



# What's New with Red Hat on System z SPARTA, January 2011



Brad Hinson  
Worldwide Lead, Linux on System z  
[bhinson@redhat.com](mailto:bhinson@redhat.com)

# What is Linux?

- Operating system and tools
  - Linux distributions
- Started by Linus Torvalds
  - Hobby project
  - University of Helsinki Finland
- Free software (GPL)


# Linux Facts

- Last year, **75%** of the Linux code was developed by **programmers working for corporations.**
- **\$7.37 billion:** projected cost to produce the 283 million lines of code which are contained in Linux Distribution **in a commercial environment.**
- IDC forecasts show that **Linux server revenue will grow by 85.5%** between 2008 and 2012 **in the non-x86 server space** equalling a four year compound annual growth rate of 16.7%.
- **Linux is Linux**, but ...features, properties and quality differ dependent on your platform

# Linux on System z Development

## Community

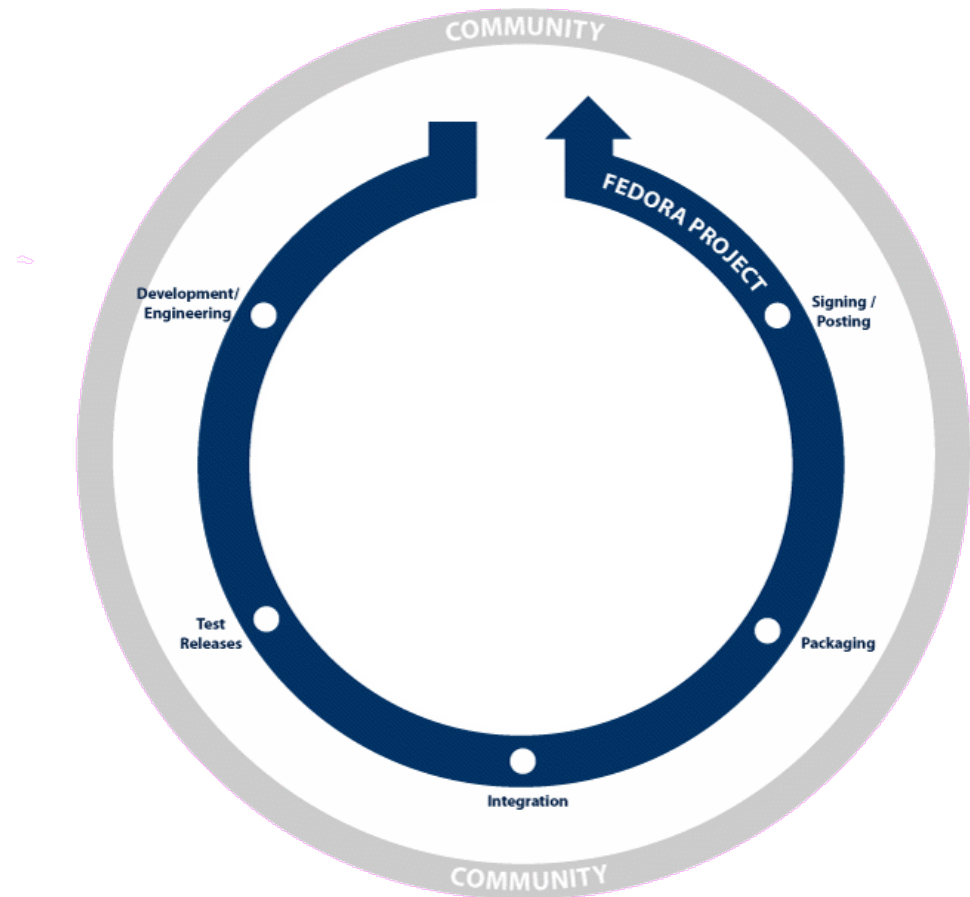
- Development with “upstream” communities
- Kernel, glibc, etc
- Collaboration with partners, IBM, open source contributors



# Linux on System z Development

## Fedora

- Bleeding edge, sets direction for RHEL technologies
- Community Supported
- Released ~6mo cycles
- [www.fedoraproject.org](http://www.fedoraproject.org)

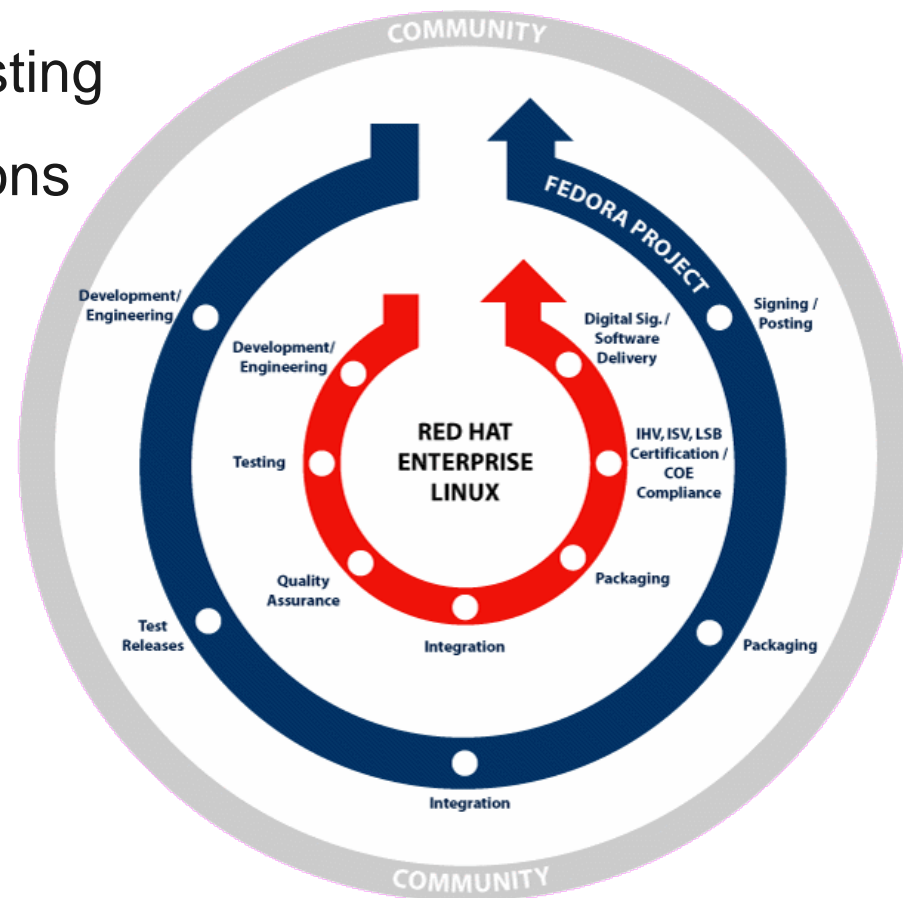




# Linux on System z Development

## Red Hat Enterprise Linux

- Stable, mature, commercial product
- Extensive Q&A, performance testing
- Hardware & Software Certifications
- 7yr maintenance
- Core ABI compatibility guarantee
- Major releases 2-3yr cycle





A photograph of a clear glass filled with water, sitting on a light-colored wooden surface. Two blue, oval-shaped buttons with the Fedora Linux logo (a stylized 'f') are placed at the base of the glass. The word 'fedora' is written in white lowercase letters to the right of the glass, with a small blue circular logo containing a white 'f' to its upper right.

fedora™

# What is fedora?



# ...A set of packages..



- 100% open source
- 100% free
- Includes the latest upstream developments
- Has over 10,000 packages



# Technical Excellence

- Work together with upstream
- A lot of our work goes to other projects



**FEATURES.**



**Who uses fedora?**



- Humanoid robots
- [Wikipedia.org](http://Wikipedia.org)
- Roadrunner, the number 2 Supercomputer in the world
- Evo Smart Console gaming console
- Over a hundred Derived distributions including Red Hat Enterprise Linux
- NASA Systems
- Many universities and web hosting companies
- Millions of other users



# Fedora

## Common Deployments

### File/Print Server

Includes technology (SAMBA) to act as a file server for Windows, Unix and Linux clients. With Add-Ons such as Resilient Storage can be used in mission critical application environments.

### Web Server

- Includes the entire LAMP stack with options for MySQL, PostgreSQL & SQLite

### Applications Server

- Includes a version of TomCat but is also the basis for any robust middleware such as JBOSS or WebSphere.

### Database Server

- Broadly deployed and can be made resilient with the High Availability Add-On

### Computational Clusters

- With special packages for HPC, and coupled with Red Hat's own MRG-Grid enables customers to build a complete HPC cluster.

### Utility Infrastructure

- With thousands of open source packages, Fedora can be used throughout the data center, including fire walls, mail handler or directory server

### Enterprise Application Platform

- For hosting most enterprise applications such as Oracle, SAP, Documentum, and SAS

### Virtualization and the “Cloud”

- With integrated virtualization and resource management, Fedora is an excellent virtualization host, and is designed to be a superior guest on any of the major hypervisors.





<http://fedoraproject.org/join-fedora>

- It is as easy as getting account in any social networking site.
- Talk to fedora contributors
- For specific questions: talk to mentors

# Red Hat Enterprise Linux on System z

The world's leading open source application platform

## On one certified platform

- Applications - Thousands of certified ISV applications
- Deployment - Including standalone or virtual servers, cloud computing, or software appliances
- Hardware - Wide range of platforms from the world's leading hardware vendors

## Comprehensive enterprise infrastructure

- Server Virtualization
- Storage Virtualization
- Systems management, provisioning, monitoring
- Push TCO by using a single Operating System from Laptop to Mainframe
- Highest security accreditations (via SELinux)
- Hardware exploitation through co-development with IBM





# Red Hat offers a comprehensive platform

Features	Red Hat	Costly vendor alternatives
Operating system	<b>Red Hat Enterprise Linux with integrated virtualization</b>	Windows, AIX, HP-UX, Solaris
HPC		Platform Computing
Multi-path I/O		EMC PowerPath
Systems provisioning		IBM Tivoli
Systems management		HP OpenView/Opware, IBM Tivoli
Directory services		eDirectory
Volume/storage management		Symantec Storage Foundation Suite
Server virtualization		VMware, Citrix

## Also in Red Hat Enterprise Linux


- Full exploitation of advanced z/VM features
- Multiple open source database options
- Production-ready LAMP stack

## Special Technical Advantages

- Proven track record of high-end scaling
- Customizable scheduler for different workloads
- Proven virtualization performance



# What are Linux users running on RHEL?



Surveys indicate customers use:

- Web Serving
- Web Application Serving
- Data Services
- Systems Development

“Best Fit” Workloads for Enterprise Linux Servers:

- **Web Application Servers:** WebSphere Application Server
- **Email and collaboration:** Domino™, Web 2.0
- **Data services:** Cognos®, DB2, Oracle, Informix®, Information Server, Information Builders WebFOCUS
- **Business critical ISV applications:** e.g. SAP
- **Development of WebSphere and Java™ applications**
- **Virtualization and security services**
- **Business connectors:** WebSphere® MQSeries®, DB2® Connect, CICS® Transaction Gateway, IMS™ Connect for Java®
- **Network Infrastructure:** FTP, NFS, DNS, etc. and Comm Server and Communications Controller for Linux, CommuniGate Pro (VoIP)
- **Applications requiring top end disaster recovery model**



# RHEL 6: Key Features

## ■ **Power Management**

- Improvements through the application stack to reduce wake ups
- Power consumption measurement by Powertop
- Power Management and adaptive system tuning by Tuned

## ■ **Next Generation Networking**

- Comprehensive IPv6 support (NFS 4, CIFS, ISATAP support)
- Redesigned QETH network driver with support for OSA data connection isolation

## ■ **Reliability, Availability, and Serviceability**

- System level enhancements from industry collaborations to make the most of hardware RAS capabilities.

# RHEL 6: Key Features

## ■ **Fine-grained Control and Management**

- Improved scheduler and better resource management in the kernel via Completely Fair Scheduler (CFS) and Control Groups (CG).

## ■ **Scalable File systems**

- New file systems btrfs and ext4 (now the default) offer robustness, scalability, and high-performance.

## ■ **Disk Storage Subsystem**

- FCP automated port discovery and Isluns utility to automatically activate all available target ports
- Support for High Performance FICON to reduce I/O overhead
- Dynamically adjustable queue depth
- I/O configuration support when running in LPAR

# RHEL 6: Key Features

## ■ **Virtualization**

- Tighter integration with z/VM for extended functionality like dynamic memory resizing, better CPU utilization, HyperPAV, and suspend/resume support.

## ■ **Enterprise Security Enhancement**

- SELinux includes improved ease of use, application sandboxing, and significantly increased coverage of system services
- SSSD provides unified access to identity and authentication services as well as caching for off-line use.
- Support for the latest Crypto Express3 accelerator and coprocessor for offloading the processing of secure data.

## ■ **Development and Runtime Support**

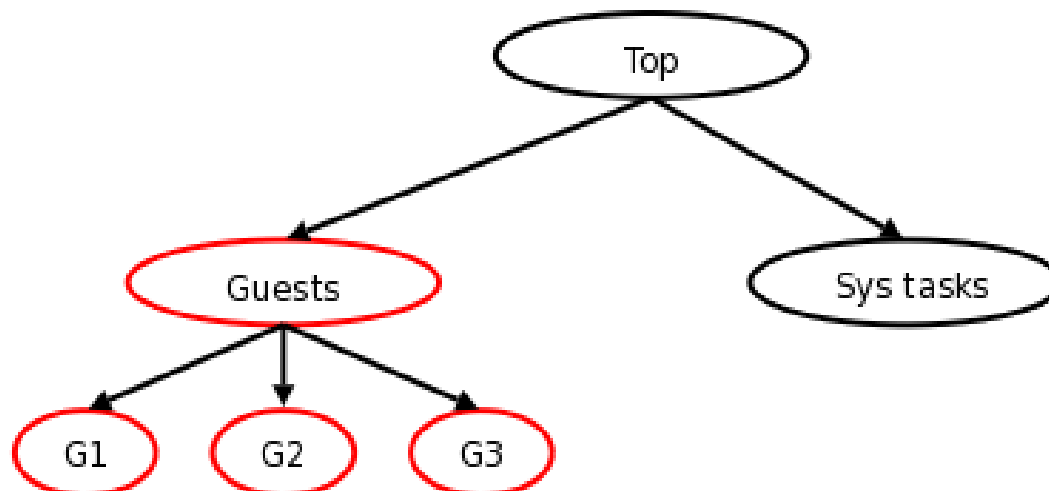
- SystemTap (allows instrumentation of a running kernel without recompilation) and ABRT (simple collection of bug information)
- Improvements to glibc (version 2.12), GDB (version 7.1), and the GCC compiler (version 4.4), which can lead to greater than 10% performance improvement.

# In Depth: Cloud Enablement with Control Groups

Problem: “I want to implement a chargeback model.”

Solution: Congrol Groups (cgroups)

- Cgroups are “process containers”. Lets you transform groups of applications into workloads





# In Depth: Cloud Enablement with Control Groups

- **Resource Limiting**

Specify limits on CPU, memory, and even file system usage

- **Prioritization**

Give mission critical workloads higher priority than others

- **Accounting**

Run report on resource utilization, i.e. for billing purposes

- **Isolation**

Separate namespaces for groups, so they don't see each other's processes, network connections or files

- **Control**

Freeze groups for checkpointing or restarting workloads

# Efficiency, Reliability & Scalability

- Efficient Scheduling
  - CFS (Completely Fair Scheduler) is a new application scheduler which operates with nanosecond granularity
  - Optimized for multi-core topologies, CFS maximizes CPU utilization with low power consumption
- Reliability
  - With support for processor degradation and other machine checks, the system can recover from fatal hardware errors with minimal disruption
  - Memory pages with errors can be declared as “poisoned” and will be avoided
- Scalability
  - Full support for up to 80 IFLs and up to 3 TB memory on z196
  - Support for up to 4 million processes per server

# Filesystems

- Ext4 replaces Ext3 as default filesystem
  - Faster, more robust, and scales to 16 TB
  - Support for thin provisioning. You can create a 16 TB file system on a single mod-3 DASD! Just add more when needed.
- NFS version 4 for network file system
  - Clustered file system with support for read/write access from multiple guests simultaneously.
  - Use VSWITCH for fast “network” access, or Hipersockets for memory-speed transfers.
- Fuse
  - Allows filesystems to run in user space, allowing testing and development of fused-based filesystems.
  - For example, use Fuse to create a cloud filesystem.

# In Depth: Next Generation BTRFS Filesystem

- Provides enhanced data integrity, better performance, and ease of use over Ext3/Ext4 filesystems.
- Key Features
  - Copy-on-write for cloning and snapshots, with filesystem seeding
    - Create a filesystem to seed other BTRFS filesystems. The original filesystem is a read-only starting point, the copy-on-write ensures the original is unchanged.
  - Online defragmentation
  - Load balancing across device nodes
  - Transparent data compression
  - In-place conversion (with rollback) from Ext3/Ext4
  - Data deduplication (under development)
- Principle developer for Ext3/Ext4 stated BTRFS is the way forward, having a number of the same design ideas that Reiser had.



# LVM, Storage & Multipath

## ■ LVM snapshot improvements



- Snapshot can now be merged back into the original logical volume, reverting changes that occurred after the snapshot.
- Practical application: “lock” your data with regular snapshots, then if necessary “undo” any changes within the last week/day/hour..

## ■ Define LVM hot spare

- Recover seamlessly from device failure of a disk or logical volume

## ■ Device Mapper Multipath

- Allows paths to be dynamically selected based on queue size or I/O time data

## ■ Performance and Integrity

- Automated I/O alignment and self-tuning
- DIF/DIX provides better integrity checks for application data

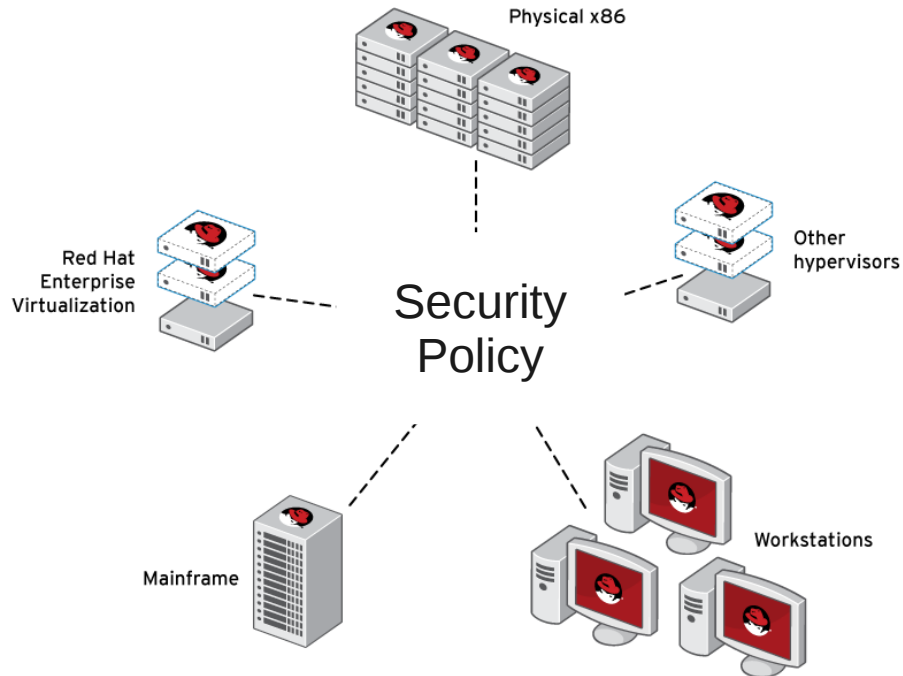


- UDP Lite tolerates partially corrupted packets to provide better service for multimedia protocols, such as VOIP, where partial packets are better than none.
- Multiqueue Networking increases processing parallelism for better performance from multiple processors and CPU cores.
- Large Receive Offload (LRO) and Generic Receive Offload (GRO) aggregate packets for better performance.
- Support for datacenter bridging includes data traffic priorities and flow control for increased quality of service.
- iSCSI partitions may be used as either root or boot filesystems.

# IPA Client (Identity, Policy, Audit)

## ■ IPA Client in Core RHEL for Centralized Security Management

- Kerberos authentication with host based access control
- Provides central storage of extended user attributes
- Enables centralized policy for applications, including SELinux policy
- Audit log aggregation services & search capabilities



# Security enhancements

## ■ **Virtualization Isolation in Conjunction with SELinux**

- Labeled NFS for filesystem isolation, guest confinement via SELinux policy enhancements

## ■ **NSS Crypto**

- Broaden the core services which utilize NSS crypto libraries
- Allows cheaper implementation of new features, ie TPM & centralized key store
- Incremental targeted conversion of: Openswan, openldap, glibc
- Add new crypto GUI for key import & establishment of trust

## ■ **Volume Encryption**

- Basic operation already present in RHEL5 – incremental centralized key management for RHEL 6

## ■ **Sectool**

- Compliance checking / intrusion detection utility – validates system admin config: file permissions, valid UIDs, reasonable passwords, etc.

## s390-tools version 1.8.2

- This package provides *the* essential tool chain for Linux on System z. It contains everything from the boot loader to dump related tools for a system crash analysis.
- New Features
  - DASD related tools: Add Large Volume Support for ECKD DASDs
  - IPL tools: Can be used to change the reipl & shutdown behavior
  - ziomon tools: Set of tools to collect data for zfcg performance analysis.
  - Isluns: List available SCSI LUNs depending on adapter or port.
  - Iszcrypt/chzcrypt: Show/modify information about zcrypt devices and configuration

# s390-tools version 1.8.2

## ■ New Features (continued)

- **cpuplugd**: Daemon that manages CPU- and memory-resources based on a set of rules. Depending on the workload CPUs can be enabled or disabled. The amount of memory can be increased or decreased exploiting the Cooperative Memory Management (CMM1) feature.
- **lschp/chchp**: Tool to show/modify channel-path states and available channel paths.
- **mon\_procd**: Daemon that writes process information data to the z/VM monitor stream.
- **vmur**: Tool to work with z/VM spool file queues (reader, punch, printer).
- **zfcpdump\_v2**: Version 2 of the zfcpdump tool. Now based on the upstream Linux kernel 2.6.23.

## ■ Plus various bug fixes

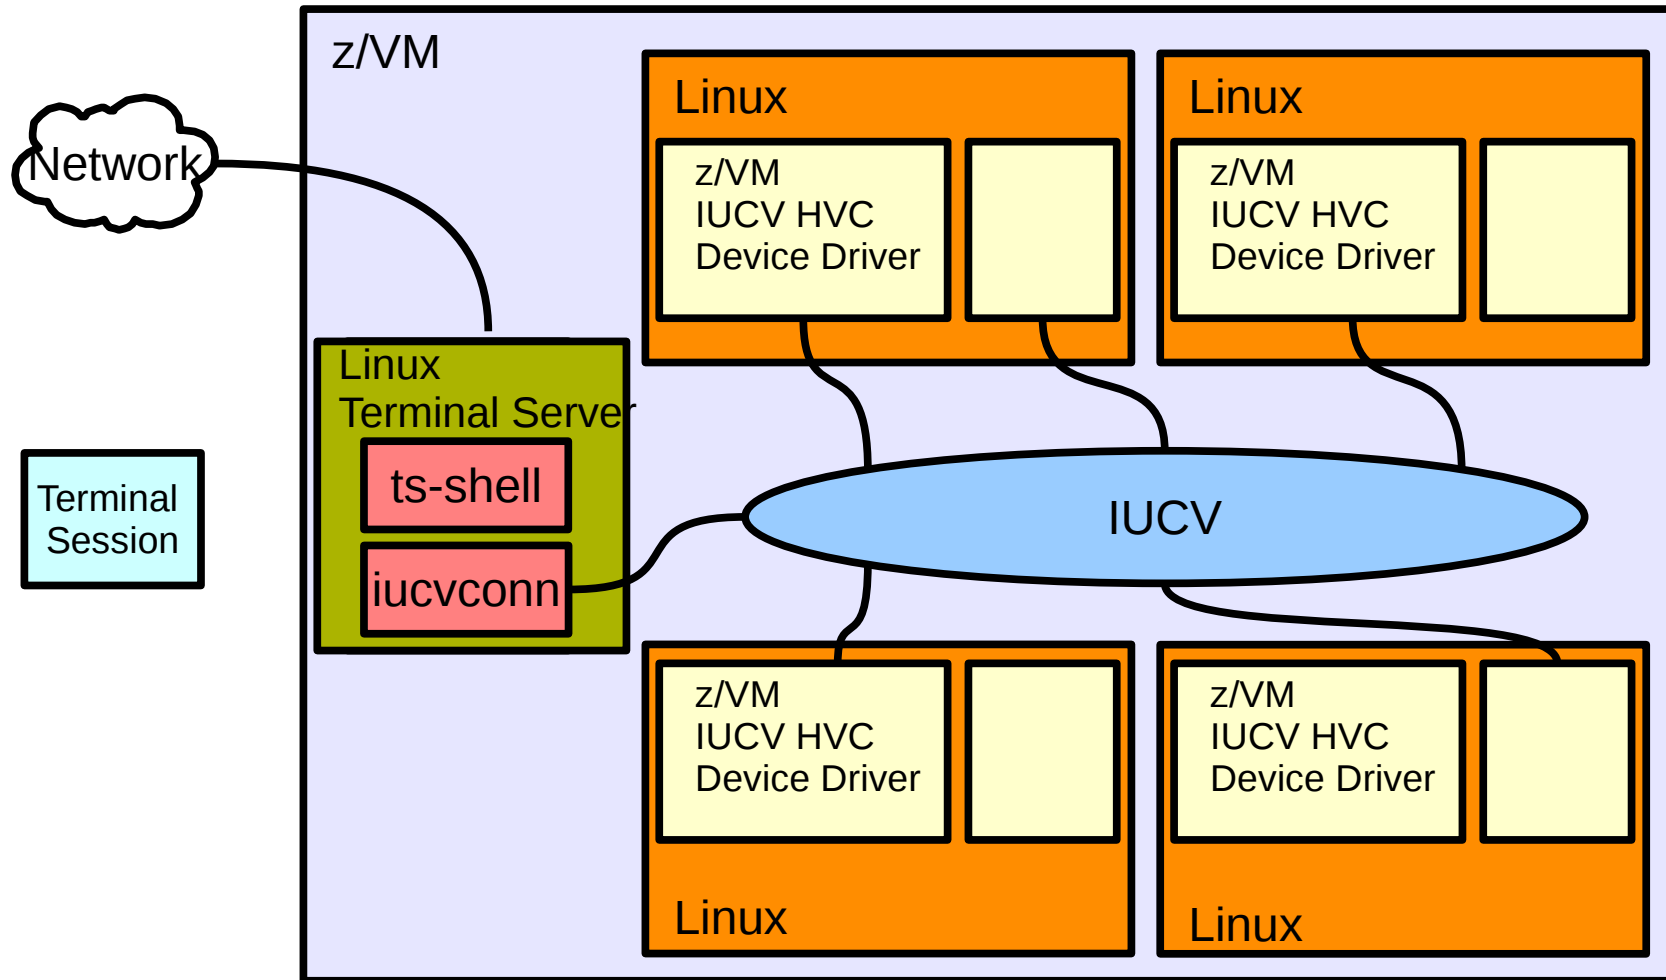


# Advanced Configurations: TTY Terminal

## ■ TTY Terminal Server Over IUCV

- Provide central access to the Linux console for the different guests of a z/VM.
- The terminal server connects to the different guests over IUCV.
- The IUCV based console is ASCII based.
- Fullscreen applications like *vi* are usable on the console.


# TTY Terminal Server Over IUCV



# Advanced Configurations: Shared Kernel

## ■ Named Saved Segments (NSS)

- Using NSS the z/VM hypervisor makes operating system code in shared real memory pages available to z/VM guest virtual machines.
- With this update, Linux guest operating systems using z/VM can boot from the NSS and be run from a single copy of the Linux kernel in memory.




# Documentation Links

- Redbook, “z/VM and Linux on IBM System z: The Virtualization Cookbook for Red Hat Enterprise Linux 6”
  - Coming soon (in final editing)
  - Draft
    - <http://www.vm.ibm.com/devpages/mikemac/SG247932.pdf>
    - <http://www.vm.ibm.com/devpages/mikemac/SG247932.tgz>
- DeveloperWorks:
  - [http://www.ibm.com/developerworks/linux/linux390/documentation\\_dev.html](http://www.ibm.com/developerworks/linux/linux390/documentation_dev.html)
- Knowledgebase:
  - <http://kbase.redhat.com/>
  - Search “s390”
- <http://www.redhat.com/z>
- <http://www.redhat.com/rhel/server/mainframe/promo/>



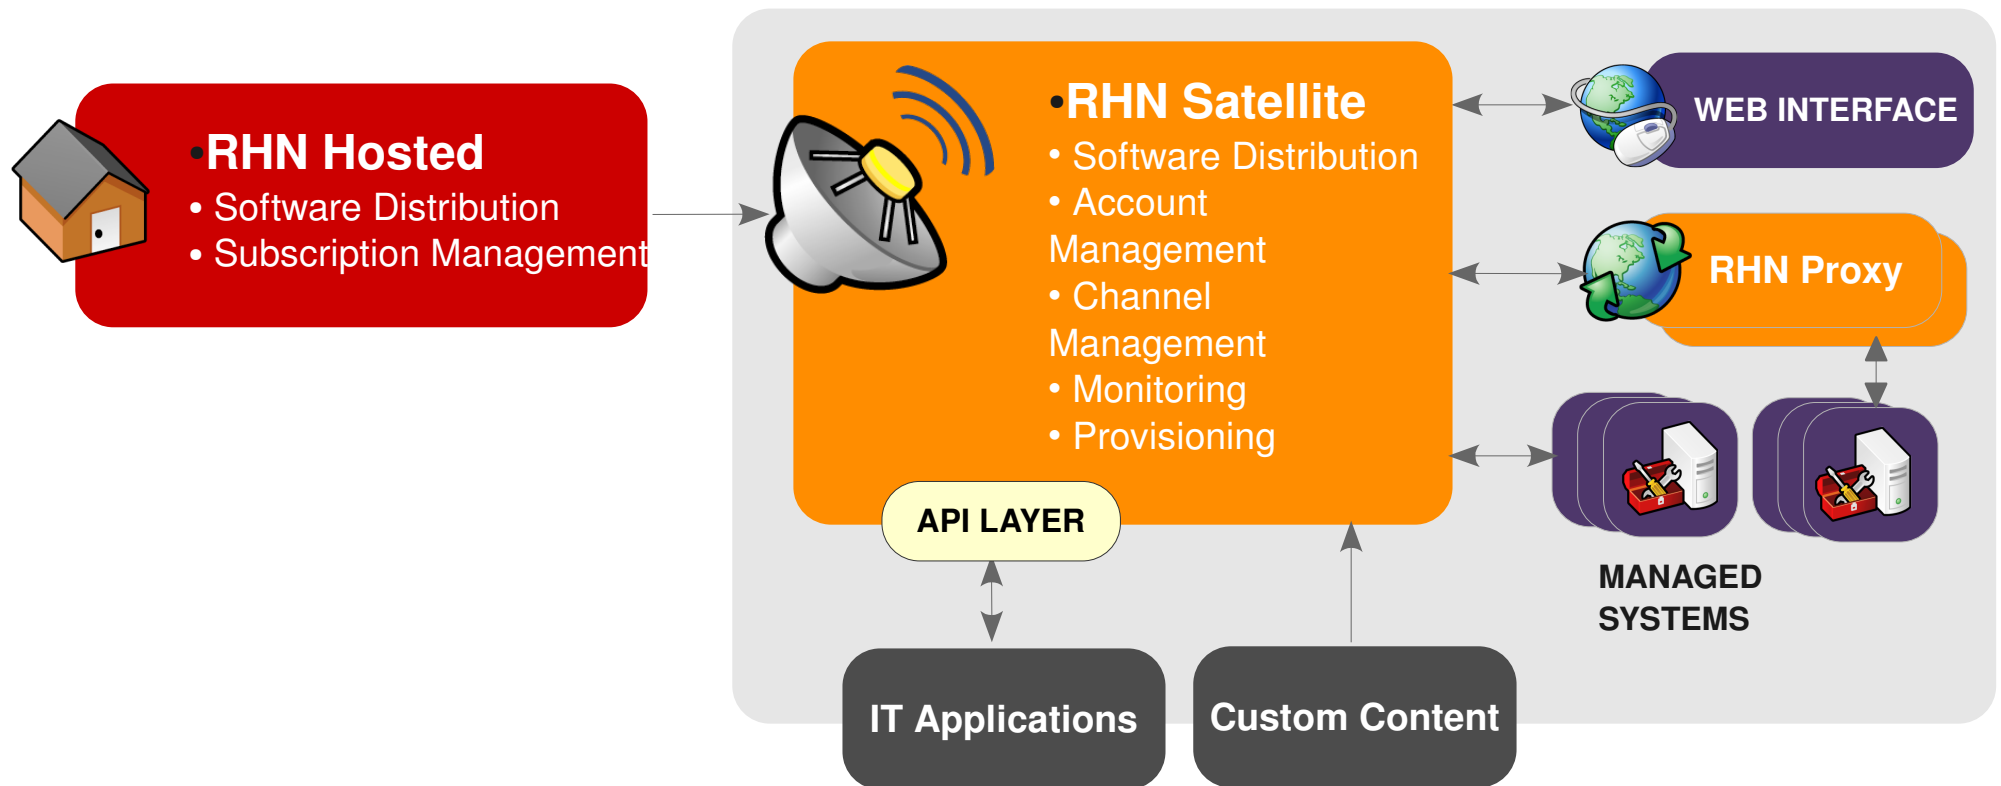
Brad Hinson  
Red Hat, Raleigh, NC  
Worldwide Lead, Linux on System z  
[bhinson@redhat.com](mailto:bhinson@redhat.com)  
+1 (919) 360-0443

# Appendix: Systems Management with RHN Satellite





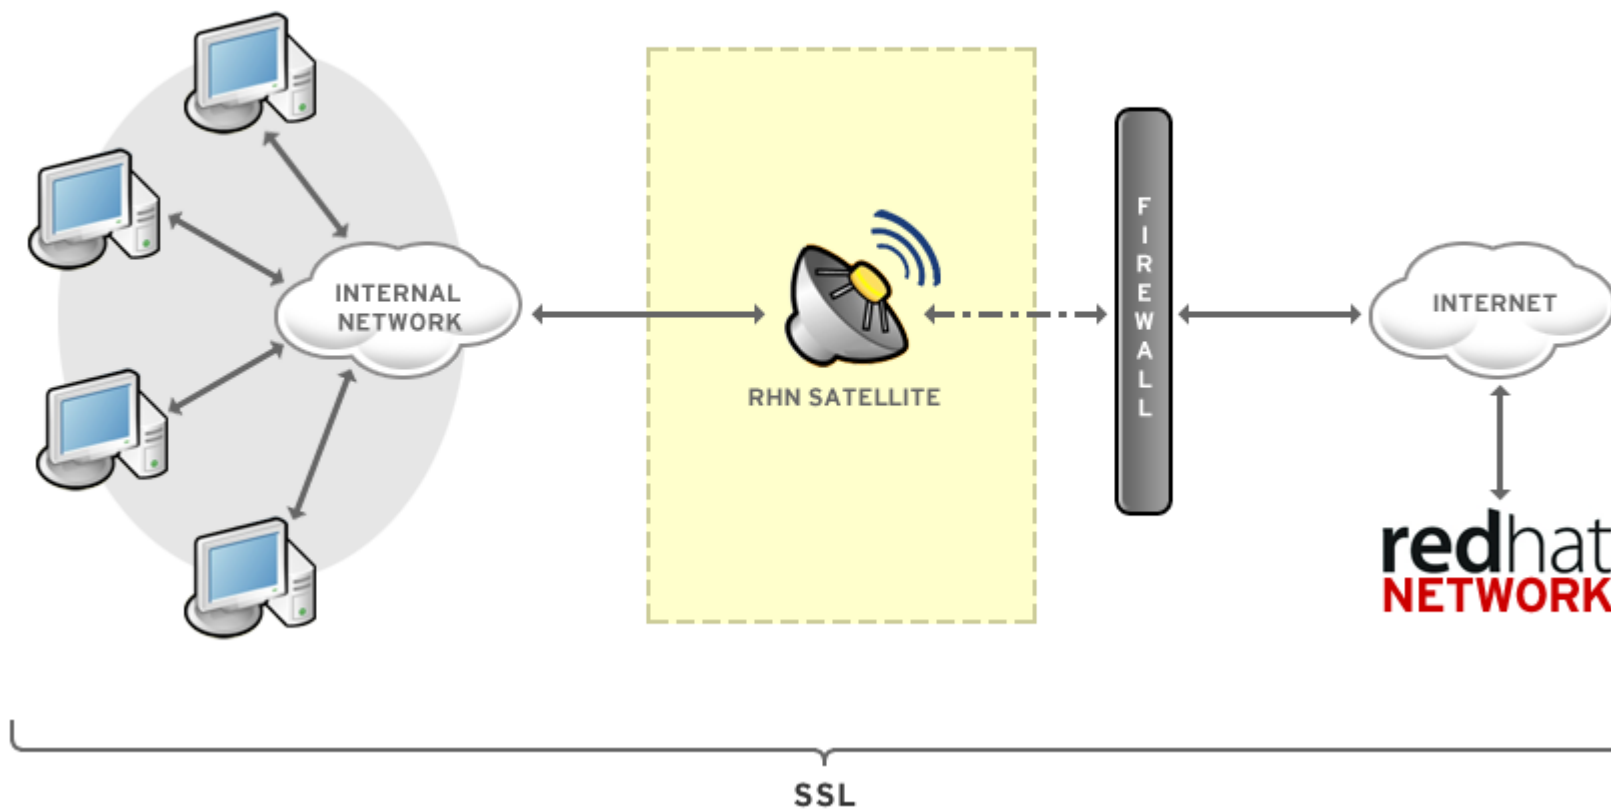
# Satellite deployment model



- Enterprise management solution – enhanced control
- Local database stores all packages, profiles, and system information
- Syncs content from RHN Hosted
- Custom content distribution
- Can run disconnected from the Internet

# RHN Satellite


**RHN SATELLITE**  
Single Satellite Topology Example



# RHN Satellite on System z

## RHN SATELLITE-PROXY

Satellite-Proxy System z Topology Example



# PXE Deployment on System z

## zPXE

- Same configuration/profile on all guests
- Read-write 191 disk not required for each guest
- All changes kept on management server
- Flexibility of kickstart
- Same principles as traditional PXE, adapted to System z
- Fits with configuration management tools (cobbler)
- Easy to update

<https://fedorahosted.org/cobbler/wiki/SssThreeNinety>

# Appendix:

## RHEL 6 on System z, Detailed Features



# Advanced Virtualization

- **Dynamic Memory Add/Remove** (kernel 2.6.27)
  - Attach and use standby memory that is configured for a logical partition or z/VM guest.
  - Memory Attach & Detach requires running Linux on System z as a VM-guest requires z/VM 5.4 plus the PTF for APAR VM64524.
- **Standby CPU Activation/Deactivation** (kernel 2.6.25)
  - Allow standby CPUs to be activated / deactivated
- **Extra Kernel Parameter for SCSI IPL** (kernel 2.6.32)
  - Modify the SCSI loader to append extra parameters specified with the z/VM VMPARM option to the kernel command line.



# Advanced Virtualization

- **hvc\_iucv: Provide IUCV z/VM User ID Filtering** (kernel 2.6.29)
  - Introduces the kernel parameter "hvc\_iucv\_allow=" that specifies a comma-separated list of z/VM user IDs.
  - If specified, the z/VM IUCV hypervisor console device driver accepts IUCV connections from listed z/VM user IDs only.
- **Suspend / Resume** (kernel 2.6.31)
  - With suspend and resume support, you can stop a running Linux on System z instance and later continue operations.
  - When Linux is suspended, data is written to a swap partition. The resume process uses this data to make Linux continue from where it left off when it was suspended.
  - A suspended Linux instance does not require memory or processor cycles.

# Suspend/Resume Support

- Ability to stop a running Linux on System z instance and later continue operations
- Memory image is stored on the swap device specified with a kernel

parameter: `resume=/dev/dasd<x>`

- ```
# grep swap /etc/fstab
/dev/dasdd1 swap swap pri=-1 0 0
/dev/dasde1 swap swap pri=-2 0 0
```

- Suspend operation is started with “echo” command to sysfs:

```
# echo disk > /sys/power/state
```

- Resume is done automatically on IPL
- Use signal (Ctrl+Alt+Del) to automatically suspend a guest:

```
ca::ctrlaltdel:/bin/sh -c "/bin/echo disk > \
/sys/power/state || /sbin/shutdown -t3 -h now"
```

# Storage Support

- **FCP Automated Port Discovery** (kernel 2.6.25)
  - Scan the connected fiber channel SAN and automatically activate all available and accessible target ports. This requires a proper SAN setup with zoning.
- **FCP SCSI Error Recovery Hardening** (kernel 2.6.30)
  - Avoid SCSI error recovery escalation in case of concurrent zfc and SCSI error recovery.
- **FCP Adjustable Queue Depth** (kernel 2.6.31)
  - Customizable queue depth for SCSI commands in zfc.

# Storage Support

## ■ **HyperPav** (kernel 2.6.25)

- HyperPav is addressing the need to access more data with good performance and high availability!
- This feature, which required a IBM DS8000™ disk storage system in average leads to a higher utilization, resulting in I/O transfer rates.
- Activated automatically when the necessary prerequisites are there (DS8000 with HyperPAV LIC, z/VM 5.3). Transparent for the Linux on System z guest

## ■ **I/O Configuration Support** (kernel 2.6.27)

- Adds the infrastructure to allow Linux system to change the I/O configuration of a System z system.
- Operations are addition, removal and reconfiguration/reassignment of I/O channels, control units and subchannels.
- This support is available only when running Linux on System z in an LPAR.

# Storage Support

- **High Performance FICON (HPF)** (kernel 2.6.29)
  - Added HPF support to the DASD Device Driver
  - HPF is an extension to the FICON architecture and is designed to improve the execution of small block I/O requests.
  - HPF streamlines the FICON architecture and reduces the overhead on the channel processors, control unit ports, switch ports, and links by improving the way channel programs are written and processed.
- **FCP SCSI Error Recovery Hardening** (kernel 2.6.30)
  - Avoid SCSI error recovery escalation in case of concurrent zfcps and SCSI error recovery.
- **DASD Large Volume Support** (kernel 2.6.29)
  - Large Volume Support is a feature that allows to use ECKD devices with more than 65520 cylinders. This features is available with DS8000 R4.0



# Networking

- **HiperSockets Network Traffic Analyzer** (kernel 2.6.33)
  - Trace HiperSockets network traffic for problem isolation and resolution.  
Supported for layer 2 and layer 3
- **Crypto Express 3** (kernel 2.6.33)
  - Support for Crypto Express3 Accelerator (CEX3A) and Crypto Express3 Coprocessor (CEX3C)
  - z/VM 6.1, 5.4, or 5.3 with the PTF for APAR VM64656 is required for z/VM guest-support.
- **Secondary Unicast Addresses** (kernel 2.6.27)
  - Allow secondary unicast MAC addresses to support MAC address based VLANs
  - This only works with an OSA interface running in layer 2 mode.

# Networking

## ■ **OSA QDIO Data Connection Isolation** (kernel 2.6.33)

- Feature (available for OSA-Express2 cards since early 2009) allows an operating system to configure the OSA adapter to prevent any direct package exchange between itself and other operating system instances that share the same OSA adapter.
- The configuration of the isolation level is done through sysfs.
- Starting with s390-tools 1.8.4 lsqeth indicates connection isolation in qeth attribute 'isolation'.

## ■ **QETH Componentization** (kernel 2.6.25)

- The qeth driver module is split into a core module and layer2-/layer3-specific modules. The default operation mode for OSA-devices is changed to layer2; for HiperSockets devices the layer3 default-mode is kept.
- For layer3 mode devices the existence of (possibly faked) ethernet-headers is guaranteed to enable smooth integration of qeth devices into Linux.

# RAS: Reliability, Availability, Serviceability

- **Automatic IPL After Dump** (kernel 2.6.30)

- Extension to the shutdown action interface which combines the actions dump and re-ipl, first a dump is taken, then a re-ipl of the system is triggered

- **Compiler Improvements** (gcc 4.4)

- The latest compiler enhancements allow a customer to recompile existing applications which can be optimized for the latest hardware generation without any changes to the source code.
- This can lead up to a > 10 % performance improvement.

- **Large Page Support** (kernel 2.6.25)


- Support for a new access method to allocate larger chunks of memory (1 MB page size), resulting in performance improvements, especially in Java based environments
- This feature exploits z10 hardware features and provides a software emulation for older systems.

# RAS: Reliability, Availability, Serviceability

- **Add Call Home Data on Halt and Panic if Running in LPAR** (kernel 2.6.32)
  - Report system failures (kernel panic) via the service element to the IBM service organization.
  - Improves service for customers with a corresponding service contract. (by default this features is deactivated)
- **CIO, DASD: Improved DASD Error Recovery** (2.6.33)
  - Improved the DASD error recovery procedures used in the early phases of IPL and DASD device initialization.

# Miscellaneous

- **Kernel VDSO Support** (kernel 2.6.29)
  - Kernel provided shared library to speed up a few system calls (gettimeofday, clock\_gettime, clock\_getres)
- **Kernel Image Compression** (kernel 2.6.33)
  - The kernel image size can be reduced by using one of three compression algorithms: gzip, bzip2 or lzma.
- **STP/ETR Support** (kernel 2.6.27)
  - Support for clock synchronization using the server time protocol (STP) or an external





# Miscellaneous

## ■ KULI (2009-06-24)

- kuli is experimental userspace sample to demonstrate that KVM can be used to run virtual machines on Linux on System z.
- This experimental proof of concept is unsupported and should not be used for any production purposes.

## ■ Oprofile

- Starting with version 0.9.4, oprofile supports sampling of Java byte code applications for Linux on System z.

## ■ Eclipse 3.3

- Starting with Eclipse 3.3, Linux on System z is officially supported.

# Kernel Scalability

## ■ Scalability Limits – Maximum Values

- Max CPUs – dynamic allocation of CPU structs (2.6.29) allows supported limit of 4096 and theoretical limit of 64K.
- Max IRQs – dynamic allocation (2.6.28) allows limit of 256
- Max memory 48-bit addressing (256TB) requires pending patch incorporation
- Max # processes – 4 million on 64 bit kernels
- Max threads per process remains at 32000 (same as RHEL5)
- Filesystem limits for ext4 – 100T (practical target – bounded by fsck time)

## ■ Scalability Features

- Split LRU VM – different eviction policies for file backed vs swap backed

# Appendix: Customer Examples



# Who is Using RHEL on System z in Production?

## Merrill Lynch

- Leading wealth management company, capital markets, advisory company
- Offices in 38 countries
- Assets > \$1.7 trillion
- 5 IFLs running WebSphere & MQ-Series

## VISA

- Largest electronic retail payment systems in the world
- 3 IFLs for VISANet – Custom Application

## Bank of Tokyo

- Largest known Linux on System z deployment in the world
- 400 IFLs for ATM network
- Mostly WebSphere

## Fratelli Carli

- Olive Oil manufacturer in Italy
- 1M orders per year, nearly 3,000 per day
- Response times < 1sec with 24GB of data



## Reduce Cost, Manage Risk and Improve Service with Linux on System z

**Selects an IBM System z and Linux to Support Growing Business in Key Emerging Market**



### **Vietnam Joint Stock Bank**

**(Vietnam):** VietinBank selected an IBM mainframe to run a new operational risk management solution, which will enable the bank to deliver premium services to its customers. "IBM System z10 is regarded as the strategic platform for mission-critical work. Nothing is more mission-critical than ensuring safe money deposits and transfers in a country like Vietnam experiencing exponential growth," said Mr. Pham Anh Tuan, Deputy Director of VietinBank.

[ibm.com/press/us/en/pressrelease/32082.wss](http://ibm.com/press/us/en/pressrelease/32082.wss)

**An IBM System z mainframe as part of a cloud computing project to help local businesses**



### **University of Bari (Italy):**

Delivering services through an IBM System z mainframe and cloud computing allows multiple entities to tap into heavy-duty computing power at minimal cost and lowers the barrier to help local businesses to benefit from this technology. The cloud solution is built on an IBM System z9® Business Class mainframe running the Linux operating system, and includes DB2, WebSphere and Tivoli® middleware from IBM.

[ibm.com/press/us/en/pressrelease/32051.wss](http://ibm.com/press/us/en/pressrelease/32051.wss)

**Enterprise-wide business transformation and lay the foundations for future growth**



### **Dongbu Insurance (South Korea):**

Four IBM System z10 Enterprise Class and Business Class mainframe servers will serve as the backbone of the system, running on z/OS® and Linux operating systems as well as IBM DB2 and WebSphere Application Server. "Based on IBM's technology performance and price, we have selected IBM as the provider of Dongbu Insurance's IT platform," said Mr Kun-gyo Lee, Chief Information Officer, Dongbu Insurance.

[ibm.com/press/us/en/pressrelease/29954.wss](http://ibm.com/press/us/en/pressrelease/29954.wss)



# More Public References and Case Studies...

## Salt River Project

<http://customers.redhat.com/2008/10/17/salt-river-project-migrates-to-red-hat-enterprise-linux-on-ibm-mainframes-for-flexibility-and-performance/>

## EDB (Nordics)

<http://customers.redhat.com/2008/06/10/edb-implements-red-hat-solutions-for-the-reliability-of-in-house-mainframe-computing/>

## Bank of New Zealand

<http://customers.redhat.com/2009/02/03/bank-of-new-zealand-reduces-carbon-footprint-with-red-hat-on-the-mainframe/>

## Allianz

<http://customers.redhat.com/2009/09/16/allianz-australia-migrates-to-red-hat-and-jboss/>

## Banco Pastor

<http://customers.redhat.com/2009/10/12/banco-pastor-reduces-costs-and-improves-scalability-with-red-hat-sap-and-ibm-solutions/>

## City of Recife, Brazil

<http://customers.redhat.com/2008/06/10/brazil-city-of-recifes-computer-technology-organization-finds-long-term-reliability-with-red-hat-solutions/>

## Handelsbanken

<http://customers.redhat.com/2009/04/22/handelsbanken-offers-21st-century-banking-services-using-ibm-and-red-hat-solutions/>

## Citibank

<http://customers.redhat.com/2010/06/17/citigroup-red-hat-innovation-award-winner/>

## Anses (JBoss and System z)

<http://customers.redhat.com/2010/06/17/anses-jboss-innovation-award-winner/>


## Fratelli Carli

<http://customers.redhat.com/2008/01/08/italian-olive-oil-manufacturer-enhances-performance-with-red-hat-enterprise-linux-5-on-ibm-mainframe-zseries/>

## City of Navarra, Spain



# A Deeper Dive – Bank of New Zealand





# A Deeper Dive – Bank of New Zealand

## Bank of New Zealand

- First production z10 in southern hemisphere
- IBM/RHT offered 1 year of proof of concept support
- 1 Year of planning, 3 days to go production!
- Converted from Solaris (E10K, v440/v480, v280Rs)

|                      |                  | x86 | System z  |
|----------------------|------------------|-----|-----------|
| SWIFT:               |                  |     |           |
| \$10,000,000,000/day | Power (kW/hr)    | 36  | 22 (38%)  |
| PCBB:                |                  |     |           |
| \$40,000,000/day     | Heat (BTUs)      | 110 | 74 (33%)  |
|                      | Space (Racks)    | 6.5 | 4.5 (31%) |
|                      | Carbon Emissions | 66  | 40 (39%)  |



# Bank of New Zealand: Price/Performance of RHEL on z with Oracle DB

## Case Study: Oracle on RHEL on System z savings

### Oracle License =

(CPU Cores) \* (Processor Factor) \* (Installed DB #) \* License

### Distributed

Oracle =  $8 * 2.7 * 50 * 7,000 = \$7,560,000$

### Mainframe

Oracle =  $1 * 1 * 100 * 7,000 = \$700,000$

~\$6.8M saved **only** in Oracle licenses!

Thank You

