

Lenovo ThinkSystem Servers & DE Series / DM Series Enterprise Storage Arrays for SAP HANA

Architecture Overview

12th December 2019

Vijayakumar S Kulageri
Irene Hopf



Abstract

This document outlines the architecture of SAP HANA solution built with Lenovo ThinkSystem® servers and DE® / DM® Series Enterprise storage systems.

This solution excels in flexibility, scalability and differentiating data protection capabilities in business critical environments. It outlines practical architectures to implement production, disaster recovery, development and testing for an analytical (e.g. SAP BW) or transactional (e.g. SAP S/4 HANA) scenario.

This whitepaper is intended for SAP customer, business partners and Lenovo Architects, who design datacenter infrastructures for SAP HANA based application landscapes.

Acknowledgements

The authors express sincere Thanks to the NetApp engineers Nils Bauer and Marco Schön from the SAP Competence Center for their contributions to this document.

If you have feedback or questions please contact us at

SAPsolutions@lenovo.com or ng-sapcc@netapp.com



Table of Contents

Abstract.....	2
Acknowledgements.....	2
Table of Contents.....	3
Executive Summary	5
Introduction	6
Lenovo ThinkSystem SR950.....	6
Lenovo ThinkSystem SR850.....	7
Lenovo ThinkSystem DE Series Storage Arrays	7
Lenovo ThinkSystem DE4000H Storage Array	7
Architecture.....	9
Lenovo ThinkSystem DE4000F All Flash Storage Array	11
Lenovo ThinkSystem DE6000H Hybrid Storage Array.....	11
Architecture.....	12
Lenovo ThinkSystem DE6000F All Flash Storage Array	13
Lenovo ThinkSystem DM Series Storage Arrays.....	14
Lenovo ThinkSystem DM5000F Unified Flash Storage Array	14
Lenovo ThinkSystem DM7000F Unified Flash Storage Array	15
Architecture	16
SAP HANA Backup.....	18
Backup Operations.....	19
60 – 100 Times Faster Backup Operations.....	19
Reduced System Downtime	20
SAP HANA Disaster Recovery	21
Storage Replication	22
Disaster Recovery Resources for Development and Testing	22

Disaster Recovery Failover Testing	23
SAP HANA System Replication	23
Lifecycle Management for SAP HANA	23
The Solution	23
Fully Automated, End-to-End Provisioning of SAP Development and Test Systems.....	24
Semiautomatic Provisioning of SAP HANA Development and Test Systems.....	25
Fast and Space-efficient SAP System Provisioning Using Storage Cloning Technology	25
Solution Components	25
Sample Scenarios.....	26
SAP BW/4 HANA Architecture.....	26
SAP S/4 HANA Architecture	27
Summary	28
Resources.....	29
Lenovo Documentation	29
NetApp Documentation	29
SAP Documentation.....	29
Legal.....	30

Executive Summary

In today's fast-paced world, you need to know what is going on with your business, your marketplace, and your customers so you can respond quickly. Many enterprises are looking to obtain real business insights through SAP HANA technology; for example, using SAP BW/4HANA and S/4HANA. However, the move to SAP HANA can be complex, so choosing your business and technology partners wisely is an important first step in this journey.

Lenovo is working closely with SAP to provide you with seamless, secure, context-aware experiences for your SAP environment. For SAP HANA, Lenovo offer customers a complete set of fully integrated solutions and services with Lenovo ThinkSystem servers and DE/DM series Enterprise Storage.

Within this document, we focus on deploying SAP HANA using the Tailored Data Center Integration (TDI) approach. With this approach, you can connect SAP HANA Database servers to SAP HANA-certified enterprise storage systems. This allows homogeneous integration into existing infrastructure, management and operational model and a faster overall HANA adoption. With this approach, you can leverage existing datacenter processes, as well as reusing existing IT infrastructure.

The SAP HANA configuration outlined in this white paper is based on Lenovo ThinkSystem servers and Lenovo ThinkSystem DE and DM Series enterprise storage systems. The solution described in this white paper combines the technologies of Lenovo ThinkSystem Mission critical servers and Lenovo DE/DM Series enterprise storage based on NetApp® ONTAP platform, to create a best in class, economic, feature-rich, and simple to operate SAP HANA implementation.

Introduction

This document describes the SAP HANA deployment architecture using SAP HANA supported Lenovo servers (ThinkSystem SR950/SR850) components and storage (ThinkSystem DE4000H/DE6000H/DM5000F/7000F) controllers. In an SAP landscape, DE storage systems support SAN environments ((iSCSI and Fibre Channel) and DM storage supports only NAS environments (NFS). It is applicable to all SAP HANA-based solutions, such as SAP BW/4HANA and S/4HANA.

Lenovo ThinkSystem SR950

Lenovo ThinkSystem SR950 is designed for your most demanding, mission-critical workloads, such as in-memory databases, large transactional databases, batch and real-time analytics, ERP, CRM, and virtualized server workloads. The powerful 4U ThinkSystem SR950 can grow from two to 8 Intel Xeon Scalable Family processors, and with 96 DIMM sockets, supports technically up to 24 TB of high-speed on DRAM & up to 36TB with the use of Intel™ Optane® DC Persistent Memory. The modular design of SR950 speeds upgrades and servicing with easy front or rear access to all major subsystems to maximize server availability.

The SR950 packs numerous fault-tolerant and high-availability features into a high-density, 4U rack-optimized design that reduces the space needed to support massive network computing operations and simplify servicing.



FIGURE 1 : LENOVO THINKSYSTEM SR950

Lenovo ThinkSystem SR850

The Lenovo ThinkSystem SR850 is a 4-socket server that features a streamlined 2U rack design that is optimized for price and performance, with best-in-class flexibility and expandability. Models of the SR850 are powered by up to four Intel Xeon Processor Scalable Family processors, each with up to 28 cores, for an efficient 4-socket solution. It can support up to 6TB memory on DRAM & with the use of Intel™ Optane® DC persistent memory it can support up to 15TB of memory. The ThinkSystem SR850's agile design provides rapid upgrades for processors and memory, and its large, flexible storage capacity helps to keep pace with data growth.



FIGURE 2 : LENOVO THINKSYSTEM SR850

The SR850 has space for 16x 2.5-inch drive bays, up to 8 of which can be configured as AnyBay drives - supporting SAS, SATA or NVMe drives. NVMe drives are high-speed, low-latency storage, ideal for storage tiering. The server also offers easy expansion, with up to nine PCIe slots (two of which are x16) and up to 2 (mirrored) M.2 drives for rapid OS boot, faster than either USB keys or SD cards.

Lenovo ThinkSystem DE Series Storage Arrays

Lenovo ThinkSystem DE4000H Storage Array

Lenovo ThinkSystem DE4000H is a scalable, hybrid entry-level storage system that is designed to provide performance, simplicity, capacity, security, and high availability for medium to large businesses. It delivers enterprise-class storage management capabilities with a wide choice of host connectivity options, flexible drive configurations, and enhanced data management features. The ThinkSystem DE4000H is a perfect fit for a wide range of enterprise workloads,

including big data and analytics, video surveillance, technical computing, backup and recovery, and other storage I/O-intensive applications.

