lmgr events

ONTAP 9.11.1 EMS reference

NetApp
February 12, 2024

Table of Contents

lmgr events

lmgr.aggr events

lmgr.aggr.CA.locks.dropped

Severity

EMERGENCY

Description

This message occurs when locks established through CIFS Continuously Available (CA) shares on an aggregate are dropped after takeover, giveback, or aggregate relocation. This causes clients to see a service disruption on the aggregate.

Corrective Action

Confirm that the aggregate is online by using the 'aggr show -aggregate [name of the aggregate] -fields state' command. Then reopen files to reestablish CA locks. For further information about correcting the problem, search the knowledgebase of the NetApp technical support support web site for the "lmgr.aggr.CA.locks.dropped" keyword.

Syslog Message

Continuously Available (CA) locks were dropped on aggregate %s because %s.

Parameters

aggregate (STRING): Name of the aggregate where CA locks were dropped.
reason (STRING): Reason for dropping CA locks, such as "remote lock synchronization is in progress".

lmgr.badstructure events

lmgr.badStructure.found

Severity

ERROR

Description

This event is issued whenever the multi-protocol lock manager finds a bad file pointer and turns off the logging of lock manipulation to preserve trace information. This information can be of use during core analysis.

Corrective Action

(None).

Syslog Message

bad file pointer found, preserving trace information

Parameters

(None).

lmgr.ca events

lmgr.CA.recons.mcc.veto

Severity

ALERT

Description

This message occurs when the system vetoes MetroCluster(tm) switchback or negotiated switchover because one of the volumes affected by the switch has CIFS Continuously Available (CA) locks that are still being reconstructed.

Corrective Action

Retry the switchback or negotiated switchover after verifying that lock reconstruction is complete by using the 'debug locks reconstruction show' diagnostic privilege command.

Syslog Message

Vetoed MetroCluster switchback or negotiated switchover because CA lock reconstruction is in progress.

Parameters

(None).

lmgr.cb events

lmgr.cb.waiting.limit.hit

Severity

ERROR

Description

This message occurs when a lock manager waiting lock callback object cannot be allocated due to reaching the limit for waiting lock callback objects in the associated pool. It is possible there are too many clients holding locks on files where other clients want locks. Clients will have problems accessing or locking files through NAS protocols while at or above the waiting lock callback object limit.

Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The number of waiting callback objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u locks %u/%u owners %u/%u

Parameters

cb_waiting_total (INT): Number of allocated waiting callback objects in the pool.
cb_waiting_limit (INT): Allocated callback waiting objects limit for the pool.
pool_name (STRING): Name of the limit pool.

files_total (INT): Number of allocated files in the pool.
files_limit (INT): Allocated files limit for the pool.
hosts_total (INT): Number of allocated hosts in the pool.
hosts_limit (INT): Allocated hosts limit for the pool.
locks_total (INT): Number of allocated locks in the pool.
locks_limit (INT): Allocated locks limit for the pool.
owners_total (INT): Number of allocated owners in the pool.
owners_limit (INT): Allocated owners limit for the pool.

lmgr.delegations events

lmgr.delegations.limit.hit

Severity

ERROR

Description

This message occurs when a lock manager WAN FlexCache® delegation object cannot be allocated due to reaching the limit for WAN FlexCache delegation objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols using a FlexCache over a WAN to an origin volume on the node generating this message while at or above the WAN FlexCache delegation object limit.

Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The number of WAN FlexCache delegation objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u locks %u/%u owners %u/%u

Parameters

delegations_total (INT): Number of allocated WAN FlexCache delegation objects in the pool.
delegations_limit (INT): Allocated WAN FlexCache delegation objects limit for the pool.
pool_name (STRING): Name of the limit pool.
files_total (INT): Number of allocated files in the pool.
files_limit (INT): Allocated files limit for the pool.
hosts_total (INT): Number of allocated hosts in the pool.
hosts_limit (INT): Allocated hosts limit for the pool.
locks_total (INT): Number of allocated locks in the pool.
locks_limit (INT): Allocated locks limit for the pool.
owners_total (INT): Number of allocated owners in the pool.
owners_limit (INT): Allocated owners limit for the pool.

lmgr.files events

lmgr.files.limit.hit

Severity

ERROR

Description

This message occurs when a lock manager file object cannot be allocated due to reaching the limit for file objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the file object limit.

Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The number of file objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: hosts %u/%u locks %u/%u owners %u/%u.

Parameters

files_total (INT): Number of allocated files in the pool.
files_limit (INT): Allocated files limit for the pool.
pool_name (STRING): Name of the limit pool.
hosts_total (INT): Number of allocated hosts in the pool.
hosts_limit (INT): Allocated hosts limit for the pool.
locks_total (INT): Number of allocated locks in the pool.
locks_limit (INT): Allocated locks limit for the pool.
owners_total (INT): Number of allocated owners in the pool.
owners_limit (INT): Allocated owners limit for the pool.

lmgr.gb events

lmgr.gb.CA.locks.lost

Severity

NOTICE

Description

This message occurs when giveback veto is overridden. Locks established through CIFS continuously available (CA) shares on the aggregate being given back are dropped because the giveback occurs before lock sync completion.

Corrective Action

(None).

Syslog Message

Dropping partner aggregate locks during lock sync on aggregate %s.

Parameters

aggr (STRING): Name of the aggregate for which CA locks are lost.

lmgr.gb.CA.locks.may.lost

Severity

ALERT

Description

This message occurs when giveback veto is overridden. Locks established using CIFS continuously available (CA) shares on the aggregate being given back might be dropped, because giveback commit occurs simultaneously with lock sync completion.

Corrective Action

(None).

Syslog Message

Continuously Available (CA) locks may be lost on aggregate %s due to forced giveback commit occurring simultaneously with lock sync completion.

Parameters

aggr (STRING): Name of the aggregate for which CA locks might be lost.

lmgr.hostfile events

lmgr.hostfile.cleared

Severity

INFORMATIONAL

Description

This message occurs when the volume lock manager hosts file is unusable. In such cases, the hosts files are cleared and reinitialized.

Corrective Action

None.

Syslog Message

Clearing lock manager file %s on volume %s due to error: %s.

Parameters

file_name (STRING): Name of the file that is being cleared.
vol (STRING): Name of the volume.
error (STRING): Error code.

lmgr.hostfile.create.abort

Severity

ERROR

Description

This message occurs when the volume lock manager fails to create the hosts file.

Corrective Action

None.

Syslog Message

Creating lock manager file %s on volume %s failed due to error: %s.

Parameters

aggr (STRING): The name of the volume.
file_name (STRING): The name of the file that was being cleared.
error (STRING): Error code.

lmgr.hostfile.init.failed

Severity

ERROR

Description

This message occurs when the volume lock manager hosts file initialization fails.

Corrective Action

None.

Syslog Message

Initializing lock manager file %s on volume %s failed due to error: %s.

Parameters

aggr (STRING): The name of the volume.
file_name (STRING): The name of the file that is being initialized.
error (STRING): Error code.

lmgr.hosts events

lmgr.hosts.limit.hit

Severity

ERROR

Description

This message occurs when a lock manager host object cannot be allocated due to reaching the limit for host objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the host object limit.

Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The number of host objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u locks %u/%u owners %u/%u.

Parameters

hosts_total (INT): Number of allocated hosts in the pool.
hosts_limit (INT): Allocated hosts limit for the pool.
pool_name (STRING): Name of the limit pool.
files_total (INT): Number of allocated files in the pool.
files_limit (INT): Allocated files limit for the pool.
locks_total (INT): Number of allocated locks in the pool.
locks_limit (INT): Allocated locks limit for the pool.
owners_total (INT): Number of allocated owners in the pool.
owners_limit (INT): Allocated owners limit for the pool.

lmgr.lock events

lmgr.lock.reclaim.aborted

Severity

ERROR

Description

This message occurs when the system aborts the volume lock reclaim process. A volume attempts to reclaim file locks when you bring the volume online after a reboot or after a takeover or giveback event. This error is usually caused by the volumes going offline during the initialization process.

