

Unbreakable Enterprise Kernel

Release Notes for Unbreakable Enterprise Kernel Release 4 Update 5



E88575-08
January 2023



Unbreakable Enterprise Kernel Release Notes for Unbreakable Enterprise Kernel Release 4 Update 5,
E88575-08

Copyright © 2020, 2023, Oracle and/or its affiliates.

Contents

Preface

Conventions	v
Documentation Accessibility	v
Access to Oracle Support for Accessibility	v
Diversity and Inclusion	v

1 New Features and Changes

Notable Changes	1-1
DTrace Improvements	1-2
Xen Hypervisor Improvements	1-3
File System Improvements	1-3
Driver Updates	1-4
Additional Notes For Driver Updates	1-5
Technology Preview	1-5
Compatibility	1-6
Header Packages for Development	1-6

2 Fixed Issues

3 Known Issues

btrfs, ext4 and xfs: Kernel panic when freeze and unfreeze operations are performed in multiple threads	3-1
btrfs	3-1
ext4	3-2
xfs	3-2
DIF/DIX is not supported for ext file systems	3-4
Console appears to hang when booting	3-4
Docker	3-5
DTrace	3-5
Error, some other host already uses address xxx.xxx.xxx.xxx	3-6

ifup-ib: line 357: /sys/class/net/ib0/acl_enabled: Permission denied error	3-6
Increased dom0 memory requirement when using Mellanox® HCAs on Oracle VM Server LXC	3-6
MSI-X interrupt allocation fails during maximum number of ixgbe/ixgbevf Virtual Function creation	3-7
NVMe devices not found under the /dev directory after PCI rescan	3-7
OFED iSER target login fails from an initiator on Oracle Linux 6	3-7
Open File Description (OFD) locks are not supported on NFSv4 mounts	3-8
SDP performance degradation	3-8
Shared Receive Queue (SRQ) is an experimental feature for RDS and is disabled by default	3-8
Unloading or removing the rds_rdma module is unsupported	3-8

4 Installation and Availability

Installation Overview	4-1
Subscribing to ULN Channels	4-1
Enabling Access to Oracle Yum Channels	4-2
Upgrading Your System	4-4
Installing the Oracle-Supported OFED Packages	4-5

Preface

[Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 5 \(4.1.12-103\)](#) provides a summary of the new features, changes, and known issues in the Unbreakable Enterprise Kernel Release 4 Update 5.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <https://www.oracle.com/corporate/accessibility/>.

For information about the accessibility of the Oracle Help Center, see the Oracle Accessibility Conformance Report at <https://www.oracle.com/corporate/accessibility/templates/t2-11535.html>.

Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry

standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

1

New Features and Changes

The Unbreakable Enterprise Kernel Release 4 (UEK R4) is Oracle's fourth major release of its heavily tested and optimized operating system kernel for Oracle Linux 6 Update 7 or later, and Oracle Linux 7 Update 1 or later, on the x86-64 architecture. It is based on the mainline Linux kernel version 4.1.12.

UEK R4 update 5 uses the 4.1.12-103.3.8 version and build of the UEK R4 kernel, which includes security and bug fixes, as well as driver updates. This kernel has been tested within environments running the latest available Oracle Linux releases: Oracle Linux 6 Update 8, Oracle Linux 6 Update 9, Oracle Linux 7 Update 3, and Oracle Linux 7 Update 4.

Oracle actively monitors upstream check-ins and applies critical bug and security fixes to UEK R4.

UEK R4 uses the same versioning model as the mainline Linux kernel version. It is possible that some applications might not understand the 4.1 versioning scheme. However, regular Linux applications are usually neither aware of nor affected by Linux kernel version numbers.

Notable Changes

- **Intel E1000 drivers updated**

This update includes a large number of patches to the Intel E1000 driver modules to support a wider range of devices, to implement performance enhancements and to resolve a number of bugs.

- **Intel i40e module updated**

This update includes a large number of upstream patches for the Intel i40e driver module to resolve a number of bugs and to bring this driver module into alignment with upstream functionality.

- **nvme module updated**

This update brings major changes to the `nvme` kernel module in the form of over a hundred upstream patches that ensure that the driver matches upstream functionality. Fixes were also included to better handle controller reset so that the request queue is properly frozen to prevent new requests from entering the queue.

- **forcedeth module added**

This update includes the addition of the popular open-source `forcedeth` kernel module that can be used as a driver for NVIDIA network interface cards.

- **smartpqi driver updated**

The Microsemi Smart Family Controller Driver (`smartpqi`) module has been updated for a large number of upstream patches. This driver enables next generation SCSI devices that implement the PQI queuing model, such as the Adaptec Series 9 controllers. The driver replaces the `aacraid` driver for these devices. Patches have been applied to bring this version of the driver up to the most current level for compatibility with this kernel release.

- **Support for IP packet duplication in `nf_tables`**

This update includes the addition of required modules to support packet duplication of IPv4 and IPv6 packets within `nf_tables`. This change provides support to clone packets to a given destination.
- **Support for netlink for OFED**

This update includes patches to the InfiniBand core code to add support for netlink from kernel to user space for the OpenFabrics Enterprise Distribution (OFED). Also included is a fix that makes the kernel OFED stack interrogate the InfiniBand address and route resolution service framework in user space to resolve path records via the netlink interface.

DTrace Improvements

A number of bug fixes and enhancements including module and utility updates are included for DTrace on UEK R4u5, bringing the current version to 0.6.1. Other notable changes include the following:

- **General bug fixes**

Numerous bug fixes have been applied to provide greater stability and better performance.
- **Function Boundary Tracing (FBT) improvements**

Fixes were applied to code for detecting whether a kernel function could be probed using FBT upon entry to make it less restrictive and to improve its ability to function on Oracle Linux 7 systems. A patch was also applied to code that handled passing the offset as an argument to FBT return probes to help ensure that FBT return probe arguments have the correct value.
- **I/O Providers**

I/O provider coverage has been expanded to add probes for block file systems based on the common block layer. Additionally, probes have been added to support the tracing of I/O read and write performed on NFS. SDT probes for the TCP and UDP protocols have also been added. These include `tcp:::send`, `tcp:::receive`, `tcp:::state-change`, `tcp:::connect-request`, `tcp:::connect-refused`, `tcp:::connect-established`, `tcp:::accept-refused`, `tcp:::accept-established`, `udp:::send` and `udp:::receive`.
- **Automatic loading of DTrace Modules**

DTrace kernel modules are automatically loaded for providers. This is achieved by pre-populating the `/etc/dtrace-modules` file with a list of default provider modules. This is implemented when `dtrace-utils` is installed on the system.
- **Improvements and fixes for CTF**

Changes have been applied to add support for bitfields in `libdtrace-ctf`, including changes to `dwarf2ctf` to emit the correct CTF, where bitfields are represented. This change enables DTrace to see the values of bitfields within structures and unions. A fix was also applied so that the Kernel CTF could start to include kernel static variables. This change also introduces a module-scope `vars_seen` hashtable into the deduplicator state, which identifies whether a variable is or is not static. Variables for which multiple static definitions exist are blacklisted, along with matching static and external variables that are declared for the same module, to avoid confusion over duplicate variable names.

