

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

	Oracle/IBM Portal, Oracle BPM, ILOG JRules	JBoss Portal Platform, JBoss jBPM, Rules Frameworks	
	BEA WebLogic, IBM Websphere	JBoss Enterprise Application Platform	
	BEA AquaLogic, IBM ESB, IBM EII	JBoss Enterprise SOA Platform, MetaMatrix Enterprise Data Services Platform	JBoss Operations Network
eDirectory, SunDS	IBM Websphere MQ, Tibco EMS	Red Hat Enterprise MRG Messaging	Red Hat Directory Server, Virtual Directory, Certificate Services
Keon, PowerBroker, Active Directory	Data Synapse, Platform	Red Hat Enterprise MRG Grid	Red Hat Enterprise IPA
HP OpenView/Opware, IBM Tivoli	EMC PowerPath	Multi-path I/O	Red Hat Network
	Veritas Storage Suite	LVM, CLVM, Global File System	
	VMware Virtualization	Red Hat Integrated Virtualization	
Citrix/VMware VDI	AIX, HP-UX, Solaris	Red Hat Enterprise Linux, Red Hat Enterprise MRG Realtime	Qumranet Solid ICE
	Veritas Cluster Suite	Red Hat Cluster Suite	

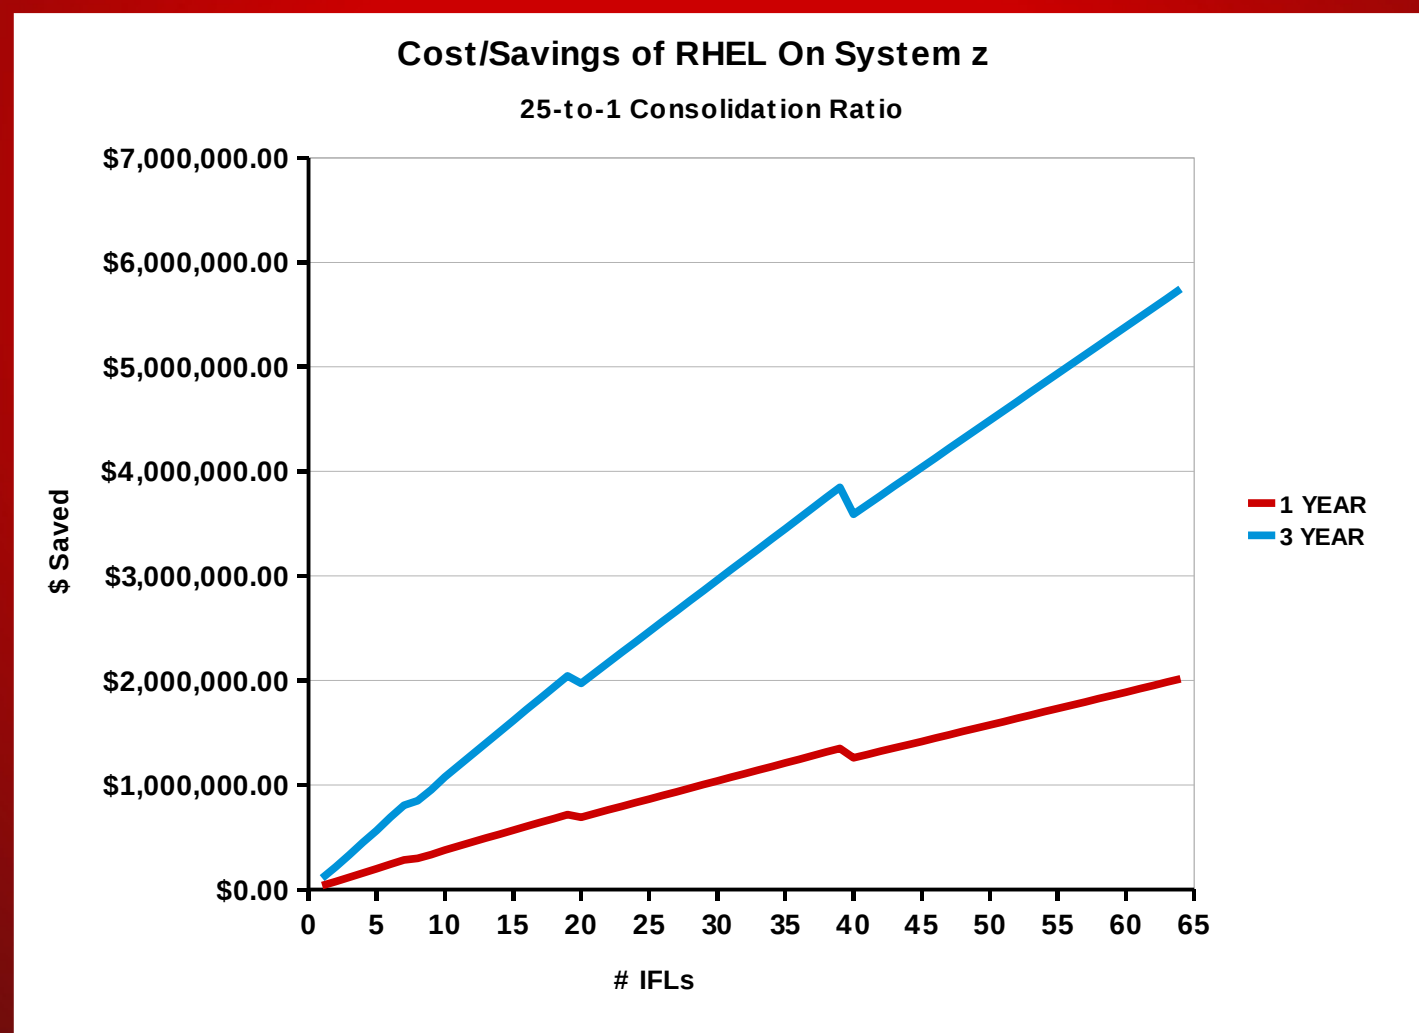
Red Hat / IBM Relationship

- 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)



Why move to System z?

RHEL Subscription Cost Elimination/Prevention





Why move to System z?

3rd Party ISV Costs

- **Oracle DB**

- MSRP * #cores * CPU_Factor

- SUN: $\$40,000 * 8 * 1.7 = \$544K$

- z9 = $\$40,000 * 1 * 1 = \$40K$

- z10 = $\$40,000 * 4 * 1 = \$160K$



Why move to System z?

Environmental “Go Green” Factors

- **Bank of New Zealand**
 - One of the top 50 largest banks in the world
 - Offices in 4 continents, 15 countries
 - Mainframe RHEL since September, 2008
 - SWIFT (\$10B/day)
 - PCBB (\$4M/day)
 - Teller Banking Applications
 - Carbon neutral by 2010



Why move to System z?

Environmental “Go Green” Factors

- **Bank of New Zealand (cont)**

- Consolidated 131 SUN servers to RHEL on z10
- Mix of small, medium, large: 280Rs, V440s, E10Ks

	SUN	RHEL & z10	
Power (kW/hr)	36	22	38% less
Heat (kBTUs/hr)	110	74	33% less
Space (Racks)	6.5	4.5	31% less
Carbon Emissions	66	40	39% less



