



Hewlett Packard
Enterprise

HPE Reference Architecture for SAP HANA TDI using HPE MSA 2052 SAN Storage and HPE ProLiant DL560

Contents

Executive summary.....	3
Introduction.....	4
Solution overview.....	6
Reference Configurations	6
Design principles.....	10
SAP HANA.....	10
Servers.....	11
Storage.....	11
Networking.....	13
Storage fabric.....	15
Management.....	17
Power.....	17
Data protection.....	17
Solution components.....	20
Hardware.....	20
Networking.....	24
Software	26
Licensing specifications.....	28
Best practices and configuration guidance for the solution.....	28
Server recommendations	28
Storage recommendations.....	29
Networking recommendations.....	32
SAP certification.....	32
Capacity and sizing	33
Storage replication considerations.....	34
Workload description.....	34
Analysis and recommendations	34
Key Reference Architecture points.....	35
Summary.....	36
Implementing a proof-of-concept.....	36
Appendix A. Bill of Materials.....	36
Appendix B. Optional rack configurations.....	39
Appendix C. Protecting SAP HANA.....	41
Appendix D. HPE StoreOnce Catalyst Plug-in for SAP HANA.....	44
Resources and additional links	46

Executive summary

When SAP announced its new database in 2008, customers were somewhat skeptical; however, today SAP HANA, an in-memory database, has demonstrated how the vision of an Internet of Things is possible. The enhanced performance and improved compression of an in-memory database has sparked the imagination of businesses across the globe. These new possibilities make SAP HANA the platform of choice for many mission-critical applications and set a new standard for databases.

Along with great possibilities, SAP HANA has created new challenges for IT organizations. With SAP HANA's enhanced performance characteristics, the rules appear to have changed around the benefits of traditional IT operations, protection, and disaster recovery. For example, if your backup solution has been solely based on Linear Tape-Open (LTO) tape speeds, can your business survive an SAP HANA outage using a process that takes an hour to restore a single terabyte of data?

Fortunately, Hewlett Packard Enterprise (HPE) and SAP have partnered for close to three decades to bring our customers high-performing, mission-critical business solutions. And SAP HANA has been at the forefront of our partnership with SAP for the last ten years. HPE is the global leader for SAP HANA and approximately 50% of all SAP HANA deployments are on HPE systems. Combining HPE servers and storage with HPE Pointnext services is key to the HPE SAP HANA solution. HPE was one of the first vendors to certify an SAP HANA appliance; and with the introduction of SAP's Tailored Datacenter Integration (TDI), HPE continues its tradition of closely working with SAP to make sure our servers and storage arrays exceed SAP HANA's Key Performance Indicators (KPIs) for certification.

As we see more organizations shift away from SAP HANA appliances to TDI integrations, customers continue to ask us for simpler solutions that are more cost effective, scale quickly, and are also highly available. Besides being a major SMB and mid-market storage contender, the HPE MSA 2052 SAN Storage system (HPE MSA 2052) using all-flash SSDs has met the rigorous SAP HANA KPIs to be a TDI-certified storage array. The HPE MSA 2052 is a simple, space-efficient, and cost-effective array with a great deal of versatility. For example, you can directly attach the HPE MSA 2052 to up to eight servers using Fibre Channel. Or, the HPE MSA 2052 can be used as part of a composable storage infrastructure, supporting multitenant data-center architectures.

And it is fast! How fast? So fast HPE customers are using the HPE MSA 2052 to load SAP HANA databases connected to HPE ProLiant DL servers, HPE Mission Critical servers, and even to HPE Superdomes. The HPE MSA 2052 can exceed 200,000 input/output operations per second (IOPS). Its ability to sequentially load an SAP HANA database upon startup is impressive.

The HPE MSA 2052 SAN Storage system has a rich set of features that can be used to protect an SAP HANA environment. Due to the database's continuously updating nature, the HPE MSA 2052's replication feature makes for an attractive protection option. Data replication to another HPE MSA storage array is significantly faster than file-based backup solutions, while providing a level of crash-consistent protection with a fast recovery option. With the HPE MSA's ability to take snapshots of replicated volumes, not only can you reduce recovery time significantly, you can also use snapshots for testing and development.

This Reference Architecture shows how to implement the HPE MSA 2052 SAN Storage system to support and protect your SAP HANA database environment, using the latest HPE ProLiant DL560 Generation10 (Gen10) servers. These 2U X86 servers are packed with memory and CPU horsepower. When coupled with the 2U HPE MSA 2052, you can bundle up to 96 TB of SAP HANA database services into a single rack.

Target audience

This Reference Architecture has been created for IT systems architects, systems engineers, database administrators, and SAP solution engineers considering SAP HANA TDI solutions.

Document purpose

This Reference Architecture describes the technical and architectural considerations necessary when using HPE MSA 2052 arrays, HPE ProLiant DL560 Gen10 servers, and their supporting infrastructure for an SAP HANA database configuration. This document is specifically written based on using the following versions of technology. The versions indicated were the available versions at the time of publication. This configuration exceeded all the key performance indicators (KPIs) set forth by SAP for production implementations of SAP HANA 2.0 TDI.

- HPE ProLiant DL560 Gen10 servers
- HPE MSA 2052 SAN Storage
- SAP HANA 2.0 SPS 02
- SUSE Linux® Enterprise Server for SAP Applications 12 SP3
- 10 Gb Ethernet
- 16 Gb Fibre Channel (FC)

Introduction

In 2008, SAP introduced SAP HANA, a new database technology that changed the database market. It was a database boasting exceptional performance, deduplication, and capabilities for on-the-fly analytics. To answer the ever-pressing need for faster data access, SAP started the move to an in-memory database model with SAP HANA by building upon a column-based data strategy, rather than the traditional row-based implementation. That was ten years ago when servers and mini computers supported 8 to 16 Cores and a maximum of 64 GB of memory.

Today, SAP HANA is at the core of mission-critical business and research applications throughout the world. Customers have multiple petabytes of data distributed and replicated throughout their environments. They are trying to keep up with availability due to an ever increasing demand for data. The sheer volume of data can be overwhelming; however, add the expectation of availability with the need to protect data from disasters and you find yourself with a very challenging IT problem.

Fortunately, HPE and SAP have partnered for close to three decades to bring our customers high-performing, mission-critical business solutions. HPE is the number one infrastructure provider for SAP HANA (Gartner Group 2018). Key to HPE SAP HANA solutions are HPE servers, storage, and services. With the introduction of SAP's Tailored Datacenter Integration for SAP HANA, HPE has seen a large shift in customers moving from SAP HANA appliances to TDI-certified custom datacenter configurations. SAP HANA customers are continually looking for better solutions to keep up with and to manage this new world of in-memory solutions.

With SAP HANA, the challenge is understanding where traditional database theories align with the new in-memory database paradigm and where they do not synchronize. For example, one of the first questions administrators ask is this: "With an in-memory database do I need a highly available storage array to persist my data?" That is a very good question. Consider the following statement from the [SAP HANA Administration Guide](#).

To maintain optimal performance, an SAP HANA database holds the bulk of its data in-memory. However, SAP HANA also uses persistent storage to provide a fallback in the event of a fault or a failure.

During normal database operation, changed data is automatically saved from memory to disk at regular savepoints. By default, savepoints are set to occur every five minutes...

Persisting data to storage for "a fault or a failure" purpose sounds like a backup. Has Tier 1 storage become a backup mechanism with SAP HANA? In some respects, it has. This is a very different perspective from traditional databases, which considered Tier 1 storage persistence as the original copy of data, and where critical analysis of your storage transactions, queue depths, and spindle counts were essential metrics to troubleshoot or determine database performance. The amount of RAM in your database server was a fraction of your database size and the server was secondary to how fast your storage array needed to be. Gone are those days when you would watch as your database server's CPU idled along at about 15% utilization while the storage array maxed itself out.

With SAP HANA, the database tables have literally changed from rows to columns. Data is deduped automatically and all data stays resident in memory. With server memory and CPUs decreasing in price, larger amounts of transactional data are being maintained in memory. Completed transactions are written to disk as log files, and database differences are written at savepoint intervals of every five minutes to make sure the database can be reloaded upon startup. Welcome to the world of SAP HANA in-memory databases. Instead of operation centers and database administrators needing to keep track of volume sizes as the key indicator of database growth, SAP HANA engineers will be watching Key Performance Indicators for memory and CPU usage.

With server-memory sizes reaching the double- and triple-digit Terabyte ranges, using a traditional backup model to recover an SAP HANA database can cripple an enterprise application. As you look into how to architect an SAP HANA environment, you will find yourself questioning traditional database backup theories. Here are a few questions you might need to ask: “How do I protect an in-memory database like SAP HANA?” “Will the traditional data protection models work for my organization?” (See Figure 1 below.) “Should I use the three-media rule of thumb for data backups?” And finally, “Do I allow my database to continually grow, or do I distribute it into multiple database nodes?” Fortunately, HPE and SAP have done the work to address these questions and make sure you have the best available solutions to support your business.

The 3-2-1 rule for best practice data protection

Protecting data against any failure...wherever it lives

3-2-1 best practice: **Three** copies of data, **two** copies on **two** different types of media, **one** copy off-site

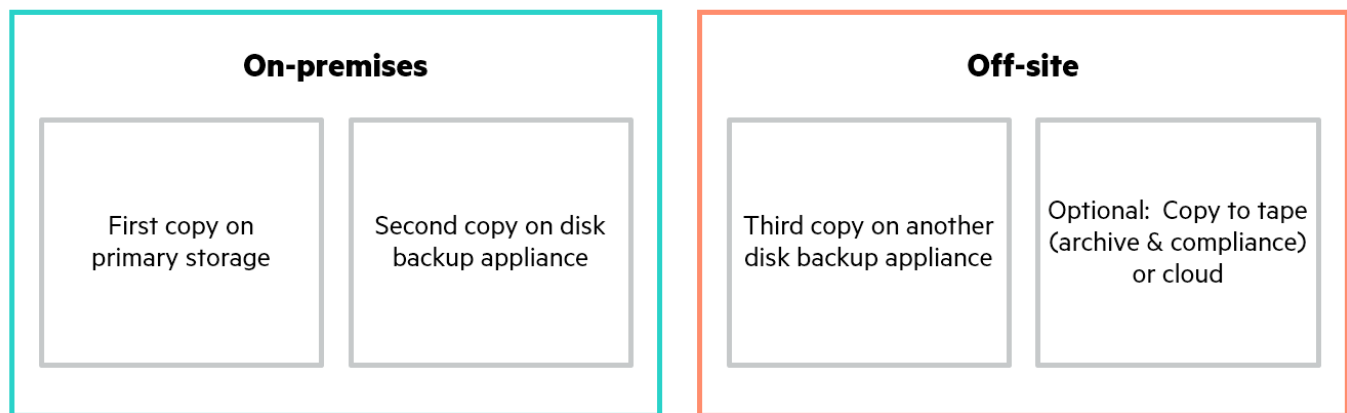


Figure 1. Traditional 3-2-1 backup best practice

This Reference Architecture addresses many of the questions customers are facing with SAP HANA, while providing the best possible TDI solution using Hewlett Packard Enterprise’s award-winning HPE ProLiant DL560 Gen10 servers, HPE MSA 2052 SAN Storage systems, HPE Fibre Channel switches, and HPE Aruba Networking equipment. This document brings together many of the best practices and insights for infrastructure configuration, operating system (OS) tuning, and protection options available for SAP HANA, as implemented in an HPE environment.

First we will look at the overall solution design, equipment, and software that makes up this HPE and SAP HANA TDI solution. We will show a simple introductory test and development implementation built upon a single HPE ProLiant server configured with direct connections to a single HPE MSA Storage array. Next, we will look at a small group of HPE ProLiant servers supported by two HPE MSA Storage arrays. Then, we will show how this architecture can scale to a data center configuration that can support multiple SAP HANA Database Containers (MDCs). We also include a review of solution components, followed by a section on best practices that HPE recommends for implementing this solution. The final section addresses capacity-planning considerations and options. Appendices have been included to provide information about options of interest relating to SAP HANA in an HPE environment.

As you study this document, we encourage you to follow the [HPE Alliances Community](#) website to get the latest announcements and developments about HPE and SAP HANA.

Solution overview

The SAP HANA Tailored Datacenter Integration (TDI) certification process was established to help customer's select equipment components that work together in a customized and supportable production environment. Although TDI-certified equipment is highly encouraged, SAP HANA does not require TDI-certified equipment for development, test, and other implementations; TDI-certified hardware is only required for production environments. HPE understands how important your business is and believes any solution we recommend should meet and exceed the expected production support levels for your solution, independent of its usage. Your business is always in production to HPE.

This Reference Architecture uses HPE ProLiant DL560 Gen10 servers in conjunction with HPE MSA 2052 SAN Storage systems to implement an SAP HANA 2.0 SPS 02 production IT environment. The solution includes redundant 10 Gb Ethernet connections to support IP networking to front-end and back-end network workloads, and 1 Gb Ethernet switches for out-of-band management and heartbeat VLANs. Storage connectivity is provided by multi-path, 16 Gb Fibre Channel (FC) switch fabrics.

HPE sees data protection as an essential part of a production environment. Protecting an SAP HANA database can be done in many ways with varying levels of Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO). This Reference Architecture shows how features of the SAP HANA Cockpit and HPE MSA 2052 can create a simple protection model for SAP HANA. This model uses file-based backups, virtual volume replication and snapshots to support backup, clone, and copy configurations. Appendix C: Protecting SAP HANA, addresses some of the challenges and pitfalls of protecting SAP HANA, while Appendix D: HPE StoreOnce Catalyst Plug-in for SAP HANA, provides information about the HPE StoreOnce Catalyst Plug-in for SAP HANA, which can provide the best data protection plan for the SAP HANA database environment with the HPE MSA 2052 Storage systems.

Reference Configurations

Introductory configuration

The introductory configuration for this Reference Architecture is included to show how a very basic server and storage array configuration can be used and expanded for testing purposes. This simple configuration includes a single HPE ProLiant DL560 Gen10 server and a single HPE MSA 2052 array. They are connected using a direct-attached FC connection, instead of a SAS or iSCSI connection because we anticipate you will want to move these two devices to a storage-switch fabric, as you attach more servers to the storage array.

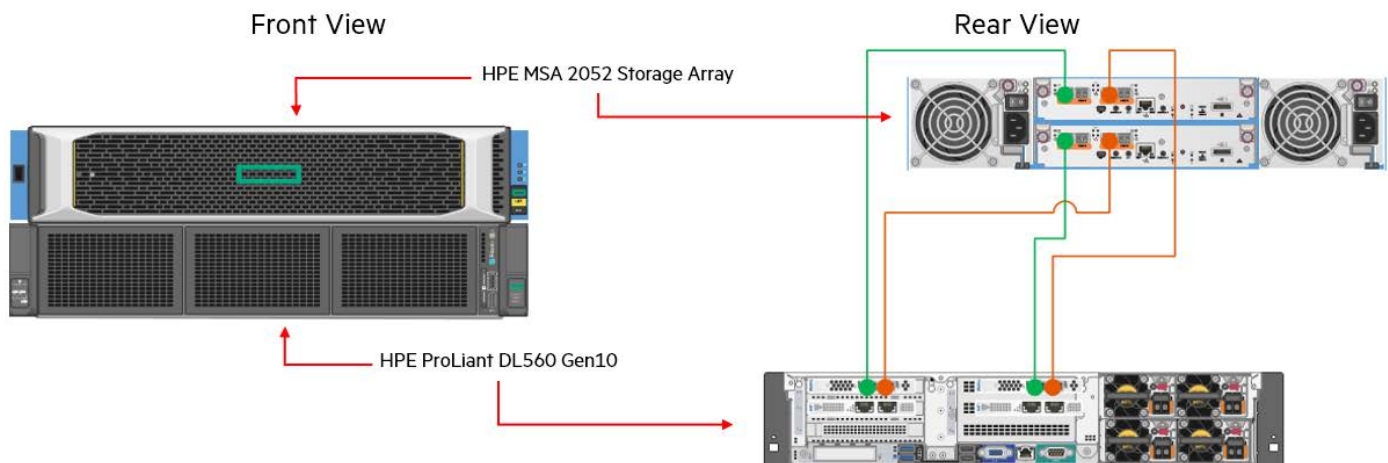


Figure 2. Introductory single-server configuration

HPE has many servers that are SAP HANA TDI certified, which could be implemented as part of this Reference Architecture. For this Reference Architecture we selected the four-processor, HPE ProLiant DL560 Gen10 server because it is an economical server for SAP HANA from both cost and data center footprint considerations. At 2U, this compact server supports both SAP HANA scale-up and scale-out configurations with memory that scales from 384 GB to 6 TB of RAM. If you are creating a composable infrastructure, this server is a great option. Other HPE compute servers compatible with SAP HANA TDI are shown in Figure 3.




Figure 3. HPE compute blocks for the SAP HANA TDI portfolio

Throughout this Reference Architecture, each server has been configured with two dual-port FC Host Bus Adapters (HBAs), allowing multiple FC ports to connect directly, or through a switch, to the HPE MSA 2052 SAN Storage arrays. This enables the SUSE SLES or Red Hat® operating system's multipath SCSI manager to manage all available paths to the array.

Note

Even though you can directly attach servers to the array through Fibre Channel, we recommend using a switch fabric in most configurations. This creates a redundant, high availability (HA), scale-up path for the HPE MSA array to support multiple servers and protection options. Because the HPE MSA 2052 supports directly attached Fibre Channel (FC) connections, this introductory configuration does not require a Fibre Channel switch; however, using a pair of FC switches will make the process of expanding this configuration easier in the future.

Six-server configuration

Although the introductory configuration is a great testing setup, you should strongly consider a group of HPE ProLiant DL560 Gen10 servers with multiple HPE MSA 2052 SAN Storage arrays for your investigation of the production features of SAP HANA. This six HPE ProLiant DL560 server grouping with two HPE MSA 2052 arrays is the configuration used for this Reference Architecture. It provides enough resources to test a production SAP HANA scale-up or scale-out configuration. With this Reference Architecture configuration, you can test multiple Worker and Standby nodes, as well as SAP's system replication. However, to really test a Disaster Recovery (DR) configuration, you will want to implement one of these configurations at two different locations.

Note

For scale-out configurations you will need an NFS source for a shared volume, which is not a feature of the HPE MSA 2052 Storage array. This six-server configuration can help you determine what data protection configuration works best for your business needs and what will fit within your budget.




Figure 4. Six-server configuration

Data center rack configuration

The final configuration of this Reference Architecture shows how two six-server configurations can support a data center rack configuration. This composable design shows how to get the best server density in a rack, while providing a component-based assembly of commodity products commonly used in service provider data centers.

