

IBM z14 Model ZR1 (M/T 3907) Technical Leadership Library

April 10, 2018 Announcement

John McLemore, z Client Architect

jemclemo@us.ibm.com

214-679-8484



Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a more complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*BladeCenter®, CICS®, DataPower®, Db2®, e business(logo)®, ESCON, eServer, FICON®, IBM®, IBM (logo)®, IMS, MVS, OS/390®, POWER6®, POWER6+, POWER7®, Power Architecture®, PowerVM®, PureFlex, PureSystems, S/390®, ServerProven®, Sysplex Timer®, System p®, System p5, System x®, z Systems®, System z9®, System z10®, WebSphere®, X-Architecture®, z13™, z13s™, z14™, z14 Model ZR1™, z Systems™, z9®, z10, z/Architecture®, z/OS®, z/VM®, z/VSE®, zEnterprise®, zSeries®, IBM Z ®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

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

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

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

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

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

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.



"I need someone well versed in the art of torture—do you know PowerPoint?"

Glossary

BPH	Bulk Power Hub
CCA	Common Cryptographic Architecture - IBM software that enables a consistent approach to cryptography on major IBM computing platforms
CPC Drawer	CPC drawer refers to the packaging of the PU and SC SCMs, Memory and PCIe Gen3 and ICA-SR fanouts
VFM	Virtual Flash Memory – Replacement for zFlash Express for z14 and zNextSF
ICA SR	Integrated Coupling Adapter Short Reach
CE LR	Coupling Express Long Reach
I/O Drawer	I/O drawer connected to InfiniBand fanouts supporting the 6 GBps InfiniBand I/O interconnect. For z13 and z13s, FICON Express8 is the only I/O feature supported in this drawer. Not available on zNextSF.
KVM	Kernel-based Virtual Machine - Open source software providing a full virtualization solution for Linux
PCIe I/O Drawer	PCIe I/O drawer connected to PCI Express Generation 2 (PCIe Gen2) 8 GBps I/O interconnect infrastructure introduced with z196/z114 or PCI Express Generation 3 (PCIe Gen3) 16 GBps PCIe I/O interconnect infrastructure used in z14 (introduced with z13 and also used in the z13s). Not available on zNextSF
PCIe+ I/O Drawer	Connected to PCI Express Generation 3 16 GBps I/O interconnect infrastructure used on z14 Model 3907; 19" form factor
RAIM	Redundant array of independent memory (RAIM). A new technology introduced with z196 designed to provide protection at the direct random access memory (DRAM), dual inline memory module (DIMM), and memory channel level
RDMA	Remote direct memory access
RG	Resource Group
RoCE	RDMA over Converged (Enhanced) Ethernet
SCH	System Control Hub
SCM	Single Chip Module. For z14, z13 and z13s, these can be either the Processor Unit (PU) or System Controller (SC) modules
SIMD	Single Instruction Multiple Data - Vector processing model providing instruction level parallelism, benefits workloads such as analytics and mathematical modeling
SMT	Simultaneous multithreading is the ability of a single physical processor (core) to simultaneously dispatch instructions from more than one hardware thread context. Because there are two hardware threads per physical processor, additional instructions can run at the same time.
SMC-D	Shared Memory Communications – Direct Memory Access over Internal Shared Memory
SMC-R	Shared Memory Communications – Remote Direct Memory Access

Glossary for I/O

Acronym	Full Name	Description / Comments
N/A	PCIe switch	Industry standard PCIe switch Application-specific integrated circuit (ASIC) used to fanout (or multiplex) the PCI bus to the I/O cards within the PCIe I/O drawer
N/A	PCIe I/O drawer	I/O drawer that supports PCIe bus I/O infrastructure; has 32 I/O card slots. Not available on z14 ZR1
PCI-IN	PCIe interconnect	Card in the I/O Drawer or the PCIe I/O drawer that contains the PCIe switch ASIC. Connects to the PCIe fanout in the CPC drawer
N/A	PCIe+ I/O drawer	New 19" form factor I/O drawer; PCIe infrastructure, has 16 I/O card slots. Available on z14 ZR1
N/A	PCIe Gen3 fanout	Card on front of CPC drawer that supports PCIe Gen3 bus; used exclusively to connect to the PCIe I/O drawer; PCIe fanout supports FICON Express8S, FICON Express16S, Crypto Express5S, OSA-Express4S, OSA-Express5S, Flash Express, 10 GbE RoCE and zEDC Express
RoCE	RDMA over CE	High speed inter communication fabric facilitating data movement between IBM Z servers and other platforms
ICA SR	ICA SR fanout	A new generation of short reach (SR) PCIe-based Coupling link, connects to ICA SRs (up to 150m) in other z14/z13/z13s systems. 8 GBps link data rate; two ports per fanout.
CE LR	Coupling Express LR	A new generation of long distance (long reach –LR) PCIe-based coupling feature residing in a PCIe I/O Drawer. Connects to another CE LR (up to 10 km unrepeated) in other z14/z13/z13s. 10 Gbps link data rate, point-to-point only.
N/A	zHyperLink Express	Short distance IBM Z I/O feature designed to work in conjunction with a FICON or High Performance FICON SAN infrastructure for interconnecting the z14 CPC directly to the I/O Bay of the DS8880
HCA3 or HCA3-O LR	HCA3-O LR fanout for 1x IFB	For 1x InfiniBand at unrepeated distances up to 10 km; 5 Gbps link data rate; <i>4 ports per fanout</i> ; may operate at 2.5 Gbps or 5 Gbps. Based upon capability of DWDM. Can communicate with an HCA2-O LR fanout; third generation Host Channel Adapter. NOT AVAILABLE on z14 ZR1
HCA3 or HCA3-O	HCA3-O fanout for 12x IFB	For 12x InfiniBand at 150 meters; supports 12x IFB and 12x IFB3 protocols; <i>increased service times when using 12x IFB3 protocol</i> ; 6 GBps link data rate; two ports per fanout; can communicate with an HCA2-O fanout; third generation Host Channel Adapter. NOT AVAILABLE on z14 ZR1

Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

Glossary for IBM Z Coupling

Acronym	Full name	Comments
AID	Adapter identification	Fanout has AID instead of a PCHID
CIB	InfiniBand Coupling	CHPID type used by z14 M0x , z13, z13s, zEC12, zBC12, z196, z114, z10, System z9; NOT supported on z14 ZR1
CS5	ICA SR Coupling	CHPID type used by z14, z13 and z13s
CL5	Coupling Express Long Reach (CE LR)	CHPID type used by z14, z13 and z13s
HCA	Host Channel Adapter	Path for communication
PSIFB	Parallel Sysplex using InfiniBand	InfiniBand Coupling Links
12x IFB	12x InfiniBand	12 lanes of fiber in each direction
1x IFB	1x InfiniBand	InfiniBand Long Reach - one pair of fiber
12x IFB3	12x InfiniBand3	Improved service times of 12x IFB on HCA3-O

Type	z14 ZR1	z14 M01, M02, M03, M04 and M05	z13/z13s	zEC12/zBC12/z196/z114
ICA SR fanout	Optical – Coupling PCIe Gen3			N/A
CE LR feature	Optical – Coupling PCIe feature in the PCIe I/O Drawer			N/A
HCA3-O fanout	Not supported	Optical - Coupling 12x InfiniBand		Optical - Coupling 12x InfiniBand
HCA3-O LR fanout	Not supported	Optical - Coupling 1x InfiniBand		Optical - Coupling 1x InfiniBand

Glossary (Miscellaneous)

- CPC : Central Processor Complex
- PU: Processor Unit = core
- SC: System Control
- FSP : Flexible Service Processor
- OSC: Oscillator Card
- CEC: Central Electronic Complex
- HCA: Host Channel Adapter
- ICA: Integrated Coupling Adapter
- 1U = 1.75 inches (1 EIA rack unit height)
- PCIe: Peripheral Component Interconnect Express
- DAT: Dynamic Address Translation
- CMS: *Conversational Monitor System*
- GMS: Group Control System
- ESA: Enterprise System Architecture
- VPD: Vital Product Data
- PSU: Power Supply Unit
- CDIMM: Centaur DIMM
- DDR4: Double Data Rate Fourth-Generation
- MES: Miscellaneous Equipment Specification
- FICON: Fibre Connection
- OSA: Open Systems Adapter
- RDMA: Remote Direct Memory Access
- RoCE: RDMA over Converged Ethernet
- zEDC: zEnterprise Data Compression
- IBF: Integrated Battery Facility
- VFM: Virtual Flash Memory

Table of Content

Naming	
Availability Dates	
IBM Z Family	
IBM z14 Model ZR1 Overview	
z14 ZR1 Details	
Processor Design and Structure	
Memory Design and Structure	
Capacity on Demand (CoD)	
Upgrades	
Virtual Flash Memory	
Capacity and Performance	
zEnterprise Data Compression (zEDC)	
I/O Overview	
I/O Details	
IBM zHyperLink Express	
I/O Innovations	
10 GbE RoCE Express2	
Coupling Express Long Reach	
Sysplex and STP	
STP Enhancements	

Cryptography Support	
Trusted Key Entry	
Reliability, Availability, Serviceability (RAS)	
SE and HMC	
Power and Cooling	
16U Reserved Feature	
eConfig	
Hypervisors – Z/VM	
Software – z/OS	
Software – z/VSE	
Software – z/TPF	
Linux and Dynamic Partition Manager	
Secure Service Container	
Statements of Direction	



IBM Z (M/T 3907) TLLB10

IBM z14 Model ZR1 availability dates – Driver Level 32

- Planned GA - May 31, 2018
- New features and functions for the IBM z14 Model ZR1 (Type number: 3907)
 - IBM z14 Model ZR1,
 - 19" Rack
 - Frame Roll MES upgrades
 - 2965 – Nxx (z13s) into 3907 – ZR1 (z14 Model ZR1)*
 - MES features available on Sept. 17th, 2018
 - Field installed features and conversions on IBM z14 that are delivered solely through a modification to the machine's Licensed Internal Code (LIC)
 - TKE 9.0 LIC (FC 0879)
 - TKE HW (new order, w/4768 Cryptographic Adapter):
 - Tower: FC 0086
 - Desktop console (KMM) for TKE (FC 0157)
 - Rack mount: FC 0085
 - HMC Tower FC 0082 for z14 (also available for z13 and z13s)
 - Desktop console (KMM) for HMC (FC 0148)
 - HMC Rack Mount FC 0083 for z14 (also available for z13 and z13s)
 - Coupling Express LR (FC 0433) on z14 (also available for z13 and z13s)
 - 16U Reserved (rack space) feature (FC 0617)



* See Speaker Notes

IBM z14 availability Dates – Driver Level 32

December 31, 2017

- IBM HMC Mobile for Z:

- <https://www-304.ibm.com/servers/resourceLink/lib03060.nsf/pages/hmcMobileApp?OpenDocument>



IBM z14
Model
ZR1



IBM Z® Generations

N-4



IBM z10 EC M/T 2097

- Announced 2/2008
- 4.4 GHz
- Up to 64 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 1.5 TB Memory



IBM z10 BC M/T 2098

- Announced 10/2008
- 3.5 GHz
- Up to 10 cfg cores (5 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 248 GB Memory

N-3



IBM zEnterprise 196 M/T 2817

- Announced 7/22/2010
- 5.2 GHz
- Up to 80 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 3 TB Memory



IBM zEnterprise 114 M/T 2818

- Announced 7/12/2011
- 3.8 GHz
- Up to 10 cfg cores (5 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 248 GB Memory

N-2



IBM zEC12 M/T 2827

- Announced 8/28/2012
- 5.5 GHz
- Up to 101 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 3 TB Memory



IBM zBC12 M/T 2828

- Announced 7/23/2013
- 4.2 GHz
- Up to 13 cfg cores (6 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 496 GB Memory

N-1



IBM z13 M/T 2964

- Announced 1/14/2015
- 5.0 GHz
- Up to 141 cfg cores
- CP, IFL, ICF, zIIP
- Up to 10 TB Memory



IBM z13s M/T 2965

- Announced 2/16/2016
- 4.3 GHz
- Up to 20 cfg cores (6 CP)
- CP, IFL, ICF, zIIP
- Up to 4 TB Memory

N



IBM z14 M/T 3906

- Announced 7/17/2017
- 5.2 GHz
- Up to 170 cfg. cores
- CP, IFL, ICF, zIIP
- Up to 32 TB cfg. Memory



IBM z14 M/T 3907
Announced 4/10/2018

- 4.5 GHz
- 19" Rack
- Up to 30 cfg cores (6 CP)
- CP, IFL, ICF, zIIP
- Up to 8 TB cfg. Memory

Ancestors ...



IBM z14 (M/T 3906 and 3907)

z14 M/T 3906



- Announced – 07/2017, GA – 09/2017
- 5 models – M05, M04, M03, M02, M01
 - Up to 170 customer configurable engines
- Sub-capacity Offerings for up to 33 CPs
- PU (Engine) Characterization: CP, IFL, ICF, zIIP, addl. SAP (No zAAPs)
- Enhanced SIMD and SMT
- On Demand Capabilities
 - CoD: CIU, CBU, On/Off CoD, CPE
- Memory – up to 32 TB
 - Up to 16 TB per LPAR (OS dependent)
 - 192 GB Fixed HSA
 - Virtual Flash Memory (zFlash Express replacement) – 1.5TB/feature
- Channel subsystem
 - Six LCSSs, up to 85 LPARs
 - Four Subchannel Sets per LCSS
- Parallel Sysplex clustering:
 - ICA SR (PCIe) Coupling,
 - InfiniBand Coupling (HCA3-O, HCA3-O LR)
 - Coupling Express Long Reach
 - Coupling Facility Control Code Level 22

- Channels
 - PCIe Gen3 16 GBps channel buses
 - FICON Express16S+ (New build) or 16S and 8S (Carry forward only)
 - IBM zHyperLink Express
 - OSA-Express6 (5S carry forward)
 - HiperSockets™ – up to 32
 - zEnterprise Data Compression (zEDC)
 - RoCE Express2, (RoCE Express - Carry forward), SMC-D
- Crypto Express6S (5S Carry forward)
- Operating Systems
 - z/OS®, z/VM®, z/VSE, z/TPF, Linux on IBM Z, the KVM hypervisor (included with Linux Distributions)
- IBM Dynamic Partition Manager (DPM)
- IBM Secure Service Container

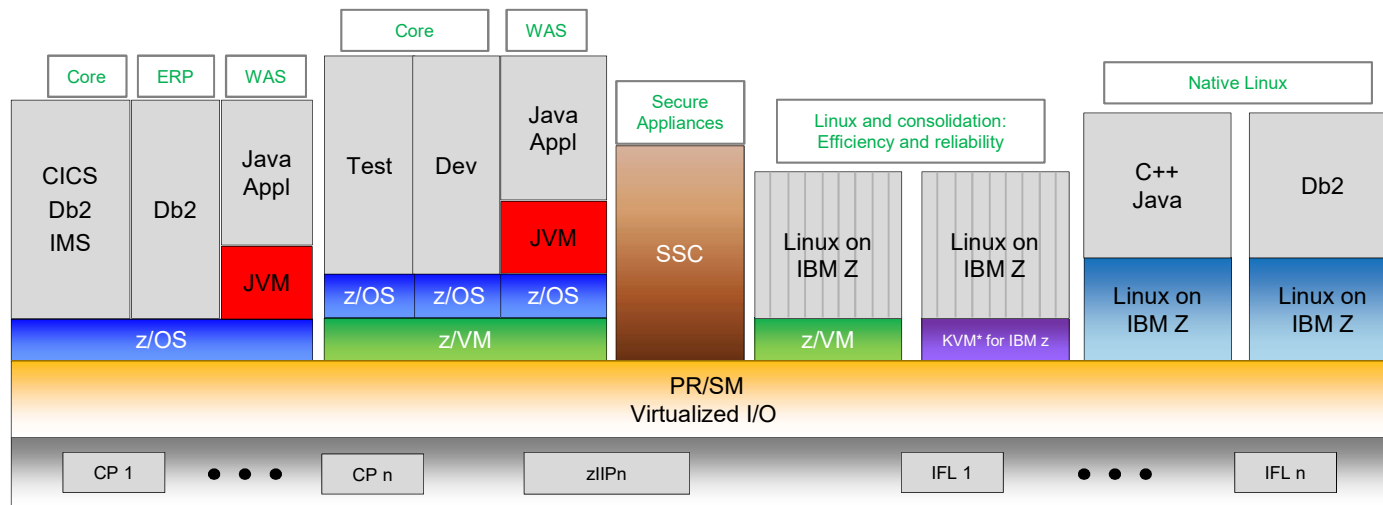
z14 M/T 3907



- Announced – 04/2018, GA – 05/2018
- One model – ZR1
 - Up to 30 customer configurable engines, Max. 6 CPs
- Sub-capacity Offerings for up to 6 CPs (156 subcapacity levels)
- PU (Engine) Characterization: CP, IFL, ICF, zIIP, addl. SAP (No zAAPs)
- Enhanced SIMD and SMT
- On Demand Capabilities
 - CoD: CIU, CBU, On/Off CoD, CPE
- Memory – up to 8 TB
 - Up to 8 TB per LPAR (OS dependent)
 - 64 GB Fixed HSA
 - Virtual Flash Memory – 512 GB/feature
- Single phase power, 19" Rack, 16U Reserved feature (FC 0617) for allowing the installation of non-Z equipment in the Z rack
- New PCIe+ I/O Drawer (up to 16 PCIe features (cards) per drawer)
- Channel subsystem
 - Three LCSSs, up to 40 LPARs
 - Three Subchannel Sets per LCSS
- Parallel Sysplex clustering:
 - ICA SR (PCIe) Coupling,
 - Coupling Express Long Reach
 - Coupling Facility Control Code Level 22
 - N-1 coupling only

IBM Z – Reliable, Scalable, Secure and Virtualized

An integrated, highly scalable computer system that allows many different pieces of work to be handled at the same time, sharing the same information as needed with protection, handling very large amounts of information for many users with security, without users experiencing any failures in service

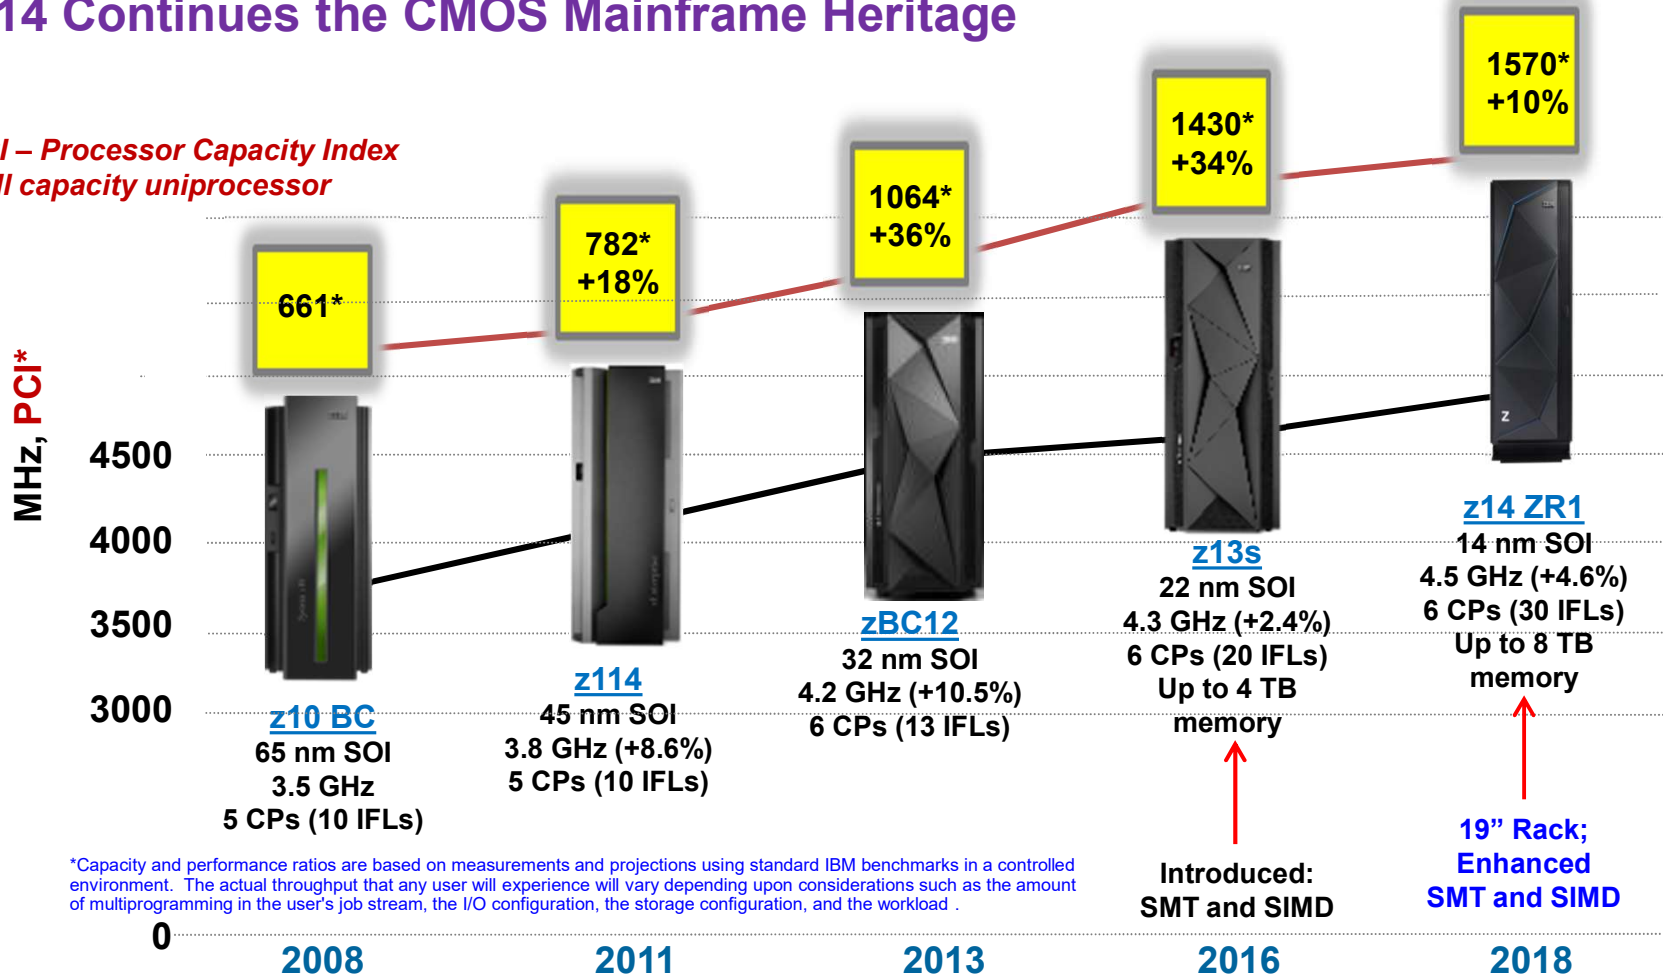


- Large scale, robust consolidation platform
- Built-in Virtualization
- 100s to 1000s of virtual servers on z/VM
- Intelligent and autonomic management of diverse workloads and system resources

* Note: See <https://www-03.ibm.com/systems/z/solutions/virtualization/kvm/>

z14 Continues the CMOS Mainframe Heritage

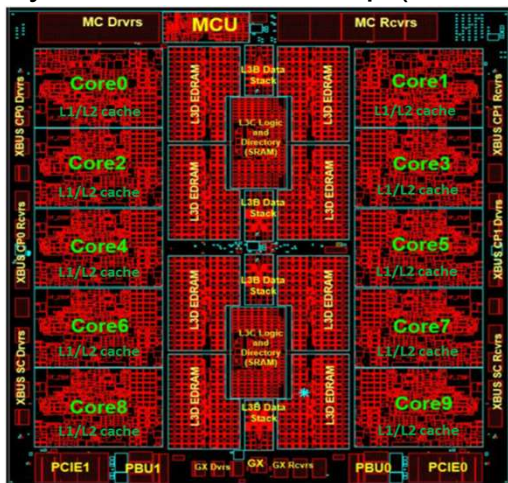
PCI – Processor Capacity Index
Full capacity uniprocessor




