Hewlett Packard Enterprise

Reference Architecture

HPE Reference Architecture for SAP HANA 2.0 on HPE SimpliVity 380 and 380G

CONTENTS

Executive summary 3
Solution overview 3
Solution components 4
HPE SimpliVity overview 4
SAP HANA 5
HPE Serviceguard for Linux 6
SAP HANA system replication 7
HPE SimpliVity RapidDR 8
HPE SimpliVity requirements for production SAP HANA 8
Design considerations and best practices 10
Configure virtual machine 12
Tune OS for SAP HANA 14
Network best practices 16
HPE SimpliVity deployment networks 16
Backup solution for SAP HANA 18
HPE SimpliVity Backup + SAP HANA log backup to file system 19
HPE SimpliVity backup with quiesce SAP HANA database 19
HPE StoreOnce Plug-in for SAP HANA 20
Disaster recovery solution for SAP HANA 21
SAP HANA system replication with HPE Serviceguard 21
HPE SimpliVity RapidDR 23
Building a comprehensive disaster recovery solution 23
SAP HANA on HPE SimpliVity additional benefits 26
Summary 28
Appendix A: Bill of materials 29
Resources and additional links 31

EXECUTIVE SUMMARY

HPE SimpliVity hyperconverged infrastructure (HCI) solution is designed from the ground up to meet the stringent price-performance, scalability, and resiliency demands of today's data-intensive applications and highly virtualized IT environments. The HPE SimpliVity solution revolutionizes the data center, by virtualizing data and assimilating all IT infrastructure and services below the hypervisor into x86 building blocks.

SAP extended its tailored data center integration (TDI) certification to support hyperconverged platforms. This enabled customers to put the production SAP HANA® database on a hyperconverged platform. Hewlett Packard Enterprise participated in the certification program and is now certified to run production SAP HANA database on HPE SimpliVity. This helps customers to deploy SAP HANA on a cost-effective, easy-to-manage, smaller footprint infrastructure with no compromise on performance.

This Reference Architecture provides guidelines to plan, prepare, and deploy SAP HANA on the HPE SimpliVity platform. This guide is intended for SAP administrators and architects, who design mission-critical SAP HANA systems to be deployed on the HPE SimpliVity platform. This paper provides the reader with specific examples of HPE SimpliVity configurations that support real-world SAP HANA deployments in virtualized environments; deployments that meet the reliability, performance, scalability, and data protection requirements of the midsize-to-large enterprise. It also highlights the unique benefits of HPE SimpliVity that complement and enhance SAP HANA.

SOLUTION OVERVIEW

HPE SimpliVity HCI platform is designed to meet the enterprise-class performance, protection, and resiliency demands of the SAP HANA in-memory database. It runs on industry-leading HPE ProLiant DL380 Gen10 servers by clustering multiple HPE SimpliVity nodes to form a shared resource pool.

This solution leverages both HPE SimpliVity 380 Gen10 and HPE SimpliVity 380G Gen10 Nodes with VMware vSphere® 6.7 hypervisor. For SAP HANA OS, Red Hat® Enterprise Linux (RHEL) 7.5 is used.

Figure 1 shows the solution diagram.

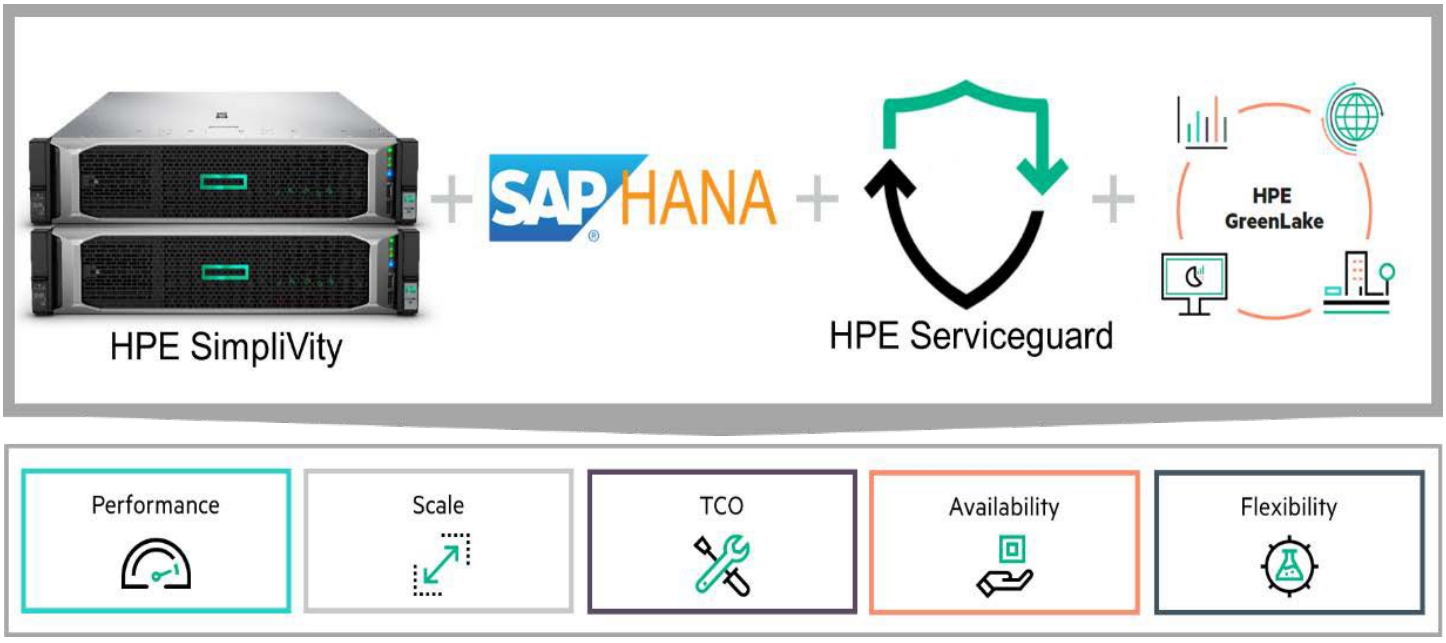

FIGURE 1. SAP HANA reference diagram

HPE SimpliVity technology delivers multiple levels of redundancy along with built-in data protection and disaster recovery capabilities. It includes the ability to withstand multiple drive failures without any loss of data, along with out-of-the-box local failover that can withstand a node outage and continue delivering high performance without data loss. HPE SimpliVity infrastructure's built-in data protection provides backup, replication, and restore capabilities for the SAP HANA database with no additional license fees.

Figure 2 shows the HPE SimpliVity network diagram.

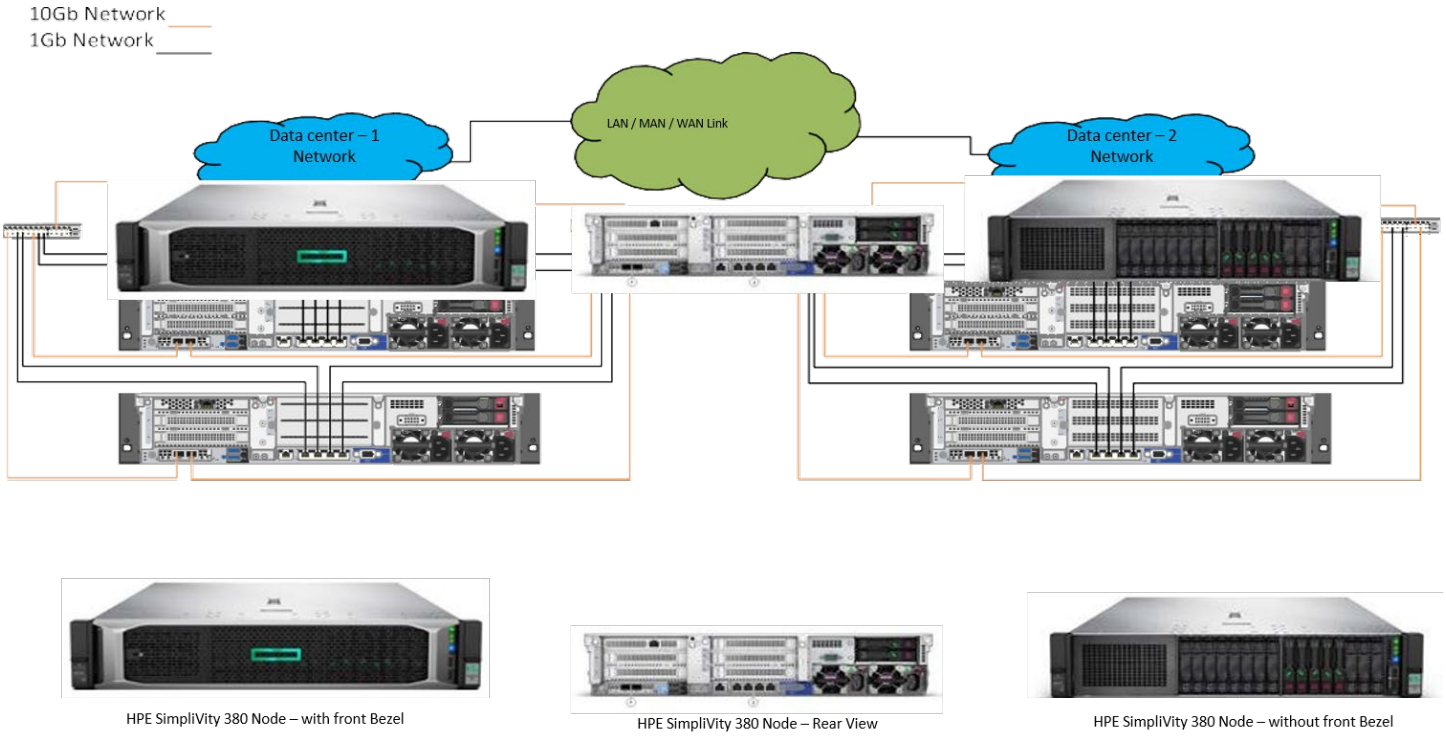

FIGURE 2. HPE SimpliVity Federation Network diagram

A common deployment for HPE SimpliVity 380 Gen10 is a two-site environment, in which HPE SimpliVity 380 Gen10 systems are deployed in each of two data centers and connected within the same Federation. This configuration provides high availability and local backup recoverability in each site. It also delivers elegant disaster recovery and simplified management as the SimpliVity systems and all associated VMs are managed from a single user interface pane within the VMware vCenter Server management Web Client.

The industry-proven HPE Serviceguard offers mission-critical high availability and disaster recovery protection with automatic failover capability to ensure continued uninterrupted availability. This is deployed in most mission-critical SAP HANA installations in the world, and it is one of the industry's best-automated failover and disaster recovery solutions for SAP HANA.

With HPE GreenLake, customers can take advantage of consumption-based IT solutions which are designed, delivered, and operated by HPE Pointnext. HPE GreenLake for SAP HANA provides the benefits of an on-prem solution, removing the unknowns associated with the network (bandwidth, latency) and security, and allowing customers to effectively have more control.

SOLUTION COMPONENTS

Hewlett Packard Enterprise provides a comprehensive solution for SAP HANA on HPE SimpliVity hyperconverged platform. This section describes the solution components.

HPE SimpliVity overview

HPE SimpliVity 380 and HPE SimpliVity 380G are hyperconverged infrastructure solutions built for the enterprise. It runs on industry-leading HPE ProLiant DL380 Gen10 servers to deliver turnkey hyperconverged infrastructure for the software-defined data center. Clustering multiple HPE SimpliVity nodes form a shared resource pool and deliver high availability, mobility, and efficient scaling of performance and capacity.

HPE SimpliVity assimilates all IT infrastructure and services below the hypervisor into a single, scalable 2U building block. Two or more instances of HPE SimpliVity building blocks form a cluster and multiple clusters form a federation – delivering a scalable pool of shared resources.

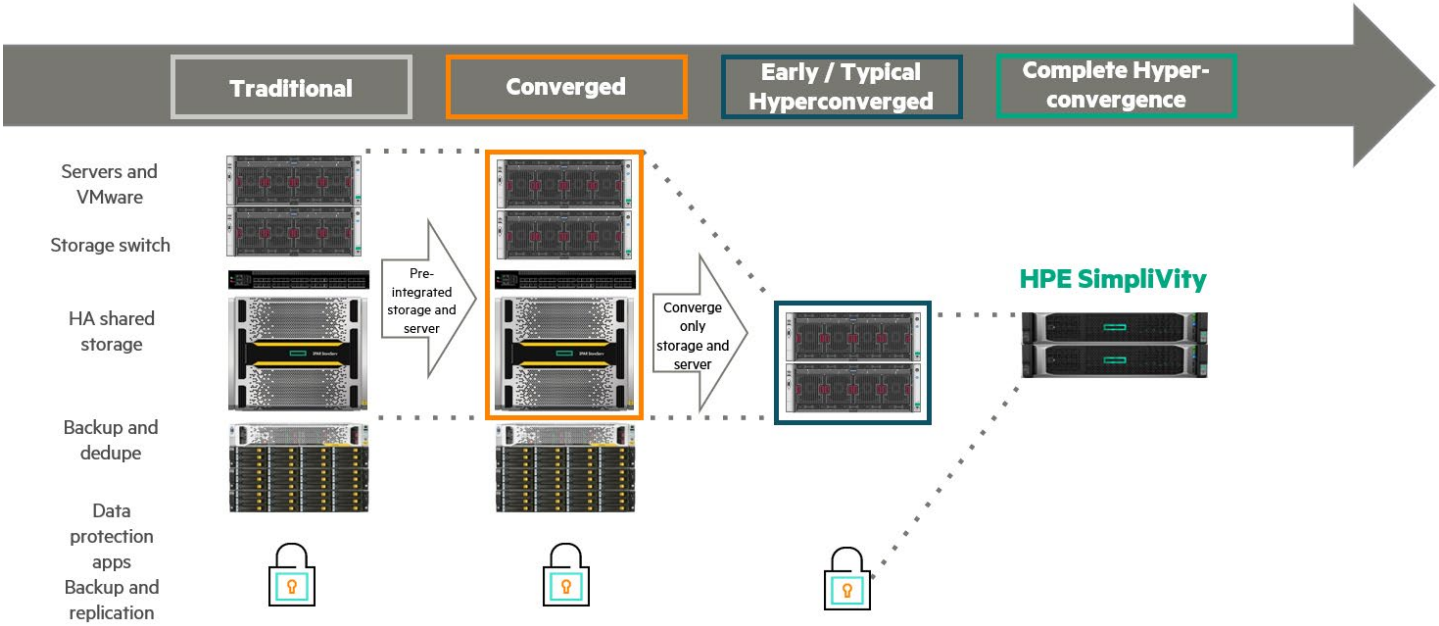

FIGURE 3. HPE SimpliVity consolidation

HPE SimpliVity offers differentiation from other converged infrastructure in its data virtualization platform, which enables the user with a single shared resource pool across multiple sites and provides highly efficient data storage and mobility. HPE SimpliVity Data Virtualization Platform consists of the HPE SimpliVity data architecture and native deduplication and compression algorithms, both of which enable accelerated data efficiency, global unified management, and built-in data protection.

