



Red Hat Enterprise Linux

7

7.1 リリースノート

Red Hat Enterprise Linux 7 リリースノート

Red Hat Customer Content
Services

Red Hat Enterprise Linux 7 7.1 リリースノート

Red Hat Enterprise Linux 7 リリースノート

Red Hat Customer Content Services

法律上の通知

Copyright © 2015 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, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

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

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

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

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

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

The OpenStack ® 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.

概要

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. 謝辞 Red Hat グローバルサポートサービスは Red Hat Enterprise Linux 7 のテストにおいて Sterling Alexander 氏および Michael Everette 氏から多大なるご協力をいただきました。

目次

前書き	5
パート I. 新機能	6
第1章 アーキテクチャー	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 インストールと起動	9
3.1. インストーラー	9
3.2. ブートローダー	13
第4章 ストレージ	14
LVM キャッシュ	14
libStorageMgmt API を使用したストレージアレイ管理	14
LSI Syncro のサポート	14
DIF/DIX サポート	15
device-mapper-multipath の構文エラーチェックおよび出力に関する改善	15
第5章 ファイルシステム	16
Btrfs ファイルシステムのサポート	16
OverlayFS	16
Parallel NFS のサポート	16
第6章 カーネル	17
Ceph ブロックデバイスのサポート	17
並列フラッシュ MCL 更新	17
動的なカーネルのパッチ適用	17
複数 CPU での Crashkernel のサポート	17
dm-era ターゲット	17
Cisco VIC カーネルドライバ	17
hwrng でのエン트로ピー管理機能の強化	17
スケジューラーの負荷分散パフォーマンスの改善	18
スケジューラーの newidle 分散の強化	18
ノードごと 1 GB のページ割り当てを行う HugeTLB のサポート	18
MCS ベースの新しいロックメカニズム	18
8KB から 16KB にプロセススタックサイズを増大	18
perf および systemtap での uprobe と uretprobe 機能の有効化	18
エンドツーエンドのデータ整合性チェック	18
32 ビットシステムでの DRBG	18
NFSv4.1 RDMA Available	18
Crashkernel の大型サイズに対するサポート	19
Kdump Supported on Secure Boot Machines	19
Firmware-assisted Crash Dumping	19
Runtime Instrumentation for IBM System z	19
Cisco usNIC Driver	19
Intel Ethernet Server Adapter X710/XL710 Driver Update	19

第7章 仮想化	20
KVM での vCPU 最大数の増加	20
QEMU、KVM、libvirt API における 5 世代 Intel Core の新命令のサポート	20
KVM ゲスト向け USB 3.0 サポート	20
dump-guest-memory コマンドの圧縮	20
Open Virtual Machine Firmware	20
Hyper-V でのネットワークパフォーマンス性が向上	20
hyperv-daemons の hypervfsd	20
libguestfs の新機能	20
フライトレコーダートレース機能	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
第8章 クラスタリング	22
Corosync の動的なトークンタイムアウト	22
Corosync のタイブレーカーの強化	22
Red Hat High Availability に関する強化	22
第9章 コンパイラーとツール	23
Linux on System z バイナリー向けパッチのホット適用のサポート	23
Performance Application Programming Interface の強化	23
OProfile	23
OpenJDK8	23
sosreport が snap の代替となる	23
リトルエンディアン 64 ビット PowerPC 向け GDB サポート	23
Tuna の機能強化	24
crash Moved to Debugging Tools	24
Accurate ethtool Output	24
Concerns Regarding Transactional Synchronization Extensions	24
第10章 ネットワーキング	25
Trusted Network Connect	25
qlcnic ドライバーの SR-IOV 機能	25
Berkeley Packet Filter	25
クロックの安定性の強化	25
libnetfilter_queue パッケージ	25
Team 機能の強化	25
Intel QuickAssist Technology ドライバー	25
PTP と NTP 間でのフェールオーバーの LinuxPTP timemaster サポート	26
ネットワーク初期スクリプト	26
TCP 遅延 ACK	26
NetworkManager	26
ネットワーク名前空間と VTI	27
MemberOf プラグイン用の代替の設定ストレージ	27
第11章 Red Hat Enterprise Linux Atomic Host	28
New features in Red Hat Enterprise Linux Atomic Host 7.1.4	28
New features in Red Hat Enterprise Linux Atomic Host 7.1.3	29
New features in Red Hat Enterprise Linux Atomic Host 7.1.2	29
第12章 Linux Containers	31
12.1. Linux Containers Using Docker Technology	31
12.2. Container Orchestration	34
12.3. Cockpit Enablement	36

12.3. Cockpit Enablement	35
12.4. Containers Using the libvirt-lxc Tooling Have Been Deprecated	36
第13章 認証と相互運用性	37
Manual Backup and Restore Functionality	37
WinSync から Trust への移行サポート	37
One-Time Password Authentication	37
Common Internet File System 向け SSSD 統合	37
認証局管理ツール	37
アクセス制御の柔軟な設定	37
特権を持たないユーザーに対するドメインへのアクセスの制限	37
自動データプロバイダー設定	38
AD および LDAP の sudo プロバイダーの使い方	38
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
第14章 セキュリティ	39
SCAP セキュリティガイド	39
SELinux ポリシー	39
OpenSSH の新機能	39
Libreswan の新機能	39
TNC の新機能	40
GnuTLS の新機能	40
第15章 デスクトップ	42
Mozilla Thunderbird	42
Quad-buffered OpenGL ステレオビジュアルのサポート	42
オンラインアカウントプロバイダー	42
第16章 サポート対応とメンテナンス	43
ABRT 承認マイクロレポート機能	43
第17章 Red Hat Software Collections	44
第18章 Red Hat Enterprise Linux for Real Time	45
パート II. Technology Previews	46
第19章 Hardware Enablement	47
第20章 Storage	48
第21章 File Systems	49
第22章 Kernel	50
第23章 Virtualization	51
第24章 Compiler and Tools	52
第25章 Networking	53
第26章 Authentication and Interoperability	54
パート III. デバイスドライバ	55
第27章 ストレージドライバの更新	56
第28章 ネットワークドライバの更新	57

第28章 ネットワークファイバーの更新	57
第29章 グラフィックドライバーの更新	58
パート IV. Deprecated Functionality	59
第30章 Deprecated Functionality in Red Hat Enterprise Linux 7	60
Windows guest virtual machine support limited	60
パート V. Known Issues	61
第31章 Installation and Booting	62
第32章 Storage	68
第33章 File Systems	69
第34章 Virtualization	70
第35章 Deployment and Tools	72
第36章 Compiler and Tools	73
第37章 Networking	74
第38章 Red Hat Enterprise Linux Atomic Host	75
第39章 Linux Containers	76
第40章 Authentication and Interoperability	78
第41章 Entitlement	81
第42章 Desktop	82
付録A 改訂履歴	83

前書き

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/>.

