

REFERENCE ARCHITECTURE

Maximizing SAP HANA Performance and Reliability with Pure Storage

A reference architecture for SAP bundled application suites on SAP HANA with Pure Storage.

Contents

Executive Summary	3
Introduction	3
Solutions Overview	3
Solutions Benefits	4
Technology Overview	5
Compute Resources	6
Network Resources	6
Pure Storage FlashArray	6
Pure Cloud Block Store	7
Pure Storage FlashBlade	7
SAP HANA	8
Technical Solution Design	8
SAP Application Suites on SAP HANA with Pure Storage	8
Enhancing Homogeneous System Copy Efficiency with Volume Snapshots	15
Protecting SAP Landscapes with Volume Snapshots	16
Disaster Recovery with ActiveDR and Asynchronous Snapshots for SAP Landscapes	17
High Availability Strategies Using Synchronous Active/Active Configurations with ActiveCluster	19
Rapid Recovery with Backint-certified Data Protection Solutions	20
Design Validation	21
Deployment	21
FlashArray	21
Host Operating System	28
Conclusion	28



Executive Summary

This document presents a comprehensive reference architecture for deploying SAP HANA on Pure Storage® solutions, addressing various business needs through specialized technical designs. The architecture integrates multiple technical solutions within a unified framework, focusing on optimizing performance, scalability, and reliability for SAP HANA environments.

The architecture outlines the deployment of SAP HANA on Pure Storage systems, detailing the compute, networking, and storage layers to ensure high performance and scalability. Solutions in this architecture include volume snapshots for robust data protection, enabling quick, efficient, and reliable backups of critical SAP data, thus ensuring minimal downtime and data loss. The architecture also combines near-synchronous replication (ActiveDR™) with periodic asynchronous snapshot replication to ensure high data availability and resilience against various disaster scenarios. ActiveCluster™ offers a zero-recovery time objective (RTO) and zero-recovery point objective (RPO) for SAP HANA environments, ensuring continuous operation and data integrity during system or site failures. Additionally, integrating SAP HANA with Pure Storage and Backint certified data protection solutions ensures rapid recovery, minimizing downtime and protecting critical data assets efficiently.

By implementing this reference architecture, organizations can achieve enhanced performance and scalability, improved efficiency, robust data protection, high availability, and cost-effective solutions. Efficient storage usage and simplified management reduce the total cost of ownership (TCO) while maintaining high performance and reliability. This strategic framework leverages Pure Storage solutions to enhance the performance, efficiency, and resilience of SAP HANA environments, supporting business continuity and growth.

Introduction

This is a reference architecture focusing on the implementation of Pure Storage products in SAP landscapes. It offers comprehensive solutions for use cases such as storage consolidation and footprint reduction, data protection, disaster recovery, high availability and system copy optimization. These solutions utilize Pure Storage FlashArray™ for primary storage in the SAP landscape and Pure Storage FlashBlade® for streaming backup solutions with third party backup software.

This composite reference architecture is focused landscapes that use bundled application suites such as S/4HANA and BW/HANA with the SAP HANA in-memory database.

Solutions Overview

This reference architecture and associated solutions within this document consist of a common SAP landscape environment consisting of multiple interconnected components (Figure 1). These components include the following:

- SAP bundled application suite software deployed on one or more hosts
- SAP HANA deployed on one or more hosts connected through block or file storage on Pure Storage FlashArray via high performance fabric (fiber channel) or ethernet networking
- Pure Cloud Block Store™ instances in public cloud providers such as Microsoft Azure or Amazon Web Services connected to public cloud virtual machines (AWS-EC2 or Azure-VM's) using iSCSI connectivity that host test, development, quality assurance or DR copies of production SAP HANA databases.
- Certified backint backup software deployed on one or more hosts connected to either FlashArray or FlashBlade storage repositories via high performance ethernet networking.



Each interconnected infrastructure component within the SAP HANA system must meet the certification requirements of [Tailored Data Center Integration](#). Pure Storage, however, is solely responsible for certifying the primary storage. It is essential to verify during procurement and deployment that all infrastructure components have SAP certification. Although the term 'hosts' usually refers to physical, bare-metal servers, it can also apply to virtual machines. Both [single host \(scale up\) and multiple-host \(scale out\) system types](#) are addressed by these solutions. These solutions do not distinguish between different [SAP HANA deployment types \(single-tenant, multi-tenant, MCOD, Virtualized and MCOS\)](#) as they attempt to address all of them with a common architectural framework.

Unless otherwise specified, the solutions do not specify a minimum number of hosts or network speed, instead these are design decisions that will be made based on supportability, capabilities, and business need. Each reference architecture can be templated without limits to any single component.




FIGURE 1 High-level overview of interconnected components and environments

Solutions Benefits

This solution offers many benefits, including:

- **Simplified data management:** Pure Storage simplifies the SAP data journey by offering streamlined migration and always-on data protection for SAP and SAP S/4HANA environments. This helps in modernizing SAP operations and accelerating digital transformation initiatives.
- **Seamless data mobility:** With Pure Storage, organizations can eliminate data gravity and achieve seamless data mobility across hybrid cloud environments. This includes consistent data services, resilience, and APIs that ensure smooth operation and integration within diverse IT landscapes.
- **Enhanced consolidation, energy efficiency, and sustainability:** Pure Storage solutions allow for the consolidation of massive SAP HANA database workloads onto fewer arrays. This consolidation leads to increased operational and energy efficiencies, reduced power consumption, and optimized space usage—all while maintaining high performance levels.
- **Enhanced data protection:** Pure Storage provides robust data protection mechanisms, including built-in data replication and space-efficient immutable snapshots (SafeMode™ Snapshots). These features protect against data loss, breaches, ransomware, and interruptions in availability, ensuring business continuity even in adverse conditions.
- **Efficient data management tools:** The use of fast and efficient snapshots allows for the automation of SAP copy, clone, and refresh processes, significantly reducing the time and effort required for these tasks. This automation also extends to backup and recovery processes, aiding in achieving low recovery-time objectives (RTO).



