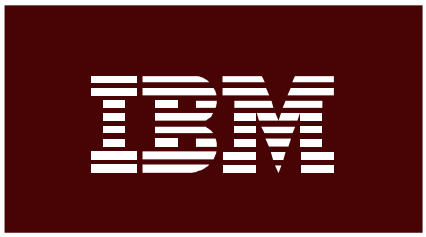

IBM and Novell:

The SUSE Linux Enterprise Server 10 Success Story

Includes SLES 10 SP1 Enhancement

White Paper

July 2007

Introduction

IBM® and Novell have collaborated to integrate new technologies in the Linux® kernel that will help customers improve efficiency and gain a competitive edge in their industries. Novell's SUSE Linux Enterprise Server (SLES) 10 delivers new functionality, improved scalability and increased performance.

IBM solutions on IBM System x™, IBM System p™, IBM System z™, and IBM BladeCenter® servers are enhanced by more than 180 IBM-requested features included in SLES 10 and over 90 IBM-requested features included in SLES 10 Service Pack 1 (SP1). These features support new IBM hardware (devices and processors), cross-platform functionality, and the interoperability of IBM Middleware and hardware. IBM has implemented customer requirements, gained community acceptance, and performed extensive testing of SLES 10 and SLES 10 SP1 on IBM platforms. IBM's collaboration with Novell has continued to advance the SLES enterprise solution with fast time to production, new technology, and superior quality.

This paper focuses on features with special value for IBM platforms and IBM contributions.

The following sections give an overview of key enhancements by major area and platform:

- Virtualization
- Kernel
- Networking
- Reliability, Availability, Serviceability
- Security
- Toolchain

For a complete overview of SLES 10 and SUSE Linux Enterprise Desktop (SLED) 10 content, visit the Novell web pages.

- For SLES 10: http://www.novell.com/news/press/item.jsp?id=1031
- For SLES 10 SP1: http://www.novell.com/news/press/novell-ships-suse-linux-enterprise-10-service-pack-1-and-new-virtual-machine-driver-pack/

Virtualization

This chapter highlights SLES 10 specific enhancements for virtualization on System p, System z, and System z9™ servers.

Virtualization on System x: Xen

One of the key highlights of SLES 10 on System x is the virtualization provided by Xen, a lightweight open source hypervisor. Xen provides a single System x servers with the capability to run multiple operating systems simultaneously in an isolated manner. This allows sparsely used machines running different operating systems and applications to be consolidated to a single machine, helping to reduce server hardware and operating costs.

Xen supports paravirtualized guest operating systems, wherein the guest operating system is modified to run on Xen. In this release, support for paravirtual guest operating systems includes:

- SLES 10 SP1
- OES 2 Linux
- OES2 NetWare 6.5 SP7
- SLED 10 SP1
- SLES 9 SP3 (technical preview)

Xen also supports unmodified guests in hardware-assisted full virtualization mode using either Intel® Virtualization Technology or AMD Virtualization. In this release, support for unmodified guests includes:

- SLES 10 SP1
- SLES 9 SP3
- MS Windows 2003
- MS Windows XP
- Red Hat Enterprise Linux 4 and 5

Support for other operating system releases are planned for future updates. For a current list of supported operating systems, visit http://support.novell.com/products/server/supported_packages/

Xen provides isolated operating environments for applications by running on independent operating system instances. Thus, the impact of a failure of one application on another is reduced without requiring the use of independent hardware. Typically, the guest architecture must match that of the hypervisor (for example, a 32-bit guest operating system on 32-bit Xen hypervisor).

