RED HAT SUMMIT

KVM Virtualization

Roadmap and Technology Update

Karen Noel
Senior Eng Manager
Red Hat, Inc.

Bhavna Sarathy
Senior Product Manager
Red Hat, Inc.

June 13, 2013

redhat.

Why we believe KVM is the best virtualization platform

Performance

KVM holds the **Top 6/11** virtual machine consolidation scores on SPECvirt (1)

Lower Cost

customers report up to **70%** savings by using KVM (2)

Cross Platform

Support and certification for leading x86_64 operating systems including **RHEL** and **Microsoft Windows** (4)

Security

EAL4+ Certification (3) plus SE Linux enabling Mandatory Access Control between virtual machines

Cloud & Virtualization Management

Red Hat Open Stack for Cloud Virtualization and Red Hat Enterprise Virtualization for data-center Virtualization

(1) Source: SpecVirt_sc2010 results: http://www.spec.org/virt_sc2010/results/specvirt_sc2010_perf.html
(2) Source: Case study on Canary Islands Government migration from VMware to RHEV: http://www.redhat.com/resourcelibrary/case-studies/canary-islands-government-migrates-telecommunications-platform-from-vmware-to-red-hat
(3) Source: http://www.redhat.com/solutions/industry/government/certifications.html
(4) Source: http://www.redhat.com/resourcelibrary/articles/enterprise-linux-virtualization-support

KVM hypervisor in multiple Red Hat products

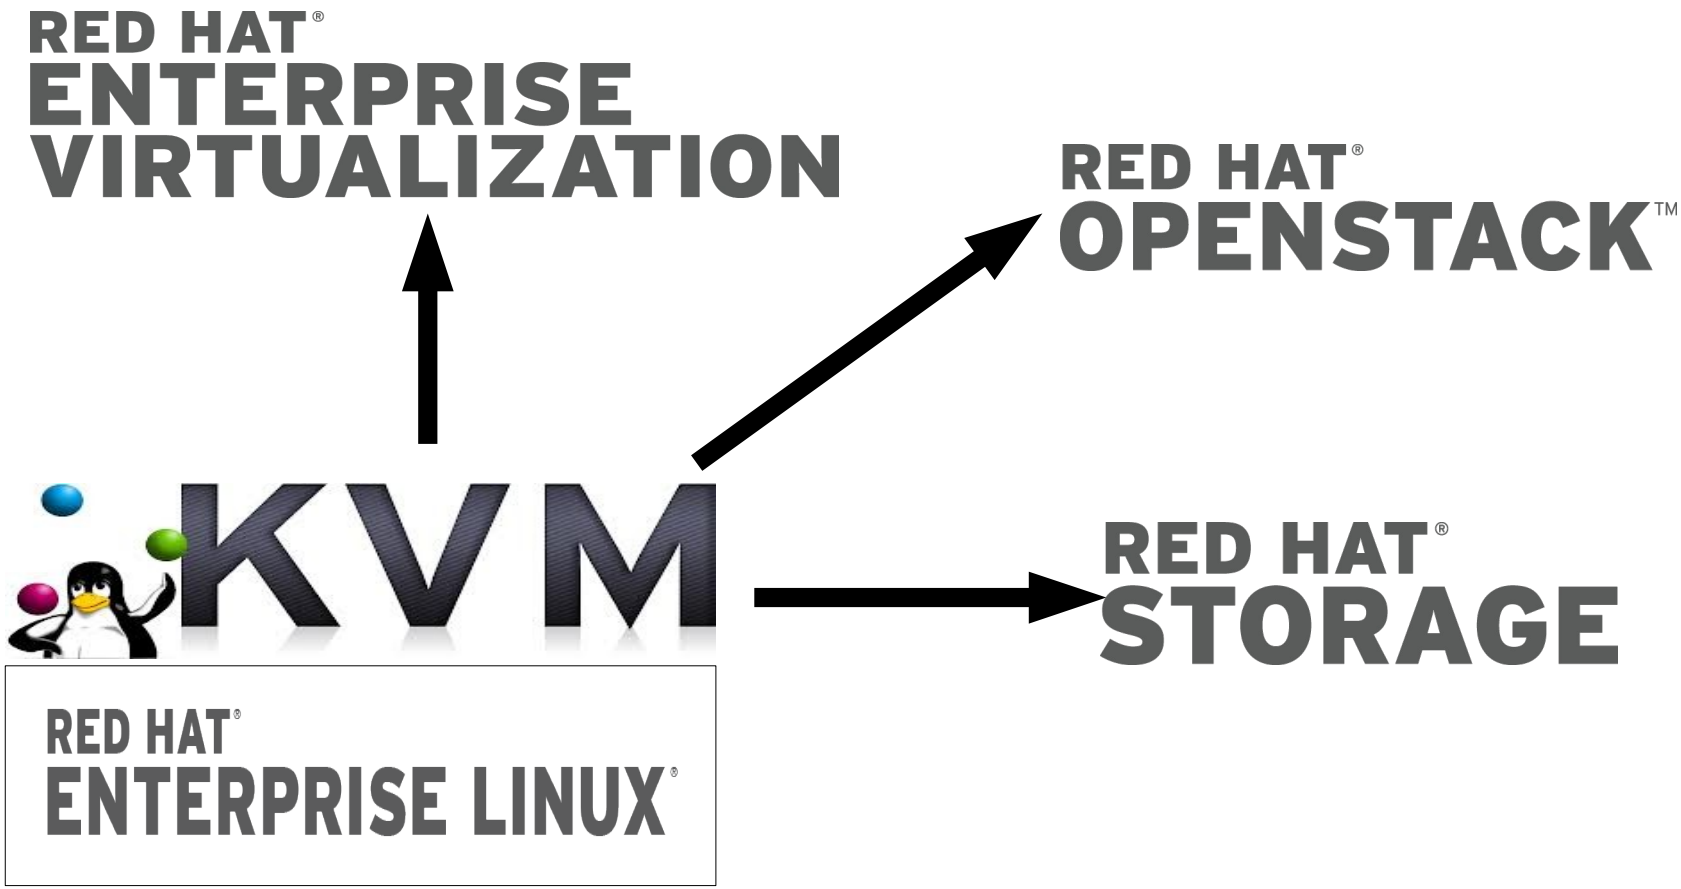

KVM is the foundation Virtualization technology in multiple Red Hat products

Red Hat Enterprise Virtualization – Hypervisor

Derived from Red Hat Enterprise Linux

SMALL FORM FACTOR, SCALABLE, HIGH PERFORMANCE

RED HAT ENTERPRISE LINUX 3 Server
RED HAT ENTERPRISE LINUX 4 Server
RED HAT ENTERPRISE LINUX 5 Server
RED HAT ENTERPRISE LINUX

Windows Server 2003
Windows Server 2008
Windows Server 2008 R2
Windows xp
Windows 7

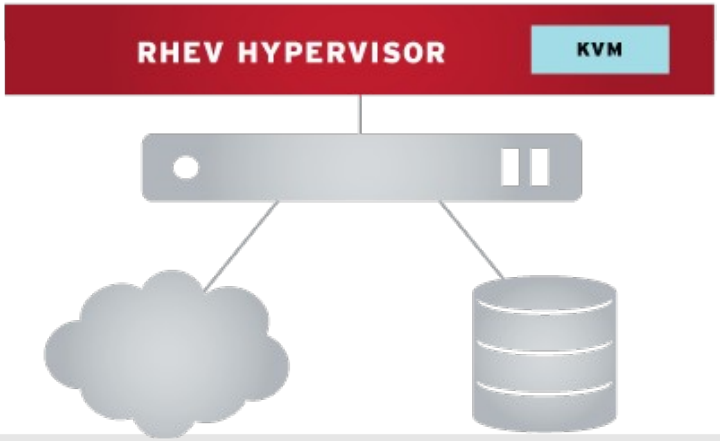

- RHEV Hypervisor
  - Prebuilt binary (ISO) with 300+ packages derived from RHEL
  - Inherits performance, scalability, security and supportability of Red Hat Enterprise Linux
  - Shares RHEL & KVM software and hardware ecosystem

