

Security-enhanced Linux

Information assurance with HP Integrity Servers and Red Hat Enterprise Linux



Security-enhanced Linux	1
Executive summary.....	3
Section 1: Security in a networked world	4
The information assurance mission	4
Common Criteria for high-security Linux environments.....	5
Multi-level Security enhancements	5
The path to Multi-Level Security	5
Red Hat Enterprise Linux 4 and HP Intel processor-based systems	5
Past approaches to Multi-Level Security	6
HP standards-based platforms.....	6
Section 2: Red Hat Enterprise Linux.....	7
Committed to the most demanding security requirements.....	7
The path forward: Red Hat Enterprise Linux 5 and HP trusted platforms.....	8
Getting started.....	9
Section 3: HP Integrity servers.....	9
Better return on investment for all ranges of security needs.....	9
Flexible capacity.....	10
Performance	11
HP Super-Scalable Processor Chipset sx2000 (Linux support in late 2006)	11
Virtualization.....	12
HP Global Workload Manager (gWLM).....	13
HP Partition Manager.....	13
Hard partitions for HP Integrity servers	13
HP Integrity Virtual Machines for HP Integrity servers.....	14
HP Pay Per Use	14
High availability	14
Enterprise Reliability, Availability and Serviceability (RAS)	14
HP Serviceguard for Linux	14
Mission-critical services for Linux	15
Simplified management.....	16
Enterprise management with HP OpenView.....	16
HP Systems Insight Manager for Linux	17

HP NetTop	17
Section 4: Investment protection	18
Commitment	18
Momentum	18
Itanium Solutions Alliance	18
HP and Linux – commitment	19
HP and Red Hat – creating information assurance value	20
Security in a networked world and HP information security	20
For more information	21

Executive summary

As part of its information assurance mission, the National Security Agency has long been involved with the computer security research community in investigating a wide range of computer security topics, including operating system security. The results of several previous research projects in this area have been incorporated in a security-enhanced Linux operating system. This version of Linux has a strong, flexible, mandatory access control architecture incorporated into the major subsystems of the kernel. Linux was chosen as the operating system platform for this work because its growing success and open development environment provided an opportunity to demonstrate that this functionality can be successful in a mainstream operating system and, at the same time, contribute to the security of a widely-used system.

HP has engaged for years in cooperative partner work and platform testing to maintain and continually update its compliance with the internationally recognized Common Criteria evaluation scheme for security across Linux platforms. HP has officially obtained Common Criteria certification for a Controlled Access Protector Protocol (CAPP) at an Evaluation Assurance Level of 3+ (EAL 3+). HP is now teaming with Red Hat and Intel to create an even more secure Linux environment for government and commercial IT deployments that require trusted operating system features that minimize the risk of damage from attacks on their networks.

The release of Red Hat Enterprise Linux 4 in February 2005 marked the first time that Mandatory Access Control was delivered as a core feature of a commercially-available, mainstream Open Source operating system. This was provided through the inclusion and full support of Security-Enhanced Linux, initially started as a research project by the NSA to add Mandatory Access Control to Linux. Security-Enhanced Linux, or SELinux as it is also known, is developed within the Open Source community and has been incorporated into the upstream 2.6 Linux kernel. It is important to note that, just as with the advent of Mandatory Access Control and Type Enforcement in Red Hat Enterprise Linux 4, the basis for the new Multi-Level Security (MLS) support will be in Red Hat Enterprise Linux 5; as before, there will be no separate "Red Hat Enterprise Trusted Linux." Red Hat Enterprise Linux 5, scheduled for release in late 2006, will be sponsored by HP and is "in evaluation" for EAL 4+ with Labeled Security Protection Profile (LSPP), Role-Based Access Control Protection Profile (RBACPP) and CAPP. Common Criteria EAL 4/LSPP replaces the older TCSEC Orange Book B1 standard for Multi-Level Security operating systems. The addition of MLS capabilities in conjunction with the Common Criteria evaluation enables the combination of Red Hat Enterprise Linux 5 and HP Integrity and HP ProLiant systems to be the foundation for a solution that represents the highest level of security to date for a Linux platform, formerly the province of a select few trusted platforms.

Global events have made it increasingly critical that government defense and intelligence agencies are equipped with the tools and technologies needed to anticipate, prepare for and correctly respond to threats or attacks, whether man-made or natural. The flexible capacity of the HP Integrity servers is architected to address these critical needs. Government departments also need to improve their service quality and make it easy for citizens and businesses to interact with them. They need to reduce the cost of delivery for existing government services and minimize the cost impact of new services. Providing an infrastructure that is highly available to satisfy these requirements is integral to HP Integrity systems value. Lastly, the availability of powerful, integrated management tools for HP Integrity servers simplifies management of IT operations so that governments can focus on their mission. With many years of enterprise computing experience and joint R&D investments with Red Hat, HP delivers enterprise-level mission-critical solutions that governments can rely on to solve complex security problems. The combined balance of stability and innovation enables HP and Red Hat to offer an exceptional, complete experience at an affordable price. Together, HP and Red Hat design, develop and deliver the Linux infrastructure technology, along with the breadth of services to implement it.

Section 1: Security in a networked world

The information assurance mission

Without a doubt, we live in a network-centric world. New information technologies arrive at lightning speed, allowing us to share information across town, across the country, or around the world faster than ever before. NSA's Information Assurance Directorate (IAD) is dedicated to providing information assurance solutions that will keep information systems safe from harm. National security depends on it.

IAD's mission involves detecting, reporting and responding to cyber threats; making encryption codes to securely pass information between systems; and embedding IA measures directly into the emerging Global Information Grid. It includes building secure audio and video communications equipment, making tamper-protected products and providing trusted microelectronics solutions. It entails testing the security of customers' systems, providing OPSEC (Operational Security) assistance and evaluating commercial software and hardware against nationally-set standards. To better meet national IA needs, IAD works with government, industry and academia.

To provide system security, end systems must be able to enforce the separation of information based on confidentiality and integrity requirements. Operating system security mechanisms are the foundation for ensuring such separation. Unfortunately, existing mainstream operating systems lack the critical security features required for enforcing separation: mandatory access control. As a consequence, application security mechanisms are vulnerable to tampering and bypass, and malicious or flawed applications can easily cause failures in system security.

