

m

IBM® System Storage™ Architecture and Configuration Guide for SAP® HANA™ Tailored Datacenter Integration

isicc@de.ibm.com

Version 2.5 – Feb 2016





Table of Contents

| | | |
|----|--|----|
| 1 | Preface | 3 |
| 2 | IBM storage architecture for SAP HANA TDI..... | 3 |
| 3 | The central theme..... | 4 |
| 4 | Requirements..... | 4 |
| 5 | Installation of SUSE SLES or Red Hat | 5 |
| 6 | SAN storage system setup | 5 |
| 7 | SAN zoning and volume mapping | 9 |
| 8 | Setup of multipath | 10 |
| 9 | Setup of file systems for DATA and LOG | 10 |
| 10 | Setup of LINUX IO scheduler (elevator) | 11 |
| 11 | Setup of /hana/shared/global.ini..... | 12 |
| 12 | Setup of shared access for the /hana/shared/ directory | 12 |
| 13 | SAP HANA TDI High Availability setup with IBM System Storage | 13 |
| 14 | SAP HANA TDI Disaster Recovery setup with IBM System Storage | 14 |
| 15 | SAP HANA TDI backup with IBM Spectrum Protect | 14 |
| 16 | Resources | 14 |
| 17 | Change History | 16 |
| 18 | Trademarks..... | 16 |
| 19 | Disclaimers..... | 17 |

1 Preface

This paper is intended as an architecture and configuration guide to setup the IBM® System Storage™ for the SAP® HANA™ tailored datacenter integration (SAP HANA TDI) within a SAN environment.

This document has been written for IT technical specialists and architects, with advanced skill levels on SUSE® Linux™ Enterprise Server or RedHat LINUX® and IBM System Storage, with a focus on architecting and setting up an IBM System Storage environment for SAP HANA TDI. The SAP HANA TDI allows the SAP customer to use external storage to attach to the SAP HANA server.

This document provides the necessary information to select, verify and connect IBM System Storage to the SAP HANA server through a SAN (fiber channel) attached storage network. The recommendations in this guideline apply to both single node and scale out configurations, as well as Intel® or IBM POWER® based SAP HANA systems.

For a list of all storage systems certified for SAP HANA production please visit:

<http://global.sap.com/community/ebook/2014-09-02-hana-hard-ware/enEN/enterprise-storage.html>

2 IBM storage architecture for SAP HANA TDI

The IBM System Storage needs to be connected through a SAN network with the SAP HANA server. IBM recommends – but not requires – the use of the IBM Spectrum Virtualize (SAN Volume Controller, SVC) as virtualization layer, enabling to centralize the management of the storage systems, as well as to provide easily high availability and disaster recovery functionality, as well as IBM FlashCopy backup and tiered storage pooling. For smaller HANA TDI deployments you can also use the IBM Storwize V7000, which is based on the same functions like the SVC, and inherits its advantages described above.

All IBM Storage Systems are certified for SAP HANA TDI production:

- IBM Storwize family, model V7000, V5000, V3700, and IBM Spectrum Virtualize
- IBM FlashSystem family, model 820, 840, 900, V840, V9000
- IBM XIV Gen 3 Storage System
- IBM DS8800 and DS8870
- IBM Spectrum Scale ESS GS2,4,6 & GL2,4,6

The setup of IBM Spectrum Scale™ (IBM GPFS™) storage system ESS (Elastic Storage Server) is documented in a separate document, located at the same document link WP102347.

For non-prod SAP HANA systems any storage system can be used.

3 The central theme

Below the generic workflow describes how to setup SAP HANA TDI with IBM storage:

- I. Verify all server & storage components and firmware requirements given by SAP are met
- II. Install SLES 11 or Red Hat operating system on all SAP HANA nodes
- III. Setup SAN storage systems
- IV. Create SAN zone for SAP HANA environment
- V. On the storage systems map the storage volumes (LUNs) to all target SAP HANA nodes
- VI. Setup of multipath
- VII. On each LINUX host create the file systems for *SAP HANA data and log*
- VIII. Setup of `/hana/shared/global.ini`
- IX. Setup NFS services for the SAP HANA shared directory
- X. Install and configure the SAP HANA software
- XI. Optional: HA, DR, and Backup considerations

4 Requirements

Before starting to deploy a SAP HANA TDI, ensure that the following key items have are in place:

- A SAN environment is required to attach IBM System Storage through fiber channel with the SAP HANA TDI nodes.
- SAP requires a storage validation for SAP HANA TDI according to their KPI (key Performance Indicators) – please contact SAP for further details.
- The SAP HANA nodes must be sized and configured according to SAPs specification (Bill of Material for SAP HANA node), in addition each SAP HANA TDI node needs to have 4 FC ports, one HBA with 4 ports, or 2 HBA with 2 ports each.

