



RHCE Loopback Virtual Meetup

- December 2, 2010

Welcome!

- Participating from Red Hat today:
 - Tim Burke, Vice President of Platform Engineering
 - Jan Mark Holzer, Senior Consulting Engineer
 - Gunnar Hellekson, Chief Technology Strategist, Public Sector
 - Erich Morisse, Solutions Architect
 - Randy Russell, Director of Certification
- Agenda
 - Some quick program updates
 - RHEL 6 Technical Overview
 - Discussion

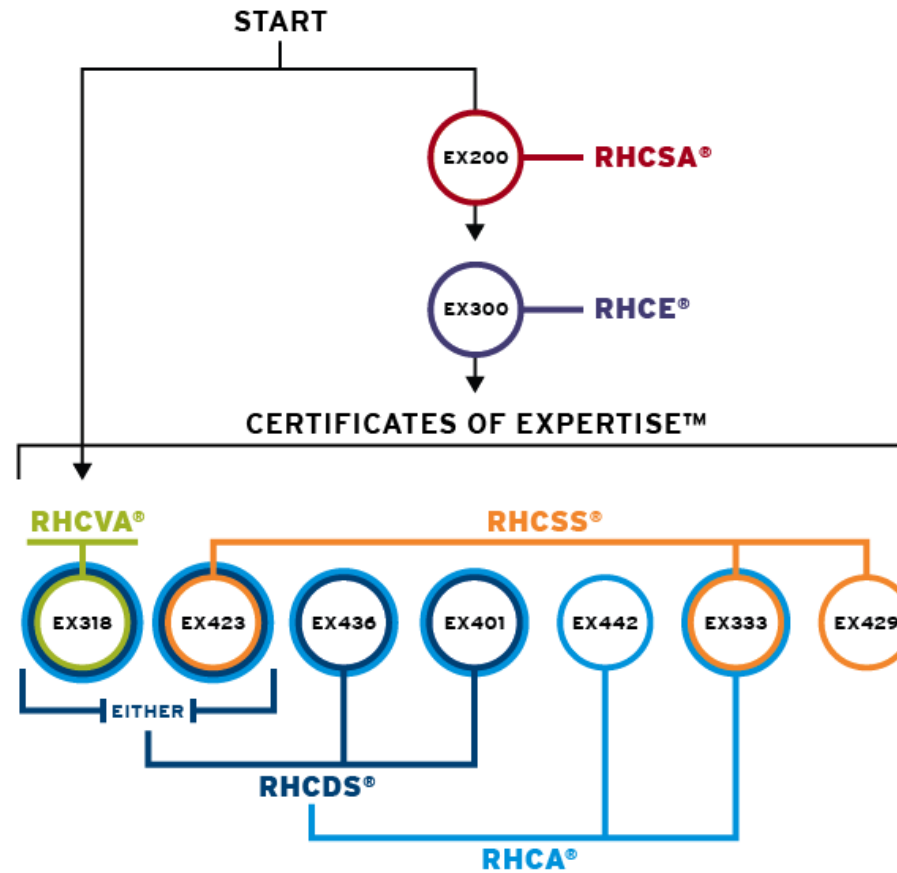


Events, etc.

- San Jose RHCE Loopback Event: December 16
- Interested in a Loopback in your area?
 - Contact us at rhceloopback@redhat.com!
- Upload your video:
 - http://www.youtube.com/results?search_query=rhceloopback



Changes to the Certification Program





Red Hat Enterprise Linux 6

Technical Overview

Tim Burke

December 2010




Major Themes



Optimized for today's IT deployments

Scale – Performance
Security – Resource control
Manageability




Virtualization

Host and Guest optimization
Cloud Foundation



Green IT

Power Management



Lifecycle

Long term stability
Future proof



Facts and Figures



Red Hat Enterprise Linux 6

Released November 10th, 2010
First new major release for 3 ½ years

Represents more than 600 person
years by Red Hat engineers

1,821 requested features included
10,587 completed builds

37GB of content
2,058 SRPMs; 21,957 binary RPMs

Red Hat engineers are based
in 26 countries

85% more packages than
Red Hat Enterprise Linux 5

14,631: resolved issues from
partner,


847 features & fixes verified
by partner QA teams

3,900 additional kernel
enhancements to 2.6.32


Kernel based on 2.6.32 with many features from .33 & .34
Red Hat is the lead developer of kernel features
Red Hat Enterprise Linux design allows smooth integration of future features



