



# Advancements with Open Virtualization & KVM

**Michael Day**

**Distinguished Engineer & Chief Virtualization Architect  
Open Systems Development, IBM**



# Discussion

- KVM's unique role as a Linux, Bare-metal Hypervisor
  - Pressuring Proprietary Hypervisors, Offering Choice, Gaining Unique Workloads
- IBM's Active KVM Development Program - 2005 and continuing
- Emerging Uses of KVM
  - Areas where KVM has a clear advantage
- Challenges ahead

# KVM Leverages Linux

As a Linux developer, it's hard for me to be that interested in Xen... When you think about it, it is really quite silly. We advocate Linux for everything from embedded systems to systems requiring real-time performance, to high-end mainframes. I trust Linux to run on my dvd player, my laptop, and to run on the servers that manage my 401k. Is virtualization so much harder than every other problem in the industry that Linux is somehow incompatible of doing it well on its own? Of course not.

-- Anthony Liguori, Qemu maintainer


# Scalability

- “With this patchset, -smp 64 flies... Amazing work!”  
-- Avi Kivity, commenting on 64-way SMP guest performance, December 27, 2009

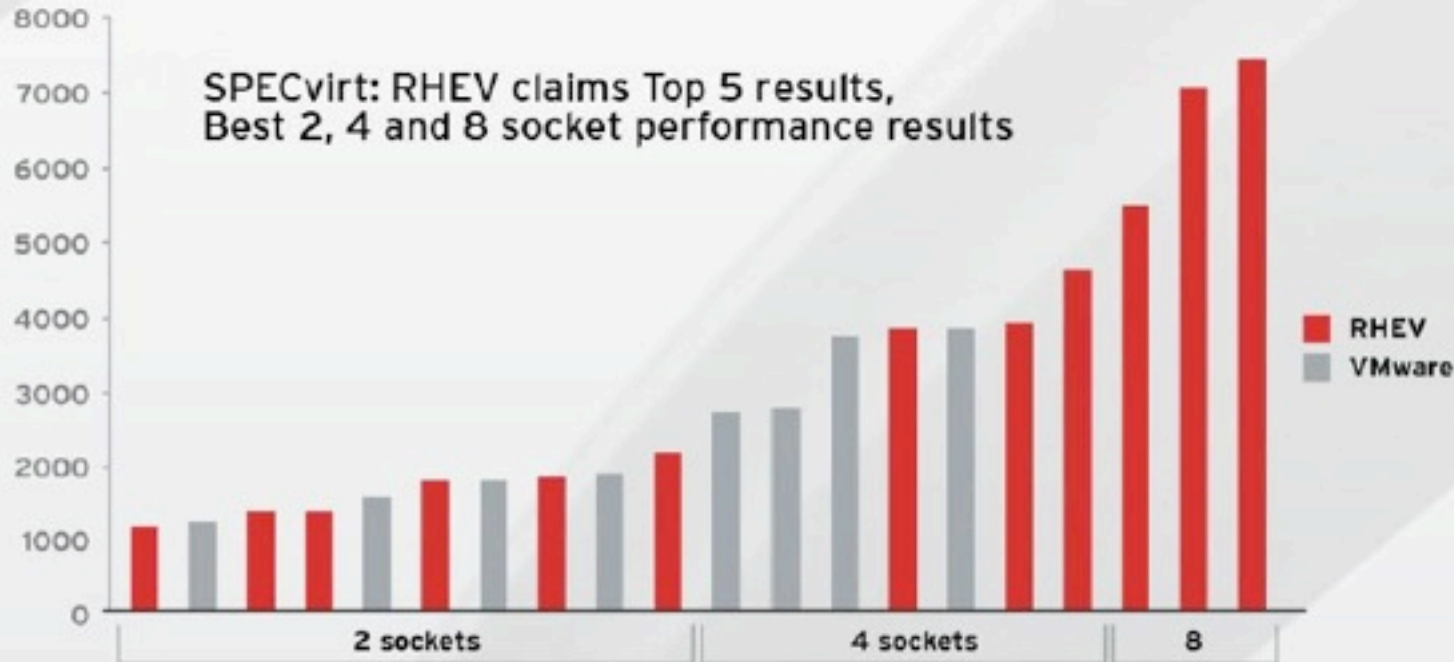
## Supported:

- 160 Cores, 4 TB RAM Per Host
- 64 Guest vCPUs, 512 GB Guest RAM
- 200 Hosts per Cluster

