

RED HAT
ON THE MAINFRAME
THE REASONS ARE COMPELLING



Current & Future State of Linux on System z

Agenda & Introduction

- **Red Hat System z Business Update**
- **RHEL 5.3 Update (released 20-JAN 2009)**
 - What's new?
 - What's new specifically for System z?
- **Future Tech / Upstream Development Efforts**

Agenda & Introduction

Shawn Wells <swells@redhat.com>

Global System z Alliance Manager

(+1) 443 534 0130

- Based in Washington, D.C.
- Global responsibility for Red Hat's System z activities

System z Business Update


Red Hat / IBM Relationship

- Cross platform relationship founded in the late 90s (when Red Hat incorporated)
- Started releasing RHEL for s390 in 2001
- Formal Linux on System z agreement & announcement in 2007
(<http://www-03.ibm.com/press/us/en/pressrelease/21513.wss>)
- Red Hat has dedicated staff to System z
(we haven't done this for any other H/W platform)



Red Hat Development & Subscription Model


Red Hat Development Model



COMMUNITY

- Development with “upstream communities”
- Kernel, glibc, Apache, etc
- Collaboration with open source community; individuals, business partners, customers

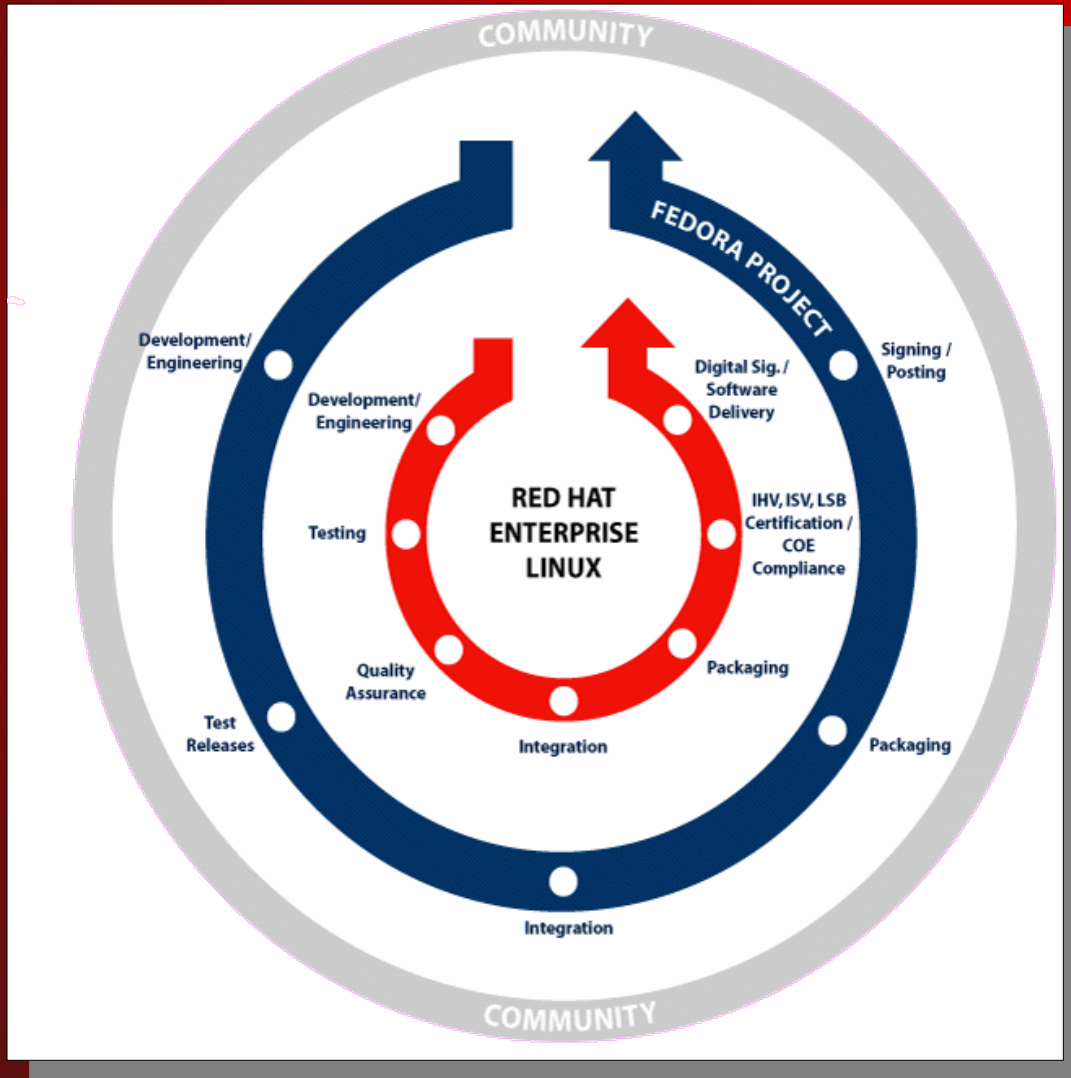
Red Hat Development Model



FEDORA

- Bleeding edge
- Sets technology direction for RHEL
- Community supported
- Released ~6mo cycles
- Fedora 8,9,10 = RHEL6

Red Hat Development Model




RHEL

- Stable, matured
- Q&A, testing
- H/W & S/W Certifications
- 7yr maintenance
- Core ABI compatibility
- Major releases 2-3yr cycle

Joint Red Hat / IBM Development

Kernel Lines Changed - 2.6.23



RHEL 5.2 Tech Deep Dive

RHEL 5.2: Technical Review

- **Accelerated in-kernel Crypto**
 - Support for crypto algorithms of z10 (SHA-512, SHA-384, AES-192, AES-256)
- **Two OSA ports per CHPID; Four port exploitation**
 - Exploit next OSA adapter generation which offers two ports within one CHPID. The additional port number 1 can be specified with the qeth sysfs-attribute “portno”

Support is available only for OSA-Express3 GbE SX and LX on z10, running in LPAR or z/VM guest (PFT for z/VM APAR VM64277 required!)

