Cisco Scale-Up Solution for SAP HANA on Cisco C880 M5 Servers with Suse Linux Enterprise Server 12 for SAP

Design and deploy a SAP HANA single-node solution based on standalone Cisco C880 M5 Servers with Suse Linux Enterprise Server 12 for SAP SP2

April 2018

Contents

Executive summary

Organizations in every industry are generating and using more data than ever before: from customer transactions and supplier delivery information to real-time user-consumption statistics. Without reliable infrastructure that can store, process, and analyze big data sets in real time, companies are unable to use this information to their advantage. The Cisco® Scale-Up Solution for SAP HANA with the Cisco Unified Computing System™ (Cisco UCS®) using the Cisco C880 M5 Server helps companies more easily harness information and make better business decisions that let them stay ahead of the competition. Our solutions help improve access to all your data to accelerate business decision making with policy-based, simplified management; lower deployment risk; and reduced total cost of ownership (TCO). Our innovations help enable you to unlock the intelligence in your data and interpret it with a new dimension of context and insight to help you gain a sustainable, competitive business advantage.

The Cisco solution for SAP HANA with the Cisco C880 M5 rack-mount server provides a robust platform for SAP HANA workloads in a single node.

Solution overview

Introduction

The Cisco C880 M5 scale-up solution provides pre-validated, ready-to-deploy infrastructure, reducing the time and complexity involved in configuring and validating a traditional data center deployment. The reference architecture detailed in this document highlights the resiliency and ease of deployment of an SAP HANA solution.

SAP HANA is SAP's implementation of in-memory database technology. The SAP HANA database takes advantage of the low-cost main memory (RAM), faster access, and data-processing capabilities of multicore processors to provide better performance for analytical and transactional applications. SAP HANA offers a multiple-engine query-processing environment that supports relational data (with both row- and column-oriented physical representations in a hybrid engine) as well as graph and text processing for semistructured and unstructured data management within the same system. As an appliance, the SAP HANA solution combines software components from SAP optimized for certified hardware. However, this solution has a preconfigured hardware setup and preinstalled software package that is dedicated to SAP HANA. In 2013, SAP introduced the SAP HANA Tailored Datacenter Integration (TDI) option. TDI offers a more open and flexible way to integrate SAP HANA into the data center by reusing existing enterprise storage hardware, thereby reducing hardware costs. With the introduction of SAP HANA TDI for shared infrastructure, the Cisco UCS Integrated Infrastructure solution provides the advantages of an integrated computing, storage, and network stack and the programmability of Cisco UCS. The TDI option enables organizations to run multiple SAP HANA production systems on a shared infrastructure. It also enables customers to run SAP application servers and the SAP HANA database hosted on the same infrastructure.

For more information about SAP HANA, see the SAP help portal: http://help.sap.com/hana/.

Audience

The intended audience for this document includes sales engineers, field consultants, professional services staff, IT managers, partner engineers, and customers deploying the Cisco solution for SAP HANA. External references are provided wherever applicable, but readers are expected to be familiar with the technology, infrastructure, and database security policies of the customer installation.

Purpose of this document

This document describes the steps required to deploy and configure a Cisco data center solution for SAP HANA. This document showcases one of the variants of Cisco's solution for SAP HANA. Although readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are important to the deployment of this solution are provided in this document.

Solution summary: Cisco C880 M5 Server

The Cisco Scale-Up Solution for SAP HANA is based on the Cisco UCS C880 M5 Server. Tables 1, 2, and 3 summarize the server specifications and show proposed disk configurations for the SAP HANA use case.

Table 1. Overview of Cisco UCS C880 M5 Server configuration

CPU specification	2.10-GHz Intel® Xeon® processor 8176 Platinum CPU Quantity: 8
Possible memory configurations	**Memory configuration** • 64-GB DDR4: Quantity 48 (3 TB) • 64-GB DDR4: Quantity 96 (6 TB) • 128-GB DDR4: Quantity 72 (9 TB) • 128-GB DDR4: Quantity 96 (12 TB)
Hard-disk drive (HDD) type and quantity	40 x 1.8-TB SAS drives (in 2 x JX40S2 JBODs)
BIOS	Release V1.0.0.0 R1.17.0 for D3858-A1x or later
Unified firmware	Release 1.14Q or later
LSI MegaRAID controller	LSI EP420e 12-Gbps SAS modular RAID controller
Network card	• Onboard Intel 1 Gigabit Ethernet controller (2 ports) • PCIe Intel 10 Gigabit Ethernet controller (8 ports)
Power supply	Redundant power supplies: Quantity 4

Table 2. Cisco C880 M5 proposed disk layout

Disk	Disk type	Drive group	RAID level	Virtual drive
JBOD 1: Slot (0 through 4)	SAS HDD	DG0	5	VD0
JBOD 1: Slot (5 through 9)	SAS HDD	DG1	5	VD1
JBOD 1: Slot (10 through 14)	SAS HDD	DG2	5	VD2
JBOD 1: Slot (15 through 19)	SAS HDD	DG3	5	VD3
JBOD 2: Slot (0 through 4)	SAS HDD	DG4	5	VD4
JBOD 2: Slot (5 through 9)	SAS HDD	DG5	5	VD5
JBOD 2: Slot (10 through 14)	SAS HDD	DG6	5	VD6
JBOD 2: Slot (15 through 19)	SAS HDD	DG7	5	VD7

Table 3. Cisco C880 M5 proposed disk configuration

Drives used	RAID type	Used for	File system
40 x 1.8-TB SAS HDD	RAID 5	Operating system	ext3
		Data file system	XFS
		Log file system	XFS
		HANA shared file system	XFS

Infrastructure overview

Cisco C880 M5 Server

The scalable Cisco C880 M5 Server (Figure 1) is an Intel-based rack server for critical business scenarios. For example, it is well suited for database management systems for medium-sized and large databases and as a consolidation platform for running an immensely large number of different applications using virtualization technologies. With its highly developed hardware and software components, the server offers a high level of data security and availability. Security and availability features include hot-pluggable HDD and solid-state disk (SSD) modules, hot-pluggable system fans and power supply units, and automatic server reconfiguration and restart (ASR&R). It is well suited for the following use cases:

- In-memory databases
- Big data analytics
- Virtualization and virtual desktop infrastructure (VDI) workloads
- Bare-metal applications

Figure 1. Cisco C880 M5 Server

Front view of the C880

Front-view detail

Rear view of the C880

Rear-view detail

Solution design

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

SAP HANA system

An 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 least one 1 Gigabit Ethernet access network and one 10 Gigabit Ethernet storage network.

Hardware requirements for the SAP HANA database

SAP defines hardware and software requirements for running SAP HANA systems. For the latest information about the CPU and memory configurations supported for SAP HANA, please refer to https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html.

Note: This document does not cover the updated information published by SAP. Additional information is available at http://saphana.com/.

File system layout

Figure 2 shows the file system layout and the required storage sizes for installing and operating SAP HANA. When installing SAP HANA on a host, specify the mount point for the installation binaries (/hana/shared/<sid>), data files (/hana/data/<sid>), and log files (/hana/log/<sid>), where sid is the instance identifier of the SAP HANA installation.

Figure 2. Proposed disk layout with partition mapping for forty 1.8-TB SAS drives

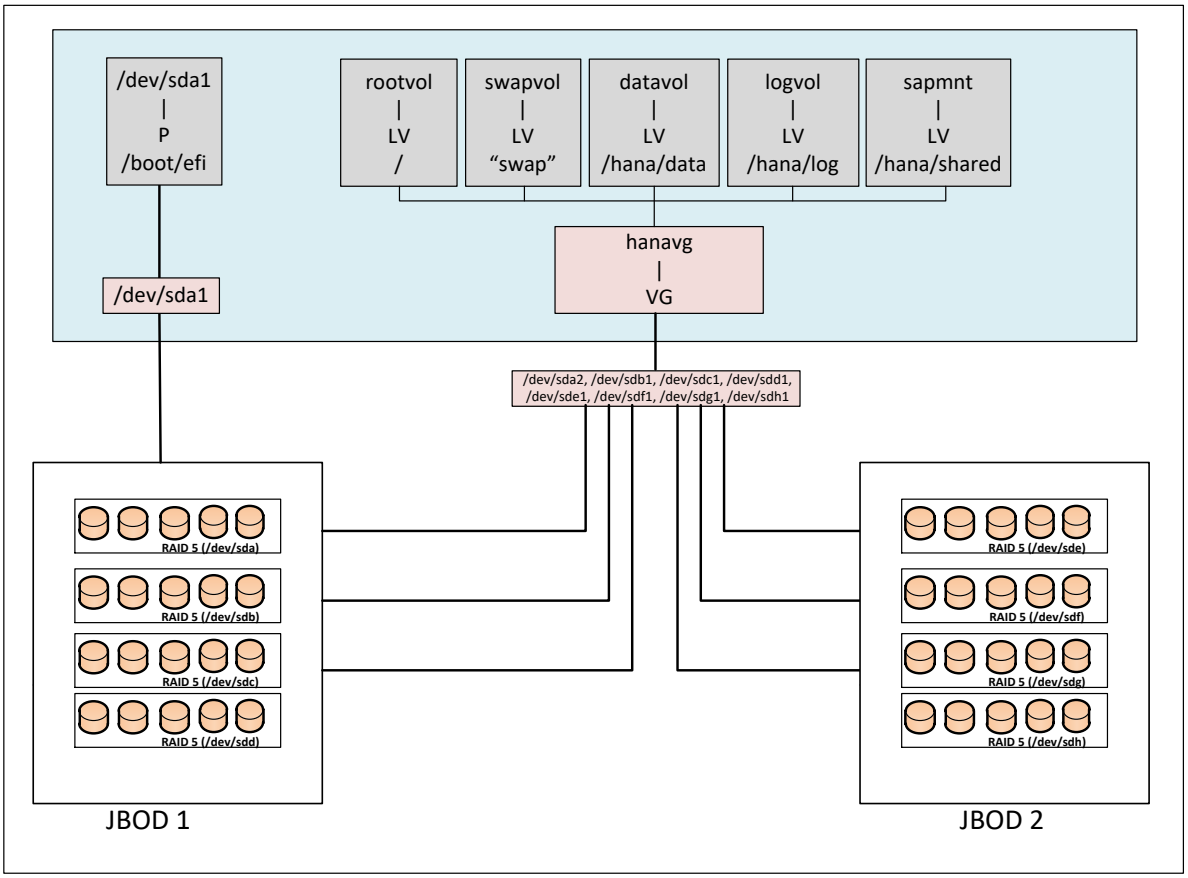

The storage size for the file system is based on the amount of memory on the SAP HANA host.

The following list shows sample file system sizes for a single-node system with 3 TB of memory:

- /hana/shared: 1 x memory (3 TB)
- /hana/data: 3 x memory (9 TB)
- /hana/log: 1 x memory (512 GB)*

* For solutions based on the Intel Xeon processor Platinum CPU, the size of the log volume must be as follows:

  - Half of the server memory for systems ≤ 256 GB of memory
  - Minimum of 512 GB for systems with ≥ 512 GB of memory

Operating system

SAP HANA supports the following operating systems:

- SUSE Linux Enterprise Server (SLES) for SAP applications
- Red Hat Enterprise Linux (RHEL) for SAP applications

Note: This document provides installation steps for SLES for SAP 12 SP2.

Deployment hardware and software

Configuration guidelines

This section is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 4 lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the configuration steps presented in this document.

Table 4. Configuration variables

Variable	Description	Customer implementation value
<<var_cimc_ip_address>>	Cisco C880 M5 Server's iRMC IP address	
<<var_cimc_ip_netmask>>	Cisco C880 M5 Server's iRMC network netmask	
<<var_cimc_gateway_ip>>	Cisco C880 M5 Server's iRMC network gateway IP address	
<<var_raid5_vd0-7_name>>	Name for virtual drive VD0 to VD7 during RAID configuration	
<<var_hostname.domain>>	SAP HANA node's fully qualified domain name (FQDN)	
<<var_sys_root-pw>>	SAP HANA node's root password	
<<var_lvm_vg_name>>	SAP HANA node's OS logical volume management (LVM) volume group name	
<<var_mgmt_ip_address>>	SAP HANA node's management and administration IP address	
<<var_mgmt_nw_netmask>>	SAP HANA node's management network netmask	
<<var_mgmt_gateway_ip>>	Cisco C880 M5 Server's management and administrative network gateway IP address	
<<var_mgmt_netmask_prefix>>	Netmask prefix in Classless Inter-Domain Routing (CIDR) notation	

Preparing the SAP HANA scale-up node

Configuring the Cisco C880 M5 integrated remote management controller

To configure the integrated remote management controller (iRMC), connect to the server with cables to the VGA and USB ports for the keyboard.

1. Connect the management network cable to the management port behind the server as shown in Figure 3.

Figure 3. Connecting to the management port

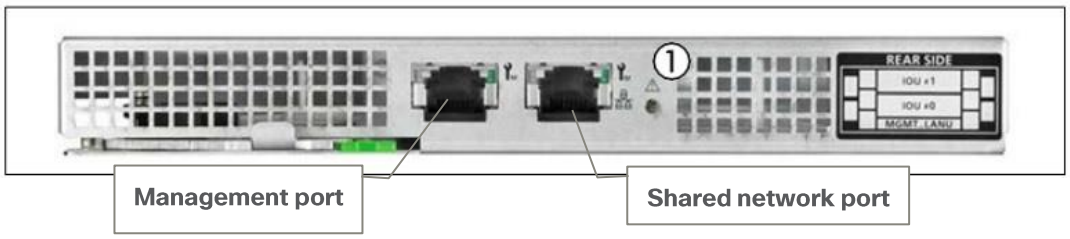

2. Power on the server and wait for the BIOS post screen.
3. When the BIOS post screen appears, press F2 for the iRMC BIOS setup (Figure 4).

Figure 4. BIOS post screen

4. On the BIOS setup screen that appears, select the Management tab (Figure 5).

Figure 5. Cisco C880 M5 BIOS setup screen (local display)

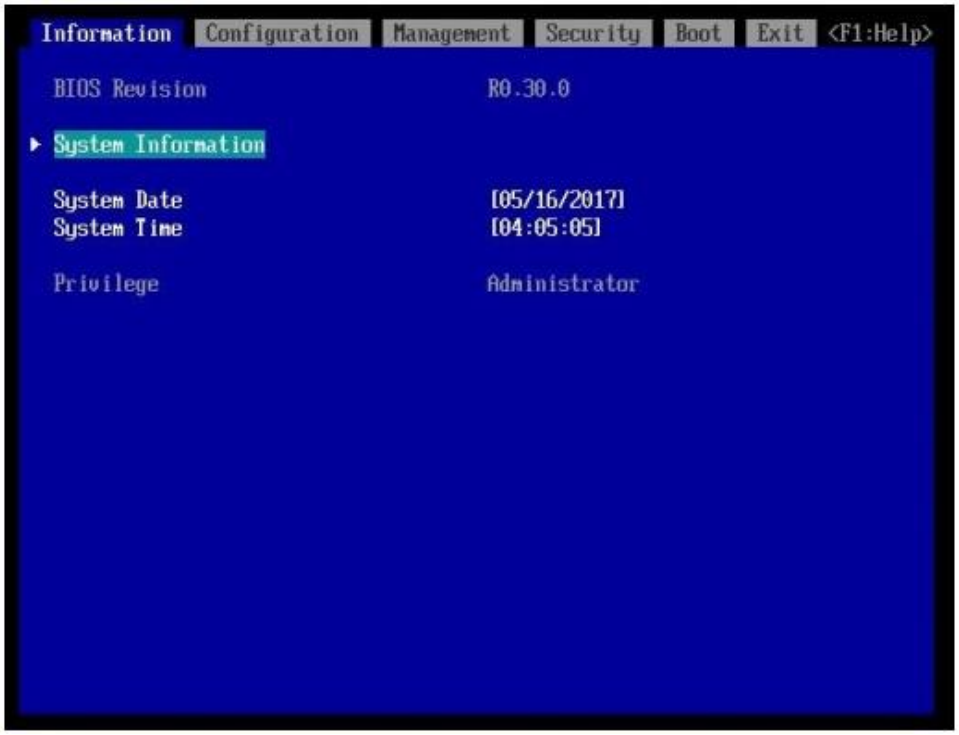

5. On the Management tab, configure the management interface IP address (Figure 6).

Figure 6. Cisco C880 M5 management interface configuration (local display)

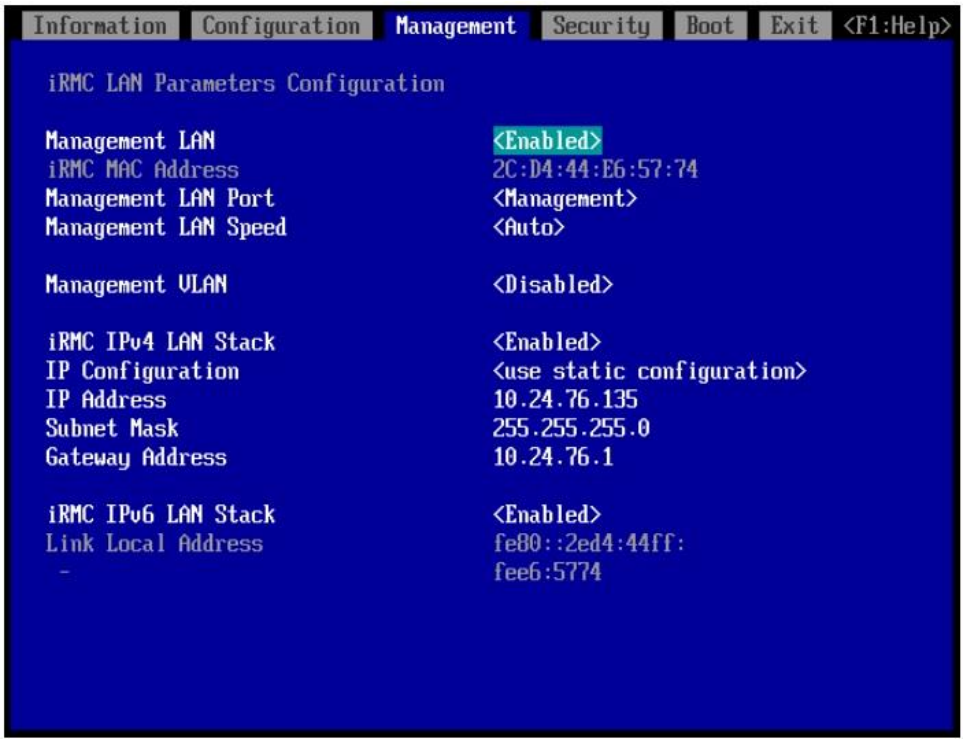

6. Make sure that Management is selected for Management LAN Port.
7. Use the console network IP address **<<var_iRMC_ip_address>>**, netmask **<<var_ iRMC _ip_netmask>>**, and gateway **<<var_ iRMC _gateway_ip>>** for the IPv4 settings of the iRMC.
8. Press F10 to save the configuration and exit the utility.
9. Open a web browser on a computer on the same network with Java and Adobe Flash installed.
10. Enter the iRMC IP address of the Cisco C880 M5 server: **http://<<var_ iRMC _ip_address>>**.
11. Enter the login credentials as updated in the Cisco Integrated Management Controller (IMC) configuration. The default username and password are **admin** and **admin** (Figure 7).

Figure 7. Logging in

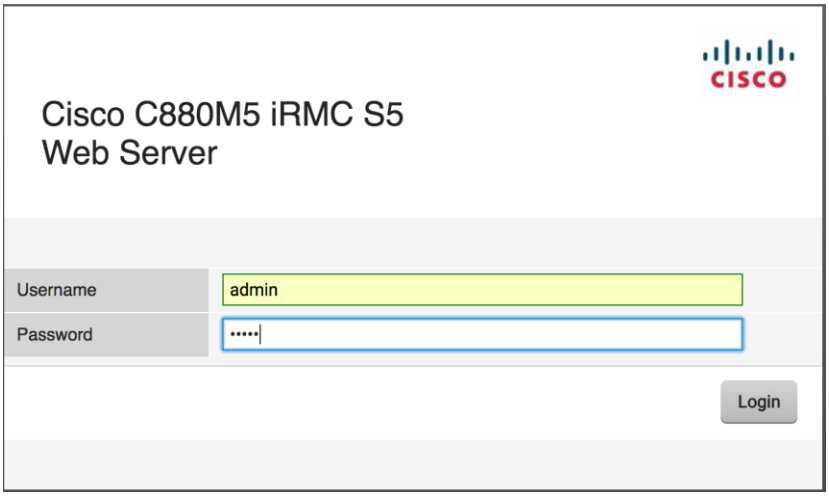

The home screen for the controller appears (Figure 8).