# VMExit Latency and Performance




## INDUSTRY LEADERSHIP: PERFORMANCE AND SCALABILITY



**INHERENT ARCHITECTURAL ADVANTAGE OF KVM,  
BASED ON THE LINUX KERNEL**

Source: [http://www.spec.org/virt\\_sc2010/12/20/2011](http://www.spec.org/virt_sc2010/12/20/2011)



**Note: IBM and HP have published KVM results**

# IBM Corporation

## IBM Corporation x3850 X5

Processor: Intel Xeon E7-4870 (80 cores, 8 chips, 10 cores/chip, 2 threads/core)

Memory: 2 TB (128 x 16 GB, Quad Rank x4 PC3-8500 CL7 ECC DDR3 1066MHz LP RDIMM)

Red Hat Enterprise Linux 6.0 (KVM)

Full result disclosure: [HTML](#)

Unprocessed data: [RAW](#)

Supporting documentation: [TGZ](#)

7067@432

Application Server

Database Server

Mail Server

Web Server

Infrastructure Server

Idle Server

IBM Corporation	-IBM Corporation x3850 X5 Processor: Intel Xeon E7-4870 (80 cores, 8 chips, 10 cores/chip, 2 threads/core) Memory: 2 TB (128 x 16 GB, Quad Rank x4 PC3-8500 CL7 ECC DDR3 1066MHz LP RDIMM) Red Hat Enterprise Linux 6.0 (KVM)  Full result disclosure: <a href="#">HTML</a> Unprocessed data: <a href="#">RAW</a> Supporting documentation: <a href="#">TGZ</a>	7067@432	+Application Server    +Web Server +Database Server        +Infrastructure Server +Mail Server              +Idle Server
-----------------	--	----------	--

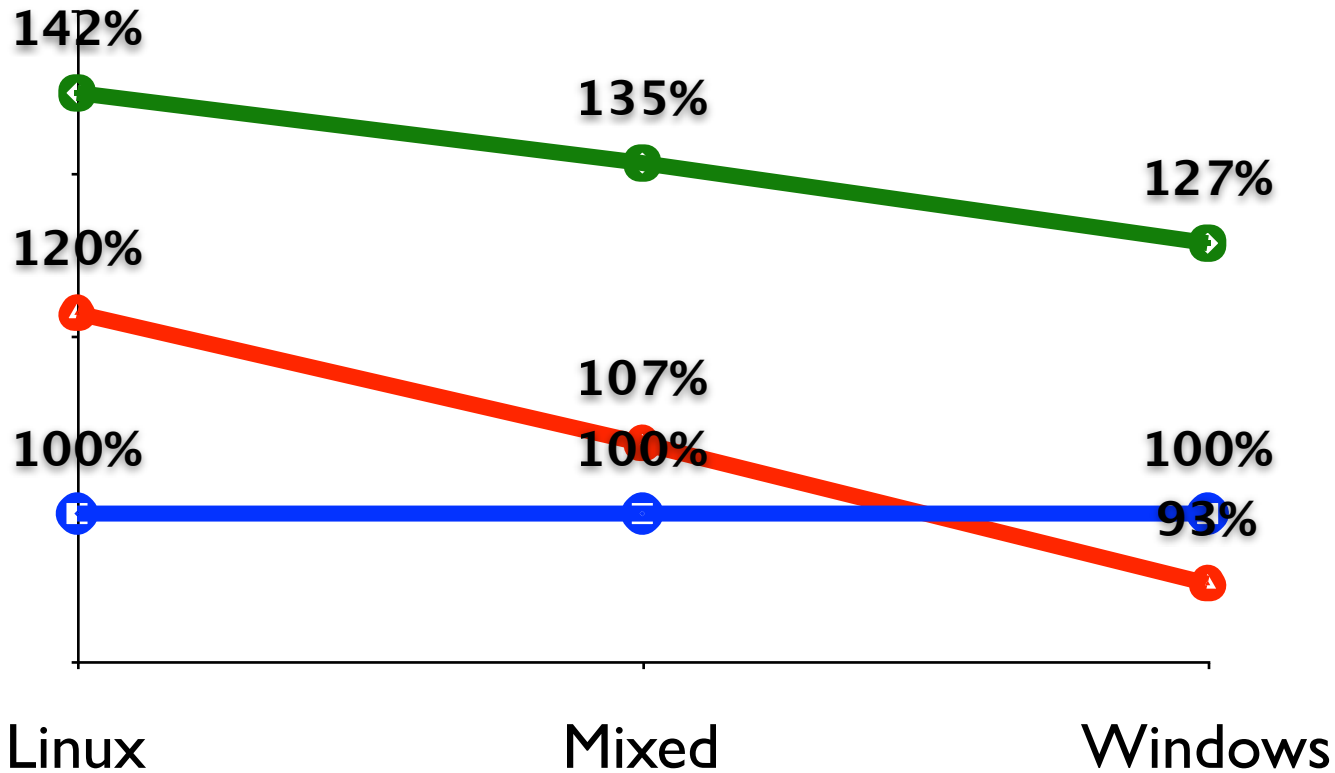
[http://www.spec.org/virt\\_sc2010/results/specvirt\\_sc2010\\_perf.html](http://www.spec.org/virt_sc2010/results/specvirt_sc2010_perf.html)

# Unprecedented I/O Performance

- September 2012:
  - 800k IOPS in a single KVM guest
  - 1.4M IOPs in 4 guests KVM
  - ([ftp://public.dhe.ibm.com/linux/pdfs/2012RHEL\\_KVM\\_Hypervisor\\_Performance\\_Brief\\_19\\_v4.pdf](ftp://public.dhe.ibm.com/linux/pdfs/2012RHEL_KVM_Hypervisor_Performance_Brief_19_v4.pdf))
- February 2013:
  - **1.2M IOPS** in a single KVM guest
  - IBM System X 3850 X5, QLogic QLE 256x HBAs, RHEL 6.4 (host and guest), virtio-blk-data-plane.



# KVM Cost Analysis by Guest Workload



[https://www.ibm.com/developerworks/mydeveloperworks/blogs/ibmvirtualization/entry/taking\\_the\\_sting\\_out\\_of\\_virtualization\\_licensing\\_costs\\_with\\_kvm](https://www.ibm.com/developerworks/mydeveloperworks/blogs/ibmvirtualization/entry/taking_the_sting_out_of_virtualization_licensing_costs_with_kvm)


(David Hsu)

# VMware kills vRAM licensing, pledges more support for OpenStack, multi-cloud world

**Summary:** *It's safe to say VMware is feeling the heat from its open source rivals.*



By [Paula Rooney](#) for [Virtually Speaking](#) | August 27, 2012 -- 17:22 GMT (10:22 PDT)

 [Follow @PaulaRooney1](#)

VMware's decision to kill its vRAM licensing model and expand its OpenStack support illustrates the significant power the open source model has amassed in the virtualization industry.

In his first official keynote as VMware's next CEO, effective Sept 1, Pat Gelsinger announced the end of VMware's confusing and pricing vRAM licensing model and said the company will expand its support for OpenStack and the multi-cloud world.

Gelsinger said a survey of 13,000 customers revealed that pricing was the most unpopular aspect of VMware's business and that he was happy to announce the end of vRAM. There will be "no vRAM or Core entitlements," said Gelsinger, who is currently president and chief operating officer of EMC.

# We are seeing marked increase in production usage of KVM among our customers

- Drivers
  - Technical progress of KVM Solutions
    - Improvements in RHEL 6.x over RHEL 5.x
    - Improvements in RHEV-M 3.0 over v2.2
  - Derivation of trust and maturity as characteristics of KVM due to time in the marketplace
  - Emerging View that KVM is a fundamental feature of the Linux Operating System
  - Economic Factors
  - Cloud computing

# Enterprise Linux Users

- KVM is an integrated feature of Enterprise Linux
- Existing Licensees already have it and already know how to use it
- KVM's administration tool-chain is integrated with the Linux tool-chain
- We see large enterprise Linux customers deploying virtual machines onto existing Linux servers rather than purchasing new VMware licenses
  - For both Linux and Windows workloads

# Scavenging/Grid


- This use case monitors Linux servers for under-utilized hosts and deploys VM images to hosts selected according to some policy
  - Energy usage
  - Time of day, workload completion
  - Dev and test
  - Distributed batch processing
- Creates an ad-hoc grid of KVM hypervisors out of dynamically selected Linux hosts
- Multiple Financial Customers

# Cloud Computing

- KVM is a natural infrastructure for Cloud Data Centers
  - KVM is efficient
    - Drives higher densities
  - KVM is secure, and has specific features that simplify multi-tenancy deployment
    - sVirt - seLinux based mandatory access control
    - Linux process control
    - cgroups hard and soft resource limits
  - KVM out-scales the competition
  - KVM is economical

# KVM and Public Clouds

- IBM's Smart Cloud Enterprise
  - Built using KVM
  - Tivoli Services Automation Manager
  - Custom integration toolkit
- Google Compute Engine
  - Built using KVM
  - Google's infrastructure
  - Custom Integration toolkit
- HP's Public Cloud
  - Build using KVM
  - OpenStack




