

Test scenario and troubleshooting of Open-E DSS V7 (up60) cluster shutdown and start up

Scenario:

Step 1.

Ensure both *Replication state* fields of *Resources pool* are **synced** (in Setup -> Failover -> Failover Manager). Otherwise, wait until synchronization is completed.

Move all resources to one node (from Node A to Node B) using **move to remote node** button in the Setup -> Failover -> Resources Pool Manager.

Step 2.

Shutdown the node without resources (Node A) from the WebGUI: Maintenance -> Shutdown -> System Shutdown.

Node B Failover status after shutdown of Node A.

The screenshot shows the Open-E DSS V7 WebGUI interface. The top navigation bar includes 'SETUP', 'CONFIGURATION', 'MAINTENANCE', 'STATUS', and 'HELP'. The breadcrumb trail indicates 'You are here: Setup > Failover'. The main content area is titled 'Failover Manager' and displays the following information:

- Cluster status:** Running - Degraded
- Important!** Please refer to [Failover. Important notes](#) help for important information related to configuration and maintenance of failover services.
- Resources pool:**
 - dssc2nb (local node) resources pool:**
 - Status: active on dssc2nb (local node)
 - Replication state: not synced
 - Persistent reservation synchronization: inactive
 - dssc2na (remote node) resources pool:**
 - Status: active on dssc2nb (local node)
 - Replication state: not synced
 - Persistent reservation synchronization: inactive
- Network statuses:**
 - Ping nodes: 0 of 1 reachable
 - Auxiliary paths: 0 of 3 reachable
- Remote node status:**
 - Remote node availability: Not reachable
 - Remote node hostname: dssc2na
 - Remote node IP: 192.168.3.21
- Maintenance mode:**
 - Status: Off

Step 3.

Shutdown Node B from the WebGUI: Maintenance -> Shutdown -> System Shutdown.

Step 4.

Power on Node B.

Step 5.

Start Failover on Node B in local start mode: Setup -> Failover -> Failover Manager -> **Start** button.

Note: Do not change Volume Replication Mode!

Note: Do not start any Volume Replication tasks!.

Wait until *Resources pool* status will be active in the Failover Manager.

Node B Failover status when running in local start mode.

The screenshot shows the Failover Manager web interface. At the top, there is a navigation bar with tabs for SETUP, CONFIGURATION, MAINTENANCE, STATUS, and HELP. Below the navigation bar, the breadcrumb path is "You are here: Setup > Failover". The main content area is titled "Failover Manager" and displays an error message:

Error
Cluster started on local node only
Could not start cluster services on both nodes, all resources served from local node only.
In order to resolve this issue please start cluster services on the remote node. Please make sure first the remote node is running and can be reached in the network by the local node.
If the remote node is not available you may also:

- Check your network settings to see if they are configured properly [Go to network settings management >](#)
- Check your network infrastructure to see if it is possible to establish connection between the local and remote node

Below the error message, the cluster status is shown as **Running - Degraded**. An **Important!** note refers to [Failover: Important notes](#) help for important information related to configuration and maintenance of failover services. A **stop** button is visible.

The **Resources pool** section shows the following details:

dssc2nb (local node) resources pool:	
Status:	active on dssc2nb (local node)
Replication state:	not synced
Persistent reservation synchronization:	inactive

dssc2na (remote node) resources pool:	
Status:	active on dssc2nb (local node)
Replication state:	not synced
Persistent reservation synchronization:	inactive

[See details >](#)

The **Network statuses** and **Remote node status** sections are also visible:

Network statuses	Remote node status
Ping nodes: 0 of 1 reachable	Remote node availability: Not reachable
See details >	Remote node hostname: dssc2na
Auxiliary paths: 0 of 3 reachable	Remote node IP: 192.168.3.21

Step 6.

Power on Node A. It should connect and join Failover automatically.

Node A Failover status after joining the cluster.