Xen 3.0.4 in the SLES 10 SP1 release provides the following features:

- Support for IBM's Secure Hypervisor Initiative components, such as the Trusted Platform Modules, which provide hardware-based security and trust, and support for other security mechanisms to provide isolation and access control between domains and system devices.
- Virtual devices presented to guest operating systems, which can protect them from or minimize the impact of failures in real devices.
- Support for up to 32-core SMP systems, and hot-plug reprovisioning of processors to adjust processor resources dynamically.
- Support for Physical Address Extensions (PAEs), which allows 32-bit operating systems to address more than 4GB of memory. 32-bit PAE mode guests are supported on 64-bit XEN.
- Checkpoint and restart functionality for paravirtual guest operating systems, as well as support of the migration of a paravirtual operating system instance to another physical server; this functionality helps minimize downtime due to server upgrades, hardware failures, and maintenance.
- Graphical and command-line virtual machine management tools for easy administration and configuration.

IBM Virtualization Manager 1.1, an extension of IBM Director which allows users to manage both physical and virtual systems from a single console, now includes Xen management. Virtualization Manager simplifies management of both VMware and Microsoft® Virtual Server environments. IBM has delivered an implementation, based on the Distributed Management Task Force's (DMTF) Common Information Model (CIM) standard for the Xen virtual environment. Virtualization Manager integrates with, and complements, VMware VirtualCenter, linking together management for physical and virtual resources. Ongoing OpenSource Xen CIM development with IBM's current Director-specific CIM development will be part of future IBM Virtualization Manager releases

Virtualization on System p servers

System p Advanced POWER™ Virtualization (APV) now scales from Linux images running on one-tenth of a processor to up to 64 processors. IBM’s implementation Micro-Partitioning™ technology, an option that is unique to POWER5™ and POWER6™ servers, brings this function to a broader class of Linux environment clients and applications. The Linux operating system has been adapted and optimized for virtualization. This virtualization of physical processors in POWER5 and POWER6 systems introduces an abstraction layer that is implemented within the hardware microcode. From an operating system perspective, a virtual processor is the same as a physical processor. The key benefit of implementing partitioning in the hardware is to allow supported operating systems to run on POWER5 and POWER6 technology with little or no changes.

POWER5 and POWER6 also allow for on-the-fly changes to virtualized resources. This dynamic approach to virtualization adds significant flexibility in managing multiple server images in a single Power Architecture™ technology-based system. For example, a virtual server can be instructed to borrow idle processing power from other virtual servers on the same system. Such servers can instantly accommodate increasing workloads without user intervention. This technology is part of APV.

Virtualization on System z9 and System z (LPAR and z/VM) servers

System z servers provide both hardware (LPAR) and software virtualization (z/VM®). Each of the two virtualization schemes is proven to be robust and to satisfy the needs of mission-critical services in an enterprise. SLES 10 can run in either an LPAR or a z/VM guest and thus inherits the flexibility, scalability and manageability benefits of System z virtualization. In particular, z/VM V5.2 improves the scalability to host servers with large memory requirements and high network and storage traffic. Using SLES 10, you can benefit from various improvements related to accurate accounting, performance, and RAS. (See "Key enhancements by platform: System z and System z9.")

Kernel Enhancements

SLES 10 provides improved features and functionality for the Linux kernel, as well as core enhancements that help to improve stability and scalability. Key kernel enhancements include the following:

- Virtual File System (VFS) enhancements to support shared subtrees of name spaces. These enhancements allow full support for Rational® ClearCase® and Multi-Version File System (MVFS) software available from IBM.
- madvise (MADV_REMOVE) and lazy paging support for shared memory allocations, which enable autonomic management of DB2® software.
- ext3 file system enhancements that allow users to add new space to an existing partition without unmounting the file system first (online resize).
- Robust Virtual Machine design to handle memory pressure handling.
- Flexible resource management (CPUSETS, NUMA enhancements, Pluggable I/O schedulers and fork/exit event notification).
- Delay accounting statistics for enabling enterprise workload manager (eWLM).
- Enhanced device driver support.
- Hotplug Memory Add Support for POWER, which allows dynamic increases of memory in the system (in particular, for virtual environments).

SLES 10 SP 1 delivers the following kernel enhancement:

- madvise (MADV_REMOVE) and lazy paging support for shared memory allocations provides better VM statistics counters and a fully enabled Self Tuning Memory Manager (STMM) feature for DB2 software.

