Red Hat Enterprise Linux6 Roadmap

Tim Burke Vice President of Linux Platform Development, Red Hat June 14, 2010





Agenda

- Major RHEL 6 release themes
- Feature detail

Disclaimers

- Describes major high level features, is not comprehensive of all RHEL 6 bound features.
- This slide-deck does not constitute formal product commitment.





Additional Info

- RH Summit this presentation delivered
 - Wed June 23 10:20-11:20
 - Thurs June 24 10:20-11:20
- Q&A followup RHEL Product Management and I
 - RHEL Campground refer to newspaper for schedule of demos
 - Wed June 23 12:00-2:00
 - Thurs June 24 11:30-12:30
- Feedback / Input
 - tburke@redhat.com





RHEL 6 Schedule

- Beta 1 public availability April 2010 via ftp and RHN for customers
- Beta 2 SOON! shortly after RH Summit
 - Jump in! We'd love to get your feedback
- Release later this year
 - Remember, RHEL 6 becomes available to all customers with active RHEL subscriptions.





Major RHEL 6 release themes

Optimized foundation OS

For large-scale, centrally-managed enterprise deployments. Lower Total Cost of Ownership. Secure. Optimized for maximum efficiency of latest generation of high corecount systems – memory, scalability, RAS, power efficiency. Resource control.

Virtualization – optimize RHEL as host or guest

Deployment, provisioning and flexibility for dynamic workloads from datacenter to desktop. Emphasis on performance, storage flexibility, security, and guest isolation.

Green IT

Through power management and dynamic guest migration

Innovation and technology leadership

With the latest enterprise ready components in Storage and File Systems, Networking, Tools, Cluster, Desktop, Installer, and Services. Providing customer access to leadership technology throughout the RHEL product lifecycle.





RHEL 6 foundation features

- Virtualization making RHEL an optimized host & guest
 - KVM
 - Industry leading virt performance, flexibility, security for both host & guest environments
 - Device I/O optimization

Improved manageability

- For large scale virtualization deployments, server & desktop
- RHEV-M (virt/cloud management) enablers
- Samba enhancements for Windows active directory and file sharing

Power management

- Efficiency lower deployment costs, reduced carbon footprint
- For virt, bare metal, laptop

JBoss

• Hardware level as well as dynamic system service startup and suspend



PRESENTED BY RED HAT

SUMIT

RHEL 6 foundation features (continued)

• RAS (Reliability, Availability, Serviceability)

- Hotplug, memory error reporting, filesystem data integrity
- Support tools automated crash detection and bug reporting infrastructure

Hardware enablement and scalability

- Maximum efficiency with latest generation of highly scalable servers with headroom to grow
- Large configurations (cpu, memory, busses, I/O), NUMA awareness
- UEFI new bios boot loader interface
- Supported architectures: x86, x86-64, PPC64, s390x

Desktop

SUMIT

- VDI- virtualized thin client, SPICE integration
- Mobility dynamic network config

JBoss

• Dis<u>play – extern</u>al monitors, multihead, projectors, docking station



PRESENTED BY RED HAT

RHEL 6 foundation features (continued)

- Innovation leveraging Red Hat's and community innovation post RHEL 5 code base
 - GCC4.4 improved C++, OpenMP, gbd for threaded debugging, new NUMA aware memory allocator
 - Desktop applications, X graphics, audio playback
 - Kernel hybrid of 2.6.32, 33, 34

Developer and diagnostic (serviceability tools)

- Systemtap
- GCC, gbd, eclipse, OpenJDK
- ABI program docs, tooling and services

RHEL ecosystem

- Kernel ABIs (facilitating 3rd party drivers)
- Driver updates for timely hardware support
- Migration guide

JBoss



ESENTED BY RED HAT





SUBSET OF FEATURE DETAIL

Kernel Resource Management

Illustrative cgroup use cases

- Database workload dedicated 90%, background backup utility 10%
- Virtualization hosting provider allows QoS (quality of service guarantees based on pricepoint)



Kernel resource management

Cgroup – Control group

- A control group provides a generic framework where several "resource controllers" can plug in and manage different resources of the system such as process scheduling, memory allocation, network traffic, or IO bandwidth.
 - Can be tracked to monitor system resource usage
 - Sysadmin can use tools to allow or deny these groups access to resources