- **Advanced security features:** The solution safeguards SAP data from disasters and maximizes business continuity. It speeds up ransomware recovery and prevents ransomware from eradicating, modifying, or encrypting data, even with admin credentials.
- **Reduced restart times:** The higher performance and lower latency of Pure Storage's platform significantly reduce restart times for SAP HANA. This results in less downtime and more consistent performance, which is critical for maintaining business operations and reducing risk.
- **Increased total cost of ownership (TCO):** The combination of Pure Storage data reduction technologies and SAP HANA Native Storage Extension (NSE) decreases the amount of data stored in memory and reduces the storage demands of the SAP HANA landscape. This reduction not only lowers storage costs but also leverages memory unloading technology to extend the lifespan of an SAP HANA system, thus contributing to a more cost-effective and sustainable operational environment.

Technology Overview

This reference architecture consists of a common environment with a set of interconnected technology components that communicate with one another over redundant high-speed ethernet connections or a fiber channel fabric. Application component inter-communication is done over ethernet while storage component and application host inter-communication are performed over either ethernet or fiber channel fabric.

FlashArray is the exclusive provider of primary storage for SAP HANA, utilizing Fiber Channel Protocol (FCP) for block storage or NFS protocol over Ethernet for file storage. For secondary storage, used for data protection or as a repository, FlashBlade supports both file and object storage over Ethernet. Additionally, FlashArray can also serve secondary storage purposes using either Fiber Channel Protocol or Ethernet.




FIGURE 2 Technology component interconnects



Compute Resources

The solutions within this document apply to physical or virtual multi-core servers with [Red Hat Enterprise Linux for SAP Applications](#) or [SUSE Enterprise Linux for SAP Applications](#) operating systems. These are referred to in this document as either physical or virtual hosts. Pure Storage does not differentiate among various virtualization platforms for SAP HANA; however, for the purposes of the solutions discussed herein, only VMware vSphere or IBM PowerVM are considered.

Matching the [SAP HANA supported hardware platforms](#) the solutions in this document apply to the following server platforms:

- Intel-based hardware platforms
- IBM Power Systems

Maximums and minimums are defined for server vendors with SAP HANA through the [TDI certification program](#).

Network Resources

The network considerations for these solutions are the same as any enterprise IT infrastructure solution: availability, performance, and extensibility. The compute resources in the solution can attach to any compatible TCP/IPv4 or TCP/IPv6 network infrastructure with the general recommendation that minimum network speeds are capable of 25GbE. For this solution, the network should be configured using dual switches to eliminate a single point of failure.

Where possible [client side \(6-balance-alb\) or switch side link aggregation \(802.3ad\)](#) should be used to achieve the highest performance and availability.

SAP sets out the requirements for networking for SAP HANA in the document [SAP HANA Network Requirements](#).

Pure Storage FlashArray

[Pure Storage FlashArray](#), a unified block and file-storage solution driven by software-defined technology, provides a seamless and reliable user experience. It incorporates data reduction capabilities without compromising performance. All Pure Storage products feature an Evergreen® subscription model, allowing for capacity and performance upgrades without requiring new storage purchases. Additionally, FlashArray empowers businesses and organizations to significantly reduce direct carbon emissions in their data storage systems, achieving up to an 80% decrease compared to competing all-flash systems and even more when compared to magnetic disks.

The FlashArray product line caters towards multiple business needs and use cases. The following [product offerings are certified for SAP HANA](#):

- **FlashArray//X™**: Certified for block storage via Fibre Channel protocol and Network File System over Ethernet. Provides latency as low as 150µs to power critical applications and business operations.
- **FlashArray//XL™**: Certified for block storage via Fiber Channel Protocol. Enterprise-grade performance and scalability for demanding workloads.
- Additionally, the following FlashArray product offerings can be used for non-production HANA environments, data repositories, backup repositories, and application binary storage:
- **FlashArray//E™**: Uses 80% less energy at 60% less cost for active data archives, file repositories, and other use cases
- **FlashArray//C™**: An all-QLC FlashArray with consistent performance at 2-4ms latency for capacity-oriented workloads



Pure Cloud Block Store

[Pure Cloud Block Store](#) provides seamless data mobility across on-premises and cloud environments with a consistent experience regardless of where data lives. It provides enterprise-grade storage features in the cloud and its industry leading data efficiency means you buy less capacity in the cloud without sacrificing agility and flexibility. Pure Cloud Block Store is available in the [Amazon Web Services Marketplace](#) and the [Microsoft Azure Marketplace](#).

Note: Pure Cloud Block Store is not certified for SAP HANA primary storage and should only be used for test, development and quality assurance landscape components.

Pure Cloud Block Store is available in two versions as shown in Figure 3: //V10 and //V20. Each version provides different capacity and performance capabilities.




FIGURE 3 The V10 and V20 CBS version. The capacity listed for each model is effective capacity with a 4:1 data reduction rate.

Pure Storage FlashBlade

[Pure Storage FlashBlade](#) is tailored to meet the demands of modern data analytics requirements. Its innovative blade architecture prioritizes speed and efficiency, ensuring exceptional performance and scalability. This simplicity and adaptability empower organizations to effortlessly scale storage infrastructure to meet the demands of intensive workloads such as rapid recovery, analytics, AI, and machine learning. With its all-flash design, FlashBlade consistently delivers high performance across various workloads, from backup to analytics and AI.

FlashBlade//S™ represents the evolution in enterprise scale-out storage, offering a blend of high density, capacity, performance, and scalability to meet the demands of modern applications.

Note: FlashBlade is not certified for SAP HANA primary storage and should only be used for scenarios such as external data repositories, backup repositories and application binary storage (HANA-Shared).

FlashBlade//E™ is a cost-efficient storage platform, aiming to provide effective performance, uncompromising reliability, and all-flash capabilities to a broader audience, especially those with budget considerations. As a cost-effective storage solution within the FlashBlade product family, it provides exceptional value for analytics.