To ensure continuous operations, HPE SimpliVity technology delivers multiple levels of redundancy along with built-in data protection and disaster recovery capabilities. It includes the ability to withstand multiple drive failures without any loss of data, along with out-of-the-box local failover that will withstand a node outage and continue delivering high performance without loss of transactions. HPE SimpliVity infrastructure's built-in data protection provides backup, replication, and granular restore capabilities for the entire virtual desktop infrastructure, with no additional license fees. Also, the optional HPE SimpliVity RapidDR feature automates and accelerates off-site disaster recovery.

All HPE SimpliVity models come with either one or two sockets. However, we must choose the 2-socket configuration for SAP HANA workloads to align with the SAP recommendation of dedicating CPU resources for SAP HANA production workload. When choosing the HPE SimpliVity model, we recommend the usage of 6000/Mixed-use drives for SAP HANA workload as they tolerate flash wear better than the 4000 series. However, customers can choose 4000 series nodes if their SAP HANA workloads are highly read-intensive.

SAP HANA

SAP HANA software is the next generation of the SAP technology platform. Originally designed to support large-scale analytics, it has been adapted to support SAP business applications. The primary innovations of SAP HANA are an in-memory database and an analytics engine that can load the entire database into memory. This capability enables businesses to rapid processing of multiple terabytes of data and dramatically reduce the time required to perform queries and complex analytics.

Although SAP HANA is an in-memory database, it still requires a persistent storage area to maintain data between restarts and after unplanned server shutdowns. SAP has implemented a certification program to ensure that hardware vendors meet the performance, scalability, and high-availability requirements of the SAP HANA software.

SAP HANA is delivered to customers in two ways: as an appliance and through the SAP HANA tailored data center integration (TDI) program. The appliance model provides a preinstalled and preconfigured hardware platform for running SAP HANA. Hardware vendors are required to

configure these systems and to allow SAP to test them to guarantee that the hardware meets SAP HANA's key performance indicators. The TDI model enables customers to combine server, network, and storage resources that have been individually certified by SAP to support SAP HANA.

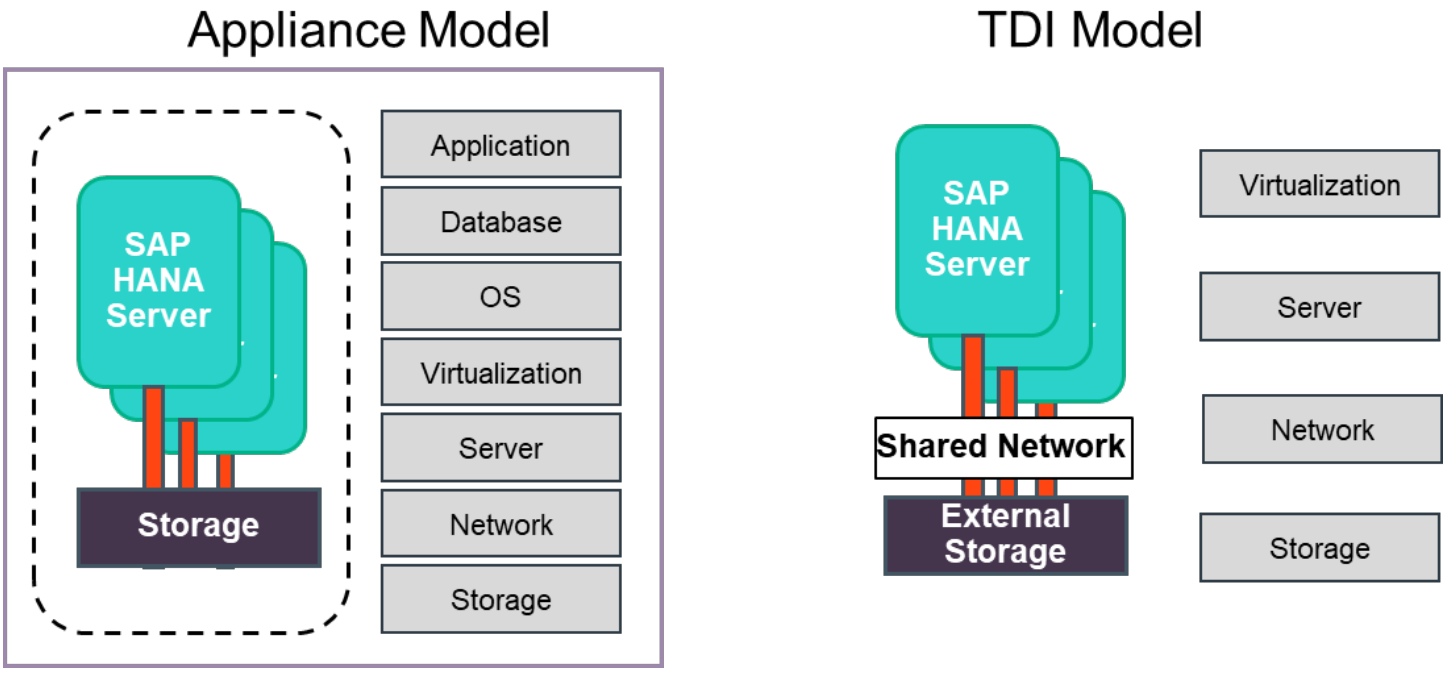

FIGURE 4. SAP HANA hardware certification models

SAP extended its tailored data center integration (TDI) certification to support hyperconverged platforms. This enabled customers to put the production SAP HANA database on a hyperconverged platform. Hewlett Packard Enterprise participated in the certification program and is now certified to run production SAP HANA database on HPE SimpliVity. This helps customers to deploy SAP HANA on a cost-effective, and easy-to-manage, smaller footprint infrastructure with no compromise on performance. The HCI certification can be found in https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/#/solutions?filters=hci&id=s:2457.

SAP HANA can be deployed in two configurations, scale-up and scale-out:

- Scale-up systems are single-server systems that run the SAP HANA database and application. In the appliance model, these systems typically shared storage. HPE SimpliVity is used for scale-up system deployment.
- Scale-out systems are multiple servers with a shared storage infrastructure that can support larger volumes of data by partitioning the SAP HANA database between the servers. Currently, only SAP Business Warehouse (BW) is supported on scale-out configurations. This option is not supported on HPE SimpliVity today.

HPE Serviceguard for Linux

HPE Serviceguard for Linux (SGLX) is a software-based application high availability and disaster recovery solution that increases the availability and uptime of your business-critical applications and minimizes the impact of unplanned outages. SGLX packages applications and other services with their associated resources and monitors the entire package for any failure.

Each package is monitored for faults related to hardware, software, OS, virtualization layer, virtual machine guests, network, and storage. When any failure is detected, SGLX shuts down the application to relocate the application or service quickly and smartly to another system with the necessary resources to bring it into production again.

SGLX uses Quorum Server arbitration mechanisms to prevent data corruption and loss in case of a split-brain situation among cluster nodes. The solution in this Reference Architecture uses VM as cluster node and running Quorum Server outside the SGLX cluster. SGLX also minimizes

planned downtimes using its Live Application Detach (LAD) and Rolling Upgrades features to perform maintenance on clusters and install upgrades for any OS and application without downtime. Take advantage of the Cluster Verification technology to find and fix cluster configuration issues before they advance and cause unplanned downtime.

This Reference Architecture has captured the test results of HPE Serviceguard for SAP HANA replication on Linux using 2 node cluster configurations spread across two sites. This manages and completely automates the takeover process of SAP HANA replication databases between primary and standby nodes. SGLX is utilized to start, stop, and monitor the databases and administer the replication between primary and standby databases. SGLX also performs automatic role management to recover from failures. In case of failures, the solution automatically recovers the SAP HANA database by promotion of the standby database instance to primary. This mode of recovery is much faster compared to the restart-DB-based solution. The databases can be located on the same premises or in geographically dispersed data centers. SGLX performs recovery point objective (RPO) sensitive automatic role management to recover from failures.

SAP HANA system replication

SAP HANA system replication is a reliable high availability and disaster recovery solution that provides continuous replication and synchronization of a HANA database to a secondary location either in the same data center or remote site. System replication is a standard SAP HANA feature. In this method, all data is replicated to the secondary site, and data is pre-loaded into memory on the secondary site which helps to reduce the recovery time objective (RTO) significantly. So, in case of a failover, the secondary site will be able to take over without even performing a HANA DB (re)start and will work as primary DB straightaway.

Once the system replication is configured properly, each HANA internal process (name server, index server, etc.) connects to its counterpart on the secondary site, and all logged changes in the primary location are replicated to the secondary site continuously through persistent redo logs.

While system replication is running, the secondary site is on standby mode with data already pre-loaded into memory, and ready to take over. SAP HANA offers different modes for the replication of the redo log.

TABLE 1. SAP HANA system replication modes

Replication Mode	Description
Synchronous in- memory (SYNCMEM)	The secondary system sends an acknowledgment back to the primary system as soon as the data is received in memory. The disk I/O speed on the secondary system doesn't influence the primary's performance. When the connection to the secondary system is lost, the primary system continues the transaction processing and writes the changes only to the local disk. Data loss can occur when primary and secondary fail at the same time if the secondary system is connected or when a takeover is executed, while the secondary system is disconnected. This option provides better performance because it is not necessary to wait for disk I/O on the secondary system, but it is more vulnerable to data loss.
Synchronous (SYNC)	The secondary system sends an acknowledgment back to the primary system as soon as the data is received and persisted to the log volumes on the disk. When the connection to the secondary system is lost, the primary system continues the transaction processing and writes the changes only to the local disk. No data loss occurs in this scenario if the secondary system is connected. Data loss can occur when a takeover is executed while the secondary system is disconnected. Additionally, this replication mode can run with a full sync option. This means that log write is successful when the log buffer has been written to the log file of the primary and the secondary systems. When the secondary system is disconnected (for example, because of network failure), the primary system suspends the transaction processing until the connection to the secondary system is reestablished. No data loss occurs in this scenario.
Asynchronous (ASYNC)	The primary system sends redo log buffers to the secondary system asynchronously. The primary system commits a transaction when it has been written to the log file of the primary system and sent to the secondary system through the network. It doesn't wait for confirmation from the secondary system. This option provides better performance because it is not necessary to wait for log I/O on the secondary system. Database consistency across all services on the secondary system is guaranteed. However, it is more vulnerable to data loss. Data changes may be lost on the takeover.

HPE SimpliVity RapidDR

HPE SimpliVity RapidDR simplifies and accelerates off-site disaster recovery through automation. The solution extends the inherent data efficiencies of HPE SimpliVity hyperconverged infrastructure, slashing recovery-point objectives (RPOs), and recovery time objectives (RTOs) from days or hours to minutes.

