



# Red Hat Enterprise Linux 7

## 7.1 发行注记

Release Notes for Red Hat Enterprise Linux 7.1



# Red Hat Enterprise Linux 7 7.1 发行注记

---

Release Notes for Red Hat Enterprise Linux 7.1

红帽 客户内容服务

## 法律通告

Copyright © 2015–2017 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## 摘要

The Release Notes document the major new features and enhancements implemented in Red Hat Enterprise Linux 7.1 and the known issues in this release. For detailed information regarding the changes between Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7, see the Migration Planning Guide. 知识库 Red Hat 全球支持服务愿藉此机会感谢 Sterling Alexander 和 Michael Everette 在 Red Hat Enterprise Linux 7 测试中做出的突出贡献。

## 目录

前言 .....	5
部分 I. 新功能 .....	6
第 1 章 构架 .....	7
1.1. RED HAT ENTERPRISE LINUX FOR POWER, LITTLE ENDIAN	7
第 2 章 HARDWARE ENABLEMENT .....	8
2.1. INTEL BROADWELL PROCESSOR AND GRAPHICS SUPPORT	8
2.2. SUPPORT FOR TCO WATCHDOG AND I2C (SMBUS) ON INTEL COMMUNICATIONS CHIPSET 89XX SERIES	8
2.3. INTEL PROCESSOR MICROCODE UPDATE	8
2.4. AMD HAWAII GPU SUPPORT	8
2.5. OSA-EXPRESS5S CARDS SUPPORT IN QETHQOAT	8
第 3 章 安装及引导 .....	9
3.1. 安装程序	9
3.2. 引导装载程序	13
第 4 章 存储 .....	14
LVM 缓存	14
使用 libStorageMgmt API 进行存储阵列管理	14
LSI Syncro 支持	14
DIF/DIX 支持	14
改进的 device-mapper-multipath 语法错误检查和输出	15
第 5 章 文件系统 .....	16
支持 Btrfs 文件系统	16
OverlayFS	16
支持平行 NFS	16
第 6 章 内核 .....	17
Ceph 块设备支持	17
共存的 Flash MCL 更新	17
动态内核补丁	17
有一个以上 CPU 的 crashkernel	17
dm-erra 目标	17
Cisco VIC 内核驱动程序	17
hwrng 中改进的熵管理	17
调度程序负载平衡性能提高	17
调度程序中提高的 newidle 平衡	17
HugeTLB 支持单节点中 1GB 大页面分配	17
新的基于 MCS 的锁定机制	18
进程栈大小从 8KB 增加到 16KB	18
在 perf 和 systemtap 中启用 uprobe 和 uretprobe	18
端-到-端数据一致性检查	18
32-位系统中的 DRBG	18
NFSv4.1 支持	18
支持大的 Crashkernel 大小	18
Kdump Supported on Secure Boot Machines	18
Firmware-assisted Crash Dumping	18
Runtime Instrumentation for IBM System z	18
Cisco usNIC Driver	18

Intel Ethernet Server Adapter X710/XL710 Driver Update	19
<b>第 7 章 虚拟化</b>	<b>20</b>
增加 KVM 中 vCPU 上限	20
QEMU、KVM 和 Libvirt API 中的第五代 Intel Core 新指令支持	20
KVM 虚拟机的 USB 3.0 支持	20
压缩 dump-guest-memory 命令	20
Open Virtual Machine Firmware	20
改进 Hyper-V 的网络性能	20
hyperv-daemons 中的 hyperfcopyd	20
libgustfs 中的新功能	20
动态记录器跟踪	21
LPAR Watchdog for IBM System z	21
RDMA-based Migration of Live Guests	21
Removal of Q35 Chipset, PCI Express Bus, and AHCI Bus Emulation	21
<b>第 8 章 集群</b>	<b>22</b>
Corosync 的动态令牌超时	22
Corosync 连接断路器改进	22
Red Hat 高可用性改进	22
<b>第 9 章 编译程序及工具</b>	<b>23</b>
System z 二进制文件中的 Linux 热补支持	23
Performance Application Programming Interface 改进	23
OProfile	23
OpenJDK8	23
使用 sosreport 替换 snap	23
Little-Endian 64-位 PowerPC 的 GDB 支持	23
Tuna 改进	23
crash Moved to Debugging Tools	23
Accurate ethtool Output	24
Concerns Regarding Transactional Synchronization Extensions	24
<b>第 10 章 联网</b>	<b>25</b>
可信的网络连接	25
qlcnic 驱动程序中的 SR-IOV 功能	25
Berkeley 数据包过滤器	25
提高的时钟稳定性	25
libnetfilter_queue 软件包	25
配对改进	25
Intel QuickAssist Technology 驱动程序	25
LinuxPTP timemaster 支持 PTP 和 NTP 之间的故障切换	25
网络 initscripts	25
TCP 延迟的 ACK	25
NetworkManager	25
网络名称空间及 VTI	26
MemberOf 插件的备选配置存储	26
<b>第 11 章 RED HAT ENTERPRISE LINUX ATOMIC HOST</b>	<b>27</b>
NEW FEATURES IN RED HAT ENTERPRISE LINUX ATOMIC HOST 7.1.4	27
NEW FEATURES IN RED HAT ENTERPRISE LINUX ATOMIC HOST 7.1.3	28
NEW FEATURES IN RED HAT ENTERPRISE LINUX ATOMIC HOST 7.1.2	28
<b>第 12 章 LINUX CONTAINERS</b>	<b>30</b>
12.1. LINUX CONTAINERS USING DOCKER TECHNOLOGY	30

12.2. CONTAINER ORCHESTRATION	33
12.3. COCKPIT ENABLEMENT	35
12.4. CONTAINERS USING THE LIBVIRT-LXC TOOLING HAVE BEEN DEPRECATED	35
<b>第 13 章 认证和互操作性</b>	<b>37</b>
Manual Backup and Restore Functionality	37
支持 WinSync 到 Trust 的迁移	37
One-Time Password Authentication	37
为通用互联网文件系统整合 SSSD	37
证书授权管理工具	37
增大访问控制精度	37
特权用户的有限域访问	37
自动数据供应者 (Automatic data provider) 配置	37
AD 和 LDAP sudo 提供者用法	37
32-bit Version of krb5-server and krb5-server-ldap Deprecated	38
SSSD Leverages GPO Policies to Define HBAC	38
Apache Modules for IPA	38
<b>第 14 章 安全性</b>	<b>39</b>
SCAP Security Guide	39
SELinux 策略	39
OpenSSH 中的新功能	39
Libreswan 的新功能	39
TNC 中的新功能	40
GNuTLS 的新功能	40
<b>第 15 章 桌面</b>	<b>41</b>
Mozilla Thunderbird	41
支持四组缓冲 OpenGL 立体视觉效果	41
在线帐户供应商	41
<b>第 16 章 支持和维护</b>	<b>42</b>
ABRT 授权的微报告	42
<b>第 17 章 RED HAT 软件集合</b>	<b>43</b>
<b>第 18 章 RED HAT ENTERPRISE LINUX FOR REAL TIME</b>	<b>44</b>
<b>部分 II. TECHNOLOGY PREVIEWS</b>	<b>45</b>
<b>第 19 章 HARDWARE ENABLEMENT</b>	<b>46</b>
<b>第 20 章 STORAGE</b>	<b>47</b>
<b>第 21 章 FILE SYSTEMS</b>	<b>48</b>
<b>第 22 章 KERNEL</b>	<b>49</b>
<b>第 23 章 VIRTUALIZATION</b>	<b>50</b>
<b>第 24 章 COMPILER AND TOOLS</b>	<b>51</b>
<b>第 25 章 NETWORKING</b>	<b>52</b>
<b>第 26 章 AUTHENTICATION AND INTEROPERABILITY</b>	<b>53</b>
<b>部分 III. 设备驱动程序</b>	<b>54</b>

---

第 27 章 存储驱动程序更新 .....	55
第 28 章 网络驱动程序更新 .....	56
第 29 章 图形驱动程序更新 .....	57
部分 IV. DEPRECATED FUNCTIONALITY .....	58
第 30 章 DEPRECATED FUNCTIONALITY IN RED HAT ENTERPRISE LINUX 7 .....	59
Symbols from libraries linked as dependencies no longer resolved by ld .....	59
Windows guest virtual machine support limited .....	59
部分 V. KNOWN ISSUES .....	60
第 31 章 INSTALLATION AND BOOTING .....	61
第 32 章 STORAGE .....	66
第 33 章 FILE SYSTEMS .....	67
第 34 章 VIRTUALIZATION .....	68
第 35 章 DEPLOYMENT AND TOOLS .....	70
第 36 章 COMPILER AND TOOLS .....	71
第 37 章 NETWORKING .....	72
第 38 章 RED HAT ENTERPRISE LINUX ATOMIC HOST .....	73
第 39 章 LINUX CONTAINERS .....	74
第 40 章 AUTHENTICATION AND INTEROPERABILITY .....	76
第 41 章 ENTITLEMENT .....	79
第 42 章 DESKTOP .....	80
附录 A. 修订历史 .....	81



## 前言

Red Hat Enterprise Linux minor releases are an aggregation of individual enhancement, security, and bug fix errata. The *Red Hat Enterprise Linux 7.1 Release Notes* document the major changes, features, and enhancements introduced in the Red Hat Enterprise Linux 7 operating system and its accompanying applications for this minor release. In addition, the *Red Hat Enterprise Linux 7.1 Release Notes* document the known issues in Red Hat Enterprise Linux 7.1.

For information regarding the Red Hat Enterprise Linux life cycle, refer to <https://access.redhat.com/support/policy/updates/errata/>.

## 部分 I. 新功能

这部分论述了 Red Hat Enterprise Linux 7.1 中引进的新功能及主要改进。

# 第 1 章 构架

Red Hat Enterprise Linux 7.1 is available as a single kit on the following architectures: [1]

- 64-位 AMD
- 64-位 Intel
- IBM POWER7+ and POWER8 (big endian)
- IBM POWER8 (big endian) [2]
- IBM System z [3]

在这个发行本中，Red Hat 为您提供服务器和系统改进，以及 Red Hat 开源体验的总体改进。

## 1.1. RED HAT ENTERPRISE LINUX FOR POWER, LITTLE ENDIAN

Red Hat Enterprise Linux 7.1 introduces little endian support on IBM Power Systems servers using IBM POWER8 processors. Previously in Red Hat Enterprise Linux 7, only the big endian variant was offered for IBM Power Systems. Support for little endian on POWER8-based servers aims to improve portability of applications between 64-bit Intel compatible systems (**x86\_64**) and IBM Power Systems.

- Separate installation media are offered for installing Red Hat Enterprise Linux on IBM Power Systems servers in little endian mode. These media are available from the **Downloads** section of the [Red Hat Customer Portal](#).
- 只有用于 POWER little endian 的 Red Hat Enterprise Linux 支持基于 IBM POWER8 处理器的服务器。
- Currently, Red Hat Enterprise Linux for POWER, little endian is supported only as a KVM guest under **Red Hat Enterprise Virtualization for Power**. Installation on bare metal hardware is currently not supported.
- **GRUB2** 引导装载程序用于安装介质和网络引导。已更新 [《安装指南》](#)，添加了使用 **GRUB2** 为 IBM Power Systems 客户端设置网络引导服务器的操作说明。
- 所有用于 IBM Power Systems 的软件包都可用于 POWER Red Hat Enterprise Linux little endian 和 big endian 变体。
- 为用于 POWER little endian 的 Red Hat Enterprise Linux 构建的软件包使用 **ppc64le** 架构代码 - 例如：`gcc-4.8.3-9.ael7b.ppc64le.rpm`。

---

[1] Note that the Red Hat Enterprise Linux 7.1 installation is supported only on 64-bit hardware. Red Hat Enterprise Linux 7.1 is able to run 32-bit operating systems, including previous versions of Red Hat Enterprise Linux, as virtual machines.

[2] Red Hat Enterprise Linux 7.1 (little endian) 是目前 **Red Hat Enterprise Virtualization for Power** 和 **PowerVM** hypervisor 唯一支持的 KVM 虚拟机。

[3] 注：Red Hat Enterprise Linux 7.1 支持 IBM zEnterprise 196 硬件或者更新的产品；不再支持 IBM System z10 大型机系统，同时将不会引导 Red Hat Enterprise Linux 7.1。

## 第 2 章 HARDWARE ENABLEMENT

### 2.1. INTEL BROADWELL PROCESSOR AND GRAPHICS SUPPORT

Red Hat Enterprise Linux 7.1 added initial support for 5th generation Intel processors (code named **Broadwell**) with the enablement of the Intel Xeon E3-12xx v4 processor family. Support includes the CPUs themselves, integrated graphics in both 2D and 3D mode, and audio support (Broadwell High Definition Legacy Audio, HDMI audio, and DisplayPort audio).

For detailed information regarding CPU enablement in Red Hat Enterprise Linux, please see the Red Hat Knowledgebase article available at <https://access.redhat.com/support/policy/intel>.

The **turbostat** tool (part of the kernel-tools package) has also been updated with support for the new processors.

### 2.2. SUPPORT FOR TCO WATCHDOG AND I2C (SMBUS) ON INTEL COMMUNICATIONS CHIPSET 89XX SERIES

Red Hat Enterprise Linux 7.1 adds support for TCO Watchdog and I2C (SMBUS) on the 89xx series Intel Communications Chipset (formerly Coletto Creek).

### 2.3. INTEL PROCESSOR MICROCODE UPDATE

CPU microcode for Intel processors in the microcode\_ctl package has been updated from version **0x17** to version **0x1c** in Red Hat Enterprise Linux 7.1.

### 2.4. AMD HAWAII GPU SUPPORT

Red Hat Enterprise Linux 7.1 enables support for hardware acceleration on AMD graphics cards using the Hawaii core (AMD Radeon R9 290 and AMD Radeon R9 290X).

### 2.5. OSA-EXPRESS5S CARDS SUPPORT IN QETHQOAT

Support for OSA-Express5s cards has been added to the **qethqoat** tool, part of the s390utils package. This enhancement extends the serviceability of network and card setups for OSA-Express5s cards, and is included as a Technology Preview with Red Hat Enterprise Linux 7.1 on IBM System z.

## 第 3 章 安装及引导

### 3.1. 安装程序

已改进 Red Hat Enterprise Linux 安装程序 **Anaconda**，提高 Red Hat Enterprise Linux 7.1 的安装过程。

界面

- 现在图形安装界面包含一个附加页面，可使用该页面在安装过程中配置 **Kdump** 内核崩溃转储机制。之前，只能在安装后使用 **firstboot** 程序进行配置，但如果没有图形界面就无法使用该程序。现在，您可以在没有图形环境的系统中将配置 **Kdump** 作为安装进程的一部分。您可以使用主安装程序菜单（**安装概述**）进入这个新页面。