Figure 8. Cisco C880 IMC home screen

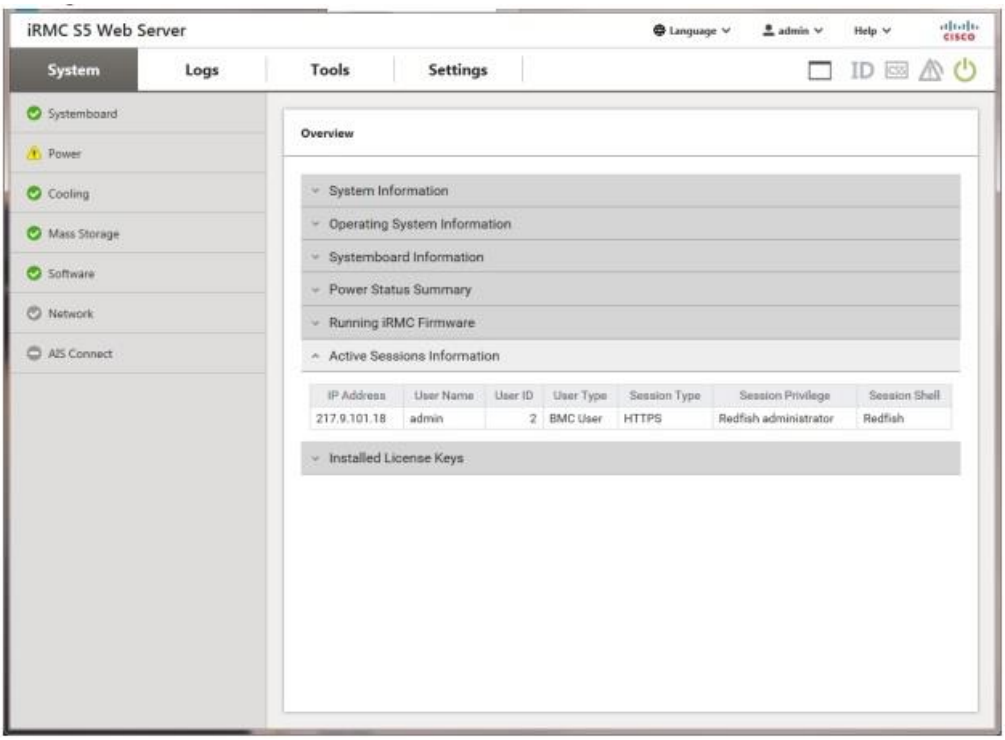

Starting Cisco C880 M5 video redirection

Follow the steps here to configure video redirection to access the console window from the browser.

1. Log in to the iRMC with the admin credentials.
2. Click the monitor icon at the upper right (Figure 9).

Figure 9. Clicking the monitor icon

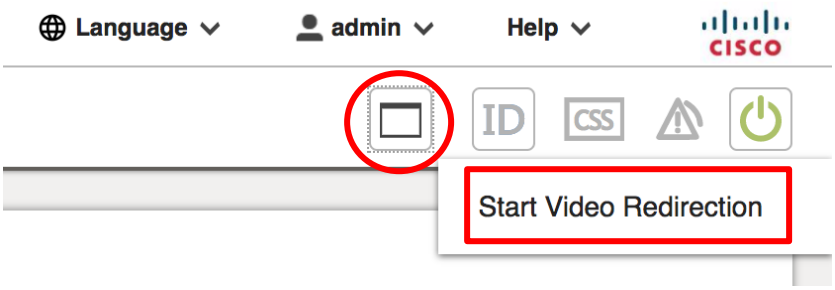

3. Click Start Video Redirection. This step opens a new Java window by granting the video console access (Figure 10).

Figure 10. Opening the video console window

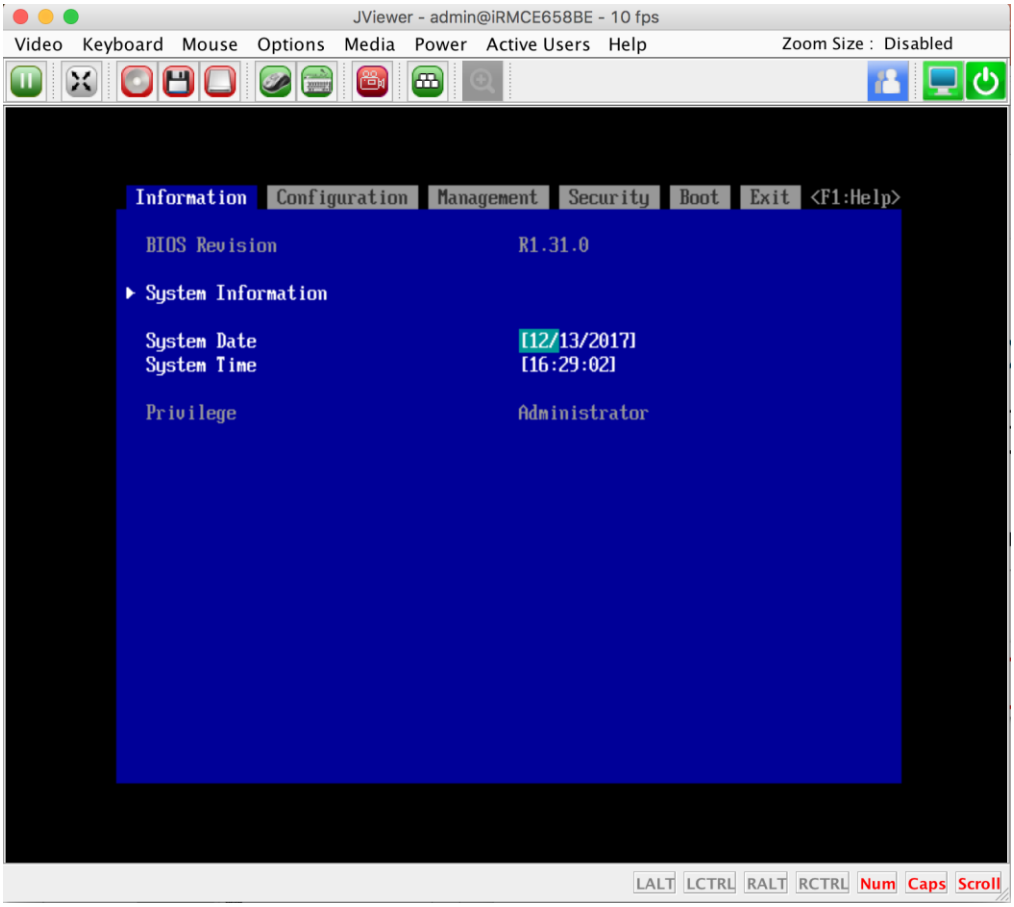

Configuring Cisco C880 M5 BIOS CPU settings

After the video console window opens, follow the steps here to configure the CPU settings.

1. From the top menu bar, select Configuration and then select CPU Configuration in the submenu (Figure 11).

Figure 11. Selecting CPU Configuration

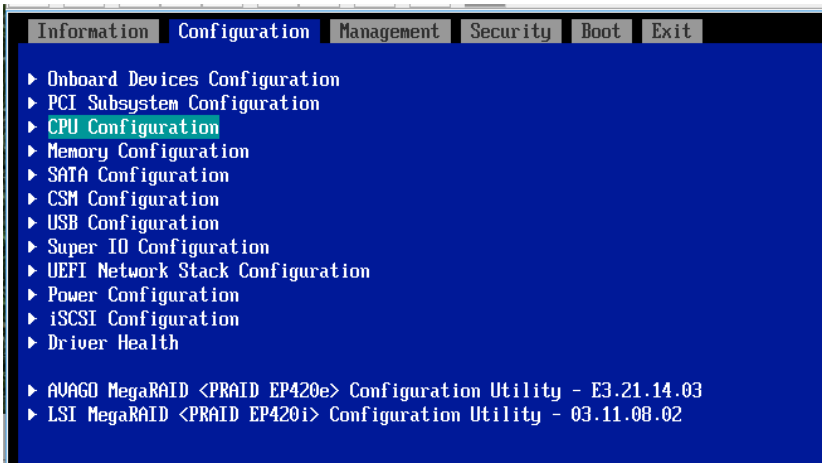

2. Configure the CPU settings as shown in Figure 12.

Figure 12. CPU configuration settings

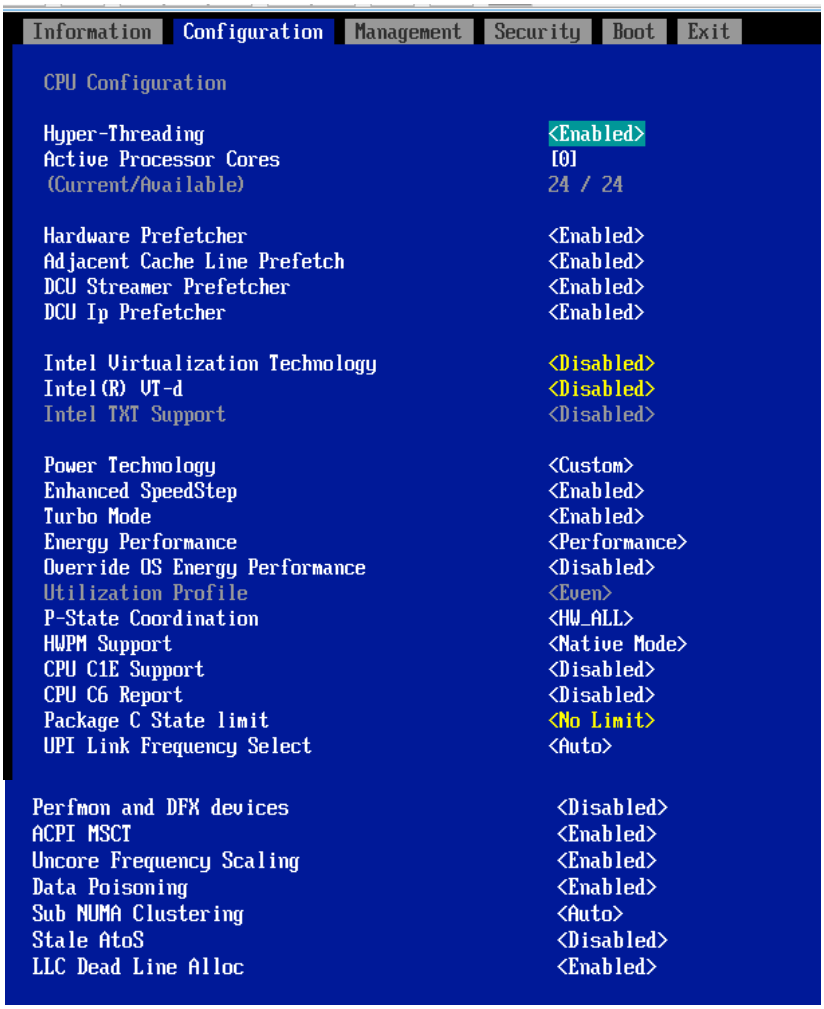

Configuring Cisco C880 M5 BIOS disk (RAID) settings for the SAP HANA file system

Now create eight virtual drives with RAID 5. Each virtual drive will have five drives associated with it. Follow the steps here until all eight virtual drives are created.

Be sure that the settings for the virtual drives are identical. If they are not, SAP HANA performance will be degraded.

Note: The steps in this section create a single RAID 5 virtual drive. Repeat these steps to create eight RAID 5 virtual drives for the appliance installation.

1. Select the Avago MegaRAID configuration utility on the Configuration screen (Figure 13).

Figure 13. Selecting the configuration utility

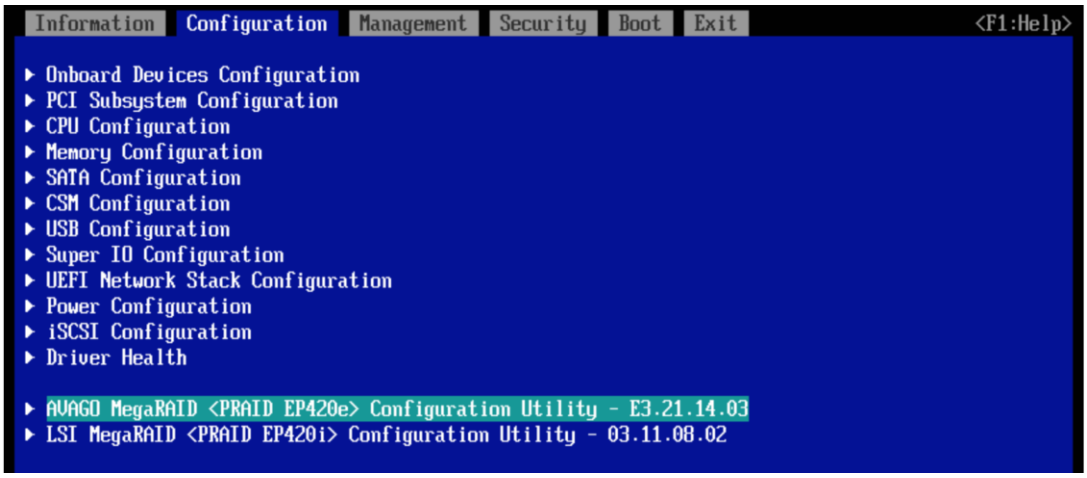

2. Select Main Menu (Figure 14).

Figure 14. Selecting the main menu

3. From the main menu, select Configuration Management (Figure 15). You will configure the first virtual drive.

Figure 15. Cisco C880 M5 RAID configuration 1

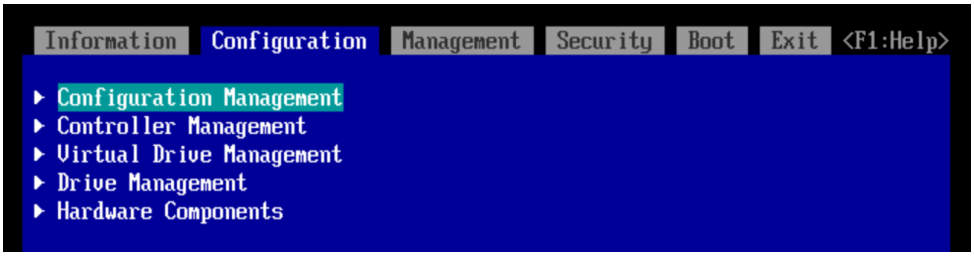

4. Select Create Virtual Drive (Figure 16).

Figure 16. Creating the virtual drive

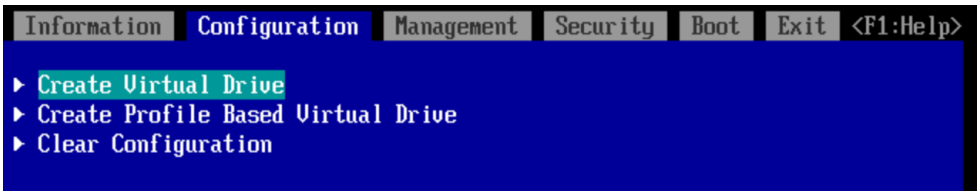

5. Select RAID5 as the RAID level and click Select Drives (Figure 17).

Figure 17. Selecting drives

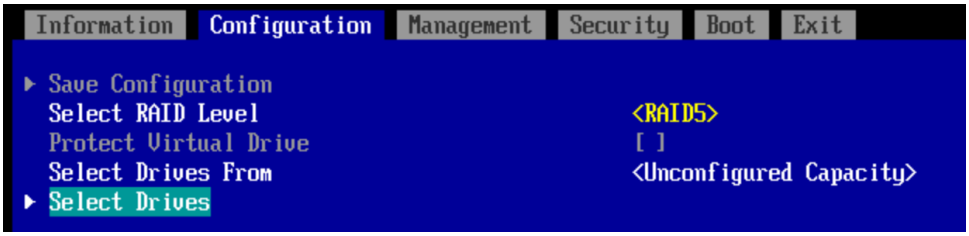

6. Select the first five drives in the list (Figure 18).

Figure 18. Selecting drives for the virtual drives

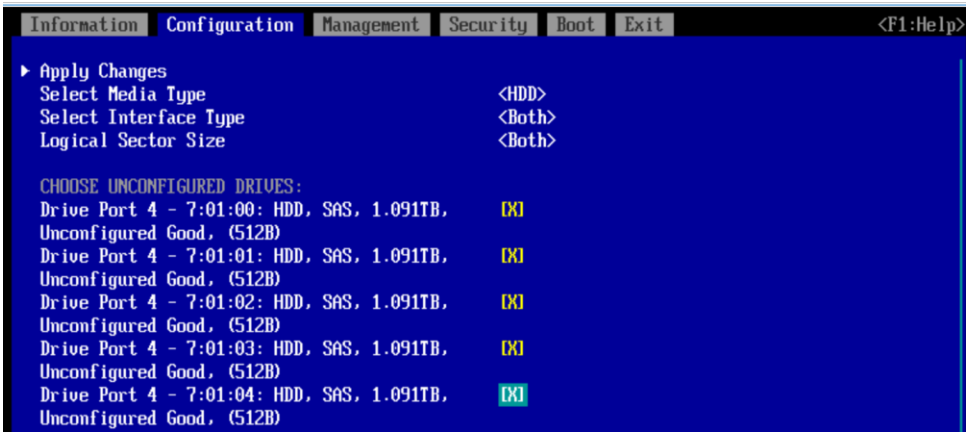

7. Select Apply Changes from the menu above the drive list.
8. Assign a virtual drive name and configure the parameters as shown in Figure 19. These parameters are very important. They will manage the performance of the SAP HANA that will be installed. Hence, be sure to set the correct values.
   a. For Strip Size, specify 256 KB.
   b. For Read Policy, specify Read Ahead.
   c. For Write Policy, specify Write Back.
   d. For I/O Policy, specify Direct.
   e. For Drive Cache, specify Unchanged.

Figure 19. Configuring virtual drive parameters

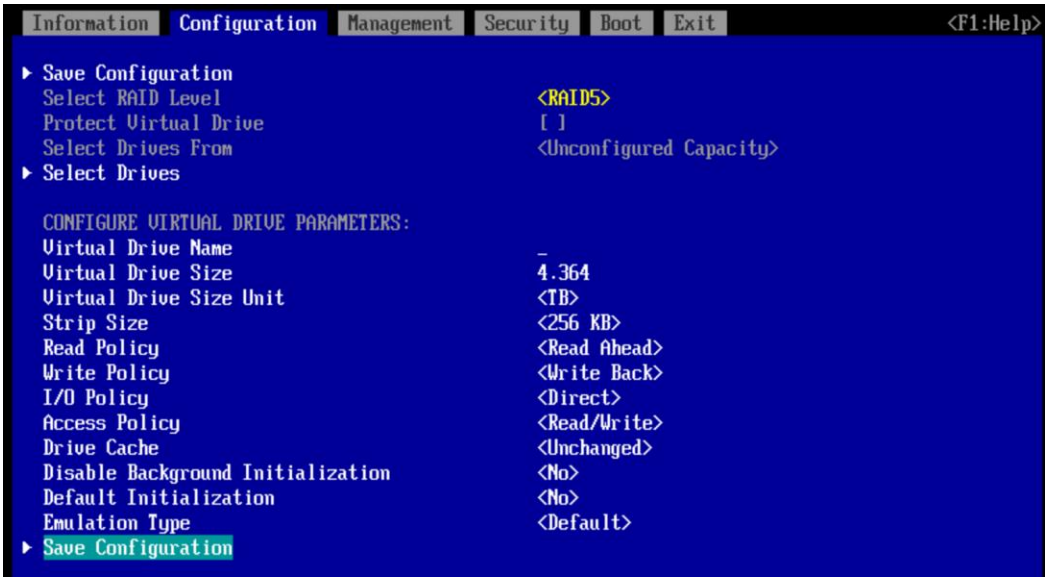

9. Click Save Configuration.
10. Select Yes to confirm the save operation (Figure 20).

Figure 20. Saving the configuration

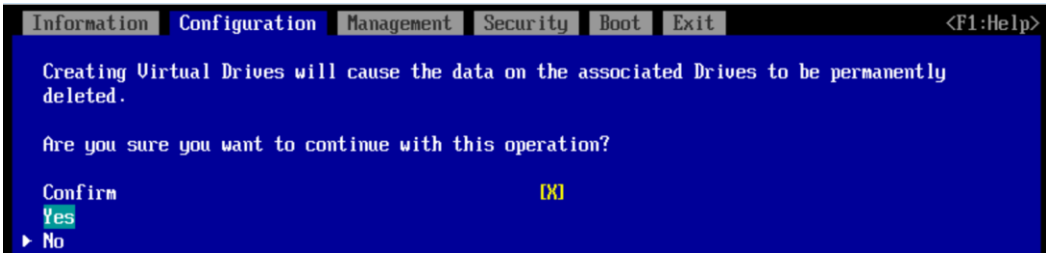

Note: You need to create eight virtual drives, each with five physical disks, so repeat the steps here until all eight virtual drives are created.

