**NAME**

auditreduce – merge and select audit records from audit trail files

**SYNOPSIS**

**auditreduce** [ *options* ] [ *audit-trail-file* . . . ]

**AVAILABILITY**

The functionality described in this man page is available only if the Basic Security Module (BSM) has been enabled. See **bsmconv**(1M) for more information.

**DESCRIPTION**

**auditreduce** allows you to select or merge records from audit trail files. Audit files may be from one or more machines.

The merge function merges together audit records from one or more input audit trail files into a single output file. The records in an audit trail file are assumed to be sorted in chronological order (oldest first) and this order is maintained by **auditreduce** in the output file.

Unless instructed otherwise, **auditreduce** will merge the entire audit trail, which consists of all the audit trail files in the directory structure *audit_root_dir*/∗/files (see **audit_control**(4) for details of the structure of the audit root). Unless stated with the -**R** or -**S** option, *audit_root_dir* defaults to **/etc/security/audit.** By using the file selection options it is possible to select some subset of these files, or files from another directory, or files named explicitly on the command line.

The select function allows audit records to be selected on the basis of numerous criteria relating to the record's content (see **audit.log**(4) for details of record content). A record must meet all of the *record-selection-option* criteria to be selected.

**Audit Trail Filename Format**

Any audit trail file not named on the command line must conform to the audit trail filename format. Files produced by the audit system already have this format. Output file names produced by **auditreduce** are in this format. It is:

*start-time*.*end-time*.*suffix*

where *start-time* is the 14-character timestamp of when the file was opened, *end-time* is the 14-character timestamp of when the file was closed, and *suffix* is the name of the machine which generated the audit trail file, or some other meaningful suffix (e.g., **all**, if the file contains a combined group of records from many machines). The *end-time* may be the literal string **not_terminated**, to indicate that the file is still being written to by the audit system. Timestamps are of the form *yyyymmddhhmmss* (year, month, day, hour, minute, second). The timestamps are in Greenwich Mean Time (GMT).

**OPTIONS**

**File Selection Options**

The file selection options indicate which files are to be processed and certain types of special treatment.

–**A** All of the records from the input files will be selected regardless of their timestamp. This option effectively disables the –**a**, –**b**, and –**d** options. This is useful in preventing the loss of records if the –**D** option is used to delete the input files after they are processed. Note, however, that if a record is *not* selected due to another option, then –**A** will not override that.

**–C** Only process complete files. Files whose filename *end-time* timestamp is **not_terminated** are not processed (such a file is currently being written to by the audit system). This is useful in preventing the loss of records if –**D** is used to delete the input files after they are processed. It does not apply to files specified on the command line.

**–D** *suffix*
Delete input files after they are processed. The files are only deleted if the entire run is successful. If **auditreduce** detects an error while reading a file, then that file is not deleted. If –**D** is specified, –**A**, –**C** and –**O** are also implied. *suffix* is given to the –**O** option. This helps prevent the loss of audit records by ensuring that all of the records are written, only complete files are processed, and the records are written to a file before being deleted. Note that if both –**D** and –**O** are specified in the command line, the order of specification is significant. The *suffix* associated with the latter specification is in effect.

**–M** *machine*
Allows selection of records from files with *machine* as the filename suffix. If –**M** is not specified, all files are processed regardless of suffix. –**M** can also be used to allow selection of records from files that contain combined records from many machines and have a common suffix (such as **all**).

**–O** *suffix*
Direct output stream to a file in the current **audit_root_dir** with the indicated suffix. *suffix* may alternatively contain a full pathname, in which case the last component is taken as the suffix, ahead of which the timestamps will be placed, ahead of which the remainder of the pathname will be placed.

If the **-O** option is not specified, the output is sent to the standard output. When **auditreduce** places timestamps in the filename, it uses the times of the first and last records in the merge as the *start-time* and *end-time*.

**–Q** Quiet. Suppress notification about errors with input files.

**–R** *pathname*
Specify the pathname of an alternate audit root directory *audit_root_dir* to be *pathname*. Therefore, rather than using **/etc/security/audit**/∗/files by default, *pathname*/∗/files will be examined instead.

**–S** *server*
This option causes **auditreduce** to read audit trail files from a specific location (server directory). *server* is normally interpreted as the name of a subdirectory of the audit root, therefore **auditreduce** will look in *audit_root_dir*/*server*/files for the audit trail files.

