

TECH NOTES

RED HAT ENTERPRISE LINUX 6 SERVER FEATURES AND FUNCTIONS SUMMARY

EFFICIENCY, SCALABILITY, AND RELIABILITY

Red Hat® Enterprise Linux® 6 evolves in concert with hardware advances, reducing system power consumption, taking advantage of hardware with greater numbers of processing and memory resources, and withstanding hardware failures better.

Red Hat Enterprise Linux 6 continues the tradition of delivering the best operating system for data center needs. Advances in system efficiency lead to higher performance as well as reduced operating costs. Increased scalability meets the need for larger systems with more CPU sockets and cores, more memory, and larger filesystems. System evolution that makes the most of emerging hardware features such as Reliability, Availability, Serviceability (RAS) capabilities keeps your systems running in the face of hardware failures that would have caused previous generations of system to crash.

Features	Functionality	Description
Efficient Scheduling	Completely Fair Scheduler (CFS)	The CFS models an abstract, idealized multi-tasking CPU, where all tasks get an equal amount of runtime. To achieve this ideal, the CFS dynamically adapts the length of time a process will run uninterrupted based on the number of currently runnable processes. This results in less context switching overhead – thereby increasing application efficiency. The CFS scheduler has more awareness of the hardware topology which is factored into its decision making such as on which processor and core to run a job. This allows for more efficient operation on today's multi-core architectures. For example, during periods of peak load, tasks are dispersed across the full set of processors. During lower periods of activity, the tasks are aggregated to a smaller number of processor cores to dynamically allow other cores to enter low power consumption operating mode.
Power Management	Tickless kernel	Time-keeping within the kernel has been changed to be more demand based rather than interrupt driven. This keeps systems in the idle state longer which in turn allows them to fall into lower power more often, resulting in a net power savings compared to previous releases.
	Active State Power Management and Aggressive Link Power Management	Enhanced system control reduces the power consumption of I/O subsystems. Administrators can actively throttle power levels to reduce consumption.
	Realtime drive access optimization	Reduces filesystem metadata write overhead.
Scalability		Red Hat Enterprise Linux 6 supports more sockets, more cores, more threads, and more memory.
File Systems	Ext4 filesystem	Ext4 is the successor to ext3, and is the default file system for RHEL 6. It provides more robust journaling than ext3 and provides hugely improved file system (fsck) check times. Ext4 is highly scalable, there is no limit on number of sub-directories within a directory, and uses delayed allocation for better performance and reduced fragmentation.
	XFS filesystem ¹	XFS is a mature and robust massively scalable filesystem that is extremely well suited for handling large files with streaming I/O. It provides capabilities for online growth, defragmentation, and native backup/restore utilities.
	NFSv4 network file system	NFSv4 allows hosts to mount file systems over a network and interact with those file systems as though they are mounted locally. NFSv4 interoperates with NFSv3 and NFSv4. Improvements include improved tracking of file state between client and server, leased based locking, file delegation that allows clients to modify cached files, compound RPCs, and support of sophisticated security protocols.
	GFS2 shared storage filesystem ²	GFS2 is a high-availability, clustered filesystem in which data is shared among GFS2 nodes with a single, consistent, and coherent view of the file system name space. Processes on different nodes may work with GFS2 files in the same way that processes on one node can share files in a local file system.

¹ Available as an Add-on

² Available as an Add-on

Features	Functionality	Description
	Write Barriers	A write barrier is a kernel mechanism that forces storage write cache flushes before and after I/Os to insure that all data has been successfully written and no re-ordering has occurred. This minimizes the chance of data corruption, particularly when spinning storage devices with volatile write caches lose power. Write barriers are enabled for ext3, ext4, XFS, and GFS2.
	Fuse support	A kernel module that allows file-systems to run in user space.
Reliability, Availability, Serviceability (RAS)	Resource scalability	Red Hat Enterprise Linux 6 supports more sockets, more cores, more threads, and more memory.
	Hot add CPU and memory	Dynamic addition of CPU and memory to preserve system uptime.
	Machine Check: Advanced Error Handling	When the underlying hardware supports it, the platform now has the ability to recover from some fatal memory, CPU, or I/O errors with minimal disruption.
	HWPoison	Memory pages are declared as “poisoned”, processes associated with the page are killed, and the kernel will avoid using the page.
High Availability³	Rich GUI based cluster management and administration	The web interface based on Conga has been re-designed for added functionality and ease of use.
	Corosync	Corosync is a cluster group communication system, provides an extremely mature, secure, high performance, and light-weight high-availability solution. NOTE: Corosync is no longer coupled with OpenAIS.
	Unfencing	Fencing is a means of removing access to resources from a cluster node that has lost contact with the cluster, thereby protecting the resource (e.g., shared storage) from uncoordinated modification. Normally, the unconnected cluster node must be rebooted to safely rejoin the cluster. However, unfencing allows a node to re-enable access itself when starting up, without administrative intervention. Extensive improvements were also made in the SCSI-3 PR reservations based fencing. Virtualized KVM guests can be run as managed services, enabling fail-over, including between physical and virtual hosts. A single cluster command can be used to manage system logs from different services, and the logs have a consistent format that is easier to parse.
	Virtualization integration	Virtualized KVM guests can be run as managed services, enabling fail-over, including between physical and virtual hosts.
	Unified Logging and Debugging	A single cluster command can be used to manage system logs from different services, and the logs have a consistent format that is easier to parse.

² Available as an Add-on

³ X86_64 CPU architecture support only. X86 support related to Red Hat Enterprise Linux guests.

UNPRECEDENTED RESOURCE MANAGEMENT

Red Hat Enterprise Linux forges new ground in managing processing, memory, storage, and network resources.

The primary responsibility of an operating system is to manage processing, memory, storage, and network resources, and Red Hat Enterprise Linux forges new ground in deriving benefits from its integral relationship with system hardware. Partitioning processing resources to applications with cgroups and associated controllers can help those applications meet SLA, reduce resource bottlenecks, and increase net system performance. Management of storage and networks includes availability and performance improvements.

Features	Functionality	Description
System Resource Allocation	cgroups (Control Groups)	Cgroups allow system tasks to be grouped so that they can be tracked, and so that other system services can apply processing controls to them. For example, cpuset (below) associates a set of CPUs and memory nodes with the tasks in a cgroup. Networking and I/O Controllers are also available for use with cgroups. Easy configuration and management of cgroups is provided by the cgexec and cgclassify commands.
	CPUSET resource controller	Partitioning: Allocate specific quantities of CPU and network resources to individual guest hosts. By partitioning and reserving computing resources, applications can avoid contention for processing power and networking resources, allowing applications to better meet SLAs.
	Memory resource controller	The memory resource controller is a cgroups-based feature that provides finer grained control of memory allocation to system administrators and applications. Memory hungry applications can be constrained to use a limited amount of memory, or conversely, critical applications can be given access to the memory that they will need so that they can execute with a deterministic outcome.
	Network controller	The network resource controller is a cgroups-based feature that schedules and shapes network traffic belonging to tasks in a cgroup. A system administrator can use this to provide a maximum bandwidth for specific cgroups.
Storage	Mirrored mirror logs	The LVM maintains a log that keeps track of which regions need to be synchronized within a set of mirrors. In RHEL 6, the mirror log is itself mirrored, providing an additional layer of high availability.
	Merging snapshots & rollback	This feature allows a user to merge a snapshot of a logical volume back into the origin logical volume. This allows system administrators to revert any changes that have occurred on a logical volume by merging back to the point preserved by a snapshot. In other words a user can take a snapshot, update the origin, try it, if one does not like it, rollback the origin to the state it was in at the time of the snapshot.
	Four Volume mirror	Up to four mirrors are supported for a volume. This was previously unspecified.
	LVM hot spare	Using this feature one can define how a mirrored logical volume behaves in the event of a device failure. This allows for more deterministic behavior in case of a device failure as well as higher availability.

Features	Functionality	Description
	Support for alternate path selection algorithms in Multipathing.	DM-Multipath introduces two new options when dynamically balancing load over paths. Paths can now be dynamically selected depending on either the queue size of each path or previous I/O time data. This offers users an ability to choose the optimal path selection algorithm based on storage configuration and traffic patterns.
	Discard support for thin provisioning	Allows a storage device to be informed when the filesystem detects that portions of a device are no longer in active use. This allows storage administrators to use thin provisioning to allocate storage on-demand, rather than allocating all storage up-front (often referred to as fat, or thick provisioning). This minimizes storage over-provisioning.
	Self tuning for storage properties	The SCSI and ATA standards have been extended to allow devices that support the SCSI extensions to provide alignment and I/O hints when queried. Linux retrieves the alignment and I/O hints from the device and presents a uniform sysfs interface for all Linux block devices. This allows for seamless tuning and I/O alignment without manual intervention.
	End-to-end data integrity check (DIF/DIX)	Better integrity check for application data.
Networking	UDP Lite support	UDP Lite is a protocol, like UDP, that uses partial checksums to protect packets so that partially corrupted packets can be delivered (rather than rejected, as in UDP). This is useful for multimedia protocols, such as VOIP, where partial packets are better than none.
	Multiqueue Networking Support	Multiqueue Networking allows network device drivers to send data packets to multiple CPU queues for processing, improving system performance by parallelizing processing across multiple processors and CPU cores.
	Large Receive Offload (LRO) / Generic Receive Offload (GRO)	Improves network performance by aggregating multiple inbound packets into a large buffer before passing the data to the CPU for processing. GRO, the new system default, uses the same aggregation technique as Large Receive Offloads, but can be applied to a wider range of transport layer protocols.
	Data Center Bridging Support	Data Center Bridging is a hardware-based enhancement to ethernet LANs that provides loss-less communications for storage data. It provides Quality of Service benefits by providing eight different priorities for data traffic, and enables priority flow control which can limit or eliminate the number of dropped packets during network stress.
	Fiber Channel over Ethernet (FCoE)	Encapsulation of Fiber Channel (FC) frames within TCP/IP packets allows FC traffic to be sent over standard Ethernet. This allows data centers to avoid purchasing a second form of networking for storage, along with its required support.
	iSCSI root/boot	iSCSI partitions may be used as either root or boot filesystems.
	IPv6 RFC support	The next-generation Internet Protocol version 6 specification (IPv6) is the successor to Internet Protocol version 4 (IPv4). Kernel subsystems and other system services, including NFSv4, CIFS, iSCSI, RDMA, and FCoE, have been upgraded to support IPv6.

DESIGNED-IN SECURITY

Red Hat Enterprise Linux provides a complete portfolio of security technologies with solutions for all facets of your system, data, and communications security challenges.

Red Hat Enterprise Linux provides a complete portfolio of security technologies with solutions for all facets of your system, data, and communications security challenges. A secure system depends on controlling access to services and data, and Red Hat Enterprise's renowned SELinux's mandatory access control has been applied to more system services, and can be used to run untrusted applications without risk to your system. Overall system security is improved by standardized deployment documentation that verifies patch installation and detects possible system compromises. Many improvements to system identity services include a new service that can aggregate, normalize, and cache information from diverse identity servers.

Features	Functionality	Description
Access Control	SELinux Kiosk Mode	Kiosk mode allows users to login to a RHEL system, where their home directory and environment is created for the duration of their session. SELinux controls user actions while they are logged in, and removes the home directory and environment when they log out, creating a secure session which cannot be accessed by other users.
	SELinux Sandbox	An SELinux sandbox can run an application in a highly restricted manner, limiting its ability to harm the host system environment. When untrusted applications (e.g., web browsers) are run this way, the system will be protected from security exploits attempted by the application.
	Lower process access rights, system defaults tuned for system hardening, libcap-ng library	Standard file and process permissions were reviewed and reduced whenever possible in order to reduce risks of privilege escalation. A new software library and a set of utilities provide additional control over process privileges and includes an easy development interface for managing reduced capabilities.
Enforcement and Verification of Security Policies	OpenScap	OpenScap defines a standard format for Security Content Automation Protocol (SCAP) content which provides a standardized approach to maintaining the security of enterprise systems. This enables automatic verification of the presence of patches, checking system security configuration settings, examining systems for signs of compromise, and auditing using Open Vulnerability and Assessment Language (OVAL®).
VPN	Openswan	Openswan includes a general implementation of IPsec that works with Cisco IPsec.
Identity and Authentication	System Security Services Daemon (SSSD)	Provides centralized access to identity and authentication remote resources, such as LDAP, through a common framework that enables caching and offline support. The new functionality is particularly useful for deployment of corporate notebooks. SSSD provides limited support for Active Directory Domain Services.
	OpenLDAP	OpenLDAP is a compliant LDAP client with high availability from N-way MultiMaster replication, and performance improvements.

STABLE APPLICATION DEVELOPMENT AND PRODUCTION PLATFORM

Red Hat Enterprise Linux is the development and deployment platform of choice with an efficient, scalable, and robust operating system, certified security and flexible deployment options on physical and virtualized platforms, and a wide variety of powerful development languages backed by excellent debugging and tuning tools.

Features	Functionality	Description
Web Infrastructure	Apache 2.2.15	This release includes many improvements, see Overview of new features in Apache 2.2.
	Squid 3.1.4	This major new version of the popular web proxy application improves manageability and includes IPv6 support.
	Memcached 1.4.4	A high-performance and highly scalable distributed memory object caching system which enhances the speed of dynamic web applications.
Java	OpenJDK 6	OpenJDK is an open source implementation of the Java Platform Standard Edition (SE) 6 specification and provides a fully open source platform for Java runtime and development. TCK-certified based on the IcedTea project. The fully open source implementation of a Java Web Browser plugin and Java web start removes the need for proprietary plugins. Tight integration of OpenJDK and Red Hat Enterprise Linux includes support for Java probes in SystemTap to support better debugging for Java.
	Tomcat 6	Tomcat is an open source and best of breed application server running on the Java platform. With support for Java Servlets and Java Server Pages (JSP), tomcat provides a robust environment for developing and deploying dynamic web applications.
Development	Ruby 1.8.7	Enables application development and deployment with Ruby on Rails 2 and Rails 3. Support for Rails 3 dependencies is new in RHEL 6.
	gcc 4.4 (C compiler)	OpenMP3 conformance for portable parallel programs, Integrated Register Allocator, Tuples, additional C++0x conformance implementations, debuginfo handling improvements.
	Libraries	Malloc optimizations include improved speed and efficiency for large blocks and NUMA considerations. Lock-free C++ class libraries. NSS crypto consolidation for LSB 4.0 and FIPS level 2. Improved automatic parallel mode in C++ library.
	gdb 7.1.29	Many improvements including C++ function, class, templates, variables, constructor / destructor improvements. Catch / throw and exception improvements. Large program debugging optimizations. Non-blocking thread debugging (threads can be stopped and continued independently).
	TurboGears 2	A powerful internet-enabled framework enables rapid web application development and deployment in Python.
Application Tuning	PHP 5.3.2 Perl 5.10.1	These updates to popular web scripting and programming languages include many improvements.
	SystemTap	SystemTap allows applications to be probed by their interaction with the kernel, providing non-intrusive debugging information.
	KTune/tuned	The tuned daemon monitors system use and uses that information to automatically and dynamically adjust system setting for better performance.

Features	Functionality	Description
	SELinux	Using SELinux in permissive mode provides developers with insight into application access to system resources, allowing access reorganization for greater security.
Databases	PostgreSQL 8.4.4	PostgreSQL is an open source object-relational database system. It is field-proven and is known for its reliability and data integrity. This release includes many improvements, see PostgreSQL 8.4 Feature List.
	MySQL 5.1.7	MySQL is the world's most popular open source database. This release includes many improvements, see What Is New in MySQL 5.1.
	SQLite 3.6.20	SQLite includes incompatible changes to the internal OS interface and VFS layers (compared to earlier releases), significant performance improvements, and many important bug fixes.
System API / ABI Stability	Runtime stability within the major release	For the full ten-year life cycle, the application runtime ABI/API is kept stable. Because of this applications are not affected by security errata or service packs and do not require re-certification.
	ABI compatibility	In addition backward compatibility for the core ABI is provided across major releases allowing existing applications to span subsequent releases.
	Extended Update Support (EUS) ⁴	EUS allows customers to standardize on a specific minor version of Enterprise Linux for 18 months instead of the normal 6 months. This reduces the amount of internal testing and revalidation the customer needs to do annually and reduces the amount of change introduced into their environment.

⁴ Available as an Add-on

INTEGRATED VIRTUALIZATION

Virtualization is a foundation for data center planning. Red Hat Enterprise Linux is an excellent virtualization host, and is designed to be a superior guest on any of the major hypervisors.

Virtualization is a foundation for data center planning. Not only is Red Hat Enterprise Linux designed to be a superior guest on any of the major hypervisors, but it can also be a virtualization host. Virtualization is integrated directly into the Red Hat Enterprise Linux kernel using KVM technology. As part of the kernel, the complete breadth of system management, security tools and certifications are available for your administrators to use.

Features	Functionality	Description
Kernel-Based Virtualization (see kernel for version number)	Deployment Flexibility	The ability to easily move guests between hosts can be used to consolidate resources onto fewer machines during quiet times, or free up hardware for maintenance downtime.
Leverages Kernel Features		<p>Red Hat Enterprise Linux 6 features form a foundation of the benefits of virtualization:</p> <ul style="list-style-type: none"> • Hardware abstraction allows applications to move from physical to virtualized environments independently of underlying hardware. • Increased scalability of CPUs and memory provides more guests per server. • Selectable I/O schedulers and support for asynchronous I/O provide flexibility and performance benefits for block storage. • Cgroups and related CPU, memory, and networking resource controls provide the ability to reduce resource contention, improve overall system performance, and help applications meet SLAs. • Reliability, Availability, and Serviceability (RAS) features minimize downtime. • Multicast bridging provides intelligent packet routing and enhanced network efficiency.
Guest Acceleration		<p>Various capabilities, including guest access to physical resources through virtual interfaces, provide advantages to guests:</p> <ul style="list-style-type: none"> • Single physical hardware resources can be shared by multiple guests. • Guest memory allocation, particularly at start-up, is significantly quicker. • Reuse (deduplication) of identical pages across virtual machines improves memory use. • Access to system interrupts via software signals provides flexibility and convenience. • The tickless kernel defines a stable time model for guests, avoiding clock drift. <p>Advanced paravirtualization interfaces, offering higher performance, include non-traditional devices such as the clock (enabled by the tickless kernel), interrupt controller, spinlock subsystem, and vmchannel.</p>
	CPU	CPU masking allows all guests to use the same type of CPU.
	SR-IOV	SR-IOV virtualizes physical I/O card resources, primarily networking, allowing multiple guests to share a single physical resource.
	Message signaled interrupts	MSIs provide an alternate way to generate an interrupt, as well as increasing the number of different interrupts (by passing a subtype as data within the interrupt). MSIs are delivered to applications as signals.
Guest Security	SVirt Mandatory Access Control	sVirt uses SELinux's Mandatory Access Control (MAC) security to limit access from a guest host to the Virtualization system itself, thereby preventing one guest from compromising another guest in a virtualized environment.
Microsoft Windows Support	Windows WHQL-certified drivers and SVVP program	Approved Windows hardware drivers support running Windows systems in virtualized environments, and meet the requirements of Microsoft's Server Virtualization Validation Program (SVVP), allowing Microsoft customers to receive technical support for Windows Server in virtualized environments.

ENTERPRISE MANAGEABILITY

Red Hat provides two categories of control for the data center operations. Our Smart Management modules and Satellite servers centralize and automate common management functions so that data centers can scale systems without burdening the staff. Red Hat also packages the platform in system administrator friendly packages that allow the tailoring of installations to application needs.

Features	Functionality	Description
Installation, Updates, and Deployment	Anaconda	Anaconda is the Red Hat Enterprise Linux installer. A system administrator can now choose to install a “minimal platform” as a starting point for creating specific server installation, or a strategy to increase security by reducing number of unnecessary software packages.
	Red Hat Network (RHN) and Satellite	Red Hat Network and RHN Satellite provides powerful systems administration capabilities such as management, provisioning and monitoring for large deployments. Administrators can use RHN Satellite to install system updates and ensure security fixes and configuration files are applied across your environment consistently.
	Installation to meet Workload Profiles	Installation options have been reorganized into “workload profiles” so that each system installation will provide the right software for specific tasks.
	Dracut	Dracut is a new toolkit for building custom ramdisk images, used for preloading modules during Linux pre-boot. Dracut minimizes the impact of underlying hardware changes, is more maintainable, makes it easier to support third party drivers, and replaces the initramfs and mkinitrd tools.
	Yum and RPM	Yum and RPM are a powerful combination of software package management tools for installing, updating, and removing packages and package groups in Red Hat Enterprise Linux. The new history command lists detailed information on the history of yum transactions, and allows undoing and redoing selected operations. This release of Yum and RPM includes significant performance improvements.
	Secure Installation and Upgrades	RPM signatures use the Secure Hash Algorithm (SHA256) for data verification and authentication, providing improved security.
	Encrypted devices	Designating encryption for devices during installation protects user and system data from unauthorized access. Key escrow allows enterprise IT departments to recover lost keys.
	Standards Based Linux Instrumentation for Manageability (SBLIM)	SBLIM consists of a set of standards-based, Web-Based Enterprise Management (WBEM) modules that use the Common Information Model (CIM) standard to gather and provide systems management information, events, and methods to local or networked consumers. This package provides a set of core providers and development tools for systems management applications.
RAS (Reliability, Availability, and Serviceability)	Automatic Bug Reporting Tool (ABRT)	Framework for unified collection of error data for automated reporting and analysis.
Routine Task Delegation	PolicyKit	Enables system administrators to configure user access to privileged operations, such adding a printer or rebooting a desktop, without granting administrative privileges. Default policy configurations are available for various computer roles.

Features	Functionality	Description
Printing	Many cups and system-config-printer changes	Provide improved printing, printer discovery, and printer configuration.
	SNMP-based monitoring of ink/toner/supply levels and printer status	Easier monitoring to enable efficient inventory management for ink/toner cartridges.
	Device discovery speed improvements	Faster printer discovery.
	Automatic PPD configuration for postscript printers	PPD option values queried from printer - available in CUPS' web interface.
Microsoft Interoperability	Samba	Comprehensive suite for file and printer sharing across multiple operating systems. The latest improvements include support for Windows 2008R2 trust relationships; Windows cross-forest, transitive trust, and one-way domain trust.
	OpenChange	The libmapi library from OpenChange allows applications to gain access to Microsoft Exchange servers using native protocols. This allows mail clients like Evolution to have tighter integration with Exchange servers.

ABOUT RED HAT

Red Hat was founded in 1993 and is headquartered in Raleigh, NC. Today, with more than 60 offices around the world, Red Hat is the largest publicly traded technology company fully committed to open source. That commitment has paid off over time, for us and our customers, proving the value of open source software and establishing a viable business model built around the open source way.

SALES AND INQUIRIES

NORTH AMERICA
1-888-REDHAT1
www.redhat.com

**EUROPE, MIDDLE EAST
AND AFRICA**
00800 7334 2835
www.europe.redhat.com
europe@redhat.com

ASIA PACIFIC
+65 6490 4200
www.apac.redhat.com
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
www.latam.redhat.com
info-latam@redhat.com