



Hewlett Packard
Enterprise

HPE ConvergedSystem 500 for SAP HANA Scale-out Configurations

Administration, Operations, and Troubleshooting Guide

Abstract

This document describes the administration, operation, and troubleshooting of the HPE ConvergedSystem 500 for SAP HANA Scale-out configurations.

Part Number: 774752-008
Published: December, 2016
Edition: 9

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel[®], Itanium[®], Pentium[®], Intel Inside[®], and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft[®] and Windows[®] are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe[®] and Acrobat[®] are trademarks of Adobe Systems Incorporated.

Java[®] and Oracle[®] are registered trademarks of Oracle and/or its affiliates.

UNIX[®] is a registered trademark of The Open Group.

Contents

- Introduction..... 5**
 - CS500 Scale-out introduction..... 5
 - SAP HANA license statement..... 6
 - CS500 Scale-out components..... 6
 - 8TB single rack configuration..... 6
 - 12TB expansion rack configuration..... 8
 - Networking..... 9
 - CS500 Scale-out storage..... 10
 - Network file servers..... 13

- Maintenance and operations tasks and responsibilities..... 14**
 - Support..... 15
 - Tools..... 15
 - SAP HANA Studio..... 15
 - YaST2..... 16
 - HPE 3PAR Management Console..... 16
 - Secure terminal emulator and file copy..... 16
 - X windows server..... 16
 - IP addresses..... 16
 - Managing licenses..... 16
 - Managing users..... 17
 - Updates..... 17
 - System power guidelines..... 17
 - Power-on sequence..... 17
 - Power off sequence..... 18
 - Backups..... 18
 - Managing data and log files..... 18
 - Creating instances volumes..... 20
 - Secure Encryption Drives on ConvergedSystem for SAP HANA appliances..... 22
 - Enabling encryption on self encrypting drives (SED) drives using the 3PAR CLI..... 24
 - Enabling encryption on SED drives using the 3PAR management GUI..... 24

- Troubleshooting..... 26**
 - Common problem resolution..... 26
 - Service notifications..... 26
 - Release Notes..... 26

- Resources..... 27**
 - References..... 27

- Safety considerations..... 28**
 - Important safety information 28
 - Symbols on equipment..... 28
 - Electrostatic discharge..... 29
 - Preventing electrostatic discharge..... 29

Grounding methods to prevent electrostatic discharge.....29

Support and other resources..... 30

Before you contact Hewlett Packard Enterprise..... 30
Proactive Care..... 30
 Requesting support for HPE ConvergedSystem and HPE Hyper Converged products..... 30
 Support Case Manager PINs..... 32
Remote support..... 33
Warranty information..... 33
Documentation feedback..... 33

Introduction

CS500 Scale-out introduction

HPE ConvergedSystem 500 for SAP HANA Scale-out Configuration provides a highly available and resilient HANA in-memory database. This appliance has a range of 2 to 16 nodes. Each node can be 1TB, 1.5TB, and 2TB yielding a database memory size of approximately 2TB to 16TB for 1TB server model up to a 4TB to 32TB for the 2TB server model. Maximum scalability can be expanded beyond 16 nodes in a range of 17-34 HANA nodes for very large solutions.

These configurations, in an appliance delivery model, provide the following benefits:

- Includes fully vetted hardware and software infrastructure architected by HPE product engineers to host HANA and certified by SAP.
- Assembled and tested as a unit before being shipped to customer datacenters.
- Installed and configured on-site by Hewlett Packard Enterprise.
- Each release is subject to a full suite of validation and verification tests.
- Provides redundant network pathways for resilient highly available access to components and data.
- Provides disk redundancy using mirroring or parity.
- Each compute engine is an HPE ProLiant DL580 Gen9 four socket server with 1TB, 1.5 TB, and 2 TB of RAM. With the 4 processors there are a total of 96 cores of high performance computing in each node. In a four node configuration, there are 384 processing cores, and in the eight node configuration there are 768 processing cores.
- With the current solution architecture, the base or starter rack can contain 4 DL580 Gen9 active HANA nodes and an optional 1 DL580 Gen9 standby node.
- Each expansion rack can contain up to 6 DL580 Gen9 HANA nodes. A sixteen HANA node solution can be enclosed within 3 racks.
- A four node base configuration includes an HPE 3PAR StoreServ 8400 four node controller with eighty 1.2TB 6G SAS disks for maximum availability and performance. Optional drives can be added in increments of 16, 32, or 48 drives for 3PAR storage backups.
- A six node expansion configuration includes an HPE 3PAR StoreServ 8400 four node controller with ninety six 1.2TB 6G SAS disks for maximum availability and performance. Optional drives can be added in increments of 16, 32, or 48 drives for 3PAR storage backups.

An in-memory database requires high performance storage for data and log storage to preserve the integrity and availability of information across shutdown and fail-over scenarios. The ConvergedSystem 500 for SAP HANA Scale-out Configuration relies on 3PAR StoreServ 8400 storage for OS boot from SAN and HANA log and data, and HPE Networking switches to best meet this requirement. Each compute engine is connected through redundant storage area network (SAN) paths to the highly available 3PAR storage unit with redundant controllers. The same 3PAR StoreServ 8400 storage is configured for HANA shared NFS file system across all HANA nodes in the cluster. The NFS file share can support 16 nodes and scaling can support up to 34 nodes (by adding additional 16 disk drives).

The SAP HANA platform, running on SUSE or Red Hat Enterprise Linux, provides the basis to dramatically increase the performance of your applications. Customers are adopting SAP HANA quickly because they value the ability to aggregate and analyze data in real time. HANA compression, in-memory caching, and connections to big data technologies provide key elements in the formula for faster response times.

SAP HANA is a database and must be integrated into your environment and applications. The purchase of a ConvergedSystem 500 for SAP HANA Scale-out Configuration includes the installation and startup of HANA. Loading operational data and integration with your applications are additional to this effort. It is important for the success of your SAP HANA implementation to engage the Hewlett Packard Enterprise sales team to plan the extraction, transformation, and loading of data into SAP HANA and integration with your applications.

SAP HANA license statement

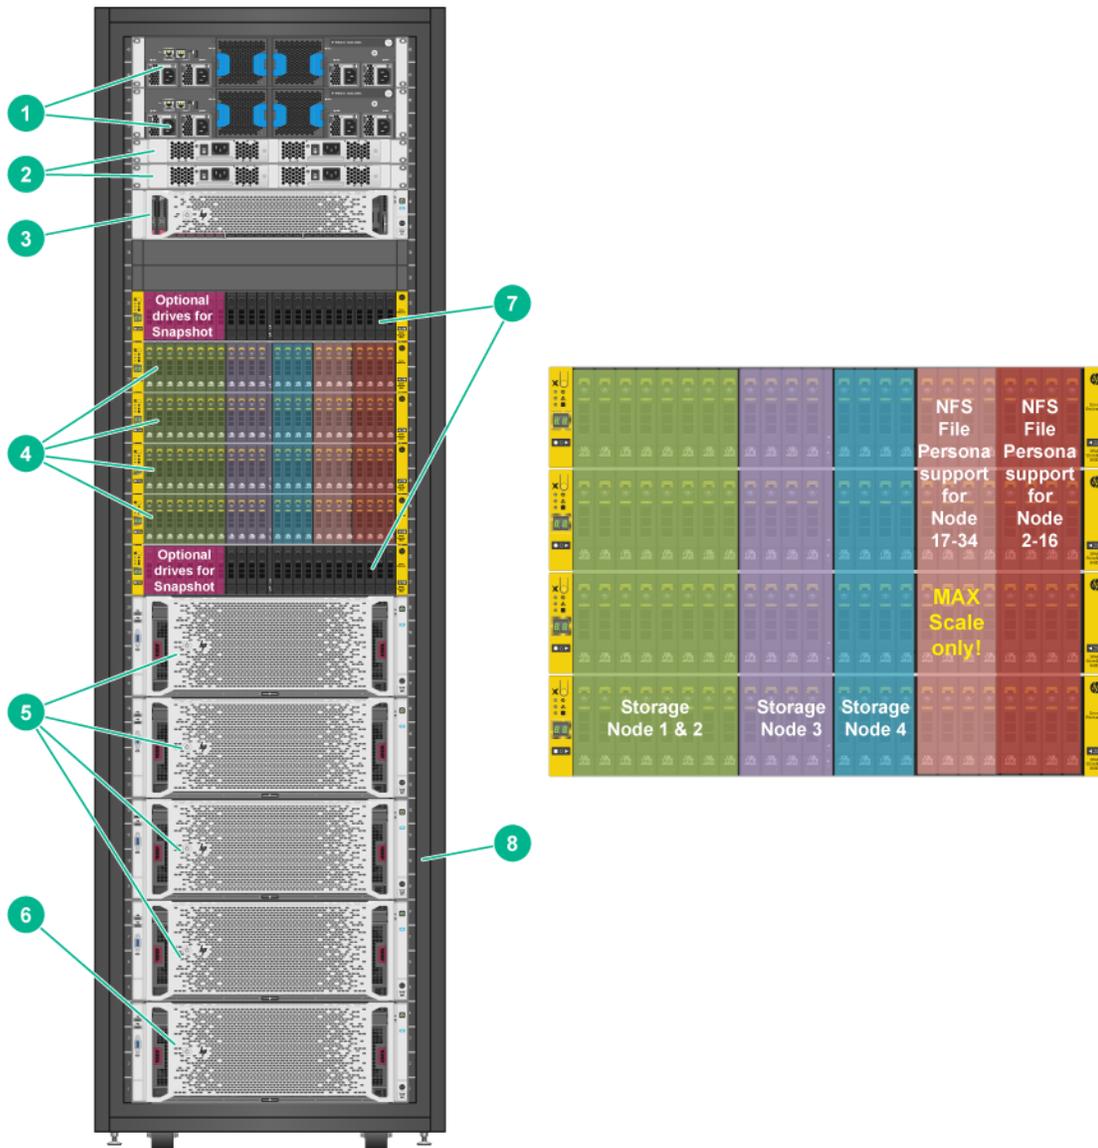
The Hewlett Packard Enterprise computer system contains a copy of the SAP HANA software by SAP. Each SAP product is subject to its respective SAP end-user license agreement. The HPE computer product license does not contain a license to SAP HANA software by SAP. Please be advised that you have no right to use and that you are not licensed to use the copy of the SAP HANA software by SAP contained in the Hewlett Packard Enterprise computer system until you have purchased or have available for use the appropriate license from SAP or its authorized distributors. Contact your responsible SAP representative to obtain the applicable license rights to use the SAP software.