Capacity-optimized and performant storage




FIGURE 4 Technical specifications of the FlashBlade product family



SAP HANA

[SAP HANA](#) is a high-performance in-memory database and application platform designed by SAP SE. It allows businesses to process large volumes of data in real-time, facilitating faster decision-making and analytics. SAP HANA's unique in-memory technology enables it to store data on main memory rather than on traditional disk storage, significantly speeding up data retrieval and processing times. This technology supports both transactional (OLTP) and analytical (OLAP) processing in a single system, eliminating the need for separate systems and reducing data redundancy and latency. SAP HANA is widely used across various industries for applications like ERP (enterprise resource planning), CRM (customer relationship management), and data warehousing, making it a versatile and critical tool in the modern data-driven business landscape.

The following bundled application suites typically used with SAP HANA:

- [SAP S/4HANA](#): An enterprise resource planning (ERP) software suite, designed to run exclusively on the SAP HANA in-memory database. It represents a major update to the SAP Business Suite, incorporating modern design principles with the SAP Fiori user experience and simplified data models. SAP S/4HANA is built to handle various business functions such as finance, sales, service, procurement, manufacturing, and supply chain in one integrated system. This reference architecture refers to the on-premises revision of S/4HANA.
- [SAP BW/HANA](#): A modern data warehousing solution from SAP that leverages the SAP HANA in-memory database to enhance speed and efficiency. This platform integrates real-time data from various sources, both SAP and non-SAP, and supports activities like data modeling, management, and staging. It features advanced analytics, a simplified user interface with SAP Fiori, and robust big data capabilities. SAP BW/4HANA is designed for high performance and is ideal for enterprises aiming to improve their data-driven decision-making processes.

Technical Solution Design

This composite reference architecture integrates multiple technical solution designs within a unified framework for SAP deployments using Pure Storage. It addresses distinct business needs through specialized architectures in the following areas:

- Implementation of SAP Application Suites on SAP HANA with Pure Storage.
- Enhancement of homogeneous system copy efficiency using volume snapshots.
- Protecting SAP landscapes with volume snapshots.
- Disaster Recovery with ActiveDR and Asynchronous Snapshots for SAP Landscapes
- High Availability strategies using synchronous Active/Active configurations with ActiveCluster.

This document details architectures aimed at delivering targeted outcomes, including the application of specific designs to comply with business rules and the enhancement of standard workflow efficiencies.

SAP Application Suites on SAP HANA with Pure Storage

This solution is a high-level overview of how Pure Storage fits into SAP Landscapes. This technical solution design is divided into separate technical layers that include compute, networking, storage and applications.

Note: SAP Landscapes are considered unified systems with different stages of use. A typical SAP landscape is divided into dedicated units for development, testing and production. This technical design considers the typical landscape units.

Additional best practice and support guidance for SAP HANA on FlashArray can be found in the [SAP HANA Implementation and Best Practices support page](#).



Compute Layer

This design uses at least three (3) sets of systems, development (DEV), quality assurance (QA) production (PROD), that include the following within them:

- One (1) or more SAP HANA hosts with either a scale up or scale out deployment configuration. The required servers are certified Intel or IBM Power servers.
- One (1) or more hosts for the SAP bundled application suite software.
- A host for [backint certified data protection software](#).
- Hosts for adjacent management suites for the such as [SAP Solution Manager](#) or [Landscape Manager](#)

All hosts, including those for the SAP application suites, SAP HANA databases and others can be deployed bare-metal or as virtual machines in a certified hypervisor.

Network Layer

Network ports on the hosts should be bonded in any one of the following configurations:

- Highest availability/lowest performance: Active/Passive (Mode 1)
- Highest performance with no switch side configuration: Active Load Balancing (Mode 6)
- Highest performance and highest availability: 802.3ad/LACP (Mode 4)

It is strongly recommended that dual switches are used and configured with an interconnect between them for availability and performance. Additional network optimizations for FlashArray file services for SAP HANA data and log storage are to set all interfaces are as follows:

- MTU - 9000

For high-performance use cases the servers and storage should be connected to the same switches or have dedicated storage networking with no routing or hops between them.

FlashArray File Services

FlashArray file services are built using physical ethernet interfaces mapped to a virtual interface. A virtual interface combines one or more physical ports on both controllers into a group for managing network communication for a specific service. Multiple network interfaces can be created on the same ports.

When deploying multiple SAP HANA nodes on a single array, it's advisable to configure multiple virtual interfaces. Assign each node to a unique virtual interface address to enhance load balancing across various ports on the array. For configurations involving 10GB or 25GB ports on the file services virtual interface, especially when numerous SAP HANA nodes are expected to share the same FlashArray, ensure ample availability of ports on each controller. Each port should be equipped with its own virtual interface to optimize performance and manageability.



The following table can be used as guidance for the requirements to meet the per node KPIs:

SAP HANA Nodes	FlashArray ports and speed (per controller)	File service virtual interfaces required
8	4 × 10Gbps	4
8	2 × 25Gbps	2
8	1 × 100Gbps	1
16	8 × 10Gbps	8
16	4 × 25Gbps	4
16	2 × 100Gbps	2

TABLE 1 Per-node KPI requirements.

Storage Layer

The storage layer consists of FlashArray, FlashBlade and Fiber Channel Fabric components.

FlashArray

FlashArray provides production storage for the SAP Landscape through block or file storage protocols. Block storage for production SAP HANA deployments can only be serviced through a fiber channel fabric. This solution design recommends that a high performant, redundant and resilient SAN topology be deployed for the SAP Landscape.

The storage devices and capacity recommendations for SAP HANA are set out as follows.

