

Unbreakable Enterprise Kernel

Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4



E86280-11
January 2023



Unbreakable Enterprise Kernel Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4,
E86280-11

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	1-3
Xen Hypervisor Improvements	1-3
btrfs Improvements	1-4
Ceph Storage	1-4
Driver Updates	1-4
Additional Notes For Driver Updates	1-6
Technology Preview	1-6
Compatibility	1-7
Header Packages for Development	1-7

2 Fixed and Known Issues

Fixed Issues	2-1
Known Issues	2-3
btrfs, ext4 and xfs: Kernel panic when freeze and unfreeze operations are performed in multiple threads	2-3
btrfs	2-4
ext4	2-4
xfs	2-5
Docker	2-6
DIF/DIX is not supported for ext file systems	2-7
DTrace	2-7
LXC	2-8
Console Appears to Hang when Booting	2-8

OFED iSER target login fails from an initiator on Oracle Linux 6	2-8
Open File Description (OFD) locks are not supported on NFSv4 mounts	2-8
Shared Receive Queue (SRQ) is an experimental feature for Reliable Datagram Sockets (RDS) and is disabled by default	2-9
Unloading or removing the rds_rdma module is unsupported	2-9
Increased dom0 memory requirement when using Mellanox® HCAs on Oracle VM Server	2-9
SDP performance degradation	2-9
DHCP fails for KVM guest on a host using i40e driver module for its network interface card	2-9
Hyper-V fcopy process fails when copying large files from host to guest	2-10

3 Installation and Availability

Installation Overview	3-1
Subscribing to ULN Channels	3-1
Enabling Access to Oracle Yum Channels	3-2
Upgrading Your System	3-4
Installing the Oracle-Supported OFED Packages	3-5

Preface

[Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#) provides a summary of the new features, changes, and known issues in the Unbreakable Enterprise Kernel Release 4 Update 4.

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.

The 4.1.12-94.1.8 release is the fourth update release for UEK R4. It 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 and Oracle Linux 7 Update 3.

Oracle actively monitors upstream checkins 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

- **Overlay File System and XFS**

This update includes a patch that prevents mounting an overlay file system on XFS where the XFS file system is not formatted to support `d_type` entries. This change helps to prevent file system corruption issues and instability in environments that might use an overlay file system, such as Docker. Since most Oracle Linux releases format the root partition with `ftype=0`, any existing overlay file systems that have been configured to run within the root partition are no longer mountable with this update. This can lead to an availability issue after upgrade. Therefore, you should ensure that appropriately formatted storage is available prior to upgrade and that the data for the overlay file system is migrated beforehand. See [xfs](#) for more information on this issue.

- **Support for Skylake Processor Family**

This update includes improved support for servers having Intel® Xeon® E3 v5, Intel® Xeon® Platinum 8100, Intel® Xeon® Gold 6100, Intel® Xeon® Gold 5100, Intel® Xeon® Silver 4100 and Intel® Xeon® Bronze 3100 families of processors (formerly known as Skylake). This release of UEK R4 is the minimum update level that should be used for systems using any of these processor families.

- **ena driver 1.1.2 added**

The `ena` (Elastic Network Adapter) driver module has been added to enable the custom EC2 network interface optimized for high throughput and low latency on EC2 instances. The driver supports a range of ENA devices, is link-speed independent (i.e., the same driver is used for 10GbE, 25GbE, 40GbE, etc.), and has a negotiated and extendable feature set.

- **fm10k driver 0.21.7-k added**

The Intel® Ethernet Switch Host Interface Driver (`fm10k`) module has been added. Patches have been applied to bring this version of the driver up to the most current level for compatibility with this kernel release.

- **macsec driver added**

The `macsec` driver enables MACsec/IEEE 802.1AE network device support. This driver provides authentication and encryption of traffic in a LAN, typically with GCM-AES-128, and optional replay protection. Patches have also been applied to bring this version of the driver up to the most current level for compatibility with this kernel release. The `iproute` package has been updated to include support for the `ip macsec` command and related functionality. Note that this functionality is currently only available on Oracle Linux 7.

