



Intel® Virtual RAID on CPU (Intel® VROC) RAID Management in Red Hat Enterprise Linux*

Configuration Guide

Revision 002

November 2023



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis. You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any Intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

All product plans and roadmaps are subject to change without notice.

The products described may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at Intel.com.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

© Intel Corporation. Intel, the Intel logo, Xeon, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

*Other names and brands may be claimed as the property of others.

Copyright© 2023, Intel Corporation. All rights reserved.

Contents

1	About This Document	5
1.1	Preconditions and General Assumptions	5
1.2	Reference OEM Platform Documentation	5
2	Managing Intel® VROC RAID.....	6
2.1	Configure /etc/mdadm.conf.....	6
2.2	Enable LEDMON Service.....	6
2.3	Enable Postfix Service.....	7
2.4	Configure LVM on Intel® VROC RAID	8
3	Monitoring Intel® VROC RAID Alerts in RHEL	9
3.1	Intel® VROC RAID Alerts in RHEL	9
3.2	Develop a Program to Handle Intel® VROC Alerts	9

Figures

Figure 2-1. Intel® VROC Email Notification Example	8
--	---

Tables

Table 3-1. Intel® VROC RAID Alerts in RHEL.....	9
---	---



Revision History

Revision	Description	Date
001	Initial release	March 2023
002	Public release with updated LVM configuration	November 2023

§§

1 About This Document

The Intel® Virtual RAID on CPU (Intel® VROC) RAID Management in Red Hat Enterprise Linux* (RHEL) Configuration Guide is aimed to guide system administrators with the best practices to configure and manage Intel® VROC in the supported RHEL OS environment. This configuration guide is based on the Intel® Virtual RAID on CPU (Intel® VROC) User Guide for Linux*. Users should reference the latest [Intel® Virtual RAID on CPU \(Intel® VROC\) User Guide for Linux*](#) for the details.

1.1 Preconditions and General Assumptions

Intel® Virtual RAID on CPU (Intel® VROC) is the integrated RAID solution on the Intel® Xeon® Scalable Processor based server platforms. It is a platform integrated RAID system which requires the platform hardware, UEFI firmware, OS drivers as well as Intel® VROC RAID management tools to cooperate together. Therefore, the following assumptions or preconditions should be met, in order to use and manage Intel® VROC RAID under RHEL correctly.

1. The server platform with Intel® VROC must be enabled.
2. Intel® VROC Standard or Premium key (license) is installed or acquired.
3. Intel® Volume Management Device (Intel® VMD) is enabled in the UEFI.
4. Intel® SATA controllers are configured in RAID mode in the UEFI (if SATA RAID is required).
5. Install the latest Linux* out-of-box driver of Intel® VROC for some specific RHEL operating systems.

Consult the server platform supplier or reference the latest [Intel® Virtual RAID on CPU \(Intel® VROC\) User Guide for Linux*](#) and [Intel® Virtual RAID on CPU \(Intel® VROC\) Release Notes for Linux*](#).

1.2 Reference OEM Platform Documentation

Refer to your OEM for a full list of available feature sets. If any of the information in this document conflicts with the support information provided by the platform OEM, the platform documentation and configurations take precedence.

Customers should always contact the place of purchase or system/software manufacturer with support questions about their specific hardware or software configuration.

2 Managing Intel® VROC RAID

2.1 Configure /etc/mdadm.conf

`/etc/mdadm.conf` is the configuration file of Intel® VROC for RAID configuration in RHEL. This configuration file provides rich parameters and options for system administrator to customize and manage Intel® VROC RAID volumes based on the needs.

The followings are the best practice actions the system administrator should take after creating the Intel® VROC RAID volumes.

1. Retrieve the information of the RAID volumes and store it in the configuration file:

```
# mdadm -Ebs > /etc/mdadm.conf
```

Note: The RAID array information in the configuration file should be updated after creating or removing a RAID volume in the system.

