

Oracle® Linux

**Release Notes for Unbreakable Enterprise Kernel Release 2
Quarterly Update 5**

ORACLE®

E48381-07
August 2015

Oracle Legal Notices

Copyright © 2013, 2015, 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.

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.

Abstract

This document contains information on Quarterly Update 5 to the Unbreakable Enterprise Kernel Release 2. 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 2 with Oracle Linux 5 and Oracle Linux 6. Oracle recommends that you read this document before installing or upgrading the Unbreakable Enterprise Kernel Release 2.

Document generated on: 2015-08-11 (revision: 3189)

Table of Contents

Preface	v
1 New Features and Changes	1
1.1 Notable Changes	1
1.2 Xen Improvements	1
1.3 Driver Updates	1
1.3.1 Storage Adapter Drivers	1
1.3.2 Network Adapter Drivers	2
1.4 Technology Preview	3
1.5 Compatibility	4
2 Fixed and Known Issues	5
2.1 Fixed Issues	5
2.2 Known Issues	5
3 Installation and Availability	13

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 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.

Chapter 1 New Features and Changes

The Unbreakable Enterprise Kernel Release 2 (UEK R2) is Oracle's second major release of its heavily tested and optimized operating system kernel for Oracle Linux 5 and Oracle Linux 6. See the initial *Oracle Unbreakable Enterprise Kernel Release 2 Release Notes* (<https://oss.oracle.com/ol6/docs/RELEASE-NOTES-UEK2-en.html>) for a detailed description of the differences between UEK R2 and the first version of the Unbreakable Enterprise Kernel.

The 2.6.39-400.209.1 release is the fifth quarterly driver update release, which also includes bug and security fixes.



Note

The version number displayed by the kernel and on the RPM packages is 2.6.39. This was done to avoid potential breakage of certain low-level utilities of the Oracle Linux distribution (also known as the *plumbing*) that potentially cannot cope with the new 3.x version scheme. Regular Linux applications are usually neither aware of nor affected by Linux kernel version numbers.

1.1 Notable Changes

- OpenFabrics Alliance (OFED) stack support for Oracle X4-2, X4-2L, and Netra X3-2 servers with Mellanox ConnectX-3 InfiniBand Host Channel Adapters.
- Oracle Virtual Networking virtual host adapter and network drivers to support Oracle SDN (Software Defined Network), previously known as Xsigo Fabric Accelerator.
- LSI Fusion-MPT SAS 3.0 driver to support up to 12 Gb/s host controllers.

1.2 Xen Improvements

Relative to quarterly update 4, several improvements have been incorporated into the Unbreakable Enterprise Kernel to support Xen usage:

- Fixes for EDD, x2apic, XenBus, and PVHVM vCPU hotplug issues.
- The indirect-descriptor feature, which increases throughput and reduces latency for block I/O.

1.3 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.

1.3.1 Storage Adapter Drivers

Adaptec

- AACRAID driver (`aacraid`) updated to 1.2-0[30200]-ms.

Cisco

- Cisco FCoE HBA Driver (`fnic`) updated to 1.5.0.45.

Emulex

- Fibre Channel HBA driver (`lpfc`) updated to 0:8.3.7.26.3p.

Version 0:8.3.7.10.4p and later of the `lpfc` driver now support the following FCoE adapters in addition to the adapters that were previously supported:

- Emulex LPe16002-M6-O 2-Port 10Gb FCoE Adapter
- Emulex LPem16002-M6-O 2-Port 10Gb FCoE Adapter
- Emulex LPe16202-X 2-Port 10Gb FCoE Adapter
- Sun Storage 16Gb FC ExpressModule Universal HBA-Emulex (#7101689 and #7101690)
- Sun Storage 16Gb FC PCIe Universal HBA-Emulex (#7101683 and #7101684)

Intel

- C600 serial attached SCSI (SAS) module (`iscsi`) updated to 1.0.0.

