

A horizontal decorative bar consisting of five colored segments: a solid red segment on the left, followed by a blue segment with a white geometric pattern, an orange segment with a white wavy pattern, a green segment with a white wavy pattern, and a purple segment with a white wavy pattern.

An Overview of the Red Hat Enterprise Linux 4 Product Family

Abstract

This white paper provides information on the family of Red Hat Enterprise Linux and Red Hat Desktop products. It describes the family' s features and benefits and also gives a brief overview of the open source layered products designed for Red Hat Enterprise Linux environments.

Revision 4b. February 2005



Table of Contents

Red Hat Enterprise Linux Family Overview.....	3
Developing the Distribution.....	3
Creation of Fedora.....	3
Creation of Red Hat Enterprise Linux.....	4
Red Hat Enterprise Linux Products.....	5
Red Hat Enterprise Linux AS.....	6
Red Hat Enterprise Linux ES.....	6
Red Hat Enterprise Linux WS.....	6
HPC with Red Hat Enterprise Linux WS.....	7
Red Hat Desktop.....	7
Product Summary.....	8
Example Configuration.....	9
Technical Features.....	9
Read Copy Update (RCU).....	10
Selectable I/O elevators.....	10
Object-Based Reverse Map VM.....	11
Generic logical CPU scheduling.....	12
Block I/O subsystem.....	12
Sys_epoll() support.....	12
Support for larger server systems.....	13
Upward Compatibility.....	13
File System Performance enhancements.....	13
Red Hat Desktop.....	13
Security.....	15
Auditing.....	17
Compiler and Library Buffer Management.....	17
Advanced Glibc memory corruption checks.....	17
Printf format string exploit prevention.....	17
GCC buffer bound checking.....	17
Standards Compliance.....	17
Development Environment.....	18
Storage Subsystem.....	18
Automounter.....	19
Networking.....	19
Feature Summary.....	19
Support Services.....	20
Red Hat Network.....	21
Application Availability.....	22
Hardware Availability.....	23
Benchmarks.....	24
Layered Products for Red Hat Enterprise Linux.....	25
Red Hat Global File System.....	25
Red Hat Cluster Suite.....	26
Comparing Red Hat Global File System and Red Hat Cluster Suite.....	26
Red Hat Application Server.....	28
Red Hat Developer Suite.....	28
Summary.....	29



Red Hat Enterprise Linux Family Overview

Since 2002, Red Hat has steadily expanded its range of open source, commercially-focused operating system and middleware products. These products provide the industry's premier Linux environment for commercial deployments.

The operating system products, sold by annual subscription under the name Red Hat Enterprise Linux, have been rapidly adopted and supported by a wide range of Independent Software Vendors (ISVs) and Original Equipment Manufacturers (OEMs). They offer excellent performance, scalability, and security, and a comprehensive array of services delivered by Red Hat and its partners. As a result, Red Hat Enterprise Linux solutions, deployed on certified commodity hardware and running a wide variety of enterprise-caliber applications, are delivering the capabilities of traditional proprietary UNIX systems but at significantly lower cost.

The initial releases of the Red Hat Enterprise Linux family, versions 2.1 and 3, are described in earlier white papers (see *An Overview of the Red Hat Enterprise Linux product family*, March 2003 and June 2004). This paper describes the latest release of the family, version 4, which was delivered in February 2005.

Developing the Distribution

As the leading provider of open source software solutions, Red Hat implements a sophisticated development process to create the Red Hat Enterprise Linux family of products. The process has two major phases:

Creation of Fedora

The Fedora Project is a Red Hat-sponsored and community-supported open source project. It serves as a proving ground for new technology that may eventually make its way into commercial Red Hat products.

The goal of the Fedora Project is to work with open source development communities to build a complete, general purpose operating system exclusively from open source software. All development is done in a public forum. Fedora Core releases are issued about 2-3 times a year and are available for free download from Red Hat servers and over 200 mirror sites worldwide. The leading-edge, rapidly-changing nature of Fedora makes it impractical for use in commercial environments, and it is not formally supported by Red Hat or its ISV/OEM partners.

The first stage in the process of creating Fedora requires defining the set of packages to be used. The number of packages to choose from in the open source arena is huge. A single code repository such as Sourceforge (www.sourceforge.net) alone has over 90,000 packages and almost 1,000,000 registered users. So package selection is a complex exercise, resulting in approximately 1500-2000 being selected. These packages are then built and integrated into a complete system, a process that requires significant engineering resources including new development, bug fixes, creation of an installation program, management utilities, documentation, and



the project management necessary to coalesce a large group of distinct projects into a usable whole.

Fedora has established itself as a highly successful free distribution and widely regarded as the de facto standard platform for applied software research and development.

Creation of Red Hat Enterprise Linux

While the creation of Fedora can be considered a first stage distillation of open source software projects into a complete distribution, the creation of Red Hat Enterprise Linux takes this process another step, the second stage distillation. In the Fedora arena, software packages enjoy significant public exposure and mature rapidly. Red Hat creates the Enterprise Linux family of products by selecting approximately 1000-1500 of the most stable Fedora packages. Those that are not selected are either not sufficiently stable, not necessary for a commercially-focused product, or provide duplicate capabilities. (For example, Fedora may include half a dozen web browsers each of which provides different quality and features. Only the best one or two will be selected for inclusion in Red Hat Enterprise Linux.)

Red Hat Enterprise Linux releases are provided approximately every 18 months and supported by Red Hat and its partners for seven years. During this time, APIs/ABIs are maintained stable so that applications continue to work for the life of the product. It is the stability offered by Red Hat Enterprise Linux that makes it practical for ISV/OEM partners to certify their products with it.

During the extended release cycle Red Hat:

- Works closely with partners and customers to ensure that the features and technologies they require are included (for example: database support features, performance features, I/O support and device drivers, etc).
- Performs extensive quality assurance testing with formal Alpha/Beta programs.
- Performs necessary internationalization, including translations.
- Develops additional (multi-lingual) documentation.
- Builds products for the required system architectures.
- Ensures that features required for necessary standards certifications (security and applications such as NIAP/CC and ISO) are provided.
- Integrates technologies required by Red Hat' sayered products (for example, clustering).

Figure 1 shows the two stage distillation process from the community projects on the outside to Fedora as the unsupported, rapidly-changing vehicle for technology development to Red Hat Enterprise Linux as the stable, mature, commercially-focused distribution in the center.

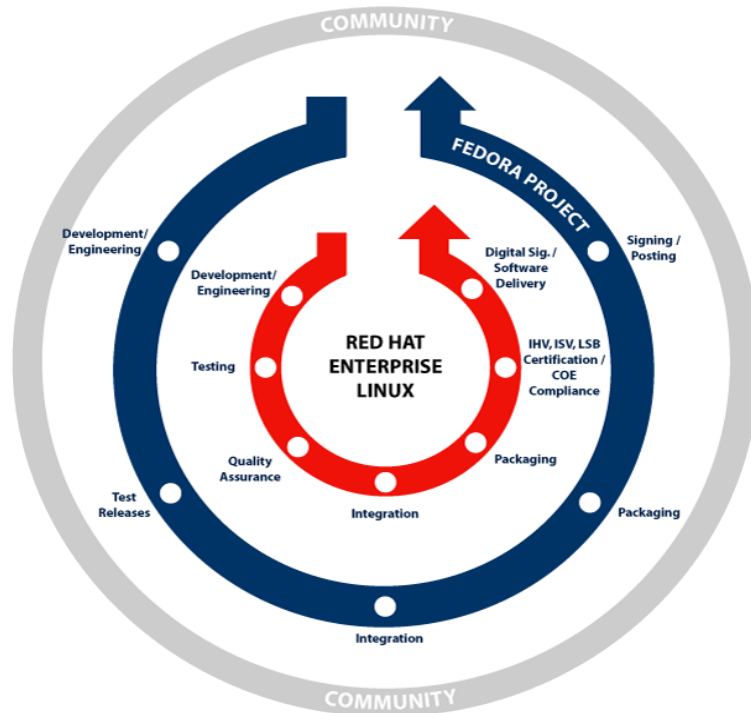


Figure 1: Distillation process from the Community to Red Hat Enterprise Linux

Red Hat Enterprise Linux Products

The Red Hat Enterprise Linux family has been designed to cover the full spectrum of corporate operating environments in a simple and consistent manner. The family is comprised of four products, two designed for server systems, two designed for client systems. There is a high level of commonality across the products, thereby ensuring that application support, user environments, and management tools are consistent. The products are primarily differentiated by the level of system architecture support, system size, and service offerings.

Red Hat Enterprise Linux supports multiple hardware architectures including:

- Intel x86-compatible (32-bit)
- Intel Itanium2 (64-bit)
- Advanced Micro Devices AMD64 (64-bit) and Intel EM64T
- IBM POWER series (eServer iSeries and eServer pSeries)
- IBM Mainframe (eServer zSeries and S/390)

Perhaps the most important feature of Red Hat's multi-architecture development process is that all implementations are built from identical source code. The primary benefit of this commonality is that all the products are completely compatible, regardless of architecture. This assists ISVs in supporting their applications on multiple architectures and also simplifies system administration and product support.



The individual members of the Enterprise Linux family remains unchanged from version 3:

- High-end server: Red Hat Enterprise Linux AS
- Entry/mid-level server: Red Hat Enterprise Linux ES
- High-end client: Red Hat Enterprise Linux WS
- General purpose client: Red Hat Desktop

An important feature of the family is that it is cleanly subsetted. That is, all the features of a low-end product are also available in a high-end product. Therefore, upgrades from one family member to another do not result in the loss of features, and server products can be deployed in client environments.

The following sections outline the major features of each Red Hat Enterprise Linux family member.

Red Hat Enterprise Linux AS

Red Hat Enterprise Linux AS (“advanced server”) is the top-of-the-line enterprise Linux solution, designed for large departmental and datacenter server deployments. Red Hat Enterprise Linux AS is the only family member that supports IBM POWER and zSeries/S-390 systems and is available with Standard and Premium Edition support. Red Hat Enterprise Linux AS is best suited for systems with more than 2 CPUs or more than 16 GB of main memory.

Typical Red Hat Enterprise Linux AS deployments would be used to support:

- Medium to large-scale databases and database applications
- Large web and application servers
- Corporate applications such as CRM, ERP, and SCM

Red Hat Enterprise Linux ES

Red Hat Enterprise Linux ES (“entry/mid server”) provides an entry-level and mid-range server operating system for the Intel x86, EM64T, Itanium2, and AMD64 markets. It supports 1-2 CPU systems with up to 16 GB of memory and is suitable for a wide range of applications—ranging from the edge-of-network to medium scale departmental deployments. It includes the same capabilities as Red Hat Enterprise Linux AS and is differentiated by its support for smaller systems and lower price. Enterprise Linux ES is available with Basic Edition and Standard Edition support.

Typical Red Hat Enterprise Linux ES deployments are used to support:

- Corporate web infrastructures
- Edge-of-network applications (DHCP, DNS, firewalls, etc.)
- Mail and file/print serving
- Small-medium database and departmental applications

Red Hat Enterprise Linux WS

Red Hat Enterprise Linux WS (“workstation”) is the high-end desktop/client



member of the Red Hat Enterprise Linux family.

Red Hat Enterprise Linux WS supports 1-2 CPU 32-bit and 64-bit Intel and AMD systems (x86, EM64T, Itanium2, and AMD64), and is ideal for “power user,” software development, and technical applications such as virtualization/rendering (CAD/CAM, EDA, etc.). It includes a full suite of desktop productivity applications for tasks such as document creation, email, instant messaging, and web browsing.

While Red Hat Enterprise Linux WS is based on the same software core as the server products, it does not include a number of network server applications (such as DNS and DHCP). Therefore it is suitable only for use in client environments. Enterprise Linux WS is available with Basic Edition and Standard Edition support.

HPC with Red Hat Enterprise Linux WS

Red Hat Enterprise Linux WS is usually the most cost effective Enterprise Linux product for use in High Performance Computing (HPC) environments. In these environments it is deployed in a headless workstation mode without a monitor, keyboard or mouse. A few common HPC-related packages are included in the Enterprise Linux family such as PVM and LAM.

Red Hat Desktop

Red Hat Desktop is the high-volume desktop/client member of the Red Hat Enterprise Linux family. It supports 32-bit Intel x86 and 64-bit Intel EM64T and AMD64 systems with one CPU and up to 4 GB of main memory. It provides the same software functionality as Red Hat Enterprise Linux WS but for smaller systems and at a lower price point. Red Hat Desktop is provided in multi-unit packages bundled with a Red Hat Network (RHN) Proxy or Satellite Server. The RHN server is used to efficiently perform desktop management functions such as the installation of updates and security patches.



Product Summary

Table 1: Summary of the Features of the Red Hat Enterprise Linux family

Feature	Red Hat Enterprise Linux AS	Red Hat Enterprise Linux ES	Red Hat Enterprise Linux WS	Red Hat Desktop
Supports Intel x86, EM64T, and AMD64 systems	Yes	Yes	Yes	Yes
Supports Intel Itanium2 systems	Yes	Yes	Yes	No
Supports IBM POWER S/390 & zSeries systems	Yes	No	No	No
Maximum CPUs supported ¹	.. ²	2	2	1
Maximum memory supported	-	16 GB	-	4 GB
Subscription to Red Hat Network	1 year	1 year	1 year	1 year
12x5 services available	Yes	Yes	Yes	N/A ³
24x7 services available	Yes	No	No	N/A
Includes desktop applications	Yes	Yes	Yes	Yes
Includes network server applications (e.g.: dhcp; dns)	Yes	Yes	No	No
Supported by leading ISV applications	Yes	Yes	Yes	Yes

¹ A processor chip with multi-core or hyper-threaded processing elements is counted as one CPU

² There is no subscription support limit, although a maximum may be imposed by hardware, software, or architectural limitations. Refer to www.redhat.com for specific details.

³ Offered with 24x7 Help Desk Escalation Support; Red Hat Network Proxy Server provided with Premium Edition support.



Example Configuration

Figure 2 shows a typical commercial intranet deployment with many small/medium servers, several high-end servers, and a High Performance Computing (HPC) compute farm.

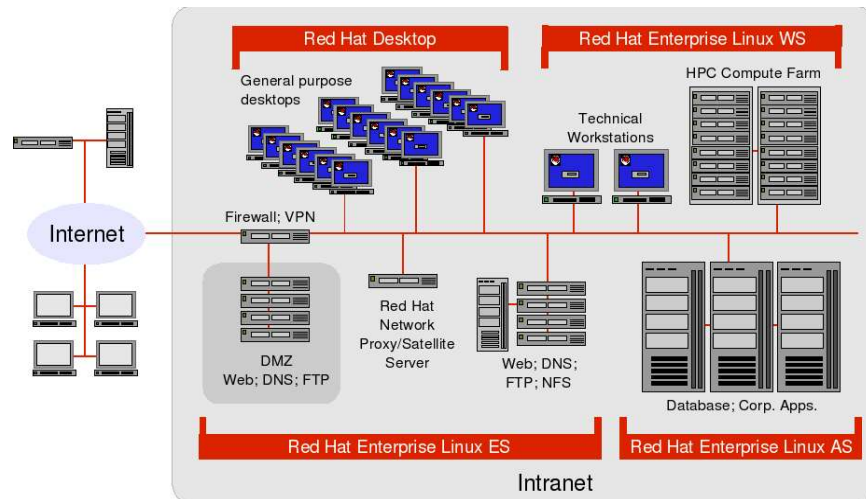


Figure 2: Typical Commercial Intranet Deployment.

The graphic shows how Red Hat Enterprise Linux family products can be deployed across a corporate IT infrastructure. Red Hat Enterprise Linux ES proves ideal for providing network services such as web servers, mail servers, file/print servers, and background network management services such as DHCP and DNS. Meanwhile Red Hat Enterprise Linux AS is used to host large-scale server applications and corporate databases. Red Hat Enterprise Linux WS is used for technical or development workstations and is also suitable for an HPC compute farm for services such as data-mining or financial modeling. Lastly, Red Hat Desktop meets the needs of the general purpose desktop user. Note that the entire environment can be provisioned, updated, and managed using the Red Hat Network Proxy Server that is included in the configuration.

Technical Features

A primary feature of Red Hat Enterprise Linux products is that they include technologies and features that provide a premier enterprise-quality computing environment. Features are selected on the basis of their appropriateness for commercial deployment (such as support for large SMP systems) and must also exhibit a high degree of reliability. This is significantly different from most Linux distributions where the focus is usually on providing the latest features as soon as possible (often at the expense of stability) and concentrating on serving low-end markets.

Red Hat Enterprise Linux v.4 was developed in close collaboration with Red Hat's major customers and ISV/OEM partners to ensure that it provides the



features they require. Development occurred over an 18 month period with almost six months dedicated to beta testing.

The kernel for Red Hat Enterprise Linux v.4 is based on the Linux 2.6.9 kernel. While many of the major features provided by the 2.6 kernel were back-ported and included in Red Hat Enterprise Linux v.3 (which was released in October 2003, based on the Linux 2.4.21 kernel), further development of these features during 2004 provides the v.4 product with additional performance and scalability.

The new kernel offers a large selection of new features, and it is beyond the scope of this paper to describe them all. However, a brief overview of a few of the latest features provides a general insight into areas of specific development and also illustrates the level of sophistication achieved by the latest Linux kernels.

Read Copy Update (RCU)

This feature provides improved performance for kernel algorithms that manipulate “read-mostly” lists. That is, lists that are generally read but with occasional writes. Examples include the Network Routing and Dentry caches. Prior to RCU, routines that traversed these lists needed to lock them from other accessors to ensure that consistency was maintained in the rare event of a list change. This prevented multiple readers from accessing the list concurrently, despite the fact that on most occasions it was safe to do so. This restricted performance in SMP systems. With RCU, multiple readers are permitted while a lock is used to ensure that there is only a single writer. List modification is carefully implemented so that a structure that is, for example, being removed from a list, is unlinked but not deallocated (essentially, it is “copied”). Any active reader(s) can continue to access the structure, while for new readers it will not be accessible. A background thread deallocates the unlinked structures when the active readers have completed their tasks. This technique permits concurrent readers, thereby improving performance while allowing writers to operate in a fully coordinated manner. Figure 3 illustrates this feature.




Figure 3: Read Copy Update (RCU) Feature

Selectable I/O elevators

Red Hat Enterprise Linux v.4 provides a number of I/O elevators that can be selected at boot time depending on the specific application environment. An I/O elevator is used to modify the order in which I/O is issued to improve the throughput or latency of the I/O subsystem. Four elevators are provided:



- NOOP scheduler. As the name suggests, this scheduler provides no I/O reordering. It is typically used in virtual system environments where the underlying host I/O subsystem will implement whichever I/O elevator is most appropriate.
- Completely Fair Queuing (CFQ) scheduler. This is the default scheduler in Red Hat Enterprise Linux v.4. It provides complete fairness by implementing a per-process I/O queue. The I/O scheduler removes one I/O from each process' queue on a round-robin basis. This ensure that each process can issue an equivalent (fair) number of I/Os.
- Deadline scheduler. This scheduler provides a per-I/O request deadline to ensure that starvation does not occur for processes that are issuing very large numbers of I/Os. This is possibly the most appropriate scheduler for databases systems, which often have centralized writer processes that issue very large numbers of write I/Os.
- Anticipatory scheduler (AS). This scheduler is possibly the most appropriate for interactive systems. It attempts to anticipate the next I/O request based on the heuristic that read I/Os tend to be synchronous and sequential while write I/Os tend to asynchronous and random. This can lead to the I/O system queuing up many write I/Os but only receiving new read I/Os when the previous read completes. As a result, when a read completes and the I/O system issues the next I/O, it is a write. To service the write, the disk heads are almost certainly required to move to another location on the disk, a process that will take 5-8mS (a seek plus the disk rotational delay). Meanwhile the reading process will usually issue another read, typically at the next sequential location on the disk. The AS scheduler will attempt to optimize this situation by delaying the issuing of pending writes at the end of a read I/O by approximately one millisecond in the anticipation of another sequential read being issued. If the read is requested it can be honored without any need for an intermediate disk seek. If a read is not issued, the queued write can be started. The cost of delaying the write is small, while the benefit to the reader will be 10-16mS (eliminating the two seeks and rotational delays caused by an off-track write).

Object-Based Reverse Map VM

Red Hat Enterprise Linux v.3 included a Reverse Map VM (Virtual Memory) feature, developed by Red Hat, which is used to locate all the process virtual addresses that map to a given physical address. This is needed when performing operations such as swapping. Without a Reverse Map VM capability, physical to virtual address translation is slow and cumbersome and significantly impacts the performance of large or memory constrained systems. The Reverse Map VM capability in Red Hat Enterprise Linux v.3 created additional memory management structures to perform the reverse translation. This provided a significant Reverse Mapping performance improvement but imposed an overhead on all systems, even those that were not memory constrained (it was high cost, high gain). During 2004 the algorithms used by Reverse Map VM were further enhanced to eliminate the additional structures and use existing memory object structures (file, process, etc). This resulted in an equivalent performance improvement but at minimal additional overhead (low cost, high gain).



Generic logical CPU scheduling

Red Hat Enterprise Linux v.3 included the O(1) scheduler back-ported from the Linux 2.5/2.6 kernel and further enhanced it by implementing support for logical, or hyper-threaded, CPUs. The standard scheduler would treat every CPU as equal and created a per-CPU compute queue. This could result in a pair of processes contending for silicon resources by being scheduled on the same hyper-threaded CPU pair, while another CPU chip was idle. The Red Hat Enterprise Linux v.3 kernel resolved this problem by creating per-hyper-thread-pair compute queues so that processes were scheduled across CPU chips prior to hyper-threaded processing elements. In Red Hat Enterprise Linux v.4 this feature has been further developed to handle the forthcoming multi-core processors. The scheduler will create compute queues correctly, based on individual CPU chips, their multiple cores, and their hyper-thread capabilities.

Block I/O subsystem

Red Hat Enterprise Linux v.2.1 and v.3 included a number of I/O features that were back-ported from the Linux 2.5/2.6 kernel. These included:

- Asynchronous I/O
- Huge Translation Buffer File System (TLBfs)
- Bounce buffer elimination
- Remap_file_pages
- O_Direct

Collectively, these features allowed significant performance improvements over the standard Linux 2.4 kernel. With the Linux 2.6 kernel they were incorporated into a completely new block I/O subsystem that also provides additional I/O scalability improvements. The new subsystem allows a larger number of I/O devices and larger filesystems to be configured. As a result Red Hat Enterprise Linux v.4 supports very large SCSI and Fibre Channel configurations, and the ext3 file system scales to 8 TB.

Other I/O enhancements include:

- Support for SATA (Serial ATA) devices. SATA is the next generation interconnect for embedded storage in low-end systems. It provides higher performance than traditional ATA devices (with a 150MB/sec transfer rate) at lower cost.
- Tagged command queuing. This feature allows multiple I/Os to be sent to a storage controller in parallel so that it can optimize how the I/Os are performed. This feature can provide noticeable performance improvement for heavy I/O loads.

Sys_epoll() support

Sys_epoll is an important new system call in the Linux kernel which provides a high efficiency polling mechanism for applications that need to wait on events that are occurring on many (potentially thousands) of file descriptors (typically, network I/O channels). With sys_epoll it is possible to eliminate heavily repeated select() and poll() calls. For networked applications this call can



result in significant performance improvement.

Support for larger server systems

For x86 systems, up to 32 logical CPUs (16 hyper-threaded CPU pairs) are supported. With Itanium2, systems with up to 64 CPUs are supported.

Upward Compatibility

An important feature of the Enterprise Linux v.4 family is that it provides forward compatibility for existing Enterprise Linux v.2.1 and v.3 systems. Compatibility libraries for v.2.1 and v.3 are included so that it is possible to run applications from these versions without rebuilding. Of course, rebuilding an application will usually result in higher performance as it will benefit from numerous improvements in the GCC compiler.

File System Performance enhancements

Red Hat Enterprise Linux v.4 includes a number of performance enhancements to its default filesystem, ext3. These include:

- Block reservations (space preallocation), which greatly improve read/write performance. (See Figure 4).
- Large directories are implemented using hash trees, resulting in much faster directory scans.
- On-demand expansion of mounted filesystems.
- Increased performance in SMP systems through synchronization (locking) improvements.

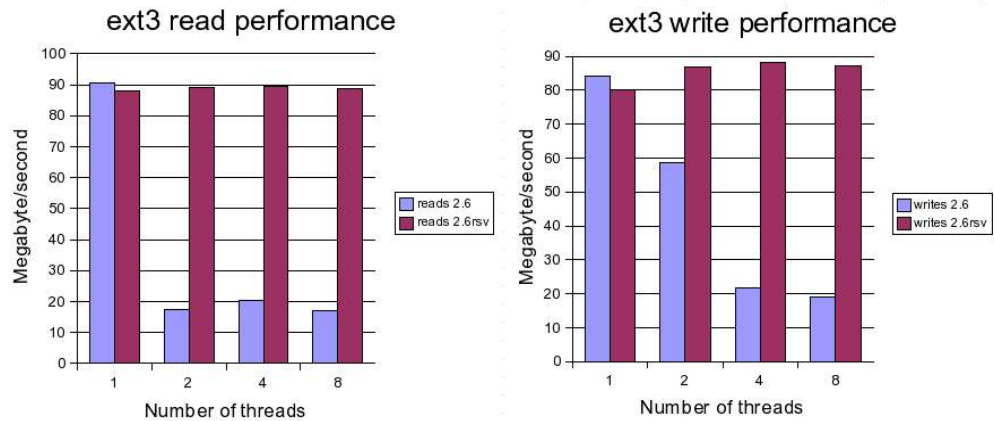


Figure 4: I/O bandwidth increase provided by block reservations (rsv) over the original Linux 2.6 ext3 filesystem.

Red Hat Desktop

The first release of Red Hat Desktop was delivered in mid-2004 and focused on providing an easily-managed and highly secure environment for multi-unit deployments (tens to hundreds of clients). Designed for customers who



centralize desktop management through their IT departments and help desks, Red Hat Desktop is typically sold as a complete solution, bundled with a Red Hat Network Proxy or Satellite Server. The Red Hat Network Proxy or Satellite Server is used to perform management tasks.

Meanwhile, Linux desktop technology continues to develop rapidly and Red Hat Desktop v.4 provides a wide variety of new features including:

- The GNOME desktop is updated to version 2.8 (from 2.2 in Red Hat Enterprise Linux v.3). Version 2.8 provides many new features such as support for plug-and-play devices (through a new Hardware Abstraction Layer and support for D-BUS), enhanced file management, and network and printer management tools.
- Inclusion of Firefox as the default web browser. Firefox is a high-performance, secure, and easily-extendable web browser. It is rapidly establishing itself as the leading alternative to Internet Explorer.
- Evolution 2.0 groupware client. Evolution provides robust email, calendaring, and contact management capabilities. It supports standards such as IMAP, POP, SMTP, LDAP, and iCalendar, interoperability with Microsoft Exchange Server, and certificate management.
- OpenOffice, the Office productivity suite included with Red Hat Desktop, has been upgraded to the latest version.
- Significant improvements in the handling of multimedia are included with HelixPlayer and RealPlayer 10 offering SMIL, MP3, Flash, and RealAudio/RealVideo support. RhythmBox provides complete music management capabilities.
- Numerous other desktop applications have been updated or included for the first time such as GAIM instant messenger, Planner project management, The GIMP v.2 image composition and editing tool, and Rdesktop RDP terminal services client.
- Cross platform interoperability has also been improved. For example, Microsoft Active Directory can be used for user login authentication, and it is possible to authenticate web-based applications with NTLM. Windows SMB file and print shares can be easily browsed from the standard desktop environment.
- Vino provides a VNC-based desktop sharing capability, which is ideal for collaboration or for use by an IT help desk when diagnosing user problems.
- As with Red Hat Desktop v.3, the new release provides a collection of third-party applications, such as Adobe Reader, Macromedia Flash, and the Citrix ICA Client. Java runtime environments from IBM and BEA are also available. Optional commercial fonts, licensed from Agfa/Monotype, improve document display quality, especially for documents that are migrated from other platforms.

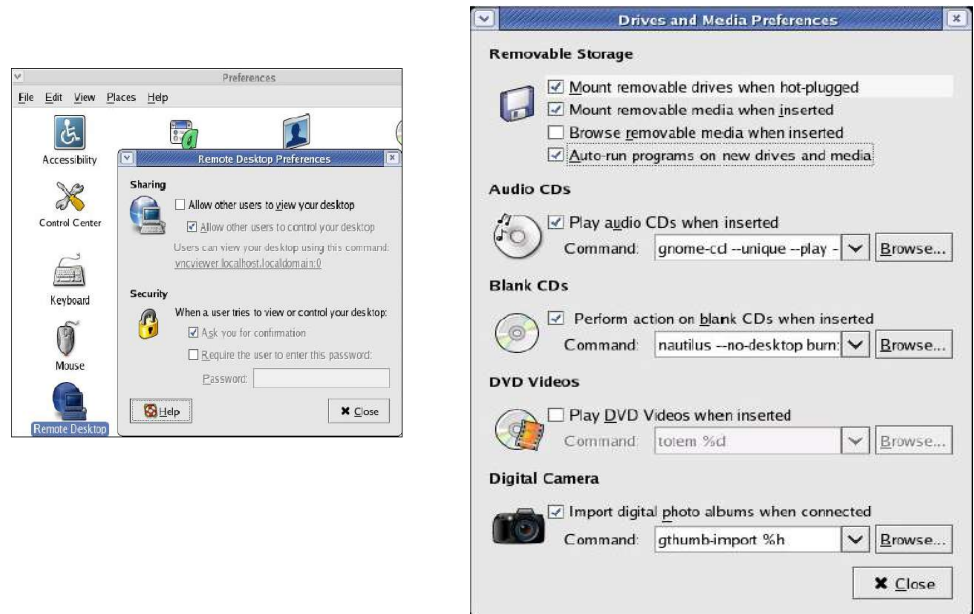


Figure 5: Typical Red Hat Desktop System Management Applications

Security

Security is a major focus of the Red Hat Enterprise Linux v.4 release. The most important new security feature is the inclusion of Security-Enhanced Linux (SELinux). This feature, developed by the US Government NSA (National Security Agency), provides a Mandatory Access Control (MAC) environment for all Red Hat Enterprise Linux systems. MAC security operates in tandem with the existing Linux security infrastructure, which provides the traditional Discretionary Access Control (DAC) environment. MAC improves the security capabilities of the system through a Security Policy that is imposed by the kernel and Role Based Access Control (RBAC).

In a traditional DAC environment, security is achieved by ensuring that applications are carefully configured and do not contain exploitable flaws. In the event that an application is compromised, it is often possible for it to damage the entire system.

In a MAC environment, a set of policy rules define what an application is permitted to do, and the kernel ensures that the rules are enforced. As a result, even a badly compromised application cannot damage the entire system. Figures 6 and 7 illustrate access control in SELinux.

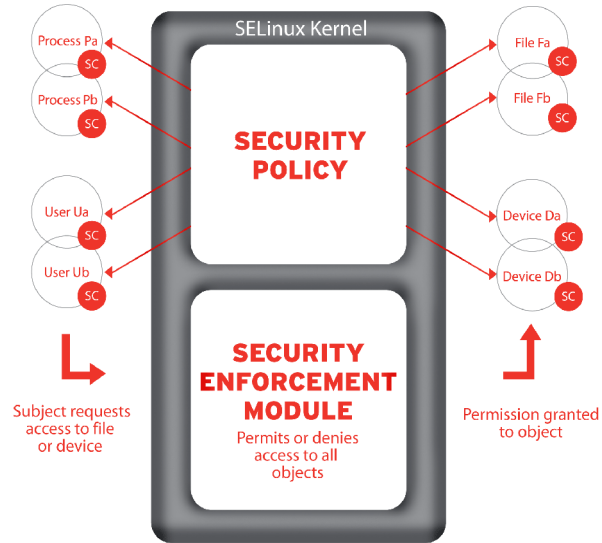
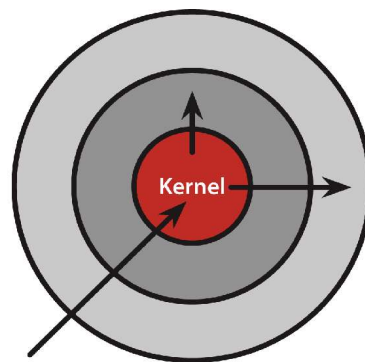
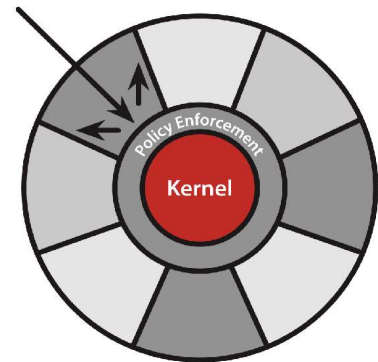


Figure 6: SELinux Access Control Mechanism



Discretionary Access Control

Once a security exploit gains access to privileged system component, the entire system is compromised.



Mandatory Access Control

Kernel policy defines application rights, firewalling applications from compromising the entire system.

Figure 7: Difference between Discretionary Access Control and Mandatory Access Control environments

It is worth noting that all the security capabilities provided by Red Hat Enterprise Linux v.3 are carried forward to the v.4 product. These include:

- File system ACL (Access Control List) support
- Position Independent Executables
- ExecShield features:
 - Support for Intel XD (eXecute Disable) and AMD NX (No eXecute) processor features
 - Support for Intel x86 Application Segmentation



Auditing

Red Hat Enterprise Linux v.4 includes a new auditing capability, “audit,” that replaces the existing LAuS feature. Audit, developed by Red Hat, has been accepted into the upstream kernel and provides an elegant, generalized capability that can audit SELinux and standard Linux events. Several reporting tools are provided, and audit also includes a bidirectional socket interface that enables other applications to interface to it (for example, snare and trace packages).⁴

Compiler and Library Buffer Management

In late 2004, Red Hat developed a new group of features that improve buffer management and security for inclusion in Red Hat Enterprise Linux v.4. At the time of writing, these features are unique to Red Hat environments.

Advanced GLIBC memory corruption checks

The GLIBC memory allocator functions now perform a set of internal sanity checks to detect double freeing of memory and heap buffer overflows. With these checks, regular application bugs and security exploit attempts that use these techniques are detected, and the program will be instantly aborted to avoid the possibility of the exploit succeeding. With these checks, double free exploits become entirely impossible, and all standard, generic heap overflow techniques are blocked.

Printf format string exploit prevention

Printf format string exploits abuse a bug in programs that have a faulty call to the standard printf() function, caused by a very rarely used formatting parameter. The printf function is now able to check that this rare formatting comes from guaranteed trusted sources and will abort the program if that is not the case, thus preventing printf format exploits entirely.

GCC buffer bound checking

An enhancement has been added to the GCC compiler such that if the size of the destination buffer can be detected at compile time, functions such as strcpy(), memcpy(), strcat() will use a checking variant of these functions that detects if the buffer will actually overflow. If that happens, the program is aborted immediately. While gcc cannot always detect the size of the destination buffer (for example, it is not possible for dynamically allocated buffers), buffer allocation errors usually occur with the types of buffer that can be detected by gcc. The result is that a large percentage of buffer overflow errors are prevented immediately.

Standards Compliance

Red Hat works closely with many industry standards groups to ensure the widest possible standards support. Red Hat Enterprise Linux v.4 is expected to complete NIAP/CC EAL 4+ (National Information Assurance Partnership; Common Criteria; Evaluation Assurance Level) certification shortly after initial release. Furthermore, to ensure easy migration of applications across Linux environments, Red Hat Enterprise Linux v.4 is designed to be Linux Standard

⁴ Audit will be available for Red Hat Enterprise Linux v.4 in the first half of 2005.



Base Runtime Environment 3 compliant. Refer to <http://www.linuxbase.org/> for information on the LSB specification.

Development Environment

Red Hat Enterprise Linux v.4 includes GCC 3.4, the latest stable development environment for application developers. Also included is a preview edition of the GCC 4.0 tool chain. GCC 3.4 provides many new features including significantly enhanced code generation, which results in improved application performance. These GCC environments provide development support for C, C++, and Fortran 95.

Storage Subsystem

To improve support for large storage subsystems, Red Hat Enterprise Linux v.4 includes LVM2 (Logical Volume Manager 2). This feature permits multiple storage devices to be combined and controlled with maximum flexibility. Storage allocation can be managed to meet application needs rather than being reliant on the underlying physical storage, and operations such as dynamically increasing the size of a filesystem are supported.

LVM2 provides numerous improvements over LVM1, which was included in Red Hat Enterprise Linux v.3. Significant redesign work has resulted in a much more stable and robust implementation with transactional metadata updates, read/write snapshots, improved storage management tools, and a host of other features. An LVM2 setup phase is incorporated into the Red Hat Enterprise Linux installation procedure (Anaconda), so that logical volumes can be configured during initial installation. Figure 8 provides a view of the new storage management GUI included with LVM2.

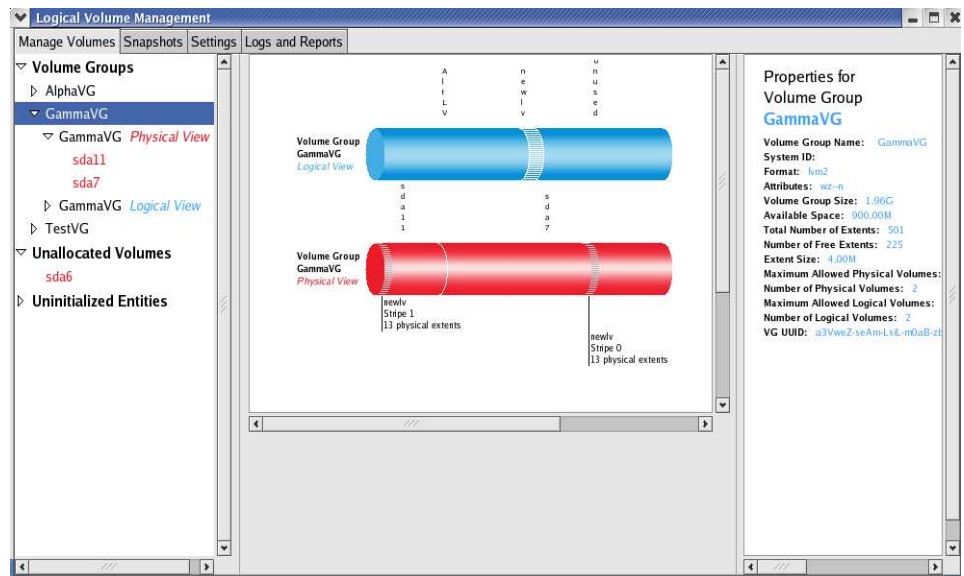


Figure 8: Storage Management GUI



Shortly after the release of Red Hat Enterprise Linux v.4, an update to LVM2 will provide support for Mirroring (RAID1). Additionally a new multi-pathing feature, "multipath," is being developed that will eventually replace the existing MD multipath driver.

A major feature of the new implementation is the clean separation between user level and kernel level functions. Kernel level functions have been encapsulated in the new Device Mapper module, which provides a generic device access layer. This is used by the user level LVM modules and also by third-party user level applications (such as IBM's *zEVMS* storage management software). Device Mapper is a Red Hat project that has been accepted into the upstream kernel. It provides a highly flexible, pluggable interface for features such as concatenation, striping, mirroring, encryption, etc.

Automounter

Red Hat Enterprise Linux v.4 includes the *autofs4* automatic device mounter. This will automatically mount filesystems as soon as a user touches them (for example, with an *ls* or *cd* command) and dismount them after a selectable idle period. The new automounter provides functionality very similar to that provided in Sun's Solaris operating system, such as multi-mounts, browsable mounts, replicated servers, and executable maps.

Networking

Numerous new networking features are provided in Red Hat Enterprise Linux v.4, including:

- Support for Network Interrupt Mitigation (referred to as NAPI, for New API). This feature combines device interrupt handling and polling to optimize the performance of heavily loaded networks. Rather than allow a network device to trigger an interrupt for every arriving packet, NAPI disables interrupts as soon as a packet is delivered. The network handler then enters a polling mode until all pending network packets are drained from the network device's receive buffers. When the last packet has been serviced, the routine then re-enables interrupts and exits normally. NAPI is most valuable for Gigabit Ethernet and other networks with high packet arrival rates.
- Support for SCTP (Stream Control Transmission Protocol). While Red Hat Enterprise Linux is primarily focused on the general commercial market, it is also suitable for use in specialized markets such as Telco. SCTP is a message-oriented, reliable transport protocol used in the Telco industry and is required by the CGL (Carrier Grade Linux) specification. SCTP provides numerous features such as multi-homing, ordered and unordered messaging, and congestion control.
- The inclusion of NFSv4 provides NFS environments with many new features such as improved performance and security, cross-platform interoperability, and full support for Windows file sharing.

Feature Summary

This list of features, though several pages long, is by no means a comprehensive summary of new features provided by Red Hat Enterprise Linux v.4. However, it demonstrates the scale and scope of the improvements.



Focusing on performance, scalability, security, application development, and standards support, Red Hat Enterprise Linux v.4 provides the world's leading enterprise-strength Linux computing environment.

Support Services

Red Hat Enterprise Linux and Red Hat Desktop solutions are sold on an annual subscription basis by Red Hat and its partners as complete packages that include the software product combined with a choice of maintenance/support services. Subscriptions are fully inclusive, providing all software updates, upgrades to new versions, unlimited users and network links (there are no Client Access Licenses), and unlimited support incidents. Three levels of maintenance and support are available as shown in Table 2.

Table 2: Maintenance and Support Levels

Support Type	Offering	Coverage	Red Hat Enterprise Linux AS/ES/WS
Basic	<ul style="list-style-type: none"> Red Hat Network 30 days Installation and Configuration support Includes Upgrades	<ul style="list-style-type: none"> 1 year 1 year; Mon-Fri 9-9 ET North America (Mon-Fri 9-5 rest of world); 1 business day response; unlimited incidents 	ES and WS only
Standard	<ul style="list-style-type: none"> Red Hat Network Software support Includes Upgrades	<ul style="list-style-type: none"> 1 year 1 year; Mon-Fri 9-9 ET North America (Mon-Fri 9-5 rest of world); 4 hour response; unlimited incidents 	AS, ES, and WS
Premium	<ul style="list-style-type: none"> Red Hat Network Software support Includes Upgrades	<ul style="list-style-type: none"> 1 year 1 year; 24x7; 1 hour response; unlimited incidents 	AS only

Maintenance services are delivered using Red Hat Network, which provides Internet-based access to updates, upgrades, patches, and errata released by Red Hat. Note that Red Hat Desktop is sold as a bundled solution with Red Hat Network Proxy Server and includes a year of Help Desk Escalation Support.

Additionally, a comprehensive selection of Professional and Consulting Services is available from Red Hat and its partners. These services are designed to help customers design, configure, deploy, and manage Red Hat



Enterprise Linux systems. Migration services from proprietary UNIX and Windows systems are also available, as are custom engineering and application development services.

Red Hat Network

Red Hat Network (RHN) is Red Hat's internet-based system maintenance and management infrastructure. RHN's core capability is to analyze all the packages on a Red Hat Enterprise Linux system and identify packages for which updates are available and resolving any dependencies that the packages require. RHN can then apply the updates as required. RHN's strength lies in being able to manage hundreds of systems automatically with features such as system grouping, automatic updates, continuous monitoring and alerts. Using RHN greatly simplifies the process of keeping multiple systems up-to-date and secure. As Red Hat continuously provides enhancements and security updates, RHN provides the mechanism to ensure that they are applied to customer systems quickly and efficiently.

All Red Hat Enterprise Linux subscriptions provide one year of Red Hat Network access, and in most cases RHN will be used to download the purchased Enterprise Linux variant and any layered products.

RHN provides several modes of operation: Hosted, Proxy, and Satellite as shown in Figures 9 and 10:

- In Hosted mode each managed system connects across the Internet to an RHN server hosted by Red Hat. The RHN server will inventory the system's packages and apply updates as required.

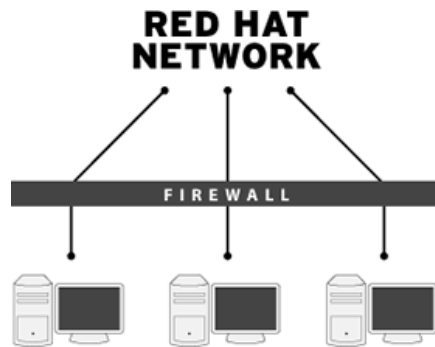


Figure 9: RHN Hosted Model

- Proxy mode is useful for reducing Internet bandwidth consumption when a customer has multiple systems. Updates are cached at the customer's site in a local Proxy server, so the customer only needs to download updates from Red Hat RHN servers once. Package inventorying is still performed by Red Hat's RHN servers, allowing Red Hat Network to notify administrators immediately if an update is available.




Figure 10: Red Hat Network Proxy Model

- Satellite mode provides a fully disconnected RHN environment, where package inventorying and caching is all performed using a Satellite Server located on the customer' site. Updates from Red Hat are performed on an as-needed basis. A Satellite server can also be used to deliver and update customer-specific applications.




Figure 11: Red Hat Network Satellite Model

Red Hat Network provides multiple management modules: Update, Management, and Provisioning. Update provides the basic ability to keep systems up-to-date and is included in all Red Hat Enterprise Linux subscriptions. Management provides additional features such as system grouping and role-based permissions. Provisioning provides bare-metal system deployment capabilities (allowing new systems to be up-and-running in the shortest possible time) and multi-state rollback. For additional information on Red Hat Network please refer to <http://rhn.redhat.com/>.

Application Availability

As described earlier, Red Hat Enterprise Linux has been enthusiastically and widely adopted by leading ISVs. At the time of writing, over 1000 applications have been certified by more than 300 ISVs. The complete list of certified



applications can be found in the online Software Catalog at www.redhat.com/apps/isv_catalog/.

A brief summary of applications supported is shown in Table 3.

Table 3: Supported applications

ISV	Supported Applications
Oracle	9i, 9iRAC, 9i Application Server, 8i, 11i eBusiness Suite, Oracle Collaboration Suite
VERITAS	Foundation Suite, VCS, NetBackup Business Server, Datacenter
BEA	WebLogic Server 7.0, WebLogic Jrocket
IBM	WebSphere v5, DB2 v8.1, various Tivoli products, Directory Server 5.1
Reuters	Reuters Market Data System
Sybase	ASE for Linux, SQL Anywhere Studio for Linux, Sybase IQ, Sybase Replication Server

Other key products from BMC, CA, Documentum, EMC, Hyperion, Legato, PeopleSoft, Rogue Wave, VMWare, and many others are also supported.

A critical feature of the Red Hat Enterprise Linux family is that ISVs only need to certify their applications once. Because all family members share a common core (kernel, development tool-chain, libraries, and so on) their application environments are identical. This means application certification on a single family member accrues to all members. (Note, however, that individual ISVs may choose to not support their applications on specific Enterprise Linux products. For example, it is usually impractical to support a large server application on a small client system).

Hardware Availability

In addition to being certified for use with a wide range of applications, Red Hat Enterprise Linux is also certified for use on a wide range of hardware platforms. Leading OEM vendors, including Bull, Dell, Fujitsu, Hitachi, HP, and IBM, verify that their server and desktop systems work correctly with Red Hat products. Several OEM partners offer Red Hat Enterprise Linux pre-loaded on selected systems. In late 2004 over 600 certified systems were listed on Red Hat's Hardware Compatibility List, at www.redhat.com/hardware/, the most extensive coverage offered by any Linux vendor. Certification also extends to peripheral hardware such as storage, network and graphics adapters, and controllers from leading vendors.



Benchmarks

For many IT organizations the availability of industry-standard, audited benchmarks is an important component in the purchase decision process. Red Hat Enterprise Linux has achieved numerous impressive benchmark results, including a number of world records. These benchmark results are documented at their associated websites and at www.redhat.com/software/rhel/benchmarks/. Benchmark results can be compared directly at the websites maintained by the relevant auditing groups:

- www.tpc.org
- www.spec.org
- www.oracle.com/benchmarks
- www.sap.com/benchmark

At the time of writing the following benchmark results are of note:

TPC-H on HP ProLiant with AMD Opteron processors

As of December 13, 2004 this benchmark is the world record performance result in the 1000 GB database category. The configuration was comprised of a cluster of eight HP ProLiant DL585s each with four AMD Opteron processors. A storage area network of HP SAN switches and HP MSA 1000 storage controllers was used. An InfiniCon InfinIO 3016 (Infiniband) switch was used as the cluster interconnect. Software used was Red Hat Enterprise Linux AS with Oracle 10g Database with Real Application Clusters (RAC) and Partitioning. The benchmark audited result was 35141 QphH@1000GB at 60\$/QphH@1000GB. This result also holds the #5 place in the price/performance rankings, being 15% more expensive than the #1 result but offering almost 12x the performance.

SPECweb99 on IBM eServer p5 570 with Red Hat Content Accelerator 3.2

This benchmark result, announced in July 2004, used an IBM eServer p5 Model 570 with Red Hat Enterprise Linux. The server was configured with 2 dual-core chips and 64 GB of main memory. The test result of 13,500 holds top place for quad CPU systems and delivered 30% more performance than the number two quad CPU result.

SPECweb99 on HP ProLiant DL580 G2 with Red Hat Content Accelerator

This benchmark result, announced in August 2004, used an HP ProLiant DL580G2 running Red Hat Enterprise Linux. The server was configured with four 3 GHz CPUs and 32 GB of main memory. The test result of 8,500 holds top place for Intel Xeon systems.



SPEComp2001 on IBM eServer p5/520

This benchmark result, announced in July 2004, used an IBM eServer p5 Model 520 with Red Hat Enterprise Linux v.3. The server was configured with a single, dual-core CPU configured with 32 GB of main memory. The test result of 5287 peak and 4758 base holds the top place for a dual CPU system, almost double the performance of the next 2 CPU system result.

TPC/C on HP Itanium cluster with Oracle 10g RAC

In December 2003, Red Hat, Oracle, and HP announced the overall performance world record result for TPC/C, based on a 16 node HP Integrity cluster running Oracle 10G RAC with Red Hat Enterprise Linux v.3 AS. The benchmark used a database of approximately 90 Terabytes, deployed on a storage subsystem that included over 2000 disks. The result was approximately 50% faster than the fastest Microsoft Windows results and 98% faster than the fastest Sun Solaris result. The benchmark also produced the lowest \$/tpmC cost of the top 10 performance-related results. This result remained the leader for almost a year before being finally displaced by a system that was almost three times the size. So, currently the result is the #2 record holder.

Layered Products for Red Hat Enterprise Linux

With Red Hat Enterprise Linux established as the leading Linux operating system for the commercial environment, Red Hat is working to extend the benefits of open source software further up the software solution stack to the middle-ware and application layers. Red Hat's strategy is to provide a set of optional layered products that can be used to enhance the standard Red Hat Enterprise Linux system. The following products, offered with full maintenance and support services, are available today:

Red Hat Global File System

Red Hat Global File System (GFS) is an open source, POSIX-compliant cluster filesystem and volume manager that executes on Red Hat Enterprise Linux servers attached to a storage area network (SAN). It works on all major server and storage platforms supported by Red Hat. The leading (and first) cluster file system for Linux, Red Hat GFS has the most complete feature set, widest industry adoption, broadest application support, and best price/performance of any Linux cluster file system today.

Red Hat GFS allows multiple Red Hat Enterprise Linux servers to simultaneously read and write a single shared filesystem on the SAN, achieving high performance and reducing the complexity and overhead of managing redundant data copies. Red Hat GFS has no single point of failure, is incrementally scalable from one to hundreds of Red Hat Enterprise Linux servers, and works with all standard Linux applications.

Red Hat GFS is integrated with Red Hat Enterprise Linux and distributed through Red Hat Network. This simplifies software installation, updates, and management. Applications such as Oracle 9i RAC and workloads in cluster



computing, file, web, and email serving can become easier to manage and achieve higher throughput and availability with Red Hat GFS.

For Red Hat Enterprise Linux v.4 a new release of Red Hat Global File System, version 6.1, is being provided. This offers numerous scalability and performance enhancements, and, perhaps most importantly, the inclusion of cluster-aware Logical Volume Management capabilities. All the features mentioned in the earlier section on Logical Volume Management operate cluster wide when deployed with Red Hat GFS. Red Hat Cluster Suite is included with all Red Hat GFS subscriptions.

Red Hat Cluster Suite

Red Hat Cluster Suite allows server systems to be clustered in a high availability configuration. High availability clustering (sometimes referred to as “failover clustering”) is a technology widely used in commercial operating system environments; it allows standard applications to be available almost continuously with automatic recovery from hardware failures and shutdowns.

Red Hat Cluster Suite is described in detail in the whitepaper, “An Overview of Red Hat Cluster Suite.” Please refer to that paper for additional information.

For Red Hat Enterprise Linux v.4, a new release of Red Hat Cluster Suite, v.4, is being provided. The release delivers several new features including:

- Red Hat Cluster Suite provides core technologies required by itself and Red Hat GFS. These include Membership Management, I/O Fencing, Services Management, and Polling. Previously, these underlying cluster technologies were duplicated by the two products with the latest release they are integrated. This simplifies both products, making them easier to deploy and maintain.
- A Distributed Lock Manager (DLM). The DLM provides a cluster-wide synchronization service that can be used by any application. The DLM is modeled after capabilities provided by the VMS and Tru64 operating systems, offering a comprehensive feature set such as hierarchical lock trees, multiple concurrency (lock) modes, blocking and asynchronous notifications, lock value blocks, and range locking. The DLM enables application providers to create fully distributed applications.
- Elimination of the requirement for shared storage. Previous releases of Red Hat Cluster Suite required a shared storage subsystem (SCSI or SAN). With Red Hat Cluster Suite v.4 this requirement has been eliminated, greatly reducing initial configuration costs. For applications which are able to easily replicate data between servers or require read-only data, it is possible to configure a simple cluster that will provide greatly increased availability for minimal cost.

Comparing Red Hat Global File System and Red Hat Cluster Suite

The crucial difference between Red Hat Cluster Suite and Red Hat Global File System relates to their data access capabilities. Cluster Suite configurations permit a single node to access a given disk partition at a time. GFS configurations permit multiple nodes to concurrently access the same disk partition and files on the partition. Cluster Suite provides high availability by failing over applications between nodes. GFS provides high availability and



scalable performance by allowing distributed applications to execute in parallel on multiple nodes. As mentioned earlier, Cluster Suite failover capabilities are also included in GFS, permitting applications that do not support parallel operation to be deployed in high availability configurations.

Red Hat Global File System and Red Hat Cluster Suite can be deployed on Red Hat Enterprise Linux AS, ES, or WS systems using Intel x86/Itanium2/EM64T and AMD AMD64 systems.

Red Hat Cluster Suite

- No shared storage
- Application failover




Figure 12: Simple configuration of Red Hat Cluster Suite with no shared storage

Red Hat Cluster Suite

- Shared storage – block based(SAN)
- No shared access – partitions accessed on a per node basis
- Servers mount GFS partitions concurrently
- Failover allows a different server to remount partition after failure and restart the application




Figure 13 (left): Cluster Suite shared storage configuration with each server accessing different disks/partitions

Red Hat Global File System

- Shared storage – block based(SAN)
- Shared access
- Servers mount GFS partitions concurrently
- Server failure does not perturb other servers
- Failover available for single instance applications




Figure 14 (right): GFS shared storage configuration with each server accessing the same disk/partition.

Figures 12, 13, and 14 show the storage access provided by the two products. It can be seen that Red Hat Cluster Suite provides a simple, low-cost, high availability capability, while Red Hat Global File System provides a



sophisticated data sharing capability that can be scaled across multiple nodes, each adding to overall application throughput.

Red Hat Application Server

Announced in mid 2004, Red Hat Application Server provides a J2EE standards compliant, open source, application server environment. Although Red Hat works closely with leading application server vendors such as BEA, Oracle, and IBM to ensure that their products are certified for use on Red Hat Enterprise Linux, many customers requested Red Hat provide a product based on existing open source projects. Consequently, Red Hat Application Server incorporates Tomcat as the Web Container and JOnAS (from the ObjectWeb consortium) as the EJB Container. Other technologies are also included such as a Struts (a framework for building Java web applications) and database resource adapters for Oracle, DB2, PostgreSQL and MySQL.

These technologies are integrated into a single product with documentation and installation procedures and offered with a full support subscription that includes security errata, updates, and upgrades.

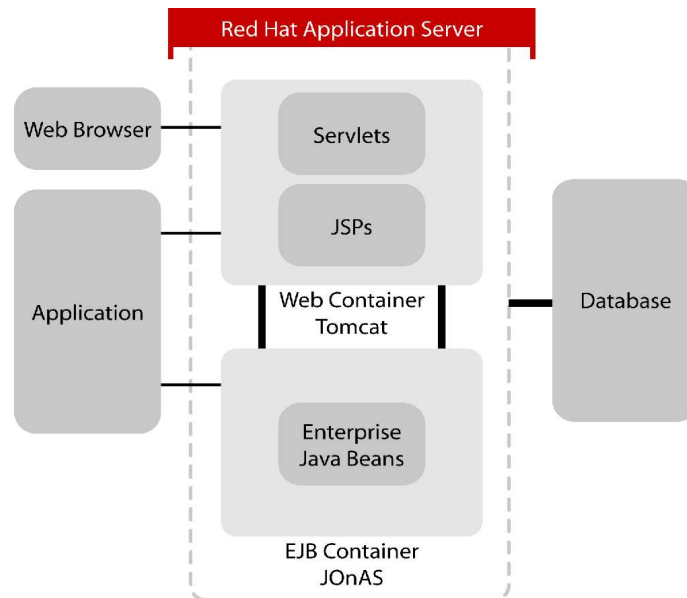


Figure 15: Application Server

Red Hat Developer Suite

Red Hat Developer Suite is a fully featured Integrated Development Environment (IDE) for application developers based on the open source Eclipse project. Eclipse provides an environment for developers to efficiently create a wide variety of applications in a rapidly growing set of languages including support for C/C++ and Java application development.

Eclipse supports a wide range of operating systems beyond Red Hat Enterprise Linux including Windows XP, Windows 2000, Windows 98, Windows ME, and Sun Solaris. This simplifies migration of developer skills and applications from these platforms to Red Hat Enterprise Linux.



Eclipse also supports a plug-in capability that allows language-specific and environment-specific extensions. The initial release of Red Hat Developer Suite includes plugins for C/C++, Java, RPM, and profiling; additional plugins will be provided as they become available. For detailed information on Eclipse refer to the project website at <http://www.eclipse.org/>.

Summary

Red Hat Enterprise Linux provides a high-quality operating system solution for the full range of commercial IT environments--from the desktop to the datacenter. It delivers excellent performance and has been enthusiastically adopted by the OEM and ISV community. It is available in certified hardware configurations with extensive commercial application support.

Enterprise Linux products are delivered with comprehensive Red Hat services, and Red Hat has partnered with ISVs and OEMs to deliver complete service solutions. For example, Oracle provides single-point-of-contact support for Red Hat Enterprise Linux. And Red Hat offers professional services to assist customers design, configure, deploy, and manage Red Hat Enterprise Linux solutions.

With Red Hat Enterprise Linux, Red Hat Desktop, and Red Hat layered products, together with certified applications and systems from Red Hat partners, customers benefit from capabilities, performance, applications, and services previously only available with proprietary solutions. When combined with the cost advantages offered by open source software, the results are compelling.

For more information or to purchase, visit www.redhat.com, or call 1-888-REDHAT1 (US and Canada) or +1-919-754-3700 (international).