Device	Size	Purpose
Installation	Installation size = Minimum 1 x RAM	Contains run-time binaries, installation scripts and other support scripts. This also contains the SAP HANA configuration files, trace files and profiles.
Backups	Backup size = (Size of Data + Size of Redo Log) x retention period	Regularly scheduled backups are written to this location.
Data	Data size = 1 x Amount of RAM	SAP HANA persists a copy of the in-memory data to this location. This is achieved by writing changed data in the form of savepoints.
Red Log	For systems <512GB, Redo Log size = 1/2 RAM For systems >512GB Redo Log size = 512GB or larger	Each transaction performed on the database is recorded to this location in the form of a redo log entry.

TABLE 2 Device and capacity recommendations for SAP HANA.



The following mount point design are put forward as a part of this architecture:

- Data - /hana/data
- Redo Log - /hana/log
- Installation - /hana/shared
- Backups - /hana/backups

Storage capacity and performance requirements for SAP bundled application suites and data protection software is outside the scope of this design.

Block Storage

Block storage on FlashArray uses the following terminology:

- Host: A server that connects to and communicates with a storage device
- Host group: A group of hosts in a management structure that connect to one or more common volumes
- Volume: A device where data is stored in fixed size blocks
- Volume group: A group of volumes in a management structure

When using block storage for the SAP HANA deployment, each device should be on a separate volume.

Scale Out deployments require a shared filesystem for the Installation device. File Services with NFS mount points can be used for this deployment type for the installation device.

Host Storage Configuration

The recommended file system to use with FlashArray Block devices is the XFS file system for both data and log volumes.

The only recommended mount option outside of the defaults is the use of noatime. Example:

```
/dev/mapper/<device> /mountpoint xfs noatime 0 0
```

File Services

File services on FlashArray uses the following terminology:

- File systems: A management object for a system containing files and directories
- Directories: A directory within a filesystem that contains files
- Directory exports: A directory exported as a share to one or more clients via a file services protocol such as SMB or NFS
- Policy: An access and control policy applied to one or more directory exports



The below guidance is set out for the use of file services (NFS) for SAP HANA:

- One file system should be created per instance.
- A single /hana/shared directory should be created within that file system. Do not use the default, root directory.
- A /hana/data and /hana/log directory should be created per node in the instance.
- Multi-Host deployment Example (Figure 4) showcases an example of how a multi-host deployment would be configured. For a 3+1 scale out configuration the following managed directories need to be created within the same file system

- 3 directories for HANA-Data
- 3 directories for HANA-Log
- 4 directories for /usr/sap/<SID>
- 1 directory for HANA-Shared




FIGURE 5 Multi-Host deployment storage subsystem configuration

A policy should be created per instance and then attached to each directory.

- Rules should include the following:
 - Clients: Only the SAP HANA nodes should be added as clients
 - Access: o-root-squash
 - Permission: rw
 - Version: NFSv3 or NFSv4 depending on requirements
- Details should include the following:
 - Type: NFS
 - Enabled: True
 - User Mapping Enabled: False (disabling NFS user mapping)
 - Version: NFSv3 or NFSv4 depending on requirements

Supported Configurations for RHEL and SUSE Enterprise Linux

The following table sets out the supported configurations for SAP HANA deployments with FlashArray file services:

Deployment Type	NFS Version	HA/DR Provider
Single Host, Multi-Host without auto failover	v3, v4	Not required
Multi-host with auto failover	v3, v4	Server-specific STONITH Implementation required.

TABLE 3 Supported configurations for SAP HANA deployments with FlashArray file services

Host Storage Configuration

The following mount point options are put forward as a part of this architecture:

- Mount point: /hana/backup
- Mount options:
 - **hard:** Sets the recovery behavior of the NFS client after an NFS request timeout.
 - **nconnect=8:** Increases performance at scale by using more TCP connections between the client and the FlashBlade file service (should be at least 8 but can be increased). NFSv3 Only.
 - **mountproto=tcp:** Specifies that the mount protocol to be used is TCP to ensure no data loss when using long running NFS connections.
 - **Example:** 10.21.236.98:/hanadata /hana/data rsize=1048576,wsiz=1048576,nconnect=8,hard,mountproto=tcp

FlashBlade

FlashBlade provides file or object storage over ethernet to any hosts for data protection or application repository purposes.

Note: FlashBlade does not provide database storage (data or log) for SAP HANA.

There are an extensive number of possibilities for the use of FlashBlade. This reference architecture provides a general template for the use of NFS filesystems for the SAP HANA backup device.

Host Storage Configuration

The following mount point and options are put forward as a part of this general template:

- Mount point: /hana/backup
- Mount options:
 - **hard:** Sets the recovery behavior of the NFS client after an NFS request timeout.
 - **nconnect=8:** Increases performance at scale by using more TCP connections between the client and the FlashBlade file service (should be at least 8 but can be increased). NFSv3 Only.
 - **mountproto=tcp:** Specifies that the mount protocol to be used is TCP to ensure no data loss when using long running NFS connections.
 - **Example:** 10.21.236.98:/hanabackup /hana/backup nfs rw,hard,fsc,nconnect=8, ,vers=3,tcp,timeo=600



SAP HANA Design Considerations

The following design considerations are set out for SAP HANA deployed on FlashArray.

Set SAP HANA Parameters with hdbparam

[SAP Note 2267798](#) sets out how a customized SAP HANA configuration file can be used to set out during the installation procedure for both scale up and scale out deployments.

With HANA 2.0 the hdbparam tool has been deprecated. See SAP Note 2399079. The parameters can now be found in Configuration → Global.ini→fileio

To get the optimal usage of the storage subsystem, set the parameters to the following values:

fileio.num_completion_queues	8
fileio.num_submit_queues	8
fileio.size_kernel_io_queue	512
fileio.max_parallel_io_requests	64
fileio.min_submit_batch_size	16
fileio.max_submit_batch_size	64
fileio.async_write_submit_active	On
fileio.async_write_submit_blocks	All
fileio.async_read_submit	on

TABLE 4 Parameter values

Competing Storage Utilization

Pure Storage FlashArray comes default with Always-On Quality of Service (QoS). There are no knobs and nothing to configure: Always-On QoS protects against noisy neighbors. Always-On QoS prevents workloads from using more than their fair share of resources on the array by efficiently throttling noisy neighbors. QoS limits in terms of either bandwidth or IOPS can be applied on a per volume basis to throttle individual workloads and ensure that no other workloads are impacted.

