

FEATURES & BENEFITS

RED HAT ENTERPRISE LINUX 6.2 - WHAT'S NEW

On December 6th, a little over the one-year anniversary of Red Hat Enterprise Linux 6 availability, Red Hat released Red Hat Enterprise Linux 6.2. As with every minor release, Red Hat Enterprise Linux 6.2 consolidates all patches, enhancements and security updates from the previous Red Hat Enterprise Linux release (6.1 in May 2011). As always, Red Hat maintains certified application compatibility, ISV and IHV support and testing for all releases.

Red Hat Enterprise Linux 6.2 delivers dramatic improvements in resource management and high availability, new features aimed at storage and file system performance, identity management functionality, and technology integration with VMware and Microsoft environments.

The key benefits for organizations deploying Red Hat Enterprise Linux 6.2 are higher levels of efficiency realized through resource management and performance optimization, along with enhanced business agility through additional security enhancements and more flexibility for virtualized and clustered environments.

Red Hat Enterprise Linux 6.2 What's New Technical Features and Benefits document
http://www.redhat.com/rhel/server/whats_new/

Red Hat Enterprise Linux 6.2 press release here <http://www.redhat.com/about/news/prarchive/2011/first-anniversary-of-red-hat-enterprise-linux-6>

HARDWARE ENABLEMENT

Reliability, Scalability and Performance

New features to support our partners

Enhancements in Reliability

- There are a number of enhancements to improve reliability that have been co-developed with our partners, Intel and AMD.
- The GCC (GNU Compiler Collection) compiler and glibc (GNU C Library) have been enhanced.
- The 32-bit library for AMD systems has been optimized to 32K memory alignment to avoid unnecessary cache validation.

Enhancement to PCI-e 3.0 and USB 3.0

Enabling system manufacturers to develop new I/O devices faster

- In Red Hat Enterprise Linux 6.2 support for PCI-e 3.0 as well as USB 3.0, enables faster and wider busses with the limited number of devices currently available on the market.
- Providing a stable, enterprise-ready support for PCI-e 3.0 and USB 3.0 provides system manufacturers and I/O developers a solid foundation for them to rapidly develop new I/O devices.

SGI UV 1000

Increased scalability in x86-64 systems

- Red Hat Enterprise Linux certification of SGI UV technology supports new levels of scalability and price performance in the largest x86-64 configuration at 1280 cores and 8TB memory.

Device Driver Support

Key functionality supporting GbE and Infiniband solutions

- Red Hat Enterprise Linux 6.2 includes support for many new 10 GbE Network adapters and Host Bus Adapters (HBAs) available from a large number of vendors, as well as support for dual-purposed Converged Network Adapters (CNAs).
- Expanded support and utilities simplify the configuration and deployment of Fibre Channel over Ethernet (FCoE) environments.
- Added support for new Infiniband-based devices as well as SR-IOV integration
- 3D support for integrated GPUs from Intel and AMD, as well as support for additional audio devices

VIRTUALIZATION AND CLOUD SUPPORT

Scalability, Performance, and More Efficient Management

More efficient management with cgroups

Processor ceilings improvements with cgroups

- Red Hat Enterprise Linux 6.2 now provides the ability to set upper limits for CPU consumption to better manage SLAs, especially for multi-tenant environments – this new feature provides better overall control of CPU allocation.
- The task scheduler in the Linux kernel (referred to as Completely Fair Scheduler or CFS) tries to allocate processor time proportionately between all the groups of tasks. This can often result in a task group getting more than its share of processor time if there are idle cycles available thus setting upper limits for consumption can better control allocation.
- Service Providers can now set maximums for CPU time associated with a virtual machine which allows them to better manage SLAs without having to “give away” CPU resources. (This capability already exists for memory.)

cgroup CPU controller

Improved scalability

- Lower overhead realized through cgroups, provides more efficient and better scaling for higher levels of guest consolidation.
- Control Groups, cgroups, are used to create and manage slices of computing resources (CPU, memory, I/O) for specific processes and can be used as part of SLA management. Large SMP systems that are hosting multiple tenants or processes need to be finely divided without limits. Enhancements in Red Hat Enterprise Linux 6.2 now provides the ability to create 100's of control groups and provide system administrators with very fine-grained control over their environment.
- System administrators now can to provision their servers at a very fine grain, making it easier to achieve very high system utilization rates for 100's of processes/tenants with no degradation of overall system performance. This is ideal for tightly managing SLAs.

cgroup I/O controller for improved performance

- The I/O controller within the cgroup framework performs better as a result of more efficient use of locks inside the kernel.

