



## **arl events**

### ONTAP EMS reference

NetApp  
November 20, 2025

This PDF was generated from <https://docs.netapp.com/us-en/ontap-ems-9151/arl-abort-events.html> on November 20, 2025. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

|   |    |
|---|----|
| arl events . . . . .                      | 1  |
| arl.abort events . . . . .                | 1  |
| arl.abort.raid . . . . .                  | 1  |
| arl.aggrdesttimeout events . . . . .      | 1  |
| arl.aggrDestTimeout . . . . .             | 1  |
| arl.aggronlinefailed events . . . . .     | 2  |
| arl.aggrOnlineFailed . . . . .            | 2  |
| arl.aggronlinetimeout events . . . . .    | 2  |
| arl.aggrOnlineTimeout . . . . .           | 2  |
| arl.aggrstart events . . . . .            | 2  |
| arl.aggrStart . . . . .                   | 2  |
| arl.connecttodenstfailed events . . . . . | 3  |
| arl.connectToDestFailed . . . . .         | 3  |
| arl.netra events . . . . .                | 3  |
| arl.netra.ca.check.failed . . . . .       | 3  |
| arl.netra.ha.dskChkFailed . . . . .       | 4  |
| arl.netra.lmgr.limit.failed . . . . .     | 4  |
| arl.netra.lmgr.reserve.fail . . . . .     | 5  |
| arl.netra.raid.failed . . . . .           | 5  |
| arl.netra.wafl.mcc.veto . . . . .         | 6  |
| arl.opaborted events . . . . .            | 6  |
| arl.OpAborted . . . . .                   | 6  |
| arl.opfinished events . . . . .           | 7  |
| arl.OpFinished . . . . .                  | 7  |
| arl.opstart events . . . . .              | 7  |
| arl.OpStart . . . . .                     | 7  |
| arl.postmigrnotok events . . . . .        | 8  |
| arl.postMigrNotOk . . . . .               | 8  |
| arl.precmt events . . . . .               | 8  |
| arl.precmt.lmgr.resync.abrt . . . . .     | 8  |
| arl.precmt.lmgr.resyncing . . . . .       | 9  |
| arl.precmt.lmgr.sync.failed . . . . .     | 9  |
| arl.precmt.lmgr.sync.forced . . . . .     | 10 |
| arl.precmt.repl . . . . .                 | 10 |
| arl.precmt.vdom.lowMem . . . . .          | 10 |
| arl.precmt.wafl.volconversion . . . . .   | 11 |
| arl.subsystemabort events . . . . .       | 11 |
| arl.subsystemAbort . . . . .              | 11 |
| arl.veto events . . . . .                 | 12 |
| arl.veto.coredump.saveinprog . . . . .    | 12 |
| arl.veto.coredump.unsavedcor . . . . .    | 12 |
| arl.veto.kmgr.keysmissing . . . . .       | 12 |
| arl.veto.lmgr.CA.volmove . . . . .        | 13 |

|                             |    |
|-----------------------------|----|
| arl.veto.lmgr.nonCA.broken  | 14 |
| arl.veto.lmgr.nonCA.locks   | 14 |
| arl.veto.lmgr.recons.left   | 15 |
| arl.veto.lmgr.syncing       | 15 |
| arl.veto.repl               | 15 |
| arl.veto.snaprestore        | 16 |
| arl.veto.volmove            | 16 |
| arl.veto.wafl.volconversion | 17 |

# arl events

## arl.abort events

### arl.abort.raid

#### Severity

ERROR

#### Description

This message occurs when a relocation of an aggregate is canceled due to activity in RAID.

#### Corrective Action

Reissue the "storage aggregate relocation start" command to initiate relocation of an aggregate to the specified destination node after the pending or in-progress operation on the aggregate is complete.

#### Syslog Message

Aggregate %s%s is %s; canceling 'storage aggregate relocation start' operation.

#### Parameters

**owner** (STRING): String indicating the owner of the affected aggregate.

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing the aggregate relocation operation.