CS500 Scale-out components

This section provides an overview of components in an 8TB base rack configuration with a single expansion rack, which is able to extend the solution to 20TB. Figure 1 provides a view of the 8TB base rack configuration with an optional 5th standby node. Figure 2 provides a view of an expansion rack for an extended 12TB configuration.

8TB single rack configuration

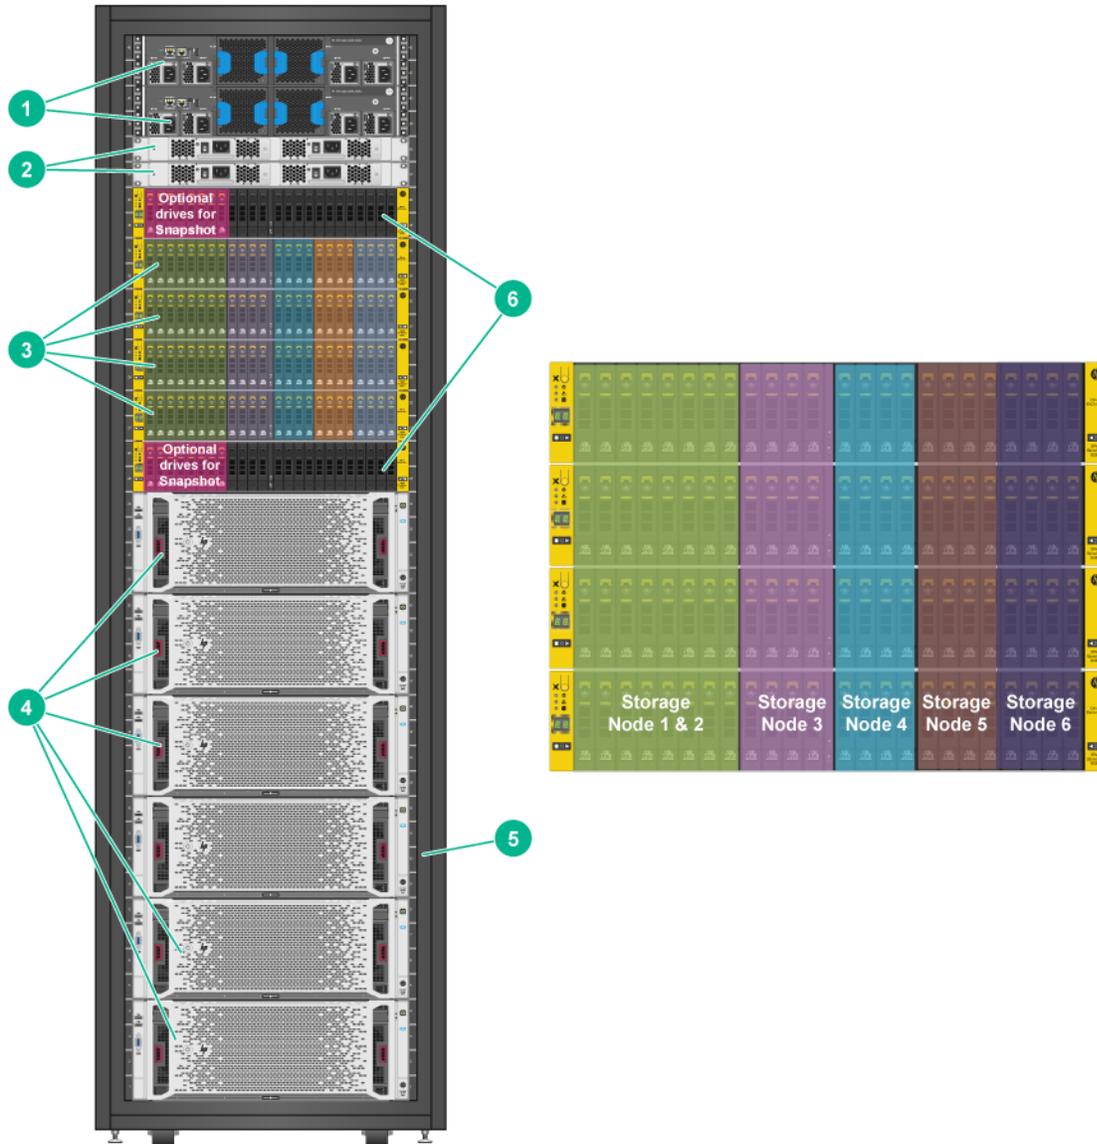
Figure 1 - Base rack view with an 8TB configuration using HPE 3PAR StoreServ and an optional 5th standby HANA node. HPE 3PAR 8400 disk configuration with respect to 4 nodes and NFS configuration.



Component	Description
1	FlexFabric 5930-4Slot Switch (2x)
2	SN6000B 48p SAN Switch (2x)
3	Management Server (HPE ProLiant DL380p Gen9) (1x)
4	HPE 3PAR 8400-4node FC Storage 80x 1.2TB HDD
5	HPE ProLiant DL580p Gen9 (4x)
6	HPE ProLiant DL580p Gen9 (1x) optional standby node
7	HPE 3PAR optional (16) drives
8	8TB Building Block (4 node)

12TB expansion rack configuration

Figure 2 - Expansion rack - 12TB configuration view with 6 HANA nodes. HPE 3PAR 8400 disk configuration with respect to 6 nodes.



Component	Description
1	FlexFabric 5930-4Slot Switch (2x)
2	8/24 Base 24p SAN switch (2x)
3	HPE 3PAR 8400-4node 4 shelves FC Storage 96x 1.2TB HDD
4	HPE ProLiant DL580p Gen9 (6x)
5	12 TB Building Block (6 node)
6	HPE optional 3PAR storage drives and enclosures

The switches provide redundant paths within the appliance among the components, and the components leverage these redundant paths by having multiple connections and failover capabilities.

The ProLiant DL580 Gen9 server provides a hosting environment for the software and server networking components. Each server is accessible through iLO and management ports. In addition to iLO and management, each DL580 Gen9 has two Mellanox 2-port 40Gb remote direct memory access (RDMA) capable network cards installed. Whenever SAP HANA is capable of providing RDMA capability the CS500 Scale-out solution will be ready to handle this networking protocol. Currently the Mellanox cards are configured for TCP Ethernet which provides 33Gb per bonded port pair. Each of the two ports are configured for bonding across cards and common ports. The servers are configured for redundancy and the networks are paired: BI/Client, ERP, and Replication have access over the 33Gbps network bond1 that provides each of three vlan networks with 10Gbps plus performance. The Internal HANA node, HANA Shared and Backup vlan networks have access across another 33Gbps network bond0 for Ethernet traffic, while data and log traffic is over a 16Gbps Fiber Channel. Internal network has a QoS setting to guarantee it will always get 10Gbps or better throughput.

