IBM

IBM zEnterprise 114

Chris Gombola
System z Client Architect

What is IBM zEnterprise System?

Re-write the rulebook and set new standards for business-centric IT with IBM System z, to be the world's premier workload-optimized platform for enterprise applications.

Our Vision:

Deliver the best of all worlds - Mainframe, UNIX®, x86 and single function processors - integrated in a single system for ultimate flexibility and simplicity to optimize service, risk, and cost across multiple heterogeneous workloads.

IBM zEnterprise: A centrally managed enterprise cloud for the flexible delivery of high value services.

- **A highly scalable heterogeneous pool of virtualized resources managed in a single system.**
- **Activate, allocate, prioritize and retire resources on demand and automate service delivery.**
- **Maximize utilization of resources for improved ROI and for low cost service delivery.**
- **Bring new levels of security, resiliency and manageability to create a cloud environment that is enterprise ready.**

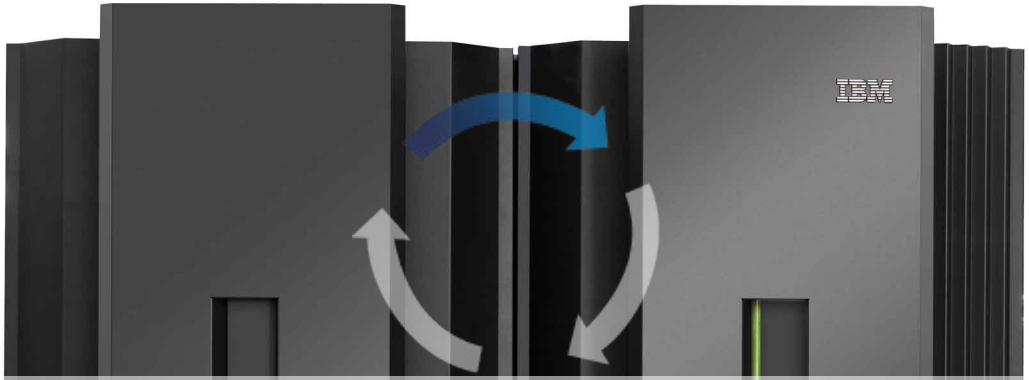

Security

Availability

Scalability

Efficiency

Virtualization

New Blades Provide Added Flexibility for Workload Deployment and Integration

- **Introducing System x Blades in the zBX**
  - Select IBM BladeCenter HX5 7873 dual-socket 16-core blade
    - The zBX web page will host the most current blade ordering information: http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfi
  - Ordered and fulfilled through System x providers and installed into the zBX by the customer
  - Blades assume System z warranty and maintenance when installed in the zBX
- **Unified Resource Manager will install an integrated hypervisor on blades in the zBX**
  - KVM-based with IBM service and support
- **Up to 112 Blades supported on zBX**
  - Ability to mix and match DataPower XI50z, POWER7 and System x blades in the same chassis for better zBX utilization
  - IBM Smart Analytics Optimizer can mix with others in same rack
  - Number of blades supported varies by type

IBM zEnterprise BladeCenter Extension (zBX)
Machine Type: 2458 Mod 002

Optimizers
- IBM Smart Analytics Optimizer
- IBM WebSphere DataPower Integration Appliance XI50z for zEnterprise

Select IBM Blades
- IBM BladeCenter PS701 Express
- IBM BladeCenter HX5 7873 blade

One to four – 42u racks – capacity for up to 112 blades

- Up to 112 PS701 Power blades
- Up to 28 HX5 System x blades
- Up to 28 DataPower XI50z blades (double-wide)
- Up to 56 IBM Smart Analytics Optimizer blades

IBM zEnterprise 114 (z114)

z114 Overview

- **Machine Type**
  - 2818
- **2 Models**
  - M05 and M10
  - Single frame, air cooled
  - Non-raised floor option available
  - Overhead Cabling and DC Power Options
- **Processor Units (PUs)**
  - 7 PU cores per processor drawer (One for M05 and two for M10)
  - Up to 2 SAPs per system, standard
  - 2 spares designated for Model M10
  - Dependant on the H/W model - up to 5 or 10 PU cores available for characterization
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs)
  - 130 capacity settings
- **Memory**
  - Up to 256 GB for System including HSA
    - System minimum = 8 GB (Model M05), 16 GB (Model M10)
    - 8 GB HSA separately managed
    - RAIM standard
    - Maximum for customer use 248 GB (Model M10)
    - Increments of 8 or 32 GB
- **I/O**
  - Support for non-PCIe Channel Cards
  - Introduction of PCIe channel subsystem
    - Up to 64 PCIe Channel Cards
  - Up to 2 Logical Channel Subsystems (LCSSs)
- **STP - optional** (No ETR)

zEnterprise 114 Models M05 and M10

- **M/T 2818 – Model M05**
  - Air cooled
  - Single Frame
  - Non-raised floor option available
  - 30 LPARs
- **Processor Units (PUs)**
  - New processor drawer design (1 CPC Drawer)
  - 7 per system
    - 2 SAPs standard
    - Up to 5 CPs
    - Up to 5 specialty engines
    - Up to 2 zIIPs/zAAPs
    - 0 spares when fully configured

- **M/T 2818 – Model M10**
  - Air cooled
  - Single Frame
  - Non-raised floor option available
  - 30 LPARs
- **Processor Units (PUs)**
  - New processor drawer design (2 CPC Drawers)
  - 14 per system
    - 2 SAPs standard
    - Up to 5 CPs
    - Up to 10 specialty engines
    - Up to 5 zIIPs/zAAPs
    - 2 dedicated spares

- **When Model M10 (requires the 2nd processor drawer)?**
  - > 5 Customer PUs
  - > 120 GB memory
  - > 4 Fanouts for additional I/O connectivity – especially PSIFB links
    - Depends - numbers vary for drawers, I/O features and PSIFB links

z114 Model M10 – Under the covers

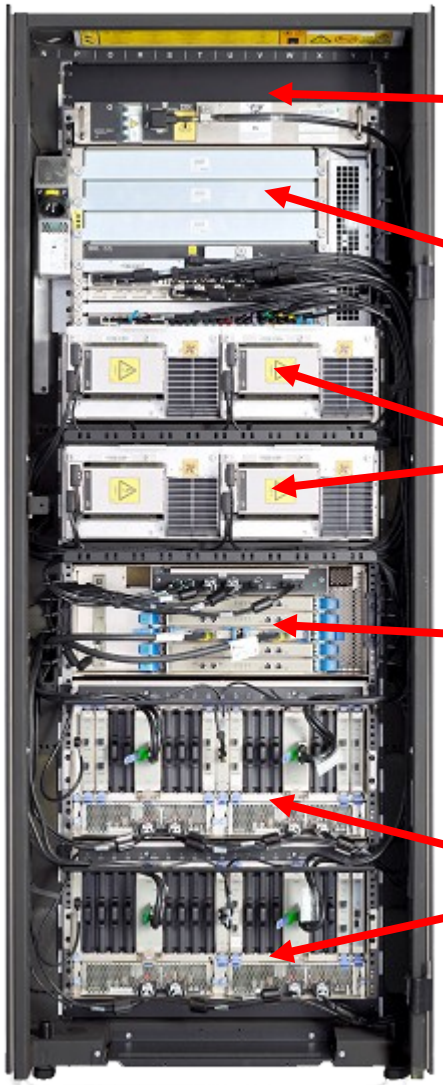

Internal Batteries (optional)

Power Supplies

2 x Processor Drawers, Memory & HCAs

I/O Drawer

PCIe I/O drawers

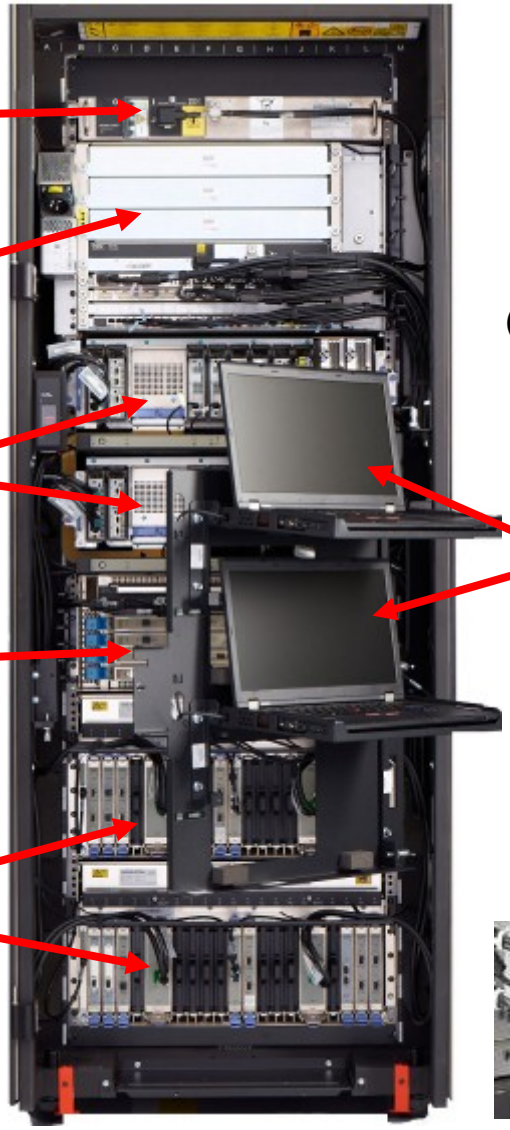

Ethernet cables for internal System LAN connecting Flexible Service Processor (FSP) cage controller cards (not shown)

2 x Support Elements

Rear View

Front View

z114 Model Structure and Upgrades

Model	CPs	IFLs Unassigned IFLs	zAAPs	zIIPs	ICFs	Std. SAPs	Add'l SAPs	Spares
M05	0-5	0-5	0-2	0-2	0-5	2	0-2	0
M10	0-5	0-10	0-5	0-5	0-10	2	0-2	2

