Technical Report

SAP HANA on IBM Power Systems and NetApp AFF Systems with NFS

Solution Guide and Best Practices

Tobias Brandl, NetApp
Carsten Dieterle, IBM
April 2020 | TR-4821

In partnership with

IBM

Abstract

IBM Power Systems™ are designed for data-intensive and mission-critical workloads like SAP HANA. IBM Power Systems simplify and accelerate SAP HANA deployments by providing four key capabilities: superior virtualization and flexibility, faster provisioning, affordable scaleability, and maximized uptime. The NetApp® AFF product family is certified for use with SAP HANA in tailored data center integration (TDI) projects and perfectly complements IBM Power Systems. This document describes best practices for a NAS (NFS) storage setup using NetApp ONTAP® with the AFF product family and IBM Power Systems.

NetApp®

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Introduction

IBM Power Systems are designed for data intensive and mission critical workloads like SAP HANA. Power Systems simplify and accelerate SAP HANA by providing four key capabilities: superior virtualization and flexibility, faster provisioning, affordable scalability, and maximum uptime. The NetApp AFF product family is certified for use with SAP HANA in tailored data center integration (TDI) projects and perfectly complements IBM Power Systems. This document describes the solution and some best practices for a NAS (NFS) storage setup using NetApp ONTAP with the AFF product family and IBM Power Systems for SAP HANA.

1.1 SAP HANA on IBM Power Systems

SAP HANA on IBM Power Systems has been rapidly adopted by customers over the past 48 months with more than 2500 clients, as shown in Figure 1.

Figure 1) Rapid adoption of SAP HANA on IBM Power systems.

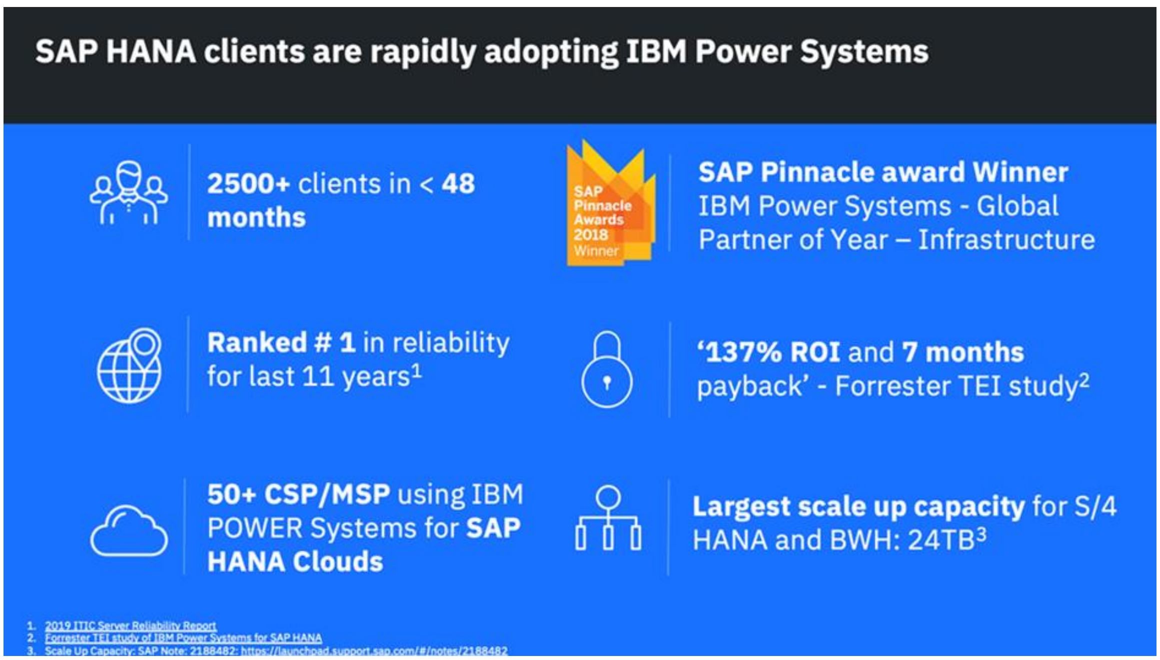

There are a number of factors that make IBM Power Systems the best platform for mission-critical SAP HANA deployments:

- Flexibility:
  - Superior built-in virtualization
  - More clients hosted on less infrastructure
  - Higher system utilization
  - Capacity on demand
  - Provision SAP HANA faster
  - Shared storage pools and superior multiple logical partition (LPAR) support
- Resiliency:
  - Redundancy built in for all critical components
  - Advanced system and memory reliability, availability, and serviceability (RAS) features

- Live partition mobility and subsequent IBM solutions like Simplified Remote Restart and PowerVM Recovery Manager
- Performance and scalability:
  - Higher core and memory bandwidth
  - Scales to large environments (scale-up, scale-out)
  - World-record, 2-tier SD benchmark [1] (https://www.sap.com/dmc/exp/2018-benchmark-directory/#/sd)
- Competitive cost:
  - Runs up to 16 SAP HANA production environments on a single server[2]
  - Supports greater workload density
  - Reduces the number of physical systems and network ports, energy consumption, and required floor space
  - Reduced cost per user

The supported IBM Power Systems are listed on SAP's certified and supported SAP HANA hardware directory.

1.2 SAP HANA on NetApp AFF Systems

NetApp solutions for SAP HANA are based on tight software integration into SAP to provide end-to-end automated workflows for SAP-relevant use cases. NetApp provides solutions for SAP that allow you to consume unique NetApp data management features with an SAP-centric view. These solutions include SAP system provision tasks as well as SAP–integrated data protection for backup and disaster recovery.

The solutions and the value proposition can be broken down into three main areas:

- Project acceleration
- Operation simplification
- Hybrid multi-cloud operations

---

[1] Result valid as of Feb 04, 2020

IBM Power Enterprise System E980 on the two-tier SAP SD standard application benchmark running SAP enhancement package 5 for the SAP ERP 6.0 application; 16 sockets / 192 cores / 1536 threads, POWER9; 3.9GHz, 8192 GB memory, 205,000 SD benchmark users running AIX® 7.2 and DB2® 10.5, Certification #: 2018055. Source: http://www.sap.com/benchmark. SAP and all SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries. All other product and service names mentioned are the trademarks of their respective companies.

[2] SAP note 2230704: https://launchpad.support.sap.com/#/notes/2230704/E

Figure 2) NetApp solution overview.

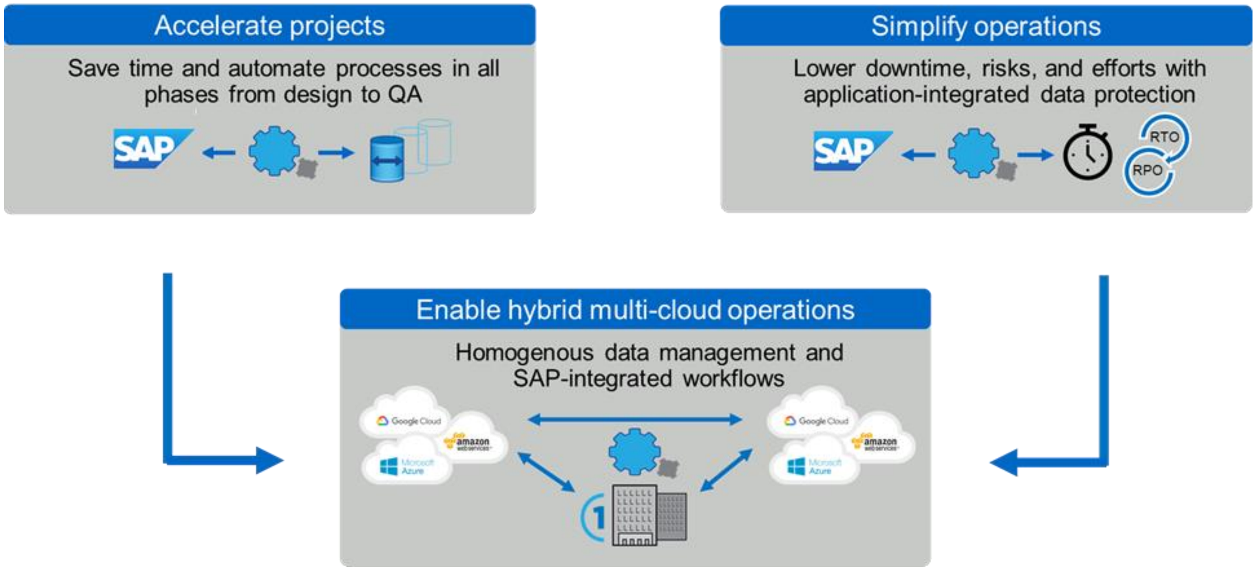

You can accelerate projects by provisioning SAP test systems in a fast and automated manner. An obvious use case is the SAP system refresh operation, for which data from a production system must be loaded into a test system. Combining storage cloning with application-integrated workflows accelerates and highly simplifies these operations. But there are also other use cases, like the handling of logical corruption, where you can spin-up a clone of your production system using any NetApp Snapshot® backup in a matter of minutes. Using this clone with data from before when the corruption occurred, you can, for example, import data that was accidently deleted into your production system. And finally, the same process can be used to test your disaster recovery workflow to make sure that you can recover from a disaster.

Simplifying operations is based on the ability to perform backup and restore operations quickly and efficiently. This is necessary to comply with service level agreements (SLAs) for production systems during normal operation, and it is also critical for upgrade projects or test cycles when backup and restore operations are part of the workflow. Instead of waiting multiple hours until the backup is finished, these projects can be accelerated, and the time saved can be used for further testing to reduce risk.

The third area of interest is support for hybrid multi-cloud operations. NetApp offers customers the flexibility to run their SAP environment on-premise, in the cloud, or distributed between on-premise and different cloud providers. Independent of the customer’s choice, the goal is to provide homogenous data management and SAP-integrated workflows.

The technological foundation for these SAP-integrated workflows is the industry-leading NetApp ONTAP storage operating system, which runs on all NetApp AFF storage systems as well as in different flavors of public cloud environments.

In addition, the NetApp AFF product family has been certified for use with SAP HANA in tailored data center integration (TDI) projects. For a complete list of NetApp certified storage solutions for SAP HANA, see the certified and supported SAP HANA hardware directory.

2 Solution and Architecture Overview

In this chapter, we discuss a reference architecture for SAP HANA on IBM Power Systems with NetApp AFF storage that has proven to be a perfect fit for large and demanding SAP HANA customer environments.

2.1 SAP HANA on IBM Power Systems with NetApp AFF Storage

The following components were tested in this solution:

- POD design for IBM Power Systems:
  - 3 x POWER9 processor-based systems in the primary datacenter, each with 192 cores and 32TB of memory (2 production, 1 backup)
  - 2 x POWER9 processor-based systems in the DR datacenter, each with 192 cores and 32TB of memory (1 production, 1 backup)
- Storage architecture:
  - 1 x NetApp AFF A700s, 72 x 7.6TB SSDs, 393TB usable – primary
  - 1 x NetApp AFF A700s, 48 x 7.6TB SSDs, 262TB usable – DR
- Network architecture:
  - 2 x ToR – Arista 7160 – LPAR customer traffic, HANA storage traffic
  - 2 x admin management switches – HMC to virtual I/O server (VIOS) direct connection, Hardware Management Console (HMC) to LPARs connection through a Customer Gateway Server (CGS)
  - 2 x FSP NW switches – HMC to FSP connection
- SAP HANA environment:
  - The POWER9 processor-based systems can host SAP HANA instances with up to 24TB of memory.
  - The AFF A700s storage system used here can host up to 28 production SAP HANA instances of any size according to the SAP HANA TDI KPIs.