By using 2U HPE ProLiant DL560 Gen10 servers and 2U HPE MSA 2052 SAN Storage arrays, a single rack can support twelve, 6 TB SAP HANA database servers. That is 72 TB of SAP HANA database server nodes in a single rack. That is a very economical solution if your environment requires multiple database nodes under 6 TB. If you need a single-node system with greater than 6 TB of database space, consider HPE's Superdome Flex, Integrity Superdome X, or MC990X systems.

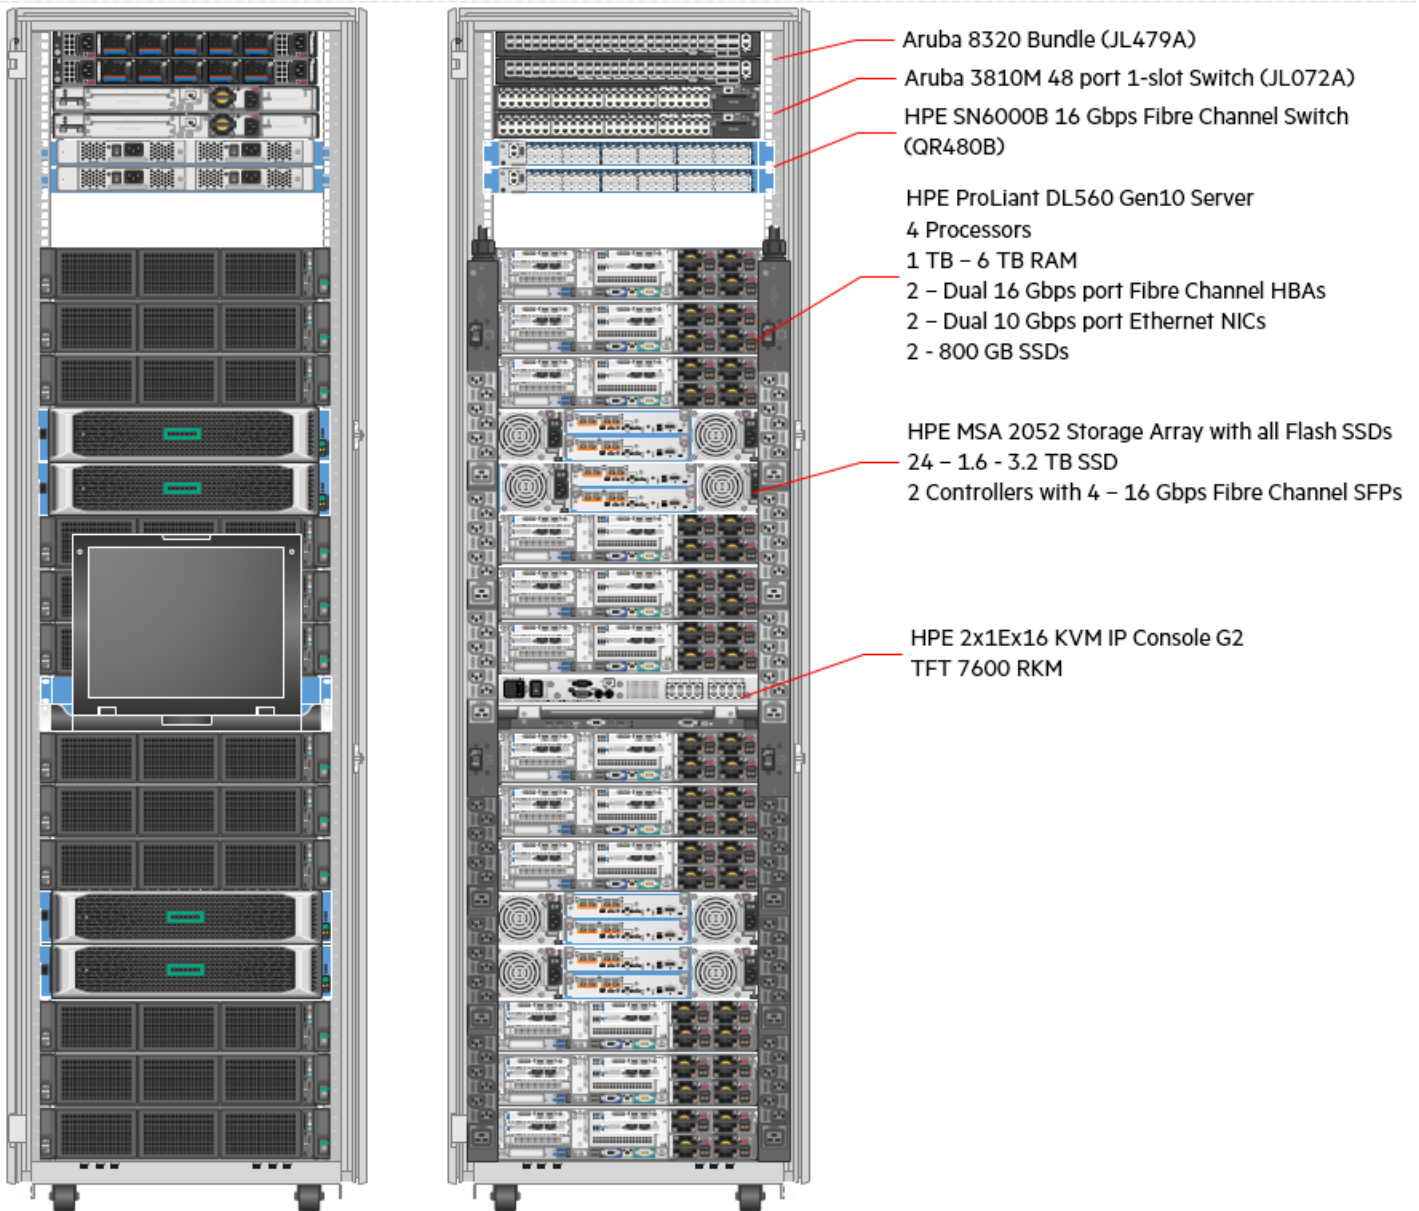


Figure 5. Twelve-server data center configuration

Design principles

When considering this SAP HANA production environment with HPE ProLiant servers and HPE MSA Storage arrays, we used the same design principles used for designing Highly Available (HA) and Disaster Recovery (DR) environments. This includes considerations for servers, storage, networking, storage fabric, management, and data protection. This system-design process ensures the overall configuration will exceed the SAP HANA TDI requirements, while maximizing how the solution will be managed, protected, and recovered by the IT engineering staff. This design includes the following elements of the system solution:

- SAP HANA
- Servers
- Storage
- Networking
- Storage Fabric
- Management
- Power
- Protection and recovery

SAP HANA

The SAP HANA TDI database environment used for this Reference Architecture is SAP HANA 2.0 SPS 02. SAP customers can implement SAP HANA in several ways. This Reference Architecture does not limit customers' ability to configure or change their SAP HANA TDI configuration.

Required services

The SAP HANA server environment requires multiple supporting services, which are part of an IT organization's standard configuration. The configuration and installation of all of these services are beyond the scope of this document. Table 1 shows the services that are expected to be available in order to support this Reference Architecture.

Table 1. SAP HANA required services

Service	Description	Supports	Delivery Method
Shared File Service	SAP HANA installations require a shared volume space where configurations are stored.	In order to support scale-out, worker and standby failover configurations, the <code>/hana/shared</code> directory must be available through a file-sharing service.	NFS
Cockpit/SAP HANA Studio	SAP HANA management is performed through the SAP Cockpit or SAP HANA Studio. These services are generally installed on separate servers for production environments.	SAP HANA system and database configuration, database management, backup management, and tenant administration.	Production environments should have a dedicated server for the SAP Cockpit or Studio. Test and Dev environments can install these services in conjunction with the SAP HANA services.
Domain Name Service (DNS)	Domain Name Services simplify the process of resolution of host to service connectivity.	SAP HANA uses DNS services for host discovery and server-to-server communications. Although not required if using static IP addressing, using DNS should be part of all production services.	For more information about Public and Private DNS Services, go to: https://developers.google.com/speed/public-dns/
Network Time Protocol (NTP)	NTP is used to keep system time clocks in sync to create a common time reference.	Although not required for SAP HANA specifically, an NTP time reference is highly recommended for keeping servers, storage arrays, and infrastructures time in sync.	For more information about NIST Internet Time Servers, go to: https://tf.nist.gov/tf-cgi/servers.cgi
Simple Mail Transfer Protocol (SMTP)	SMTP is standard protocol used for many email services.	SAP HANA uses SMTP to forward alerts and notification messages.	SMTP Mail Service, Sendmail, Apache, and Exchange.
Authentication Service (DS/AD)	Authentication services provide user and group access controls.	Authentication services provided to user and group are critical to the security and access control of all SAP HANA applications.	Directory Services (DS), Active Directory (AD), and LDAP.

Servers

The HPE ProLiant DL560 Gen10 server is included in this Reference Architecture for the following reasons:

- SAP HANA TDI certified server
- Certified for both SAP HANA Scale-up and Scale-out configurations
- RAM expansion up to 6 TB
- Sufficient I/O slots to support networking and storage fabric redundancies
- Redundant power protection
- Out-of-band management
- Small 2U Rack footprint

In the HPE Storage Solutions lab, each HPE ProLiant DL560 Gen10 server tested with this Reference Architecture was configured with a minimum of 1.5 TB of RAM using 64 GB DIMMS and four Intel® Xeon® 8180 (28 Core) series processors. Using four Intel Xeon 8180M processors and 128 GB DIMMS, the HPE ProLiant DL560 can max out at 6 TB of RAM. This provides an excellent scale up capability for your server as your SAP HANA database grows.

An onboard quad Gb Ethernet card provides network connectivity for SAP management, heartbeat, and shared access VLANs. Using PCIe Slots 2 and 5 on the HPE ProLiant DL560 Gen10 server we included two-dual port 16 Gb Fibre Channel HBAs for redundant connections to the storage fabric. (See Storage fabric for details about connectivity design.) PCIe slots 3 and 6 contain dual port 10 Gb Ethernet Cards for redundant connections to front-end and back-end networks. (See Networking for details about configuration and VLAN design.)

Two HPE 800 GB Solid State Drives (SSD) are mirrored and used for boot and swap volumes. This configuration provides adequate volume space to support the SUSE Enterprise Server for SAP OS installation, SAP HANA runtime executables, and swap space for a 6 TB server. Although it is possible to boot from SAN or boot from LAN, HPE and SAP both recommend local storage for both boot and swap volumes.

Storage

The HPE MSA 2052 SAN Storage array is included in this Reference Architecture for the following reasons:

- SAP HANA TDI certified
- Storage array with all flash SSDs
- Exceptionally fast database load performance
- Array-to-array virtual volume replication
- Simplicity of setup and management
- Redundant power protection
- Small 2U rack footprint

This Reference Architecture recommends using two HPE MSA 2052 SAN Storage arrays for production environments: one as the primary storage persistence for SAP HANA databases and a second HPE MSA 2052 as a remote replication array. In the six-server configuration we recommend, three servers use one array as primary storage and the other three servers use the second array as primary storage with replication to the opposite arrays as shown in Figure 6.

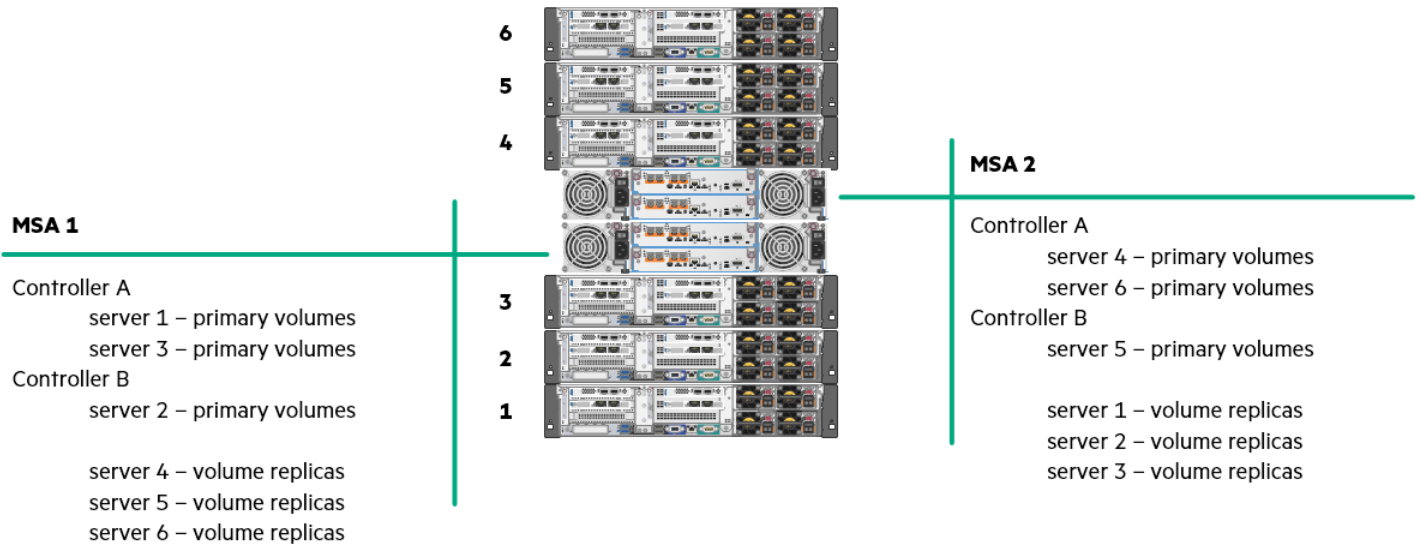


Figure 6. HPE MSA Controller balancing schema

The HPE MSA 2052 is a great utility array when used as a single application array in an SAP HANA database environment. If your configuration requires a mix of different types of application servers like Oracle databases or an Exchange services, HPE recommends you consider the HPE 3PAR storage offering. This arrays is better suited for mixed-application services and also offers increased levels of availability and data services, including HPE InfoSight predictive analytics.

HPE All-Flash Storage Portfolio for SAP HANA

– Comprehensive offering for your TDI environment




Figure 7. HPE All-Flash Storage portfolio for SAP HANA

Each HPE MSA 2052 in this architecture is configured with two quad-port multi I/O Storage Controllers. This allows up to eight FC connections to the storage fabric or eight iSCSI ports to an Ethernet network, or a combination of four FC and four iSCSI ports. The FC/iSCSI configuration is a great option for using iSCSI for remote replication over Ethernet.

Each storage enclosure holds 24 Small Form Factor (SFF) drives. This Reference Architecture uses the HPE MSA 2052 SAN Storage array with all-flash SSDs for both the primary and secondary storage arrays. Using twenty-four, 3.2 TB SSDs, each array can hold 72 TB of flash storage to support the data, logs, shared volumes, and snapshot.

To add to your protection plan, consider adding a hybrid HPE MSA 2050 SAN Storage array to your configuration. A secondary copy of snapshots can be accomplished using the HPE MSA's volume replication to a secondary location. Using a Hybrid HPE MSA can provide greater storage capacity at a more cost-effective price point to support longer-term storing of storage snapshots. Using the HPE MSA's FC/iSCSI connection option makes the remote replication over iSCSI simple, while keeping the primary array access over Fibre Channel.

The standard HPE MSA 2050 Hybrid SAN Storage array uses two 800 GB SSDs for caching and traditional spinning HDDs for storing data. With expansion enclosures, the HPE MSA 2050 can support up to 192 SFF drives, providing 512 TB of disk space for fast online accessibility to any of your snapshots. For more specifics about the best practices for using the HPE MSA 2050/2052 Storage Array, refer to the [HPE MSA Best Practices Guide](#).

Although snapshots provide a level of protection, you should not consider them a formal method of backup. The best backup protection option for SAP HANA with the HPE MSA Storage arrays is the HPE StoreOnce 6600 backup appliance integrated with the HPE Catalyst Plugin for SAP. See Appendix C: Protecting SAP HANA for more information.

Networking

The networking for this Reference Architecture is designed around SAP HANA traffic isolation and networks targeted to support SAP HANA node management. The VLAN design layout also takes into consideration the option of adding HPE Serviceguard/LX for SAP (SG/LX) support to your SAP HANA servers.

Using the onboard quad 1 GbE interfaces and two dual port 10 GbE interfaces on the HPE ProLiant DL560 Gen10, the Reference Architecture follows the design pattern of network connectivity used by the HPE SAP HANA appliances. This provides redundant 10 GbE to traditional front-end and back-end networks, while providing dedicated management, heartbeat, and shared networks. Figure 8 shows the wiring diagram used for GbE networking for this Reference Architecture.

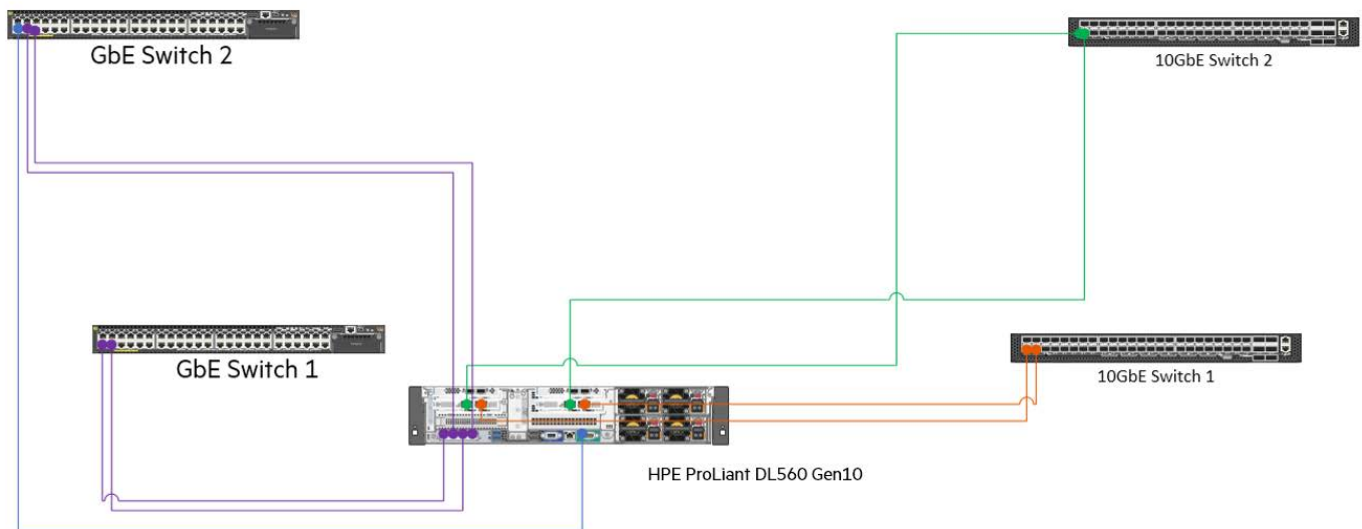


Figure 8. GbE wiring diagram for the HPE ProLiant DL560 Gen10

Table 2 outlines the networks and general addressing used for this solution. See Figure 9 below for a visual model of the VLANs and their connections.