- **Model structure based on number of drawers**
- **M05 sparing based on prior Business Class (BC) offerings – no dedicated spares**
- **M10 sparing based on Enterprise Class (EC) offerings – dedicated spares**
  - SAP and PU Allocation/Sparing in the M10
    - Default assignment is one SAP per drawer; one Spare per drawer. Spill and fill CP low to high; spill and fill specialty engines high to low
    - Two defective PUs may cause the default assignment to spill and fill into the second processor drawer. LPAR has the capability to request PU of a specified type to be grouped together in a book/drawer (i.e. LPAR may change the default assignment)
- ***Disruptive* upgrade from M05 to M10**
  - No model downgrades
- **Upgrades from z9 BC and z10 BC into either model M05 or M10**
- **Only the M10 will upgrade to z196 Model M15 (Air cooled only)**

z114 Upgrade paths

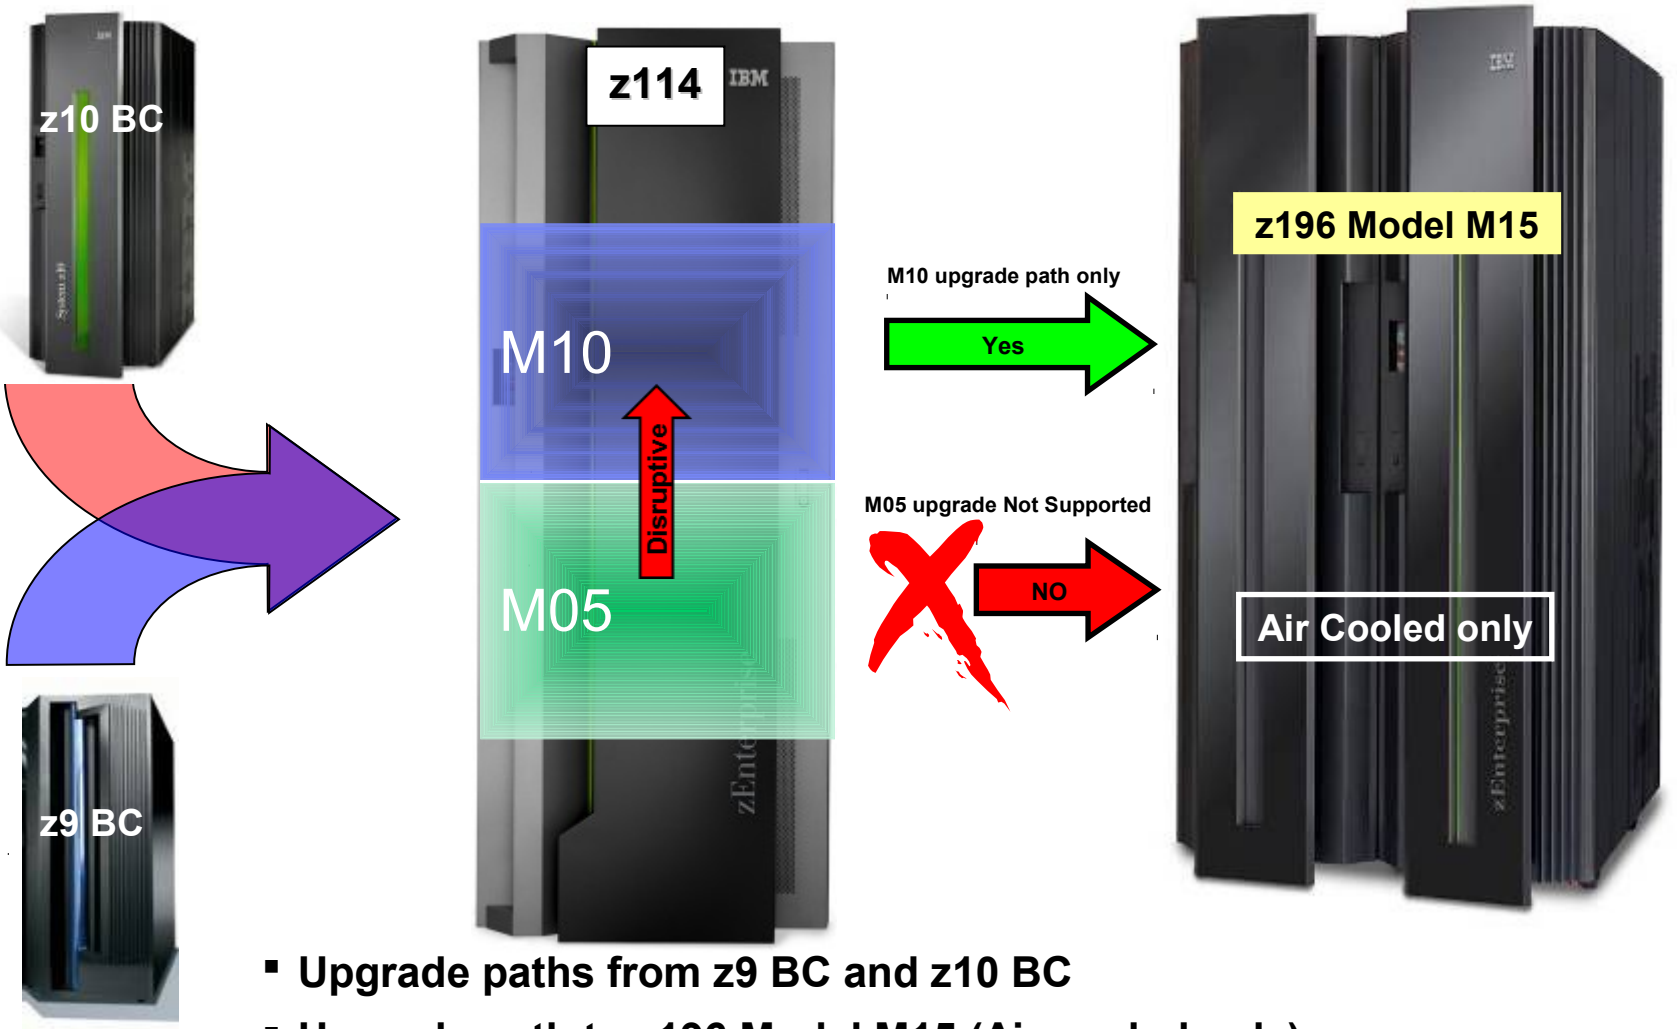

- Upgrade paths from z9 BC and z10 BC
- Upgrade path to z196 Model M15 (Air cooled only)
- Disruptive upgrade M05 to M10 and from M10 to z196 M15

z114 RAIM Memory

- **Memory technology introduced on the z196 is used on the z114**
  - Redundant Array of Memory (RAIM) which in the Disk industry is known as RAID
  - Protection from Unscheduled Incident Repair Actions (UIRAs) caused by a DIMM failure
    - DIMM failures include all components on the DIMM
    - Portions of the memory controller or card failure isolated to one memory channel
- **Flexible memory option not available on z114**
  - Used for Enhanced Book Availability on z196

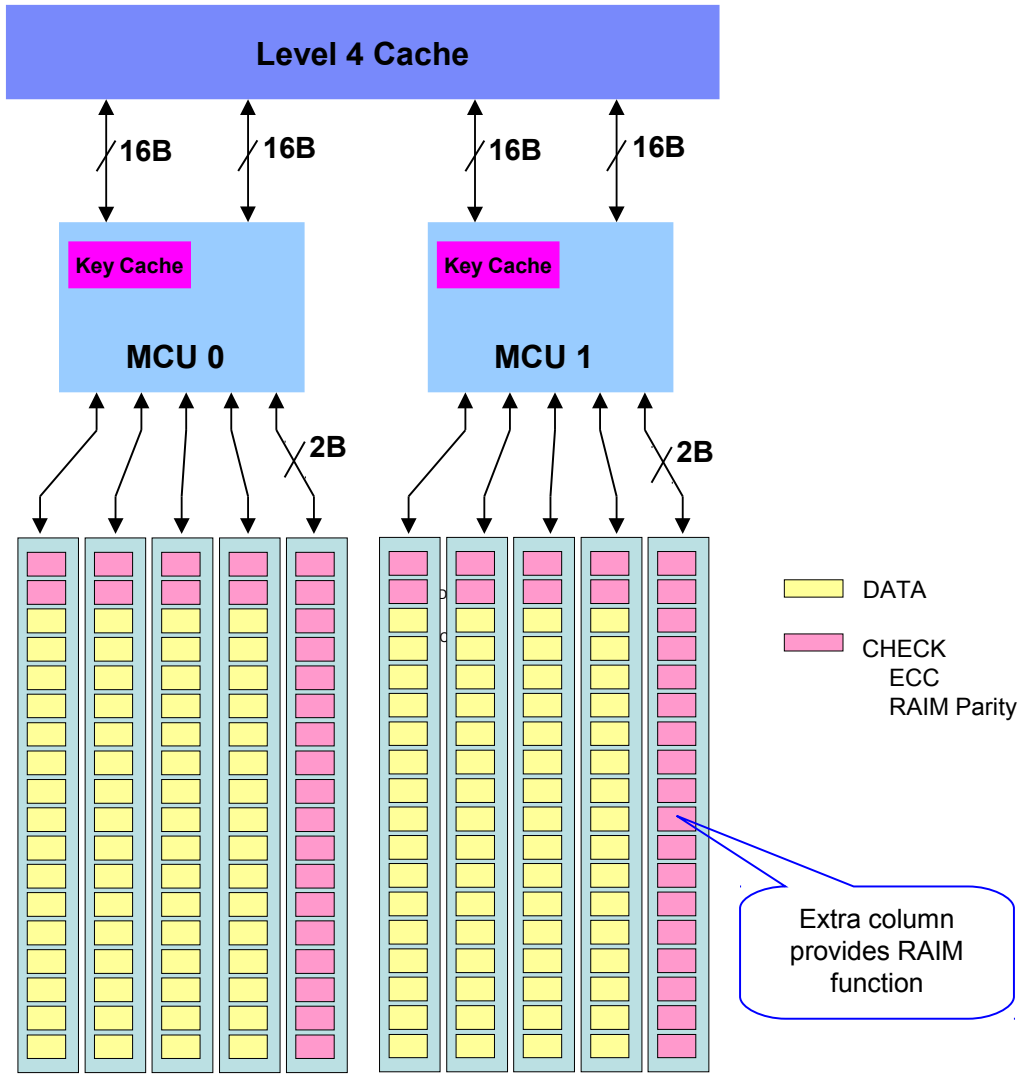

System z HSA Comparisons

- **HSA significantly larger than pre-z10 Servers**
- **Fixed HSA, does not affect customer purchased memory**
- **Size of HSA on prior Servers (dependant on defined configuration)**
  - Multiprise® 2000 From 12 MB up to 40 MB
  - 9672 G4 From 48 up to 64 MB
  - Multiprise 3000 From 38 MB up to 136 MB
  - 9672 G5/G6 From 64 MB up to 192 MB
  - z800 From 160 MB up to 256 MB
  - z900 From 288 MB up to 512 MB
  - z890 From 768 MB up to 1.9 GB
  - z990 From 1 GB MB up to 2 GB
  - z9 BC From 896 MB up to 2.7 GB
  - z9 EC From 1.2 GB up to 4.2 GB
  - z10 BC 8 GB – Fixed
  - z10 EC 16 GB – Fixed
  - z196 16 GB – Fixed
  - z114 8 GB - Fixed
- **HSA Estimator on Resource Link not relevant**

FICON Express8S – SX and 10KM LX in the PCIe I/O drawer

- **For FICON, zHPF, and FCP environments**
  - CHPID types: FC and FCP
    - 2 PCHIDs/CHPIDs
- **Auto-negotiates to 2, 4, or 8 Gbps**
- **Increased performance compared to FICON Express8**

