




Red Hat on IBM System z

Shawn Wells <sdw@redhat.com>
Global System z Platform Manager

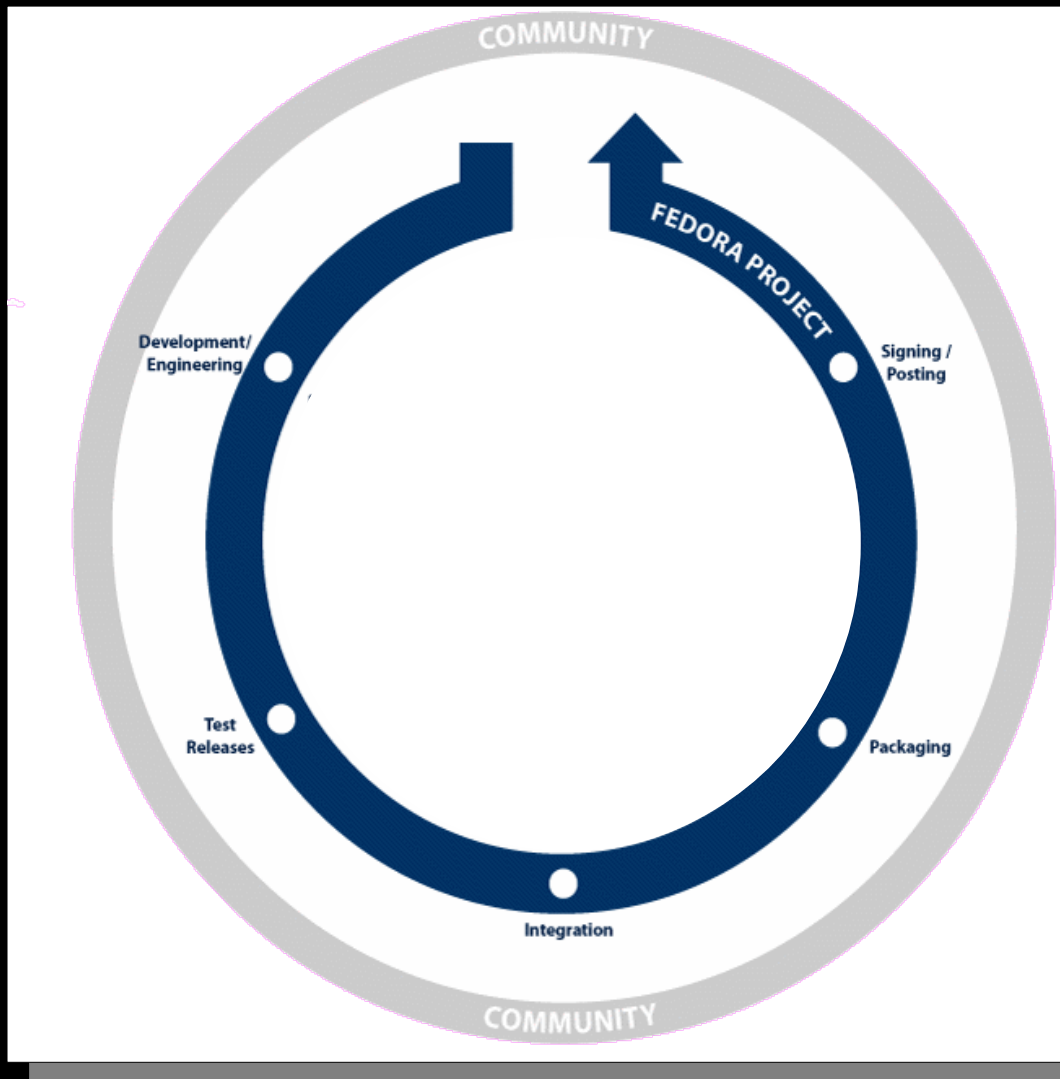
Red Hat Development Model



COMMUNITY

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

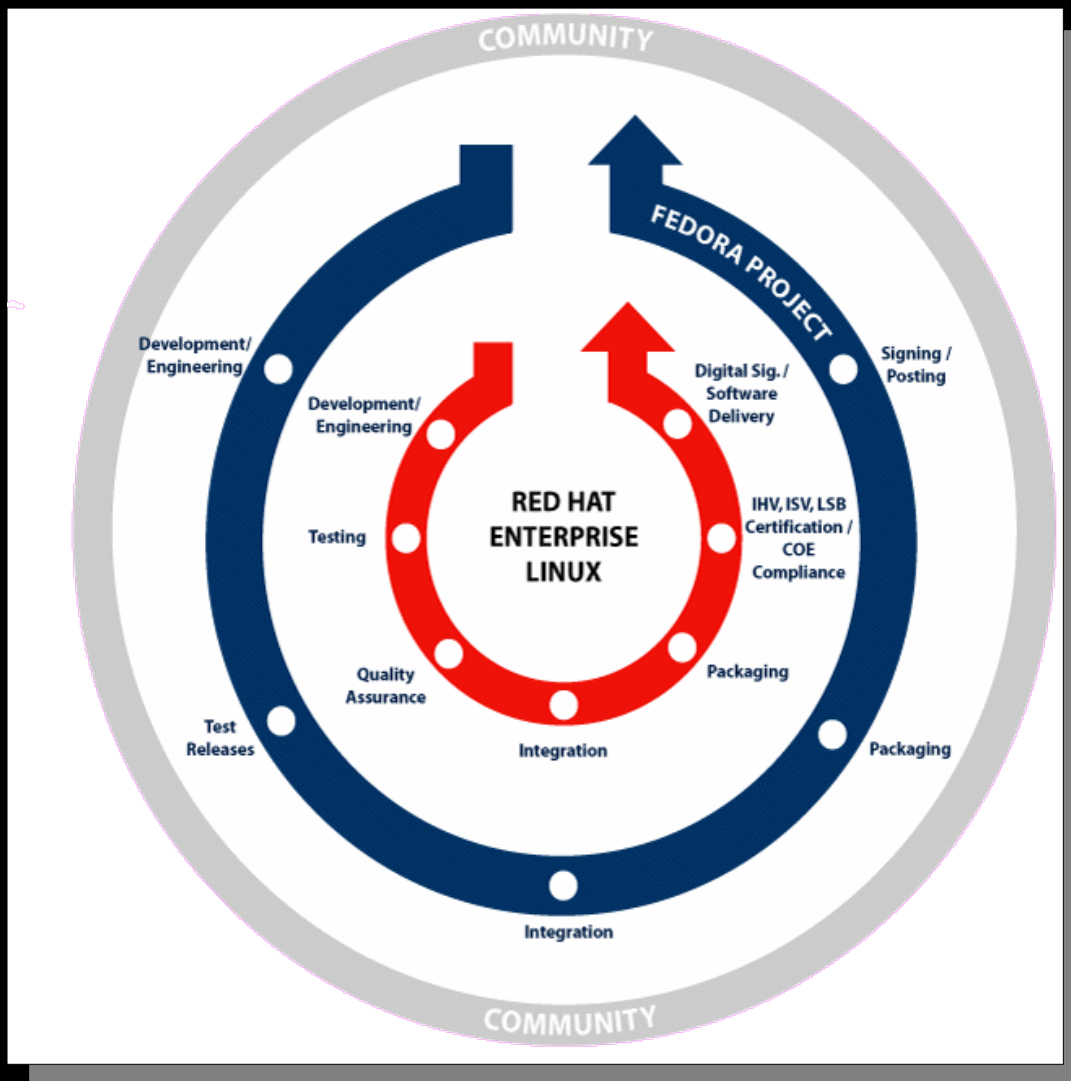
Red Hat Development Model



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

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 5.3 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!
 - apparently Sine Nomine ported this to s390x. Investigating upstreaming & inclusion into RHEL
- 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 Icainfo
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