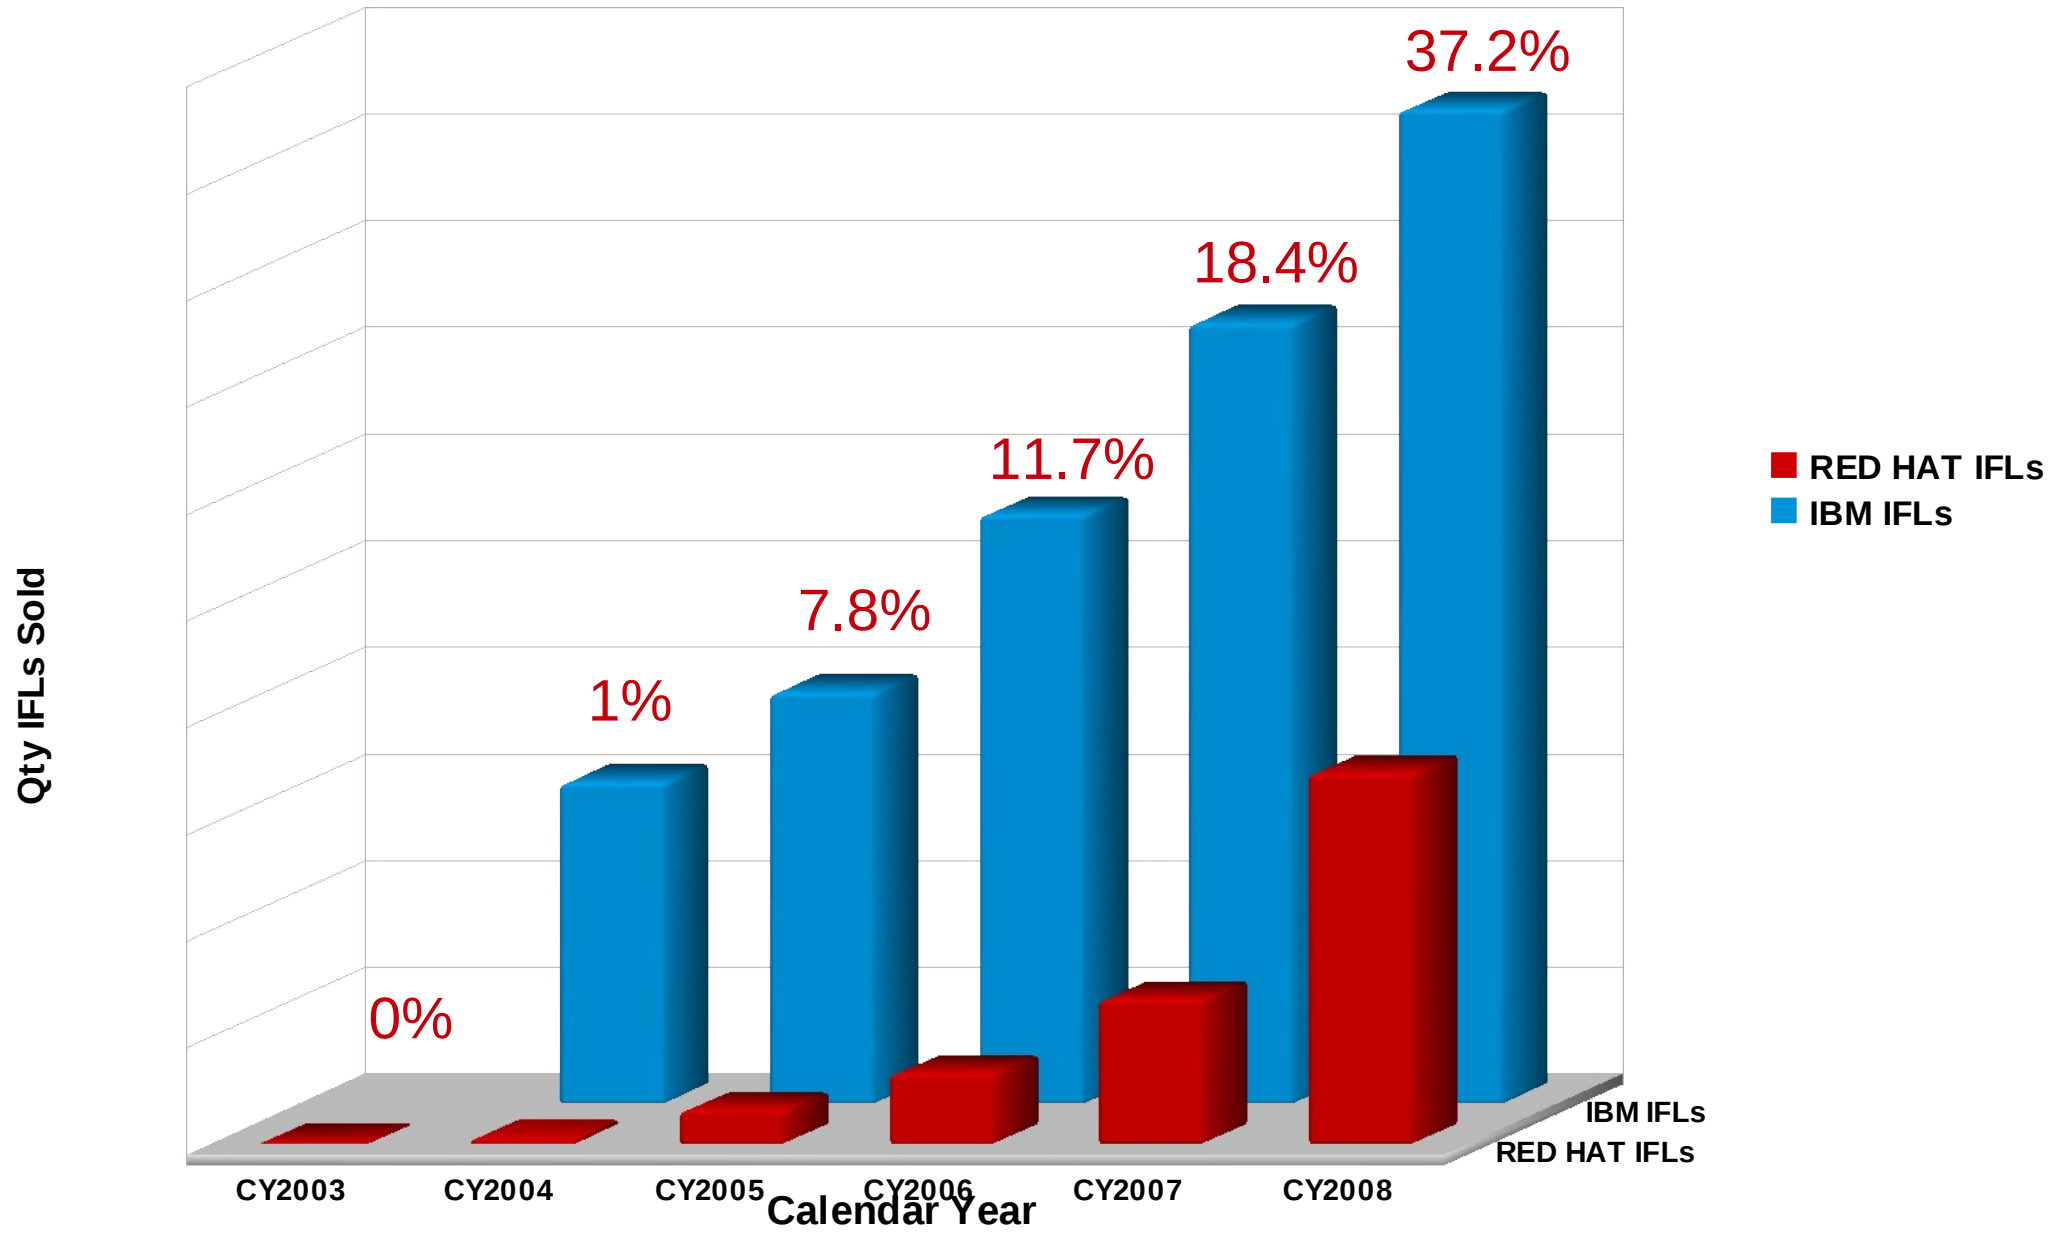
Why move to System z?

Performance

- **104 SUN cores to 7 z10 EC IFLs**
 - Large US Government customer
 - 700M+ rows in Oracle
 - 104 SUN cores, incl prod/dev/test,
 - Processing time 3 days
 - 7 z10 EC IFLs, incl prod/dev/test,
 - Processing time < 15 minutes, peak utilization of 55%