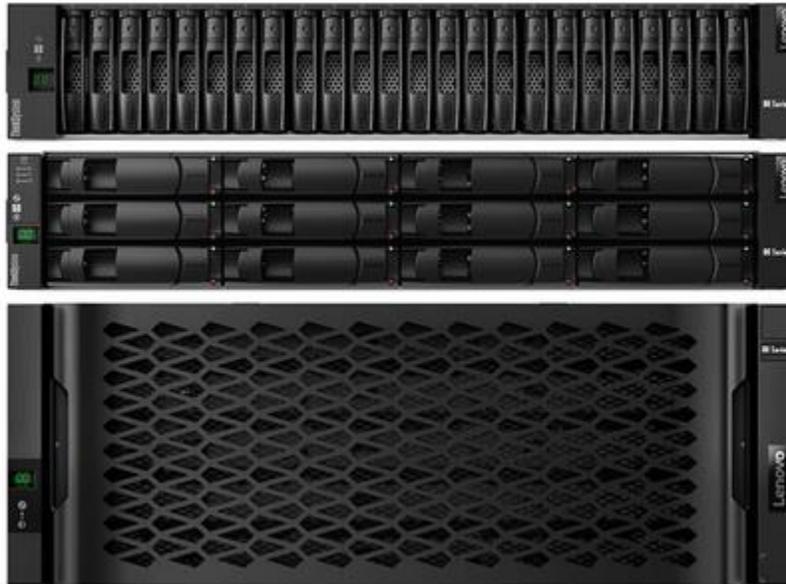


FIGURE 3 : LENOVO THINKSYSTEM DE4000H 2U24 SFF (TOP), 2U12 LFF (MIDDLE), AND 4U60 LFF (BOTTOM)

ThinkSystem DE4000H models are available in a 2U rack form-factor with 24 small form-factor (2.5-inch SFF) drives (2U24 SFF), 12 large form-factor (3.5-inch LFF) drives (2U12 LFF), or a 4U rack form-factor with 60 LFF drives (4U60 LFF) and include two controllers, each with 8 GB or 32 GB cache for a system total of 16 GB or 64 GB. Universal 1/10 Gb iSCSI or 4/8/16 Gb Fibre Channel (FC) ports provide base host connectivity, and the host interface cards provide additional 1/10 Gb iSCSI or 4/8/16 Gb FC, 12 Gb SAS, 10/25 Gb iSCSI, or 8/16/32 Gb FC connections.

The ThinkSystem DE4000H Storage Array scales up to 192 drives with the attachment of Lenovo ThinkSystem DE120S 2U12, DE240S 2U24 SFF, and DE600S 4U60 LFF Expansion Enclosures. It also offers flexible drive configurations with the choice of 2.5-inch (SFF) and 3.5-inch (LFF) form factors, 10 K rpm SAS and 7.2 K rpm NL SAS hard disk drives (HDDs), and SAS solid-state drives (SSDs).

Architecture

Using FC Protocol – SAN

SAP HANA nodes are connected to storage controllers by using a redundant 16 Gigabit FC (also available with 8/32 Gbps) connections. A redundant SAN switch fabric is required to provide a fault-tolerant SAP HANA node to storage connectivity in case of switch or FC HBA (Host Bus Adapter) failure.

Similar to the NAS storage, the total number of hosts allowed to be connected to single SAN storage is dependent on many KPI's. Please refer to [SAP HANA Hardware Directory](#) for the actual & latest details.

Figure 4 shows an example configuration with four SAP HANA ThinkSystem SR850 nodes with different SAP HANA roles connected to DE4000H.

The architecture can be scaled in two dimensions:

- By scaling *up* the storage system and attaching additional SAP HANA nodes and storage capacity to the storage, as long as the storage controllers provide enough performance to meet the KPIs.
- By scaling *out* the storage system and adding more storage controllers and capacity for the additional SAP HANA nodes

Being a hybrid storage DE4000H brings in the best of both the worlds for SAP workloads, where in SAP HANA database itself can be hosted on an all flash Dynamic disk pool (DDP) and the non-HANA workloads like application servers, non-HANA database systems can be hosted on a HDD volume as well.

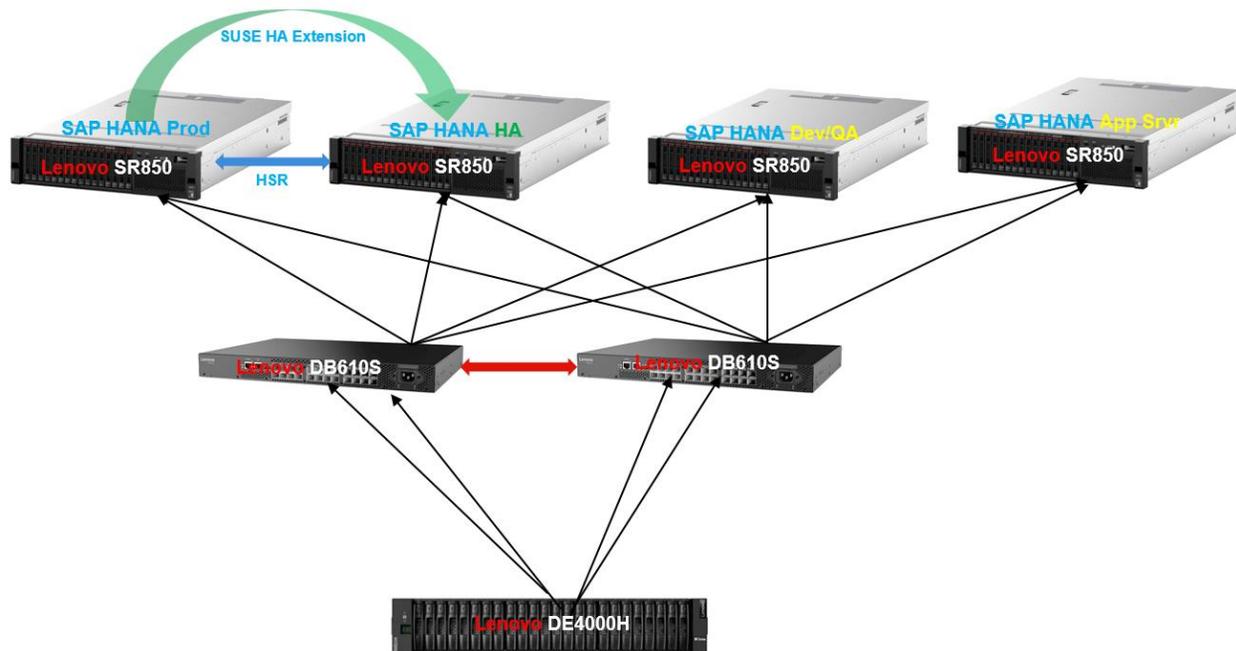


FIGURE 4 : SAP HANA LANDSCAPE USING DE4000H

In the above example which depicts a SAP HANA landscape, a single DE4000H hosts an SAP HANA production database server which has a redundant pair (high availability) enabled by SAP HANA system replication and an automated failover design which is based on SUSE® High Availability extension which is based on pacemaker clustering software.

The supporting instances like Dev & QA are also running the same storage array which makes a data refresh process like SAP system copy easy to manage. In addition to the DB servers, the DE4000H array can be used to host the application stack as well, in this type of deployment all considerations regarding performance aspects of a HANA production DB and an application server needs to be considered. For instance, Lenovo recommends a DDP (Dynamic disk pool which helps improve performance and availability with significantly faster rebuild time and reduced exposure to multiple drive failures by allowing data and built-in spare capacity to be distributed across all physical drives in the storage pool) with a minimum 8+2+1 disk configuration.