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



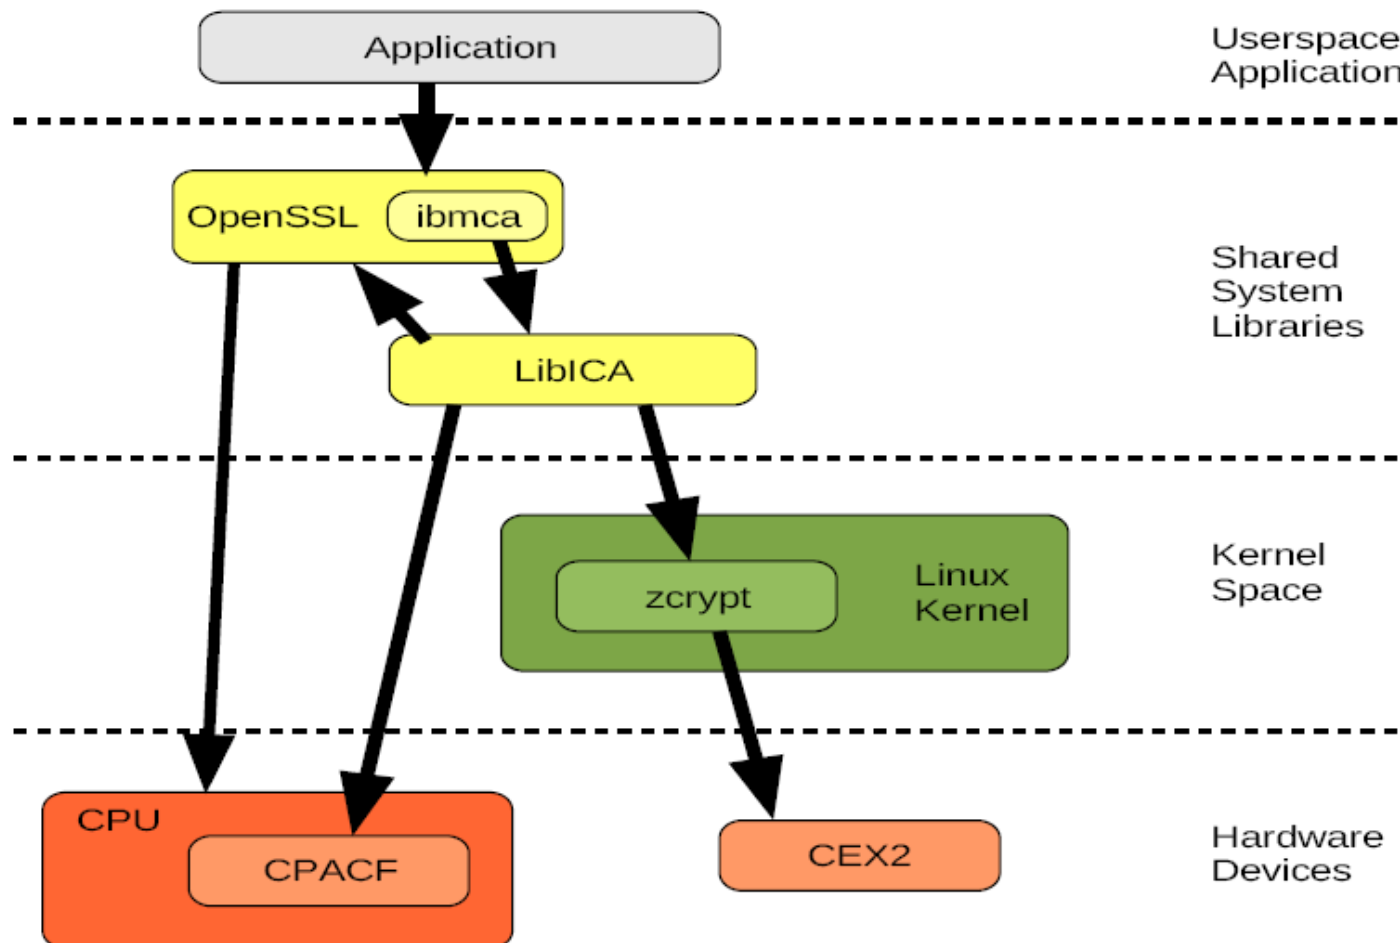
RHEL 5.3: Technical Review

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 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

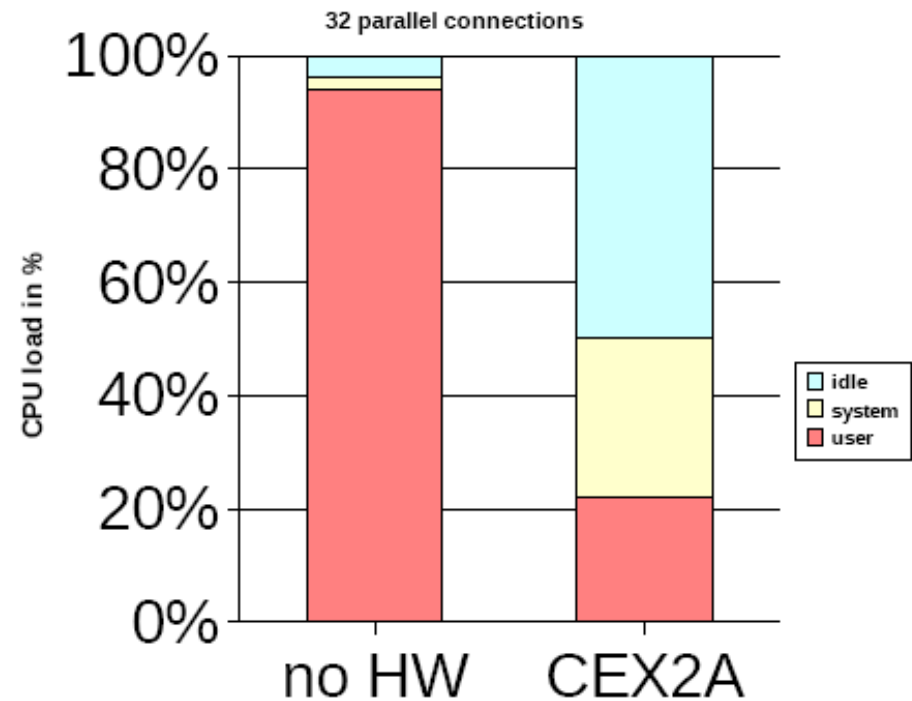

Long Numbers Generation
(Red Hat BugZilla [439440](#), IBM BugZilla [43340](#))



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

FCP – Enhanced Trace Facility
(Red Hat BugZilla 439482, IBM BugZilla [43384](#))

- Detailed troubleshooting information for LUNs

```
[root@h0530014 s390dbf]# ll zfcpl_0.0.5914_rec  
total 0
```

```
--w----- 1 root root 0 Sep 12 08:11 flush  
-r----- 1 root root 0 Sep 12 08:11 hex_ascii  
-rw----- 1 root root 0 Sep 12 08:11 level  
-rw----- 1 root root 0 Sep 12 08:11 pages  
-r----- 1 root root 0 Sep 12 08:11 structured
```

```
[root@h0530014 zfcpl_0.0.5914_rec]# cat structured
```

```
timestamp      01221199894:240391062  
cpu            01  
tag            trigger  
hint           online  
id             85  
reference      0x0000000000000000  
erp_action     0x0000000019a0d9d8  
requested      4  
executed       4  
wwpn           0x0000000000000000  
fcpl_lun       0x0000000000000000  
adapter_status 0x41000124  
port_status    0x00000000  
unit_status    0x00000000
```

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 5.4: 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

[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 `sclp_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 `sclp_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 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

RHEL 6.0 Tech

(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 VMPPARM
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

