



HPE PRIMERA ALL-FLASH FOR SAP HANA ENTERPRISE STORAGE

Using HPE Primera All-Flash for SAP HANA
Tailored Datacenter Integration configurations

CONTENTS

Introduction.....	3
Scope.....	3
Audience.....	3
HPE Primera All-Flash Storage overview.....	3
Enterprise storage for SAP HANA TDI.....	4
Scalability and sizing.....	5
Configuration guidance.....	5
Setup and installation requirements.....	5
SAP HANA Fibre Channel storage API connector.....	6
Further server requirements.....	7
Installation and configuration.....	8
HPE Primera Storage Virtual Volume definition.....	8
Quality of Service.....	10
HPE data protection solutions for SAP HANA.....	10
Summary.....	12



INTRODUCTION

The era of intelligence has arrived. Enterprises want data handling and processing to be as seamless and intuitive as possible. SAP® is driving its focus on the same lines towards what it calls the “Intelligent Enterprise.” SAP aims to embed machine learning algorithms and artificial intelligence technologies in SAP applications so that enterprises can continue to learn and adapt as the data landscape changes, providing insights and patterns to make the right decisions at the right time.

To achieve this advancement, enterprises need to make sure that their underlying infrastructure that runs SAP HANA® complements the capabilities of SAP applications. The underlying infrastructure should support the seamless identification, analysis, transactions, prioritization, and recommendations provided by the applications in the most efficient and intelligent way possible.

HPE Primera is the ideal storage solution for the demanding needs of SAP HANA. HPE Primera storage solutions have the architecture to deliver exceptional performance, extreme availability, and data services for mission-critical HANA environments.

SAP HANA requires storage to tackle a wide range of I/O so arrays can achieve the desired performance levels and scale with efficiency. Most storage offerings can accomplish this task if loaded up with a large number of drives. The question then becomes how efficiently an array can deliver the necessary performance and scalability.

HPE Primera introduces support for Non-Volatile Memory Express (NVMe) NAND SSDs. This technology ensures improved performance and lower latencies for mission-critical workloads such as SAP HANA. SAP HANA Tailored Data Center Integration (TDI) solutions from Hewlett Packard Enterprise (HPE) include compute blocks and certified storage. HPE Compute Blocks for SAP HANA TDI are available for several Intel® Xeon® architectures and platforms. Certified SAP HANA TDI storage from HPE includes HPE Primera, HPE Nimble Storage, HPE 3PAR StoreServ, HPE MSA, and HPE XP7 Storage platforms.

Scope

This white paper explains how to integrate HPE Primera All-Flash Storage systems with SAP HANA TDI using SAP specifications for mission-critical applications. It describes key solution components, implementation details, and configuration recommendations needed for SAP HANA key performance indicators (KPIs). Information included is in accordance with the SAP HANA-HWC-ES-1.1 certification scenario. For more information on the SAP HANA Hardware Configuration Check Tool, refer to SAP Note [1943937](#).

Audience

The intended audience of this technical white paper is IT professionals seeking to design and implement enterprise storage for an SAP HANA TDI environment. Readers should have a functional understanding of SAP HANA concepts and technologies.

HPE PRIMERA ALL-FLASH STORAGE OVERVIEW

The HPE Primera offers Tier 0 flash arrays for massive consolidation of most demanding workloads, including SAP HANA. This product family can handle up to 1.5 million IOPS at 44 GB/s sub-millisecond latencies and offers more than 4 PB of usable capacity.

NVMe is an industry-standard specification and interface for accessing and communicating with non-volatile storage media attached via a PCIe bus. Compared to SCSI, NVMe offers better utilization of hardware and software components to achieve higher levels of concurrency in I/O and has a significantly optimized and streamlined command set.

NVMe offers a large number of deep queues, specifically up to 64K queue pairs with queue depth of 64K. The media behind both NVMe and SAS interfaces is the same NAND drive technology. However, the difference is that SAS disks connect to a SAS controller to serve I/O whereas NVMe disks attach directly to the PCIe bus. This configuration results in lower latencies while using the NVMe interface.

HPE Primera is a distributed system, which means it can provision NVMe as well as SAS SSDs. Figure 1 shows a sample layout of the drives in a four-node Primera system.