Red Hat: Leading Development



Performance




- *Performance enhancements in every component*

Component	Feature
CPU/Kernel	NUMA – Ticketed spinlocks; Completely fair scheduler; Extensive use of Read Copy Update (RCU)
Memory	Caching and NUMA enhancements Large memory optimizations: Transparent HugePages
Virtualization	Large SMP virtual machines: guests with up to 64 CPUs Block: Async-I/O; Network: Kernel network support
Disk/Network	Per LUN flush daemons Multi-queue Network drivers




Performance: SMP Scalability



64 CPU Scalability - Stream Benchmark

Intel EX 64-cpu, 128GB, FC




- Excellent, linear scalability; minimal virtualization overhead


Note: With this h/w, at 32 CPUs socket bandwidth is saturated
64-cpu result is with hyperthreading enabled



Performance: NUMA




RHEL6 Non-Uniform-Memory-Access
(NUMA diff w/ Stream Benchmark on Intel EX)




- NUMA enhancements deliver valuable performance gains



Performance: NUMA Applications



RHEL6 NUMA Application Performance




- NUMA enhancements deliver valuable performance gains



RHEL5.5 to RHEL6 AMQP TCP/IP msg/sec


RHEL5 vs RHEL6 (preliminary)




10 Gbit Ethernet (Mellanox)



Performance: Network





Comparing RHEL5.5 Infiniband with RHEL6 10Gb with RoCE





- RoCE support is unique to Red Hat (RDMA over Converged Ethernet)





Scalability: CPUs



Scalability: Memory



Scalability: File Systems



Ext4: 16 TB XFS: 100 TB



Green IT




- *Reduces power consumption at multiple layers in the software stack*

Component	Feature
Utilities	Power audit to reduce power consumption, E.G. convert utilities to event-based algorithms
File System	Intelligent drive spin-down File metadata I/O reduction: relatime mount option
Kernel	Tickless kernel enables extended low power states for idle systems
CPU	Core/CPU idling in lightly loaded SMP systems; applies for virtual guests
I/O	Dynamic power adjustment to PCIe & SATA links via ASPM & ALPM



Green IT: Power Management




Green IT: Tools



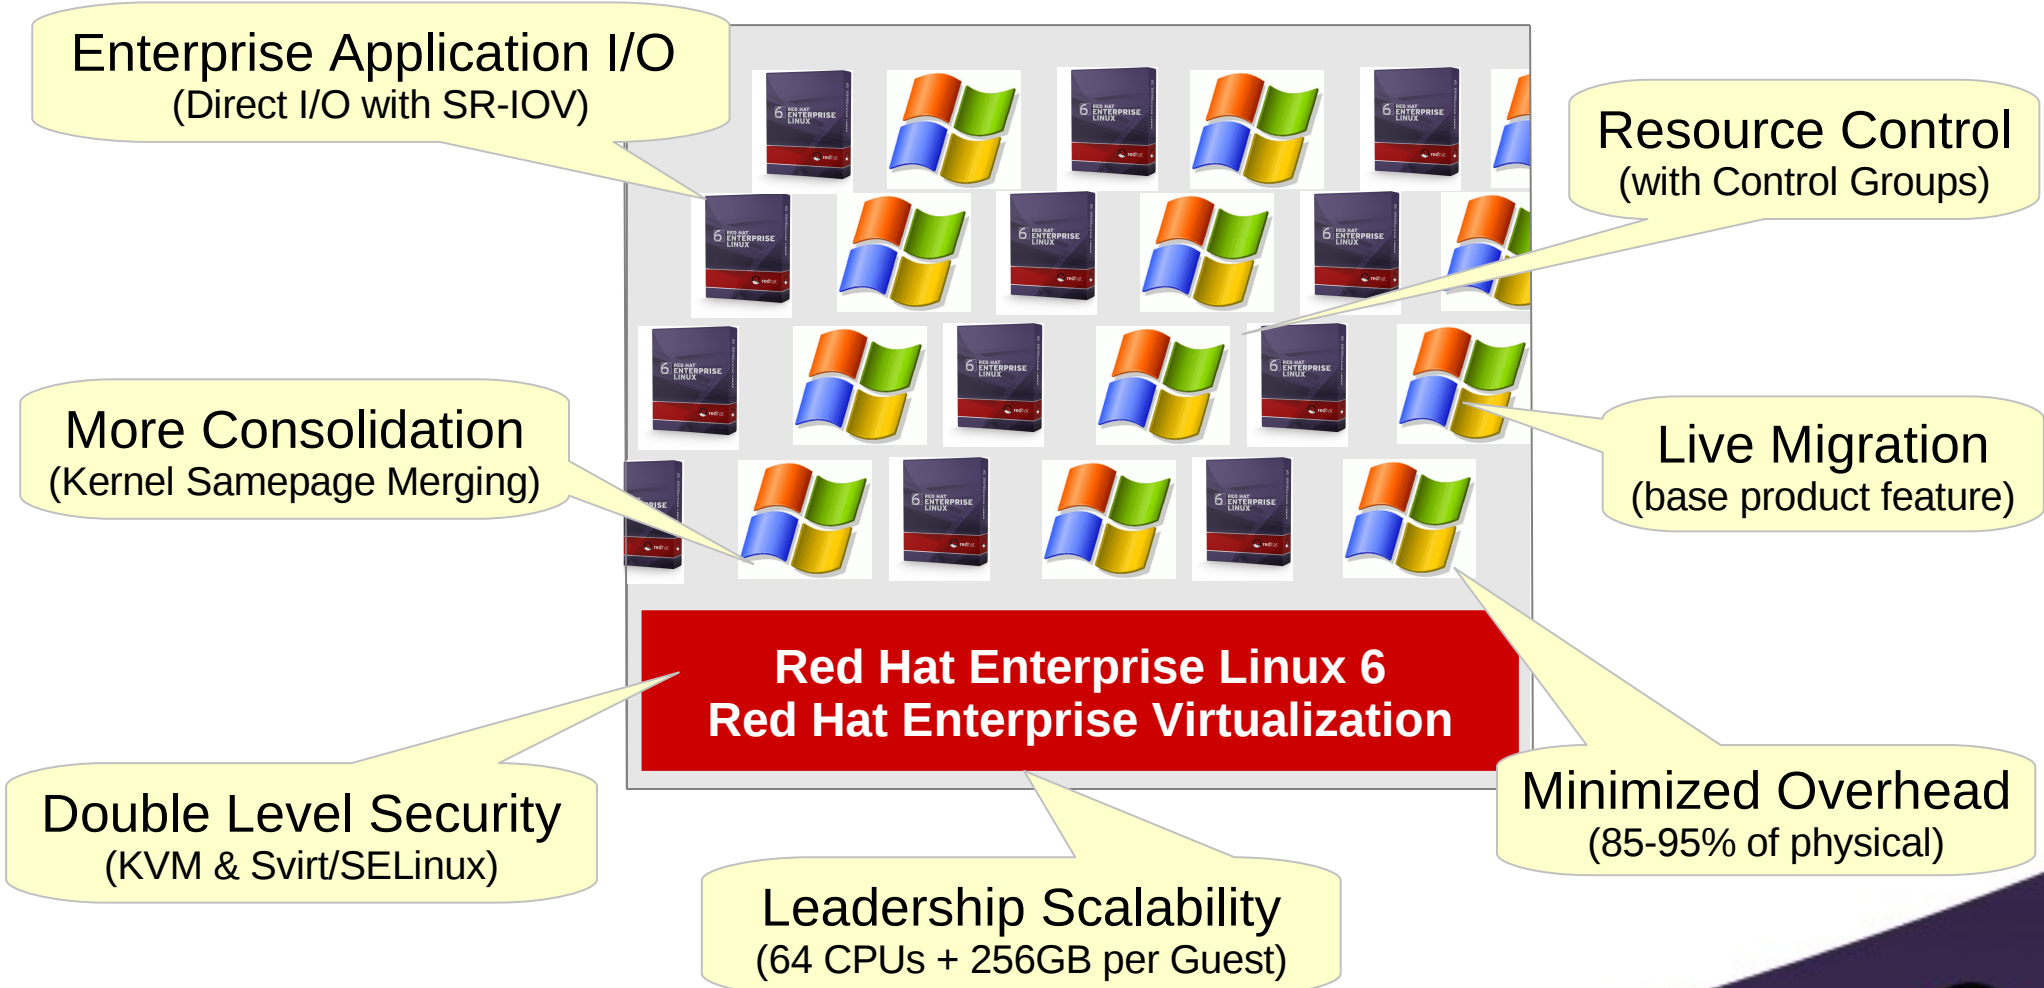
- *Enable developers and system administrators to minimize power consumption*
 - Powertop
 - Identifies power hungry applications and system services
 - Tuned - adaptive tuning daemon
 - Power down idle peripherals
 - Latency policy scripts
 - Provides a variety of power tuning profiles
 - Documentation to application developers and system administrators on tips and tricks