If the application layer is being hosted on the same array then the recommended way to size the array is to consider a separate DDP or a VG (Volume Group) with required amount of storage space on HDD's or SSD's.

Lenovo ThinkSystem DE4000F All Flash Storage Array

Lenovo ThinkSystem DE4000F is a scalable, all flash entry-level storage system that is designed to provide performance, simplicity, capacity, security, and high availability for medium to large businesses. It delivers enterprise-class storage management capabilities with a wide choice of host connectivity options, flexible drive configurations, and enhanced data management features. The ThinkSystem DE4000F is a perfect fit for a wide range of enterprise workloads, including big data and analytics, video surveillance, technical computing, and other storage I/O-intensive applications.



FIGURE 5 : LENOVO THINKSYSTEM DE4000F 2U24 SFF ENCLOSURE

ThinkSystem DE4000F models are available in a 2U rack form-factor with 24 small form-factor (2.5-inch SFF) drives (2U24 SFF) and include two controllers, each with 8 GB or 32 GB cache for a system total of 16 GB or 64 GB. Universal 1/10 Gb iSCSI or 4/8/16 Gb Fibre Channel (FC) ports provide base host connectivity, and the host interface cards provide additional 1/10 Gb iSCSI or 4/8/16 Gb FC, 12 Gb SAS, 10/25 Gb iSCSI, or 8/16/32 Gb FC connections.

The ThinkSystem DE4000F Storage Array scales up to 192 solid-state drives (SSDs) with the attachment of Lenovo ThinkSystem DE240S 2U24 SFF Expansion Enclosures.

Lenovo ThinkSystem DE6000H Hybrid Storage Array

Lenovo ThinkSystem DE6000H is a scalable, hybrid mid-range storage system that is designed to provide high performance, simplicity, capacity, security, and high availability for medium to large businesses. The ThinkSystem DE6000H delivers enterprise-class storage management capabilities in a performance-optimized system with a wide choice of host connectivity options, flexible drive configurations, and enhanced data management features. The ThinkSystem DE6000H is a perfect fit for a wide range of enterprise workloads, including big data and analytics, video surveillance, technical computing, backup and recovery, and other storage I/O-intensive applications.



FIGURE 6 : LENOVO THINKSYSTEM DE6000H 2U24 SFF (TOP) AND 4U60 LFF (BOTTOM) ENCLOSURES

ThinkSystem DE6000H models are available in a 2U rack form-factor with 24 small form-factor (2.5-inch SFF) drives (2U24 SFF) or a 4U rack form-factor with 60 LFF drives (4U60 LFF) and include two controllers, each with 16 GB or 64 GB cache for a system total of 32 GB or 128 GB. Universal 10 Gb iSCSI or 4/8/16 Gb Fibre Channel (FC) ports provide base host connectivity, and the host interface cards provide additional 12 Gb SAS, 10/25 Gb iSCSI, or 8/16/32 Gb FC connections.

The ThinkSystem DE6000H Storage Array scales up to 240 (base configuration) or 480 (optional upgrade) drives with the attachment of Lenovo ThinkSystem DE240S 2U24 SFF and DE600S 4U60 LFF Expansion Enclosures. It also offers flexible drive configurations with the choice of 2.5-inch (SFF) and 3.5-inch (LFF) form factors, 10 K rpm SAS and 7.2 K rpm NL SAS hard disk drives (HDDs), and SAS solid-state drives (SSDs).

Architecture

DE6000H with a total controller cache up-to 128GB makes a best fit for a landscape where end users would need a storage subsystem which not only caters for a high write I/O performance of SAP HANA database systems, but suffices the high performance asks of SAP Application servers & non-SAP applications which can be co-existing with SAP landscape.

With a total of 480 drive upgrade option, DE6000H will take care of any capacity based applications like SAP DMS where performance is not a criteria but a cost-effective storage is.

Figure 7 shows an example architecture of using DE6000H in an enterprise SAP landscape.

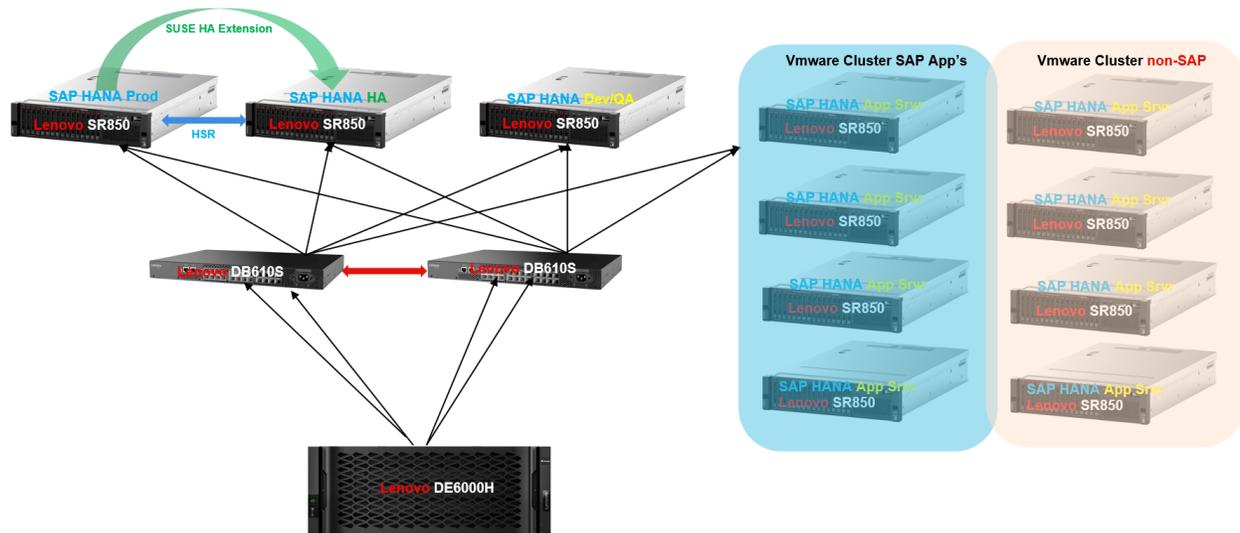


FIGURE 7 : DE6000H FOR SAP & NON-SAP LANDSCAPES

In the above example, Lenovo DE6000H hosts a SAP HANA Database Prod, HA & Dev/QA landscapes & equally a larger ecosystem of SAP Applications servers & non-SAP workload.

Lenovo ThinkSystem DE6000F All Flash Storage Array

Lenovo ThinkSystem DE6000F is a scalable, all flash mid-range storage system that is designed to provide high performance, simplicity, capacity, security, and high availability for medium to large businesses. The ThinkSystem DE6000F delivers enterprise-class storage management capabilities in a performance-optimized system with a wide choice of host connectivity options and enhanced data management features. The ThinkSystem DE6000F is a perfect fit for a wide range of enterprise workloads, including big data and analytics, video surveillance, technical computing, and other storage I/O-intensive applications.



FIGURE 8 : LENOVO THINKSYSTEM DE6000F 2U24 SFF ENCLOSURE

ThinkSystem DE6000F models are available in a 2U rack form-factor with 24 small form-factor (2.5-inch SFF) drives (2U24 SFF) and include two controllers, each with 64 GB cache for a

system total of 128 GB. Universal 10 Gb iSCSI or 4/8/16 Gb Fibre Channel (FC) ports provide base host connectivity, and the host interface cards provide additional 12 Gb SAS, 10/25 Gb iSCSI, or 8/16/32 Gb FC connections.