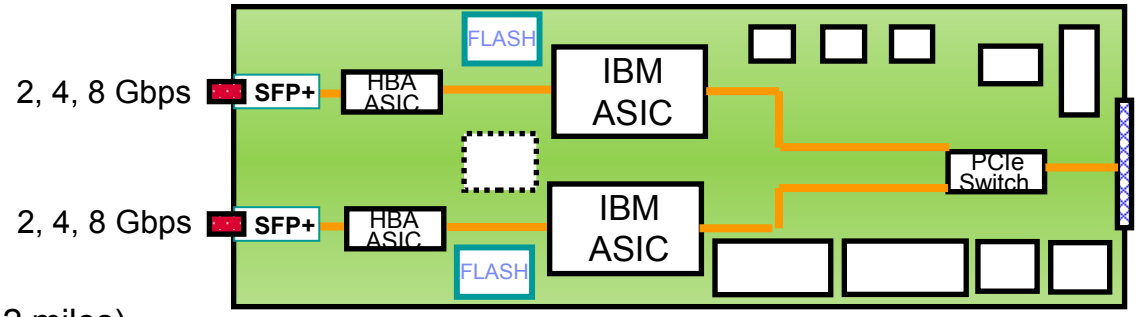

FC 0409 – 10KM LX, FC 0410 – SX

- **10KM LX - 9 micron single mode fiber**
  - Unrepeated distance - 10 kilometers (6.2 miles)
  - Receiving device must also be LX
- **SX - 50 or 62.5 micron multimode fiber**
  - Distance variable with link data rate and fiber type
  - Receiving device must also be SX
- **2 channels of LX or SX (no mix)**
- **Small form factor pluggable (SFP) optics**
  - Concurrent repair/replace action for each SFP

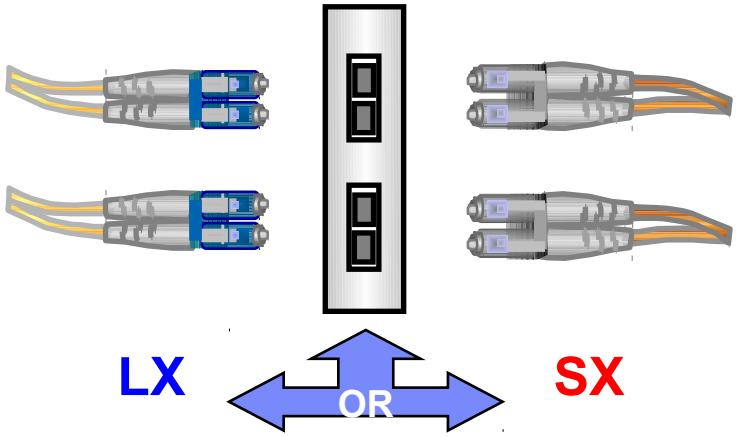

OSA-Express4S GbE and 10 GbE fiber for the PCIe I/O drawer

- **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

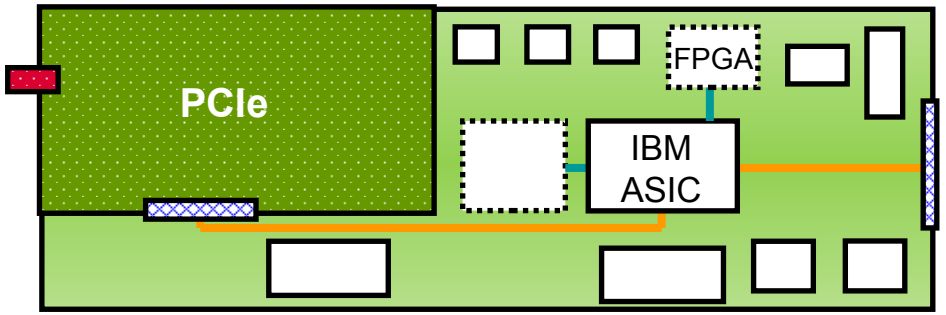

FC 0406 – 10 GbE LR, FC 0407 – 10 GbE SR

- **Gigabit Ethernet (GbE)**
  - CHPID types: OSD (no CHPID=OSN)
  - Single mode (LX) or multimode (SX) fiber
  - Two ports of LX or two ports of SX
    - 1 PCHID/CHPID

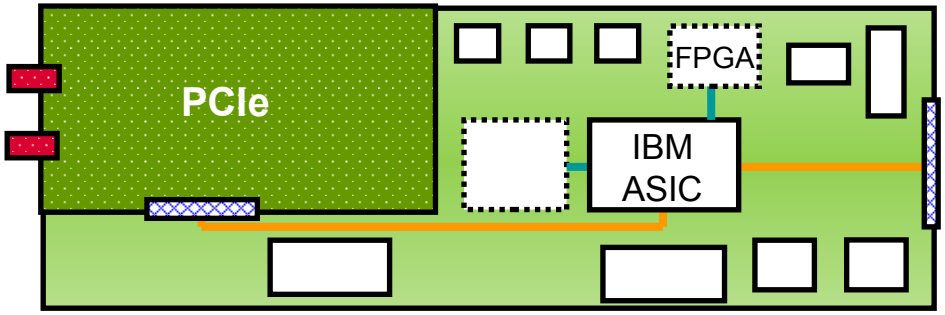

FC 0404 – GbE LX, FC 0405 – GbE SX

- **Small form factor optics – LC Duplex**

OSA-Express3 1000BASE-T 'copper' for I/O Drawer

- **Auto-negotiation to 10, 100, 1000 Mbps**
- **Double the port density of OSA-Express2**
- **Reduced latency & improved throughput**
  - Ethernet hardware data router
- **Improved throughput – standard & jumbo frames**
  - New microprocessor
  - New PCIe adapter
- **Usage of ports**
  - OSC, OSD, OSE can exploit four ports
  - OSM, Port 0 only
  - CHPID type OSN is for LPAR-to-LPAR communication. Does not use ports

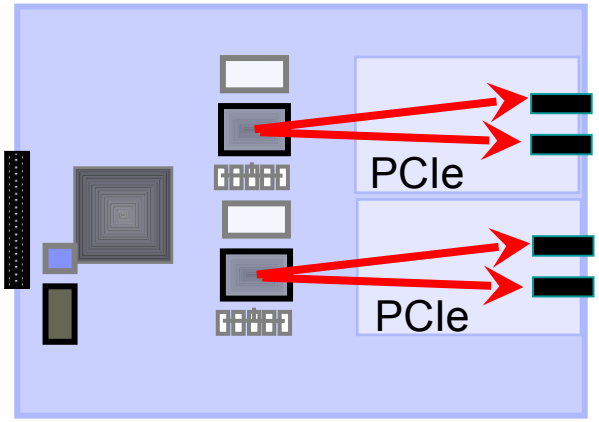

CHPID shared by two ports

4 ports - FC 3367 1000BASE-T

2 ports - FC 3369 for z10 BC and z114 only

	OSA-Express2	OSA-Express3
Microprocessor	448 MHz	667 MHz
PCI bus	PCI-X	PCIe G1

Mode	CHPID	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, IPL CPCs, and LPARs, 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
OSA for NCP (LP-to-LP)	OSN	NCPs running under IBM Communication Controller for Linux (CDLC)
Unified Resource Manager	OSM	Connectivity to intranode management network (INMN) from z196/z114 to zBX

Open Systems Adapter CHPIDs types: OSM (Express3 1000BASE-T)) and OSX (Express3 and 4S 10 GbE)

- **Two new OSA CHPID types to support new types of z196/z114 networks**
- **A z196/z114 System can have up to 6 types of OSA CHPID's**
  - External (customer managed) networks
    - Defined as OSC,OSD,OSE, & OSN
    - Existing customer provided and managed OSA ports used for access to the current customer external networks - no changes
  - Intranode management network (INMN)
    - Defined as CHPID type OSM, OSA-Express for Unified Resource Manager
      - When the PCIe adaptor on 1000BASE-T is defined as CHPID type OSM, the second port cannot be used for anything else
    - OSA-Express3 1000BASE-T configured as CHPID type OSM for connectivity to INMN from z196/z114 to Unified Resource Manager functions
    - OSA connection via the Bulk Power Hub (BPH) on the z196/z114 to the Top of the Rack (TORs) switches on zBX
  - Intraensemble data networks (IEDN)
    - Defined as CHPID OSX, OSA-Express for zBX
    - OSA-Express3 or 4S 10 GbE configured as CHPID type OSX for connectivity and access control to IEDN from z196/z114 to zBX
- **Functions Supported:**
  - Dynamic I/O support
  - HCD
  - CP Query capabilities
  - Ensemble Management for these new channel paths and their related subchannels.
- **z/VM 5.4 – CHPID types OSX and OSM cannot be varied online**