Xen Hypervisor Improvements

This update includes several patches for Xen that are backported from changes in the Linux 4.9 and 4.10 kernel releases. Many of these patches improve performance, error handling and provide general bug fixes. Other notable fixes and improvements include the following:

- A fix has been applied to the Xen blkfront driver to resolve a race condition that could cause I/O to hang when the ring buffer was full and the hardware queue was stopped.
- A fix has been applied to the Xen blkback driver to update the xenstore node hotplug status to provide notification of a busy block device that cannot be detached because it is mounted inside a guest. This can prevent block-detach hangs when detaching a block device that is in use.
- Performance enhancements were introduced in the Xen netback and netfront drivers to improve network throughput on chatty network activities that stress IRQ processing. This can enhance database performance in a virtual machine to more closely match the performance on bare metal. Other netback and netfront driver fixes were also applied to improve handling for network stress
- An update of the Xen PCIfront driver is available, which supports vNUMA functionality so that a guest can mark specific PCI devices as belonging to NUMA nodes.

File System Improvements

The following improvements were implemented on the file systems supported by the current version of UEK:

btrfs Updates

- Patches were applied to support the ability to delete a device by specifying device ID. This feature enables functionality in the latest `btrfs-progs` package that is available for Oracle Linux 7 update 4 and later.
- Upstream patches were applied to better handle dedupe which was changing the `mtime` on files. This caused some applications, such as `rsync`, to incorrectly assert that file contents had changed. The patch helps to ensure that dedupe is transparent to the user.
- A fix has been included for an issue that caused a crash when removing a file where `flush_space()` did not overwrite the return value for successful chunk allocation.

XFS Updates

- Upstream patches have been applied to provide for the ability to reconfigure XFS error handling. This includes changes to behavior on `ENOSPC` and error handling during unmount.
- Upstream patches were applied to resolve a very rare issue that caused XFS to emit warning messages when under extremely high I/O load when available memory is limited.
- Upstream patches were applied to resolve some XFS warning issues when clearing SUID and SGID bits and to better handle the removal of security labels.

OCFS2 Updates

Various patches were applied to improve cluster locking and tracking of locks to prevent deadlocks.

Driver Updates

The Unbreakable Enterprise Kernel supports a wide range of hardware and devices. In close cooperation with hardware and storage vendors, several device drivers have been updated or added by Oracle, per the information in the following table.

Driver	Version	Description
aacraid	1.2-1[41010]-ms + Patches	Adaptec Advanced Raid Products
be2net	11.4.0.0 + Patches	Broadcom/Emulex OneConnect 10Gbps NIC Driver
bnx2x	1.713.10 + Patches	QLogic BCM57710/57711/57711E/57712/57712_MF/57800/57800_MF/57810/57810_MF/57840/57840_MF Driver
bnxt_en	1.7.0 + Upstream Patches	Broadcom/Emulex BCM573xx NIC Driver
e1000	7.3.21-k8-NAPI + Patches	Intel(R) PRO/1000 Network Driver
e1000e	3.2.6-k + Patches	Intel(R) PRO/1000 Network Driver
fnic	1.6.0.34	Cisco FCoE HBA Driver
forcedeth	Upstream patches	Reverse Engineered nForce ethernet driver
i40e	2.1.7-k	Intel® Ethernet Connection XL710 Network Driver
i40evf	2.1.7-k	Intel® XL710 X710 Virtual Function Network Driver
igb	5.4.0-k	Intel® Gigabit Ethernet Network Driver
igbvf	2.4.0-k	Intel® Gigabit Virtual Function Driver
ixgbe	5.0.0-k	Intel® 10 Gigabit PCI Express Network Driver
ixgbev	3.2.2-k + Upstream Patches	Intel® 82599 Virtual Function Driver
lpfc	11.4.0.2	Broadcom/Emulex LightPulse Fibre Channel SCSI Driver
megaraid_sas	07.701.17.00-rc1	Avago MegaRAID SAS Driver
mgag200	Upstream Patches	Matrox GP200e4 Device Driver
mpt3sas	15.100.00.00	LSI MPT Fusion SAS 3.0 Device Driver
nf_dup_ipv4	1.0 + Upstream Patches	IPv4/IPv6 packet duplication support for nf_tables
nf_dup_ipv6		
nft_dup_ipv4		
nft_dup_ipv6		
nvme-core	1.0 + Patches	NVM Express (NVMe) Core Driver

Driver	Version	Description
qed	8.10.10.21	QLogic FastLinQ 4xxxx Ethernet Driver and Core Module
qede		
smartpqi	1.0.4-100	Microsemi Smart Family Controller Driver
xen-blkback	Upstream Patches and Bug Fixes	Xen Virtual Block Device and Xen Virtual Network Device drivers
xen-blkfront		
xen-netback		
xen-netfront		
xen-pcifront		
xscore	6.0.r8044	Xsigo Oracle Virtual Network Core Driver
xsvhba	6.0.r8044	Xsigo Oracle Virtual Network VHBA Driver
xsvnic	6.0.r8044	Xsigo Oracle Virtual Network XSVNIC Network Driver
xve	6.0.r8044	Xsigo Oracle Virtual Network Virtual Ethernet Driver

Some InfiniBand drivers are patched with bug fixes and minor enhancements in this release.

Additional Notes For Driver Updates

The following notes are included at the request of a vendor for the listed driver:

- `lpfc`: Locked optics support is enabled for LPE32000 HBAs and all variant HBAs of this architecture. With this capability, these HBAs will detect and enable both Avago or Emulex certified SFP and QSFP optics.