For redundancy, some configurations may include a spare or standby node. For example, in a four node starter rack environment, four nodes may be active and an optional node is available as a standby node for activation when one of the original nodes fail, or the database instance on that server fails. SAP HANA may then select the standby node and activate the database instance on it. The standby node then becomes an active member in the cluster and mounts the data and log files from the previously active and now failed node or instance.

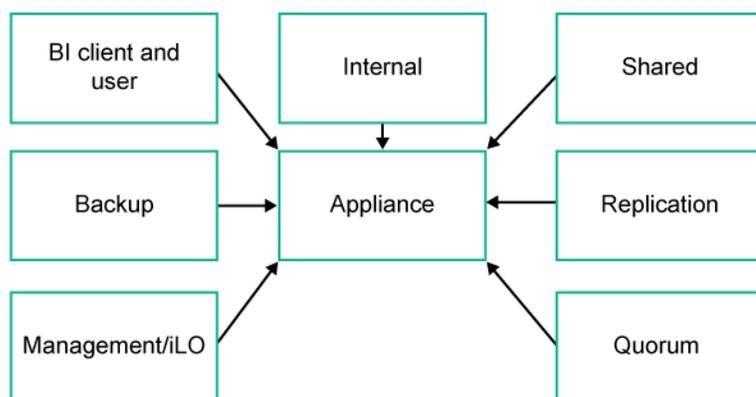
Networking

Hewlett Packard Enterprise collects a set of parameters for each customer installation with a few details about setup requirements. One of the requirements is to collect networking details such as addresses and the type and quantity of networks required for a specific customer implementation.

A distributed architecture requires solid networking. Hewlett Packard Enterprise has architected this solution with 10GbE throughput among the nodes and NFS servers, and 16Gb Fibre Channel SAN connections to the HPE 3PAR storage. There are a total of seven networks configured in the appliance. Two networks, internal and shared, are private networks and are internal to the cluster. Five external customer networks may connect to the appliance, as shown in Figure 3.

Five external networks may connect from the customer network to the appliance, as shown in Figure 3.

Figure 3 - SAP HANA appliance network customer touch points.



All data and BI-connections are 10GbE throughput capable. These selections help isolate the traffic types onto their LANs as best fits your environment. Management and iLO networks are merged on the same management network. If there is a requirement to keep them as separate managed networks, then the factory or Hewlett Packard Enterprise consulting can configure them accordingly. These connections are typically set

up at the Hewlett Packard Enterprise factory prior to delivery of the appliance or can be configured at installation time.

Table 1: Locations within the appliance to connect to the customer network.

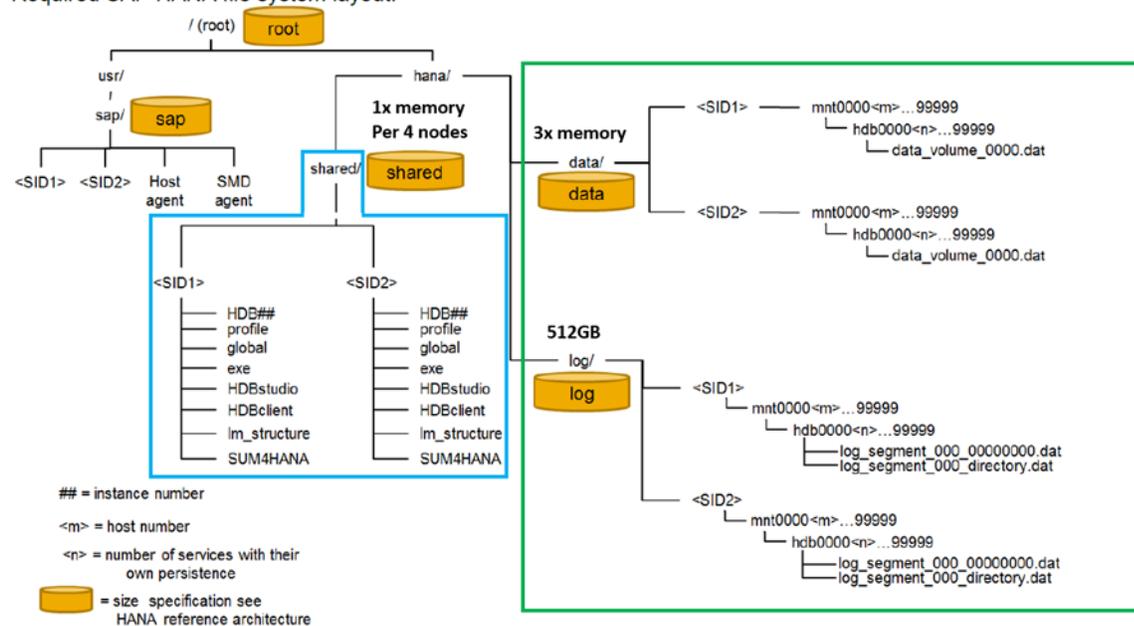
Traffic	Vlan(s)
Bi client and user	10G fiber
Data	10G fiber
Backup	10G fiber
Replication	10G fiber
Quorum	1G fiber
Internal and shared	10G fiber (internal to the appliance only)
Management iLO	1G fiber

CS500 Scale-out storage

The file system layout for the ConvergedSystem 500 for SAP HANA Scale-out Configuration with 3PAR is shown in Figure 4. The operating system is configured to boot from 3PAR SAN at RAID1; the shared directory is mounted from the 3PAR storage where redundancy is built into a four internal node solution; and the data and log directories are mounted using the storage API directly by HANA. The configuration used for data and log files is known as non-shared storage, which means that each server and HANA instance has control over its own data and log files.

Figure 4: SAP file system layout with edits for deployment overlay.

Required SAP HANA file system layout:



Data and log directories are under the control of HANA and are not entered in the /etc/fstab file where Linux file systems are normally entered for mounting. The connections to data and log are described in the HANA global.ini file. The path to the file is dependent on the system identifier (SID) of the SAP HANA installation: /hana/<SID>/global/hdb/custom/config/global.ini. For more information on the SAP storage API that enables this direct mounting of the data and log files, see SAP HANA Storage Requirements in "References."


```
[root@hana11-int ~]# /sbin/multipath -ll | /usr/bin/grep "3PARdata,VV" | /usr/bin/sort
bootLUN (360002ac00000000000000003300018eb3) dm-0 3PARdata,VV
r1_data1.1 (360002ac00000000000000006f00018eb3) dm-8 3PARdata,VV ←
r1_data1.2 (360002ac00000000000000007000018eb3) dm-12 3PARdata,VV
r1_data1.3 (360002ac00000000000000007100018eb3) dm-13 3PARdata,VV
r1_data1.4 (360002ac00000000000000007200018eb3) dm-9 3PARdata,VV
r1_log1.1 (360002ac00000000000000006500018eb3) dm-5 3PARdata,VV
r1_log1.2 (360002ac00000000000000006600018eb3) dm-11 3PARdata,VV
r1_log1.3 (360002ac00000000000000006700018eb3) dm-6 3PARdata,VV
r1_log1.4 (360002ac00000000000000006800018eb3) dm-7 3PARdata,VV
```

The next step is to use LVM and the user friendly names to create the physical volumes volume groups as shown in the output in figure 6.

Figure 7: Results of LVM creation of devices on the HANA servers

```
pvs
/dev/mapper/r1_data1.1      vg_data1_11  lvm2 a--    6.00t    0 ←
/dev/mapper/r1_data1.2      vg_data1_12  lvm2 a--    6.00t    0
/dev/mapper/r1_data1.3      vg_data1_13  lvm2 a--    6.00t    0
/dev/mapper/r1_data1.4      vg_data1_14  lvm2 a--    6.00t    0
/dev/mapper/r1_log1.1       vg_log1_11   lvm2 a--   519.75g  0
/dev/mapper/r1_log1.2       vg_log1_12   lvm2 a--   519.75g  0
/dev/mapper/r1_log1.3       vg_log1_13   lvm2 a--   519.75g  0
/dev/mapper/r1_log1.4       vg_log1_14   lvm2 a--   519.75g  0

lvs
lvoll vg_data1_11  -wi-ao----  6.00t ←
lvoll vg_data1_12  -wi-a-----  6.00t
lvoll vg_data1_13  -wi-a-----  6.00t
lvoll vg_data1_14  -wi-a-----  6.00t
lvoll vg_log1_11   -wi-ao----  519.75g
lvoll vg_log1_12   -wi-a-----  519.75g
lvoll vg_log1_13   -wi-a-----  519.75g
lvoll vg_log1_14   -wi-a-----  519.75g
lvoll vg_ursap     -wi-a-----  63.00g
lvoll vg_ursapsid1 -wi-a-----  63.75g
```

The last step in getting the storage volumes ready for HANA use is to create the filesystems using mkfs.xfs.