QoS limits can also be applied to a group of volumes ensuring a consistent experience for all tenants of the array by offering one performance limit setting (MB/s) to configure for the group. This also ensures that tenants receive consistent performance as new tenants are added.

Operating other workloads on an array with SAP HANA installed is possible by using QoS rate limiting and ensuring that the volumes used for SAP HANA have all the IOPS and bandwidth required to complete any operations regardless of the other workloads on the storage system.

If it is required that each volume needs to have a QoS set, all that needs to occur is that a user navigates to the Storage view in the FlashArray GUI, under the volume heading a volume or volume group is selected and the QoS rate limiting is set for that volume.



Enhancing Homogeneous System Copy Efficiency with Volume Snapshots

Overview

This use case focuses on leveraging volume snapshots to efficiently copy databases between different environments within the SAP landscape, such as production (PROD), development (DEV), and quality assurance (QA). Traditionally, system copies are performed using backups or row-level copies, which can be time-consuming and resource-intensive. By utilizing volume snapshots, the process of copying databases becomes significantly faster and more efficient. Volume snapshots enable near-instantaneous, space-efficient copies of the entire database, thereby minimizing downtime and reducing the overhead associated with traditional copy methods. This approach enhances overall system copy efficiency, streamlining operations and improving productivity across SAP environments.

Solution Benefits

Volume snapshots with FlashArray offer a simplified and highly efficient solution for system copies. The creation of snapshots is instantaneous and consumes no additional storage capacity, as they only store changes made since the last snapshot, making them extremely space-efficient. Furthermore, copying a snapshot to a new volume is also instantaneous, eliminating the lengthy data transfer processes typical of traditional methods. This efficiency not only reduces downtime but also accelerates development and testing cycles.

The snapshot technology in FlashArray ensures that all snapshots reduce unnecessary data duplication, optimizing storage usage while maintaining high performance. By providing a quick, resource-efficient, and easy-to-manage solution, FlashArray volume snapshots significantly enhance the overall productivity and agility of SAP environments.

Technical Solution

Snapshots are created by capturing the state of one or more volumes at a point in time without duplicating the contents. After the point of creation, they only store the changes made since the last snapshot, significantly reducing the storage consumption. When a new copy is required, the snapshot can be quickly cloned to create a new volume, drastically cutting down the time needed for data transfers and minimizing downtime.




FIGURE 6 Copying volumes with snapshots between SAP HANA instances.



For landscapes not consolidated into a single FlashArray, asynchronous replication can be used to copy snapshots between higher and lower environments.




FIGURE 7 Asynchronous Replication for torn landscapes.

More information on best practices and how-to guidance for the use of volume snapshots with SAP HANA can be found in the [Snapshots and Data Protection](#) support documentation.

Protecting SAP Landscapes with Volume Snapshots

Overview

This use case focuses on employing volume snapshots to protect SAP landscapes, particularly those with SAP HANA databases. Volume snapshots serve as an effective data protection strategy by allowing for quick, efficient, and reliable backups of critical SAP data. These snapshots can be utilized for recovery in the event of data corruption, system failure, or other disruptive incidents, ensuring minimal downtime and data loss.

Solution Benefits

Volume snapshots provide several significant benefits when used to protect SAP landscapes, especially those incorporating SAP HANA databases. Firstly, they enable rapid recovery, allowing SAP HANA environments to be restored to a prior state swiftly, which minimizes downtime and reduces business disruption. This frequent and strategic snapshot reduces the risk of substantial data loss by capturing consistent, point-in-time snapshots of the data, reflecting recent changes. Additionally, volume snapshots are more storage-efficient than traditional backups, as they don't require a full data copy, making them a cost-effective solution for data protection. The operational simplicity of managing snapshots allows for easy automation and seamless integration into existing data management frameworks. Importantly, the non-disruptive nature of snapshots ensures that they can be implemented without affecting the ongoing performance of SAP HANA systems, thus maintaining continuous operation without operational slowdowns during backup processes.



Technical Solution

With great similarity to the content for the [technical solution for homogeneous system copies](#), this solution uses identical technologies and principles.




FIGURE 8 Volume snapshots being created at regular intervals results in multiple recovery points.




FIGURE 9 Volumes can be recovered from any of the multiple recovery points.

Disaster Recovery with ActiveDR and Asynchronous Snapshots for SAP Landscapes

Overview

This use case outlines a robust disaster recovery strategy for SAP landscapes, utilizing both ActiveDR and asynchronous snapshot replication. ActiveDR provides near-synchronous replication that continuously captures and replicates data to a secondary system with minimal latency. Asynchronous snapshot replication complements this by capturing periodic snapshots, which can then be replicated to another FlashArray or offloaded to external storage solutions such as AWS S3, Azure Blob Storage, or generic NFS systems. This dual approach ensures high data availability and resilience, safeguarding against data loss in various disaster scenarios.

Solution Benefits

The integration of ActiveDR and asynchronous snapshot replication offers a robust set of benefits for disaster recovery within SAP environments. By employing ActiveDR near-synchronous replication, businesses can significantly minimize data loss by ensuring continuous data protection to a secondary location with minimal latency. Asynchronous snapshot replication provides added layers of security by allowing periodic snapshots that can be replicated to another FlashArray or offloaded to various storage solutions like AWS S3, Azure Blob Storage, or NFS. This strategy not only enhances flexibility in data recovery options but also helps reduce operational downtime, enabling rapid restoration of services in the event of a disaster. Moreover, the ability to leverage cost-effective cloud storage for snapshot offloading presents a scalable and economical solution, accommodating growing data needs while ensuring compliance and performance requirements are met. This dual approach not only enhances data resilience but also supports scalable, cost-effective disaster recovery planning.