Global RHEL on System z Market Share


Updated 11/2008





Red Hat Development & Subscription Model

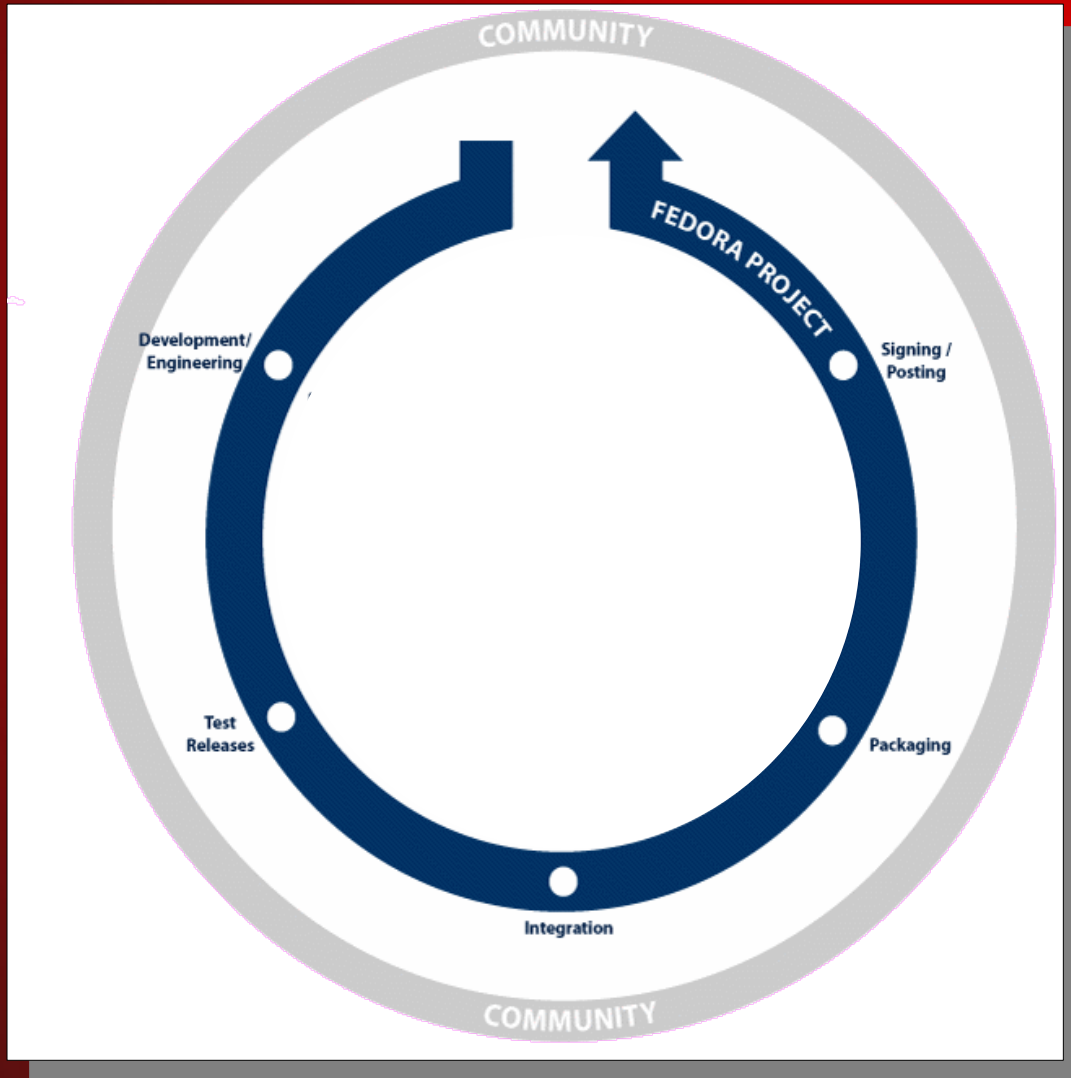
Red Hat Development Model



COMMUNITY

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

Red Hat Development Model

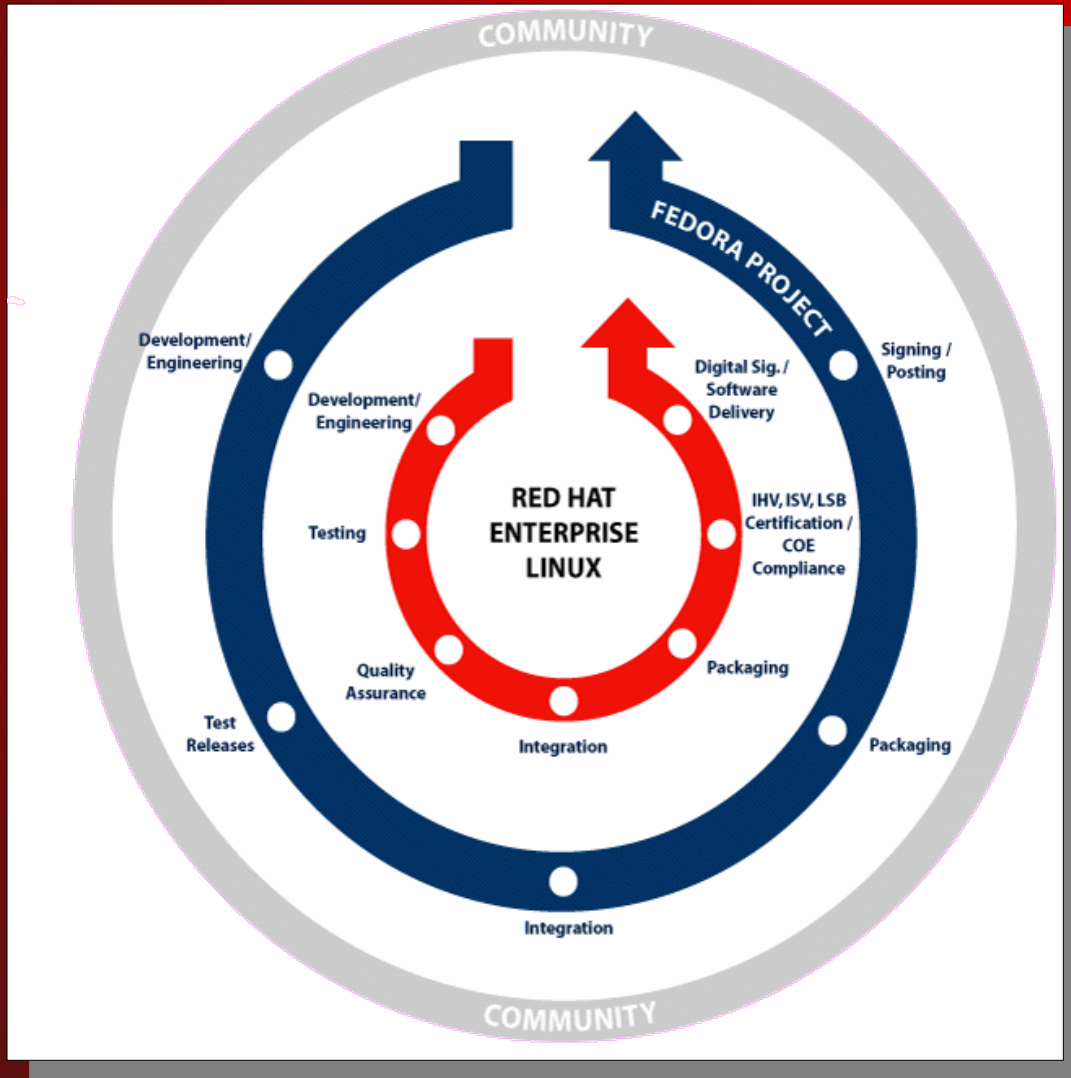


FEDORA

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

Fedora for System z @


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

Red Hat Development Model



- **RHEL 4.7** GA planned for July 21st.
- 4.8 as an extension of hardware enablement planned for early 2009.
- Limited hardware enablement focused on making platforms run RHEL4.
- Phase 2 at least until Q4 / 2009 (depends on RHEL6 schedule, TBD).
- 4.9 as the last minor release, then transition into Phase 3.

- **RHEL 3** in Maintenance phase (no new hardware, critical changes only).
- New hardware still available through virtualization
- Improved para-virtualized drivers are released.

- **RHEL 5** public release on March 14th, 2007
- Introducing Virtualization and the Advanced Platform.
- 5.2 released May 21, 2008
- 5.3 January 2009, schedule extended to better sync with hardware schedules.