LSI

- LSI Fusion-MPT base driver (`mptbase`) updated to 4.28.20.02.
- LSI Fusion-MPT `ioctl` driver (`mptctl`) updated to 4.28.20.02.
- LSI Fusion-MPT Fibre Channel host driver (`mptfc`) updated to 4.28.20.02.
- LSI Fusion-MPT IP Over Fibre Channel driver (`mptlan`) updated to 4.28.20.02.
- LSI Fusion-MPT SAS driver (`mptsas`) updated to 4.28.20.02.
- LSI Fusion-MPT SCSI host driver (`mptscsih`) updated to 4.28.20.02.
- LSI Fusion-MPT SPI host driver (`mptspi`) updated to 4.28.20.02.
- LSI Fusion-MPT SAS 2.0 driver (`mpt2sas`) updated to 17.00.01.00.
- LSI Fusion-MPT SAS 3.0 driver (`mpt3sas`) version 3.00.00.00 added. Supports up to 12 Gb/s host controllers.

MegaRAID

- MegaRAID SAS driver (`megaraid_sas`) updated to 06.600.18.00.

Oracle

- Core services module driver (`xscore`) version 6.0.r7269 added. Required by all other Oracle Virtual Networking modules.
- Virtual HBA driver (`xsvhba`) version 6.0.r7269 added.

QLogic

- iSCSI driver (`qla4xxx`) updated to 5.03.00.03.06.02-uek2. Supports Open-iSCSI.

1.3.2 Network Adapter Drivers

Broadcom

- Tigon3 Ethernet adapter driver (`tg3`) updated to 3.131d.

Emulex

- Blade Engine 2 10Gbps adapter driver ([be2net](#)) updated to 4.6.63.0u.

Intel

- PRO/1000 PCI-Express Gigabit network adapter driver ([e1000e](#)) updated to 2.4.14.
- Gigabit Ethernet network adapter driver ([igb](#)) updated to 4.3.0.
- Gigabit Linux driver ([igbvf](#)) updated to 2.3.2. Provides 82576-based virtual function devices on kernels that support Single Root I/O Virtualization (SR-IOV).
- 10 Gigabit PCI-Express network adapter driver ([ixgbe](#)) updated to 3.15.1.

Oracle

- Core services module driver ([xscore](#)) version 6.0.r7269 added. Required by all other Oracle Virtual Networking modules.
- Virtual Ethernet driver ([xve](#)) added.
- Virtual NIC driver ([xsvnic](#)) version 0.316.0.r7269 added.

QLogic

- 1/10 GbE Converged/Intelligent Ethernet Adapter driver ([qlcnic](#)) updated to 5.2.43.

1.4 Technology Preview

The following features included in the Unbreakable Enterprise Kernel Release 2 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. (This technology preview feature is available for Oracle Linux 6 only.)

- **DTrace**

DTrace is a comprehensive dynamic tracing framework that was initially developed for the Oracle Solaris operating system and which is being ported to Linux by Oracle. DTrace provides a powerful infrastructure to permit administrators, developers, and service personnel to concisely answer arbitrary questions about the behavior of the operating system and user programs in real time. DTrace feature previews are published as a separate set of kernel packages in UEK R2. DTrace support is integrated with the kernel in the UEK Release 3 (UEK R3) distribution. (This technology preview feature is available for Oracle Linux 6 only.)

- **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.

- **Linux Containers (lxc)**

Based on the Linux Cgroups and name spaces functionality, containers allow you to safely and securely run multiple applications or instances of an operating system on a single host without risking them interfering with each other. Containers are lightweight and resource-friendly, which saves both rack space and power. In order to get started with containers, you need to install the `lxc` package, which is included in the package repository of the Unbreakable Enterprise Kernel. (This technology preview feature is available for Oracle Linux 6 only.)

- **Transcendent memory**

Transcendent Memory (tmem for short) 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/>.

1.5 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 2 and no re-certifications are needed for RHEL certified applications.

From UEK R2 quarterly update 3 (2.6.39-400) onward, support for IB, OFED, and RDS is integrated into the kernel. The OFED userland RPMs continue to be provided, but the `kernel-ib` and `ofa-kernel` RPMs are no longer required.

The kernel ABI remains unchanged in all updates to UEK R2 subsequent to quarterly update 3.

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. However, to allow the introduction of new drivers, there might be instances where changes must be made. Before installing this update, verify the support status of this release with your application vendor.

Chapter 2 Fixed and Known Issues

This chapter describes the fixed and known issues for the Unbreakable Enterprise Kernel Release 2.