## arl.aggrdesttimeout events

### arl.aggrDestTimeout

#### Severity

NOTICE

#### Description

This message occurs during an aggregate relocation operation, when the source node does not receive notification from the destination node, within the specified time, that the aggregate is online.

#### Corrective Action

Use the "storage aggregate show" command to verify that the aggregate and its volumes are online on the destination node.

#### Syslog Message

Node did not receive notification from the destination in %llu milliseconds that aggregate %s is online after relocation.

#### Parameters

**time** (LONGINT): Timeout, in milliseconds, for notification from the destination node that an aggregate is online after relocation.

**aggrName** (STRING): Name of the aggregate that was relocated.

## arl.aggronlinefailed events

### arl.aggrOnlineFailed

#### Severity

EMERGENCY

#### Description

This message occurs when the RAID or WAFL® subsystem fails to bring online the aggregate or its containing volumes after relocation.

#### Corrective Action

Check the EMS logs for messages from the RAID and WAFL subsystems to determine the exact cause of the failure and the corrective action.

#### Syslog Message

Bringing online of aggregate %s or its volumes failed after relocation with error %d.

#### Parameters

**aggrName** (STRING): Name of the aggregate that was relocated.

**error** (INT): Error that occurred while bringing online the aggregate or its volumes after relocation.

## arl.aggronlinetimeout events

### arl.aggrOnlineTimeout

#### Severity

ALERT

#### Description

This message occurs when a node fails to bring online the aggregate and its containing volumes within the specified time after relocation.

#### Corrective Action

Use the "storage aggregate show" command to verify that the aggregate and its volumes are online.

#### Syslog Message

Node did not bring online the aggregate %s and its volumes after relocation for %llu milliseconds.

#### Parameters

**aggrUUID** (STRING): UUID of the aggregate that was relocated.

**time** (LONGINT): Timeout, in milliseconds, to bring online the aggregate and its volumes after relocation.

## arl.aggrstart events

### arl.aggrStart

#### Severity

NOTICE

## Description

This message occurs when relocation of a storage failover aggregate starts.

## Corrective Action

(None).

## Syslog Message

Starting relocation of aggregate %s at time %lu.

## Parameters

**aggrName** (STRING): Name of the aggregate being relocated.

**time** (LONGINT): Time at which relocation of the aggregate started.

**aggrUUID** (STRING): UUID of the aggregate being relocated.

# arl.connecttodestfailed events

## arl.connectToDestFailed

### Severity

ALERT

## Description

This message occurs when the local node cannot connect over the cluster network to the d-blade of the destination node while relocating an aggregate.

## Corrective Action

1. Verify that the destination node is up.
2. Use the "network interface show" command to verify that the cluster network interfaces on the local node and the destination node are both up.
3. Retry the aggregate relocation command.

## Syslog Message

Failed to connect with the destination node %s (d-blade ID: %s) over the cluster network while relocating aggregate %s.

## Parameters

**destName** (STRING): Name of the destination node.

**destDbladeid** (STRING): D-blade ID of the destination node.

**aggrName** (STRING): Name of the aggregate being relocated.

# arl.netra events

## arl.netra.ca.check.failed

### Severity

ERROR

## Description

This message occurs during the relocation of an aggregate, when the destination node cannot reach the object stores.

## Corrective Action

Verify that your intercluster LIF is online and functional by using the 'network interface show' command. Also, check network connectivity to the object store server by using the 'ping' command over the destination node intercluster LIF. Additionally, verify that the configuration of your object store has not changed and that login and connectivity information is still accurate by using the 'aggregate object-store config show' command. It is possible to override this error by using the 'override-destination-checks' parameter of the relocation command. For more information or assistance, contact NetApp technical support.

## Syslog Message

Relocation of aggregate '%s' (uuid: %s) failed due to %s preventing object store access on the destination node.

## Parameters

**vol** (STRING): Name of the aggregate.

**aggr\_uuid** (STRING): UUID of the aggregate.

**reason** (STRING): Activity that is preventing the aggregate relocation operation.

## arl.netra.ha.dskChkFailed

### Severity

ALERT

### Description

This message occurs when relocation of an aggregate fails because the destination node cannot see all disks belonging to the aggregate.