2. Add the global hot-spare policy:

```
# echo "POLICY domain=DOMAIN path=* metadata=ismm action=spare-same-slot" >> /etc/mdadm.conf
```

3. Add the auto-rebuilding policy for scenarios when replacing the failed RAID member drive with a brand-new drive:

```
# mdadm --udev-rules > /etc/udev/rules.d/65-md-bare.rules  
# udevadm control --reload
```

4. Add an email address for the Intel® VROC alert notifications:

```
# echo "MAILADDR <admin@email_address>" >> /etc/mdadm.conf
```

5. Add a user-defined program to handle alerts/events of Intel® VROC RAID volumes:

```
# echo "PROGRAM /usr/sbin/$program_name" >> /etc/mdadm.conf
```

Note: Reference next chapter for the sample code.

6. Restart the Intel® VROC RAID monitor service:

```
# systemctl restart mdmonitor.service
```

7. [Optional] Configure the RAID resync speed upper limit to a higher value.

```
# sysctl -w dev.raid.speed_limit_max=5000000
```

Note: The increase of maximum resync speed may impact the application's I/O performance during RAID rebuilding/recovery process.

2.2 Enable LEDMON Service

LEDMON service is a daemon for LED management of Intel® VROC RAID in RHEL. It is by default disabled in a newly installed RHEL operating system. System administrators should use the below command to enable it for every OS booting and activate it right away.

```
# systemctl enable --now ledmon.service
```

Use the below command to check the status of the LEDMON service:

```
# systemctl status ledmon.service
```

The following is the example of the expected result, which shows the LEDMON service is enabled and actively running.

```
● ledmon.service - Enclosure LED Utilities
   Loaded: loaded (/usr/lib/systemd/system/ledmon.service; enabled; vendor preset:
disabled)
   Active: active (running) since Mon 2023-02-20 20:11:13 EST; 2s ago
     Main PID: 18387 (ledmon)
        Tasks: 1 (limit: 409308)
       Memory: 724.0K
      CGroup: /system.slice/ledmon.service
             └─18387 /usr/sbin/ledmon -foreground

Feb 20 20:11:13 localhost systemd[1]: Started Enclosure LED Utilities.
```

2.3 Enable Postfix Service

Postfix is the email service in RHEL. The email notification of Intel® VROC alerts relies on the postfix email service. The postfix service is by default disabled in a newly installed RHEL operating system. System administrators should use the below command to enable it for every OS booting and activate it right away.

```
# systemctl enable --now postfix.service
```

Use the below command to check the status of the postfix service:

```
# systemctl status postfix.service
```

The following is an example of the expected result, which shows the postfix service is enabled and actively running:

```
● postfix.service - Postfix Mail Transport Agent
   Loaded: loaded (/usr/lib/systemd/system/postfix.service; enabled; vendor preset:
disabled)
   Active: active (running) since Thu 2023-03-02 03:58:18 EST; 9min ago
     Process: 1232347 ExecStart=/usr/sbin/postfix start (code=exited, status=0/SUCCESS)
     Process: 1232344 ExecStartPre=/usr/libexec/postfix/chroot-update (code=exited,
status=0/SUCCESS)
     Process: 1232339 ExecStartPre=/usr/libexec/postfix/aliasesdb (code=exited,
status=0/SUCCESS)
     Main PID: 1232415 (master)
        Tasks: 4 (limit: 409308)
       Memory: 6.8M
      CGroup: /system.slice/postfix.service
             └─1232415 /usr/libexec/postfix/master -w
                 └─1232416 pickup -l -t unix -u
                     └─1232417 qmgr -l -t unix -u
                         └─1232422 tlsmgr -l -t unix -u
```

After activating the postfix service, along with the email address already configured in the `/etc/mdadm.conf` file as shown in [section 2.1 step 4](#), the system administrator will receive email notifications when Intel® VROC alerts occur. The following is the example of an email notification of a Fail event of Intel® VROC RAID, which indicates `/dev/nvme7n1` is failed and `/dev/md/vol10` is degraded.

