

SAP HANA TDI on Dell EMC Isilon All-Flash Scale-Out NAS

Best Practices for Isilon F800 Storage Configuration

December 2018

H17525

Validation Guide

Abstract

This validation guide describes storage configuration best practices for SAP HANA in Tailored Data Center (TDI) deployments on Isilon F800 all-flash scale-out NAS.

Dell EMC Solutions

Copyright

The information in this publication is provided as is. Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners. Published in the USA 12/18 Validation Guide H17525.

Dell Inc. believes the information in this document is accurate as of its publication date. The information is subject to change without notice.



Contents

- Executive summary.....4
- Isilon F800 storage overview.....7
- Design principles and recommendations for SAP HANA on Isilon F800 NAS.....11
- Configuring and installing an SAP HANA scale-out cluster on Isilon F800 NAS.....19
- Implementing STONITH with the HA/DR provider for SAP HANA.....32
- Conclusion.....41
- References.....42

Executive summary

Introduction

SAP HANA is an in-memory data platform that can be deployed on-premises (locally) or in the cloud. Organizations use the SAP HANA platform to analyze large volumes of data and develop and deploy applications in real time. The SAP HANA database is at the core of this real-time data platform.

SAP HANA combines SAP software components that are optimized on proven hardware that is provided by SAP partners. Two different models are available for on-premises deployment, as shown in the following figure.

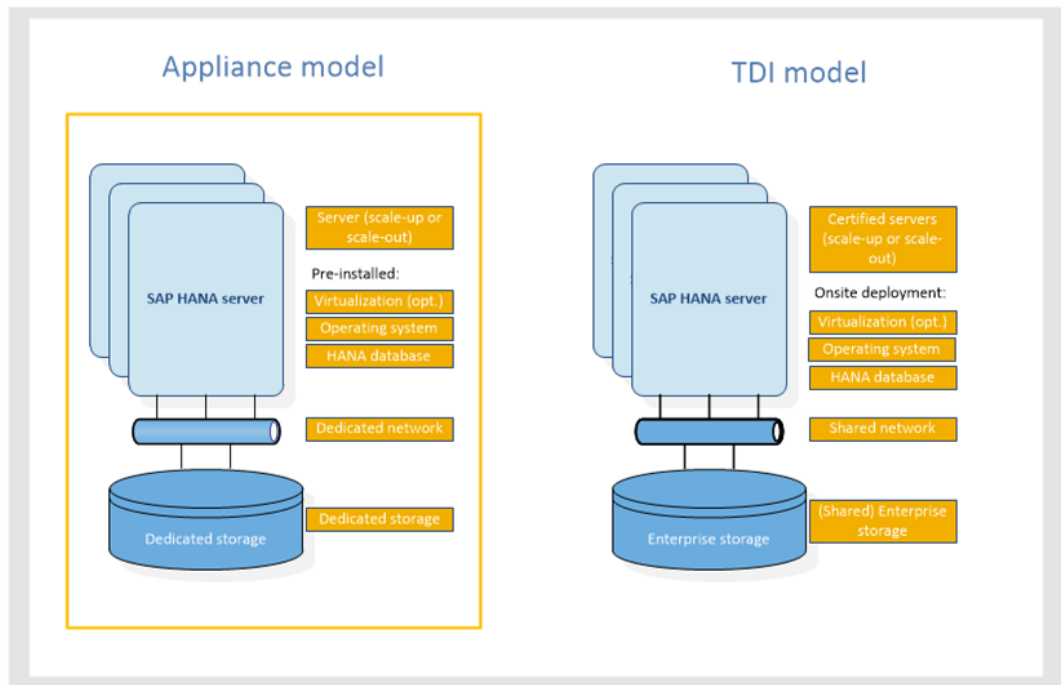


Figure 1. SAP HANA appliance model and TDI model comparison

Appliance model

By default, an SAP HANA appliance includes integrated storage, compute, and network components. The appliance is certified by SAP, built by one of the SAP HANA hardware partners, and shipped to customers with all of its software components preinstalled, including the operating systems and SAP HANA software.

The SAP HANA appliance model presents customers with the following limitations:

- Limited choice of servers, networks, and storage
- Inability to use existing data center infrastructure and operational processes
- Little knowledge and control of the critical components in the SAP HANA appliance
- Inability to respond rapidly to unexpected growth demands without incurring higher infrastructure upgrade costs because of fixed memory sizes for SAP HANA server and storage capacities.

TDI model

In a TDI model, the SAP HANA servers must be certified by SAP and meet the SAP HANA requirements, but the network and storage components, including storage arrays, can be shared in customer environments. Organizations can integrate SAP HANA seamlessly into existing data center operations such as disaster recovery, data protection, monitoring, and management. This integration reduces the time-to-value, costs, and risk of an overall SAP HANA adoption.

Solution overview

SAP has certified Dell EMC Isilon F800 all-flash scale-out NAS as meeting all performance and functional requirements for SAP HANA TDI deployments. This certification means that customers can use Isilon F800 NAS storage for SAP HANA TDI deployments in a fully supported environment with their existing data center infrastructures.

Note: See the [SAP-Certified Dell Solutions website](#) for a complete list of the Dell EMC servers that SAP has certified.

We used the SAP HANA Hardware Configuration Check Tool (hwcct) to perform extensive testing on the Isilon F800 storage system in accordance with the SAP HANA-HWC-ES-1.1 certification scenario. From the test results, we derived storage configuration recommendations for the Isilon F800 system that meet SAP performance requirements and ensure the highest availability for database persistence on disk.

Note: SAP recommends that customers run the hwcct tool in their environment to ensure that their SAP HANA TDI implementation meets the SAP performance criteria.

Key benefits

The TDI solution increases server and network vendor flexibility while reducing hardware and operational costs. Organizations using SAP HANA TDI on Isilon F800 all-flash scale-out NAS storage can:

- Integrate SAP HANA into an existing data center.
- Avoid the significant risks and costs that are associated with operational change by using their existing operational processes, skills, and tools.
- Transition easily from an appliance-based model to the Isilon F800 all-flash NAS-based TDI architecture while relying on Dell EMC Professional Services to minimize risk.
- Use the performance and scale benefits of Isilon F800 all-flash NAS storage to obtain real-time insights across the business.
- Take advantage of all-flash technology for the SAP HANA persistence and benefit from reduced SAP HANA startup, host auto-failover, and backup times.
- Deploy any SAP HANA scenario with Isilon F800 scale-out storage in single-host or multiple-host configurations.

Document purpose

This validation guide describes a solution that uses the SAP HANA platform in a TDI deployment scenario on Isilon all-flash enterprise NAS. The guide provides configuration recommendations that are based on SAP requirements for high availability and key performance indicators (KPIs) for data throughput and latency for the TDI deployment.

Topics that the guide addresses include:

- Introduction to the key technologies in the solution
- Design principles and configuration recommendations for Isilon F800 all-flash scale-out NAS storage with SAP HANA
- Example of an SAP HANA scale-out installation using Isilon NAS storage

Note: This validation guide describes SAP HANA TDI deployments in physical environments. If you plan to use SAP HANA in VMware virtualized environments on vSphere, see the [VMware Virtualized SAP HANA with Dell EMC Storage Deployment Guide](#)

Audience

This document is for system integrators, system or storage administrators, customers, partners, and members of Dell EMC Professional Services who want to configure an Isilon all-flash scale-out NAS storage system to use in an SAP HANA TDI environment.

We value your feedback

Dell EMC and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the [Dell EMC Solutions team](#) with your comments.

Authors: Werner Katzenberger, Donagh Keeshan, Aighne Kearney

The following page of the Everything SAP at Dell EMC space on the Dell EMC Communities website provides links to additional documentation for this solution: [SAP Info Hub for Ready Solutions](#)

Isilon F800 storage overview

Dell EMC Isilon storage systems offer storage amounts on a “grow-as-you-go” basis, eliminating the need for overprovisioning. Isilon hardware platforms are built on the innovative Isilon scale-out architecture, which is designed to provide faster access to massive amounts of data while dramatically reducing cost and complexity.

The new generation of Isilon hardware systems enables you to consolidate and support a wide range of file workloads. With a highly dense design that contains four nodes within a single 4U chassis, the Isilon F800 storage system delivers exceptional performance and efficiency for your most demanding data applications and workloads.

System specifications

The following table shows the system specifications for the Isilon F800 storage system.

Table 1. Isilon F800 all-flash chassis and cluster specifications

Attribute/option	1.6 TB SSD	3.2 TB SSD	15.4 TB SSD
Chassis capacity	96 TB	192 TB	924 TB
SSDs (2.5 in.) per chassis	60	60	60
Operating system	Isilon OneFS version 8.1.2.0 or later		
Number of nodes per chassis	4	4	4
CPU type per node	Intel Xeon Processor E5-2697A v4		
ECC memory per node	256 GB		
Front-end networking per node	2 x 10 GbE (SFP+ or twin-ax copper) or 2 x 40 GbE (QSFP+)		
Infrastructure networking per node	2 InfiniBand connections supporting QDR links or 2 x 40GbE (QSFP+)		
Cluster attributes	Isilon F800 all-flash		
Number of chassis	1 to 36		
Number of nodes	4 to 144		
Cluster capacity	96 TB to 33 PB		
Rack units	4 to 144		

Performance

The Isilon F800 all-flash storage system delivers up to 250,000 IOPS and 15 GB/s aggregate throughput in a single chassis configuration. You can scale performance up to 9M IOPS and 540 GB/s of aggregate throughput in a single Isilon cluster. For front-end networking, you have a choice of dual 10 GbE or dual 40 GbE ports. For back-end networking, you can choose from high-performance InfiniBand QDR or 40 GbE networking.