Technical Solution

There are several technical solutions that can be implemented to implement disaster recovery strategies for SAP HANA in an SAP environment using FlashArray data services. Resiliency should be established from business requirements that focus on recovery point (RPO's) and recovery target (RTO) objectives.

Asynchronous replication is the copying of volume snapshots from one FlashArray or Pure Cloud Block Store instance to another to achieve as low as a five minute RPO, ActiveDR provides continuously active replication capabilities for block storage volumes from one FlashArray or Cloud Block Store instance to another, with an objective to implement near zero




FIGURE 10 Asynchronous replication and snapshot offload workflows

ActiveDR and asynchronous replication can be combined to provide the best of both scenarios where near zero and five minute RPOs can happen simultaneously. This is achieved by adding a protection group with a snapshot-and-replicate schedule to an existing ActiveDR Pod. ActiveDR can be combined with SAP HANA system replication to efficiently extend third site resiliency.




FIGURE 11 Near Synchronous replication with ActiveDR.

High Availability Strategies Using Synchronous Active/Active Configurations with ActiveCluster

Overview

ActiveCluster offers a robust solution designed for SAP HANA storage availability, providing a seamless approach to business continuity specifically tailored for business-critical systems. It facilitates the replication of storage data, ensuring operational resilience and continuity during system and site failure scenarios. Suitable for both single and multiple host SAP HANA systems within Tailored Datacenter Integration (TDI) deployments, ActiveCluster stands out by delivering a zero-recovery time objective (RTO) alongside an Active/Active configuration, which means that there is no delay between a failure occurrence and its resolution. Additionally, the solution can be configured to achieve a zero-recovery point objective (RPO), further enhancing the system's resilience by continuously mirroring data so that no data loss occurs in the event of a disaster.

Solution Benefits

Implementing FlashArray with ActiveCluster for SAP HANA offers a robust solution for ensuring high availability and data protection. This approach features a zero recovery time objective (RTO) and recovery point objective (RPO), meaning there is no downtime or data loss in the event of a system failure, crucial for environments where continuous access is critical. The integration simplifies management through a user-friendly interface and reduces costs by eliminating the need for additional software licenses. It also offers scalability to meet growing business demands without compromising performance. Additionally, ActiveCluster enhances data security by supporting geographically dispersed data replication, which protects against site-specific risks. Overall, this solution provides a cost-effective, efficient, and secure method to maintain continuous operations and safeguard critical data in SAP HANA environments.

Technical Solution

The technical underpinnings of ActiveCluster include the deployment of two FlashArray storage arrays, which can be configured in various modes such as optimized or active/active. In synchronous mode, data changes are immediately mirrored between the primary and secondary arrays, ensuring data integrity and continuity without waiting for disk write confirmation, which allows for instantaneous failover and access without any impact on performance.




FIGURE 12 Storage replication with ActiveCluster maintaining availability for an SAP HANA instance

For scenarios requiring distance tolerance or involving multiple data sites, asynchronous replication can be implemented to allow changes to be captured periodically to a secondary or tertiary array, facilitating disaster recovery with minimal impact on primary site operations.



Rapid Recovery with Backint-certified Data Protection Solutions

Overview

Organizations using SAP HANA require robust data protection solutions to ensure rapid recovery in the event of data loss or corruption. The combination of SAP HANA with Pure Storage and Backint certified data protection solutions offers a powerful reference architecture designed to meet these needs. This solution ensures that businesses can minimize downtime and protect their critical data assets efficiently.

Solution Benefits

The integration of SAP HANA with Pure Storage and Backint certified data protection solutions offers a multitude of benefits that are crucial for modern businesses. This architecture ensures rapid recovery, minimizing downtime and enabling swift restoration of data to maintain business continuity and meet service level agreements (SLAs). Pure Storage's all-flash arrays deliver exceptional performance, ensuring that backup and recovery processes do not impact the operational efficiency of the SAP HANA environment. The solution is highly scalable, allowing organizations to expand their data protection capabilities seamlessly as their data volumes grow, thus adapting to evolving business needs without significant reconfiguration. Additionally, the integration simplifies the backup and recovery process, reducing complexity and administrative overhead, which allows IT teams to focus on strategic initiatives. Furthermore, leveraging the efficiency and performance of Pure Storage results in a cost-effective solution, reducing the total cost of ownership (TCO) through savings on storage costs, minimized downtime, and lower administrative expenses.

Technical Solution

The implementation of SAP HANA, Pure Storage, and Backint certified data protection solutions involves several key components and processes:

- Data protection solution: SAP HANA is integrated with the 3rd party data protection solution. These solutions need to be [SAP HANA Backint certified](#) and also support media repositories such as 3rd party files, object or locally attached block storage inside media servers.
- Pure Storage® FlashBlade file or object storage or FlashArray block or file storage connected to by the Data Protection ISV Solution.
- SAP HANA instance: One or more SAP HANA instance, either scale up or out, connected to by the data protection solution.




FIGURE 13 Backint backup workflow with data protection, SAP HANA and storage target components



Design Validation

This design is validated using the [SAP HANA Hardware and Cloud Measurement Tools \(HCMT\)](#). These tools help customers and partners optimize their hardware or cloud systems before deploying SAP HANA or applying for the SAP HANA certification.

During storage certification two major aspects are measured using this tooling:

- Performance measured as an outcome in terms of the number of SAP HANA nodes that consistently meet the performance KPIs
- Scale out failover using a storage connector. When a node in an SAP HANA landscape goes offline, the storage connector is used to fence the storage for that node, move it over to a standby node and then bring that node online.

The number of nodes achieved per FlashArray model varies and is subject to change as new updates are released. For further details on the number of supported nodes per model see the [Certified and supported SAP HANA Hardware directory](#).

Deployment

This section provides detailed guidance for deploying SAP HANA on FlashArray, including installation, and configuration of key components.

