



HPE Solution Architecture for SAP HANA on HPE Superdome Flex with D3710 direct attached storage

Built with Intel Xeon Scalable 2nd Generation processors (Cascade Lake)

CONTENTS

Executive summary.....	3
Introduction.....	4
HPE persistent memory module.....	4
Solution overview.....	4
Hardware.....	5
HPE Superdome Flex Server hardware components.....	18
Software.....	19
Introduction to Intel Xeon Scalable Processors.....	20
HPE persistent memory module.....	21
Memory allocation goal – Namespace.....	23
Region.....	23
Solution configuration.....	25
HPE Superdome Flex Scale-up with D3710 Direct Attached Storage Server configuration for SAP HANA.....	25
Network interface card.....	25
Storage.....	26
Persistent memory configuration.....	27
SAP HANA Cockpit view of persistent memory.....	30
Capacity and sizing.....	38
Network requirements.....	39
Configured networks.....	40
Tuning the SAP HANA operating system.....	40
Workload description.....	40
Summary.....	41
Resources and additional links.....	42



EXECUTIVE SUMMARY

Hewlett Packard Enterprise is [#1 in SAP HANA server deployments](#) with approximately 40% of the market share. Ranked #1 in scale-up and scale-out capacity, Hewlett Packard Enterprise commands a wide margin, and leadership, in total SAP® systems deployed, SAP HANA® appliances, tailored data center integration (TDI), SAP BW/4HANA, and SAP S/4HANA.

With nearly 25,000 customers running SAP applications on HPE hardware, thousands of SAP HANA customers worldwide, and unmatched migration experience, Hewlett Packard Enterprise brings a unique understanding of SAP and SAP HANA environments that benefit customers of all sizes and with all types of workload.

Hewlett Packard Enterprise offers a differentiated infrastructure portfolio for SAP HANA designed to address a broad range of customer needs beginning with production systems for small and medium businesses to highly scalable solutions for mission-critical environments, including the application tier.

HPE Pointnext Services provides truly comprehensive SAP solutions for the Hybrid Cloud leveraging experience from over 1500 SAP HANA consulting projects and supporting over 25,000 SAP customers.

SAP HANA is a modern technology that offers an in-memory database, enables a radically different application architecture, and provides a new philosophy with regards to data model simplicity. SAP HANA offers a new way of solving current and future challenges with enterprise applications, such as real time data analysis of large amounts of data. More customers are transitioning to SAP HANA and are accelerating its use as a mission-critical platform.

SAP HANA requires memory for processing data in real time, as well as persistent storage to hold the data on media. In the context of an SAP HANA integration, either shared storage or local disks can be leveraged. Shared or external storage based on Fibre Channel or Network File System (NFS) is required to get certified as SAP HANA Enterprise Storage. NFS is not required for certifying local disk connectivity with SAP, instead, a vendor verification for implementing the SAP HANA persistence layer on local disks has to be completed.

This vendor verification is required as a proof that the SAP HANA KPIs are met, as outlined in the [SAP HANA Frequently Asked Questions](#) document (credentials required).

An SAP HANA appliance deployment utilizes specific, preconfigured hardware (compute, network, and storage) and software (OS, SAP HANA) certified by SAP as appliance and implemented by Hewlett Packard Enterprise as a plug-and-play solution.

Hewlett Packard Enterprise and SAP work together to improve the enterprise capabilities of IT organizations, to meet their business needs, safe guarding their current investments in data center infrastructure, tools, and operational processes.

This document describes implementations of HPE Superdome Flex Scale-up with D3710 Direct Attached Storage configurations based on Intel® Xeon® Scalable 2nd Generation processor (Cascade Lake) Server. The Superdome Flex Scale-up D3710 Direct Attached Storage supports a wide variety of memory and disk configurations that SAP supports.

Target audience: This document has been created for Enterprise Solution Architects, Deployment or Implementation Engineers, and others wishing to learn more about the implementation of HPE Solution for SAP HANA using Superdome Flex Scale-up with D3710 Direct Attached Storage.

Document purpose: This document describes the architectural and technical overview required while using HPE Superdome Flex Scale-up with Direct Attached Storage (D3710) to deploy in-memory SAP HANA Database based on Intel® Xeon® Scalable 2nd Generation processors.



INTRODUCTION

SAP HANA is a column-oriented, relational database management system that combines online analytical processing (OLAP) and online transaction processing (OLTP) in one database system. Its primary function as a database server is to store and retrieve data as requested by applications. In-memory processing fundamentally utilizes the system's RAM (main memory) to store data, instead of using disk storage. This allows the application to provide instantaneous results from customer transactions and data analysis.

A columnar database management system (DBMS) boosts system performance and speed when working with aggregates. This is achieved by reading only the data that is required to satisfactorily fulfill a query, through effective compression of similar columnar data. SAP HANA effectively deploys the combination of in-memory computing and columnar database management system, thus becoming one of the best-suited tools for data warehousing applications.

It combines the features of in-memory computing and columnar DBMS, a duo that tremendously enhances the speed of operation when performing real-time analysis or running real-time applications.

This document provides SAP HANA information in terms of storage performance. Additionally, this document outlines the advantages of HPE Superdome Flex Scale-up with D3710 Direct Attached Storage, details best practices, and gives artifact information for deploying SAP HANA on these compute blocks.

Testing for this document showed that SAP HANA TDI KPIs are fully met with SAP HANA on HPE Superdome Flex Scale-up with D3710 Direct Attached Storage. Even the stricter KPIs for the SAP HANA appliance can be fulfilled for up to two SID instances. Testing has been carried out in a standard lab environment using the Hardware Configuration Check Tool (HWCCT) for SAP HANA single-node server. Testing results show that single-node configuration with four multi-tenants under one production SID can fulfill the TDI KPIs.

Having HPE Superdome Flex Scale-up with D3710 Direct Attached Storage configured with a more affordable storage configuration of SSDs allows to run more than one SAP HANA production system. Up to three SID have been tested and the results fulfill the SAP HANA KPIs. But as more SID are tested, the IO latency gets much higher and the bandwidth gets reduced even though the KPIs are still met.

HPE Superdome Flex Scale-up with D3710 Direct Attached Storage provides a flexible, simple, and efficient option for the SAP HANA workload. This allows you to consolidate and standardize infrastructure for mixed workloads, while still meeting high-performance requirements.

HPE persistent memory module

Persistent memory describes technologies that store data in dual in-line memory module (DIMM) format (memory that resides on the memory bus), that is accessed using memory instructions, and is non-volatile (data remains even without power). Hewlett Packard Enterprise is delivering this technology in a new HPE persistent memory module format, featuring Intel Optane™ DC persistent memory. Throughout this paper, the term "persistent memory" refers to the HPE persistent memory module.

Persistent memory possesses properties of both dynamic random access memory (DRAM) and storage, but it doesn't replace DRAM. It can be viewed as a tier between DRAM and flash storage.

What really makes Intel persistent memory technology stand out is that it is fast enough for the processor to access directly without stopping for the block I/O required for traditional storage.

Applications can access persistent memory-resident data structures in place, as they do with traditional memory, eliminating the need to page blocks of data back and forth between memory and storage. Persistent memory cannot replace DRAM memory completely and for SAP HANA, must be used in conjunction with DRAM to operate.

To leverage this lower-latency, higher-speed data access, we need a software architecture that allows applications to connect up with ranges of persistent memory.

SOLUTION OVERVIEW

HPE Superdome Flex Scale-up with D3710 Direct Attached Storage is built with the Intel Xeon Scalable family processors. These Hewlett Packard Enterprise servers deliver security, agility, and flexibility without compromise. Deploy this secure platform for diverse workloads in space-constrained environments.

Configurations and combinations mentioned in this document have been tested in a standard lab environment:



4S servers

- [HPE Superdome Flex Scale-up with 4S/1.5 TB and 6x 3.2 TB SSDs](#)
- [HPE Superdome Flex Scale-up with 4S/3 TB and 12x 3.2 TB SSDs](#)
- [HPE Superdome Flex Scale-up with 4S/4.5 TB and 12x 3.2 TB SSDs using Persistent Memory Modules](#)
- [HPE Superdome Flex Scale-up with 4S/6 TB and 12x 3.2 TB SSDs](#)
- [HPE Superdome Flex Scale-up with 4S/6 TB and 12x 3.2 TB SSDs using Persistent Memory Modules](#)
- [HPE Superdome Flex Scale-up with 4S/7.6 TB and 16x 3.2 TB SSDs using Persistent Memory Modules](#)
- [HPE Superdome Flex Scale-up with 4S/9 TB and 16x 3.2 TB SSDs using Persistent Memory Modules](#)

8S servers

- [HPE Superdome Flex Scale-up with 8S/3 TB and 12x 3.2 TB SSDs](#)
- [HPE Superdome Flex Scale-up with 8S/6 TB and 12x 3.2 TB SSDs using Persistent Memory Modules](#)
- [HPE Superdome Flex Scale-up with 8S/9 TB and 16x 3.2 TB SSDs using Persistent Memory Modules](#)
- [HPE Superdome Flex Scale-up with 8S/12 TB and 21x 3.2 TB SSDs](#)

12S servers

- [HPE Superdome Flex Scale-up with 12S/4.5 TB and 12x 3.2 TB SSDs](#)
- [HPE Superdome Flex Scale-up with 12S/9 TB and 16x 3.2 TB SSDs](#)

The configurations above have been chosen based on the following criteria used during test:

- Number of processors: 4, 8, and 12
- Types of processors used for Appliance or TDI configuration
- Number of SSDs used from minimum to maximum
- Different size (64 GB, 128 GB DDR4, 128 GB PMEM and 256 GB PMEM) DIMMS used during these testing

Supported memory combinations are mentioned in the [Capacity and sizing](#) section of this document.