Capacity

The Isilon F800 all-flash storage system enables you to scale raw storage capacity from 96 TB to 924 TB in a single 4U chassis, and up to 33 PB in a single Isilon cluster. Each

Isilon F800 chassis houses 60 solid state drives (SSDs) with a capacity choice of 1.6 TB, 3.2 TB, or 15.4 TB per drive.

Efficiency

Isilon scale-out storage delivers up to 80 percent storage utilization compared with approximately 50 percent for traditional platforms. Isilon SmartDedupe data deduplication software enhances storage efficiency to reduce your physical storage requirements. Isilon's policy-based, automated tiering options, including Isilon SmartPools and CloudPools software, enable you to optimize storage resources and further lower costs.

Simplicity

You can install an Isilon F800 all-flash cluster, and configure it online, in a short amount of time. To meet your organization's requirements, you can add another Isilon F800 chassis to scale performance and capacity.

Flexibility

Powered by the OneFS operating system, Isilon F800 all-flash technology supports all major protocols and data access methods, including NFS, SMB, HDFS, HTTP, FTP, and OpenStack Swift. Therefore, you can support a wide range of unstructured data applications and workloads on a single platform.

Data protection

Isilon F800 all-flash storage is highly resilient and offers N+1 through N+4 redundancy. With Isilon storage systems, you can also choose from a variety of efficient and proven enterprise data backup and disaster recovery options.

Security

Isilon F800 all-flash technology offers a broad range of robust security options, including FIPS 140-2 level 2 self-encrypting drives, role-based access control (RBAC), secure access zones, SEC 17a-4 compliant WORM data immutability, and integrated file system auditing support.

Isilon OneFS

The Isilon OneFS operating system provides the intelligence behind all Isilon scale-out storage solutions. As part of a storage solution designed for exceptional ease of use, the OneFS operating system can help you accelerate processes and workflows, while providing the highest levels of data protection and scaling easily to handle massive growth.

The following figure shows the features of the OneFS operating system.

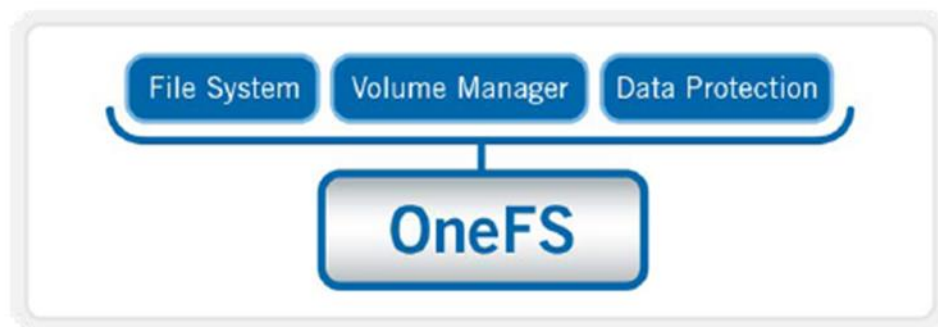


Figure 2. Isilon OneFS operating system

The Isilon OneFS operating system works with the Isilon scale-out NAS nodes, which are referred to as a “cluster.” A single Isilon cluster consists of multiple nodes, storage media, and rack-mountable enterprise appliances containing memory, CPU, networking, Ethernet or low-latency InfiniBand interconnects, and disk controllers. Each node in the distributed cluster has compute as well as storage or capacity capabilities.

With the Isilon F800 storage system, a single chassis of four nodes in a 4U form factor is required to create a cluster, which currently scales up to 144 nodes. Each node or chassis that is added to a cluster increases aggregate disk, cache, CPU, and network capacity. OneFS uses each of the hardware building blocks so that the whole becomes greater than the sum of the parts. The RAM is grouped together into a single coherent cache, enabling I/O on any part of the cluster to benefit from data cached anywhere. A file system journal ensures that writes are safe across power failures. Spindles and CPU are combined to increase throughput, capacity, and IOPS as the cluster grows for access to one file or to multiple files.

Isilon network

Two types of networks are associated with a cluster: internal or back-end and external or front-end.

Back-end network

All intra-node communication in a cluster is performed across a dedicated back-end network, comprising either 40 GbE or low-latency QDR InfiniBand. This back-end network, which is configured with redundant switches for high availability, acts as the backplane for the cluster, enabling each node to act as a contributor in the cluster while also isolating node-to-node communication to a private, high-speed, low-latency network. The back-end network uses IP for node-to-node communication.

Front-end network

Clients connect to the cluster by using the 10 GbE or 40 GbE Ethernet connections that are available on all nodes. Because each node provides its own Ethernet ports, the amount of network bandwidth available to the cluster scales linearly with performance and capacity. The Isilon cluster supports standard network communication protocols to a customer network, including NFS, SMB, HTTP, FTP, HDFS, and OpenStack Swift. Furthermore, the OneFS operating system provides full integration with both IPv4 and IPv6 environments.

Note: 1 GbE client connections are not supported in SAP HANA deployments.

Complete cluster view

The following figure depicts the complete cluster architecture: software, hardware, and network working together along with servers to provide a distributed single file system that can scale dynamically as throughput needs or workloads and capacity needs change in a scale-out environment.

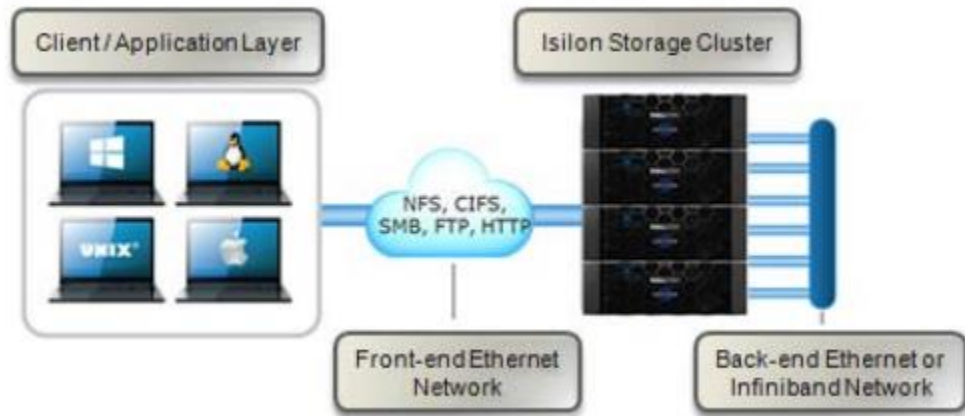


Figure 3. Complete cluster view

Design principles and recommendations for SAP HANA on Isilon F800 NAS

Overview

SAP HANA production systems in TDI environments must meet the SAP performance KPIs. This section describes system requirements, general considerations, and best-practice recommendations for connecting SAP HANA to Isilon scale-out NAS storage systems.

SAP HANA capacity requirements

Every SAP HANA node requires storage devices and capacity for:

- Operating system boot image
- SAP HANA installation
- SAP HANA persistence (data and log)
- Backup

In general, every directory that is created on the Isilon system has access to the total disk capacity of the Isilon cluster. You can apply SmartQuotas by using the Isilon quota-management tool if disk capacity for specific directories must be limited. For more information, see [SAP HANA installation](#).

Operating system boot image

The operating system (either SUSE Enterprise Linux for SAP Applications or Red Hat Enterprise Linux for SAP HANA), must be installed on local disks on the SAP HANA certified servers.

SAP HANA installation

Every SAP HANA node requires access to a file system mounted under the local `/hana/shared/` mount point for installation of the SAP HANA binaries and configuration files, traces, and logs. An SAP HANA scale-out cluster requires a single shared file system, which must be mounted on every node. You can create the installation directory in the Isilon system.

To calculate the size of the `/hana/shared/` file system, use the formulas that are provided in the SAP paper [SAP HANA Storage Requirements](#). Version 2.10 of the paper provides the following formulas:

Single node (scale-up):

$$\text{Size}_{\text{installation}}(\text{single-node}) = \text{MIN}(1 \times \text{RAM}; 1 \text{ TB})$$

Multinode (scale-out):

$$\text{Size}_{\text{installation}}(\text{scale-out}) = 1 \times \text{RAM}_{\text{of_worker}} \text{ per } 4 \text{ worker nodes}$$

You can limit the disk size of the `/hana/shared` directory by using SmartQuotas. The following Isilon command limits the capacity of the `/hana/shared` directory `/ifs/home/hana/shared` to 400 GB:

```
isi quota quotas create /ifs/home/hana/shared directory --hard-threshold 400G --enforced true
```

SAP HANA persistence (data and log)

For the SAP HANA data and log directories, Dell EMC recommends enabling access to the overall Isilon cluster capacity for performance reasons and not limiting the disk capacity using SmartQuotas. If for any reason the disk capacity must be limited, use the following information for sizing the SmartQuotas.