RHEL 6.0: Works In Progress

463604	[LTC 6.0 FEAT] 201179:Decimal Floating Point support in gcc backend - PFPO
463605	[LTC 6.0 FEAT] 201178:Decimal floating point support in gcc backend (HW support)
463606	[LTC 6.0 FEAT] 201177:Decimal floating point support in gcc backend (SW support)
463455	[LTC 6.0 FEAT] 200785:Improved handling of dynamic subchannel mapping
463516	[LTC 6.0 FEAT] 200855:Linux on System z: Enhanced Linux System Layout
463518	[LTC 6.0 FEAT] 200789:ETR support
463532	[LTC 6.0 FEAT] 201000:FICON: Hyper PAV enablement
463535	[LTC 6.0 FEAT] 200975:System z support for processor degradation
463537	[LTC 6.0 FEAT] 200994:Linux CPU Node Affinity
463551	[LTC 6.0 FEAT] 201064:Standby cpu activation/deactivation.
463552	[LTC 6.0 FEAT] 201065:Modularization of qdio and thin interrupt layer
463553	[LTC 6.0 FEAT] 201067:In-kernel crypto generic algorithm fallback
463559	[LTC 6.0 FEAT] 201140:Standby memory add via SCLP
463584	[LTC 6.0 FEAT] 201159:System z ZFCP Performance Statistics
463603	[LTC 6.0 FEAT] 201168:Linux Support for Dynamic Memory Detach
463607	[LTC 6.0 FEAT] 201176:Merge CTCMPC into CTC device driver for upstream integration
463677	[LTC 6.0 FEAT] 201579:Linux struct page elimination

RHEL 6.0: Works In Progress

463681	[LTC 6.0 FEAT] 201542:FCP - Enhanced Trace Facility
463813	[LTC 6.0 FEAT] 201751:HiperSockets Layer3 support for IPv6
463608	[LTC 6.0 FEAT] 201175:Support for Eclipse IDE Plattform
463704	[LTC 6.0 FEAT] 201739:CMM2 Lite
463178	[LTC 6.0 FEAT] 201680:Adjust Anaconda Swap recommendations to Linux on System z specifics
463186	[LTC 6.0 FEAT] 201689:Installer - FCP automatic port discovery
463218	[LTC 6.0 FEAT] 200092:(big) Install from SCSI/FCP attached CD/DVD
463544	[LTC 6.0 FEAT] 201085:cio_ignore entry in generic.prm for LPARs
463824	[LTC 6.0 FEAT] 201755:Shutdown actions tools
464236	[LTC 6.0 FEAT] 201180:Decimal Floating Point Support in glibc
463575	[LTC 6.0 FEAT] 201157:ZFCP Performance Statistics - blktrace
463662	[LTC 6.0 FEAT] 201465:Web 2.0 - Inclusion of package rubygems
463668	[LTC 6.0 FEAT] 201470:Web 2.0 - Inclusion of package Apache MyFaces Core (JSR 252 - JSF implementation - recommended Version 1.1.x)
463669	[LTC 6.0 FEAT] 201466:Web 2.0 - Inclusion of package rubygem-rails

Network Performance Considerations

The Environment

- HiperSockets
 - Four Linux LPARs (2 IFLs ea.) to one z/OS Server LPAR (4 & 8 Cps)
 - z/OS V1R2
 - TCP buffer sizes increased from 65335B to 131071B, allowing for more unacknowledged data
 - L3 traffic only
 - IPv4, z990, Large Send (requires OSA, card-to-card, and IPv4)

The Environment

- **Throughput & Response time charts**
 - Transactions per second for interactive workloads in milliseconds
 - Megabytes per second for streams-type workloads
 - 1MB = 1,048,576 bytes/sec

- **Throughput & CPU-Milliseconds Charts**
 - # of milliseconds of processing per transaction
 - Streams-type workloads, # milliseconds of processing per megabyte