## KVM and Private Clouds

- Large Financial Customers are building private clouds using Moab with xCAT and KVM
- Moab and xCAT have become fully-featured cloud infrastructures, while maintaining their scalability and performance as HPC cluster managers
- These tools support VMware and KVM equally well and make them plug-compatible
  - Customers are using this to phase-in more KVM usage over time
  - Postponing or canceling new VMware license purchases



## KVM and Large Enterprises

- The largest enterprises can save tens-of-millions of dollars per year just by reducing growth in new VMware licenses, or by slowing the renewal of existing VMware licenses
- They are finding it easy to identify Virtual Machine use cases on the margin to deploy to KVM instead of VMware
- This allows them to phase in more KVM over time
- Over three years the savings can be very significant

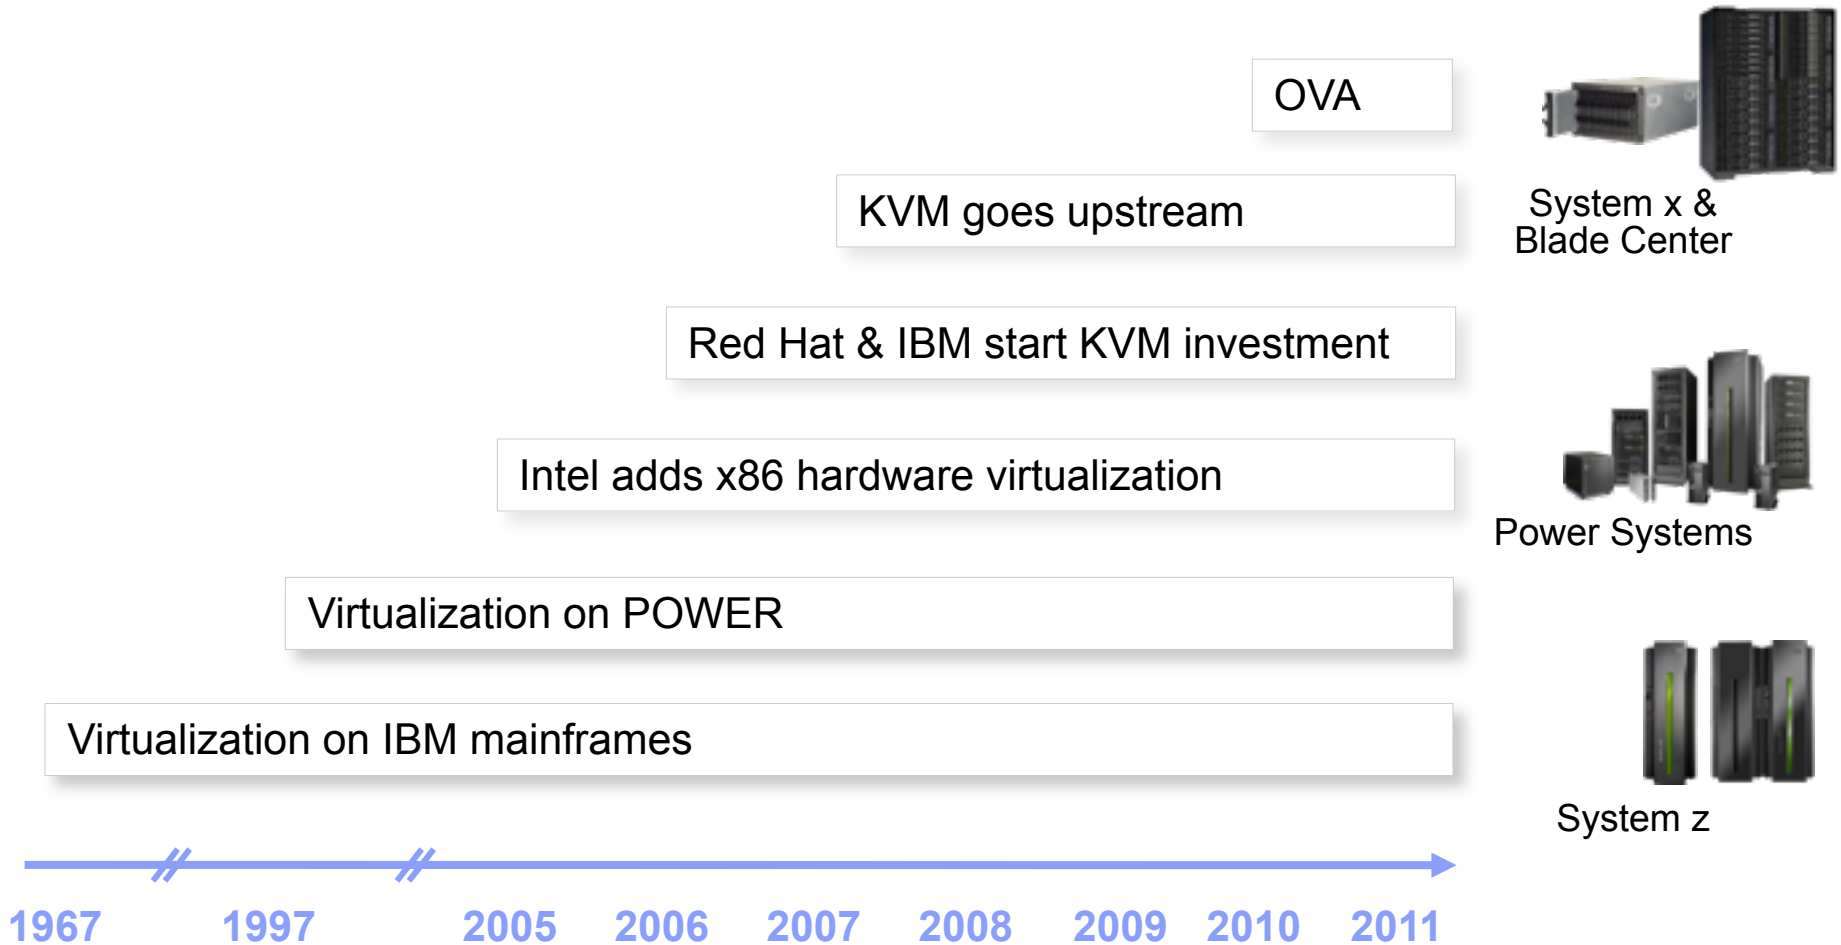


## Gartner - Reconsidering Heterogeneous x86 Server Virtualization, Sept. 2012

In fact, KVM is the only hypervisor in our last two surveys to notch gains in both the number of overall users **and** in the number of users adopting it as their standard go-to hypervisor. - Gabriel Consulting Group - "Hyperversity Rages On"

# IBM and Virtualization

*A brief history of virtualization*



## IBM's Market Goals for KVM

- IBM would like for KVM to be a significant presence in the marketplace
- To provide a technically excellent hypervisor and overall competitive open-source solution stack to the market
- To influence the direction of hypervisor development in a positive way
- To leverage our hardware platforms in a timely manner

### IBM Contributes to the Community

- IBM developers contributing to 100+ Linux and Open Source projects
- Develop closely with Red Hat and Novell
- Developers sharing technical knowledge on <http://planet-ltc.org>

### IBM Supports Linux as a Tier 1 OS

- All IBM Systems, SW, and Middleware run on and are certified for Linux
- Driving performance parity with IBM's own operating systems
- Making contributions in security, RAS, scalability, performance, management