Figure 3) Rack view of SAP HANA POD design with IBM Power Systems.

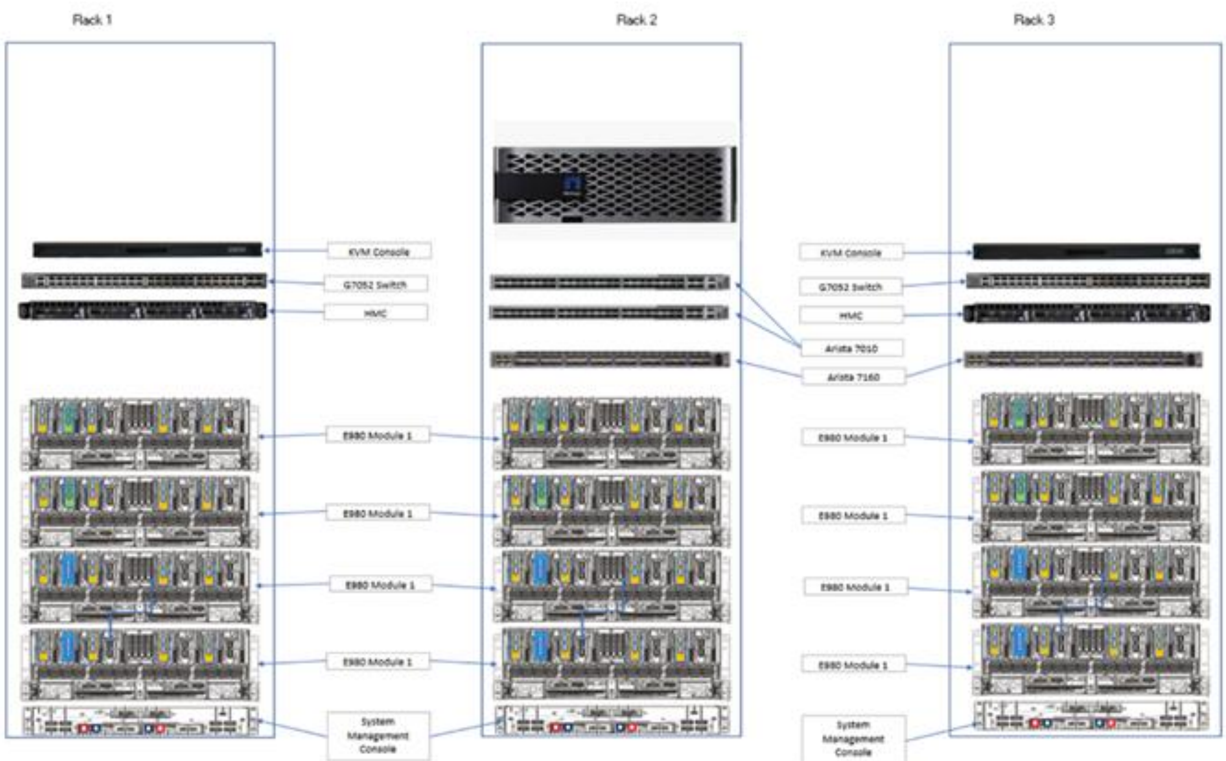

Storage for the SAP HANA databases is connected to the LPAR by NFS. The boot disks are provisioned as iSCSI LUNs on the NetApp system and exported through the VIOS as virtual disks (vSCSI disks) to the client LPARs as shown in Figure 4. A more detailed description of PowerVM VIOS best practices can be found in section 3.5.

Figure 4) Storage connectivity.

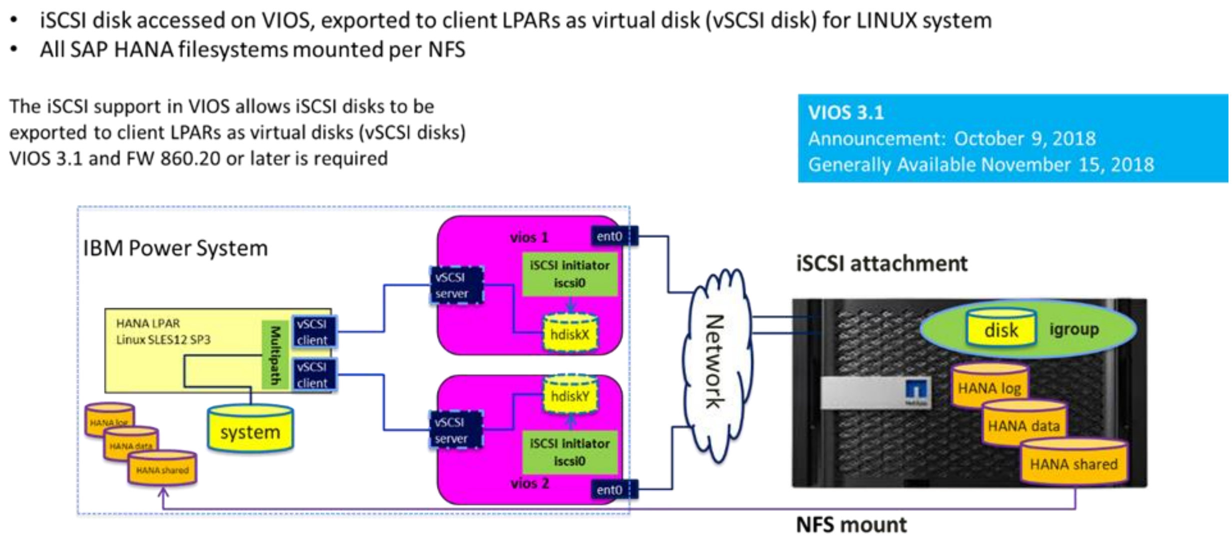

The backup Power9 server in the primary datacenter is used for maintenance tasks. This could be during maintenance of one of the two production servers, for example, when a new firmware release with a power cycle is installed. Before the firmware upgrade, the running SAP HANA instances are live migrated with Live Partition Mobility (LPM) to the Backup Power9 server. This LPM operation has no effect on the running application and the user does not get a notice. Notably, LPM must use virtualized resources for storage area network (SAN) and local area network (LAN) adapters.

To virtualize network adapters, IBM introduced a technology called single root I/O virtualization (SR-IOV). SR-IOV directly virtualizes an adapter, allowing one or more VMs or LPARs to share the adapter. Although this concept could include FC adapters, IBM currently only supports this capability for Ethernet adapters. SR-IOV virtual functions (virtual ports) can be shared by either a native OS of the correct version or by the VIOS itself. IBM offers a variety of SR-IOV-capable adapters. The number and speed of physical ports available and the number of logical ports (virtual functions) that those adapters support is adapter dependent.

SR-IOV offers various advantages, primarily concerned with system overhead and throughput. Although SR-IOV can use a VIOS, SR-IOV does not require it, so the cycles needed to perform Ethernet virtualization, which can be significant, are saved. Performing traditional virtual Ethernet with a VIOS requires copying packets from the hypervisor to the VIOS and back. This copy function limits the amount of Ethernet throughput that is possible, especially on high performance Ethernet adapters. A 10Gb Ethernet adapter shared through a VIOS cannot attain its rated speed. Sharing it directly through SR-IOV improves this throughput. These considerations become more significant as adapters move from 10Gb to 40Gb and 100Gb.

The SR-IOV adapters support QoS capabilities. A portion of guaranteed bandwidth can be assigned to each virtual function. If that bandwidth is not consumed, it is available to other virtual functions sharing the port. In addition, no QoS setting must be set. If no QoS settings are configured, then the ports are shared without any QoS capabilities.

You can configure SR-IOV adapters as dedicated or in SR-IOV mode. Although link aggregation is supported, you can only configure one virtual function per port when using this option.

- Dedicated or shared SR-IOV network adapters support virtual or native Ethernet.

- Fibre Chanel over Ethernet converged network adapter (CNA) capability is not currently supported when an adapter is in SR-IOV mode.
- LPM is limited to virtual functions (VFs) assigned to a shared Ethernet adapter or the new vNIC support
- I/O bandwidth controls—QoS setting. The desired minimum percent of physical port bandwidth assigned to a logical port.
- Multiple VFs share a physical function

A virtual network interface controller (vNIC) is a new PowerVM virtual networking technology that delivers enterprise capabilities and simplifies network management. It is an efficient, high-performance technology that, when combined with an SR-IOV NIC, provides QoS capabilities for bandwidth control at the virtual NIC level. vNIC significantly reduces virtualization overhead, resulting in lower latencies and fewer server resources (CPU or memory) required for network virtualization.

Previously, PowerVM network virtualization primarily relied on a shared Ethernet adapter in the VIOS and a virtual switch in the PowerVM hypervisor to connect the virtual Ethernet adapters with the physical network infrastructure. This approach provides a great deal of flexibility to enable network connectivity for client LPARs. However, a shared Ethernet adapter-based virtual networking solution incurs layered software overhead. For example, multiple data copies are necessary from the time a packet is committed for transmission on virtual Ethernet adapters to the time the packets are queued on the physical NIC for transmission. The same issues apply for the receipt of packets.

In the meantime, the PCI industry has developed the SR-IOV standard for hardware-based virtualization technology. An SR-IOV adapter allows creation of multiple virtual replicas of a PCI function (also known as a VF), and each VF can be assigned to an LPAR independently. The SR-IOV VF operates with little software intervention providing superior performance with very little CPU overhead. The host Ethernet adapter (HEA), introduced with POWER6 based systems, was an early implementation of such a hardware virtualization solution.

In 2014, support was added for SR-IOV adapters for select models of POWER7+ systems and more recently for POWER8 based systems. Although a dedicated SR-IOV VF provides a significant performance advantage, this configuration does not allow LPM, which can be a major drawback. With this new technology, LPM is supported for SR-IOV VFs that are assigned to vNICs. This process is made possible because the SR-IOV VF is assigned to the VIOS directly and is used by the client LPAR. Because the SR-IOV VF or logical port resides in the VIOS instead of the client LPAR, the LPAR is LPM-capable.

Besides the optimized data path, the vNIC device supports multiple transmit and receive queues like many high-performance NIC adapters. These design points enable vNIC to achieve performance that is comparable to a direct-attached logical port, even for workloads dominated by small packets. Link aggregation technologies such as the IEEE 802.3ad/802.1ax Link Aggregation Control Protocol (LACP) and active-back approaches (for example, AIX Network Interface Backup (NIB), IBM i VIPA, or Linux Active-Backup bonding mode) are supported for failover with some limitations. In the case of LACP, the backing logical port must be the only VF on the physical port. This restriction is not specific to vNIC; it applies to the direct-attached VF as well.

Figure 5) Network virtualization - SR-IOV with vNIC technology.