パート I. 新機能

本パートでは Red Hat Enterprise Linux 7.1 に導入された新機能および主要な機能拡張について記載しています。

第1章 アーキテクチャー

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

- ※ 64 ビット AMD
- ※ 64 ビット Intel
- ※ IBM POWER7+ and POWER8 (big endian)
- ※ IBM POWER8 (リトルエンディアン) [2]
- ※ IBM System z [3]

本リリースでは、サーバーおよびシステムに対する各種の改善の他、オープンソースより得られた技術をまとめて提供しています。

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](#).
- ※ Red Hat Enterprise Linux for POWER (リトルエンディアン) で対応しているのは IBM POWER8 プロセッサベースのサーバーのみになります。
- ※ 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** ブートローダーはインストールメディアおよびネットワーク起動で使用されます。[インストールガイド](#) 内の **GRUB2** を使った IBM Power Systems クライアント用ネットワーク起動サーバーの設定に関する記事を更新しています。
- ※ IBM Power Systems 向けのソフトウェアパッケージはすべて Red Hat Enterprise Linux for POWER のリトルエンディアンおよびビッグエンディアンのいずれにも使用できます。
- ※ Red Hat Enterprise Linux for POWER (リトルエンディアン) 向けにビルドされるパッケージのアーキテクチャーコードは **ppc64le** になります (*gcc-4.8.3-9.ael7b.ppc64le.rpm* など)。

[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 (リトルエンディアン) は現在、**Red Hat Enterprise Virtualization for Power** および **PowerVM** ハイパーバイザー配下の KVM ゲストとしてしか対応していません。

[3] Red Hat Enterprise Linux 7.1 が対応するのは IBM zEnterprise 196 ハードウェアまたはそれ以降になります。IBM System z10 メインフレームのシステムには対応しなくなるため Red Hat Enterprise Linux 7.1 は起動しなくなります。

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 インストールと起動

3.1. インストーラー

Red Hat Enterprise Linux 7.1 のインストールプロセスを改善するため Red Hat Enterprise Linux インストーラー **Anaconda** の機能が強化されています。

インターフェース

- ※ グラフィカルなインストーラーインターフェースに **Kdump** カーネルクラッシュダンプ機能メカニズムの設定を行う画面が追加されます。今まではインストール終了後の **firstboot** ユーティリティで設定することになり、グラフィカルインターフェースがない場合はアクセスできませんでした。今後はグラフィカル環境以外でもシステムへのインストールプロセスの一部として **Kdump** を設定できるようになります。新しい画面はメインのインストーラーメニュー (インストール概要) からアクセスすることができます。

