

# Oracle Linux 7

## Release Notes for Oracle Linux 7.4



E88149-17  
April 2026



Oracle Linux 7 Release Notes for Oracle Linux 7.4,  
E88149-17

Copyright © 2022, 2026, Oracle and/or its affiliates.

# Contents

## Preface

---

Conventions	i
Documentation Accessibility	i
Access to Oracle Support for Accessibility	i
Diversity and Inclusion	ii

## 1 System Requirements and Limits

---

File System, Storage, and Address Space Limitations	1
---	---

## 2 Shipped Kernels

---

## 3 New Features and Changes

---

Booting	1
Desktop	1
Development Tools	2
File Systems	2
btrfs: Deprecated in RedHat Compatible Kernel (RHCK)	3
xfs: d_type support (ftype=1) enabled by default on newly formatted partitions	3
Installation	4
Kernel	5
Significant Changes to Kernel Entries and Parameters	7
Networking	10
Packaging	10
Security	11
Server and Services	12
Spacewalk Client Registration	13
Storage	13
Multipathing Improvements and Changes	14
Support Tools	15
Virtualization	15
Technology Preview	16

## 4 Fixed Issues

---

File Systems Issues Fixed	1
btrfs: UEK R3 incompatibility issue fixed	1
btrfs: Empty symbolic link after fsync of parent directory issue fixed	1
xfs: Kernel panic in multi-block buffer logging code issue fixed	1
btrfs: Kernel warning after snapshot is created with incorrect qgroup fixed	1
Co-Existing IPv4 and IPv6 VxLAN issue for UEK R4u2 fixed	1
kmod fixes for module compatibility checking	2

## 5 Known Issues

---

Installation Issues	1
Configuring Encryption and /boot During Installation	1
Network Installation	2
Installation on an iSCSI Disk	2
Installation on an HP 3PAR TPVV	2
Unable to boot after installation on systems using the Broadcom MegaRAID 9460 controller	2
Upgrade Issues	3
cgconfig and cgrd packages must be restored separately	3
Postupgrade fails if web proxy is required	3
Using old version of yum causes dependency errors	3
Package Conflicts	3
rdma-core and infiniband-diags when installing oracle-ofed-release	3
dovecot-devel.i686 and dovecot-devel.x86_64	3
ipa-server-dns.x86_64 and freeipa-server-dns	4
ipa-admintools.x86_64 and freeipa-admintools	4
rear.x86_64 and rear.noarch	4
PackageKit.i686 and PackageKit.x86_64	4
sssd-common.i686 and sssd-common.x86_64	5
File System Related Bugs	5
AutoFS: AMD map browsable_dirs option does not work unless it is set in the [amd] section of autofs.conf	5
btrfs: Send operation causes soft lockup on large deduped file	6
btrfs, ext4 and xfs: Kernel panic when freeze and unfreeze operations are performed in multiple threads	6
btrfs: qgroup reserve space leaks	6
btrfs: Incorrect exclusive reference count after cloning file between subvolumes	6
btrfs: Kernel oops when unmounting during a quota rescan or disable	6
btrfs: Kernel oops when removing shared extents using qgroup accounting	6

ext4: System hangs on unmount after an append to a file with negative i_size	7
ext4: hang occurs during dynamic expansion of inode size	7
ext4: System hang when processing corrupted orphaned inode list	7
xfs: Oracle Linux 7 Update 4 is incompatible with UEK R3 where file systems are formatted using XFS with d_type enabled	7
xfs: Directory readahead completions can hang the system after unmount	7
xfs: System hangs on unmount after a buffered append to a file with negative i_size	8
xfs, ext4: IO error during DIO/AIO write results in disk content corruption	8
Ceph: SSL certificate verification error when accessing the Ceph Object Gateway in SSL mode	8
Automatic Bug Reporting Tool	8
Auto-completion of Commands in the bash Shell	9
crashkernel=auto setting on UEK R3	9
grubby sets incorrect saved entry	9
grubby fatal error upgrading Kernel when /boot is on a btrfs subvolume	9
Hebrew LaTeX fonts	10
InfiniBand Issues	10
ifup-ib: line 357: /sys/class/net/ib0/acl_enabled: Permission denied error	10
Kdump might fail due to an incorrect InfiniBand Adapter M3 Firmware version	10
Changing the IPoIB mode of an InfiniBand interface	10
Disabling an InfiniBand CA port generates warnings	11
Intel QuickAssist Acceleration Technology	11
Database installation and operation fails if RemoveIPC=yes is configured for systemd	11
Oracle ASM fails to initialize with SELinux in Enforcing mode	12
Multipath messages related to zram on UEK R3	12
Unable to create Oracle Linux 7 LXC containers on NFS	12
Oracle Linux 7 guests on Oracle VM and Xen	13
Hyper-V related services fail to start on Oracle Linux 7 Update 4 guest with UEK R4 kernel under Windows Hyper-V Server	13
Per-CPU Allocation Fails when Loading kvm_intel module with UEK R3	13
OpenSSH does not update login records with ssh client's host name	13
Geneve network driver support not available in UEK releases	13
net_prio control group not supported on UEK R3	14
NetworkManager fails to set the default gateway and route for interfaces configured with DHCP on UEK R3	14
NetworkManager unable to add IPv6 addresses to interfaces on UEK R3	14
Network connection icon reports incorrect state for interfaces	15
Power button defaults to ACPI Suspend	15
Cockpit web interface fails to display subscription status	15
32-bit RDMA packages are installed when upgrading a system with rdma-core installed	15

## 6 Installation and Availability

---

Oracle-Supported OFED Packages	2
Installing or Upgrading Oracle-Supported OFED Packages for UEK R4	2
Upgrading to Oracle Linux 7 Update 4 with the Oracle-Supported OFED Packages for UEK R4 installed	2
Upgrade using ULN	3
Upgrade using the Oracle Linux Yum Server	3

## 7 Package Changes from the Upstream Release

---

Removed Packages	1
Modified Packages	3
New Packages	6
Modified Optional Packages	8
Packages Added by Oracle	8

## 8 Removed Modules

---

# Preface

## Warning

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

[Oracle Linux 7: Release Notes for Oracle Linux 7.4](#) provides a summary of the new features and known issues in Update 4 for Oracle Linux 7. This document may be updated after it is released.

## Conventions

The following text conventions are used in this document:

Convention	Meaning
<b>boldface</b>	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.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

## Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <https://www.oracle.com/corporate/accessibility/>.

## Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

## Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

# 1

## System Requirements and Limits

### **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

You can install Oracle Linux 7 on x86-64 systems with up to 2048 logical CPUs and 64 TB of memory. The theoretical upper limit is 5120 logical CPUs and 64 TB of memory, but Oracle hasn't tested this configuration. A minimum of 2 logical CPUs and 1 GB of memory per logical CPU is recommended. Although the minimum disk space required for installation is 1GB, a minimum of 5 GB is recommended.

## File System, Storage, and Address Space Limitations

The following table describes the maximum file size and maximum file system size for the `btrfs`, `ext4`, and `XFS` file systems. File system limitations are affected by kernel versions and features, and by the architecture of the system where Oracle Linux is installed. The values depicted here are estimates based on the known variables that may affect the maximum theoretical value that can be achieved. The theoretical values may be higher than those depicted here, and the actual achievable values may be below the values shown depending on hardware and the kernel version used.

File System Type	Maximum File Size	Maximum File System Size
<code>btrfs</code>	8 EiB	8 EiB
<code>ext4</code>	16 TiB	1 EiB
<code>XFS</code>	8 EiB	8 EiB

The limits described here for `ext4` are higher than recommended and may prove unstable. If you intend to work with systems where you are working toward higher file system sizes or file sizes, it is recommended that you use either `btrfs` or `XFS`.

The maximum supported size for a bootable logical unit number (LUN) is 50 TB. GPT and UEFI support are required for LUNs that are larger than 2 TB.

The maximum size of the address space that is available to each process is 128 TB.

# 2

## Shipped Kernels

 **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

Oracle Linux 7 Update 4 ships with the following kernel packages:

**kernel-3.10.0-693.e17**  
Red Hat Compatible Kernel (RHCK).

**kernel-uek-4.1.12-94.3.9.e17uek**  
Unbreakable Enterprise Kernel Release 4 update 4 (UEK R4u4), which is the default kernel.

This release of Oracle Linux is tested as a bundle as shipped on the installation media image. When installed from the installation media image, the minimum kernel version supported is the one included in the image. Downgrading kernel packages is not supported, unless recommended by Oracle Support.

The kernel source code for the shipped kernel is available after the initial release via a public git source code repository at <https://oss.oracle.com/git/?p=linux-uek.git>.

# 3

## New Features and Changes

### **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

This section describes new features and changes in Update 4 for Oracle Linux 7.

For details of the new features and changes in the initial release of Oracle Linux 7, see [Oracle Linux 7: Release Notes for Oracle Linux 7](#).

## Booting

This section describes booting features in this release, including improvements, changes, and bug fixes.

- **UEFI Secure Boot**

You can install and use Oracle Linux 7 on systems that have UEFI Secure Boot enabled. A system in Secure Boot mode loads only those boot loaders and kernels that have been signed by Oracle. Oracle has updated the `kernel` and `grub2` packages to sign them with a valid Extended Validation (EV) certificate. The EV certificate has been compiled into the `shim` binary and has been signed by Microsoft. This feature is fully supported on Oracle Linux 7 update 4.

If you have previously enabled Secure Boot while it was available under a technology preview, ensure that the `shim`, `grub2` and `kernel` packages are updated as an atomic operation if you intend to upgrade the system. If all of these packages are not updated, the Secure Boot process might break and must be disabled until a full system upgrade has been completed. (Bug ID 24616226)

- **Updated `shim-signed` package**

The `shim-signed` package is updated to include numerous bug fixes and enhancements over the previously shipped version.

## Desktop

The following desktop features, improvements, and changes are included in this release:

- **GNOME desktop updated to 3.22.3**

This version of the GNOME desktop includes several improvements and bug fixes, including the following:

- Desktop notifications overhauled
- Built-in integration with world clocks and media players

- Automatic screen brightness adjustment capabilities (for systems with an integrated light sensor)
- Standard dialog for documenting key keyboard shortcuts for several applications
- Setting panels improvements (printer, mouse, touchpad, keyboard shortcuts)
- Option for renaming multiple files simultaneously
- Undo support for trash
- Built-in support for compressed files and Google Drive
- **Added `xorg-x11-drv-libinput` driver to X.Org input drivers**

After you install `xorg-x11-drv-libinput`, you can remove the `xorg-x11-drv-synaptics` driver, which enables you to access to some of the improved input device handling features that are offered by `libinput`.
- **`cloud-init` package moved to Base channel**

The `Cloud-init` tool handles the early initialization of a system using metadata that is provided by the environment. You typically use `cloud-init` to configure servers that are booted in a cloud environment, such as OpenStack or Amazon Web Services.

## Development Tools

The following development tools have been updated and improved:

- **`demidecode` package version updated to 3.0**

The updated version of the `demidecode` includes several bug fixes and hardware enablement improvements.
- **TLS version restriction capability added to `IO::Socket::SSL` Perl module**

For improved security, the `Net::SSLey` Perl module has been updated to enable the explicit specification of TLS version 1.1 or 1.2, and the `IO::Socket::SSL` module has been updated accordingly.

