

NEC iStorage M series for SAP HANA Tailored Datacenter Integration

(iStorage M110/M310/M510/M710/M5000)

Configuration and Best Practice Guide

Copyright © 2015 NEC Corporation. All Rights Reserved.

The information in this document is for informational purposes only and may contain typographical errors and technical inaccuracies. This document is provided as is and NEC Corporation makes no representations or warranties of any kind.

Product, brands, and trade names used in this document are trademark or registered trademarks of their respective holders.

Table of contents

| | |
|---|----|
| Introduction | 4 |
| Overview of NEC iStorage..... | 5 |
| Selection of NEC iStorage M series configuration..... | 6 |
| Selection of NEC iStorage M series model..... | 6 |
| Connection to the host server | 6 |
| Disk drives..... | 7 |
| RAID level | 7 |
| Cache memory | 7 |
| Estimation of required disk size | 7 |
| SAP HANA data | 7 |
| SAP HANA log..... | 8 |
| OS, /usr/sap..... | 8 |
| SAP HANA shared..... | 8 |
| Backup area..... | 9 |
| Example of disk size calculation..... | 9 |
| Estimation of number of disk drives | 9 |
| Installation and Configuration..... | 11 |
| Access Control | 11 |
| Virtual Cache Partitioning | 11 |
| Path redundancy | 11 |
| File system | 12 |
| SAP HANA Storage Connector API | 12 |
| Performance check..... | 13 |
| Conclusion | 14 |
| References | 15 |

Introduction

This document is the configuration and best practice guide to use NEC iStorage M series with SAP HANA Tailored Datacenter Integration (TDI).

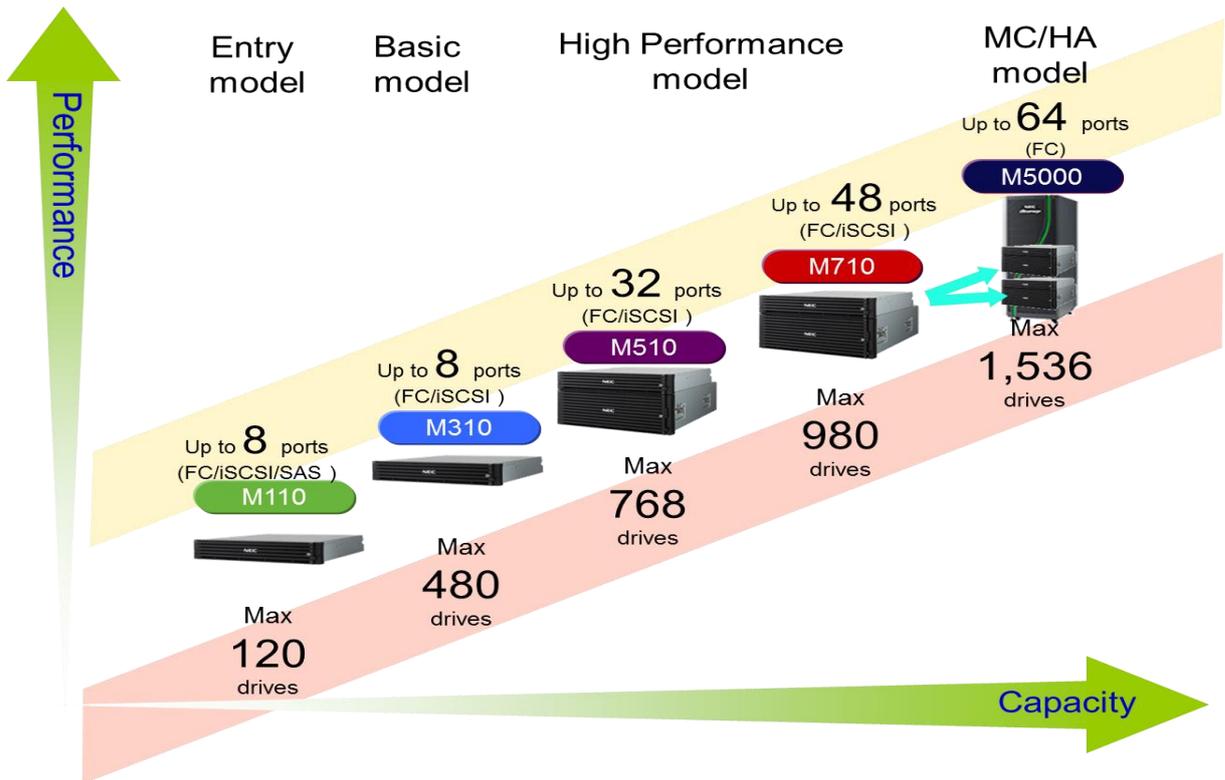
SAP HANA was initially provided as an appliance server and storage for SAP HANA DB was included in an appliance configuration. With TDI SAP offers an additional approach to use SAP certified Enterprise Storage as SAP HANA storage, providing a more flexible system.