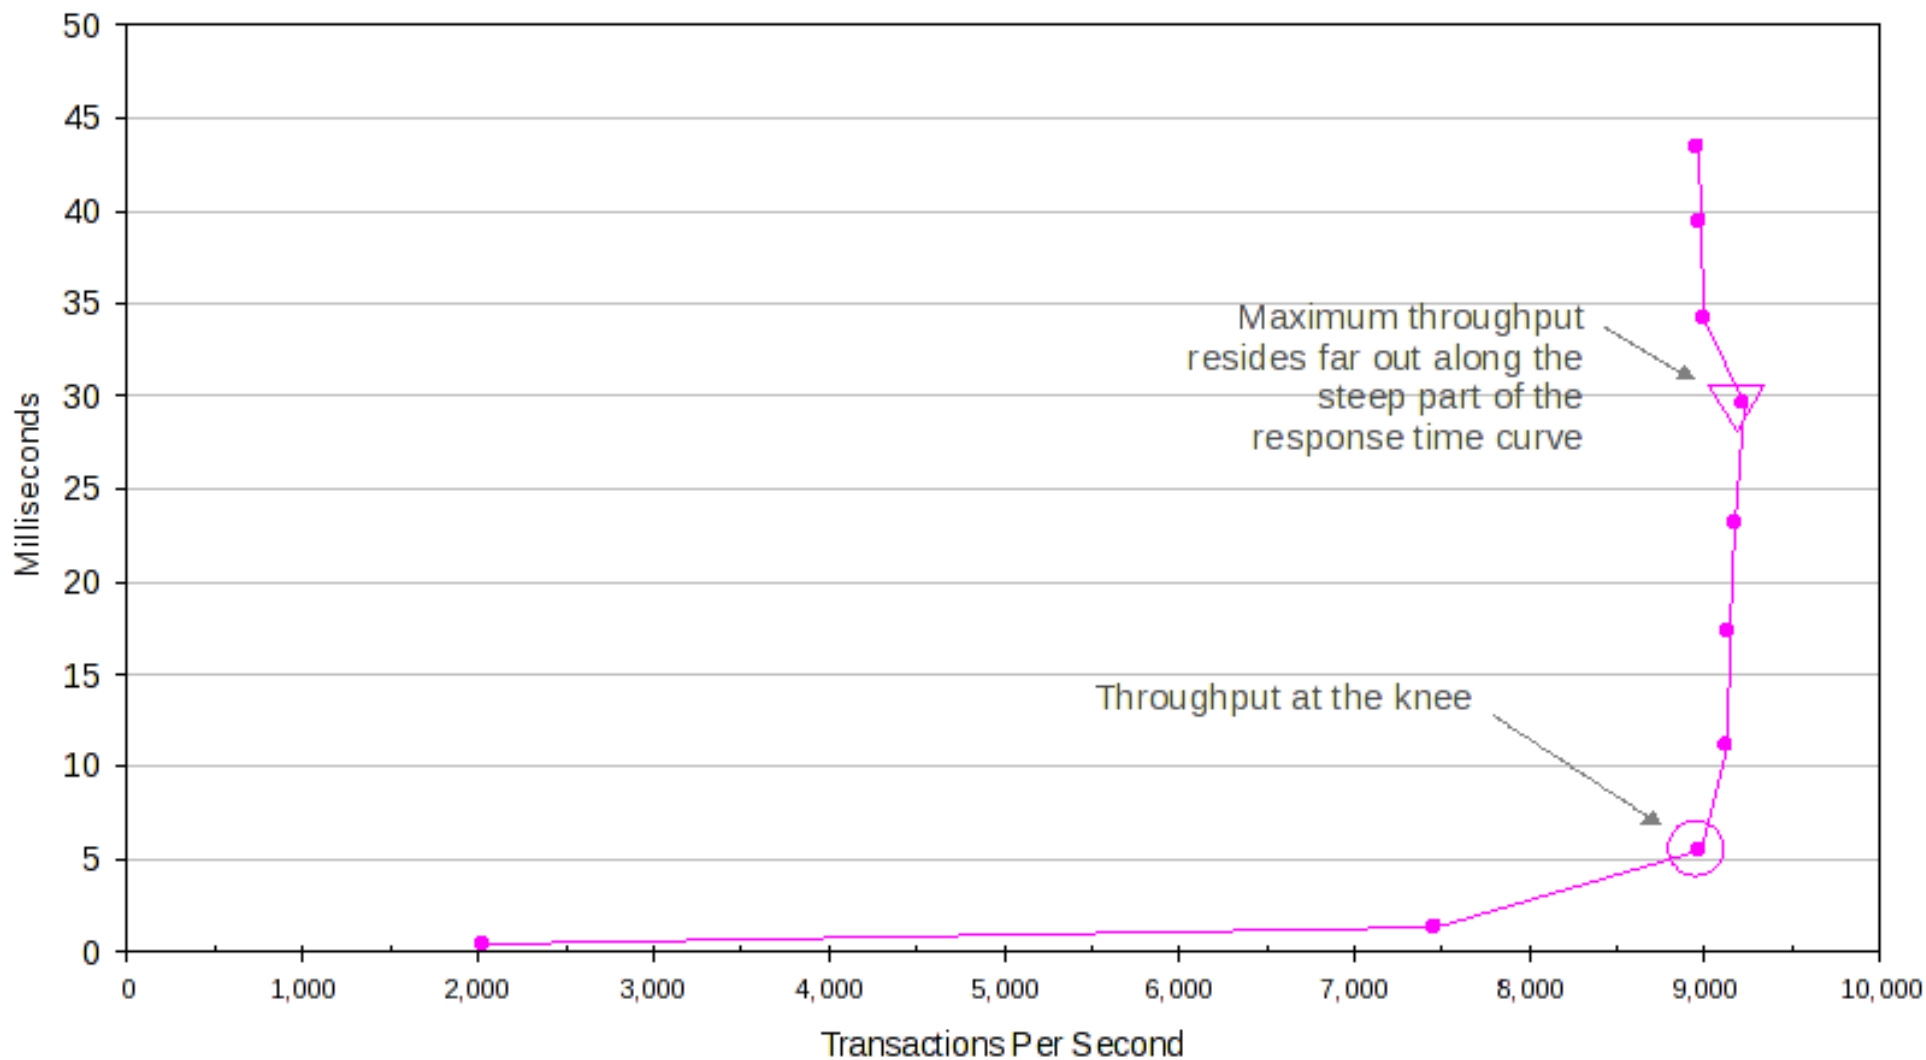
The Environment

WORKLOAD	DESCRIPTION
D cra	CRR 200B/15000B (A 200-byte Request-message from the client to the server followed by a 15000-byte Response-message from the server to the client. A TCP-connection is established for each RR.)
rra	RR 200B/1000B (A 200-byte Request-message from the client to the server followed by a 1000-byte Response-message from the server to the client. All RRs operate under that same TCP-connection.)
rrb	RR 200B/15000B
rr1	RR 1B/1B
rr64k	RR 64KB/64KB
sor	Streams Outbound (20B from the client to the server, followed by 20MB from the server to the client)
sir	Streams Inbound (20MB/20B)
mix	Mixe of <i>sor</i> and <i>sir</i> streams

OSA Response Time

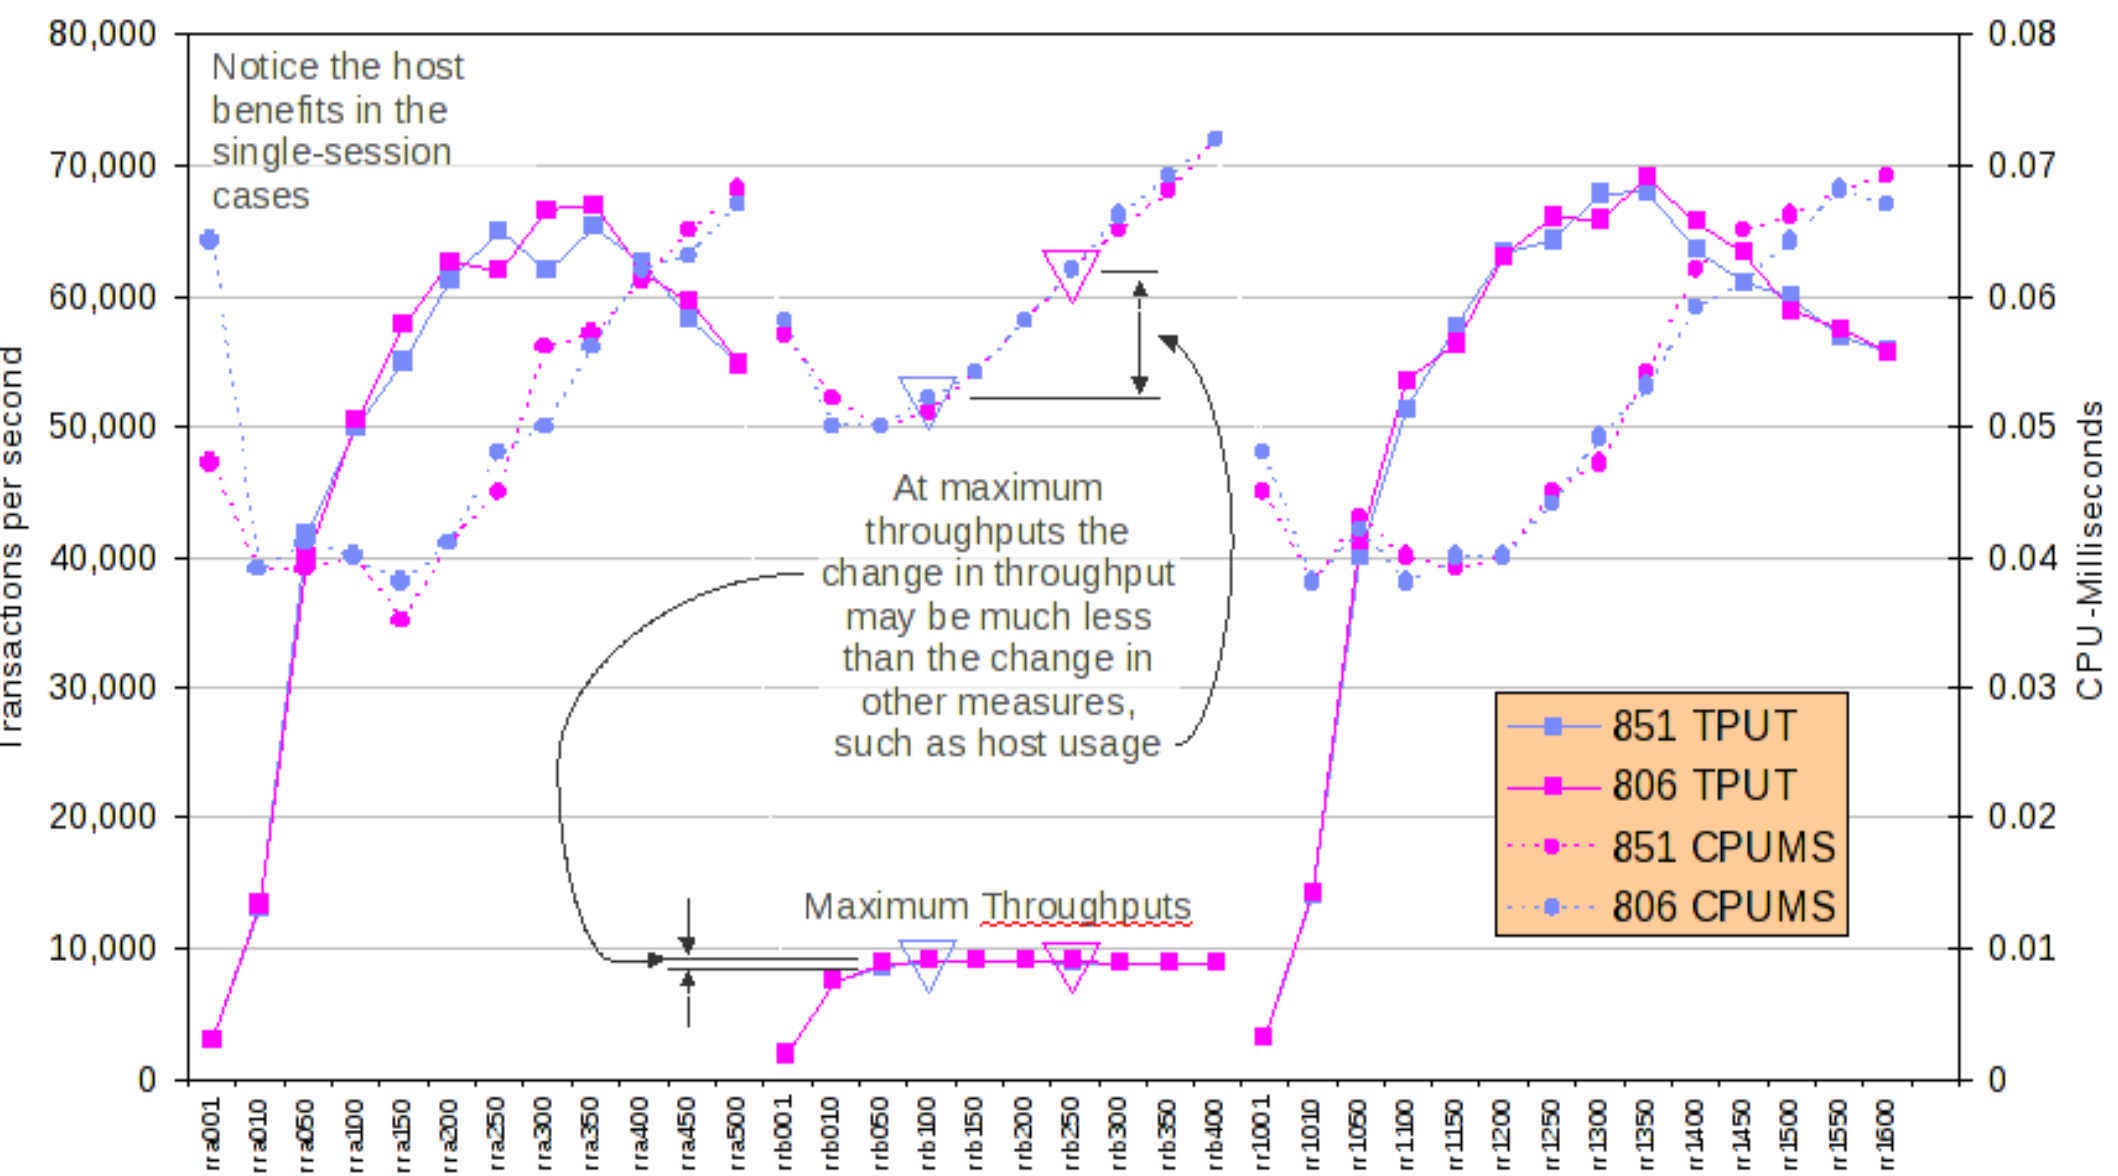
200 byte request, 15000 byte response, same tcp session

