Dell PowerEdge Systems Red Hat Enterprise Linux 8

Release Notes

December 2024
Rev. A26

DELL Technologies

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

1

Release summary

Red Hat Enterprise Linux (RHEL) 8 is the follow-on Operating System (OS) release to the RHEL 7 operating system. Major enhancements in RHEL 8 are improvements in security and stability.

Topics:

Version

8.4

Release date

May 2021

Priority and recommendations

RECOMMENDED: Dell EMC recommends applying this update during your next scheduled update cycle. The update contains feature enhancements or changes that help keep your system software current and compatible with other system modules (firmware, BIOS, drivers, and software).

2

Compatibility

Red Hat Enterprise Linux 8 is available on the 64-bit Intel architecture.

Topics:

System configuration requirements

For detailed system configuration requirements for RHEL 8, see the documentation at Red Hat Support.

Memory

The following table lists the system memory requirements for the x86_64 architecture of RHEL 8.

Table 1. Memory requirements for x86_64 architecture

Memory	Size
Minimum recommended system memory	1.5 GB per logical CPU
Maximum certified system memory	6 TB

3

New and enhanced in RHEL 8 release

- Simple content structure available through BaseOS and Application Stream repository
- More life cycle options and frequent application updates
- YUM4 package manager is based on the DNF technology
- Increased security with memory
- Better management of System Processes
- Support for Control Group v2 mechanism
- Multiple packages of the kernel package
- Distributed with Stratis, a new local storage manager, which provides easy to use CLI and well designed API
- System-wide cryptographic policies and support for OpenSSL 1.1.1 and TLS 1.3 cryptographic standard
- New TCP/IP stack with Bottleneck Bandwidth and RTT (BBR) congestion control

NOTE: For more information on what is new and enhanced in this release, see Red Hat Documentation.

4

Important notes

In the earlier versions of Red Hat Enterprise Linux, the default network interface naming scheme was **biosdevname**. From RHEL 8 onwards, **systemd** is the default naming scheme for network interface naming.

NOTE: You can turn on the **biosdevname** naming scheme during installation or post installation by passing the kernel command line parameter **biosdevname=1**.

5

Fixes

Topics:

- The anaconda installer crashes while autoconfiguring disk partitions
- The system hangs, and a kernel soft lockup error is observed when OMSA is installed
- I/O related errors are observed in dmesg
- RHEL 8.8 and RHEL 9.2 failed to boot
- Unattended operating system installation fails in RHEL using server configuration profile method
- RHEL 8.6 drops into maintenance mode while booting
- RHEL 8.6 fails to boot when UEFI secure boot is enabled BIOS
- Lspci utility lists HBA350i MM/ HBA350i MM LP as generic SAS devices
- RHEL 8.5 boot fails with Intel Corporation Ethernet Controller E810
- NetworkManager may restart unexpectedly when creating greater than 256 VLAN devices configured with DHCP IP
- RHEL 8.3 installer does not automatically locate the source installation repository when only inst.stage2=hd boot option is used
- The output of the systemctl status command displays the status as thawing
- RHEL 8 does not discover FCoE LUNs connected over Broadcom BCM57XXX NICs
- System may drop into emergency mode during boot process
- The version field in the output of the modinfo command for certain networking drivers is null
- System hangs when Intel tboot is used to boot the operating system
- BIOS update does not complete when an update is performed using the Linux .BIN files
- Dmesg shows drm related call trace in RHEL 8.3
- Operating system crashes on servers with NVIDIA GPGPUs
- Dmesg and /var/log/messages display AMD-Vi related messages
- The status of the NetworkManager service may be inactive when RHEL 8.3 is rebooted
- Operating system crashes on AMD Rome CPU-based systems and with Intel E810 NIC
- The lvcreate command requests a response from the user when -wipesignature -yes parameters are passed
- The mdmonitor service displays an error during operating system installation
- The dmidecode utility displays the slot type as <OUT OF SPEC> for PCIe Gen 4 NVMe slots
- The mcelog utility logs only decoding architectural errors message in var/log/messages
- Disk drives part of MD RAID are not listed as installation destination by the installer
- Dell EMC OpenManage Storage Services utility fails to reconfigure the virtual disk
- Guest VMs with SRIOV VFs assigned take a long time to power on, and libvirt related errors are observed
- Dmesg displays Integrity Measurement Architecture (IMA) driver related-messages during system boot
- After every reboot, the network interface name changes
- Red Hat Enterprise Linux Version 8 installation wizard creates a duplicate bonding interface
- Servers with the AMD Rome processor display a CCP initialization failure message in dmesg
- PowerEdge servers with the AMD Rome processor fail to detect an NVMe drive after multiple hot plugs
- Operating system enters the dracut shell during boot
- System crashes when rebooted with SR-IOV-enabled QLogic cards
- After system reboot, Disk data format (DDF) devices are not listed in /proc/mdstat
- Updating NVMe firmware using the nvme-cli utility displays an error in dmesg
- Fatal error BDF 02:00.0 is detected with BCM574xx NICs
- NVMe devices are not detected after hot-plugging
- Linux operating system fails to detect the Intel x710 card
- Dmidecode displays OUT OF SPEC in Slot Type and Slot Length of SMBIOS system slots
- Custom partitioning fails with FC LUN
- When booting the system from iSCSI with Mellanox CX-4 and CX-5 adapters, the system reports csum failure message
- Red Hat Enterprise Linux 8 kernel panic is observed due to fatal hardware error

The anaconda installer crashes while autoconfiguring disk partitions

Description

On a drive with older GPT metadata, the anaconda installer crashes while autoconfiguring disk partitions. The following messages are observed in syslog every time a disk is discovered:

```
WARNING kernel:GPT:Primary header thinks Alt. header is not at the
end of the disk.
WARNING kernel:GPT:209715199 != 282407679
WARNING kernel:GPT:Alternate GPT header not at the end of the disk.
WARNING kernel:GPT:209715199 != 282407679
WARNING kernel:GPT: Use GNU Parted to correct GPT errors.
```

Cause

When a disk is discovered by the operating system, the contents of the disk are analyzed. A GPT partition table is found, and the location of the alternate GPT partition table is checked. Since the alternate GPT partition table is not found at the end of the disk, the above mentioned messages are observed in **syslog**. This scenario can occur when a disk is extended after a GPT partition table is already created on it.

Solution

Issue is fixed in python-blivet-3.2.2-8.el8 which is in RHEL 8.5 or later.

Systems affected

All Dell EMC PowerEdge servers

Applies to

Red Hat Enterprise Linux 8.2 and later

Tracking number

165679

The system hangs, and a kernel soft lockup error is observed when OMSA is installed

Description

On the Dell PowerEdge system configured with PERC 12 and with the mpi3mr out of box driver installed, or when the user installs OM-SrvAdmin-Dell-Web-LX-11.0.1.0-5487.RHEL8.x86_64 (OMSA) the following call trace is observed when the system is in use: `kernel:watchdog: BUG: soft lockup - CPU#30 stuck for 22s! [kworker/30:1H:931].`

Issue is observed when September firmware version 8.4.10.0.18-5 (DPN 2DVNW) and December firmware version 8.4.0.0.18-29 (DPN 93K10) driver combination is used. For more information, see Red Hat Enterprise Linux Knowledge Base article 6985596.

Workaround

Use inbox mpi3mr driver instead of OOB driver or Use December FY24 firmware and December FY24 PERC 12 OOB driver.

Resolution

The issue has been resolved in RHEL 8.7 z stream kernel. (kernel-core-4.18.0-425.10.1.el8_7).

Systems affected

All Dell PowerEdge Servers with a high number of CPU cores.

Applies to

Red Hat Enterprise Linux 8.5 and later.