Committing the changes to the BIOS

After the BIOS settings and RAID configurations are applied, you need to commit the changes to the BIOS and exit the BIOS utility to proceed with the OS installation.

1. From the top menu, select Exit.
2. Select “Commit settings and Exit” (Figure 21).

Figure 21. Committing the settings and exiting

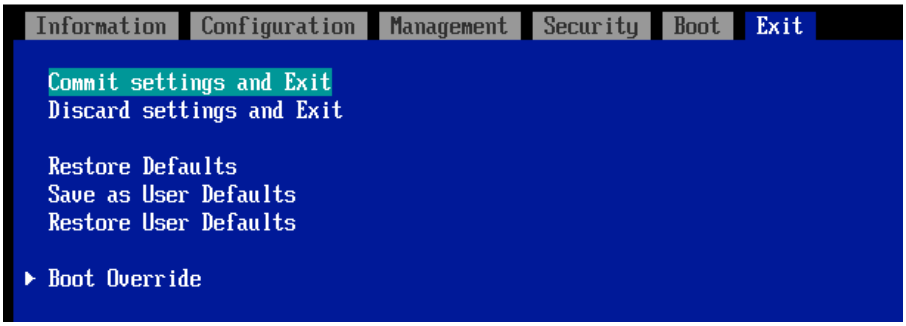

3. The server will reboot. You can now mount the operating system ISO file on the console.

Installing the operating system

The following procedure shows the steps for installing SLES 12 for SAP SP2 on local drives. Keep the SLES DVD ISO file handy.

1. In the video console window, click the CD icon (Figure 22).

Figure 22. Clicking the CD icon

2. Browse for the SLES 12 for SAP Applications SP2 DVD ISO image and click Connect CD/DVD (Figure 23).

Figure 23. Connecting the system

3. After the system boots, press F12 to display the boot menu (Figure 24).

Figure 24. Boot menu

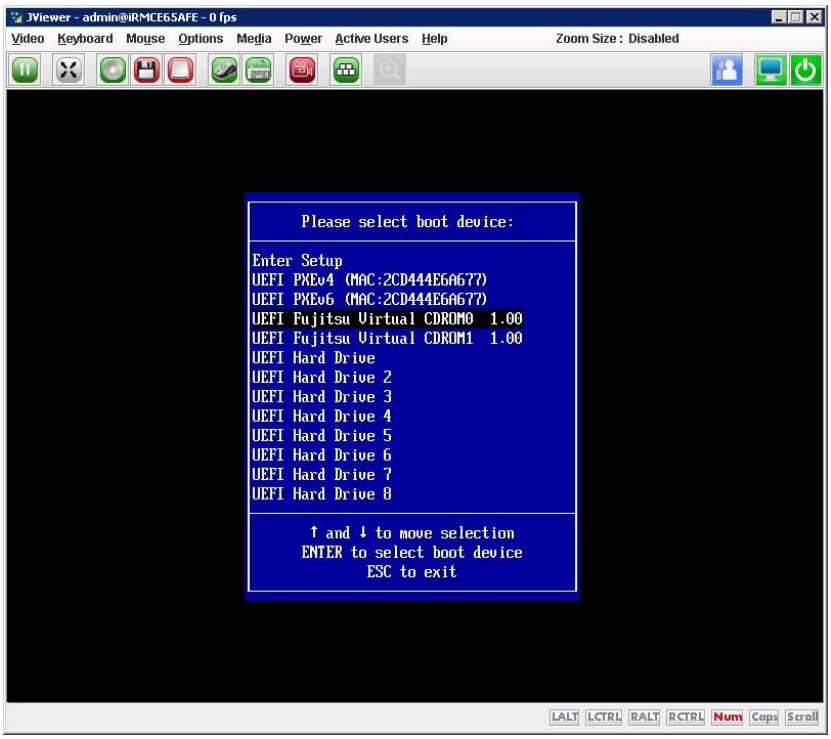

4. Choose UEFI Fujitsu Virtual CDROM0.
5. The ISO will start booting and display the Installation menu. Choose the Installation option from the menu (Figure 25).

Figure 25. Starting the installation process

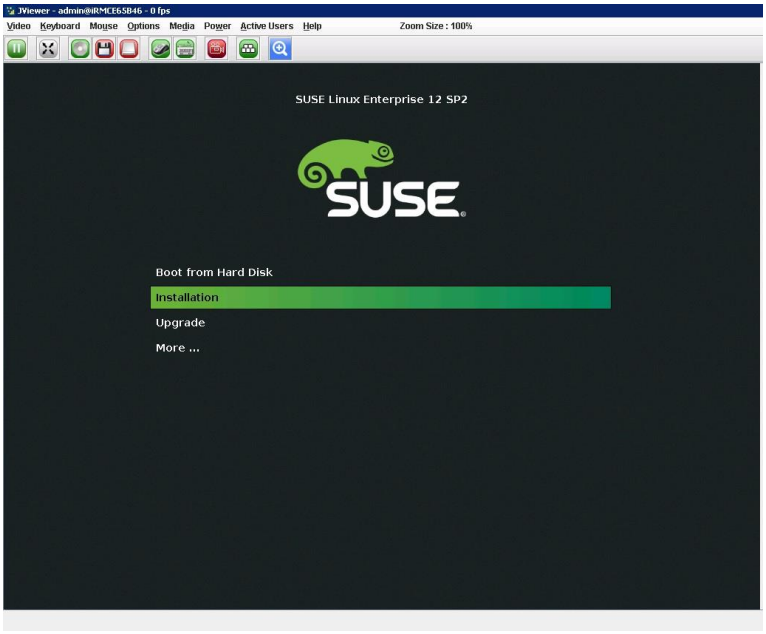

6. On the Language, Keyboard and License Agreement page, select the language of your choice and keyboard layout, select "Agree to license terms," and click Next.
7. On the Network Settings page, click Next. You will return to the network configuration as part of the post-installation tasks.
8. On the Registration page, select Skip Registration. You will register later as part of the post-installation tasks.
9. On the Product Installation Mode page, select the "Proceed with standard SLES for SAP Applications installation" option (Figure 26).

Figure 26. Product Installation Mode page

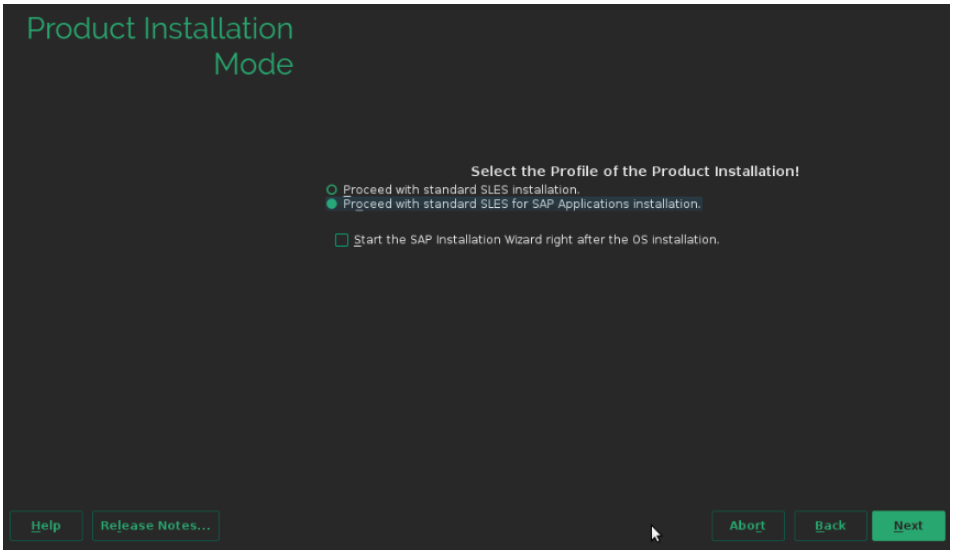

10. On the Add-On Products page, click Next. You will not install any additional add-on products.
11. You need to partition the disk. Review the Suggested Partitioning page, which contains an initial proposal for the partitions (Figure 27)

Figure 27. Suggested Partitioning page with initial partitioning proposal

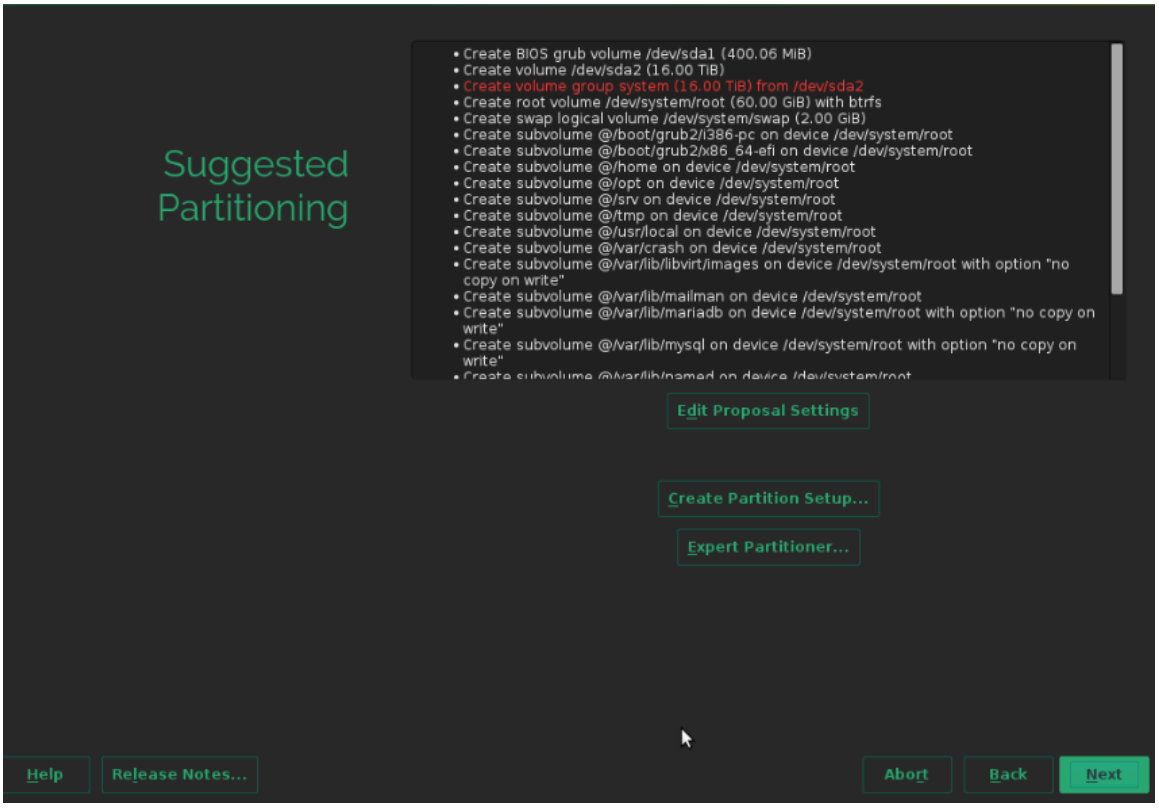

12. Click Expert Partitioner and on the left choose System View > Linux > Hard Disks > sda.
13. Clear the suggested partitions if there are any. The hard drive list should look like Figure 28.

Figure 28. Expert Partitioner drives list

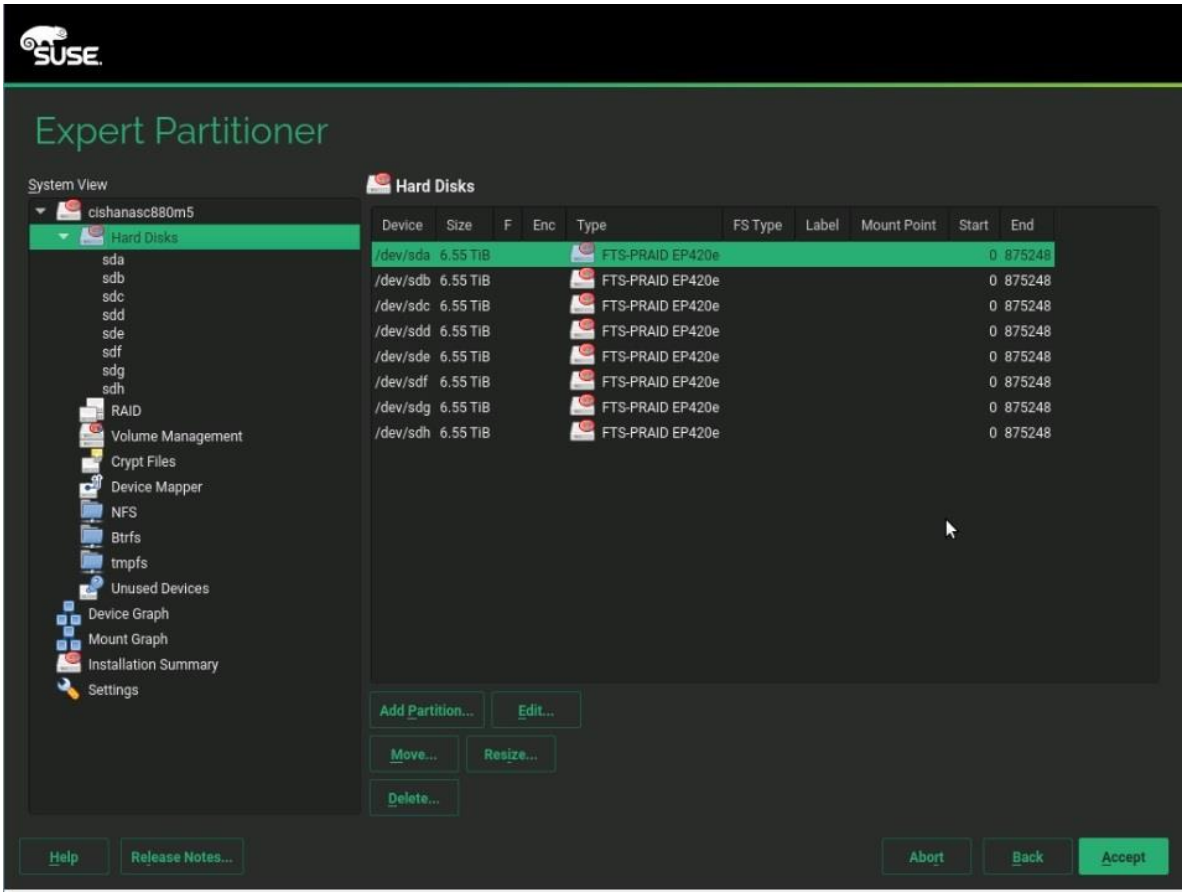

14. From the unpartitioned device sda, do the following:
    a. Create a 200-MB /boot/efi partition (/dev/sda1) from the disk device available (/dev/sda).
    b. Create another partition (/dev/sda2) assigning the rest of the available space in the device (/dev/sda). Assign this partition to the Linux LVM, thus making it a physical volume.
    c. Create a volume group (hanavg) and assign the eight available physical volumes (/dev/sda2, /dev/sdb, /dev/sdc, /dev/sdd, /dev/sde, /dev/sdf, /dev/sdg, and /dev/sdh) to it.
    d. Create a logical volume for the / file system with a size of 100 GB and formatted for ext3.
    e. Create a swap volume with a size of 2 GB.

15. On the Expert Partitioner page, choose device /dev/sda and click Add Partition (Figure 29).

Figure 29. Adding a new partition

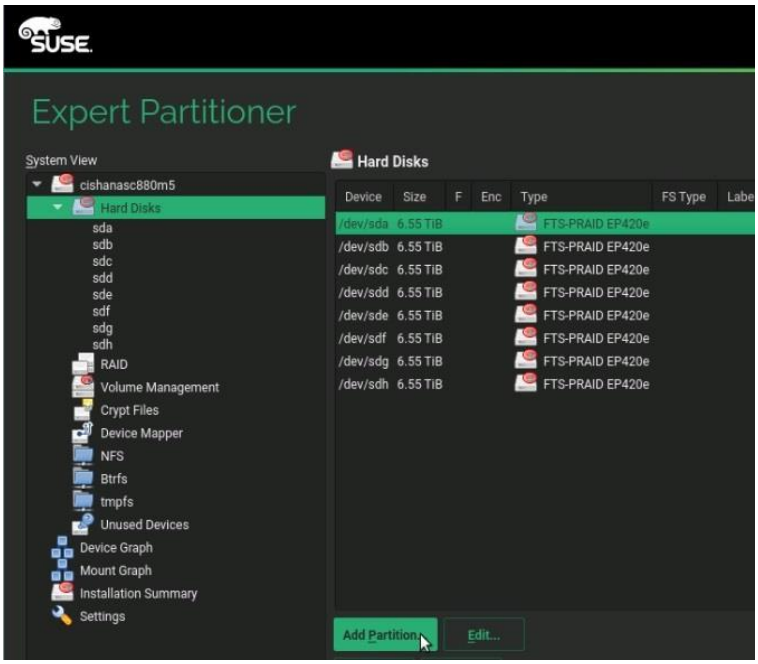

16. Create a partition with a size of 200 MB for /boot/efi (Figure 30).

Figure 30. Adding a partition: Specifying a new partition size

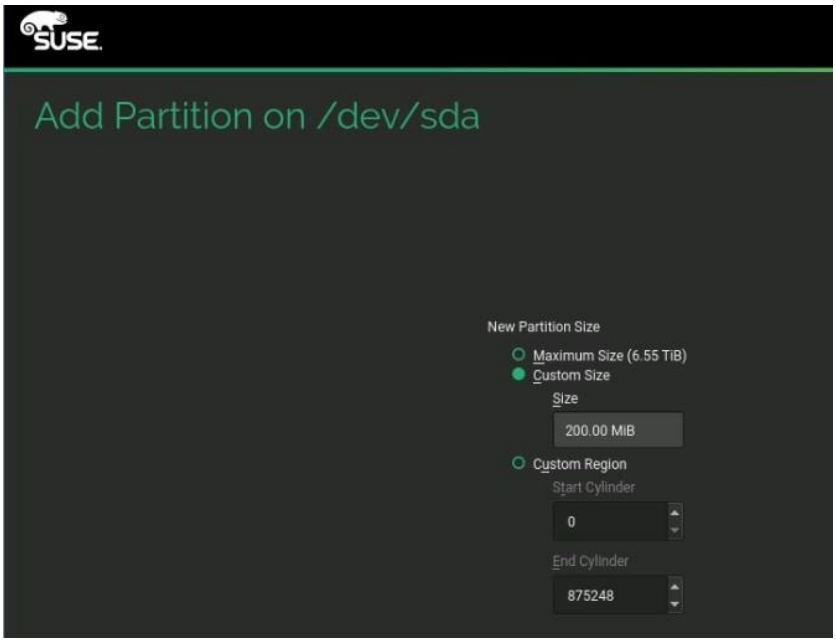

17. Click Next. Select EFI Boot Partition Role (Figure 31).

Figure 31. Adding a partition: Specifying the role

18. Click Next. You will now set the formatting and mounting options (Figure 32).

Figure 32. Adding a partition: Specifying formatting and mounting options

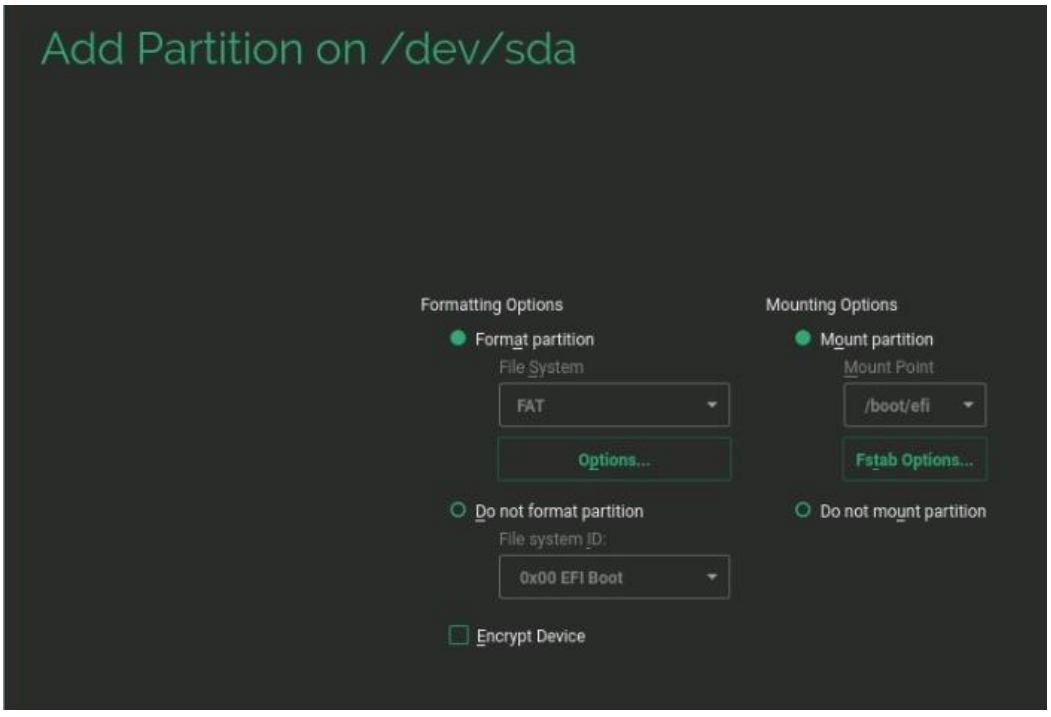

19. Choose the FAT file system and choose /boot as the mount point. Click Finish.
20. Click “Add another partition.” Specify the partition size (Figure 33).