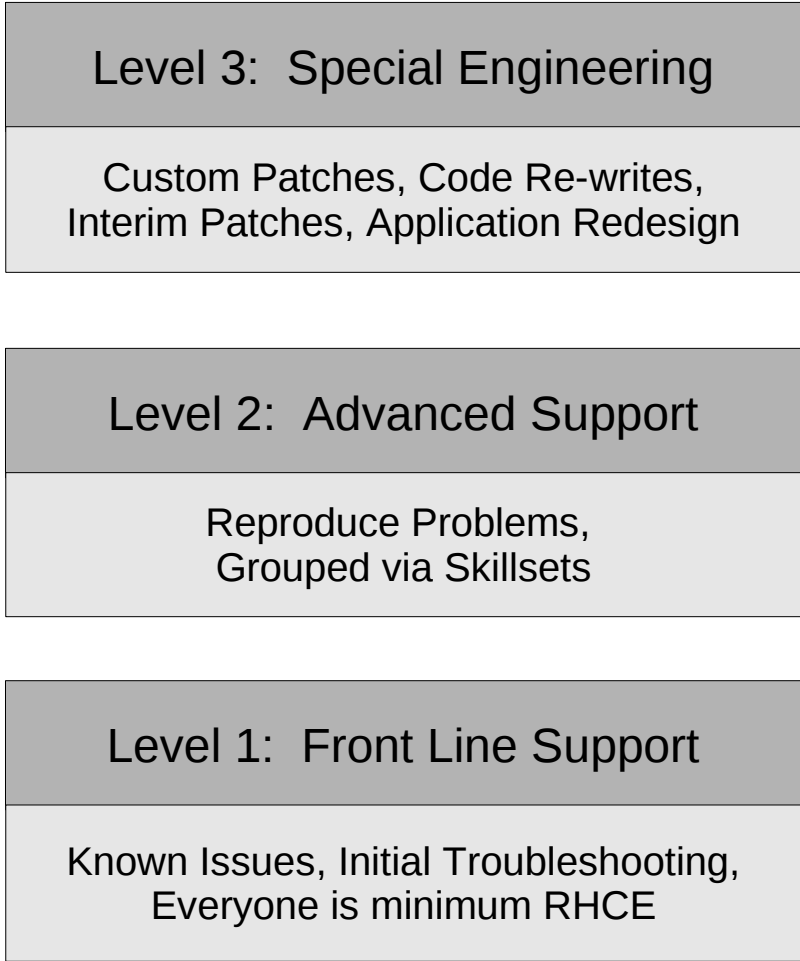
RHEL .next

- No fixed schedule for next major release yet.
- Currently in planning and requirements phase.
- Planning is driven by customer and partner requirements.
- Further announcements to come.

Linux on System z Support

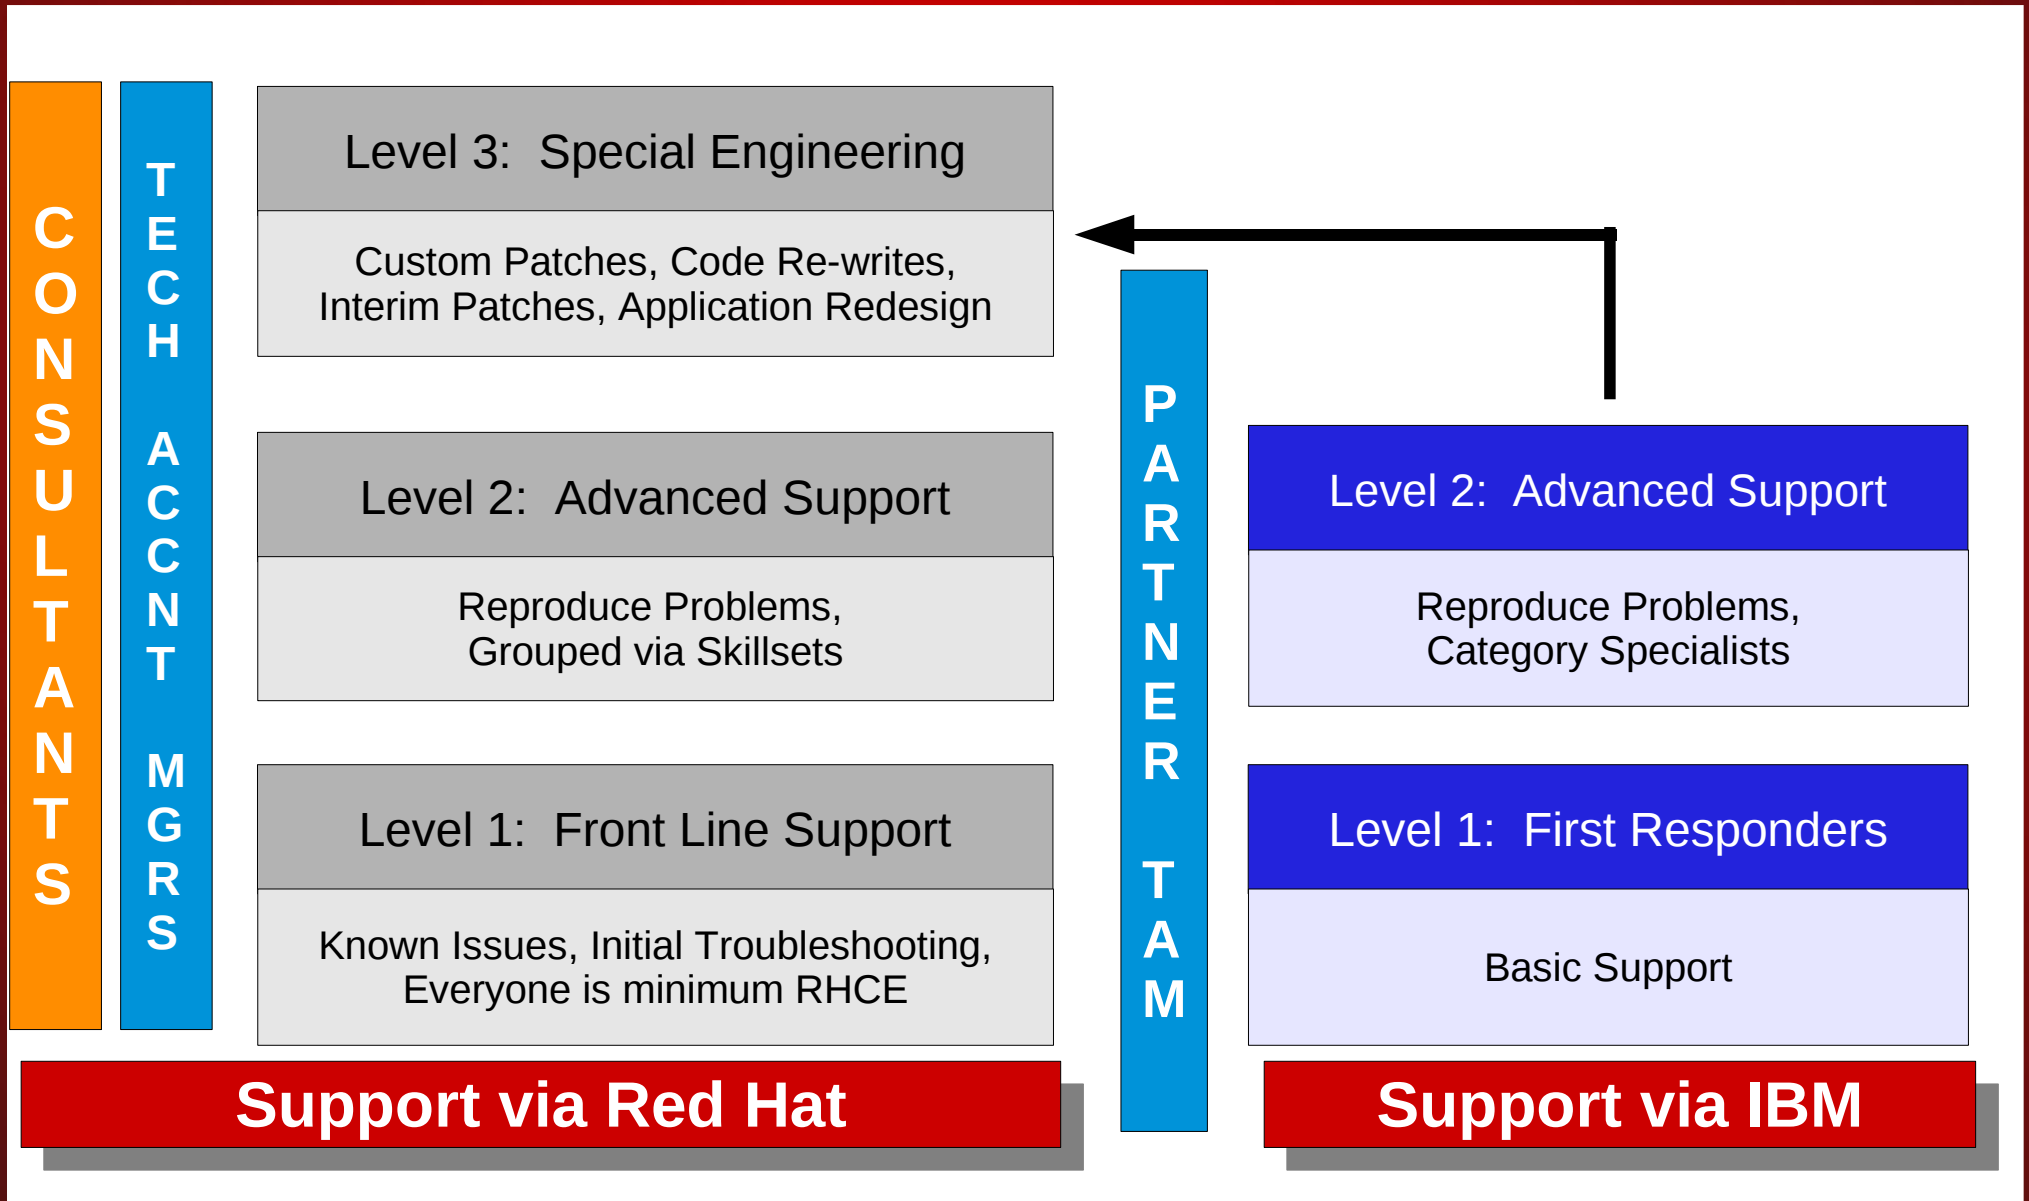
CONSULTANTS

TECH ACCOUNT MGRS



Support via Red Hat

Linux on System z Support



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

RHEL 5.2: Technical Review

- **Audit Subsystem**
 - Support for process-context based filtering
 - More filter rule comparators
- **Address Space Randomization**
 - Address randomization of multiple entities – including stack & mmap() region (used by shared libraries) (2.6.12; more complete implementation than in RHEL4)
 - Greatly complicates and slows down hacker attacks

