SUSE Linux Enterprise

What's new with SLES10 and SP1 for System z

Mark Post
Technical Support Engineer IV
Linux Deployment SWAT Team
mpost@novell.com

Novell®

z/VM Maintenance Needed

- Read the release notes on the installation media. It has important information about z/VM maintenance that is required if you're running older releases. (No, I don't want to hear about who's still running 3.1 ;)
- Check the with the IBM Support Center for any recommended z/VM maintenance that is Linux-related
  - Almost always some virtual networking, etc. fixes are available

Updating from SLES9

- Updates from SLES 9 to SLES 10 are supported starting from one of the following bases:
  - SLES 9 GA
  - SLES 9 SP3
  - SLES 9 SP3 plus the latest patches from the maintenance web site
  - To verify whether one of the above variants is installed, you can use the tool “SPident -vv” to show the current level of your system.
- Update a system by starting the SLES 10 installation system and choosing “Update” instead of “New installation.” **Don't** try to do it from YaST.

SUSE Linux Enterprise 10

- Novell Customer Center
  - One location to manage your subscriptions, software updates and online support requests for your SUSE Linux Enterprise products.
- Online Updates
  - Still use YaST, but online_update command is gone
    - Replaced by rug command (yet another learning curve)
  - System registration is required
    - Can be done at install time, or later via YaST
  - Seriously look at using YUP (YUM Update Proxy) to mirror updates locally
  - I'll be publishing a document on this Real Soon Now [tm]

Deprecated Features

- **31-bit operating system** (“IBM made me do it”) 31-bit applications are supported via the 31-bit emulation layer
- CLAW
- CTC, virtual and real
- IUCV network device - Note that the IUCV-infrastructure is **not** being deprecated
- JFS (the kernel module is there, just no support in YaST) Use EXT3 or XFS instead

Deprecated Features

- **Native** FBA DASD devices. The DASD device driver will continue to support FBA for the following scenarios:
  - Virtual FBA devices (using FBA channel programs or DIAG access method)
  - FBA emulated SCSI devices (z/VM minidisk support)
- Support for ttys over CTC connections. (The ctctty code never really worked and complicated things in the CTC driver.)

Major Package Versions

- Kernel 2.6.16 + IBM updates
- s390-tools 1.5.3 (Contains vmcp, but we still ship cpint)
- gcc 4.1.0
- binutils 2.6.91.0.5
- glibc 2.4
- Apache 2.2.3
- Samba 3.0.22
- IBM Java 1.4.2
- Check IBM developerWorks for info on “restrictions”
  http://www.ibm.com/developerworks/linux/linux390/october2005_restrictions.html

New Features

(* = new in SP1)

- Accurate CPU time accounting (no more weirdness when viewing data from inside the Linux guest).
- vmconvert utility (supports 64-bit z/VM dump format)
- ctcmpc: network device driver supporting SNA MPC CTC-channels - needed to run Communications Server for Linux on System z.
- Support for OSA-Express2 OSN - Open Systems Adapter for NCP (CDLC support)
- * Channel path measurements and channel path measurement characteristics per channel path.

New Features

- 3590 tape device driver now Open Source - no longer an OCO module. No remaining OCO modules!
- * 3592 tape device support
- * Support 3592 hardware encryption for channel attached tape drives.
- Support for the zfcp HBA API - needed for SAN discovery tool san_disc
- SCSI over FCP n-port ID virtualization (NPIV)
- New zfcp statistics interface - used to gather performance data

New Features

- * Newer versions of the System z FCP HBA provide additional measurement data, including a so-called 'channel latency' (roughly the time a request spent in the HBA) and a 'fabric latency' (the time a request spent outside the System z machine). The new zfcp driver allows users to accumulate these latencies separately for requests with outbound/inbound/no data transfer. These statistics are useful when there is a need to further break down overall latencies as seen by Linux.

New Features

- * zcrypt replaces z90crypt with a new, re-designed crypto device driver. zcrypt features:
  - support for PCICC, PCICA, PCIXCC, CEX2C, CEX2A cards
  - support for clear key and secure key cryptographic functions
  - user space interface compatible with z90crypt
  - a modular design utilizing the Linux device model
  - improved performance by using a poll thread and an advanced load balancing algorithm

Only "secure key cryptographic functions" is new compared to the existing z90crypt support.