Tracking number

277161

I/O related errors are observed in dmesg

Description

On Dell EMC PowerEdge systems equipped with PERC or HBA and installed with RHEL 8.8 and started performing I/O on EXT4 file system drives using the I/O tool, the following error messages were displayed by Dmesg after some time of stress.

```
EXT4-fs (sdc): Delayed block allocation failed for inode 157 at
logical offset 1024 with max blocks 4 with error 74
EXT4-fs (sdc): This should not happen!! Data will be lost
```

Workaround	Free up some space in the file system or extend the file system.
Resolution	The error is expected when the file system is full. For more information, see Red Hat Enterprise Linux Knowledge Base article 330523.
Systems affected	Dell PowerEdge system XE9680
Applies to	Red Hat Enterprise Linux 8.6
Tracking number	278631

RHEL 8.8 and RHEL 9.2 failed to boot

Description	On Dell PowerEdge systems configured with the AMD Siena SP6 Processor. RHEL 8.8 and RHEL 9.2 operating system cannot normally boot into the operating system after installation and intermittently entering to emergency mode when system profile settings are set to performance mode in the BIOS.
Workaround	While the system is booting, press 'e' to enter the kernel command line: 1. Add **nomodeset** after rhgb quiet to the kernel command line and press ctrl+x to save and boot into the operating system. 2. Remove **rhgb** from the kernel command line and press ctrl+x to save and boot into the operating system.
Resolution	The issue is resolved in Red Hat Enterprise Linux 8.9 and 9.3.
Systems affected	Dell PowerEdge system C6615
Applies to	Red Hat Enterprise Linux 8.5 and later.
Tracking number	273714

Unattended operating system installation fails in RHEL using server configuration profile method

Description	On Dell PowerEdge system, during the installation of Red Hat Enterprise Linux by importing server configuration profile(SCP), installer is not reading the ks.cfg file, that is stored in a block device with the file system label OEMDRV exposed by life cycle controller to the operating system. As a result, unattended installation stops at language selection wizard.
Workaround	Pass **inst.ks=hd:LABEL=OEMDRV:/ks.cfg** kernel command-line parameter. For more information, see Red Hat Enterprise Linux Knowledge Base article 6999834.
Resolution	Issue is resolved in RedHat Enterprise Linux 8.9.
Systems affected	All Dell PowerEdge systems
Applies to	RedHat Enterprise Linux 8.0 and later.
Tracking number	245112

RHEL 8.6 drops into maintenance mode while booting

Description	On the Dell EMC PowerEdge system with high number of CPU cores and system profile set to performance, the system fails to boot into the operating system and drops into the maintenance mode.
Workaround	Pass **module_blacklist=acpi_cpufreq** kernel command line parameter.
Resolution	Issue is resolved in RedHat Enterprise Linux 8.8.
Systems affected	All Dell EMC PowerEdge Servers with high number of CPU cores.

Applies to	RedHat Enterprise Linux 8.5 and later.
Tracking number	218734

RHEL 8.6 fails to boot when UEFI secure boot is enabled BIOS

Description	On a system which has a operating system with shim version 15.7 or later installed, trying to install RHEL 8.6 fails with error message `Verification failed: (0x1A) Security Violation` shown on console. The behavior is as expected.
Workaround	To install RHEL 8.6 on a system showing the message `Verification failed: (0x1A) Security Violation`: Option 1: See, Red Hat Enterprise Linux Knowledge Base article 7010515. Option 2: ssh to iDRAC and execute `racadm systemerase bios`. **NOTE:** The command may reset all BIOS settings to default.
Systems affected	All Dell PowerEdge systems
Applies to	RedHat Enterprise Linux 8.6.
Tracking number	260940

Lspci utility lists HBA350i MM/ HBA350i MM LP as generic SAS devices

Description	On a system configured with HBA350i MM/HBA350i MM LP Controllers lspci utility lists the HBA350i MM/HBA350i MM LP devices as generic SAS devices Fusion-MPT 12GSAS/PCIe Secure SAS38xx instead of HBA350i MM/HBA350i MM LP.
Cause	The PCI vendor/device IDs of HBA350i MM/HBA350i MM LP controller are not present in PCI ID database in Red Hat Enterprise Linux 8.5.
Workaround	Not Available. It is a cosmetic issue.
Solution	This issue is resolved in Red Hat Enterprise Linux 8.6.
Systems affected	All Dell EMC PowerEdge servers supporting HBA350i MM/HBA350i MM LP controller.
Applies to	Red Hat Enterprise Linux 8.5
Tracking number	222570

RHEL 8.5 boot fails with Intel Corporation Ethernet Controller E810

Description	RHEL 8.5 boot process fails when the system has Intel Corporation Ethernet Controller E810 plugged in. During the boot process the splash screen hangs. The PCI Device ID of E810 NIC is 8086:1593. For more information, see Red Hat Enterprise Linux Knowledge Base article 6745731.
Workaround	Issue can be worked around by removing **rhgb** keyword from grub command line.
Solution	Issue is resolved in RedHat Enterprise Linux 8.6.
Systems affected	All Dell EMC PowerEdge Servers supporting Intel Corporation Ethernet Controller E810.

Applies to	RedHat Enterprise Linux 8.4 and RedHat Enterprise Linux 8.5.
Tracking number	217700

NetworkManager may restart unexpectedly when creating greater than 256 VLAN devices configured with DHCP IP

Description	On RHEL 8.3, when more than 256 VLAN devices are created and configured with DHCP IP, NetworkManager utility may restart unexpectedly resulting in not creating expected number of VLAN devices and some VLAN interfaces may not receive DHCP IP.
Workaround	1. Increase the limit on number of files that NetworkManager can open as below: a. Create file/etc/systemd/system/NetworkManager.service.d/set-limit-nofile.conf b. Add following lines to /etc/systemd/system/NetworkManager.service.d/set-limit-nofile.conf [Service] LimitNOFILE=65536 c. systemctl daemon-reload d. systemctl restart NetworkManager or reboot the system after step (b). 2. Disable IPv6 configuration for the VLAN interfaces at the time of connection activation. Instead of disabling IPv6 in two steps, disable it in a single step as shown below. a. Merge steps b) and c) to a single step d). b. nmcli con add type vlan ifname <interface name> con-name <connection name> dev <device name> id <vlan id> c. nmcli con modify <connection name> ipv6.method disable d. nmcli con add type vlan ifname <interface name> con-name <connection name> dev <device name> id <vlan id> ipv6.method disabled
Solution	This issue is resolved in Red Hat Enterprise Linux 8.4.
Systems affected	All Dell EMC PowerEdge systems
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	191829

RHEL 8.3 installer does not automatically locate the source installation repository when only inst.stage2=hd boot option is used

Description	When both stage2 and the installation repository are on the same internal drive or external drive (USB hard drive), RHEL 8.3 installer does not automatically locate the source installation repository when only inst.stage2=hd boot option is used.
Solution	Modify the boot option **inst.stage2=hd:LABEL=<NAME>** to **inst.repo=hd:LABEL=<NAME>**.
Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8.3
Tracking number	182181

The output of the systemctl status command displays the status as thawing

Description	The output of the **systemctl status <service name>** command displays the status **Active: active (running) (thawing)** instead of **Active: active (running)** when a service is reloaded. This issue is a cosmetic issue and can be ignored.
Solution	The issue is resolved in systemd version 239-41.el8_3.x86_64 and later.
Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8.3
Tracking number	178127