*Capacity and performance ratios are based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload.

z14 ZR1 System Design Changes (M/T 3907)

- 14 nm Processor Lithography with improved SIMD, SMT
- 10 Cores per PU SCM design (active cores: 5, 6, 7, 8, or 9)
- 19 inch rack
- Single CPC* Drawer
- 1, 2, or 4 PU SCMs
- Integrated I/O with PCIe Direct Attach
- Single System Controller Chip (SC SCM)



- New PCIe+ I/O Drawer
- Crypto Express6S
- OSA-Express6S
- FICON Express16S+
- 10 GbE RoCE Express2
- IBM zHyperLink Express
- Coupling Express Long Reach
- Operating environment :
 - ASHRAE Class A3



* See Speaker Notes

z14 ZR1 at a Glance



- **Machine Type**
 - 3907
- **ONE Model**
 - ZR1
 - Designed in a 19" Rack, air cooled
 - Non-raised floor support
 - Top Exit and/or Bottom Exit Cabling features
 - PDU based power
- **Processor Units (PUs) – Feature driven – 6 CPs max.**
 - Four CPC drawer features (one SC SCM):
 - Max 4 – Single PU SCM, max 4 client configurable cores (8 PUs)
 - Max 12 – Two PU SCMs, max 12 client configurable cores (16 PUs)
 - Max 24 – Four PU SCMs, max 24 client configurable cores (28 PUs)
 - Max 30 – Four PU SCMs, max 30 client configurable cores (34 PUs)
 - Two standard SAPs per system
 - One Integrated firmware processor (IFP)
 - One spare per system (all CPC drawer features)
 - 156 subcapacity settings
- **Memory**
 - Up to 8 TB including (Max memory depends on the number of installed PU SCMs):
 - System minimum = 64 GB
 - 64 GB fixed HSA separately managed
 - RAIM standard
 - Maximum for customer use 8128 GB
 - Virtual Flash Memory – part of main memory
 - 512 GB per feature, up to four features
- **I/O**
 - PCIe Gen3 channel subsystem
 - Up to 64 PCIe features
 - Up to three (3) Logical Channel Subsystems (LCSSs)
 - NO InfiniBand fanouts
- **STP* - optional (No ETR)**

* See Speaker Notes

z14 ZR1 Functions and Features (Driver Level 32)



System, Processor, Memory
One model, one CPC drawer, four available sizes
10 core 14nm PU SCM (5, 6, 7, 8, or 9 active cores per PU SCM)
Up to 30 configurable PUs as CPs, zIIPs, IFL, ICFs, or optional SAPs (up to 6 CPs)
Increased uni-processor capacity
156 Capacity settings
19" Rack, ASHRAE class A3 (for Data Center requirements relief)
Enhanced SMT (for IFLs and zIIPs only) and SIMD
Enhanced processor / cache design with bigger cache sizes
Up to 8 TB of Memory protected by Redundant Array of Independent Memory (RAIM)
16U Reserved (rack space) feature
Up to 40 LPARs
IBM Dynamic Partition Manager
Secure Service Container
LPAR Group Absolute Capping
CPUMF sampling w/o PE Mode enablement

I/O Subsystem, Parallel Sysplex, STP, Security
Up to eight (8) PCIe Gen3 I/O fanouts with 16 GBps Busses
New PCIe+ I/O Drawer (up to 4 per system, up to 64 PCIe features),
3 LCSSs, 3 Subchannel Sets per LCSS
32K I/O Devices per channel for all FICON features
FICON Express16S+ and FICON Enhancements
zHyperLink Express
10 GbE RoCE Express2
Shared Memory Communications - Direct Memory Access over Internal Shared Memory (ISM) – SMC-D
Virtual Flash Memory (512 GB per feature, up to four features)
CFCC Level 22
Crypto Express6S and Crypto enhancements

RAS, Other Infrastructure Enhancements	
Keyboard Video Monitor Switch, single display console	Ethernet switches replace SCHs
STP Enhancements - Configuration	Rack-Mounted Support Elements (CPC rack)
Key Locks for doors	Tower & Rack-mounted HMCs and TKEs
Support for ASHRAE Class A3 datacenter	TKE 9.0 LICC

IBM Z Systems Comparison (Single Frame Servers)



	zBC12 H06	zBC12 H13	z13s N10	z13s N20 ⁽¹⁾	z14 ZR1 (four feature sizes)
Uniprocessor Performance ⁽¹³⁾	1064 PCI		1430 PCI		1570 PCI
Frequency	4.2 GHz		4.3 GHz		4.5 GHz (+4.65%)
z/OS Capacity ⁽¹³⁾	50 – 4958 PCI		80 to 7123 PCI		88 to 8036 PCI
Total System Memory	240 GB	496 GB	1TB	4TB	2 to 8 TB, feature based ⁽²⁾
Configurable Specialty Engines	6	13	10	20	Max4, Max12, Max24, Max30 ⁽²⁾
Configurable CPs	0 – 6		0 to 6		0 to 4 (Max4), 0 to 6
LPARS/LCSS	30/2		40/3		40/3
HiperSockets	32		32		32
PCIe I/O drawer, I/O drawer, PCIe + I/O Drawer, Max	2, 1, 3, 0	2, 1, 3, 0	1, 1, 2 ⁽³⁾ , 0	2, 1, 3 ⁽³⁾ , 0	0, 0, 0, 4 ⁽⁴⁾
I/O slots per I/O drawers/ PCIe I/O drawers/ PCIe+ I/O drawers	8/32/0		8/32/0		0/0/16 ⁽⁴⁾
FICON® Channels	128		64 ⁽⁵⁾	128 ⁽⁵⁾	128 ⁽⁵⁾
OSA Ports	96		64 ⁽⁵⁾	96 ⁽⁵⁾	96 ⁽⁵⁾
IFB host bus Bandwidth, PCIe Bandwidth	6.0 GB/sec (IFB), 8.0 GB/sec (PCIe Gen2)		6.0GB/sec(IFB), 16.0 GB/sec (PCIe Gen3)		N/A 16.0 GB/sec (PCIe Gen3)
ISC-3, PSIFB, ICA SR, CE LR	32, 8 -16 0, 0	32, 16- 32 0, 0	0 ⁽⁶⁾ , 4-8 ⁽⁷⁾ , 8, 32	0 ⁽⁵⁾ , 16-32 ⁽⁸⁾ , 16, 32	0, 0, 16 ⁽⁹⁾ , 32 ⁽⁹⁾
zIIP/zAAP Maximum Qty	Up to 4 / 3	Up to 8 / 6	Up to 6 ⁽¹⁰⁾ / 0	Up to 12 ⁽¹⁰⁾ / 0	Up to 12 ⁽¹⁰⁾ / 0
IFL Maximum Qty ⁽¹³⁾	6 (4958 PCI)	13 (8733 PCI)	10 (10767 PCI)	20 (18335 PCI)	30 (29493 PCI)
ICF Maximum Qty	6	13	10	20	30
Capacity Settings	156	156	156	156	156
Upgradeable	H06 to H13, H13 to zEC12 H20 (Radiator-based air cooled only) or to z13s		Within z13s (N10 to N20), N20 to z13 N30 (Radiator-based /air cooled only)		z13s to z14 ZR1, Within z14 ZR1, feature based.

See next page for Notes

What is new?

- The IBM z14 Model ZR1 symmetric multiprocessor (SMP) system is the next step in an evolutionary trajectory that began with the introduction of the IBM System/360 in 1964. Over time, the design was adapted to the changing requirements that were dictated by the shift toward new types of applications that clients depend on.
- The z14 ZR1 has ultra-high frequency, large high-speed buffers (caches) and memory, superscalar processor design, with enhanced out-of-order core execution, enhanced simultaneous multithreading (SMT) and single-instruction multiple-data (SIMD) execution, and flexible configuration options. It is the next implementation of IBM Z systems to address the challenges of continuously evolving IT environment.
- The microprocessor of z14 ZR1 uses the same design as z14 M0x (M/T 3906). The difference is the frequency - 4.5 GHz vs. 5.2 GHz for z14 M0x.
- The drawer packaging is completely new, as the z14 ZR1 frame is now a 19" rack. The PU chips are air cooled (vs. closed loop water cooled for z14 M0x) as the lower frequency generates less heat.
- In addition to microprocessor evolution (described later in this section), design changes also include*:
 - New 19" form factor PCIe+ I/O drawer
 - Rear side of the rack cabling **only** with improved cable management
 - No BPA, no IBF, no HVDC option, single phase AC power only (200 ~ 240 VAC), no EPO switch
 - No InfiniBand fanouts (PCIe only)

* See Speaker Notes

z14 ZR1 Processor Enhancements: Guarded Storage Facility (GSF) for Pause-Less (Java) Garbage Collection

- **Problem:**
 - When garbage collection occurs today, all threads running under a JVM must stop
 - Customers are consolidating from multiple to single JVM environments to increase productivity and save money.
 - The consolidation effort generates heap sizes >100GB where garbage collection pauses can take minutes!
 - Long pause times cause transactional application failures and SLA violations.
- **Solution:**
 - Define flexible new architecture that provides hardware assisted read barriers for guarded storage involved in a garbage collection/compaction event.
 - Whenever a Pointer is loaded from memory, the pointer is checked against a pending GC, and in case of a “hit”, the control flow is redirected
 - The Dynamic Runtime can then assist in GC-ing the pointed-to object, before resuming the SW thread.
 - Software exploitation of fast hardware barrier detection and acceleration will allow application threads to run concurrently during the majority of garbage collection events
- **Impact (A less-pausing (Pause-Less) approach, not a no-pause approach):**
 - Reduces worst case latency impacts for critical applications like financial trading platforms
 - Maintains SLAs, and keeps IBM Z servers in our customers modernization roadmaps.

z14 M/T3907 Model ZR1 Processor Allocation/Usage – zIIP to CP 2:1 ratio

CPC Drawer Feature	CPs	IFLs	zIIP	ICFs	IFPs	Std SAPs	Add'l SAPs	Spares
Max30 (FC 0639)	0 - 6	0 – 30	0-12	0-30	1	2	0-2	1
Max24 (FC 0638)	0 - 6	0 – 24	0-12	0-24	1	2	0-2	1
Max12 (FC 0637)	0 - 6	0 – 12	0-8	0-12	1	2	0-2	1
Max4 (FC 0636)	0 - 4	0 – 4	0-2	0-4	1	2	0-2	1

- z14 M/T 3907 features **one** model: ZR1
 - z14 ZR1 can be ordered in four CPC drawer sizes, feature code driven, as shown in the table above.
 - The max. number of installable PCIe fanouts is feature dependent
 - Maximum installable memory is also feature dependent
 - The maximum number of logical ICFs or logical CPs supported in a CF logical partition is 16
 - The integrated firmware processor (IFP) is used for PCIe I/O support functions
1. At least one CP, IFL, or ICF must be purchased in every machine
 2. Two zIIPs may be purchased for each CP purchased if PUs are available. This remains true for sub-capacity CPs and for “banked” CPs.
 3. The IFP is conceptually an additional, special purpose SAP, used to run the Resource Groups used for managing the “native” PCIe features*

* See speaker notes

IBM z Integrated Information Processor (zIIP) on the z14

- The IBM z14 continues to support the z Integrated Information Processor (zIIP) which can take advantage of the optional simultaneous multithreading (SMT) technology capability. SMT allows up to two active instruction streams per core, each dynamically sharing the core's execution resources.
 - With the multithreading function enabled, the performance capacity of the zIIP processor is expected to provide up to **25%** more capacity than without SMT enabled
- The rule for the CP to zIIP purchase ratio is that for every CP purchased, up to two zIIPs may be purchased
- zAAP eligible workloads such as Java and XML, can run on zIIPs using zAAP on zIIP processing.
- zAAPs are not supported on the z14



Reminder - What workloads are eligible to run on zIIPs?

- Centralized data serving eligible for zIIP: Portions of BI, ERP, and CRM remote connectivity to Db2 V8 and later, as well as portions of long running parallel queries and select utilities.
- Network encryption on zIIP – IPsec network encryption/ decryption (starting with z/OS V1.8)
- XML parsing – z/OS XML System Services eligible on zAAP or zIIP (starting with z/OS V1.9, V1.8 and V1.7 with maintenance)
- Remote mirror – zIIP-assisted z/OS Global Mirror function (starting with z/OS V1.9)
- HiperSockets™ – HiperSockets Multiple Write operation for outbound large messages (starting with z/OS V1.9)
- Business Intelligence – IBM Scalable Architecture for Financial Reporting provides a high-volume, high performance reporting – can be eligible for zIIP processing.
- Intra-server communications – z/OS CIM Server processing eligible for zIIP (starting with z/OS V1.11).
- Db2 sort utility – Db2 utilities sorting fixed-length records using IBM's memory object sorting technique
- “zAAP on zIIP” capability – Optimize the purchase of a new zIIP or maximize your investment in existing zIIPs.
- Select Tivoli® products – for DASD scans and Performance Expert/ Performance Monitor
- Select RMF™ processing – (z/OS V2.1 and later) small portion of RMF monitoring eligible for zIIP
- Java – for WebSphere® Application Server and Java technology-based applications
- Supported for z/OS running under z/VM as a guests
- Select XML System services workloads
- Select ISV applications

z14 ZR1 Memory Design

- One Memory Controller (MCU) per processor chip with five memory channels, one DIMM per channel (No DIMM Cascading).
- The fifth channel in each MCU enables memory to be implemented as a Redundant Array of Independent Memory (RAIM). This technology has significant error detection and correction capabilities. Bit, lane, DRAM, DIMM, socket, and complete memory channel failures can be detected and corrected, including many types of multiple failures. So, RAIM takes 20% of DIMM capacity. (There is no non-RAIM option.)
- DIMM carry-forward not supported
- DIMM sizes used are 32, 64, 128, 256 and 512 GB with five DIMMs of the same size included in a memory feature. (160, 320, 640, 1280 and 2560 GB RAIM array size respectively)
- One, two, three, or four features (5, 10, 15 or 20 DIMMs)
- Features with different DIMMs sizes can be mixed in the same drawer
- 160 to 10240 GB RAIM which equates to 128 to 8192 GB addressable memory (CPC drawer size feature dependent)
- Addressable memory is required for partitions and Hardware System Area (HSA)

z14 ZR1 Purchased Memory Offering Ranges and Increments

- Purchased Memory – Memory available for assignment to LPARs
- Hardware System Area – Standard 64 GB of addressable memory for system use outside customer memory
- Standard Memory – Provides minimum physical memory required to hold customer purchase memory plus 64 GB HSA
- Preplanned Memory* – Provides additional physical memory needed for a concurrent upgrade (LIC CC change only) to a preplanned target customer memory

Model	Standard Memory GB
Max4	64 - 1984
Max12	64 - 4032
Max24	64 - 8128
Max30	64 - 8128

Memory Increment (GB)	Offered Memory Sizes (GB)
8	64, 72, 80, 88, 96
32	128, 160, 192, 256, 288, 320, 352, 384
64	448, 512, 576
128	704, 832, 960
256	1216, 1472, 1728, 1984, 2240 ... 4032
512	4544, 5056 ... 8128

* See speaker notes

z14 ZR1 Plan Ahead Memory

- Provides the capability for concurrent memory upgrades changing Licensed Internal Code (LIC CC).
 - Memory DIMMs are pre-installed to support a specified target preplanned memory size
- Preplanned memory features are chargeable
 - The charge is calculated as 50% of the cost to upgrade to the MAXIMUM standard memory ALLOWED by any LARGER DIMM HARDWARE that must be added to enable the selected Plan Ahead target, NOT on 50% of the cost of upgrade to the target the client wanted unless that happens to be the maximum supported by the larger DIMM hardware that has to be added.
(“Additional hardware” means larger hardware than needed for base memory.)
 - Features: FC 1993 – 8GB Preplanned Memory, FC 1996 – 16GB Preplanned Memory
- Preplanned memory activation is chargeable
 - Subsequent memory upgrade orders will “consume” Preplanned Memory first
 - Charged when Preplanned Memory is enabled by concurrent LIC upgrade at 50% of the cost if the Preplanned Features were not present.
- Notes:
 - Preplanned Memory is NOT temporary, On Demand Memory. Why? Because Memory LIC downgrade is disruptive.
 - **The maximum orderable amount of plan ahead memory is 2 TB**

IBM z14 ZR1 Memory Granularity (Increment Size)

- Memory Granularity for Assignment to an LPAR and for Configuration ON and OFF is the LPAR's Increment Size
 - z14 ZR1 physical increment size is fixed at 1 GB (Increased from 512 MB on z13s)
 - Central memory granularity is virtualized as a multiple of 1024 MB (1 GB) for each LPAR based on the size of the larger of its two Elements: Initial Central and Reserved Central.
 - Increment Size increases with Element size because operating systems that support memory reconfiguration (z/OS and z/VM) support only up to 512 Increments per Element.
 - Notes:
 - z14 ZR1 does NOT support Expanded Memory
 - **Action:** Review MVS™ RSU parameter. Change RSU (if non-zero) to MB or GB values, not as a number of Increments to avoid problems if LPAR memory grows. A MB or GB value will be rounded up if not equal to the partition's Increment Size.
- **NOTE:** For z/OS V2R3 with z14 ZR1, a minimum of processor storage of 8GB for z/OS LPAR and 2GB for z/OS on z/VM is required.

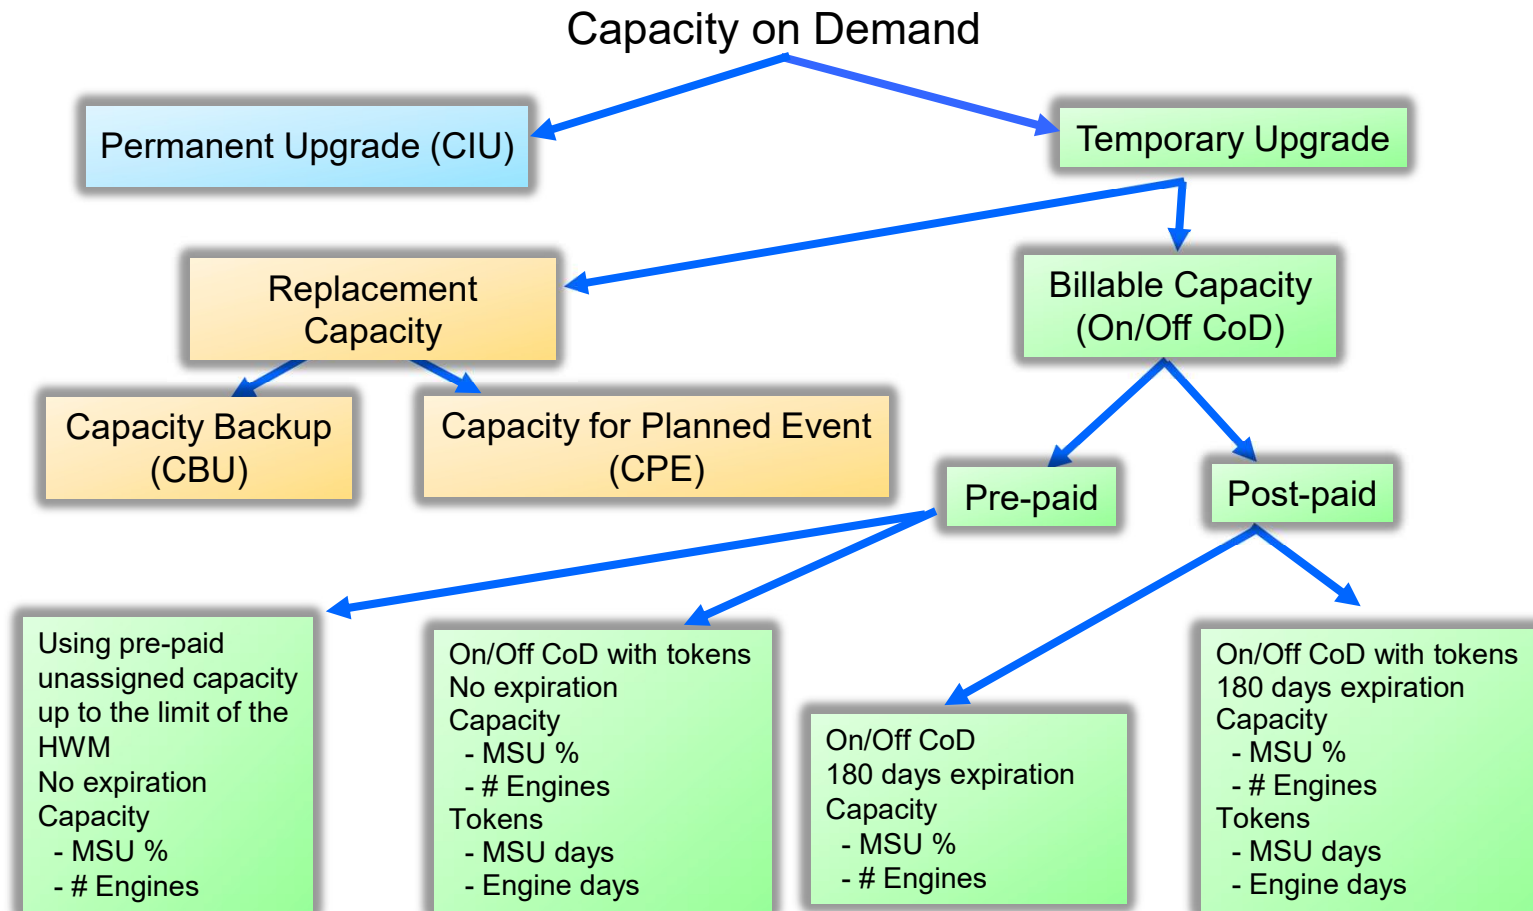
Capacity on Demand (CoD) Offerings

- **On-line Permanent Upgrade**
 - Permanent upgrade performed by customer (previously referred to Customer Initiated Upgrade - CIU)
- **Capacity Backup (CBU)**
 - For disaster recovery
 - Concurrently add CPs, IFLs, ICFs, zIIPs, SAPs
 - Pre-paid
- **Capacity for Planned Event (CPE)**
 - To replace capacity for short term lost within the enterprise due to a planned event such as a facility upgrade or system relocation
 - Predefined capacity for a fixed period of time (3 days)
 - Pre-paid
- **On/Off Capacity on Demand (On/Off CoD)**
 - Production Capacity
 - Supported through software offering – Capacity Provisioning Manager (CPM)
 - Payment:
 - Post-paid or Pre-paid by purchase of capacity tokens
 - Post-paid with unlimited capacity usage
 - On/Off CoD records and capacity tokens configured on Resource Link
- **Customer Initiated Upgrade (CIU)**
 - Process/tool for ordering temporary and permanent upgrades via Resource Link
 - Permanent upgrade support:
 - Un-assignment of currently active capacity
 - Reactivation of unassigned capacity
 - Purchase of all PU types physically available but not already characterized
 - Purchase of installed but not owned memory

z14 Basics of Capacity on Demand

Upgrade to z14: Installed On Demand records: zAAPs are converted to zIIPs and the record is migrated staged.

Upgrade to z14: Staged On Demand records: Records with zAAPs are rejected. Others are migrated staged.



Virtual Class Memory (VFM): What is it – High Level

- The "storage class memory" provided by Flash Express (FC 0402 and #0403) adapters is replaced with Virtual Flash Memory (VFM) which is part of the main memory
- VFM is offered as a priced hardware feature.
 - Customer can buy one to four "units" of VFM via eConfig at initial purchase.
- A "unit" will be 512 GB on z14 ZR1 , FC 0614*
- Much simpler management of VFM resource (HMC task)
- No hardware repair and verify (no cables, no adapters)
- Better performance since no I/O to a PCIe attached adapter takes place.
- RAS: Memory protected by RAIM and ECC (internal / main memory)

Note: Use cases and exploitation for VFM have not been changed (e.g. z/OS paging, CF shared queue overflow), they just transparently benefit from the changes in the hardware implementation.