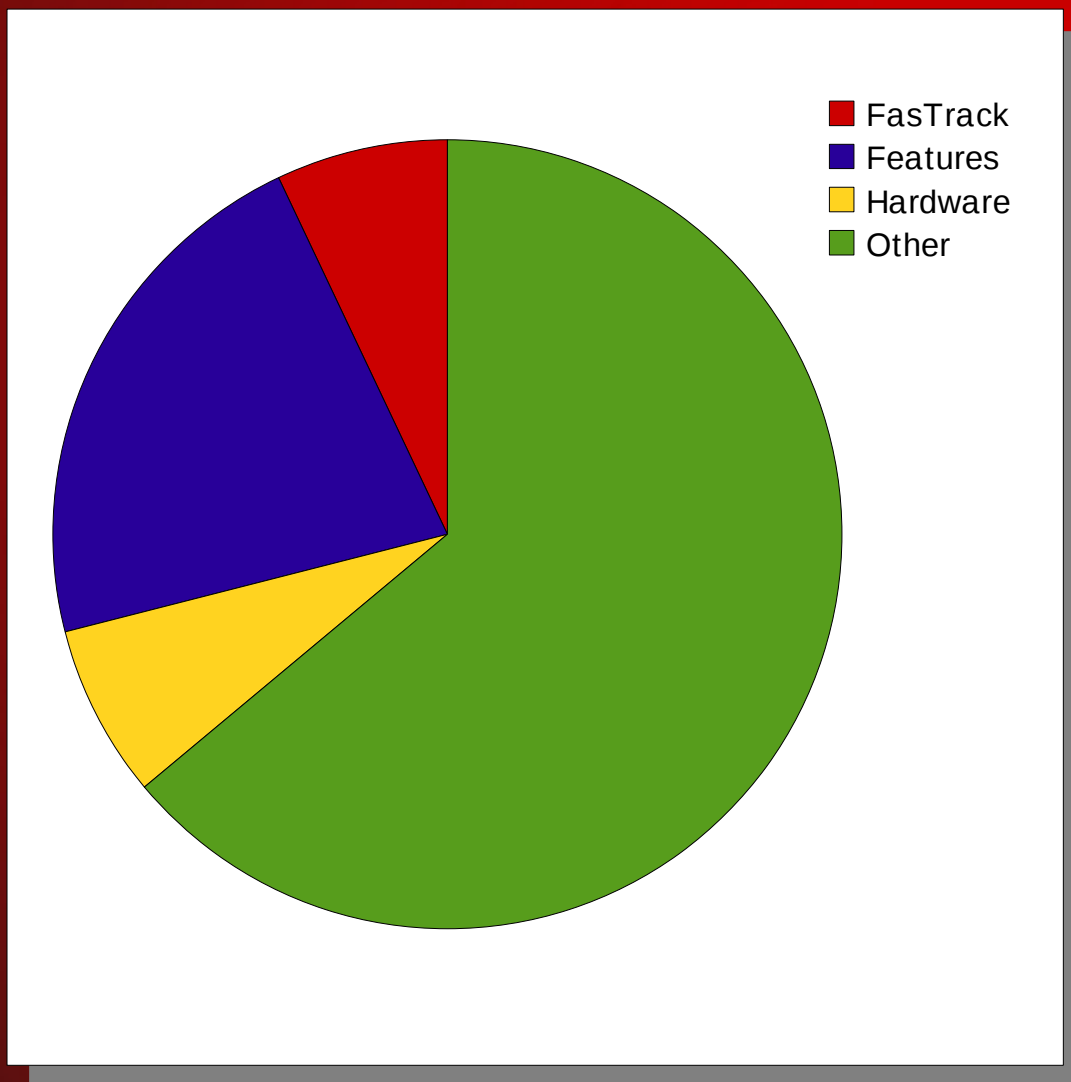
RHEL 5.2: Technical Review

- **High Resolution Timers**

- Provide fine resolution and accuracy depending on system configuration and capabilities - used for precise in-kernel timing

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: Technical Review

xDR System Initialization for LPAR Clients

(Red Hat BugZilla [184770](#), IBM BugZilla [37874](#))

- This requirement enables a new version of the "GDPS/PPRC Multiplatform Resiliency" disaster recovery solution.
- This new version will support site failover and Hyperswap (transparent storage subsystem failover) to Linux running in a zSeries LPAR
- (in a next step) non-zSeries Linux images attached to an ESS

