

# NVMe Surprise Removal on Dell EMC PowerEdge servers running Linux operating systems

## Abstract

This white paper describes the support for Non-Volatile Memory Express (NVMe) Surprise Removal on Dell EMC PowerEdge servers running supported Enterprise Linux operating systems.

October 2020

## Revisions

Date	Description
October 2020	Initial release

## Acknowledgements

Author: Narendra K

Support: Austin Bolen, Gurupreet Kaushik, Sherry Keller

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 27/10/2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

# Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents .....	3
Executive summary.....	4
1 Introduction.....	5
1.1 Audience and scope .....	5
1.2 Terminology .....	5
1.3 Command-line utilities used for verifying surprise removal of NVMe devices.....	5
2 Surprise removal of NVMe devices .....	6
2.1 Supported and unsupported scenarios for surprise removal of NVMe devices .....	6
2.2 Identifying the NVMe device slot and verifying surprise removal.....	6
2.3 Platform and operating system support summary.....	7
3 Known issues with NVMe surprise removal .....	8
3.1 SUSE Linux Enterprise Server Service Pack 2 .....	8
3.1.1 MD RAID layer is not notified of the surprise removal of Samsung NVMe devices.....	8
3.1.2 Status of the RAID 0 logical volume is displayed as Available when one of the members of the RAID array is surprise removed.....	8
3.1.3 LVM does not activate a free physical volume when one of the NVMe devices is surprise removed .....	8
4 Summary .....	9
5 References .....	10

## Executive summary

NVMe devices are being used more widely, and features such as surprise removal are important to the continuous availability of the server and serviceability needs. Surprise removal allows you to remove a device from the server without prior notification. This white paper outlines the best practices that are to be followed for the surprise removal of NVMe devices running supported Linux operating systems on supported Dell EMC PowerEdge servers. Both supported and unsupported scenarios and known issues encountered while performing surprise removal on Linux operating systems are documented in this white paper.

# 1 Introduction

As NVMe devices are being used more widely, they must provide enterprise functionality such as surprise removal that you rely on.. Surprise removal enhances the serviceability of NVMe devices by eliminating additional steps required to prepare the devices for orderly removal and ensures availability of servers by eliminating server downtime.

## 1.1 Audience and scope

The intended audience for this white paper includes IT administrators and those using hot-pluggable NVMe devices on Dell EMC PowerEdge servers running supported enterprise Linux operating systems.

## 1.2 Terminology

**Hot insertion:** Connecting the NVMe device to the server when the Linux operating system is booted up.

**Surprise removal:** Removing the NVMe device from the Linux operating system without notifying the operating system beforehand.

**Orderly removal:** Removing the NVMe device from the server after completing the prerequisites, such as suspending all processes accessing the NVMe device and quiescing all I/O operations accessing the NVMe device.

**Hot swap:** Replacing an existing NVMe device with a new NVMe device from the same or different vendor while the host operating system is booted. Hot swap is a surprise removal or orderly removal followed by a hot insertion operation with a different NVMe device.

## 1.3 Command-line utilities used for verifying surprise removal of NVMe devices

The following command-line utilities that are available in the enterprise Linux operating systems are used to verify hot-plug operations:

- nvme-cli
- lspci
- lsblk

## 2 Surprise removal of NVMe devices

### 2.1 Supported and unsupported scenarios for surprise removal of NVMe devices

The following table describes the supported and unsupported scenarios while performing surprise removal of NVMe devices.

Table 1 Supported and unsupported scenarios for surprise removal of NVMe devices

Supported scenarios	Unsupported scenarios
<p>Surprise removal of a single NVMe device at a time is supported.</p> <p>The following requirements ensure successful surprise removal of NVMe devices:</p> <ul style="list-style-type: none"> <li>• Surprise removal must be performed within one-second period, as a slower surprise removal may cause the operating system to crash.</li> <li>• To avoid an operating system crash, a fifteen-second time interval should be provided between successive hot-plug operations to ensure that the operating system, applications, and drivers have enough time to fully handle the operation.</li> </ul>	<ul style="list-style-type: none"> <li>• Performing surprise removal of the drive that has the operating system installed or the drive that has a swap partition.</li> <li>• Performing surprise removal when the operating system is booting up.</li> <li>• Performing surprise removal of an NVMe device when another NVMe device is being hot inserted, or within 15 seconds of another NVMe device being hot inserted.</li> <li>• Performing surprise removal of two or more NVMe devices serially without a fifteen second time interval between the surprise removals.</li> <li>• Surprise removal of an NVMe device that is either directly or partially assigned to a virtual machine.</li> </ul>

**Note:** Specific solutions may have additional requirements to perform successful surprise removal. For more information, see your solution documentation.

### 2.2 Identifying the NVMe device slot and verifying surprise removal

This section describes a scenario where `/dev/nvme0n1` is the device to be surprise removed. The slot numbers used in this section are specific only to this use case.