Hardware

Table 1 shows the hardware components of HPE Superdome Flex Scale-up with D3710 Direct Attached Storage Server used during testing of this document. These hardware components are common in different configurations for SAP HANA solutions.

TABLE 1. Hardware components used for this document

Components	Quantity	Purpose
HPE Superdome Flex Base Chassis	1	Server Chassis
MegaRAID9361-4i LSI00415 (SGL) LP MD2 4 1 x SFF-8643	1	Controller
HPE 3154-8e 8-port External RAID Controller	1	Controller
HPE Superdome Flex 8280, 8276 (with M/L) Proc Kit	4	Processor Kit
HPE Ethernet 10/25Gb 2P 640SFP28 Adapter	4	NIC
HPE Ethernet 1Gb 4-port 331T Adapter	1	NIC
HPE Ethernet 10Gb 2P 562T Adapter/ Copper	4	NIC



NOTE

During the testing of the configuration mentioned in this paper, 8280, 8280M/L, 8276, 8276M/L processor models have been used. Depending on the memory requirement, Platinum, Gold, or Silver processor can be used. In case of non-availability of “M” processors, “L” processors can be used. Also, an extra 3.2TB SSD can be added for a spare.

Figure 1 shows the front view with Solid State Drive, and the rear view shows PCI Cards and other ports for external connections and power supply of HPE Superdome Flex Scale-up with D3710 Server.




FIGURE 1. HPE Superdome Flex 4s Base Chassis Server with 16 SSD External D3710 Direct Attached Storage

B. Front View details of HPE Superdome Flex with external D3710 Direct Attached Storage

C. Rear View details of HPE Superdome Flex with external D3710 Direct Attached Storage



NOTE

SAP has determined that all persistent memory configurations are TDI. For a system to have the CPU frequency and cores for an SAP Appliance, the server processor must be at top bin or second bin: 8280, 8280M/L, 8276, 8276M/L. Any server not using these processors, but using Silver, Gold, or Platinum processors and at least 8 cores per processor is only supported under a TDI deployment.

Hardware configuration of Superdome Flex with 4s/1.5 TB DRAM memory configuration

Figure 2 shows 4s/1.5 TB memory configuration with 2x 1.2 TB internal OS SAS HDDs and 6x 3.2 TB SSDs. This configuration is designed for pre-defined implementation or minimum required storage space for TDI implementation.

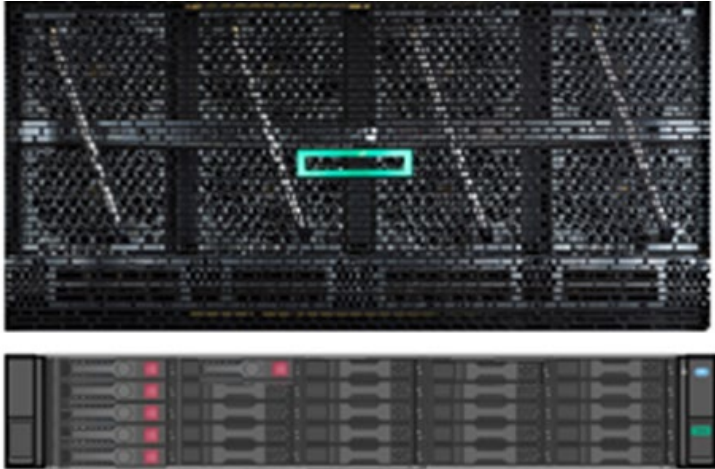


FIGURE 2. HPE Superdome Flex 4s Base Chassis Server with external D3710 with 6 SSD

Table 2 describes the hardware components specific to this configuration.

TABLE 2. Hardware components used for 4s/1.5 TB configuration in addition to Table 1.

Components	Quantity	Component Description
Processor	4	Intel Xeon Platinum 8276 CPU @ 2.60GHz
Memory – DRAM	24	HPE 64 GB 4Rx4 PC4-2933Y-L Smart Kit
SSD - Storage	6	HPE 3.2 TB SAS MU SFF SC DS SSD



Hardware configuration of HPE Superdome Flex with 4s/3 TB DRAM Memory Configuration

Figure 3 shows 4s/3 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 12x 3.2 TB external SAS SSDs. This configuration is designed for pre-defined implementation or minimum required storage space for TDI implementation.




FIGURE 3. HPE Superdome Flex 4s Base Chassis Server with external D3710 with 12 SSD

Table 3 describes the hardware components specific to this configuration.

TABLE 3. Hardware components used for 4s/3 TB configuration in addition to Table 1

Components	Quantity	Component Description
Processor	4	Intel Xeon-Platinum 8276 CPU @ 2.60GHz
Memory – DRAM	48	HPE 64 GB 4Rx4 PC4-2933Y-L Smart Kit
SSD – Storage	12	HPE 3.2 TB SAS MU SFF SC DS SSD

Hardware configuration of Superdome Flex with 4s/4.5 TB Memory Configuration with Intel Persistent Memory

Figure 4 shows 4s/4.5 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 12x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with a minimum storage requirement for TDI as defined by SAP.




FIGURE 4. HPE Superdome Flex 4s Base Chassis Server with external D3710 with 12 SSD



Table 4 describes the hardware components specific to this configuration.

TABLE 4. Hardware components used for 4s/4.5 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	4	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
Memory – Persistent	24	HPE 128GB 2666 MT/s Persistent Mem Mod
Memory – DRAM	24	HPE 64GB 4Rx4 PC4-2933Y-L Smart Kit
SSD – Storage	12	HPE 3.2TB SAS MU SFF SC DS SSD

Hardware configuration of Superdome Flex with 4s/6 TB DRAM Memory Configuration

Figure 5 shows 4s/6 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 12x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with minimum storage requirement as defined by SAP.



FIGURE 5. HPE Superdome Flex 4s Base Chassis Server with external D3710 with 12 SSD

Table 5 describes the hardware components specific to this configuration.

TABLE 5. Hardware components used for 4s/6 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	4	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
Memory – DRAM	48	HPE 128 GB 8Rx4 PC4-2933Y-L 3DS Smart Kit
SSD – Storage	12	HPE 3.2 TB SAS MU SFF SC DS SSD



Hardware configuration of Superdome Flex with 4s/6 TB Memory Configuration with Intel Persistent Memory

Figure 6 shows 4s/6 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 12x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with minimum storage requirement as defined by SAP.



FIGURE 6. HPE Superdome Flex 4s Base Chassis Server with external D3710 with 12 SSD

Table 6 describes the hardware components specific to this configuration.

TABLE 6. Hardware components used for 4s/6 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	4	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
Memory – Persistent	24	HPE 128 GB 2666 MT/s Persistent Mem Mod
Memory – DRAM	24	HPE 128 GB 8Rx4 PC4-2933Y-L 3DS Smart Kit
SSD – Storage	12	HPE 3.2 TB SAS MU SFF SC DS SSD

Hardware configuration of Superdome Flex with 4s/7.6 TB Memory Configuration with Intel Persistent Memory

Figure 7 shows 4s/7.6 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 16x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with minimum storage requirement as defined by SAP.

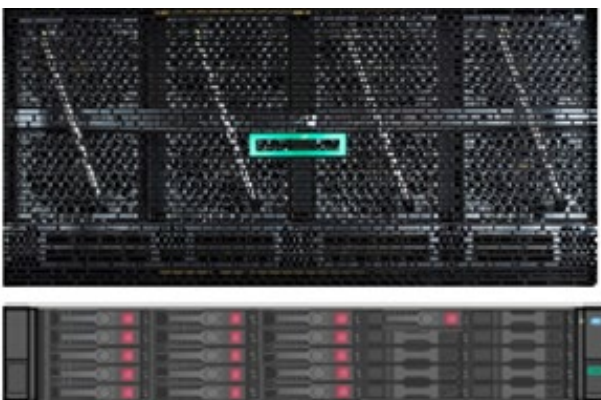


FIGURE 7. HPE Superdome Flex 4S Base Chassis Server with external D3710 with 16 SSD



Table 7 describes the hardware components specific to this configuration.

TABLE 7. Hardware components used for 4s/7.6 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	4	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
Memory – Persistent	24	HPE 256 GB 2666 MT/s Persistent Mem Mod
Memory – DRAM	24	HPE 64 GB 4Rx4 PC4-2933Y-L Smart Kit
SSD – Storage	16	HPE 3.2 TB SAS MU SFF SC DS SSD

Hardware configuration of Superdome Flex with 4S/9TB Memory Configuration with Intel Persistent Memory

Figure 8 shows 4s/9 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 16x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with minimum storage requirement as defined by SAP.



FIGURE 8. HPE Superdome Flex 4S Base Chassis Server with external D3710 with 16 SSD

Table 8 describes the hardware components specific to this configuration.

TABLE 8. Hardware components used for 4s/9 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	4	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
Memory – Persistent	24	HPE 256 GB 2666 MT/s Persistent Mem Mod
Memory – DRAM	24	HPE 128 GB 8Rx4 PC4-2933Y-L 3DS Smart Kit
SSD – Storage	16	HPE 3.2 TB SAS MU SFF SC DS SSD



Hardware configuration of Superdome Flex with 8s/3 TB DRAM Memory Configuration

Figure 9 shows 8s/3 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 12x 3.2 TB external SAS SSDs. This configuration is designed for Pre-Defined implementation or minimum required storage space for TDI implementation.



FIGURE 9. HPE Superdome Flex 8s Base Chassis Server with external D3710 with 12 SSD

Table 9 describes the hardware components specific to this configuration.