Table 2. HPE SAP HANA networking recommendations

Network name	Purpose
Intranet	This is the private intranet that exposes the SAP application services to users.
Client network	The network used by SAP HANA clients to connect to and make queries through SQL or SAP application interfaces to access the SAP HANA analytics and database.
SG/LX network	This is the HPE Serviceguard quorum network used to monitor and initiate failover.
HANA shared network	This is a private network used for sharing information between SAP HANA internals. This network is most commonly used for NFS traffic between SAP HANA nodes in a scale-out deployment.
HANA system replication	This is the SAP HANA system replication network.
Data network	The data network is for data traffic primarily from SAP application servers.
Backup network	This is a dedicated private network used for backing up SAP HANA data and log volumes using the HPE StoreOnce Catalyst Plugin for SAP HANA. It can also be used with other file-based backup services to protect other types of system data.
Management network	This network supports connections to management interfaces for the infrastructure and server equipment. This includes the management ports on switches, arrays, server iLOs, and any managed device in the environment. This also includes a management interface on the SAP HANA servers.
Storage replication network	This is a private data network for HPE MSA NAS Storage replication between storage arrays using iSCSI. This is typically used for MSA storage replication to a remote location over iSCSI.

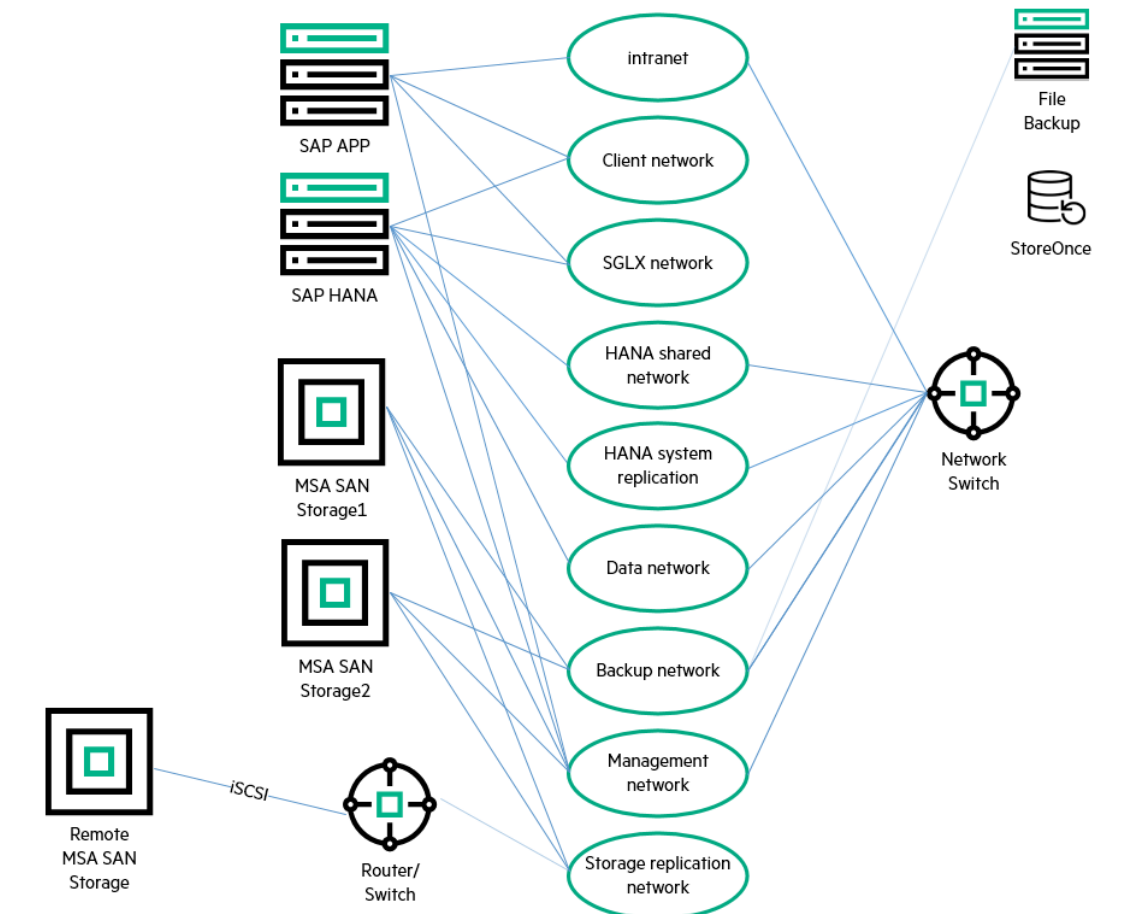


Figure 9. Networking segments for SAP HANA Reference Architectures

Storage fabric

The storage fabric design for this Reference Architecture follows the time-tested Fibre Channel dual-fabric design. This means two completely separate FC fabrics are created in the infrastructure. This highly available architecture specifies that each SAP HANA server will be configured with Multipath IO, supporting a balanced FC traffic pattern. Each storage fabric supports multiple FC peer-to-peer connections from the server to the storage arrays.

The architecture is based on two HPE SN6000B 16 Gb Fibre Channel switches. Each switch creates a single fabric. Each server is configured with two dual-port FC HBA cards. Port 0 of each HBA is connected to FC switch 1. And Port 1 of each HBA connects to FC switch 2.

The HPE MSA 2052 SAN Storage arrays are connected with an equal distribution of two or four FC ports to each FC switch. When four HPE MSA FC ports are used, ports A1 and B1 are mapped to FC switch 1. Ports A2 and B2 are mapped to FC switch 2. When eight HPE MSA FC ports are used, ports A1, A3, B1, and B3 are mapped to FC switch 1, and ports A2, A4, B2, and B4 are mapped to FC switch 2.

Figure 10 shows the recommended pattern for connecting the Fibre Channel HBAs for the storage fabric.

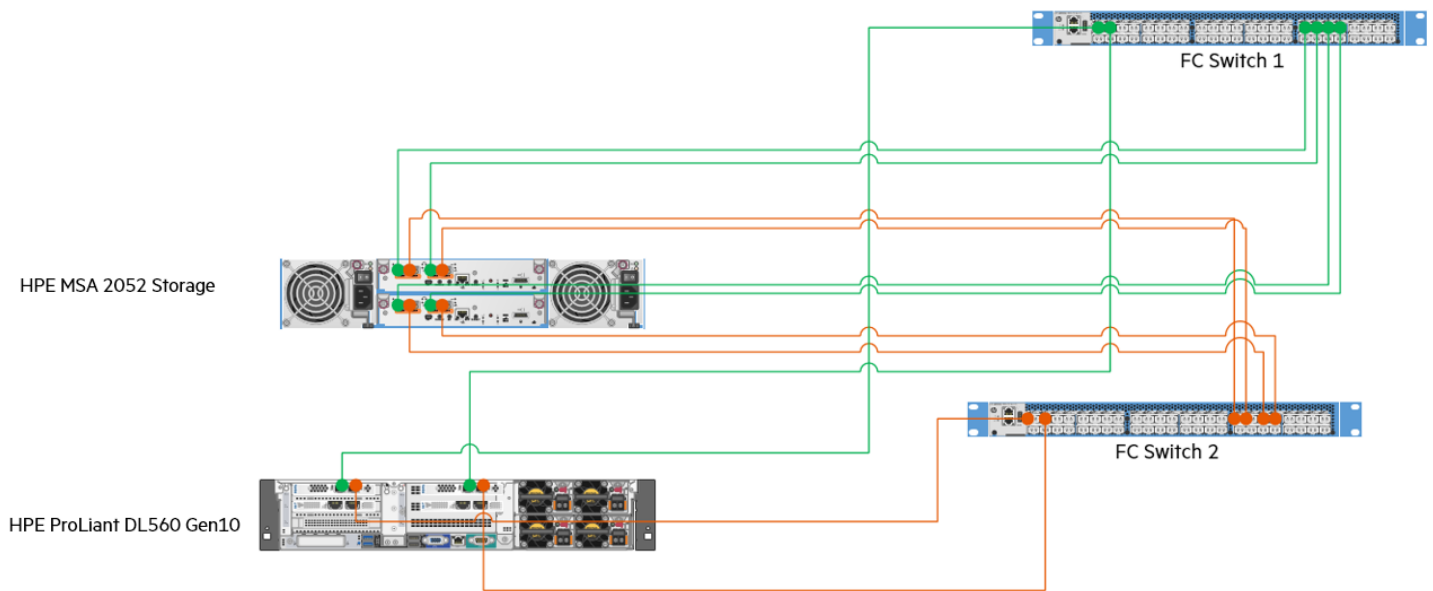


Figure 10. Eight- port server-to-storage array configuration over Fiber Channel fabrics

Note

The HPE MSA Storage array does not allow selective configuration of ports in a mix of FC and iSCSI configurations. Ports are assigned in pairs. Typically ports A1, A2, B1, and B2 are assigned to Fibre Channel. Ports A3, A4, B3, and B4 are assigned to iSCSI. A common error is to intermix iSCSI and FC FSP adapters.

Figure 11 shows the recommended pattern for connecting the server and storage array GbE connections to network switches.

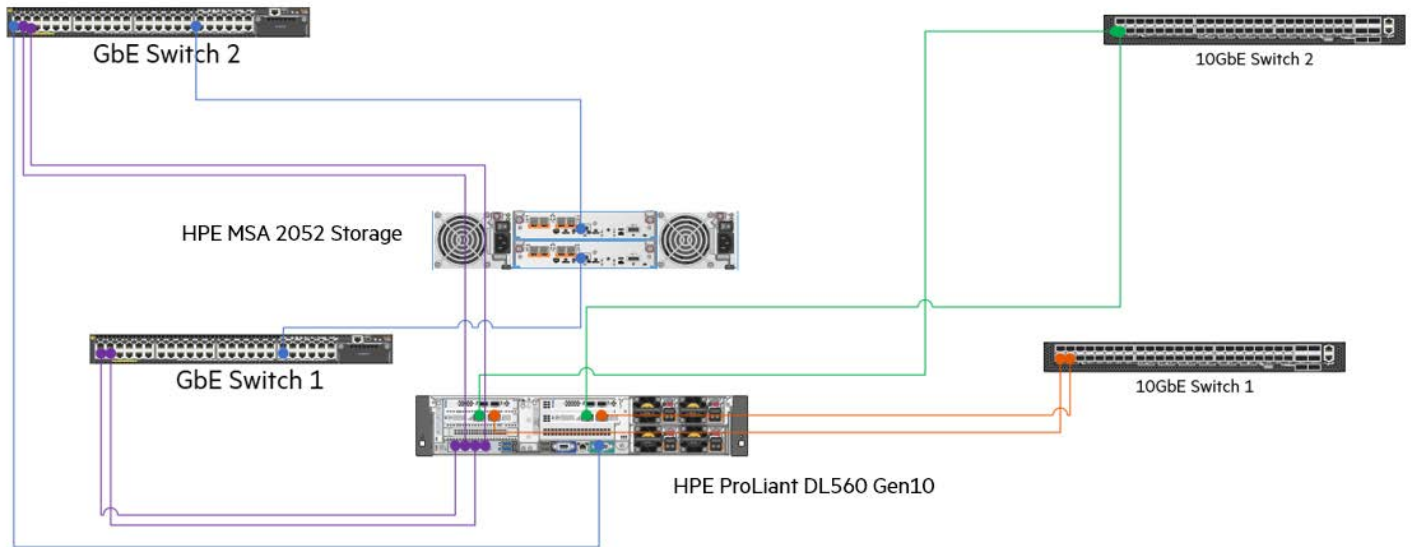


Figure 11. Server and storage array GbE switch connections

Figure 12 shows the HPE MSA Storage iSCSI connections to GbE Ethernet switches

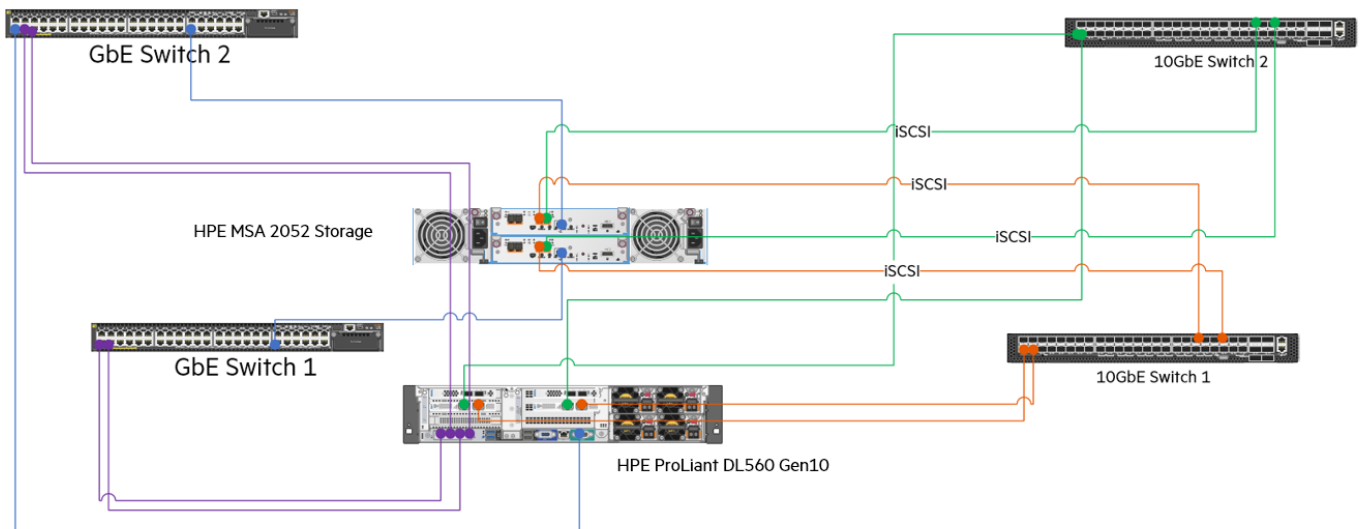


Figure 12. HPE MSA Storage iSCSI connections to GbE Ethernet switches

As a best practice, the FC fabric should be configured using peer-to-peer zoning between servers and storage arrays. Although some storage best practices recommend using Dynamic Peer Port Zoning to create secure peer-to-peer zones, HPE recommends using manually configured peer zones for production environments to minimize the potential of unwanted zone re-configuration change events.

HPE recommends configuring FC switches with their own unique domain IDs and defining each switch initially as a single fabric. This makes sure that when the switches are connected to other FC fabrics in the future, they will quickly adapt to the new fabric. Many enterprises have existing FC fabrics that use core/distribution switches and Inter-Switch Links (ISL) to extend the fabric. This centralized fabric model consolidates and simplifies zone and port management.

Management

A key design principle of this Reference Architecture is simplifying management for IT professionals. There is nothing more frustrating than having an issue with a system and not being able to fix it unless you are onsite. Therefore, included in this design, we consider support for remote management to all devices and the ability to remotely access all your equipment as an essential best practice.

All the equipment in this Reference Architecture supports secure out-of-band management capabilities—most notably, HPE award-winning Integrated Lights Out (iLO) remote management firmware. Each HPE ProLiant DL560 Gen10 server comes with an iLO 5, which provides secure remote console, software updates, alert monitoring, and power on/off capabilities to each server. The HPE MSA 2052 SAN Storage array includes HTTPS-protected management console and an SSH-supported Command Line Interface (CLI).

HPE also highly recommends you consider the option of including the HPE 2x1Ex16 KVM IP Console and the HPE TFT 7600 RKM. This rack-mounted keyboard and video monitor (KVM) provides remote-console access to supporting devices in the rack, as well as access by onsite support engineers to connected devices.




Figure 13. Network Management wiring diagram

Power

Following an HA design pattern for power is inherent in this Reference Architecture. However, even though the design takes into consideration redundant power sources, the data center power plant has to implement a fully operational redundant power infrastructure to ensure a true HA power protection plan.

Data protection

Key to any IT solution is the ability to back up and recover systems from faults and failures. Most enterprises have multiple contingency plans to protect and recover from system failures. Because the SAP HANA database creates some unique challenges for an IT organization, it is important to align this Reference Architecture with your organization's requirements for Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO). For this Reference Architecture, we used two HPE MSA 2052 arrays and MSA volume replication to create a fast recovery model to preserve data to a crash-consistent state that aligns with SAP HANA functionality. This replication in conjunction with snapshots gave us the ability to test the same data in a clone/copy pattern in a way similar to common test and development environments. For a more in-depth discussion of data protection for SAP HANA, see Appendix C: Protecting SAP HANA.

Building protective contingency plans

For this Reference Architecture, we chose to use the following protection model for a simple three-tier contingency plan to provide a crash-consistent recovery model:

1. MSA storage-based replication for fast snapshot recovery of volumes for quick point-in-time recovery.
2. Fast storage recovery with added protection using SAP Management Suite-file based recovery for data and log backups.
3. Remote hybrid storage-based replication to lower cost media.

The first of these recovery models (option 1 above) supports a quick recovery model; however, replication restrictions on the HPE MSA array limit the times between replication syncs to 30 minutes. This means, by itself, this contingency option provides a worst-case, crash-consistent RPO of 30 mins. The second option expands protection against data loss by utilizing storage replication in conjunction with SAP Management Suites' ability to back up log files. As a result, these two contingency tiers provides a protection model with an RPO potential of fifteen minutes and an RTO of less than one hour when manually executed. The final contingency tier provides for a secondary copy of our snapshots using a secondary MSA volume replication that allows the restoration of snapshots of data, log, and backup volumes.