RHEL 5.2: Technical Review

- **SELinux per-package access controls**
 - Replaces old packet controls
 - Adds secmark support to core networking
- **Add nf_conntrack subsystem**
 - Allows IPv6 to have stateful firewall capability


Enables analysis of whole streams of packets, rather than only checking the headers of individual packets

Audit Subsystem Support for process-context based filtering

More filter rule comparators

RHEL 5.3 Tech Deep Dive

RHEL 5.3 Overview



~150 additions, ~3,400 BugZillas

- **7% FasTrack**
Early release of low impact fixes
- **7% Hardware Enablement**
New chipsets & processor feature support
- **21% New Features**
Feature requests from customers & partners
- **65% "Other"**
Feature enhancements, Bug fixes, Documentation

RHEL 5.3: Technical Review

- **Highlights**

- Added RAID 4/5/10 in dm-raid
- DHCPv6 Support
- Inclusion of OpenJDK
 - Full open source JDK for Java 1.6 support
 - Tested with Java SE 1.0 Technical Compatibility Kit (TCK) ==> 100%
 - x86 and x86_64 architectures only!
- Root (/) and SWAP encryption support in the installer

RHEL 5.3: Technical Review

- **Highlights, cont**
 - Improved Audit & Logging
 - TTY Input audit support

RHEL 5.3: System z Specific

BugZilla ID	Summary
46327	stage1: sshd error loading shared lib: libfipscheck.so.1
184770	LTC18425-62140: (big) xDR system Initialization for LPAR Clients
472788	rhel 5.3 snapshot3 scsi mpath install failed on z9bc lpar
439479	LTC:5.3:201474:Include gcc 4.3 as Add-On for latest z10 instruction set support
439440	LTC:5.3:201160:Long Random Numbers Generation
439441	LTC:5.3:201158:Selective Logging of ECKD DASD devices
439482	LTC:5.3:201542:FCP - Enhanced Trace Facility
447379	LTC:5.3:200994:Linux CPU Node Affinity
463917	unable to find DASD drives to install
439484	LTC:5.3:201490:Libica Library: Integration of lcainfo
43946	LTC:5.3:201360:OSA 2 Ports per CHPID Support - Installer Enhancements
466474	[RHEL5.3] *** glibc detected *** /usr/bin/python: double free or corruption (!prev): 0x000 0000080d55e90 ***
466305	cosmetic error message: failure in nl_set_device_mtu
466291	anaconda silently omits uninitialized disk

RHEL 5.3: Technical Review

GCC 4.3 Inclusion (latest z10 instruction support)

(Red Hat BugZilla [439479](#), IBM BugZilla [43379](#))