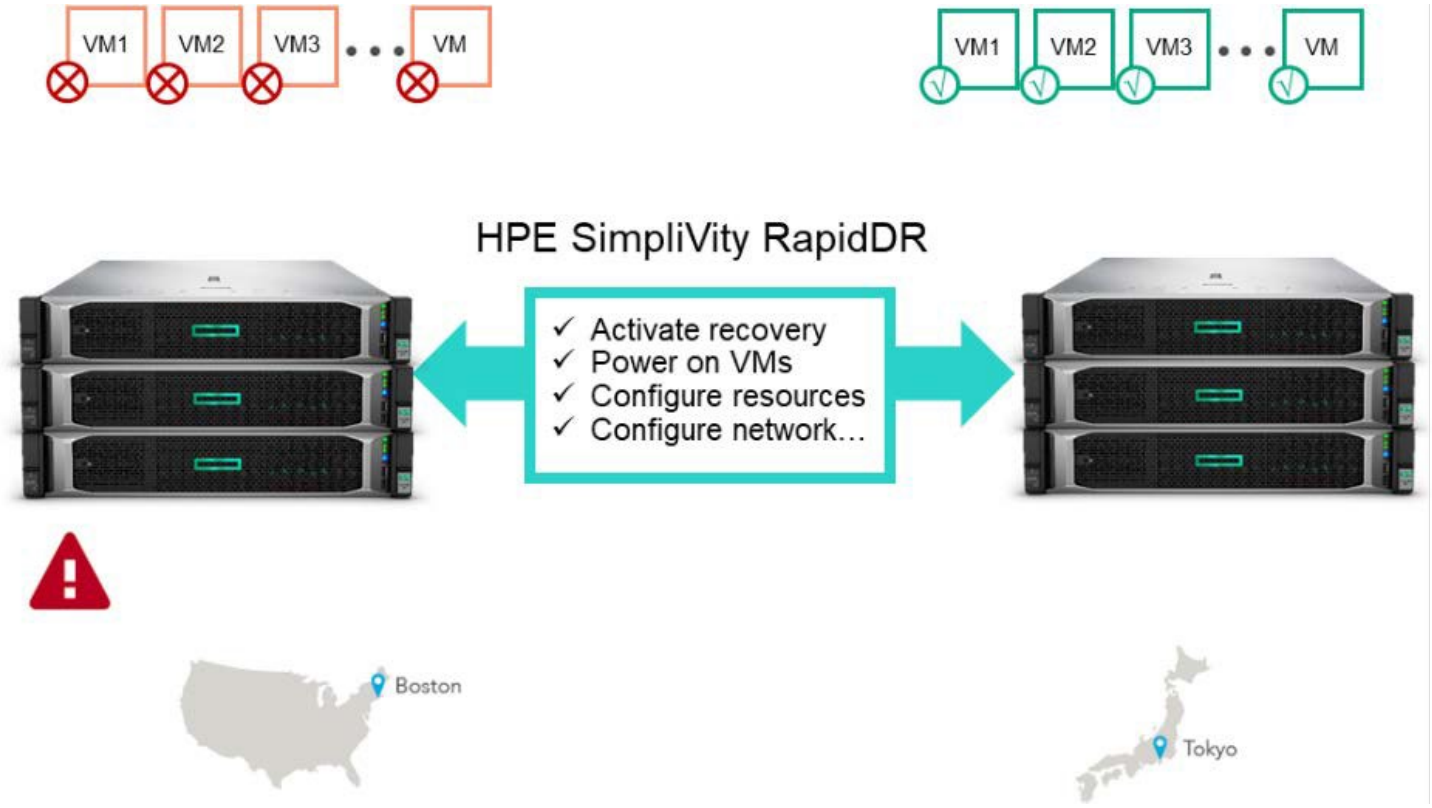

FIGURE 5. HPE SimpliVity RapidDR

HPE SimpliVity RapidDR eliminates complicated DR scripting processes and complex runbooks that consume IT time, resources, and budget. An intuitive, VM-centric GUI guides the system administrator through the configuration process streamlining setup. A configuration such as VMs to restore at the destination site, startup order, networking for VMs can be automated using HPE SimpliVity RapidDR.

HPE SIMPLIVITY REQUIREMENTS FOR PRODUCTION SAP HANA

SAP has a strict requirement to certify the HPE SimpliVity HCI platform to place the production SAP HANA database. Prior to deploying an SAP HANA production on HPE SimpliVity platform, it is important to follow the prerequisite to build out a fully supported infrastructure. Currently, SAP HANA scale-out deployments are not supported. HPE SimpliVity 380 and HPE SimpliVity 380G are certified to run production SAP HANA scale-up deployments. Table 2 and Table 3 provide supported configurations for both the models respectively.

TABLE 2. HPE SimpliVity 380 requirement for production SAP HANA

Specification	Supported Element
HPE SimpliVity platform	HPE SimpliVity 380 Gen10 (Hardware accelerated)
HPE SimpliVity Disk Type	6000 series (Recommended) or 4000 series drives
HPE SimpliVity version	4.0.1
Supported hypervisor	VMware ESXi™ 6.7 U3
Supported processor type	Intel® Cascade lake

Specification	Supported Element
Number of sockets per node	2
Min number of cores per socket	8
Maximum memory per node	Up to 3 TB
SAP HANA version	SAP HANA 2.0
Number of disks per node	5/9/12
SAP HANA deployment type	Scale-up

TABLE 3. HPE SimpliVity 380G requirement for production SAP HANA

Specification	Supported Element
HPE SimpliVity platform	HPE SimpliVity 380G Gen10 (Software-optimized)
HPE SimpliVity Disk Type	Mixed-use (Recommended) or Read-intensive drives
HPE SimpliVity version	4.1.0
Supported hypervisor	VMware ESXi™ 6.7 U3
Supported processor type	Intel® Cascade lake
Number of sockets per node	2
Min number of cores per socket	8
Maximum memory per node	Up to 3 TB
SAP HANA version	SAP HANA 2.0
Number of disks per node	6/8/12/16
SAP HANA deployment type	Scale-up

For a production SAP HANA instance on HPE SimpliVity, a minimum of two nodes is required (up to 96 nodes are supported in a single HPE SimpliVity federation with ESXi 6.7 U3), while one node is reserved for high availability.

DESIGN CONSIDERATIONS AND BEST PRACTICES

SAP HANA is a resource-intensive application. Improper sizing will impact SAP application performance to the end-users., which will have a negative impact on both the IT services and business outcomes. It is important to plan and design a solution that meets the customer's requirements. This section will cover design considerations and best practices.

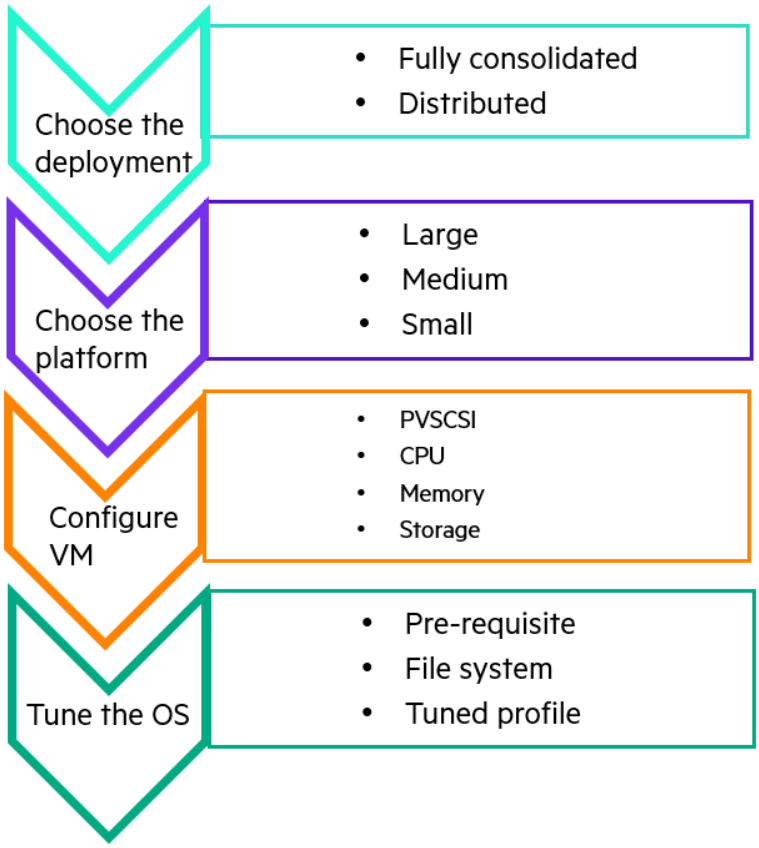

FIGURE 6. SAP HANA best practice flow

Choose the right deployment

Once the requirements to deploy SAP HANA are gathered, one of the main decisions is to identify what type of deployment suits their business needs. This decision is based on the customer's resource requirement between an SAP application and HANA database, as well as the protection service-level agreement. HPE SimpliVity provides several deployment options to implement SAP HANA and application workload.

Option 1: Single Cluster deployment

This option allows installing the SAP HANA database and application on the same HPE SimpliVity nodes. This configuration allows us to consolidate all the SAP components into a single cluster. The HANA DB will be run on a dedicated HPE SimpliVity node socket and the application will run on the other socket.

NOTE

The HPE SimpliVity OmniStack Virtual Controller is hosted on socket 0 by default and this socket cannot be shared with the SAP HANA database instance.

Figure 7 shows the consolidated configuration.

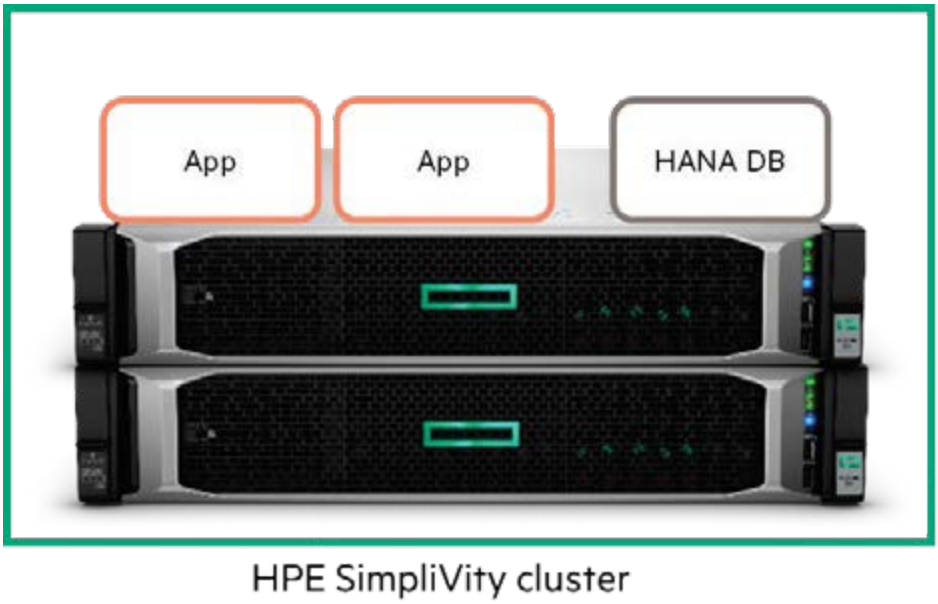

FIGURE 7. Fully consolidated configuration

Option 2: Multi-Cluster deployment

This option allows installing the SAP HANA database and application on two separate HPE SimpliVity clusters. This configuration allows sizing the HPE SimpliVity clusters for the SAP HANA database and applications separately. Also, it enables us to have different protection plans for SAP HANA and its applications.

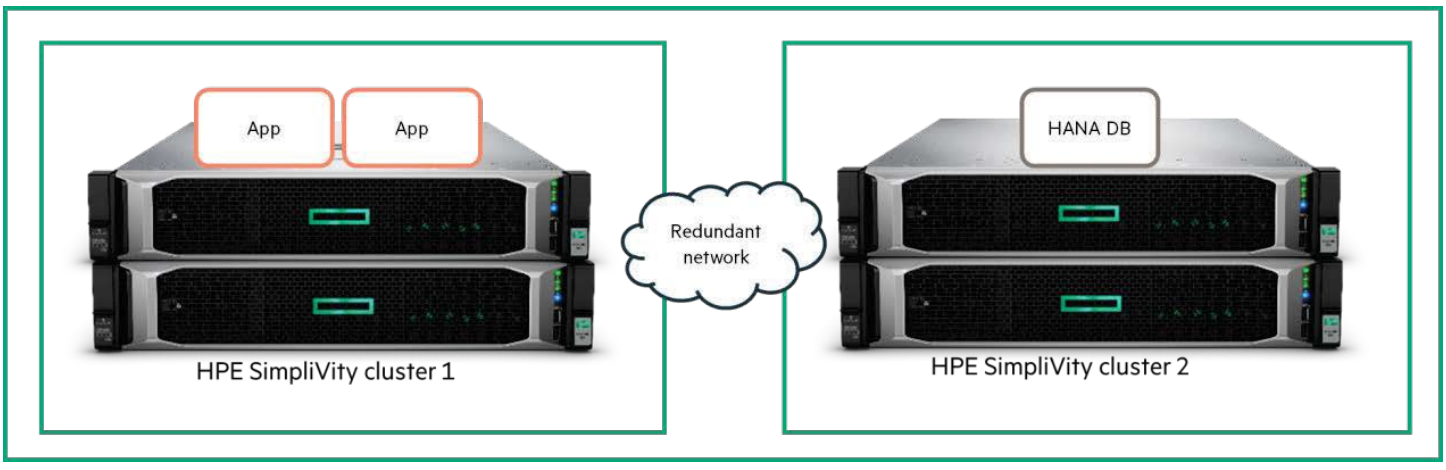

FIGURE 8. Distributed configuration

Size the right HPE SimpliVity node

HPE SimpliVity nodes come in different sizes. This allows customers to buy the right node with proper sizing for memory, CPU, and storage capacity. Currently, we recommend running one SAP HANA database per node irrespective of the form-factor as the certification is achieved for one SAP HANA instance per node. We don't recommend HPE SimpliVity 380 Gen10 x-small nodes for the SAP HANA database since the disks are not designed for a high write-intensive workload.

HPE SimpliVity models come with one and two sockets. However, we must choose a 2-socket configuration for SAP HANA to align with the SAP recommendation of dedicating CPU resources for SAP HANA production workload. We also recommend the usage of 6000 or Mixed-use drives for SAP HANA workload as they tolerate flash wear better than the 4000 series. However, customers can choose read-intensive (RI) or 4000 series nodes if their SAP HANA workloads are read-intensive.

SAP workloads are unique and the best configuration and sizing for SAP HANA can be different for each TDI scenario. To determine the SAP HANA requirements, you have these options:

- Use SAP tools like the SAP Quick Sizer
- SAP sizing reports for existing systems, (e.g. the reports described in SAP Note 1872170 - ABAP on HANA sizing report and SAP Note 2296290 - New Sizing Report for SAP BW/4HANA)
- Contact your local HPE Representative

Tables 4 and 5 show HPE SimpliVity nodes and recommended configuration.