Memory resource controller

- Isolates the memory behavior of a group of tasks cgroup from the rest of the system (including paging). It can be used to:
 - Isolate an application or a group of applications
 - Create a cgroup with limited amount of memory

Cgroup scheduler

- CFS Hierarchical proportional fair scheduler (SCHED_OTHER)
- Static priority scheduler with constant bandwidth limits (SCHED_FIFO)





Kernel resource management (continued)

I/O controller

• Designate portion of I/O bandwidth (based on controller queue depth)

Network controller

 Define classes & queues between generic network layer and NIC. Tagging packets with class identifier with different priorities, placing outbound packets in different queues for traffic shaping.

Libcgroup

SUMIT

PRESENTED BY RED HAT

- SELinux policy
- Cgroup creation, deletion, move and configuration management.
- Rules based automatic task placement, PAM module, daemon, uid/gid based rules

Illustrative cgroup use cases

JBoss

- Database workload dedicated 90%, background backup utility 10%
- Virtualized hosting provider allows QoS (quality of service guarantees based on pricepoint)



Power management

• **Objective** — reduced deployment costs through efficiency

Kernel

- Tickless kernel fewer interrupts, more idle time to drop to lower power states x86/x86-64 only
- ASPM (Active State Power Management) PCI Express reduced power states on inactivity
- ALPM (Aggressive Link Power Management) SATA links in low power mode when no I/O pending
- Energy efficient turbo and deep C states
- **Relatime** drive access optimization reducing filesystem metadata write overhead
- Graphics power management

System services / daemons

Intelligent drive spin down

JBoss

WORLD

- Application audit and redesign where necessary to be event based rather than needless polling
- TuneD adaptive tuning daemon powerdown idle peripherals & latency policy scripts.
 Providing a variety of power tuning pre-canned profiles

Virtualization management – RHEV-M integration

 Systems which are powered off are the most efficient. Workload consolidation with power management automatic migration policy

PRESENTED BY RED HAT

SUMIT



Power monitoring tools

Powertop

Identifies power hungry applications and system services

Battery Life Tool Kit (BLTK)

- Workload test framework for typical laptop use cases
- Tuned adaptive tuning daemon
 - Power down idle peripherals
 - Latency policy scripts
 - Provides a variety of power tuning profiles
- Documentation to application developers on tips and tricks





Runtime idle power consumption



PRESENTED BY RED HAT

Reliability, availability, servicability (RAS)

• **RAS hardware enhancements - Objective** – improve availability by coupling advanced error recovery with enhanced logging/reporting of errors

CPU and memory hot add

• Single socket hot add

Machine checks

- Finer granularity memory errors kill isolated processes or fail IO rather than panic system
- Corrected machine check interrupt (CMCI) handling, interrupt based rather than polling, so more rapid correction handling (Nehalem)
- New mcelog for x86/64 architecture
- Uncorrectable memory error recovery for Nehalem EX

HWPoison

SUMIT

- OS declares a page "poisoned", kills the processes associated with the page and avoids using the page in the future
- Enhanced error reporting for PCI devices (PCI-AER & APEI)
 - Enabling fine grained reporting, response and recovery of faulty components



PRESENTED BY RED HAT

JBoss

WORLD

Reliability, Availability, Serviceability (RAS)

- **Ext4** filesystem repair time performance enhancements (4 times faster than ext3)
- **DIF/DIX** storage integrity via hardware checksums providing end-to-end checksumming from application to storage platter.

ABRT (Automated Bug Reporting Tool)

- A more consistent way to report system exception conditions
- Will provide faster turnaround on problem resolution especially for duplicates, allowing better trend analysis
- Highly configurable plugin architecture allows choice of destination
 - "Phone home" bug reports to Bugzilla in the RHEL6 Beta
 - Or save locally, or send email
 - Connected to Red Hat Customer Portal by GA
 - Easy to develop "Send to local helpdesk"
- Applet/syslog event informing user of new crashes
 - GUI for reporting bugs





Kernel scalability

- Objective providing scaling headroom anticipating many years of upcoming hardware generations. Tested and supported limits will likely grow over the course of product lifespan.
- Scalability features enhancements in algorithms. Applicable to bare-metal and virtualized guests
 - Split LRU VM different eviction policies for file backed vs swap backed
 - Ticket spinlocks fixes NUMA starvation
 - CFS scheduler better NUMA balancing
 - UEFI boot loader install & boot support on > 2TB disk partitions
- Virtualization scalability accommodates running older releases on newer hardware. ie, RHEL 4 guests on RHEL 6 host.





