Hewlett Packard Enterprise HPE Support Center

HPE 3PAR StoreServ 8000 Storage - Adding Drives and Expansion Drive Enclosures

Information about drive enclosure upgrades

3PAR StoreServ 8000 Storage products include 3PAR licensing that enables all functionalityassociated with the system. A failure to register the license key might limit access and restrictupgrading of your system. Before you proceed with upgrading, verify that all applicable licensesassociated with the system are registered.

For assistance with registering Hewlett Packard Enterprise software licenses, see the HewlettPackard Enterprise Support website.

Click here to access the Hewlett Packard Enterprise Support Center website .

There are two types of drive enclosures that are used for expansion:

- HPE 3PAR StoreServ Storage SFF 2.5-inch Drive Enclosure
- HPE 3PAR StoreServ Storage LFF 3.5-inch Drive Enclosure

- The number of drive enclosures attached to a specific controller node pair should bedetermined by the desired RAID set size and HA Cage protection requirements. Driveenclosures should be added and configured to meet the HA Cage protection requirementfor a specific controller node pair, and also consider the RAID set requirement of the customer.
- The distribution of drive enclosures between DP-1 and DP-2 of the controller node shouldbe done to achieve maximum balance across the ports.
- When adding both 2U and 4U drive enclosures, they should be mixed on SAS chains (DP-1and DP-2), added in pairs across controller node pairs on a 4-node system, and balancedacross SAS ports on each controller node pair.

top

Drive enclosure expansion limits

NOTE: Drives in the node enclosure are connected internally through DP-1.

Enclosure	Expansion Limits
8200	2-node - 9 drive enclosures total • Node DP1 = 4 drive enclosures connected • Node DP2 = 5 drive enclosures connected
8400	2-node - 11 drive enclosures total • Node DP1 = 5 drive enclosures connected • Node DP2 = 6 drive enclosures connected 4-node - 22 drive enclosures total • Node DP1 = 10 drive enclosures connected • Node DP2 = 12 drive enclosures connected
8440	2-node - 19 drive enclosures total • Node DP1 = 9 drive enclosures connected • Node DP2 = 10 drive enclosures connected 4-node - 38 drive enclosures total • Node DP1 = 18 drive enclosures connected • Node DP2 = 20 drive enclosures connected
8450	2-node - 9 drive enclosures total • Node DP1 = 4 drive enclosures connected • Node DP2 = 5 drive enclosures connected 2-node - 18 drive enclosures total • Node DP1 = 8 drive enclosures connected • Node DP2 = 10 drive enclosures connected

top

Information about drive upgrades

User can install additional drives to upgrade partially populated drive enclosures.

WARNING: If the 3PAR StoreServ Storage is enabled with the Data-at-Rest (DAR) encryptionfeature, only use the self-encrypting drives (SED). Using a non-self-encrypting drive might causeerrors during the upgrade process.

NOTE: SSDs have a limited number of writes that can occur before reaching the SSD's writeendurance limit. This limit is generally high enough so wear out will not occur during the expectedservice life of a 3PAR StoreServ Storage under the great majority of configurations, I/O patterns,and workloads. 3PAR StoreServ Storage tracks all writes to SSDs and can report the percent ofthe total write endurance limit that has been used. This allows any SSD approaching the writeendurance limit to be proactively replaced before they are automatically spared out. An SSD hasreached the maximum usage limit once it exceeds its write endurance limit. Following the productwarranty period, SSDs that have exceeded the maximum usage limit will not be repaired orreplaced under Hewlett Packard Enterprise Support contracts.

- In HA-Cage configurations the first expansion drive enclosure added to a system must be populated with the samenumber of drives as the controller node enclosure.
- The drives must be identical pairs.
- The same number of drives and type should be added to all of the drive enclosures in thesystem.
- The minimum addition to a 2-node system without expansion drive enclosures is two identicaldrives.
- The minimum addition to a 4-node system without expansion drive enclosures is four identicaldrives.

3PAR StoreServ 8000 Storage 2.5-inch Drive Enclosure Drive Placement

Drive pairs should be placed in the lowest available slot numbers.