*See speaker notes

What is it – details

- Data persistence is not supported (no change from Flash Express*)
- No change to the existing customer operating system interface i.e. existing architecture is unchanged
- Allocation of VFM storage moves to LPAR activation since LPAR hypervisor “owns” management of partition memory.
- Customer specifies initial and maximum amount of VFM
- VFM can be added or deleted by operating systems using existing “storage class memory” after partition activation
- Simplification in management of VFM since there is no hardware adapter to manage
- VFM allocations and definitions for all partitions can be viewed through “Storage Information” panel

[*See speaker notes](#)

Managing VFM – Considerations

- Concurrent add of VFM fully supported subject to LICCC and hardware installed
- Concurrent downgrade not supported, as in “regular memory”
- Plan ahead memory feature for VFM is available
- New HMC and SE interfaces for VFM management (vs. Flash Express)
- VFM allocations are specified on the Storage Tab*
 - Initial and Maximum values are specified in 16GB increments
 - Values specified can be between 0 and the amount of entitled VFM
 - Profiles imported from earlier releases set the Initial and Maximum VFM values to 0
- VFM allocations cannot be dynamically changed. The partition(s) must be activated (or re-activated) for VFM allocations to go into effect
- Storage Information Task
 - Base Storage Allocation Tab includes Entitled and Allocated VFM
 - Logical Partition Storage Allocation Tab displays active partitions and associated initial, current and maximum VFM values
- SNMP, BCPii and Web Services APIs support is provided
- **VFM is not supported in Dynamic Partition Manager mode**

*See speaker notes

Migration considerations

- VFM is a REPLACEMENT for Flash Express and has slightly different handling/configuration characteristics.
- There are some situations that cannot be handled non-disruptively.
- Pre-planning is therefore critical !
 - Question: “Hot plug” no longer available to increase amount of installed Flash Memory (VFM). So how does the customer increase the amount of VFM for a partition?
 - Answer: “Over-commitment” of VFM at partition activation time is possible. The maximum amount of VFM can be set to the LICCC value in each partition.
- Allows more storage to be added to partitions subject to amount currently not assigned (same rule applied to Flash Express implementation)
- If total amount of VFM storage allocated to all active partitions is equal to the LICCC value, but the sum of active partition maximums is larger than the installed amount, customer may be able to concurrently add VFM and increase allocations without re-activating partitions.

z14 Model ZR1 Capacity and Performance Planning

z14 ZR1 Performance: Design Highlights – Primary Performance Drivers

- Memory subsystem
 - Focused on keeping data "closer" to the processor unit
 - Larger L1, L2, and L3 caches
 - One **unified** L4 shared by L3s
 - Twice (2x) configurable memory (8 TB versus 4 TB for z13s)
- Processor
 - Improved IPC with microarchitecture enhancements
 - Pipeline optimizations and improved branch prediction
 - Cache redesigned to use virtual TLB1 and reduce TLB2 miss latency
 - Four faster dynamic address translation engines versus one for z13
 - Improved SMT for zIIPs, IFLs, and SAPs (new for z14)
 - SIMD architecture extensions for analytics and register-based decimal operations
 - New Guarded Storage Facility enables near "pause-less" garbage collection for Java
 - Improved crypto co-processor with 4x to 6x performance improvement for AES
 - Up to 10 processor units (cores) per chip versus 8 on z13
- Up to 30 configurable processor units (cores) versus 20 on z13s
- 156 sub-capacity settings (same as z13s)
- HiperDispatch
 - Exploits new chip configuration
 - Required for SMT on zIIPs
- PR/SM
 - 40 customer partitions (same as z13s)
 - Up to 170 LCPs and 8 TB memory per partition (actual limits are OS dependent)
 - Improved resource allocation algorithms based on z13s experience

z14 ZR1 Performance: Design Highlights – Additional Details

- **Processor microarchitecture improvements**
 - **4.5 GHz z14 SF vs 4.3 GHz z13s (+4.65%)**
 - Merged L1/TLB1 – eliminates TLB1 miss penalty, inlined TLB2 access on L1 miss mitigates TLB2 access penalty
 - Four HW-implemented translation engines on z14 SF vs one picocoded engine on z13s
 - 2x CRSTE (combined region segment table entry) and 1.25x PTE (page table entry) growth
 - Branch prediction improvements; 33% BTB1-and-2 growth, new perceptron predictor and simple call-return stack
 - Pipeline optimization; improved instruction delivery, faster branch wakeup, reduced execution latency, improved OSC prediction, ...
 - Optimized 2nd generation SMT; improved thread balancing, multiple outstanding translations, optimized hang avoidance mechanisms
 - Improved Hot Cache Line handling; dynamic throttling – e.g., “XI strong-arming” and a Hot Line Table
- **Cache Improvements**
 - L1 I-Cache increased from 96K to 128K per Core (1.33x)
 - L2 D-Cache increased from 2MB to 4MB per Core (2x)
 - L3 Cache increased from 64MB to 128MB per CP (2x)
 - L4 Cache changed from 480MB L4 + 224MB NIC to 672MB L4, *however only one L4 per drawer on z14 SF vs two on z13s*
 - L4 sequential prefetch
- **Storage Hierarchy Improvements**
 - Single SC Topology (Reduced system latencies)
 - Redesigned System protocols (Reduce contention points)
 - New L3 fetch miss resumption on L4 sequential pre-fetch conflict
 - Bus Feeds and Speeds
 - Wider processor store bus
 - Improved Hot Cache Line handling; contention affinity, single SC
- **Software / Firmware Improvements**
 - PR/SM: Improved Memory Affinity, improved logical partition placement algorithms based on z13s experience
 - New Translation Management Facility (avoid expensive ops)
 - New per-work-unit dispatch cache footprint metrics (cache efficiency)
 - Improved Hot Cache Line handling (new instructions)

z14 ZR1 Capacity Matrix

Z01	Z02	Z03	Z04	Z05	Z06
Y01	Y02	Y03	Y04	Y05	Y06
X01	X02	X03	X04	X05	X06
W01	W02	W03	W04	W05	W06
V01	V02	V03	V04	V05	V06
U01	U02	U03	U04	U05	U06
T01	T02	T03	T04	T05	T06
S01	S02	S03	S04	S05	S06
R01	R02	R03	R04	R05	R06
Q01	Q02	Q03	Q04	Q05	Q06
P01	P02	P03	P04	P05	P06
O01	O02	O03	O04	O05	O06
N01	N02	N03	N04	N05	N06
M01	M02	M03	M04	M05	M06
L01	L02	L03	L04	L05	L06
K01	K02	K03	K04	K05	K06
J01	J02	J03	J04	J05	J06
I01	I02	I03	I04	I05	I06
H01	H02	H03	H04	H05	H06
G01	G02	G03	G04	G05	G06
F01	F02	F03	F04	F05	F06
E01	E02	E03	E04	E05	E06
D01	D02	D03	D04	D05	D06
C01	C02	C03	C04	C05	C06
B01	B02	B03	B04	B05	B06
A01	A02	A03	A04	A05	A06
1-way	2-way	3-way	4-way	5-way	6-way
Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine

PCI: Processor Capacity Index

- Improvement over z13s
 - A01 Uniprocessor: 88 PCIs (+ 10%)
 - Z01 Uniprocessor: 1,570 PCI (+ 10%)
 - Largest z/OS capacity – Z06: 8,036 PCI (+ 10%)
 - Full Linux capacity – 30-way: 29,493 PCI (+ 54%)
- Granularity levels similar to zBC12 to facilitate upgrades and incremental growth
- 26 capacity levels x 6 CPs = 156 settings
- Nomenclature: **XY**
 - X = Capacity level (A-Z)
 - YY= Number of processors
 - A00 = ICF or IFL only (zero CPs)
- Processor Value Unit (PVU) for IFL = 100
- Any to any capacity upgrade/downgrade capability within the Model
- CBU capability from smallest to largest capacities within the Model
- On/Off CoD within the Model

→ 24 more specialty engines (w/ FC 0639 – Max30)

Capacity Planning Tools Updated for z14 ZR1

- zPCR* Processor Capacity Reference for IBM Z
- zCP3000 Performance Analysis and Capacity Planning for IBM Z
- zTPM IBM Tivoli Performance Modeler
- zBNA* IBM Z Batch Network Analyzer
- CP3KEXTR* z/OS Data Extraction Tool for CPS Tools
- CP3KVMXT* z/VM Data Extraction Tool for CPS Tools
- zSoftCap* Software Migration Capacity Planning Aid for IBM Z
- zSCON Server Consolidation on IBM Z
- zPSG Processor Selection Guide for IBM Z
- zMCAT IBM Z Migration Capacity Analysis Tool

* Note: Available to customers

zEnterprise Data Compression (zEDC)

Deploying zEDC – Review

- Operating system requirements
 - Requires z/OS 2.1 (with PTFs) and later with the zEDC Express for z/OS feature
 - z/OS V1.13 and V1.12 offer software decompression support only
 - z/VM V6.3¹ and later support for z/OS V2.1 and later guest exploitation
- Server requirements
 - Available on zEC12, zBC12, z13, z13s and z14
 - zEDC Express feature for PCIe I/O drawer (FC 0420)
 - Each feature can be shared across up to 15 LPARs
 - Up to 8 features available on zEC12/zBC12/z13/z13s/z14 ZR1 (M/T 3907)
 - Up to 16 features available on z14 M0x (M/T 3906)
 - Recommended high availability configuration per server is four features³
 - This provides up to 4 GB/s of compression/decompression
 - Provides high availability during concurrent update (half devices unavailable during update)³
 - Recommended minimum configuration per server is two features
 - Steps for installing zEDC Express in an existing zEC12/zBC12/z13/z13s/z14
 - Apply z/OS Service; Hot plug zEDC Express; update your IODF, and Dynamic Activate



¹ z/VM 6.3 not supported on z14 ZR1

² For the full zEDC benefit, zEDC should be active on ALL systems that might access or share compressed format data sets. This eliminates instances where software inflation would be used when zEDC is not available.

³ Availability rules relaxed for zEDC installed on a system with 4 RGs (z13 D27, z13s, with MCL, z14 D32)

QSAM/BSAM Data Set Compression with zEDC*

- Reduce the cost of keeping your sequential data online
- zEDC compresses data up to 4X, saving up to 75% of your sequential data disk space
- Capture new business opportunities due to lower cost of keeping data online
- Better I/O elapsed time for sequential access
- Potentially run batch workloads faster than either uncompressed or QSAM/BSAM current compression
- Sharply lower CPU cost over existing compression
- Enables more pervasive use of compression
- Up to 80% reduced CPU cost compared to tailored and generic compression options
- Simple Enablement
- Use a policy to enable the zEDC

Disclaimer: Based on projections and/or measurements completed in a controlled environment. Results may vary by customer based on individual workload, configuration and software levels.

Example Use Cases

SMF Archived Data can be stored compressed to increase the amount of data kept online up to 4X

zSecure output size of Access Monitor and UNLOAD files reduced up to 10X and CKFREEZE files reduced by up to 4X

Up to 5X more *XML* data can be stored in sequential files

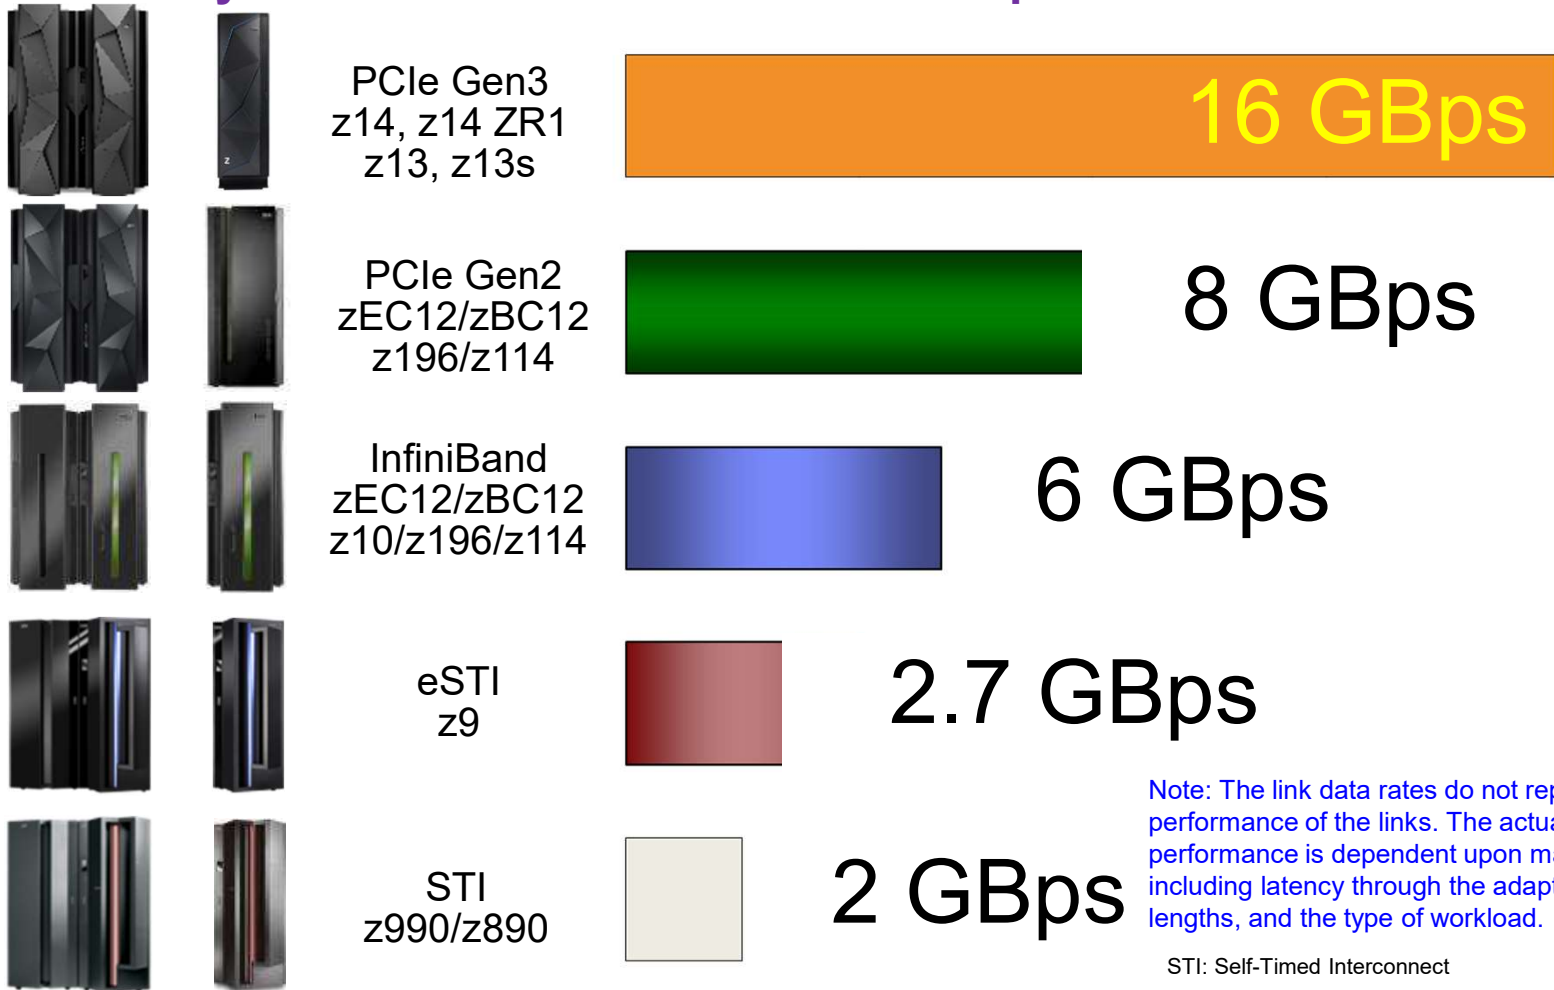
The IBM Employee Directory was stored in up to 3X less space

z/OS SVC and Stand Alone DUMPs can be stored in up to 5X less space

IBM Offers Multiple Compression Technologies for the Mainframe

Type	Optimized for:	Performance Overhead	Supported data	Frequency of access post compression
CMPSC compression on IBM Z processor chip	Optimal for Db2 or select DFSMS files	On Chip, relatively little CPU overhead and less I/O, Fast	<ul style="list-style-type: none"> • Db2 - Optimized for row-wise access to data is required • DFSMS files – for VSAM and non-VSAM extended format data sets 	Often
Other software compression (zlib, or similar)	Most compression uses industry std today. Used by many file types	Higher CPU - software instructions executed. Note: if Java then eligible for zIIP (or zAAP)	Any. De facto standard for almost any type of data.	Often
Tape HW compression	Tape compression – optimized for use with large files, archival purposes	Performed by the tape subsystem	Any.	Often – Rare (application dependent)
Archival / Backup	Archive data and data to backup/copy	CPU overhead, longer wall clock time	DFSMSHsm, DFSMSdss	Often – Rare (application dependent)
Real time compression	IBM NAS storage	No performance degradation	SVC	Designed for active primary data.
zEDC Express	Active, for cross platform data exchange. Enables compression of active and inactive data	Processing on zEDC Express-expect minimal CPU overhead, low latency	<ul style="list-style-type: none"> • SMF though logger • zlib compatible • Java • BSAM/QSAM Extended format • SOD DFSMSHsm /dss • Encryption Facility 	Frequent access required. Useful for files that previously used software compression as well

IBM Z I/O Subsystem Internal Bus Interconnect Speeds



Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

STI: Self-Timed Interconnect

“Native PCIe” FUNCTION definition, assignment and mapping

- Conceptually similar to channel (CHPID) or I/O device definition with different rules
- FUNCTION Definition in HCD or HCM to create IOCP input
 - Uniquely identified by a hexadecimal FUNCTION Identifier (FID) in the range 000 – FFF
 - NOT assigned to a Channel Subsystem so ANY LPAR can be defined to a FUNCTION.
 - Has a PARTITION parameter that dedicates it to ONE LPAR or allows reconfiguration among a group of LPARs. (A FUNCTION can NOT be defined as shared.)
 - If the intended PCIe hardware supports multiple partitions, has a decimal Virtual Function Identifier (VF=) in the range 1 – n, where n is the maximum number of partitions the PCIe feature supports.
 - May have other parameters specific to the PCIe feature.
For Example, RoCE Express features require a Physical Network Identifier (PNETID=).
- FUNCTION Mapping to hardware
 - Assign a Physical Channel Identifier (PCHID=) to identify the hardware feature in a specific PCIe I/O drawer and slot to be used for the defined FUNCTION.
 - Methods:
 - Manually using the configurator (eCONFIG) “AO Data” report
 - With assistance using the CHPID Mapping tool with eConfig Configuration Report File (CFR) input

Notes:

1. Unlike CHPIDs, multiple FUNCTIONS can be mapped to the SAME PCHID. This is conceptually similar to mapping multiple InfiniBand coupling CHPIDs to the same adapter and port.
2. CE LR is defined in IOCDS similar to PSIFB, uses PCHID instead of AID to identify the physical card and cannot have a “native” definition for the PCHID (no FUNCTION statement)

Application and Activation of Machine Change Levels (MCLs) Licensed Internal Code (LIC) Serviceability

- MCL Application: By design, all MCLs can be applied concurrently including:
 - During MCL Bundle application
 - During Enhanced Driver Maintenance (Concurrent Driver Upgrade). For example, for z13 at Driver 22, the original GA level, upgrade from Driver 22 to Driver 27, the z13s and z13 GA2 LIC driver.
- MCL Activation: By design and with planning, MCLs can be activated concurrently:
 - A large majority of MCLs activate concurrently when applied
 - A few MCLs are “Pended” for scheduled activation because activation is disruptive in some way (Recent History: Most commonly seen for traditional OSA-Express features)
 - Activate traditional I/O Feature Pended MCL – LIC on the hardware feature
 - Display Pending MCLs using HMC function or Resource Link Machine Information Reports
 - Activate using HMC function on a feature basis by PCHID one at a time – disruptive: CONFIG the CHPID OFF to all sharing LPARs, activate, and then CONFIG ON to all
 - Activate Native PCIe Pended MCL – LIC on a hardware feature OR Resource Group (RG) LIC
 - Display Pending MCLs using HMC function or Resource Link Machine Information Reports
 - Feature LIC: Activate using HMC function on a one feature (PCHID) at a time basis - disruptive: CONFIG FUNCTIONS mapped to the feature OFF to all LPARs, activate, and then CONFIG ON
 - RG LIC: Activate using HMC function to each RG in turn – disruptive to all PCHIDs in the RG: CONFIG all FUNCTIONS mapped to all PCHIDs in RG1 OFF, activate, then CONFIG ON. Repeat for all PCHIDs in RG2, RG3, RG4
 - Note: For hardware that does not need CHPID or a FUNCTION definition, for example Crypto Express, an a different method that is specific to the feature is used.
 - Alternative: Apply and activate all Pended MCLs disruptively with a scheduled Power On Reset (POR)

What is IBM zHyperLink™?

- zHyperLink Express is a direct connect short distance IBM Z I/O feature designed to work in conjunction with a FICON or High Performance FICON SAN infrastructure
- IBM zHyperLink™ dramatically reduces latency by interconnecting the z14 CPC directly to the I/O Bay of the DS8880
- zHyperLink improves application response time, cutting I/O sensitive workload response time in half without significant application changes.



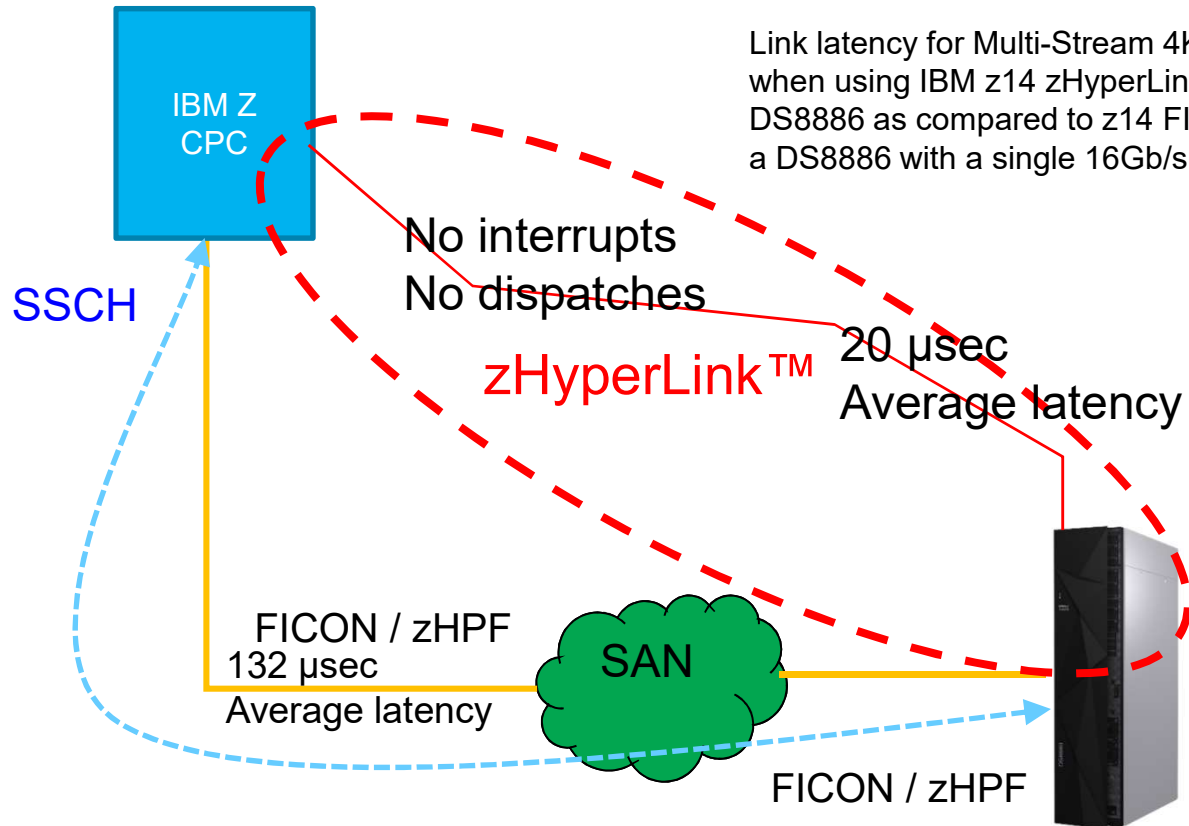
Why?