## Corrective Action

If the destination is the source node's high-availability (HA) partner, identify the missing disks by using the "storage failover show -fields local-missing-disks, partner-missing-disks" command. If the destination is not the source node's HA partner, use the "storage disk show -aggregate" command to list the disks belonging to the aggregate, and then verify that those disks are visible from the destination by using the "storage disk show" command.

## Syslog Message

Relocation of aggregate %s failed from node %s because the destination node cannot see all disks belonging to the aggregate.

## Parameters

**aggr** (STRING): Name of the aggregate that was not relocated.

**node** (STRING): Name of the node that failed the relocation.

## arl.netra.lmgr.limit.failed

### Severity

ERROR

### Description

This message occurs during the relocation of an aggregate, when the destination node fails the lock count limit checks.

## Corrective Action

Ensure that the destination node supports the identified lock count limits, or select a different destination node and retry the operation.

## Syslog Message

Aggregate '%s': Lock limits check failed on destination node. %s

## Parameters

**vol** (STRING): Name of the aggregate.

**reason** (STRING): Activity that is preventing the aggregate relocation operation.

## arl.netra.lmgr.reserve.fail

### Severity

ERROR

### Description

This message occurs when an aggregate is being relocated, and the partner node fails to reserve lock resources needed for this aggregate.

## Corrective Action

Take the appropriate action based on the reason supplied in the logged message: - 'Partner reconstructing locks': Retry the relocation after reconstruction is complete on the partner node. You can check the reconstruction progress by using the 'debug locks reconstruction show' diagnostic privilege command. - 'Partner does not have sufficient lock resources': Free up lock resources on the partner and retry the relocation. You can free lock resources on the partner as client applications using locks are closed for volumes contained on that node. You can also free lock resources by moving volumes with locks from the partner node to a different node in the cluster. - 'Memory allocation failure': The system is running low on memory. Monitor memory statistics on the node and retry the relocation when more memory is available. - For other reason messages, or for further assistance, contact NetApp technical support.

## Syslog Message

Aggregate '%s': Lock reservation failed. Reason: %s.

## Parameters

**aggr** (STRING): Name of the aggregate being relocated.

**reason** (STRING): Reason the lock reservation failed.

## arl.netra.raid.failed

### Severity

ERROR

### Description

This message occurs during migration of an aggregate as part of aggregate relocation, when one of the aggregate checks fails on the destination node, which include duplicate aggregate name, duplicate Universally Unique Identifier (UUID), aggregate count, and capacity limits.

## Corrective Action

1. If checks fail due to aggregate count or capacity limits, then relocate this aggregate to an alternate destination node that has enough capacity to receive the aggregate.
2. If checks fail due to duplicate aggregate name or UUID, delete the duplicate aggregate created from previous "storage failover

giveback" or "storage aggregate relocation" operations. Contact NetApp technical support for assistance with deletion of the unwanted duplicate aggregate.

### Syslog Message

Aggregate '%s' (UUID = %s, type = %s, home\_owner\_id = %llu, dr\_home\_owner\_id = %llu): RAID aggregate migration checks failed on destination node %s.

### Parameters

**vol** (STRING): Name of the aggregate.  
**aggregate\_uuid** (STRING): UUID of the aggregate.  
**raid\_type** (STRING): RAID type of the volume.  
**home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's home owner.  
**dr\_home\_owner\_id** (LONGINT): NVRAM system ID of the aggregate's disaster recovery (DR) home owner.  
**reason** (STRING): Activity that is preventing the aggregate relocation or giveback operation.

## arl.netra.wafl.mcc.veto

### Severity

ERROR

### Description

This message occurs during the relocation of an aggregate, when one or more online left-behind disaster recovery(DR) aggregates are found on the destination node and the "node-object-limit" option is set to off.

### Corrective Action

Check whether there exists any online left-behind DR aggregates on the destination node from a previous switchover operation. If such an aggregate exists, then perform the corrective actions specified in the earlier EMS messages for the left-behind aggregate to return it to its original owner, and then retry the operation. If you cannot perform the corrective action, then use the "override-destination-checks" option in the "relocation" command to force the relocation.

