

Oracle® Linux

**Release Notes for Unbreakable Enterprise Kernel Release 4
Update 1**

ORACLE®

E73420-06
May 2018

Oracle Legal Notices

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

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

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

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

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

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

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

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

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

Abstract

This document contains information on Update 1 to the Unbreakable Enterprise Kernel Release 4. This document may be updated after it is released. To check for updates to this document, and to view other Oracle documentation, refer to the Documentation section on the Oracle Technology Network (OTN) Web site:

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

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

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

Table of Contents

Preface	v
1 New Features and Changes	1
1.1 Notable Changes	1
1.2 LXC Improvements	1
1.3 Xen Improvements	1
1.4 btrfs Improvements	2
1.5 DTrace Improvements	2
1.6 Driver Updates	3
1.7 Technology Preview	3
1.8 Compatibility	4
1.9 Header Packages for Development	4
2 Fixed and Known Issues	7
2.1 Fixed Issues	7
2.2 Known Issues	8
2.2.1 btrfs	8
2.2.2 DIF/DIX is not supported for ext file systems	9
2.2.3 Docker	9
2.2.4 DTrace	9
2.2.5 LXC	9
2.2.6 Console Appears to Hang when Booting	10
2.2.7 OFED iSER feature causes Out-Of-Memory error on Oracle Linux 6	10
2.2.8 XFS quotas are disabled after a read-only remount on Oracle Linux 6	10
2.2.9 Open File Description (OFD) locks are not supported on NFSv4 mounts	10
3 Installation and Availability	11
3.1 Installation Overview	11
3.2 Subscribing to ULN Channels	11
3.3 Enabling Access to Oracle Linux Yum Server Channels	12
3.4 Upgrading Your System	14
3.5 Installing the Oracle-Supported OFED Packages	14

Preface

The *Oracle Linux Unbreakable Enterprise Kernel Release Notes* provides a summary of the new features, changes, and fixed and known issues in the Unbreakable Enterprise Kernel Release 4 Update 1.

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 4 (UEK R4) is Oracle's fourth major release of its heavily tested and optimized operating system kernel for Oracle Linux 6 Update 7 or later, and Oracle Linux 7 Update 2 or later, on the x86-64 architecture. It is based on the mainline Linux kernel version 4.1.12. This release also updates drivers and includes bug and security fixes.

The 4.1.12-37.2.1 release is the first update release for UEK R4. 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 R4.

UEK R4 uses the same versioning model as the mainline Linux kernel version. It is possible that some applications might not understand the 4.1 versioning scheme. However, regular Linux applications are usually neither aware of nor affected by Linux kernel version numbers.

1.1 Notable Changes

- **bcache-tools package added.** The block layer cache feature, `bcache`, which was introduced with the release of UEK R4, requires the `bcache-tools` package to configure this facility. The package is built and included for both Oracle Linux 6 and Oracle Linux 7.
- **crash updated to 7.1.4.** The `crash` utility that is used to analyze crash dump data or live systems has been updated to version 7.1.4. The package is built and included for both Oracle Linux 6 and Oracle Linux 7.
- **libfuse updated to 2.9.4.** New features in `libfuse` have been backported from the mainline `fuselib` (3.0). UEK4 does not use fuse 3.0 because it breaks backward compatibility on the binary level as well as the source level. New features include security context options, new caching logic for network-based fuse file systems, support for asynchronous direct I/O and writeback cache, `readdirplus`, and more. The update applies to both Oracle Linux 6 and Oracle Linux 7.
- **xfsplogs updated to 3.2.3.** New features in XFS for CRC require version 3.2.3 of the `xfsplogs` package. The package is updated for both Oracle Linux 6 and Oracle Linux 7.

1.2 LXC Improvements

The LXC userspace package has been updated from 1.0.7 to 1.1.5 to support all of the UEK R4 features and ongoing container initiatives. The package is updated for both Oracle Linux 6 and Oracle Linux 7.

A fix is included to resolve an issue where TCP packet corruption could occur when running a veth interface inside a container.

