



Imgr events

ONTAP EMS reference

NetApp
November 18, 2025

This PDF was generated from <https://docs.netapp.com/us-en/ontap-ems/Imgr-aggr-events.html> on November 18, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Imgr events	1
Imgr.aggr events	1
Imgr.aggr.CA.locks.dropped	1
Imgr.badstructure events	1
Imgr.badStructure.found	1
Imgr.ca events	2
Imgr.CA.recons.mcc.veto	2
Imgr.cb events	2
Imgr.cb.waiting.limit.hit	2
Imgr.clients events	3
Imgr.clients.st.limit.hit	3
Imgr.delegations events	3
Imgr.delegations.limit.hit	3
Imgr.enable events	4
Imgr.enable.remote.lock.space.error	4
Imgr.files events	5
Imgr.files.limit.hit	5
Imgr.gb events	5
Imgr.gb.CA.locks.lost	5
Imgr.gb.CA.locks.may.lost	6
Imgr.hostfile events	6
Imgr.hostfile.cleared	6
Imgr.hostfile.create.abort	7
Imgr.hostfile.init.failed	7
Imgr.hosts events	7
Imgr.hosts.limit.hit	7
Imgr.lock events	8
Imgr.lock.reclaim.aborted	8
Imgr.lock.type.counts	8
Imgr.locks events	9
Imgr.locks.limit.hit	9
Imgr.locks.sync.complete	10
Imgr.locks.sync.start	10
Imgr.locks.threshold.hit	10
Imgr.lockstate events	11
Imgr.lockstate.cleanup.fail	11
Imgr.mcc events	12
Imgr.mcc.oplock.recall.fail	12
Imgr.meta events	12
Imgr.meta.corrupted	12
Imgr.mismatch events	13
Imgr.mismatch.file.objects	13
Imgr.move events	13

Imgr.move.dump.lock.counts	13
Imgr.nl events	14
Imgr.nl.open.limit.hit	14
Imgr.notify events	14
Imgr.notify.host.skip.write	14
Imgr.owners events	15
Imgr.owners.limit.hit	15
Imgr.owners.threshold.hit	15
Imgr.partner events	16
Imgr.partner.locks.outOfSync	16
Imgr.partner.locks.resynched	17
Imgr.persistent events	17
Imgr.persistent.file.open.limit.hit	17
Imgr.persistent.file.open.limit.warn	18
Imgr.persistent.open.limit.hit	18
Imgr.precommit events	19
Imgr.precommit.oplock.recall.fail	19
Imgr.precommit.oplock.recall.fail.oam	19
Imgr.precommit.oplock.recall.forced	20
Imgr.precommit.oplock.recall.timeout	20
Imgr.pruning events	21
Imgr.pruning.delayed	21
Imgr.recons events	21
Imgr.recons.CA.lock.dropped	21
Imgr.scrub events	22
Imgr.scrub.lock.space.error	22
Imgr.sendhome events	22
Imgr.sendhome.veto.bootarg.override	22
Imgr.session events	23
Imgr.session.failed	23
Imgr.sf events	23
Imgr.sf.up.ready	23
Imgr.sh events	24
Imgr.sh.oplock.recall.fail	24
Imgr.vol events	24
Imgr.vol.lock.reclaim.partial	24

Imgr events

Imgr.aggr events

Imgr.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 "Imgr.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".

Imgr.badstructure events

Imgr.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).

Imgr.ca events

Imgr.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).

Imgr.cb events

Imgr.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 Imgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr_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.

Imgr.clients events

Imgr.clients.st.limit.hit

Severity

ERROR

Description

This message occurs when the number of NFSv4 clients reaches the max limit of available per client object allocations, thereby prohibiting allocations for any new NFSv4 clients. This possibly happens due to system being overloaded. Client objects are used to efficiently keep track of all the locks per client.

Corrective Action

The per client lcs objects are introduced as part of NFSV4 session trunking. If you do not intend to use session trunking, disable it using "nfs modify -v4.1-trunking disabled" command. Use the "statistics show -raw true -object Imgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr_ng -node <node-name>" command for displaying all lock manager statistics on the specified node.

Syslog Message

The number of NFSV4 clients %u has reached the limit %u, exhausting the client objects available for allocation for the pool %s.

Parameters

clients_st_total (INT): Number of session trunking capable clients in total in the pool.
clients_st_limit (INT): Max supported limit for the session trunking capable clients in the pool.
pool_name (STRING): Name of the limit pool.

Imgr.delegations events

Imgr.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 Imgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr_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.

Imgr.enable events

Imgr.enable.remote.lock.space.error

Severity

NOTICE

Description

This message occurs when the enabling of the remote lock functionality is unsuccessful due to insufficient storage space.

Corrective Action

Run "Imgr enable_remote_lock <DSID> -multipart true" command from the diagnostic privilege level for enabling remote lock functionality manually.

Syslog Message

Remote lock functionality could not be enabled on volume '%s%s%s' (DSID %llu) due to insufficient storage space. Please consider freeing up some space and then attempt to enable it manually.

Parameters

vol (STRING): Name of the volume.
app (STRING): Application UUID.
volident (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.
volDsid (LONGINT): Data set identifier (DSID) of the volume.

Imgr.files events

Imgr.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 Imgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr_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.

Imgr.gb events

Imgr.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.

Imgr.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.

Imgr.hostfile events

Imgr.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.

Imgr.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.

Imgr.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.

Imgr.hosts events

Imgr.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 Imgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr_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.