### Syslog Message

Aggregate '%s': One or more online left-behind DR aggregates were found on the destination node and the "node-object-limit" option is set to off.

### Parameters

**vol** (STRING): Name of the aggregate.

## arl.opaborted events

### arl.OpAborted

### Severity

ERROR

### Description

This message occurs when an aggregate relocation operation is aborted.

### Corrective Action

Refer to other generated error messages to identify the exact cause of failure and the corrective action.

## Syslog Message

Aggregate relocation operation aborted by node %s because %s.

## Parameters

**node** (STRING): Name of the node that aborted relocation.

**msg** (STRING): Error that prevented relocation.

# arl.opfinished events

## arl.OpFinished

### Severity

NOTICE

### Description

This message occurs when the aggregate relocation operation from the source node to the specified destination is finished. Use the "storage aggregate relocation show" command to check whether all aggregates were relocated successfully.

### Corrective Action

(None).

## Syslog Message

Aggregate relocation operation from the source node %s to the destination node %s finished in %llu milliseconds. 'override-veto' set to %s, and 'override-destination-checks' set to %s.

## Parameters

**srcName** (STRING): Name of the source node.

**destName** (STRING): Name of the destination node.

**time** (LONGINT): Time, in milliseconds, taken to complete the relocation operation.

**override\_veto** (STRING): Flag that indicates whether the system should override veto checks for a relocation operation. This flag corresponds to the "-override-veto" parameter of the "storage aggregate relocation start" command. When set to true, it might result in the relocation proceeding even if the node detects outstanding issues that would make aggregate relocation dangerous or disruptive.

**override\_dest\_checks** (STRING): Flag that indicates whether the system should override checks done on the destination node. This flag corresponds to the "-override-destination-checks" parameter of the "storage aggregate relocation start" command. When set to true, it can be used to force relocation of aggregates even if the destination has outstanding issues. Note that this could make the relocation dangerous or disruptive.

# arl.opstart events

## arl.OpStart

### Severity

NOTICE

### Description

This message occurs when the aggregate relocation operation starts. After this message is displayed, aggregates are relocated to the specified destination serially.

## Corrective Action

(None).

## Syslog Message

Starting relocation to destination %s at time %llu. 'override-vetoes' set to %s, and 'override-destination-checks' set to %s.

## Parameters

**destName** (STRING): Name of the destination node.

**time** (LONGINT): Time at which relocation started.

**override\_vetoes** (STRING): Flag that indicates whether the system should override veto checks for a relocation operation. This flag corresponds to the "-override-vetoes" parameter of the "storage aggregate relocation start" command. When set to true, it might result in the relocation proceeding even if the node detects outstanding issues that would make aggregate relocation dangerous or disruptive.

**override\_dest\_checks** (STRING): Flag that indicates whether the system should override checks done on the destination node. This flag corresponds to the "-override-destination-checks" parameter of the "storage aggregate relocation start" command. When set to true, it can be used to force relocation of aggregates even if the destination has outstanding issues. Note that this could make the relocation dangerous or disruptive.

# arl.postmigrnotok events

## arl.postMigrNotOk

### Severity

ERROR

### Description

This message occurs when the specified subsystem encounters an error during the post-migration phase of the aggregate relocation (ARL) operation.

## Corrective Action

Check the event log for a subsystem-specific reason as to why the subsystem failed the post-migration phase of the relocation operation. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by that message.

## Syslog Message

Relocation of aggregate '%s' encountered an error during post-migration processing. The error was encountered by the subsystem '%s'.

## Parameters

**aggr** (STRING): Name of the aggregate on which the subsystem attempted the relocation operation.

**subsystem** (STRING): Name of the subsystem that failed the relocation operation.

# arl.precmt events

## arl.precmt.lmgr.resync.abrt

### Severity

NOTICE

## Description

This message occurs during aggregate relocation, when automatic lock resynchronization with the node's high-availability (HA) partner for files on this aggregate was aborted.

## Corrective Action

(None).