TABLE 9. Hardware components used for 8s/3 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	8	Intel Xeon Platinum 8280 CPU @ 2.60GHz
Memory – DRAM	48	HPE 64 GB 4Rx4 PC4-2933Y-L Smart Kit
SSD – Storage	12	HPE 3.2 TB SAS MU SFF SC DS SSD



Hardware configuration of Superdome Flex with 8s/6 TB Memory Configuration with Intel Persistent Memory

Figure 10 shows 8s/6 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 12x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with minimum storage requirement as defined by SAP.



FIGURE 10. HPE Superdome Flex 8s Base Chassis Server with external D3710 with 12 SSD

Table 10 describes the hardware components specific to this configuration.

TABLE 10. Hardware components used for 8s/6 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	8	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
Memory – Persistent	24	HPE 128 GB 2666 MT/s Persistent Mem Mod
Memory – DRAM	24	HPE 128 GB 8Rx4 PC4-2933Y-L 3DS Smart Kit
SSD – Storage	12	HPE 3.2 TB SAS MU SFF SC DS SSD



Hardware configuration of Superdome Flex with 8s/9 TB Memory Configuration with Intel Persistent Memory

Figure 11 shows 8s/9 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 16x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with minimum storage requirement as defined by SAP.



FIGURE 11. HPE Superdome Flex 8s Base Chassis Server with external D3710 with 16 SSD

Table 11 describes the hardware components specific to this configuration.

TABLE 11. Hardware components used for 8s/9 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	8	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
Memory – Persistent	24	HPE 256 GB 2666 MT/s Persistent Mem Mod
Memory – DRAM	24	HPE 128 GB 8Rx4 PC4-2933Y-L 3DS Smart Kit
SSD – Storage	16	HPE 3.2 TB SAS MU SFF SC DS SSD



Hardware configuration of Superdome Flex with 8s/12 TB DRAM Memory Configuration

Figure 12 shows 8s/12 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 21x 3.2 TB external SAS SSDs. This configuration is designed for pre-defined implementation or minimum required storage space for TDI implementation.




FIGURE 12. HPE Superdome Flex 8s Base Chassis Server with external D3710 with 21 SSD

Table 12 describes the hardware components specific to this configuration.

TABLE 12. Hardware components used for 8s/12 TB configuration in addition to Table 1

Components	Quantity	Purpose
Processor	8	Intel Xeon Platinum 8280L CPU @ 2.60GHz
Memory – DRAM	48 x 2 Chassis	HPE 128 GB 8Rx4 PC4-2933Y-L 3DS Smart Kit
SSD – Storage	21	HPE 3.2 TB SAS MU SFF SC DS SSD



Hardware configuration of Superdome Flex with 12s/4.5 TB DRAM Memory Configuration

Figure 13 shows 12s/4.5 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 12x 3.2 TB external SAS SSDs. This configuration is for TDI implementation with minimum storage requirement as defined by SAP. Table 13 describes the hardware components specific to this configuration.



FIGURE 13. HPE Superdome Flex 12s Base Chassis Server with external D3710 and RMC with 12 SSD

Table 13 shows the hardware components used for 12s/4.5 TB configuration.

TABLE 13. Hardware components used for 12s/4.5 TB configuration in addition to Table 1

Components	Quantity	Details
Processor	12	Intel Xeon Platinum 8276L CPU @ 2.60GHz
Memory – DRAM	72	HPE 64 GB 4Rx4 PC4-2933Y-L Smart Kit
SSD – Storage	12	HPE 3.2 TB SAS MU SFF SC DS SSD
RMC	1	1U External Rack Management Controller (RMC)



Hardware configuration of Superdome Flex with 12s/9 TB DRAM Memory Configuration

Figure 14 shows 12s/9 TB DRAM memory configuration with 2x 1.2 TB internal OS SAS HDDs and 16x 3.2 TB external SAS SSDs. This configuration is designed for Pre-Defined implementation or minimum required storage space for TDI implementation. Table 14 describes the hardware components specific to this configuration.




FIGURE 14. HPE Superdome Flex 12s Base Chassis Server with external D3710 and RMC with 16 SSD

Table 14 shows the hardware components used for 12s/9 TB configuration.

TABLE 14. Hardware components used for 12s/9 TB configuration in addition to Table 1

Components	Quantity	Details
Processor	12	Intel Xeon Platinum 8276L CPU @ 2.60GHz
Memory – DRAM	72	HPE 128 GB 8Rx4 PC4-2933Y-L 3DS Smart Kit
SSD – Storage	16	HPE 3.2 TB SAS MU SFF SC DS SSD
RMC	1	1U External Rack Management Controller (RMC)



HPE Superdome Flex Server hardware components

Servers

HPE Superdome Flex is a breakthrough system for powering critical applications and enabling real-time analytics using in-memory computing.

#1 in scale-up and scale-out capacity, Superdome Flex utilizes a unique, modular, “building block” architecture that will never outgrow, without overprovisioning.

HPE Superdome Flex server is included for the following reasons:

- Certified Server for SAP HANA implementation
- PCI and PCIe slots available for Storage and Network Configurations
- 16 half-height IO slots, or 8 full-height + 4 half-height IO slots, per four-socket chassis
- Superdome Flex Analysis Engine for better diagnostics and mission-critical reliability
- Redundant Power Supply

HPE 9361-4i MegaRAID Controller

The HPE 9361-4i MegaRAID Controller, supporting 12 Gbps SAS and PCIe 3.0, provides enterprise-class storage performance, reliability, security and efficiency needed to address your evolving data storage needs.

Features

- Increases performance with 1 GB Flash-Backed Write Cache and reduces latency in write-intensive applications that require heavy logging such as databases.
- Enhances performance in storage environments where data is read repeatedly from a large number of SSDs by using HPE SSD Smart Path to create a faster I/O path between the controller and SSD.
- Standard support for the most popular RAID levels and accessories, including RAID 5, RAID 6, and CacheVault™ Flash cache protection module, further strengthen the data protection capabilities of the MegaRAID SAS 9361-4i
- Connect up to 128 3 Gbps, 6Gb/s, or 12Gb/s SATA and SAS devices

HPE 3154-8e RAID Controller

The HPE 3154-8e RAID Controller, Ideal for enabling 12 Gbps storage capabilities in performance-hungry server and workstation platforms, without compromising proven Microsemi reliability. Provides high I/O transaction and high bandwidth processing solutions that reduce energy consumption and maintenance costs.

Features

- Accelerates storage with up to 4 GB of high speed DRAM cache with integrated cache protection.
- The drives can be used as physical disks (HBA mode) or hardware RAID 0, 1, 10, 5, 50, 6, 60. Mixed mode allows devices connected to the same adapter to be used in RAID and HBA modes simultaneously
- Adapter dynamic power management to save up to 30% power
- Supports up to 256 SAS or SATA devices using SAS expanders

Storage

A large variety of storage options are supported with TDI configurations, but the all flash local storage shown in this paper comes with 6.4 TB, 3.2 TB, 1.6 TB, and 800 GB SSDs. These SAS SSD drives are of different capacities, but it is best to use same size disks for flexibility of RAID configurations.



Network interface cards

SAP HANA services communicate via different network channels. It is a best practice to have a well-defined network topology to control and limit network bandwidth to SAP HANA, only to those communication channels required for specified services, and to apply appropriate additional security and performance measures as necessary.

Hewlett Packard Enterprise provides variety of network interface cards, from 1 Gbps through 40 Gbps network bandwidth to fulfill the SAP HANA database requirement.

Non-volatile memory

Disk drives perform as system storage and DRAM performs as system memory. Flash storage is cheaper, but there is lag in fetching data from storage to DRAM and to the processor. Even the fastest SSD technology can't compete with DRAM for high-speed bandwidth or overall read/write performance. Though DRAM is the fastest resilient entity nearest to CPU, it lacks persistency against power.

The HPE Persistent Memory Kit featuring Intel Optane DC Persistent Memory, together with traditional volatile DRAM DIMMs, can help businesses extract greater value from data-intensive applications such as in-memory analytics, database, content delivery networks, and high performance computing (HPC). This also provides fast, high-capacity, cost-effective memory, and storage to transform big data workloads and analytics by enabling data to be stored, moved, and processed quickly.

Software

Table 15 provides minimum software requirements for this solution.

TABLE 15. Minimum requirements for operating system and software

Software	Version
Operating System	SUSE Linux Enterprise Server 12 SP4, 15 SPO Red Hat Enterprise Linux Server 7 Update 6
SAP HANA	2.0 SPS 04 Rev 40
SAP HANA HW Config Check Tool	2.00.012.05

SAP HANA

SAP HANA is a modern technology by SAP and is offered as an in-memory database, enabling a radically different application architecture and a new philosophy with regards to the data model simplicity. SAP HANA offers a new way of solving current and future challenges with enterprise applications. More customers are transitioning to SAP HANA and are accelerating on utilizing it as a mission-critical platform. In response, SAP has planned to convert all existing customers running traditional SAP applications to HANA-based applications by 2025.

For this paper, SAP HANA 2.0 SPS 04 Rev40 was used.

Hardware Configuration Check Tool

To run SAP HANA as a productive instance, minimum performance criteria defined by SAP have to be met. The Hardware Configuration Check Tool (HWCCT) determines with tests and reports if the hardware you intend to use for single host or scale-out configurations meet these performance criteria.