Imgr.lock events

Imgr.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.

Imgr.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.

Imgr.locks events

Imgr.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 Imgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr_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.

Imgr.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 "Imgr.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 occurring over all aggregates on the node.

Imgr.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 occurring over all aggregates on the node.

Imgr.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.

Imgr.mcc events

Imgr.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.

Imgr.meta events

Imgr.meta.corrupted

Severity

ERROR

Description

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

Corrective Action

It is possible 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.

Imgr.mismatch events

Imgr.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.

Imgr.move events

Imgr.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.

Imgr.nl events

Imgr.nl.open.limit.hit

Severity

NOTICE

Description

This message occurs when the number of mirrored open lock count reaches the maximum supported limit.

Corrective Action

Run the "vserver locks show-mirrored-open-count -detailed" command to check the mirrored open lock count and try closing the unused files to bring this count down to the currently supported maximum mirrored open locks.

Syslog Message

The number of mirrored open locks count %u has reached the limit %u for 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

nl_locks_total (INT): Number of mirrored open locks.

nl_locks_limit (INT): Maximum supported mirrored open locks.

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.

Imgr.notify events

Imgr.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.

Imgr.owners events

Imgr.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 Imgr_ng" command for displaying all lock manager resource counters. Use the "statistics show -raw true -object Imgr_ng -instance <node-name>" command for displaying the lock manager statistics for a node's limit pool. Use the "statistics show -raw true -object Imgr_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.

Imgr.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.
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.
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 (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.

snapshot_id (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.

Imgr.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.

snapshot_id (INT): Snapshot(tm) copy in which the file exists, 0 for the active file system.

Imgr.persistent events

Imgr.persistent.file.open.limit.hit

Severity

NOTICE

Description

This message is triggered when the persistent open lock count for a single file reaches its maximum supported limit.

Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. From the diagnostic privilege level, execute "Imgr test show_ca_locks_count <file_path>" to inspect the persistent open lock count for the file. If necessary, close unused opens to reduce the count to the maximum limit of persistent open locks per file.

Syslog Message

The file with id %u on volume '%s' (DSID %llu) has hit its maximum persistent open lock count, fully utilizing the supported limit of %u.

Parameters

file_id (INT): File inode number.
vol_name (STRING): Name of the volume.
vol_dsid (LONGINT): Data set identifier (DSID) of the volume.
locks_limit (INT): Maximum supported persistent open locks count for a file.

Imgr.persistent.file.open.limit.warn

Severity

NOTICE

Description

This message is triggered when the persistent open lock count for a single file reaches 80% its maximum supported limit.

Corrective Action

Use the "volume file show-inode" command with the file ID and volume name information to find the file path. From the diagnostic privilege level, execute "Imgr test show_ca_locks_count <file_path>" to inspect the persistent open lock count for the file. If necessary, close unused opens to reduce the count to the maximum limit of persistent open locks per file.

Syslog Message

The persistent open lock count for file with id %u on volume '%s' (DSID %llu) has reached %u, constituting 80 percent of its maximum supported limit of %u.

Parameters

file_id (INT): File inode number.
vol_name (STRING): Name of the volume.
vol_dsid (LONGINT): Data set identifier (DSID) of the volume.
locks_total (INT): Number of persistent open lock count for file.
locks_limit (INT): Maximum supported persistent open locks count for a file.

Imgr.persistent.open.limit.hit

Severity

NOTICE

Description

This message occurs when the number of persistent open lock count reaches the maximum supported limit.

Corrective Action

Run the "vserver locks show-mirrored-open-count -detailed" command to check the persisted open lock count and try closing the unused files to bring this count down to the currently supported maximum persistent open locks.

Syslog Message

The number of persisted open locks count %u has reached the limit %u for 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

pl_locks_total (INT): Number of persisted open locks.
pl_locks_limit (INT): Maximum supported persisted open locks.

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.

Imgr.precommit events

Imgr.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.

Imgr.precommit.oplock.recall.fail.oam

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 relocation to proceed: (1) Force the sendhome pre-commit phase. It can be done by specifying the '-override-destination-checks true' 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. To allow aggregate giveback to proceed: (1) 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.

Imgr.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.

Imgr.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 be ignored by specifying the '-override-veto 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.

Imgr.pruning events

Imgr.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.

Imgr.recons events

Imgr.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.

Imgr.scrub events

Imgr.scrub.lock.space.error

Severity

NOTICE

Description

This message occurs when the internal scrubbing of Remote Lock Entry Metafile (RLEM) fails due to insufficient storage space.

Corrective Action

Run "Imgr scrub_remote_lock <DSID> -multipart true" command from the diagnostic privilege level for scrubbing the Remote Lock Entry Metafile (RLEM) manually.

Syslog Message

The system was unable to scrub the Remote Lock Entry Metafile (RLEM) automatically on volume '%s%s%s' (DSID %llu). Please scrub it manually.

Parameters

vol (STRING): Name of the volume.

app (STRING): Application UUID.

volident (STRING): Unique identifier of the volume in cases where the volume name alone is insufficient.

volDsid (LONGINT): Data set identifier (DSID) of the volume.

Imgr.sendhome events

Imgr.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.

Imgr.session events

Imgr.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.

Imgr.sf events

Imgr.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

Imgr.sh events

Imgr.sh.oplock.recall.fail

Deprecated

This EMS was replaced by the new "Imgr.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.

Imgr.vol events

Imgr.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.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.