RHEL 8 does not discover FCoE LUNs connected over Broadcom BCM57XXX NICs

Description	On systems with Broadcom BCM57XXX NICs installed, the following cases are observed: 1. When Red Hat Enterprise Linux 8 is installed on FCoE Logical Unit (LUN) as the installation destination, LUN is not discovered. 2. While upgrading Red Hat Enterprise Linux 7 to Red Hat Enterprise Linux 8, and if the current operating system is installed on FCoE LUN, then the operating system fails to boot. 3. When Red Hat Enterprise Linux 7 is upgraded to Red Hat Enterprise Linux 8, and if FCoE LUN is connected to the system as data LUN, then LUN is not discovered. **CAUTION: Do not upgrade to Red Hat Enterprise Linux 8 if FCoE LUNs are attached to Broadcom BCM57XXX NICs. Boot failure or data inaccessibility may result.**
Cause	From Red Hat Enterprise Linux 8 onwards, the software FCoE is deprecated. When software FCoE support is not enabled in the operating system, the bnx2fc driver uses the libfcoe module. Since FCoE works over VLAN 802.1Q, the relevant change needs to be implemented in 802.1Q network module. As these changes were never done, Red Hat Enterprise Linux 8 fails to discover FCoE LUN connected over Broadcom BCM57XXX NICs.
Workaround	There is no workaround for this issue.
Solution	This issue is resolved in Red Hat Enterprise Linux 8.3.
Systems affected	Dell EMC PowerEdge M630, FC630, M830, FC830, R630, R730, R730xd, R930, R640, R740, R940
Applies to	Red Hat Enterprise Linux 8.0
Tracking number	131903

System may drop into emergency mode during boot process

Description	System may drop into emergency mode during boot process if CPU Power Management is set to Maximum Performance in BIOS. For more information, see Red Hat Enterprise Linux Knowledge Base article 6130261.
Workaround	Pass **module_blacklist=acpi_cpufreq** to the boot loader.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.5.
Systems affected	Dell PowerEdge R6515, Dell PowerEdge R7515, Dell PowerEdge R6525, Dell PowerEdge R7525, and Dell PowerEdge C6525.
Applies to	Red Hat Enterprise Linux 8.4

Tracking number	200841

The version field in the output of the modinfo command for certain networking drivers is null

Description	The version field in the output of the **modinfo** command for certain networking drivers is null. As a result, the version field in the output of **ethtool-i** command is set to the version of the kernel. The output of the **modinfo** command is similar to the following:
	`modinfo -F version i40e`
	`<no output>`

Workaround	Use the following command to retrieve the version of the network driver: `modinfo <driver name> | grep rhelversion` For more information, see *Certain kernel drivers do not display their version section* under Know Issues.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.5.
Systems affected	All Dell EMC PowerEdge servers.
Applies to	Red Hat Enterprise Linux 8.4
Tracking number	197095

System hangs when Intel tboot is used to boot the operating system

Description	In legacy BIOS boot mode, system hangs when tboot version 1.9.12-2 is used to boot the operating system. The issue is observed with both TPM v2.0 and v1.2.
Workaround	Use tboot version 1.9.10-1 from RHEL 8.2 GA. For more information, see Red Hat Enterprise Linux Knowledge Base article 6013091.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.5.
Systems affected	All Dell EMC PowerEdge systems.
Applies to	Red Hat Enterprise Linux 8.3 and later.
Tracking number	197339

BIOS update does not complete when an update is performed using the Linux .BIN files

Description	BIOS update does not complete when an update is performed using the Linux .BIN files.
Cause	Interaction with Intel Management Engine Interface (Intel MEI) results in cold reboot instead of warm reboot.
Workaround	Exclude the mei and mei_me drivers.
Solution	This issue is resolved in Red Hat Enterprise Linux 8.4.
Systems affected	Dell EMC PowerEdge systems R240, R340, T140, and T340
Applies to	Red Hat Enterprise Linux 8.3

Tracking number	195178

Dmesg shows drm related call trace in RHEL 8.3

Description	On a Dell EMC PowerEdge MX740C system installed with RHEL 8.3, following call trace is observed in Dmesg: `WARNING: CPU: 102 PID: 1722 at drivers/gpu/drm/ drm_gem_vram_helper.c:576` For more information, see Red Hat Enterprise Linux Knowledge Base article 5899381.
Workaround	Not available. The message can be ignored.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.4.
Systems affected	Dell EMC PowerEdge MX740C, Dell EMC PowerEdge MX750C.
Applies to	Red Hat Enterprise Linux 8.3
Tracking number	194811

Operating system crashes on servers with NVIDIA GPGPUs

Description	Operating system crashes when installing Red Hat Enterprise Linux 8.3 on servers with NVIDIA GPGPUs. For more information, see Red Hat Enterprise Linux Knowledge Base article 5853331.
Workaround	Prevent nouveau driver from loading, by passing `modprobe.blacklist=nouveau` parameter during the installation.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.4.
Systems affected	All Dell EMC PowerEdge systems which support Nvidia GPGPUs
Applies to	Red Hat Enterprise Linux 8.3
Tracking number	189163

Dmesg and /var/log/messages display AMD-Vi related messages

Description	Dmesg and /var/log/messages display the following messages on Dell EMC PowerEdge servers with 256 core CPUs and Dell EMC PowerEdge Express Flash Enterprise NVMe Agnostic (AGN) device or Dell EMC Express Flash Enterprise NVMe CD6 or Dell EMC Express Flash Enterprise NVMe CM6 device: `AMD-Vi: Failed to allocate IRTE` The message indicates that only a total of 255 IO queues are enabled instead of 256.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.4.
Systems affected	Dell EMC PowerEdge C6525, Dell EMC PowerEdge R6525, and Dell EMC PowerEdge R7525.
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	171631

The status of the NetworkManager service may be inactive when RHEL 8.3 is rebooted

Description	The status of the NetworkManager service may be inactive when the operating system is rebooted. The issue is observed when the Performance Co-Pilot (PCP) package is installed on the operating system. The command **systemctl status NetworkManager** can be used to check the status of the service. For more information, see Red Hat Enterprise Linux Knowledge Base article 5394191.
Workaround	Disable the **pmlogger.service** using the following command: `#sudo systemctl disable pmlogger.service`
Solution	The issue is resolved in Red Hat Enterprise Linux 8.4.
Systems affected	All Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.3
Tracking number	185625

Operating system crashes on AMD Rome CPU-based systems and with Intel E810 NIC

Description	Operating system crashes when installing Red Hat Enterprise Linux 8.3 on AMD Rome CPU based systems and with Intel E810 NIC if the system has 256 logical CPUs.
Workaround	Pass nr_cpus=255 kernel command line parameter during and after installation. For more information, see Red Hat Enterprise Linux Knowledge Base article 5597881.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.4
Systems affected	Dell EMC PowerEdge C6525, Dell EMC PowerEdge R6525, and Dell EMC PowerEdge R7525.
Applies to	Red Hat Enterprise Linux 8.3
Tracking number	179019

The lvcreate command requests a response from the user when -wipesignature -yes parameters are passed

Description	The **lvcreate** command requests a response from the user when **-wipesignature -yes** parameters are passed. When -yes parameter is passed, **lvcreate** is not expected to prompt the user for a response.
Workaround	Use the **wipefs** command to clear the signatures before reusing the space for Logical Volume Manager (LVM).
Solution	The issue is resolved in Red Hat Enterprise Linux 8.4.
Systems affected	All Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	179326

The mdmonitor service displays an error during operating system installation

Description	When the operating system is installed on a virtual disk that is created by Dell EMC S150 RAID Controller, the following error message is observed during installation: `[FAILED] Failed to start Software RAID monitoring and management` This issue is a cosmetic issue and can be ignored. The operating system installation completes successfully.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.4
Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8.0 and later.
Tracking number	178369

The dmidecode utility displays the slot type as <OUT OF SPEC> for PCIe Gen 4 NVMe slots