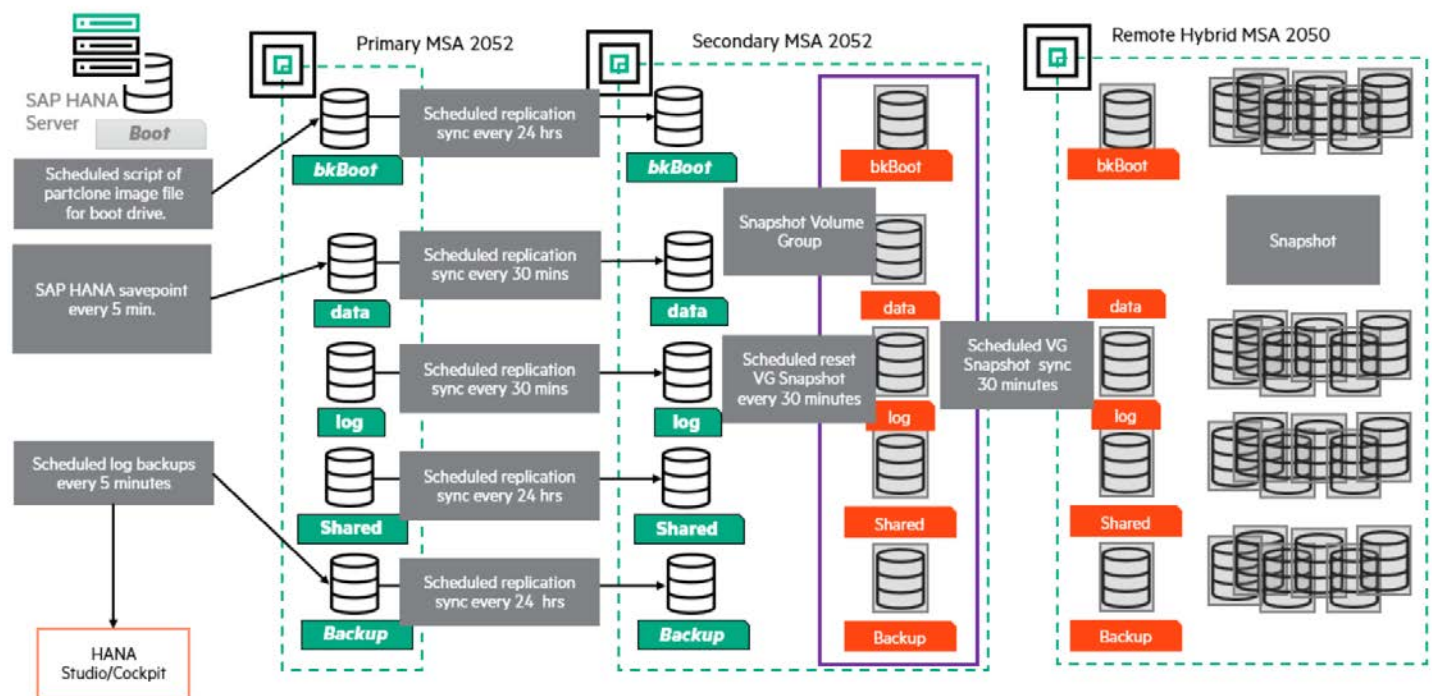


Figure 14. All three contingencies for storage-based protection

MSA storage-based replication for fast snapshot recovery

The first recovery contingency plan is based on a common sense practice. That being IT administrators would rather try to quickly mount a point-in-time snapshot to attempt a reload to a crash-consistent state, to determine if the database is viable—rather than perform a full-volume restore operation only to find the data is unusable. To set up this configuration, we used the HPE MSA 2052 volume-replication feature in conjunction with scheduled SAP HANA data savepoint operations. By using the MSA storage replication, SAP HANA recovery savepoints that occur every five minutes can be preserved on the MSA as snapshots. These snapshots can then be exported to the host and mounted quickly for the SAP HANA-required volumes.

This recovery model requires two HPE MSA 2052 Storage arrays: a primary array, where a set of virtual volumes for the SAP data, logs, and backup volumes are created, and a secondary array used to save volume replicas. On the primary array a replication set is created for the SAP HANA volumes, which is pointed to the secondary array as its Peer Connection. For this Reference Architecture test environment, a sync replication schedule of once every thirty minutes for the data volume and every thirty minutes for the log volume with an offset of fifteen minutes

from when the data volume was created. This recovery plan provided a minimum RPO of thirty minutes. This worst case scenario occurs only if a failure event happens just before the replication sync operation.

The screenshot displays the Hewlett Packard Enterprise Storage Management Utility (SMU) interface for an MSA 2050 SAN. The main window shows a list of volumes under the heading "VOLUMES". The table below represents the data shown in the interface:

Group	Name	Pool	Type	Size	Allocated
SAP7VG1	sap7data	A	base	3999GB	107GB
SAP7VG1	sap7log	A	base	5119GB	57210MB
SAP7VG2	sap7backup	B	base	71999GB	4068.4MB
SAP7VG2	sap7shared	B	base	9999GB	509GB
SAP7VG2	sap7boot	B	base	3999GB	633.3MB
-	sap8data	A	base	29999GB	801GB
-	sap8log	A	base	-	-
-	sap8usr	A	base	-	-
-	sap8backup	B	base	-	-
-	sap8_data_from_sap8_data	A	base	-	-
-	sap7backup_0949	B	snapshot	-	-
-	sap8backup_0895	B	snapshot	-	-
-	sap7backup_0951	B	snapshot	-	-
-	sap7backup_0952	B	snapshot	-	-
-	sap7backup_0953	B	snapshot	-	-
-	sap7backup_0954	B	snapshot	-	-
-	sap7backup_0955	B	snapshot	-	-
-	sap7backup_0956	B	snapshot	-	-
-	sap7backup_0945	B	snapshot	-	-
-	sap7boot_0037	B	snapshot	-	-
-	sap7backup_0946	B	snapshot	-	-
-	sap7backup_0947	B	snapshot	-	-
-	sap7backup_0948	B	snapshot	-	-
-	sap8log_0941	A	snapshot	-	-
-	sap7boot_0039	B	snapshot	3999GB	633.3MB
-	sap7boot_0040	B	snapshot	3999GB	6201MB
-	sap8log_0941	A	snapshot	5119GB	57210MB

A "Volume Information" pop-up window is open for the selected volume "sap8data", showing the following details:

- Name: sap8data
- Type: base
- Pool: A
- Group: -
- Class: Virtual
- Size: 29999GB
- Allocated Size: 801GB
- Owner: Preferred: A, Current: A
- Serial Number: 00c0ff23d62100000a9955a03c000000
- Volume Copy Job: 0
- Write Policy: write-back
- Optimization: standard
- Read-ahead Size: Adaptive
- Tier Affinity: Performance

Figure 15. HPE MSA 2052 base and snapshot volume view

Fast storage recovery with added protection

The second recovery contingency plan enhances the first plan by adding a scheduled SAP Suite Management log backup every five minutes, which is saved to a backup volume. Adding the SAP HANA log-backup protection to the recovery plan provides a five-minute incremental protection of the transaction logs.

This plan requires creating a virtual volume on the HPE MSA for file backups and exporting it to the SAP HANA server. Add this volume to a replication set on the secondary HPE MSA 2052 array in the same fashion as the data and log volumes. Next, create a schedule to backup SAP HANA log files every five minutes, using the SAP Management Suite or the SAP Cockpit management utility, and set the backup-log configuration to point to the backup volume's mount point. By setting the backup to every five minutes and the replication sync schedule to every 30 minutes, the worst scenario use case is an RPO of thirty minutes. The recovery process for this type of failure requires the following: 1) mount the latest snapshots for the data, log, and backup volumes; 2) start the database; 3) use the SAP management utilities to recover the catalog if necessary and restore any missing log transactions. Even though this plan relies on a log-file restore process, the time to mount and restore from the mounted backup volume is much less than finding the right backup set and restoring a file set from another backup source.

Another benefit of this protection model is the ability to quickly roll back your data volume to a specific point-in-time, whether this is for development purposes or to troubleshoot the database.

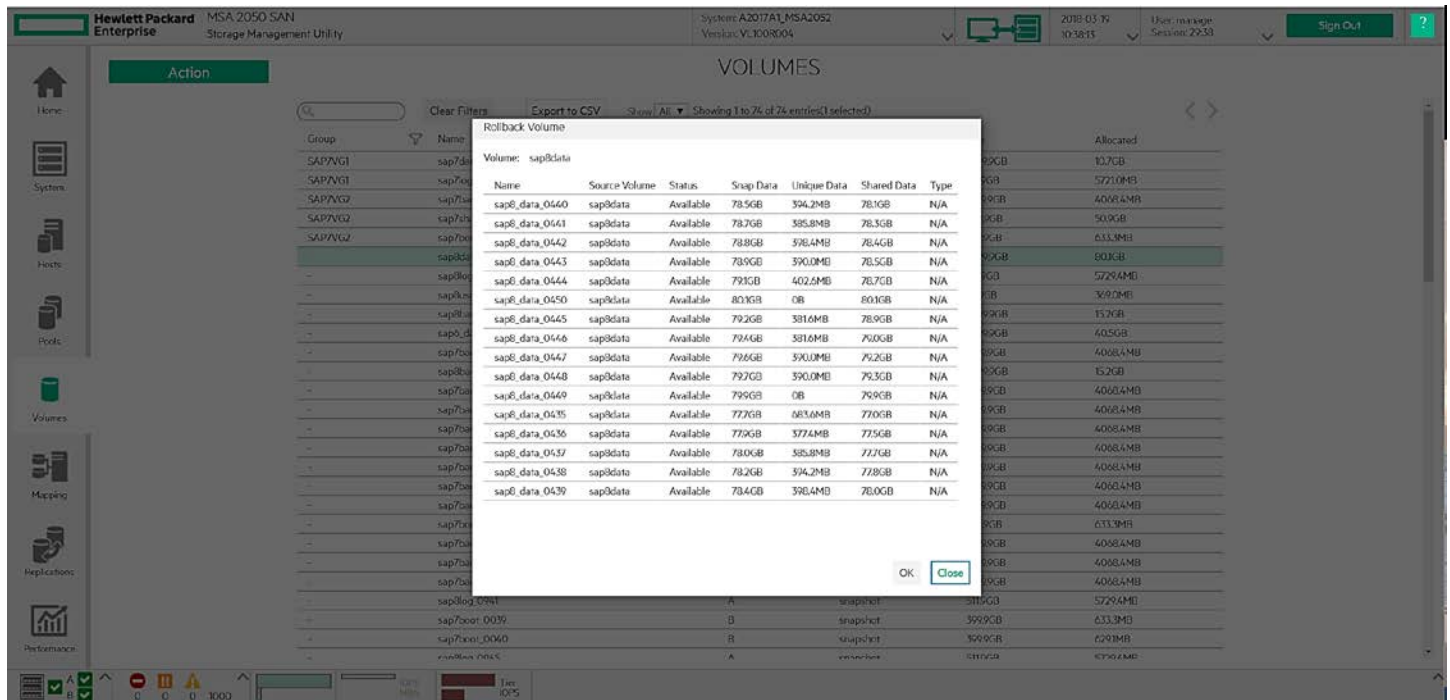


Figure 16. HPE MSA 2052 volume rollback

Remote hybrid storage-based replication

With the first two protection contingency plans, we created replication sets that saved the maximum snapshot copies of the base volumes that HPE MSA allows at the time of synchronization. This includes 16 snapshots, meaning we have eight hours of data backups for the data volume. For the log and backup volumes, we have saved eight hours of changes. This means we can mount volumes to roll back to a point-in-time within the past eight hours to attempt a recovery. The third protection contingency provides the capability to create a secondary copy of snapshots to a remote hybrid storage array, using spinning media to increase the time we can recover from mounted volumes. By implementing an HPE MSA 2050 Hybrid SAN Storage array, we set up a snapshot of each of our snapshot volumes and then scheduled a reset of each volume on a 24-hour schedule. These volumes replicated to the HPE MSA Hybrid 2050 array on a daily basis, and provided sixteen days of the latest volume snapshots. Although this secondary MSA storage-based protection model provides a simple offsite storage capability, HPE highly recommends you consider the HPE StoreOnce backup appliance with the HPE StoreOnce Catalyst Plug-in for SAP to create a longer-term storage alternative. See Appendix D: HPE StoreOnce Catalyst Plug-in for SAP HANA for details.

Solution components

Hardware

When considering server hardware for an SAP HANA database environment, you look at servers quite differently from traditional database servers. SAP HANA maintains the complete data model in-memory. All read and write requests are served from memory. The persistence of the database to storage is only for the purpose of loading the database when shutting the database down and bringing the database back up. This in-memory database model changes some of the assumptions that traditional database system design has been based upon.

HPE ProLiant DL560 Gen10 server

The HPE ProLiant DL560 Gen10 server is ideal for business-critical workloads, featuring virtualization and server consolidation, with business-processing and high-performance enhancements. Optimized for 4 Processor data-intensive applications, the HPE ProLiant DL560 Gen10 offers scalability and reliability in a 2U chassis that is built to conserve data center space.

Supporting Intel Xeon Scalable processors, the HPE ProLiant DL560 Gen10 server offers superior processing power, with up to 6 TB of memory, and I/O of up to eight PCIe 3.0 slots. Plus, the HPE ProLiant DL560 Gen10 server supplies the intelligence and simplicity of automated management through HPE OneView and HPE Integrated Lights Out 5 (iLO 5) technology.

Key features

- Scalable 4-processor performance in a dense 2U form factor:
 - Up to four Intel Xeon Platinum (8100 series) and Gold (6100 and 5100 series) processors, which provide up to 68% more processor performance¹ and 27% more cores² than the previous generation.
 - Up to 48 DIMM slots, which support up to 6 TB for 2666 MT/s DDR4 HPE SmartMemory (128 GB/6 TB DIMMs available). HPE DDR4 SmartMemory improves workload performance and power efficiency, while reducing data loss and downtime with enhanced error handling.
- Flexible new generation expandability and reliability for multiple workloads:
 - Configurable processor tray, allowing scale up from two to four processors, saving upfront costs. The drive cage design supports up to 24 SFF SAS/SATA with a maximum of 12 NVMe drives.
 - Up to eight PCIe 3.0 expansion slots for graphical processing units (GPUs), storage fabric cards, and networking cards, offering increased I/O bandwidth and expandability.
 - Up to four, 96% efficient, HPE 800W or 94% efficient 1600W Flexible Slot Power Supplies³, which enable higher power redundant configurations and flexible voltage ranges. The slots provide the capability to trade-off between 2+2 power supplies or can be used as extra PCIe slots.
- Secure and reliable:
 - HPE iLO 5 uses HPE Silicon Root of Trust technology to protect your servers from attacks, detect potential intrusions, and recover essential server firmware securely. With Runtime Firmware Validation enabled by HPE iLO Advanced Premium Security Edition, server firmware is checked every 24 hours to verify validity and credibility of essential system firmware.
 - Secure recovery allows server firmware to roll back to the last known good state or factory settings after detection of compromised code.
 - Additional security options available with Trusted Platform Module (TPM) prevent unauthorized access to the server and safely store artifacts used to authenticate the server platforms. Additionally, the HPE Gen10 Chassis Intrusion Security Kit logs and alerts whenever the server hood is removed.
- Agile infrastructure management to accelerate IT service delivery:
 - HPE OneView provides infrastructure management for automation simplicity across servers, storage, and networking.
 - For server lifecycle management, a suite of embedded and downloadable tools is available. These tools include Unified Extensible Firmware Interface (UEFI), Intelligent Provisioning; HPE iLO 5 (to monitor and manage), HPE iLO Amplifier Pack, Smart Update Manager (SUM), and HPE Service Pack for ProLiant (HPE SPP).

¹ HPE measurement: Up to 68% performance increase of Intel Xeon Platinum versus previous generation E5-4600 v4 average performance based on generational gains on HPE servers comparing 4-socket Intel Xeon Platinum 8180 to E5-4699 v4 Family processors. Any difference in system hardware, software design, or configuration may affect actual performance. May 2017

² Up to 27% performance increase of Intel Xeon Platinum versus previous generation comparing 4-socket Intel Xeon Platinum 8180 (28 cores) to E5-4669 v4 (22 cores). Calculation 28 cores / 2x cores = 1.27 = 27%. May 2017.

³ 1600W Power supplies only support high line voltage (200V AC to 240V AC)

Figure 17 shows the HPE ProLiant DL560 Gen10 server. Go to the HPE ProLiant DL560 Gen10 QuickSpecs for information at h20195.www2.hp.com/v2/GetDocument.aspx?docname=a00008181enw.



Figure 17. HPE ProLiant DL560 Gen10 server (front facing)

HPE MSA 2052 SAN Storage

The hybrid flash HPE MSA 2052 SAN Storage system with the new Gen10 ProLiant Branding is designed for affordable application acceleration that is ideal for small and remote office deployments. The HPE MSA 2052 provides the combination of simplicity, flexibility to grow now and into the future, and advanced features not expected in an entry-priced array. Start with 1.6 TB of flash capacity included and scale as needed with any combination of solid state disks (SSD) and high-performance Enterprise or lower-cost Midline SAS-based drives.

HPE MSA SAN Storage has been the industry-leading entry-storage SAN platform for the past eight years, with nearly 500,000 storage systems sold worldwide. Delivering performance in excess of 200,000 IOPS, the hybrid flash HPE MSA 2052 SAN Storage system can save up to 40% with an all-inclusive software suite and 1.6 TB of flash capacity included.

Key features

- The HPE MSA 2052 performs in excess of 200,000 IOPS for affordable application acceleration.
 - Delivers 2x IOPS performance than the previous generation, HPE MSA 2042.
 - Saves 40% of cost with all-inclusive software suite.
 - 1.6 TB of SSD capacity included.
- Advanced data services with no experience required.
 - Easy to install, easy to use, easy to maintain—no storage expertise necessary.
 - Automated tiering dynamically responds to workload changes.
- Keep your business running with expanded protection features.
 - New virtualized snapshot technology makes protection and instant recovery a snap.
 - Remote replication with FC and iSCSI supports affordable disaster recovery.
- Grow flexibly now and into the future.
 - Data-in-place upgrades protect drive investments and eliminate data migrations.
 - Start small and scale as needed with any combination of SSD and Enterprise or Midline SAS drives.

Note

The HPE MSA 2052 is the only SAP HANA TDI-certified storage array that uses Solid State Drives (SSD).

Figure 18, Figure 19, and Figure 20 show the HPE MSA2052 hardware. Go to the HPE MSA 2052 Storage QuickSpecs for information at www.hpe.com/h20195/v2/GetDocument.aspx?docname=a00008277enw.




Figure 18. HPE MSA 2052 (front facing, bezel, SFF drives)



Figure 19. HPE MSA 2052 (rear facing, dual controllers)



Figure 20. HPE MSA 2052 (front facing, drive shelf)

Fibre Channel Switches – 16 Gb SAN – SN6000B switches and SN1000Q HBAs

The HPE SN6000B Fibre Channel Switch meets the demands of hyper-scale, private cloud storage environments by delivering market-leading 16 Gbps Fibre Channel technology and capabilities that support highly virtualized environments. Designed to enable maximum flexibility and investment protection, the HPE SN6000B switch is configurable in 24, 36, or 48 ports and supports 4, 8, 10, or 16 Gbps speeds in an efficiently designed 1U package. This switch also provides a simplified deployment process and a point-and-click user interface, making it both powerful and easy to use. The HPE SN6000B switch offers low-cost access to industry-leading Storage Area Network (SAN) technology, while providing pay-as-you-grow scalability to meet the needs of an evolving storage environment.