The results of several previous research projects in this area have been incorporated in a security-enhanced Linux operating system. This version of Linux has a strong, flexible mandatory access control architecture incorporated into the major subsystems of the kernel. The operating system provides a mechanism to enforce the separation of information based on confidentiality and integrity requirements. This allows threats of tampering and bypassing of application security mechanisms to be addressed and enables the confinement of damage that can be caused by malicious or flawed applications.

Linux was chosen as the operating system platform for this work because its growing success and open development environment provided an opportunity to demonstrate that this functionality can be successful in a mainstream operating system and, at the same time, contribute to the security of a widely-used system. Additionally, the integration of these security research results into Linux may encourage additional operating system security research that may lead to further improvements in system security.

Providing a complete security solution for Linux goes well beyond the scope of this work, and security-enhanced Linux is not an attempt to correct any flaws that may currently exist in Linux. Instead, it is simply an example of how Mandatory Access Controls that can confine the actions of any process, including a super-user process, can be added to Linux. The focus of this work has not been on system assurance or other security features, such as security auditing, although these elements are also important for a secure system.

The security mechanisms implemented in the system provide flexible support for a wide range of security policies and make it possible to configure the system to meet a variety of security requirements. The release includes a general-purpose security policy configuration designed to meet a number of security objectives as an example of how this may be done. The flexibility of the system allows the policy to be modified and extended to customize the security policy as required for any installation.

While there is still much work needed to develop a complete security solution, security-enhanced Linux presents a good starting point to bring valuable security features to Linux and a foundation to build upon this work within the Linux community.

Common Criteria for high-security Linux environments

HP has engaged for years in cooperative partner work and platform testing to maintain and continually update its compliance with the internationally recognized Common Criteria standard for high security across Linux platforms. This is an international effort where over 20 countries came together to reach agreement on security levels. HP has officially obtained Common Criteria certification for a Controlled Access Protector Protocol (CAPP) at an Evaluation Assurance Level of 3+ (EAL3+).

Multi-level Security enhancements

HP is now teaming with Red Hat and Intel to create an even more secure Linux environment for government and commercial IT deployments that require trusted operating system features that minimize the risk of damage from network attacks. Several Red Hat Enterprise Linux versions have achieved Common Criteria evaluation on HP systems and workstations. Currently, Red Hat Enterprise Linux 4 is the latest version to be evaluated for EAL3+ with the Controlled Access Protection Profile (CAPP) on HP Integrity and HP ProLiant systems. In progress is Labeled Security Protection Profile (LSPP) and Role Based Access Control (RBAC) certification at EAL4+, resulting in an Open Source Common Criteria-certified, multi-level security platform using a standard commercial Linux offering. This level of security functionality is targeted for integration into Red Hat Enterprise Linux from HP and Red Hat running on industry-standard HP Integrity and ProLiant servers.

This raises the bar for highly secure, trusted Linux customer deployments while also facilitating ISV and developer integration efforts, since unlike alternative trusted operating systems, a specialized operating system and proprietary hardware are not required.

The path to Multi-Level Security

Red Hat Enterprise Linux 4 and HP Intel processor-based systems


Organizations requiring Multi-Level Security have faced many challenges. Implementing Multi-Level Security has traditionally been an expensive process, made more difficult by the fact that traditional Multi-Level Security and Mandatory Access Control capabilities have not been available previously in a mainstream operating system. As most of the trusted operating systems that provided these capabilities have failed and exited the market, a new built-in solution is required.

A multi-user system within the military and intelligence community may have users with different levels of clearance, each trying to access information with different levels of classification. The operating system used in these environments plays a critical role in enforcement of privileges. It must allow users to access the data for which they have been cleared while ensuring they do not access information they are not authorized to view. Because of the nature of sensitive information, and the potentially grave consequences that could be expected should that information be leaked, the operating system must enforce these restrictions irrespective of the actions of programs, users or administrators. A Multi-Level Security, or MLS, operating system is one that permits the appropriate information flow and enforces these restrictions.

The classic model for Multi-Level Security is Bell-LaPadula (BLP). The fundamental precepts of this model can be summarized in the phrase “no read up, no write down.” A user on the system must not be allowed to “read up” — that is, read information that is of a higher level than they have been cleared for e.g., a user with Confidential clearance must not read information classified as Secret or Top Secret. In addition, users must not be able to read information in compartments for which they have not been granted access. These restrictions also extend to programs being run by that user on the system. In addition, “no write down” means users and programs on the system must not be able to write information classified at one level down to a lower level.


Mandatory Access Control architecture

Flexibility within the mainstream 2.6 Linux Kernel allows for multiple policies



Discretionary access control

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



Mandatory access control

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

Multi-Level Security/Bell-LaPadula is one implementation of Mandatory Access Control (MAC). On a Mandatory Access Control operating system, a security policy on the system is the final arbiter for all access-control decisions. A flawed, compromised or misconfigured program must not be able to work around the security policy, nor may the accidental or intentional actions of a user or even a system administrator. This is to be contrasted with a Discretionary Access Control (DAC) model, where users and programs have discretion over their objects on the system, such as their files. On a MAC system the system security policy, not the user, would make the final determination as to who and what can read, write and execute certain files.

Past approaches to Multi-Level Security

Traditionally, Multi-Level Security through the Bell-LaPadula model has been implemented in what are commonly known as trusted operated systems, so called because the operating system is trusted to enforce the strictures of information flow amongst multiple users and multiple security levels on the system. Examples of trusted operating systems include Trusted Tru64, Trusted HP-UX, Trusted AIX and Trusted Solaris.

The commercial availability of these trusted operating systems has been less than successful; this has been true for the manufacturers of these products and for those who have attempted to implement them. The one major exception to this rule has been the comparative success of Trusted Solaris within the military and intelligence community. There are several reasons for their lack of adoption; key among them is the fact that these trusted operating systems have been separate from their mainstream commercial counterparts. As such they have lagged behind their cousins in core operating system features and capabilities. Moreover, the ecosystem of certified hardware and applications is smaller, as is the pool of available expertise to properly implement and administer these solutions. All of these reasons, in addition to the initial cost of the software itself, have made these systems extremely expensive to deploy and maintain.

HP standards-based platforms

HP Integrity and HP ProLiant servers offer the flexibility, performance and value that are ideal building blocks for today's highly competitive government environment demands. Flexibility begins with IT solutions that incorporate the skills, resources and hardware necessary to support heterogeneous

computing environments. HP Integrity and ProLiant servers run the industry's leading operating systems, enabling rapid, seamless deployment of new solutions and applications.