Description	The dmidecode utility displays the Slot Type field of Type 9 record as `<OUT OF SPEC>` for PCIe Generation 4 NVMe slots.
Cause	The dmidecode utility does not support the Slot Type value of 0x24 that is defined in System Management BIOS (SMBIOS) version 3.4.
Solution	The issue is resolved in SUSE Linux Enterprise Server 15 Service Pack 2 maintenance update dmidecode version dmidecode-3.2-9.6.1.
Systems affected	Dell PowerEdge R6515, Dell PowerEdge R7515, Dell PowerEdge R6525, Dell PowerEdge R7525, and Dell PowerEdge C6525.
Applies to	SUSE Linux Enterprise Server 15 Service Pack 1 and later.
Tracking number	171945

The mcelog utility logs only decoding architectural errors message in var/log/messages

Description	The mcelog utility logs only **decoding architectural errors** message in /var/log/messages: `mcelog: Family 6 Model 106 CPU: only decoding architectural errors` `mcelog: Family 6 Model 106 CPU: only decoding architectural errors`
Solution	This issue is resolved in Red Hat Enterprise Linux 8.3. For more information, see Red Hat Enterprise Linux Knowledge Base article 6032341.
Systems affected	Dell PowerEdge R750, Dell PowerEdge R650, Dell PowerEdge C6525, Dell PowerEdge MX750C.
Applies to	Red Hat Enterprise Linux 8.2
Tracking number	166937

Disk drives part of MD RAID are not listed as installation destination by the installer

Description	During installation, if a disk drive which was previously part of MD RAID volume is used, it is not listed as an installation destination by the operating system installer.
Solution	The drive must be formatted. For more information, see Red Hat Enterprise Linux Knowledge Base article 5801081.
Systems affected	All Dell PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	192006

Dell EMC OpenManage Storage Services utility fails to reconfigure the virtual disk

Description

On a system with multiple guest VMs running storage IO and network IO operations, Dell EMC OpenManage Storage Services utility fails to reconfigure the virtual disk that is connected to the PERC H755 or PERC H755N card. The dmesg displays following call trace indicating page allocation failure for high order memory allocation requests:

```
dsm_sa_datamgrd: page allocation failure: order:6,
mode:0x6000c0(GFP_KERNEL), nodemask=(null),cpuset=/,mems_allowed=0-1
```

dmesg also shows following messages:

```
megaraid_sas 0000:01:00.0: Failed to alloc kernel SGL buffer for IOCTL
megaraid_sas 0000:01:00.0: Failed to alloc kernel SGL buffer for IOCTL
```

The issue may occur when the operating system is under pressure for high-order memory.

Workaround

Reconfigure the virtual disk from the Integrated Dell Remote Access Controller (iDRAC).

Cause

By default, the 'cache' setting for VMs is set to writeback. This results in host memory being used for buffer cache when storage IO is being performed by guest VMs.

Solution

Modify the 'cache' setting for guest VMs to 'none'. For more information, see Red Hat Enterprise Linux Knowledge Base article 5905701.

In the XML profile of guest VMs, change the following:

```
<disk type='file' device='disk'>
      <driver name='qemu' type='qcow2'/>
```

To

```
<disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' cache='none'/>
```

Systems affected	All Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.3
Tracking number	194332

Guest VMs with SRIOV VFs assigned take a long time to power on, and libvirt related errors are observed

Description

On a system with multiple guest VMs running storage I/O and network IO operations, VMs with SRIOV VFs assigned take a long time to power on. When the issue occurs, following error message is logged in /var/log/messages continuously until the guest VM successfully powers on:

```
Timed out during operation: cannot acquire state change lock (held by
monitor=remoteDispatchDomainCreate)
```

Cause

By default, the cache setting for VMs is set to write back. This results in host memory being used for buffer cache when storage I/O is being performed by guest VMs.

Solution

Modify the cache setting for guest VMs to none. For more information, see Red Hat Enterprise Linux Knowledge Base article 5886891.

In the XML profile of guest VMs, change the following:

```
<disk type='file' device='disk'>
      <driver name='qemu' type='qcow2'/>
```

To

```
<disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' cache='none'/>
```

Systems affected All Dell EMC PowerEdge servers

Applies to Red Hat Enterprise Linux 8.3

Tracking number 194002

Dmesg displays Integrity Measurement Architecture (IMA) driver related-messages during system boot

Description

Dmesg displays the following IMA driver related messages during system boot when the TPM 2.0 chip is enabled with a SHA-256 algorithm:

```
ima: Allocated hash algorithm: sha1
ima: Error Communicating to TPM chip
```

Cause The Integrity Measurement Architecture (IMA) driver currently supports the SHA-1 algorithm only.

Solution Issue is resolved in Red Hat Enterprise Linux 8.3

Systems affected Dell EMC PowerEdge yx4x and yx5x servers

Applies to Red Hat Enterprise Linux 8.1 and later

Tracking number 165260

After every reboot, the network interface name changes

Description

Network interfaces for NICs plugged into the PCI slot are named incorrectly. The issue could manifest where interfaces are named as ethN instead of the expected ensXfY format, or the network interface is incorrectly named. For example, the network interface is named ens5f0 when it is plugged into slot number 4.

Cause	There is an issue with _SUN ACPI Slot numbers in system firmware.
Solution	Update the system BIOS to 1.4.8 version or later.
Systems affected	• PowerEdge R6525: Impacted slots are 1, and 2 • PowerEdge R7525: Impacted slots are 1, 2, 3, 6, 4, 5, 7, and 8
Applies to	Red Hat Enterprise Linux 8.1
Tracking number	148546

Red Hat Enterprise Linux Version 8 installation wizard creates a duplicate bonding interface

Description	During Red Hat Enterprise Linux Version 8.1 installation, after creating a bonding and VLAN interface with the bonding interface as the parent interface, the system automatically creates a duplicate bonding, VLAN, and virtual network interfaces named as None. Later, the system encounters an unknown error, and further network configurations are not allowed. The duplicate entries are created when the bonding interface or VLAN interface does not obtain an IP address. You can resume the operating system installation by restarting anaconda service, `systemctl restart anaconda` in the shell.
Solution	Issue is resolved in Red Hat Enterprise Linux 8.3.
Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8.1
Tracking number	128135

Servers with the AMD Rome processor display a CCP initialization failure message in dmesg

Description	Servers with AMD Rome processor display the following messages in dmesg:
	`CCP initialization failed`
	`PSP Initialization failed`
	The CCP message indicates that the CCP device is not available.
	The PSP message indicates its dependency on the SEV feature.

Solution	The issue is resolved in Red Hat Linux Enterprise Server 8.2.
Systems affected	Dell PowerEdge R6515, Dell PowerEdge R6525, Dell PowerEdge R7515, Dell PowerEdge R7525, and Dell PowerEdge C6525.
Applies to	Red Hat Linux Enterprise Server 8.0, Red Hat Linux Enterprise Server 8.1
Tracking number	144920

PowerEdge servers with the AMD Rome processor fail to detect an NVMe drive after multiple hot plugs

Description	The PowerEdge servers with the AMD Rome processor fail to detect an NVMe drive after multiple hot plugs.
Cause	Issue in pciehp driver.

Workaround

NOTE: The output of each command varies based on your system configuration.

Follow the steps:

1. Identify the parent device to which the NVMe device is connected using the following command:

```
lspci -t

-+-[0000:e0]-+-00.0
|            +-00.2
|            +-01.0
|            +-02.0
|            +-03.0
|            +-03.1-[e2]----00.0
```

   From the above snippet, the device `e2:00.0` is the NVMe device and the device `e0:03.1` is the parent device. The parent device is used for the steps that are described subsequently.