This release also includes security and bug fixes.

1.3 Xen Improvements

A large number of bug fixes and security fixes have been backported for this release to include the features available in Xen 4.3.

A fix has been applied to the Xen `pciback` driver to allow it to take control of devices and to prevent native drivers from loading. This change makes it easier for developers to do PCI passthrough. The driver has also been compiled directly into the kernel and does not require a separate module.

1.4 btrfs Improvements

- **btrfs-progs updated to 4.2.2-3.** The new version of `btrfs-progs` checks the compatibility of features that are supported by the kernel when using the `mkfs` and `btrfs-convert` commands. These commands are able to warn users about unsupported features. This update also fixes several bugs that were present in previous releases.
- **Handling of inline extents improved.** Methods handling the truncation of files that have inline and compressed extents have been improved to reduce the amount of metadata space required and to provide a security fix. Also, methods handling the cloning of inline extents have been fixed to prevent file corruption and potential data loss. A patch was also applied to fix cloning operations for compressed extents by checking the source clone offset and length to ensure that it falls within the source file's boundaries. (Bug IDs 22307083, 22598448, 22466327, 22589051)
- **Send operation in cloning code fixed.** A fix for a bug where snapshots would fail to send or receive across different file systems was implemented. The fix was applied to the send operation in the cloning code to handle the case where a snapshot is taken of a received subvolume and that is then used as the parent to try to receive on a different file system, due to differing UUIDs. (Bug ID 22580612)
- **COW Page Size Extents Improved.** A fix was applied to help to reduce unnecessary COW page size extent by only zeroing out the end of the page if the range requested for file allocation exceeds the size of the inode. (Bug ID 22573877)
- **File corruption triggered by faulty cloning operations fixed.** A fix was applied to ensure that clone operations on files that share extents with other files do not cover regions in the source file that may point to different extents. (Bug ID 22579887)

1.5 DTrace Improvements

- **Accurately pair signal probes for non-ignored unhandled signals.** DTrace was improved to more accurately pair signal probes for non-ignored unhandled signals, by using the signal number, obtained from the original signal in the target task, in the signal-handled probe. Previously, the signal was translated into a `SIGKILL` because the signal effectively causes the process to be terminated. This change improves the accuracy of exit reporting.
- **Sleep code performance improvements and fixes.** Code that DTrace used to sleep between operations has been updated to improve performance during these loops and to avoid endless looping.
- **Dynamically allocated text in Oracle VM VirtualBox drivers workaround.** VirtualBox dynamically allocates space for executable text within the kernel. When these text addresses were encountered by stack trace back code, they were dropped or the stack trace was terminated. A patch was applied to DTrace to apply a workaround which allows DTrace to assume that executable pages in `vmalloc` or module areas as valid text addresses.
- **Missing psinfo for tasks run with PID 0 fixed.** A bug which caused some memory access errors when attempting to access the `psinfo` for a task running with PID 0 is resolved. (Bug ID 22561297)
- **Multiple probe firings of a Statically Defined Tracing (SDT) probe in a function now supported.** A patch was applied to allow DTrace to handle situations where a single SDT probe occurs more than once in a function, usually as a result of the compiler duplicating code fragments that contain an SDT probe. In these situations multiple probe firings were not previously reported, however the patch resolves this issue.

1.6 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.