The test tool has three modules:

- Landscape test
- File system test
- Network test

For proofing, the KPIs for SAP HANA on HPE Superdome Flex Scale-up with Direct Attached Storage D3710 Storage Module, the File system test has been executed. This test simulates and measures SAP HANA data and log IO throughput and latency in various block sizes on the IO subsystem. The test is performed by creating a test file, enlarging it to a specified size between 1 GB and 32 GB, overwrite the file, and finally read it. This is performed with different block sizes: 4 KB, 16 KB, 64 KB, 16 MB, and 64 MB for data volume, and 4 KB, 16 KB, and 1 MB for log volumes. Details for downloading and using the HWCCT are described in SAP Note 1943937—Hardware Configuration Check Tool—Central Note (SAP S-User required). For this white paper, HWCCT version 2.00.012.05 was used.



SAP HANA Hardware and Cloud Measurement Tools (HCMT)

SAP HANA hardware and cloud measurement tools are a new set of tools that can help the user to measure and analyze hardware or cloud systems performance parameters required before deploying SAP HANA. Going ahead, HCMT is the preferred tool to be used while applying for SAP HANA Certification for Hardware.

HCMT is now replacing HWCCT tool. The HCMT tool consist of the following components:

- SAP HANA hardware and cloud measurement tool
- SAP HANA hardware and cloud measurement analysis

For more details refer SAP Note: [2493172](#)

Introduction to Intel Xeon Scalable Processors

Different compute platforms offer unique benefits and capabilities. It is critical to understand how platforms based on Intel Xeon processors can address the particular requirement of the compute environment. The processor number is one of several factors, along with processor brand, specific system configurations, and system-level benchmarks, to be considered when choosing the right processor for computing needs.

Each Intel Xeon processor is associated to a number, which can help user to identify the feature available with the Intel Xeon processor.

An alphanumeric number for Intel Xeon scalable processors can explain the scheme based on feature, performance, and generation with brand.

- First digit - performance and feature level
- Second digit - processor generation
- Third and fourth digit - SKU numbers
- Last alpha suffix - memory capacity

Figure 15 shows Intel Xeon Platinum Series High Bin Cascade Lake Processor with 2 TB memory capacity for one processor.




FIGURE 15. Intel Xeon Platinum processor numbering explanation

Figure 16 shows Intel Xeon Gold Series Skylake with extended reliability.




FIGURE 16. Intel Xeon Gold processor numbering explanation



Table 16 provides an alphanumeric string which describes the characteristics of Intel Xeon processor.

TABLE 16. Alphanumeric string which describes the characteristics of Intel Xeon processor

SKU Level	SKU Descriptions	
8	Platinum	
6, 5	Gold	
Processor Generation		
1	1st Gen Intel Xeon scalable processors (Skylake)	
2	2nd Gen Intel Xeon scalable processors (Cascade Lake)	
Memory Capacity	Skylake	Cascade Lake
No Suffix	768 GB	1.0 TB
M	1.5 TB	2.0 TB
L	3.0 TB	4.0 TB

NOTE

In case of non-availability of **M** processor, **L** processors can be used.

HPE persistent memory module

HPE persistent memory features Intel Optane DC Persistent Memory, together with traditional volatile DRAM DIMMs. It provides fast, high-capacity, cost-effective memory, and storage to transform big data workloads and analytics by enabling data to be stored persistently in memory, moved, and processed quickly.

HPE persistent memory modules are available in the following categories and support a maximum 1.5 TB per processor socket persistent memory capacity:

- 128 GB
- 256 GB

The DRAM capacity to persistent memory capacity ratio supported by HPE for SAP HANA environment is as follows:

- 1 : 1
- 1 : 2
- 1 : 4



Persistent memory architecture

Figure 17 shows the layer of server hardware to the application exposing the layer of persistent memory.

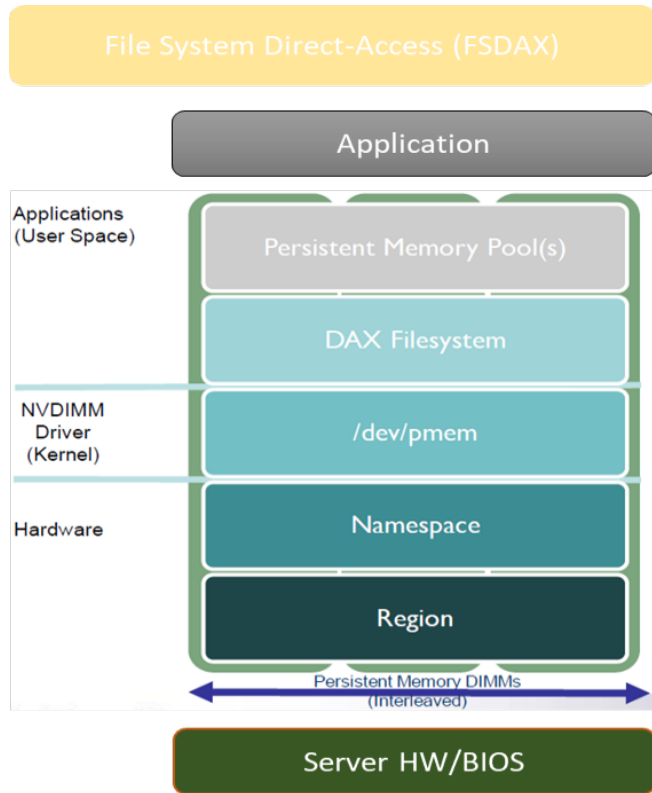


FIGURE 17. Exposing persistent memory to application



Persistent Memory Development Kit (PMDK)

The Persistent Memory Development Kit (PMDK), is a collection of libraries and tools designed to support development of persistent-memory-aware applications. The open source PMDK project currently supports ten libraries, which are targeted at various use cases for persistent memory with language support for C, C++, Java*, and Python*.

The PMDK is designed to solve persistent memory challenges and facilitate the adoption of persistent memory programming. It offers developers well-tested, production-ready libraries and tools in a comprehensive implementation of the Storage Networking Industry Association Non-Volatile Memory (SNIA NVM) programming model.

What is DAX?

In traditional file system model, the paging cache is used during read and write from and to file system and maintains extra copy of buffer in form of paging mechanism. Direct Access (DAX) enables direct access to files stored in persistent memory or on a block device.

DAX removes the extra copy operation by performing read and write directly to the storage device. It is also used to provide the pages that are mapped into a user space by a call to mmap. For more information, refer to [Direct Access for Files](#).

If the system holds block devices which are like memory, the paging cache would be unnecessary. DAX removes this unnecessary copy by performing read and write directly from the storage block devices. To map the files, the storage devices are mapped directly into user spaces.

Block devices, which supports DAX, allow file system creation. The DAX code supports files with a block size equal to the Kernel's PAGE_SIZE. So it is important to specify a block size while creating the file system. During mount, it is important to use “-o dax” and required to use the option in /etc/fstab. This is the only file system supported by SAP HANA for persistent memory.

Memory allocation goal – Namespace

Create a memory allocation goal from the server BIOS utility, which will be read after the next reboot in order to map the Intel Optane DC Persistent Memory Module (DCPMM) capacity into the system address space. In general, persistent memory can operate in different profiles like regular DRAM, persistent memory, or mixed mode, where persistent memory can have persona of DRAM and persistent memory. With respect to SAP HANA only persistent memory profile is supported and the other two options are not supported. Created goal will define how much percentage of persistent memory will be allocated to persistent memory with persistent storage and as DRAM. The task to create and allocate the percentage of persistent memory capacity into the system address space is called the Allocation of Goal. It can be from 0 to 100%.

Region

A region is a portion of system memory presented from one or more DCPMMs.

Non-interleaved region

Regions are created in non-interleaved sets. Interleaving can be 1 to n-way mapping. Figure 18 shows the non-interleaved memory.

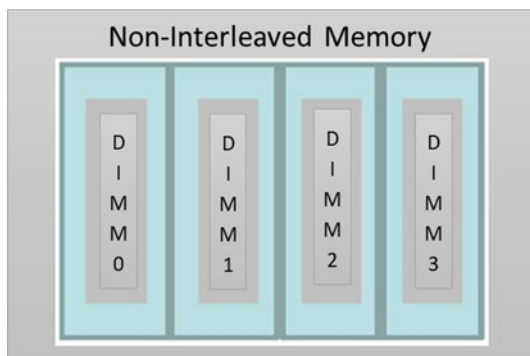


FIGURE 18. Non-interleaved memory diagram



Interleaved region

Interleave region creates contiguous physical address space and provides striped reads/writes for better throughput. The memory block addresses will be continuous in numbers. This interleaved region design is recommended by SAP HANA. Figure 19 shows the interleaved memory.

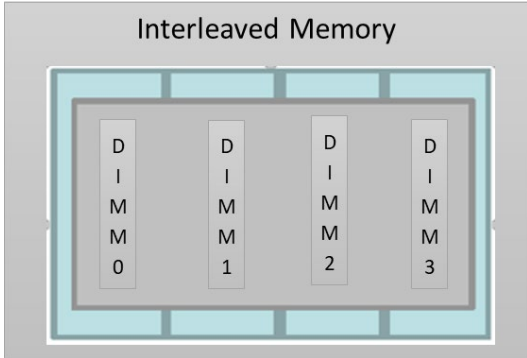


FIGURE 19. Interleaved memory diagram

Namespaces

A namespace is a portion of a region. Similar to hard drives, the raw capacity of region can be partitioned into one or more logical devices called namespaces. Figure 20 shows the creation of namespace on persistent memory.