5 Installation of SUSE SLES or Red Hat

Install *SUSE LINUX Enterprise Server™ 11*, *SUSE LINUX Enterprise Server for SAP™ 11* – Service Pack 3 or higher, or Red Hat Enterprise Linux for SAP HANA on local disks or on SAN attached storage. If SAN attached storage is used, layout the volumes to match the best practices guideline provided by the Linux distributors.

If necessary upgrade the kernel to minimum **level 3.0.101**

6 SAN storage system setup

This section covers the evaluated storage options.

6.1 Recommended IBM System Storage combinations

To achieve the required storage performance according to SAP KPIs (key performance indicators) for SAP HANA TDI, one of the following configurations listed should be chosen – nevertheless all IBM systems listed can be used as stand-alone system as well, used for DATA and LOG.

IBM Spectrum Virtualized is referred as SVC.

| Storage virtualization | SAP HANA data | SAP HANA log |
|------------------------|-----------------|--------------------------|
| SVC | IBM XIV Gen 3 | IBM FlashSystem |
| SVC | IBM Storwize | IBM FlashSystem |
| SVC | IBM FlashSystem | IBM FlashSystem |
| SVC | IBM DS8000 | IBM FlashSystem / DS8000 |

Because IBM SVC, and for smaller configurations IBM Storwize, provisions *volumes* (LUNs) as a single volume type and single storage system, it will ease the administration of storage pool tiers with Flash, SSD, and HDD types; and optional tasks like setup of storage high availability, FlashCopy base backup, storage based mirroring, or storage maintenance. If the IBM storage system will be direct attached the SAP HANA TDI nodes, then the following table specifies the recommended disk types.

| | SAP HANA data | SAP HANA log |
|-----------------|---------------|--------------|
| IBM Storwize | HDD | SSD |
| IBM DS8000 | HDD | SSD |
| IBM XIV | Module | Module |
| IBM FlashSystem | Flash Card | Flash Card |

6.1 Storage sizing through SAPmagic tool

The IBM storage systems should be configured (sized) through the IBM tool SAPmagic; this tool is available to IBM Sales and IBM Business Partner on IBM Tech-Docs.

Note: The performance sizing described in this chapter is only required for productive HANA nodes (HANA databases for the SAP production systems) and not for the non-production HANA Nodes (e.g. QA, Test, Dev, Sandbox Systems). For the non-production systems only a capacity sizing is required, DATA capacity is roughly three times the RAM, and LOG capacity is one time the RAM size of the HANA node -- please check with the latest SAP HANA installation guide.

6.2 Required IBM System Storage HW configurations

This chapter provides some general sizing guidelines.

All components of the SAN infrastructure must be configured to use a minimum 8GB link speed. All listed storage configurations apply to active SAP HANA TDI nodes only, no additional storage configuration is required for any stand-by node.

If SVC stretched cluster needs to be configured, all listed storage configuration must be doubled.

Maximal number of HANA production node support per IBM System Storage.

| IBM System Storage | # nodes |
|---------------------------------------|----------|
| Storwize V3700 | 6 |
| Storwize V5000 | 8 |
| Storwize V7000 | 10 |
| Storwize V7000 cluster with 4 systems | 40 |
| FlashSystem V9000 | 40 |
| FlashSystem 900 | 10 |
| FlashSystem 840/V840 | 10 |
| FlashSystem 820 | 10 |
| DS8870 | 16 |
| DS8884/6/8 | 16/32/40 |
| XIV Gen 3 | 12 |
| SVC 2145-DH8, one node pair | 10 |
| SVC 2145-DH8, four node pairs | 40 |
| SVC 2145-CG8, one node pair | 5 |
| SVC 2145-CG8, four node pairs | 20 |



IBM Storwize

- HANA DATA: For every 1 node 8 HDD as RAID 5, 10 or 15krpm
- HANA LOG: For every 4 nodes 8 SSD as RAID 5
- HDD & SSD capacity size can be chosen arbitrary to match SAP HANA capacity requirements
- For a small number of nodes, and small RAM size the LOG RAID can be built as RAID 10, with 2, 4, or 6 SSD – depending on actual sizing.