Table 1.1 Updated Drivers in UEK R4 U1

Manufacturer	Driver	Version	Description
Broadcom	be2iscsi	10.6.0.2	Broadcom/Avago OneConnect Open-iSCSI Driver
Broadcom	be2net	11.0.0.0	Broadcom/Emulex OneConnect NIC Driver
Broadcom	lpfc	11.0.0.13	Broadcom/Emulex LightPulse Fibre Channel SCSI driver
Broadcom	megaraid_sas	Upstream patches	Broadcom/Avago SAS based MegaRAID driver
PMC-Sierra	hpsa	3.4.10-0	PMC-Sierra Driver for HP Smart Array Controller
Intel Corporation	i40e	1.4.25-k	Intel(R) Ethernet Connection XL710 Network Driver
Intel Corporation	i40evf	1.4.15-k	Intel(R) XL710 X710 Virtual Function Network Driver
Microsoft Corporation	hv_vmbus hv_netvsc hv_storvsc hv_utils hv_balloon hyperv_fb	Upstream patches	Microsoft Hyper-V Drivers
QLogic	qla2xxx	8.07.00.33.40.0-k	QLogic Fibre Channel HBA Driver
VMware, Inc.	vmw_vsock_vmci_transport	1.0.2.0-k	VMCI transport for Virtual Sockets
VMware, Inc.	vmxnet3	1.4.2.0-k	VMware vmxnet3 virtual NIC driver

Additional Notes For Driver Updates

The following notes are included at the request of a vendor for the listed driver:

- [lpfc](#): Locked optics support is enabled for LPE32000 HBAs and all variant HBAs of this architecture. With this capability, these HBAs will detect and enable both Avago or Emulex certified SFP and QSFP optics.

For driver rev 11.0.0.13 and higher unqualified optics will be disabled, the link will not come up, an error message is written to the log file and the lpfc driver will display this message:

```
3176 Port Name [wwpn] Unqualified optics - Replace with Avago optics for Warranty and Technical support
```

1.7 Technology Preview

The following features included in the Unbreakable Enterprise Kernel Release 4 are still under development, but are made available for testing and evaluation purposes.

- **Ceph File System and Object Gateway Federation**

Ceph presents a uniform view of object and block storage from a cluster of multiple physical and logical commodity-hardware storage devices. Ceph can provide fault tolerance and enhance I/O performance by replicating and striping data across the storage devices in a Storage Cluster. Ceph's monitoring and self-repair features minimize administration overhead. You can configure a Storage Cluster on non-identical hardware from different manufacturers.

The Ceph File System (CephFS) and Object Gateway Federation features of Ceph are in technology preview.

- **DCTCP (Data Center TCP)**

DCTCP enhances congestion control by making use of the Explicit Congestion Notification (ECN) feature of state-of-the-art network switches. DCTCP reduces buffer occupancy and improves throughput by allowing a system to react more intelligently to congestion than is possible using TCP.

- **DRBD (Distributed Replicated Block Device)**

A shared-nothing, synchronously replicated block device (*RAID1 over network*), designed to serve as a building block for high availability (HA) clusters. It requires a cluster manager (for example, pacemaker) for automatic failover.

- **Kernel module signing facility**

Applies cryptographic signature checking to modules on module load, checking the signature against a ring of public keys compiled into the kernel. GPG is used to do the cryptographic work and determines the format of the signature and key data.

- **NFS over RDMA interoperation with ZFS and Oracle Solaris**

NFS over RDMA does not yet fully interoperate with ZFS and Oracle Solaris. NFS over RDMA for NFS versions 3 and 4 is supported for Oracle Linux systems using the Oracle InfiniBand stack and is more efficient than using NFS with TCP over IPoIB. Currently, only the Mellanox ConnectX-2 and ConnectX-3 Host Channel Adapters (HCAs) pass the full Connectathon NFS test suite and are supported.

- **NFS server-side copy offload**

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

- **Server-side parallel NFS**

Server-side parallel NFS (pNFS) improves the scalability and performance of an NFS server by making file metadata and data available on separate paths.

1.8 Compatibility

Oracle Linux maintains user-space compatibility with Red Hat Enterprise Linux, which is independent of the kernel version running underneath the operating system. Existing applications in user space will continue to run unmodified on the Unbreakable Enterprise Kernel Release 4 and no re-certifications are needed for RHEL certified applications.

To minimize impact on interoperability during releases, the Oracle Linux team works closely with third-party vendors whose hardware and software have dependencies on kernel modules. The kernel ABI for UEK R4 will remain unchanged in all subsequent updates to the initial release. In this release, there are changes to the kernel ABI relative to UEK R3 that require recompilation of third-party kernel modules on the system. Before installing UEK R4, verify its support status with your application vendor.