New Features

- * Collaborative Memory Management Stage II. Support for CMMA in z/VM APAR VM63856 and z/VM 5.3 reduces hypervisor paging I/O overhead. See 'Improved memory management for Linux guests' in IBM U.S. Software Announcement Letter 207-019 "IBM z/VM V5.3 - Improving scalability, security, and virtualization technology IBM".
  The Linux support for CMM2 now has to be activated via IPL option cmma=on (default is cmma=off).
  - This is a change from the previous CMM2 add-on patch which activated the support automatically when the necessary prerequisites were available.

New Features

- DIAG 250 support for 64-bit systems (z/VM 5.2/5.3)
- Special DASD request flag "REQ_FAILFAST" to return requests immediately, in case the device is not operational
- Multiple Subchannel Sets (MSS) support on z9-109 allows an increased number of subchannels
- GuestLAN Sniffer for use with z/VM 5.2/5.3 tracing capability on a guest LAN or VSWITCH
- oprofile now supports recording of call graphs for in kernel profiling sessions

New Features

- "V=V QDIO Pass-thru" exploits z/VM 5.2 "QDIO Enhanced Buffer-State Management" (QEBSM) for OSA, HiperSockets, and FCP on z9-109, z990, and z890 - allowing a Linux guest to initiate QDIO operations directly to the applicable channel, without interception by z/VM.

  Linux QDIO driver recognizes (and reacts appropriately) when required hardware and z/VM support (QEBSM) are available and enabled.

New Features

- * Dynamic toggling of qeth and qdio performance statistics.
  The kernel configuration options CONFIG_QETH_PERF_STATS and CONFIG_QDIO_PERF_STATS are replaced by switches to activate/deactivate performance data services dynamically at runtime. The switches are located in
  /sys/bus/ccwgroup/drivers/qeth/<device_bus_id>/performance_stats
  and
  /sys/bus/ccw/qdio_performance_stats

New Features

- * Directed yield of spin locks

  Uses the new diagnose 0x9c in the spin lock implementation to yield the remaining time slice of the virtual CPU that tries to acquire a lock to the virtual CPU that is the current holder of the lock. This reduces the amount of CPU time spent in the spin lock loop for contended locks.

  Using this support requires IBM System z9 and z/VM 5.2 with PTF for APAR VM63952, or z/VM 5.3.

New Features

- * Add sysfs attribute 'status' and generate uevent(CHANGE) for DASD CCW devices. Previously, device status was not available in user-space. As a result, it was not possible for user-space applications to implement appropriate processing.
- * New DASD feature for ERP related logging. It is now possible to enable/disable ERP related logging without a re-compile and re-IPL. An additional sysfs-attribute 'erplog' allows switching the logging non-disruptively.

New Features

- DASD extended error reporting facility that allows you to get detailed information about certain problems in DASD I/O. This information can be used to implement fail-over applications that can recover from these problems.
- * Support for Parallel Access Volumes (PAV) added to the DASD driver. This allows concurrent access to DASD devices with multiple channel programs.
  - Note that this support is new for running Linux on an LPAR; Linux running as a guest under z/VM already has PAV support.

New Features

- * Hypervisor file system "s390_hypfs"
  This new virtual file system is used to access PR/SM LPAR hypervisor performance data returned by DIAG 204. For example, information about physical and logical CPUs is provided. All data is put into different virtual files which can be accessed from user space. All values are represented as ASCII strings.
- * Hypervisor file system "s390_hypfs" for z/VM. (This is an extension of hypfs for LPAR DIAG 204). Data returned by DIAG 2fc is exported using the s390_hypfs when Linux is running under z/VM. Information about CPUs and memory is provided.

New Features

- * Character device driver for writing z/VM monitor records.
  Provides an interface which allows applications to write application-specific monitor data (APPLDATA) to the z/VM monitor stream.
- * Data execution protection for user processes. Also known as "Enhanced Linux system layout" (esl).
  Provides execution protection for code in user space data segments which 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.

New Features

- * Support s390 Pseudo Random Number Generator

  Starting with the z9 the CPU Cryptographic Assist Facility comes with an integrated Pseudo Random Number Generator. The generator creates random numbers by an algorithm similar to the ANSI X9.17 standard. The pseudo-random numbers can be accessed via a character device driver node called /dev/prandom. Similar to dev/urandom any amount of bytes can be read from the device without blocking.

New Features

- * Re-IPL with alternate parameters
  It is now possible to specify a new CCW or FCP re-IPL (reboot) device under Linux. When the user calls reboot, the system automatically IPLs the specified device.
- * Dump on panic
  It is now possible to specify a CCW or FCP dump device which is IPLed in case of a kernel panic.