The ThinkSystem DE6000F Storage Array scales up to 192 solid-state drives (SSDs) with the attachment of Lenovo ThinkSystem DE240S 2U24 SFF Expansion Enclosures.

Lenovo ThinkSystem DM Series Storage Arrays

Lenovo ThinkSystem DM5000F Unified Flash Storage Array

Lenovo ThinkSystem DM5000F is a unified, all flash storage system that is designed to provide performance, simplicity, capacity, security, and high availability for medium enterprises. Powered by the ONTAP software, ThinkSystem DM5000F delivers enterprise-class storage management capabilities with a wide choice of host connectivity options and enhanced data management features. The ThinkSystem DM5000F is a perfect fit for a wide range of enterprise workloads, including big data and analytics, artificial intelligence, engineering and design, enterprise applications, and other storage I/O-intensive applications.

ThinkSystem DM5000F models are 2U rack-mount controller enclosures that include two controllers, 64 GB RAM and 8 GB battery-backed NVRAM (32 GB RAM and 4 GB NVRAM per controller), and 24 SFF hot-swap drive bays (2U24 form factor). Controllers provide universal 1/10 GbE NAS/iSCSI or 8/16 Gb Fibre Channel (FC) ports for host connectivity.



FIGURE 9 : LENOVO THINKSYSTEM DM5000F

A single ThinkSystem DM5000F Storage Array scales up to 144 solid-state drives (SSDs) with the attachment of Lenovo ThinkSystem DM240S 2U24 SFF Expansion Enclosures. Up to 12 DM5000F Storage Arrays can be combined into a clustered system in a NAS environment, or up to 6 DM5000F Storage Arrays can be combined into a clustered system in a SAN environment.

Lenovo ThinkSystem DM7000F Unified Flash Storage Array

Lenovo ThinkSystem DM7000F is a scalable, unified, all flash storage system that is designed to provide high performance, simplicity, capacity, security, and high availability for large enterprises. Powered by the ONTAP software, ThinkSystem DM7000F delivers enterprise-class storage management capabilities with a wide choice of host connectivity options, flexible drive configurations, and enhanced data management features, including support for NVMe over Fabrics. The ThinkSystem DM7000F is a perfect fit for a wide range of enterprise workloads, including big data and analytics, artificial intelligence, engineering and design, hybrid clouds, and other storage I/O-intensive applications.



FIGURE 10 : LENOVO THINKSYSTEM DM7000F

ThinkSystem DM7000F models are 3U rack-mount controller enclosures that include two controllers, and 256 GB RAM and 16 GB battery-backed NVRAM (128 GB RAM and 8 GB NVRAM per controller). Universal 1/10 GbE NAS/iSCSI or 4/8/16 Gb Fibre Channel (FC) ports and 1/10 GbE RJ-45 ports provide base host connectivity, with an option for additional 1/10 GbE or 40 GbE NAS/iSCSI, or 8/16/32 Gb FC connections with the adapter cards.

A single ThinkSystem DM7000F Storage Array scales up to 384 SFF solid-state drives (SSDs) with the attachment of Lenovo ThinkSystem DM240S 2U24 SFF Expansion Enclosures. Up to 12 DM7000F Storage Arrays can be combined into a clustered system in a NAS environment, or up to 6 DM7000F Storage Arrays can be combined into a clustered system in a SAN environment.

All the above-mentioned Server & Storage models are fully supported and are listed in the SAP HANA HW directory on [this](#) page.

Architecture

Using NFS protocol - NAS

SAP HANA nodes are connected to storage controllers by using a redundant Ethernet (10GbE/40GbE) infrastructure. Data communication between SAP HANA nodes and storage controllers is based on the NFS protocol, Fibre Channel and iSCSI is not supported. A

redundant switching infrastructure is required to provide a fault-tolerant SAP HANA node to storage connectivity in case of switch or network interface card (NIC) failure.

The maximum number of SAP HANA nodes attached to the storage is defined by the SAP HANA performance requirements. The number of required disk shelves is determined by the capacity and performance requirements of the SAP HANA systems. The capacity requirements depend on the number of SAP HANA nodes and the RAM size of each node. The storage partitions of the SAP HANA nodes are distributed to the storage controllers. Please refer to [SAP HANA Hardware Directory](#)

- **Note:** The storage and SAP HANA node-to-node communication is separated. Only the storage communication is depicted in Figure 11.

Figure 11: Example configuration with four SAP HANA Lenovo ThinkSystem SR950 nodes attached to a two-node storage cluster.

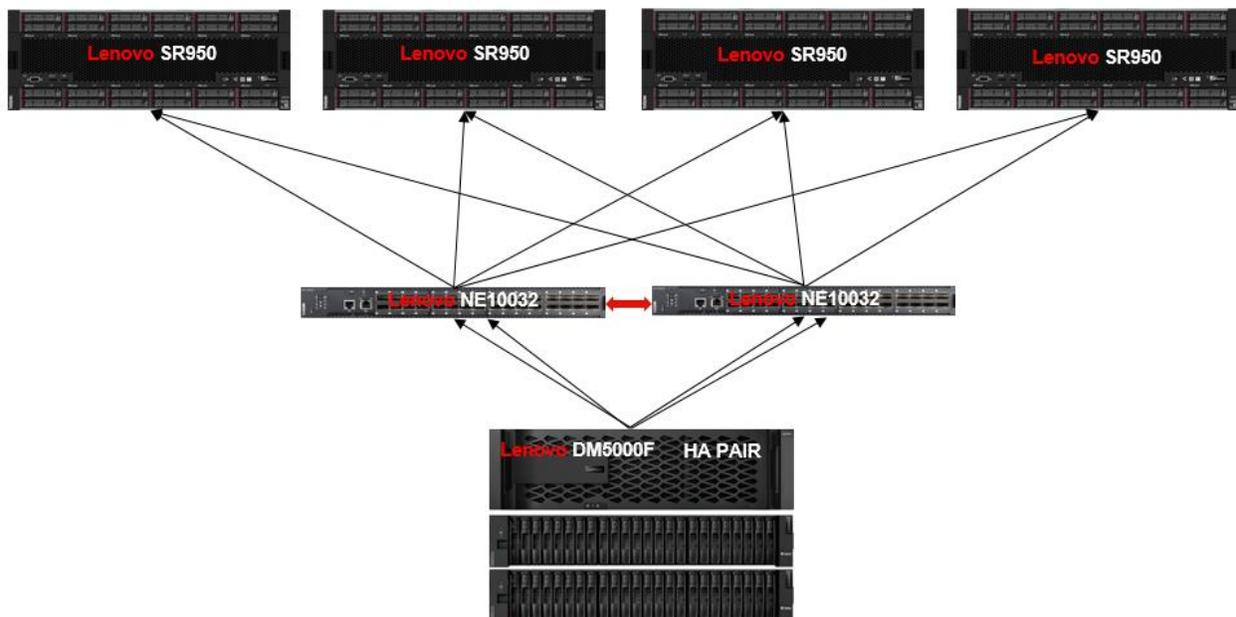


FIGURE 11: EXAMPLE CONFIGURATION WITH FOUR SAP HANA LENOVO THINKSYSTEM SR950 NODES

The architecture can be scaled in two dimensions:

- By scaling *up* the storage system and attaching additional SAP HANA nodes and storage capacity to the storage, as long as the storage controllers provide enough performance to meet the KPIs.
- By scaling *out* the storage system and adding more storage controllers and capacity for the additional SAP HANA nodes

Both, scaling up and scaling out the environment can be performed without any interruption to the running SAP HANA environment, making use of the non-disruptive operations (NDO) built into ONTAP.

