



Using Red Hat Enterprise Linux 7.5 with NetApp ONTAP

ONTAP SAN Host

Amanda Stroman, Sean Daffy, Madhulika Kola, Reenu Jojo
March 12, 2021

Table of Contents

- Using Red Hat Enterprise Linux 7.5 with NetApp ONTAP 1
- Installing the Linux Unified Host Utilities 1
- SAN Toolkit 1
- SAN Booting 2
- Multipathing 2
- Recommended Settings 3
- Known Problems and Limitations 5
- Release Notes 9

Using Red Hat Enterprise Linux 7.5 with NetApp ONTAP

Installing the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the [NetApp Support Site](#) in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the [NetApp Interoperability Matrix Tool](#) to verify which one you need.

Installing the Linux Unified Host Utilities is strongly recommended, but not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

Before you begin

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the [NetApp Support Site](#) to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
controller(7mode/E-Series)/          device      host          lun
vserver(cDOT/FlashRay)  lun-pathname filename  adapter  protocol  size
Product
-----
data_vserver            /vol/vol1/lun1  /dev/sdb  host16  FCP
120.0g  cDOT
data_vserver            /vol/vol1/lun1  /dev/sdc  host15  FCP
120.0g  cDOT
data_vserver            /vol/vol2/lun2  /dev/sdd  host16  FCP
120.0g  cDOT
data_vserver            /vol/vol2/lun2  /dev/sde  host15  FCP
120.0g  cDOT
```

SAN Booting

Before you begin

If you decide to use SAN booting, it must be supported by your configuration. You can use the [NetApp Interoperability Matrix Tool](#) to verify that your OS, HBA, HBA firmware and the HBA boot BIOS, and ONTAP version are supported.

Steps

1. Map the SAN boot LUN to the host.
2. Verify multiple paths are available.

Remember, multiple paths will only be available after the host OS is up and running on the paths.

3. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

For information on how to enable the HBA BIOS, see your vendor-specific documentation.

4. Reboot the host to verify the boot is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.5 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 7.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array Configuration

In All SAN Array (ASA) configurations, all paths to a given Logical Unit (LUN) are active and optimized. This means I/O can be served through all paths at the same time, thereby enabling better performance.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 11:0:7:1    sdfi   130:64   active ready running
  |- 11:0:9:1    sdiy   8:288    active ready running
  |- 11:0:10:1   sdml   69:464   active ready running
  |- 11:0:11:1   sdpt   131:304  active ready running
```



Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Non-ASA Configuration

For non-ASA configuration there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb 8:i6 active ready running
  |- 12:0:0:0 sdz 65:144 active ready running
```



Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file using the command:

```
touch /etc/multipath.conf
```

The first time you create this file, you might need to enable and start the multipath services.

```
# systemctl enable multipathd
# systemctl start multipathd
```

There is no requirement to add anything directly to `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.

You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices.

Replace the `<DevId>` with the WWID string of the device you want to exclude. Use the following command to determine the WWID:

```

blacklist {
    wwid <DevId>
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9] *"
    devnode "^hd[a-z]"
    devnode "^cciss.*"
}

```

Example

In this example, `sda` is the local SCSI disk that we need to blacklist.

Steps

1. Run the following command to determine the WWID:

```

# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833

```

2. Add this WWID to the blacklist stanza in the `/etc/multipath.conf`:

```

blacklist {
    wwid 360030057024d0730239134810c0cb833
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9] *"
    devnode "^hd[a-z]"
    devnode "^cciss.*"
}

```

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that may be overriding default settings.

The table below shows the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in `multipath.conf` that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs may not work as expected. These defaults should only be overridden in consultation with NetApp and/or OS vendor and only when the impact is fully understood.

| Parameter | Setting |
|--------------------------------|---|
| <code>detect_prio</code> | yes |
| <code>dev_loss_tmo</code> | "infinity" |
| <code>failback</code> | immediate |
| <code>fast_io_fail_tmo</code> | 5 |
| <code>features</code> | "3 queue_if_no_path pg_init_retries 50" |
| <code>flush_on_last_del</code> | "yes" |
| <code>hardware_handler</code> | "0" |

| Parameter | Setting |
|----------------------------|------------------|
| no_path_retry | queue |
| path_checker | "tur" |
| path_grouping_policy | "group_by_prio" |
| path_selector | "service-time 0" |
| polling_interval | 5 |
| prio | "ontap" |
| product | LUN.* |
| retain_attached_hw_handler | yes |
| rr_weight | "uniform" |
| user_friendly_names | no |
| vendor | NETAPP |

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
    path_checker      readsector0
    no_path_retry     fail
}

devices {
    device {
        vendor        "NETAPP  "
        product        "LUN.*"
        no_path_retry  queue
        path_checker   tur
    }
}
```