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

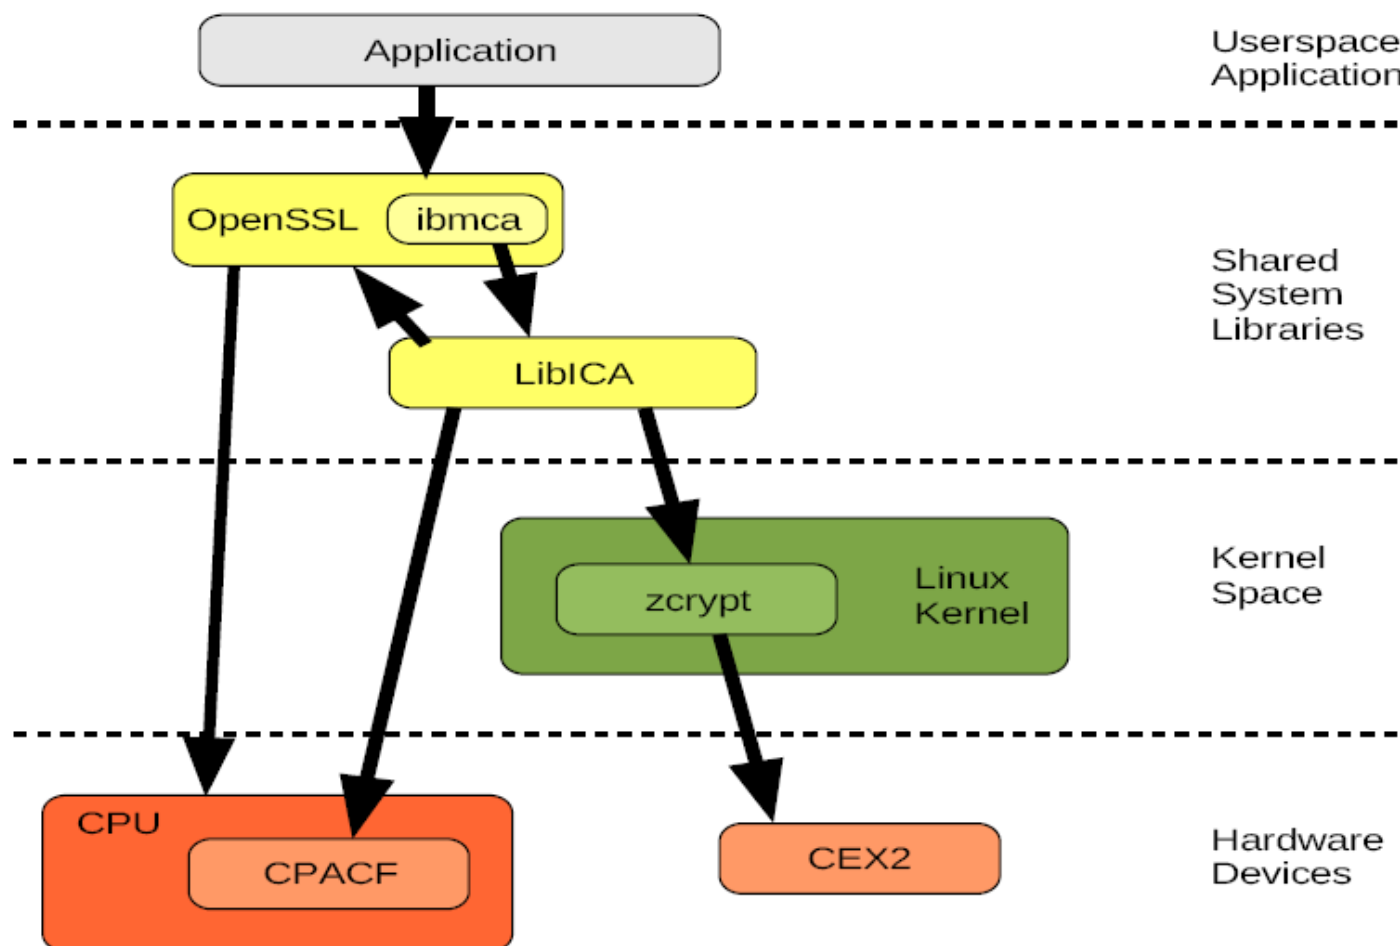
Long Numbers Generation

(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 SWAP 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](#))


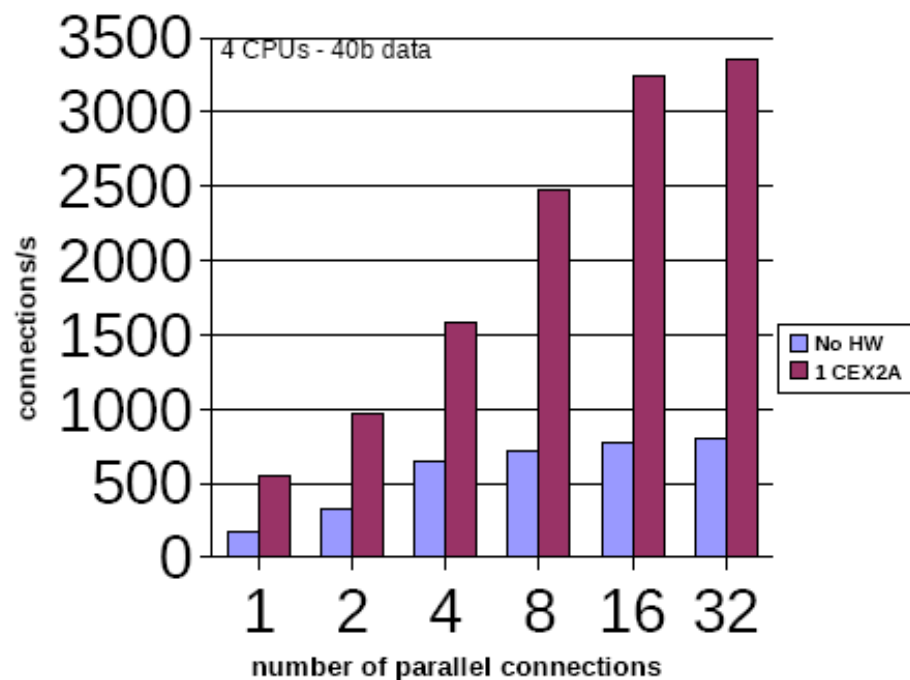


RHEL 5.3: Technical Review

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

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.

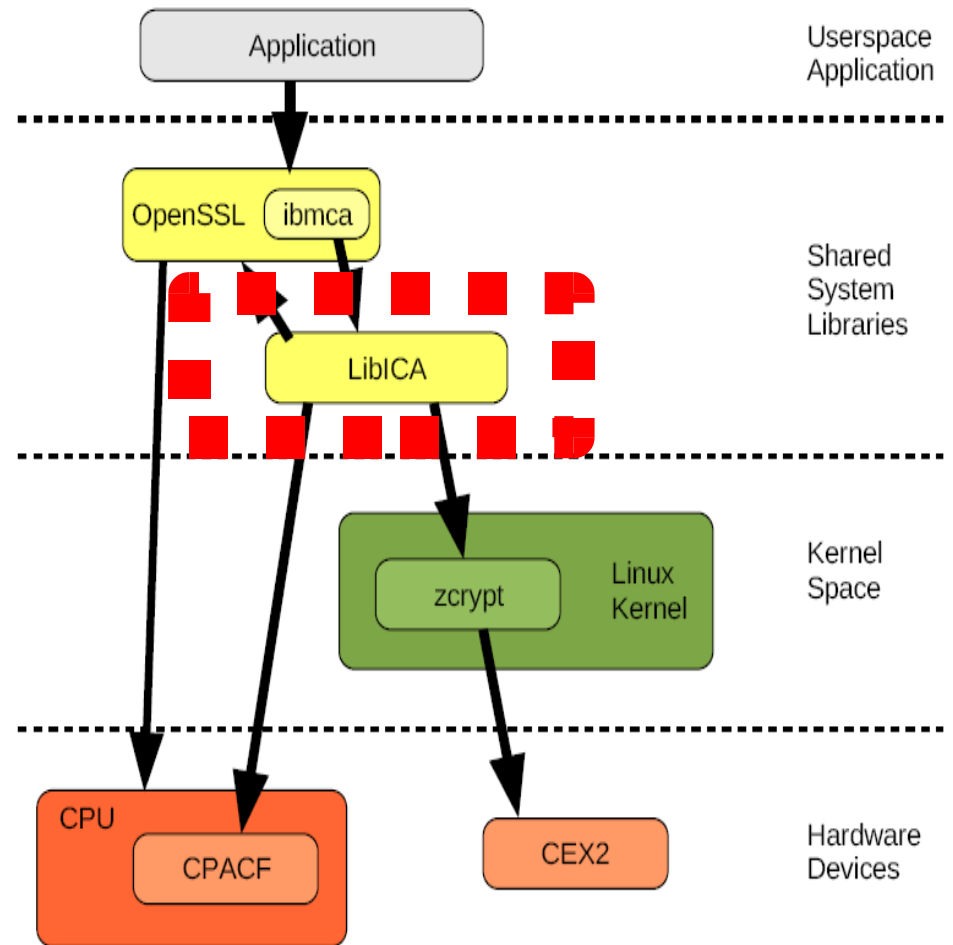
English Version: You dedicate 2 z10 IFLs to a RHEL5 VM. We can then pin applications to specific cores, or to IFLs in their entirety.

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.

An underwater scene with a blue background, coral reefs, and several fish swimming. The text is overlaid on this scene.

RHEL 5.4 Tech Deep Dive (Planned Features)