SAP HANA Backup

These days, your company needs continuous, uninterrupted availability of your SAP applications. Backing up your SAP databases is a critical task and can significantly affect the performance of your production SAP system. The time that it takes to restore and recover SAP systems is also a concern.

Another challenge that your company faces is logical corruption, which can be caused by software errors, human errors, or sabotage. The worst case is logical corruption in an SAP landscape, where applications exchange data with each other. If you restore a single SAP system to a point in time before the corruption occurred, the result is data loss, and your SAP landscape is no longer synchronized. Instead of restoring the SAP system, to mitigate the logical corruption, you need a clone of the production system that's based on data that was stored before the logical corruption occurred. This allows to fix the corruption on the application or database layer.

With SAP HANA backup and restore operations, your organization faces the following challenges:

- Long backup operations with performance degradation on production SAP systems
- Unacceptable system downtime due to long restore and recovery operations
- Shrinking backup windows because of the criticality of the applications
- The need for a flexible solution to mitigate logical corruption

With Lenovo DM® storage solutions that run ONTAP® data management software, in combination with SnapCenter® data protection software, you can meet all those challenges. In addition, with the Snapshot™ technology that is included in ONTAP software, you can create backups or execute restore operations of any size dataset in a matter of seconds. SAP HANA

supports the use of storage-based Snapshot copies as a valid backup operation with documented interfaces.

Backup Operations

SnapCenter and the plug-in for SAP HANA use ONTAP Snapshot technology and the SAP HANA SQL backup interface to give you an SAP-integrated backup solution. SnapCenter gives you automated workflows for backup operations, including retention management for data backups, for log backups, and for the SAP HANA backup catalog.

And for long-term retention, SnapCenter manages the optional replication of application-consistent backups to an off-site secondary location. Your off-site backup storage can be either a physical storage system on the premises or a Cloud Volumes ONTAP instance that runs in Amazon Web Services (AWS) or in Microsoft Azure.

Figure 12 shows an overview of the solution architecture.

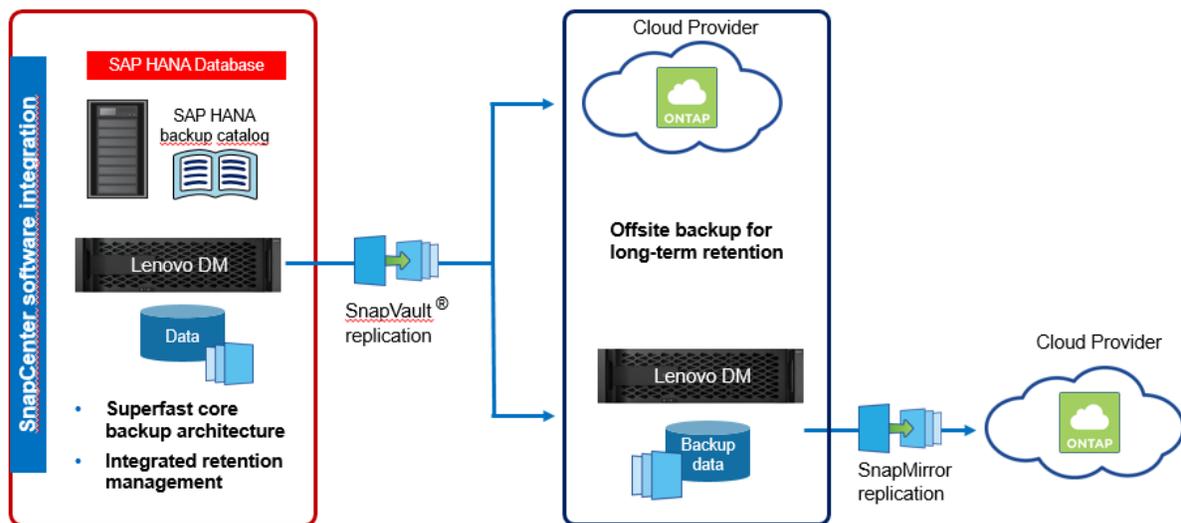


FIGURE 12 : APPLICATION INTEGRATED BACKUP ARCHITECTURE

60 – 100 Times Faster Backup Operations

Evaluation of customer data has shown that for SAP HANA, the average backup time with Snapshot copies is in the range of a few minutes. In the customer scenario in Figure 13, a complete backup for a 2.3TB database took 2 minutes and 11 seconds.

The largest contributor to the overall backup duration is the time that SAP HANA needs to write the synchronized backup savepoint. The amount of time that is required to write the savepoint is a function of the memory of the SAP HANA system and the activity on the system. The storage Snapshot operation is performed in a matter of seconds, independent of the size of the database

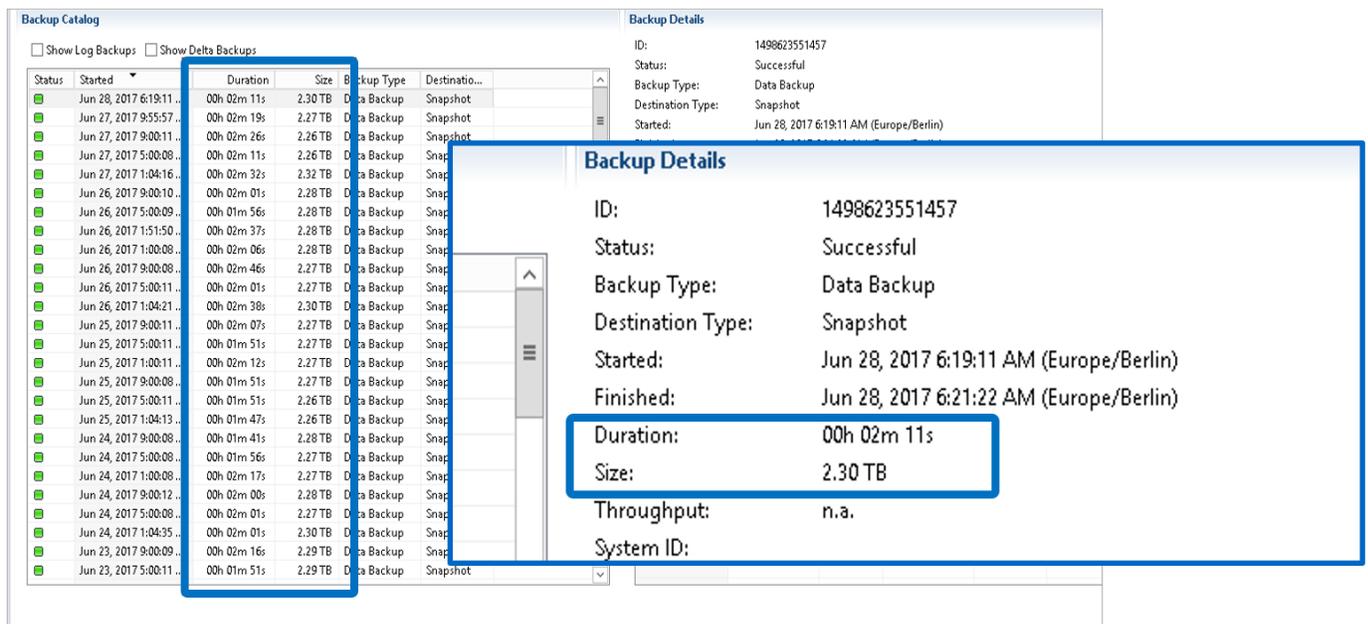


FIGURE 13 : CUSTOMER EXAMPLE: 2.3TB BACKED UP IN 2 MIN 11 SEC

Reduced System Downtime

Because Snapshot copy-based backup operations are super-fast and do not affect system performance, you can schedule multiple Snapshot copy backups daily instead of creating a single daily backup as with traditional streaming backup technology. When a restore and recovery operation is necessary, your system downtime is significantly reduced by two key features. By using SnapRestore® data recovery technology on the storage layer, the restore operation is executed in mere seconds. And because a higher backup frequency results in fewer database logs that need to be applied, the forward recovery is also accelerated.