FlashArray

This deployment guidance covers the core steps in storage provisioning for SAP HANA scale up instance. It assumes the array has been deployed but that specific configuration for the technology components used for this reference architecture has not been performed.

Block Storage over FCP

Follow the steps below to provision block storage for an SAP HANA scale up instance with FlashArray. It assumes that a host has already been created on FlashArray and the associated fiber channel connectivity and zoning has been completed.

1. In the FlashArray graphical user interface navigate to the **Storage** view and select the **Volumes** tab. To create a new volume, select the **+** in the top right-hand corner of the volumes list.





FIGURE 14 The FlashArray graphical user interface, Storage view



- In the dialog box that appears, provide a name for the volume and provide a provisioned size (capacity) the volume should conform to. Once complete select **Create**. Repeat this step for all volumes for the SAP HANA instance.



The 'Create Volume' dialog box is shown. It has a title bar with a close button. The 'Pod or Volume Group' field is set to 'none'. The 'Name' field is 'SAPHANA-Data'. The 'Provisioned Size' field is '4096' with a unit dropdown set to 'G'. Below these fields are two optional sections: 'QoS Configuration (Optional)' and 'Protection Configuration (Optional)', both with downward arrows. At the bottom are three buttons: 'Create Multiple...', 'Cancel', and 'Create' (which is highlighted with a red rectangle).

FIGURE 15 The Create Volume Dialog

- Review the volumes and provisioned capacities as required.

Volumes ^

Space

QoS

Details


14 of 4 +

Name ^	Size	Virtual	Unique	Snapshots	Reduction	
HANA						
SAPHANA-Data	4 T	0.00	0.00	0.00	1.0 to 1	
SAPHANA-Host-SuSE-Boot	384 G	792 G	1.50 G	0.00	3.9 to 1	
SAPHANA-Log	512 G	0.00	0.00	0.00	1.0 to 1	
SAPHANA-Shared	384 G	0.00	0.00	0.00	1.0 to 1	

Destroyed (0) ^

FIGURE 16 The list of volumes for the SAP HANA instance.

- This step assumes the host has already been created. To connect the volumes to a host, once the volumes are created navigate to the **Hosts** tab. In the hosts tab select the SAP HANA scale up host, in its context view select the three ellipses in the top right-hand corner of **Connected Volumes** and select **Connect**.



The screenshot shows the 'Storage' section of the Pure Storage interface. The 'Hosts' tab is selected, showing a list of hosts. The 'SAP-HANA-Host' is selected, and its context view is displayed. The 'Connected Volumes' section shows a list of volumes. A red rectangle highlights the 'Connect...' button in the context menu. The 'Host Ports' section shows a list of ports. The 'Details' section shows 'CHAP Credentials', 'Personality' (ESXi), and 'Preferred Arrays'.

FIGURE 17 The context view of the SAP HANA Host.

5. In the dialog box that appears select the volumes to connect to the host and then select connect.




FIGURE 18 The Connect Volumes to Host dialog box

6. Once the volumes are connected review the Connected Volumes list in the Hosts context view

Connected Volumes ^				14 of 4	
Name ▲	Shared	LUN			
SAPHANA-Data	True	2			
SAPHANA-Host-SuSE-Boot	True	1			
SAPHANA-Log	True	3			
SAPHANA-Shared	True	4			

FIGURE 19 The connected volumes in the host's view

File Storage over NFS

Follow the below steps to provision file (NFS) storage for an SAP HANA scale up instance with FlashArray.

1. FlashArray file services use a virtual interface with an assigned IP address. The support documentation for [Configuring and Enabling the Virtual Interface for File](#) provides guidance for how to perform this step. Once configured there should be a virtual interface for the file service listed in the network interfaces list.

```

pureuser@flasharray> purenetwork eth list
Name      Enabled Subinterfaces Type Subnet Address      Mask      Gateway      MTU  MAC      Speed      Ser-
vices
ct0.eth0  True   physical -      10.21.227.131 255.255.255.0 10.21.227.1 1500 24:a9:37:05:0d:8a 1.00 Gb/s  manage-
ment -
ct0.eth1  True   physical -      10.21.227.133 255.255.255.0 10.21.227.1 1500 24:a9:37:05:0d:8b 1.00 Gb/s  manage-
ment -
ct0.eth4  True   physical -      10.21.227.135 255.255.255.0 10.21.227.1 9000 24:a9:37:05:0d:8f 25.00 Gb/s  isc-
si -
ct0.eth5  True   physical -      10.21.227.137 255.255.255.0 10.21.227.1 9000 24:a9:37:05:0d:8e 25.00 Gb/s  isc-
si -
ct1.eth0  True   physical -      10.21.227.132 255.255.255.0 10.21.227.1 1500 24:a9:37:05:0d:36 1.00 Gb/s  manage-
ment -
ct1.eth1  True   physical -      10.21.227.134 255.255.255.0 10.21.227.1 1500 24:a9:37:05:0d:37 1.00 Gb/s  manage-
ment -
ct1.eth4  True   physical -      10.21.227.136 255.255.255.0 10.21.227.1 9000 24:a9:37:05:0d:3b 25.00 Gb/s  isc-
si -
ct1.eth5  True   physical -      10.21.227.138 255.255.255.0 10.21.227.1 9000 24:a9:37:05:0d:3a 25.00 Gb/s  isc-
si -
filevif   True   vif -      10.21.227.139 255.255.255.0 10.21.227.1 1500 3e:f7:99:f2:d5:5b 25.00 Gb/s  file
ct0.eth4
ct1.eth4
ct0.eth5
ct1.eth5
replbond  False  bond -      -      -      -      -      1500 6a:f4:9b:a8:96:e9 0.00 b/s  repli-
cation -
vir0      True   vif -      10.21.227.130 255.255.255.0 10.21.227.1 1500 6a:3f:b2:a1:d4:af 1.00 Gb/s  manage-
ment -
vir1      False  vif -      -      -      -      -      1500 e2:3a:79:2e:ab:f7 1.00 Gb/s  manage-
ment -

```

- It is advised that export policies and required quotas are set prior to configuring filesystems. Using the FlashArray graphical user interface navigate to the Storage view and select **Policies**. Identify the section for Export Policies and select the **+** in the top right-hand corner to create a new policy.