OSA-E2 10 GbE C-to-C on z9 - OSA firmware 806
RRb (200/15000) Response Time - 1492B




OSA Response Time

Transactions per Second vs CPU Time



OSA Response Time


Streams Throughput & Server CPU-Milliseconds - 1492B



OSA Response Time

Streaming Data @ 57344B/sec


z9 - HiperSockets over OSA-E2 10 GbE C-to-C (806)
Streams Throughput



OSA Response Time

Streaming Data @ 57344B/sec

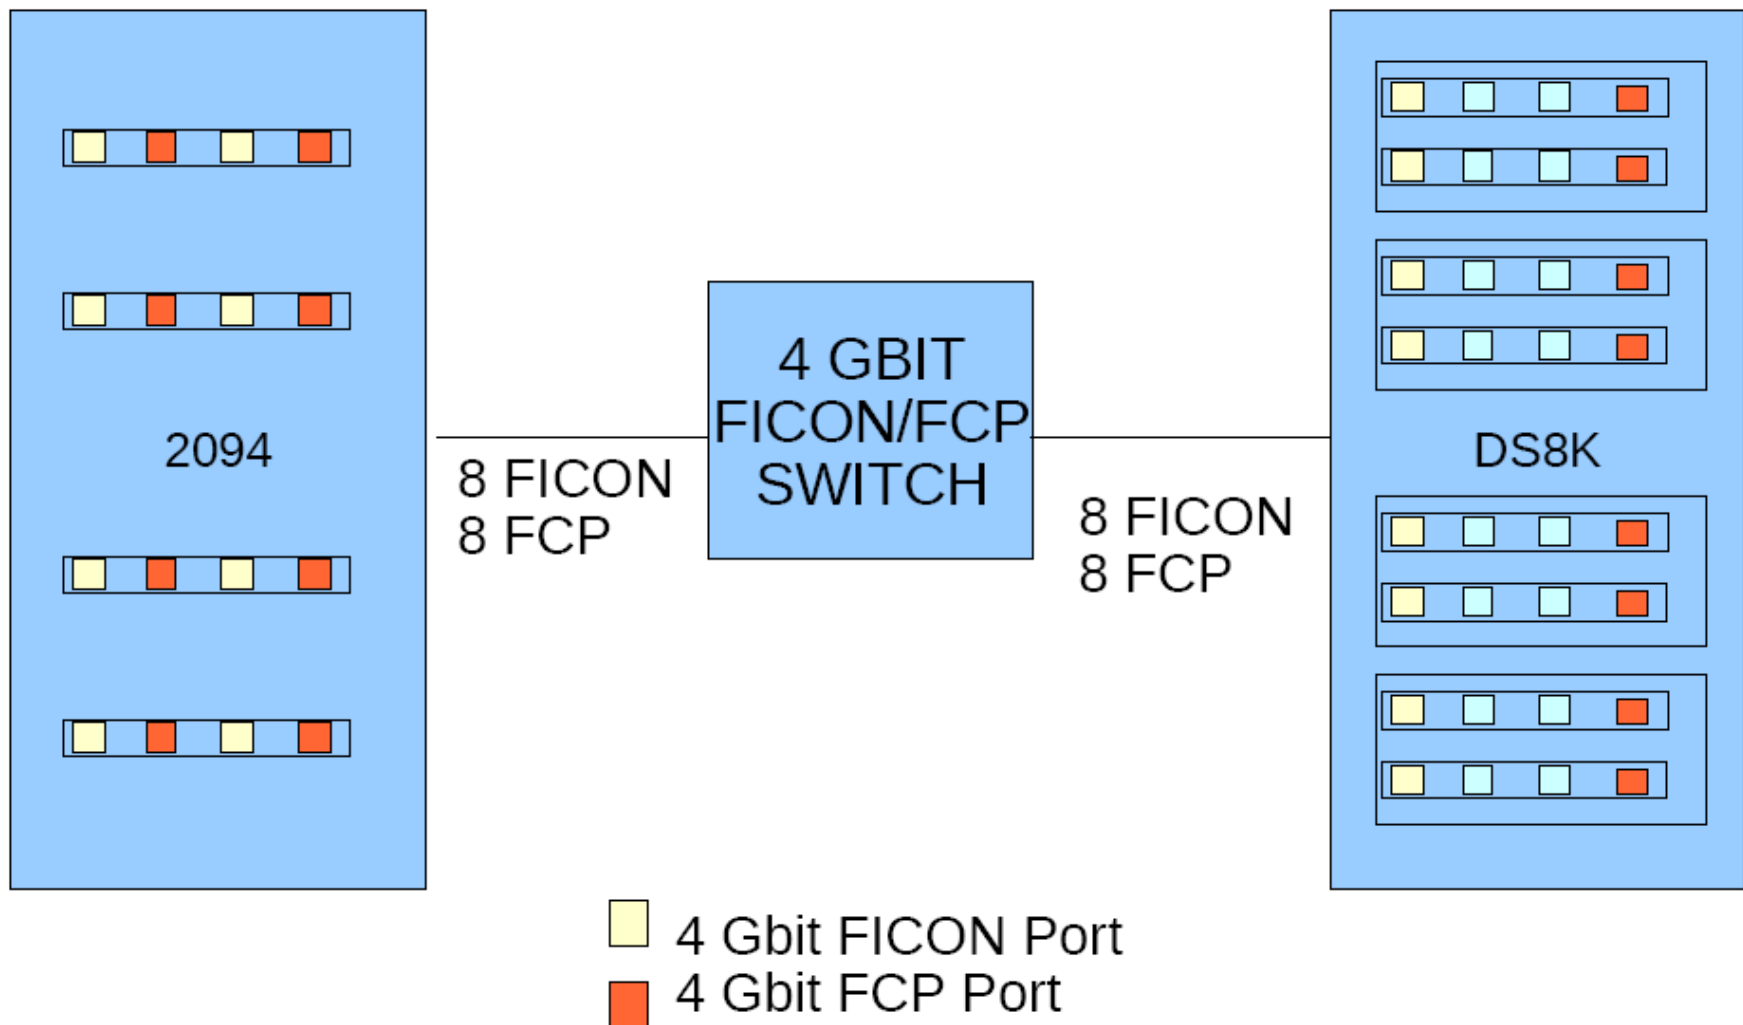
z9 - HiperSockets over OSA-E2 10 GbE C-to-C (806)
Streams Server CPU-Milliseconds