For driver rev 11.0.0.13 and higher unqualified optics will be disabled, the link will not come up, an error message is written to the log file and the `lpfc` driver will display this message:

```
3176 Port Name [wwpn] Unqualified optics - Replace with Avago optics for Warranty and Technical support
```

- `mpt3sas`: As of UEK R4 update 4, the `mpt2sas` driver has been merged with the `mpt3sas` driver to provide a single driver module that supports both SAS 2.0 and SAS 3.0 HBAs. Changes have been applied to `dracut` to correctly handle the module aliases for the migration to a single driver module.

Technology Preview

The following features that are included in the Unbreakable Enterprise Kernel Release 4 are still under development, but are made available for testing and evaluation purposes:

- **DCTCP (Data Center TCP)**

DCTCP enhances congestion control by making use of the Explicit Congestion Notification (ECN) feature of state-of-the-art network switches. DCTCP reduces buffer

occupancy and improves throughput by allowing a system to react more intelligently to congestion than is possible using TCP.

- **DRBD (Distributed Replicated Block Device)**
A shared-nothing, synchronously replicated block device (*RAID1 over network*), designed to serve as a building block for high availability (HA) clusters. It requires a cluster manager (for example, pacemaker) for automatic failover.
- **Kernel module signing facility**
Applies cryptographic signature checking to modules on module load, checking the signature against a ring of public keys compiled into the kernel. GPG is used to do the cryptographic work and determines the format of the signature and key data.
- **NFS over RDMA interoperation with ZFS and Oracle Solaris**
NFS over RDMA does not yet fully interoperate with ZFS and Oracle Solaris. NFS over RDMA for NFS versions 3 and 4 is supported for Oracle Linux systems using the Oracle InfiniBand stack and is more efficient than using NFS with TCP over IPoIB. Currently, only the Mellanox® ConnectX-2 and ConnectX-3 Host Channel Adapters (HCAs) pass the full Connectathon NFS test suite and are supported.
- **Server-side parallel NFS**
Server-side parallel NFS (pNFS) improves the scalability and performance of an NFS server by making file metadata and data available on separate paths.

Compatibility

Oracle Linux maintains user-space compatibility with Red Hat Enterprise Linux (RHEL), which is independent of the kernel version running underneath the operating system. Existing applications in user space will continue to run unmodified on the Unbreakable Enterprise Kernel Release 4 and no re-certifications are needed for RHEL certified applications.

To minimize impact on interoperability during releases, the Oracle Linux team works closely with third-party vendors whose hardware and software have dependencies on kernel modules. The kernel ABI for UEK R4 will remain unchanged in all subsequent updates to the initial release. In this release, there are changes to the kernel ABI relative to UEK R3 that require recompilation of third-party kernel modules on the system. Before installing UEK R4, verify its support status with your application vendor.

Header Packages for Development

As of UEK-3.8-QU2, the `kernel-uek-headers` package is no longer built and distributed. There are three kernel packages that might be useful for development purposes. The `kernel-headers` package forms part of the API for user space programs. The `kernel-devel` package is used for standard RHCK development and module compilation. The `kernel-uek-devel` package is used for UEK development and module compilation. Neither the `kernel-uek-headers`, nor the `kernel-headers` packages, are needed for kernel development.

The `kernel-headers` package provides the C header files that specify the interface between user-space binaries or libraries and UEK or RHCK. These header files define

the structures and constants that you need to build most standard programs or to rebuild the `glibc` package.

The `kernel-devel` and `kernel-uek-devel` packages provide the kernel headers and makefiles that you need to build modules against UEK and RHCK.

To install the packages required to build modules against UEK and the C header files for both UEK and RHCK:

```
# yum install kernel-uek-devel-`uname -r` kernel-headers
```

2

Fixed Issues

This chapter describes fixed issues for the Unbreakable Enterprise Kernel Release 4.

! Important:

Run the `yum update` command regularly to ensure that the latest bug fixes and security errata are installed on your system.

- **aacraid driver issue that caused a kernel panic fixed**
An issue that caused a kernel panic when the Adaptec RAID controller was loaded and the system attempted to perform an asynchronous scan of the SCSI disks is fixed by an upstream patch to the `aacraid` driver. The fix initializes the SCSI shared tag map when the driver is loaded.
- **i40e driver module fixed for DHCP requests when bridged with a virtio network interface**
An issue in the `i40e` driver module that caused DHCP to fail for a KVM guest when bridged with a `virtio` network interface has been fixed.
- **lpfc driver module fixed for an issue connecting to targets that send PRLI under P2P mode**
An issue in the `lpfc` driver module has been patched to enable HBAs to connect to targets that send a Process Login (PRLI) in P2P mode. The issue prevented the resending of the PRLI and resulted in the connection being rejected.
- **mgag200 driver issue that caused random system hangs fixed**
An issue in the Matrox GP200e4 driver that caused random system hangs is fixed by an upstream patch to the driver.
- **xfs: Directory readahead completions fixed for unmount issue**
A bug that caused a system to hang if an XFS formatted file system is unmounted suddenly after mount has been fixed. A mechanism has been added to track all in-flight, asynchronous buffers using per-cpu counters and to explicitly exclude buffers from the I/O accounting.
- **Memory management code for virtual memory area adjustment fixed for use-after-free bug**
A use-after-free issue in the Linux memory management code that resulted when memory allocation failed in `vma_adjust()` is fixed. This bug was replicated when running `fdisk -l`, which caused a kernel panic.
- **RDS patched to fix memory allocation for scatter-gather based on message sizes**

Reliable Datagram Sockets (RDS) protocol support has been patched to fix a bug that was resulting in page allocation failures, by reducing the memory footprint in `rds_sendmsg`. Since InfiniBand supports scatter-gather in hardware, RDS fragment sizes don't need to match page size, which was causing unnecessary pressure on the memory allocation system. The fix only allocates the buffer according to `RDS_MAX_FRAG_SIZE`.

- **RDS patched to fix a portstate transition issue that caused ports to get stuck in the `INIT` state**

RDS is patched for a rare race condition between RDS module initialization and links coming up late, that resulted in the state not getting initialized correctly. In most cases this problem would not occur but could trigger if a host and switch were rebooted at the same time. The issue appeared after reboot as one or more RDS InfiniBand ports never transitioned out of `INIT` state. If the initialization did not complete correctly, subsequent port failovers would fail and disrupt cluster failover and failback robustness.

- **InfiniBand code patched to resolve a race condition that occurred when accessing FMR resources**

Code for InfiniBand was patched to resolve a race condition that resulted when accessing Fast Memory Registration (FMR) resources in the InfiniBand FMR pool. This resulted in a bad syncing issue that could cause a system crash or reboot due to deadlock on a hash list lookup.