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

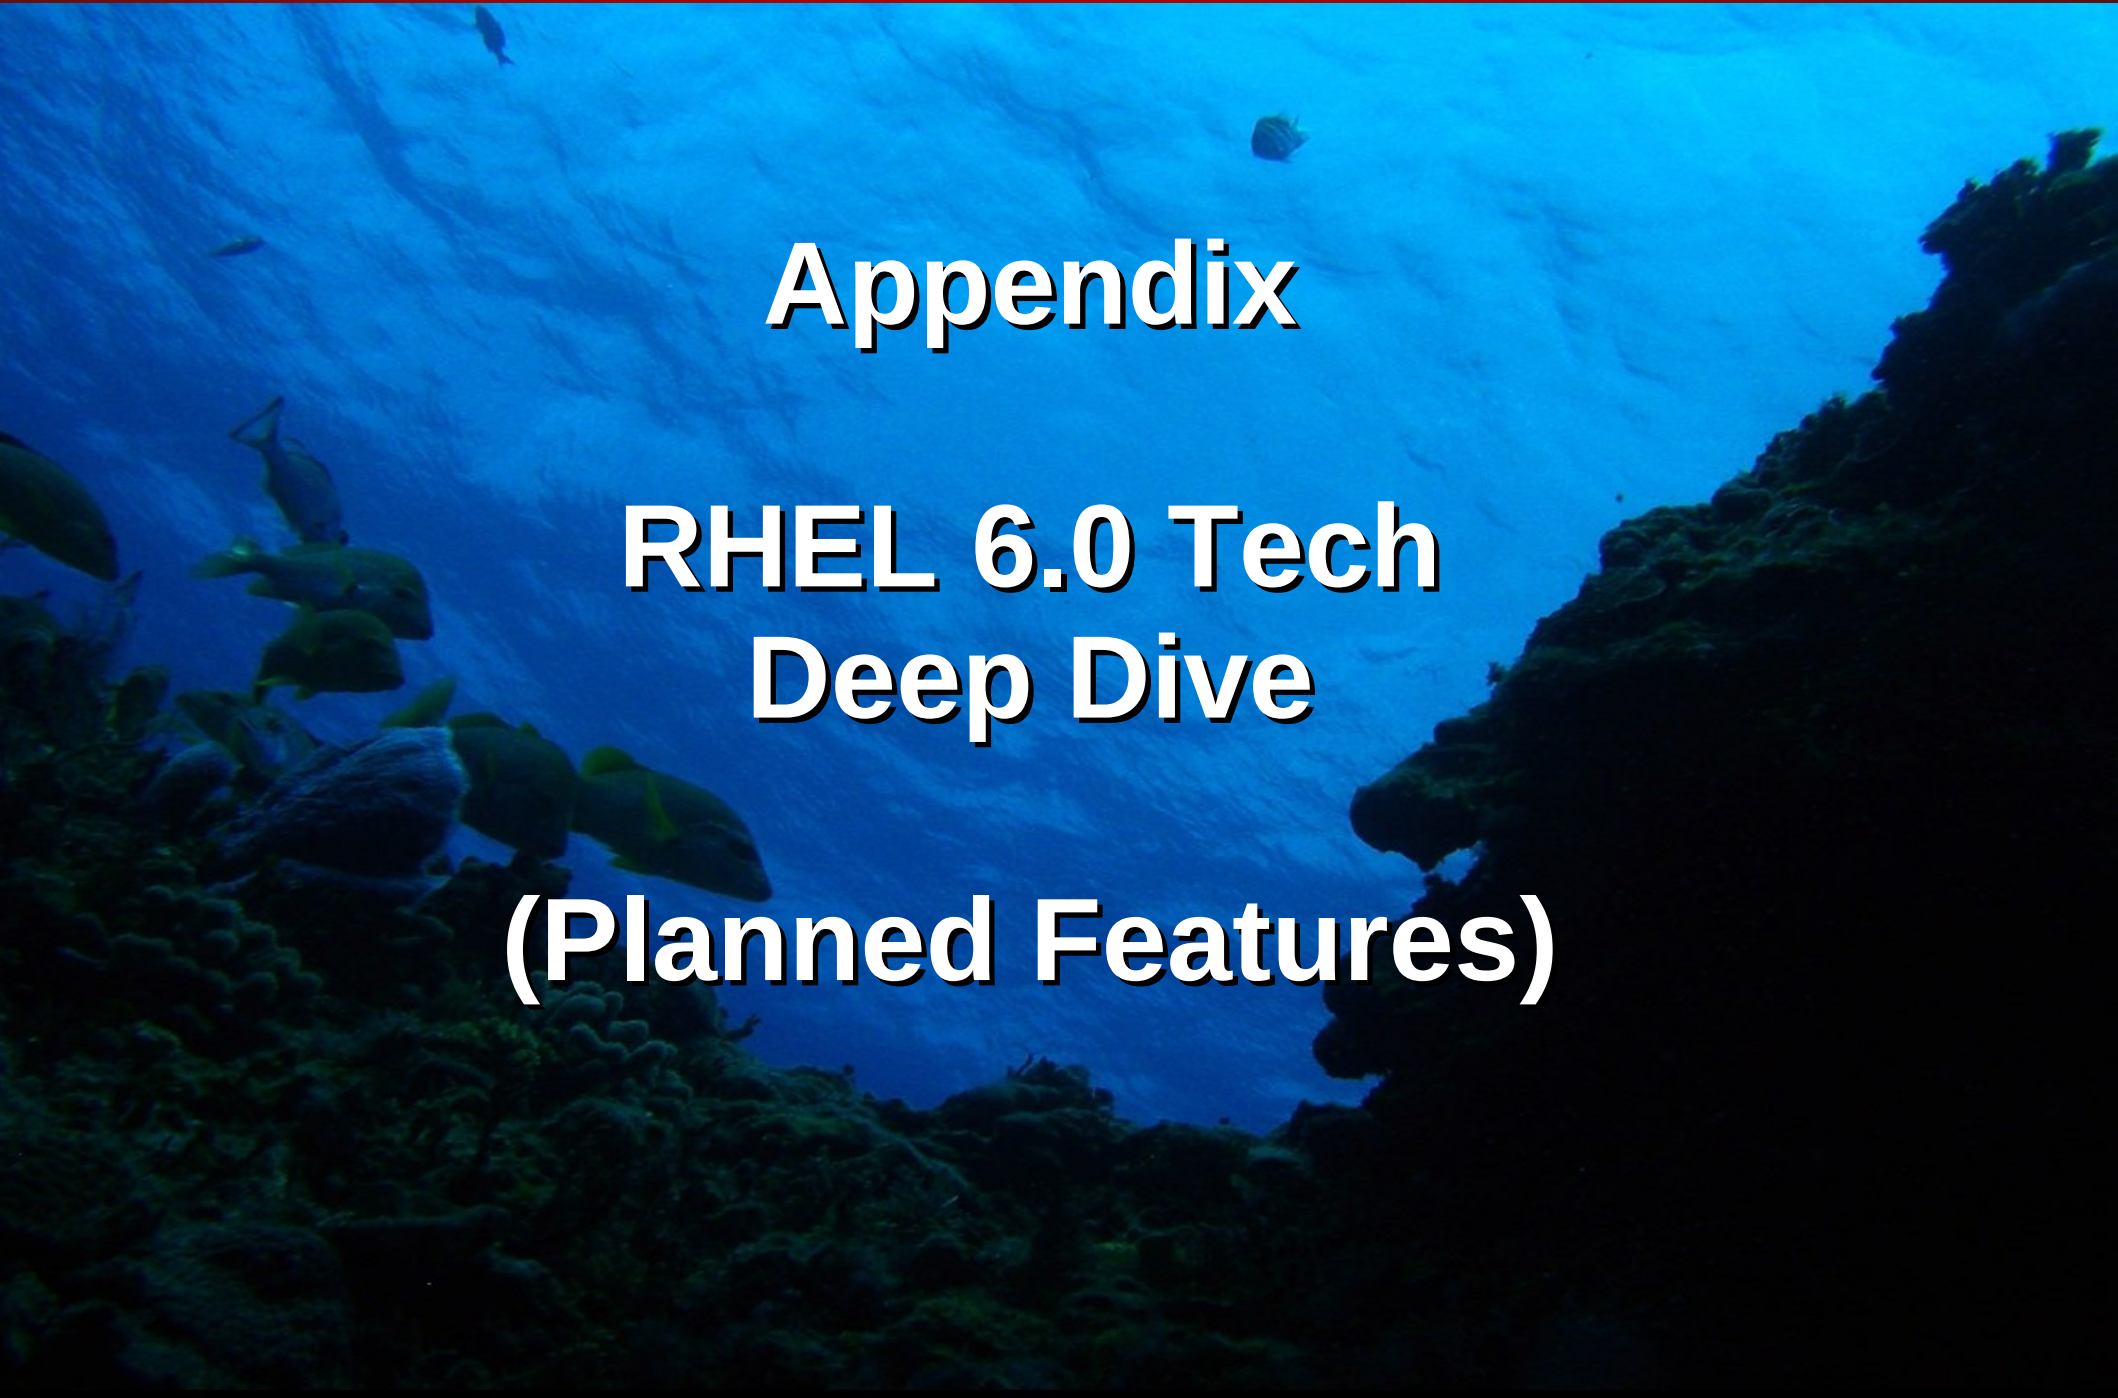
BugZilla	Feature Description
475345	[LTC 5.4 FEAT] Change list of Anaconda network alternatives to indicate supported devices on System z [201679]
475346	[LTC 5.4 FEAT] Improve checking mechanisms and workflow of Linux on System z Anaconda install process [201676]
475350	[LTC 5.4 FEAT] Dialog defaults for Linux on System z specific Anaconda [201677]
475358	[LTC 5.4 FEAT] Adjust Anaconda Swap recommendations to Linux on System z specifics [201680]
475520	[LTC 5.4 FEAT] Intuitive dump device configuration workflow and dialogue [201624]
475675	[LTC 5.4 FEAT] cio_ignore entry in generic.prm for LPARs [201085]
475677	[LTC 5.4 FEAT] Firstboot for System z [201092]
461288	[EMC 5.4 feat] Require kernel support to issue Control I/O to CKD dasd on EMC Symmetrix arrays
472936 [SEC]	extension of linuxrc.s390: improved workflow, dialog defaults, indicate supported network devices
474679	[LTC 5.4 FEAT] Dynamic CPU hotplug daemon for System z [201132]
474942	[LTC 5.4 FEAT] Add vmconvert option to vmur tool [201758]
475333	[LTC 5.4 FEAT] FCP - Performance Data collection & analysis (userspace) [201591]
475552	[LTC 5.4 FEAT] FCP - Performance data reports [201730]
475557 [SEC]	[LTC 5.4 FEAT] DS8000 Disk Encryption [201740]