## Syslog Message

Automatic lock resynchronization was aborted for aggregate %s on forced relocation.

## Parameters

**aggr** (STRING): Name of the aggregate for which resynchronization was aborted.

## arl.precmt.lmgr.resyncing

### Severity

ERROR

## Description

This message occurs during aggregate relocation, when automatic lock resynchronization with the node's high-availability (HA) partner is in progress for files on this aggregate.

## Corrective Action

Retry the relocation after verifying that lock synchronization is complete by using the 'debug locks auto-resync-status -aggregate <aggregate>' diagnostic privilege command. If lock state disruption for all CA locks on the aggregate is acceptable, retry the aggregate relocation with the '-override-vetoed true' option.

## Syslog Message

Could not relocate aggregate %s because automatic lock resynchronization was in progress.

## Parameters

**aggr** (STRING): Name of the aggregate with locks not yet in sync.

## arl.precmt.lmgr.sync.failed

### Severity

ERROR

## Description

This message occurs when relocation of an aggregate has started but mirrored lock state cannot be created locally for the relocating aggregate. The relocate operation is aborted.

## Corrective Action

Take the appropriate action based on the reason supplied in the logged message: - 'Timed Out': Retry the aggregate relocation with the '-override-destination-checks' option. Overriding the check might result in a longer administrative outage period during which no administrative operations can be performed for volumes on this aggregate. - 'Fatal Error': Contact NetApp technical support.

## Syslog Message

Could not relocate the aggregate %s because creating locally mirrored lock state failed.

## Parameters

**aggr** (STRING): Name of the aggregate being relocated.

**reason** (STRING): Reason the local mirrored lock state creation failed.

## arl.precmt.Imgr.sync.forced

### Severity

NOTICE

### Description

This message occurs during aggregate relocation, when local mirrored lock state creation is allowed to proceed without any time restrictions. This might result in a longer administrative outage period during which no administrative operations can be performed for volumes on this aggregate.

### Corrective Action

(None).

### Syslog Message

Local mirrored lock state creation for relocating aggregate %s is proceeding without time restrictions.

## Parameters

**aggr** (STRING): Name of the aggregate being relocated.

## arl.precmt.repl

### Severity

NOTICE

### Description

This message occurs when all replication transfers are aborted because of aggregate relocation.

### Corrective Action

(None).

### Syslog Message

Replication transfers aborted due to aggregate relocation.

## Parameters

(None).

## arl.precmt.vdom.lowMem

### Severity

ERROR

### Description

This message occurs when an aggregate relocation fails due to the unavailability of Data ONTAP® memory on the source node.

## Corrective Action

Free some memory and retry the aggregate relocation.

## Syslog Message

Relocation of aggregate %s failed due to unavailability of Data ONTAP® memory on the source node.

## Parameters

**aggr** (STRING): Name of the aggregate that was not relocated.

# arl.precmt.wafl.volconversion

## Severity

ERROR

## Description

This message occurs when an active volume conversion is running on one or more volumes on the source aggregate. Pre-commit is aborted.

## Corrective Action

When volume conversion is complete, re-run the command 'storage aggregate relocation start'.

## Syslog Message

Aggregate relocation/takeover is blocked as Volume Conversion is in progress on one of the volumes on aggregate %s.

## Parameters

**aggregate\_name** (STRING): The name of the aggregate containing the volume on which conversion is in progress.

# arl.subsystemabort events

## arl.subsystemAbort

## Severity

ERROR

## Description

This message occurs when the specified subsystem aborts relocation of the aggregate.

## Corrective Action

Check the syslog/EMS output for a subsystem-specific reason for aborting the relocation operation. The corrective action is subsystem-specific and is detailed in the corrective action portion of the message. Follow the corrective action specified by the subsystem and then reissue the 'relocation' command. If you cannot perform the corrective action, then use the 'override-vetoes' option in the 'relocation' command to force the relocation.

## Syslog Message

The relocation operation of '%s' was aborted by '%s'.

