

Cisco Scale-Up Solution for SAP HANA on Cisco C880 M5 Servers with Red Hat Enterprise Linux for SAP HANA



Design and deploy an SAP HANA single-node solution based on standalone Cisco C880 M5 Servers with Red Hat Enterprise Linux for SAP HANA 7.3

April 2018

Contents

Executive Summary	3
Solution Overview	3
Introduction	3
Audience	3
Purpose of this document	3
Solution summary: Cisco C880 M5 Server.....	4
Infrastructure overview	5
Cisco C880 M5 Server.....	5
Solution design	7
SAP HANA system.....	7
Hardware requirements for the SAP HANA database.....	7
Deployment hardware and software.....	9
Configuration guidelines	9
Preparing the SAP HANA scale-up node	10
Configuring the Cisco C880 M5 integrated remote management controller	10
Starting Cisco C880 M5 video redirection	14
Configuring Cisco C880 M5 BIOS CPU settings.....	15
Configuring Cisco C880 M5 BIOS disk (RAID) settings for the SAP HANA file system.....	17
Committing the changes to the BIOS	20
Installing the operating system	21
Performing post-installation OS customization.....	43
Installing SAP HANA.....	51
Important SAP Notes.....	51
SAP HANA post-installation checkup.....	52
Tuning the SAP HANA performance parameters.....	53
SAP HANA operation and maintenance	54
Monitoring SAP HANA.....	54
For more information	60
Appendix: Solution variables used for this document	61

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 prevalidated, 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 <ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> 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:

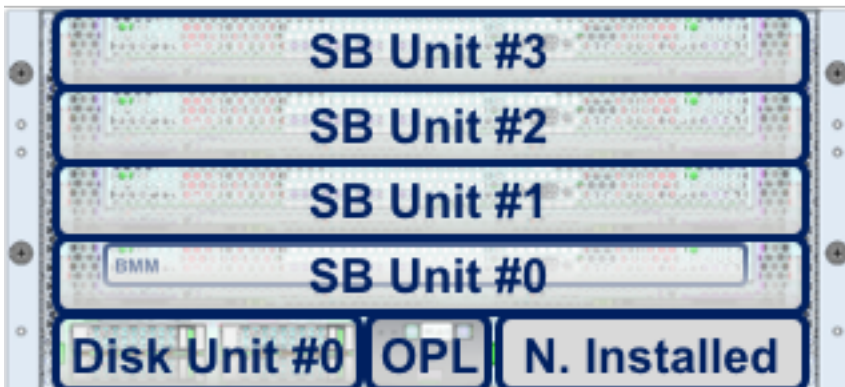
- In-memory databases
- Big data analytics
- Virtualization and virtual desktop infrastructure (VDI) workloads
- Bare-metal applications

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

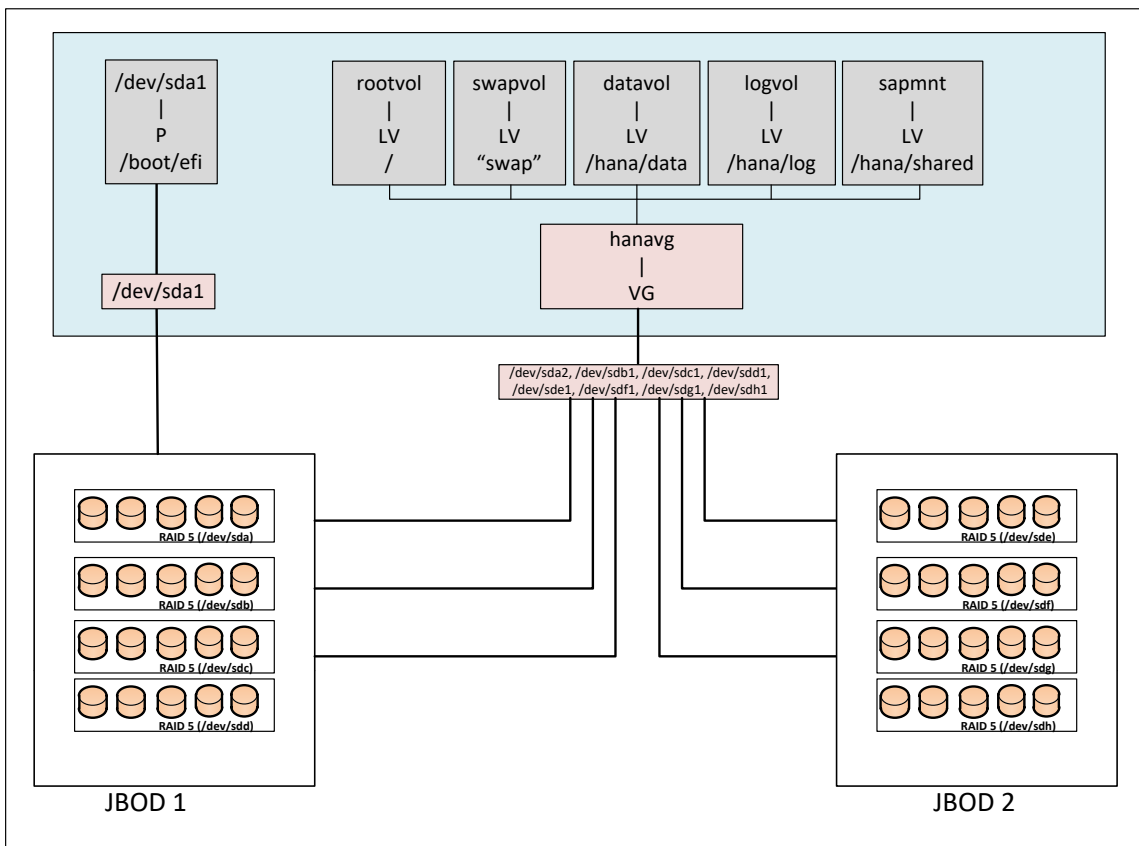
SAP defines hardware and software requirements for running SAP HANA systems. For the latest information about the CPU and memory configurations supported for SAP HANA, please refer to <https://global.sap.com/community/ebook/2014-09-02-hana-hardware/enEN/index.html>.

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 with 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 \leq 256 GB of memory
- Minimum of 512 GB for systems with \geq 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 HANA

Note: This document provides installation steps for RHEL for SAP HANA 7.3.

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



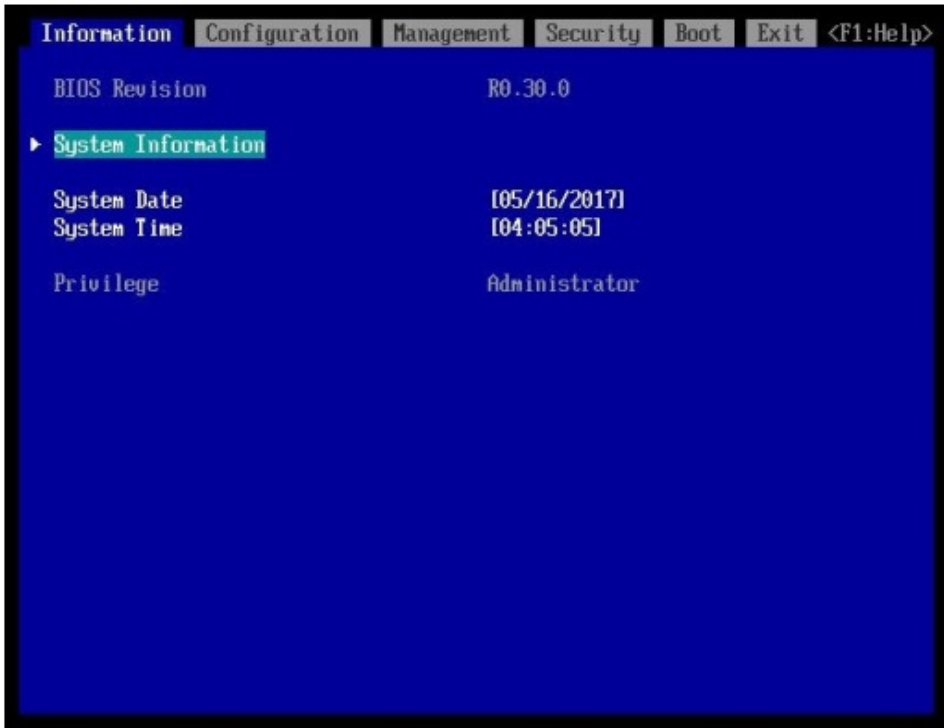
2. Power on the server and wait for the BIOS post screen.
3. When the BIOS post screen appears, press F2 for the iRMC BIOS setup (Figure 4).

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)



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

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

```

Information Configuration Management Security Boot Exit <F1:Help>
iRMC LAN Parameters Configuration

Management LAN <Enabled>
iRMC MAC Address 2C:D4:44:E6:57:74
Management LAN Port <Management>
Management LAN Speed <Auto>

Management VLAN <Disabled>

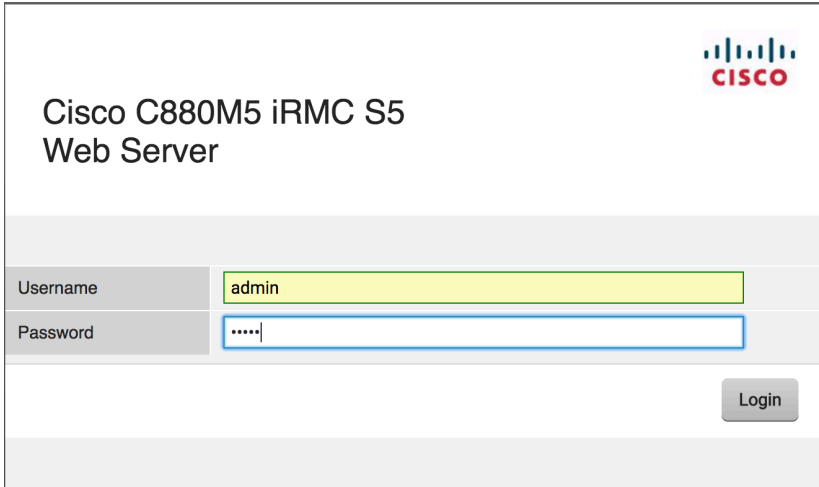
iRMC IPv4 LAN Stack <Enabled>
IP Configuration <use static configuration>
IP Address 10.24.76.135
Subnet Mask 255.255.255.0
Gateway Address 10.24.76.1

iRMC IPv6 LAN Stack <Enabled>
Link Local Address fe80::2ed4:44ff:
- fee6:5774

```

- Make sure that Management is selected for Management LAN Port.
- 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.
- Press F10 to save the configuration and exit the utility.
- Open a web browser on a computer on the same network with Java and Adobe Flash installed.
- Enter the iRMC IP address of the Cisco C880 M5 server: http://<var_iRMC_ip_address>.
- 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).

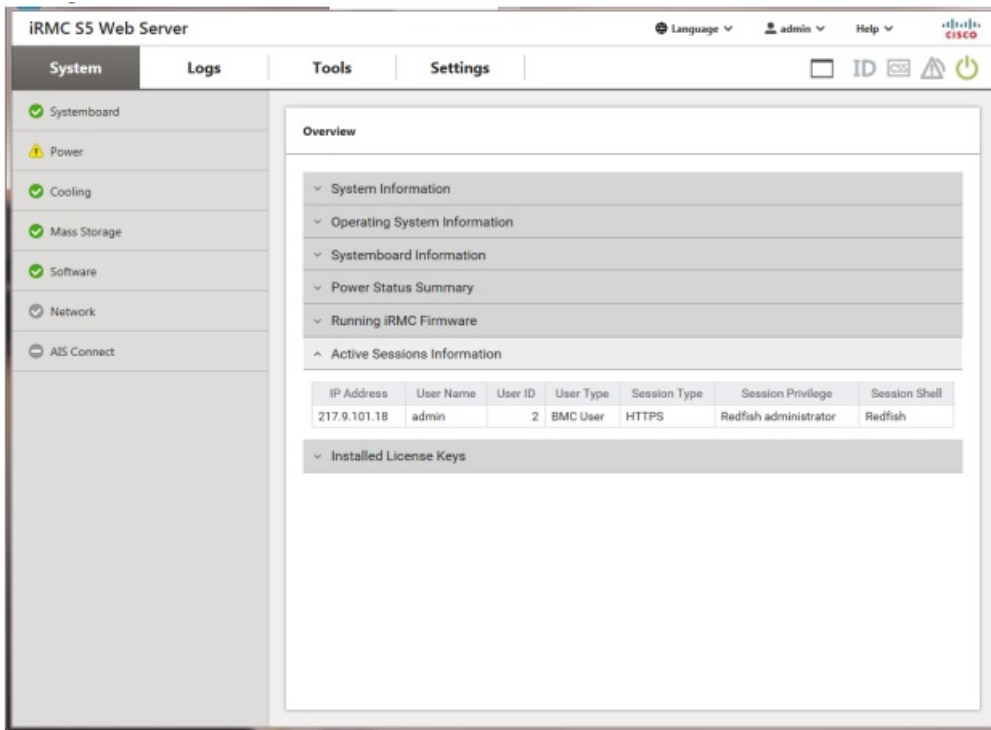
Figure 7. Logging in



The login page for the Cisco C880M5 iRMC S5 Web Server. It features the Cisco logo in the top right corner. The main heading is "Cisco C880M5 iRMC S5 Web Server". Below the heading, there are two input fields: "Username" with the value "admin" and "Password" with masked characters ".....". A "Login" button is located at the bottom right of the form area.

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

Figure 8. Cisco C880 IMC home screen



The home screen of the iRMC S5 Web Server. The title bar reads "iRMC S5 Web Server" and includes navigation links for "Language", "admin", and "Help", along with the Cisco logo. The main interface is divided into a left sidebar and a main content area. The sidebar, under the "System" tab, lists various components: Systemboard (checked), Power (warning icon), Cooling (checked), Mass Storage (checked), Software (checked), Network (checked), and AIS Connect (disabled). The main content area is titled "Overview" and contains several expandable sections: System Information, Operating System Information, Systemboard Information, Power Status Summary, Running iRMC Firmware, and Active Sessions Information. The Active Sessions Information section is expanded to show a table of active sessions.

IP Address	User Name	User ID	User Type	Session Type	Session Privilege	Session Shell
217.9.101.18	admin	2	BMC User	HTTPS	Redfish administrator	Redfish

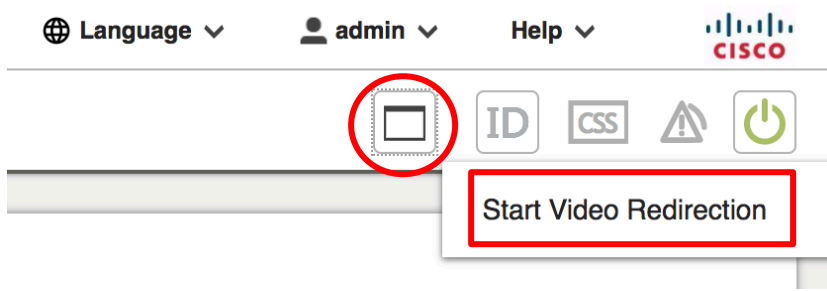
Below the table is an expandable section for "Installed License Keys".