- New business requirements demand fast and consistent application response times
 - New storage technologies like Flash storage are driven by market requirements of low latency
 - Low latency helps organizations to improve customer satisfaction, generate revenue and address new business opportunities
 - Low latency drove the high adoption rate of new and enhanced I/O technologies including zHyperWrite, FICON Express16s, HyperPAV, and zHPF
- Clients can save application development costs with a low latency infrastructure
 - It is expensive to completely re-engineer existing applications to meet latency objectives
 - Low Latency I/O enables shorter batch windows and a more resilient I/O infrastructure with predictable and repeatable I/O performance.
- Improving transaction latency provides value

In laboratory measurements using a IBM z14 zHyperLink Express attached to a DS8886, Multi-Stream 4K Read link latency is reduced by up to 5x as compared to z14 FEx16S+ attached to a DS8886 with a single 16Gb Host Adapter.

Disclaimer: This result was measured in a controlled environment running an I/O driver program under z/OS on a z14. The actual link latency that any user will experience may vary. Link latency only accounts for a portion of I/O latency as experienced by an application. It reflects part of the benefit that zHyperLink Express provides for z/OS, middleware, and client applications.

IBM Z - Today and Future I/O



Link latency for Multi-Stream 4K Reads is reduced by up to **5x** when using IBM z14 zHyperLink Express attached to a DS8886 as compared to z14 FICON Express 16S+ attached to a DS8886 with a single 16Gb/s Host Adapter

Notes:

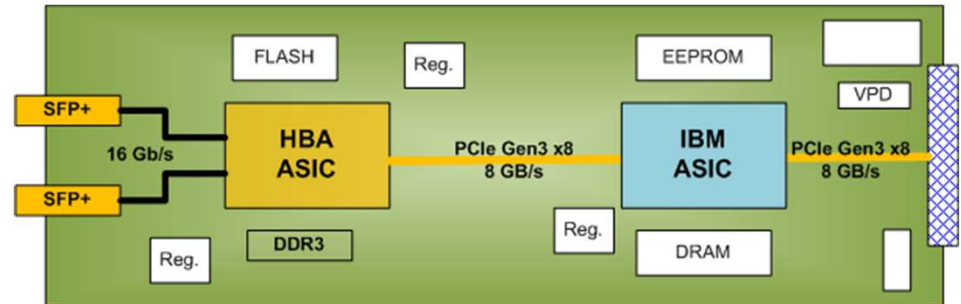
- Typically over 80% cache hit ratio on random reads.
- 100% cache hit on writes.

Disclaimer: Performance results are extrapolated from benchmark measured in a controlled environment. The actual link latency that any user will experience may vary. Link latency only accounts for a portion of I/O latency as experienced by an application. It reflects part of the benefit that zHyperLink Express provides for z/OS, middleware, and client applications.


FICON Express16S+

NEW


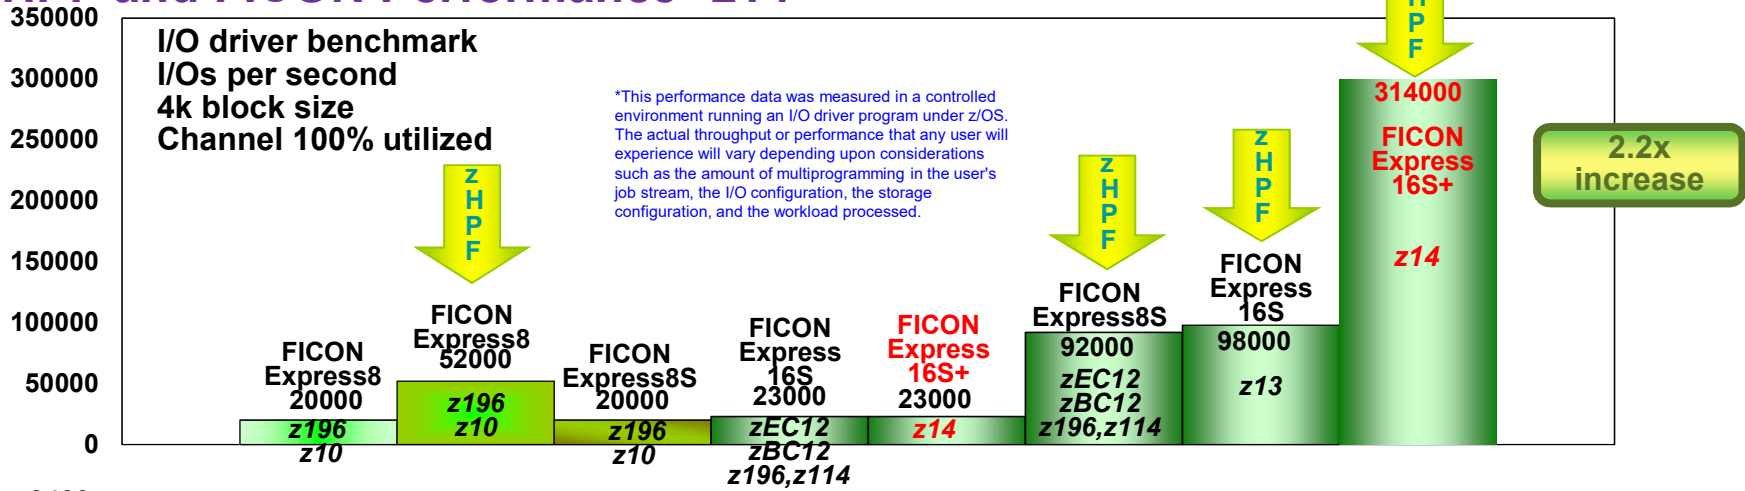
- **For FICON, zHPF, and FCP**
 - CHPID types: FC and FCP
 - **Both ports must be same CHPID type**
 - 2 PCHIDs / CHPIDs
- **Auto-negotiates to 4, 8, or 16 Gbps**
 - 2 Gbps connectivity not supported
 - FICON Express8S will be available for 2Gbps (carry forward only)
- **Increased performance compared to FICON Express16S**
- **Small form factor pluggable (SFP) optics**
 - Concurrent repair/replace action for each SFP
 - 10KM LX - 9 micron single mode fiber
 - Unrepeated distance - 10 kilometers (6.2 miles)
 - SX - 50 or 62.5 micron multimode fiber
 - Distance variable with link data rate and fiber type
- **Two channels of LX or SX (no mix)**



FC 0427 – 10KM LX, FC 0428 – SX



zHPF and FICON Performance* z14

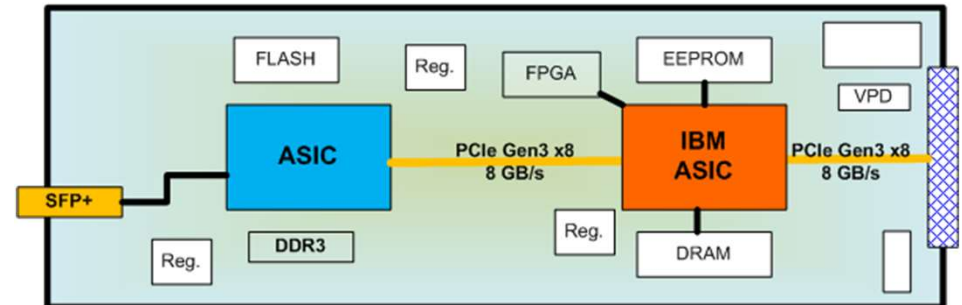


OSA-Express6S Fiber Optic Features*

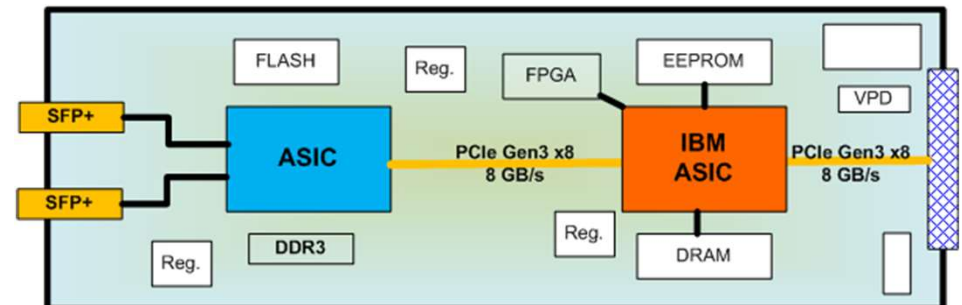
- 10 Gigabit Ethernet (10 GbE)
 - CHPID types: OSD, OSX
 - Single mode (LR) or multimode (SR) fiber
 - One port of LR or one port of SR
 - 1 PCHID/CHPID
 - Small form factor pluggable (SFP+) optics
 - LC duplex

- Gigabit Ethernet (1 GbE)
 - CHPID types: OSD (OSN not supported)
 - Single mode (LX) or multimode (SX) fiber
 - Two ports of LX or two ports of SX
 - 1 PCHID/CHPID
 - Small form factor pluggable (SFP+) optics
 - Concurrent repair/replace action for each SFP
 - LC Duplex

NEW



FC 0424 – 10 GbE LR, FC 0425 – 10 GbE SR




FC 0422 – GbE LX, FC 0423 – GbE SX



OSA-Express6S 1000BASE-T Ethernet feature



- PCIe form factor feature supported by PCIe I/O drawer
 - One two-port CHPID per feature
 - Half the density of the OSA-Express3 version
- Small form factor pluggable (SFP+) transceivers
 - Concurrent repair/replace action for each SFP
- Exclusively Supports: Auto-negotiation to 100* or 1000 Mbps and full duplex only on Category 5 or better copper
 - No 10Mbps
 - RJ-45 connector
 - Operates at “line speed”



FC 0426

Connector = RJ-45

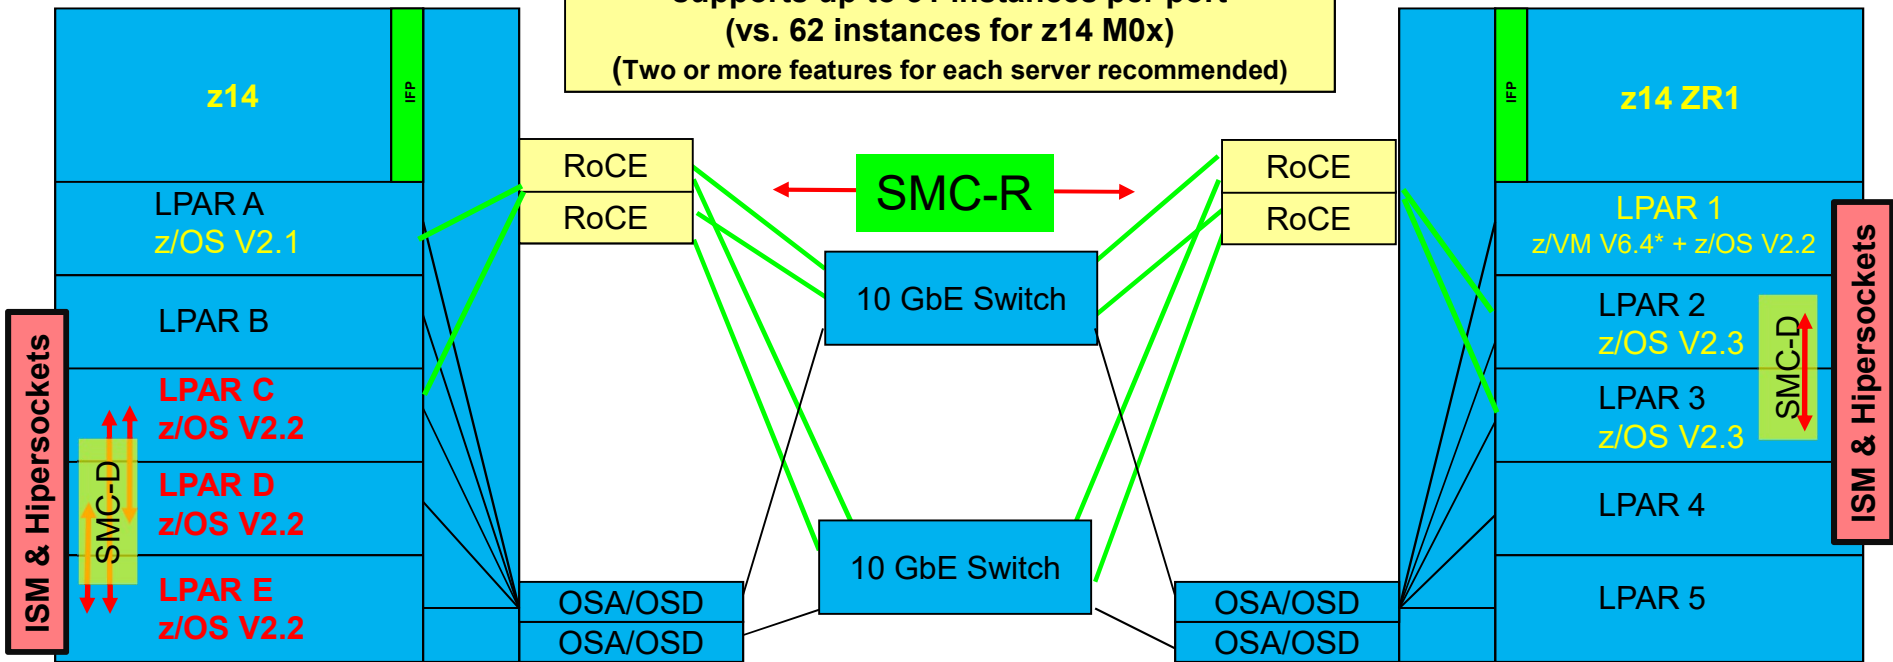
CHPID TYPE Support:

Operation Mode	TYPE	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, OS system console operations
QDIO	OSD	TCP/IP traffic when Layer 3, Protocol-independent when Layer 2
Non-QDIO	OSE	TCP/IP and/or SNA/APPN/HPR traffic
Unified Resource Manager	OSM	Connectivity to intranode management network (INMN)
OSA for NCP (LP-to-LP)	OSN	NCPs running under IBM Communication Controller for Linux (CCL)

* OSA-Express6S 1000BASE-T adapters (#0426) will be the last generation of OSA 1000BASE-T adapters to support connections operating at 100 Mb/second link speed. Future OSA-Express 1000BASE-T adapter generations will support operation only at 1000 Mb/second (1Gb/s) link speed.

z14 ZR1: Shared Memory Communications – (SMC-R and SMC-D)

On z14 ZR1, each 10 GbE RoCE FC 0412 supports up to 31 instances per port (vs. 62 instances for z14 M0x)
(Two or more features for each server recommended)



- This configuration allows SMC-D connectivity among LPAR C, LPAR D, and LPAR E.
- SMC-D within one machine is better than using HiperSockets alone.
- For LPAR to LPAR, HiperSockets or OSD connections are required to establish the SMC-D communication.
- ISM = Internal Shared Memory
- No additional hardware purchase required.
- z/VM Guest support

Coupling Express Long Reach Long Distance Coupling

Overview

- New adapter and channel type for long distance coupling
 - Feature Code #0433
- Coupling Express LR adapter in PCIe/PCIe+ I/O drawer
 - 32 features per system for z14 and z13 (Two ports per feature)
 - 16 features for z14 ZR1 and z13s (Two ports per feature)
 - Long-distance optics/fiber
 - Fiber is same single-mode fiber as used for ISC and PSIFB-1x (9/125 μ m)
 - 10km unrepeated distance, up to 100 km with qualified DWDM, [More than 100 km requires RPQ 8P2781](#).
 - 10 Gbps link speed*
 - Point-to-point only (no switching)
- Single PCHID identifies card/slot
- Four channels (CHPIDs) per port
 - Each channel is identified by VCHID/CSS.CHPID
 - CHPID type (CL5)

***Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.**

What is Long Distance Coupling using RoCE and Why?

- RoCE (RDMA over Converged Ethernet)
 - All InfiniBand methods (Send/Receive, atomic-ops, RDMA) are defined
 - Definition is done by the IBTA (<http://www.ibta.org>)
 - Concepts of QP, CQ, and EQ are carried forward *
 - UD, RC service types are carried forward
 - Point-to-point topology ONLY (Subnet management/administration functions removed)
 - Support similar to RoCE Express
 - Physical Function owned by PCI Support Partition (PSP)
 - Resource groups like RoCE Express/ RoCE Express2
 - Adapter and Link initialization
 - Diagnostics and Service
 - Single Virtual Function owned by IOP coupling firmware
 - STP support
 - Using CQE timestamps
 - No Going-Away-Signal (GAS)
- Why the new adapter?
 - Provides a non-InfiniBand based solution for long distance coupling connectivity.

* See Speaker notes

Channel Definition

- Defined in IOCDs similar to PSIFB
- Defines PCHID instead of AID to identify the physical card
 - CHPID
PATH=(CSS(0),34),TYPE=CIB,CPATH=35,AID=0B,PORT=1,PART=LP1,CSYSTEM=BIGBUBBA
 - CHPID
PATH=(CSS(0),34),TYPE=CL5,CPATH=35,PCHID=234,PORT=1,PART=LP1,CSYSTEM=BIGBUBBA
- Cannot have a “native” definition for the PCHID (no FUNCTION statement)
 - Illegal for CL5 PCHID (no FIDs) --> FUNCTION FID=5,PCHID=234,PART=((LP1),(LP1))
- Connected channel path is specified for the same reasons
 - Still allow for CSYSTEM in order to allow switches in the future
 - Still optional for point-to-point link

Section Content

- z14 ZR1 Coupling Enhancements
 - CE LR Coupling Links and Coupling Link Migration
 - Large Memory
 - Coupling Facility (CF) Processor Scalability
 - CF List Notification Enhancements
 - Coupling Link Constraint Relief
- CF structure encryption
- Asynchronous duplexing of CF lock structures

Parallel Sysplex Coupling Links – ICA SR

- IBM Integrated Coupling Adapter (ICA SR) – FC 0172
 - Coupling Connectivity into the Future (Short Distance)
 - ICA SR is Recommended for Short Distance Coupling z13/z13s to z13/z13s and beyond
 - Coupling channel type: CS5
 - Performance similar to Coupling over Infiniband 12X IFB3 protocol
 - PCIe Gen3, Fanout in the CPC drawer, 2-ports per fanout, up to 150m;
 - 8 GigaBytes per second (GBps)**.
 - z13/z13s/z14(M0x and ZR1) to z13/z13s/z14(M0x and ZR1) and up connectivity
 - Maximum configurations supported*:
 - 40 links per z13 CPC; Up to 4 CHPIDs per port, 8 buffers (i.e. 8 subchannels) per CHPID
 - 80 links per z14 CPC; Up to 4 CHPIDs per port, 8 buffers (i.e. 8 subchannels) per CHPID
 - 16 links per z14 ZR1 CPC; Up to 4 CHPIDs per port, 8 buffers (i.e. 8 subchannels) per CHPID
 - 32 links per z13s CPC; Up to 4 CHPIDs per port, 8 buffers (i.e. 8 subchannels) per CHPID
 - ICA requires new cabling for single MTP connector; cables - 150m: OM4; 100m OM3
 - Differs from 12X Infiniband split Transmit/Receive connector;
 - Available as of z13 GA1

****Note:** The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

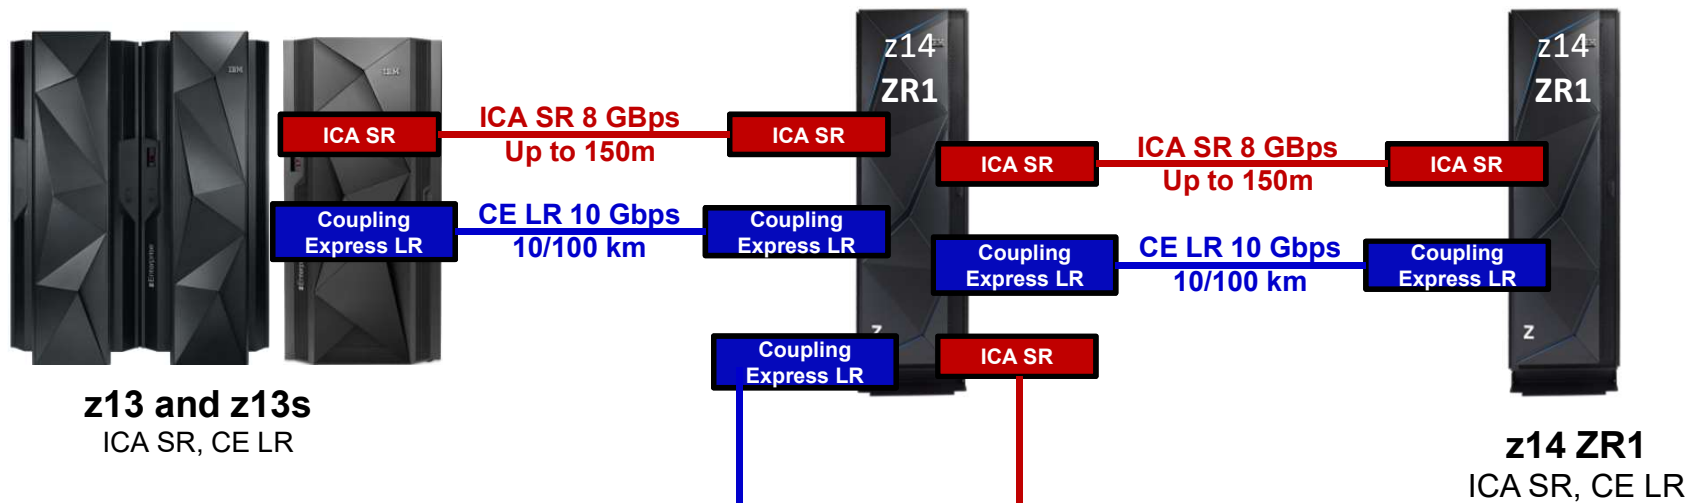
* See Speaker notes

Parallel Sysplex Coupling Links – CE LR

- Coupling Express Long Reach (CE LR) FC 0433
 - Coupling Connectivity into the Future (Long Distance)
 - Coupling Express LR is recommended for Long Distance Coupling z13/z13s to z13/z13s and up
 - New coupling channel type: CL5
 - Performance is similar to Coupling over InfiniBand 1x
 - PCIe I/O drawer required for CL5 adapter – even for standalone CF usage
 - Feature (2-port card) with Coupling Optics and Firmware
 - 10 Gbps*, Up to 4 CHPIDs per port, 32 buffers (i.e. 32 subchannels) per CHPID
 - Distance: 10 km Unrepeated; up to 100 km with a qualified DWDM
 - More than 100 km requires RPQ 8P2781.
 - Maximum configuration supported:
 - 64 ports (32 features) for z14 M0x and z13 , 32 ports (16 features) for z14 ZR1 and z13s
 - Point-to-Point just like InfiniBand 1X and ISC-3;
 - CANNOT be utilized in a switched environment
 - Cabling: Utilizes same 9u, Single Mode fiber type as 1X IFB and ISC-3
 - z13 GA2+, z13s GA+ and z14 availability
 - No Going Away Signal for STP

*Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

z14 ZR1 Coupling Connectivity



z13 and z13s
ICA SR, CE LR

z14 ZR1
ICA SR, CE LR

Coupling Express LR (CE LR)
10 Gbps, 10/100 km

z13, z13s, z14 to z13/z13s/z14 Connectivity ONLY

Integrated Coupling Adapter (ICA SR)
8 Gbps, up to 150 m

z13, z13s, z14 to z13/z13s/z14 Connectivity ONLY



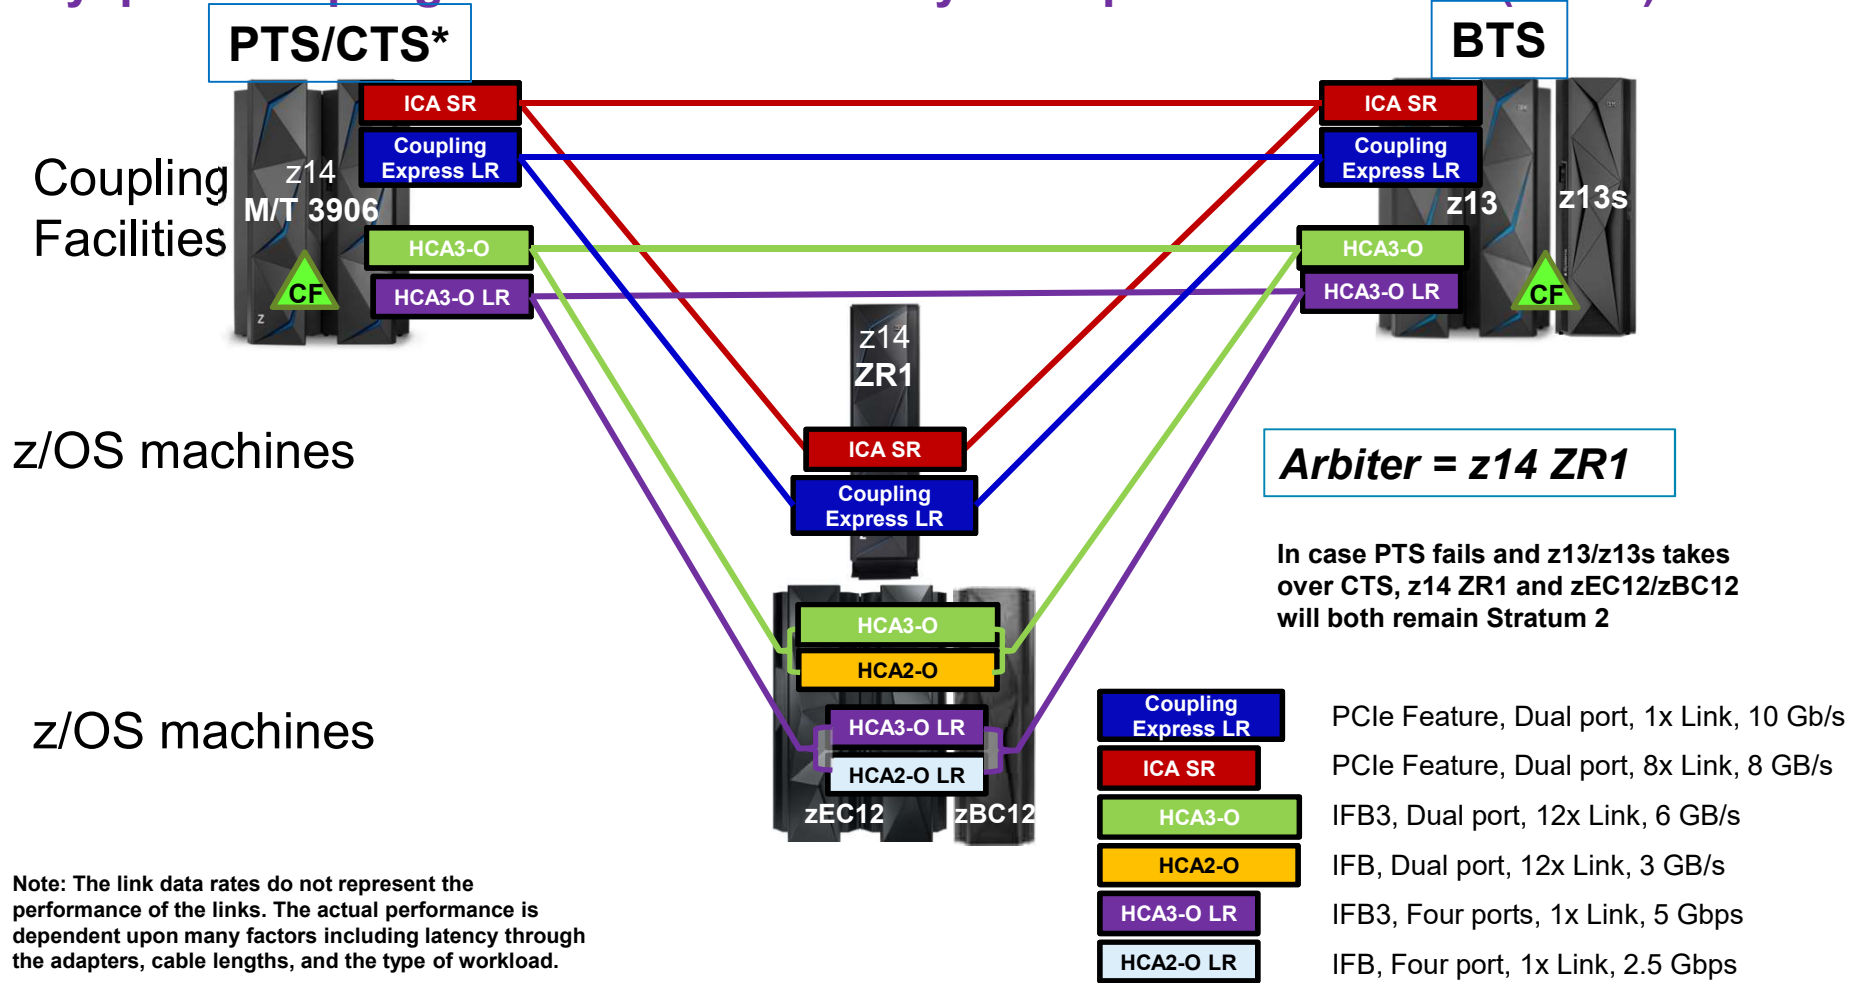
z196, z114, and older CPCs CANNOT coexist in the same Parallel Sysplex or STP CTN with z14 ZR1 (no coupling connectivity)
zEC12 or zBC12 can coexist in the same Parallel Sysplex with z14 ZR1 **only** if the CPC hosting the CFs has coupling connectivity to both the zEC12/zBC12 and z14 ZR1 CPCs

IC (Internal Coupling Link):
Only supports IC-to-IC connectivity

HCA2-O and HCA2-O LR and ISC-3 are NOT supported on z13, z13s and z14 M/T 3906
HCA3-O and HCA3-O LR are NOT supported on z14 ZR1 (M/T 3907)

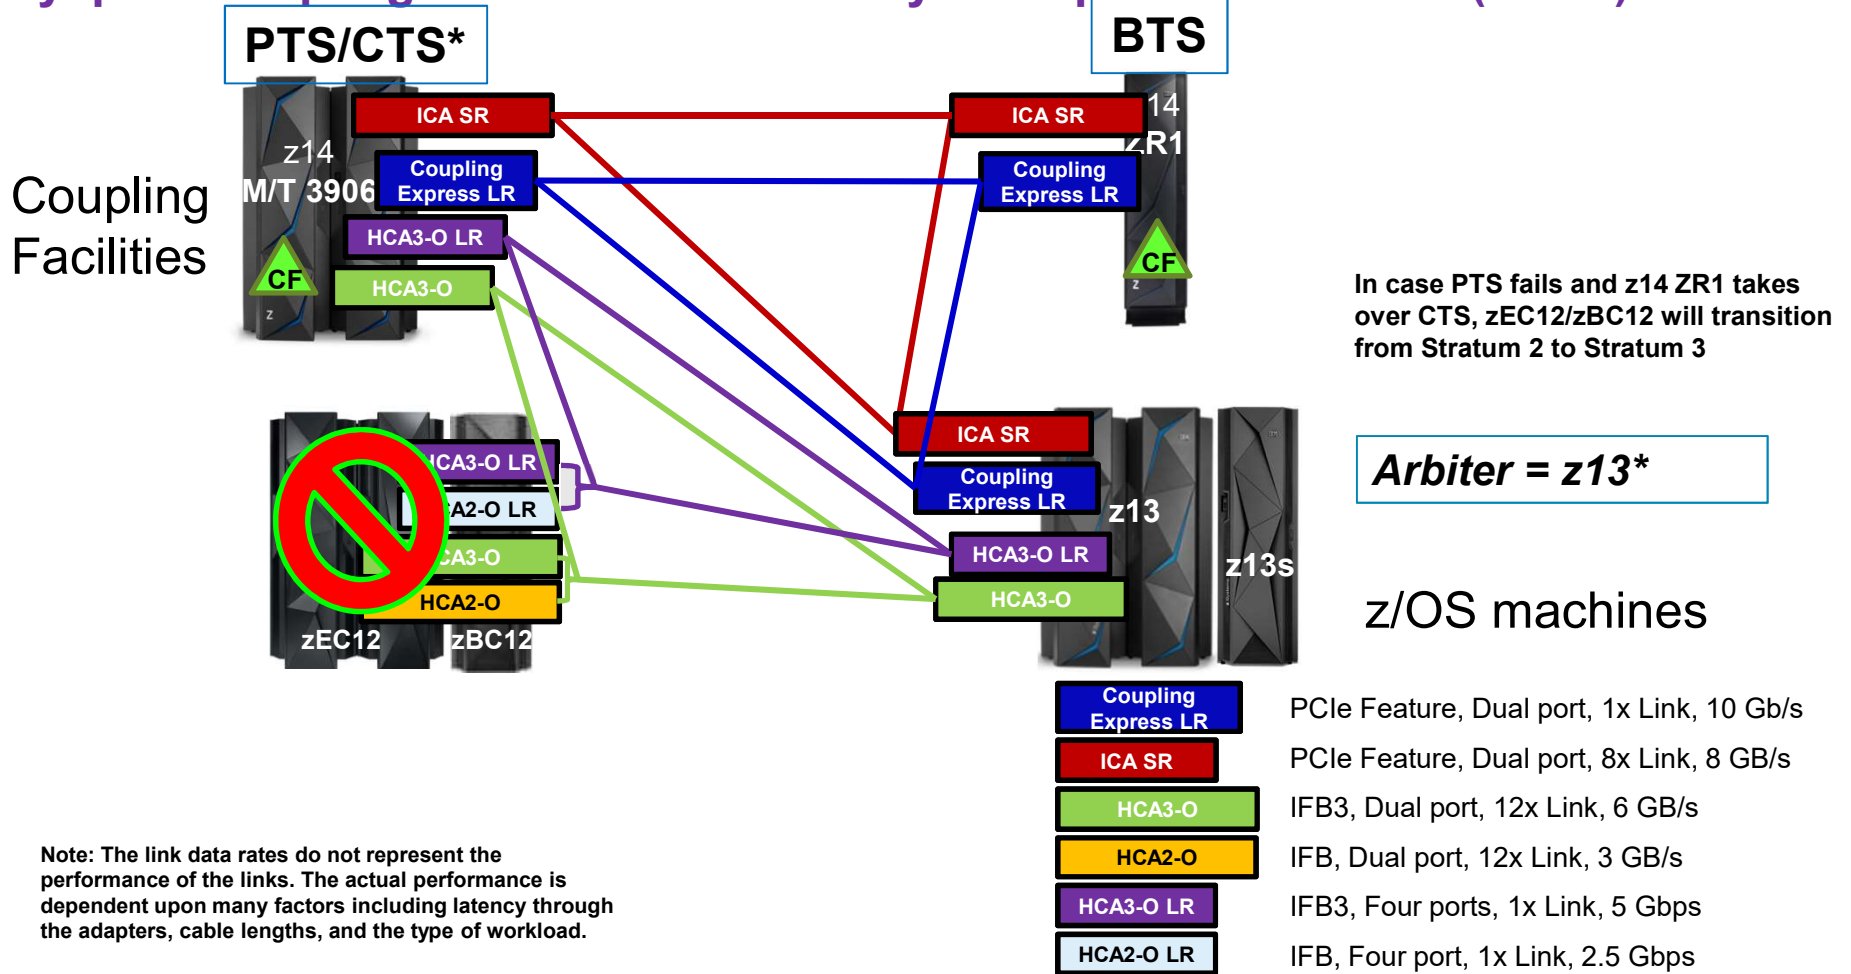
Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

Sysplex Coupling and STP Connectivity example with z14 ZR1 (1 of 3)

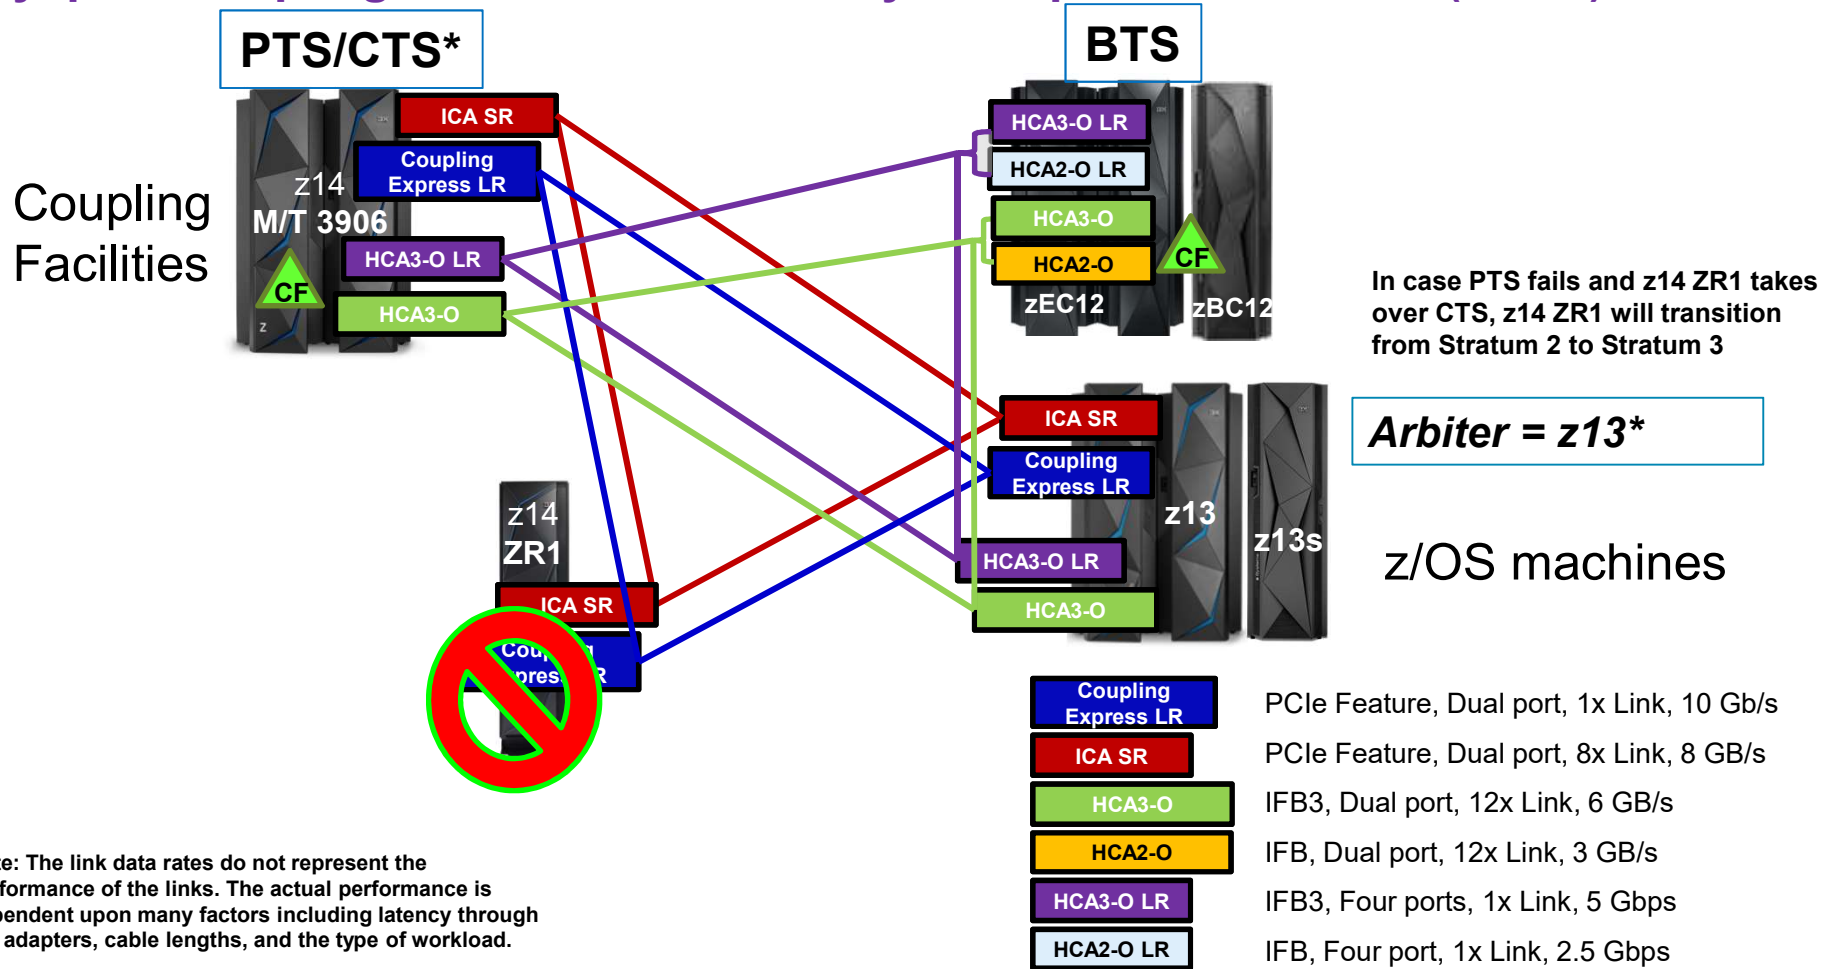


Note: The link data rates do not represent the performance of the links. The actual performance is dependent upon many factors including latency through the adapters, cable lengths, and the type of workload.

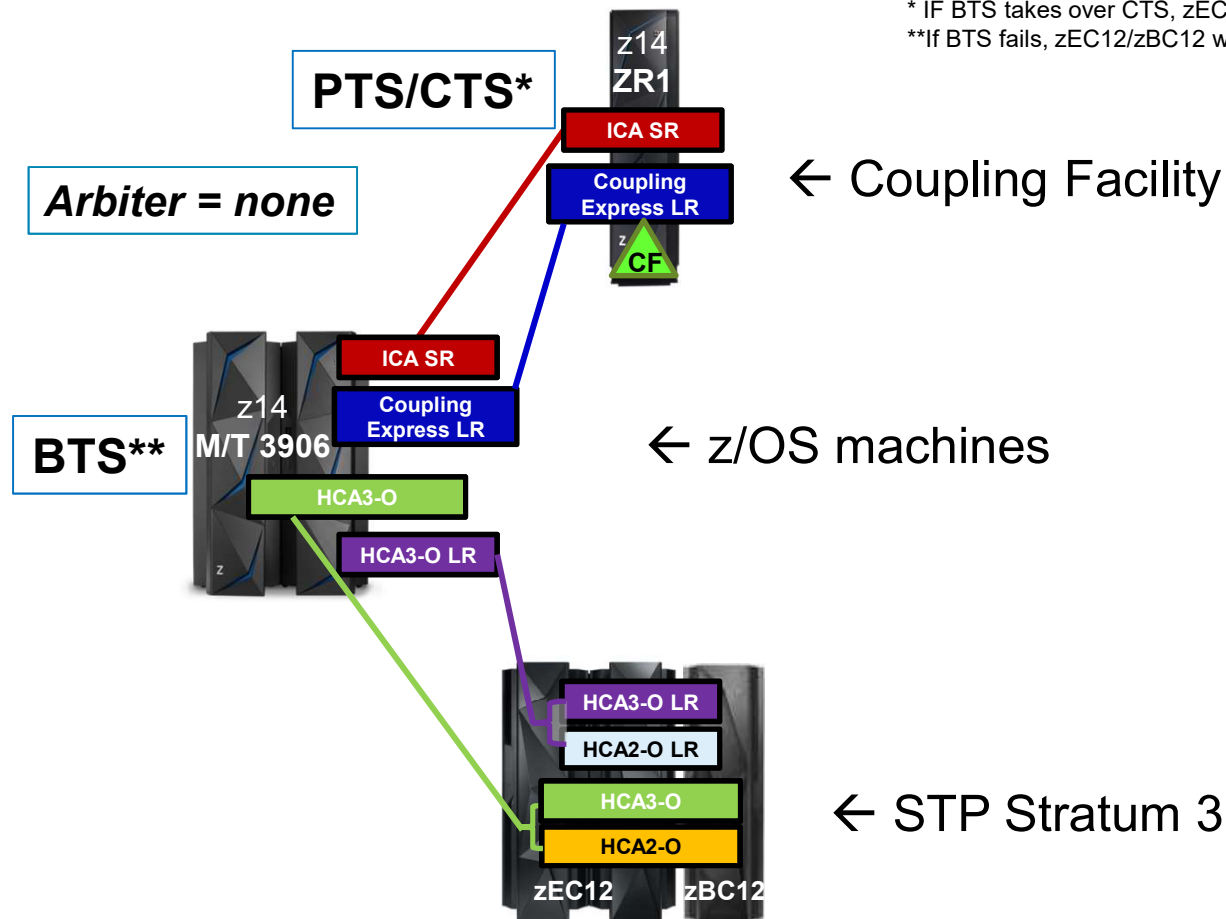
Sysplex Coupling and STP Connectivity example with z14 ZR1 (2 of 3)



Sysplex Coupling and STP Connectivity example with z14 ZR1 (3 of 3)



More STP Considerations



* IF BTS takes over CTS, zEC12/zBC12 will become Stratum 2
 **If BTS fails, zEC12/zBC12 will become Stratum 0 (unsynchronized)

Coupling Link Considerations for z14 ZR1

- z14 ZR1 supports only PCIe-based coupling (ICA SR and CE LR)
 - Maximum number of physical ICA SR coupling links (ports) is 16 per CPC
 - However, long distance coupling requires PCIe+ I/O Drawer for hosting CE LR features, thus reducing the max. number of ICA SR features.
- While physical coupling links can be (and most often are) shared across images and sysplexes, there are some customers who configure dedicated physical CF links
 - “Stacked” consolidated sysplexes, on a set of physical CECs, with dedicated connectivity provided for each sysplex
 - These kinds of dedicated/isolated configurations drive requirements for higher maximum limits on both physical links and logical CHPIDs
- CPCs that host standalone CFs tend to have the highest per-CPC consumption of coupling connectivity resources
 - CFs are “focal points” for both z/OS-to-CF and CF-to-CF sysplex connectivity, as well as STP timing roles
- z14 M/T 3906 is the last machine supporting the Infiniband coupling. Additional coupling link configuration complexity is expected for transitioning to ICA SR and CE LR coupling.
- Customer benefits:
 - Customers will have sufficient physical coupling link connectivity they need to build the kind of configurations they require
 - While meeting their requirements for sysplex scalability, dedication/isolation, robust redundancy and availability, technology transition/migration, performance and bandwidth, capacity and throughput etc.,

INTERMISSION

- <https://youtu.be/d0-pLcgq-2M>
- <https://youtu.be/ewsaYde59C0>

STP Enhancements for z14 ZR1

See also previous section (Coupling Connectivity)

Additional Stratum Level – STP Stratum 4

- What is it?
 - Allows CPCs to operate as part of a CTN at STP Stratum level 4
 - 3 steps away from the Current Time Server (CTS).
 - Previously the limit was Stratum level 3
 - 2 steps away from the Current Time Server (CTS).
- Why ?
 - During reconfiguration operations, customers may be rolling systems.
 - When they remove systems that have critical roles, it exposes some other systems to becoming unsynchronized.
 - Today customers incur the additional expense and difficulty of leasing or purchasing extra coupling adapters, laying extra fiber, and going through multiple reconfigurations in order to avoid this problem
- Benefits
 - Avoids the additional complexity and expense of system reconfiguration
- Migration considerations (Compatibility Issues / Requirements)
 - All systems that may become exposed to this situation, should have this change installed

WARNING!

- This additional stratum level should only be used as a temporary state during reconfiguration.
- Customers should not run with machines at Stratum level 4 for extended periods of time because of the lower quality of the time synchronization.