- **IPoIB patched to improve `ioctl` handling**

Code for InfiniBand was patched to improve `ioctl` handling. The code was unable to fully handle the IPoIB data structure and was triggering a warning message in `/var/log/messages`:

```
ib0: ioctl fail to copy request data
```

The patch ensures that only unsupported structures trigger the warning and only if debug is enabled.

- **Hyper-V `fcopy` processes fixed to allow copying large files from host to guest**

Patches and driver updates applied to the Hyper-V driver modules have resolved an issue that caused trouble for Oracle Linux guests, running UEK R4 and hosted on Microsoft® Windows 2012 R2 Hyper-V systems, when copying large files from the host to the guest virtual machine.

- **Hyper-V service start-up issue resolved**

Patches and driver updates applied to the Hyper-V driver modules have resolved an issue that prevented the `hypervkvpd` and `hypervvssd` services from starting after reboot. Note that in previous release notes for Oracle Linux 7, the recommended workaround for this issue was to downgrade packages and to modify `/etc/yum.conf` to exclude the Hyper-V packages from future updates. In this update release of UEK R4, the issue is resolved and you can safely upgrade Hyper-V packages to their most current version.

- **Xen problem hot adding CPU for live Guest VM fixed**

A patch was applied to the Xen event handling code to fix the potential for soft lock-ups when adding CPUs to a guest VM.

- **Xen networking fix for stalling receive during network stress**

A patch was applied to the Xen netfront driver to resolve an issue that caused virtual NICs to stall when memory usage was high and where available memory was limited. This high load and out-of-memory scenario caused the allocation of `sk_buff` in the `xennet_alloc_rx_buffers()` to fail, resulting in a stall.

- **IPv4 code fixed for an issue that crashed `dom0` running on Xen**

An upstream patch that was applied to the IPv4 code to use coalescing during defragmentation was reverted to fix an issue that caused `dom0` to crash in the Xen hypervisor when handling some network traffic. The original patch helped to reduce overhead and simplify processing for the receiving socket, however it also destroyed the geometry of the original packet fragments. In some situations, it is possible that fastpath processing can result in fragment mismatches.

3

Known Issues

This chapter describes the known issues in this update.

btrfs, ext4 and xfs: Kernel panic when freeze and unfreeze operations are performed in multiple threads

Freeze and unfreeze operations that are performed across multiple threads on any supported file system can cause the system to hang and the kernel to panic. This problem is the result of a race condition that occurs when the unfreeze operation is triggered before it is actually frozen. The resulting unlock operation attempts a write operation on a non-existent lock, resulting in the kernel panic. (Bug ID 25321899)

btrfs

The following are known `btrfs` issues:

- **Send operation causes soft lockup on large deduped file**

Using `btrfs send` on a large deduped file results in a soft lockup or out-of-memory issue. This problem occurs because the `btrfs send` operation cannot handle a large deduped file containing file extents that are all pointing to one extent, as these types of file structures create tremendous pressure for the `btrfs send` operation.

To prevent this issue from occurring, do not use `btrfs send` on systems with less than 4 GB of memory. (Bug ID 25306023)
- **Kernel oops when unmounting during a quota rescan or disable**

Operations that trigger a quota rescan or to disable the quota on a mounted file system cause a kernel oops message when attempting to unmount the file system. This can cause the system to hang. (Bug ID 22377928)
- **Kernel oops when removing shared extents using qgroup accounting**

The removal of shared extents where quota group (qgroup) accounting is used can result in a kernel oops message. This relates to an issue where inaccurate results are obtained during a back reference walk, due to missing records when adding delayed references. (Bug ID 21554517)
- **No warning when balancing file system on RAID**

The `btrfs filesystem balance` command does not warn that the RAID level can be changed under certain circumstances, and does not provide the choice of cancelling the operation. (Bug ID 16472824)
- **Disk space requirement to perform all btrfs operations**

The copy-on-write nature of `btrfs` means that every operation on the file system initially requires disk space. It is possible that you cannot execute any operation on a disk that has no space left, and even removing a file might not be possible. In the case that there is no space to store metadata, an `ENOSPC` error is returned. In this situation, run `sync`

before retrying an operation, as this can clear a background writeback that might be reserving metadata space. Another potential workaround is to add a disk or a file backed loop device using the `btrfs device add` command. The mechanism that is used to store data and metadata might lead to some confusion on the information returned by tools like `df`. Sometimes, metadata might fill all of the disk space allocated for this purpose, even while there is still space available for data. In this case, the file system is unbalanced and the problem can be resolved by performing a `btrfs fi balance` operation. See https://btrfs.wiki.kernel.org/index.php/Problem_FAQ#I_get_.22No_space_left_on_device.22_errors.2C_but_df_says_I.27ve_got_lots_of_space.

- **Double count of overwritten space in `qgroup show`**

When you overwrite data in a file, starting somewhere in the middle of the file, the overwritten space is counted twice in the space usage numbers that `btrfs qgroup show` displays. Using the `btrfs quota rescan` does not help fix this issue either. (Bug ID 16609467)

- **Sector size should match page size**

If you use the `-s` option to specify a sector size to `mkfs.btrfs` that is different from the page size, the created file system cannot be mounted. By default, the sector size is set to be the same as the page size. (Bug ID 17087232)

- **Location of `btrfs-progs` and `btrfs-progs-devel` packages**

The `btrfs-progs` and `btrfs-progs-devel` packages for use with UEK R4 are made available in the `ol6_x86_64_UEKR4` and `ol7_x86_64_UEKR4` ULN channels and the `ol6_UEKR4` and `ol7_UEKR4` channels on the Oracle Linux Yum Server. In UEK R3, these packages were made available in the `ol6_x86_64_latest` and `ol7_x86_64_latest` ULN channels and the `ol6_latest` and `ol7_latest` channels on the Oracle Linux Yum Server.

ext4

The following are known `ext4` issues:

- **System hangs when processing corrupted orphaned inode list**

If the orphaned inode list is corrupted, the inode might be processed repeatedly, resulting in a system hang. For example, if the orphaned inode list contains a reference to the bootloader inode, `ext4_iget()`, returns a bad inode, it can result in a processing loop that can hang the system. (Bug ID 24433290)

