NAME

audit_get_features, audit_set_feature − query or change kernel audit features

SYNOPSIS

#include <libaudit.h>

uint32_t audit_get_features(void);
int audit_set_feature(int *fd***, unsigned** *feature***, unsigned** *value***, unsigned** *lock***);"**

DESCRIPTION

audit_get_features() returns a bitmap describing which kernel audit features are supported. The bitmap is cached internally and retrieved from the kernel on the first call.

audit_set_feature() changes a feature bit for the kernel using the descriptor *fd* which must be an open audit netlink socket. *feature* selects the bit to modify. If*value* is nonzero the feature is enabled, otherwise it is disabled. If *lock* is nonzero the feature setting is locked until reboot.

The feature bits currently defined are:

AUDIT_FEATURE_BITMAP_BACKLOG_LIMIT
Kernel supports changing the backlog queue depth.

AUDIT_FEATURE_BITMAP_BACKLOG_WAIT_TIME
Kernel supports delaying syscalls when the queue is full.

AUDIT_FEATURE_BITMAP_EXECUTABLE_PATH
Kernel will include the executable path on EXECVE records.

AUDIT_FEATURE_BITMAP_EXCLUDE_EXTEND
Exclude rules may be used with more fields than just message type.

AUDIT_FEATURE_BITMAP_SESSIONID_FILTER
Session identifier filtering is supported.

AUDIT_FEATURE_BITMAP_LOST_RESET
Allows resetting the lost event counter.

AUDIT_FEATURE_BITMAP_FILTER_FS
Kernel supports file system field filtering.

RETURN VALUE

audit_get_features returns the feature bitmap or 0 if feature queries are unsupported. **audit_set_feature** returns <= 0 on error, otherwise it is the netlink sequence id number.

SEE ALSO

audit_request_features(3), **audit_reset_lost**(3), **audit_open**(3).

AUTHOR

Steve Grubb