New Features

- * Enable kprobe support
  Kprobes allows you to trap at almost any kernel address and execute a callback function. register_kprobe() establishes a probepoint and specifies the callback. Kprobes is useful for kernel debugging, non-intrusive instrumentation and testing.
- * Two new fields added to the appldata_os record, steal time and CPU ID. The steal time reports the time stolen by hypervisor scheduling and the CPU ID fixes a problem which occurred with CPU hotplug and the per-CPU data in the appldata_os record.
- Lots of fixes in qdio, qeth, lcs, dasd, cio (common i/o) kernel, zfcp, zcrypt/z90crypt

Still Pending

- Built-in NSS creation
  - Can still be done manually
- Clock synchronization to an external time reference (ETR)
- AF_IUCV protocol support

SUSE Linux Enterprise Server 10 What's not in

- **JBoss** – The JBoss code is not included in SUSE Linux Enterprise, but can be downloaded and is still supported through additional subscription from Novell or JBoss.
- **GFS** – Storage stack from Red Hat. The High-availability Storage Foundation that ships in SUSE Linux Enterprise is a superior solution.
- **Reiser4, Samba v4** – Not enterprise ready for this release.
- **SELinux** – Security MUST be easy. AppArmor is much easier to configure and deploy than SELinux and provides the same level of security.

SUSE Linux Enterprise Server Roadmap

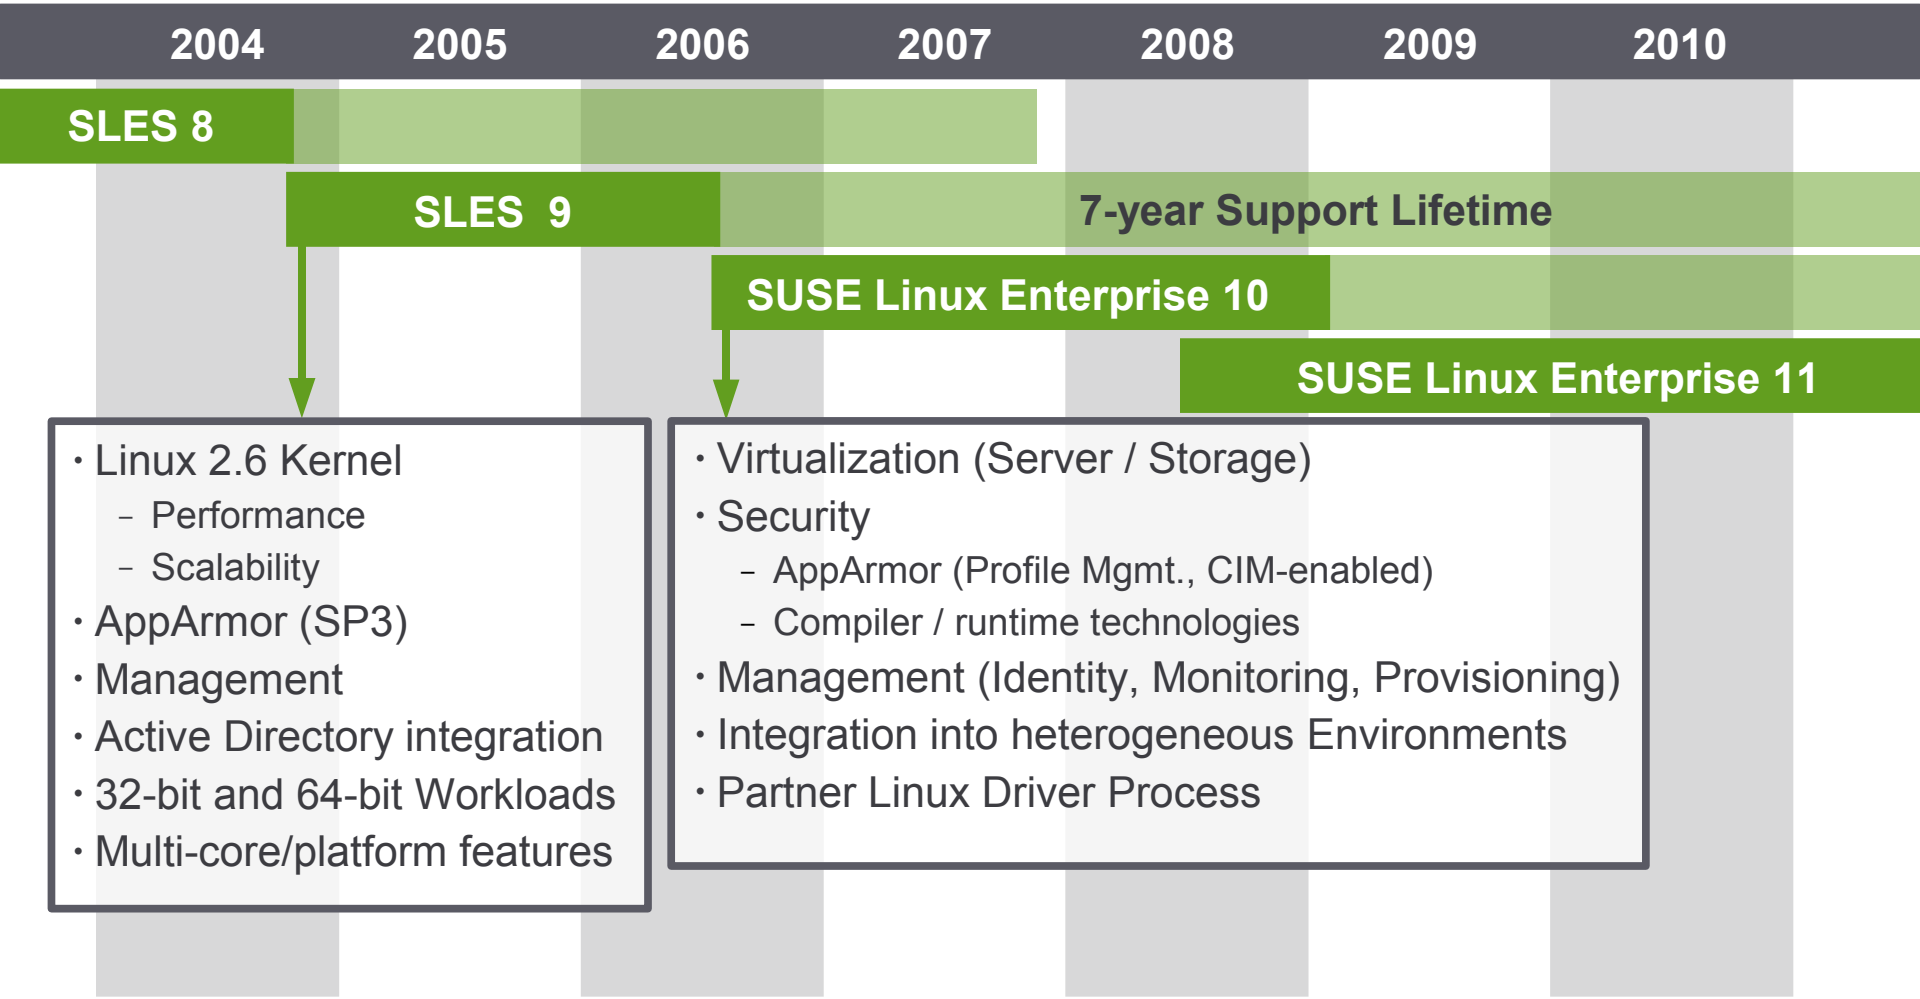

SUSE Linux Enterprise Server Roadmap

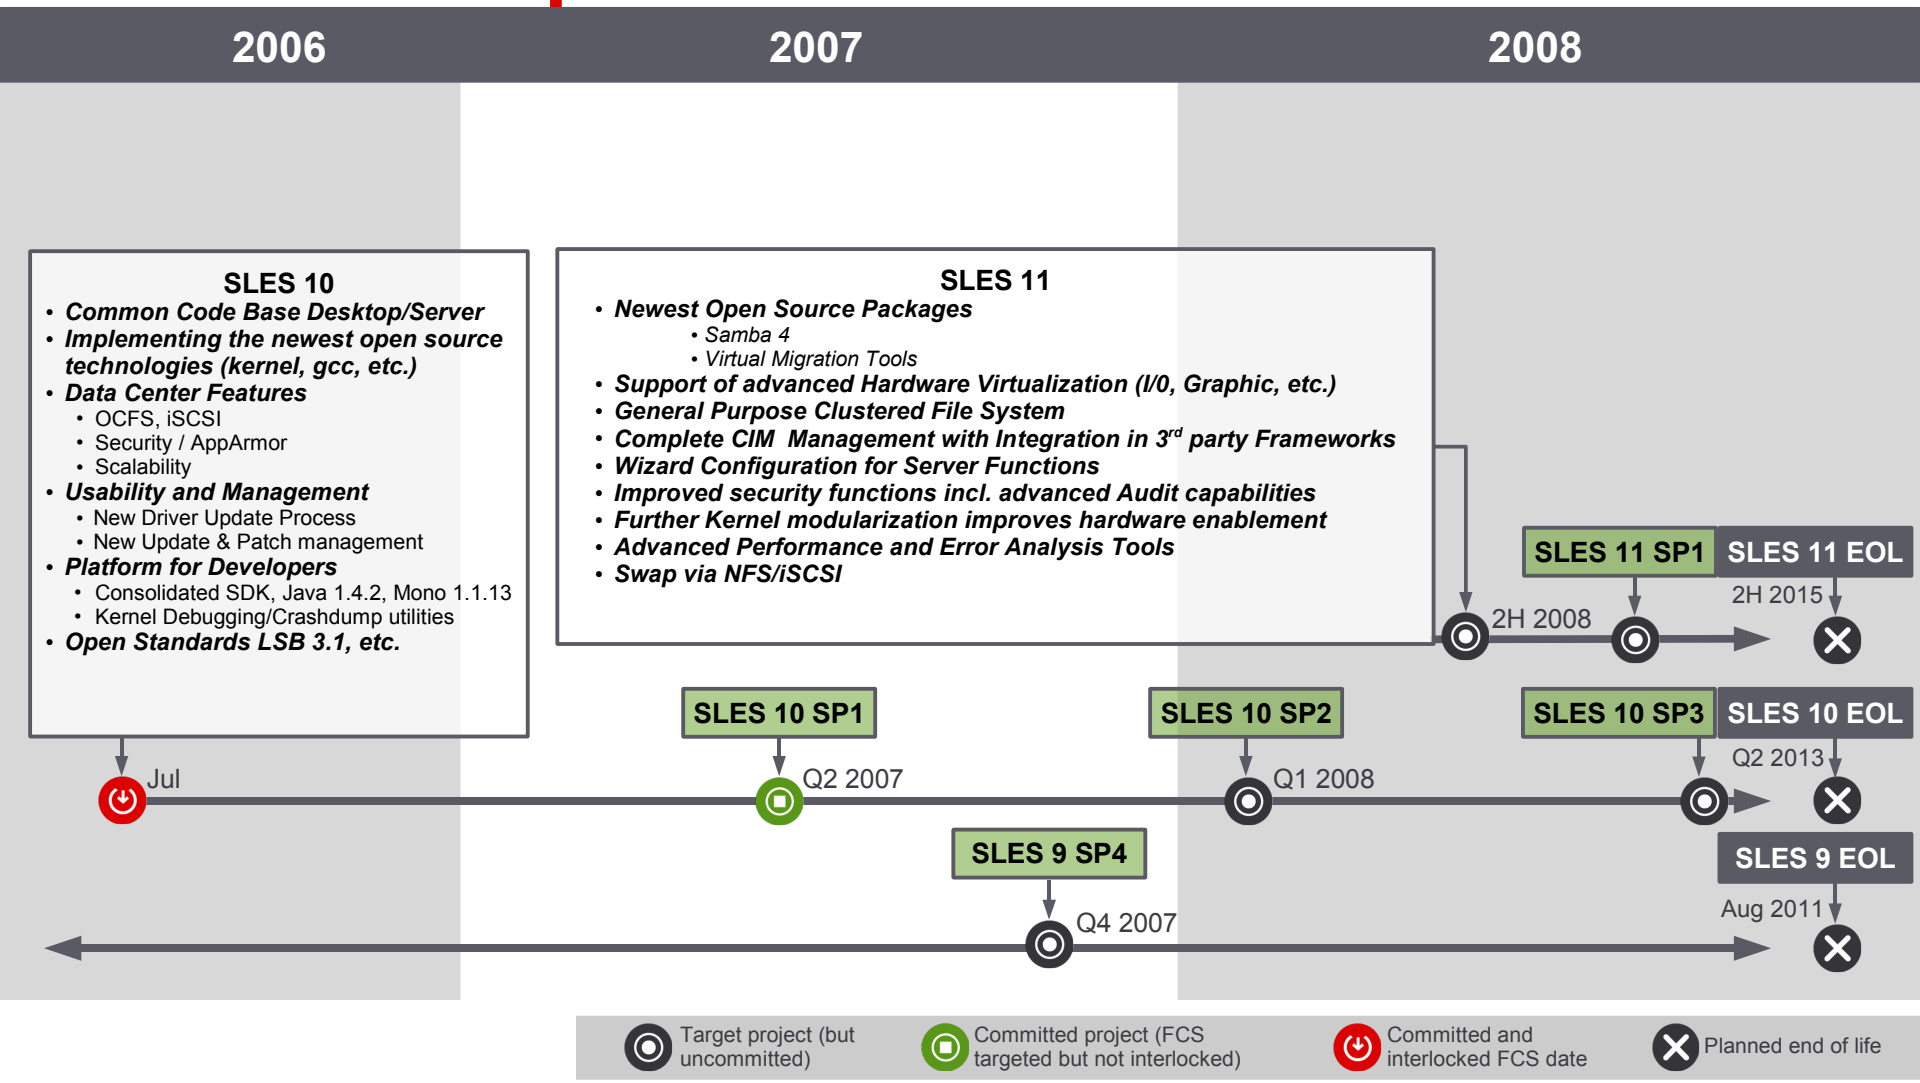

Subscription Pricing Server

Support Included	Basic	Standard	Priority
1 Year	$349	$799	$1,499
3 Year	$873	$1,998	$3,748
1 Year	$11,999	$15,000	$18,000
3 Year	$29,998	$37,500	$45,000

Non-mainframe

- Valid for
  - x86 (32-bit), x86-64 (AMD64 & EM64T), Itanium 2 (IA), IBM POWER (PPC)
- Up to 32 processor sockets per server
- No additional cost per virtual image
- No additional cost for AppArmor, Xen, High Availability Storage Infrastructure

Mainframe

- No longer priced by type of mainframe
- Priced per engine (standard CP or IFL)
- SUSE Linux Enterprise Server 10 only has 64-bit for mainframe, no 31-bit (as per IBM request)
- Prior version rights include 31-bit for SUSE Linux Enterprise Server 8 & SUSE Linux Enterprise Server 9

Support Levels

- Basic
  - “Installation Assistance”
  - http://support.novell.com/linux/linux_install_sppt.html
- Standard
  - Web, email and telephone