- **System hangs on unmount after an append to a file with negative `i_size`**

While it is invalid for a file system to load an inode with a negative `i_size`, it is possible to create a file like this and append to it. However, doing so causes an integer overflow in the routine's underlying writeback, resulting in the kernel locking up. (Bug ID 25565527)

xfs

The following are known `xfs` issues:

- **File system corruption occurs after direct I/O writes**

A race condition that results in post-eof blocks being used for direct I/O writes causes a corruption in the file system. If a file release occurs during a file extending direct I/O write, it is possible to mistake the post-eof blocks for speculative preallocation and incorrectly truncate them from the inode. This issue is unlikely to be reproduced in real-world workloads. (Bug ID 26128822)

- **Invalid corrupted file system error resulting from a problem with log recovery on v5 superblocks**

A problem with log recovery on v5 superblocks that causes the metadata LSN not to update for buffers that it writes out, can result in a corruption error similar to the following:

```
[1044224.901444] XFS (sdcl): Metadata corruption detected at
xfs_dir3_block_write_verify+0xfd/0x110 [xfs], block 0x1004e90
[1044224.901446] XFS (sdcl): Unmount and run xfs_repair
...
[1044224.901460] XFS (sdcl): xfs_do_force_shutdown(0x8) called from line 1249
of file fs/xfs/xfs_buf.c. Return address = 0xfffffffffa07a8910
[1044224.901462] XFS (sdcl): Corruption of in-memory data detected. Shutting
down filesystem
[1044224.901463] XFS (sdcl): Please umount the filesystem and rectify the
problem(s)
[1044224.904207] XFS (sdcl): log mount/recovery failed: error -117
[1044224.904456] XFS (sdcl): log mount failed"
```

This problem is encountered because the log attempts to replay a buffer update that is no longer valid due to subsequent replayed updates. The result is a corruption error, when in fact, the file system is fine. (Bug ID 25380003)

- **System hangs on unmount after a buffered append to a file with negative `i_size`**

While it is invalid for a file system to load an inode with a negative `i_size`, it is possible to create a file like this, and in the case where a buffer appends to it, an integer overflow in the routine's underlying writeback results in the kernel locking up. A direct append does not cause this behavior. (Bug ID 25565490)

- **System hangs during `xfs_fsr` on two-extent files with speculative preallocation**

During an `xfs_fsr` process on extents that are generated by speculative preallocation, the code that determines whether all of the extents fit inline miscalculates because the `di_nextents` call that is used does not account for these extents. This results in corruption of the in-memory inode, and ultimately the code attempts to move memory structures using incorrectly calculated ranges. This causes a kernel panic. (Bug ID 25333211)

- **XFS quotas are disabled after a read-only remount on Oracle Linux 6**

Quotas are disabled on XFS if the file system is remounted with read-only permissions on Oracle Linux 6. (Bug ID 22908906)

- **Overlay file system is unable to mount on XFS where there is no `d_type` support**

Overlay file systems rely on a feature known as `d_type` support. This feature is a field within a data structure that provides some metadata about files in a directory entry within the base file system. Overlay file systems use this field to track many file operations such as file ownership changes and whiteouts. `d_type` support can be enabled in XFS when the file system is created, by using the `-n ftype=1` option. When `d_type` support is not enabled, an overlay file system might become corrupt and behave in unexpected ways. For this reason, this update release of UEK R4 prevents the mounting of an overlay file system on an XFS base, where `d_type` support is not enabled.

The `root` partition on Oracle Linux is automatically formatted with `-n ftype=0`, where XFS is selected as the file system. Thus, for backward compatibility reasons, if you have overlay file systems in place already and these are not hosted on alternate storage, you must migrate them to a file system that is formatted with `d_type` support enabled.

To check that the XFS file system is formatted correctly:

```
# xfs_info /dev/sdb1 |grep ftype
```

Replace `/dev/sdb1` with the path to the correct storage device. If the information returned by this command includes `ftype=0`, you must migrate the overlay data held in this directory to storage that is formatted correctly.

To correctly format a new block device with the XFS file system with support for overlay file systems, do:

```
# mkfs -t xfs -n ftype=1 /dev/sdb1
```

Replace `/dev/sdb1` with the path to the correct storage device. It is essential that you use the `-n ftype=1` option when you create the file system.

If you do not have additional block storage available, it is possible to create an XFS file system image and loopback that can be mounted. For example, to create a 5 GB image file in the `root` directory, you could use the following command:

```
# mkfs.xfs -d file=1,name=/OverlayStorage,size=5g -n ftype=1
```

To temporarily mount this file, you can enter:

```
# mount -o loop -t xfs /OverlayStorage /mnt
```

Adding an entry in `/etc/fstab` to make a permanent mount for this storage, might look similar to the following:

```
/OverlayStorage /mnt xfs loop 0 0
```

This configuration can help as a temporary solution to solve upgrade issues. However, using a loopback mounted file system image as a form of permanent storage is not recommended for production environments. (Bug ID 26165630)

DIF/DIX is not supported for ext file systems

The Data Integrity Field (DIF) and Data Integrity Extension (DIX) features that have been added to the SCSI standard are dependent on a file system that is capable of correctly handling attempts by the memory management system to change data in the buffer while it is queued for a write.

The `ext2`, `ext3` and `ext4` file system drivers do not prevent pages from being modified during I/O which can cause checksum failures and a "Logical block guard check failed" error. Other file systems such as XFS are supported. (Bug ID 24361968)

Console appears to hang when booting

When booting Oracle Linux 6 on hardware with an ASPEED graphics controller, the console might appear to hang during the boot process after starting `udev`. However,

the system does boot properly and is accessible. The workaround is to add `nomodeset` as a kernel boot parameter in `/etc/grub.conf`. (Bug ID 22389972)

Docker

The following are known Docker issues:

- **Running `yum install` within a container on an overlayfs file system can fail with the following error:**

```
Rpmdb checksum is invalid: dCDPT(pkg checksums): package_name
```

This error can break Dockerfile builds but is expected behavior from the kernel and is a known issue upstream (see <https://github.com/docker/docker/issues/10180>.)

The workaround is to run `touch /var/lib/rpm/*` before installing the package.