Storage Performance Considerations

The Environment

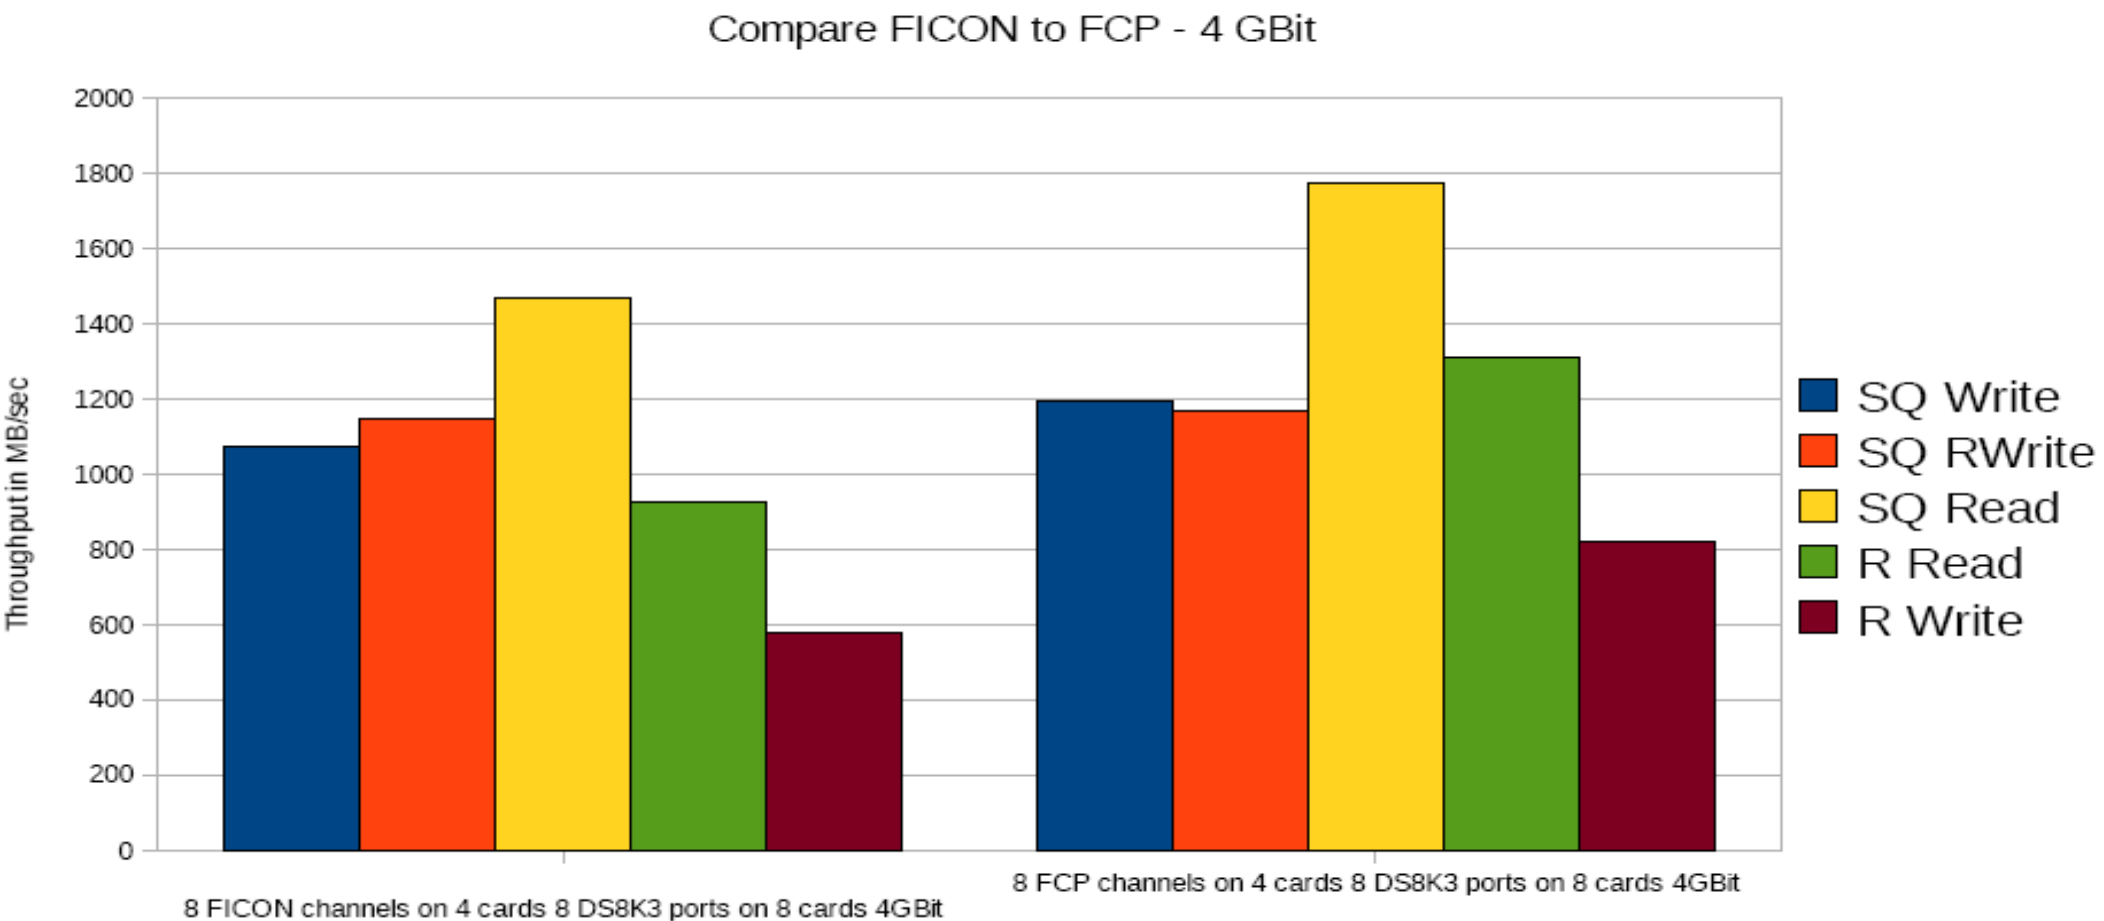
- Configuration for 4Gbps disk I/O measurements



Disk I/O Performance with 4Gbps links

– FICON vs FCP

- Throughput for sequential write is similar
- FCP throughput for random I/O is 40% higher