Starting Cisco C880 M5 video redirection

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

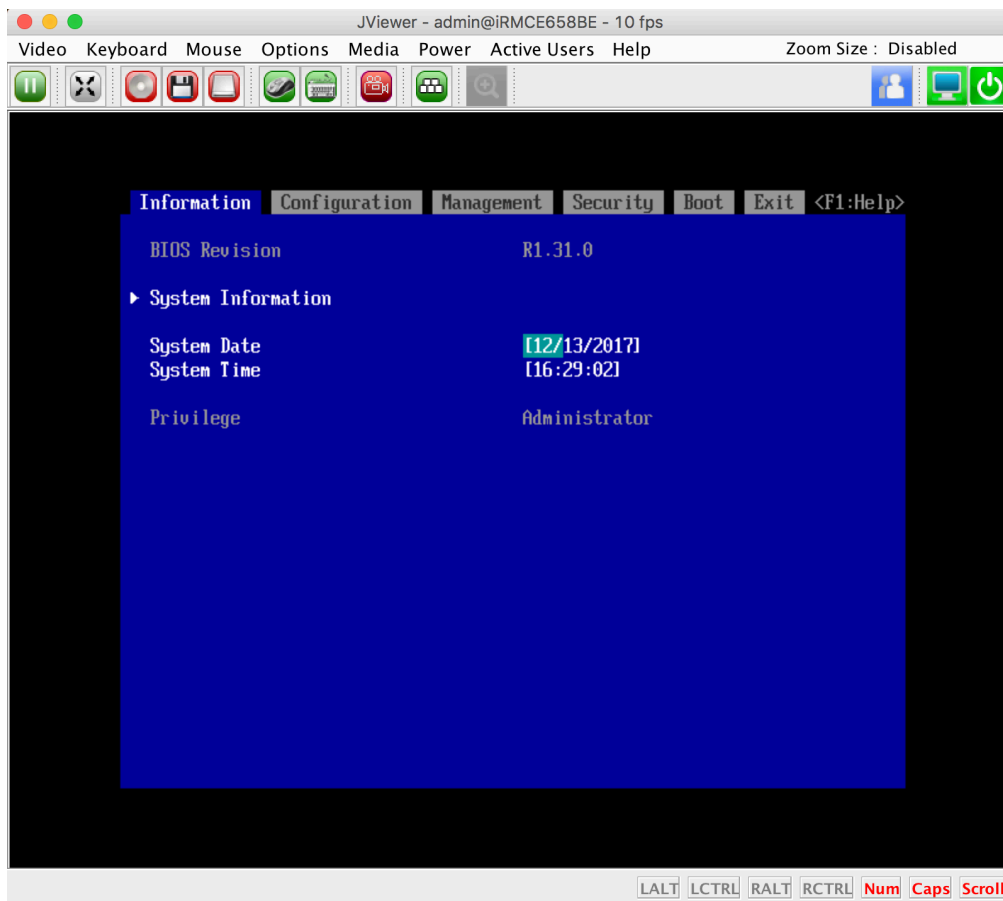
1. Log in to the iRMC with the admin credentials.
2. Click the monitor icon at the upper right (Figure 9).

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

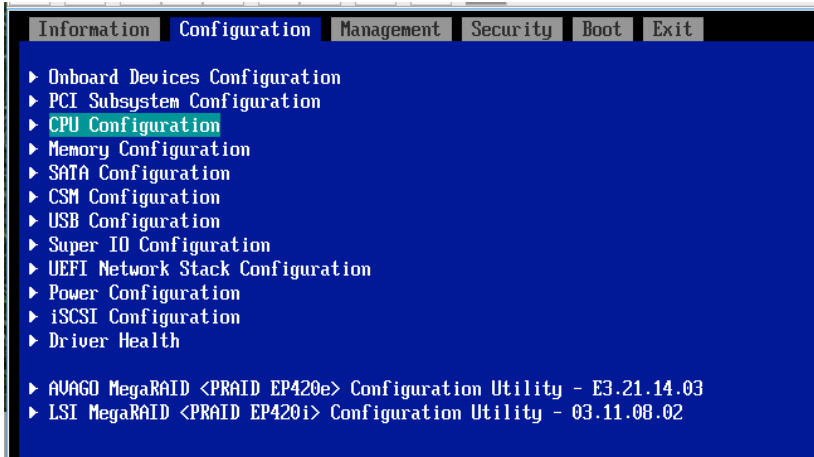


Configuring Cisco C880 M5 BIOS CPU settings

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

```

Information Configuration Management Security Boot Exit
CPU Configuration
Hyper-Threading <Enabled>
Active Processor Cores I01
(Current/Available) 24 / 24

Hardware Prefetcher <Enabled>
Adjacent Cache Line Prefetch <Enabled>
DCU Streamer Prefetcher <Enabled>
DCU Ip Prefetcher <Enabled>

Intel Virtualization Technology <Disabled>
Intel(R) VT-d <Disabled>
Intel TXT Support <Disabled>

Power Technology <Custom>
Enhanced SpeedStep <Enabled>
Turbo Mode <Enabled>
Energy Performance <Performance>
Override OS Energy Performance <Disabled>
Utilization Profile <Even>
P-State Coordination <HW_ALL>
HWP Support <Native Mode>
CPU C1E Support <Disabled>
CPU C6 Report <Disabled>
Package C State limit <No Limit>
UPI Link Frequency Select <Auto>

Perfmon and DFX devices <Disabled>
ACPI MSCT <Enabled>
Uncore Frequency Scaling <Enabled>
Data Poisoning <Enabled>
Sub NUMA Clustering <Auto>
Stale AtoS <Disabled>
LLC Dead Line Alloc <Enabled>

```


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

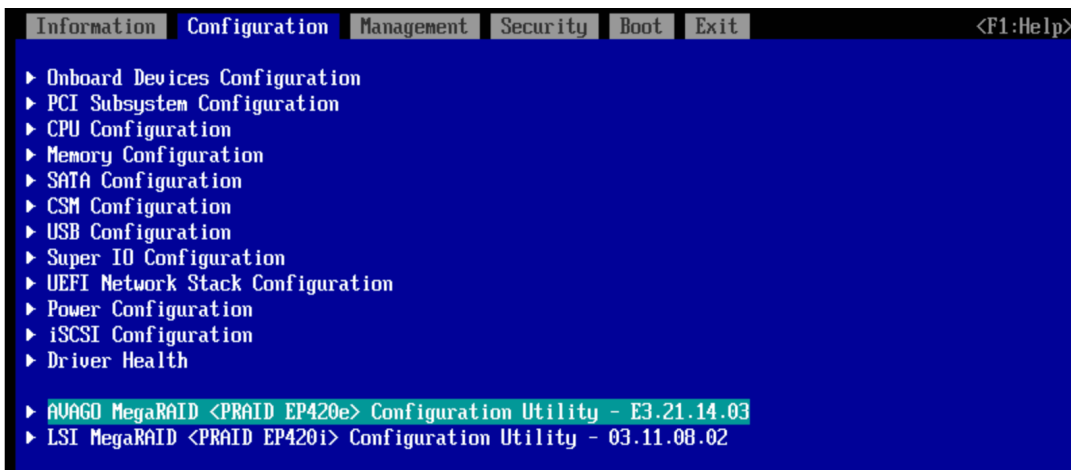
Now create eight virtual drives with RAID 5. Each virtual drive will have five drives associated with it. Follow the steps here until all eight virtual drives are created.

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



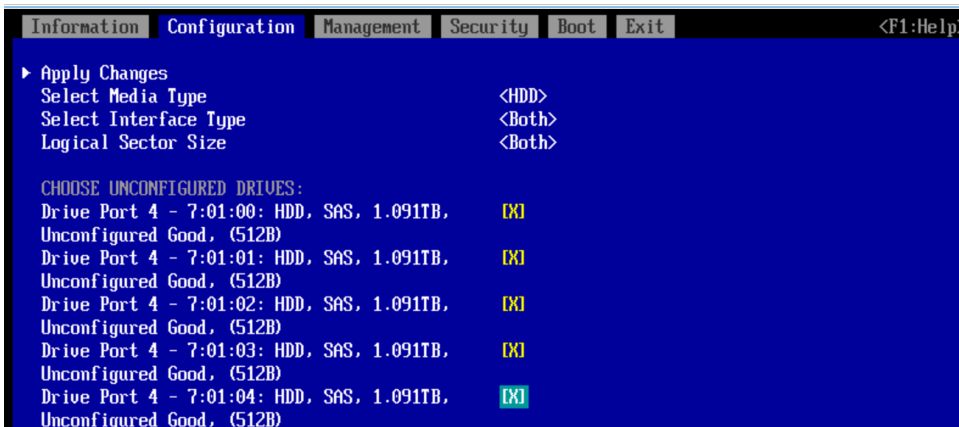
5. Select RAID5 as the RAID level and click Select Drives (Figure 17).

Figure 17. Selecting drives



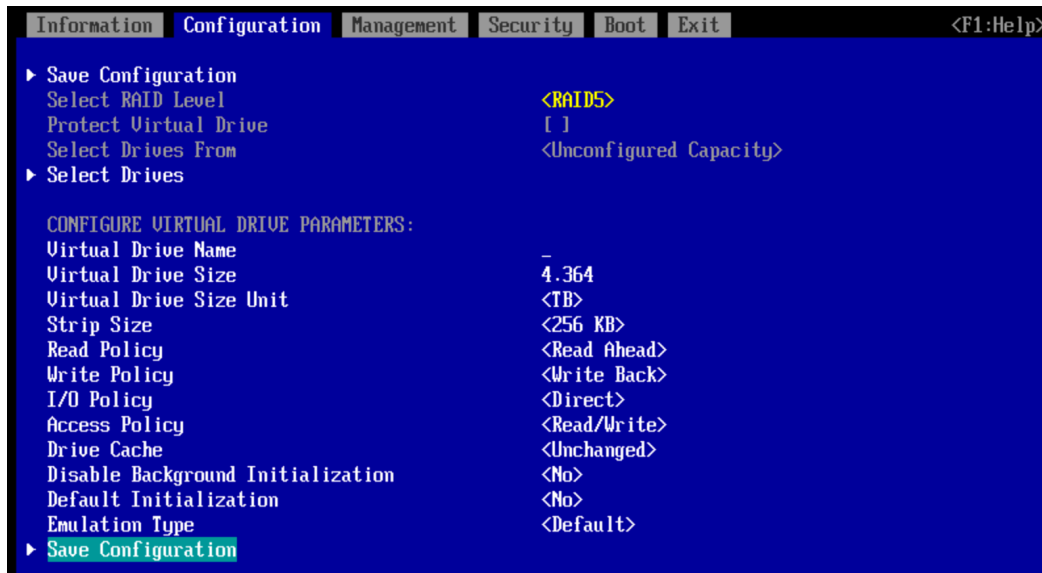
6. Select the first five drives in the list (Figure 18).

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



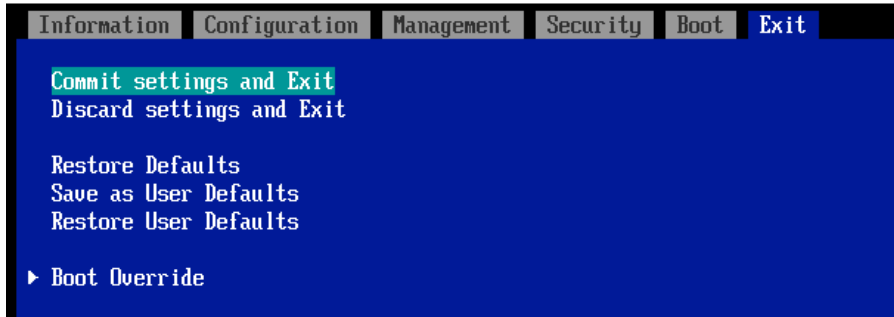
Note: You need to create eight virtual drives, each with five physical disks, so repeat the steps here until all eight virtual drives are created.

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



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



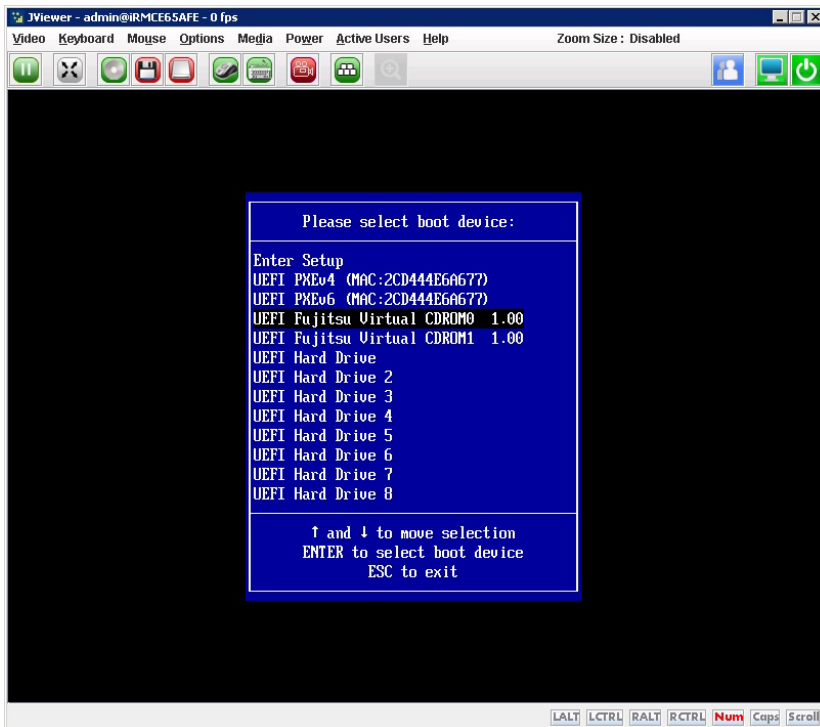
2. Browse for the SLES 12 for SAP Applications SP2 DVD ISO image and click Connect CD/DVD (Figure 23).