KVM Architecture: Integrated Virtualization

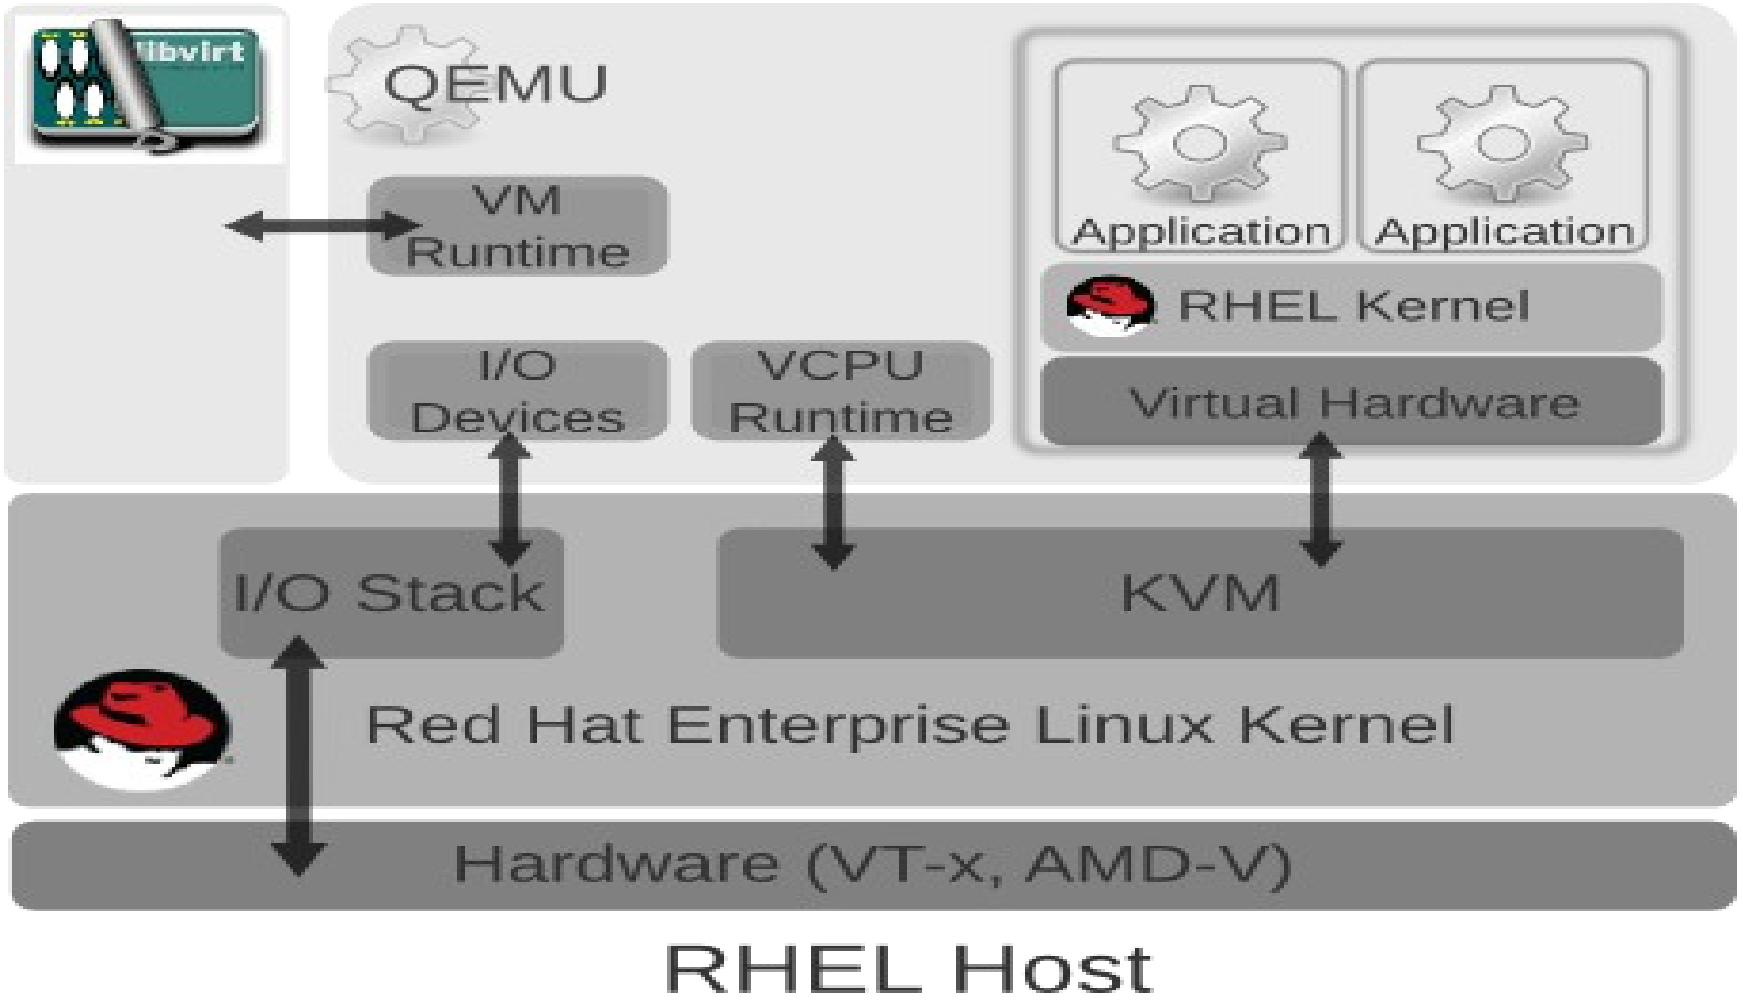

KVM I/O Architecture

Emulated Devices

- Native drivers
- Compatibility over performance

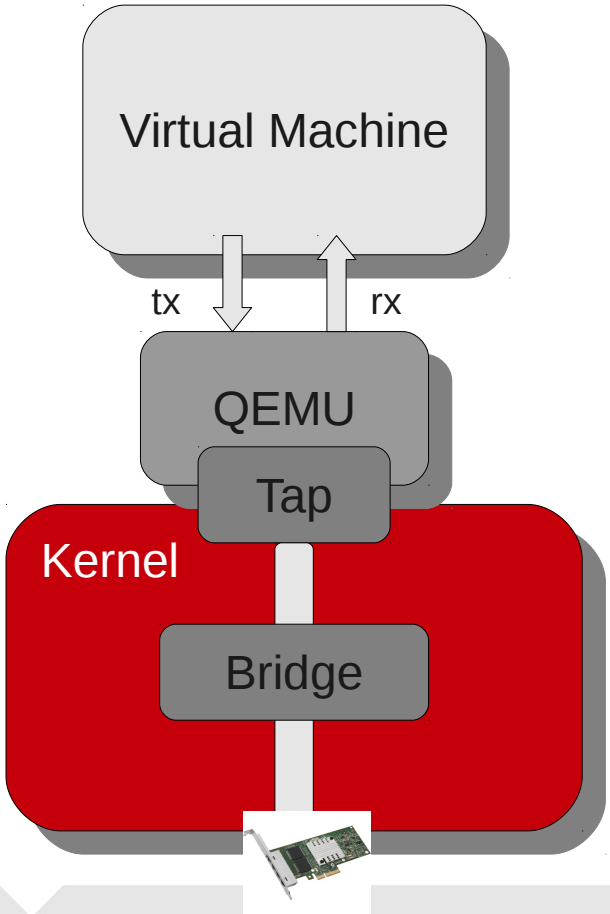

Virtio Devices

- Paravirtualized
- Performance over compatibility

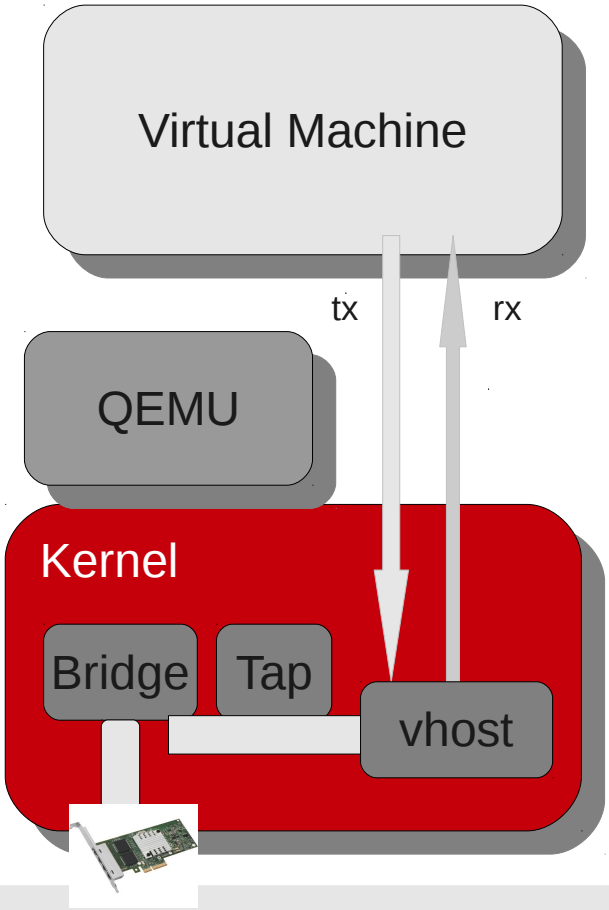

Device Assignment

- Native drivers
- Compatibility and Performance

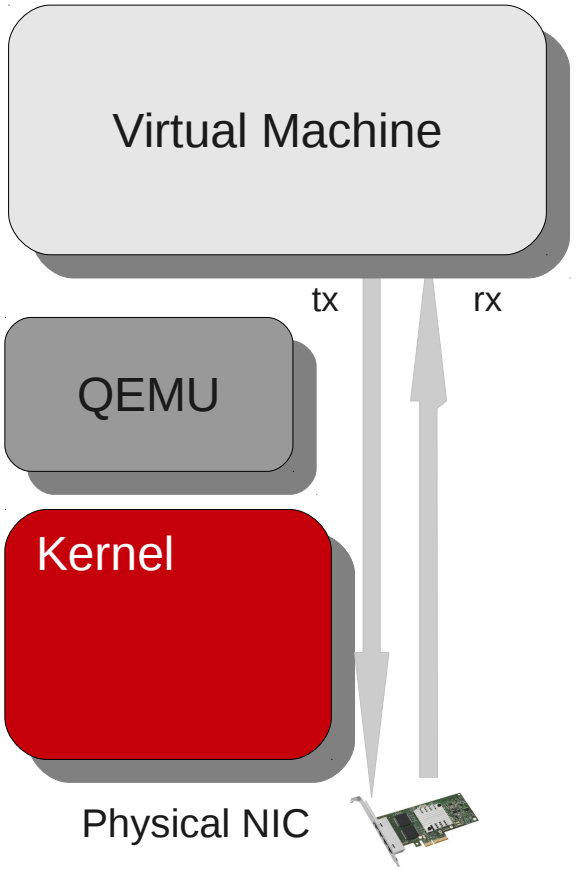

Let's have a look at RHEL6

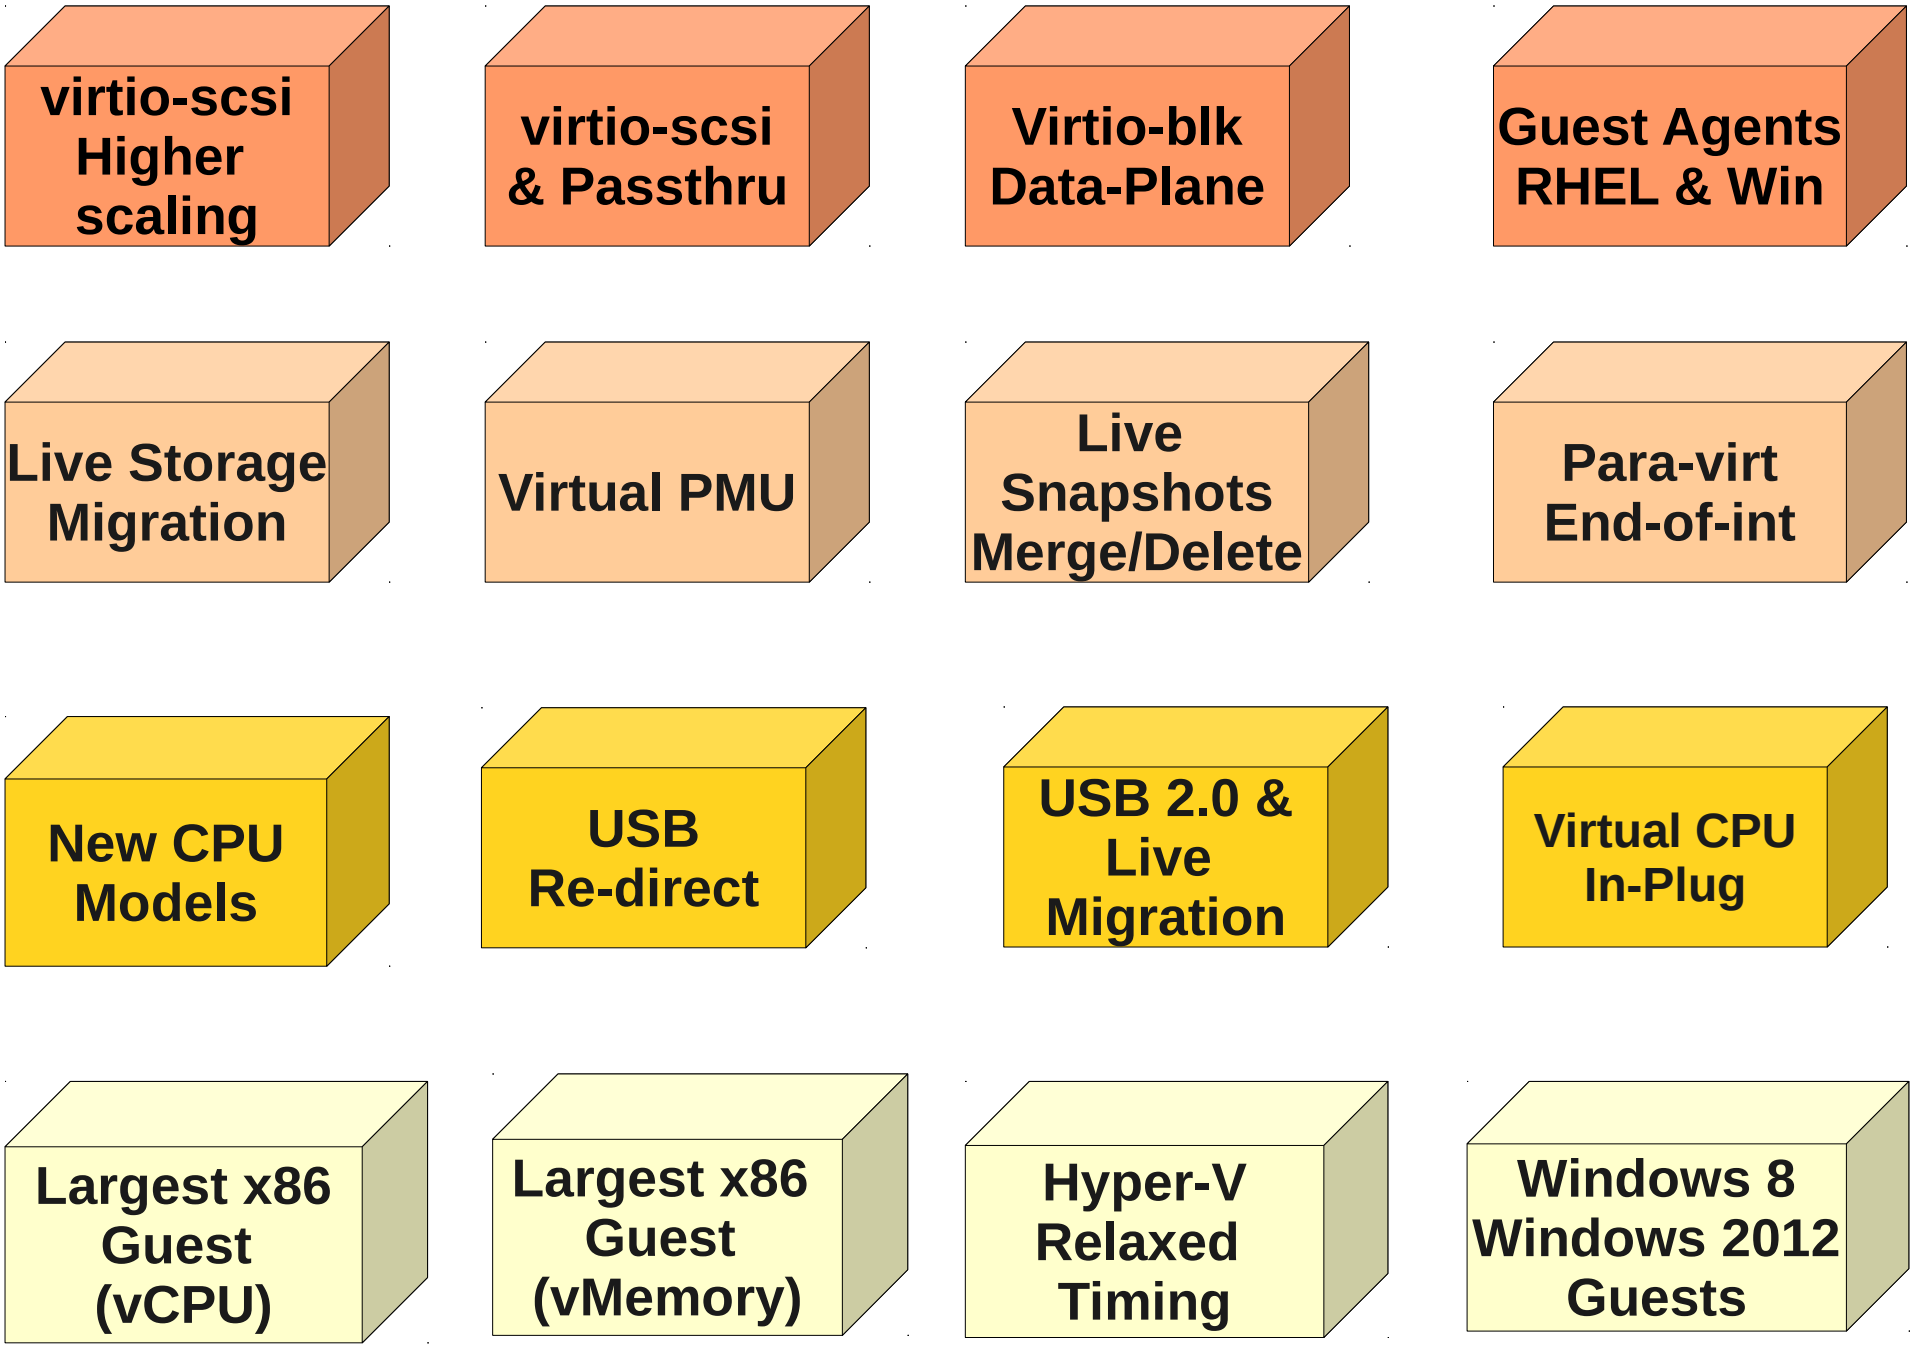
virtio-scsi Higher scaling
virtio-scsi & Passthru
Virtio-blk Data-Plane
Guest Agents RHEL & Win
Live Storage Migration
Virtual PMU
Live Snapshots Merge/Delete
Para-virt End-of-int
New CPU Models
USB Re-direct
USB 2.0 & Live Migration
Virtual CPU In-Plug
Largest x86 Guest (vCPU)
Largest x86 Guest (vMemory)
Hyper-V Relaxed Timing
Windows 8 Windows 2012 Guests

These were all introduced in RHEL 6.3 and RHEL6.4 alone

Virtualization with KVM
Performance

SPECvirt_sc2010: *RHEL 6 KVM Posts Industry Leading Results*

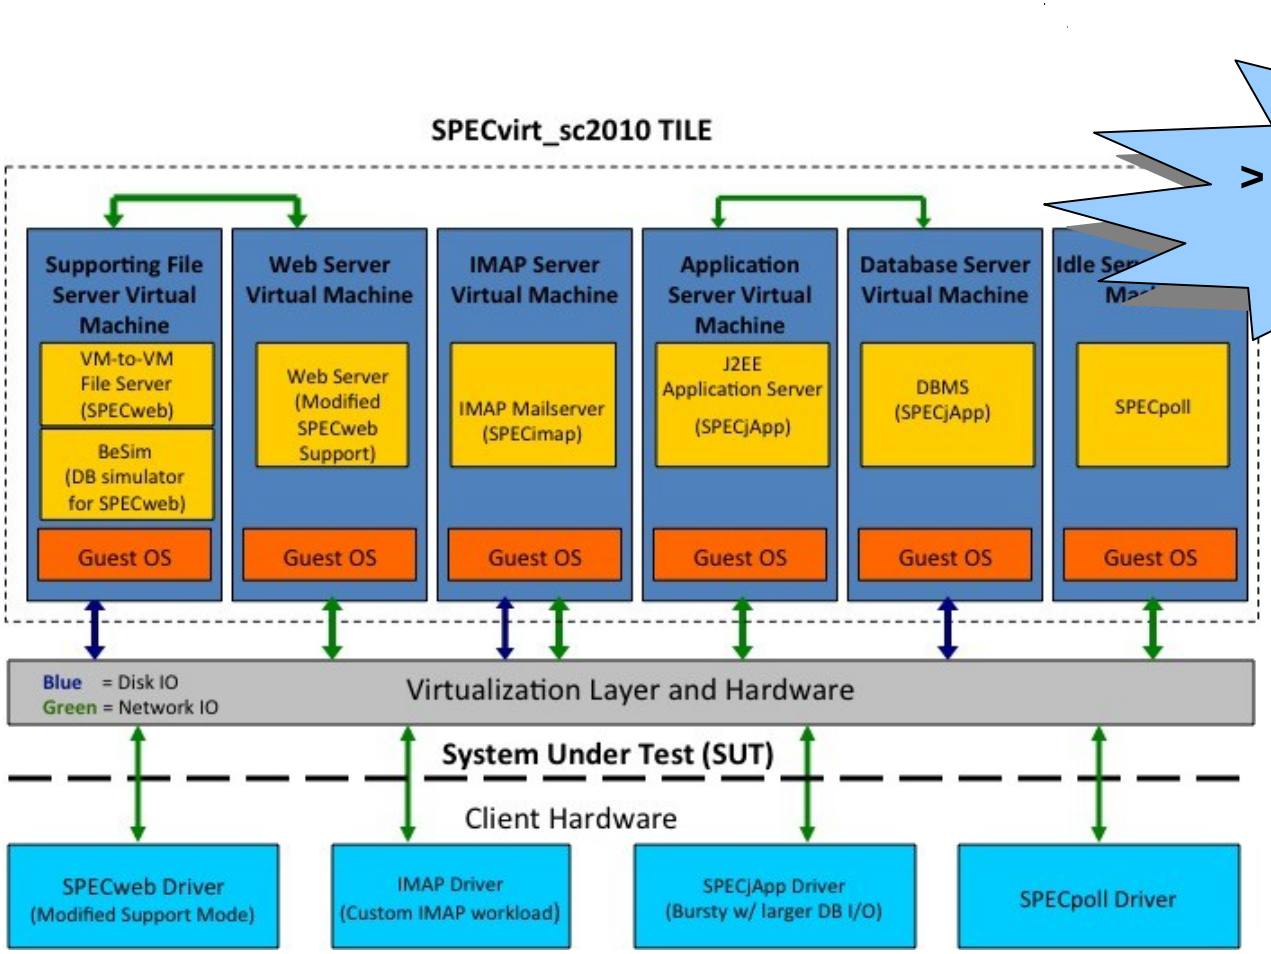

> 1 SPECvirt Tile/core

RHEL (KVM) key enablers:

- ✓ SR-IOV
- ✓ Huge Pages
- ✓ NUMA
- ✓ Node Binding

http://www.spec.org/virt_sc2010/results/

Best SPECvirt_sc2010 Scores by CPU Cores

(As of May 30, 2013)

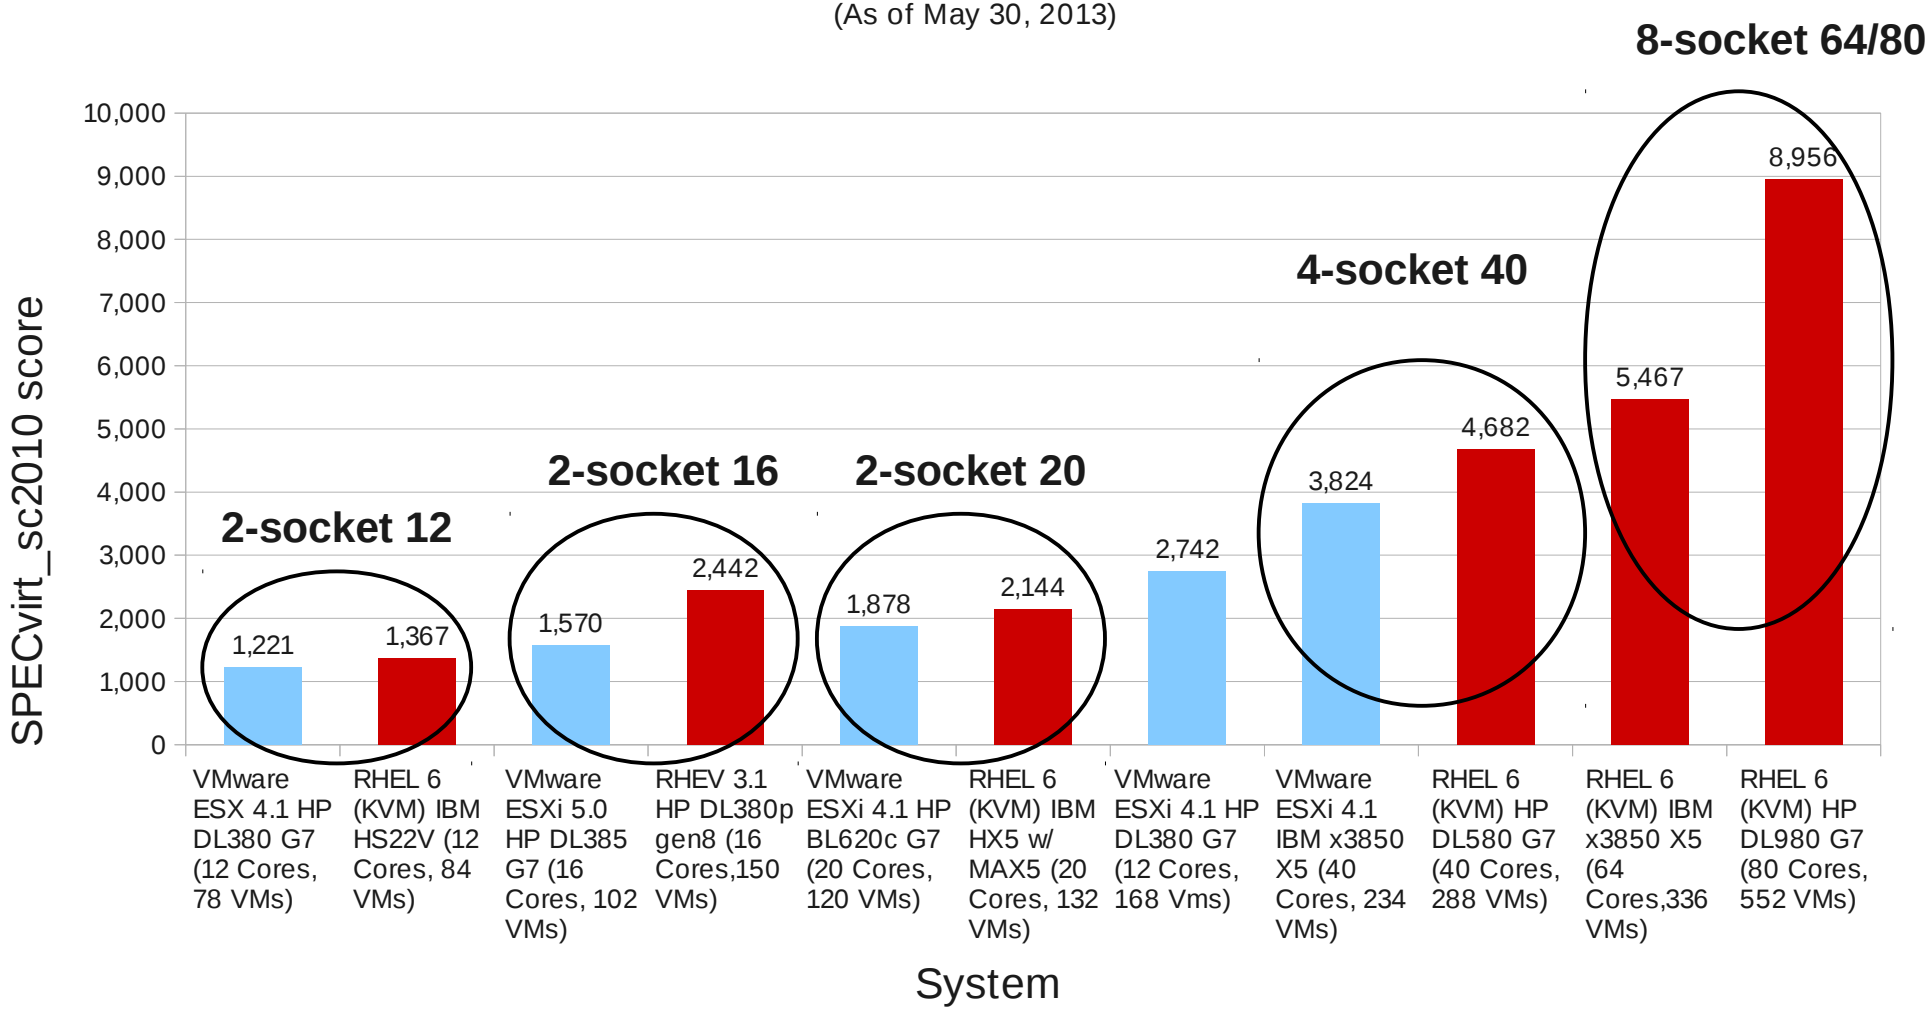

Comparison based on best performing Red Hat and VMware solutions by cpu core count published at www.spec.org as of May 17, 2013. SPEC® and the benchmark name SPECvir_sct® are registered trademarks of the Standard Performance Evaluation Corporation. For more information about SPECvirt_sc2010, see www.spec.org/virt_sc2010/.

Best SPECvirt_sc2010 Results by Number of Sockets

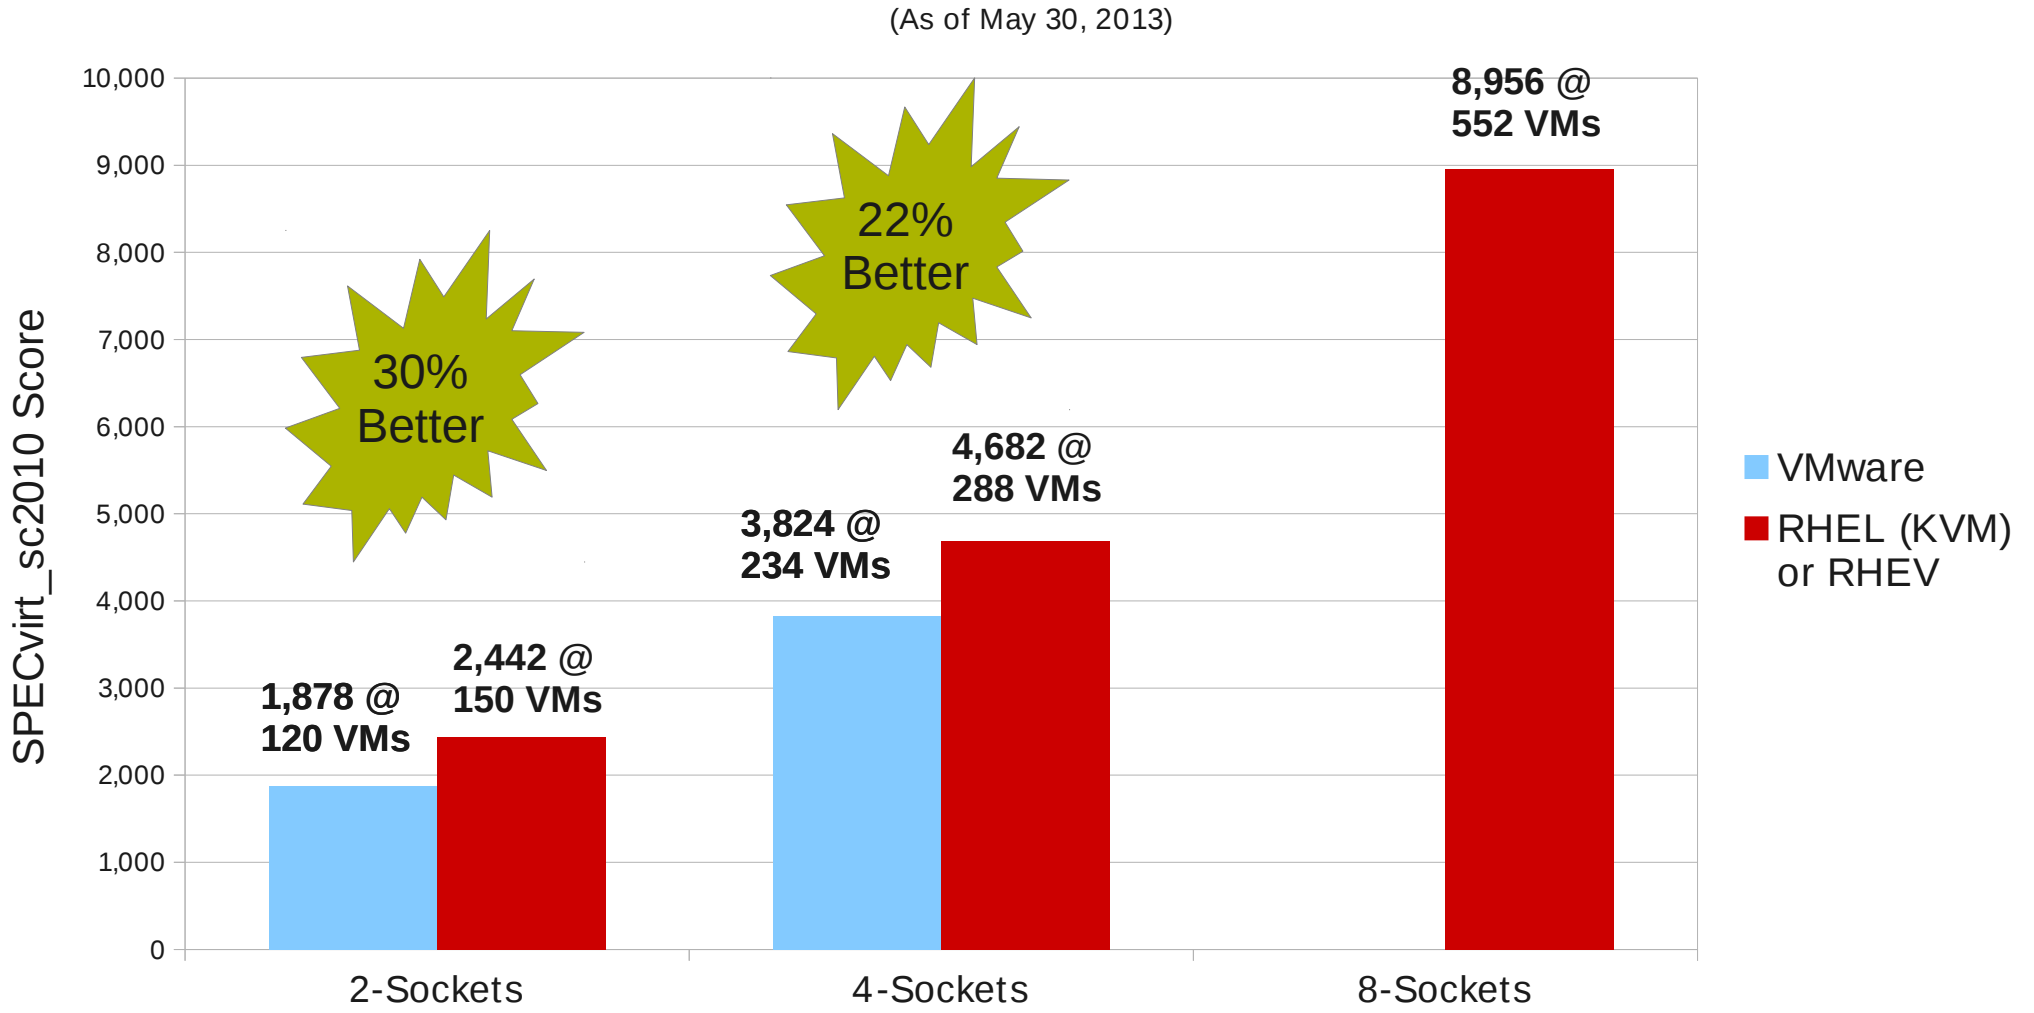

Comparison based on best performing Red Hat and VMware solutions by cpu socket count published at www.spec.org as of May 30, 2013. SPEC® and the benchmark name SPECvirt_sc® are registered trademarks of the Standard Performance Evaluation Corporation. For more information about SPECvirt_sc2010, see www.spec.org/virt_sc2010/.

SPECvirt_sc2013 Tile

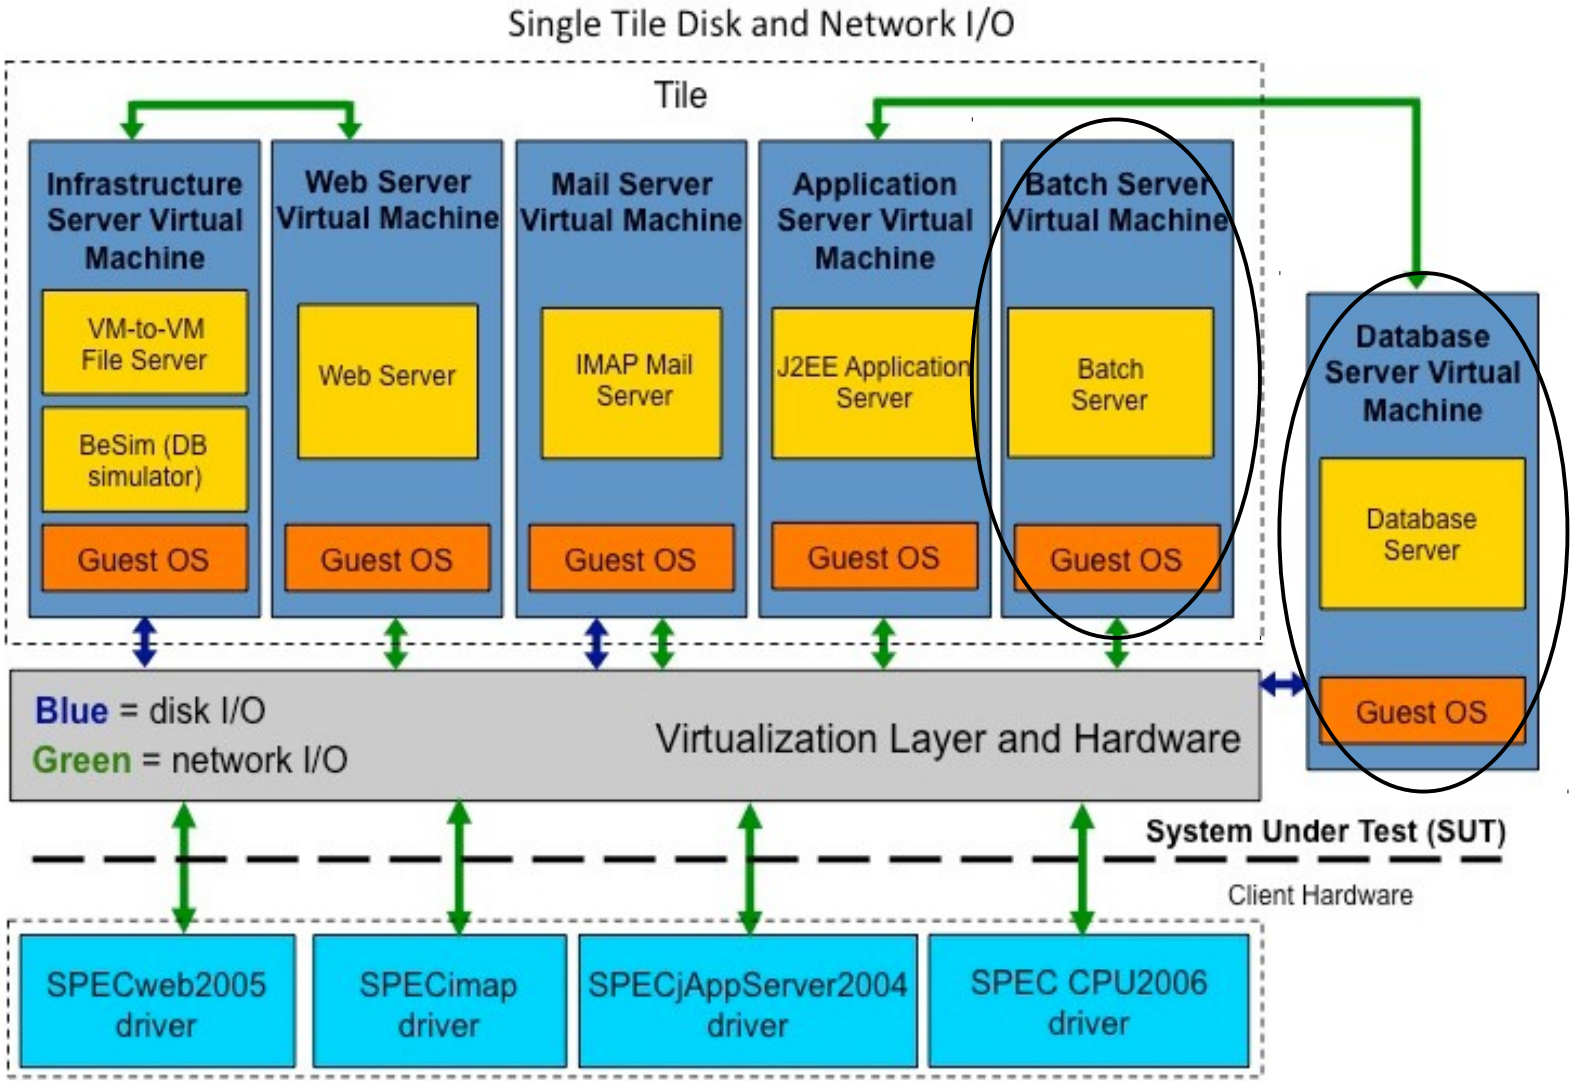

SPECvirt_sc2013 Scores by Number of Sockets

(As of May 30, 2013)

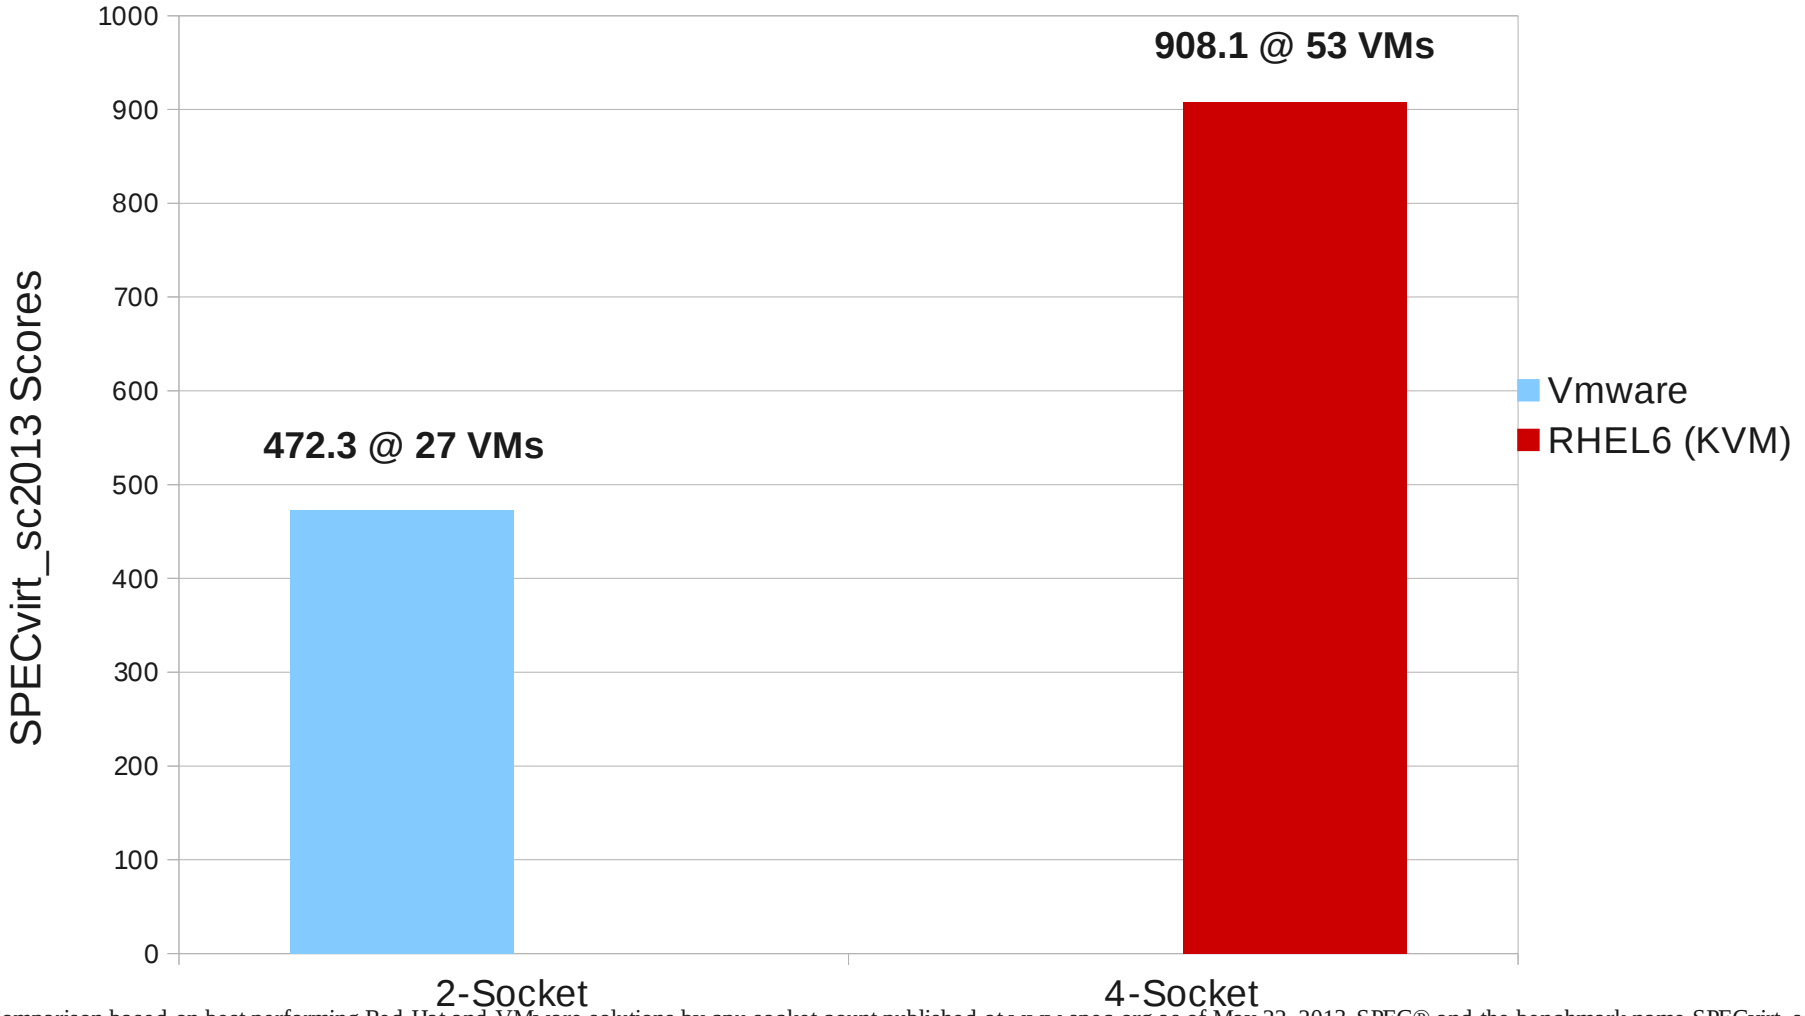

Comparison based on best performing Red Hat and VMware solutions by cpu socket count published at www.spec.org as of May 22, 2013. SPEC® and the benchmark name SPECvirt_sc
Performance Evaluation Corporation. For more information about SPECvirt_sc2013, see www.spec.org/virt_sc2013/.

Virtualization – Performance

Leadership Virtual Disk I/O performance (RHEL 6.4)

1,577,684 IOPS!

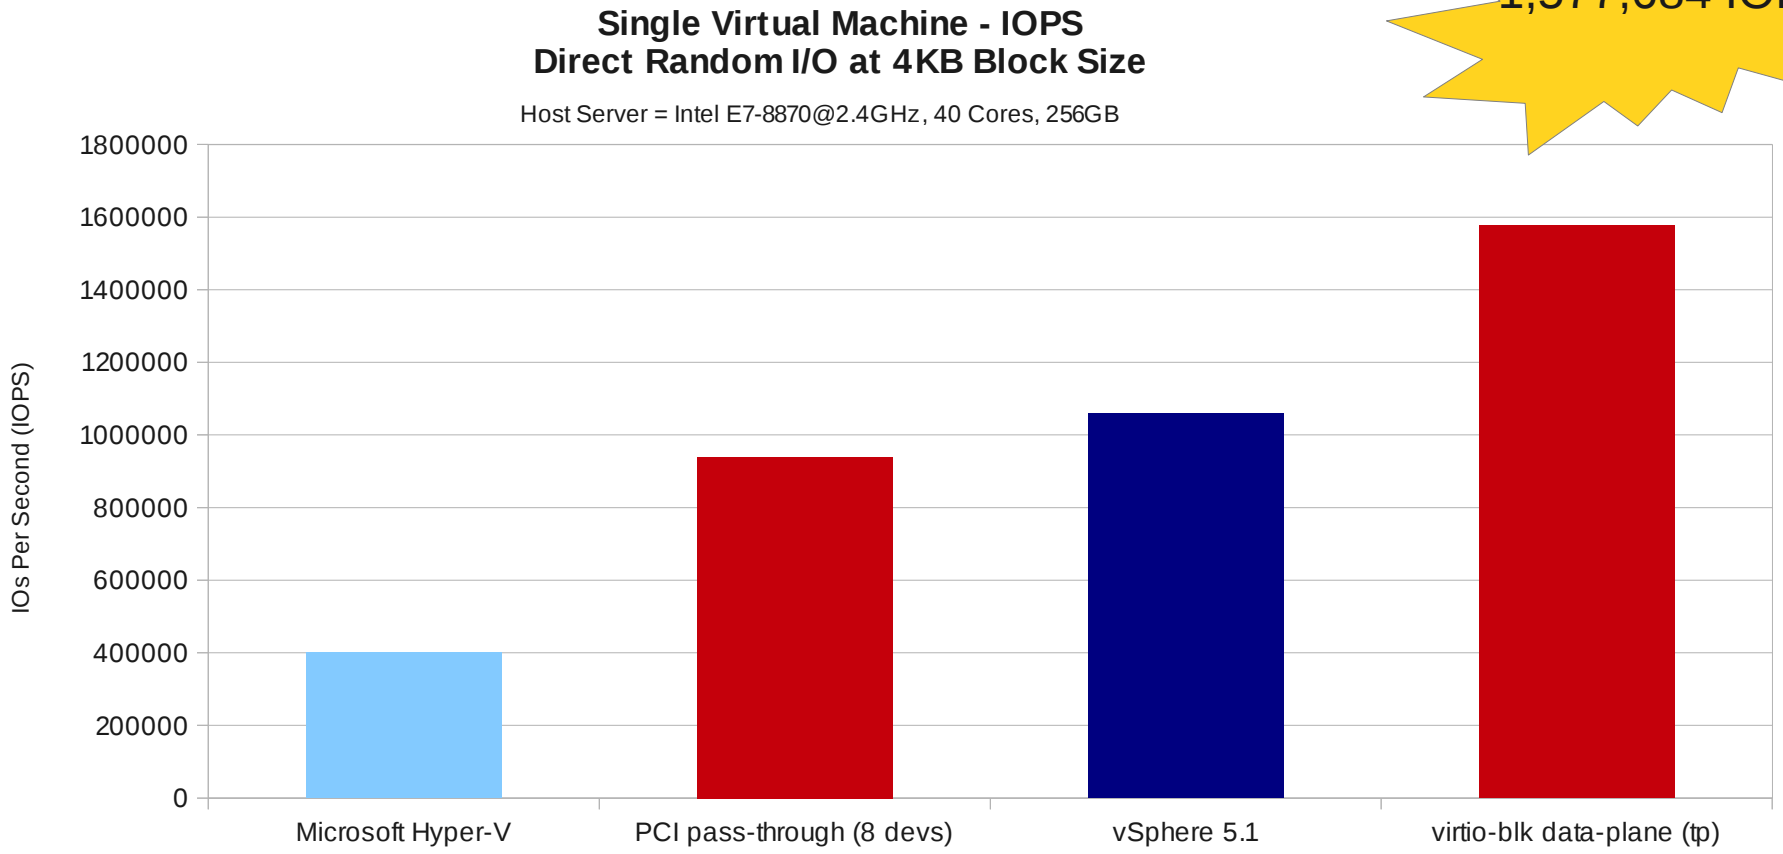

http://www.redhat.com/resourcelibrary/reference-architectures/leadership-io-performance-with-red-hat-enterprise-linux-kvm-hypervisor-on-ibm-x385-ex5

Virtualization – Performance

- Only virtualized TPC-C result (RHEL 6.4 with KVM)
  - Virtualized was approx. 88% of bare-metal
  - Great $/tpmC: $0.51
  - Feb 2013: 1,320,082 tpmC, 0.51 $/tpmC, RHEL6.4 w/KVM, DB2 ESE 9.7
  - Apr. 2012: 1,503,544 tpmC, 0.53 $/tpmC, RHEL 6.2, DB2 ESE 9.7

See the details for these results at: http://www.tpc.org/tpcc/results/tpcc_last_ten_results.asp.
Results referenced are current as of May 30, 2013. To view all TPC results, visit www.tpc.org
http://www.redhat.com/resourcelibrary/reference-architectures/leadership-virtualized-tpc-c-benchmark-using-red-hat-enterprise-linux-kvm

Let's look at
Scalability

Once we accept our limits, we go beyond them.

Albert Einstein

Single Guest Scalability – Virtual CPUs

RHEL6.4 vCPU = 2.5 x * vSphere 5.1 vCPU limits

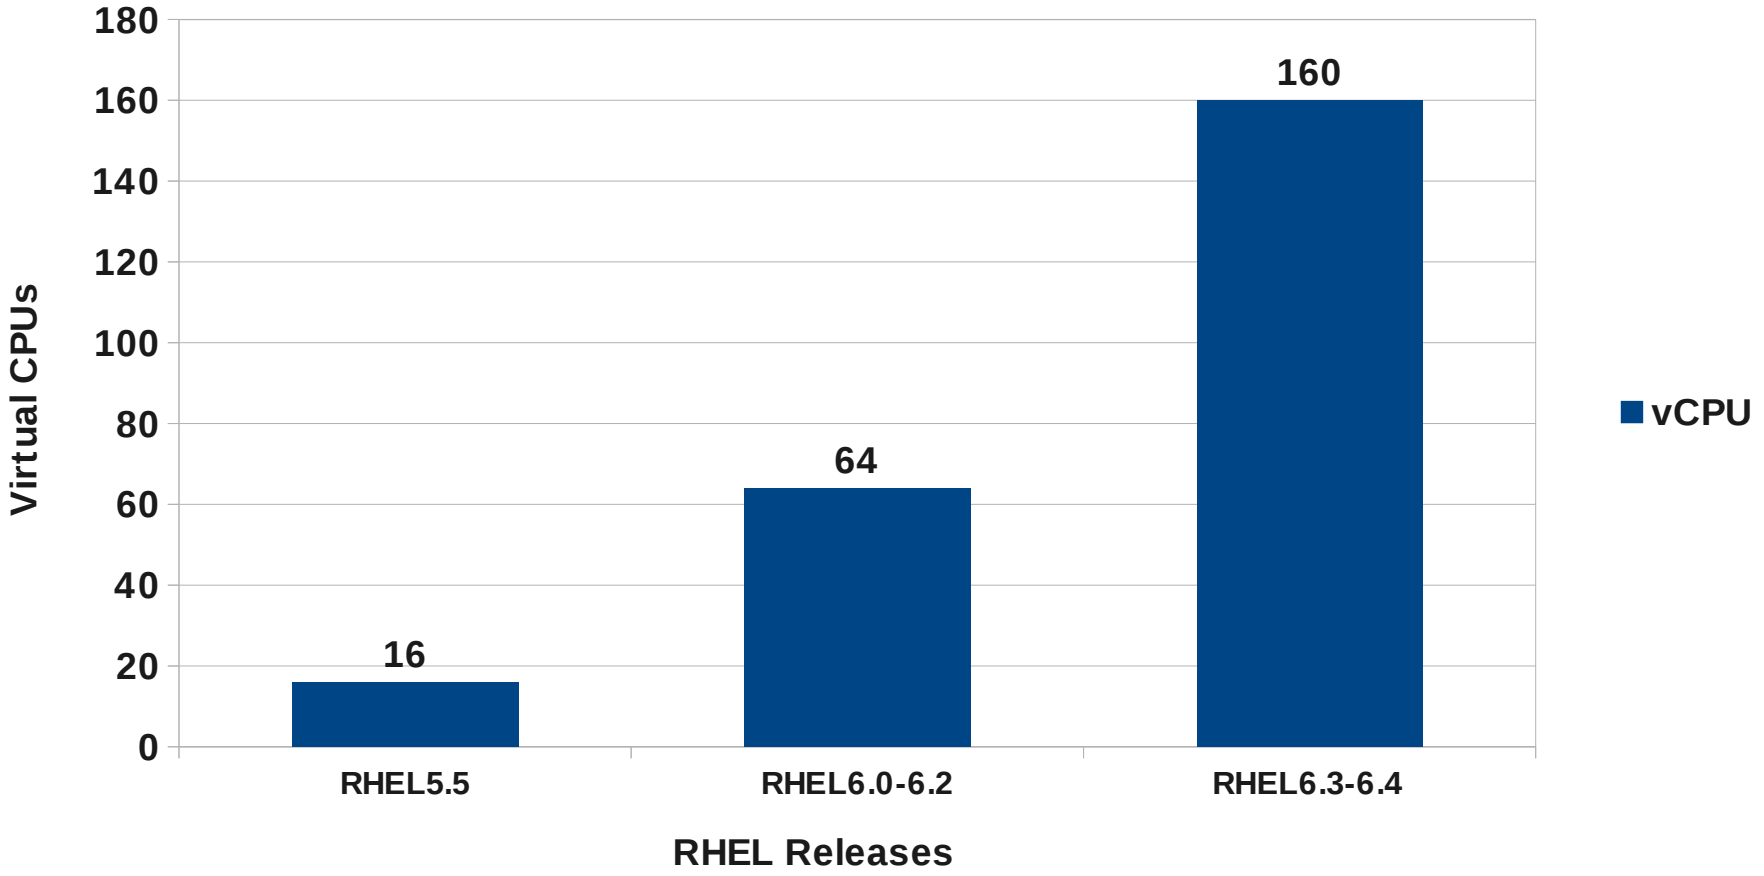

Single Guest Scalability – Virtual Memory

RHEL6.4 vMemory = 2 x * vSphere 5.1 vMemory limits

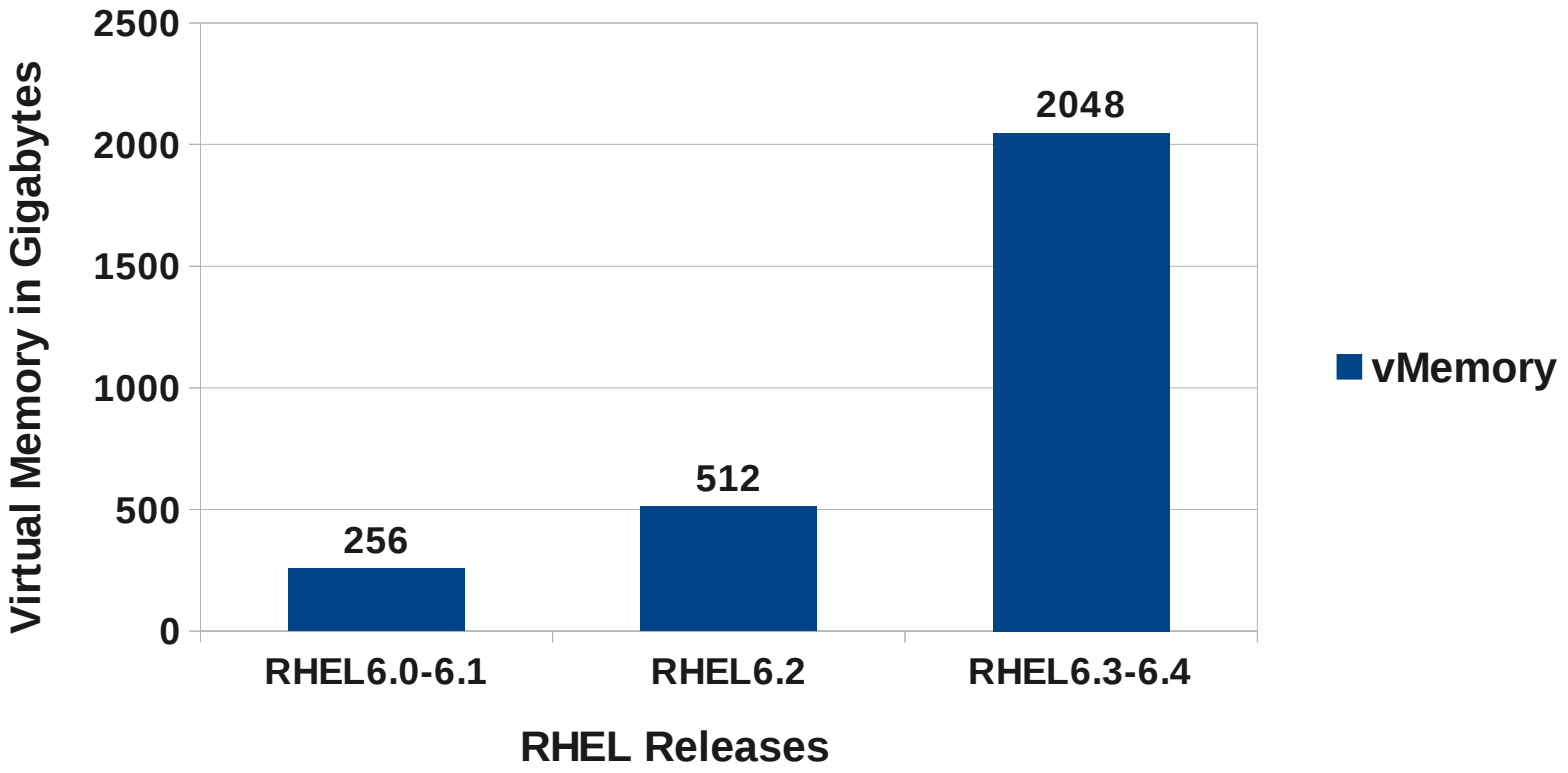

http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf

The very important topic of

Security

KVM Security and Isolation - sVirt

Attacks on guests are isolated from the host kernel,any associated storage, as well as other virtual machines

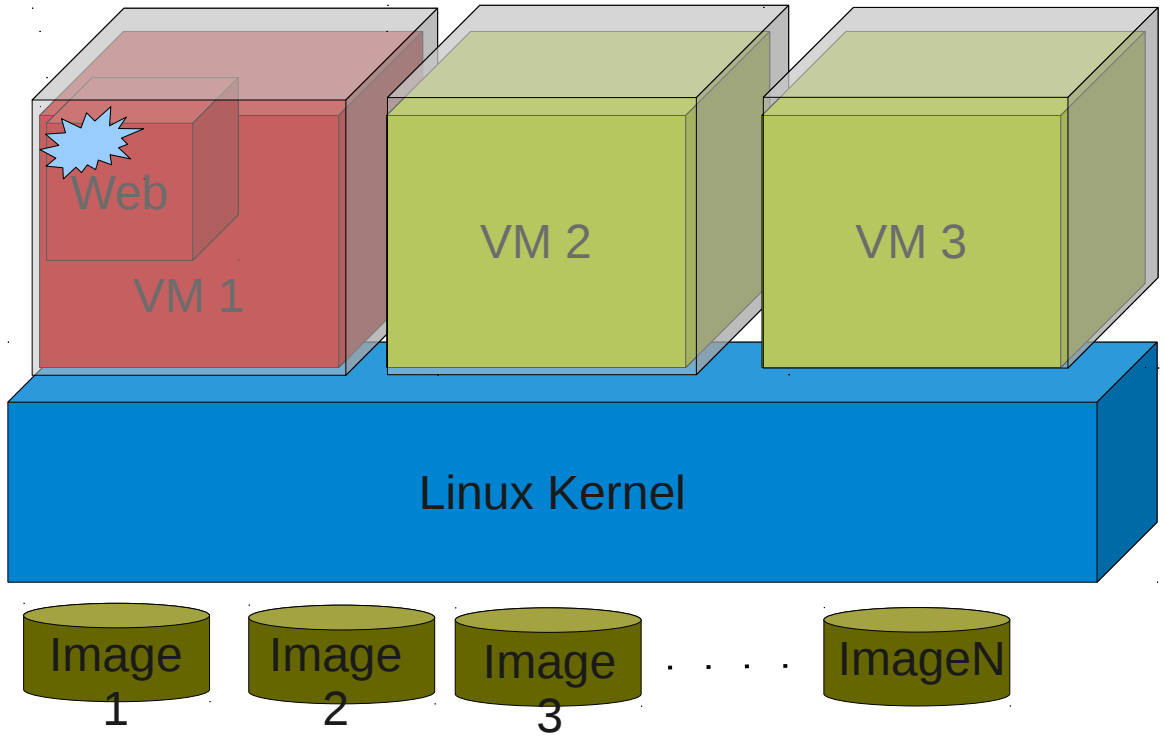

RHEL 6.5: Cryptography

- Para-Virtual Random Number Generator (RNG)
  - Provide true randomness in the guest for cryptographic purposes
  - RHEL host or RHEV Hypervisor feeds entropy to the virtual machines
  - Helps alleviate entropy starvation in guests

Certification status

http://www.redhat.com/solutions/industry/government/certifications.html

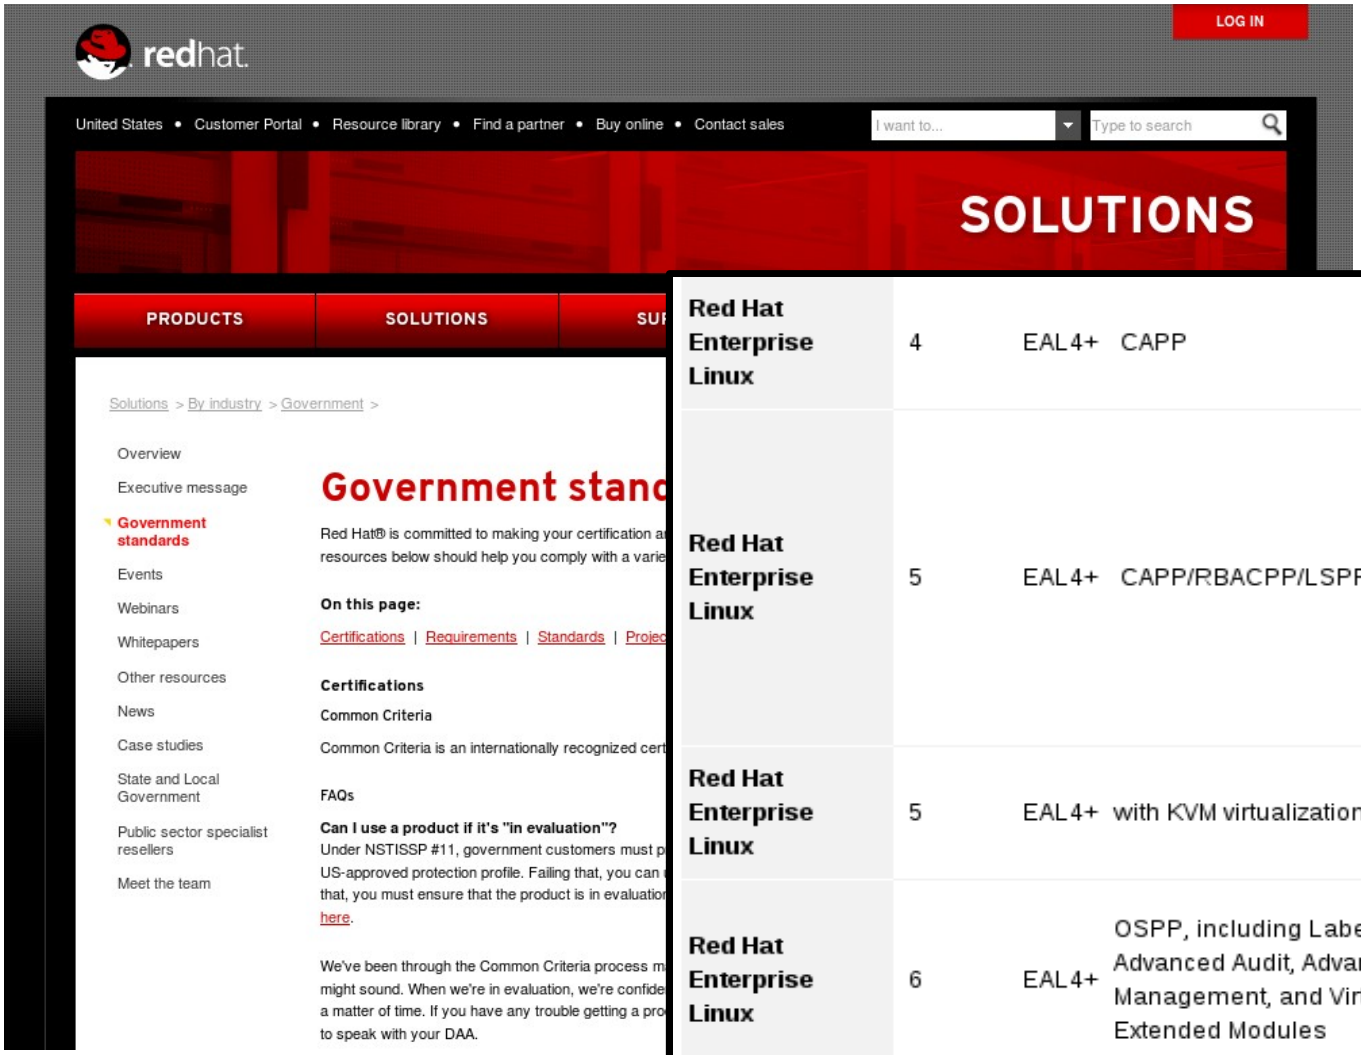

Red Hat Enterprise Linux	4	EAL4+	CAPP	IBM (report, target) [PDFs]	Evaluated
Red Hat Enterprise Linux	5	EAL4+	CAPP/RBACPP/LSPP	Dell (report, target) HP (report, target) IBM (report, target) SGI (report, target) [PDFs]	Evaluated
Red Hat Enterprise Linux	5	EAL4+	with KVM virtualization	IBM (report, target)	Evaluated
Red Hat Enterprise Linux	6	EAL4+	OSPP, including Labeled Security, Advanced Audit, Advanced Management, and Virtualization Extended Modules	Dell HP IBM SGI (report, target)	Evaluated
Red Hat Enterprise Linux	6	EAL4+	OSPP, including Labeled Security, Advanced Audit, Advanced Management	IBM (report, target)	Evaluated

Mission Critical

Mission Critical Systems / Software

Mission critical refers to any factor of a system (software, hardware) whose failure will result in the failure of business operations

KVM is hardened to run mission critical workloads

Integrated Virtualization

KVM leverages RHEL kernel, trusted for mission critical workloads

Hardware Abstraction

KVM has the highest single guest scalability (vCPU = 160, vMemory = 2TB)

Security

EAL4+ Certification (1) plus SE Linux enabling Mandatory Access Control between virtual machines

Server Consolidation

KVM runs heavy IO workloads such as large databases and holds the **Top 6/11** virtual machine consolidation scores on SPECvirt (2)

Resource Management

Cgroups helps manage resources in virtual environments too

(1) Source: http://www.redhat.com/solutions/industry/government/certifications.html
(2) Source: SpecVirt_sc2010 results: http://www.spec.org/virt_sc2010/results/specvirt_sc2010_perf.html

What else?

RHEL 6.4 Features

RHEL6.4: virtio-scsi

New storage architecture for KVM!

- virtio-scsi device = SCSI host bus adapter (HBA)
- Allows arbitrary number of devices per guest
  - Virtual hard drives and CDs
  - Pass-through physical SCSI device
- Supports SCSI pass-through and SCSI reservations
- Rich features – depends on the target, not virtio-scsi
- Drop-in physical disk replacement, friendlier for P2V and V2V
- RHEL6.4 guests and Windows guests (excluding XP)

RHEL6.4: virtio-blk “data-plane” mode

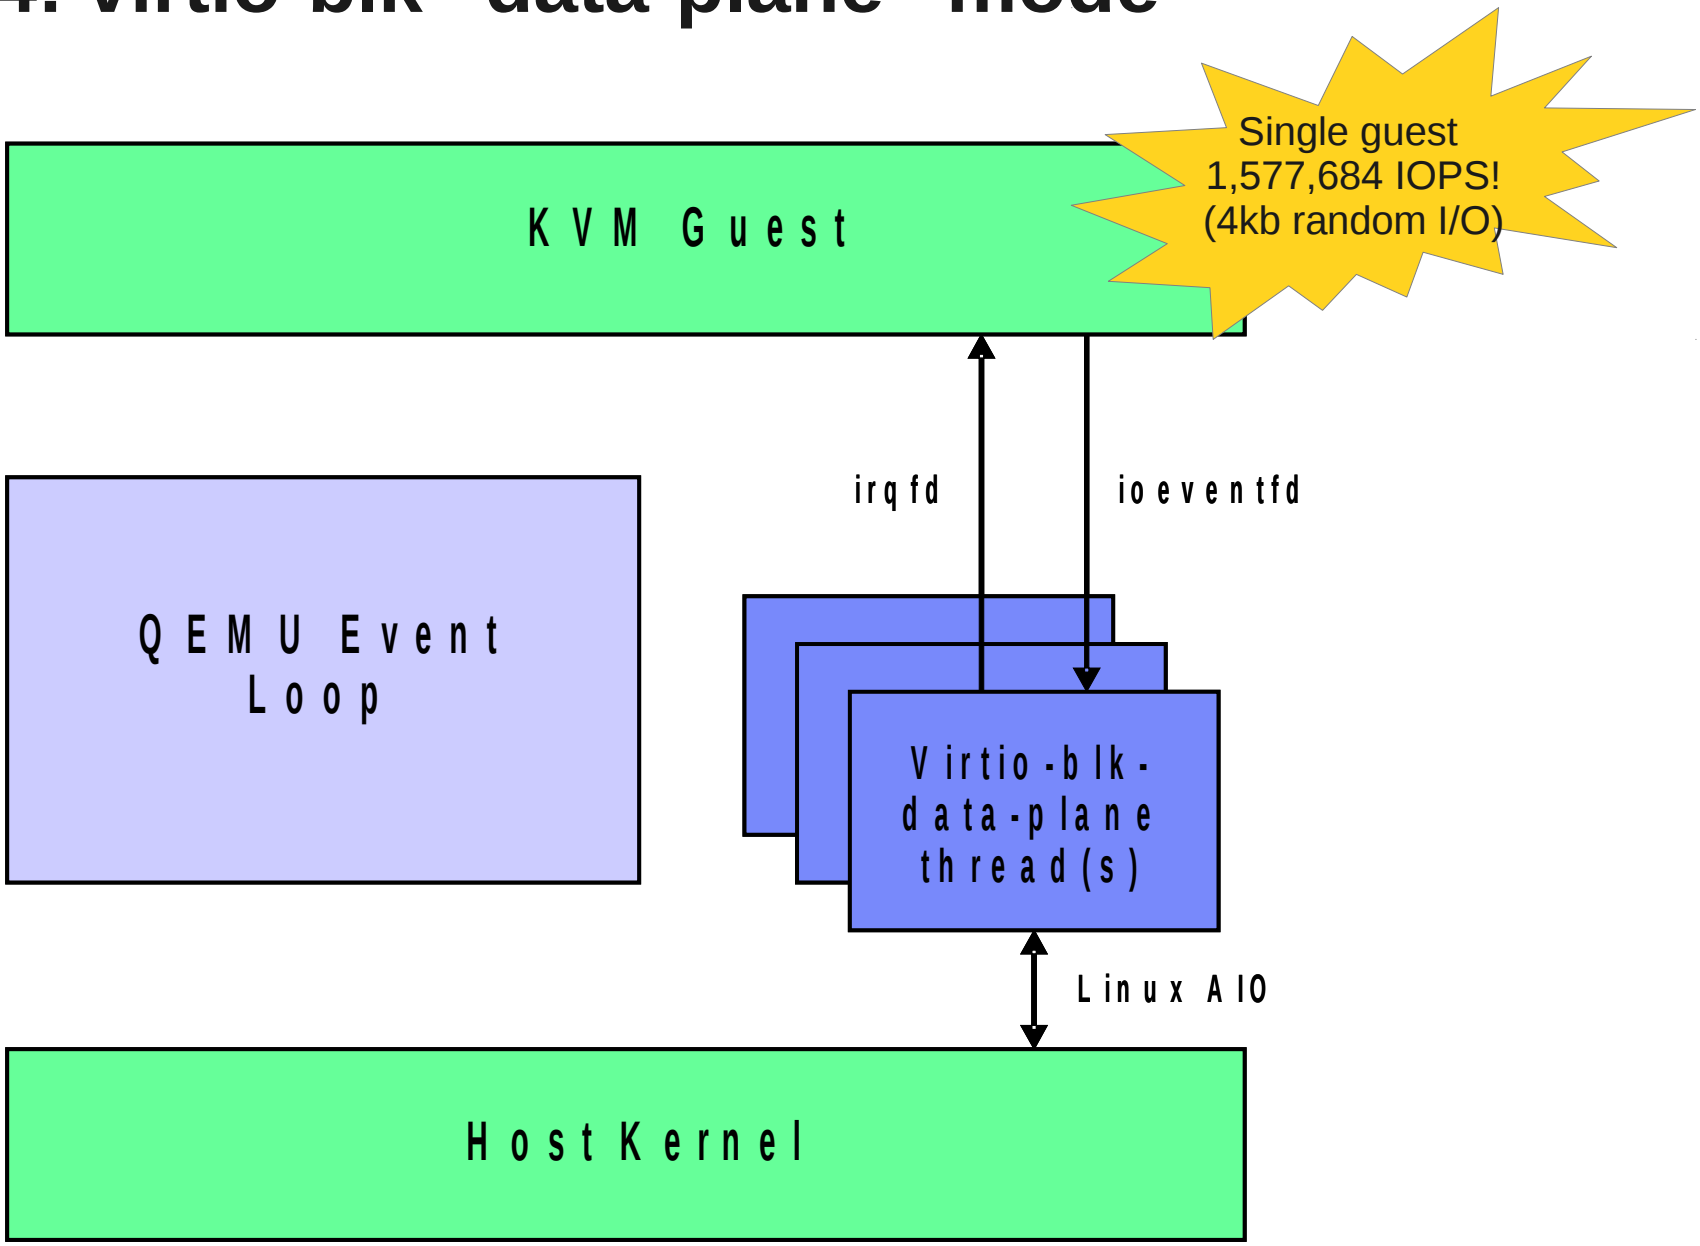

RHEL6.4: Para-virt End-of-Interrupt (PV-EOI)

Improved performance!

- Optimization for interrupt-intensive workloads
  - Up to 10% less CPU usage in some scenarios
- Reduces the number of context switches between the VM and the hypervisor.

Works out-of-the-box with all I/O types

- Particularly useful for high incoming network traffic
- Guest OS = RHEL 6.4 (PV calls in guest kernel)

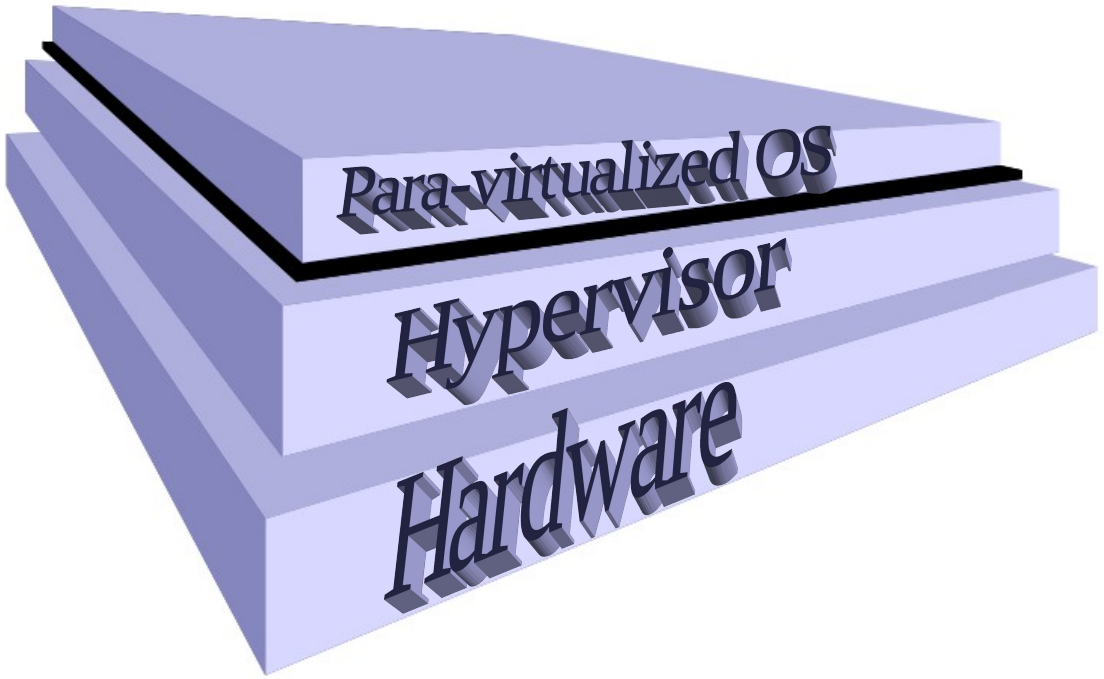

RHEL6.4: New Virtualized CPU models

- **New virtualized CPU models – Intel and AMD**
  - Intel 4th generation Intel Core processors (Haswell, Ivy Bridge)
  - AMD Opteron Series 6300 (Abu Dhabi, Seoul)
  - Performance!
    - Leverage new processor features
    - Use new instructions
  - New CPU model definitions in KVM
    - RHEL host support – new virt features
    - Virtualized guest benefits, too

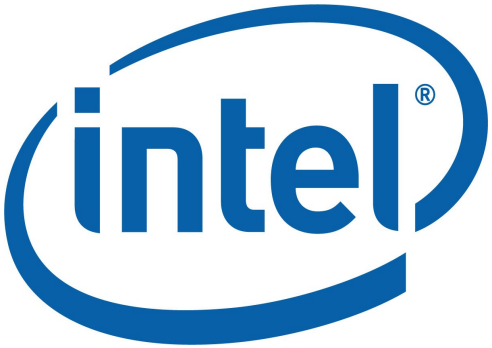

RHEL6.4: USB 2.0 Improvements

- USB 2.0 redirection – with Spice
- USB 2.0 Live Migration Support
  - System admins can relocate VDI desktops for efficient load balancing, transparent to the end user!
- Migration of VMs with USB devices attached - mostly useful for “migration to file” (save VM)

Reliability Availability Serviceability

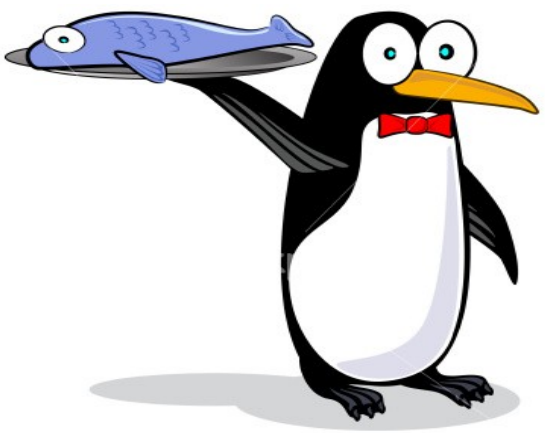

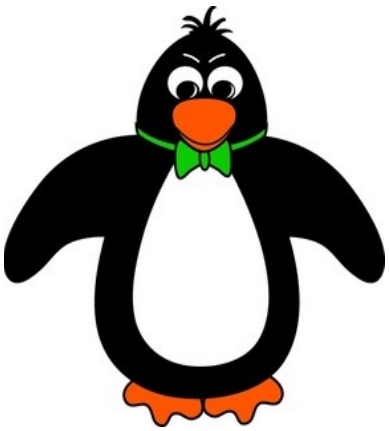

Virtual Reliability, Availability, Serviceability (vRAS)

What is vRAS?

- Maintain the Service Application Level Agreement (SLA)
- Meet workload demands, spikes at peak hours
- Maintain running virtual services and applications
- Perform live operations on running VMs and storage with no downtime

Virtual RAS Features

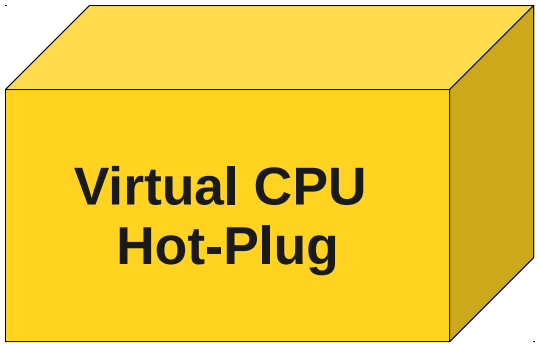

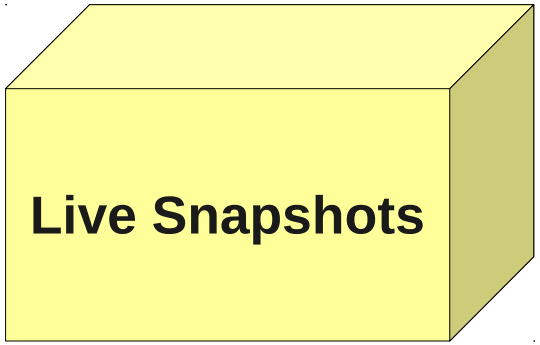

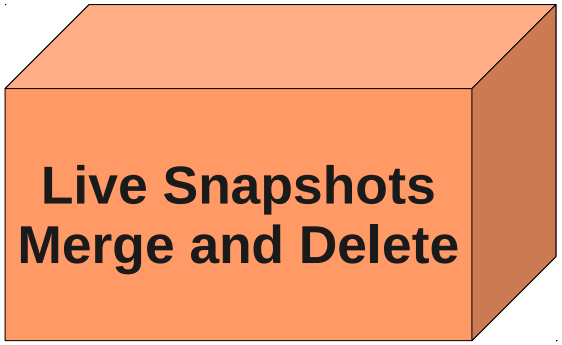

RHEL 6.5: Virtual CPU Hot-Plug

- Flexibility - add additional vCPUs to running VMs
- No application downtime to adjust VM's compute capacity

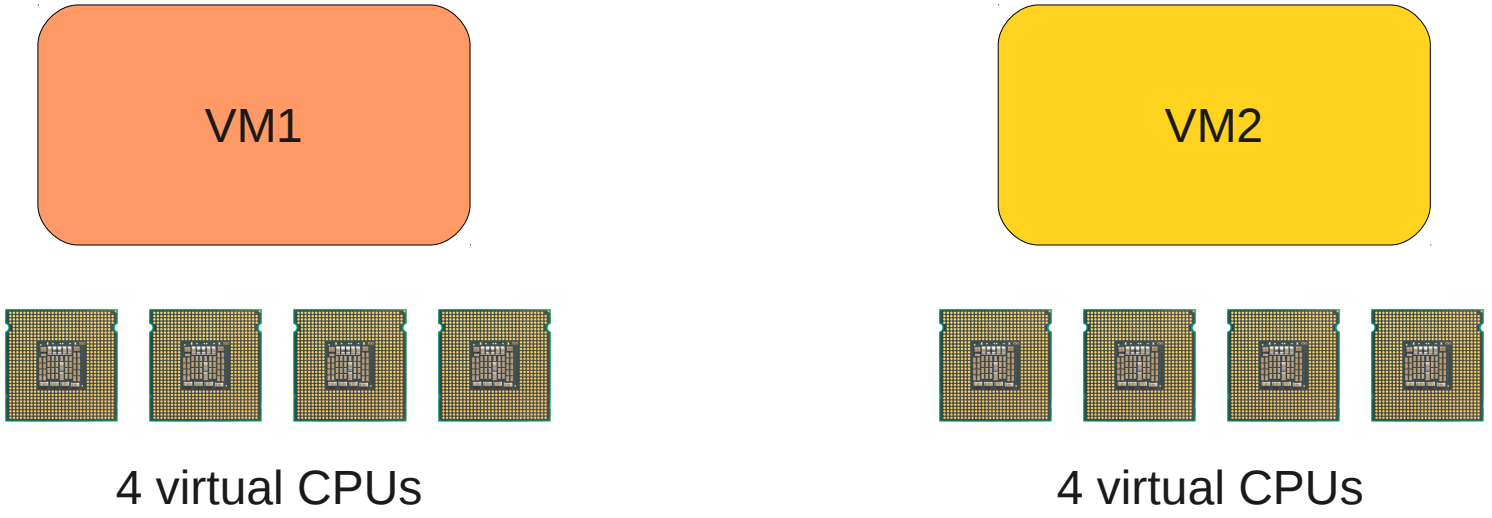

RHEL 6.5: Virtual CPU Hot In-Plug

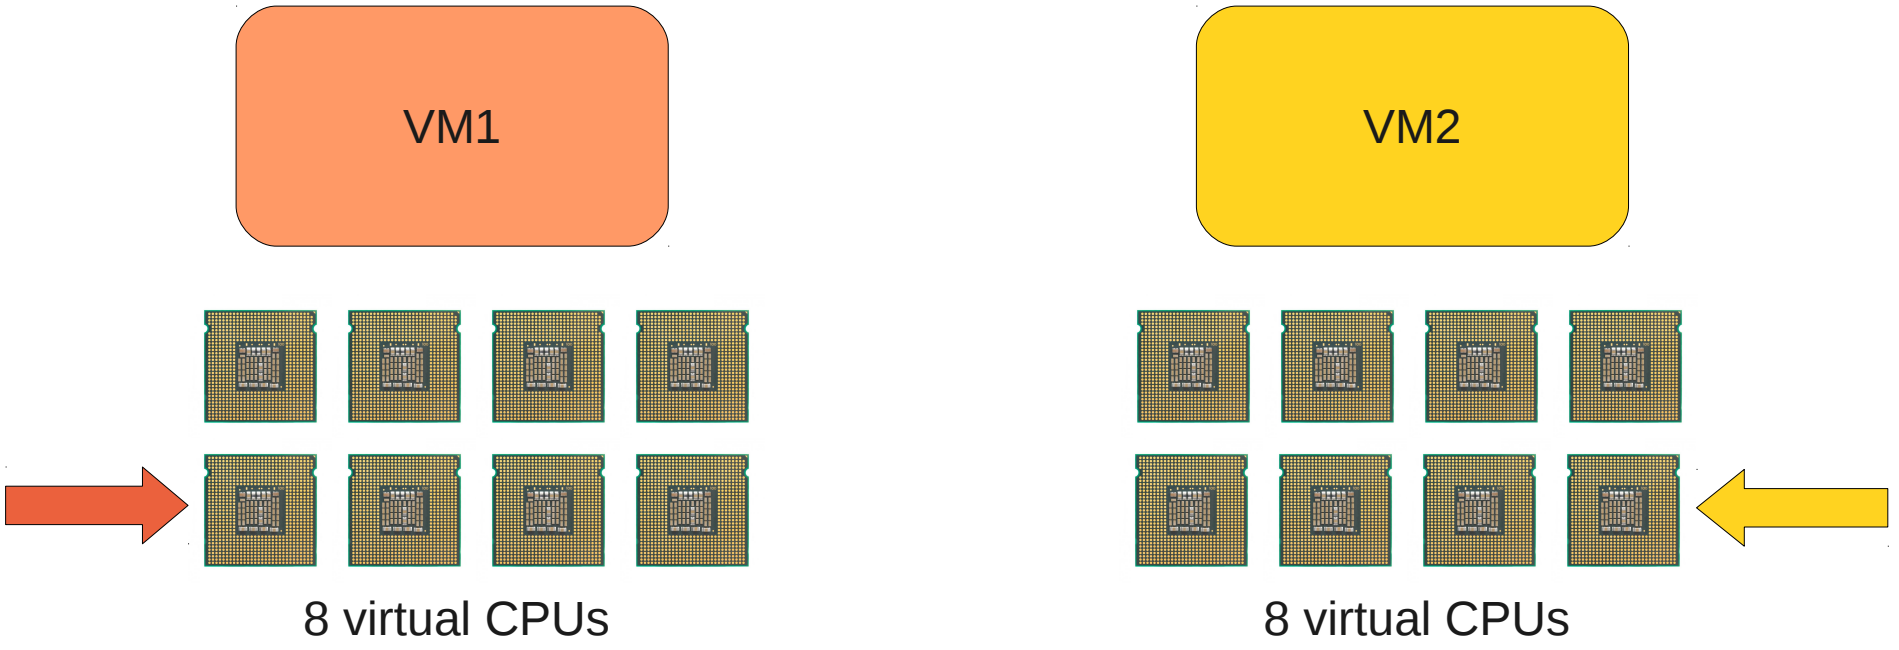

RHEL 6.5: Virtual CPU Enable/Disable

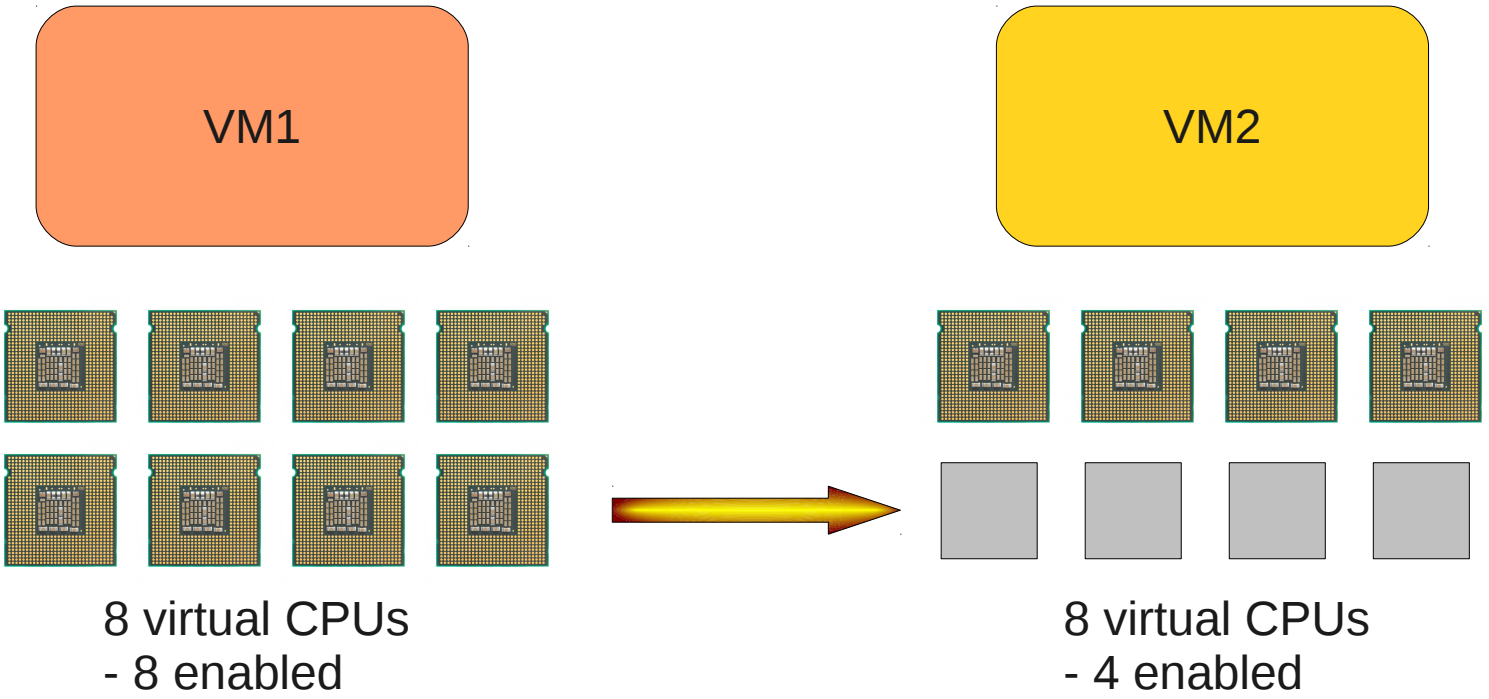

RHEL 6.5: Virtual CPU Enable/Disable

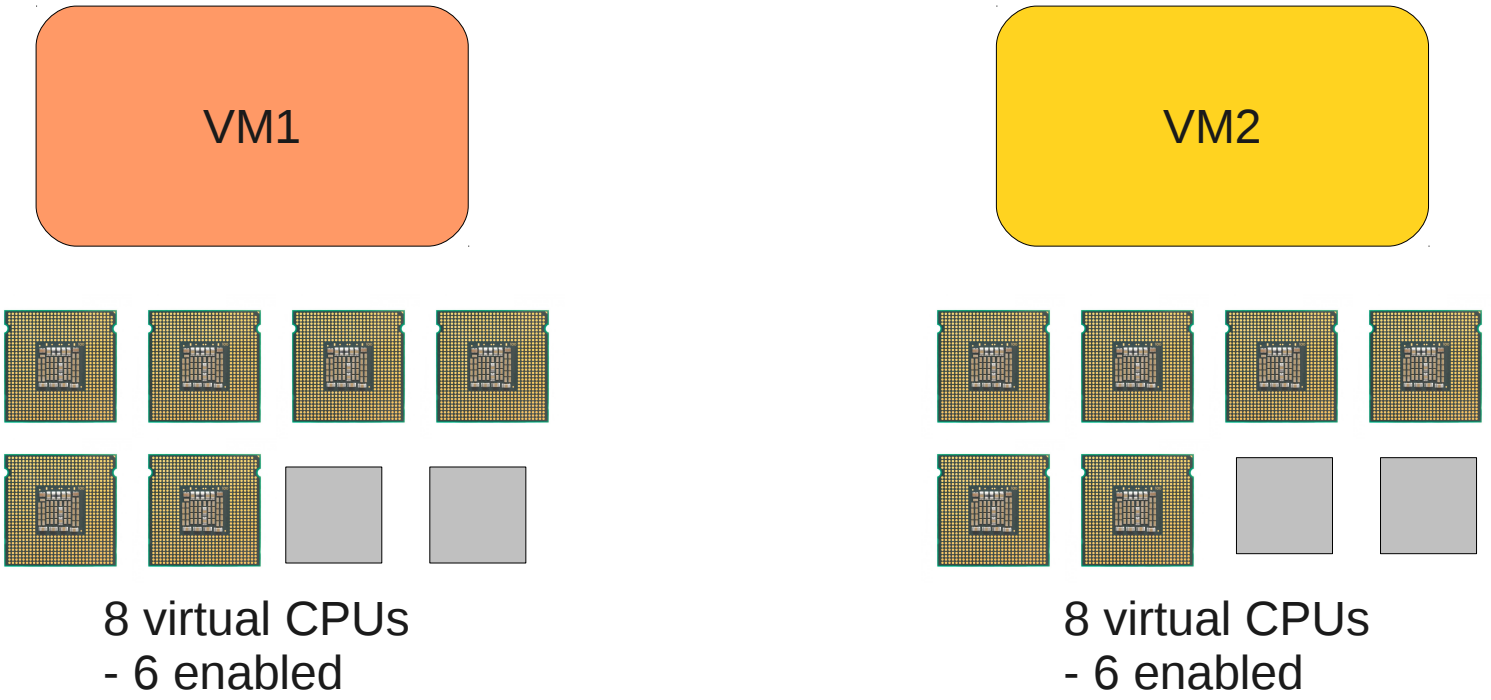

KVM-enabled features in RHEV 3.1

RHEV 3.2 Storage data-center features
Live Snapshots, Merge & Delete

- Snapshot a VM while the guest is running, preserving state and data of the VM at a given point in time
- Sample use cases:
  - Data-center admin saves snapshot of a running VM prior to disruptive upgrades to system
  - Backup scenario with periodic incremental snapshot/backup

Simple image chain:

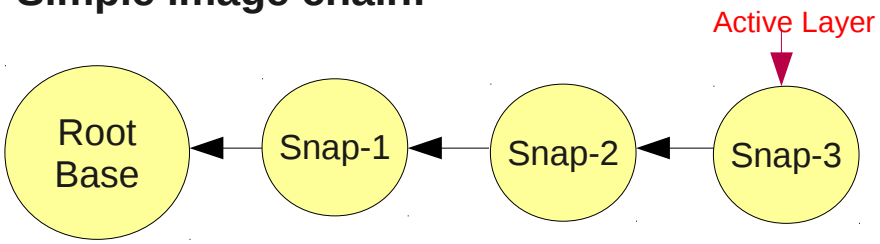

Merge toward root:

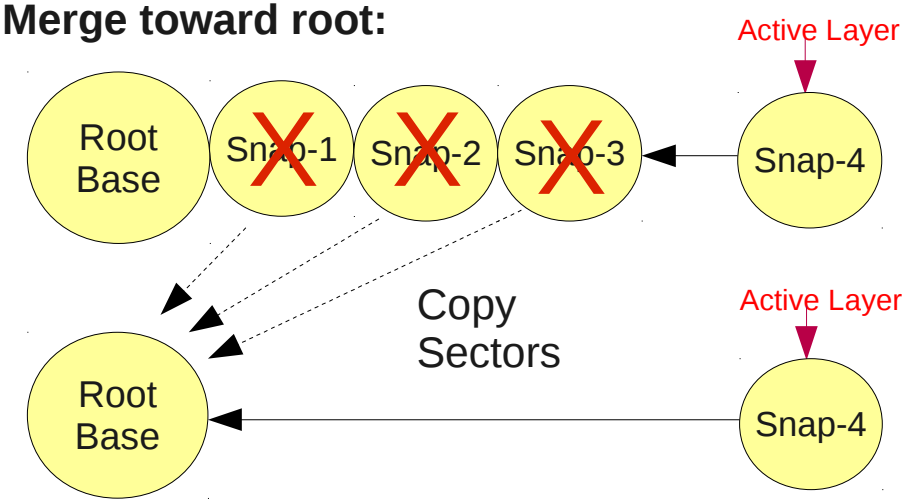

Merge toward active layer:

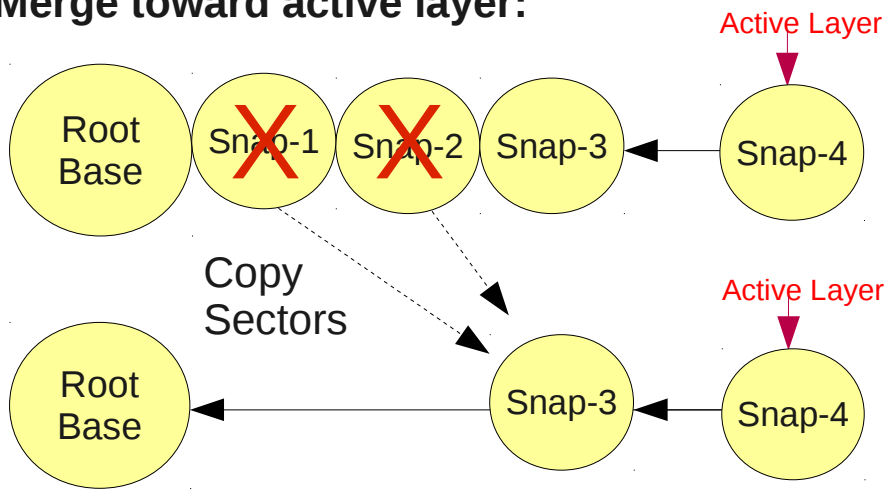

RHEV 3.1 Storage data-center features
Live Storage Migration

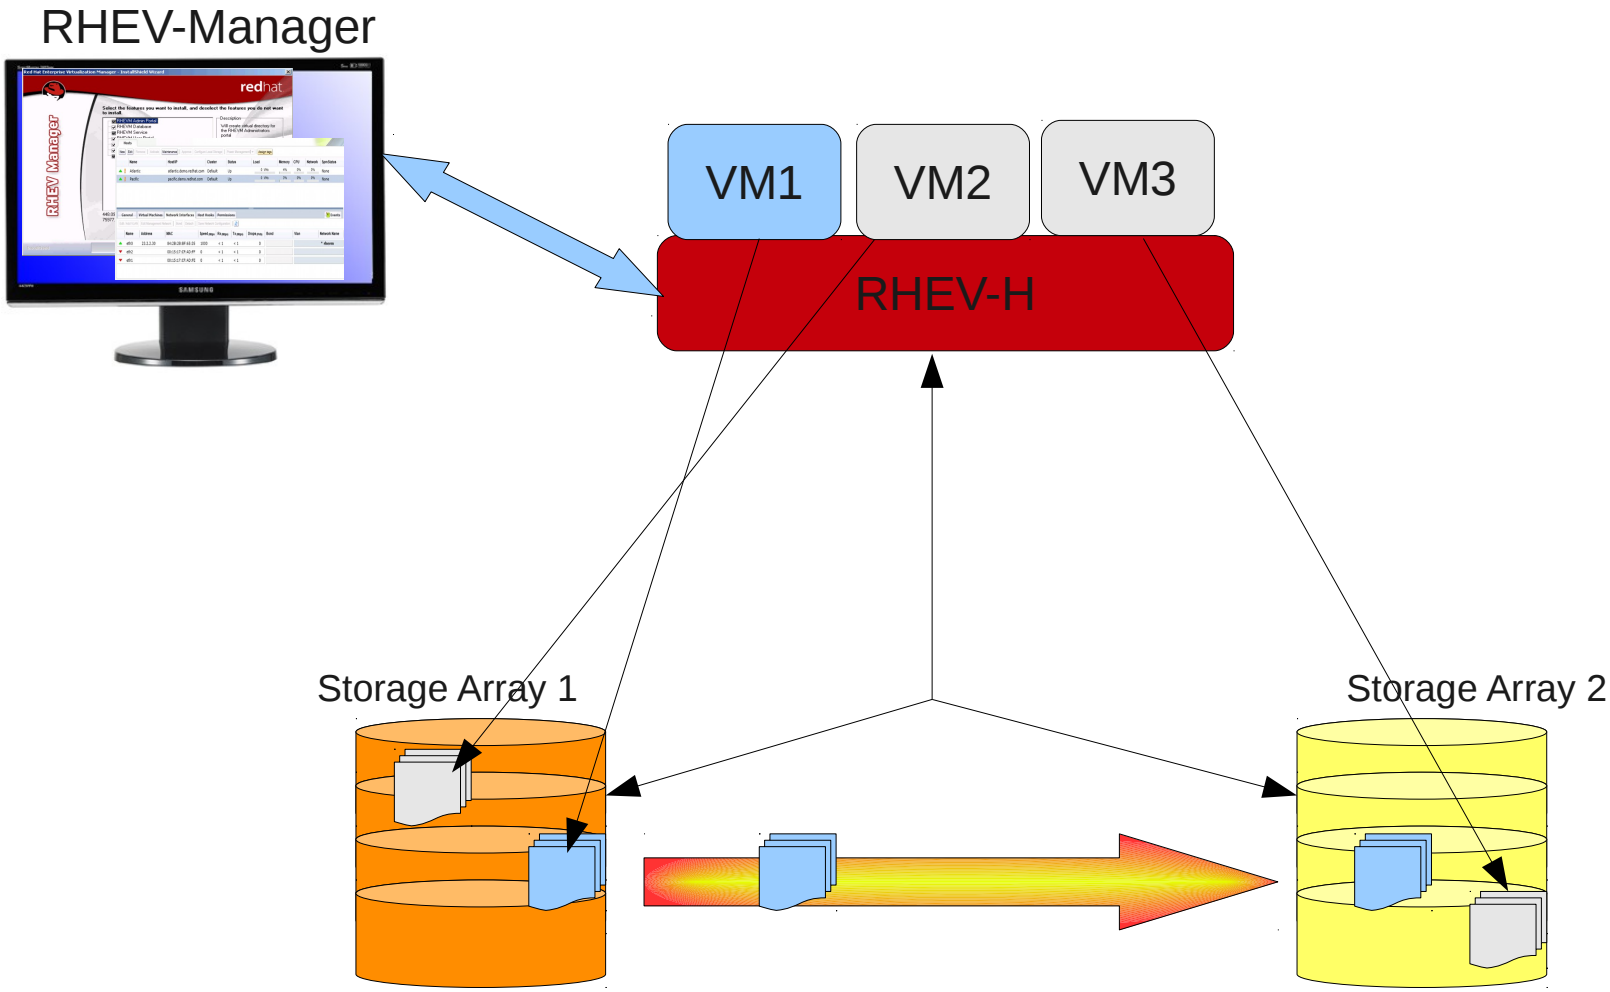

Migrate a VM's storage files across storage arrays – no application downtime!

Let's have a look at
Red Hat KVM Futures

KVM Futures – 2H 2013

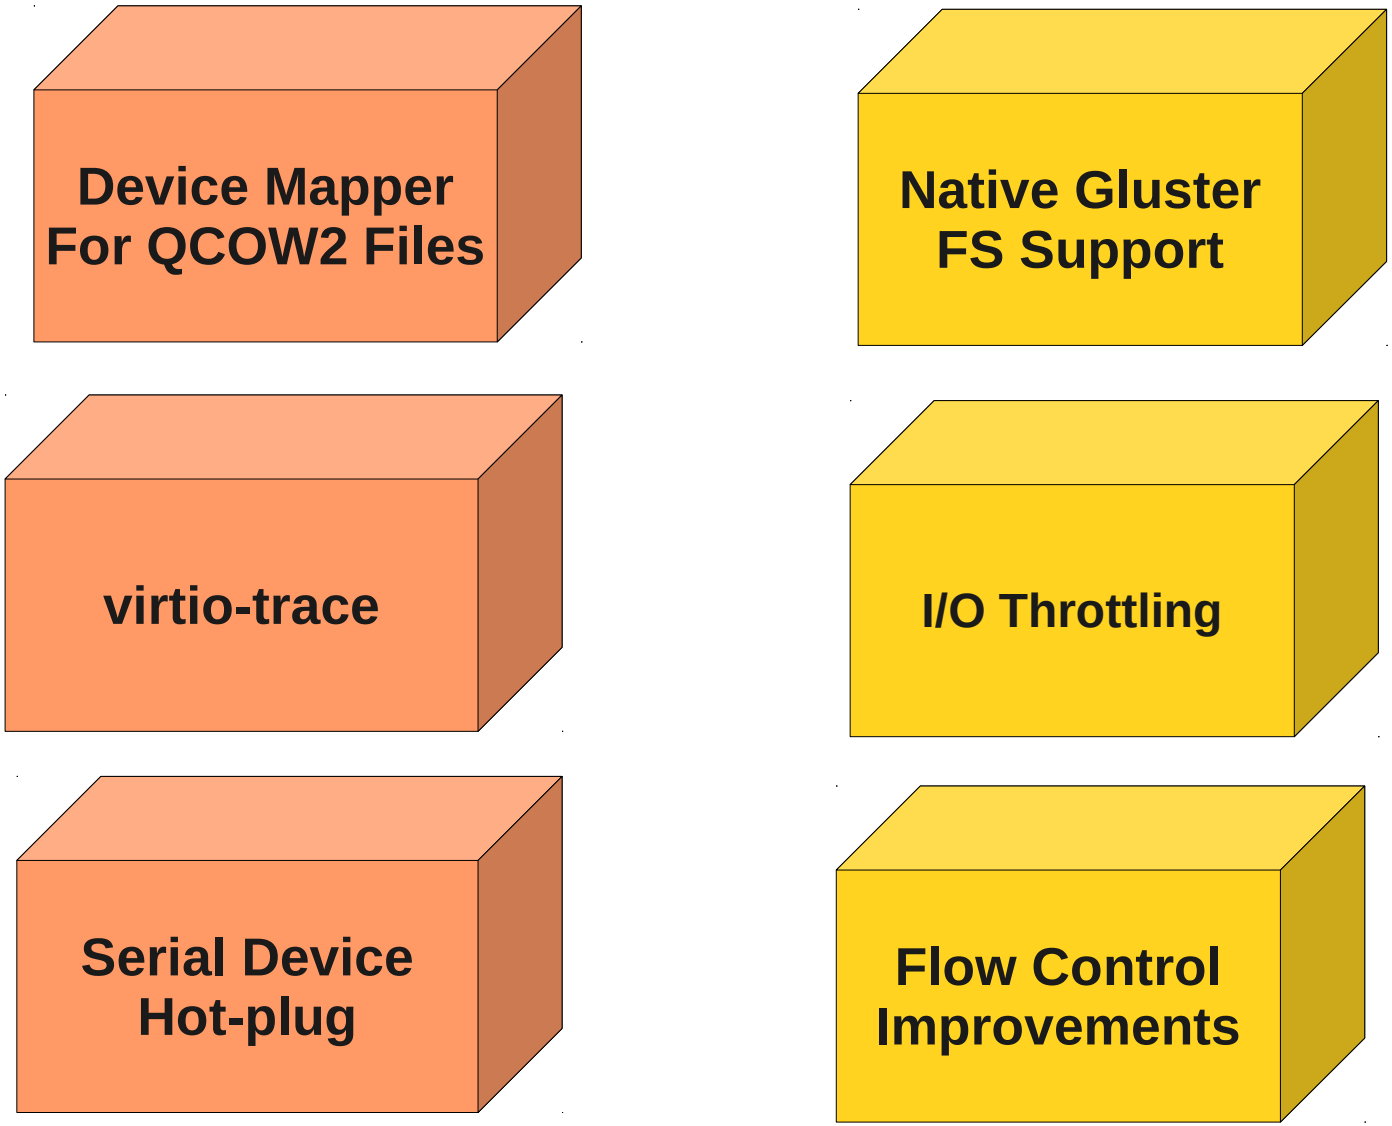

KVM Futures – 2014

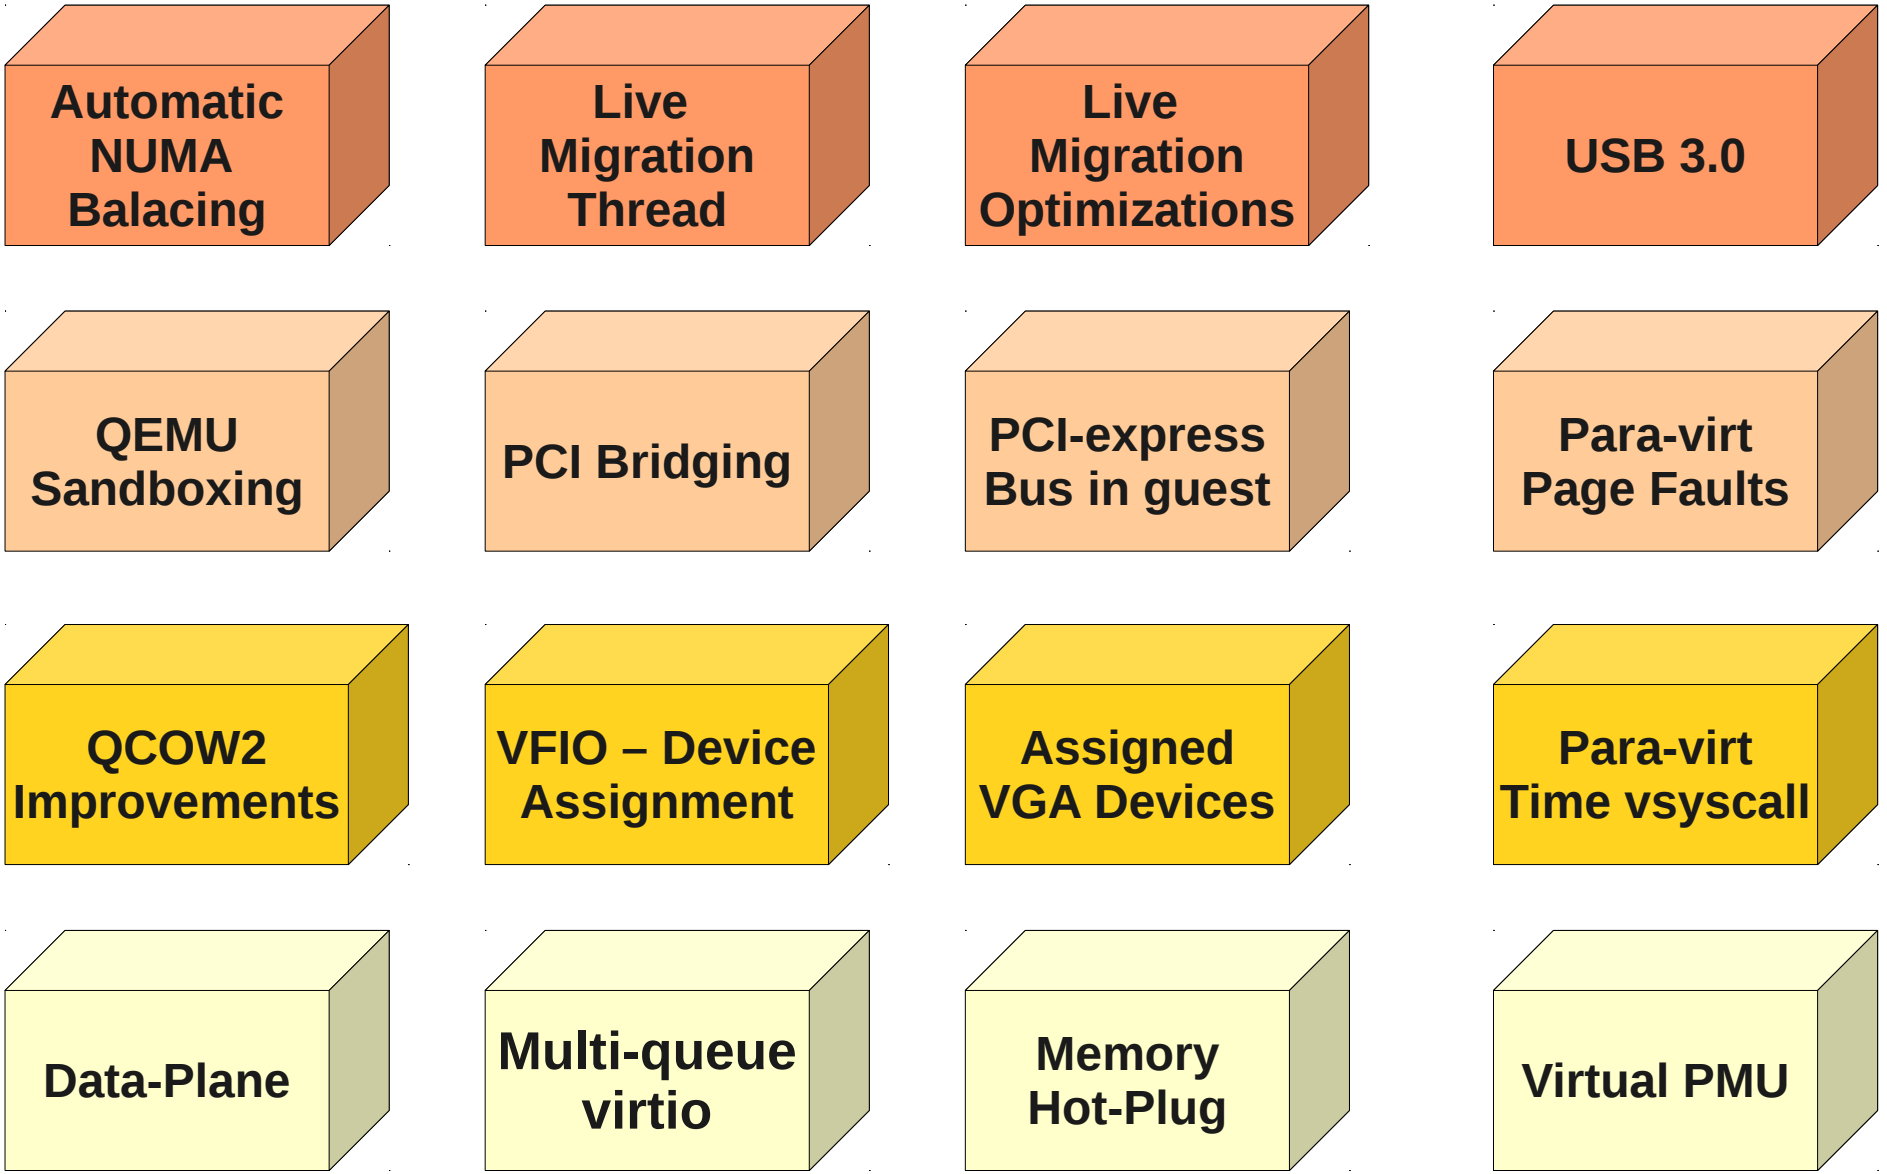

KVM Community and Ecosystem

KVM : Strong hypervisor foundation for RHEV, Open Stack, IBM Cloud and many others

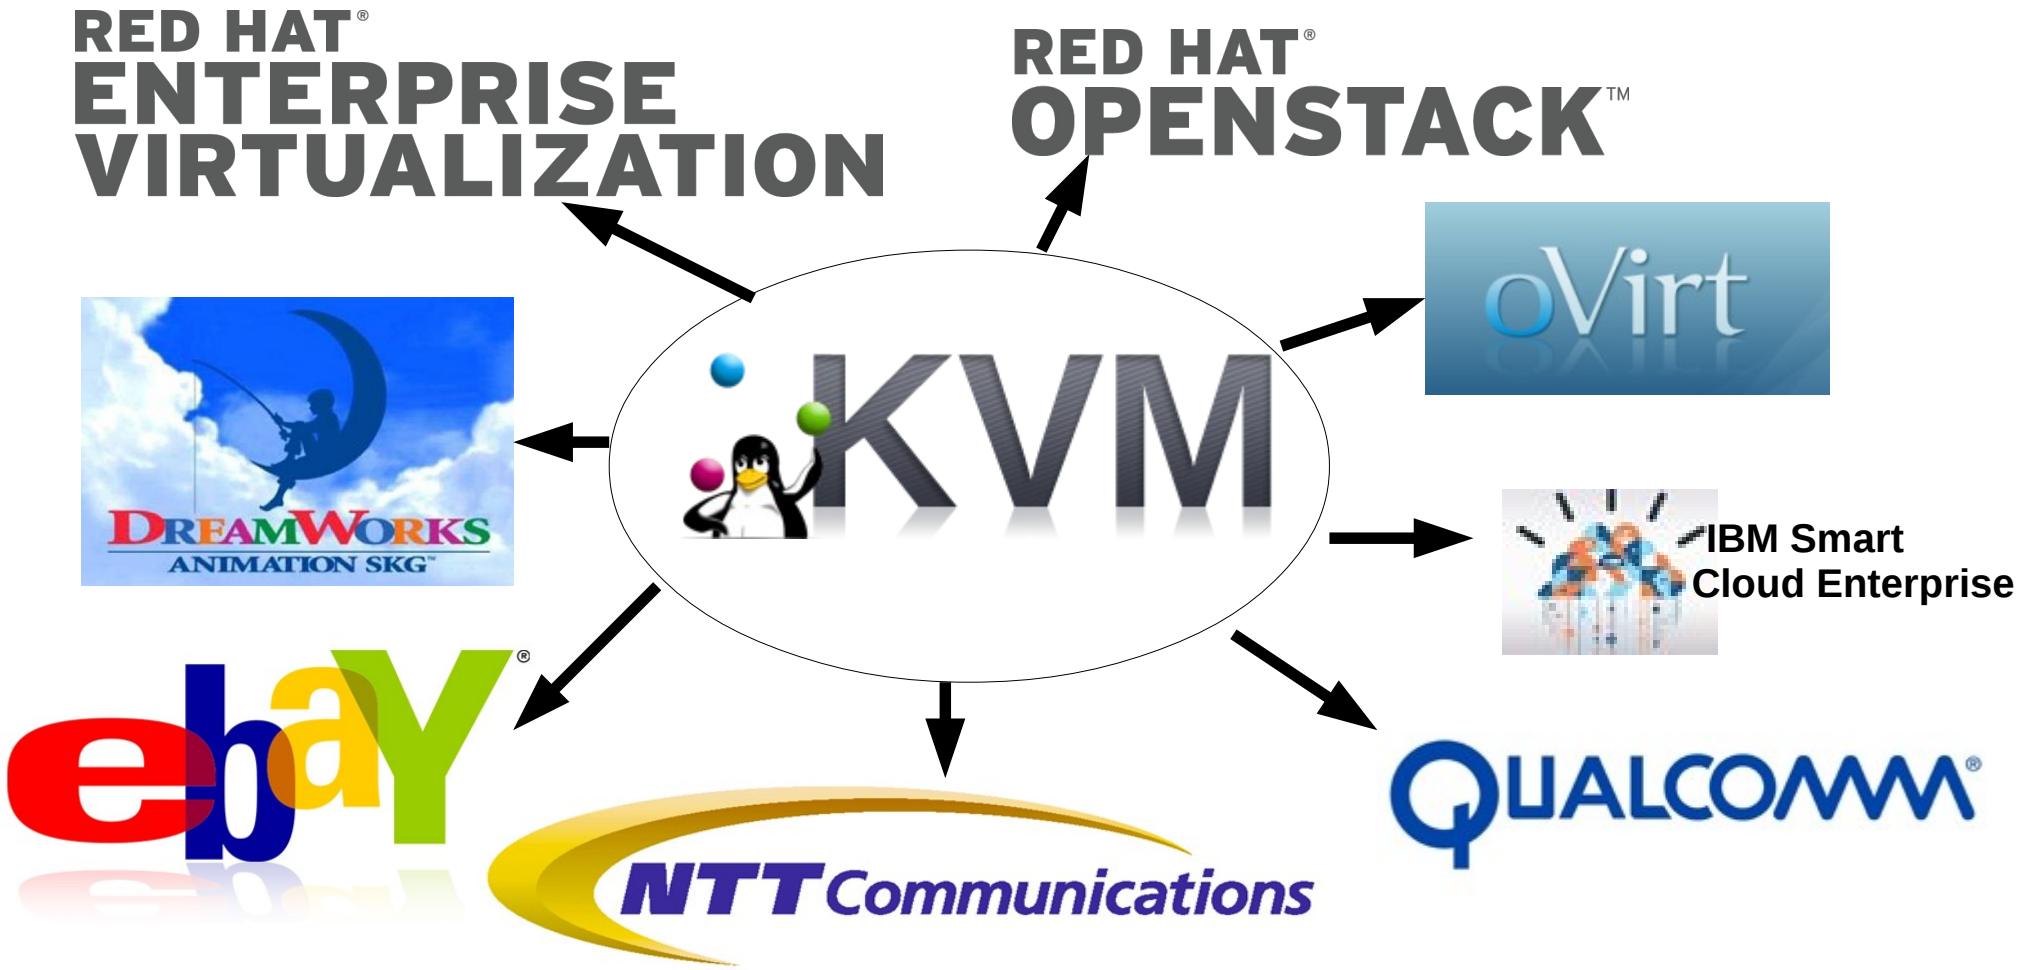

Open Virtualization Alliance

An alliance that includes leading…

- Virtualization
- Datacenter and
- Cloud Solution Providers

Governing Members

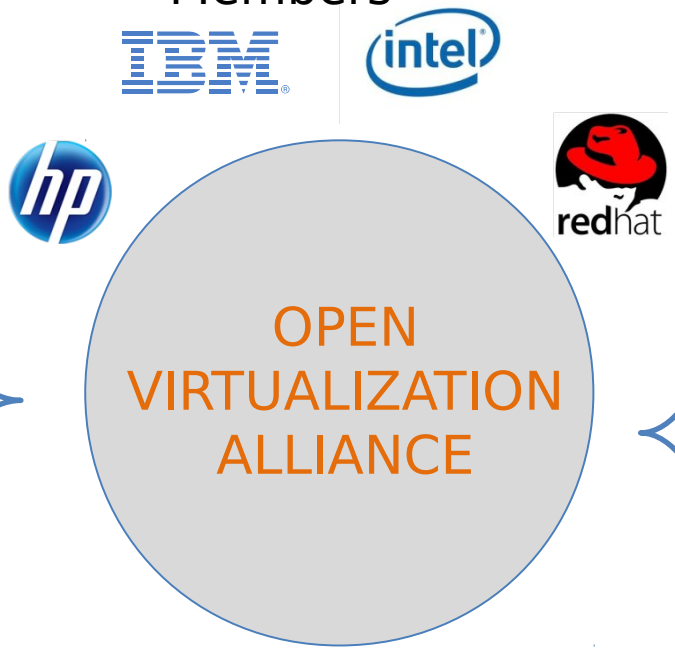

244 Members and counting!

to help…

- Increase overall **awareness** of KVM
- Drive the **adoption** of KVM based solutions
- Foster an **ecosystem** of third-party solutions around KVM
- Promote **Best Practices** and highlight **Customer Successes**

Bringing the Community, Vendors and Users Together

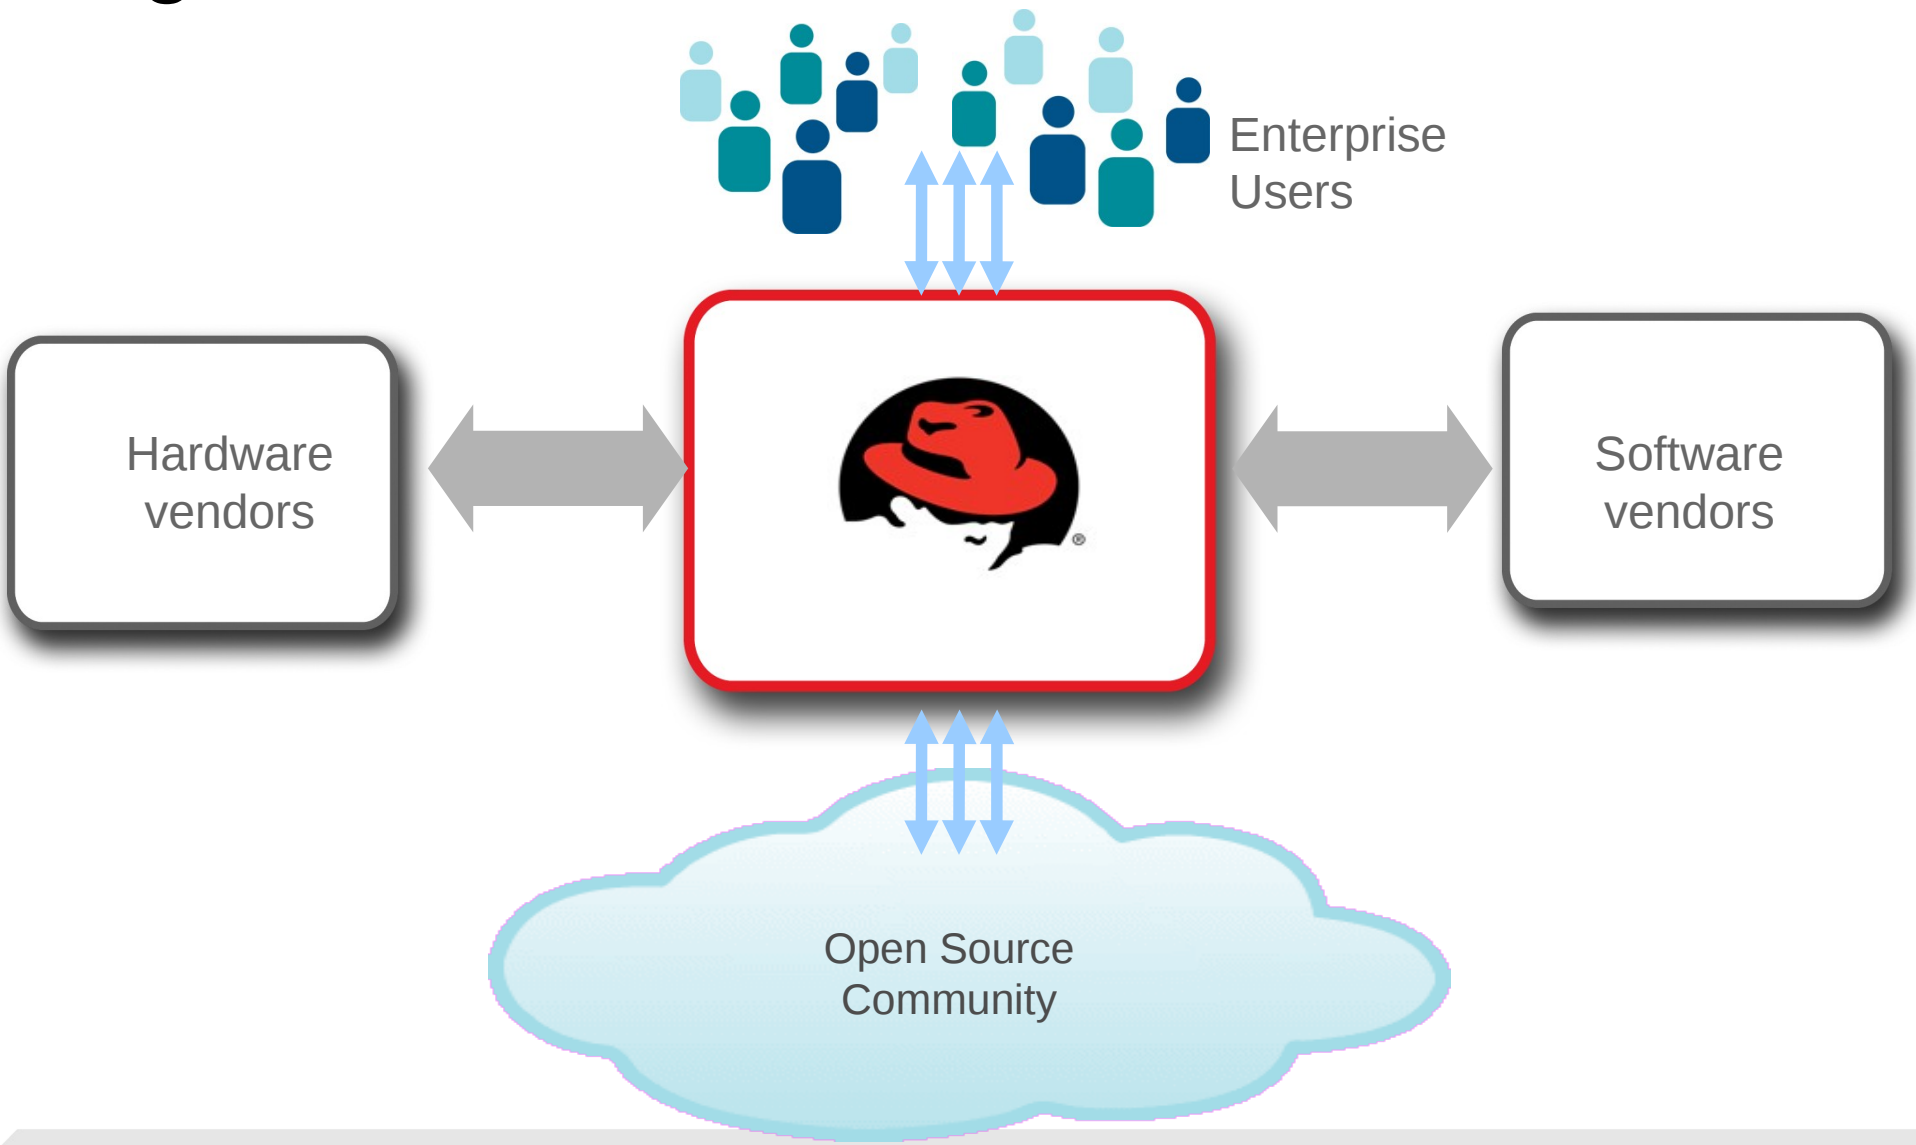

What makes KVM the best Virtualization platform?

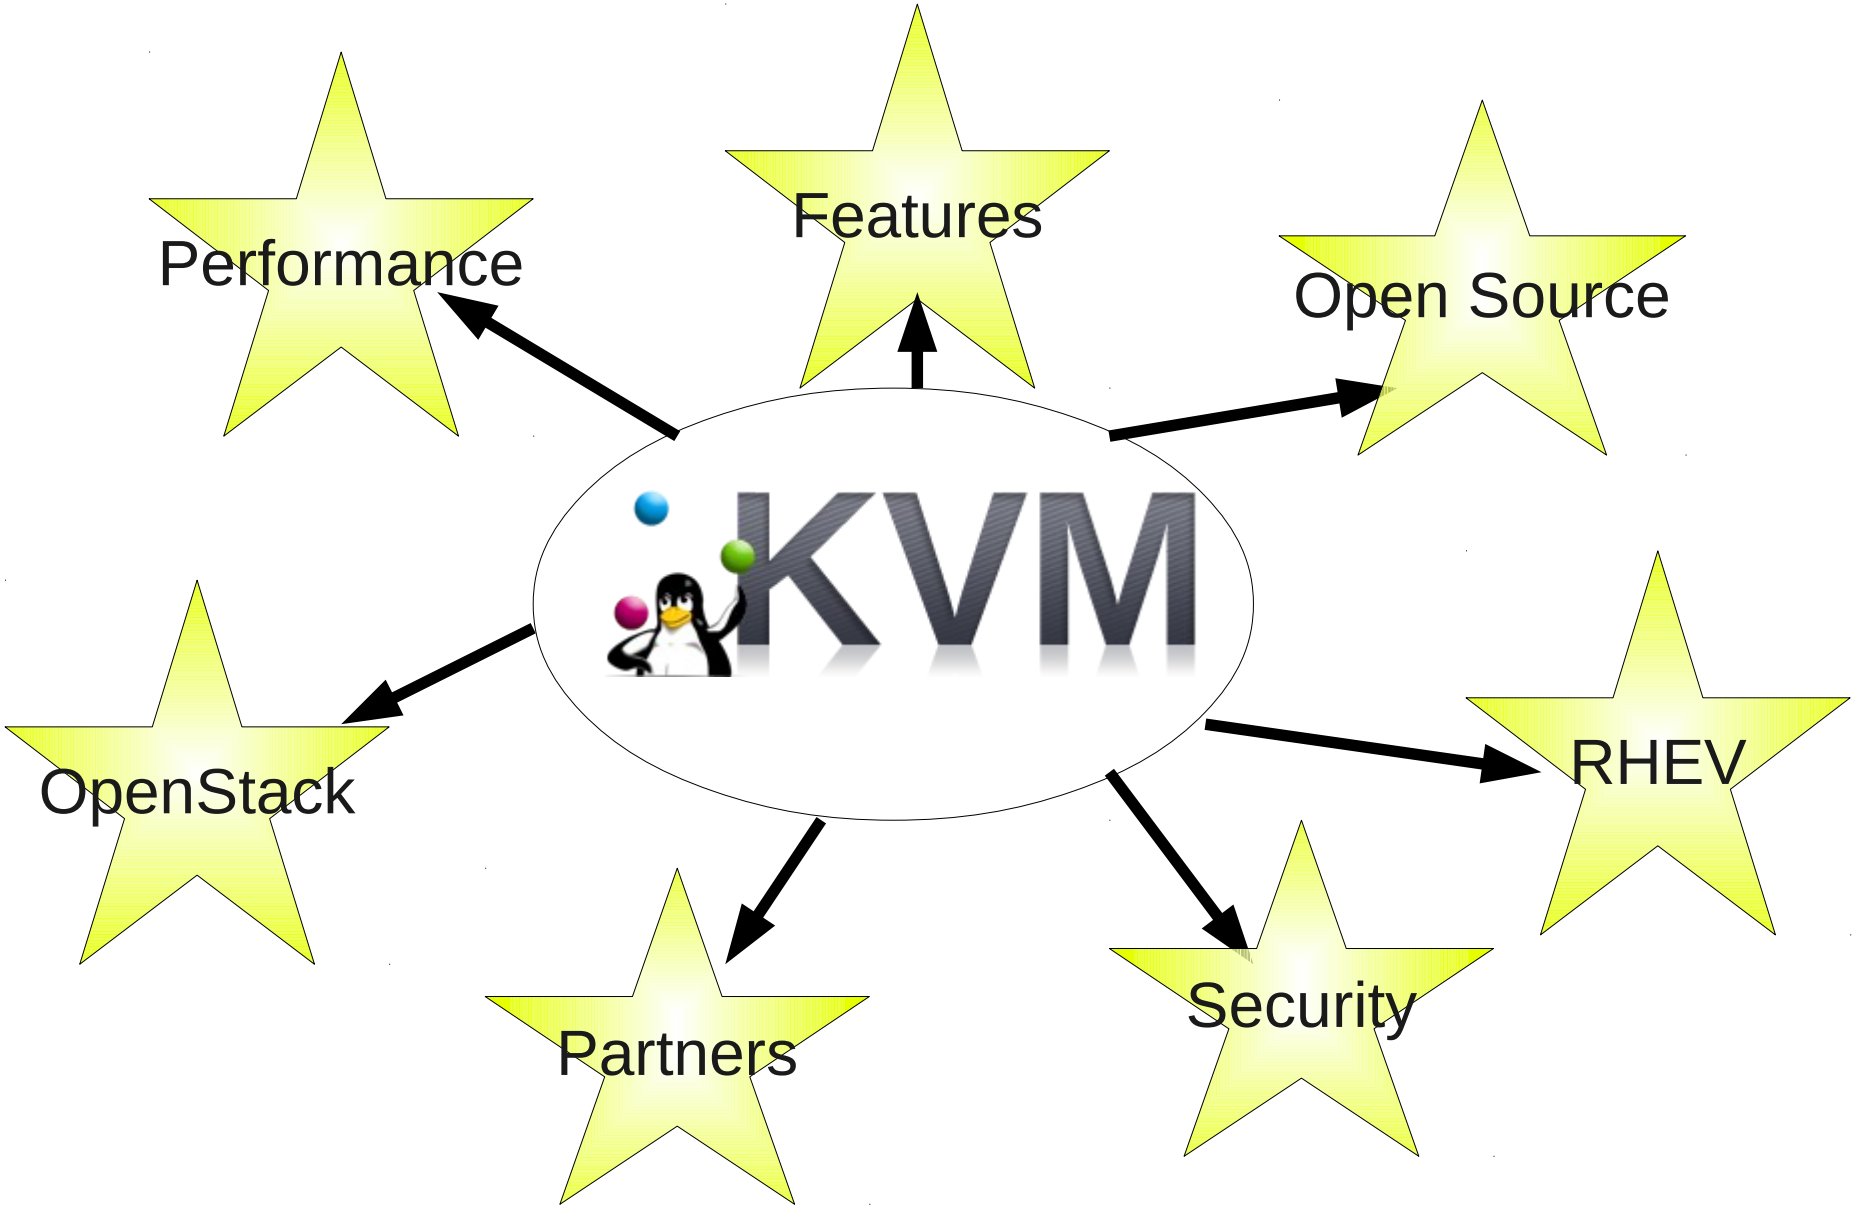

Questions?

Related Summit Sessions

- Hypervisor Technology Comparison & Migration
  - Fri 9:45, Room 313
- Partner Pavilion
  - Demo of the latest KVM features
  - In the RHEV booth

Thank you