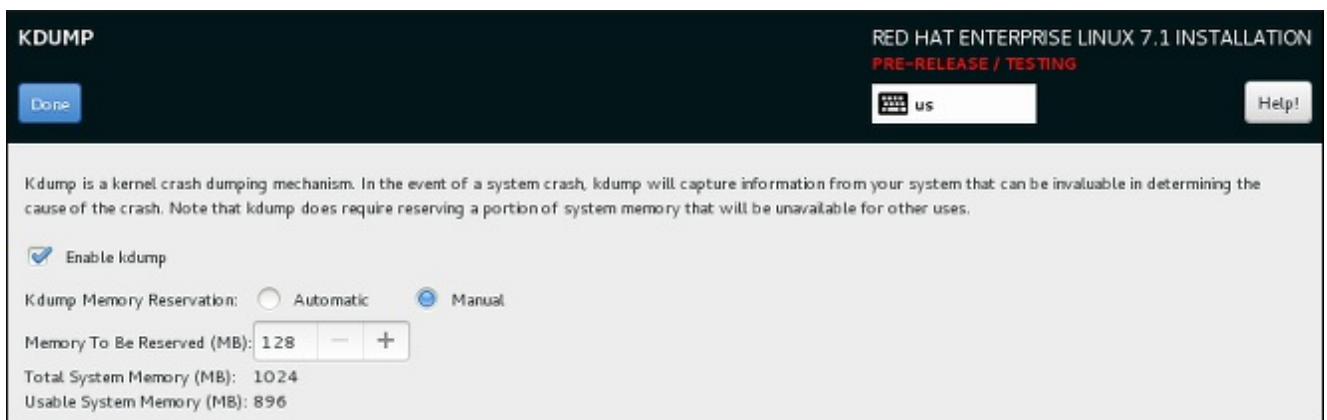


図3.1 新しい Kdump 画面

- ※ 手作業でパーティションの設定を行う画面はより使いやすいようデザインが変更されています。一部コントロールの画面上の位置が移動しています。

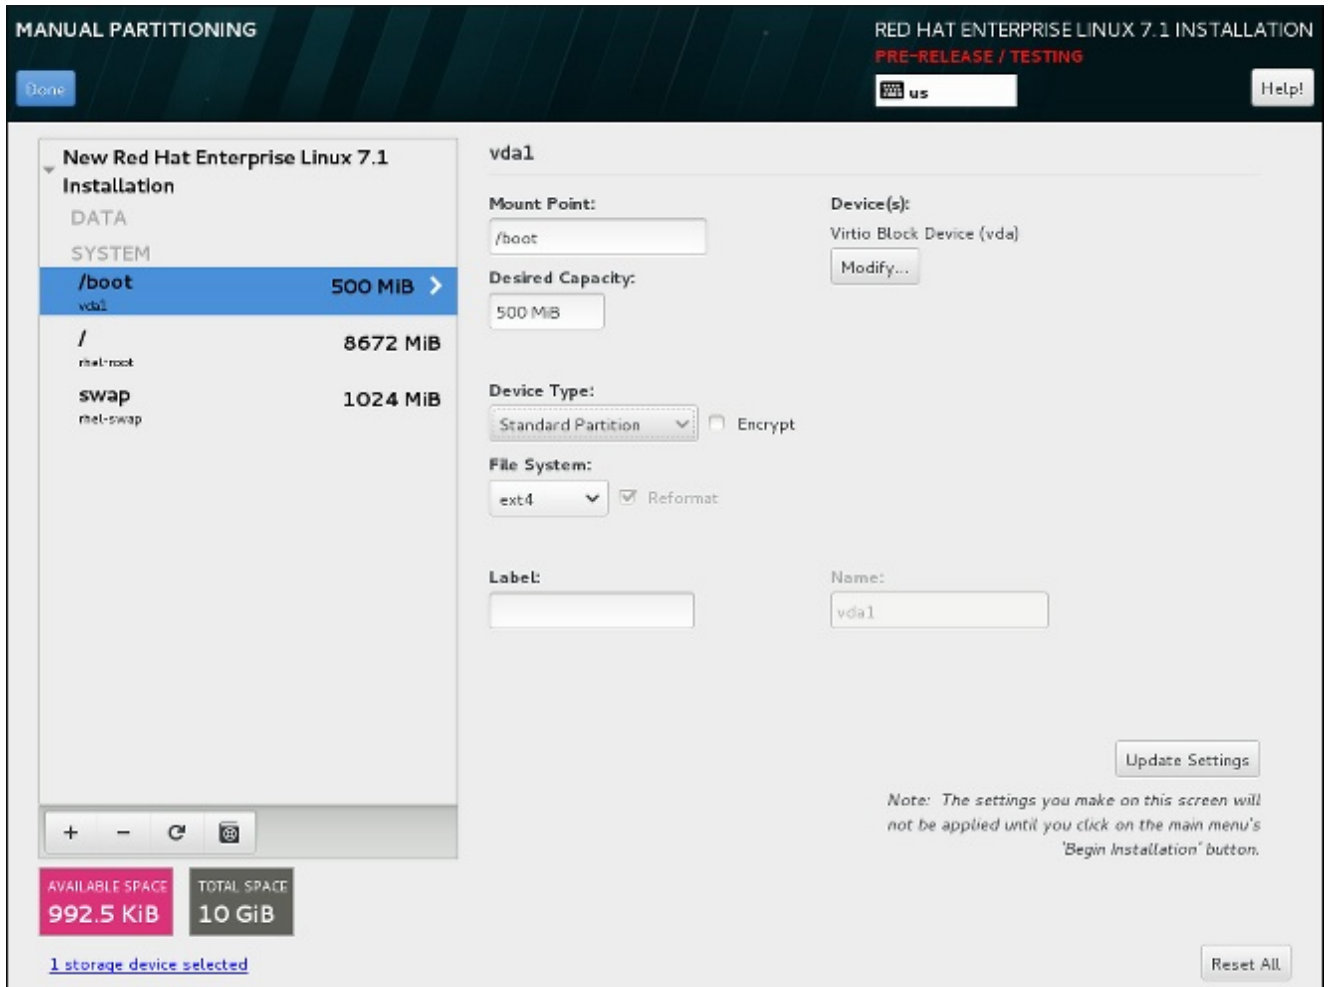


図3.2 変更後の手動パーティション設定画面

- ※ 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』.

ブリッジ設定に関するキックスタートの新オプションがいくつか追加されています。以下に詳細を示します。

- ※ インストーラーではログの表示に複数のコンソールを使用しなくなります。すべてのログは仮想コンソール 1 (**ttty1**) の **tmux** ペインに表示されるようになります。インストール中、ログにアクセスする場合は **Ctrl+Alt+F1** を押して **tmux** に切り替えます。別のウィンドウへの切り替えは **Ctrl+b X** で行います (X には画面下部に表示されるウィンドウ番号を使用)。

グラフィカルなインターフェースに戻る場合は **Ctrl+Alt+F6** を押します。

- ※ **Anaconda** のコマンドラインインターフェースで詳細なヘルプが参照できるようになります。ヘルプを表示させるには **anaconda** パッケージがインストールされているシステムで **anaconda -h** を使用します。コマンドラインインターフェースを使用するとインストール済みシステムでインストーラーを実行させることができるようになるため、ディスクイメージのインストールの際に便利です。

キックスタートコマンドとオプション

- ※ 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.

デフォルトで使用できるプロファイルは **default** と **thin-performance** です。/etc/lvm/profile ディレクトリー内で定義します。詳細については **lvm(8)** の man ページを参照してください。

- ※ 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**.
- ※ **liveimg** コマンドで tar ファイルおよびディスクイメージいずれからのインストールにも対応するようになります。tar アーカイブにはインストールメディアの root ファイルシステムを含ませ、ファイル名の末尾は **.tar**、**.tbz**、**.tgz**、**.txz**、**.tar.bz2**、**.tar.gz**、**.tar.xz** のいずれかにしておく必要があります。
- ※ 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** コマンドに新しいオプション **--fstype** が追加されています。キックスタートファイルで自動パーティション設定を使用する場合、このオプションを使うとデフォルトのファイルシステムタイプを変更することができますようになります (**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.



注記

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](#).

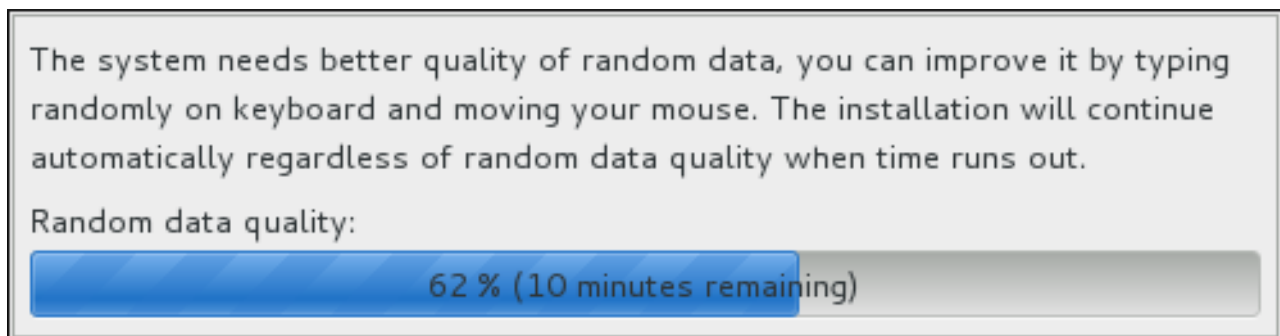


図3.3 Gathering Entropy for Encryption

グラフィカルインストーラーにヘルプがビルトイン

インストーラーのグラフィカルインターフェースおよび **初期設定** ユーティリティー内の各画面の右上に **ヘルプ** ボタンが表示されるようになります。このボタンをクリックすると、**Yelp** ヘルプブラウザを使って現在の画面に適切な [インストールガイド](#) のセクションが表示されます。




図3.4 Anaconda built-in help

3.2. ブートローダー

IBM Power Systems 用のインストールメディアでは以前まで使用していた **yaboot** ブートローダーではなく **GRUB2** ブートローダーを使用するようになります。Red Hat Enterprise Linux for POWER のビッグエンディアンの場合、推奨しているのは **GRUB2** ですが **yaboot** を使用することもできます。新たに導入されるリトルエンディアンの場合は **GRUB2** で起動する必要があります。

[インストールガイド](#) 内の **GRUB2** を使った IBM Power Systems 用ネットワーク起動サーバーの設定に関する記事を更新しています。

第4章 ストレージ

LVM キャッシュ

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 **lvm(7)** manual page for information on creating cache logical volumes.

キャッシュ論理ボリュームの使用に際して次のような制約があるので注意してください。

- ※ キャッシュ論理ボリュームは最上位レベルのデバイスにしてください。シンプールの論理ボリューム、RAID 論理ボリュームのイメージ、その他サブ論理ボリュームタイプなどには使用できません。
- ※ The cache LV sub-LVs (the origin LV, metadata LV, and data LV) can only be of linear, stripe, or RAID type.
- ※ キャッシュ論理ボリュームのプロパティの作成後の変更はできません。キャッシュのプロパティを変更する場合は、キャッシュを削除してから適切なプロパティでキャッシュを作成しなおしてください。

libStorageMgmt API を使用したストレージアレイ管理

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 Filer (ontap 7-Mode)
- ※ Nexenta (nstor 3.1.x のみ)
- ※ SMI-S (ベンダーは以下の通り)
 - HP 3PAR
 - OS リリース 3.2.1 またはそれ以降
 - EMC VMAX および VNX
 - Solutions Enabler V7.6.2.48 またはそれ以降
 - SMI-S Provider V4.6.2.18 ホットフィックスキットまたはそれ以降
 - HDS VSP Array 非組み込みプロバイダー
 - Hitachi Command Suite v8.0 またはそれ以降

libStorageMgmt の詳細については [ストレージ管理ガイドで該当する章](#) を参照してください。

LSI Syncro のサポート

Red Hat Enterprise Linux 7.1 では LSI Syncro CS の HA-DAS (high-availability direct-attached storage) アダプターを有効にするため **megaraid_sas** ドライバーにコードを含ませています。**megaraid_sas** ドライバーはこれまでの有効なアダプターに対しては完全対応になりますが、Syncro CS に使用する場合はテクノロジープレビューとしてご利用いただくことになります。本アダプターのサポートについては LSI、システ

ムインテグレーターまたはシステムのベンダーより直接提供されることになります。Red Hat Enterprise Linux 7.1 に Syncro CS を導入されるお客様にはぜひ Red Hat および LSI へフィードバックを頂けるようご協力をお願いします。LSI Syncro CS ソリューションの詳細については <http://www.lsi.com/products/shared-das/pages/default.aspx> を参照してください。

DIF/DIX サポート

DIF/DIX が SCSI 標準に新たに追加されました。Red Hat Enterprise Linux 7.1 ではテクノロジープレビューになります。DIF/DIX により一般的に使用されている 512 バイトのディスクブロックのサイズが 512 バイトから DIF (Data Integrity Field) を追加した 520 バイトに増えます。この追加された DIF で書き込み時に HBA (ホストバスアダプター) で計算されるデータブロックのチェックサム値を格納します。ストレージデバイスはデータを受け取る時にこのチェックサムを確認してからデータとチェックサムの両方を保存します。読み取り時は逆にストレージデバイスおよび受取側の HBA にチェックサムを検証させることができます。

詳細については [Storage Administration Guide \(ストレージ管理ガイド\)](#) で DIF/DIX を有効にした場合のブロックデバイスに関するセクションを参照してください。

device-mapper-multipath の構文エラーチェックおよび出力に関する改善

`device-mapper-multipath` ツールは `multipath.conf` ファイルの検証に関する信頼性を高めるため改善が行われています。このため、`multipath.conf` ファイルに解析できない行が含まれていると `device-mapper-multipath` によりエラーが報告され、誤った解析を避けるため解析できない行を無視するようになります。

また、`multipathd show paths format` コマンドに以下のようなワイルドカードが追加されています。

- ※ ホスト名および目的のファイバーチャンネルワールドワイドノード名に `%N` と `%n` が追加されています。
- ※ ホスト名および目的のファイバーチャンネルワールドワイドポート名に `%R` と `%r` が追加されています。

特定ファイバーチャンネルのホストやターゲット、そのポートとマルチパスとの関連付けがより容易になり、ストレージ構成に関する管理がより効率的に行えるようになります。

第5章 ファイルシステム

Btrfs ファイルシステムのサポート

Red Hat Enterprise Linux 7.1 では**Btrfs** (B-Tree) ファイルシステムはテクノロジープレビューとしての対応になります。このファイルシステムでは高度な管理や信頼性、拡張性を提供しています。こうした機能によりユーザーによるスナップショットの作成、圧縮や統合デバイス管理を可能にします。

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 **xf**s and **gfs2** file systems is not supported.
- ✦ SELinux is not supported, and to use OverlayFS, it is required to disable enforcing mode.

Parallel NFS のサポート

Parallel NFS (pNFS) は NFS 4.1 標準の一部になります。複数クライアントによるストレージデバイスへの直接アクセスおよび並列アクセスを可能にします。pNFS アーキテクチャーにより NFS サーバーの一般的な負荷の一部に関して拡張性やパフォーマンス性を向上させることができます。

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 では、将来的には新しい pNFS レイアウトのタイプを認定し、完全サポートできるようパートナーおよびオープンソースプロジェクトと共に取り組み続けていきます。

pNFS の詳細については <http://www.pnfs.com/> を参照してください。

第6章 カーネル

Ceph ブロックデバイスのサポート

Red Hat Enterprise Linux 7.1 のカーネルに **libceph.ko** と **rbd.ko** のモジュールが追加されています。この RBD カーネルモジュールにより Linux ホストが Ceph ブロックデバイスを通常のディスクデバイスエントリとして認識できるようになるため、ディレクトリーにマウントして **XFS** や **ext4** などの標準ファイルシステムでフォーマット化することができるようになります。

Red Hat Enterprise Linux 7.1 では現在、CephFS モジュールの **ceph.ko** には対応していないため注意してください。

並列フラッシュ MCL 更新

IBM System z アーキテクチャーの Red Hat Enterprise Linux 7.1 ではマイクロコードレベルの更新 (MCL) が有効になります。更新はフラッシュストレージメディアに対する I/O 動作に影響を与えることなく適用でき、変更されたフラッシュハードウェアサービスレベルについてユーザーに通知を行います。

動的なカーネルのパッチ適用

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.

複数 CPU での Crashkernel のサポート

Red Hat Enterprise Linux 7.1 では複数 CPU を使用する場合、**crashkernel** を起動できるようになります。この機能はテクノロジープレビューとしての対応になります。

dm-era ターゲット

Red Hat Enterprise Linux 7.1 では dm-era デバイスマッパーのターゲットをテクノロジープレビューとして導入しています。「era」と呼ばれるユーザー定義の期間内に書き込みが行われたブロックの追跡を行います。それぞれの era ターゲットインスタンスでは現在の era は単調増加する 32 ビットカウンターとして管理されます。このターゲットによりバックアップソフトウェアが最後のバックアップ後に変更が行われたブロックの追跡を行えるようになります。また、キャッシュコンテンツの部分的な無効化を行うことでベンダーのスナップショットにロールバックした際のキャッシュの一貫性を復元することもできます。dm-era は主に dm-cache ターゲットと組み合わせた使用が求められます。

Cisco VIC カーネルドライバー

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 でのエントロピー管理機能の強化

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.

スケジューラーの負荷分散パフォーマンスの改善

今まではスケジューラーの負荷分散コードは待機中の全 CPU に対して負荷分散を行っていました。Red Hat Enterprise Linux 7.1 では待機中のすべての CPU ではなく、負荷分散されるべき期限に達している待機中の CPU のみに負荷分散を行います。これにより待機していない CPU での負荷分散率を低減、これに伴いスケジューラーの不必要な作業量が減るためパフォーマンスが向上することになります。

スケジューラーの **newidle** 分散の強化

実行可能なタスクがある場合、**newidle** 分散コード内でのタスクの検索を停止するようスケジューラーの動作に変更が加えられたため、パフォーマンス性が向上されます。

ノードごと 1 GB のページ割り当てを行う HugeTLB のサポート

Red Hat Enterprise Linux 7.1 ではランタイム時の膨大なページ割り当てに対するサポートを追加しています。これにより 1GB **hugetlbfs** のユーザーがランタイム時に 1 GB を割り当てるべき NUMA (Non-Uniform Memory Access) ノードを指定できるようになります。

MCS ベースの新しいロックメカニズム

Red Hat Enterprise Linux 7.1 では新しいロックメカニズム MCS ロックを導入します。大規模なシステムでの **spinlock** のオーバーヘッドを大幅に低減させ、Red Hat Enterprise Linux 7.1 での **spinlocks** の効率性を全般的に向上させます。

8KB から 16KB にプロセススタックサイズを増大

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.

perf および systemtap での uprobe と uretprobe 機能の有効化

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

エンドツーエンドのデータ整合性チェック

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 ビットシステムでの DRBG

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

NFSv4.1 Available

As a Technology Preview, the NFSv4.1 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.

Crashkernel の大型サイズに対するサポート

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 仮想化

KVM での vCPU 最大数の増加

KVM ゲストでの仮想 CPU (vCPU) の最大対応数が 240 に増えています。これによりユーザーがゲストに割り当てることができる仮想処理ユニットの量が増大するため、パフォーマンス性が向上します。

QEMU、KVM、libvirt API における 5 世代 Intel Core の新命令のサポート

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 ゲスト向け USB 3.0 サポート

Red Hat Enterprise Linux 7.1 では USB 3.0 のホストアダプター (xHCI) エミュレーションをテクノロジープレビューとして追加することで USB サポートの強化を図っています。

dump-guest-memory コマンドの圧縮

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

Red Hat Enterprise Linux 7.1 では Open Virtual Machine Firmware (OVMF) がテクノロジープレビューとして利用できます。OVMF とは AMD64 および Intel 64 のゲスト向けの UEFI セキュアブート環境です。

Hyper-V でのネットワークパフォーマンス性が向上

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 の hypervfcopyd

hypervfcopyd デーモンが *hyperv-daemons* パッケージに追加されています。**hypervfcopyd** は Hyper-V 2012 R2 ホストで稼働させる Linux ゲストのファイルコピーサービス機能の実装です。これによりホストから Linux ゲストにファイルをコピーすることができるようになります (VMBUS 経由)。

libguestfs の新機能

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.

フライトレコーダートレース機能

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.

フライトレコーダートレース機能の設定方法および使い方については [Virtualization Deployment and Administration Guide](#) を参照してください。

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 クラスタリング

Corosync の動的なトークンタイムアウト

Corosync Cluster Engine に **token_coefficient** オプションが追加されています。 **token_coefficient** の値の使用は **nodelist** セクションが指定され少なくとも 3 ノードが含まれている場合に限りです。このような状況の場合、トークンのタイムアウトは次のように計算されます。

```
[トークン + (ノード数 - 2)] * token_coefficient
```

新しいノードが追加される度に手作業でトークンのタイムアウトを変更しなくてもクラスタで測定できるようになります。デフォルト値は 650 ミリ秒ですが、0 秒に設定してこの機能を効率的に取り除くこともできます。

Corosync でノードの動的な追加や削除が処理できるようになります。

Corosync のタイブレーカーの強化

Corosync の **auto_tie_breaker** 定足数機能でタイブレーカーノードの設定や変更がより柔軟に行えるようオプションが提供されています。クラスタの同数割れが発生した場合に定足数を維持するノードの一覧を選択するか、定足数を維持するノードの ID が最小の ID か最大の ID かを選択することができるようになります。

Red Hat High Availability に関する強化

Red Hat Enterprise Linux 7.1 リリースの場合、**Red Hat High Availability Add-On** では次の機能に対応します。各機能の詳細については『High Availability Add-On Reference』を参照してください。

- ※ **pcs resource cleanup** コマンドで全リソースの **failcount** およびリソースの状態をリセットできるようになります。
- ※ **pcs resource move** コマンドの **lifetime** パラメーターを使って、このコマンドで作成されるリソースの制限の有効期間を指定することができます。
- ※ **pcs acl** コマンドを使って、クラスタ設定への読み取り専用アクセスまたは読み取りおよび書き込みアクセスをアクセス制御一覧 (ACL) でローカルユーザーに許可することができます。
- ※ **pcs constraint** コマンドでは汎用リソースオプションに加え、特定の制約オプションの設定についても対応するようになります。
- ※ **pcs resource create** コマンドでは作成したリソースが自動的に起動しないよう **disabled** パラメーターに対応します。
- ※ **pcs cluster quorum unblock** コマンドでは定足数の確立時にクラスタがすべてのノードを待機しないようにします。
- ※ **pcs resource create** コマンドの **before** と **after** パラメーターを使ってリソースグループの順番を設定することができます。
- ※ **pcs config** コマンドの **backup** と **restore** オプションを使ってクラスタの設定を tarball にバックアップしたり、バックアップから全ノードに設定ファイルを復元したりすることができます。

第9章 コンパイラーとツール

Linux on System z バイナリー向けパッチのホット適用のサポート

GCC (GNU Compiler Collection) では Linux on System z バイナリー向けにマルチスレッド化したコードのオンラインパッチ適用に関するサポートを実装しています。パッチのホット適用に特定の機能を選択して有効にする場合は「機能の属性」を使用します。全機能のパッチホット適用を有効にする場合は `-mhotpatch` コマンドラインオプションを使用します。

パッチのホット適用を有効にするとソフトウェアのサイズとパフォーマンスに悪影響を与えます。このため、パッチホット適用のサポートは特定の機能に限定し全機能に対して行わないことを推奨しています。

Linux on System z バイナリー向けのパッチホット適用のサポートは Red Hat Enterprise Linux 7.0 のテクノロジープレビューでした。Red Hat Enterprise Linux 7.1 リリースでは完全対応になります。

Performance Application Programming Interface の強化

Red Hat Enterprise Linux 7 には **Performance Application Programming Interface (PAPI)** が収録されています。PAPI は最近のマイクロプロセッサに搭載されているハードウェアパフォーマンスカウンターに対するクロスプラットフォームインターフェースの仕様になります。カウンターは複数レジスターの小さな集合体の形で存在しイベントをカウントします。イベントとはプロセッサの機能に関連して発生する特定の信号を指します。こうしたイベントを監視してアプリケーションのパフォーマンス分析やチューニングに利用します。

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.

sosreport が snap の代替となる

廃止予定だった **snap** ツールが *powerpc-utils* パッケージから削除されました。snap の機能は **sosreport** ツールに統合されています。

リトルエンディアン 64 ビット PowerPC 向け GDB サポート

Red Hat Enterprise Linux 7.1 では GDB (GNU Debugger) での 64 ビット PowerPC リトルエンディアンアーキテクチャー向けサポートを実装します。

Tuna の機能強化

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 ネットワーキング

Trusted Network Connect

Red Hat Enterprise Linux 7.1 では Trusted Network Connect をテクノロジープレビューとして導入しています。Trusted Network Connect は TLS、802.1X、IPsec など既存のネットワークアクセス制御 (NAC) ソリューションと併用することで endpoint posture assessment を統合します。つまり、エンドポイントのシステム情報を収集します (オペレーティングシステムを構成している設定、インストールしているパッケージ、その他、整合性測定と呼ばれている)。エンドポイントがネットワークにアクセスする前に、Trusted Network Connect を使用してこうした測定をネットワークアクセスポリシーに対して検証します。

qlcnic ドライバーの SR-IOV 機能

SR-IOV (Single-Root I/O virtualization) のサポートが **qlcnic** ドライバーにテクノロジープレビューとして追加されています。この機能のサポートは QLogic から直接提供されます。QLogic および Red Hat へのフィードバックをぜひお願いします。qlcnic ドライバー内の他の機能は引き続き完全対応になります。

Berkeley Packet Filter

トラフィック分類技術をベースとした BPF (Berkeley Packet Filter) のサポートが Red Hat Enterprise Linux 7.1 に追加されています。BPF はパケットソケットのパケットフィルタリング、セキュアコンピューティングモード (seccomp) でのサンドボックス機能のパケットフィルタリング、Netfilter 内のパケットフィルタリングで使用されます。BPF には最も重要なアーキテクチャー向けの just-in-time 実装およびフィルター構築用のリッチ構文が備わっています。

クロックの安定性の強化

以前はティックレスカーネル機能を無効にするとシステムクロックの安定性が大幅に改善されることがテスト結果で示されていました。このカーネルティックレスモードはカーネルの起動オプションパラメーターに **nohz=off** を追加すると無効にすることができます。ただし、最近、Red Hat Enterprise Linux 7.1 のカーネルに適用された改善によりシステムクロックの安定性が大幅に改善され、ほとんどの使用状態で **nohz=off** を付けた場合と付けない場合でのクロックの安定性における違いは以前に比べかなり小さくなっているはずです。ティックレスカーネルは **PTP** や **NTP** を使った時間同期を適用する場合に便利な機能になります。

libnetfilter_queue パッケージ

libnetfilter_queue パッケージが Red Hat Enterprise Linux 7.1 に追加されています。**libnetfilter_queue** はユーザースペースのライブラリーでカーネルのパケットフィルターによってキュー待ちに置かれているパケットに対して API を提供します。これによりカーネルの **nfnetlink_queue** サブシステムからのキュー待ちパケットの受信、パケットの解析、パケットヘッダーのリライト、変更されたパケットの再インジェクトなどができるようになります。

Team 機能の強化

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 ドライバー

Red Hat Enterprise Linux 7.1 に QAT (Intel QuickAssist Technology) ドライバーが追加されています。システムにハードウェアのオフロード暗号化機能を追加する QuickAssist ハードウェアは QAT ドライバーで有効になります。

PTP と NTP 間でのフェールオーバーの LinuxPTP timemaster サポート

Red Hat Enterprise Linux 7.1 では *linuxptp* パッケージがバージョン **1.4** に更新されています。多くのバグ修正や機能強化が施され、特に **timemaster** アプリケーションを使った **PTP** ドメインと **NTP** ソース間のフェールオーバーに対応するようになっています。ネットワーク上に複数の **PTP** ドメインがある場合、または **NTP** へのフォールバックが必要な場合、**timemaster** プログラムを使用するとシステムクロックを利用可能なすべてのタイムソースに対して同期させることができます。

ネットワーク初期スクリプト

Red Hat Enterprise Linux 7.1 では VLAN のカスタム名のサポートが追加されています。GRE トンネルでの **IPv6** のサポートが改善、追加され、内部アドレスが再起動後も維持されるようになります。

TCP 遅延 ACK

Red Hat Enterprise Linux 7.1 では設定可能な TCP 遅延 ACK のサポートが *iproute* パッケージに追加されています。**ip route quickack** コマンドで有効にすることができます。

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.

Red Hat Enterprise Linux 7.1 でボンディングオプションの **lACP_rate** に対応するようになります。**NetworkManager** が強化されスレーブインターフェースを持つマスターインターフェースの名前を変更する際にデバイス名の変更が容易になります。

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 **nmc11** command-line utility, including the ability to provide passwords when connecting to Wi-Fi or 802.1X networks.

ネットワーク名前空間と VTI

Red Hat Enterprise Linux 7.1 ではネットワーク名前空間を持つ VTI (*virtual tunnel interfaces*) のサポートが追加されています。パケットがカプセル化された場合またはカプセル化が外された場合、VTI からのトラフィックを異なる名前空間で渡し合うことができるようになります。

MemberOf プラグイン用の代替の設定ストレージ

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 **kubect1** 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 認証と相互運用性

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](#).

WinSync から Trust への移行サポート

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](#).

Common Internet File System 向け SSSD 統合

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](#).

認証局管理ツール

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.

アクセス制御の柔軟な設定

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.

特権を持たないユーザーに対するドメインへのアクセスの制限

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](#).

自動データプロバイダー設定

ipa-client-install コマンドではデフォルトで **SSSD** を sudo サービスのデータプロバイダーとして設定するようになります。この動作は **--no-sudo** オプションを使用すると無効にすることができます。また、Identity Management クライアントのインストール用 NIS ドメインを指定する **--nisdomain** オプション、NIS ドメイン名を設定しないようにする **--no_nisdomain** オプションがそれぞれ追加されています。いずれのオプションも使用しないと IPA ドメインが使用されます。

AD および LDAP の sudo プロバイダーの使い方

AD プロバイダーは Active Directory サーバーへの接続に使用するバックエンドです。Red Hat Enterprise Linux 7.1 では AD sudo プロバイダーと LDAP プロバイダーとの併用はテクノロジープレビューとしての対応になります。AD sudo プロバイダーを有効にするには **sudo_provider=ad** 設定を **sssd.conf** ファイルのドメインセクションに追加します。

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 セキュリティ

SCAP セキュリティガイド

Red Hat Enterprise Linux 7.1 に *scap-security-guide* パッケージが収録され安全性に関するガイダンス、安全基準、安全性に関する検証メカニズムなどの説明が提供されています。安全性を強化する多数の実践的なアドバイスから構成されるガイダンスは *Security Content Automation Protocol (SCAP)* 内をご覧ください。**SCAP Security Guide** には規定の安全方針要件に関するシステム安全準拠のスキャンを実施するにあたって必要なデータが含まれています (書面による詳細と自動テスト (プローブ) の両方が含まれる)。**SCAP Security Guide** では、テストの自動化により便利なだけでなく信頼できるシステム準拠の定期的な確認法を提供しています。

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 ポリシー

Red Hat Enterprise Linux 7.1 では SELinux ポリシーが修正されています。これまで独自の SELinux ポリシーを持たず **init_t** ドメインで実行されていたサービスは新たに追加された **unconfined_service_t** ドメインで実行されるようになります。詳細は Red Hat Enterprise Linux 7.1 の [SELinux User's and Administrator's Guide](#) 内の [Unconfined Processes](#) の章を参照してください。

OpenSSH の新機能

OpenSSH ツールセットがバージョン 6.6.1p1 に更新され、暗号に関連する新機能がいくつか追加されています。

- ※ Daniel Bernstein **Curve25519** の elliptic-curve **Diffie-Hellman** を使うキー交換に対応するようになります。このメソッドはデフォルトで与えられ、サーバーおよびクライアントの両方で対応します。
- ※ **Ed25519** elliptic-curve 署名スキームをパブリックキータイプとして使用するためのサポートが追加されています。ユーザーキー、ホストキーの両方に使用できる **Ed25519** では **ECDSA** や **DSA** より高い安全性とパフォーマンスを提供します。
- ※ **bcrypt** キー誘導機能 (KDF) を使用する新しいプライベートキーフォーマットが追加されています。**Ed25519** にはこのフォーマットがデフォルトで使用されますが、別のキータイプを要求することも可能です。
- ※ 新しい転送暗号 **chacha20-poly1305@openssh.com** が追加されています。これにより Daniel Bernstein **ChaCha20** ストリーム暗号と **Poly1305** メッセージ暗証コード (MAC) が結合されます。

Libreswan の新機能

IPsec VPN の **Libreswan** 実装が新機能や機能強化が追加されたバージョン 3.12 に更新されています。

- ※ 新しい暗号が追加されています。

- ※ **IKEv2** support has been improved.
- ※ 中間的証明書チェーンのサポートが **IKEv1** と **IKEv2** に追加されています。
- ※ 接続処理機能に改善が施されています。
- ※ OpenBSD、Cisco、Android などのシステムとの相互運用性が改善されています。
- ※ **systemd** のサポートが改善されています。
- ※ ハッシュ化された **CERTREQ** およびトラフィック統計用のサポートが追加されています。

TNC の新機能

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.
- ※ 新しい TPMRA ワークアイテムの実装により Attestation IMV のサポートが改善されています。
- ※ SWID IMV 搭載の JSON ベースの REST API にサポートが追加されています。
- ※ 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.
- ※ 既存の **IF-TNCCS (PB-TNC、IF-M (PA-TNC))** プロトコルおよび **OS IMC/IMV** の組み合わせでバグ修正が行われています。

GnuTLS の新機能

SSL、TLS、DTLS プロトコルの **GnuTLS** 実装がバージョン 3.3.8 に更新され、新機能や機能強化が追加されています。

- ※ **DTLS 1.2** のサポートが追加されています。
- ※ *Application Layer Protocol Negotiation* (ALPN) のサポートが追加されています。
- ※ elliptic-curve 暗号スイートのパフォーマンスが向上されています。
- ※ 新しい暗号スイート **RSA-PSK** と **CAMELLIA-GCM** が追加されています。
- ※ *Trusted Platform Module* (TPM) 標準のネイティブサポートが追加されています。
- ※ **PKCS#11** スマートカードおよび *hardware security modules* (HSM) のサポートがいくつかの点で改善されています。

-
- * *FIPS 140* 安全標準 (*Federal Information Processing Standards*) への準拠がいくつかの点で改善されています。

第15章 デスクトップ

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.

Quad-buffered OpenGL ステレオビジュアルのサポート

GNOME Shell および **Mutter** コンポジット型ウィンドウマネージャーでは対応ハードウェアでの quad-buffered OpenGL ステレオビジュアルの使用が可能になります。機能を正しく動作させるには NVIDIA ディスプレイドライバーバージョン 337 またはそれ以降をインストールしておく必要があります。

オンラインアカウントプロバイダー

GSettings の新しいキー `org.gnome.online-accounts.whitelisted-providers` が **GNOME** オンラインアカウント (*gnome-online-accounts* パッケージで提供される) に追加されています。このキーにより起動時の読み込みが明示的に許可されているオンラインアカウントプロバイダーの一覧が提供されます。キーを指定すると、管理側で適切なプロバイダーを有効にする、または選択的に他のプロバイダーを無効にすることができます。

第16章 サポート対応とメンテナンス

ABRT 承認マイクロレポート機能

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.

[本機能の詳細](#) についてはカスタマーポータルを参照してください。

第17章 Red Hat Software Collections

AMD64 および Intel 64 アーキテクチャーの Red Hat Enterprise Linux 6 と Red Hat Enterprise Linux 7 の全対応リリースにインストールして使用できる動的なプログラミング言語、データベースサーバー、関連パッケージなどを提供する Red Hat のコンテンツセットが Red Hat Software Collections です。

Red Hat Software Collections で配信される動的な言語、データベースサーバー、その他のツールは Red Hat Enterprise Linux で提供されるデフォルトのシステムツールに取って代わるものではありません。また、こうしたツールで優先的に使用されるわけでもありません。

Red Hat Software Collections では **scl** ユーティリティーをベースとした代替パッケージングメカニズムを使用してパッケージの並列セットを提供しています。このセットにより Red Hat Enterprise Linux で代替パッケージのバージョンが使用できるようになります。**scl** ユーティリティーを使用すると実行したいパッケージバージョンをいつでも選択することができます。



重要

Red Hat Software Collections のライフサイクルおよびサポートは Red Hat Enterprise Linux に比べ短くなります。詳細は [Red Hat Software Collections Product Life Cycle \(Red Hat Software Collections 製品ライフサイクル\)](#) を参照してください。

Red Hat Developer Toolset は Red Hat Software Collections の一部になり、別の Software Collection として収納されます。Red Hat Enterprise Linux プラットフォームで作業する開発者向けに設計され、GNU Compiler Collection、GNU Debugger、Eclipse 開発プラットフォーム、その他の開発ツール、デバッグツール、パフォーマンス監視用ツールなどの現在のバージョンを提供しています。

Red Hat Software Collections のセットに収納されているコンポーネント、システム要件、既知の問題、使い方、各 Software Collection の詳細などについては [Red Hat Software Collections のドキュメント](#) を参照してください。

Red Hat Software Collections の一部となる Red Hat Developer Toolset に収納されているコンポーネント、インストール、使い方、既知の問題など詳細については [Red Hat Developer Toolset のドキュメント](#) を参照してください。

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](#).

パート 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 **qethgoat**, see [「OSA-Express5s Cards Support in qethgoat」](#)

第20章 Storage

- ※ **Targetd** plug-in from the **libStorageMgmt** API, see [「libStorageMgmt API を使用したストレージレイ管理」](#)
- ※ LSI Syncro CS HA-DAS adapters, see [「LSI Syncro のサポート」](#)
- ※ DIF/DIX, see [「DIF/DIX サポート」](#)

第21章 File Systems

- ✦ **Btrfs** file system, see [「Btrfs ファイルシステムのサポート」](#)
- ✦ **OverlayFS**, see [「OverlayFS」](#)

第22章 Kernel

- ※ **kpatch**, see [「動的なカーネルのパッチ適用」](#)
- ※ **crashkernel** with more than one CPU, see [「複数 CPU での Crashkernel のサポート」](#)
- ※ **dm-era** device-mapper target, see [「dm-era ターゲット」](#)
- ※ Cisco VIC kernel driver, see [「Cisco VIC カーネルドライバー」](#)
- ※ NFSoRDMA Available, see [「NFSoRDMA Available」](#)
- ※ Runtime Instrumentation for IBM System z, see [「Runtime Instrumentation for IBM System z」](#)
- ※ Cisco usNIC Driver, see [「Cisco usNIC Driver」](#)
- ※ Intel Ethernet Server Adapter X710/XL710 Driver Update, see [「Intel Ethernet Server Adapter X710/XL710 Driver Update」](#)

第23章 Virtualization

- ✦ USB 3.0 host adapter (xHCI) emulation, see [「KVM ゲスト向け USB 3.0 サポート」](#)
- ✦ Open Virtual Machine Firmware (OVMF), see [「Open Virtual Machine Firmware」](#)
- ✦ LPAR Watchdog for IBM System z, see [「LPAR Watchdog for IBM System z」](#)

第24章 Compiler and Tools

- ※ Accurate ethtool Output, see [「Accurate ethtool Output」](#)

第25章 Networking

- ✦ Trusted Network Connect, see [「Trusted Network Connect」](#)
- ✦ SR-IOV functionality in the **q1cnic** driver, see [「q1cnic ドライバーの SR-IOV 機能」](#)

第26章 Authentication and Interoperability

- ※ Use of AD sudo provider together with the LDAP provider, see [「AD および LDAP の sudo プロバイダーの使い方」](#).
- ※ Apache Modules for IPA, see [「Apache Modules for IPA」](#).

パート III. デバイスドライバー

本章では Red Hat Enterprise Linux 7.1 で更新された全デバイスドライバーをすべて記載しています。

第27章 ストレージドライバーの更新

- ※ **hpsa** ドライバーがバージョン 3.4.4-1-RH1 にアップグレードされています。
- ※ **qla2xxx** ドライバーがバージョン 8.07.00.08.07.1-k1 にアップグレードされています。
- ※ **qla4xxx** ドライバーがバージョン 5.04.00.04.07.01-k0 にアップグレードされています。
- ※ **qlcnic** ドライバーがバージョン 5.3.61 にアップグレードされています。
- ※ **netxen_nic** ドライバーがバージョン 4.0.82 にアップグレードされています。
- ※ **qlge** ドライバーがバージョン 1.00.00.34 にアップグレードされています。
- ※ **bnx2fc** ドライバーがバージョン 2.4.2 にアップグレードされています。
- ※ **bnx2i** ドライバーがバージョン 2.7.10.1 にアップグレードされています。
- ※ **cnic** ドライバーがバージョン 2.5.20 にアップグレードされています。
- ※ **bnx2x** ドライバーがバージョン 1.710.51-0 にアップグレードされています。
- ※ **bnx2** ドライバーがバージョン 2.2.5 にアップグレードされています。
- ※ **megaraid_sas** ドライバーがバージョン 06.805.06.01-rc1 にアップグレードされています。
- ※ The **mpt2sas** ドライバーがバージョン 18.100.00.00 にアップグレードされています。
- ※ The **ipr** ドライバーがバージョン 2.6.0 にアップグレードされています。
- ※ **kmod-lpfc** パッケージが Red Hat Enterprise Linux 7 に追加されています。ファイバーチャンネル (FC) アダプターおよびイーサネット経由のファイバーチャンネル (FCoE) アダプターで **lpfc** ドライバーを使用する際の安定性が高まります。**lpfc** ドライバーがバージョン 0:10.2.8021.1 にアップグレードされています。
- ※ **be2iscsi** ドライバーがバージョン 10.4.74.0r にアップグレードされています。
- ※ **nvme** ドライバーがバージョン 0.9 にアップグレードされています。

第28章 ネットワークドライバーの更新

- ※ **bna** ドライバーがバージョン 3.2.23.0r にアップグレードされています。
- ※ **cxgb3** ドライバーがバージョン 1.1.5-ko にアップグレードされています。
- ※ **cxgb3i** ドライバーがバージョン 2.0.0 にアップグレードされています。
- ※ **iw_cxgb3** ドライバーがバージョン 1.1 にアップグレードされています。
- ※ **cxgb4** ドライバーがバージョン 2.0.0-ko にアップグレードされています。
- ※ **cxgb4vf** ドライバーがバージョン 2.0.0-ko にアップグレードされています。
- ※ **cxgb4i** ドライバーがバージョン 0.9.4 にアップグレードされています。
- ※ **iw_cxgb4** ドライバーがバージョン 0.1 にアップグレードされています。
- ※ **e1000e** ドライバーがバージョン 2.3.2-k にアップグレードされています。
- ※ **igb** ドライバーがバージョン 5.2.13-k にアップグレードされています。
- ※ **igbvf** ドライバーがバージョン 2.0.2-k にアップグレードされています。
- ※ **ixgbe** ドライバーがバージョン 3.19.1-k にアップグレードされています。
- ※ **ixgbev** ドライバーがバージョン 2.12.1-k にアップグレードされています。
- ※ **i40e** ドライバーがバージョン 1.0.11-k にアップグレードされています。
- ※ **i40evf** ドライバーがバージョン 1.0.1 にアップグレードされています。
- ※ **e1000** ドライバーがバージョン 7.3.21-k8-NAPI にアップグレードされています。
- ※ **mlx4_en** ドライバーがバージョン 2.2-1 にアップグレードされています。
- ※ **mlx4_ib** ドライバーがバージョン 2.2-1 にアップグレードされています。
- ※ **mlx5_core** ドライバーがバージョン 2.2-1 にアップグレードされています。
- ※ **mlx5_ib** ドライバーがバージョン 2.2-1 にアップグレードされています。
- ※ **ocrdma** ドライバーがバージョン 10.2.287.0u にアップグレードされています。
- ※ **ib_ipoib** ドライバーがバージョン 1.0.0 にアップグレードされています。
- ※ **ib_qib** ドライバーがバージョン 1.11 にアップグレードされています。
- ※ **enic** ドライバーがバージョン 2.1.1.67 にアップグレードされています。
- ※ **be2net** ドライバーがバージョン 10.4r にアップグレードされています。
- ※ **tg3** ドライバーがバージョン 3.137 にアップグレードされています。
- ※ **r8169** ドライバーがバージョン 2.3LK-NAPI にアップグレードされています。

第29章 グラフィックドライバーの更新

- ※ **vmwgfx** ドライバーがバージョン 2.6.0.0 にアップグレードされています。

パート 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

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).

パート 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_paranoid** 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](#).



注記

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 **zip1** 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 drumps 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_Rela)' 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.

付録A 改訂履歴

改訂 1.0-26	Mon Aug 01 2016	Lenka Špačková
Added a note about limited support for Windows guest virtual machines to Deprecated Functionality.		
改訂 1.0-25	Fri Jun 03 2016	Lenka Špačková
Added an OpenSSL known issue, which is valid for all Red Hat Enterprise Linux 7 minor releases.		
改訂 1.0-24	Thu May 26 2016	Lenka Špačková
Fixed the "Intel Broadwell Processor and Graphics Support" note in the Hardware Enablement chapter.		
改訂 1.0-22	Wed Apr 20 2016	Jiří Herrmann
Added a deprecated Virtualization feature.		
改訂 1.0-21	Wed Oct 14 2015	Lenka Špačková
Fix in the Architectures chapter.		
改訂 1.0-20	Mon May 04 2015	Radek Bíba
Update of the Red Hat Enterprise Linux 7.1 Release Notes.		
改訂 1.0-13	Tue Mar 03 2015	Milan Navrátil
Red Hat Enterprise Linux 7.1 リリースノートの公開		
改訂 1.0-9.3	Tue Jan 27 2015	Noriko Mizumoto
翻訳および査読完了		
改訂 1.0-9.1	Tue Jan 27 2015	Noriko Mizumoto
翻訳ファイルを XML ソースバージョン 1.0-9 と同期		