In addition, HP Integrity servers achieve strong performance across multiple workloads and operating environments by leveraging the capabilities of the Intel Itanium 2 microprocessor. This powerful processor is integrated with the HP Scalable Processor Chipset zx1 in entry-level servers and the HP Super-Scalable Processor Chipset sx2000 (Linux support available in late 2006) in mid-range and high-end servers—significantly increasing memory and I/O subsystem scalability.

Section 2: Red Hat Enterprise Linux

Committed to the most demanding security requirements

The release of Red Hat Enterprise Linux 4 in February 2005 marked the first time that Mandatory Access Control was delivered as a core feature of a commercially-available, mainstream operating system. This was provided through the inclusion and full support of Security-Enhanced Linux, initially started as a research project by the NSA to add Mandatory Access Control to Linux. Security-Enhanced Linux, or SELinux as it is also known, is developed within the Open Source community and has been incorporated into the upstream 2.6 Linux kernel.

SELinux in Red Hat Enterprise Linux 4 provides Type Enforcement (TE), another approach to Mandatory Access Control besides the aforementioned Multi-Level Security/Bell-LaPadula model. Whereas the Bell-LaPadula model of Multi-Level security is focused on data confidentiality at the sake of integrity, Type Enforcement does not force users to make that compromise. SELinux with Type Enforcement is a more flexible approach for users building systems that provide high levels of both data confidentiality and integrity. Type Enforcement adopts the principle of least privilege, where an application is given only enough permission to function as intended, but no more. Under a Type Enforcement model, applications run in separate areas, known as domains, and as such are isolated from one another and from the underlying operating system. These domains are defined by the security policy on the system. A flaw or misconfiguration in an application protected by Type Enforcement is isolated within that application's domain.

Recognizing that the advantages of Mandatory Access Control have a wider application than the traditional customer base of the trusted operating systems, and desiring to avoid the mistakes of the past, Red Hat intentionally made Security-Enhanced Linux a core feature of Red Hat Enterprise Linux 4. There is no separate Red Hat Enterprise Trusted Linux product nor is the purchase of "Trusted Extensions" or a similar add-on product required in order to implement and use SELinux. Customers have access to a comprehensive ecosystem of certified HP hardware and third-party applications, as well as the latest advantages in core operating system functionality. These are the benefits of deploying a rapidly-growing, mainstream operating system.

By default, Red Hat Enterprise Linux has SELinux enabled and implemented under a targeted policy. This policy is designed to draw a balance between usability on the one hand and increased security on the other. The targeted policy confines several network-listening daemons—services such as Web servers, DNS servers and mail servers that are constantly under attack—that generally do not impact an end user's ability to run other applications and otherwise use the system.

Because SELinux and Type Enforcement offer a flexible security model, allowing one to tailor the policy to the actual functional needs and security requirements of users on the system, it allows for the implementation of other policies is allowed. On the one hand, users may choose to disable SELinux completely and run in a traditional Discretionary Access Control model; on the other hand, they have the ability to use an optional strict policy, a superset of the targeted policy that confines many more applications and places additional limits on what end users can do on the system.


The path forward: Red Hat Enterprise Linux 5 and HP trusted platforms

While Type Enforcement, owing to its flexibility, is arguably better suited for most uses, there remain several critical instances where traditional Multi-Level Security is still required. These are typically scenarios involving a large number of classification levels where there are multiple permutations of several different compartments and sensitivity levels and multiple users with varying levels of clearance.

In order to enable these environments, Red Hat Enterprise Linux 5 plans to add support for Multi-Level Security using the Bell-LaPadula model in conjunction with Type Enforcement. This MLS policy will be one of the available policy options, in addition to the default targeted policy and optional strict policy currently available for Red Hat Enterprise Linux 4. Security-Enhanced Linux will serve as the foundation for the MLS policy. As with the original work that was done to implement SELinux in Red Hat Enterprise Linux 4, the initial developments are being done within the Fedora Project. Fedora Core 5, released in March 2006, will have much of the core infrastructure necessary to support the MLS policy in place. This will enable users and developers to begin testing and exercising the code paths used by MLS. The approach is to refine these capabilities so that they are production-ready for Red Hat Enterprise Linux 5.

The highest-level security capability for a Linux platform

- **HP Integrity and ProLiant systems with Red Hat Enterprise Linux 5 at EAL4+ (planned for late 2006)**
 - Labeled Security Protection Profile (LSPP)
 - Role Based Access Control Protection Profile (RBAC)
- **Full support within mainstream Red Hat OS**
 - Not a separate trusted product or trusted extensions
 - No dead-end proprietary costs
- **Development environment available now**
 - RHEL4 and Fedora Core 5 w/ SELinux core feature
 - Training to gain hands-on experience
 - Porting from proprietary trusted environments
 - HP configuration guides



It is important to note that, just as with the advent of Mandatory Access Control and Type Enforcement in Red Hat Enterprise Linux 4, the basis for the new MLS support will be in Red Hat Enterprise Linux 5; as before, there will be no separate "Red Hat Enterprise Trusted Linux". Red Hat Enterprise Linux 5, scheduled for release in late 2006, will be sponsored by HP and soon be "in evaluation" for EAL 4+ with CAPP, Labeled Security Protection Profile (LSPP) and Role-Based Access Control Protection Profile (RBACPP). Common Criteria EAL 4/LSPP replaces the older TCSEC Orange Book B1 standard for Multi-Level Security operating systems. The addition of MLS capabilities in conjunction with the Common Criteria evaluation enables the combination of Red Hat Enterprise Linux 5 and HP Integrity and HP ProLiant systems to be the foundation for a solution that meets the requirements of Director of Central Intelligence Directive (DCID) 6/3 Protection Level 4 (PL4) for Top Secret and Below Interoperability (TSABI). This represents the highest level of security capability to date for a Linux platform, formerly the province of a select few of the trusted platforms.