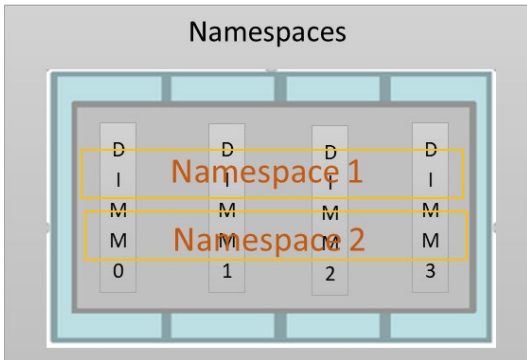


FIGURE 20. Namespace creation on persistent memory

Persistent memory tools

Operating system command line tools

ipmctl: ipmctl is a utility for configuring and managing persistent memory kit. It provides the following functionalities:

- Discover DCPMMs on the platform
- Provision of the platform memory configuration
- View and update the firmware on DCPMMs
- Configure data-at-rest security on DCPMMs
- Monitor DCPMM health
- Track performance of DCPMMs
- Debug and troubleshoot DCPMMs



ndctl: Non-Volatile Device Control ndctl is utility for managing the libnvdimm kernel subsystem. The LIBNVDIMM subsystem defines a kernel device model and control message interface for platform NVDIMM Firmware Interface Table (NFIT). Operations supported by the tool include, provisioning capacity (namespaces), as well as enumerating/enabling/disabling the devices (dimms, regions, namespaces) associated with a persistent memory bus.

SOLUTION CONFIGURATION

This section describes the solution configuration and an overview on how it has been built and what specific configuration has been carried out during the building process.

HPE Superdome Flex Scale-up with D3710 Direct Attached Storage Server configuration for SAP HANA

Table 17 lists the server hardware components used for this configuration.

TABLE 17. List of major hardware components for the HPE Superdome Flex Scale-up with D3710 Direct Attached Storage Server

Components	Details
HPE Superdome Flex Scale-up	Intel Xeon Scalable 2nd Generation processor (Cascade Lake)
Memory DRAM	HPE 64 GB 4Rx4 PC4-2933Y-L Smart Kit
Persistent Memory	HPE 128 GB 2666 MT/s Persistent Mem Mod
Processor	Intel Xeon Platinum 8280M/L CPU @ 2.60GHz
MegaRAID Controller	MegaRAID9361-4i LSI00415 (SGL) LP MD2 4 1 x SFF-8643
External RAID Controller	HPE 3154-8e 8-port External RAID Controller
HDD	2x 1.2 TB 12G 10K SAS HDD RAID1 (or) 2x800 GB 12G SAS SSD RAID1 (Internal Disks for Operation System)
SSD	12x 3.2 TB 12G SAS MU-3 SFF SC SSD – RAID 5
External Storage	HPE D3710 Enclosure
Network	HPE Ethernet 10/25Gb 2P 640SFP28 Adapter HPE Ethernet 1Gb 4-port 331T Adapter HPE Ethernet 10Gb 2P 562T Adapter/ Copper

Network interface card

Components of the SAP HANA landscape communicate via different network channels. It is recommended by SAP to have a well-defined network topology to control, limit, and isolate the network access and usage to only those communication channels required for the particular implementation scenarios.

Various options are available for network interface cards to get the required bandwidth for SAP HANA. The best suitable combination can be selected.

Network interface card option for HPE Superdome Flex Scale-up with D3710 Direct Attached Storage Server

- HPE Ethernet 10/25Gb 2-port 640SFP28 Adapter
- HPE Ethernet 1Gb 4-port 331T Adapter
- HPE Ethernet 10Gb 2-port 562T Adapter/ Copper



Storage

SAP HANA DATA, LOG, and shared volumes have been created on direct-attached storage HPE D3710 through HPE 3154-8e External SmartRAID Controller. It has 12x 3.2 TB of SAS SSDs in RAID 5.

Figure 21 shows SSDs installed on port of the storage controller.

```
MSCC SmartRAID 3154-8e in Slot 2248

Array A

physicaldrive CN0:1:1 (port CN0:box 1:bay 1, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:2 (port CN0:box 1:bay 2, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:3 (port CN0:box 1:bay 3, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:4 (port CN0:box 1:bay 4, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:5 (port CN0:box 1:bay 5, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:6 (port CN0:box 1:bay 6, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:7 (port CN0:box 1:bay 7, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:8 (port CN0:box 1:bay 8, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:9 (port CN0:box 1:bay 9, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:10 (port CN0:box 1:bay 10, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:11 (port CN0:box 1:bay 11, SAS SSD, 3.2 TB, OK)
physicaldrive CN0:1:12 (port CN0:box 1:bay 12, SAS SSD, 3.2 TB, OK)
```

FIGURE 21. SSDs installed on the system

NOTE

For the HPE External SmartRAID controller, the following properties need to be modified:

- Smartpath to be set to disabled on Array
- Caching is required to be enabled on DATA and LOG Logical Drives

This can be performed through HPE Smart Storage Administration CLI (SSACLI).



PERSISTENT MEMORY CONFIGURATION

This section describes the steps to configure the persistent memory goals, regions, and namespaces.

Figure 22 shows the allowed maximum memory per socket and the currently mapped memory to each socket.

SocketID	MappedMemoryLimit	TotalMappedMemory
0x0000	4608.0 GiB	1139.0 GiB
0x0001	4608.0 GiB	1138.5 GiB
0x0002	4608.0 GiB	1140.0 GiB
0x0003	4608.0 GiB	1140.0 GiB

FIGURE 22. Maximum memory per socket and mapped memory

Figure 23 shows the available persistent memory to the system, their ID, capacity, health state, lock state, and firmware version. The system is provisioned with 24x 128 GB of persistent memory DIMMs. The system holds 24x 64 GB LRDIMM. The DRAM to persistent memory ratio is 1:2.

DimmID	Capacity	HealthState	ActionRequired	LockState	FWVersion
0x0001	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x0011	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x0021	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x0101	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x0111	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x0121	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x1001	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x1011	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x1021	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x1101	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x1111	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x1121	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x2001	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x2011	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x2021	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x2101	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x2111	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x2121	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x3001	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x3011	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x3021	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x3101	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x3111	126.4 GiB	Healthy	0	Disabled	01.02.00.5395
0x3121	126.4 GiB	Healthy	0	Disabled	01.02.00.5395

FIGURE 23. Listing out persistent memory available on OS



You can see the created persistent memory goal and the regions. Figure 24 shows the persistent memory configured in AppDirect Memory Interleaving Mode, each region available, and the free capacity memory.

SocketID	ISetID	PersistentMemoryType	Capacity	FreeCapacity	HealthState
0x0000	0xd4dc7f48f8be2ccc	AppDirect	756.0 GiB	0.0 GiB	Healthy
0x0001	0x63be7f48f9db2ccc	AppDirect	756.0 GiB	756.0 GiB	Healthy
0x0002	0x002a7f48cbc52ccc	AppDirect	756.0 GiB	756.0 GiB	Healthy
0x0003	0x72587f4869d42ccc	AppDirect	756.0 GiB	756.0 GiB	Healthy

FIGURE 24. Listing persistent memory goal

Kernel device model of the persistent memory region is created at the server hardware level. Figure 25 shows the Kernel representing the persistent memory region. The command output will show all available regions.

```
[
  {
    "dev": "region1",
    "size": 811748818944,
    "available_size": 811748818944,
    "type": "pmem",
    "numa_node": 1,
    "iset_id": 7187322006736153804,
    "persistence_domain": "memory_controller"
  },
]
```

FIGURE 25. Creation of persistent memory region

Figure 26 shows the namespace creation for Region0. Once the namespace is created, the operating system kernel assigns a kernel block device name, UUID, and the mode of operation. In this case, the mode of operation is **FSDAX**. It is required to repeat the procedure for each region.

```
{
  "dev": "namespace0.0",
  "mode": "fsdax",
  "size": "744.19 GiB (799.06 GB)",
  "uuid": "68b042ac-7eb2-482a-adac-d3db8b03f6ea",
  "raw_uuid": "08e7395d-511a-44a7-8392-bd0f2eeeb468",
  "sector_size": 512,
  "blockdev": "pmem0",
  "numa_node": 0
}
```

FIGURE 26. Persistent memory namespace creation

After all namespaces have been created, the kernel will create the block devices for each persistent memory namespaces and that will appear as block device on operating system. Figure 27 shows fdisk for persistent memory block devices.



```
Disk /dev/pmem0: 744.2 GiB, 799063146496 bytes, 1560670208 sectors
Disk /dev/pmem1: 744.2 GiB, 799063146496 bytes, 1560670208 sectors
Disk /dev/pmem2: 744.2 GiB, 799063146496 bytes, 1560670208 sectors
Disk /dev/pmem3: 744.2 GiB, 799063146496 bytes, 1560670208 sectors
```

FIGURE 27. fdisk for persistent memory block devices

Figure 28 shows persistent memory devices available on the server and details of the persistent memory.

```
pmem0          259:0      0 744.2G  0 disk /hana/pmem/DB1/nvm0
pmem1          259:1      0 744.2G  0 disk /hana/pmem/DB1/nvm1
pmem2          259:2      0 744.2G  0 disk /hana/pmem/DB1/nvm2
pmem3          259:3      0 744.2G  0 disk /hana/pmem/DB1/nvm3
```

FIGURE 28. Listing the block device for persistent memory

Create the file system for all of persistent memory block devices as shown in Figure 29. Repeat these steps for all persistent memory devices. Format the persistent memory for the block size which is supported by kernel for persistent memory namespaces.

```
meta-data=/dev/pmem0      isize=512    agcount=4, agsize=48770944 blks
=                sectsz=4096   attr=2, projid32bit=1
=                crc=1      finobt=1, sparse=0, rmapbt=0, reflink=0
data        =                bsize=4096   blocks=195083776, imaxpct=25
=                sunit=0     swidth=0 blks
naming      =version 2      bsize=4096   ascii-ci=0 ftype=1
log         =internal log   bsize=4096   blocks=95255, version=2
=                sectsz=4096 sunit=1 blks, lazy-count=1
realtime    =none          extsz=4096   blocks=0, rtextents=0
```