Kernel scalability limits - x86-64

Parameter	RHEL5 Support Limit	RHEL6 Support Limit	RHEL6 Theoretical Limit
CPUs	64 (192 – platform dependent)	4096	4096
Memory – Physical addressing	1T	8 T (pending testing)	64TB
Memory – process virtual address space (note – hardware dependent both RHEL5&6)	128T user 64T kernel	128T user 128T kernel	128TB
IRQs	239	33024	33024
# of processes	32000	32000 (larger pending testing)	4 million
KVM guest memory	512	Same as bare metal	Same as bare metal
KVM guest cpus	32	64 (pending testing)	64

PRESENTED BY RED HAT



Determinism & realtime enhancements

- Some capabilities from Red Hat in MRG-realtime kernel (currently shipping as a RHEL 5 layered product) mainstreamed in RHEL 6
 - **Determinism** Ability to schedule priority tasks predictability and consistently
 - Priority Ensure highest priority applications are not blocked by low priority High Resolution Timers
 - **Timer** Microsecond precision not timer interrupt ~millisecond precision
- **CFS scheduler** (completely fair scheduler)
 - Provides fair interactive response times by equally sharing available cycles rather than fixed quantum of timeslice
 - Includes modular scheduler framework realtime task scheduler first
 - Priority inheritance algorithm prevents low priority processes from blocking higher priority by temporarily boosting priority to allow completion
- There will be a separate MRG-realtime offering for RHEL 6
 - Includes threaded interrupts and features not yet incorporated upstream
 - Allows rapid kernel innovation in supported product offering





RHEL 6 virtualization

Advancements to make virtualization ubiquitous:

Easier to deploy and manage

- Better control of resource allocation
- Migration among non-identical hardware

Performance close to bare metal

- Allowing all classes of workloads to benefit from virtualization flexibility

 allowing a "run anywhere" deployment strategy.
- Scalability I/O, memory, CPU
- More direct device access by guests, avoiding hypervisor overhead
- Heterogeneous Includes focus on Windows guests as well as Linux

• Secure

- Further guest isolation when cohabitating
- Compatibility of RHEL ecosystem
 - Consistent application environment Obviate need for applications to be aware of virt vs bare-metal, vs cloud deployment
- KVM's tight kernel integration avails the majority of kernel features to virtualized guests. Examples: cgroups, CFS scheduler, timer precision.
 SUMT Additionally allows paravirtualization of clock, interrupt controllers, etc.



PRESENTED BY RED HAT

Virtualization – virtual memory enhancements

Transparent hugepages

- Hugepages is a mechanism to efficiently manage large memory allocations (ie 2MB) as a unit rather than as small 4K chunks. Often 4X more efficient memory handling.
- Historically, hugepages suffered usability challenges as it required system startup time pre-allocation and is not swappable.
- Transparent hugepages obviates need to manually reserve memory. Dynamically allocates hugepage VM mappings for large allocation requests. Provides migration flexibility.
- Flexible policy controls, can enable per guest or process group.
- Beneficial to applications requesting large memory chunks.
- Highly beneficial to KVM hypervisor to more efficiently manage and allocate guest memory.

• Extended Page Table (EPT) age bits

- Enhancements in paging/swapping algorithm, nested page table support
- Allows host to make smarter choices in swapping when under memory pressure
- Allows swapping of transparently allocated hugepages by breaking up into smaller pages.



PRESENTED BY RED HAT

SUMIT

JBoss

WORLD

Virtualization – performance enhancements

KSM – Kernel Shared Memory Swapping

- Identifies duplicate pages, consolidating duplicates. Major example use case is Windows zero'ing all of memory at startup.
- Previously shared pages were not swapable
- New in RHEL6 is ability to swap KSM shared pages, beneficial to alleviate memory pressure in overcommit situations.