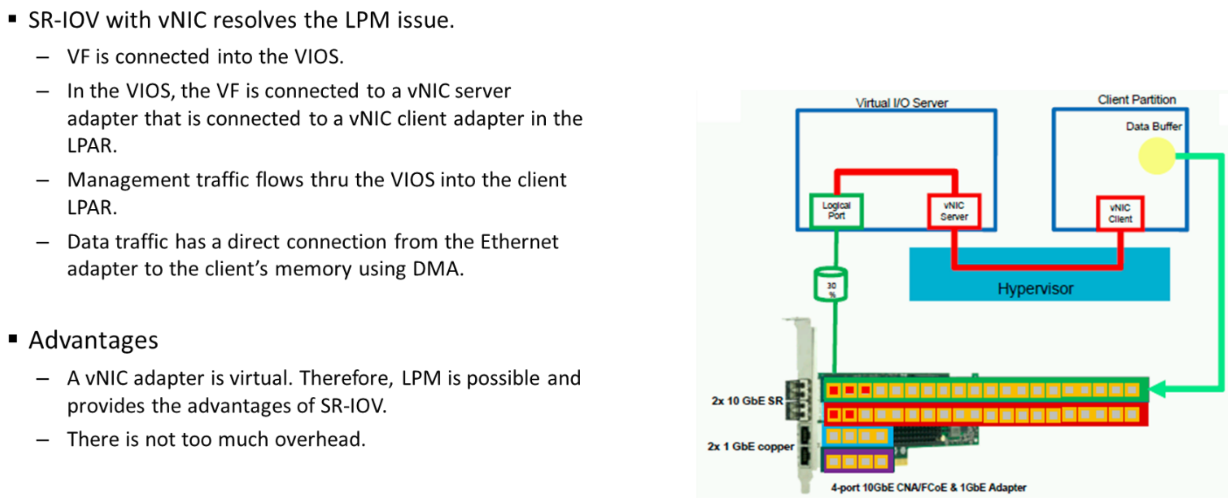

Note: For additional information, see the [IBM Power Systems SR-IOV Technical Overview and Introduction].

For network connectivity through NFS and iSCSI disk attachment, SR-IOV-capable 25GB network adapters are used in combination with the vNIC technology described above. The required network bandwidth for iSCSI disks is not that challenging, therefore we use these adapters for iSCSI connectivity as well for LPM traffic when needed.

The most critical network bandwidth in an SAP HANA installation is the data traffic in the SAP HANA data and SAP HANA log filesystems. For these filesystems, we define separate two-port 25GB network adapters and use the same vNIC technology to connect these filesystems with NFS across two VIOS servers for redundancy.

2.2 SAP HANA Disaster Recovery

Studies have shown that business application downtime has a significant negative effect on enterprise businesses. Such downtime not only has a financial effect, but it can also affect the company's reputation, staff morale, and customer loyalty. Unfortunately, not all companies have a comprehensive disaster recovery policy.

Running SAP HANA on NetApp storage gives customers access to additional features that extend and improve the built-in data protection and disaster recovery capabilities of SAP HANA. There are two main NetApp ONTAP technologies that can play an important role in the overall disaster recovery planning, NetApp SnapMirror® and NetApp MetroCluster™. Figure 6 shows how those technologies compare to HANA System Replication in terms of RPO and RTO.

Figure 6) Disaster recovery options.

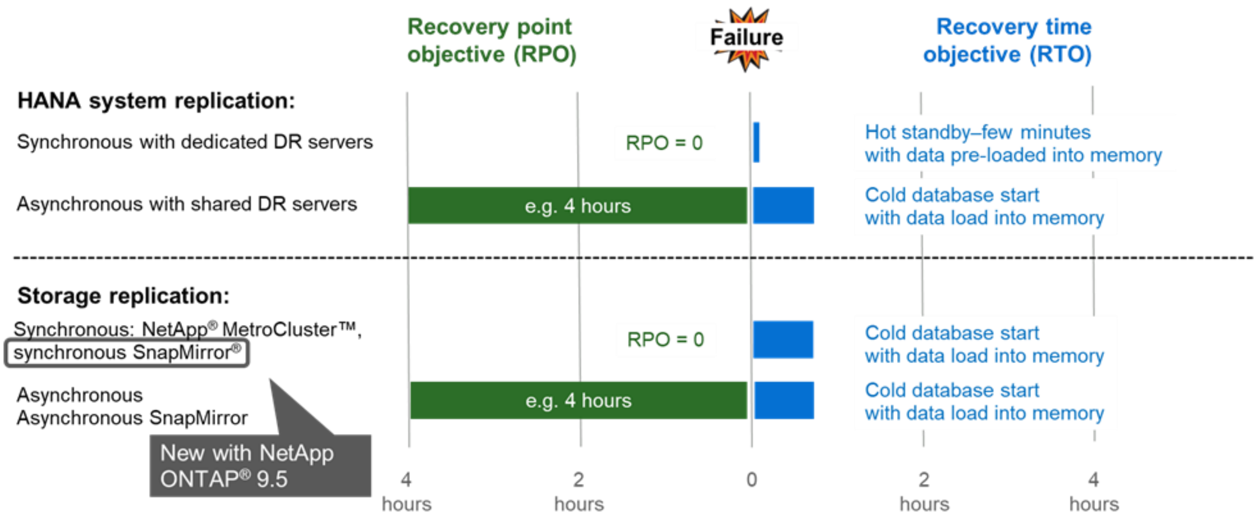

Having a closer look at different SnapMirror configurations, it becomes obvious that, depending on how SAP HANA data volumes, log volumes, and log backups are replicated, different RPO times are achievable (Figure 7).

Figure 7) Disastery recovery comparison using NetApp SnapMirror replication.

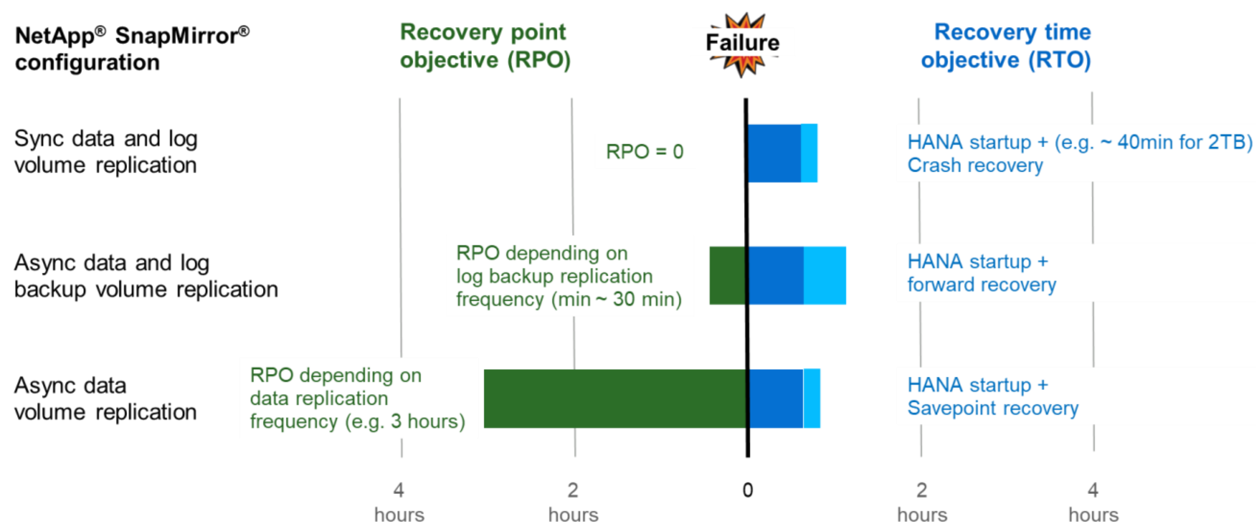

To develop a comprehensive disaster recovery policy, you must understand the business application requirements and technical capabilities needed for data protection and disaster recovery.

This overview section briefly explained the NetApp replication options for disaster recovery protection so that you can select metrics that support your business needs. For a detailed description of disaster recovery options using NetApp SnapMirror technology, see TR-4646: SAP HANA Disaster Recovery with Asynchronous Storage Replication.

2.3 SAP HANA Backup and Recovery

Companies require continuous, uninterrupted availability for their SAP applications. They expect consistent performance levels in the face of ever-increasing volumes of data and the need for routine maintenance tasks such as system backups. Performing backups of SAP databases is a critical task and can have a significant performance effect on the production SAP system.

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 a 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 the cost to business.

The following points summarize the challenges facing SAP backup and recovery:

- **Performance effects on production SAP systems.** Typically, traditional copy-based backups create a significant performance drain on production SAP systems because of the heavy loads placed on the database server, the storage system, and the storage network.
- **Shrinking backup windows.** Conventional backups can only be made when few dialog or batch activities are in process on the SAP system. The scheduling of backups becomes more difficult when SAP systems are in use around the clock.
- **Rapid data growth.** Rapid data growth and shrinking backup windows require ongoing investment in backup infrastructure. In other words, you must procure more tape drives, additional backup disk space, and faster backup networks. You must also cover the ongoing expense of storing and managing these tape assets. Incremental or differential backups can address these issues, but this arrangement results in a very slow, cumbersome, and complex restore process that is harder to verify. Such systems usually increase recovery time objective (RTO) and recovery point objective (RPO) times in ways that are not acceptable to a business.
- **Increasing cost of downtime.** Unplanned downtime of an SAP system typically affects business finances. A significant part of any unplanned downtime is consumed by the requirement to restore and recover the SAP system. Therefore, the desired RTO dictates the design of the backup and recovery architecture.
- **Backup and recovery time for SAP upgrade projects.** The project plan for an SAP upgrade includes at least three backups of the SAP database. These backups significantly reduce the time available for the upgrade process. The decision to proceed is generally based on the amount of time required to restore and recover the database from the previously created backup. Rather than just restoring a system to its previous state, a rapid restore provides more time to solve problems that might occur during an upgrade.

SAP HANA supports different methods for database backups:

- File-based backup to a file system, typically an NFS share
- Backups using the SAP HANA BACKINT API and certified third-party backup tools
- Storage-based NetApp Snapshot™ copy backups

To choose the best method, you must understand the infrastructure and performance effect as well as the additional required features of the selected HANA backup method. The following subsections provide a few examples.

File-Based and Stream-Based Backups

With file-based backups or stream-based backups using the BACKINT API, the SAP HANA database server reads the data from the primary storage. The database server then either writes the data to an NFS share or streams the data to a backup server using the third-party backup tool. Both approaches have a significant effect on the performance of the SAP HANA database in the following ways:

- Additional CPU load at the SAP HANA database server
- Additional I/O load at the primary storage
- The load on the backup network

In addition, the backup run time, specifically for larger databases, can also be significant, resulting in lower operation speed during backup. The restore and recovery process can also be a challenge because of the long run time. The SAP HANA Administration Guide lists these aspects in a comparison of the different SAP HANA backup types.

Storage-Based Snapshot Backups

NetApp storage-based Snapshot backups address the challenges discussed earlier. Independent of the size of the SAP HANA database, a Snapshot backup is executed within a few seconds instead of hours. The backup is executed at the storage layer, and there is no effect on the performance of the SAP HANA database. Also, the restore process occurs in a matter of seconds, which has a significant effect on the RTO if a restore operation is required.

NetApp SnapCenter with the SAP HANA plug-in can facilitate an automated and fully integrated HANA backup based on Snapshot copies, including the automation of SAP HANA block integrity checks by using file-based backups.

SnapCenter also handles the scheduling and housekeeping of backups on the storage and within the SAP HANA backup catalog based on flexible, configurable retention policies. In addition, non-database files can be secured with SnapCenter. Figure 8 shows an overview of the SnapCenter solution.

Figure 8) SnapCenter overview.