Figure 23. Connecting the system




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

Figure 24. Boot menu



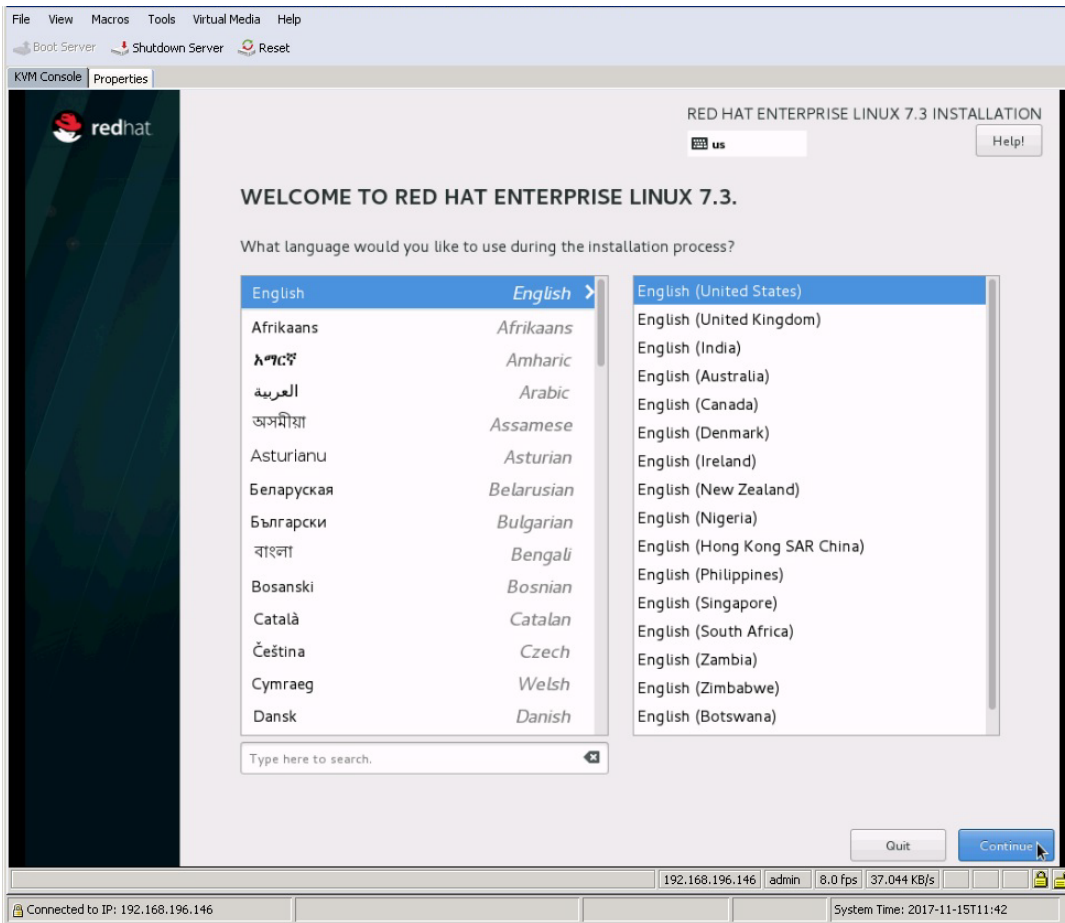
4. Choose UEFI Fujitsu Virtual CDROM0. The installation process begins (Figure 25).

Figure 25. Beginning the installation process



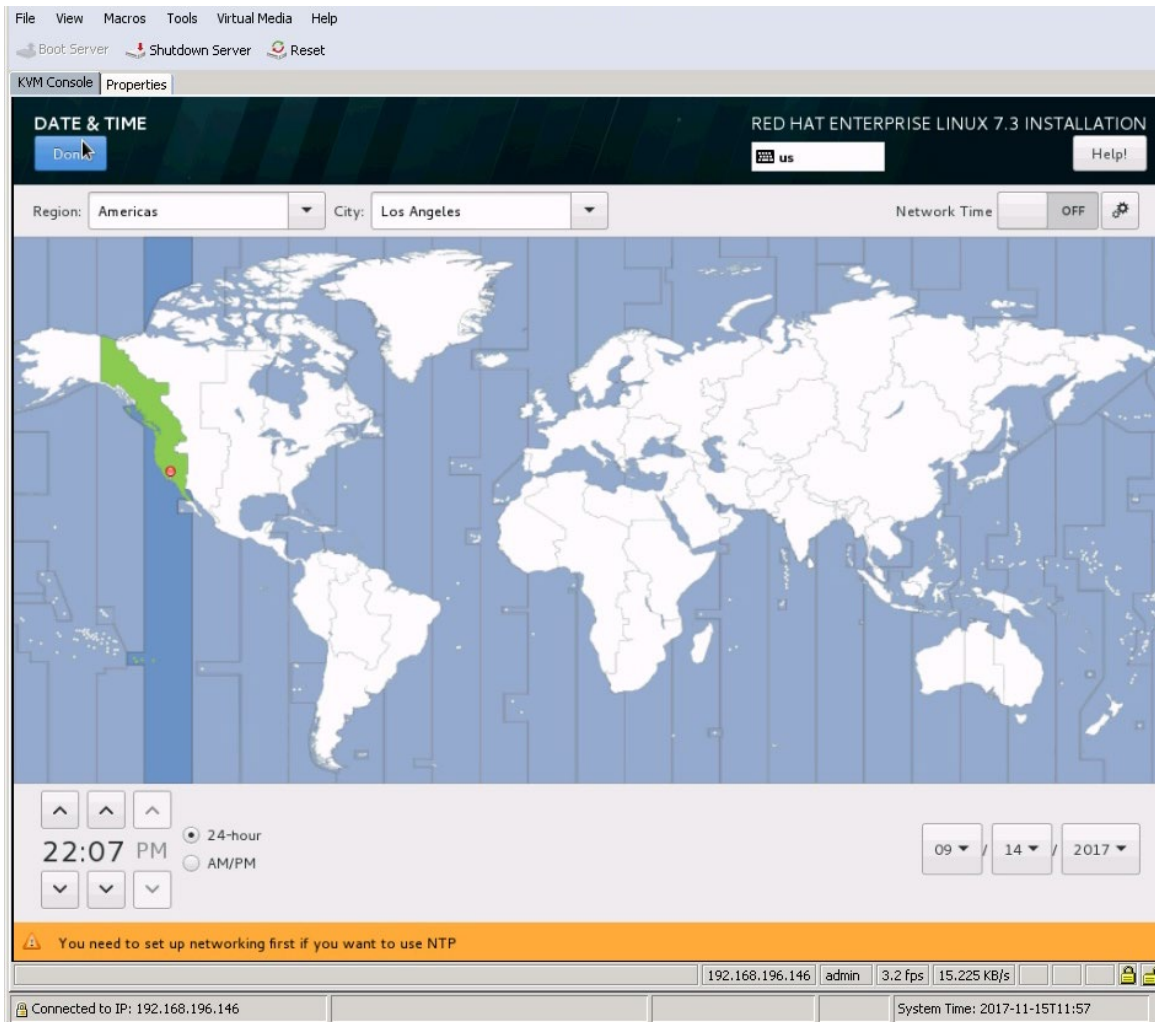
5. Select the language and keyboard layout you want (Figure 26).

Figure 26. Selecting the language and keyboard layout



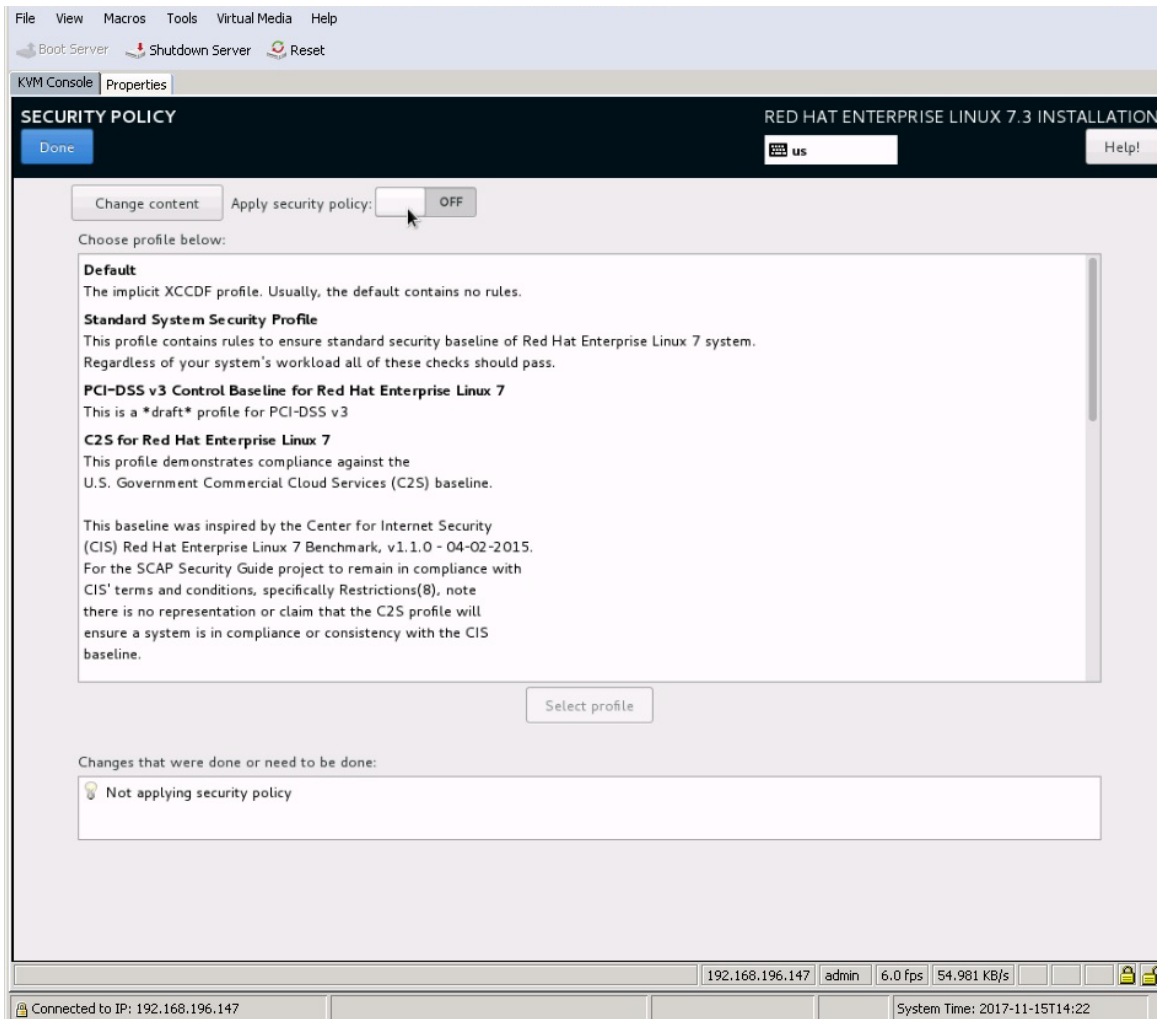
6. Click Continue. The central Installation Summary page appears. Here, you need to configure various features. Choose Localization > Date & Time. Choose the appropriate region and city (Figure 27). You will configure the Network Time Protocol (NTP) later. Click Done.

Figure 27. Setting the date and time



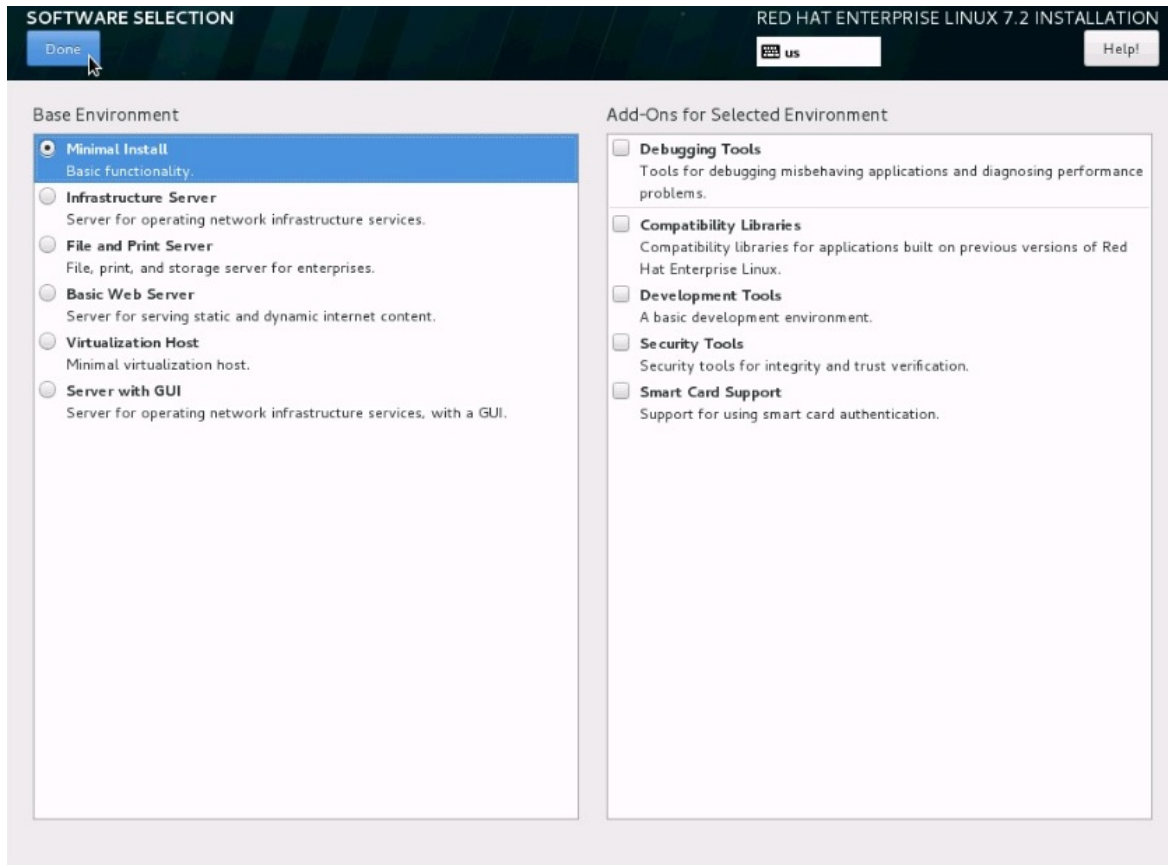
7. Choose Security > Security Policy. Turn off the security policy (Figure 28).

Figure 28. Setting security policy



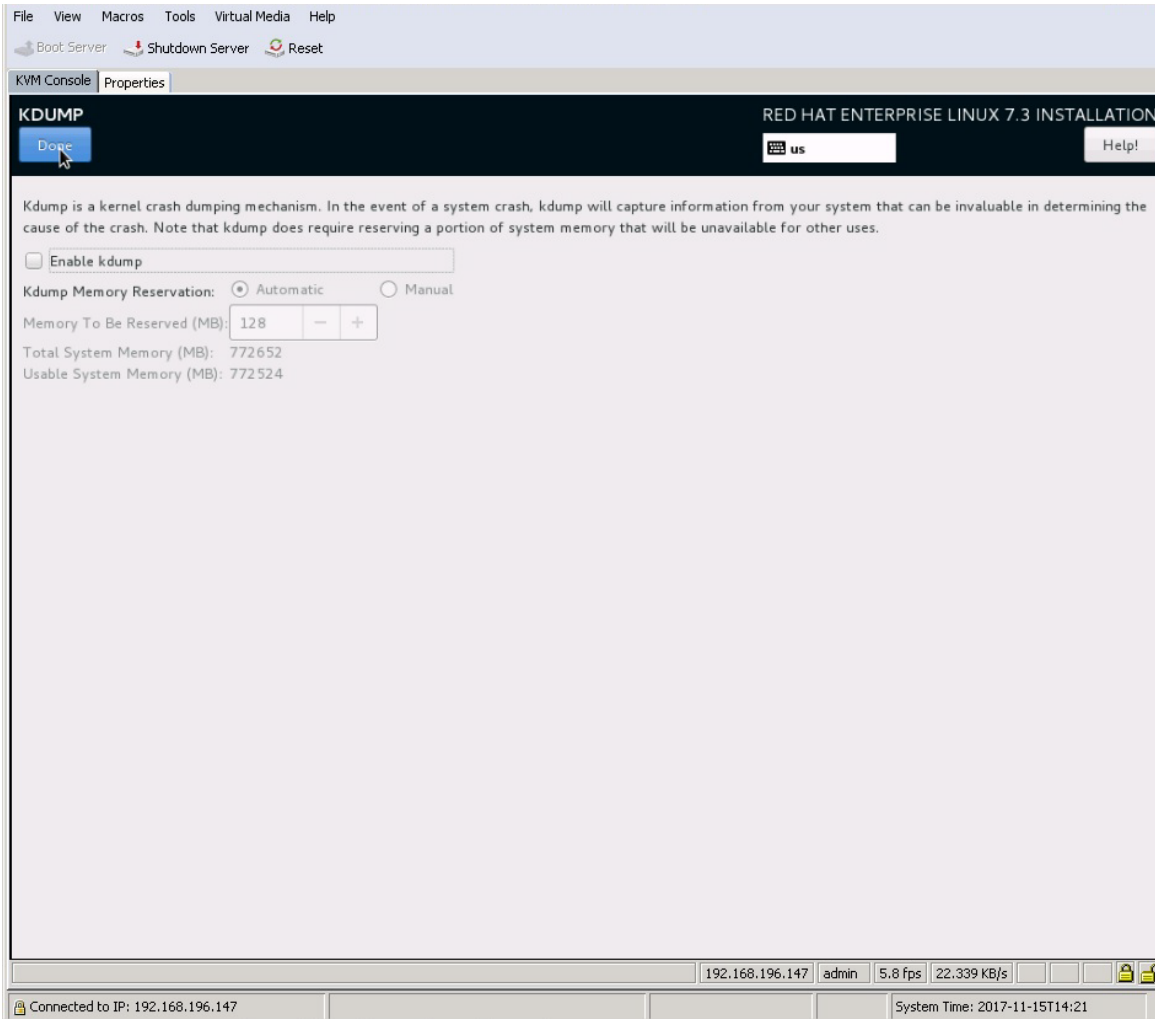
8. Select Software Selection. Retain the default selection: Minimal Install (Figure 29).

Figure 29. Software Selection page



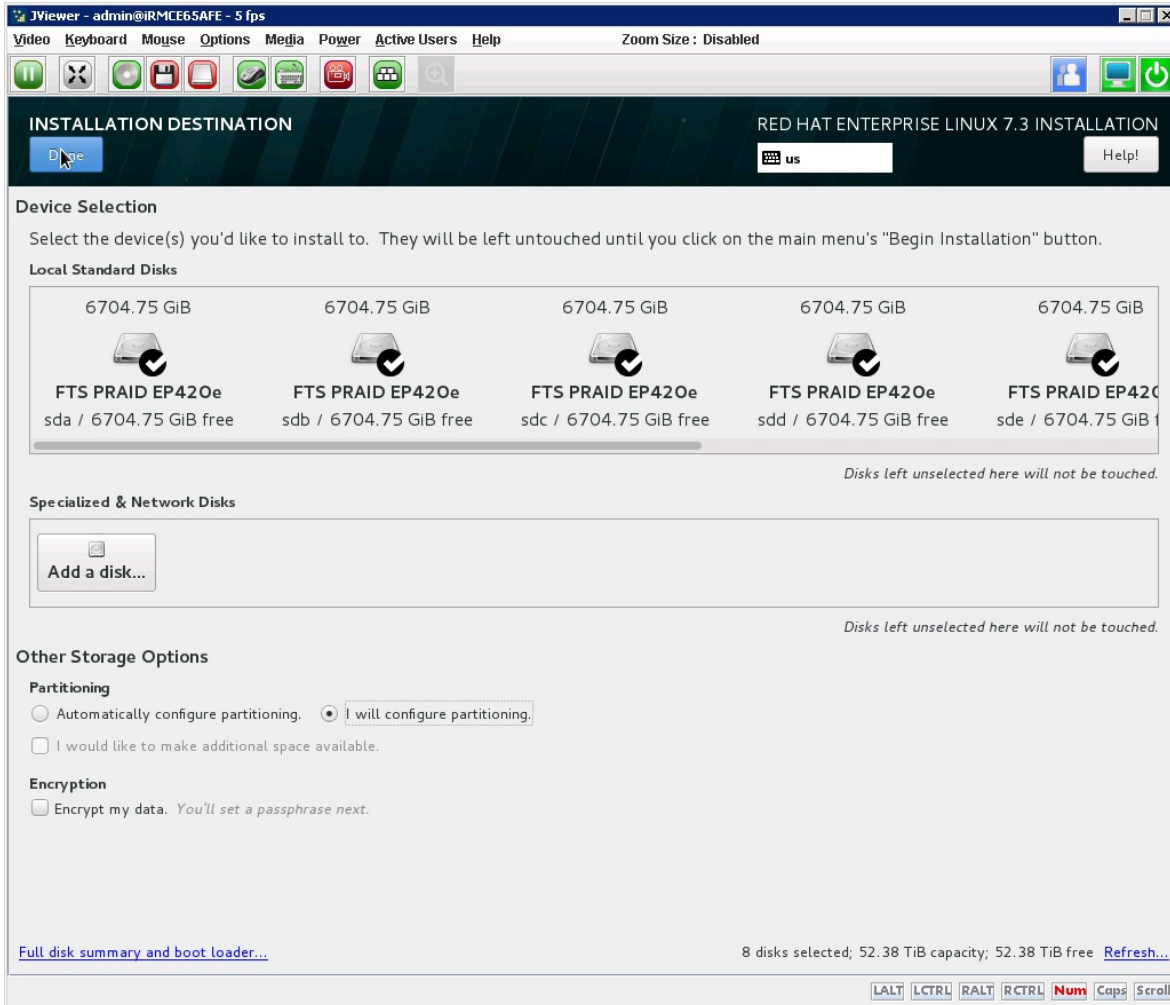
9. Select KDUMP. Deselect the Enable Kdump option to disable it (Figure 30).

Figure 30. Disabling kdump



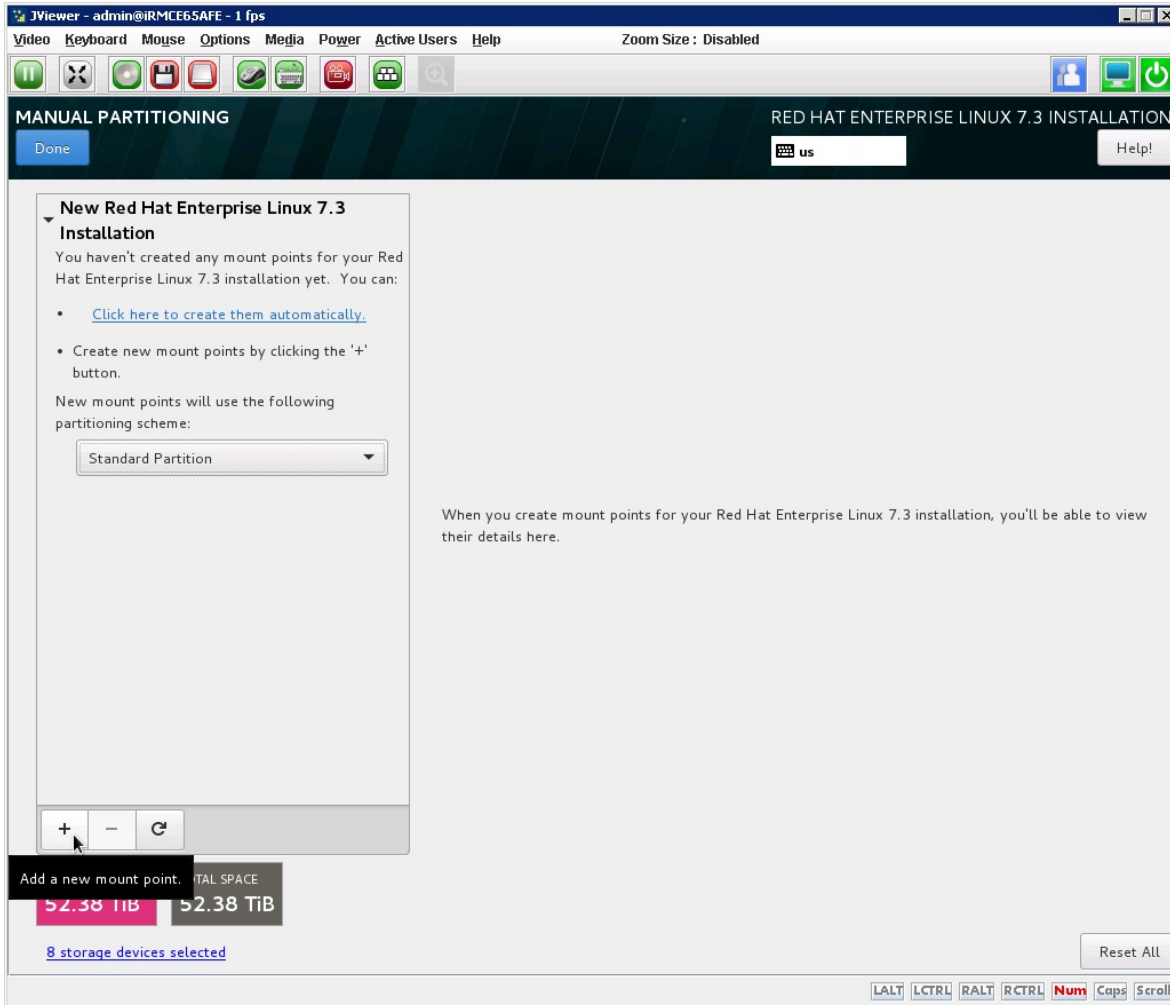
10. Choose System > Installation Destination. Under the other storage options, select the option to manually configure the disk partition layout: "I will configure partitioning." (Figure 31). Click Done.

Figure 31. Installation Destination page



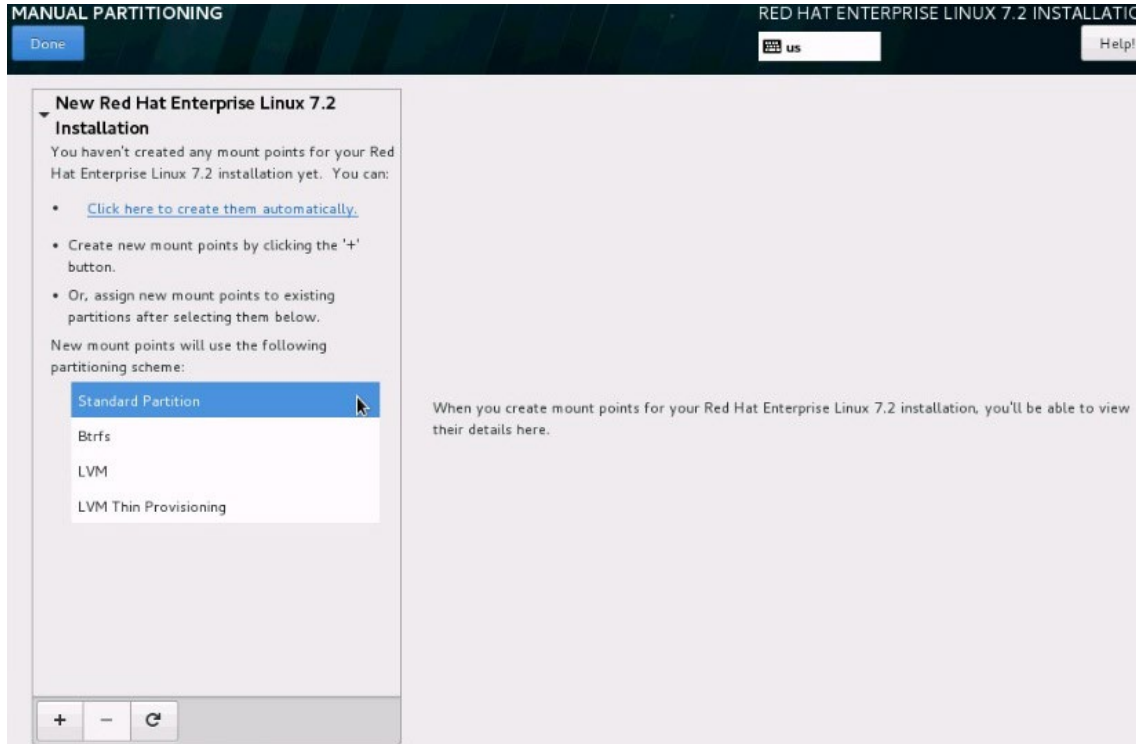
11. On the Manual Partitioning page that appears, Scroll and select all eight virtual drives and then click Done at the top left of the page (Figure 32).

Figure 32. Manually creating partitions



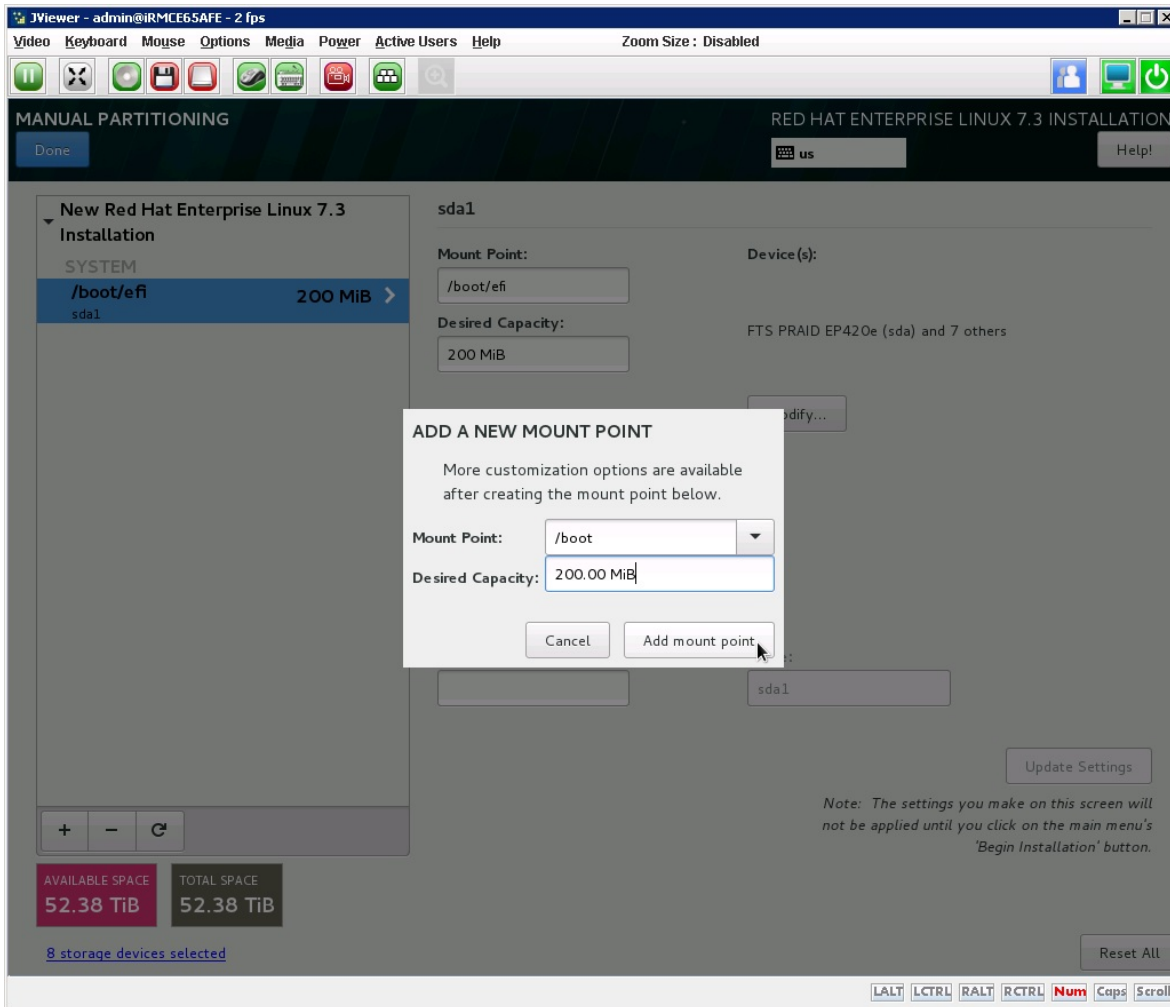
12. You will first create the /boot partition with a standard partition scheme. Change the default partition scheme from Logical Volume Manager (LVM) to Standard Partition (Figure 33).

Figure 33. Choosing the Standard Partition type



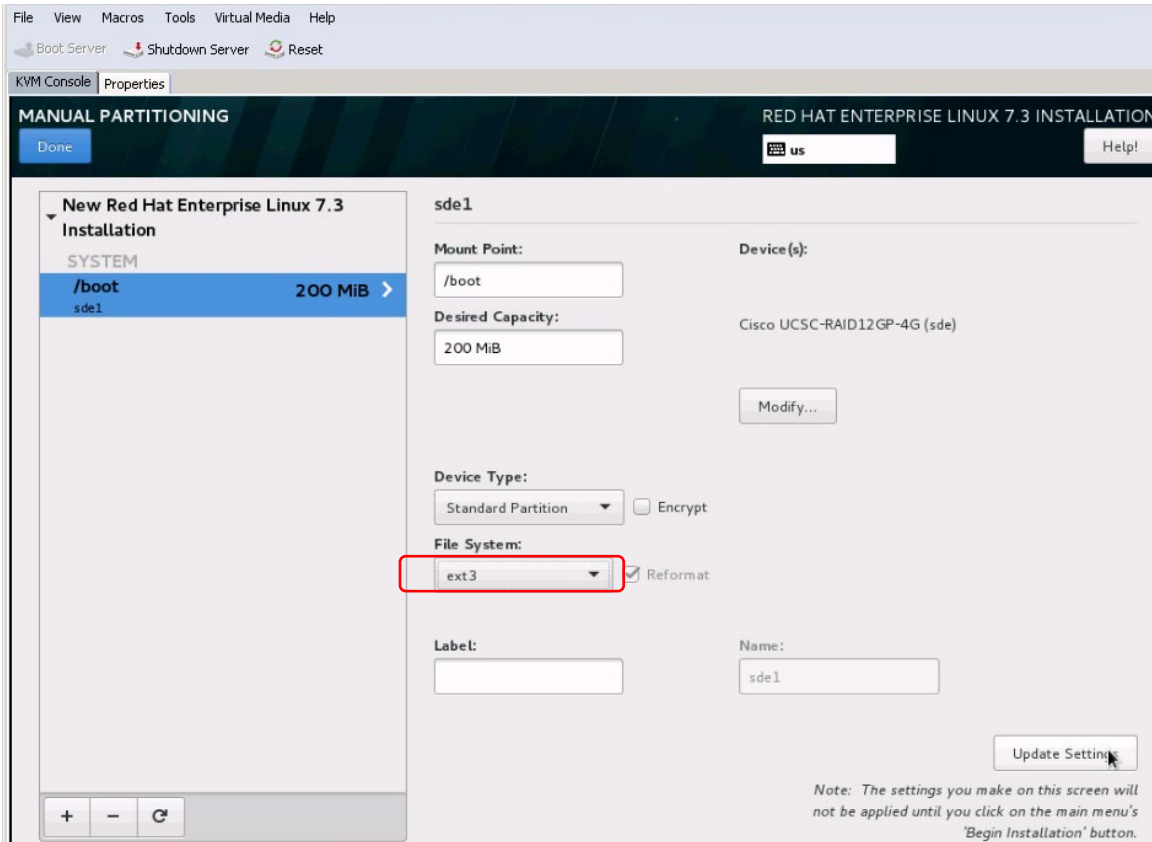
13. Click the + button and create a /boot partition with a size of 200 MiB. Click “Add mount point” (Figure 34).

Figure 34. Entering mount point and capacity information



14. Change the file system type from the default XFS to ext3 (Figure 35).

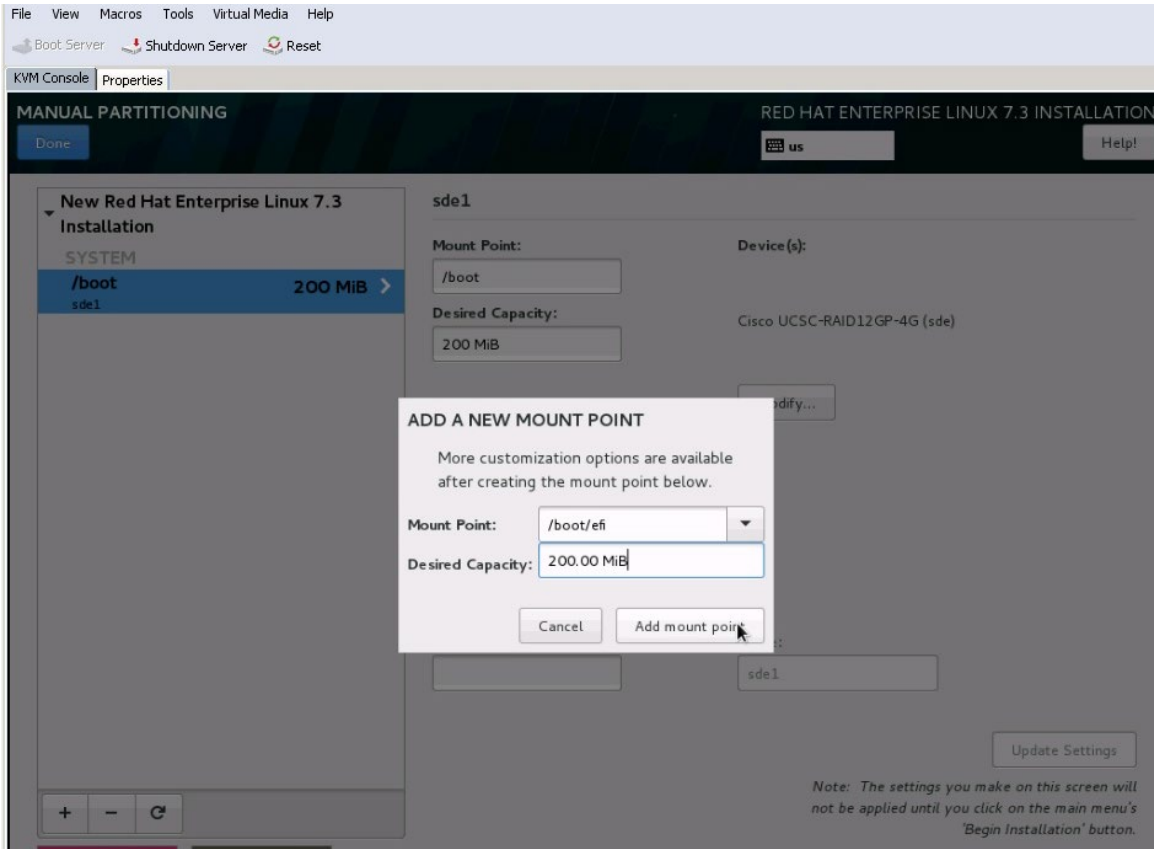
Figure 35. Modifying the file system type



15. Create /boot/efi partition of 200MiB

Click the + button, select /boot/efi for the mount point, enter 200 MiB as the desired capacity, and click “Add mount point” (Figure 36).

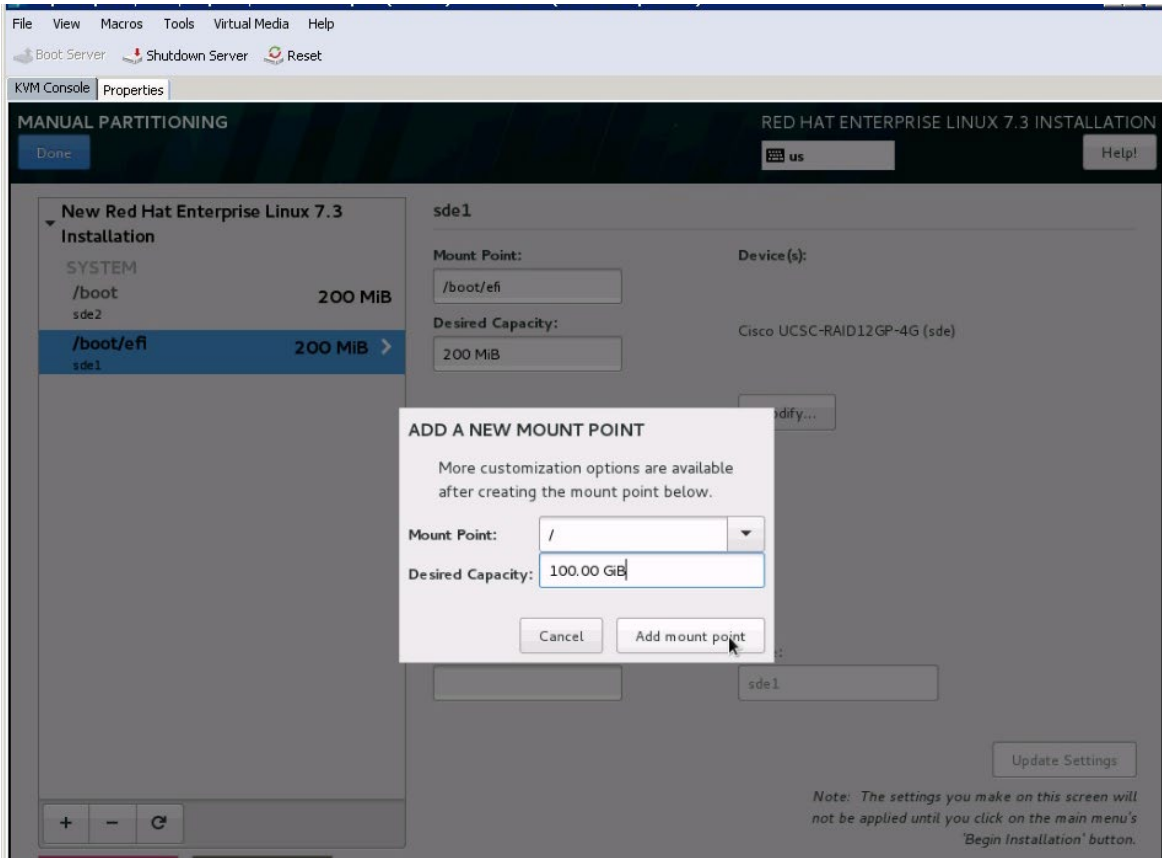
Figure 36. Creating /boot/efi partition



16. After you define the /boot and /boot/efi partitions, you will assign the remaining disk space to the LVM as a volume group (VG) and then carve out a root volume, swap volume, and HANA system-related volumes.

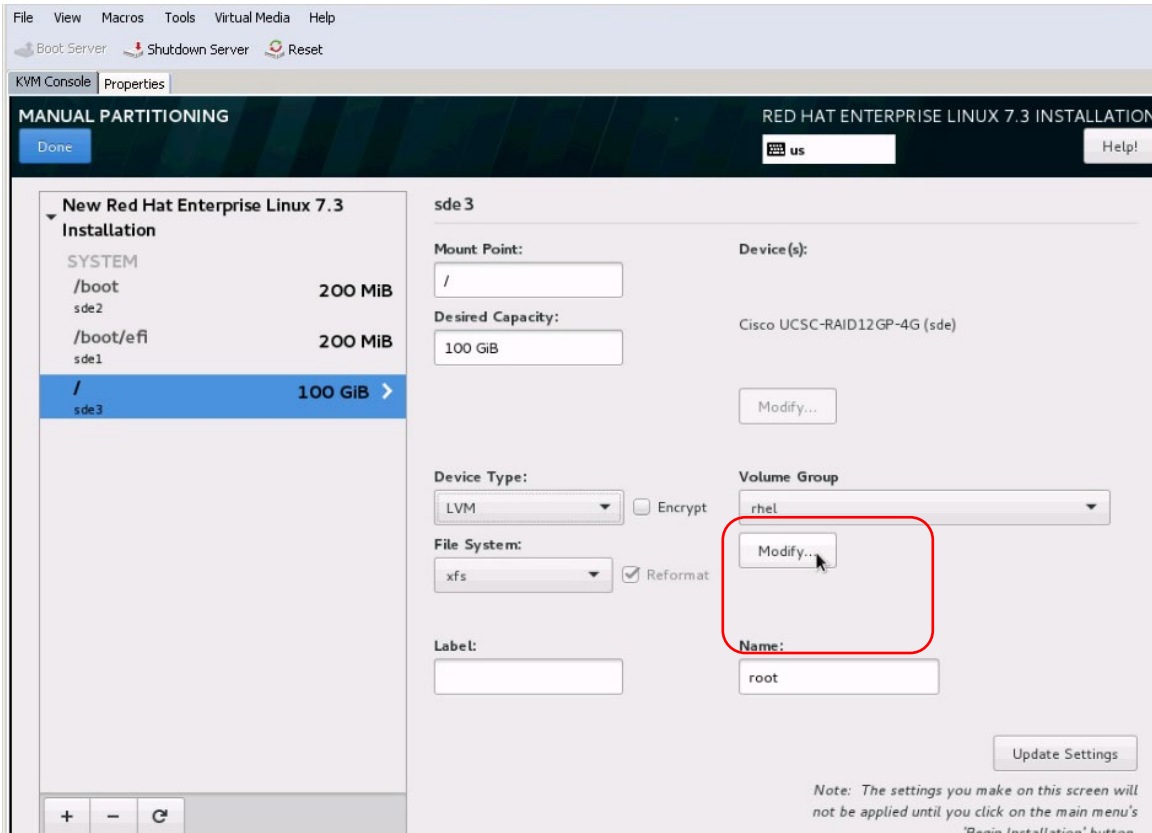
Click the + button, select “/” for the mount point, enter **100 GiB** as the desired capacity, and click “Add mount point” (Figure 37).

Figure 37. Creating the root file system with 100 GB



17. Click Modify to change the device type (Figure 38).

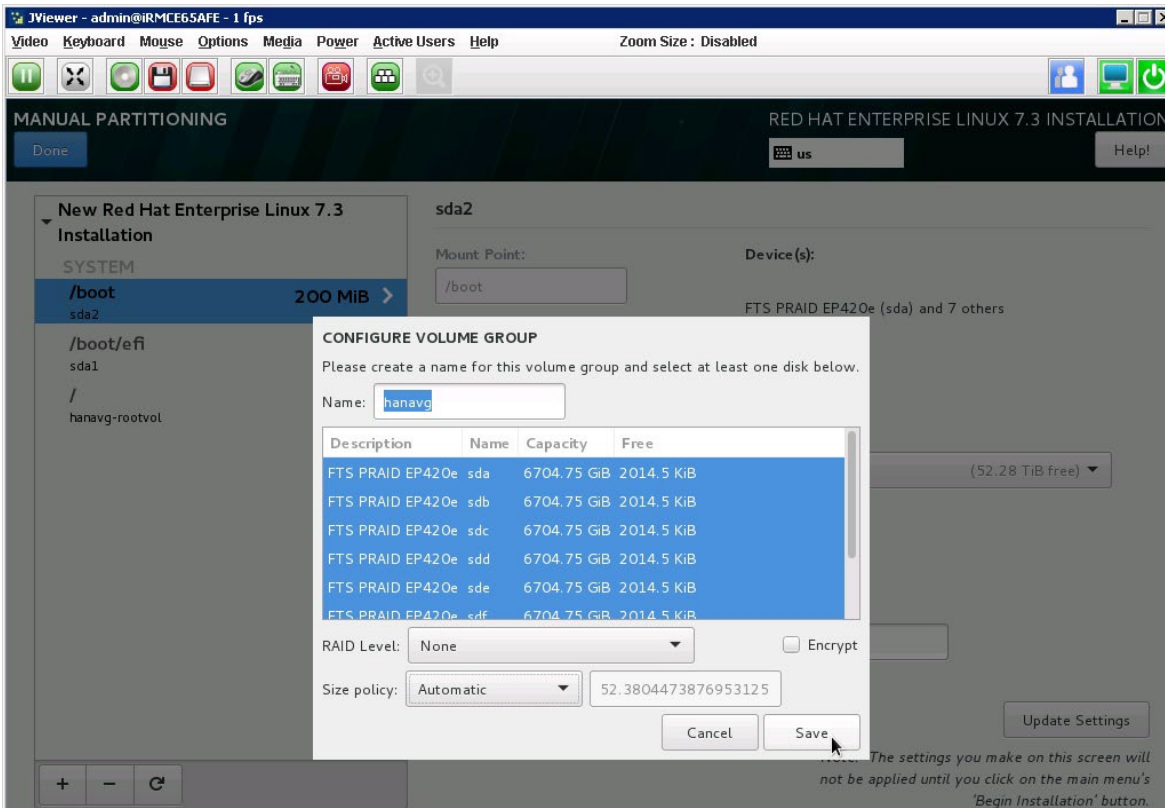
Figure 38. Preparing to change the device type to LVM



18. Change the Device Type selection from Standard Partition to LVM.

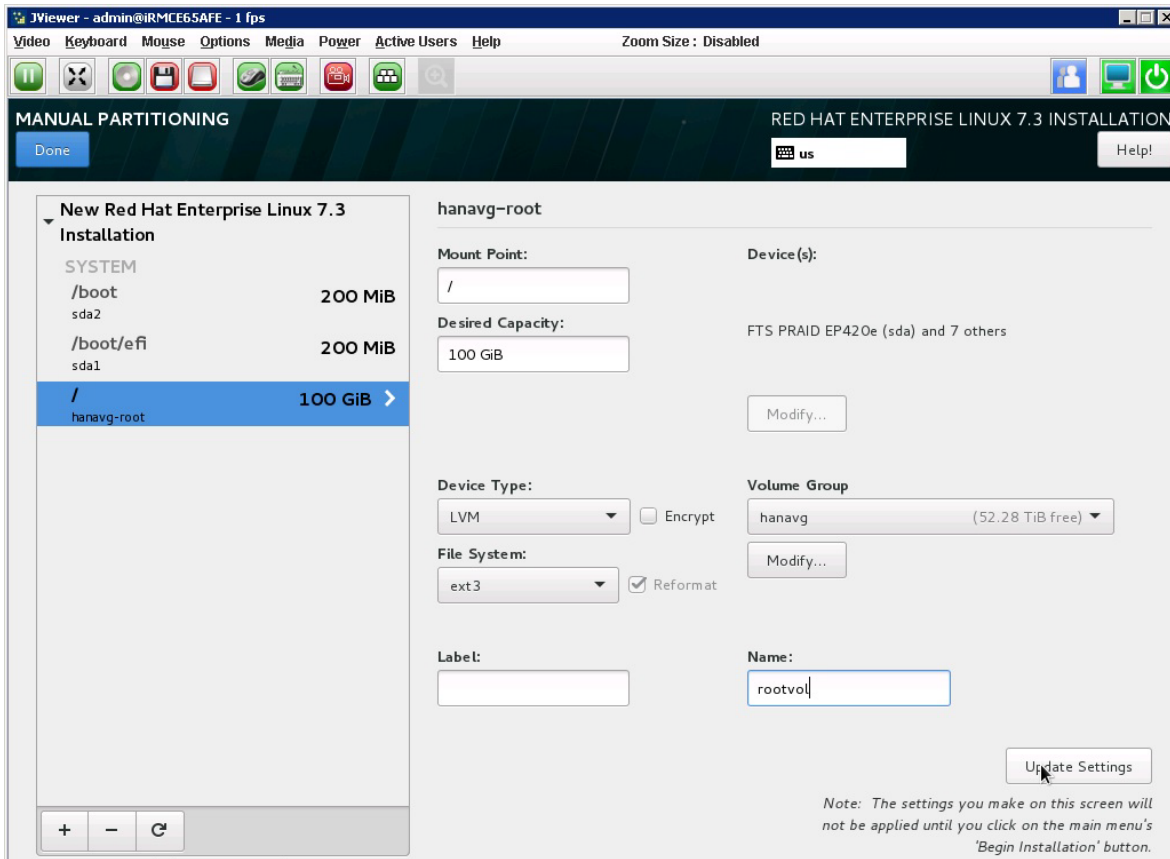
19. Change the name of the volume group from the default rhel to hanavg (Figure 39). Then click Save.

Figure 39. Configuring the volume group



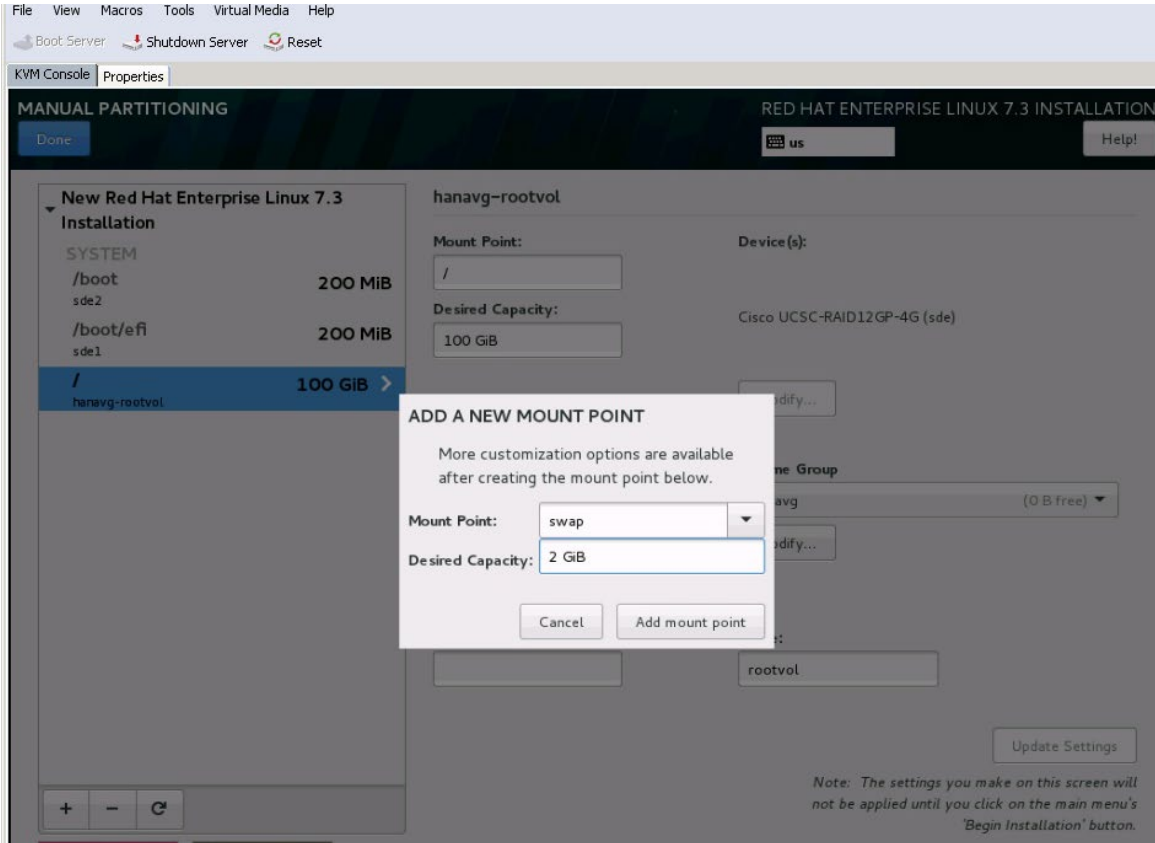
20. Change the file system type to ext3 and change the name to **rootvol**. Click Update Settings (Figure 40).

Figure 40. Updating the file system type and volume group name



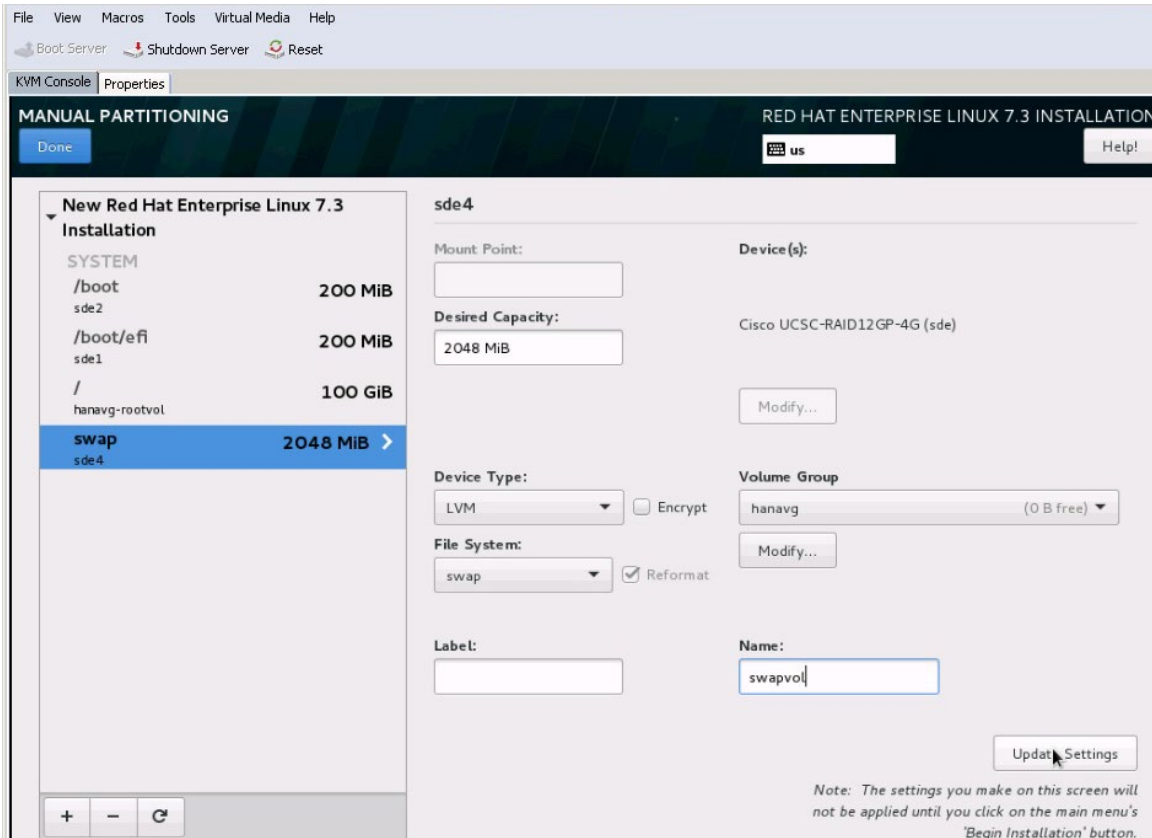
You will now create a 2 GiB swap volume. Click the + button, select “swap” for the mount point, enter 2 GiB as the desired capacity, and click “Add mount point” (Figure 41).

Figure 41. Creating a swap volume



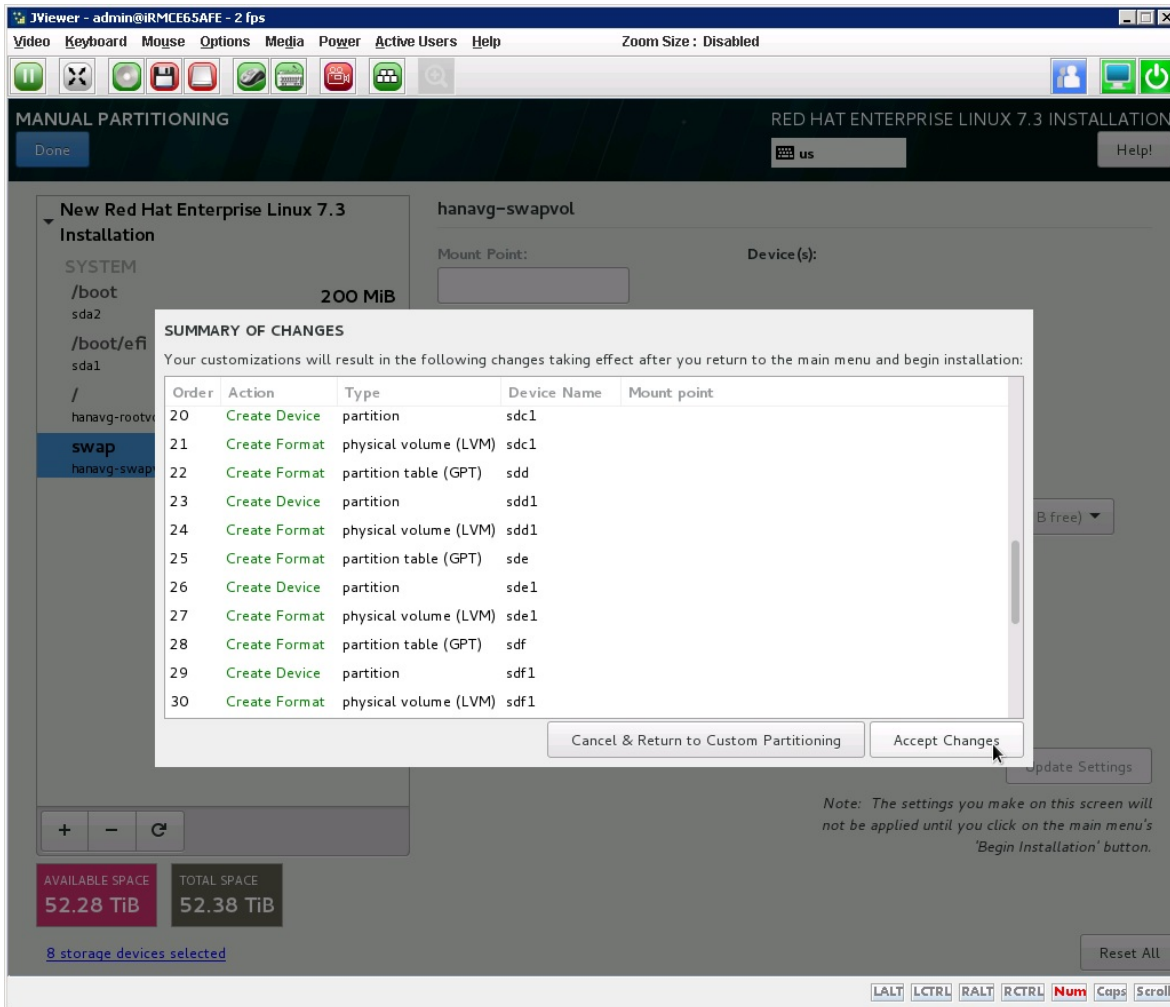
21. Change the device type to LVM, verify that hanavg is the selected volume group, and change the volume group name to `swapvol` (Figure 42).

Figure 42. Updating swap volume properties



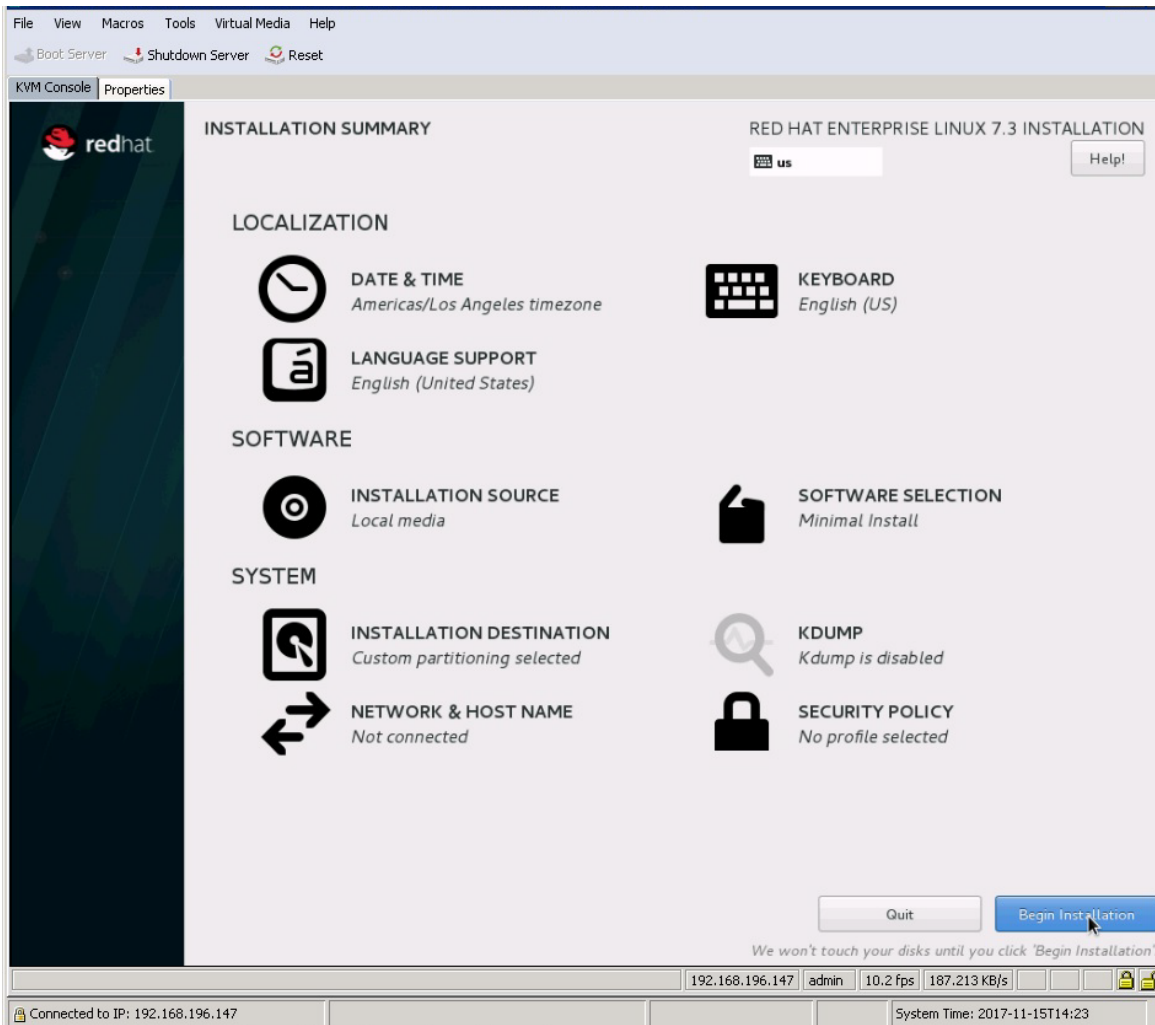
22. Click Done. A summary of the changes will appear. Click Accept Changes (Figure 43).

Figure 43. Summary of changes for manual partition configuration



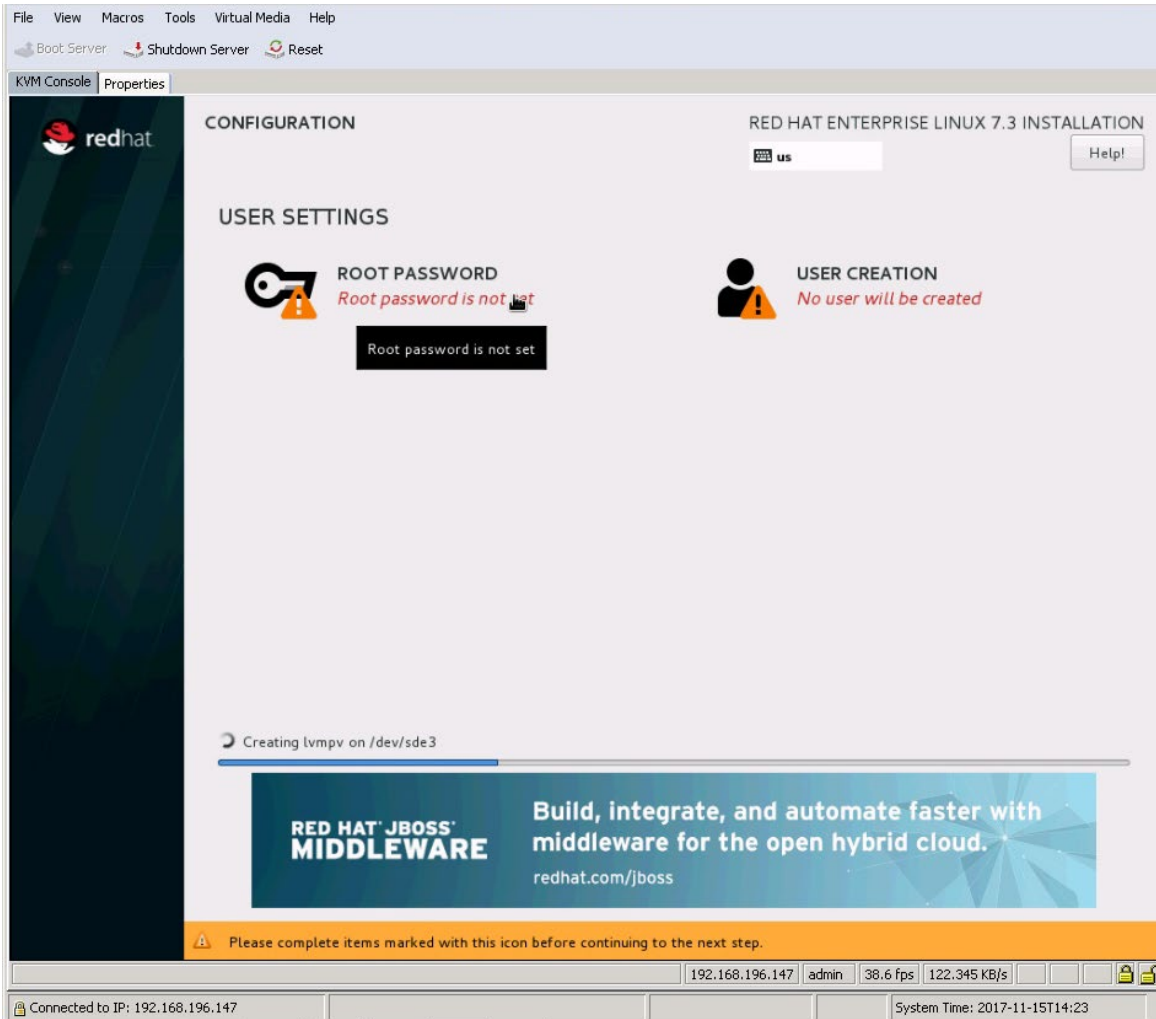
23. On the Installation Summary page that appears, click Begin Installation (Figure 44).

Figure 44. Beginning the installation



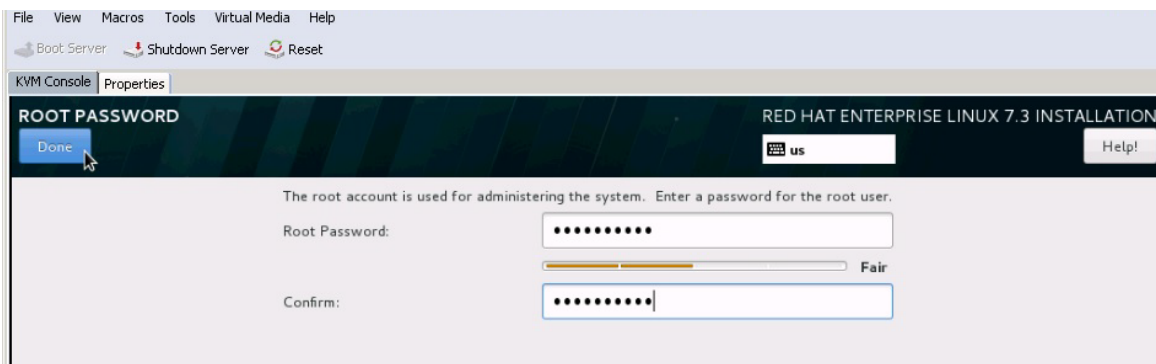
24. As the installation progresses, set the root password (Figure 45).

Figure 45. Setting the root password



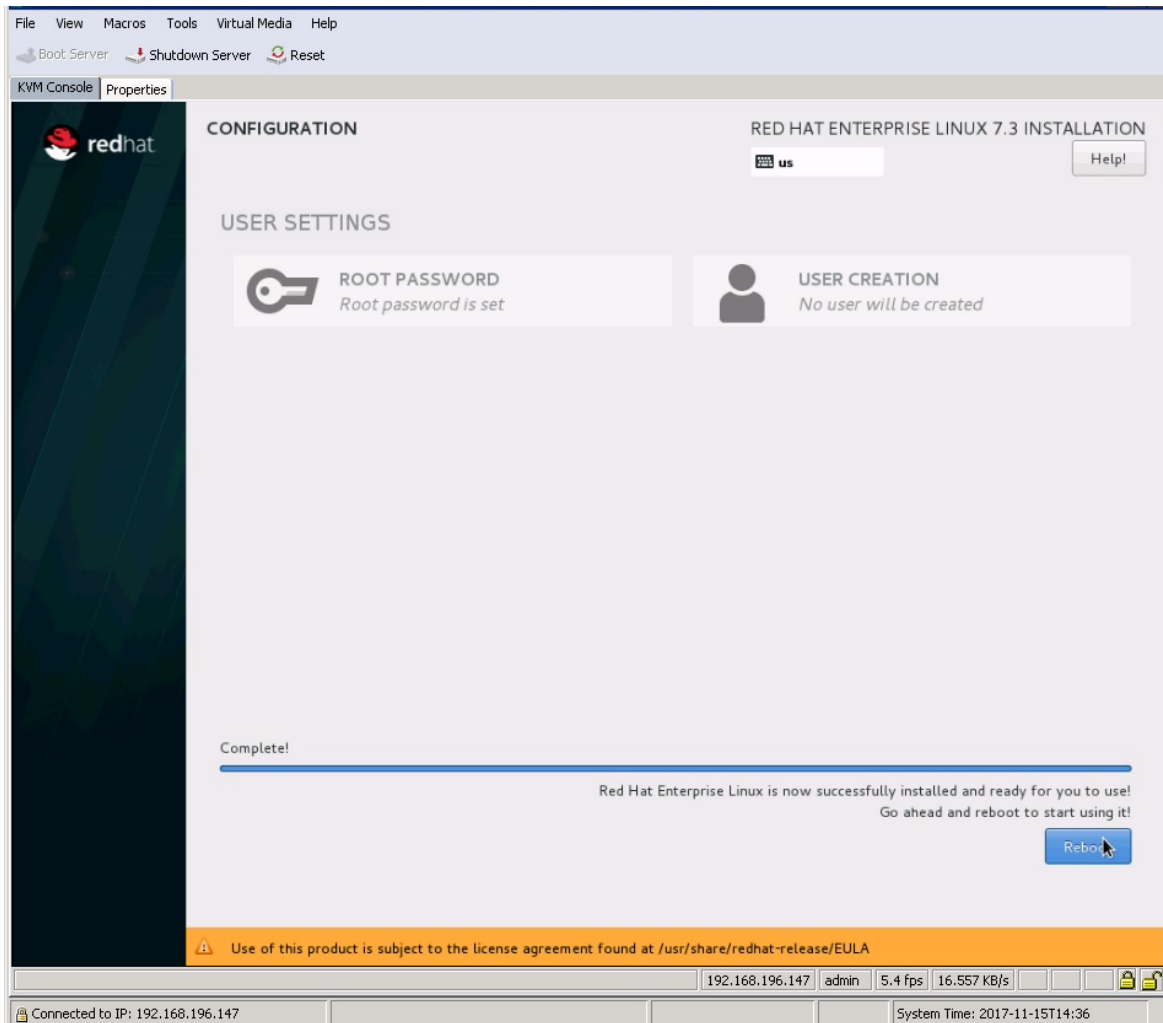
25. Enter and confirm the root password (Figure 46).

Figure 46. Entering and confirming the root user password



26. The Installation completes (Figure 47). Click Reboot.

Figure 47. Finishing the configuration



Performing post-installation OS customization

Follow the steps presented here to customize the server in preparation for SAP HANA installation.

Customizing the host name

1. Use the video console to log in to the installed system as the user **root** with the password `<<var_sys_root-pw>>`.
2. Update the `/etc/hosts` file with an entry matching the host name and IP address of the system (Figure 48).

Figure 48. Sample hosts file

```
[root@cishana01 ~]# more /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1        localhost localhost.localdomain localhost6 localhost6.localdomain6
173.36.215.118 cishana01.custdom.local cishana01
[root@cishana01 ~]#
```

3. Verify that the host name is displayed correctly.

The operating system must be configured so that the short name of the server is displayed with the command **hostname -s**, and the fully qualified host name is displayed with the command **hostname -f**. Figure 49 shows sample output.

Figure 49. Sample **hostname** command output

```
[root@cishana01 ~]# hostname
cishana01.custdon.local
[root@cishana01 ~]# hostname -s
cishana01
[root@cishana01 ~]# hostname -f
cishana01.custdon.local
[root@cishana01 ~]# hostname -d
custdon.local
```

Configuring the network

The Cisco C880 M5 Server comes with onboard 1-Gbps NICs and 10 Gigabit Ethernet PCI cards. In addition to the administration and management networks, you can optionally have networks for backup, client access, etc. You can configure additional networks based on customer-specific requirements and use cases.

1. To get an overview of the Ethernet interface configuration, use the **ip addr** command. Figure 50 shows sample output.

Figure 50. Sample **ip addr** command output

```
[root@cishana01 ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: enp72s0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:8b:31:40 brd ff:ff:ff:ff:ff:ff
3: enp73s0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:8b:31:41 brd ff:ff:ff:ff:ff:ff
4: enp135s0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:5b:de:20 brd ff:ff:ff:ff:ff:ff
5: enp136s0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:5b:de:21 brd ff:ff:ff:ff:ff:ff
6: enp65s0f0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 88:1d:fc:39:f2:12 brd ff:ff:ff:ff:ff:ff
7: enp65s0f1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 88:1d:fc:39:f2:13 brd ff:ff:ff:ff:ff:ff
    inet 173.36.215.118/24 brd 173.36.215.255 scope global enp65s0f1
        valid_lft forever preferred_lft forever
    inet6 fe80::8a1d:fcff:fe39:f213/64 scope link
        valid_lft forever preferred_lft forever
8: enp80s0f0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 88:1d:fc:39:f2:16 brd ff:ff:ff:ff:ff:ff
9: enp80s0f1: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 88:1d:fc:39:f2:18 brd ff:ff:ff:ff:ff:ff
[root@cishana01 ~]#
```

In RHEL 7, **systemd** and **udev** support a number of different naming schemes. By default, fixed names are assigned based on firmware, topology, and location information: for instance, **enp72s0**, as seen in Figure 50.

With this naming convention, names stay fixed even if hardware is added or removed. However, the names are often more difficult to read than traditional kernel-native **ethX** names: for instance, **eth0**.

Another method for naming network interfaces, biosdevnames, is also available with the installation.

2. Configure the boot parameters `net.ifnames=0 biosdevname=0` to disable both approaches to use the original kernel-native network names.
3. Also, you can disable IPv6 support at this time because this solution uses IPv4. You can accomplish this by appending `ipv6.disable=1` to `GRUB_CMDLINE_LINUX` as shown in Figure 51.

Figure 51. Sample grub file with CMDLINE parameter additions

```
[root@cishana01 ~]# vi /etc/default/grub
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="rd.lvm.lv=hanavg/rootvol rd.lvm.lv=hanavg/swapvol rhgb quiet net.ifnames=0 biosdevname=0 ipv6.disable=1"
GRUB_DISABLE_RECOVERY="true"
```

4. Run the `grub2-mkconfig` command to regenerate the `grub.cfg` file (Figure 52):

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

Figure 52. Updating the grub configuration

```
[root@cishana01 ~]# grub2-mkconfig -o /boot/grub2/grub.cfg
Generating grub configuration file ...
Found linux image: /boot/vmlinuz-3.10.0-327.28.3.el7.x86_64
Found initrd image: /boot/initramfs-3.10.0-327.28.3.el7.x86_64.img
Found linux image: /boot/vmlinuz-3.10.0-327.el7.x86_64
Found initrd image: /boot/initramfs-3.10.0-327.el7.x86_64.img
Found linux image: /boot/vmlinuz-0-rescue-2bee164fc0474586a395513638954a08
Found initrd image: /boot/initramfs-0-rescue-2bee164fc0474586a395513638954a08.img
done
[root@cishana01 ~]#
```

5. Reboot the system to make the changes take effect:

```
# reboot
```

6. After the reboot, use the KVM console to log in to the installed system as the user `root` with the password `<<var_sys_root-pw>>`.

- Run the `ip addr` command to see the interfaces in the traditional kernel-native ethX nomenclature (Figure 53).

Figure 53. Checking the interface status with the `ip addr` command

```
[root@cishana01 network-scripts]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
2: eth0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:8b:31:40 brd ff:ff:ff:ff:ff:ff
3: eth1: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:8b:31:41 brd ff:ff:ff:ff:ff:ff
4: eth2: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:5b:de:20 brd ff:ff:ff:ff:ff:ff
5: eth3: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 84:b8:02:5b:de:21 brd ff:ff:ff:ff:ff:ff
6: eth4: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 88:1d:fc:39:f2:12 brd ff:ff:ff:ff:ff:ff
7: eth5: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 88:1d:fc:39:f2:13 brd ff:ff:ff:ff:ff:ff
8: eth6: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 88:1d:fc:39:f2:16 brd ff:ff:ff:ff:ff:ff
9: eth7: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc mq state DOWN qlen 1000
    link/ether 88:1d:fc:39:f2:18 brd ff:ff:ff:ff:ff:ff
[root@cishana01 network-scripts]#
```

- A close observation of the output reveals that the previous IP address setting was lost due to the changes in the interface naming you just implemented. You will have to again find the interface that has the uplink connectivity. Check the link status using the `ethtool` command to ascertain the interface that is connected to the management network (Figure 54).

Figure 54. Using the `ethtool` command to check the link status

```
[root@cishana01 network-scripts]# for i in `seq -w 0 7`; do ethtool eth$i | grep 'Link detected';done
Link detected: no
Link detected: no
Link detected: no
Link detected: no
Link detected: no
Link detected: yes
Link detected: no
Link detected: no
[root@cishana01 network-scripts]#
```

- Assign `<<var_mgmt_ip_address>>` as the IP address and enter `<<var_mgmt_ip_mask>>` as the subnet mask for the available interface (eth5 in the example in Figure 54). You can use this configuration temporarily until you port this interface to a high-availability bond device and create another interface with Cisco virtual interface card (VIC) 10-Gbps ports.
- Go to the network configuration directory and create a configuration for eth5 as shown in the following example:

```
#cd /etc/sysconfig/network-scripts
#vi ifcfg-eth5
DEVICE=eth5
TYPE=Ethernet
ONBOOT=yes
BOOTPROTO=static
IPV6INIT=no
```

```
USERCTL=no
NM_CONTROLLED=no
IPADDR=<<var_mgmt_ip_addr>>
NETMASK=<<var_mgmt_nw_netmask>>
```

11. Add the default gateway:

```
#vi /etc/sysconfig/network
NETWORKING=yes
HOSTNAME=<<var_hostname.domain>>
GATEWAY=<<var_mgmt_gateway_ip>>
```

Configuring the network time

Be sure that the time on all components used for SAP HANA is synchronized. Use the same NTP configuration on all systems:

```
#vi /etc/ntp.conf
...
server <NTP-SERVER1 IP>
server <NTP-SERVER2 IP>

#service ntpd stop
#ntpdate ntp.example.com
#service ntpd start
#chkconfig ntpd on
#chkconfig ntpdate on
```

Configuring the Domain Name System

Configure the Domain Name System (DNS) based on the local requirements.

A sample configuration is shown here. Add the DNS IP address if it is required to access the Internet.

```
#vi /etc/resolv.conf

DNS1=<<IP of DNS Server1>>
DNS2=<<IP of DNS Server2>>
DOMAIN= <<Domain_name>>
```

Updating the Red Hat system and customizing the OS for SAP HANA

Before you can customize the OS for SAP HANA, you need to update the Red Hat system.

1. Update the Red Hat repository.

To patch the system, you must first update the repository. Note that the installed system doesn't include any update information. To patch the Red Hat system, the system must be registered and attached to a valid subscription. The following code will register the installation and update the repository information:

```
#subscription-manager register --auto-attach
Username: <<username>>
Password: <<password>>
```

2. To list the repositories subscribed to, use the following command:

```
#yum repolist
```

Update only the OS kernel and firmware packages to the latest release that appeared in RHEL 7.3. Set the release version to 7.3:

```
#subscription-manager release --set=7.3
```

3. Apply the latest updates for RHEL 7.3. Typically, the kernel is updated as well.

```
#yum -y update
```

4. Reboot the machine and use the new kernel.

5. Install the base package group:

```
#yum -y groupinstall base
```

6. Install dependencies in accordance with the SAP HANA Server Installation and Update Guide. Install the numactl package if the benchmark HWCCT is to be used.

```
#yum install gtk2 libicu xulrunner sudo tcsh libssh2 expect cairo graphviz iptraf-ng krb5-
workstation krb5-libs libpng12 nfs-utils lm_sensors rsyslog openssl PackageKit-gtk3-module
libcanberra-gtk2 libtool-ltdl xorg-x11-xauth numactl xfsprogs net-tools bind-utils screen
compat-sap-c++-6 compat-sap-c++-5
```

7. Disable SELinux.

To ensure that SELinux is fully disabled, modify the file `/etc/selinux/config`:

```
# sed -i 's/\(SELINUX=enforcing\|SELINUX=permissive\) /SELINUX=disabled/g' /etc/selinux/config
```

For compatibility reasons, four symbolic links are required:

```
#ln -s /usr/lib64/libssl.so.0.9.8e /usr/lib64/libssl.so.0.9.8
#ln -s /usr/lib64/libssl.so.1.0.1e /usr/lib64/libssl.so.1.0.1
#ln -s /usr/lib64/libcrypto.so.0.9.8e /usr/lib64/libcrypto.so.0.9.8
#ln -s /usr/lib64/libcrypto.so.1.0.1e /usr/lib64/libcrypto.so.1.0.1
```

8. Configure `tuned` to use the profile `sap-hana`:

```
# yum install tuned-profiles-sap-hana tuned
# systemctl start tuned
# systemctl enable tuned
# tuned-adm profile sap-hana
```


9. Disable the abort and crash dump features:

```
# systemctl disable abrt-d
# systemctl disable abrt-ccpp
# systemctl stop abrt-d
# systemctl stop abrt-ccpp
```

a. Disable core file creation. To disable core dumps for all users, open `/etc/security/limits.conf` and add the following lines:

```
* soft core 0
* hard core 0
```

b. Enable the `sapsys` group to create an unlimited number of processes:

```
echo "@sapsys soft nproc unlimited" > /etc/security/limits.d/99-sapsys.conf
```

10. To avoid problems with the firewall during SAP HANA installation, you can disable the firewall completely with the following commands:

```
# systemctl stop firewalld
# systemctl disable firewalld
```

11. Configure the network time and date.

Make sure that NTP and its utilities are installed and that `chrony` is disabled:

```
# yum -y install ntp ntpdate
# systemctl stop ntpd.service
# systemctl stop chronyd.service
# systemctl disable chronyd.service
```

a. Edit `/etc/ntp.conf` and make sure that the server lines reflect your NTP servers:

```
# grep ^server /etc/ntp.conf
server ntp.example.com
server ntp1.example.com
server ntp2.example.com
```

b. Force an update to the current time:

```
# ntpdate ntp.example.com
```

c. Enable and start the NTP and date (NTPD) service:

```
# systemctl enable ntpd.service
# systemctl start ntpd.service
# systemctl restart systemd-timedated.service
```

- d. Double-check that the NTP service is enabled:

```
# systemctl list-unit-files | grep ntp
ntpd.service enabled
ntpdate.service disabled
```

- e. The **ntpdate** script adjusts the time according to the NTP server every time the system comes up. This process occurs before the regular NTP service is started and helps ensure an exact system time even if the time deviation is too large to be compensated for by the NTP service.

```
# echo ntp.example.com >> /etc/ntp/step-tickers
# systemctl enable ntpdate.service
```

Tuning the OS for SAP HANA—Adapting SAP Notes

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

12. Apply the following SAP Notes settings as instructed.

[2292690—SAP HANA DB: Recommended OS settings for RHEL 7](#)

13. Optionally, remove old kernels after the OS update:

```
package-cleanup --oldkernels --count=1
```

14. Reboot the server after applying the SAP notes

```
# reboot
```

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

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** i.e **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
hdbcompilesrv, HDB Compilesrv, GREEN, Running, 2016 04 13 08:51:56, 866:37:31, 41837
hdbindexsrv, HDB Indexsrv, GREEN, Running, 2016 04 13 08:52:00, 866:37:27, 41863
hdbnamesrv, HDB Namesrv, 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-bwl:/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-
bwl
skladm      41691  41683  0.0 582888 290988 \_ /usr/sap/SKL/HDB00/cishana01-
bwl/trace/hdb.sapSKL_HDB00 -d -nw -f /usr/sap/SKL/HDB00/cishana01-bwl/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-bwl -D -u skladm
skladm@cishana01-bwl:/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 Secure Shell (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: SKL, Status: Error
cishana01:~ #
```

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

```
cishana01:~ # /usr/sap/hostctrl/exe/saphostctrl -function ListInstances
Inst Info : SKL - 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

19.07.3016 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 HANA appliance, enter the following command:

```
cishana01:~ # su - skladm
```

```
cishana01:/usr/sap/SKL/HDB00> HDB info
USER      PID  PPID %CPU   VSZ   RSS COMMAND
skladm   61208 61207  1.6  13840  2696 -sh
skladm   61293 61208  0.0  11484  1632  \_ /bin/sh /usr/sap/SKL/HDB00/HDB info
skladm   61316 61293  0.0   4904   872   \_ ps fx -U skladm -o user,pid,ppid,pcpu,vsz,rss,args
skladm    8532   1  0.0  20048  1468 sapstart pf=/hana/shared/SKL/profile/SKL_HDB00_cishana01
skladm    8545  8532  1.5 811036 290140  \_ /usr/sap/SKL/HDB00/cishana01/trace/hdb.sapHAN_HDB00 -d
-nw -f /usr/sap/SKL/HDB00/cis
skladm   11579  8545  6.6 16616748 1789920  \_ hdbnameserver
skladm   11580  8545  1.5 5675392 371984  \_ hdbpreprocessor
skladm   11581  8545 10.9 18908436 6632128  \_ hdbindexserver
skladm   11582  8545  8.7 17928872 3833184  \_ hdbstatisticsserver
skladm   11583  8545  7.4 17946280 1872380  \_ hdbxsengine
skladm   11584  8545  0.0 203396 16000  \_ sapwebdisp_hdb
pf=/usr/sap/SKL/HDB00/cishana01/wdisp/sapwebdisp.pfl -f /usr/sap/H
skladm   11585  8545  1.5 15941688 475708  \_ hdbcompileserver
skladm    8368   1  0.0 216268 75072 /usr/sap/SKL/HDB00/exe/sapstartsrv
pf=/hana/shared/SKL/profile/SKL_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 skladm
cishana01:/usr/sap/SKL/HDB00> HDB info
USER      PID  PPID %CPU   VSZ   RSS COMMAND
skladm    61477 61476 2.0  13840 2692 -sh
skladm    61562 61477 0.0  11484 1632 \_ /bin/sh /usr/sap/SKL/HDB00/HDB info
skladm    61585 61562 0.0   4904  872  \_ ps fx -U skladm -o user,pid,ppid,pcpu,vsz,rss,args
skladm    8368  1  0.0 216784 75220 /usr/sap/SKL/HDB00/exe/sapstartsrv
pf=/hana/shared/SKL/profile/SKL_HDB00_cishana01 -D -u
cishana01:/usr/sap/SKL/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 - skladm
cishana01:/usr/sap/SKL/HDB00> HDB stop
hdbdaemon will wait maximal 300 seconds for NewDB services finishing.
Stopping instance using: /usr/sap/SKL/SYS/exe/hdb/sapcontrol -prot NI_HTTP -nr 00 -function
StopWait 400 2

19.07.3016 19:10:37

Stop

OK
```

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

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

StartService
Impromptu CCC initialization by 'rscpCInit'.
    See SAP note 1266393.

OK

OK

Starting instance using: /usr/sap/SKL/SYS/exe/hdb/sapcontrol -prot NI_HTTP -nr 00 -function
StartWait 2700 2

19.07.3016 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 UCS C480 M5 server's IMC IP address	173.36.215.117
<<var_cimc_ip_netmask>>	Cisco UCS C480 M5 server's IMC network netmask	255.255.255.0
<<var_cimc_gateway_ip>>	Cisco UCS C480 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 UCS C480 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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)