But if *server* contains any '/' characters, it is the name of a specific directory not necessarily contained in the audit root. In this case, *server*/files will be consulted.

This option allows archived files to be manipulated easily, without requiring that they be physically located in a directory structure like that of **/etc/security/audit**.

**–V** Verbose. Display the name of each file as it is opened, and how many records total were written to the output stream.

**Record Selection Options**

The record selection options listed below are used to indicate which records are written to the output file produced by **auditreduce.**

Note: Multiple arguments of the same type are not permitted.

**–a** *date-time*
: Select records that occurred at or after *date-time*. The *date-time* argument is described under **Option Arguments**, below. *date-time* is in local time. The **–a** and **–b** options can be used together to form a range.

**–b** *date-time*
: Select records that occurred before *date-time*.

**–c** *audit-classes*
: Select records by audit class. Records with events that are mapped to the audit classes specified by *audit-classes* are selected. Audit class names are defined in **audit_class**(4). The *audit-classes* can be a comma separated list of *audit flags* like those described in **audit_control**(4). Using the *audit flags,* one can select records based upon success and failure criteria.

**–d** *date-time*
: Select records that occurred on a specific day (a 24-hour period beginning at 00:00:00 of the day specified and ending at 23:59:59). The day specified is in local time. The time portion of the argument, if supplied, is ignored. Any records with timestamps during that day are selected. If any hours, minutes, or seconds are given in *time,* they are ignored. **–d** can not be used with **–a** or **–b.**

**–e** *effective-user*
: Select records with the specified *effective-user.*

**–f** *effective-group*
: Select records with the specified *effective-group.*

**–g** *real-group*
: Select records with the specified *real-group.*

**–j** *subject-ID*
: Select records with the specified *subject-ID* where *subject-ID* is a process ID.

**–m** *event*
: Select records with the indicated *event*. The *event* is the literal string or the *event* number.

**–o** *object_type=objectID_value*
: Select records by object type. A match occurs when the record contains the information describing the specified *object_type* and the object ID equals the value specified by *objectID_value.* The allowable object types and values are as follows:

**file**=*pathname*

Select records containing file system objects with the specified *pathname,* where *pathname* is a comma separated list of regular expressions. If a regular expression is preceeded by a tilda (˜), files matching the expression are excluded from the output. For example, the option file="˜/usr/openwin,/usr,/etc" would select all files in **/usr** or **/etc** except those in **/usr/openwin.** The order of the regular expressions is important because auditreduce processes them from left to right, and stops when a file is known to be either selected or excluded. Thus the option file= **/usr**, **/etc**, **˜/usr/openwin** would select all files in **/usr** and all files in **/etc.** Files in **/usr/openwin** are not excluded because the regular expression **/usr** is matched first. Care should be given in surrounding the *pathname* with quotes so as to prevent the shell from expanding any tildas.

**msgqid**=*ID*
: Select records containing message queue objects with the specified *ID* where *ID* is a message queue ID.

**pid=***ID*
: Select records containing process objects with the specified *ID* where *ID* is a process ID. Note: Process are objects when they are receivers of signals.

**semid=***ID*
: Select records containing semaphore objects with the specified *ID* where *ID* is a semaphore ID.

**shmid=***ID*
: Select records containing shared memory objects with the specified *ID* where *ID* is a shared memory ID.

**sock=***port_number* | *machine*
: Select records containing socket objects with the specified *port_number* or the specified *machine* where *machine* is a machine name as defined in **hosts**(4).

**−r** *real-user*
: Select records with the specified *real-user.*

**−u** *audit-user*
: Select records with the specified *audit-user.*

When one or more *filename* arguments appear on the command line, only the named files are processed. Files specified in this way need not conform to the audit trail filename format. However, −**M**, −**S**, and −**R** may not be used when processing named files. If the *filename* is ‘‘−’’ then the input is taken from the standard input.

**Option Arguments**

*audit-trail-file*
: An audit trail file as defined in **audit.log**(4). An audit trail file not named on the command line must conform to the audit trail file name format. Audit trail files produced as output of **auditreduce** are in this format as well. The format is:

*start-time . end-time . suffix*

*start-time* is the 14 character time stamp denoting when the file was opened. *end-time* is the 14 character time stamp denoting when the file was closed. *end-time* may also be the literal string **not_terminated**, indicating the file is still be written to by the audit daemon or the file was not closed properly (a system crash or abrupt halt occurred). *suffix* is the name of the machine that generated the audit trail file (or some other meaningful suffix; e.g. **all** would be a good suffix if the audit trail file contains a combined group of records from many machines).