IBM FlashSystem V9000

- One building block supports up to 10 HANA nodes
- The maximum number of SAP HANA nodes is 40 by using 4 building blocks
- One building block is one IBM FlashSystem V9000 Control Enclosure (AC2) and one to four FlashSystem 900 (AE2) with 12 Flash Cards

IBM FlashSystem 900

- Supports up to 10 HANA nodes

IBM FlashSystem 820, 840, V840

- Supports up to 10 HANA nodes

IBM XIV Storage System

- 15 Modules supporting 12 HANA nodes

IBM DS8000 family up to 16 nodes

- HANA DATA: For every 2 HANA nodes 1 HDD Drive Set (16 disks), RAID 5, 10 or 15 krpm
- HANA LOG: For every 8 HANA nodes 1 SSD Drive Set (16 disks), RAID 5
- HDD & SSD capacity size can be chosen arbitrary to match SAP HANA capacity requirements

IBM DS8884/6/8 family up to 40 nodes

- HANA DATA & LOG: For every 4 HANA nodes 1 SSD Drive Set (16 disks RAID 5) or one Flash Card.
- SSD or Flash capacity size can be chosen arbitrary to match SAP HANA capacity requirements

6.3 Definitions of storage pools and volumes

The SAP HANA Data and Log have different I/O characteristics reflected by the different storage systems configurations as described earlier. This has to be applied by pooling the appropriate storages accordingly.

For all storage systems:

- Create 2 volumes for HANA DATA or 1 volume for HANA LOG, for each node
- These volumes would also server as MDISKS if SVC will be used

IBM SVC storage pool and VDISK definitions

- Create separate pools for DATA and LOG
- Create 2 VDISKS for HANA DATA and 1 or 2 VDISKS for HANA LOG, for each node
- Switch off VDISK caching for pools with FlashSystem MDISK

6.4 Layout of storage pools and storage volumes

The following diagrams show the relationship between Linux file system, Logical Volumes (VDISKS), and Managed Disks (physical volumes) of a SVC served by various storage subsystems in the backend.

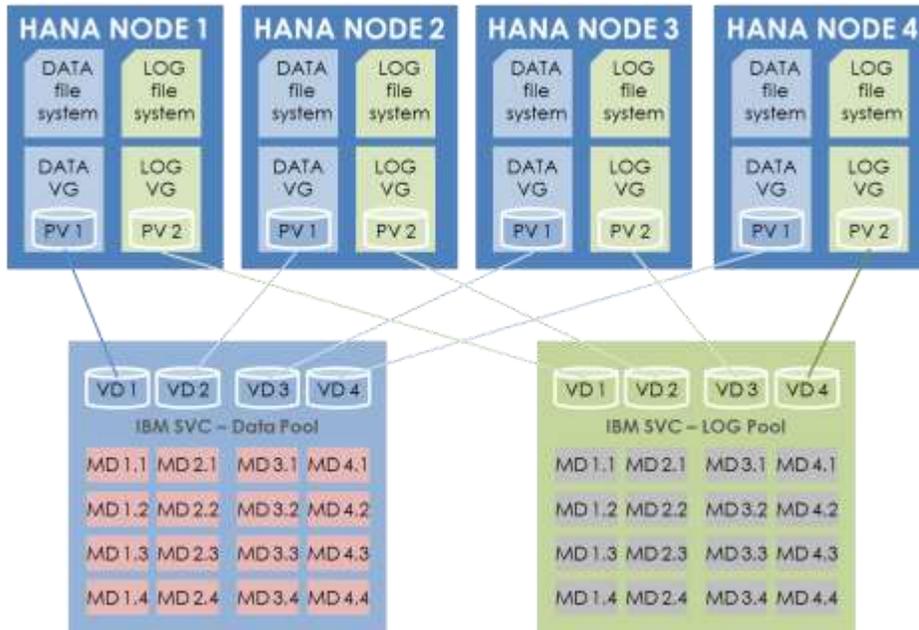


Figure 1: SVC storage pool (MDG) setup

Figure 1 shows the SVC pool setup for data and log. The virtual disks (VD) are assigned to the HANA TDI nodes, the managed disks will be built from the underlying physical storage as pictures in Figure 2 to Figure 4.