z114 continues the CMOS Mainframe heritage

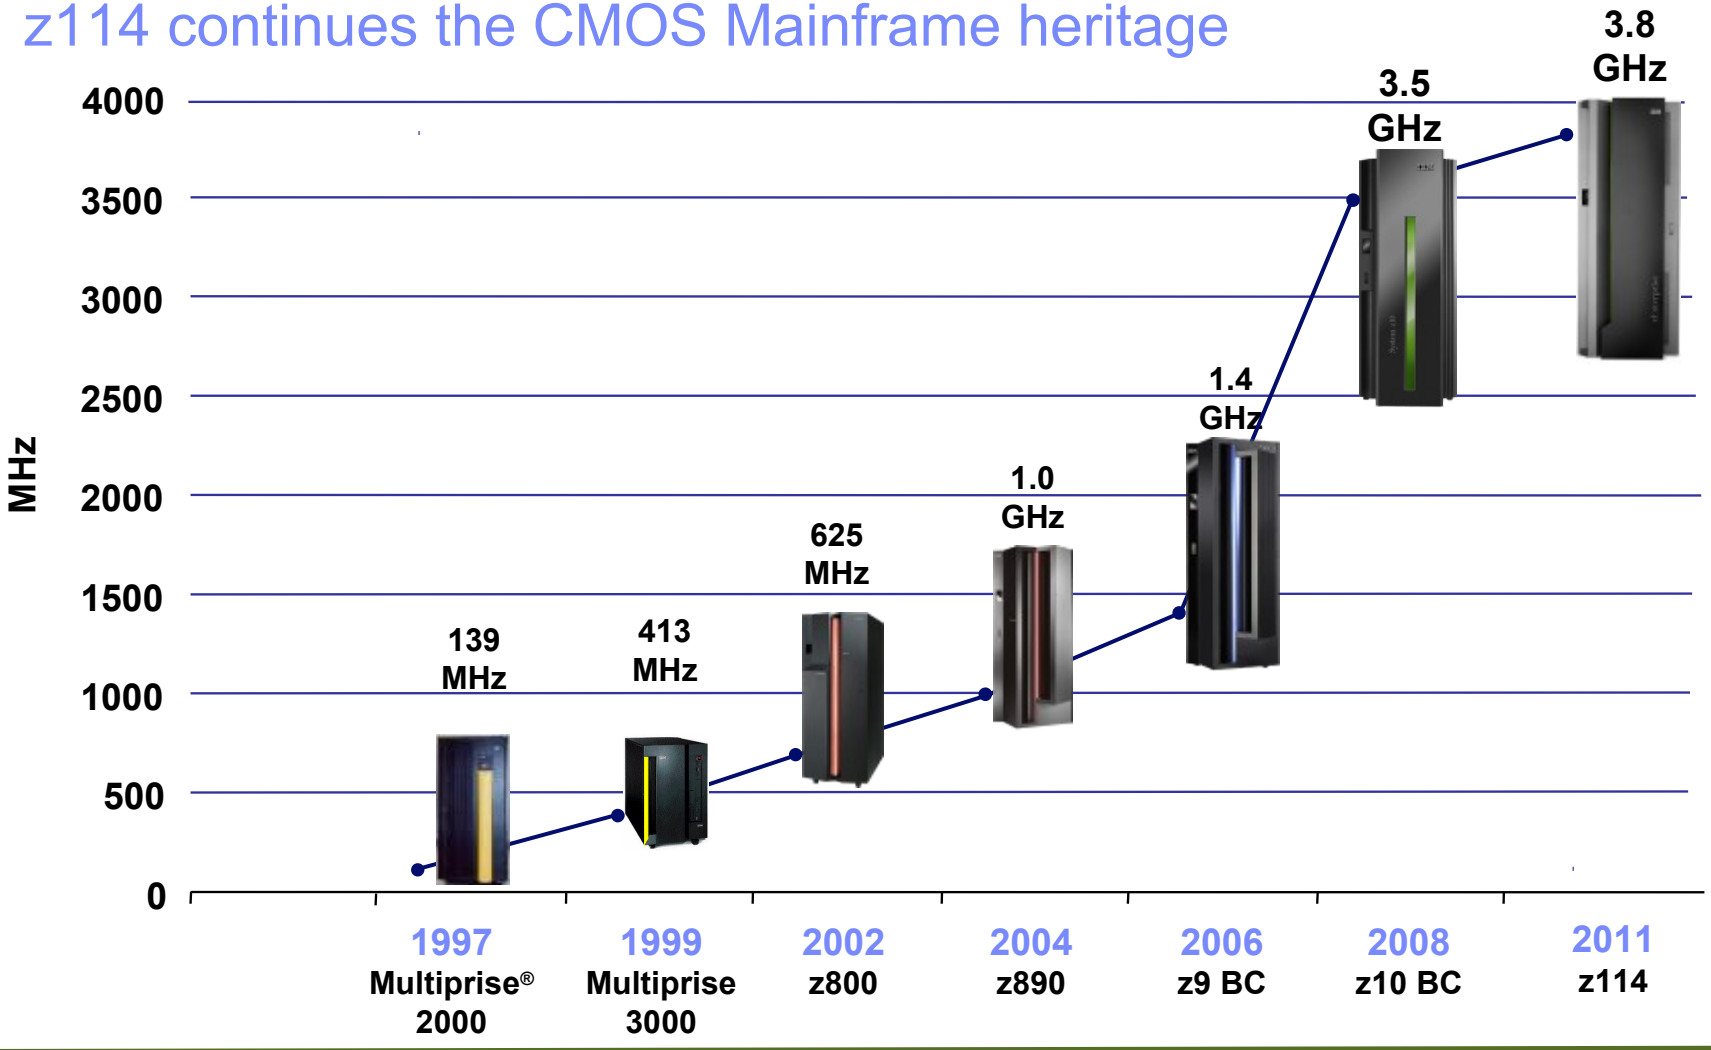

- Multiprise 2000 – 1st full-custom Mid-range CMOS S/390
- Multiprise 3000 – Internal disk, IFL introduced on midrange

- z800 - Full 64-bit z/Architecture®
- z890 - Superscalar CISC pipeline
- z9 BC - System level scaling

- z10 BC - Architectural extensions
  - Higher frequency CPU
- z114 – Additional Architectural extensions and new cache structure

zEnterprise System Out of Order (OOO) Value

- **OOO yields significant performance benefit for compute intensive apps through**
  - –Re-ordering instruction execution
    - Later (younger) instructions can execute ahead of an older stalled instruction
  - –Re-ordering storage accesses and parallel storage accesses
- **OOO maintains good performance growth for traditional apps**

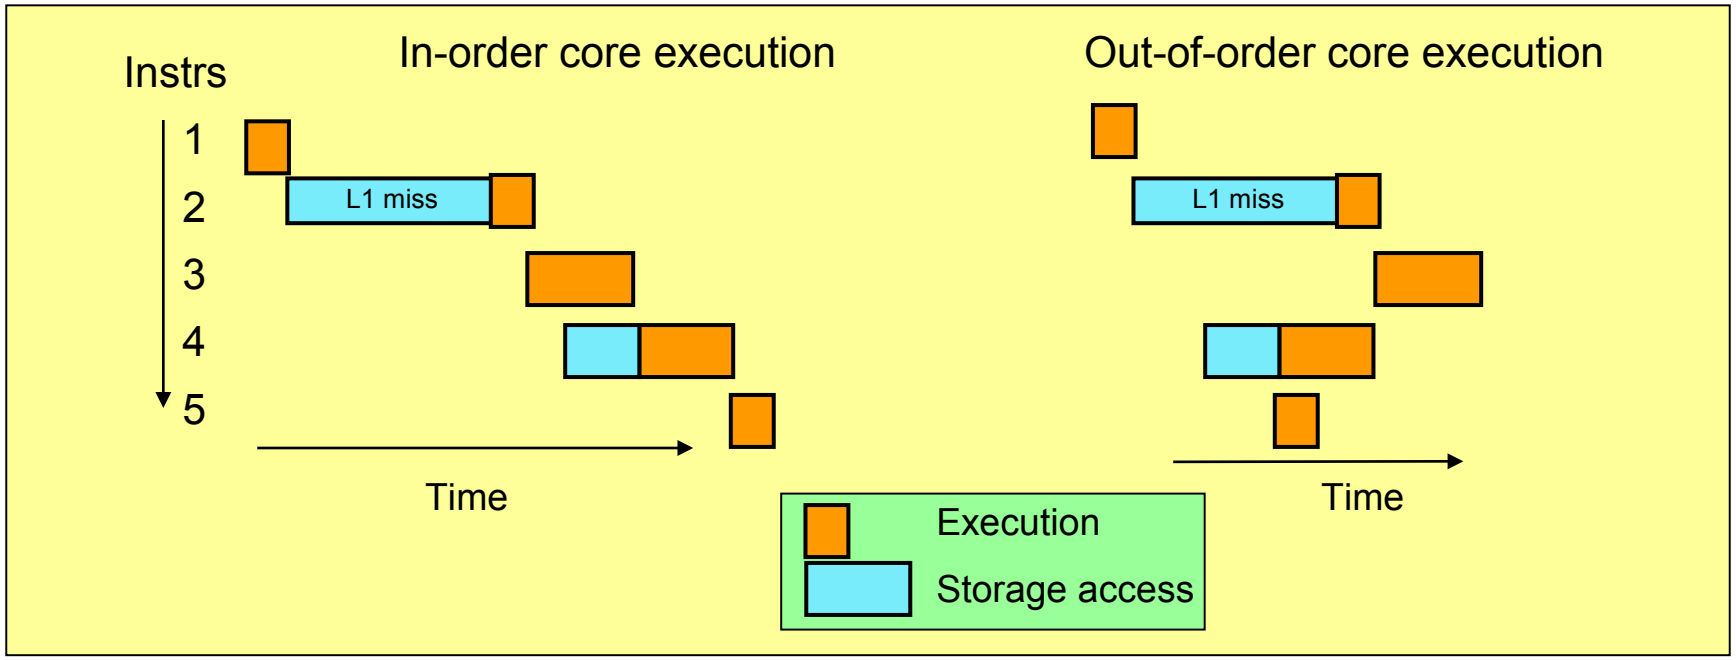

zEnterprise 114 performance and scalability

Think both Inside and Outside the box!

- **Faster Processors**
  - 18% Improvement in Uniprocessor speed
  - 12% Improvement in overall System capacity
- **Architectural equivalence to the z196**
- **More economic delivery of equivalent capacity**

Think System z Qualities of Service!

- **zEnterprise 114 Multiple Architecture Ensembles**
  - Tightly integrated heterogeneous systems
  - Robust, n+1 configurations
  - Highly virtualized
  - Managed end-to-end
  - Improved economies of scale, efficiency, and price performance

