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

location of this document:
http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102347

escc.isv-services@de.ibm.com

Version 2.22 – July 2020

Content

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:

https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/enterprise-storage.html#categories=certified%23International%20Business%20Machines%20Corporation

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 FlashSystem, which is based on the same functions of SVC, and inherits its advantages described above.

IBM Storage Systems are certified for SAP HANA TDI production:

- IBM Spectrum Virtualize (SVC)
- IBM FlashSystem family, model FS900, FS9100, V9000, A9000R
- IBM FlashSystem 5000/7000 family (formerly Storwize)
- IBM DS8880, DS8880F, DS8900F

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

© Copyright IBM Corporation 2020

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 or Red Hat operating system on all SAP HANA nodes

 III.   If required: setup NFS services for the SAP HANA shared directory

 IV.   Setup SAN storage systems

   V.   Create SAN zone for SAP HANA environment

 VI.   On the storage systems map the storage volumes (LUNs)
to all target SAP HANA nodes

 VII.   Perform Linux setup and tuning, e.g. multipath, IO scheduler, …

VIII.   On each LINUX host create the file systems for *SAP HANA data, log, shared*

 IX.   Setup of /hana/shared/global.ini for HANA scale out or HANA node failover

   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 system needs to have min 4 FC ports, recommended are 8 FC ports.

Naming: The term "node" is use for a single HANA instance: if for instance 4 HANA instances are deployed on a single IBM POWER server, or 4 HANA instances are deployed into four single LPAR, then this reflects 4 nodes.

5   Installation of SUSE SLES or Red Hat

Install the Linux operating System accordingly to the SAP PAM for the chosen SAP HANA version, see SAP note https://launchpad.support.sap.com/#/notes/1944799

6   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 setup 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:

> Highly Available NFS service with DRBD and Pacemaker with SUSE Linux Enterprise High Availability Extension.

Setup these systems according to the appropriate implementation guides, and make the 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(...) node3(...)
```

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