Getting started

System integrators and developers need not wait until the commercial availability of Red Hat Enterprise Linux 5, scheduled for release in late 2006, to begin the migration of their Multi-Level Security applications. Since Security-Enhanced Linux will be used as the foundation for the MLS capabilities, an understanding of how SELinux works and how to write SELinux policy is a necessary first step. Red Hat Enterprise Linux 4 and Fedora Core 3 and above include SELinux as a core feature of the operating system and can be used to gain hands-on experience with the technology. Formal training is also available from Red Hat, including an "Introduction to SELinux and Red Hat Targeted Policy" and a more advanced class on "Red Hat Enterprise SELinux Policy Administration."

In addition to the particulars of SELinux policy writing, Red Hat offers training and support for application developers currently working in a proprietary UNIX environment desiring assistance with porting and developing applications for Red Hat Enterprise Linux. This includes a class on "Red Hat Linux Application Development and Porting" and Premium Developer Support.

Professional Services to assist with the migration and porting effort are also available from HP. For information on a single point of contact for Red Hat Enterprise Linux products running on HP systems, see <http://www.hp.com/go/linuxservices>

Section 3: HP Integrity servers

Better return on investment for all ranges of security needs

While Multi-Level Security is focused on integrity of the data, users often require additional capabilities to allow them to build systems that provide high levels of both data confidentiality and integrity. Within these requirements, there are some common themes that federal and civilian agencies are faced with. First, they are on a drive to consolidate and simplify their IT environment, free up staff and achieve government budget restrictions. Too often, 75% of their budget and resources are tied up in maintenance and in maintaining applications or servers in their data centers. Second, they are striving to bring their high-value, mission-critical projects online and respond faster to changes in government requirements. Two out of three IT projects never actually reach completion, and much of the reason for that is that mission requirements change before the project can be completed rendering the result "meaningless." And thirdly, every agency must continuously meet the high service-level expectations of their users and their constituents. And it's not just response time or high availability. Increasingly, it involves security and government compliance.

To meet these needs, HP focuses on delivering three key capabilities to help agencies achieve a better return on IT on Linux with HP Integrity solutions. Flexible capacity – not just fast, individual boxes – provides the capacity to meet the workloads that are at hand, whether they're a batch or a query or transaction, and the ability to adapt quickly to change. And virtualization, the utilization of resources and in-box upgrades – all apply to making the capacity of HP systems flexible, so that the capacity of their data centers meets their mission needs. The second capability is high availability. Availability must be built-in from the ground up, from systems, to operating systems, to integration with applications. HP is the leader here, and as agencies deploy HP servers and storage, the availability, stability and security of their environment improves as a result.

Finally, there is the need for simplified management. It should be easier to maintain the systems for mission-critical deployments. It should take fewer people to maintain those stable systems and it should be easier to set up and operate a virtualized environment. So the investments that HP is making in HP Systems Insight Manager, HP OpenView and Virtual Server Environment are all contributing to that better return on investment. In addition, the supportability of the systems based on services for a longer deployment lifecycle ensure mission continuity.

One of the major thrusts of the HP investment is towards delivering a ROI that fits customers' budget constraints. The Integrity platform from HP helps accomplish this through the longevity of the architecture that uses standards-based technology; industry-leading multi-OS capability to easily repurpose hardware and virtualization to deliver increased utilization; leadership in high availability and Linux support; an upgrade path that does not require a forklift every 30 months; and automated, intelligent management. These capabilities work together in HP Integrity server solutions to drive down the cost of operations.

HP Integrity servers: the best platform for your government challenges



Flexible capacity

Global events have made it increasingly critical that government defense and intelligence agencies are equipped with the tools and technologies needed to anticipate, prepare for and correctly respond to threats or attacks, whether man-made or natural. The flexible capacity of the HP Integrity servers is architected to address these critical needs. HP Integrity servers are designed to deliver real-world performance and flexibility for today's diverse workloads, making mission critical projects more effective in meeting government needs.

Adjust to mission changes in real time:

- Move resources to where they are needed, easily and quickly
- Handle workloads at peak times
- Leverage spare CPU cycles

Have the capacity for all workloads:

- Handle transactions, batch and queries
- Enough resources for ad-hoc workloads
- Immediate response for real-time demands

Scale systems as mission needs change:

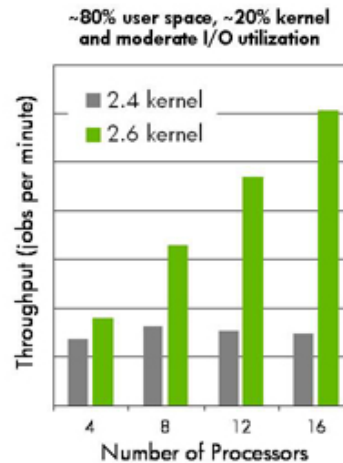
- Leverage modular building blocks
- Repurpose systems as application needs evolve
- Pay for contingency capacity if and when needed

Performance

HP continues to produce world-record-breaking results on Intel Itanium 2 processor-based HP Integrity servers running Linux, and the benchmarks with Red Hat Enterprise Linux and industry leaders such as Oracle are no different. Scale-up support on HP Integrity servers for Red Hat Enterprise Linux 5 shows that Linux 2.6 kernels perform better with near-linear scaling to ~16 processors for many workloads. Compute-intensive workloads, such as floating-point intensive encryption/decryption typically scale well beyond original estimates.

Scale up support on HP Integrity for Red Hat Enterprise Linux 4

- 2.6 kernels perform better with near linear scaling to ~16 processors for many workloads
 - Compute intensive workloads typically scale well beyond
- Characterization of commercial database workloads is planned



Oracle on HP Integrity servers currently hold the #1 Linux TPC-C benchmark, held for 25 months and counting, the #1 Oracle on Linux 4-processor TPC-C, held for 14 months and counting, and the #1 Oracle on Linux TPC-C price/performance, held for 16 months and counting. Gartner ranks the HP Integrity as the #1 Linux servers for Oracle DBMS in the December 2005 Server Scorecard Evaluation Model, Version 1.0.