2. Read the **Slot Status Register** in the PCI Express Capability structure by running the following command:

```
setpci -s e0:03.1 CAP_EXP+0x1a.w
```

   Output: 0148 (value returned)

3. Clear the event bits that are impacted by running the following command:

```
setpci -s e0:03.1 CAP_EXP+0x1a.w=0x0108
```

4. Re-read the **Slot Status Register** to confirm that event bits are cleared by running the following command:

```
setpci -s e0:03.1 CAP_EXP+0x1a.w
```

   Output: 0040 (value returned)

5. Unplug the drive and then plug-in the drive after clearing the event bits.

Solution The issue is resolved in Red Hat Enterprise Linux 8.2 z-stream kernel version kernel-4.18.0-193.13.2.el8_2.x86_64 and later

Systems affected PowerEdge R6515, PowerEdge R7515, PowerEdge R6525, PowerEdge C6525, and PowerEdge R7525

Applies to Red Hat Enterprise Linux 8

Tracking number 155501, 155503

Operating system enters the dracut shell during boot

Description When Red Hat Enterprise Linux 8.1 is installed on PowerEdge RAID Controller (PERC) virtual disk (VD) and the system is rebooted, during this boot process the system enters the dracut shell with the following error message:

```
scsi_alloc_sdev: Allocation failure during SCSI scanning
```

Cause The SCSI disk discovery failure is due to an issue in the block device allocation code.

Solution Update the kernel to 4.18.0-147.5.1.el8_1.x86_64 version or later.

Systems affected All Dell EMC PowerEdge servers

Applies to Red Hat Enterprise Linux 8.1

Tracking number 160374

System crashes when rebooted with SR-IOV-enabled QLogic cards

Description	When the system is rebooted with SR-IOV-enabled QLogic cards, the system crashes and generates a coredump (vmcore) in **/var/crash**.
Cause	The system crash is due to an issue in the qede driver.
Workaround	Disable SR-IOV before rebooting the system. To disable SR-IOV, run the following command: `echo0 > /sys/class/net/network inteface name/device/sriov_numfs`
Solution	The issue is resolved in Red Hat Enterprise Linux 7.8 and later. For more information, see SUSE Linux Enterprise Server Knowledge Base article 4991311.
Systems affected	All Dell PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	151479, 152995, 152997

After system reboot, Disk data format (DDF) devices are not listed in /proc/mdstat

Description	Disk data format (DDF)-based MD virtual disk (VD) devices are not listed in the `/proc/mdstat` file, if there are more than one MD VD devices, after rebooting the system.
Cause	Not available
Workaround	After creating the MD VD devices and before rebooting the system, run the following command: `mdadm --detail --scan >> /etc/mdadm.conf` **NOTE:** Verify the file before system reboot to confirm that the MD VD devices created earlier are not altered. For more information, see Red Hat Knowledge Base article 4990891.
Systems affected	All Dell PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	152384

Updating NVMe firmware using the nvme-cli utility displays an error in dmesg

Description	When updating the NVMe firmware using **nvme-cli**, the error message **Get FW SLOT INFO log error** is displayed. This issue occurs when the drive firmware receives a request to activate immediately without reset. This message is not observed when you update the firmware using a Dell-provided firmware update package. There is no functionality loss.
Cause	Not available
Solution	The issue is resolved in Red Hat Enterprise Linux 8.2
Systems affected	Dell EMC PowerEdge yx4x and yx5x servers

Applies to	Red Hat Enterprise Linux 8.1
Tracking number	158536

Fatal error BDF 02:00.0 is detected with BCM574xx NICs

Description	A fatal error on bus number 02, device number 00, and function number 0 (BDF 2:0.0) is observed in the following configurations: • When a graceful shutdown action is initiated. • In Gnome settings, under the **Suspend & Power Button** section, the **When the Power Button is pressed** option is set to **Suspend**.
Cause	The NIC continues to perform DMA operations even after device goes into a D3hot state.
Workaround	Set the **When the Power Button is pressed** option to **Power Off** or **Nothing**.
Solution	The issue is resolved in Red Hat Enterprise Linux 8.2
Systems affected	All Dell EMC PowerEdge servers that support BCM574xx series NICs.
Applies to	Red Hat Enterprise Linux 8.1
Tracking number	158506

NVMe devices are not detected after hot-plugging

Description	When an NVMe drive is hot-plugged, it is not detected at the PCIe and NVMe level. Running the **dmesg** command, displays the following message along with a call trace of the dead lock:
	`nvme nvme1: I/O 11 QID 0 timeout, completion polled`
	Later, the following message is displayed:
	`INFO: task irq/32-pciehp:588 blocked for more than 120 seconds`
	This issue is observed only when the NVMe device is hot plugged after booting the operating system. The NVMe drives that are plugged in before booting the operating system are detected.

Cause	Not available
Workaround	Not available
Solution	The issue is resolved in Red Hat Enterprise Linux 8.2
Systems affected	PowerEdge R740xd, PowerEdge R7515, and PowerEdge R7425
Applies to	Red Hat Enterprise Linux 8.0
Tracking number	130372

Linux operating system fails to detect the Intel x710 card

Description	The i40e in-box driver is not enabled to detect an Intel x710 NIC. NIC details can be identified by running the command **lspci**.

The output of the command is as follows:

```
Ethernet controller: Intel Corporation Ethernet Controller x710 for
10GBASE-T
Vendor ID: 8086
Device ID: 15FF
```

Cause	The i40e in-box driver is not enabled to detect an Intel x710 NIC.
Workaround	Install the out-of-box drivers.
Solution	The issue is resolved in SUSE Linux Enterprise Server 15 SP2, SUSE Linux Enterprise Server 15 SP1 Maintenance Update kernel-default-4.12.14-197.29.1 version and later.
Systems affected	Dell PowerEdge R6515, Dell PowerEdge R6525, Dell PowerEdge R7515, Dell PowerEdge C6525, and Dell PowerEdge R7525.
Applies to	Red Hat Enterprise Linux 8.1
Tracking number	146136, 146448, 146451, 152855

Dmidecode displays OUT OF SPEC in Slot Type and Slot Length of SMBIOS system slots

Description	The execution of `dmidecode` command displays an **OUT OF SPEC** message in the **Slot Type** and **Slot Length** fields of the SMBIOS system slots (Type 9) structure.
Cause	The dmidecode tool does not recognize the 0x5h and 0x6h values in the **Slot Length** field and 0x20h value in the **Slot Type** field in SMBIOS specification version 3.0.0.
Workaround	Not available
Solution	The issue is resolved in Red Hat Enterprise Linux 8.2
Systems affected	Dell PowerEdge R6515, Dell PowerEdge R7515, Dell PowerEdge R6525, and Dell PowerEdge R7525
Applies to	Red Hat Enterprise Linux 8.0 kernel-4.18.0-80.7.1.el8 or later
Tracking number	150263, 157470, 157474

Custom partitioning fails with FC LUN

Description	While installing the Linux operating system on FC LUN, if **Custom Partitioning** is selected, partition creation fails in: • `/home` under automatic option • `/ and /home` under manual standard partition This issue is not observed if **Automatic Partition** is selected in the installation wizard.
Cause	Not available
Workaround	Not available
Solution	The issue is resolved in Red Hat Enterprise Linux 8.2
Systems affected	All Dell PowerEdge servers
Applies to	Red Hat Enterprise Linux 8
Tracking number	155150

When booting the system from iSCSI with Mellanox CX-4 and CX-5 adapters, the system reports csum failure message

Description

When booting the system from iSCSI with Mellanox CX-4 and CX-5 adapters, the following message is displayed in the `dmesg` and **/var/log/messages** reports:

```
localhost kernel: ibft0: hw csum failure
```

This message can be ignored.

Cause

Some networking switches which are sending frames with nonzero padding bytes maybe causing the issue.

Workaround

Install out-of-box drivers.

Solution

The issue is resolved in SUSE Linux Enterprise Server 15 SP2.