Corrective Action

None;When the volumes come back online the reclaim process will be restarted.

Syslog Message

Aborting Lock manager grace period initialization on volume %s due to error: %s.

Parameters

aggr (STRING): The name of the aggregate.
error (STRING): Error code.

lmgr.lock.type.counts

Severity

NOTICE

Description

This message provides additional information about the current count of lock manager objects when a lock manager pool limit is reached. The combination of the counts from a message indicating a limit has been reached and this message provide context for investigating and resolving the limit issues.

Corrective Action

(None).

Syslog Message

The number of lock manager object types in pool %s are: share locks %u byte locks %u waiting locks %u callback waiting %u pfs (clustered Data ONTAP FlexCache delegations) %u cached locks %u WAN FlexCache delegations %u

Parameters

pool_name (STRING): Name of the limit pool.
share (INT): Number of allocated share locks in the pool.
byte (INT): Number of allocated byte locks in the pool.
waiting (INT): Number of allocated waiting locks in the pool.
cb_waiting (INT): Number of allocated cb_waiting locks in the pool.
pfs (INT): Number of allocated clustered Data ONTAP® FlexCache® delegations in the pool.
cached (INT): Number of allocated clustered Data ONTAP FlexCache cached locks in the pool.
delegations (INT): Number of allocated WAN FlexCache delegations in the pool.

lmgr.locks events

lmgr.locks.limit.hit

Severity

ERROR

Description

This message occurs when a lock manager lock object cannot be allocated due to reaching the limit for lock objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the lock object limit.

Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The number of lock objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u owners %u/%u.

Parameters

locks_total (INT): Number of allocated locks in the pool.
locks_limit (INT): Allocated locks limit for the pool.
pool_name (STRING): Name of the limit pool.
files_total (INT): Number of allocated files in the pool.

files_limit (INT): Allocated files limit for the pool.
hosts_total (INT): Number of allocated hosts in the pool.
hosts_limit (INT): Allocated hosts limit for the pool.
owners_total (INT): Number of allocated owners in the pool.
owners_limit (INT): Allocated owners limit for the pool.

lmgr.locks.sync.complete

Severity

NOTICE

Description

This message occurs when synchronization of client file locks ("locks sync") with the partner node completes over the aggregates noted in the previous "lmgr.locks.sync.start" message. Lock sync occurs for locks that were established using CIFS continuously available (CA) shares.

Corrective Action

(None).

Syslog Message

Lock sync completed on %s.

Parameters

aggr (STRING): Lock sync occurs in two cases: 1. Aggregate relocation 2. Giveback In the case of aggregate relocation, the aggregate name is provided. In the case of giveback, "aggr" indicates that lock sync is occuring over all aggregates on the node.

lmgr.locks.sync.start

Severity

NOTICE

Description

This message occurs when synchronization of client file locks ("locks sync") with the partner node starts over one or more aggregates. Lock sync occurs for locks that were established using CIFS continuously available (CA) shares.

Corrective Action

(None).

Syslog Message

Lock sync started on %s.

Parameters

aggr (STRING): Lock sync occurs in two cases: 1. Aggregate relocation 2. Giveback In the case of aggregate relocation, the aggregate name is provided. In the case of giveback, "aggr" indicates that lock sync is occuring over all aggregates on the node.

lmgr.locks.threshold.hit

Severity

ERROR

Description

This message occurs when lock manager lock object has reached the threshold limit for lock objects in the associated pool. The system might become overloaded. Clients will have problems accessing files through NAS protocols while at or above the lock object limit.

Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The current number of lock objects %u, with maximum limit of %u has reached the threshold limit percentage of %u for pool "%s". The current number/limit of other lock manager objects are: files %u/%u, hosts %u/%u and owners %u/%u.