This guide is applied for productive SAP HANA systems. For non-productive SAP HANA systems, performance for storage system is not required, so you can choose any disk drive configuration.

Overview of NEC iStorage

NEC iStorage M series is a series of optimum disk array systems for open systems and suitable for use in SAN environments. NEC iStorage M series address the need for high performance and high availability, advanced eco-friendly performance, easy installation and operation, and economic efficiency,

NEC iStorage M series offers a diverse product lineup consisting of entry model M110 Disk array, basic model M310 Disk array, high performance model M510 Disk array and M710 Disk array, and mission critical and high availability model M5000.



M5000: Only a Japanese market.

Selection of NEC iStorage M series configuration

To select and configure NEC iStorage M series for SAP HANA TDI you need to follow the guide described here.

Selection of NEC iStorage M series model

You can use only SAP certified enterprise storage. You can find certified storages at:

<http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html>

At the time of this document is written, M110, M310, M510, M710, and M5000 are certified. M110 is dual controller model only.

You need to decide which model of iStorage you use from following points.

- Number of SAP HANA and non-HANA systems connected to.

There are upper limits of supported SAP HANA nodes for each model. These limits are decided to keep good performance for SAP HANA.

If non-HANA DB systems share the storage, those systems are also counted.

You must take care that the workload of non-HANA system do not affect the workload of SAP HANA, for example, by separating disk pool and FC ports for SAP HANA and non-HANA systems.
- Number of disk drives

Number of disk drives for one HANA system is determined from the memory size of HANA system.

From a performance point of view, minimum number of disk drive is limited.

The table below shows the supported maximum number of HANA nodes for each iStorage model.

| Model | Max. HANA nodes |
|-------|-----------------|
| M110 | 3 |
| M310 | 8 |
| M510 | 16 |
| M710 | 32 |
| M5000 | 40 |

If a non-HANA DB system is connected, it must be counted as one HANA node.

Connection to the host server

iStorage has Fibre channel(FC), iSCSI and SAS interfaces to connect to the host server, but for TDI you can use only FC interface. iSCSI and SAS are not allowed.

Server and storage must be connected with 2 or more FC ports for high availability.

You can use a FC switch, but it is recommended to connect storage and server directly to separate work load of other HANA system and non-HANA system.

Disk drives

The types of disk drives you can use for HANA data and log volume are

- ✧ SAS HDD/10000rpm
- ✧ SAS HDD/15000rpm
- ✧ SSD

You can select any size of disk drive from supported drives on each iStorage model

RAID level

You must configure disk drives with redundancy. RAID5/50 or RAID6/60 is recommended. You can use RAID10, too.

Cache memory

It is recommended to use maximum available cache for each model.

Estimation of required disk size

In SAP HANA systems, disk spaces are required for:

- ✧ OS, swap, /usr/sap
- ✧ SAP HANA shared
- ✧ SAP HANA data
- ✧ SAP HANA log
- ✧ backup

Among these, performance of SAP HANA data and SAP HANA log areas are very important and you must follow the guide of this document.

You may need additional disk space. For example, if an ABAP instance is also installed on the SAP HANA server, more space will be required.

The size of these disk areas are calculated from the memory size (RAM) of SAP HANA. The value used for calculation is the memory size used by SAP HANA, not physical memory size of SAP HANA server. For example, your sizing requires 1.3TB memory for SAP HANA, and you want to deploy SAP HANA on 2TB memory server, the disk size should be calculated by using 1.3TB.

SAP HANA data

The size of data volume is

1.2 x net disk space for data(including additional free space of 20%) or

1 x RAM.

net disk space for data is equivalent to the value calculated using the SAP Quick Sizer or one of the sizing reports(SAP note 1793345, 1872170, 1736976).

If SAP HANA requires more data area, you can add disks later.

Note: SAP HANA appliance has 3 x RAM for SAP HANA data.

SAP HANA log

The size of the log volume depends on the RAM.