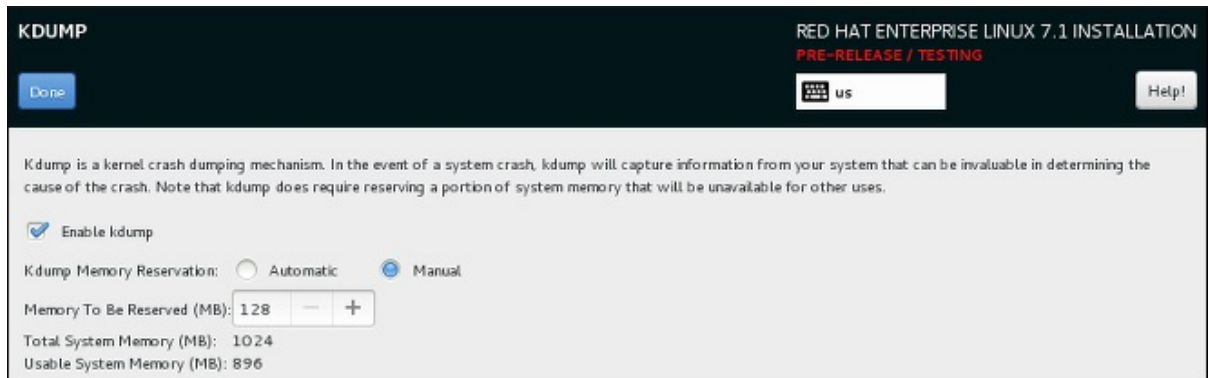


图 3.1. 新 Kdump 页面

- 已重新设计手动分区页面，以改进用户体验。在该页面中将一些控制按钮移动到不同的位置。

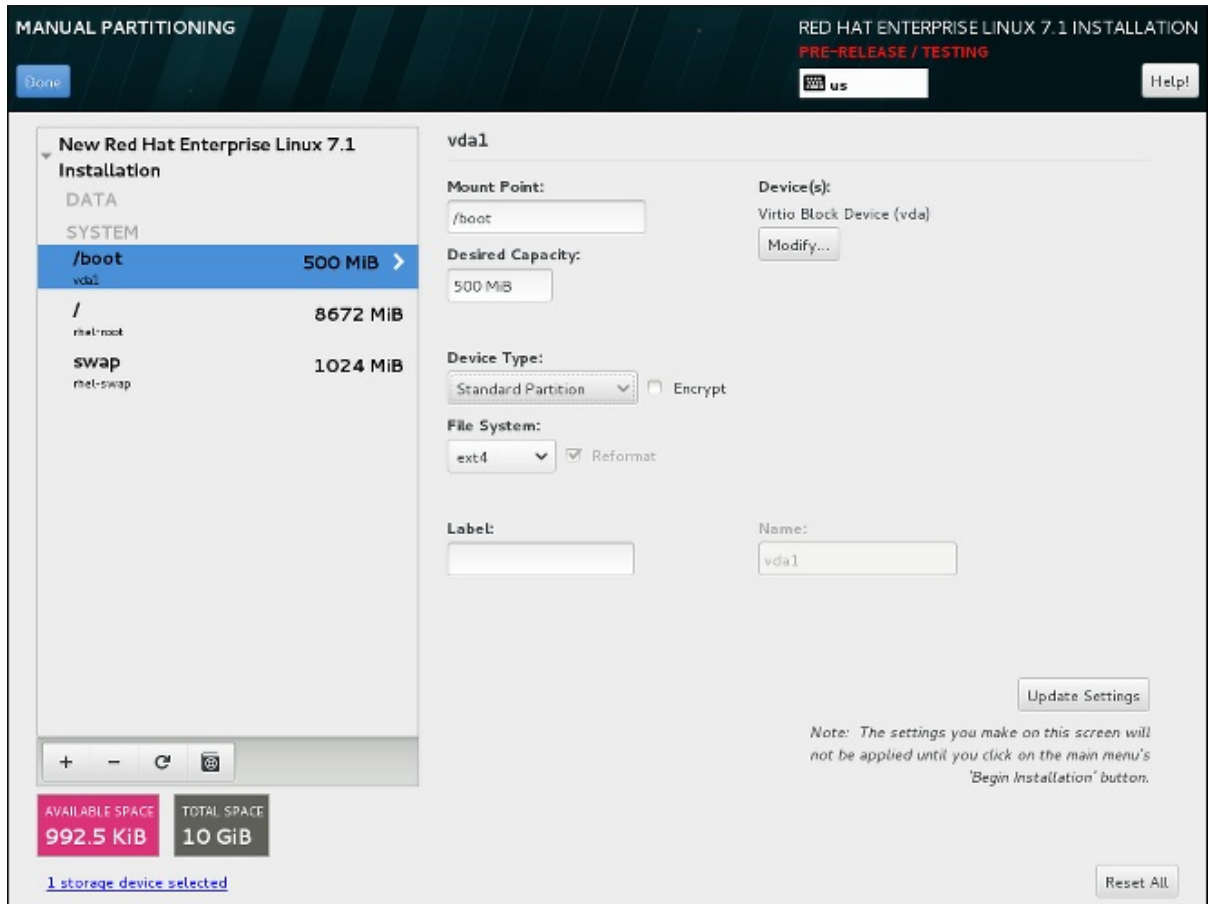


图 3.2. 重新设计的手动分区页面

- You can now configure a network bridge in the **Network & Hostname** screen of the installer. To do so, click the **+** button at the bottom of the interface list, select **Bridge** from the menu, and configure the bridge in the **Editing bridge connection** dialog window which appears afterwards. This dialog is provided by **NetworkManager** and is fully documented in the *Red Hat Enterprise Linux 7.1 Networking Guide*.

在桥接配置中添加了几个新的 Kickstart 选项，详情如下。

- 安装程序不再使用多控制台显示日志。相反，所有日志都将在虚拟控制台 1 (**tty1**) 中的 **tmux** 边框中显示。要在安装过程中查看日志，请按 **Ctrl+Alt+F1** 切换至 **tmux**，然后使用 **Ctrl+b X** 在不同窗口间切换（使用该页面底部显示的具体窗口号替换 X）。

按 **Ctrl+Alt+F6** 切换回图形界面。

- 现在 **Anaconda** 的命令行界面包含完整帮助信息。要查看帮助信息，请在安装了 **anaconda** 软件包的系统中使用 **anaconda -h** 命令。该命令行界面可让您在已安装的系统中运行该安装程序，有利于磁盘映像安装。

## Kickstart 命令及选项

- The **logvol** command has a new option, **--profile=**. This option enables the user to specify the configuration profile name to use with thin logical volumes. If used, the name will also be included in the metadata for the logical volume.

默认情况下，可使用的侧写为在 **/etc/lvm/profile** 目录中定义的 **default** 和 **thin-performance**。详情请查看 **lvm(8)** man page。

- The behavior of the **--size=** and **--percent=** options of the **logvol** command has changed. Previously, the **--percent=** option was used together with **--grow** and **--size=** to specify how much a logical volume should expand after all statically-sized volumes have been created.

Since Red Hat Enterprise Linux 7.1, **--size=** and **--percent=** can not be used on the same **logvol** command.

- The **--autoscreenshot** option of the **autostep** Kickstart command has been fixed, and now correctly saves a screenshot of each screen into the **/tmp/anaconda-screenshots** directory upon exiting the screen. After the installation completes, these screenshots are moved into **/root/anaconda-screenshots**.
- 命令 **liveimg** 现在支持 tar 文件以及磁盘映像。tar 归档必须包含安装介质 root 文件系统，且该文件名必须以 **.tar**、**.tbz**、**.tgz**、**.txz**、**.tar.bz2**、**.tar.gz** 或者 **.tar.xz** 结尾。
- Several new options have been added to the **network** command for configuring network bridges:
  - When the **--bridgeslaves=** option is used, the network bridge with device name specified using the **--device=** option will be created and devices defined in the **--bridgeslaves=** option will be added to the bridge. For example:

```
network --device=bridge0 --bridgeslaves=em1
```

- The **--bridgeopts=** option requires an optional comma-separated list of parameters for the bridged interface. Available values are **stp**, **priority**, **forward-delay**, **hello-time**, **max-age**, and **ageing-time**. For information about these parameters, see the **nm-settings(5)** man page.
- **autopart** 命令有一个新选项 **--fstype**。这个选项可让您在 Kickstart 文件中使用自动分区时更改默认的文件系统类型 (**xfs**)。
- Several new features have been added to Kickstart for better container support. These features include:
  - The new **--install** option for the **repo** command saves the provided repository configuration on the installed system in the **/etc/yum.repos.d/** directory. Without using this option, a repository configured in a Kickstart file will only be available during the installation process, not on the installed system.
  - The **--disabled** option for the **bootloader** command prevents the boot loader from being installed.
  - The new **--nocore** option for the **%packages** section of a Kickstart file prevents the system from installing the **@core** package group. This enables installing extremely minimal systems for use with containers.



### 注意

Please note that the described options are useful only when combined with containers. Using these options in a general-purpose installation could result in an unusable system.

## Entropy Gathering for LUKS Encryption

- If you choose to encrypt one or more partitions or logical volumes during the installation (either during an interactive installation or in a Kickstart file), **Anaconda** will attempt to gather 256 bits of entropy (random data) to ensure the encryption is secure. The installation will continue after 256 bits of entropy are gathered or after 10 minutes. The attempt to gather entropy happens at the beginning of the actual installation phase when encrypted partitions or volumes are being created. A dialog window will open in the graphical interface, showing progress and remaining time.

The entropy gathering process can not be skipped or disabled. However, there are several ways to speed the process up:

- If you can access the system during the installation, you can supply additional entropy by pressing random keys on the keyboard and moving the mouse.
- If the system being installed is a virtual machine, you can attach a *virtio-rng* device (a virtual random number generator) as described in the [Red Hat Enterprise Linux 7.1 Virtualization Deployment and Administration Guide](#).

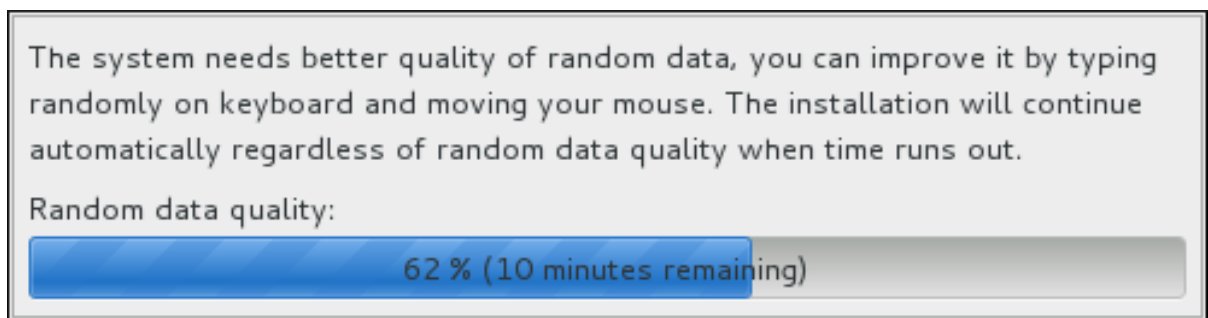


图 3.3. Gathering Entropy for Encryption

#### 图形安装程序中的内置帮助信息

安装程序图形界面中的每个页面以及 **Initial Setup** 程序现在都在其右上角有一个 **帮助** 按钮。点击此按钮使用 **Yelp** 帮助浏览器打开 [《安装指南》](#) 中与当前页面相关的部分。






图 3.4. Anaconda built-in help

### 3.2. 引导装载程序

现在 IBM Power Systems 安装介质使用 **GRUB2** 引导装载程序，而不是之前提供的 **yaboot**。用于 POWER big endian 变体的 Red Hat Enterprise Linux 首选使用 **GRUB2**，但仍可以使用 **yaboot**。新引进的 little endian 变体需要使用 **GRUB2** 引导。

已更新《安装指南》，添加了使用 **GRUB2** 为 IBM Power Systems 设置网络引导服务器的操作说明。

## 第 4 章 存储

### LVM 缓存

As of Red Hat Enterprise Linux 7.1, LVM cache is fully supported. This feature allows users to create logical volumes with a small fast device performing as a cache to larger slower devices. Please refer to the **lvmt(7)** manual page for information on creating cache logical volumes.

注：使用缓存逻辑卷（LV）有以下限制：

- 缓存 LV 必须是顶层设备。不能在精简池 LV、RAID LV 映像以及其他子 LV 类型中使用。
- The cache LV sub-LVs (the origin LV, metadata LV, and data LV) can only be of linear, stripe, or RAID type.
- 生成缓存 LV 后就无法更改其属性。要更改缓存属性，请删除该缓存，并使用所需属性重新生成缓存。

### 使用 **libStorageMgmt** API 进行存储阵列管理

Since Red Hat Enterprise Linux 7.1, storage array management with **libStorageMgmt**, a storage array independent API, is fully supported. The provided API is stable, consistent, and allows developers to programmatically manage different storage arrays and utilize the hardware-accelerated features provided. System administrators can also use **libStorageMgmt** to manually configure storage and to automate storage management tasks with the included command-line interface. Please note that the **Targetd** plug-in is not fully supported and remains a Technology Preview. Supported hardware:

- NetApp 过滤器（数据库 7-模式）
- Nexenta（仅限于 nstor 3.1.x）
- SMI-S，适用于以下销售商：
  - HP 3PAR
    - OS 发行本 3.2.1 或者之后的版本
  - EMC VMAX 和 VNX
    - Solutions Enabler V7.6.2.48 或者之后的版本
    - SMI-S Provider V4.6.2.18 热补套件或者之后的版本
  - HDS VSP Array 非内置供应商
    - Hitachi Command Suite v8.0 或者之后的版本

有关 **libStorageMgmt** 详情请参考 [《存储管理指南》](#) 的相关章节。

### LSI Syncro 支持

Red Hat Enterprise Linux 7.1 包含 **megaraid\_sas** 驱动程序中代码用于启用 LSI Syncro CS 高可用直接附加组件（HA-DAS）适配器。尽管在以前启用的适配器中完全支持 **megaraid\_sas** 驱动程序，但在 Syncro CS 中使用这个驱动程序仍处于技术预览阶段。对这个适配器的支持将直接由 LSI、您的系统集成商或者系统销售商提供。我们鼓励在 Red Hat Enterprise Linux 7.1 中部署 Syncro CS 的用户为 Red Hat 和 LSI 提供反馈意见。有关 LSI Syncro CS 解决方案的详情请参考 <http://www.lsi.com/products/shared-das/pages/default.aspx>。

### DIF/DIX 支持

DIF/DIX 是 SCSI 标准的新增内容，同时在 Red Hat Enterprise Linux 7.1 中也是技术预览。DIF/DIX 将通常使用的 512 字节磁盘块大小从 512 字节增大到 520 字节，添加了数据完整性字段（DIF）。DIF 在进行写入操作时为主机总线适配器（HBA）计算的数据块保存 checksum 值。该存储设备在收到 checksum 后确认，并同时保存该数据和 checksum。相反，当执行读取操作时，可在收到 HBA 后，由该存储设备确认 checksum。

详情请参考《[存储管理指南](#)》中的《启用了 DIF/DIX 的块设备》一节。

### 改进的 `device-mapper-multipath` 语法错误检查和输出

改进了 `device-mapper-multipath` 工具，让 `multipath.conf` 文件更可靠。这样 `multipath.conf` 就可以控制那些无法解析的行，`device-mapper-multipath` 会报告错误并忽略这些行以避免错误解析。

另外在 `multipathd show paths format` 命令中添加了以下通配符表达式：

- %N 和 %n 分别用于主机和目标光纤世界范围节点名称。
- %R 和 %r 分别用于主机和目标光纤世界范围端口名称。

现在更方便将多路径与具体光纤主机、目标及其端口关联，让用户更有效地管理器存储配置。

## 第 5 章 文件系统

### 支持 Btrfs 文件系统

在 Red Hat Enterprise Linux 7.1 中将 **Btrfs** (B-Tree) 文件系统作为技术预览支持。这个文件系统提供高级管理、可靠性及可伸缩功能。它可让用户生成快照，启用压缩和集成的设备管理。

### OverlayFS

The **OverlayFS** file system service allows the user to "overlay" one file system on top of another. Changes are recorded in the upper file system, while the lower file system remains unmodified. This can be useful because it allows multiple users to share a file-system image, for example containers, or when the base image is on read-only media, for example a DVD-ROM.

In Red Hat Enterprise Linux 7.1, OverlayFS is supported as a Technology Preview. There are currently two restrictions:

- It is recommended to use **ext4** as the lower file system; the use of **xfs** and **gfs2** file systems is not supported.
- SELinux is not supported, and to use OverlayFS, it is required to disable enforcing mode.

### 支持平行 NFS

平行 NFS (pNFS) 是 NFS v4.1 标准的一部分， 可让客户端直接且平行访问存储设备。pNFS 架构可提高 NFS 服务器一些常规负载的可延伸性和性能。

pNFS defines three different storage protocols or layouts: files, objects, and blocks. The client supports the files layout, and since Red Hat Enterprise Linux 7.1, the blocks and object layouts are fully supported.

Red Hat 继续致力于与合作伙伴及开源项目合作， 对新的 pNFS 布局类型进行认证， 并在将来提供更多布局类型的全面支持。

有关 pNFS 详情请参考 <http://www.pnfs.com/>。

## 第 6 章 内核

### Ceph 块设备支持

在 Red Hat Enterprise Linux 7.1 内核中添加了 **libceph.ko** 和 **rbd.ko** 模块。这些 RBD 内核模块允许 Linux 主机将 Ceph 块设备视为常规磁盘设备条目，可以挂在到某个目录并使用标准文件系统格式化，比如 **XFS** 或者 **ext4**。

注：Red Hat Enterprise Linux 7.1 目前不支持 CephFs 模块 **ceph.ko**。

### 共存的 Flash MCL 更新

在 IBM System z 架构中的 Red Hat Enterprise Linux 7.1 启用微代码级升级（MCL）。可在不影响 flash 存储介质 I/O 操作的情况下应用这些升级，并通知用户更改了 flash 硬件服务等级。

### 动态内核补丁

Red Hat Enterprise Linux 7.1 introduces **kpatch**, a dynamic "kernel patching utility", as a Technology Preview. The **kpatch** utility allows users to manage a collection of binary kernel patches which can be used to dynamically patch the kernel without rebooting. Note that **kpatch** is supported to run only on AMD64 and Intel 64 architectures.