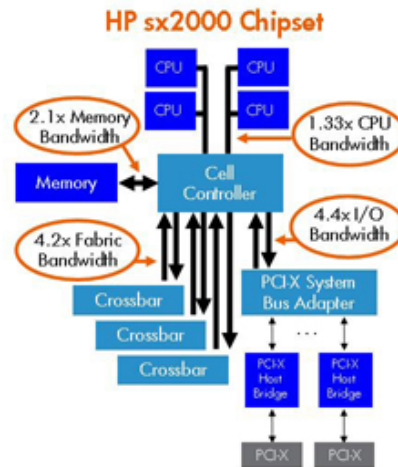
HP Super-Scalable Processor Chipset sx2000 (Linux support in late 2006)

In addition, HP Integrity servers achieve strong performance across multiple workloads and operating environments by leveraging the capabilities of the Intel Itanium 2 microprocessor. This powerful processor is integrated with the HP Scalable Processor Chipset zx1 in entry-level servers and the HP Super-Scalable Processor Chipset sx2000 in mid-range and high-end servers—significantly increasing memory and I/O subsystem scalability. The HP Super-Scalable Processor Chipset sx2000 is the latest example of HP innovation designed to deliver more simplicity, agility and value to midrange and high-end HP Integrity server customers. The HP sx2000 chipset adds greater performance, scalability, reliability and manageability to these Itanium 2-based systems—and support for future processors—with larger bandwidths and lower latencies. The HP sx2000 chipset makes the Intel Itanium 2 9M processor more powerful. It provides interconnectivity between processors, memory, and I/O cards, turning this group of components into a high-performance computer system. The HP sx2000 chipset is a set of five very large-scale integration (VLSI) components that enhance the performance, availability and manageability of the Intel Itanium 2 processors that power the HP Integrity rx7640 Server, the HP Integrity rx8640 Server and the HP Integrity Superdome. These components consist of a cell controller, a memory buffer, a crossbar switch, a PCI-X system bus adapter and a PCI-X host bridge.

The HP sx2000 chipset's well-balanced architecture is designed to achieve significantly greater performance from existing Intel Itanium 2 9M processors. It provides more than twice the memory bandwidth and four times the I/O and fabric bandwidths of the HP sx1000 chipset—along with as much as 30% reduction in cache-to-cache, memory and I/O latencies. In addition to providing greater performance and availability, the HP sx2000 chipset offers memory enhancements based on DDR2 memory technology, increased memory address space and planned support for larger DIMM sizes in the future as well as double chip-sparing error-correcting code (ECC) to increase availability and reliability.

Larger bandwidths and reduced latencies create greater balanced performance

- HP sx2000 chipset - perfect complement to Intel® Itanium® 2 processors
 - Bandwidth – up to 4.4x
 - System latencies – up to 45% reduction
 - Memory capacity – up to 4x (DDR2 SDRAMs with support for 2Gb when available)
 - I/O cache size – 50% increase
- Perfect for any customer workloads
 - Transactional processing
 - Analytical processing
 - Technical workloads
- Upgrade to Intel's next-generation dual-core Itanium® 2 processor
 - Double the compute density with twice as many cores per box



The HP sx2000 chipset delivers I/O advancements and support for a wide variety of I/O cards that enhance system performance and availability. The cache size of the PCI-X System Bus Adapter has been increased by 50% so that the increased bandwidth requirements of the I/O cards can be satisfied. Plus, the cache hides the non-uniform memory access time for different lines that the I/O card has requested from memory for DMA reads or writes. This allows data to be streamed to and from the I/O card without added delays. High-speed links have been re-designed with self-healing qualities inside the HP sx2000 chipset. These new link technologies provide greater reliability by tolerating transient errors and greater availability by tolerating a hard error in one of the channels that make up a link. If the chip finds a problem in one of the channels of the link, it will swap in the spare channel and continue running at full bandwidth. If a link or crossbar chip fails completely, the HP sx2000 chipset minimizes downtime by allowing the partitions affected by the failure to be immediately rebooted using the unaffected fabrics.

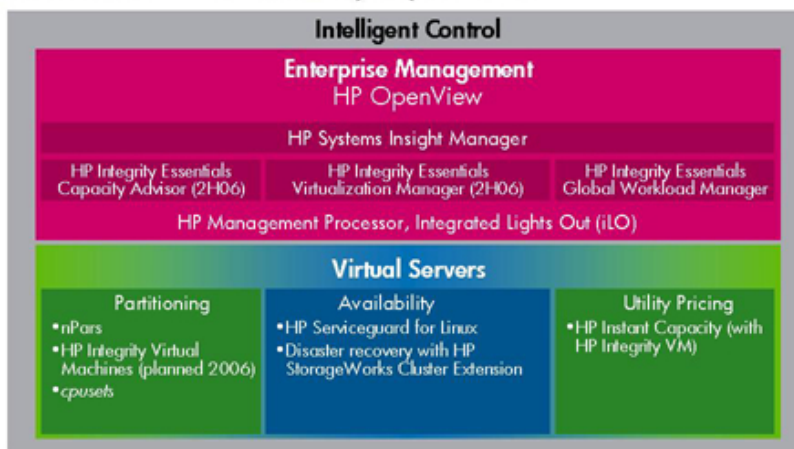
Virtualization

By reducing complexity, increasing resource utilization and lowering costs, government IT managers can acquire the flexibility to devote more of their attention to new mission opportunities and less on maintenance and management. Virtualization is an approach to IT that pools and shares resources (servers, storage, networking, applications and clients) so that supply automatically meets demand. Virtualization delivers a unique opportunity for IT to simplify the IT infrastructure while “doing more with less,” i.e., meeting budget constraints while increasing utilization of system resources and decreasing cost at the same time. The Virtual Server Environment for Linux is an integrated solution for HP Integrity servers that controls and virtualizes server resources for the optimum utilization that

mission priorities demand in real time, allowing a greater return on budget investments. Within this environment, virtual servers automatically grow and shrink based on the service level objectives set for each application they host. Through tight integration with high availability and partitioning, service levels can be maintained in the event of unexpected downtime. The low-latency massive bandwidth architecture of HP Integrity servers is designed to reduce latency and delivers the response times to meet service-level requirement for latency-sensitive applications like business intelligence and quick transactions. Unlike IBM, which has off-chip 3rd level cache and requires more CPUs and performs poorly for business intelligence, or Sun, which only recently added on-chip cache in UltraSparc IV, HP Integrity servers have massive I/O bandwidth that can be used in parallel to give excellent database performance, and Itanium 2 cache, currently at 9MB and expanding on Montecito to 'share' 24 MB.

Simplified management of Linux on HP Integrity

Virtual Server Environment (VSE) for Linux



HP Global Workload Manager (gWLM)

HP Global Workload Manager (gWLM) is the policy engine that dynamically allocates virtual server resources to specific applications. It integrates with resource management groups and virtual partitioning techniques to improve server utilization while maintaining service levels across many servers. If any resources don't meet a service-level agreement at any time, gWLM will immediately re-allocate the appropriate resource.