Note that this issue is fixed in any Oracle Linux images available on the Docker Hub or Oracle Container Registry, but the issue could still be encountered when running any container based on a third-party image. (Bug ID 21804564)

- **Docker can fail where it uses the `overlay2` storage driver on XFS-formatted storage**

A kernel patch has been applied to prevent overlay mounts on XFS if the `f_type` is not set to 1. This fix resolves an issue where XFS did not properly support the whiteout features of an overlay filesystem if `d_type` support was not enabled. If the Docker Engine is already using XFS-formatted storage with the `overlay2` storage driver, an upgrade of the kernel can cause Docker to fail if the underlying XFS file system is not created with the `-n f_type=1` option enabled. The root partition on Oracle Linux 7 is automatically formatted with `-n f_type=0` where XFS is selected as the file system. Therefore, if you intend to use the `overlay2` storage driver in this environment, you must format a separate device for this purpose. (Bug ID 25995797)

- **Docker can fail where it uses the `overlay2` storage driver and SELinux is enabled**

If the Docker Engine is configured to use the `overlay2` storage driver and SELinux is enabled and set to Enforcing mode, Docker containers are unable to function properly and permissions errors are encountered. If you intend to use Docker with the `overlay2` storage driver, you must set SELinux to Permissive mode. (Bug ID 25684456)

DTrace

The following are known DTrace issues:

- Argument declarations with USDT probe definitions cannot be declared with derived types such as `enum`, `struct`, or `union`.
- The following compiler warning can be ignored for USDT probe definition arguments of type `string` (which is a D type but not a C type):

```
provider_def.h:line#: warning: parameter names (without types) in function declaration
```

- Multi-threaded processes under `ustack()`, `usym()`, `uaddr()` and `umod()`, which perform `dlopen()` in threads other than the first thread might not have accurate symbol resolution for symbols introduced by `dlopen()`. (Bug ID 20045149)

Error, some other host already uses address xxx.xxx.xxx.xxx

The following error message might be triggered in certain instances:

```
Error, some other host already uses address xxx.xxx.xxx.xxx
```

The following are the two instances in which this error message might be triggered:

- When active-bonding is enabled, and you run the `ifup ib-interface` command.
- When you run the `service rdma start` command.

You can ignore this message, as in both cases, the InfiniBand interface is brought up successfully. (Bug IDs 21052903, 26639723)

ifup-ib: line 357: /sys/class/net/ib0/acl_enabled: Permission denied error

Running `ifup ib-interface` or `service network restart` reports the following error:

```
/etc/sysconfig/network-scripts/ifup-ib: line 357: /sys/class/net/ib0/
acl_enabled: Permission denied
```

This error is reported, even though the InfiniBand interface is brought up successfully.

The workaround for this issue is to change from using the older configuration method, where you manipulate `sysfs` files to the newer `ibacl` tools that are provided. (Bug ID 26197105)

Increased dom0 memory requirement when using Mellanox® HCAs on Oracle VM Server

Oracle VM Servers running UEKR4u2 and upward in `dom0` require at least 400MB more memory to use the Mellanox® drivers. This memory requirement is a result of the default size of the SRQ count being increased from 64K to 256K in later versions of the kernel and the `scale_profile` option is now enabled by default in the `mlx_core` module.

In the case where out-of-memory errors are observed in `dom0`, the maximum `dom0` memory size should be increased. Alternative workarounds might involve manually setting the module parameters for the `mlx4_core` driver. To set these parameters, edit `/etc/modprobe.d/mlx4_core.conf` and set `scale_profile` to 0. Alternately, set `log_num_srq` to 16. The preferred resolution to this issue is to increase the memory allocated to `dom0` on an Oracle VM Server. (Bug ID 23581534)

LXC

The following are known LXC issues:

- **The `lxc-net` service does not always start immediately after installation on Oracle Linux 6**

The `lxc-net` service does not always start immediately after installation on Oracle Linux 6, even though this action is specified as part of the RPM post-installation script. This can prevent the `lxcbr0` interface from coming up. If this interface is not up after installation, you can manually start it by running `service lxc-net start`. (Bug ID 23177405)

- **LXC read-only `ip_local_port_range` parameter**

With `lxc-1.1` or later and UEK R4, `ip_local_port_range` is a read-writable parameter under `/proc/sys/net/ipv4` in an Oracle Linux container rather than being read-only. (Bug ID 21880467)

MSI-X interrupt allocation fails during maximum number of `ixgbe/ixgbevf` Virtual Function creation

The Intel `ixgbe/ixgbevf` and Qlogic `qla2xxx` drivers compete for MSI-X resources. As a result, if both drivers are used in a system, and an attempt is made to create the maximum number of Virtual Function (VF) devices that are allowed for the `ixgbe/ixgbevf` driver, an interrupt allocation failure occurs during the creation of the last VF device.

Note that you can create and use up to, but not including, the maximum number of VF devices that are allowed for the `ixgbe/ixgbevf` without encountering this issue. (Bug ID 25952728)

NVMe devices not found under the `/dev` directory after PCI rescan

After removing the PCI bus of NVMe Express (NVMe) adapter card devices and running a rescan of the PCI bus, no NVMe adapter card devices are found under the `/dev` directory.

The workaround for this issue is to also remove the PCI slot that the NVMe adapter card device is plugged into before running a rescan of the PCI bus. (Bug ID 26610285)

OFED iSER target login fails from an initiator on Oracle Linux 6

An Oracle Linux 6 system with the `oracle-Ofed-release` packages installed and an iSER (iSCSI Extensions for RDMA) target configured, fails to login to the iSER target as an initiator. On the Oracle Linux 6 initiator machine, the following behavior is typical:

```
# iscsiadm -m node -T iqn.iser-target.t1 -p 10.196.100.134 --login
Logging in to [iface: default, target: iqn.iser-target.t1, portal:
10.196.100.134,3260] (multiple)
iscsiadm: Could not login to [iface: default, target: iqn.iser-target.t1,
portal: 10.196.100.134,3260].
iscsiadm: initiator reported error (8 - connection timed out)
iscsiadm: Could not log into all portals
```

This is expected behavior resulting from an errata fix for CVE-2016-4564, to protect against a write from an invalid context.

(Bug ID 23615903)

Open File Description (OFD) locks are not supported on NFSv4 mounts

NFS is not designed to handle OFD locking. (Bug ID 22948696).

SDP performance degradation

The Sockets Direct Protocol (SDP), which was designed to provide an RDMA alternative to TCP over InfiniBand networks, is known to suffer from performance degradation on more recent kernels such as UEK R4u2 and later. There is no active development on this protocol.