FIGURE 29. Filesystem creation on persistent memory block devices

Figure 30 shows that the persistent memory is mounted with DAX support.

```
/dev/pmem1 on /hana/pmem/ type xfs (rw,noatime,nodiratime,attr2,dax,inode64,noquota)
/dev/pmem3 on /hana/pmem/ type xfs (rw,noatime,nodiratime,attr2,dax,inode64,noquota)
/dev/pmem0 on /hana/pmem/ type xfs (rw,noatime,nodiratime,attr2,dax,inode64,noquota)
/dev/pmem2 on /hana/pmem/ type xfs (rw,noatime,nodiratime,attr2,dax,inode64,noquota)
```

FIGURE 30. Persistent memory with mount attributes

Figure 31 shows the available sizes for different persistent memory devices.



```

/dev/pmem1          xfs      744G  792M  744G  1% /hana/pmem/
/dev/pmem3          xfs      744G  792M  744G  1% /hana/pmem/
/dev/pmem0          xfs      744G  792M  744G  1% /hana/pmem/
/dev/pmem2          xfs      744G  792M  744G  1% /hana/pmem/
    
```

FIGURE 31. Persistent memory mount on OS with sizes

With this, configuration of non-volatile memory on the HPE Superdome Flex Scale-up with D3710 Direct Attached Storage Server with Intel Xeon Scalable 2nd Generation processor (Cascade Lake) server is complete. The above figures show how the persistent memory is configured and mounted on the operating system. There are four persistent memory regions and each belongs to a processor socket. The size of each persistent memory region and with respective namespace is around 744 GB. Each region and its respective namespace is visible to the operating system as a mounted block device.

When you boot the server configured with persistent memory and have persistent memory block devices mount through /etc/fstab, the system goes into maintenance mode. This is because the system is trying to mount persistent memory mount points before they are ready to mount.

You can solve this by changing the default timeouts for starting of units.

To change the default timeouts for starting of units:

1. Edit the /etc/systemd/system.conf file.
2. Uncomment the following parameter: #DefaultTimeoutStartSec=90s.
3. Increase the default timeout in seconds to a greater value.

The following values were used for testing:

- For 256 GB of persistent memory, the value is DefaultTimeoutStartSec=360s
- For 128 GB of persistent memory, the value is DefaultTimeoutStartSec=180s

SAP HANA Cockpit view of persistent memory

This section shows how to get information about persistent memory using SQL statement in SAP HANA Cockpit. The tables referred here are built in SAP HANA database. These are the example tables that give overview about the persistent memory and its usage.

NOTE

Screen shots provided in this document are for reference ONLY. They may differ in real time scenario based on the hardware and software configuration.



1. The “m_cs_tables” table contains information about all database users and their respective tables. Space utilized by a particular table can be checked using SAP HANA cockpit. Run the following SQL statement against “m_cs_tables”:

```
select * from m_cs_tables
```

The output of the SQL statement appears as shown in Figure 32. This output gives a complete overview of the persistent memory space consumed by all schema and tables present in DB, along with other details.

HOST	PORT	SCHEMA_NAME	TABLE_NAME	PERSISTENT_MEMORY_SIZE_IN_TOTAL	PERSISTENT_MERGE	PERSISTENT_MEMORY
sle-cl1	30101	_SYS_REPO	ACTIVATION_HELPER2	0	TRUE	FALSE
sle-cl1	30101	_SYS_REPO	ACTIVE_CONTENT_TE	0	TRUE	FALSE
sle-cl1	30101	_SYS_REPO	ACTIVE_CONTENT_TE	43	TRUE	FALSE
sle-cl1	30101	_SYS_REPO	ACTIVE_OBJECT	6481466	TRUE	FALSE
sle-cl1	30101	_SYS_REPO	ACTIVE_OBJECTCROS	2815	TRUE	FALSE
sle-cl1	30101	_SYS_REPO	ACTIVE_OBJECT_TEXT1	0	TRUE	FALSE
sle-cl1	30101	_SYS_REPO	ACTIVE_OBJECT_TEXT1	0	TRUE	FALSE
sle-cl1	30101	_SYS_REPO	ACTIVE_TAGS	0	TRUE	FALSE
sle-cl1	30101	_SYS_TASK	BEST_RECORD_GROU	0	TRUE	FALSE
sle-cl1	30101	_SYS_TASK	BEST_RECORD_RESUI	0	TRUE	FALSE
sle-cl1	30101	_SYS_TASK	BEST_RECORD_STRAI	0	TRUE	FALSE
sle-cl1	30101	_SYS_TASK	BEST_RECORD_STRAI	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_ATTRIBUTE	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_CUBES	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_DIMENSION	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_SOURCES	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_VARIABLES	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_VARIABLE_I	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_VARIABLE_I	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_ALL_VARIABLE_I	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_CONFIGURATION	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_DESCRIPTIONS	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_PROPERTIES	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_VARIABLE_ASSK	0	TRUE	FALSE
sle-cl1	30101	_SYS_BI	BIMC_VERSION_INFOR	0	TRUE	FALSE
sle-cl1	30101	_SYS_RT	CDS_ANNOTATION	0	TRUE	FALSE
sle-cl1	30101	_SYS_RT	CDS_ANNOTATION_AS	0	TRUE	FALSE
sle-cl1	30101	_SYS_RT	CDS_ANNOTATION_VA	0	TRUE	FALSE

FIGURE 32. Persistent memory utilization table from SAP HANA Cockpit

2. Figure 33 shows the utilization of persistent memory by SAP HANA database on their respective mount point for each persistent memory namespace tied to its respective processor. It also shows the Direct Access attached to these persistent memory block devices that are formatted using XFS file system.

HOST	PORT	VOLUME_ID	NUMA_NODE_INDEX	PATH	FILESYSTEM_TYPE	IS_DIRECT_ACCESS_SUPPORTED	TOTAL_SIZE	USED_SIZE
sle-cl1	30101	0	0	/hana/pmem/NV1/nvm0i	xfs	true	3182020825088	12977487872
sle-cl1	30101	1	1	/hana/pmem/NV1/nvm1i	xfs	true	3182020825088	52109312
sle-cl1	30101	2	2	/hana/pmem/NV1/nvm2i	xfs	true	3182020825088	1076236288
sle-cl1	30101	3	3	/hana/pmem/NV1/nvm3i	xfs	true	3182020825088	30916018176

FIGURE 33. View of persistent memory block device usage from SAP HANA Cockpit



3. Run the following SQL statement to view the persistent memory consumed by HANA DB in bytes:

```
select SUM (PERSISTENT_MEMORY_SIZE_IN_TOTAL) from m_cs_tables
```

The output of the SQL statement appears as shown in Figure 34.

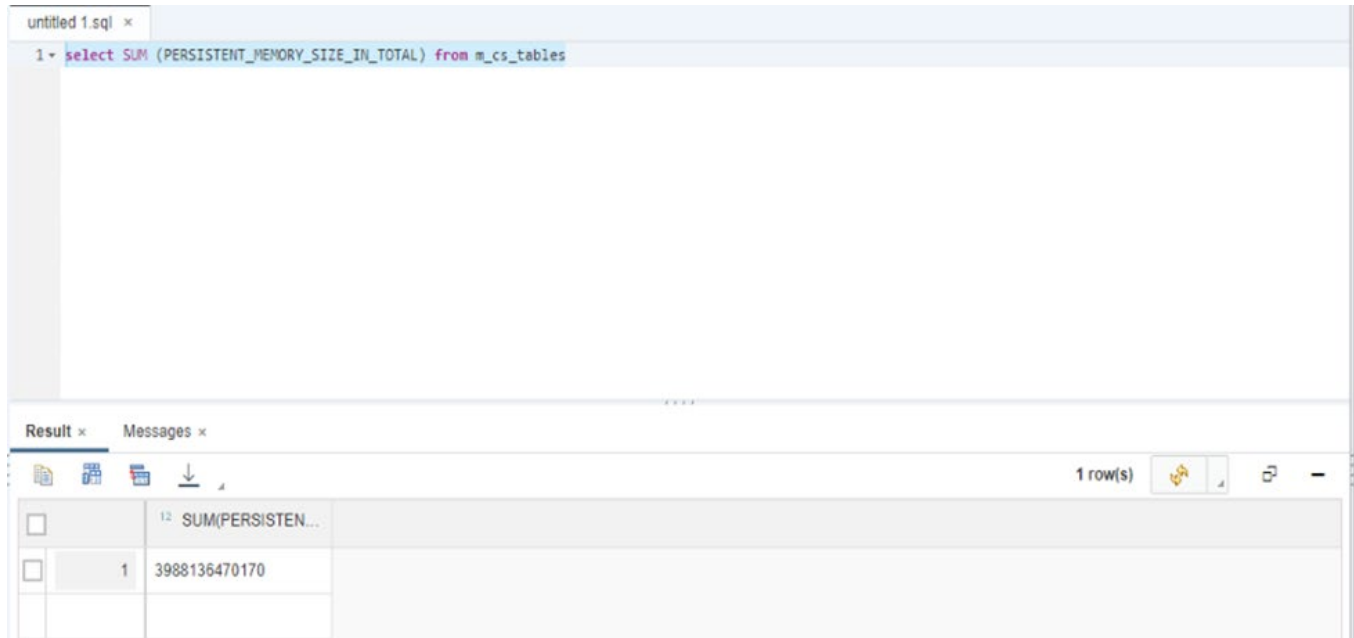


FIGURE 34. Persistent memory consumption by HANA database



- 4. Run the following SQL statement to check space consumption for the table in database that resides in persistent memory for a specific user and schema:

```
select * FROM m_cs_tables WHERE SCHEMA_NAME = 'admin/schema' AND TABLE_NAME = 'table_name'
```

The output of the SQL statement appears as shown in Figure 35.