FIGURE 20 The policies tab in the storage view

- In the Create Export Policy dialog, provide a Name for the new policy and ensure User Mapping is disabled. Once complete select **Create**.




FIGURE 21 The Create Export Policy Dialog



3a. Selecting the new policy from the list of policies will open its configuration view. A new access rule needs to be added to the policy by selecting the **+** in the top right hand corner of the **Rules** section.




FIGURE 22 The policy configuration view

3b. In the Add Rule for Policy dialog, specify the various NFS access requirements. **No-root-squash** is required.




FIGURE 23 The Add Rule for Policy dialog

4. Navigate to the **File Systems** tab and then identify the File Systems section, selecting the **+** in the top right-hand corner to create a new file system. This filesystem will contain all the directories for the SAP HANA instance.




FIGURE 24 The File Systems tab in the Storage view

5. In the Create File System dialog provide a name for the new file system and then select **Create**.




FIGURE 25 The Create File System Dialog

6. Once the File System has been created it should show a root directory in the **Directories** section. Select the **+** in the top right-hand corner to create a new directory in the File System created in step 5.





FIGURE 26 The Directories section with the default root directory for the new filesystem


7. In the dialog that appears select the File System created in step 5, then provide a name and path for the directory. Select **Create** when complete. Repeat for all directories that need to be created for the SAP HANA instance.



The 'Create Directory' dialog box contains three input fields: 'File System' with the value 'HANA-New', 'Name' with the value 'Data', and 'Path' with the value '/data'. At the bottom are 'Cancel' and 'Create' buttons.

FIGURE 27 The Create directory dialog


8. Once the directories have been created they can be seen in the Directory section.



Directories ^			General	Space	Usage	14 of 4	+	:
Name ^	Path	File System						
HANA-New								
HANA-New:Data	/data	HANA-New						
HANA-New:HANA-Log	/log	HANA-New						
HANA-New:HANA-Shared	/shared	HANA-New						
HANA-New:root	/	HANA-New						
Destroyed (0) v								

FIGURE 28 The directory section with the new Directories created in step 6 and 7 listed.


9. Once the directories have been created identify the section for **Directory Exports**. Select the **+** in the top right-hand corner to create a new directory export.



The 'Directory Exports' section shows a table with columns: Name, Directory, Path, Policy, Type, and Enabled. A red box highlights a '+' button in the top right corner. Below the table, it says 'No directory exports found.'

FIGURE 29 The Directory Exports section

10. In the Create Exports dialog select the directory to create an export for, assign it a policy (the same policy created in step 2 and 3), and provide the export with a name and then select **Create**. Repeat this step for each directory being exported to the SAP HANA instance.



The 'Create Exports' dialog box contains four input fields: 'Directory' with the value 'HANA-New:Data', 'NFS Policy' with the value 'hana-new-policy', 'SMB Policy' with the value '-- Select a SMB policy --', and 'Export Name' with the value 'hana-new-data'. Below the 'Export Name' field is a small text note: 'Name used to mount this path for clients to access'. At the bottom are 'Cancel' and 'Create' buttons.

FIGURE 30 The Create Exports dialog



11. Review the Directory Exports list and ensure the exports and policies meet the requirements for the SAP HANA instance.

Directory Exports						1-3 of 3 +
Name ▲	Directory	Path	Policy	Type	Enabled	
hana-new						
hana-new-data	HANA-New:Data	/data	hana-new-policy	nfs	true	🗑
hana-new-log	HANA-New:HANA-Log	/log	hana-new-policy	nfs	true	🗑
hana-new-shared	HANA-New:HANA-Shared	/shared	hana-new-policy	nfs	true	🗑

FIGURE 31 The Directory Exports Section.

Host Operating System

SAP HANA can be deployed on Red Hat Enterprise Linux (RHEL) and SUSE Enterprise Linux (SLES). For general information on how to configure each operation system please review the below SAP Notes.

- [SAP Note 1944799: SAP HANA Guidelines for SLES Operating System Installation](#)
- [SAP Note 2009879: SAP HANA Guidelines for Red Hat Enterprise Linux \(RHEL\) Operating System](#)

Further information on supported operating systems and revisions for SAP HANA can be found in SAP Note 2235581.

The following optimizations can be applied to the host operating system to maximize the performance and experience of the storage:

- Follow the [Linux Recommended Settings](#) guide to optimize the FlashArray block storage deployment.
- Tune the system with [saptune for SUSE](#) or [RHEL tuned profiles for SAP](#).

Conclusion

This reference architecture provides a robust and comprehensive framework for deploying SAP HANA on Pure Storage solutions, addressing critical business needs for performance, scalability, and reliability. By integrating advanced technical solutions such as volume snapshots, ActiveDR, asynchronous snapshot replication, and ActiveCluster, the architecture ensures high data availability, robust data protection, and rapid recovery capabilities. These solutions collectively minimize downtime, safeguard data integrity, and enhance operational efficiency, making them indispensable for maintaining business continuity in SAP HANA environments.

The strategic implementation of these technologies not only optimizes storage usage and reduces the total cost of ownership but also simplifies management, allowing organizations to focus on their core business objectives. By leveraging the capabilities of Pure Storage, businesses can achieve a resilient, high-performing, and scalable SAP HANA infrastructure that supports growth and adapts to evolving demands.

It empowers organizations to enhance their SAP HANA deployments, ensuring they are well-equipped to handle the challenges of modern data management and business operations. The solutions presented in this paper provide a clear path to achieving a dependable, efficient, and future-proof SAP HANA environment, driving business success and continuity.