The screenshot displays the Failover Manager web interface. At the top, there is a navigation bar with tabs for SETUP, CONFIGURATION, MAINTENANCE, STATUS, and HELP. The current page is 'Failover', indicated by the breadcrumb 'You are here: Setup > Failover'. The main content area is titled 'Failover Manager' and shows the following information:

- Cluster status: Running - OK**
- Important!** Please refer to [Failover: Important notes](#) help for important information related to configuration and maintenance of failover services. A red 'stop' button is located below this message.
- Resources pool**
 - dssc2na (local node) resources pool:**
 - Status: **active on dssc2nb (remote node)**
 - Replication state: **synced**
 - Persistent reservation synchronization: **active**
 - dssc2nb (remote node) resources pool:**
 - Status: **active on dssc2nb (remote node)**
 - Replication state: **synced**
 - Persistent reservation synchronization: **active**
- [See details »](#)
- Network statuses**
 - Ping nodes: **1 of 1 reachable**
 - [See details »](#)
 - Auxiliary paths: **3 of 3 reachable**
 - [See details »](#)
- Remote node status**
 - Remote node availability: **Reachable**
 - Remote node hostname: **dssc2nb**
 - Remote node IP: **192.168.3.22**
 - [See details »](#)
- Maintenance mode**
 - Status: **Off**

Step 7.

Move resources to the original state (from Node B to Node A) using on Node B: **move to remote node** button in the Setup -> Failover -> Resources Pool Manager.

Troubleshooting:

Case 1.


In Step 6 after Node A join, when on both nodes, *Cluster Status* will show up as **Running – Degraded** and all *Resources pool* fields show **active** and **synced**.

The screenshot displays the 'Failover Manager' interface. At the top, the 'open-e' logo and 'ENTERPRISE CLASS STORAGE OS for EVERY BUSINESS' are visible, along with 'DATA STORAGE SOFTWARE V7'. The navigation menu includes 'SETUP', 'CONFIGURATION', 'MAINTENANCE', 'STATUS', and 'HELP'. The breadcrumb trail shows 'You are here: Setup > Failover'. The main content area shows the 'Failover Manager' title and a 'Cluster status: Running - Degraded' message. Below this, an 'Important!' note refers to 'Failover: Important notes' help. A red 'stop' button is present. The 'Resources pool' section is divided into two parts: 'dssc2nb (local node) resources pool' and 'dssc2na (remote node) resources pool'. Both show 'Status: active on dssc2nb (local node)', 'Replication state: synced', and 'Persistent reservation synchronization: active'. A 'See details >' link is provided. The 'Network statuses' section shows 'Ping nodes: 1 of 1 reachable' and 'Auxiliary paths: 3 of 3 reachable', with 'See details >' links. The 'Remote node status' section shows 'Remote node availability: Reachable', 'Remote node hostname: dssc2na', and 'Remote node IP: 192.168.3.21', with a 'See details >' link. The 'Maintenance mode' section shows 'Status: Off'.

Solution:

1. Continue with Step 7 (Move resources to their original configuration).
2. On console tools, start Maintenance Mode by selecting **Ctrl+Alt+X -> Cluster Maintenance Mode**

3. On GUI, Stop the cluster in Setup -> Failover -> Failover Manager. On the confirmation prompt, you must choose **stop and enable VIPs** button



4. Start the cluster again.

NOTE: The maintenance mode can be enabled in production permanently.

Case 2.

In Step 6 after Node A join, when on Node A *Cluster Status* will show up as **Stopped** (and on Node B as **Running – Degraded**) and in *Resources pool*, *Replication state* shows **not synced** (on at least one resource pool). Also, Event Viewer on Node A shows “**Failed to join the cluster, incorrect configuration (error code: 13).**” error message.

Node A (joining node):

The screenshot shows the Failover Manager interface. At the top, there is a header with the title "Failover Manager" and three icons: a refresh icon, a home icon, and a help icon. Below the header is a red error banner with a white 'x' icon. The error message reads: "Error Cluster services are stopped. Please make sure that remote node is running and can be reached in the network by the local node and start cluster services manually." Below the error banner, the main content area shows "Cluster status: Stopped" in red. Underneath, there is an "Important!" note with a link to "Failover: Important notes" and a red "start" button. The "Resources pool" section is divided into two parts: "dssc2nb (local node) resources pool:" and "dssc2na (remote node) resources pool:". Each part lists "Status:", "Replication state:", and "Persistent reservation synchronization:". For both pools, the status is "active on dssc2na (remote node)", the replication state is "not synced", and the reservation synchronization is "inactive". There are "See details »" links for both pools. The "Network statuses" section is split into two columns: "Network statuses" and "Remote node status". Under "Network statuses", "Ping nodes:" is "1 of 1 reachable" and "Auxiliary paths:" is "3 defined". Under "Remote node status", "Remote node availability:" is "Reachable", "Remote node hostname:" is "dssc2na", and "Remote node IP:" is "192.168.3.21". There are "See details »" links for both columns. Finally, the "Maintenance mode" section shows "Status:" as "Off".

