invoking the command.
ip	The IP address of the client machine.
stream1	The user’s workspace stream.
stream2	The workspace’s parent (backing) stream. For <i>promote</i> commands, this is the stream being promoted to.
stream3	The new name of the workspace or stream.
depot	Name of depot targeted by the command.
objectType	Name of object targeted by the command.
objectName	Type of object targeted by the command: 1 =reference tree; 2 =workspace; 3 =stream; 5 =user; 6 =group
user	AccuRev username being modified.
newKind	New user kind (dispatch or cm).
newName	New AccuRev username.
fromClientPromote	(<i>promote</i>) The number of the AccuWork issue record entered by the user, when prompted by the transaction-level integration or the change-package-level integration.
changePackagePromote	(<i>promote</i>) A set of <changePackageID> subelements, specifying the change packages (that is, issue records) specified in the user’s command. These forms of the <i>promote</i> command generate a <changePackagePromote> element: <ul style="list-style-type: none"> • <i>promote -I</i> (user specifies issue record number(s) on the command line) • <i>promote</i> (user prompted to specify issue record number(s) by the transaction-level or change-package-level integration) • <i>promote -Fx</i> (user specifies a set of issue records with an XML file) The user can also specify issue record(s) through the AccuRev GUI.

Parameter	Description
issues	<p>(<i>promote</i>) The <issues> element contains <issue> subelements for all issues that are being promoted, whether or not the issue was explicitly selected for the promote (as can be the case when promoting by file, when one or more issues can be implicitly selected). Each <issue> subelement specifies the following information:</p> <ul style="list-style-type: none"> • Issue number • Workflow stage the issue will be associated with after it transitions from its current workflow stage • The issue’s current workflow stage • Whether or not the issue will be complete in the destination stream upon completion of the promote operation
comment	Comment string specified by the user. If the comment spans multiple lines, line-terminators are embedded in the string, but the final line does not have a line-terminator.
elements	<p>(<i>add, co, keep</i>) A set of <elem> subelements, each specifying one element processed by the user’s command. For the <i>add</i> and <i>keep</i> commands, hierType is an <elem> attribute, which indicates the exclusive file locking state specified with the -E command-line option: serial or parallel.</p>
output_file	<p>(<i>XML-format parameters file only</i>) The name of the file that the server_preop_trig script creates and AccuRev reads. See Controlling Element Type and Exclusive File Locking with a “server_preop_trig” Script, below.</p>

Controlling Element Type and Exclusive File Locking with a “server_preop_trig” Script

The trigger parameters file sent to a *server_preop_trig* script contains an **<output_file>** element — for example:

```
<triggerInput>
  <hook>server_preop_trig</hook>
  <output_file>cache/0_0.out</output_file>
  ...
```

The script can create a file at this relative pathname (it doesn’t exist when the trigger fires), in order to control the element type and/or exclusive file locking state of some or all of the elements processed by the user command.

The XML element named `<elements>` in the trigger parameters file contains the data that the script needs to generate the output file— for example:

```
<elements>
  <elem
    count="0"
    eid="0"
    elemType="text"
    hierType="parallel">/dir03/sub05/able.txt</elem>
  <elem
    count="1"
    eid="0"
    elemType="text"
    hierType="parallel">/dir03/sub05/baker.bin</elem>
  <elem
    count="2"
    eid="0"
    elemType="text"
    hierType="parallel">/dir03/sub05/carr.doc</elem>
</elements>
```

For each AccuRev element to be processed, `<elements>` contains information about how the new version of the element will be created — unless the script intervenes. This includes both the element type (`elemType` attribute) and the exclusive file locking state (`hierType` attribute).

Note: `<elemList>` contains a subset of the data in `<elements>`, and can be safely ignored.

Suppose the example code above was passed to the `server_preop_trig` script by the `add` command, which the user invoked to place three files under version control: `able.txt`, `baker.bin`, and `carr.doc`. And suppose that the script decides to specify that:

- Elements `baker.bin` and `carr.doc` are to be placed in the `serial` exclusive file locking state.
- The first version of `baker.bin` is to have the `binary` element type.

In this case, the output file should contain the following code:

```
<elemList>
  <elem count="1" hierType="serial" elemType="binary"></elem>
  <elem count="2" hierType="serial"></elem>
</elemList>
```

Notes:

- The top-level XML element in the output file is `<elemList>`, not `<elements>`.
- Each `<elem>` XML subelement identifies an AccuRev element through the `count` attribute (representing the position on the command-line); no element pathname is required.
- The value of the `hierType` attribute must be either `serial` or `parallel`.

- An `<elem>` is required only for AccuRev elements whose exclusive file locking state is to be changed from the default (or with a *keep* command, to be changed from its existing state). Thus, there need not be an `<elem>` for file `able.txt`, which is to be created in the default locking mode, `parallel`.
- The number of `<elem>`s need not match the number of AccuRev elements being processed by the command; if there are “too many”, the final `<elem>`s are silently ignored; if there are “too few”, the final AccuRev elements get the default processing.
- A *server_preop_trig* script can coexist with a *pre-create-trig* script, both of them making element-type specifications. The *pre-create-trig* script must specify an element type for every new element; this is not a requirement for the *server_preop_trig* script. If both scripts specify an element type for the same element, the *server_preop_trig* script “wins”.

Suppressing Transition Execution on Promote

A workflow can optionally specify a transition to be executed when an issue meeting one or more conditions is promoted into a stream. For example, your workflow might execute a “ready for QA” transition when an issue in the workflow stage Complete is promoted into your integration stream.

You can suppress transition execution:

- Using the `-q` option for the *promote* command. See the *AccuRev® CLI User’s Guide* for more information.
- Using the *server_preop_trig* script, as described here.

The trigger parameters file sent to a *server_preop_trig* script contains an `<output_file>` element — for example:

```
<triggerInput>
  <hook>server_preop_trig</hook>
  <command>promote</command>
  <output_file>cache/0_0.out</output_file>
  ...
```

The script can create a file at this relative pathname (it doesn’t exist when the trigger fires), in order to control transition execution on some or all of the elements processed by the user command.

The XML element named `<issues>` in the trigger parameters file contains the data that the script needs to generate the output file— for example:

```
<issues>
  <issue id="11" destination_stage="WIP" current_stage="NEW"
  workflow="Enhancement" complete="true"/>
  <issue id="12" destination_stage="WIP" current_stage="NEW"
  workflow="Enhancement" complete="false"/>
</issues>
```

The `<issue>` subelement contains information about how the issue will be modified when it is promoted—unless the script intervenes.

Suppose the example code above was passed to the `server_preop_trig` script by the `promote` command, which the user invoked to promote issue 11 and issue 12. And suppose that the script decides to specify that trigger execution for issue 11 should be suppressed. In this case, the output file should contain the following code:

```
<ISSUES>
  <ISSUE id="11" apply_transition="false" />
</ISSUES>
```

The transition for issue 12 is not executed under any circumstances because it is incomplete.