	TABLE_NAME	PERSISTENT_MEMORY_SIZE_IN_TOTAL	PERSISTENT_MERGE	PERSISTENT_M...
1	TABLE1_000	0	TRUE	TRUE
2	TABLE1_000	0	TRUE	TRUE
3	TABLE1_000	22153114022	TRUE	TRUE
4	TABLE1_000	0	TRUE	TRUE

FIGURE 35. Persistent memory utilization by a table



5. Run the following SQL statement to check the corresponding non-uniform memory access (NUMA) node, file block persistent memory size in use at column level of a table for a particular user.

```
select schema_name, table_name, column_name, PERSISTENT_MEMORY_SIZE_IN_TOTAL, NUMA_NODE_INDEX,
PERSISTENT_MEMORY_FILE_NAME from m_cs_columns where schema_name='ADMIN1'
```

From the above table, we can note the PERSISTENT_MEMORY_FILE_NAME, which we can look up into HANA nvm file systems.

The “m_cs_columns” table also has following fields that can help us to analyze the PMM usage:

- PERSISTENT_MEMORY_SIZE_IN_TOTAL
- MEMORY_SIZE_IN_TOTAL
- MEMORY_SIZE_IN_MAIN
- MEMORY_SIZE_IN_DELTA

The output of the SQL statement appears as shown in Figure 36.

RB	SCHE...	RB	TABLE_N...	RB	COLUMN_NAME	12	PERSISTENT_MEMORY_SIZE_I...	12	NUMA_NODE_INDEX	RB	PERSISTENT_MEMORY_FILE_NAME
531	ADMIN1		TABLE12_000		KVALUE		9000228		0		8604-000000fe000099c8-00000003-000000c9-00000001_167.fileblo
743	ADMIN1		TABLE13_000		KVALUE		9000228		2		9260-000000fe000099ee-00000003-000000c9-00000001_181.fileblo
955	ADMIN1		TABLE14_000		KVALUE		9000228		2		9887-000000fe00009a1c-00000003-000000c9-00000001_195.fileblo
322	ADMIN1		TABLE11_000		C_01		442454892		1		7826-000000fe0000999e-00000003-000000cc-00000001_153.fileblo
323	ADMIN1		TABLE11_000		C_02		442454892		1		7828-000000fe0000999e-00000003-000000cd-00000001_153.fileblo
324	ADMIN1		TABLE11_000		C_03		442454892		1		7835-000000fe0000999e-00000003-000000ce-00000001_153.fileblo
325	ADMIN1		TABLE11_000		C_04		442454892		1		7843-000000fe0000999e-00000003-000000cf-00000001_153.fileblo
326	ADMIN1		TABLE11_000		C_05		442454892		1		7847-000000fe0000999e-00000003-000000d0-00000001_153.fileblo
327	ADMIN1		TABLE11_000		C_06		442454892		1		7851-000000fe0000999e-00000003-000000d1-00000001_153.fileblo
328	ADMIN1		TABLE11_000		C_07		442454892		1		7855-000000fe0000999e-00000003-000000d2-00000001_153.fileblo
329	ADMIN1		TABLE11_000		C_08		442454892		1		7859-000000fe0000999e-00000003-000000d3-00000001_153.fileblo
330	ADMIN1		TABLE11_000		C_09		442454892		1		7863-000000fe0000999e-00000003-000000d4-00000001_153.fileblo

FIGURE 36. Persistent memory file name with respect to NUMA node



- 6. We can also check capacity, usage of persistent memory volumes configured per NUMA node with pre-defined view by running the following command:

```
select * from sys.m_persistent_memory_values
```

The output of the SQL statement appears as shown in Figure 37.

ID	AB	PORT	VO...	NUMA...	PATH	IS_DIRECT_AC...	TOTAL_SIZE	USED_SIZE
1	aep	30103	3	0	/hana/pmем/DB1/nvm0/mnt00001/hdb00003.000	true	1597348057088	864023105536
2	aep	30103	3	1	/hana/pmем/DB1/nvm1/mnt00001/hdb00003.000	true	1597348057088	1063407611904
3	aep	30103	3	2	/hana/pmем/DB1/nvm2/mnt00001/hdb00003.000	true	1597348057088	931525140480
4	aep	30103	3	3	/hana/pmем/DB1/nvm3/mnt00001/hdb00003.000	true	1597348057088	1129870405632
5	aep	30107	2	0	/hana/pmем/DB1/nvm0/mnt00001/hdb00002.000	true	1597348057088	864023105536
6	aep	30107	2	1	/hana/pmем/DB1/nvm1/mnt00001/hdb00002.000	true	1597348057088	1063407611904
7	aep	30107	2	2	/hana/pmем/DB1/nvm2/mnt00001/hdb00002.000	true	1597348057088	931525140480
8	aep	30107	2	3	/hana/pmем/DB1/nvm3/mnt00001/hdb00002.000	true	1597348057088	1129870405632

FIGURE 37. Persistent memory utilization against each file path



- 7. Use `m_persistent_memory_volume_data_files` view to get information about metadata statistics of files created by SAP HANA for data storage on persistent memory values as shown here:

With this view, we can check the numa node information in which the particular file block is present. This will be helpful to know the numa node location of a particular column of a table, as we can file block name from `m_cs_columns`. Figure 38 shows persistent memory and metadata.

```
1 select * from sys.m_persistent_memory_volume_data_files
2
```

ROWID	HOST	PORT	VOLUME...	NUMA...	FILE_NAME	IS_MAPPED	SIZE
10	aep	30103	3	0	20-00000fe0000017-0000003-00000d0-0000000	true	92
29	aep	30103	3	0	2133-00000fe00009858-0000003-00000c9-0000000	true	900228
32	aep	30103	3	0	440-00000fe00009824-0000003-00000c9-0000000	true	900228
47	aep	30103	3	0	464-00000fe00009827-0000003-00000c9-0000000	true	900228
57	aep	30103	3	0	469-00000fe00009803-0000003-00000c9-0000000	true	900228
201	aep	30103	3	0	5073-00000fe0000990b-0000003-00000c9-0000000	true	900228
230	aep	30103	3	0	1491-00000fe00009846-0000003-00000c9-0000000	true	900228
235	aep	30103	3	0	1326-00000fe00009834-0000003-00000c9-0000000	true	900228
298	aep	30103	3	0	4291-00000fe000098d4-0000003-00000c9-0000000	true	900228
403	aep	30103	3	0	2611-00000fe0000989e-0000003-00000c9-0000000	true	900228
423	aep	30103	3	0	2684-00000fe0000988a-0000003-00000c9-0000000	true	900228
466	aep	30103	3	0	2868-00000fe0000988e-0000003-00000c9-0000000	true	900228
541	aep	30103	3	0	7067-00000fe00009970-0000003-00000c9-0000000	true	900228
579	aep	30103	3	0	11198-00000fe00009a66-0000003-00000c9-0000000	true	900228
603	aep	30103	3	0	7902-00000fe000099b4-0000003-00000c9-0000000	true	900228

FIGURE 38. Persistent memory and metadata



- Get the volume statistics, operations performed on NUMA node, such as total create counts and size of the create request. You can use the pre-defined view, for example, "m_persistent_memory_volume_statistics". Figure 39 shows persistent memory utilization by a table.

NOTE

With the view "m_persistent_memory_volume_statistics", you can also get commit, delete, destroy, load, mapped and active events count and size of these operations.

```
1 select * from sys.m_persistent_memory_volume_statistics
2
```

	12	12	12	12	12
	PORT	VOLUME_ID	NUMA_NODE_INDEX	TOTAL_CREATE_COUNT	TOTAL_CREATE_SIZE
1	30103	3	0	3998	863997155732
2	30103	3	1	7067	1063380163503
3	30103	3	2	4498	931505869919
4	30103	3	3	6678	1129859701292
5	30107	2	0	0	0
6	30107	2	1	0	0
7	30107	2	2	0	0
8	30107	2	3	0	0

FIGURE 39. Persistent Memory utilization by a table

- Get the number of active core per NUMA nodes and core IDs associated with a NUMA node using the pre-defined view M_NUMA_NODES:

Select * from sys.m_numa_nodes

Figure 40 shows the output of the SQL statement .

```
1 select * from sys.m_numa_nodes
```

	46	12	12	12	12	46	46	46
	HOST	PORT	NUMA_NODE_ID	NUMA_NODE_INDEX	ACTIVE_LOGICAL_CORE_COUNT	LOGICAL_CORE_IDS	NUMA_NODE_DISTANCES	NEIGHBOUR_NUMA_NODE_IDS
1	aep	30103	0	0	56	0-27,112-139	10,21,21,21	1,2,3
2	aep	30103	1	1	56	28-55,140-167	21,10,21,21	0,2,3
3	aep	30103	2	2	56	56-83,168-195	21,21,10,21	0,1,3
4	aep	30103	3	3	56	84-111,196-223	21,21,21,10	0,1,2
5	aep	30107	0	0	56	0-27,112-139	10,21,21,21	1,2,3
6	aep	30107	1	1	56	28-55,140-167	21,10,21,21	0,2,3
7	aep	30107	2	2	56	56-83,168-195	21,21,10,21	0,1,3
8	aep	30107	3	3	56	84-111,196-223	21,21,21,10	0,1,2

FIGURE 40. Persistent memory utilization and NUMA nodes



CAPACITY AND SIZING

The following recommendation about the sizing for HANA TDI implementation is based on the formula derived by SAP for HANA implementation in TDI environments using HPE Superdome Flex Server. It is sized to offer optimum performance for SAP HANA Implementation on HPE Compute Environment.