Figure 33. Adding a partition: Specifying the partition size

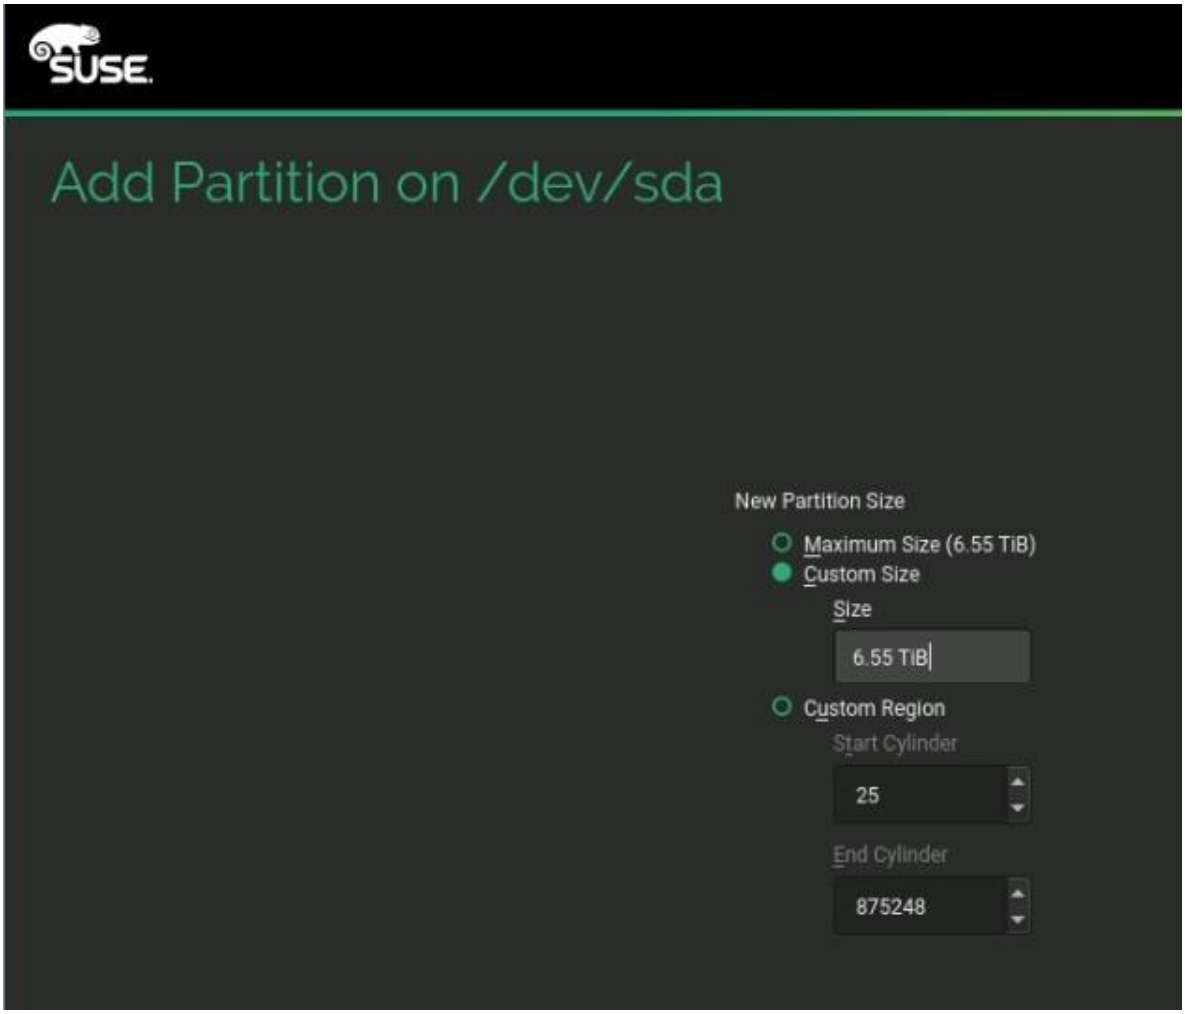

21. Allocate the rest of available space for the partition.
22. Click Next (Figure 34).

Figure 34. Adding a partition: Specifying the role

23. Assign the partition with file system ID 0x8E Linux LVM (Figure 35).

Figure 35. Adding a partition: Specifying formatting and mounting options

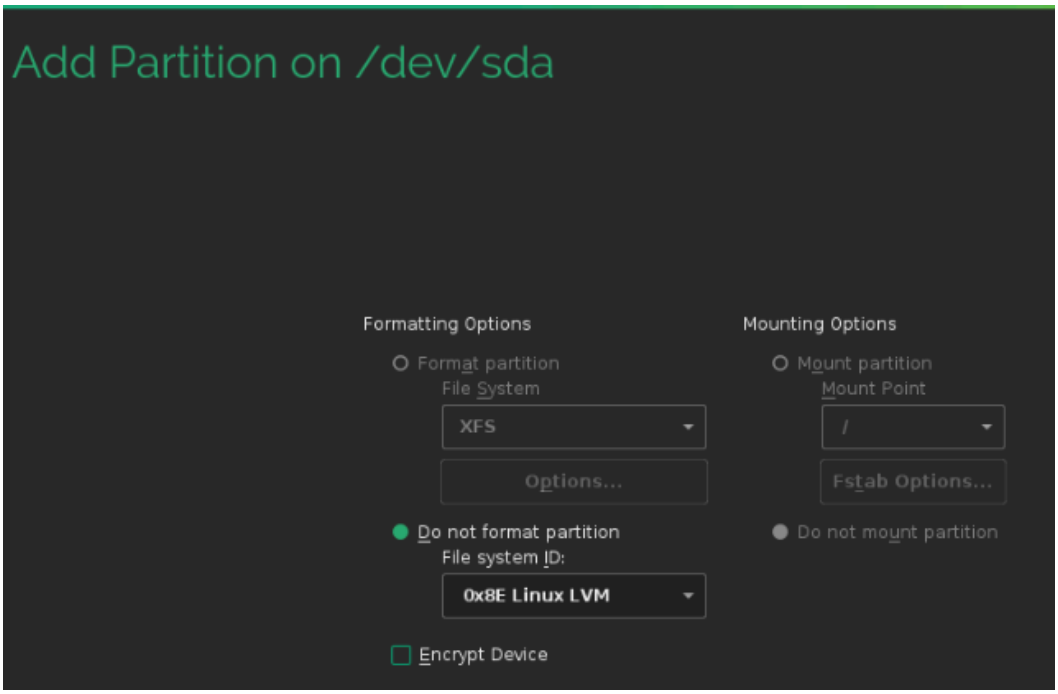

24. Click Finish.
25. Next, add a volume group. Select Volume Management in the System View pane at the left. Choose Add > Volume Group (Figure 36).

Figure 36. Expert Partitioner: Volume management > Add > Volume Group

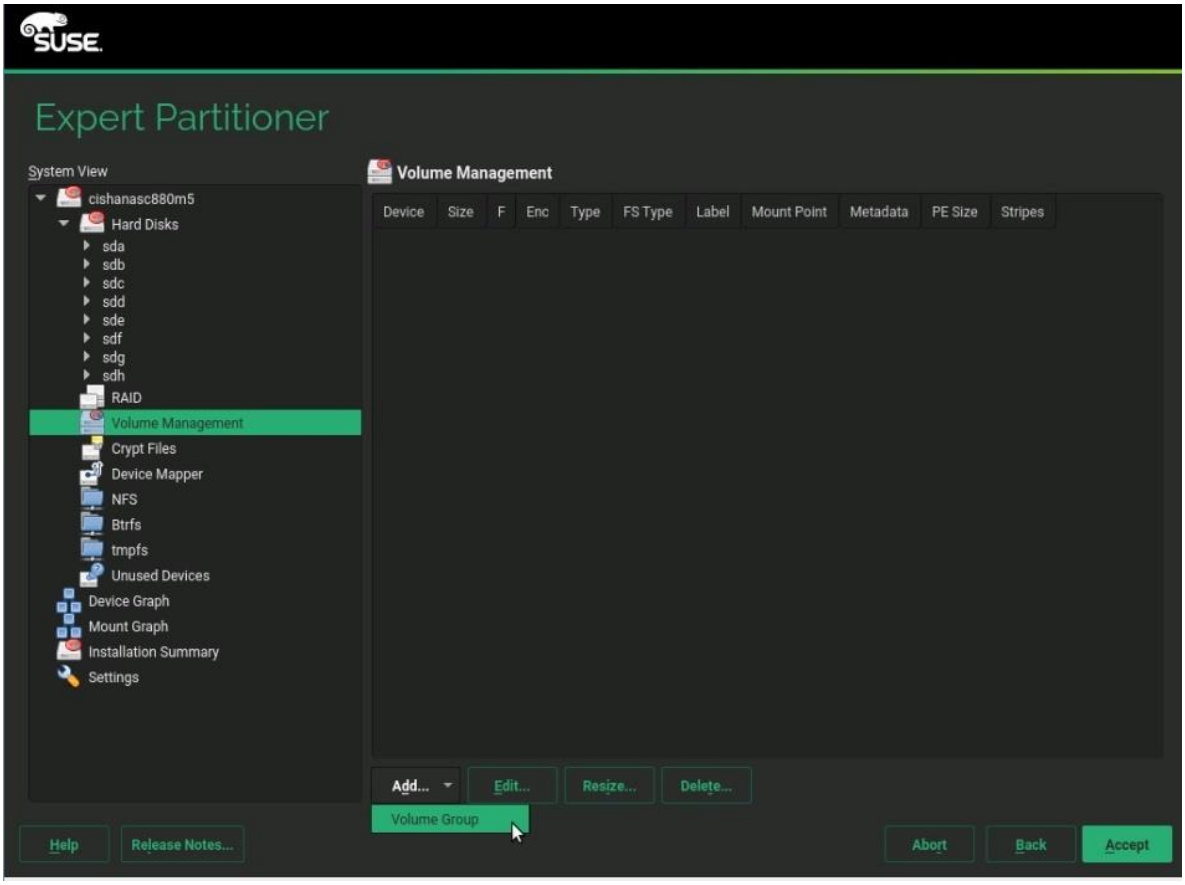

26. Provide a name for the volume group, select all the available physical volumes, and click Add All (Figure 37).

Figure 37. Add Volume Group: Selecting available physical volumes

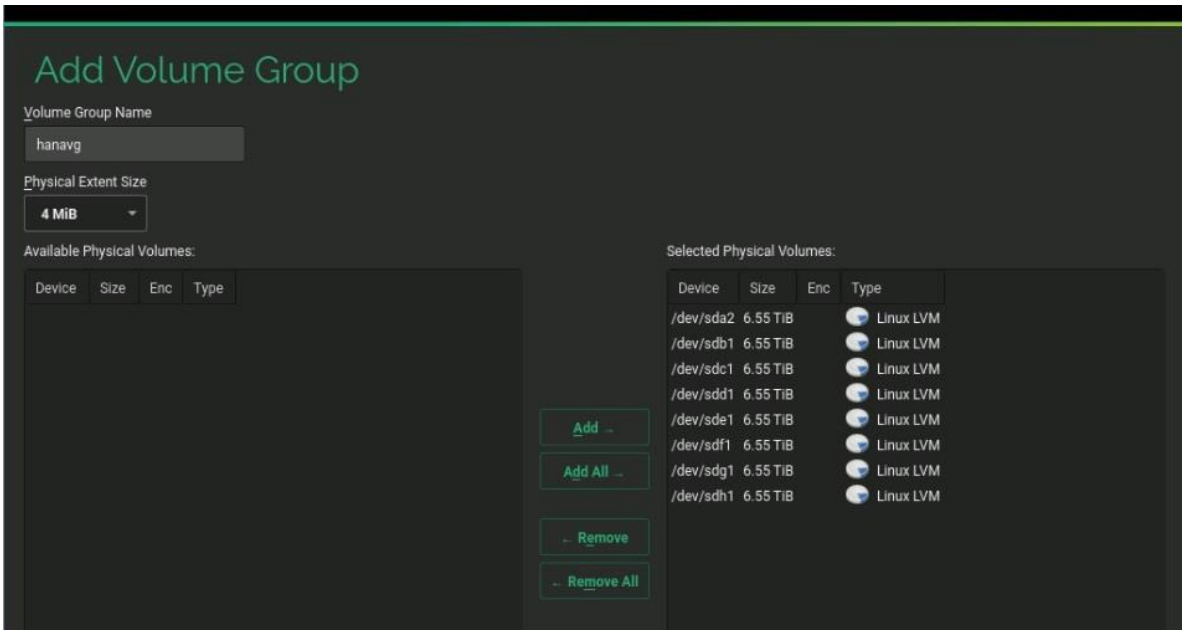

27. Click Finish.
28. Next, add a logical volume. Under Volume Management, click Add and select Logical Volume (Figure 38).

Figure 38. Expert Partitioner: Adding a logical volume

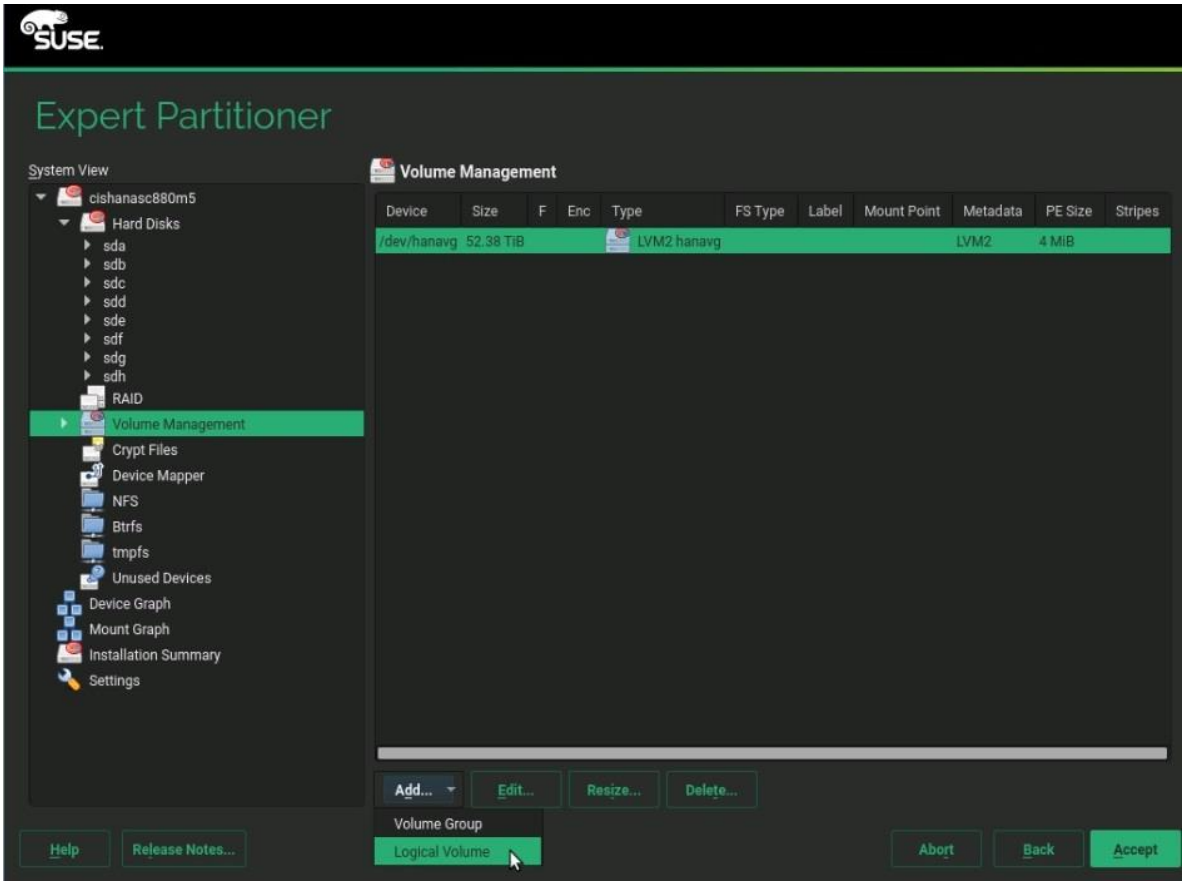

29. Add a logical volume for / with the name **rootvol** in the volume group (Figure 39).

Figure 39. Add Logical Volume: Specifying the name

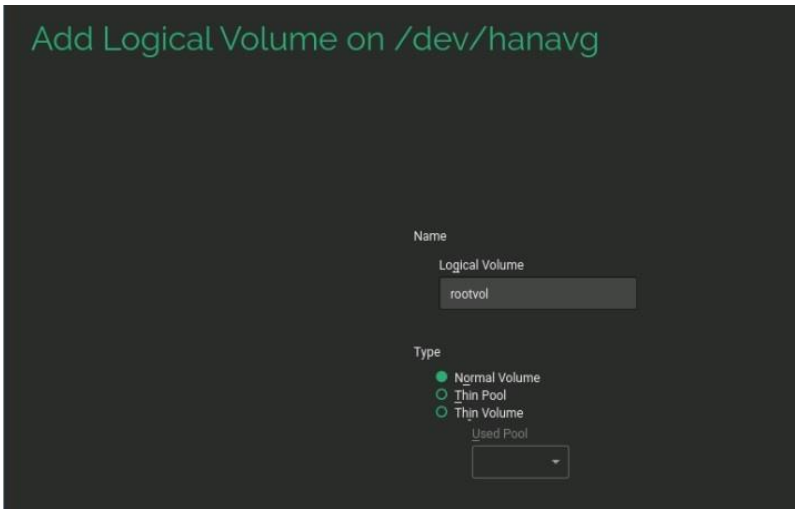

30. Click Next. Specify a size of 100 GB and a stripe of 256 KB (Figure 40).

Figure 40. Add Logical Volume: Specifying the size and stripe

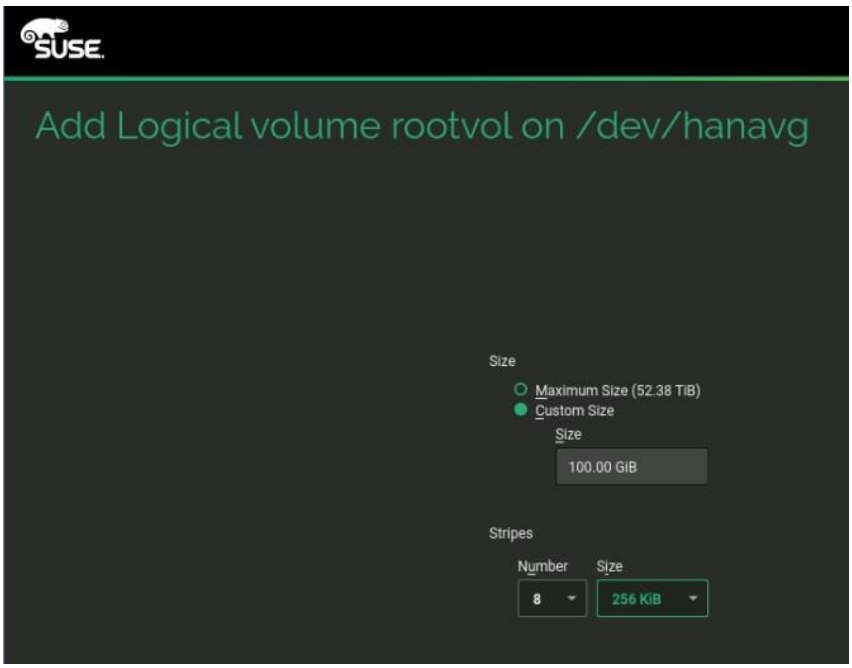

31. Click Next. Specify the operating system as the role (Figure 41).

Figure 41. Add Logical Volume: Specifying the role

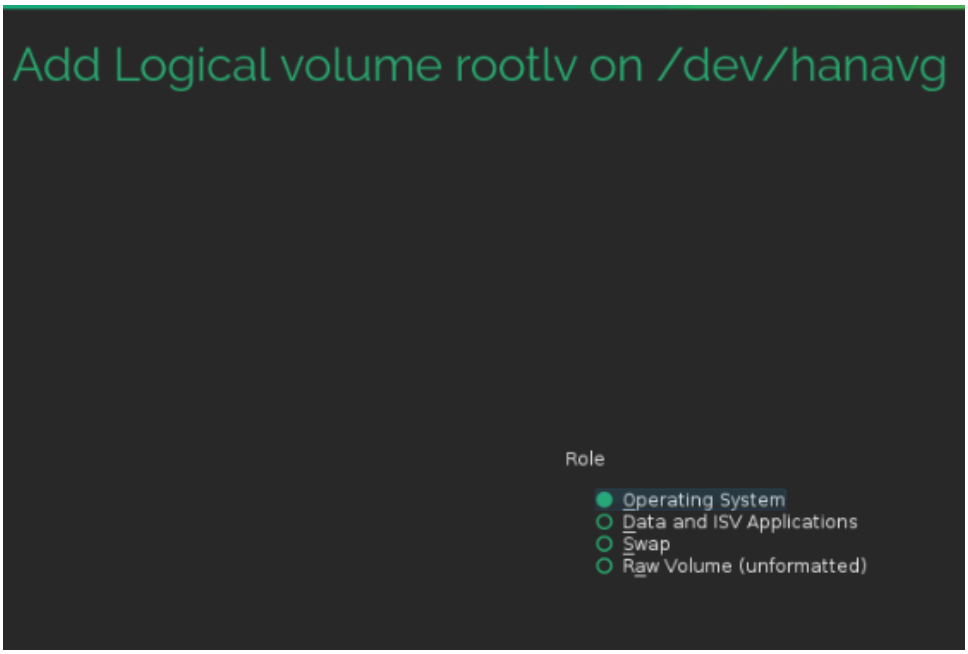

32. Click Next. Format the 100-GB logical volume rootlv for the ext3 file system and assign "/" as the mount point (Figure 42).