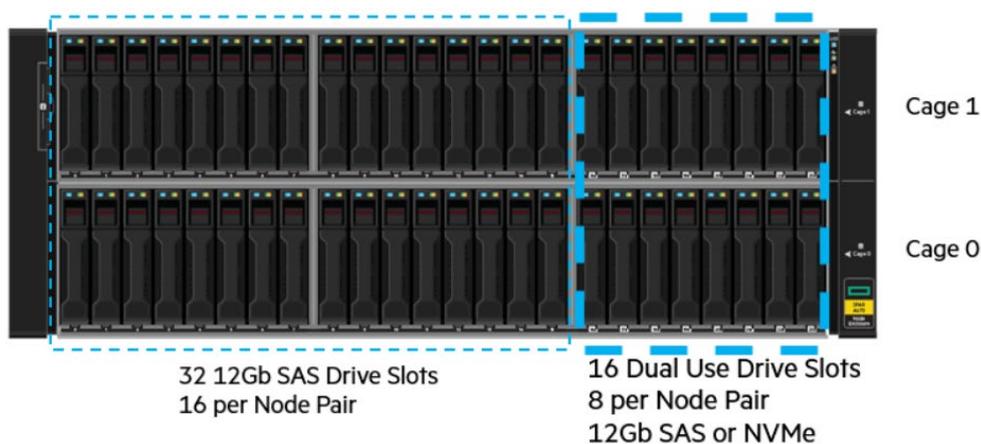


FIGURE 1. NVMe and SAS drive layout in HPE Primera

Each node pair has 16 slots from 0 to 15 for SAS SSDs and eight slots from 16 to 23 for NVMe SSDs. You can configure the HPE Primera system with SAS SSDs, NVMe SSDs, or both depending on your requirements. However, if you want to use NVMe SSDs together with SAS SSDs, there must be eight NVMe SSDs and a minimum of 16 SAS SSDs per node pair. SAS expansion shelves can be added as needed. SAS and NVMe SSDs ship in supported capacities of 1.92 TB, 3.84 TB, 7.68 TB and 15.36 TB. All SSDs drives are fully hot-plug compliant and offer options including non-encrypting, FIPS encrypting, and FIPS encrypting drives with TAA compliance.

HPE Primera All-Flash arrays are designed for 100% availability with full hardware redundancy. They support availability objectives for SAP HANA environments.¹ Enhanced storage capabilities provide continuous data access and fine-grained Quality of Service (QoS) controls to ensure predictable service levels for all HANA and non-HANA applications without physical partitioning of resources. HPE Primera also greatly improves the speed and robustness of data protection with application-aware and storage-integrated backup and recovery software.



FIGURE 2. HPE Primera flash-optimized portfolio

Enterprise storage for SAP HANA TDI

In an SAP HANA TDI environment, the storage requirements specified in the [SAP HANA Storage Requirements](#) white paper must be fulfilled. All internal disks can be removed from the configuration because the log and data volumes reside on the enterprise storage array. Appropriate host bus adapters (HBAs) are required within the server system to provide I/O access to the HPE Primera All-Flash Storage array.

NOTE

For more information about the validated hardware platforms, refer to the [Certified SAP HANA Hardware Directory](#).

¹ hpe.com/content/dam/hpe/download/pdf/gated/4AA5-2846ENW.pdf



Scalability and sizing

You should base HPE Primera capacity sizing for SAP HANA TDI on the sizing rules given in the [SAP HANA Storage Requirements](#) white paper from SAP. You should apply these sizing rules based on the memory needed for the SAP HANA database.

SAP HANA host scaling for the various HPE Primera models based on SAS SSDs is described in Table 1.

TABLE 1. Scalability of HPE Primera All-Flash arrays for SAP HANA TDI configurations

Primera model	Maximum number of SAP HANA hosts
HPE Primera 630	20
HPE Primera 650 (two-node)	30
HPE Primera 670 (two-node)	36
HPE Primera 650 (four-node)	60
HPE Primera 670 (four-node)	72

Note that information in Table 1 is for SAS SSD configurations based on current SAP HANA KPIs with HWC-ES-1.1. With NVMe SSDs alone, a maximum of 16 SAP HANA nodes can be supported on eight NVMe SSDs in a two-node Primera system and 32 SAP HANA nodes on 16 NVMe SSDs in a four-node Primera system.