Figure 1: 3PAR StoreServ 8000 Storage 2.5-inch Drive Enclosure drive placement order

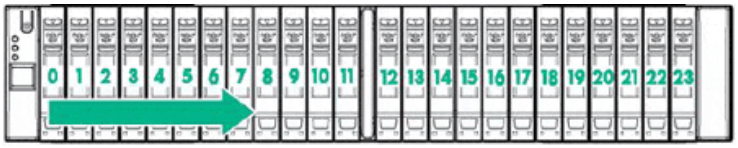

3PAR StoreServ 8000 Storage 3.5-inch Drive Enclosure Drive Placement

Drive pairs should be populated in columns and in the lowest available vertical slots in thatcolumn.

Figure 2: 3PAR StoreServ 8000 Storage 3.5-inch Drive Enclosure drive placement order

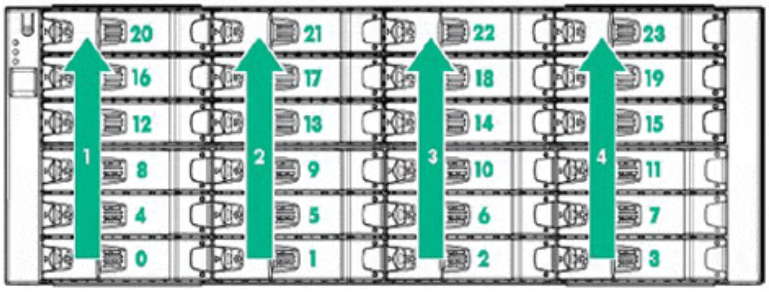

NOTE: For optimal utilization and performance, drive quantities and types should be configuredsymmetrically across all drive enclosures in the system.

top

Adding drives

There are five processes for adding drives:

- Checking initial status
- Inserting drives
- Checking status
- Checking progress
- Completing the upgrade

1. In the SSMC main menu, select **Storage Systems** > **Systems**. A list of storage systems isdisplayed in the list pane.
2. In the **Systems** filter, select the storage system.
3. In the detail pane, select the **Configuration** view.
4. In the **Physical Drives** panel, click the total physical drives hyperlink. The **Physical Drives**screen is displayed.
5. In the list pane, select a physical drive to display its properties in the detail pane.
6. Install the drives.

The display refreshes periodically, and you should see the inserted drives as **New** in the**State** column. They are ready to be admitted into the system, which occurs automatically.

Within six minutes (depending on the system load and the size of the upgrade), the **State**of the new drives changes to **Normal**, and the system starts to initialize the chunklets toready for use.

Chunklet initialization can take several hours to complete and the output of the availablecapacity is displayed.

NOTE: The system can be used normally, but newly added capacity must be initializedbefore it can be allocated.

top

Adding expansion drive enclosures

Before you begin: Connect to the service processor and start an SPMAINT session.

1. From the SPMAINT home page, enter **7** for **Interactive CLI for a StoreServ**, and thenchoose the system.
2. When prompted, enter **y** to enable maintenance mode for the system.
3. Enter `checkhealth -detail` to verify the current state of the system.
4. Install the expansion drive enclosure.
5. Add the storage drives to the front of the drive enclosure bays, starting from bay 0 andcontinuing sequentially.
6. Connect the mini-SAS cables from the appropriate controller node SAS ports to the I/Omodules in the rear of the drive enclosure according to guidelines.
7. Label the mini-SAS cables according to guidelines.
8. Install the 580 W PCMs into the drive enclosure.
9. Connect the power cables to the power supplies and secure them with cable straps.
10. On the PCMs, press the power switch to ON.

    NOTE: It might take up to five minutes for the drive enclosure to fully power up with thestorage drives available to the cluster.

11. Verify that the drive enclosure, fan module, I/O modules, drive enclosure, power supplies,and drive status LEDs are lit green and operating normally.
12. To verify that the upgrade is successful, issue the `showcage` command to identify the newcage ID of the newly installed drive enclosure.
13. Enter `showport` to verify that the connected ports are ready.
14. Enter `showcage -d` followed by `<cageID>` of the new drive enclosure, to confirm the drivesin the new drive enclosure are spun-up and seen by the system.
15. Enter `showpd` to verify that the new drives appear and the disk state is normal.
16. Enter `checkhealth -detail` to verify the current state of the system.
17. Enter `exit` and press **Enter** to continue to the main menu. Select **X** to exit from the 3PARService Processor Menu and to log out of the session.

top