User return notifiers

 Allows register caching and avoids needlessly preserving register states during context switching (expensive operations) when optional components like floating point are not currently utilized.





Virtualization – scalability enhancements

- **SMP** kernel synchronization enhancements (more fine grained)
 - Benefit: Scalable to 64 cpus per guest (vs 16 on RHEL5)
 - RCU kernel locking utilizing a lock-free mechanism
 - Guest spin lock detector causes guest to yield if spinning on same instruction too long (another guest not running may hold lock).
 - Intel & AMD hardware primitives "PLE exit" optimize
- Guest hotplug CPU, disk & net
 - Allows virtual CPUs to be added/removed to running guests
 - Can also add/remove disk and network devices
 - Memory hotplug not currently supported
- x2apic
 - A virtual interrupt controller allowing direct guest access, obviating need for KVM emulation overhead.





Virtualization – network optimizations

- Vhost-net a ring buffer abstraction between guest/host
- (Performance) Much of the network implementation moved into kernel (from Qemu user space) for optimization. Fewer context switches and vmexits. Increases multithreading.
- Raw socket mode for SRIOV
 - Previously networking interrupts handled through software bridging in "tap mode"
 - Bypasses bridge allowing logical NICs assigned to guests direct PCI passthrough access. While optimized this ties guest to specific hardware and limits migration flexibility.
 - (Migration Flexibility) The vhost-net abstraction makes SRIOV allocation transparent to guest, allowing migration, even among non-identical systems.
- Network boot using gpxe providing a more modern environment for pxe network booting





Virtualization – storage enhancements

- AIO asynchronous IO allows initiating large number of IO operations (ie database workloads)
 - RHEL5 provides AIO emulation Qemu spawns individual threads per IO operation
 - Utilizing native AIO infrastructure yields 20% improvement in many IO intensive workloads
- External ring buffers used in host/guest interfaces
 - Allows more concurrent IO operations to be in progress, not limited by finite buffer descriptors
 - Doesn't consume extra buffer space when not needed
- **Block alignment** storage topology awareness
 - Interrogates underlying storage hardware and pass through optimal alignment and physical sector size to guests. Ie, in support of 4K sectors. Requires storage device commands providing the info.
 - Allows optimal filesystem layout and application aware IO optimizations.





Virtualization – migration enhancements

• Static PCI slots

- Minor differences in device ordering preclude migration, especially PCI slot numbering among different versions of host/guest (ie RHEL4.7 on 5.5)
- This capability enables logical assignment of PCI slots, preserving across migration ensuring consistency of device namespace.

CPU capability enumeration

- Providing accurate physical CPU type to applications & libraries allows usage of optimization instructions, example SSE4 in recent instruction set.
- Allows optimization of application performance with dyanamic adaptation to match capabilities among migration domains.

Vhost over SRIOV

- Logically separate physical / virtual device assignment. Guest sees virtual device.
- SRIOV optimizations previously were hardwired to specific units, precluding migration in many cases.
- Host dynamically binds SRIOV resources to guests, allowing migration among nonidentical systems.





Virtualization – svirt – SELinux virtual guest containment





PRESENTED BY RED HAT

Virtualization – Xen interaction

- RHEL 5 includes both Xen & KVM hypervisors
 - RHEL 5 can accommodate RHEL 6 guests on either hypervisor
 - Xen accommodates PV (paravirt) & FV (fully virtualized) Linux guests
- RHEL 6 includes KVM hypervisor
- Migration tool provided to convert RHEL 5 Xen guests to KVM format to run on RHEL 6





