DM Series Root Aggregate and Cluster Setup Guide

DM3000/DM5000/DM7000

First Edition (February 2020)

© Copyright Lenovo 2020.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925

Table of Contents

1. Overview of DM Series Setup Process 4

2. Create Root Aggregates and Cluster 5

Appendix A: Installing Data ONTAP OS at boot menu in the serial console 8

Appendix B: Disabling AutoSupport 9

Appendix C: Verifying HA mode on nodes 10

Appendix D: Setting the shelf ID on shelves 11

1. Overview of DM Series Setup Process

Notice:

By following this guide, all data currently stored on the DM system will be wiped. It is important to verify the pre-requisites before continuing with the process. If a pre-requisite is not followed, the re-configuration process may not complete successfully.

Prior to beginning the cluster setup and being able to use a DM stack, it is necessary to setup the root aggregates first.

- For All Flash models, the root aggregate will be striped across the drives using Advanced Drive Partitioning (ADP) and a small partition on each SSD will be reserved and used in the root aggregate. Four partitions in total will be created across the drives. Two root aggregate partitions (one for each node) and two data partitions (one for each node).
- For Hybrid models running ONTAP 9.4 FW, the root aggregate will reserve 3 HDDs to serve as exclusive access for the root aggregate for each node.
- For Hybrid models running ONTAP 9.5 and later FW, the root aggregate will be created with a root-data partition. This will normally consist of 5 drives per controller on a DM3000H (ten drives total).

Once the disks have been assigned and root aggregates configured, cluster setup should be performed before using the system. If the DM system was used previously in a cluster or disk ownerships have been assigned, it will be necessary to clear the disk ownerships and initialize the disks to re-do the root aggregates.

The cabling for expansion shelves does matter when it comes to creating root aggregates. In All Flash models, the first SSDs detected will be the drives chosen for the root aggregates. For Hybrid models, if the DM Series stack has 24 or less SSDs, the default drives selected for the root aggregates will be the first HDDs it detects within the stack. For stacks with 25 or greater SSDs, the first SSDs detected will be used for the root aggregates.

Once the root aggregates are created, it is important to keep the stack with the head shelf and expansions shelves together so the root aggregate can be kept intact.

2. Create Root Aggregates and Cluster

Pre-requisites:

- Both nodes in the DM head shelf are in ha mode. Refer to Appendix C if not
- The shelf IDs have been assigned and are unique. Refer to Appendix D if not
- The DM head shelf and expansion shelves, if present, should all be powered **off** before beginning this process

1. Follow cabling procedures for standalone DM system (without any expansions) or a DM system with expansion shelves. Refer to sections 2 & 3 in the Installation and Setup Poster for cabling guidelines, also found on the Data Center Support site (https://datacentersupport.lenovo.com)

   Notes:
   - Cabling of expansion shelves does matter. Refer to Overview of DM Series Setup Process for more details.
   - For DM7000, it is necessary to attach at least 1 expansion shelf so the system has disks.
   - Ensure the two micro USB serial connections from the two nodes are connected to the PC.
     1. PC should be running Windows 10 or later
     2. The COM Settings are 115200, 8-N-1
2. Power on the expansion shelves if present. If not, power on the head shelf and skip to step 4
3. Wait 1 minute and then power on the DM head shelf
4. When the message stating AUTOBOOT is starting appears, hit Ctrl-C to interrupt the boot and get to the loader prompt; type `boot_ontap menu` to access the boot menu. Do this on both nodes. Will be presented with the following:

```
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)?
```

   Note:
   A set boot parameter may cause the node(s) to halt at the loader prompt before getting to the boot menu. In this case, run `boot_ontap menu` at the loader prompt to continue the boot process and stop at the boot menu.
5. Select Option 9 from the boot menu for configuring advanced drive partitioning and the following will be presented:

```
**********************************************
* Advanced Drive Partitioning Boot Menu Options *
**********************************************

(9a) Unpartition all disks and remove their ownership information.
```

```
(9b) Clean configuration and initialize node with partitioned disks.
(9c) Clean configuration and initialize node with whole disks.
(9d) Reboot the node.
(9e) Return to main boot menu.

Selection (9a-9e)?:
```