Networking

IBM's contributions to the new networking functionality of SLES 10 include the following:

- Network File System version 4 (NFSv4) is a standards-based state-oriented network file system protocol with built-in security support. NFSv4 support, which is available through the SLES 10 kernel, has been further stabilized with patch contributions from IBM.
- IPv6 Advanced Sockets API enables "advanced" IPv6 applications to access features, such as interface identification options and IPv6 extension headers that are not addressed in the basic sockets API (RFC 2133). The IPv6 Advanced Sockets API also enables hop-by-hop options, destination options and other features making it compliant with RFC 3542.
- The Stream Control Transport Protocol (SCTP) connectx() API in SLES 10 provides a faster way to set up an association with a multi-homed server in a situation where the client is aware of the multiple IP addresses of the server. SCTP allows the caller to specify multiple addresses at which a peer can be reached and let the kernel do the retries.

Other SCTP enhancements include Multiple SCTP bug-fixes, performance enhancements and updates to the API to bring it closer to the latest SCTP sockets API draft.

SLES 10 SP 1 delivers the following networking enhancement:

- The zebra route management daemon, which is part of the quagga package, now includes support for route matching using quagga's "route-map" and "prefix-list" features. These features allow the filtering of routes that are installed by other routing daemons in the package as well as the modification of matching route properties. As part of this feature, system administrators can now use routes that were acquired by dynamic routing protocols to set the preferred source address to be used by applications. This feature enables the use of Virtual IP Addresses (VIPA) for TCP sessions to provide transparent network failover.

Reliability, Availability, Serviceability

With SLES 10, numerous RAS enhancements are available to IBM solutions. Kdump, a new crash dump utility, is a First Failure Data capture (FFDC) mechanism that is supported across multiple architectures. Unlike other crash dump utilities, Kdump is reliable because it doesn't depend on the dying kernel to capture the dump. SLES 10 includes the enhanced analysis tool, Crash, for advanced dump analysis capabilities.

IBM brand hardware capabilities, error recovery, and remote management and service capabilities, are enabled in SLES 10. IBM solutions with SLES 10 deliver improved error handling at the hardware level and memory parity error checks. IBM has further improved the enhanced error handling (EEH) capabilities on System z and System p servers. On Power PC® processor-based hardware, EEH helps to detect and recover from a wide assortment of PCI device and bus errors. System z servers now have better instrumentation to analyze problems with SCSI/FCP and networking in guest LANs. IBM platforms also have the ability to manage and service systems remotely with the help of remote management cards in the system.

SLES 10 provides advanced performance analysis through the SystemTap debugging tool. SystemTap offers a safe and easy scripting language for system administrators to use in analyzing performance problems that originate anywhere from application programs to device drivers. The tool is also sophisticated enough to support the on-demand probing required in the debugging sessions of experienced system programmers. Because SystemTap contributes no overhead when not in use and very low overhead when in use, it is a highly accurate tool for real-time performance analysis of production systems.

SLES 10 SP1 enhances the capability of SystemTap by adding support for the S390 platform.

Security

SLES 10 delivers security enhancements to provide advanced cryptographic performance, hardware-enabled security, workload isolation, data confidentiality, and regulatory compliance. IBM delivers the following contributions:

- Enhancements to openCryptoki, a PKCS#11 implementation that exploits the new embedded System z instructions for improved cryptographic performance.
- A new dynamic OpenSSL engine that exploits the new embedded System z instructions for improved cryptographic performance within OpenSSL.

Support for the Trusted Computing Group (TCG) version 1.1 Trusted Platform Module (TPM). The TPM device enables smart card-like support to provide secure key storage to reduce software attack vectors.