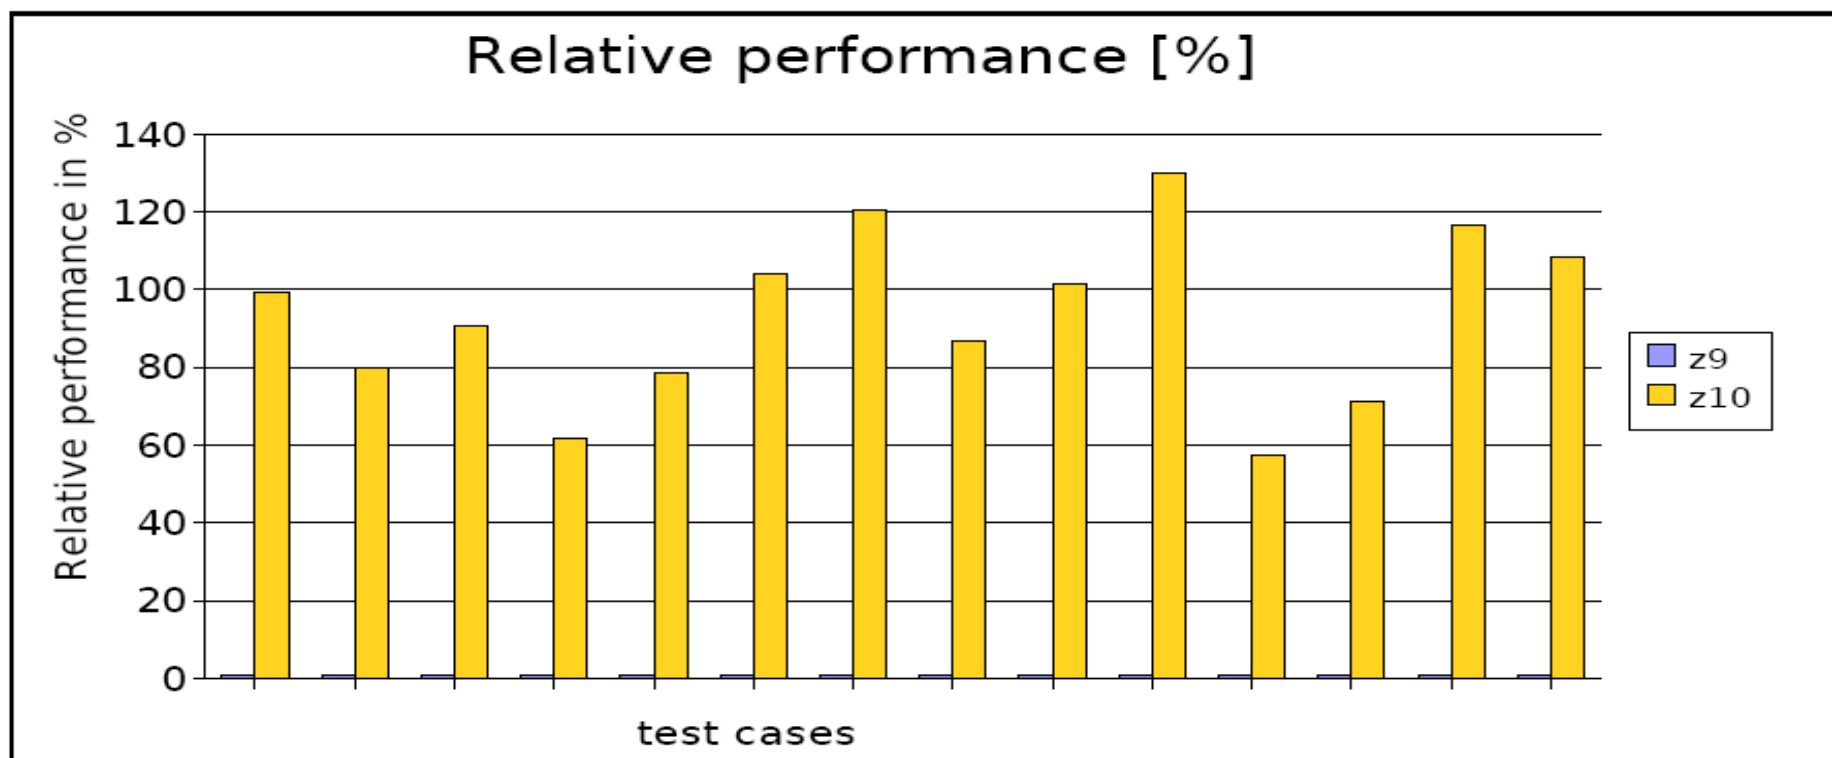
- Includes the following z10 specific patches to GCC
 - Introduce TARGET_MEM_CONSTRAINT macro
 - Introduce 'enabled' insn attribute
 - S/390: Exploit the 'enabled' insn attribute
 - S/390: Replace 'm' with 'RT' constraints
 - S/390: Add the -march=z10/-mtune=z10 options for z10
 - S/390: Support the new instructions introduced with z10
 - S/390: z10 pipeline description
 - PR36822 recog: Reorder extra memory constraint checks for inline assemblies
 - S/390: Fix -march=z9-ec -msoft-float

RHEL 5.3: Technical Review

GCC 4.3 Inclusion (latest z10 instruction support)

(Red Hat BugZilla [439479](#), IBM BugZilla [43379](#))

- Includes the following z10 specific patches to GCC
 - Overall improvement with z10 versus z9: 1.9x



RHEL 5.3: Technical Review


Long Numbers Generation

(Red Hat BugZilla [439440](#), IBM BugZilla [43340](#))

- Provides access to the random number generator on the crypto card in order to meet high volume random number requirements
- Frequently useful when high amount of SSL handshakes occur (JBoss, WebSphere, etc), or encryption/decryption (remember, encrypted memory is now supported!)
- Specific performance numbers not available at this time from Red Hat... but we do have IBMs.

RHEL 5.3: Technical Review

Long Numbers Generation (Red Hat BugZilla [439440](#), IBM BugZilla [43340](#))





RHEL 5.3: Technical Review

Long Numbers Generation

(Red Hat BugZilla [439440](#), IBM BugZilla [43340](#))

- The number of handshakes is up to 4x higher with HW support.
- In the 32 connections case we save about 50% of the CPU resources



RHEL 5.3: Technical Review

Enablement of ECKD DASD Sense Data

(Red Hat BugZilla [439441](#), IBM BugZilla [43339](#))