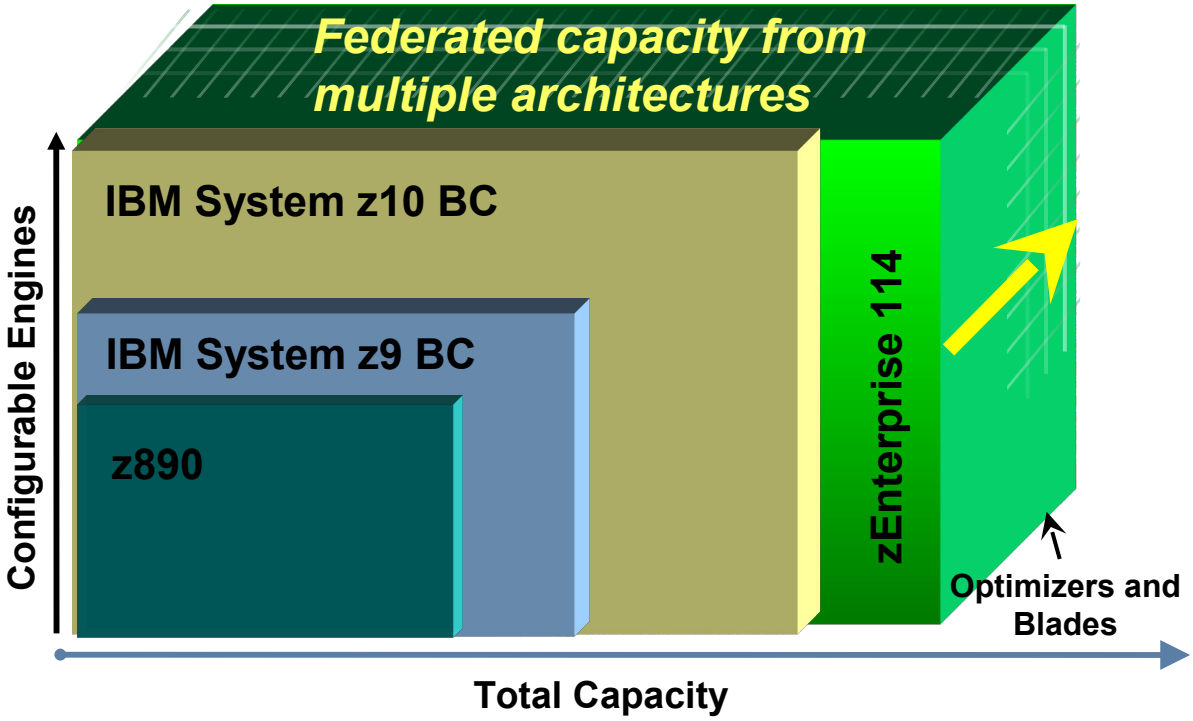

System Comparisons

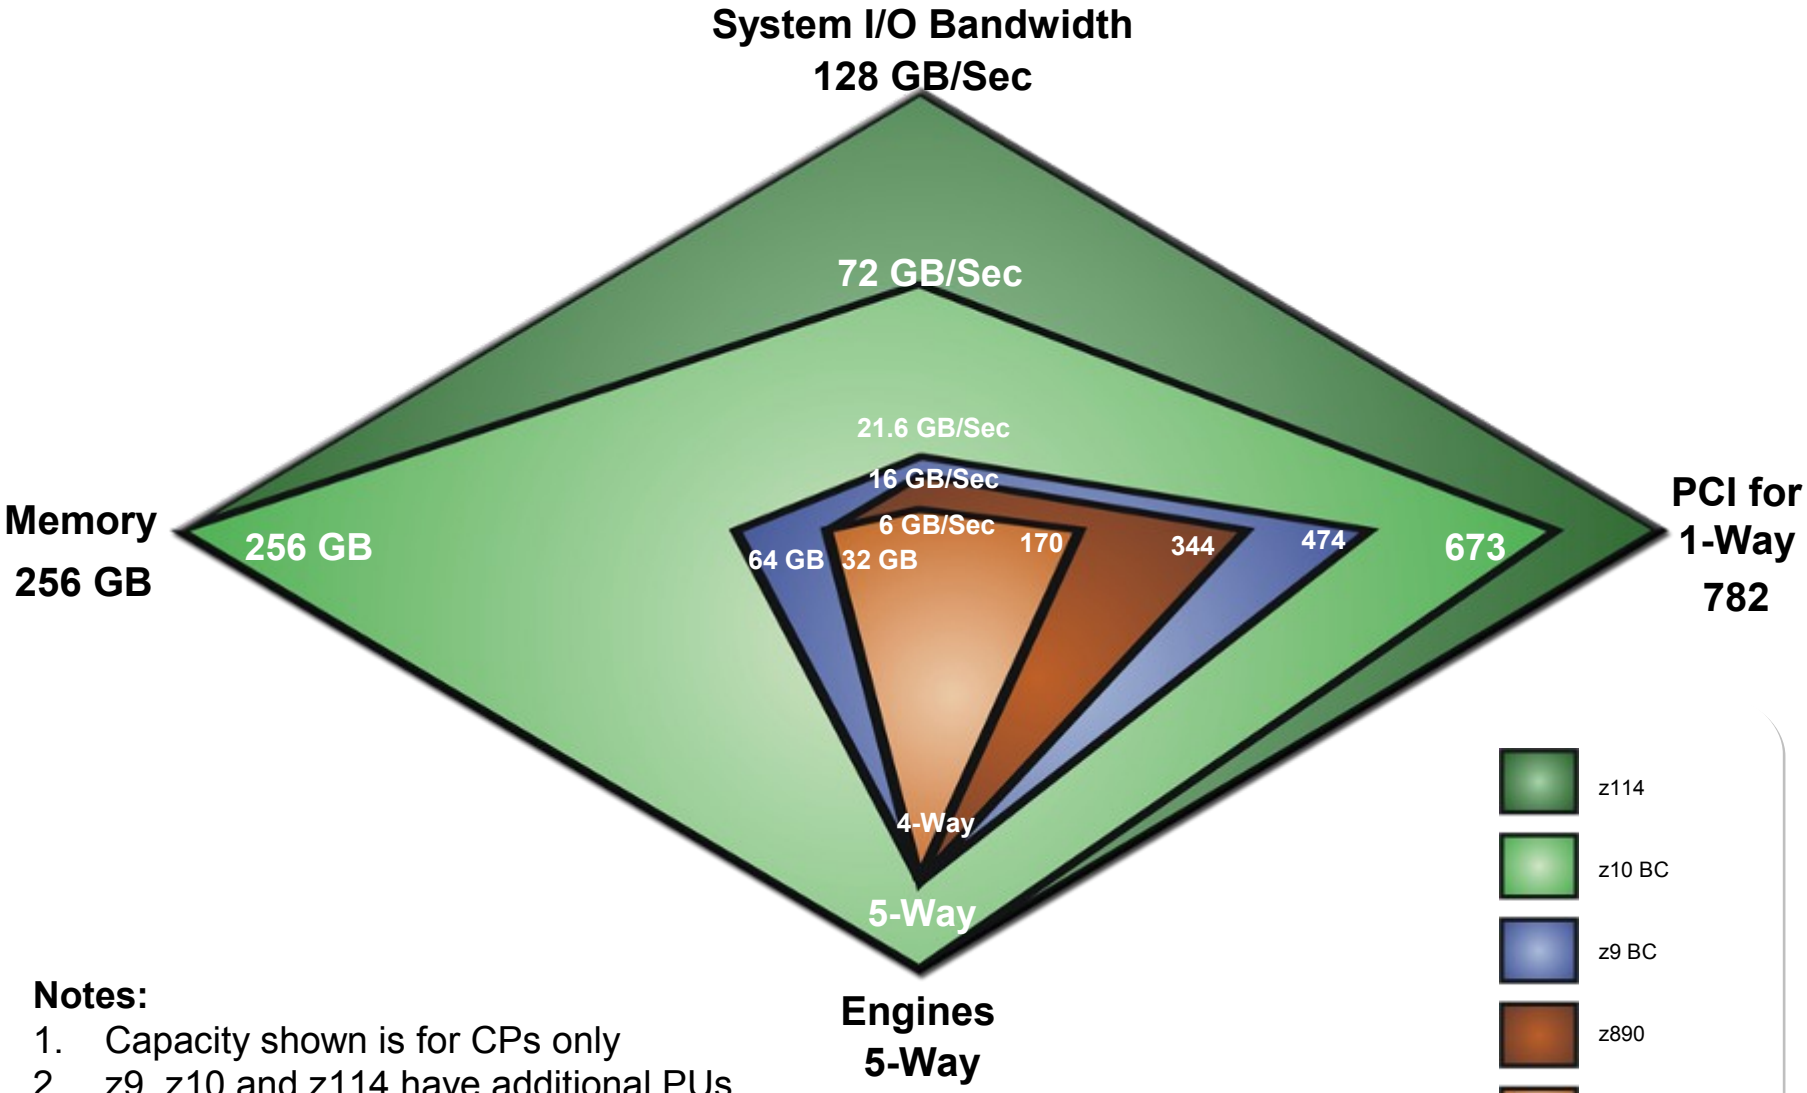

Notes:

1. Capacity shown is for CPs only
2. z9, z10 and z114 have additional PUs which can be used as Speciality Engines

z114 Capacity Matrix

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

z114

- Granularity levels similar to z10 BC to facilitate upgrades and incremental growth
- Nomenclature: **XYY**
  - **X = Capacity level**
  - **YY= Number of processors**
  - **A00 = ICF or IFL only**
- 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
- Linux only and ICF only servers
- Model M10 provides specialty engine scale out capabilities

Additional engines available on the M10

Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine

z114 Sub-capacity Processor Granularity

- **The z114 has 26 CP capacity levels (26 x 5 = 130)**
  - Up to 5 CPs at any capacity level
    - All CPs must be the same capacity level
- **The one for one entitlement to purchase one zAAP and/or one zIIP for each CP purchased is the same for CPs of any speed.**
  - All specialty engines run at full speed
  - Processor Unit Value for IFL = 100

Number of z114 CPs	Base Ratio	Ratio z10 BC to z114
1 CP	z10 BC Z01	1.18
2 CPs	z10 BC Z02	1.16
3 CPs	z10 BC Z03	1.14
4 CPs	z10 BC Z04	1.13
5 CPs	z10 BC Z05	1.12

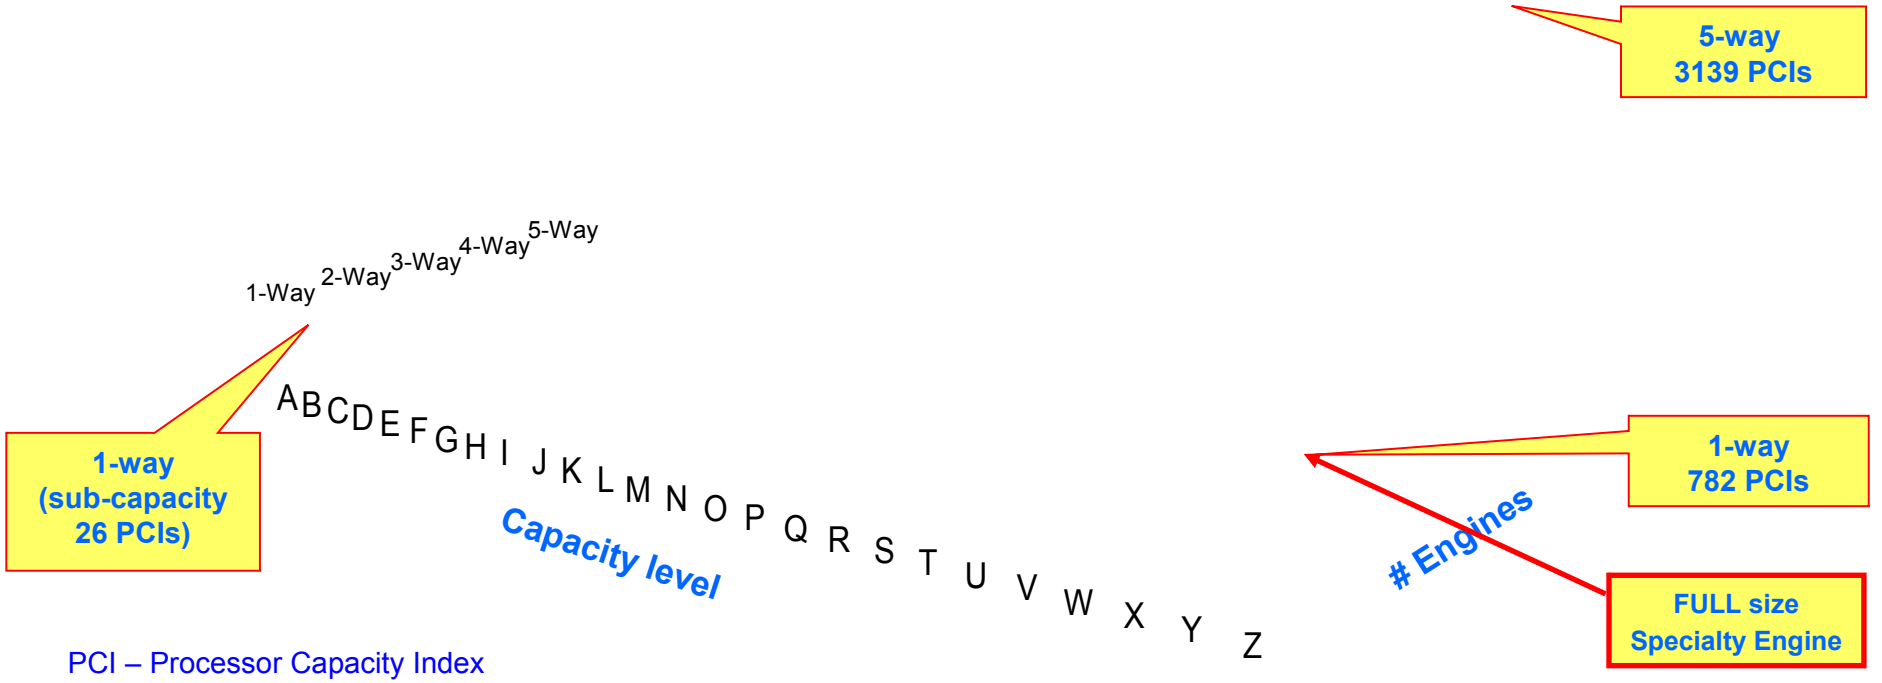

PCI – Processor Capacity Index

System z I/O Subsystem Internal Bus Interconnect Speeds (GBps)