SLES10 SP1 delivers the following security enhancements:

- A new openCryptoki token that further exploits the IBM 4764 PCI-X Cryptographic Coprocessor (PCIXCC) card capabilities on System z machines. This new token takes advantage of the cryptographic functions and encrypted key values of the PCIXCC card.
- Continuing on the success of previous common criteria certifications, IBM and Novell are working to certify SLES10 SP1 across multiple IBM hardware platforms. For the latest status on common criteria certification at the CAPP/EAL4+ level, visit: http://www.niap-ccevs.org/cc-scheme/in_evaluation.cfm

Toolchain

SLES 10 brings a number of toolchain-related (gcc, binutils, glibc) security and performance enhancements.

For 32-bit Power PC, SLES 10 provides the -msecure-plt option and makes it the default. The msecure-plt option allows applications to be divided into pure read-only executable code and read-write no-execute data areas. This is important to preventing so-called buffer overrun attacks.

SLES 10 provides improved signal handling for 32- and 64-bit Power PC so that the run-time stack does not need to be executable.

The gcc compiler added the -fstack-protector feature that detects buffer overflows on the stack.

The gcc-4.1.1 compiler provides numerous code generation improvements over previous gcc-3.3 and gcc-3.4 compilers. Using SPECint2000 tests, Power PC gains are estimated at 10% improvement for 32-bit when comparing SLES 9 with gcc 3.3, and 13% for 64-bit when comparing SLES 10 with gcc 4.1 on a 4-core POWER5 system. SLES 10 also offers improved VMX/Altivec support for the JS20 Blade with autovectorization features of gcc 4.1.

SLES10 SP1 updates gcc-4.1.1 with POWER6 specific tuning options (-mcpu=power6, -mtune=power6). gdb was updated to the 6.6 level.

Key Enhancements by Platform

This chapter presents key enhancements for System p, System x, System z, and System z9 hardware platforms.

System p (Power Architecture)

Linux on POWER processor-based servers provides supported applications with the following qualities:

- high performance, high scaling, and high capacity within a full family of servers
- high reliability
- leadership virtualization capabilities
- low operating costs compared to alternative server architectures

To help clients provide their companies, customers, and shareholders with competitive advantages and top value, IBM drives innovation and value into their hands through the proven, open, and powerful computing platform of Linux on POWER. SLES 10 on Power Architecture provides significant hardware reliability and serviceability, an efficient native virtualization environment, and critical system scalability (with regard to number of processors, amount of physical memory, and I/O bandwidth).

System p servers, using IBM POWER5, POWER5+™, and POWER6 processors, help customers simplify their IT infrastructures at a low cost and increase flexibility while helping them to improve overall performance and efficiency of operations. IBM System p5™ servers running SLES10 have achieved industry-leading benchmark results for business and scientific performance. For more information on performance results, see: http://www-03.ibm.com/press/us/en/pressrelease/20745.wss.

SLES 10 delivers the following features for POWER:

- 64-core SMP allows execution of single large OS images that require large amounts of computing power operating on a single large shared memory. This feature is useful in certain high performance computing applications.
- NUMA optimization improves the performance of workloads that run across large (greater than 4-8 processor cores, by specifying that memory may only be accessed by local processors.
- The iSCSI software initiator allows access to remote iSCSI storage devices across standard (Ethernet) network connections, providing a low-cost remote storage solution.
- IBM's XL C/C++ Advanced Edition V8.0.1 for Linux and XL Fortran Advanced Edition V10.1.1 for Linux offers numerous enhancements for use on SLES 10. The latest IBM compilers offer