When creating a new `IO::Socket::SSL`, you can restrict the TLS version to 1.1 or 1.2 by setting the `SSL_version` option to `TLSv1_1` or `TLSv1_2`, respectively. Alternatively, you can specify the `TLSv11` and `TLSv12` options. Note that these values are case-sensitive.
- **TLS version restriction capability added to `Net::SSLey` Perl module**

For improved security, the `Net::SSLey` Perl module has been updated to enable the explicit specification of TLS version 1.1 or 1.2. To restrict the TLS version, set the `Net::SSLey::ssl_version` variable to `11` or `12`, respectively.
- **TLS version specification capability added to `wget`**

Previously, the `wget` command used the highest TLS version (1.2) by default. In this update, the `wget` command has been enhanced to enable you to explicitly select the TLS protocol minor version by specifying either the `--secure-protocol=TLSv1_1` or `--secure-protocol=TLSv1_2` options with the `wget` command.

## File Systems

The following file systems features have been updated and improved:

- **`autofs` browse options added for `amd` format maps**

You can now add mount point sections to the `autofs` configuration for `amd` format mounts, similarly to how automount points are configured in `amd`, without the need to also add a corresponding entry to the master map. This improvement helps to avoid having incompatible master map entries in the `autofs` master map within shared multi-vendor environments.

You can use the `browsable_dirs` option in either the `autofs [amd]` configuration section, or following the `amd` mount point sections. You can also use the `browsable` and `utimeout` map options of `amd` type `auto` map entries.

For information about an issue related to using the `browsable_dirs` option, see [AutoFS: AMD map browsable\\_dirs option does not work unless it is set in the \[amd\] section of autofs.conf](#).

- **Capability for adding mount request log entries in `autofs` configuration**

By enabling the adding of a mount request log identifier to the mount request log entries in the `autofs` configuration, you can quickly filter entries for specific mount requests. The improvement makes searching logs easier.

- **`rpc.idmapd` capability for obtaining NFSv4 ID domains from the Domain Name System (DNS)**

In the event that an NFSv4 ID map domain name is not configured on the system, this feature enables the NFS `idmapping` library to attempt to obtain the proper domain name by performing a DNS lookup of a special TXT record. If the TXT record is not present, it uses other heuristics to obtain the proper domain name.

- **Added support for Kerberos authentication for NFSoRDMA client and server**

This improvement enables you to use `krb5`, `krb5i`, and `krb5p` authentication with NFS over RDMA (NFSoRDMA) features, for both client and server. You can now use Kerberos with NFSoRDMA to securely authenticate each Remote Procedure Call (RPC) transaction.

**Note**

To use Kerberos with NFSoRDMA, you must install the `nfs-utils` package, version 1.3.0-0.36 or higher.

- **`SEEK_DATA` and `SEEK_HOLE` Options for FUSE `lseek` System Call**

The `SEEK_DATA` and `SEEK_HOLE` are now available for the Filesystem in Userspace (FUSE) `lseek` system call when using the RedHat Compatible Kernel (RHCK). Use the `SEEK_DATA` option to adjust the file offset to the next location in the file that contains data. Use the `SEEK_HOLE` option to adjust the file offset to the next hole in the file, greater than or equal to the offset. Note that this functionality is not available in UEK at the time of this update release.

## btrfs: Deprecated in RedHat Compatible Kernel (RHCK)

As of Oracle Linux 7 update 4, `btrfs` is deprecated in the RHCK. With UEK R4, `btrfs` is fully supported.

## xfs: `d_type` support (`f_type=1`) enabled by default on newly formatted partitions

For systems installed with the Oracle Linux 7 Update 4 installer, when formatting a device using XFS, `d_type` support is enabled automatically, which means all XFS-formatted partitions are created using the `f_type=1` parameter as the default. Whereas, in previous Oracle Linux 7

updates, `f_type=0` was the default parameter, meaning `d_type` was disabled or off and XFS-formatted partitions were created using `f_type=0` as the default.

The `d_type` functionality exposed by this feature enables the file system to store additional metadata that is critical for overlay file system types.

## Installation

Several changes, bug fixes and improvements have been made to the installation process in this update release. These include:

- **Change to kickstart parameters to support specification of RAID chunk size**

Changes were implemented in the installer to enable the ability to set RAID chunk sizing in a kickstart file using the `--chunksize` parameter. This update allows tuning for performance when using RAID.

- **Added kickstart support for thin LVM snapshots during installation**

The new kickstart `snapshot` command creates an LVM thin volume snapshot before or during installation. To use this functionality, specify all of the required parameters for the command. For example:

```
snapshot <origin_vg/origin_lv> --name=<snapshot_name> --when=<pre-install | post-install>
```

- **Change to automatic partitioning behavior for LVM thin pools**

Changes to automatic partitioning behavior where LVM thin pools are created during installation are important to note.

LVM thin pools created with automatic partitioning reserve 20% of the volume group size and require a minimum of 1GiB and a maximum of 100 GiB.

The `logvol --thinpool --grow` command causes the thin pool to grow to the maximum possible size. To reserve space for the volume group, use the `volgroup --reserved-space` or `volgroup --reserved-percent` command to specify the amount of space to keep available for the volume group.

- **Added kickstart option to disable the creation of a /home partition**

The `--nohome` option can be used with the `autopart` command in a kickstart installation to prevent the creation of a partition designated for `/home` use.

- **Added support for loading driver disks from hard disk or USB device**

Support has been added to enable loading a driver disk from a hard disk or USB device. This can be triggered either via kickstart or as a boot option. To use this option you must set the label for the device where the driver disk RPM files are stored. To load a driver from the specified driver disk, use:

```
driverdisk LABEL=<LABEL>:<driver.rpm>
```

Substitute `<LABEL>` with the label that you set for the device and substitute `<driver.rpm>` with the driver RPM file name.

To specify the driver disk as a boot option, use:

```
inst.dd=hd:LABEL=<LABEL>:<driver.rpm>
```

Substitute `<LABEL>` with the label that you set for the device and substitute `<driver.rpm>` with the driver RPM file name.

- **Added support for IP over InfiniBand (IPoIB) in text mode installation**

The text mode installer now supports IPoIB network interfaces during a manual installation. IPoIB interface status information and configuration options are available.

- **Improvements to cater for multiple network locations for stage2 or kickstart files to provide failover during installation**

The installer is now capable of handling multiple `inst.stage2` and `inst.ks` boot options where those options point to alternate network locations. This caters to a scenario where the network location for either stage2 or kickstart file is not available and a failover may be required for installation to continue. Options are processed sequentially until all location options are exhausted. If a file system is specified as one of the locations for either of these options only the last location specified is used, regardless of whether that location is a file system or URL.

- **Improved debug functionality for Anaconda installation issues**

The new `inst.debug` boot option can be used to start the Anaconda installer in debug mode. This option stores log files for `lsblk`, `dmesg` and `lvm dump` in the `/tmp/pre-anaconda-logs` directory to help with debugging installation issues.

- **Fix to enable Lorax to ignore SSL errors**

The `lorax` tool, which is used to create an Anaconda installer `boot.iso` and the release tree and related metadata, has the new `--noverifyssl` command line switch to disable SSL certificate verification, allowing the tool to be used with systems using self-signed certificates.

## Kernel

The following changes are specific to the RedHat Compatible Kernel (RHCK). For more information, refer to latest versions of the release notes for Unbreakable Enterprise Kernel Release 4 at [Unbreakable Enterprise Kernel documentation](#).

- **crash package version updated to 7.1.9**

The updated version of the `crash` utility includes a number of bug fixes and enhancements from the previous version.

- **New dbxtool package**

The `dbxtool` package provides a command-line interface (CLI) and a one-shot `systemd` service for applying UEFI Secure Boot DBX updates.

- **fjes driver updated to version 1.2**

The updated version of the `fjes` driver includes a number of bug fixes and enhancements from the previous version.

- **Added getrandom system call to kernel**

The `getrandom` system call has been added to the kernel. As a result, the user space can now request randomness from the same non-blocking entropy pool that is used by `/dev/urandom`. In addition, the user space can block until at least 128 bits of entropy has been accumulated in that pool.

- **Changes to hardware utility tools to correctly identify recently released hardware**

The PCI, USB, and vendor device identification files have been updated. As a result, the hardware utility tools can now correctly identify recently released hardware.

- **Added i40e support for trusted and untrusted virtual functions**

The `i40e` NIC driver now includes support for both trusted and untrusted virtual functions.

- **Addition of the Intel Cache Allocation Technology**

The Intel Cache Allocation Technology enables the software to restrict cache allocation to a defined subset of cache. The defined subset can overlap with other subsets.
- **Jitter Entropy Random Number Generator included**

The Jitter Entropy Random Number Generator (RNG) is responsible for collecting entropy through CPU timing differences for the kernel. By default, this RNG is available through the `algif_rng` interface. The generated numbers can be added back to the kernel through the `/dev/random` file, which makes these numbers available to other `/dev/random` users, thus making the operating system have more sources of entropy available.
- **macsec driver added**

The `macsec` driver enables support for the MACsec/IEEE 802.1AE network device. This driver provides authentication and encryption of traffic in a LAN, typically with GCM-AES-128 and optional replay protection. Patches have also been applied to bring this version of the driver up to the most current level for compatibility with this kernel release. The `iproute` package has also been updated to include support for the `ip macsec` command and related functionality.
- **makedumpfile updated to version 2.0.14-1**

This version of the `makedumpfile` utility includes a number of bug fixes and enhancements from the previous version.
- **NVMe driver updated to version 4.10**

The updated version of the `NVMe` driver includes a number of bug fixes and enhancements from the previous version.
- **nvme-cli package version updated to 1.1**

The updated version of the `nvme-cli` utility includes support for Nonvolatile Memory Express (NVMe). With NVMe support, you can find targets over Remote Direct Memory Access (RDMA) and connect to these targets.
- **Added perf support for uncore events on Intel Xeon v5**

The `perf` performance analysis tool now includes support for uncore events on the Intel Xeon v5 server CPU. These events provide additional performance monitoring information.
- **Random driver (/dev/random) displays messages pertaining to urandom pool initialization**

The random driver (`/dev/random`) now prints a message when the non-blocking pool that is used by `/dev/urandom` is initialized.
- **Change to spinlock implementation in the kernel**

The spinlock implementation in the kernel has changed from ticket spinlocks to queued spinlocks on AMD64 and Intel 64 architectures. Because queued spinlocks are more scalable than the ticket spinlocks, system performance is improved, especially on Symmetric Multi Processing (SMP) systems with large number of CPUs. The performance now increases more linearly with an increasing number of the CPUs.

**Note**

Note that because of this change in the spinlock implementation, kernel modules that are built on Red Hat Enterprise Linux 7 might not be loadable on kernels from earlier releases. Kernel modules released in Red Hat Enterprise Linux (RHEL) versions earlier than 7.4 are loadable on the kernel that is released in RHEL 7.4.

- **Added functionality for `switchdev` infrastructure and `mlxsw` driver**

The following functionality has been added in this update:

- Ethernet switch device driver model (`switchdev` infrastructure)

Switch devices can now offload forwarding data plane from the kernel.

- `mlxsw` driver support

The following switch hardware is supported by the `mlxsw` driver: Mellanox SwitchX-2 (slow path only), Mellanox SwitchIB and SwitchIB-2, and Mellanox Spectrum.