```
nfsserver:/hana/shared /hana/shared nfs4 rw,soft,intr,rsize=8192,wsize=8192 0 0
```

7   IBM System Storage sizing & configuration

This section covers the evaluated storage options.

7.1   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 Virtualize is referred as SVC.

Because IBM SVC, and for smaller configurations IBM FlashSystem, 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.

IBM recommends the use of read intensive SSD or Flash Systems for any kind of work-load; no HDD should be used anymore for online data.

Near line HDD (NL HDD) are suitable for archive data, or for backup as last-storage-pool.

7.2   NSE Considerations

SAP HANA Native Storage Extension (NSE) is an SAP HANA feature to keep so called warm data on disk instead of loading it into memory. NSE enabled data is loaded on demand only.

NSE requires to store the data on fast storage with low IO latencies. Therefore, the storage pool providing /hana/data should be formed of the fastest available drives of the storage system. For FlashSystem, we recommend using at least FlashCore Modules (FCM) or even better Storage Class Memory (SCM).

To avoid any movement of such data to lower storage tiers, ensure also that *IBM Storage Easy Tier* is disabled on the storage system.

7.3   IBM System Storage sizing

This chapter provides some <u>general</u> sizing guidelines.

All components of the SAN infrastructure must be configured to use a minimum 8GB link speed with 4 links. All listed storage configurations apply to active SAP HANA TDI nodes only; no additional storage configuration is required for any stand-by system. If SVC stretched cluster needs to be configured, all listed storage configuration must be doubled.

Maximal number of HANA production systems support per IBM System Storage:

IBM System Storage	# HANA Prod
IBM FlashSystem 5010(H)/5030(H)/5100(H)	4/18*/26*
IBM FlashSystem 7200(H) 824	112*
IBM FlashSystem 9200(R)	164*
IBM FlashSystem A9000(R)	48
IBM FlashSystem 900	12
DS8910F (rack mounted) /DS8910F / DS8950F	16 / 16 / 48
SVC 2145/2147 SV2, one / four node pairs	25 / 100

* with max. supported IOGroups

7.4   RI SSD, Flash Drives, Flash Core Modules (FCM)

Since 2016, IBM offers RI SDD for the formerly Storwize (now: IBM FlashSystem) product line, currently the capacity ranges from approx. 2 TB up to 15 TB per disk. RI stands for Read Intensive.

These RI SDD are specified with one drive-write per day (DWPD), meaning, one full overwrite of the entire capacity per day at maximum.

During lifetime of all kind of flash technology (RI SSD up to FCM) the capacity utilization should not exceed 80%. If the utilization approaches this limit, additional capacity of the same kind and device capacity need to be added.

7.5   IBM storage performance sizing for HANA

This section provides sizing guideline how to size IBM storage for HANA.

If not stated otherwise (and if applicable), use distributed RAID 6.

The statement *"Supports up to xx HANA systems"* refer to SAP production systems, in case of non-prod systems, like Dev/Test/Sand much more HANA system can be used.

IBM SAN Volume Controller

- SV2: Each IO group (two nodes) can handle 25 HANA nodes, a four IO-group cluster is able to handle 100 HANA systems
- SV1: Each IO group (two nodes) can handle 25 HANA nodes, a four IO-group cluster is able to handle 100 HANA systems

IBM FlashSystem 5010 / 5030

- 12NVMe SSD for every 6 HANA systems

IBM FlashSystem 5100 / 7200

- 12 NVMe FCM for every 13/18 HANA systems

IBM FlashSystem A9000 / A9000R

- One building block supports up to 12 HANA systems
- One building block is one Flash enclose and two Grid controllers
- The maximum number of SAP HANA nodes is 48 by using 4 building blocks

IBM FlashSystem 900