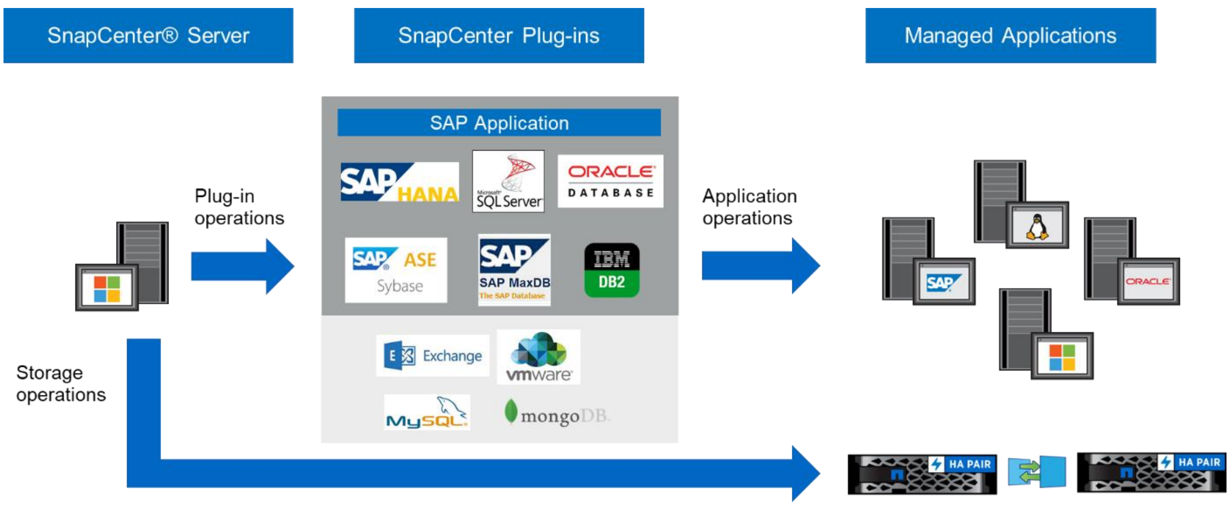

The workflow for backing up a SAP HANA database with SnapCenter is shown in Figure 9.

Figure 9) SAP HANA backup workflow with SnapCenter.

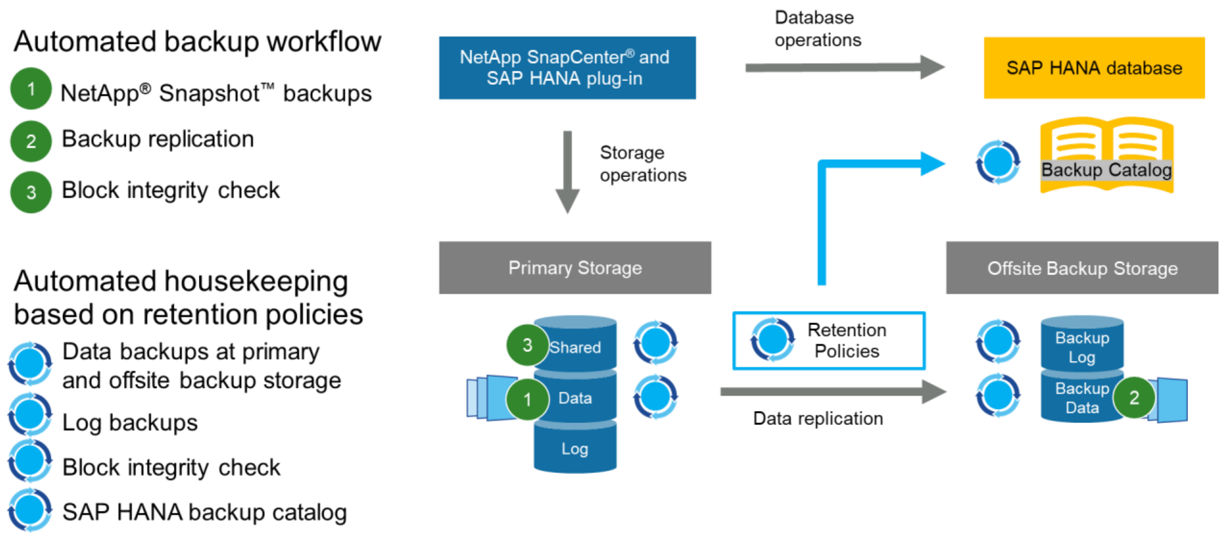

For more information, see the [TR-4614: SAP HANA Backup and Recovery with SnapCenter installation and configuration guide](#).

Note: The SAP HANA plug-in for SnapCenter is supported on Windows and Linux (Intel x86) operating systems, but not on IBM Power. However, a central communication host that has the SAP HANA plug-in installed and that communicates with the different SAP HANA databases running on IBM Power is supported and recommended in such environments. See TR-4614 for more details about the central communication host for the SAP HANA plug-in in SnapCenter.

2.4 SAP HANA Lifecycle Management with SAP Landscape Management

SAP Landscape Management (LaMa) enables SAP system administrators to automate SAP system operations, including end-to-end SAP system copy and refresh operations. SAP LaMa is an SAP software product that allows infrastructure providers such as NetApp and IBM Power to integrate their products into an SAP environment. With such integration, you can use the value added by NetApp and IBM Power from within the SAP LaMa GUI.

NetApp offers the NetApp Storage Services Connector (SSC), which allows SAP LaMa to directly access technologies such as NetApp FlexClone® and NetApp SnapMirror data replication. These technologies help minimize storage use and shorten the time required to create SAP system clones and copies.

IBM Power is deeply integrated as well in SAP Landscape Management operations, SAP LaMa is supported with IBM storage solutions in combination with IBM PowerVC.

These capabilities are available to customers who run their own on-premises data center or private cloud. They are also available to customers planning a hybrid cloud solution by integrating public cloud providers such as Amazon Web Services (AWS), Microsoft Azure, or Google Cloud Platform into their overall data center concept. SAP LaMa together with NetApp SSC can bridge the gap between on-premises systems and the cloud by defining clear data ownership and providing the tools to move systems seamlessly between clouds.

SAP LaMa can be used to manage SAP systems that run on any kind of infrastructure that supports SAP applications, including the following:

- Standard physical servers in an on-premises data center
- Cloud-like infrastructure that uses converged systems such as the FlexPod® platform, the Cisco and NetApp data center solution.

- Virtual environments such as PowerVM, VMware, Hyper-V, and Linux KVM
- Cloud infrastructures such as AWS, Microsoft Azure, Google Cloud Platform, and IBM Cloud

You must meet the following prerequisites to manage SAP systems with SAP LaMa:

- SAP LaMa must communicate with the SAP Host Agent running on the physical or virtual host. SAP Host Agent is installed automatically during SAP system installation. However, it can be configured manually to include hosts in SAP LaMa management that do not run SAP software, such as web servers.
- To communicate with NetApp storage systems, SAP LaMa must communicate with NetApp SSC. The NetApp SSC is a small Java-based application that runs on a central host in the SAP environment. NetApp recommends using an Intel x86-based environment for the SSC installation, because it has not yet been qualified to run on a Power Linux OS. For more information about NetApp SSC, see the NetApp SSC for SAP LaMa site.
- In cloud-like multitenant environments, SAP LaMa must be able to reach all systems by using host names with DNS name resolution. This requirement also applies if SAP LaMa extends beyond data center boundaries by integrating external systems hosted at a service provider or in a public cloud extension.
- To use all SAP LaMa features, install systems following adaptive design principles (see the SAP Landscape Management community page for more details). However, a classically installed SAP system can benefit from the central management functions in SAP LaMa too.

The following figure shows a typical on-premises data center setup. SAP LaMa can integrate any SAP system, including classical NetWeaver-based SAP systems and SAP HANA running on all supported operating systems (for example, IBM Power Systems using NFS attached storage).

Figure 10) SAP LaMa system landscape.

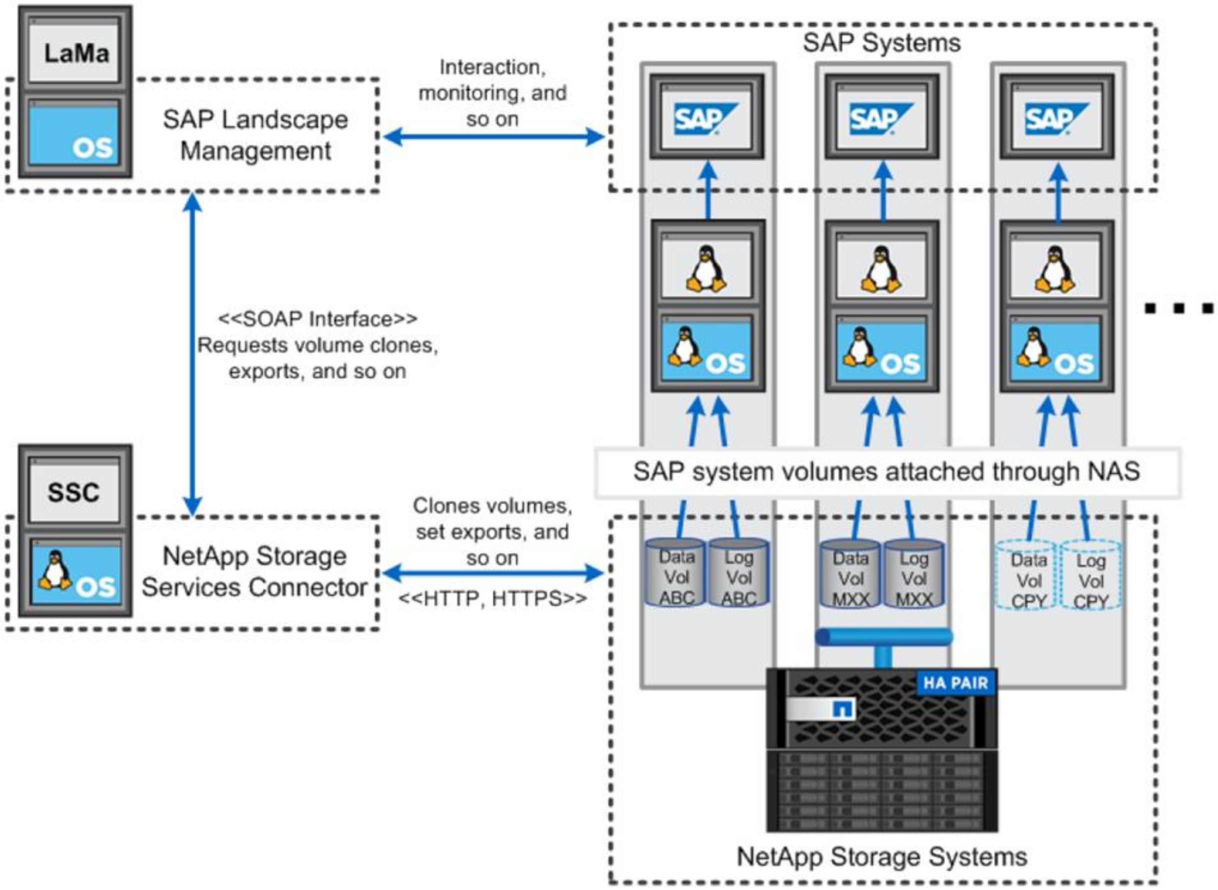

3 Infrastructure Sizing and Configuration Best Practices

3.1 Compute Sizing Overview

When SAP HANA on IBM Power Systems was introduced in August 2015, there were a number of restrictions in the deployment of SAP HANA instances. After the initial start of SAP HANA on IBM Power Systems, the compute sizing on IBM POWER8 processor-based systems was performed with a core-to-memory sizing approach, which was valid until the announcement of TDI Phase 5 in October 2017. With TDI Phase 5, SAP and IBM sunset the core-to-memory sizing and adjusted the compute to the workload, resulting in a SAP Application Performance Standard (SAPS) based sizing as described in the next section.

3.2 TDI Phase 5 – SAP Application Performance Standard-Based Workload Sizing

SAP made a important announcement in October 2017 at SAP TechEd in Las Vegas: TDI Phase 5 SAP Application Performance Standard (SAPS)-based sizing for SAP HANA workloads. Since its debut, SAP HANA has been sized based on a strict memory to core ratio determined by SAP based on workloads and platform characteristics (for example, the generation of processor, MHz, interconnect technology, and so on). This made sense in the early days when not much was known about the loads that customers were likely to experience, and SAP did not have much experience with SAP HANA on IBM Power

Systems. Over time, with a few exceptions, CPU loads tended to be far lower than the ratios might have predicted.