The following is an example of creating the first log and data volume:

```
/sbin/mkfs.xfs -f -i size=2048 -b size=4k -s size=4k /dev/mapper/vg_log1_11-
lvoll
/sbin/mkfs.xfs -f -i size=2048 -b size=4k -s size=4k -l su=256k /dev/mapper/
vg_data1_11-lvoll
```

This is repeated for all the data and log volumes for each rack of HANA servers.

The SAP HANA global.ini file is created with the lvm naming and not the WWN as shown below:[storage]

```
ha_provider=hdb_ha.fcClientLVM
partition_*_* prType = 5
partition_1_data lvmname = vg_data1_11-lvoll
partition_1_data mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swaloc
partition_1_log lvmname = vg_log1_11-lvoll
```

```
partition_1_log mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc
```

NOTE:

When using LVM for RHEL / SLES the global.ini file requires fcClientLVM as the storage type.

The data and log files are mounted to the individual servers. While all entries are visible to all servers, only the appropriate data and log file are mounted for that server and instance. They are visible to allow an alternate, or standby, server to mount the data and log file for a failed server or instance during the recover process. This recovery process is described in SAP HANA Storage Requirements. For more information, see "References."

To provide redundant paths, these paths are configured using multi-pathing. For more information on multi-pathing, see SAP HANA Storage Requirements "Managing Multipath I/O for Devices."

A device's WWID [WWN], user-friendly name, and alias name persist across reboots, and are the preferred way to identify the device. The use of the LVM devices is recommended for consistency across nodes.

Network file servers

The /hana/shared directory is mounted as a NFS client on each DL580 Gen9 using redundant NFS hosted on 3PAR File Persona storage server; two 3PAR 8400 internal nodes are configured for high availability. One 3PAR node is active and the other one is standby. These two 3PAR nodes connect redundantly to the HPE FlexFabric 5930 switch using 10Gbs Ethernet, for HANA node access to the HANA shared data. If a node in the 3PAR fails the standby node seamlessly takes over.

For more information on 3PAR File Persona, see "References."

Maintenance and operations tasks and responsibilities

❗ **IMPORTANT:**

As a best practice, Hewlett Packard Enterprise advises you to follow data center best practices to connect Hewlett Packard Enterprise servers, storage, and switches for power, networking, and fiber channel. Doing so will minimize downtime and inconveniences.

❗ **IMPORTANT:**

Review the terms of your support contract to help you understand how best to work with Hewlett Packard Enterprise and SAP to support your environment and to understand the tasks necessary to support the appliance in your environment.

Each component is administered using the appropriate methods described in the manuals for that device. The following tables provide further information for the administration and operation of this HPE Converged System appliance as a whole. This information complements the *SAP HANA Technical Operations Manual*, Section 2.1, "Roles and Responsibilities".

Provisioning and setup tasks

Task	Responsible
Installation of Hardware	Hewlett Packard Enterprise
Installation of Operating System	Hewlett Packard Enterprise
Installation of SAP HANA Platform	Hewlett Packard Enterprise
Adding additional SAP HANA database instances (MCOS)	Customer
Data Source Connectivity	Customer

Maintenance tasks

Task	Responsible
Patching of Firmware	Hewlett Packard Enterprise
Patching of Operating System	Customer ¹ (Optional HPE Technical Services)
Patching of SAP HANA Platform components	Customer
Patching of peripheral components	Customer
Upgrade of Operating System	Hewlett Packard Enterprise
Update of SAP HANA platform components	Customer

¹ *The customer is generally responsible for maintenance of the SAP HANA system. If the customer has a special support agreement with the hardware partner, maintenance may be the responsibility of the hardware partner.*

Operations tasks

Task	Responsible
General administration of SAP HANA database	Customer
Backup and Recovery	Customer
SAP HANA system replication	Customer
SAP HANA System Monitoring	Customer
SAP HANA Database Monitoring	Customer
Installation of third-party software components	Customer
Support	Customer
Issue Handling Process	SAP, Customer ¹

¹ SAP is the main point of contact and distributes all issues within the support organization by default, as is the case for other SAP applications. If the customer has defined special support agreements with the hardware partner, the customer should contact the hardware partner directly in case of obvious hardware or operating system issues. If no agreements have been made, neither SAP nor the hardware partner is responsible for the installation, maintenance, and possible adjustment of external software installed on the SAP HANA system.

During installation and startup, the hardware, operating system, and SAP HANA Platform are installed in the Factory Express and On-site startup services.

During provisioning and startup, Hewlett Packard Enterprise, another party, or the customer may add additional SAP HANA database instances and data source connectivity.

During ongoing maintenance and operations, the maintenance and update activities are dependent upon the support agreement with Hewlett Packard Enterprise and SAP.

Hewlett Packard Enterprise SAP Consulting is available to help with additional configuration requirements such as high availability, disaster recovery, backup and recovery strategies, connectivity with other SAP solutions.

Support

Hewlett Packard Enterprise provides support services to help protect your investment in this appliance. Contacts for support services for the software and firmware updates for the appliance are provided in your support agreement.

Tools

Tools must be installed on the central management server (CMS) before you can run some tasks. These tools are subject to the site security requirements and network protocol. Follow all security protocols, and if approved, apply the approved security and anti-virus processes before exposing the system to the external network, and before connecting to an internet domain or downloading software to the system.

SAP HANA Studio

SAP HANA Studio is the tool used for general administration and monitoring functions related to SAP HANA. The system includes a Linux version of SAP HANA Studio that is installed on the delivered appliance.

To install SAP HANA Studio, see the SAP HANA Studio Installation Guide. You also need to download Java JRE to run this tool on the CMS. Many administrative functions are provided by SAP HANA Studio. Follow the instructions available in the SAP HANA Administration Guide and the SAP HANA Update and Configuration Guide, as needed. For more information, see "References."

YaST2

YaST2 (Yet another Setup Tool) is the operating system setup and configuration tool for the SLES distribution. Use this tool to perform normal operating system maintenance. This tool is not available in RedHat distributions. Please consult RedHat documentation for specific maintenance procedures for operating system maintenance.

HPE 3PAR Management Console

The HPE 3PAR Management Console is for use by system and storage administrators who monitor and direct system configurations and resource allocation for 3PAR Storage Systems. The storage system is set up in the factory for proper operation with SAP HANA.

There is also an additional management tool for monitoring the 3PAR environment including the 3PAR File Persona. The new 3PAR StoreServ Management Console (SSMC) streamlines management tasks using a simple, clean view of the array. SSMC was designed through a user-centered methodology, which focuses on the common tasks for mainstream users. This fresh approach to storage management better utilizes the IT operations staff and allows them to concentrate on other areas of the data center.

Secure terminal emulator and file copy

From an appropriate source, obtain and install a secured terminal emulator (ssh) and secured file copy (scp) on the CMS.

These tools enable access to the LAN, SAN, and other management ports and downloading configuration files.

X windows server

From an appropriate source, obtain and install an X windows server on the CMS. This tool enables some graphical based access.

IP addresses

The IP addresses for the components in the appliance are listed in the Smart CID.

Managing licenses

For SAP HANA, see the “Managing SAP HANA Licenses” section in the SAP HANA Administration Guide.

For SLES for SAP, this process should be handled upon delivery of the system. Hewlett Packard Enterprise offers the SLES for SAP subscriptions for electronic delivery. For the registration process, perform the following steps:

1. Receive by email your Hewlett Packard Enterprise License Entitlement Certificate that contains your Entitlement Order Number.
2. Go to the Hewlett Packard Enterprise Software Licensing Portal to retrieve your SUSE activation token located on the **Hewlett Packard Enterprise website**.
3. To activate your subscriptions after you retrieve your activation token, see the **SUSE website**.