KVM Settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known Problems and Limitations

| NetApp Bug ID | Title | Description | Bugzilla ID |
|---------------|--|--|-------------|
| 1139053 | Kernel disruption occurs on RHEL7.5 with QLogic QLE2672 16GB FC during storage failover operations | <p>During storage failover operations on the RHEL7U5 kernel with QLogic QLE2672 16GB fibre channel host bus adapter, the kernel disruption occurs due to a panic in the kernel. The kernel panic causes RHEL 7.5 to reboot, which leads to an application disruption. The kernel panic generates the vmcore file under the /var/crash/directory if kdump is configured. The vmcore file is used to understand the cause of the failure. In this case, the panic was observed in the "get_next_timer_interrupt+440" module which is logged in the vmcore file with the following string: "[exception RIP: get_next_timer_interrupt+440]" After the kernel disruption, you can recover the operating system by rebooting the host operating system and restarting the application as required.</p> | 1542564 |

| NetApp Bug ID | Title | Description | Bugzilla ID |
|-------------------------|--|--|-------------------------|
| 1138536 | Kernel disruption occurs on RHEL7U5 with QLogic QLE2742 32GB FC during storage failover operations | During storage failover operations on the Red Hat Enterprise Linux (RHEL) RHEL7U5 kernel with QLogic QLE2742 HBA, kernel disruption occurs due to a panic in the kernel. The kernel panic leads to a reboot of the operating system, causing an application disruption. The kernel panic generates the vmcore file under the /var/crash/ directory if kdump is configured. When the kernel panics, you can use the vmcore file to investigate the reason for the failure. The following example shows a panic in the bget_next_timer_interrupt+440b module. The panic is logged in the vmcore file with the following string: "[exception RIP: get_next_timer_interrupt+440]" You can recover the operating system by rebooting the host OS and restarting the application as required. | 1541972 |

| NetApp Bug ID | Title | Description | Bugzilla ID |
|---------------|---|--|-------------|
| 1148090 | Kernel disruption occurs on RHEL 7.5 with QLogic QLE2742 32GB FC HBA during storage failover operations | During storage failover operations on the Red Hat Enterprise Linux (RHEL) 7.5 kernel with a QLogic QLE2742 Fibre Channel (FC) host bus adapter (HBA), a kernel disruption occurs due to a panic in the kernel. The kernel panic causes RHEL 7.5 to reboot, which leads to an application disruption. If the kdump mechanism is enabled, the kernel panic generates a vmcore file located in the /var/crash/ directory. You can analyze the vmcore file to determine the cause of the panic. In this instance, when storage failover with the QLogic QLE2742 HBA event occurs, the "native_queued_spin_lock_slowpath+464" module is affected. You can locate the event in the vmcore file by finding the following string: "[exception RIP: native_queued_spin_lock_slowpath+464]" After the kernel disruption, you can reboot the Host OS and recover the operating system, and then you can restart the applications as required. | 1559050 |

| NetApp Bug ID | Title | Description | Bugzilla ID |
|-------------------------|--|--|-------------------------|
| 1146898 | Kernel disruption occurs on RHEL 7.5 with Emulex HBAs during storage failover operations | <p>During storage failover operations on a Red Hat Enterprise Linux (RHEL) 7.5 system with Emulex LPe32002-M2 32-GB FC host bus adapters (HBAs), a disruption in the kernel occurs. The kernel disruption causes a reboot of the operating system, which in turn causes an application disruption. If you configure kdump, the kernel disruption generates the vmcore file under the <code>/var/crash/</code> directory. You can use the vmcore file to determine the cause of the failure. In the following example, you can see the disruption in the</p> <pre>"lpfc_hba_clean_txcmplq+368" module. This disruption is logged in the vmcore file with the following string: "[exception RIP: lpfc_hba_clean_txcmplq+368]"</pre> <p>After the kernel disruption, reboot the host OS to recover the operating system. Restart the application as required.</p> | 1554777 |

Release Notes

ASM Mirroring

ASM mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate fail group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM does not mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See [Oracle Databases on ONTAP](#) for further information.

Copyright Information

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.