Figure 2-1. Intel® VROC Email Notification Example


2.4 Configure LVM on Intel® VROC RAID

Logical Volume Manager (LVM) is a tool for logical volume management in Red Hat Enterprise Linux*. Following the same method and command lines of LVM logical volume management in RHEL, Intel® VROC RAID supports LVM logical volume management on top of it. For details of LVM configuring and managing logical volumes in RHEL, please reference the latest [Red Hat Enterprise Linux Administration Guide](#).

The following examples show the steps and command lines of creating logical volumes on top of an Intel® VROC RAID volume.

1. Creating the Physical Volume (PV) on an Intel® VROC RAID volume (`/dev/md126`):

```
# pvcreate /dev/md126
```

2. Creating the Volume Group (VG):

```
# vgcreate myTestVG /dev/md126
```

3. Creating the Logical Volumes (LV):

```
# lvcreate -L 10G -n cache myTestVG
# lvcreate -L 200G -n data myTestVG
```

After LVM is created on the Intel® VROC RAID, users should follow the below steps to correctly configure the LVM configuration file, in order to make LVM works properly.

1. Editing `/etc/lvm/lvm.conf` file with following changes:

```
fw_raid_component_detection = 1
external_device_info_source = "udev"
```

2. Rebuilding the `initramfs`:

```
# dracut --force --verbose
```



3 Monitoring Intel® VROC RAID Alerts in RHEL

Intel® VROC is running a monitoring service in Red Hat Enterprise Linux* for managing Intel® VROC RAID arrays. That service is monitoring all the Intel® VROC RAID arrays and report alert events through email notification when it is correctly configured, as well as sending error or warning messages to the system log file.

In addition, the Intel® VROC Linux* monitoring service allows a customized program to be hooked for handling any alerts or events the service reports. The following sections will share the Intel® VROC RAID alerts as well as the guide to develop your customized program to handle them.

3.1 Intel® VROC RAID Alerts in RHEL

Intel® VROC reports follow RAID alerts through the monitoring service in RHEL. Users can develop their own program to hook the monitoring service in order to receive and handle those RAID alerts.

Table 3-1. Intel® VROC RAID Alerts in RHEL

VROC Alert/Event	Severity	Description
Fail	Critical	A member drive of Intel® VROC RAID is faulty.
FailSpare	Critical	The spare drive which was being rebuilt to has failed.
DeviceDisappeared	Critical	A RAID volume is disappeared (or removed).
DegradedArray	Critical	A newly noticed RAID array appears to be degraded.
RebuildStarted	Warning	A degraded RAID started rebuilding (recovery).
RebuildNN	Warning	Notification of rebuilding progress, NN is two-digit number (e.g., 20, 40, ...) which indicates the rebuild has passed that many percent of the total.
RebuildFinished	Warning	The rebuilding of a degraded RAID is completed or aborted.
SparesMissing	Warning	One or more spare drives defined in the <code>mdadm.conf</code> file is disappeared (or removed).
SpareActive	Information	The spare drive which was being rebuilt to has been successfully rebuilt and made active.
NewArray	Information	A new RAID array has been detected.
MoveSpare	Information	A spare drive has been moved from one array to another.

3.2 Develop a Program to Handle Intel® VROC Alerts

The Intel® VROC RAID monitoring service allows any user-defined program to be hooked to receive and handle the RAID alerts. System administrator should configure the `/etc/mdadm.conf` file to add the user-defined program which will be called by the monitoring service whenever an alert is detected. Refer to [Section 2.1 step 5](#) for the details of how to add a user-defined program.

The Intel® VROC RAID monitoring service will call the user-defined program and pass two or three parameters to it when an alert occurs. The first parameter is the event name. The second parameter is the RAID volume device name. There will be the third parameter when the event is related to the spare device or RAID member device.