Virtualization




- *Enhancements to make Virtualization ubiquitous*




Virtualization



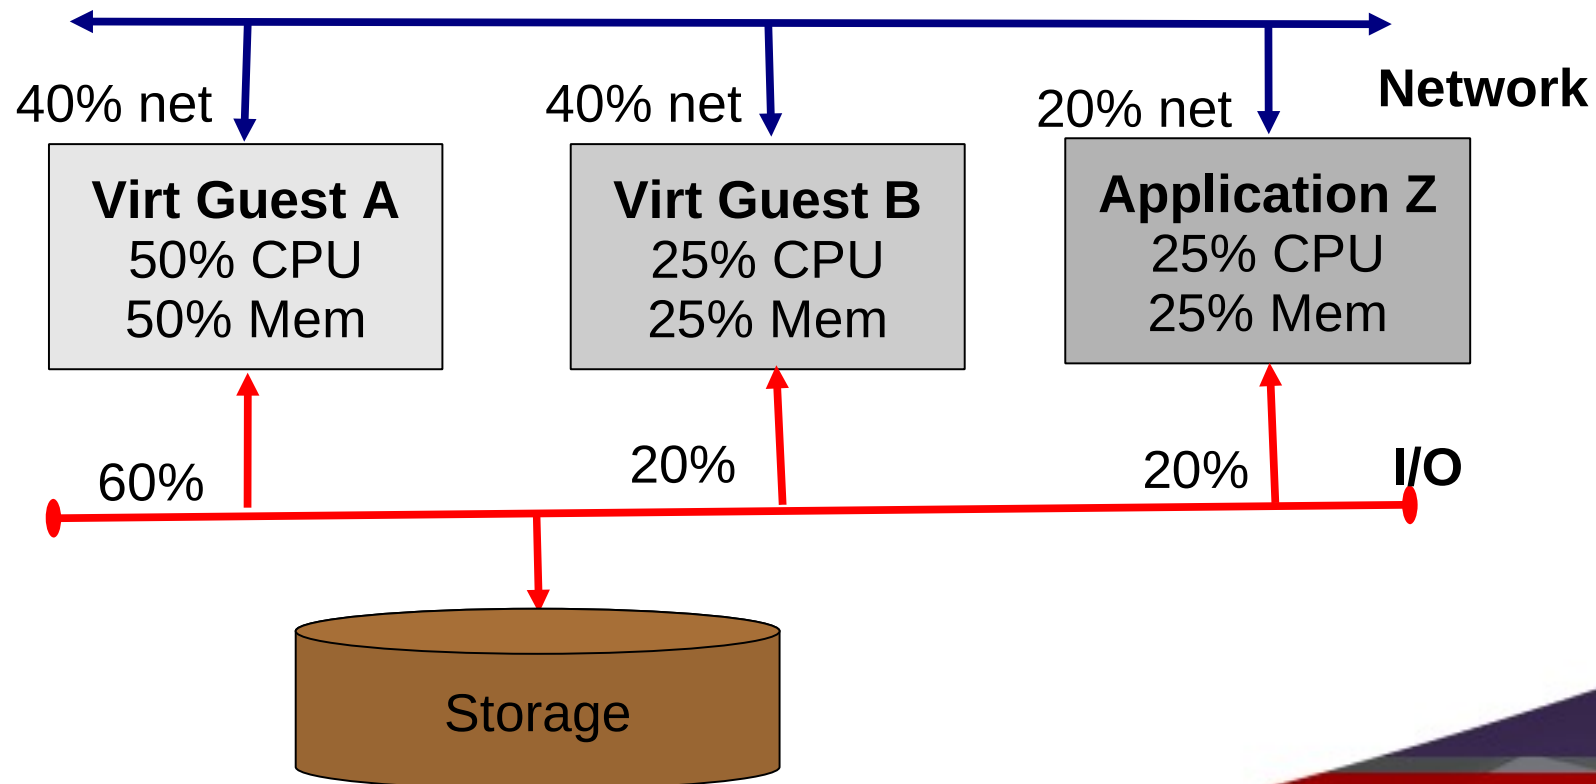
- *Deliver enhancements to make Virtualization ubiquitous*
- Performance:
 - Commonly 85%-95% of bare metal, including I/O bound workloads
- Scalability:
 - Host: 96 cores; 2 TB RAM
 - Guest: 64 CPUs; 256 GB RAM
- Advanced capabilities:
 - Live Migration; CPU/Mem resource control
 - Memory page sharing (KSM); SR-IOV; VT-D; SE-Linux security
- Hypervisor integrated into the Linux kernel
 - All features accrue to Red Hat Enterprise Virtualization, which also provides sophisticated management capabilities




Virtualization: Resource Management




- *Ability to manage large system resources effectively*
 - Control Group (Cgroups) for CPU/Memory/Network/Disk
 - Benefit: guarantee Quality of Service
 - Ideal for: Virtualization/Cloud deployments



Virtualization: Svirt Security



- *Applying security labels to individual guest virtual machines and their resources*
 - Guest Isolation achieved with SELinux Mandatory Access Controls (MAC)
 - Protect against untrusted Guest VM
 - Protect against Host misconfiguration
 - Prevents unauthorized access of Guests/Host
 - Builds on existing, proven security mechanisms & controls



Reliability, Availability, Serviceability (RAS)



- *Enable replacement of high-end RISC/Unix*
 - Based on new hardware/software capabilities
- 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 (up to 10x faster than RHEL 5)
 - E.G. Fsync for 1TB filesystem (45 million files): Ext3 = 1 hour, Ext4 = 6 minutes.
- DIF/DIX: End-to-end data checksumming
- ABRT: Automated bug reporting with “phone home”



Development



- *Complete, modern, development environment*
- Red Hat is a technical leader on all these initiatives
- Extensive compiler/toolchain:
 - OpenJDK: runtime and open Browser/Web Start plugins
 - Gcc 4.4: OpenMP3 conformance & C++ enhancements
 - Glibc: NUMA/malloc speed enhancements
 - GDB debugger: multithread/process debugging
 - Systemtap: Linux Dtrace for kernel/user/Java
 - Eclipse: CDT for Java, C, C++



OpenJDK




Security



- *Capabilities for new deployment models, virtual servers & clouds*
 - Sandbox: control of applications within an SELinux domain
 - Kiosk: run secure kiosk Live OS with SELinux enforcement
 - XACE: secure independent MLS windows environment
 - SSSD: robust client for LDAP & Kerberos authentication
 - Key Escrow: secure storage & recovery of encryption keys
 - Minimal Install: lowest possible attack surface for core OS
 - Security Standards
 - NSS: FIPS 140-2 certified cryptographic library
 - SHA256: release management tools meeting FISMA requirements
 - OpenSwan: RFC 5114 and Cisco VPN compliant



Networking



- *Enhancements cover multiple interconnects, transmission algorithms, and broad RFC support*
 - Performance optimizations
 - 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
 - 10GbE Driver support – on card switch and 8-16 pci devices
 - Data Center Bridging (DCB) support
 - 802.1p VLAN tags; 802.1Qaz (grouping); 802.1Qbb (flow control)
 - FCoE (Fibre Channel over Ethernet)
 - Performance improvements throughout the storage stack
 - RDMA support – over 10GbE (RoCE) & Infiniband
 - Ideal for low-latency messaging & large packets




Filesystems & Storage



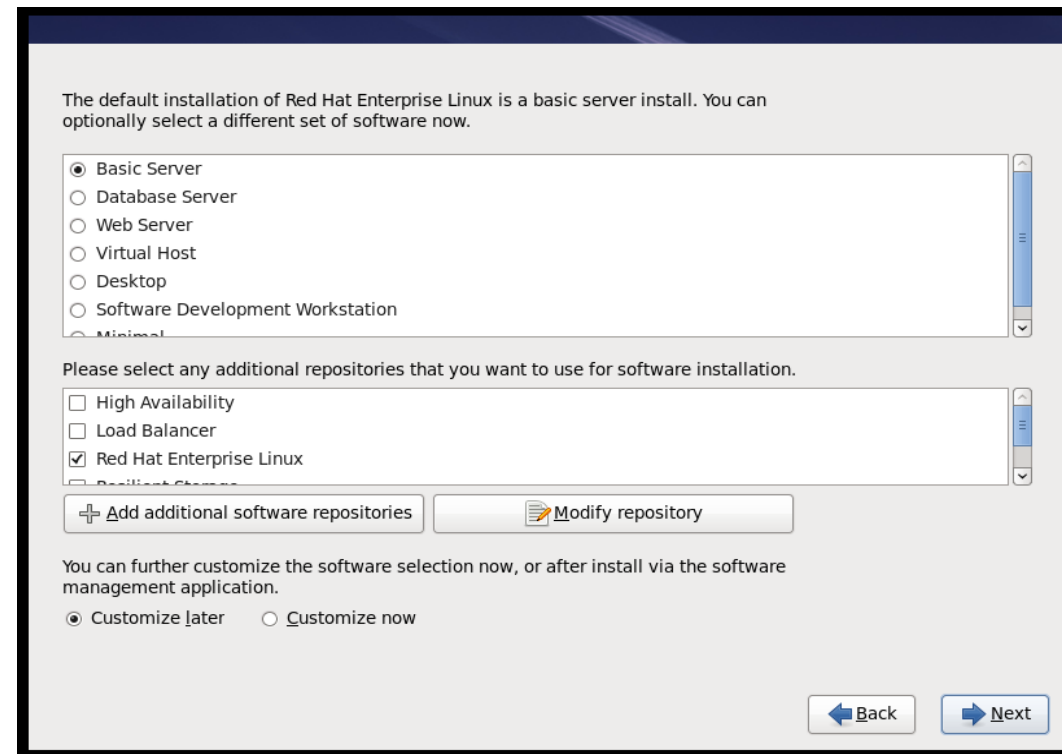
- *Multiple file systems for different capabilities*
 - Ext4: new system default file system - faster, more scalable
 - XFS and GFS: options for very large storage, and clustered environments
 - NFS 4.0/1: enhanced encryption and IPv6
- *Storage enhancements for data integrity, SSDs, Virtualization*
 - DIF/DIX: data checksumming from HBA <-> disk
 - I/O alignment & size awareness: allows storage access optimizations
 - Virtual server storage optimizations: SR-IOV, NPIV, VSAN
 - Thin provisioning & Block discard: assist SSD wear leveling
 - Extensive LVM enhancements: online resize; multipathing; mirroring; snapshots; encryption



Interoperability and Deployment



- Installation using Workload Profiles/Personalities
 - Minimal install option provides minimal security attack surface
- SHA256 support for strengthened Red Hat package integrity
- Microsoft Interoperability
 - Client support for Windows 2008 R2 active directory
 - File/Print (Samba) file sharing
 - Ipv6 & Windows 7 domain support
 - Encryption between client & server
- Smart Management Add-on provides manageability, monitoring and provisioning



Compatibility



- Most applications will transfer directly from Red Hat Enterprise Linux 5 to Red Hat Enterprise Linux 6
- Compatibility libraries for older Red Hat Enterprise Linux versions are provided for applications that need them
- Documentation & white-papers provided to simplify migration:
 - Planning guide
 - Configuration tips for system administrators
 - Documentation on configuration file changes
 - Software packaging mapping across releases
 - Consulting services for planning & executing migration



Desktop




- *Complete update of desktop application set*
- Significant advances:
 - Laptop/mobile/wireless
 - Graphics (KMS, DR12, OpenGL 3D)
 - Hardware – multi-monitor, dynamic input devices, bluetooth + 3G wireless, smartcard
 - PackageKit – lockdown support, application security levels
 - MS Exchange 2007/2010 interop (OpenChange & Evolution)
- GNOME 2.28 and KDE 4.3.0
 - Red Hat is the largest contributor to GNOME development

GNOME Contributions			
	Company	Commits	%
1	None	101823	23.45
2	Unknown	73558	16.94
3	Red Hat	70790	16.30
4	Novell	45349	10.44
5	Collabora	21684	4.99
6	Intel	11160	2.57
7	Fluendo	10218	2.35
8	Lanedo	10090	2.32
9	Independent	8922	2.05
10	Sun	8862	2.04
...			
16	Canonical	4487	1.03

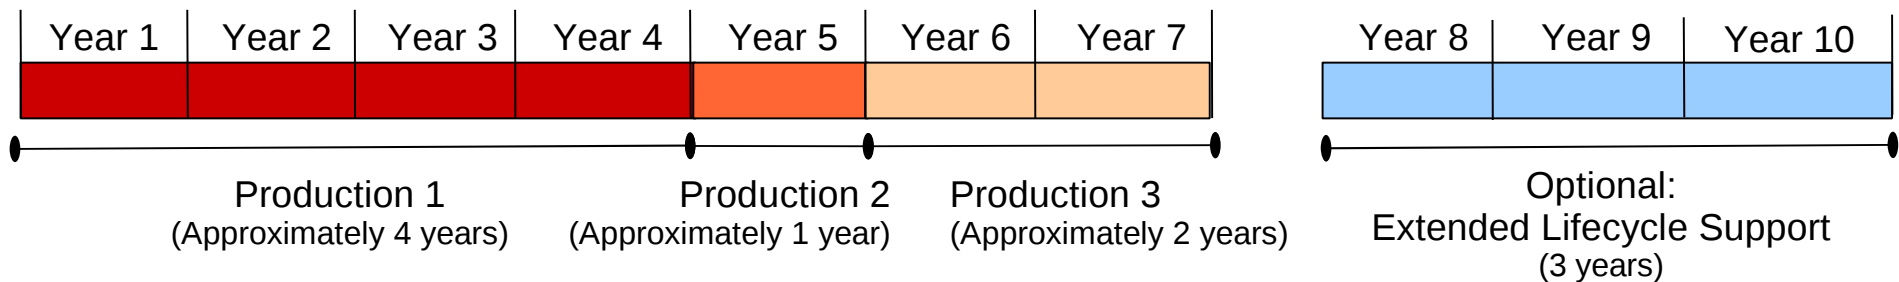
Source: Nealy Consulting
GNOME Census, August 2010



Product/Support Lifecycle



- Designed to allow:
 - Long-term stable deployment
 - No hardware/application re-certification after updates
 - ISV/application ecosystem: stable User/Kernel API/ABIs maintained for product life
 - Allows ongoing feature updates with continuous stability
- 7 year core lifecycle; additional 3 years optionally available



Summary



Red Hat Enterprise Linux 6 provides:

Extensive enhancements covering all dimensions of an enterprise-class operating system

Benefits for end-users, private & public IT organizations, developers & hardware suppliers

Capabilities that scale from the laptop to SMB, mid-range, enterprise & mainframe deployments

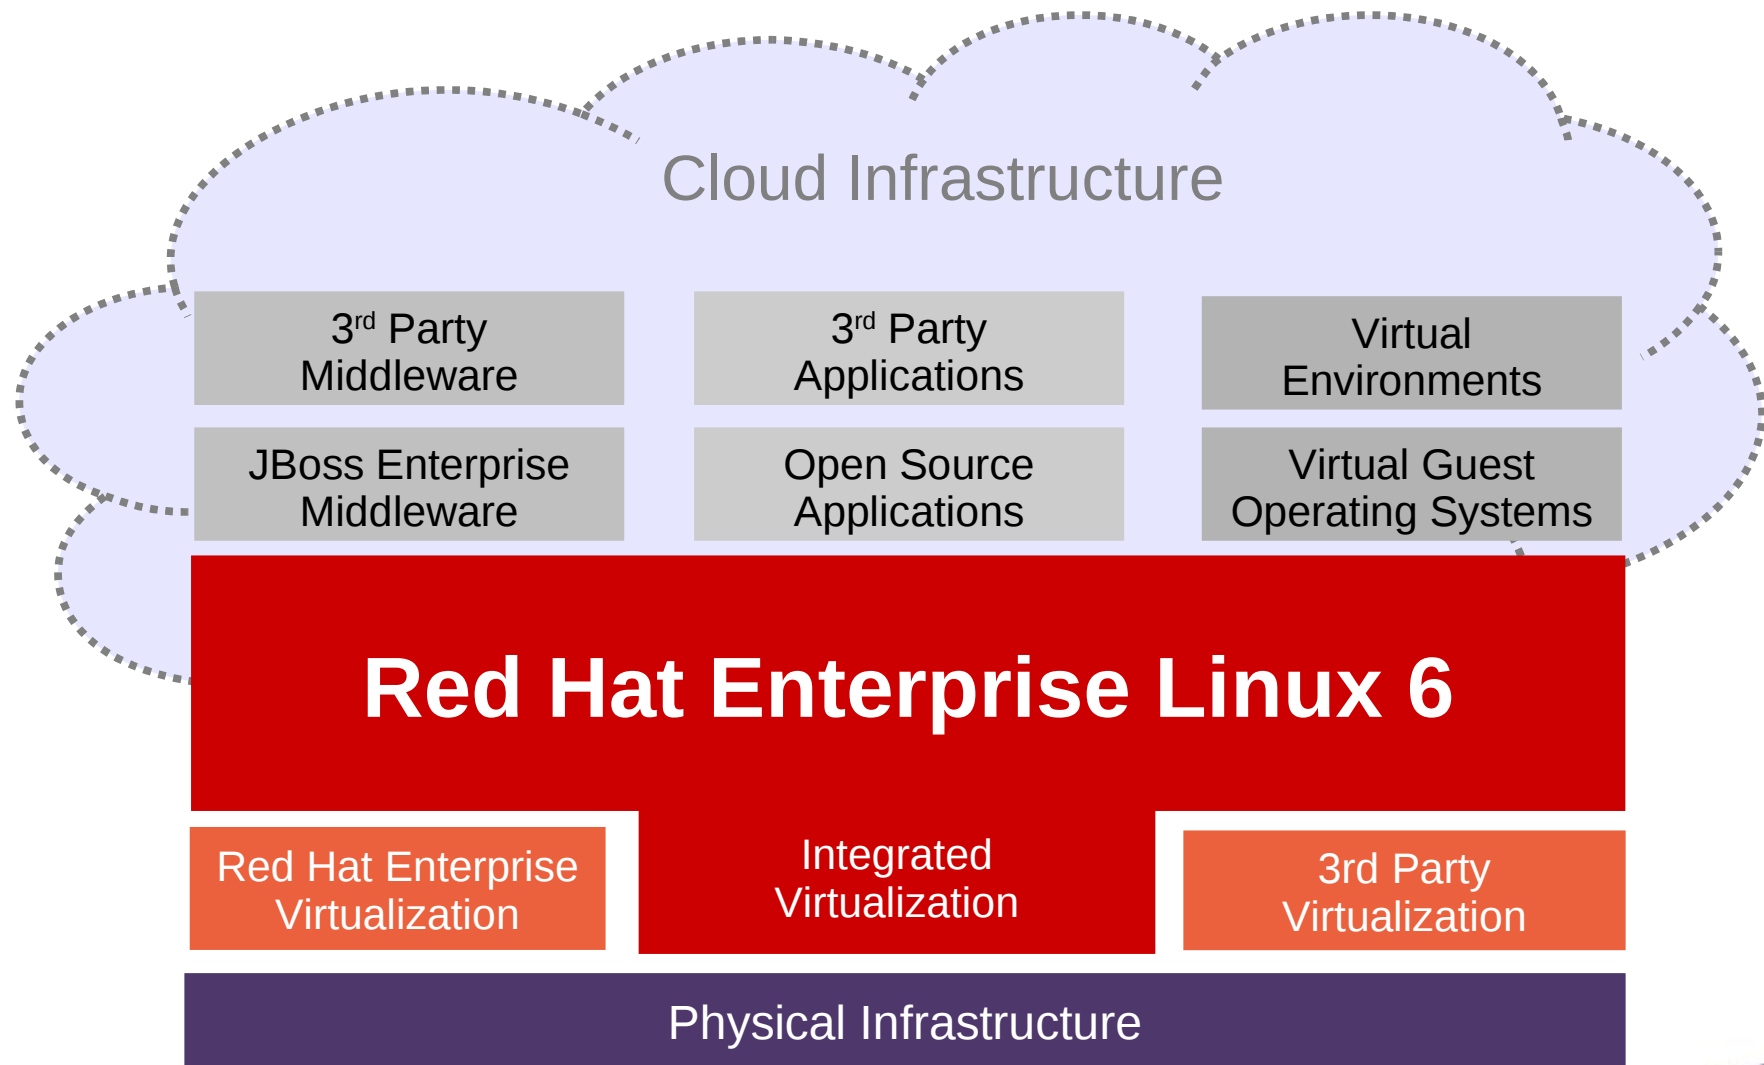
A long-term, stable application environment for physical, virtual and cloud deployments

An easy-to-configure product packaging model, with high quality services

Delivered by the world's leading developer & supplier of open source solutions



At the Heart of IT





THANK YOU!

Questions