Storage management

- Topology awareness I/O (alignment and chunk size) based on info from the storage device. This is in dm, LVM, md, and utilities such as parted and mkfs standardized interfaces to obtain alignment and optimal I/O stripe width.
- DIF/DIX scsi data integrity commands (checksum) superior integrity
 - End-to-end data integrity check (SCSI DIF/DIX). Initially this extra checksum will be from the HBA to the storage. Added to applications and filesystems in the future. (requires storage hardware providing this capability)
 - Initially targeted at database use case on raw partition & HBA to storage in filesystem
- **FCoE** (fibre channel over ethernet) on specialized adapters (Emulex, QLogic, Cisco), and on standard NICs. FCoE install & boot support with DCB.
- **iSCSI** root/boot, including target
- Thin provisioning (virtual storage overcommit) via "discard" command in LVM & filesystem. Requires storage device capability. Improves SSD wear leveling.
- Block Discard. Optimizes thin provisioning in the storage device, and improves wear leveling of SSDs. Currently used by XFS and ext4. Currently not usable with LVM/DM/multipath or md
- **SRIOV, NPIV** driver virtualization IO accelerators guest direct access
- VSAN virt SAN fabric based on NPIV, each guest has a separate WID,
 Summer-guess control



PRESENTED BY RED HAT

Storage management

- LVM/DM (Logical Volume Manager / Device Mapper)
 - LVM hot spare, a disk or group of disks used to replace a failing disk
 - Online resize of mirrored & multipath volumes
 - Snapshot scalability enhancements for virtualization
 - Multipath enhancements
 - Dynamic multipath load balancing. Path selection based on queue depth, or I/O service time
 - Mirroring enhancements
 - Snapshot of snapshot mirror
 - Mirrored mirror log. Improves mirror log availability, to avoid the need for a resynch after a failure
 - Cluster mirror
 - Snapshot merge. Provides the ability to "rollback" changes that were made since the snapshot was taken. (Additional work to integrate this with Anaconda and yum will be post 6.0).
 - dm-crypt enhancements. Selectable hash algorithm for LUKS header, new cryptsetup commands, new libcryptsetup with versioned API.
 - Remote Replication tech. preview.





Filesystem – larger & faster

- Ext4 will be the default file system and scale to 16TB
- XFS optional offering to support extremely large file systems > 16TB, up to 100TB. Tuned for larger servers & high end arrays.
- NFS
 - NFS4.0 clients will default to use NFS4.0 (tunable via mount or config file)
 - NFS4.1 enhanced support for referrals, delegation & failover.
 - Support for enhanced encryption types for kerbrerized NFS
 - Added IPv6 support
- GFS2 optional
 - Targeting high availability clusters of 2-16 nodes.
 - Clustered samba (CTDB) parallel (concurrent) servers for scalability & availability
- Filesystem utilities enhancements:
 - Creation tools warn about unaligned partitions, and new partitions are created on aligned boundaries with preferred block sizes. (hardware dependent)

ced Jodge barriers for increased data reliability – for ext3, ext4, GFS2, xfs



PRESENTED BY RED HAT

SUMI

Networking improvements

- **10GbE** Driver support on card switch and 8-16 pci devices.
 - Virtual guest can access the full NIC directly SR-IOV enhancement – ie a single virt guest can saturate a 10GbE link.
 - Data Center Bridging (DCB) support in ixgbe driver
 - Uses 802.1p VLAN priority tags to schedule and control traffic rates
 - Uses 802.1Qaz (priority grouping) and 802.1Qbb (priority flow control) to physically separate traffic flows that coexists on the same physical link
- **FCoE** (Fibre Channel over Ethernet)
 - Working on performance improvements throughout the storage stack (locking changes in the block and SCSI layers, improved interrupt handling).
- **RDMA** support over 10GbE & Infiniband
 - Add IPv6 support





PRESENTED BY RED HAT

Networking improvements

- Major new features post RHEL5
 - Ipv6 Mobility support, RFC 3775
 - UDP-lite support, benefits multimedia protocols such as Voice Over IP, RFC 3828
 - Add Mutiqueue hardware support API
 - Large Receive Offload in network devices
 - Network controller for cgroup
 - Add multi-queue, DDR scheduler
 - Add TCP Illinois and YeAH-TCP congestion control algorithms
- General Networking Stack Performance improvement
 - RCU (read copy update) SMP locking optimization adoption in networking stack
 - Use RCU for the UDP hash lock
 - Convert TCP & DCCP has tables to use RCU.
 - RCU handling for Unicast packets
 - Multi-CPU rx to pull in from the wire faster
 - Multi-queue xmit networking for multiple transmit queues devices
 - New monitor tools for dropped packets, tc and dropwatch



JBoss



System services enhancements

- Dracut replacement for initramfs, mkinitrd
 - Better long-term supportability of storage configurations
 - Can automatically add raid members via udev rules
 - Allows change in hardware setup without needing to recreate initramfs
- NetworkManager iSCSI, FCoE config, IPv6, Bridging
- Upstart flexible system service startup infrastructure
- CUPS printing enhancements
 - SNMP-based monitoring of ink/toner/supply levels and printer status
 - Device discovery speed improvements (backends now run in parallel)
 - Automatic PPD configuration for PostScript printers (PPD options values queried from printer) -- available in CUPS web interface
- Portreserve
 - Avoids network port allocation failures for network services





Compiler / Tools

OpenJDK

- TCK certified based on IcedTea project
- Fully open source implementation of Java Web Browser plugin & Java web start avoids need for proprietary plugins.

• GCC 4.4

- OpenMP3 conformance for portable parallel programs
- IRA (Integrated Register Allocator)
- Tuples to enable future features and improve memory footprint
- Additional C++0x conformance implementations
- Improved automatic parallel mode in C++ library
- Lock free C++ class libraries
- Debuginfo handling improvements





Compiler / Tools

- Glibc C runtime library
 - Malloc optimizations Improved speed and efficiency for large blocks & NUMA considerations
 - Lock free C++ class libraries, ie ring buffer, fifo MRG dependency
 - NSS crypto consolidation for LSB 4.0 & FIPS level 2

GDB debugger

- Focus on C/C++ debugging on native Linux
- C++ function, class, templates, variables, constructor / destructor improvements
- Catch / throw and exception improvements
- Large program debugging optimizations
- Pretty printing of C++ values
- Non-blocking thread debugging Threads can be stopped or continued independently
- Asynchronous interaction respond to commands while program is running
- Limited multi-process debugging
- System call tracing
- Improved name lookups
- Constructor/destructor manipulation

WORLD

Additional python scripting capabilities

PRESENTED BY RED HAT

SUMI



Compiler / Tools (continued)

• Binutils – Intel XSAVE, AVX (Advanced Vector Extensions)

Systemtap

- Mangling-aware support for probing c++ applications
- Java probing (expecting to use DTrace markers)
- Java backtracing
- Remote probing (staprun service)
- Pre-configured Kernel tracepoints
- User space static markers/probes
- Eclipse CDT (x86, x86-64 only)
 - Java, C, C++ editing, debugging, and compilation support
 - C and C++ memory and CPU profiling with Valgrind and OProfile
- ElfUtils dwarf compression





Investment Guarantee Initiatives

Kernel ABI

- Some kABI symbols are maintained for the life of the release the " kABI whitelist"
- Kernel modules restricted to these symbols need no recompilation during release life
- Tools available to validate kernel module against the list
- yum kABI plugin warns and optionally prevents installation of non-kABI conforming modules (those that use non-whitelisted symbols)

• Driver Update program improvements

- Hardware support for new devices, critical fixes
- Now simple RPMs on a disk, with installation support

Migration Services

• Available through Red Hat Support

JBoss

- Designed to help sysadmins migrate older systems to RHEL6
- Tools to run on RHEL5 system, generates report of necessary migration tasks



PRESENTED BY RED HAT

SUM

Manageability, Interoperability, Security

- **OpenChange** client-side library
 - Implements Microsoft Exchange MAPI protocols
 - Evolution plugin allows native access to Microsoft Exchange

• Samba 3.5

- New graphical user interface for joining Windows Domains
- Windows 2008 (R2) trust relationships, Windows 7 domain members
- Encrypted SMB transport between Samba client and server
- Full support for Windows cross-forest, transitive trusts and one-way domain trusts
- Active Directory signing/sealing policy
- IPv6 support

SHA256 support

- verify integrity of software throughout the build system, rpm, yum, createrepo, RHN, and satellite
- Updated management clients for enterprise infrastructures
 - Openwsman





Security features

- System Security Services Daemon (SSSD) for identification and authentication
 - A more robust and better performing LDAP and kerberos client
 - Caches network identities for offline authentication and NSS lookups for client speedups and reduction of server load
 - Overcomes deficiencies of current implementation of NSS LDAP (connection pooling)
 - Multiple identity domains & providers allowing clients to authenticate correctly in mixed identity environments
 - Provides proxy provider to load existing NSS modules
 - Integrated with authconfig
- Common nss_ldap daemon
- NSS Softoken Centralized key and certificate store
 - FIPS 140-2 compliant, new Algorithm support (SEED, CAMILLA)
 - PKIX certificate verification, support for CRL distribution points
- OpenSWAN RFC 5114 and Cisco VPN support
- Minimal Install to lower the attack surface by installing only necessary packages
- Key Escrow for LUKS and Anaconda

WORLD

to says disk encryption keys of LUKS volumes, during installation or later



PRESENTED BY RED HAT

SUMI

Security – SELinux

- sVirt
 - Builds upon existing process-based security mechanisms to prevent unauthorized access of guests/host in a virtualized environment. Central to sVirt design is the integration of Mandatory Access Control (MAC) with virtualization. Guests are isolated using a MAC security policy which helps contain hypervisor breaches. The intent of this program is to limit the guests to specific storage devices, network ports, etc. with little or no configuration.

Confined Users

Allow SELinux to control what a particular user can access on the system by assigning the user an SELinux role, including 'Guest', 'Xguest', 'User', 'Staff' and 'Unconfined' (default). Allows policy writer to confine administrator activities. Confined administrator users can transition to other role when executing root commands. We ship the webadm which allows UID=0 processes to only edit files that are labeled as apache content.

• XACE

SUMI

 XACE (X Access Control Extension) allows X extensions to perform access checks. The concept is identical to the Linux Security Module (LSM) in the Linux Kernel. XACE is currently used in MLS environments to prevent data transfer between windows running at different levels.

• Sandbox / Kiosk

• The SELinux sandbox enables administrators to run any application within a tightly confined SELinux domain. Using the sandbox, administrators can test the processing of untrusted content without damaging the system. Kiosk mode suitable for deployment in public locations, schools, libraries.



WORLD



High availability clustering

- Modular split of cluster functions per upstream development
 - Corosync cluster engine, config database, IPC
 - Openais application management framework, cluster checkpoint service
- Common cluster logging system utilized by all cluster infrastructure components
- Compatibility mode provided to allow staged migration from RHEL5
- KVM fully supported as hypervisor for running highly available virtual machines
 - fence_virt/virtd replaces fence_xvm/xvmd as a pluggable virtualization capable fencing agent
- Running clusters inside of KVM guests is provided as tech preview
 - Using either RHEL5 or 6 as host as well as guest

JBoss

ESENTED BY RED HAT

SUM

- Cluster infrastructure supports IPv6 and SELinux in enforcing/targeted mode
- CTDB active/active samba failover offering layered on GFS2 (Technical Preview)
- Pacemaker introduced and integrated with the cluster stack (Technical Preview)



Desktop features

- Plymouth = fully graphical boot sequence infrastructure
 - Cleaner, less "noisy" boot sequence (requires kernel mode setting)
 - Integrated passwords for encrypted volumes
 - Driver work ongoing ie Intel, ATI, Nvidia (via new Nouveau driver)
- Kernel mode setting (KMS) for graphics drivers (Intel & ATI)
 - Enables fast user switching without re-initializing the display, similar for transitions from text mode consoles to graphic
 - Enables kernel panic messages to be displayed
- PulseAudio
 - Glitch free audio playback rewritten to be timer based rather than interrupt based -> reduced power consumption, minimization of drop-outs and flexible adjustment of the latency
- GNOME 2.28
 - Empathy instant messaging client replaces Pidgin
 - Ekiga 3.0 video conferencing applet
 - Enhanced file browser (tabbed & compact views)
 - Enhanced multihead configuration



X graphics enhancements

- Kernel mode setting (KMS) for graphics drivers (Intel & ATI)
 - Enables fast user switching without re-initializing the display
 - Also fast transitions from text mode consoles to graphic
 - nables kernel panic messages to be displayed
- DRI2 –accelerated off-screen rendering, video support, OpenGL 3D
- Hardware auto detection
- Enhanced external monitor support for laptops





SUMMIT JBoss WORLD

PRESENTED BY RED HAT

LEARN. NETWORK. EXPERIENCE OPEN SOURCE.

www.theredhatsummit.com