Although the library for this protocol is still available for this kernel, support is limited. You should consider using TCP on top of IP over InfiniBand as a more stable alternative. (Bug ID 22354885)

Shared Receive Queue (SRQ) is an experimental feature for RDS and is disabled by default

The SRQ function that optimizes resource usage within the `rds_rdma` module is experimental and is disabled by default. A warning message is displayed when you enable this feature by setting the `rds_ib_srq_enabled` flag. (Bug ID 23523586).

Unloading or removing the `rds_rdma` module is unsupported

Once the `rds_rdma` module has been loaded, you cannot remove the module using either `rmmmod` or `modprobe -r`. Unloading of the `rds_rdma` module is unsupported and can trigger a kernel panic. Do not set the `module_unload_allowed` flag for this module. (Bug ID 23580850).

4

Installation and Availability

You can install the Unbreakable Enterprise Kernel Release 4 on Oracle Linux 6 Update 7 or later, or Oracle Linux 7 Update 1 or later, running either the Red Hat compatible kernel or a previous version of the Unbreakable Enterprise Kernel. If you are still running an older version of Oracle Linux, first update your system to the latest available update release.

The Unbreakable Enterprise Kernel Release 4 is supported on the x86-64 architecture, but not on x86.

Installation Overview

If you have a subscription to Oracle Unbreakable Linux support, you can obtain the packages for Unbreakable Enterprise Kernel Release 4 by registering your system with the Unbreakable Linux Network (ULN) and subscribing it to additional channels. See [Subscribing to ULN Channels](#).

If your system is not registered with ULN, you can obtain most of the packages from the Oracle Linux yum server. See [Enabling Access to Oracle Yum Channels](#).

Having subscribed your system to the appropriate channels on ULN or the Oracle Linux yum server, upgrade your system. See [Upgrading Your System](#).

After upgrading to UEK R4, you can replace any existing OFED packages with the Oracle-supported OFED packages, see [Installing the Oracle-Supported OFED Packages](#).

Subscribing to ULN Channels

The kernel image and user-space packages are available on the following ULN channels for Oracle Linux 6:

- `ol6_x86_64_latest` (latest user-space packages for Oracle Linux 6 other than DTrace, OFED, and DRBD packages)
- `ol6_x86_64_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `ol6_x86_64_UEKR4_DTrace_userspace` (`dtrace-utils*`)
- `ol6_x86_64_UEKR4_OFED` (latest OFED tools packages)
- `ol6_x86_64_mysql-ha-utils` (`drbd84-utils`)

The kernel image and user-space packages are available on the following ULN channels for Oracle Linux 7:

- `ol7_x86_64_latest` (all of the latest user-space packages for Oracle Linux 7 other than DTrace, OFED, and DRBD packages)
- `ol7_x86_64_latest_optional` (the latest optional user-space packages for Oracle Linux 7 other than DTrace, OFED, and DRBD packages)
- `ol7_x86_64_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)

- `ol7_x86_64_UEKR4_DTrace_userspace` (`dtrace-utils*`)
- `ol7_x86_64_UEKR4_OFED` (latest OFED tools packages)
- `ol7_x86_64_mysql-ha-utils` (`drbd84-utils`)

The following procedure assumes that you have already registered your system with ULN.

To subscribe your system to a channel on ULN:

1. Log in to <https://linux.oracle.com> with your ULN user name and password.
2. On the Systems tab, click the link named for the system in the list of registered machines.
3. On the System Details page, click **Manage Subscriptions**.
4. On the System Summary page, select each required channel from the list of available channels and click the right arrow to move the channel to the list of subscribed channels.

For Oracle Linux 6, subscribe the system to the `ol6_x86_64_latest` and `ol6_x86_64_UEKR4` channels. If required, you can also add the channels for the DTrace, OFED, and DRBD packages. You do not need to subscribe the system to the `ol6_x86_64_UEK_latest` or `ol6_x86_64_UEKR3_latest` channels.

For Oracle Linux 7, subscribe the system to the `ol7_x86_64_latest` and `ol7_x86_64_UEKR4` channels. If required, you can also add the channels for the DTrace, OFED, and DRBD packages. You do not need to subscribe the system to the `ol7_x86_64_UEKR3` channel.

5. Click **Save Subscriptions**.

For information about using ULN, see [Oracle Linux: Unbreakable Linux Network User's Guide for Oracle Linux 6](#) and [Oracle Linux 7](#).

Enabling Access to Oracle Yum Channels

On the Oracle Linux yum server at <https://yum.oracle.com/>, the kernel image and user-space packages are available on the following channels.

For Oracle Linux 6:

- `ol6_latest` (latest user-space packages for Oracle Linux 6 other than the OFED tool packages)
- `ol6_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `ol6_UEKR4_OFED` (latest OFED tools packages)

For Oracle Linux 7:

- `ol7_latest` (latest user-space packages for Oracle Linux 7 other than the OFED tool packages)
- `ol7_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `ol7_UEKR4_OFED` (latest OFED tools packages)

 **Note:**

To be able to install UEK R4, enable the appropriate `ol6_UEKR4` or `ol7_UEKR4` channel and disable the `ol6_UEKR3_latest` or `ol7_UEKR3` channel.

The DTrace utility and DRBD (Distributed Replicated Block Device) packages are not available on the Oracle Linux Yum Server.

To enable access to the Oracle Linux 6 channels on the Oracle Linux Yum Server, create entries such as the following in `/etc/yum.conf` or in a repository file in the `/etc/yum/repos.d` directory:

```
[ol6_latest]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=https://yum.oracle.com/repo/OracleLinux/OL6/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1

[ol6_UEK_latest]
name=Latest Unbreakable Enterprise Kernel for Oracle Linux $releasever ($basearch)
baseurl=https://yum.oracle.com/repo/OracleLinux/OL6/UEK/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0

[ol6_UEKR4]
name=Latest Unbreakable Enterprise Kernel Release 4 for Oracle Linux $releasever
($basearch)
baseurl=https://yum.oracle.com/repo/OracleLinux/OL6/UEKR4/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1

[ol6_playground_latest]
name=Latest mainline stable kernel for Oracle Linux 6 ($basearch) - Unsupported
baseurl=https://yum.oracle.com/repo/OracleLinux/OL6/playground/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0

[ol6_UEKR4_OFED]
name=OFED supporting tool packages for Unbreakable Enterprise Kernel Release 4 on
Oracle Linux 6 ($basearch)
baseurl=https://yum.oracle.com/repo/OracleLinux/OL6/UEKR4/OFED/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0
priority=20
```

To enable a channel, set the value of the `enabled` parameter for the channel to 1.

To disable a channel, set the value of the `enabled` parameter for the channel to 0.