6. On one node, select option 9a to unpartition disks and remove the ownership information
7. Once the process is complete, move to the alternate node and repeat steps 5 and 6
8. On one node, clean the configuration and initialize the disks with the nodes
    a. For All Flash models, select option `9b`
    b. For Hybrid models using full drives, select option `9c`
    c. For Hybrid models using ADP, select option `9b`
9. Once the process is complete and the node has finished rebooting, move to the alternate node and repeat step 8.
   Note: Node should be finished rebooting if you see the cluster setup wizard prompt.
10. To perform a Data ONTAP OS install at this point, please refer to Appendix A on the process
    Note: This step is only necessary for nodes without an ONTAP image. ONTAP updates can be performed later once the cluster has been setup
11. Retrieve applicable licenses for the nodes
    Note: Licenses can be applied in the cluster setup wizard for the classic GUI. In the v4 GUI, licenses are applied after the cluster has been setup
    a. Navigate to https://fod2.lenovo.com/lkms/angular/app/pages/index.htm#/welcome
    b. Sign in using your Lenovo ID
    c. Select Features on Demand
    d. Then Select Retrieve History
    e. Select Search Type. Then change dialogue box to “Search History via machine type serial number”
    f. Type In your serial number in the Search value box in the format of: `MTMSerial Number`
    g. Click continue and the licenses will be retrieved
    h. You can then have all of the licenses exported to a file and emailed to you
12. Once both nodes show the cluster setup wizard, on one node:
    a. Acknowledge the AutoSupport message
       Note: To disable AutoSupport after the cluster has been setup, refer to Appendix B
    b. Select e0M as the port to host the node management port
    c. Enter an IP address, subnet mask and gateway
13. Continue the cluster setup by opening a web browser and going to `https://<node_mgmt_ip_address>`
14. Follow and complete the cluster setup wizard
    Notes:
    - For the classic GUI, at the bottom of the first menu, enter the license keys in a comma separated list for both nodes to apply the license keys previously retrieved from step 11. Each node will have its own set of license keys
    - Starting with ONTAP 9.7, a new GUI will be presented (v4). To access the classic GUI, look at the top banner and click “Return to classic version”
15. Upon completion of the wizard:
    a. Classic GUI: Click on the button to open the DM Series Storage Manager through the cluster management IP address to complete the process

b. New (v4) GUI: The webpage will automatically redirect to Storage Manager once the cluster has been setup

Appendix A: Installing Data ONTAP OS at boot menu in the serial console

1. Attach the included micro USB serial cables to both nodes and connect to PC

   Notes:

   - PC should be running Windows 10 or later
   - The COM Settings are 115200, 8-N-1
2. Power on all shelves in the DM stack
   a. If there are expansion shelves, power them on first and wait about one minute, then power on the head shelf
   b. If there are no expansion shelves, proceed to power on the head shelf

   Note: For DM7000 models, there is at least one expansion shelf to power on for the root aggregates
3. When the message stating AUTOBOOT is starting appears, hit Ctrl-C to interrupt the boot and get to the loader prompt; type `boot_ontap menu` to access the boot menu. Do this on both nodes
4. Select option 7 for installing new software first
5. Acknowledge the prompt and enter `e0M` as the temporary network port to download the ONTAP bundle through. Node will reboot after setting the temporary network port
6. Upon reboot, the node will boot to continue the ONTAP update without needing to interfere
7. Configure the temporary network port with an IP, subnet mask and default gateway
8. Enter the URL where the ONTAP update bundle is located at as well as credentials for access if necessary (for example, ftp site location)

   Notes:

   - ONTAP update bundle can be downloaded from the Lenovo Datacenter Support Site (https://datacentersupport.lenovo.com)
   - Ensure the remote location of the ONTAP bundle is accessible with the temporary network settings applied to the port from step 5
9. Upon successful connection the node will download and install the new ONTAP image
10. Enter "`y`" for the two prompts after installation success to set this image as the default image to boot to and reboot the node
11. Repeat steps 2-10 on alternate node if not already done and power cycle the DM head shelf before proceeding with any further steps

Appendix B: Disabling AutoSupport

To disable AutoSupport through CLI:

1. Check if AutoSupport is enabled on the nodes: `autosupport show`

```
cluster::> autosupport show
Node                  State     From          To            Mail Hosts
--------------------- --------- ------------- ------------- ----------
node-01           enable    Postmaster    -             mailhost
node-02           enable    Postmaster    -             mailhost
2 entries were displayed.
```

2. Disable AutoSupport on a node: `autosupport modify -node <node_name> -state disable`

```
cluster::> autosupport modify -node node-01 -state disable

Notice: AutoSupport configuration will be modified for all nodes in the
cluster.
```