Key features

- Delivers 16 Gbps performance with up to 48 ports in an energy-efficient, 1U form factor, providing maximum flexibility for diverse deployment and cooling strategies.
- Features Ports on Demand (PoD) capabilities for fast, easy, and cost-effective scaling from 24 to 48 ports in 12-port increments.
- Provides a flexible, simple, and easy-to-use SAN solution with industry-leading technology.
- Supports highly virtualized, private-cloud storage with multitenancy and non-stop operations.
- Offers best-in-class port density and scalability for midrange-enterprise SAN switches, along with redundant, hot-pluggable components and non-disruptive software upgrades.
- Metro cloud connectivity features, including integrated DWDM and dark fiber support (optional license).
- In-flight compression and encryption included, providing efficient link utilization and security.
- Yields exceptional price/performance value, exceeding comparable Ethernet storage-based alternatives.
- Unique bundle model: The HPE StoreFabric SN6000B 16Gb Bundled Fibre Channel switch model facilitates ease of ordering by offering the HPE SN6000B 48/24 Fibre Channel switch along with 24 16Gb optical transceivers in one package. The HPE StoreFabric SN6000B 16Gb Bundled Fibre Channel switch model is offered at a lower price than the sum of the list prices of the HPE SN6000B 48/24 Fibre Channel Switch and the 16Gb optical transceivers.

Figure 21 shows the back view of the HPE SN6000B Fibre Channel switch. For more information, go to the HPE SN6000B Fibre Channel switch QuickSpecs at www.hpe.com/h20195/v2/GetDocument.aspx?docname=c04111492.



Figure 21. HPE SN6000B Fibre Channel switch (back view)

Networking

The networking for this Reference Architecture uses two switches to follow the best practices observed and designed in the HPE ConvergedSystem 500 and HPE ConvergedSystem 900 SAP HANA appliances. Network segregation of traffic for application servers, system management, replication, backup, and clustering heartbeat networks have all been taken into consideration. We use two Aruba 8320 switches as top-of-rack switches connecting to core data center routing switches. Multiple ISL trunk ports were used to connect to an OSPF IP routing core.

In conjunction with the Aruba 8320 switches, we designed into this Reference Architecture networking configuration a lower-cost Aruba switch to support management and server iLOs connections. This separate Ethernet fabric for management purposes also provides the benefit of controlled security access and traffic isolation from production networks. The Aruba 3810M series switches are fully managed at a price point appropriate for the management requirements of your infrastructure resources.

Aruba 8320 Switch series

The Aruba 8320 Switch series, a campus core and aggregation switch, offers an innovative approach to dealing with the demands of mobile and cloud computing, while meeting requirements for application security and scalability. The Aruba 8320 is based on ArubaOS-CX, a modern software system that automates and simplifies complex network tasks, delivers enhanced fault tolerance, and facilitates zero-service disruption. Key innovations in ArubaOS-CX are its micro-services style modular architecture, REST APIs, Python scripting capabilities, and the Aruba Network Analytics Engine.

Key features

- The high-speed, fully distributed architecture provides up to 2.5 Tbps switching capacity to meet the demands of bandwidth-intensive applications and always-on networking.
- The compact 1U switch series delivers line rate 10 GbE and 40 GbE connectivity with very low latency for support of full Internet routes.
- The ArubaOS-CX automates using REST APIs and Python scripts.
- The Aruba Network Analytics Engine has the ability to monitor and troubleshoot network, system, application, and security-related issues.
- Advanced Layer 2 and Layer 3 features provide robust security and QoS.
- Deployment can be optimized with a model that provides 32 ports of 40 GbE, and models with 48 ports of 10 GbE (SFP/SFP+ or 10GBASE-T) with 6 ports of 40 GbE.
- The built-in time series database enables developers to create software modules for historical troubleshooting, as well as analysis of historical trends, and to predict and avoid future problems due to scale, security and performance bottlenecks.
- Individual software modules can be monitored and restarted independently, and upgraded for higher availability.
- Redundant power supplies and fans feature resiliency and high availability.

Figure 22 shows a front view of an Aruba 8320 switch. Go to the Aruba 8320 Switch Series at <https://www.hpe.com/us/en/product-catalog/networking/networking/pip.aruba-8320-switch-series.1010323113.html> for more information.



Figure 22. Aruba 8320 switch (front facing)

Aruba 3810 Switch Series

The Aruba 3810 Switch Series, an industry leading mobile campus access solution, provides performance and resiliency for enterprises, small-to-medium businesses, and branch office networks. With HPE Smart Rate multi-gigabit ports for high-speed 802.11ac devices, this advanced Layer 3 switch delivers a better application experience with low latency, virtualization with resilient stacking technology, and line rate 40 GbE for maximum back haul capacity. The Aruba 3810 is easy to deploy and manage with advanced security and network management tools, such as Aruba ClearPass Policy Manager, Aruba AirWave, and cloud-based Aruba Central.

Key features

- HPE Smart Rate delivers high-speed multi-gigabit capacity and full Power over Ethernet Plus (PoE+) provisioning, supporting multi-gigabit Ethernet speeds (1, 2.5, 5, and 10 gigabit Ethernet) with existing cabling.
- Dual, redundant, hot-swappable power supplies and innovative backplane stacking technology provide resiliency and scalability in a convenient 1U form factor.
- A built-in, programmable REST API interface provides configuration automation for mobile-first campus networks.
- Backplane stacking technology delivers high-performance stacking with up to 336 Gb/s of stacking throughput and increased resiliency. Up to ten Aruba 3810 switches can be stacked in a ring topology or up to five can be stacking in a mesh topology.
- Flexible traffic controls, such as policy-based routing, Quality of Service (QoS), and Access Control Lists (ACLs), manage end-to-end application priorities.
- Concurrent 802.1X, MAC and Web Authentication modes enable security and policy-driven application authentication. Advanced distributed denial of service (DDOS) protection, such as DHCP Snooping, IP Source Guard, and Address Resolution Protocol (ARP) protection, prevents unwanted traffic.

- Tunneled Node transports network traffic on a per-port basis to the Aruba controller with authentication and network policies applied and enforced at the controller. Industry standard MACsec provides switch-to-switch link level security.
- Out-of-band Ethernet management port keeps management traffic segmented from your network data traffic.⁴

Figure 23 shows a front view of an Aruba 8310 switch. Go to the Aruba 8310M Switch Series at <https://www.hpe.com/us/en/product-catalog/networking/networking-switches/pip.aruba-3810-switch-series.1008605435.html> for more information.



Figure 23. Aruba 8310 switch (front facing)

Software

A production SAP HANA solution involves many interactive services, each presenting unique challenges and requiring different resources. This section outlines the underlying software and services required for a successful SAP HANA environment.

Operating systems

As of this publication date, SAP has certified SUSE Linux Enterprise Server (SLES) 12 and Red Hat Enterprise Linux (RHEL) 7 as supported operating systems for SAP HANA 2.0. We chose to use the latest supported release of SUSE SLES 12 for SAP Applications SP3, to test and validate this Reference Architecture.

SUSE Linux Enterprise Server for SAP Applications

SUSE Linux Enterprise Server for SAP Applications was the only operating system supported by SAP when SAP HANA was first released. This Reference Architecture was built using the SUSE Enterprise foundation using the SUSE Linux Enterprise Server for SAP Applications bundle. This SUSE bundle includes the customized installers, helper wizards, and scripts that led to SUSE winning SAP's HANA Innovation Award in 2017. The specialized installation wizard, in-memory security hardening, and automated SAP HANA scale-out and scale-up installation support simplify the process of getting your SAP HANA solution up and running.

Each HPE ProLiant DL560 Gen10 in this Reference Architecture is configured with SUSE Linux Enterprise Server for SAP Applications bundles to create highly available clusters to support the SAP HANA database and core SAP applications. More information can be found at [suse.com/products/sles-for-sap](https://www.suse.com/products/sles-for-sap).

Note

If you are not using SUSE Linux Enterprise Server for SAP Applications, see SAP Note 1944799 - SAP HANA Guidelines for SLES Operating System Installation and 2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12 for specific details about SUSE OS customizations for SAP HANA.

Red Hat Enterprise Linux for SAP solutions

As of this writing, SAP HANA 2.0 SPS02 is fully supported by Red Hat Enterprise Linux 7.4. Although RHEL was not used to build this Reference Architecture, HPE fully supports a similar configuration using RHEL. See the online SAP Notes for RHEL configuration details and recommendations if your environment requires the use of Red Hat Linux. You can find these recommendations in the SAP Notes repository.

⁴ For more information, go to [Hewlett Packard Enterprise Aruba 3810M Switch Series](#) or [Aruba: a Hewlett Packard Enterprise Company: 3810 Switch Series](#).

HPE Serviceguard

Although not part of this Reference Architecture, a valuable addition to your SAP HANA configuration is the use of HPE Serviceguard LX for SAP. HPE Serviceguard is a clustering software, which provides High Availability (HA) and Disaster Recovery (DR) capabilities for applications deployed in distributed environments. HPE Serviceguard provides enterprise HA and DR support to HP-UX and Linux operating systems. The HPE Serviceguard product has customizations for SAP HANA to provide automated mechanisms to detect failures and initiate recovery procedures to ensure continuous availability of SAP application services. Additionally, it employs precise control measures to protect the integrity of application data during the execution of remedial procedures in multi-node cluster configurations.

Other HA and DR architectures based exclusively on SAP System Replication rely on system administrators to make the determination when a failure has occurred and manually initiate the failover. This dependence on human intervention can result in significant delays and can result in data loss, and even worse—corruption of data. Enterprise class HA and DR with HPE Serviceguard in conjunction with SAP System Replication provides a significant improvement over using SAP System Replication by itself. When protected with HPE Serviceguard HANA-clustered configuration, if there is a catastrophic failure of the primary (production) system, HPE Serviceguard performs an automatic, unattended failover to a secondary system. Based on the type of failure, HPE Serviceguard takes appropriate recovery actions as part of a failover operation. Because HPE Serviceguard integrates with SAP System Replication, the failover only needs to promote the already running secondary instance to a primary role. This significantly reduces downtime because there is no need to restart the SAP HANA database at the secondary system. Additionally, HPE Serviceguard fails over the IP addresses used to access the SAP HANA database and services.

When a primary system is restored after a failure, HPE Serviceguard starts it in a secondary role to restore the HA and DR environment to a fully protected state without any manual intervention.

Protection from database corruption: There are several opportunities for data corruption in an HA and DR cluster. The HPE Serviceguard cluster has an advanced data-access control mechanism that protects the data integrity for the SAP HANA environment, preventing split-brain scenarios. This condition occurs when both primary and secondary nodes are running as a primary instances at the same time, resulting in data corruption. HPE Serviceguard uses a quorum server to monitor network communication between the primary and secondary system to prevent this scenario.

The A.12.00.20 release of HPE Serviceguard software introduced a new feature called Site Aware Disaster Tolerant Architecture (SADTA), which extends disaster-tolerance capabilities to include scale-out SAP HANA systems. HPE Serviceguard software makes use of sites to manage the failover of application processes spanning multiple systems from one site to another. SADTA enables coordination between scale-out SAP HANA nodes running on separate servers at a given site. Failover is coordinated to invoke only when a sufficient number of servers have failed or a database outage has been detected.

SADTA also allows the failover of HANA partition instances within a site to provide various active-standby configuration options available to increase the protection from failures at site level.

Increased uptime in case of site failures with the unique Smart Quorum feature: The 12.00.30 release of HPE Serviceguard software introduced the Smart Quorum feature, which increases the availability of critical workloads. This feature can be deployed only in clusters configured with site-aware failover capability. The Smart Quorum provides proper coexistence with SAP's built-in HANA auto-host failover mechanism and allows the combination of SAP High Availability with HPE Serviceguard Disaster Recovery technologies. For more information on how to configure this feature, refer to [Managing HPE Serviceguard for Linux A.12.00.51](#) and [Managing HPE Serviceguard Extension for SAP for Linux Version B.06.00.80](#).

Licensing specifications

HPE MSA 2050/2052 SAN Storage system

Remote Snap requires a purchasable license in order to implement the software. On HPE MSA 2052 Storage, the ADS Software Suite includes HPE MSA Remote Snap Software; however, this software requires a license key from HPE and must be installed on an HPE MSA 2052 array to enable Remote Snap.

Storage fabric

To configure ISL Trunking between the edge Fibre Channel switches and a core director class switch, HPE highly recommends purchasing the Power-Pak licenses for your HPE SN6000B switches.

Best practices and configuration guidance for the solution

Beside those items presented in the Design section above, the following sections outline specific configurations and recommendations to help create a successful implementation of this Reference Architecture.

Server recommendations

In this section, we identify specific server configurations and recommendations for successfully implementing this Reference Architecture.

BIOS settings

HPE recommends the following BIOS settings for power management and processor options when setting up the HPE ProLiant Gen10 servers for SAP HANA. Refer to the user guide: *UEFI System Utilities User Guide for HPE ProLiant Gen10 Servers and HPE Synergy* http://h20566.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=1009955118&docLocale=en_US&docId=emr_na-a00016407en_us

Virtualization and processor options

- Virtualization Technology: Disabled
- Intel Hyper threading Options: Enabled
- Intel Turbo Boost Technology: Enabled
- Boot mode : UEFI

Power management

- HPE Power Profile: Maximum Performance
- HPE Power Regulator: HP_Dynamic_Power_Savings_Mode

Advanced power options

- Intel QPI Link Power Management: Enabled
- Minimum Processor Idle Power Core State: No C-states
- Date & time : UTC, GMT
- ASR status : Disabled

Boot volume

In this Reference Architecture, we recommend two Solid State Drives (SSD) using RAID 1 to mirror volumes for booting the operating system and providing swap space. Some SAP clients have used `boot from SAN` or `network boot` for these volumes; however, with the SAP HANA configuration, HPE recommends using onboard SSDs to eliminate any extra traffic over the network or FC fabric.

Multiple methods of protecting the boot OS can be implemented using free utilities like `dd` and `partclone` to very expensive backup products. In this Reference Architecture, we have used a very simple implementation of `partclone` to create clone image files to back up the boot OS. This simple `partclone` script creates a clone image file to a backup volume on the HPE MSA SAN Storage array, which is replicated for archival purposes. Whatever your protection method, make sure you have a valid copy to simplify the restoration of the SAP HANA server.

Storage recommendations

In this section, we identify specific storage configurations and recommendations to successfully implement this Reference Architecture.

Disk and file formats

SAP recommends a separate disk partition for data and log requirements; however, with the virtualized volumes prevalent with storage array some experts go so far as to recommend that the two partitions reside on two different storage delivery systems. For this Reference Architecture, we followed SAP's recommended separation by first creating separate virtual volumes for the data and log volumes using the Linux XFS disk format. We also created virtual volumes for the **shared**, **usr**, and **backup** directories.

During our testing, we also used the SUSE Enterprise Server for SAP Applications Installation Wizard to do a comparison using their recommended layout of disk volumes. This revealed the SUSE installation wizard-made sizing determinations based on the amount of memory in the server. The following shows Logical Volume Manager (LVM) layout for the volume group and logical volumes for our 3 TB server.

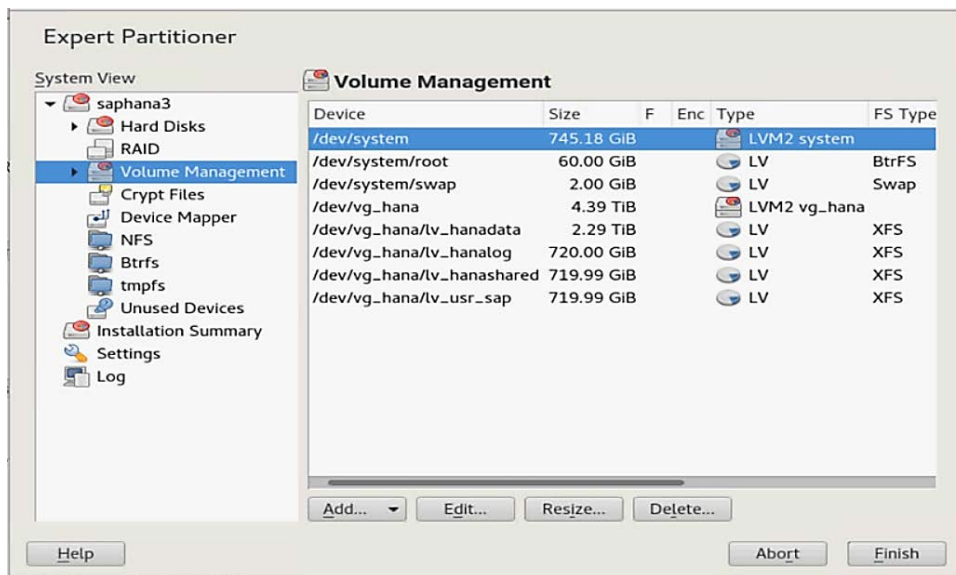


Figure 24. SLES for SAP Applications Wizard LVM pattern

Although the Linux purist might be opposed to this method of configuration for many reasons, including never using a wizard installation for a production system, we found this pattern of disk configuration supported a simplified storage design that matched well with the storage replication pattern used for this Reference Architecture.

RAID levels

This Reference Architecture follows the HPE MSA 2052 best practice of using RAID 6 for the creation of Virtual Disk Groups. Each array in the configuration included 12 SSDs assigned to disk groups for each controller pool. Each Virtual Disk Group was assigned the Tier property setting of Performance with the oversubscribing property set to enabled. Although we tested the HPE MSA in both enabled and disabled mode, there was no noticeable change in performance. Because the SAP HANA database environment is more predictable based on the memory size of the server, oversubscription can be more predictable in this environment.