SAP HANA Disaster Recovery

Lenovo has developed a full portfolio of technologies and tools to help IT organizations build or adapt their disaster recovery plans to respond to all business demands.

These Lenovo & ONTAP technologies constitute an extraordinarily versatile disaster recovery solution for SAP HANA on the market. The solution includes SnapMirror® replication, MetroCluster™ software, and FlexClone® thin-cloning technology. The solution supports:

- Asynchronous and synchronous storage replication
- Replication of non-database data, such as application server binaries
- Use of disaster recovery resources for development and testing
- Use of replicated data to refresh development and testing systems
- Disaster recovery testing based on cloning
- SAP HANA System Replication

The following table compares the disaster recovery use cases and highlights the most important features.

	Storage Replication		SAP HANA System Replication	
	Snap Mirror	Metro Cluster	With Data Preload	Without Data Preload
RTO	Low to medium, depending on database startup time	Low to medium, depending on database startup time	Very low	Low to medium, depending on database startup time
RPO	Synchronous or asynchronous replication	Synchronous replication	Synchronous or asynchronous replication	Synchronous or asynchronous replication
Servers at DR site can be used for dev/test	Yes	Yes	No	Yes
Replication of non -database data	Yes	Yes	No	No
DR data can be used for refresh of dev/test systems	Yes	No	No	No
DR testing without affecting RTO and RPO	Yes	Yes	No	No
DR configuration effort	For each storage volume used by the databases	All databases on the storage system are automatically replicated	For each database	For each database

TABLE 1: COMPARISON OF THE DISASTER RECOVERY (DR) USE CASES

Storage Replication

Storage replication is suitable for low to medium RTO requirements, where it is acceptable for the SAP HANA database to be started and for data to be loaded into memory after a disaster recovery failover. Storage replication is also used to replicate non-database data, such as SAP application server binaries.

SnapMirror® data replication software provides synchronous and asynchronous replication. The replication is configured on the storage volume level.

MetroCluster high-availability and disaster recovery storage software provides synchronous replication that works on the storage system level. All databases that reside on the storage system are automatically replicated without any additional configuration effort.

Disaster Recovery Resources for Development and Testing

With storage replication, the servers at the disaster recovery site can be used for development and testing during normal operation. When you use a SnapMirror based solution, the disaster recovery site can be either on premises or in the cloud, and the replicated data can be used for performing a development and testing & system refresh. Figure 8 shows an example of the same.

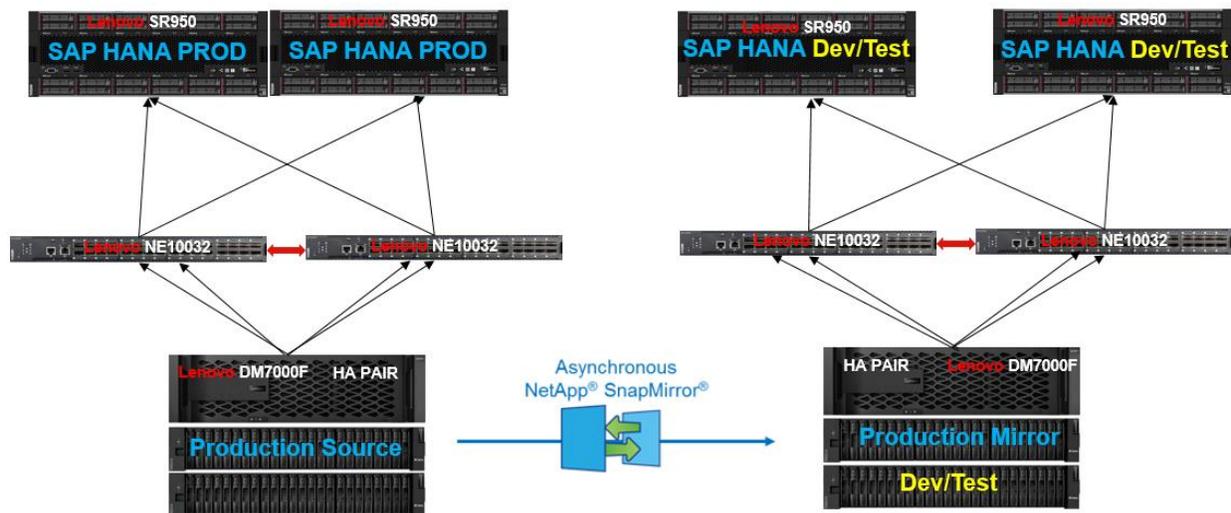


FIGURE 14 : SAP HANA DR WITH STORAGE REPLICATION

Disaster Recovery Failover Testing

Every organization must test its disaster recovery plan. This testing shows whether the system reacts as stipulated in the disaster recovery plan and documentation. With FlexClone technology, you can execute a disaster recovery failover test without influencing or interrupting the ongoing replication to the disaster recovery site. In this way, FlexClone lets you run a test without influencing the RTO or RPO.

SAP HANA System Replication

To achieve very low RTO values with SAP HANA, you must use SAP HANA System Replication with data preloaded into the memory of the dedicated secondary server at the disaster recovery site. Although operating System Replication without preloaded data allows you to use the secondary server at the disaster recovery site for development and testing, this approach increases the RTO in a failover, because the database needs to be started and data needs to be loaded into memory. Also, SAP HANA System Replication does not replicate non-database files. Therefore, to increase resiliency of the whole system, System Replication is typically combined with storage-based replication such as SnapMirror or MetroCluster for non-database data.

Lifecycle Management for SAP HANA

Today, enterprises are looking at ways to increase competitiveness by accelerating projects and speeding time to market. To reach this goal, they need to improve the lifecycle of their enterprise applications by automating tasks and simplifying processes. Traditional SAP lifecycle management approaches to development and test-system provisioning are primarily based on manual processes. These manual processes are often error prone and time consuming, delaying innovation and the ability to respond to business requirements.

The main challenges IT organizations face today are:

- Slow implementation of new features
- Lack of automation
- Lost productivity through lack of integration among orchestration tools

The Solution

Lenovo is addressing these challenges by providing a lifecycle management solution that is fully integrated into the tools that SAP administrators use for day-to-day operations, such as SAP Landscape Management (LaMa). The goal is to simplify the provisioning workflow from preprocessing to postprocessing, including all the software and storage layer tasks needed to

create a copy of the production system. With this solution, administrators can create a development and test environment in a couple of mouse clicks, resulting in improved lifecycle management.