3. Verify AutoSupport was disabled: `autosupport show`

```
cluster::> autosupport show
Node                  State     From          To            Mail Hosts
--------------------- --------- ------------- ------------- ----------
node-01           disable   Postmaster    -             mailhost
node-02           disable   Postmaster    -             mailhost
2 entries were displayed.
```

To disable AutoSupport through GUI:

Classic GUI:

1. Click Settings on the top banner in the top right
2. Navigate to AutoSupport under Setup in the navigation pane on the left
3. Click Disable to disable AutoSupport

New (v4) GUI:

1. In the navigation pane on the left, go to Cluster -> Settings
2. On the AutoSupport tile click on the 3 dots for settings
3. Click Disable to disable AutoSupport

Appendix C: Verifying HA mode on nodes

1. Attach the included micro USB serial cables to both nodes and connect to PC
   Notes:
   - You must use a PC running Windows 10 and above
   - The COM Settings are 115200, 8-N-1
2. Power on all the shelves in the DM stack:
   a. If there are expansion shelves, power them on first and wait about 1 minute, then power on the head shelf
   b. If there are no expansion shelves, proceed to power on the head shelf
   Note: For DM7000 models, there is at least 1 expansion shelf to power on for the root aggregates
3. When the message stating AUTOBOOT is starting appears, hit Ctrl-C to interrupt the boot and get to the loader prompt; type `boot_ontap menu` to access the boot menu. Do this on both nodes
4. Select option 5 for maintenance mode boot and acknowledge the warning
5. Should be presented with a prompt like: `*>`
6. Enter the command: `ha-config show` and verify ha mode is enabled
   a. `*> ha-config show`
   b. `Chassis HA configuration: ha`
   c. `Controller HA configuration: ha`
7. If necessary, set the mode to ha by executing:
   a. `ha-config modify chassis ha`
   b. `ha-config modify controller ha`
8. Exit the maintenance menu by entering halt on the command line
9. Repeat steps 4-8 on the alternate node

Appendix D: Setting the shelf ID on shelves

Before starting:

A valid shelf ID is 00 through 99.
Shelf IDs must be unique within a head shelf and the expansion shelves (if present).

- If you have a DM3000/DM5000, shelf IDs must be unique across the head shelf itself and any expansion shelves
- If you have a DM7000, shelf IDs must be unique across the expansion shelf containing the root aggregates and any expansion shelves

You must power cycle a shelf in order for the shelf ID to take effect.

- The amount of time you wait before turning the power back on depends on the state of ONTAP, as described later in this procedure.

1. Power on the disk shelf if it is not already on
2. Remove the left end cap (expansion shelf) or remove the bezel (head shelf) to locate the button near the shelf LEDs
3. Change the first number of the shelf ID by pressing and holding the orange button until the first number on the digital display blinks, which can take up to three seconds
   Note: If the ID takes longer than three seconds to blink, press the button again, making sure to press it in all the way
4. Press the button to change the number until you reach a desired number from 0 to 9
5. Change the second number of the shelf ID by pressing and holding the button until the second number on the digital display blinks, which can take up to three seconds
6. Press the button to advance the number until you reach a desired number from 1 to 9
7. Set the shelf ID and exit the programming mode by pressing and holding the button until the second number stops blinking, which can take up to three seconds
   - Both numbers on the digital display start blinking and the amber LED on the operator display panel illuminates after about five seconds, alerting you that the pending disk shelf ID has not yet taken effect
8. Power cycle the disk shelf to make the shelf ID take effect

You must turn off both power switches, wait the appropriate amount of time, and then turn them back on to complete the power cycle:

- If ONTAP is not yet running or you are hot-adding a shelf (that has not yet been cabled to the system), wait at least 10 seconds
- If ONTAP is running (controllers are available to serve data), and all disk drives in the shelf are unowned, spares, or part of offlined aggregate(s), wait at least 70 seconds

9. Replace the left end cap or bezel
10. Repeat steps 1-9 for each additional disk shelf
11. Verify that your system does not have duplicate shelf IDs

When two or more disk shelves have the same ID, the system assigns the duplicate disk shelf a soft ID number equal to or greater than 100. You must change the soft ID (duplicate) number

a. Run Config Advisor to check for duplicate shelf ID alerts or run the storage shelf show -fields shelf-id command to see a list of shelf IDs already in use including any duplicate IDs
b. If your stack has any duplicate shelf IDs, change the duplicate shelf IDs by repeating steps 1-9