Parameters

locks_total (INT): Number of allocated locks in the pool.
locks_limit (INT): Allocated locks limit for the pool.
threshold_percentage (INT): Threshold locks limit for the pool.
pool_name (STRING): Name of the limit pool.
files_total (INT): Number of allocated files in the pool.
files_limit (INT): Allocated files limit for the pool.
hosts_total (INT): Number of allocated hosts in the pool.
hosts_limit (INT): Allocated hosts limit for the pool.
owners_total (INT): Number of allocated owners in the pool.
owners_limit (INT): Allocated owners limit for the pool.

lmgr.lockstate events

lmgr.lockstate.cleanup.fail

Severity

ERROR

Description

This message occurs when the system cannot clean up the lock state associated with the active file system. The stale lock state can lead to problems accessing files through NAS protocols.

Corrective Action

Contact NetApp technical support for assistance.

Syslog Message

Lock state cleanup failed for '%s' during '%s' on volume '%s%s%s' (DSID %llu) with error %d.

Parameters

    lock_type (STRING): Type of locks being cleaned up.
    cleanup_reason (STRING): Reason to clean up locks.
    vol (STRING): Name of the volume.
    app (STRING): Application UUID.
    volident (STRING): Unique identifier for the volume.
    volDsid (LONGINT): Data set identifier (DSID) of the volume.
    code (INT): Error code returned while trying to clean up lock state.

lmgr.mcc events

lmgr.mcc.oplock.recall.fail

Severity

    ERROR

Description

    This message occurs when the system vetoes MetroCluster™ switchback or negotiated switchover because CIFS oplock recall fails. Some oplocks could not be recalled because of CPU overload, slow CIFS client response, or a large number of open files.

Corrective Action

    Recall oplocks for Vservers participating in negotiated switchover or switchback operations by using the command 'volume qtree oplocks -volume <vol-name> -qtree <qtree-name> -vserver <vserver-name> -oplock-mode disable' for volumes or qtrees with oplocks enabled. Then retry the MetroCluster operation, and reenable oplocks after the operation is complete. If the error persists, then contact NetApp technical support for assistance.

Syslog Message

    Vetoed MetroCluster switchback or negotiated switchover because CIFS oplock recall failed with internal error (%d).

Parameters

    error (INT): Error number.

lmgr.meta events

lmgr.meta.corrupted

Severity

    ERROR

Description

    This event occurs when the system detects that the lock manager host metafile is corrupted.

Corrective Action

    It is possibile that the corruption affected other parts of the file system. Run wafliron on the volume to verify file system consistency.

Syslog Message

Corrupted lock state file in volume %s: reinitializing.

Parameters

path (STRING): Volume name.

lmgr.mismatch events

lmgr.mismatch.file.objects

Severity

NOTICE

Description

This message occurs when the current file object of a lock does not match the file object found through the file handle. This might be caused by a SnapMirror® promote operation where the Master Data Set ID (MSID) is remapped.

Corrective Action

None.

Syslog Message

(None).

Parameters

volume (STRING): Name of the volume, based on the file handle.
lockid (STRING): Lock Universal Unique Identifier (UUID) of the sharelock.
caller (STRING): WAFL® spinnp operation that encountered this event.

lmgr.move events

lmgr.move.dump.lock.counts

Severity

NOTICE

Description

This message occurs when a volume move operation fails the cutover phase, due to volume quiesce timer expiration, and not all locks were transferred before the failure.

Corrective Action

Run the "vserver locks show -vserver <vserver> -volume <volume>" command to check the locks and the client IP addresses they are associated with. For NFS v4 locks, run instead the "vserver locks nfsv4 show -fields client-name" command in all nodes to get the same information. Reduce locks so that the number of locks remaining in this volume is fewer than the number scanned, and then attempt the volume move cutover again.

Syslog Message

The volume move of volume with DSID %llu did not cut over due to volume quiesce timer expiration. The volume move has scanned %d out of %d locks in the volume.