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

This update includes minor fixes for various software issues as well as security fixes for several CVEs.

- A kernel panic could occur at boot time if an Emulex LPe16XXX FCoE adapter or a Sun Storage version of such an adapter was installed and the `lpfc` driver was not updated to version 0:8.3.7.10.4p or later. The version of the `lpfc` driver provided in this update corrects the issue.

2.2 Known Issues

This section describes known issues in this update.

ACPI

On some systems you might see ACPI-related error messages in `dmesg` similar to the following:

```
ACPI Error: [CDW1] Namespace lookup failure, AE_NOT_FOUND
ACPI Error: Method parse/execution failed [_SB_.OSC|\|\|\|]
ACPI Error: Field [CDW3] at 96 exceeds Buffer [NULL] size 64 (bits)]>
```

These messages, which are not fatal, are caused by bugs in the BIOS. Contact your system vendor for a BIOS update. (Bug ID 13100702)

ASM

Calling the `oracleasm init` script, `/etc/init.d/oracleasm`, with the parameter `scandisks` can lead to error messages about missing devices similar to the following:

```
oracleasm-read-label: Unable to open device "device": No such file or directory
```

However, the device actually exists. You can ignore this error message, which is triggered by a timing issue. Only use the `init` script to start and stop the `oracleasm` service. All other options, such as `scandisks`, `listdisk`, and `createdisk`, are deprecated. For these and other administrative tasks, use `/usr/sbin/oracleasm` instead. (Bug ID 13639337)

bnx2x driver

When using the `bnx2x` driver in a bridge, disable Transparent Packet Aggregation (TPA) by including the statement `options bnx2x disable_tpa=1` in `/etc/modprobe.conf`. (Bug ID 14626070)

Btrfs

- Together with the UEK R2 kernel in quarterly update 4, a new version of the user-space `btrfs-progs` package (`btrfs-progs-0.20-1.4`) was provided on the `ol6_latest` channel. This package enabled command-line access to new btrfs features. If you install this version or later of the `btrfs-progs` package on a system that does not have an upgraded kernel, most of the new command functionality supported by the package fails with the error `Inappropriate ioctl for device`. The exception

is the `-l` option to `mkfs.btrfs`, which requires version 2.6.39-400.109.1 or later of the Unbreakable Enterprise Kernel. Do not use this option with a non-upgraded kernel as correct functioning of the file system cannot be guaranteed. Although the leaf size appears to be set successfully, there is a risk of data corruption if you subsequently use the resulting file system. (Bug ID 16920640)

- 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. This bug is present in the user-space `btrfs-progs` package (`btrfs-progs-0.20-1.4`). (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. This bug is present in the user-space `btrfs-progs` package (`btrfs-progs-0.20-1.4`). (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)
- A failing RAID1 disk can result in a kernel panic with the error message:

```
BTRFS error (device (null)) in btree_writepage_io_failed_hook:3662: \
IO failure (Error occurred while writing out btree at offset).
```

(Bug ID 16262571)

- The `btrfs filesystem defragment` command exits with an exit code of 20 even if it succeeds. (Bug ID 13714531)
- 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)
- 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. (Bug ID 16278563)
- 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.

- The functionality to limit the space that is available to a quota group before compressing the subvolume is not yet implemented. The `-c` option (limit the space after compression) to the `btrfs qgroup limit` command is implicitly enabled. (Bug ID 16569387)
- 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>.
- The seed-device functionality of btrfs causes a kernel panic when the `btrfs device add` command is run. There is currently no known workaround for this issue. (Bug ID 17334251)

CPU microcode update failures on PVM/PVHVM guests

When running Oracle Linux with UEK R2, 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)

Emulex OneConnect UCNA Firmware Flashing

The `request_firmware` interface to the Emulex Ethernet driver for Emulex OneConnect adapters supports flash updating of the UCNA firmware image.

Oracle Linux 5 systems with the Unbreakable Enterprise Kernel require the `ethtool-6-4.0.1.el5.arch.rpm` version of the `ethtool` package, where `arch` is `i386` or `x86_64` as appropriate. The package is available from the `ol5_i386_UEK_latest` and `ol5_x86_64_UEK_latest` channels on the Unbreakable Linux Network (Advisory [ELBA-2013-2544](#), released Aug 29, 2013) or from Oracle Public Yum at <http://public-yum.oracle.com/repo/OracleLinux/OL5/latest>.