- **smartpqi driver 0.9.13-370 added**

The Microsemi Smart Family Controller Driver (`smartpqi`) module has been added. 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.

- **intel_idle driver updated for Skylake processor family**

The `intel_idle` driver that loads on specific Intel® processors in lieu of the legacy ACPI `processor_idle` driver, to provide greater efficiency and stability for Linux on these processors, has been updated to include improved support for servers having Intel® Xeon® E3 v5, Intel® Xeon® E5 v5 and Intel® Xeon® E7 v5 families of processors (formerly known as Skylake). This driver is compiled into the kernel.

- **vmmouse driver updated for Secure Boot**

The `vmmouse` driver has been updated with patches that enable it to work when Secure Boot is enabled.

- **Hardware enablement for latest HPE SGI UV platforms**

The kernel is updated to apply upstream patches that enable hardware support for an HPE system utilizing SGI UV technology (HPE acquired SGI). The patches are limited to HPE system specific sections of the kernel or are properly guarded to only impact HPE systems.

- **Userfaultfd support added to kernel**

The kernel is updated to support `userfaultfd`, a new feature that allows on-demand paging from user space and allows user space applications to take control of various memory page faults. The code for this takes advantage of several patches, including a back port of patches for `hugetlbfs`.

- **RDS improvements for cluster scalability on shared-port HCAs**

The code for Reliable Datagram Sockets (RDS) has been updated to better handle large-scale clustering on shared-port InfiniBand HCAs. These changes include a modification to the code for active bonding to decouple the queue from the Reliable Datagram Sockets (RDS) code, and improvements to the RDS code for sync point failback scenarios to handle a staggered network interface readiness issue across the cluster that resulted in increased load. These changes better facilitate larger clusters and improve performance of RDS active bonding.

DTrace Improvements

A number of bug fixes and enhancements including module and utility updates are included for DTrace on UEK R4u4 bringing the current version to 0.6.0. Notable changes include:

- **Function Boundary Tracing (FBT)**

Function Boundary Tracing is now available in DTrace. Work includes:

- Handling of memory, data and instruction access traps.
- Improved output for Kernel OOPS debugging, including DTrace probe context and last probe processed.
- New user stack walker implementation within the kernel for consistency and better access to low level structures like page tables; and a field to track the end of the user stack.
- FBT black lists are updated with functions that may be invoked during probe processing, including functions starting with `insn_*` or `inat_*` for instruction analysis.

- **Typed SDT arguments**

Changes have been made to the kernel, provided DTrace modules and to the userspace utilities for DTrace to better facilitate working with typed SDT arguments. Previous versions of DTrace required updates to the DTrace module for added SDT and perf probes. This update allows kernel developers to add probes without updating the module, relying instead on DTrace's ability to dynamically compute argument types to navigate structures. Additionally, the `DTRACE_PROBE()` macro no longer needs an argument-counting suffix making it much easier to add SDT provider `DTRACE_PROBE()` macro wrappers.

- **Process Grabbing**

Several fixes have been applied to improve DTrace process grabbing, including the option now to 'self-grab' to obtain better backtraces and to allow DTrace consumers to monitor the operation of non-DTrace components using `-p`.

- **Added Providers**

New network providers have been added, including an IP Provider supporting `ip:::send`, `ip:::receive`, `ip:::drop-in` and `ip:::drop-out`, with parameters compatible with other implementations and appropriate translators. IPv4 and IPv6 are both supported.

- **Is-enabled probes**

DTrace supports 'is-enabled' probes, which are conditionals that allow you to execute different code if a probe is enabled, so you can avoid collecting expensive data needed only for certain probes when those probes are inactive.

- **General bugfixes**

Numerous bugfixes have been applied to provide greater stability and better performance.

Xen Hypervisor Improvements

Many upstream patches have been applied and back-ported from the Linux 4.7 and Linux 4.8 kernels. These include fixes for PCI passthrough, changes for grant device performance

improvements and fixes for compilation warnings. Most significantly, a fix has been applied that allows the kernel to run as a `kexec` kernel.

Patches were also applied for the `xen-netback` and `xen-netfront` virtual network driver modules to improve stability during migrations and correct error handling.