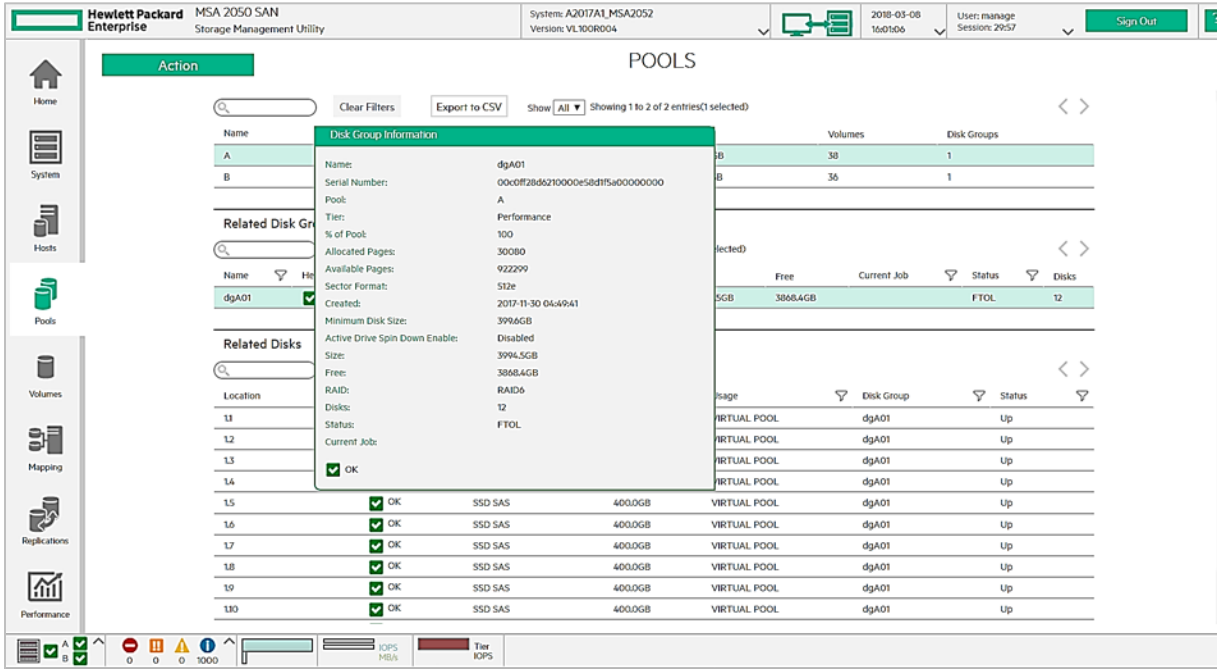


Figure 25. Storage Management Utility view of pools and virtual disk groups

MSA Storage replication

One layer of protection for this Reference Architecture is based on the use of the HPE MSA array’s volume replication feature, using peer connections between HPE MSA arrays. Part of this replication configuration includes two HPE MSA 2052 arrays setup with Peer Connections between them and Replication Sets for the data, log, shared, backup, and boot/backup volumes. The [HPE MSA Storage Array User Guide](#) recommends setting up your Peer Connection to use all local and remote ports for replication. When creating the virtual volumes for these supporting volumes, HPE recommends balancing Virtual Volumes between Controller A and B, as shown in Figure 26.




Figure 26. Storage volume balancing for a six-server pod

The HPE MSA Storage array is an active/active controller array. This means that both controllers are active at all times. Virtual volumes can be accessed through either controller; however, one controller is the primary owner for the virtual volume. During the volume-creation process, each virtual volume is assigned a priority owner controller, as shown in Figure 27.

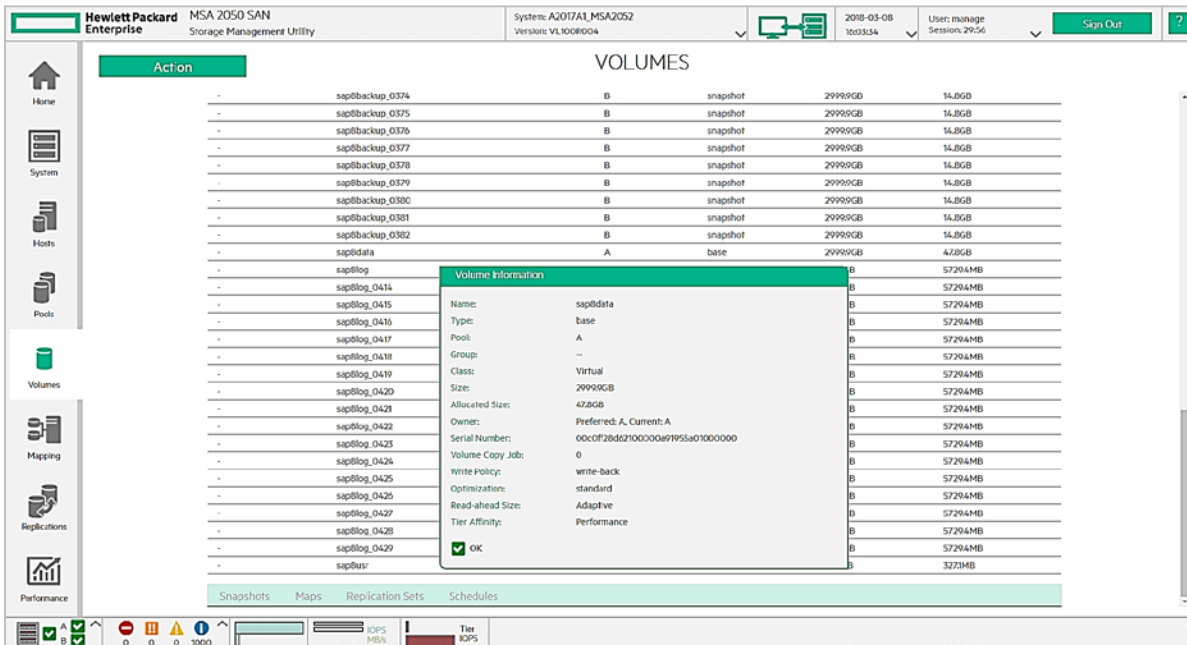


Figure 27. Storage Management Utility view of pools and virtual volumes

By default, the Storage Management Utility (SMU) tries to balance controller prioritization. During the creation of a replication set, a priority controller owner is assigned in a similar way for the controller on the peer HPE MSA 2052 array. When following the suggested six-server configuration outlined for this Reference Architecture, HPE recommends having replication owned by the B controller on both HPE MSA 2052 arrays. This can be done during the creation of the replication set, as shown in Figure 28.

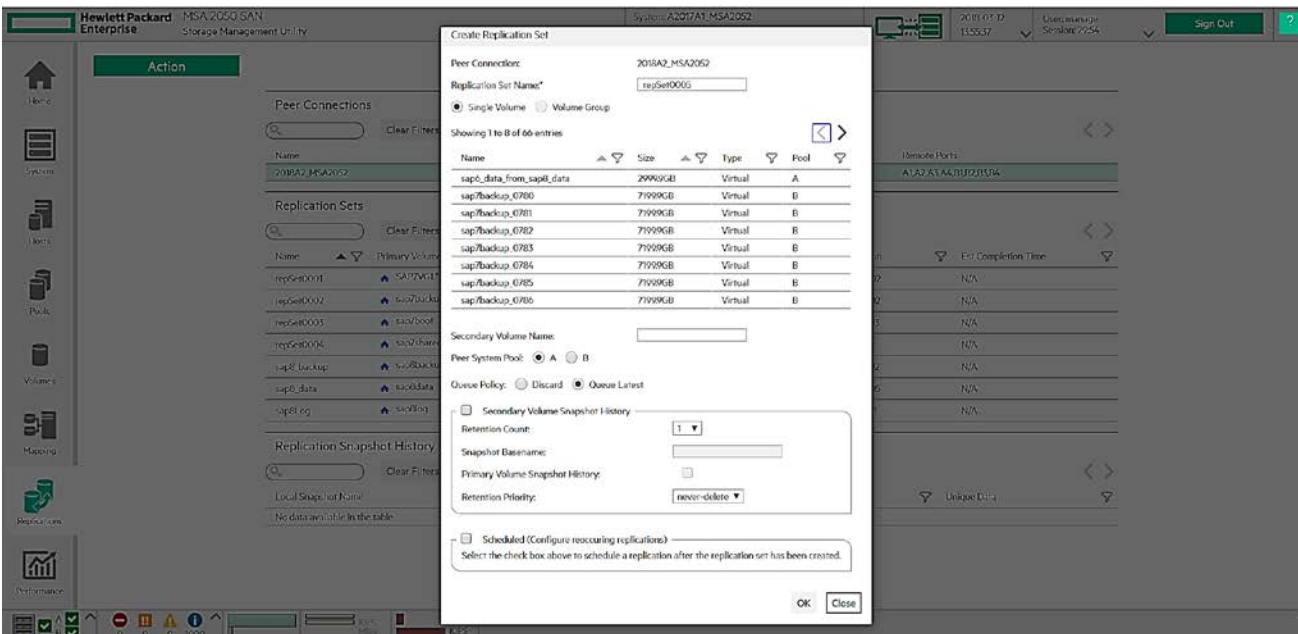


Figure 28. Storage Management Utility create replication view

NFS shared volume

By default, the installation of SAP HANA creates a `/hana/shared` directory where the configuration and executables are stored for the database instance. This shared directory becomes the key to creating Worker and Slave nodes to support Scale-Out configurations. In order to support this type of configuration, a shared disk is required to perform the proper file locking/disk fencing to ensure the data is not overwritten. Most installations use an NFS service to share this data directory.

Unlike HPE 3PAR arrays, the HPE MSA SAN Storage array does not provide a feature to expose an NFS service as part of the storage array. In order to support an SAP HANA Scale-Out configuration, you need an NFS service. Setting up, configuring, and managing these services is beyond the scope of this Reference Architecture and is left to the discretion of your IT organization.

To understand more about the HPE 3PAR File Persona services, see the following link to the HPE 3PAR product documentation: [Technical Overview of HPE 3PAR File Persona Software](#)

Networking recommendations

NIC teaming

As shown in the design section above, the networking requirements for this Reference Architecture follow the same design patterns found in the HPE Converged SAP HANA Appliances. These include two, dual 10 GbE network interfaces to provide front-end and back-end network connectivity. This is done by configuring the Aruba switches to use a compatible load-balancing mode that doesn't send the same MAC address out to each interface.

On the operation systems side, configure SLES using NIC teaming. NIC teaming is the new implementation of link aggregation or bonding.

When setting up the networking teaming configuration, we followed the best practices outlined in the SUSE Linux Enterprise Server Documentation under [Setting Up Team Devices for Network Teaming](#). This includes creating a configuration file under `/etc/sysconfig/network/` with configuration information similar to information shown on the right.

This Reference Architecture recommends two separate networks for the dual 10 GbE NICs: one for a front-end network and the other for a back-end network. Generally this dual network configuration is used with 3 Tier business applications; however, some companies are choosing to configure all four interfaces as a NIC team to support a single-network load to their SAP HANA databases.

```
STARTMODE=auto
BOOTPROTO=static
IPADDR="192.168.109.2/24"
TEAM_RUNNER=activebackup
TEAM_PORT_DEVICE_0="eth4"
TEAM_PORT_DEVICE_1="eth5"
TEAM_LW_NAME=ethtool
TEAM_LW_ETHTOOL_DELAY_UP="10"
TEAM_LW_ETHTOOL_DELAY_DOWN="10"
```

Dedicated VLANs versus VLAN tagging

For SAP HANA production configurations, HPE does not recommend using VLAN tagging for network interfaces unless your configuration is based on a virtual server environment.

Network Time Service

HPE highly recommends you use a common Network Time Protocol (NTP) source for all the systems in your SAP HANA solution. This is critical in environments where SAP's system replication services are used and where dependencies on storage-replication synchronization are required. This includes time sources for networking and storage fabrics.

SAP certification

The SAP HANA HW configuration check tool is a framework that provides tests and reports to determine if the hardware you intend to use meets the minimum requirements. Refer to the following links for details:

For *SAP HANA Hardware Configuration Check Tool for Tailored Data Center Integration*, go to the following link:
<https://help.sap.com/viewer/6b94445c94ae495c83a19646e7c3fd56/2.0.01/en-US/2f334531d3314262aa7c605f8f5f02c1.html>

For specific HPE models that have been certified, go to the following link:
<https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/appliances.html>

Capacity and sizing

This Reference Architecture has been created to support scalability and expandability. This section covers considerations for sizing SAP HANA hosts, storage persistence, and storage replication requirements.

There are three principle elements for sizing for SAP HANA: host memory, CPU performance, and storage persistence. Although SAP has several methods for capacity planning and right-sizing of SAP HANA databases, including exceptions for BWoH and BW4H, Table 3 lists the supported servers HPE has certified with SAP and their supported deployment classifications.

Table 3. SAP HANA TDI supported server information

Server	Max Memory	CPUs/Maximum Cores	Max Supported Configuration* per Node	Supported Deployment Type
HPE ProLiant DL360/DL380/BL460/SY480 Gen10 server	3 TB	2/28	2 CPUs / 3 TB	Scale Up
HPE ProLiant DL560/SY660 Gen10 server	6 TB	4/28	4 CPUs / 6 TB	Scale Up, Scale Out, BWoH/BW4H/DM/SoH/S4H
HPE ProLiant DL580 Gen9/SS680	12 TB	4/24	4 CPUs / 4 TB	Scale Up, Scale Out, BWoH/BW4H/DM/SoH/S4H
HPE Integrity MC990 X server	6 TB-48 TB	4-32/24	24 CPUs / 24 TB	Scale Up, Scale Out, BWoH/BW4H/DM/SoH/S4H
HPE Superdome X**	48 TB	2-16/24	16 CPUs / 24 TB	Scale Up, Scale Out, BWoH/BW4H/DM/SoH/S4H
HPE Superdome Flex	48 TB	4-32/28	16 CPUs / 12 TB	Scale Up, Scale Out, BWoH/BW4H/DM/SoH/S4H

* For the latest certified configurations go to <https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/appliances.html>

**Hardware partitioning is available to create multiple nodes.

The [SAP HANA Storage Requirements](#) document was last updated February 2017 and it is anticipated the general formulas will be updated as database sizes continue to increase. However, database sizing can be confusing when you are trying to find out what rules to use, especially because general rules are communicated for several different HANA products. Customers anticipating using this Reference Architecture should consult the [SAP S/4HANA Sizing Guidelines](#) website to determine their specific storage requirements.

Table 4 shows the most current SAP rules that we have used for storage sizing for SAP HANA for this Reference Architecture.

Table 4. SAP minimum storage-sizing rules

Volume or Directory	SAP HANA minimum storage preference requirements	
Data volume	1 X RAM	
Log volume	RAM < 512 GB	½ x RAM
	RAM > 1 TB	512 GB
Installation share	RAM < 512 GB	1 x RAM
	RAM > 512 GB	1 TB
usr/sap	50 GB	
Backup catalog	2 X Data Volume Size	
swap	512 GB – 1 TB	192 GB
	1 TB – 2 TB	224 GB
	2 TB – 4 TB	256 GB
	4 TB – 8 TB	288 GB
	RAM > 8 TB	320 GB

See SAP Note 1597355 – Swap-space recommendation for Linux.

Storage replication considerations

Determining the amount of disk storage necessary to support the storage replication outlined in the contingency plans of this Reference Architecture is subject to many external variables: size of volume partitions, number of volumes, number of transactions/hour, size of transactions, percent of data change, snapshot retention policies, and the server-to-storage array ratio you plan on using. Tables 5 outlines the sizing calculations used for sizing the MSA for this Reference Architecture.

Table 5. Storage array sizing for the HPE MSA 2052 based on SAP HANA replication sizing

Server-to-storage array ratio	2 to 1	3 to 1	4 to 1*	6 to 1*	8 to 1*
RAM size	3 TB	3 TB	3 TB	3 TB	3 TB
Data volume	3 TB	3 TB	3 TB	3 TB	3 TB
Log volume	512 GB	512 GB	512 GB	512 GB	512 GB
Shared volume	1 TB	1 TB	1 TB	1 TB	1 TB
Backup volume**	2 TB	2 TB	2 TB	2 TB	2 TB
Backup of Boot volume	512 GB	512 GB	512 GB	512 GB	512 GB
Storage for # of nodes	~ 14 TB	~ 21 TB	~ 28 TB	~ 42 TB	~ 56 TB*
Total Snapshot space ***	~ 32 TB	~ 48 TB	~ 64 TB	~ 96 TB*	~ 128 TB*
Estimated storage need	46 TB	69 TB	92 TB*	138 TB*	180 TB*
Volume Replica usage****	10/32	15/32	20/32	30/32	40/32****
MAX storage cap	72 TB	72 TB	72 TB	72 TB	72 TB

* Although the HPE MSA 2052 is TDI certified for up to 8 nodes, the server-to-array configurations shown in this table illustrate the array is oversubscribed based on the storage replication pattern shown in this Reference Architecture when using theoretical maximum values. It also shows the number of volumes replications would be a limiting factor if using the same replication pattern used for testing.

** Backup volume is based on SAP HANA file backups of the transaction logs.

*** Based on data used in the testing of this Reference Architecture, each snapshot on the MSA was less than 1% of the volume's actual data size. However to approximate the possible max disk usage, the value of 2% of the actual volume replication size was used for calculation purposes. This is a theoretical maximum storage size since the actual snapshot sizes depend upon the actual data saved to the volume and the rate of change of data. The total storage snapshot space used for this table is calculated by (5 volumes x 16 snapshots x (total volume sizes x 2%)) x # of nodes.

**** The MSA supports a total of 32 volume replicas. The number-to-max replica value is based on a volume replica created for the 5 volumes created for each node: data, log, shared, backup, and boot backup volumes. Note: In a scale-out configuration a single NFS shared volume is used and would not be included in this count for each node.

Workload description

Tests for this Reference Architecture solution have been generated through industry-trusted database creation tools and the SAP HANA Hardware Configuration Certification Tool. In conjunction with these tests, the SAP SHINE demo database and SAP database tools were used to create a 1 TB database.

Analysis and recommendations

This Reference Architecture has outlined a configuration of HPE ProLiant DL560 Gen10 servers and multiple HPE MSA 2050/2052 arrays to support the SAP HANA 2.0 TDI requirements. As this document has specified, this combination of 2U devices can provide a great starting point for SAP HANA test and development environments, as well as provide a dense set of servers and storage to meet a very large production SAP HANA scale-up and scale-out environment.

Although we did not show this configuration in conjunction with SAP system replication and scale-out configurations, there is nothing in this Reference Architecture that would limit you from implementing these types of configurations.

For a server configuration choice, we selected the HPE ProLiant DL560 Gen10 for the following reasons:

- TDI certified for scale-up and scale-out
- Able to scale memory to 6 TB of RAM
- Supports enough IO device slots for redundant 10 GbE NICs and 16 GB FC HBAs
- Small rack footprint (Only two units high)

If your SAP HANA application solution requires an in-memory databases larger than 6 TB, you should consider some of the other HPE SAP HANA TDI-certified servers. These include the HPE Superdome Flex, HPE Integrity Superdome X, and HPE Integrity MC990 X. See Table 3. SAP HANA TDI supported server information. More information about these servers can be found in the QuickSpecs section of [HPE Solutions for SAP HANA Tailored Data Center Integration \(TDI\)](#).

From a storage perspective, this Reference Architecture recommends the HPE MSA 2052 SAN storage system because the HPE MSA 2052 has the following features:

- TDI certified for up to 8 nodes
- Active/Active controller backplane
- All Flash SSD storage array
- Super-fast read streaming capability
- Economic cost
- Simple array management
- Small footprint (only two units high)

HPE recommends the MSA 2052 as a single application array when used for SAP HANA. If your environment requires multi-application uses from a storage array environment, HPE recommends you consider looking at the HPE 3PAR solutions.

This Reference Architecture provides a storage-replication model for protection and recovery using volume replication and snapshots. This configuration can meet RTO and RPO requirements for most businesses. Although this option does not include a deduplication capability, the recovery time is very fast. This HPE MSA storage replication feature was one of the factors that made it necessary to conservatively limit the server to storage array ratio of three servers to a single array. Although there are several ways to configure a set of servers to the arrays, including an 8:1 ratio, using the more conservative ratio found in this Reference Architecture provides a better balance of rack usage, storage space, and management functionality. Two other rack configuration options have been provided in Appendix B. Optional rack configurations.