Figure 42. Add Logical Volume: Specifying formatting and mounting options

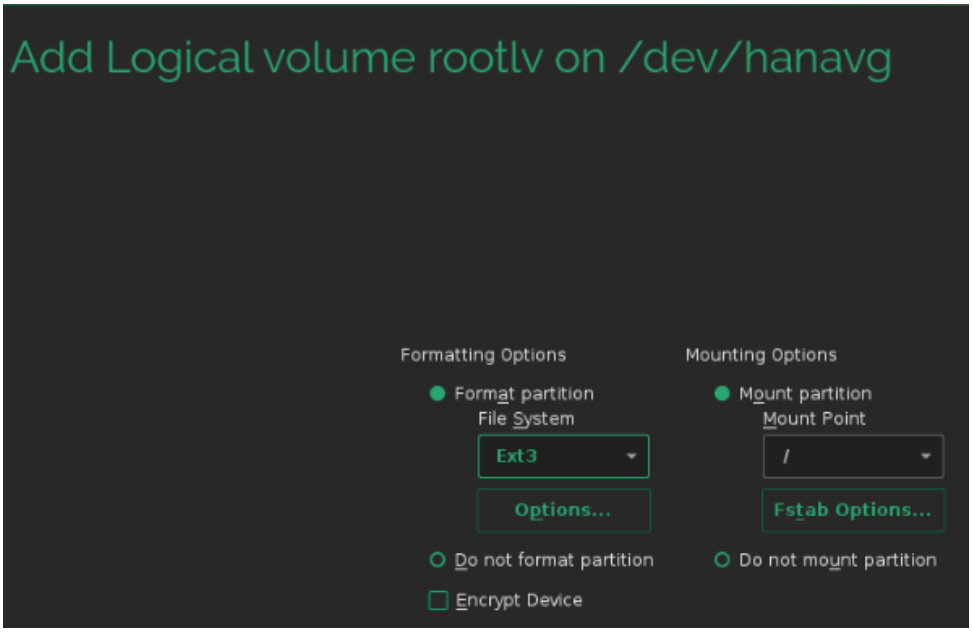

33. Click Finish.
34. Next create a swap volume with a size of 2 GB. Under Volume Management, click Add and choose Logical Volume (Figure 43).

Figure 43. Expert Partitioner: Choosing Volume Management and adding another logical volume

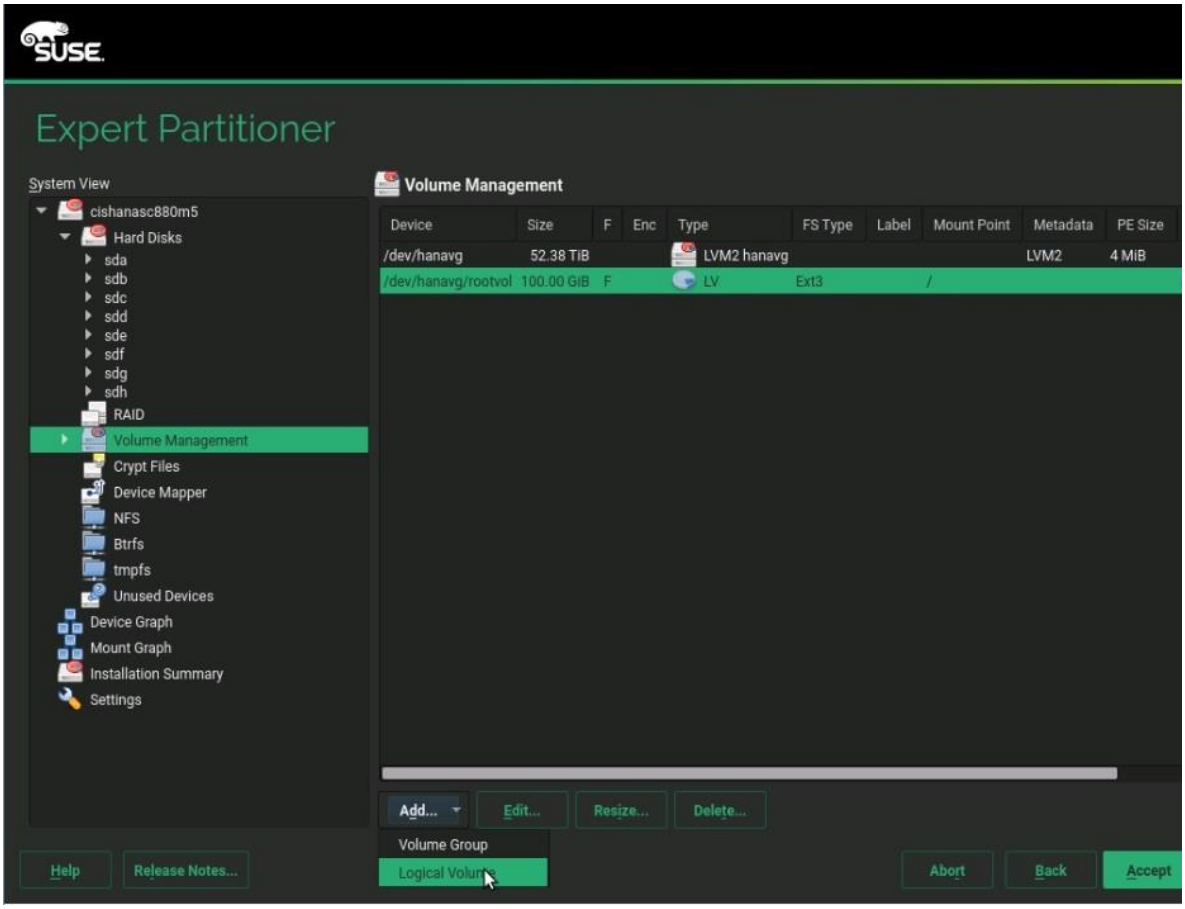

35. Name the logical volume for swap files **swapvol** (Figure 44). Then click Next.

Figure 44. Add Logical Volume: Specifying the name and type

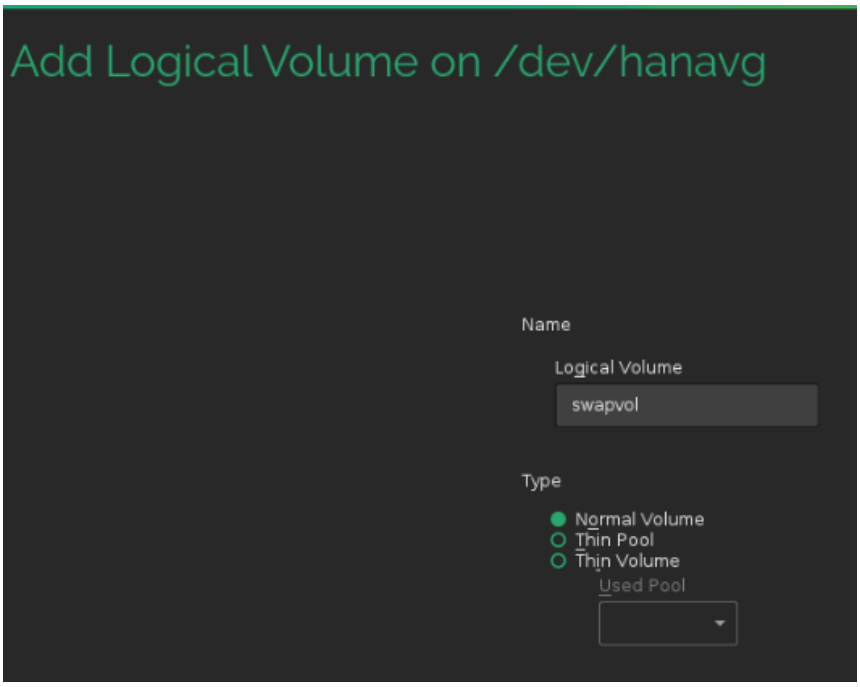

36. Assign 2 GB of space (Figure 45). Then click Next.

Figure 45. Add Logical Volume: Specifying the size and stripe information

37. Select Swap as the role (Figure 46). Then click Next.

Figure 46. Add Logical Volume: Specifying the role

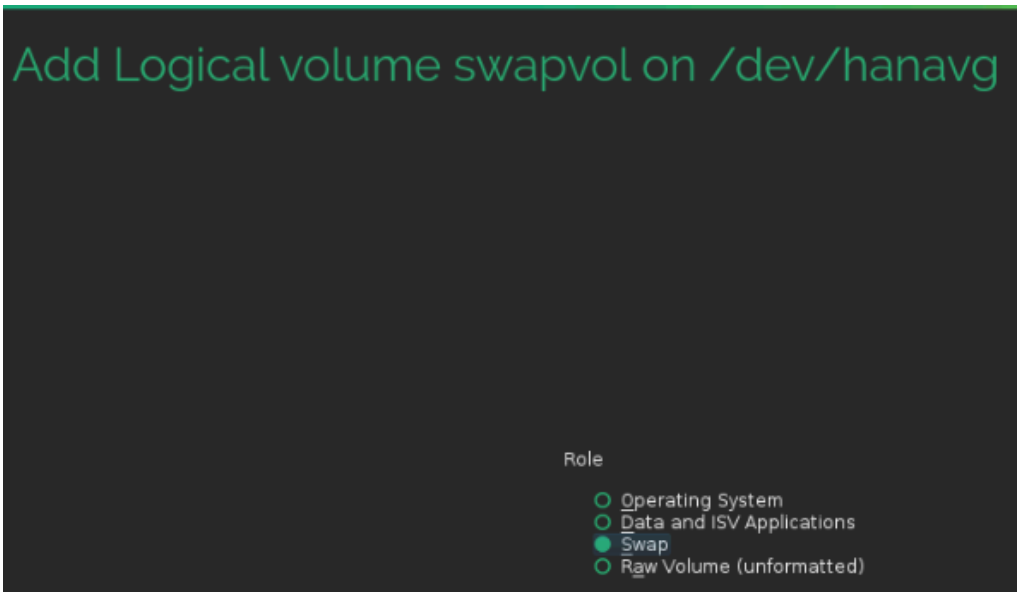

38. Specify the formatting and mounting options as shown in Figure 47.

Figure 47. Add Logical Volume: Specifying formatting and mounting options

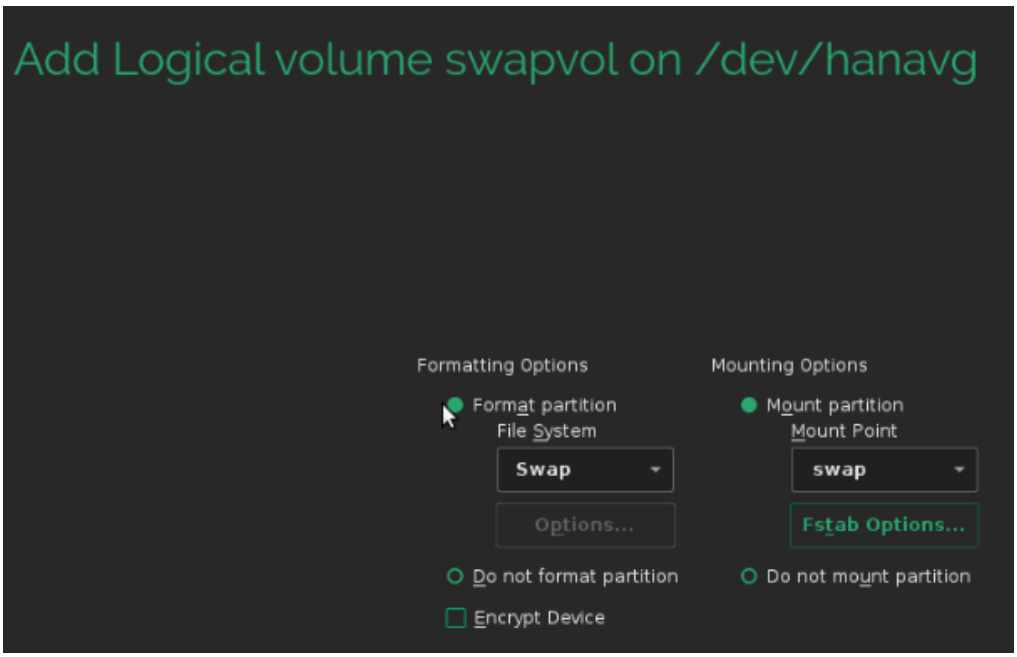

39. Click Finish. A summary of your settings appears (Figure 48).

Figure 48. Expert Partitioner: Volume Management summary

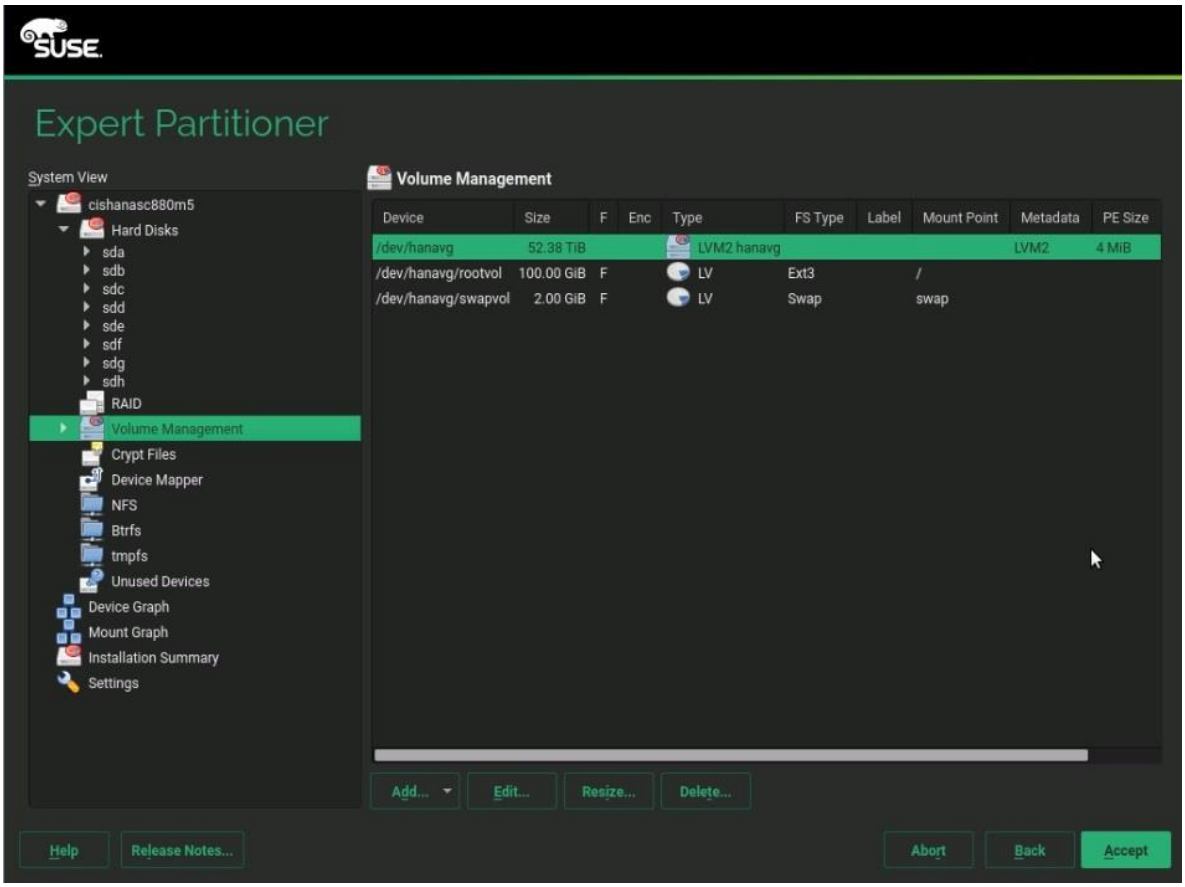

40. Click Accept to create the partitions (Figure 49).

Figure 49. Suggested Partitioning: Configuration created

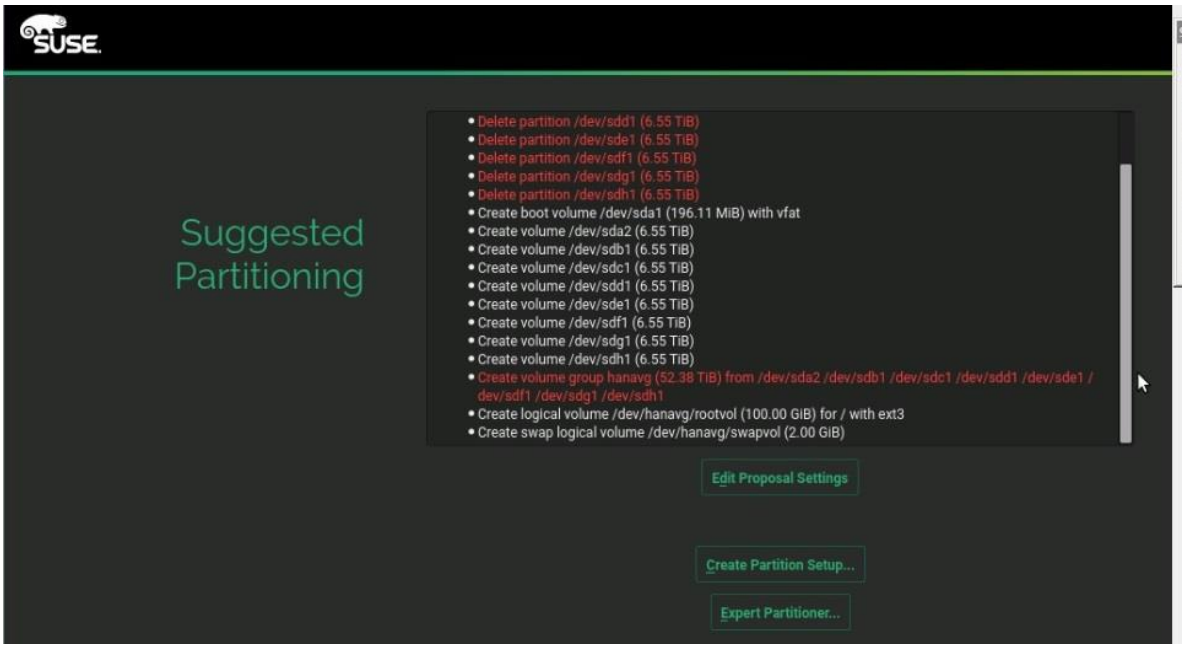

41. Review the updated partition information and click Next.
42. Now set the clock and time zone. Choose Clock and Time Zone, select the appropriate time zone, and set the hardware clock to UTC.
43. Set the password for the system administrator root by entering an appropriate password **<<var_sys_root-pw>>**.

