



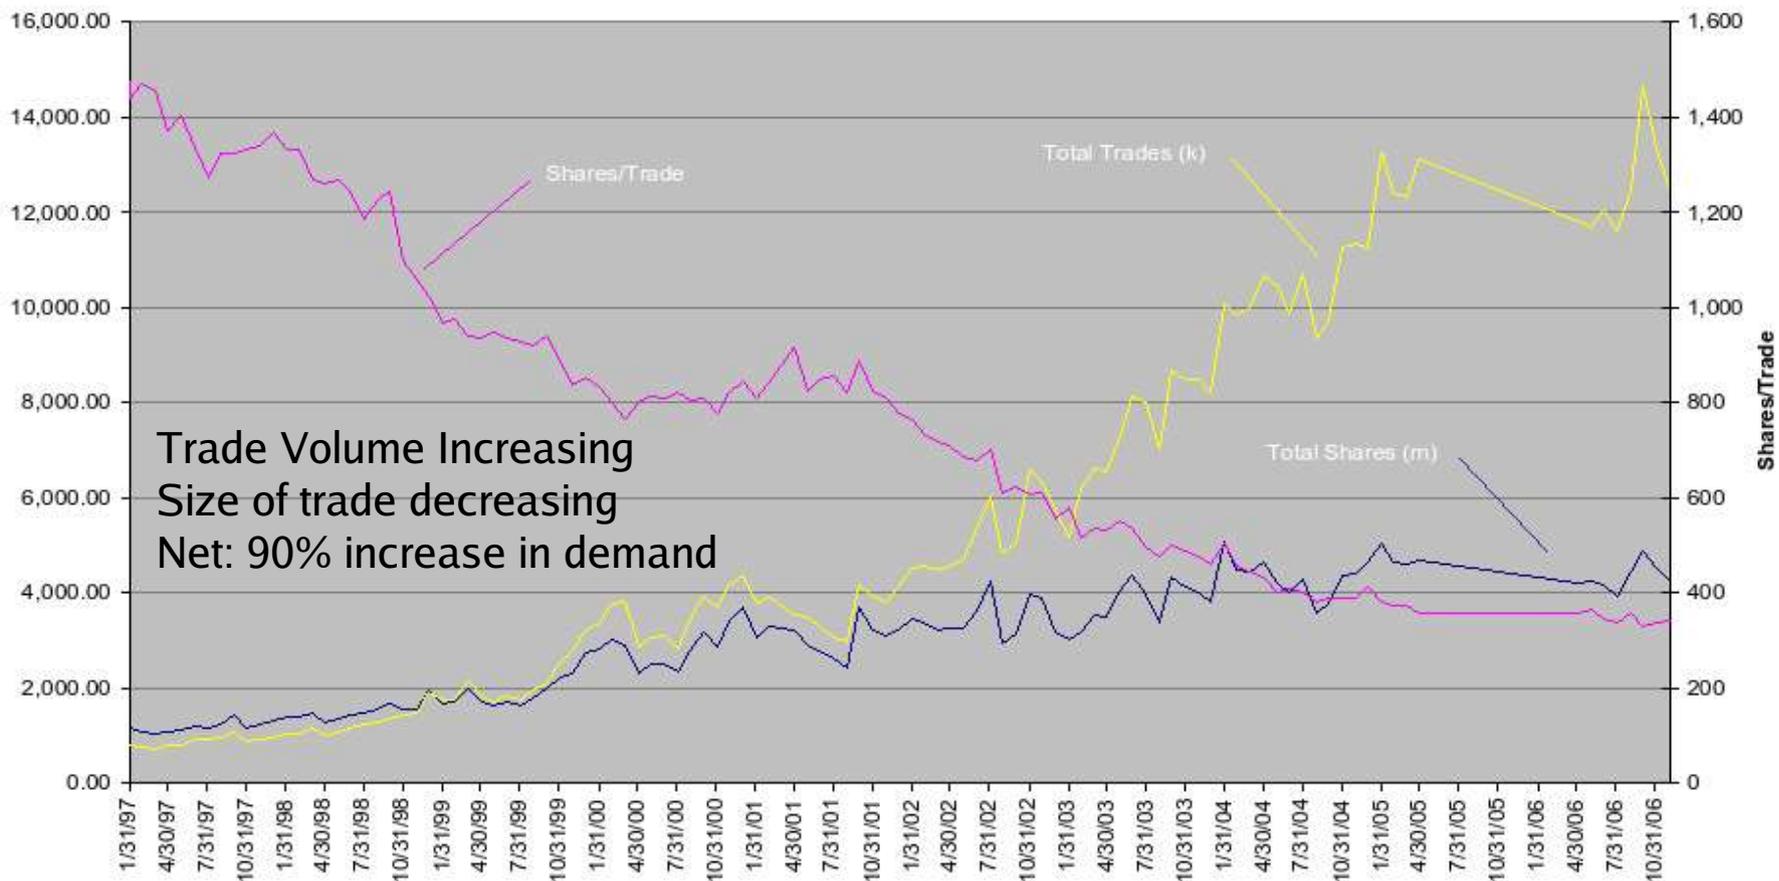
Red Hat Enterprise Linux 5 Technology Overview

Rick Ring
Senior Solutions Architect, Federal

Increasing IT Demands

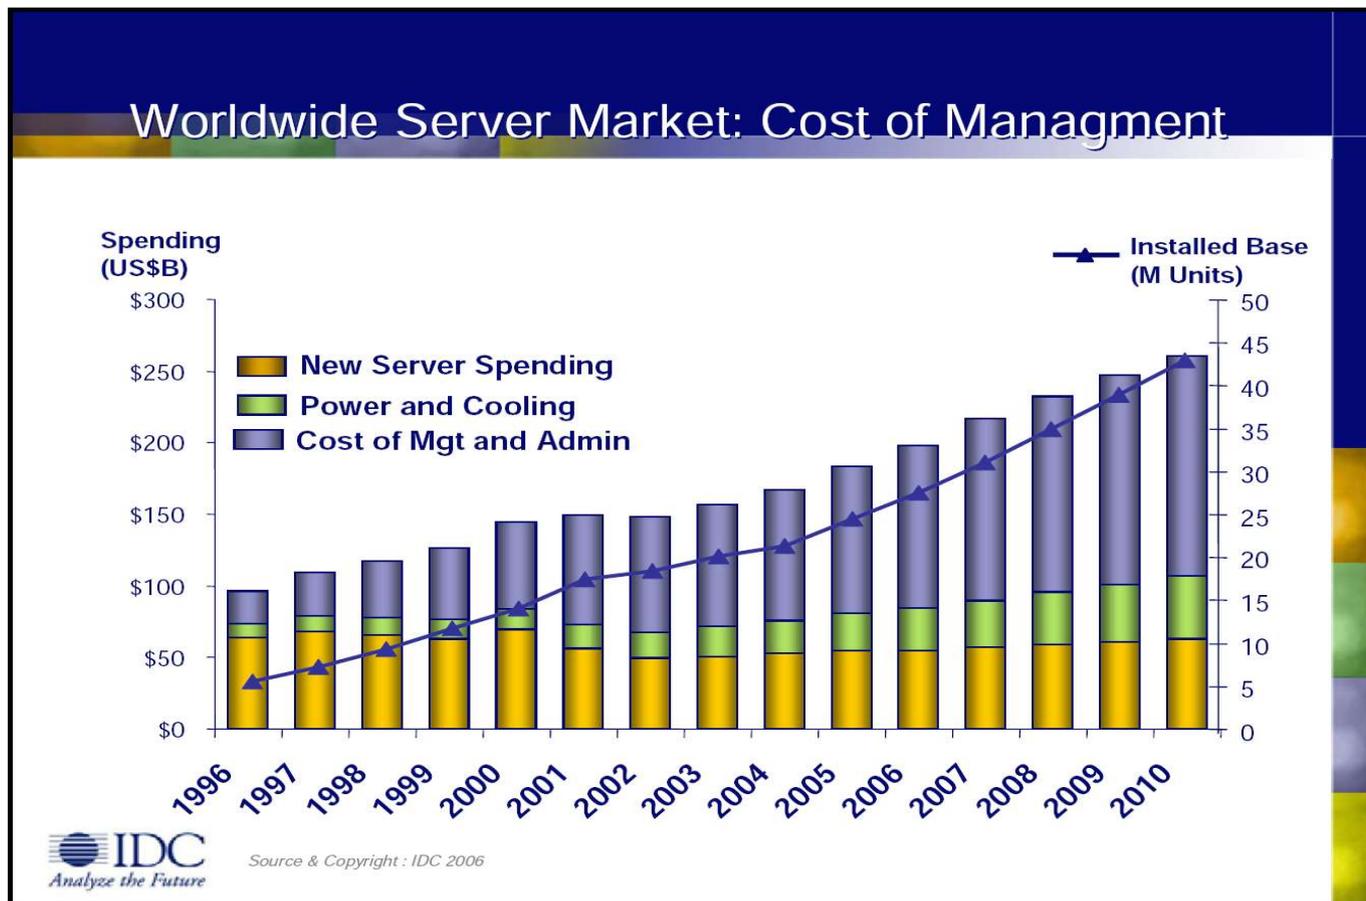
- Example: Chart shows transaction growth in Capital Markets
- Similar dynamic in other sectors

Traffic Multiplying Factor: Volume Increasing as Size Decreases



Opex Now Outweighs Capex

- Open source solutions reduce both cost components:
 - Reduced acquisition cost (h/w, s/w, deployment)
 - Reduced operating costs (e.g. server/admin ratio, security outages, etc.)



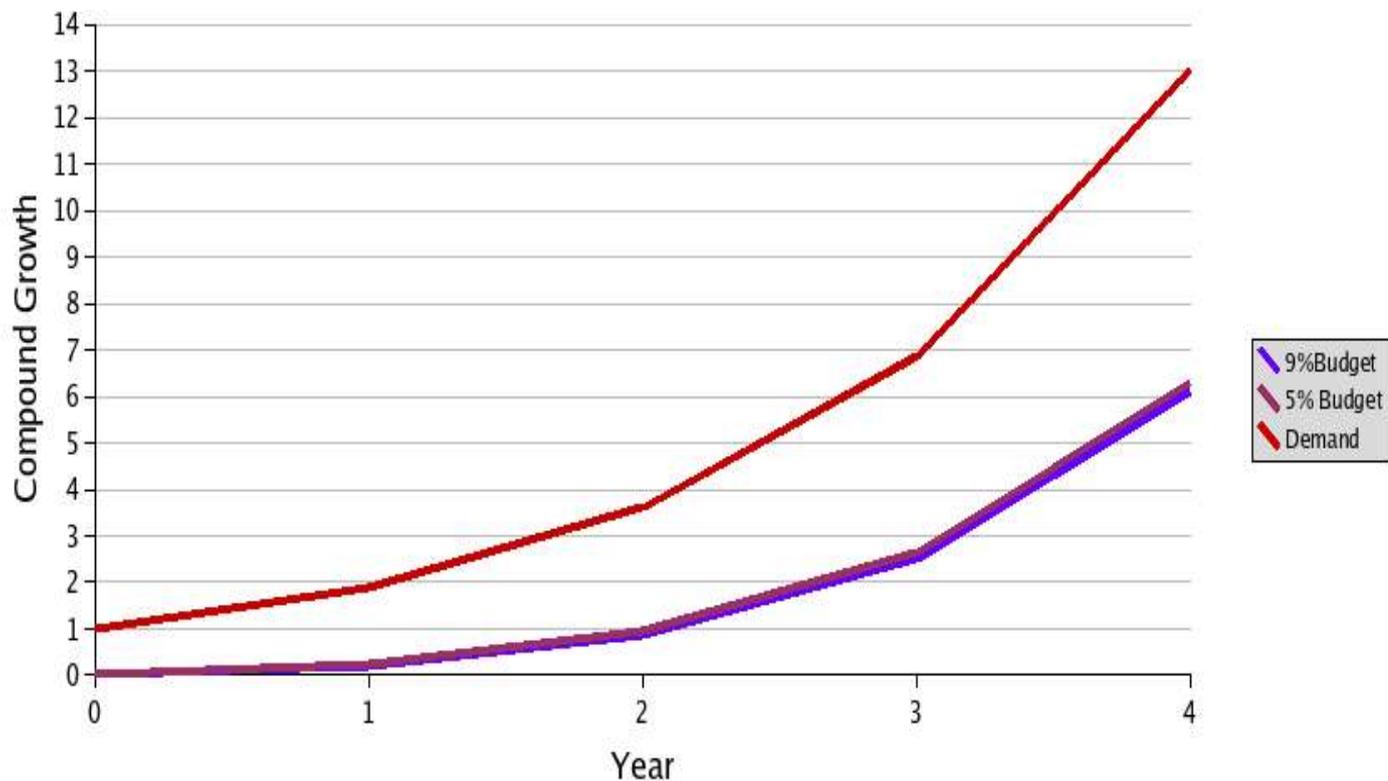
The GAP

■ Assumptions:

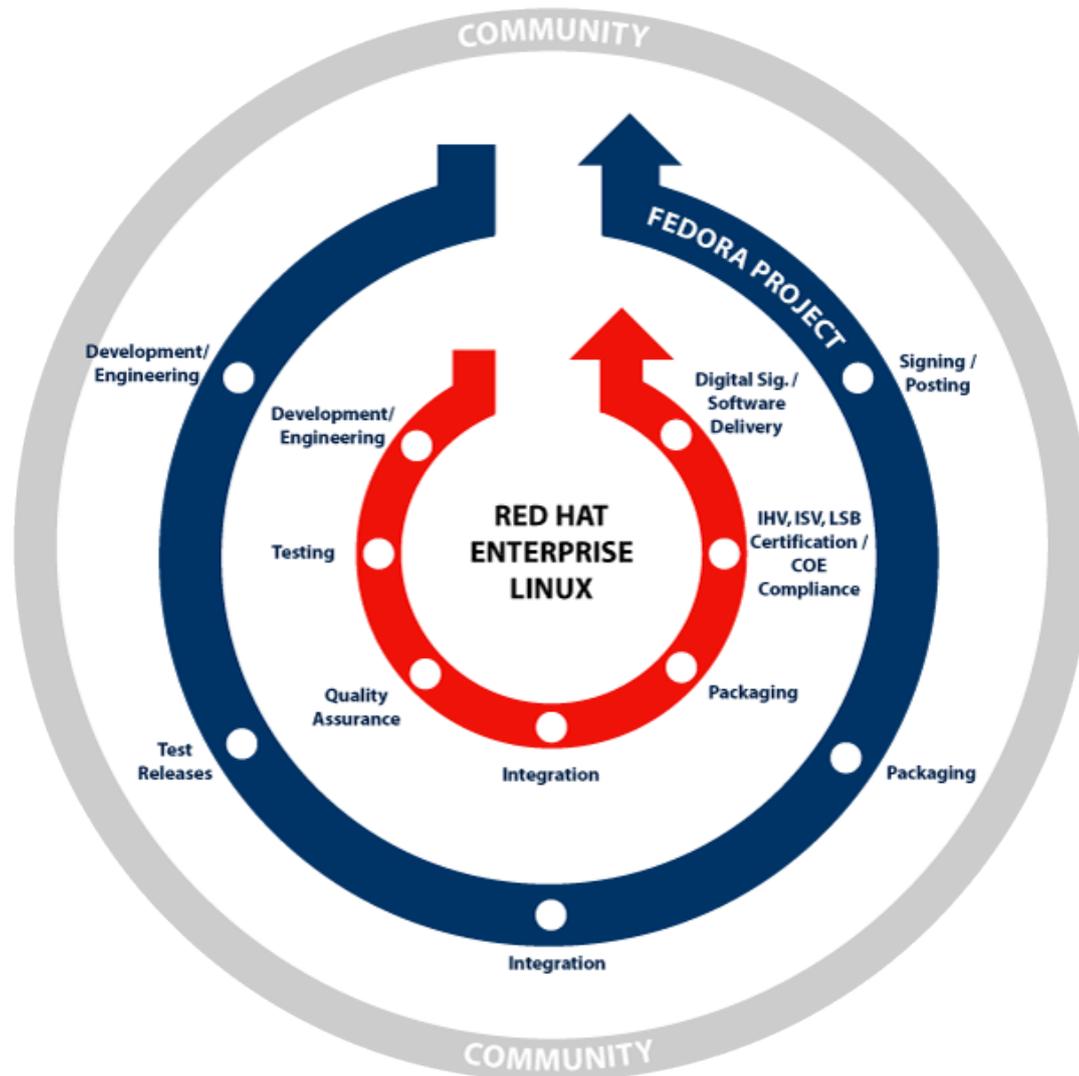
- Demand growing at 90%
- Moore's law increasing capacity by 60%
- NET: Budgets cannot fill the gap

■ Result:

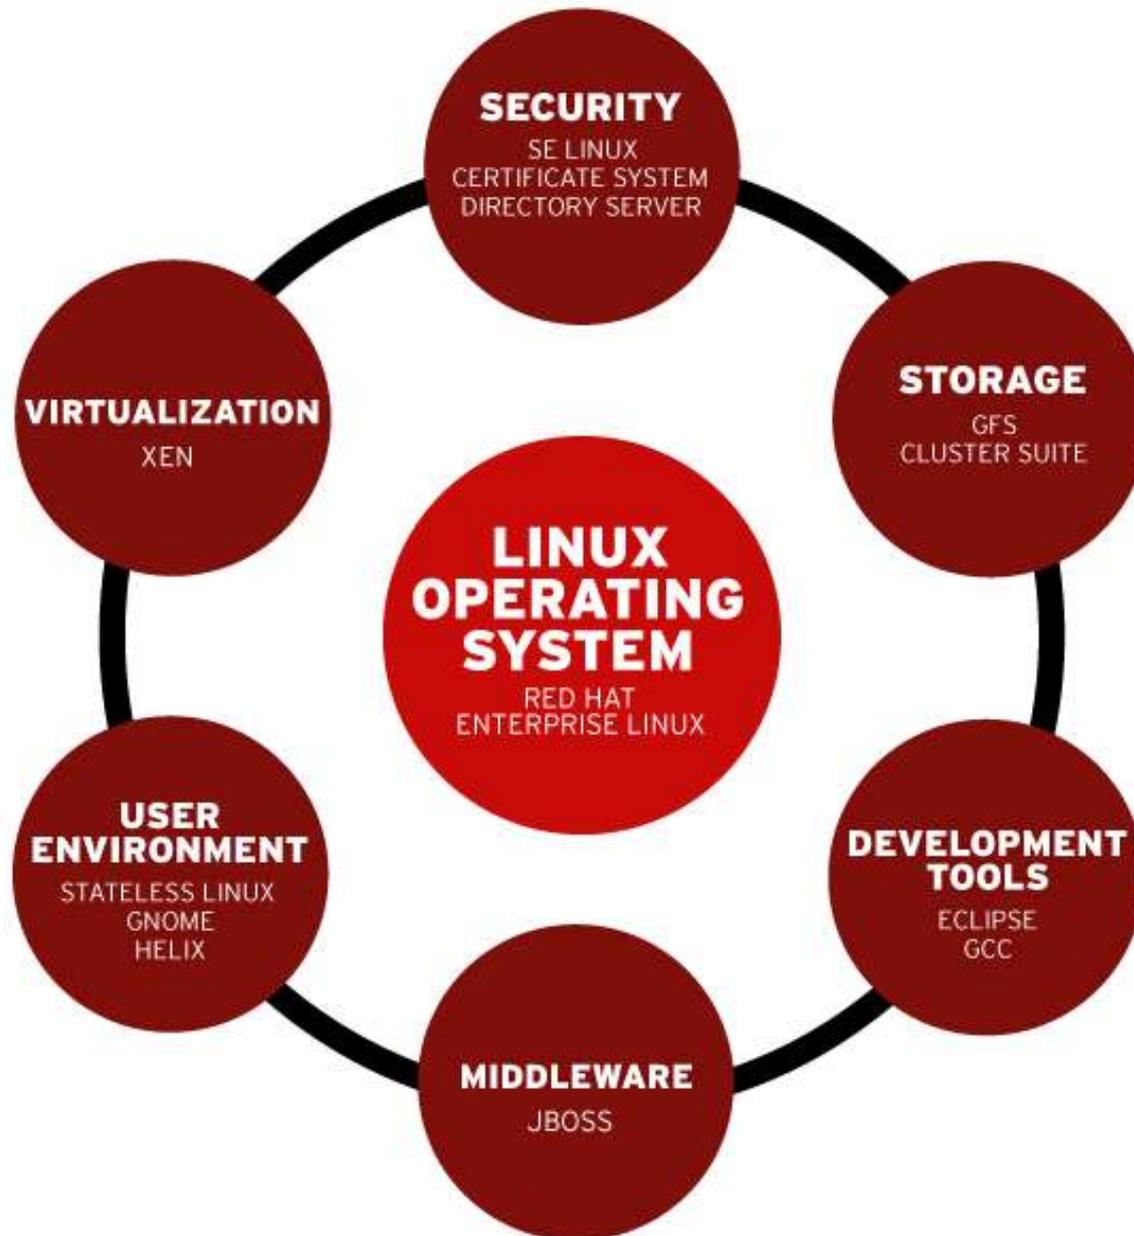
- Need a new approach
- Open source cost model bridges the gap



Red Hat development model

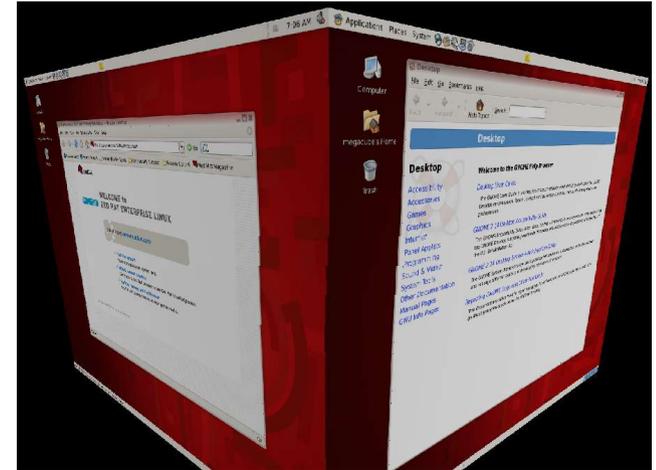


Enterprise open source opportunity today



What's new in Version 5?

- Enterprise Linux Advanced Platform
- Integrated virtualization
- Industry leading performance and price/performance
- Enhanced and easy to use security
- Improved networking and interoperability
- Enhanced development tools
- New SLAs
- Enhanced client



RED HAT
ENTERPRISE LINUX ADVANCED PLATFORM

RED HAT
ENTERPRISE LINUX DESKTOP

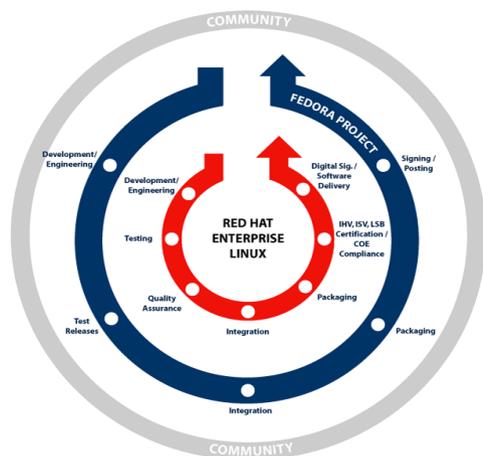
Red Hat Enterprise Linux 5: Feature Summary

- Red Hat Enterprise Linux 5 comprises more than 1200 components
- Over two years of development
- Technology created by:
 - Red Hat
 - Partners
 - Community

Virtualization
Server virtualization is provided in the base server product & is available for the client product
Storage and extended server virtualization is provided with the Virtualization Platform option
Red Hat Network support for virtualization
Virt-Manager, libvirt/virsh management tools
Packaging
Replacement of previous AS, ES and WS products with a single server and a single client
New Options provide additional server and client product capabilities
Kernel & Performance
Red Hat Enterprise Linux is based on the Linux 2.6.18 kernel
Support for multi-core processors
Broad range of new hardware support
Updated crash dump capability provided by Kexec/Kdump
Support for Intel network accelerator technology (IOAT)
Numerous enhancements for large SMP systems
Enhanced pipe buffering
IPv4/IPv6 fragmentation offload & buffer management
Dynamically switchable per-queue I/O schedulers
Kernel buffer splice capability for improved I/O buffer operations
Security
SELinux enhancements include Multi-Level Security and targeted policies for all services
SEtroubleshooter GUI simplifies SELinux management
Integrated directory & security capabilities
IPSEC enhancements improve security and performance
ExecShield enhancements, such as a call frame Canary word, strengthen hacker defenses
New Audit features provide powerful new search/reporting and realtime monitoring

Red Hat Enterprise Linux 5: Feature Summary

- Features exposed to extensive testing with Fedora Core 4/5/6
 - Ensures high quality



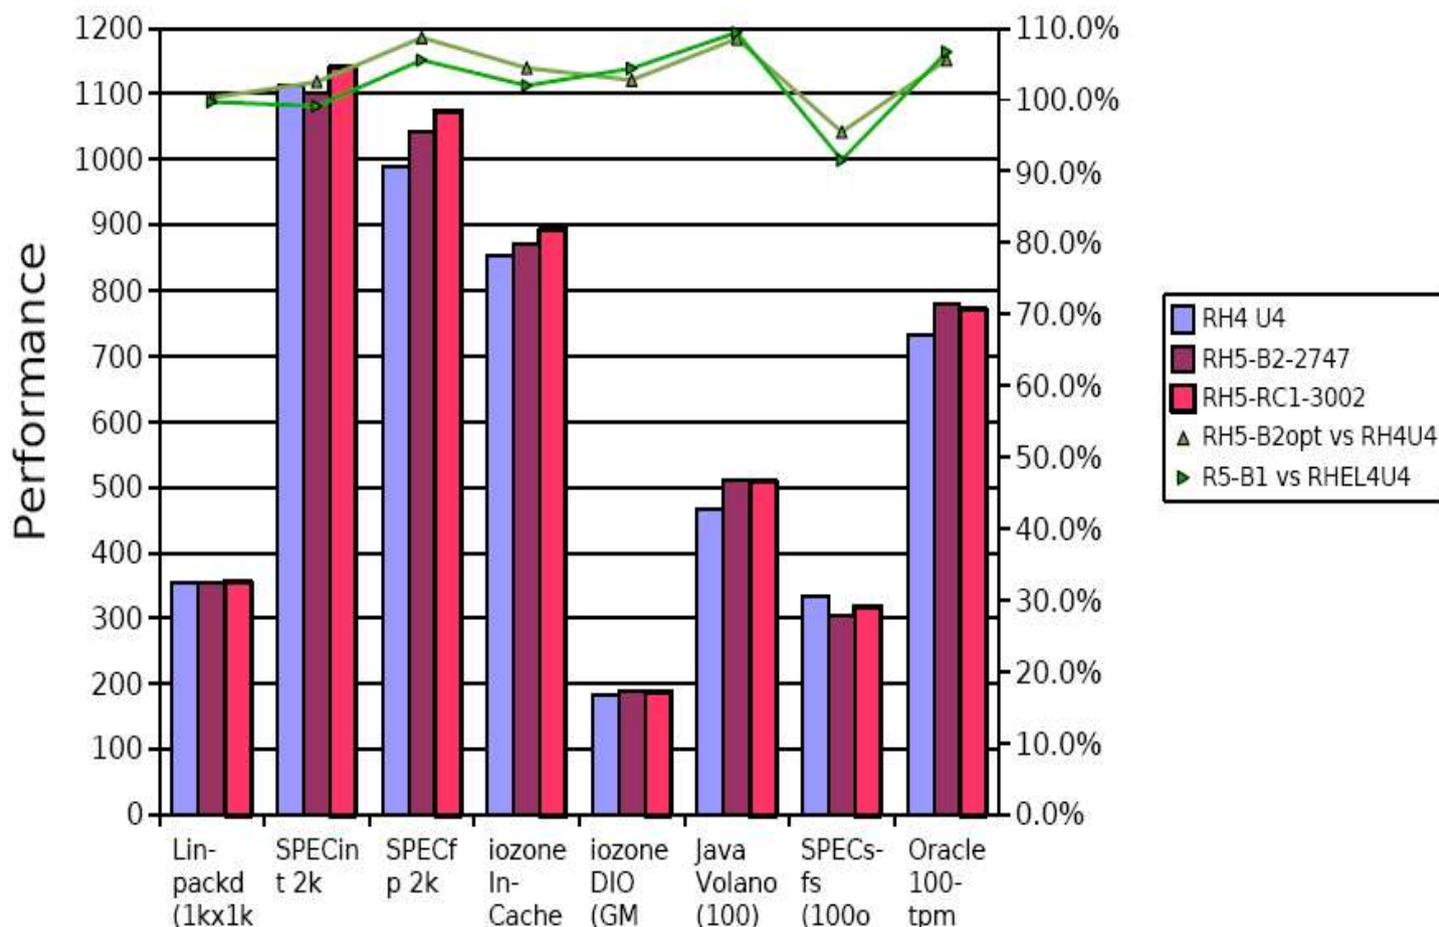
- Application interfaces held stable for life of product

Networking & Interoperability
Network storage enhancements include Autofs, FS-Cache and iSCSI support
IPv6 support and conformance enhancements
Improved Microsoft file/print and Active Directory integration
Desktop
Desktop enhancements provide updated configuration tools, applications and laptop support
Foundational Stateless Linux features (X autoconfigure, NetworkManager, etc)
Improved ACPI and laptop support
Smart card login - with PKI/Kerberos authentication
Integrated multi-media support
Enhanced plug and play hardware support (cameras, printers, scanners, etc)
Network Manager provides automatic wired and wireless network configuration
Enhanced graphics using AIGLX/Compiz (with fading, transparency, etc)
Development Environment
Enhanced application development tools including SystemTap profiler and Frysk debugger
GCC 4.1 and glibc 2.4 toolchain
Storage
Support for root device multipath IO (MPIO) improves availability
Single system/guest version of Red Hat Global File System included in the base product
Block device data encryption support
Management
Numerous installer improvements make system configuration simpler
Yum/Pup-based updaters for Red Hat Network
Conga cluster & storage management (with Virtualization Platform)

Early RHEL5 testing...

- Shows up to 10% performance improvement in most tests
 - Sybase shows 12% improvement
- Ongoing work expected to further increase performance
- SPECsfs enhancements underway
 - RC is within 2%
 - Lock contention in SUNrpc model

RHEL5 RC1 vs B2 vs RHEL4 U4



RED HAT ENTERPRISE LINUX ADVANCED PLATFORM

Delivers: Lower Capital Costs

- Eliminate the purchase of third party server software, saving \$3k - 300k/server, including:
 - VMWare, EMC PowerPath, Clustering / Availability software
 - Veritas filesystem, Polyserve, RAC
 - (and of course, JBoss delivers further capital cost reductions)
- Consolidate and purchase fewer servers.
 - Utilize unused server and storage capacity
 - Purchase across applications for average, not peak, demand
 - Reduce redundancy necessary to achieve availability goals
- No per-instance OS charge
- Better price/performance means fewer servers needed

RED HAT ENTERPRISE LINUX ADVANCED PLATFORM

Delivers: Lower Operating Costs

- Lifecycle management: application encapsulation
- Administration overhead: sysadm efficiency ratio, server consolidation, reduced time to deploy servers
- Lower power consumption: dynamic resource allocation, speed scaling
- Reduced facilities cost: dynamic resource allocation, better performance
- Lower cooling cost: dynamic resource allocation, speed scaling
- Reduced overtime cost: elimination of outage windows, availability
- Faster time to solution: dynamic resource allocation, virtual sandbox

RED HAT ENTERPRISE LINUX ADVANCED PLATFORM

Delivers: Better Service Delivery

- Increased availability: redundancy, failover, migration, elimination of planned downtime
- Better response times: dynamic infrastructure, load balancing
- Faster application and service delivery: eliminate purchasing cycle, dynamic provisioning, live migration
- ... and much more with JBoss



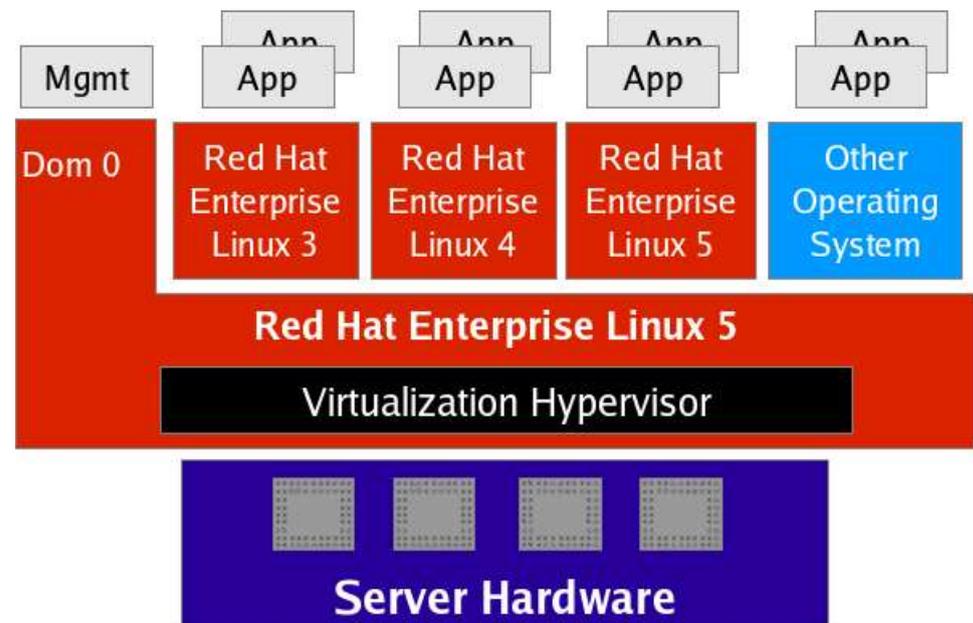
Virtualization

Virtualization Use Cases

- Datacenter Consolidation
 - Primary driver: increase utilization (primarily CPU, memory)
 - Thin host, workload in guests
 - Special case: security isolation (E.g. hosting providers)
 - Server-oriented
- Development and Testing
 - Driver: multiple development and testing environments, isolation from main workspace in host. Fault injection.
 - Typically main workload in host, testing in guests
 - Client and Server oriented
- Hardware Abstraction
 - Driver: new hardware with older guest OS
 - Support via virtual device drivers
 - Typically only few guests (maybe only 1)
 - Server-oriented

Full (hardware assisted) & Para-Virtualization

- Red Hat Enterprise Linux 5 will support hardware and software virtualization scenarios:
 - Fully virtualized on Intel VT & AMD-V systems
 - Allows guest to be Red Hat Enterprise Linux 3, 4 as well as other operating systems
 - Para-virtualized all systems
 - Red Hat Enterprise Linux 5.0
 - Red Hat Enterprise Linux 4.5
 - Support for x86, x86_64 at product release
 - Support for IA64 at 5.1

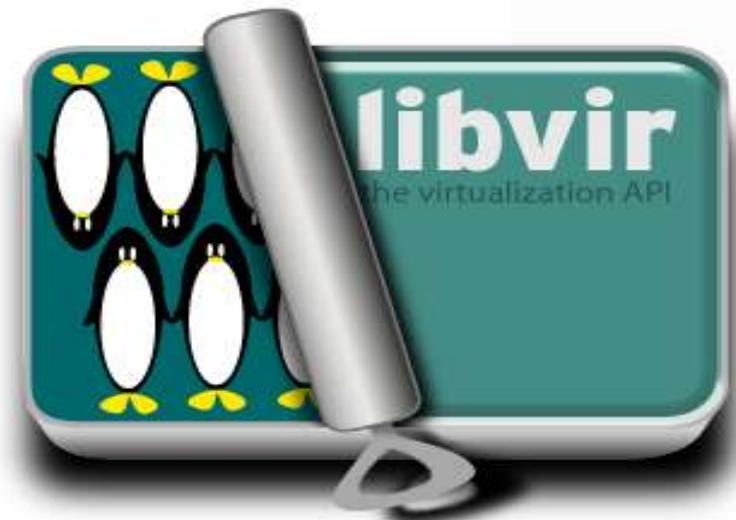


Virtualization: Management Infrastructure

- Local VM control and configuration
- Resource Monitoring & Policy Definition
- Deployment Support
 - Cobbler: command line boot/provisioning server; supports PXE
 - Koan: Kickstart-over-a-network

libvirt

- Stable API for tool/app development
 - CIM providers; Python, C bindings, scriptable
- Hypervisor agnostic (Xen, QEMU, ...)
- Local VM functionality
 - Start, stop, pause, ...
 - Support for hot and cold migration
- <http://www.libvirt.org>





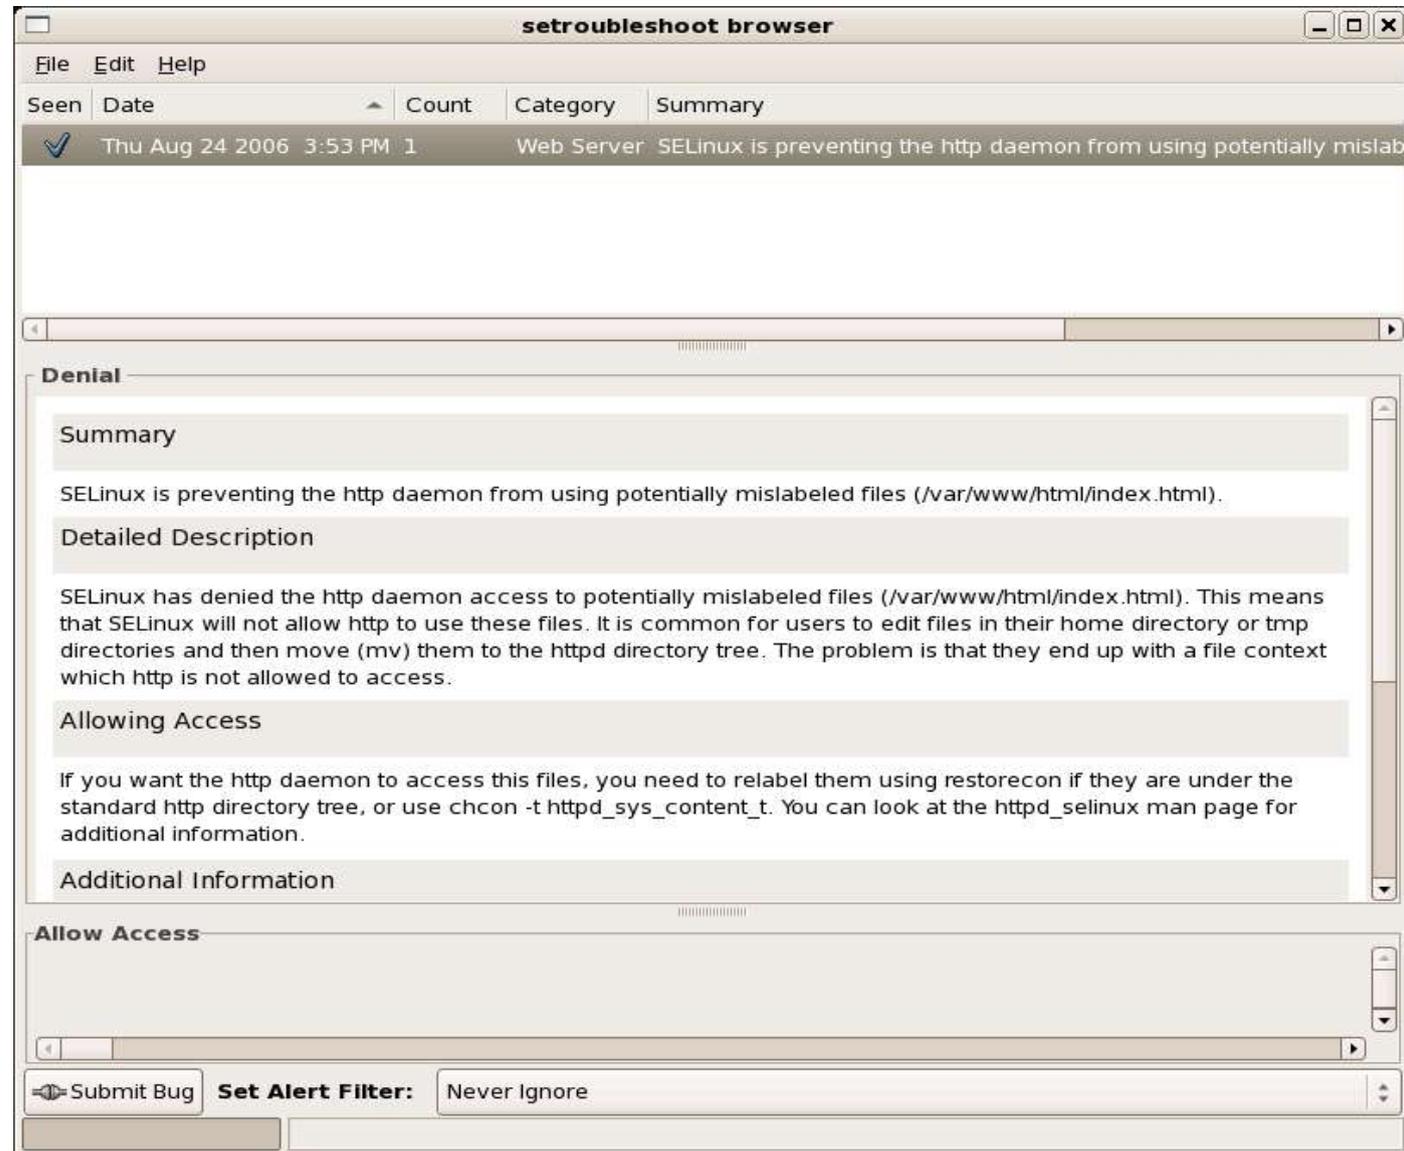
Security

Security: SELinux Enhancements

- Expanded SELinux targeted policy coverage
 - Provides coverage for all core system services, versus 11 in Red Hat Enterprise Linux 4
 - Includes support for Multi Level Security (MLS) enforcement model
 - In addition to existing RBAC and TE models
- An additional level of protection against security exploits
 - Fine-grained policies via kernel-enforced mandatory access controls
 - Limits the scope of security vulnerabilities
 - Beyond what any other general-purpose OS can deliver

Security: SELinux Enhancements

- New Setroubleshooter provides clear, easy-to-understand, GUI-based, security violation notifications
- Over 60 events defined today



The screenshot shows the 'setroubleshoot browser' window. At the top, there is a menu bar with 'File', 'Edit', and 'Help'. Below the menu bar is a table with columns: 'Seen', 'Date', 'Count', 'Category', and 'Summary'. A single row is visible with a checkmark in the 'Seen' column, the date 'Thu Aug 24 2006 3:53 PM', a count of '1', the category 'Web Server', and the summary 'SELinux is preventing the http daemon from using potentially mislabeled files'. Below the table, the main content area is titled 'Denial'. It contains several sections: 'Summary' with the text 'SELinux is preventing the http daemon from using potentially mislabeled files (/var/www/html/index.html).', 'Detailed Description' with a paragraph explaining that SELinux has denied the http daemon access to these files and that this is common for users who move files from their home directory or tmp directories to the httpd directory tree, resulting in a file context that http is not allowed to access. 'Allowing Access' with instructions to relabel files using 'restorecon' or 'chcon' and to refer to the 'httpd_selinux' man page for more information. 'Additional Information' which is currently empty. At the bottom of the window, there is a 'Submit Bug' button, a 'Set Alert Filter:' dropdown menu set to 'Never Ignore', and a search input field.

Security: SELinux Enhancements

- Greatly improved logging, with easy-to-decipher information

OLD: Red Hat Enterprise Linux 4 /var/log/messages entry

```
time->Thu Aug 24 15:50:58 2006
type=AVC_PATH msg=audit(1156449058.917:552):
path="/var/www/html/index.html"
type=SYSCALL msg=audit(1156449058.917:552): arch=40000003 syscall=196
success=no exit
t=-13 a0=8d4d4d0 a1=bf5e97c a2=434ff4 a3=2008171 items=0 ppid=23799
pid=23805 auid=3267 uid=48 gid=48 euid=48 suid=48 fsuid=48 egid=48 sgid=48 fsgid=48
tty=(none) comm="httpd" exe="/usr/sbin/httpd" subj=user_u:system_r:httpd_t:s0 key=(null)
type=AVC msg=audit(1156449058.917:552): avc: denied { getattr } for pid=23805 com
m="httpd" name="index.html" dev=dm-0 ino=6260297
scontext=user_u:system_r:httpd_t:s0
tcontext=system_u:object_r:user_home_t:s0 tclass=file
```

NEW: Red Hat Enterprise Linux 5 /var/log/messages entry

```
Aug 24 15:53:10 localhost /usr/sbin/setroubleshootd: SELinux is
preventing /usr/sbin/httpd "getattr" access to /var/www/html/index.html.
See audit.log for complete SELinux messages.
```

Security: Binary Code Protection

- Execshield enhancements provide additional armoring against most common kinds of security exploits
- Introducing stack “canary” word feature to detect overflow exploits
- Core packages built with new FORTIFY_SOURCE GCC option which implements run-time bounds checking to prevent buffer overflow exploits

Security: Identity Management

- Native support for Identity management in conjunction with Red Hat Directory Server and Red Hat Certificate System
- Integration of Identity & Certificate Management capabilities with Red Hat Enterprise Linux and community applications
 - Clear and secure architecture
 - Addition of Enterprise Security Client (smartcard, physical token, support)
 - Centralized key management for core desktop applications
 - system login, web browser, email, SSH
- Integration of certificate-based security and Kerberos infrastructure via PKInit
- Enables centralized management of users and rights
- Enables “Single Sign-On” user experience



Red Hat Enterprise Linux 5

General New Features

Storage Improvements

- NFSv4 Improvements
 - More complete implementation of the specification
 - Delegation (aka lease), increased client caching
 - Server migration (failover)
 - Improved security integration
 - Kerberos authentication
 - 2 different encryption options, header-only & payload
 - Performance improvements – FS Cache integration
- iSCSI Software Target under consideration for Red Hat Enterprise Linux 5.1
- Ext3 enhancements for speed and scalability
- More complete automounter with Autofs5
- Volume Management, Multipathing and SAN integration improvements.
- Single node GFS2 in base OS
 - Scales beyond the current Ext3 8TB limit and provides performance enhancements.

File System Features

- EXT3
 - Ext3 block reservation & on-line growth (2.6.10 & RHEL4)
 - Extended Attributes in the body of large inode
 - Saves space and improves performance (2.6.11)
 - Increases maximum ext3 file-system size from 8TB to 16TB (2.6.18)
- Device mapper multipath support (RHEL4)
- ACL support for NFSv3 and NFSv4 (2.6.13)
- NFS
 - Support large reads and writes on the wire (2.6.16)
 - Linux NFS client supports transfer sizes up to 1MB

Runtime Environment

- GCC 4.1
 - Including 4.2 backport of OpenMP
 - More complete Java 1.4 in gcj and class libraries, Fortran95 support
- Glibc 2.5, Libstdc++ 4.1
- SystemTap, Oprofile and Frysks enhancing serviceability
 - Also designed for optimization of production environments
- Backwards compatibility for Red Hat Enterprise Linux 3 and 4
 - Userspace applications that are compiled for Red Hat Enterprise Linux 3 or 4 are expected to continue to work unmodified in Red Hat Enterprise Linux 5
 - Additional compatibility options via unchanged stack in virtualized environment
- ISVs certify in guest environment
 - Kernel ISVs are a special case

File System Features

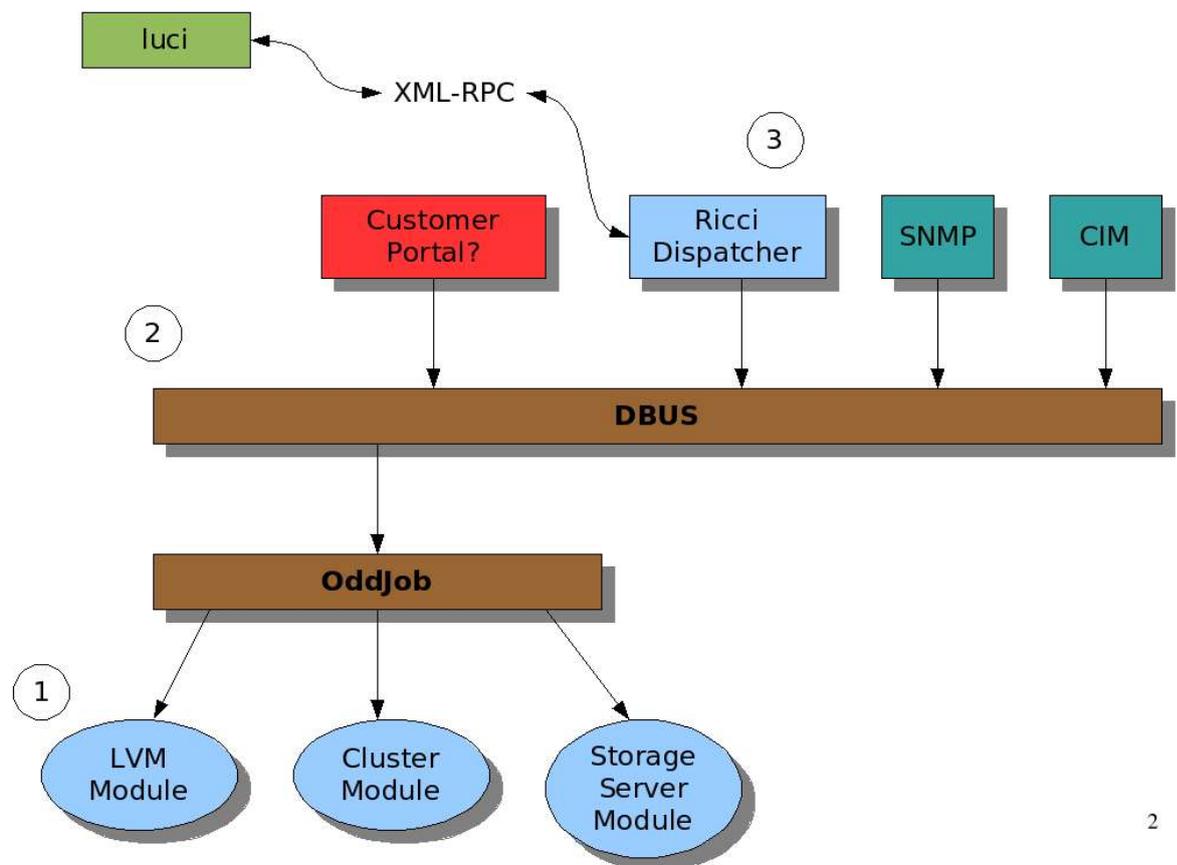
- VFS changes
 - "Shared subtree" enhancement (2.6.15). Ref: <http://lwn.net/Articles/159077/>
 - Read Copy Update used for improved performance & locking
- Big CIFS update: (2.6.15)
 - Lots of performance improvements
 - Kerberos and CIFS ACL support
- Autofs4
 - Updated to provide direct mount support for userspace autofs (2.6.18)
- FS-Cache
 - Core enablers (2.6.18)
 - Tech preview until 5.1



Conga: Cluster and Storage Management

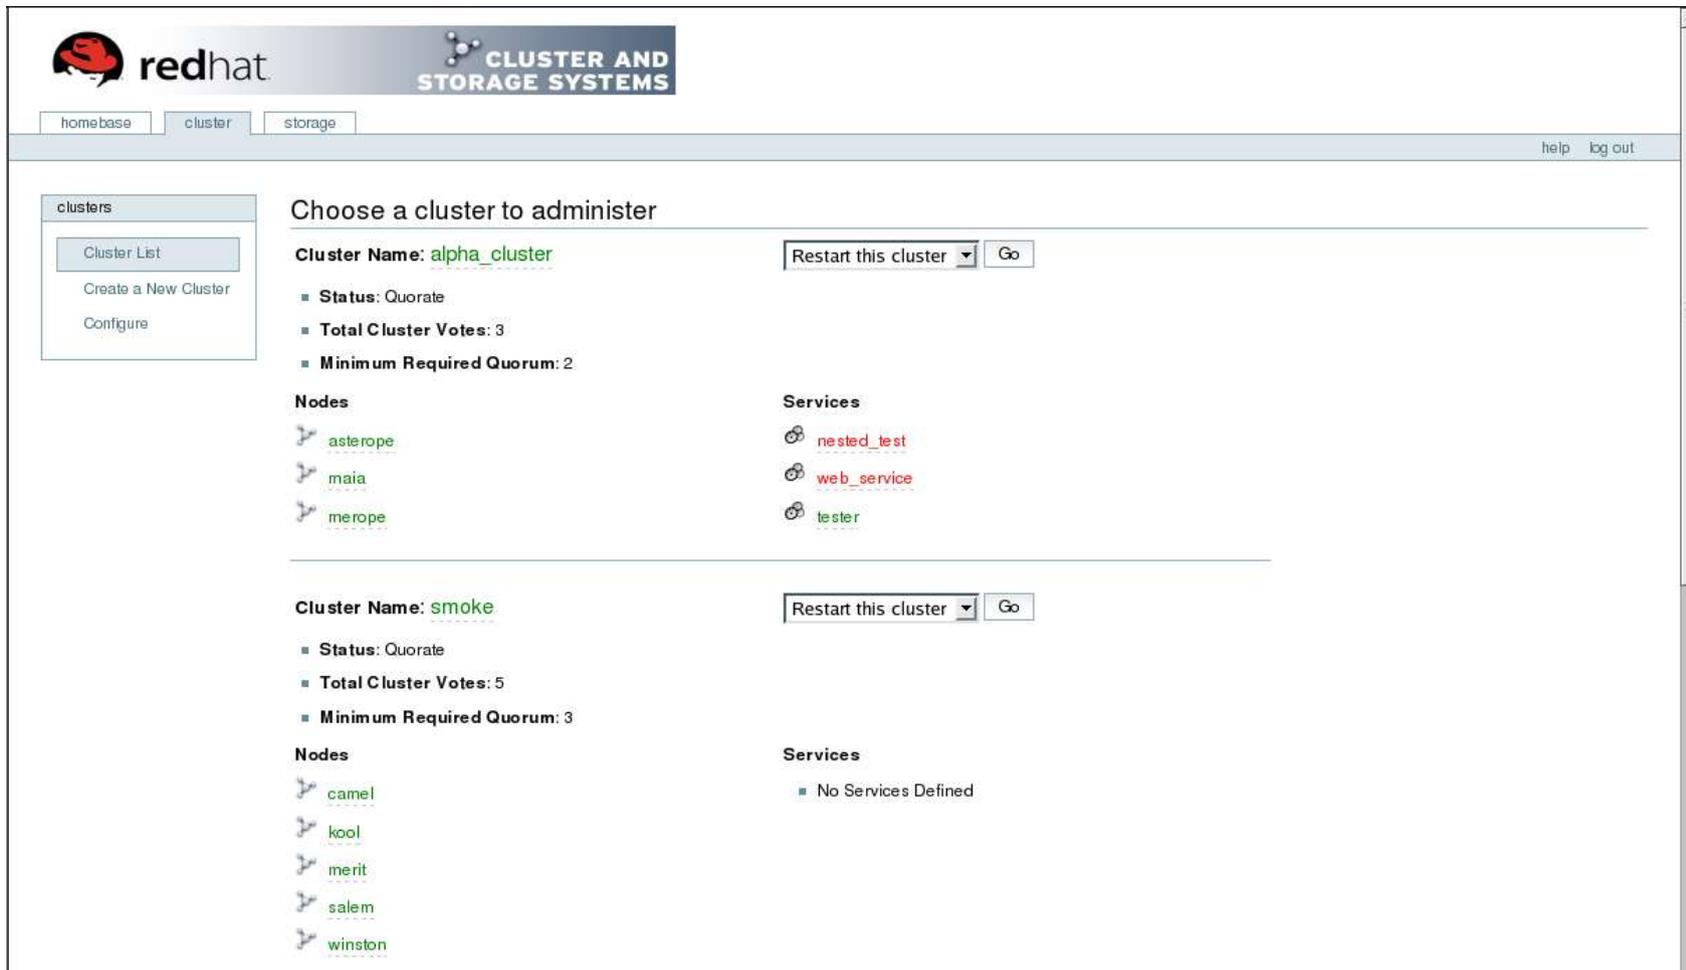
Cluster & Storage Management: Conga

- New web/XML based management tool for cluster and storage management
- Provides multi-system and multi-guest management
 - Installation, configuration, control
 - Flexible, extensible
- Included with Advanced Platform



Conga: Cluster Management

- Controls multiple clusters
- Comprehensive features



The screenshot displays the Red Hat Conga Cluster Management web interface. At the top, the Red Hat logo and 'redhat' text are on the left, and 'CLUSTER AND STORAGE SYSTEMS' is on the right. Below this, there are navigation tabs for 'homebase', 'cluster', and 'storage', with 'cluster' selected. On the far right of the top bar are links for 'help' and 'log out'.

The main content area is titled 'Choose a cluster to administer'. On the left, a sidebar menu under 'clusters' includes 'Cluster List', 'Create a New Cluster', and 'Configure'. The main area shows two cluster entries:

- Cluster Name: alpha_cluster**
 - Status: Quorate
 - Total Cluster Votes: 3
 - Minimum Required Quorum: 2
 - Nodes:** asterope, maia, merope
 - Services:** nested_test, web_service, tester
- Cluster Name: smoke**
 - Status: Quorate
 - Total Cluster Votes: 5
 - Minimum Required Quorum: 3
 - Nodes:** camel, kool, merit, salem, winston
 - Services:** No Services Defined

Each cluster entry includes a 'Restart this cluster' dropdown menu and a 'Go' button.

Conga: Storage Management

- Controls multi-system and multi-guest logical volumes
- Comprehensive features

The screenshot displays the Conga storage management interface for the system `salem.lab.msp.redhat.com`. The interface includes a navigation menu on the left with options like "System List", "Hard Drives", "Partition Tables", "Volume Groups", and "New Volume Group". The main area shows a graphical view of the "Smoke_Cluster" volume group, consisting of a blue cylinder representing logical volumes and a red cylinder representing physical volumes. A portion of the physical volume is labeled "Unused Space". Below this, a form titled "Unused Space - Creating New Logical Volume" is shown, with the following fields:

Unused Space - Creating New Logical Volume	
Logical Volume Name	<input type="text" value="new_lv"/>
Volume Group Name	Smoke_Cluster
Size	<input type="text" value="395.24"/> (0.00 - 395.24) GB
Clustered	true
Content <input type="text" value="GFS2 - Global FS v.2"/>	
Cluster Name	smoke
Unique GFS Name	<input type="text" value="unique_gfs_name"/>
Journal Size	<input type="text" value="33554432"/>
Block Size	<input type="text" value="4.0 KB"/>
Mountpoint	<input type="text"/>
Mount	<input type="text" value="false"/>
List in /etc/fstab	<input type="text" value="false"/>
Mountable	true
Locking Protocol	dlm
Number Of Journals	<input type="text" value="5"/> (1 - 128)
Clustered	<input type="text" value="true"/>



Red Hat Enterprise Linux 5

Stateless Linux

Stateless Linux: Overview

- The Stateless Linux initiative enables systems to be set up as replaceable appliances, with no important local state
 - For example, a system administrator can set up a network of hundreds of desktop client machines as clones of a master system
 - Virtual machine instances can be deployed rapidly as clones of a pre-configured master
 - Stateless systems are kept synchronized with the master system
- Benefits:
 - Management Cost Reduction & Improved Scalability
 - Manage/provision/update multiple systems as though they were one
 - Easy replacement of failed systems
 - Improved data security and management
 - Centralize system state and application data
 - Centralized control of disconnected systems
 - Simple Laptop management



Stateless Linux: Servers and Clients

- A Stateless Linux server provides the following:
 - It stores each of the prototype (master) systems
 - It stores snapshots in time of those systems
 - It stores home directories for the users of the client systems
 - Post RHEL5 release:
 - It maintains an LDAP directory containing information on which client machines should be running which snapshots of which prototype system
- Several Stateless Linux clients are possible:
 - Diskless clients, which boot directly from a snapshot stored on the server
 - Caching clients, which boot from a copy of a snapshot, cached locally on a hard drive
 - Post RHEL5 release:
 - Live CD clients, which boot from a copy of a snapshot burned onto a CD

Stateless Linux: Technologies

- The Stateless Linux initiative requires a number of technology enhancements, including:
 - Read-only root filesystem support
 - All hardware is configured automatically at boot, without pre-existing configuration
 - Network Manager, X subsystem
 - Network root/network boot
 - NFS & iSCSI
 - Support for using local storage as a cache's block device (dependent on upstream acceptance)
 - Support for swapping to local storage
 - LiveCD support (post RHEL5 release)
 - Client booting of server-resident, LVM-based snapshots
 - CoW (Copy on Write) instances, automatically re-synched with master at reboot

Stateless Linux: Rollout

- Project is currently in Technology Preview, to be delivered to two major phases:
 - Core stateless technologies with simple management tools
 - Improved management tools (post RHEL5 release)
- Removing state from the system is equivalent to removing and automating as many “Settings” as possible (network, video, etc) and converting the rest to “Preferences”
- Initial usage models:
 - Virtual server: boot/run entirely from network/SAN. No local storage required
 - Corporate desktops with local storage: stateless images are synced to the local disk and run from there
 - Desktops – local swap, read-only root synced to the local disk, and network homedirs
 - Laptops – local swap, read-only root synced; local homedirs synced back to centralized network storage for backups
- See: <http://fedoraproject.org/wiki/StatelessLinux>



Red Hat Enterprise Linux Client

Red Hat Enterprise Linux Desktop 5

- Updated desktop environment and applications
 - **OpenOffice 2.0**
 - Oasis OpenDocument format is now the default
 - Improved interoperability with Microsoft formats
 - Addition of Base – new database tool (MS Access)
 - **Firefox 1.5**
 - Continuing to drive innovation for web browsers
 - **GNOME 2.16**
 - New applications



Red Hat Enterprise Linux Desktop 5

- Upgraded laptop support
 - **New power management features**
 - Suspend / Resume / Hibernate
 - **Reduced power usage in battery mode**
 - **Better userspace notification of system status**



Red Hat Enterprise Linux Desktop 5

- Network management
 - **Network Manager**
 - Automatically detects and configures network connections
 - Seamlessly switches between network connections
 - Provides feedback on network state
 - Integrates with VPN clients
 - Supports WPA/WEP



Red Hat Enterprise Linux Desktop 5

- Modern graphics and user experience
 - **Modularized X.org**
 - Easier to test and deploy X components
 - Driver auto-detection
 - **Composite-based architecture**
 - New effects
 - Cleaner graphics experience
 - **Graphics accelerating**
 - Multiplexed access to direct 3D rendering



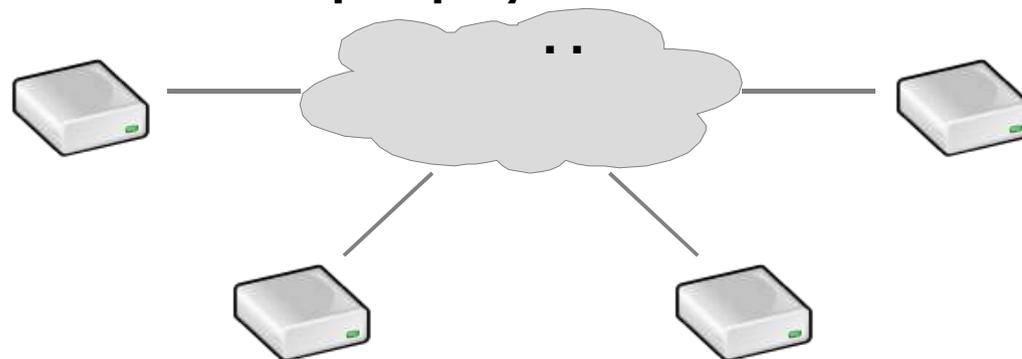
Red Hat Enterprise Linux Desktop 5

- Management enhancements beyond RHN
 - **Stateless Linux**
 - Revolutionary new image-based management system
 - **Single Sign On**
 - Improved authentication
 - **Sabayon**
 - User profile management

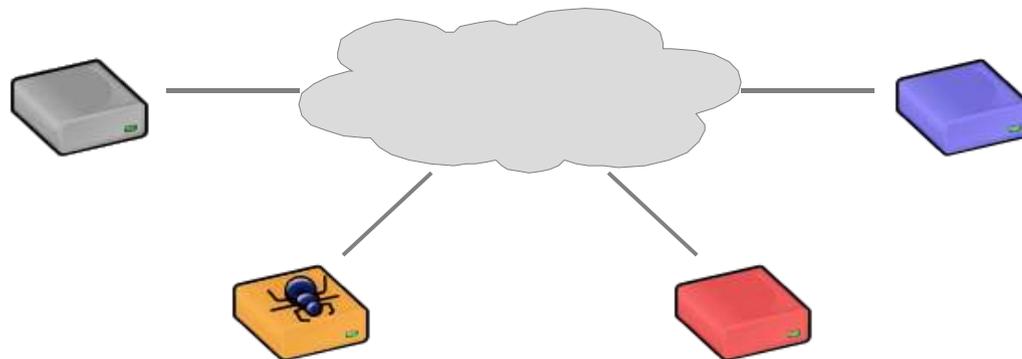
Red Hat Enterprise Linux Desktop 5

- Stateless Linux

Desktop deployment starts as .



. . . but six months later becomes . . .

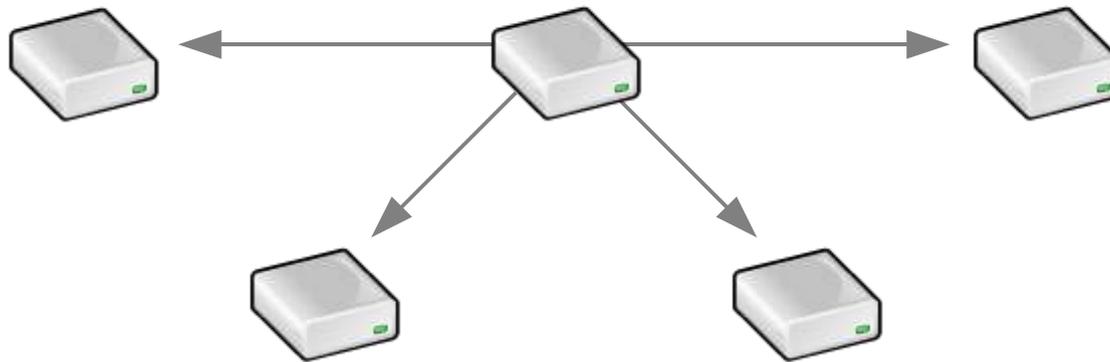


. . . and forever after is labor intensive to manage

Red Hat Enterprise Linux Desktop 5

- Stateless Linux

Image-based model (changeable image on master server, local filesystems are read-only) . . .



. . . single image deployed across all systems leads to lower management effort and improved security

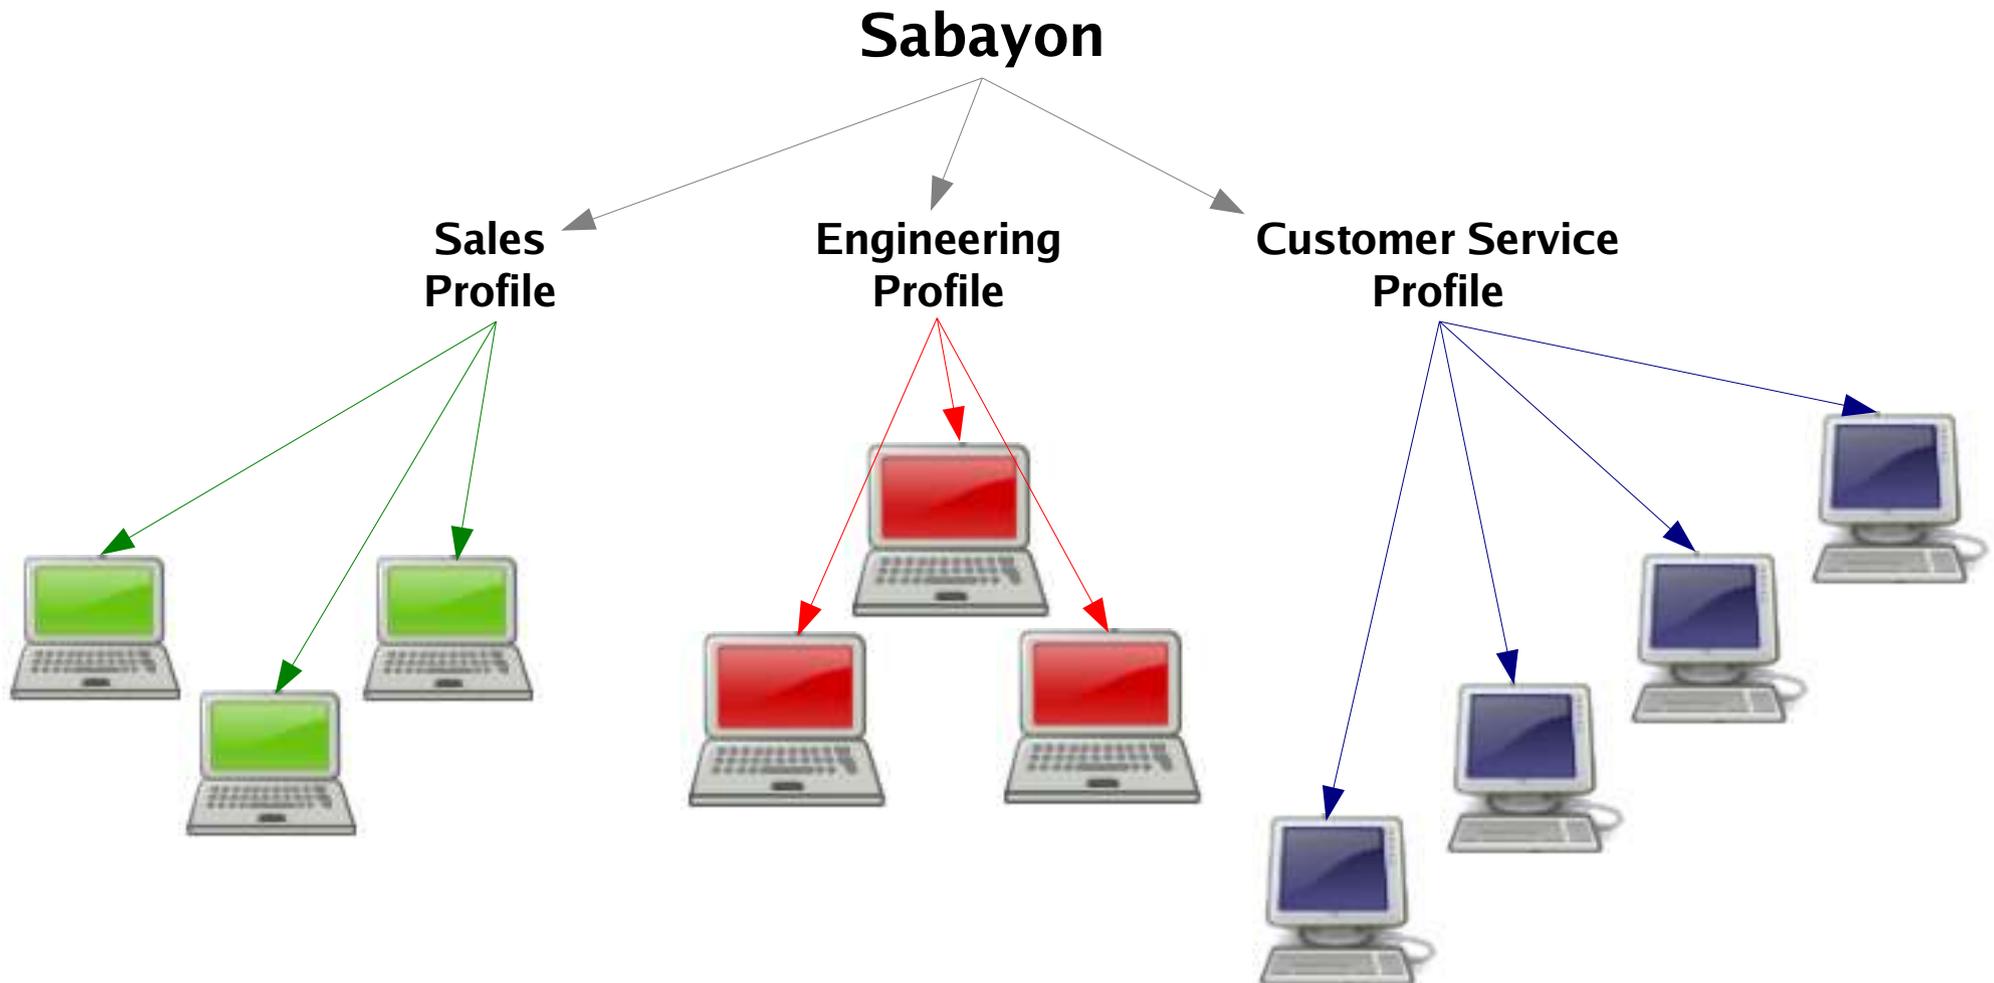
Red Hat Enterprise Linux Desktop 5

- Single Sign On
 - **Improved integration of security components**
 - Easier set up for allowing users to authenticate only once
 - **SmartCard integration**
 - Ability to authenticate using a SmartCard with a PIN
 - **Secure access to a host of applications and services**
 - GNOME KeyRing
 - File sharing (SMB, NFS)
 - Firefox and Thunderbird (NSS)
 - SSH
 - Shared printers



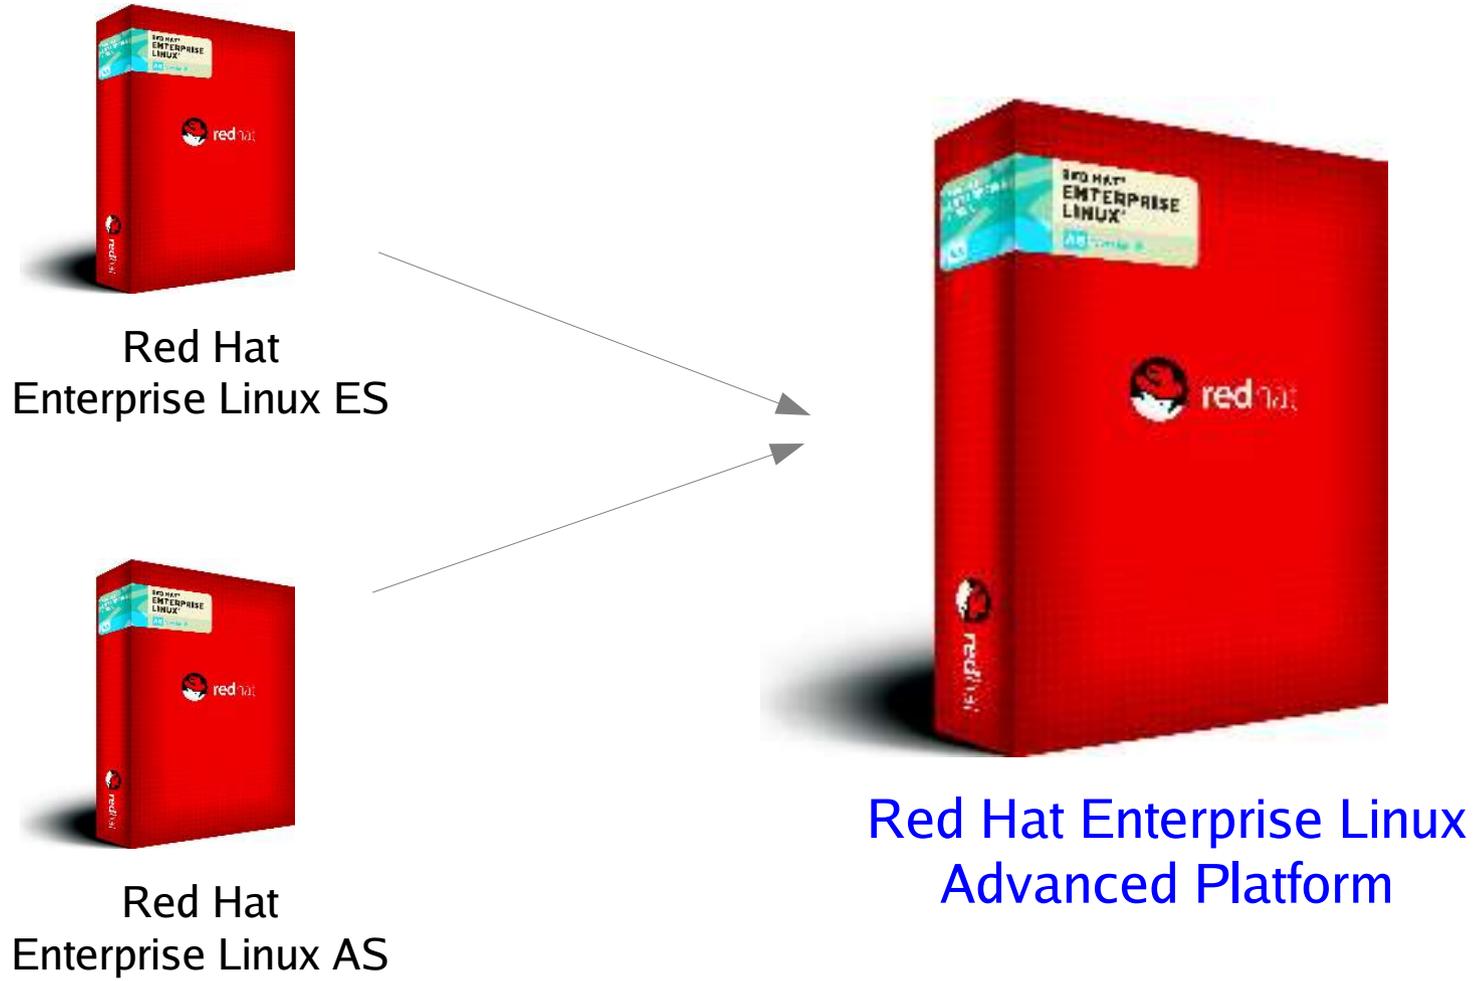
Red Hat Enterprise Linux Desktop 5

- Sabayon



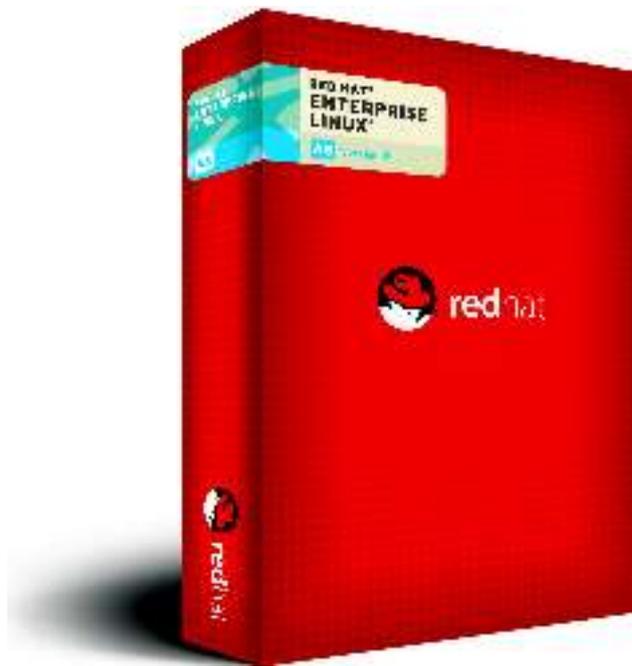


Red Hat Enterprise Linux Packaging

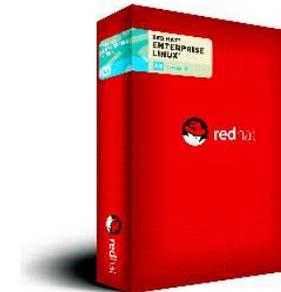


Feature Differences:

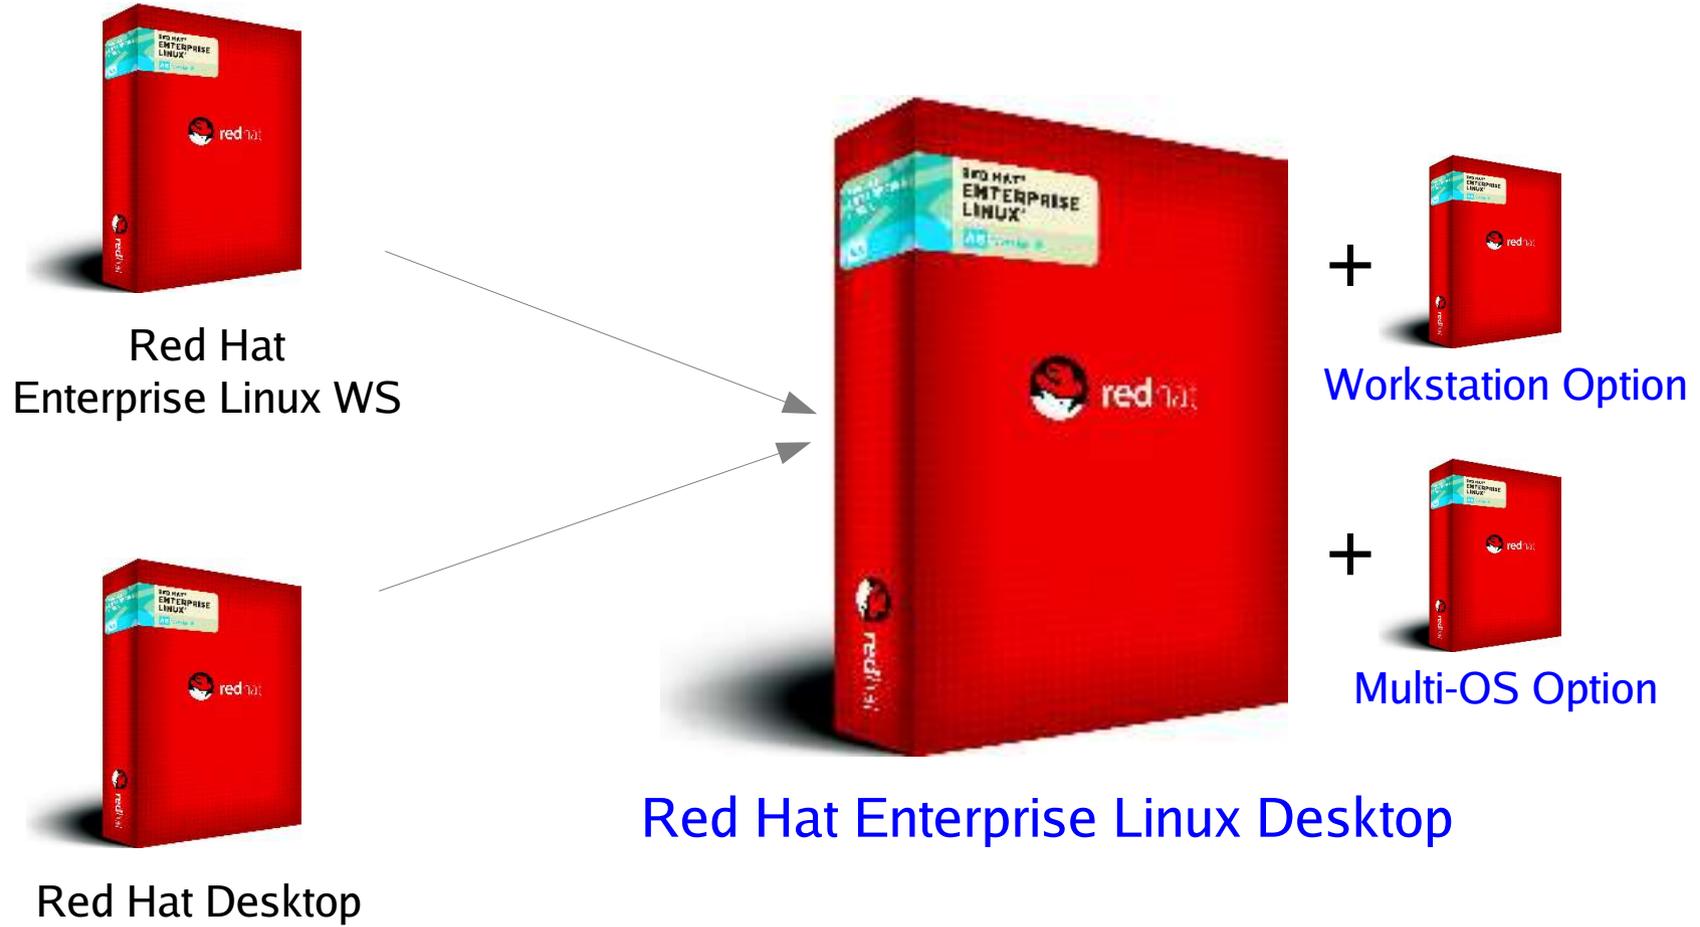
	<u>Sockets</u>	<u>Virtual Guests</u>	<u>GFS</u>	<u>Cluster Suite</u>
RHEL (1-2 socket)	1-2	4	No	No
RHEL Advanced Platform	unlimited	unlimited	Yes	Yes



Red Hat Enterprise Linux
Advanced Platform

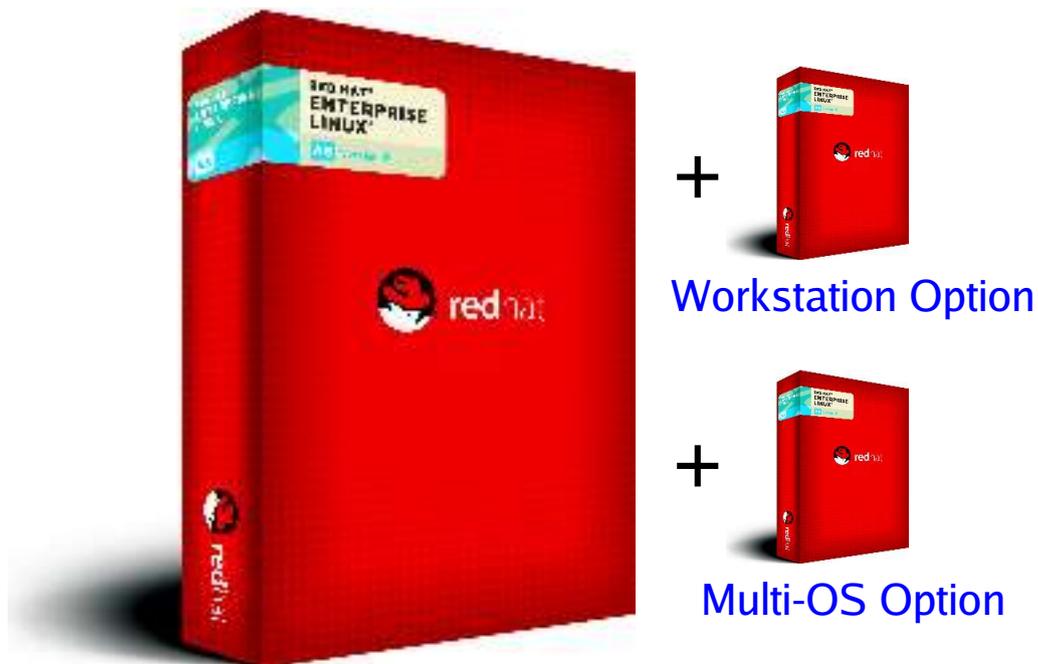


Red Hat
Enterprise Linux
(1-2 socket)



Feature Differences:

	<u>Sockets</u>	<u>Memory</u>	<u>Added Packages</u>	<u>Virtualization</u>
RHEL Desktop	1	4 GB	-----	No
Workstation Option	2	unlimited	server/developer	No
Multi-OS	---	----	-----	Yes

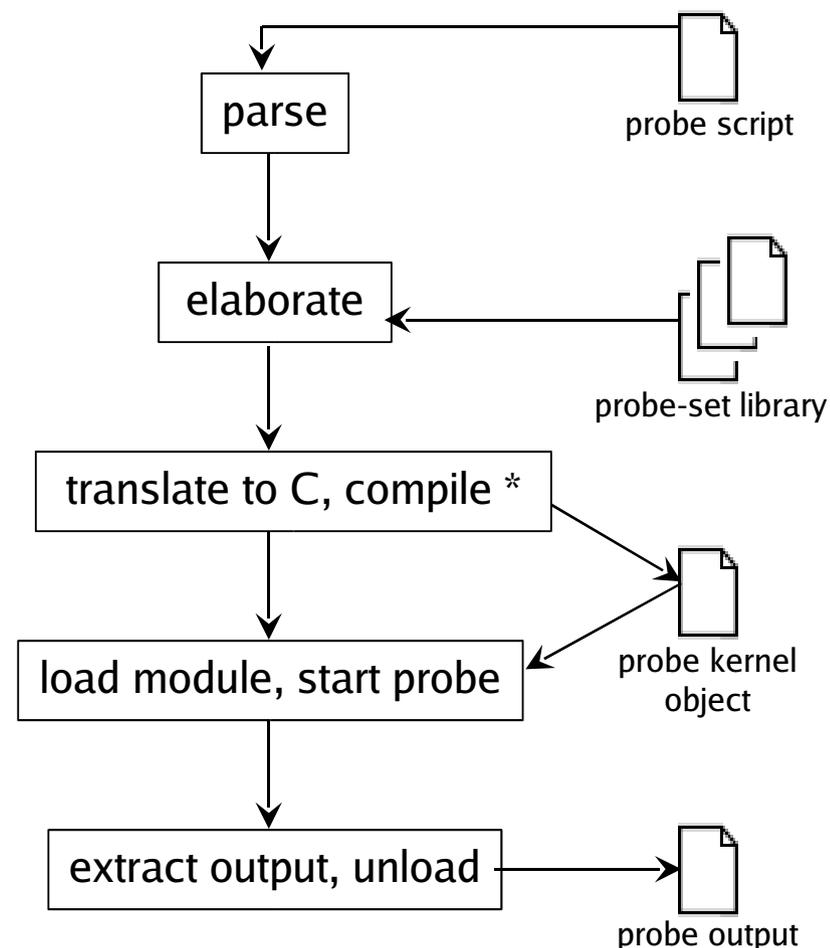


Red Hat Enterprise Linux Desktop
 (includes RHN Management and Provisioning)



Profiling Tools: SystemTap

- Red Hat, Intel, IBM & Hitachi collaboration
- Linux answer to Solaris Dtrace
- Dynamic instrumentation
- Tool to take a deep look into a running system:
 - Assists in identifying causes of performance problems
 - Simplifies building instrumentation
- Current snapshots available from: <http://sources.redhat.com/systemtap>
 - Source for presentations/papers
- Kernel space tracing today, user space tracing under development
- Technology preview status until 5.1



* Solaris Dtrace is interpretive

Dtrace vs SystemTap

- | | |
|--|--|
| Fixed probe points in kernel | ■ Probes at any location in kernel |
| Fixed pool of accessible data | ■ Extract any (debugger-visible) data |
| D language with limited capabilities | ■ Scripting language with control structures |
| Interpreted | ■ Compiled |
| Providers: Statically compiled code inserted in kernel or applications | ■ Tapsets: External scripts defining libraries of probe points |