optimization and performance-tuning features to exploit all Power PC systems, including the latest POWER5, POWER5+, POWER6, and Power PC 970 hardware architectures. The -qarch and -qtune compiler options include new suboptions to provide code tuning for the new POWER5+ processors. For more information on the IBM compilers, visit:
    - http://www-306.ibm.com/software/awdtools/xlcpp/features/linux/xlcpp-linux.html
    - http://www-306.ibm.com/software/awdtools/fortran/xlfortran/features/linux/xlf-linux.html
- Post-Link Optimizations for Linux on POWER push performance improvements to applications by reorganizing the compiled application code to reduce page faults, cache misses, and other impediments to performance. For more information about Post-Link Optimizations, visit alphaWorks® at: http://www.alphaworks.ibm.com/tech/fdprpro
- Portions of the operating system (such as kernel, libraries, tools) have been tuned to provide better performance on POWER systems.
- VMX exploitation improves performance for applications that require “vector operations” by providing a single instruction that operates on multiple data items at once. This capability is available in systems built around 64-bit Power PC 970FX and POWER6 processors.
- Large page support improves performance by using a smaller number of translations that the processor must address (and can cache in the processor) for a given large contiguous range of virtual memory.
- Transparent large page support is provided through the new libhugetlbfs feature. The transparent large pages feature allows applications to take advantage of the Power 16MB large pages with no application source code changes. Mallocs can be backed with large pages with the setting of environment variables. An application’s .bss, .data, and .text segments can be backed simply by relinking the executable.
- Device Driver error log analysis provides a mechanism for analyzing the errors generated by a device and recommending replacement actions.
- iSCSI TCP/IP offload adapter support provides a high performance mechanism for communicating with iSCSI storage devices by offloading the work of the TCP/IP stack to the adapter rather than the operating system.
- The ltrace, itrace, and oProfile tools allow tracing and performance analysis of the operating system and applications to determine where performance problems are occurring so that they can be addressed. The itrace tool provides input to the IBM Performance Simulator for Linux on POWER. The IBM Performance Simulator for Linux on POWER is also available from alphaWorks at: http://www.alphaworks.ibm.com/tech/simppc
- Physical memory add allows more physical memory to be added to a Linux partition without rebooting the partition. This allows some adjustment to the size of a partition to accommodate larger workloads. Physical memory cannot be removed from a Linux partition without rebooting the partition. Processors can be added and removed from a Linux partition without rebooting it.
- Serial Attach SCSI are a new type of SCSI storage device that use serial rather than parallel communication paths. SCSI storage devices are transitioning from parallel to serial interface types.
- PCI-Express support allows BladeCenter JS21 systems to use a faster, higher bandwidth path to the I/O devices, providing better I/O performance and scalability.
- Trusted Computing Specification Support provides a mechanism for applications to validate that they are running in an approved environment for processing digital media and other protected content.
- Non-executable stack and heap is a feature that prevents applications from executing code in their stack or heap areas. This feature thwarts the common security problem of buffer overflow attacks, where a security attack causes the overwrite and overflow of application code residing in a buffer that is either in the stack or the heap.
- High performance time system calls allow timestamps to be generated more quickly, improving the performance of transactional systems like databases.
- 64-bit tracing tools allow tracing of 64-bit applications, easing the debugging and tuning of 64-bit applications.
- Improved overall processor performance through libraries that tuned for specific processor types. Since different processors use unique instruction sets and instruction scheduling mechanisms, they experience different performance. This feature causes the fastest running code to be engaged for the actual processor in use.

SLES 10 SP 1 features enhanced processing for POWER6 systems. This includes the following:

- When a processor fails, it is replaced with an unused processor or a processor from a shared pool, and execution continues. (A degradation in overall system performance may result.) This feature increases the reliability of POWER6 systems and improves system availability through reduced downtime.
- PCI-Express support allows POWER6 systems to use a faster, higher bandwidth path to the I/O devices, helping to provide better I/O performance and scalability.
- This service pack adds POWER6 tuned libraries.
- The Post-Link Optimization tool is updated to support POWER6. For more information about Post-Link Optimizations, visit alphaWorks at: http://www.alphaworks.ibm.com/tech/fdprpro

System x

SLES 10 offers improved functionality and increased scalability on System x servers. System x servers utilize Intel and AMD processors, which represent the majority of Linux installments worldwide. The System x platform differentiates itself from other Intel/AMD systems by providing world-class customer satisfaction and assurance with respect to robustness, reliability, availability, and stability. System x includes features that are typically only found in higher-end mainframe environments. SLES 10 leverages these high-end features to provide a superior hardware/software environment.