Parameters

volDsid (LONGINT): Data set identifier (DSID) of the volume.
scanned_locks (INT): Number of locks in the volume that were scanned before the failure.
total_locks (INT): Total number of locks in the volume.

lmgr.notify events

lmgr.notify.host.skip.write

Severity

INFORMATIONAL

Description

This message occurs when the volume lock manager cannot write to the notify hosts file due to an error, for example, insufficient space. In such cases, writing to the notify hosts file is skipped.

Corrective Action

(None).

Syslog Message

Skipped writing to notify hosts file %s on volume %s due to error: %s.

Parameters

file_name (STRING): Name of the file that writing was skipped on.
vol (STRING): Name of the volume.
error (STRING): Error code.

lmgr.owners events

lmgr.owners.limit.hit

Severity

ERROR

Description

This message occurs when a lock manager owner object cannot be allocated due to reaching the limit for owner objects in the associated pool. It is possible that the system is overloaded. Clients will have problems accessing files through NAS protocols while at or above the owner object limit.

Corrective Action

Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The number of owner objects %u has reached the limit %u for the pool %s. The current number/limit of other lock manager objects are: files %u/%u hosts %u/%u, locks %u/%u.

Parameters

> **owners_total** (INT): Number of allocated owners in the pool.
> **owners_limit** (INT): Allocated owners limit for the pool.
> **pool_name** (STRING): Name of the limit pool.
> **files_total** (INT): Number of allocated files in the pool.
> **files_limit** (INT): Allocated files limit for the pool.
> **hosts_total** (INT): Number of allocated hosts in the pool.
> **hosts_limit** (INT): Allocated hosts limit for the pool.
> **locks_total** (INT): Number of allocated locks in the pool.
> **locks_limit** (INT): Allocated locks limit for the pool.

lmgr.owners.threshold.hit

Severity

> ERROR

Description

> This message occurs when a lock manager owner object has reached the threshold limit for owner objects in the associated pool.The system might become overloaded. Clients will have problems accessing files through NAS protocols while at or above the owner object limit.

Corrective Action

> Move volumes to a node with a lower utilization of lock manager resources. Lock manager statistics can be displayed from the diagnostic privilege level. Use the "statistics show -raw true -object lmgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object lmgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object lmgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

> The current number of owner objects %u with maximum limit of %u has reached the threshold limit percentage of %u for pool "%s". The current number/limit of other lock manager objects are: files %u/%u, hosts %u/%u and locks %u/%u.

Parameters

> **owners_total** (INT): Number of allocated owners in the pool.
> **owners_limit** (INT): Allocated owners limit for the pool.
> **thresold_percentage** (INT): Threshold locks limit for the pool.
> **pool_name** (STRING): Name of the limit pool.
> **files_total** (INT): Number of allocated files in the pool.
> **files_limit** (INT): Allocated files limit for the pool.
> **hosts_total** (INT): Number of allocated hosts in the pool.
> **hosts_limit** (INT): Allocated hosts limit for the pool.
> **locks_total** (INT): Number of allocated locks in the pool.
> **locks_limit** (INT): Allocated locks limit for the pool.

lmgr.partner events

lmgr.partner.locks.outOfSync

Severity

ERROR

Description

This message occurs when a partner node's mirrored lock state goes out of sync with the actual lock state.

Corrective Action

Under normal conditions, no action is necessary. There is an automatic synchronize process that will attempt to resynchronize the lock state. You can wait for this process to synchronize the lock state, or you can manually resynchronize the lock state by running the 'debug locks mirrored sync' command in diag mode, with the Vserver, volume, and file ID supplied in this message.

Syslog Message

Mirrored lock state is out of sync for %s%s/%s (DSID %08x) for file %d Snapshot copy %d (mirroring status %d communication status %d)

Parameters

vol (STRING): Volume name.
app (STRING): Application UUID.
volident (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.
dsid (LONGINT): Data set identifier.
fileid (INT): File identifier.
snapid (INT): Snapshot(tm) copy in which the file exists, 0 for the active file system.
mirroring_status (INT): Error status from the mirroring operation.
csm_status (INT): Communication status from the mirroring operation.

lmgr.partner.locks.resynched

Severity

INFORMATIONAL

Description

This message occurs when a partner node's mirrored lock state is resynchronized with the actual lock state after having been out of sync previously.