HP Partition Manager

Partition Manager for Linux on Integrity Servers provides system administrators with a convenient Web-based interface to manage nPars (hard partitions) on HP servers. Administrators can create, modify and delete nPars and perform high-availability checks to ensure new configurations are highly available on HP Integrity servers. Partition Manager allows them to visualize and configure "big picture" views of all hardware resources associated with the nPars.

Hard partitions for HP Integrity servers

HP nPartition (nPars) system capabilities enable the configuration of a single-server complex as one large system or as multiple smaller systems. These Linux hard partitions provide complete electrical and software isolation so that any fault within one partition cannot impact any other partition—making it ideal for organizations in which high availability is critical.

HP Integrity Virtual Machines for HP Integrity servers

Planned availability in 2006 for Linux, HP Integrity Virtual Machines will increase resource utilization by offering sub-CPU virtual machines with shared I/O across Linux, HP-UX 11i and Windows environments, with dynamic resource allocation.

HP Pay Per Use

The HP Linux Pay Per Use (PPU) program is designed for organizations with widely varying or unpredictable demand for computing resources. With PPU, customers have access to reserve capacity, but don't pay for that capacity when it's not in use. Their bill is comprised of a fixed base amount plus a variable amount that is linked to their usage. The key to the HP PPU offering is its metering technology. All of the HP PPU server solutions employ the same metering appliance (Utility Meter) to communicate usage data to HP. The Utility Meter is a separate piece of hardware that resides inside the data center and on the network.

High availability

Government departments need to improve their service quality to enable inter-agency and global collaboration as well as make it easy for citizens and businesses to interact with them. They need to reduce the cost of delivery for existing government services and minimize the cost impact of new services. Providing an infrastructure that is highly available to satisfy these requirements is integral to HP Integrity systems. Integrated availability solutions are built-in for HP Integrity servers to keep applications up and running and preventing agency operations from being affected by IT outages... whatever the cause.

Support continuous mission critical operations:

- Maximize application availability to the users
- Minimize planned and unplanned downtime
- Recover quickly from a disaster
- Preserve data integrity
- Optimize user response time

Enterprise Reliability, Availability and Serviceability (RAS)

HP Integrity servers running Linux offer agencies an extremely reliable, industry-standard alternative to proprietary RISC systems. Reliability is built into every server with features such as enhanced Machine Check Architecture (MCA), double-chip spare memory for error checking and correction (ECC) and Dynamic Resiliency Memory (DRM) providing such features as dynamic page de-allocation.

Governments' IT investments in HP Integrity servers are protected by design. To handle increased user demand as their mission grows, Integrity servers provide them the ability to perform fast, in-box technology upgrades for greater computing capacity. Because HP Integrity servers are standards-based, they can be quickly and seamlessly integrated into heterogeneous, legacy infrastructures—and they can run multiple operating systems on the same platform.

HP Serviceguard for Linux

HP Serviceguard for Linux is an industry-leading high availability clustering solution that protects mission-critical applications from a wide variety of hardware and software failures. HP Serviceguard for Linux brings these robust technologies to the Linux operating environment, providing efficient, continuous access to mission-critical applications, information and services, as well as backup and data recovery.

HP Serviceguard for Linux groups Linux processes into highly-available packages. When a server, application, service, network or other resource fails, HP Serviceguard automatically transfers control of that node's package to another node within the cluster—which maintains service availability with minimal interruption. Through an end-to-end system of alerts, error detection and dynamic resource

allocation, HP Serviceguard for Linux safeguards cluster operations to provide a high degree of availability.

HP offers extended Linux clusters with HP StorageWorks Cluster Extension EVA (CLX) for inter-site replication, management of site failover and failback, and finally, geographical dispersion with automatic disaster tolerance. CLX is middleware between HP Continuous Access EVA and Serviceguard that extends otherwise local cluster solutions over metropolitan distances. CLX EVA automates server cluster failover operations by synchronizing failover of EVA storage with replication management for a disaster tolerant solution for HP Serviceguard for Linux. CLX EVA provides automatic failover capabilities by making the decision on whether it is appropriate for the cluster software to failover the application from one cluster node to another.

In addition, leveraging the robust remote mirroring capability of HP Continuous Access XP, Cluster Extension XP will ensure business continuity by confirming that data will be available at a remote location up to metropolitan-wide distances, all without downtime or performance impact. Without this confirmation, cluster failover will occur independently from HP Continuous Access XP remote mirroring capability, thus requiring a labor-intensive, manual process to monitor both the cluster failover and as well as the remote mirroring process.

Automatic failover/failback reduces the complexity involved in a disaster scenario. CLX seamlessly integrates into MC/SG for Linux and utilizes the fast failback functionality of HP Continuous Access XP to provide automatic, fast and efficient failover and reliable recovery.

HP Continuous Access XP also can be deployed over a wide range of network connectivity, ranging from direct ESCON or fiber channel connects, or extended over wider distances through Dense Wave Division Multiplexing (DWDM) or converters.

Mission-critical services for Linux

HP and Red Hat take a cooperative approach by building quality into the process. Rigorous joint-testing ensures customers can trust that their Open Source environment has been integrated and fully tested. HP Services professionals take a collaborative approach to help organizations reduce IT complexity. They work with customers to design, deploy, integrate and manage an agile IT infrastructure that responds to change and more closely aligns IT with mission goals. HP offers a full range of service and support capabilities that extend throughout the lifecycle of the HP Integrity server and Linux environment. HP has recently introduced the Mission Critical Workshop and Mission Critical Design Service for organizations that wish to engage with HP at the beginning of the solution lifecycle. These enterprise-level approaches to needs assessment and design leverage the breadth and depth of HP's systems expertise, and both are offered across a number of operating system platforms in addition to Linux.

HP has over 6,500 trained Linux professionals with software and tools, ranging from root-cause analysis to system hardening, as well as online productivity and serviceability tools.

Many HP customers have also turned to HP Solution Centers for assistance in planning and deploying their solutions. With more than 80 locations worldwide, HP Solution Centers offer a risk-free environment—as well as tools and expertise—for developing, testing and validating a variety of technologies. HP Solution Centers help customers make informed, cost-effective decisions and prepare them to manage and support their solutions effectively—before making an investment.

HP also offers nine Linux courses, including two courses for implementing various security functions with Linux. Information is available at <http://www.hp.com/education/courses/u8630s.html>

Professional Services to assist with the migration and porting effort are also available from HP. For information on a single point of contact for Red Hat Enterprise Linux products running on HP systems, see <http://www.hp.com/go/linuxservices>

HP Consulting & Integration has worked extensively with government organizations in secure environments, such as the joint development for NSA SELinux-based NetTop™ security software. HP

C&I provides the SELinux Assessment Service and SELinux Security Policy Development Service to assist organizations in this area. Within HP Services, Linux is supported as Tier 1 platform for HP Integrity and ProLiant servers and includes architecture planning, consulting, migration and training. Managing facilities can be performed onsite, remote or outsourced with Worldwide Service support for SLAs, 24x365 or 9x5. High availability options are capable of meeting whatever the mission requirement dictates.

Simplified management

For government agencies, driving business process efficiency and predictable costs is as important as increasing the quality of security and ensuring compliance. In higher education establishments, high-performance grid computing, research and visualization are critical. In both cases, the ability to manage diverse computing environments is the key to optimize the use of their datacenter facilities. Powerful, integrated management tools for HP Integrity servers simplify management of IT operations and enable users to focus on their mission.

Improve productivity of system managers and DBAs:

- Make managing IT as consistent as possible
- Implement updates and patches faster and easier
- Automate to save time and eliminate human error

Optimize the use of datacenter facilities:

- Allow added capacity without adding floor space
- Keep power and cooling costs under control

Leverage industry best practices for efficient IT operations:

- Streamline IT service operations
- Make vs. buy decisions

Enterprise management with HP OpenView

HP OpenView and HP Systems Insight Manager lay a solid foundation upon which a broad IT service management solution can be developed. HP OpenView and HP Systems Insight Manager deliver a consolidated and centralized management platform for an entire enterprise, helping to reduce operating costs, increase availability and ensure efficient management of critical business resources.

HP OpenView Identity Management takes a unique approach to the enormous and vital task of identity management across—and beyond—the enterprise. Whereas most solutions focus on connecting users to resources, HP Identity Management focuses on the agency itself, structuring identity management to model the agency and the processes that support it. Resources are grouped into services, and users are then connected to the services they need to do their jobs. Provisioning and de-provisioning are handled through an intuitive, GUI interface. This service-oriented architecture greatly reduces complexity from day-to-day identity management tasks, allowing users more control and freeing IT for more strategic tasks. HP Identity Management functionality includes authorization and access control, authentication, Web single sign-on, federated identity as well as password management and synchronization.

With HP OpenView as the primary interface to manage the availability and performance of critical business services, integration with HP Systems Insight Manager enables administrators to correlate HP hardware status with the availability of business service-levels and obtain in-depth data for more accurate root cause analysis and faster problem resolution.

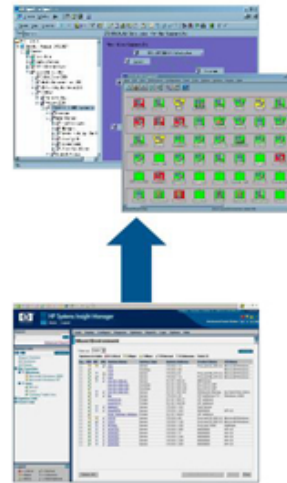
Simplified management stack from hardware to enterprise level

HP OpenView products

- Manage all levels of the IT service solution stack: hardware, networks and applications
- Maximize availability and performance of IT services
- Identity management
- Heterogeneous, vendor-independent management deployed at the enterprise level

HP Systems Insight Manager

- Optimize uptime and administration of platforms
- Reduce hardware lifecycle costs
- Best product for managing HP hardware
- Deployed at departmental and site level



HP delivers tight integration between Systems Insight Manager and OpenView using HP-developed connector components for OpenView Network Node Manager (NNM), OpenView Operations (OVO) and OpenView Service Desk (OVSD). This enables users of NNM and OVO and OVAD to view in-depth hardware data from their existing OpenView consoles.

HP Systems Insight Manager for Linux

HP Systems Insight Manager (HP SIM) is a robust platform lifecycle-management solution that improves efficiency and stability, increases staff utilization, enhances service levels and reduces downtime. As a key component of HP's industry-leading management portfolio, HP SIM is the foundation of HP's Linux management strategy. HP SIM is the central administration point to manage physical and virtualized environments, for all OS including Linux, and for virtually all HP hardware, including clients, storage, printers, network switches, racks and power products. It also leverages industry standards to enable management of non-HP platforms.

Manage the desktop with HP NetTop

The best defense against external threats is to never let them through the door. This principle, long applied in defense and intelligence communities, has become the standard in high-assurance computing environments as well, and is typically implemented through a physical "air gap" in which separate networks are physically disconnected from one another. Thus, users needing access to multiple security domains must employ a separate workstation for each domain.

HP NetTop is an information assurance solution architecture that transforms a single workstation into a high-assurance platform running multiple operating systems concurrently, with total domain isolation between each OS and its attached network. Operating systems such as Windows® 2000, Windows XP, Linux and others execute in their own isolated "Virtual Machine" (VM) vault on the HP NetTop system, and run the same office productivity and other software applications as standalone PCs. Users switch between VMs with a mouse click. HP NetTop provides a policy-driven barrier between VMs to prevent data leakage between security domains. To system administrators and their enterprise software systems, VMs are indistinguishable from stand-alone workstations on the network. Applications like Exchange, SMS and HP OpenView work transparently with HP NetTop VMs.

One HP NetTop workstation replaces multiple workstations in environments where data isolation cannot be compromised. In these environments, users typically access different security domains/networks from different workstations—leaving the IT organization to deal with the cost, clutter and complexity.

The HP NetTop solution, licensed from NSA, is based on NSA's NetTop research project and is supported by a suite of HP services that tailor HP NetTop to the needs of a unique enterprise. HP NetTop complies with NSTISSP No. 11, and is certified by the NSA Information Assurance Directorate to provide domain separation.

HP NetTop is backed by highly trained consultants and technical specialists with industry-recognized security certifications and a worldwide security solutions delivery and support team that provide in-depth expertise across today's security landscape, including technologies for firewalls, VPNs, authentication, identity management and access control. HP recognizes that unique security needs may require a custom solution, so the HP NetTop solution is offered as an integrated system of software and services to tailor HP NetTop to a specific environment. HP offers a menu of services to assess customer environments, tailor a security policy solution, pilot and roll out the solution and provide documentation, administrator training and continued software upgrades and support.

Section 4: Investment protection

Commitment

With the incorporation of Security-Enhanced Linux in Red Hat Enterprise Linux 4, HP and Red Hat took an unprecedented step by adding Mandatory Access Control into a growing, flexible security operating system. Its flexible Type Enforcement and targeted security policy capabilities extended the benefits of Mandatory Access Control to a wider class of users. This will enable users of the trusted operating systems to take advantage of the benefits that the HP trusted Linux platform has offered UNIX users in the past, namely an operating system that has UNIX-like capabilities but with the ability to run on lower-cost, higher-performance, standards-based hardware

HP and Red Hat—an industry-leading partnership: With many years of enterprise computing experience and joint R&D investments with Red Hat, HP delivers enterprise-level mission-critical solutions that governments can rely on to solve complex security problems. The combined balance of stability and innovation enables HP and Red Hat to offer an exceptional, complete experience at an affordable price. Together, HP and Red Hat design, develop and deliver the Linux infrastructure technology, along with the breadth of services to implement it.

Momentum

HP, Intel and Red Hat partnered together at FOSE 2006 to launch the Federal Open Source Alliance to deliver agile, economic, supported Open Source solutions on industry-standard platforms.

The HP and Intel alliance is strategic for both the companies and their customers, with four generations of Intel Itanium 2 processor-based products in definition and development and over 1000 Intel software engineers working on Itanium 2 processor-based tools, compilers and ecosystem. In conjunction with this, HP has committed to a \$1 billion investment per year to accelerate the adoption of HP Integrity Servers based on Intel Itanium 2 processors.

Itanium Solutions Alliance

HP and Intel, along with the top enterprise and technical computing providers Bull, Fujitsu, Fujitsu Siemens Computers, Hitachi, NEC, SGI and Unisys announced on September 26, 2005, the formation of the Itanium Solutions Alliance. Leading software vendors BEA, Microsoft, Novell, Oracle, Red Hat and SAS also joined this organization as charter members; and early members including

Hyperion, Informatica, MIT Systems, MSC Software, Sybase, Symantec, TIBCO and Trend Micro. This organization will boost availability of Itanium solutions through delivery of a suite of enabling programs targeted at enterprise and technical computing developers. With over 70,000 data center deployments of Itanium systems within the last four years, formation of the organization underscores the Alliance participants' collective long-term commitment to the expansion of Itanium solutions within enterprise and technical computing environments. The new Alliance delivers a suite of enabling programs designed to accelerate the availability of Itanium solutions, including Developer Days, the Itanium Solutions Center Network and the Itanium Solutions Catalog. All programs will provide software developers with resources to speed application optimization for Itanium solutions as well as ultimately providing end customers with a richer set of solution alternatives. For example, the Itanium Solutions Center Network is comprised of 20 global facilities hosted by Alliance member companies. The establishment of this program facilitates on-demand engineering support for tuning application environments. The centers will offer extensive support tools from founding sponsor companies as well as technical assistance from member company experts.

For customers, this will mean faster time-to-implementation of the applications most critical to their new business needs. Now customers can reap the benefits of industry-standard technology that provides lower cost, higher levels of system performance and application integration as well as a broader choice of hardware and software vendors committed to the Itanium architecture. IDC estimates that Itanium server shipments will grow by over 65 percent CAGR and the market share revenue will rise to 10 percent by 2009 to nearly \$6.6 billion.

HP and Linux – commitment

HP has invested as a supporter of the Linux environment by providing organizational leadership and sponsorships (Linux International, Open Source Dev Labs and others), a dedicated HP Linux and Open Source Lab, leading in the Eclipse Development Environment, and providing Intel Itanium (IA-64) kernel maintenance. HP was instrumental in Defining Carrier Grade Linux (telecom), driving the Open SSI Cluster Project and providing ongoing extensive support of the SAMBA, Apache and Debian Projects. The HP SELinux contributions include auditing, labeled printing and networking.

HP and Red Hat... creating information assurance value

CUSTOMERS		
Common Criteria Certified	Cost-effective Open Solutions	Mainstream OS with SELinux Core
	 <p>Security in a networked world teamed for the most demanding security requirements today</p>	
Investment	Commitment	Long-term Support
	 <p>Information security managing risks, enabling effectiveness</p>	

HP and Red Hat – creating information assurance value

Security in a networked world and HP information security

Creating real value that is based on an open operating environment and delivered from desktop to datacenter is what the combination of Red Hat and HP is all about. There is tight linkage between the HP information security initiative, known for managing risks and enabling effectiveness while bringing flexible capacity, high availability and simplified management to the synchronization of the mission and IT, and the Red Hat “Security in the networked world” initiative for making better decisions, faster by enabling the movement of information throughout secure environments. Together, HP and Red Hat are delivering mission-critical value that provides the information-assurance environment that governments demand.

For more information

Additional Resources

Dr. Rick Smith. "Introduction to Multilevel Security."
<http://www.cs.stthomas.edu/faculty/resmith/r/mls/index.html>

James Morris. "An Overview of Multilevel Security and LSPP under Linux."
http://www.livejournal.com/users/james_morris/5020.html

For the latest Fedora Core 5 schedule, refer to
<http://fedora.redhat.com/About/schedule/>

For the latest information on the HP and Red Hat Enterprise Linux and Common Criteria visit
http://niap.nist.gov/cc-scheme/in_evaluation.html#r
<http://www.hp.com/go/linuxsecurity>
<http://www.hp.com/go/linuxservices>
http://www.redhat.com/en_us/USA/home/solutions/government/commoncriteria/

For more information see
<https://www.redhat.com/training/security/courses/rhs427.html>

For more information see
<https://www.redhat.com/training/security/courses/rhs429.html>

For more information see
<https://www.redhat.com/training/developer/courses/rhd256.html>

For more information see
<https://www.redhat.com/support/offerings/premium.html>

© 2006 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

Red Hat is a trademark of Red Hat, Inc. Linux is a U.S. registered trademark of Linus Torvalds. UNIX is a registered trademark of The Open Group.

XXXX-XXXXEN, 4/2006