Oracle Linux 6 systems with the Unbreakable Enterprise Kernel do not require a package update.



Note

You can update the firmware image while the UCNA is online and passing network or storage traffic. However, you must reboot the system for the new firmware image to take effect.

To update the UCNA firmware image:

1. Copy the firmware image file (for example, `be3flash.ufi`) to the `/lib/firmware` directory:

```
# cp be3flash.ufi /lib/firmware
```

2. Start the update process:

```
# ethtool -f ethN be3flash.ufi 0
```

where `ethN` is the name of the interface.

3. Reboot the system to enable the new firmware image to take effect.

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)

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.

InfiniBand warning messages when disabling a switch port

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)

libfprint

The following message might appear in `dmesg` or `/var/log/messages`:

```
WARNING! power/level is deprecated; use power/control instead.
```

The USB subsystem in UEK R2 deprecates the `power/level sysfs` attribute in favor of the `power/control` attribute. The `libfprint` fingerprinting library triggers this warning via `udev` rules that try to use the old attribute first. You can safely ignore this warning. The setting of the appropriate power level still succeeds. (Bug ID 13523418)

Multipathed root disk

At boot time, the root file system is not writable when the multipathing service starts. As a result, the system cannot automatically generate the file `/etc/multipath/wwids` and the console displays a message similar to the following:

```
Cannot open file [/etc/multipath/wwids] readonly: No such file or directory
```

You can safely ignore this warning. The subsequent post-boot operation of `multipathd` and `device-mapper-multipath` is not affected.

To prevent this warning message from recurring, enter the following command as `root` to generate `/etc/multipath/wwids` after the system has booted:

```
# multipath -l
```

(Bug ID 17395420, 16076888)

NFSv4 and open(2) flag combinations that can cause a process to hang

If you specify both the `O_RDONLY` and `O_TRUNC` flags when opening a regular file in a mounted NFS version 4 file system, the calling process hangs. The workaround is not to use this combination of flags with regular files. `O_TRUNC` is intended for use with regular files where the open mode allows writing. (Bug ID 17412390)

Nouveau kernel driver is not compatible with NVIDIA graphics driver

After upgrading to UEK R2, the NVIDIA driver upgrade script does not correctly blacklist the Nouveau kernel driver. To blacklist the driver, append `rdblacklist=nouveau nouveau.modeset=0` to the kernel boot parameters in `/boot/grub/grub.conf`.

NUMA warning messages on a non-NUMA system

You can safely ignore the following warning messages in `dmesg` and `/var/log/messages` if you see them on a non-NUMA system:

```
kernel: NUMA: Warning: node ids are out of bound, from=-1 to=-1 distance=10
hcid[4293]: Register path:/org/bluez fallback:1
kernel: No NUMA configuration found
```

(Bug ID 13711370)

pcspkr driver error message

You can safely ignore the following error message:

```
Error: Driver 'pcspkr' is already registered, aborting...
```

The message arises from an alias conflict between `snd-pcsp` and `pcspkr`. To prevent the message from being displayed, add the following line to `/etc/modprobe.d/blacklist.conf`:

```
blacklist snd-pcsp
```

(Bug ID 10355937)

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.

Soft lockup errors when booting

When upgrading or installing the UEK R2 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_pgfld-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)

Support for large memory 32-bit systems

Releases of Oracle Linux prior to Oracle Linux 5 supplied a *hugemem* kernel to allow a system to address up to 64 GB of memory in 32-bit mode. The *hugemem* kernel is no longer available in Oracle Linux 5 and later releases.

The Unbreakable Enterprise Kernel (UEK) supports a maximum of 16 GB of memory for 32-bit kernels on bare metal and hardware virtualized machine (HVM) systems, and 8 GB for fully paravirtualized machine (PVM) systems. 32-bit PVM guest operating systems must be located in the first 128 GB of physical memory on the host.

The Red Hat Compatible Kernel (RHCK) has the same limitations, except that PVM systems can have up to 16 GB of memory. The limitation of 8 GB for PVM on UEK was chosen for reasons of reliability.