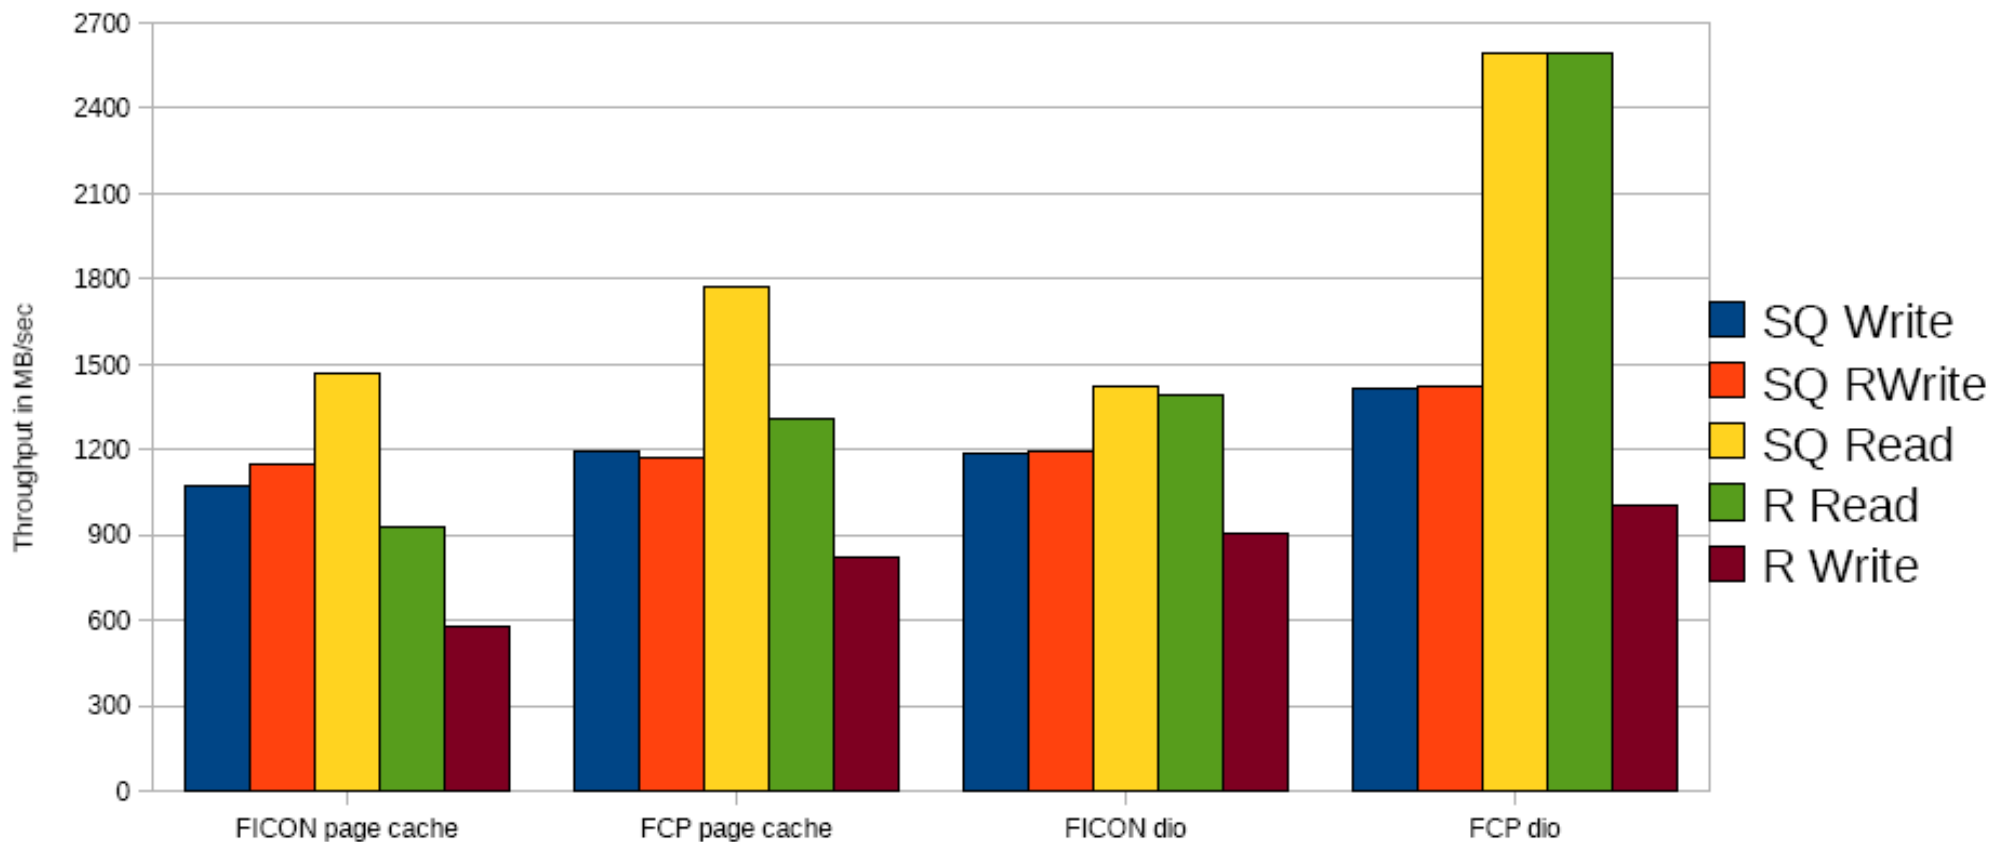


Disk I/O Performance with 4Gbps links

– FICON vs FCP / direct I/O

- Bypassing RHEL page cache improves throughput for FCP by 2x, FICON up to 1.6x
- Read operations much faster on FCP

Compare FICON to FCP - 4 GBit





Crypto Performance Considerations

Crypto Express2 Accelerator (CEX2A)


SSL handshakes

- The number of handshakes is up to 4x higher with h/w assist
- In 32 connections case, we save 50% of CPU



Crypto Performance – WebSEAL SSL Access

- Improvement by h/w support



- The connection from the client to the WebSEAL server runs encrypted using SSL (AES-128)
- Scaling the size of the requested page
- uses mostly CPACF
- Improvement up to factor 2.4 for hardware encryption versus software encryption

Memory Management Performance Considerations

CMM1 & CMMA

- **Cooperative Memory Management (CMM1)**
 - Ballooning technique
 - When z/VM detects memory constraints, it tells RHEL guests to release page frames (done by issuing a DIAG X'10')
 - When memory is freed, it alerts requesting RHEL guest