*date-time*
: The *date-time* argument to −**a**, −**b**, and −**d** can be of two forms: An absolute *date-time* takes the form:
*yyyymmdd* [ *hh* [ *mm* [ *ss* ]]]

  where *yyyy* specifies a year (with 1970 as the earliest value), *mm* is the month (01-12), *dd* is the day (01-31), *hh* is the hour (00-23), *mm* is the minute (00-59), and *ss* is the second (00-59). The default is 00 for *hh*, *mm* and *ss*.

  An offset can be specified as: **+***n***d** | **h** | **m** | **s** where *n* is a number of units, and the tags **d**, **h**, **m**, and **s** stand for days, hours, minutes and seconds, respectively. An offset is relative to the starting time. Thus, this form can only be used with the −**b** option.

*event*
: The literal string or ordinal event number as found in **audit_event**(4). If *event* is not found in the **audit_event** file it is considered invalid.

*group*
: The literal string or ordinal group ID number as found in **group**(4). If *group* is not found in the **group** file it is considered invalid. *group* may be negative.

*pathname*
: A regular expression describing a pathname.

*user*
: The literal username or ordinal user ID number as found in **passwd**(4). If the username is not found in the **passwd** file it is considered invalid. *user* may be negative.

## EXAMPLES

**praudit**(1M) is available to display audit records in a human-readable form.

This will display the entire audit trail in a human-readable form:

**% auditreduce | praudit**

If all the audit trail files are being combined into one large file, then deleting the original files could be desirable to prevent the records from appearing twice:

**% auditreduce −V −D /etc/security/audit/combined/all**

This will print what user **milner** did on April 13, 1988. The output will be displayed in a human-readable form to the standard output:

**% auditreduce −d 19880413 −u milner | praudit**

The above example may produce a large volume of data if **milner** has been busy. Perhaps looking at only login and logout times would be simpler. The −**c** option will select records from a specified class:

**% auditreduce −d 19880413 −u milner −c lo | praudit**

To see **milner**'s login/logout activity for April 13, 14, and 15 the following is used. The results are saved to a file in the current working directory. Note that the name of the output file will have **milnerlo** as the *suffix*, with the appropriate timestamp prefixes. Note that the long form of the name is used for the −**c** option:

**% auditreduce −a 19880413 −b +3d −u milner −c login_logout −O milnerlo**

To follow **milner**'s movement about the file system on April 13, 14, and 15 the **chdir** record types could be viewed. Note that in order to get the same time range as the above example we needed to specify the −**b** time as the day after our range. This is because **19880416** defaults to midnight of that day, and records before that fall on **0415**, the end-day of the range.

**% auditreduce −a 19880413 −b 19880416 −u milner −m AUE_CHDIR | praudit**

In this example the audit records are being collected in summary form (the login/logout records only). The records are being written to a summary file in a different directory than the normal audit root to prevent the selected records from existing twice in the audit root.

**% auditreduce −d 19880330 −c lo −O /etc/security/audit_summary/logins**

If activity for user ID 9944 has been observed, but that user is not known to the system administrator, then the following example will search the entire audit trail for any records generated by that user. **auditreduce** will query the system as to the current validity of ID 9944, and print a warning message if it is not currently active:

**% auditreduce −O /etc/security/audit_suspect/user9944 −u 9944**

## FILES

**/etc/security/audit/***server***/files/**∗ location of audit trails, when stored

## SEE ALSO

**bsmconv**(1M), **praudit**(1M), **audit.log**(4), **audit_class**(4), **audit_control**(4), **group**(4), **hosts**(4), **passwd**(4)

## DIAGNOSTICS

**auditreduce** will print out error messages if there are command line errors and then exit. If there are fatal errors during the run **auditreduce** will print an explanatory message and exit. In this case the output file may be in an inconsistent state (no trailer or partially written record) and **auditreduce** will print a warning message before exiting. Successful invocation returns **0** and unsuccessful invocation returns **1**.

Since **auditreduce** may be processing a large number of input files, it is possible that the machine-wide limit on open files will be exceeded. If this happens, **auditreduce** will print a message to that effect, give information on how many file there are, and exit.

If **auditreduce** prints a record's timestamp in a diagnostic message, that time is in local time. However, when filenames are displayed, their timestamps are in GMT.

## BUGS

Conjunction, disjunction, negation, and grouping of record selection options should be allowed.