Key Reference Architecture points

- HPE ProLiant DL560 Gen10 server
 - SAP HANA TDI certified scale-up and scale-out.
 - 6 TB of memory
 - 2U rack economy
- HPE MSA 2052 SAN Storage system
 - SAP HANA TDI-certified storage solution for up to eight nodes
 - All SSD flash array
 - Active/Active controller backplane
 - Fast read streaming for SAP HANA startup
 - 2U rack economy
- Protection and recovery
 - Storage-based replication
 - Simple snapshot-based recovery

Summary

With the SAP in-memory database, SAP HANA has changed the landscape of the database world. SAP HANA has also changed the way IT architects and engineers look at their system solutions. This Reference Architecture has covered many of the questions and concerns IT organizations have to consider when implementing an SAP HANA database environment. We have shown how the HPE ProLiant DL560 Gen10 SAP HANA TDI-certified server is an excellent choice for creating any of the SAP HANA environments from test and development to production. We have also shown how the HPE MSA 2052 SAN Storage system is a great storage option for persisting and preserving the SAP HANA database efficiently, and at an economy of price and space. We invite you to contact HPE's award winning Sales, Support, and HPE Pointnext teams to find out how we can help you solve your toughest SAP HANA challenges.

Implementing a proof-of-concept

As a matter of best practice for all deployments, HPE recommends implementing a proof-of-concept using a test environment that matches as closely as possible to your planned production environment. In this way, appropriate performance and scalability characterizations can be obtained and used to design the best solutions for your environment. For help with a proof-of-concept, contact an HPE Pointnext representative (hpe.com/us/en/services/consulting.html) or your HPE solutions partner.

Appendix A. Bill of Materials

The following Bill of Materials (BOM) contain a list of the TDI-certified equipment and software licenses used in this Reference Architecture. Electronic software license delivery is now available in most countries. HPE recommends purchasing electronic products over physical products (when available) for faster delivery and for the convenience of not tracking and managing confidential paper licenses.

Note

Part numbers are at time of publication/testing and subject to change. This BOM does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please contact your HPE Reseller or HPE Sales Representative for more information at hpe.com/us/en/services/consulting.html

Table 6. Bill of Materials

Qty	Part number	Description
HPE ProLiant DL560 Gen10 server		
1	Q7F13A	HPE DL560 SAP HANA Scale-up Compute Node
1	872766-L21	HPE DL560 Gen10 Intel Xeon-Platinum 8180 (2.5GHz/28-core/205W) FIO Processor Kit
3	872766-B21	HPE DL560 Gen10 Intel Xeon-Platinum 8180 (2.5GHz/28-core/205W) Processor Kit
48	815101-B21	HPE 64GB (1x64GB) Quad Rank x4 DDR4-2666 CAS-19-19-19 Load Reduced Smart Memory Kit
1	872235-B21	HPE DL560 Gen10 8SFF HDD Bay2 Kit
2	872376-B21	HPE 800GB SAS 12G MU SFF SC DS SSD
1	870548-B21	HPE DL Gen10 x8/x16/x8 Riser Kit
1	872253-B21	HPE DL560 Gen10 x8/x8 Tertiary Riser Kit
1	P01366-B21	HPE 96W Smart Storage Battery (up to 20 Devices) with 14.5mm Cable Kit
1	804405-B21	HPE Smart Array P408e-p SR Gen10 (8 External Lanes/4GB Cache) 12G SAS PCIe Plug-in Controller
1	869081-B21	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular LH Controller
1	873444-B21	HPE DL5x0 Gen10 12Gb SAS Expander Card Kit with Cables
1	629135-B22	HPE Ethernet 1Gb 4-port 331FLR Adapter
1	647594-B21	HPE Ethernet 1Gb 4-port 331T Adapter
2	817753-B21	HPE Ethernet 10/25Gb 2-port 640SFP28 Adapter
2	830272-B21	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
2	Q0L14A	HPE SN1200E 16Gb 2p FC HBA
1	867809-B21	HPE Gen10 2U Bezel Kit

Qty	Part number	Description
1	872261-B21	HPE DL5x0 Gen10 System Insight Display Kit
1	872222-B21	HPE DL5x0 Gen10 CPU Mezzanine Board Kit
1	Q7F23A	HPE DL560 SAP HANA Scale-up DP Stg Trk
1	Q7F22A	HPE DL560 SAP HANA Scale-up Base Cfg Trk
1	Q7F15A	HPE SAP HANA BC/TDI D3710 Storage
1	BD505A	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
1	D7S27A	HPE Smart Array SR SmartCache (Single Key/Multiple Servers) LTU
2	NOU73A	SUSE Linux Enterprise Server SAP 2 Sockets or 1-2 VM 3yr Subscription 24x7 Support Flexible LTU
1	Q7F17A	HPE CS500 SAP HANA DP SLES Trk
1	Q7F20A	HPE CS500 SAP HANA Scale-up SLES Trk
1	H1K92A3	HPE 3Y Proactive Care 24x7 SVC
1	H1K92A3 R2M	HPE iLO Advanced Non Blade - 3yr Support
1	H1K92A3 2J6	HPE SmartStorage Pack Support
1	H1K92A3 XXF	HPE DL560 HANA Scale-up Comp Nd Support
1	H1K92A3 XXH	HPE SAP HANA BC/TDI D3710 Strg Support
2	H1K92A3 Y12	HPE Package 2 for Quoting Support
2	H1K92A3 YMJ	HPE SLES SAP 1-2 VM 3yr Flx LTU Support
HPE MSA 2052 SAN Storage		
1	Q1J03A	HPE MSA 2052 SAN DC SFF Storage
2	C8R24B	HPE MSA 16Gb SW FC SFP 4pk XCVR
24	N9X91A	HPE MSA 1.6TB 12G SAS MU 2.5in SSD
1	HA114A1	HPE Installation and Start up Service
HPE Storage Fabric		
2	QR480B	HPE SN6000C 48-port 16Gb FC Switch
2	K2Q16A 05Y	2.4m Jumper (IEC320 C13/C14, M/F CEE 22)
24	C8S72A	HPE C-series 16Gb FC SW SFP+ Transceiver
24	QK735A	HPE Premier Flex LC/LC OM4 2f 15m Cbl
HPE Networking		
2	JL479A	Aruba 8320 48p 10G SFP/SFP+ and 6p 40G QSFP+ with X472 5 Fans 2 Power Supply Switch Bundle
	J9150D	Aruba 10G SFP+ LC SR 300m OM3 MMF Transceiver
	JH231A	HPE X142 40G QSFP+ MPO SR4 Transceiver
2	JL072A	HPE Aruba 3810M 48G

Table 7. Recommended options

Qty	Part number	Description
HPE ProLiantDL560 Gen10 server		
481	Q7F13A	HPE DL560 SAP HANA Scale-up Compute Node
HPE MSA 2052 SAN Storage		
24	N9X92A	HPE MSA 3.2TB 12G SAS MU 2.5in SSD
HPE Management System		
1	AZ870A	HPE TFT7600 KVM Console US Kit
1	AF621A	HPE 2x1Ex16 KVM IP Cnsl G2 VM CAC SW

Table 8. Required additional components

Qty	Part number	Description
HPE ProLiant DL560 Gen10 server		
8	C7535A	HPE Ethernet 7ft CAT5e RJ45 M/M Cable
4	QK733A	HPE Premier Flex LC/LC OM4 2f 2m Cbl
HPE MSA 2052 SAN Storage		
8	QK733A	HPE Premier Flex LC/LC OM4 2f 2m Cbl
HPE Storage Fabric		
8	C8S72A	HPE C-series 16Gb FC SW SFP+ Transceiver (for ISL up-links)
8	QK736A	HPE Premier Flex LC/LC OM4 2f 30m Cbl
HPE Networking		
8	QK734A	HPE Premier Flex LC/LC OM4 2f 15m Cbl
12	JD097C	HPE X240 10G SFP+ SFP+ 3m DAC Cable

Appendix B. Optional rack configurations

Using the basic components described in this Reference Architecture, the following rack configurations have been put together to support similar TDI configurations. These adaptations are based on the subscription ratios outlined in the capacity and sizing section. All of these configurations are within the eight node SAP HANA TDI certification of the HPE MSA 2052.

This first rack diagram demonstrates how three, four-server configurations fit in a rack providing a 2:1 server to storage array ratio.

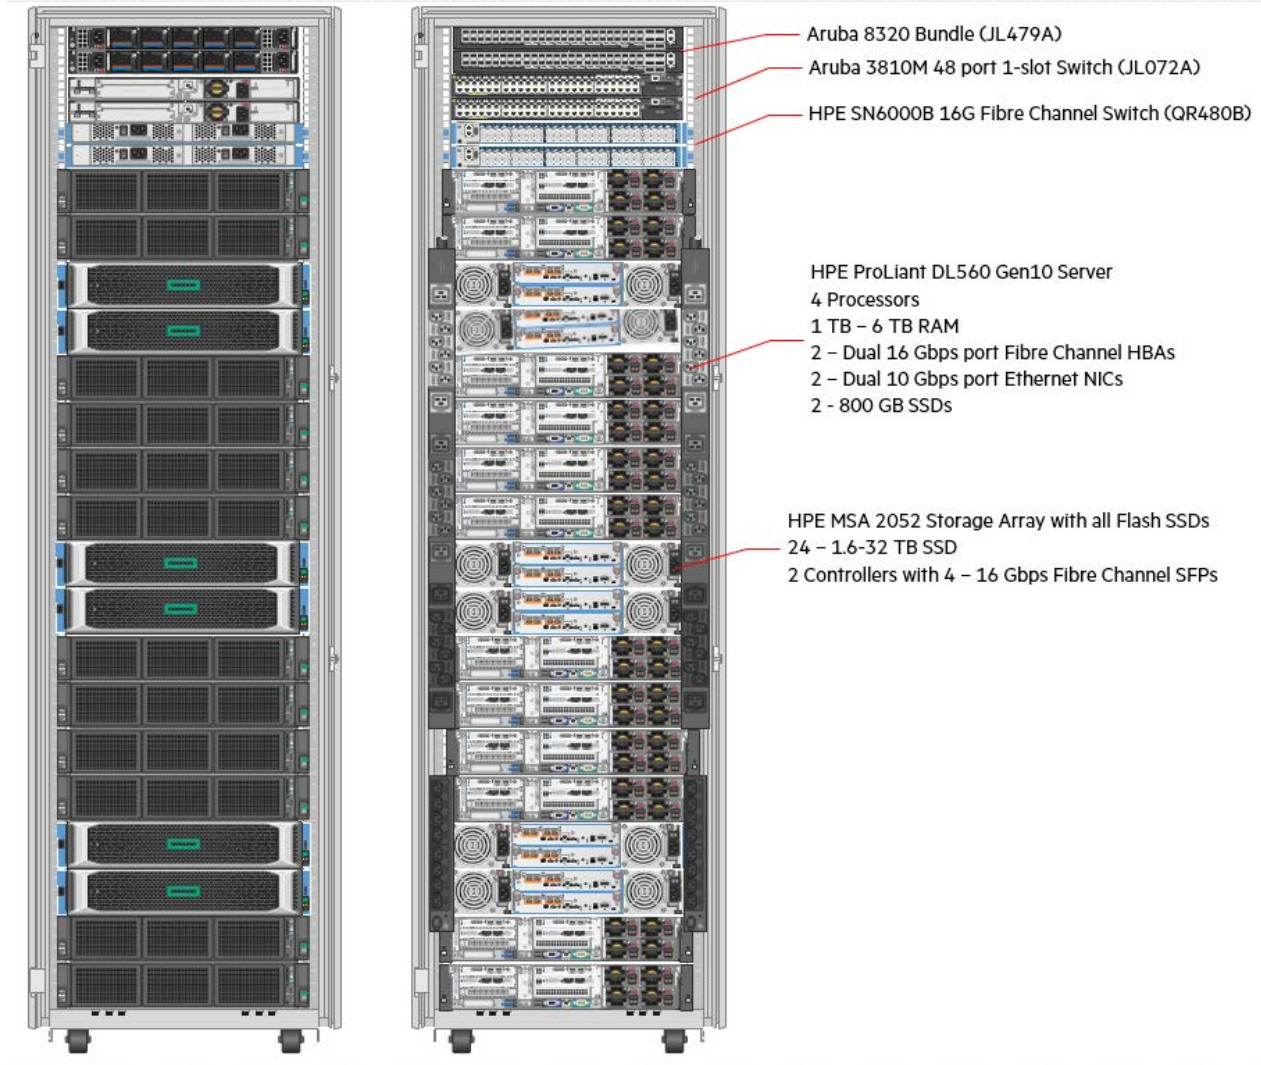


Figure 29. 2:1 server-to-storage array data center rack layout

Table 9. Bill of Materials

Qty	Part number	Description
HPE ProLiant DL560 Gen10 server		
12	Q7F13A	HPE DL560 SAP HANA Scale-up Compute Node
6	Q1J03A	HPE MSA 2052 SAN DC SFF Storage
2	QR480B	HPE SN6000C 48-port 16Gb FC Switch
2	JL479A	Aruba 8320 48p 10G Switch Bundle
2	JL072A	HPE Aruba 3810M 48G Switch

This second rack diagram demonstrates how sixteen servers provided an 8:1 server-to-storage array ratio. This configuration significantly oversubscribes the capabilities of the HPE MSA array, which would be a concern for a production environment; however it would support a test and development environment where the number of replications would not be as high.

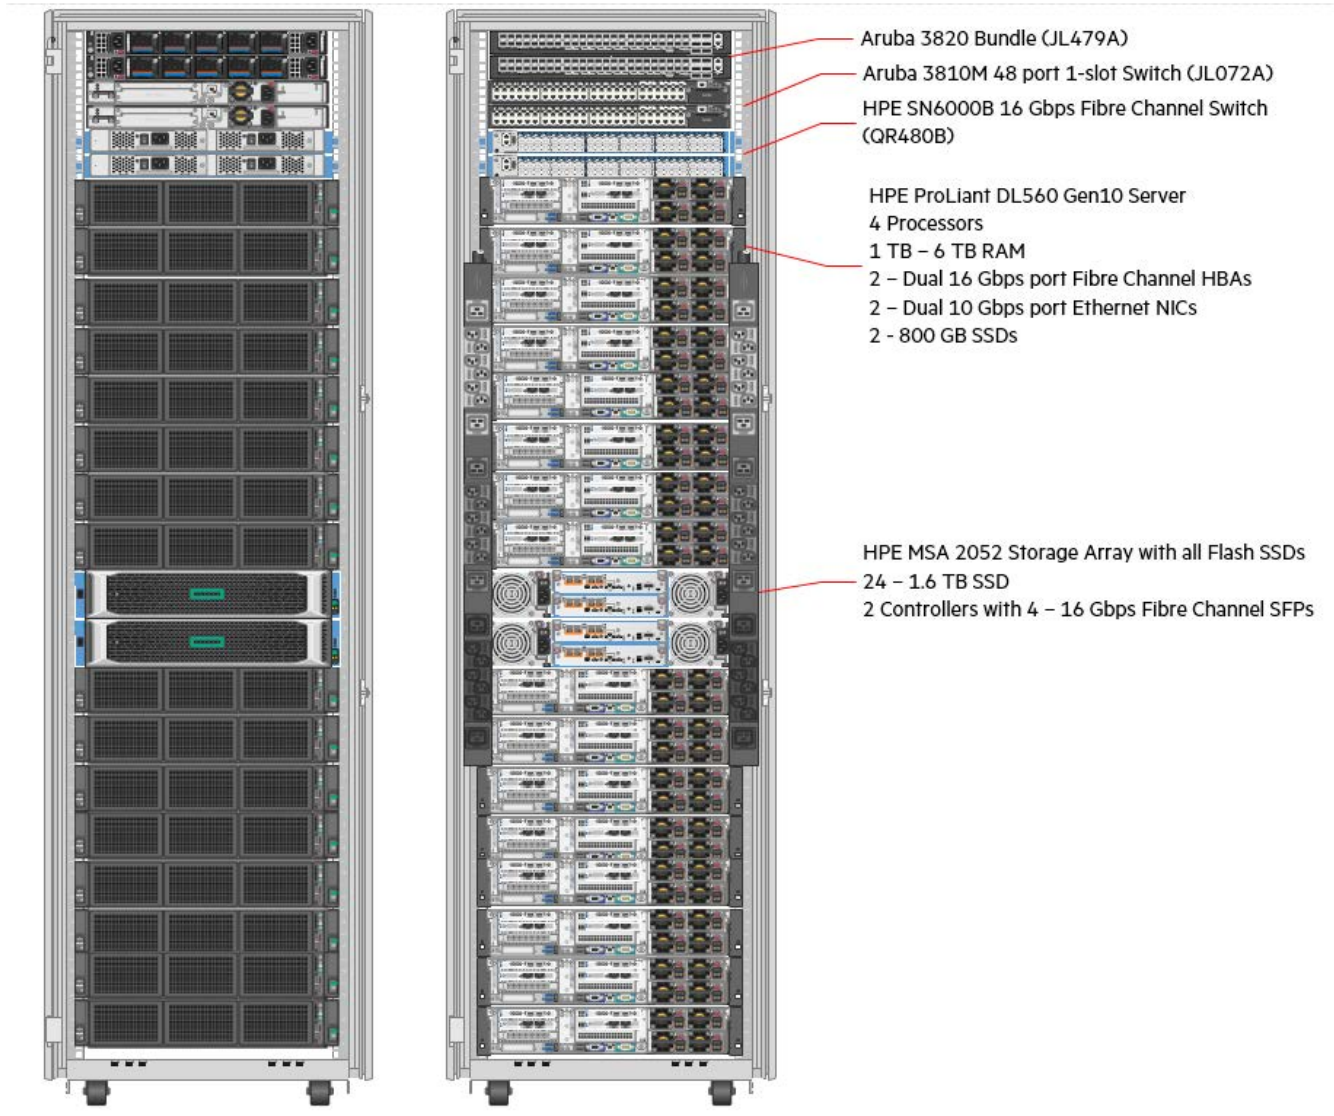


Figure 30. 8:1 server to storage array data center rack layout

Table 10. Bill of Materials

Qty	Part number	Description
HPE ProLiant DL560 Gen10 server		
16	Q7F13A	HPE DL560 SAP HANA Scale-up Compute Node
2	Q1J03A	HPE MSA 2052 SAN DC SFF Storage
2	QR480B	HPE SN6000C 48-port 16Gb FC Switch
2	JL479A	Aruba 8320 48p 10G Switch Bundle
2	JL072A	HPE Aruba 3810M 48G Switch

Appendix C. Protecting SAP HANA

Recovery Time Objective (RTO) is the duration of time in which a business service must be restored after a disruption in order to meet the required service availability agreed upon by the business community. Think of this as the maximum time your business service can be down.

