

Oracle® Linux

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

ORACLE®

E41136-05
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 4 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	2
1.3 Driver Updates	2
1.3.1 Storage Adapter Drivers	2
1.3.2 Network Adapter Drivers	3
1.3.3 Miscellaneous Drivers	4
1.4 Technology Preview	4
1.5 Compatibility	5
2 Fixed and Known Issues	7
2.1 Fixed Issues	7
2.2 Known Issues	7
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. It is based on the mainline Linux 3.0 version 3.0.36. 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, which was based on Linux 2.6.32.

The 2.6.39-400.109.1 release is the fourth 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

- The integrated OpenFabrics Alliance (OFED) 1.5.5 stack, which supports the following InfiniBand hardware on systems with an x86_64 architecture:
 - Mellanox ConnectX-2 InfiniBand Host Channel Adapters
 - Sun InfiniBand QDR Host Channel Adapter PCIe #375-3696
- The following kernel-level features are implemented for the btrfs file system:
 - Allow metadata blocks to be larger than the page size (4 KB) and up to 64 KB in size.
 - A new `ioctl` call, `BTRFS_IOC_DEVICES_READY`, to read device-readiness status.
 - Defragmentation does not undo sharing of data blocks between snapshots. Previously, defragmentation could result in more disk space usage because shared data blocks were duplicated.
 - New `ioctl` calls, `BTRFS_IOC_GET_LABEL` and `BTRFS_IOC_SET_LABEL`, to get or set a file system label.
 - Speed improvements for `fsync` direct I/O, and concurrent, multithreaded reads.
 - Support for file-hole punching by calling `fallocate()` with the `FALLOC_FL_PUNCH_HOLE` flag. Use the `FIEMAP` `ioctl` to detect holes. The `SEEK_HOLE` and `SEEK_DATA` options to the `lseek()` system call are not supported.
- The user-space `btrfs-progs` package (`btrfs-progs-0.20-1.4`) will be provided on the `ol6_latest` channel for installation by using `yum update`. Command-line access to the new features in this package is not supported on kernel versions prior to this update. If you install this update but not the updated user-space `btrfs-progs` package, programmatic access is possible for those features that present an `ioctl` or other programming interface.

The updated `btrfs-progs` package supports the following features in addition to the kernel-level features:

- The `-l` option to `mkfs.btrfs` specifies the leaf size of a file system. The default value of `4k` (4 KB) is recommended for most uses. A value of `16k` or `32k` (16 or 32 KB) can improve performance for some workloads by reducing metadata fragmentation. The maximum supported size is `64k` (64 KB).
- The `btrfs filesystem label device newlabel` command supports labeling or relabeling of an existing file system. The limitations are that the file system can only have one device and you must first unmount the file system. The `btrfs filesystem show` command displays a file system's label.
- Devices can be replaced at run time by using the `btrfs replace start` command.
- The `btrfs filesystem balance` command is able change btrfs RAID profiles dynamically.
- Collection of device statistics (numbers of read and write failures, checksum errors, and corrupted blocks) by using the `btrfs device stats` command.
- Defragmentation can be cancelled by pressing `Ctrl-C` or by killing the defragmentation process.

1.2 Xen Improvements

Several improvements have been incorporated into the Unbreakable Enterprise Kernel to support Xen usage:

- Numerous bug fixes and performance improvements.
- Optimizations for `xen/privcmd` for ARM and PVH by a new hypercall (`add_to_physmap_range`).
- Support for the balloon driver in Xen ARM.
- Enhancements to permit PVHVM backend drivers and allow dom0 functionality to be moved to guests.
- Implementation of the persistent grant extension to the block ring protocol to improve block protocol scalability for large numbers of physical cores and guests.
- Xen Processor Aggregator Device (PAD) added.

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

ATA over Ethernet

- ATA over Ethernet (AoE) driver (`aoe`) updated to 47q.

Broadcom

- NetXtreme II iSCSI driver (`bnx2i`) updated to 2.7.6.1d.
- NetXtreme II Fibre Channel over Ethernet driver (`bnx2fc`) updated to 2.3.4.

Cisco

- Cisco FCoE HBA Driver (`fnic`) updated to 1.5.0.41.

Emulex

- Blade Engine 2 Open-iSCSI driver ([be2iscsi](#)) updated to 10.0.467.0o.
- Fibre Channel HBA driver ([lpfc](#)) updated to 0:8.3.7.10.4p.