Sense Key	Name	Description
0h	No Sense	Indicates there is no specific Sense Key information to be reported for the disc drive. This would be the case for a successful command or when the ILI bit is one.
1h	Recovered Error	Indicates the last command completed successfully with some recovery action performed by the disc drive. When multiple recovered errors occur, the last error that occurred is reported by the additional sense bytes. Note: For some Mode settings, the last command may have terminated before completing.
2h	Not Ready	Indicates the logical unit addressed cannot be accessed. Operator intervention may be required to correct this condition.
3h	Medium Error	Indicates the command terminated with a non-recovered error condition, probably caused by a flaw in the medium or an error in the recorded data.
4h	Hardware Error	Indicates the disc drive detected a nonrecoverable hardware failure while performing the command or during a self test.
5h	Illegal Request	Indicates an illegal parameter in the command descriptor block or in the additional parameters supplied as data for some commands (Format Unit, Mode Select, and so forth). If the disc drive detects an invalid parameter in the Command Descriptor Block, it shall terminate the command without altering the medium. If the disc drive detects an invalid parameter in the additional parameters supplied as data, the disc drive may have already altered the medium. This sense key may also indicate that an invalid IDENTIFY message was received. This could also indicate an attempt to write past the last logical block.
6h	Unit Attention	Indicates the disc drive may have been reset.
7h	Data Protect	Indicates that a command that reads or writes the medium was attempted on a block that is protected from this operation. The read or write operation is not performed.
9h	Firmware Error	Vendor specific sense key.
h	Aborted Command	Indicates the disc drive aborted the command. The initiator may be able to recover by trying the command again.
Ch	Equal	Indicates a SEARCH DATA command has satisfied an equal comparison.
Dh	Volume Overflow	Indicates a buffered peripheral device has reached the end of medium partition and data remains in the buffer that has not been written to the medium.
Eh	Miscompare	Indicates that the source data did not match the data read from the medium.

RHEL 5.3: Technical Review

CPU Node Affinity

(Red Hat BugZilla [447379](#), IBM BugZilla [44875](#))

- Newer hardware (System z10 EC) supports an interface which can be used to get information about the CPU topology of an LPAR.
 - This can be used to optimize the Linux scheduler which bases its decisions on which process gets scheduled to which CPU on the CPU topology.
 - This feature should increase cache hits and therefore overall performance as well.

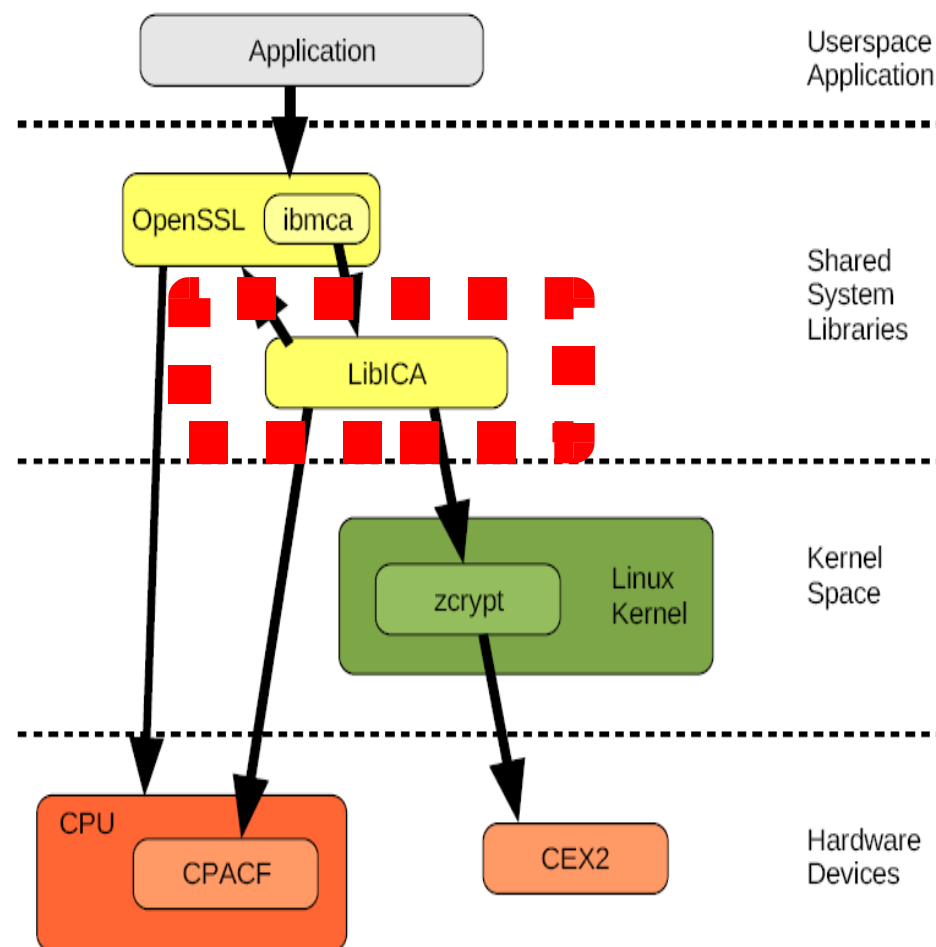
English Version: You dedicate 2 z10 IFLs to a RHEL5 VM. We can then pin applications to specific cores, or to IFLs in their entirety.

RHEL 5.3: Technical Review

Integration of icainfo into libICA

(Red Hat BugZilla [439484](#), IBM BugZilla [43383](#))

- icainfo is a part of the SHA & AES enhancements. It shows the customer which CPACF instructions are available in their system.
- libica allows customer applications to speed up cryptographic operations by using the CP Assist for Cryptographic Function (CPACF) facility.
- A new tool called 'icainfo' allows the customer to display a list of all CPACF operations supported by libica.
- This is helpful to verify that CPACF is correctly enabled on a particular system.



RHEL 5.3: Technical Review

OSA 2 Ports per CHPID Installer Support

(Red Hat BugZilla [439461](#), IBM BugZilla [43371](#))

- Anaconda now supports both ports on CHPID for OSA Express3 cards.
 - The installer will prompt for the port number in the initial stage of the installation.
 - The value provided for the port also affects installed network interface startup script. When port 1 is selected, the value "portno=1" is added to OPTIONS parameter of ifcfg-eth* file.