Features that are supported by the `mlxsw` driver include the following:

- \* Per port jumbo frames
- \* Speed setting, state setting, statistics
- \* Port splitting together with splitter cables
- \* Port mirroring
- \* QoS: 802.1p, Data Center Bridging (DCB)
- \* Access Control Lists (ACLs) using TC flower offloading

Note that this feature is introduced as a Technology Preview.

- Layer 2 and Layer 3 features:

Layer 2:

- \* Virtual local area networks (VLANs)
- \* Spanning Tree Protocol (STP)
- \* Link Aggregation (LAG) using team or bonding offloading
- \* Link Layer Discovery Protocol (LLDP)

Layer 3 now includes the unicast feature.

You can configure these features by using the standard tools that are provided by the `iproute` package, which has also been updated in this release.

## Significant Changes to Kernel Entries and Parameters

The following is a summary of significant changes in the kernel that is shipped with the RHCK for Oracle Linux 7.4. Included are new or updated `proc` entries, `sysctl` and `sysfs` default values, boot parameters, kernel configuration options, as well as other notable behavior changes.

**Table 3-1 Updated /proc/sys/kernel Entries**

Kernel Entry	Description	Format
hung_task_panic	Controls the behavior of the kernel when an unresponsive task is detected. This file occurs if CONFIG_DETECT_HUNG_TASK is enabled.	{ "0"   "1" } 0 - Continue operation (Default behavior). 1 - Panic immediately.
hung_task_check_count	Provides the upper bound on the number of tasks that are checked. This file occurs if CONFIG_DETECT_HUNG_TASK is enabled.	N/A
hung_task_timeout_secs	Checks interval. Reports a warning in case that a task in D state is not scheduled for longer time than this value. This file occurs if CONFIG_DETECT_HUNG_TASK is enabled.	0 - Infinite timeout. No checking done.
hung_task_warning	Provides the maximum number of warnings to report during a check interval. When this value is reached, no more warnings will be reported. This file occurs if CONFIG_DETECT_HUNG_TASK is enabled.	-1 - Reports an infinite number of warnings.
panic_on_rcu_stall	When set to 1, calls the panic() function after RCU stall detection messages. This is useful to define the root cause of RCU stalls using a vmcore.	0 - Do not panic when RCU stall takes place (Default behavior). 1 - Panic after printing RCU stall messages.

Files in the `/proc/sys/user` directory can be used to override the default limits for the number of namespaces and other objects that have per-user namespace limits. These limits are used to stop programs that malfunction and attempt to create a high number of objects. The default values of these limits are adjusted so that any program in normal operation cannot reach them.

**Table 3-2 Updated /proc/sys/user Entries**

Updated file	Description
max_cgroup_namespaces	Maximum number of cgroup namespaces that any user in the current user namespace can create.
max_ipc_namespaces	Maximum number of ipc namespaces that any user in the current user namespace can create.
max_mnt_namespaces	Maximum number of mount namespaces that any user in the current user namespace can create.
max_net_namespaces	Maximum number of network namespaces that any user in the current user namespace can create.

**Table 3-2 (Cont.) Updated /proc/sys/user Entries**

Updated file	Description
max_pid_namespaces	Maximum number of pid namespaces that any user in the current user namespace can create.
max_user_namespaces	Maximum number of user namespaces that any user in the current user namespace can create.
max_uts_namespaces	Maximum number of user namespaces that any user in the current user namespace can create.

**Table 3-3 Kernel Parameter Changes**

Kernel Parameter	Description and Format
acpi_force_table_verification [HW,ACPI]	Enables table checksum verification during early stage. By default, disabled due to x86 early mapping size limitation.
acpi_no_static_ssdt [HW,ACPI]	Disables the installation of static SSDTs at early boot time. By default, SSDTs contained in the RSDT/XSDT are installed automatically and they appear in the <code>/sys/firmware/acpi/tables</code> directory.  This option turns off this feature. Specifying this option does not affect dynamic table installation which installs SSDT tables to the <code>/sys/firmware/acpi/tables/dynamic</code> directory.
irqaffinity= [SMP]	Sets the default irq affinity mask.  Formats: <i>cpu number, ..., cpu number</i> <i>cpu number-cpu number</i> Or, you can use a positive range in ascending order or a mixture: <i>cpu number, ...,cpu number-cpu number</i>
nokaslr [KNL]]	Disables installation of static SSDTs at early boot time. By default, SSDTs contained in the RSDT/XSDT are installed automatically and they appear in the <code>/sys/firmware/acpi/tables</code> directory.  Disables kernel and module base offset Address SpaceLayout Randomization (ASLR) if <code>CONFIG_RANDOMIZE_BASE</code> is set.
nohibernate	Disables hibernation and resume.
crash_kexec_post_notifiers	Runs <code>kdump</code> after running <code>panic-notifiers</code> and dumping <code>kmsg</code> .
[PCI] hpbussize= <i>nn</i>	Provides the minimum amount of additional bus numbers reserved for buses below a hotplug bridge (Default is 1).

**Table 3-3 (Cont.) Kernel Parameter Changes**

Kernel Parameter	Description and Format
<code>pcie_port_pm=[PCIE]</code>	PCIe port power management handling. Format: { "off"   "force" }  off - Disables power management of all PCIe ports. 1 - Enabled power management of all PCIe ports.
<code>sunrpc.svc_rpc_per_connection_limit=[NF S,SUNRPC]</code>	Limits the number of requests for the server to process in parallel from a single connection( Default value is 0 (no limit)).

## Networking

Networking features, changes, and bug fixes in this release include the following.

- iproute package includes changing bridge port options**  
 In this update, changing bridge port options, such as `state`, `priority`, and `cost`, are included in the `iproute` package. This change enables you to use the `iproute` package as an alternative to the `bridge-utils` package.
- Load Balancing and High Availability**  
 Oracle Linux 7 includes the Keepalived and HAProxy technologies for balancing access to network services while maintaining continuous access to those services.  
  
 Keepalived uses the IP Virtual Server (IPVS) kernel module to provide transport layer (Layer 4) load balancing, redirecting requests for network-based services to individual members of a server cluster. IPVS monitors the status of each server and uses the Virtual Router Redundancy Protocol (VRRP) to implement high availability.  
  
 HAProxy is an application layer (Layer 7) load balancing and high availability solution that you can use to implement a reverse proxy for HTTP and TCP-based Internet services.  
  
 For more information, see [Oracle Linux 7: Administrator's Guide](#).
- Support for MACsec (802.1AE) added to NetworkManager**  
 The `wpa_supplicant` utility now supports the Media Access Control Security (MACsec) encryption 802.1AE, which enables MACsec to be used in configuration by default. This change provides a convenient way to deploy MACsec.
- Packages related to rdma consolidated into rdma-core version 13**  
 Several packages that are related to the `rdma` package have been upgraded and consolidated into a single source package, `rdma-core` version 13.

## Packaging

The following packaging additions and changes are included this release.

- payload\_gpgcheck Option Added to yum**

The new `payload_gpgcheck` option enables `yum` to perform a GNU Privacy Guard (GPG) signature check on the payload sections of packages. This capability provides enhanced security and integrity when installing packages.

Before, when the `gpgcheck` option was used, `yum` only checked package headers. In the event that the payload data were tampered with or somehow corrupted, and an RPM unpacking error occurred, the package would only be partially installed. As a result, the operating system could be inconsistent or in a vulnerable state. You can use the `payload_gpgcheck` option with the `gpgcheck` or `localpkg_gpgcheck` option to prevent this problem from occurring.

Note that using the `payload_gpgcheck` option is the same as manually running the `rpm -K` command on downloaded packages.

## Security

This section describes new, changed, and improved security features.

- **New NBDE security packages**

The following new security packages are provided for the Network Bound Disk Encryption (NBDE) feature. NBDE enables you to encrypt root volumes of hard drives on physical machines without requiring you to manually enter a password when the systems are rebooted.

- `clevis` – Is a pluggable framework for automated decryption. You can use `clevis` to provide an automated decryption of data or even an automated unlocking of LUKS volumes. The `clevis` package provides the client side of the NBDE project.
- `jose` – Is a C-language implementation of the Javascript Object Signing and Encryption standards. The `jose` package is a dependency of the `clevis` and `tang` packages.
- `luksmeta` – LUKSMeta is a simple library for storing metadata in the LUKSv1 header. The `luksmeta` package is a dependency of the `clevis` and `tang` packages.
- `tang` – Is a server for binding data to a network presence. the `tang` package includes a daemon that provides cryptographic operations for binding to a remote service. The `tang` package provides the server side of the NBDE project.

- **New `http-parser` package**

The `http-parser` package provides a utility for parsing HTTP messages (both requests and responses). The parser is designed for use in performance HTTP applications. The parser does not make any system calls or allocations, does not buffer data, and can be interrupted at any time. Depending on your architecture, the parser only requires about 40 bytes of data, per message stream.

- **New `usbguard` package**

The USBGuard software framework provides system protection against intrusive USB devices by implementing basic *allowlisting* and *blocklisting* capabilities that are based on device attributes. To enforce a user-defined policy, USBGuard uses the Linux kernel USB device authorization feature.

The USBGuard framework provides the following components:

- **Daemon** – Is the component with an inter-process communication (IPC) interface that is used for dynamic interaction and policy enforcement.
- **Command-line interface** – Is the component that interacts with a running USBGuard instance.

- Rule language – Is the component that is used for writing USB device authorization policies.
- C++ API – Is the component that interacts with the daemon component that is implemented in a shared library.
- **Updated security package versions**

The versions of the following security package have been updated. The updated version provides a number of new features, improvements, and bug fixes:

  - audit version updated to 2.7.6
  - libica version updated to 3.0.2
  - libreswan version updated to 3.20
  - opensc version updated to 0.16.0
  - openssh version updated to 7.4
  - openssl version updated to 1.0.2k
  - openssl-ibmca version updated to 1.3.0
- **Modification to openSSH to use SHA-2 for public key signatures**

By default, the algorithm for public key signatures that is used in this release is SHA-2. Note that SHA-1 is available for backward compatibility purposes *only*.
- **pmrfc3164 replaces pmrfc3164sd in rsyslog**

The `pmrfc3164sd` module, which is used for parsing logs in the BSD `syslog` protocol format (RFC 3164), has been replaced by the official `pmrfc3164` module in this update.

 **Note**

Because the `pmrfc3164` module does not fully cover `pmrfc3164sd` functionality, the `pmrfc3164sd` module is still available in `rsyslog`. However, whenever possible, you should use the new `pmrfc3164` module, as the `pmrfc3164sd` module is no longer supported.

## Server and Services

The following server and services improvements and changes have been made:

- **New libfastjson package**

The `libfastjson` library replaces the `json-c` library for `rsyslog` in this update. The `libfastjson` library includes a limited feature set that provides significantly improved performance, compared to `json-c`.
- **New cache configuration options for mod\_nss**

New options for controlling caching of Offensive Security Certified Professional (OCSF) responses have been added to the `mod_nss` module.

You can use these new options to control the following:

  - Time to wait for OCSF responses.
  - Size of the OCSF cache.