Corrective Action

The mirrored lock state is now synchronized between the partner nodes. You can verify the mirror state of the locks by running the 'debug locks mirrored show' command in diag mode, with the Vserver, volume, and file ID supplied in this message.

Syslog Message

Mirrored lock state is now in sync for %s/%s (DSID %08x) for file %d Snapshot copy %d.

Parameters

vol (STRING): Volume name.
volident (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.
dsid (LONGINT): Data set identifier.
fileid (INT): File identifier.
snapid (INT): Snapshot(tm) copy in which the file exists, 0 for the active file system.

lmgr.precommit events

lmgr.precommit.oplock.recall.fail

Severity

ALERT

Description

This message occurs when the system fails aggregate giveback or relocation because SMB oplock or NFSv4 delegation recall fails. Giveback or relocation cannot proceed until all oplocks are recalled successfully or the lock manager pre-commit phase is forced, in which case cached writes in clients might be lost.

Corrective Action

There are two options for allowing aggregate giveback or relocation to proceed: (1) Force the sendhome pre-commit phase. For aggregate relocation, this can be done by specifying the '-override-destination-checks true' option. For giveback, this can be done by specifying the '-require-partner-waiting false' option. (2) Recall oplocks for Vservers on aggregates participating in the operation by using the 'volume qtree oplocks -volume <vol-name> -qtree <qtree-name> -vserver <vserver-name> -oplock-mode disable' command for volumes or qtrees with oplocks enabled. Then retry the operation, and reenable oplocks after the operation is complete. If the error persists, contact NetApp technical support for assistance.

Syslog Message

Aggregate "%s" failed giveback or relocation because lock manager oplock recall failed. Error: %s.

Parameters

aggrname (STRING): Aggregate name.
error (STRING): Error details.

lmgr.precommit.oplock.recall.forced

Severity

NOTICE

Description

This message occurs when the lock manager oplock recall fails during aggregate giveback or relocation, but the pre-commit phase has been forced to proceed. It is possible that cached writes in clients might have been lost as a result.

Corrective Action

(None).

Syslog Message

Lock manager oplock recall failed for aggregate "%s" during giveback or relocation. The failure abort was overridden. Error: %s.

Parameters

aggrname (STRING): Aggregate name.
error (STRING): Error details.

lmgr.precommit.oplock.recall.timeout

Severity

ALERT

Description

This message occurs when the lock manager oplock recall fails during aggregate giveback or relocation because the recalls could not be completed within the internal timeout window. The timeout typically can expire due to CPU overload, slow client response, or a large number of open files.

Corrective Action

The timeout can can be ignored by specifying the '-override-vetoes true' option in the giveback or relocation command.

Syslog Message

Lock manager oplock recall failed to be completed for aggregate "%s" within %d seconds.

Parameters

aggrname (STRING): Aggregate name.
timeout (INT): Timeout value.

lmgr.pruning events

lmgr.pruning.delayed

Severity

NOTICE

Description

This message occurs when there is a delay in pruning and removing unused entries in the volume lock manager hosts file. The current hosts file contents remain; unused entries were not removed.

Corrective Action

(None).

Syslog Message

Pruning lock manager on volume %s for file %s delayed due to error: %s.

Parameters

volume (STRING): Name of the volume.
file_name (STRING): Name of the host file being pruned.
error (STRING): Error description.

lmgr.recons events

lmgr.recons.CA.lock.dropped

Severity

ERROR

Description

This message occurs when locks established through CIFS Continuously Available (CA) shares are dropped after a takeover, giveback, or aggregate relocation. This is because the multi-protocol lock manager finds bad mirrored lock state for CA locks and skips reconstructing them. This causes a client to see a service disruption on the file to which the CA locks belong.

Corrective Action

Reopen the file to reestablish CA locks.

Syslog Message

A Continuously Available (CA) lock for volume %s%s/%s (DSID %08x) file %d snapshot copy %d reconnect key %d is dropped because of %s.

