

Oracle® Linux

**Release Notes for Unbreakable Enterprise Kernel Release 4
Update 2**

ORACLE®

E78976-06
May 2018

Oracle Legal Notices

Copyright © 2018, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Abstract

This document contains information on Update 2 to the Unbreakable Enterprise Kernel Release 4. This document may be updated after it is released. To check for updates to this document, and to view other Oracle documentation, refer to the Documentation section on the Oracle Technology Network (OTN) Web site:

<http://www.oracle.com/technology/documentation/>

This document is intended for users and administrators of Oracle Linux. It describes potential issues and the corresponding workarounds you may encounter while using the Unbreakable Enterprise Kernel Release 4 with Oracle Linux 6 or Oracle Linux 7. Oracle recommends that you read this document before installing or upgrading the Unbreakable Enterprise Kernel Release 4.

Document generated on: 2018-05-30 (revision: 139)

Table of Contents

Preface	v
1 New Features and Changes	1
1.1 Notable Changes	1
1.2 Xen Improvements	2
1.3 btrfs Improvements	2
1.4 Driver Updates	2
1.5 Technology Preview	4
1.6 Compatibility	5
1.7 Header Packages for Development	5
2 Fixed and Known Issues	7
2.1 Fixed Issues	7
2.2 Known Issues	8
2.2.1 btrfs	8
2.2.2 DIF/DIX is not supported for ext file systems	8
2.2.3 Docker	9
2.2.4 DTrace	9
2.2.5 LXC	9
2.2.6 Console Appears to Hang when Booting	9
2.2.7 mlx4_ib Insertion Error when RDMA Starts	9
2.2.8 OFED iSER feature causes Out-Of-Memory error on Oracle Linux 6	10
2.2.9 XFS quotas are disabled after a read-only remount on Oracle Linux 6	10
2.2.10 Open File Description (OFD) locks are not supported on NFSv4 mounts	10
2.2.11 Shared Receive Queue (SRQ) is an experimental feature for Reliable Datagram Sockets (RDS) and is disabled by default	10
2.2.12 Unloading or removing the <code>rds_rdma</code> module is unsupported	10
2.2.13 Sufficient memory must be allocated when using multiple Oracle Infiniband HCAs on Oracle VM Server	10
2.2.14 Increased dom0 memory requirement when using Mellanox HCAs on Oracle VM Server	10
2.2.15 SDP performance degradation	11
3 Installation and Availability	13
3.1 Installation Overview	13
3.2 Subscribing to ULN Channels	13
3.3 Enabling Access to Oracle Linux Yum Server Channels	14
3.4 Upgrading Your System	16
3.5 Installing the Oracle-Supported OFED Packages	16

Preface

The *Oracle Linux Unbreakable Enterprise Kernel Release Notes* provides a summary of the new features, changes, and fixed and known issues in the Unbreakable Enterprise Kernel Release 4 Update 2.

Audience

This document is written for system administrators who want to use the Unbreakable Enterprise Kernel with Oracle Linux. It is assumed that readers have a general understanding of the Linux operating system.

Related Documents

The latest version of this document and other documentation for this product are available at:

<http://www.oracle.com/technetwork/server-storage/linux/documentation/index.html>.

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.
<code>monospace</code>	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

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Chapter 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 4.1.12-61.1.6 release is the second 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 and Oracle Linux 7 Update 2.

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.