Recovery Point Objective (RPO) is the maximum period of time during which data can be lost from an IT service due to an incident. You can think of this in terms of how many hours, minutes, or seconds of data you can risk losing.

These two objectives can have a significant impact upon the complexity and cost of your data protection design. Your business should weigh these objectives with your business goals to help determine if the cost outweighs the risks. For example, let us say your call center takes 10,000 orders an hour with an average order value of \$250. That creates a potential downtime cost and a recovery-point cost of \$2.5 million/hour. Based on this simple example, you can see how quickly a 10 minute outage could result in a data loss of 1,666 orders with an average cost of \$41,500. The lost opportunity costs associated with not being able to take orders or having to take orders by hand with an associated re-entry time adds to that cost. Now consider what would happen if you lost one hour of data. This would be devastating to most companies! Because of downtime cost, businesses can justify spending millions of dollars on data protection implementations.

Failures have been traditionally divided into the following six protection categories: file loss, file corruption, hardware failure, power failure, malware, and compliance, as shown in Figure 31. These six categories have been commonly mitigated through the use of primary and secondary storage appliances. However, as was pointed out in the introduction, these traditional protection methods most likely will not meet your SAP HANA protection requirements.

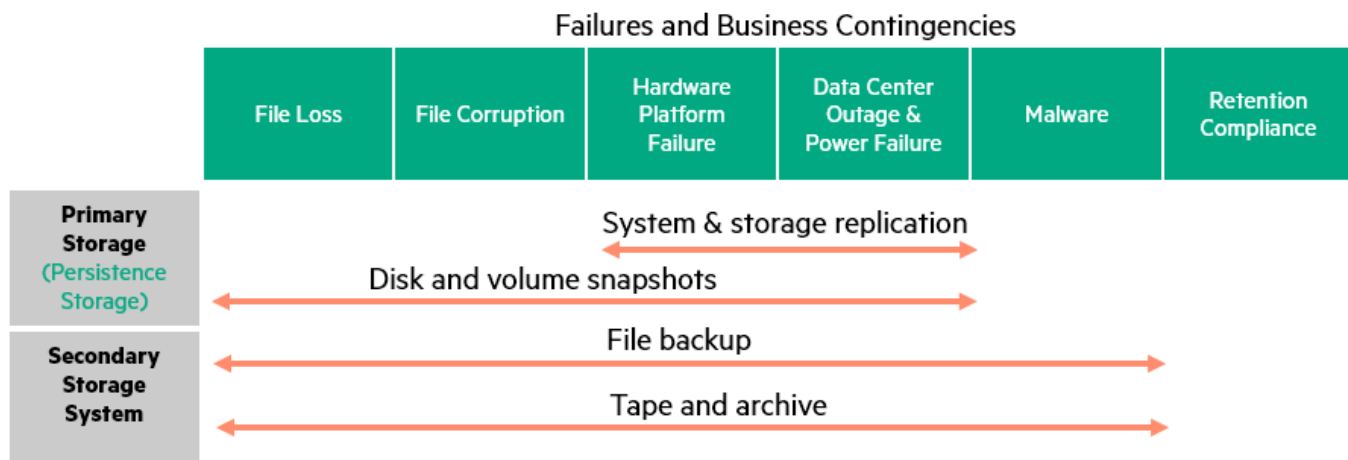


Figure 31. Data protection and failure coverage

Determining what RTO and RPO your business model requires is important to the success of any SAP HANA data protection plan. The HPE Pointnext services team can help you look at your business processes to make sure you have a solid plan to support your business objectives.

Protecting SAP HANA RPO

How much data can you lose? Every time we ask this question to customers, they almost always reply: **none!** However, our sales teams report most SAP clients implement an RPO of less than 15 minutes for production systems. Protecting the RPO of your SAP HANA data involves understanding how the database works.

As was pointed out in the Introduction, there are two key elements of the database persistence: the data and the log volumes. The data volume fulfills the persistence role of the SAP HANA database when the data in memory has been updated. Data **savepoints** occur every five minutes, unless transaction loads are high enough to force a savepoint action.

By contrast, the log volume is constantly saving in-memory transactions to protect for up-to-the-moment database changes. When the SAP HANA database services start up, the data volume is loaded first, followed by transaction log entries reloaded from the last known savepoint. Failure of either of these two volumes requires a restore process.

For this Reference Architecture, protecting the SAP HANA database is performed through the SAP management Suite (SAP Studio) or SAP Cockpit management tools. This restore process relies extensively on the backup catalog, which tracks the file based backups of the data and log volumes.

Table 11. Failure and recovery coverage for SAP HANA

Protection Methods	File loss	File corruption	Hardware failure	Power failure	Malware
SAP system replication*	No	No	Yes	Yes	No
Worker/Standby	No	No	Yes	Yes	No
Storage Replication	Yes	Yes	Yes	Yes	No
Backup/Backint**	Yes	Yes	Yes	Yes	No**

*Automation of system recovery can be accomplished with HPE Serviceguard LX for SAP

**Protection against malware can be accomplished using the HPE Catalyst Plugin for SAP in conjunction with a HPE StoreOnce system.

SAP HANA protection options

Any IT organization trying to protect a production SAP HANA environment will want a set of proven contingency plans to protect their SAP databases. Although traditional backup theory, like the 3-2-1 model referenced earlier, might satisfy requirements for small instances of SAP HANA, implementations of SAP HANA databases larger than 1 TB can make you question the feasibility of traditional recovery times. The RTO and RPO for these large databases becomes more difficult to manage because of the size of these databases.

When looking at the protection of the SAP HANA environment, we often speak in terms of protecting data in a crash-consistent or an application-consistent state. The normal operation of SAP HANA database persistence follows a crash-consistent model. When an SAP HANA node crashes, the current state of the data volume is recovered, followed by the reloading of the transaction log files from the last savepoint made to the data volume. This creates an application-consistent state of the database. A protection model based on preserving to this type of state must protect the data and log volumes.

What protection method should you use?

The best data protection method is based on a plan that meets your organizations objectives. This Reference Architecture identifies multiple methods of providing different levels and types of protection for your SAP HANA solution. HPE encourages you to talk with your channel partner, HPE pre-sales, or the HPE Pointnext team to determine the best protection method for your company.

What protection methods we chose

In this Reference Architecture, we propose protection that comes with various contingency plans. That is why we have included high availability (HA) design principles, while keeping the architecture compatible with disaster recovery (DR) configurations and implementation patterns. We have implemented many forms of protection. This includes a crash-consistent model, while limiting the scope to follow the most commonly implemented RPO and RTO objectives reported by our customers: an RPO data loss of less than 15 minutes and an RTO downtime of less than one hour. These two metrics have guided the architectural decisions used to design, implement, and test this configuration—using the best options available at the time of this publication.

What data needs to be protected in an SAP HANA solution?

When looking at the SAP HANA database solution, there are many key hardware and software components that need to be protected. In this Reference Architecture, we focus specifically on the protection of SAP HANA and how to recover the database server. Protection of other system components is beyond the scope of this paper and follows the more traditional protection practices for infrastructure devices.

The SAP HANA database server includes six different sources of data to protect. This includes SAP HANA data, logs, shared configuration files, and backups, SAP HANA user data, and the boot OS. Table 12 shows each of these data protection requirements and possible methods for protecting data.

Table 12. Protection and recovery for SAP HANA

Source	Description	Typical Location	Protection Options	Recovery
Data	The data directory is where the SAP HANA database data is persisted to disk. Every five minutes the latest changes in the database are updated to this directory. Note: An SAP HANA full data backup does not include logs and configuration settings.	/hana/data	SAP HANA data backup, SAP HANA data snapshots, Backint for SAP HANA, storage snapshots, and storage replication.	Restore from data backup, mount storage snapshot, or mount remote-replication volume. Replay of logs to restore point-in-time data.
Log	The log directory is where individual database transactions are logged as they are updated to the in-memory database. These transaction files are persisted to the log directory until logs are stored. They are preserved as long as there is storage space.	/hana/log	SAP HANA log backup, Backint for SAP HANA, storage replication.	Restore from log backup or mount remote-replication volume.
Shared	The shared directory stores the copy of the SAP HANA database server configuration and is typically part of a shared-storage configuration using NFS or some other sharing technology.	/hana/shared	Traditional file backup, storage replication.	Restore from file-based backup source or mount remote-replication volume.
Backups and catalogs	The SAP HANA management suite includes a backup utility that tracks the backup of data and log volumes. The catalog for this historical data is kept by default in the /usr/sap/backup/ directory.	/usr/sap/backup	Traditional file backup, storage replication.	Restore from file-based backup source or mount remote-replication volume.
User	User and SAP HANA user executable files for administering and managing the SAP HANA database.	/usr/sap	Traditional file backup, storage replication.	Restore from file-based backup source or mount remote-replication volume.
OS Boot	The primary boot volume for the operating system.	/	Traditional file backup, mirrored drives, partition cloning, storage replication.	Restore from file-based backup source or mount remote replication volume. Important: Although the HPE MSA SAN Storage arrays support a boot-from - SAN option, HPE does not recommend using the FC resources for the OS boot volume in conjunction with SAP HANA.

Appendix D. HPE StoreOnce Catalyst Plug-in for SAP HANA

Although not included as part of this Reference Architecture, HPE offers a fully integrated solution to protect the SAP HANA environment. The HPE StoreOnce Catalyst Plug-in for SAP provides integration to HPE StoreOnce backup appliances through the `Backint` interface. This integration is summarized in Figure 32.

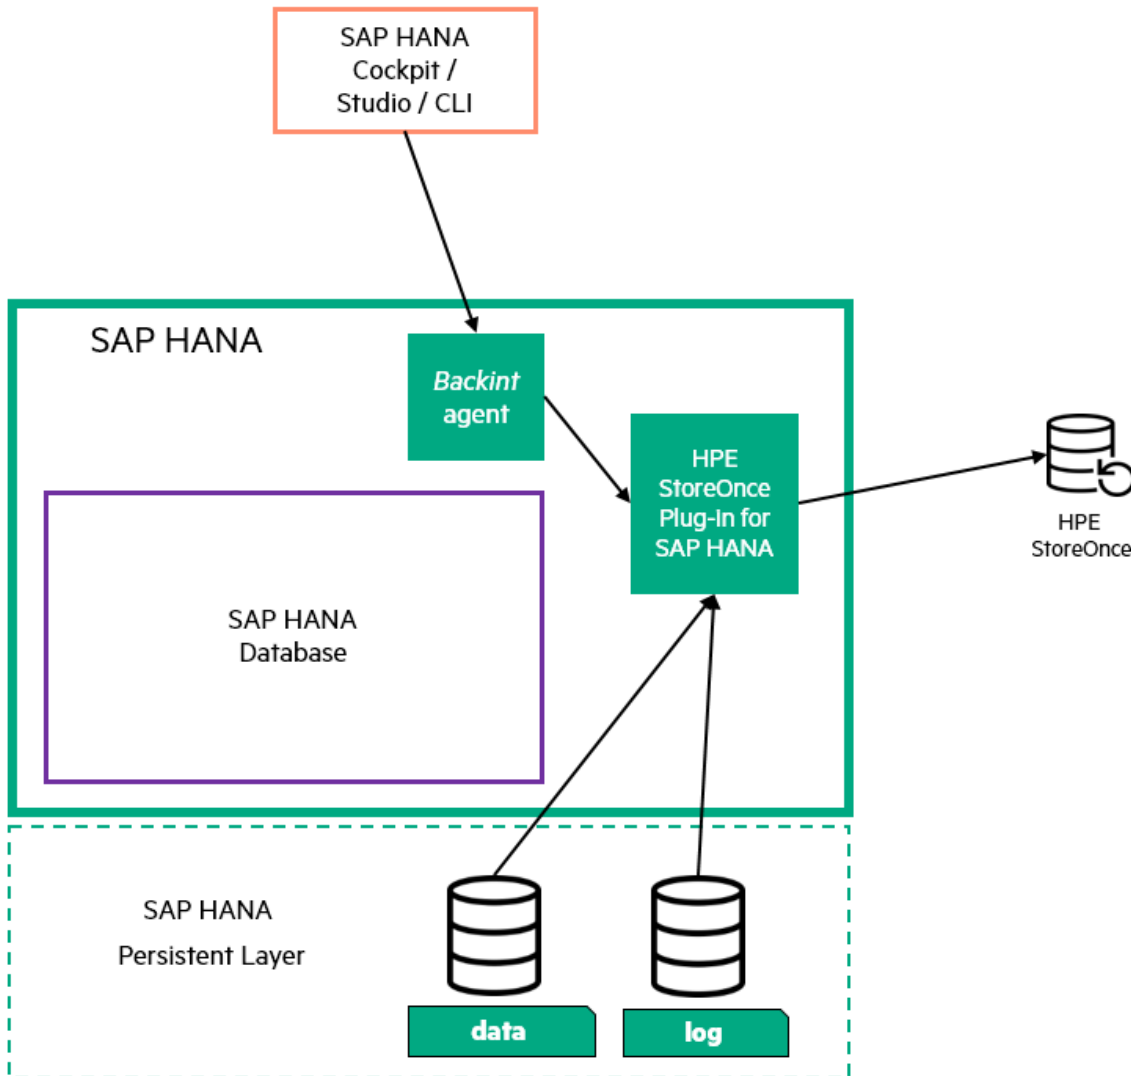


Figure 32. HPE StoreOnce Catalyst Plug-in for SAP HANA

The HPE StoreOnce Catalyst Plug-in for SAP HANA leverages the SAP Backint interface to backup and restore the SAP HANA database to an HPE StoreOnce backup appliance. Integrated with the SAP Management Cockpit and SAP HANA Studio, the HPE StoreOnce Catalyst Plug-in enables an optimized protection method for the SAP HANA database that is fully under the control of the database administrator (DBA). With the Catalyst's integration with the SAP HANA management systems, backup and restore can be performed directly to and from a Catalyst storage point on an HPE StoreOnce backup appliance. The result is flexible, high-performance protection, managed by the SAP HANA DBA. The HPE StoreOnce Catalyst Plug-in for SAP HANA can help you meet the protection requirements of your SAP HANA databases, while simplifying organization-wide, data protection processes.

The screenshot displays the SAP HANA Studio interface for configuring the Catalyst Plug-in. The title bar reads "Backup SYSTEMDB@HN1 (SYSTEM) SAP on 28.1" with a "Last Update: 1:11:05 PM" timestamp. The navigation tabs include "Overview", "Configuration", and "Backup Catalog". The main content area is titled "Backint Settings" and contains the following sections:

- Backint Settings:** A heading with a sub-instruction: "Configure the connection to a third-party backup tool by specifying a parameter file for the Backint agent." Below this, the "Backint Agent:" field is set to `/hana/shared/HN1/hpe/HPE-Catalyst-SAP-HANA-Plugin/bin/backint`.
- Data Backup:** A section with a "Backint Parameter File:" field set to `/hana/shared/HN1/hpe/HPE-Catalyst-SAP-HANA-Plugin/config/plugin.conf`. A checkbox "Use the same parameter file for data backup and log backup." is present and unchecked.
- File-Based Data Backup Settings:** A section with explanatory text: "The default destination is used unless you specify a different destination. If you specify a new destination, ensure that the directory already exists before you start a data backup. For improved data safety, we recommend that you specify an external backup destination." The "Destination:" field is set to `/usr/sap/HN1/HDB00/backup/data`. Below this, there is a checkbox "Limit Maximum File Size" which is unchecked, and a "Maximum File Size:" field with a dropdown arrow.
- Log Backup:** A section with a "Backint Parameter File:" field set to `/hana/shared/HN1/hpe/HPE-Catalyst-SAP-HANA-Plugin/config/plugin_log.conf`.
- Log Backup Settings:** A section with "Destination Type:" radio buttons for "File" and "Backint", with "Backint" selected. The "Destination:" field is set to `/usr/sap/HN1/SYS/global/hdb/backint/SYSTEMDB`. The "Backup Interval:" is set to "15" with a "Minutes" dropdown. A warning icon and text state: "If you disable automatic log backup, the log area will continue to fill. A full log area will cause the database to hang." A checkbox "Enable Automatic Log Backup" is checked.

Figure 33. SAP HANA Studio Catalyst Plugin integration view

The HPE StoreOnce Plug-in for SAP HANA must be installed using the SAP HANA operating system user. The plug-in has guided installation and configuration for fast setup. This plug-in has a zero-cost use license and can be downloaded from the [HPE StoreOnce Catalyst Plug-in for SAP HANA](#) web site for free. However, an HPE StoreOnce Catalyst license must be purchased and installed for each appliance that hosts the targeted HPE StoreOnce stores. For more information, go to [HPE Reference Architecture 2.0 for SAP HANA backup and recovery using the StoreOnce Catalyst Plug-in for SAP HANA 1.0.1](#) or the [HPE StoreOnce Catalyst Plug-in for SAP HANA](#) website.

Resources and additional links

HPE Reference Architectures, hpe.com/info/ra

HPE Servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE Technology Consulting Services, hpe.com/us/en/services/consulting.html

HPE MSA Remote Snap Software, <https://h20195.www2.hpe.com/v2/Getdocument.aspx?docname=4AA1-0977ENW>

HPE StoreOnce Catalyst Plug-in for SAP HANA, <https://h20392.www2.hpe.com/portal/swdepot/displayProductInfo.do?productNumber=StoreOnceForSAPHANA>

SAP HANA Storage Requirements, <https://assets.cdn.sap.com/sapcom/docs/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.pdf>

SAP S/4HANA Sizing Guidelines, https://help.sap.com/viewer/p/SAP_S4H_SIZING

SAP on SUSE, <https://wiki.scn.sap.com/wiki/display/ATopics/SAP+on+SUSE>

SUSE Linux Enterprise Server for SAP Applications 12, https://www.suse.com/documentation/sles-for-sap-12/singlehtml/book_s4s/book_s4s.html

SAP HANA Administration Guide, <https://archive.sap.com/documents/docs/DOC-60368>

SAP Notes

1736976 - Sizing Report for BW on HANA

2296290 - New Sizing Report for SAP BW/4HANA

2235581 - SAP HANA: Supported Operating Systems

1944799 - SAP HANA Guidelines for SLES Operating System Installation

2205917 - SAP HANA DB: Recommended OS settings for SLES 12 / SLES for SAP Applications 12

Blogs

The Ultimate Guide to Effective Sizing of SAP HANA, <https://blogs.sap.com/2017/03/07/the-ultimate-guide-to-effective-sizing-of-sap-hana/>

To help us improve our documents, please provide feedback at hpe.com/contact/feedback

Learn more at

hpe.com/storage



Sign up for updates



© Copyright 2018 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 registered trademarks of SAP AG in Germany and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Oracle is a registered trademark of Oracle and/or its affiliates. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries.

a00047300enw, June 2018