IBM Power Systems customers have been almost universal in their concern about poor utilization. These customers have historically driven high utilization, often over 65%. Power has up to five times the memory bandwidth per socket of x86 systems (without compromising reliability) and very wide and parallel interconnect paths with very low latencies. SAP HANA has never been offered as an appliance on Power Systems, instead being offered only using a Tailored Datacenter Infrastructure (TDI) approach. As a result, customers typically view on-premise Power Systems as a utility. They should be able to use them as they see fit and drive as much workload through them as possible while maintaining the SLAs that their end users require.

IBM's virtualization solution, PowerVM, enabled SAP customers to run multiple production workloads (up to 16 on large systems) or a mix of up to 15 production workloads. They could use a shared pool of CPU resources on which a mix of VMs could run, including non-production HANA, application servers, and non-SAP and other OS workloads (for example, AIX and IBM i). In this mixed mode, some of the excess CPU resource not used by the production workloads could be used by the shared-pool workloads. This helped drive up utilization somewhat, but not enough for many. For the latest updates for SAP HANA on IBM Power Systems, see SAP notes 2055470 and 2188482.

Using a SAPS-based methodology is a good start and might require fewer cores for the same workload than would have been calculated previously based on the memory-to-core ratio. You now have the option to reduce CPU resource allocations to specific workloads, which means that an even more significant reduction in CPU allocation might be possible. This can result in a more efficient use of CPU resources, more capacity available to other workloads, and/or the ability to size systems with less resources to drive down system costs. Either way, SAPS can improve the TCO by reducing the number and size of systems with their associated datacenter and personnel costs.

The newest sizing report for SAP BW/4 HANA (at the SAP support portal) provides the customer a SAPS-based result and the CPU category. CPU category L reflects the old core-to-memory (CtM) ratio. Category M is 50% of the CtM, and category S is equal to 25% of the CtM.

Keep in mind that the memory requirement is still the primary sizing factor for a HANA database. For more information, see the SAP HANA on Power Planning and Operations Guide.

3.3 SAP HANA and Shared Processor Pool support

In August 2019, SAP released support for operating SAP HANA production systems in shared-processor-pool-based LPARs on IBM Power Systems. Prior to this, shared pool LPARs were supported for non-production HANADB workloads (quality assurance, test, and so on), as well as associated production workloads (for example, the application server tier). Before this support announcement, up to 16 production HANA DBs in dedicated (or dedicated donating) LPARs could co-reside on the same physical Power System. This capability was provided to minimize any negative effect on production HANA database performance caused by CPU resource sharing, increased non-uniform memory access (NUMA) effects by dynamically assigning CPU resources, or cross-LPAR interference (noisy neighbors).

Meanwhile, IBM and SAP teams have performed testing and optimizations at the OS and application level to show that customers can also exploit the shared processor pool capability of PowerVM for performance-sensitive production SAP HANA databases. This section describes the preferred use cases and effects when running SAP HANA LPARs in shared processor pools versus in dedicated mode. It also provides configuration guidance from numerous validation runs. SAP Note 2055470 HANA on POWER Planning and Installation Specifics was updated with a section about running SAP HANA in shared processor pool partitions (SPLPARs).

Shared Processor Pool Partitions in PowerVM

SPLPARs group the cores that are allocated in one shared processor pool for each Power System. The shared pool can be logically segmented into multiple virtual shared pools, which is primarily used for software licensing optimization. For an SAP HANA database, it does not matter whether it resides in the default shared pool or in one of the virtual shared pools. Therefore, we do not distinguish between the two possibilities in this report.

One partition has a fractional number of cores (min = 0.05 cores, max = cores in pool). Their sum per LPAR is called the capacity entitlement and determines the minimum guaranteed compute resources. The virtual processor parameter limits the number of physical cores an uncapped SPLAPR can consume above the entitlement if free CPU cycles are available in the pool. In conjunction with simultaneous multi-threading (SMT) mode, this parameter defines the highest degree of compute parallelism for an application in an SPLPAR.

SPLPARs can either share their unused cores or alternatively can borrow a fractional number of cores if there is idle capacity available in the resource pool. This happens automatically within microseconds and provides elasticity within a shared pool. Borrowing CPU cycles can be enabled (uncapped) or disabled (capped). If uncapped, an LPAR can ideally access all cores within the (virtual) shared pool. Distribution of free compute cycles among LPARs is determined by the individual LPAR weight (weighting factor), so that LPARs running critical workloads (for example, HANA production) can consume more idle CPU cycles than less important systems.

From an operational perspective, it is important that the SPLPAR parameters (entitlement, capping, or weight) can be dynamically adjusted. Therefore, entitled capacity can be adjusted while SAP HANA instances stay fully operational. Dedicated LPARs cannot re-use idle cycles from others LPAR cores. Resource adjustments can also be performed by adopting LPAR parameters. These changes require a HANA restart so that HANA recognizes any topology changes. Dedicated LPARs provide the advantages of minimized execution latency and regular memory-CPU affinity. Dedicated-donating LPARs can donate idle CPU cycles into a shared pool, without being part of the hypervisor's time-slicing process with the potential latencies introduced by it. Therefore, they are the preferred choice for large production HANA systems that require a higher level of resource isolation and can share compute resources to other LPARs during periods of low utilization. The following picture shows the different possible LPAR modes in a PowerVM environment.

Figure 11) LPAR layout.

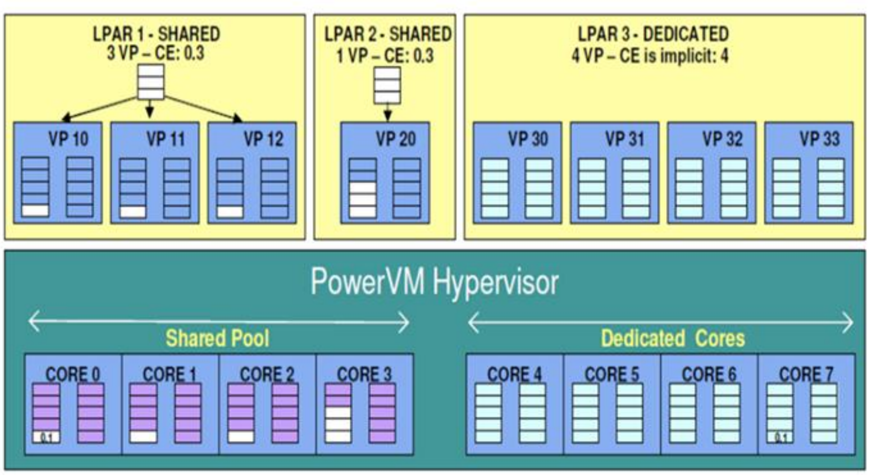

For initial production support, SAP allows shared LPAR sizes that fit onto a single socket of an IBM Power server and its attached memory, which Linux OS reports as an NUMA node. This configuration preserves memory locality to physical cores within shared LPARs. With increasing feedback from customers and cloud service providers (CSPs), the number of supported configurations will be expanded. The required test series for POWER9 systems were completed in December 2019, and you can also use production systems in SPLPARs, including pure OLAP scenarios. Regularly check SAP Note 2055470 for the most current status.

In the last four years, clients with IBM Power Systems were able to virtualize with more granularity and scalability, but the cores assigned to each of the SAP HANA production VMs on IBM Power Systems were in dedicated donating mode. However, this form of virtualization is still much better than what our x86 competition could offer with only limited support for dedicated cores. Shared processor pools are not a new feature for IBM Power Systems. Rather, it has been available for the last few generations of Power processor-based systems. However, in the world of SAP HANA deployment, sharing CPU cycles dynamically from a processor pool is new, and no other industry player can provide this level of flexibility.

What are the advantages of shared processor pools for SAP HANA workloads?

With this support announcement from SAP, clients can now use the dynamic resource elasticity of the processors in a shared processor pool. This process happens autonomously according to customer-defined rules, and it can be adjusted dynamically for active workloads. This process also improves the TCO because it improves resource utilization and agile deployment. IBM and SAP teams have performed intensive testing, and our clients that were part of early adoption program have seen great benefits using shared pool LPAR technology.

IBM has validated both POWER8-based and POWER9-based servers.

Shared processor pools further simplify and optimize the deployment of SAP HANA workloads. The client can now share the processor cores between not only SAP HANA production environments but also with SAP application servers, non-production workloads, and other non-SAP workloads. Clients can also define the priority of the workloads, and the system autonomously manage the CPU resources, resulting in better consolidation and utilization of resources.

3.4 Storage Sizing Overview

The following section provides an overview of the required performance and capacity considerations for sizing a storage system for SAP HANA.

Note: Contact NetApp or your NetApp partner sales representative to support the storage sizing process and to create a properly sized storage environment.

SAP has defined a static set of storage key performance indicators (KPIs). These KPIs are valid for all production SAP HANA environments independent of the memory size of the database hosts and the applications that use the SAP HANA database. These KPIs are valid for single-host, multiple-host, Business Suite on HANA, Business Warehouse on HANA, S/4HANA, and BW/4HANA environments. Therefore, the current performance sizing approach depends only on the number of active SAP HANA hosts that are attached to the storage system.

Note: Storage performance KPIs are only required for production SAP HANA systems.

SAP provides a performance testing suite called SAP HANA Hardware and Cloud Optimization Tools (HCOT; formerly called Hardware Configuration Check Tool [HWCCT]). This tool must be used to validate storage performance for the number of active SAP HANA hosts attached to the storage.

The storage vendor defines the maximum number of SAP HANA hosts that can be attached to a specific storage model while fulfilling the required storage performance KPIs from SAP for production SAP HANA systems. As an example, the A700s system used in the reference architecture in chapter 2.1 can support up to 28 SAP HANA nodes in production according to the KPIs.

The capacity requirements for SAP HANA are defined in the [SAP HANA Storage Requirements Whitepaper](). SAP HANA installations require three different volumes for the database, the data volume, the log volume, and the HANA shared volume. As a rule of thumb, the total required capacity for those three volumes is 2.5 times the RAM size of the SAP HANA database host. Additional volumes might be required for file-based data and log backups and exports.

Note: The capacity sizing of the overall SAP landscape with multiple SAP HANA systems must be determined by using SAP HANA storage sizing tools from NetApp. Contact NetApp or your NetApp partner sales representative to validate the storage sizing process for a properly sized storage environment.

3.5 Network Setup Best Practices

iSCSI support for IBM Power Systems was officially announced in the VIOS 3.1 announcement on October 9th 2018 (GA on November 15th 2018). iSCSI support in VIOS 3.1 allows iSCSI disks to be exported to client LPARs as virtual disks (vSCSI disks). Requirements for such an attachment are VIOS 3.1 and FW 860.20 or later on the IBM Power server.

iSCSI Network for Boot LUNs

Figure 12) iSCSI network.