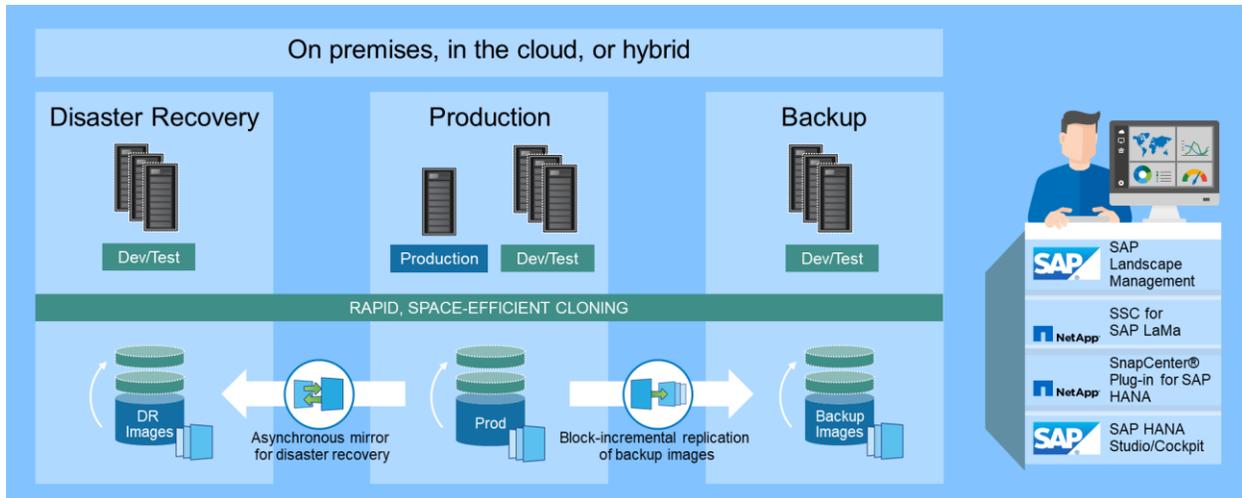


FIGURE 15: OVERVIEW OF SAP LANDSCAPE MANAGEMENT USING ONTAP® FEATURES

Fully Automated, End-to-End Provisioning of SAP Development and Test Systems

Fundamental to full automation of SAP development and test-system provisioning is a deep integration of the storage layer into SAP LaMa. This integration is achieved by deploying the Storage Services Connector (SSC), which acts as a gateway between SAP LaMa and the storage systems. The benefit is that administrators can automate the full workflow from preprocessing to postprocessing within the standard SAP LaMa UI. For instance, storage operations such as cloning can be executed directly from SAP LaMa. SAP LaMa accesses the FlexClone feature through the SSC and can create a copy of a production system on the storage systems.

To increase the solution's flexibility, has added to SSC cloud support. SSC can provision storage clones in the cloud on hyperscalers such as Amazon Web Services (AWS) and Microsoft Azure where Cloud Volumes ONTAP® software is running.

Semiautomatic Provisioning of SAP HANA Development and Test Systems

If you have different administrators taking care of tasks on the three different layers—SAP application, database, and storage—you can use SAP LaMa to control and automate the tasks for SAP systems. Then you can use SnapCenter® software for all the other tasks on the storage and database layer. The SAP HANA plug-in for SnapCenter brings application consistency to the solution and can be used to automate the SAP development and test-system provisioning, including the required steps on the HANA database layer.

Fast and Space-efficient SAP System Provisioning Using Storage Cloning Technology

FlexClone is the part of the solution that creates fast system copies. Traditional copies can take many hours to make. With FlexClone technology, even the largest volumes can be cloned in a matter of seconds. Through innovative technology, a clone uses a small amount of space for metadata and then only consumes..., and then only consumes additional space as data is changed or added. These clones can be created from either the production, disaster recovery, or backup storage system.

Solution Components

The lifecycle management solution is based on three components:

- **SnapCenter.** SnapCenter software is a unified, scalable platform for data protection. SnapCenter provides centralized control and oversight, allowing you to use a single tool to manage backup, restore, and clone operations for a variety of applications and databases., such as SAP HANA.
- **Storage Services Connector.** Implementation of the SAP storage adapter in SAP LaMa.
- **FlexClone.** FlexClone technology allows you to make fast, space-efficient copies of your storage volumes and LUNs. With FlexClone, you can create as many copies of your full production dataset as you need. If a test corrupts the data, you can start again in seconds. Developers and test engineers spend less time waiting for access to datasets and more time doing productive work.

Sample Scenarios

In this section, we would like to bring the two perspectives of server and storage together and give insight into two typical types of SAP HANA projects.

We consider a SAP BW/4 HANA project with development and test chain and a SAP Suite on HANA or S/4HANA with development and test chain with system replication to implement high availability.

SAP BW/4 HANA Architecture

The following example architecture is for a 3TB BW/4 HANA system based on SAP HANA scale-out (3 workers + 1 standby). There is one production system (PRD) with a development instance (DEV) on the left and an asynchronous mirrored system (PRD failover) on the right.

Asynchronous mirroring is performed with SnapCenter and asynchronous storage mirroring of consistent database images. Quality assurance (QAS) and test (TST) system on the right side as well. Those systems are refreshed from the PRD-failover system using cloning technology (FlexClone). The development and test systems can be provided as either virtual or physical.

Either way, an NFS-attached clone provides the data for those systems. Virtual machines attach their data/log by mounting directly from the storage controller.

Figure 16, below, illustrates a 3+1 scale-out configuration with 3 worker nodes.

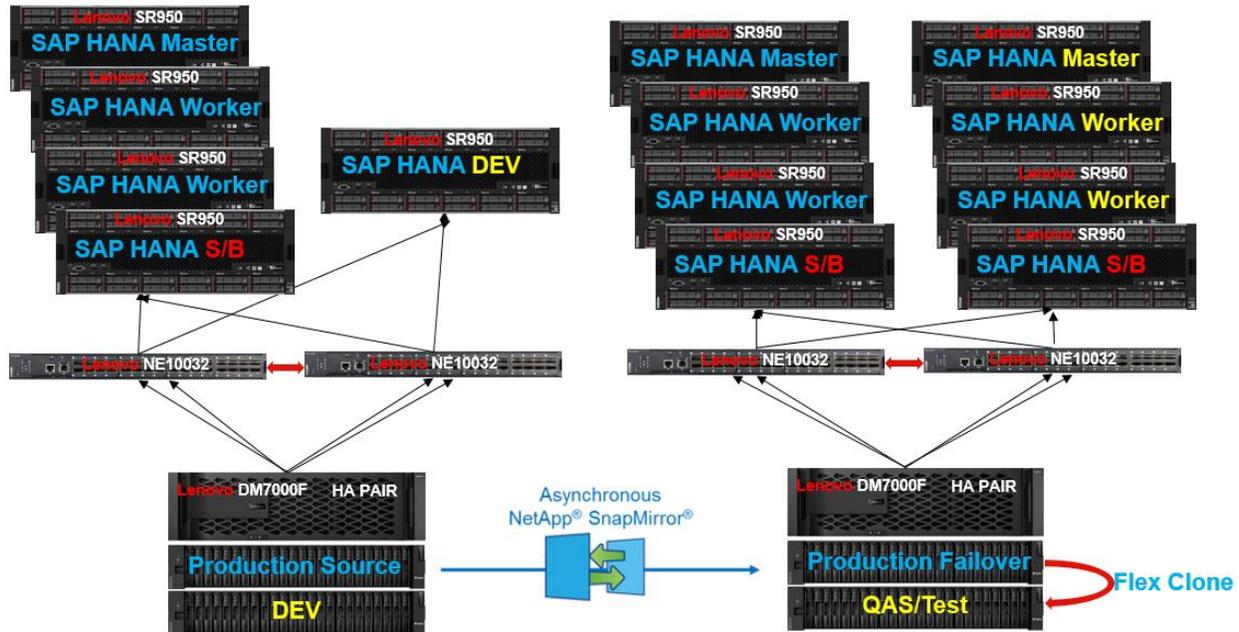


FIGURE 16 : A 3+1 BW/4 HANA SCALE OUT

On both sides, qualified DM5000F/7000F arrays could be used. In the following. Within this architecture, we implement:

- Snapshot-based storage backup for system PRD.
- Asynchronous storage replication from PRD to PRD-failover.
- Development on the left (DC)
- Quality assurance and test on the right (DR)

SAP S/4 HANA Architecture

As a second architecture example, we outline S/4HANA or a Suite-on-HANA system. This system is replicated synchronously using SAP HANA system replication from the production system (PRD) to the failover system (PRD-failover). On the left side, the development system (DEV) is located. The PRD-failover system, the quality assurance (QAS) and test system (TST) are on the right side. Those systems are refreshed using SAP System copy of PRD-failover system. In this setup, we make use of DM7000F storage.

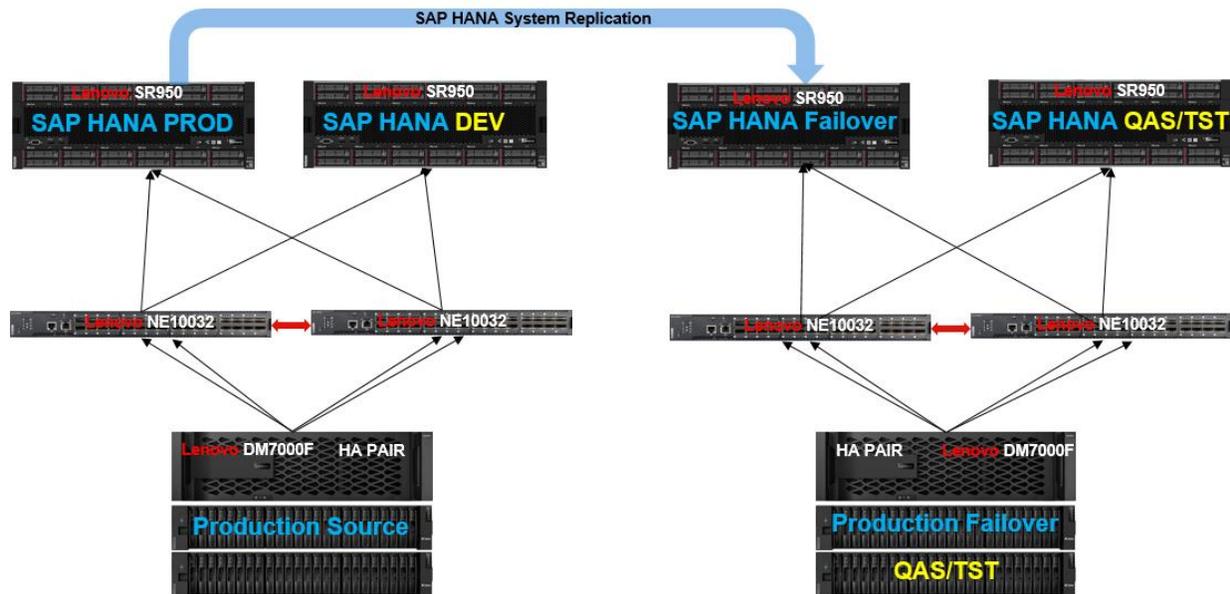


FIGURE 17 : SAP S/4 HANA BUSINESS CONTINUITY USING HSR

Within this architecture, we implement:

- Snapshot -based storage backup for system PRD
- Development on the left
- Quality assurance and test on the right
- There is sufficient headroom available to provide snapshot copies on both sites.

Summary

This document highlights Lenovo enterprise-class servers and enterprise storage components for TDI-based SAP HANA deployments. This allows customers to operate SAP HANA with their operational best practice as an existing sourcing strategy. We outlined the scalability on the server as well as the storage side to address growing demands. It is a simple architecture for virtual and physical SAP HANA deployments covering important development, availability and data protection needs—right out of the box. With many SAP software stacks turning to HANA, the ability to scale up and scale with the correct number of SAP HANA systems is critical.

Resources

Lenovo Documentation

[Lenovo ThinkSystem DM7000F Unified Flash Storage Array](#)

[Lenovo ThinkSystem DM5000F Unified Flash Storage Array](#)

[Lenovo ThinkSystem DE Series Storage Array's](#)

[ThinkSystem SR950 Performance Leadership](#)

[Lenovo ThinkSystem SR850 Server](#)

[Lenovo ThinkSystem SR950 Server](#)

[In-memory Computing with SAP HANA on Lenovo Systems](#)

NetApp Documentation

[SAP HANA Backup and Recovery with SnapCenter](#)

[Automating SAP System Copies Using the SnapCenter SAP HANA Plug-In](#)

[SAP HANA System Replication Backup and Recovery with SnapCenter](#)

[SAP HANA Disaster Recovery with Asynchronous Storage Replication](#)

[Automating SAP System Copies using the SnapCenter 4.0 SAP HANA Plug-In](#)

[Integrating NetApp ONTAP Systems with SAP Landscape Management](#)

SAP Documentation

[SAP HANA Server Installation Guide](#)

[SAP HANA](#)

[SAP HANA High Availability Paper](#)

[SAP HANA HW Directory](#)

SAP Tailored Datacenter Integration [Overview](#) and [FAQ](#)

Legal

© Copyright 2019, Lenovo

This document outlines a general product direction and should not be relied on in making a purchase decision. This document is not subject to your license agreement or any other agreement with Lenovo, NetApp, or SAP. Lenovo, NetApp, or SAP has no obligation to pursue any course of business outlined in this document or to develop or release any functionality mentioned in this document. This document and the strategy and possible future developments by Lenovo, NetApp, and SAP are subject to change and may be changed by the respective companies at any time, for any reason, without notice. Lenovo, NetApp, and SAP assume no responsibility for errors or omissions in this document, except if such damages were caused by Lenovo, NetApp, or SAP intentionally or through gross negligence.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of Lenovo, NetApp, and SAP.

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. This information could include technical inaccuracies or typographical errors.

Any performance data contained herein was determined in a controlled environment; therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Lenovo, the Lenovo logo, and ThinkSystem are trademarks of Lenovo in the United States, other countries, or both.

Intel and Xeon are trademarks of Intel Corporation in the United States, other countries, or both.

NetApp, the NetApp logo, ONTAP, FlexClone, Snapshot, Snap Creator, Snap Mirror, and Snap Vault are trademarks of NetApp in the United States, other countries, or both. SAP, the SAP logo, and HANA are trademarks of SAP in the United States, other countries, or both.

Other trademarks used herein are the properties of their respective owners.