Failover Manager

Error
Cluster services are stopped

Please make sure that remote node is running and can be reached in the network by the local node and start cluster services manually.

Cluster status: Stopped

Important! Please refer to [Failover: Important notes](#) help for important information related to configuration and maintenance of failover services.

start

Resources pool

dssc2nb (local node) resources pool:
Status: **active on dssc2na (remote node)**
Replication state: **not synced**
Persistent reservation synchronization: inactive

dssc2na (remote node) resources pool:
Status: **active on dssc2na (remote node)**
Replication state: **not synced**
Persistent reservation synchronization: inactive

[See details »](#)

Network statuses		Remote node status	
Ping nodes:	1 of 1 reachable	Remote node availability:	Reachable
See details »		Remote node hostname:	dssc2na
Auxiliary paths:	3 defined	Remote node IP:	192.168.3.21
See details »		See details »	

Maintenance mode

Status: Off

Node B (active node):

Failover Manager [Refresh] [Help]

Cluster status: Running - Degraded

Important! Please refer to [Failover: Important notes](#) help for important information related to configuration and maintenance of failover services.

stop

Resources pool

dssc2na (local node) resources pool:
Status: **active on dssc2na (local node)**
Replication state: **not synced**
Persistent reservation synchronization: inactive

dssc2nb (remote node) resources pool:
Status: **active on dssc2na (local node)**
Replication state: **not synced**
Persistent reservation synchronization: inactive

[See details »](#)

Network statuses	Remote node status
Ping nodes: 1 of 1 reachable See details »	Remote node availability: Reachable
Auxiliary paths: 0 of 3 reachable See details »	Remote node hostname: dssc2nb
	Remote node IP: 192.168.3.22 See details »

Maintenance mode

Status: Off

Solution:

1. Please stop all running replication tasks on Node B (Configuration -> Volume manager -> Volume replication)

Replication tasks manager [Refresh] [Help]

Info
Following replication task(s) can't be removed because such operation is allowed only for tasks not added to the cluster:
• t0
• t1_reverse
In order to have a task completely removed from the system it has to be removed from the cluster first.

Name	Start time	Action
t0	2017-09-19 14:01:55	[Play] [Stop] [Delete]
t1_reverse	2017-09-19 14:01:45	[Play] [Stop] [Delete]

2. Ensure all volumes (which are used in cluster) on Node A are marked as Destination (Configuration -> Volume manager -> Volume replication) and all volumes on Node B are marked as Source. If not, make appropriate changes.

Node A (joining node):

Volume replication mode					⌂	↻	?
Logical Volume	Init	Source	Destination	Clear metadata			
lv0000	done	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			
lv0001	done	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>			

apply

Node B (active node):

Volume replication mode					⌂	↻	?
Logical Volume	Init	Source	Destination	Clear metadata			
lv0000	done	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>			
lv0001	done	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>			

apply

3. Start all replication tasks on Node B (Configuration -> Volume manager -> Volume replication)
4. Ensure *Replication state* for both pools in Failover Manager (Setup -> Failover -> Failover Manager) on any node shows **syncing in progress** or **synced**.
5. Start cluster on Node A (joining node) in Failover Manager (Setup -> Failover -> Failover Manager)

Possible cause of problem:

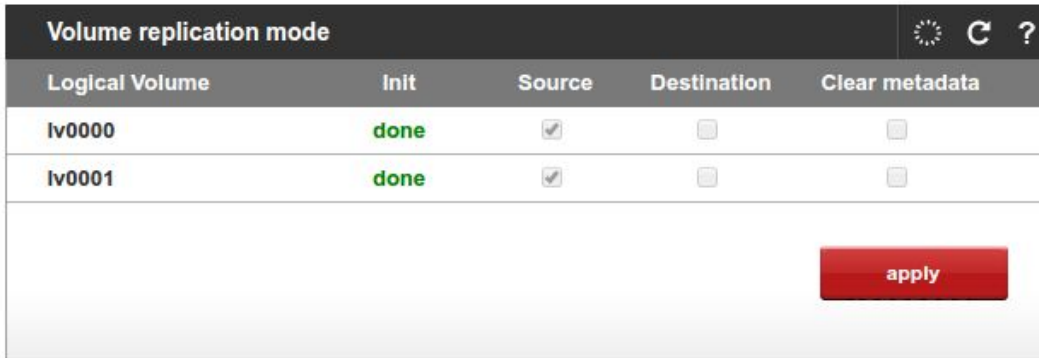
- during Node A joining (Step 6), connection errors on replication level or network issues occurred,
- user has changed volume replication mode on node B while node A is off (after Step 2) - eg. change of one volume from Source to Destination and back to Source

Case3.

In Volume replication mode (Configuration -> Volume manager), if Source / Destination check boxes are "grey out" and it is not possible to change destination to source. (See the screenshot below.)

Workaround is to click on task-start button and refresh browser (F5)

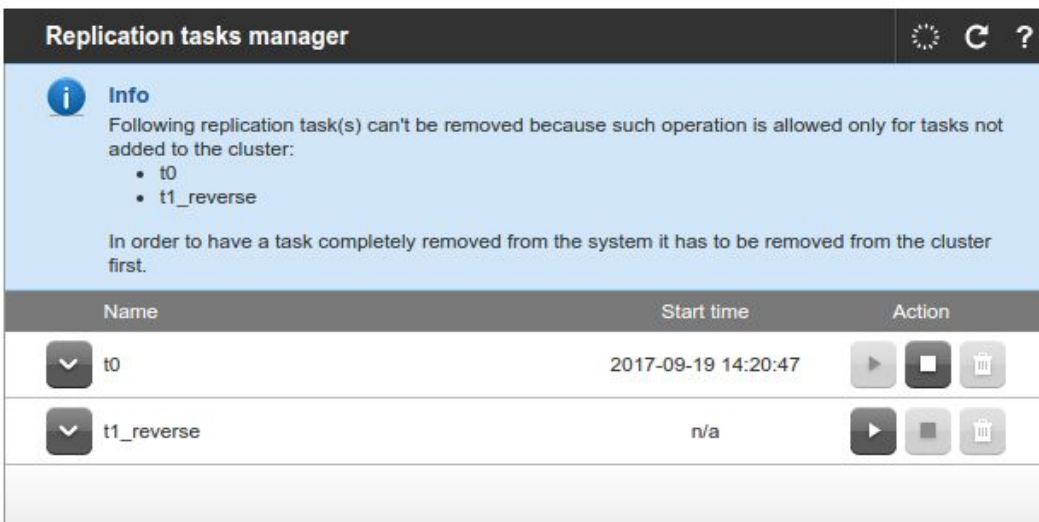
Note: This is an old known bug. This bug was fixed and should NOT happen in up60



The screenshot shows a web interface titled "Volume replication mode". It features a table with the following columns: "Logical Volume", "Init", "Source", "Destination", and "Clear metadata". There are two rows of data:

Logical Volume	Init	Source	Destination	Clear metadata
lv0000	done	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
lv0001	done	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Below the table is a red "apply" button.









The screenshot shows a web interface titled "Replication tasks manager". It contains an information box and a table of tasks:

Info
Following replication task(s) can't be removed because such operation is allowed only for tasks not added to the cluster:

- t0
- t1_reverse

In order to have a task completely removed from the system it has to be removed from the cluster first.

Name	Start time	Action
t0	2017-09-19 14:20:47	  
t1_reverse	n/a	  

General case.

In some cases, your web browser may cache information from the GUI and show unexpected information. In such case you can refresh the browser using ctrl + F5 keys or refer to your web browser manual.