### 有一个以上 CPU 的 crashkernel

Red Hat Enterprise Linux 7.1 启用在多个 CPU 中引导 **crashkernel** 的功能。这个功能作为技术预览支持。

### dm-era 目标

Red Hat Enterprise Linux 7.1 引进了 dm-era 设备映射器目标作为技术预览。dm-era 可跟踪在用户定义的时间段（即“era”）内写入的块。每个 era 目标事务将当前 era 作为单调增长的 32 位计数器维护。这个目标允许备份软件，以便追踪自上次备份后产生变化的块。它还可启用缓存的部分失效内容，以便在返回经销商快照后恢复缓存一致性。dm-era 目标主要与 dm-cache 目标配对。

### Cisco VIC 内核驱动程序

The Cisco VIC Infiniband kernel driver has been added to Red Hat Enterprise Linux 7.1 as a Technology Preview. This driver allows the use of Remote Directory Memory Access (RDMA)-like semantics on proprietary Cisco architectures.

### hwrng 中改进的熵管理

The paravirtualized hardware RNG (hwrng) support for Linux guests via virtio-rng has been enhanced in Red Hat Enterprise Linux 7.1. Previously, the **rngd** daemon needed to be started inside the guest and directed to the guest kernel's entropy pool. Since Red Hat Enterprise Linux 7.1, the manual step has been removed. A new **khwrngd** thread fetches entropy from the **virtio-rng** device if the guest entropy falls below a specific level. Making this process transparent helps all Red Hat Enterprise Linux guests in utilizing the improved security benefits of having the paravirtualized hardware RNG provided by KVM hosts.

### 调度程序负载均衡性能提高

之前，调度程序负载均衡代码在所有闲置 CPU 之间平衡负载。在 Red Hat Enterprise Linux 7.1 中，只有需要对 CPU 进行负载均衡时才会代表闲置 CPU 执行闲置负载均衡。这个新行为降低了非闲置 CPU 的负载均衡率，进而减少了该调度程序的非必要任务负载，从而提高其性能。

### 调度程序中提高的 newidle 平衡

修改了该调度程序的行为，如果有可运行的任务，则会停止在 **newidle** 平衡代码中搜索任务，从而提高性能。

### HugeTLB 支持单节点中 1GB 大页面分配

Red Hat Enterprise Linux 7.1 添加在运行时支持超大页面分配的功能，可让 1GB **hugetlbfs** 用户指定在运行时为哪个节点分配非一致内存访问（NUMA）节点。

### 新的基于 **MCS** 的锁定机制

Red Hat Enterprise Linux 7.1 引进了新的锁定机制，MCS 锁。这个新锁定机制极大减少大系统中的 **spinlock** 负担，让 **spinlock** 在 Red Hat Enterprise Linux 7.1 中更有效地运行。

### 进程栈大小从 **8KB** 增加到 **16KB**

Since Red Hat Enterprise Linux 7.1, the kernel process stack size has been increased from 8KB to 16KB to help large processes that use stack space.

### 在 **perf** 和 **systemtap** 中启用 **uprobe** 和 **uretprobe**

In Red Hat Enterprise Linux 7.1, the **uprobe** and **uretprobe** features work correctly with the **perf** command and the **systemtap** script.

### 端-到-端数据一致性检查

End-To-End data consistency checking on IBM System z is fully supported in Red Hat Enterprise Linux 7.1. This enhances data integrity and more effectively prevents data corruption as well as data loss.

### 32-位系统中的 **DRBG**

In Red Hat Enterprise Linux 7.1, the deterministic random bit generator (DRBG) has been updated to work on 32-bit systems.

### **NFS**oRDMA Available

As a Technology Preview, the NFSoRDMA service has been enabled for Red Hat Enterprise Linux 7.1. This makes the **svcrdma** module available for users who intend to use Remote Direct Memory Access (RDMA) transport with the Red Hat Enterprise Linux 7 NFS server.

### 支持大的 **Crashkernel** 大小

The **Kdump** kernel crash dumping mechanism on systems with large memory, that is up to the Red Hat Enterprise Linux 7.1 maximum memory supported limit of 6TB, has become fully supported in Red Hat Enterprise Linux 7.1.

### **Kdump** Supported on Secure Boot Machines

With Red Hat Enterprise Linux 7.1, the Kdump crash dumping mechanism is supported on machines with enabled Secure Boot.

### Firmware-assisted Crash Dumping

Red Hat Enterprise Linux 7.1 introduces support for firmware-assisted dump (fadump), which provides an alternative crash dumping tool to kdump. The firmware-assisted feature provides a mechanism to release the reserved dump memory for general use once the crash dump is saved to the disk. This avoids the need to reboot the system after performing the dump, and thus reduces the system downtime. In addition, fadump uses of the kdump infrastructure already present in the user space, and works seamlessly with the existing kdump init scripts.

### Runtime Instrumentation for IBM System z

As a Technology Preview, support for the Runtime Instrumentation feature has been added for Red Hat Enterprise Linux 7.1 on IBM System z. Runtime Instrumentation enables advanced analysis and execution for a number of user-space applications available with the IBM zEnterprise EC12 system.

### Cisco usNIC Driver

Cisco Unified Communication Manager (UCM) servers have an optional feature to provide a Cisco proprietary User Space Network Interface Controller (usNIC), which allows performing Remote Direct Memory Access (RDMA)-like operations for user-space applications. As a Technology Preview, Red Hat Enterprise Linux 7.1 includes the **libusnic\_verbs** driver, which makes it possible to use usNIC devices via standard InfiniBand RDMA programming based on the Verbs API.

## Intel Ethernet Server Adapter X710/XL710 Driver Update

The **i40e** and **i40evf** kernel drivers have been updated to their latest upstream versions. These updated drivers are included as a Technology Preview in Red Hat Enterprise Linux 7.1.

## 第 7 章 虚拟化

### 增加 KVM 中 vCPU 上限

KVM 虚拟机中支持的虚拟 CPU (vCPU) 上限已增至 240。这增加了用户可为虚拟机分配的虚拟处理单元数，因此有可能提高其性能。

### QEMU、KVM 和 Libvirt API 中的第五代 Intel Core 新指令支持

In Red Hat Enterprise Linux 7.1, the support for 5th Generation Intel Core processors has been added to the QEMU hypervisor, the KVM kernel code, and the **libvirt** API. This allows KVM guests to use the following instructions and features: ADCX, ADOX, RDSFEED, PREFETCHW, and supervisor mode access prevention (SMAP).

### KVM 虚拟机的 USB 3.0 支持

Red Hat Enterprise Linux 7.1 通过添加 USB 3.0 主机适配器 (xHCI) 模拟作为技术预览提供改进的 USB 支持。

### 压缩 **dump-guest-memory** 命令

Since Red Hat Enterprise Linux 7.1, the **dump-guest-memory** command supports crash dump compression. This makes it possible for users who cannot use the **virsh dump** command to require less hard disk space for guest crash dumps. In addition, saving a compressed guest crash dump usually takes less time than saving a non-compressed one.

### Open Virtual Machine Firmware

在 Red Hat Enterprise Linux 7.1 中 Open Virtual Machine Firmware (OVMF) 作为技术预览提供。OVMF 是用于 AMD64 和 Intel 64 虚拟机的 UEFI 安全引导环境。

### 改进 Hyper-V 的网络性能

Several new features of the Hyper-V network driver have been introduced to improve network performance. For example, Receive-Side Scaling, Large Send Offload, Scatter/Gather I/O are now supported, and network throughput is increased.

### hyperv-daemons 中的 **hyperfcopyd**

在 hyperv-daemons 软件包中添加了 **hyperfcopyd** 守护进程。**hyperfcopyd** 是在 Hyper-V 2012 R2 主机中为所运行的 Linux 虚拟机使用的文件复制服务功能。它可让主机将文件 (通过 VMBUS) 复制到 Linux 虚拟机中。

### libguestfs 中的新功能

Red Hat Enterprise Linux 7.1 introduces a number of new features in **libguestfs**, a set of tools for accessing and modifying virtual machine disk images. Namely:

- **virt-builder** – a new tool for building virtual machine images. Use **virt-builder** to rapidly and securely create guests and customize them.
- **virt-customize** – a new tool for customizing virtual machine disk images. Use **virt-customize** to install packages, edit configuration files, run scripts, and set passwords.
- **virt-diff** – a new tool for showing differences between the file systems of two virtual machines. Use **virt-diff** to easily discover what files have been changed between snapshots.
- **virt-log** – a new tool for listing log files from guests. The **virt-log** tool supports a variety of guests including Linux traditional, Linux using journal, and Windows event log.
- **virt-v2v** – a new tool for converting guests from a foreign hypervisor to run on KVM, managed by libvirt, OpenStack, oVirt, Red Hat Enterprise Virtualization (RHEV), and several other targets. Currently, **virt-v2v** can convert Red Hat Enterprise Linux and Windows guests running on Xen



and VMware ESX.

### 动态记录器跟踪

Support for flight recorder tracing has been introduced in Red Hat Enterprise Linux 7.1. Flight recorder tracing uses **SystemTap** to automatically capture qemu-kvm data as long as the guest machine is running. This provides an additional avenue for investigating qemu-kvm problems, more flexible than qemu-kvm core dumps.

有关如何配置和使用动态记录器跟踪的详细步骤请参考《虚拟化部署和管理指南》。

### LPAR Watchdog for IBM System z

As a Technology Preview, Red Hat Enterprise Linux 7.1 introduces a new watchdog driver for IBM System z. This enhanced watchdog supports Linux logical partitions (LPAR) as well as Linux guests in the z/VM hypervisor, and provides automatic reboot and automatic dump capabilities if a Linux system becomes unresponsive.

### RDMA-based Migration of Live Guests

The support for Remote Direct Memory Access (RDMA)-based migration has been added to **libvirt**. As a result, it is now possible to use the new **rdma://** migration URI to request migration over RDMA, which allows for significantly shorter live migration of large guests. Note that prior to using RDMA-based migration, RDMA has to be configured and **libvirt** has to be set up to use it.

### Removal of Q35 Chipset, PCI Express Bus, and AHCI Bus Emulation

Red Hat Enterprise Linux 7.1 removes the emulation of the Q35 machine type, required also for supporting the PCI Express (PCIe) bus and the Advanced Host Controller Interface (AHCI) bus in KVM guest virtual machines. These features were previously available on Red Hat Enterprise Linux as Technology Previews. However, they are still being actively developed and might become available in the future as part of Red Hat products.

## 第 8 章 集群

### Corosync 的动态令牌超时

在 **Corosync Cluster Engine** 中添加 **token\_coefficient** 选项。只有指定 **nodelist**，且至少有三个节点时才会使用 **token\_coefficient**。在这种情况下，按以下方法计算令牌超时：

```
[token + (amount of nodes - 2)] * token_coefficient
```

这样就可以在每次添加新节点时，无需手动更改令牌超时即可按比例调整集群。默认值为 650 毫秒，但可将其设定为 0，即完全不使用这个功能。

这个功能可让 **Corosync** 处理动态添加和删除的节点。

### Corosync 连接断路器改进

改进了 **Corosync** 的 **auto\_tie\_breaker** 仲裁功能，提供更灵活的配置，并可以修改连接断路器节点。用户现在可以在对称集群分割时选择一组保留仲裁的节点，或者选择由最小节点 ID 或者最大节点 ID 保留的仲裁。

### Red Hat 高可用性改进

在 Red Hat Enterprise Linux 7.1 发行本中，**Red Hat High Availability Add-On** 支持以下功能。有关这些功能的详情请查看《*High Availability Add-On 参考*》手册。

- 现在 **pcs resource cleanup** 命令可以重置资源状态以及所有资源的 **failcount**。
- 您可以为 **pcs resource move** 命令指定 **lifetime** 参数，给出这个命令限制该资源的时限。
- 您可以使用 **pcs acl** 命令设定本地用户权限，使其有使用访问控制列表（ACL）只读或者读写集群配置的权限。
- 除常规资源选项外，**pcs constraint** 命令现在支持具体限制选项。
- **pcs resource create** 命令支持 **disabled** 参数，表示不会自动启动要创建的资源。
- **pcs cluster quorum unblock** 命令防止集群在建立仲裁时等待所有节点。
- 您可以使用 **pcs resource create** 命令的 **before** 和 **after** 参数配置资源组顺序。
- 您可以使用 tarball 备份集群配置，并在所有节点中，使用 **pcs config** 命令的 **backup** 和 **restore** 选项在所有节点中恢复集群配置。

## 第 9 章 编译程序及工具

### System z 二进制文件中的 Linux 热补支持

GNU 编译器集合 (GCC) 支持为 System z 二进制文件提供多线程代码补丁。使用 "function attribute" 可为热补选择具体功能，使用 **-mhotpatch** 命令行选项启用所有功能的热补。

启用热补对软件大小和性能有负面影响。因此建议为具体功能使用热补，而不是为所有功能都提供热补服务。

System z 二进制文件中的 Linux 热补支持在 Red Hat Enterprise Linux 7.0 为技术预览。在 Red Hat Enterprise Linux 7.1 发行本中现在提供全面支持。

### Performance Application Programming Interface 改进

Red Hat Enterprise Linux 7 包括 **Performance Application Programming Interface (PAPI)**。PAPI 是现代微处理器中硬件性能计数器的跨平台接口规格。这些计数器是一小组暂存器，可计算与具体处理器功能关联信号同时出现的时间。监控这些事件可在不同方面帮助分析和调整应用程序性能。

In Red Hat Enterprise Linux 7.1, PAPI and the related **libpfm** libraries have been enhanced to provide support for IBM POWER8, Applied Micro X-Gene, ARM Cortex A57, and ARM Cortex A53 processors. In addition, the events sets have been updated for Intel Xeon, Intel Xeon v2, and Intel Xeon v3 processors.

### OProfile

**OProfile** is a system-wide profiler for Linux systems. The profiling runs transparently in the background and profile data can be collected at any time. In Red Hat Enterprise Linux 7.1, **OProfile** has been enhanced to provide support for the following processor families: Intel Atom Processor C2XXX, 5th Generation Intel Core Processors, IBM POWER8, AppliedMicro X-Gene, and ARM Cortex A57.

### OpenJDK8

Red Hat Enterprise Linux 7.1 features the java-1.8.0-openjdk packages, which contain the latest version of the Open Java Development Kit, OpenJDK8, that is now fully supported. These packages provide a fully compliant implementation of Java SE 8 and may be used in parallel with the existing java-1.7.0-openjdk packages, which remain available in Red Hat Enterprise Linux 7.1.

Java 8 brings numerous new improvements, such as Lambda expressions, default methods, a new Stream API for collections, JDBC 4.2, hardware AES support, and much more. In addition to these, OpenJDK8 contains numerous other performance updates and bug fixes.

### 使用 **sosreport** 替换 **snap**

从 powerpc-utils 软件包中删除已弃用的 **snap** 工具，将其归纳整合到 **sosreport** 工具中。

### Little-Endian 64-位 PowerPC 的 GDB 支持

Red Hat Enterprise Linux 7.1 在 GNU Debugger (GDB) 中支持 64-位 PowerPC little-endian 架构。

### Tuna 改进

**Tuna** is a tool that can be used to adjust scheduler tunables, such as scheduler policy, RT priority, and CPU affinity. In Red Hat Enterprise Linux 7.1, the **Tuna** GUI has been enhanced to request root authorization when launched, so that the user does not have to run the desktop as root to invoke the **Tuna** GUI. For further information on **Tuna**, see the [Tuna User Guide](#).

### crash Moved to Debugging Tools

With Red Hat Enterprise Linux 7.1, the crash packages are no longer a dependency of the abrt packages. Therefore, crash has been removed from the default installation of Red Hat Enterprise Linux 7 in order to keep the installation minimal. Now, users have to select the **Debugging Tools** option in the Anaconda installer GUI for the crash packages to be installed.

### Accurate ethtool Output

As a Technology Preview, the network-querying capabilities of the **ethtool** utility have been enhanced for Red Hat Enterprise Linux 7.1 on IBM System z. As a result, when using hardware compatible with the improved querying, **ethtool** now provides improved monitoring options, and displays network card settings and values more accurately.

### Concerns Regarding Transactional Synchronization Extensions

Intel has issued erratum [HSW136](#) concerning Transactional Synchronization Extensions (TSX) instructions. Under certain circumstances, software using the Intel TSX instructions may result in unpredictable behavior. TSX instructions may be executed by applications built with the Red Hat Enterprise Linux 7.1 GCC under certain conditions. These include the use of GCC's experimental Transactional Memory support (**-fgnu-tm**) when executed on hardware with TSX instructions enabled. Users of Red Hat Enterprise Linux 7.1 are advised to exercise further caution when experimenting with Transaction Memory at this time, or to disable TSX instructions by applying an appropriate hardware or firmware update.