Note: When installing under z/VM, you can add either PORTNO=0 (to use port 0) or PORTNO=1 (to use port 1) to the CMS configuration file to avoid being prompted for the mode.

Planned Features

RHEL 5.4, RHEL6 & Beyond

RHEL 5.4: Works In Progress

- This list includes items currently under development, and is **not** a commitment to include features.
 - Is there something you must have? Let us know! It only took two customer request to back-port NPIV into RHEL 4.8. Your feedback matters!
 - If you have a BugZilla account (it's free!), [you can use this link](#) to view latest information
 - Don't have an account? Sign up at <http://bugzilla.redhat.com/>
- **Expected ETA: Mid-Late 2009**

RHEL 5.4: Works In Progress

- [200790] Kernel NSS support
- [200975] System z support for processor degradation
- [201801] Linux to add Call Home data - kernel part
- [201169] Automatic IPL after dump - kernel part
- [201174] Crypto Device Driver use of Thin Interrupts
- [201590] FCP - Performance Data collection – kernel part
- [201729] FCP - Performance data collection - blktrace
- [201726] Extra kernel parameter via VMPARM
- [201734] TTY terminal server over IUCV – kernel part
- [201747] Shutdown actions interface - kernel part
- [201751] HiperSockets Layer3 support for IPv6
- [201753] Provide service levels of HW & Hypervisor in Linux

RHEL 5.4: Works In Progress

[201801] Linux to add Call Home data - kernel part



- **IBM Link:** https://bugzilla.linux.ibm.com/show_bug.cgi?id=50307
- **Red Hat Link:** https://bugzilla.redhat.com/show_bug.cgi?id=475820

- **Description:**
 - ▶ The `scpl_cpi` sysfs interface allows the association of a set of descriptive data called "Control Program Identification" (CPI) with an operating system instance (currently LPAR only). This information is not persistent and has to be set once per IPL. CPI data will be visible at the HMC/SE. CPI information is also referred to as "Call Home data" in some contexts.
 - ▶ CPI information could previously be set by loading the `scpl_cpi` kernel module. The new CPI interface in sysfs complements the existing one. The new interface was included in Linux kernel 2.6.25.
 - ▶ More details under:
 - "Control program identification" in latest version of Device Drivers, Features, and Commands at: http://www.ibm.com/developerworks/linux/linux390/development_documentation.html

- **Business Value:**
 - ▶ RAS Improvement by decreasing mean time to resolve problems in case of System Failure.

- **Contents:**
 - ▶ Patch in kernel (only s390 specific)

- **Work responsibilities:**
 - ▶ Already tested by IBM internally
 - ▶ IBM will backport the patch from upstream.
 - ▶ Red Hat to accept/integrate/activate it in Kernel, and document it for RHEL 5.4.
 - ▶ Red Hat should verify that it works, since Red Hat has the HW to do functional test when integrating it.
 - ▶ IBM will test it again when it is delivered in RHEL 5.4 test phase.

RHEL 5.4: Works In Progress

[201676] Improve checking mechanisms and workflow of Linux on System z Anaconda install process
[201677] Dialog defaults for Linux on System z specific Anaconda
[201679] Change list of Anaconda network alternatives to indicate supported devices on System z



■ IBM Links:

- ▶ https://bugzilla.linux.ibm.com/show_bug.cgi?id=50342
- ▶ https://bugzilla.linux.ibm.com/show_bug.cgi?id=50343
- ▶ https://bugzilla.linux.ibm.com/show_bug.cgi?id=50344

■ Red Hat Link:

- ▶ https://bugzilla.redhat.com/show_bug.cgi?id=475346
- ▶ https://bugzilla.redhat.com/show_bug.cgi?id=475350
- ▶ https://bugzilla.redhat.com/show_bug.cgi?id=475345

■ Description:

- ▶ Provides a better syntactic, semantic and probing workflow, as well as meaningful error messages, to improve the user experience when installing RHEL on System z.

■ Business Value:

- ▶ These changes will improve the installation experience for customers by making the installation workflow more usable and efficient, which will result in an improvement of the customer satisfaction.

■ Contents:

- ▶ Patch to linuxrc.390 (part of anaconda) and new file (with list of network device types)

■ Work responsibilities:

- ▶ Already tested by IBM internally
- ▶ IBM will provide the patch
- ▶ Red Hat to accept/integrate it in Kernel, and document it for RHEL 5.4.
- ▶ Red Hat should verify that it works, since Red Hat has the HW to do functional test when integrating it.
- ▶ IBM will test it again when it is delivered in RHEL 5.4 test phase.