The SAP HANA in-memory database requires disk storage for the following purposes:

- Data:
 - To maintain the persistence of the in-memory data on disk to prevent data loss resulting from a power outage
 - To enable a host auto-failover, where a standby SAP HANA host takes over the in-memory data of a failed worker host in scale-out installations
- Log: To log information about data changes (redo log)

Every SAP HANA scale-up node and scale-out (worker) node requires two disk directories on the Isilon storage system—to save the in-memory database on disk (data) and keep a redo log (log). The size of these directories depends on the anticipated total memory requirement of the database and the RAM size of the node. To prepare the disk sizing, see the SAP paper [SAP HANA Storage Requirements](#). Version 2.10 of the paper states that you can calculate the size of the data directory by using the following formula:

$$\text{Size}_{\text{data}} = 1.2 \times \text{net disk space for data}$$

The net disk space is the anticipated total memory requirement of the database plus 20 percent free space.

If the database is distributed across multiple nodes in a scale-out cluster, divide the net disk space by the number of SAP HANA worker nodes in the cluster. For example, if the net disk space is 2 TB and the scale-out cluster consists of four worker nodes, then every node must have a 616 GB data volume assigned to it ($2 \text{ TB} / 4 = 512 \text{ GB} \times 1.2 = 616 \text{ GB}$).

If the net disk space is unknown at the time of storage sizing, Dell EMC recommends using the RAM size of the node plus 20 percent free space to calculate the capacity of the data file system.

The size of the log directory depends on the RAM size of the node. The [SAP HANA Storage Requirements](#) paper provides the following formulas to calculate the minimum size of the log directory:

$$[\text{systems} \leq 512\text{GB}] \text{ Size}_{\text{redolog}} = 1/2 \times \text{RAM}$$

$$[\text{systems} > 512\text{GB}] \text{ Size}_{\text{redolog}(\text{min})} = 512 \text{ GB}$$

Backup

SAP HANA supports backup to a file system or the use of SAP-certified third-party tools. Dell EMC supports data-protection strategies for SAP HANA backup using Dell EMC Data Domain systems and Dell EMC NetWorker software. Although you can back up an SAP HANA database to a directory on the Isilon storage system, Dell EMC does not recommend backing up the SAP HANA database to the storage system on which the

primary persistence resides. If you plan to back up SAP HANA to a directory on a different Isilon storage system, see the information about sizing the backup file system in the [SAP HANA Storage Requirements](#) paper. The capacity depends not only on the data size and the frequency of change operations in the database but also on the number of backup generations that are kept on disk.

Storage network considerations

This section describes storage network considerations relating to NAS connectivity, the network topology, SAP HANA hosts, and switches.

NAS connectivity

The Network File System (NFS) protocol is used when SAP HANA clients mount remote directories on the Isilon storage system. The two major versions of NFS are NFSv3 and NFSv4. NFSv4 is designed as a stateful protocol to provide additional functions such as file locking. NFSv3 is the most popular protocol.

Note: Because Isilon has been certified by using NFSv3, this guide only addresses the NFSv3 protocol.

As a best practice, connect the SAP HANA hosts to the Isilon system by using a redundant private storage network infrastructure.

For more information about configuring and troubleshooting Isilon scale-out NAS external networking, see the [Isilon Network Design Considerations White Paper](#). The paper also provides information about the network architecture design to address availability considerations (no single point of failure) and bandwidth requirements.

For information about implementing NFS with Isilon, see [Isilon OneFS NFS Design Considerations and Best Practices](#).

Network topology

The following figure shows a sample network topology, with eight SAP HANA hosts that are connected over a dedicated private storage network to a four-node Isilon F800 storage system.

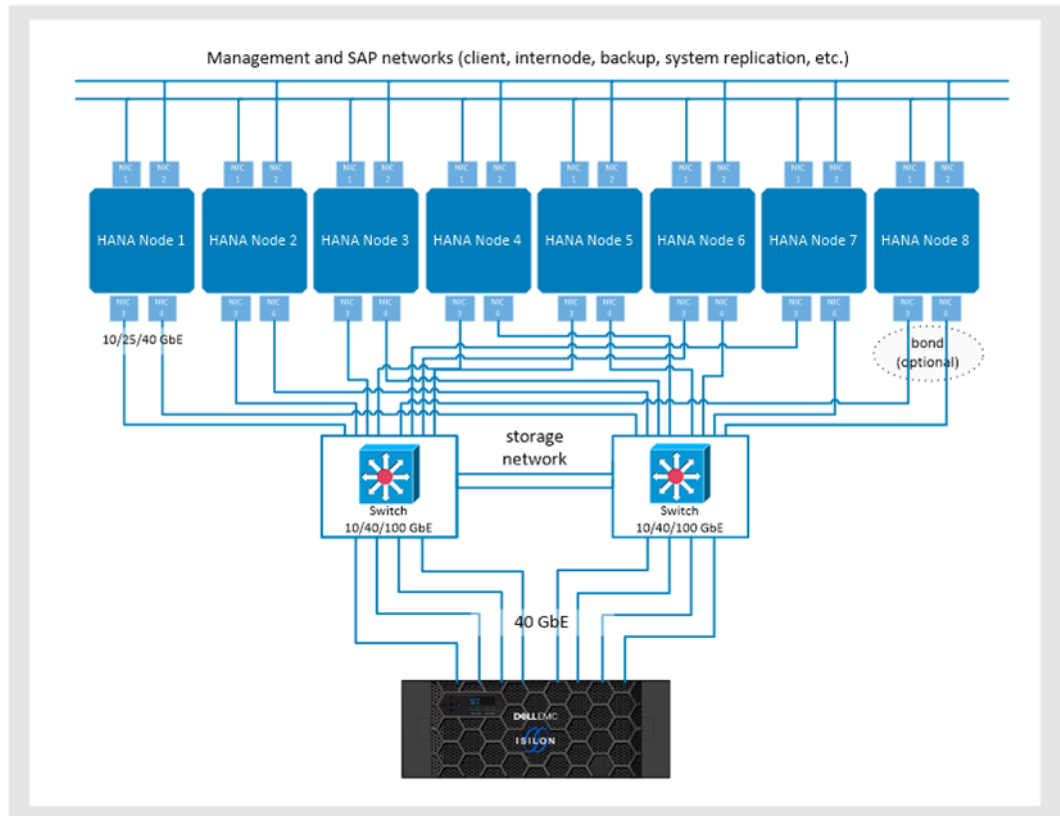


Figure 4. Sample network topology

All the components in the storage network require configurations using MTU size 9000 (jumbo frames).

Storage network connections for SAP HANA hosts

SAP HANA hosts require network interfaces for the management network and the SAP HANA-specific networks for client connections, internode, system replication, and backup. In addition, each SAP HANA host requires a dedicated network connection to the storage network switches with 10 GbE or 40 GbE link speed. The connection should be redundant, that is, it should have at least two NICs. The 25 GbE network interfaces on the SAP HANA hosts require four 25 GbE NICs bonded to a 100 GbE aggregate. Switches that support 100 GbE are also required.

With 10 or 40 GbE on the SAP HANA hosts, you can optionally configure active-active interface groups on the SAP HANA hosts to use network bonding (bonds), where multiple network interfaces are aggregated into a single logical bonded interface. Hosts that are configured with bonds require Multi-Chassis Link Aggregation (MLAG) and Link Aggregation Control Protocol (LACP) on the switches.

The following figure shows a sample network topology, with the two network ports bonded on the SAP HANA host and MLAG with LACP on the switches.