In the previous example, access is enabled to the `ol6_latest` and `ol6_UEKR4` channels but not to the `ol6_UEK_latest`, `ol6_playground_latest` and `ol6_UEKR4_OFED` channels.

To enable access to the Oracle Linux 7 channels, create entries such as the following:

```
[ol7_latest]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=https://yum.oracle.com/repo/OracleLinux/OL7/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1

[ol7_UEKR4]
name=Latest Unbreakable Enterprise Kernel Release 4 for Oracle Linux $releasever
($basearch)
baseurl=https://yum.oracle.com/repo/OracleLinux/OL7/UEKR4/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1

[ol7_UEKR4_OFED]
name=OFED supporting tool packages for Unbreakable Enterprise Kernel Release 4
on Oracle Linux 7 ($basearch)
baseurl=https://yum.oracle.com/repo/OracleLinux/OL7/UEKR4/OFED/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0
priority=20
```

In this example, access is enabled to the `ol7_latest` and `ol7_UEKR4` channels, but not to the `ol7_UEKR4_OFED` channel.

You can find more information about installing the software at <https://yum.oracle.com/>, from where you can download a copy of a suitable repository file (<https://yum.oracle.com/public-yum-ol6.repo> or <https://yum.oracle.com/public-yum-ol7.repo>).

Upgrading Your System

To upgrade your system to UEK R4:

1. After enabling access to the appropriate channels, including `ol6_UEKR4` or `ol7_UEKR4`, on the Oracle Linux yum server or `ol6_x86_64_UEKR4` or `ol7_x86_64_UEKR4` on ULN, run the following command:

```
# yum update
```
2. After upgrading the system, reboot it, selecting the UEK R4 kernel (version 4.1.12) if this is not the default boot kernel.

See the *Boot and Service Configuration* chapter of [Oracle Linux 7: Administrator's Guide](#) for more information on updating the default boot kernel on Oracle Linux 7.

See https://docs.oracle.com/en/operating-systems/oracle-linux/6/admin/ol_bootconf.html for more information on updating the default boot kernel on Oracle Linux 6.

For instructions on how to install the Oracle-supported OFED packages after upgrading to UEK R4, see [Installing the Oracle-Supported OFED Packages](#).

If you are upgrading from Oracle Linux 7 Update 3 to Oracle Linux 7 Update 4 and you already have Oracle-supported OFED packages for UEK R4 installed on your system, follow the upgrade procedures that are described in [Oracle® Linux 7: Release Notes for Oracle Linux 7.4](#).

If you have questions regarding configuring or using `yum` to install updates, refer to [Oracle Linux: Unbreakable Linux Network User's Guide for Oracle Linux 6 and Oracle Linux 7](#).

The kernel's source code is available via a public git source code repository at <https://oss.oracle.com/git/?p=linux-uek.git;a=summary>.

Installing the Oracle-Supported OFED Packages

The following procedure describes how to install the OFED packages that are provided by Oracle, including how to remove any existing OFED packages.

To install the OFED packages that are provided by Oracle:

1. If your system is registered with ULN, subscribe the system to the `ol6_x86_64_UEKR4_OFED` or `ol7_x86_64_UEKR4_OFED` channel on ULN as appropriate.

By default, the `ol7_x86_64_UEKR4` and `ol7_x86_64_latest` channels are enabled when you register an Oracle Linux 7 system with ULN; and the `ol6_x86_64_UEKR4` and `ol6_x86_64_latest` channels are enabled when you register an Oracle Linux 6 system with ULN. Check that these channels are still enabled before you begin installing the OFED packages provided by Oracle.

If you want to install the packages from the Oracle Linux Yum Server, edit the `yum` repository file at `/etc/yum.repos.d/public-yum-olN.repo` and enable the `ol6_UEKR4_OFED` or `ol7_UEKR4_OFED` repository as appropriate. If there is no `ol6_UEKR4_OFED` or `ol7_UEKR4_OFED` repository in the `yum` repository file, do the following:

- a. Move the existing `yum` repository file to a backup file, for example:

```
# mv /etc/yum.repos.d/public-yum-olN.repo /etc/yum.repos.d/public-yum-olN.repo.bck
```

- b. Download the latest `yum` repository file for Oracle Linux 6 or Oracle Linux 7 from the Oracle Linux yum server.

```
# wget -O /etc/yum.repos.d/public-yum-olN.repo https://yum.oracle.com/public-yum-olN.repo
```

- c. Edit the downloaded `yum` repository file and enable the `ol6_UEKR4_OFED` or `ol7_UEKR4_OFED` repository as appropriate.

2. If you are running Oracle Linux 7, stop and disable the `rdma.service` service.

```
# systemctl stop rdma.service
# systemctl disable rdma.service
```

3. Remove any existing OFED packages:

```
# yum remove 'ibacm*'
# yum remove 'ib-bonding*'
# yum remove 'ibutils*'
# yum remove 'infiniband-diags*'
# yum remove 'libibcm*'
# yum remove 'libibmad*'
# yum remove 'libibumad*'
# yum remove 'libibverbs*'
# yum remove 'libmlx4*'
# yum remove 'librdmacm*'
# yum remove 'libsdp*'
# yum remove 'mstflint*'
```

```
# yum remove 'ofed-docs*'
# yum remove 'ofed-scripts*'
# yum remove 'opensm*'
# yum remove 'perftest*'
# yum remove 'qperf*'
# yum remove 'sdpnstat*'
# yum remove 'rdma-*'
# yum remove 'rds-tools*'
```

4. Clean all yum cached files from all enabled repositories:

```
# yum clean all
```

5. Run one of the following commands, based on server type:

- For a bare metal server, install the OFED packages for UEK R4 as follows:

```
# yum install oracle-ofed-release
```

- For a server that will function as a guest, install the OFED packages for UEK R4 as follows:

```
# yum install oracle-ofed-release-guest
```

6. Enable the RDMA service by entering the following command:

```
# chkconfig rdma on
```

Each UEK release requires a different set of OFED packages. If you change the kernel on your system to a UEK release earlier than UEK R4, remove the existing UEK R4-based OFED packages before installing the correct packages for the new kernel by running the following command:

```
# yum remove --setopt=clean_requirements_on_remove=1 oracle-ofed-release
```

 **Caution:**

Downgrading UEK versions is not advisable, except for testing purposes.

To update OFED packages that are already installed for UEK R4, run this command:

```
# yum update oracle-ofed-release
```