Systems affected

All Dell PowerEdge systems

Applies to

Red Hat Enterprise Linux 8

Tracking number

156922

Red Hat Enterprise Linux 8 kernel panic is observed due to fatal hardware error

Description

On Dell EMC PowerEdge servers with Operating System Red Hat Enterprise Linux 8, kernel panic is observed when system is booting, shut down, or rebooting. When the kernel panic occurs, the following message is displayed:

```
[    1.410443] {1}[Hardware Error]: Hardware error from APEI Generic
Hardware Error Source: 3
[    1.419667] {1}[Hardware Error]: event severity: fatal
[    1.425399] {1}[Hardware Error]: Error 0, type: fatal
[    1.431130] {1}[Hardware Error]: section_type: PCIe error
[    1.437347] {1}[Hardware Error]: port_type: 4, root port
[    1.443465] {1}[Hardware Error]: version: 1.16
[    1.448617] {1}[Hardware Error]: command: 0x0143, status: 0x4010
[    1.455514] {1}[Hardware Error]: device_id: 0000:80:02.0
[    1.461633] {1}[Hardware Error]: slot: 4
[    1.466201] {1}[Hardware Error]: secondary_bus: 0x82
[    1.471932] {1}[Hardware Error]: vendor_id: 0x8086, device_id:
0x6f04
[    1.479312] {1}[Hardware Error]: class_code: 000406
[    1.484948] {1}[Hardware Error]: bridge: secondary_status: 0x2000,
control: 0x0003
[    1.493588] Kernel panic - not syncing: Fatal hardware error!
```

Solution

Update the system BIOS.

- For yx3x servers, update BIOS to 2.10.5 version.
- For yx4x servers, update BIOS to 2.2.11 version.

Download the latest BIOS version from Drivers Support.

Systems affected

All Dell EMC PowerEdge systems supported by Red Hat Enterprise Linux 8. For more information, see Red Hat Enterprise Linux Certification Matrix for Dell EMC PowerEdge Servers.

Applies to

Red Hat Enterprise Linux 8

Tracking number

129907

6

Known issues

Topics:

Broadcom NIC port naming is incorrect in RHEL 8.9

Description	On a Dell PowerEdge system that is configured with Broadcom Inc. subsidiaries BCM57414 and RHEL 8.9, the NIC port naming is showing different values for port1 and port2 as below from the ifconfig output:
	`enp156s0: flags=4099<UP, BROADCAST, MULTICAST> mtu 1500`
	`ether 40:a6:07:99:2a:04 txqueuelen 1000 (Ethernet)`
	
	`ens1: flags=4099<UP, BROADCAST, MULTIMAST> mtu 1500`
	`ether 40:a6:67:99:2a:00 txqueuelen 1000 (Ethernet)`

Workaround	Specify naming scheme as kernel command line parameter according to user requirements to ensure compatibility with scripts and configurations that rely on a specific naming convention. For example, add "net.naming-scheme=v243" kernel command line parameter.
Systems affected	All Dell PowerEdge systems
Applies to	Red Hat Enterprise Linux 8.9 and later.
Tracking number	291224

Unable to add gid with error=-14 in dmesg

Description	On a Dell PowerEdge system configured with Broadcom BCM57508 2x100G QSFP PCIE and RHEL 8.10. Multiple error messages are observed as below when trying to create bonding using BroadcomBCM57508 NIC interfaces. This message does not affect the bonding functionality. `"__ib_cache_gid_add: unable to add gid fe80:0000:0000:0000:5e6f:69ff:fef5:af90 error=-14"`
Workaround	N/A
Systems affected	All Dell PowerEdge systems. The issue is not observed in RHEL 9.4.
Applies to	Red Hat Enterprise Linux 8.10
Tracking number	304789

snmpd output the message "systemstats_linux: unexpected header length in /proc/net/snmp. 237 != 224"

Description	On a Dell PowerEdge system installed with RHEL 8.10 and OMSA, when snmp walk is executed the following message is displayed. The Add Linux 6.7 compatibility parsing /proc/net/snmp upstream patch fixes the issue. `systemstats_linux: unexpected header length in /proc/net/snmp. 237 != 224`
Cause	A change was introduced in the kernel that added the OutTransmits field to Ip in /proc/net/snmp. This broke the net-snmp.
Workaround	Not available
Systems affected	All Dell PowerEdge systems
Applies to	Red Hat Enterprise Linux 8.10
Tracking number	305474

Unable to switch to portrait mode with Matrox G200eW3 controller

Description	On a Dell PowerEdge system installed with RHEL 8.8 or later, the I/O (mouse and keyboard) does not respond and the operating system is unavailable when trying to change display settings from landscape to other. The orientation of the screen changes, yet the operating system remains stuck. The mouse and keyboard do not respond, making it impossible to select the "keep changes" option. Therefore, after 20 s, the screen orientation resets to the default landscape mode.
Cause	Wayland is disabled by default for Matrox chipsets on RHEL 8 and the X11 backend has a issue and fails to work in portrait mode with Matrox chipsets.
Workaround	Configure GDM to use the Wayland graphics backend. For more information, see Red Hat Enterprise Linux Knowledge Base article 6231831.
Systems affected	All Dell Power Edge Systems
Applies to	Red Hat Enterprise Linux 8.8 and later.
Tracking number	288724

The dmesg displays qat related messages

Description	On a PowerEdge system installed with the RHEL 8.6 operating system, the `Failed to load MMP firmware Qat_4xxx_mmp.bin` message is observed in dmesg. The message indicates that the firmware files required for Intel qat to function are not installed in the system under `/lib/firmware/`. To avoid the error messages, qat needs to be configured correctly.
Systems affected	All Dell PowerEdge systems
Applies to	Red Hat Enterprise Linux 8.6
Tracking number	275032

Booting into operating system fails on systems with TPM 1.2 chips

Description	On systems with TPM 1.2 chips, booting into operating system fails when **TPM Security** field is set to **On without Pre-boot Measurements** in BIOS.
Cause	When the option **On without Pre-boot Measurements** is set in BIOS, the shim utility cannot write to TPM PCR registers. Shim considers this as a fatal error and fails to boot.
Workaround	In BIOS, set **TPM Security** field to **On with Pre-boot Measurements**.
Systems affected	All Dell PowerEdge systems supporting TPM 1.2
Applies to	RedHat Enterprise Linux 8 and later.
Tracking number	209250, 209177

FCoE session is not reestablished after MX9116N switch is rebooted

Description	On Dell EMC PowerEdge MX740C system containing Qlogic QL41000 series CNA device configured for FCoE, FCoE session is not restored automatically when the MX9116N switch is rebooted.
Cause	Issue in qedf driver related to handling the switch reboot scenario.
Workaround	To reestablish FCoE session execute the command: `rescan-scsi-bus.sh -i --hosts=<host number>` The <host number> refers to the host number associated with the lost FCoE session and this can be identified from dmesg.
Systems affected	All Dell EMC systems supporting Qlogic QL41000 series CNA devices.
Applies to	Red Hat Enterprise Linux 8.2
Tracking number	190870

Dmesg displays error messages when NVMe device is surprise removed

Description	Dmesg or /var/log/messages show the following error messages after an NVMe device is unbound from the NVMe driver and surprise removed:
	`kernel: pcieport 0000:b0:06.0: Timeout waiting for Presence Detect`
	`kernel: pcieport 0000:b0:06.0: link training error: status 0x8001`
	`kernel: pcieport 0000:b0:06.0: Failed to check link status`

Solution	The issue is a cosmetic issue and can be ignored.
Cause	The error that is displayed is due to an issue with the pciehp driver.
Systems affected	Dell EMC PowerEdge R740XD and Dell EMC PowerEdge R7525.
Applies to	Red Hat Enterprise Linux 8.2 and later
Tracking number	180987

Status of the RAID 0 logical volume is displayed as Available when one of the members of the RAID array is surprise removed

