Inspur solution for SAP HANA installation guide

Contents

1. Solution overview

The Inspur solution for SAP HANA with the NF8260M5 servers has been certified by SAP. The SAP HANA platform based on NF8260M5 provides a scalable database with advanced analytical capabilities and intuitive application-development tools in an in-memory data platform. SAP HANA supports Intel® Optane™ Data Center Persistent Memory Module (PMEM). Persistent memory modules are nonvolatile memory modules that bring together the low latency of memory and the persistence of storage. Persistent memory modules provide faster access to data and retain data across power cycles, based on the mode.

Intel® Optane™ DC persistent memory represents an entirely new means of managing data for demanding workloads such as the SAP HANA platform. Intel® Optane™ DC persistent memory is nonvolatile, so in-memory databases such as SAP HANA do not have to completely reload all data from persistent storage to memory, and it runs at near-DRAM speeds, maintaining today's performance expectations. It also delivers greater data density than memory technologies, which enables additional innovation and simpler IT landscapes. With its persistence, performance, and lower cost per gigabyte than conventional memory, Intel® Optane™ DC persistent memory can help reduce total cost of ownership (TCO), reshape the way that businesses tier their data for database systems, and open new use cases for the speed and power of the SAP HANA platform.

Memory for databases is currently small, expensive, and volatile. Intel® Optane™ DC persistent memory is denser, more affordable, and persistent, and it performs at speeds close to that of memory. These features of Intel® Optane™ DC persistent memory can help lower TCO through reduced downtime and simplified data-tiering operations. These same features can also make SAP HANA in-memory databases economically viable for a wider range of use cases. Intel® Optane™ DC persistent memory provides near-DRAM in-memory computing speed in a form factor similar to that of dual in-line memory modules (DIMMs) at a lower price per gigabyte than DRAM. Support for Intel® Optane™ DC persistent memory is available with the next-generation Intel® Xeon® processor Scalable family.

Because it is nonvolatile, Intel® Optane™ DC persistent memory enables you to keep the data in the SAP HANA platform loaded in main memory, even when power is off. Because you don't have to reload the data back into memory after downtime, restart time for the SAP HANA platform is greatly reduced. Intel® Optane™ DC persistent memory is available in a form factor called persistent memory modules, which are similar to DIMMs, but with greater capacity than is available with conventional, volatile memory.

This document describes the steps required to deploy and configure a Inspur solution for SAP HANA with Intel® Optane™ Data Center Persistent Memory Module (PMEM).

2. Solution design

This section describes the SAP HANA system requirements defined by SAP and the architecture of the Inspur solution for SAP HANA.

2.1. SAP HANA system

SAP HANA scale-up system on a single server is the simplest of the SAP HANA installation types. You can run an SAP HANA system entirely on one host and then scale the system up as needed. All data and processes are located on the same server and can be accessed locally. For this option the network must have at least one 1 Gigabit Ethernet access network.

Intel® Optane™ PMEM is supported on servers equipped with second-generation Intel Xeon Gold processors and Intel Xeon Platinum processors. Two primary modes are supported: App Direct mode, including Block over App Direct mode, and Memory mode. App Direct mode is the only mode that is currently supported by SAP HANA 2.0 SPS 03+. In App Direct mode, the persistent memory modules appear as byte-addressable memory resources that are controlled by SAP HANA 2.0 SPS 03+. In this mode, the persistent memory space is controlled directly by SAP HANA.

2.2. Ratio of PMEMs to DIMMs

Intel® Optane™ PMEMs must be installed with DRAM DIMMs in the same system. The persistent memory modules will not function without any DRAM DIMMs installed. In two-, four-, and eight-socket configurations, each socket contains two IMCs. Each memory controller is connected to three double data rate (DDR) memory channels that are then connected to two physical DIMM persistent memory slots. In this configuration, a maximum of 12 memory slots per CPU socket can be configured with a combination of Intel® Optane™ PMEMs and DRAM DIMMs.

SAP HANA 2.0 SPS 03 currently supports various capacity ratios between Intel Optane PMEMs and DIMMs. Ratio examples include the following:

- 1:1 ratio: A single 128-GB Intel Optane PMEM is matched with a single 128-GB DDR4 DIMM.
- 2:1 ratio: A 128-GB Intel Optane PMEM is matched with a 64-GB DDR4

DIMM.

- 4:1 ratio: A 256-GB Intel Optane PMEM is matched with a 64-GB DRAM DIMM.

Different-sized Intel Optane PMEMs and DIMMs can be used together as long as supported ratios are maintained as following table.

Memory configuration (PMEM+ DRAM)	CPU type	Capacity (GB)	Ratio of PMEMs to DIMMs
128GB PMEMM + 64GB DRAM	M or L	4608	2:1
128GB PMEMM + 128GB DRAM	M	6144	1:1
256GB PMEMM + 64GB DRAM	M	7680	4:1

2.3. SAP HANA Solution Specification

NF8260M5 specification-1:

PLATFORM		Inspur NF8260M5
PROCESSOR		4 * Intel® Xeon® Platinum 8276M or 8280M or 8276L or 8280L processors
MEMORY	**Total**	4608GB (DRAM/PEME Ratio 1:2)
	PMEM	24 * 128GB
	DRAM	24 * 64GB
BOOT drive		2* Intel® SSD D3-S4510 Series (480 GB, 2.5 in SATA 6 Gb/s, 3D2, TLC) at 480 GB or larger
HANA data		3* Intel® SSD DC P4610 Series (3.2TB, 2.5in PCIe 3.1 x4, 3D2, TLC) at 3.2 TB or larger
SOFTWARE		OS(minimum or higher): SLES 12 SP4 for SAP
		SAP HANA 2.0 SPS04 or later
		Software optimizations as described in SAP HANA manuals
Network		1 * Integrated 1 GbE (Management network) optional: 2* 10 Gbit Intel® C620 Series Chipset with integrated Intel® Ethernet Network Adapter X722 (e.g. used in Scale-Out or HANA system replication)

Table 1: Detailed configuration-1 of NF8260M5.

NF8260M5 specification-2:

PLATFORM		Inspur NF8260M5
PROCESSOR		4 * Intel® Xeon® Platinum 8276M or 8280M or 8276L or 8280L processors
MEMORY	**Total**	6144GB (DRAM/PEME Ratio 1:1)
	PMEM	24 * 128GB
	DRAM	24 * 128GB
BOOT drive		2* Intel® SSD D3-S4510 Series (480 GB, 2.5 in SATA 6 Gb/s, 3D2, TLC) at 480 GB or larger

HANA data	4* Intel® SSD DC P4610 Series (3.2TB, 2.5in PCIe 3.1 x4, 3D2, TLC) at 3.2 TB or larger
SOFTWARE	OS(minimum or higher): SLES 12 SP4 for SAP
	SAP HANA 2.0 SPS04 or later
	Software optimizations as described in SAP HANA manuals
Network	1 * Integrated 1 GbE (Management network) optional: 2* 10 Gbit Intel® C620 Series Chipset with integrated Intel® Ethernet Network Adapter X722 (e.g. used in Scale-Out or HANA system replication)

Table 2: Detailed configuration-2 of NF8260M5.

NF8260M5 specification-3:

PLATFORM		Inspur NF8260M5
PROCESSOR		4 * Intel® Xeon® Platinum 8276M or 8280M or 8276L or 8280L processors
MEMORY	**Total**	7680GB (DRAM/PEME Ratio 1:4)
	PMEM	24 * 256GB
	DRAM	24 * 64GB
BOOT drive		2* Intel® SSD D3-S4510 Series (480 GB, 2.5 in SATA 6 Gb/s, 3D2, TLC) at 480 GB or larger
HANA data		5* Intel® SSD DC P4610 Series (3.2TB, 2.5in PCIe 3.1 x4, 3D2, TLC) at 3.2 TB or larger
SOFTWARE		OS(minimum or higher): SLES 12 SP4 for SAP
		SAP HANA 2.0 SPS04 or later
		Software optimizations as described in SAP HANA manuals
Network		1 * Integrated 1 GbE (Management network) optional: 2* 10 Gbit Intel® C620 Series Chipset with integrated Intel® Ethernet Network Adapter X722 (e.g. used in Scale-Out or HANA system replication)

Table 3: Detailed configuration-3 of NF8260M5.

2.4. Sizing persistent storage

The storage size for the file system is based on the amount of memory (DRAM + Intel Optane PMEM) on the SAP HANA host. For a single-node system with 4608GB of memory (1.5TB DRAM + 3TB Intel Optane PMEM), the recommended file system sizes are as follows:

- /hana/data >= 1 x memory (DRAM + Intel Optane PMEM) = 1 x 4608GB = 4608GB
- /hana/log >= 512 GB
- /hana/shared >= 1 TB

2.5. Operating system

SAP HANA with Intel Optane PMEM is supported by the following operating systems:

- SUSE Linux Enterprise Server (SLES) for SAP Applications
  - SLES for SAP Applications 12 SP4
  - SLES for SAP Applications 15
  - SLES for SAP Applications 15 SP1
- Red Hat Enterprise Linux (RHEL)
  - RHEL 7.6 for SAP Solutions
  - RHEL 8 for SAP Solutions

2.6. Recommendations for Persistent Memory Configuration

Customers must carefully evaluate the sizing requirements of the applications to use different ratios, and we highly recommend running SAP's Quick Sizer* tool and utilize the SQL statements for memory assessment. (According to document at https://launchpad.support.sap.com/#/notes/2786237)

SAP HANA quick sizer:

https://apps.support.sap.com/sap/bc/bsp/sap/qs_oberflaeche/hana.do?saplanguage=en&bsp-language=en

3. Installation Planning for SAP HANA scale-up node

3.1. Hardware planning

The hardware plan used in this installation guide is shown in the following table：

PLATFORM		Inspur NF8260M5
PROCESSOR		4 x Intel® Xeon® Platinum 8276M
MEMORY	**Total**	4608GB (PMEM/DRAM Ratio 2:1)
	PMEM	24 * 128GB
	DRAM	24 * 64GB
STORAGE	**BOOT drive**	2* 480GB Intel® SSD D3-S4510, RAID-1
	HANA data	3* 3.2TB Intel® SSD DC P4610, RAID-5
SOFTWARE		OS: SLES 12 SP4 for SAP
		SAP HANA 2.0 SPS04

3.2. Software planning

The software planning is shown in the following table 1.

Software name (version)	Source	Description
SLE-12-SAP4-DVD-x86_64-GM-DVD.iso	https://www.suse.com	OS
SAP HANA 2.0 SPS04	http://www.sap.com	SAP HANA Database
ipmctl utility	https://github.com/intel/ipmctl	PMEM management tools
ndctl utility	https://github.com/pmem/ndctl	

3.3. Built-in disk capacity planning

NF8260M5 internal storage planning as follows:

- A 2 * 480GB SATA SSD configured as RAID-1 is used to install the operating system.
- A 3 * 3.2TB 2.5-inch configured as RAID-5 is used to store "data, shared, log" volumes of the HANA database.

Partition	File System Type	Partition Size	Description
/boot/efi	VFAT	500 MB	Pool from 2* 480GB SSD，RAID-1
swap	SWAP	50 GB	Pool from 2* 480GB SSD，RAID-1
/	XFS	400GB	Pool from 2* 480GB SSD，RAID-1
/hana/data	XFS	4.5TB	Pool from 3*3.2TB SSD, RAID-5
/hana/shared	XFS	1TB	Pool from 3*3.2TB SSD, RAID-5
/hana/log	XFS	512GB	Pool from 3*3.2TB SSD, RAID-5
/usr/sap	XFS	50GB	Pool from 3*3.2TB SSD, RAID-5

4. BIOS Configuration

This section describes how to configure the BIOS parameters of the server to meet performance requirements for SAP HANA operation.

Step 1. When the screen is shown as in Figure 4-1 during the server startup, press "Delete" to open the BIOS setup screen.

Figure 4-1. Inspur Logo Screen

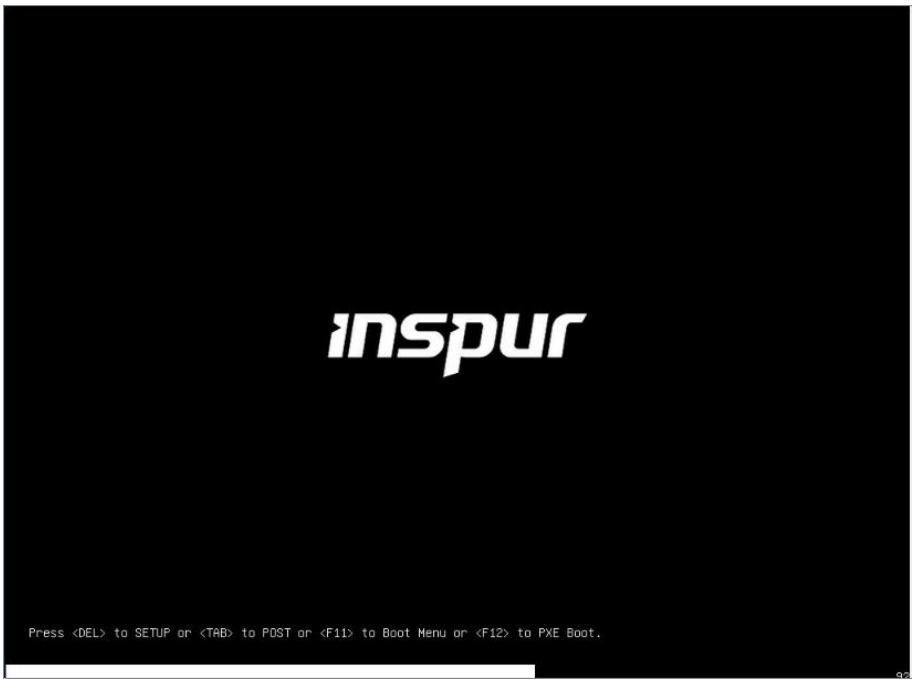

Step 2. In the BIOS menu list, select "Processor-> Advanced Power Management Configuration" as shown in Figure 4-2.

Figure 4-2. Processor Interface

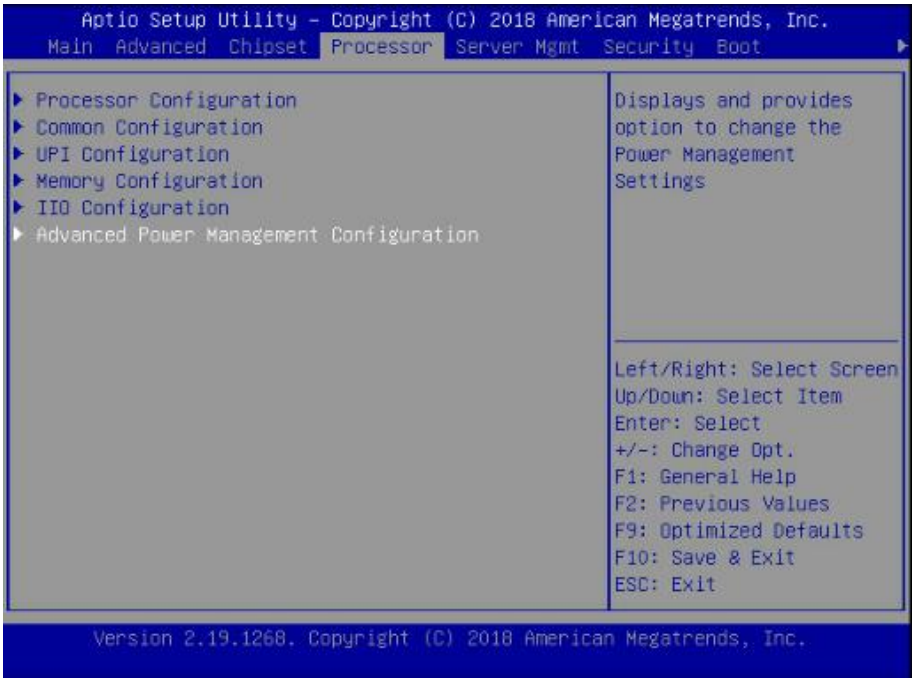

Step 3. In the Advanced Power management Configuration interface, select "CPU P State Control" to setup p-state parameters as shown in Figure 4-3 and Figure 4-4.

Figure 4-3. CPU Power Management interface

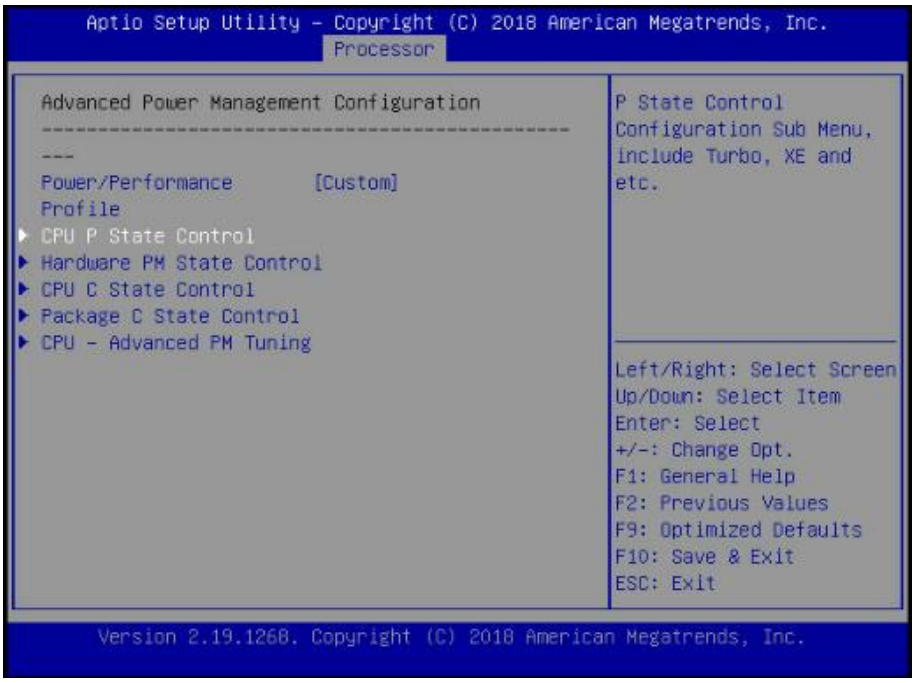

Figure 4-4. CPU P State Control interface

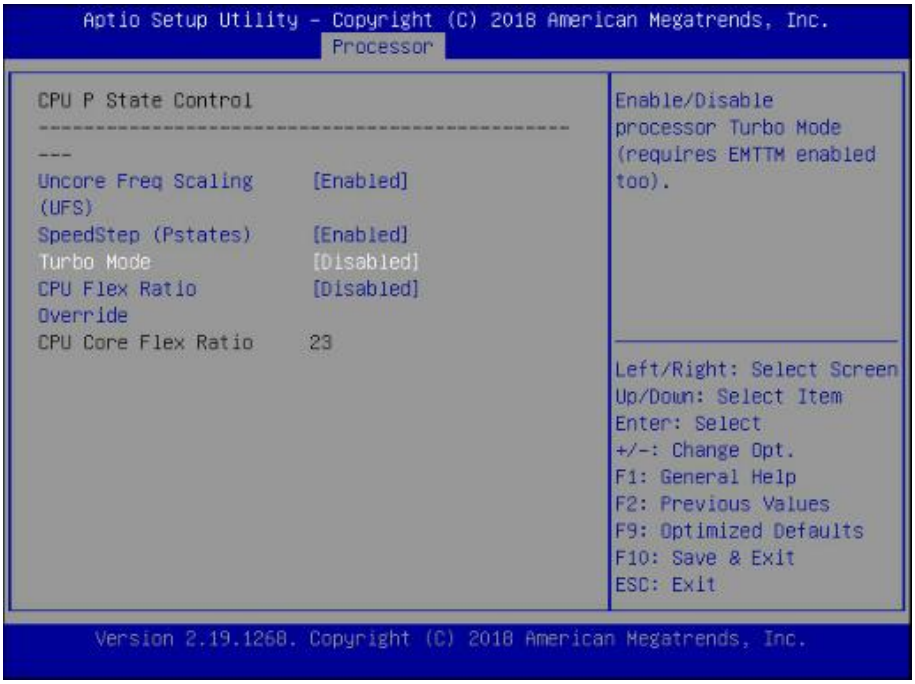

Step 4. Return to the Advanced Power management Configuration interface, select "Package C State Control" to confugre C-State parameter to "C0/C1 state" as shown in Figure 4-5 and Figure 4-6.

Figure 4-5. CPU Power Management interface

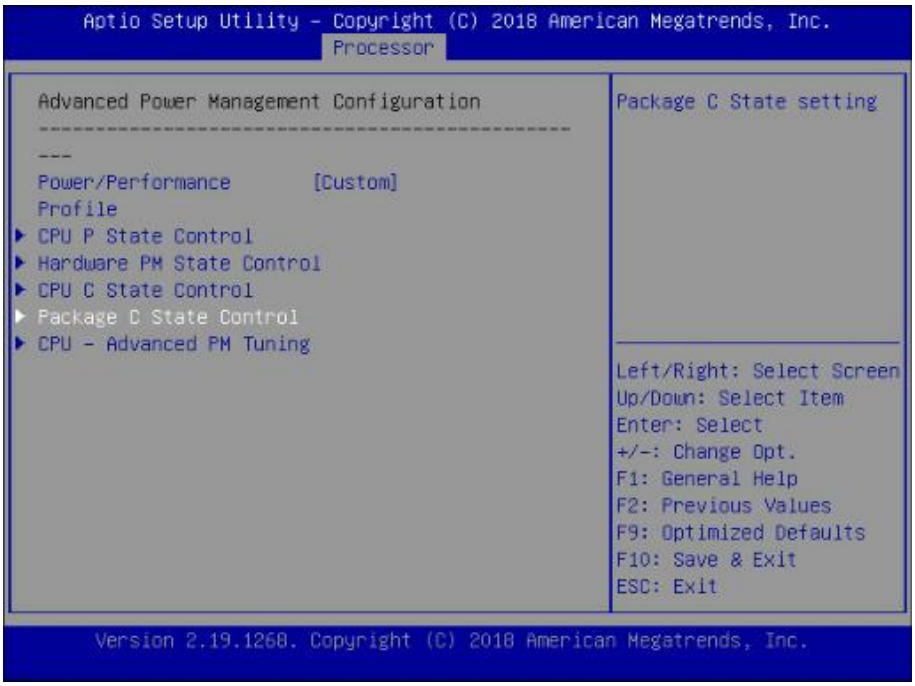

Figure 4-6. CPU C State Control interface

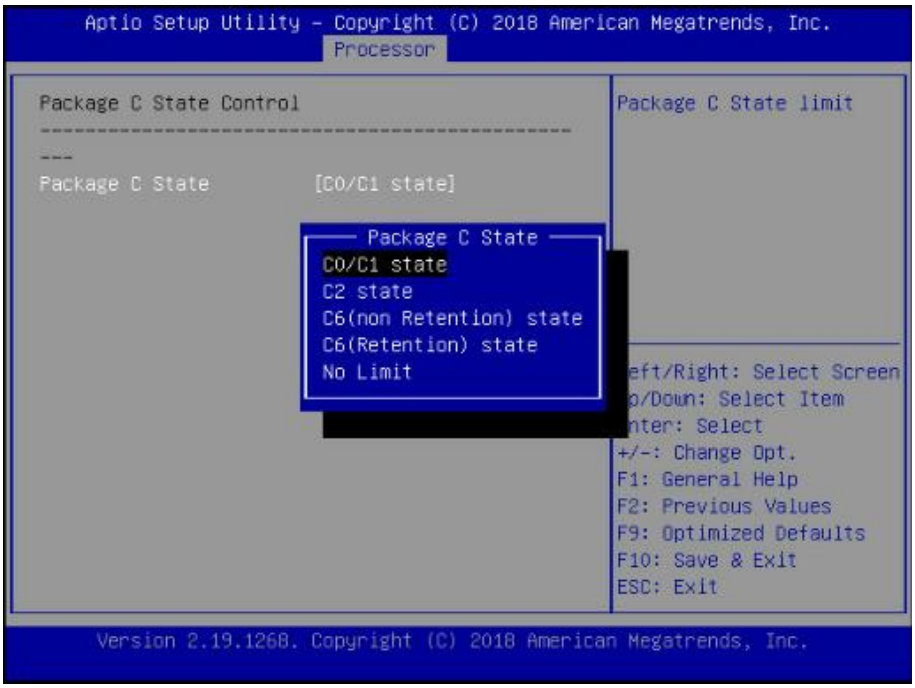

Step 5. Return to the Advanced Power management Configuration interface, select "CPU Advanced PM Tuning ->Energy Perf BIAS" as shown in Figure 4-7 and Figure 4-8.

Figure 4-7. CPU Power Management interface

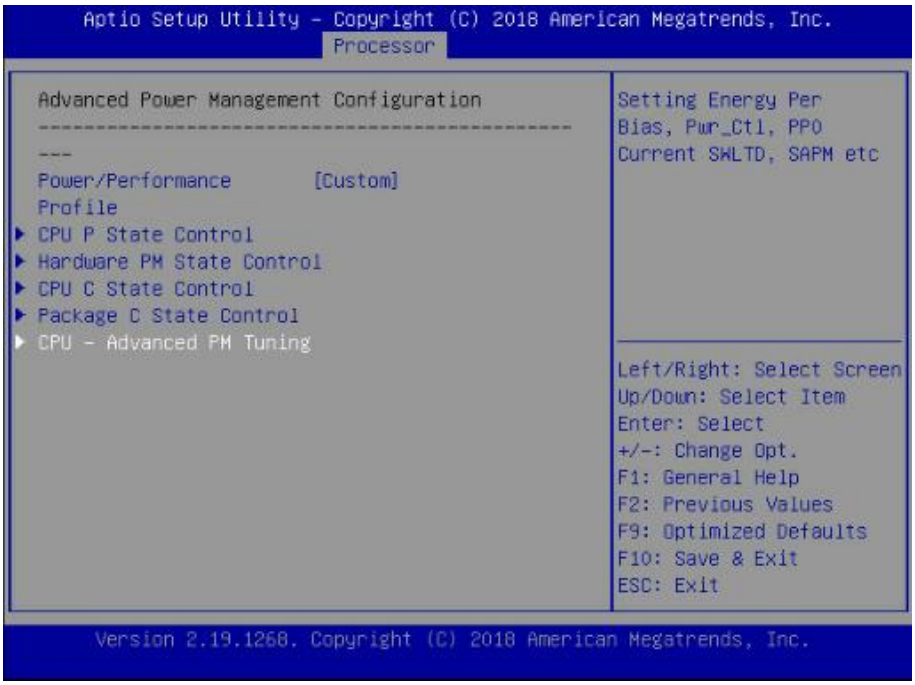

Figure 4-8. CPU Advanced PM Tuning interface

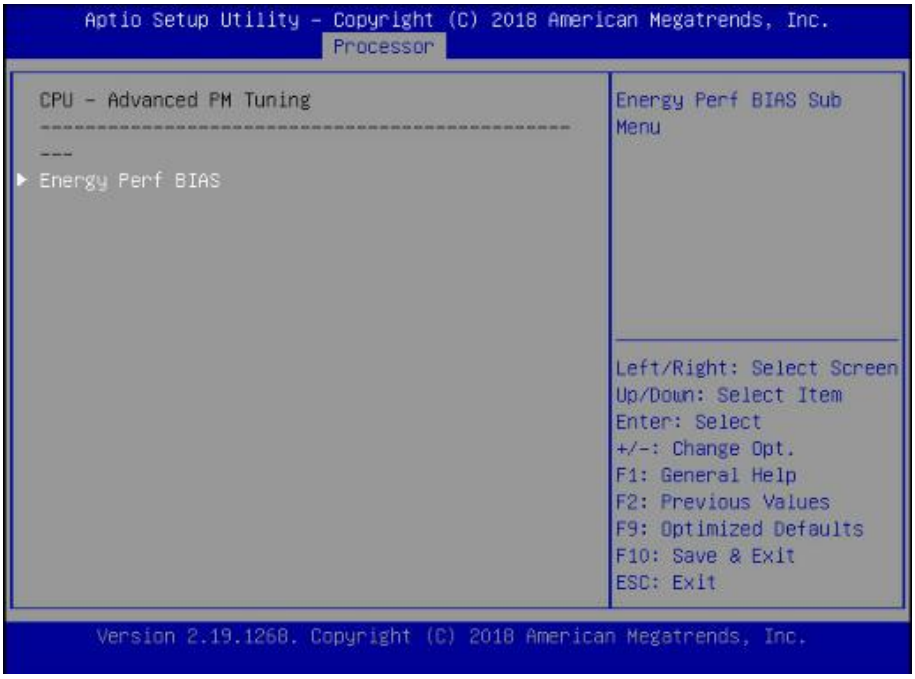

Step 6. In the Energy Perf BIAS interface, set parameter to "OS Control EPB" as shown in Figure 4-9.

Figure 4-9. Energy Perf BIAS interface

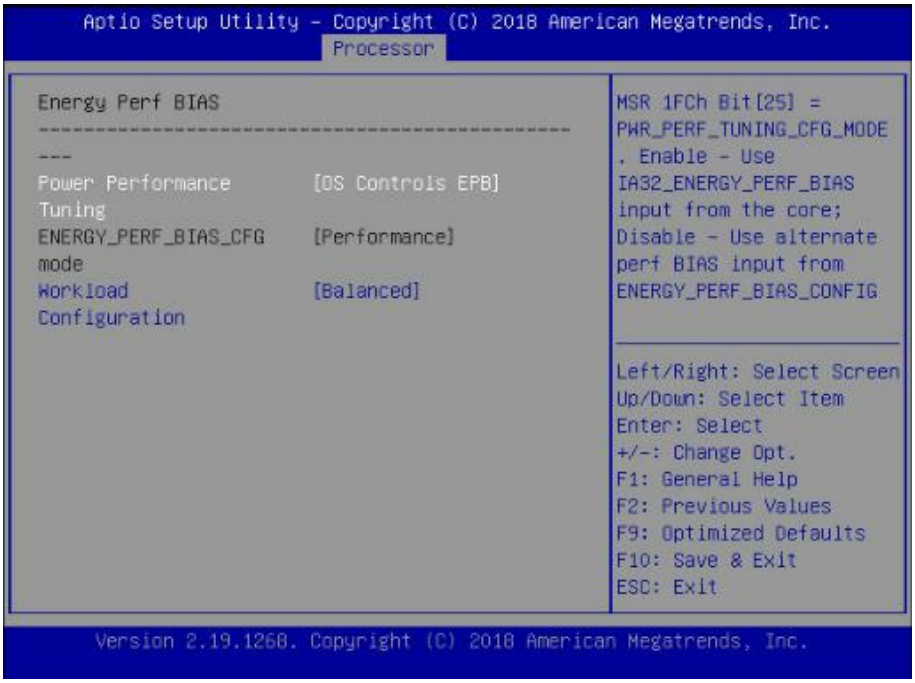

At this point, the BIOS parameters configuration is completed.

5. RAID configuration

This section describes how to configure disk raid for OS and SAP Hana data volumes.

5.1. Create virtual drive for OS

Step 1. During the server startup, press "Delete" to open the BIOS setup screen.

Step 2. In the BIOS interface, select "Advanced->AVAGO MegaRAID<AVAGO MegaRAID SAS 9361-8i> ".

Step 3. Select "Configure ->Create Virtual Drive ".

Step 4. Configure "Select RAID Level" as "RAID1", and select "Select Drives".

Step 5. Set the status of 480GB disks to "Enable" and apply Changes.

Step 6. Save settings and press F10 to exit.

5.2. Create virtual drive for SAP HANA

Step 1. During the server startup, press "Delete" to open the BIOS setup screen.

Step 2. Enter into the "Advanced->Intel(R) Virtual RAID on CPU" interface

Step 3. Select the appropriate number of disks according to the actual situation and select the RAID5 level.

> ***Note:*** *The default size of Capacity is ~95% of the actual disk capacity and needs to be modified manually.*

Step 4. Select "Create Volume" and Press the enter key to confirm.

Step 5. Press F10 to save the configuration and exit.

6. Operating System Configuration

This section describes how to install SLES 12 SP4.

6.1. OS installation

Step 1. Select "Installation" and press "Enter" in the installation interface, as shown in the figure below.

Figure 6-1. Installation interface

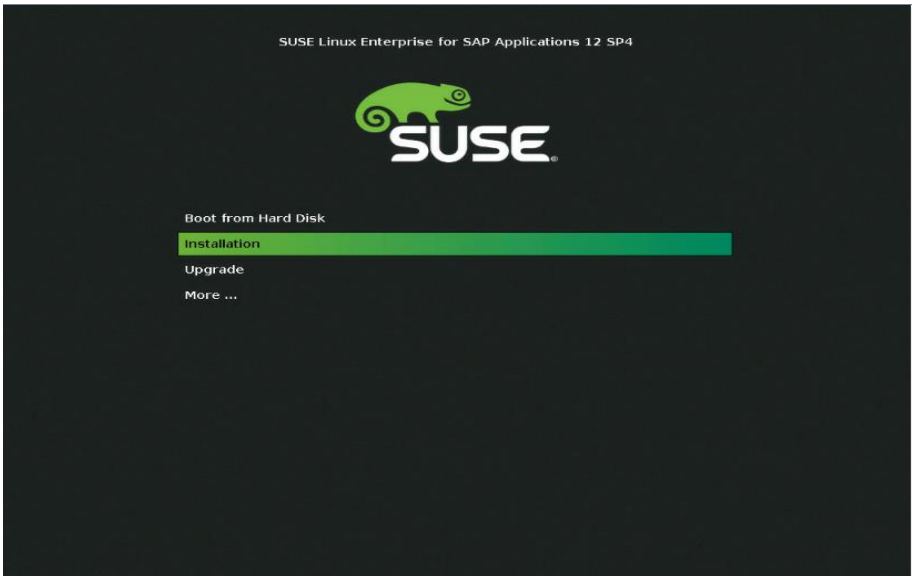

Figure 6-2. Initialization interface

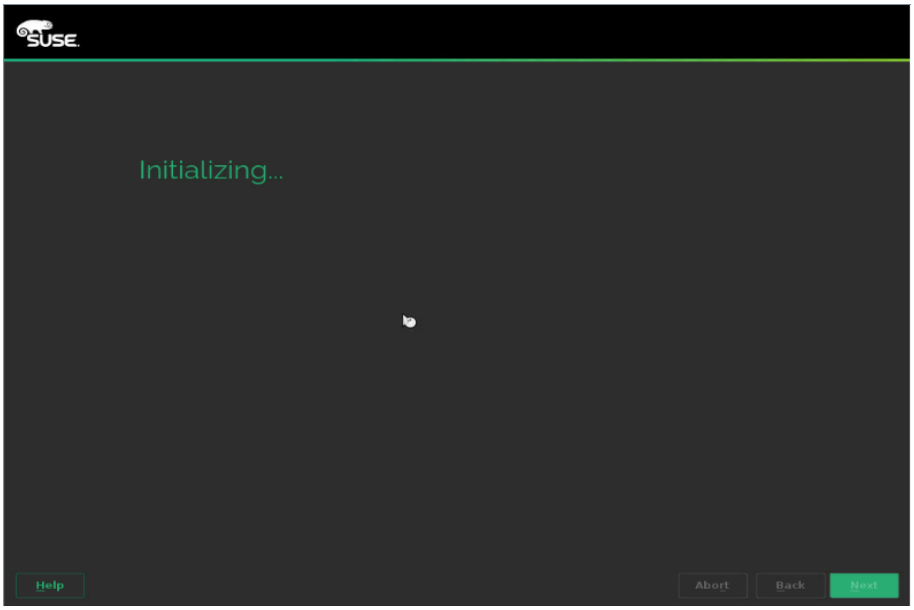

Step 2. Check "I Agree to the License Terms" and click "Next" as shown in the figure below.

Figure 6-3. Check "I Agree to the Licese Terms"

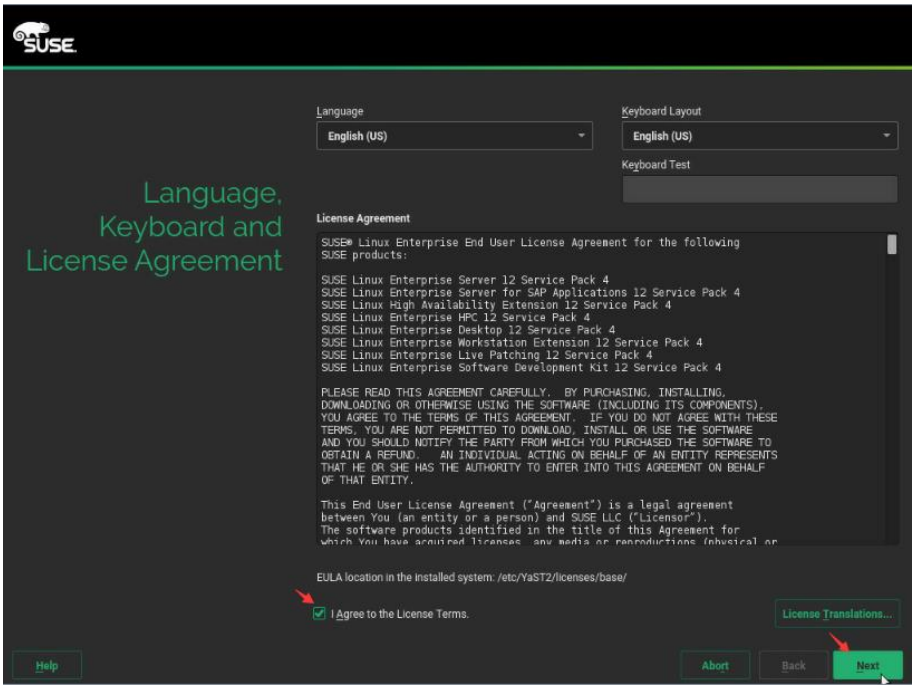

Step 3. Skip Registration as shown in Figure 6-4.

Figure 6-4. Skip Registration

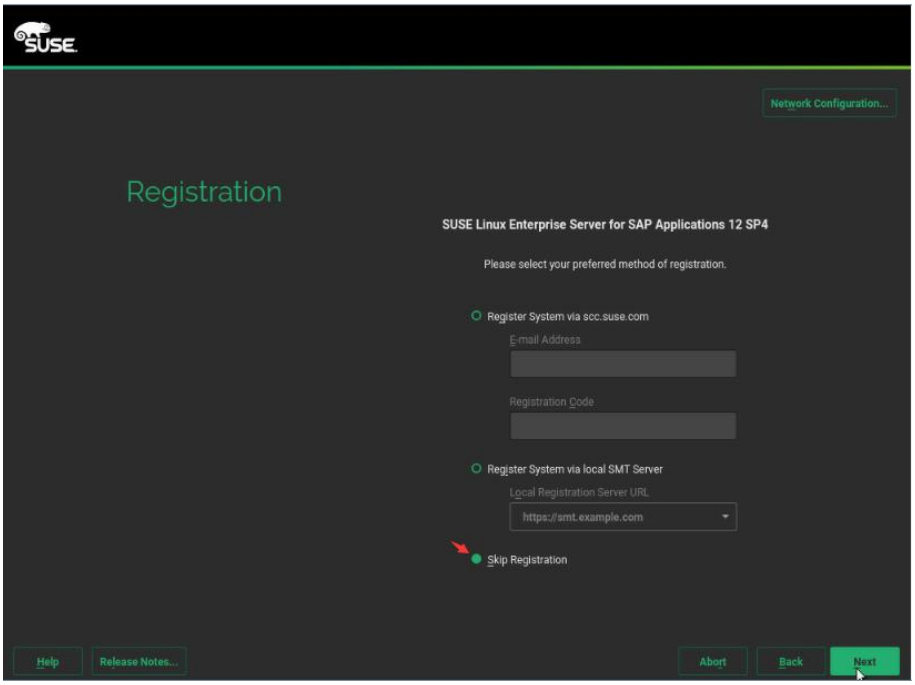

Step 4. Choose "SUSE Linux Enterprise Server for SAP Application" and click "Next" as shown in Figure 6-5.

Figure 6-5. Figure 4-6 Choose installation model

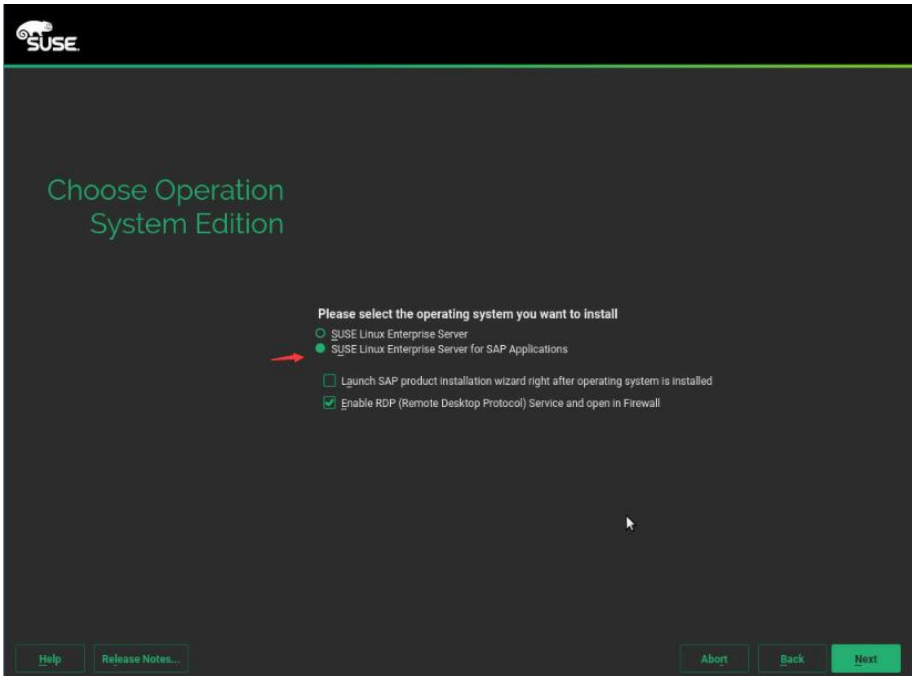

Step 5. Click "Next" as shown in figure 6-6.

Figure 6-6. Add On Product Interface

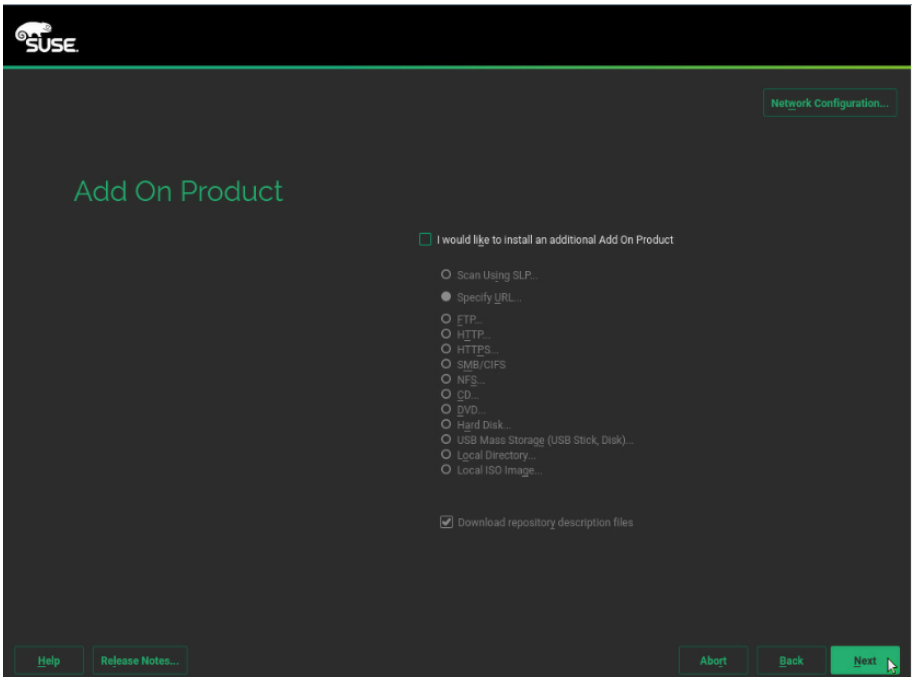

Step 6. Enter into the installation settings page as shown in Figure 6-7 and click "Expert Partitioner".

Figure 6-7. Suggested Partitioning

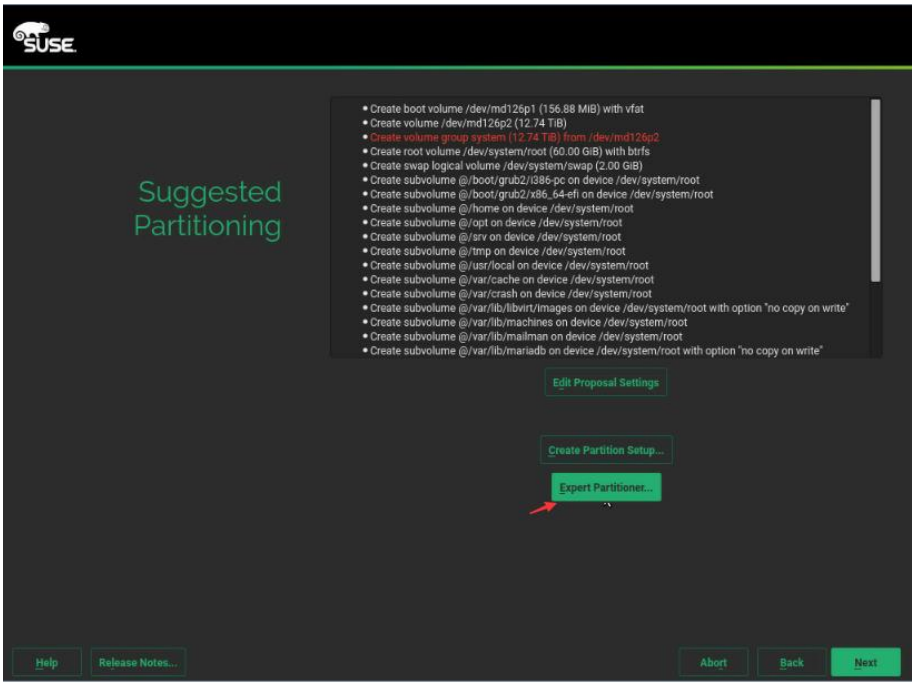

Step 7. Adjust the size of the partition according to the actual demand, as shown in the figure below.

Figure 6-8. Expert Partitioner

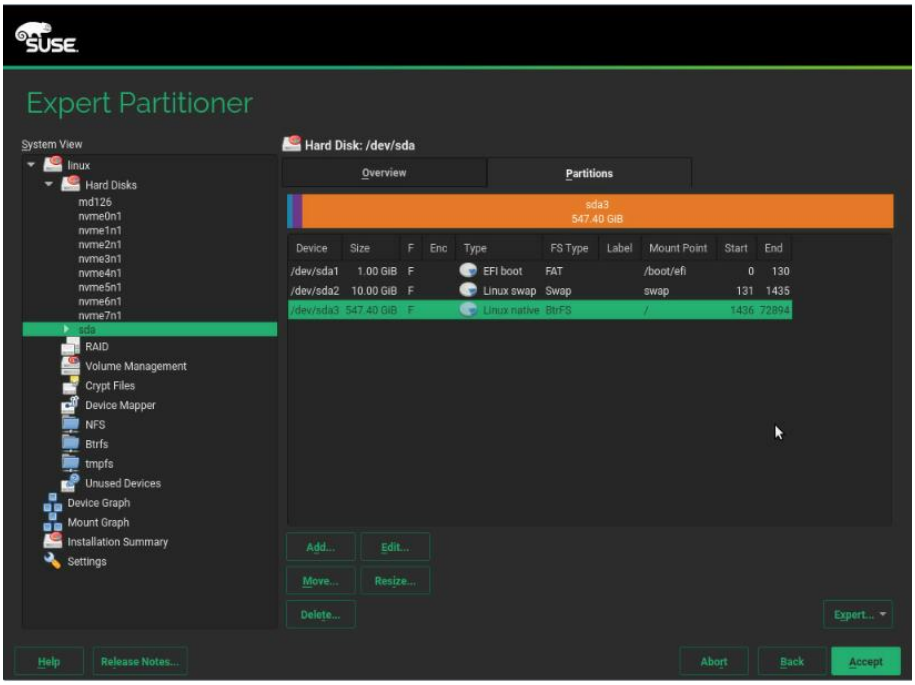

Step 8. Select time zone and click "Next" as shown in Figure 6-9.

Figure 6-9. Colck and Time Zone

Step 9. Set password for root user.

Figure 6-10. Set password

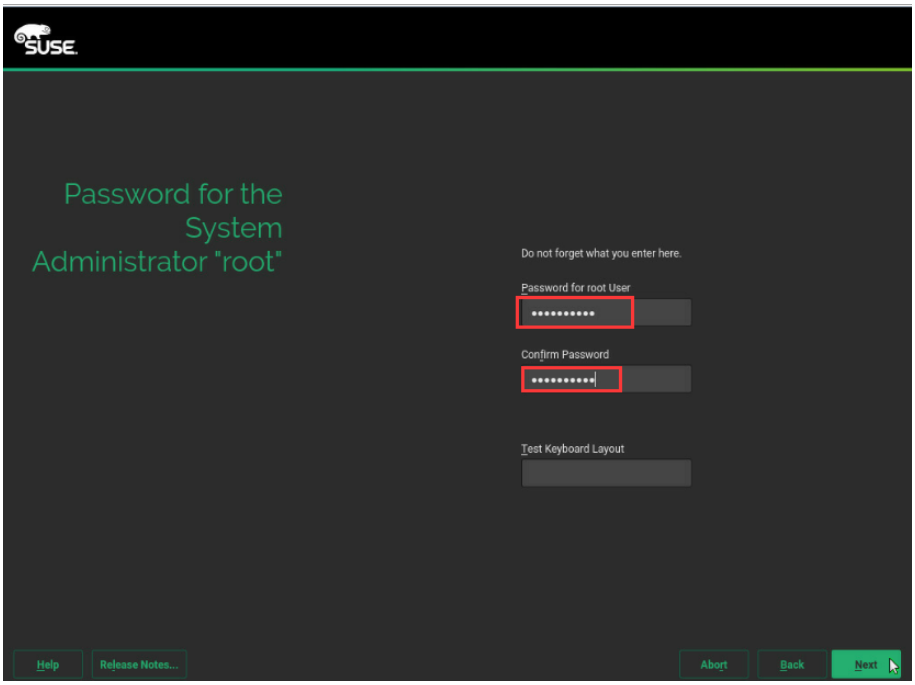

Step 10. Click "Software" and choose "SAP HANA Server Base".

Figure 6-11. Choose Software

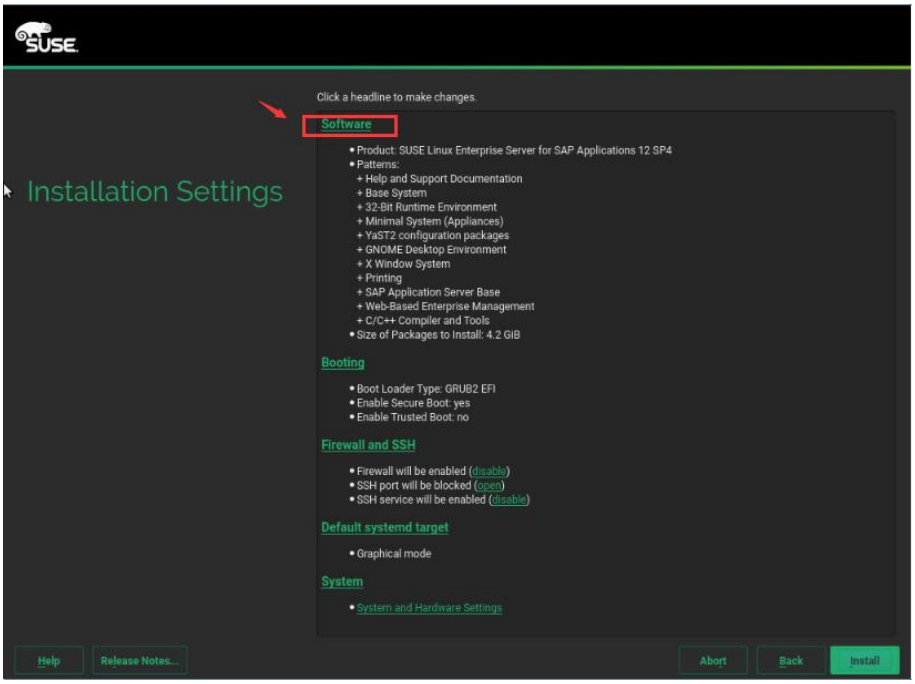

Figure 6-12. Software selection

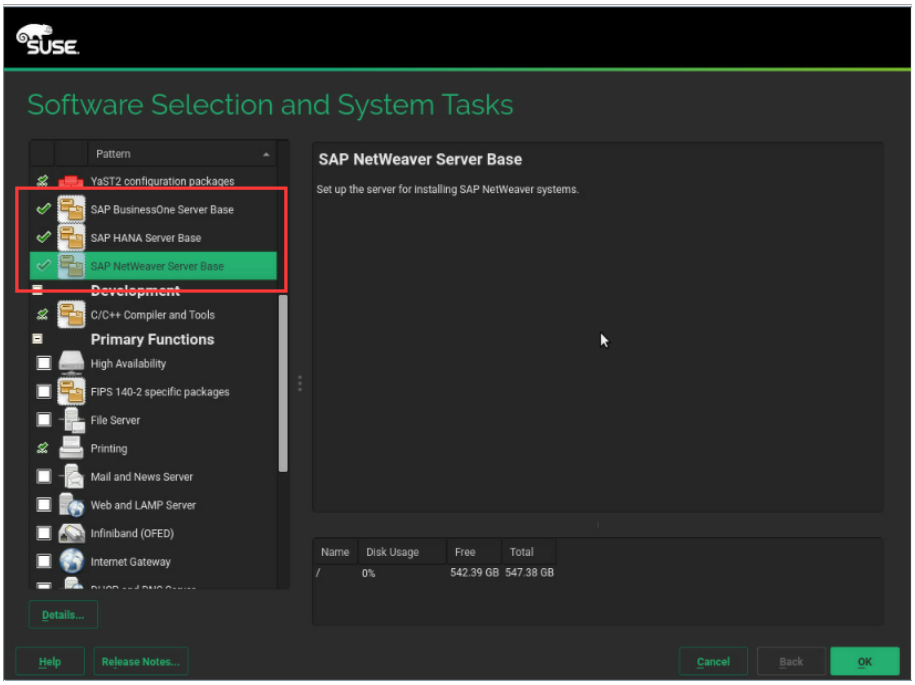

Step 11. Disable Firewall and install.

Figure 6-13. Disable Firewall

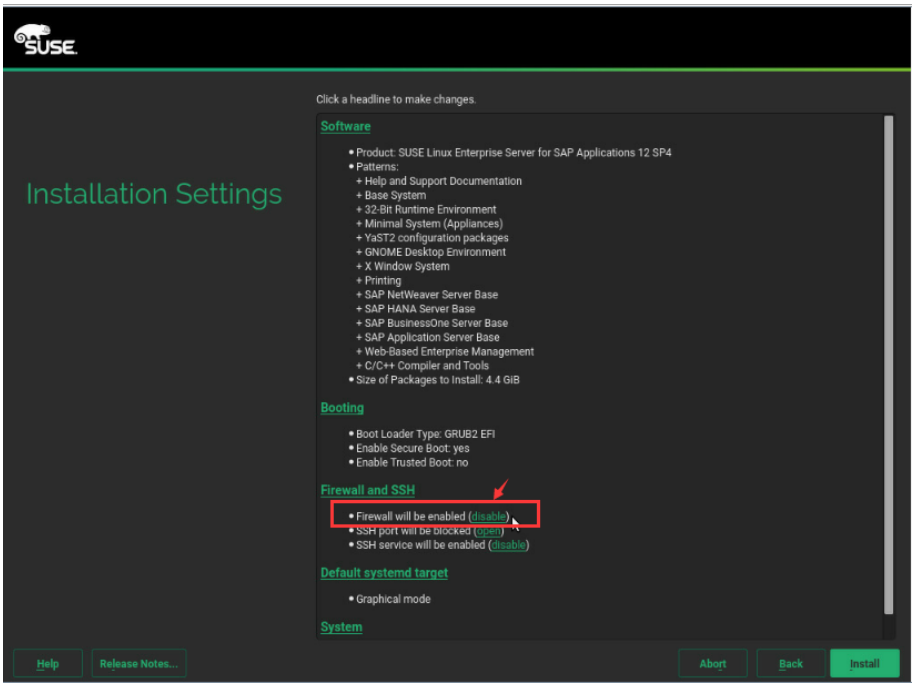

Figure 6-14. Confirm Install

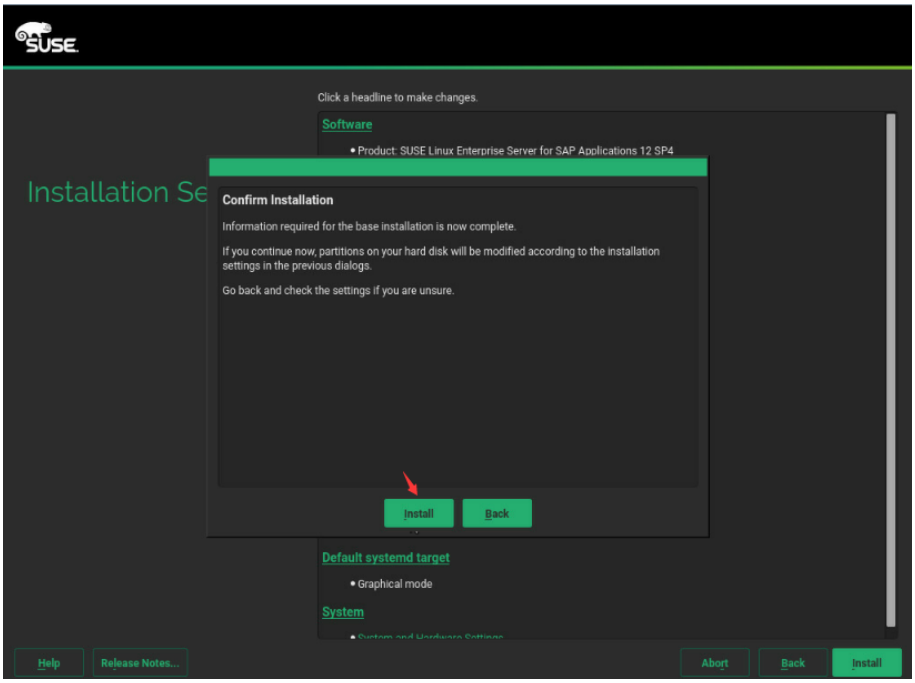

6.2. Host name configuration

Step 1. Add the following information in "/etc/hosts" (establish mapping of each node IP address and host name).

```
10.152.18.216 NF8260M5
```

6.3. saptune service configuration

This section introduces the advanced saptune configuration, tuning and checking tool included with SLES for SAP Applications 12 SP1 and later. Saptune can prepare the operating system based on implementing specific SAP notes, groups of SAP notes per solution (e.g. SAP HANA, SAP NetWeaver).

Step 1. Configure the tuned service to start with the operating system startup.

```
systemctl enable tuned.service
```

Step 2. Create a directory with the name of the profile: /etc/tuned/saptune/.

```
mkdir /etc/tuned/saptune/
```

Step 3. Copy the tuned configuration: 'cp /usr/lib/tuned/saptune/tuned.conf /etc/tuned/saptune/'.

```
cp /usr/lib/tuned/saptune/tuned.conf /etc/tuned/saptune/
```

Step 4. Edit the '/etc/tuned/saptune/tuned.conf' file file and modify the value of the script.

```
script = /usr/lib/tuned/saptune/script.sh
```

Step 5. Execute the following command to active saptune service:

```
saptune daemon start
```

Step 6. Adjust saptune solution and note:

```
saptune solution apply HANA
saptune note apply SUSE-GUIDE-01
saptune note apply SUSE-GUIDE-02
```

6.4. SAP HANA Data volumes configuration

Step 1. Run "pvcreate" command to create physical devices.

```
NF8260M5:~ # pvcreate /dev/md126
```

Step 2. Run "vgcreate" command to group two physical volume into VG0

```
vgcreate vg0 /dev/md126
```

Step 3. Run "lvcreate" command to create "data, shared, log" logical volume for SAP HANA.

```
lvcreate -L 4.5TB -n data vg0
lvcreate -L 1TB -n shared vg0
lvcreate -L 512GB -n log vg0
lvcreate -L 50GB -n usr_sap vg0
```

Notice:

Please set the LVM volume size according to the actual situation.

6.5. Mounting SAP Data volume

Step 1. Format all the volumes。

```
mkfs.xfs -f /dev/vg0/data
mkfs.xfs -f /dev/vg0/shared
mkfs.xfs -f /dev/vg0/log
mkfs.xfs -f /dev/vg0/usr_sap
```

Step 2. Run "mkdir" command to create mount directory which is needed by SAP HANA installation.

```
mkdir /hana /hana/log /hana/data /hana/shared /usr/sap
```

Step 3. Use root user to add mount directory list into "/etc/fstab" file, ensure that the "/hana/data, /hana/log, /hana/shared" file system automatically mount at system startup.

```
echo "/dev/mapper/vg0-data /hana/data xfs defaults 0 0">>
/etc/fstab
echo "/dev/mapper/vg0-shared /hana/shared xfs defaults 0 0">>
/etc/fstab
```

```
echo "/dev/mapper/vg0-log /hana/log xfs defaults 0 0">>
/etc/fstab
echo "/dev/mapper/vg0-usr_sap /usr/sap xfs defaults 0 0">>
/etc/fstab
```

6.6. Intel Optane PMEM configuration

Follow the github links below to install the latest version of ipmctl and ndctl on the SAP HANA Linux server.

- For ipmctl utility https://github.com/intel/ipmctl
- For ndctl utility library https://github.com/pmem/ndctl

Step 1. The default **create -goal** command creates an interleaved region configured for App Direct mode. Here is an example of output from the **ipmctl create -goal** command:

```
ipmctl create -goal PersistentMemoryType=AppDirect
```

Step 2. Showing the regions created.

```
ipmctl show -region
```

Step 3. Creating a name space for each region.

Use the ndctl create-namespace command to create a name space for each region. You must run this command for each CPU in the server. Here is an example of output from the ndctl create-namespace command on a server with four CPUs.

Filesystem-DAX mode is the default name-space mode. If you specify ndctl create-namespace with no options, a block device (/dev/pmemX[.Y]) is created that supports the DAX capabilities of Linux file systems. DAX removes the page cache from the I/O path and allows mmap(2) to establish direct mappings to persistent memory media. In this mode, applications can either directly load and access storage using a persistent memory region or continue to use a storage API, thus requiring no changes to the application.

```
ndctl create-namespace --mode=fsdax --region=region0
ndctl create-namespace --mode=fsdax --region=region1
ndctl create-namespace --mode=fsdax --region=region2
ndctl create-namespace --mode=fsdax --region=region3
```

Step 4. Use the **ndctl list** command to list all the active name spaces. Here is an example of output from the **ndctl** list command.

```
ndctl list
 [
{
```

```
"dev":"namespace3.0",
"mode":"fsdax",
"map":"dev",
"size":756.00 GiB,
"uuid":"43002f2c-b37c-4cec-9474-d3d8b1223e65",
"blockdev":"pmem3"
},
{
"dev":"namespace2.0",
"mode":"fsdax",
"map":"dev",
"size":756.00 GB,
"uuid":"45e0fc9e-149c-4616-b308-eb10eecd5e19",
"blockdev":"pmem2"
},
{
"dev":"namespace1.0",
"mode":"fsdax",
"map":"dev",
"size":756.00 GB,
"uuid":"9375a814-ac10-498a-9e73-3e28e7242519",
"blockdev":"pmem1"
},
{
"dev":"namespace0.0",
"mode":"fsdax",
"map":"dev",
"size":756.00 GB,
"uuid":"83425d72-c451-4eb7-b450-8dc3f4b1978a",
"blockdev":"pmem0"
}
]
```

Step 5. Creating the file system and mounting the persistent memory nodules.

Use this set of commands to create the file system and mount the persistent memory modules. This example uses a server with four CPUs. It therefore has four regions.

```
mkfs.xfs /dev/pmem0
mkfs.xfs /dev/pmem1
```

```
mkfs.xfs /dev/pmem2
mkfs.xfs /dev/pmem3
mkdir -p /hana/pmem/nvmem0
mkdir -p /hana/pmem/nvmem1
mkdir -p /hana/pmem/nvmem2
mkdir -p /hana/pmem/nvmem3
mount -t xfs -o dax /dev/pmem0 /hana/pmem/nvmem0
mount -t xfs -o dax /dev/pmem1 /hana/pmem/nvmem1
mount -t xfs -o dax /dev/pmem2 /hana/pmem/nvmem2
mount -t xfs -o dax /dev/pmem3 /hana/pmem/nvmem3
```

6.7. System parameters configuration

Step 1. Stop and disable irqbalance service:

```
systemctl stop irqbalance.service
systemctl disable irqbalance.service
```

Step 2. Create a boot.sh file in the /etc/ directory.

```
vi /etc/rc.d/boot.local
```

Step 3. Add the following to the boot.sh file.

```
#!/bin/bash
echo never >/sys/kernel/mm/transparent_hugepage/enabled
echo 0 >/sys/kernel/mm/ksm/run
echo 0 >/proc/sys/kernel/numa_balancing
for i in `ls  /sys/block/|grep md`
do
   if [ -f /sys/block/$i/md/sync_speed_max ];then
     echo 5000000 > /sys/block/$i/md/sync_speed_max
     echo 2000000 > /sys/block/$i/md/sync_speed_min
   fi
   if [ -f /sys/block/$i/md/group_thread_cnt ];then
     echo 8 > /sys/block/$i/md/group_thread_cnt
echo 4096 > /sys/block/$i/md/stripe_cache_size
   fi
done
sleep 10
mount -t xfs -o dax /dev/pmem0 /hana/pmem/nvmem0
mount -t xfs -o dax /dev/pmem1 /hana/pmem/nvmem1
```

```
mount -t xfs -o dax /dev/pmem2 /hana/pmem/nvmem2
mount -t xfs -o dax /dev/pmem3 /hana/pmem/nvmem3
```

Step 4. Add executable permissions for the boot.sh file.

```
chmod a+x /etc/rc.d/boot.local
```

Step 5. Create a service for the boot.sh script.

```
vi /etc/systemd/system/script01.service
```

Add the following:

```
[Unit]
Description=Run boot script
After=network.target

[Service]
Type=oneshot
ExecStart=/bin/bash -c "/etc/boot.sh"

[Install]
WantedBy=multi-user.target
```

Step 6. Configure the service to start automatically when the server starts.

```
systemctl enable script01.service
```

7. SAP HANA DATABASE Installation

Use the official SAP documentation, which describes the installation process with and without the SAP unified installer. For the SAP

HANA installation documentation, see **SAP HANA Server Installation Guide**. All other SAP HANA administration documentation is available at SAP HANA Administration Guide.

