System z Update

2013-01-15

Brad Hinson
bhinson@redhat.com
Solution Architect

Product News

Current Lifecycle Milestones

https://access.redhat.com/support/policy/updates/errata/

Red Hat Enterprise Linux 2.1 – End of Maintenance was May 31, 2009

Red Hat Enterprise Linux 3 – End of Maintenance was October 31, 2010

Red Hat Enterprise Linux 4 – End of Maintenance was March 31, 2012

- **RHEL 5**

  GA Date: March 14, 2007
  End of Phase 1: Q4, 2012
  End of Phase 2: Q1, 2014
  End of Production: **March 31, 2017**

- **RHEL 6**

  GA Date: November 10, 2010
  End of Phase 1: Q2, 2016
  End of Phase 2: Q2, 2017
  End of Production: **November 30, 2020**

Release News

- RHEL 5.9 released 1/8
- RHEL 6.3 released 6/21/12
- MRG 2.2 released 9/24/12 (RHEL 6.x Realtime, Linux 3.x kernel)
- Red Hat Storage Appliance released
  - Gluster acquisition Nov 2011

Solutions Overview

RED HAT SOLUTIONS

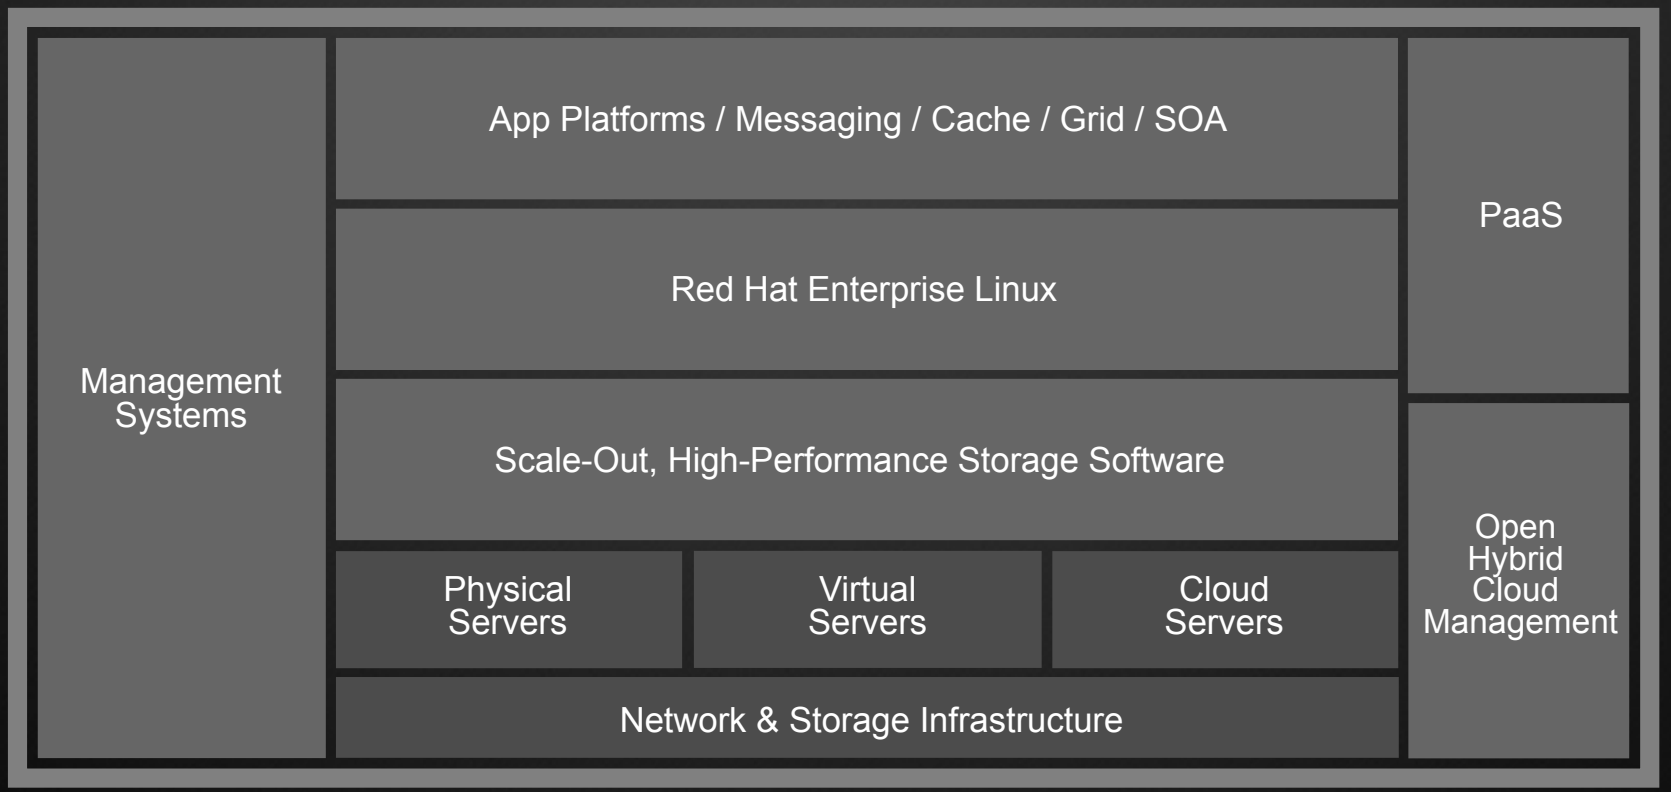

> I wouldn't be in the SaaS business if it weren't for Red Hat. If we'd had to make the capital investment in infrastructure that HP-UX or Solaris required, the margins for achieving profitability […] just wouldn't be there.
>
> —MARK TIRSCHWELL, CTO

WALLSTREET SYSTEMS

RED HAT SOLUTIONS

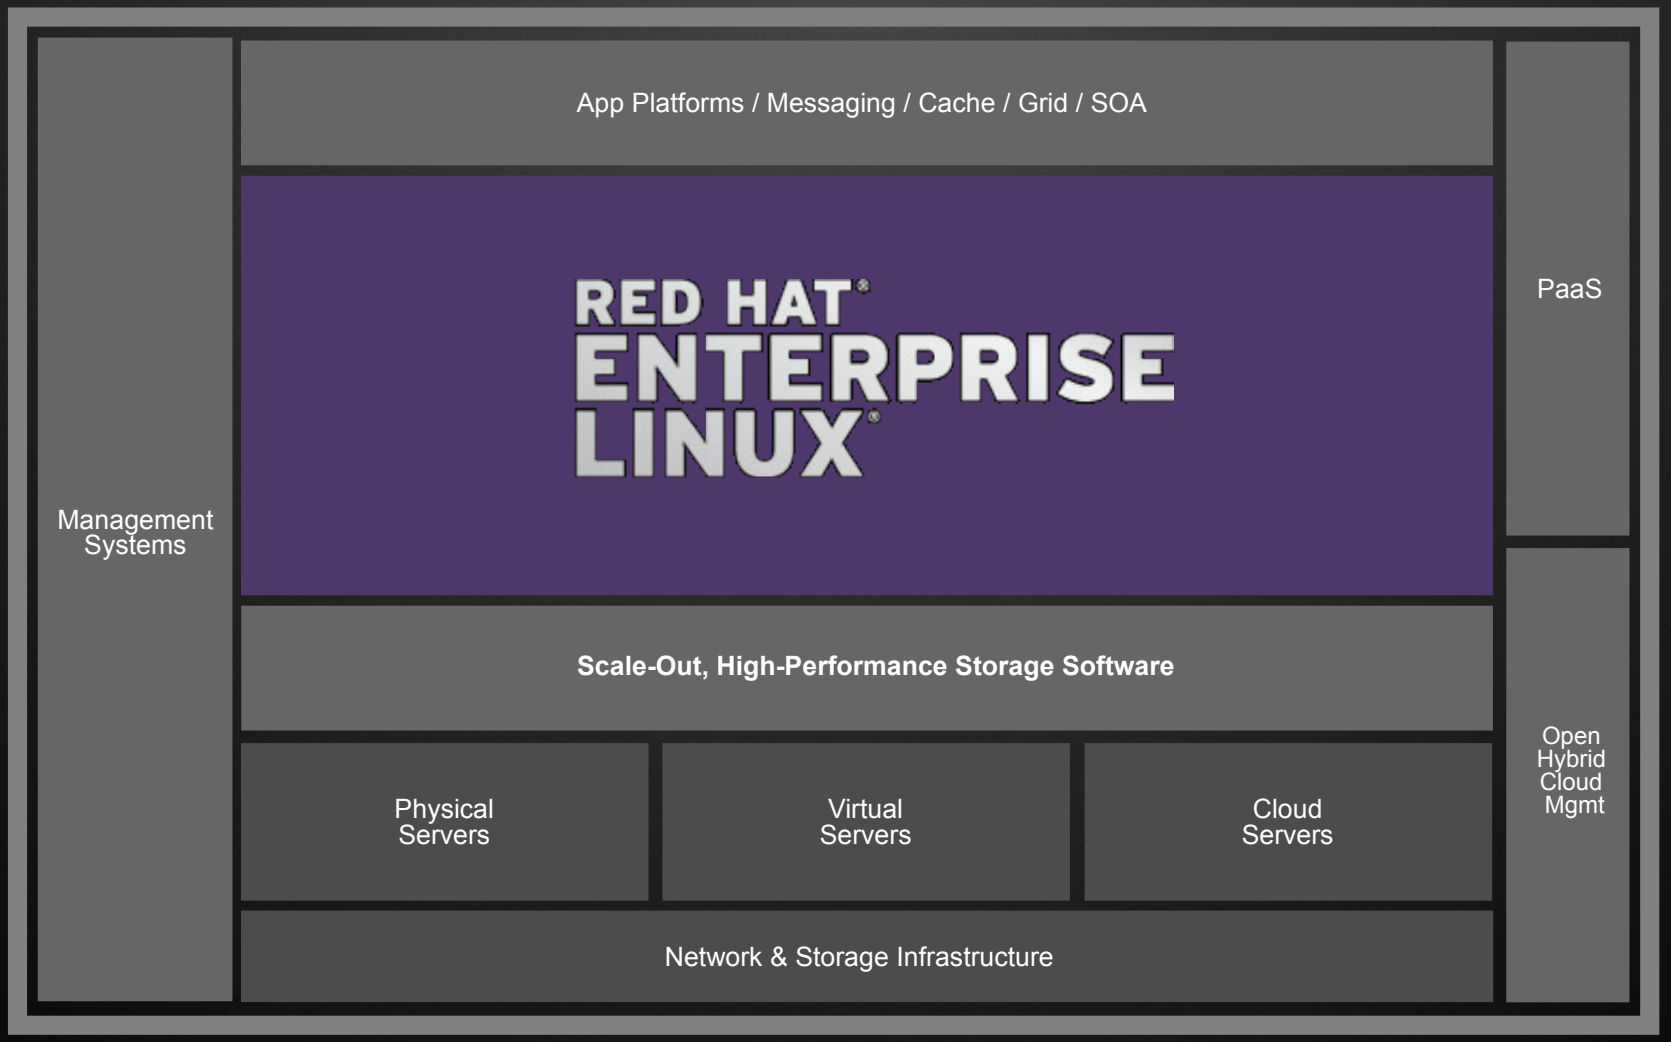

"The choice to invest in Red Hat was largely based on its commitment to the ongoing development of the platform & its strong support capabilities, particularly in reference to mainframe support. —LYLE JOHNSTON, Infrastructure Architect

bnz

Red Hat Enterprise Linux 5.9 Highlights

RHEL 5.9

- GA: Jan 8
- Mostly bug fix release
- Biggest System z feature
  - Enable HyperPAV for parallel I/O to ECKD DASD
  - Licensed feature in z/VM

RHEL 5.9 Highlights

- *Security, standards, and certification*
  - Tighter security, password management, FIPS for RAID
- *New developer tools*
  - SystemTap improvements, faster compiles
- *Enhanced application support*
  - Samba 3.6 (full SMB2 support), updated rsyslog
- *Virtualization*
  - Hyper-V x86 drivers for improved performance
- *Enhanced subscription management*
  - Manage subscriptions locally or with Subscription Asset Manager

Red Hat Enterprise Linux 6

RHEL 6.4

- Beta available since Dec 4, 2012
- Target GA: Feb 21

RHEL 6.4 Features: Security

- Support for CryptoExpress4S in EC12
  - Support for the PCIe I/O drawer configuration
  - Support for digital signatures with new PKCS #11
  - Support for American Express EMV (Europay, Mastercard, Visa) cards

RHEL 6.4 Features: Performance

- Kernel instruction support to improve s/390 Java performance
  - Runtime instrumentation using realtime signals
  - Allow user-space processes to use transactional execution.

RHEL 6.4 Features: Performance

- Support new storage device: SCM via EADM subchannels
  - With this feature, Linux can access Storage Class Memory as a block device.
- HugeTLBfs support
  - System z support added to libhugetlbfs, a library which provides easy access to huge pages of memory. It is a wrapper for the hugetlbfs file system.

RHEL 6.4 Features: Performance

- Optimized zlib compression library
  - Optimize the existing compression library zlib by using dedicated SSE instructions and optimized compile options. zlib is used by Java (decompression of class files), Cognos (PDF generation), TSM (backup) and for Linux installations (binaries compressed in RPMs)

RHEL 6.4 Features: Tools

- Update to Valgrind 3.8.0
  - Valgrind has proven to be a valuable tool debugging user-space memory management problems.
- Update lscpu/chcpu
  - This simplifies and enhances the usability of CPU hotplug by providing additional functionality via scripts

RHEL 6.4 Features: Reliability

- Compare system dump with boot system
  - Add z/VM Live Guest Relocation detection
  - Better support for system recovery after LGR or suspend/resume
- Full support for 2-stage dump framework
  - Previously tech preview

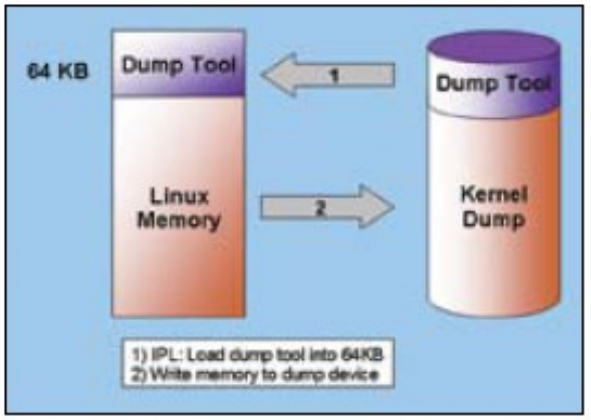

Figure 2: Standalone Dump Tool for Channel-Attached Devices

RHEL 6.4 Features: FCP

- Implement data routing for FCP
  - Enable FCP to pass data directly from memory to SAN (data routing) when memory on the adapter card is blocked by large and slow I/O requests.
- End-to-end data consistency checking
  - The T10 Technical Committee introduced an enhancement to the SCSI standard (SPC-4, SBC-3) to protect against errors in user data blocks. This introduces the zfcp-specific part in the Linux on System z I/O stack for E2E data consistency checking.

Fundamental RHEL 6 Improvements

CPUs

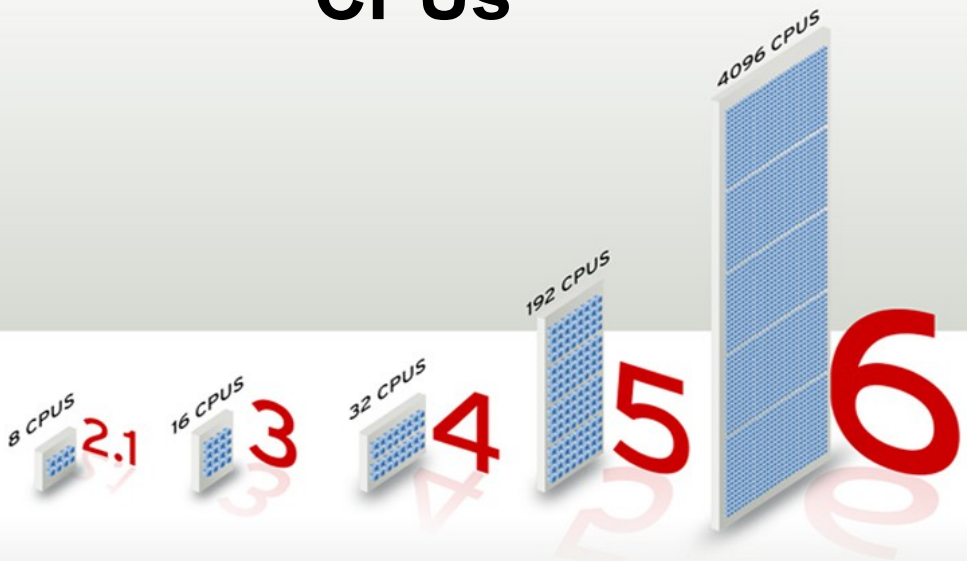

Memory

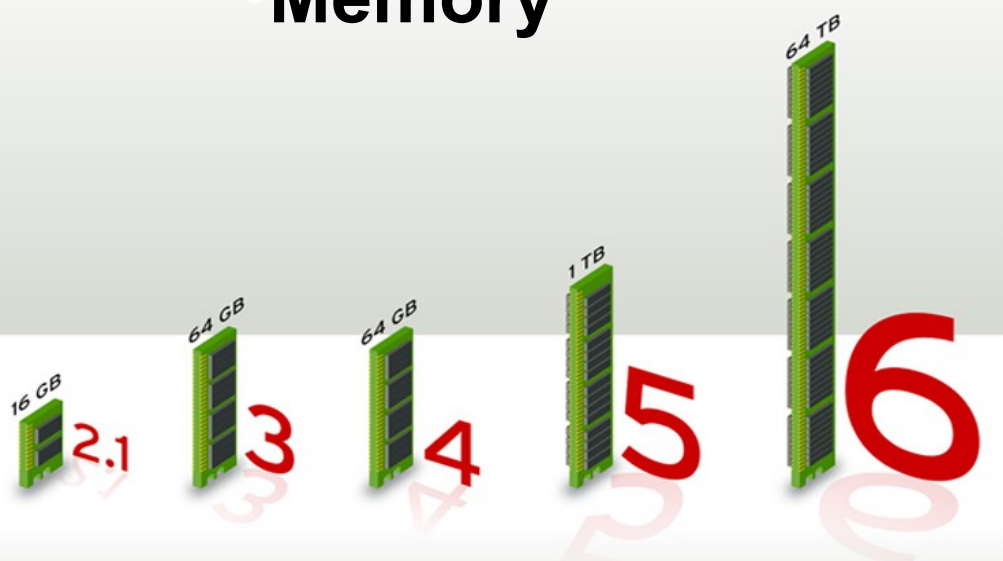

File Systems

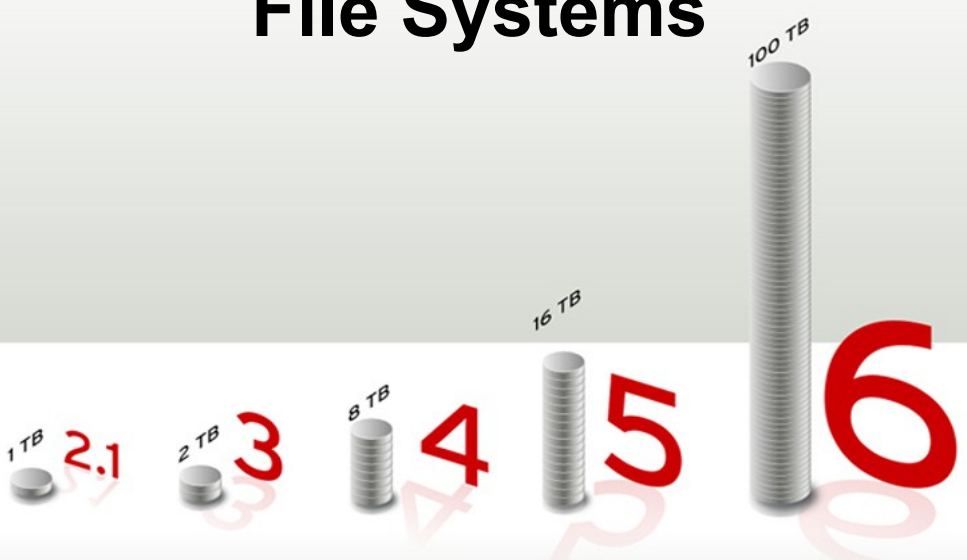

Scalability

Fundamental RHEL 6 Improvements

Network

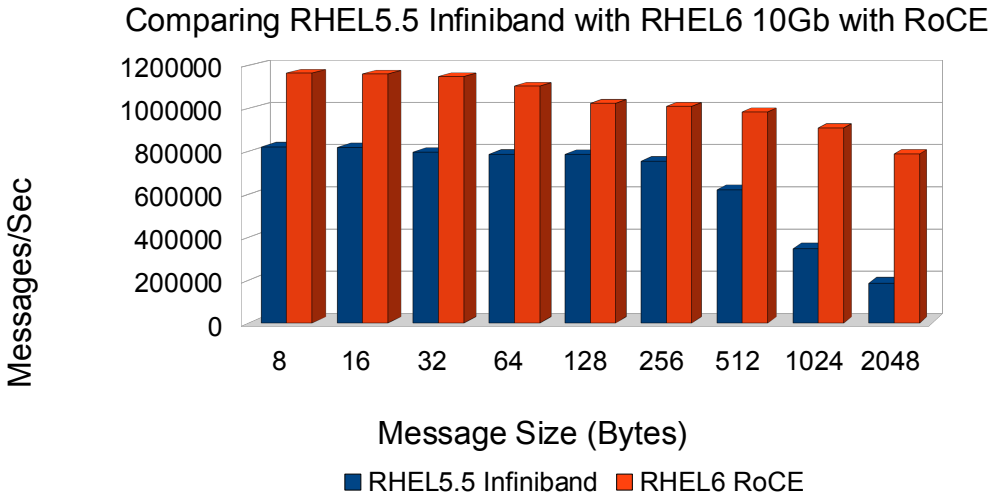

Filesystem

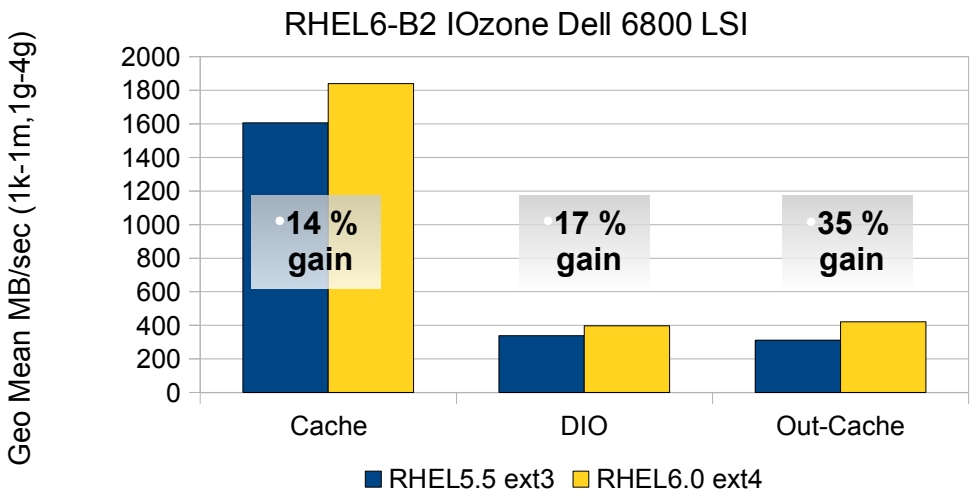

Power Consumption

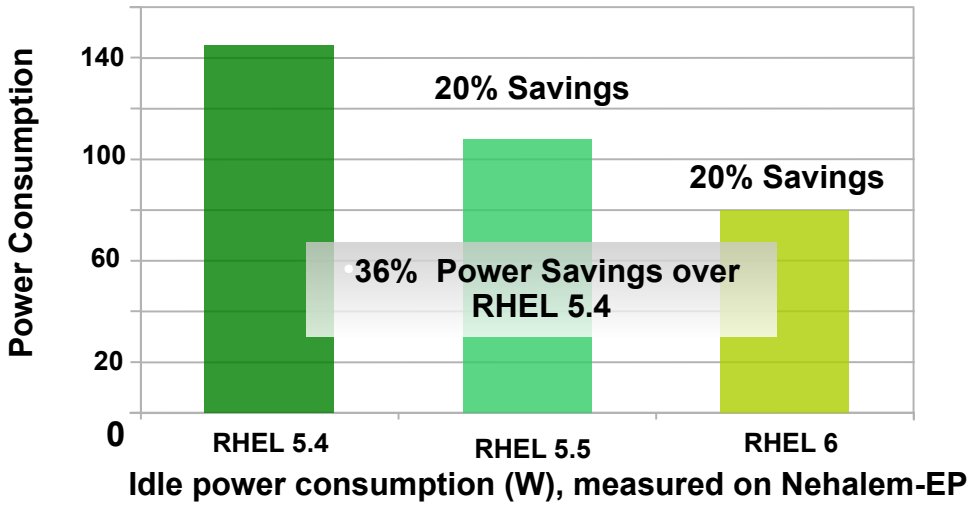

Performance

RHEL 6 Enhancements

- CPU Scheduler: Completely Fair Scheduler (CFS)
  - priority determined purely by CPU use
    (vs. complex heuristics in O(1) scheduler)
  - better SMP and NUMA balancing code
  - Less corner cases, more deterministic
  - includes micro accounting for scheduler timekeeping

RHEL 6 Enhancements

- Control Groups Resource Management (CGroups)
  - Dynamic allocation of:
    - processes, memory, storage & network
    - Reduce resource contention
    - Meet SLA's
    - Increase predictability & performance.

RHEL 6.3

Per Interface Network Priority (net_prio)

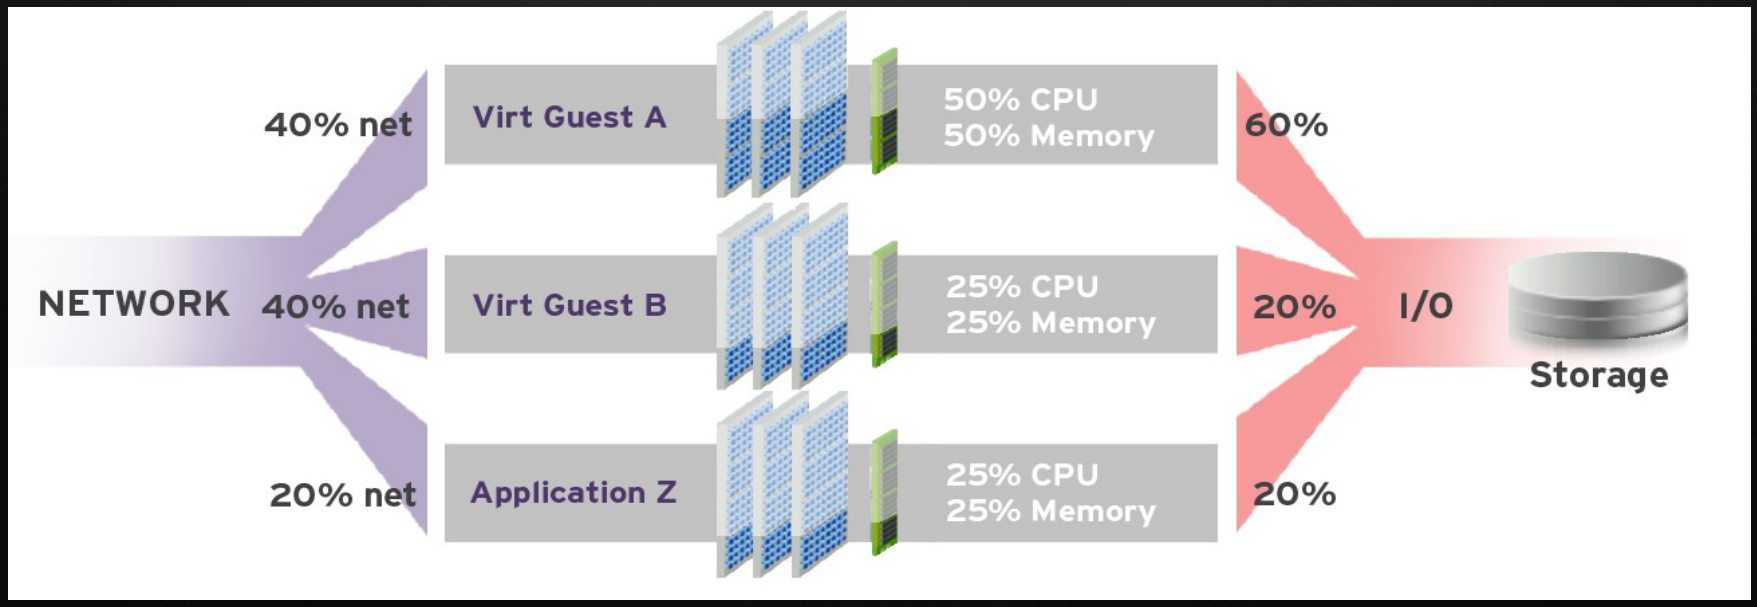

Red Hat Enterprise Linux 6: Tech Preview News

Remote File Systems

- Current technology: NFS (network file system)
  - Advantage: provides clustering capability, replaces GFS
  - Disadvantage: NFS server is single point of failure

Today's NFS Architecture

Remote File Systems

- Parallel NFS (pNFS) *(Technology Preview in 6.2)*
  - www.pnfs.com/docs/DMG_Parallel_NFS.pdf

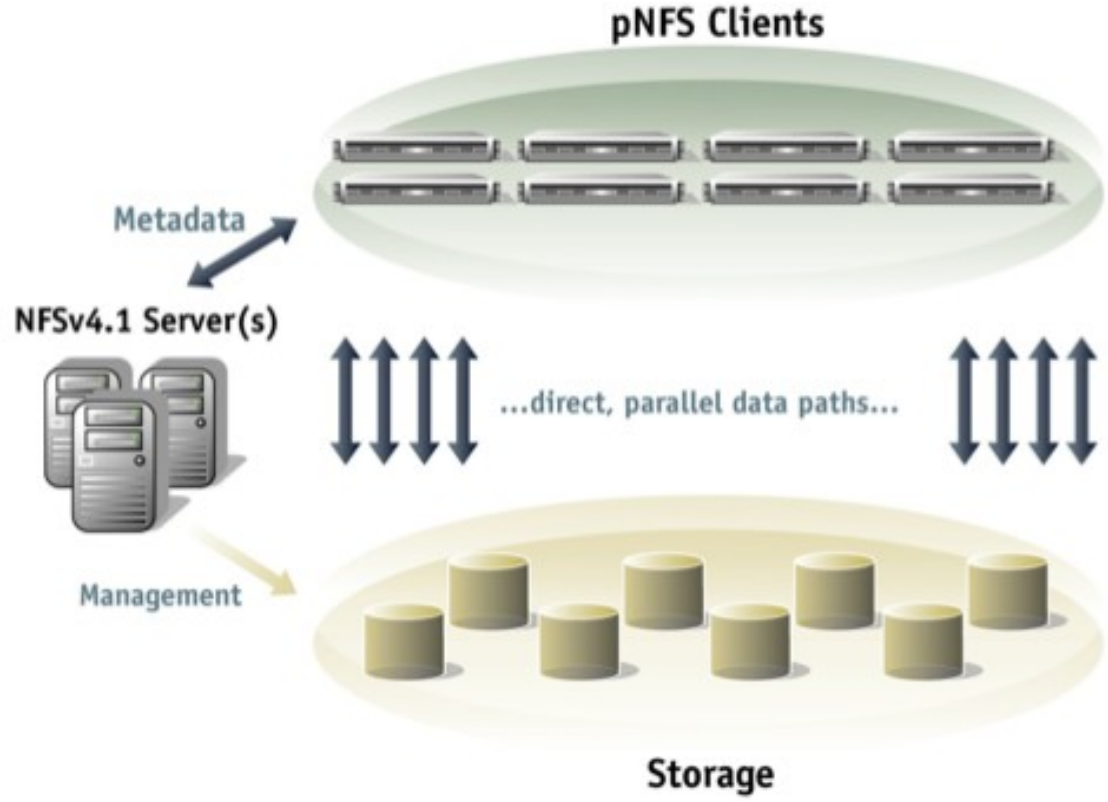

Parallel NFS architecture

Red Hat Enterprise Linux 7 Planning

RHEL 7 Schedule

- Alpha 2 delivered to Partners Jul 17th, 2012
- Public Beta Q2 2013
- Public GA TBD

RHEL 7 Kernel & Libraries

- Alpha 2: kernel-3.3.0-0.20
  - Current dev snapshot is kernel-3.6.0-0.29
- gcc-4.7.2-8
- gdb-7.4
- glibc-2.16-17
- boost-1.50.0-4

RHEL 7 Installer Plans

- Totally rewritten user interface / kickstart generator
  - Hub and spoke model for simplification and streamlining – less time answering questions – 3 screens for a standard install
- Memory footprint for installation reduced to 512 MB,
  - makes smaller guests supportable
- Stage 1 loader functionality merged into dracut so install environment boots same way as installed system

RED HAT SOLUTIONS

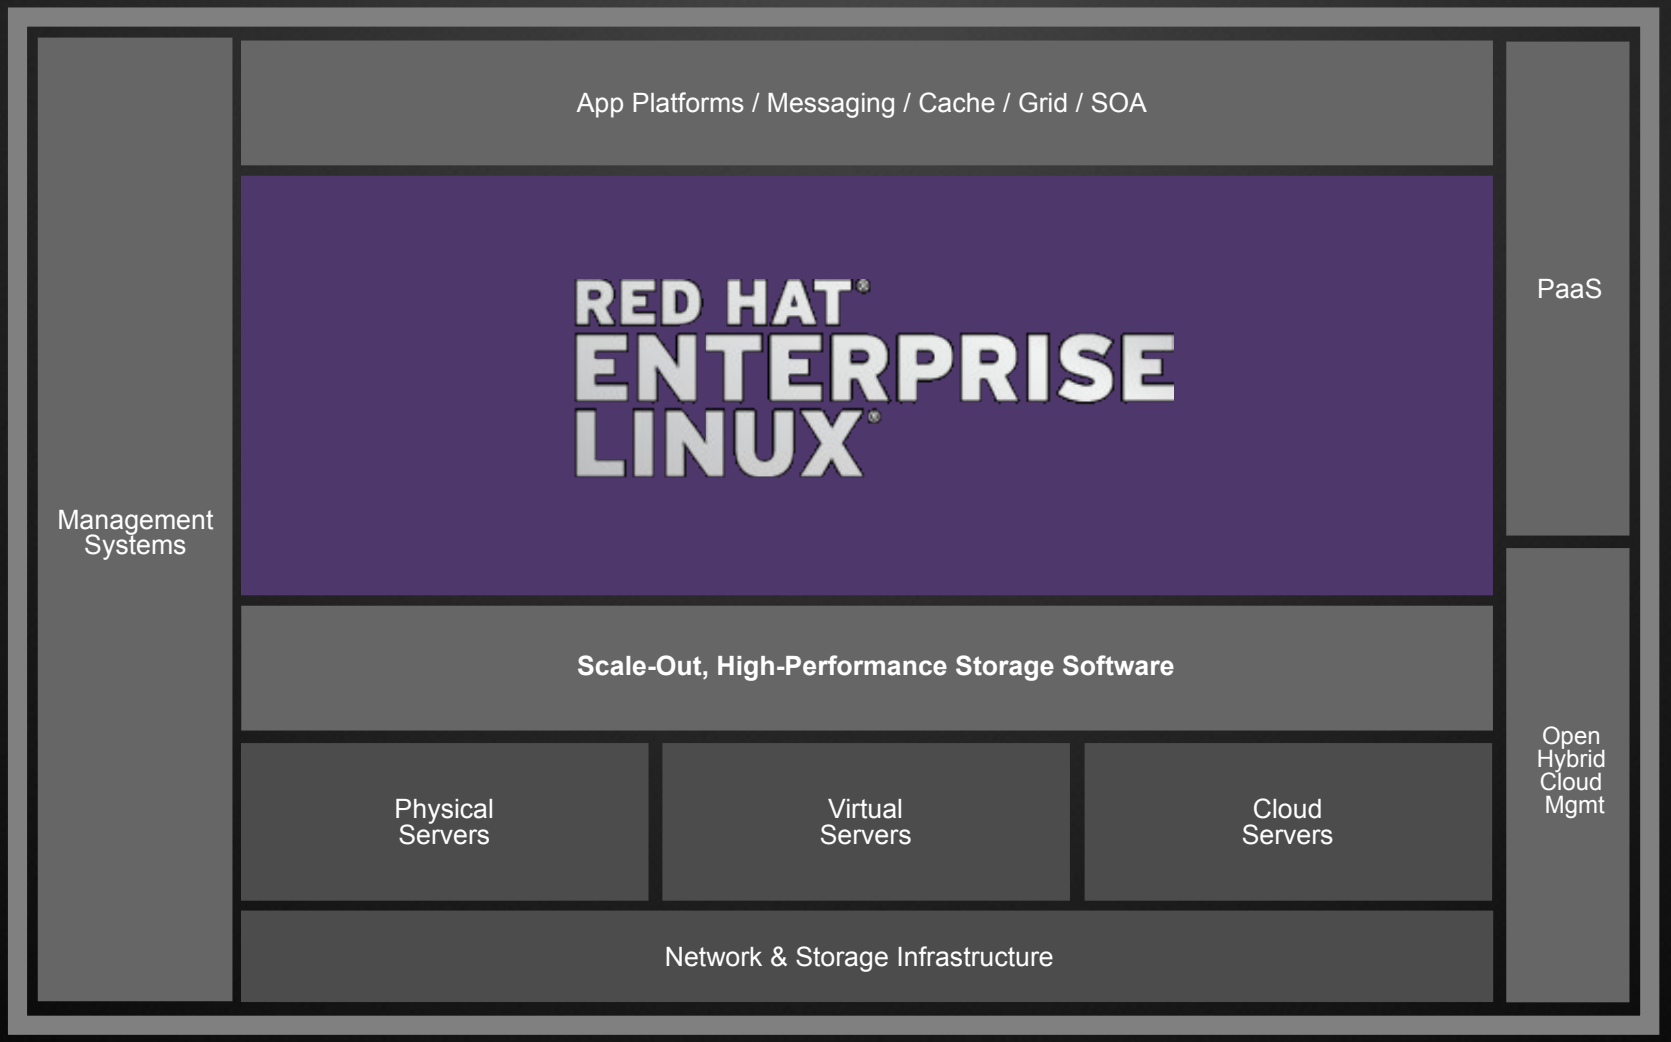

"The choice to invest in Red Hat was largely based on its commitment to the ongoing development of the platform & its strong support capabilities, particularly in reference to mainframe support. —LYLE JOHNSTON, Infrastructure Architect

bnz

RED HAT SOLUTIONS

THE UNIVERSITY OF TEXAS MD Anderson Cancer Center
Making Cancer History®

"We're running a world-class hospital. We need a mainstream enterprise OS with enterprise support. Red Hat delivers that support with Red Hat Enterprise Linux. —CHRISTOPHER GRAY, Manager Server Technical Support

WHAT IS RED HAT STORAGE?

Open, Unified, and Extensible Scale-out
Network-Attached Storage (NAS) and
Object Storage Software Solution for On-premise,
Virtualized and Cloud Environments

RED HAT STORAGE DESIGN GOALS

Scale out

- Elimination of metadata
- Effective distribution of data to achieve scalability and flexibility

Linear Scaling

- Capacity – scale up vertically
- Performance – scale out horizontally

Elasticity

- Flexibly adapt to the growth or reduction of data in the enterprise
- Add or remove resources to/from storage pool with zero application disruption

Deployment Agnostic

- Deploy on-premise, in the public cloud or a hybrid setup.

Must run on commodity hardware

- Industry standard servers.

  No-purpose built hardware.

WHAT IS IN RED HAT STORAGE?

- RED HAT ENTERPRISE LINUX
- XFS
- GLUSTER FS
- RED HAT STORAGE CONSOLE MANAGEMENT STATION

PHYSICAL SERVER: 2 CPU WITH LOCAL DISKS

or

VIRTUAL SERVER: AMAZON AWS, OR VIRTUALIZED RED HAT ENTERPRISE LINUX

A PRE-INTEGRATED, PRE-VERIFIED AND READY TO RUN SOFTWARE PLATFORM

SOURCED BY CUSTOMER

RED HAT STORAGE USER PERSPECTIVE (DISTRIBUTED VOLUMES)

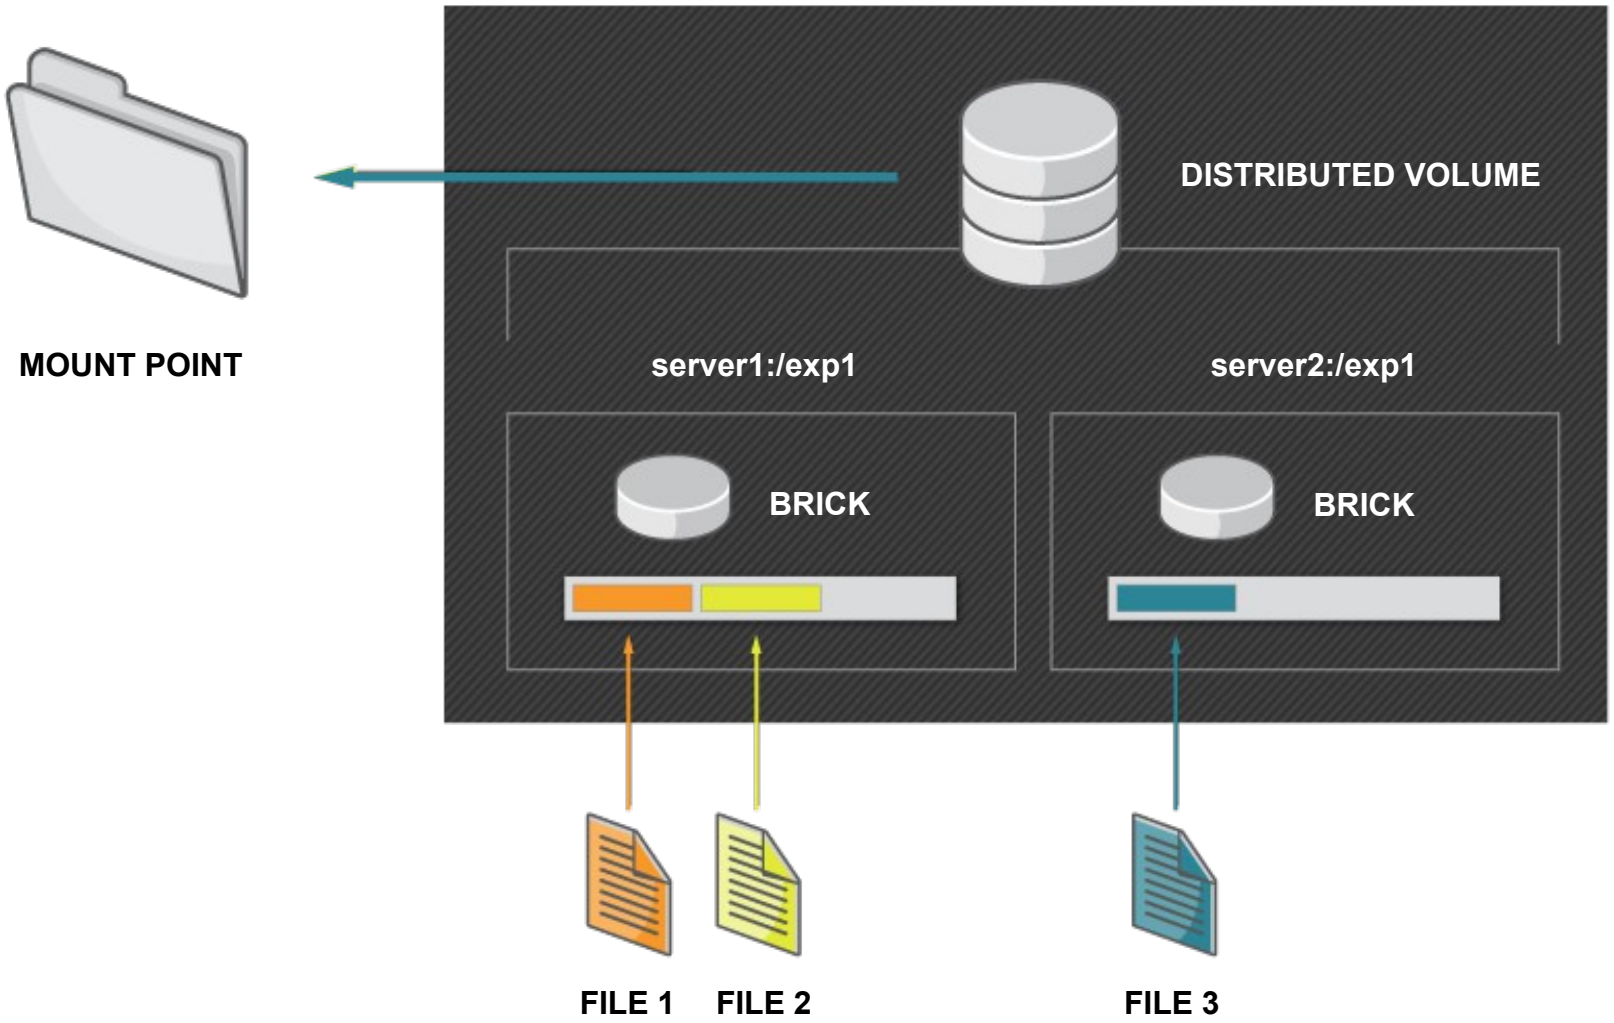

RED HAT STORAGE USER PERSPECTIVE (REPLICATED VOLUMES)

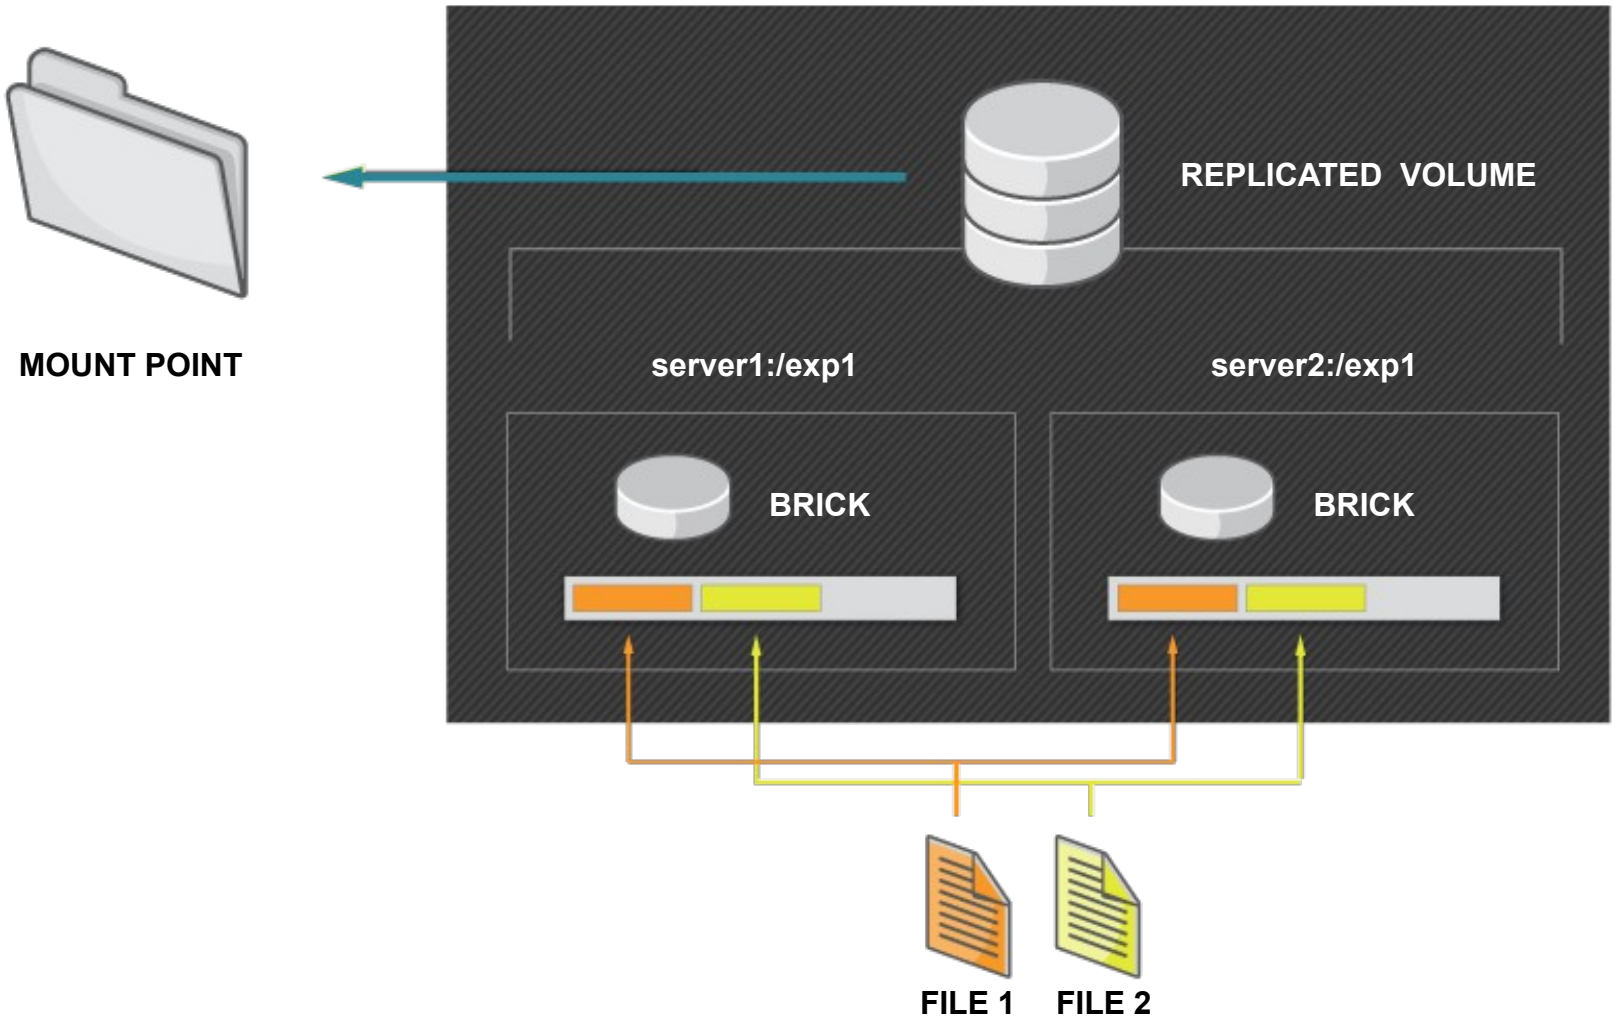

WHAT CAN BE STORED IN RED HAT STORAGE?

FILES AND FOLDERS

Documents, Photos,Videos, and Images

10101110101010101
10101011001010111
10101110100010 10
01011010010 10101

BIG DATA

Log files, RFID data

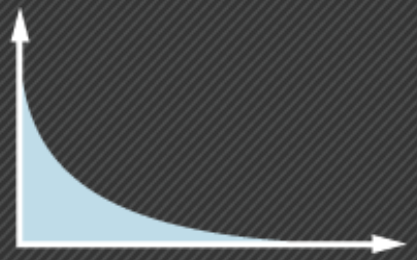

OBJECTS

Long Tail Data

RED HAT SOLUTIONS

THE UNIVERSITY OF TEXAS MD Anderson Cancer Center
Making Cancer History®

"We're running a world-class hospital. We need a mainstream enterprise OS with enterprise support. Red Hat delivers that support with Red Hat Enterprise Linux. —CHRISTOPHER GRAY, Manager Server Technical Support

RED HAT SOLUTIONS

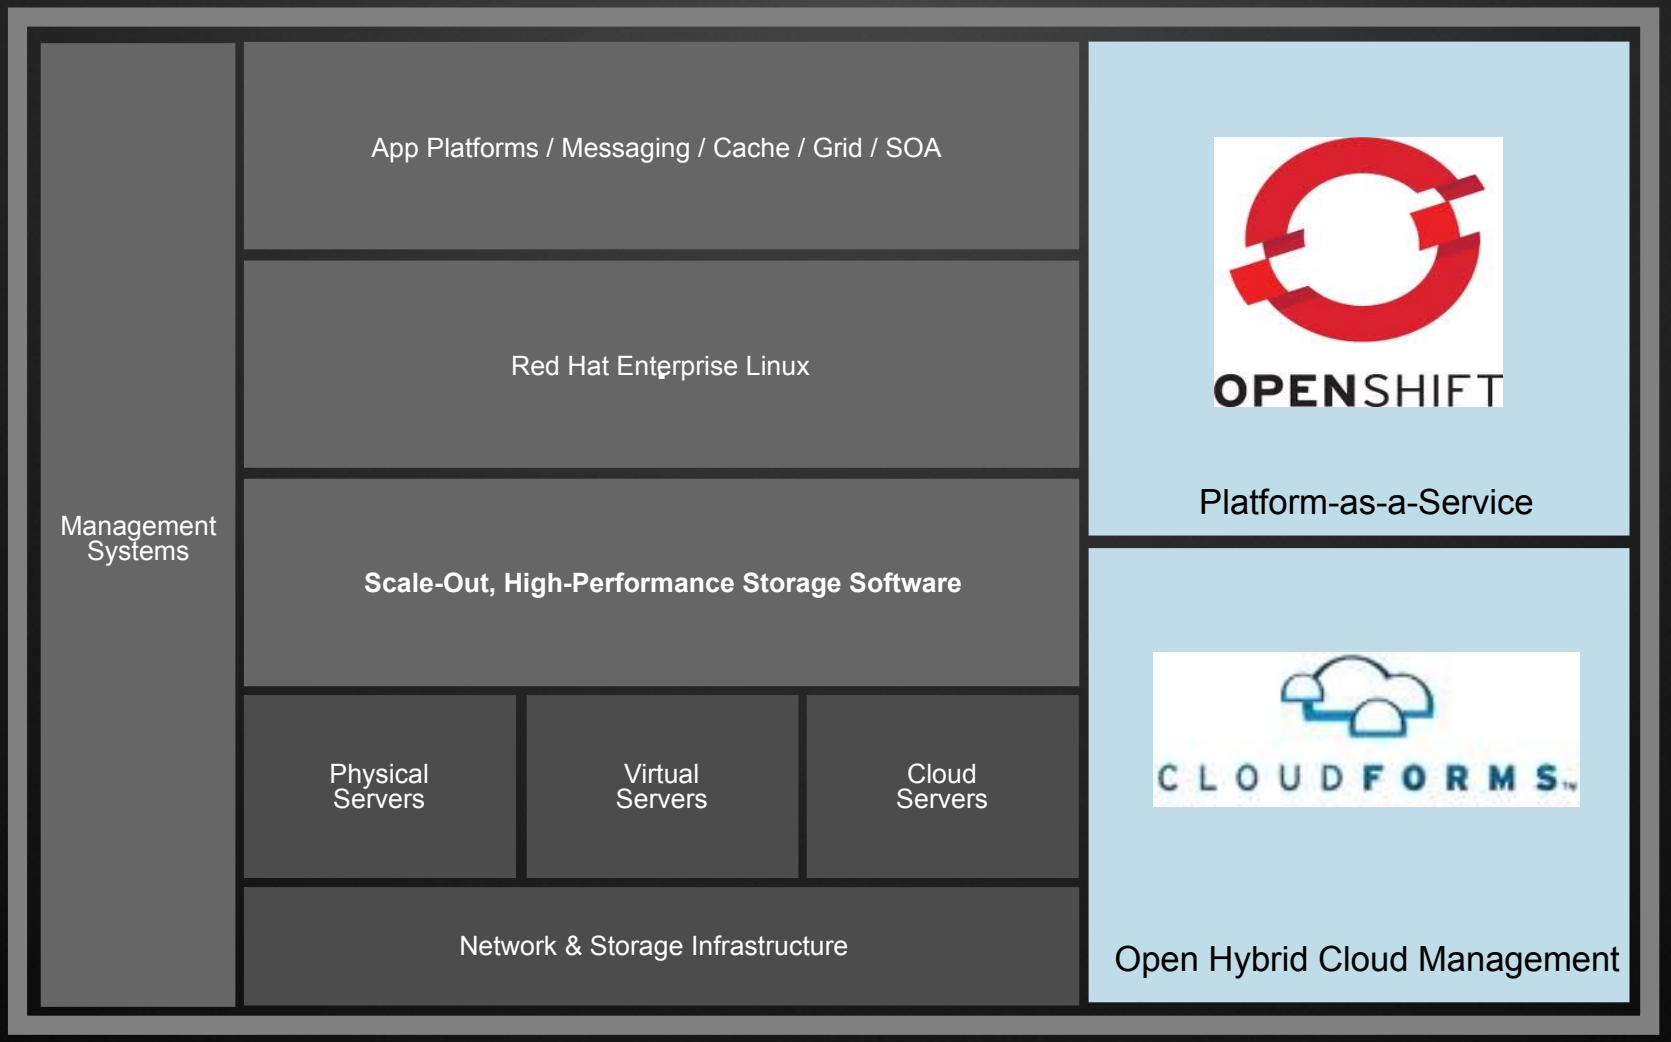

voddler

"The powerful combination of Red Hat Enterprise Linux & Red Hat Enterprise Virtualization made Red Hat the obvious solution to virtualize our key business-critical systems…

—ROGER KEMMLER, IT manager

Cloud Service Models

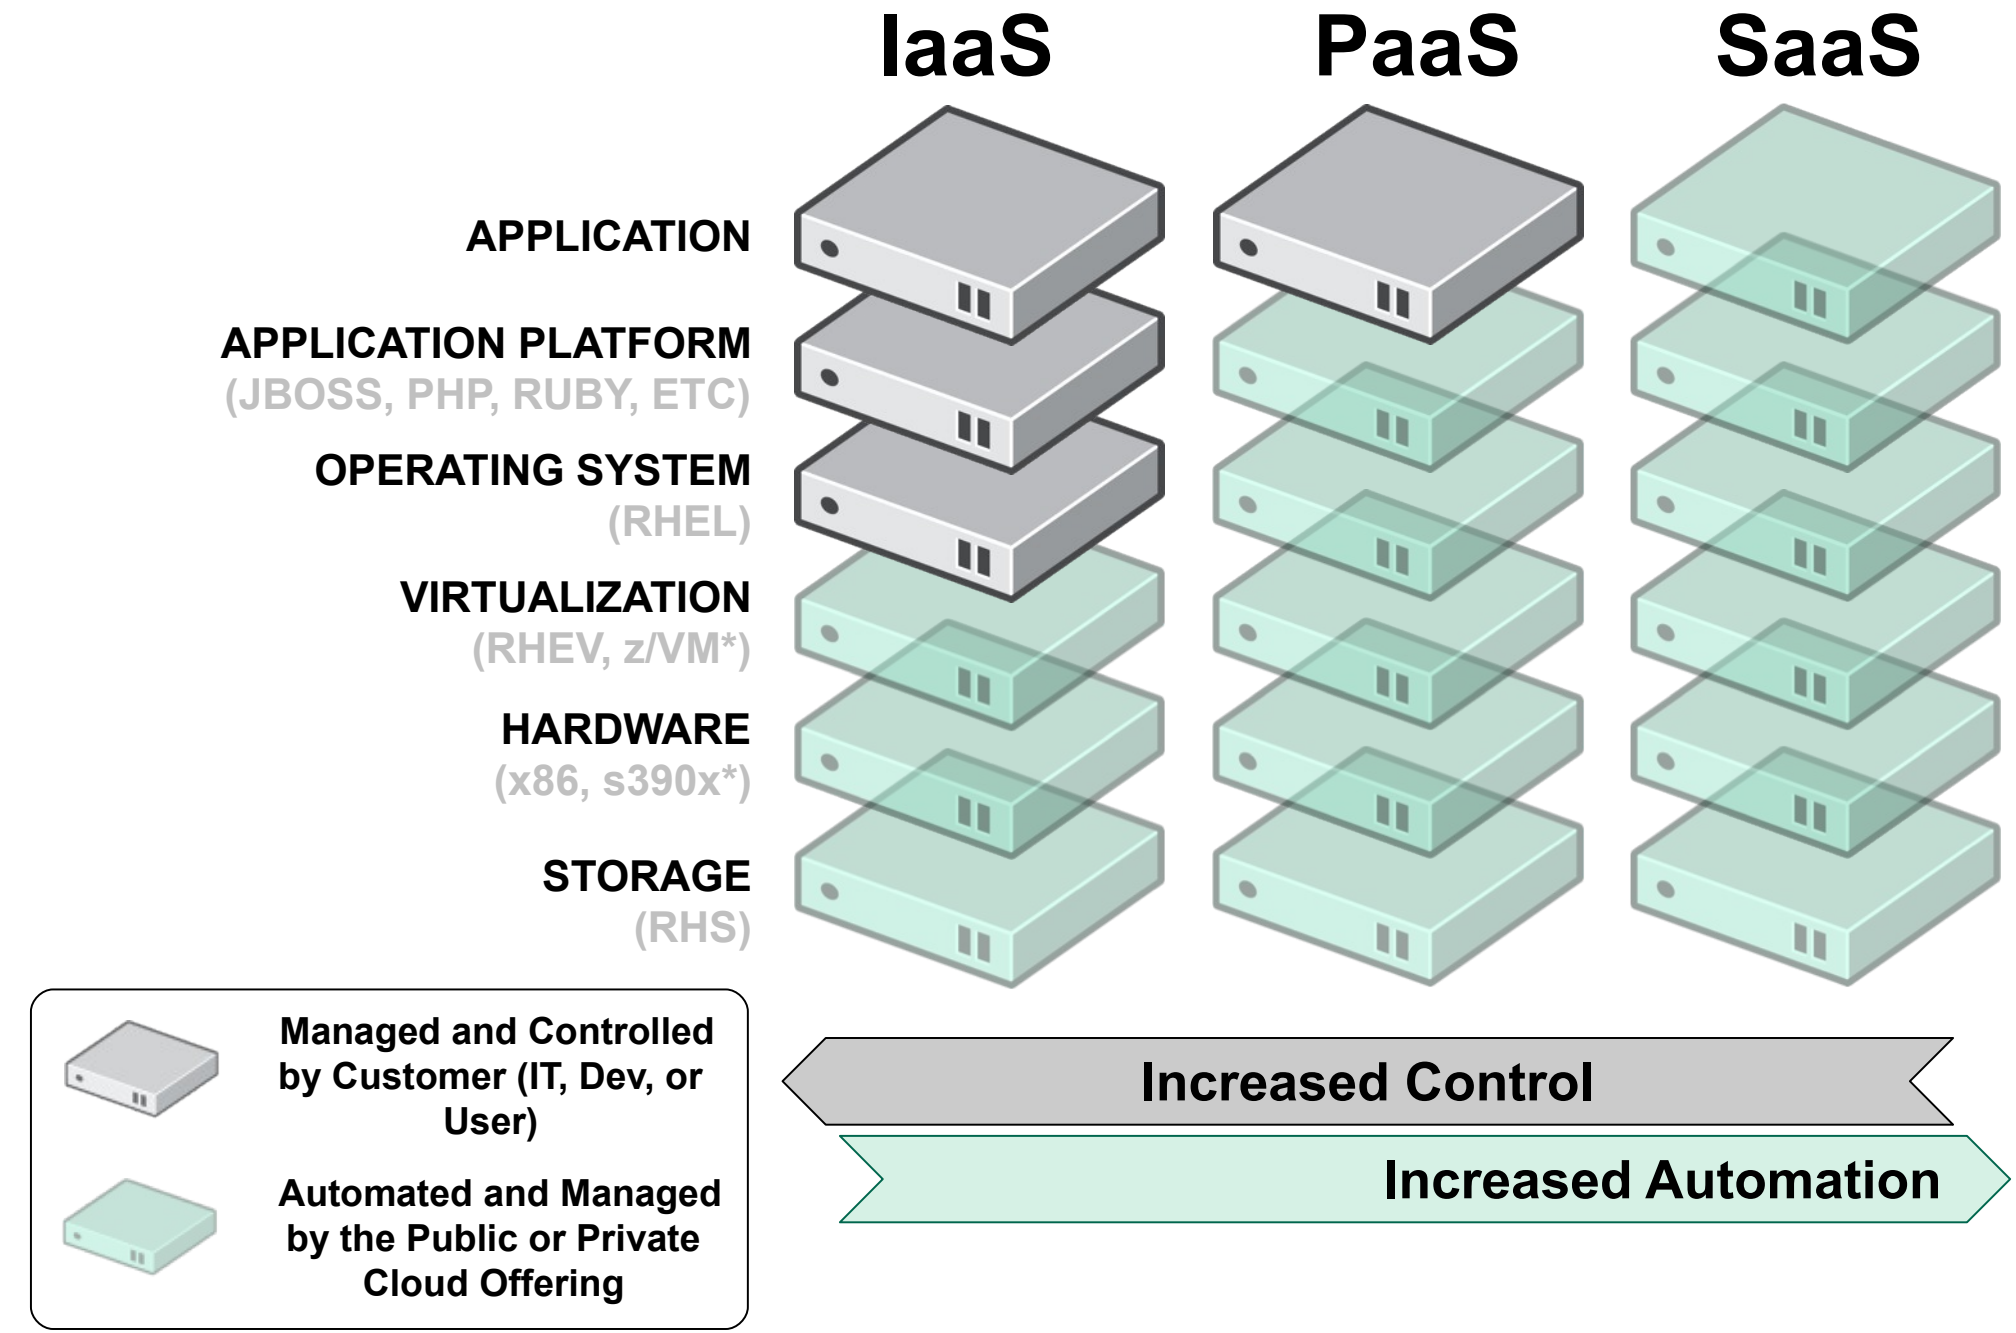

PaaS = Platform as a Service

A Cloud Application Platform

Code

Deploy

Enjoy

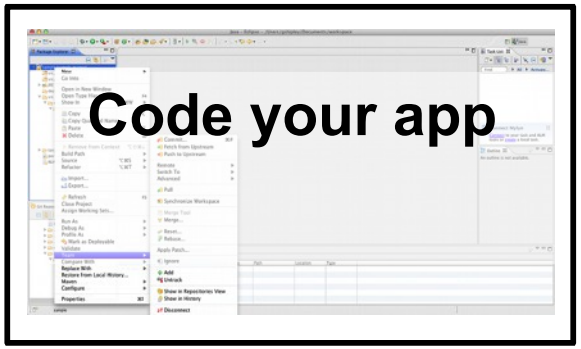

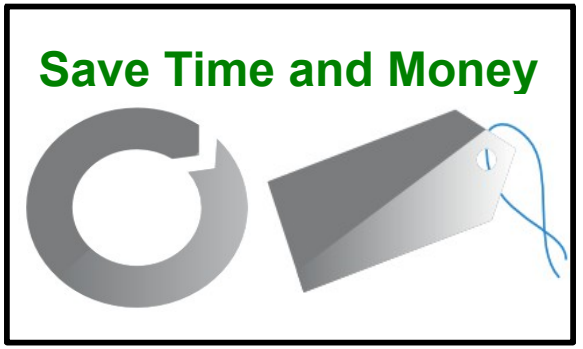

The Foundation of OpenShift is Red Hat Enterprise Linux

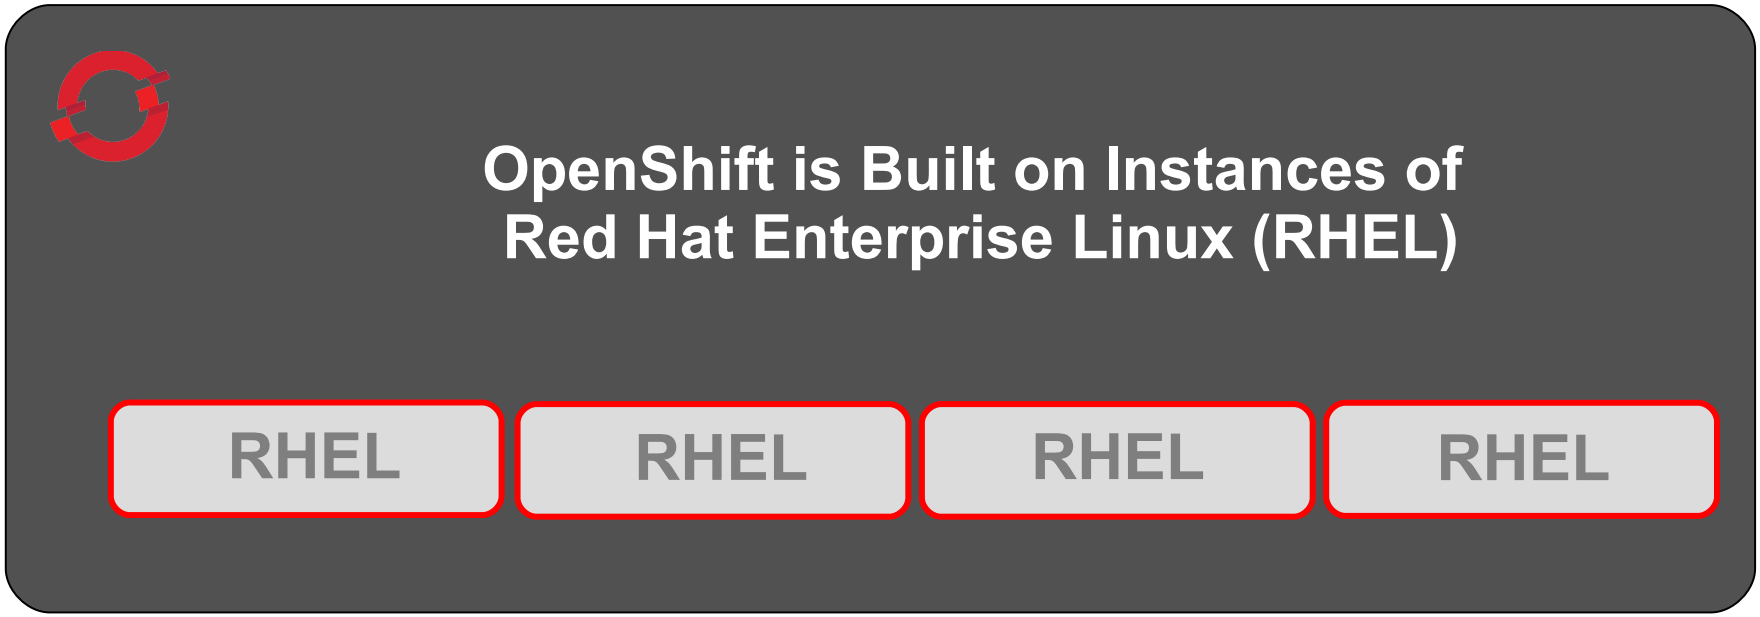

Unique SELinux Approach Enables Security and Multi-tenancy

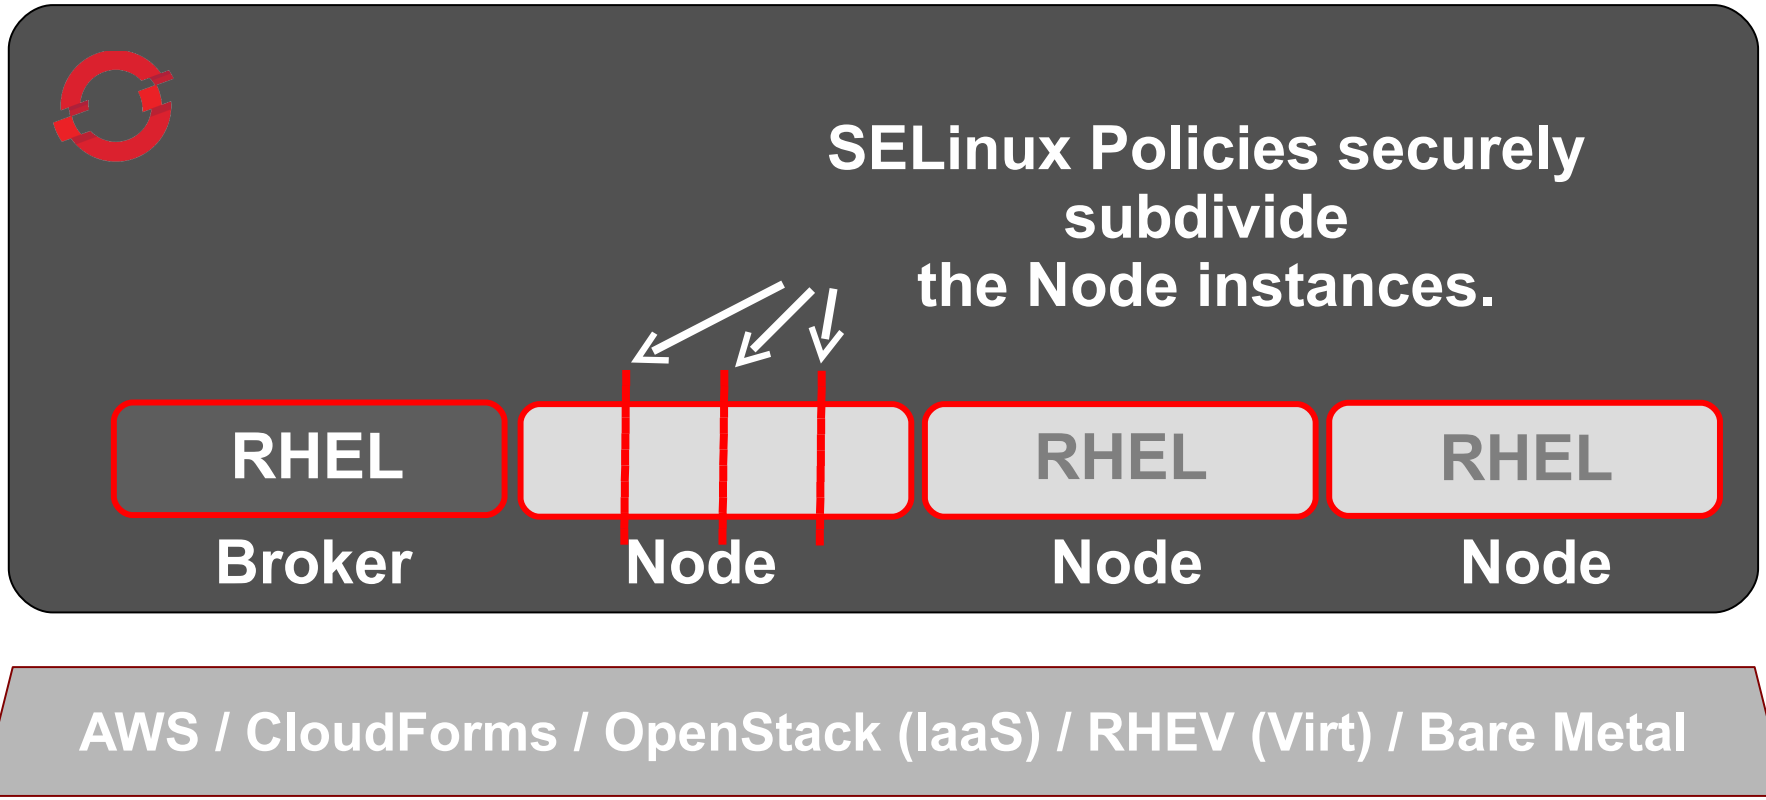

OpenShift User Applications Run in OpenShift Gears

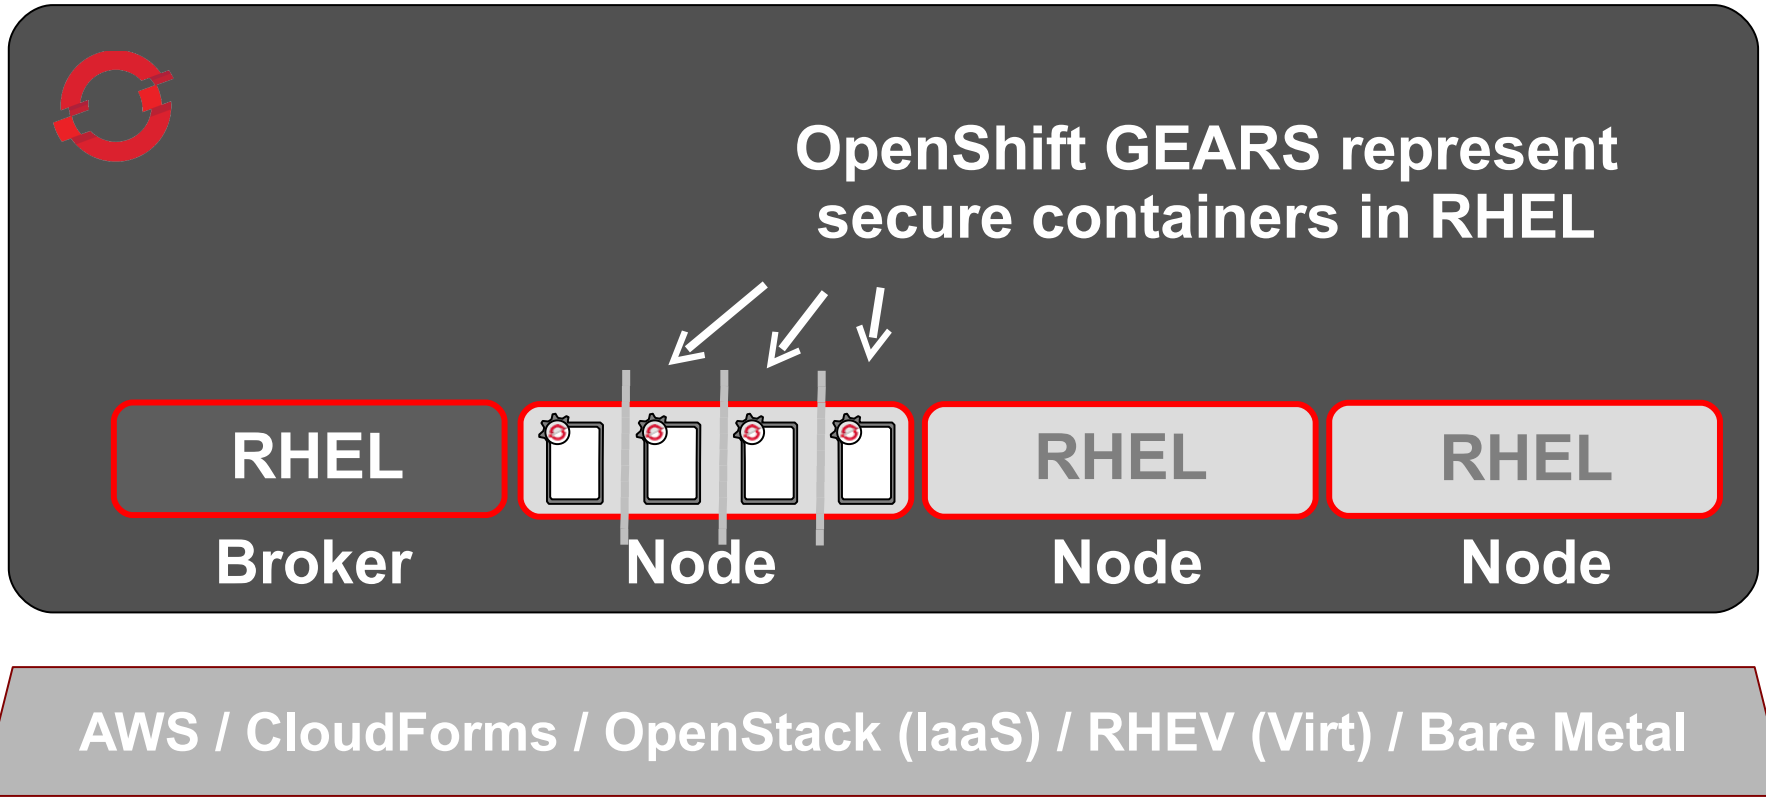

OpenShift Automates Gear Configuration via Cartridges

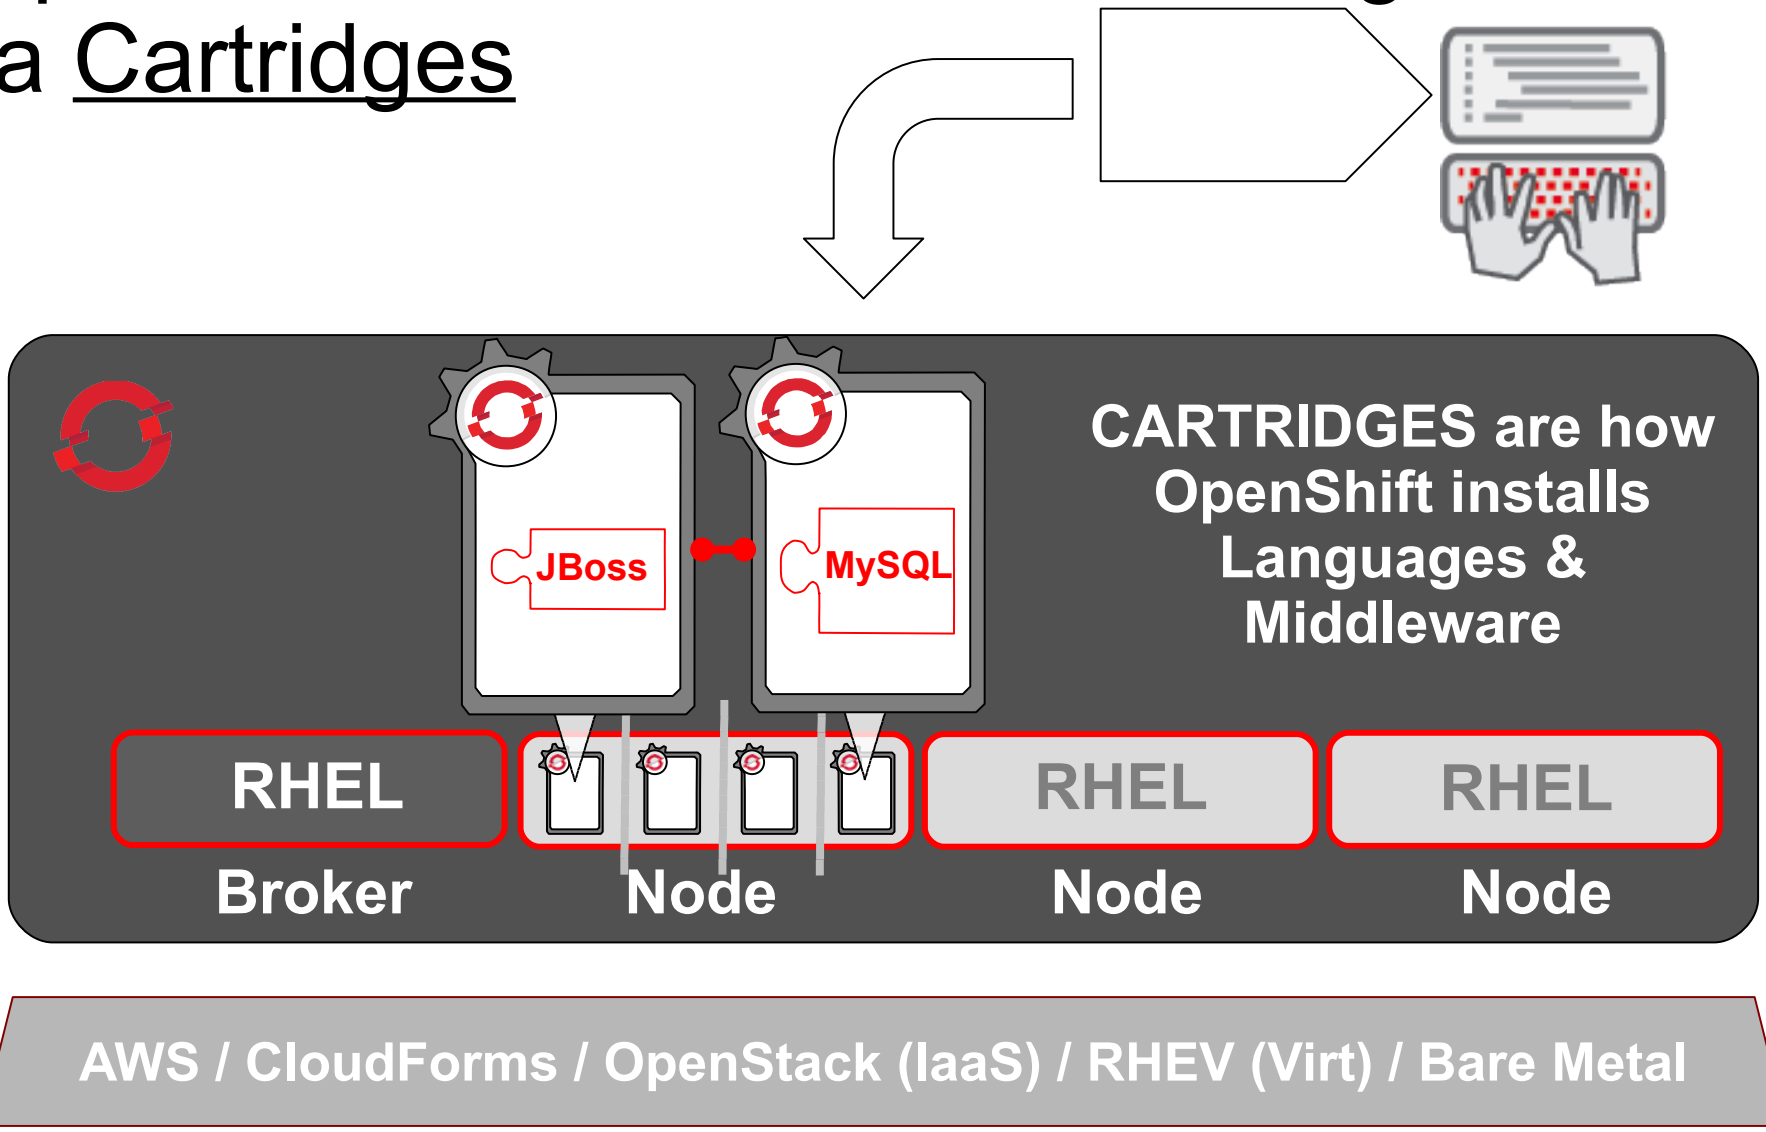

OpenShift Cartridge System Enables User-Built Cartridges

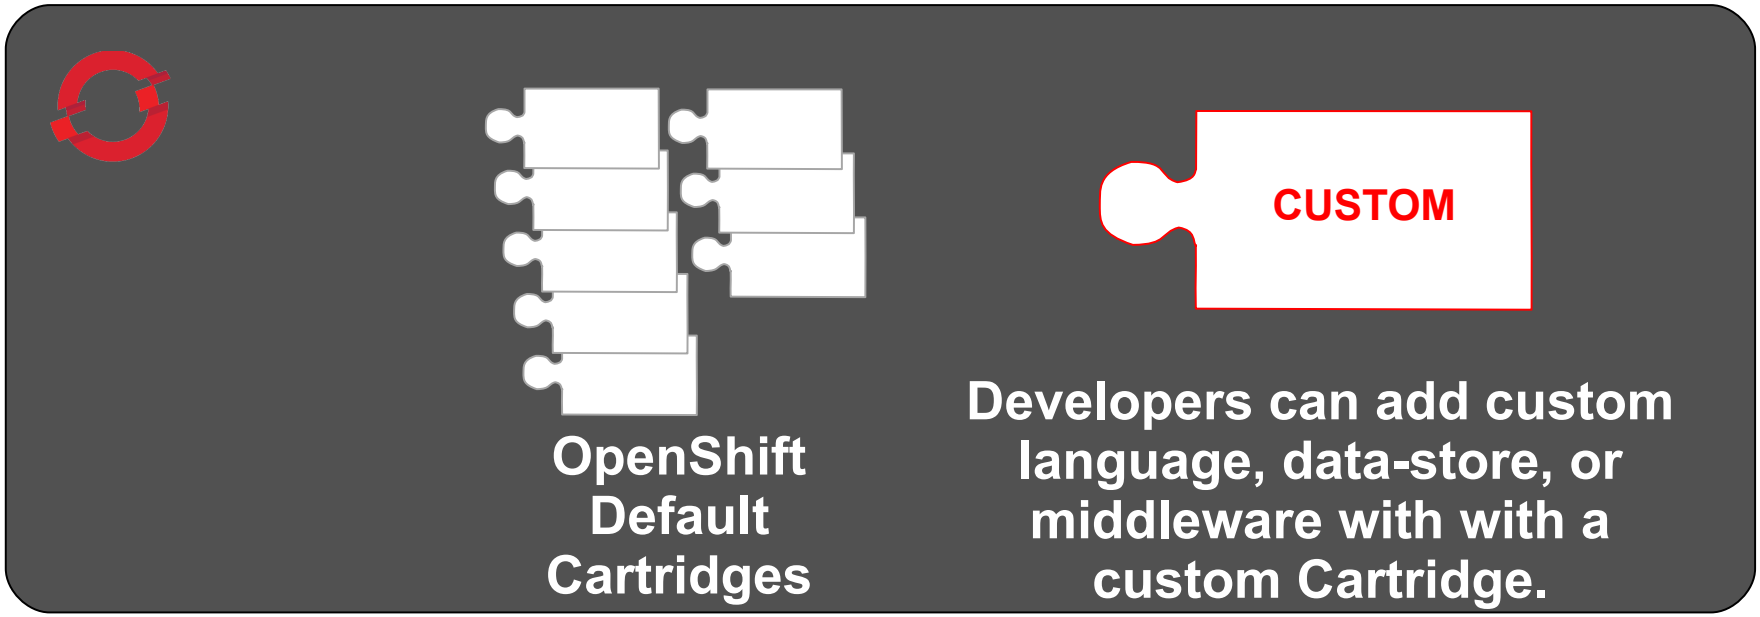

Why OpenShift?

1. **Strength. OpenShift is built on proven Red Hat technologies.**
2. **Freedom. In OpenShift, work the way you want.**
   - Choice of Interface: Web Console, Command-line, or IDE
   - Choice of Middleware: Java(EE6), Ruby, Node.js, PHP, Python, etc.
   - Choice of Cloud: Public, Private, or Hybrid Cloud
   - Choice of Elasticity: Automatic application scaling when needed
1. **Openness. OpenShift's open source software stack ensures application portability and No Lock-In.**

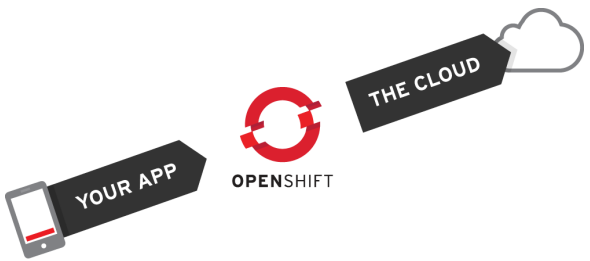

If you would like to learn more, visit :

http://www.redhat.com/rhel/server

http://www.redhat.com/cloud

Brad Hinson
bhinson@redhat.com
Solution Architect

Appendix: RHEL 6, RHEL 7 Features

RHEL 6 Enhancements

- **Logical Volume Manager (LVM)**
  - Snapshot Merge
    - lvconvert merge snapshot into original volume
    - Restore previous (point in time) states of LVM.

RHEL 6.3

RAID 4 5 & 6
Thin Provisioned Lvs
FcoE Target

- Filesystems & Storage
  - Ext4
    - Default filesystem for RHEL6
    - Faster & more scalable (vs. ext3)
    - Vastly improved fsck time (vs. ext3)
  - XFS & GFS: options for very large storage, and clustered environments

RHEL 6.3

GFS2 perf boost
O_Direct in FUSE

RHEL 6 Enhancements

- VM Scalability Improvements
  - Lockless page cache
    - Multiple CPUs can look up pages from the same file simultaneously
  - Split LRU VM
    - Split file backed, mem/swap backed & mlocked pages onto their own LRU lists
    - Better targeted page scanning
    - Different eviction policies for file backed & mem/swap backed pages
    - Pageout code scales to larger amounts of memory

RHEL 6 Enhancements

- Networking
  - Multi-queue transmit & multi-CPU receive for NUMA scalability
  - RCU SMP locking optimization across networking stack
  - SR-IOV enables a virtual server to saturate a 10GbE link
  - Virt - Raw socket mode – kernel net I/O avoids prior context switch

RHEL 6 Enhancements

- **Reliability Availability Serviceability (RAS)**
  - Advanced error recovery/reporting
  - CPU and memory hot add
  - Machine Check Architecture
    - Intelligent recovery from CPU/memory errors
    - Enhanced error reporting for PCI devices (PCI-AER & APEI)
- Rapid file system recovery (10x faster than RHEL 5)
  - E.G. Fsck for 1TB filesystem (45 million files)
  - RHEL5 Ext3 = 1 hour, RHEL6 Ext4 = 6 minutes.

RHEL 6 Enhancements

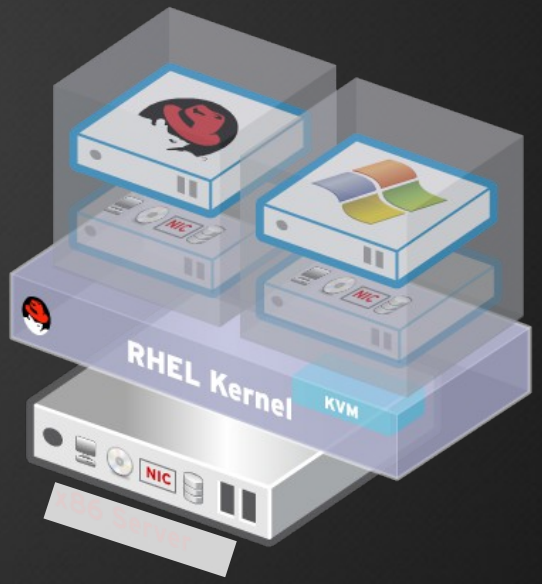

- **Virtualization**
  - **Performance:**
    - Commonly 85%-95% of bare metal, including I/O bound workloads
  - **Scalability:**
    - Host: 128/4096 cores; 2TB/64TB RAM
    - Guest: 64vCPU; 256 MB RAM
  - **Advanced capabilities:**
    - Live Migration; CPU/Mem resource control
    - Memory page sharing (KSM); SR-IOV; VT-D; SVirt security
  - **Hypervisor integrated into the Linux kernel**
  - All features accrue to Red Hat Enterprise Virtualization, which also provides sophisticated management capabilities

RHEL 6.3

Guest: 160 vCPU
Guest: 2TB RAM

RHEL 7 Debugging Mechanisms

- Perf and Oprofile Updates
- Hardware Error Reporting Mechanism (HERM) (RHEL7)
  - merge of mcelog and edac
  - Add support for reporting APEI events
- Kexec Kdump Supportability
  - Multipath and FcoE target device support (RHEL7)

Linux Containers (LXC)

- Linux Container
  - LXC builds up from chroot to implement complete virtual systems, adding resource management and isolation mechanisms to Linux's existing process management infrastructure.
    - Control Groups/Controller support
    - Name Spaces (NS) support
  -

RHEL 7 Virtualization

- **Scalability** – Biggest x86 guest
- **Performance** – KVM wins in all specVirt categories
- **RAS** – SLA, online resource provisioning, etc...
- **Maintenance** – Serve & Protect
- **Exceptional Features** – Same OS for host & guest
- **Enterprise, Cloud –** KVM addresses all scenarios

RHEL 7 Virtualization RAS

- **VCPU hotplug** (RHEL6.3 TP)
- **Memory hotplug** (RHEL7)
- **Live Snapshots** (RHEV3.1)
- **Live Block Migration** (RHEV3.1)
- **VM Power Management** (RHEL6.3 TP)
- **Direct LUN pass through** (RHEL6.4)
- **VPMU** (RHEL6.3 TP)
- More
  -

RHEL 7 Virtualization Coming

- **Non Uniform Memory Access Optimization**
  - Numad (RHEL6.3 TP)
  - AutoNuma / SchedNuma (RHEL 7)
    - Heavy review needed
    - Possible solution may be blended approach
    - Upstream acceptance pending
- **MultiQueue virtual NICs** (RHEL7)
- **Zero copy networking** (RHEL7)
- **Virtio-SCSI: new blok layer** (RHEL6.3 TP / RHEL7)
- **Handful of new paravirt optimization**
  -

Linux Containers (LCX) Roadmap

- RHEL6 TP due to many security concerns
- RHEL 7
  - Effort focused here, esp. on security and scalability
- Fedora 17 (Proof of concept)
  - User space: Libvirt integration with selinux/svirt
  - Kernel: task counter to prevent fork bomb
- Fedora 18
  - Focus on SELinux policies & security improvements
  - Audit messages location inside containers
  - Add extended attribute to cgroup fs
  - Namespace scalability testing