SLES 10 provides the following features to further expand System x solutions:

- Dual core support (two processor cores per socket) is provided for both Intel and AMD processor packages.
- IBM Extended X-Architecture® (EXA) chipset support enables support for large SMP NUMA-based IBM servers, primarily with 64-bit kernels. The EXA chipset provides the I/O bridge controller and the cache/memory and scalability controller for high-end NUMA platforms.
- Global timesource is provided to synchronize timekeeping for NUMA systems. Support for both the High Precision Event Timer (HPET) and the ACPI PM timer are used to track time consistently in environments with multiple clock sources.
- Intelligent Platform Management (IPMI) support using the OpenIPMI driver and user level interfaces (such as ipmitool) provides the ability to manage and monitor platform hardware in-band or out-of-band. It includes IPMI 2.0 compliance, extensions for SoL, VLAN, and DHCP firewall support.
- Serial Attached SCSI (SAS) support, which is SCSI protocol with a serial (as opposed to older parallel bus-based) interface, provides improved I/O bandwidth and performance.
- Advanced Configuration and Power Interface (ACPI) extensions for system reset, PM timer, PCI hotplug, power management, and processor throttling,
- Support for legacy-free System x blades.
- New System x device support is provided, including ServeRAID™ 8i/8k driver (aacraid) support and synchronization, Broadcom 5706/5708 (bnx2) NIC support, Broadcom 5714/5715 (tg3) NIC support and nVidia ATI graphics accelerators.
- Execute Disable (XD), an Intel processor feature that protects against execution of malicious software at the hardware level, is provided.
- Active PCI – IBM hot plug features that extend existing PCI hot plug capabilities – PCI slot/device identification and notification handler extensions.
- iSCSI allows for remote access to storage devices over standard ethernet connections.
- PCI Express - offers extended PCI configuration space, improved bandwidth and faster speeds to improve overall performance and scalability.
- System x 4-node x460 with support for up to 64 logical processors and 256GB memory.
- NUMA support, which enables the kernel and applications (through export APIs) to make intelligent decisions with respect to memory placement and process scheduling on systems that do not have uniform memory characteristics (for example, local and remote memory latencies differ).
- Improved Power Management capabilities are included that provide power cost savings by throttling processor frequency on unused and under-used processors. Support for Demand Based Switching (Intel) and PowerNow! (AMD)

SLES 10 SP1 delivers the following features for System x servers:

- New device support for System x servers: Broadcom 5722 (tg3) NIC, Broadcom 5790 (bnx2) NIC, NetXen 10Gb NIC support (netxen_nic), Adaptec 29320LPE (aic79xx), LSI Logic 1078 SAS (mptsas, megaraid_sas), Qlogic QLE234X (qla2xxx), Qlogic QLA40XX (qla4xxx), Qlogic iSCSI support (qla3xxx), Quantum GoVault (ahci, ata_piix), Intel PRO/1000 quad port NIC (e1000), extensions and device hardening for the Adaptec Razor SAS controller (aic94xx).
- Support for new Intel and AMD processors and chipsets used in recent System x servers.
- IOMMU support for IBM Calgary PCI Host Bridge controller. Provides improved data protection.
- iSCSI Boot Firmware Table (iBFT) support.
- Latest OpenIPMI driver and IPMI utilities. Provides support for multiple BMCs in a single system.
- Improved hot add memory support. Available on x260, x366, x460 (single-node), x3800, x3850, x3950 (single-node).
- Improvements to dmraid, device mapper providing an open source alternative to some on-disk proprietary RAID formats (e.g., Adaptec HostRAID).

System z and System z9

SLES 10 on System z servers provides a cost-efficient and reliable scale-up and scale-out application-hosting environment specifically meant to augment z/OS® enterprise deployments. To serve this goal, SLES 10 exploits the latest features of System z9, the newest System z architecture level.