RAM is 512GB or less: 1/2 of RAM

RAM is more than 512GB: minimum 512GB

OS, /usr/sap

These are local to each server and can be put on local disks of the server or storage.

The minimum requirements for each area are:

OS installation: 10GB

/usr/sap: 50GB

SAP HANA shared

This is mounted on /hana/shared/<SID> and SAP HANA is installed here. In a single node system, this area can be put on a server local disk or on the storage. In a scale out system, this area is accessed from all HANA nodes and implemented using a shared file system such as OCFS2 or NFSv4.

For single-node SAP HANA systems, the recommended disk space for /hana/shared/<SID> is:

RAM is 1024GB or less: size of shared is 1 x RAM

RAM is 1024GB or more: size of shared is 1024GB

For scale-out SAP HANA systems, the recommended disk space for /hana/shared/<SID> depends on the number of worker nodes. Per each four worker nodes of a given scale-out system, a disk space of 1x RAM of one worker is recommended:

size of shared is 1 x RAM of worker per 4 worker nodes

Examples:

1 ~ 4 x 512GB worker node: size of /hana/shared/<SID> is 1 x 512GB = 512GB

2 ~ 8 x 512GB worker node: size of /hana/shared/<SID> is 2 x 512GB = 1024GB

9 ~ 12 x 1024GB worker node: size of /hana/shared/<SID> is 3 x 1024GB = 3072GB

13 ~ 16 x 1024GB worker node: size of /hana/shared/<SID> is 4 x 1024GB = 4098GB

Backup area

This is a disk space to keep backup files of HANA DB. Prepare necessary space according to your backup policy. It is recommended to use another disk pool from HANA data and log disk pool.

Example of disk size calculation

Examples of disk space calculation are shown below.

RAM=256GB:

OS:10GB + swap:50GB + /usr/sap:50GB + HANA data(1xRAM):256GB + HANA log:128GB + HANA shared:256GB = 750GB

RAM=1.3TB with additional disk space:

OS,swap,/usr/sap:200GB + HANA data(3xRAM):3994GB + HANA log:512GB + HANA shared:1024GB = 5730GB

Estimation of number of disk drives

From performance requirements, use minimum of 10 HDDs (RAID5/50 or RAID6/60) for one SAP HANA node (worker node).

When three SAP HANA nodes use one storage, minimum $3 \times 10 = 30$ HDDs are required.

Examples of number of disks are shown below.

RAM=256GB:

750GB of disk space is required from the above calculation. 10 x 600GB HDD provides enough disk space and use this as one disk pool of RAID5/50(4+P).

RAM=1.3TB with additional disk space::

5730GB of disk space is required from above calculation. 14 x 600GB HDD are necessary for this size. You can use 14 HDDs as one disk pool of RAID5/50(4+P), but 15 x 600GB HDD is recommended because 15 is multiple of 5(4+P).

In case of 1.2TB HDD, less than 10 disks of RAID5/50(4+P) is enough, but you should use 10 x 1.2TB HDDs as one disk pool.

When HANA with RAM=256GB and HANA with RAM=1.3TB shares one storage:

Prepare one disk pool of 10 x 600GB HDD for 256GB HANA and one disk pool of 15 x 600GB for 1.3TB HANA.

Instead you can use one disk pool of 20(=2 node x 10 HDD) or more disks as one disk pool for both HANA. Total 6480GB (750GB+5730GB) is required for two HANA system and 20 x 600GB HDD (RAID5/50(4+P)) is enough for this size.

When multiple HANA nodes share one disk pool, the number of HANA nodes for one disk pool should not exceed two. Also HANA and non-HANA system must not share one disk pool.

Installation and Configuration

Please consider the following when you install SAP HANA on server and storage.

SAP HANA installation in TDI must be done by who is certified as “SAP Certified Technology Specialist – SAP HANA Installation” (E_HANAINS141).

Access Control

The iStorage AccessControl is a function of iStorage. It associates logical disks on storage and server. The logical disks are only visible from associated servers.

Virtual Cache Partitioning

iStorage VirtualCachePartitioning virtually divides primary cache of disk arrays to multiple areas and assigns each divided area to a virtual machine or business application. This ensures stable performance by restricting the cache memory used by each virtual machine or business application and excluding intervention.

Path redundancy

The storage and server must be connected with multiple FC paths and setup of path redundancy is required.