44. Review the installation settings (Figure 50).

Figure 50. Installation Settings: Overview

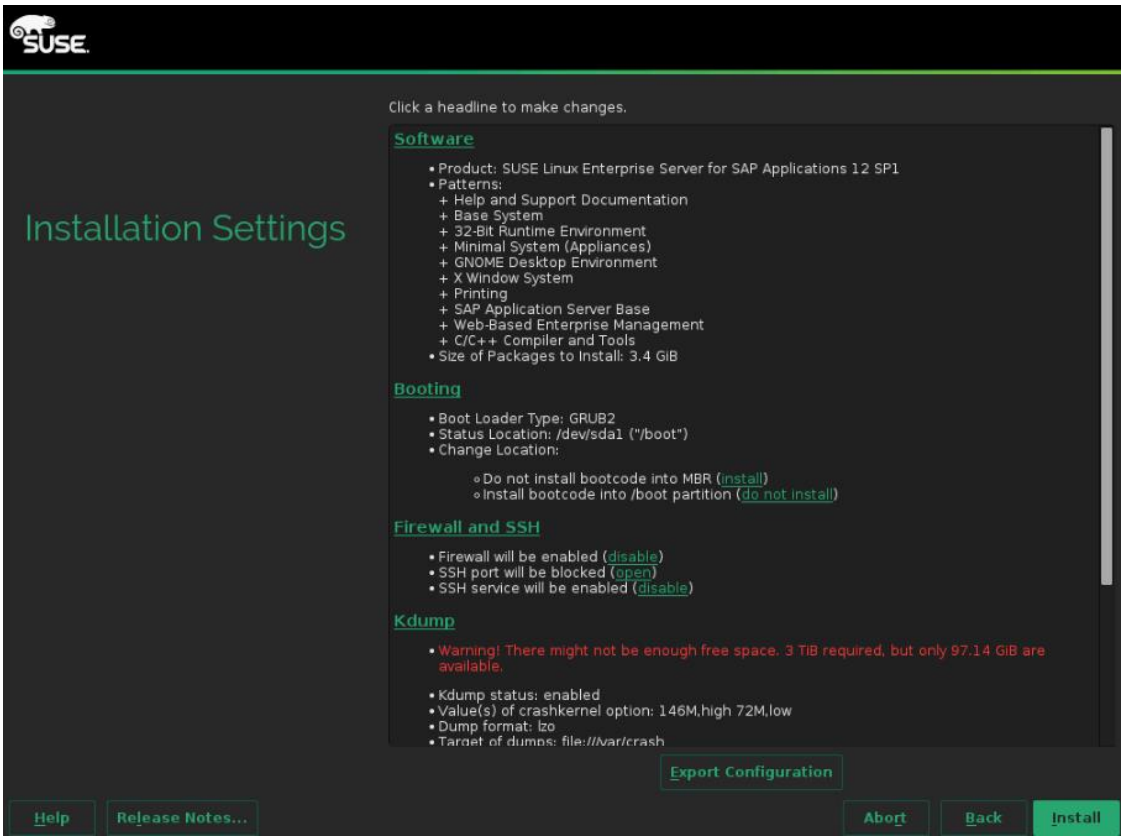

Customizing the installation

You can now customize the installation.

1. On the Installation Settings Screen, click Software. Then make the following changes (Figure 51):
   a. Deselect Gnome Desktop Environment.
   b. Select SAP HANA Server Base.
   c. Select C/C++ Compiler and Tools.

Figure 51. Customizing software selection and system tasks

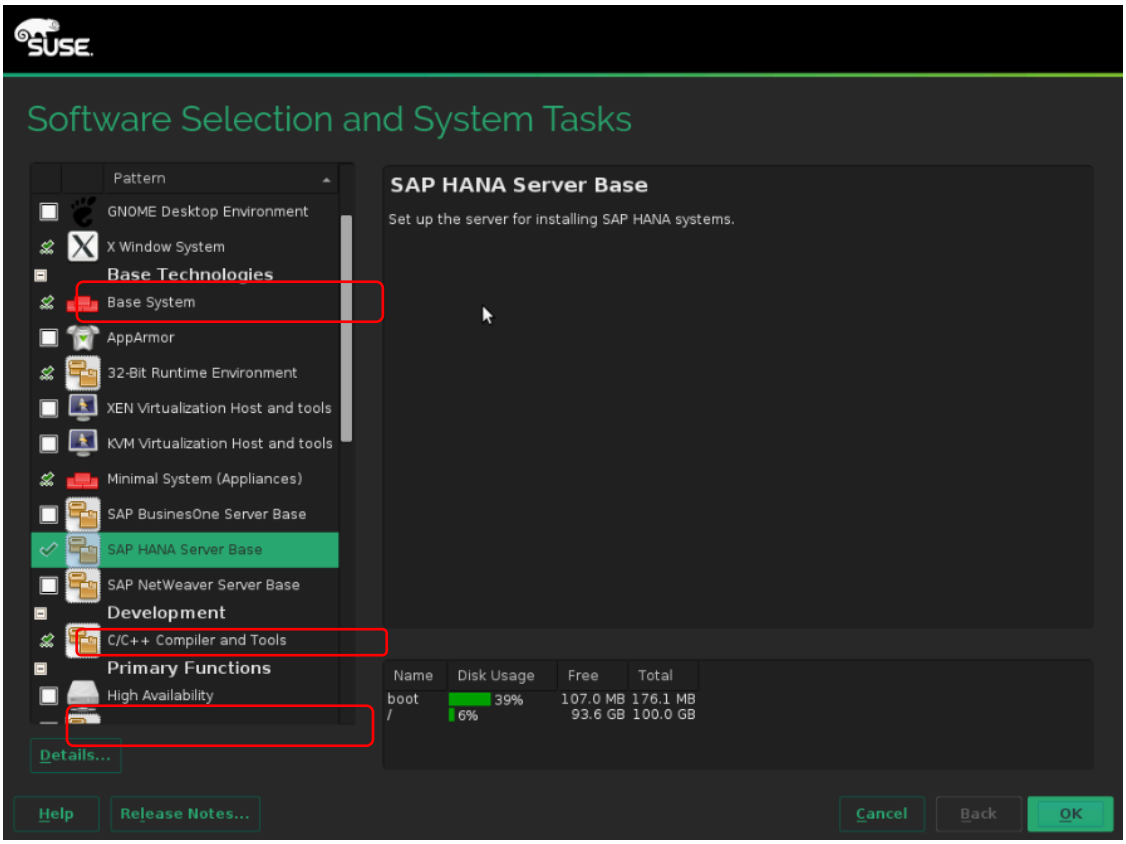

2. Click OK.
3. On the Installation Settings Screen, click Firewall and SSH. For the firewall, click disable. This step will automatically enable Secure Shell (SSH) service (Figure 52).

Figure 52. Firewall and SSH customized

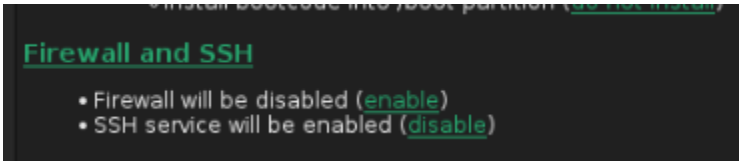

4. On the Installation Settings Screen, click Kdump. Select Disable Kdump (Figure 53).

Figure 53. Disabling Kdump

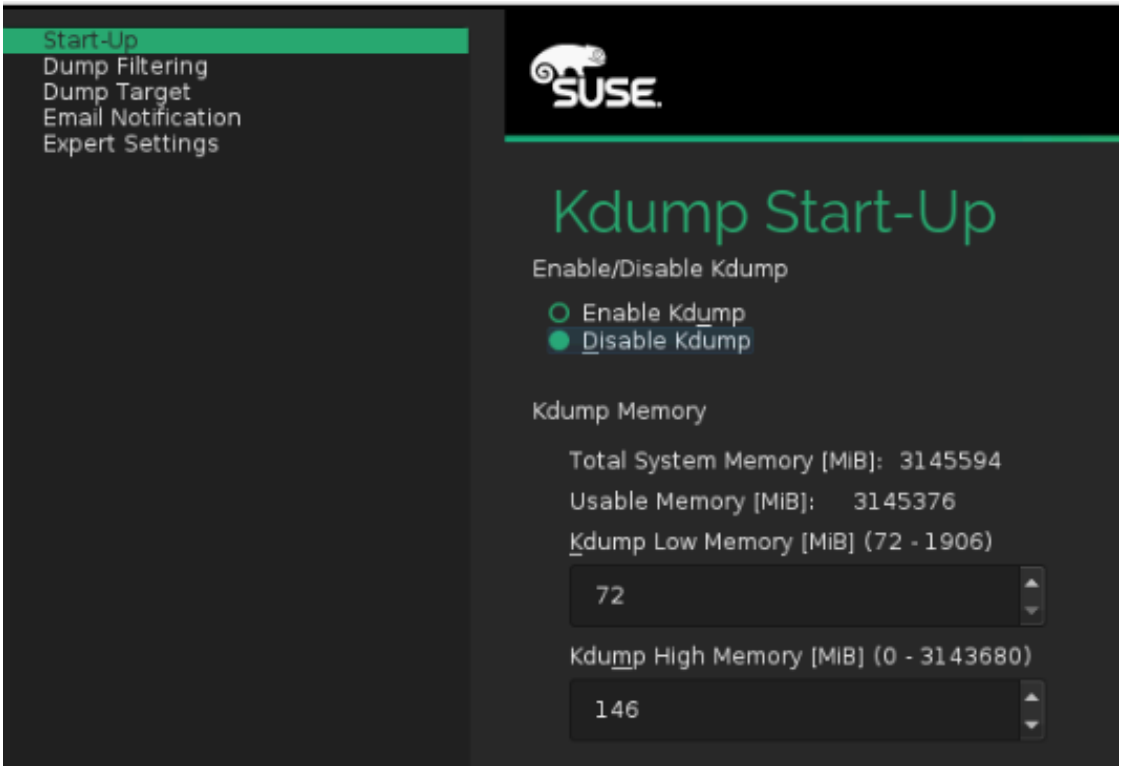

5. Click OK.

6. On the Installation Settings Screen, click “Default systemd target.” Select “Text mode” (Figure 54).

Figure 54. Setting the default systemd target to text mode

7. Click OK.
8. On the Installation Settings Screen, click Clone System Configuration (Figure 55).

Figure 55. Selecting Clone System Configuration

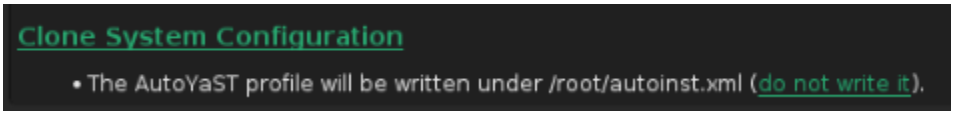

9. Click “do not write it” (Figure 56).

Figure 56. Changing the Clone System Configuration setting

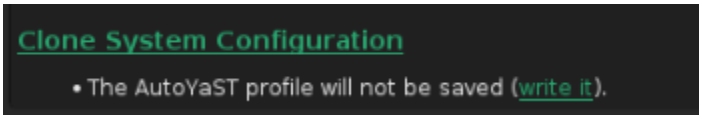

10. Leave the Booting and System default selections unchanged (Figure 57).

Figure 57. Installation Settings: Final selections

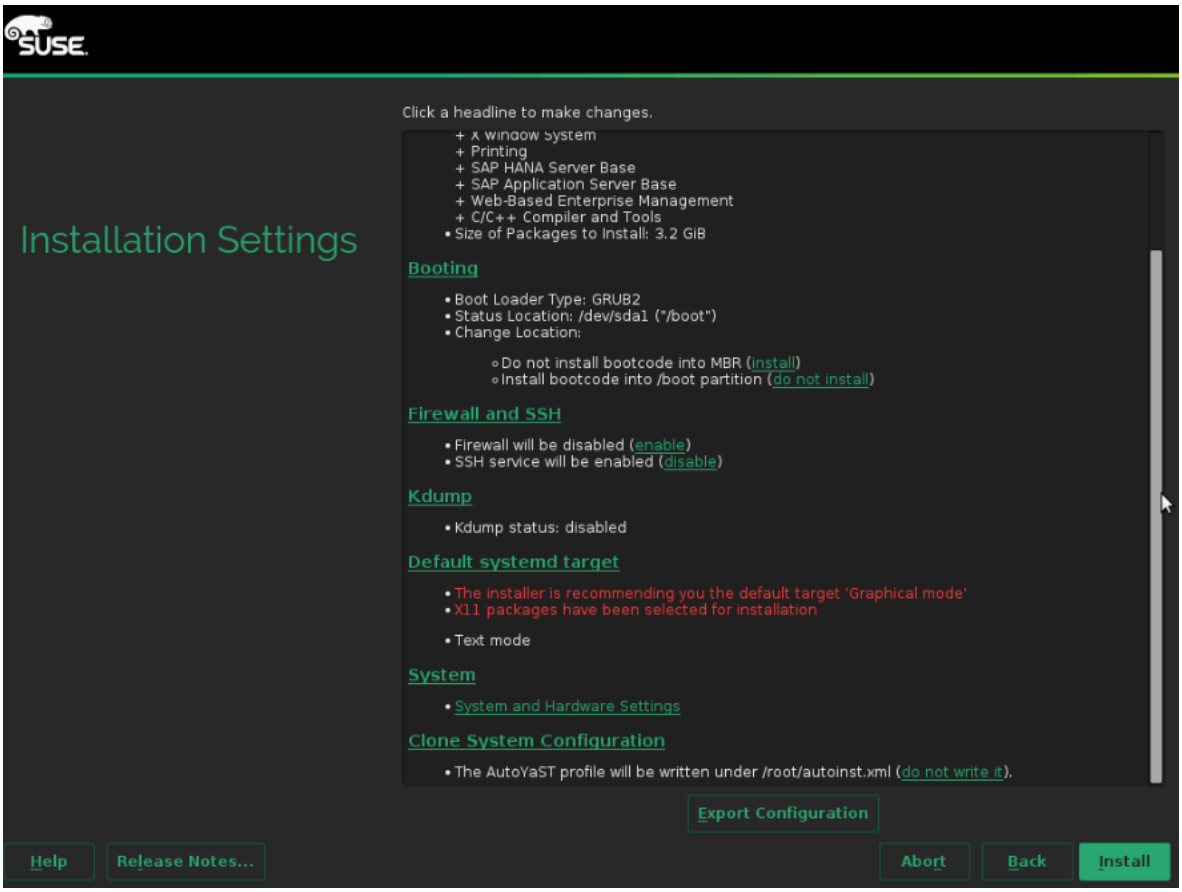

11. Click Install and select Install again at the subsequent Confirm Installation prompt. The installation starts, and you can monitor the status (Figure 58).

Figure 58. Performing the installation

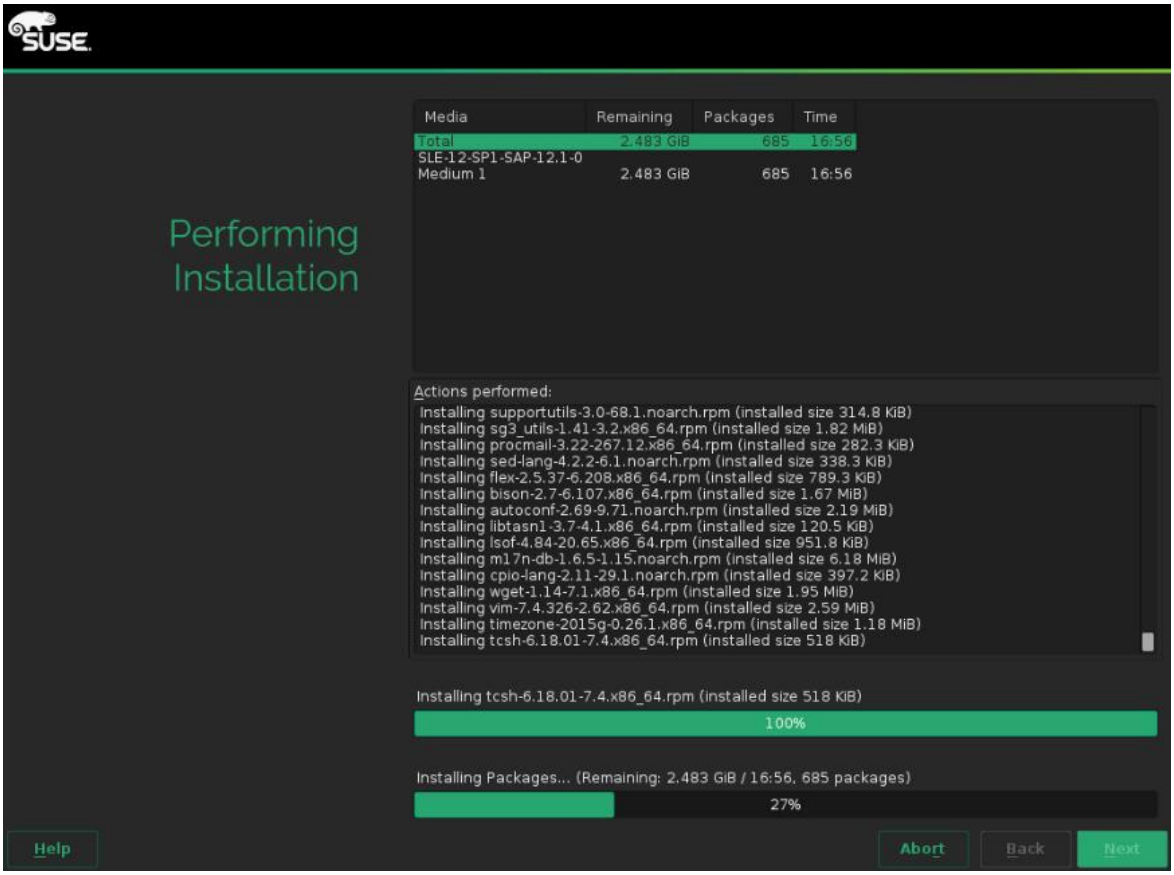

12. A reboot prompt will appear. Click OK (Figure 59).

Figure 59. Reboot alert when the in installation is complete

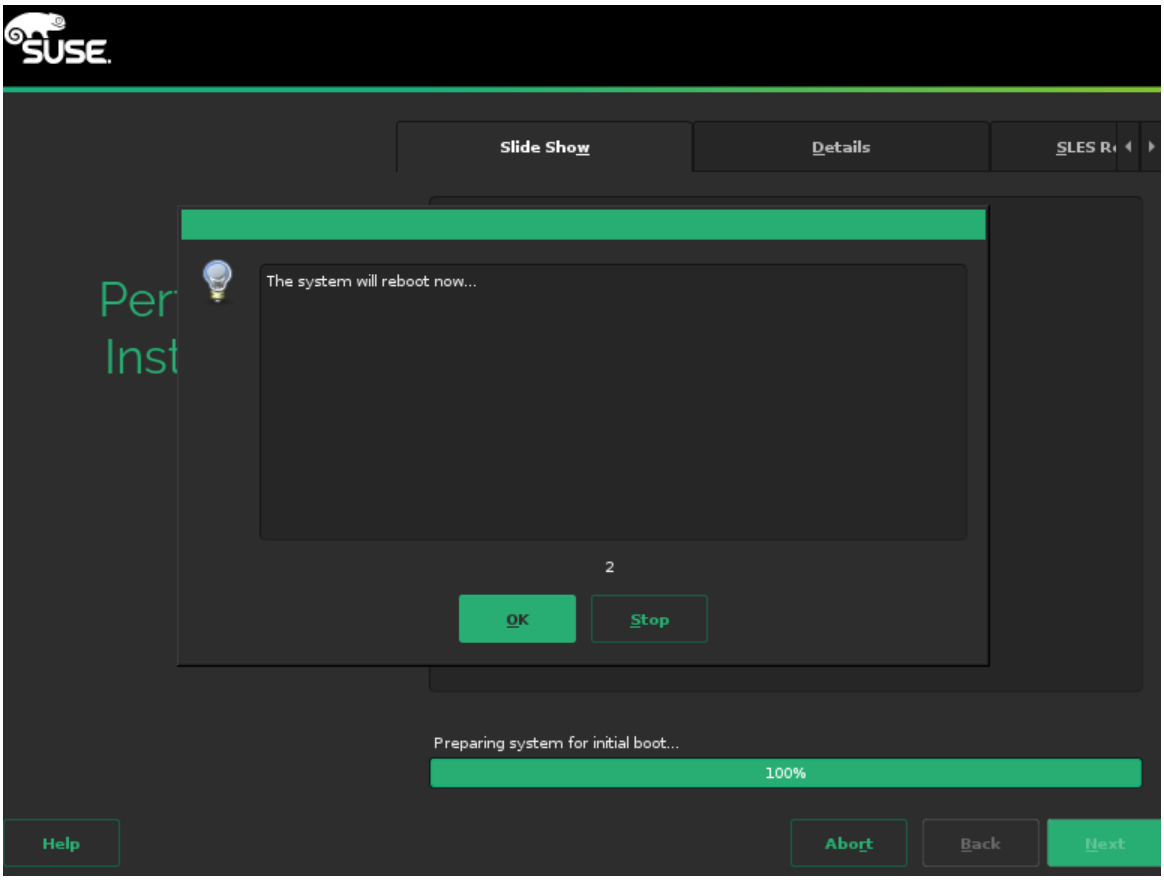

The system will reboot and boot from the disk when it restarts (Figure 60).