## Parameters

**aggr** (STRING): Name of the aggregate on which the subsystem attempted the relocation operation.

**subsystem** (STRING): Name of the subsystem that aborted the relocation operation.

# arl.veto events

## arl.veto.coredump.saveinprog

### Severity

INFORMATIONAL

### Description

This message occurs when the coredump subsystem vetoes the use of storage failover (SFO) during takeover because a coredump save is in progress; cluster failover (CFO) will be used instead. The coredump save must be complete before the system can perform a takeover.

### Corrective Action

(None).

### Syslog Message

Takeover use of SFO vetoed because a coredump save was in progress.

## Parameters

(None).

## arl.veto.coredump.unsavedcor

### Severity

ERROR

### Description

This message occurs when aggregate relocation (ARL) is vetoed because there is an unsaved kernel core file on one of the aggregate disks.

### Corrective Action

Use the "node coredump save" command to save any unsaved cores, or use the "-override-vetoes true" option when relocating aggregates.

### Syslog Message

Aggregate relocation vetoed due to an unsaved coredump %s on disk %s ("%s").

## Parameters

**corefile** (STRING): Name of the kernel core file.

**disk** (STRING): Name of the disk containing the kernel core file.

**panic** (STRING): Panic string of the unsaved kernel core file that caused the aggregate relocation (ARL) veto.

## arl.veto.kmgr.keysmissing

## Severity

ERROR

## Description

This message occurs when an aggregate relocation fails due to the unavailability of volume encryption keys for the encrypted volumes of the aggregate on the destination node.

## Corrective Action

Wait a few minutes, and then try the aggregate relocation again. If the problem persists, run the "security key-manager external restore" command for external key manager, or the "security key-manager onboard sync" command for Onboard Key Manager. The commands retrieve the encryption keys for the encrypted volumes of the aggregate from the key manager. Then try the aggregate relocation again.

## Syslog Message

Relocation of aggregate %s failed due to unavailability of volume encryption keys for the encrypted volumes of the aggregate on the destination node %s. Details: %s.

## Parameters

**aggr** (STRING): Name of the aggregate that was not relocated.

**node** (STRING): Name of the destination node that is missing encryption keys.

**details** (STRING): Detailed reason for failure.

## arl.veto.lmgr.CA.volmove

## Severity

ERROR

## Description

This message occurs when an aggregate relocation has started but cannot proceed because a volume move to or from this aggregate is in the cutover phase and the volume being moved has continuously available (CA) locks. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

## Corrective Action

Retry the relocation after volume move exits the cutover phase. To check for volume move state, run the 'volume move show' command.

## Syslog Message

Could not complete aggregate relocation because volume move of '%s%s%s%s' is in the cutover phase and this volume has CA locks.

## Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

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

## arl.veto.lmgr.nonCA.broken

### Severity

NOTICE

### Description

This message occurs when an aggregate relocation is forced that causes non-continuously available (non-CA) locks on the volume to be dropped. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

### Corrective Action

(None).

### Syslog Message

Dropped non-CA locks on volume %s%s%s%s due to forced aggregate relocation of aggregate %s.

### Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

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

**aggrname** (STRING): Aggregate name.

## arl.veto.lmgr.nonCA.locks

### Severity

ERROR

### Description

This message occurs when an aggregate relocation has started but cannot proceed because non-continuously available (non-CA) locks are present on the volume. CA locks are established by opens through CIFS CA shares for regular files on read-write volumes that reside in storage failover (SFO) aggregates. These locks are mirrored to the node's high-availability (HA) partner to support the nondisruptive property of CA shares. The rest of the locks are classified as non-CA locks and are not mirrored to the node's HA partner.

### Corrective Action

Based on how resilient they are to failures, applications must either gracefully close sessions over which non-CA locks are established or accept lock state disruption. To determine the open files that have these sessions established, run the 'vserver cifs session file show -hosting-aggregate "aggregate name" -continuously-available No' command. "aggregate name" is the aggregate for which relocation is being attempted. If lock state disruption for all existing non-CA locks is acceptable, retry the aggregate relocation by using the '-override-vetoes true' option.

### Syslog Message

Could not complete aggregate relocation because of non-CA locks on volume %s%s%s%s aggregate %s.