- Minimum and maximum duration for an item's presence in cache, including not caching at all.
- **Server and service package version updates**

The following package versions have been updated. These updated versions include various enhancements and bug fixes:

  - chrony version updated to 3.1
  - rear version updated to 2.0
  - rsyslog version updated to 8.24.0
  - tuned version updated to 2.8.0
- **Change to default state file path for logrotate**

To prevent confusion and potential mismatching of paths, the default state file path that is used by `logrotate` has been changed to match the state file path that is used by the `logrotate cron job`. As a result, `logrotate` now uses `/var/lib/logrotate/logrotate.status` as the default state file path in both scenarios.
- **Removed `nss_pcache` options**

The `nss_pcache` pin-caching service no longer shares the Network Security Services (NSS) database of the `mod_nss` Apache module because `nss_pcache` does not need access to the tokens. Also, options for the NSS database and the prefix have been removed and are now handled automatically by `mod_nss`.
- **Expanded support in `openwsman` for disabling SSL protocols**

The `openwsman` utility has been updated to include a new configuration file option for listing disabled protocols. The new option enables you to specifically disable particular SSL protocols.
- **Deprecated `openldap-server`**

Starting with Oracle Linux 7.4, the `openldap-server` package is deprecated and new versions of this package will not be included in the next major release of Oracle Linux. Consider using an alternate LDAP server application included with Oracle Linux, such as the 389 Directory Server.

## Spacewalk Client Registration

It is not necessary to install the Spacewalk client before registering an Oracle Linux 7 Update 4 system with a Spacewalk server. Instead, you can use the `rhreg_ks` command, specifying the CA certificate file for the server, the server URL, and the activation key to be associated with the system.

For detailed instructions, see the Spacewalk 2.6 for Oracle Linux Client Life Cycle Management Guide at [Oracle® Linux Manager & Spacewalk for Oracle® Linux Documentation](#). (Bug ID 20656368)

## Storage

This update includes the following storage features, improvements, and changes.

- **LVM commands for reducing RAID logical volume size added**

As of this update, you can use the Logical Volume Manager (LVM) commands, `lvreduce` or `lvresize`, to reduce the size of a RAID logical volume.

- **Added support in LVM for RAID takeover and reshaping**

LVM now fully supports RAID takeover, which enables users to convert a RAID logical volume from one RAID level to another RAID level. Note that this feature was previously only available as a Technology Preview. In addition, LVM now provides support for RAID reshaping, which enables you to reshape properties such as the RAID algorithm, stripe size, and number of images.

**Note**

The new RAID types that are added by means of RAID takeover or reshape are not supported in older kernel versions. These RAID types include the following: `raid0`, `raid0_meta`, `raid5_n`, and `raid6_{ls,rs,la,ra,n}_6`. Creating or converting to these RAID types on RHCK for Oracle Linux 7.4 cannot activate the logical volumes on systems that are running previous releases.

- **Capability for changing region size of RAID logical volume added**

You can now change the region size of a RAID logical volume using the `-R/--regionsize` option of the `lvconvert` command. You must also change the old default value set by the `activation.raid_region_size = N` parameter in the existing `lvm.conf` file or the old value will still be applied when you create new logical volumes

## Multipathing Improvements and Changes

The following are new, improved, or changed Multipathing features:

- **New `detect_checker` multipath parameter**

The Multipath feature now supports the `detect_checker` parameter in the `multipath.conf` defaults and devices sections. If the parameter is set, multipath detects whether device supports the Asymmetric Logical Unit Access (ALUA) mode. If so, multipath overrides the configured `path_checker` and uses the Test Unit Ready (TUR) checker instead. The `detect_checker` option enables devices with an optional ALUA mode to be correctly auto configured, regardless of the device's current mode.

- **Support added to `device-mapper-multipath` for `max_sectors_kb` configuration parameter**

The `device-mapper-multipath` resource includes a new `max_sectors_kb` parameter in the defaults, devices, and multipaths sections of the `multipath.conf` file. This new parameter enables you to set the `max_sectors_kb` device queue parameter to the specified value on all underlying paths of a multipath device before the multipath device is first activated.

When a multipath device is created, it inherits the `max_sectors_kb` value from the path devices. Manually raising or lowering this value for the multipath device can cause multipath to create I/O operations that are larger than the path devices allow. The addition of the `max_sectors_kb` `multipath.conf` parameter provides a way to set these values before a multipath device is created on top of the path devices, thus preventing invalid sized I/O operations from being passed down.

- **New `disabled_changed_wwid` multipath configuration parameter**

The Multipath feature now includes a new `disable_changed_wwid` parameter that you can set in the default section of the `multipath.conf` file. When this parameter is set, `multipathd` notes whenever a path device changes its `wwid` while it is in use, and then disables access to that device until its `wwid` returns its previous value.

- **New multipathd commands for resetting device statistics**

In this update, two new `multipathd` commands are introduced: `multipathd reset multipaths stats` and `multipathd reset multipath dev stats`. You use these commands to reset the device statistics that `multipathd` tracks for all devices, or a specified device, respectively. This capability enables you to reset device statistics after making changes to them.
- **New `remove_retries` multipath configuration value**

You can now control the number of times that the `multipath` command tries to remove a multipath device that is busy. You enable this capability by changing the `remove_retries` configuration value from its default value of 0, as when the value is set to 0, `multipath` will not retry any failed removes.
- **Warning messages printed when `multipathd` is not running**

The `multipathd` daemon now prints a warning message if you run a `multipath` command that creates or lists multipath devices while `multipathd` is not running.

## Support Tools

Oracle Linux 7 includes tools to assist with the resolution of runtime issues. Notable features and changes in this update are as follows:

- **Kdump Configuration During Installation**

It is now possible to configure Kdump during a non-graphical installation. For limitations on using the `crashkernel=auto` setting, see [crashkernel=auto setting on UEK R3](#).
- **makedumpfile Support for Large Memory Images**

`makedumpfile` can now use `sadump` format for dumps of more than 16 TB of physical memory.
- **Kpatch Removed**

The upstream Kpatch RPM has been removed from Oracle Linux. Customers who want to patch their running kernel with zero downtime should evaluate Oracle's Ksplice technology, which is included at no additional cost with Oracle Linux Premier support. For more information, see [Oracle Linux: Ksplice User's Guide](#).

## Virtualization

This section describes new, improved, and updated virtualization features.

- **KVM and QEMU support for new features in 2nd Generation Xeon and Xeon Phi processors**

The Kernel-based Virtual Machine (KVM) modules and the QEMU hypervisor are now capable of supporting the new features that are present in 2nd Generation Xeon and Xeon Phi processors. KVM guests can use the `avx512_4vnniw` and `avx512_4fmops` instructions if they are enabled in the virtual machine CPU configuration.
- **Configuring MTU settings on KVM guest interfaces added**

In this update, you have the ability to configure MTU settings on KVM guest interfaces.
- **libvirt changed to use generic PCIe root ports in QEMU**
- **libvirt version updated to 3.2.0**

This update makes it possible to install and uninstall specific `libvirt` storage sub-drivers, thereby reducing the installation footprint. In addition, you can now configure the `/etc/nsswitch.conf` file to instruct the Name Services Switch (NSS) to automatically resolve names of KVM guests to their network addresses.

- **Added support in KVM for MCE**

Support for Machine Check Exception (MCE) has been added to the KVM kernel modules. It is now possible to use the Local MCE (LMCE) feature of Intel Xeon v5 processors in KVM guest virtual machines. LMCE can deliver MCE to a single processor thread, instead of broadcasting to all threads, which ensures the machine check does not impact the performance of more vCPUs than is needed. As a result, the software load is reduced when processing MCE on machines with a large number of processor threads.

- **Improved `virt-v2v` installation of QXL drivers**

The `virt-v2v` implementation of QXL driver installation in Windows guest virtual machines has been improved. This change ensures that QXL drivers are installed correctly on these guests.

## Technology Preview

Features that are currently under technology preview when using UEK R4u4 are described in [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#).

For RHCK, the following features are currently under technology preview:

- Systemd:
  - Importd features for container image imports and exports
- File Systems:
  - DAX (Direct Access) for direct persistent memory mapping from an application. This is under technical preview for the ext4 and XFS file systems.
  - Block and object storage layouts for parallel NFS (pNFS).
  - SCSI layout for parallel NFS (pNFS), including support for both client and server configurations.
  - OverlayFS remains in technical preview.
- Kernel:
  - Heterogeneous memory management (HMM).
  - User namespace (security features for isolating Linux containers from the host).
  - 10GbE RoCE Express for RDMA.
  - `ocrdma` and `libocrdma` packages for RDMA over RoCE.
  - No-IOMMU mode virtual I/O feature.
- Networking:
  - Support for a Cisco proprietary User Space Network Interface Controller in UCM servers provided in the `libusnic_verbs` driver
  - Cisco VIC InfiniBand kernel driver that provides similar functionality to RDMA on proprietary Cisco architectures.
  - Trusted Network Connect support.

- Single-Root I/O virtualization (SR-IOV) in the `qlcnict` driver.
- `nftables` and `libnftnl` network filtering and classification functionality
- Storage:
  - Multi-queue I/O scheduling for SCSI (`scsi-mq`). This functionality is disabled by default.
  - The plug-in for the `libStorageMgmt` API used for storage array management. The `libStorageMgmt` API is now fully supported, but the plug-in is under technology preview.
  - DIF/DIX for data integrity checking on SCSI devices, other than certain, specified native HBA and storage hardware. Oracle supports DIF/DIX with UEK R4.

## Compatibility

Oracle Linux maintains user-space compatibility with Red Hat Enterprise Linux, which is independent of the kernel version that underlies the operating system. Existing applications in user space will continue to run unmodified on the Unbreakable Enterprise Kernel Release 4 (UEK R4) 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. UEK R4 contains 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.

# 4

## Fixed Issues

### **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

This chapter describes issues that are fixed in Oracle Linux 7 Update 4.

Note that additional issues specific to the kernel that you are using might also be resolved. If you are using the default UEK R4u4, please see [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#). If you are using an alternate UEK release or update, please refer to the appropriate release notes for this kernel version, available at [Unbreakable Enterprise Kernel documentation](#).

## File Systems Issues Fixed

The following file systems issues are fixed in this release:

### btrfs: UEK R3 incompatibility issue fixed

An incompatibility issue related to btrfs updates and UEK R3, where a btrfs formatted `/root` partition failed to boot, has been fixed. (Bug ID 24840489)

### btrfs: Empty symbolic link after fsync of parent directory issue fixed

The issue that occurs after an fsync of a symbolic link's parent directory, followed by an attempt to mount the file system after a system crash or outage, which results in an empty symbolic link, has been fixed. (Bug ID 23748445)

### xfs: Kernel panic in multi-block buffer logging code issue fixed

The XFS bug in the multi-block buffer logging code that caused a kernel panic at log push time due to invalid regions being set in the buffer log format bitmap has been fixed. (Bug ID 24400444)

### btrfs: Kernel warning after snapshot is created with incorrect qgroup fixed

The bug that caused the file system to crash and a kernel warning when a snapshot was created with an incorrect qgroup has been fixed. (Bug ID 24716895)

## Co-Existing IPv4 and IPv6 VxLAN issue for UEK R4u2 fixed

The issue, where VxLANs that are configured for both IPv4 and IPv6 cannot exist on the same host because of an inability to bind the VxLAN tunnel on the same port, and due to the way in which IPv6 sockets lists for IPv4 traffic has been fixed. This fix is made available in UEK R4u4, which is the default in UEK release in Oracle Linux 7 Update 4. (Bug ID 24579830)

## kmod fixes for module compatibility checking

An issue in the upstream kmod package that was causing the kmod `weak-modules` script to incorrectly validate whether a module is compatible with the Kernel Application Binary Interface (KABI) is fixed in this release of Oracle Linux. Changes include performing the check at the same time when the `--dry-run` option is used and improvements to handling of printing of non-compatibility messages. This issue caused problems for users attempting to install or load ACFS drivers on RedHat Linux 7.4 systems. (Bug ID 26320387)

# 5

## Known Issues

### **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

This chapter describes the known issues for Oracle Linux 7 Update 4.

Note that additional issues specific to the kernel that you are using may also be present. If you are using the default UEK R4u4, please see [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#). If you are using an alternate UEK release or update, please refer to the appropriate release notes for this kernel version, available at [Unbreakable Enterprise Kernel documentation](#).

## Installation Issues

The following sections describe issues that might be encountered during installation.

### Configuring Encryption and /boot During Installation

During installation, if you select **Encrypt my data** on the Installation Destination screen and then perform manual partitioning, the **Encrypt** check box is not shown as selected on the Manual Partitioning screen. This check box refers to encryption that you can configure on a file system type that supports encryption or on an LVM logical volume that contains the file system. If you click **Modify**, the **Encrypt** check box on the Configure Volume screen is shown as selected for the volume, meaning that the encryption will be applied at the level of the underlying block device.

For LVM, selecting **Encrypt my data** encrypts the LVM physical volume and all the logical volumes that it contains. If you do not select **Encrypt my data**, you can encrypt the logical volume by selecting the **Encrypt** check box on the Manual Partitioning screen or encrypt the physical volume by selecting the **Encrypt** check box on the Configure Volume screen.

For btrfs, encryption can only be applied to the block device that contains the file system, including its subvolumes. For example, enabling encryption for the `/home` subvolume of a btrfs root file system implicitly enables encryption for the root file system itself. You can only select the **Encrypt** check box on the Configure Volume screen. As btrfs does not support encryption at the file-system level, you cannot select the **Encrypt** check box on the Manual Partitioning screen for a btrfs file system.

Do not select the **Encrypt** check box or a **BTRFS**, **LVM**, or **LVM Thin Provisioning** device type for `/boot`. The `/boot` file system must be configured on a standard partition and should be of type `ext4` or `XFS`.

## Network Installation

Attempting to perform a network installation without configuring a network interface to use DHCP to obtain its IP settings or with static IP settings results in the error `Error in Installation Source`.

For example, if you use a feature such as a remote console or Lights-out management to access a boot ISO, the network configuration of the embedded server manager might not be available when you select the installation location. The workaround is to use the graphical installer to configure the network settings manually before configuring the installation location. (Bug ID 19047736)

## Installation on an iSCSI Disk

When installing on an iSCSI disk, add either `ip=ibft` or `rd.iscsi.ibft=1` to the boot command line and specify at least one MBR or GPT-formatted disk as an installation target. Otherwise, the installation fails with the error message `No valid boot loader target device found`. (Bug ID 22076589)

## Installation on an HP 3PAR TPVV

If you have not applied a Thin Persistence license to an HP 3PAR storage array, installation fails to create a file system on a thin provisioned virtual volume (TPVV). This license is required to support the low-level SCSI UNMAP command for storage reclamation. If you do not have a suitable license, the workaround is to use a fully provisioned virtual volume (FPVV) instead of a TPVV. (Bug ID 22140852)

## Unable to boot after installation on systems using the Broadcom MegaRAID 9460 controller

Some systems having Intel® Xeon® E3 v5, Intel® Xeon® Platinum 8100, Intel® Xeon® Gold 6100, Intel® Xeon® Gold 5100, Intel® Xeon® Silver 4100 and Intel® Xeon® Bronze 3100 families of processors (formerly known as Skylake) may use the updated Broadcom MegaRAID 9460 RAID controller that depends on the MegaRAID\_SAS v7.x driver. This driver is not available on the installation media for this release. If Oracle Linux 7 is installed on a RAID volume attached to this controller the system is not able to boot.

To install the correct driver modules for this hardware, you must download and prepare a Driver Update Disk. This Driver update disk contains the following updated driver modules:

- megaraid\_sas 07.701.17.00-rc1
- mpt3sas 15.100.00.00
- smartpqi 1.0.4-100

Updated modules are provided for both UEK and RHCK as required.

You can download the Driver Update Disk from the Oracle Software Delivery Cloud at <https://edelivery.oracle.com/>. Search for 'Oracle Linux 7.4' and select the software to add it to your basket. Click on the 'Selected Software' basket and click Continue. Accept the Oracle Standard Terms and Restrictions and click Continue. You can select the Driver Update Disk files that are part of this media pack to download:

- V952636-01.zip Readme for Driver Update Disk

- V952635-01.iso Driver Update Disk for Oracle Linux 7 x86\_64

Instructions for preparation and installation are covered in [Oracle Linux 7: Installation Guide](#) at: <https://docs.oracle.com/en/operating-systems/oracle-linux/7/install/ol7-install-dud.html>

(Bug ID 26426929)

## Upgrade Issues

You might encounter the following issues when upgrading from Oracle Linux 6 (`_latest`) to Oracle Linux 7 Update 4.

### cgconfig and cgroup packages must be restored separately

The `libcgroup` package in Oracle Linux 7 does not include the `cgconfig` and `cgroup` control group services. To restore these services on an upgraded system, install the `libcgroup-tools` package. (Bug ID 19177606)

### Postupgrade fails if web proxy is required

The postupgrade scripts fail if a proxy is required to access Oracle Linux yum server. (Bug ID 19169163)

### Using old version of yum causes dependency errors

The `redhat-upgrade-tool-cli` utility requires that you install version 3.2.29-43.0.1 or later of the `yum` package on the Oracle Linux 6 system that you want to upgrade. If you use an earlier version of the `yum` package, the upgrade tool fails with dependency errors. (Bug ID 18648783)

## Package Conflicts

The following are known package conflicts for packages distributed by Oracle for Oracle Linux 7 through ULN or the Oracle Linux yum server.

### rdma-core and infiniband-diags when installing oracle-ofed-release

There is a conflict between the `rdma-core` and `infiniband-diags` packages.

The conflict occurs when installing the `oracle-ofed-release` package, as RPM detects the conflict and attempts to install `rdma-core-*` packages instead of RDMA, which results in additional errors related to dependencies.

To avoid the conflict, use the `--exclude=rdma-core*` yum option when performing an install or upgrade of the OFED packages and install the `yum-plugin-priorities` package from the `ol7_optional_latest` yum repository on the Oracle Linux yum server. See [Oracle-Supported OFED Packages](#) for more information. (Bug ID 26309256)

### dovecot-devel.i686 and dovecot-devel.x86\_64

The `dovecot-devel.i686` and `dovecot-devel.x86_64` packages in the `ol7_x86_64_optional_latest` ULN channel conflict. Attempting to install both packages results in a transaction check error:

Transaction check error:

```
file /usr/include/dovecot/config.h conflicts between attempted installs of
dovecot-devel-1:2.2.10-7.el7.i686 and dovecot-devel-1:2.2.10-7.el7.x86_64
```

There are bitsize differences between the identified file. You may only install one of these packages on the same system at once. (Bug ID 25057633)

## ipa-server-dns.x86\_64 and freeipa-server-dns

There is a conflict between the `ipa-server-dns.x86_64` package and the `freeipa-server-dns` package in the `ol7_x86_64_latest` ULN channel. The `.x86_64` version of the `ipa-server-dns` package has been superseded by a `.noarch` package for Oracle Linux 7 update 3.

To avoid the conflict you should exclude the `ipa-server-dns.*.x86_64` package in your Yum configuration. See [Oracle Linux 7: Administrator's Guide](#) for more information on how to exclude packages. (Bug ID 25054687)

## ipa-admintools.x86\_64 and freeipa-admintools

There is a conflict between the `ipa-admintools.x86_64` package and the `freeipa-admintools` package in the `ol7_x86_64_latest` ULN channel. The `.x86_64` version of the `ipa-admintools` package has been superseded by a `.noarch` package for Oracle Linux 7 update 3.

To avoid the conflict you should exclude the `ipa-admintools.*.x86_64` package in your Yum configuration. See [Oracle Linux 7: Administrator's Guide](#) for more information on how to exclude packages. (Bug ID 25054687)

## rear.x86\_64 and rear.noarch

The `.noarch` version of the `rear` package in the `ol7_x86_64_latest` ULN channel has been superseded by a `.x86_64` package for Oracle Linux 7 update 3.

To avoid the conflict you should exclude the `rear.*.noarch` package in your Yum configuration. See [Oracle Linux 7: Administrator's Guide](#) for more information on how to exclude packages. (Bug ID 25054687)

## PackageKit.i686 and PackageKit.x86\_64

The `PackageKit.i686` package from the `ol7_x86_64_optional_latest` ULN channel conflicts with the `PackageKit.x86_64` package in the `ol7_x86_64_u3_base` channel. Attempting to install both packages results in a transaction check error:

Transaction check error:

```
file /usr/lib/python2.7/site-packages/packagekit/__init__.pyc from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
```

```
file /usr/lib/python2.7/site-packages/packagekit/__init__.pyo from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
```

```
file /usr/lib/python2.7/site-packages/packagekit/backend.pyc from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
```

```
file /usr/lib/python2.7/site-packages/packagekit/backend.pyo from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
```

```
file /usr/lib/python2.7/site-packages/packagekit/enums.pyc from install of
PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
```

```

PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/enums.pyo from install of
PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/filter.pyc from install of
PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/filter.pyo from install of
PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/misc.pyc from install of
PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/misc.pyo from install of
PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/package.pyc from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/package.pyo from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/progress.pyc from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64
  file /usr/lib/python2.7/site-packages/packagekit/progress.pyo from install
of PackageKit-1.0.7-6.0.1.el7.i686 conflicts with file from package
PackageKit-1.0.7-6.0.1.el7.x86_64

```

You may only install one of these packages on the same system at once. You should exclude the PackageKit.i686 package in your Yum configuration. See [Oracle Linux 7: Administrator's Guide](#) for more information on how to exclude packages. (Bug ID 24963661)

## sssd-common.i686 and sssd-common.x86\_64

The sssd-common.i686 package conflicts with the sssd-common.x86\_64 package in the ol7\_x86\_64\_optional\_base ULN channel. Attempting to install both packages results in a transaction check error:

```

Transaction check error:
  file /usr/share/systemtap/tapset/sssdp.stp conflicts between attempted
installs of sssd-common-1.14.0-14.el7.i686 and
sssd-common-1.14.0-14.el7.x86_64

```

You may only install one of these packages on the same system at once. You should exclude the sssd-common.i686 package in your Yum configuration. See [Oracle Linux 7: Administrator's Guide](#) for more information on how to exclude packages. (Bug ID 24963661)

## File System Related Bugs

The following file systems issues are related to Oracle Linux 7 Update 4.

**AutoFS: AMD map browsable\_dirs option does not work unless it is set in the [amd] section of autofs.conf**

The autofs package, used to automatically mount file systems as they are required, includes some support for the browsable\_dirs option when using an AMD format mount map, however

this option only works if you manually set it in the [ amd ] section of the `autofs.conf` configuration file. See `autofs.conf(5)` and `/usr/share/doc/autofs-5.0.7/README.amd-maps` for more information. (Bug ID 26363401)

## btrfs: Send operation causes soft lockup on large deduped file

Using `btrfs send` on a large deduped file results in a soft lockup or out-of-memory issue. This problem occurs because the `btrfs send` operation cannot handle a large deduped file containing file extents that are all pointing to one extent, as these types of file structures create tremendous pressure for the `btrfs send` operation.

To prevent this issue from occurring, do not use `btrfs send` on systems with less than 4 GB of memory. (Bug ID 25306023)

## btrfs, ext4 and xfs: Kernel panic when freeze and unfreeze operations are performed in multiple threads

Freeze and unfreeze operations performed across multiple threads on any supported file system can cause the system to hang and the kernel to panic. This issue is the result of a race condition that occurs when the unfreeze operation is triggered before it is actually frozen. The resulting unlock operation attempts a write operation on a non-existent lock resulting in the kernel panic. (Bug ID 25321899)

## btrfs: qgroup reserve space leaks

Several bugs surrounding the way in which quota groups (qgroups) reserve space result in leaks. This includes an issue where leaks are caused by rewriting to dirty ranges, resulting in a "pwrite64: Disk quota exceeded" error. (Bug ID 22483655)

## btrfs: Incorrect exclusive reference count after cloning file between subvolumes

The count for exclusive references is incorrect after cloning a file between two subvolumes. This issue is related to quota groups and the way in which some code is implemented. (Bug ID 22456419)

## btrfs: Kernel oops when unmounting during a quota rescan or disable

Operations that trigger a quota rescan or to disable the quota on a mounted file system cause a kernel oops message when attempting to unmount the file system. This can cause the system to hang. (Bug ID 22377928)

## btrfs: Kernel oops when removing shared extents using qgroup accounting

The removal of shared extents where quota group (qgroup) accounting is used can result in a kernel oops message. This relates to an issue where inaccurate results are obtained during a back reference walk due to missing records when adding delayed references. (Bug ID 21554517)

## ext4: System hangs on unmount after an append to a file with negative `i_size`

While it is invalid for a file system to load an inode with a negative `i_size`, it is possible to create a file like with a negative `i_size` and append to it. However, doing so causes an integer overflow in the routines underlying writeback, which results in the kernel locking up. (Bug ID 25565527)

## ext4: hang occurs during dynamic expansion of inode size

A hang occurs with the `ext4` file system during the dynamic expansion of inode size when using the inode's `i_extra_size` field. (Bug ID 25718971)

## ext4: System hang when processing corrupted orphaned inode list

If the orphaned inode list is corrupted the inode may be processed repeatedly resulting in a system hang. For example, if the orphaned inode list contains a reference to the bootloader inode, `ext4_iget()` returns a bad inode resulting in the processing loop that can hang the system. (Bug ID 24433290)

## xfs: Oracle Linux 7 Update 4 is incompatible with UEK R3 where file systems are formatted using XFS with `d_type` enabled

All `xfs` file systems that are created with the Oracle Linux 7 Update 4 installer have `d_type` support enabled automatically and are formatted with the `fotype=1` option. The UEK R3 kernel is incompatible with this option and does not boot on systems installed with the Oracle Linux 7 Update 4 installer, where default file system formatting is selected because the UEK R3 kernel cannot mount any `xfs` file system that is created using `fotype=1`.

Upgrades are unaffected, as previous updates of Oracle Linux 7 formatted disks use the `fotype=0` option. However, when using UEK R3 on Oracle Linux 7 Update 4 and later, be aware that you must explicitly set the `fotype=0` option when formatting the disk with XFS.

When performing a kickstart installation, if you intend to use UEK R3 on the system, you can manually specify alternate file system options for formatting. If you want to continue to use XFS with UEK R3, you must explicitly set the formatting option to `fotype=0` in your kickstart configuration. (Bug ID 26176688)

## xfs: Directory readahead completions can hang the system after unmount

Directory readahead can hang the system if the file system is unmounted suddenly after a mount. If a directory readahead is delayed for long enough, buffer I/O completion might occur after the unmount has completed. The asynchronous nature of directory readahead I/O means that when the readahead I/O completion occurs, core data structures could have been freed, causing the completion to run into invalid memory accesses., which can result in a kernel panic and system hang. (Bug ID 25550712)

## xfs: System hangs on unmount after a buffered append to a file with negative `i_size`

While it is invalid for a file system to load an inode with a negative `i_size`, it is possible to create a file like this. In the case where a buffer appends to the file, an integer overflow in the routines underlying writeback result in the kernel locking up. Note that a direct append does not cause this behavior. (Bug ID 25565490)

## xfs, ext4: IO error during DIO/AIO write results in disk content corruption

Disk content is corrupted when a Direct IO (DIO) or Asynchronous IO (AIO) write to an unwritten extent fails due to an IO error. (Bug ID 24393811)

## Ceph: SSL certificate verification error when accessing the Ceph Object Gateway in SSL mode

If you configured your Ceph Object Gateway service to enable SSL and you opted to use a self-signed certificate, you may encounter SSL certificate verification errors when you attempt to access the service in SSL mode.

The example Python scripts provided in the [Ceph Storage for Oracle Linux Release 2.0 Release Notes](#) and used to test the Ceph Object Gateway service, require Python libraries that have been updated in Oracle Linux 7 update 4. Libraries such as `urllib2`, a dependency for `python-boto`, have been updated to include much stricter SSL validation and verification, and may return an SSL certificate verification error when connecting over HTTPS to a service that uses a self-signed certificate.

If you choose to use a self-signed certificate, you can copy the CA certificate to the client system's certificate bundle. For example:

```
cat custom.crt >> /etc/pki/tls/certs/ca-bundle.crt
```

Alternately, use the program's environment to specify the path to additional trusted CA certificates in PEM format. The environment variables `SSL_CERT_FILE` and `SSL_CERT_DIR` can be used to specify additional trusted CA certificates. For example:

```
SSL_CERT_FILE=/root/ceph/custom.pem python script.py
```

(Bug ID 26451186)

## Automatic Bug Reporting Tool

The automated reporting daemons and features provided by the Red Hat Automatic Bug Reporting Tool (ABRT) are not supported with Oracle Linux

ABRT packages and associated files, such as `libreport`, are included in the distribution to satisfy package dependencies and can be used to generate local bug reports but the features to automatically upload these reports are not supported. For technical assistance, contact Oracle Support by using the My Oracle Support portal or by telephone.

## Auto-completion of Commands in the bash Shell

Pressing the Tab key to complete commands automatically in the `bash` shell works for some commands such as `ls` but not for other commands such as `export`. You can use the following workaround to enable auto completion for all commands:

1. Remove the `bash-completion` package:

```
sudo yum remove bash-completion
```

2. Run the `complete -r` command in the shell. To make this command persistent, you could put it in `$HOME/.bashrc`.

(Bug ID 19248362)

## crashkernel=auto setting on UEK R3

If you enable the `crashkernel=auto` kernel parameter for UEK R3 to simplify Kdump configuration, both `dmesg` output and `/proc/cmdline` show `crashkernel=NNMM@0M`. This is the expected behavior for the implementation, where `@0M` implies the `auto` setting. The `crashkernel=auto` parameter is not supported for Xen. (Bug ID 17616874)

## grubby sets incorrect saved entry

If `grubby` is used to remove a kernel menu entry from the GRUB 2 configuration, the value of the default entry in `/etc/grub2/grub.cfg` is incorrect.

The workaround for this issue is to set the value of `GRUB_DEFAULT` in `/etc/default/grub` to the correct entry and then use `grub2-mkconfig` to regenerate the `/etc/grub2/grub.cfg` file. Or, use the `yum` command to remove the kernel packages. (Bug ID 19192278)

## grubby fatal error upgrading Kernel when `/boot` is on a btrfs subvolume

If `/boot` is hosted on a btrfs subvolume, GRUB 2 is unable to correctly process the `initramfs` and `vmlinuz` pathnames. This problem occurs when you update or install a new kernel and `grubby` attempts to update the GRUB 2 configuration. In the case where you are running a fresh installation of Oracle Linux 7 Update 4 and you upgrade the RHCK or UEK kernel, the following error is displayed:

```
grubby fatal error: unable to find a suitable template
```

When the system is rebooted, after the kernel update, the system boots to the old kernel.

Similarly, when upgrading from Oracle Linux 7 Update 3 to Oracle Linux 7 Update 4, if the `/boot` directory is hosted on a btrfs subvolume, the system boots to the old Oracle Linux 7 Update 3 kernel after the upgrade is complete.

The workaround to this problem is to use `grub2-mkconfig` to recreate `/etc/grub2/grub.cfg` immediately after the kernel has been installed or upgraded, as shown in this example:

```
sudo grub2-mkconfig -o /boot/grub2/grub.cfg
```

Obtain a listing of the kernel menu entries in the generated configuration:

```
grep -P "submenu|^menuentry" /boot/grub2/grub.cfg | cut -d '"' -f2
```

From the listing, select the kernel entry that you wish to run as the default kernel and set this entry as the default using the following command, substituting *menu entry title* with the title of the kernel entry that you identified in the listing:

```
sudo grub2-set-default "menu entry title"
```

You can use the `grub2-editenv list` command to check that the `saved_entry` has been updated with the selected kernel menu title.

Reboot and use `uname -a` to check that the correct kernel is running when the system is rebooted.

(Bug ID 22750169)

## Hebrew LaTeX fonts

Installing the `tex-fonts-hebrew` package fails unless you first install all `texlive*` packages.  
(Bug ID 19059949)

## InfiniBand Issues

The following are issues that you might encounter when using InfiniBand devices.

### ifup-ib: line 357: /sys/class/net/ib0/acl\_enabled: Permission denied error

Running `ifup ib-interface` or `service network restart` on an Oracle Linux 7 Update 4 system reports the following error:

```
/etc/sysconfig/network-scripts/ifup-ib: line 357: /sys/class/net/ib0/acl_enabled:
Permission denied
```

This error is reported, even though the InfiniBand interface is brought up successfully.

The workaround for this issue is to change from using the older configuration method, where you manipulate `sysfs` files to the newer `ibac1` tools that are provided. (Bug ID 26197105)

### Kdump might fail due to an incorrect InfiniBand Adapter M3 Firmware version

Kdump might fail on Oracle Linux 7 Update 4 if the Oracle Dual Port QDR InfiniBand Adapter M3 Firmware version 2.31.5350 is installed.

To prevent this issue from occurring, update the Oracle Dual Port QDR InfiniBand Adapter M3 Firmware version to at least 2.31.5350. (Bug ID 26351183)

### Changing the IPoIB mode of an InfiniBand interface

The IPoIB driver supports the use of either connected mode or datagram mode with an interface, where datagram mode is the default mode. Changing the mode of an InfiniBand interface by echoing either `connected` or `datagram` to `/sys/class/net/ibN/mode` is not supported for UEK R3. It is also not possible to change the mode of an InfiniBand interface while it is enabled if you are running UEK R3.

To change the IPoIB mode of an InfiniBand interface on a UEK R3 system:

1. Edit the `/etc/sysconfig/network-scripts/ifcfg-ibN` configuration file, where *N* is the number of the interface:
  - To configure connected mode, specify `CONNECTED_MODE=yes` in the file.
  - To configure datagram mode, either specify `CONNECTED_MODE=no` in the file or do not specify this setting at all (datagram mode is enabled by default).

**Note**

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

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

```
sudo ifdown ibN
sudo ifup ibN
```

**Note**

This issue is resolved in UEK R4.

(Bug ID 17479833)

## Disabling an InfiniBand CA port generates warnings

You might see the following warning messages if you use the `ibportstate disable` command to disable an InfiniBand CA or router 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)

## Intel QuickAssist Acceleration Technology

UEK R3 does not support the QAT driver that allows cryptographic capabilities to be offloaded to QuickAssist hardware.

## Database installation and operation fails if `RemoveIPC=yes` is configured for `systemd`

If `RemoveIPC=yes` is configured for `systemd`, interprocess communication (IPC) is terminated for a non-system user's processes when that user logs out. This setting, which is intended for laptops, can cause software problems on server systems. For example, if the user is a database software owner such as `oracle` for Oracle Database, this configuration can cause database installation to fail or database services to crash.

By default, Oracle Linux 7 Update 4 configures `RemoveIPC=no` in `/etc/systemd/logind.conf` to prevent systemd from terminating IPC. However, if you have touched this file before updating your system to Oracle Linux 7 Update 4, the update installs the new version of the file as `/etc/systemd/logind.conf.rpmnew` and does not set `RemoveIPC=no` in `/etc/systemd/logind.conf`. To avoid database crashes, set `RemoveIPC=no` in `/etc/systemd/logind.conf` and run `systemctl reboot` to reboot the system. (Bug ID 22224874)

## Oracle ASM fails to initialize with SELinux in Enforcing mode

The `/etc/init.d/oracleasm` script fails if SELinux is in `Enforcing` mode. This interface is deprecated. Instead, use the global `oracleasm` userspace tool installed in `/usr/sbin/oracleasm` to run any `oracleasm` operations while SELinux is enabled.

(Bug ID 18513404)

## Multipath messages related to `zram` on UEK R3

Running the `multipath -ll` command under UEK R3 produces messages such as the following:

```
zram0: No fc_host device for 'host'  
zram0: No fc_host device for 'host'  
zram0: No fc_remote_port device for 'rport--1:-1-0'
```

You can ignore these message as there is no effect on multipath functionality. You can prevent the messages from occurring by blocklisting the `zram` device in `/etc/multipath.conf`.

Note that this issue is fixed for RHCK and UEK R4 u2 an later, as `zram` support is compiled into separate kernel modules that can be loaded as needed. The warning messages reappear if the `zram` kernel module is loaded for either of these kernels.

(Bug ID 20300644)

## Unable to create Oracle Linux 7 LXC containers on NFS

The creation of Oracle Linux 7 containers fail when the root file system (`/container`) is hosted on an NFS share. RPM fails to set capabilities while attempting to install some packages. For instance, when attempting to create an Oracle Linux 7 Update 4 container, the installation fails while installing the `iputils` package:

```
Error unpacking rpm package iputils-20121221-7.el7.x86_64  
error: unpacking of archive failed on file /usr/bin/ping: cpio: cap_set_file  
error: iputils-20121221-7.el7.x86_64: install failed
```

Similar issues are seen when attempting to install the `initscripts` and `systemd` packages while creating an Oracle Linux 7 Update 3 container.

This issue occurs on both NFSv3 and NFSv4.

Oracle Linux 6 containers are unaffected.

(Bug ID 25024258)

## Oracle Linux 7 guests on Oracle VM and Xen

Oracle Linux 7 guests are supported for both hardware virtualization (HVM) and hardware virtualization with paravirtual drivers (PVHVM) on Oracle VM 3. Oracle Linux 7 guests in a paravirtualized domain (PVM) on Oracle VM or other Xen-based hypervisors are not supported.

Oracle Linux 7 guests of any type are not supported on Oracle VM 2. (Bug IDs 18712168, 18667813, 18266964)

## Hyper-V related services fail to start on Oracle Linux 7 Update 4 guest with UEK R4 kernel under Windows Hyper-V Server

If you are running UEK R4 or an earlier release on Oracle Linux 7, the `hypervkvpd` and `hypervvssd` services fail to start if the Hyper-V packages are at version `0-0.29.20160216git.el7` or later.

To avoid this issue, upgrade to UEK R4U5 or later.

(Bug ID 24745861)

## Per-CPU Allocation Fails when Loading `kvm_intel` module with UEK R3

Per-CPU allocation fails when the `kvm_intel` module is loaded with UEK R3. Messages such as the following are logged:

```
kvm_intel: Could not allocate 48 bytes percpu data
PERCPU: limit reached, disable warning
```

There is no current workaround for this issue. (Bug ID 18459498)

## OpenSSH does not update login records with ssh client's host name

By default, after installing or upgrading to Oracle Linux 7 Update 4, OpenSSH does not update login records such as `/var/run/utmp` and other files with the ssh client's host name. This behavior is expected.

If you want to revert to the previous behavior, where login records are updated with the ssh client's host name, edit the `/etc/ssh/sshd_config` file and uncomment the "UseDNS yes" line. See the `sshd_config(5)` man page for more information. (Bug ID 26286750)

## Geneve network driver support not available in UEK releases

The `ip` and `iproute` commands included with Oracle Linux 7 Update 4 include support for Geneve-capable devices. The module for this driver is included with the RHCK but is not included in UEK R4. The commands to set, add or view Geneve devices are only functional when used with the RHCK. (Bug ID 24652835)

## net\_prio control group not supported on UEK R3

The Network Priority cgroup subsystem (`net_prio`) is not currently supported for use with UEK R3. Attempting to use the module with UEK R3 results in error messages such as the following:

```
modprobe: FATAL: Module netprio_cgroup not found
mount: special device cgroup does not exist.
```

(Bug ID 18966564)

## NetworkManager fails to set the default gateway and route for interfaces configured with DHCP on UEK R3

When running UEK R3 on Oracle Linux 7 update 4, `NetworkManager` fails to set the default gateway and route for network interfaces that are configured with DHCP. This can result in network interfaces not behaving correctly.

The workaround for this issue is to disable `NetworkManager` control over interfaces that are configured for DHCP.

To disable `NetworkManager` for an interface, edit the network script for the interface in `/etc/sysconfig/network-scripts/ifcfg-dev` and add the parameter `NM_CONTROLLED=no`, as shown in the following example:

```
echo "NM_CONTROLLED=no" >> /etc/sysconfig/network-scripts/ifcfg-eno4
```

(Bug ID 26268996)

## NetworkManager unable to add IPv6 addresses to interfaces on UEK R3

After upgrading to Oracle Linux update 4, or when running UEK R3 on Oracle Linux 7 update 4, network interfaces that are configured for IPv6 might not be brought up by `Network Manager`.

Errors similar to the following appear in the system log:

```
<error> platform-linux: do-add-ip6-address[2: fe80::210:e0ff:fe5f:920c]:
failure 22 (Invalid argument)
<error> platform-linux: do-add-ip6-address[5: fd00:1:1:24::456]:
failure 22 (Invalid argument)
```

It is possible to manually add the IPv6 address to the interface using the `ip addr add` command.

This issue is apparent regardless of whether IPv6 is configured statically, assigned dynamically via DHCP, or configured via Stateless Address Autoconfiguration (SLAAC).

The workaround for this issue is to disable `Network Manager` for interfaces where IPv6 must be configured for an interface. To disable `Network Manager` for an interface, edit the network script for the interface in `/etc/sysconfig/network-scripts/ifcfg-dev` and add the parameter `NM_CONTROLLED=no`, as shown in the following example:

```
echo "NM_CONTROLLED=no" >> /etc/sysconfig/network-scripts/ifcfg-eno4
```

(Bug ID 24848072)

## Network connection icon reports incorrect state for interfaces

The network connection icon might report an active network interface as being disconnected. This behavior is seen for the `root` user but not for other users. Command-line utilities such as `ip link` and `ifconfig` report the correct state. (Bug ID 19060089)

## Power button defaults to ACPI Suspend

By default, Oracle Linux 7 in graphical (GUI) console mode treats the hardware power button as equivalent to the ACPI "Sleep" button, which puts the system into low-power sleep mode. This behavior is specific to GNOME desktop environment.

In previous Oracle Linux versions, the hardware power button initiated a system shutdown. To make Oracle Linux 7 do the same, create a file named `/etc/dconf/db/local.d/01-shutdown-button` with the following content:

```
[org/gnome/settings-daemon/plugins/power]
button-power='shutdown'
```

Then run the following command:

```
sudo dconf update
```

You must log out of the desktop environment and log back in for the new setting to take effect. (Bug ID 25597898)

## Cockpit web interface fails to display subscription status

The Cockpit web interface may fail to display subscription status when you click on the Subscriptions menu option. An error is returned:

```
Couldn't get system subscription status. Please ensure subscription-manager
is installed.
```

There is no subscription-manager package available. (Bug ID 26581257)

## 32-bit RDMA packages are installed when upgrading a system with `rdma-core` installed

When upgrading from a system where the `rdma-core.noarch` package is installed, 32-bit versions of the packages and many dependencies are also installed, unnecessarily. This is because the original version of the package is obsoleted and during upgrade the package is replaced with both the `rdma-core.i686` and `rdma-core.x86_64` versions of the package, along with those packages' dependencies.

To work around the issue, run the `yum update` command with the `--exclude=*.i686` option:

```
sudo yum update --exclude=*.i686
```

(Bug ID 28217831)

# 6

## Installation and Availability

### **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

You can download a full Oracle Linux 7 Update 4 installation media image from the Oracle Software Delivery Cloud at <https://edelivery.oracle.com/linux>. You can also obtain the latest Oracle Linux 7 packages from the Unbreakable Linux Network (ULN) and the Oracle Linux yum server.

You can install additional software for Oracle Linux 7 by subscribing to the different channels on ULN or by enabling the required repositories within the Oracle Yum configuration. To explore the channels that are available to you on ULN, login to <https://uln.oracle.com/> and view the Channels option. To view the Oracle Yum repositories available for Oracle Linux 7, visit <https://yum.oracle.com/oracle-linux-7.html>.

### **Note**

The Oracle Linux yum server does not provide equivalent repositories for some channels that are available on ULN. These channels provide non-open source packages.

If you are installing the update on a system on which you have previously installed the Oracle-supported OFED packages, see [Upgrading to Oracle Linux 7 Update 4 with the Oracle-Supported OFED Packages for UEK R4 installed](#) for instructions on how to update these packages during the upgrade.

UEK R4 Update 4 is the default boot kernel for fresh installations of Oracle Linux 7 Update 4. For more information, see [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#).

For systems that are running UEK R3 or UEK R4 and are subscribed to the `o17_x86_64_UEKR3` or `o17_x86_64_UEKR4` channel on ULN, or the `o17_x86_64_UEKR3` or `o17_x86_64_UEKR4` repository on the Oracle Linux yum server, upgrade to the latest UEK release as follows:

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

```
sudo yum update
```

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

2. Reboot the system.

```
sudo systemctl reboot
```

**! Important**

Oracle Linux 7 Update 4 updates many major subsystems. To ensure that your updated systems function correctly, reboot them after updating.

## Oracle-Supported OFED Packages

The following information describes additional steps specific to Oracle Linux 7 update 4 that may be required to install or upgrade the Oracle-supported OFED packages for UEK R4. This section also describes steps to upgrade an Oracle Linux system where the Oracle-supported OFED packages for UEK R4 are already installed.

(Bug ID 19177152)

## Installing or Upgrading Oracle-Supported OFED Packages for UEK R4

For instructions for installing or upgrading OFED packages with UEK R4 u4, see [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#).

**! Important**

Before installing or upgrading the Oracle-supported OFED packages on Oracle Linux 7 update 4, you must perform some preparation steps to ensure that the system is up to date and that any potential package conflicts can be avoided.

If you are installing or upgrading from ULN follow the steps described in [Upgrade using ULN](#) before continuing with the standard installations instructions described in [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#).

If you are installing or upgrading from Oracle Linux yum server, follow the steps described in [Upgrade using the Oracle Linux Yum Server](#) before continuing with the standard installation instructions described in [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#).

Note that when you install the `oracle-ufed-release` or `oracle-ufed-release-guest` package, you must use the `--exclude=rdma-core*` option to avoid a potential package conflict. For example:

```
sudo yum install oracle-ufed-release --exclude=rdma-core*
```

## Upgrading to Oracle Linux 7 Update 4 with the Oracle-Supported OFED Packages for UEK R4 installed

The following procedures describe how to upgrade an Oracle Linux 7 system to Oracle Linux 7 Update 4 on a system where the Oracle-supported OFED packages for UEK R4 are already present.

## Upgrade using ULN

The following procedure describes how to use ULN to perform an upgrade.

### Note

This procedure assumes that you have registered the system with ULN. See [Oracle Linux: Unbreakable Linux Network User's Guide for Oracle Linux 6 and Oracle Linux 7](#).

1. Subscribe the system to the `ol7_x86_64_optional_latest`, `ol7_x86_64_UEKR4_OFED`, `ol7_x86_64_UEKR4`, and `ol7_x86_64_latest` channels on ULN. By default, the `ol7_x86_64_UEKR4` and `ol7_x86_64_latest` channels are enabled when you register an Oracle Linux 7 system with ULN.
2. Edit the `/etc/yum/pluginconf.d/rhnplugin.conf` file and add the following lines to the end of the file:

```
[ol7_x86_64_UEKR4_OFED]
priority=20
```
3. Install the `yum-plugin-priorities` package from the `ol7_x86_64_optional_latest` channel:

```
sudo yum install yum-plugin-priorities
```
4. Apply Oracle Linux 7 Update 4 to the system and use the `--exclude=rdma-core*` option to avoid a potential package conflict:

```
sudo yum update --exclude=rdma-core*
```

Any Oracle-supported OFED packages for UEK R4 that were already present are updated from the `ol7_x86_64_UEKR4_OFED` channel. If you do not have any of the OFED packages installed, you can proceed to install the packages as described in the installation instructions provided in [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#).

## Upgrade using the Oracle Linux Yum Server

The following procedure describes how to use the Oracle Linux yum server to perform an upgrade.

1. If you updated the server from Oracle Linux 6 or the initial release of Oracle Linux 7, ensure that your system is up to date and that you have transitioned to use the modular yum repository configuration by installing the `oraclelinux-release-el7` package and running the `/usr/bin/ol_yum_configure.sh` script.

```
sudo yum install oraclelinux-release-el7
sudo /usr/bin/ol_yum_configure.sh
```
2. Enable the `ol7_optional_latest` and `ol7_UEKR4_OFED` repositories. By default, `ol7_latest` and `ol7_UEKR4` are already enabled, but you should enable these repositories as well if they are not already enabled.

```
sudo yum-config-manager --enable ol7_latest ol7_UEKR4 ol7_optional_latest
ol7_UEKR4_OFED
```
3. Install the `yum-plugin-priorities` package from the `ol7_optional_latest` repository:

```
sudo yum install yum-plugin-priorities
```

4. To apply Oracle Linux 7 Update 4 to the system and use the `--exclude=rdma-core*` option to avoid a potential package conflict:

```
sudo yum update --exclude=rdma-core*
```

Any Oracle-supported OFED packages for UEK R4 that were already present are updated from the `o17_UEKR4_OFED` repository on the Oracle Linux yum server. If you do not have any of the OFED packages installed, you can proceed to install the packages as described in the installation instructions provided in [Unbreakable Enterprise Kernel: Release Notes for Unbreakable Enterprise Kernel Release 4 Update 4 \(4.1.12-94\)](#).

# 7

## Package Changes from the Upstream Release

### **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

The following sections list the changes to source packages from the upstream release.

### Removed Packages

The following packages from the upstream release have been removed:

- anaconda-user-help
- dtc
- kernel-aarch64
- kpatch
- libcxl
- libehca
- libica
- libreport-plugin-rhtsupport
- libreport-rhel
- librtas
- libservicelog
- libvpd
- libzfcphbaapi
- lsvpd
- opal-prd
- openssl-ibmca
- paflib
- powerpc-utils
- powerpc-utils-python
- ppc64-diag
- ppc64-utils
- publican-redhat

- python-rhsm
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-as-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-bn-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-de-DE
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-en-US
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-es-ES
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-fr-FR
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-gu-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-hi-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-it-IT
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-ja-JP
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-kn-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-ko-KR
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-ml-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-mr-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-or-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-pa-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-pt-BR
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-ru-RU
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-ta-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-te-IN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-zh-CN
- Red\_Hat\_Enterprise\_Linux-Release\_Notes-7-zh-TW
- redhat-access-gui
- redhat-access-insights
- redhat-access-plugin-ipa
- redhat-logos
- redhat-support-lib-python
- redhat-support-tool
- SLOF
- s390utils
- servicelog
- subscription-manager
- subscription-manager-migration-data
- virt-who
- yaboot

## Modified Packages

The following source packages from the upstream release have been modified. Note that a single source package might generate multiple binary packages, each of which would also be modified:

- abrt
- abrt-java-connector
- akonadi
- anaconda
- apr-util
- autofs
- basesystem
- btrfs-progs
- clutter
- cockpit
- coreutils
- crash
- dbus
- dhcp
- dracut
- firefox
- flatpak
- fuse
- fwupdate
- gperftools
- grub2
- grubby
- gstreamer
- hivex
- httpd
- initial-setup
- initscripts
- ipa
- iproute
- irqbalance
- iscsi-initiator-utils
- java-1.7.0-openjdk

- kabi-yum-plugins
- kde-settings
- kdeplibs
- kexec-tools
- kmod
- ksc
- libdbi-drivers
- libfprint
- libguestfs
- libreoffice
- libreport
- libreswan
- libxml2
- libxslt
- linux-firmware
- lorax
- lvm2
- mkbootdisk
- mysql-connector-odbc
- net-snmp
- nfs-utils
- nss-pam-ldapd
- ntp
- opa-ff
- opa-fm
- openscap
- open-vm-tools
- oracleasm
- os-prober
- osinfo-db
- PackageKit
- pcs
- perl-DBD-MySQL
- perl-XML-Parser
- plymouth
- policycoreutils
- postfix

- pykickstart
- python
- python-blivet
- qt3
- rear
- redhat-bookmarks
- redhat-indexhtml
- redhat-lsb
- redhat-release-server
- redhat-rpm-config
- redhat-upgrade-dracut
- redhat-upgrade-tool
- redland
- rhn-client-tools (updated to support ULN)
- rhnlib
- rhnsd
- rpmdevtools
- scap-security-guide
- scap-workbench
- shim
- shim-signed
- selinux-policy
- setroubleshoot
- setroubleshoot-plugins
- sos
- system-config-date
- system-config-kickstart
- systemd
- tog-pegasus
- wireshark
- xdg-desktop-portal
- xfsprogs
- xsane
- xulrunner
- yum
- yum-rhn-plugin

## New Packages

The following packages are new for Update 4, relative to Update 3 of Oracle Linux 7:

- `autoconf-archive`
- `brasero`
- `cdrdao`
- `clevis`
- `cloud-utils-growpart`
- `clutter-gst3`
- `cockpit`
- `compat-cheese314`
- `compat-glade315`
- `compat-gnome-desktop314`
- `compat-grilo02`
- `compat-libmediaart0`
- `dbxtool`
- `dconf-editor`
- `flatpak`
- `fwupd`
- `fwupdate`
- `gcab`
- `genwqe-tools`
- `gnome-devel-docs`
- `gsound`
- `gspell`
- `http-parser`
- `intel-cmt-cat`
- `iperf3`
- `jose`
- `keycloak-httpd-client-install`
- `libburn`
- `libfastjson`
- `libgepub`
- `libgexiv2`
- `libgpod`
- `libinput`

- libisofs
- LibRaw
- libusbmuxd
- libXfont2
- libxkbcommon
- llvm-private
- luksmeta
- mallard-rng
- nss-pem
- nvmetcli
- osinfo-db
- osinfo-db-tools
- ovmf
- pcre2
- perl-Perl4-CoreLibs
- python-jsonpatch
- python-mutagen
- python-oauthlib
- python-prettytable
- python-requests-oauthlib
- rdma-core
- rhythmbox
- shotwell
- si-units
- tang
- tpm2-tools
- tpm2-tss
- tss2
- unit-api
- uom-lib
- uom-parent
- uom-se
- uom-systems
- usbguard
- vulkan
- webkitgtk4
- xdg-desktop-portal

- `xdg-desktop-portal-gtk`
- `xorg-x11-drv-libinput`

## Modified Optional Packages

The following optional source packages have been modified. Note that a single source package might generate multiple binary packages, each of which would also be modified:

- `gnu-efi`
- `golang-github-syndtr-gocapability`
- `pesign`
- `publican`
- `sanlock`
- `jetty-artifact-remote-resources`
- `jetty-parent`
- `jetty-toolchain`
- `thunderbird`
- `uom-lib`

## Packages Added by Oracle

The following packages have been added:

- `dtrace-modules`
- `inotify-tools`
- `kernel-uek`
- `libdtrace-ctf`
- `lxc`
- `ocfs2-tools`
- `oracleasm-support`
- `oraclelinux-release`
- `oracle-logos`
- `oracle-database-server-12cR2-preinstall`
- `reflink`
- `uname26`
- `yum-plugin-ulninfo`

# 8

## Removed Modules

 **Warning**

Oracle Linux 7 is now in Extended Support. See [Oracle Linux Extended Support](#) and [Oracle Open Source Support Policies](#) for more information.

Migrate applications and data to Oracle Linux 8 or Oracle Linux 9 as soon as possible.

The following modules have been removed from UEK R4 for Oracle Linux 7 compared with UEK R4 for Oracle Linux 6:

- encrypted-keys
- usbserial
- xhci-pci
- xhci-hcd
- opencores-kbd
- max7359\_keypad
- adp5588-keys
- mcs5000\_ts
- rotary\_encoder
- 3w-xxxx
- scsi\_dh\_rdac
- scsi\_dh\_emc
- scsi\_dh\_alua
- scsi\_dh\_hp\_sw
- hid-magicmouse
- cpufreq\_powersave
- cpufreq\_conservative
- cciss
- rsxx
- dmi-sysfs
- kvaser\_pci
- ems\_pci
- sja1000\_platform
- fealnx

- ns83820
- natsemi
- via-rhine
- via-velocity
- axnet\_cs
- ne2k-pci
- 8390
- pcnet\_cs
- forcedeth
- ath5k
- orinoco\_pci
- wl12xx
- wlcore
- gpio-sch
- configfs
- autofs4
- af\_alg
- seqiv
- pkcs7\_message
- pkcs7\_test\_key
- algif\_hash
- ctr
- algif\_skcipher
- aes-x86\_64
- can-gw