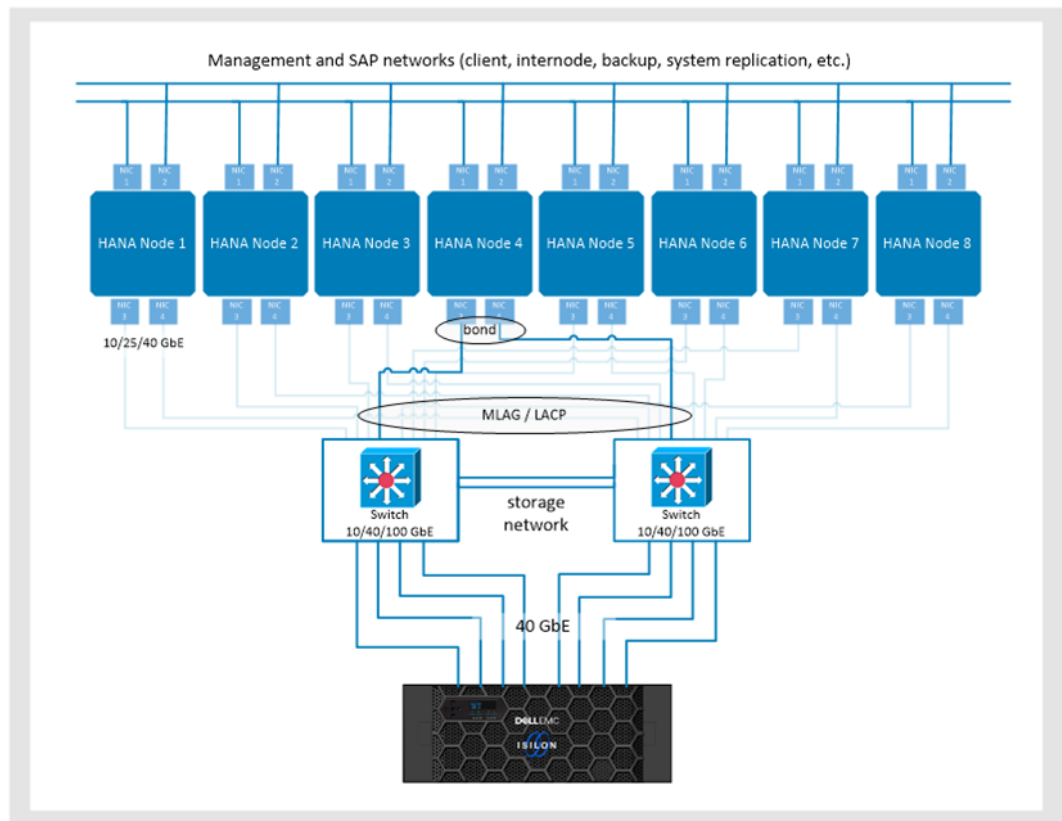


Figure 5. Sample network topology with host bond and MLAG/LACP on the switches

Storage network switches

The switches in the storage network must support 40 GbE connections to the Isilon storage system as well as link speeds of 10 GbE, 25 GbE, or 40 GbE from the SAP HANA hosts.

With 25 GbE host networks that are bonded to 100 GbE aggregates, the switches must support 100 GbE. You must have a redundant dedicated storage network in place that has enough ports to connect all of the Isilon external network interfaces as well as two network connections per SAP HANA host.

If the data center has a hierarchical network design with core, distribution, and access switches, the link from the core to the distribution to the access switches must be capable of processing the throughput from the access switches with the Isilon connections.

Isilon network interfaces

Each Isilon node provides two interfaces for external networking. A four-node Isilon cluster provides eight external network interfaces. For optimal performance, the SAP HANA client connections and the mounts of the SAP HANA disk partitions (data and log) must be evenly distributed across all available Isilon interfaces. You can achieve this distribution by using dedicated Isilon interface IP addresses for the mounts of each SAP HANA disk partition. In SAP HANA scale-out deployments, the mounts on the SAP HANA clients are:

```
<Isilon-IP-address1>:/ifs/home/hana/<SID>/data01 →  
/hana/data/<SID>/mnt00001  
<Isilon-IP-address1>:/ifs/home/hana/<SID>/log01 →  
/hana/log/<SID>/mnt00001  
<Isilon-IP-address2>:/ifs/home/hana/<SID>/data02 →  
/hana/data/<SID>/mnt00002  
<Isilon-IP-address2>:/ifs/home/hana/<SID>/log02 →  
/hana/log/<SID>/mnt00002  
<Isilon-IP-address3>:/ifs/home/hana/<SID>/data03 →  
/hana/data/<SID>/mnt00003  
<Isilon-IP-address3>:/ifs/home/hana/<SID>/log03 →  
/hana/log/<SID>/mnt00003
```

and so on.

With multiple single-host (scale-up) SAP HANA installations, each SAP HANA client must use a dedicated Isilon IP address for the mounts, as follows:

SAP HANA scale-up 1:

```
<Isilon-IP-address1>:/ifs/home/hana/<SID1>/data →  
/hana/data/<SID1>/mnt00001  
<Isilon-IP-address1>:/ifs/home/hana/<SID1>/log →  
/hana/log/<SID1>/mnt00001
```

SAP HANA scale-up 2:

```
<Isilon-IP-address2>:/ifs/home/hana/<SID2>/data →  
/hana/data/<SID2>/mnt00001  
<Isilon-IP-address2>:/ifs/home/hana/<SID2>/log →  
/hana/log/<SID2>/mnt00001
```

SAP HANA scale-up 3:

```
<Isilon-IP-address3>:/ifs/home/hana/<SID3>/data →  
/hana/data/<SID3>/mnt00001  
<Isilon-IP-address3>:/ifs/home/hana/<SID3>/log →  
/hana/log/<SID3>/mnt00001
```

and so on.

SAP HANA I/O patterns

The SAP HANA persistent devices use different I/O patterns. For more information, see [SAP HANA Storage Requirements](#).

Data volume

Access to the data volume is primarily random, with blocks ranging from 4 KB to 64 MB. The data is written asynchronously with parallel I/Os to the data file system. During normal operations, most of the I/Os to the data file system are writes and data is read from the file system only during database restarts, SAP HANA backups, host auto-failover, or a column store table load or reload operation.

Log volume

Access to the log volume is primarily sequential, with blocks ranging from 4 KB to 1 MB. SAP HANA keeps a 1 MB buffer for the redo log in memory. When the buffer is full, it is synchronously written to the log volume. When a database transaction is committed before the log buffer is full, a smaller block is written to the file system. Because data is written synchronously to the log volume, a low latency for the I/O to the storage device is important, especially for the smaller 4 KB and 16 KB block sizes. During normal database operations, most of the I/Os to the log volume are writes and data is read from the log volume only during database restart, high availability failover, log backup, or database recovery.

For information about optimizing SAP HANA for specific storage environments, see [Optimizing file I/O after the SAP HANA installation](#) on page 29.

Isilon F800 scalability for SAP HANA

We performed tests on an Isilon F800 storage system by using the SAP hwcct tool for HANA-HWC-ES 1.1 certification. Based on the test results, the following table provides guidelines for estimating the initial number of SAP HANA production hosts that can be connected to a specific Isilon F800 all-flash scale-out NAS storage system.

Table 2. Isilon F800 scalability

Number of Isilon chassis/nodes	Number of SAP HANA worker hosts
1 chassis, 4 nodes	8
2 chassis, 8 nodes	16

Note: Scalability beyond 16 SAP HANA worker hosts requires a customer-specific validation using the SAP HANA hwcct tool.

Depending on the workload, the number of SAP HANA hosts that can be connected to an Isilon F800 storage system in a customer environment can be higher or lower than the number specified in Table 2. Use the SAP HANA hwcct tool with scenario HANA-HWC-ES 1.1 to validate the SAP HANA performance and determine the maximum number of SAP HANA hosts on a single storage array.

Competing workloads

When you share physical hardware in a Dell EMC Isilon cluster with other workloads, competition exists for the following services:

- CPU
- Memory
- Network bandwidth
- Disk I/O
- Disk capacity

Quality of service

Quality of service (QoS) addresses physical hardware performance characteristics that can be measured, improved, and sometimes guaranteed. Characteristics that are measured for QoS include throughput rates, CPU usage, and disk capacity.

Although access zones do not provide logical QoS guarantees to these resources, you can partition the resources between access zones on a single cluster. The following table describes a number of ways to partition resources to improve QoS.

Table 3. Resource assignment guidelines

Feature	Description
NICs	You can assign specific NICs on specific nodes to an IP address pool that is associated with an access zone. By assigning these NICs, you can determine the nodes and interfaces that are associated with an access zone, enabling the separation of CPU, memory, and network bandwidth.
SmartPools	Node hardware equivalence classes separate SmartPools, usually into multiple tiers of high, medium, and low performance. The data written to a SmartPool is written only to the disks in the nodes of that pool. Associating an IP address pool with only the nodes of a single SmartPool enables partitioning of disk I/O resources.
SmartQuotas	Through SmartQuotas, you can limit disk capacity by user or group or in a directory. By applying a quota to an access zone's base directory, you can limit the disk capacity that is used in that access zone.

In general in SAP HANA environments, Isilon nodes and resources are logically dedicated to SAP HANA performance. With additional workloads, Dell EMC recommends that you add more Isilon nodes to the cluster with logically separated resources.

Configuring and installing an SAP HANA scale-out cluster on Isilon F800 NAS

Overview

This section describes how to configure and install an SAP HANA scale-out cluster on a Dell EMC Isilon F800 NAS system. These procedures entail:

- Creating and configuring the persistent storage (data and log) on an Isilon F800 NAS system for an SAP HANA scale-out cluster with two worker nodes and one standby node (2+1)
- Preparing the SAP HANA hosts
- Installing the SAP HANA cluster by using the SAP database life cycle management command-line tool hdblcm

Configuring the Isilon F800

The OneFS operating system centralizes cluster management through a web administration interface and a command-line interface (CLI). Both interfaces provide methods to:

- Activate licenses
- Check the status of nodes
- Configure the cluster
- Upgrade the system
- Generate alerts
- View client connections
- Track performance
- Change various settings.