RHEL 5.4: Works In Progress

BugZilla	Feature Description
475558	[LTC 5.4 FEAT] TTY terminal server over IUCV (userspace) [201735]
475564	[LTC 5.4 FEAT] Shutdown actions interface (userspace) [201748]
475569	[LTC 5.4 FEAT] Shutdown actions tools [201755]
475571	[LTC 5.4 FEAT] Large image dump on DASD [201752]
475670	[LTC 5.4 FEAT] Program directed IPL support - no XML in system dumper [200782]
477189	[LTC 5.4 FEAT] Pick up latest version of s390-tools
474646	[LTC 5.4 FEAT] Kernel NSS support - kernel part [200790]
474664	[LTC 5.4 FEAT] System z support for processor degradation [200975]
474688	[LTC 5.4 FEAT] Automatic IPL after dump (kernel) [201169]
475530	[LTC 5.4 FEAT] Extra kernel parameter via VMPARM [201726]
475551	[LTC 5.4 FEAT] TTY terminal server over IUCV (kernel) [201734]
475563	[LTC 5.4 FEAT] Shutdown actions interface (kernel) [201747]
475570	[LTC 5.4 FEAT] Provide service levels of HW & Hypervisor in Linux [201753]
475572	[LTC 5.4 FEAT] HiperSockets Layer3 support for IPv6 [201751]

RHEL 5.4: Works In Progress

475820	[LTC 5.4 FEAT] Linux to add Call Home data [201167]
477188	[LTC 5.4 FEAT] ETR support
475343	[LTC 5.4 FEAT] Provide CMS script for initial IPL under z/VM [201594]
475548	[LTC 5.4 FEAT] FCP - Performance data collection (blktrace) [201729]
475669	[LTC 5.4 FEAT] snIPL SCSI LOAD for LPAR [200787]
472764	let mkinitrd default to recreating the initrd for the currently running kernel
474912 [SEC]	[LTC 5.4 FEAT TRACKER] Web 2.0
474917	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package 'mod_security' [201558]
474924	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package memcached [201469]
474925	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package Apache MyFaces Core
474926	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package perl-CGI-Session [201471]
474927	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package mysql-connector-java [201472]
474928	[LTC 5.4 FEAT] Web 2.0 - Inclusion of packages 'rubygems-actionwebservice' and 'rubygems-tzinfo' [201556]
474929	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package 'rubygems-rake' [201554]
474930	[LTC 5.4 FEAT] Web 2.0 - Inclusion of packages 'rubygems-actionpack', 'rubygems-activerecord', 'rubygems-activesupport', 'rubygems-actionmailer' [201555]
474932	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package rubygems [201465]
474933	[LTC 5.4 FEAT] Web 2.0 - Inclusion of package rubygem-rails [201466]
475334	[LTC 5.4 FEAT] FCP - Performance Data collection (kernel) [201590]
468172 [SEC]	FEAT: 201085: cio_ignore entry in generic.prm for LPARs



Appendix

**RHEL 6.0 Tech
Deep Dive**

(Planned Features)

RHEL 6.0: 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: Early 2010**

RHEL 6.0: Works In Progress

462973	[LTC 6.0 FEAT] 201679:Change list of Anaconda network alternatives to indicate supported devices on System z
462974	[LTC 6.0 FEAT] 201677:Dialog defaults for Linux on System z specific Anaconda
462975	[LTC 6.0 FEAT] 201676:Improve checking mechanisms and workflow of Linux on System z Anaconda install process
463177	[LTC 6.0 FEAT] 201686:Installer - HiperSockets MAC Layer Routing Support
463180	[LTC 6.0 FEAT] 201687:Installer - QETH Componentization
463184	[LTC 6.0 FEAT] 201690:Installer - FCP LUN discovery tool
463187	[LTC 6.0 FEAT] 201688:Installer migration - Merge CTCMPC into CTC device driver
463831	[LTC 6.0 FEAT] 201764:Installer enhancement - FICON Hyper PAV enablement
463564	[LTC 6.0 FEAT] 201092:Firstboot for System z
462976	[LTC 6.0 FEAT] 201674:Pick up latest version of s390-tools
462977	[LTC 6.0 FEAT] 201675:Pick up latest version of libica
463208	[LTC 6.0 FEAT] 201730:FCP - Performance data reports
463560	[LTC 6.0 FEAT] 201132:Dynamic CPU hotplug daemon for System z
463688	[LTC 6.0 FEAT] 201591:FCP - Performance Data collection & analysis (userspace)
463707	[LTC 6.0 FEAT] 201735:TTY terminal server over IUCV (userspace)

RHEL 6.0: Works In Progress

463806	[LTC 6.0 FEAT] 201748:Shutdown actions interface (userspace)
463812	[LTC 6.0 FEAT] 201752:Large image dump on DASD
463822	[LTC 6.0 FEAT] 201757:Automatic IPL after dump (userspace)
463823	[LTC 6.0 FEAT] 201758:Add vmconvert option to vmur tool
463826	[LTC 6.0 FEAT] 201754:Extend lstage to support SCSI tapes
463650	[LTC 6.0 FEAT] 201303:Provide a utmp format that is compatible between 32 and 64 bit.
463795	[LTC 6.0 FEAT] 201184:Provide DFP hardware accelerated libgcc
463796	[LTC 6.0 FEAT] 201183:System z optimizations for gcc 2007
463830 [SEC]	[LTC 6.0 FEAT] 201765:Compiler- Architecture Level Set for IBM System z9 and newer
463541	[LTC 6.0 FEAT] 201066:QETH Componentization
463583	[LTC 6.0 FEAT] 201162:CMM2 Merge for Upstream Integration (Full version)
463601	[LTC 6.0 FEAT] 201171:FCP Automatic Port Discovery
463602	[LTC 6.0 FEAT] 201169:Automatic IPL after dump
463678	[LTC 6.0 FEAT] 201546:FCP - code cleanup stage 2
463679	[LTC 6.0 FEAT] 201545:FCP - code cleanup stage 1

RHEL 6.0: Works In Progress

463689	[LTC 6.0 FEAT] 201590:FCP - Performance Data collection (kernel)
463692	[LTC 6.0 FEAT] 201593:Sysplex Timer Protocol Support
463694	[LTC 6.0 FEAT] 201592:Exploitation of DCSSs above 2G
463695	[LTC 6.0 FEAT] 201723:Kernel Message Catalog autogeneration - Stage 1: infrastructure
463696	[LTC 6.0 FEAT] 201728:Secondary unicast addresses for qeth layer2 devices
463697	[LTC 6.0 FEAT] 201725:Pre-allocated headers for HiperSockets (qeth driver)
463698	[LTC 6.0 FEAT] 201727:Kernel Message Catalog autogeneration - Stage 3: DASD, tape, QETH and CIO
463699	[LTC 6.0 FEAT] 201726:Extra kernel parameter via VMPARM
463700	[LTC 6.0 FEAT] 201724:Kernel Message Catalog autogeneration - Stage 2: all s390 drivers and s390 arch. code except for DASD, tape, CIO and QETH
463706	[LTC 6.0 FEAT] 201736:Support for HiperSockets Sniffer
463708	[LTC 6.0 FEAT] 201734:TTY terminal server over IUCV (kernel)
463710	[LTC 6.0 FEAT] 201743:FCP - SCSI error recovery hardening
463799	[LTC 6.0 FEAT] 201747:Shutdown actions interface (kernel)
463804	[LTC 6.0 FEAT] 201750:HiperSockets enhanced SIGA
463805	[LTC 6.0 FEAT] 201749:I/O dynamic configuration support
463811	[LTC 6.0 FEAT] 201753:Provide service levels of HW & Hypervisor in Linux
463825	[LTC 6.0 FEAT] 201756:Linux support for dynamic memory attach

Open Discussion / Q&A