  - 12x5 access, 4 hour target response time, unlimited requests
- Priority
  - Web, email and telephone
  - 24x7 access, 1 hour target response time, unlimited requests

High Availability Storage Foundation

SUSE Linux Enterprise integrates Oracle Cluster File System 2 (OCFS2), Enterprise Volume Management System 2 (EVMS2) and Heartbeat v2 clustering services to deliver an entirely open source high availability storage solution

- EVMS2
  - Cluster-aware volume manager
  - Single, unified system for handling all storage management tasks
  - Unparalleled flexibility and extensibility
- OCFS2
  - Symmetrical parallel cluster file system
  - Optimized for fast access to large files
  - Joint development with Oracle
- Heartbeat v2
  - Powerful resource dependency model based on XML
  - Modular design with new cluster resource manager
  - Large clusters: 16 nodes tested, no inherent limit
  - Resources actively monitored for health

Storage Features

DRBD

- Distributed Replicated Block Device
- Integrated with Heartbeat 2

Multipathing

- Improved multipathing support
  - I/O Load balancing
  - HA failover

NFS

- NFS v4 Client and Server
  - Performance improvements
  - Improved security

iSCSI

- iSCSI target and initiator
- “turn SUSE Linux Enterprise Server into a SAN box”

Complete Server Functionality

- File and Print Services
  - Samba 3 with improved integration into Windows environments (authentication against Active Directory)
  - NFS v4
  - CUPS
  - Apple File Protocol (AFP) - NetATalk
  - Support for Linux, UNIX, Windows and Macintosh clients
- Mail Services
  - SMTP / Postfix / Sendmail
  - IMAP / POP / Cyrus
  - SASL
- Security
  - SUSEFirewall
  - ClamAV Anti-Virus Protection
  - Snort
  - VPN with FreeS/WAN

Middleware Services

- Apache Web Server 2.2.0 and scripting support (PHP5, PERL, Python, Ajax, Ruby on Rails, etc.)
- Java Runtime Environment 1.4.2
- Geronimo (Java Application Server)
- Tomcat
- Mono 1.1

Network Services

- DNS / DHCP
- NTP
- NIS
- OpenLDAP
- SLP

Databases

- MySQL
- PostgreSQL
- ODBC connectivity

Questions ?

Novell®

Your Linux is Ready™


Novell®