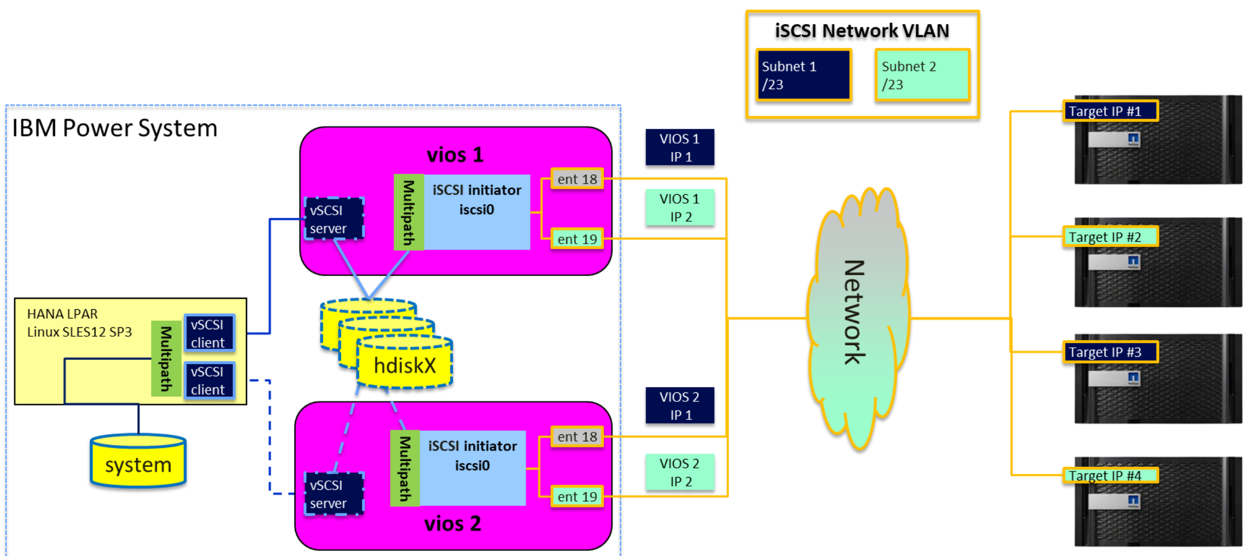

On Each VIOS:
Four iSCSI targets populated into ODM for iscsi0 (mkisci)
<Target IP #1> 3260 iqn.xxxx-yy.com.netapp:zzzzzzzzz
<Target IP #2> 3260 iqn.xxxx-yy.com.netapp:zzzzzzzzz
<Target IP #3> 3260 iqn.xxxx-yy.com.netapp:zzzzzzzzz
<Target IP #4> 3260 iqn.xxxx-yy.com.netapp:zzzzzzzzz

Figure 13) iSCSI configuration.

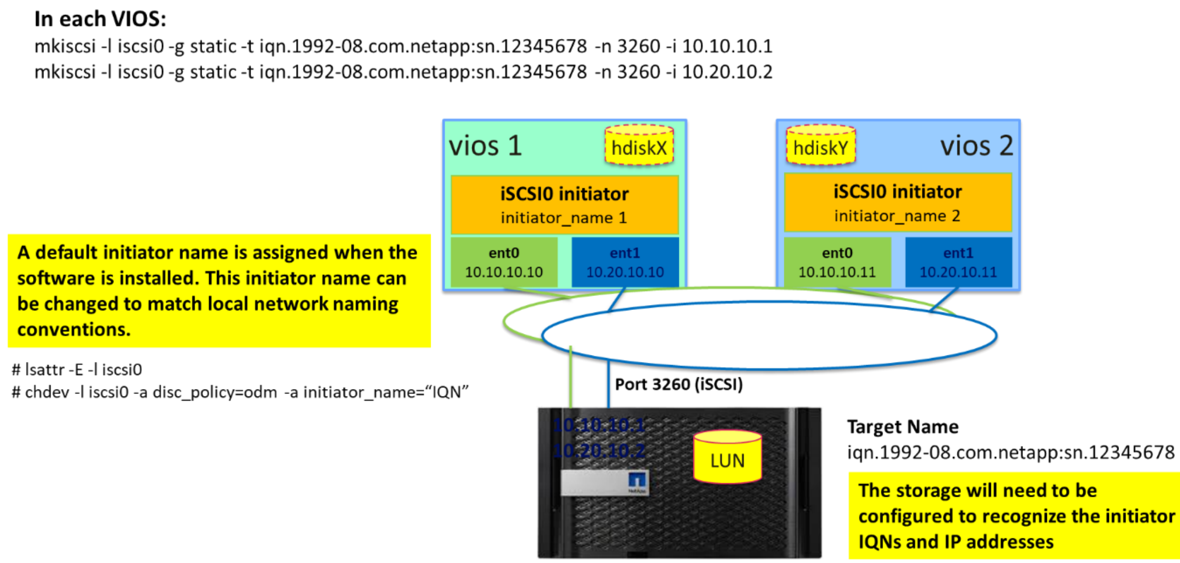

Figure 14) vSCSI configuration.

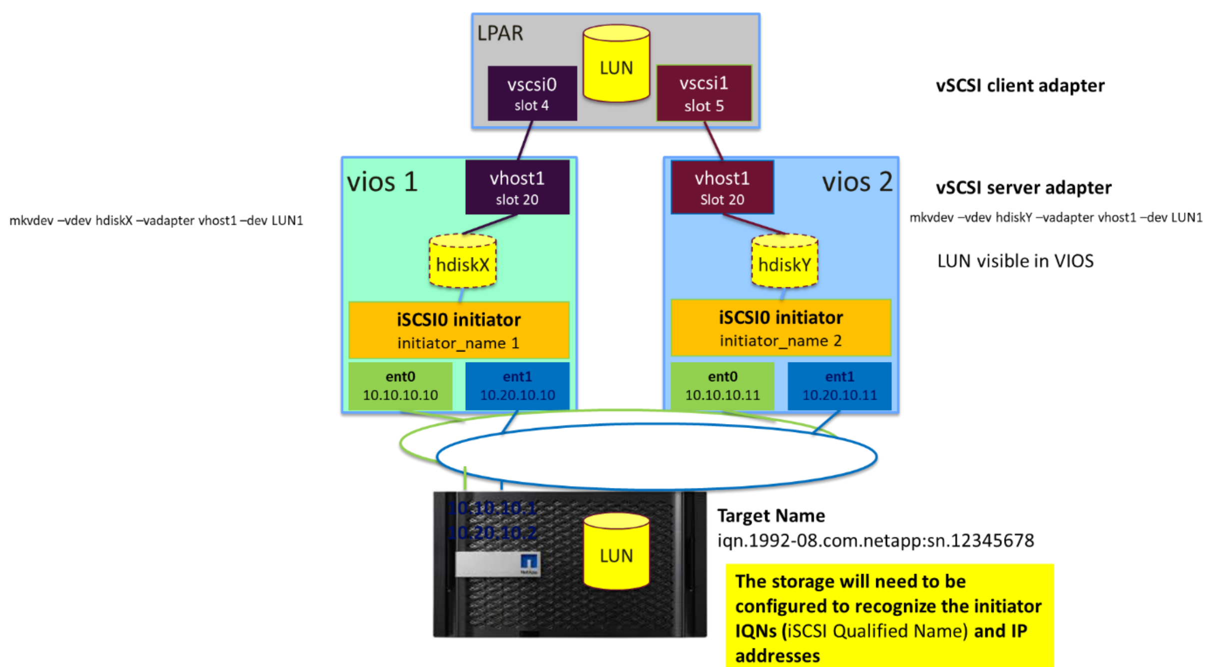

NFS Network for SAP HANA

All SAP application data (for example, executables and configuration files), database data, and log files are accessed through NFS volumes mounted directly on each Power Linux LPAR following the SAP adaptive computing principle. In order to meet the required performance KPIs for SAP HANA, the following guidelines must be considered when planning the storage network.

- You must use a dedicated 10GbE or faster storage network to connect the SAP HANA hosts to the storage controller.

- Disable flow control at all physical ports used for the storage traffic on the storage network switch and host layer.
- Each SAP HANA host must have a redundant network connection with a minimum of 10Gb of bandwidth.
- You must enable jumbo frames with an MTU size of 9,000 on all network components between the SAP HANA hosts and the storage controllers and the VIOS, including Platform Largesend Segmentation Offload (PLSO). See the IBM Network Configuration Guide for more details.

The following picture shows an example with four SAP HANA hosts attached to a storage-controller HA pair using a 10GbE network. Each SAP HANA host has an active-passive connection to the redundant fabric.

At the storage layer, four active connections are configured to provide 10Gb throughput for each SAP HANA host. In addition, one spare interface is configured on each storage controller.

At the storage layer, a broadcast domain with an MTU size of 9000 is configured, and all required physical interfaces are added to this broadcast domain. This approach automatically assigns these physical interfaces to the same failover group. All logical interfaces that are assigned to these physical interfaces are added to this failover group.

Figure 15) Network configuration example.

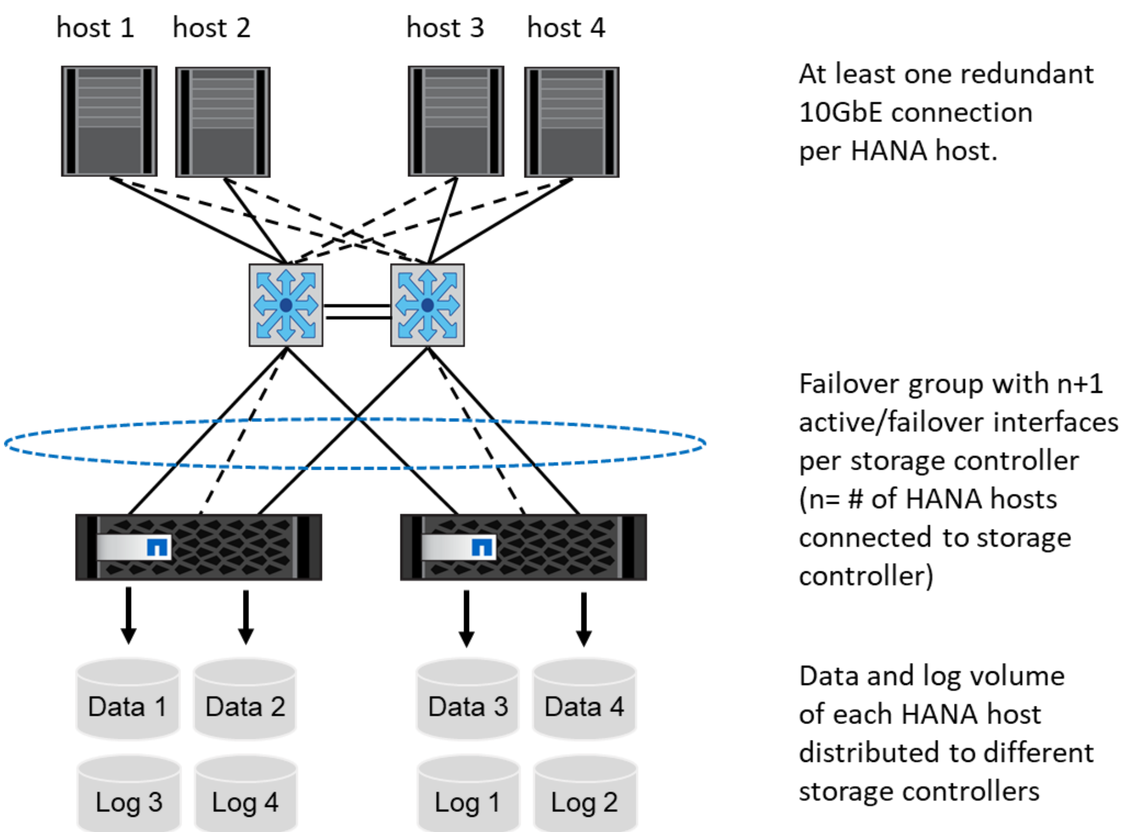

In general, it is also possible to use active-active interface groups on the servers (bonds) and the storage systems (for example, LACP and ifgroups). With active-active interface groups, make sure that the load is equally distributed between all interfaces within the group. The load distribution depends on the functionality of the network switch infrastructure.