For more information on SAP HANA Tailored data center integration (TDI), refer [SAP HANA TDI - Storage Requirements](#).

Use the best practices and the sizing recommendation to achieve optimum performance.

Disk space requirement for SAP HANA installation in appliance and in TDI environments

SAP HANA database installation is always done on shared filesystem /hana/shared/<SID>. As the database size grows more, the traces will be written on HANA installation disk.

For more information, refer [SAP HANA Master Guide for HANA 2.0 SPS04](#).

Storage sizing chart for appliance implementation

As per customer requirements, storage sizing can be changed for appliance implementation.

Table 18 shows the overall formula used to estimate the minimum storage requirement for appliance implementation.

TABLE 18. Formula to derive the storage requirement for appliance configurations

Logical Drive	Purpose	Sizing Formula
Logical Drive – 1	Operating System	1 x MegaRAID9361-4i 2 x 1.2 TB 12G 10K SAS HDD RAID1 or 2 x 800 GB 12G SAS SSD RAID1 (Internal Disks for the OS)
Logical Drive – 2	/usr/sap	50 GB
Logical Drive – 3	LOG	[systems ≤ 512 GB] Sizeredolog = 1/2 x RAM [systems > 512 GB] Sizeredolog (min) = 512 GB
Logical Drive – 4	SHARED	Sizeinstallation (single-node) = MIN (1 x RAM)
Logical Drive – 5	DATA	Sizedata = 3 x anticipated net data size on disk

Table 19 shows the minimum storage requirement for appliance implementation. The servers in this paper used RAID 5 or 50, and it is recommended.

TABLE 19. Minimum storage requirement for appliance configuration and sizing with different sizes of disks

Total Memory (GB)	For Appliance, Minimum space Required	800 GB All SSD	1.6 TB ALL SSD	3.2 TB ALL SSD	6.4 TB ALL SSD
1536	6962	12 - R50	6 - RAID 5	4 - RAID 5	NA
3072	13106	NA	12 - RAID 50	6 - RAID 5	NA
4608	19250	NA	NA	9 - RAID 5	5 – RAID 5
6144	25394	NA	NA	12 - RAID 50	6 – RAID 5
9216	37682	NA	NA	16 - RAID 50	8 – RAID 5



Storage sizing chart for TDI implementation

For TDI implementation, storage sizing can be changed per customer requirements.

Table 20 shows the overall formula used to estimate the minimum storage requirement for a TDI implementation.

TABLE 20. Formula to derive TDI Storage requirement

Logical Drive	Purpose	Sizing Formula
Logical Drive – 1	Operating System	1 x MegaRAID9361-4i 2 x 1.2 TB 12G 10K SAS HDD RAID1 or 2 x 800 GB 12G SAS SSD RAID1 (Internal Disks for the OS)
Logical Drive – 2	/usr/sap	50 GB
Logical Drive – 3	LOG	[systems ≤ 512 GB] Sizeredolog = 1/2 x RAM [systems > 512 GB] Sizeredolog (min) = 512 GB
Logical Drive – 4	SHARED	Sizeinstallation (single-node) = MIN (1 x RAM; if RAM ≤ 1 TB) Sizeinstallation (single-node) = MIN (if RAM > 1 TB; 1 TB)
Logical Drive – 5	DATA	Sizedata = 1.2 x anticipated net data size on disk

Table 21 shows the minimum storage requirement for TDI implementation. The servers in this paper used RAID 5 or 50, and it is recommended.

TABLE 21. Minimum storage requirement for SAP HANA TDI implementation

Total Memory (GB)	For TDI, Minimum Space Required	800 GB All SSD	1.6 TB ALL SSD	3.2 TB ALL SSD	6.4 TB ALL SSD
768	2508	5 - RAID 5	3 - RAID 5	3 - RAID 5	NA
1536	3686	7 - RAID 5	4 - RAID 5	3 - RAID 5	NA
3072	5529	9 - RAID 5	5 - RAID 5	3 - RAID 5	NA
4608	7372	12 - RAID 50	7 - RAID 5	4 - RAID 5	NA
6144	9215	15 - RAID 50	8 - RAID 5	5 - RAID 5	3 - RAID 5
9216	12902	NA	12 - RAID 50	6 - RAID 5	4 - RAID 5

NETWORK REQUIREMENTS

SAP recommends three network zones for SAP HANA implementation and achieving optimal performance. It also suggests to allocate the minimum requirements for basic network bandwidth and throughput requirements.

SAP recommends configuring with a minimum 10GbE full-duplex for all network zone with full redundancy. For management, it will be 1GbE with high availability.

- Client Network Zone
 - SQL Client Communication
 - HTTP Client Communication
 - Management Network
- Internal Network Zone
- Internode Network
 - System Replication
 - Storage Replication
 - Backup Network



- Storage Network Zone
 - Shared/NFS Network
 - Fibre Channel Network

Configured networks

Components of the SAP HANA landscape communicate via different network channels. It is recommended by SAP to have a well-defined network topology to control, limit, and isolate the network access and usage to only those communication channels required for the particular implementation scenarios. Table 22 shows the network bond or VLAN for SAP HANA.

TABLE 22. Network bond/VLAN for SAP HANA purpose

Network Zone	Purpose	Bandwidth and Bond/VLAN
DATA 10	Purposed for data traffic from other SAP applications. Generally, this network is connected to SAP application servers and other source systems, which are supposed to interact with HANA landscape.	10 Gbps / Bond 2 NIC
DATA 1	Same as DATA 10 with 1 Gbps	1 Gbps / Bond 2 NIC
Management network	To connect to landscape and handle the management task	1 Gbps / Bond 2 NIC
Backup	Dedicated backup link to take the backup of HANA landscape log, data, and shared volumes	10 Gbps / Bond 2 NIC
User/BI	Purposed to external connection for SAP HANA client to run SQL queries and/or for external application to connect to SAP HANA for fetch the analytical data	10 Gbps / Bond 2 NIC
Replication	Network connections network two HANA Landscape to replicate the HANA Landscape from primary to secondary HANA system	10 Gbps / Bond 2 NIC
Quorum for SG	Required if HPE Serviceguard is used to automate failover process across 2 systems, used to manage the failover and failback of HANA Landscape	10 Gbps / Bond 2 NIC

To provide high availability, it is important to create BONDS between more than one physical NIC. If the high bandwidth NIC is used for VLAN configuration under BOND, it is again important to restrict the maximum utilization of each VLAN.

TUNING THE SAP HANA OPERATING SYSTEM

- For recommended OS settings for SLES 12 / SLES for SAP Applications 12, refer SAP Note: [2205917](#)
- For recommended OS settings for RHEL 7, refer SAP Note: [2292690](#)
- To optimize the Network Configuration on HANA- and OS-Level, refer SAP Note: [2382421](#)
- For supported operating system, refer SAP Note: [2235581](#)
- To prepare SLES for SAP environments, refer SAP Note: [1275776](#)

Workload description

HPE Superdome Flex Scale-up with D3710 Direct Attached Storage Server with Intel Xeon Scalable 2nd Generation processor (Cascade Lake) Server is suitable for following workload. Refer the SAP Note for specific workload.

- BW on HANA
- Suite on HANA



SUMMARY

This document describes the hardware for SAP HANA in the HPE Superdome Flex environment using D3710 Direct attached storage. This document also talks about SAP Hardware Configuration Check Tool briefly.

HPE Superdome Flex Servers provide a flexible, simple, and efficient option for SAP HANA TDI environments:

- Flexible: Customers can choose the best storage option for their environment, either shared storage (HPE 3PAR, Nimble, MSA, XP, or 3rd-party HANA certified storage), internal drives or local disks (D3710) can be leveraged.
- Simple: Reduce complexity by using common composable infrastructure and manage it.
- Efficient: Consolidate and standardize infrastructure for mixed workloads, while still meeting high-performance requirements.



RESOURCES AND ADDITIONAL LINKS

HPE Reference Architectures, <https://www.hpe.com/docs/reference-architecture>

HPE Servers, [hpe.com/servers](https://www.hpe.com/servers)

HPE Storage, [hpe.com/storage](https://www.hpe.com/storage)

HPE Superdome Flex Servers, [hpe.com/superdomeflex](https://www.hpe.com/superdomeflex)

HPE D3000 Enclosures, [hpe.com/d3000](https://www.hpe.com/d3000)

HPE Networking, [hpe.com/networking](https://www.hpe.com/networking)

HPE Technology Consulting Services, [hpe.com/us/en/services/consulting.html](https://www.hpe.com/us/en/services/consulting.html)

SAP HANA Tailored data center integration (TDI), [SAP HANA TDI - Storage Requirements](#)

SAP HANA database installation, [SAP HANA Master Guide for HANA 2.0 SPS04](#)

To help us improve our documents, please provide feedback at [hpe.com/contact/feedback](https://www.hpe.com/contact/feedback).

© Copyright 2020-2024 Hewlett Packard Enterprise Development LP. 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.

SAP and SAP HANA are the trademarks of SAP SE in Germany and in several other countries. Intel, Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. CacheVault is trademark of LSI. All third-party trademarks are the property of their respective owners.

a50000936enw, June 2024