Graphical representation of the CTN

- What is it?
 - Improved user interface to STP controls.
- Why? - Customer demand
 - Customer took a complete outage while performing system maintenance
 - They intentionally took down the Current Time Server, not realizing the backup time server was already down.
 - The customer believes that having a visual display of the CTN status would have avoided this outage
- Benefits
 - It provides a clearer view of the CTN, and easier to perform operations
 - Simple to understand
 - Much more intuitive
- Migration considerations* (Compatibility Issues / Requirements)
 - Changes to HMC supports previous CPCs as well

* See Speaker notes

Manage System Time menu →



Home Manage System Time - MU... [X]

Time: 11:41:39 AM Date: 3/7/18 Time Zone: UTC-5:00 Current time details Status: ✔

100%

CTN ID CTLSOMRS [edit]

Membership Unrestricted CTN

STRATUM 1

CETUS PTS CTS

STRATUM 2

LEPUS BTS

MUSCA

STP ACTIONS

- [Add systems to CTN](#)
- [Configure External Time Source](#)
- [Deconfigure CTN](#)
- [Export CTN data \(.xls\)](#)
- [Modify assigned server roles](#)
- [Remove systems from CTN](#)
- [Set CTN member restriction](#)
- [Setup new CTN](#)

Diagnostic Actions [dropdown]

EXTERNAL TIME SOURCE

NTP

192.168.70.46

NTP

192.168.70.65

STRATUM 1

PTS CTS CETUS

STRATUM 2

BTS

LEPUS

MUSCA

STRATUM 3


z14 ZR1 Cryptography

Note: All functions require ICSF WD17 (HCR77C1) for exploitation unless otherwise noted.

WD17 supports z/OS V2R1, V2R2, and V2R3.

Overview


- General availability in z14
 - Based on IBM PCIe Cryptographic Coprocessor (PCIeCC), available as machine type model 4768
- Enhanced Card Performance
 - Main PPC 476 (Qty 2 in lock-step) @ 1.2Ghz (vs. Crypto Express5S @ 800MHz)
 - PCI Express Gen 2 (vs. Crypto Express5S @ Gen 1)
 - DDR3 1600
 - Persistent Memory Management for faster boot time (FPGA)
 - New Miniboot Implementation for certification compliance
 - Allows easier transition of algorithms when certification requirements change
- Enhanced Public Key Cryptography Algorithms performance
- Upgraded secure module tamper detection technology
 - Improved thermal capabilities for increased performance
 - Continued support for temperature and voltage detection



Security Certifications

- Physical Security Standards in progress/planned:
 - ✓ FIPS 140-2 level 4
 - ✓ Common Criteria EP11 EAL4
 - ✓ Payment Card Industry (PCI) HSM
 - ✓ German Banking Industry Commission (GBIC, formerly DK)

Note: PCI-HSM certification is new for Crypto Express6S. The others also apply to Crypto Express5S.



SHA-3, SHAKE algorithms

- SHA-3 was standardized by NIST in 2015 (FIPS Pub. 202) as an alternative to SHA-2
- SHA-2 is still acceptable and there is no indication that SHA-2 is vulnerable or that SHA-3 is more or less vulnerable than SHA-2.
- CPACF is introducing support for the four SHA-3 hashing algorithms
 - SHA3-224
 - SHA3-256
 - SHA3-384
 - SHA3-512
- And the two extendable output functions as described by the standard
 - SHAKE-128
 - SHAKE-256
- The SHAKE variants are known as an extendable output functions and allow generating hash values of user-specified lengths. Note however that they are not approved as hash functions.
- These algorithms are exposed to applications via the ICSF One-Way Hash Generate callable service (CSNBOWH)

True Random Number Generation (TRNG)

- CPACF is adding support for TRNG
- TRNG is an improvement over Deterministic RNG in the sense that the numbers generated are more random. However, this comes with a performance penalty.
- ICSF will use the best of both worlds and take a hybrid approach to random number generation. TRNG will be used to seed a DRNG which will then be used to generate random numbers.
- Random numbers may be obtained/used via the following ICSF callable services:
 - PKCS #11 Pseudo-Random Function (CSFPPRF)
 - Random Number Generate (CSNBRNG, CSNBRNGL)
 - PKCS#11 services when generating a clear key


z14 ZR1 CPACF Performance

CPACF encryption rates for like modes and data sizes on z14 ZR1 are up to six times (6x) faster than z13s. This is based on preliminary internal IBM lab measurements on a standalone dedicated system in a controlled environment and compared to the z13s. Results may vary.

HMC Highlights

New SE console (KVM switch + tray + KMM)

Section Content

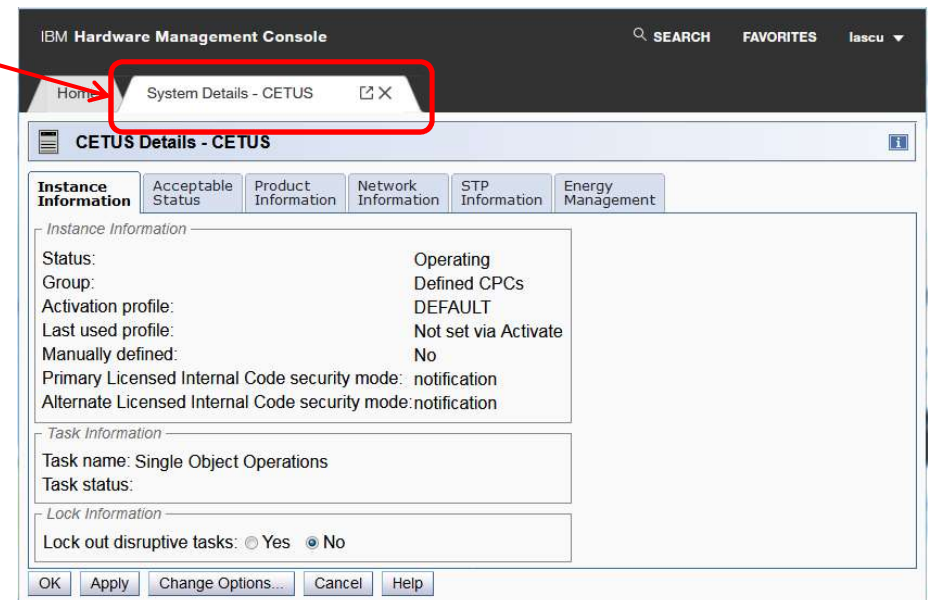
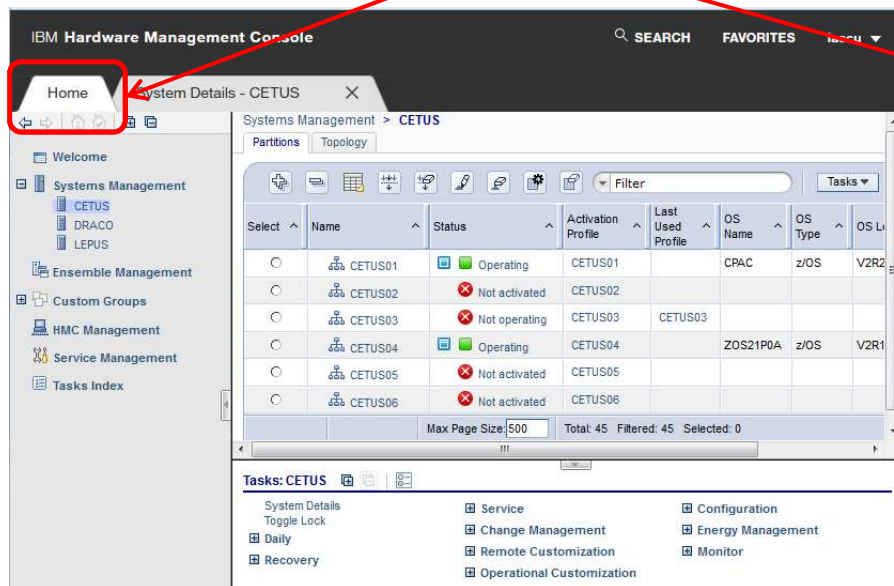


Section Content

- Workspace Direction/Enhancements – No more Classic UI
- Enhanced security for z14 HMC and SEs
- Java Applet Removal for z14
- Firmware management enhancements
- IBM Enhanced Support Facility – Changes to Call-Home infrastructure
- Crypto PCI-HSM Compliance Level
- Global OSA/SF
- LPAR Resource Assignment
- New characters allowed for load parameters
- FTP Through HMC
- Secure Console-to-Console Communications
- SNMP/BCPii API Enhancements
- Remote Browser IP Address Limiting
- Manage System Time (STP interface enhancements)
- Multi-factor Authentication (MFA for HMC authentication)

Workspace Enhancements

- **Classic UI Style NO longer available on HMC/SE 2.14.0 and later**
 - z13/z13s Statement of Direction - **Fulfilled**
- **Workspace Enhancements to aid in transition to Tree Style UI**
 - Changes to address user concerns
 - A **new masthead** for the HMC and SE will help users quickly find and launch tasks
 - Tasks will now open in **tabs** within the user interface (instead of separate browser windows) to make finding and managing running tasks easier.



Firmware Integrity Monitoring* - Overview

- Firmware Integrity Monitoring in support of NIST Computer Security Standard 800-147B
- Current IBM Z Firmware protected on delivery => Digital Signatures
- Firmware Integrity Monitoring* => protection against subsequent tampering
 - BIOS Secure boot of SE/HMC/TKE Embedded Operating System
 - Signature/hash verification of SE/HMC IBM Z Firmware
 - Continuously monitors integrity of files
 - Code measurements stored in TPM (Trusted Platform Module) on SE/HMC
 - Provides security logs for internal analysis
 - Trusted 3rd Party Validation (IBM Resource Link via zRSF data)
 - Analyzes periodic call home measurement data
 - Initiates challenge/response to verify authenticity of the data
 - Display of local console data analysis, Resource Link analysis, and notification of lack of receiving console data (console locked or blocked network reporting of data).

[*See Speaker Notes](#)



Firmware Integrity Monitoring - Customer Choices

- Modes of Operation => New HMC/SE task configuration
 - Monitoring: Logging/Hardware Message, mirror/backup block
 - Monitoring and Protecting: Logging/Hardware Message & 'Blue Screen' console lockup
- SE Firmware Integrity Monitoring is enabled on all z14 ZR1 systems
 - Sufficient to comply with the NIST Standard 800-147B for z14 ZR1 CPC firmware
- HMC Firmware Integrity Monitoring
 - Not required for z14 ZR1 CPC NIST 800-147B compliance but same technology used as SE
 - Will be enabled on new z14 ZR1 HMC Feature Codes (new build hardware only)
 - Provides tamper protection monitoring and reporting
 - MES HMC hardware to z14 ZR1 HMC level cannot utilize HMC Firmware Integrity Monitoring

z14 Java Applet Removal

- 2.13.1: HMC Tasks no longer Java Applet based implementations
 - Operating System Messages
 - Integrated 3270 Console
 - Integrated ASCII Console
 - Text Console

- 2.14.0: NEW: Open IOCDS Source option on Input/Output Configuration task
 - With HMC/SE 2.14.0 can now edit IOCDS source directly on HMC console
 - Alternative with remote browsing still available:
 - Use Export Source File option on Input/Output Configuration task
 - Use editor on your own workstation
 - Use Import Source file option to put back onto Support Element

Firmware Management Enhancements

EDM (Enhanced Driver Maintenance) / CDU (Concurrent Driver Upgrade) to bundle boundary

- Prior to HMC/SE 2.14.0,
 - All available licensed internal code change fixes (MCLs) which are available in IBM Support System will be applied after a EDM/CDU if the zRSF (IBM Z Remote Support Facility) is used for fix retrieval
 - This could be limited to certain fix levels by using a MCL (SUL) media for fix retrieval
- Starting with HMC/SE 2.14.0,
 - Customers will be able to specify the target bundle they want their GA2 machine to go to.
 - IBM Z Service Representatives will be able to specify this target bundle on an HMC panel
 - the CDU will wind up at that GA2 bundle at the completion of the CDU.

IBM Enhanced Support Facility – Overview*

- First introduced with the zEC12/zBC12
- Provides enhanced call home capabilities for IBM Z
 - Improved problem data uploads
 - Improved fix data downloads
- For zEC12/zBC12 and z13 systems, call homes to both the IBM Enhanced Support Facility, or to the legacy IBM support facility were allowed
- z14 M0x and z14 ZR1 processors will no longer connect to the legacy IBM Support facility
 - z14 HMCs supporting older z10EC/z10BC - z13 systems will still require connections to the legacy IBM Support Facility
 - z14 HMCs supporting only z14 systems only require IBM Enhanced Support Facility connections
 - A connection to the IBM Enhanced Support Facility is required to support z14 HMCs regardless of the CPC they are supporting.

[*See Speaker Notes](#)

IBM Enhanced Support Facility – Changes

- The IBM Enhanced support facility is deployed redundantly in 3 different locations.
- **Customer firewall should be updated** to allow access to each of the locations
- Firewall rules to allow connection the legacy IBM Support facility should remain as long as the customer has pre-z14 machines
- Both the IBM Enhanced support facility and the legacy IBM support facility support IPV4 or IPV6 connections.
- If the customer uses a proxy server, then the proxy server configuration may need to be updated

IP addresses for call home

- IBM Enhanced support facility
 - IPV4
 - 129.42.56.189
 - 129.42.60.189
 - 129.42.54.189
 - IPV6
 - 2620:0:6c0:200:129:42:56:189
 - 2620:0:6c2:200:129:42:60:189
 - 2620:0:6c4:200:129:42:54:189
- Legacy IBM support Facility
 - IPV4
 - 129.42.26.224
 - 129.42.42.224
 - 129.42.50.224
 - IPV6
 - 2620:0:6c0:1::1000
 - 2620:0:6c2:1::1000
 - 2620:0:6c4:1::1000

IBM Enhanced Support Facility - Proxy server configuration

- Proxy configuration is done via the Customized Outbound Connectivity task
- Configuration allows for either HMC or Proxy to resolve DNS addresses
- If HMC does DNS lookup, then proxy must accept IP addresses shown on prior page.
- If HMC passes DNS address to proxy (recommended), then proxy needs to support connections to the following host names:
 - IBM Enhanced Support Facility : esupport.ibm.com
 - Legacy IBM Support Facility: www-945.ibm.com

LPAR Resource Assignment – Background

- It has become increasingly important to know which physical processors are allocated for logical partitions.
- With today's machines, software running in a logical partition can obtain the partition's CPU topology with the STSI instruction, but not of the entire system
- Today the only way to get a complete picture of all the logical partition topologies is to take a non-disruptive LPAR dump and have development run formatting execs against the extracted data.

LPAR Resource Assignment Information – What is it?

- A file can now be generated containing the current active logical partitions and associated processor information (i.e. mapping) along with a number of previous mappings
- This file can be generated and sent to IBM using the Transmit Service Data task
- BCPii API support is provided to obtain this data
- Web Services APIs are not supported
- If necessary, a Customer Engineer may be requested to manually generate this file; generation of this file is non-disruptive
- The information contained in this file can then be used for diagnostic purposes

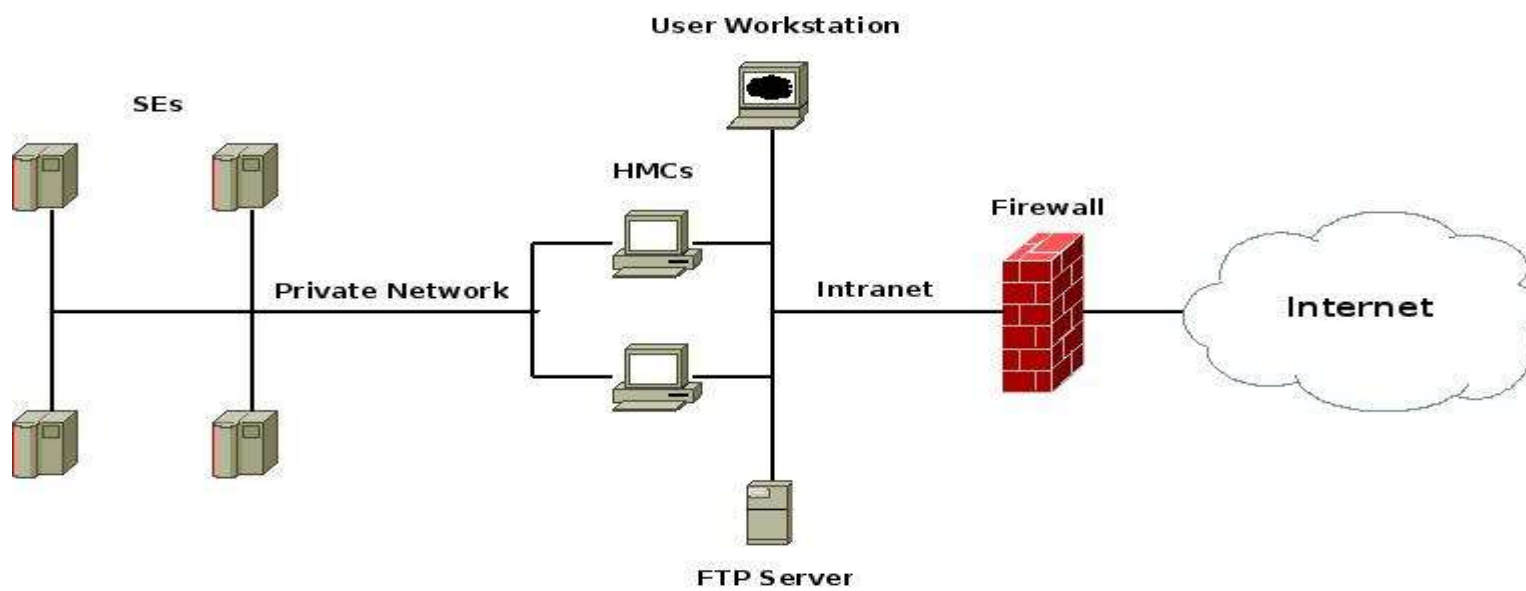
New characters allowed for load parameters

- Load parameter for load task, load profile and image profile now allows three new characters: @#\$
 - Old character set permitted: (z13 and older)
 - A-Z, 0-9, "." (period) and " " (blanks).
 - New character set permitted: (Starting with z14)
 - A-Z, 0-9, "." (period), " " (blanks), "@", "#" and "\$"
- New characters supported with BCPii, SNMP, and web services APIs

FTP Through HMC

- For maximum security, we recommend that customers keep their IBM Z on an isolated network.
- This prevents FTP operations originating from the SE from working.
 - Customer has to either put their FTP server on the IBM Z network, or (more likely) put their IBM Z on their intranet network.
- Starting in z14, all FTP operations originating from the SE will be proxied through a managing HMC.
 - The HMC performs the FTP operation on the SE's behalf and returns the results.
 - Secure FTP server credentials will need to be imported to one or more managing HMCs.
 - The IBM Z server will need to be managed by at least one HMC to allow FTP operations to work.

FTP Through HMC – Recommended Network Topology



FTP Through HMC– FTP Enhancements

- Adding support for FTP, FTPS and SFTP
 - FTPS is SSL based and uses certificates to authenticate servers.
 - SFTP is SSH based and uses SSH keys to authenticate servers.
 - Username and passwords are required for client authentication in all 3 protocols.
 - Key/Certificate Management tasks are already in place on the HMC.
- All three protocols are supported when proxy-ing through an HMC.
- All tasks that previously supported FTP will now support all 3 protocols.
- Many tasks are being updated with FTP support where previously only removable media was an option.

- FTP server requirements
 - Must support 'Passive' data connections
 - A server configuration allowing the client to connect on an ephemeral port
- FTPS server requirements
 - Must operate in 'Explicit' mode
 - Allows a server to offer both secure and unsecured connections
 - Must support 'Passive' data connections
 - Must support secure data connections
- SFTP server requirements
 - Must support password based authentication

FTP Through HMC – List of updated tasks

- System Input/Output Configuration Analyzer
- Analyze Console Internal Code
- Change Console Internal Code
- Save Upgrade Data
- Configure Backup Settings/Backup Crit. Console Data
- Load from Removable Media or Server
- FCP Configuration
- OSA Advanced Facilities → Manual Config Options
- Export WWPNS
- Audit log Scheduled Operations
- Export/Import IOCDS
- Save/Restore Customizable Console Data
- Crypto Config
- View/Archive Security Logs

Secure Console-to-Console Communications

- Older IBM Z HMC consoles have been kept secure against industry reported vulnerabilities like POODLE, BarMitzvah and FREAK but use anonymous cipher suites.
- Network security scanners can detect this and many customers have requested a resolution.
- In z14 ZR1 we are moving away from anonymous cipher suites and will begin using an industry standard based password driven cryptography system.
- The existing Domain Security Settings will be used to provide authentication as well as high quality encryption.
- Because of these changes, we now recommend that customers use unique Domain Security settings to provide maximum security.
 - The new system will provide greater security than anonymous cipher suites, even if the default settings are used.

Secure Console-to-Console Communication Changes to Domain Security

- To allow greater flexibility in password selection we have increased the password limit to 64 characters.
 - Passwords will be case sensitive and allow some special characters.
- If communications with older systems is needed, the previous password limits must be followed.
 - Minimum of 6 characters
 - Maximum of 8 characters
 - Only Uppercase and numeric characters allowed
- Increased password limits will be limited to z14 only shops.
 - 8 randomly generated characters will provide sufficient security in this system.
 - The increased password limits allow greater flexibility with human generated password strategies.
 - Word based passphrases
 - Padded passwords/phrases

Secure Console-to-Console Communications Backwards Compatibility

- By default, a z14 HMC will be able to fall back to using SSL to communicate with older HMCs
 - As a consequence, network scanners will still flag the use of anonymous SSL ciphers.
- A setting will be provided in the “Customize Console Services” task to completely disable anonymous cipher suite support.
 - This will prevent z14 HMCs from talking to any older HMCs, but will also allow network scanners to run clean.
- The z14 HMC will always be able to communicate with older SEs since anonymous cipher suites were not used for those connections.
 - Again, communications with older systems will require a password 6-8 characters long.