For further configuration and sizing information, contact your HPE sales or presales representative.

CONFIGURATION GUIDANCE

You can configure HPE Primera to support a scale-up as well as a scale-out SAP HANA setup. It supports both bare-metal and virtualized SAP HANA configurations. Storage sizing does not differ when running a bare-metal or virtualized HANA instance.

Setup and installation requirements

The configuration shown in Figure 3 outlines the certification configuration with one HPE Primera all-flash array and three HPE ProLiant DL servers running SUSE Linux® Enterprise Server (SLES) 12 SP4. High availability on a Fibre Channel level is granted by separated fabrics using two Brocade SAN switches. The HPE Primera virtual volumes are thin-provisioned in a RAID 6 configuration.



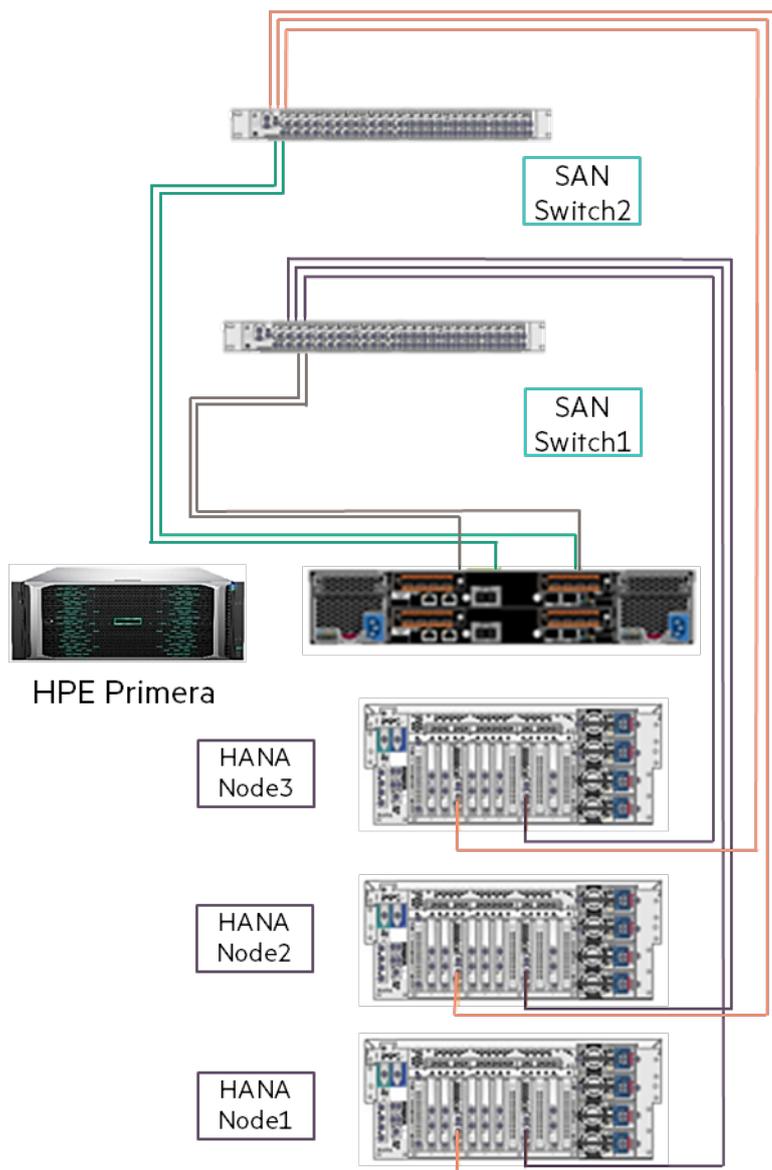


FIGURE 3. System setup for the HPE Primera enterprise storage certification

SAP HANA Fibre Channel storage API connector

The SAP HANA implementation with HPE Primera All-Flash arrays uses the SAP HANA Fibre Channel storage API connector to access the SAP HANA persistence layer. This SAP HANA built-in Fibre Channel client provides direct Fibre Channel access to the SAP HANA data and log files for the SAP HANA nodes and enables high throughput and low-latency access to the SAP HANA database.