Note: Using `promote -q` overrides any transition execution behavior specified in the `server_preop_trig` script `<output_file>` element.

Format of the “server_admin_trig” Trigger Parameters File

The parameters file passed to a `server_admin_trig` script is in XML format: The set of subelements under the `<triggerInput>` element depends on the user’s command. The following table provides a summary. For full details, see the sample `server_admin_trig` script in the `examples` directory in the AccuRev installation area.

Parameter	Description
hook	Type of trigger: <code>server_admin_trig</code> .
command	The user command. See <i>Server-Side Triggers</i> on page 71 for a list of commands that can fire this trigger.
principal	AccuRev username of person invoking the command.
ip	The IP address of the client machine.
stream1	The stream targeted by the user command.
stream2	The parent (backing) stream of <code>stream1</code> .
stream3	(chws, chstream) The new name of the workspace or stream.
depot	Name of depot in which the AccuWork issue database resides.
objectName	(remove, rmws, reactivate) Name of object targeted by the command.
objectType	(remove, rmws, reactivate) Type of object targeted by the command: 1 =reference tree; 2 =workspace; 3 =stream; 5 =user; 6 =group
streamType	(mkstream, mksnap) The kind of stream being created by the command: regular , passthru , or snapshot .
user	(chuser, chpasswd) AccuRev username being modified.
group	(ismember, addmember, rmmember) AccuRev group name

Parameter	Description
newKind	(chuser) New user kind (dispatch or cm).
newName	(chuser, chgroup) New AccuRev username or group name.
fromClientPromote	The number of the AccuWork issue record entered by the user, when prompted by the transaction-level integration or the change-package-level integration.
changePackagePromote	A set of <changePackageID> subelements, specifying the change packages (that is, issue records) specified in the user's command.
comment	Comment string specified by the user. If the comment spans multiple lines, line-terminators are embedded in the string, but the final line does not have a line-terminator.
elemList elements	A set of <elem> subelements, each specifying one element processed by the user's command.

Format of the “server_dispatch_post” Trigger Parameters File

The parameters file passed to a *server_dispatch_post* script is in flat-text format. The order of the parameters in the table below is the order in which they appear in the file.

Parameter	Description
hook	Type of trigger: server_dispatch_post .
depot	Name of depot targeted by the command.
stream1	The stream targeted by the user command.
fromClientPromote	Two SPACE-separated fields: <ul style="list-style-type: none"> The number of the transaction that created the previous version of this issue record. (The number 0 indicates that this is a newly created issue record.) The issue number
transNum	The number of the transaction that created this new version of the issue record.
transTime	The time at which transNum was created.
comment	Comment string specified by the user. If the comment spans multiple lines, line-terminators are embedded in the string, but the final line does not have a line-terminator.
principal	AccuRev username of person invoking the command.
elemList elements	A set of <elem> subelements, each specifying one element processed by the user's command.

Encoding of Element Lists

In both kinds of trigger parameters files, each element is listed by its path relative to the depot's top-level directory:

```
/tools/cont.sh
```

The path begins with a slash in order to simplify constructing the element's full pathname on the client machine: just append the given element pathname to the **topDir** pathname (the top-level directory of the user's workspace tree).

In the flat-text trigger parameters file, the elements (or elements-to-be) to be processed by the user command are listed, one per line, at the end of the file:

```
/tools/cont.sh
/tools/end.sh
/tools/start.sh
                                <-- end-of-file of trigger parameters file
```

(Unlike the set of comment lines, there is no need to supply an element count; an end-of-file condition signals the end of the element list.)

- For **pre-create-trig** and **pre-promote-trig**, the element pathname appears alone on the line.
- For **pre-keep-trig**, each element is followed by the version-ID of the version about to be created (with the workspace specified by name), followed by the element-type code:

```
/dir07/sub04/file02.txt rack_dvt_john/3 2
```

As always the element-type coding is: 1=directory, 2=text file, 3=binary file, 4=ptext file.

- For **server-post-promote-trig**, each element is followed by the real version-ID of the promoted version (with the workspace specified by number), followed by the element-type code:

```
/doc/Chapter_03.rtf 9/7 4
```

In the XML-format *server_preop_trig* trigger parameters file, the element paths are encoded as **<elem>** sub-elements of the **<elemList>** element:

```
<elemList>
  <elem>/tools/cont.sh</elem>
  <elem>/tools/end.sh</elem>
  <elem>/tools/start.sh</elem>
</elemList>
```

Encoding of Command Comments

In the flat-text trigger parameters file, the user's comment is indicated by a line-count (0 or greater), followed by the lines of the comment, if any:

```
4                                <-- number of comment lines to follow
adding some files
this multi-line
```

```
comment has
four lines
```

In the XML-format trigger parameters file, the user's comment is encoded as the contents of the `<comment>` element: a single string. For a multi-line comment, this string has line-terminators embedded:

```
<comment>adding some files <-- embedded line-terminator
this multi-line           <-- embedded line-terminator
comment has                <-- embedded line-terminator
four lines</comment>
```

Note that the final line-terminator is automatically stripped from all comment strings.

The sample set of trigger scripts includes a Perl script for each kind of trigger. The script's comments include a detailed description of the layout of the parameters file for that kind of trigger.

Trigger Script Contents

A trigger script can send email, start a build, update a Web site, or perform many other tasks. In particular, you can run AccuRev commands to get more information. One common use of the *server-post-promote-trig* trigger is to run the *hist* command using the transaction number of the promotion, generating the list of promoted elements for inclusion in an email notification.

Trigger Script Exit Status

The exit status (return value) of a *pre-create-trig*, *pre-keep-trig*, *pre-promote-trig*, *server_preop_trig*, or *server_admin_trig* script is important:

- A zero exit status indicates success: the AccuRev command is allowed to proceed.
- A non-zero exit status indicates failure: the AccuRev command is canceled and the depot remains unchanged.

File Handling by Trigger Scripts

A trigger script can overwrite its parameters file (after reading it, presumably). This provides a way for the script to communicate with the AccuRev command or with a "downstream" script:

- The parameters file for a *pre-keep-trig* script ends with a series of lines, one per element to be kept:

```
<pathname-of-element> <version-ID> <element-type>
```

`<pathname-of-element>` is not a full file system pathname, but starts at the workspace's top-level directory (which is included earlier in the parameters file). `<version-ID>` is the new version to be created for that element. `<element-type>` is the numeric code 1, 2, 3, or 4, as described above. Note that different versions of an element can have different types.

See sample trigger script **addheader.pl** in the **examples** subdirectory of the AccuRev installation directory.

- The parameters file for a *pre-promote-trig* script ends with a series of lines, one per element to be promoted:

<pathname-of-element>

<pathname-of-element> is not a full file system pathname, but starts at the workspace’s top-level directory (which is included earlier in the parameters file).