PCIe
z196/z114
Sep 2011

8 GBps

InfiniBand
z10/z196/z114

6 GBps

STI
z9

2.7 GBps

STI
z990/z890

2 GBps

STI
z900/z800

1 GBps

STI: Self-Timed Interconnect

z114 and z196 GA2 Parallel Sysplex Coupling Connectivity

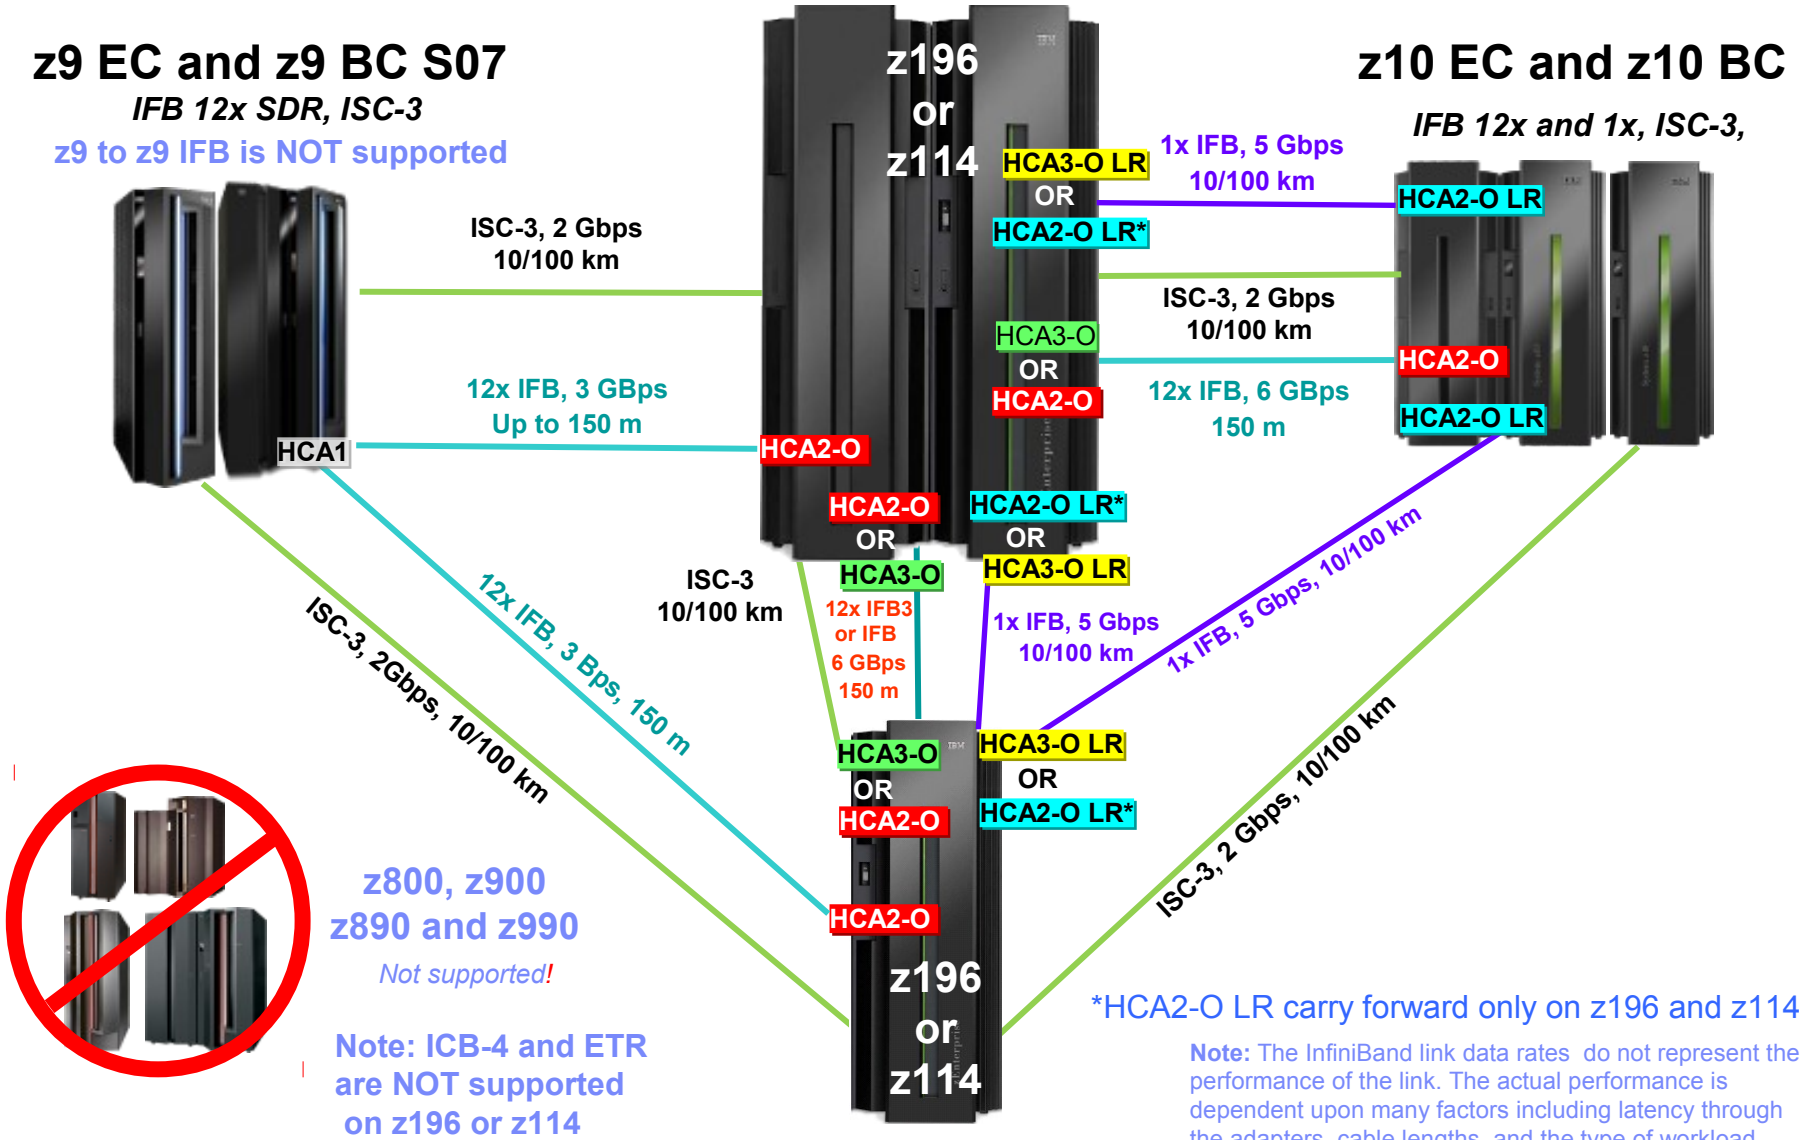

System z Crypto History

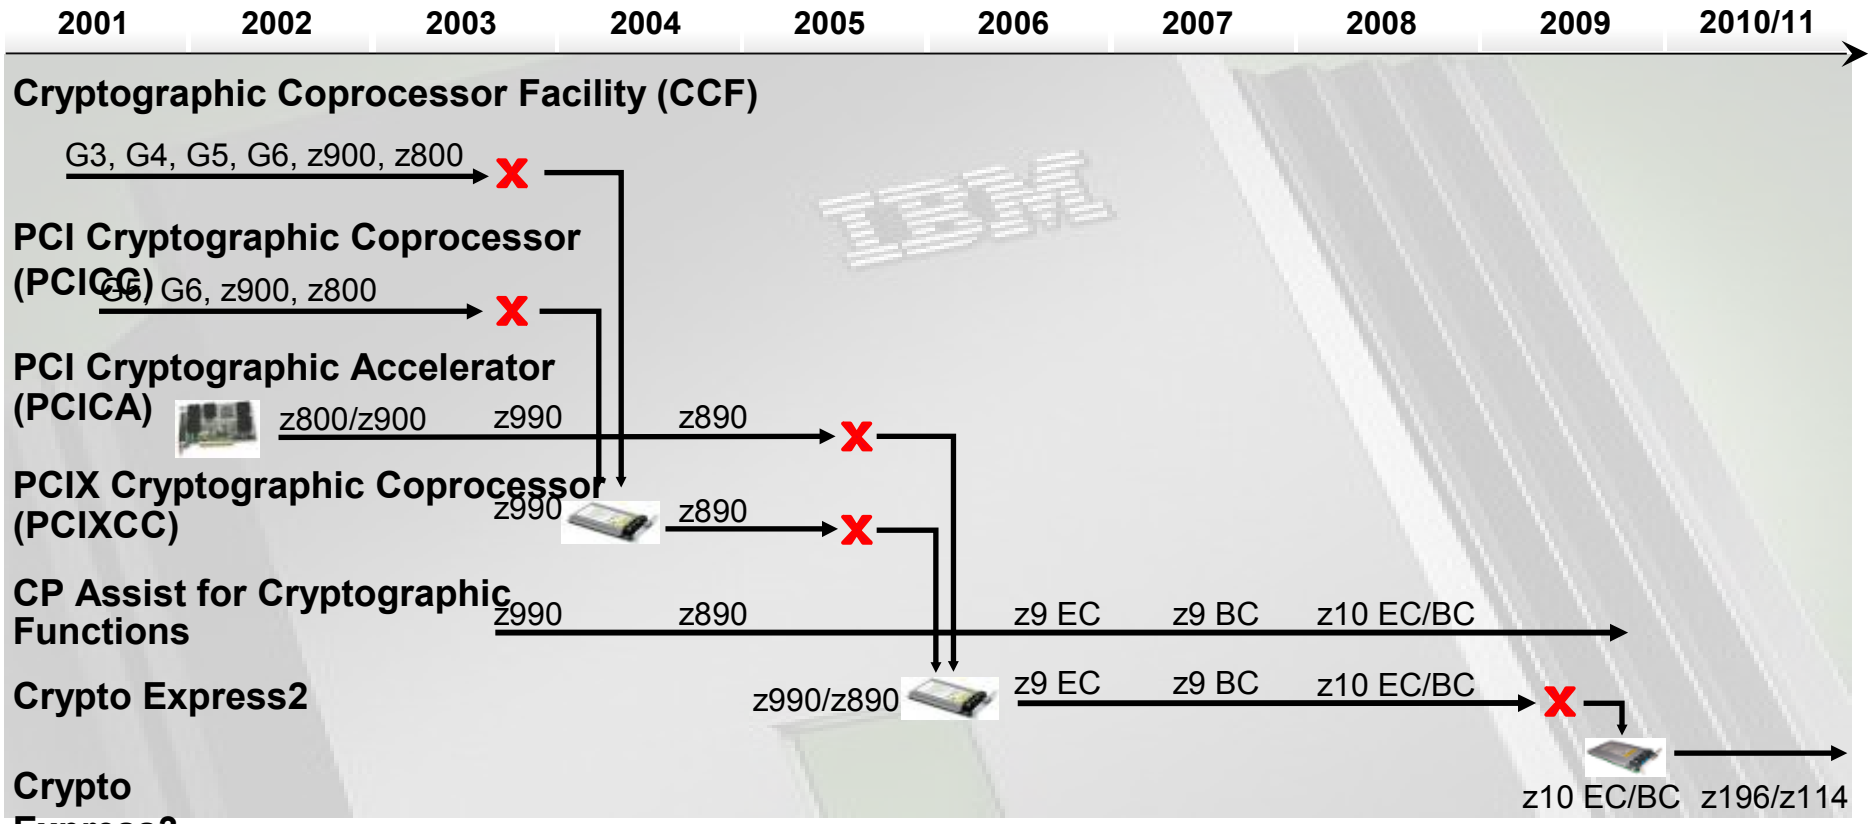

- **Cryptographic Coprocessor Facility – Supports "Secure key" cryptographic processing**
- **PCICC Feature – Supports "Secure key" cryptographic processing**
- **PCICA Feature – Supports "Clear key" SSL acceleration**
- **PCIXCC Feature – Supports "Secure key" cryptographic processing**
- **CP Assist for Cryptographic Function allows limited "Clear key" crypto functions from any CP/IFL**
  - NOT equivalent to CCF on older machines in function or Crypto Express2 capability
- **Crypto Express2 – Combines function and performance of PCICA and PCICC**
- **Crypto Express3 – PCI-e Interface, additional processing capacity with improved RAS**

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, zAAPs, 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 to satisfy periods of peak demand for computing resources
  - 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

Optional High Voltage DC Power

- **Using high voltage DC power can save power by eliminating DC to AC and AC to DC conversion losses**
- **The bulk power in z196/z114 is modified to support HVDC so the only differences in shipped HW to implement the option are the DC line cords:**
  - This adds DC line cord feature codes
  - Expect that nominal DC supply voltages supported will include:
    - 570V (500 - 600V)
    - 380V (330- 420V)

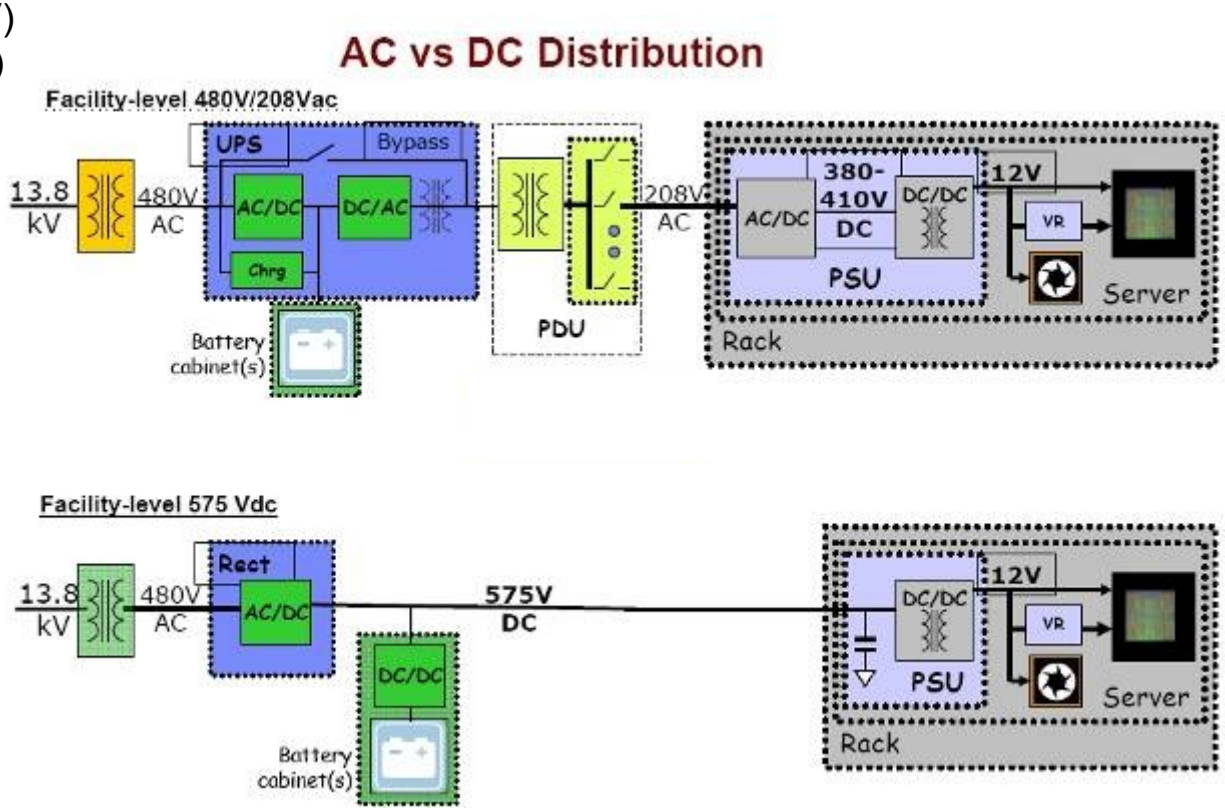

z/OS Support Summary

Release	z9 EC WdfM	z9 BC WdfM	z10 EC	z10 BC	z196	z114	End of Service	Coexists with z/OS
z/OS V1.8	X	X	X	X	X	X	9/09[1]	V1.10
z/OS V1.9	X	X	X	X	X	X	9/10[1]	V1.11
z/OS V1.10	X	X	X	X	X	X	9/11[2]	V1.12
z/OS V1.11	X	X	X	X	X	X	9/12*	V1.13*
z/OS V1.12	X	X	X	X	X	X	9/13*	V1.14*
z/OS V1.13	X	X	X	X	X	X	9/14*	V1.15*

Notes:

[1]The IBM Lifecycle Extension for z/OS provides the ability for customers to purchase extended defect support for that release of z/OS for up to 24 months after the z/OS release's end of service date.

[2] z/OS 1.10 EOS is September 29, 2011. Requires Lifecycle Extension after this date

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

System z z/VM & z/VSE Support Summary

		z9 EC WdfM	z9 BC WdfM	z10 EC	z10 BC	z196	z114	Ship Date	End of Market	End of Service
z/VSE*	4.2	x	x	x	x	x	X	10/08	TBD	TBD
	4.3	x	x	x	x	x	X	4Q10	TBD	TBD
	5.1*	x	x	x	x	x	X	4Q11	TBD	TBD
z/VM	5.4	x	x	x	x	x	X	09/08	TBD	9/13*(1)
	6.1	No	No	x	x	x	x	10/09	TBD	4/13*

Note: z/VSE V4 is designed to exploit 64-bit real memory addressing, but does not support 64-bit virtual addressing. z/VSE V5.1 has been previewed.

Crypto Express3 requires z/VSE V4.2 or later.

1) End of Service date for z/VM 5.4 extended

* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Note: z/VM 6.1 Implements a new Architecture Level Set (ALS) available only on IBM System z10 servers or later

System z Linux Support






	z9 EC WdfM	z9 BC WdfM	z10 EC	z10 BC	z196	z114	Availability Date
RHEL 5	x	x	x	x	x	x	03/2007
RHEL 6	x	x	x	x	x	x	11/2010
SLES 10	x	x	x	x	x	x	08/2006
SLES 11	x	x	x	x	x	x	03/2009

	End of Production Ph 1	End of Production Ph 2	End of Production Ph 3
RHEL 5 support*	4Q 2011	4Q 2012	03/312014
RHEL 6 support*	4Q 2014	4Q 2015	11/30/3017

	General support	Extended support	Self support
SLES 10 support*	07/31/2013	07/31/2016	07/31/2016
SLES 11 support*	03/31/2016	03/31/2019	03/31/2019

- **For latest information and details contact your Linux distributor**
- Recommendation: use RHEL 6 or SLES 11 for new projects
- For latest information about supported Linux distributions on System z refer to: http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html

* SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux
Support dates may be changed by Linux distributors

zEnterprise 114 in a nutshell:

- **Integration with the zEnteprise BladeExtension**
- **Two Models**
- **Increased capacity in a single footprint**
  - 12s0 technology
  - out-of-order instruction processing
  - higher clock frequency
  - larger cache
- **Robust Memory**
- **Upgrades**
  - Investment protection with upgrades from two previous families
    - z10 BC
    - z9 BC
  - Upgradeability to z196 (M15)
- **I/O Improvements**
  - new I/O features and New I/O Drawer

IBM zEnterprise 114 (z114) and z196 GA2 Key Dates for Hardware

- **Planned Availability Dates:**
  - **September 09, 2011 – Driver 93**
    - IBM zEnterprise 114 Models M05 and M10 new builds
    - Field installed features and conversions that are delivered solely through a modification to the machine's Licensed Internal Code (LIC):
      - LIC only upgrades for any PU type and memory
      - LIC for CBU and On/Off Capacity upgrade for installed z114s
      - z114 Capacity Setting downgrades
    - IBM System z9® Business Class (z9® BC) upgrades to z114
    - IBM System z10™ Business Class (z10 BC™) upgrades to z114
    - NEW z196s at GA2 level
      - MES for new function/features for installed z196s
    - TKE Support for LIC 7.1
  - **September 26, 2011**
    - Unified Resource manager functions:
      - Manage suite (#0019) enhancements
      - Automate/Advanced Management Firmware Suite (#0020) enhancements
      - Manage Firmware System x Blade (#0042)
      - Advanced Management Firmware System x Blade (#0046)
  - **October 21, 2011**
    - Add a zBX as an MES to the installed z114
  - **December 31, 2011**
    - z114 Model M05 to M10 upgrades
    - Feature adds (memory, I/O, RPQs and zBX features)
  - **Future***
    - Microsoft Windows (64-bit only) for select IBM System x Blades in a zBX

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

Thank You!

zEnterprise Backups

What is a zEnterprise Ensemble?