Parameters

vol (STRING): Volume name.
app (STRING): Application UUID.
volident (STRING): Volume identifier to uniquely identify the volume in cases where the volume name alone is insufficient.
dsid (LONGINT): Data set identifier.
fileid (INT): File identifier.
snapid (INT): Snapshot(tm) copy in which the file exists, 0 for the active file system.
reconnect_key (INT): Open identifier to uniquely identify a durable or persistent open on a volume
reason (STRING): Description of the reason for dropping a CA lock.

lmgr.sendhome events

lmgr.sendhome.veto.bootarg.override

Severity

NOTICE

Description

This message occurs during an SFO aggregate giveback or relocation. It implies that the automated nondisruptive update procedure will override any veto by the lock-manager, due to the presence of a bootarg.

Corrective Action

(None).

Syslog Message

Because bootarg.gb.override.lmgr.veto is set to true, the ANDU procedure will override any lock-manager veto regarding the giveback of the "%s" aggregate.

Parameters

aggregate (STRING): Name of the aggregate.

lmgr.session events

lmgr.session.failed

Severity

INFORMATIONAL

Description

This event occurs when a SpinNP session has failed, which causes session-scoped (e.g. CIFS) locks on the D-blade to be dropped. Only affects locks acquired using the failed session. This will cause clients to see a service disruption until they re-connect to the server.

Corrective Action

None — clients will automatically try to reconnect to the server, re-locking files as necessary.

Syslog Message

Dropped session-scoped locks from session %s:%s:%s due to failed session.

Parameters

initiator (STRING): The blade ID of the initiator.
uniquifier (STRING): The session uniquifier.
target (STRING): The blade ID of the target.

lmgr.sf events

lmgr.sf.up.ready

Severity

NOTICE

Description

This message occurs when the lock manager non-disruption subsystem for CIFS continuously available (CA) shares notifies the high-availability (HA) module that transition into the up state is allowed. After the lock manager has sent this notification, the HA module will allow certain failover options to proceed such as aggregate relocation and giveback.

Corrective Action

(None).

Syslog Message

Lock manager allowed high availability module to transition to the up state for the following reason: %s.

Parameters

reason (STRING): Reason for allowing transition to occur. Valid reasons include: 1. Lock-sync complete 2. Lock-sync timed out 3. CSM Invalid 4. Partner down 5. Mirroring disabled from bootarg 6. SMB3 disabled

lmgr.sh events

lmgr.sh.oplock.recall.fail

Deprecated

This EMS was replaced by the new "lmgr.precommit.oplock.recall.*" variants below and is no longer

needed.

Severity

ALERT

Description

This message occurs when the system vetoes aggregate giveback or relocation because CIFS oplock recall fails. Some oplocks could not be recalled because of CPU overload, slow CIFS client response, or a large number of open files. Giveback or relocation cannot proceed until all oplocks are recalled successfully or the veto is overridden, in which case cached writes in clients might be lost.

Corrective Action

Recall oplocks for Vservers on aggregates participating in the operation by using the command 'volume qtree oplocks -volume <vol-name> -qtree <qtree-name> -vserver <vserver-name> -oplock-mode disable' for volumes or qtrees with oplocks enabled. Then retry the operation, and reenable oplocks after the operation is complete. If the error persists, contact NetApp technical support for assistance.

Syslog Message

Aggregate "%s" vetoed giveback or relocation because CIFS oplock recall failed. Internal error: %s.

Parameters

aggrname (STRING): Aggregate name.
error (STRING): Error details.

lmgr.vol events

lmgr.vol.lock.reclaim.partial

Severity

NOTICE

Description

This message occurs when a volume cannot reclaim certain file locks when the system attempts to bring it online. This can occur after a reboot, a takeover, a giveback or an MCC event.

Corrective Action

(None).

Syslog Message

Lock reclaims for volume %s are partial due to the following error: %s.

Parameters

vol (STRING): The name of the volume.
error (STRING): Error code.