The `xen-blkback` and `xen-blkfront` drivers for block device support have been updated to fix issues related to resume after migration. Other fixes include upstream patches for stability and performance.

btrfs Improvements

Some upstream patches for various bug fixes have been applied to btrfs to improve stability. These include patches to allow `fsync` to work after cloning into a file and a fix to allow reads to be issued during the `down_interval` if `drop_writes` is configured. .

A patch was applied to the `btrfs_abort_transaction` code to properly identify blocks that are not part of the current Copy-on-Write transaction. .

Another patch was applied to the input/output control code to account for the possibility that the file descriptor used for a snapshot source may be located on another file system type.

Ceph Storage

Ceph Storage 2.0 is based on the Ceph Community Jewel release (v10.2.2) and is described in the Ceph Storage for Oracle Linux Release 2.0 Release Notes available at https://docs.oracle.com/cd/E52668_01/E66514/html/index.html

Note that Ceph Storage 2.0 is not supported for Oracle Linux 6. Support for earlier versions of Ceph Storage on Oracle Linux 6 are also discontinued in UEK R4u4.

At the time of this update release, the **Ceph File System** (CephFS) feature of Ceph is in technology preview. To track technology previews and supported features for Ceph Storage, please refer to the release notes.

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.

Driver	Version	Description
aacraid	1.2-1[41010]-ms + Patches	Adaptec Advanced Raid Products
be2iscsi	11.2.1.0 + Patches	Broadcom/Avago OneConnect Open-iSCSI Driver
be2net	11.1.0.0 + Patches	Broadcom/Emulex OneConnect 10Gbps NIC Driver
bnx2	2.2.6 + Patches	QLogic BCM5706/5708/5709/5716 Driver
bnx2fc	2.10.3 + Patches	QLogic FCoE Driver

Driver	Version	Description
bnx2i	2.7.10.1 + Patches	QLogic NetXtreme II BCM5706/5708/5709/57710/57711/57712/57800/57810/57840 iSCSI 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	Broadcom BCM573xx network driver
ena	1.1.2	Amazon Elastic Network Adapter Driver
enic	2.3.0.31	Cisco VIC Ethernet NIC Driver
fm10k	0.21.7-k	Intel® FM10K Driver
hpilo	1.5.0	HPE Integrated Lights-Out (iLO) Driver
hpsa	3.4.18-0	PMC-Sierra Driver for HP Smart Array Controller
hpwdt	1.4.0 + Patches	HPE Software Watchdog Driver
i40e	1.6.21-k	Intel® Ethernet Connection XL710 Network Driver
i40evf	1.6.21-k	Intel® XL710 X710 Virtual Function Network Driver
ixgbevf	3.2.2-k	Intel® 82599 Virtual Function Driver
lpfc	11.2.0.5	Broadcom/Emulex LightPulse Fibre Channel SCSI driver
macsec	1.0.0	MACsec IEEE 802.1AE Driver
megaraid_sas	06.812.07.00-rc1	Avago MegaRAID SAS Driver
mpt3sas	14.101.00.00	LSI MPT Fusion SAS 3.0 Device Driver
nvme	1.0 + Patches	NVM Express (NVMe) Device Driver
qed	8.10.10.20	QLogic FastLinQ 4xxxx Ethernet Driver and Core Module
qede		
qla2xxx	8.07.00.38.40.0-k + Patches	QLogic Fibre Channel HBA Driver
qlcnlc	5.3.65	QLogic 1/10 GbE Converged/ Intelligent Ethernet Driver
qlge	1.00.00.35	QLogic 10 Gigabit PCI-E Ethernet Driver
smartpqi	0.9.13-370	Microsemi Smart Family Controller Driver

Driver	Version	Description
xen-blkback xen-blkfront xen-netback xen-netfront	Upstream Patches	Xen Virtual Block Device and Xen Virtual Network Device drivers
vmxnet3	1.4.a.0-k	VMware vmxnet3 virtual NIC driver
xscore	6.0.r8043	Xsigo Oracle Virtual Network core driver
xsvhba	6.0.r8043	Xsigo Oracle Virtual Network VHBA driver
xsvnic	6.0.r8043	Xsigo Oracle Virtual Network XSVNIC network driver
xve	6.0.r8043	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`: 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 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 interoperoperation 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.