Description	When Logical Volume Manager (LVM) is used to create a RAID 0 array and a member of the RAID array is surprise removed, the lvdisplay command shows the logical volume (LV) status as 'Available'.
Solution	Use the command lvs -o +lv_health_status to check the status of the RAID array. The command displays the output Partial when a member of the RAID array is removed.
Systems affected	All Dell EMC PowerEdge systems supporting NVMe Surprise Removal.
Applies to	Red Hat Enterprise Linux 8.2 and later
Tracking number	175865

/proc/mdstat and mdadm -D commands display incorrect statuses when two NVMe devices are surprise removed from a RAID 5 MD array

Description	When two of three NVMe devices are surprise removed from a RAID 5 MD array, the command `cat/ proc/mdstat` displays the array status incorrectly as active. Similarly, when the status of the MD RAID is queried using the `mdadm -D /dev/mdN` command, the number of active and working devices that are displayed is two. Only the status of the array reported is incorrect. However, when I/O operations are performed, I/O errors are observed as expected.
Cause	When the number of devices that are surprise removed exceeds the number of devices that are required for the array to function, the MD status is not updated.
Systems affected	All Dell EMC PowerEdge systems supporting NVMe Surprise Removal.
Applies to	Red Hat Enterprise Linux 8.2 and later
Tracking number	182820

Dell Controlled Turbo feature is not functional

Description	Dell Controlled Turbo is not functional when the system profile setting in BIOS is set to performance or custom.
Cause	The intel_pstate driver may interfere with the processor settings which results in Dell Controlled Turbo feature being not functional.
Solution	Prevent the intel_pstate driver from loading. For more information on how to prevent the module from loading, see Red Hat Enterprise Linux Knowledge Base article 41278.
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	167802

Caps Lock key-press is not registered on the Dell PowerEdge iDRAC virtual console

Description	The Caps Lock key-press is not registered when the Dell PowerEdge iDRAC virtual console is used. The console displays lowercase letters when the Caps Lock key is enabled.
Workaround	Do the following: 1. Create the file /etc/udev/rules.d/99-kbd.rules and add the rule: `ACTION=="add", SUBSYSTEM=="leds", ENV{DEVPATH}=="*/input*::capslock", ATTR{trigger}="kbd-ctrllock"` 2. Run the command: `udevadm control --reload-rules` 3. Disconnect and reconnect the Dell PowerEdge iDRAC virtual console.
Systems affected	All Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0, Red Hat Enterprise Linux 8.1, and Red Hat Enterprise Linux 8.2
Tracking number	176749

Advanced Configuration and Power Interface (ACPI) error messages displayed in dmesg

Description	The following ACPI-related error messages are displayed in dmesg and /var/log/messages:
	`[ 10.265354] ACPI Error: No handler for Region [SYSI] (ffff8f04edb31a20) [IPMI] (20130517/evregion-162)`
	`[ 10.265359] ACPI Error: Region IPMI (ID=7) has no handler (20130517/exfldio-305)`
	`[ 10.265368] ACPI Error: Method parse/execution failed [\_SB_.PMI0._GHL] (Node ffff8f106dea07e0), AE_NOT_EXIST (20130517/psparse-536)`
	`[ 10.265376] ACPI Error: Method parse/execution failed [\_SB_.PMI0._PMC] (Node ffff8f106dea0720), AE_NOT_EXIST (20130517/psparse-536)`
	`[ 10.265381] ACPI Exception: AE_NOT_EXIST, Evaluating _PMC (20130517/power_meter-753)`
	The error messages are related to the module dependency and module load order of the acpi_power_meter and acpi_ipmi drivers.

Solution	The messages are informational and can be ignored. For more information, see Red Hat Enterprise Linux Knowledge Base article 49225.
Applies to	Red Hat Enterprise Linux 8 and later

Tracking number	172165

Drivers available in OEMDRV drive are not installed during the operating system installation

Description	Intermittently, during the installation of Red Hat Enterprise Linux, the operating system cannot enumerate the USB drive at the stage of driver disk probing. This prevents the drivers from LC/OEMDRV from being installed.
Workaround	Use the kernel boot parameter command inst.dd=LABEL=OEMDRV to install the drivers in the OEMDRV drive.
Systems affected	Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8 and later
Tracking number	168569

The Mellanox IB devices are listed under an incorrect device category on Red Hat Enterprise Linux 8

Description	The Mellanox IB devices are listed under the **Bluetooth** section of the **Network settings** tab in Red Hat Enterprise Linux 8 and Red Hat Enterprise Linux 8.1 operating systems. This is a cosmetic issue and can be ignored.
Workaround	Use the nmcli tool or the nm-connection-editor to display available devices.
Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8 and later
Tracking number	123915

The lspci utility is unable to read Vital Product Data (VPD) from QLogic QLE2692 adapter

Description

The lspci utility is unable to read Vital Product Data (VPD) from QLogic QLE2692 card and displays the following error message:

```
# lspci -vvv -s <PCI bus/device/func>
[…]
Capabilities: [88] Vital Product Data
pcilib: sysfs_read_vpd: read failed: Input/output error
Not readable
```

This issue is a cosmetic issue and can be ignored.

Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8 and later
Tracking number	123915

Driver dependency mismatch errors reported while installing out-of-box drivers on Red Hat Enterprise Linux 8.x

Description	Driver dependency mismatch errors are reported when `inst.dd` installation methods (flash drive, driver package, HTTP, NFS, FTP) are used to install the out-of-box (OOB) drivers during operating system installation. The dependency errors are not observed after the initial installation as the operating system will use the out-of-box (OOB) drivers. The issue is industry wide and is not specific to Dell EMC PowerEdge servers.
Workaround	Block the installation of the inbox versions of the OOB drivers.
Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8 and later
Tracking number	176624

Dmesg displays TPM and nvdimm related-messages in Red Hat Enterprise Linux 8.1

Description	When Red Hat Enterprise Linux Version 8.1 is booted with the TPM 2.0 chip enabled, dmesg displays the following messages in var/log/message:
	`No TPM handle discovered.`
	`failed to open file /etc/ndctl/keys/nvdimm-master.blob: No such file or directory`
	There is no impact to the functionality of the TPM chip if it is enabled in the system BIOS. This is a cosmetic error and can be ignored.

Systems affected	Dell EMC PowerEdge yx4x and yx5x servers
Applies to	Red Hat Enterprise Linux 8.1 and later
Tracking number	164988

Link Up message is observed when the NVMe device slot is powered off and the device is unplugged

Description	The message **Link Up** is observed in the dmesg log when the NVMe device slot is powered off and the device is unplugged by running the following command: `echo 0 > /sys/bus/pci/slots/<slot number>/power` The expected message is **Link Down** because the device is unplugged.
Cause	This is due to an issue with the pciehp driver.
Workaround	This issue is a cosmetic issue and can be ignored. The NVMe device is recognized when plugged into the slot.
Systems affected	Dell EMC PowerEdge R740xd and Dell EMC PowerEdge R7425
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	160773

Mellanox InfiniBand adapters are listed in Bluetooth

Description	Mellanox InfiniBand adapters are listed in Gnome, under the **Network** > **Settings** > **Bluetooth** section.
Cause	Not available
Workaround	Not available. This issue is a cosmetic issue. Use other tools like **nmcli**, and **nm-connection-editor** to view and configure the InfiniBand interfaces.
Systems affected	All Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0 and later
Tracking number	133715

iscsiadm output displays STATIC in the iface.bootproto field when the network interface is configured to DHCP

Description	When the system is booting from iSCSI with the Emulex OneConnect device that is configured with DHCP IP, the **iface.bootproto** field is listed as **STATIC** in the command **iscsiadm -m fw** output. There is no functionality loss.
Cause	Not available
Workaround	Not available
Systems affected	Dell PowerEdge yx3x servers that support Emulex OneConnect
Applies to	Red Hat Enterprise Linux 8.1 and later
Tracking number	147877