HPE SimpliVity 380 Gen10

- **CPU -** 2P Intel® Xeon® Scalable Processors 8 to 28 cores
- **Memory** - 144GB to 3072GB per node selectable

TABLE 4. Choosing HPE SimpliVity 380 node type

Model	Storage kit	Maximum SAP HANA production instances/node
Small	5 x 1.92TB	1
Medium	9 x 1.92 TB	1
Large	12 x 1.92TB	1

HPE SimpliVity 380G Gen10

- **CPU -** 2P Intel® Xeon® Scalable Processors 8 to 28 cores
- **Memory** - 128GB to 3072GB per node selectable

TABLE 5. Choosing HPE SimpliVity 380G node type

Model	Storage kit	Maximum SAP HANA production instances/node
x6	6 x 1.92TB	1
x8	8 x 1.92TB	1
x12	12 x 1.92TB	1
x16	16 x 1.92TB	1

Refer to the HPE SimpliVity 380 Gen10 Quickspecs for more information.

Configure virtual machine

SAP recommends dedicating server resources to production HANA VM for optimum performance. This section provides best practice guidelines on how to achieve resource dedication along with the number of configurations to improve the overall SAP HANA performance.

Use multiple PVSCSI for virtual disk configuration

SAP HANA VM requires several file systems to place data, log, shared data, root, and local storage. Paravirtual SCSI (PVSCSI) adapters are high-performance storage adapters that can result in greater throughput and lower CPU utilization and are best suited for VMs running database systems such as SAP HANA. For optimum disk access performance, create a primary SCSI adapter for the file system for root and local storage,

and multiple separate SCSI adapters for the SAP HANA data, log, and shared data. Using multiple virtual SCSI controllers parallelizes the units of work in a database and significantly increases the throughput and lowers the access time.

Here is the VM configuration of the tested environment.

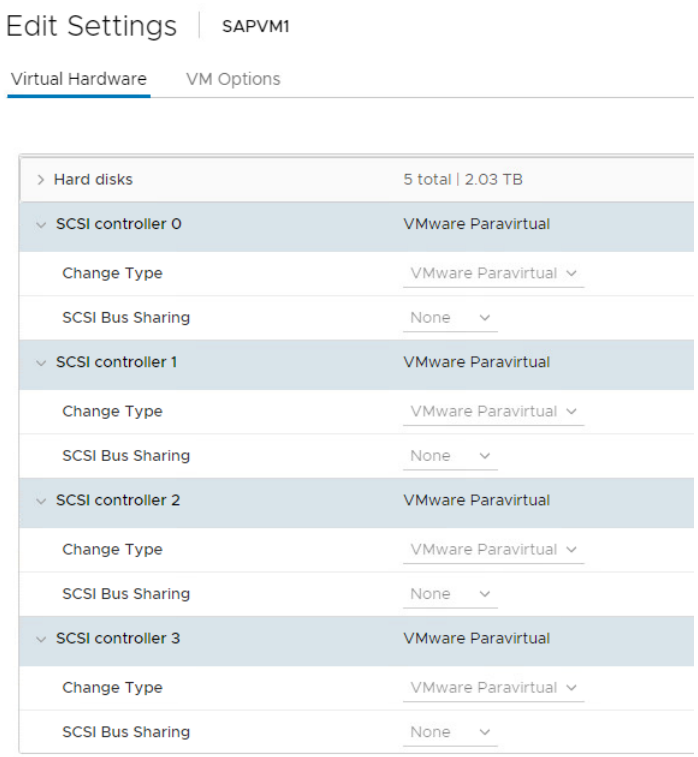

FIGURE 9. PVSCSI configuration

Disk consideration

Use a dedicated PVSCSI controller for the different SAP HANA volumes to ensure high IO bandwidth and low latency. To parallelize the unit workloads, create several VM disks for each HANA system. For best performance, we recommend one each for the OS, usr/sap, hana/shared, SAP HANA log, and SAP HANA data. Use the SAP Quick sizer for the right sizes of the SAP HANA data and log devices. The disk sizes listed in Table 6 are examples and can be adapted according to the results of the SAP Quick Sizer and the SAP HANA TDI Storage Requirements.

TABLE 6. Example of SAP TDI disk size recommendation

Mountpoint	Disk size used during current validation	SAP disk size recommendation	Controller
/ (root)	50 GB	min 10 GB	PVSCSI controller 0
/usr/sap	100 GB	min 50 GB	PVSCSI controller 1
/hana/shared	256 GB	min 1 x RAM up to 1 TB	PVSCSI controller 1
/hana/data	1.2 TB	min 1 x RAM in total	PVSCSI controller 2
/hana/log	500 GB	min 0.5 x RAM in total, up to 512 GB	PVSCSI controller 3

Do not share SAP HANA production CPU socket with other workloads

SAP enforces to dedicate a server CPU socket for SAP HANA production workload. This ensures that the production HANA database has enough resources and it delivers predictable performance. HPE OmniStack Virtual Controller (OVC) uses socket 0 of an HPE SimpliVity node. Place SAP HANA production virtual machines on socket 1, so it will not be sharing resources with HPE OVC or other virtual machines. The

following parameter is configured on the SAP HANA production virtual machine to allocate an unused socket for SAP workload. On the virtual machine advanced configuration, set numa.nodeAffinity to 1.

HPE OVC reserves CPU cores on Socket 0 (four to six on HPE SimpliVity 380 and six to eight on HPE SimpliVity 380G) to provide storage, management, and data protection functionalities. The remaining unused CPU cores on socket 0 can be used for non-production SAP HANA VMs or other application workloads VMs.

Allocate multiple sockets to production VM

With VMware as the hypervisor, HPE SimpliVity nodes provide several CPU allocation options to meet small to large SAP HANA production CPU requirements. A minimum of eight physical cores must be configured as per the SAP recommendation. Depending on the workload, a minimum of 8 cores to the whole socket can be allocated for SAP HANA VM. Ensure that all the CPU cores of the allocated sockets are reserved for SAP HANA VM.

When allocating sockets to VM, use numa.nodeAffinity parameter to set the correct socket number in the VMware configuration. The following link docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.resmgmt.doc/GUID-A80A6337-7B99-48C8-B024-EE47E2366C1B.html provides detailed steps on allocating CPU sockets to a VM.

Enable hyperthreading

SAP HANA is a latency-sensitive workload with higher memory utilization and moderate processor utilization. This type of workload requires higher interthread communication. For this reason, VMware recommends using hyperthreads with fewer NUMA nodes instead of full physical cores spread over multiple NUMA nodes.

When hyperthreading is enabled, each physical CPU core will have two threads (thread 0, 1). The default behavior of VMware vSphere® scheduler is to assign a vCPU to an idle physical CPU core. Therefore, hyperthreads (thread 1) are used only when all thread 0 of all physicalCPU cores are already utilized. To allocate vCPU of the same physical CPU (thread 0, 1) to the virtual machine, configure numa.vcpu.preferHT=TRUE or Numa.PreferHT=1 in the advanced configuration on the virtual machine. See the kb.vmware.com/s/article/2003582 for details.

Allocate memory within NUMA boundary

SAP HANA heavily utilizes memory. Allocating memory outside the boundaries of a NUMA node causes the virtual machine to fetch memory from a remote location, and this causes higher latency. SAP HANA VMs should be sized to stay within a NUMA node for optimum performance. Besides the performance implication, NUMA node sharing is not supported by VMware for SAP HANA production workload.

Allocate different memory per CPU socket

Since SAP HANA does not allow HPE SimpliVity socket 0 to be used for production SAP HANA, you can configure different memory for each socket. For example, you can allocate the maximum memory (768 GB) for socket 1 and use it to put SAP HANA production database and assign lower memory on socket 0 and used that to put test or dev SAP HANA database. This will allow allocating resources optimally.

Tune OS for SAP HANA

SAP HANA production is supported to run on SUSE Linux® Enterprise and Red Hat® Enterprise Linux. HPE SimpliVity HCI Certification testing used Red Hat Enterprise Linux 7.4. This section provides operating system level setup and tuning prior to installation of production SAP HANA database.

Install OS and required packages

Installing the operating system and the required package will make the SAP HANA deployment smooth. Follow these high-level steps to install the OS:

1. Install minimum Red Hat OS and enable subscription.
2. Install the base package group and other additional packages required for running SAP HANA.
3. Install GCC 6 C++ Runtime Compatibility Libraries for SAP HANA 2.0.
4. Configure hostname, network, time, and date.

Configure storage and file system

A minimum of four storage file systems must be created for SAP HANA VM to host data, log, shared data, root, and local storage. The default file system of XFS is recommended. For sizing the file system, use the SAP Quick Sizer tool for more accurate sizing.

Configure tuned profile

Apply the recommended SUSE Linux Enterprise Server or Red Hat Enterprise Linux OS settings for SAP HANA, such as saptune or tuned. Details can be found in the following resources:

- Linux Operating System with SAP HANA Reference Guide
- SAP HANA on VMware vSphere

A tuning service can adapt the operating system to perform better under certain workloads by setting a tuning profile. Red Hat has developed specifically tuned profiles to optimize the performance of SAP HANA on RHEL. For optimum performance, download the tuned-profiles-sap-hana- 2.8.0-5.el7_4.2 or newer rpm and install it on the SAP HANA VM.

The following steps will enable the tuned profile:

```
# yum install tuned-profiles-sap-hana# systemctl start tuned
# systemctl enable tuned
# tuned-adm profile sap-hana
```

Follow SAP Note 2292690 for additional performance tuning best practices for Red Hat 7.x.

Other settings

Add the following kernel boot arguments to the default grub config located at `/etc/default/grub`:

- numa_balancing=disable
- transparent_hugepage=never
- intel_idle.max_cstate=1
- processor.max_cstate=1
- vmw_pvscsi.cmd_per_lun=254
- vmw_pvscsi.ring_pages=32

SAP system monitoring

To ensure SAP support in an SAP HANA VM environment, it's mandatory to set up the enhanced SAP system monitoring as described in SAP Note 1409604.

1. On each vSphere ESXi host, enable the SAP monitoring by navigating to **Configure** > **System** > **Advanced System Settings** and set the parameter Misc.GuestLibAllowHostInfo = 1, as shown in Figure 10.

FIGURE 10. ESXi host setting of Misc:GuestLibAllowHostinfo

2. Enable SAP monitoring on each SAP HANA VM via **VM settings** > **VM Options** > **Advanced** > **Edit Configuration** > **Add Configuration** Parameters and add the parameter tools.guestlib.enableHostInfo = TRUE as shown in Figure 11.

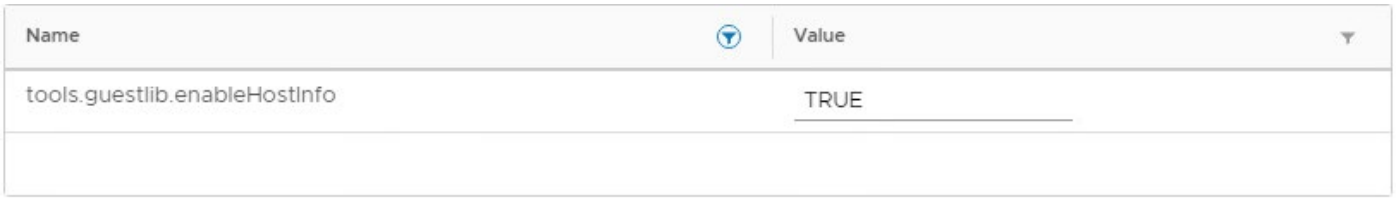

FIGURE 11. VM setting tools.guestlib.enableHostInfo

3. Make sure that the latest Open VMware Guest Tools are installed on each SAP HANA VM.

NETWORK BEST PRACTICES

This section describes various network configuration guidance for SAP HANA on HPE SimpliVity solutions.

HPE SimpliVity deployment networks

The HPE SimpliVity deployment network usually contains 4 networks known as, Management network, VM network, Federation network, and Storage network. Federation and Storage networks always use 10Gbe interface and Management and VM networks may either use 10Gbe network or 1Gbe network based on application workload needs.

NOTE

HPE SimpliVity 380 Gen10 node supports 10/25GbE – 2 Port FlexLOM as well. If a production workload needs higher bandwidth, 10/25GbE could be ordered instead of 10GbE – 2 Port FlexLOM.

Figure 12 depicts the configuration used in this Reference Architecture, where the SAP HANA database client-side network is connected to a 10 Gbe network port group called VM Network – 2, and the Serviceguard Heartbeat private network is configured on a separate port group with 10 Gbe NIC teaming interface, called as PvtNet. A separate 10Gbe network is used for SAP HANA system replication network traffic.

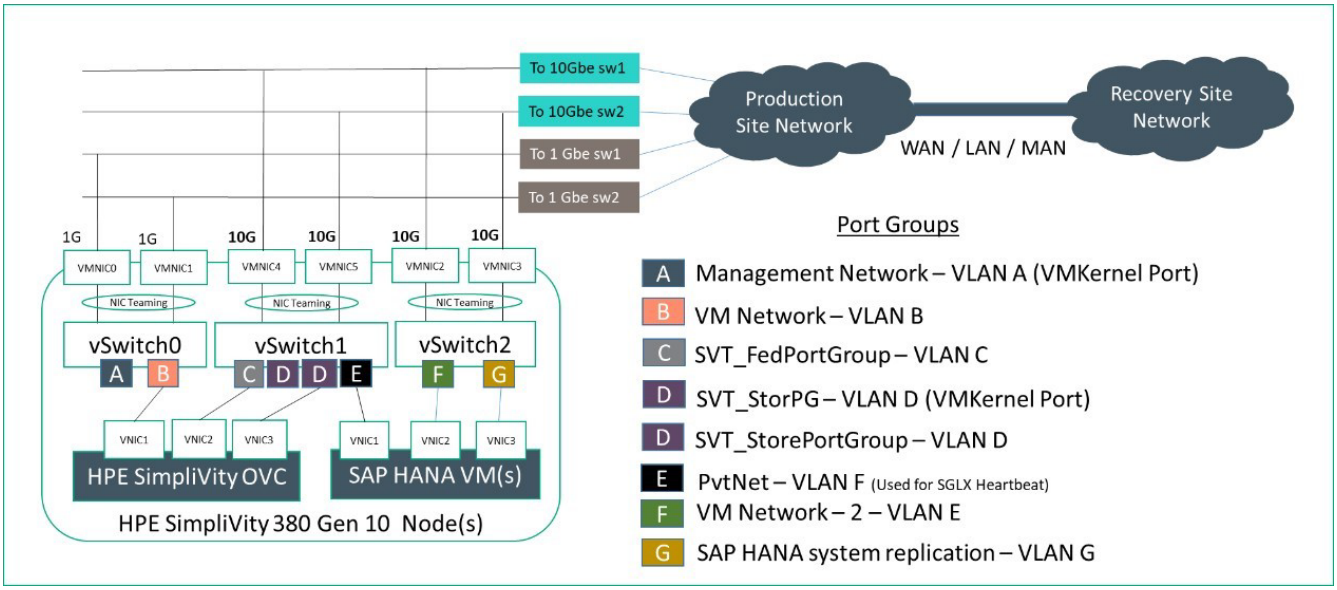

FIGURE 12. VMware vSwitch network configuration

The storage network handles the below traffic:

- HPE SimpliVity datastore access to VMs.

The Federation network handles the below traffic:

- Communication between HPE SimpliVity nodes within an HPE SimpliVity Cluster.
- VM replication between HPE SimpliVity nodes within an HPE SimpliVity Cluster.
- If a network route is available, VM level backup traffic goes between HPE SimpliVity Clusters via the Federation network within an HPE SimpliVityFederation.

NOTE
Backup traffic by default is transported over the Federation network if a network route is available between the Federation networks of HPE SimpliVity hyperconverged nodes. If there is no route available, which is generally the case with remote data centers, the backups are transported over the management network. Therefore, based on the backup schedules and rate of change of data on the VMs being protected, sufficient bandwidth should be available on the management network.

HPE SimpliVity node - NIC Teaming
NIC teaming helps to increase network capacity for the virtual switch hosting the team and provides passive failover if the hardware fails or it loses power. To use NIC teaming, you must uplink two or more adapters to a virtual switch and the following settings need to be set for every NIC team in the setup @ vCenter > Simplivity Federation > Click **Hosts** to open Objects tab > Right-click **host** and click **Settings** > Click **Networking** subtab > Click the name of the switch > Click the pencil icon to open the Edit Settings dialog box:

- **Load balancing:** Route based on the original virtual port
- **Network failover detection:** Link status only
- **Notify switches:** Yes
- **Failback:** Yes

In this configuration, a VM will have its network running on a designated physical NIC only and can failover to another NIC within the NIC teaming, in case of physical NIC or switch failure, and it does not require any special physical switch side configuration. Network bandwidth for a VM network is limited to a single physical NIC within the NIC team.

HPE SimpliVity – Arbiter service network
The Arbiter is a service that runs in a Windows system and acts as a witness to maintain quorum for an HPE SimpliVity Cluster to ensure data availability and data consistency, should an HPE SimpliVity node fail or become inaccessible. HPE SimpliVity servers communicate with the Arbiter over the management network using both UDP and TCP on port 22122. Round trip latency between the Arbiter and the SimpliVity nodes should be no more than 300 ms.

A single Arbiter can witness all HPE SimpliVity Clusters in an HPE SimpliVity Federation. The Arbiter is a dependency for HPE SimpliVity Clusters but if an Arbiter fails all SimpliVity features continue to function, and workloads remain available. All nodes within the same HPE SimpliVity Cluster must communicate with the same Arbiter.

The Arbiter service can be installed on a physical server or virtual machine and there is no need to install it on a dedicated server. The Arbiter service can be safely installed on a server running other services such as vCenter, Active Directory Server, DNS, etc., but it cannot be run on an HPE SimpliVity datastore in an HPE SimpliVity Cluster which it is witnessing.

For more information, refer to the HPE SimpliVity deployment guide.

HPE Serviceguard – Data and heartbeat network
To avoid a single point of failure, HPE Serviceguard for Linux recommends you deploy a highly available network configuration with redundant heartbeats and data networks. Use VMware NIC teaming at the host level for all networks used by the applications that run on VMs and do not use NIC teaming at the guest level. Heartbeat roundtrip network latency measured using ping should be no more than 200 ms.

NOTE
While applying the cluster configuration, Serviceguard just verifies the network requirement at the cluster node (VM) level only and does not verify at the hypervisor level. Therefore, you may see a warning message about the heartbeat network not being redundant, while applying the cluster configuration with a single heartbeat network connected to a vSwitch with NIC teaming configured. You can safely ignore the message and continue the cluster configuration.

You may also configure every subnet including the data network as a Serviceguard heartbeat network. This would avoid the above error message and also improves the heartbeat network redundancy.

HPE Serviceguard Quorum network

The Serviceguard Quorum Server provides arbitration services for Serviceguard clusters when a cluster partition is discovered. Ideally, the Quorum Server and the cluster (s) must communicate over a subnet that does not handle other traffic. This helps to ensure that the Quorum Server is available when it is needed. If this is not practical, or if the communication must cover a long distance (for example, if the Quorum Server is serving an Extended Distance cluster, such as SAP HANA system replication), heavy network traffic or network delays could cause Quorum Server timeouts. You can reduce the likelihood of timeouts by increasing the Quorum Server timeout interval; use the QS_TIMEOUT_EXTENSION parameter in the cluster configuration file.

If a subnet that connects the Quorum Server to a cluster is also used for the cluster heartbeat, configure the heartbeat on at least one other network, so that both Quorum Server and heartbeat communication are not likely to fail at the same time.

SAP HANA system replication network

In SAP HANA system replication, each SAP HANA instance communicates on the service level with a corresponding peer in the secondary SAP HANA system to persist the same data and logs as on the primary system. The services in the secondary system operate in live replication mode: that is, all secondary system services communicate constantly with their primary counterparts, replicating and persisting data and logs and typically preloading data into memory.

SAP HANA system replication must meet the following two minimal requirements regarding throughput and latency:

- **Throughput:** It must be possible to transport the size of the persistently stored data within one day from the primary to the secondary data center.
- **Latency:** The redo log shipping time for 4 KB log buffers must be less than a millisecond or in a low single-digit millisecond range – depending on the application requirements (relevant for synchronous replication only).

A network connection of 10Gbit/s between data centers is recommended by SAP for the system replication traffic.

BACKUP SOLUTION FOR SAP HANA

SAP HANA is a mission-critical application. Companies expect continuous, uninterrupted availability for their SAP HANA environment. Therefore, backup of the SAP HANA database is a critical task and can have a significant performance effect on the application that uses them. Backup windows are shrinking, while the amount of data to be backed up is increasing. Therefore, it is difficult to find a time when backups can be performed with minimal effect on business processes. The time needed to restore and recover SAP systems is a concern because downtime for SAP production and nonproduction systems must be minimized to reduce data loss and cost to the business.

HPE SimpliVity provides three backup options for customers:

HPE SimpliVity Backup + SAP HANA log backup to file system

This method recommends taking HPE SimpliVity backup on the SAP HANA VM and periodic log backup to a separate backup volume. HPE SimpliVity comes with a built-in backup capability at no extra charge. This backup creates a fully independent logical copy of the virtual machine. There is no link or dependencies to the original VM. Therefore, operations on the original VM do not affect the backup or the restored copy of the VM. HPE SimpliVity built-in backups are extremely quick because they were designed into the platform to keep the backups simple while maintaining performance. The backups are de-duplicated and compressed within the SimpliVity data virtualized platform.

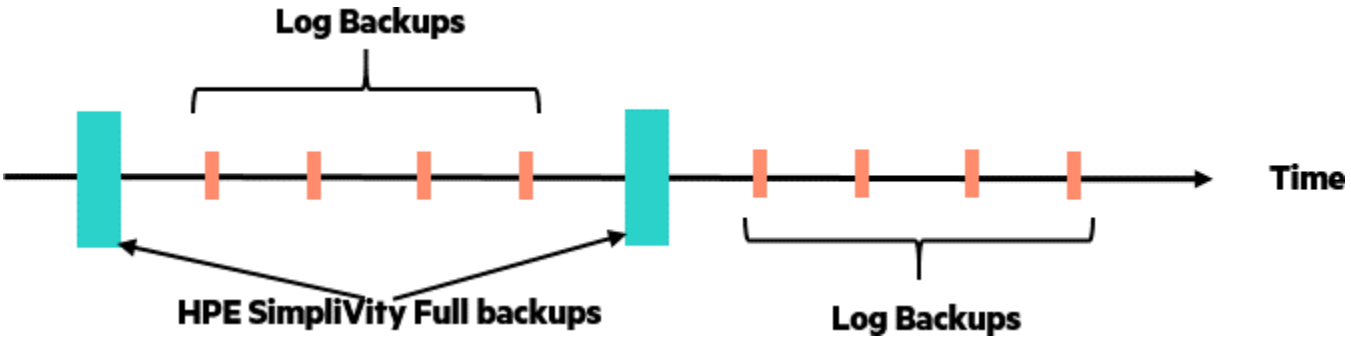

FIGURE 13. SAP HANA backup schedule

Using HPE SimpliVity backup will capture the data and log volume of the SAP HANA database. However, it does not trigger any flushing of data from the memory to the disks. Therefore, the backup will be crash-consistent. The log backup will capture all the redo log entries that are contained in the log volume. The log backup can be performed using HANA studio or cockpit. We recommend creating a backup volume (on HPE SimpliVity datastore) on the HANA VM to place the log backups in that volume. A typical backup schedule will be to take daily or hourly HPE SimpliVity full backup and take multiple log backups to reduce RPOs.

HPE SimpliVity backup with quiesce SAP HANA database

HPE SimpliVity built-in backup will capture the data and log volume of the SAP HANA database. However, it does not trigger the flushing of data from the memory to the disks to provide an application-consistent state. To put the database and log volumes in a consistent state, SAP HANA cockpit and studio provide capabilities by creating an internal database snapshot to put the database volume in a consistent state.

Currently, SAP HANA 2.0 SPS 3 supports a single-tenant database to be backed up with this snapshot backup method.

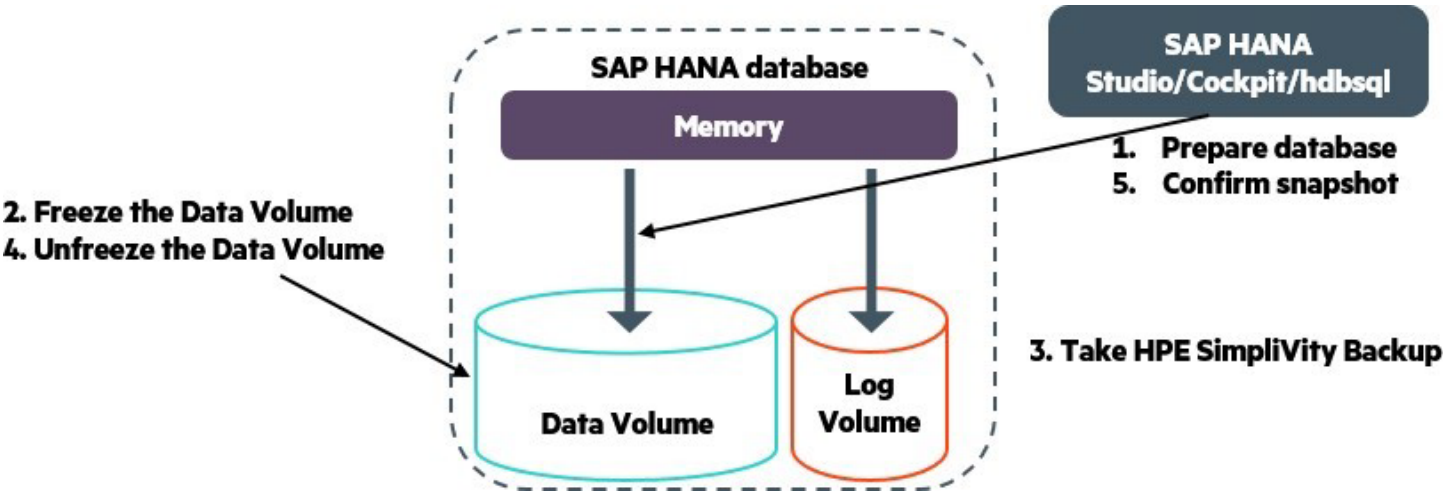

FIGURE 14. SAP HANA application consistent backup steps

Here are the steps to create an application-consistent backup:

1. First, you trigger the creation of an internal data snapshot in SAP HANA using SAP HANA Studio or cockpit or SQL commands ("prepare database"). This will prepare the data volume with an application-consistent state by creating an internal snapshot.
2. On the SAP HANA database server, freeze the data volume file system by executing "`xfs_freeze -f <Data Volume mount point>`" (for xfs file system).
3. Take a full HPE SimpliVity backup of the SAP HANA server VM.
4. On the SAP HANA database server, unfreeze the data volume file system by executing "`xfs_freeze -u <Data volume mount point>`" (for xfs file system).
5. Confirm the snapshot as successful, using SAP HANA Studio or cockpit or SQL commands. This is necessary to include the snapshot in SAP HANA's backup catalog.

HPE StoreOnce Plug-in for SAP HANA

The HPE StoreOnce Plug-in for SAP HANA is a backint for SAP HANA-compatible intelligent connector between SAP HANA and HPE StoreOnce backup target, which provides DBAs with a fast and efficient backup and recovery of SAP HANA databases. The StoreOnce Plug-in for SAP HANA is installed directly on SAP HANA database servers and integrates with SAP HANA Studio to initiate backup and restore operations. The plug-in provides an option to either deduplicate the SAP HANA database backup at the source i.e., on the SAP HANA nodes, or the target StoreOnce appliance. The advantage of source-side deduplication is that less backup data flows through the network, whereas the target-side deduplication reduces the deduplication load on the SAP HANA application servers. The plug-in communicates with the HPE StoreOnce Catalyst.

The HPE StoreOnce plugin provides the following benefits:

- Efficient resource consumption—Optimal network bandwidth and efficient backup processing time using source-side deduplication enable sending only the unique data to HPE StoreOnce, thereby optimizing network utilization, and significantly reducing bandwidth requirements.
- Reduced backup footprint—By improving the deduplication ratio, backups can be stored longer with reduced storage capacity compared to file backup implementations
- Increased cost savings—By enabling direct backups to HPE StoreOnce, no third-party data protection software is required. The HPE StoreOnce Catalyst Plug-in for SAP HANA controls the backup via backint to make a direct copy of the log backup savepoints to HPE StoreOnce, without the need to pass through a backup media server.
- Control and simplicity—HPE StoreOnce Catalyst Plug-in for SAP HANA is installed directly onto the SAP HANA nodes with only a few clicks. The SAP HANA DBAs can backup, restore, and delete their database backups from SAP HANA Studio directly. Backup is completely under the control of the application admin.

Table 7 shows the backup solution features and benefits.

TABLE 7. Backup solution comparison

Backup and restore Qualities	HPE SimpliVity Backup + SAP HANAlog backup	HPE SimpliVity Backup with quiesce SAP HANA database	HPE StoreOnce Plug-in forSAP HANA
Application/Crash consistent?	Crash consistent with log backup	Application consistent	Application consistent
RPO	Minutes	Minutes	Minutes
Cost	$	$	$$
Complexity	Low	Medium	High
SAP-certified Backup?	No	No	Yes
RTO	Minutes	Minutes	Minutes to hours
Backup data location?	Within SimpliVity Federation	Within SimpliVity Federation	StoreOnce Appliance

DISASTER RECOVERY SOLUTION FOR SAP HANA

Protecting and recovering any database environment is proving ever more challenging in the face of large and increasing data volumes, rising business demands, cost efficiency, and sustainable growth. SAP HANA is no exception. Organizations deploying large-scale SAP HANA environments struggle with protecting the rapidly growing and changing mission-critical data. According to SAP, disaster recovery is a major pillar of data center readiness in an SAP HANA environment. SAP is a mission-critical application to many enterprises and downtime on this application has a significant negative impact on the business. Such downtime not only has a financial impact, but also can affect the company's reputation, staff morale, and customer loyalty.

To develop a comprehensive disaster recovery policy, customers must understand the business application requirements and technical capabilities needed for data protection and disaster recovery. There are two metrics to measure the business application requirement:

- **Recovery point objective (RPO):** Maximum tolerable period in which operational data is lost without the ability to recover. This is your business continuity plan's maximum allowable threshold for data loss. The RPO is expressed backward in time (that is, into the past) from the point the failure occurs.
- **Recovery time objective (RTO):** Maximum permissible time it takes to recover the system after a disaster (or disruption) for system operations to resume. This objective can include the time for trying to fix the problem without recovery options, the recovery itself, and the testing of services before handing it over to the business users.

It is important to evaluate the business requirements that define the RPO and RTO because these requirements have a significant impact on the technical options and cost.

HPE SimpliVity provides the following disaster recovery solution for the SAP HANA environment.

SAP HANA system replication with HPE Serviceguard

The HPE Serviceguard extension for SAP (SGeSAP) provides an out-of-box unattended disaster recovery solution for SAP HANA deployments on the HPE SimpliVity platform. SGeSAP provides SAP-specific modules, service monitors, cluster resources, cluster deployment, and cluster verification tools as well as a shared library that makes SAP startup framework cluster-aware.

In an SGeSAP configuration for SAP HANA the following packages are configured:

- A primary package that makes sure that one of the systems works as a production system. This package is installed on the production SAP HANA server.
- A secondary package that makes sure that the other system replicates. This package is installed on the standby disaster recovery SAP HANA server.

The HANA instances are installed on the production HPE SimpliVity node and do not move between cluster nodes during package failover. Instead of failover, the mode of secondary is switched to primary and vice versa. The cluster controls client access capabilities to the production system and coordinates the takeover and role-reversal operations that promote primary instances or demote secondary instances.

Full heartbeat loss between sites makes it impossible to judge whether the remote site is still operational. Either the heartbeat loss is caused by pure interconnectivity issues that leave the remote system up or running or it is caused by a disaster that has stopped the whole remote system, uncertainty of what happened to the other site exists in both directions. This is a potentially dangerous situation if a HANA primary instance in one site does not shut down properly and a HANA secondary instance in the second site becomes the active primary replacement due to an automatically triggered takeover operation. With this, inconsistencies can occur. Thus, HANA system replication clusters must be connected to Serviceguard Quorum Services. The quorum server software is delivered as part of Serviceguard, but it must be installed on a separate server or in a separate cluster.

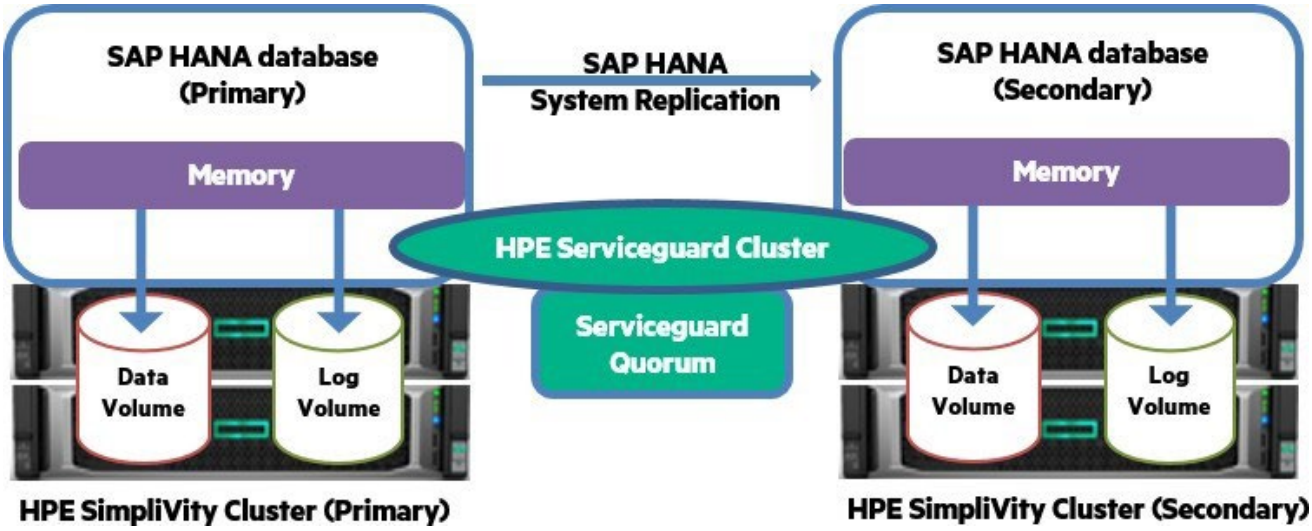

FIGURE 15. SAP HANA application consistent backup steps

SAP HANA system replication requires an active primary database system and a standby secondary database system. If the secondary database system is located near the primary system, it serves as a failover solution for planned downtime or unplanned downtime (server failures, software failures, storage corruption, and so on). If the secondary system is located remotely, it provides a disaster recovery solution. Using system replication, the database content of the production system is replicated from primary to secondary SAP HANA database system.

Depending on customer requirements, SAP HANA offers different modes for replication:

- **Synchronous:** Secondary system sends an acknowledgment back to the primary as soon as data are received and persisted to disk.
- **Synchronous in-memory:** Secondary system sends an acknowledgment back to the primary as soon as data is received (this might lead to performance increase depending on disk speed).
- **Asynchronous:** As per the design of asynchronous replication, the primary does not wait until the secondary sends an acknowledgment.

The synchronous and synchronous in-memory requires that the primary and secondary clusters should be less than 100 km and the network round trip time be less than 5ms. However, in asynchronous mode, primary and secondary sites can be beyond 100 km.

HPE SimpliVity RapidDR

HPE SimpliVity RapidDR can help reduce the complexity of a system replication DR solution by automating the complex DR steps on any level. It is designed for DR of a complete site or data center failure. It supports both unidirectional and bidirectional failover. The solution extends the inherent data efficiencies of HPE SimpliVity hyperconverged infrastructure, slashing recovery-point objectives (RPOs), and recovery time objectives (RTOs) from days or hours to minutes.

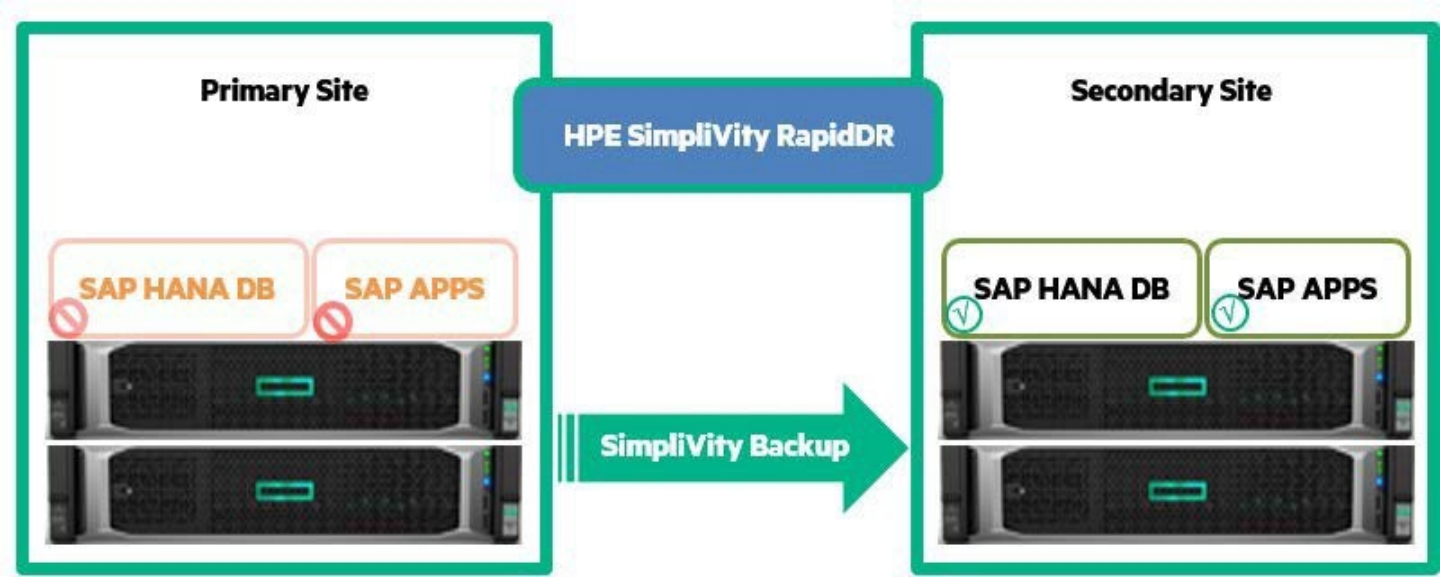

FIGURE 16. SAP HANA with RapidDR configuration

Figure 16 shows an example SAP landscape protected by HPE SimpliVity RapidDR. The VMs running on the primary site contain all the required infrastructure and SAP components such as SAP HANA database and SAP application servers, as in an SAP Business Suite implementation. The SAP VMs can be replicated depending on RPO needs through HPE SimpliVity built-in backup. HPE SimpliVity backup can be accomplished as low as every minute to replicate to the secondary site and the backup will be crash consistent. However, the recovery point objective will be as low as 5 minutes since the SAP HANA data is automatically saved from the memory to disk at a regular savepoint interval of every 5 minutes.

Here are the high-level steps that need to be configured for the solution:

1. Deploy HPE SimpliVity cluster on the production site and DR site.
2. Deploy SAP business suite components on the production site (SAP HANA server, SAP application servers, etc.).
3. Set up a backup plan on the HPE SimpliVity cluster between primary and secondary sites.
4. On the RapidDR, Create Recovery plans to failover the whole SAP business application components.
5. On the RapidDR, use the test recovery feature to validate the recovery plan to ensure that the recovery plan works.
6. When disaster happens, execute the recovery plan.
7. Ensure that the secondary site is up and running with SAP business suite components.

Building a comprehensive disaster recovery solution

Failures are inevitable. Planning a comprehensive high availability solution for SAP HANA requires an evaluation of the impact of potential failures, the company's tolerance, and requirements for different RPOs. There are multiple solutions for high availability. Many of these options can be combined to provide different levels of availability. This solution uses all the technologies mentioned in the previous disaster recovery solution to protect multi-site failures.

A comprehensive disaster recovery solution should include protection against the local site as well as a remote site. This solution will address both failures. The initial setup will have these high-level steps:

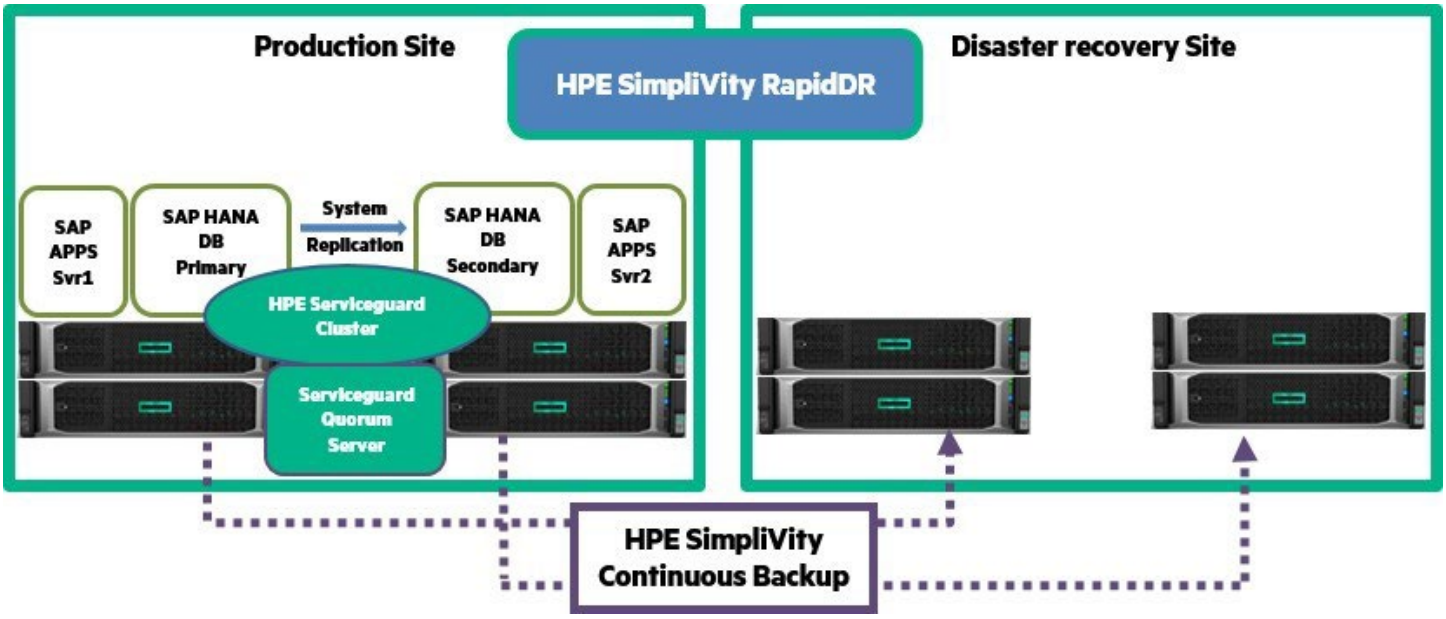

FIGURE 17. Disaster Recovery Solution

1. Deploy primary SAP HANA and application server on HPE SimpliVity clusters on the production site.

NOTE
HPE SimpliVity RAID and RAIN implementation along with the vSphere high availability feature will protect against failures at every level of the data center—from individual components, such as network adapters and HBA card to a total node failure. In the event of such hardware failure, the SAP VMs are automatically restarted in the remaining nodes of the cluster. However, this does not protect from multiple failures.

2. Deploy secondary SAP HANA and redundant application server on another HPE SimpliVity cluster on the production site.
3. Create asynchronous (or synchronous in-memory) replication between the primary and secondary SAP HANA servers using SAP system replication.

4. Setup Serviceguard cluster between the primary and secondary SAP HANA node on the production site. Setup a quorum server on the production site to monitor the SAP HANA clusters.

NOTE

The secondary HPE SimpliVity Cluster on the production site with SAP system replication and HPE Service guard will protect against multiple failures on the production site by quickly switching over the SAP HANA servers to the secondary site. For high availability on the application servers, we recommend configuring multiple application servers and place them in primary and secondary HPE SimpliVity clusters on the production site.

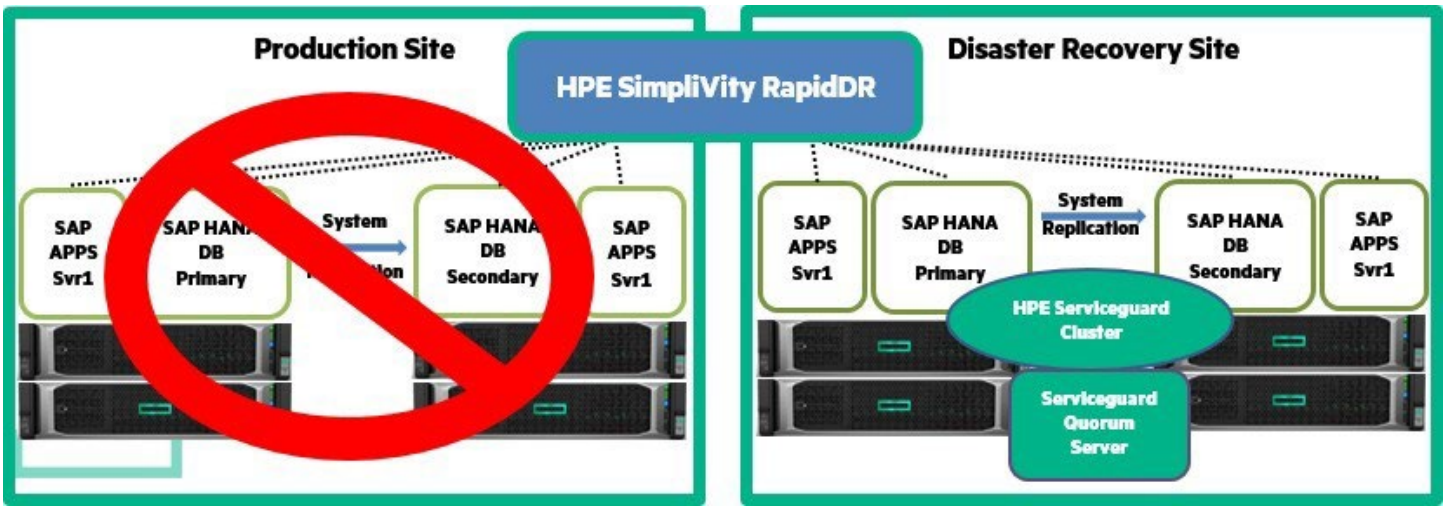

FIGURE 18. Setup Serviceguard cluster and a quorum server on the production site

5. Deploy two HPE SimpliVity clusters in the disaster recovery site and set up backup policies to backup SAP HANA and application. Ensure that primary production site backup is the point to one cluster and secondary production site backup is the point to the second cluster on the disaster recovery site.
6. Deploy HPE SimpliVity RapidDR and set up failover plans. Two plans need to be created to failover each production cluster component toits disaster recovery site cluster.

When the entire site fails, the HPE SimpliVity RapidDR will failover the entire SAP HANA and SAP application servers to the disaster recovery site. Once the failover to the disaster recovery site is completed, SAP system replication along with the service guard will protect against any failures in the cluster on the disaster recovery site.

Table 8 shows the disaster recovery solution features and benefits.

TABLE 8. Disaster recovery solution comparison

Disaster recovery Qualities	SAP HANA System replication with Serviceguard	HPE SimpliVity RapidDR	Comprehensive disaster recovery
RTO between sites	Very short	Medium	Very short
RPO	0	As low as 5 minutes	0
Cost	$$	$	$$$
Complexity	Medium	Low	High
Can failover within the Site?	Yes, with HPE SimpliVity VM restart. There will be small downtime	Yes, with HPE SimpliVity VM restart. There will be small downtime	Yes with 0 downtime.

Disaster recovery Qualities	SAP HANA System replication with Serviceguard	HPE SimpliVity RapidDR	Comprehensive disaster recovery
Can use DR resource for Test/Dev?	No. Need to dedicate SAP HANA server on the DR site	Yes.	No. Need to dedicate SAP HANA server on the DR site
Performance ramp after failover	Seconds if synchronous in-memory replication is selected.	Hours since it needs to be loaded in the memory	Failover within Site: Seconds if synchronous in-memory replication is selected. Failover between Sites: Hours if asynchronous replication is selected since it needs to be loaded in the memory.

SAP HANA ON HPE SIMPLIVITY ADDITIONAL BENEFITS

SAP HANA is a critical component of the data layer within an enterprise architecture. Many of the world's most critical applications use SAP HANA Server to store and process their data. The following scenarios provide additional benefits of deploying SAP HANA on HPE SimpliVity.

Deploying pre-production deployment of SAP HANA Database in seconds for test and development environments

Pre-production SAP HANA workloads benefit from HPE SimpliVity inherent data efficiencies and Global Unified Management capabilities. With HPE SimpliVity, the system administrator can clone VMs in just seconds, with a few mouse clicks, to easily spin up SAP HANA QA, test, or dev environments. HPE SimpliVity performs inline deduplication, compression, and optimization at inception, before data hits the disk, eliminating redundancy and overhead, making optimal use of storage capacity. If application-consistent cloning is desired, it is recommended to take the clones' with quiesce scripts.

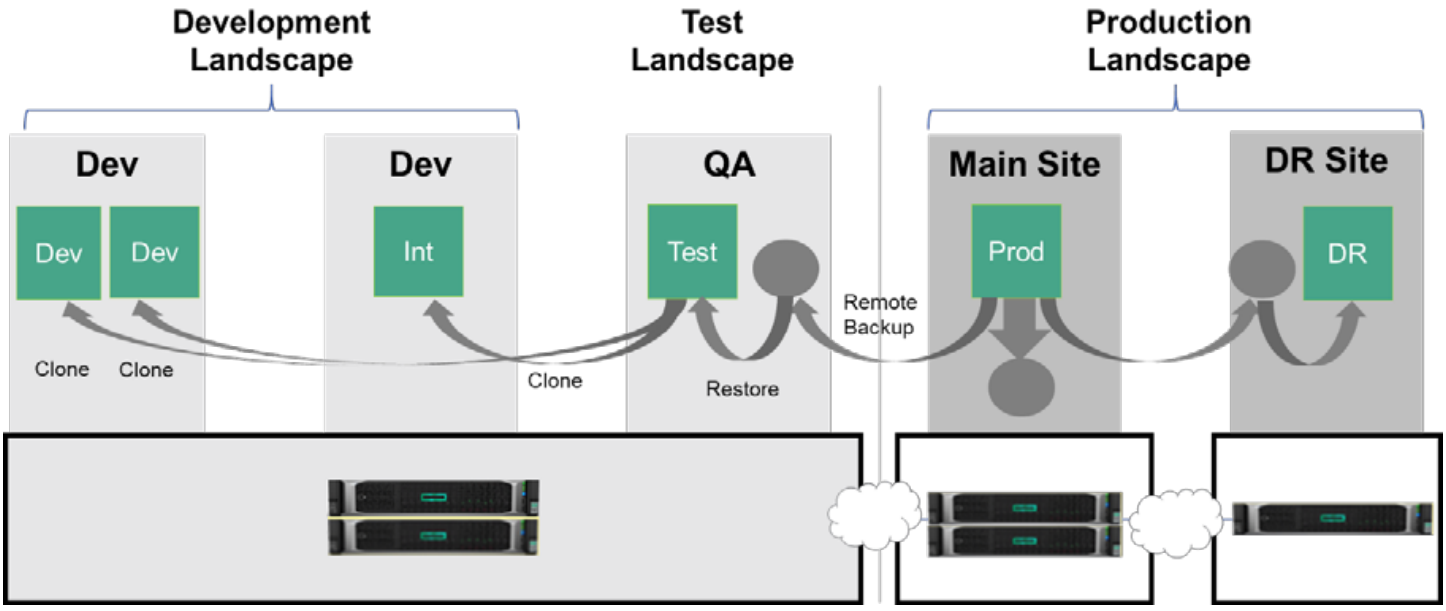

FIGURE 19. SAP HANA re-production with HPE SimpliVity clones

Below are the steps to clone a VM using the HPE SimpliVity Clone feature, and it allows you to either create an application-consistent clone or crash-consistent clone.

1. From the VMware vCenter, right-click the SAP HANA VM to be cloned.
2. Click All HPE SimpliVity Actions.
3. Select the Clone Virtual Machine option.
4. Enter a name for Clone virtual machine.
5. Click Next.

6. Select the backup option desired.
7. Click Finish to create an SAP HANA server clone.

Data center consolidation

With its revolutionary hyperconverged infrastructure platform, HPE SimpliVity offers a better approach for data center infrastructure to support SAP HANA environments. In a typical TDI model, several components require independent management like storage, server, and hypervisor. HPE SimpliVity provides SAP HANA infrastructure within a box approach with simple management. This is like the SAP HANA appliance model. However, the SAP HANA appliance is very expensive.

Because HPE SimpliVity consolidates the infrastructure and functions below the hypervisor, it eliminates the need for about a dozen discrete infrastructure and software products, which dramatically simplifies IT. Designed and optimized for the virtual environment, it enables us with dramatic improvements to the management, protection, and performance of virtual workloads – all at a fraction of the cost and extreme reduction in complexity compared to today's traditional infrastructure stack. Overall HPE SimpliVity provides appliance-like benefits with TDI lower cost to SAP HANA customers.