For RHEL for SAP, this process should be handled upon delivery of the system. Hewlett Packard Enterprise offers the RHEL for SAP subscriptions for electronic delivery. For the registration process, perform the following steps:

1. Receive by email your Hewlett Packard Enterprise License Entitlement Certificate that contains your Entitlement Order Number.
2. Go to the My Hewlett Packard Enterprise Licensing Portal to retrieve your RHEL activation key located on the **Hewlett Packard Enterprise website**.
3. Once you retrieve your activation key, go to the **Red Hat website**.

Managing users

The default user identifiers and passwords for the components in the appliance are listed in the Customer Intent Documentation (CID).

There are two user sets for SAP HANA and SLES, they are managed accordingly.

For SAP HANA, see the SAP HANA Administration Guide, "Managing SAP HANA Users".

For SLES, see the YaST User and Group Administration dialog.

For RHEL, see "Managing Users and Groups" in the Red Hat Enterprise Linux Deployment Guide.

For the other components, see the appropriate user guide.

Updates

Hewlett Packard Enterprise Support offers a variety of update and upgrade services to keep your HPE ConvergedSystem for SAP HANA appliance running at peak performance. Contact your local Hewlett Packard Enterprise sales representative for services details.

System power guidelines

Hewlett Packard Enterprise recommends following these guidelines before powering your system on or off.

- Do not adjust the position of the cabinet when the power is on.
- Never power off or remove a controller node without confirming that the service LED is solid blue.
- Wait until all controller nodes are in a halted state before powering the system down. Failure to do so could cause the system to view the shutdown as uncontrolled and place the system in a "checkld" state upon power up. This can seriously impact host access to data.
- The system takes approximately five minutes to become fully operational providing it was gracefully shut down. If the system was powered off abruptly, powering on could take considerably longer.

Power-on sequence

1. Power on the cabinets.
2. Power on the switches.
3. Power on the 3PAR storage.

Verify that all drive chassis LEDs are solid green and all controller node status LEDs are blinking green once per second.

NOTE:

The storage system takes approximately five minutes to become fully operational, if it was gracefully shut down. If the system was powered off abruptly, powering on could take longer.

4. Power on the 3PAR File Persona NFS server.
5. Power on the HANA nodes and wait for them to fully boot.
6. Verify that all components are started.
7. Start HANA using the SAP HANA Studio.

Power off sequence

Procedure

1. Stop HANA using SAP HANA Studio and wait for the HANA processes to stop.
2. For each HANA node, shut down RHEL / SLES and wait until servers have completely shut down.
3. Power off the 3PAR File Persona NFS server as the 3paradm user using the following commands:
 - a. Connect the maintenance PC to the SP using the serial connection.
 - b. Start a `spmaint` session.
 - c. In the 3PAR Service Processor menu, select option 4. StoreServ Product Maintenance.
 - d. Select option 6. Halt a StoreServ cluster/node, then select the preferred system.
 - e. Confirm all prompts to halt the system.
 - f. Enter x to return to the 3PAR Service Processor menu.
4. Power off the 3PAR Storage.
 - a. Connect the maintenance PC to the SP using the serial connection.
 - b. Start a `spmaint` session.
 - c. In the 3PAR Service Processor menu, select option 4. StoreServ Product Maintenance.
 - d. Select option 6. Halt a StoreServ cluster/node, then select the preferred system.
 - e. Confirm all prompts to halt the system.
 - f. Enter x to return to the 3PAR Service Processor menu.
5. Turn off the 3PAR Storage power switch.
6. Power off the switches.
7. Power off the cabinets.

Backups

Use the site recommended tools and procedures to back up the HANA nodes and NFS servers, and SAP HANA Studio or site recommended tools to back up the database. If you require more information on the backup and recovery options, contact Hewlett Packard Enterprise.

The restoration DVD also contains tar files to enable you to restore to a factory ship state.

Managing data and log files

Review the file system layout in the Storage section. The following table shows the 3PAR storage assigned for a four node system using RAID 1 with 64 x 1.2TB storage configuration. The “hanasharedfpg.1” Virtual Volume (VV) for NFS has a size of 12TB with the 16 x 1.2TB storage using RAID 5 configuration.

Table 2: Four node storage configuration using RAID 1 with 64 x 1.2TB storage configuration

Name	LUN ID	Prov	V size (GB)	RAID level
R1_OS.1	0	full	256	1
R1_OS.2	0	full	256	1
R1_OS.3	0	full	256	1
R1_OS.4	0	full	256	1
R1_USRSAP.1	71	full	64	5
R1_USRSAP.2	72	full	64	5

Table Continued

Name	LUN ID	Prov	V size (GB)	RAID level
R1_USRSAP.3	73	full	64	5
R1_USRSAP.4	74	full	64	5
R1_LOG1.1	101	full	520	5
R1_LOG1.2	102	full	520	5
R1_LOG1.3	103	full	520	5
R1_LOG1.4	104	full	520	5
R1_DATA1.1	111	full	6,144	5
R1_DATA1.2	112	full	6,144	5
R1_DATA1.3	113	full	6,144	5
R1_DATA1.4	114	full	6,144	5
R1_USRSAPSID1.1	121	full	64	5
R1_USRSAPSID1.2	122	full	64	5
R1_USRSAPSID1.3	123	full	64	5
R1_USRSAPSID1.4	124	full	64	5

If requested a development or a test environment, the factory can add up to three thin provisioned volumes for up to three HANA instances to the 3PAR storage. The following table shows an example of the thin provisioned volumes for instance 2 and 3 for a four node system. Create the volume sizes based on the memory size of the HANA SID to be used. Normally the sizes are smaller than the full memory capacity. All test/dev instances must be sized so that the total of all instances does not exceed the total memory/storage of the appliance. These instances are configured as thin provisioning. In Table 3, for DATA.<INSTANCE#> as the suffix of LOG or DATA for additional volumes.

Table 3: Four node storage configuration thin provisioned volumes for instance 2 and 3 for a four node system

Instance	Name	Lun ID	Provisioning	V size (GB)
2	R1_LOG2.1	201	Thin	520
2	R1_LOG2.2	202	Thin	520
2	R1_LOG2.3	203	Thin	520
2	R1_LOG2.4	204	Thin	520
2	R1_DATA.2.1	211	Thin	1,536
2	R1_DATA.2.2	212	Thin	1,536
2	R1_DATA.2.3	213	Thin	1,536
2	R1_DATA.2.4	214	Thin	1,536
2	R1_LOG3.1	301	Thin	520
3	R1_LOG3.2	302	Thin	520
3	R1_LOG3.3	303	Thin	520

Table Continued

Instance	Name	Lun ID	Provisioning	V size (GB)
3	R1_LOG3.4	304	Thin	520
3	R1_DATA3.1	311	Thin	1,536
3	R1_DATA3.2	312	Thin	1,536
3	R1_DATA3.3	313	Thin	1,536
3	R1_DATA3.4	314	Thin	1,536

NOTE:

For ten node appliance 3PAR storage names, the second rack names are changed from R1 to R2. For a sixteen node, the third rack names are change from R1 to R3.

NOTE:

Table 3 is an example of multi SID environment. The volume size might vary in the customer environment.

Creating instances volumes

The following procedure is used to create data and log volumes for development and testing. These volumes are already created in the factory. If you need to recreate the volumes, use the following steps.

On 3PAR, add data and log volumes for non-production systems using the current hostset. Use the following steps. It is not necessary to add SYS, or OS volumes, only the data and log volumes. The commands used are explained in the 3PAR Command Line Reference. These commands can also be run in the 3PAR management console using a GUI.

To review the current configuration, use the following commands:

- showvv
- showvvset
- showvlun
- showcpg
- showhostset
- showhost

Perform steps 1 through 4 on the 3PAR.

Procedure

1. Review the hostset using the showhostset command, as the hostset is already created. Start with creating additional common provisioning groups. Review the existing CPGs with the showcpg -sdg command. If there is an existing CPG that contains RAID 1 volumes already created use that CPG; otherwise provide a CPG name to host the RAID 1 thin volumes.