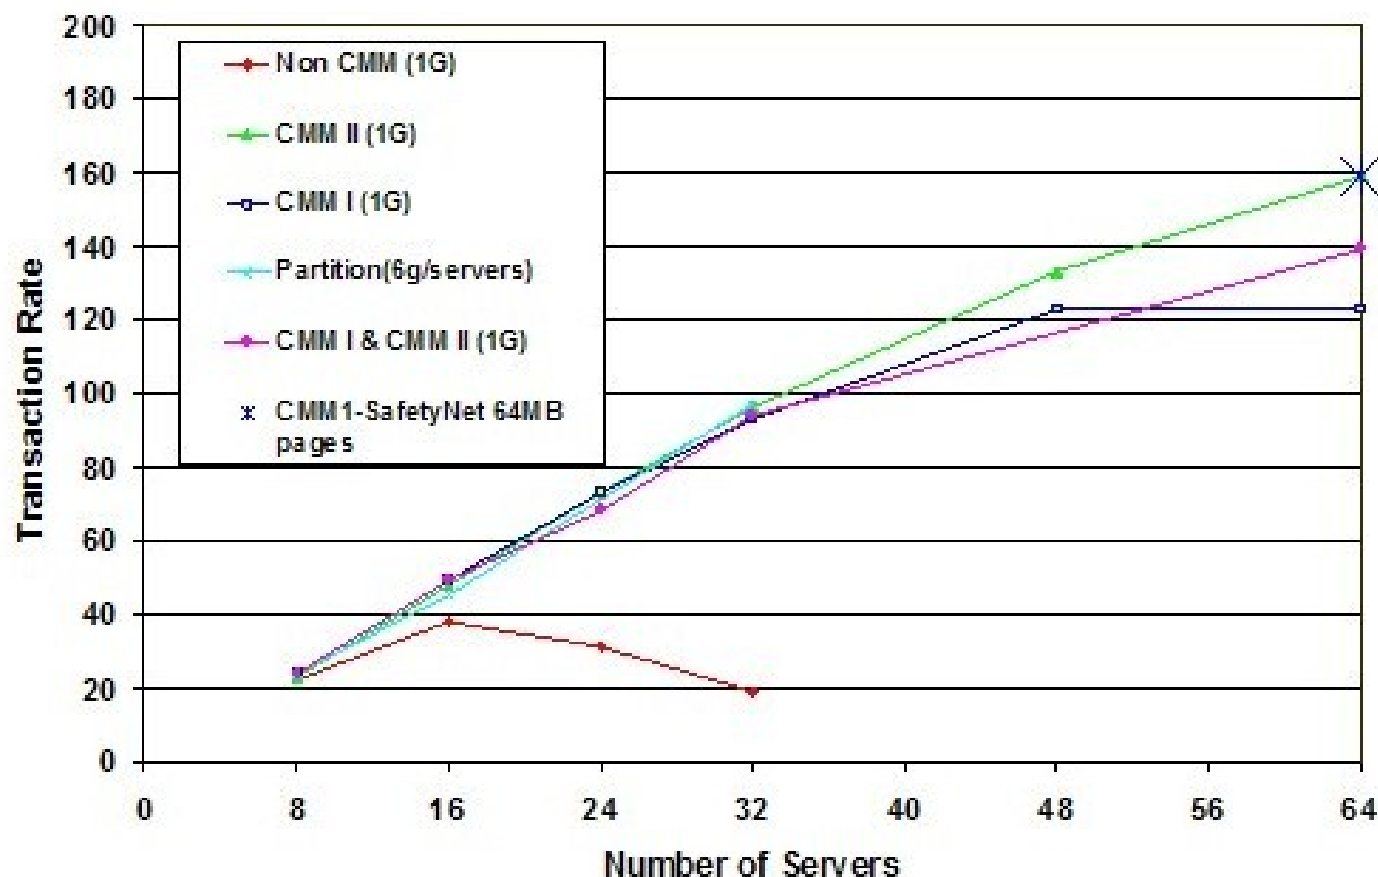
CMM1 & CMMA

- Collaborative Memory Management Assist (CMMA)
 - Page status technique
 - Stable (S)
page has essential content
 - Unused (U)
no useful content and any access to the page will cause an addressing exception
 - Volatile (V)
page has useful content. CP can discard the page anytime.
 - Potentially Volatile (P)

CMM & CMMA Performance

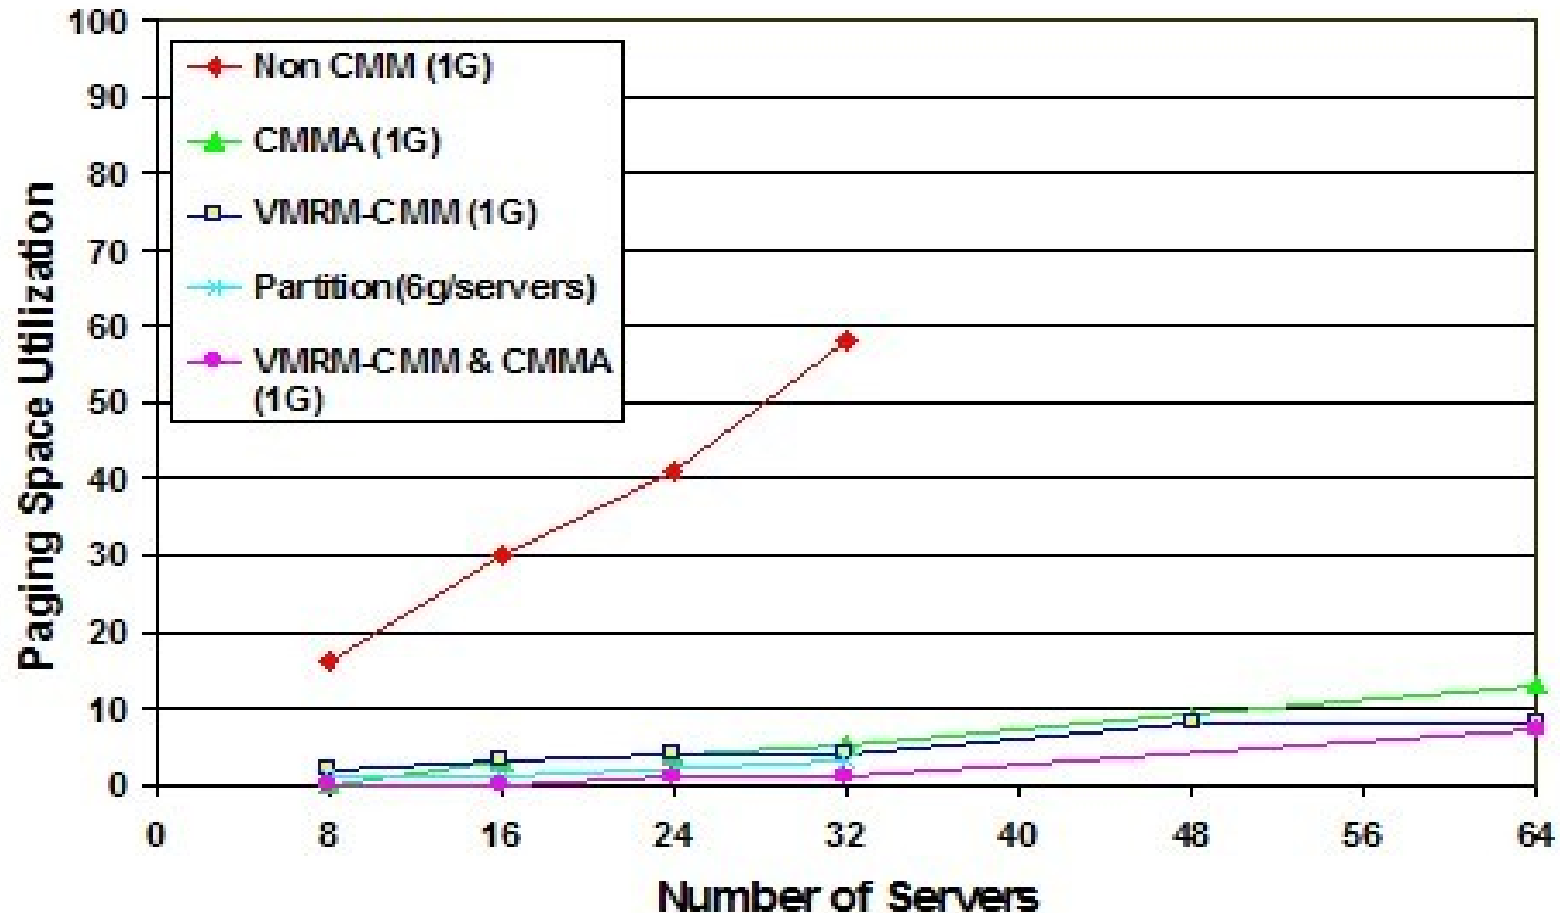
Transaction Rate vs. Number of Servers

for various Storage Management Products using Apache servers with a virtual storage size as shown in parenthesis in the legend



CMM & CMMA Performance

Paging Space Utilization vs. Number of Servers



References

Who's Actually Doing This?

Who is Actually Doing This?

■ Bank of Tokyo

- Largest known Linux on System z deployment in the world
- 275 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
- 3 IFLs for VISANet – Custom Application

Who is Actually Doing This?

■ Bank of New Zealand

- SWIFT: \$10,000,000,000/day
- PCBB: \$40,000,000/day

- 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)

Who is Actually Doing This?

■ Bank of New Zealand

- 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!

Who is Actually Doing This?

- Bank of New Zealand

	x86	System z
Power (kW/hr)	36	22 (38%)
Heat (BTUs)	110	74 (33%)
Space (Racks)	6.5	4.5 (31%)
Carbon Emissions	66	40 (39%)

What Else Should I Know?

Red Hat offers Proof of Concept support

- Free 180 day evaluation
 - For large or strategic deals, can be “supported” eval
- Dedicated Global Solution Architect (Shawn Wells, swells@redhat.com)
- Dedicated local resources, both business and technical

What Else Should I Know?

Red Hat has dedicated Linux on System z support staff

Level 1

- Front line support (basic troubleshooting, gather information, resolve known issues)
- All L1 specialists are RHCE
- Dedicated group for RHEL on zSeries machines

Level 2

- Advanced troubleshooting, callbacks, triage with customers & partners

Level 3

- Work with engineering to develop patches

TAM (Technical Account Manager)

- Provides Account Management and Pro-active Support to customers and Partners. Feature requests from Partners are handled by TAMs.