A 32-bit system uses the PAE (physical address extension) memory feature to map physical memory beyond 4 GB into the 32-bit address space that is available to each process. A 64-bit system can address memory beyond 4 GB without requiring an extra layer of memory abstraction.

Oracle Linux on x86_64 includes 32-bit libraries, which allow applications built for both 64-bit and 32-bit Linux to run on the same system. This capability provides scalability to virtually unlimited memory sizes, while retaining the ability to run 32-bit applications. Oracle recommends this configuration for any system with more than 4 GB of memory. (Bug ID 16974301)

Transparent Huge Pages

This update removes the Transparent Huge Pages (THP) feature. Following extensive benchmarking and testing, Oracle found that THP caused a performance degradation for some workloads of between 5 and 10%. 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 update, 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 17279055)

udev

A message similar to the following might be recorded in `dmesg` or `/var/log/messages` at boot time:

```
udevd (pid): /proc/pid/oom_adj is deprecated, please use /proc/pid/oom_score_adj instead.
```

The `udev` process uses the deprecated `oom_adj` kernel interface to prevent it from being killed if the system runs short of memory. You can safely ignore the message as the action still succeeds. To prevent the message from occurring, install the package `udev-147-2.42.el6.arch.rpm` or higher for Oracle

Linux 6, or [udev-095-14.29.0.1.el5.arch.rpm](#) or higher for Oracle Linux 5. (Bug ID 13655071, 13712009)

Virtualization

- When booting UEK R2 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 R2 Quarterly Update 5 kernel causes a disparity between the date and time as displayed by `date` and `hwclock`. 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 R2 quarterly update 5. 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)

- Under Oracle VM Server 2.2.2, increasing the memory that is assigned to an Oracle Linux 6 Update 4 x86_64 HVM guest running UEK R2 Quarterly Update 5 causes the guest to reboot. There is no known workaround for this issue. (Bug ID 17440635)

XFS

- Reading from or writing to a `zero` or `null` device on a read-only XFS file system fails with the following error:

```
XFS_IOC_FSGETGEOMETRY: Inappropriate ioctl for device
```

(Bug ID 16970090)

- Newly created directories in an XFS file system do not inherit their group setting from the parent directory on which the `setgid` bit is set. (Bug ID 17423815)

Chapter 3 Installation and Availability

The Unbreakable Enterprise Kernel Release 2 Quarterly Update 5 can be installed on Oracle Linux 5 Update 8 or newer, as well as Oracle Linux 6 Update 2 or newer, both 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 kernel images are available as binary RPM packages from dedicated channels on Oracle's Unbreakable Linux Network (ULN) as well as the Oracle Public Yum repository. Four channels are available:

- Oracle Linux 5 (x86): `o15_i386_UEK_latest`
- Oracle Linux 5 (x86_64): `o15_x86_64_UEK_latest`
- Oracle Linux 6 (x86): `o16_i386_UEK_latest`
- Oracle Linux 6 (x86_64): `o16_x86_64_UEK_latest`

If your system is registered with ULN, make sure you subscribe it to the appropriate `UEK_latest` channel.

For Oracle Public Yum, the appropriate `UEK_latest` channel is automatically enabled in the yum respiratory file under `/etc/yum.repos.d` when you install Oracle Linux 5 update 9 or later and Oracle Linux 6 update 3 or later.

To upgrade an existing Oracle Linux 5 or Oracle Linux 6 installation to the latest UEK R2, enable the appropriate `UEK_latest` channel and run `yum update`.

For detailed instructions on how to download and install the Unbreakable Enterprise Kernel on Oracle Linux see the *Getting Started with the Unbreakable Enterprise Kernel for Oracle Linux* document on the Oracle Technology Network (OTN) at <http://www.oracle.com/technetwork/articles/servers-storage-admin/uek-rel2-getting-started-1555632.html>.

If you have questions regarding configuring or using `yum` to install updates, refer to the *Oracle Linux Administrator's Solutions Guide* at http://docs.oracle.com/cd/E37670_01/index.html.

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.

The kernel's source code is available via a public git source code repository at <http://oss.oracle.com/git/?p=linux-uek-2.6.39.git>.