The following is the bash script example of a user-defined program, which shows how to handle the alerts Intel® VROC reports. In this example, the program is just printing the event messages to the `/tmp/vroc_alerts.log` file when it receives an event. Users can develop their own method of handling those Intel® VROC alerts.

```
#!/bin/bash

event=$1
md_device=$2
device=$3

case $event in
  DegradedArray)
    msg="$md_device is running in the Degraded MODE"
    ;;
  DeviceDisappeared)
    msg="$md_device has disappeared"
    ;;
  Fail)
    msg="$md_device had a failed member device: $device"
    ;;
  FailSpare)
    msg="$md_device: Spare device ($device) FAIL during rebuild"
    ;;
  RebuildStarted)
    msg="Recovery/Rebuilding of $md_device has started"
    ;;
  Rebuild??)
    msg="$md_device REBUILD is now $(echo $event|sed 's/Rebuild//')% complete"
    ;;
  RebuildFinished)
    msg="Rebuild of $md_device is completed or aborted"
    ;;
  SpareActive)
    msg="$device has become an ACTIVE COMPONENT of $md_device"
    ;;
  NewArray)
    msg="$md_device has been detected"
    ;;
  MoveSpare)
    msg="SPARE device $device has been MOVED to a new array :$md_device"
    ;;
  SparesMissing)
    msg="$md_device is MISSING one or more SPARE devices"
    ;;
  TestMessage)
    msg="TEST MESSAGE generated for $md_device"
    ;;
esac

# In this example, we just send the event message to the tmp log.
echo "[$(date -u)] $msg" >> /tmp/vroc_alerts.log
```

System administrators can place the user-defined program into the `/usr/sbin` directory. For example: `/usr/sbin/vroc_linux_events_handler.sh`.

Then, edit the `/etc/mdadm.conf` file and add following line:

```
PROGRAM /usr/sbin/vroc_linux_events_handler.sh
```

Here is the example of the `mdadm.conf` file with the user-defined program added:

Monitoring Intel® VROC RAID Alerts in RHEL

```
# cat /etc/mdadm.conf
ARRAY metadata=imsm UUID=f69f9275:68fce440:3420da7a:48e2a723
ARRAY /dev/md/vol0 container=f69f9275:68fce440:3420da7a:48e2a723 member=0
UUID=06c1975e:2c160226:ef62cbc6:b42e4570
POLICY domain=DOMAIN path=* metadata=imsm action=spare-same-slot
PROGRAM /usr/sbin/vroc_linux_events_handler.sh
```

The following are the messages from the above example program printed to the log file when receiving the Intel® VROC alerts.

```
# cat /tmp/vroc_alerts.log
[Tue Feb 21 01:59:28 UTC 2023] Rebuild of /dev/md/vol0 is completed or aborted
[Tue Feb 21 01:59:28 UTC 2023] /dev/md/vol0 has disappeared
[Tue Feb 21 02:13:47 UTC 2023] /dev/md/vol0 REBUILD is now 21% complete
[Tue Feb 21 02:29:53 UTC 2023] /dev/md/vol0 REBUILD is now 40% complete
[Tue Feb 21 02:43:53 UTC 2023] /dev/md/vol0 REBUILD is now 60% complete
[Tue Feb 21 02:56:54 UTC 2023] /dev/md/vol0 REBUILD is now 80% complete
[Tue Feb 21 03:10:08 UTC 2023] Rebuild of /dev/md/vol0 is completed or aborted
[Wed Feb 22 02:44:28 UTC 2023] /dev/md/vol0 had a failed member device: /dev/nvme7n1
[Wed Feb 22 02:47:01 UTC 2023] Recovery/Rebuilding of /dev/md/vol0 has started
[Wed Feb 22 02:57:35 UTC 2023] Rebuild of /dev/md/vol0 is completed or aborted
```