1.1 Notable Changes

- **Ceph is no longer under technology preview.** Ceph presents a uniform view of object and block storage from a cluster of multiple physical and logical commodity-hardware storage devices. Ceph can provide fault tolerance and enhance I/O performance by replicating and striping data across the storage devices in a Storage Cluster. Ceph's monitoring and self-repair features minimize administration overhead. You can configure a Storage Cluster on non-identical hardware from different manufacturers. Ceph for Oracle Linux is based on the Ceph Community Jewel release (v10.2.2). It includes the Object Store, Block Device, Storage Cluster, and Object Gateway components of Ceph. For more information, see the [Ceph Storage for Oracle Linux Release 2.0 Release Notes](#).
- **sif driver 1.0.0.1 added.** The `sif` (Scalable Interface) driver module that enables PCIe-based Infiniband HCAs (Host Channel Adapter), such as the Oracle Dual-port QDR IB Adapter M4 and the Oracle Dual-port EDR IB Adapter, has been added in this update.
- **qed and qede 8.7.1.20 drivers added.** The `qed` and `qede` driver modules that enable QLogic FastLinQ 4xxxx Ethernet devices, have been added in this update.
- **bnxt_en driver 1.2.0 added.** The `bnxt_en` driver module that enables Broadcom BCM573xx network devices, has been added in this update.
- **RDS performance enhancements and bug fixes.** A large number of performance enhancements and bug fixes have been included in this release for RDS (Reliable Datagram Sockets). Of particular note are the following fixes and enhancements:
 - Enhancement patches have been applied for RDS to support the larger message size catered for by Infiniband. This change improves RDS performance over Infiniband transport by increasing the RDS fragment size while maintaining backward compatibility for previous RDS versions.
 - Performance improvements were made to the `rds_fmr_flush` workqueue to take advantage of targeted CPU scheduling on NUMA-aware systems.
 - Old code and unused features within the code were removed to ease the maintenance burden, tighten up the code and improve code security.
 - Several patches were merged into the code to improve diagnostics and supportability.
 - Improvements were made to the RDS code for better compatibility with Infiniband HCA drivers.

- **Important NVMe driver updates were applied and tested.** The Non-Volatile Memory Express driver was updated and patched with many new fixes that were of particular interest to existing Oracle customers. Several rigorous tests were run against the updated driver.
- **DTrace enhancement to export `perf_events` tracepoints as SDT probes.** Upstream Linux `perf_events` tracepoints are now automatically exported as SDT probes for DTrace.

1.2 Xen Improvements

A number of bug fixes and security fixes have been backported for this release to include the features available in Xen 4.4, Xen 4.5 and Xen 4.6.

A fix was applied to improve the code that handled cycling through the E820 map sent to the bootloader by the BIOS during boot. This fix resolved a potential mismatch error that resulted in a Dom0 kernel crash during boot on some systems.

A new feature allowing the dynamic configuration of VBD (Virtual Block Device) I/O buffers has been applied to allow for better performance tuning per virtual block device, depending on specific application requirements. The values for these buffer configuration settings are available in `/sys/devices/vbd-xxx/max_indirect_segs`, `/sys/devices/vbd-xxx/max_ring_page_order` and `/sys/devices/vbd-xxx/max_queues`.

1.3 btrfs Improvements

- **Improvements to extent accounting.** Extent accounting has been improved to provide fixes for a number of bugs. Improvements have also been made to handle scenarios where there is no disk space, but an `rm` operation is in action to free up disk space.
- **`fsync` issues fixed.** A fix was applied to correct an issue that could corrupt the filesystem after a hard link was removed and `fsync` was called. Another fix was applied to prevent a potential data loss issue that could result through calling `fsync` after an inode was renamed to match a removed inode and a power outage occurred. .
- **Infinite loop during an incremental send fixed.** A fix was applied to correct an ENOMEM error that resulted from an infinite loop when building path strings during an incremental send. This issue was resolved by applying a patch that checked if an orphanized directory inode rename required delay.
- **Improvements to transaction handling for hard and symbolic links.** Fixes were applied to improve transaction handling for the creation of hard and symbolic links. One of these fixes resolved a memory leak that could occur when creating a hard link.

1.4 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.

Table 1.1 Updated Drivers in UEK R4 U2

Driver	Version	Description
<code>aacraid</code>	1.2-1[40709]-ms + Patches	Adaptec Advanced Raid Products
<code>be2iscsi</code>	11.0.0.0	Broadcom/Avago OneConnect Open-iSCSI Driver
<code>be2net</code>	11.0.0.0 + Patches	Broadcom/Emulex OneConnect 10Gbps NIC Driver
<code>bfa</code>	3.2.25.0	QLogic BR-series Fibre Channel HBA Driver

Driver	Version	Description
bnx2x	1.713.10	QLogic BCM57710 / 57711 / 57711E / 57712 / 57712_MF / 57800 / 57800_MF / 57810 / 57810_MF / 57840 / 57840_MF Driver
bnxt_en	1.2.0	Broadcom BCM573xx network driver
bonding	3.7.1 + Patches	Ethernet Channel Bonding Driver
fnic	1.6.0.24 + Patches	Cisco FCoE HBA Driver
hpsa	3.4.16-0	PMC-Sierra Driver for HP Smart Array Controller
i40e	1.5.16-k	Intel(R) Ethernet Connection XL710 Network Driver
i40evf	1.5.10-k	Intel(R) XL710 X710 Virtual Function Network Driver
ixgbe	4.4.0-k	Intel(R) 10 Gigabit PCI Express Network Driver
ixgbev	2.12.1-k + Patches	Intel(R) 82599 Virtual Function Driver
lpfc	11.1.0.4	Broadcom/Emulex LightPulse Fibre Channel SCSI driver
mpt3sas	13.100.00.00	LSI MPT Fusion SAS 3.0 Device Driver
nvme	1.0 + Patches	Non-Volatile Memory Express Driver (Fultondale)
qed	8.7.1.20	QLogic FastLinQ 4xxxx Ethernet Driver and Core Module
qedf		
qla2xxx	8.07.00.38.40.0-k	QLogic Fibre Channel HBA Driver
qlcnic	5.3.64	QLogic 1/10 GbE Converged/Intelligent Ethernet Driver
sif	1.0.0.1	Oracle SIF Infiniband HCA driver
tg3	3.137 + Patches	Broadcom Tigon3 Ethernet driver
vmw_vsock_vmci_transport	1.0.4.0-k	VMCI transport for Virtual Sockets
vmxnet3	1.4.7.0-k	VMware vmxnet3 virtual NIC driver
xen-blkback	Upstream Patches	Xen Virtual Block Device and Xen Virtual Network Device drivers
xen-blkfront		
xen-netback		
xen-netfront		
xscore	6.0.r8020	Xsigo Oracle Virtual Network core driver
xsvhba	6.0.r8020	Xsigo Oracle Virtual Network VHBA driver
xsvnic	6.0.r8020	Xsigo Oracle Virtual Network XSVNIC network driver
xve	6.0.r8020	Xsigo Oracle Virtual Network Virtual Ethernet driver

Some InfinBand 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.

1.5 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.

- **Ceph File System**

The Ceph File System (CephFS) feature of Ceph is in technology preview.

- **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.

- **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.

1.6 Compatibility

Oracle Linux maintains user-space compatibility with Red Hat Enterprise Linux, 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.

1.7 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
```

Chapter 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.

2.1 Fixed Issues

The following issues have been fixed in this update.

- **Improvements to robust recovery of RDS over TCP sockets.** Reliable Datagram Sockets (RDS) that use TCP as an underlying carrier protocol previously lacked the ability to recover state in the instance where the TCP connection was reset. This issue is patched to improve the robustness of RDS over TCP.
- **Improvements to facilitate Infiniband transport flow control for RDS.** Several patches have been applied to fix and improve RDS flow control when using Infiniband transport. Previously, Infiniband flow control was always disabled and the `rds_ib_sysctl_flow_control` flag was ignored.
- **NFS over RDMA on IPv6 support fixed.** Fixes were applied to resolve an issue that was causing NFS over RDMA on IPv6 to fail due to an inability to handle the `rdma6` identifier used as the netid for this protocol configuration.
- **NFS over RDMA client mount point hang during cluster failover fixed.** A fix was applied to resolve an issue that was causing NFS over RDMA to hang for a client mount point when a clustered server fails over.
- **Backported fix for OCFS2 cluster nodes self-fencing when storage is unavailable.** A patch that fixed an issue where OCFS2 fences nodes when the disk-heartbeat timeout issue was reached has been applied to allow nodes to take a quorum vote when the heartbeat timeout is reached. If the node can see all other nodes, and the storage is not accessible from all nodes, then the disk-heartbeat timeout is reset by a calculated amount.
- **API compatibility issue resolved for Mellanox OFED libraries.** Several library packages have been updated to add support for the XRC (eXtended Reliable Connected) API to resolve a compatibility issue with the Mellanox OFED libraries.
- **Server reboot issue triggered by large writes on an iSCSI LUN is fixed.** A server reboot issue triggered by a write of over 100 MB on an iSCSI LUN has been fixed. This issue resulted in an error similar to the following:

```
iscsid: Kernel reported iSCSI connection 3:0 error (1010 - ISCSI_ERR_BAD_ITT:
Received invalid initiator task tag from target) state (3)
```

- **Kernel panic at `__blk_bios_map_sg()` resolved.** A bug which resulted in a kernel panic at `__blk_bios_map_sg()` has been resolved. This issue seemed to result when a cloned request was retried on other queues without checking the queue limits of the queue. This sometimes resulted in calculations for `nr_phys_segments` to be wrong leading to a crash in `scsi_init_sgtable()`. A fix has been applied to always check queue limits for cloned requests.
- **ENobuf error and RDS socket data sending hang resolved.** An issue that resulted in and ENobuf error and caused send timeouts on RDS sockets has been resolved. The issue was identified

as resulting from confusion during an RDS congestion bitmap update where a race condition occurred if multiple threads running on different CPUs updated the same set of 64 bits that map to 64 different ports. The operation was made atomic to fix the CPU cache issue that resulted.

- **Vendor ID update to correctly identify newer QLogic NX2 HBAs.** Patches were applied to the `iscsi-initiator-utils` to correctly identify QLogic NX2 HBAs.
- **IPoIB send hang resulting from concurrent access to `tx_outstanding` fixed.** An issue that caused IPoIB send hangs to occur as a result of concurrent writes to `tx_outstanding`. The value of `tx_outstanding` is used to determine the stop/open on the IPoIB interface send queue. If the value of `tx_outstanding` is incorrect, IPoIB send can hang. The fix required that access to `tx_outstanding` is serialized.
- **`ext4` patch applied to fix preallocated space after truncation.** A bug which caused an `ext4` filesystem to report the incorrect preallocated space size after an operation to truncate up past `i_size` has been fixed. The patch prevents the trimming of `fallocated` blocks that are beyond the new `i_size` and that are flagged with the `KEEP_SIZE` flag. The `truncate` code was reworked to make it easier to handle and to ensure that `i_disksize` is properly set when truncating up.

2.2 Known Issues

This section describes the known issues in this update.

2.2.1 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. The workaround is to run `sync` before retrying the operation. If this does not help, remount the file system with the `-o nodatacow` option and delete some files to free up space. See <https://btrfs.wiki.kernel.org/index.php/ENOSPC>.
- 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` Oracle Linux yum server repositories. 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` Oracle Linux yum server repositories.

2.2.2 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)

2.2.3 Docker

- 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. (Bug ID 21804564)

2.2.4 DTrace

- Argument declarations for probe definitions cannot be declared with derived types such as `enum`, `struct`, or `union`.
- The following compiler warning can be ignored for 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)

2.2.5 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)

2.2.6 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)

2.2.7 mlx4_ib Insertion Error when RDMA Starts

You can safely ignore the following error when the Oracle-supported RDMA service starts:

```
Loading OpenIB kernel modules:insmod: error inserting
'/lib/modules/3.8.13-68.3.4.el6uek.x86_64/kernel/drivers/infiniband/hw/mlx4/
mlx4_ib.ko': -1 File exists
```

The message indicates that the `mlx4_ib` module is already loaded. (Bug ID 21410136)

2.2.8 OFED iSER feature causes Out-Of-Memory error on Oracle Linux 6

An Oracle Linux 6 system with the `oracle-ufed-release` packages installed and an iSER (iSCSI Extensions for RDMA) target configured, hangs and generates an Out-Of-Memory error when the initiator logs into the iSER target disk. The login fails and the target system hangs. The following error is displayed on the remote console:

```
Out of memory: Kill process ...
```

The cause is under investigation and there is no known workaround at this time. (Bug ID 22230403)

2.2.9 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)

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

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

2.2.11 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).

2.2.12 Unloading or removing the `rds_rdma` module is unsupported

Once the `rds_rdma` module has been loaded, you cannot remove the module using either `rmmod` 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).

2.2.13 Sufficient memory must be allocated when using multiple Oracle Infiniband HCAs on Oracle VM Server

Oracle Infiniband HCAs that require the `sif` (Scalable Interface) driver to be loaded must have sufficient memory available to support each interface. The `sif` driver requires at least 30 MB of memory for the driver itself and then 50 MB of memory per interface in SR-IOV PF and VF conditions. This is particularly important when loading this driver on Oracle VM Server 3.4.1 and above, since the default memory allocation for dom0 may not be sufficient to support multiple psif instances and could result in "Out of memory" errors. To resolve this issue, increase the memory allocated to dom0 on Oracle VM Server. (Bug ID 23523713)

2.2.14 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)

2.2.15 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. 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)

Chapter 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.

3.1 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 [Section 3.2, “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 [Section 3.3, “Enabling Access to Oracle Linux Yum Server Channels”](#).

Having subscribed your system to the appropriate channels on ULN or Oracle Linux yum server, upgrade your system. See [Section 3.4, “Upgrading Your System”](#).

After upgrading to UEK R4, you can replace any existing OFED packages with the Oracle-supported OFED packages, see [Section 3.5, “Installing the Oracle-Supported OFED Packages”](#).

3.2 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 <http://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 the *Oracle Linux Unbreakable Linux Network User's Guide* at http://docs.oracle.com/cd/E37670_01/index.html.

3.3 Enabling Access to Oracle Linux Yum Server Channels

At the Oracle Linux yum server at <http://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=http://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=http://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=http://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 ($basearch) - Unsupported
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/playground/latest/$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 ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/UEKR4/OFED/$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=http://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=http://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 ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL7/UEKR4/OFED/$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 <http://yum.oracle.com/>, from where you download a copy of a suitable repository file (<http://yum.oracle.com/public-yum-ol6.repo> or <http://yum.oracle.com/public-yum-ol7.repo>).

3.4 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 http://docs.oracle.com/cd/E52668_01/E54669/html/ol7-bootconf.html for more information on updating the default boot kernel on Oracle Linux 7.

See http://docs.oracle.com/cd/E37670_01/E41138/html/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 [Section 3.5, “Installing the Oracle-Supported OFED Packages”](#).

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

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>.

3.5 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.

If you want to install the packages from the Oracle Linux yum server:

- 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 http://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 transitioning from an earlier UEK version, such as UEK R2 or UEK R3, you must remove any existing OFED packages and clean any cached files from all previously enabled repositories. You do not need to perform this step if you are already running a version of UEK R4.

```
# 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*'
```

Clean all yum cached files from all enabled repositories:

```
# yum clean all
```

3. Install the OFED packages for UEK R4.

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

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

```
# chkconfig rdma on
```



Note

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, use the following command to 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-ofed-release
```

Downgrading UEK versions is not advisable except for testing purposes.

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

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