OneFS minimum requirements for SAP HANA deployments

For SAP HANA deployments, OneFS 8.1.2 is the minimum required version. The following Roll Up DA patch is also required:

[RUP-DA](Protocols/Userspace) Roll-Up v1.0 (8.1.2.0)

Note: A field support engineer can download and install the DA patches from [this location](#).

To achieve optimal performance for SAP HANA environments with Isilon all-flash scale-out NAS storage systems, you must apply some global settings to the Isilon system. These settings only become adjustable with the Roll-Up v1.0 (8.1.2.0) patch.

After the patch is installed, use the OneFS 8.1.2 CLI to adjust the settings. Log in to the OneFS CLI as root and run the following commands:

```
isi_sysctl_cluster efs.rbm.dwt_threads=4

/usr/likewise/bin/lwregshell set_value
"[HKEY_THIS_MACHINE\\Services\\lwio\\Parameters\\Drivers\\nfs]"
SchedulerFhaMaxConcurrentIO 22

/usr/likewise/bin/lwsm shutdown
```

The `/usr/likewise/bin/lwsm shutdown` command triggers a restart of the service and the new settings become active.

The `isi_sysctl_cluster` command ensures that the settings are applied after a node or cluster reboot.

For more information, see the [Isilon OneFS CLI Administration Guide](#).

Create a directory

The OneFS web administration interface is used to create and configure all storage directories and NFS exports for the SAP HANA scale-out cluster. For more information, see the [OneFS Isilon Web Administration Guide](#).

Follow these steps to create a directory:

1. Log in to the OneFS web administration interface.
2. Select **File System > File System Explorer**.

In File System Explorer, you can manage the content that is stored on your Isilon F800 cluster. You can navigate the Isilon file system (`/ifs`), add directories, and manage file and directory properties including data protection, I/O optimization, and UNIX permissions.

You can also use the File System Explorer to create a directory in the `/ifs` directory tree. In our example, we created the directories `/ifs/home/hana` and `/ifs/home/hana/SID`.

3. Go to the directory to which you want to add the directory and click **Create Directory**.
4. In the **Directory Name** field, enter the name `hana`, assign the relevant permissions in the **Permissions** area, and then click **Create Directory**.

We used “HIB” for our directory name

5. Go to the new `hana` directory and repeat steps 3 to 4 to create the subdirectories “shared” and “SID.”

For our 2+1 SAP HANA cluster, we need to create two data directories and two log directories within the HIB directory for the SAP HANA persistence, as shown in the following figure.




Figure 6. Creating data and log directories

6. For each data and log directory, choose the following directory properties:
 - **Protection Settings**—Disable **Managed by Smartpools** and keep the defaults.
 - **I/O optimization Settings**—Keep **Enable Smartcache**.
 - **Data Access Pattern**—Select **Optimize** for concurrent access.

For optimal performance, Dell EMC recommends turning off Endurant Cache (EC) for the SAP HANA data directories only.

EC is the OneFS caching mechanism for synchronous writes to improve latency. With the SAP HANA persistence, writes to data are asynchronous and writes to log are synchronous. EC is, therefore, left turned on for the log directories.

EC operates in combination with the OneFS write cache (coalescer). By default, EC and the coalescer are turned on. To turn them on or off by using the Isilon OneFS CLI, open a Secure Shell (SSH) connection to any Isilon node in the cluster by using the node's IP address.

7. For each directory, enable or disable the coalescer and EC by running the following command:

```
isi set -c [on|off|endurant_all|coal_only] <directory_name>
```

8. To enable the coalescer but switch off EC on the data directories, go to the location of the data directories and run the following command:

```
isi set -R -c coal_only data1
```

```
isi set -R -c coal_only data2
```

The following figure shows the command for disabling EC along with the optimum settings for the SAP HANA data and log directories.

```

Isilon-92N-1-1# pwd
/ifs/home/hana/HIB
Isilon-92N-1-1# isi get -d *
POLICY      LEVEL PERFORMANCE COAL  FILE
default     4x concurrency on   data1/
default     4x concurrency on   data2/
default     4x concurrency on   log1/
default     4x concurrency on   log2/
default     4x concurrency on   shared/
Isilon-92N-1-1# isi set -R -c coal_only data1
Isilon-92N-1-1# isi set -R -c coal_only data2
Isilon-92N-1-1# isi get -d *
POLICY      LEVEL PERFORMANCE COAL  FILE
default     4x concurrency noEC data1/
default     4x concurrency noEC data2/
default     4x concurrency on   log1/
default     4x concurrency on   log2/
default     4x concurrency on   shared/
Isilon-92N-1-1# █

```

Figure 7. Disabling Endurant Cache for the data directories

For more information about EC and the coalescer, see the [Isilon OneFS CLI Administration Guide](#).

NFS exports

Use the OneFS web administration interface to create NFS export rules that define the mount points (paths) that are available to NFS clients and how the server performs with these clients. Each rule must have at least one path and can include additional paths. You can also specify that all subdirectories of the given path or paths are mountable. Otherwise, only the specified paths are exported and child directories are not mountable. An export rule can specify a particular set of clients, enabling you to restrict access to certain mount points or apply a unique set of options to these clients. If the rule does not specify any clients, then the rule applies to all clients that connect to the server.

Configure NFS file sharing

Enable or disable the NFS service and set the lock protection level and security type. These settings are applied across all node in the cluster. You can also change the settings for individual NFS exports that you define. Because Isilon storage is only certified for use with NFSv3 for SAP HANA deployments, you must enable the NFS export service and NFSv3:

1. Log in to the OneFS web administration interface and select **Protocols > UNIX Sharing (NFS) > Global Settings**.
2. Enable **NFS Export Service** and **NFSv3**, as shown in the following figure.

Figure 8. Enabling NFSv3 global settings

Create NFS exports

To share the SAP HANA directories with the SAP HANA clients, create NFS exports. In our example, we created the following two NFS exports:

- SID directory with the HIB path `/ifs/home/hana/HIB` containing the data and log directories. We enabled mount access to all the subdirectories.
- The `hana/shared` directory with the path `/ifs/home/hana/shared`.

To create NFS exports:

1. Select **Protocols > UNIX Sharing (NFS) > NFS Exports**.
2. Click **Create Export**.
3. For the **Directory Paths** setting, browse to or enter the name of the directory that you want to export: `/ifs/home/hana/HIB`
4. (Optional) In the **Description** field, enter a comment that describes the export.
5. Specify the NFS clients that are allowed to access the export.
Provide root access to allow the SAP HANA clients to mount the exports. The client can be identified by host name, IPv4 or IPv6 address, subnet, or netgroup.
6. Select an export permissions setting.
7. Choose **Enable mount access to subdirectories**. Allow subdirectories below the path or paths to be mounted.
8. Specify user and group mappings.
Dell EMC recommends keeping the defaults.
9. For the **Security Flavors** setting, keep the default setting, **UNIX**.
10. Click **Save Changes**.
The new NFS export is created.

Preparing the SAP HANA nodes

This procedure entails:

- Configuring the `/etc/fstab` file on each SAP HANA client and mounting all the SAP HANA data, log, and shared directories from the Isilon storage to the SAP HANA nodes
- Installing an SAP HANA scale-out instance with the SAP HANA `hdbclm` tool by using the NAS storage directories that you previously created.

Prerequisites

The configuration example in this guide assumes that the following basic installation and configuration operations are complete on the SAP HANA nodes:

- The operating system is installed and properly configured according to the SAP recommendations. Our example uses SUSE Linux 12 SP3 for SAP applications.
- The directories to be used for mount points have been created for `/hana/shared` and also for the SAP HANA data directories `/hana/data/SID/mnt0000x` and log directories `/hana/log/SID/mnt0000x` on each of the SAP HANA clients.
- All network settings and bandwidth requirements for internode communications are configured according to the SAP requirements.
- SSH keys have been exchanged between all SAP HANA nodes.
- System time synchronization has been configured using an NTP server.
- The SAP HANA installation DVD ISO file is downloaded from the SAP website and made available on a shared file system.

Note: SAP HANA can be installed only on certified server hardware. A certified SAP HANA expert must perform the installation.

Prepare the `/etc/hosts` files

To achieve optimal performance, ensure that the SAP HANA mounts are evenly distributed across all available Isilon interface IPs in the `/etc/hosts` file. Because each Isilon node in a cluster provides two interfaces, each SAP HANA host's storage partition (data and log directories) must mount with a different interface IP balance across the Isilon nodes. For more information, see [Isilon network interfaces](#) on page 15.