**Note:** Surprise removing an NVMe device that is in use may result in data loss. It is recommended that you create a data backup before surprise removing the NVMe device.

To perform surprise removal of an NVMe device:

1. Use the command `nvme list` to list the NVMe devices connected to the server.
2. Use the command `nvme list-subsys` to retrieve the PCI bus/device/function number of the `/dev/nvme0n1` device.
3. Determine the PCIe slot number using the PCI bus/device/function number and surprise remove the NVMe device from slot 22.

```
localhost:~ # cat /sys/bus/pci/devices/0000\:3d\:00.0/label
PCIe SSD in Slot 22 Bay 1
```

Figure 1 Determining the PCIe slot number of the /dev/nvme0n1

4. To verify that the operating system successfully unregisters the device:
  - a. Use the command `nvme list` to list the connected devices and verify that the /dev/nvme0n1 is not listed.
  - b. Use the command `lspci` to verify PCIe device 0000:3d:00.0 is not listed.
  - c. Use the command `lsblk` to verify that the /dev/nvme0n1 is not listed.

---

**Note:** The operating system might crash if subsequent hot-plug operations are not performed at time intervals of at least fifteen seconds.

---

## 2.3 Platform and operating system support summary

The following table lists the Dell EMC PowerEdge servers and the Linux operating systems that support NVMe surprise removal.

Table 2 Supported Dell EMC PowerEdge servers and Linux operating systems that support NVMe surprise removal

Dell EMC PowerEdge generation	SUSE Linux Enterprise Server Service Pack 2	
	Supported	Unsupported
Intel Skylake and Cascade Lake SP CPU based yx4x servers	<ul style="list-style-type: none"> <li>• Hot insertion</li> <li>• Orderly removal</li> <li>• Surprise removal</li> </ul>	
AMD Naples CPU based yx4x servers	<ul style="list-style-type: none"> <li>• Hot insertion</li> <li>• Orderly removal</li> <li>• Surprise removal</li> </ul>	
AMD Rome CPU based yx5x servers	<ul style="list-style-type: none"> <li>• Hot insertion</li> <li>• Orderly removal</li> </ul>	<ul style="list-style-type: none"> <li>• Surprise removal</li> </ul>

---

**Note:** Linux upstream kernel version 5.7 and later have hot-plug related patches that enhance hot-plug user experience.

---

## 3 Known issues with NVMe surprise removal

The following section describes the known issues encountered when surprise removal is performed on servers running supported Linux operating systems.

### 3.1 SUSE Linux Enterprise Server Service Pack 2

#### 3.1.1 MD RAID layer is not notified of the surprise removal of Samsung NVMe devices

**Description:** When a virtual disk is created on the MD RAID layer using Samsung NVMe device, the MD RAID layer is not notified of the surprise removal of the NVMe drive. The output of the `mdadm -D` command displays an incorrect status of the MD RAID virtual disk. The issue is observed on Dell Express Flash PM1725a, PM1725b, Enterprise NVMe agnostic devices. The array status reporting is incorrect. When I/O operations are performed, I/O errors are observed as expected and the file-system changes to read-only.

**Cause:** The issue is observed on handling devices which showcase multipath capability.

**Workaround:** Pass the `multipath=N` module parameter to the `nvme_core` driver.

#### 3.1.2 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.

#### 3.1.3 LVM does not activate a free physical volume when one of the NVMe devices is surprise removed

**Description:** When one of the members of a RAID 1 LVM array is surprise removed, the LVM does not replace the removed device with a free physical volume (PV) that is available in the volume group.

**Cause:** The issue is related to the handling of failover logic in the LVM.

**Workaround:** The command `lvconvert --repair` can be used to add the free PV to the RAID 1 LVM array.

**Solution:** The issue is resolved in the following Program Temporary Fix: [www.ptf.suse.com/sle-modulebasesystem-15-sp2/20119/x86\\_64/20200820](http://www.ptf.suse.com/sle-modulebasesystem-15-sp2/20119/x86_64/20200820).



## 4 Summary

This white paper describes the concept of NVMe surprise removal and provides guidance on how to perform surprise removal on supported enterprise Linux operating systems on supported Dell EMC PowerEdge servers. The step-by-step instructions for performing NVMe surprise removal are documented with guidelines to be followed for successful surprise removal of NVMe devices. This document will be updated if there is a change in the support offered for surprise removal or if there are any major enhancements to the scenarios involving this feature. Further known issues related to surprise removal will be updated on the respective release notes document published on the operating system documentation page of [www.dell.com/support](http://www.dell.com/support).

## 5 References

- [Dell Express Flash NVMe PCIe SSD User's Guide](#)
- [SUSE Linux Enterprise Server Certification Matrix for Dell EMC PowerEdge Servers](#)
- [Dell EMC PowerEdge Systems Running SUSE Linux Enterprise Server 15 Release Notes](#)