7.1. Preparation

Installation preparation information：

Parameters	Value	Description
HANA System ID	Example: ANA	HANA System ID, three letters or numbers, the first must be letters

Instance Number	00	HANA Instance Number is used for distinguishing HANA system in a server
Database Mode	single_container or multiple_containers	Single container or multiple container
System Usage	Production，test，development，custom	Database is used for production, test or development. Customized Database is used for test and development
System Password	xxxx	Database password, set by users
Restrict maximum memory	for example:256GB	To install multiple databases on a server and limit the memory allocation for different databases

7.2. SAP HANA database installation

Step 1. Upload the SAP HANA database media to the server.

Step 2. Install database using SAP HANA **hdblcm** tool.

```
./hdblcm

SAP HANA Lifecycle Management - SAP HANA Database 2.00.040.00.1553674765
************************************************************
**********


Scanning software locations...
Detected components:
   SAP HANA Database (2.00.040.00.1553674765) in /soft/hanadb/SAP_HANA_DATABASE/server
   SAP HANA Database Client (2.3.78.1521836270) in /soft/hanadb/SAP_HANA_CLIENT/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 ' 2.00.040.00.1553674765' will be
installed.

Select additional components for installation:

  Index | Components | Description
  -------------------------------------------------------------
-------------------
  1     | all        | All components
  2     | server     | No additional components
  3     | client     | Install SAP HANA Database Client version
2.3.78.1521836270

Enter comma-separated list of the selected indices [3]: 1
Enter Installation Path [/hana/shared]: (keep default, press
enter)
Enter Local Host Name [NF8260M5]: (keep default, press enter)
Do you want to add hosts to the system? (y/n) [n]: (keep default,
press enter)
Enter SAP HANA System ID: ANA (SID of HANA system)
Enter Instance Number [00]: (keep default, press enter)
Enter Local Host Worker Group [default]: (keep default, press
enter)

  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]: 4  (Based on the purposes
of the database)
Enter Location of Data Volumes [/hana/data/ANA]: (keep default,
press enter)
Enter Location of Log Volumes [/hana/log/ANA]: (keep default,
press enter)
Restrict maximum memory allocation? [n]: (keep default, press
enter)
Enter Certificate Host Name For Host 'NF8260M5' [NF8260M5]:
(keep default, press enter)
Enter SAP Host Agent User (sapadm) Password: ( input user
password)
Confirm SAP Host Agent User (sapadm) Password: ( input user
password)
Enter System Administrator (anaadm) Password: ( input user
password)
Confirm System Administrator (anaadm) Password: ( input user
password)
Enter System Administrator Home Directory [/usr/sap/ANA/home]:
(keep default, press enter)
Enter System Administrator Login Shell [/bin/sh]: (keep default,
press enter)
Enter System Administrator User ID [1001]: (keep default, press
enter)
Enter System Database User (SYSTEM) Password: ( input system
user password)
Confirm System Database User (SYSTEM) Password: ( input system
user password)
Restart system after machine reboot? [n]: (keep default, press
enter , not restart)

Summary before execution:
==========================
```

```
SAP HANA Database System Installation
   Installation Parameters
      Remote Execution: ssh
      Database Isolation: low
      Installation Path: /hana/shared
      Local Host Name: NF8260M5
      SAP HANA System ID: ANA
      Instance Number: 00
      Local Host Worker Group: default
      System Usage: custom
      Location of Data Volumes: /hana/data/ANA
      Location of Log Volumes: /hana/log/ANA
      Certificate Host Names: NF8260M5 -> NF8260M5
      System Administrator Home Directory: /usr/sap/ANA/home
      System Administrator Login Shell: /bin/sh
      System Administrator User ID: 1001
      ID of User Group (sapsys): 79
      SAP HANA Database Client Installation Path:
/hana/shared/ANA/hdbclient
   Software Components
      SAP HANA Database
         Install version 2.00.040.00.1553674765
         Location: /soft/hanadb/SAP_HANA_DATABASE/server
      SAP HANA Database Client
         Install version 2.3.78.1521836270
         Location: /soft/hanadb/SAP_HANA_CLIENT/client

Completely installation is shown as follows.
Do you want to continue? (y/n): y (enter y to start the
installation)
```

Step 3. Modify the permissions of the pmem directory according to the SAP HANA system user. In this guide, the SAP HANA instance is named ANA. After the database installation is complete, the **<sid> adm** user will be generated in the system, and its group is **sapsys**.

```
chown -R anaadm:sapsys /hana/pmem
```

Step 4. Setting the SAP HANA base path to use persistent memeory.

The directory that SAP HANA uses as its base path must point to the XFS file system. Define the base path location with the configuration parameter

basepath_persistent_memory_volumes in the persistence section of the SAP HANA global.ini file. This section can contain multiple locations separated by semicolons. Changes to this parameter require a restart of SAP HANA services.

```
[persistence]
basepath_datavolumes = /hana/data/ANA
basepath_logvolumes = /hana/log/ANA
basepath_persistent_memory_volumes=/hana/pmem/nvmem0;/hana/pmem/nvmem1;/hana/pmem/nvmem2;/hana/pmem/nvmem3
```

7.3. Performing an SAP HANA post-installation checkup

For an SAP HANA system installed with <SID> set to **ANA** and the system number <nr> set to **00**, log in as **anaadm** and run the commands presented here.

Step 1. Commands for checking SAP HANA services

```
anaadm@NF8260M5:/usr/sap/ANA/HDB00>
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetProcessList
```

Step 2. Commands for checking SAP HANA database information

```
anaadm@NF8260M5:/usr/sap/ANA/HDB00> HDB info
```

Step 3. Commands for stopping SAP HANA database.

```
anaadm@NF8260M5:/usr/sap/ANA/HDB00>
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function StopSystem HDB
```

Step 4. Commands for starting SAP HANA database.

```
anaadm@NF8260M5:/usr/sap/ANA/HDB00>
/usr/sap/hostctrl/exe/sapcontrol -nr 00 -function StartSystem HDB
```

8. Important SAP Notes

Read the following SAP Notes before you start the installation. These SAP Notes contain the latest information about the

installation, as well as corrections to the installation documentation.

The latest SAP Notes can be found at SAP Notes and Knowledge base.

SAP HANA IMDB notes

- SAP Note 1514967: SAP HANA: Central note
- SAP Note 2298750: SAP HANA Platform SPS 12 Release Note

- SAP Note 1523337: SAP HANA database: Central note
- SAP Note 2000003: FAQ: SAP HANA
- SAP Note 2380257: SAP HANA 2.0 Release Notes
- SAP Note 1780950: Connection problems due to host name resolution
- SAP Note 1755396: Released disaster tolerant (DT) solutions for SAP HANA with disk replication
- SAP Note 2519630: Check whether power save mode is active
- SAP Note 1681092: Support for multiple SAP HANA databases on a single SAP HANA appliance
- SAP Note 1514966: SAP HANA: Sizing the SAP HANA database
- SAP Note 1637145: SAP BW on HANA: Sizing the SAP HANA database
- SAP Note 1793345: Sizing for Suite on HANA
- SAP Note 2399079: Elimination of hdbparam in HANA 2
- SAP Note 2186744: FAQ: SAP HANA Parameters
- SAP Note 2786237 - Sizing SAP HANA with Persistent Memory

Linux notes:

- SAP Note 2684254: SAP HANA DB: Recommended OS settings for SLES 15 and SLES for SAP Applications 15
- SAP Note 2235581: SAP HANA: Supported operating systems
- SAP Note 1944799: SAP HANA guidelines for the SLES operating system
- SAP Note 1731000: Non-recommended configuration changes
- SAP Note 1557506: Linux paging improvements
- SAP Note 1740136: SAP HANA: Wrong mount option may lead to corrupt persistency
- SAP Note 2382421: Optimizing the network configuration on the HANA and OS levels

Third-party software notes:

- SAP Note 1730928: Using external software in an SAP HANA appliance
- SAP Note 1730929: Using external tools in an SAP HANA appliance
- SAP Note 1730930: Using antivirus software in an SAP HANA appliance
- SAP Note 2031547: Using backup tools with Backint for SAP HANA

SAP HANA virtualization notes:

- SAP Note 1788665: SAP HANA running on VMware vSphere virtual machines