- **NFS server-side copy offload**

NFS server-side copy offload is an NFS v4.2 feature that reduces the overhead on network and client resources by offloading copy operations to one or more NFS servers rather than involving the client in copying file data over the network.

- **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 may 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 and Known Issues

This chapter describes the fixed and known 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.

Fixed Issues

The following issues have been fixed in this update.

- **Severe NVMe storage performance degradation issue fixed**

A bug that caused significant performance degradation when using NVMe devices in conjunction with earlier UEK R4 releases has been resolved. The issue caused sequential reads with abnormal request sizes that severely limited I/O regardless of block size. NVMe performance related patches were applied to the kernel and the issue has been resolved.

- **ocfs2: Slow journal replay fixed**

A minor bug in the `ocfs2` code caused journal replay for a dead node to take longer than necessary as all blocks from the dead node's journal inode were loaded from disk to memory to avoid a stale cache, however performance is enhanced if only the cached blocks are reloaded. A patch was applied to improve recovery performance.

- **ocfs2: Fixed issue releasing disk space after file deletion**

An issue that caused `ocfs2` to not release disk space after a large number of files were deleted has been fixed. The patches applied fixed the code that extended credits while there are free cached blocks and while flushing the truncate log.

- **ocfs2: Fixed issue with unknown option 'ExecRestart' in section 'Service' for the `o2cb.service` file**

An invalid entry in the `o2cb.service` file that was included in the `ocfs2-tools` user space package had been resolved. The problem caused the following message to appear in the `systemd` status or in `/var/log/messages`:

```
systemd[1]: [/usr/lib/systemd/system/o2cb.service:11] Unknown lvalue
'ExecRestart' in section 'Service'
```

The problem is now resolved in `ocfs2-tools-1.8.6-9.el7` and later.

- **Kernel panic during storage device reset when using the `lpfc` driver module**

A bug that caused a kernel panic during a storage device reset was fixed in the `lpfc` driver module. The issue appeared during both when an `sg_reset` command was issued and when SCSI EH (Error Handling) triggered a reset. The patch fixes this issue.

- **Hyper-V clock source changed to use TSC**

An upstream fix that changed the Hyper-V clock source to use the Time Stamp Counter (TSC), for greater efficiency in kernel operations that involve reading time stamps, has been backported into this release.

- **Hyper-V storage driver performance improvements**

Upstream updates to the `storvsc` Hyper-V storage driver were included to provide moderate performance improvement of I/O operations for certain workloads.

- **Hyper-V fix for guest reboot on failover issue**

A bug that caused virtual machines running on Hyper-V to reboot during a graceful node failover, so that live migration was unsuccessful, has been fixed. This problem was caused as a result of a failure to check that all heartbeat and vmbus messages were correctly processed. A patch was applied and the problem is resolved.

- **Hyper-V fix for incorrect receive checksum offloading in the `netvsc` driver**

A bug in the Hyper-V `netvsc` driver caused TCP packets with a bad receive checksum to be passed up the stack to the application layer, potentially causing data corruption. The fix included in this update causes packets with incorrect checksums to be dropped with an error.

- **Update for network bonding to fix `primary_reselect` with failure**

The `primary_reselect` option used to define the reselection policy for the primary slave in a network bond would not behave correctly when set to `failure` or `2`. The issue would result in the primary slave becoming active again when it had recovered. The expected, and documented, behavior is that the primary slave should not become active until the current active slave is down. A patch was applied to the bonding code to change the `bond_find_best_slave` functionality to avoid traversing members if the primary interface is not a candidate for failover or reselection and the current active slave is still up.

- **Fix for SCSI code that caused a kernel crash when a target node in an HA pair was rebooted**

Several patches were applied to the `scsi` driver code to fix an issue that caused the kernel to crash when a target node in an HA pair was rebooted on a SAN booted LUN configured for multipath. The issue resulted when the SCSI target device was marked for removal and there was a delay in getting it into the `DEL` state. This could cause the same target to get marked for removal twice. The applied patches resolved the issue.

- **Fix applied for Mellanox® `m1x4` driver to resolve the "Node crashed at `cache_alloc_refill+0x1ab`" error**

A bug that allowed multiple work queues to be allocated to the same `id_map_ent` structure in the `m1x4` driver code was patched. This issue could cause a kernel crash if a worker routine cleaned up and freed the structure. The patch checks that previously queued work on the structure has been successfully cancelled before new work is queued on the same structure.

- **`timer` code patched for race condition that could result in a kernel oops**

A patch was applied to fix an issue in the `timer` code that resulted in a kernel oops when a timer was migrated to an alternate CPU and had been left unlocked on the original CPU. The fix performs a proper migration and does the appropriate checks and locks during the migration to prevent the race condition.

- **Race condition in freeing aging forwarding tables in `xsgo` driver fixed**

A fix was applied to the `xsgo` driver code to detect and avoid a potential race condition while accessing the forwarding table during the deletion of an aged forwarding entry. This issue could cause nodes to reboot inadvertently.

- **Race condition issue in the `optrom` functions in the `qla2xxx` driver fixed**

A race condition that triggered when a thread modified the `optrom` buffer, in the `qla2xxx` driver, at the same time that another thread attempted to read from it was patched in this update. This issue was fixed by getting a mutex lock before checking the `optrom` state. The problem could result in kernel panic and inadvertent system reboots.

- **`netxen` driver patched for incorrect error handling**

The QLogic/NetXen (1/10) GbE Intelligent Ethernet Driver was patched to fix an error handling issue that prevented the `netxen_rom_fast_read()` function from ever returning -1. Additional vendor patches were also applied.

- **RDS patched to fix QoS threshold calculation**

When the Reliable Datagram Sockets (RDS) protocol was placed under loads that caused it to drop packets, the `qos_threshold_exceeded` parameter was not incremented because the RDMA payloads were calculated incorrectly. This caused the Quality of Service (QoS) functionality to fail. A patch was applied to fix this calculation so that the QoS threshold could be enforced.

- **RDMA package updated to fix `m1x4_ib` insertion error when RDMA starts**

A bug that caused a benign insertion error when the RDMA service started was fixed in the `rdma-3.10-3.0.25` package. With this update release, the RDMA package is updated to `rdma-3.10-3.0.31` to provide several further bug fixes and code improvements.

- **NFSv4 issue with client incrementing the lock sequence number on `NFS4ERR_MOVED` fixed**

A change in the NFSv4 specification meant that when the NFS client connected to a server based on the newer specification and sent a lock to the source and got an `NFS4ERR_MOVED` response, if it resent the lock to the destination, it would generate a bad sequence ID error: `NFS4ERR_BAD_SEQID`. The UEK R4 NFS client adheres to RFC 3530, which does not cover `NFS4ERR_MOVED`. A patch was applied to better adhere to RFC 7530 and to prevent the client from incrementing the lock sequence ID after receiving an `NFS4ERR_MOVED` from the server.

Known Issues

This section 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 performed across multiple threads on any supported file system can cause the system to hang and the kernel to panic. This 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 `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)
- 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; 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 may 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 used to store data and metadata may lead to some confusion on the information returned by tools like `df`. Sometimes, metadata may 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 for more information.
- 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)
- 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)
- 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

- System hangs when processing corrupted orphaned inode list
If the orphaned inode list is corrupted the inode may 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 resulting in the 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. This causes an integer overflow

in the routines underlying writeback, which results in the kernel locking up. (Bug ID 25565527)

- A hang occurs with the `ext4` file system during the dynamic expansion of inode size when using the inode's `i_extra_size` field. (Bug ID 25718971)

xf

- Directory readahead completions can hang the system after unmount

Directory readahead can hang the system if the file system is unmounted suddenly after mount. If a directory readahead is delayed for long enough, buffer I/O completion may occur after the unmount has completed. The asynchronous nature of directory readahead I/O means that when the readahead I/O completion occurs, core data structures may have been freed, causing completion to run into invalid memory accesses. This can result in a kernel panic and system hang. (Bug ID 25550712)

- 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.

```
[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/xf/xf_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"
```

The problem is that the log attempts to replay a buffer update that is no longer valid due to subsequent replayed updates. This results in 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 routines underlying writeback, result 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 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 may 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.

Since the root partition on Oracle Linux is automatically formatted with `-n ftype=0` where XFS is selected as the file system, for backward compatibility reasons, if you have overlay file systems in place already and these are not hosted on alternate storage, you must migrate these 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 mount this. 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
```

An entry in `/etc/fstab`, to make a permanent mount for this storage, may 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)

Docker

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

```
Rpmdm 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)

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)

DTrace

- 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 may not have accurate symbol resolution for symbols introduced by the `dlopen()`. (Bug ID 20045149)

LXC

- **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)

Console Appears to Hang when Booting

When booting Oracle Linux 6 on hardware with an ASPEED graphics controller, the console may 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)

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).

Shared Receive Queue (SRQ) is an experimental feature for Reliable Datagram Sockets (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).

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 is because the default size of the SRQ count was 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 may involve manually setting the module parameters for the `mlx4_core` driver. To do this, 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 Oracle VM Server. (Bug ID 23581534)

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 R4 U2 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)

DHCP fails for KVM guest on a host using `i40e` driver module for its network interface card

The `i40e` driver module does not function correctly when bridged with a `virtio` network interface that makes a DHCP request from within a KVM guest. Although the DHCP request is sent through the `virtio` network interface, the request does not reach the network beyond the `i40e` network interface card. This regression is the result of a patch that was applied to enable VSI broadcast in promiscuous mode instead of adding a broadcast filter. The fix is under investigation. (Bug ID 25825419)

Hyper-V fcopy process fails when copying large files from host to guest

Oracle Linux guests, running UEK R4 and hosted on Microsoft® Windows 2012 R2 Hyper-V systems, can have trouble completing a file copy when using the `Guest services` facility, along with the `hypervfcopyd` service, to copy a large file from the host to the guest virtual machine. This issue can be seen when copying a file larger than 10GB from the host to the guest mount point and results in an error message appearing in the host log, similar to the following:

```
Copy-VMFile : The operation cannot be performed while the virtual machine is
in its current state. The name of the virtual machine is TestVM and its ID is
9ebdc189-439a-4db2-b33f-05f3b07726bf.
At line:1 char:1 + Copy-VMFile "TestVM" -SourcePath "E:\ISO\largefile.7z"
-DestinationPath "/mnt" -C ...
...
+ CategoryInfo          : InvalidResult: (:) [Copy-VMFile],
VirtualizationOperationFailedException
+ FullyQualifiedErrorId :
InvalidState,Microsoft.HyperV.PowerShell.Commands.CopyVMFileCommand
```

The issue can be reproduced across a variety of file system types. The issue is intermittent and does not appear during every copy attempt. (Bug ID 25866707 and 25866691)

3

Installation and Availability

You can install 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:

- `o16_x86_64_latest` (latest user-space packages for Oracle Linux 6 other than DTrace, OFED, and DRBD packages)
- `o16_x86_64_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `o16_x86_64_UEKR4_DTrace_userspace` (`dtrace-utils*`)
- `o16_x86_64_UEKR4_OFED` (latest OFED tools packages)
- `o16_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:

- `o17_x86_64_latest` (all of the latest user-space packages for Oracle Linux 7 other than DTrace, OFED, and DRBD packages)
- `o17_x86_64_latest_optional` (the latest optional user-space packages for Oracle Linux 7 other than DTrace, OFED, and DRBD packages)
- `o17_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

At 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 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 this 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 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 to upgrade the system to UEK R4:

```
# 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 <https://docs.oracle.com/en/operating-systems/oracle-linux/7/admin/ol7-bootconf.html> 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 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

This section describes how to install the OFED packages provided by Oracle, including how to remove any existing OFED packages.

To install the OFED packages 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-ufed-release
```

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

```
# yum install oracle-ufed-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:

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

 **Caution:**

Downgrading UEK versions is not advisable except for testing purposes.

To update OFED packages already installed for UEK R4, simply run:

```
# yum update oracle-ufed-release
```