For example, create a common provisioning group (CPG) using RAID1 and magazine layout as follows:

```
createcpg -ssz 2 -ss 512 -t r1 -ha port -p -mg 0-15 -cg 0-3 -devtype FC
HANA_CPG_r1
```

2. Create the data and log volumes according to the sizes shown in Table 3. You may use the showvv command to review the existing virtual volumes. Examples of these commands follow; repeat them for each volume required for the instance.

```
createvv -i 101 -tpvv -wait 0 HANA_CPG_r1 R1_LOG1.1 520G
```

```
createvv -i 111 -tpvv -wait 0 HANA_CPG_r1 R1_DATA1.1 <DATASIZE>G
```

3. Export the virtual volumes to the HANA nodes by creating vluns using the following commands. You should use the hostset defined in step 2. To review the hostsets, use the showhostset command.
 - a. Display the current vluns using the following command. Locate the highest vlun and start numbering (nextvlun) after the highest vlun displayed.

```
showvlun
```

- b. Create a vlun for each data volume just created.

Increase nextvlun after each createvlun.

```
createvlun -f R1_DATA2.1 nextvlun set:HostSet_R<RID>
```

- c. Create a vlun for each log volume just created.

Increase nextvlun after each createvlun.

```
createvlun -f R1_LOG2.1 nextvlun set:HostSet_R<RID>
```

4. On the 3PAR, run the showvv command with the following parameters to extract the WWN for each volume. The VV_WWN is used in setting up the HANA nodes to use these volumes.

The following example also shows the size of the volumes.

```
showvv -showcols Name,VV_WWN,VSize_MB
```

The following is an example of the output:

```
Name VV_WWN VSize_ME
R1_DATA2.1 60002AC00000000000000000150000C43E 1572864
R1_DATA2.2 60002AC00000000000000000160000C43E 1572864
R1_DATA2.3 60002AC00000000000000000170000C43E 1572864
R1_DATA2.4 60002AC00000000000000000180000C43E 1572864
R1_LOG2.1 60002AC00000000000000000B0000C43E 532480
R1_LOG2.2 60002AC00000000000000000C0000C43E 532480
R1_LOG2.3 60002AC00000000000000000D0000C43E 532480
R1_LOG2.4 60002AC00000000000000000E0000C43E 532480
```

5. Perform steps 6 through 8 on each server.

The volumes must appear and be initialized on the HANA node. The data and log entries must be added to a new global.ini file for each new instance created.

6. Create the LVM volumes (only created on the first node in the appliance since all storage is seen by all nodes).

For example:

```
/sbin/pvcreate --metadatatype 2 -ff /dev/mapper/r2_data1.1
```

```
/sbin/vgcreate vg_data2_11 --physicalextentsize 256M /dev/mapper/r1_data2.1
```

```
/sbin/lvcreate -l 100%VG --name lv011 vg_data2_11
```

```
/sbin/pvcreate --metadatatype 2 -ff /dev/mapper/r2_log1.1
```

```
/sbin/vgcreate vg_log2_11 --physicalextentsize 32M /dev/mapper/r2_log1.1
```

```
/sbin/lvcreate -l 100%VG --name lv011 vg_log2_11
```

7. On the servers, rescan the scsi bus and restart the multi-path to assure the virtual volumes are visible to RHEL / SLES. Once complete, check the scsi volumes and compare with the export using the 3PAR management console as described in the Storage section.

For example:

```
rescan-scsi-bus.sh
service multipathd restart
ls -l /dev/disk/by-id
```

8. Create the file systems and the mount points. Use the following example commands to perform these tasks. The scsi-3... volume names are dependent upon the lvm logical volumes (vluns) created in step 6.

```
# make.xfs for each data and log lun, for example
/sbin/mkfs.xfs -f -i size=2048 -b size=4k -s size=4k -l su=256k /dev/mapper/
vg_data2_11-lv011
/sbin/mkfs.xfs -f -i size=2048 -b size=4k -s size=4k /dev/mapper/vg_log2_11-
lv011
# make the mount directories for each data and log filesystem associated with
an instance.
```

For example:

```
mkdir /hana/data/INx mkdir /hana/log/INx
chmod 777 /hana/data/INx /hana/log/INx
```

9. Use the example output from step 6 and prepare the storage section of the global.ini file. The storage section should look similar to the following example:

```
[storage] ha_provider=hdb_ha.fcClientLVM partition_*_* prType = 5
partition_1_data lvmname = vg_data2_11-lv011 partition_1_data mountoptions =
-t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc partition_1_log lvmname =
vg_log2_11-lv011 partition_1_log mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc partition_2_data lvmname =
vg_data2_12-lv011 partition_2_data mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc partition_2_log lvmname =
vg_log2_12-lv011 partition_2_log mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc partition_3_data lvmname =
vg_data2_13-lv011 partition_3_data mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc partition_3_log lvmname =
vg_log2_13-lv011 partition_3_log mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc partition_4_data lvmname =
vg_data2_14-lv011 partition_4_data mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc partition_4_log lvmname =
vg_log2_14-lv011 partition_4_log mountoptions = -t xfs -o
inode64,nobarrier,logbufs=8,logbsize=256k,swalloc
```

10. Install another HANA instance using the newly exported virtual volumes.

Secure Encryption Drives on ConvergedSystem for SAP HANA appliances

ConvergedSystem 500 for SAP HANA supports secure encryption drives (SED) through the 3PAR storage array.

! **IMPORTANT:**

Backups, when running with an External Key manager, are for configuration information ONLY! While this file is still important, as it is necessary to recover from a disaster, the keys are stored ONLY on the EKM, and must be backed up independently. Also, when configuring the EKM and before encryption is enabled, a backup filename is not required.

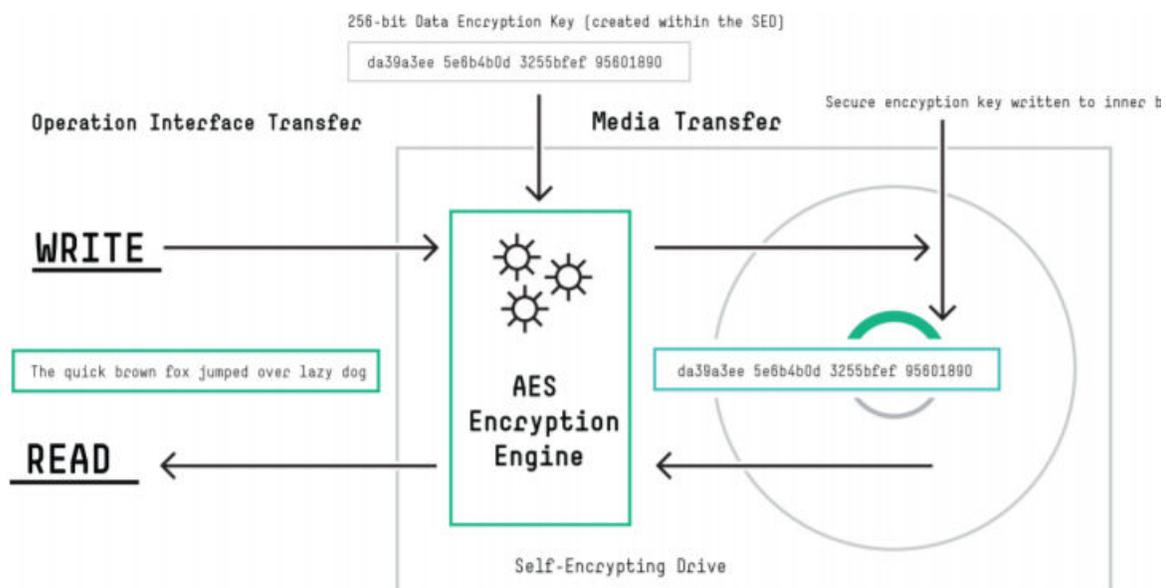
NOTE:

You must have a valid HPE 3PAR DAR license to enable the encryption key.

NOTE:

When a disk is admitted into or dismissed from the system, a unique internal encryption key is created so that all data is destroyed. Neither the encryption key nor data are ever exposed outside the array.

The following illustration shows how encryption works in an SED. As shown, the encryption engine is located in the drive and therefore requires no additional hardware or software to encrypt the data stored on the drive.



When using SEDs, observe the following:

Encrypting drives:

- SED drives encrypt data even if encryption is not enabled on the 3PAR 8400 array.
- Enabling the encryption key will protect the drives by locking the disks to the 3PAR 8400 array.
- A single key is used for drives in the entire 3PAR 8400 array.
- Encryption must be enabled on site. It is not enabled when installed at the factory.
- It is a best practice to enable encryption during installation.
- You can enable encryption on the array from a nonencrypted state and not lose any data. However, once encrypted the array must always be encrypted. Removing encryption will erase the secure encryption band on each of the physical drives and render all data unintelligible.
- Mixing encrypted and nonencrypted drives is not supported.
- 3PAR virtualization is separate from disk encryption.