3.6 Storage Controller Setup Best Practices

This section provides an overview of best practices for storage controller setup. See TR-4435 SAP HANA on NetApp AFF with NFS Configuration Guide for more details.

Storage Efficiency and Data Security Features

All storage efficiency features enabled by default on AFF systems are supported for SAP HANA deployments and should remain enabled. Storage efficiency savings up to 30% have been observed for SAP HANA deployments. You can also use additional features like NetApp Volume Encryption and storage quality of service.

Storage Configuration

The following overview summarizes the required storage configuration steps. Each step is covered in detail in TR-4435. We assume that the storage hardware setup, ONTAP software installation, and connection of the 10GbE or 40GbE storage ports to the network are already completed.

1. Check the correct SAS stack configuration.
2. Create and configure the required aggregates.
3. Create a storage virtual machine (SVM).
4. Create logical interfaces (LIFs).
5. Create volumes within the aggregates.
6. Set the required volume options.
7. Set the required options for NFSv3 or for NFSv4.
8. Mount volumes to the namespace and set export policies.

Volume Configuration for SAP HANA

In accordance with the SAP HANA Storage Requirements whitepaper, the following storage layout is required:

- SAP HANA data files at `/hana/data/<SID>`
- SAP HANA log files at `/hana/log/<SID>`
- SAP HANA installation binaries at `/hana/shared/<SID>`
- Additional SAP software required by SAP HANA at `/usr/sap`

Figure 16 shows the volume configuration recommendation of four single-host SAP HANA systems. The data and log volumes of each SAP HANA system are distributed to different storage controllers. For example, volume `SID1_data_mnt00001` is configured on controller A, and volume `SID1_log_mnt00001` is configured on controller B. The file systems for `/hana/shared/SID1` and `/usr/sap` are combined into the `SID1_shared` volume. Dedicated subfolders need to be created in the volume to achieve this. See TR-4435 for more details about this configuration.

Figure 16) Volume layout for SAP HANA single host system.

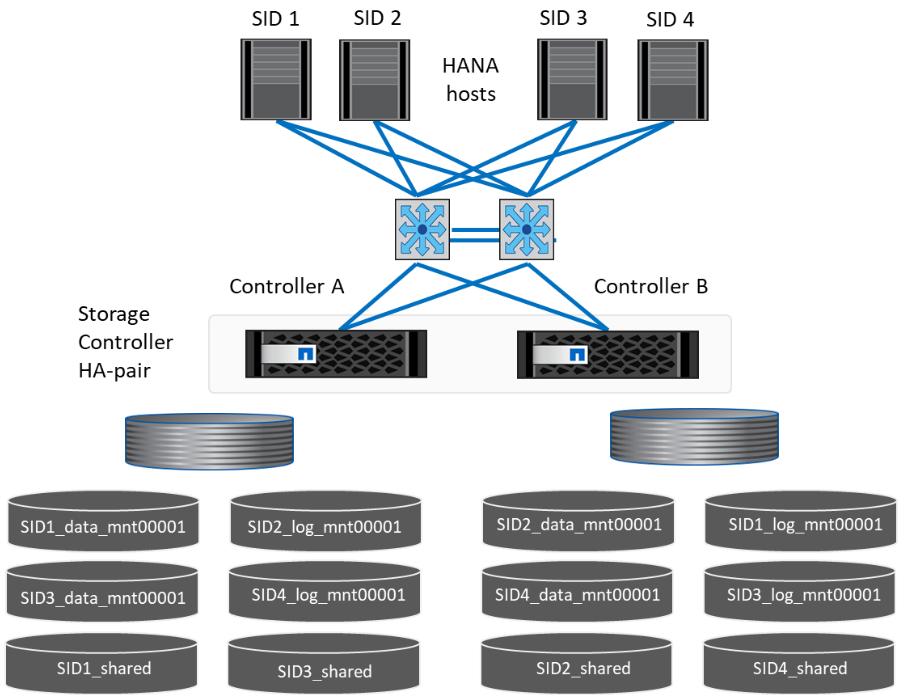

For SAP HANA multiple-hosts systems, the volume layout is very similar. The main difference is that dedicated data and log volumes are required for the master node and each SAP HANA worker node. For example, you could use `SID_data_mnt00001`, `SID_data_mnt00002`, and `SID_data_mnt00003` as data volumes for a 3+1 SAP HANA multiple-host system. Ideally, the data and log volumes should be distributed onto the different controllers in a manner similar to the single-host configuration shown in Figure 16. Finally, the volumes can be mounted on their specific mount points in the filesystem of the host. The following table shows the mount points for a single-host system as an example:

Table 1) Mount points for single-host systems.

Junction Path	Directory	Mount Point at HANA Host
SID_data_mnt00001		/hana/data/SID/mnt00001
SID_log_mnt00001		/hana/log/SID/mnt00001
SID_shared	usr_sap shared	/usr/sap/SID /hana/shared/SID

3.7 PowerVM Setup Best Practices

The following section describes how to integrate PowerVM environments with the NetApp AFF systems.

PowerVM VIOS 3.1 and NetApp Integration

The new PowerVM VIOS 3.1 is based on AIX 7.2 TL3. The VIOS was rebased to the latest version of AIX for modernization, security, and resiliency. AIX 7.2 is a clean, state-of-the-art base OS with many features that are beneficial for PowerVM virtualization.

The key features of version 3.1 are as follows:

- **Cloud ready.** iSCSI (network storage) virtualization is supported in the following ways:

- iSCSI support in VIOS allows iSCSI disks to be exported to client LPARs as virtual disks (vSCSI disks).
- Enables MPIO support for the iSCSI initiator. With MPIO support, users can configure and create multiple paths to an iSCSI disk similar to what is available and supported for other storage protocols.
- Both VIOS vSCSI and AIX iSCSI device drivers have been modified to support iSCSI disks in VIOS.
- Requires FW level 860.20 or later on POWER8 systems and FW level 910.00 or later on POWER9 systems.

- **Modernization.** Changes include the following:
  - Device driver enhancements in AIX 7.2 (for example performance, efficiency, and RAS)
  - Accelerator enablement is only available on AIX 7.2, VIOS on a base that allows the use of accelerators for virtualization.
  - Native compatibility mode (POWER8, POWER9).
- **Security and resiliency.** A smaller footprint enables shorter maintenance windows in the field. Storage multipathing improvements include the following:
  - Improved transient and sick-but-not-dead error handling and improved command time-out handling.
  - SSP enhancements.
  - Support for virtual IP address (VIPA) with multipath routing. This functionality allows the storage pool to use multiple networks for communication redundancy. Therefore, if one network goes down, then the storage pool can use the other network.
  - Storage pool log retention for improved RAS capability, and the maintenance of multiple log files to better preserve debug data. A log-wrapping issue was also resolved, which was causing premature wrap of the log.
  - The hardening of disk quorum (meta-root replica set) and manager disk challenges when mirroring. These changes provide various fixes for disk failure scenarios with a mirrored storage pool.
  - Various changes needed to avoid cluster-wide outages. Several issues were resolved that would result in a hung or inaccessible storage pool.
- **Simplified migration to VIOS 3.1.** The viosupgrade tool on the NIM server supports easy migration from VIOS 2.2.x to VIOS 3.1.
  - The viosupgrade tool backs up the virtual and logical configuration data, installs the specified image, and restores the virtual and logical configuration data of the virtual I/O server from the NIM server.
  - Supported from AIX 7.2 TL-03 SP-02 (7200-03-02).
  - bosinst type of installation is supported for new and complete installation.
  - altdisk type of installation is supported for alternate disk `alt_mksysb_install`.
  - Up to 30 VIO server installations can be triggered in parallel.
- viosupgrade tool on VIOS for smooth and easy migration from VIOS 2.2.6.x to VIOS 3.1 version.
  - It is an auto-migration process. This VIOS tool performs the operation of self-Backup of virtual and logical configuration data, Self-Installation of the specified image, and Self-Restore of the virtual and logical configuration data of the Virtual I/O Server.
  - Supported from VIOS 2.2.6.32 or later for smooth VIOS 3.1 migration.
  - Installation is of alt_disk_mksysb type.
- Availability and other details:
  - General availability date: Nov 9th 2018

- Supported hardware: POWER9, POWER8, and POWER7+
- USB flash drive install support
- VIOS recognized solutions

Virtual I/O Server Overview

The VIOS is part of the PowerVM® Editions hardware feature, and is located in a logical partition. This software facilitates the sharing of physical I/O resources between client logical partitions within the server. The VIOS provides virtual SCSI target support, virtual Fibre Channel support, a shared Ethernet adapter, and PowerVM Active Memory™ Sharing capability to client logical partitions within the system. The VIOS also provides suspend, resume, and remote restart features to AIX®, IBM® i, and Linux client logical partitions within the system.

As a result, you can perform the following functions on client logical partitions:

- Share SCSI devices, Fibre Channel adapters, and Ethernet adapters.
- Expand the amount of memory available to logical partitions and suspend and resume logical partition operations by using paging space devices.
- LPM (Live Partition Mobility), SSR (Simplified Remote Restart), and IBM VM RM (Virtual Machine Recovery Manager)

An exclusive, dedicated logical partition is required for the VIOS software. You can use the VIOS to perform the following functions:

- Sharing of physical resources between logical partitions on the system
- Creating logical partitions without requiring additional physical I/O resources
- Creating more logical partitions than there are I/O slots or physical devices available with the ability for logical partitions to have dedicated I/O, virtual I/O, or both
- Maximizing the use of physical resources on the system
- Helping to reduce the SAN infrastructure

You can manage the VIOS and client logical partitions by using the HMC and the VIOS command-line interface.

PowerVM® Editions includes the installation media for the VIOS software. The VIOS enables the sharing of physical I/O resources between client logical partitions within the server. When you install the VIOS in a logical partition on a system that is managed by the HMC, you can use the HMC and the VIOS command-line interface to manage the VIOS and client logical partitions.

When you install the VIOS on a managed system and there is no HMC attached to the managed system when you install the VIOS, then the VIOS logical partition becomes the management partition. The management partition provides the web-based Integrated Virtualization Manager system management interface and a command-line interface that you can use to manage the system.

For the most recent information about devices that are supported on the VIOS and to download VIOS fixes and updates, see the [Fix Central website.]

The VIOS contains the following primary components:

- Virtual SCSI
- Virtual networking

The following sections provide a brief overview of each of these components.

Virtual SCSI

Physical adapters with attached disks or optical devices on the VIOS logical partition can be shared by one or more client logical partitions. The VIOS offers a local storage subsystem that provides standard SCSI-compliant LUNs. VIOS can export a pool of heterogeneous physical storage as a homogeneous pool of block storage in the form of SCSI disks.

Unlike typical storage subsystems that are physically located in the SAN, the SCSI devices exported by the VIOS are limited to the domain within the server. Although the SCSI LUNs are SCSI compliant, they might not meet the needs of all applications, particularly those that exist in a distributed environment.

The following SCSI peripheral-device types are supported:

- Disks backed up by logical volumes
- Disks backed up by physical volumes
- Disks backed up by files
- Optical devices (DVD-RAM and DVD-ROM)
- Optical devices backed up by files
- Tape devices

Virtual Networking

The VIOS provides the following virtual networking technologies.

Table 2) Virtual networking technologies on the VIOS.