A new Fibre Channel Protocol (FCP), SCSI over Fibre Channel, Host Bus Adapter Virtualization Technology (N_Port ID Virtualization: logical WWPNs) on System z9 now allows sharing of FCP adapters in a fully SCSI standards-compliant way. This comprises SAN access right management and disk sharing through virtual HBAs.

SLES 10 network adapters have been enhanced to enable the Communication Controller for Linux products on System z9 for 374x NCP virtualization. To provide optimal performance in a security-critical environment, SLES 10 supports the new cryptographic adapters of System z servers and the new crypto instructions for AES and SHA-256 of the System z9 processors, the latter being not only accessible to applications, but also to the kernel.

SLES 10 contains miscellaneous virtual server enhancements that relate to pre-System z9 machines.

- Linux images can use the clock comparator for virtual elapsed time (as opposed to having users consult a wall clock), a feature that is more hypervisor-friendly.
- Network traffic can be analyzed on virtual network connections (Guest LANs) and there is support for concurrent I/O through multiple paths for ESCON/FICON (channel)-attached disk storage subsystems using the System z Parallel Access Volume technology.

With SLES 10 execute in place, support is fully integrated in the ext2 file system. System z customers can leverage this technology by running multiple similar servers in the z/VM hypervisor. It allows for efficiently shared memory for executables and library code across multiple Linux guests running in the same z/VM system.

Further SLES 10 improvements include enhancements for z/OS disaster recovery integration for planned outages both for Linux in LPARs and Linux operating with FCP I/O attachments, and support of IBM Director 5.1 to provision and manage Linux images in z/VM.

SLES 10 SP 1 delivers the following features for IBM System z:

- Support for Pseudo Random Number Generator (PRNG) integrated into the CPU Cryptographic Assist Facility with the IBM System z9.
- Support for secure key cryptographic functions.
- Support for IBM 3592 tape encryption allowing users to archive encrypted data and to share data securely with z/OS.

- Gathering hardware-provided metrics for an FCP channel, relating to FCP performance (latencies) and FCP channel usage; this feature requires IBM FICON Express4 and FICON Express2 features on z9 EC and z9 BC.
- Collaborative Memory Management Stage II provides support for the Collaborative Memory Management Assist (CMMA) in z/VM 5.3, reducing hypervisor paging I/O overhead. The Linux support has to be activated per IPL-option cmma=on. (The default is cmma=off.)
- Data execution protection for user processes provides execution protection for code in user space data segments that are marked as not executable. It can prevent, for example, stack overflow exploits and generally makes a system insensitive to buffer-overflow attacks in user space.
- Simple network IPL - Linux Image Control (snIPL) for LPAR / VM) is a command line tool that serves as a remote control using basic System z Support Element (SE) functions or basic z/VM system management functions. It can also be used as a virtual power switch in a clustering environment.

IBM Corporation
Marketing Communications
Systems Group
Route 100
Somers, New York 10589

Produced in the United States of America
August 2007

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries.

The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only. Any statements about support or other commitments may be changed or cancelled at any time without notice. This Information is provided "AS IS" without warranty of any kind.

IBM, the IBM logo, alphaWorks, BladeCenter, ClearCase, DB2, Micro-Partitioning, Power Architecture, POWER, POWER5, POWER5+, POWER6, Power PC, Rational, ServeRAID, System p, System p5, System x, System z, System z9, X-Architecture, z/OS, z/VM are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at http://www.ibm.com/legal/copytrade.shtml.

Intel is a registered trademark of Intel Corporation in the United States and/or other countries.

Linux is a trademark of Linux Torvalds in the United States, other countries or both.

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

Copying or downloading the images contained in this document is expressly prohibited without the written consent of IBM.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

SLES 10 is a non-IBM product and is not warranted by IBM. Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

LXW03001-USEN-00

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven, ClusterProven or BladeCenter Interoperability Program products. Support for these third-party (non-IBM) products is provided by non-IBM Manufacturers.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

The IBM home page on the Internet can be found at http://www.ibm.com.

The System p home page on the Internet can be found at http://www.ibm.com/systems/p