Make Linux Better

Enable IBM Products

Customer Collaboration

Expand Linux Reach

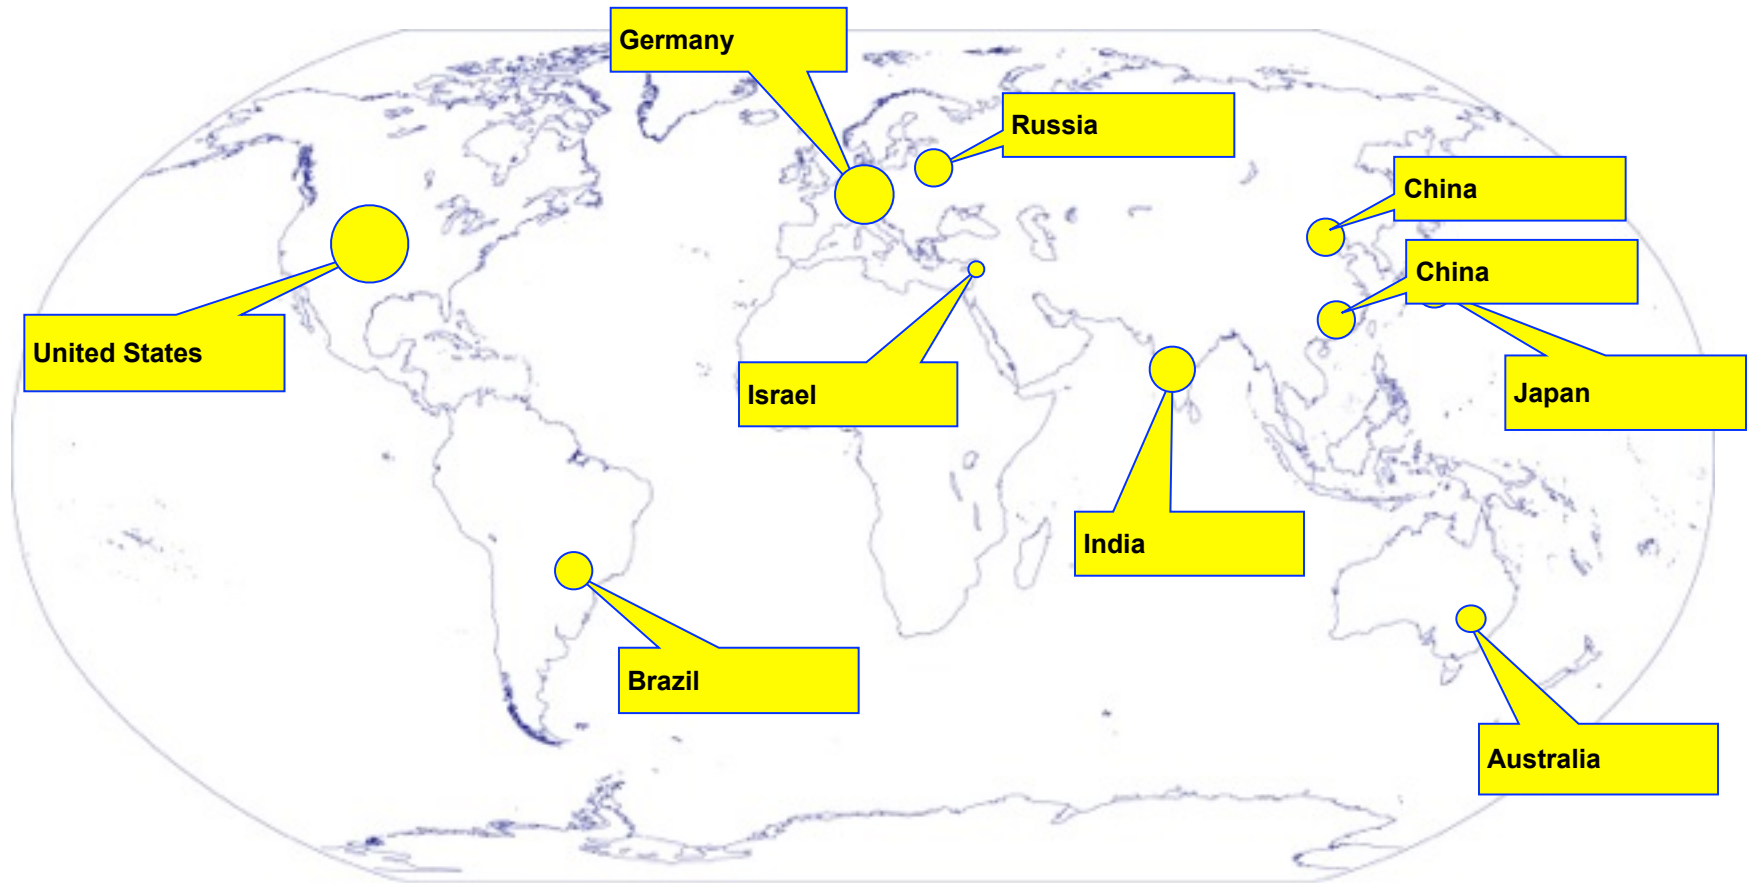


### IBM Collaborates with Customers

- Specialized and very detailed knowledge of IBM Systems and Software
- The LTC works with customers on unique proof of concept projects
- Scale Out File Services (SoFS)
- Real Time Linux and Java for Raytheon

### IBM Enables Linux for New Markets

- Working with groups such as the Linux Foundation to address new workloads
- Expanding and providing leadership in:
  - Blue Cloud Computing
  - SOA / Web 2.0 / SaaS
  - Distributed computing and HPC

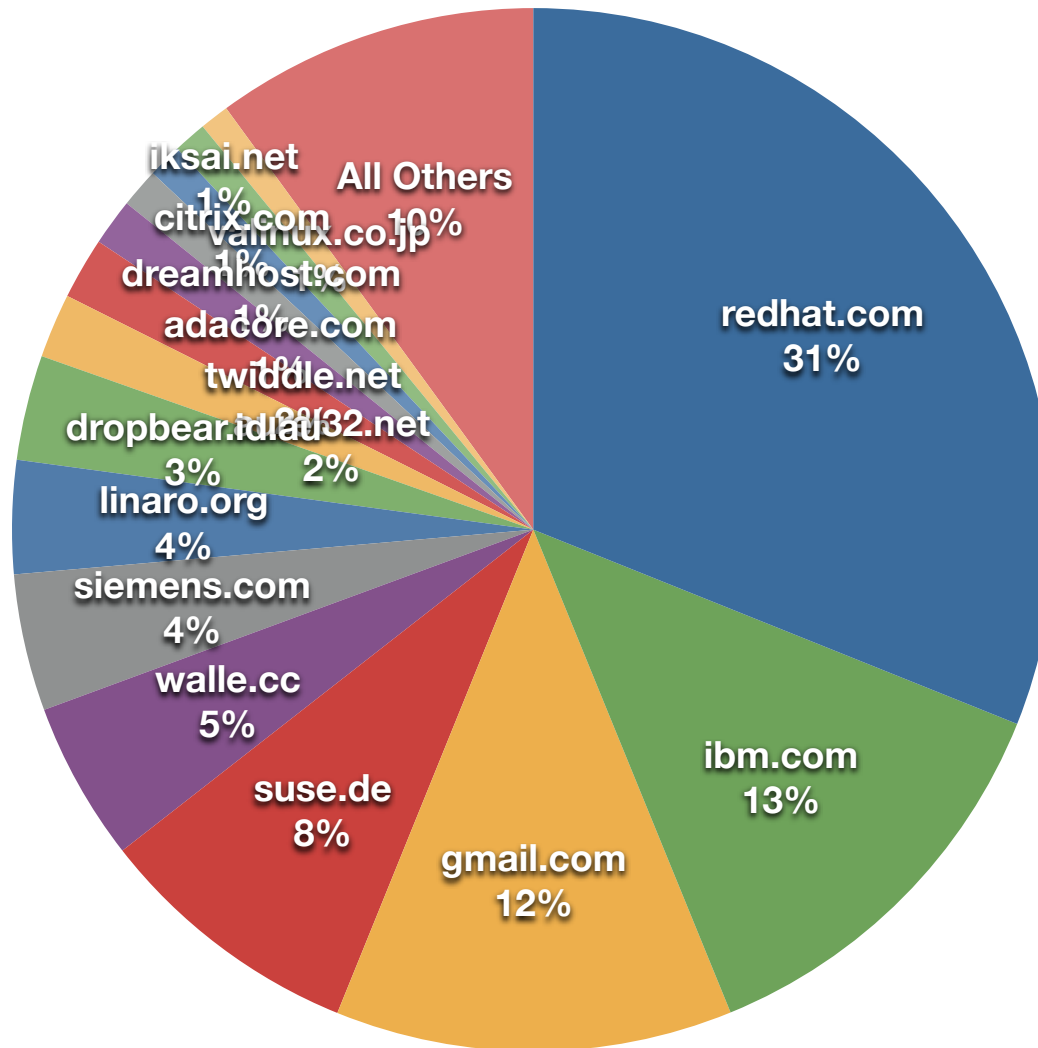


# IBM's KVM Development Focus

- Support for IBM's Public Cloud
  - oVirt
  - OpenStack
  - Qemu
  - I/O Performance
- Platform Support
- Security
- Performance

**IBM has more than  
65 developers  
dedicated to KVM**

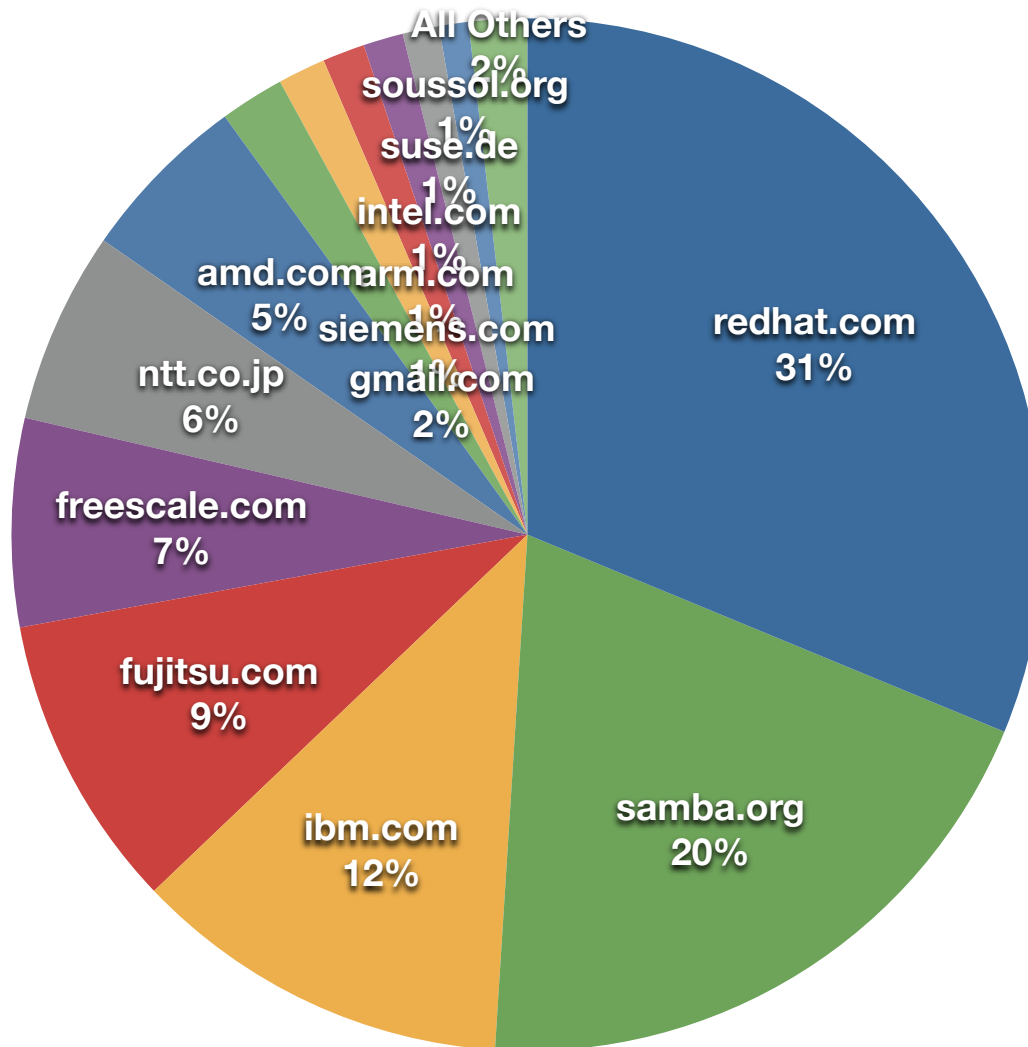
## 2011 QEMU Source Changes by Domain



<http://code.ncultra.org/2012/10/148/>




## 2011 KVM Source Code Changes by Domain




<http://code.ncultra.org/2012/10/148/>

## 2008-11 Source Code Changed Lines



<http://code.ncultra.org/2012/10/148/>

## 2008-11 Development Email Traffic



<http://code.ncultra.org/2012/10/148/>

# KVM for Power and s390 Platforms

- Increased upstream development activity from IBM
- Reflects lessons learned from Linux and KVM

```
[mdday@shoshone arch]\ $ pwd
/Users/mdday/src/kvm/arch
[mdday@shoshone arch]\ $ ls
Kconfig          blackfin          h8300             m68k              openrisc          score            um
alpha            c6x              hexagon          microblaze        parisc           sh              unicore32
arm              cris             ia64             mips              powerpc          sparc           x86
avr32            frv             m32r            mn10300          s390            tile            xtensa
[mdday@shoshone arch]\ $
```





**Challenges - we “pretty much” know where they are**

# Thank You

## Mike Day

[mdday@us.ibm.com](mailto:mdday@us.ibm.com)

**Virtualization @ IBM Blog**

<https://www.ibm.com/developerworks/mydeveloperworks/blogs/ibmvirtualization/>

**Life in Code**

<http://code.ncultra.org>

# Legal

## Trademarks and Disclaimers

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml):

IBM, the IBM logo, BladeCenter, Calibrated Vectedored Cooling, ClusterProven, Cool Blue, POWER, PowerExecutive, Predictive Failure Analysis, ServerProven, System p, System Storage, System x, System z, WebSphere, DB2 and Tivoli are trademarks of IBM Corporation in the United States and/or other countries. For a list of additional IBM trademarks, please see <http://ibm.com/legal/copytrade.shtml>.

The following are trademarks or registered trademarks of other companies:

Java and all Java based trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries or both Microsoft, Windows, Windows NT and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both. Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. UNIX is a registered trademark of The Open Group in the United States and other countries or both. Linux is a trademark of Linus Torvalds in the United States, other countries, or both. Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc. InfiniBand is a trademark of the InfiniBand Trade Association.

Other company, product, or service names may be trademarks or service marks of others.

## NOTES:

Linux penguin image courtesy of Larry Ewing ([lewing@isc.tamu.edu](mailto:lewing@isc.tamu.edu)) and [The GIMP](#)

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Information is provided "AS IS" without warranty of any kind.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices are suggested US list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Any proposed use of claims in this presentation outside of the United States must be reviewed by local IBM country counsel prior to such use.

The information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this