A *pre-promote-trig* script can overwrite its parameters file, in order to communicate with a *server-post-promote-trig* script: the *first line* of the overwritten parameters file becomes the value of the **from_client_promote** parameter in the *server-post-promote-trig* script.

See sample trigger script **client_dispatch_promote_custom.pl** in the **examples/dispatch** subdirectory of the AccuRev installation directory, along with **server_post_promote.pl** in the **examples** subdirectory. **Note:** The **server_post_promote** trigger is a separate trigger action than the Change-Package-Level Integration which is specifically called only for workspace-to-stream promotes (see *Change-Package-Level Integration* below for more detail).

A trigger script can also send data to STDOUT and STDERR. If the command for which the trigger fired was executed in the AccuRev CLI, this data appears in the user’s command window. If a GUI command caused the trigger to fire, the script’s exit status determines whether the user sees the STDOUT/STDERR data: in the “failure” case (non-zero exit status), the data is displayed in an error-message box; in the “success” (zero exit status) case, the data is discarded.

Trigger Script Execution and User Identities

When a trigger script executes on a client machine, it runs under the identity of the AccuRev user who entered the command. Since the user himself is registered (i.e. has a principal-name) in the AccuRev user registry, there won’t be any authentication problems if the trigger script runs AccuRev commands that access the repository.

When a trigger script executes on the server machine, it runs under the AccuRev user identity of the AccuRev Server itself. Methods for setting an AccuRev username for the AccuRev Server process are described in *AccuRev User Identity of the Server Process* on page 10.

We recommend against using the *login* command to set the AccuRev username in the script itself. You would have to include the password for the AccuRev username in the script; this presents a significant security risk.

‘Administrative Users’ in Trigger Scripts

The sample Perl trigger scripts supplied by AccuRev provide a very simple implementation of the “administrative user” concept: a user is permitted to perform certain operations only if his username is recorded in the **administrator** hash defined in the script:

```
$administrator{"derek"} = 1;
$administrator{"allison"} = 1;
...
if ( ! defined($administrator{$principal}) ) {
```

```
print TIO "Execution of '$command' command disallowed:\n";  
...
```

The Trigger Log File

When a trigger script runs on the AccuRev server machine — for a *server-post-promote-trig*, *server_preop_trig*, or *server_admin_trig* trigger — an invocation line is written to file **trigger.log** in the **logs** subdirectory of the repository's **site_slice** directory:

```
##### [2004/06/28 20:50:42] running: "C:\Program Files\AccuRev\bin\pst_pro.bat" ...
```

If the script produces console output (STDOUT and/or STDERR), this output is also sent to the **trigger.log** file. As with other server log files, the **trigger.log** file is “rotated” periodically, to keep active logs from growing too large.

Integrations Between AccuRev and AccuWork

This section describes two similar, but separate facilities that integrate AccuRev's configuration management functionality with its issue management (AccuWork) functionality. Both integrations record information about the *Promote* transaction in a user-specified AccuWork issue record. One of them uses an issue record's change package as the point of integration; the other uses a particular field in the issue record as the point of integration.

Change-Package-Level Integration

When a *promote* command is executed in a user's workspace (but not in a higher-level dynamic stream), the change-package-level integration records all the promoted versions on the Changes subtab of a user-specified AccuWork issue record.

Enabling the Integration

The change-package-level integration is enabled on a depot-by-depot basis. Open the AccuWork Schema Editor for a particular depot's issue database, and go to the Change Packages subtab.

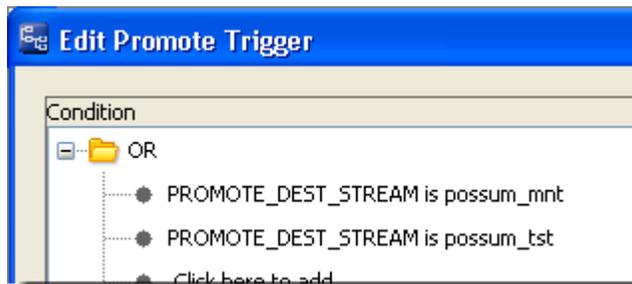
format for *Show Issues* and *Show Diff by Issues* results

one of these queries is executed when a *Promote* happens

and its results are displayed in this format

Filling in the lower section, “Change Package Triggers”, enables the integration for that particular depot. This section is structured as a set of condition/query pairs. One of the queries will be selected for execution at *promote*-time. If the first condition in the left column is satisfied, the first query in the right column will be executed; otherwise the second condition will be evaluated, and if it’s satisfied, the second query will be executed; and so on.

Each clause of a condition performs a test on the *Promote* destination stream. For example, this condition is satisfied if the user is promoting to either of the streams **possum_mnt** or **possum_tst**:



The query corresponding to each condition can be any AccuWork query, which selects a set of issue records.

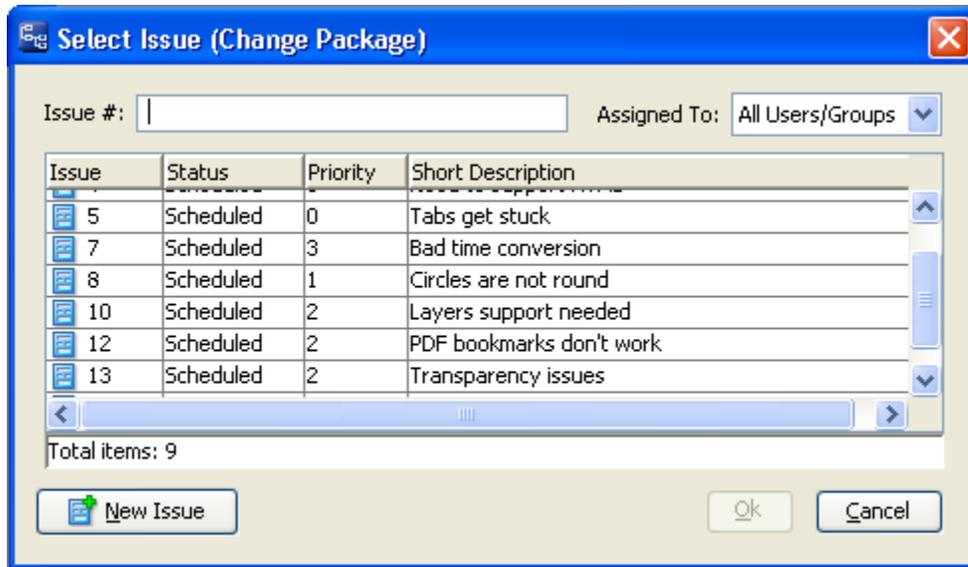
The Change Package Triggers section also specifies the format of each query’s results table — the fields to appear as columns, the column widths, the order of the columns (fields), and the sort order of the rows (records).

Triggering the Integration

Once the integration is enabled for a depot, it is triggered whenever a user performs a *promote* command in a workspace associated with that depot.

If the *Promote* command is invoked through the AccuRev GUI:

1. One of the AccuWork queries specified in the Change Package Triggers section is executed, selecting a certain set of issue records.
2. The GUI displays those records to the user in the results table format specified in the Change Package Triggers section.
3. The user selects one or more of the issue records. There is also a *New Issue* button, which enables the user to create a new issue record “on the fly”.
4. The command completes its work.
5. The versions involved in the command are recorded in the change package section (Changes page) of the selected issue record(s).



If the *promote* command is invoked through the AccuRev CLI:

1. Just as with the GUI, one of the AccuWork queries specified in the Change Package Triggers section is executed, selecting a certain set of issue records.
2. The user is prompted to specify an issue record:

```
Please enter issue number ?
```

Users can bypass this prompt by specifying an issue number with the *-I* option:

```
> accurev promote -I 4761 chap01.doc
Validating elements.
Promoting elements.
Promoted element \.\doc\chap01.doc
```

You can specify multiple issue records with the *-I* option like this:

```
... -I "4761 4795 5006" ...
```

Note: attempting to update an existing change package entry might cause a “change package gap” or “change package merge required” situation, either of which cancel the *promote* command. For example:

```
Promoting elements.
Change package gap: /doc/chap01.doc
```

You can handle a change package gap by adding the *-g* option to the *promote* command. This combines the new and existing change package entries by “spanning the gap”:

```
> accurev promote -I 4761 -g chap01.doc
Validating elements.
Promoting elements.
Promoted element \.\doc\chap01.doc
```

There is no way to have the *promote* command automatically handle a “change package merge required” situation. You must either perform a merge on the element to be promoted, or remove the existing change package entry for that element.

For more on “change package gap” and “change package merge required” situations, see *Updating Change Package Entries* of the *AccuRev Concepts Manual*.

3. Assuming the user-specified issue record is one of those selected by the query, the command completes its work, and the promoted versions are recorded in the change package section of the selected issue record.

What happens if the user specifies no issue record or a non-existent issue record, such as 99999 or 0? In both the GUI and CLI cases:

- If the user enters “0” (or equivalently, makes a blank or non-numeric entry), AccuRev checks whether issue record #0 is among the issues selected by the query executed in Step 1.

Note: the query *can* select issue record #0, even though it doesn’t exist — for example with this clause:

```
issueNum equal to 0
```

- If the query does select issue record #0, the user’s command completes but no information is sent to the issue database. This provides a way for the user to bypass the integration when performing the *promote* command.
- If the query does not select issue record #0, the user’s command is cancelled, and no information is sent to the issue database.
- If the user specifies a non-existent issue record, such as “99999”, the command is cancelled, and no information is sent to the issue database.

Transaction-Level Integration

The integration between configuration management and issue management at the transaction level records the number of each *promote* transaction in a particular field of a user-specified issue record.

Enabling the Integration

The transaction-level integration is enabled on a depot-by-depot basis, by setting the depot’s *pre-promote-trig* trigger. For example:

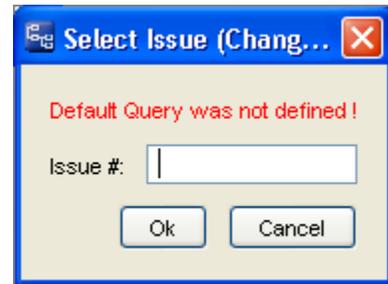
```
accurev mktrig -p kestrel pre-promote-trig client_dispatch_promote
```

Note that “client_dispatch_promote” is simply a keyword, not the name of a script file. The integration is implemented by two cooperating routines, one built into the AccuRev client software, one built into the AccuRev server software.

Triggering the Integration

Once the integration is enabled for a depot, it is activated whenever a user executes the *Promote* command in any workspace or dynamic stream.

1. The depot’s default query, as defined on the Queries tab (*Issues > Queries*), is executed and the results are displayed to the user.
2. The user selects one of the issue records. Note that if no default query is defined for the depot, the user is prompted to type an issue record number.
3. The *promote* command completes its work, propagating the versions to the backing stream.
4. The *promote* transaction number is recorded in the **affectedFiles** field of the selected issue record. (This change to the issue record is, itself, recorded as a transaction, of kind **dispatch**.)



If the user enters “0” or makes a blank entry, the *promote* command completes but no change is made to any issue record. This provides a way for the user to bypass the integration.

Over time, the **affectedFiles** field of a given issue record can accumulate a SPACE-separated list of *Promote* transaction numbers.

Implementation and Customization of the Transaction-Level Integration

When the transaction-level integration is activated, processing takes place on both the AccuRev client machine and the AccuRev server machine:

- The client-side processing — querying the AccuWork issue database and prompting the user to specify an issue record — is structured as a *pre-promote-trig* trigger routine built into the AccuRev client software.
- The server-side processing — updating of the AccuWork issue record — is structured as a *server-post-promote-trig* trigger routine built into the AccuRev server software.

You enable the integration by setting the *pre-promote-trig* trigger with the “client_dispatch_promote” keyword, as described above. You don’t need to explicitly set a *server-post-promote-trig* trigger script.

In most cases, you’ll want to avoid setting a *server-post-promote-trig* trigger script, just letting the built-in routines do their work. But suppose that after a *Promote*, you want the server machine to perform operations in addition to those defined in the transaction-level integration — for example, updating reference trees and sending email messages. In such cases:

1. Create a script that performs the server-side part of the transaction-level integration, along with the desired additional processing. Start with the sample script **server_dispatch_promote_custom.pl**, which is located in the **examples/dispatch** subdirectory of the AccuRev installation directory.
2. Place the script in the AccuRev **bin** directory.
3. Use a **mktrig** command to make the script the depot's **server-post-promote-trig** trigger script.

Further customizations of the transaction-level integration are possible. For example, you might want the user to be able to specify several issue records, not just one. Or you might want to link **promote** commands in one depot with the AccuWork issue database in another depot. Or you might want to update an issue record field other than **affectedFiles**. In such cases, you'll want to dispense with the built-in "client_dispatch_promote" routines altogether:

1. Start with the sample script **client_dispatch_promote_custom.pl** (in the **examples/dispatch** subdirectory), and create your own script for use as a **pre-promote-trig** script to execute on the client.
2. As described above, start with the sample script **server_dispatch_promote_custom.pl** (in the **examples/dispatch** subdirectory), and create your own script for use as a **server-post-promote-trig** script to execute on the server.

If Both Integrations are Enabled

Both the change-package-level and transaction-level integrations can be enabled for a given depot at the same time. In this case, a user performing a **Promote** command in a workspace is prompted to specify an issue record just once, not twice. The prompting for an issue record by the change-package-level integration takes place as usual. That issue record is then updated by both integrations.

Even if both integrations are enabled, a **Promote** command performed in a dynamic stream (not a workspace) activates the transaction-level integration, not the change-package-level integration.

Disabling Triggers

The procedure for disabling a trigger depends on how the trigger was created, as summarized in the following table:

Table 1: Procedures for disabling triggers

If you created the trigger by:	Disable it with:	Comment:
Using the Schema Editor.	The Schema Editor.	Change Package-level integration.

Table 1: Procedures for disabling triggers

If you created the trigger by:	Disable it with:	Comment:
Using the <i>mktrig</i> command.	The <i>rmtrig</i> command.	<p>Applies to:</p> <ol style="list-style-type: none"> 1. Transaction-level integration (or “affectedFiles triggers”) 2. <ul style="list-style-type: none"> <code>pre-create-trig</code> <code>pre-keep-trig</code> <code>pre-promote-trig</code> <code>server-post-promote-trig</code>
Putting scripts in a “triggers” directory.	<p>Rename the script, or the entire directory.</p> <p>Depot specific triggers: <code>.../storage/depots/depot name/triggers</code></p> <p>Global repository (site_slice) triggers: <code>.../storage/site_slice/triggers</code></p>	<p>Server-side triggers</p> <p><i>Depot specific trigger:</i> <code>server_preop_trig</code></p> <p><i>Global repository site_slice trigger:</i> <code>server_admin_trig</code></p> <p>NOTE: If you rename the script, you must modify the name of the script, not just the extension. For example, to disable <code>server_admin_trig</code>, it is NOT sufficient to rename the script from <code>server_admin_trig.bat</code> to <code>server_admin_trig.BAK</code>. You must modify the filename itself, such as <code>server_admin_trig_BAK.bat</code>, otherwise AccuRev will return an error that the trigger <code>server_admin_trig</code> cannot be executed.</p>

The ‘maintain’ Utility

This document describes AccuRev’s **maintain** utility, an administrative tool for occasional use under the guidance of an AccuRev Support Services representative. Before executing a **maintain** command, you must first stop the AccuRev Server.

The **maintain** commands that require a database administrator password will prompt you it. For security reasons, you cannot specify a database administrator password on the command line.

The **maintain** program is located in the AccuRev **bin** directory. If the command-line client program, **accurev**, is on your search path, then so is **maintain**.

Each of the **maintain** commands is described in the next section. Following that are “how to” sections involving use of **maintain** commands.

Specifying a Database Admin Username and Password

Several **maintain** options require a database administrator username and password. For security reasons, you cannot enter a database administrator password on the command line when invoking a **maintain** command. You can only provide a password when prompted by **maintain**, or through the encrypted DB_PASS entry in *ac-install/bin/acerver.cnf*. Here are the rules for providing this information to a **maintain** command:

- If you specify a value for *<db-admin>* when starting the **maintain** command, **maintain** will prompt you to enter the database administrator password.
- If you do not specify a value for *<db-admin>* when starting the **maintain** command, **maintain** will attempt to connect using the DB_USER and the encrypted DB_PASS values specified in *ac-install/bin/acerver.cnf*. In the case of standard installations, the command will fail since, for security reasons, the database role specified by DB_USER typically does not (and should not) have database administrative privileges. You will then have to specify a value for *<db-admin>* on the command line, and be prompted for the password.
- If you try to specify a value for the database administrator password when starting the **maintain** command on the command line, **maintain** will display an error message.

‘maintain’ Command Reference

chpasswd

```
maintain chpasswd <user> <new-password>
```

Changes the password stored in the AccuRev repository for an existing principal-name (named AccuRev user). To remove a user’s password, use two consecutive double-quote characters as the *<new-password>* parameter:

```
maintain chpasswd derek ""
```

chslice

```
maintain chslice <slice-number> <new-location>
```

Changes the location of an existing slice in the repository. Use the *show slices* command to get a listing of slice numbers and their current locations.

Note: this command does not physically move the slice. See the *chslice* command in the *AccuRev CLI User's Guide* for more information.

chuser

```
maintain chuser <user-ID> <new-username>
```

Changes the principal-name (AccuRev username) of an existing user. You specify the user by the unique numeric user-ID, which is immutable. This command is similar to the **accurev chuser** command.

concheck

```
maintain concheck [ site | diag ] [<db-admin>]
```

Tests AccuRev's ability to connect to the database.

When **site** is specified, the command verifies the existence of the AccuRev database as well.

When **diag** is specified, the command performs a series of diagnostics (server initialization, database connection, query execution, and database creation) and displays timing information for each one.

Also see *Specifying a Database Admin Username and Password* on page 99.

dbcheck

```
maintain dbcheck [<depot-name>]
```

Performs certain checks on the whole AccuRev database or a specified *depot-name*:

- Checks the connectivity to the AccuRev database.
- Checks the database-schema version information.
- Checks the access to the **streams** table.
- For each active depot, checks access to the depot schema, the database-schema version, and the monotonicity of transaction numbers and element-IDs.

Note: stop the AccuRev Server before running this command.

dbupgrade

```
maintain dbupgrade [<db-admin>]
```

Upgrades an AccuRev 4.7.x or higher database to an AccuRev 5.x database. This process may take several minutes or several hours, depending upon the size of the database being upgraded.

The dbupgrade wizard will prompt you through the options and values that you need to provide, including database admin name and password, and whether or not to perform a “dry run” upgrade without writing to the database, to evaluate and address the results of a dbupgrade.

During the upgrade, checks and corrects (if necessary) the integrity of information in AccuRev depots. These integrity checks focus primarily on the **trans** (transactions), **trans_entry** (transaction-entries), **ver** (version), **vvir** (virtual-version), and **anc** (ancestry) tables. See “Using the ‘maintain dbupgrade’ Command” in the current Accurev *Installation and Release Notes* for the most up-to-date information about using **dbupgrade**.

Also see *Specifying a Database Admin Username and Password* on page 99.

mkdbuser

```
maintain mkdbuser <db-admin>
```

Makes a database user using the DB_USER and DB_PASS values from the **acserver.cnf** file in the AccuRev **bin** directory. The AccuRev Server uses this username to communicate with the database.

mksite

```
maintain mksite [<db-admin>]
```

This command creates an empty database for use with AccuRev 5.0 and higher.

Also see *Specifying a Database Admin Username and Password* on page 99.

restore

```
maintain restore <backup-file-spec> [<db-admin>]
```

Provides a facility for restoring AccuRev metadata from backup (see *Backup/Restore of the AccuRev Repository* on page 103).

The *<backup-file-spec>* can be either a file name (if the backup was saved to \$BACKUP_LOC) or the full path and name of the backup file.

Also see *Specifying a Database Admin Username and Password* on page 99.

rmdepot

```
maintain rmdepot <depot-name>
```

Removes a depot from the AccuRev repository. All streams, snapshots, and workspace streams are also removed from the repository. (Workspace trees are *not* removed.) For details, see *Removing a Depot from the AccuRev Repository* on page 103.

rmsite

```
maintain rmsite <db-admin>
```

Removes the PostgreSQL database used by AccuRev 5.0 and higher.

Also see *Specifying a Database Admin Username and Password* on page 99.

setcnf

```
maintain setcnf <parameter> [<value>]
```

Writes the <parameter> and <value> specified to the **acserver.cnf** file in the AccuRev **bin** directory in the form **PARAMETER=value**. If the <parameter> is **DB_PASS**, this command encrypts the plain-text password given as well.

server_properties

```
maintain server_properties [ update ]
```

Displays the following server properties from both the **acserver.cnf** file and the AccuRev database: **SITE_SLICE_LOC**, **MASTER_SERVER**, and **PORT**. In addition, the values from each location are compared and the command displays whether or not they match.

When the **update** option is specified, the values that do not match are written from the **acserver.cnf** file into the database.

*Important! Whenever you edit the values for **SITE_SLICE_LOC**, **MASTER_SERVER**, or **PORT** in the **acserver.cnf** file, you **must** also run the **maintain server_properties update** command before the new values will take effect.*

show slices

```
maintain show slices [ all ]
```

For each AccuRev depot, displays the slice number (**Slice#**) and the full pathname to the directory within the repository that stores the data for that depot (**Location**). Use the **all** parameter to include deactivated (removed) slices in the listing. Use the **chslice** command to change the location of a slice.

su

```
maintain su [ { -a | -r } <accountname> ]
```

Add or remove a superuser account for working with element-level security (“EACL”s). If run with no arguments, **maintain su** displays a list of any superusers. See [Element-Level Security \(EACLs\)](#) on page 59 for details.

vercheck

```
maintain vercheck [{ -c | -q } [ -e <eid> ] <depot-name>
```

Checks the storage containers in the specified depot’s **data** directory tree, to verify that a storage container file (**.sto**) exists for each file version recorded in the depot database. It also reports occurrences of “crc mismatch” problems: the actual checksum (CRC) of a **.sto** file does not match the checksum recorded in the corresponding version record.

In addition, you can correct “crc mismatch” problems, using these options:

- **-q** option: for each file with a “crc mismatch”, (step 1) compute the checksum of the **.sto** file, and (step 2) replace the “c:” value in the version record with this newly computed value.

- **-c** option: like **-q**, start with this step for each file: (step 0) change the **.sto** file by removing all its CR characters — that is, all bytes with the value **0x0D**.

You can restrict the processing to versions of a particular element with the **-e** option. **vercheck** fixes all the versions of the element that have a “crc mismatch” problem, leaving other versions as is.

Backup/Restore of the AccuRev Repository

An **accurev** command (**backup**) and a **maintain** command (**restore**) are involved in the scheme for backing up and restoring the AccuRev repository with a minimum of disruption to development activities.

The **accurev backup** command can be executed while the AccuRev Server is running and development activities are ongoing. For more information, see *Backing Up the Repository* on page 3.

At any time after you’ve executed an **accurev backup** command and verified that it has completed successfully, you can restore the repository to its state at the time the **backup** command was started using **maintain restore**. This is an offline procedure — the AccuRev Server must be stopped when you run it. This procedure is documented in *Restoring the Repository* on page 4.

Removing a Depot from the AccuRev Repository

This section describes a procedure for removing a depot completely from the AccuRev repository. Removing a depot:

- Deletes every version of every file and directory in the depot.
- Deletes the entire history of the depot — all transactions involving the depot and its elements.

Removing a depot does not affect any of the workspaces or reference trees that contain copies of the depot’s elements.

Before You Begin

We strongly recommend that you preserve a backup copy of the AccuRev data repository before deleting any depots. See *Backing Up the Repository* on page 3. Much of a depot’s data is stored in its slice of the repository. Use the command **accurev show slices** to determine the pathname of a depot’s slice; you’ll need it in Step 4 below.

Depot Removal Procedure

The following procedure must be performed on the machine where the AccuRev repository resides.

1. Stop the AccuRev Server:

- UNIX/Linux: use the **acservctl** utility, located in the AccuRev **bin** directory:

```
acservctl stop
```
- Windows: use the **Services** control panel, or enter the command **net stop accurev** in a Command Prompt window.

2. Remove the depot using the **maintain** command, which is located in the AccuRev **bin** directory.

```
maintain rmdepot <depot-name>
```

For safety, the **rmdepot** command goes through two confirmation steps, including having you retype the depot name.

The **rmdepot** command removes the depot's records from the repository database (everything except its name and slice information). See Step 5 if you want to rename the deleted depot to reuse the name. It will only appear in an **accurev show -fi depots** command listing (the **-fi** option includes removed items)

3. Restart the AccuRev Server:

- UNIX/Linux: use the **acservctl** utility, located in the AccuRev **bin** directory:

```
acservctl start
```
- Windows: use the **Services** control panel, or enter the command **net start accurev** in a Command Prompt window.

4. Remove the depot's directory from the AccuRev **storage** directory tree. Be careful not to remove any other depot's directory! If you're not sure where this information is located, use the commands **accurev show -fi depots** and **accurev show -fi slices** to determine the pathname. (These commands require the AccuRev Server to be running.)

Reusing a Depot's Name

5. If you want to reuse the depot's name with **accurev mkdepot**, you must first rename the deleted depot with **accurev chdepot -p <depot-name> <new-name>**.

License Management

This chapter describes the AccuRev standard/flexible license manager, and also touches upon the older, proprietary named “keys.txt” license file-based manager used prior to AccuRev Release 4.9.

Note: The named keys.txt license manager is no longer supported as of AccuRev 5.4.

License Types

AccuRev supports three license types:

- *named* -- the traditional, proprietary AccuRev license [**not supported as of AccuRev 5.4**]
- *standard* -- a “floating” license that is valid for one week
- *flexible* -- a “floating” license that is valid for one day (24 hours)

In addition, the standard/flexible options give you the choice using or installing a license server on the local machine, or using an existing license server on a remote machine.

The standard/flexible license manager has been implemented so that changes are generally invisible to users, while new functionality provides more flexibility to AccuRev administrators. These features include:

- support of standard or flexible licenses
- simpler administration of replica server licenses

To Install and Enable Licenses

When you install AccuRev, you specify one of three license options:

- Local Named User Licensing (e.g., keys.txt) [**not supported as of AccuRev 5.4**]
- Local Standard/Flexible Licensing (e.g., accurev.lic)
- Remote Standard/Flexible Licensing

When you choose `Local Named User Licensing` or `Local Standard/Flexible Licensing`, only the selected manager will be configured and enabled, but files for both managers will be installed.

If necessary, you can switch from one license manager to the other by acquiring the appropriate license file and modifying the configuration files.

For example, if you installed `Local Named User Licensing` (e.g., keys.txt) and wish to switch to either standard or flexible licenses:

1. Contact AccuRev to request standard/flexible licenses for your site.

2. Install the new “accurev.lic” file as specified in the instructions you receive with the file (typically, you place this file in `ac-install/bin`).
3. Start the standard/flexible license manager server. The easiest way to do this once is to enter `rlm` on the command line:

```
> ac-install/bin/rlm
```

For the long-term, you will want to configure your system to start the standard/flexible license manager automatically at boot time.

For an example of how to do this on UNIX/Linux, see the example script at `ac-install/extras/unix/accurev_rlm`.

For Windows, you can create a service from the command line:

```
> rlm -install_service -dlog "ac-install\bin\rlm_license.log"
   -service_name "AccuRev License Manager"
```

For details, see “The License Server” section of the on-line *Reprise License Manager End-User Guide* at http://www.reprisesoftware.com/RLM_Enduser.html

4. Edit `ac-install/bin/acserver.cnf` to contain the following lines to specify the standard/flexible license manager:

```
### LICENSE_MANAGER = keys.txt
LICENSE_MANAGER = reprise
LICENSE_SERVER = <license_server_hostname>
LICENSE_PORT = 2375
```

Notes: LICENSE_PORT (default 2375) is the port that AccuRev uses to communicate with the RLM server, and should not be confused with the `ac-port` value (typically 5050) that AccuRev clients use to communicate with the AccuRev server. LICENSE_PORT is the same port that comes with your accurev.lic license file. If you are familiar with the older keys.txt license files, it is important to realize that this is different than the AccuRev PORT value that came with those licenses. For information about changing these port values, see the “License Management” chapter in the *AccuRev Administrator’s Guide*.

The possible values for the LICENSE_MANAGER setting are:

- `keys.txt` — Use the traditional named AccuRev license model (**not supported as of AccuRev 5.4**).
 - `reprise` — Use the newer third-party standard/flexible (accurev.lic) license model.
5. Restart the AccuRev server.

Types of Standard/Flexible Licenses

The standard/flexible license manager supports two product licenses:

- Enterprise, for AccuRev SCM plus AccuWork
- AccuWork-only

In addition, AccuRev provides the option of replica licenses, which allow one of the product licenses above to be used on a replica server.

Finally, the product licenses can be one of two types:

- a seven day *standard* license, for users who need guaranteed access to AccuRev for an extended period of time.
- a one day *flexible* license, which can be used by multiple users needing less-frequent, non-guaranteed access to AccuRev. The duration of the flexible license is for a period of 24-hours.

In this release, you have *either* standard or flexible licenses. You cannot currently mix licenses on the same server.

Replication Server Licenses

When you use the standard/flexible license manager, installation and administration of replication licenses are now automatic — you no longer need to install a license file on the replication server. You just need to calculate how many users you need to support on the master server plus each replica server, plus one additional license for each replica server to be used as a “machine license”. Purchase enough licenses to cover the total number of users who will be accessing AccuRev as well licensing the appropriate number of those users who will be accessing via a Replica server. As long as the license file on the master server covers all of your users, you are all set.

When a user logs into a replica server, the replica server contacts the master server and verifies that the user is permitted to have both a replica license as well as either an Enterprise or an AccuWork license. The master server and replica server negotiate the needed licenses, with no additional attention required by the AccuRev administrator.

The license server ensures that all required licenses are available before checking them out, to ensure that one license does not get used for a replica session that ultimately cannot be successfully started due to unavailability of another license.

Standard Licenses vs. Flexible Licenses

An important feature of the license manager is the ability for a company to select 7-day standard licenses or 24-hour flexible licenses depending on the usage pattern of the end users.

For example, assume that your company employs 100 software developers. If your developers use AccuRev every day to get their job done, you probably would want to get standard licenses which ensure that they will always be able to have access to AccuRev whenever they need it.

However, if your developers typically work on several other projects which may not be under AccuRev control, or if you have other employees such QA or documentation people who may only require occasional access to AccuRev, you may find flexible licenses are more cost-effective.

Rather than purchasing 100 standard licenses to accommodate every user who might ever require access to AccuRev-controlled files, you might instead purchase fewer flexible licenses that can be used as needed. The optimal number of licenses would, of course, need to be determined by you and your AccuRev representative, based on your specific work patterns and requirements.

How Flexible Licenses Work

Flexible and standard licenses are differentiated by a “minimum check-out time” setting. Note that flexible licenses are NOT returned and made available to other users as soon as a user logs out. Nor are standard licenses tied forever to a particular user name.

Instead, standard licenses have a relatively long minimum check-out period and flexible licenses have a relatively short one. At this time, standard licenses have a seven-day check-out period, while flexible licenses have a 24-hour check-out period.

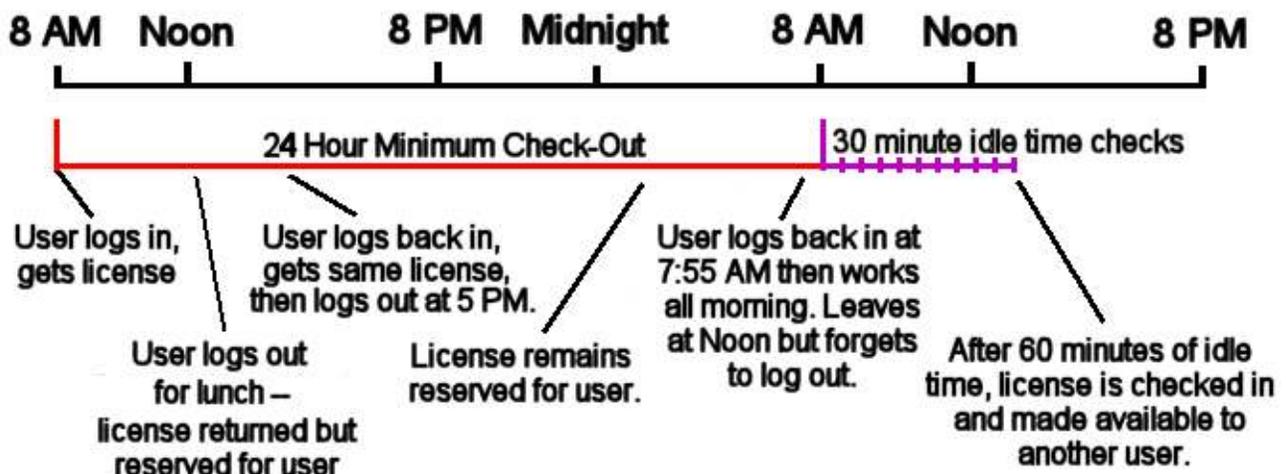
This means that a user who is granted a standard license is guaranteed access to that license for a full week. Even if that user logs out and doesn’t log back in for two days, that license will not be granted to anybody else.