Red Hat Enterprise Virtualization 3.0

- Red Hat Enterprise Virtualization 3.0, currently in beta, will benefit from the performance, scalability, hardware enablement, KVM hypervisor enhancement, and advanced technologies enabled by Red Hat Enterprise Linux 6.2.
- Red Hat's Enterprise Virtualization provides a complete infrastructure for managing hundreds of physical hosts and thousands of virtual machines running Red Hat Enterprise Linux or Microsoft Windows guests for server and desktop virtualization use cases. More information about Red Hat Enterprise Virtualization 3.0 can be found at <http://www.redhat.com/rhev3>.

KVM Updates

Virtual CPU timeslice sharing for multiprocessor guests

- Virtual CPU timeslice sharing for multiprocessor guests is a new feature in Red Hat Enterprise Linux 6.2. Scheduler changes within the kernel now allow for virtual CPUs inside a guest to make more efficient use of the timeslice allocated to the guest, before processor time is yielded back to the host. This change is especially beneficial to large SMP systems that have traditionally experienced guest performance lag due to inherent lock holder preemption issues. In summary, this new feature eliminates resource consuming system overhead so that a guest can use more of the CPU resources assigned to them much more efficiently.
- CPU resources allocated to a guest can be partially consumed due to pre-emption, especially in multi-socket systems. Red Hat Enterprise Linux 6.2 was enhanced to ensure a guest can minimize pre-emption and use as much of the CPU resources allocated to them before the guest hits its pre-determined CPU limits.
- Guests will now use more of the CPU resources assigned to them which will increase the efficiency of the guest and the overall system by eliminating some resource-consuming system overhead.

Network performance improvements with KVM

- KVM network performance is a critical requirement for Virtualization and cloud-based products and solutions. Improvements have been made in KVM to more efficiently process a variety of workloads that generate small (less than 4K) messages. New functionality added has significantly reduced networking related CPU utilization on the host on which the guest is located.
- In addition, network drivers in KVM (virtio-net) have been enhanced to process UDP traffic faster. This is achieved with the help of more efficient UDP checksum validation. Also in cases where the UDP checksum has been validated by the host, it is not necessary to also validate in the guest KVM guests no longer validate the UDP checksum if that action has been performed by the host NIC.

Guest debugging configuration

- The ability to configure guests for debugging, while already present in KVM, is now available via the libvirt API and the virsh command line tool. This functionality enables an easier method of debugging virtualized guests and simplifies virtualization administration. It is now possible to trigger non-maskable interrupts (NMI) or initiate a back trace to debug guest environments.

Pre-boot guest environment auditing

- IT administrators want the ability to audit and monitor their virtualized environments, from creation to runtime and retirement. Red Hat Enterprise Linux 6.2 now enables administrators the ability to monitor and audit the entire boot sequence including the boot process before a guest is actually booted. This is achieved using the newly introduced sgabios component and the existing virsh console.

Guest memory pinning KVM

- Bare-metal systems have enjoyed the ability to pre-assign, a.k.a., pin down CPUs and memory in NUMA systems for the greatest performance and minimal delays due to waiting for these resources to be freed up. Prior Red Hat Enterprise Linux releases allowed CPU pinning in virtualized guests. With the added capability to pin down memory in KVM, this same capability has now been delivered for virtualized environments, delivering greater performance for critical virtualized workloads.

Linux Containers

- Linux containers provide a flexible approach to application runtime containment on bare-metal without the need to fully virtualize the workload. This release provides application level containers to separate and control the application resource usage policies via cgroup and namespaces. This release introduces basic management of container life-cycle by allowing for creation, editing and deletion of containers via the libvirt API and the virt-manager GUI.
- Linux Containers provides a means to run applications in a container, a deployment model familiar to UNIX administrators. Also provides container life-cycle management for these containerized applications through a graphical user interface (GUI) and user space utility (libvirt).
- Linux Containers is in Technology Preview at this time.

IDENTITY MANAGEMENT

Reduced Administrative Overhead with Increased Interoperability

- A new feature referred to as Identity Management in Red Hat Enterprise Linux 6.2 provides the administrative tools to quickly install, configure and manage server authentication and authorization in Linux enterprise environments.
- This feature provides the option to inter-operate with Microsoft Active Directory. Enterprises can now manage Linux environments easily and cost-effectively. Centralized identity management and host-based access control reduces administrative overhead, streamlines provisioning and improves overall security.

FILE SYSTEM

Performance and Reliability

Parallel NFS (pNFS) in Technology Preview

Transfer rates faster than NFS support

- pNFS is an extension of NFS that provides significantly larger data transfer rates compared to the traditional NFS architectures. This is achieved by processing the meta-data (a typical bottleneck for I/O throughput) separately from the actual data. By separating the meta-data and the data paths, pNFS allows clients to access storage devices directly and in parallel. pNFS is a part of standard NFS version 4.1 specification.
- Support for pNFS is limited to client-side functionality for file layout only.



XFS - delayed logging of meta-data enhanced performance

Available in the Scalable File System Add-On

- XFS now implements a delayed logging of meta-data feature. This new feature will significantly improve performance for highly paralleled meta-data intensive workloads such as thousands of small files in a directory. Since this is now a common use case, the default mount options have also been updated to use delayed logging.

Asynchronous writes in CIFS

- A 200% improvement in writing to the file system can be realized with this new feature
- The CIFS (Common Internet File System) client has historically dealt with synchronous writes which can be time consuming for distributed workloads that generate a lot of writes to the file system. To improve the performance of distributed workloads or applications with large transactions, the CIFS client now supports parallel writes.
- Linux Containers provides a means to run applications in a container, a deployment model familiar to UNIX administrators. Also provides container life-cycle management for these containerized applications through a graphical user interface (GUI) and user space utility (libvirt).

ext4 File System

Faster and simplified initialization

- Key functionality added to support ext4 file system provides customers the option to delay time-consuming initialization of the data structure until after the filesystem is mounted. Typically the creation of an ext4 filesystem consumes a lot of time when you try to initialize the data structure (inode) and mount the filesystem at the same time.



Clustered Samba on GFS2

Available in the Resilient Storage Add-On

- New feature improves throughput for Samba services, and increased availability for large Samba environments.
- Customers now have the ability to run clustered Samba services in an active-active setup using a shared GFS2 filesystem on a Red Hat cluster infrastructure. Prior releases only supported active-passive configurations, so only a single instance of Samba provided services within the cluster.

STORAGE



iSCSI for RDMA

Available in the High Performance Network Add-On

- Low latency and high throughput for 10Gb Ethernet solutions eliminates the need for expensive Infiniband hardware with dedicated fabrics.
- RDMA improves performance by avoiding unnecessary data copying at the network interface. RDMA is implemented in Infiniband, and in certain 10GbE NICs.
- iSCSI Extensions for RDMA (iSER) allows you to run the iSCSI storage protocol on RDMA on these interfaces. This allows you to take advantage of the same low latency and high bandwidth provided by Infiniband and RDMA-enabled NICs.
- iSER support added for both the software initiator (iscsi-initiator-utils) and for the software target (scsi-target-utils).

Reduced activation time for LVM devices

- Improved efficiency on flash based solid state disks (LVM devices)
- LVM support of activation/deactivation time. It is becoming commonplace to have multiple logical volumes rather than partitions for storage devices since they are easier to create and can be resized as needed without much effort. Just as important is the ability to bring these volumes on-and off-line very quickly especially for high LVM density environments.

LVM discard enhancements

- Enhanced discard commands that allow the operating system to mass information to the storage device about block ranges that are no longer in use. This enhancement improves the efficiency of solid state flash based devices, and thinly-provisioned storage devices. Red Hat Enterprise Linux 6.2 adds more discards in conjunction with LVM freeing operations.
- Improved efficiency on flash based solid state disks.

WWID installer support

Faster and more efficient identification of newly installed storage devices

- Support for device identification using World Wide Identifier (WWID). During installation, Fibre Channel and Serial Attach SCSI (SAS) devices can be now specified by a World Wide Name (WWN) or World Wide Identifier (WWID) for unattended installations.
- When a storage device is attached to a server using multiple physical paths for redundancy or improved performance, WWN for any of these paths is sufficient to identify the device.

NETWORKING

Faster and More Efficient Data Transmission

Stream Control Transmission Protocol (SCTP)

- Support for Transmission Control Protocol (TCP) which has long been the standard method for delivering and handling network traffic in an ordered fashion. Stream Control Transmission Protocol (SCTP), which separates messages and their control information into separate chunks, facilitates the streaming of related content (such as video and voice) with lower latency while also enhancing the security of transmitted information.
- SCTP delivers fast, reliable transmission of related but independent content such as voice, video and text over a single end-to-end connection (multi-streaming).
- Instead of relying on applications to retry a transmission for a failed network connection, SCTP supports rerouting to another network connection without interrupting the application.

Transmit Packet Steering (XPS)

Improved network throughput by up to 30% in performance tests conducted by Red Hat

- Networking traffic can sometimes be slowed down waiting for an unassigned CPU to complete the transaction. Transmit Packet Steering (XPS) allows an administrator to pre-assign a CPU to handle the network transmission requests which will remove this delay in waiting for the first available CPU to respond. This will complement the Receive-side Packet Steering that was previously delivered. This pre-assignment also has the benefit of preserving memory cache efficiency – particularly useful in modern NUMA systems.
- The ability to dedicate CPU resources for managing networking traffic results in improved network throughput by up to 30% while providing finer grain control over more resources.

HIGH AVAILABILITY



HA support for VMware

Available in the High Availability Add-On

- When an enterprise deploys their applications to run in a Red Hat Enterprise Linux 6.2 guest hosted by VMware, applications can be clustered for High Availability (HA). For virtual machines, this release includes full support for use of GFS2 which is a shared disk file system for Linux system clusters.
- Applications running in a Red Hat Enterprise Linux guest hosted by VMware can now be configured for High Availability (HA). This is achieved with the aid of GFS2. GFS2 is a shared disk file system for Linux clusters.
- This new feature delivers High Availability for large virtualized Red Hat Enterprise Linux deployments on VMware, bring this in line with what is already achieved with KVM.



Administrative UI enhancements

Available in the High Availability Add-On

- Simplified, intuitive User Interface (UI) for cluster configuration has been updated to include Role-Based Access Control (RBAC) that supports Luci, the web-based administrative UI configuring clusters. Luci has been updated to include the following:
 - Role-based access control (RBAC): Enables fine-grained access levels by defining user classes to access specific cluster operations.
 - Improved response times for destructive operations in a cluster.
 - Simplified, intuitive User Interface for cluster configuration and Role-Based Access Control.



UDP-Unicast support

Available in the High Availability Add-On

- UDP-Unicast support improvements makes it easier and faster to set up a clustering infrastructure. UDP-unicast in contrast to IP multicasting, offers a similar approach to cluster configuration and is an established protocol for cluster communication.
- To date IP multicasting has been the only supported option for a cluster transport. IP multicasting is inherently complex to configure and often requires re-configuration of network switches. UDP-unicast in contrast offers a simpler approach to cluster configuration and is an established protocol for cluster communication.



Watchdog integration with fence_scsi

Available in the High Availability Add-On

- This feature eliminates the need for manual intervention to reboot the node after it has been fenced using fence_scsi.
- The Watchdog feature is a general timer service available in Linux that can be used to periodically monitor system resources. Fence agents, commonly used in clustered deployment, have now been integrated with watchdog such that the watchdog service can reboot a cluster node after it has been fenced using fence_scsi.

Error detection and reporting

- Improved Automated Bug Reporting Tool (ABRT)
- The ABRT is a framework for collecting and analyzing application errors and system exceptions. Improvements in this release include easier configuration of plug ins and settings, a more consistent way to store error reports, a more secure environment as most of the processing is performed with a non-privileged account and more stability for plug ins.
- Customers will have better tools to either self-diagnose problems or send error reports to Red Hat for expedited resolution for problem reports.

DESKTOP AND GRAPHICS

- The X server has been re-based in this release. Updating the X server will increase system stability through the isolation of the system display drivers and will provide a better base for new features. Overall improved support for newer workstation optional hardware, multiple displays and new input devices.
- Improved support for Wacom graphic tablets commonly used by the graphic artists and the animation creators. It is now no longer necessary to reconfigure device settings after it has been unplugged and plugged back in. Easier installation and support for multiple devices is made possible with this feature enhancement.
- Tablet enhancements includes Improved support for Wacom graphic tablets commonly used by graphic artists and animators. Installation and support for multiple devices is much easier with Red Hat Enterprise Linux 6.2.

Improved Network Manager support for wireless roaming

- The Network Manager in Red Hat Enterprise Linux 6.2 has been enhanced to provide background scanning for wireless networks.
- Improved roaming in enterprise WiFi networks while consuming less system resources.

System monitor enhancements

- Today, a 4 socket server can provide up to 64 CPUs and larger systems can achieve up to 2048 CPUs. Virtualization can lead to very large virtual CPU (vCPU) counts. The gnome-system-monitor can now monitor systems has been enhanced to monitor greater than 64 CPUs.
- System administrators want and need the ability to monitor all of their resources to ensure they are meeting SLAs. This is especially true in high CPU count environments.

DEVELOPMENT TOOLS

New CPU optimizations

- Compilation and runtime performance optimizations matching the most recent CPUs.
- Gcc, glibc, gdb, oprofile – include optimizations to support new chip set capabilities – Intel Westmere & Ivy Bridge, AMD Bulldozer & Piledriver.

Systemtap enhancements

- Systemtap is a highly flexible tracing and monitoring tool. Recent enhancements in Red Hat Enterprise Linux 6.2 include the ability to work remotely over SSH and to provide more complete back traces.
- Increased ability to utilize systemtap tracing in virtualization and cloud deployments brings guest analysis more on parity with the capabilities of bare metal diagnosis.

Java runtime updates

- Newer OpenJDK provides hundreds of JDK class library and Hot Spot improvements, including performance and graphics rendering gains.
- IcedTea-Web improvements include faster performance in Firefox/Chromium and applet initialization, Proxy Auto Configuration file support, Web Start uses default browser, and more.
- Java is a rapidly evolving runtime environment, demanding the latest enterprise ready versions.

PERFORMANCE BENCHMARKS AND SCALABILITY LEADERSHIP

Red Hat already increases the IT capability of tens of thousands of enterprises while saving them hundreds of millions of dollars. Proven to perform and scale, Red Hat continues to top marks in Benchmark Performance tests with our partners.

Please reference our reference architecture page to view the details of all our benchmarks and industry publicized tests. http://www.redhat.com/rhel/resource_center/reference_architecture

Some of our more recent achievements include:

- IDC Standardize on Red Hat Enterprise Linux – lower TCO than Free or Mixed environment
- Fusion IO SSD performance: 1.3 million iops per second on Red Hat Enterprise Linux 6 – better than Oracle's published results
- SpecVirt leadership: Red Hat Enterprise Linux 6 continues leadership with a new benchmark on large systems and now has the #1 position in 2, 4 and 8 socket systems. Every single SpecVirt submission is running Red Hat Enterprise Linux guests, including VMware results.
- Proven filesystem improvements with AIM7 benchmarks comparing ext3, ext4 and XFS (Scalable File System Add-On)
- Great scaling and speed running SAS and processing over 90,000 jobs/minute
- Largest Linux result for SAP to date with new SAP Two-Tier SD benchmark on an HP 8 socket system
- Faster messaging and now iSCSI data movement using 10Gb Ethernet with HPN (RDMA over Ethernet)
- Red Hat Enterprise Linux 6 is 20%-30% lower idle power than Red Hat Enterprise Linux 5

FEATURES & BENEFITS

EXECUTIVE SUMMARY

Red Hat demonstrates its commitment to continuously delivering value and innovation to our customers with the availability of Red Hat Enterprise Linux 6.2. This release builds upon the rich enterprise platform that has set the standard in flexibility, efficiency and control that our customers worldwide rely on for their business. For more information on Red Hat Enterprise Linux visit <http://www.redhat.com/rhel>.

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
---	--	--	--