NAME

audit_add_rule_data − Add new audit rule

SYNOPSIS

#include <libaudit.h>

int audit_add_rule_data(int *fd***, struct audit_rule_data ****rule***, int** *fla gs***, int** *action***);"**

DESCRIPTION

audit_add_rule_data adds an audit rule previously constructed with audit_rule_fieldpair_data(3) to one of several kernel event filters. The filter is specified by the *fd* file descriptor and the *flags* argument. Possible values for *flags* are:

- AUDIT_FILTER_USER - Apply rule to userspace generated messages. This is the user filter. Normally all user space originating events are accepted. Rules on this filter are typically written to block specific events.
- AUDIT_FILTER_TASK - Apply rule at task creation (not syscall). This is the task filter. It's normally used to exclude an application from being audited.
- AUDIT_FILTER_EXIT - Apply rule at syscall exit. This is the main filter that is used for syscalls and filesystem watches. Normally all syscall do not trigger events, so this is normally used to specify events that are of interest.
- AUDIT_FILTER_EXCLUDE - Apply rule at audit_log_start. This is the exclude filter which discards any records that match. The action type is ignored for this filter, defaulting to "never".
- AUDIT_FILTER_FS - Apply rule when adding PATH auxiliary records to SYSCALL events. This is the filesystem filter. This is used to ignore PATH records that are not of interest.

The rule's action has two possible values:

- AUDIT_NEVER - Do not build context if rule matches.
- AUDIT_ALWAYS - Generate audit record if rule matches.

RETURN VALUE

The return value is <= 0 on error, otherwise it is the netlink sequence id number. This function can have any error that sendto would encounter.

SEE ALSO

audit_rule_fieldpair_data(3), **audit_delete_rule_data**(3), **auditctl**(8).

AUTHOR

Steve Grubb.