However, a flexible license is guaranteed to its user for a much shorter period — in this example, a single 24-hour day. If that user needs it at 9 AM and is finished with it by 11:00 AM, the license will not become available to another user until 9 AM the next morning. On the other hand, it will be returned and made available to another user after only a single day— much sooner than the week it takes for a standard license to become available to other users.

Note that the minimum check-out time does *not* mean that the license will suddenly return to the available license pool at the end of that period. If you are still actively using the license when the minimum check-out time is reached, your license will be automatically renewed as long as you do not exceed a minimal idle time setting. The minimal idle time is set to 60 minutes by default. During this period, AccuRev sends a “heartbeat” to the license manager every thirty minutes if the user is still active. (User administrators can extend the 60 minute period with RLM administration functions, though they cannot shorten it.) Once the user has been inactive for 60 minutes, the license is checked in and made available to somebody else.

The key concept to note here is: *A license key, whether standard or flexible, will not time-out and be made available to different user as long as the original user is actively working with AccuRev.*

The figure below illustrates how these time-outs affect an AccuRev license. In this example, the times are representative of a flexible license. A standard license would be similar, except the minimum check-out time would be seven days.



In addition to these license manager settings, licenses are affected by the following actions:

- log outs — When the user logs out of AccuRev, the license is checked in. However, if the log out occurs within the minimum check out time, the same license will be returned if that user logs in again.
- session time outs — If the AccuRev session times out, the license is checked in. If and when the user logs in to renew the session, a license is checked out again. This may be a new license, or the original license, depending on the minimum check out time. In any case, no more than one license is checked out to renew the session.
- multiple computers—If a user logs in from different physical computers or clients during the course of the day, the license manager keeps track of this and will not use more than one license to accommodate these multiple log ins.

License Files, Options Files, and Configuration File

A new license file, `accurev.lic`, replaces the traditional `keys.txt` license file. Like `keys.txt`, `accurev.lic` is a simple text file that you receive from AccuRev when you purchase your licenses.

For general information about the contents of this file, please see the third-party documentation referenced in the next section.

For information about where to place the file and any edits that you may need to make to it, please see the AccuRev installation documentation and the documentation that you receive with the license file.

Note that the `acserver.cnf` `LICENSE_PORT` (default value 2375) specifier introduced with the standard/flexible licensing differs from the existing AccuRev `PORT` value (default value 5050) which clients use to connect to the AccuRev host. The proprietary “`keys.txt`” license file specifies the AccuRev `PORT`. The standard/flexible “`accurev.lic`” license file specifies the `LICENSE_PORT` value at the end of the `HOST` line.

If you need to change the standard/flexible `LICENSE_PORT` value, you must make that change not only to the `HOST` line in `accurev.lic`, but also to the `acserver.cnf` configuration file on the server as well. Then you need to restart the AccuRev and license manager servers.

If you need to change the AccuRev `PORT` value under the standard/flexible license manager (which under the proprietary license manager requires a new “`keys.txt`” file), you need to:

- edit the `PORT` line in `acserver.cnf`
- restart the server
- run the **`maintain server_properties update`** command See *The ‘maintain’ Utility* on page 99 for more information.

Then update your clients to connect to the new port.

ISV options and `accurev.opt`

The license manager provides the ability to define an “ISV options” file that allows you to tweak certain parameters, such as extending the idle time allowed before a license gets checked back in, or maintaining a list of users who are prevented from obtaining a license. By default, AccuRev installs an ISV options file named `accurev.opt` in the same folder as the license manager executable (`rlm.exe`), typically `<ac_install>/bin`. This file has a timeout setting line “`TIMEOUT 120`”, to enable automatic idle license check-in.

A common modification of this file is to reserve licenses for specific users so that they stay checked out beyond the usual expiration period.

For example, if you are an administrator and a principal developer with the user name “jsmith” is working on an important project and requires special access to a dedicated license, you can specify a RESERVE line in the `accurev.opt` file such as:

```
RESERVE 1 accurev-ent user jsmith
```

Make sure not to accidentally modify the existing “`TIMEOUT 120`” when you edit `accurev.opt`.

For more information about ISV files, see the third-party documentation listed in the next section.

Multiple AccuRev Servers

With the standard/flexible license manager, it is easy to configure two or more AccuRev Servers to use the same license manager. You simply need to edit the `acserver.cnf` file on each AccuRev server to point to the same license server.

This means, for example, that if you have 100 AccuRev licenses and two AccuRev servers, each AccuRev server could have 50 simultaneous users, or 90 and 10, or 25 and 75, or any other combination that adds up to 100.

Configuration Example

Here are two sample `acserver.cnf` files, showing how to configure two AccuRev servers to share the same AccuRev standard/flexible license server on a machine named “zappa”.

AccuRev Server #1 (a Windows XP laptop named “volman”):

```
C:\Program Files\AccuRev\bin>cat acserver.cnf
MASTER_SERVER = volman
PORT = 5050
SITE_SLICE_LOC = C:\Program Files\AccuRev\storage\site_slice
DEPOTS_DEFAULT = C:\Program Files\AccuRev\storage\depots
# log level 2 or 3 is recommended by accurev support team
LOG_LEVEL = 3
REPLICATION_ENABLED = true
# LICENSE_MANAGER = keys.txt
LICENSE_MANAGER = reprise
LICENSE_SERVER = zappa
LICENSE_PORT = 2375
```

AccuRev Server #2 (a Linux box named “kaylan”):

```
$ cat acserver.cnf

MASTER_SERVER = kaylan
PORT = 4800
SITE_SLICE_LOC = /home/dvanvliet/accurev_49/storage/site_slice
DEPOTS_DEFAULT = /home/dvanvliet/accurev_49/storage/depots
# log level 2 or 3 is recommended by accurev support team
LOG_LEVEL = 2
# if server is started by root, uncomment to run as another user.
# need to 'chmod' all files under SITE_SLICE and DEPOTS_DEFAULT
# when changing users.
#USER = nobody
#GROUP = nobody
LICENSE_MANAGER = reprise
LICENSE_SERVER = zappa
LICENSE_PORT = 2375
```

License Manager Administration Page

The license manager provides a web-based administration tool that you can access at port 5054 on the license server machine. For example:

```
http://<license_server_machine>:5054
```

If you are logged into the license server machine, just use `http://localhost:5054`

See the third-party documentation referenced below for more information.

Note: AccuRev administrators should consider whether they wish to restrict access to this page using network security.

Additional Documentation

The third-party license manager provider makes an *RLM End-User Manual* available on their web site: http://www.reprisesoftware.com/RLM_Enduser.html

In spite of its title, this document is actually targeted towards administrators at end user *sites*, rather than at end users themselves.