### Parameters

**owner** (STRING): Volume owner.

**vol** (STRING): Volume name.

**app** (STRING): Application UUID.

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

**agrname** (STRING): Aggregate name.

## arl.veto.Imgr.recons.left

### Severity

ERROR

### Description

This message occurs when an aggregate relocation has started, but cannot proceed because reconstruction of client file locks from mirrored lock information is not yet complete.

### Corrective Action

Retry the relocation a few times after checking the status of reconstruction by using the 'debug locks reconstruction show' diagnostic privilege command. If the locks reconstruction operation is not completed, contact NetApp technical support.

### Syslog Message

Could not relocate aggregate '%s' because lock reconstruction on node %s is in progress.

### Parameters

**aggregate** (STRING): Name of the aggregate being relocated.

**node** (STRING): Name of the node reconstructing locks.

## arl.veto.Imgr.syncing

### Severity

ERROR

### Description

This message occurs when an aggregate relocation has started but cannot proceed because synchronization of client file locks ("locks sync") with the partner is not yet complete.

### Corrective Action

Retry the relocation after verifying that lock synchronization is complete, by using the 'debug locks lock-sync show' diagnostic privilege command.

### Syslog Message

Could not relocate aggregate '%s' because locks sync from node %s to its partner node %s is in progress.

### Parameters

**aggregate** (STRING): Name of the aggregate being relocated.

**node** (STRING): Name of the node performing locks sync with its partner.

**partner** (STRING): Name of the partner node with which locks sync is being performed.

## arl.veto.repl

### Severity

NOTICE

## Description

This message occurs when an aggregate relocation is vetoed by a critical replication transfer (for example, volume move in the cutover phase).

## Corrective Action

Either abort the replication transfer that is preventing the aggregate relocation from being completed, retry the 'aggregate relocation start' command with the '-override-vetoes true' option, or wait for the transfer to finish.

## Syslog Message

Could not complete giveback because a replication transfer with UUID %s involving the aggregate %s is in progress.

## Parameters

**transferId** (STRING): UUID of the SnapMirror® replication transfer that vetoed the aggregate relocation.  
**aggr** (STRING): Name of the aggregate.

## arl.veto.snaprestore

### Severity

ERROR

## Description

This message occurs when the system cannot perform an aggregate relocation because a volume snaprestore operation is active. The relocation is aborted.

## Corrective Action

Retry aggregate relocation after the snaprestore operation is finished.

## Syslog Message

%s %s%s%s%s is running snaprestore. Canceling aggregate relocation.

## Parameters

**type** (STRING): Type of object (volume or aggregate).  
**owner** (STRING): Volume owner.  
**vol** (STRING): Volume name.  
**app** (STRING): Application UUID.  
**volident** (STRING): Unique volume identity when the volume name itself is insufficient.

## arl.veto.volmove

### Severity

ERROR

## Description

This message occurs when an active volume move operation that cannot be automatically aborted prevents aggregate relocation from starting. When the volume move operation is completed, aggregate relocation can be retried using the 'storage aggregate relocation start' command.

## Corrective Action

When the volume move operation is complete, reissue the 'storage aggregate relocation start' command. To abort the volume move immediately, include the '-override-vetoes true' command option.

## Syslog Message

Unable to relocate aggregate while volume move for volume (DSID: %llu, NAME: %s) on aggregate %s is in progress.

## Parameters

**volume\_dsid** (LONGINT): Source/destination volume Data Set ID (DSID).  
**vol\_name** (STRING): The name of the volume.  
**aggr\_name** (STRING): The name of the containing aggregate.

## arl.veto.wafl.volconversion

### Severity

ERROR

### Description

This message occurs when an active volume conversion against one or more of the volumes on the source aggregate vetoed aggregate relocation.

### Corrective Action

When volume conversion is complete, re-run the command 'storage aggregate relocation start'.

## Syslog Message

Aggregate relocation is blocked as Volume Conversion is in progress on one of the volumes on aggregate %s.

## Parameters

**aggregate\_name** (STRING): The name of the aggregate containing the volume on which conversion is in progress.

## 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.