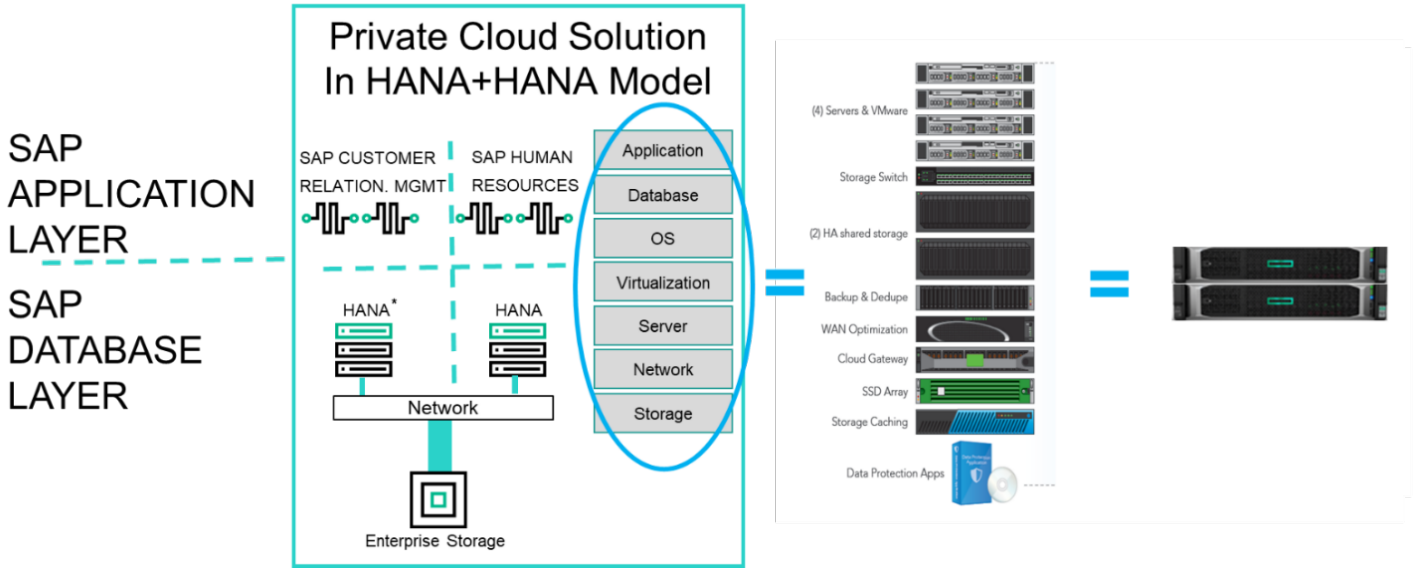

FIGURE 20. HPE SimpliVity data solutions

Unified management

Remote and branch offices of corporate headquarters, typically connected via a wide-area network (WAN), receive some level of technical support from headquarter locations. Due to the wide distribution of corporate data across sites, backup and recovery continue to challenge companies with remote and branch offices. Many remote sites cannot afford to have an expert that can do all the SAP HANA administration, backup, and protection tasks.

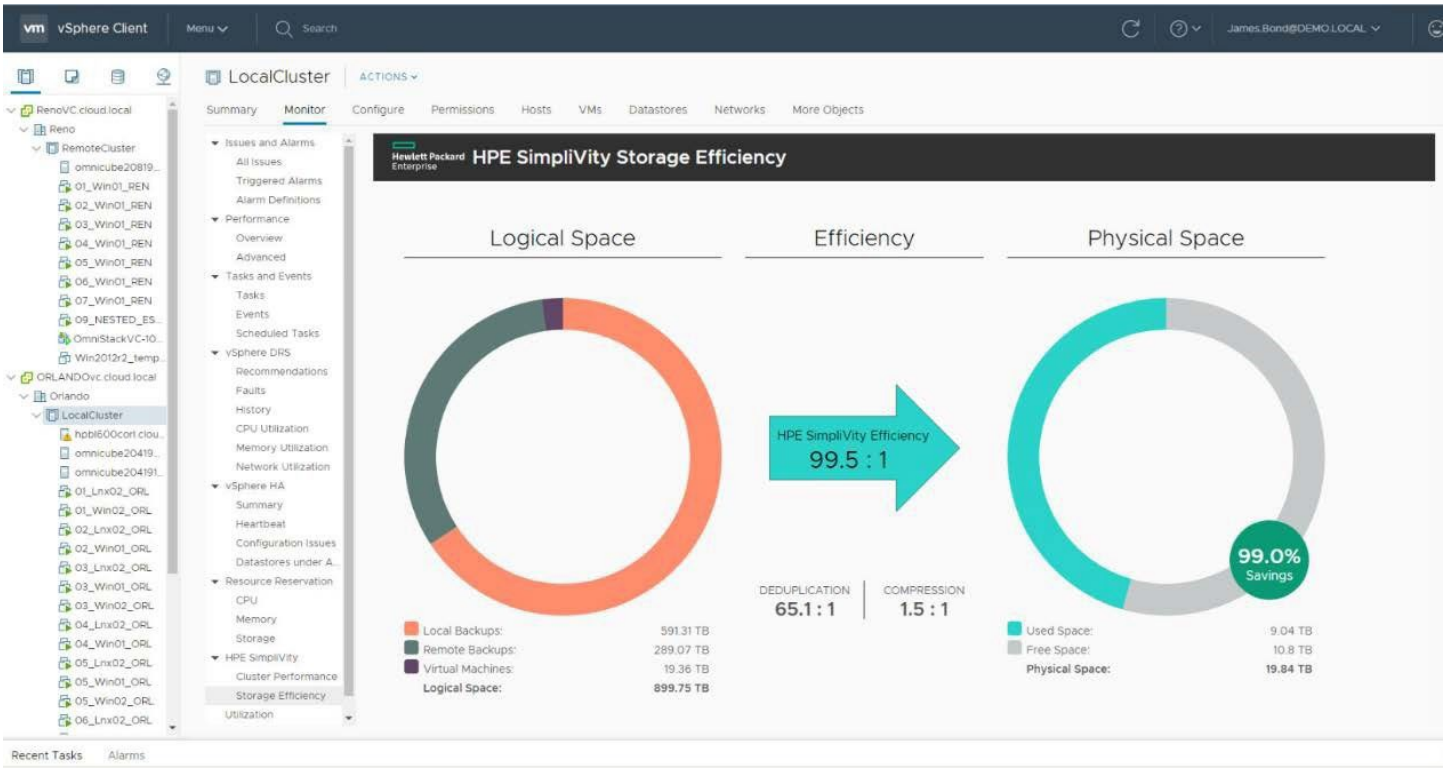

FIGURE 21. HPE SimpliVity unified management

HPE SimpliVity provides single unified management across its federation. This will enable data center administration of remote sites from primary sites. Besides, HPE SimpliVity provides single management for backup and recovery of primary and remote sites from the same VMware vCenter management. This has proven a significant reduction in cost and time to manage the infrastructure as well as perform backup and disaster recovery operations leaving more resources to focus on new application development.

SUMMARY

SAP HANA is critical to the operation of many businesses. Unexpected downtime, poor performance, or data loss can lead to lost revenue, diminished user productivity, and customer dissatisfaction. HPE SimpliVity hyperconverged infrastructure provides a highly resilient and scalable operating environment for SAP HANA. This Reference Architecture has the following features:

- Design considerations and best practices for SAP HANA on HPE SimpliVity.
- Networking best practices.
- Disaster recovery solution for SAP HANA using SAP HANA system replication with HPE Serviceguard for Linux and the advantages of integrating Serviceguard with SAP HANA system replication.
- Backup and restore options using HPE SimpliVity built-in backup and StoreOnce plugin.
- Rapid provisioning of Test or Dev SAP HANA Database using HPE SimpliVity cloning feature.
- Global Unified Management: HPE SimpliVity VM-centric approach simplified administrative tasks of Federation clusters.

- Data consolidation of SAP HANA environment with TDI mode with the appliance-like integration for ease of administration.

APPENDIX A: BILL OF MATERIALS

The following BOMs contain electronic license to use (E-LTU) parts. Electronic software license delivery is now available in most countries. Hewlett Packard Enterprise recommends purchasing electronic products over physical products (when available) for faster delivery and for the convenience of not tracking and managing confidential paper licenses. For more information, please contact your reseller or an HPE representative.

NOTE

Part numbers are at the time of publication or testing and subject to change. The bill of materials does not include complete support options or otherrack and power requirements. If you have questions regarding ordering, please consult with your HPE Reseller or HPE Sales Representative for more details. hpe.com/us/en/services/consulting.html.

TABLE A1. Bill of Materials (per HPE SimpliVity 380 Gen10 node)

Product#	Qty	Product Description
Q8D81A	1	HPE SimpliVity 380 Gen10 Node
Q8D81A 001	1	HPE SimpliVity 380 Gen10 VMware Solution
P02498-L21	1	Intel Xeon-Gold 5218 (2.3GHz/16-core/125W) FIO Processor Kit for HPE ProLiant DL380 Gen10
P02498-B21	1	Intel Xeon-Gold 5218 (2.3GHz/16-core/125W) Processor Kit for HPE ProLiant DL380 Gen10
P02498-B21 0D1	1	Factory Integrated
R4C24A	2	HPE SimpliVity 512G (8x64G) LRDIMM Kit
R4C24A 0D1	2	Factory Integrated
Q5V88A	1	HPE SimpliVity 380 for 6000 Series Large Storage Kit
P01366-B21	1	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
P01366-B21 0D1	1	Factory Integrated
804338-B21	1	HPE Smart Array P816i-a SR Gen10 (16 Internal Lanes/4GB Cache/SmartCache) 12G SAS Modular Controller
804338-B21 0D1	1	Factory Integrated
817749-B21	1	HPE Ethernet 10/25Gb 2-port FLR-SFP28 MCX4121A-ACFT Adapter
817749-B21 0D1	1	Factory Integrated
830272-B21	2	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
830272-B21 0D1	2	Factory Integrated
BD505A	1	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	1	Factory Integrated
Q8A64A	1	HPE OmniStack 2P Large SW
733664-B21	1	HPE 2U Cable Management Arm for Easy Install Rail Kit
733664-B21 0D1	1	Factory Integrated
867809-B21	1	HPE Gen10 2U Bezel Kit

TABLE A2. Bill of Materials (per HPE SimpliVity 380G Gen10 node)

Product#	Qty	Product Description
R2G96A	1	HPE SimpliVity 380 Gen10 G Node
R2G97A	1	HPE SimpliVity 380 Gen10 Node with VMware Tracking
R2G97A 0D1	1	Factory Integrated
P02504-L21	1	Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) FIO Processor Kit for HPE ProLiant DL380 Gen10
P02504-B21	1	Intel Xeon-Gold 6238 (2.1GHz/22-core/140W) Processor Kit for HPE ProLiant DL380 Gen10
P02504-B21 0D1	1	Factory Integrated
R4C22A	2	HPE SimpliVity 256GB (4x64GB) DDR4-2933 Load Reduced Memory Kit
R4C22A 0D1	2	Factory Integrated
826691-B21	1	HPE DL38X Gen10 SFF Box1/2 Cage/Backplane Kit
826691-B21 0D1	1	Factory Integrated
872475-K21	2	HPE 300GB SAS 12G Enterprise 10K SFF (2.5in) SC 3yr Wty Digitally Signed Firmware HDD
872475-K21 0D1	2	Factory Integrated
R2Z28A	4	HPE SimpliVity 4x1.92TB SFF SSD Kit
R2Z28A 0D1	4	Factory Integrated
817718-B21	2	HPE Ethernet 10/25Gb 2-port SFP28 BCM57414 Adapter
817718-B21 0D1	2	Factory Integrated
P01366-B21	1	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
P01366-B21 0D1	1	Factory Integrated
804338-B21	1	HPE Smart Array P816i-a SR Gen10 (16 Internal Lanes/4GB Cache/SmartCache) 12G SAS Modular Controller
804338-B21 0D1	1	Factory Integrated
817709-B21	1	HPE Ethernet 10/25Gb 2-port FLR-SFP28 BCM57414 Adapter
817709-B21 0D1	1	Factory Integrated

TABLE A3. Bill of Material for HPE Serviceguard for Linux

Product#	Qty	Product Description
R1T32AAE	4	HPE Serviceguard for Linux x86 Enterprise 1yr Subscription 24x7 Support PSL E-LTU (One SGLX LTU per Socket)

RESOURCES AND ADDITIONAL LINKS

Architecture Guidelines and Best Practices for Deployments of SAP HANA on VMware vSphere, vmware.com/content/dam/digitalmarketing/vmware/en/pdf/whitepaper/sap_hana_on_vmware_vsphere_best_practices_guide-white-paper.pdf

Certified and Supported SAP HANA® Hardware directory https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/#/solutions?filters=hci

SAP HANA DB: Recommended OS settings for RHEL 7 SAP Note 2292690

SAP HANA guidelines for Red Hat SAP Note 2009879

SAP HANA on VMware vSphere Virtualization, wiki.scn.sap.com/wiki/display/VIRTUALIZATION/SAP+HANA+on+VMware+vSphere HPE Reference Architectures, hpe.com/info/ra

SAP HANA on HPE SimpliVity 380 Field Guide, https://psnow.ext.hpe.com/doc/a00065065enw

SAP HCI certification best practice for HPE SimpliVity 380 and 380G, https://psnow.ext.hpe.com/doc/a00061670enw

LEARN MORE AT

hpe.com/us/en/solutions/sap-hana.html

hpe.com/info/simplivity

a50004836enw, August 2024

Hewlett Packard Enterprise