In the 2+1 SAP HANA system with SID HIB example, the mount points are included in the `/etc/fstab` file on each of the SAP HANA clients, including any standby hosts, as shown in the following figure.

```
#HANA SID HIB
192.168.40.150:/ifs/home/hana/shared /hana/shared nfs rw,vers=3,hard,timeo=600,rs
size=1048576,ws
ize=1048576,intr,noatime,nolock 0 0
192.168.40.150:/ifs/home/hana/HIB/data01 /hana/data/HIB/mnt00001 nfs rw,vers=3,
hard,timeo=600,rs
ize=1048576,ws
ize=1048576,intr,noatime,nolock 0 0
192.168.40.151:/ifs/home/hana/HIB/data02 /hana/data/HIB/mnt00002 nfs rw,vers=3,
hard,timeo=600,rs
ize=1048576,ws
ize=1048576,intr,noatime,nolock 0 0
192.168.40.150:/ifs/home/hana/HIB/log01 /hana/log/HIB/mnt00001 nfs rw,vers=3,ha
rd,timeo=600,rs
ize=1048576,ws
ize=1048576,intr,noatime,nolock 0 0
192.168.40.151:/ifs/home/hana/HIB/log02 /hana/log/HIB/mnt00002 nfs rw,vers=3,ha
rd,timeo=600,rs
ize=1048576,ws
ize=1048576,intr,noatime,nolock 0 0
```

Figure 9. Example of the `/etc/fstab` file being used to mount SAP HANA directories

The operating system command `mount -a` mounts all the directories from the Isilon system and must be run on each SAP HANA client.

After all the file systems are mounted, you are ready to install the SAP HANA scale-out cluster.

Installing an SAP HANA scale-out cluster

Our example uses the `hdblcm` command-line tool to install the SAP HANA 2+1 scale-out cluster. For more information about this tool, see the [SAP HANA Studio Installation and Upgrade Guide](#).

After the SAP HANA installation DVD ISO file has been extracted on a shared software repository file system that is mounted on all hosts, run the following command from the extracted installation folder to begin the installation:

```
//software/SAP_HANA2_SPS03_r33_IM/51053381/DATA_UNITS/HDB_LCM_LINUX_X86_64/hdblcm
```

```
SAP HANA Lifecycle Management - SAP HANA Database
2.00.033.00.1535711040
```

```
*****
```

```
Scanning software locations...
```

```
Detected components:
```

```
    SAP HANA Database (2.00.033.00.1535711040) in
    /SAPShare/software/SAP_HANA2_SP3/SAP_HANA1_SPS03_r33_IM/51053381/D
    ATA_UNITS/HDB_SERVER_LINUX_X86_64/server
```

```
    SAP HANA Database Client (2.3.119.1535661774) in
    /SAPShare/software/SAP_HANA2_SP3/SAP_HANA1_SPS03_r33_IM/51053381/D
    ATA_UNITS/HDB_CLIENT_LINUX_X86_64/client
```

```
Choose an action
```

Index	Action	Description
1	install	Install new system
2	extract_components	Extract components
3	Exit (do nothing)	

```
Enter selected action index [3]: 1
```

```
SAP HANA Database version '2.00.033.00.1535711040' will be
installed.
```

```
Select additional components for installation:
```

Index	Components	Description

Configuring and installing an SAP HANA scale-out cluster on Isilon F800 NAS

```
1      | all          | All components
2      | server       | No additional components
3      | client       | Install SAP HANA Database Client version
2.3.119.1535661774
4      | studio      | Install SAP HANA Studio version
2.3.41.000000
5      | smartda     | Install SAP HANA Smart Data Access version
2.00.3.000.0
6      | xs          | Install SAP HANA XS Advanced Runtime
version 1.0.88.12598
7      | afl         | Install SAP HANA AFL (incl.PAL,BFL,OFL)
version 2.00.033.0000.1535724035
8      | eml         | Install SAP HANA EML AFL version
2.00.033.0000.1535724035
9      | epmmds     | Install SAP HANA EPM-MDS version
2.00.033.0000.1535724035
```

```
Enter comma-separated list of the selected indices [3]:
Enter Installation Path [/hana/shared]:
Enter Local Host Name [r640-isi01]:
Do you want to add hosts to the system? (y/n) [n]: y
Enter comma-separated host names to add: r640-isi02,r640-isi03
Enter Root User Name [root]:
Collecting information from host 'r640-isi02'...
Collecting information from host 'r640-isi03'...
Information collected from host 'r640-isi03'.
Information collected from host 'r640-isi02'.
```

```
Select roles for host 'r640-isi02':
```

Index	Host Role	Description
1	worker	Database Worker
2	standby	Database Standby

```

...
Enter comma-separated list of selected indices [1]: 1
Enter Host Failover Group for host 'r640-isi02' [default]:
Enter Storage Partition Number for host 'r640-isi02' [<<assign
automatically>>]:
Enter Worker Group for host 'r640-isi02' [default]:
Select roles for host 'r640-isi03':

  Index | Host Role          | Description
  -----
  1      | worker            | Database Worker
  2      | standby           | Database Standby
...
Enter comma-separated list of selected indices [1]: 2
Enter Host Failover Group for host 'r640-isi03' [default]:
Enter Worker Group for host 'r640-isi03' [default]:
Enter SAP HANA System ID: HIB
Enter Instance Number [00]:
Enter Local Host Worker Group [default]:

  Index | System Usage | Description
  -----
  1      | production   | System is used in a production
environment
  2      | test         | System is used for testing, not
production
  3      | development  | System is used for development, not
production
  4      | custom       | System usage is neither production, test
nor development
Select System Usage / Enter Index [4]:
Enter Location of Data Volumes [/hana/data/HIB]:
Enter Location of Log Volumes [/hana/log/HIB]:
Restrict maximum memory allocation? [n]:

```

Configuring and installing an SAP HANA scale-out cluster on Isilon F800 NAS

```
Enter Certificate Host Name For Host 'r640-isi01' [r640-isi01]:
Enter Certificate Host Name For Host 'r640-isi02' [r640-isi02]:
Enter Certificate Host Name For Host 'r640-isi03' [r640-isi03]:
Enter System Administrator (hibadm) Password:
Confirm System Administrator (hibadm) Password:
Enter System Administrator Home Directory [/usr/sap/HIB/home]:
Enter System Administrator Login Shell [/bin/sh]:
Enter System Administrator User ID [1001]:
Enter System Database User (SYSTEM) Password:
Confirm System Database User (SYSTEM) Password:
Restart system after machine reboot? [n]:
.....
Do you want to continue? (y/n): y
Installing components...
Installing SAP HANA Database...
    Preparing package ....
    Creating System...
    Extracting software...
    Installing package
    Starting SAP HANA Database system...
        All server processes started on host 'r640-isi01' (worker).
    Importing delivery units...
    Adding 2 additional hosts in parallel
    Adding host 'r640-isi03'...
    Adding host 'r640-isi02'...
        r640-isi02: Adding host 'r640-isi02' to instance '00'...
        r640-isi03: Adding host 'r640-isi03' to instance '00'...
        r640-isi02: Starting SAP HANA Database...
        r640-isi03: Starting SAP HANA Database...
        r640-isi03: All server processes started on host 'r640-
isi03' (standby).
```

```

r640-isi02:    All server processes started on host 'r640-
isi02' (worker).

Installing Resident hdblcm...

Installing SAP HANA Database Client...

...

Registering SAP HANA Database Components on Local Host...

  Regenerating SSL certificates...

  Deploying SAP Host Agent configurations...

Updating SAP HANA Database Instance Integration on Remote Hosts...

  Updating SAP HANA Database instance integration on host 'r640-
isi02'...

  Updating SAP HANA Database instance integration on host 'r640-
isi03'...

Creating Component List...

```

SAP HANA Database System installed

```

Log file written to '/var/tmp/hdb_HIB_hdblcm_install_2018-11-
06_21.03.03/hdblcm.log' on host 'r640-isi01'.

```

Optimizing file I/O after the SAP HANA installation

The base layer of SAP HANA provides two file I/O interfaces:

- **Simple File**—Used for small, simple I/O requests on configuration files, traces, and so on. The simple file interface uses lightweight, platform-independent wrappers around system calls.
- **FileFactory and File**—Used for large, complex streams of I/O requests on the data and log volumes and for backup and recovery. This interface uses synchronous and asynchronous I/O operations.

Configure the SAP HANA file I/O layer with configuration parameters to optimize file I/Os for a specific storage array and file system. The NFSv3 file system is used with Dell EMC Isilon F800 storage systems for the SAP HANA persistence.

After the SAP HANA persistence is installed on the Isilon directories, set the following file I/O layer parameters for optimal I/O processing:

- `max_parallel_io_requests=256`
- `num_submit_queues=8`
- `num_completion_queues=4`

SAP HANA 1.0

After the initial SAP HANA installation is complete, set the parameters by running the SAP HANA `hdbparam` command as `<sid>adm` in the Linux shell.

```
su - <sid>adm
```

```
hdbparam -p          # lists current parameter setting
hdbparam --paramset fileio [DATA].max_parallel_io_requests=256
hdbparam --paramset fileio [LOG].max_parallel_io_requests=256
hdbparam --paramset fileio [DATA].async_read_submit=on
hdbparam --paramset fileio [LOG].async_read_submit=on
hdbparam --paramset fileio [DATA].async_write_submit_blocks=all
hdbparam --paramset fileio [LOG].async_write_submit_blocks=all
hdbparam --paramset fileio [DATA].num_submit_queues=8
hdbparam --paramset fileio [LOG].num_submit_queues=8
hdbparam --paramset fileio [DATA].num_completion_queues=4
hdbparam --paramset fileio [LOG].num_completion_queues=4
```

SAP HANA 2.0

Starting with SAP HANA 2.0, the `hdbparam` command-line tool was deprecated. Instead, the parameters are defined in `global.ini > [fileio]`.