LSI

- LSI Fusion-MPT base driver ([mptbase](#)) updated to 4.28.20.01.
- LSI Fusion-MPT `ioctl` driver ([mptctl](#)) updated to 4.28.20.01.
- LSI Fusion-MPT Fibre Channel host driver ([mptfc](#)) updated to 4.28.20.01.
- LSI Fusion-MPT IP Over Fibre Channel driver ([mptlan](#)) updated to 4.28.20.01.
- LSI Fusion-MPT SAS driver ([mptsas](#)) updated to 4.28.20.01.
- LSI Fusion-MPT SCSI host driver ([mptscsih](#)) updated to 4.28.20.01.
- LSI Fusion-MPT SPI host driver ([mptspi](#)) updated to 4.28.20.01.
- LSI Fusion-MPT SAS 2.0 driver ([mpt2sas](#)) updated to 16.05.01.00.

MegaRAID

- MegaRAID SAS driver ([megaraid_sas](#)) updated to 06.505.02.00.

Mellanox



Note

These drivers were first released in UEK R2 Quarterly Update 3.

- ConnectX core driver ([mlx4_core](#)) released at 1.0-ofed1.5.5 (x86_64 only). Handles low-level functions such as device initialization and firmware commands processing, and controls resource allocation so that the InfiniBand and Ethernet functions can share a device without interfering with each other.
- ConnectX Ethernet driver ([mlx4_en](#)) updated to 1.5.10 (Jan 2013) (x86_64 only). Handles Ethernet-specific functions and plugs into the netdev mid-layer.
- ConnectX InfiniBand driver ([mlx4_ib](#)) released at 1.0-ofed1.5.5 (x86_64 only). Handles InfiniBand-specific functions.

QLogic

- Fibre Channel HBA driver ([qla2xxx](#)) updated to 8.05.00.03.39.0-k.
- iSCSI driver ([qla4xxx](#)) updated to 5.03.00.02.06.02-uek2. Supports Open-iSCSI.

1.3.2 Network Adapter Drivers

Broadcom

- NetXtreme II network adapter driver ([bnx2](#)) updated to 2.2.3n.
- NetXtreme II 10Gbps network adapter driver ([bnx2x](#)) updated to 1.76.54.

- Converged Network Interface Card core driver (`cnic`) updated to 2.5.16g.
- Tigon3 Ethernet adapter driver (`tg3`) updated to 3.129d.

Emulex

- Blade Engine 2 10Gbps adapter driver (`be2net`) updated to 4.6.62.0u.

Intel

- PRO/1000 PCI-Express Gigabit network adapter driver (`e1000e`) updated to 2.3.2-NAPI.
- Gigabit Ethernet network adapter driver (`igb`) updated to 4.1.2.
- 10 Gigabit PCI-Express network adapter driver (`ixgbe`) updated to 3.14.5.
- 10 Gigabit Server Adapter virtual function driver (`ixgbevf`) updated to 2.8.7. The kernel must support Single Root I/O Virtualization (SR-IOV).

QLogic

- 1/10 GbE Converged/Intelligent Ethernet Adapter driver (`qlcnict`) updated to 5.2.29.45.
- QLE81xx network adapter driver (`qlge`) updated to v1.00.00.32.

Realtek PCI Express Gigabit Ethernet controller

- Realtek PCI Express Gigabit Ethernet controller (`r8169`) updated to 2.3LK-NAPI.

VMware

- VMware VMXNET3 virtual ethernet driver (`vmxnet3`) updated to 1.1.30.0-k.

1.3.3 Miscellaneous Drivers

HP

- HP ProLiant Channel Interface Device Driver for iLO (`hpilo`) updated to 1.4.

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.

- **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.

- **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.

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

The following issues have been fixed in this update.

- A kernel panic could occur during path failover if one of the paths to a multipathed SCSI device was disabled. (Bug ID 16684527)
- Several fixes have been applied to correct race conditions within cgroup movement of newly created, forking, and waking processes. (Bug ID 13740515)
- Several fixes have been applied to the OVM API.
- A fix has been applied to prevent a kernel panic that could occur when moving OCFS2 extents during defragmentation. (Bug ID 16631951)
- The updated `btrfs-progs` package (`btrfs-progs-0.20-1.4`), which will be provided on the `ol6_latest` channel, fixes a bug where the `btrfs subvolume get-default` command listed all existing subvolumes instead of only the default subvolume. (Bug ID 13815433)
- Security fixes for several CVEs, including CVE-2013-2094.

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 this quarterly update, a new version of the user-space `btrfs-progs` package (`btrfs-progs-0.20-1.4`) will be provided on the `ol6_latest` channel. This package enables command-line access to new btrfs features. If you install this latest version 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 `fsync` 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)

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)

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 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)

udev

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

```
udev (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 and 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 4 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)

Chapter 3 Installation and Availability

The Unbreakable Enterprise Kernel Release 2 Quarterly Update 4 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>.