Removing drives:

- Extraction of a SED disk drive within the array without the (License Key Manager) LKM or Enterprise Key Manager (EKM) enabled could allow an unauthorized user access to the data stored on the drive. When the drive is unlocked, all I/O behaves exactly as a non-SED disk drive.
- When an SED drive with encryption enabled is removed from the array it is locked. The data cannot be accessed without the key.

For more information, see:

- "Data-at-rest encryption overview " in the [HPE StoreServ Management Console 3.0 User Guide](#).
- [HPE 3PAR StoreServ Data-At-Rest Encryption](#) Technical White Paper.

Enabling encryption on self encrypting drives (SED) drives using the 3PAR CLI

Prerequisites

You have a valid encryption license.

This procedure uses the 3PAR OS Command Line Interface (CLI). For additional encryption commands, including, `rekey`, `status`, `restore`, and `backup`, see the [HPE 3PAR Command Line Interface User Guide](#).

Procedure

1. Access the 3PAR StoreServ CLI and enter:

```
setlicense
```

2. Enter the license key.

3. To validate that the licence has been applied, enter:

```
showlicense
```

4. To enable encryption, enter:

```
controlencryption enable
```

5. To verify encryption has been enabled, enter:

```
controlencryption status
```

```
cli% controlencryption status -d
Licensed Enabled BackupSaved State SeqNum Keystore FIPS non-SEDs
FailedDisks
yes      yes      no      normal      0 ---      ---
0        0
```

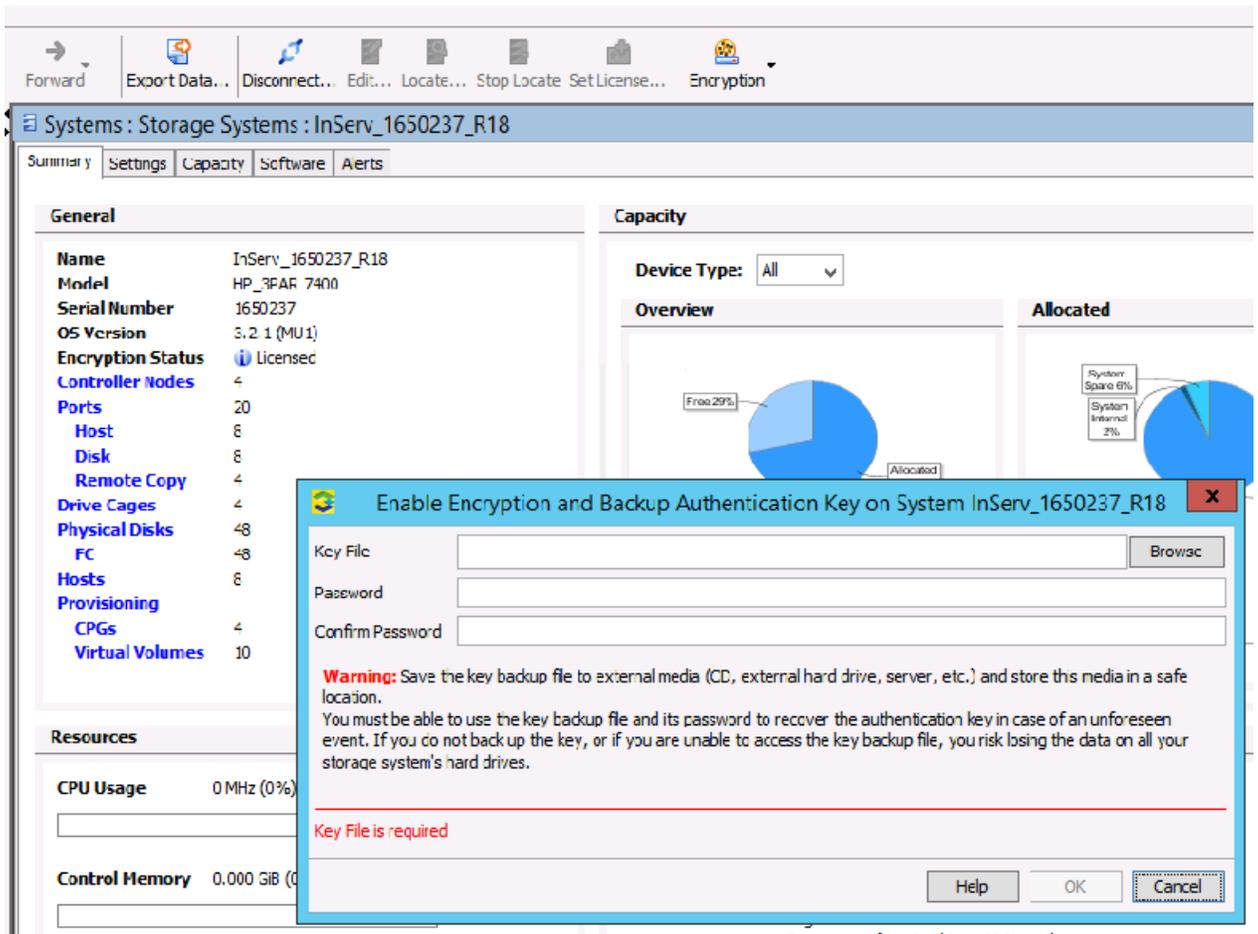
Enabling encryption on SED drives using the 3PAR management GUI

Prerequisites

You have a valid encryption license.

Procedure

1. Select **Systems** in the Manager Pane, click **Encryption** in the Main Toolbar, then select **Enable**.
The Enable Encryption and Back Up Authentication Key dialog appears.



2. Enter the file name and location for the key file, or click **Browse** to navigate to a file, then click **Save** to return to the dialog.

The Key File text box is populated with the selected path to the file.

3. Enter a password (the length must be between 6 and 31 characters).
4. Confirm the password.
5. Click **OK** to begin the backup procedure and enable encryption.

If the encryption is successful, a confirmation dialog appears, and the status bar updates with a similar message. This process is executed as a task and run in the background. The confirmation dialog shows that the task has launched successfully.

6. To confirm that encryption is enabled, visit the Tasks & Schedules Manager for detailed information on the task.

System	ID	Type	Name	Status	User	Progress	Start Time	Finish Time	Duration	Priority
InServ1650237_R18	7210	Encryption Change	Encryption enable	Completed	operator	100%	May 20, 2015 10:58:47 PDT	May 20, 2015 10:59:52 PDT	5 seconds	--

Troubleshooting

Common problem resolution

The following list provides some recommendations for starting the troubleshooting process.

Before calling for service, perform the following basic troubleshooting procedures. If the following procedures do not solve your problem, follow the instructions outlined in "[Proactive Care electronic case logging](#)."

Power and connectivity issues

Isolate the problem to the component level by checking power and connectivity.

- Check any interlock or interconnect LEDs that may indicate a component is not connected properly.
- Be sure all power cords are securely connected.
- Be sure all network cables are properly aligned and securely connected.
- If a device has latches, be sure they are completely closed and locked.

Server and storage issues

- For server and storage, use the Service Pack for ProLiant tools to check for basic operational statistics.
- For networking, isolate to the network and port using ping and troubleshoot from this point.
- Ensure that all the files systems, including HANA data and log, have enough free space.

HANA Studio tools (database issues)

- HANA Studio has monitoring capability. Use this tool to discover issues with the HANA software.

Service notifications

To search for service notifications, see the [Hewlett Packard Enterprise Support Center](#). Enter the product name or number, and then click **Go**. Select **Advisories, Bulletins and Notices**. The complete list of documents is displayed.

Release Notes

When comparing this version with the previous generation of the scale-out appliance, the following features are highlighted:

- Replacing the blades and enclosures with DL580 Gen8 servers
- Replace the DL580 Gen8 with DL580 Gen9 servers
- Reduce the number of 3PAR Service Processor to one per rack
- Remove the 3PAR Service Processor and replace with Virtual Service Processors
- Reduce the number of 3PAR 7400s to one per rack
- Added a management LAN switch to rack one
- Removed the added management LAN switch where all networking is handled though the FlexFabric 5930 switches
- Added replication and backup networks
- The SAN switch model changed to SN6000B
- Replace 2x DL380 Gen8 with HPE Serviceguard for NFS with 3PAR File Persona
- Added additional HANA servers per rack.

