redhat.

RED HAT TECHNICAL SYMPOSIUM

NSA R&E Symposium Center
Monday November 8th, 2010
1300-1600

RED HAT® ENTERPRISE LINUX®

AGENDA

1:00-1:20	Software Central Opening Notes
1:20-2:20	Red Hat Enterprise Linux 6 Update
2:20-2:30	Break
2:30-3:30	Red Hat in the Virtualized Environment & Security
3:30-4:00	Q&A Panel w/ Red Hat Technologists

RED HAT ENTERPRISE LINUX 6 UPDATE

Shawn D. Wells
Technical Director, Intelligence Programs
sdw@redhat.com / 443-534-0130

RED HAT ENTERPRISE LINUX 6 UPDATE
RHEL6 FOUNDATION FEATURES & THEMES

- **Trusted data center platform**

  Ideally positioned to provide non-disruptive path forward.

- **Application and infrastructure performance, scalability and security**

  Support for varied workloads and advanced hardware capabilities.

- **Making IT agile across physical, virtual and Cloud**

  The right operating system across any environment.

RED HAT ENTERPRISE LINUX 6 UPDATE

RHEL6 FOUNDATION FEATURES & THEMES

- **Improved manageability**
  For large scale virtualization deployments, server & desktop. Samba enhancements for Windows active directory and file sharing

- **Power management**
  Efficiency for lower deployment costs, reduced carbon footprint for virt, bare metal, desktop. Hardware level as well as dynamic system service startup and suspend.

- **RAS (Reliability, Availability, Serviceability)**
  Hotplug, memory error reporting, filesystem and data integrity. Support tools such as automated crash detection and bug reporting.

RED HAT ENTERPRISE LINUX 6 UPDATE
RHEL6 FOUNDATION FEATURES & THEMES

- **Hardware enablement and scalability**
  Maximum efficiency, large configurations (cpu, memory, busses, I/O), NUMA awareness, new BIOS boot loader interface

  Supported architectures: x86, x86_64, PPC64, s390x

RED HAT ENTERPRISE LINUX 6 UPDATE
SYSTEMS MANAGEMENT INTEGRATION

- ***Red Hat Network Satellite***
  - Install & provision new systems
  - Update existing
  - Manage configuration files & maintain over time
  - Monitoring system metrics
  - Multiple managed Satellites on ELA

RED HAT ENTERPRISE LINUX 6 UPDATE
SYSTEMS MANAGEMENT INTEGRATION

- ***Red Hat Network Satellite 5.4***
  - RPMs will be now be encrypted with 256-bit keys, up from 128-bit.
  - Find & remove stale instances that are no longer being used
  - Flex Guest / Floating Entitlements
  - Centrally manage SELinux Context on remote files
  - Expanded APIs for 3rd party integration

RED HAT ENTERPRISE LINUX 6 UPDATE
SYSTEMS MANAGEMENT INTEGRATION

- ***Satellite Deployment Model***

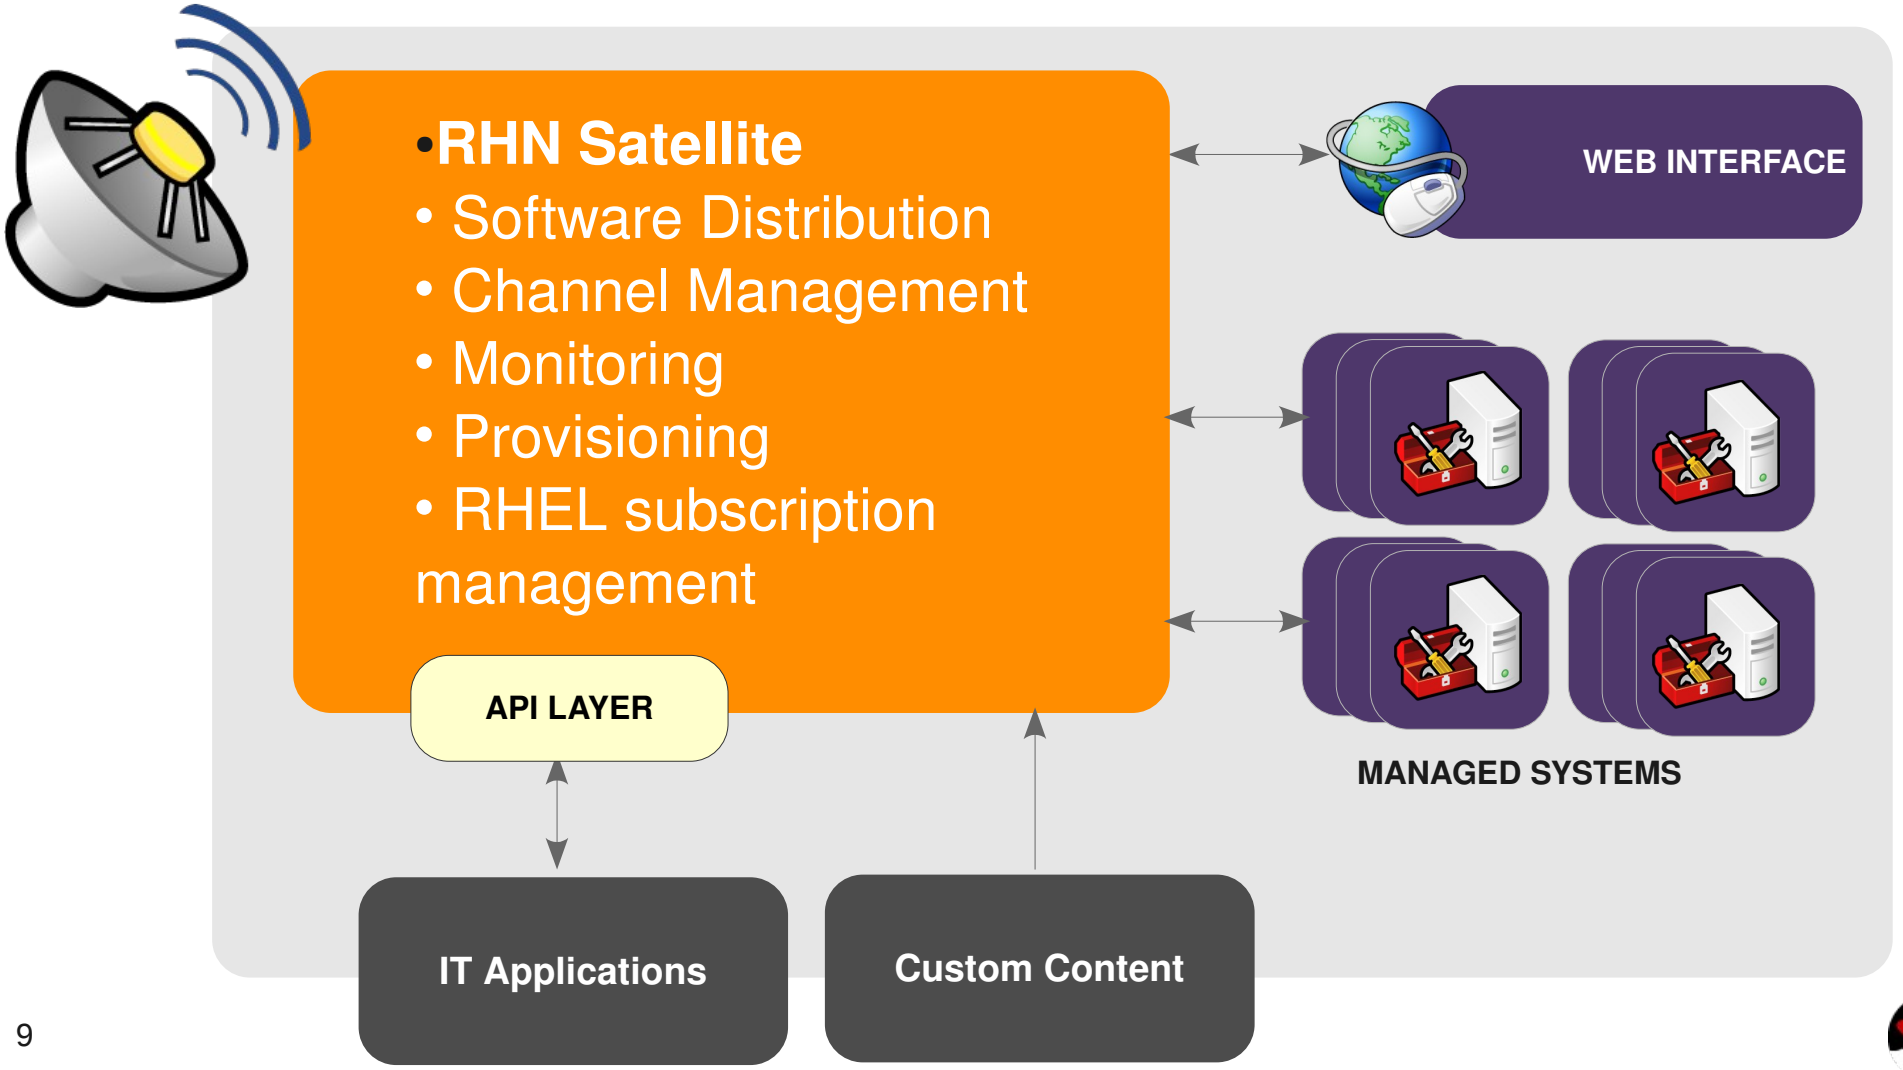

RED HAT ENTERPRISE LINUX 6 UPDATE
SYSTEMS MANAGEMENT INTEGRATION

- ***Satellite Deployment Model: Single Satellite***

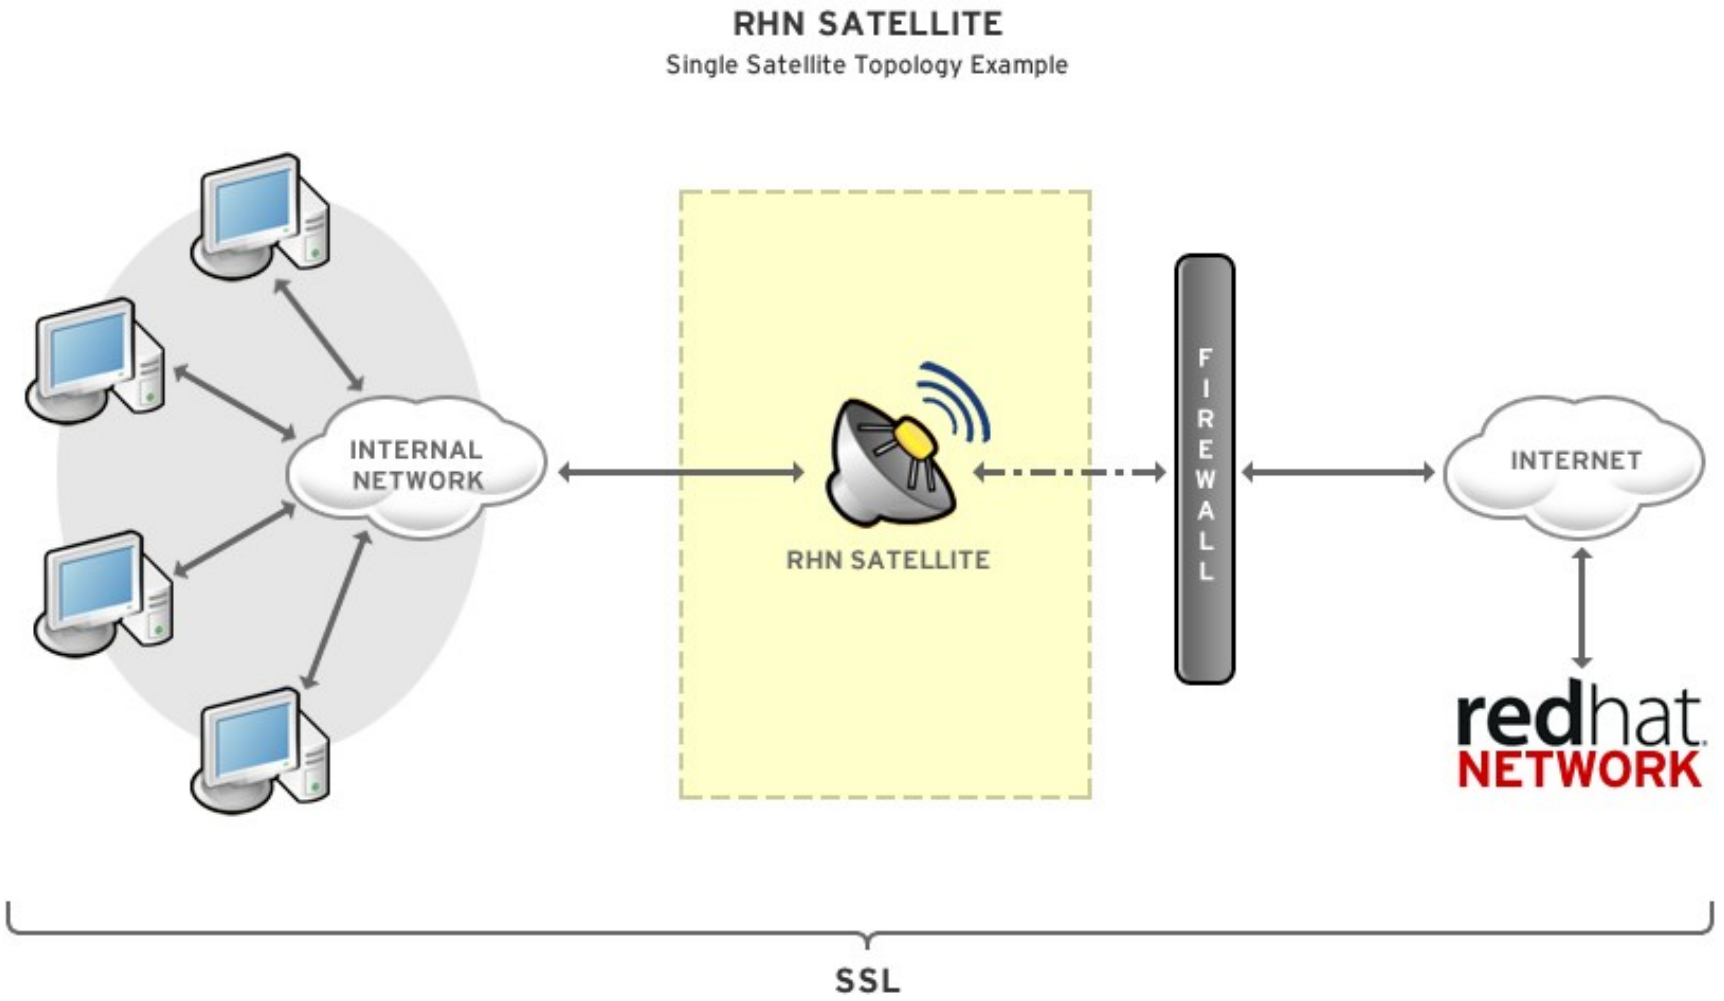

RED HAT ENTERPRISE LINUX 6 UPDATE
SYSTEMS MANAGEMENT INTEGRATION

- ***Satellite Deployment Model: Multi Tiered Satellite***

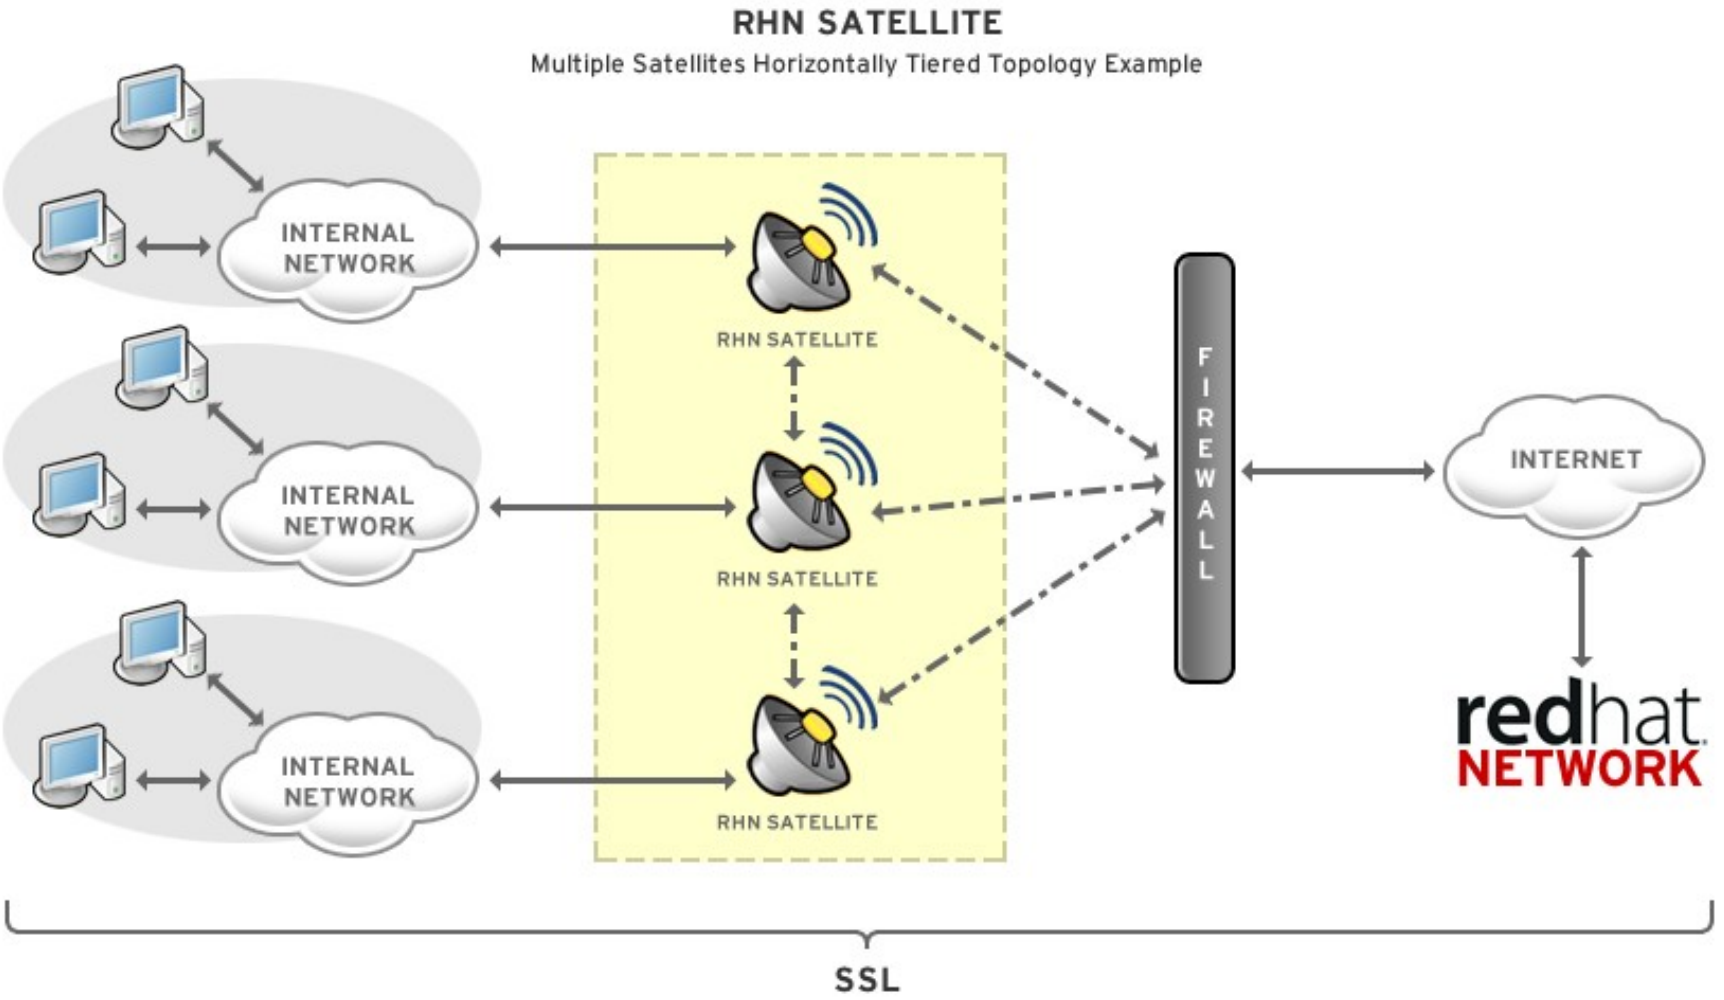

RED HAT ENTERPRISE LINUX 6 UPDATE

EFFICIENCY, SCALABILITY, RELIABILITY

Red Hat Enterprise Linux 6 evolves in concert with hardware advances, reducing system power consumption, taking advantage of hardware with greater numbers of processing and memory resources, and withstanding hardware failures better.

RED HAT ENTERPRISE LINUX 6 UPDATE
EFFICIENCY, SCALABILITY, RELIABILITY

- New scheduler (Completely Fair Scheduler)
- Dynamic addition of processor and memory
- More robust error reporting for PCIe devices (PCIe AER)
- Isolation of memory hardware failures with minimal downtime

RED HAT ENTERPRISE LINUX 6 UPDATE
EFFICIENCY, SCALABILITY, RELIABILITY

- Mature file systems to cater to varied usage and performance characteristics.
  - Ext4: Now default, scales to 16TB
  - GFS2: Scales to 25TB
  - XFS: Scales up to 100TB, tuned for storage arrays
  - NFSv4
- Configured to reduce chances of data corruption for low-end locally attached storage.

RED HAT ENTERPRISE LINUX 6 UPDATE
EFFICIENCY, SCALABILITY, RELIABILITY

- Power management
  - Tickless kernel
  - User-space tools – powertop
  - Dynamic throttling of power to devices
  - Relatime drive optimization

RED HAT ENTERPRISE LINUX 6 UPDATE
EFFICIENCY, SCALABILITY, RELIABILITY

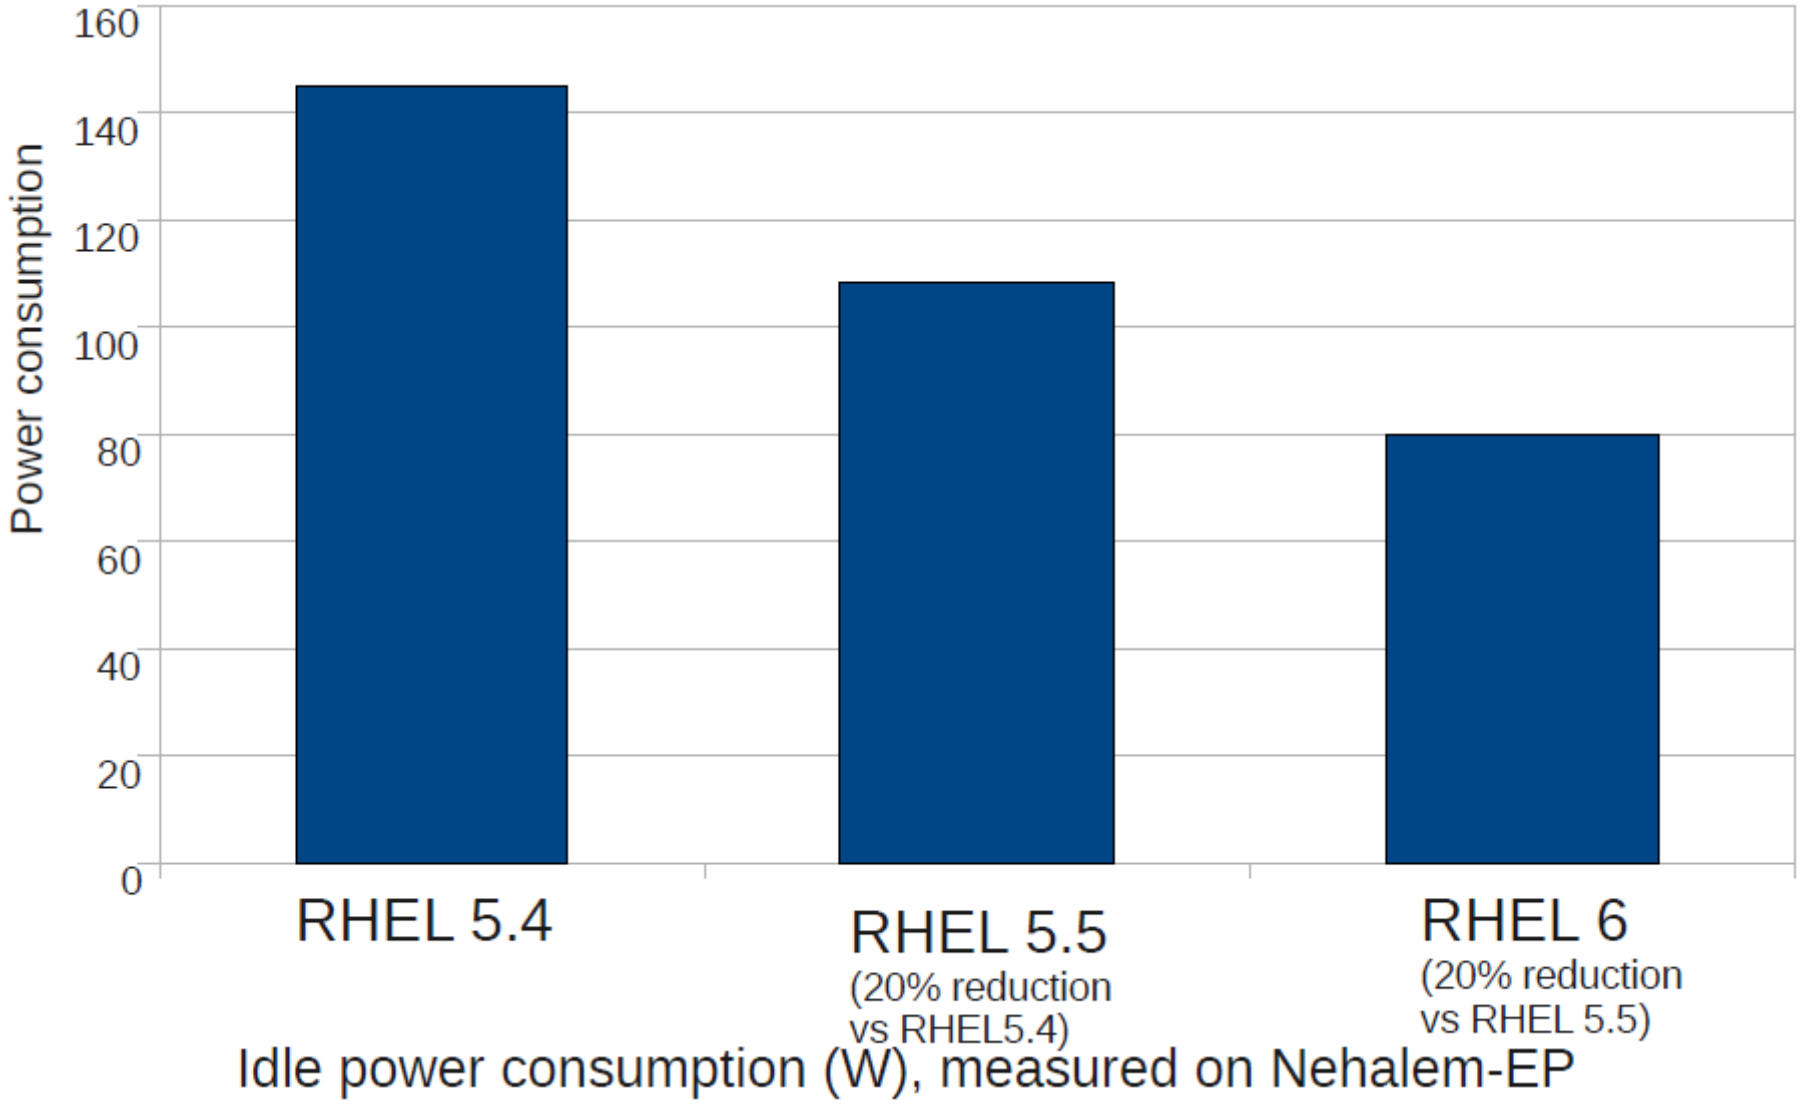

RED HAT ENTERPRISE LINUX 6 UPDATE
EFFICIENCY, SCALABILITY, RELIABILITY

```
     PowerTOP version 1.11      (C) 2007 Intel Corporation

Cn                Avg residency          P-states (frequencies)
C0 (cpu running)        (24.7%)            1.67 Ghz     9.6%
C0                0.0ms ( 0.0%)            1333 Mhz     1.0%
C1 mwait          0.0ms ( 0.0%)            1000 Mhz    89.4%
C2 mwait          1.1ms (75.3%)


Wakeups-from-idle per second : 689.3    interval: 10.0s
no ACPI power usage estimate available

Top causes for wakeups:
  55.4% (678.6)     <kernel core> : hrtimer_start_range_ns (tick_sched_timer)
  34.0% (416.0)       <interrupt> : PS/2 keyboard/mouse/touchpad
   3.0% ( 37.1)       <interrupt> : uhci_hcd:usb5, i915@pci:0000:00:02.0
   1.3% ( 16.3)      simpress.bin : hrtimer_start_range_ns (hrtimer_wakeup)
   1.0% ( 12.7)     <kernel core> : hrtimer_start (tick_sched_timer)
   0.9% ( 10.5)   thunderbird-bin : hrtimer_start_range_ns (hrtimer_wakeup)

Suggestion: Enable USB autosuspend by pressing the U key or adding
usbcore.autosuspend=1 to the kernel command line in the grub config

 Q - Quit    R - Refresh    U - Enable USB suspend
```


RED HAT ENTERPRISE LINUX 6 UPDATE
EFFICIENCY, SCALABILITY, RELIABILITY

Objective: Providing scaling headroom anticipating many years of upcoming hardware generations. Tested and supported limits will likely grow over the course of product lifespan.

RED HAT ENTERPRISE LINUX 6 UPDATE
EFICIENCY, SCALABILITY AND RELIABILITY

- ***Kernel scalability limits - x86_64***

Parameter	RHEL5 Supported Limit	RHEL6 Supported Limit	RHEL6 *Theoretical Limit*
CPUs	64	128	4096
Memory – Physical addressing	1TB	2TB	64TB
Memory – Process virtual address space (note – hardware dependent	2TB	128TB	128TB
IRQs	293	33,024	33,024
# of processes	32,000	32,000 (larger pending testing)	4 million
KVM guest memory	512GB	8TB	64TB
KVM guest CPU	32	64 (pending testing)	64 (pending testing)

RED HAT ENTERPRISE LINUX 6 UPDATE
EFICIENCY, SCALABILITY AND RELIABILITY

- ***File systems and storage limits - x86_64***

	Maximum filesize	Maximum filesystem size
EXT3	2TB	16TB
EXT4	16TB	16TB
GFS	16TB	16TB
GFS2	25TB	25TB
XFS	100TB	100TB

RED HAT ENTERPRISE LINUX 6 UPDATE
DETERMINISM & REALTIME ENHANCEMENTS

- Some capabilities from Red Hat in MRG-realtime kernel (currently shipping as layered product) mainstreamed in RHEL6.
  - **Determinism –** Ability to schedule priority tasks predictably and consistently
  - **Priority** – Ensure highest priority applications are not blocked by lower priorities
  - **Timer** – Microsecond precision not timer interrupt, ~millisecond precision

RED HAT ENTERPRISE LINUX 6 UPDATE
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. Interfaces standardized to obtain alignment and optimal I/O stripe width.
- **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
- **NPIV –** n_Port ID Virtualization

RED HAT ENTERPRISE LINUX 6 UPDATE
STORAGE MANAGEMENT: *LVM/MD*

- **LVM hot spare;** a disk or group of disks used to replace one failing
- **Online resize** or 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**
  - Mirrored mirror log, avoids need for re-sync after failure
- **Selectable hash algorithm** for LUKS header, new cryptsetup commands, new libcryptsetup

SOFTWARE RELEASE VERSIONING & SUBSCRIPTIONS

- **Red Hat Enterprise Linux Advanced Platform (RHEL AP)**
  - > 2 CPUs
  - Includes Global File System (GFS), Red Hat Cluster Suite (RHCS)
- **Red Hat Enterprise Linux Enterprise Server (RHEL ES)**
  - Less than 2 CPUs
  - Does not include GFS & RHCS

All servers, bare metal or virtual, must have an **active** Red Hat subscription.

SOFTWARE RELEASE VERSIONING & SUBSCRIPTIONS

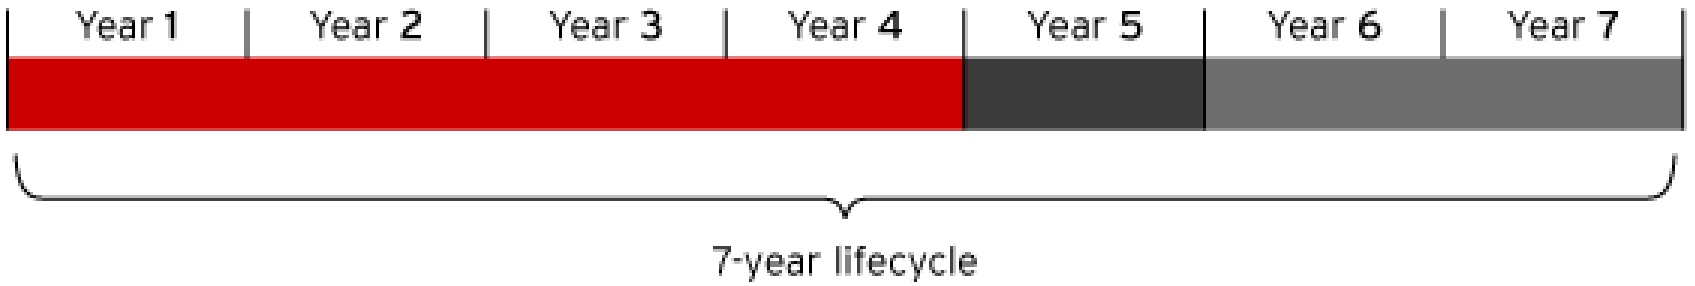

- **Production 1: Ongoing, active development of features and hardware enablement for inclusion into RHEL.**

- **Production 2: New software functionality is not available during this phase. The focus for minor releases during this life cycle phase lies on resolving defects with a minimum priority of high.**

- **Production 3: No new functionality, new hardware enablement or updated installation images are planned. Commonly known as "Maintenance Mode."**

SOFTWARE RELEASE VERSIONING & SUBSCRIPTIONS

- **General Availability:** March 14, 2007
- **End of Production 1 phase:** Q4 of 2011
- **End of Production 2 phase:** Q4 of 2012
- **End of Production 3 phase:** March 31, 2014
- **End of Extended Life Cycle phase:** March 31, 2017

Latest information available at:
http://www.redhat.com/security/updates/errata/

CONSUMING NSA ENTERPRISE SUBSCRIPTIONS
HOW & WHERE TO GET THE SOFTWARE

- Two on-site badged technical consultants at NBP1
  - Jeff Weatherford: jweatherford@redhat.com
    (u) 240-373-0842
  - Joe Glenn jglenn@redhat.com
    (u) 410-854-3104

Highside contact information on SearchLight or “go redhat”

RED HAT IN THE VIRTUALIZED ENVIRONMENT

Chris Runge
Technical Director, U.S
crunge@redhat.com / 703-748-2202

RED HAT VIRTUALIZATION
EVOLUTION OF x86 VIRTUALIZATION

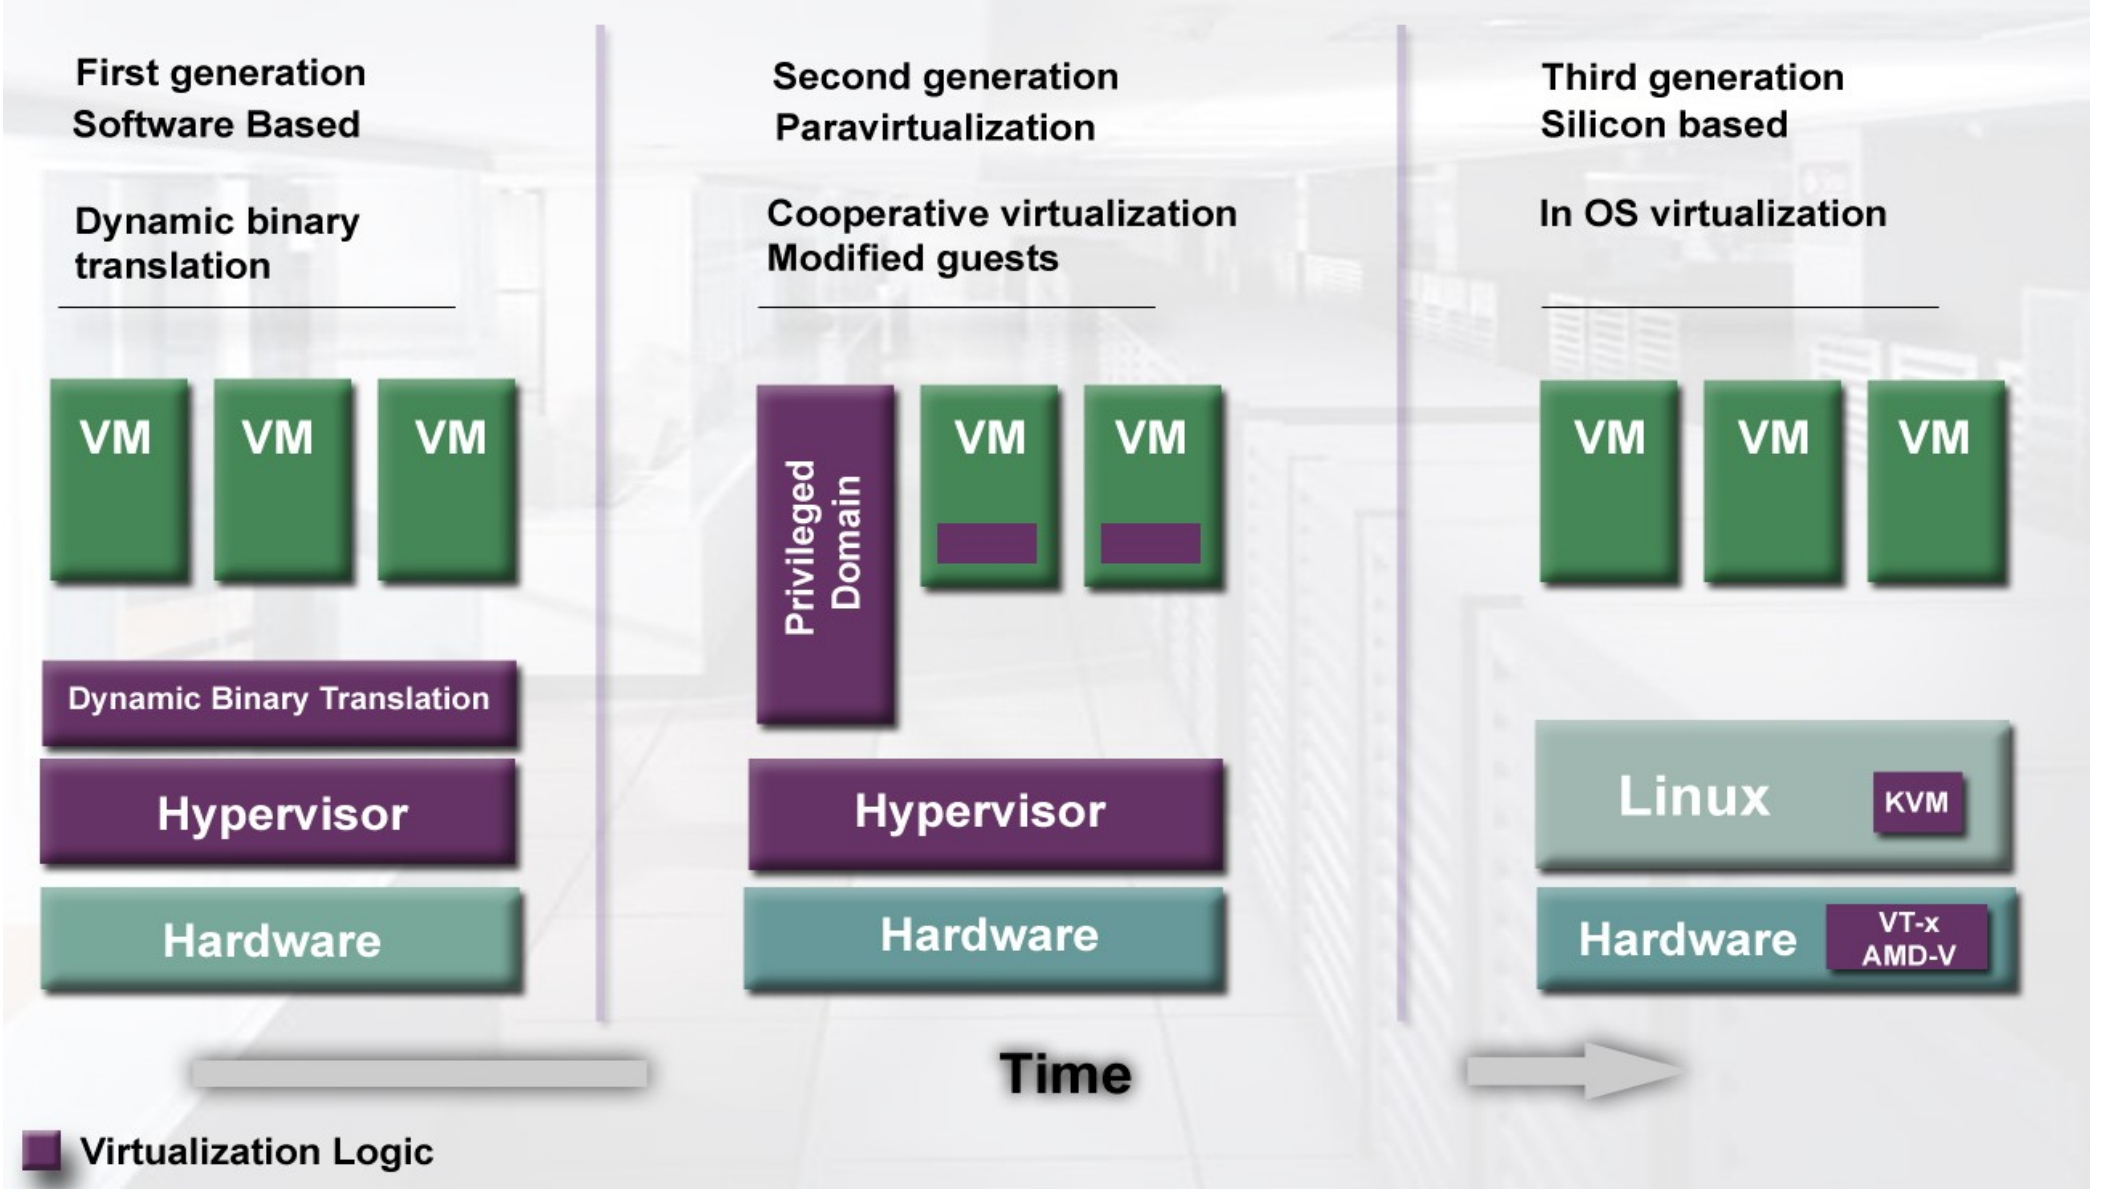

RED HAT IN THE VIRTUALIZED ENVIRONMENT
KERNEL-BASED VIRTUAL MACHINE (KVM)

Included in Linux kernel since 2006

Added to RHEL 5.4 and included in RHEL 6

- Xen supported in RHEL 5 through 2014

Available on x86_64 architecture

Requires Intel VT-X or AMD-V CPU capabilities

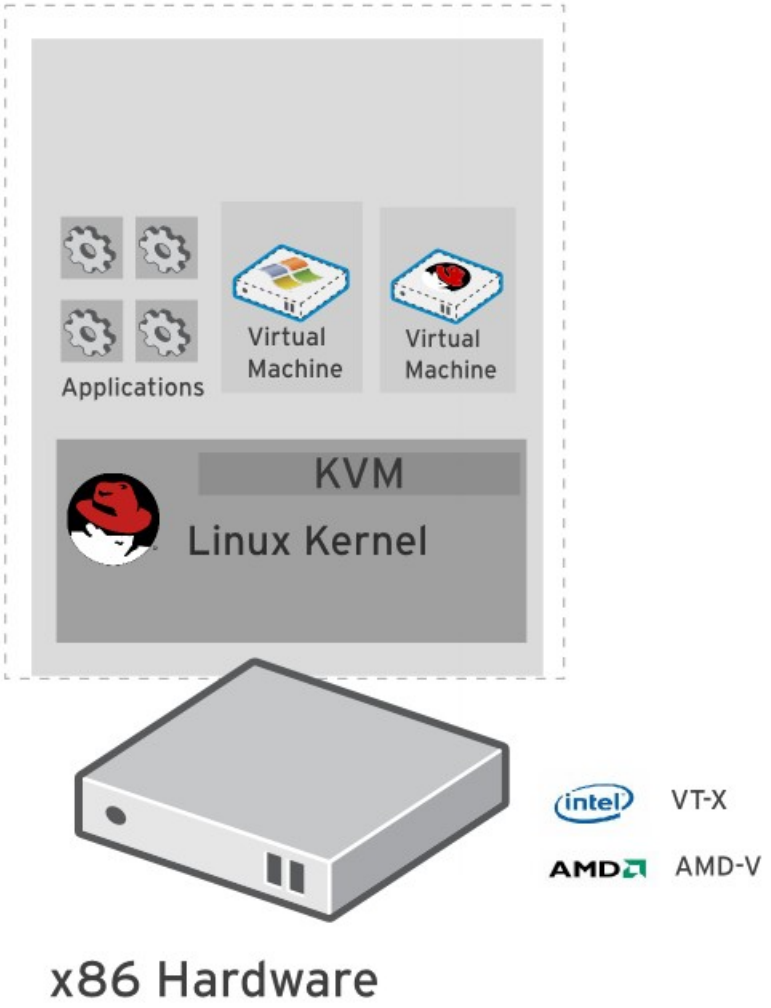

RED HAT IN THE VIRTUALIZED ENVIRONMENT
KVM GUEST SUPPORT

Runs Linux, Windows and other operating system guests

RHEL guests supported on third-party hypervisors:

- Microsoft Hyper-V
- VMWare

Microsoft certified drivers ensure compliance and support (WHQL and SVVP)

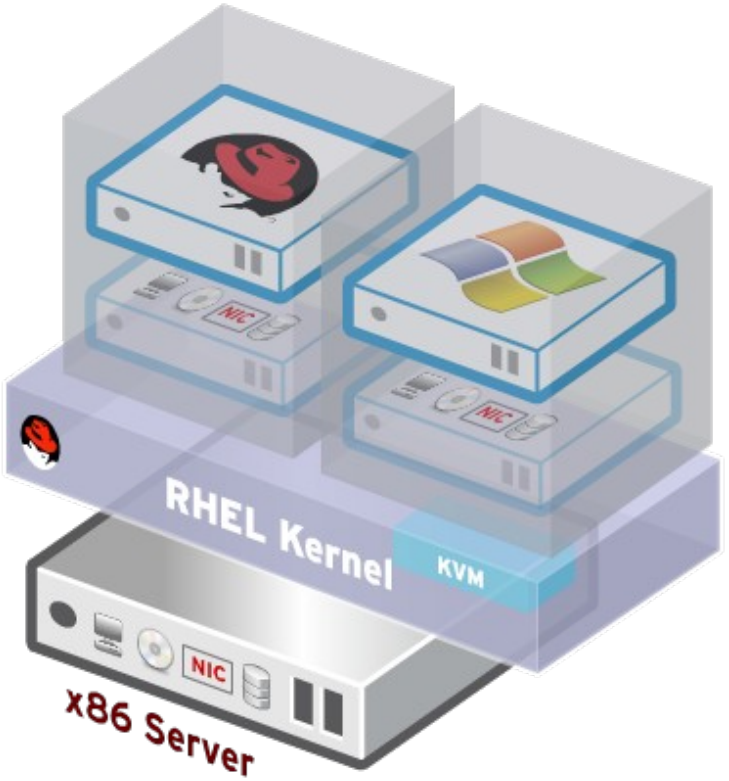

RED HAT IN THE VIRTUALIZED ENVIRONMENT
KVM ADVANTAGES

The OS is the hypervisor

Same platform for bare-metal, virtualization, and cloud

KVM is a Linux kernel module

VM's run as Linux processes

Simplifies certification

KVM architecture provides high “feature-velocity”

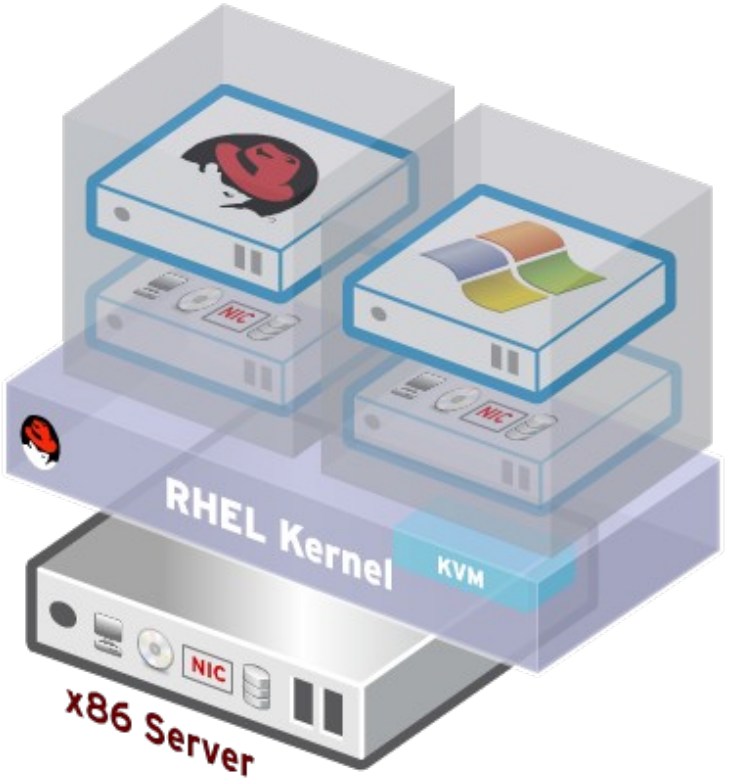

RED HAT IN THE VIRTUALIZED ENVIRONMENT
KVM FEATURE-VELOCITY: SCALABILITY

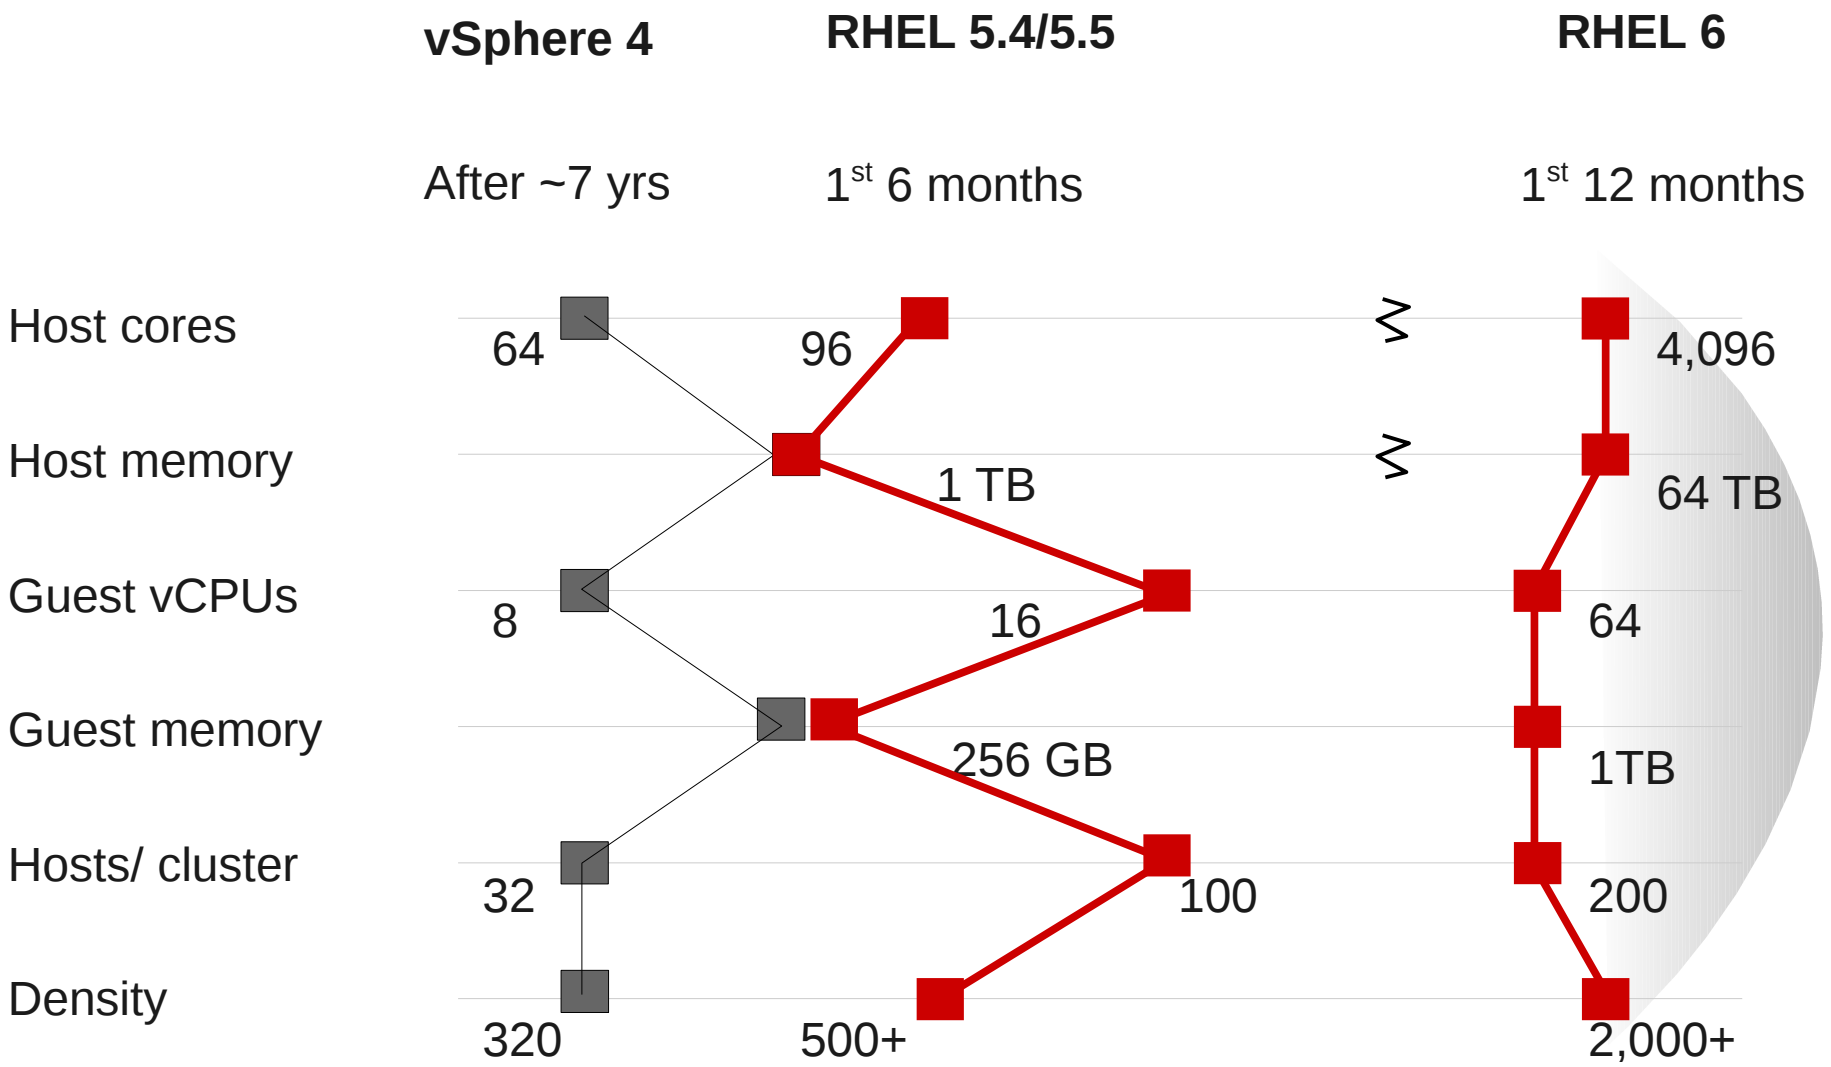

RED HAT IN THE VIRTUALIZED ENVIRONMENT
KVM ADVANCED FEATURES

Kernel Same-Page Merging (KSM)

Memory Page Sharing

Securely shares identical memory pages between virtual machines

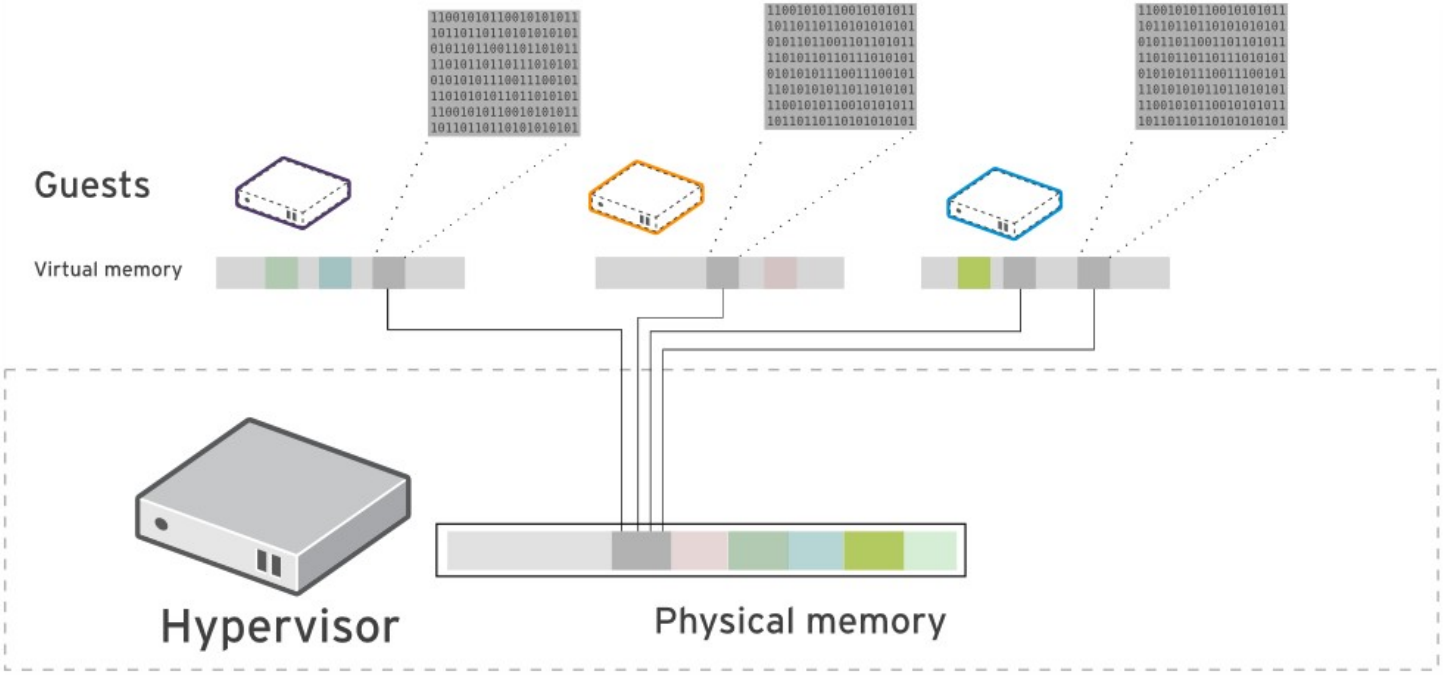

RED HAT IN THE VIRTUALIZED ENVIRONMENT
KVM ADVANCED FEATURES

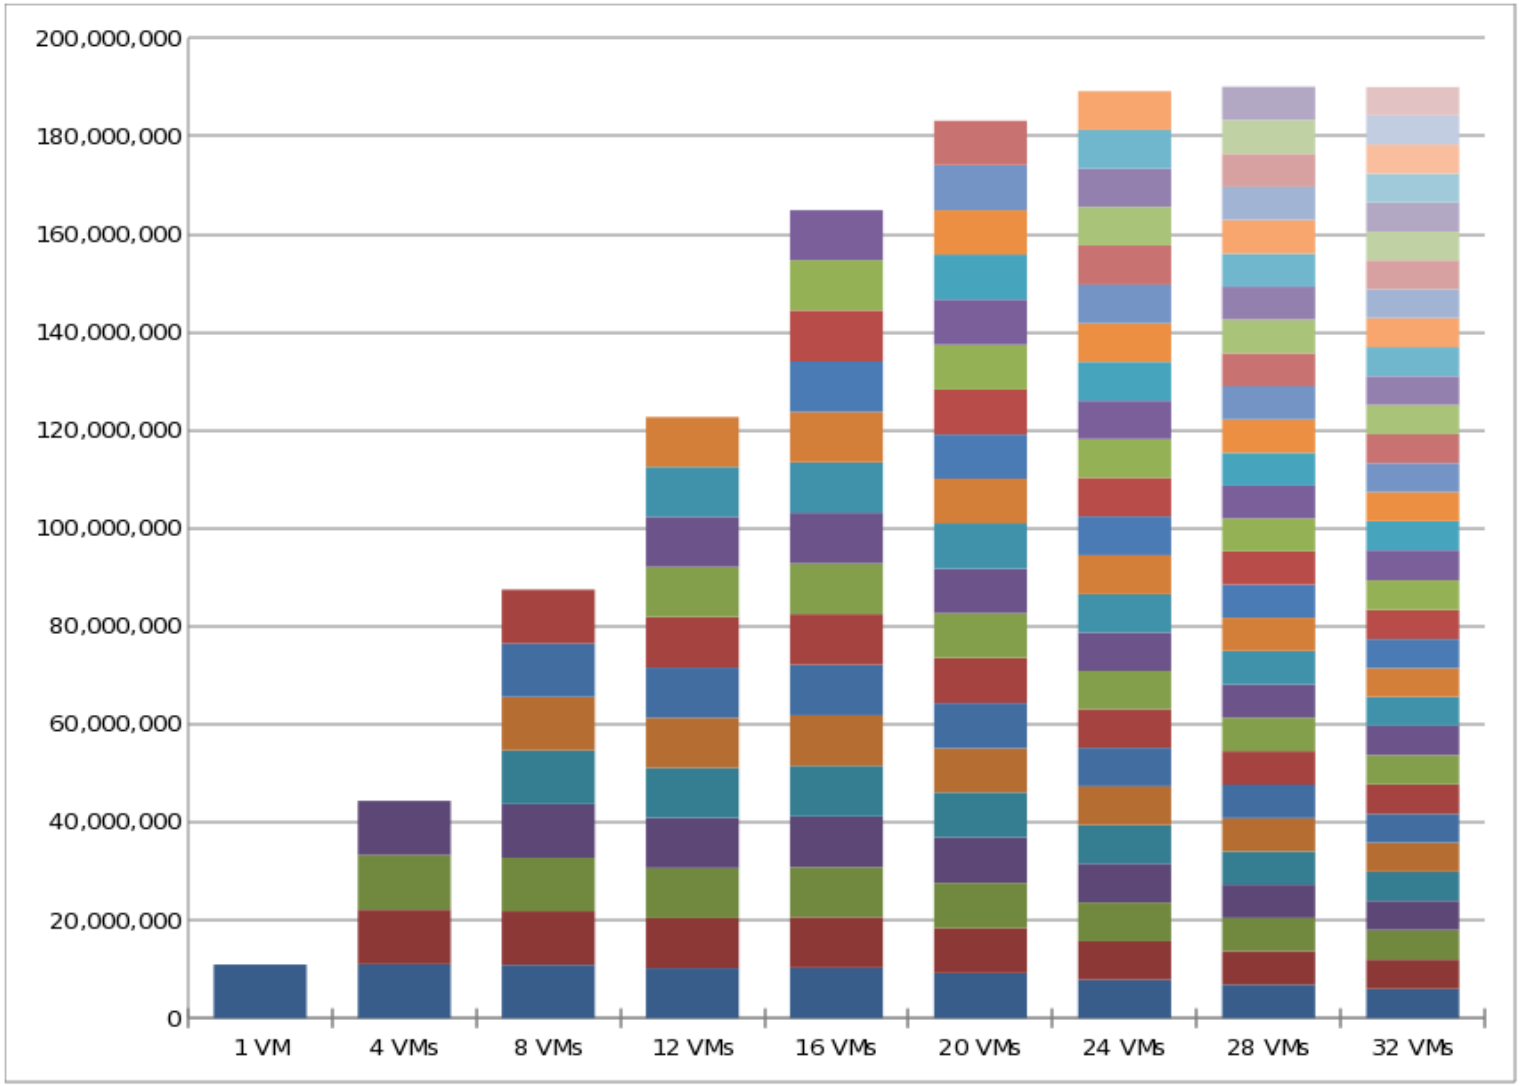

Kernel Same-Page Merging (KSM)

Enterprise Java workload benchmark

Intel Xeon Processor X5550 with 24GB RAM

Running multiple 3GB Windows 2003 VMs

Scaling up to 200% over-commit

RED HAT IN THE VIRTUALIZED ENVIRONMENT
KVM ADVANCED FEATURES

Thin Provisioning

Allocate storage only when needed

Oversubscribe storage

Transparent to VM

Improve storage utilization

Reduced storage costs

Works with NFS, iSCSI, and FC

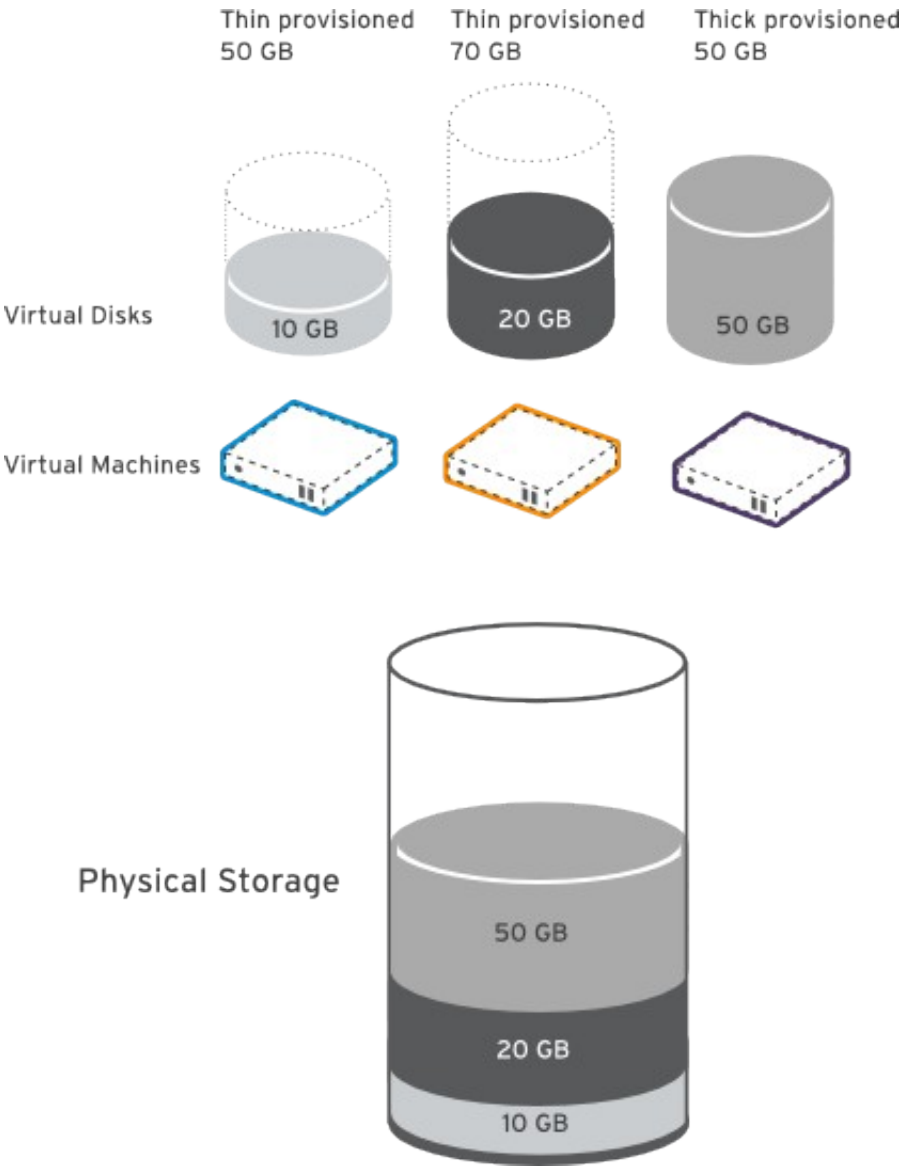

RED HAT IN THE VIRTUALIZED ENVIRONMENT
VIRTUALIZATION IN RHEL 6

Increased scalability

- Guest/host CPU and memory
  - Host: 4096 cores, 64 TB
  - Guest: 64 vCPUs, 1 TB

Increased performance

- Improved guest/host memory management
- Networking improvements
- Storage improvements

Increased Security

- sVirt

Migration Tools Available

- v2v – convert Xen VM's to KVM

RED HAT IN THE VIRTUALIZED ENVIRONMENT
FOUNDATION FOR THE CLOUD

PHASE 1: CONSOLIDATE

VIRTUALIZE YOUR SERVERS

Virtualize your physical hardware to achieve higher utilization, consolidation, and flexibility.

Virtualization increases the utilization of physical servers and provides a foundation for cloud computing.

Virtualization technology included in RHEL6.

PHASE 2: AUTOMATE

BUILD A PRIVATE CLOUD

As you expand your use of virtualization, build a private cloud to manage the scale and complexity.

A private cloud abstracts multiple instances of virtual resources into elastic pools of computation with self-provisioning and scalable services.

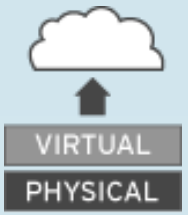

PHASE 3: UTILITY

ADD ADDITIONAL CLOUDS

As you expand your use of cloud computing, add additional clouds delivered as a utility to increase capacity and lower costs.

Red Hat's cloud architecture lets you manage and integrate various virtualization systems and cloud providers together. This allows you to leverage additional clouds as a utility.

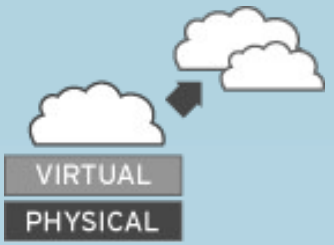

RED HAT IN THE VIRTUALIZED ENVIRONMENT

FOUNDATION FOR THE CLOUD: MANY COMPONENTS ALREADY IN NSA INFRASTRUCTURE

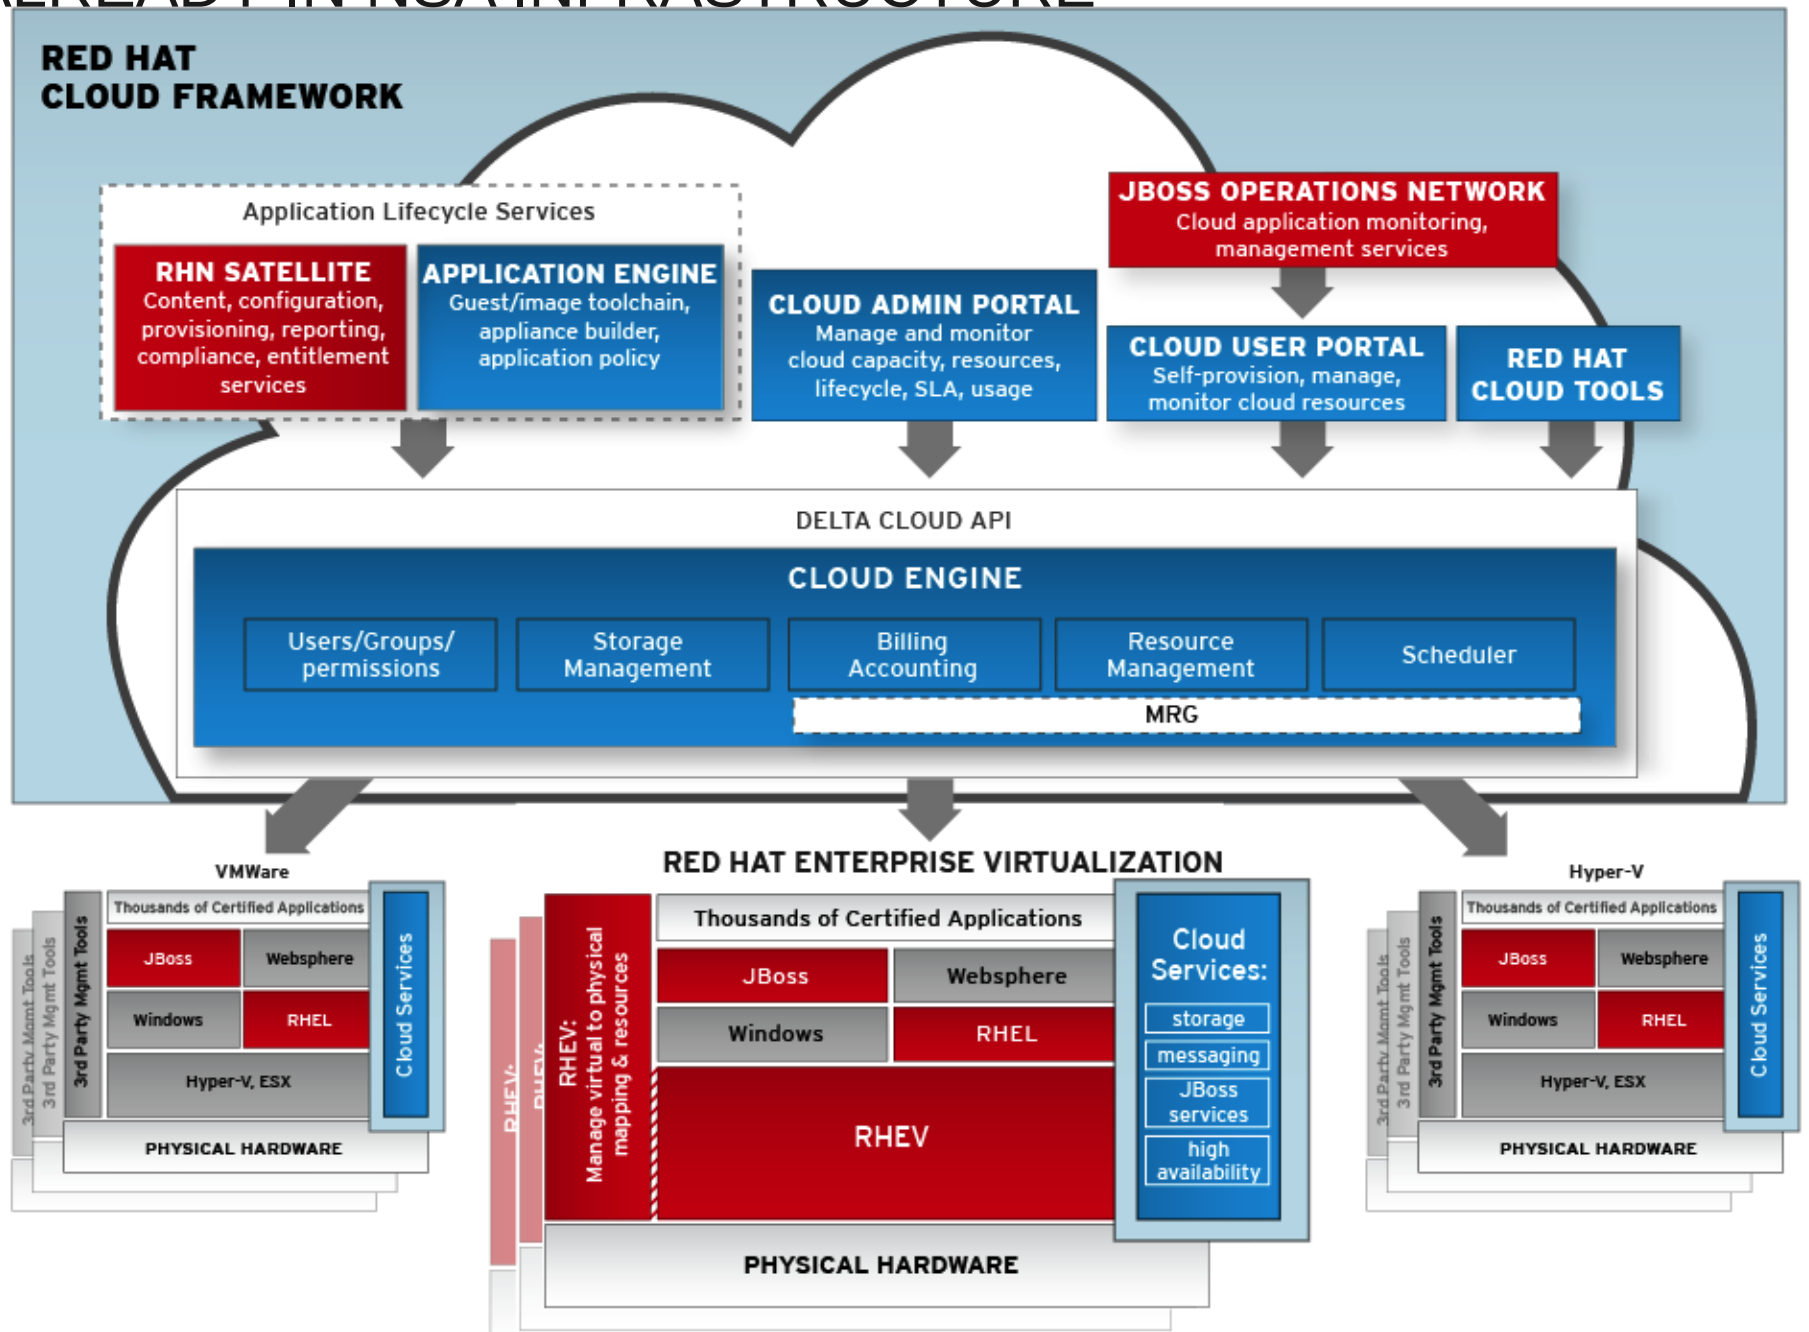

RED HAT SECURITY UPDATE

Gunnar Hellekson
CTO, Red Hat Public Sector
gunnar.hellekson@redhat.com / 202-507-9027

SECURITY
SELINUX

- Confined users
  - Role-based controls to limit system access for users.
- Sandbox
  - Untrusted applications can be run confined to prevent compromising the entire system
- X Access Control Extension (XACE).
- SELinux Kiosk Mode
  - Creation of a user session environment that is valid for a limited time.

SECURITY
MLS DESKTOP

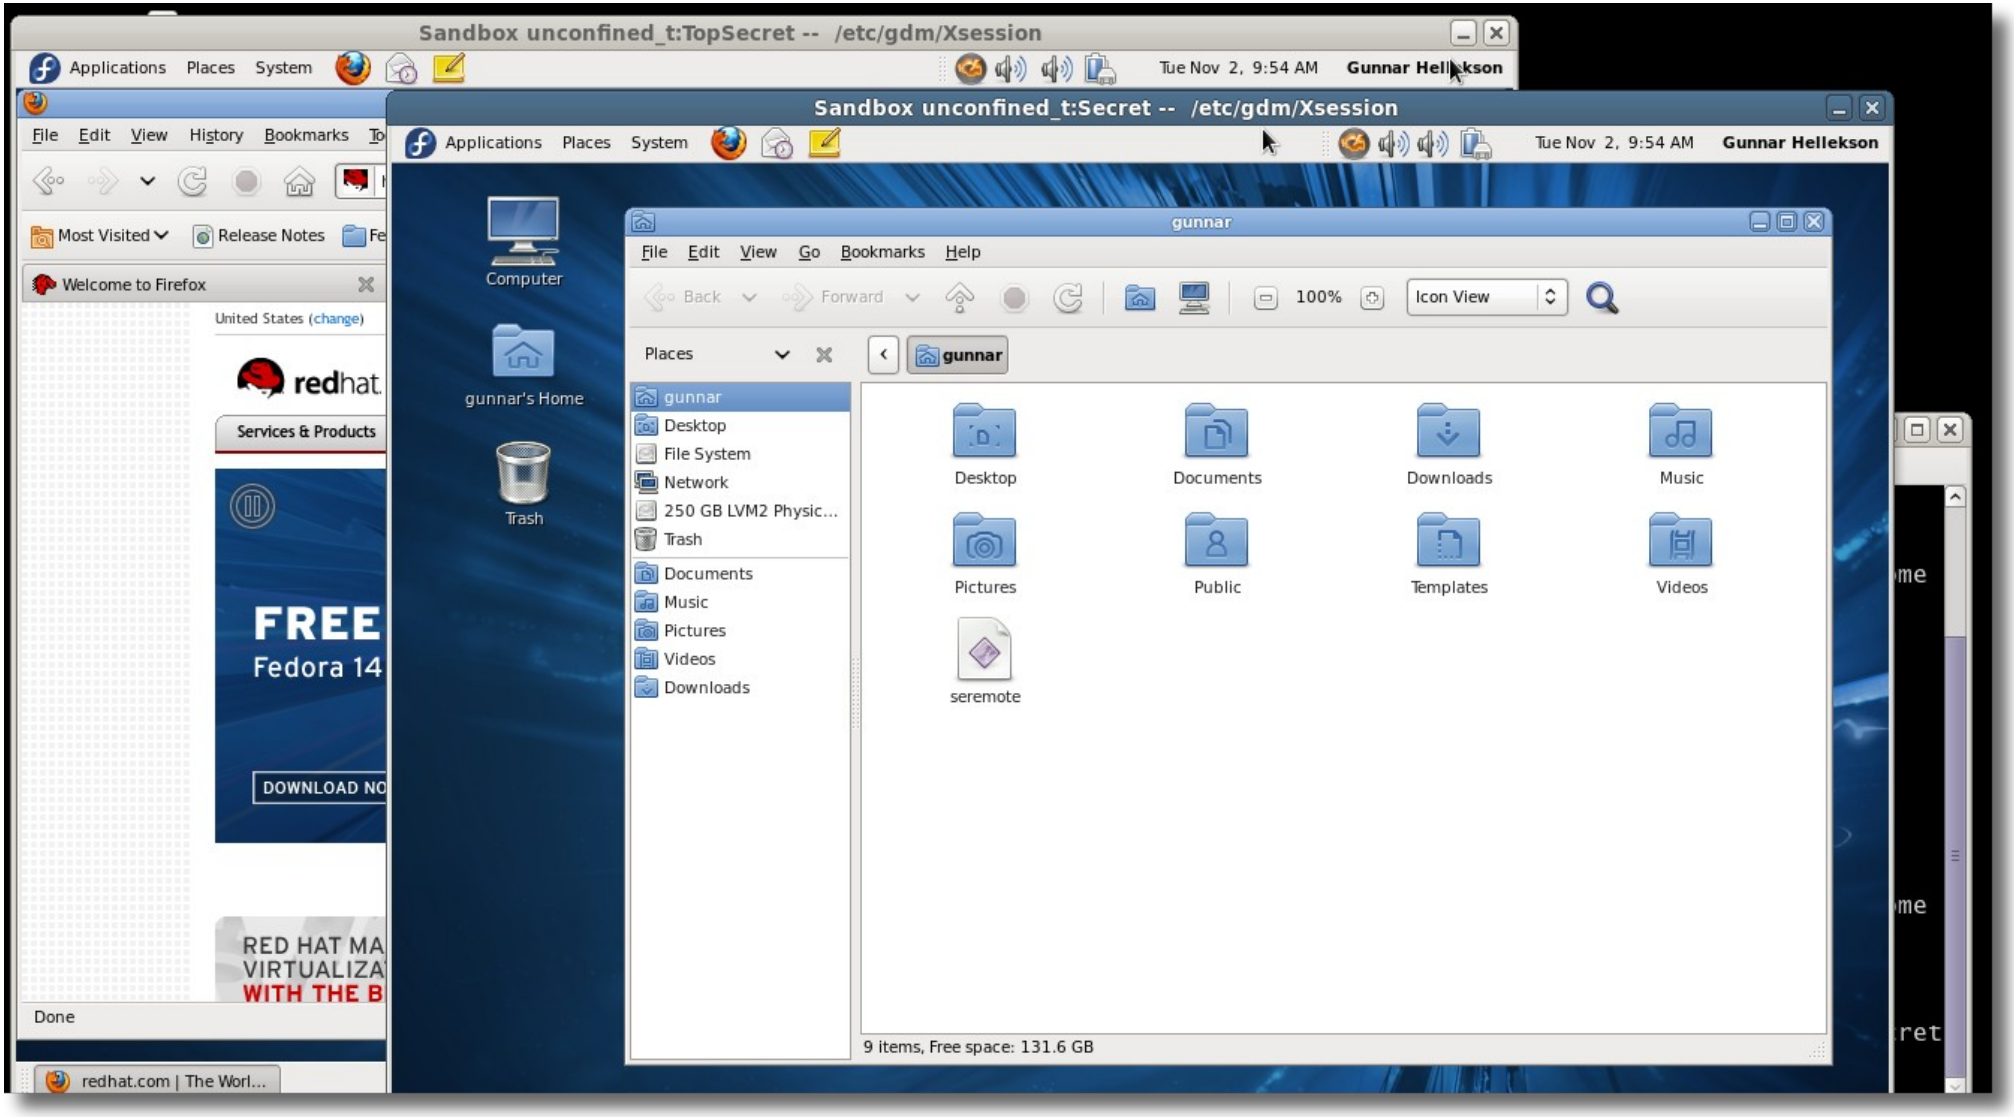

SECURITY
VIRTUALIZATION

RHEV and KVM inherit the security features of Linux and RHEL

SELinux security policy infrastructure

- Provides protection and isolation for virtual machines and host
- Compromised virtual machine cannot access other VMs or host

sVirt Project

- Sub-project of NSA's SELinux community. Provides “hardened” hypervisors
- Multilevel security. Isolate guests
- Contain any guest breaches

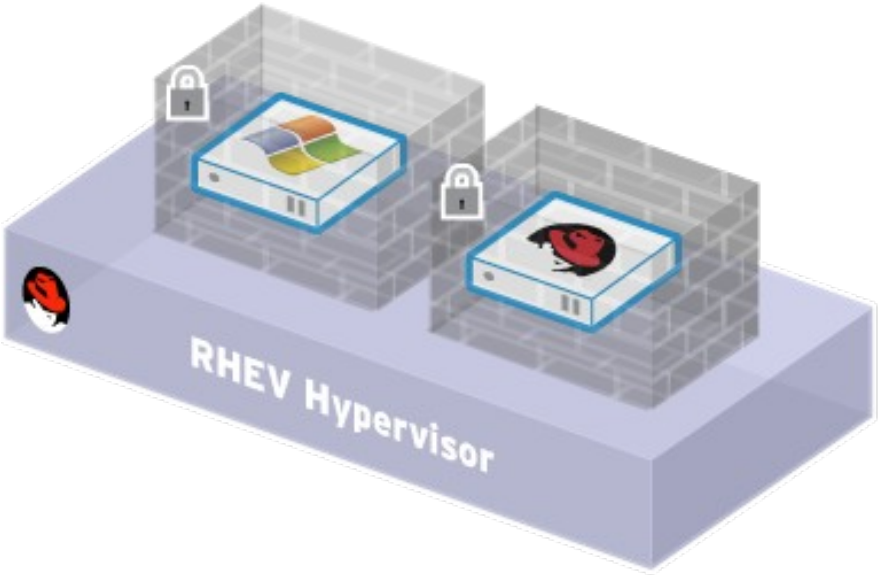

SECURITY
SCAP

SCAP and TNC

- Provides continuous lifecycle management for virtual machines and their hosts
- Compromised virtual machine can be quarantined
- Allows a provider's security SLA to be enforced by customers

OpenSCAP

- An open source implementation of SCAP
- Included in RHEL 6

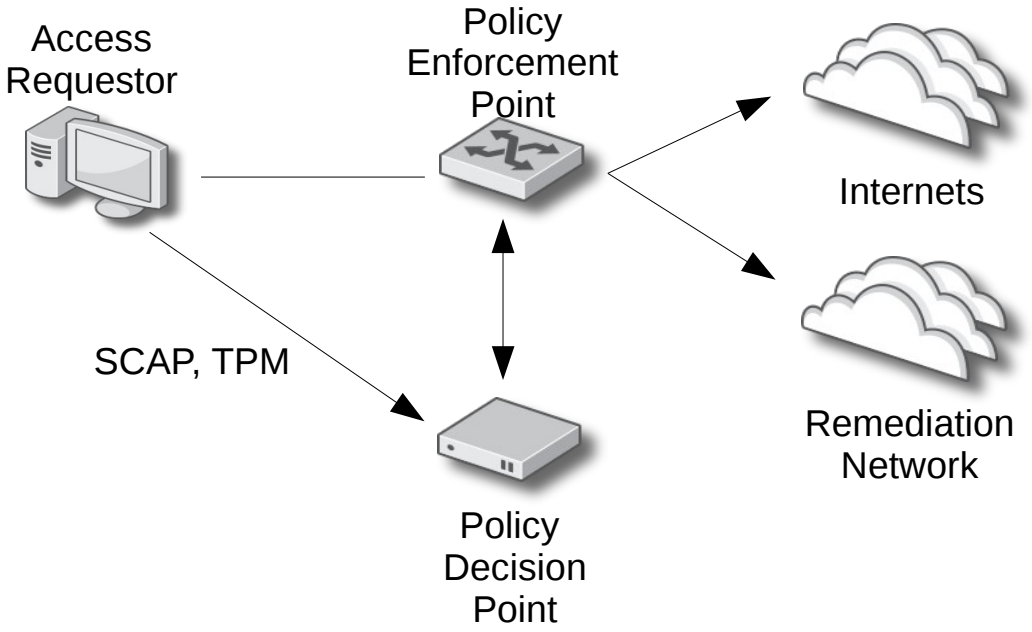

SECURITY
COMMON CRITERIA UPDATE

RHEL 5.6

- Refreshed to include KVM
- Custom protection profile

RHEL 6

- Will include sVirt, IPSec
- Targeting EAL 4
- Again, custom protection profile

Q&A PANEL SESSION