RHEL 6.0: Works In Progress

- Old [200322] xDR system Initialization for LPAR Clients
- Old [200092] Install from SCSI/FCP attached CD/DVD
- Old [201085] cio_ignore entry in generic.prm for LPARs
- Old [201092] Firstboot for System z
- Old [201360] OSA 2 Ports per CHPID Support – Installer and tools Enhancements
- [201594] Provide CMS script for initial IPL under z/VM
- [201624] Intuitive dump device configuration workflow and dialogue
- [201676] Improve checking mechanisms and workflow of Linux on System z Anaconda install process
- [201677] Dialog defaults for Linux on System z specific Anaconda
- [201679] Change list of Anaconda network alternatives to indicate supported devices on System z
- [201680] Adjust Anaconda Swap recommendations to Linux on System z specifics
- [201167] Linux to add Call Home data - Installer and tool enhancement
- [201674] Pick up latest version of s390-tools
- [200787] snIPL SCSI LOAD for LPAR
- [201666] ZFCP Performance Statistics – blktrace
- [201729] FCP - Performance data collection – blktrace
- [201598] FCP - HBA API followup for upstream
- [201675] Pick up latest version of libica
- [201744] Cleanup of libICA Crypto library
- [201465] Web 2.0 - Inclusion of package rubygems
- [201466] Web 2.0 - Inclusion of package rubygem-rails
- [201469] Web 2.0 - Inclusion of package memcached
- [201470] Web 2.0 - Inclusion of package Apache MyFaces Core (JSR 252 - JSF implementation - recommended Version 1.1.x)
- [201471] Web 2.0 - Inclusion of package perl-CGI-Session
- [201472] Web 2.0 - Inclusion of package mysql-connector-java
- [201554] Web 2.0 - Inclusion of package 'rubygem-rake'
- [201555] Web 2.0 - Inclusion of packages 'rubygem-actionpack', 'rubygem-activerecord', 'rubygem-activesupport', 'rubygem-actionmailer'
- [201556] Web 2.0 - Inclusion of packages 'rubygem-actionwebservice' and 'rubygem-tzinfo'
- [201558] Web 2.0 - Inclusion of package 'mod_security'

RHEL 6.0: Works In Progress

[200092] Install from SCSI/FCP attached CD/DVD



■ IBM Links:

- ▶ https://bugzilla.linux.ibm.com/show_bug.cgi?id=48131
- ▶ https://bugzilla.linux.ibm.com/show_bug.cgi?id=15609
- ▶ https://bugzilla.linux.ibm.com/show_bug.cgi?id=37839

■ Red Hat Link:

- ▶ https://bugzilla.redhat.com/show_bug.cgi?id=463218
- ▶ https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=184648

■ Description:

- ▶ It has been implemented in RHEL 5.3.
- ▶ Build Linux distribution such that it can be installed from a CD or DVD that is via SCSI/FCP directly attached a zSeries.
 - Build an image with "El Torito" Multiboot format.
 - When this image is loaded, it should be possible to choose from menu "Install from CD/DVD", getting all the packages from the CD/DVD.

■ Business Value:

- ▶ This greatly simplifies the installation process for zSeries customers that have SCSI/FCP attached CD or DVD drives, avoiding the need of a network installation server.