Figure 60. Booting from the hard disk

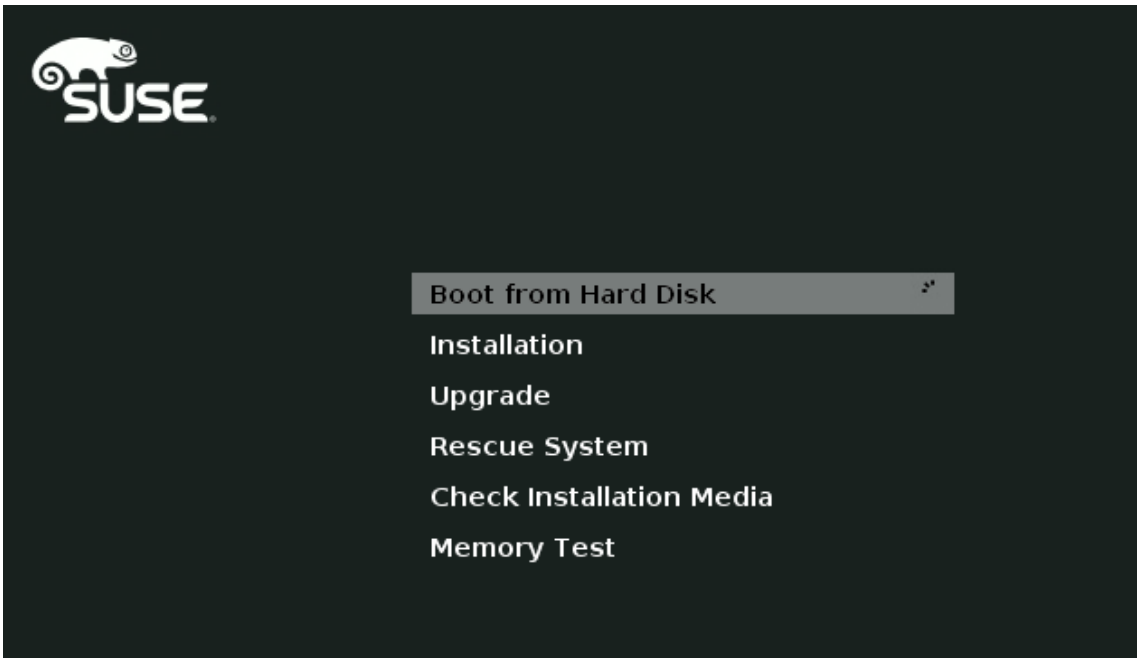

The login prompt will then appear (Figure 61).

Figure 61. Login prompt

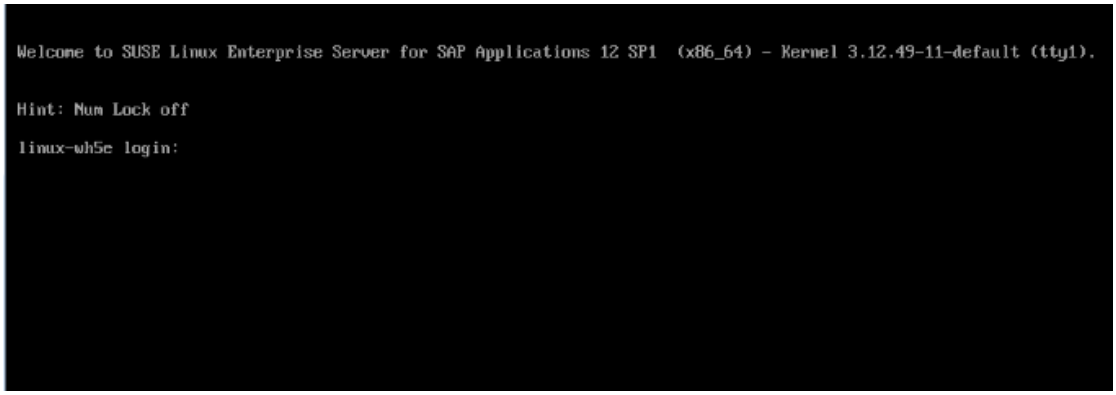

13. Use the KVM console to log in to the installed system as the user **root** with the password **<<var_sys_root-pw>>** (Figure 62).

Figure 62. Logging in as the root user

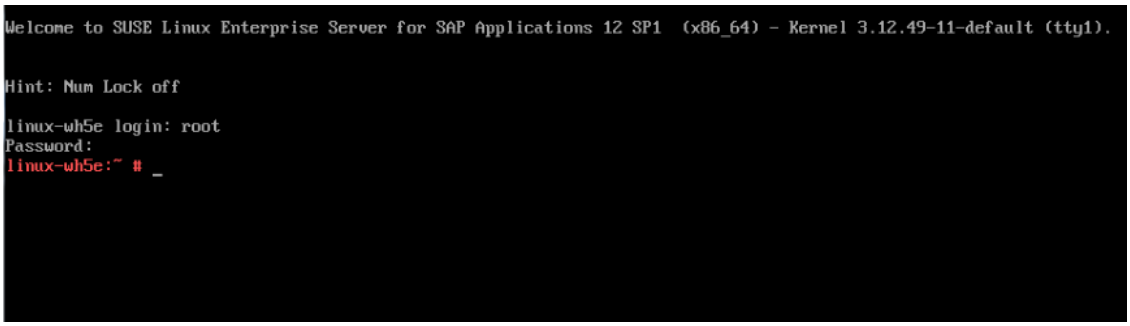

14. You will now configure the host name and disable IPv6. Enter the following command to display the YaST Control Center (Figure 63):

    #yast2

Figure 63. YaST Control Center

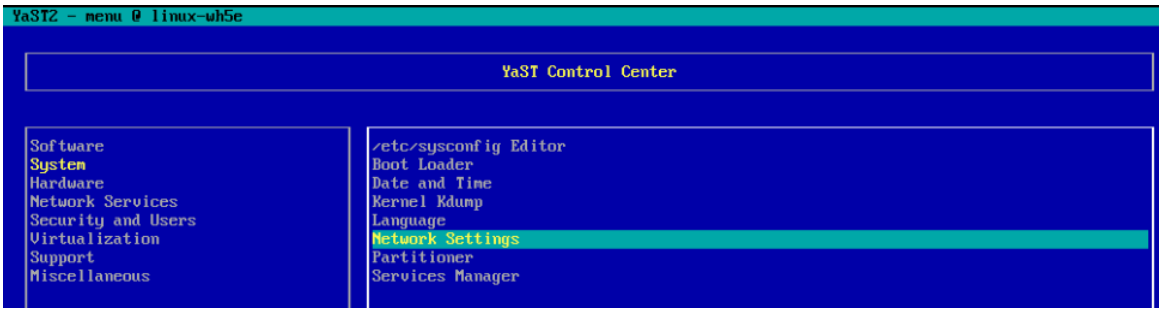

15. Select System > Network Settings and press Alt+S to select the Hostname/DNS tab (Figure 64).

Figure 64. YaST Control Center: Hostname/DNS settings

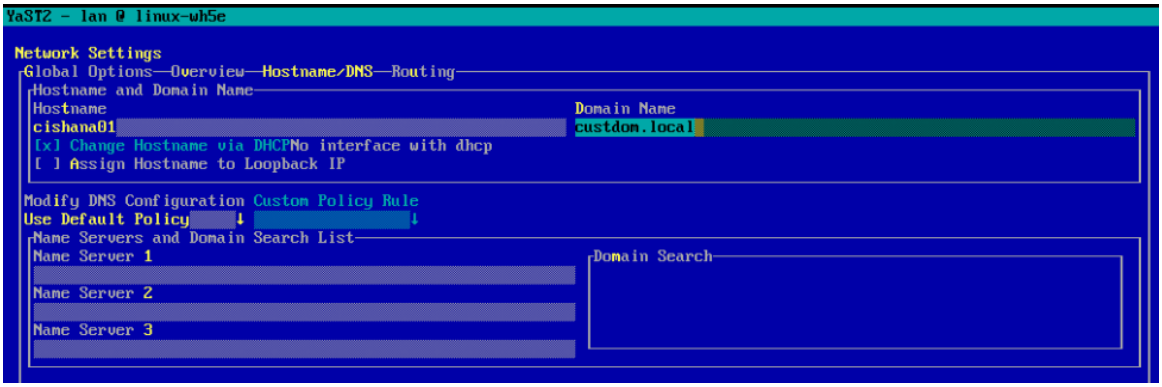

16. Enter the host domain name **<<var_hostname.domain>>**. Also enter the DNS server address of your network for resolution, if necessary. Then press Alt+O.
17. Press Alt+G to go to the Global Options tab. Disable IPv6 by deselecting the Enable IPv6 option as shown in Figure 65. Note that changing the IPv6 setting requires a reboot to make the changes take effect.

Figure 65. YaST: Configuring the IPv6 setting

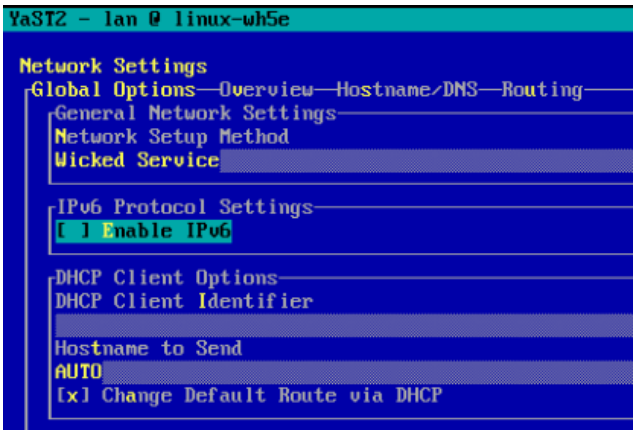

18. Press Alt+O to save the network configuration. The press Alt+Q to quit the YaST Control Center.
19. Reboot the system to make the IPv6 selection and the host-name settings take effect:
    #reboot
20. Identify the Ethernet interface port that is connected to the top-of-the-rack (ToR) switch. For now, you can use that port for management connectivity to the host. You can also check the port by using the **ifconfig** command, as shown in the sample in Figure 66.

Figure 66. Network interface configuration

```
Welcome to SUSE Linux Enterprise Server for SAP Applications 12 SP1  (x86_64) - Kernel 3.12.49-11-default (tty1).


cishana01 login: root
Password:
Last login: Thu Aug 11 03:41:42 on tty1
cishana01:~ # for i in `seq -w 0 7`; do ifconfig eth$i up;done
cishana01:~ # for i in `seq -w 0 7`; do cat /sys/class/net/eth$i/operstate;done
down
down
down
down
down
up
down
down
cishana01:~ # _
```

21. Assign **<<var_mgmt_ip_address>>** as the IP address and enter **<<var_mgmt_nw_netmask>>** as the subnet mask for the available interface (for example, **eth5** as in the example in Figure 66. You can use this configuration temporarily until you we port it to a high-availability bond device and create another with the Cisco VIC's 10-Gbps ports.

22. Go to the network configuration directory and create a configuration for eth5:

```
#cd /etc/sysconfig/network
#vi ifcfg-eth5
BOOTROTO='static'
IPADDR='<<var_mgmt_ip_address>>'
NETMASK='<<var_mgmt_nw_netmask>>'
NETWORK=''
MTU=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

23. Add the default gateway:

```
#cd /etc/sysconfig/network
# vi routes
default <<var_mgmt_gateway_ip>> - -
```

Note: Be sure that the system has access to the Internet or to a SUSE update server to install the patches.

24. Verify /etc/hosts as shown in the sample in Figure 67.

Figure 67. Verifying /etc/hosts

```
cishana01:~ # more /etc/hosts
#
# hosts         This file describes a number of hostname-to-address
#               mappings for the TCP/IP subsystem.  It is mostly
#               used at boot time, when no name servers are running.
#               On small systems, this file can be used instead of a
#               "named" name server.
# Syntax:
#
# IP-Address  Full-Qualified-Hostname  Short-Hostname
#

127.0.0.1       localhost

# special IPv6 addresses
::1             localhost ipv6-localhost ipv6-loopback

fe00::0         ipv6-localnet

ff00::0         ipv6-mcastprefix
ff02::1         ipv6-allnodes
ff02::2         ipv6-allrouters
ff02::3         ipv6-allhosts
173.36.215.118  cishana01.custdom.local cishana01
cishana01:~ #
```

25. Set up a proxy service so that the appliance can reach the Internet (Figure 68):

    #yast2

Figure 68. YaST: Proxy configuration

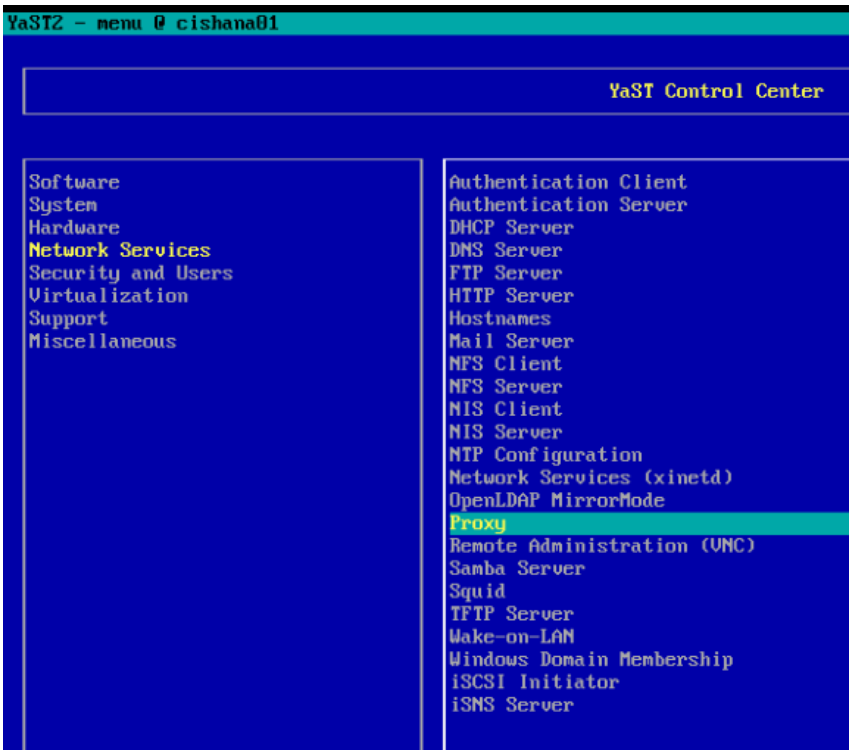

26. Enter the proxy server and port as shown in the sample configuration in Figure 69. Select OK and then quit YaST to save the configuration.

Figure 69. YaST: Proxy Configuration page

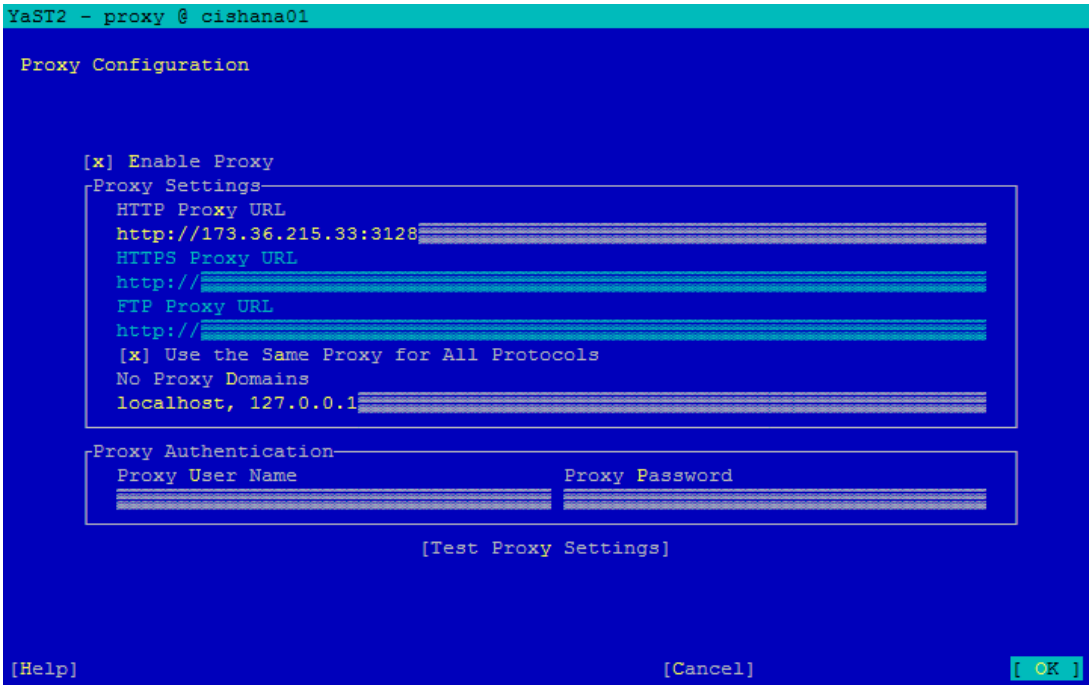

27. Register the system with SUSE to get the latest patches. For more information, refer to the SUSE knowledge base article https://www.suse.com/de-de/support/kb/doc?id=7016626.

    The system must have access to the Internet to proceed with this step.

    ```
    #SUSEConnect -r <<registration_code>> -e <<email_address>>
    ```

28. Update the system with the following command. Again, the system must have access to the Internet to proceed with this step.

    ```
    #zypper update
    ```

29. Follow the on-screen instructions to complete the update process. Reboot the server and log in to the system again.

Performing post-installation OS customization

To optimize the use of HANA database (HDB) with SLES for SAP 12 SP2,

1. Apply the following SAP Notes settings as instructed.
   2205917-SAP HANA DB-Recommended OS settings for SLES12-SLES for SAP Applications 12

2. Append the following entries to the /etc/init.d/boot.local file:

```
cpupower frequency-set -g performance
##
echo never > /sys/kernel/mm/transparent_hugepage/enabled
##
cpupower set -b 0
##
echo 0 > /sys/kernel/mm/ksm/run
##
echo deadline > /sys/block/sda/queue/scheduler
echo deadline > /sys/block/sdc/queue/scheduler
echo deadline > /sys/block/sdd/queue/scheduler
echo deadline > /sys/block/sde/queue/scheduler
echo deadline > /sys/block/sdf/queue/scheduler
echo deadline > /sys/block/sdg/queue/scheduler
echo deadline > /sys/block/sdh/queue/scheduler
echo deadline > /sys/block/sdi/queue/scheduler
##
echo 4096 > /sys/block/sda/queue/nr_requests
echo 4096 > /sys/block/sdc/queue/nr_requests
echo 4096 > /sys/block/sdd/queue/nr_requests
echo 4096 > /sys/block/sde/queue/nr_requests
echo 4096 > /sys/block/sdf/queue/nr_requests
echo 4096 > /sys/block/sdg/queue/nr_requests
echo 4096 > /sys/block/sdh/queue/nr_requests
echo 4096 > /sys/block/sdi/queue/nr_requests
##
echo 4096 > /sys/block/sda/queue/read_ahead_kb
echo 4096 > /sys/block/sdc/queue/read_ahead_kb
echo 4096 > /sys/block/sdd/queue/read_ahead_kb
echo 4096 > /sys/block/sde/queue/read_ahead_kb
echo 4096 > /sys/block/sdf/queue/read_ahead_kb
echo 4096 > /sys/block/sdg/queue/read_ahead_kb
echo 4096 > /sys/block/sdh/queue/read_ahead_kb
echo 4096 > /sys/block/sdi/queue/read_ahead_kb
```

3. Reboot the server:

   #reboot

Preparing SAP HANA data, log, and shared file systems

To prepare the file systems, start by carving out logical volumes for the data, log, and hana shared files. Then create the file systems. Then update /etc/fstab and mount the volumes.

The following commands and sizes are provided as examples.

1. Use the following command to check for the available physical volume (PV):

   #pvdisplay

   Figure 70 shows the results.

Figure 70. Checking for the physical volume

```
[root@cishana01 ~]# pvdisplay
  --- Physical volume ---
  PV Name               /dev/sda2
  VG Name               hanavg
  PV Size               16.00 TiB / not usable 0
  Allocatable           yes
  PE Size               4.00 MiB
  Total PE              4193418
  Free PE               4167306
  Allocated PE          26112
  PV UUID               zvb8AK-Dyf8-fkyU-Xznt-1XHx-RgmZ-9q6qI6
```

2. Use the following command to check for the available volume group (VG) hanavg:

   #vgdisplay

   Figure 71 shows the results.

Figure 71. Checking for the volume group

```
[root@cishana01 ~]# vgdisplay
  --- Volume group ---
  VG Name               hanavg
  System ID
  Format                lvm2
  Metadata Areas        1
  Metadata Sequence No  3
  VG Access             read/write
  VG Status             resizable
  MAX LV                0
  Cur LV                2
  Open LV               2
  Max PV                0
  Cur PV                1
  Act PV                1
  VG Size               16.00 TiB
  PE Size               4.00 MiB
  Total PE              4193418
  Alloc PE / Size       26112 / 102.00 GiB
  Free  PE / Size       4167306 / 15.90 TiB
  VG UUID               4sIS1V-tpWq-cSse-qjaY-0V1V-0dFu-38jYrV
```

3. Create logical volumes (LVs) for the data, log, and hana shared file systems:

```
lvcreate –name <<lvname>>  -I<<stripesize>> -L<<volume-size>>  <<parent-vg-name>>
```

```
# lvcreate -L 9T -i 8 -I 256k -n datavol hanavg
```

```
# lvcreate -L 512G -i 8 -I 256k -n logvol hanavg
```

```
# lvcreate -L 3T -i 8 -I 256k -n sapmnt hanavg
```

4. Verify the created logical volumes:

```
# lvdisplay
```

5. Create file systems in the data, log, and hana shared volumes:

```
# mkfs.xfs -f -b size=4096 /dev/hanavg/datalv
```

```
# mkfs.xfs -f -b size=4096 /dev/hanavg/loglv
```

```
# mkfs.xfs -f -b size=4096 /dev/hanavg/sapmnt
```

6. Create mount directories for the data, log, and hana shared file systems:

```
#mkdir -p /hana/data
```

```
#mkdir -p /hana/log
```

```
#mkdir -p /hana/shared
```

7. Mount options vary from the default Linux settings for XFS for SAP HANA data and log volumes. The following is a sample /etc/fstab entry. Make sure that you use the same mount options for the data and log file systems as shown in the example.

```
/dev/mapper/hanavg-rootvol              /                       ext3    defaults      1 1
UUID=fc2e52c4-e6f6-4e9a-9ad1-86aeb3369942     /boot/efi                 ext3    defaults      1 2
/dev/mapper/hanavg-swapvol      swap            swap    defaults      0 0
/dev/hanavg/datavol /hana/data     xfs
nobarrier,noatime,nodiratime,logbufs=8,logbsize=256k,async,swalloc,allocsize=262144k    1 2
/dev/hanavg/logvol /hana/log    xfs
nobarrier,noatime,nodiratime,logbufs=8,logbsize=256k,async,swalloc,allocsize=131072k    1 2
/dev/hanavg/sapmnt /hana/shared       xfs   defaults    1 2
```

8. Use the following command to mount the file systems:

```
#mount -a
```

9. Use the **df –h** command to check the status of all mounted volumes (Figure 72).

Figure 72. Checking the status of mounted volumes

```
[root@cishana01 ~]# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/mapper/hanavg-rootlv
                       99G  882M   93G   1% /
/dev/sda1             485M   47M  414M  11% /boot
tmpfs                 1.5T     0  1.5T   0% /dev/shm
/dev/mapper/hanavg-datalv
                      9.0T   33M  9.0T   1% /hana/data
/dev/mapper/hanavg-loglv
                      512G   33M  512G   1% /hana/log
/dev/mapper/hanavg-sharedlv
                      3.0T   33M  3.0T   1% /hana/shared
[root@cishana01 ~]#
```

10. Change the directory permissions **before** you installing SAP HANA. Use the **chown** command on each SAP HANA node after the file systems are mounted:

```
#chmod -R 777 /hana/data
#chmod -R 777 /hana/log
#chmod -R 777 /hana/shared
```

Installing SAP HANA

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

For the SAP HANA installation documentation, see the SAP HANA Server Installation Guide.

All other SAP installation and administration documentation is available at http://service.sap.com/instguides.

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: https://service.sap.com/notes.

SAP HANA in-memory database (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 1890444: HANA system slow due to CPU power save mode

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 Business Warehouse (BW) on HANA: Sizing the SAP HANA database

SAP Note 1793345: Sizing for Suite on HANA

Linux notes

SAP Note 2205917: SAP HANA DB: Recommended OS settings for SLES 12 and SLES for SAP Applications 12

SAP Note 2235581: SAP HANA: Supported operating systems

SAP Note 1944799: SAP HANA guidelines for the SLES operating system

SAP Note 1731000: Nonrecommended configuration changes

SAP Note 1557506: Linux paging improvements

SAP Note 1726839: SAP HANA database: Potential crash when using XFS file system

SAP Note 1740136: SAP HANA: Wrong mount option may lead to corrupt persistency

SAP Note 1829651: Time-zone settings in SAP HANA scale-out landscapes

SAP application notes

SAP Note 1658845: SAP HANA database hardware check

SAP Note 1637145: SAP BW on SAP HANA: Sizing SAP in-memory database

SAP Note 1661202: Support for multiple applications on SAP HANA

SAP Note 1681092: Support for multiple SAP HANA databases on one SAP HANA appliance (also known as multi-SID)

SAP Note 1577128: Supported clients for SAP HANA 1.0

SAP Note 1808450: Homogenous system landscape for on SAP BW on SAP HANA

SAP Note 1976729: Application component hierarchy for SAP HANA

SAP Note 1927949: Standard behavior for SAP logon tickets

SAP Note 1577128: Supported clients for SAP HANA

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 1730932: Using backup tools with Backint for SAP HANA

SAP HANA virtualization notes

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

SAP HANA post-installation checkup

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

Checking the SAP HANA services

```
skladm@cishana01:/usr/sap/SKL/HDB00> /usr/sap/hostctrl/exe//sapcontrol -nr 00 -function
GetProcessList


19.05.2016 11:29:27
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2016 04 13 08:51:49, 866:37:38, 41691
hdbcompileserver, HDB Compileserver, GREEN, Running, 2016 04 13 08:51:56, 866:37:31, 41837
hdbindexserver, HDB Indexserver, GREEN, Running, 2016 04 13 08:52:00, 866:37:27, 41863
hdbnameserver, HDB Nameserver, GREEN, Running, 2016 04 13 08:51:50, 866:37:37, 41711
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2016 04 13 08:51:56, 866:37:31, 41839
hdbwebdispatcher, HDB Web Dispatcher, GREEN, Running, 2016 04 13 08:53:11, 866:36:16, 42431
hdbxsengine, HDB XSEngine, GREEN, Running, 2016 04 13 08:52:00, 866:37:27, 41865
skladm@cishana01-skl:/usr/sap/SKL/HDB00>
```

Checking the HANA database information

```
skladm@cishana01:/usr/sap/SKL/HDB00> HDB info
USER         PID    PPID %CPU    VSZ   RSS COMMAND
skladm     59578   59577  0.0 108472  1944 -sh
skladm     59663   59578  0.0 114080  2020  \_ /bin/sh /usr/sap/SKL/HDB00/HDB info
skladm     59692   59663  0.0 118048  1596      \_ ps fx -U skladm -o user,pid,ppid,pcpu,vsz,rss,args
skladm     41683       1  0.0  22188  1640 sapstart pf=/hana/shared/SKL/profile/SKL_HDB00_cishana01-skl
skladm     41691   41683  0.0 582888 290988  \_ /usr/sap/SKL/HDB00/cishana01-
skl/trace/hdb.sapSKL_HDB00 -d -nw -f /usr/sap/SKL/HDB00/cishana01-skl/daemon.ini
skladm     41711   41691  0.3 54292416 2058900      \_ hdbnameserver
skladm     41837   41691  0.1 4278472 1243356      \_ hdbcompileserver
skladm     41839   41691  0.2 11773976 8262724      \_ hdbpreprocessor
skladm     41863   41691  6.2 22143172 18184604      \_ hdbindexserver
skladm     41865   41691  0.5 8802064 2446612      \_ hdbxsengine
skladm     42431   41691  0.1 4352988 823220      \_ hdbwebdispatcher
skladm     41607       1  0.0 497576 23232 /usr/sap/SKL/HDB00/exe/sapstartsrv
pf=/hana/shared/SKL/profile/SKL_HDB00_cishana01-skl -D -u skladm
skladm@cishana01-skl:/usr/sap/SKL/HDB00>
```

Tuning the SAP HANA performance parameters

After SAP HANA is installed, tune the parameters as explained in the following SAP Notes and shown in Table 5.

Table 5. SAP HANA performance tuning parameters

Parameters and file system	Data file system	Log file system
max_parallel_io_requests	256	Default
async_read_submit	On	On
async_write_submit_blocks	All	All
async_write_submit_active	Auto	On

SAP Note 2399079: Elimination of hdbparam in HANA 2

SAP Note 2186744: FAQ: SAP HANA parameters

SAP HANA operation and maintenance

SAP HANA operation and maintenance processes are described in detail in many related SAP documents. For a complete list of the documentation available, see http://help.sap.com/hana.

This document summarizes only a few important operation and maintenance procedures. Most of the procedures described in this document are command-line interface (CLI) procedures and are independent of any GUI requiring an X terminal or other GUI front end (Microsoft Windows PC, Linux desktop, etc.). CLI procedures can be started using the KVM or any SSH tool such as PuTTY (for Windows) or Terminal (for Mac OS), or any Linux terminal window to connect to the SAP HANA database system (the appliance).

Monitoring SAP HANA

Three easy CLI methods are available to check the running SAP HANA database.

saphostagent

1. Start a shell and connect to the SAP HANA system as the root user.

```
cishana01:~ # /usr/sap/hostctrl/exe/saphostctrl -function ListDatabases
Instance name: HDB00, Hostname: cishana01, Vendor: HDB, Type: hdb, Release: 1.00.60.0379371
  Database name: HAN, Status: Error
cishana01:~ #
```

2. Get a list of installed HANA instances or databases.

```
cishana01:~ # /usr/sap/hostctrl/exe/saphostctrl -function ListInstances
 Inst Info : HAN - 00 - cishana01 - 740, patch 17, changelist 1413428
cishana01:~ #
```

3. Using this information (system ID [SID] and system number), you can use **sapcontrol** to gather more information about the running HANA database.

sapcontrol

1. In a shell, use the **sapcontrol** function **GetProcessList** to display a list of running HANA OS processes.

```
cishana01:~ # /usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetProcessList

19.07.2016 14:54:45
GetProcessList
OK
name, description, dispstatus, textstatus, starttime, elapsedtime, pid
hdbdaemon, HDB Daemon, GREEN, Running, 2016 07 15 11:57:45, 98:57:00, 8545
hdbnameserver, HDB Nameserver, GREEN, Running, 2016 07 15 12:05:27, 98:49:18, 11579
hdbpreprocessor, HDB Preprocessor, GREEN, Running, 2013 08 15 12:05:27, 98:49:18, 11580
hdbindexserver, HDB Indexserver, GREEN, Running, 2016 07 15 12:05:27, 98:49:18, 11581
hdbstatisticsserver, HDB Statisticsserver, GREEN, Running, 2016 07 15 12:05:27, 98:49:18, 11582
hdbxsengine, HDB XSEngine, GREEN, Running, 2016 07 15 12:05:27, 98:49:18, 11583
sapwebdisp_hdb, SAP WebDispatcher, GREEN, Running, 2016 07 15 12:05:27, 98:49:18, 11584
hdbcompileserver, HDB Compileserver, GREEN, Running, 2016 07 15 12:05:27, 98:49:18, 11585
```

You see processes such as **hdbdaemon**, **hdbnameserver**, and **hdbindexserver** that belong to a running HANA database.

2. You can also get a system instance list, which is more useful for a scale-out appliance.

```
cishana01:~ # /usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetSystemInstanceList

19.07.3016 15:03:12
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
cishana01, 0, 50013, 0, 0.3, HDB, GREEN
```

HDB info

Another important tool is the **HDB** command, which needs to be issued by the <SID>adm user: the OS user who owns the HANA database.

As the root user on the SAP HANA appliance, enter the following command:

```
cishana01:~ # su - hanadm
```

```
cishana01:/usr/sap/HAN/HDB00> HDB info
USER       PID  PPID %CPU    VSZ   RSS COMMAND
hanadm   61208 61207  1.6  13840  2696 -sh
hanadm   61293 61208  0.0  11484  1632  \_ /bin/sh /usr/sap/HAN/HDB00/HDB info
hanadm   61316 61293  0.0   4904   872      \_ ps fx -U hanadm -o user,pid,ppid,pcpu,vsz,rss,args
hanadm    8532     1  0.0  20048  1468 sapstart pf=/hana/shared/HAN/profile/HAN_HDB00_cishana01
hanadm    8545  8532  1.5 811036 290140  \_ /usr/sap/HAN/HDB00/cishana01/trace/hdb.sapHAN_HDB00 -d
-nw -f /usr/sap/HAN/HDB00/cis
hanadm   11579  8545  6.6 16616748 1789920      \_ hdbnameserver
hanadm   11580  8545  1.5 5675392 371984      \_ hdbpreprocessor
hanadm   11581  8545 10.9 18908436 6632128      \_ hdbindexserver
hanadm   11582  8545  8.7 17928872 3833184      \_ hdbstatisticsserver
hanadm   11583  8545  7.4 17946280 1872380      \_ hdbxsengine
hanadm   11584  8545  0.0 203396 16000      \_ sapwebdisp_hdb
pf=/usr/sap/HAN/HDB00/cishana01/wdisp/sapwebdisp.pfl -f /usr/sap/H
hanadm   11585  8545  1.5 15941688 475708      \_ hdbcompileserver
hanadm    8368     1  0.0 216268 75072 /usr/sap/HAN/HDB00/exe/sapstartsrv
pf=/hana/shared/HAN/profile/HAN_HDB00_cishana01 -D -u
```

This command produces output similar to that from the **sapcontrol GetProcessList** function, with a bit more information about the process hierarchy.

Starting and stopping SAP HANA

Before you stop the SAP HANA appliance, you must be able to stop and start the HANA database. You can use the commands shown here.

sapcontrol

You can use the **sapcontrol** functions **StartSystem** and **StopSystem** to start and stop a HANA database.

Stop the system with the **StopSystem** function.

```
cishana01:~ # /usr/sap/hostctrl/exe/sapcontrol -nr 00 -function StopSystem HDB

19.07.3016 15:05:35
StopSystem
OK
```

Use the following command to verify that the database has stopped.

```
cishana01:~ # /usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetSystemInstanceList

19.07.3016 15:05:58
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
cishana01, 0, 50013, 0, 0.3, HDB, YELLOW
```

Wait for the status to be **GRAY**.

```
cishana01:~ # /usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetSystemInstanceList

19.07.3016 15:07:52
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
cishana01, 0, 50013, 0, 0.3, HDB, GRAY
```

You can also use the **HDB info** command.

```
cishana01:~ # su -l hanadm
cishana01:/usr/sap/HAN/HDB00> HDB info
USER       PID  PPID %CPU    VSZ   RSS COMMAND
hanadm   61477 61476  2.0  13840  2692 -sh
hanadm   61562 61477  0.0  11484  1632  \_ /bin/sh /usr/sap/HAN/HDB00/HDB info
hanadm   61585 61562  0.0   4904   872      \_ ps fx -U hanadm -o user,pid,ppid,pcpu,vsz,rss,args
hanadm    8368     1  0.0 216784 75220 /usr/sap/HAN/HDB00/exe/sapstartsrv
pf=/hana/shared/HAN/profile/HAN_HDB00_cishana01 -D -u
cishana01:/usr/sap/HAN/HDB00>
```

You can start the database again with the sapcontrol command **StartSystem** function.

```
cishana01:~ # /usr/sap/hostctrl/exe/sapcontrol -nr 00 -function StartSystem HDB
19.07.3016 15:08:48
StartSystem
OK
```

To check the system status, use the sapcontrol command **GetSystemInstanceList** function. Wait for the status to be **GREEN**.

```
cishana01:~ # /usr/sap/hostctrl/exe/sapcontrol -nr 00 -function GetSystemInstanceList

19.07.3016 15:10:19
GetSystemInstanceList
OK
hostname, instanceNr, httpPort, httpsPort, startPriority, features, dispstatus
cishana01, 0, 50013, 0, 0.3, HDB, GREEN
```

HDB

You can use the **HDB start** and **stop** commands to stop and start the HANA database.

Use **HDB stop** to stop the database.

```
cishana01:~ # su - hanadm
cishana01:/usr/sap/HAN/HDB00> HDB stop
hdbdaemon will wait maximal 300 seconds for NewDB services finishing.
Stopping instance using: /usr/sap/HAN/SYS/exe/hdb/sapcontrol -prot NI_HTTP -nr 00 -function
StopWait 400 2



19.07.2016 19:10:37
Stop
OK
```

In contrast to **sapcontrol**, this command waits until the database is stopped or started.

```
cishana01:/usr/sap/HAN/HDB00> HDB start



StartService
Impromptu CCC initialization by 'rscpCInit'.
  See SAP note 1266393.
OK
OK
Starting instance using: /usr/sap/HAN/SYS/exe/hdb/sapcontrol -prot NI_HTTP -nr 00 -function
StartWait 2700 2



19.07.2016 19:11:20
Start
OK
```

Downloading revisions

To download revisions, you need to connect to the service marketplace and select the software download area to search for available patches.

Refer to http://help.sap.com/hana/SAP_HANA_Master_Update_Guide_en.pdf for update procedures for SAP HANA.

For more information

For more information about SAP HANA, see https://hana.sap.com/abouthana.html.

For a list of certified and supported SAP HANA hardware, see https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html.

Appendix: Solution variables used for this document

Before starting the configuration process, you need to collect some specific configuration information. Table 6 provides information to help you assemble the required network and host address, numbering, and naming information. This worksheet can also be used as a "leave behind" document for future reference.

Table 6. Solution variables used for this document

Variable	Description	Value used in the lab for this document
<<var_cimc_ip_address>>	Cisco C880 M5 server's IMC IP address	173.36.215.117
<<var_cimc_ip_netmask>>	Cisco C880 M5 server's IMC network netmask	255.255.255.0
<<var_cimc_gateway_ip>>	Cisco C880 M5 server's IMC network gateway IP address	173.36.215.1
<<var_raid50_vd_name>>	Name for virtual drive VD0 during RAID configuration	ucs_hana
<<var_hostname.domain>>	SAP HANA node FQDN	cishana01.custdom.local
<<var_sys_root-pw>>	SAP HANA node's root password	Saphana1!
<<var_lvm_vg_name>>	SAP HANA node's OS LVM volume group name	hanavg
<<var_mgmt_ip_address>>	SAP HANA node's management and administration IP address	173.36.215.118
<<var_mgmt_nw_netmask>>	SAP HANA node's management network netmask	255.255.255.0
<<var_mgmt_gateway_ip>>	Cisco C880 M5 server's management and administration network gateway IP address	173.36.215.1
<<var_mgmt_netmask_prefix>>	Netmask prefix in CIDR notation	24

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-740470-00 04/18