Virtual Networking Technology	Description
Shared Ethernet adapter	A Shared Ethernet adapter is a layer-2 Ethernet bridge that connects physical and virtual networks together. It allows logical partitions on the virtual local area network (VLAN) to share access to a physical Ethernet adapter and to communicate with systems outside the server. By using a shared Ethernet adapter, logical partitions on the internal VLAN can share the VLAN with stand-alone servers. On POWER7®-processor-based systems and later, you can assign a logical host Ethernet port of a logical Host Ethernet adapter, sometimes referred to as integrated virtual Ethernet, as the real adapter of a shared Ethernet adapter. A host Ethernet adapter is a physical Ethernet adapter that is integrated directly into the GX+ bus on a managed system. Host Ethernet adapters offer high throughput, low latency, and virtualization support for Ethernet connections. The shared Ethernet adapter on the VIOS supports IPv6. IPv6 is the next generation of internet protocol and is gradually replacing the current internet standard, IPv4. The key IPv6 enhancement is the expansion of the IP address space from 32 bits to 128 bits, providing virtually unlimited, unique IP addresses.
Shared Ethernet adapter failover	Shared Ethernet adapter failover provides redundancy by configuring a backup shared Ethernet adapter on a different VIOS logical partition that can be used if the primary shared Ethernet adapter fails. The network connectivity in the client logical partitions continues without disruption.
Link aggregation (or EtherChannel)	A ink aggregation (or EtherChannel) device is a network port-aggregation technology that allows several Ethernet adapters to be aggregated. The adapters can then act as a single Ethernet device. Link aggregation provides more throughput over a single IP address than would be possible with a single Ethernet adapter.

Virtual Networking Technology	Description
VLAN	A VLAN allows the physical network to be logically segmented.

Using virtual SCSI, client logical partitions can share disk storage and tape or optical devices that are assigned to the VIOS logical partition.

Disk, tape, USB mass storage, or optical devices attached to physical adapters in the VIOS logical partition can be shared by one or more client logical partitions. The VIOS is a standard storage subsystem that provides standard LUNs that are compliant with SCSI. The VIOS can export a pool of heterogeneous physical storage as a homogeneous pool of block storage in the form of SCSI disks. The VIOS is a localized storage subsystem. Unlike typical storage subsystems that are physically located in the SAN, the SCSI devices exported by the VIOS are limited to the domain within the server. Therefore, although the SCSI LUNs are SCSI-compliant, they might not meet the needs of all applications, particularly those that exist in a distributed environment.

The following SCSI peripheral device types are supported:

- Disk backed by logical volume
- Disk backed by physical volume
- Disk backed by file
- Disk backed by a logical unit in shared storage pools
- Optical CD-ROM, DVD-RAM, and DVD-ROM
- Optical DVD-RAM backed by file
- Tape devices
- USB mass storage devices

Virtual SCSI is based on a client-server relationship. The VIOS owns the physical resources and the virtual SCSI server adapter, and acts as a server, or SCSI target device. The client logical partitions have a SCSI initiator referred to as the virtual SCSI client adapter and access the virtual SCSI targets as standard SCSI LUNs. You can configure the virtual adapters and virtual disk resources by using the HMC or Integrated Virtualization Manager. The configuration and provisioning of virtual disk resources can be performed by using the HMC or the VIOS command line.

Physical disks owned by the VIOS can be exported and assigned to a client logical partition, added to a shared storage pool, or partitioned into parts, such as logical volumes or files. The logical volumes and files can then be assigned to different logical partitions. Therefore, by using virtual SCSI, you can share adapters and disk devices. Logical units in logical volumes and file-backed virtual devices prevent the client partition from participating in Live Partition Mobility.

To make a physical volume, logical volume, or file available to a client logical partition, it must be assigned to a virtual SCSI server adapter on the VIOS. The client logical partition accesses its assigned disks through a virtual-SCSI client adapter. The virtual SCSI client adapter recognizes standard SCSI devices and LUNs through this virtual adapter.

On the VIOS, for logical units in shared storage pools, you can thin-provision a client virtual-SCSI device for better storage space utilization. In a thin-provisioned device, the used storage space might be greater than the actual used storage space. If the blocks of storage space in a thin-provisioned device are unused, then the device is not entirely backed by physical storage space.

With thin provisioning, the storage capacity of the storage pool can be exceeded, which raises a threshold-exceeded alert. To determine that a threshold alert has occurred, check the errors listed in the HMC serviceable events or the VIOS system error log by running the errlog command in the VIOS command line. To recover after the threshold has been exceeded, you can add physical volumes to the storage pool. You can verify that the threshold is no longer exceeded by checking the HMCserviceable events or the VIOS system error log.

On the VIOS, multiple applications running on the virtual client can manage reservations on virtual disks of the client by using the persistent reserves standard. These reservations persist across hard resets, logical unit resets, or initiator target nexus loss. Persistent reservations that are supported by logical devices from the VIOS shared storage pools support the required features for the SCSI-3 persistent reserves standard.

On the VIOS, you can thick-provision a virtual disk. In a thick-provisioned virtual disk, you can allocate or reserve storage space while initially provisioning the virtual disk. The allocated storage space for the thick-provisioned virtual disk is guaranteed. This operation ensures that there are no failures because of a lack of storage space. By using thick-provisioning, virtual disks have faster initial access time because the storage is already allocated.

The following figure shows a standard virtual SCSI configuration.

Figure 17) Virtual SCSI configuration.

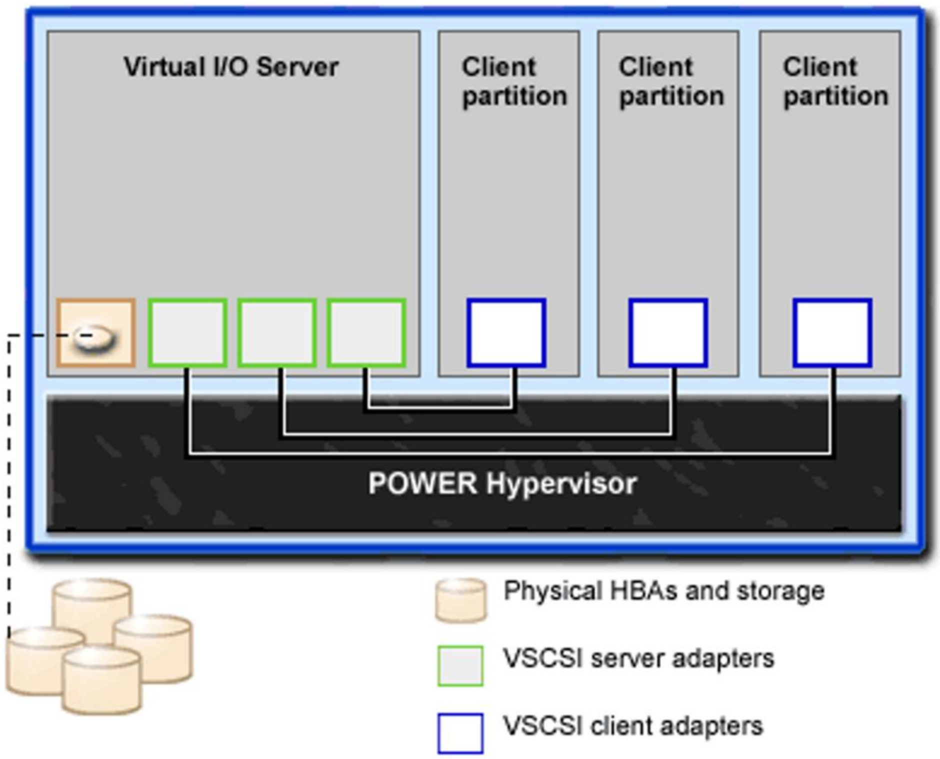

Note: The VIOS must be fully operational for the client logical partitions to be able to access virtual devices.

Installing and Customizing NetApp AIX Host Utilities

To enable and attach a NetApp storage server to an IBM Power9 server, you must complete the following steps:

1. Install a VIOS 3.1 image on the Power9 server.
2. Install the required NetApp device drivers and the utilities on the VIOS server. The following packages from the NetApp AIX Host Utilities must be installed:

```
# lslpp -l |grep NetApp
NetApp.MPIO_Host_Utilities_Kit.config              6.0.0.0  NetApp MPIO PCM Host Utilities
NetApp.MPIO_Host_Utilities_Kit.iscsi               6.0.0.0  NetApp MPIO PCM Host Utilities
NetApp.MPIO_Host_Utilities_Kit.pcmodm              6.0.0.0  NetApp MPIO PCM Host Utilities
```

```
NetApp.SAN_toolkit.sanlun                             6.0.0.0  NetApp SAN Toolkit sanlun
NetApp.SAN_toolkit.scripts                            6.0.0.0 NetApp SAN Toolkit Scripts
NetApp.iSCSI_Host_Utilities_Kit.LUN.msg.en_US         6.0.0.0  NetApp iSCSI Host Utilities
```

3. Configure the attached iscsi(x) interfaces with smit.

```
# lsattr -El iscsi0
disc_filename  /etc/iscsi/targets                     Configuration file                    False
disc_policy    odm                                    Discovery Policy                      True
initiator_name iqn.xxx:isvioa282.hostid.0a428131      iSCSI Initiator Name                  True
isns_srvnames  auto                                   iSNS Servers IP Addresses             True
isns_srvports                                         iSNS Servers Port Numbers             True
isw_err_recov  fast_fail                              iSCSI Network Error Recovery Policy   True
max_targets    16                                     Maximum Targets Allowed               True
max_xfer_size  0x80000                        Maximum transfer size                 True
num_cmd_elems  200                                    Maximum number of commands to queue   True
```

Note: Be careful to use the provided IQNs, and make sure to use 3260 as the default iSCSI port number.

You must set the following parameters on the iscsi0 device:

- disc_policy: odm
- nitiator_name: # Important: Use the IEEE standard (yyyy-mm)
- isw_err_recov: fast_fail

```
chdev -l iscsi0 -a disc_policy=odm
chdev -l iscsi0 -a isw_err_recov=fast_fail
chdev -l isci0 -a initiator_name=<iqn…>
```

4. Configure the iSCSI targets with smit.

```
# lsiscsi
iscsi0    198.18.18.175    3260 iqn.1992-08.com.netapp:sn.2214298335b211e9876200a098d99942:vs.14
iscsi0    198.18.18.178    3260 iqn.1992-08.com.netapp:sn.2214298335b211e9876200a098d99942:vs.14
```

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- All Flash FAS Documentation Resources
  https://www.netapp.com/us/documentation/all-flash-fas.aspx
- AFF and FAS System Documentation Center
  https://docs.netapp.com/platstor/index.jsp
- TR-4435: SAP HANA on NetApp AFF with NFS Configuration Guide
  https://www.netapp.com/us/media/tr-4435.pdf
- TR-4018: Integrating NetApp ONTAP Systems with SAP Landscape Management
  https://www.netapp.com/us/media/tr-4018.pdf
- TR-4646: SAP HANA Disaster Recovery with Asynchronous Storage Replication
  https://www.netapp.com/us/media/tr-4646.pdf
- TR-4614: SAP HANA Backup and Recovery with SnapCenter installation and configuration guide
  https://www.netapp.com/us/media/tr-4614.pdf
- SAP Note 2296290 - New Sizing Report for SAP BW/4HANA (SAP Login required)
  https://launchpad.support.sap.com/#/notes/2296290
- SAP HANA Storage Requirements
  https://www.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html
- SAP HANA on IBM Power Systems and IBM System Storage – Guides
  https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102502

- IBM Power Systems SR-IOV - Technical Overview and Introduction
  http://www.redbooks.ibm.com/redpapers/pdfs/redp5065.pdf

Version History

Version	Date	Document Version History
Version 1.0	March 2020	Initial release

Refer to the [Interoperability Matrix Tool (IMT)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp