



STORAGE SIZING FOR SAP HANA TDI ON VMAX3

Sriram Chandrasekar

Advisory Systems Engineer Analyst

Dell

sriram.chandrasekar@dell.com

Table of contents

1 Purpose	3
2 Scope	3
3 Audience	3
4 Introduction	3
4.1 HANA Deployment Models	4
4.2 HANA terminology	4
4.3 SAP HANA TDI installation types	4
4.4 SAP HANA Deploy Environment	5
4.5 SAP HANA storage components	5
5 Sizing SAP HANA storage components	5
6 VMAX ³ storage recommendations for HANA	6
6.1 Drive configuration for HANA persistence	6
6.2 Drive configuration for HANA shared filesystem	7
7 Determining VMAX ³ model and number of engines	7
8 VMAX ³ FA-port requirements	8
9 VMAX ³ Storage Resource Pool (SRP) considerations	8
9.2 Multi SRP (mSRP) scenario	8
10 HANA IO profile	8
11 Determining VMAX ³ Cache	9
12 Sizing example	9
Summary	11
References	11

NOTE: This article was written prior to May 2016. Software and hardware updates are not incorporated into this article. As a result, some of these recommendations/configurations may no longer be valid.

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell EMC's views, processes or methodologies.

1 Purpose

The purpose of the paper is to describe how to build a firsthand solution for SAP HANA in a Tailored Datacenter Integration (TDI) environment on a VMAX³ storage array. The storage configuration recommendations are based on the HANA best practices for VMAX³ and performance that are required to meet SAP Key Performance indicators (KPIs) for a HANA TDI environment.

2 Scope

This paper provides a quicker approach to size a VMAX³ solution for HANA TDI environments and covers:

- Introduction to key components of SAP HANA
- Sizing HANA components
- Best practice considerations for VMAX³ sizing

This paper does not cover sizing requirements for Backup, Replication and HANA appliance model.

3 Audience

This paper will help Presales System Engineers and sales professional who do not have prior experience in SAP HANA to arrive at a firsthand VMAX³ solution quickly and accurately. This can also be a handy document for members of Dell EMC professional services to validate if best practices are followed for HANA TDI deployments on VMAX³ storage.

4 Introduction¹

SAP HANA is an in-memory data platform that is best suited for performing real-time analytics, and developing and deploying real-time applications (see Figure 1). There is an increasing need for the data to be recent, preferably in real-time and requires high speed and fast response time. SAP HANA is unique and can address this effectively. The data of HANA worker nodes is kept in RAM, where all the database operations take place. This means the operations are performed at memory speeds and the latency due to disk bottlenecks are completely avoided.

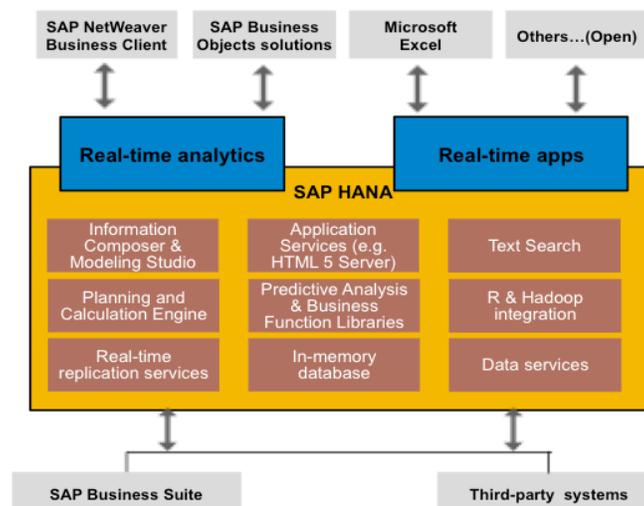


Figure 1: SAP HANA - platform for a new class of real-time analytics and applications¹

¹ Content and image extracted from SAP community network

4.1 HANA Deployment Models

SAP HANA can be deployed in 2 different models:

- Appliance Model
- Tailored Datacenter Integration (TDI) model

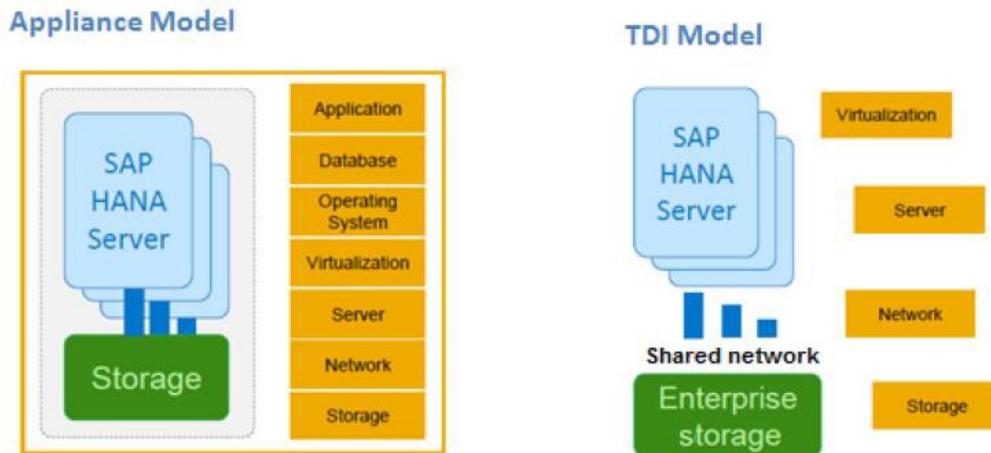


Figure 2: Appliance model v/s TDI model²

In this paper we will discuss only HANA TDI model sizing. TDI model will help customers make best use of their existing infrastructure and blend into their existing datacenter operations such as management, regular monitoring, data protection and DR. This will help achieve lower TCO and time to value.

4.2 HANA terminology

HANA worker node: Primary HANA host which processes the data

HANA standby node: A standby HANA host that will be used during a failover in case worker node goes down

For example, 3+1 node configuration means ---> 3 worker nodes and 1 standby node

4.3 SAP HANA TDI installation types

The HANA database can be installed as a stand-alone node (scale-up) or in cluster (scale-out).

In a **scale-up installation** (single-node), the RAM is local to the node and the database needs to fit into the RAM. This installation is good for random workloads such as online transaction processing (OLTP).

In a **scale-out installation** (multiple-nodes), database records are distributed across the RAM of multiple hosts. These deployments have high in-memory capacities and are best suited for online transaction processing (OLAP) type workloads.

² Image extracted from EMC whitepaper "STORAGE CONFIGURATION BEST PRACTICES FOR SAP HANA TDI ON EMC VMAX AND VMAX³ STORAGE SYSTEMS"

4.4 SAP HANA Deploy Environment

SAP HANA is deployed under Production and non-Production environments. Only production HANA nodes need to be compliant with SAP KPI's. The non-production HANA nodes do not have any performance requirements and therefore require no specific drive recommendations and need not be compliant with SAP KPI's.

So how do we differentiate between HANA production and non-Production nodes?

A SAP landscape can consist of multiple installations of the same application. For example: a production installation, a pre-production installation, a test installation, a DR installation, a QA installation. It is hard to define which installation is treated as production and which as non-production. Only the customer can make this decision.

For example, while test and QA may not be that important, a HANA DR application or a pre-production installation may be treated as a production. You need to know how many HANA nodes are treated as production.

4.5 SAP HANA storage components

Following are the key storage components for any HANA deployments. They have to be carefully planned for adequate storage resources.

4.5.1 HANA Persistence (Data and Log)

Data represents the anticipated total memory requirement. HANA persists a copy of the in-memory data in the form of savepoint blocks and writes separately to its own savepoint files every 5 minutes. (/hana/data)

Log is maintained to ensure the recovery of database with zero data loss in case of faults (/hana/log).

4.5.2 HANA shared filesystems

HANA installation binaries, configuration files and trace files are stored under /hana/shared and made available to all nodes.

5 Sizing SAP HANA storage components

5.1 Sizing for HANA Persistence

The following details are required to determine the usable storage capacity of HANA persistence.

- Number of HANA worker nodes (**N**)
- RAM size of each HANA worker node(**R**)

Please note, for storage capacity calculations only the worker nodes are considered. Standby nodes are used only to achieve High Availability in HANA deployments.

Total capacity required for persistence = Total capacity for Data + Total capacity for Log

Total capacity for Data

$$\text{➤ } (N) * (R) * 1.2(\text{growth factor})^3$$

³ Calculations are based on SAP whitepaper "SAP HANA Storage Requirements"

Total capacity for Log

- For Systems $\leq 512\text{GB} = (\frac{1}{2} * R) * (N)^3$
- For system $> 512\text{GB} = R * (N)^3$

For example, assume the requirement is 4+1 node with 1TB RAM each

Total capacity for Data = $4 * 1 * 1.2 = 4.8\text{TB}$

Total capacity for Log = $512\text{GB} * 6 = 2\text{TB}$

Total capacity required for HANA persistence = 6.8TB

5.2 Sizing for HANA shared filesystems

5.2.1 For scale-up systems

For scale-up systems, this can be local to the node and can be a block device.

HANA shared filesystem size for single-node = **MIN (1 * R; 1TB)**³

For example,

- Size of HANA shared filesystem for a Single-node with 512GB RAM = 512GB
- Size of HANA shared filesystem for a Single-node with 2TB RAM = 1TB

5.2.2 For scale-out systems

For scale-out systems, this is shared across all the nodes in the cluster.

HANA shared filesystem size for scale-out node = **1 * RAM per worker node for every 4 worker nodes**³

For example,

- Size of HANA shared filesystem for 4+1 with 1TB per node = $1 * 1\text{TB} = 1\text{TB}$
 - Size of HANA shared filesystem for 6+1 with 1TB per node = $2 * 1\text{TB} = 2\text{TB}$
- 6 VMAX³ storage recommendations for HANA

6.1 Drive configuration for HANA persistence

The drive recommendation for HANA persistence should satisfy both the performance required to meet SAP KPI's and the capacity requirements.

6.1.1 Drives required to meet Performance

The performance tests run to certify VMAX³ for HANA showed that 12 x 10K drives per node in RAID 1 configuration meet the SAP key performance indicators (KPIs). This is on the basis that each HANA node is expected to produce 1200 IOPS on average and, assuming a single 10K drive can deliver 100 IOPS on average, 12 x 10K drives are required to support the performance required for a single HANA node.

For example, a 6+1 node configuration requires $6 * 12 = (72)$ 10K drives to meet the SAP KPI's for performance.

³ Calculations are based on SAP whitepaper "SAP HANA Storage Requirements"

6.1.2 Drives required to meet Capacity

However, we may need additional drives to meet the HANA persistence capacity requirements. You can choose between 300GB, 600GB or 1.2TB drives based on the capacity required.

For example, a 2+1 node configuration with 1TB RAM each requires 4.5TB for HANA persistence. Here 24 x 10K drives would be sufficient to meet the performance requirements but we need to choose the right drive size to meet the HANA persistence capacity. In this case, using 24 x 300GB 10K in RAID 1 will not meet the HANA persistence capacity requirements, so 24 x 600GB 10K drives in RAID 1 is required

6.2 Drive configuration for HANA shared filesystem

The capacity for the “shared” filesystem is always required. If you have multiple scale-up installations, each node will need a filesystem called /hana/shared where it stores the binaries. This filesystem is just local to the node and resides on a block filesystem (XFS, ext3).

In a scale-out installation, the /hana/shared must be a single filesystem shared across multiple nodes. The best approach is to use NFS (NAS) but some other vendors use cluster filesystems like OCFS2 for this purpose.

In VMAX³, eNAS is used to present the mount points for HANA shared libraries. There is no performance requirement for HANA shared libraries and therefore no specific drive recommendations. They can co-exist along with the drives recommended for HANA persistence.

Note that some customers may already have the infrastructure setup for NAS. i.e. the NAS capacity can be presented from a different storage array or they can use NFS server exports using the block storage in VMAX³. Therefore, it is important to check with customers on how they want to approach this because introducing eNAS in the VMAX³ solution is going to incur additional cost.

7 Determining VMAX³ model and number of engines

VMAX³ model and the number of engines required is driven by the number of HANA production nodes. Table below shows the VMAX3 models and the estimated maximum number of HANA nodes that can be connected according to the number of available number of engines.

VMAX3 model	Engines	Maximum HANA nodes
100K	1	12
	2	20
200K	1	16
	2	28
	3	40
	4	52
400K	1	20
	2	32
	3	44
	4	56
	5	68
	6	80
	7	92
	8	104

Figure 3: VMAX³ scalability for HANA nodes⁴

For example, a HANA 14+1 configuration would require minimum 100K 2 engines or 200K 1 engine.

⁴ Image extracted from EMC whitepaper “Storage configuration best practices for HANA TDI on VMAX³ storage arrays”

8 VMAX³ FA-port requirements

HANA requires dedicated FA-ports to achieve full I/O performance. The FA-ports have to be distributed across all available engines and balanced between the directors.

see Figure 4 below to determine the required number of FA-ports

HANA worker nodes	Required FA-ports
1-4	2
5-8	4
9-16	8
17-20	12

Figure 4: VMAX³ FA-ports for HANA nodes⁴

For example, a 8 node HANA on a 2-engine VMAX³ should be connected to 4 ports on each engine (2 per director) to have balanced activity across engines.

9 VMAX³ Storage Resource Pool (SRP) considerations

9.1 Single SRP scenario

If the VMAX³ array is dedicated only to SAP HANA then the solution can be modeled as a Single SRP. This solution can be run through VMAX³ Sizer and the Sizer XML can be imported into “My Quotes” to generate the quote instantly.

9.2 Multi SRP (mSRP) scenario

Arrays holding HANA plus other workloads call for an mSRP solution. For example, consider a scenario where the customer requests a VMAX³ sizing for HANA production, HANA non-production and non-HANA (general workloads). The SRPs should be configured as:

- SRP 1 for HANA Production workload
- SRP2 for HANA non-Production and non-HANA (general) workloads

Note that the non-Production nodes and non-HANA (general) workloads can coexist in the same SRP because the non-Production nodes need not be compliant with HANA best practices and KPI's as discussed earlier.

Also note that a separate SRP is not required for the eNAS (HANA binaries) because they do not have any performance requirements. They could potentially coexist with eNAS for non-HANA workloads in the same SRP.

The mSRP configurations cannot be built using VMAX Sizer and they have to be quoted in “My Quotes” manually.

10 HANA I/O profile

HANA does not use a fixed block size. Depending on the device type (data or log), HANA uses block sizes from small 4K up to large 1M blocks. It all depends on the application on top of HANA and the data structure. But, an 80% write and 20% read ratio with the largest possible block size (128K) would be ideal to use.

⁴ Image extracted from EMC whitepaper “Storage configuration best practices for SAP HANA TDI on EMC VMAX3 storage arrays”

11 Determining VMAX³ Cache

Please make sure to run the VMAX³ cache calculator @ <https://hubv1.corp.emc.com/services/cachecalc> to ensure the VMAX³ system has enough cache to support the capacity requirements.

12 Sizing example

Customer **XYZ** wanted to size a VMAX³ for the following HANA requirements:

- *No of worker nodes: 16*
- *RAM per node: 512GB*
- *All the nodes will be in cluster*
- *All the nodes will be run as Production*
- *New VMAX³ will be dedicated to the HANA*
- *Include eNAS in the solution for HANA shared filesystem*

HANA component sizing

$$\begin{aligned}\text{Total capacity for Data} &= (\text{Number of worker nodes} * \text{RAM per node}) * 1.2 \\ &= 16 * 512 * 1.2 \\ &= 9,830\text{GB}\end{aligned}$$

$$\begin{aligned}\text{Total capacity for Log (for systems } \leq 512\text{GB)} &= (\text{RAM per node} * \frac{1}{2}) * \text{Number of worker nodes} \\ &= (512 * \frac{1}{2}) * 16 \\ &= 4,096\text{GB}\end{aligned}$$

$$\begin{aligned}\text{Total capacity required for persistence} &= \text{Total capacity for Data} + \text{Total capacity for Log} \\ &= 9830 + 4096 \\ &= 13,926\text{GB}\end{aligned}$$

Note: The customer requirement states “All the nodes are in cluster (which implies scale-out nodes)
HANA shared filesystem size for scale-out node = 1 * RAM per worker node for every 4 worker nodes

$$\begin{aligned}&= 1 * 512 * 4 \\ &= 2,048\text{GB}\end{aligned}$$

VMAX³ Storage configuration

Drive recommendation for HANA persistence Performance

$$\begin{aligned}\text{Number of 10K drives required to meet SAP KPI's for Performance} &= 12 * \text{Number of worker nodes} \\ &= 12 * 16 \\ &= 192 \text{ drives}\end{aligned}$$

Drive recommendation for HANA persistence capacity

300GB, 600GB and 1.2TB drives are the available choices.

In this case using 300GB drives will be more than sufficient to meet the HANA persistence capacity (13,926GB). You will notice that 192 x 300GB 10K drives in RAID 1 results in capacity of around 25,700GB but the required HANA persistence capacity is only 13,296GB. The additional capacity is required to meet the HANA KPI's for performance.

Recommendation= 192 x 300GB 10K drives in RAID 1

eNAS for HANA shared filesystem

Because the customer wanted to include eNAS in the solution, we would configure 2,048GB for eNAS with 2 Data Movers (one primary and another standby) and required GigE ports. The excess capacity available in the persistence can be used to present the eNAS capacity.

Recommendation: 2 Data Movers with required GigE ports

Determining VMAX³ Model and # of engines

As per the “VMAX³ scalability for HANA nodes table” 16 HANA nodes would require minimum VMAX 100K 2-engine or VMAX 200K 1-engine.

The better choice would be to pick the VMAX 200K 2-engine for the following reasons:

- The VMAX 100K 2-engine can support only a maximum of 16 nodes and if the customer wanted to add more nodes in future they cannot scale further with a 100K
- The VMAX 200K 1-engine is already at its maximum with 16 nodes, leaving no room for further growth.

Recommendation: VMAX 200K 2 engines

Determining VMAX³ FA-Ports

As per the “VMAX³ FA-ports for HANA nodes” table 16 HANA nodes would require 8 dedicated FA-ports. In order to have a balanced activity across engines, use 8 ports on each engine (4 per director).

Note that these ports need to be reserved for HANA only to achieve full I/O performance of HANA nodes

Recommendation: 8 dedicated FA-ports

VMAX³ Storage Resource pool

Customer has clearly expressed the intent that the VMAX³ will be dedicated to HANA environment, so isolation of workloads is not required and a single storage resource pool is recommended.

Recommendation: Single Storage Resource Pool (SRP)

VMAX³ Cache

The VMAX³ cache calculator shows “400GB” for the actual cache required. Because this is a 2 engine configuration and 512GB is minimal cache slot available per engine, the total cache required for this solution will be 1TB (512GB per engine)

Recommendation: 1TB cache (512GB per engine)

VMAX³ solution summary

- VMAX 200K 2-engines
- 1TB cache
- eNAS – 2 Data Movers
- 192 x 300GB 10K drives
- 8 dedicated FA-ports
- Single storage resource pool

Summary

The approach discussed in this paper will help presales system engineers/sales to effectively position a VMAX³ solution for any SAP HANA TDI requirement. The procedure we have described are excerpts gathered from SAP/EMC best practice publications on HANA and conversations with the SAP specialist group. This will enable system engineers who do not have prior knowledge on SAP HANA to build a first-hand HANA solution with clear recommendations.

References

1. SAP HANA Storage requirements-
https://hcp.sap.com/content/dam/website/saphana/en_us/Technology%20Documents/Storage%20Whitepaper%202%204.pdf
2. Storage configuration best practice for SAP HANA TDI on VMAX and VMAX³ systems-
<https://www.emc.com/collateral/white-paper/h12342-storage-config-best-practices-sap-hana-tdi-vmax-vmax3-wp.pdf>

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL EMC 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 Dell EMC software described in this publication requires an applicable software license.

Dell, EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries.