

Oracle® Linux

**Release Notes for Unbreakable Enterprise Kernel Release 3
Quarterly Update 3**

ORACLE®

E55611-04
August 2017

Oracle Legal Notices

Copyright © 2014, 2017, 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 Quarterly Update 3 to the Unbreakable Enterprise Kernel Release 3. 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 3 with Oracle Linux 6 or Oracle Linux 7. Oracle recommends that you read this document before installing or upgrading the Unbreakable Enterprise Kernel Release 3.

Document generated on: 2017-08-22 (revision: 4698)

Table of Contents

Preface	v
1 New Features and Changes	1
1.1 Notable Changes	1
1.2 Driver Updates	1
1.2.1 Storage Adapter Drivers	1
1.2.2 Network Adapter Drivers	2
1.2.3 Miscellaneous Drivers	2
1.3 Technology Preview	2
1.4 Compatibility	2
1.5 Header Packages for Development	3
2 Fixed and Known Issues	5
2.1 Fixed Issues	5
2.2 Known Issues	6
3 Installation and Availability	15
3.1 Upgrading to the Latest UEK Release	15
3.2 Switching an Oracle Linux 6 System to UEK R3 (ULN)	16
3.3 Switching an Oracle Linux 6 System to UEK R3 (Oracle Yum Server)	17
3.4 Upgrading OFED Packages on Oracle Linux 6 Systems	19

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 3.

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.

Chapter 1 New Features and Changes

Table of Contents

1.1 Notable Changes	1
1.2 Driver Updates	1
1.2.1 Storage Adapter Drivers	1
1.2.2 Network Adapter Drivers	2
1.2.3 Miscellaneous Drivers	2
1.3 Technology Preview	2
1.4 Compatibility	2
1.5 Header Packages for Development	3

The Unbreakable Enterprise Kernel Release 3 (UEK R3) is Oracle's third major release of its heavily tested and optimized operating system kernel for Oracle Linux 6 and 7 on the x86-64 architecture. It is based on the mainline Linux kernel version 3.8.13.

The 3.8.13-44 release is the third quarterly update release for UEK R3. It includes security and bug fixes, as well as driver updates.

Oracle actively monitors upstream checkins and applies critical bug and security fixes to UEK R3.

UEK R3 uses the same versioning model as the mainline Linux kernel version. It is possible that some applications might not understand the 3.x versioning scheme. If an application does require a 2.6 context, you can use the `uname26` wrapper command to start it. However, regular Linux applications are usually neither aware of, nor affected by, Linux kernel version numbers.

1.1 Notable Changes

- Enhancements to the Huge Translation Lookaside Buffer (HugeTLB) to improve page-fault scalability.
- Added support for the `SO_REUSEPORT` socket option.
- Added support for XFS extended attributes.
- Added support for Intel Avoton and Haswell-EX CPUs.
- Updated `intel_idle` driver to support Intel Atom Processor C2000, N2000, and Bay Trail series.
- Bug fixes to support Oracle Linux guests running on Microsoft Azure or Hyper-V.
- Bug fixes for ext4, btrfs, and OCFS2 file systems.

1.2 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 by Oracle.

1.2.1 Storage Adapter Drivers

LSI

- MegaRAID Serial Attached SCSI (SAS) driver (`megaraid_sas`) updated to 06.803.02.00-rc1.
- MPT Fusion Serial Attached SCSI (SAS) 2.0 driver (`mpt2sas`) updated to 18.100.00.00.

- MPT Fusion Serial Attached SCSI (SAS) 3.0 driver ([mpt3sas](#)) updated to 04.100.00.00.

QLogic

- iSCSI host bus adapter (HBA) driver ([qla4xxx](#)) updated to 5.04.00.06.06.02-uek3.

1.2.2 Network Adapter Drivers

Intel

- PRO/1000 PCI Express Gigabit network adapter driver ([e1000e](#)) updated to 3.1.0.2-NAPI.
- Gigabit Ethernet network adapter driver ([igb](#)) updated to 5.2.5.
- 10 Gigabit PCI Express network adapter driver ([ixgbe](#)) updated to 3.21.2.
- 10 Gigabit PCI Express Virtual Function driver ([ixgbevf](#)) updated to 2.14.2.

QLogic

- 1/10 GbE Converged/Intelligent Ethernet adapter driver ([qlcnice](#)) updated to 5.3.59.

1.2.3 Miscellaneous Drivers

- Added the Random Number Generator (RNG) driver ([tpm-rng](#)) for Trusted Platform Module (TPM) devices (Oracle Linux 7 only).

1.3 Technology Preview

The following features included in the Unbreakable Enterprise Kernel Release 3 are still under development, but are made available for testing and evaluation purposes.

- **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.

- **Transcendent memory**

Transcendent Memory (*tmem*) provides a new approach for improving the utilization of physical memory in a virtualized environment by claiming underutilized memory in a system and making it available where it is most needed. From the perspective of an operating system, tmem is fast pseudo-RAM of indeterminate and varying size that is useful primarily when real RAM is in short supply. To learn more about this technology and its use cases, see the Transcendent Memory project page at <http://oss.oracle.com/projects/tmem/>.

1.4 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 3 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 R3 will remain unchanged in all subsequent updates to the initial release. In this release, there are changes to the kernel ABI relative to UEK R2 that require recompilation of third-party kernel modules on the system. Before installing UEK R3, verify its support status with your application vendor.

1.5 Header Packages for Development

The `kernel-headers` packages provide 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.

Chapter 2 Fixed and Known Issues

Table of Contents

2.1 Fixed Issues	5
2.2 Known Issues	6

This chapter describes the fixed and known issues for the Unbreakable Enterprise Kernel Release 3.



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.

bnx2x Driver

- When using the `bnx2x` driver in a bridge, the bug that required you to disable Transparent Packet Aggregation (TPA) by including the statement `options bnx2x disable_tpa=1` in `/etc/modprobe.conf` has been fixed. (Bug ID 14626070)

btrfs

- The `btrfs subvolume delete` command no longer displays "Directory not empty" error when you try to delete the default subvolume. (Bug ID 17661944)

CIFS

- If an error occurs on a mounted CIFS file system, the error no longer prevents a Linux system from shutting down. (Bug ID 18447168)

DTrace

- DOF ELF objects generated using the `-G` option to the `dtrace` command no longer get the `EXEC` flag set for the stack. (Bug ID 19217436)
- Order-only prerequisites are now supported in the kernel module build process when generating the `sdtstub` file for SDT probes in modules. (Bug ID 18906444)

ext4

- The bug that caused hangs with intensive direct random writing to files opened for synchronous I/O (`O_DIRECT`, `O_SYNC`) with unwritten extent conversions has been fixed. (Bug ID 18389351)

OCFS2

- OCFS2 no longer hangs as the result of lost pending messages in the TCP queue, when the idle timeout triggers a TCP disconnect and reconnect. (Bug ID 18363396)

XFS

- Attempting to add many extended attributes to a file in an XFS file system under UEK R3 no longer results in a kernel panic or a system reboot. (Bug ID 18504299).

- Reading from or writing to a special file, such as `/dev/null`, that resides on an XFS file system no longer results in an error. (Bug ID 19207712)

2.2 Known Issues

This section describes the known issues in this update.

Updating Oracle Linux Fails if the `kernel-uek` Package Cannot Be Updated

By default, the installation of Oracle Linux includes the `dtrace-modules` package for UEK R3. This package requires a specific `kernel-uek` version. However, a `yum update` fails if it cannot update the `kernel-uek` package when the `installonly_limit` of three updates is reached as the `dtrace-modules` package does not allow the existing `kernel-uek` packages to be removed. `yum` displays an error message similar to the following example:

```
--> Finished Dependency Resolution
Error: Package: kernel-uek-debug-3.8.13-55.1.1.el6uek.x86_64 (public_ol6_UEKR3_latest)
       Requires: kernel-firmware = 3.8.13-55.1.1.el6uek
```

The workaround is to remove any existing `dtrace-modules` packages before updating the `kernel-uek` package, for example:

```
# for package in `rpm -qa | grep dtrace-modules`; do yum remove -y $package; done
```

When you have removed all `dtrace-modules` packages, you can update Oracle Linux, including the UEK R3 kernel. If you want to use DTrace with UEK R3, reinstall the `dtrace-modules` package for the current kernel:

```
# yum install dtrace-modules-`uname -r`
```

(Bug ID 21669543)

ACPI

The following messages indicate that the BIOS does not present a suitable interface, such as `_PSS` or `_PPC`, that the `acpi-cpufreq` module requires:

```
kernel: powernow-k8: this CPU is not supported anymore, using acpi-cpufreq instead.
modprobe: FATAL: Error inserting acpi_cpufreq
```

There is no known workaround for this error. (Bug ID 17034535)

btrfs

- If you use the `--alloc-start` option with `mkfs.btrfs` to specify an offset for the start of the file system, the size of the file system should be smaller but this is not the case. It is also possible to specify an offset that is higher than the device size. (Bug ID 16946255)
- The usage information for `mkfs.btrfs` reports `raid5` and `raid6` as possible profiles for both data and metadata. However, the kernel does not support these features and cannot mount file systems that use them. (Bug ID 16946303)
- 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)
- Converting an existing `ext2`, `ext3`, or `ext4` root file system to `btrfs` does not carry over the associated security contexts that are stored as part of a file's extended attributes. With SELinux enabled and set

to enforcing mode, you might experience many permission denied errors after reboot, and the system might be unbootable. To avoid this problem, enforce automatic file system relabeling to run at bootup time. To trigger automatic relabeling, create an empty file named `.autorelabel` (for example, by using `touch`) in the file system's `root` directory before rebooting the system after the initial conversion. The presence of this file instructs SELinux to recreate the security attributes for all files on the file system. If you forget to do this and rebooting fails, either temporarily disable SELinux completely by adding `selinux=0` to the kernel boot parameters, or disable enforcing of the SELinux policy by adding `enforcing=0`. (Bug ID 13806043)

- Commands such as `du` can show inconsistent results for file sizes in a btrfs file system when the number of bytes that is under delayed allocation is changing. (Bug ID 13096268)
- 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>.
- Btrfs has a limit of 237 or fewer hard links to a file from a single directory. The exact limit depends on the number of characters in the file name. The limit is 237 for a file with up to eight characters in its file name; the limit is lower for longer file names. Attempting to create more than this number of links results in the error `Too many links`. You can create more hard links to the same file from another directory. Although the limitation of the number of hard links in a single directory has been increased to 65535, the version of `mkfs.btrfs` that is provided in the `btrfs-progs` package does not yet support the compatibility flag for this feature. (Bug ID 16285431)
- If you run the `btrfs quota enable` command on a non-empty file system, any existing files do not count toward space usage. Removing these files can cause usage reports to display negative numbers and the file system to be inaccessible. The workaround is to enable quotas immediately after creating the file system. If you have already written data to the file system, it is too late to enable quotas. (Bug ID 16569350)
- The `btrfs quota rescan` command is not currently implemented. The command does not perform a rescan and returns without displaying any message. (Bug ID 16569350)
- 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. (Bug ID 16609467)
- If you run `btrfsck --init-csum-tree` on a file system and then run a simple `btrfsck` on the same file system, the command displays a Backref mismatch error that was not previously present. (Bug ID 16972799)
- 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 subvolume delete` command may result in a "Directory not empty" error. This error message is incorrect. The actual reason that the subvolume cannot be deleted is that the subvolume is configured as the default subvolume. The default subvolume is the subvolume that is mounted when no subvolume is specified with the `mount` command. Before you can delete the subvolume, you need to configure a different default subvolume using the `btrfs subvolume set-default` command. (Bug ID 17661944)
- Defragmentation can break data block sharing. Due to the copy-on-write design of btrfs, snapshots initially share the same data blocks of the original subvolume. However, when either the snapshot or the subvolume is defragmented, this sharing can be undone, resulting in a higher disk space usage.

CPU Microcode Update Failures on PVM or PVHVM Guests

When running Oracle Linux 6 with UEK R3, you might see error messages in `dmesg` or `/var/log/messages` similar to this one:

```
microcode: CPU0 update to revision 0x6b failed.
```

You can ignore this warning. You do not need to upgrade the microcode for virtual CPUs as presented to the guest. (Bug ID 12576264, 13782843)

DHCP Lease is not Obtained at Boot Time

If DHCP lease negotiation takes more than 5 seconds at boot time, the following message is displayed:

```
ethX: failed. No link present. Check cable?
```

If the `ethtool ethX` command confirms that the interface is present, edit `/etc/sysconfig/network-scripts/ifcfg-ethX` and set `LINKDELAY=N`, where `N` is a value greater than 5 seconds (for example, 30 seconds). Alternatively, use NetworkManager to configure the interface. (Bug ID 16620177)

dm-nfs Module Obsolete

In UEK R2, the `dm-nfs` module provided the ability to create a loopback device for a mounted NFS file or file system. For example, the feature allowed you to create the shared storage for an Oracle 3 VM cluster on an NFS file system. The `dm-nfs` module provided direct I/O to the server and bypassed the `loop` driver to avoid an additional level of page caching. The `dm-nfs` module is not provided with UEK R3. The `loop` driver can now provide the same I/O functionality as `dm-nfs` by extending the AIO interface to perform direct I/O. To create the loopback device, use the `losetup` command instead of `dmsetup`.

DTrace

- Using `kill -9` to terminate `dtrace` can leave breakpoints outstanding in processes being traced, which might sooner or later kill them.
- 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
```

ERST Message

You can safely ignore the following message that might be displayed in `syslog` or `dmesg`:

```
ERST: Failed to get Error Log Address Range.
```

The message indicates that the system BIOS does not support an Error Record Serialization Table (ERST). (Bug ID 17034576)

Ext4

The inline data feature that allows the data of small files to be stored inside their inodes is not yet available. The `-O inline_data` option to the `mkfs.ext4` and `tune2fs` commands is not supported. (Bug ID 17210654)

Firmware Warning Message

You can safely ignore the following firmware warning message that might be displayed on some Sun hardware:

```
[Firmware Warn]: GHES: Poll interval is 0 for generic hardware error source:
1, disabled.
```

(Bug ID 13696512)

Huge Pages

One-gigabyte (1 GB) huge pages are not currently supported for the following configurations:

- HVM guests
- PV guests
- Oracle Database

Two-megabyte (2 MB) huge pages have been tested and work with these configurations.

(Bug ID 17299364, 17299871, 17271305)

I/O Scheduler

The Unbreakable Enterprise Kernel uses the [deadline](#) scheduler as the default I/O scheduler. For the Red Hat Compatible Kernel, the default I/O scheduler is the [cfq](#) scheduler.

ioapic Failure Messages

You can safely ignore messages such as `ioapic: probe of 0000:00:05.4 failed with error -22`. Such messages are the result of the `ioapic` driver attempting to re-register I/O APIC PCI devices that were already registered at boot time. (Bug ID 17034993)

InfiniBand

- You might see the following warning messages if you use the `ibportstate disable` command to disable a switch port:

```
ibwarn: [2696] _do_madrpc: recv failed: Connection timed out
ibwarn: [2696] mad_rpc: _do_madrpc failed; dport (Lid 38)
ibportstate: iberror: failed: smp set portinfo failed
```

You can safely ignore these warnings. (Bug ID 16248314)

- The Internet Protocol over InfiniBand (IPoIB) driver supports the use of either connected mode or datagram mode with an interface, where datagram mode is the default mode. Changing the mode of an InfiniBand interface by echoing either `connected` or `datagram` to `/sys/class/net/ibN/mode` is not supported. It is also not possible to change the mode of an InfiniBand interface while it is enabled.

To change the IPoIB mode of an InfiniBand interface:

1. Edit the `/etc/sysconfig/network-scripts/ifcfg-ibN` configuration file, where `N` is the number of the interface:
 - To configure connected mode, specify `CONNECTED_MODE=yes` in the file.

- To configure datagram mode, either specify `CONNECTED_MODE=no` in the file or do not specify this setting at all (datagram mode is enabled by default).



Note

Before saving your changes, make sure that you have not specified more than one setting for `CONNECTED_MODE` in the file.

2. To enable the specified mode on the interface, use the following commands to take down the interface and bring it back up:

```
# ifdown ibN
# ifup ibN
```

(Bug ID 17479833)

- When the `rds_ib_srq` parameter for the `rds_rdma` module is enabled and the module is in use (for example when running the `rds-stress` tool), restarting the `rdma` service (which reloads the `rds_rdma` module) generates error messages visible in `dmesg` or `/var/log/messages`. (Bug ID 18243427)

Linux Containers (LXC)

- The correct operation of containers might require that you completely disable SELinux on the host system. For example, SELinux can interfere with container operation under the following conditions:
 - Running the `halt` or `shutdown` command from inside the container hangs the container or results in a `permission denied` error. (An alternate workaround is to use the `init 0` command from inside the container to shut it down.)
 - Setting a password inside the container results in a `permission denied` error, even when run as `root`.
 - You want to allow `ssh` logins to the container.

To disable SELinux on the host:

1. Edit the configuration file for SELinux, `/etc/selinux/config` and set the value of the `SELINUX` directive to `disabled`.
 2. Reboot the host.
- The `root` user in a container can affect the configuration of the host system by setting some `/proc` entries. (Bug ID 17190287)
 - Using `yum` to update packages inside the container that use `init` scripts can undo changes made by the Oracle template.
 - Migrating live containers (`lxc-checkpoint`) is not yet supported.
 - Oracle Database is not yet supported for use with Linux Containers. The following information is intended for those who want to experiment with such a configuration.

The following `/proc` parameter files may only be set on the host and not for individual containers:

- `/proc/sys/fs/aio-max-nr`
- `/proc/sys/net/core/rmem_default`

- `/proc/sys/net/core/rmem_max`
- `/proc/sys/net/core/wmem_default`
- `/proc/sys/net/core/wmem_max`
- `/proc/sys/net/ipv4/ip_local_port_range`

Setting the parameters in the host to the Oracle recommended values sets them for all containers and allows the Oracle database to run in a container. For more information, see [Configuring Kernel Parameters and Resource Limits](#). (Bug ID 17217854)

- If a `Device or resource busy - failed to set memory.use_hierarchy to 1` error message is displayed when you start a container, edit the `/etc/cgconfig.conf` file and add the following:

```
group . {
memory {
memory.use_hierarchy = "1";
}
}
```

(Bug ID 19237222)

Missing Documentation for PCIe Kernel Parameters

The following documentation for some PCIe kernel parameters is not included in the `kernel-parameters.txt` file in the `kernel-uek-doc` package:

<code>pcie_bus_tune_off</code>	Disable PCIe MPS (Max Payload Size) tuning and use the BIOS-configured MPS defaults.
<code>pcie_bus_safe</code>	Set every device's MPS to the largest value supported by all devices below the root complex.
<code>pcie_bus_perf</code>	Set device MPS to the largest allowable MPS based on its parent bus. Also set MRRS (Max Read Request Size) to the largest supported value (no larger than the MPS that the device or bus can support) for best performance.
<code>pcie_bus_peer2peer</code>	Set every device's MPS to 128B, which every device is guaranteed to support. This configuration allows peer-to-peer DMA between any pair of devices, possibly at the cost of reduced performance. This also guarantees that hot-added devices will work.

(Bug ID 18706504)

RDMA Does Not Load the `mlx4_ib` Module

If you enable the OFED stack and the RDMA service but the version of the RDMA package is lower than `rdma-3.10-3.0.2.el6`, the RDMA service does not load the `mlx4_ib` module automatically.

To configure the RDMA service to load the `mlx4_ib` module at boot time:

1. Edit `/etc/rdma/rdma.conf` and set the entry `MLX4_LOAD=yes` in this file.
2. To make the change take effect, restart the RDMA service or reboot the system.

sched_yield() Settings for CFS

For the Unbreakable Enterprise Kernel, `kernel.sched_compat_yield=1` is set by default. For the Red Hat Compatible Kernel, `kernel.sched_compat_yield=0` is used by default.

Slow Performance With Multipath Devices

Starting with UEK R2, the device mapper has had the capability to check whether the underlying storage device has advertised the need to flush the data that resides in the device's cache to its non-volatile storage. For a data integrity operation, such as `fsync` and `sync`, the operation will now need to include the time to flush the device's cache (if it is advertised). Such an operation will appear to be slower when compared to a previous older kernel, however this is the correct behavior. (Bug ID 17823743)

Soft Lockup Errors When Booting

When upgrading or installing the UEK R3 kernel on fast hardware, usually with SAN storage attached, the kernel can fail to boot and `BUG: soft lockup` messages are displayed in the console log. The workaround is to increase the baud rate from the default value of 9600 by amending the kernel boot line in `/boot/grub/grub.conf` to include an appropriate console setting, for example:

```
console=ttyS0,115200n8
```

A value of 115200 is recommended as smaller values such as 19200 are known to be insufficient for some systems (for example, see http://docs.oracle.com/cd/E19045-01/blade.x6220/820-0048-18/sp.html#0_pgfid-1002490). If the host implements an integrated system management infrastructure, such as ILOM on Sun and Oracle systems or iLO on HP systems, configure the integrated console baud rate to match the setting for the host system. Otherwise, the integrated console is likely to display garbage characters. (Bug ID 17064059, 17252160)

Transparent Huge Pages

This release removes the Transparent Huge Pages (THP) feature. Following extensive benchmarking and testing, Oracle found that THP caused a performance degradation of between 5 and 10% for some workloads. This performance degradation was a result of a slower memory allocator code path being used even when the applications were not using THP. When the fact that huge pages are not swappable was taken into account, the positive effect that THP should provide was outweighed by its negative effects.

After installing this UEK release, you cannot enable THP (for example, by specifying kernel boot parameters). The THP settings under `/sys/kernel/mm/transparent_hugepage` have also been removed. A future update might contain an updated THP implementation which resolves the performance issue.



Note

This change does not affect support for applications that use explicit huge pages (for example, Oracle Database).

(Bug ID 16823432)

User Namespaces

The kernel functionality (`CONFIG_USER_NS`) that allows unprivileged processes to create namespaces for users inside which they have root privileges is not currently implemented because of a clash with the implementation of XFS. This functionality is primarily intended for use with Linux Containers. As a result, the `lxc-checkconfig` command displays `User namespace: missing`. (Bug ID 16656850)

Virtualization

- When booting UEK R3 as a PVHVM guest, you can safely ignore the following kernel message:

```
register_vcpu_info failed:
    err=-38
```

(Bug ID 13713774)

- Under Oracle VM Server 3.1.1, migrating a PVHVM guest that is running the UEK R3 kernel causes a disparity between the date and time as displayed by `date` and `hwclock`. To prevent this from occurring, upgrade to Oracle VM Server 3.2.1 or later. The workaround post migration is either to run the command `hwclock --hctosys` on the guest or to reboot the guest. (Bug ID 16861041)
- On virtualized systems that are built on Xen version 3, including all releases of Oracle VM 2 including 2.2.2 and 2.2.3, disk synchronization requests for ext3 and ext4 file systems result in journal corruption with kernel messages similar to the following being logged:

```
blkfront: barrier: empty write xvda op failed
blkfront: xvda: barrier or flush: disabled
```

In addition, journal failures such as the following might be reported:

```
Aborting journal on device xvda1
```

The workaround is to add the mount option `barrier=0` to all ext3 and ext4 file systems in the guest VM before upgrading to UEK R3. For example, you would change a mount entry such as:

```
UUID=4e4287b1-87dc-47a8-b69a-075c7579eaf1 / ext3 defaults 1 1
```

so that it reads:

```
UUID=4e4287b1-87dc-47a8-b69a-075c7579eaf1 / ext3 defaults,barrier=0 1 1
```

This issue does not apply to Xen 4 based systems, such as Oracle VM 3. (Bug ID 17310816)

X.509 Certificates for Module Verification

The system reports a message similar to the following if there is a problem loading an in-kernel X.509 module verification certificate at boot time:

```
Loading module verification certificates
X.509: Cert 0c21da3d73dcdabaffc799e3d26f3c846a3afdc43 is not yet valid
MODSIGN: Problem loading in-kernel X.509 certificate (-129)
```

This error occurs because the hardware clock lags behind the system time as shown by `hwclock`, for example:

```
# hwclock
Tue 20 Aug 2013 01:41:40 PM EDT -0.767004 seconds
```

The solution is to set the hardware clock from the system time by running the following command:

```
# hwclock --systohc
```

After correcting the hardware clock, no error should be seen at boot time, for example:

```
Loading module verification certificates
MODSIGN: Loaded cert 'Slarti: Josteldalsbreen signing key:
0c21da3d73dcdabaffc799e3d26f3c846a3afdc43'
```

(Bug ID 17346862)

XFS

- In some circumstances, `xfsdump` can fail when the file system to be backed up is specified as a mount point, for example:

```
xfsdump: ERROR: /mnt/myxfs/ does not identify a file system
```

The workaround is to specify the file system by its device name, for example `/dev/sdb`. (Bug ID 18483275)

- In a virtual machine, the intensive direct I/O on files which are being fragmented by `fallocate` and `punch_hole` in parallel may result, in some circumstances, in data corruption with the error `Invalid or incomplete multibyte or wide character`. (Bug ID 18711409)
- The following message might be displayed when you use the `xfs_growfs` command:

```
xfs_growfs: cannot find mount point for path `mount_point': Success
```

The message is caused by an incorrect mounted project path listed in the XFS project quota initialization file, `/etc/projects`. If you are not using project quotas, remove the `/etc/projects` file. (Bug ID 18886520)

Chapter 3 Installation and Availability

Table of Contents

3.1 Upgrading to the Latest UEK Release	15
3.2 Switching an Oracle Linux 6 System to UEK R3 (ULN)	16
3.3 Switching an Oracle Linux 6 System to UEK R3 (Oracle Yum Server)	17
3.4 Upgrading OFED Packages on Oracle Linux 6 Systems	19

Unbreakable Enterprise Kernel Release 3 (UEK R3) is only supported on the x86-64 architecture.

Oracle Linux 7

UEK R3 is the default boot kernel for fresh installations of Oracle Linux 7.

Oracle Linux 6

You can install UEK R3 on Oracle Linux 6 Update 4 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.

Starting with Oracle Linux 6 Update 5 for x86-64, UEK R3 is the default boot kernel for fresh installations of Oracle Linux 6.

For systems that are currently running a previous version of the Unbreakable Enterprise Kernel (UEK R2) or the Red Hat compatible kernel (RHCK), you can switch to UEK R3 at any time. For details, see:

- [Section 3.2, “Switching an Oracle Linux 6 System to UEK R3 \(ULN\)”](#)
- [Section 3.3, “Switching an Oracle Linux 6 System to UEK R3 \(Oracle Yum Server\)”](#)

Kernel Source Code

The kernel's source code is available from a public git source code repository at:
<https://oss.oracle.com/git/?p=linux-uek3-3.8.git>.

3.1 Upgrading to the Latest UEK Release

For systems that are currently running UEK R3, you upgrade to the latest UEK release as follows:

1. Ensure the system is configured to receive updates from the correct channels or repositories.

For Oracle Linux 6, subscribe to either the `o16_x86_64_UEKR3_latest` channel on ULN, or the `o16_UEKR3_latest` repository on the Oracle Yum Server.

For Oracle Linux 7, subscribe to either the `o17_x86_64_UEKR3` channel on ULN, or the `o17_UEKR3` repository on the Oracle public yum server.

2. Upgrade all packages on the system, including kernel packages.

```
# yum update
```

By default, the boot manager automatically enables the most recent kernel version so you do not need to change your GRUB or GRUB 2 configuration.

3. Reboot the system.

Oracle Linux 6:

```
# shutdown -r now
```

Oracle Linux 7:

```
# systemctl reboot
```

3.2 Switching an Oracle Linux 6 System to UEK R3 (ULN)

If you have a subscription to Oracle Unbreakable Linux support, you can obtain the packages for Unbreakable Enterprise Kernel Release 3 (UEK R3) by registering your system with the Unbreakable Linux Network (ULN) and subscribing it to additional channels.

Before you begin:

- Ensure the system is registered with ULN.

For information about registering with ULN, see the [Oracle Linux Unbreakable Linux Network User's Guide](#).

- Check that the system meets the requirements for installing UEK R3.

For details, see [Chapter 3, Installation and Availability](#).

- Remove some of the existing OpenFabrics Enterprise Distribution (OFED) packages.

You only need to do this if you have installed any OFED packages on your system and you want to replace them with the latest OFED tools packages. The packages have to be removed manually, see [Section 3.4, "Upgrading OFED Packages on Oracle Linux 6 Systems"](#).

To Switch a System to UEK R3:

1. Using a browser, log in at <http://linux.oracle.com> with the ULN user name and password that you used to register the system.
2. On the Systems tab, click the link named of your 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 in the **Available Channels** list and click the right arrow to move the channel to the **Subscribed Channels** list.

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

Channel Name and Label	Description
Oracle Linux 6 Latest (x86_64) ol6_x86_64_latest	All packages released for Oracle Linux 6 (x86_64) including the latest errata packages. (x86_64).
Unbreakable Enterprise Kernel Release 3 for Oracle Linux 6 (x86_64) - Latest ol6_x86_64_UEKR3_latest	Latest packages for Unbreakable Enterprise Kernel Release 3 for Oracle Linux 6 (x86_64). Contains the <code>kernel-uek*</code> , <code>dtrace-modules-*</code> , <code>libdtrace-*</code> , and <code>uname26</code> packages.

Channel Name and Label	Description
Oracle Linux 6 Dtrace Userspace Tools (x86_64) - Latest	The latest DTrace userspace tools for Oracle Linux 6 (x86_64).
ol6_x86_64_Dtrace_userspace_latest	Contains the <code>dtrace-utils*</code> packages.
OFED supporting tool packages for Unbreakable Enterprise Kernel on Oracle Linux 6 (x86_64)	Latest OpenFabrics Enterprise Distribution (OFED) supporting tools for the Unbreakable Enterprise Kernel (UEK) on Oracle Linux 6 (x86_64).
ol6_x86_64_ofed_UEK	
HA Utilities for MySQL and Oracle Linux 6 (x86_64)	Management Utilities for MySQL HA with Oracle Linux 6.
ol6_x86_64_mysql-ha-utils	Contains the <code>drbd84-utils</code> package.

As a minimum, you should subscribe the system to the [ol6_x86_64_latest](#) and the [ol6_x86_64_UEKR3_latest](#) channels. If required, you can also add the channels for the DTrace, OFED, and DRBD packages.



Caution

Take care not to select the [ol6_x86_64_UEK_BETA](#) channel.

Because you are switching to the latest UEK kernel, you no longer need to subscribe the system to the previous UEK R2 ([ol6_x86_64_UEK_latest](#)) channel.

- When you have finished selecting channels, click **Save Subscriptions** and log out of ULN.
- Log in as `root` on the system.
- Upgrade all packages on the system, including kernel packages.

```
# yum update
```

By default, the boot manager automatically enables the most recent kernel version so you do not need to change your GRUB configuration.

- Reboot the system.

```
# shutdown -r now
```

3.3 Switching an Oracle Linux 6 System to UEK R3 (Oracle Yum Server)

If your system is not registered with ULN, you can obtain most of the packages for Unbreakable Enterprise Kernel Release 3 (UEK R3) from Oracle Yum Server by subscribing it to additional repositories.

Before you begin:

- Check that the system meets the requirements for installing UEK R3.

For details, see [Chapter 3, Installation and Availability](#).

- Remove some of the existing OpenFabrics Enterprise Distribution (OFED) packages.

You only need to do this if you have installed any OFED packages on your system and you want to replace them with the latest OFED tools packages. The packages have to be removed manually, see [Section 3.4, “Upgrading OFED Packages on Oracle Linux 6 Systems”](#).

To Switch a System to UEK R3:

1. Log in as `root` on the system.
2. Change directory to `/etc/yum.repos.d`.

```
# cd /etc/yum.repos.d
```



Note

This assumes that `yum` on your system is configured to find repository files in the default `/etc/yum.repos.d` directory.

3. Download the Oracle Linux 6 repository configuration file, <http://yum.oracle.com/public-yum-ol6.repo>.

For example:

```
# wget http://yum.oracle.com/public-yum-ol6.repo
```

The `/etc/yum.repos.d` directory is updated with the repository configuration file.

4. Enable the required repositories by editing the `public-yum-ol6.repo` file.

You enable or disable repositories in the file by setting the value of the `enabled` directive to 1 or 0 as required.

The kernel image and user-space packages are available on the following Oracle Yum Server repositories:

Repository	Description
<code>ol6_latest</code>	All packages released for Oracle Linux 6 (x86_64) including the latest errata packages.
<code>ol6_UEKR3_latest</code>	Latest packages for Unbreakable Enterprise Kernel Release 3 for Oracle Linux 6. Contains the <code>kernel-uek*</code> , <code>dtrace-modules-*</code> , <code>libdtrace-*</code> , and <code>uname26</code> packages.
<code>ol6_ofed_UEK</code>	Latest OpenFabrics Enterprise Distribution (OFED) supporting tools for the Unbreakable Enterprise Kernel (UEK) on Oracle Linux 6 (x86_64).

As a minimum, you should enable the `ol6_latest` and the `ol6_UEKR3_latest` repositories.



Note

The DTrace and DRBD utility packages are not available on Oracle Yum Server.

Because you are switching to the latest UEK kernel, you can disable the previous UEK R2 (`ol6_UEK_latest`) repository.

In the following example, the `ol6_UEKR3_latest` repository is enabled, and the `ol6_UEK_latest` repository is disabled:

```
[ol6_UEKR3_latest]
name=Latest Unbreakable Enterprise Kernel for Oracle Linux $releasever ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/UEKR3/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
```

5. Upgrade all packages on the system, including kernel packages.

```
# yum update
```

By default, the boot manager automatically enables the most recent kernel version so you do not need to change your GRUB configuration.

6. Reboot the system.

```
# shutdown -r now
```

3.4 Upgrading OFED Packages on Oracle Linux 6 Systems

If you have enabled the `ol6_ofed_UEK` channel, you must remove any existing OFED packages for the 32-bit x86 architecture before you can upgrade the remaining OFED packages on your system. You must also completely remove and reinstall the `ibutils` packages. The latest version of the `ibutils` package no longer depends on an `ibutils-libs` package as the libraries are now included in `ibutils` itself.

Use the following command to remove any non-upgradable packages for the x86 architecture:

```
# rpm -e infiniband-diags \
libibcm \
libibcm-devel \
libibmad \
libibmad-devel \
libibumad \
libibumad-devel \
libibverbs \
libibverbs-devel \
libmlx4 \
librdmacm \
librdmacm-devel \
opensm-devel \
opensm-libs \
ibacm-devel
```

Use the following commands to remove the existing `ibutils` and `ibutils-libs` packages and install the new `ibutils` package:

```
# rpm -e ibutils ibutils-libs
# yum install ibutils
```