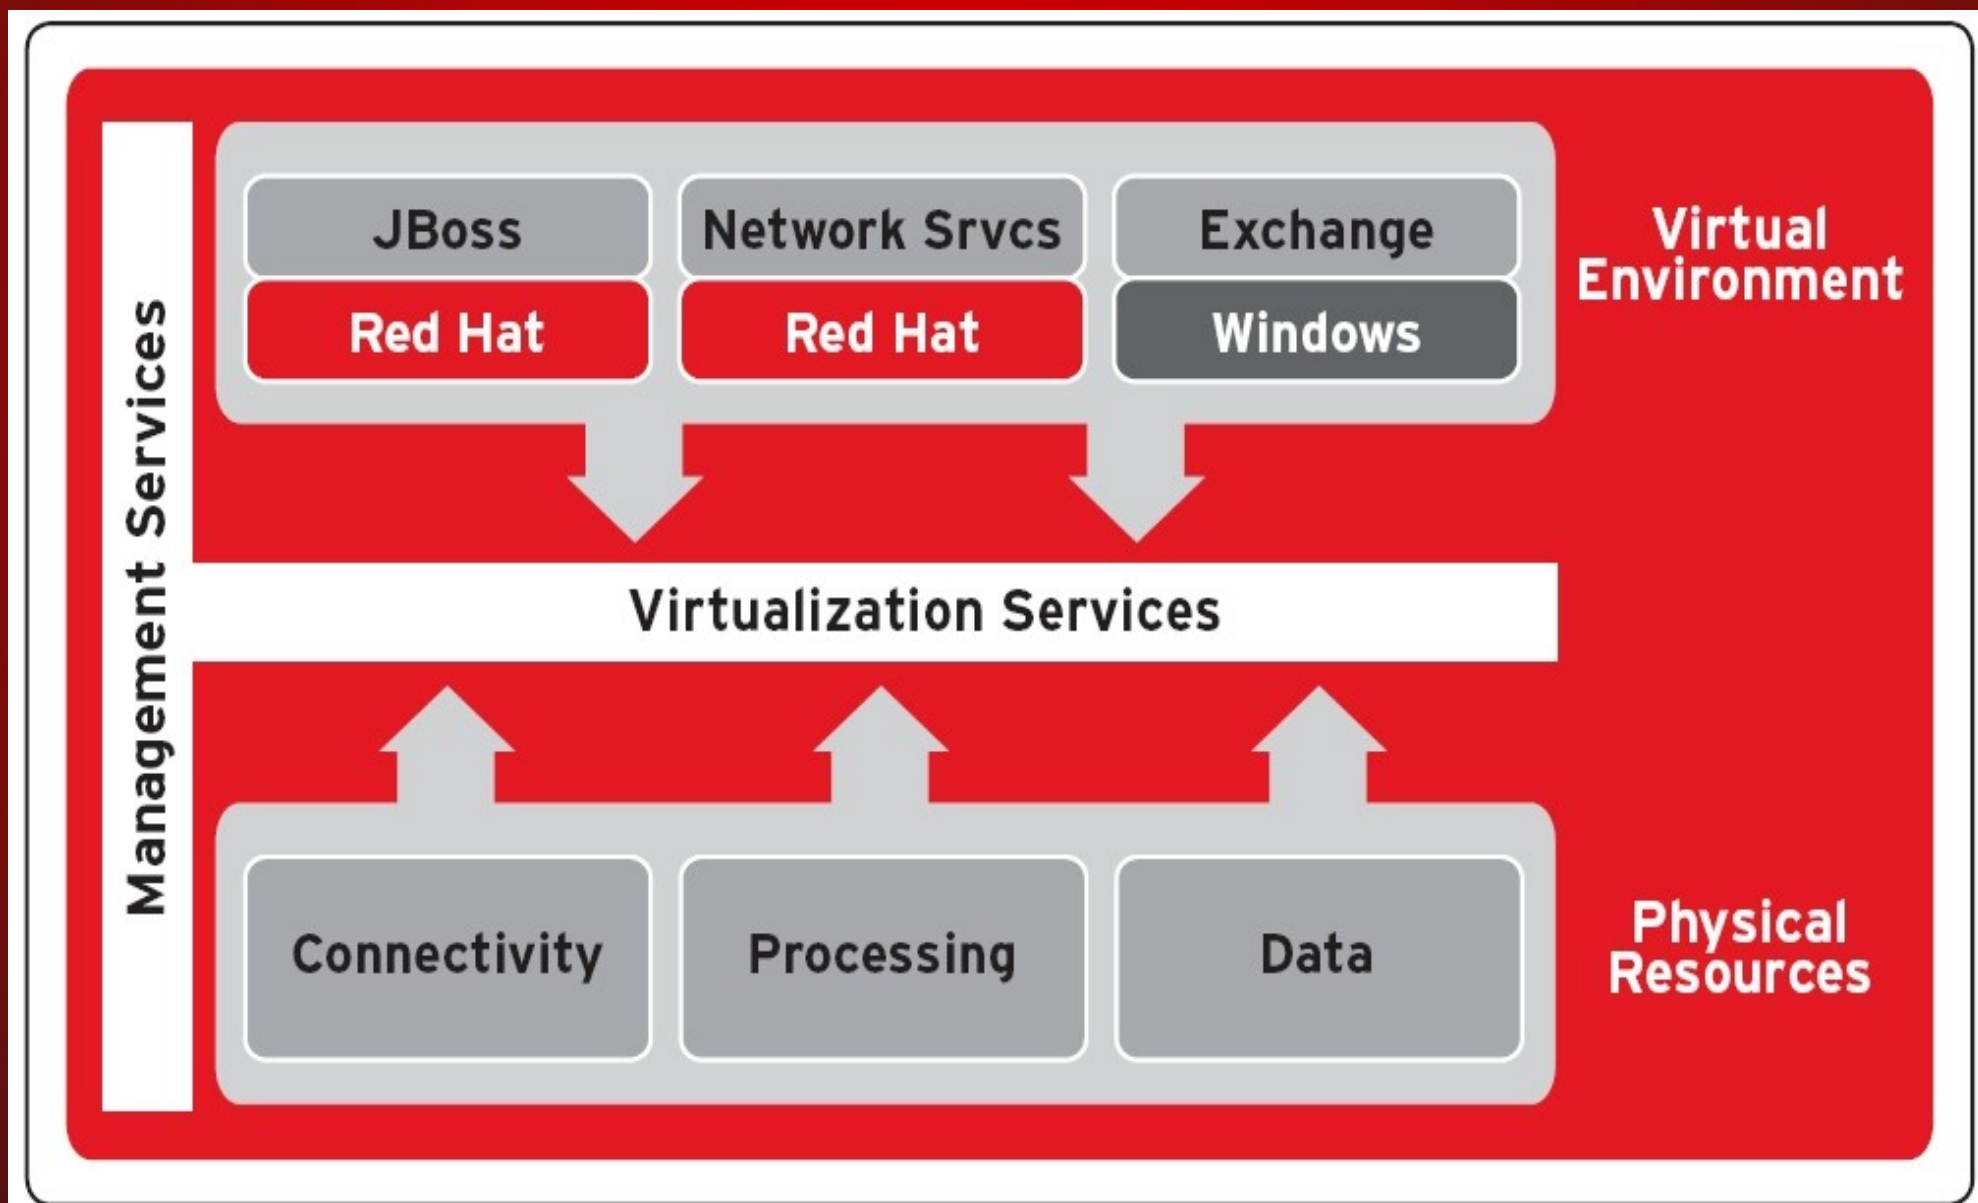
■ Contents:

- ▶ Build image bootable.
- ▶ Patch in anaconda (installer) to be able to install from CD/DVD.
- ▶ Update documentation