Resources

References

HPE SAP HANA

- [HPE ConvergedSystem 500 for SAP HANA Scale-out Configurations solution architecture white paper](#)
- [SAP HANA Platform Documentation](#)
- [SAP HANA Administration Guide](#)
- [SAP HANA Technical Operations Manual](#)
- [SAP HANA Update and Configuration Guide](#)
- [SAP HANA Storage Requirements](#)
- [SAP HANA Studio Installation Guide](#)
- [HPE ConvergedSystem 500 for SAP HANA Scale-up Configurations solution architecture white paper](#)
- For more detailed information, see SAP Note 1514967 – SAP HANA 1.0: Central Note and related notes.

HPE ProLiant Servers, storage and miscellaneous

- [ProLiant DL580 Gen9](#)
- [HPE Insight Remote Support](#)

HPE 3PAR StoreServ

- [HPE 3PAR StoreServ](#)
- [HPE 3PAR Management Console Users Guide](#)
- [HPE 3PAR Command Line Reference](#)
- [HPE 3PAR StoreServ 8000 Storage Site Planning Manual](#)
- [HPE 3PAR StoreServ 8000 Storage Installation Guide](#)

HPE 5930 Switch

- [HPE 5930 Switch Series manuals](#)
- [Fundamentals Command Reference](#)
- [HPE SN6000B 16 GB 48-port/24-port Active Fibre Channel SwitchHP 8/24 SAN Switch](#)
- [Brocade Fabric OS v7.3 Release Notes](#)
- [HPE SN6000 Fibre Channel Switch Command Line Interface Guide](#)

SUSE Linux Enterprise Server for SAP

- [SLES11](#)

RedHat Linux Enterprise Server for SAP

- [RedHat](#)

Safety considerations

Important safety information

Familiarize yourself with the safety information in the following sections before troubleshooting the server.



Important safety information

Before servicing this product, read the Important Safety Information document provided with the server.

Symbols on equipment

The following symbols might be found on the equipment to indicate the presence of potentially hazardous conditions.



This symbol indicates the presence of hazardous energy circuits or electric shock hazards. Refer all servicing to qualified personnel.

WARNING: To reduce the risk of injury from electric shock hazards, do not open this enclosure. Refer all maintenance, upgrades, and servicing to qualified personnel.



This symbol indicates the presence of electric shock hazards. The area contains no user or field serviceable parts. Do not open for any reason.

WARNING: To reduce the risk of injury from electric shock hazards, do not open this enclosure.



This symbol on an RJ-45 receptacle indicates a network interface connection.

WARNING: To reduce the risk of electric shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.



This symbol indicates the presence of a hot surface or hot component. If this surface is contacted, the potential for injury exists.

WARNING: To reduce the risk of injury from a hot component, allow the surface to cool before touching.



This symbol indicates that the component exceeds the recommended weight for one individual to handle safely.

WARNING: To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manual material handling.



These symbols, on power supplies or systems, indicate that the equipment is supplied by multiple sources of power.

WARNING: To reduce the risk of injury from electric shock, remove all power cords to disconnect power from the system completely.

Electrostatic discharge

Preventing electrostatic discharge

To prevent damaging the system, be aware of the precautions you must follow when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor may damage system boards or other static-sensitive devices. This type of damage may reduce the life expectancy of the device.

To prevent electrostatic damage:

- Avoid hand contact by transporting and storing products in static-safe containers.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-free workstations.
- Place parts on a grounded surface before removing them from their containers.
- Avoid touching pins, leads, or circuitry.
- Always be properly grounded when touching a static-sensitive component or assembly.

Grounding methods to prevent electrostatic discharge

Several methods are used for grounding. Use one or more of the following methods when handling or installing electrostatic-sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm ± 10 percent resistance in the ground cords. To provide proper ground, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.
- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an authorized reseller install the part.

For more information on static electricity or assistance with product installation, contact an authorized reseller.

Support and other resources

Before you contact Hewlett Packard Enterprise

Be sure to have the following information available before you call Hewlett Packard Enterprise:

- The ConvergedSystem solution support product number and serial number (found on the warranty label)
- Active Health System log (HPE ProLiant Gen8 or later products)

Hewlett Packard Enterprise recommends providing an Active Health System log of the 7 days before the failure was detected. For more information, see the *HPE iLO 4 User Guide* or *HPE Intelligent Provisioning User Guide* on the [Hewlett Packard Enterprise website](#).

- Technical support registration number (if applicable)
- Product serial number
- Product model name and number
- Product identification number
- Applicable error messages
- Add-on boards or hardware
- Third-party hardware or software
- Operating system type and revision level

Proactive Care

ConvergedSystems are sold with Proactive Care as a minimum support level. You will have been provided a Service Agreement ID (SAID). You can engage support using the following methods.

- Proactive Care electronic case logging (recommended)
- Proactive Care phone support

Requesting support for HPE ConvergedSystem and HPE Hyper Converged products

Prerequisites

- An active HPE Passport account.
- An active support contract and the contract Service Agreement ID (SAID).
- A Support Case Manager PIN is saved in **Edit SCM Settings**.

NOTE:

Customers in Germany, Austria, Switzerland, France, Japan, South Korea, and China should follow country-specific procedures for local language support.

Procedure

1. On the [HPESC website](#):

1. Click **My HPE Support sign-in** and log in with your HPE Passport ID.

My HPE Support Center



My HPE Support sign-in

The Get drivers, manuals, parts & solutions screen is displayed.

Get drivers, manuals, parts & solutions

Select your HPE product

Recent products:

[HPE Hyper Converged 380](#)

Enter a product name or number

(e.g. ProLiant DL360p)

OR

My HPE Support Center

[Sign-out](#)

[Edit your profile](#)

[View my recent cases](#)

[View my contracts & warranties](#)

[Manage my contracts & warranties](#)

More support options

 [Community forums](#)

 [Warranty Check](#)

 [Chat with HPE](#)

 [Submit or manage support cases](#)

2. Under **More support options**, click **Submit or manage support cases**.

The Support Case Manager screen is displayed.

3. Under **Submit a case**, enter your Service Agreement Identifier (SAID) then click **Submit case**.

The Support Case Manager is displayed.

4. In the **Action** column, click the **Submit a case** button in the solution row.

The Case details page is displayed.

5. Enter detailed information about your request then click **Submit**.

The Contact & equipment location Information screen is displayed.

2.

6. Verify, change, or enter the information in the Contact & equipment location Information screen then click **Submit**.

NOTE: The PIN, in the Support Case Manager PIN field, is the default PIN set in your SCM Settings. If necessary, edit the PIN field to contain the correct Support Case Manager PIN.

For a list of PINs, see [Support Case Manager PINs](#).

Contact & equipment location Information

Enter information below or select prefills from the dropdown lists. You may modify the contact information. Any changes made here are effective for this case only. For permanent changes, please contact your Hewlett Packard Enterprise representative listed on your Hewlett Packard Enterprise customer support service documentation. Enter any additional contact information into the "problem description" under case details. [?](#)

Contact	Contact (2)	Equipment	Ship to
Contact	Select a contact <input type="button" value="v"/>	Equipment contact name	<input type="text"/>
First name	<input type="text"/>	Equipment contact phone	<input type="text"/>
Last name	<input type="text"/>	Company name	<input type="text"/>
Phone number	<input type="text"/> Ext <input type="text"/>	Address line 1	<input type="text"/>
Alternate Phone	<input type="text"/>	Address line 2	<input type="text"/>
Email address	<input type="text"/>	Mailstop	<input type="text"/>
Support Case Manager PIN	<input type="text"/>	City/Town	<input type="text"/>
		State/Province	<input type="text"/>
		Zip/Postal code	<input type="text"/>
		Country/Region	<input type="text"/> <input type="button" value="v"/>

Your request is submitted. A member of the support team will contact you within two hours of receiving your case.

Support Case Manager PINs

Solution	PIN
ConvergedSystem product with SAP HANA	SAPHANA
All other ConvergedSystem and Hyper Converged products	HP_CLOUD_SOL

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

For more information and device support details, go to the [Insight Remote Support website](#).

Warranty information

To view the warranty for your product, see the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* document, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional warranty information

HPE ProLiant and x86 Servers and Options	<u>www.hpe.com/support/ProLiantServers-Warranties</u>
HPE Enterprise Servers	<u>www.hpe.com/support/EnterpriseServers-Warranties</u>
HPE Storage Products	<u>www.hpe.com/support/Storage-Warranties</u>
HPE Networking Products	<u>www.hpe.com/support/Networking-Warranties</u>

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.