1.9 Header Packages for Development

As of UEK-3.8-QU2, the `kernel-uek-headers` package is no longer built and distributed. There are three kernel packages that may be useful for development purposes. The `kernel-headers` package

forms part of the API for user space programs. The `kernel-devel` package is used for standard RHCK development and module compilation. The `kernel-uek-devel` package is used for UEK development and module compilation. Neither the `kernel-uek-headers`, nor the `kernel-headers` packages, are needed for kernel development.

The `kernel-headers` 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.

To install the packages required to build modules against UEK and the C header files for both UEK and RHCK:

```
# yum install kernel-uek-devel-`uname -r` kernel-headers
```

Chapter 2 Fixed and Known Issues

This chapter describes the fixed and known issues for the Unbreakable Enterprise Kernel Release 4.



Important

Run the `yum update` command regularly to ensure that the latest bug fixes and security errata are installed on your system.

2.1 Fixed Issues

The following issues have been fixed in this update.

- **OFED packages updated to release 1.0.0-8.** The packages required for Oracle OFED release 1.0.0-8 have been updated for both Oracle Linux 6 and Oracle Linux 7.
- **kexec-tools package update to version 2.0.7-38.0.2.** The `kexec-tools` package has been updated to version 2.0.7-38.0.2 to resolve an issue that caused iSCSI boot to fail on some systems.
- **Intel xHCI controller reset code (`xhci`) fixed.** A workaround fix, provided by Intel, has been applied to the xHCI controller reset code to prevent an intermittent boot hang. (Bug ID 22805913)
- **Kernel panic when `bio_endio()` is called for devices with write cache enabled (WCE) fixed.** A fix was applied to code in the kernel block IO controller to prevent kernel panic when a write or sync takes place on a device with WCE (Write Cache Enabled). (Bug ID 22820562)
- **ALSA usb-audio fixed to avoid freeing `umidi` object twice.** A fix was applied to the ALSA `usb-audio` code which resulted in a double-free flaw in the USB MIDI driver. (Bug ID 22740866)
- **RDMA operations report success when the `bcopy` message fails fixed.** A fix was applied to delay the signaling of a completion event for the RDMA (Remote Direct Memory Access) operation until the `bcopy` send event completes. (Bug ID 22847528)
- **Branch offset adjustment on backjumps after patching `ctx` expansion fixed.** When `ctx` access is used, the kernel needs to expand instructions so that branch offsets are adjusted for forward and backward jumps. A fix was applied to correctly calculate the delta required for the backward jumps. (Bug ID 22740787)
- **i40e driver fixed to revert a commit that caused the Ethernet interface to flap when bonded.** A fix was applied to the `i40e` driver to revert an earlier code commit that caused the Ethernet interface to flap when in a bond configuration. (Bug ID 22519254)
- **XFS data corruption while performing EOF zeroing fixed.** A fix was applied to XFS for an issue that resulted in data corruption due to I/O lock management and pending Direct I/O not being drained before EOF zeroing. (Bug ID 22682207)
- **SCSI/bdi code reverted to fix multiple issues with backing devices.** A previous change to the SCSI/bdi code was reverted to resolve multiple issues with backing devices, including a race condition that occurred during bdi reregistration. This fix is in line with changes to the mainline kernel. (Bug ID 22679087)
- **libfuse I/O error for a `chmod` on DBFS fixed.** A fix was applied to `libfuse` to resolve an issue that resulted in an I/O error when running a `chmod` operation on files or directories on DBFS. (Bug ID 22637182)
- **libfuse fixed for a bug where the `numa` mount option for DBFS was not shown in `/proc/mounts`.** A fix was applied to `libfuse` to resolve an issue where the `numa` mount option was not visible in `/proc/mounts` if it was used to perform a mount of a DBFS. (Bug ID 22637071)

- **Block driver I/O errors and performance degradation in some environments fixed.** A fix that was applied to resolve a regression in the kernel block driver code that caused performance degradation and I/O errors due to the size of an I/O request at the block driver level. (Bug ID 22611290)
- **OCFS2 system reboot issue due to a driver error fixed.** A fix that was applied to the ocfs2 driver so that when the journal commit receives an error from the device driver, the journal changes to an abort state that stops the commit thread. (Bug ID 22611290)
- **OCFS2 heartbeat stability and mount failure issue fixed.** A fix that was applied to the OCFS2 cluster heartbeat code to increase unsteady iterations and improve heartbeat stability. (Bug ID 22611290)
- **NFS server crash triggered by 1MB NFS WRITE fixed.** A fix has been applied for a bug in NFS server code that caused a general protection fault if an NFS client triggered a 1MB NFS WRITE operation on an NFS/RDMA mount.
- **Mellanox InfiniBand driver updated to use kvfree() instead of kfree() to free buffers.** A fix has been applied to the mlx4 driver to use kvfree() rather than kfree() to free buffers, since kvfree() can handle memory allocated by both kcalloc() and vmalloc().
- **RDS fencing operation issue fixed for parallel write operations that caused connections to drop.** A fix has been applied to the mlx4 driver code that was causing the fencing operation for RDS (Reliable Datagram Sockets) to drop connections as a result of a process that recycled keys too quickly. This problem affected Oracle Database processes using shared memory buffers that had pending RDMA operations.
- **RDS enhancements for connection tracking and latency tracing.** Enhancements to code for RDS connection tracking have been applied to make it easier to debug connection shutdown and reconnect behavior. A tracker variable keeps the latest value of the connection drop source making it easier to identify the reason for a connection drop. Tracking is also applied to determine the length of time that a reconnection attempt takes to fail. An interface has been applied to trace receive message latency.
- **Block Loop Driver Updated To Log ENOSPC And EDQUOT Errors.** The block loop driver has been modified to write masking ENOSPC and EDQUOT errors before they get masked to EIO, to assist with debugging on loop devices.
- **btrfs quota enable command allows correct file space usage calculation.** A bug in the btrfs quota enable command which caused existing files not to count toward space usage on a non-empty file system has been fixed. (Bug ID 16569350)
- **The btrfs quota rescan command has been implemented and is working.** The btrfs quota rescan command has been implemented in this release. The command correctly performs a rescan and returns a message to indicate whether the scan was successful. (Bug ID 16569350)

2.2 Known Issues

This section describes the known issues in this update.

2.2.1 btrfs

- The `btrfs filesystem balance` command does not warn that the RAID level can be changed under certain circumstances, and does not provide the choice of cancelling the operation. (Bug ID 16472824)
- The copy-on-write nature of btrfs means that every operation on the file system initially requires disk space. It is possible that you cannot execute any operation on a disk that has no space left; even removing a file might not be possible. The workaround is to run `sync` before retrying the operation. If this

does not help, remount the file system with the `-o noatacow` option and delete some files to free up space. See <https://btrfs.wiki.kernel.org/index.php/ENOSPC>.

- When you overwrite data in a file, starting somewhere in the middle of the file, the overwritten space is counted twice in the space usage numbers that `btrfs qgroup show` displays. Using the `btrfs quota rescan` does not help fix this issue either. (Bug ID 16609467)
- If you use the `-s` option to specify a sector size to `mkfs.btrfs` that is different from the page size, the created file system cannot be mounted. By default, the sector size is set to be the same as the page size. (Bug ID 17087232)
- The `btrfs-progs` and `btrfs-progs-devel` packages for use with UEK R4 are made available in the `ol6_x86_64_UEKR4` and `ol7_x86_64_UEKR4` ULN channels and the `ol6_UEKR4` and `ol7_UEKR4` Oracle Linux yum server repositories. In UEK R3, these packages were made available in the `ol6_x86_64_latest` and `ol7_x86_64_latest` ULN channels and the `ol6_latest` and `ol7_latest` Oracle Linux yum server repositories.

2.2.2 DIF/DIX is not supported for ext file systems

The Data Integrity Field (DIF) and Data Integrity Extension (DIX) features that have been added to the SCSI standard are dependent on a file system that is capable of correctly handling attempts by the memory management system to change data in the buffer while it is queued for a write.

The ext2, ext3 and ext4 file system drivers do not prevent pages from being modified during I/O which can cause checksum failures and a "Logical block guard check failed" error. Other file systems such as XFS are supported. (Bug ID 24361968)

2.2.3 Docker

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

```
RpmdB checksum is invalid: dCDPT(pkg checksums): package_name
```

This error breaks Dockerfile builds and is a known issue upstream (see <https://github.com/docker/docker/issues/10180>.)

The workaround is to run `touch /var/lib/rpm/*` before installing the package. (Bug ID 21804564)

2.2.4 DTrace

- Argument declarations for probe definitions cannot be declared with derived types such as `enum`, `struct`, or `union`.
- The following compiler warning can be ignored for probe definition arguments of type `string` (which is a D type but not a C type):

```
provider_def.h:line#: warning: parameter names (without types) in function declaration
```

- Multi-threaded processes under `ustack()`, `usym()`, `uaddr()` and `umod()` which perform `dlopen()` in threads other than the first thread may not have accurate symbol resolution for symbols introduced by the `dlopen()`. (Bug ID 20045149)

2.2.5 LXC

- **The `lxc-net` service does not always start immediately after installation on Oracle Linux 6.** The `lxc-net` service does not always start immediately after installation on Oracle Linux 6, even though

this action is specified as part of the RPM post-installation script. This can prevent the `lxcbr0` interface from coming up. If this interface is not up after installation, you can manually start it by running `service lxc-net start`. (Bug ID 23177405)

- **LXC read-only `ip_local_port_range` parameter.** With `lxc-1.1` or later and UEK R4, `ip_local_port_range` is a read-writable parameter under `/proc/sys/net/ipv4` in an Oracle Linux container rather than being read-only. (Bug ID 21880467)

2.2.6 Console Appears to Hang when Booting

When booting Oracle Linux 6 on hardware with an ASPEED graphics controller, the console may appear to hang during the boot process after starting `udev`. However, the system does boot properly and is accessible. The workaround is to add `nomodeset` as a kernel boot parameter in `/etc/grub.conf`. (Bug ID 22389972)

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

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

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

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

2.2.8 XFS quotas are disabled after a read-only remount on Oracle Linux 6

Quotas are disabled on XFS if the file system is remounted with read-only permissions on Oracle Linux 6. (Bug ID 22908906)

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

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

Chapter 3 Installation and Availability

You can install Unbreakable Enterprise Kernel Release 4 on Oracle Linux 6 Update 7 or later, or Oracle Linux 7 Update 1 or later, running either the Red Hat compatible kernel or a previous version of the Unbreakable Enterprise Kernel. If you are still running an older version of Oracle Linux, first update your system to the latest available update release.

The Unbreakable Enterprise Kernel Release 4 is supported on the x86-64 architecture but not on x86.

3.1 Installation Overview

If you have a subscription to Oracle Unbreakable Linux support, you can obtain the packages for Unbreakable Enterprise Kernel Release 4 by registering your system with the Unbreakable Linux Network (ULN) and subscribing it to additional channels. See [Section 3.2, “Subscribing to ULN Channels”](#).

If your system is not registered with ULN, you can obtain most of the packages from Oracle Linux yum server. See [Section 3.3, “Enabling Access to Oracle Linux Yum Server Channels”](#).

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

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

3.2 Subscribing to ULN Channels

The kernel image and user-space packages are available on the following ULN channels for Oracle Linux 6:

- `ol6_x86_64_latest` (latest user-space packages for Oracle Linux 6 other than DTrace, OFED, and DRBD packages)
- `ol6_x86_64_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `ol6_x86_64_UEKR4_DTrace_userspace` (`dtrace-utils*`)
- `ol6_x86_64_UEKR4_OFED` (latest OFED tools packages)
- `ol6_x86_64_mysql-ha-utils` (`drbd84-utils`)

The kernel image and user-space packages are available on the following ULN channels for Oracle Linux 7:

- `ol7_x86_64_latest` (latest user-space packages for Oracle Linux 7 other than DTrace, OFED, and DRBD packages)
- `ol7_x86_64_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `ol7_x86_64_UEKR4_DTrace_userspace` (`dtrace-utils*`)
- `ol7_x86_64_UEKR4_OFED` (latest OFED tools packages)
- `ol7_x86_64_mysql-ha-utils` (`drbd84-utils`)

The following procedure assumes that you have already registered your system with ULN.

To subscribe your system to a channel on ULN:

1. Log in to <http://linux.oracle.com> with your ULN user name and password.
2. On the Systems tab, click the link named for the system in the list of registered machines.
3. On the System Details page, click **Manage Subscriptions**.
4. On the System Summary page, select each required channel from the list of available channels and click the right arrow to move the channel to the list of subscribed channels.

For Oracle Linux 6, subscribe the system to the `ol6_x86_64_latest` and `ol6_x86_64_UEKR4` channels. If required, you can also add the channels for the DTrace, OFED, and DRBD packages. You do not need to subscribe the system to the `ol6_x86_64_UEK_latest` or `ol6_x86_64_UEKR3_latest` channels.

For Oracle Linux 7, subscribe the system to the `ol7_x86_64_latest` and `ol7_x86_64_UEKR4` channels. If required, you can also add the channels for the DTrace, OFED, and DRBD packages. You do not need to subscribe the system to the `ol7_x86_64_UEKR3` channel.

5. Click **Save Subscriptions**.

For information about using ULN, see the *Oracle Linux Unbreakable Linux Network User's Guide* at http://docs.oracle.com/cd/E37670_01/index.html.

3.3 Enabling Access to Oracle Linux Yum Server Channels

At the Oracle Linux yum server repository at <http://yum.oracle.com/>, the kernel image and user-space packages are available on the following channels for Oracle Linux 6:

- `ol6_latest` (latest user-space packages for Oracle Linux 6 other than the OFED tool packages)
- `ol6_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `ol6_UEKR4_OFED` (latest OFED tools packages)

For Oracle Linux 7:

- `ol7_latest` (latest user-space packages for Oracle Linux 7 other than the OFED tool packages)
- `ol7_UEKR4` (`kernel-uek*`, `dtrace-modules-*`, and `libdtrace-*`)
- `ol7_UEKR4_OFED` (latest OFED tools packages)



Note

To be able to install UEK R4, enable the appropriate `ol6_UEKR4` or `ol7_UEKR4` channel and disable the `ol6_UEKR3_latest` or `ol7_UEKR3` channel.

The DTrace utility and DRBD packages are not available on Oracle Linux yum server.

To enable access to the Oracle Linux 6 channels on Oracle Linux yum server, create entries such as the following in `/etc/yum.conf` or in a repository file in the `/etc/yum.repos.d` directory:

```
[ol6_latest]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
```

```
gpgcheck=1
enabled=1

[ol6_UEK_latest]
name=Latest Unbreakable Enterprise Kernel for Oracle Linux $releasever ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/UEK/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0

[ol6_UEKR4]
name=Latest Unbreakable Enterprise Kernel Release 4 for Oracle Linux $releasever ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/UEKR4/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1

[ol6_playground_latest]
name=Latest mainline stable kernel for Oracle Linux 6 ($basearch) - Unsupported
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/playground/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0

[ol6_UEKR4_OFED]
name=OFED supporting tool packages for Unbreakable Enterprise Kernel Release 4 on Oracle Linux 6 ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL6/UEKR4/OFED/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0
priority=20
```

To enable a channel, set the value of the `enabled` parameter for the channel to 1.

To disable a channel, set the value of the `enabled` parameter for the channel to 0.

In this example, access is enabled to the `ol6_latest` and `ol6_UEKR4` channels but not to the `ol6_UEK_latest`, `ol6_playground_latest` and `ol6_UEKR4_OFED` channels.

To enable access to the Oracle Linux 7 channels, create entries such as the following:

```
[ol7_latest]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL7/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1

[ol7_UEKR4]
name=Latest Unbreakable Enterprise Kernel Release 4 for Oracle Linux $releasever ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL7/UEKR4/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1

[ol7_UEKR4_OFED]
name=OFED supporting tool packages for Unbreakable Enterprise Kernel Release 4 on Oracle Linux 7 ($basearch)
baseurl=http://yum.oracle.com/repo/OracleLinux/OL7/UEKR4/OFED/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0
priority=20
```

In this example, access is enabled to the `ol7_latest` and `ol7_UEKR4` channels but not to the `ol7_UEKR4_OFED` channel.

You can find more information about installing the software at <http://yum.oracle.com/>, from where you download a copy of a suitable repository file (<http://yum.oracle.com/public-yum-ol6.repo> or <http://yum.oracle.com/public-yum-ol7.repo>).

3.4 Upgrading Your System

To upgrade your system to UEK R4:

1. After enabling access to the appropriate channels, including `ol6_UEKR4` or `ol7_UEKR4`, in the Oracle Linux yum server repository or `ol6_x86_64_UEKR4` or `ol7_x86_64_UEKR4` on ULN, run the following command to upgrade the system to UEK R4:

```
# yum update
```

2. After upgrading the system, reboot it, selecting the UEK R4 kernel (version 4.1.12) if this is not the default boot kernel.

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

See http://docs.oracle.com/cd/E37670_01/E41138/html/ol_bootconf.html for more information on updating the default boot kernel on Oracle Linux 6.

For instructions on how to install the Oracle-supported OFED packages after upgrading to UEK R4, see [Section 3.5, “Installing the Oracle-Supported OFED Packages”](#).

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

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

3.5 Installing the Oracle-Supported OFED Packages

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

To install the OFED packages provided by Oracle:

1. If your system is registered with ULN, subscribe the system to the `ol6_x86_64_UEKR4_OFED` or `ol7_x86_64_UEKR4_OFED` channel on ULN as appropriate.

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

- a. Move the existing yum repository file to a backup file, for example:

```
# mv /etc/yum.repos.d/public-yum-olN.repo /etc/yum.repos.d/public-yum-olN.repo.bck
```

- b. Download the latest yum repository file for Oracle Linux 6 or Oracle Linux 7 from Oracle Linux yum server.

```
# wget -O /etc/yum.repos.d/public-yum-olN.repo http://yum.oracle.com/public-yum-olN.repo
```

- c. Edit the downloaded yum repository file and enable the `ol6_UEKR4_OFED` or `ol7_UEKR4_OFED` repository as appropriate.

2. Remove any existing OFED packages.

```
# yum remove 'ibacm*'
# yum remove 'ib-bonding*'
# yum remove 'ibutils*'
# yum remove 'infiniband-diags*'
# yum remove 'libibcm*'
# yum remove 'libibmad*'
# yum remove 'libibumad*'
# yum remove 'libibverbs*'
# yum remove 'libmlx4*'
# yum remove 'librdmacm*'
# yum remove 'libsdp*'
# yum remove 'mstflint*'
# yum remove 'ofed-docs*'
# yum remove 'ofed-scripts*'
# yum remove 'opensm*'
# yum remove 'perfctest*'
# yum remove 'qperf*'
# yum remove 'sdpnstat*'
# yum remove 'rdma-*'
# yum remove 'rds-tools*'
```

3. Clean all yum cached files from all enabled repositories.

```
# yum clean all
```

4. Install the OFED packages for UEK R4.

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

5. If the SDP feature is required, install the `libsdp` package.

```
# yum install libsdp
```

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

```
# chkconfig rdma on
```



Note

Each UEK release requires a different set of OFED packages. If you change the kernel on your system to a UEK release other than UEK R4, use the following command to remove the existing UEK R4-based OFED packages before installing the correct packages for the new kernel:

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