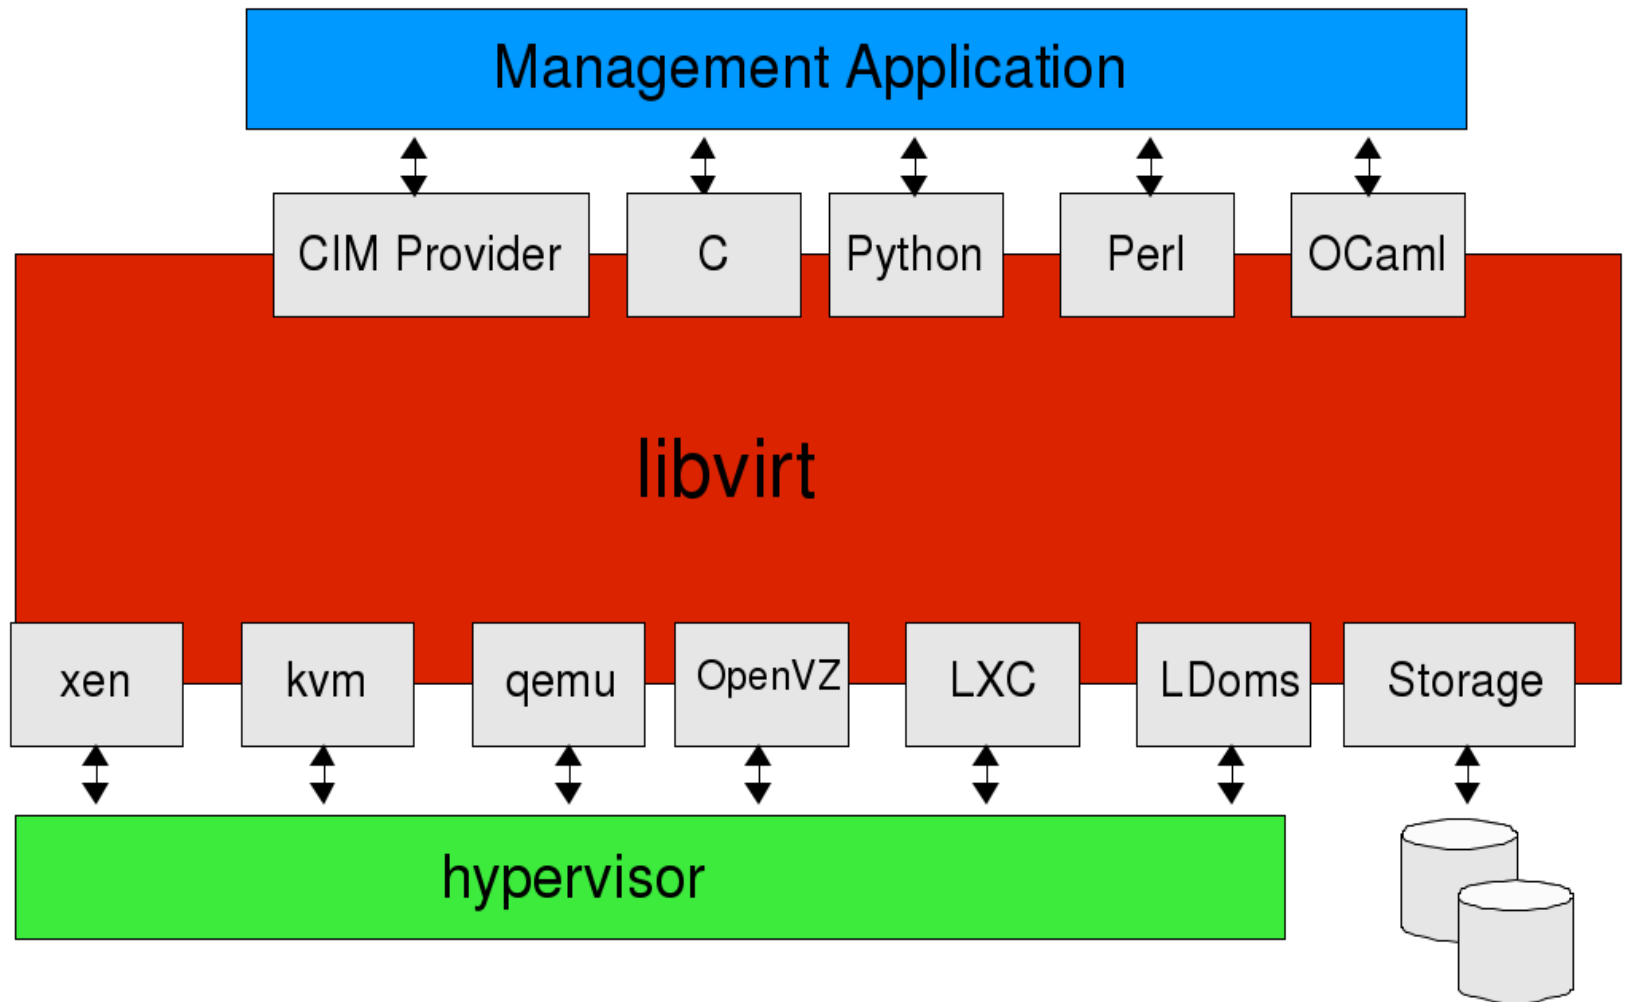

■ Work responsibilities:

- ▶ Red Hat to implement changes in build and anaconda (installer) and documentation for RHEL 6.
- ▶ Red Hat has the HW to do functional test when implemented.
- ▶ IBM will test it again when it is delivered in RHEL 6 test phase.

RHEL 6.0: Works In Progress



RHEL 6.0: Works In Progress



Open Discussion / Q&A