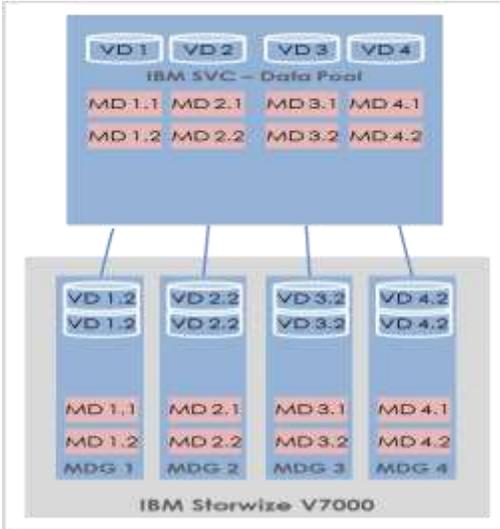


Figure 2: Storwize V7000 setup for DATA or LOG

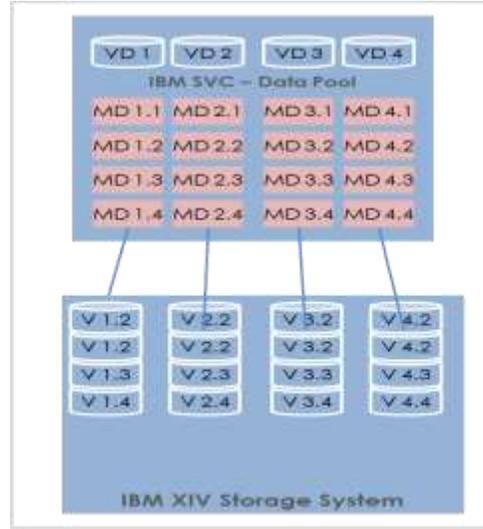


Figure 3: XIV setup for DATA or LOG

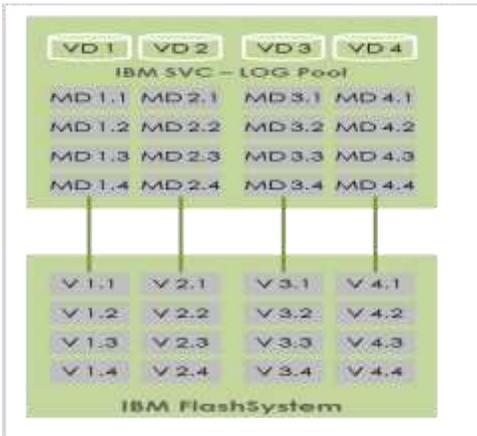


Figure 4: FlashSystem setup for DATA or LOG

7 SAN zoning and volume mapping

Apply standard, best practices zoning within the SAN.

If HANA High Availability node-fail over is use then, and only then it is required to map all Data and Log volumes to **all** SAP Hana Nodes.

8 Setup of multipath

Setup multipath according SLES or Red Hat documentation, enable user friendly volume names via alias in `/etc/multipath.conf`:

```
defaults {
    polling_interval    30
    failback            immediate
    no_path_retry      5
    user_friendly_names no
}

devices {
    device {
        vendor          "IBM"
        product         "2145"
        failback        immediate
        no_path_retry   5
        rr_min_io       10
        dev_loss_tmo    120
        path_checker     tur
        path_grouping_policy group_by_prio
        path_selector   "round-robin 0"
        prio             const
        rr_weight       uniform
    }
}
```

9 Setup of file systems for DATA and LOG

Create the DATA and LOG file system needed by SAP HANA according to the *SAP HANA Storage Requirements* documentation from SAP.