Using the SQL console in SAP HANA Studio, use the following commands to set the parameters `max_parallel_io_requests=256`, `num_submit_queues=8` and `num_completion_queues=4` in the `global.ini` file:

```
ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'max_parallel_io_requests[DATA]') = '256';

ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'max_parallel_io_requests[LOG]') = '256'

ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'num_submit_queues[DATA]') = '8';

ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'num_submit_queues[LOG]') = '8';

ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'num_completion_queues[DATA]') = '4';

ALTER SYSTEM ALTER CONFIGURATION ('global.ini', 'SYSTEM') SET
('fileio', 'num_completion_queues[LOG]') = '4' WITH RECONFIGURE
```

The following figure shows the `fileio` section of `global.ini` after the parameters are set:

fileio		
async_read_submit	on	
async_write_submit_active	on	
async_write_submit_blocks	all	
max_parallel_io_requests	64	
max_parallel_io_requests[data]		● 256
max_parallel_io_requests[log]		● 256
max_submit_batch_size	64	
min_submit_batch_size	16	
num_completion_queues	1	
num_completion_queues[data]		● 4
num_completion_queues[log]		● 4
num_submit_queues	1	
num_submit_queues[data]		● 8
num_submit_queues[log]		● 8
size_kernel_io_queue	512	

Figure 10. Fileio section of the global.ini file in SAP HANA Studio

Both `async_read_submit=on` and `async_write_submit_blocks=all` are set by default during installation. For more information, see SAP Note [2399079—Elimination of hdbparam in SAP HANA 2](#) (access requires an SAP username and password).

Note: The instructions that are provided in this guide for tuning file I/O parameters are based on SAP HANA 1.0 and SAP HANA 2.0 SPS03. See the latest SAP HANA documentation for any updates.

Implementing STONITH with the HA/DR provider for SAP HANA

Overview

This section applies to multihost SAP HANA scale-out instances and the host auto-failover. On failover, the database on the standby host must have read and write access to the files of the failed active host. To avoid the risk of corrupting these files, the failed host must not be able to write to them. This concept is known as fencing.

When using shared file systems such as Isilon NAS storage and NFSv3, you can implement the “shoot the other node in the head” (STONITH) method to achieve proper fencing capabilities and ensure that locks are always freed.

In such a setup, use the Storage Connector API to invoke the STONITH calls. During failover, the SAP HANA master host calls the STONITH method of the custom storage connector with the hostname of the failed host as an input value.

A mapping of hostnames to management network addresses is maintained, which is used to send a reboot signal to the server over the management network. When the host comes up again, it automatically starts in standby host role.

The STONITH example in this guide uses the Intelligent Platform Management Interface (IPMI) protocol in bare-metal deployments with Dell EMC PowerEdge servers.

Enable IPMI over LAN

For Dell EMC PowerEdge servers, you must configure IPMI over LAN for iDRAC to enable IPMI commands over LAN channels to any external systems. Unless IPMI over LAN is configured, external systems cannot communicate with the iDRAC server using IPMI commands.

To configure IPMI over LAN:

1. In the **iDRAC Settings Utility**, select **Connectivity**.
The **iDRAC Settings Network** page appears.
2. In **IPMI Settings**, choose **Enabled** from the **Enable IPMI Over Lan** drop-down list.
3. Click **Apply**.
4. Repeat the preceding steps for each host used in the SAP HANA scale-out instance.

The `/etc/hosts` file using a standard naming convention, as shown in the following figure, maintains a mapping of hostnames to IPMI IP addresses to be used in STONITH.

```
# IPMI IP mapping to hostnames
10.110.79.61    r640-isi01-ipmi
10.110.79.62    r640-isi02-ipmi
10.110.79.63    r640-isi03-ipmi
```

Figure 11. IPMI IP mapping to hostnames in the `/etc/hosts` file

5. Verify that the IPMI tool is working on each SAP HANA host by running the following command as root:

```
ipmitool power status -H r640-isi01-ipmi -U root -P xxxx
```

If IPMI is working successfully, `Chassis Power is on` is returned.

6. Set the `set-user-ID` bit for `ipmitool` to enable `sidadm` execution permissions:

```
chmod u+s /usr/bin/ipmitool
```

Create a custom HA/DR STONITH provider

To create your own HA/DR provider, complete the following steps and then add the hook method you want to use. Our example uses STONITH.

For information about the HA/DR provider, see the [SAP HANA Administration Guide](#).

1. As `sidadm`, create a directory for the HA/DR provider.

The directory should be within the `/hana/shared` storage of the SAP HANA installation but outside the `<SID>` directory structure. Our example uses the `/hana/shared/HANA_Hooks` location.

2. Copy the `exe/python_support/hdb_ha_dr/HADRDummy.py` from an installed SAP HANA system to the new location.

For example, you can copy the file and rename it to:

```
/hana/shared/HANA_Hooks/HA_STONITH_Hook.py
```

3. Customize the new file by renaming the Python class to the name of the file `HA_STONITH_Hook`.

Within the `HA_STONITH_Hook.py` file, we also customized `def __init__()`, `def about` and the STONITH hook `def stonith`, as shown in the following code sample:

```
"""
```

```
Sample for a HA/DR hook provider.
```

When using your own code in here, please copy this file to location on `/hana/shared` outside the HANA installation.

This file will be overwritten with each `hdbupd` call! To configure your own changed version of this file, please add to your `global.ini` lines similar to this:

```
[ha_dr_provider_<HA_STONITH_Hook>]
provider = <HA_STONITH_Hook>
path = /hana/shared/HANA_Hooks
execution_order = 50
```

For all hooks, 0 must be returned in case of success.

```
"""
```

```

from hdb_ha_dr.client import HADRBase, Helper
import os, time

class HA_STONITH_Hook(HADRBase):
    def __init__(self, *args, **kwargs):
        # delegate construction to base class
        super(HA_STONITH_Hook, self).__init__(*args, **kwargs)

    def about(self):
        return {"provider_company" :      "DelleMC",
                "provider_name" :       "HA_STONITH_Hook", #
provider name = class name
                "provider_description" :  "Dell EMC IPMI stonith
for HANA",
                "provider_version" :    "1.0"}

    def startup(self, hostname, storage_partition,
system_replication_mode, **kwargs):
        self.tracer.debug("enter startup hook; %s" % locals())
        self.tracer.debug(self.config.toString())

        self.tracer.info("leave startup hook")
        return 0

    def shutdown(self, hostname, storage_partition,
system_replication_mode, **kwargs):
        self.tracer.debug("enter shutdown hook; %s" % locals())
        self.tracer.debug(self.config.toString())

        self.tracer.info("leave shutdown hook")
        return 0

    def failover(self, hostname, storage_partition,
system_replication_mode, **kwargs):
        self.tracer.debug("enter failover hook; %s" % locals())
        self.tracer.debug(self.config.toString())

        self.tracer.info("leave failover hook")
        return 0

```

```

def stonith(self, failingHost, **kwargs):
    self.tracer.debug("enter HANA HA stonith hook; %s" %
locals())
    self.tracer.debug(self.config.toString())

    self.tracer.info( "Stonith - rebooting failing host %s" %
failingHost)
    ipmi_host = "%s-ipmi" % failingHost
    # default credentials in command example, update credentials
for your environment
    power_off = "ipmitool power off -I lanplus -H %s -U root -P
calvin" % ipmi_host
    power_status = "ipmitool power status -I lanplus -H %s -U
root -P calvin" % ipmi_host
    power_on = "ipmitool power on -I lanplus -H %s -U root -P
calvin" % ipmi_host

#Power off failing host and check power status
(code, output) = Helper._runOsCommand(power_off)
    self.tracer.info(output)
    time.sleep(10)
(code, output) = Helper._runOsCommand(power_status)
self.tracer.info(output)
    if 'Power is off' in output:
        msg = "Successfully powered off %s" % failingHost
            self.tracer.info(msg)
        rc = 0
    elif 'Power is on' in output:
        msg = "failed to power off %s, will try again" %
failingHost
            self.tracer.info(msg)
        (code, output) = Helper._runOsCommand(power_off)
        self.tracer.info(output)
        time.sleep(10)
        (code, output) =
Helper._runOsCommand(power_status)
        self.tracer.info(output)
        if 'Power is off' in output:
            msg = "Successfully powered off %s" %
failingHost
                self.tracer.info(msg)