You can use the native Linux multipathd functionality. The settings of multipathd are in the file /etc/multipath.conf. This is an example used in a RHEL 6.5 OS environment.

```
defaults {
#   polling_interval    5
#   path_selector       "service-time 0"
#   path_grouping_policy  failover
#   prio                alua
#   path_checker        directio
#   rr_min_io           1000
#   rr_weight            uniform
#   failback            immediate
#   no_path_retry       queue
#   user_friendly_names yes
}
blacklist {
    device {
        vendor "*"
        product "*"
    }
}
blacklist_exceptions {
    device {
        vendor "NEC"
        product "*"
    }
}
```

```

}
multipaths {
    multipath {
        wwid    258c2323c1e740012
        alias   hana_log1
    }
    multipath {
        wwid    258c2323c1e740013
        alias   hana_data1
    }
}

```

Note: value of wwid is specific to each logical disk.

iStorage StoragePathSavior is the software product to enable path redundancy and workload balancing. You can also use this instead of native Linux multipathd. When you use SAP HANA Storage Connector API described later, then you must use the native Linux multipathd.

File system

The xfs file system must be used for SAP HANA data and log volume.

You can format disk by mkfs.xfs command.

```

mkfs.xfs /dev/mapper/hana_log1
mkfs.xfs /dev/mapper/hana_data1

```

SAP HANA Storage Connector API

In a SAP HANA scale out system, which has more than one worker and standby nodes, SAP HANA data and log disks are mounted and unmounted by the SAP HANA storage connector API for Fibre Channel (fcClient). While the disk is mounted, a SCSI-3 Persistent Reservation is written to the device and it ensures that other nodes do not access those disks (I/O fencing).

The storage connector API is controlled in the storage section of SAP HANA global.ini file. This is an example of it.

```

[storage]
ha_provider = hdb_ha.fcClient
partition_*_*__prtype = 5
partition_*_data__mountoptions = -o inode64,largeio,swalloc
partition_*_log__mountoptions = -o inode64,largeio,swalloc
partition_1_data__alias = hana_data1
partition_1_log__alias = hana_log1
partition_2_data__alias = hana_data2
partition_2_log__alias = hana_log2

```

When you install SAP HANA, create a global.ini file and start the SAP HANA installation with --storage-cfg=/some/path option. /some/path is a directory where global.ini file exists.

Performance check

After you completed the installation of a productive system, you must check that performance of HANA data and log disks fulfills TDI KPIs defined by SAP. This is done by using HWCCT (Hardware Configuration Check tool). The official tool documentation is contained in the SAP HANA Administration Guide. The documentation contains the latest KPIs for data throughput and latency. For latest changes of the documentation, see the PDF attachment of SAP Note 1943937.

Additionally, SAP Active Global Support offers the HANA Going-Live Check which – among other tests –conducts a data throughput test using the very same “SAP HANA Hardware Configuration Check Tool”.

Conclusion

NEC iStorage M series are developed on over a half century history of NEC storage and fifteen years history of iStorage. Using SAP HANA with NEC iStorage M series provides best solution for your business.

References

- NEC Storage(Global)
<http://www.nec.com/en/global/prod/storage/>
- NEC iStorage(Japan)
<http://jpn.nec.com/istorage/>
- SAP HANA Tailored Data Center Integration FAQ
<https://scn.sap.com/docs/DOC-62942>
- SAP HANA TDI - Storage Requirements
<http://scn.sap.com/docs/DOC-62595>
- Certified and Supported SAP HANA Hardware Directory
<http://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html>
- SAP HANA Administration Guide
http://help.sap.com/hana/SAP_HANA_Administration_Guide_en.pdf
2.8 Hardware Checks for Tailored Data Center Integration
- SAP note 1900823: SAP HANA Storage Connector API
<https://service.sap.com/sap/support/notes/1900823>
- SAP note 1943937: Hardware Configuration Check Tool - Central Note
<https://service.sap.com/sap/support/notes/1943937>
- Sizing Approaches for SAP HANA
<https://websmp210.sap-ag.de/~sapidb/011000358700000050632013E>
- SAP Quick Sizer tool
<http://service.sap.com/quicksizing>
- SAP Note 1793345: Sizing for SAP Suite on HANA
<https://service.sap.com/sap/support/notes/1793345>
- SAP Note 1872170: Suite on HANA memory sizing report
<https://service.sap.com/sap/support/notes/1872170>
- SAP Note 1736976: Sizing Report for BW-on-HANA
<https://service.sap.com/sap/support/notes/1736976>