## 第 10 章 联网

### 可信的网络连接

[Docker Images](#) 引进了可信网络连接功能作为技术预览。可信网络连接可用于现有网络访问控制 (NAC) 解决方案, 比如 TLS、802.1X 或者 IPsec 整合端点态势评估, 即收集端点系统信息 (比如操作系统配置设置, 安装的软件包及其他, 总称为完整性测量)。在允许该端点访问该网络前, 使用可信网络连接, 根据网络访问策略确认这些测量。

### qlcnic 驱动程序中的 SR-IOV 功能

已在 **qlcnic** 中添加单一 Root I/O 虚拟化 (SR-IOV) 支持作为技术预览。对这个功能的支持直接由 QLogic 提供, 同时鼓励用户为 Red Hat 提供反馈意见。仍全面支持 qlcnic 驱动程序中的其他功能。

### Berkeley 数据包过滤器

在 Red Hat Enterprise Linux 7.1 中添加了基于 *流量分类器* 的 Berkeley 数据包过滤器 (BPF)。BPF 是用于数据包套接字的数据包过滤, *安全计算模式* (seccomp) 的沙箱, 以及 Netfilter。BPF 对大多数架构都足够使用, 且有丰富的构建过滤器的句法。

### 提高的时钟稳定性

之前的测试结果表示禁用无缝内核功能可显著提高系统时钟的稳定性。在内核引导选项参数中添加 **nohz=off** 即可禁用内核无缝模式。但在 Red Hat Enterprise Linux 7.1 中应用的最新改进极大提高了系统稳定性, 目前对于大多数用户来说, 使用或者不使用 **nohz=off** 对系统时钟稳定性的影响并不大。这对使用 **PTP** 和 **NTP** 的时间同步应用程序非常有益。

### libnetfilter\_queue 软件包

在 Red Hat Enterprise Linux 7.1 中添加了 libnetfilter\_queue 软件包。**libnetfilter\_queue** 是用户空间存储库, 为使用内核数据包过滤器排队的数据包提供 API。您可以使用它从内核 **nfnlqueue** 子系统中接收排队的数据包、解析数据包、重新编写数据包标头、以及将更改的数据包重新放回队列中。

### 配对改进

The libteam packages have been updated to version **1.15** in Red Hat Enterprise Linux 7.1. It provides a number of bug fixes and enhancements, in particular, **teamd** can now be automatically re-spawned by **systemd**, which increases overall reliability.

### Intel QuickAssist Technology 驱动程序

在 Red Hat Enterprise Linux 7.1 中添加了 Intel QuickAssist Technology (QAT) 驱动程序。QAT 驱动程序可启用 QuickAssist 硬件, 可在系统中添加硬件的卸载加密功能。

### LinuxPTP timemaster 支持 PTP 和 NTP 之间的故障切换

在 Red Hat Enterprise Linux 7.1 中已将 linuxptp 软件包更新至版本 **1.4**。它提供大量 bug 修复和改进, 尤其是使用 **timemaster** 应用程序的 **PTP** 域和 **NTP** 源。需要在网络中有多个 **PTP** 域, 或者返回 **NTP** 时, 可使用 **timemaster** 程序将所有可用时间源于系统时钟同步。

### 网络 initscripts

在 Red Hat Enterprise Linux 7.1 中添加了自定义 VLAN 名称支持。添加了 GRE 通道的 **IPv6** 支持改进; 内部地址现可在重启后保存。

### TCP 延迟的 ACK

Red Hat Enterprise Linux 7.1 的 iproute 软件包添加了可配置 TCP 延迟 ACK 支持。您可以使用 **ip route quickack** 命令启用此功能。

### NetworkManager

NetworkManager has been updated to version **1.0** in Red Hat Enterprise Linux 7.1.

The support for Wi-Fi, Bluetooth, wireless wide area network (WWAN), ADSL, and **team** has been split into separate subpackages to allow for smaller installations.

To support smaller environments, this update introduces an optional built-in Dynamic Host Configuration Protocol (DHCP) client that uses less memory.

A new NetworkManager mode for static networking configurations that starts NetworkManager, configures interfaces and then quits, has been added.

NetworkManager provides better cooperation with non-NetworkManager managed devices, specifically by no longer setting the IFF\_UP flag on these devices. In addition, NetworkManager is aware of connections created outside of itself and is able to save these to be used within NetworkManager if desired.

In Red Hat Enterprise Linux 7.1, NetworkManager assigns a default route for each interface allowed to have one. The metric of each default route is adjusted to select the global default interface, and this metric may be customized to prefer certain interfaces over others. Default routes added by other programs are not modified by NetworkManager.

Improvements have been made to NetworkManager's IPv6 configuration, allowing it to respect IPv6 router advertisement MTUs and keeping manually configured static IPv6 addresses even if automatic configuration fails. In addition, WWAN connections now support IPv6 if the modem and provider support it.

Various improvements to dispatcher scripts have been made, including support for a pre-up and pre-down script.

现在 Red Hat Enterprise Linux 7.1 支持捆绑选项 **lACP\_rate**。已改进 **NetworkManager**，方便在使用辅助接口重新命名主接口时的设备重命名。

A priority setting has been added to the auto-connect function of **NetworkManager**. Now, if more than one eligible candidate is available for auto-connect, **NetworkManager** selects the connection with the highest priority. If all available connections have equal priority values, **NetworkManager** uses the default behavior and selects the last active connection.

This update also introduces numerous improvements to the **nmcli** command-line utility, including the ability to provide passwords when connecting to Wi-Fi or 802.1X networks.

## 网络名称空间及 VTI

在 Red Hat Enterprise Linux 7.1 中添加了对附带网络名称空间的**虚拟通道接口** (VTI) 支持。这样可在数据包压缩或者解压缩时将来自 VTI 的流量分散到不同的名称空间。

## MemberOf 插件的备选配置存储

The configuration of the **MemberOf** plug-in for the Red Hat Directory Server can now be stored in a suffix mapped to a back-end database. This allows the **MemberOf** plug-in configuration to be replicated, which makes it easier for the user to maintain a consistent **MemberOf** plug-in configuration in a replicated environment.

## 第 11 章 RED HAT ENTERPRISE LINUX ATOMIC HOST

Included in the release of Red Hat Enterprise Linux 7.1 is Red Hat Enterprise Linux Atomic Host - a secure, lightweight, and minimal-footprint operating system optimized to run Linux containers. It has been designed to take advantage of the powerful technology available in Red Hat Enterprise Linux 7. Red Hat Enterprise Linux Atomic Host uses SELinux to provide strong safeguards in multi-tenant environments, and provides the ability to perform atomic upgrades and rollbacks, enabling quicker and easier maintenance with less downtime. Red Hat Enterprise Linux Atomic Host uses the same upstream projects delivered via the same RPM packaging as Red Hat Enterprise Linux 7.

Red Hat Enterprise Linux Atomic Host is pre-installed with the following tools to support Linux containers:

- **Docker** - For more information, see [Get Started with Docker Formatted Container Images on Red Hat Systems](#).
- **Kubernetes, flannel, etcd** - For more information, see [Get Started Orchestrating Containers with Kubernetes](#).

Red Hat Enterprise Linux Atomic Host makes use of the following technologies:

- **OSTree and rpm-OSTree** - These projects provide atomic upgrades and rollback capability.
- **systemd** - The powerful new init system for Linux that enables faster boot times and easier orchestration.
- **SELinux** - Enabled by default to provide complete multi-tenant security.

### NEW FEATURES IN RED HAT ENTERPRISE LINUX ATOMIC HOST 7.1.4

- The `iptables-service` package has been added.
- It is now possible to enable automatic "command forwarding" when commands that are not found on Red Hat Enterprise Linux Atomic Host, are seamlessly retried inside the RHEL Atomic Tools container. The feature is disabled by default (it requires a RHEL Atomic Tools pulled on the system). To enable it, uncomment the **export** line in the `/etc/sysconfig/atomic` file so it looks like this:

```
export TOOLSIMG=rhel7/rhel-tools
```

- The **atomic** command:
  - You can now pass three options (**OPT1**, **OPT2**, **OPT3**) to the **LABEL** command in a Dockerfile. Developers can add environment variables to the labels to allow users to pass additional commands using **atomic**. The following is an example from a Dockerfile:

```
LABEL docker run ${OPT1}${IMAGE}
```

This line means that running the following command:

```
atomic run --opt1="-ti" image_name
```

is identical to running

```
docker run -ti image_name
```

- You can now use **\${NAME}** and **\${IMAGE}** anywhere in your label, and **atomic** will substitute it with an image and a name.
- The **\${SUDO\_UID}** and **\${SUDO\_GID}** options are set and can be used in image **LABEL**.
- The **atomic mount** command attempts to mount the file system belonging to a given container/image ID or image to the given directory. Optionally, you can provide a registry and tag to use a specific version of an image.

## NEW FEATURES IN RED HAT ENTERPRISE LINUX ATOMIC HOST 7.1.3

- Enhanced **rpm-OSTee** to provide a unique machine ID for each machine provisioned.
- Support for remote-specific GPG keyring has been added, specifically to associate a particular GPG key with a particular OSTree remote.
- the **atomic** command:
  - **atomic upload** – allows the user to upload a container image to a docker repository or to a Pulp/Crane instance.
  - **atomic version** – displays the "Name Version Release" container label in the following format: **ContainerID;Name-Version-Release;Image/Tag**
  - **atomic verify** – inspects an image to verify that the image layers are based on the latest image layers available. For example, if you have a **MongoDB** application based on **rhel7-1.1.2** and a **rhel7-1.1.3** base image is available, the command will inform you there is a later image.
  - A **dbus** interface has been added to verify and version commands.

## NEW FEATURES IN RED HAT ENTERPRISE LINUX ATOMIC HOST 7.1.2

The **atomic** command-line interface is now available for Red Hat Enterprise Linux 7.1 as well as Red Hat Enterprise Linux Atomic Host. Note that the feature set is different on both systems. Only Red Hat Enterprise Linux Atomic Host includes support for OSTree updates. The **atomic run** command is supported on both platforms.

- **atomic run** allows a container to specify its run-time options via the **RUN** meta-data label. This is used primarily with privileges.
- **atomic install** and **atomic uninstall** allow a container to specify install and uninstall scripts via the **INSTALL** and **UNINSTALL** meta-data labels.
- **atomic** now supports container upgrade and checking for updated images.

The **iscsi-initiator-utils** package has been added to Red Hat Enterprise Linux Atomic Host. This allows the system to mount iSCSI volumes; Kubernetes has gained a storage plugin to set up iSCSI mounts for containers.

You will also find *Integrity Measurement Architecture* (IMA), **audit** and **libwrap** available from **systemd**.





## 重要

Red Hat Enterprise Linux Atomic Host is not managed in the same way as other Red Hat Enterprise Linux 7 variants. Specifically:

- The **Yum** package manager is not used to update the system and install or update software packages. For more information, see [Installing Applications on Red Hat Enterprise Linux Atomic Host](#).
- There are only two directories on the system with write access for storing local system configuration: **/etc/** and **/var/**. The **/usr/** directory is mounted read-only. Other directories are symbolic links to a writable location - for example, the **/home/** directory is a symlink to **/var/home/**. For more information, see [Red Hat Enterprise Linux Atomic Host File System](#).
- The default partitioning dedicates most of available space to containers, using direct Logical Volume Management (LVM) instead of the default loopback.

For more information, see [Getting Started with Red Hat Enterprise Linux Atomic Host](#).

Red Hat Enterprise Linux Atomic Host 7.1.1 provides new versions of **Docker** and **etcd**, and maintenance fixes for the **atomic** command and other components.

## 第 12 章 LINUX CONTAINERS

### 12.1. LINUX CONTAINERS USING DOCKER TECHNOLOGY

#### Red Hat Enterprise Linux Atomic Host 7.1.4 includes the following updates:

The docker packages have been upgraded to upstream version 1.7.1, which contains various improvements over version 1.7, which, in its turn, contains significant changes from version 1.6 included in Red Hat Enterprise Linux Atomic Host 7.1.3. See the following change log for the full list of fixes and features between version 1.6 and 1.7.1:

<https://github.com/docker/docker/blob/master/CHANGELOG.md>. Additionally, Red Hat Enterprise Linux Atomic Host 7.1.4 includes the following changes:

- - FirewallD is now supported for docker containers. If firewallD is running on the system, the rules will be added via the firewallD passthrough. If firewallD is reloaded, the configuration will be re-applied.
  - Docker now mounts the cgroup information specific to a container under the **/sys/fs/cgroup** directory. Some applications make decisions based on the amount of resources available to them. For example, a Java Virtual Machines (JVMs) would want to check how much memory is available to them so they can allocate a large enough pool to improve their performance. This allows applications to discover the maximum amount of memory available to the container, by reading **/sys/fs/cgroup/memory**.
  - The **docker run** command now emits a warning message if you are using a device mapper on a loopback device. It is strongly recommended to use the **dm.thinpooldev** option as a storage option for a production environment. Do not use **loopback** in a production environment.
  - You can now run containers in systemd mode with the **--init=systemd** flag. If you are running a container with systemd as PID 1, this flag will turn on all systemd features to allow it to run in a non-privileged container. Set **container\_uid** as an environment variable to pass to systemd what to store in the **/etc/machine-id** file. This file links the journald within the container to the external log. Mount host directories into a container so systemd will not require privileges then mount the journal directory from the host into the container. If you run journald within the container, the host **journalctl** utility will be able to display the content. Mount the **/run** directory as tmpfs. Then automatically mount the **/sys/fs/cgroup** directory as read-only into a container if **--systemd** is specified. Send proper signal to systemd when running in systemd mode.
  - The search experience within containers using the **docker search** command has been improved:
    - You can now prepend indices to search results.
    - You can prefix a remote name with a registry name.
    - You can shorten the index name if it is not an IP address.
    - The **--no-index** option has been added to avoid listing index names.
    - The sorting of entries when the index is preserved has been changed: You can sort by **index\_name**, **start\_count**, **registry\_name**, **name** and **description**.
    - The sorting of entries when the index is omitted has been changed: You can sort by **registry\_name**, **star\_count**, **name** and **description**.

- You can now expose configured registry list using the Docker info API.

### Red Hat Enterprise Linux Atomic Host 7.1.3 includes the following updates:

- `docker-storage-setup`
  - `docker-storage-setup` now relies on the Logical Volume Manager (LVM) to extend thin pools automatically. By default, 60% of free space in the volume group is used for a thin pool and it is grown automatically by LVM. When the thin pool is full 60%, it will be grown by 20%.
  - A default configuration file for `docker-storage-setup` is now in **`/usr/lib/docker-storage-setup/docker-storage-setup`**. You can override the settings in this file by editing the **`/etc/sysconfig/docker-storage-setup`** file.
  - Support for passing raw block devices to the docker service for creating a thin pool has been removed. Now the `docker-storage-setup` service creates an LVM thin pool and passes it to docker.
  - The chunk size for thin pools has been increased from 64K to 512K.
  - By default, the partition table for the root user is not grown. You can change this behavior by setting the **`GROWPART=true`** option in the **`/etc/sysconfig/docker-storage-setup`** file.
  - A thin pool is now set up with the **`skip_block_zeroing`** feature. This means that when a new block is provisioned in the pool, it will not be zeroed. This is done for performance reasons. One can change this behavior by using the **`--zero`** option:
 

```
lvchange --zero y thin-pool
```
  - By default, docker storage using the devicemapper graphdriver runs on loopback devices. It is strongly recommended to not use this setup, as it is not production ready. A warning message is displayed to warn the user about this. The user has the option to suppress this warning by passing this storage flag **`dm.no_warn_on_loop_devices=true`**.
- Updates related to handling storage on Docker-formatted containers:
  - NFS Volume Plugins validated with SELinux have been added. This includes using the NFS Volume Plugin to NFS Mount GlusterFS.
  - Persistent volume support validated for the NFS volume plugin only has been added.
  - Local storage (HostPath volume plugin) validated with SELinux has been added. (requires workaround described in the docs)
  - iSCSI Volume Plugins validated with SELinux has been added.
  - GCEPersistentDisk Volume Plugins validated with SELinux has been added. (requires workaround described in the docs)

### Red Hat Enterprise Linux Atomic Host 7.1.2 includes the following updates:

- `docker-1.6.0-11.el7`
  - A completely re-architected Registry and a new Registry API supported by Docker 1.6 that enhance significantly image pulls performance and reliability.

- A new logging driver API which allows you to send container logs to other systems has been added to the docker utility. The **--log driver** option has been added to the **docker run** command and it takes three sub-options: a JSON file, syslog, or none. The **none** option can be used with applications with verbose logs that are non-essential.
- Dockerfile instructions can now be used when committing and importing. This also adds the ability to make changes to running images without having to re-build the entire image. The **commit --change** and **import --change** options allow you to specify standard changes to be applied to the new image. These are expressed in the Dockerfile syntax and used to modify the image.
- This release adds support for custom cgroups. Using the **--cgroup-parent** flag, you can pass a specific cgroup to run a container in. This allows you to create and manage cgroups on their own. You can define custom resources for those cgroups and put containers under a common parent group.
- With this update, you can now specify the default ulimit settings for all containers, when configuring the Docker daemon. For example:

```
docker -d --default-ulimit nproc=1024:2048
```

This command sets a soft limit of 1024 and a hard limit of 2048 child processes for all containers. You can set this option multiple times for different ulimit values, for example:

```
--default-ulimit nproc=1024:2408 --default-ulimit nofile=100:200
```

These settings can be overwritten when creating a container as such:

```
docker run -d --ulimit nproc=2048:4096 httpd
```

This will overwrite the default nproc value passed into the daemon.

- The ability to block registries with the **--block-registry** flag.
- Support for searching multiple registries at once.
- Pushing local images to a public registry requires confirmation.
- Short names are resolved locally against a list of registries configured in an order, with the docker.io registry last. This way, pulling is always done with a fully qualified name.

### Red Hat Enterprise Linux Atomic Host 7.1.1 includes the following updates:

- docker-1.5.0-28.el7
  - IPv6 support: Support is available for globally routed and local link addresses.
  - Read-only containers: This option is used to restrict applications in a container from being able to write to the entire file system.
  - Statistics API and endpoint: Statistics on live CPU, memory, network IO and block IO can now be streamed from containers.
  - The **docker build -f *docker\_file*** command to specify a file other than Dockerfile to be used by docker build.

- The ability to specify additional registries to use for unqualified pulls and searches. Prior to this an unqualified name was only searched in the public Docker Hub.
- The ability to block communication with certain registries with **--block-registry=<registry>** flag. This includes the ability to block the public Docker Hub and the ability to block all but specified registries.
- Confirmation is required to push to a public registry.
- All repositories are now fully qualified when listed. The output of **docker images** lists the source registry name for all images pulled. The output of **docker search** shows the source registry name for all results.

For more information, see [Get Started with Docker Formatted Container Images on Red Hat Systems](#)

## 12.2. CONTAINER ORCHESTRATION

**Red Hat Enterprise Linux Atomic Host 7.1.5 and Red Hat Enterprise Linux 7.1 include the following updates:**

- `kubernetes-1.0.3-0.1.gitb9a88a7.el7`
  - The new `kubernetes-client` subpackage which provides the **kubectl** command has been added to the `kubernetes` component.
- `etcd-2.1.1-2.el7`
  - **etcd** now provides improved performance when using the peer TLS protocol.

**Red Hat Enterprise Linux Atomic Host 7.1.4 and Red Hat Enterprise Linux 7.1 include the following updates:**

- `kubernetes-1.0.0-0.8.gitb2dafda.el7`
  - You can now set up a Kubernetes cluster using the Ansible automation platform.

**Red Hat Enterprise Linux Atomic Host 7.1.3 and Red Hat Enterprise Linux 7.1 include the following updates:**

- `kubernetes-0.17.1-4.el7`
  - `kubernetes` nodes no longer need to be explicitly created in the API server, they will automatically join and register themselves.
  - NFS, GlusterFS and Ceph block plugins have been added to Red Hat Enterprise Linux, and NFS support has been added to Red Hat Enterprise Linux Atomic Host.
- `etcd-2.0.11-2.el7`
  - Fixed bugs with adding or removing cluster members, performance and resource usage improvements.
  - The **GOMAXPROCS** environment variable has been set to use the maximum number of available processors on a system, now `etcd` will use all processors concurrently.
  - The configuration file *must* be updated to include the **-advertise-client-urls** flag when setting the **-listen-client-urls** flag.

## Red Hat Enterprise Linux Atomic Host 7.1.2 and Red Hat Enterprise Linux 7.1 include the following updates:

- kubernetes-0.15.0-0.3.git0ea87e4.el7
  - Enabled the v1beta3 API and sets it as the default API version.
  - Added multi-services.
  - The Kubelet now listens on a secure HTTPS port.
  - The API server now supports client certificate authentication.
  - Enabled log collection from the master pod.
  - New volume support: iSCSI volume plug-in, GlusterFS volume plug-in, Amazon Elastic Block Store (Amazon EBS) volume support.
  - Fixed the NFS volume plug-in \* configure scheduler using JSON.
  - Improved messages on scheduler failure.
  - Improved messages on port conflicts.
  - Improved responsiveness of the master when creating new pods.
  - Added support for inter-process communication (IPC) namespaces.
  - The **--etcd\_config\_file** and **--etcd\_servers** options have been removed from the **kube-proxy** utility; use the **--master** option instead.
- etcd-2.0.9-2.el7
  - The configuration file format has changed significantly; using old configuration files will cause upgrades of etcd to fail.
  - The **etcdctl** command now supports importing hidden keys from the given snapshot.
  - Added support for IPv6.
  - The etcd proxy no longer fails to restart after initial configuration.
  - The **-initial-cluster** flag is no longer required when bootstrapping a single member cluster with the **-name** flag set.
  - etcd 2 now uses its own implementation of the Raft distributed consensus protocol; previous versions of etcd used the goraft implementation.
  - Added the **etcdctl** import command to import the migration snap generated in etcd 0.4.8 to the etcd cluster version 2.0.
  - The **etcdctl** utility now takes port 2379 as its default port.
- The cadvisor package has been obsoleted by the kubernetes package. The functionality of cadvisor is now part of the kubelet sub-package.

Red Hat Enterprise Linux 7.1 includes support for orchestration Linux Containers built using docker technology via kubernetes, flannel and etcd.

## Red Hat Enterprise Linux Atomic Host 7.1.1 and Red Hat Enterprise Linux 7.1 include the following updates:

- etcd 0.4.6-0.13.el7 - a new command, **etcdctl** was added to make browsing and editing etcd easier for a system administrator.
- flannel 0.2.0-7.el7 - a bug fix to support delaying startup until after network interfaces are up.

For more information see [Get Started Orchestrating Containers with Kubernetes](#).

## 12.3. COCKPIT ENABLEMENT

### Red Hat Enterprise Linux Atomic Host 7.1.5 and Red Hat Enterprise Linux 7.1 include the following updates:

- The **Cockpit Web Service** is now available as a privileged container. This allows you to run Cockpit on systems like Red Hat Enterprise Linux Atomic Host where the `cockpit-ws` package cannot be installed, but other prerequisites of Cockpit are included. To use this privileged container, use the following command:

```
$ sudo atomic run rhel7/cockpit-ws
```

- Cockpit now includes the ability to access other hosts using a single instance of the Cockpit Web Service. This is useful when only one machine is reachable by the user, or to manage other hosts that do not have the Cockpit Web Service installed. The other hosts should have the `cockpit-bridge` and `cockpit-shell` packages installed.
- The authorized SSH keys for a particular user and system can now be configured using the "Administrator Accounts" section.
- Cockpit now uses the new **storaged** system API to configure and monitor disks and file systems.

### Red Hat Enterprise Linux Atomic Host 7.1.2 and Red Hat Enterprise Linux 7.1 include the following updates:

- `libssh` – a multiplatform C library which implements the SSHv1 and SSHv2 protocol on client and server side. It can be used to remotely execute programs, transfer files, use a secure and transparent tunnel for remote programs. The Secure FTP implementation makes it easier to manager remote files.
- `cockpit-ws` – The **cockpit-ws** package contains the web server component used for communication between the browser application and various configuration tools and services like **cockpitd**. **cockpit-ws** is automatically started on system boot. The `cockpit-ws` package has been included in Red Hat Enterprise Linux 7.1 only.

## 12.4. CONTAINERS USING THE LIBVIRT-LXC TOOLING HAVE BEEN DEPRECATED

The following `libvirt-lxc` packages are deprecated since Red Hat Enterprise Linux 7.1:

- `libvirt-daemon-driver-lxc`
- `libvirt-daemon-lxc`

- libvirt-login-shell

Future development on the Linux containers framework is now based on the docker command-line interface. libvirt-lxc tooling may be removed in a future release of Red Hat Enterprise Linux (including Red Hat Enterprise Linux 7) and should not be relied upon for developing custom container management applications.



## 第 13 章 认证和互操作性

### Manual Backup and Restore Functionality

This update introduces the **ipa-backup** and **ipa-restore** commands to Identity Management (IdM), which allow users to manually back up their IdM data and restore them in case of a hardware failure. For further information, see the [ipa-backup\(1\)](#) and [ipa-restore\(1\)](#) manual pages or the documentation in the [Linux Domain Identity, Authentication, and Policy Guide](#).

### 支持 WinSync 到 Trust 的迁移

This update implements the new **ID Views** mechanism of user configuration. It enables the migration of Identity Management users from a **WinSync** synchronization-based architecture used by **Active Directory** to an infrastructure based on Cross-Realm Trusts. For the details of **ID Views** and the migration procedure, see the documentation in the [Windows Integration Guide](#).

### One-Time Password Authentication

One of the best ways to increase authentication security is to require two factor authentication (2FA). A very popular option is to use one-time passwords (OTP). This technique began in the proprietary space, but over time some open standards emerged (HOTP: RFC 4226, TOTP: RFC 6238). Identity Management in Red Hat Enterprise Linux 7.1 contains the first implementation of the standard OTP mechanism. For further details, see the documentation in the [System-Level Authentication Guide](#).

### 为通用互联网文件系统整合 SSSD

A plug-in interface provided by **SSSD** has been added to configure the way in which the **cifs-utils** utility conducts the ID-mapping process. As a result, an **SSSD** client can now access a CIFS share with the same functionality as a client running the **Winbind** service. For further information, see the documentation in the [Windows Integration Guide](#).

### 证书授权管理工具

The **ipa-cacert-manage renew** command has been added to the Identity management (IdM) client, which makes it possible to renew the IdM Certification Authority (CA) file. This enables users to smoothly install and set up IdM using a certificate signed by an external CA. For details on this feature, see the [ipa-cacert-manage\(1\)](#) manual page.

### 增大访问控制精度

It is now possible to regulate read permissions of specific sections in the Identity Management (IdM) server UI. This allows IdM server administrators to limit the accessibility of privileged content only to chosen users. In addition, authenticated users of the IdM server no longer have read permissions to all of its contents by default. These changes improve the overall security of the IdM server data.

### 特权用户的有限域访问

The **domains=** option has been added to the **pam\_sss** module, which overrides the **domains=** option in the [/etc/sss/sss.conf](#) file. In addition, this update adds the **pam\_trusted\_users** option, which allows the user to add a list of numerical UIDs or user names that are trusted by the **SSSD** daemon, and the **pam\_public\_domains** option and a list of domains accessible even for untrusted users. The mentioned additions allow the configuration of systems, where regular users are allowed to access the specified applications, but do not have login rights on the system itself. For additional information on this feature, see the documentation in the [Linux Domain Identity, Authentication, and Policy Guide](#).

### 自动数据供应者 (Automatic data provider) 配置

现在 **ipa-client-install** 命令默认将 **SSSD** 配置为 sudo 服务的数据供应者。可使用 **--no-sudo** 选项禁用这个行为。此外还添加了 **--nisdomain** 选项为身份管理客户端安装指定 NIS 域名，添加了 **--no\_nisdomain** 选项以避免设置 NIS 域名。如上述两个选项均未使用，则使用 IPA 域。

### AD 和 LDAP sudo 提供者用法

AD 提供者是与 Active Directory 服务器连接的后端程序。在 Red Hat Enterprise Linux 7.1 中将 AD sudo 提供者与 LDAP 提供者一同使用作为技术预览支持。要启用 AD sudo 提供者，请在 **sssd.conf** 文件的 domain 部分添加 **sudo\_provider=ad** 设置。

### 32-bit Version of krb5-server and krb5-server-ldap Deprecated

The 32-bit version of **Kerberos 5 Server** is no longer distributed, and the following packages are deprecated since Red Hat Enterprise Linux 7.1: `krb5-server.i686`, `krb5-server.s390`, `krb5-server.ppc`, `krb5-server-ldap.i686`, `krb5-server-ldap.s390`, and `krb5-server-ldap.ppc`. There is no need to distribute the 32-bit version of `krb5-server` on Red Hat Enterprise Linux 7, which is supported only on the following architectures: AMD64 and Intel 64 systems (**x86\_64**), 64-bit IBM Power Systems servers (**ppc64**), and IBM System z (**s390x**).

### SSSD Leverages GPO Policies to Define HBAC

SSSD is now able to use GPO objects stored on an AD server for access control. This enhancement mimics the functionality of Windows clients, allowing to use a single set of access control rules to handle both Windows and Unix machines. In effect, Windows administrators can now use GPOs to control access to Linux clients.

### Apache Modules for IPA

A set of Apache modules has been added to Red Hat Enterprise Linux 7.1 as a Technology Preview. The Apache modules can be used by external applications to achieve tighter interaction with Identity Management beyond simple authentication.

## 第 14 章 安全性

### SCAP Security Guide

在 Red Hat Enterprise Linux 7.1 中添加了 `scap-security-guide` 软件包，提供安全指导、基线及相关验证机制。该指南由 *安全内容自动化协议* (SCAP) 指定，该协议由一组可实践硬性建议组成。**SCAP Security Guide** 控制根据所述安全策略要求执行安全合规扫描的必要数据，其中包括写入描述及自动测试（探测）。通过自动测试，**SCAP Security Guide** 可提供常规确定系统合规性的便利及可靠的方法。

The Red Hat Enterprise Linux 7.1 version of the **SCAP Security Guide** includes the *Red Hat Corporate Profile for Certified Cloud Providers (RH CCP)*, which can be used for compliance scans of Red Hat Enterprise Linux Server 7.1 cloud systems.

Also, the Red Hat Enterprise Linux 7.1 `scap-security-guide` package contains SCAP datastream content format files for Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7, so that remote compliance scanning of both of these products is possible.

The Red Hat Enterprise Linux 7.1 system administrator can use the **oscap** command line tool from the `openscap-scanner` package to verify that the system conforms to the provided guidelines. See the `scap-security-guide(8)` manual page for further information.

### SELinux 策略

在 Red Hat Enterprise Linux 7.1 中修改了 SELinux 策略，之前在 `init_t` 域中不使用 SELinux 策略的服务现在是在新添加的 `unconfined_service_t` 域中运行。详情请查看 Red Hat Enterprise Linux 7.1 《[SELinux 用户及管理员指南](#)》中《[自由进程](#)》一章。

### OpenSSH 中的新功能

已将 OpenSSH 工具组更新至版本 6.6.1p1，其中添加了几个与加密有关的新功能：

- 现在支持使用 Daniel Bernstein **Curve25519** 中的椭圆曲线 **Diffie-Hellman** 交换密钥。现在默认在支持此方法的服务器和客户端中提供这个方法。
- 添加对使用 **Ed25519** 椭圆曲线签名方案作为公钥类型的支持，该方案可用于用户和主机密钥，提供比 **ECDSA** 和 **DSA** 更好的安全性及良好性能。
- 添加了使用 **bcrypt** 密钥衍生功能 (KDF) 的新私钥格式。默认在 **Ed25519** 密钥中使用此格式，但其他密钥类型也可能要求使用这个格式。
- 添加了新的传输密码 **chacha20-poly1305@openssh.com**。它由 Daniel Bernstein 的 **ChaCha20** 流密码及 **Poly1305** 信息认证代码 (MAC) 组成。

### Libreswan 的新功能

IPsec VPN 的 **Libreswan** 实施已更新至版本 3.12，该更新添加了几个新功能和改进：

- 添加新密码。
- **IKEv2** support has been improved.
- 在 **IKEv1** 和 **IKEv2** 中添加了中间证书链支持。
- 改进连接处理。
- 改进与 OpenBSD、Cisco 和 Android 系统的互操作性。
- 改进了 **systemd** 支持。
- 添加了散列 **CERTREQ** 和流量统计支持。

## TNC 中的新功能

The Trusted Network Connect (TNC) Architecture, provided by the `strongimcv` package, has been updated and is now based on **strongSwan 5.2.0**. The following new features and improvements have been added to the TNC:

- The **PT-EAP** transport protocol ([RFC 7171](#)) for Trusted Network Connect has been added.
- The Attestation *Integrity Measurement Collector* (IMC)/*Integrity Measurement Verifier* (IMV) pair now supports the IMA-NG measurement format.
- 通过使用新的 TPMRA 工作项改进认证 IMV 支持。
- 为使用 SWID IMV 基于 JSON 的 REST API 添加支持。
- The SWID IMC can now extract all installed packages from the `dpkg`, `rpm`, or `pacman` package managers using the `swidGenerator`, which generates SWID tags according to the new ISO/IEC 19770-2:2014 standard.
- The `libtls TLS 1.2` implementation as used by **EAP-(T)TLS** and other protocols has been extended by AEAD mode support, currently limited to **AES-GCM**.
- Improved (IMV) support for sharing access requestor ID, device ID, and product information of an access requestor via a common `imv_session` object.
- 修复了现有 **IF-TNCCS (PB-TNC、IF-M、PA-TNC)** 协议及 **OS IMC/IMV** 对中的几个 bug。

## GnuTLS 的新功能

将 **SSL、TLS 和 DTLS** 协议中的 **GnuTLS** 实施更新至版本 3.3.8, 提供大量新功能和改进:

- 添加 **DTLS 1.2** 支持。
- 添加 *应用程序层协议谈判* (ALPN) 支持。
- 改进椭圆曲线密码套件性能。
- 添加了新密码套件 **RSA-PSK** 和 **CAMELLIA-GCM**。
- 添加了内置 *可信平台模块* (TPM) 标准支持。
- 以多种方式改进了对 **PKCS#11** 智能卡和 *硬件安全模块* (HSM) 的支持。
- 以多种方式改进了对 *FIPS 140* 安全标准 (*联邦信息处理标准*) 的遵循。

---

## 第 15 章 桌面

### Mozilla Thunderbird

**Mozilla Thunderbird**, provided by the `thunderbird` package, has been added in Red Hat Enterprise Linux 7.1 and offers an alternative to the **Evolution** mail and newsgroup client.

### 支持四组缓冲 OpenGL 立体视觉效果

**GNOME Shell** 和 **Mutter** 组成了窗口管理程序，现在可让您在支持的硬件中使用四组缓冲 OpenGL 立体视觉效果。您必须安装 NVIDIA 显示驱动程序版本 337 或者更新的版本方可使用这个功能。

### 在线帐户供应商

在 **GNOME Online Accounts** 中添加了新的 **GSettings** 密钥 **org.gnome.online-accounts.whitelisted-providers**（由 `gnome-online-accounts` 软件包提供）。该密钥提供在线帐户供应商列表，您可在启动时载入该列表。指定这个密钥后，系统管理员就可以启用正确的供应商或者选择性禁用其他供应商。

## 第 16 章 支持和维护

### ABRT 授权的微报告

In Red Hat Enterprise Linux 7.1, the **Automatic Bug Reporting Tool (ABRT)** receives tighter integration with the Red Hat Customer Portal and is capable of directly sending micro-reports to the Portal. **ABRT** provides a utility, **abrt-auto-reporting**, to easily configure user's Portal credentials necessary to authorize micro-reports.

The integrated authorization allows **ABRT** to reply to a micro-report with a rich text which may include possible steps to fix the cause of the micro-report. For example, **ABRT** can suggest which packages are supposed to be upgraded or offer Knowledge base articles related to the issue.


有关 [这个功能的详情](#) 请查看客户门户网站。

## 第 17 章 RED HAT 软件集合

Red Hat Software Collections 是一个 Red Hat 内容套件，可提供一组可在 AMD64 和 Intel 64 架构中的 Red Hat Enterprise Linux 6 和 Red Hat Enterprise Linux 7 支持发行本中安装和使用的动态编程语言、数据库服务器及相关软件包。

Red Hat Software Collections 发布的动态语言、数据库服务器及其他工具既不能替换 Red Hat Enterprise Linux 提供的默认系统工具，也不能作为这类工具的首选。

Red Hat Software Collections 使用基于 **scl** 程序的备用打包机制提供一组平行软件包。这个软件包组可让您在 Red Hat Enterprise Linux 中使用备选软件包版本。用户可使用 **scl** 程序选择在任意时间要运行的软件包版本。



### 重要

Red Hat Software Collections 比 Red Hat Enterprise Linux 的生命周期和支持时限更短。详情请查看 [《Red Hat Software Collections 产品生命周期》](#)。

Red Hat Developer Toolset 现在是 Red Hat Software Collections 的一部分，其中包括独立软件集合。Red Hat Developer Toolset 旨在让开发人员在 Red Hat Enterprise Linux 平台中工作。它提供 GNU 编译程序集合、GNU Debugger、Eclipse 开发平台以及其他开发、调试和性能监控工具的最新版本。

有关本集合所包含内容、系统要求、已知问题、用法及具体的软件集合详情请查看 [《Red Hat Software Collections 文档》](#)。

有关中国软件集合所包含文档、安装、用法、已知问题等等内容详情请查看 [《Red Hat Developer Toolset 文档》](#)。

## 第 18 章 RED HAT ENTERPRISE LINUX FOR REAL TIME

Red Hat Enterprise Linux for Real Time is a new offering in Red Hat Enterprise Linux 7.1 comprised of a special kernel build and several user space utilities. With this kernel and appropriate system configuration, Red Hat Enterprise Linux for Real Time brings deterministic workloads, which allow users to rely on consistent response times and low and predictable latency. These capabilities are critical in strategic industries such as financial service marketplaces, telecommunications, or medical research.

For instructions on how to install Red Hat Enterprise Linux for Real Time, and how to set up and tune the system so that you can take full advantage of this offering, refer to the [Red Hat Enterprise Linux for Real Time 7 Installation Guide](#).



## 部分 II. TECHNOLOGY PREVIEWS

This part provides an overview of Technology Previews introduced or updated in Red Hat Enterprise Linux 7.1.

For more information on Red Hat Technology Previews, see <https://access.redhat.com/support/offerings/techpreview/>.

## 第 19 章 HARDWARE ENABLEMENT

- OSA-Express5s Cards Support in **qethqoat**, see [第 2.5 节 “OSA-Express5s Cards Support in qethqoat”](#)

## 第 20 章 STORAGE

- **Targetd** plug-in from the **libStorageMgmt** API, see [“使用 libStorageMgmt API 进行存储阵列管理”一节](#)
- LSI Syncro CS HA-DAS adapters, see [“LSI Syncro 支持”一节](#)
- DIF/DIX, see [“DIF/DIX 支持”一节](#)

## 第 21 章 FILE SYSTEMS

- **Btrfs** file system, see [“支持 Btrfs 文件系统”一节](#)
- **OverlayFS**, see [“OverlayFS”一节](#)

---

## 第 22 章 KERNEL

- **kpatch**, see [“动态内核补丁”一节](#)
- **crashkernel** with more than one CPU, see [“有一个以上 CPU 的 crashkernel”一节](#)
- **dm-era** device-mapper target, see [“dm-era 目标”一节](#)
- Cisco VIC kernel driver, see [“Cisco VIC 内核驱动程序”一节](#)
- NFSoRDMA Available, see [“NFSoRDMA Available”一节](#)
- Runtime Instrumentation for IBM System z, see [“Runtime Instrumentation for IBM System z”一节](#)
- Cisco usNIC Driver, see [“Cisco usNIC Driver”一节](#)
- Intel Ethernet Server Adapter X710/XL710 Driver Update, see [“Intel Ethernet Server Adapter X710/XL710 Driver Update”一节](#)

## 第 23 章 VIRTUALIZATION

- USB 3.0 host adapter (xHCI) emulation, see [“KVM 虚拟机的 USB 3.0 支持”一节](#)
- Open Virtual Machine Firmware (OVMF), see [“Open Virtual Machine Firmware”一节](#)
- LPAR Watchdog for IBM System z, see [“LPAR Watchdog for IBM System z”一节](#)

## 第 24 章 COMPILER AND TOOLS

- Accurate ethtool Output, see [“Accurate ethtool Output”一节](#)

## 第 25 章 NETWORKING

- Trusted Network Connect, see [“可信的网络连接”一节](#)
- SR-IOV functionality in the **qlcnic** driver, see [“qlcnic 驱动程序中的 SR-IOV 功能”一节](#)



## 第 26 章 AUTHENTICATION AND INTEROPERABILITY

- Use of AD sudo provider together with the LDAP provider, see [“AD 和 LDAP sudo 提供者用法”一节](#)
- Apache Modules for IPA, see [“Apache Modules for IPA”一节](#)

## 部分 III. 设备驱动程序

本章提供了在 Red Hat Enterprise Linux 7.1 中更新的所有设备的完整列表。

## 第 27 章 存储驱动程序更新

- 已将 **hpsa** 驱动程序升级至版本 3.4.4-1-RH1。
- 已将 **qla2xxx** 驱动程序升级至版本 8.07.00.08.07.1-k1。
- 已将 **qla4xxx** 驱动程序升级至版本 5.04.00.04.07.01-k0。
- 已将 **qlcnic** 驱动程序升级至版本 5.3.61。
- 已将 **netxen\_nic** 驱动程序升级至版本 4.0.82。
- 已将 **qlge** 驱动程序升级至版本 1.00.00.34。
- 已将 **bnx2fc** 驱动程序升级至版本 2.4.2。
- 已将 **bnx2i** 驱动程序升级至版本 2.7.10.1。
- 已将 **cnic** 驱动程序升级至版本 2.5.20。
- 已将 **bnx2x** 驱动程序升级至版本 1.710.51-0。
- 已将 **bnx2** 驱动程序升级至版本 2.2.5。
- 已将 **megaraid\_sas** 驱动程序升级至版本 06.805.06.01-rc1。
- 已将 **mpt2sas** 驱动程序升级至版本 18.100.00.00。
- 已将 **ipr** 驱动程序升级至版本 2.6.0。
- 在 Red Hat Enterprise Linux 7 中添加了 **kmod-lpfc** 软件包，该软件包可保证在附带光纤（FC）和以太网光线（FCoE）中使用 **lpfc** 驱动程序时有更好的稳定新。已将 **lpfc** 驱动程序升级至版本 0:10.2.8021.1。
- 已将 **be2iscsi** 驱动程序升级至版本 10.4.74.0r。
- 已将 **nvme** 驱动程序升级至版本 0.9。

## 第 28 章 网络驱动程序更新

- 已将 **bnx** 驱动程序升级至版本 3.2.23.Or。
- 已将 **cxgb3** 驱动程序升级至版本 1.1.5-ko。
- 已将 **cxgb3i** 驱动程序升级至版本 2.0.0。
- 已将 **iw\_cxgb3** 驱动程序升级至版本 1.1。
- 已将 **cxgb4** 驱动程序升级至版本 2.0.0-ko。
- 已将 **cxgb4vf** 驱动程序升级至版本 2.0.0-ko。
- 已将 **cxgb4i** 驱动程序升级至版本 0.9.4。
- 已将 **iw\_cxgb4** 驱动程序升级至版本 0.1。
- 已将 **e1000e** 驱动程序升级至版本 2.3.2-k。
- 已将 **igb** 驱动程序升级至版本 5.2.13-k。
- 已将 **igbvf** 驱动程序升级至版本 2.0.2-k。
- 已将 **ixgbe** 驱动程序升级至版本 3.19.1-k。
- 已将 **ixgbev** 驱动程序升级至版本 2.12.1-k。
- 已将 **i40e** 驱动程序升级至版本 1.0.11-k。
- 已将 **i40evf** 驱动程序升级至版本 1.0.1。
- 已将 **e1000** 驱动程序升级至版本 7.3.21-k8-NAPI。
- 已将 **mlx4\_en** 驱动程序升级至版本 2.2-1。
- 已将 **mlx4\_ib** 驱动程序升级至版本 2.2-1。
- 已将 **mlx5\_core** 驱动程序升级至版本 2.2-1。
- 已将 **mlx5\_ib** 驱动程序升级至版本 2.2-1。
- 已将 **ocrdma** 驱动程序升级至版本 10.2.287.Ou。
- 已将 **ib\_ipoib** 驱动程序升级至版本 1.0.0。
- 已将 **ib\_qib** 驱动程序升级至版本 1.11。
- 已将 **enic** 驱动程序升级至版本 2.1.1.67。
- 已将 **be2net** 驱动程序升级至版本 10.4r。
- 已将 **tg3** 驱动程序升级至版本 3.137。
- 已将 **r8169** 驱动程序升级至版本 2.3LK-NAPI。

## 第 29 章 图形驱动程序更新

- 已将 **vmwgfx** 驱动程序升级至版本 2.6.0.0。

## 部分 IV. DEPRECATED FUNCTIONALITY

This part provides an overview of functionality that has been deprecated in all minor releases up to Red Hat Enterprise Linux 7.1.

## 第 30 章 DEPRECATED FUNCTIONALITY IN RED HAT ENTERPRISE LINUX 7

### Symbols from libraries linked as dependencies no longer resolved by `ld`

Previously, the `ld` linker resolved any symbols present in any linked library, even if some libraries were linked only implicitly as dependencies of other libraries. This allowed developers to use symbols from the implicitly linked libraries in application code and omit explicitly specifying these libraries for linking.

For security reasons, `ld` has been changed to not resolve references to symbols in libraries linked implicitly as dependencies.

As a result, linking with `ld` fails when application code attempts to use symbols from libraries not declared for linking and linked only implicitly as dependencies. To use symbols from libraries linked as dependencies, developers must explicitly link against these libraries as well.

To restore the previous behavior of `ld`, use the `-copy-dt-needed-entries` command-line option. (BZ#[1292230](#))

### Windows guest virtual machine support limited

As of Red Hat Enterprise Linux 7, Windows guest virtual machines are supported only under specific subscription programs, such as Advanced Mission Critical (AMC).

## 部分 V. KNOWN ISSUES

This part describes known issues in Red Hat Enterprise Linux 7.1.



## 第 31 章 INSTALLATION AND BOOTING

### anaconda component, BZ#1067868

Under certain circumstances, when installing the system from the boot DVD or ISO image, not all assigned IP addresses are shown in the network spoke once network connectivity is configured and enabled. To work around this problem, leave the network spoke and enter it again. After re-entering, all assigned addresses are shown correctly.

### anaconda component, BZ#1085310

Network devices are not automatically enabled during installation unless the installation method requires network connectivity. As a consequence, a traceback error can occur during Kickstart installation due to inactive network devices. To work around this problem, set the **ksdevice=link** option on boot or add the **--device=link** option to the **ks.cfg** file to enable network devices with active links during Kickstart installation.

### anaconda component, BZ#1185280

An interface with IPv6-only configuration does not bring up the network interface after manual graphical installation from an IPv6 source. Consequently, the system boots with the interface set to **ONBOOT=no**, and consequently the network connection does not work. Select the **Automatically connect to network** check box if available, or use kickstart with a command as follows:

```
network --noipv4 --bootproto=dhcp --activate
```

In both cases IPv6 will be configured to be active on system start.

If the network interface is set to IPv4 **and** IPv6 configuration, and is installed from an IPv6 address, after installation it will be configured to be active on system start (**ONBOOT=yes**).

### anaconda component, BZ#1085325

The **anaconda** installer does not correctly handle adding of FCoE disks. As a consequence, adding FCoE disks on the **anaconda** advance storage page fails with the following error message:

```
No Fibre Channel Forwarders or VN2VN Responders Found
```

To work around this problem, simply repeat the steps to add the FCoE disks; the configuration process produces the correct outcome when repeated. Alternatively, run the **lldpad -d** command in the **anaconda** shell before adding the FCoE disks in the **anaconda** user interface to avoid the described problem.

### anaconda component, BZ#1087774

The source code does not handle booting on a **bnx2i** iSCSI driver correctly. As a consequence, when installing Red Hat Enterprise Linux 7.1, the server does not reboot automatically after the installation is completed. No workaround is currently available.

### anaconda component, BZ#965985

When booting in rescue mode on IBM System z architecture, the second and third rescue screens in the rescue shell are incomplete and not displayed properly.

### anaconda component, BZ#1190146

When the **/boot** partition is not separated and the **boot=** parameter is specified on the kernel command line, an attempt to boot the system in the FIPS mode fails. To work around this issue, remove the **boot=** parameter from the kernel command line.

#### **anaconda component, BZ#1174451**

When the user inserts a space character anywhere between nameservers while configuring the nameservers in the **Network Configuration** dialog during a text-mode installation, the installer terminates unexpectedly. To work around this problem, if you want to configure multiple nameservers during the **Network Configuration** step of the installation, enter them in a comma-separated list without spaces between the nameservers. For example, while entering **1.1.1.1, 2.1.2.1** with a space in this situation causes the installer to crash, entering **1.1.1.1,2.1.2.1** without a space ensures the installer handles configuring multiple nameservers correctly and does not crash.

#### **anaconda component, BZ#1166652**

If the installation system has multiple iSCSI storage targets connected over separate active physical network interfaces, the installer will hang when starting iSCSI target discovery in the **Installation Destination** screen.

The same issue also appears with an iSCSI multipath target accessible over two different networks, and happens no matter whether the **Bind targets to network interfaces** option is selected.

To work around this problem, make sure only one active physical network interface has an available iSCSI target, and attach any additional targets on other interfaces after the installation.

#### **anaconda component, BZ#1168169**

When using a screen resolution of less than 1024x768 (such as 800x600) during a manual installation, some of the controls in the **Manual Partitioning** screen become unreachable. This problem commonly appears when connecting to the installation system using a VNC viewer, because by default the VNC server is set to 800x600.

To work around this issue, set the resolution to 1024x768 or higher using a boot option. For example:

```
linux inst.vnc inst.resolution=1024x768
```

For information about **Anaconda** boot options, see the [Red Hat Enterprise Linux 7.1 Installation Guide](#).

#### **dracut component, BZ#1192480**

A system booting with iSCSI using IPv6 times out while trying to connect to the iSCSI server after about 15 minutes, but then connects successfully and boots as expected.

#### **kernel component, BZ#1055814**

When installing Red Hat Enterprise Linux 7 on UEFI-based systems, the Anaconda installer terminates unexpectedly with the following error:

```
BootLoaderError: failed to remove old efi boot entry
```

To work around this problem, edit the **Install Red Hat Enterprise Linux 7** option in the boot menu by pressing the **e** key and append the **efi\_no\_storage\_paranoia** kernel parameter to the end of the line that begins with **linuxefi**. Then press the **F10** key to boot the modified option and start installation.

**sg3\_utils component, BZ#1186462**

Due to the conversion of the `iprutils` package to use **systemd** instead of legacy init scripts, the **sg** driver is no longer loaded during system boot. Consequently, if the **sg** driver is not loaded, the `/dev/sg*` devices will not be present.

To work around this issue, manually issue **modprobe sg** or add it to an init script. Once the **sg** driver is loaded, the `/dev/sg*` devices will be present and the **sg** driver may be used to access SCSI devices.

**anaconda component, BZ#1072619**

It is not possible to use read-only disks as hard drive installation repository sources. When specifying the **inst.repo=hd:device:path** option ensure that `device` is writable.

**kernel component, BZ#1067292, BZ#1008348**

Various platforms include BIOS or UEFI-assisted software RAID provided by LSI. This hardware requires the closed-source **megasr** driver, which is not included in Red Hat Enterprise Linux. Thus, platforms and adapters that depend on **megasr** are not supported by Red Hat. Also, the use of certain open-source RAID alternatives, such as the **dmraid** Disk Data Format 1 (DDF1) capability, is not currently supported on these systems.

However, on certain systems, such as IBM System x servers with the ServeRAID adapter, it is possible to disable the BIOS RAID function. To do this, enter the UEFI menu and navigate through the **System Settings** and **Devices and I/O Ports** submenus to the **Configure the onboard SCU** submenu. Then change the SCU setting from **RAID** to **nonRAID**. Save your changes and reboot the system. In this mode, the storage is configured using an open-source non-RAID LSI driver shipped with Red Hat Enterprise Linux, such as **mptsas**, **mpt2sas**, or **mpt3sas**.

To obtain the **megasr** driver for IBM systems, refer to the [IBM support page](#).

Certain Cisco Unified Computing System (UCS) platforms are also impacted by this restriction. However, it is not possible to disable the BIOS RAID function on these systems. To obtain the **megasr** driver, refer to the [Cisco support page](#).

**注意**

The described restriction does not apply to LSI adapters that use the **megaraid** driver. Those adapters implement the RAID functions in the adapter firmware.

**kernel component, BZ#1168074**

During CPU hot plugging, the kernel can sometimes issue the following warning message:

```
WARNING: at block/blk-mq.c:701__blk_mq_run_hw_queue+0x31d/0x330()
```

The message is harmless, and you can ignore it.

**kernel component, BZ#1097468**

The Linux kernel Non-Uniform Memory Access (NUMA) balancing does not always work correctly. As a consequence, when the **numa\_balancing** parameter is set, some of the memory can move to an arbitrary non-destination node before moving to the constrained nodes, and the memory on the destination node also decreases under certain circumstances. There is currently no known workaround available.

**kernel component, BZ#1087796**

An attempt to remove the **bnx2x** module while the **bnx2fc** driver is processing a corrupted frame causes a kernel panic. To work around this problem, shut down any active FCoE interfaces before executing the **modprobe -r bnx2x** command.

#### kernel component, BZ#915855

The QLogic 1G iSCSI Adapter present in the system can cause a call trace error when the **qla4xx** driver is sharing the interrupt line with the USB sub-system. This error has no impact on the system functionality. The error can be found in the kernel log messages located in the **/var/log/messages** file. To prevent the call trace from logging into the kernel log messages, add the **nousb** kernel parameter when the system is booting.

#### kernel component, BZ#1164997

When using the **bnx2x** driver with a BCM57711 device and sending traffic over Virtual Extensible LAN (VXLAN), the transmitted packets have bad checksums. Consequently, communication fails, and **UDP: bad checksum** messages are displayed in the kernel log on the receiving side. To work around this problem, disable checksum offload on the **bnx2x** device using the **ethtool** utility.

#### kernel component, BZ#1164114

If you change certain parameters while the Network Interface Card (NIC) is set to **down**, the system can become unresponsive if you are using a **qlge** driver. This problem occurs due to a race condition between the New API (NAPI) registration and unregistration. There is no workaround currently available.

#### system-config-kdump component, BZ#1077470

In the **Kernel Dump Configuration** window, selecting the **Raw device** option in the **Target settings** tab does not work. To work around this problem, edit the **kdump.conf** file manually.

#### yaboot component, BZ#1032149

Due to a bug in the **yaboot** boot loader, upgrading from Red Hat Enterprise Linux 6 to Red Hat Enterprise Linux 7 can fail on the IBM Power Systems with an **Unknown or corrupt filesystem** error.

#### util-linux component, BZ#1171155

The **anaconda** installer cannot handle disks with labels from the IBM AIX operating systems correctly. As a consequence, an attempt to install Red Hat Enterprise Linux on such a disk fails. Users are advised to not use disks with AIX labels in order prevent the installation failures.

#### kernel component, BZ#1192470

If you attempt to perform an in-place upgrade from Red Hat Enterprise Linux 6.6 running on IBM System z architecture to Red Hat Enterprise Linux 7.1 and have the **kernel-kdump** package installed on Red Hat Enterprise Linux 6.6, the **kdump** boot record is not removed. Consequently, the upgrade fails when the **zipl** utility is called. To work around this problem, remove the **kdump** boot record from the **/etc/zipl.conf** file before performing the upgrade.

#### anaconda component, BZ#1171778

Setting only full name and no user name for a new user in text installation does not require root password to be set. As a consequence, when such a user is configured and no root password is set, the user is not able to log in either, and neither is root. There is also no straightforward way to create a user or set the root password after such an installation since initial-setup crashes due to this bug. To work around this problem, set the root password during installation or set the user name for the user during text installation.

**python-blivet component, BZ#1192004**

The installer terminates unexpectedly if you set up partitioning before adding an iSCSI disk and then set up partitioning again. As a consequence, it is impossible to successfully complete the installation in this situation. To work around this problem, reset storage or reboot before adding iSCSI or FCoE disks during installation.

**anaconda component, BZ#1168902**

The **anaconda** installer expects a **ks.cfg** file if booting with the **inst.ks=cdrom:/ks.cfg** parameter, and enters the emergency mode if the **ks.cfg** file is not provided within several minutes. With some enterprise servers that take a long time to boot, Anaconda does not wait long enough to enable the user to provide the **ks.cfg** file in time.

To work around this problem, add the **rd.retry** boot parameter and use a large value. For example, using **rd.retry=86400** causes a time-out after 24 hours, and using **rd.retry=1<<15** should, in theory, time out after about 34 years, which provides the user with sufficient time in all known scenarios.

**subscription-manager component, BZ#1158396**

The **Back** button used in the **firstboot** utility is not working properly. It is often disabled, and if it is enabled, pressing it has no effect. Consequently, during **Subscription Management Registration**, clicking **Back** does not return you to the previous panel. If you want to go back, enter an invalid server or invalid credentials and click **Done**. After this, either an **Unable to reach the server** dialog or an **Unable to register the system** dialog appears at the top of the initial **firstboot** panel. Dismiss the error dialog, and choose the **No, I prefer to register at a later time** option.

**kernel component, BZ#1076374**

The GRUB2 bootloader supports network booting over the Hypertext Transfer Protocol (HTTP) and the Trivial File Transfer Protocol (TFTP). However, under heavy network traffic, network boot over HTTP is very slow and may cause timeout failures. If this problem occurs, use TFTP to load the kernel and initrd images. To do so, put the boot files in the TFTP server directory and add the following to the **grub.cfg** file where **1.1.1.1** is the address of the TFTP server:

```
insmod tftp
set root=tftp,1.1.1.1
```

**anaconda component, BZ#1164131**

The Driver Update Disk loader does not reconfigure network devices if they have already been configured. Consequently, installations that use a Driver Update Disk to replace an existing, functional network driver with a different version will not be able to use the network to fetch the installer runtime image.

To work around this problem, use the provided network driver during the installation process and update the network driver after the installation.

## 第 32 章 STORAGE

### **kernel component, BZ#1170328**

When the Internet Small Computer System Interface (iSCSI) target is set up using the iSCSI Extensions for RDMA (iSER) interface, an attempt to run a discovery over iSER fails. Consequently, in some cases, the target panics. Users are advised to not use iSER for discovery but use iSER only for the login phase.

### **kernel component, BZ#1185396**

When using the server as an iSER-enabled iSCSI target and connection losses occur repeatedly, the target can stop responding. Consequently, the kernel becomes unresponsive. To work around this issue, minimize iSER connection losses or revert to non-iSER iSCSI mode.

### **kernel component, BZ#1061871, BZ#1201247**

When a storage array returns a CHECK CONDITION status but the sense data is invalid, the Small Computer Systems Interface (SCSI) mid-layer code retries the I/O operation. If subsequent I/O operations receive the same result, I/O operations are retried indefinitely. For this bug, no workaround is currently available.

## 第 33 章 FILE SYSTEMS

### **kernel component, BZ#1172496**

Due to a bug in the ext4 code, it is currently impossible to resize ext4 file systems that have 1 kilobyte block size and are smaller than 32 megabytes.

## 第 34 章 VIRTUALIZATION

### netcf component, BZ#1100588

When installing Red Hat Enterprise Linux 7 from sources other than the network, the network devices are not specified by default in the interface configuration files. As a consequence, creating a bridge by using the **iface-bridge** command in the **virsh** utility fails with an error message. To work around the problem, add the **DEVICE=** lines in the **/etc/sysconfig/network-scripts/ifcfg-\*** files.

### grub2 component, BZ#1045127

Nesting more than 7 PCI bridges is known to cause segmentation fault errors. It is not recommended to create more than 7 nested PCI bridges.

### kernel component, BZ#1075857

The kernel **sym53c8xx** module is not supported in Red Hat Enterprise Linux 7. Therefore, it is not possible to use an emulated Small Computer System Interface (SCSI) disk when Red Hat Enterprise Linux is running as a guest on top of the Xen hypervisor or Amazon Web Services (AWS) Elastic Compute Cloud (EC2). Red Hat recommends to use paravirtualized devices instead.

### kernel component, BZ#1081851

When the **xen\_emulated\_unplug=never** or **xen\_emulated\_unplug=unnecessary** options are passed to the guest kernel command line, an attempt to hot plug a new device to the Xen guest does not work. Running the **xl** command in the host succeeds but no devices appear in the guest. To work around this issue, remove the aforementioned options from the guest kernel command line and use paravirtualized drivers to allow hot plugging. Note that **xen\_emulated\_unplug=never** and **xen\_emulated\_unplug=unnecessary** are supposed to be used for debugging purposes only.

### kernel component, BZ#1035213

After multiple hot plugs and hot unplugs of a SCSI disk in the Hyper-V environment, the disk in some cases logs an error, becomes unusable for several minutes, and displays incorrect information when explored with the **partprobe** command.

### kernel component, BZ#1183960

A prior Intel microcode update removed the Hardware Lock Elision (HLE) and Restricted Transactional Memory (RTM) features from 4th Generation Intel Core Processors, Intel Xeon v3 Processors, and some 5th Generation Intel Core Processors. However, after performing a live migration of a KVM guest from a host containing a CPU without the microcode update to a host containing a CPU with the update, the guest may attempt to continue using HLE and RTM. This can lead to applications on the guest terminating unexpectedly with an **Illegal Instruction** error. To work around this problem, shut down the guest and perform a non-live migration if moving from a CPU with HLE and RTM to a CPU without the features. This ensures that HLE and RTM are unavailable on the guest after the migration, and thus prevents the described crashes.

### systemd component, BZ#1151604, BZ#1147876

Due to an unintended incompatibility between QEMU and the pSeries platform, the **systemd-detect-virt** and **virt-what** commands cannot properly detect PowerKVM virtualization on IBM Power Systems. There is currently no known workaround.

### kernel component, BZ#1153521

When the kernel shared memory (KSM) feature is enabled with the **merge\_across\_nodes=1** parameter, KSM ignores memory policies set by the **mbind()** function, and may merge pages from some memory areas to Non-Uniform Memory Access (NUMA) nodes that do not match the policies.



To work around this issue, disable KSM or set the **merge\_across\_nodes** parameter to **0** if using NUMA memory binding with QEMU, as this leads to NUMA memory policies configured for the KVM VM working as expected.

## 第 35 章 DEPLOYMENT AND TOOLS

### systemd component, BZ#1178848

The **systemd** service cannot set **cgroup** properties on **cgroup** trees that are mounted as read-only. Consequently, the following error message can occasionally appear in the logs:

```
Failed to reset devices.list on /machine.slice: Invalid argument
```

You can ignore this problem, as it should not have any significant effect on your system.

### systemd component, BZ#978955

When attempting to start, stop, or restart a service or unit using the **systemctl [start|stop|restart] NAME** command, no message is displayed to inform the user whether the action has been successful.

### subscription-manager component, BZ#1166333

The Assamese (as-IN), Punjabi (pa-IN), and Korean (ko-KR) translations of **subscription-manager**'s user interface are incomplete. As a consequence, users of **subscription-manager** running in one of these locales may see labels in English rather than the configured language.

### systemtap component, BZ#1184374

Certain functions in the kernel are not probed as expected. To work around this issue, try to probe by a statement or by a related function.

### systemtap component, BZ#1183038

Certain parameters or functions cannot be accessed within function probes. As a consequence, the **\$parameter** accesses can be rejected. To work around this issue, activate the **systemtap** prologue-searching heuristics.

## 第 36 章 COMPILER AND TOOLS

### java-1.8.0-openjdk component, [BZ#1189530](#)

With Red Hat Enterprise Linux 7.1, the java-1.8.0-openjdk packages do not provide "java" in the RPM metadata, which breaks compatibility with packages that require **Java** and are available from the Enterprise Application Platform (EAP) channel. To work around this problem, install another package that provides "java" in the RPM metadata before installing java-1.8.0-openjdk.

## 第 37 章 NETWORKING

### rsync component, BZ#1082496

The **rsync** utility cannot be run as a socket-activated service because the **rsyncd@.service** file is missing from the rsync package. Consequently, the **systemctl start rsyncd.socket** command does not work. However, running **rsync** as a daemon by executing the **systemctl start rsyncd.service** command works as expected.

### InfiniBand component, BZ#1172783

The **libocrdma** package is not included in the default package set of the InfiniBand Support group. Consequently, when users select the InfiniBand Support group and are expecting RDMA over Converged Ethernet (RoCE) to work on Emulex OneConnect adapters, the necessary driver, **libocrdma**, is not installed by default. On first boot, the user can manually install the missing package by issuing this command:

```
~]# yum install libocrdma
```

As a result, the user will now be able to use the Emulex OneConnect devices in RoCE mode.

### vsftpd component, BZ#1058712

The **vsftpd** daemon does not currently support ciphers suites based on the Elliptic Curve Diffie-Hellman Exchange (ECDHE) key-exchange protocol. Consequently, when **vsftpd** is configured to use such suites, the connection is refused with a **no shared cipher SSL** alert.

### arptables component, BZ#1018135

Red Hat Enterprise Linux 7 introduces the **arptables** packages, which replace the **arptables\_jf** packages included in Red Hat Enterprise Linux 6. All users of **arptables** are advised to update their scripts because the syntax of this version differs from **arptables\_jf**.

### openssl component, BZ#1062656

It is not possible to connect to any Wi-Fi Protected Access (WPA) Enterprise Access Point (AP) that requires MD5-signed certificates. To work around this problem, copy the **wpa\_supplicant.service** file from the **/usr/lib/systemd/system/** directory to the **/etc/systemd/system/** directory and add the following line to the **Service** section of the file:

```
Environment=OPENSSL_ENABLE_MD5_VERIFY=1
```

Then run the **systemctl daemon-reload** command as root to reload the service file.



#### 重要

Note that MD5 certificates are highly insecure and Red Hat does not recommend using them.

## 第 38 章 RED HAT ENTERPRISE LINUX ATOMIC HOST

### dracut component, BZ#1160691

Red Hat Enterprise Linux Atomic Host 7.1.0 allows configuring encrypted root installation in the Anaconda installer, but the system will not boot afterwards. Choosing this option in the installer is not recommended.

### dracut component, BZ#1189407

Red Hat Enterprise Linux Atomic Host 7.1.0 offers iSCSI support during Anaconda installation, but the current content set does not include iSCSI support, so the system will not be able to access the storage. Choosing this option in the installer is not recommended.

### kexec-tools component, BZ#1180703

Due to some parsing problems in the code, the kdump utility currently saves the kernel crash dumps in the `/sysroot/crash/` directory instead of in `/var/crash/`.

### rhel-server-atomic component, BZ#1186923

Red Hat Enterprise Linux Atomic Host 7.1.0 does not currently support systemtap, unless the host-kernel-matching packages which contain kernel-devel and other packages are installed into the rheltools container image.

### rhel-server-atomic component, BZ#1193704

Red Hat Enterprise Linux Atomic Host allocates 3GB of storage to the root partition, which includes the docker volumes. In order to support more volume space, more physical storage must be added to the system, or the root Logical Volume must be extended. The [Managing Storage with Red Hat Enterprise Linux Atomic Host](#) section from the Getting Started with Red Hat Enterprise Linux Atomic Host article describes the workaround methods for this issue.

### rhel-server-atomic component, BZ#1186922

If the `ltrace` command is executed inside a Super-Privileged Container (SPC) to trace a process that is running on Red Hat Enterprise Linux Atomic Host, the `ltrace` command is unable to locate the binary images of the shared libraries that are attached to the process to be traced. As a consequence, `ltrace` displays a series of error messages, similar to the following example:

```
Can't open /lib64/libwrap.so.0: No such file or directory
Couldn't determine base address of /lib64/libwrap.so.0
ltrace: ltrace-elf.c:426: ltelf_destroy: Assertion `(&lt;e->plt_relocs)->elt_size == sizeof(GElf_Rel)'
failed.
```

### rhel-server-atomic component, BZ#1187119

Red Hat Enterprise Linux Atomic Host does not include a mechanism to customize or override the content of the host itself, for example it does not include a tool to use a custom kernel for debugging.

### rhel-server-atomic component, BZ#1187119

Red Hat Enterprise Linux Atomic Host does not include a mechanism to customize or override the content of the host itself, for example it does not include a tool to use a custom kernel for debugging.

## 第 39 章 LINUX CONTAINERS

### docker component, BZ#1193609

If docker is setting up loop devices for docker thin pool setup, docker operations like docker deletion and container I/O operations can be slow. The strongly recommended alternative configuration is to set up an LVM thin pool and use it as storage back-end for docker. Instructions on setting up an LVM thin pool can be found in the [lvmthin\(7\)](#) manual page. Then modify the `/etc/sysconfig/docker-storage` file to include the following line to make use of the LVM thin pool for container storage.

```
DOCKER_STORAGE_OPTIONS= --storage-opt dm.thinpooldev=<pool-device>
```

### docker component, BZ#1190492

A Super-Privileged Container (SPC) that is launched while some application containers are already active has access to the file system trees of these application containers. The file system trees reside in device mapper "thin target" devices. Since the SPC holds references on these file system trees, the docker daemon fails to clean up the "thin target" (the device is still "busy") at the time when an application container is terminated. As a consequence, the following error message is logged in the journal of systemd:

```
Cannot destroy container {Id}: Driver devicemapper failed to remove root filesystem {Id}: Device is Busy
```

where `{Id}` is a placeholder for the container runtime ID, and a stale device mapper "thin target" is left behind after an application container is terminated.

### docker component, BZ#1190492

A Super-Privileged Container (SPC) that is launched while some application containers are already active has access to the file system trees of these application containers. The file system trees reside in device mapper "thin target" devices. Since the SPC holds references on these file system trees, the docker daemon fails to clean up the "thin target" (the device is still "busy") at the time when an application container is terminated. As a consequence, the following error message is logged in the journal of systemd:

```
Cannot destroy container {Id}: Driver devicemapper failed to remove root filesystem {Id}: Device is Busy
```

where `{Id}` is a placeholder for the container runtime ID, and a stale device mapper "thin target" is left behind after an application container is terminated.

### docker component, BZ#1188252

The docker daemon can occasionally terminate unexpectedly while a Super-Privileged Container (SPC) is running. Consequently, a stale entry related to the Super-Privileged Container is left behind in `/var/lib/docker/linkgraph.db`, and the container cannot be restarted correctly afterwards.

### gdb component, BZ#1186918

If the GNU debugger (GDB) is executing inside a Super-Privileged Container (SPC) and attaches to a process that is running in another container on Red Hat Enterprise Linux Atomic Host, GDB does not locate the binary images of the main executable or any shared libraries loaded by the process to be debugged. As a consequence, GDB may display error messages relating to files not being present, or being present but mismatched, or GDB may seem to attach correctly but then subsequent commands may fail or display corrupted information. A workaround is to specify the `sysroot` and `file` prior to issuing the command, as follows:

```
set sysroot /proc/PID/root  
file /proc/PID/exe  
attach PID
```

## 第 40 章 AUTHENTICATION AND INTEROPERABILITY

### **bind-dyndb-ldap** component, BZ#1139776

The latest version of the **bind-dyndb-ldap** system plug-in offers significant improvements over the previous versions, but currently has some limitations. One of the limitations is missing support for the LDAP rename (MODRDN) operation. As a consequence, DNS records renamed in LDAP are not served correctly. To work around this problem, restart the **named** daemon to resynchronize data after each MODRDN operation. In an Identity Management (IdM) cluster, restart the **named** daemon on all IdM replicas.

### **ipa** component, BZ#1187524

The **userRoot.ldif** and **ipaca.ldif** files, from which Identity Management (IdM) reimports the back end when restoring from backup, cannot be opened during a full-server restore even though they are present in the tar archive containing the IdM backup. Consequently, these files are skipped during the full-server restore. If you restore from a full-server backup, the restored back end can receive some updates from after the backup was created. This is not expected because all updates received between the time the backup was created and the time the restore is performed should be lost. The server is successfully restored, but can contain invalid data. If the restored server containing invalid data is then used to reinitialize a replica, the replica reinitialization succeeds, but the data on the replica is invalid.

No workaround is currently available. It is recommended that you do not use a server restored from a full-server IdM backup to reinitialize a replica, which ensures that no unexpected updates are present at the end of the restore and reinitialization process.

Note that this known issue relates only to the full-server IdM restore, not to the data-only IdM restore.

### **ipa (slapi-nis)** component, BZ#1157757

When the Schema Compatibility plug-in is configured to provide Active Directory (AD) users access to legacy clients using the Identity Management (IdM) cross-forest trust to AD, the 389 Directory Server can under certain conditions increase CPU consumption upon receiving a request to resolve complex group membership of an AD user.

### **ipa** component, BZ#1186352

When you restore an Identity Management (IdM) server from backup and re-initialize the restored data to other replicas, the Schema Compatibility plug-in can still maintain a cache of the old data from before performing the restore and re-initialization. Consequently, the replicas might behave unexpectedly. For example, if you attempt to add a user that was originally added after performing the backup, and thus removed during the restore and re-initialization steps, the operation might fail with an error, because the Schema Compatibility cache contains a conflicting user entry. To work around this problem, restart the IdM replicas after re-initializing them from the master server. This clears the Schema Compatibility cache and ensures that the replicas behave as expected in the described situation.

### **ipa** component, BZ#1188195

Both anonymous and authenticated users lose the default permission to read the **facsimiletelephonenumber** user attribute after upgrading to the Red Hat Enterprise Linux 7.1 version of Identity Management (IdM). To manually change the new default setting and make the attribute readable again, run the following command:

```
ipa permission-mod 'System: Read User Addressbook Attributes' --includedattrs facsimiletelephonenumber
```



-

**ipa component, BZ#1189034**

The **ipa host-del --updatedns** command does not update the host DNS records if the DNS zone of the host is not fully qualified. Creating unqualified zones was possible in Red Hat Enterprise Linux 7.0 and 6. If you execute **ipa host-del --updatedns** on an unqualified DNS zone, for example, *example.test* instead of the fully qualified *example.test.* with the dot (.) at the end, the command fails with an internal error and deletes the host but not its DNS records. To work around this problem, execute **ipa host-del --updatedns** command on an IdM server running Red Hat Enterprise Linux 7.0 or 6, where updating the host DNS records works as expected, or update the host DNS records manually after running the command on Red Hat Enterprise Linux 7.1.

**ipa component, BZ#1193578**

Kerberos libraries on Identity Management (IdM) clients communicate by default over the User Datagram Protocol (UDP). Using a one-time password (OTP) can cause additional delay and breach of Kerberos timeouts. As a consequence, the **kinit** command and other Kerberos operations can report communication errors, and the user can get locked out. To work around this problem, make communication using the slightly slower Transmission Control Protocol (TCP) default by setting the **udp\_preference\_limit** option to **0** in the **/etc/krb5.conf** file.

**ipa component, BZ#1170770**

Hosts enrolled to IdM cannot belong to the same DNS domains as the DNS domains belonging to an AD forest. When any of the DNS domains in an Active Directory (AD) forest are marked as belonging to the Identity Management (IdM) realm, cross-forest trust with AD does not work even though the trust status reports success. To work around this problem, use DNS domains separate from an existing AD forest to deploy IdM.

If you are already using the same DNS domains for both AD and IdM, first run the **ipa realmdomains-show** command to display the list of IdM realm domains. Then remove the DNS domains belonging to AD from the list by running the **ipa realmdomains-mod --del-domain=wrong.domain** command. Un-enroll the hosts from the AD forest DNS domains from IdM, and choose DNS names that are not in conflict with the AD forest DNS domains for these hosts. Finally, refresh the status of the cross-forest trust to the AD forest by reestablishing the trust with the **ipa trust-add** command.

**ipa component, BZ#988473**

Access control to Lightweight Directory Access Protocol (LDAP) objects representing trust with Active Directory (AD) is given to the **Trusted Admins** group in Identity Management (IdM). In order to establish the trust, the IdM administrator should belong to a group which is a member of the **Trusted Admins** group and this group should have relative identifier (RID) 512 assigned. To ensure this, run the **ipa-adtrust-install** command and then the **ipa group-show admins --all** command to verify that the **ipantsecurityidentifier** field contains a value ending with the **-512** string. If the field does not end with **-512**, use the **ipa group-mod admins --setattr=ipantsecurityidentifier=SID** command, where *SID* is the value of the field from the **ipa group-show admins --all** command output with the last component value (-XXXX) replaced by the **-512** string.

**sssd component, BZ#1024744**

The OpenLDAP server and the 389 Directory Server (389 DS) treat grace logins differently. 389 DS treats them as the number of grace logins *left*, while OpenLDAP treats them as the number of grace logins *used*. Currently, SSSD only handles the semantics used by 389 DS. As a result, when using OpenLDAP, the grace password warning can be incorrect.

**sssd component, BZ#1081046**

The **accountExpires** attribute that SSSD uses to see whether an account has expired is not replicated to the global catalog by default. As a result, users with expired accounts can be allowed to log in when using GSSAPI authentication. To work around this problem, the global catalog support can be disabled by specifying **ad\_enable\_gc=False** in the **sssd.conf** file. With this setting, users with expired accounts will be denied access when using GSSAPI authentication. Note that SSSD connects to each LDAP server individually in this scenario, which can increase the connection count.

#### **sssd component, BZ#1103249**

Under certain circumstances, the algorithm in the Privilege Attribute Certificate (PAC) responder component of the SSSD service does not effectively handle users who are members of a large number of groups. As a consequence, logging from Windows clients to Red Hat Enterprise Linux clients with Kerberos single sign-on (SSO) can be noticeably slow. There is currently no known workaround available.

#### **sssd component, BZ#1194345**

The SSSD service uses the global catalog (GC) for initgroup lookups but the POSIX attributes, such as the user home directory or shell, are not replicated to the GC set by default. Consequently, when SSSD requests the POSIX attributes during SSSD lookups, SSSD incorrectly considers the attributes to be removed from the server, because they are not present in the GC, and removes them from the SSSD cache as well.

To work around this problem, either disable the GC support by setting the **ad\_enable\_gc=False** parameter in the **sssd-ad.conf** file, or replicate the POSIX attributes to the GC. Disabling the GC support is easier but results in the client being unable to resolve cross-domain group memberships. Replicating POSIX attributes to the GC is a more systematic solution but requires changing the Active Directory (AD) schema. As a result of either one of the aforementioned workarounds, running the **getent passwd user** command shows the POSIX attributes. Note that running the **id user** command might not show the POSIX attributes even if they are set properly.

#### **samba component, BZ#1186403**

Binaries in the **samba-common.x86\_64** and **samba-common.i686** packages contain the same file paths but differ in their contents. As a consequence, the packages cannot be installed together, because the RPM database forbids this scenario.

To work around this problem, do not install **samba-common.i686** if you primarily need **samba-common.x86\_64**; neither in a kickstart file, nor on an already installed system. If you need **samba-common.i686**, avoid **samba-common.x86\_64**. As a result, the system can be installed, but with only one architecture of the **samba-common** package at a time.

## 第 41 章 ENTITLEMENT

**subscription-manager** component, [BZ#1189006](#)

The **Save** button in the **Proxy Configuration** dialog is available only in English. When **Proxy Configuration** is displayed in a different language, the **Save** button is always rendered in English.

## 第 42 章 DESKTOP

### **spice component, BZ#1030024**

Video playback on a Red Hat Enterprise Linux 7.1 guest with GNOME Shell is sometimes not detected as a video stream by **spice-server**. The video stream is therefore not compressed in such a case.

### **gobject-introspection component, BZ#1076414**

The **gobject-introspection** library is not available in a 32-bit multilib package. Users who wish to compile 32-bit applications that rely on GObject introspection or libraries that use it, such as **GTK+** or **GLib**, should use the mock package to set up a build environment for their applications.

### **kernel component, BZ#1183631**

Due to a bug, the X.Org X server running on a Lenovo T440s laptop crashes if the laptop is removed from a docking station while an external monitor is attached. All applications running in the GUI are terminated, which leads to potential loss of unsaved data. To work around this problem, detach the laptop from the docking station while the laptop's lid is closed, or unplug all monitors from the docking station first.

### **firefox component, BZ#1162691**

The **icedtea-web** Java plugin does not load in Firefox when running on Red Hat Enterprise Linux for POWER, little endian, architecture. Consequently, Java Web Start (javaws) does not work in this environment. Firefox supports NPAPI plugins for Intel P6, AMD64 and Intel 64 systems, PowerPC platform (32bit), and ARM architectures. All other architectures are not supported by Firefox at the moment and there is no plan to extend it.

---

## 附录 A. 修订历史

<b>修订 1.0-27</b>	<b>Mon Oct 30 2017</b>	<b>Lenka Špačková</b>
Added information on changes in the <b>ld</b> linker behavior to Deprecated Functionality.		
<b>修订 1.0-26</b>	<b>Mon Aug 01 2016</b>	<b>Lenka Špačková</b>
Added a note about limited support for Windows guest virtual machines to Deprecated Functionality.		
<b>修订 1.0-25</b>	<b>Fri Jun 03 2016</b>	<b>Lenka Špačková</b>
Added an OpenSSL known issue, which is valid for all Red Hat Enterprise Linux 7 minor releases.		
<b>修订 1.0-24</b>	<b>Thu May 26 2016</b>	<b>Lenka Špačková</b>
Fixed the "Intel Broadwell Processor and Graphics Support" note in the Hardware Enablement chapter.		
<b>修订 1.0-22</b>	<b>Wed Apr 20 2016</b>	<b>Jiří Herrmann</b>
Added a deprecated Virtualization feature.		
<b>修订 1.0-21</b>	<b>Wed Oct 14 2015</b>	<b>Lenka Špačková</b>
Fix in the Architectures chapter.		
<b>修订 1.0-20</b>	<b>Mon May 04 2015</b>	<b>Radek Bíba</b>
Update of the Red Hat Enterprise Linux 7.1 Release Notes.		
<b>修订 1.0-13</b>	<b>Tue Mar 03 2015</b>	<b>Milan Navrátil</b>
发布 Red Hat Enterprise Linux 7.1 发行注记。		
<b>修订 1.0-9.10</b>	<b>Wed Jan 29 2015</b>	<b>Leah Liu</b>
Chinese (simplified) translation of the Red Hat Enterprise Linux 7.1 Release Notes.		