```

```

        rc = 0
    elif 'Power is on' in output:
        msg = "unable to power off %s - Please CHECK" %
failingHost
            self.tracer.info(msg)
        return 1
    #Power back on the failed host
    if rc == 0:
        (code, output) = Helper._runOsCommand(power_on)
        time.sleep(10)
        self.tracer.info(output)
        (code, output) =
Helper._runOsCommand(power_status)
        self.tracer.info(output)
        if 'Power is on' in output:
failingHost
            msg = "successfully powered on %s" %
                self.tracer.info(msg)
            rc = 0
        elif 'Power is off' in output:
            msg = "unable to power on %s - will try again to
power on" % failingHost
                self.tracer.info(msg)
            (code, output) = Helper._runOsCommand(power_on)
            self.tracer.info(output)
            time.sleep(10)
            (code, output) =
Helper._runOsCommand(power_status)
            self.tracer.info(output)
            if 'Power is off' in output:
Please CHECK" % failingHost
                msg = "unable to power on %s -
                    self.tracer.info(msg)
                    rc = 1
                elif 'Power is on' in output:
                    msg = "Successfully powered on
%s - Please CHECK" % failingHost
                        self.tracer.info(msg)
                    rc = 0
            self.tracer.info("leaving HANA HA stonith hook")

```

```

return rc

def preTakeover(self, isForce, **kwargs):
    """Pre takeover hook."""
    self.tracer.info("%s.preTakeover method called with
isForce=%s" % (self.__class__.__name__, isForce))

    if not isForce:
        # run pre takeover code
        # run pre-check, return != 0 in case of error => will
abort takeover
        return 0
    else:
        # possible force-takeover only code
        # usually nothing to do here
        return 0

def postTakeover(self, rc, **kwargs):
    """Post takeover hook."""
    self.tracer.info("%s.postTakeover method called with
rc=%s" % (self.__class__.__name__, rc))

    if rc == 0:
        # normal takeover succeeded
        return 0
    elif rc == 1:
        # waiting for force takeover
        return 0
    elif rc == 2:
        # error, something went wrong
        return 0

def srConnectionChanged(self, parameters, **kwargs):
    self.tracer.debug("enter srConnectionChanged hook; %s" %
locals())

    # Access to parameters dictionary
    hostname = parameters['hostname']

```

```

port = parameters['port']
database = parameters['database']
status = parameters['status']
databaseStatus = parameters['database_status']
systemStatus = parameters['system_status']
timestamp = parameters['timestamp']
isInSync = parameters['is_in_sync']
reason = parameters['reason']
siteName = parameters['siteName']

self.tracer.info("leave srConnectionChanged hook")
return 0

def srReadAccessInitialized(self, parameters, **kwargs):
    self.tracer.debug("enter srReadAccessInitialized hook; %s"
% locals())

    # Access to parameters dictionary
    database = parameters['last_initialized_database']
    databasesNoReadAccess =
parameters['databases_without_read_access_initialized']
    databasesReadAccess =
parameters['databases_with_read_access_initialized']
    timestamp = parameters['timestamp']
    allDatabasesInitialized =
parameters['all_databases_initialized']

    self.tracer.info("leave srReadAccessInitialized hook")
    return 0

def srServiceStateChanged(self, parameters, **kwargs):
    self.tracer.debug("enter srServiceStateChanged hook; %s" %
locals())

    # Access to parameters dictionary
    hostname = parameters['hostname']
    service = parameters['service_name']
    port = parameters['service_port']
    status = parameters['service_status']
    previousStatus = parameters['service_previous_status']

```

```

timestamp = parameters['timestamp']
daemonStatus = parameters['daemon_status']
databaseId = parameters['database_id']
databaseName = parameters['database_name']
databaseStatus = parameters['database_status']

self.tracer.info("leave srServiceStateChanged hook")
return 0

```

Install the HA/DR provider script

Add, configure, and monitor your custom provider scripts in SAP HANA Studio. After the HA/DR provider script is created, install it on an SAP HANA system by adding a section called `[ha_dr_provider_<classname>]` to the `global.ini` file with the following parameters:

- **Provider**—The class name.
- **Path**—The location of the script.
- **Execution_order**—The ordering of the HA/DR provider if there is more than one. This value is a number between 1 and 99.

For example:

```

[ha_dr_provider_<HA_STONITH_Hook>]
provider = <HA_STONITH_Hook>
path = /hana/shared/HANA_Hooks
execution_order = 50

```

Verify the installation of the HA/DR provider script `HA_STONITH.Hook.py`

All scripts are loaded during the startup phase of the name server. The name server trace files are monitored as general information about the `ha_dr_provider` and trace information and return codes are collected.

Perform host auto-failovers to ensure that the failovers work as expected and that the STONITH has been executed to the failed host. The following figure shows sample output from the nameserver tracefile following a host auto-failover and successful execution of the STONITH.

```
[28229]{-1}[-1/-1] 2018-11-13 00:08:53.499907 i failover          TREXNameServer.cpp(05934)
: failover from failed host r640-isi02 to host r640-isi01 started
[30171]{-1}[-1/-1] 2018-11-13 00:08:53.500693 i ha_dr_provider      PythonProxyImpl.cpp(00940)
: calling HA/DR provider HA_STONITH_Hook.stonith(failing_host=r640-isi02)
[30171]{-1}[-1/-1] 2018-11-13 00:08:53.502382 i ha_dr_HA_STONITH   HA_STONITH_Hook.py(00061)
: Stonith - rebooting failing host r640-isi02
[30171]{-1}[-1/-1] 2018-11-13 00:08:53.502487 i ha_dr_HA_STONITH   HA_STONITH_Hook.py(00067)
: ipmitool power off -I lanplus -H r640-isi02-ipmi -U root -P calvin
[30171]{-1}[-1/-1] 2018-11-13 00:08:59.970498 i ha_dr_HA_STONITH   HA_STONITH_Hook.py(00085)
: successfully powered off r640-isi02
[30171]{-1}[-1/-1] 2018-11-13 00:09:09.653074 i ha_dr_HA_STONITH   HA_STONITH_Hook.py(00106)
: successfully powered on r640-isi02
[30171]{-1}[-1/-1] 2018-11-13 00:09:09.653217 i ha_dr_HA_STONITH   HA_STONITH_Hook.py(00116)
: leaving HANA HA stonith hook
[28229]{-1}[-1/-1] 2018-11-13 00:09:09.915674 i failover          TREXNameServer.cpp(05987)
: failover from failed host r640-isi02 to host r640-isi01 succeeded
```

Figure 12. Ha_dr_provider output from the master nameserver tracefile

Conclusion

Summary

SAP has certified Dell EMC Isilon F800 all-flash scale-out NAS storage for use in SAP HANA installations on production and nonproduction systems and on single-node (scale-up) and multi-node (scale-out) systems. This new TDI enterprise storage solution is the first certified NAS solution from Dell EMC to complement the many SAP HANA-certified SAN storage offerings from Dell EMC.

Using SAP HANA in TDI deployments with Isilon F800 all-flash scale-out NAS storage provides many benefits, including reducing hardware and operational costs, lowering risk, improving availability and performance, and increasing hardware vendor flexibility.

Findings

During our tests with SAP HANA on Isilon F800 all-flash scale-out NAS storage systems, we observed the following:

- The SAP HANA-HWC-ES 1.1 certification scenario makes higher demands of disk configuration.
- SAP HANA production installations on Isilon systems require SSDs for the SAP HANA persistence.
- Using SSDs for the SAP HANA persistence provides significant benefits, including:
 - Reduced SAP HANA startup and host auto-failover times
 - Reduced SAP HANA backup and recovery times
- SAP HANA client connections must be evenly distributed across the Isilon nodes interfaces to ensure optimal performance.

References

Dell EMC documentation

The following Dell EMC documentation provides additional and relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell EMC representative.

- [Isilon All-Flash Scale-Out NAS Storage Specification Sheet](#)
- [Isilon Network Design Considerations](#)
- [Isilon OneFS NFS Design Considerations and Best Practices](#)
- [Isilon OneFS Version 8.1.0 CLI Administration Guide](#)
- [Isilon OneFS Web Administration Guide](#)
- [Dell EMC Isilon OneFS Best Practices](#)

For SAP Ready Solution documentation, see [SAP Info Hub for Ready Solutions](#) in the Everything SAP at Dell EMC Community forum.

SAP documentation

The following SAP documentation provides additional and relevant information:

- [SAP HANA Master Guide](#)
- [SAP HANA Server Installation and Update Guide](#)
- [SAP HANA Studio Installation and Update Guide](#)
- [SAP HANA Technical Operations Manual](#)
- [SAP HANA Administration Guide](#)
- [SAP HANA Storage Requirements](#)

Web resources

- [SAP HANA Platform](#)
- [SAP HANA One](#)
- [SAP HANA Enterprise Cloud](#)
- [SAP HANA Tailored Data Center Integration \(TDI\)](#)
- [SAP HANA Tailored Data Center Integration - Frequently Asked Questions](#)
- [SAP HANA Hardware Directory](#)

Note: The following documentation requires SAP login credentials.

- [SAP Note 1943937 - Hardware Configuration Check Tool - Central Note](#)
- [SAP Note 2399079—Elimination of hdbparam in HANA 2](#)