For scale-out scenarios, a highly available shared NFS service must be available for the SAP HANA configuration, log, and trace files to be stored.

Built-in high availability

The SAP HANA Fibre Channel client used with the HPE Primera Storage All-Flash system configuration provides a highly available SAP HANA deployment. If one SAP HANA node fails, the defined standby node requests access to the data and log devices of the failed node. The standby node automatically recovers the SAP HANA persistence of the failed node to enable continued SAP HANA operations.



Further server requirements

Along with the required hardware configuration, some operating system layer settings must be configured to use HPE Primera efficiently.

Multipath implementation

To access a block device from an SAP HANA physical Linux server, you need to install and configure multipathing capabilities. You need to edit the multipathing configuration in the `/etc/multipath.conf` file as follows:

```
# The defaults section
defaults {
    polling_interval 10
}
# For Primera device
devices {
    device {
        vendor          "3PARdata"
        product         "VV"
        no_path_retry   0
        path_grouping_policy group_by_prio
        prio            "alua"
        hardware_handler "1 alua"
        path_selector   "service-time 0"
        rr_weight       uniform
        rr_min_io_rq    1
        path_checker    tur
        failback        immediate
    }
}
```

Udev tuning

To set the optimal Linux I/O scheduler parameters for HPE Primera, you need to create a new udev rule. On SLES 11 SP4 or Red Hat® Enterprise Linux (RHEL) 7, create a file named `/etc/udev/rules.d/10-Primera.rules` with the following entry:

```
ACTION=="add|change", KERNEL=="dm-*", \
PROGRAM="/bin/sh -c 'grep 3PARdata /sys/block/$name/slaves/*/device/vendor'", \
ATTR{queue/nr_requests}="4096", ATTR{queue/scheduler}="noop", \
ATTR{queue/max_sectors_kb}="4096", ATTR{queue/add_random}="0", \
ATTR{queue/rq_affinity}="2"
```

On SLES 15, SLES 11 SP5, or RHEL 8, create a file named `/etc/udev/rules.d/10-Primera.rules` with the following entry:

```
ACTION=="add|change", KERNEL=="dm-*", \
PROGRAM="/bin/sh -c 'grep 3PARdata /sys/block/$name/slaves/*/device/vendor'", \
ATTR{queue/nr_requests}="4096", ATTR{queue/scheduler}="none", \
ATTR{queue/max_sectors_kb}="4096", ATTR{queue/add_random}="0", \
ATTR{queue/rq_affinity}="2"
```



INSTALLATION AND CONFIGURATION

You must complete the following initial HPE Primera All-Flash storage configuration tasks before installing SAP HANA enterprise storage TDI.

HPE Primera Storage Virtual Volume definition

For each SAP HANA server, you must define and export a data and a log virtual volume (VV) to all SAP HANA servers. The first step is to create a common provisioning group (CPG). For example, you can use the following command to create a CPG named "HANA_data":

```
cli% createcpg HANA_data
```

After a CPG is created, the next step is to create the data and log volumes. These volumes should be thinly provisioned. Thinly provisioned volumes are known as "TPVVs." For example, you can use the following command to create a 1 TB data volume named "HANA_data":

```
cli% createvv -tpvv HANA_data 1t
```

After you create the data and log volumes, they should be presented to the HANA servers.

Creating a file system

After scanning the operating system on the SAP HANA host to make the exported VVs visible as Linux devices, create a file system on top of these Linux devices. HPE uses the XFS file system for the SAP HANA data and log files.

The following is a sample command to create a file system on one of the disk volumes:

```
# mkfs.xfs /dev/disk/by-id/scsi-360002ac0000000000000000200003c7e
/dev/disk/by-id/scsi-360002ac0000000000000000200003c7e
```

You must complete this step on all devices dedicated to SAP HANA data and log files.

In every SAP HANA scale-out environment, there must be an NFS shared storage directory for the SAP HANA instance. To offer this shared storage, one option is to use a Linux-based high availability cluster (for example, two physical or virtual compute nodes) providing an NFS export.

The file system can be mounted on a dedicated interface for storage traffic with an entry into the /etc/fstab on each host so this file system is mounted automatically. For example:

```
haix:/hana/shared /hana/shared nfs rw,bg,vers=3,hard,rsize=65536,wsiz=65536,nointr,actimeo=0,noexec 0 0
```

SAP HANA installation on HPE Primera Storage

The preparation before starting an SAP HANA scale-out installation procedure is the mapping of the exported HPE Primera VVs to the SAP HANA data and log partitions. This is done in a customized global.ini file to be placed in the SAP HANA installation directory /usr/sap/install.

The following output is from an HPE Primera Storage Virtual LUN that is identified as an operating system device:

```
# ls /dev/disk/by-id/scsi-360002ac00000000000000001600003c7e
/dev/disk/by-id/scsi-360002ac0000000000000000200003c7e
```

Following is a sample of the global.ini file required for the SAP HANA installation:

```
[storage]
ha_provider = hdb-ha.fcClient
ha_provider_path=/sapmnt/install
partition_*_*__mountpoints = -t xfs -o inode64,nobarrier,logbufs=8,logbsize=256k,swalloc
partition_1_data__wwid = 360002ac0000000000000000200003c7e
partition_1_log__wwid = 360002ac0000000000000000600003c7e
partition_2_data__wwid = 360002ac0000000000000000300003c7e
```



```
partition_2_log__wwid = 360002ac0000000000000000700003c7e
partition_3_data__wwid = 360002ac0000000000000000400003c7e
partition_3_log__wwid = 360002ac0000000000000000800003c7e
partition_*_*__prType = 5
```

After you implement these prerequisites, you can then execute the SAP HANA standard installation procedure to create an SAP HANA scale-out instance on top of the HPE Primera array.

HANA certification parameter settings

The SAP HWCCT test toolkit requires specific parameters to be set for the fsperf executable to test the I/O subsystem. The parameter settings listed in Table 2 have been used to run the performance part of the SAP HANA TDI Enterprise Storage certification.

TABLE 2. SAP HANA certification parameter values

SAP HANA file system I/O parameters (default)	Setting
async_read_submit: (on)	on
async_write_submit_active: (auto/on)	on
async_write_submit_blocks: (all)	all
size_kernel_io_queue: (512)	512
max_parallel_io_requests: (256)	64
min_submit_batch_size: (16)	16
max_submit_batch_size: (64)	64
num_submit_queues: (1)	1
num_completion_queues: (1)	1

You can change parameters for further tuning from the default after an initial SAP HANA installation is complete by using the HANA hdbparam command as <sid>adm:

```
su - <sid>adm
sh1adm@blade1:/usr/sap/SH1/HDB00> hdbparam --paramset fileio.max_parallel_io_requests=64
sh1adm@blade1:/usr/sap/SH1/HDB00> hdbparam --paramset fileio.size_kernel_io_queue=512
```

For SAP HANA 2.0 installations, to change from default parameters after the initial HANA installation is complete, you can modify the global.ini file as shown:

```
[fileio]
max_parallel_io_requests=64
size_kernel_io_queue=512
```

Refer to [SAP Note 2399079 - Elimination of hdbparam in HANA 2](#) for more information on how to configure the hdb parameter settings.

HPE Primera All-Flash for SAP HANA TDI configurations in mixed-workload environments

With HPE Primera Storage, you can securely partition resources within a shared infrastructure to pool physical storage resources for lower storage costs without compromising security or performance.

The HPE Primera Storage platform was built from the ground up to deliver multitenant capacity that supports massive consolidation with ultra-high performance. The multi-controller scalability and extreme flexibility built into HPE Primera All-Flash make deploying and maintaining separate storage silos to deliver different QoS levels a thing of the past. Unlike application-centric approaches to storage, one-click autonomic rebalancing on HPE Primera Storage enables you to enhance QoS levels without service disruption, pre-planning, or the need to purchase separate arrays to support different service levels. To support multiple tenants and workloads, HPE Primera Storage provides secure administrative segregation of users, hosts, and application data.



Quality of Service

HPE Primera Priority Optimization software enables you to select service levels for applications, virtual domains, and workloads as business requirements dictate. It protects mission-critical applications such as SAP HANA by specifying service-level objectives (SLOs) and service-level caps for I/O per second, bandwidth, and latency. If these goals are not met, Priority Optimization software automatically adjusts the service levels of lower-priority applications and workloads to help ensure necessary QoS levels for your highest-priority applications.

Key features and benefits

HPE Primera All-Flash Storage offers the following enterprise-level functions:

- Helps ensure QoS for your most mission-critical applications; protects from poor performance
- Creates and modifies threshold limits, including IOPS, bandwidth, and latency to meet business needs
- Aligns QoS with mission-critical apps as **high** and others as **medium** or **low** priority
- Configures thresholds by volume set with real-time enforcement measured in seconds or sub-seconds
- Provides easy configuration with autonomic SLOs for a VVset
- Enables you to control and balance the distribution of the powerful HPE Primera Storage disk bandwidth across multiple production and test and development workloads on a single system
- Avoids resource contention in the storage system by throttling the I/O of one or more workloads to allow the performance of others
- Manages user application performance expectations even with only one workload present

HPE DATA PROTECTION SOLUTIONS FOR SAP HANA

Data protection is a critical component of a robust environment for SAP HANA. Although persistent storage is in place to handle memory failures or server power loss, storage failures, corruption, or even natural disasters can occur. A robust data protection strategy is needed to mitigate these risks. A solid plan for backup and recovery of data, log files, operating system data, and configuration files must be in place.

SAP HANA offers two types of backups, which are both needed to recover the database to a specific point in time:

- Data backups can be triggered manually or scheduled in the SAP HANA Studio or SAP HANA Cockpit, or by issuing SQL commands. A data backup effectively replicates a database savepoint to the backup destination.
- Log backups occur automatically when a log segment (a log segment is represented by a file on disk with a fixed size) fills up or a configurable time threshold is exceeded. The log segment is copied to the backup destination. This may happen in parallel to a data backup.

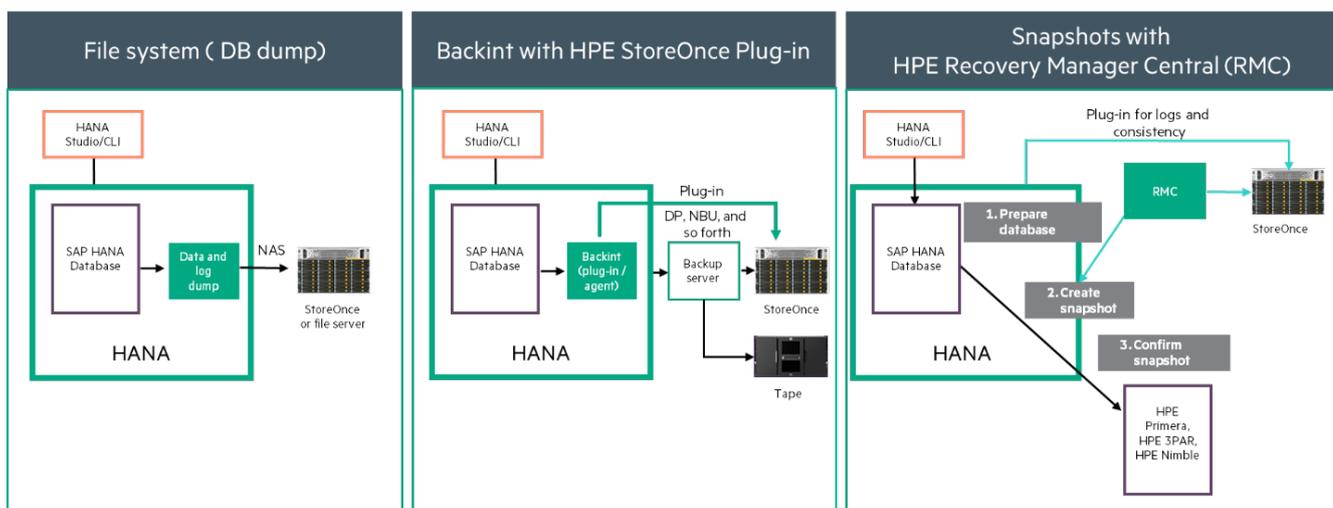


FIGURE 4. HPE SAP HANA backup and recovery option



SAP HANA offers three options to back up the database:

- **File system:** You can use an external shared file system as the backup target for all nodes of an SAP HANA installation. This the easiest way to implement a backup solution within an SAP HANA environment. SAP recommends that the shared file system for the backup not use the same storage as the database.² A file-system based solution can be used for log and data backups. HPE offers HPE StoreOnce Backup Systems, which can be configured as an NFS target. Learn more about [HPE StoreOnce data protection](#).
- **Backint:** You can use Backint, which is essentially an API SDK available for backup vendors to implement, to access an enterprise backup solution (EBS) certified with SAP HANA. Leveraging Backint, the HPE StoreOnce Catalyst plug-in for SAP HANA enables optimized protection for SAP HANA databases. It enables a DBA to back up and restore directly to or from a Catalyst store on an HPE StoreOnce backup appliance. The result is flexible, high-performance protection, managed by the DBA, that can be configured to meet the protection needs of specific databases independently or in addition to organization-wide data protection processes. The plug-in is integrated with Backint for SAP HANA to enable backups to be efficiently transferred from the SAP HANA database to the HPE StoreOnce backup target.

After the backup target is created, backup, restore, and other data protection tasks can be executed through SAP HANA Studio or the SAP HANA CLI. HPE StoreOnce plug-in for SAP HANA must be installed using the SAP HANA operating system user, not the root user. The plug-in has guided installation and configuration for fast setup. This plug-in has a zero-cost license to use but you must purchase and install a Catalyst license for all appliances that host the target HPE StoreOnce stores. It is supported for use with all HPE StoreOnce appliances running the required software version. Consult the [HPE StoreOnce Catalyst for SAP HANA technical white paper](#) for more information.

- **Data snapshot:** With SPS 07 and later, SAP HANA can use storage snapshots for backups. Third-party tools can put an SAP HANA database into a prepare mode before creating a storage snapshot of the data volume. The benefit of using a storage snapshot is that it can be created with minimal impact to the SAP HANA system. Additionally, a restore from a snapshot is faster than a recovery from a file or Backint. This snapshot-based backup solution is designed for data volume only. SAP does not support application-consistent snapshots for the log volume. Note that additional storage capacity³ is required to implement a snapshot backup solution.

HPE Recovery Manager Central Software (RMC) is a snapshot-based copy data management solution that leverages HPE Primera Virtual Copy capability. Storage snapshots offer an alternative to keeping the SAP HANA database in online backup mode for an extended period of time. Using RMC, you can copy the data snapshot from HPE Primera to HPE StoreOnce. This process is called "Express Protect." You can retain this Express Protect copy on HPE StoreOnce for a longer period of time than the snapshot on HPE Primera. If you want to restore the SAP HANA database to a previous point in time, you can use either the snapshot or the Express Protect copy. You can also use these copies to create a new SAP HANA instance on the same HPE Primera array or on a new HPE Primera array connected to the RMC instance.

RMC creates an initial full Express Protect copy on the HPE StoreOnce system. All subsequent backup operations transfer only the changed data.

NOTE

For more information, read the [SAP HANA Copy Data Management using HPE Recovery Manager Central](#) white paper.

² For more information, refer to SAP Note 1820529.

³ SAP estimates a 10% daily change rate.



SUMMARY

HPE Primera All-Flash arrays are an industry-leading offering for SAP HANA Enterprise Storage TDI solutions. They are high-performing, efficient, and highly available solutions with integrated data protection.

Additional links

HPE Intelligent Data Platform

<https://www.hpe.com/us/en/storage/intelligent-storage.html>

Data protection storage

<https://www.hpe.com/us/en/storage/data-protection-solutions.html>

Mission-critical systems

<https://www.hpe.com/us/en/servers/mission-critical.html>

LEARN MORE AT

[hpe.com/us/en/storage/mission-critical-applications.html](https://www.hpe.com/us/en/storage/mission-critical-applications.html)

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates

© Copyright 2020 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.

Intel Xeon is the trademark of Intel Corporation in the U.S. and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. SAP and SAP HANA are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All third-party marks are property of their respective owners.