- Supports up to 12 HANA systems
- One Flash Module supports one HANA node, it recommended to use always 12 Modules per FlashSystem

IBM FlashSystem 9200

- Supports up to 41 HANA production nodes
- It is recommended to use full equipped units only, with 24 NVMe Flash Core Modules.

IBM DS8000 family

- 16 Flash Cards for every 8 nodes

7.6   Sizing tool MyFlash4HANA

IBM provides a tool to size in IBM storage systems for a given/planed HANA environment. This tool is available only to IBM and IBM Business Partner.

Please be so kind and ask your IBM sales representative or IBM Business Partner for a detailed IBM Storage sizing for SAP HANA.

7.7   Definitions of storage pools and volumes

Even though SAP HANA Data and Logs have different I/O characteristics, both workloads can be placed on the same SSD / Flash storage pool since Flash technology can handle any type of workload. With DRAID 6 there is no need to place DATA and LOGs in two different pools to protect against drive failure.

For all storage systems:

- Create 4 volumes for HANA DATA and 4 volumes for HANA LOG,
  for each node, recommendation: thin provisioned
- Use minimum 8 VD if used as MDISKs for SVC, increase number by four, in case more capacity or more performance are needed (8,12,16, …)
  Recommendation: full provisioned
- Switch on Volume / VDISK caching

7.8   The use of Volume Compression (in DRP or RtC)

Real time Compression (RtC) and Data Reduction Pools (DRP) are both features to enable data compression on VDisk / Volume level. Because the data of the HANA database are "just" integers, these data cannot be compressed much. Compression should not be used.

## 7.9   The use of Data Deduplication in Data Reduction Pools (DRP)

The IBM storage systems provide the option to deduplicate data on VDISK / Volume level. Because the HANA database has an internal deduplication feature, additonal benefit of VDisk deduplication is not expected. Deduplication should not be used.

7.10   The use of IBM inline compression

The IBM Flash Core Modules, available for IBM FlashSystem, provides HW bases compression. These FCM can and should be used for SAP HANA.

7.11   Layout of storage pools and storage volumes

The following diagrams show the relationship between Linux file system, LINUX LV, VG, and PV, and the Virtual Volumes (VD, VDISK) and Managed Disks (MD, MDISK, RAIDs/Arrays) of a Storage Systems.

© Copyright IBM Corporation 2020
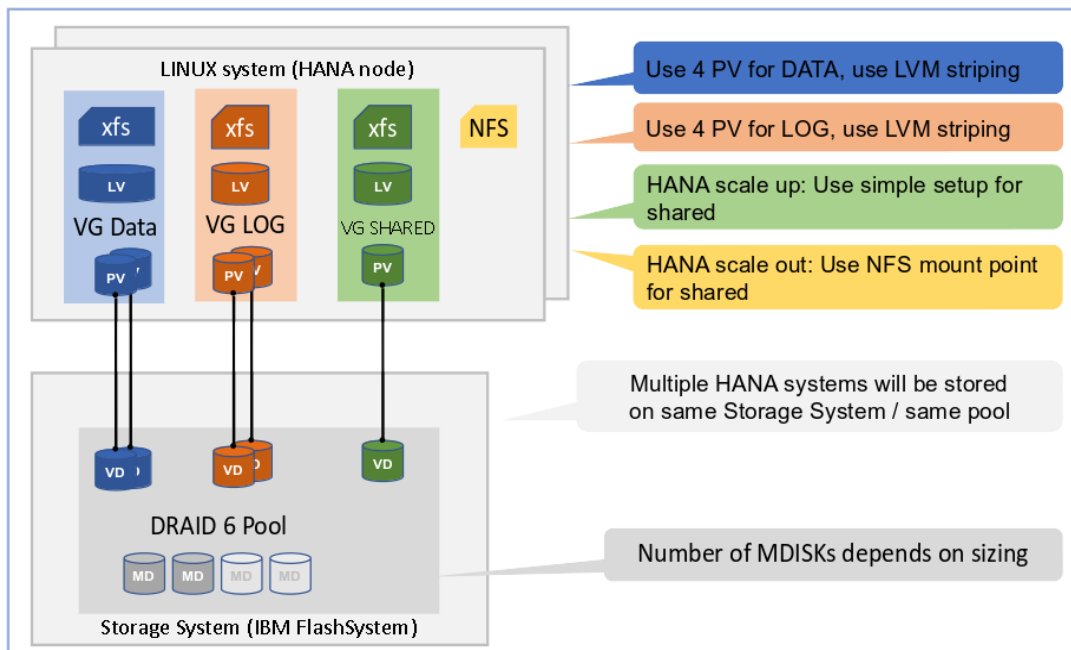
7.12   IBM SVC stretched cluster configuration

IBM SVC stretched cluster configuration with VDISK mirror is the standard implementation to gain best high availability (HA) solution for storage. The following picture illustrates the relationship between LINUX physical volumes with IBM Virtual and IBM Managed Disks.
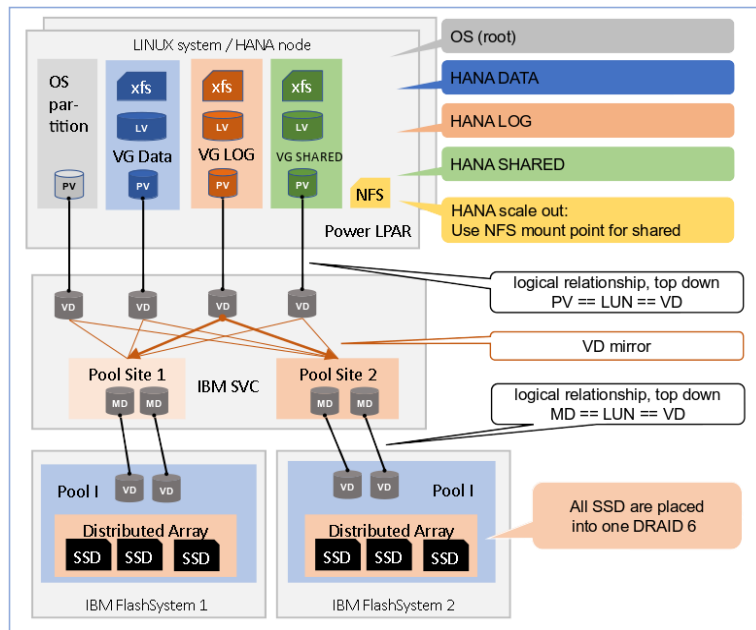
More details can be found in the IBM Redbook " Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize":

7.13   Competing storage utilization

Every IBM storage system can be shared between SAP HANA (production) and any other SAP or non-SAP (production) workload.

To do so proper sizing, the entire storage infrastructure needs be evaluated, including the configurations of RAID controller, number of HDD, SSD, or Flash Modules, number and type of HBA and FC ports/links.

8   Linux setup & tuning

This section provides guides about the required setup and tuning, additional information can be found in guide "SUSE Linux Enterprise Server (12.x) for SAP Applications Configuration Guide for SAP HANA":
https://launchpad.support.sap.com/#/notes/1944799

8.1   multipath

Here is the recommended configuration /etc/multipath.conf for IBM FlashSystem 5000 / 7000 /9000 family, including IBM SVC:

```
devices {
        device {
                vendor                  "IBM"
                product                 "2145"
                path_grouping_policy    group_by_prio
                prio                    "alua"
                path_checker            "tur"
                path_selector           "service-time 0"
                failback                "immediate"
                rr_weight               "priorities"
                no_path_retry           5
                rr_min_io_rq            32
                dev_loss_tmo            600
                fast_io_fail_tmo        5
        }
}
```

For using the SAP HANA fcClientLVM we recommend these settings to avoid hangs in the takeover:

- fast_io_fail_tmo = 5
  The fast_io_fail_tmo parameter sets the length of time to wait before failing I/O when a link problem is detected. I/O fails that want to reach the driver.
- path_selector: service time 0
  A service-time oriented load balancer that balances I/O on paths according to the latency. This optimizes the SAP HANA Log I/O which is latency sensitive (as stated by Novell).

Using alias for UUIDs in multipath.conf

The alias setting overrules the user_friendly_names setting.

Using alias makes managing and identifying volumes easier. If you want to use alias, please make sure that in the case of a multi-node environment, you fully understand the requirement and handling of consistent multipath device names across all nodes: The multipath.config must be identical on all nodes (LINUX systems) belonging to the HANA instance.

```
multipaths {
      multipath {
              wwid            36006048000028350131253594d303030
              alias           HANA_ANA_DATA_1
      }
      multipath {
              wwid            36006048000028350131253594d303041
              alias           HANA_ANA_LOG_1
}
```

8.2   IO device tuning

The following tunable parameters can be considered in case I/O performance issues have been detected. Do not change these in advance. Consult a performance specialist to find a vital combination:

- Increase /sys/block/<device>/queue/nr_requests if the default (128) results in blocked I/O submission. This will indirectly help to optimize the blocking inside SAP HANA.
- Increase 'rr_min_io_rq 32' in multipath.conf
- Increase the queue depth of devices, e.g.
  echo 64 > cat /sys/bus/scsi/devices/<device>/queue_depth

8.3   IO scheduler (kernel parameter elevator)

The default scheduler cfq delivers low IO performance for multiple (parallel) streaming read processes. The NOOP scheduler is recommended for setups with devices that do I/O scheduling themselves, such as intelligent storage or in multi-pathing environments. The DEADLINE scheduler can provide a superior throughput over the CFQ I/O scheduler in cases where several threads do reads and writes, and fairness is not an issue. For example, for several parallel readers from a SAN.

We recommend using the deadline scheduler for SAP HANA.

For more information, see SUSE System Analysis and Tuning Guide at
https://www.suse.com/documentation/sles12/

9 SAN zoning and volume mapping

Apply standard, best practices zoning within the SAN.

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

10 Setup of file systems for DATA and LOG

Create the DATA, LOG, and SHARED 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 ⏪☺ lvreate ⏪☺ mkfs.xfs
- Use telling 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:

DATA

1. Use 4 PV for Volume Group DATA
2. Create a volume group with an extent size of 1MB and an alignment size of 1MB
   vgcreate vgcreate -s 1M --dataalignment 1M VG_ANA_DATA_1 /dev/mapper/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 -b size=4096 -s size=4096 /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/data/ANA

LOG:

1. Use 4 PV for Volume Group LOG:
2. Create a volume group with an alignment size of 1MB
   vgcreate -s 1M --dataalignment 1M VG_ANA_LOG_1 /dev/mapper/WWID_1 ...
3. Create a logical volume with a stripe size of 256KB
   lvcreate -i 4 -I 256K -l 100%VG -n LV_ANA_LOG_1 VG_ANA_LOG_1
4. Create the file system
   mkfs.xfs -b size=4096 -s size=4096 /dev/mapper/VG_ANA_LOG_1-LV_LOG_1
5. Mount the file system
   mount /dev/mapper/VG_ANA_LOG_1-LV_ANA_LOG_1 /hana/log/ANA

SHARED (local, non-failover):

1. Use 1 PV for Volume Group SHARED
2. Create a volume group:
   vgcreate -s 1M --dataalignment 1M VG_SHARED /dev/mapper/WWID_1
3. Create a logical volume (without striping)
   lvcreate -l 100%VG -n LV_SHARED VG_SHARED
4. Create the file system
   mkfs.xfs -b /dev/mapper/VG_SHARED-LV_SHARED
5. Mount the file system
   mount /dev/mapper/VG_SHARED-LV_SHARED /hana/shared

If HANA scale-out or HA node fail-over is used, do not add the DATA & LOG file system to /etc/fstab, mounting will be done by SAP HANA.
Add the /hana/shared file system always to fstab.

11 VMware

The same setup applies like for LINUX native installation as described in chapter 10; the only deviation is that VMware™ ESX© server controls the SAN multipath, not the LINUX OS. The volumes are typically accessible under the path:

/dev/sd... instead of /dev/mapper/...

Please have a look in the IBM Redbook how to setup VMware VVols with IBM storage:

12 Setup of global.ini

A dedicated global.ini file (located in /hana/shared/...) is only needed if you want to use the node fail-over functionality provided by SAP HANA software and for HANA scale-out; else the HANA installer will create a global.ini file.

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

```
[communication]
listeninterface = .global

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

[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
```

13 fileio parameter

IBM storage systems do not require a fileio setting (SAP note 2399079) in case of HANA 2.0. – If you experience performance issues with the SAP tool hwcct then use these setting:

```
async_write_submit_active :  on
async_write_submit_blocks : all
async_read_submit         :  on
```

14 SAP HANA 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, IBM recommends to setup the storage environment as a SVC stretched cluster with symmetric VDISK mirroring.

15 SAP HANA Disaster Recovery setup with IBM System Storage

The disaster recovery capabilities of SAP HANA are documented on :

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

- SAP HANA – High Availability

If IBM Global Mirror (a-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.

16 SAP HANA backup with IBM Spectrum Protect

The product *IBM Spectrum Protect for Enterprise Resource Planning* includes support to backup SAP HANA online.

Please see current product information how to install, configure, and run this integrated solution:

17 IBM FlashCopy backup for SAP HANA

To minimize the backup and restore time IBM offers a solution that combines SAP HANA SNAPSHOT technology with IBM Spectrum Virtualize FlashCopy:

18 Change History

V1	Initial version
V1.2	Update of chapter 10 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	Changed sizing rules, added recommendation for multipath and VG/LV settings.
V2.5	Changed /etc/multipath setting
V2.6	Added fileio parameter for hdbparam
V2.7	Added A9000R & DS8880; info on LINUX LVM settings for HANA LOG
V2.8	Added chapter on Storage RtC, changes in storage sizing
V2.9	Added paragraph about VMware VVols

V2.10	Updated multipath section
V2.11	Update default disk type to RI SSD.
V2.12	Storwize: New sizing, 4 PV for LOG VG
V2.13	Add SVC 2145/2147 SV1
V2.14	Add FS9100, updated sections 7.7 and 7.11,
V2.15	Add V7000G3, minor update to multipath
V2.16	Add V5000E, adding FCM and DRP
V2.17	Minor changes in supported storage systems
V2.18	Minor changes in supported storage systems
V2.19	Adding new IBM FlashSystem models, update DRP section
V2.20	Updating max. numbers of supported systems for FlashSystem, changed DMMP config
V2.21	Minor change: update contact email address
V2.22	Adding NSE section

19 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:

IBM documentations:
This document is ID on

SAP HANA on IBM Power Systems and IBM System Storage – Guides

IBM FlashSystem product manuals:

Installation, configuration and usage of IBM Spectrum Protect for ERP (V8)

SAP documentation:
SAP HANA FAQ

SUSE documentation:

20 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

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

Spectrum Storage, System Storage, FlashSystem, Spectrum Protect, Spectrum Virtualize, Spectrum Accelerate, Spectrum Scale, GPFS, IBM, POWER, Redbooks.

The following terms are trademarks of other companies:

SAP, R/3, ABAP, SAP NetWeaver, 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.

21 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.