SNMP/BCPii API Functional Enhancements

- Virtual Flash Memory
 - Query support for Image
 - Query/update support for Image profiles

- Secure Service Container
 - Query/update support for Image profiles

- Performance
 - Performance improvements for query operations
 - Profiles, LPAR Controls/Security data, EC/MCL data
 - Applies to SNMP requests to HMC or SE and BCPii requests

SNMP/BCPii API BCPii Security Enhancements – Security Controls

- Previous controls
 - “Cross Partition Authority” used to control partitions that can send BCPii requests
 - Limited to SNMP community name setup for receive controls
- New “BCPii Permission” control
 - Enable/disable send capability
 - Enable/disable receive capability
 - Enable for all senders
 - Enable for list of origin partitions
 - Can be specified for CPC via System Details task (HMC or SE)
 - Can be specified for Images via Image profile and Change LPAR Security task (HMC or SE)
- Upgrades of “Cross Partition Authority” setting
 - “Cross Partition Authority” setting remains the same
 - Enabled maps to “BCPii Permission” of send enabled, receive enabled for all origins
 - Disabled maps to “BCPii Permission” of send/received disabled

Remote Browser IP Address Limiting

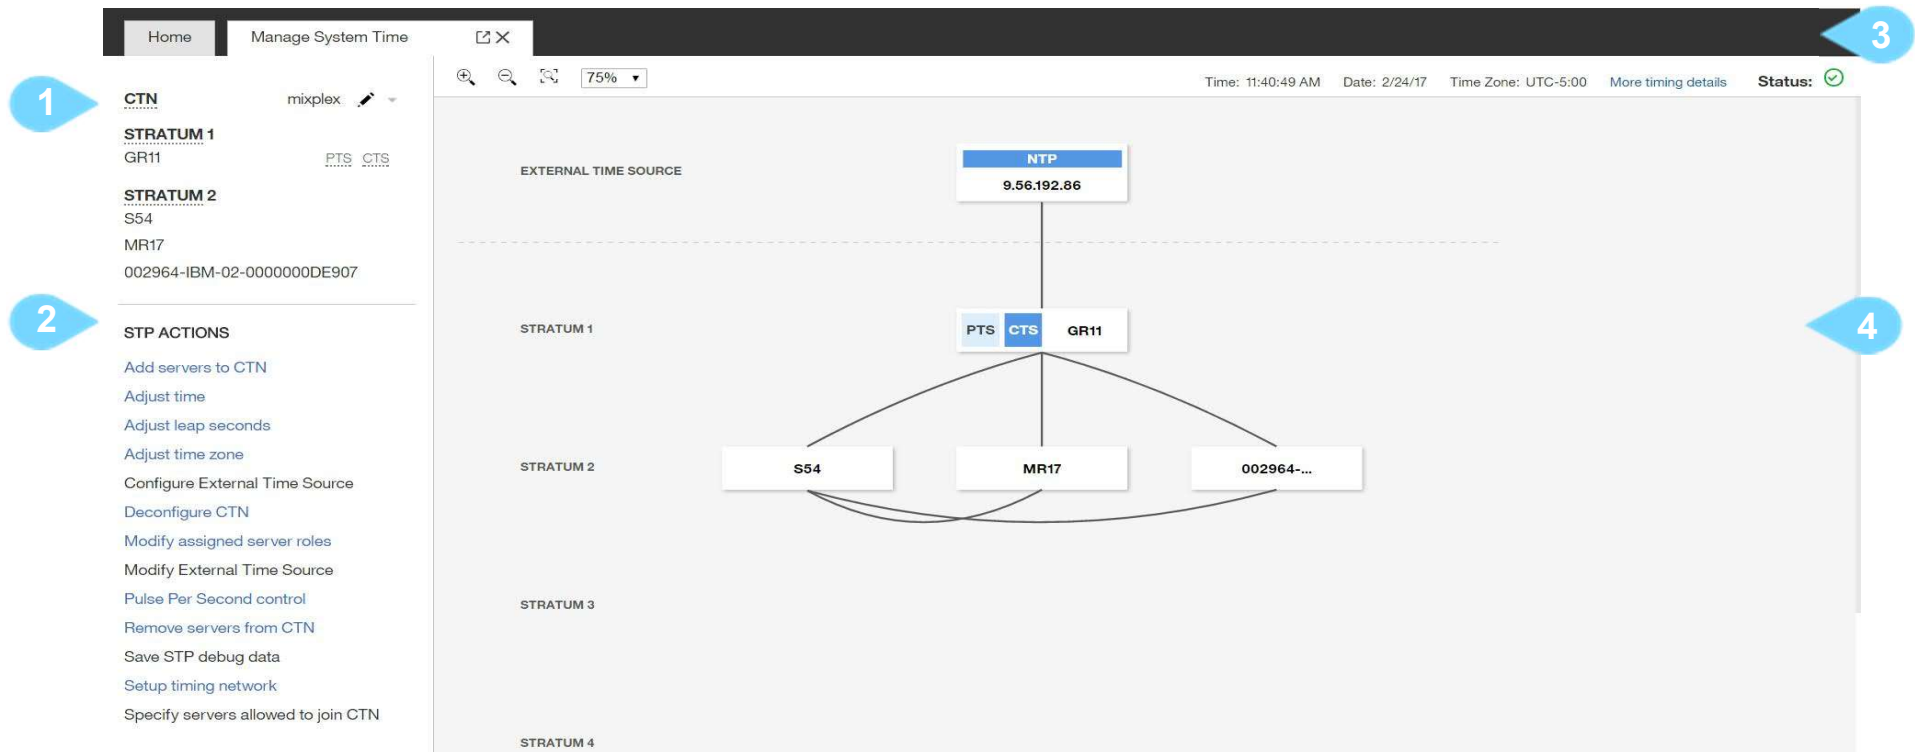
- Additional controls for remote browser access
- Global enable/disable for remote access still available
- Can limit the valid remote browser IP address
 - Can be expressed as:
 - Individual IP address
 - 1.2.3.4
 - IP address and mask for a groups of addresses
 - 1.2.3.0/24

Manage System Time - Highlights

- IBM Z administrator is guided through a system time management workflow reducing need to refer to external documentation.
 - Inline definition of technical terms eliminates need to look up documentation to find out definitions.
 - Detailed instructions and guidelines are provided within task workflow.
- IBM Z administrator gets to see visual representation
 - current system time networks shown in topological display.
 - preview of any reconfiguration action is shown in topological display.
- IBM Z administrator's confidence for system time management operations is enhanced.
 - Errors surfaced in the visualization. User can drill down to find out details.
 - Reconfiguration workflows clearly show potential pitfalls and errors before change is applied.
 - Workflows steps are design to prevent mistakes.
- System time management performed for all managed systems from a single task on HMC. Previously, this was done by launching separate task for each system.

Manage System Time – Main Task

1. Global CTN Details.
2. STP Actions – launches wizards.
3. Topology Toolbar with global CTN timing details.
4. CTN Topology.



z14 ZR1 HMC MFA – Overview

• What is Multi-Factor Authentication?

- Method of user authentication that requires multiple pieces of evidence

- Factor – component used to prove identity
 - Something you know
 - e.g., username, password, PIN
 - Something you have
 - e.g., badge, smartphone, keyfob, USB stick, key, bank card
 - Something you are (biometrics)
 - e.g., fingerprint, eye iris, voice, typing speed/pattern

- Online banking, ATMs, Google, Facebook, etc.

• HMC MFA:

- Optional
- Configurable on a per-user, per-template basis
- Local and remote GUI logon
- Web Services APIs
- Locally-authenticated users
- LDAP-authenticated users
- Pattern/Template users
- Standalone solution required
 - No network connectivity
 - Partitions not running (e.g., LDAP, RACF)
- User-supplied smartphone
- HMC, SE and TKE

z14 ZR1 HMC MFA – Two factors

- 1st factor: HMC/SE/TKE userid and password
- 2nd factor: smartphone that provides a “multi-factor authentication code”
 - TOTP – Time-based One-Time Password
 - Defined by RFC 6238
 - A cryptographic hash function that combines a secret key with the current time to generate a one-time-use password
 - Can be used only once
 - Changes every 30 seconds
 - HMC/SE/TKE accepts current, previous and next TOTP
 - Compatible with existing freely-available smartphone apps
 - e.g., Google Authenticator and FreeOTP

New Console for the Support Elements

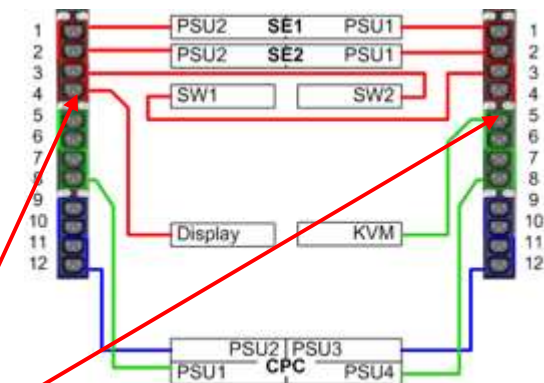
- The Support Elements console is now implemented in a 1U rack tray containing:
 - One Keyboard and pointing device
 - One flat panel display (with its power supply)
 - One Keyboard, Monitor, Mouse (KMM) switch that is connected to both Support Elements
- The KMM tray (installed at rack location A22B) slides out and provides access to the keyboard, pointing device and monitor.



KMM (front view)



Retracted KMM
(front view)



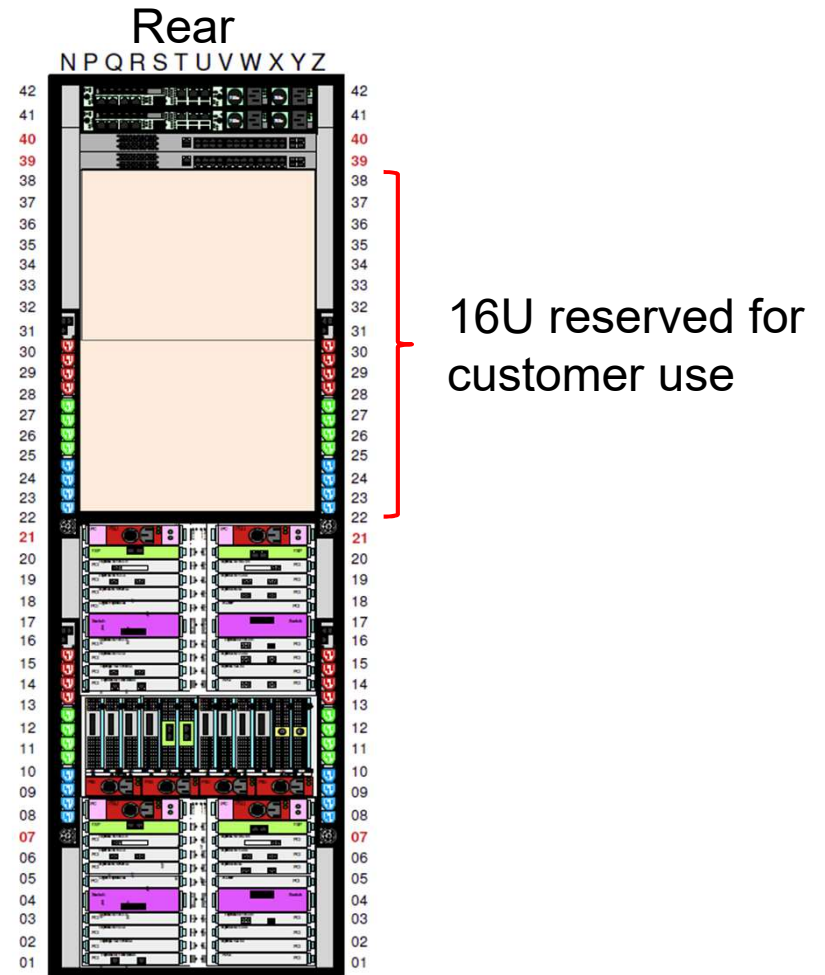
Rear view

Note: During PDU maintenance, the display or KMM switch power cables may need to be relocated to the PDU on the other side of the rack to ensure power availability.

16U Reserved Feature Detail

- Feature code is based on the ability to allocate (and use) contiguous 16U space within system configuration in lieu of 3rd and 4th PCIe+ I/O Drawers.
- The 16U Reserved Rack Space feature is documented in the M/T 3907 IMPP document.
- Some examples* of equipment that can be hosted in the space allocated through the 16U Reserved feature:
 - SAN switches
 - Network Switches
 - Rack mountable HMC and TKE
 - Storage devices
 - etc.,

* Note: All equipment installed in the space allocated for the 16U Reserved feature must comply with requirements described in the IBM 3907 Installation Manual for Physical Planning, GC28-6974



z14 ZR1 Power Infrastructure Design

- IBM z14 ZR1 power characteristics:
 - Packaged in a 42U 19” rack
 - Single Phase (200 – 240 VAC)
 - Intelligent Power Distribution Units (PDUs) – two or four PDUs per system, depending on configuration, switchable (Ethernet controlled)
 - Rating per power cord: 200 – 240 VAC, 30A
 - Redundant power supply units (PSUs) for:
 - Support Elements (1+1)
 - PCIe+ I/O Drawers (1+1)
 - CPC Drawer (two or four PSUs, redundant, configuration dependent)
 - Redundant Ethernet Switches
 - Each Ethernet switch has one built in power supply, the entire switch is a FRU
 - Max. supported rack power – 9600W
 - Rack Airflow – front to rear
 - All cabling in the rear of the rack
- IBM z14 ZR1 does **not** provide:
 - Bulk Power Regulator (BPR)
 - External EPO switch
 - Internal Battery Feature (IBF)
 - Three phase power
 - 480 VAC
 - HVDC (high voltage DC)
 - Balanced power
- Notes:
 - Power is also provided for hosting non-Z equipment in the z14 ZR1 rack (FC 0617)
 - Power specifications for “hosted” equipment are provided in the *IBM z14 ZR1 Installation manual for Physical Planning*, GC28-6974
 - When servicing equipment hosted in the z14 ZR1 rack an assessment must be carried out to avoid power outages to the other equipment installed in the same rack.

z14 ZR1 Software Support – z/VM

Section Content

- z14 ZR1 Support (M/T 3907)
- z14 ZR1 Compatibility Support
- z14 ZR1 Exploitation Support
- z/VM 6.4 Overview
- z/VM 6.4 2017 Enhancements
 - 1Q/2Q Security Enhancements
 - 3Q Additional Enhancements
 - 4Q Additional Enhancements

Section Content

- z14 ZR1 Support (M/T 3907)
- z14 ZR1 Compatibility Support
- z14 ZR1 Exploitation Support
- z/VM 6.4 Overview
- z/VM 6.4 2017 Enhancements
 - 1Q/2Q Security Enhancements
 - 3Q Additional Enhancements
 - 4Q Additional Enhancements

z/VM Release Status Summary



z/VM Level	GA	End of Service	End of Marketing	Minimum Processor Level	Maximum Processor Level	Security Level
7.1	Preview				--	--
6.4	Nov 11, 2016			IBM System z196 & z114®	-	
6.3	7/2013	12/2017 ^[1]	11/2016	IBM System z10®	-	EAL 4+ OSPP-LS
6.2	12/2011	07/2017 ^[2]	7/2013	IBM System z10®	z13	-
5.4	9/2008	12/2017 ^[3]	3/2012	IBM eServer zSeries 800& 900	zEC12	-

4 Releases in Service Dec 2016
1 Release in Service Jan 2018

Marketed & Serviced

Serviced, but not Marketed

End of Service & Marketing

^[1] Announced February 3, 2015
^[2] Announced February 2, 2016
^[3] Announced August 2, 2016

Updated: "Why Migrate?" presentation
<http://www.vm.ibm.com/devpages/bitner/presentations/gotovm64.pdf>

z/VM Support for z14 ZR1 (M/T 3907)

- Only z/VM 6.4 supports the z14 ZR1
 - z/VM 6.3 EoS is 12/31/2017 and is **not** supported on the z14 ZR1
- z14 PTFs must be applied before IPL-ing an existing z/VM 6.4 image on z14 ZR1; see <http://www.vm.ibm.com/service/redalert/#SAPLZ14>
 - VM65942 - z/VM support for z14 with driver D/T3906
 - VM65856 - ESA/390 Removal Support for z/VM V630
 - No additional PTFs are required beyond z14 support to IPL on a z14 ZR1
 - z/VM 6.4 will not IPL without the z14 compatibility support PTFs
- z/VM 6.4 can be installed directly on the z14 ZR1 with an image obtained from IBM after August 25, 2017
- The only additional PTFs which are required:
 - VM66015: VMHCD support
 - VM66017: VMHCD LinuxONE support
 - IOCP, EREP, VMHCM, and HLASM z14 support APARs provide z14 ZR1 support
- Updates will be available on the z/VM web page at announce: <http://www.ibm.com/vm>
- PSP bucket z/VM subset should also be reviewed; at announce
 - Upgrade 3907DEVICE
 - Subset 3907/ZVM

Section Content

- z14 ZR1 Support (M/T 3907)
- z14 ZR1 Compatibility Support
- z14 ZR1 Exploitation Support
- z/VM 6.4 Overview
- z/VM 6.4 2017 Enhancements
 - 1Q/2Q Security Enhancements
 - 3Q Additional Enhancements
 - 4Q Additional Enhancements

z/VM z14 Compatibility Support - APARs

- z14 compatibility APARs:
 - VM65942 and VM66071: z14 compatibility support / CEX6S crypto Adapter support
 - PI73016: TCP/IP Stack and NETSTAT support for OSA-Express6S
 - VM65952: EREP support
 - VM65865: Dynamic I/O support for CL5 CHPIDs
 - PI62275 and PI65715: HLASM support
 - VMHCD: VM65843 GA1; VM66009 LinuxONE; VM65849 Coupling Express
 - VMHCM: VM65266 GA1; VM65208 Coupling Express LR
 - IOCP: VM65939 GA1; VM65880 Coupling Express LR

- Additional Information available on the z/VM web page: <http://www.ibm.com/vm>


- Review hardware PSP buckets:
 - 3906DEVICE z/VM subset 3906/ZVM
 - 3907DEVICE z/VM subset 3907/ZVM

z/VM z14 Compatibility Support

- Crypto:
 - Crypto Express6S dedicated to or virtualized for guest exploitation; Both interfaces will include clear key ECC support; dedicated also allows secure key
 - CPACF enhancements (MSA6/SHA3, MSA7/True RNG, MSA8/GCM)
- SIMD and compression enhancements
- Miscellaneous-Instruction-Extensions Facility 2
- Vector Enhancements Facility 1 and Vector Packed Decimal Facility
- Memory management improvements: greater concurrency in address translation.
- NETSTAT OSAINFO support for OSA Express6S
- 10 GbE RoCE Express2 guest attachment (guest exploitation)
- Dynamic I/O support for zHyperLink (Dynamic I/O only), 10 GbE RoCE Express2, OSA Express6S, FICON Express 16S/16S+, Regional Crypto
- ESA/390 compatibility mode support, sufficient for CMS and GCS guests only
- CPUMF Counter enhancements “Diag x’204’ Busy indication facility”

SSI Cluster and z14 M0x and z14 ZR1 Machines

- APAR VM65976 is needed on ALL members (6.2, 6.3 and/or 6.4) of an SSI cluster prior to bringing up ANY member on a z14 machine.
- APAR VM65867 is required on ALL z/VM V6.2 and V6.3 systems in an SSI cluster that includes a z/VM V6.4 system.



z14 items z/VM will not support, including Guest Exploitation

- z/VM does not currently have plans to support the following z14 functions:
 - Time Slice Instrumentation
 - TOD Multiple Epochs
 - zHyperLink I/O (Dynamic I/O support only; No host nor guest exploitation)
 - Virtual Flash Memory (zVFM) (Flash Express is not supported by z/VM)

- Still no z/VM support for the following:
 - Runtime Instrumentation (RI)
 - zHPF Extended distance II
 - Enhanced DAT Facility 2
 - CHSC & I/O: Guest Support for Multiple-Subchannel-Sets
 - Flash Express Adapter

z/VM Statement of Directions with the z14 announce

- **Stabilization of z/VM V6.3 support**

The z14 is planned to be the last IBM Z high-end server and z13s is planned to be the last mid-range IBM Z server supported by z/VM V6.3 and the last IBM Z server that will be supported when z/VM V6.3 is running as a guest (second level).

***** Also applicable to the LinuxONE servers *****

- **Future z/VM release guest support**

z/VM V6.4 will be the last z/VM release supported as a guest of z/VM V6.2 or older releases.

- **Disk-only support for z/VM dumps**

z/VM V6.4 will be the last z/VM release to support tape as a media option for stand-alone, hardabend, and snap dumps. Subsequent z/VM releases will support dump to ECKD DASD or FCP SCSI disks only.

Section Content

- z14 ZR1 Support (M/T 3907)
- z14 ZR1 Compatibility Support
- **z14 ZR1 Exploitation Support**
- z/VM 6.4 Overview
- z/VM 6.4 2017 Enhancements
 - 1Q/2Q Security Enhancements
 - 3Q Additional Enhancements
 - 4Q Additional Enhancements

z/VM 6.4 z14 Exploitation (Available as of December 15, 2017)

- APAR VM65986: Guest exploitation for the Instruction Execution Protection Facility: z14 functionality to improve the security of programs running on IBM Z by allowing virtual memory elements to be identified as containing only data. If an attempt is made to fetch an instruction from an address in such an element, or if an address in such an element is the target of an execute-type instruction a Protection Exception will occur.
- APAR VM65987: Guest exploitation for Pause-Less Garbage Collection: z14 functionality designed to improve the performance of garbage-collection processing by various languages, in particular Java. This facility can be used to detect when application threads make reference to potentially stale data, allowing invocation of a handler to resolve and update such references. As a result, the pause time associated with these phases will be significantly reduced, leading to more consistent application response time.
- APAR VM65993: Encrypted paging, in support of pervasive encryption: Ciphering will occur as data moves between active memory and a paging volume owned by CP. Can be enabled and disabled dynamically.

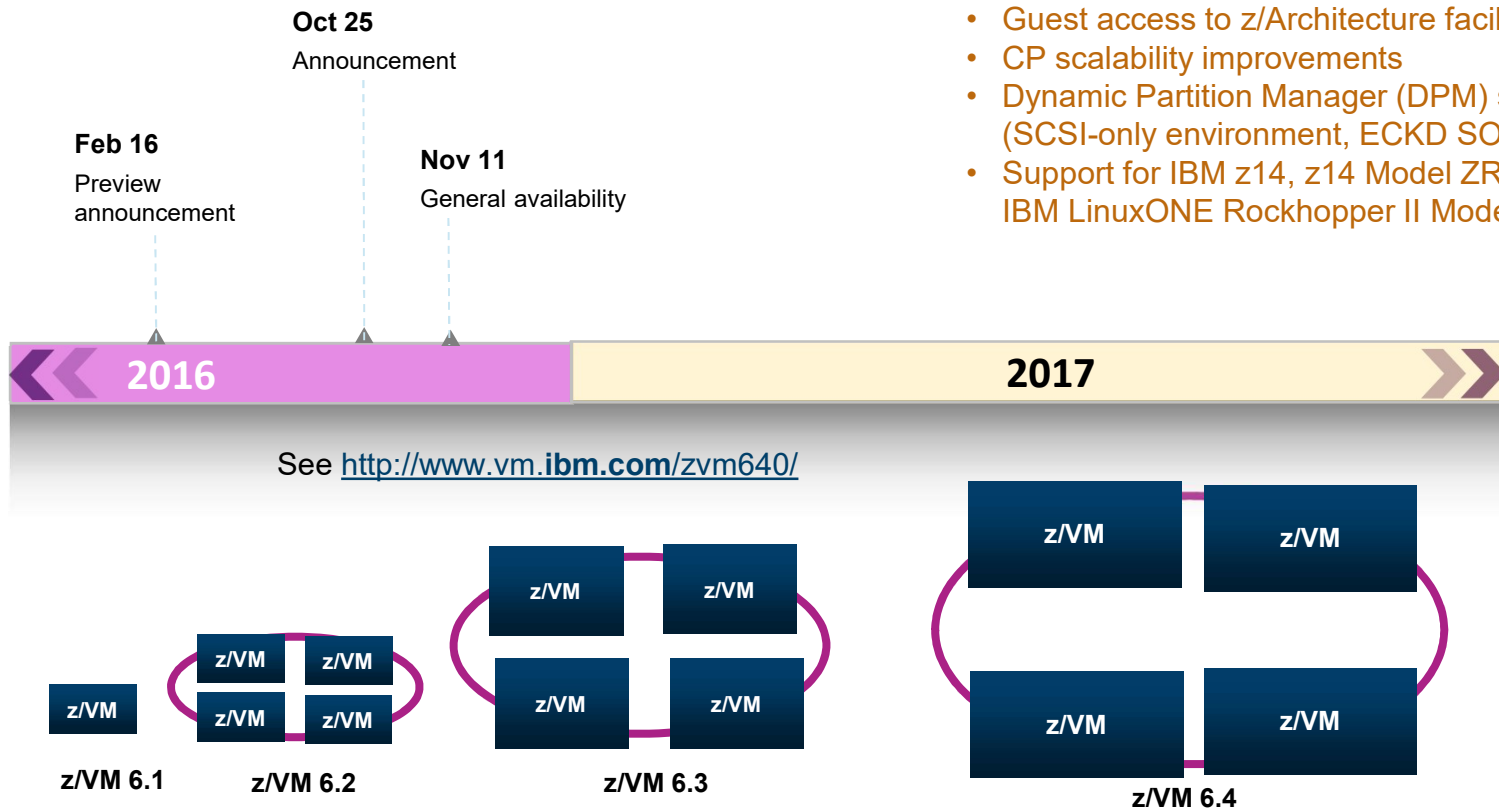
Section Content

- z14 ZR1 Support (M/T 3907)
- z14 ZR1 Compatibility Support
- z14 ZR1 Exploitation Support
- **z/VM 6.4 Overview**
- z/VM 6.4 2017 Enhancements
 - 1Q/2Q Security Enhancements
 - 3Q Additional Enhancements
 - 4Q Additional Enhancements

z/VM Version 6 Release 4 - Highlights

Designed for Clients of Today and Tomorrow

- Increased central storage 2 TB
- Dynamic SMT
- HyperPAV for paging
- Guest access to z/Architecture facilities
- CP scalability improvements
- Dynamic Partition Manager (DPM) support (SCSI-only environment, ECKD SOD)
- Support for IBM z14, z14 Model ZR1, IBM LinuxONE Rockhopper II Model LR1



z/VM Continuous Delivery

- IBM has adopted a continuous delivery (CD) practice for z/VM
- To deliver new functions via PTFs to assist enterprises in receiving function more quickly
- IBM intends to deliver new functions via PTFs continuously, while continuing to recognize base stability as a critical requirement
- Clients will receive the same world-class support and assistance from IBM to which they are accustomed, as part of the standard z/VM Software Subscription and Support (S&S)
- Sponsor users will guide design to ensure support addresses client requirements
- To learn about new enhancements going out through continuous delivery and the proposed schedules see the Continuous Delivery page:
 - <http://www.vm.ibm.com/newfunction/>
- To be notified when this function becomes available see the New Function APARs page:
 - <http://www.vm.ibm.com/service/vmnmfapar.html>

Sub-capacity pricing terms for z/VM and select z/VM-based programs

- Sub-capacity pricing for the z/VM V6 operating environment is available to clients running z/VM Version 6 Release 3 or higher. Software pricing at less than full machine capacity can provide more flexibility and improved cost of computing as a client manages the volatility and growth of new workloads.
- Through the implementation of sub-capacity pricing for select z/VM programs, clients can:
 - Pay for z/VM programs based on defined workload requirements and not necessarily the full engine capacity of the machine
 - Add hardware capacity for new workloads, capabilities, and functions, such as KVM, and not have IBM software pricing for z/VM programs automatically increase for existing workloads
 - Buy new hardware capacity for future growth with no immediate increase to their IBM software bill provided that no additional software capacity is required at that time
 - Potentially benefit from improved price/performance as workloads grow
- ILMT and instructions for its ordering and installation and use are available at: <https://www-01.ibm.com/software/passportadvantage/ibmlicensemetrictool.html>
- Installation instructions for the z/VM Hypervisor Proxy are available at: <https://www-03.ibm.com/systems/z/resources/swprice/subcap/zvm.html>

z/VM Cloud Enablement

z/VM is moving from the Cloud Management Appliance (CMA) to a new set of Cloud Enablement APIs. These APIs will allow customers to enjoy a broader set of features in their vertically integrated Cloud solutions from IBM or partners.