- Use the **xfs** file system type with a 4KB block size for the DATA and LOG file systems
- To use LVM 2 use this schema:
vgcreate → lvcreate → mkfs.xfs
- Use "useful" names for volume groups and logical volumes – do not use the hyphen (minus) sign '-' as part of the name, because the LINUX mapper uses this sign as separator between VG and LV name.
- Create Volume Groups and Logical Volumes for DATA and LOG, for each HANA mount point (as listed in `global.ini`) create one Volume Group with exactly one Logical Volume.

It has turned out that these settings are most beneficial for SAP HANA workload in combination with SVC:

1. Use 4 PV per Volume Group
2. Create a volume group with an extent size of 4MB and an alignment size of 1MB

```
vgcreate vgcreate -s 4M --dataalignment 1M VG_ANA_DATA_1 /dev/map-  
per/WWID_1 /dev/mapper/WWID_2 ...
```

3. Create a logical volume with a stripe size of 256KB

```
lvcreate -i 4 -I 256K -l 100%VG -n LV_ANA_DATA_1 VG_ANA_DATA_1
```

4. Create the file system

```
mkfs.xfs /dev/mapper/VG_ANA_DATA_1-LV_ANA_DATA_1
```

5. Mount the file system

```
mount /dev/mapper/VG_ANA_DATA_1-LV_ANA_DATA_1 /hana/ANA/data
```

Do this for LOG accordingly.

Do **not** add the file system to /etc/fstab, mounting will be done by SAP HANA.

More information on IO tuning can be found in this whitepaper:

LINUX IO performance tuning for IBM System Storage

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102584>

10 Setup of LINUX IO scheduler (elevator)

The default scheduler cfq delivers low IO performance for multiple (parallel) streaming read processes. It seems that deadline delivers slightly better performance than noop, special during multiple (parallel) IO operations.

We recommend to use the **deadline** scheduler.

11 Setup of /hana/shared/global.ini

Enter the Logical Volume names as shown in directory /dev/mapper into /hana/shared/global.ini :

```
[communication]
listeninterface = .global

[persistence]
basepath_datavolumes = /hana/ANA/data
basepath_logvolumes  = /hana/ANA/log

[storage]
ha_provider           = hdb_ha.fcClientLVM
partition_*_*_prtype = 5
partition_1_data__lvmname = VG_ANA_DATA_1-LV_ANA_DATA_1
partition_1_log__lvmname  = VG_ANA_LOG_1-LV_ANA_LOG_1
partition_2_data__lvmname = VG_ANA_DATA_2-LV_ANA_DATA_2
partition_2_log__lvmname  = VG_ANA_LOG_2-LV_ANA_LOG_2
```

The [storage] section is only needed, if you want to use the node fail-over functionality provided by SAP HANA software. If you use the system replicating capability leave this section empty.

12 Setup of shared access for the /hana/shared/ directory

SAP HANA **scale-out** and HANA HA node fail-over require that the directory /hana/shared/ must be accessible from all nodes (else this directory can be set-up as local file system from type xfs).

The shared access can be achieved most easily through NFS. If such setup does not already exist, one of the following options can be used to implement it:

- IBM Storwize V7000 Unified
- High Availability NFS service with DRBD and IBM Tivoli System Automation for Multiplatform (SA MP) with SUSE Linux Enterprise High Availability Extension 11.
- Highly Available NFS service with DRBD and Pacemaker with SUSE Linux Enterprise High Availability Extension 11.
- GPFS file system instead of NFS

Setup these systems according to the appropriate implementation guides (see chapter [Resources](#)), and make the central, shared NFS file system available on all HANA nodes, including the spare nodes.

Basic setup of NFS Server, add this line to /etc/exports

```
/hana/shared node1(fsid=0,crossmnt,rw,no_root_squash, sync,no_subtree_check)
node2(fsid=0,crossmnt,rw, ...) node3( ...)
```

Basic setup of the HANA nodes (NFS clients), add this line to /etc/fstab

```
nfs-server:/hana/shared /hana/shared nfs rw,soft,intr,rsize=8192,wsiz=8192 0 0
```

13 SAP HANA TDI High Availability setup with IBM System Storage

The SAP HANA system gains high availability through an N+1 concept – one or more server act as standby SAP HANA node – with dedicated storage for each active node. If one active SAP HANA node fails, the SAP HANA cluster software initiates a failover to the standby node, and the standby node will mount the data and log from the failed node.

For more details please read the SAP document [SAP HANA – High Availability](#).

To improve storage availability the ISICC recommends to setup the storage environment as a SVC stretched cluster with symmetric VDISK mirroring, as shown in Figure 5.

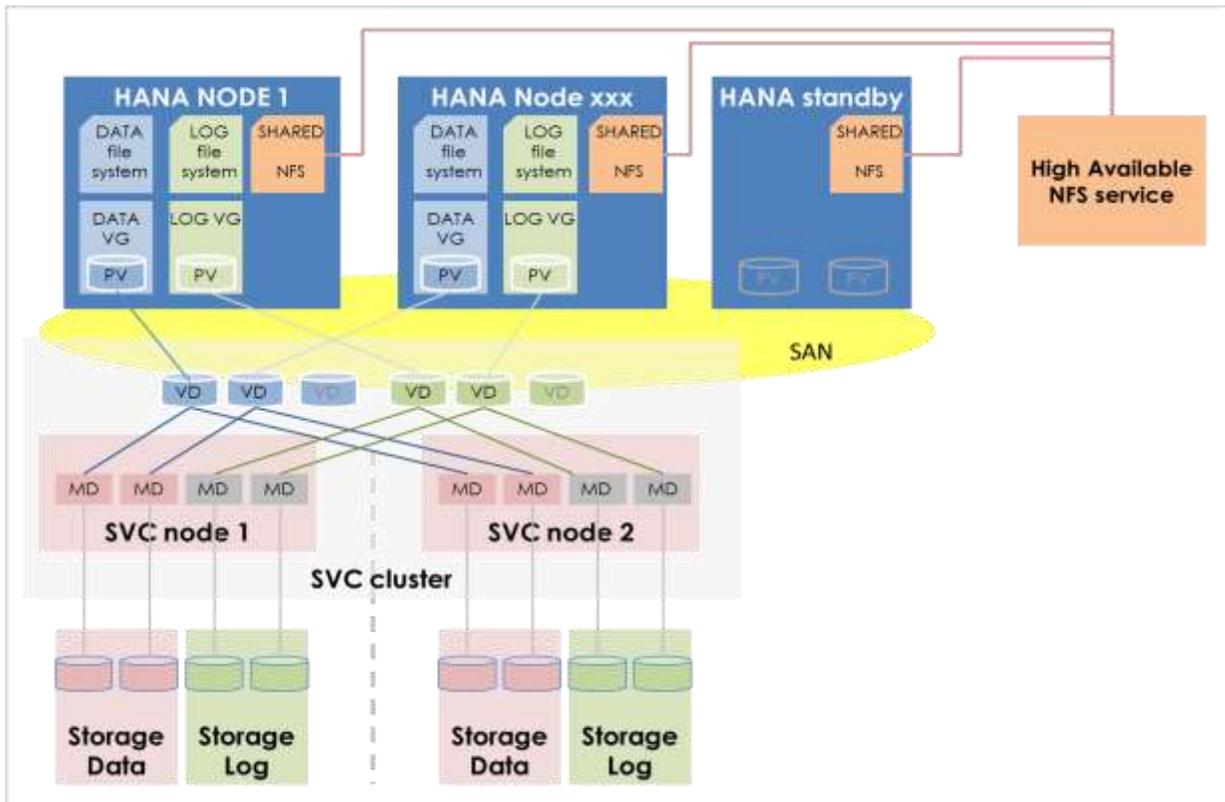


Figure 5: SAP HANA TDI HA setup with SVC stretched cluster for data redundancy

14 SAP HANA TDI Disaster Recovery setup with IBM System Storage

The disaster recovery capabilities of SAP HANA are documented on www.saphana.com:

- SAP HANA Scale-Out, High Availability & Disaster Tolerance
- SAP HANA – High Availability

If IBM Metro Mirror (synchronous) is used for storage replication for SAP HANA TDI, all DATA and all LOG volumes of all nodes from one HANA system must be included in **one** storage **Consistency Group**.

That implies (when using native attached storage) that both – Data and Log – must come from the same storage unit.

15 SAP HANA TDI backup with IBM Spectrum Protect

IBM Spectrum Protect, also known as Tivoli Storage Manager for Enterprise Resource Planning (TSM for ERP) includes the package Data Protection for SAP HANA®.

Please see current [product documentation](#) how to install, configure, and run this integrated solution.

16 Resources

For any product or documentation provided by SAP please contact SAP.

For any product or documentation provided by SUSE please contact SUSE.

For any product or documentation provided by IBM regarding SAP send an email to isicc@de.ibm.com.

IBM documentations:

This document is ID [WP102347](#) on www.ibm.com

IBM Storwize product manuals:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S7003318>

Implementing the IBM SAN Volume Controller and FlashSystem 820

<http://www.redbooks.ibm.com/abstracts/sg248172.html>



Installation, configuration and usage of Tivoli Storage Manager for Enterprise Resource Planning V6.4 Data Protection for SAP HANA®

<http://www-01.ibm.com/support/docview.wss?uid=swg21608240>

IBM Tivoli SA for MP Installation and Configuration Guide

http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.samp.doc_3.2.1/HALICG21.pdf

SAP documentation:

Overview - SAP HANA tailored data center integration

<http://www.saphana.com/docs/DOC-3633>

FAQ - SAP HANA tailored data center integration

<http://www.saphana.com/docs/DOC-3634>

Introduction to High Availability for SAP HANA

<http://www.saphana.com/docs/DOC-2775>

IBM/SAP Whitepaper "High-end customer IT Landscapes based on SAP HANA"

<http://www.saphana.com/docs/DOC-3211>

Access to sapjam and documentation will be provided by SAP – the listed links only work after login: login first, then access the links.

SAP HANA Fiber Channel Storage Connector Admin Guide

SAP HANA Storage Requirements (Storage Whitepaper.pdf)

<https://jam4.sapjam.com/wiki/show/60549>

SAP HANA reference architecture

<https://jam4.sapjam.com/wiki/show/48495>

SAP HANA Bill of Material

<https://jam4.sapjam.com/wiki/show/48202>

Novell / SUSE documentation:

List of all SLES 11 documentations

<https://www.suse.com/documentation/sles11/>

Highly Available NFS Storage with DRBD and Pacemaker with SUSE Linux Enterprise High Availability Extension 11

https://www.suse.com/documentation/sle_ha/single-html/book_sleha_techguides/book_sleha_techguides.html



17 Change History

| | |
|------|--|
| V1 | Initial version |
| V1.2 | Update of chapter 0 to clarify files system setup |
| V1.3 | Included new Storwize V5000 |
| V1.4 | Specified NFS section to be used only for /hana/shared/ |
| V1.5 | Minor updates – add some clarifications |
| V1.6 | Add setup of multipath and global.ini |
| V1.7 | Rework of document |
| V1.8 | Added LINUX LVM2 & list of certified IBM storage systems |
| V2.1 | Added IBM DS8000 system as certified for SAP HANA TDI |
| V2.2 | Added IBM FlashSystem 900 & V9000 |
| V2.3 | Changes sizing rules, added recommendation for multipath and VG/LV settings. |
| V2.5 | Changed /etc/multipath setting |

18 Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®, BladeCenter®, DB2®, Global Business Services®, Global Technology Services®, GPFST™, IBM SmartCloud®, IBM®, Intelligent Cluster™, Passport Advantage®, POWER®, PureFlex™, RackSwitch™, Redbooks®, Redpaper™, Redbooks (logo) ®, System Storage®, System x®, System z®, Tivoli®, z/OS®.

The following terms are trademarks of other companies:

SAP, R/3, ABAP, BAPI, SAP NetWeaver, Duet, PartnerEdge, ByDesign, SAP BusinessObjects Explorer, StreamWork, SAP HANA, the Business Objects logo, BusinessObjects, Crystal Reports, Crystal Decisions, Web Intelligence as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany or an SAP affiliate company.

Intel Xeon, Intel, Itanium, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

SUSE is a registered trademark of SUSE Germany and Novell USA, other countries, or both. SLES is a trademark of SUSE Germany and Novell USA, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

19 Disclaimers

This information was developed for products and services offered in Germany.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to: IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION 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. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.