When system reboots, system stops responding at the end of the reboot process

Description	When the system reboots, system stops responding at the end of the reboot process. When the system is in this state, it does not respond to keyboard or mouse device and iDRAC virtual console displays the **No signal** message.
Workaround	To recover the system, perform any one of the following: • Using iDRAC GUI, on the **Dashboard** page, click **Graceful Shutdown** drop-down, and then select **Reset System(warm reboot)**. • Restart the system by pressing the Power button.
Solution	Update system BIOS to 1.4.8 version.
Systems affected	Dell PowerEdge R6515, Dell PowerEdge R6525, Dell PowerEdge R7515, and Dell PowerEdge C6525.
Applies to	Red Hat Enterprise Linux 8
Tracking number	141837

Unable to shut down RHEL 8 when you select Graceful shutdown option or when you press power button on the server

Description	When you select the **Graceful shutdown** option from any system management interface such as iDRAC or press the power button, the system goes to suspended state and stops all the tasks. So, firmware updates through iDRAC interface that require restart take longer, and the system is hard reset upon iDRAC watchdog time expiration.
Workaround	Change the power settings by performing the following steps: 1. In the **Settings** pane, click **Power**. 2. In the **Suspend & Power Button** section, from the **When the Power Button is pressed** list, select **Power Off**. **NOTE:** By default, **Suspend** option is selected. **NOTE:** The preceding workaround may not work when the system is locked. GNOME prevents the accidental shutdown when system is locked. As a security precaution, GNOME does not allow any power related actions such as: • Shutdown from any system management interfaces like iDRAC • Pressing the power button unless the user is active and logged in to the operating system.
Systems affected	N/A
Applies to	Red Hat Enterprise Linux 8.0
Tracking number	137495

iSCSI LUN not discovered during RHEL 8 installation

Description	When Red Hat Enterprise Linux 8 is installed on an iSCSI LUN, the LUN is not discovered automatically.
Cause	If the Network interface displays the status as **UP** but displays **NO-CARRIER**, it indicates that **dracut** has reduced the wait time to run the network interface from 7 seconds to 5 seconds. This reduction in time causes **dracut** to assume that the network interface is not connected, which results in the LUN not being discovered.
Workaround	• During the iSCSI installations, in addition to the boot parameter **rd.iscsi.ibft=1**, add the following boot parameter: **rd.net.timeout.carrier=7** • After the iSCSI installation, during boot, add the following boot parameter: **rd.net.timeout.carrier=7**
Systems affected	All Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0
Tracking number	124792

RHEL 8 installation fails on systems with Emulex OneConnect card

Description	When RHEL 8 is installed on an iSCSI LUN that is configured on Emulex OneConnect card on Dell EMC 13G systems, the installation fails with a **UnicodeDecodeError**. This error occurs because the Anaconda installer is unable to decode the iSCSI initiator name.

Workaround	There is no workaround available for this issue, but the LUN can be used as DATA LUN.
Systems affected	All Dell EMC 13G systems
Applies to	Red Hat Enterprise Linux 8.0
Tracking number	111760

Switching between runlevels fails

Description	Switching from runlevel 5 to 3 fails.
Workaround	Switch to different TTYs (Ctrl+Alt+F2/F3) to continue using the system.
Systems affected	All Dell EMC PowerEdge servers
Applies to	Red Hat Enterprise Linux 8.0
Tracking number	124039

7

Limitations

- Unattended operating system installation of Red Hat Enterprise Linux 8.0 and later updates using iDRAC Server Configuration Profile (SCP) method is not supported up to RHEL 8.8.
- The `auth` and `authconfig` kickstart commands require the AppStream repository.
- The `--interactive` option of the `ignoredisk` kickstart command does not work in RHEL 8.
- RHEL 8 system becomes unresponsive when many devices are connected.
- Physical memory hot plugging does not work.
- Database servers are not installable in parallel.

For more information about the limitations of RHEL 8, see Red Hat Enterprise Linux Documentation.

8

Resources and support

This section provides information about the documentation resources for your server.

Table 2. Additional documentation resources for your server

Task	Document	Location
Setting up your server	For information about installing the server into a rack, see the Rack documentation included with your rack solution or the *Getting Started Guide* that is shipped with your server.	PowerEdge Manuals
	For information about turning on the server and the technical specifications of your server, see the *Getting Started Guide* that is shipped with your server.	PowerEdge Manuals
Configuring your server	For information about the iDRAC features, configuring and logging in to iDRAC, and managing your server remotely, see the Integrated *Dell Remote Access Configuration Tool User's Guide*	iDRAC Manuals
	For information about installing the operating system, see operating system documentation.	Operating System Manuals
	For information about understanding Remote Access Controller Admin (RACADM) subcommands and supported RACADM interfaces, see the RACADM Command Line Reference Guide for iDRAC.	iDRAC Manuals
	For information about updating drivers and firmware, see Download drivers and firmware topic in this document.	Drivers
Managing your server	For information about server management software offered by Dell, see the Dell *Systems Management Overview Guide*.	OpenManage Manuals
	For information about setting up, using, and troubleshooting OpenManage, see the Dell *OpenManage Server Administrator User's Guide*.	OpenManage Manuals
	For information about installing, using, and troubleshooting Dell OpenManage Essentials, see the Dell OpenManage Manuals >	OpenManage Manuals

Table 2. Additional documentation resources for your server (continued)

Task	Document	Location
	OpenManage Essentials *User's Guide*.	
	For information about installing and using Dell SupportAssist, see the Dell *SupportAssist Enterprise User's Guide*.	Software Serviceability Tools
	For understanding the features of Dell Lifecycle Controller (LC), see the *Dell Lifecycle Controller User's Guide*.	iDRAC Manuals
	For information about partner programs enterprise systems management, see the *OpenManage Connections Enterprise Systems Management* documents.	ESM Manuals
	For information about viewing inventory, performing configuration, and monitoring tasks, remotely turning on or off servers, and enabling alerts for events on servers and components using the Dell Chassis Management Controller (CMC), see the *Chassis Management Controller User's Guide* .	ESM Manuals
Working with the Dell PowerEdge RAID controllers	For information about understanding the features of the Dell PowerEdge RAID controllers (PERC) and deploying the PERC cards, see the Storage controller documentation.	Storage Controller Manuals
Understanding event and error messages	For information about checking the event and error messages generated by the system firmware and agents that monitor server components, see the Dell *Event and Error Message Reference Guide for 14th Generation Dell PowerEdge Servers* .	OpenManage Manuals
Troubleshooting your system	For information about identifying and troubleshooting the PowerEdge server issues, see the *PowerEdge Servers Troubleshooting Guide*.	PowerEdge Manuals

Topics:

- Download drivers and firmware

Download drivers and firmware

It is recommended that you download and install the latest BIOS, drivers, and systems management firmware on your system.

Ensure that you clear the web browser cache before downloading the drivers and firmware.

1. Go to Drivers.
2. In the **Drivers & Downloads** section, enter the Service Tag of your system in the **Enter a Dell Service Tag, Dell Product ID, or Model** text box, and then click the right arrow button.

   NOTE: If you do not have the Service Tag, click **Detect PC** to allow the system to automatically detect your Service Tag.

3. Click **Drivers & Downloads**.
   A list of applicable downloads is displayed.
4. Download the drivers or firmware to a USB drive, CD, or DVD.

9

Contacting Dell EMC

Dell provides several online and telephone-based support and service options. Availability varies by country, region, and product, and some services may not be available in your area.

To contact Dell for sales, technical assistance, or customer service issues, see Contact Dell.

If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or the product catalog.