- Switch from an IBM-provided OpenStack and xCAT solution (CMA) to supporting partner-provided cloud solutions via new z/VM Cloud Enablement APIs
- The Cloud Enablement APIs will allow third-party OpenStack or other Infrastructure as a Service (IaaS) solutions and may run in the same Linux guest, or a different one, from those solutions
- The code is being developed in open source and is enabled for third-party contributions, see <https://github.com/mfcloud/python-zvm-sdk>
- The target availability is 1H2018
- As of Oct. 10, 2017, the CMA DVD and electronic download are no longer provided when ordering z/VM

Section Content

- z14 ZR1 Support (M/T 3907)
- z14 ZR1 Compatibility Support
- z14 ZR1 Exploitation Support
- z/VM 6.4 Overview
- **z/VM 6.4 2017 Enhancements**
 - 1Q/2Q Security Enhancements
 - 3Q Additional Enhancements
 - 4Q Additional Enhancements

1Q17 / 2Q17 z/VM 6.4 SPEs

- APAR VM65930, VM65982: RACF Security Policy Enhancements – Functional RACF enhancements providing improved security and usability for security administrators and auditors
- APAR PI72106: Crypto Express APVIRT for TLS/SSL Server – Enables connectivity from the TLS/SSL Server to crypto adapters for improved performance and reduced CPU overhead
- APAR VM65929: SCSI Container XIV Support – Enables multiple I/O requests to be issued concurrently to EDEVICES using XIV, which may improve performance. Particularly benefits EDEVICED paging I/O or volumes containing multiple minidisks
- APAR VM65872: SSI Distributed IUCV – SSI configuration and administration improvements; remove restrictions on distributed IUCV connections in an SSI cluster while allowing distributed IUCV policy to be changed dynamically
- APAR VM65989: Dump Processing Enhancements – The amount of time it takes for z/VM to write a hard abend or snap dump will be improved via changes to the I/O program used to write central memory to z/VM spool space located on 3390 DASD. This APAR also provides support for reducing dump size of SNAPDUMPs by excluding PGMBKs.

3Q17 z/VM 6.4 SPEs

- APARs VM65925, VM65926, VM65931: NICDEF Security Controls – Improved configuration and authorization support when coupling a virtual NIC to a VSwitch. Puts configuration and authorization into the z/VM User Directory for greater security control.
- APAR VM65942: z14 compatibility support, including alternate subchannel set dynamic I/O support. z/VM 6.4 provides dynamic I/O commands within the z/VM hypervisor for supporting define, modify, and delete of a device within either the default or alternate subchannel set. HCD support is currently not available.
- APAR VM65943, VM64945: EAV Minidisk Support for 3390-A DASD devices supporting 28-bit cylinder addresses allowing larger non-fullpack minidisks to reside anywhere on the volume, including beyond the current restriction of the 64K cylinder boundary (0-65519), and up to the 1 terabyte limit currently supported.
 - With documented restrictions
- APAR VM65988: the z/VM hypervisor has been enhanced to manage its spinlocks more efficiently and thereby reduce system overhead. This enhancement contributes to improved performance and throughput, and thereby helps to improve overall system capacity by allowing additional work to be performed. While most workload configurations will benefit to some extent, the improvements are greatest for workloads using large numbers of logical CPUs. The support is available on all hardware supported by z/VM V6.4. The design for some spin locks takes advantage of efficiencies unique to z14.
- APAR PI80912: The z/VM TCP/IP FTP client is configured to determine the party initially responsible for initiating data connections, improve compliance with RFC 2428, and reflect the active/passive FTP mode and the EPSV4 setting.

Additional 4Q17 z/VM 6.4 SPEs

- APAR VM66063: PR/SM LPAR Management Time Relief -- Provide two new options to affect processor unparking heuristics, which can reduce PR/SM dispatch contention arising from unproductive use of low-entitlement logical cores. (PTF available as of November 2017)
- APAR VM65918: Multi-VSwitch Link Aggregation Enhancements focusing on improving the area of load balancing to leverage both horizontal and vertical growth in both single and cross virtual switch networking configurations. (PTF available as of December 2017)

1Q18 z/VM 6.4 SPEs

- APARs VM66105 and VM66109: provide support for the IBM Z Hypervisor Performance Manager for z/VM (IBM ZHPM), which enables goal-oriented CPU management. The combination of z/VM and IBM ZHPM support enhances the capability to control access to CPU resources to achieve workload goals. IBM ZHPM executes within a Linux service virtual machine. IBM is currently soliciting sponsor users interested in exercising IBM ZHPM as part of a beta test program.
- APARs VM66098 and VM66108: Provide support for DS8880 z-Thin Provisioning and Extent Space-Efficient (ESE) volumes, enabling guest support for thin-provisioned volumes and removing the restriction that CPOWNER volumes cannot be defined on thin-provisioned volumes. [PTFs available as of March 28, 2018.](#)
- APAR VM66095: Provides new monitor records for the z/VM SCSI subsystem. The new information provided relates to the utilization of FCP devices (subchannel) and CHPIDs (physical channels) for EDEVICE activity. This provides z/VM clients information that can be used to determine if their existing EDEVICE configuration is optimal or should be altered for better performance.
- APAR VM66085: Provides new and updated performance reports within the Performance Toolkit Feature in support of HyperPAV Paging. These new reports include new information which will help clients tune the z/VM HyperPAV Paging Subsystem.

z/OS Support Summary



Release	z9 EC z9 BC WdfM OoS	z10 EC z10 BC WdfM	z196 z114 WdfM	zEC12 zBC12 WdfM	z13 z13s	z14 M0x z14 ZR1	End of Service	Extended Defect Support ¹
z/OS V1.13	X	X	X	X	X	X ¹	9/16	9/19*
z/OS V2.1	X	X	X	X	X	X	9/18*	9/21*
z/OS V2.2		X	X	X	X	X	9/20*	9/23*
z/OS V2.3				X	X	X	9/22*	9/25

Notes:

¹ The IBM Software Support Services for z/OS V1.13, offered as of October 1, 2016, provides the ability for customers to purchase extended defect support service for z/OS V1.13

* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

WdfM – Server has been withdrawn from Marketing
 OoS -- Out of support

Legend

IBM Software Support Services required for z/OS support
Generally supported

IBM Z server support ...

VSE Release	z9	z10	z196 / z114 / zEC12 zBC12 / z13 / z13s	z14	VSE EoM	VSE EoS
z/VSE V6.2	No	No	Yes	Yes	tbd	tbd
z/VSE V6.1	No	Yes	Yes	Yes	tbd	tbd
z/VSE V5.2	Yes	Yes	Yes	Yes	03/13/2017	10/31/2018
z/VSE V5.1	Yes	Yes	Yes	Yes	05/23/2014	06/30/2016
z/VSE V4.3	Yes	Yes	Yes	No	06/25/2012	10/31/2014
z/VSE V4.2	Yes	Yes	Yes	No	10/26/2010	10/31/2012
z/VSE V4.1	Yes	Yes	Yes	No	10/17/2008	04/30/2011
z/VSE V3.1	Yes	Yes	Yes	No	05/31/2008	07/31/2009
VSE/ESA V2.7	Yes	Yes	Yes	No	09/30/2005	02/28/2007
VSE/ESA V2.6	Yes	Yes	Yes	No	03/14/2003	03/31/2006

z/TPF V1.1 Support

- z/TPF V1 .1 (with PTFs) supports z14
 - z/TPF will not be updated to exploit new functionality, toleration only
 - z/Architecture mode only (ESA/390 TPF mode has been eliminated)
- Implicit benefits on z14:
 - Per-engine performance
 - FICON Express16S+ increased bandwidth
 - OSA Express6S improved network performance
 - Crypto Express6S improved performance
 - Improved performance diagnostic capabilities with expanded CPUMF functionality
- z/TPF V1 .1 with all APARs released prior to June 30, 2017 continues to support the following servers:
 - z13, z13s, zEC12, zBC12, z196, z114, z10 EC, z10 BC
- z/TPF APARs released after June 30, 2017 are only supported on
 - z14, z13, z13s, zEC12, zBC12

What is new in z/TPF with z14



- Dynamic CPU capacity
 - z/TPF will be able to utilize additional processors immediately without reactivating the LPAR or IPL'ing the z/TPF system
 - z/TPF environment will now be able to respond to sudden changes in demand which would previously require a system outage to accommodate the changes
- HiperDispatch enhancements
 - In installation prior to z14, z/TPF workload is evenly distributed across all available processors, even in low utilization situations. This causes cache and core contention with other LPARs. When z/TPF is running in a shared processor configuration, the achieved MIPS will be higher when z/TPF is utilizing a minimum set of processors.
 - In low utilization periods, [z/TPF will now minimize the processor footprint](#), by compressing TPF workload onto a minimal set of I-streams (engines), thus reducing the impact on other LPARs, allowing the entire CPC to operate more efficiently.
 - As a consequence, z/OS and z/VM will experience less contention from the z/TPF system when the z/TPF system is operating at periods of low demand.
- FICON Express16S+ exploitation
 - FICON Express16S+ channels will allow z/TPF processors to reduce the number of channels.
 - z/TPF will be able to handle the same IO workload using a reduced number of channels as compared to a z13. Reducing the number of channels decreases the complexity of the environment and reduces the amount of supporting hardware required to accomplish the same IO rates.

Virtualization and Linux on IBM Z

Dynamic Partition Manager (DPM)
KVM for IBM Z news
Linux Support for z14 - Overview

IBM Dynamic Partition Manager (DPM)

- DPM is a new administrative mode, introduced for managing Linux-based deployments for IBM z14, IBM z13™, IBM z13s™, IBM LinuxONE Emperor™ II (Emperor II) IBM LinuxONE Emperor™ (Emperor), IBM LinuxONE Rockhopper™ (Rockhopper), and IBM LinuxONE Rockhopper II .
- A system can be configured *either* in **DPM** mode or in **PR/SM** mode (POR is required to switch modes).
- DPM functions (high level):
 - Create, provision and manage partitions (processor, memory, adapters)
 - Manage network and storage configuration and allocation
 - Monitor and troubleshoot the environment
- Three things to remember about IBM Dynamic Partition Manager (DPM):
 - Fast -Much faster than managing with HCD and/or HCM. From hours to minutes. Does not require prior IBM Z knowledge.
 - Easy–Intuitive user interface. No need for multiple administrators with different skills or tools. Do NOT expect First In Enterprise Linux clients to adopt the legacy way.
 - Powerful–The same low overhead PR/SM hardware virtualization without the complexity. It supports dynamic configuration changes with just a few clicks of the mouse. It provides a foundation for bare metal Cloud.

DPM functions available via Graphical and API interfaces

IBM Hardware Management Console

SEARCH FAVORITES octavian

Home Partition Details - M90-L...

Partition Details - M90-LNXT02

General

Status

Controls

Processors

Memory

Network

Storage

Accelerators

Cryptos

⚠ Boot

Related Tasks

Stop

System Details

Manage Adapters

Monitor System

Processors

Processor mode: Shared Dedicated

Processors: 7

1 18 35 52 69 Processors

2 36 70 104 138 Threads

Processing weight: 100

999 Very High

900

700 High

500 Medium

300 Low

100 Very Low

1

Enforce weight capping

Enforce absolute processor capping

Number of processors (0.01-255.0): 1

Manage Processor Sharing

Processors

--- Entitled

Dedicated

M90-LNXT02

Shared

Shared Processors

Virtual/Physical: 58.82%

M90-LNXT02

M90-KVMP10

M90-LNXT01

M90-LNXP11

M90-LNXP10

Active Processing Weights

M90-LNXT02

M90-LNXT01

M90-KVMP10

M90-LNXP11

M90-LNXP10

Memory

OK Apply Cancel Help



Technical Specifications

- **IBM z14 M0x, z14 ZR1, z13 or z13s, IBM LinuxONE Emperor, Emperor II, Rockhopper or Rockhopper II**
 - Feature Code #0016: Hardware Requirements for IBM DPM
 - Two dedicated OSA-Express5/6S 1000BASE-T Ethernet
- **Supported Features**
 - FICON Express (Type FC and FCP)
 - ECKD DASD support
 - SCSI disk support
 - OSA Express5S, 6S
 - Crypto Express5S, 6S
 - zEDC Express
 - RoCE Express RoCE Express2 (as NIC)
 - Hipersockets
- **Hypervisors and Operating Systems**
 - KVM and/or Linux on Z & LinuxONE
 - ***z/VM support for FCP and FICON (type FC) configurations with z/VM 6.4***
 - ***IBM Secure Service Container Appliances***
- **Functions supported:**
 - Import an existing IOCP file
 - Limited ECKD paths
 - FCP / FICON adapter configuration
- **No support (yet) for:**
 - GDPS Virtual Appliance
 - FICON CTC (channel to channel links)
 - FICON tape

DPM References

- IBM Resource Link: <http://www.ibm.com/servers/resourceink>
 - Hardware Management Console
 - Driver Exception Letter
 - HMC Web Services API
- zHMC ClientPython Library
 - <https://github.com/zhmcclient>
- Live Virtual Classroom
 - <http://www.vm.ibm.com/education/lvc/LVC0309.pdf>
 - Demo Recording (registration information may ne required):
<https://attendee.gotowebinar.com/recording/7179881671264077572>

New delivery strategy for KVM on IBM Z and IBM LinuxONE

- Delivery model is changing; see:
- <https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03110USEN&>
- Delivery of KVM hypervisor and Linux operating system can be done more efficiently through our Linux distribution partners
 - Helps to simplify delivery of open source infrastructure
 - Linux and KVM now provided from a single source
 - Easier obtaining and installing of KVM
 - Distributions include open source cloud infrastructure components
 - Simplifies enabling of other technologies, such as container management and cloud management.



IBM Secure Service Container


IBM Secure Service Container Overview - What is it?

- IBM Secure Service Container (SSC*)
 - provides the base infrastructure to create an **IBM Z Appliance** including operating system, middleware, SDK and firmware support
 - Deploy an appliance (that provides a function or a service) in minutes instead of days
 - Protect the workload from being accessed by a system administrator or an external attacker
- Objective
 - To streamline the IBM Z application experience so it is comparable to installing an app on a mobile device.
- Terminology
 - **IBM Z Appliance** is an integration of operating system, middleware and Software components that work autonomously and provides core services and infrastructure focusing on consumability and security
 - **SSC Partition** is a LPAR type running an **IBM Z Appliance** based on **Secure Service Container**. The Partition type provides main aspects of the security features.

*Formerly IBM Z Appliance Container Infrastructure - zACI)

Traditional approach -- heterogeneous deployment

Multiple appliances in the network (e.g. accelerators, network security appliances, health monitors) scattered over various HW servers and/or virtual machines.

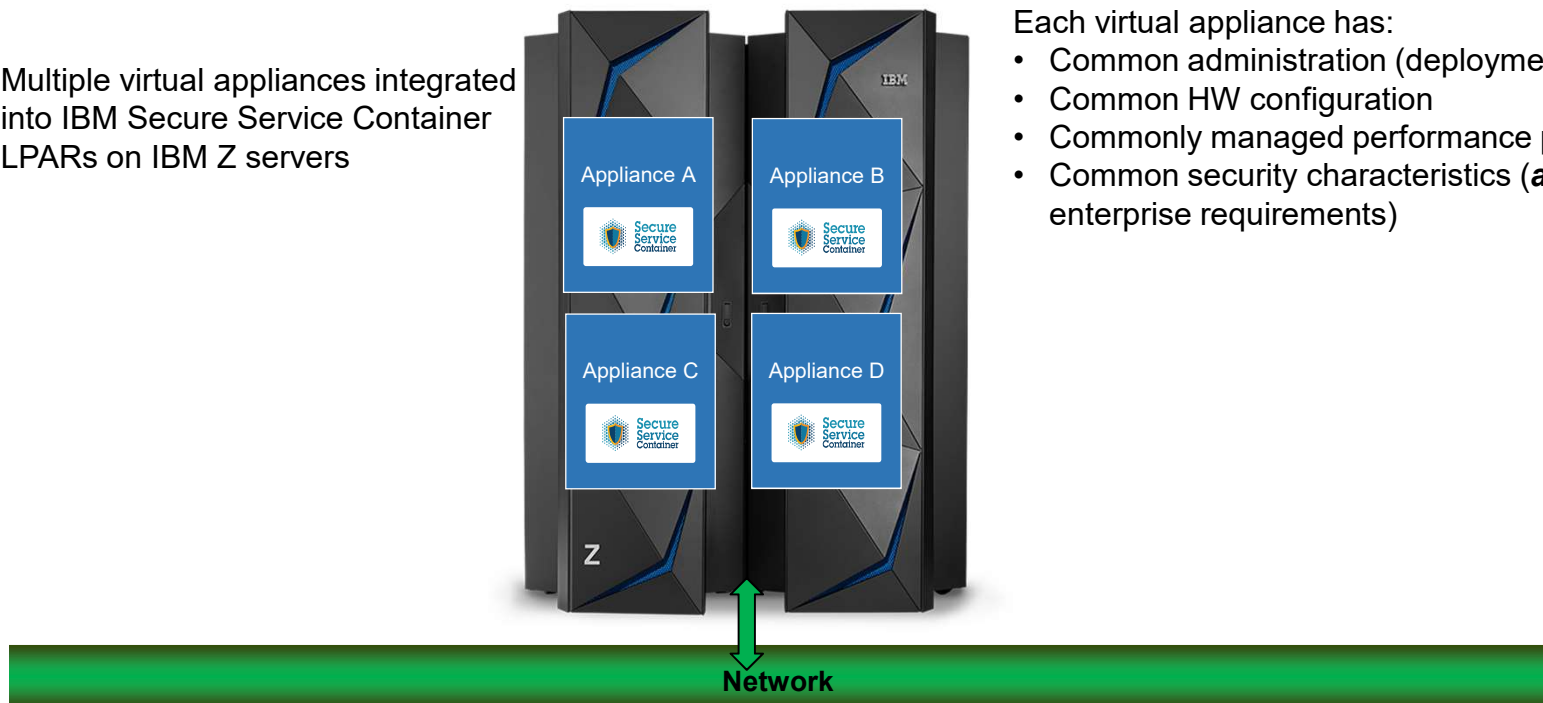


Each appliance has:

- Separate administration (and deployment)
- Different HW configuration requirements
- Different performance profile and management
- Different security characteristics (**requiring alignment** with enterprise requirements)

IBM Z and Appliances Integration

Multiple virtual appliances integrated into IBM Secure Service Container LPARs on IBM Z servers



Each virtual appliance has:

- Common administration (deployment)
- Common HW configuration
- Commonly managed performance profiles
- Common security characteristics (**aligned** with enterprise requirements)

Statements of Direction

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remain at our sole discretion.

Statements of General Direction

- **Future z/VM release guest support:** z/VM V6.4 will be the last z/VM release supported as a guest of z/VM V6.2 or older releases.
- **Disk-only support for z/VM dumps:** z/VM V6.4 will be the last z/VM release to support tape as a media option for stand-alone, hardabend, and snap dumps. Subsequent releases will support dumps to ECKD DASD or FCP SCSI disks only.
- **IBM intends to deliver IMS exploitation of IBM z14 and DS8880 zHyperLink™ WRITE operations.** zHyperLink Express is a direct connect short distance IBM Z I/O adapter designed to work in conjunction with a FICON or High Performance FICON SAN infrastructure.
- **IBM z14 will be the last z Systems and IBM Z high-end server to support FICON Express8S:** z14 will be last z Systems and IBM Z high-end server to support FICON Express8S (#0409 and #0410) channels. Enterprises should begin migrating from FICON Express8S channels to FICON Express16S+ channels (#0427 and #0428). FICON Express8S will not be supported on future high-end IBM Z servers as carry forward on an upgrade.

Statements of General Direction (Cont.)

- **OSA-Express6S 1000BASE-T adapters:** OSA-Express6S 1000BASE-T adapters (#0426) will be the last generation of OSA 1000BASE-T adapters to support connections operating at 100 Mb/second link speed. Future OSA-Express 1000BASE-T adapter generations will support operation only at 1000 Mb/second (1Gb/s) link speed.
- **Secure Service Container for IBM Cloud Private on Linux on Z and LinuxONE:** IBM intends to deliver IBM Cloud Private on Linux on Z in a Secure Service Container. IBM Cloud Private in Secure Service Container on Linux on Z and LinuxONE is designed to offer unmatched security of IBM Cloud Private nodes with workload isolation, pervasive encryption of data and executable binaries for container-based applications, and protection from privileged users to mitigate breaches and leaks from internal or external threats, ransomware, and malware.
- **IBM Storage and IBM Z Collaboration on 16u Reserved (#0617) Applications:** IBM Storage intends to deliver a high performance, high-availability, ECKD/Distributed format flash storage that can be mounted in a z14 ZR1 or LinuxONE Rockhopper II server rack with the associated 16U Reserved (FC #0617) space inside the 19" rack. . This new storage solution opens the opportunities for clients who require the high-availability, performance and Z synergy functions of DS8880, and are looking for ways to increase their data center flexibility.

Statements of General Direction (Cont.)

- **New Db2 Analytics Accelerator deployment option on the IBM Z infrastructure:**
The Db2 Analytics Accelerator for z/OS extends IBM Z and Db2 for z/OS to form a hybrid environment, which is capable of running both transactional and analytical SQL query workload. This extension is currently delivered via two deployment options: either as an appliance, based on IBM PureData System for Analytics, or as a hosted cloud environment. In addition to these two form factors, IBM intends to deliver a new Db2 Analytics Accelerator deployment option on the IBM Z infrastructure. This would further extend deployment options available to Db2 Analytics Accelerator clients. This new deployment option would allow for deeper integration with the IBM Z infrastructure. Clients would benefit with the flexibility to deploy the form factor that best suits their requirements to enable unified homogeneity of service, support and operations and deeper integration with their processes, e.g. for their disaster recovery.
- **IBM intends to deliver VSAM exploitation of z14 and DS8880 zHyperLink Express.** zHyperLink Express is a short distance mainframe attach link designed for up to 10x lower latency than High Performance FICON.

End of Presentation

Questions and Comments

please send to:

jemcleo@us.ibm.com