Cisco UCS C890 M5 Rack Server for SAP HANA with Red Hat Linux Enterprise Server for SAP Solutions

Last Updated: September 24, 2021

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 datasets in real time, companies are unable to use this information to their advantage.

The Cisco® solution for SAP HANA using the high-performance Cisco Unified Computing System™ (Cisco UCS®) C890 M5 Server is well-suited for mission-critical business scenarios such as for medium- to large-size SAP HANA database management systems including high availability and secure multitenancy. The Cisco solution provides a robust platform for SAP HANA workloads as a single-node (scale-up) SAP HANA instance.

The SAP HANA platform provides a scalable database with advanced analytical capabilities and intuitive application-development tools in an in-memory data platform. SAP HANA Tailored Data Center Integration (TDI) supports nonvolatile Intel Optane Persistent Memory (PMem) that brings together the low latency of memory and the persistence of storage. PMem modules provide faster access to data and retain data across power cycles.

Red Hat Enterprise Linux (RHEL) for SAP Solutions combines the reliability, scalability, and performance of Red Hat Enterprise Linux with additional content specific to operating SAP landscapes. Based on an intelligent operating system, Red Hat delivers a stable, security-focused, high-performance foundation for critical systems. Consistency across bare-metal, virtualized, hybrid, and cloud infrastructure enables you to deploy and manage both traditional and cloud-native workloads using the same tools and interfaces.

Solution Overview

This section introduces the SAP HANA on Cisco UCS C890 M5 rack server solution.

Introduction

The Cisco Unified Computing System™ (Cisco UCS®) C890 M5 Server for SAP HANA solution provides a 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 using the Red Hat Enterprise Linux (RHEL) for SAP Solutions operating system.

The SAP HANA in-memory database combines transactional and analytical SAP workloads and hereby takes advantage of the low-cost main memory, data-processing capabilities of multicore processors, and faster data access. The Cisco UCS C890 M5 rack server is equipped with the second-generation Intel® Xeon™ scalable processors, and it supports DDR4-only memory modules or mixture with Intel Optane Persistent Memory.

The solution combines Cisco hardware and software components optimized for SAP HANA workloads as preconfigured and validated architecture. The Cisco UCS C890 M5 provides the advantages of an integrated computing and storage stack while the standard configuration already enables the server for SAP HANA TDI deployments. SAP HANA TDI 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.

Audience

The audience for this document includes, but is not limited to, field consultants, professional services staff, partner engineers, and customers who want to deploy SAP HANA on a Cisco UCS C890 M5 rack server. 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 provides high-level steps to deploy and configure SAP HANA on a Cisco C890 M5 rack server. It showcases one of the variants of this Cisco solution for SAP HANA. Although readers of this document are expected to have access to the knowledge bases of Red Hat and SAP and sufficient knowledge to install and configure the products used, configuration details important to the solution deployment are provided.

Solution summary

The Cisco C890 M5 rack server for SAP HANA scale-up solution is prevalidated with local storage technology. Details about the system architecture and design considerations are available in the Cisco UCS C890 M5 rack server for SAP HANA white paper accompanying this document. It is a defined set of hardware and software that serves as an integrated foundation for nonvirtualized workloads.

All components are connected and configured according to best practices of Cisco, Red Hat, and SAP SE. The Cisco UCS C890 M5 rack server data sheet summarizes the server specifications and Table 1 shows the proposed configurations for the specific SAP HANA use cases.

Table 1. Cisco UCS C890 M5 rack server configuration overview

		Quantity	DDR4 capacity	PMem capacity	Usable capacity
CPU specifications	Intel Xeon Platinum Processor 8280L	8			
Memory configurations	64-GB DDR4	96	6 TB	–	6 TB
	128-GB DDR4	96	12 TB	–	12 TB
	256-GB DDR4	96	24 TB	–	24 TB*
	64-GB DDR4/128-GB PMem	48 + 48	3 TB	6 TB	9 TB
	128-GB DDR4/128-GB PMem		6 TB	6 TB	12 TB
	64-GB DDR4/256-GB PMem		3 TB	12 TB	15 TB*
	128-GB DDR4/256-GB PMem		6 TB	12 TB	18 TB*
	256-GB DDR4/256-GB PMem		12 TB	12 TB	24 TB*
	128-GB DDR4/512-GB PMem		6 TB	24 TB	30 TB*
	256-GB DDR4/512-GB PMem		12 TB	24 TB	36 TB*
Internal storage for SAP HANA volumes	3.2-TB SAS-3 disks (RAID 5)	10	6 TB	–	28.8 TB
	6.4-TB SAS-3 disks (RAID 5)	10	12 TB	–	57.6 TB

* In bare-metal installations, SAP HANA supports a maximum of 12 TB (24 TB with Intel PMem). In a virtualized environment, SAP HANA supports a maximum of 6 TB per single virtual machine.

Technology overview

This section provides a technical overview of the computing and storage components and the solution requirements important for SAP HANA operation for this environment.

Solution architecture

The high-performance 8-socket Cisco C890 M5 rack server is based on the second-generation Intel Xeon scalable processors. It is well-suited for mission-critical business scenarios such as for medium- to large-size SAP HANA database management systems including high availability and secure multi-tenancy.

The baseline server configuration uses eight 25-GE uplink network ports and four 32-GB Fibre Channel (FC) Host bus adapters (HBAs). One 12-GBps MegaRAID SAS controller manages two internal, RAID-1-protected SAS drives to host the RHELfor SAP Solutions operating system. A second 12-GBps MegaRAID SAS controller manages at a minimum 10 (up to 16) internal, RAID-5 protected and hot-swappable SAS disks to host the persistent data of the SAP HANA database.

Solution requirements

Several hardware and software requirements must be fulfilled before you set up and install an SAP HANA server. SAP increases flexibility and allows alternative configurations to the SAP HANA appliances with the six phases of SAP HANA TDI, which are described in SAP HANA Tailored Data Center Integration – Overview.

CPU architecture

The 8-socket Cisco UCS C890 M5 rack server is available with a second-generation Intel Xeon Scalable Platinum processor (either the 8276 or the 8280L processor) with 28 cores per CPU socket each. Both CPU models are certified for SAP HANA but are specified for a different amount of addressable memory size.

The smaller platinum 8276 processor addresses a maximum memory size of 1 TB and is available for the 6-TB main memory configuration only. The recommended platinum 8280L processor addresses a maximum memory size of 4.5 TB and is available for all memory configurations listed in Table 1.

Memory platform support

SAP HANA supports the following DDR4 DRAM memory configurations:

- The SAP HANA 2.0 memory per-socket ratio is up to 768 GB for SAP NetWeaver Business Warehouse (BW) and DataMart.

- The SAP HANA 2.0 memory per-socket ratio is up to 1536 GB for SAP Business Suite on SAP HANA (SoH).

With SAP HANA 2.0 SPS 03 or later, Intel Optane PMem modules are supported. In App Direct mode, the PMem modules appear as byte-addressable memory resources that are controlled by SAP HANA.

The use of Intel Optane PMem requires a homogeneous, symmetrical assembly of DRAM and PMem modules with maximum use of all memory channels per processor. Presently the total usable memory capacity for SAP HANA is up to 24 TB.

Various capacity ratios between PMem and DRAM memory modules are supported. The ratios are heavily dependent on the data model and data distribution. Cisco highly recommends that you perform a correct SAP sizing exercise before you consider Intel Optane PMem for SAP HANA.

Ratio examples include the following:

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

- 1:2 ratio: A single 256-GB Intel Optane PMem is matched with a single 128-GB DDR4 DIMM.

- 1:4 ratio: A singe 512-GB Intel Optane PMem is matched with a single 128-GB DDR4 DIMM.

You can use different-sized Intel Optane PMem and DDR4 DIMMs together as long as you maintain supported ratios. The FAQ: SAP HANA Persistent Memory provides more information if you are interested in using PMem in SAP HANA environments.

Operating system

The Cisco C890 M5 rack server is certified for SAP HANA starting with the Red Hat Enterprise Linux (RHEL) for SAP Solutions 8.1 operating system release. To evaluate compatibility information between Linux operating system release and SAP HANA platform releases, refer to SAP Note 2235581 – SAP HANA: Supported Operating Systems.

Disk size and file-system requirements

The RHEL standard installation uses Logical Volume Management (LVM) as the device type and the highly scalable, high-performance XFS file system as the default and recommended file system on RHEL. Cisco recommends joining the root directory "/" and the SAP instance directory "/usr/sap", which contains the

SAP HANA executable programs, instance profiles, and SAP HANA configuration files. The minimum space requirement when joining both directories is 62 GB.

You should account for the recommendations for the [SAP HANA Storage Requirements](#) when you configure the SAP HANA data, log, and shared directories. If no application-specific sizing program data is available, the recommended size for the SAP HANA volumes is directly related to the total memory required for the system.

For example, for an SAP HANA scale-up appliance with 12-TB main memory, the recommended SAP HANA file system sizes are as follows:

- /hana/data = 3 x memory = 3 x 12 TB = 36 TB

- /hana/log = 512 GB

- /hana/shared = 1 x memory = 12 TB

Server hardware configuration

It is beyond the scope of this document to explain the Cisco UCS infrastructure installation and connectivity. All physical hardware must be racked according to the specific hardware installation guides. This deployment guide assumes the cabling is complete and the server is powered off prior to starting the initial configuration.

The server contains a security feature for the BMC firmware stack and is shipped with an individual, unique, and preprogrammed default password for the ADMIN user. Cisco recommends that you change the password after the first logon to the IPMI Console (IPMI-Over-LAN) management web interface.

Server deployment

This section provides the steps for the deployment of the SAP HANA scale-up environment on a Cisco UCS C890 M5 rack server using RHEL for SAP Solutions 8.1.

Naming convention

This guide uses the following conventions for commands that you enter at the command-line interface (CLI). Commands to enter at a CLI prompt follow:

`# hostname`

Angle brackets (<>) indicate a mandatory character string that you need to enter like you would enter a variable pertinent to the customer environment or a password.

`# ip add <ip_address/mask> dev <interface>`

For example, to assign 192.168.1.200/255.255.255.0 to network interface eno1:

`# ip add 192.168.1.200/255.255.255.0 dev eno1`

BIOS configuration

The following configuration procedure is recommended for the Cisco UCS C890 M5 server for an SAP HANA appliance bare-metal installation.

Connect to the KVM console

1. Connect the keyboard, video, mouse (KVM) cable available in the accessory box to a laptop and the KVM connector on the low right rear side of the server for the initial basic input/output (BIOS) and IPMI configuration.
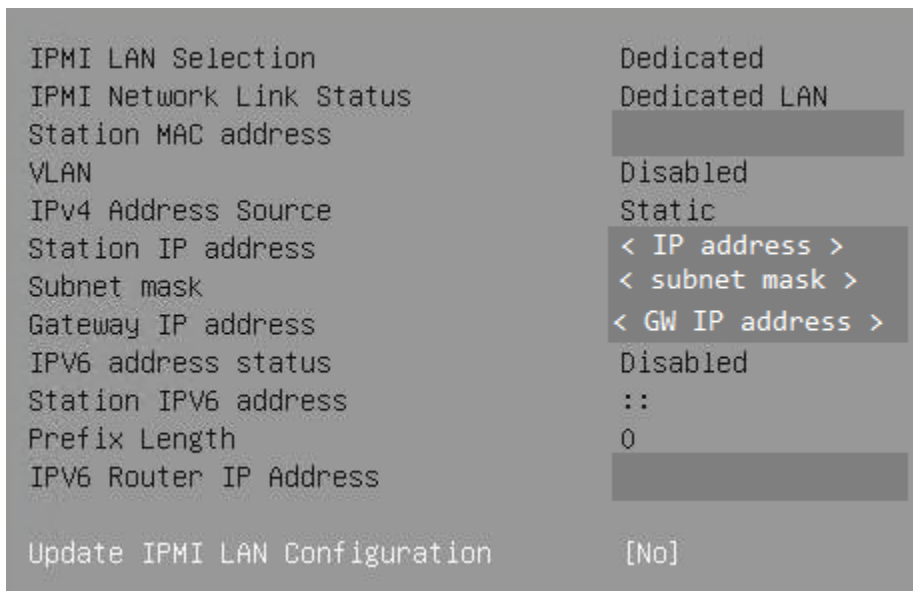
2. When physically connected to the server press the Power button located in the top right corner on the front side of the server.

3. When the Cisco logo appears, press <DEL> to enter the AMI BIOS setup utility in the KVM console.

Main menu

Set the system date and system time in the main screen of the BIOS.

IPMI menu (BMC network configuration)

4. Change to the IPMI top menu and use the arrow keys to select the BMC Network Configuration sub-menu.

5. Update the IPMI LAN Configuration and provide an IP address, subnet mask, and gateway IP address.

```
IPMI LAN Selection              Dedicated
IPMI Network Link Status        Dedicated LAN
Station MAC address
VLAN                            Disabled
IPv4 Address Source             Static
Station IP address              < IP address >
Subnet mask                     < subnet mask >
Gateway IP address              < GW IP address >
IPV6 address status             Disabled
Station IPV6 address            ::
Prefix Length                   0
IPV6 Router IP Address

Update IPMI LAN Configuration   [No]
```

6. Move to the Save & Exit top menu, Save the changes; then reset the server.

```
Save Changes and Reset
```

7. When the Cisco logo appears, press <DEL> to enter the AMI BIOS setup utility.

IPMI Console management

Disconnect the KVM cable and connect to the IPMI LAN IP address with your preferred browser. Log in with the user ADMIN and the unique password that came with the server.

1. To modify the initial password, select the menu Configuration > Users.

2. Select the row with the ADMIN user and press the Modify User button to change the initial password.

Modify User

Enter the new information for the user below and press Modify. Press Cancel to return to the user list.

User Name: ADMIN
Change Password: ☐
Password:
Confirm Password:
Network Privileges: Administrator ▾

[Save] [Cancel]

Password requirements:
- Require password length: 8 to 20 characters
- Password can not be reverse of the user name
- Password must include characters from at least 3 of the listed character classes
- Allowed character classes
  - a - z
  - A - Z
  - 0 - 9
  - Special characters

3. Select the menu Remote Control > Remote Console.

4. Press the Launch Console button to access the remote console.

5. The remote console opens in a new window displaying the AMI BIOS setup utility.

Advanced menu (CPU configuration)

1. Use the arrow keys to select the CPU Configuration sub-menu.

2. For the bare-metal installation, disable Intel Virtualization Technology.

```
Intel Virtualization Technology      [Disable]
```

3. Select Advanced Power Management Configuration at the bottom of the processor configuration page.

4. Change Power Technology to Custom.

```
Advanced Power Management Configuration
---------------------------------------------------------
Power Technology                  [Custom]
Power Performance Tuning          [OS Controls EPB]
Energy Performance BIAS Setting   [Balanced Performance]
Super Performance Mode            [Disable]
```

5. Select CPU C State Control and disable the Autonomous Core C-State and CPU C6 report.

```
CPU C State Control

Autonomous Core C-State           [Disable]
CPU C6 report                     [Disable]
Enhanced Halt State (C1E)         [Enable]
```

6. Select Package C State Control and change the Package C State value to C0/C1 state.

```
Package C State Control

Package C State                        [C0/C1 state]
```

Advanced menu (memory RAS configuration)

Modern servers, including Cisco UCS M5 servers, provide increased memory capacities that run at higher bandwidths and lower voltages. These trends, along with higher application memory demands and advanced multicore processors, contribute to an increasing probability of experiencing memory errors.

Fundamental Error-Correcting Code (ECC) capabilities and scrub protocols were historically successful at handling and mitigating memory errors. As memory and processor technologies advance, reliability, availability, and serviceability (RAS) features must evolve to address new challenges. You can improve the server resilience with the Adaptive Double Device Data Correction (ADDDC Sparing) feature, which mitigates correctable memory errors.

7. Use the arrow keys to select the Chipset Configuration sub-menu.

8. Continue with North Bridge and Memory Configuration.

9. Keep all values and select Memory RAS Configuration at the bottom of the screen.

10. Enable ADDDC Sparing.

```
ADDDC Sparing                          [Enable]
```

Advanced menu (PCIe/PCI/PnP configuration)

In the default configuration, the Option Rom Settings are set to Legacy. You can access configuration utilities of the add-on cards by pressing a key combination such as <CTRL>-<R> for the MegaRAID cards during initialization of the card itself. Changing the BIOS settings to EFI will enable access to the configuration utilities directly from the BIOS after the next reboot.

11. Use the arrow keys to select the PCIe/PCI/PnP Configuration sub-menu.

12. Change the PCI Devices Option Rom from Legacy to EFI.

```
PCI Devices Option Rom Setting

Onboard Video OPROM                    [Legacy]
SIOM Module OPROM                      [EFI]
NVMe Module OPROM                      [Legacy]
GPU Module OPROM                       [Legacy]
HDD Module OPROM                       [EFI]

PCIE Module SL01 x16 OPROM             [EFI]
PCIE Module SL02 x16 OPROM             [EFI]
PCIE Module SL03 x16 OPROM             [EFI]
PCIE Module SL04 x16 OPROM             [EFI]
PCIE Module SL05 x16 OPROM             [EFI]
```

13. Press <ESC> and select optionally the network stack configuration to disable the PXE and HTTP boot support if not required.

```
Network Stack                          [Disabled]
```

Boot menu

14. Switch to the Boot screen in the top menu. Confirm that the entry for boot mode select is UEFI.

```
Boot mode select                          [UEFI]
```
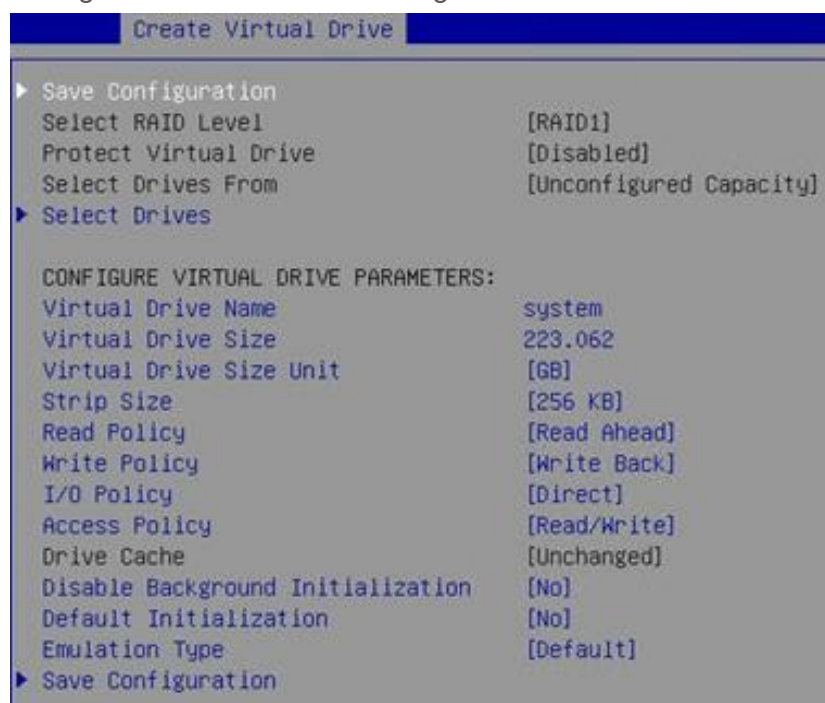
Save and exit menu

15. Save all changes and reset the server to activate the changes.

```
Save Changes and Reset
```

After the system initialization, the Cisco logo displays. Press <DEL> to enter the AMI BIOS setup utility again.

MegaRAID 3108 configuration

1. In the Advanced menu of the Setup Utility, select the AVAGO 3108 MegaRAID Configuration Utility.

2. Select Configure from the Actions section and create a virtual drive.

3. Select RAID level 1 and the two unconfigured SATA drives to be added.

4. Apply the changes and provide a virtual drive name like, for example, system. Keep all other settings on default. Save the configuration.

```
        Create Virtual Drive

▶ Save Configuration
  Select RAID Level             [RAID1]
  Protect Virtual Drive         [Disabled]
  Select Drives From            [Unconfigured Capacity]
▶ Select Drives

  CONFIGURE VIRTUAL DRIVE PARAMETERS:
  Virtual Drive Name            system
  Virtual Drive Size            223.062
  Virtual Drive Size Unit       [GB]
  Strip Size                    [256 KB]
  Read Policy                   [Read Ahead]
  Write Policy                  [Write Back]
  I/O Policy                    [Direct]
  Access Policy                 [Read/Write]
  Drive Cache                   [Unchanged]
  Disable Background Initialization  [No]
  Default Initialization        [No]
  Emulation Type                [Default]
▶ Save Configuration
```

MegaRAID 9361 configuration

1. From the Advanced screen, select the AVAGO 9361-16i MegaRAID Configuration Utility.

2. Select Configure from the Actions section and create a virtual drive.

3. Select RAID level 5 and all unconfigured SAS drives to be added.

4. Apply the changes and provide a virtual drive name like, for example, hana. Change the strip size to 128 KB and the Write Policy to Write Back. Keep all other settings on default and save the configuration.
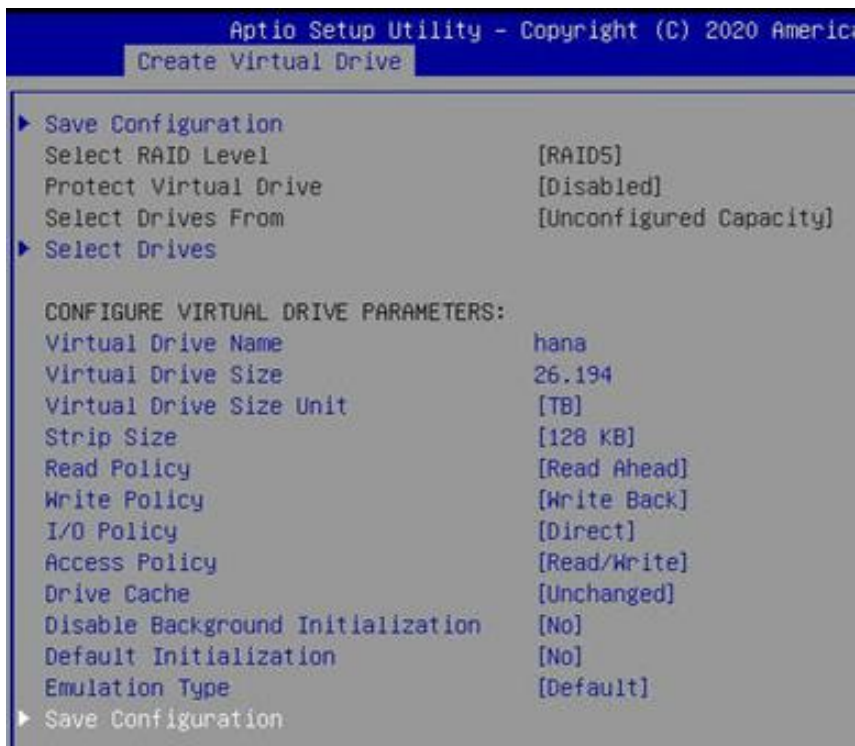
Intel Optane DC PMem configuration

If the Cisco UCS C890 M5 server is not equipped with Intel PMem modules, the BIOS menu entry is hidden.

1. From the Advanced screen select Intel Optane DC Persistent Memory Configuration.

2. Confirm all 48 DIMMs have been detected and are healthy.

For SAP HANA bare-metal installations, no further BIOS configuration is required. The Intel Optane PMem modules are managed using the management utilities from the Linux command line. Return to the main screen and keep the server in the AMI BIOS setup utility.

SAP HANA with Intel Optane PMem on VMware vSphere

For SAP HANA VMware-based installations, the Intel Optane PMem modules require initial configuration in the BIOS to enable ESXi to detect the PMem namespaces properly. Further configuration in VMware vCenter ([VMware KB article 78094](#)) and configuration changes using the management utilities from the Linux command line within the virtual machine are required. Additional information is available from [SAP note 2913410](#) – SAP HANA about VMware vSphere with PMem.

For the initial configuration of the PMem modules for VMware, follow these steps:

1. In the Advanced screen, select Intel Optane DC Persistent Memory Configuration.

```
Version: 1.0.0.3506
Select an action below.

Detected DIMMs:                        48
All DIMMs are healthy.

▶ DIMMs
▶ Regions
▶ Namespaces
▶ Total capacity
▶ Diagnostics
▶ Preferences
```

2. Select Regions and scroll to the end of the screen. Select Create goal config and press <Enter>.

```
Select the scope of the new region then set the desired
sizes.

Create goal config for:           [Platform]

Reserved [%]:                      0
Memory Mode [%]:                   0

Persistent memory type:            [App Direct]

Namespace Label version:           [1.2]

▶ Create goal config
▶ Back to Regions menu
▶ Back to main menu
```

3. Keep the defaults to create a goal for App Direct mode. Select Create goal config and press <Enter>.

4. Save the changes and reset the server to initialize the Intel Optane PMem modules.

5. When the Cisco logo appears, press <DEL> to enter the AMI BIOS setup utility.

6. Move to the Advanced screen and select Intel Optane DC Persistent Memory Configuration.

7. Select Namespaces and then Create namespace.

```
Select a namespace to view more information.

There are no namespaces defined in system.

▶ Create namespace
▶ Back to main menu
```

8. Select Name and press <Enter> to provide a name for the new namespace and select the region ID accordingly; for example, namespace1.0 for Region ID 0x0001.

```
Create a namespace.
Name
Region ID                         [0x0001]
Mode                              [None]
Capacity input                    [Remaining]
Units                             [GiB]
Capacity                          3012.0

▶ Create namespace
▶ Back to Namespaces
▶ Back to main menu
```

9. Keep the mode on None, thereby creating a namespace in raw mode.

10. Select Create namespace.

```
Select a namespace to view more information.

NamespaceID   Name            Health Status
▶ 0x00000101   namespace1.0    Healthy

▶ Create namespace
▶ Back to main menu
```

11. Repeat steps 7 through 9 and create a namespace for each of the eight region IDs.

```
Select a namespace to view more information.

NamespaceID   Name            Health Status
▶ 0x00000101   namespace1.0    Healthy
▶ 0x00000201   namespace2.0    Healthy
▶ 0x00000301   namespace3.0    Healthy
▶ 0x00000401   namespace4.0    Healthy
▶ 0x00000501   namespace5.0    Healthy
▶ 0x00000601   namespace6.0    Healthy
▶ 0x00000701   namespace7.0    Healthy
▶ 0x00000801   namespace8.0    Healthy

▶ Create namespace
▶ Back to main menu
```

12. Select Save from the Save & Exit top menu and press <Enter> to save the changes.

13. Mount the ESXi ISO installation image to proceed with the VMware installation after the next reboot.

Red Hat Enterprise Linux for SAP Solutions installation

Red Hat Enterprise Linux 8 introduces the concept of Application Streams. Multiple versions of user space components are now delivered and updated more frequently than the core operating system packages, providing greater flexibility to customize Red Hat Enterprise Linux without affecting the underlying stability of the platform or specific deployments. Consult the RHEL 8 installation and configuration guide for instructions for downloading the appropriate RHEL for SAP Solutions installation image and follow the installation workflow. During the installation workflow, ensure that you apply the best practices listed in SAP note 2772999 – Red Hat Enterprise Linux 8.x: Installation and Configuration.

The following supplement RHEL information is available from the SAP notes system:

- SAP Note 2369910: SAP Software on Linux: General Information

- SAP Note 2772999: Red Hat Enterprise Linux 8.x: Installation and Configuration

- SAP Note 2777782: SAP HANA DB: Recommended OS settings for RHEL 8

- SAP Note 2886607: Linux: Running SAP applications compiled with GCC 9.x

- SAP Note 2382421: Optimizing the Network Configuration on HANA- and OS-Level

At the release time of this document, SAP HANA 2.0 is supported up to RHEL 8.2 and the server is SAP HANA certified for RHEL 8.1 and 8.2. Download the RHEL 8 release from the Red Hat Customer Portal.

1. In the remote console window, select Virtual Media in the top menu.

2. Keep the device type ISO image and choose the RHEL for SAP Solutions 8.2 ISO installation image.

3. Click the mount button to mount the image and close the window.

4. In the BIOS Setup Utility go to the Save & Exit top menu, select Save Changes, and Reset to restart the server.

The Cisco UCS C890 M5 server will automatically boot from the mapped ISO installation source, and the RHEL installation wizard will start. Install and configure RHEL 8 for SAP Solutions according to the Red Hat Configuration Guide.

Follow the installation workflow.

- Use English as the installation and system default language and a keyboard layout according to your local preferences.

- Manually partitioning of the system disk is strongly recommended, but it is possible to join the /usr/sap and the root / directory as well.

- In the network and hostname configuration screen, enter the short name of the server <hostname> and not the fully qualified hostname.

- Select the correct time zone and ensure that the date and time are set correctly. If applicable, configure a local Network Time Protocol (NTP) server.

Initial setup

After the first login, modify the system to fit into the existing infrastructure and to follow the SAP best practices for RHEL 8 for SAP Solutions. Register the host and apply the RHEL release lock to stay with an SAP HANA-supported RHEL release when patching:

```
# subscription-manager release --set=8.<1|2>
```

Disable all repositories:

```
# subscription-manager repos --disable="*"
```

Enable the following repositories if you intend to use the server for the SAP HANA database only:

```
# subscription-manager repos \
--enable="rhel-8-for-$(uname -m)-baseos-e4s-rpms" \
--enable="rhel-8-for-$(uname -m)-appstream-e4s-rpms" \
--enable="rhel-8-for-$(uname -m)-sap-solutions-e4s-rpms" \
--enable="ansible-2.9-for-rhel-8-$(uname -m)-rpms"
```

The following additional software packages are required for SAP HANA environments:

```
# yum -y install ansible bind-utils compat-sap-c++9 expect krb5-workstation \
  krb5-libs libaio libatomic libcanberra-gtk2 libnsl libibverbs libicu libtool-ltdl \
  nfs-utils numactl openssl PackageKit-gtk3-module psmisc rsyslog sudo tcsh tuned \
  tuned-profiles-sap-hana uuidd xfsprogs xorg-x11-xauth
```

For supportability of SAP HANA purposes, consider installing the following packages as well:

```
# yum -y install cairo graphviz iptraf-ng lm_sensors net-tools
```

Install the following group and the SAP System Role:

```
# yum group install Server
# yum -y install rhel-system-roles-sap
```

Apply the most current security patches:

It is highly recommended that you fully patch and reboot the Cisco UCS server before proceeding with the SAP HANA preconfiguration steps:

```
# yum -y update
```

Server configuration with RHEL System Roles for SAP

With the Ansible Engine having been installed, use the Ansible playbook to configure the host.

1. Create the file sap-hana.yml (end the stream with <CTRL-D>):

```
# cat >> sap-hana.yml
  ---
  - hosts: localhost
    vars:
      ansible_connection: local
      sap_preconfigure_fail_if_reboot_required: no
      sap_hana_preconfigure_fail_if_reboot_required: no
      sap_hana_preconfigure_update: yes
    roles:
```

```
    - role: sap-preconfigure

    - role: sap-hana-preconfigure
```

2. Configure the server according to the applicable SAP notes for SAP HANA, and run the ansible playbook:

```
# ansible-playbook sap-hana.yml
```

3. Reboot the server.

Configure systemd tmpfiles behavior

Ensure that important lock files and sockets in /tmp will not be deleted by systemd-tmpfiles (end the stream with <CTRL-D>):

```
# cat >> /etc/tmpfiles.d/sap.conf

# systemd.tmpfiles exclude file for SAP

# SAP software stores some important files in /tmp which should not be deleted \
  automatically

# Exclude SAP socket and lock files

x /tmp/.sap*

# Exclude HANA lock file

x /tmp/.hdb*lock

# Exclude TREX lock file

x /tmp/.trex*lock
```

SAP HANA partitions

The SAP HANA disk size requirements in the following configuration example can vary based on the available main memory or backup/snapshot configuration, for example.

From the command line, list the available block devices to check for empty devices:

```
# lsblk
NAME           MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda              8:0     0  26.2T  0 disk
sdb              8:16    0 223.1G  0 disk
├─sdb1           8:17    0   512M  0 part /boot/efi
├─sdb2           8:18    0 172.6G  0 part /
├─sdb3           8:19    0    50G  0 part /home
└─sdb4           8:20    0     2G  0 part [SWAP]
```

In this example, the operating system uses /dev/sdb, and device /dev/sda is available. Create a new logical volume management (LVM) volume group and add three logical volumes from the available space in the group following the SAP HANA sizing requirements.

Initialize the physical RAID volume:

```
# pvcreate /dev/sda
```

Create a new logical volume group:

```
# vgcreate hanavg /dev/sda
```

Create three logical volumes on the volume group. For example:

```
# lvcreate -L 18T -n datavol hanavg
# lvcreate -L 512G -n logvol hanavg
# lvcreate -L 6T -n sharedvol hanavg
```

Create a XFS file system for the three new logical volumes:

```
# mkfs.xfs -f -b size=4096 /dev/hanavg/logvol
# mkfs.xfs -f -b size=4096 /dev/hanavg/datavol
# mkfs.xfs -f -b size=4096 /dev/hanavg/sharedvol
```

Create the mount points:

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

From the command line, list the available block devices to confirm the configuration:

```
# lsblk -f
NAME                   FSTYPE      LABEL UUID         MOUNTPOINT
sda                    LVM2_member       JI15YZ…
├─hanavg-sharedvol xfs                   39efac…
├─hanavg-datavol   xfs                   676e1b…
└─hanavg-logvol    xfs                   0531ce…
sdb
├─sdb1                 vfat                4A24-A…     /boot/efi
├─sdb2                 xfs                 5191a0…     /
├─sdb3                 xfs                 743c6a…     /home
└─sdb4                 swap                1fa4e8…     [SWAP]
```

The Red Hat default is to use the Universally Unique Identifiers (UUIDs) instead of device names. Create persistent mount points by adding the volumes to /etc/fstab (end the stream with <CTRL-D>).

```
# cat >> /etc/fstab
UUID=0531ce…                        /hana/log       xfs   defaults   0  0
UUID=676e1b…                        /hana/data      xfs   defaults   0  0
UUID=39efac…                        /hana/shared    xfs   defaults   0  0
```

Finally, mount the logical volumes:

```
# mount -av
```

Disable OS-based memory error monitoring

Linux supports two features related to error monitoring and logging, Error Detection and Correction (EDAC) and mcelog. Both are common in most recent Linux distributions. Cisco recommends disabling EDAC-based error collection to allow all error reporting to be handled in firmware.

Disable EDAC by adding the kernel option *edac_report=off* to the kernel command line.

Mcelog is enabled by default in most recent Linux distributions such as Red Hat Enterprise Linux 8. If you prefer to collect all diagnostic and fault information from OS resident tools, mcelog is recommended. Firmware logs may be incomplete when OS logging is enabled.

Storage Command Line tool installation

The solution has been tested and certified with the RHEL inbox driver. Use the Storage Command Line Tool (STorCLI) to monitor and manage the two virtual RAID groups. Select Management Software and Tools in the download area of https://www.broadcom.com/products/storage/raid-controllers/megaraid-sas-9361-16i and download the current StorCLI software package. Use the RPM package manager to install the MegaRAID StorCLI tool. Show a summary of controller and controller-associated information using the following command:

```
# /opt/MegaRAID/storcli/storcli64 show
```

Fix the BIOS boot order priority

Reboot the server from the CLI. Connect to the IPMI address of the Cisco UCS C890 M5 server with your preferred browser. Use the ADMIN user to log on to IPMI.

1. From the top menu, select Remote Control and Remote Console from the drop-down menu. The Launch Console button opens a new window with the console redirection.

2. After the system initialization, the Cisco logo displays. Press <DEL> to enter the AMI BIOS setup utility.

3. In the BIOS setup utility, change to the boot top menu and select UEFI Hard Disk Drive BBS Priorities.

4. Open Boot Option #1 and select Red Hat Enterprise Linux from the drop-down menu.

5. If the Boot Option UEFI Hard Disk: Red Hat Enterprise Linux is not listed at the first position of the boot order, move it to the first position.

6. Change to the Save & Exit top menu and quit the setup utility with Save Changes and Reset.

Persistent Memory configuration

Configure and manage Intel Optane PMem modules from the Linux command line. First install the management utilities ipmctl and ndctl required to manage the libnvdimm (non-volatile memory device) sub-system in the Linux kernel. For additional information, refer to the [IPMCTL User Guide](#) and the [NDCTL User Guide](#).

Open a Secure Shell (SSH) Protocol prompt as root to install the host tools and follow these steps:

1. EPEL packages assume that the "codeready-builder" repository is enabled

```
# subscription-manager repos --enable "codeready-builder-for-rhel-8-$(arch)-rpms"
```

2. Enable the EPEL 8 repository or download the required rpm file from
   https://dl.fedoraproject.org/pub/epel/8/Everything/x86_64/Packages/ manually.

```
# yum install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm

# yum info ipmctl

# yum -y install ipmctl
```

3. Install the ndctl utility:

```
# yum -y install ndctl
```

4. Verify that the PMem modules have been discovered and the software can communicate with them.

```
# ipmctl show -dimm

DimmID  | Capacity    | LockState | HealthState | FWVersion
================================================================
 0x0001 | 502.599 GiB | Disabled  | Healthy     | 01.02.00.5435
 0x0011 | 502.599 GiB | Disabled  | Healthy     | 01.02.00.5435
 0x0021 | 502.599 GiB | Disabled  | Healthy     | 01.02.00.5435

 .

 .

 .

 0x7121 | 502.599 GiB | Disabled  | Healthy     | 01.02.00.5435
```

5. Add a UDEV rule to ensure persistent device naming. Update the initrd as well:

```
# vi /etc/udev/rules.d/60-persistent-storage.rules

# PMEM devices

KERNEL=="pmem*", ENV{DEVTYPE}=="disk", ATTRS{uuid}=="?*", SYMLINK+="disk/by-id/pmem-
$attr{uuid}"

KERNEL=="pmem*", ENV{DEVTYPE}=="disk", ATTRS{uuid}=="?*", SYMLINK+="disk/by-
uuid/$attr{uuid}"

# dracut -f
```

6. The ipmctl command default is to provision the Intel PMem in App Direct mode with interleaving
   enabled which increases the throughput of write and reads to the persistent memory area. Create
   the goal:

```
# ipmctl create -goal

The following configuration will be applied:

 SocketID | DimmID | MemorySize | AppDirect1Size | AppDirect2Size
================================================================
 0x0000   | 0x0001 | 0.000 GiB  | 502.000GiB     | 0.000 GiB
 0x0000   | 0x0011 | 0.000 GiB  | 502.000GiB     | 0.000 GiB
 0x0000   | 0x0021 | 0.000 GiB  | 502.000GiB     | 0.000 GiB

 .

 .

 .

Do you want to continue? [y/n]
```

7. Confirm with Y and reboot the server to apply the new memory allocation goals.

8. Verify that regions have been created:

```
# ipmctl show -region
SocketID | ISetID    | Persistent | Capacity     | FreeCapacity | HealthState
         |           | MemoryType |              |              |
================================================================================
 0x0000   | 0x03927… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
 0x0001   | 0x7e2a7… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
 0x0002   | 0xf78a7… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
 0x0003   | 0xe5607… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
 0x0004   | 0xf3967… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
 0x0005   | 0xc5f27… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
 0x0006   | 0x7e007… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
 0x0007   | 0x558c7… | AppDirect  | 3012.000 GiB | 0.000 GiB    | Healthy
```

9. Create a namespace for each of the eight regions.

```
# ndctl create-namespace --continue
```

10. The default syntax creates a namespace with direct access (DAX) support, mode fsdax, required by SAP HANA to use Intel Optane PMem properly. Verify that the namespace has been created successfully:

```
# ndctl list
[
  {
    "dev":"namespace1.0",
    "mode":"fsdax",
    "map":"dev",
    "size":3183575302144,
    "uuid":"de793a2f-5763-4d82-b2c7-59b73721a3e9",
    "sector_size":512,
    "align":2097152,
    "blockdev":"pmem1"
  },
  .
  .
  .
]
```

11. The assigned device name for the namespaces is usually /dev/pmemX, where X denotes the NUMA node mapping of the underlying hardware. Construct an XFS file system on the block devices:

```
# for i in {0..7}; do mkfs.xfs -f -d su=2m,sw=1 -m reflink=0 /dev/pmem$i; done
```

12. Create directories and mount the block devices using the DAX file system option:

```
# for i in {0..7}; do mkdir -p /hana/pmem/nvmem$i; done
# for i in {0..7}; do mount -t xfs -o dax,lazytime /dev/pmem$i \
  /hana/pmem/nvmem$i; done
```

13. Change the permission of the mount points:

```
# chmod 755 /hana/pmem/nvmem*
```

14. Identify the UUID of the persistent memory namespaces:

```
# ndctl list |egrep 'uuid|blockdev'
  "uuid":"de793a2f-5763-4d82-b2c7-59b73721a3e9",  "blockdev":"pmem1"
  "uuid":"2a2b109c-9b3d-463a-8d12-6a51237b4251",  "blockdev":"pmem3"
  "uuid":"d7276772-af1d-496e-a5b2-3f9d2a19f6c4",  "blockdev":"pmem5"
  "uuid":"24057307-8d32-4aa5-b1dc-858006ed4c83",  "blockdev":"pmem7"
  "uuid":"1466813e-b789-452d-bb27-464795f9a32f",  "blockdev":"pmem0"
  "uuid":"98d295c1-8bf9-4b75-a62c-1e42f337f9f0",  "blockdev":"pmem2"
  "uuid":"b0b8f716-e23a-4e02-8d11-b471b348d373",  "blockdev":"pmem4"
  "uuid":"2b0ade1b-8b64-4a55-90d6-1f69ee0695c0",  "blockdev":"pmem6"
```

15. Add the mount points to /etc/fstab to persist them:

```
# vi /etc/fstab
UUID=1466813e-b789-452d-bb27-464795f9a32f /hana/pmem/nvmem0 xfs dax,lazytime 1 2
UUID=de793a2f-5763-4d82-b2c7-59b73721a3e9 /hana/pmem/nvmem1 xfs dax,lazytime 1 2
UUID=98d295c1-8bf9-4b75-a62c-1e42f337f9f0 /hana/pmem/nvmem2 xfs dax,lazytime 1 2
UUID=2a2b109c-9b3d-463a-8d12-6a51237b4251 /hana/pmem/nvmem3 xfs dax,lazytime 1 2
UUID=b0b8f716-e23a-4e02-8d11-b471b348d373 /hana/pmem/nvmem4 xfs dax,lazytime 1 2
UUID=d7276772-af1d-496e-a5b2-3f9d2a19f6c4 /hana/pmem/nvmem5 xfs dax,lazytime 1 2
UUID=2b0ade1b-8b64-4a55-90d6-1f69ee0695c0 /hana/pmem/nvmem6 xfs dax,lazytime 1 2
UUID=24057307-8d32-4aa5-b1dc-858006ed4c83 /hana/pmem/nvmem7 xfs dax,lazytime 1 2
```

The Linux-specific setup of Intel Optane PMem is now complete. After you finish the SAP HANA installation, you must configure SAP HANA in order to use Intel Optane PMem.

SAP HANA installation

All version-specific SAP installation and administration documentation is available from the SAP HANA Help portal: https://help.sap.com/hana. Please refer to the official SAP documentation that describes the different SAP HANA installation options.

Note: Review all relevant SAP notes related to the SAP HANA installation on recent changes.

SAP HANA 2.0 platform installation

The official SAP documentation describes in detail how to install the HANA software and its required components. All required file systems are mounted before the installation.

SAP HANA installation

Download and extract the SAP HANA Platform 2.0 software to an installation sub-folder of your choice. Follow the installation workflow of the SAP HANA Database Lifecycle Manager (hdblcm) and provide the user passwords when asked for:

1. Change to the folder <path>/DATA_UNITS/HDB_LCM_Linux_X86_64.

2. Adapt the following command for the SAP SID, SAP system ID number, hostname, and required components:

```
# ./hdblcm --action install --components=server,client --install_hostagent \
 --number <SAP System ID> --sapmnt=/hana/shared --sid=<SID> --hostname=<hostname>
 --certificates_hostmap=<hostname>=<map name>
```

3. Switch user to <sid>adm.

4. Confirm that SAP HANA is started:

```
# sapcontrol -nr <SAP System ID> -function GetProcessList
```

Configure SAP HANA to use Intel Optane PMem

In a multi-tenant database architecture, create sub-directories in the mounted /hana/pmem/nvmem* directories. For example:

```
# mkdir /hana/pmem/nvmem*/DB1
# mkdir /hana/pmem/nvmem*/DB2
```

Adapt the directory permissions and assign the database user and group. For example:

```
# chmod 755 /hana/pmem/nvmem*/DB*
# chown db1adm:sapsys /hana/pmem/nvmem*/DB1
# chown db2adm:sapsys /hana/pmem/nvmem'/DB2
```

For a single instance, assign the database user and group to the /hana/pmem/nvmem* mount points.

Specify the HANA PMem base path

Enable SAP HANA to use the nonvolatile memory and add the DAX volumes to the SAP HANA configuration file global.ini. Configure the required SAP HANA parameter *basepath_persistent_memory_volumes* in the persistence section. Separate multiple locations by a semicolon.

```
(<sid>adm) # cdglo


(<sid>adm) # vi hdb/custom/config/global.ini
[persistence]
basepath_datavolumes = /hana/data/<SID>
basepath_logvolumes = /hana/log/<SID>
basepath_persistent_memory_volumes =
/hana/pmem/nvmem0;/hana/pmem/nvmem1;/hana/pmem/nvmem2;/hana/pmem/nvmem3;
hana/pmem/nvmem4;/hana/pmem/nvmem5;/hana/pmem/nvmem6;/hana/pmem/nvmem7
```

For a multi-tenant database server, you must configure the parameter separately for both tenants:

```
(db1adm) # cdglo


(db1adm) # vi hdb/custom/config/global.ini
[persistence]
basepath_datavolumes = /hana/data/DB1
basepath_logvolumes = /hana/log/DB1
```

```
basepath_persistent_memory_volumes =
/hana/pmem/nvmem0/DB1;/hana/pmem/nvmem1/DB1;/hana/pmem/nvmem2/DB1;/hana/pmem/nvmem3/DB1;/han
a/pmem/nvmem4/DB1;/hana/pmem/nvmem5/DB1;/hana/pmem/nvmem6/DB1;/hana/pmem/nvmem7/DB1


(db2adm) # cdglo


(db2adm) # vi hdb/custom/config/global.ini
[persistence]
basepath_datavolumes = /hana/data/DB2
basepath_logvolumes = /hana/log/DB2
basepath_persistent_memory_volumes =
/hana/pmem/nvmem0/DB2;/hana/pmem/nvmem1/DB2;/hana/pmem/nvmem2/DB2;/hana/pmem/nvmem3/DB2;/han
a/pmem/nvmem4/DB2;/hana/pmem/nvmem5/DB2;/hana/pmem/nvmem6/DB2;/hana/pmem/nvmem7/DB2
```

Restart the SAP HANA database instance to activate the parameter change. SAP HANA automatically verifies the DAX-enabled file system during startup; all column store tables will use persistent memory by default afterward.

SAP HANA system replication

The SAP HANA system replication (HSR) mechanism helps ensure high availability of the SAP HANA database; you can combine it with a disaster recovery setup. In system replication mode, the content of the SAP HANA database is transferred synchronously to a secondary location in the same or a different data center.

The SAP HANA System Replication - SAP Help Portal provides detailed information about the functions and configuration of SAP HANA system replication. Red Hat Enterprise Linux for SAP Solutions is optimized in various ways for SAP applications. Red Hat provides documentation about how to configure and manage high-availability clusters and how to configure the SAP HANA resource agents to set up an SAP HANA scale-up system with Performance Optimized System Replication:

- Configuring and managing high-availability clusters

- Automating SAP HANA Scale-Up System Replication using the RHEL HA Add-On

One critical element of the cluster implementation is the configuration of a node fencing mechanism that avoids a cluster split-brain situation when the two cluster nodes are not able to communicate with each other any longer. For two cluster nodes with dedicated storage like in this SAP HANA appliance configuration, a shared STONITH block device (SBD) is not available for use as a node fencing mechanism. Instead, configure an IPMI interface as the STONITH cluster resource

IPMI interface STONITH resource

First, you need to create an IPMI user in the management web interface (IPMI Console). Connect to the IPMI address of the Cisco UCS C890 M5 server with your preferred browser. Use the ADMIN user to log on to IPMI.

1. From the top menu, select Configuration and Users from the drop-down menu.

2. Select an empty slot and press the Add User button. Add a new user and provide that user with the privilege level operator. This level allows all baseboard-management-control (BMC) commands except commands that can modify out-of-band interface settings.

Add New User

Enter the new user information below and press Add. Press Cancel to return to the user list.

User Name:	cisco
Password:	•••••••••
Confirm Password:	•••••••••
Network Privileges:	Operator ▾

Password requirements:
- Require password length: 8 to 20 characters
- Password can not be reverse of the user name
- Password must include characters from at least 3 of the listed character classes
- Allowed character classes
  - a - z
  - A - Z
  - 0 - 9
  - Special characters

Save | Cancel

Configure the IPMI interface as a STONITH resource and use the IPMI user credentials to authenticate.

For example:

```
# crm configure edit
...
primitive STONITH-<hostname> stonith:external/ipmi \
        params hostname=<hostname> ipaddr=<IPMI IP address> userid=cisco \
  passwd=<IPMI cisco password> interface=lanplus \
...
```

IPMI watchdog driver

Red Hat Linux Enterprise ships with several modules that provide hardware-specific watchdog drivers and uses the SBD daemon as the software component that "feeds" the watchdog. Configure the correct hardware watchdog if you plan to use SBD instead of the IPMI STONITH resource. The default Intel driver module iTCO_wdt doesn't interact with the baseboard management controller (BMC) and needs to be replaced.

The BMC runs independent of the server processors, so if a processor, memory, or any other hardware failure occurs, BMC still provides its services unless there is a full power outage.

Loading the correct kernel module

List the watchdog modules that are currently loaded in the kernel:

```
# lsmod | egrep "(wd|dog)"
  iTCO_wdt                16384  0
  iTCO_vendor_support     16384  1 iTCO_wdt
```

Unload the kernel module from the system:

```
 # rmmod iTCO_wdt iTCO_vendor_support
```

Blacklist the iTCO watchdog driver to fully disable it (end the stream with <CTRL-D>).

```
# cat >> /etc/modprobe.d/50-blacklist.conf

# unload the iTCO watchdog kernel module
```

```
blacklist iTCO_wdt
blacklist iTCO_vendor_support
```

Next install the [OpenIPMI](#) software package and configure the IPMI watchdog timer.

```
# yum install OpenIPMI ipmitool


# cp -p /etc/sysconfig/ipmi /etc/sysconfig/ipmi.orig


# vi /etc/sysconfig/ipmi
IPMI_SI=yes
DEV_IPMI=yes
IPMI_WATCHDOG=yes
IPMI_WATCHDOG_OPTIONS="timeout=20 action=reset nowayout=0 panic_wdt_timeout=15"
IPMI_POWEROFF=no
IPMI_POWERCYCLE=no
IPMI_IMB=no
```

Note: If a kernel core dump capture is active, increase the panic_wdt_timeout value to allow the core dump to be written to disk.

Configure the systemd watchdog timer and uncomment the following lines in the file system.conf:

```
# vi /etc/systemd/system.conf
…
RuntimeWatchdogSec=10
ShutdownWatchdogSec=60
```

Enable the change and restart the systemd daemon:

```
# systemctl daemon-reexec
```

Start the IPMI driver and enable startup during boot time:

```
# systemctl start ipmi
# systemctl enable ipmi
```

Verify that the IPMI kernel modules are loaded:

```
# lsmod |grep ipmi
ipmi_watchdog          32768  1
ipmi_ssif              40960  0
ipmi_si                69632  1
ipmi_devintf           20480  0
ipmi_msghandler       114688  4 ipmi_devintf,ipmi_si,ipmi_watchdog,ipmi_ssif
```

Verify that the watchdog timer is running and active:

```
# ipmitool mc watchdog get
Watchdog Timer Use:     SMS/OS (0x44)
Watchdog Timer Is:      Started/Running
Watchdog Timer Logging: On
```

```
Watchdog Timer Action:  Hard Reset (0x01)

Pre-timeout interrupt:  None

Pre-timeout interval:   0 seconds

Timer Expiration Flags: (0x10)

                        * SMS/OS

Initial Countdown:      10.0 sec

Present Countdown:      8.4 sec
```

Now a watchdog device is present on the system as well:

```
# ls -l /dev/watchdog

crw------- 1 root root 10, 130 Apr 28 18:09 /dev/watchdog
```

Finally, test the watchdog configuration and trigger a kernel crash. The system should reboot either after the time set by the parameter panic_wdt_timeout runs out or after 5 minutes (BMC timeout).

```
# echo c > /proc/sysrq-trigger
```

Conclusion

The Cisco solution for SAP HANA using the high-performance Cisco UCS C890 M5 Server is well-suited for mission-critical business scenarios such as for medium- to large-size SAP HANA database management systems including high availability and secure multitenancy. The integrated infrastructure solution provides a robust platform for SAP HANA workloads as a single-node (scale-up) SAP HANA appliance.

The 8-socket server delivers exceptional performance and is sized, configured, and deployed to match the SAP HANA appliance key performance metrics that SAP SE demands. With standard configuration options including 16 hot-swappable 2.5-inch SAS-3 drive bays (with RAID cards), eight 25-Gbps Small Form-Factor Pluggable (SFP) Ethernet ports, and four 32-Gbps Fibre Channel ports, this server is well-suited for mission-critical SAP HANA workloads where performance and memory size are critical attributes.

Red Hat Enterprise Linux (RHEL) for SAP Solutions combines the reliability, scalability, and performance of Red Hat Enterprise Linux with additional content specific to operating SAP landscapes. Based on an intelligent operating system, Red Hat delivers a stable, security-focused, high-performance foundation for critical systems. Consistency across bare-metal, virtualized, hybrid, and cloud infrastructure allows you to deploy and manage both traditional and cloud-native workloads using the same tools and interfaces.

The SAP HANA platform provides a scalable database with advanced analytical capabilities and intuitive application-development tools in an in-memory data platform. SAP HANA Tailored Data Center Integration (TDI) supports nonvolatile Intel Optane Persistent Memory that brings together the low latency of memory and the persistence of storage. Persistent memory modules provide faster access to data and retain data across power cycles.

Software- and hardware-based design options to meet your high-availability and disaster recovery requirements are available, and you can combine them with SAP application monitoring using Cisco AppDynamics® technology to detect problems in real time before they affect customers.

For more information

For additional information, refer to the resources listed in this section.

Cisco Unified Computing System

- Cisco UCS C890 M5 Rack Server for SAP HANA
  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/ucs-c890m5-sap-hana-wp.html

- Cisco UCS C890 M5 Rack Server data sheet
  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/datasheet-c78-744647.html

- Cisco UCS for SAP HANA with Intel Optane DC Persistent Memory Module
  https://www.cisco.com/c/dam/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/whitepaper-c11-742627.pdf

- Cisco UCS M5 Memory Technical Overview – Memory RAS Features
  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/technical-overview-c17-743902.html

- Configure Cisco AppDynamics Monitoring Solution for SAP Applications
  https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/appd-sap-monitoring-wp.html

SAP HANA platform

- SAP HANA platform on SAP Help portal:
  https://help.sap.com/viewer/p/SAP_HANA_PLATFORM

- SAP HANA TDI: Overview:
  https://www.sap.com/documents/2017/09/e6519450-d47c-0010-82c7-eda71af511fa.html

- SAP HANA TDI storage requirements:
  https://www.sap.com/documents/2015/03/74cdb554-5a7c-0010-82c7-eda71af511fa.html

- SAP HANA network requirements:
  https://www.sap.com/documents/2016/08/1cd2c2fb-807c-0010-82c7-eda71af511fa.html

- SAP HANA System Replication – SAP Help Portal:
  https://help.sap.com/viewer/4e9b18c116aa42fc84c7dbfd02111aba/2.0.04/en-US/afac7100bc6d47729ae8eae32da5fdec.html

- SAP HANA high availability with SAP HANA system replication and RHEL cluster:
  https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/sap-applications-on-cisco-ucs/whitepaper-c11-735382.pdf

Interoperability matrixes

- Cisco UCS hardware compatibility matrix:
  https://ucshcltool.cloudapps.cisco.com/public/

- SAP Note 2235581: SAP HANA: Supported Operating Systems:
  https://launchpad.support.sap.com/#/notes/2235581

- SAP certified and supported SAP HANA hardware:
  https://www.sap.com/dmc/exp/2014-09-02-hana-hardware/enEN/index.html

- VMware Compatibility Guide:
  [VMware Compatibility Guide - Cisco UCS C890 M5](#)

Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](#) at [https://cs.co/en-cvds](#).

221120.7   09/21