- A zEnterprise ensemble is a collection of 1 to 8 z196/z114 CPCs with/without zBX managed collectively by the Unified Resource Manager as a single logical virtualized system using the HMC
- A zEnterprise node is a z196/z114 CPC with 0 to 4 racks up to 2 BladeCenters per rack
- Blade based fit-for-purpose Solutions
- Integrated Advanced Virtualization Management
- Implements well-defined external interface to Data Center Service Management functions
- Virtual Resource Management and Automation

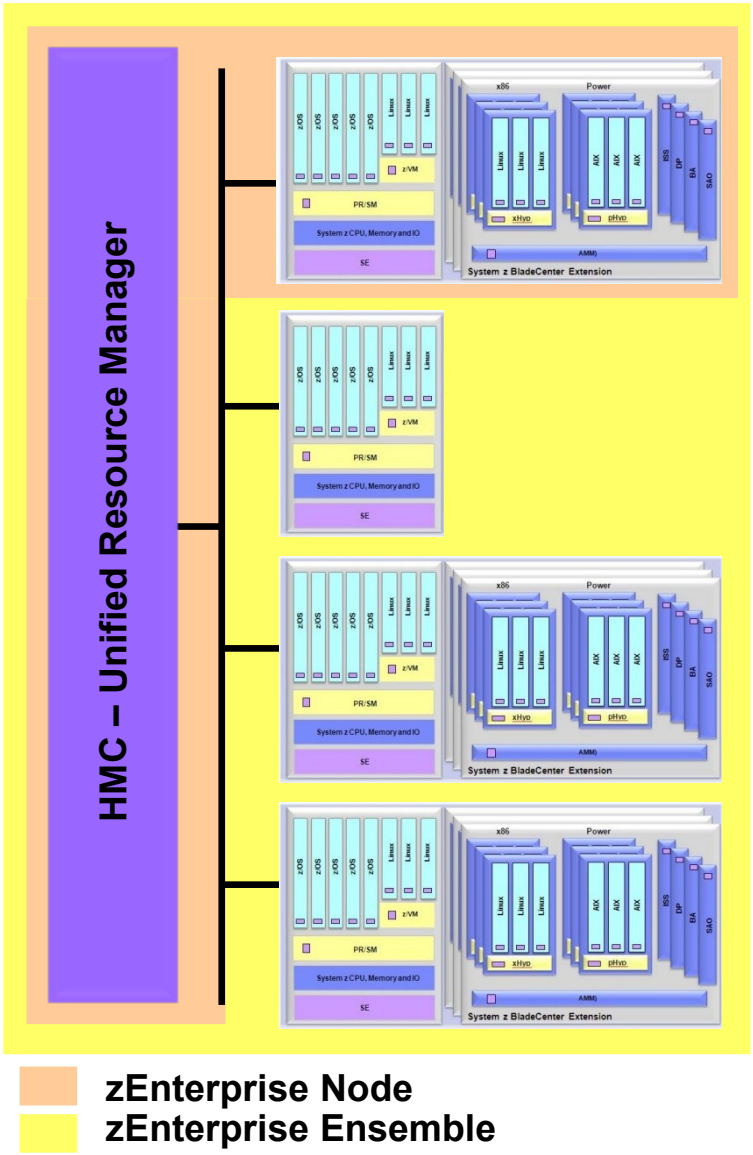

zEnterprise – What are the INMN, IEDN and Customer networks

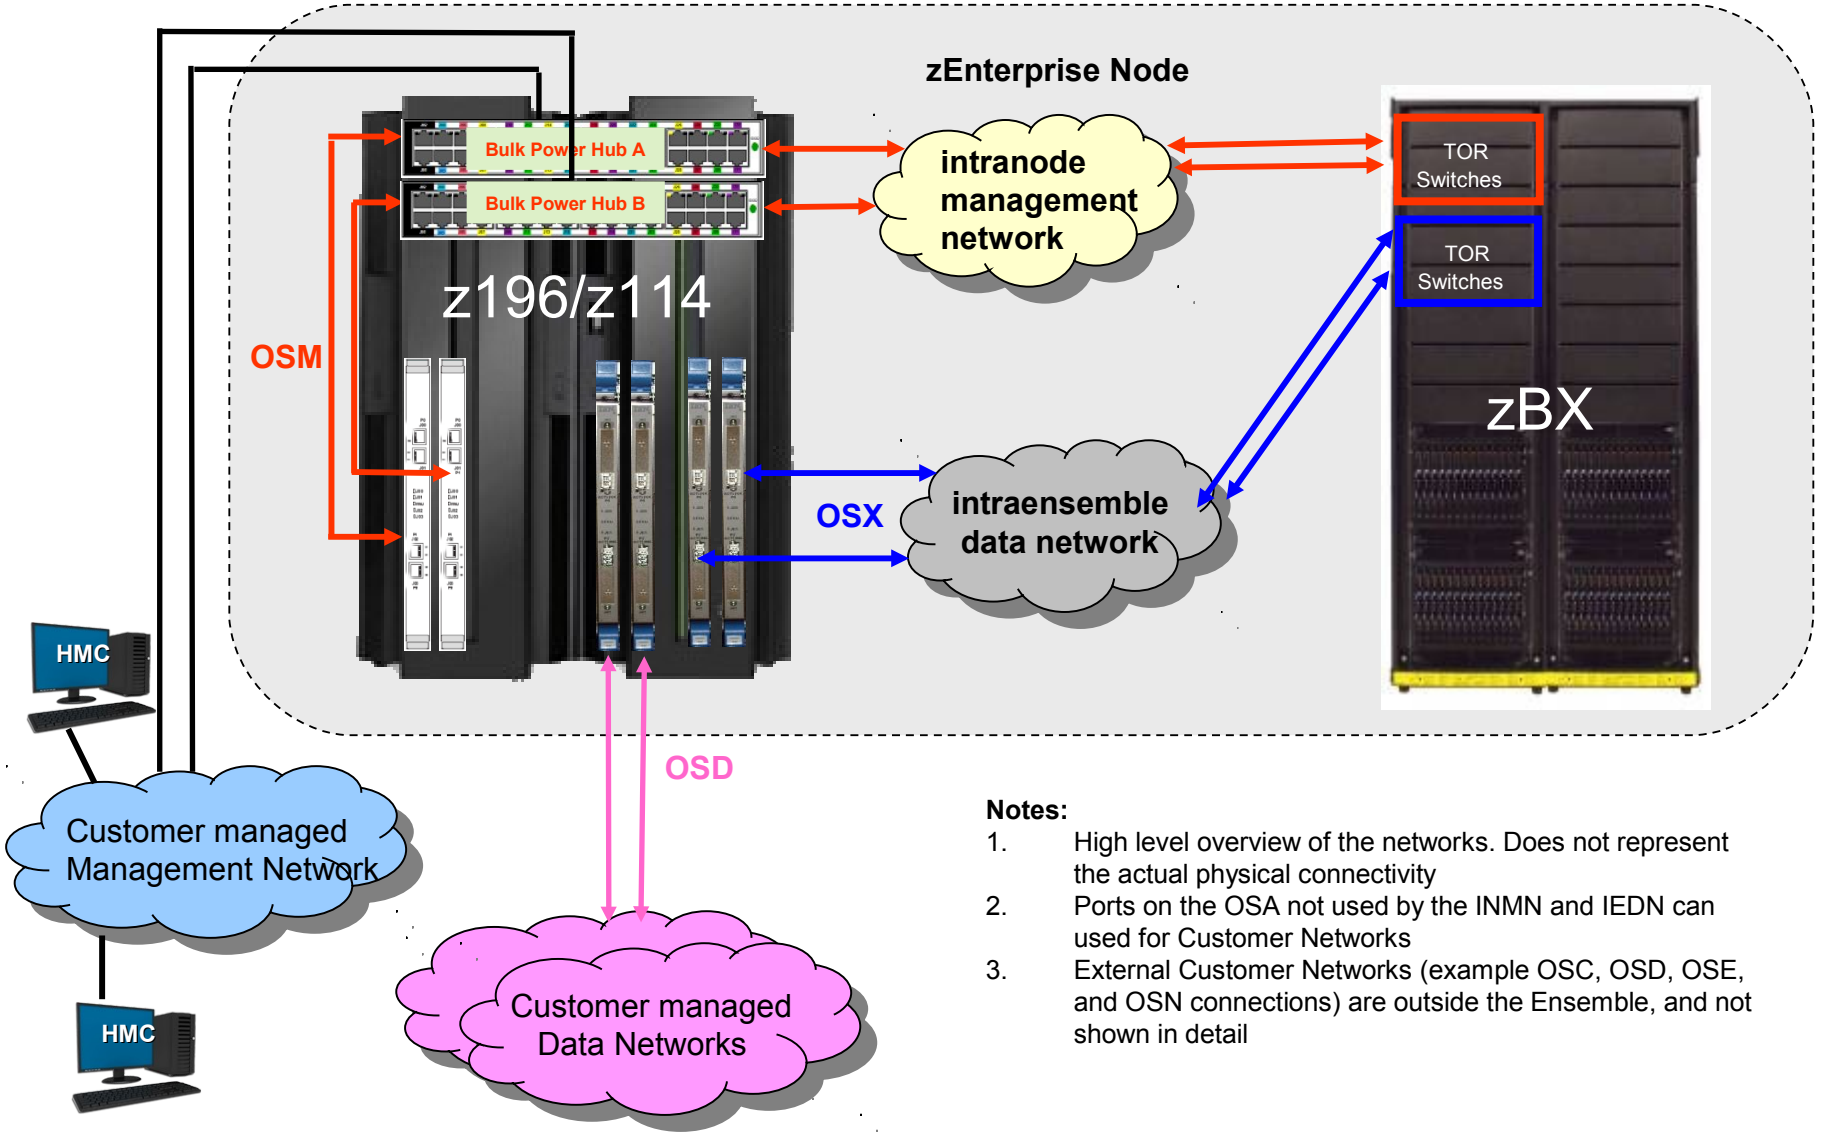

Notes:

1. High level overview of the networks. Does not represent the actual physical connectivity
2. Ports on the OSA not used by the INMN and IEDN can used for Customer Networks
3. External Customer Networks (example OSC, OSD, OSE, and OSN connections) are outside the Ensemble, and not shown in detail

zBX Hardware Details

zBX Infrastructure

- Rack
- Top-of-Rack Switch
- Blade Center Chassis
- Ethernet & FC Cables
- Switches (HSS, ESM, FC)
- Power Dist. Units
- Opt: Heat Exchanger, Power cord types

Blades

IBM
Smart Analytics Optimizer

POWER7 Blades

IBM System x Blades*

DataPower XI50z

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

Putting zEnterprise System to the task

Use the smarter solution to improve your application design

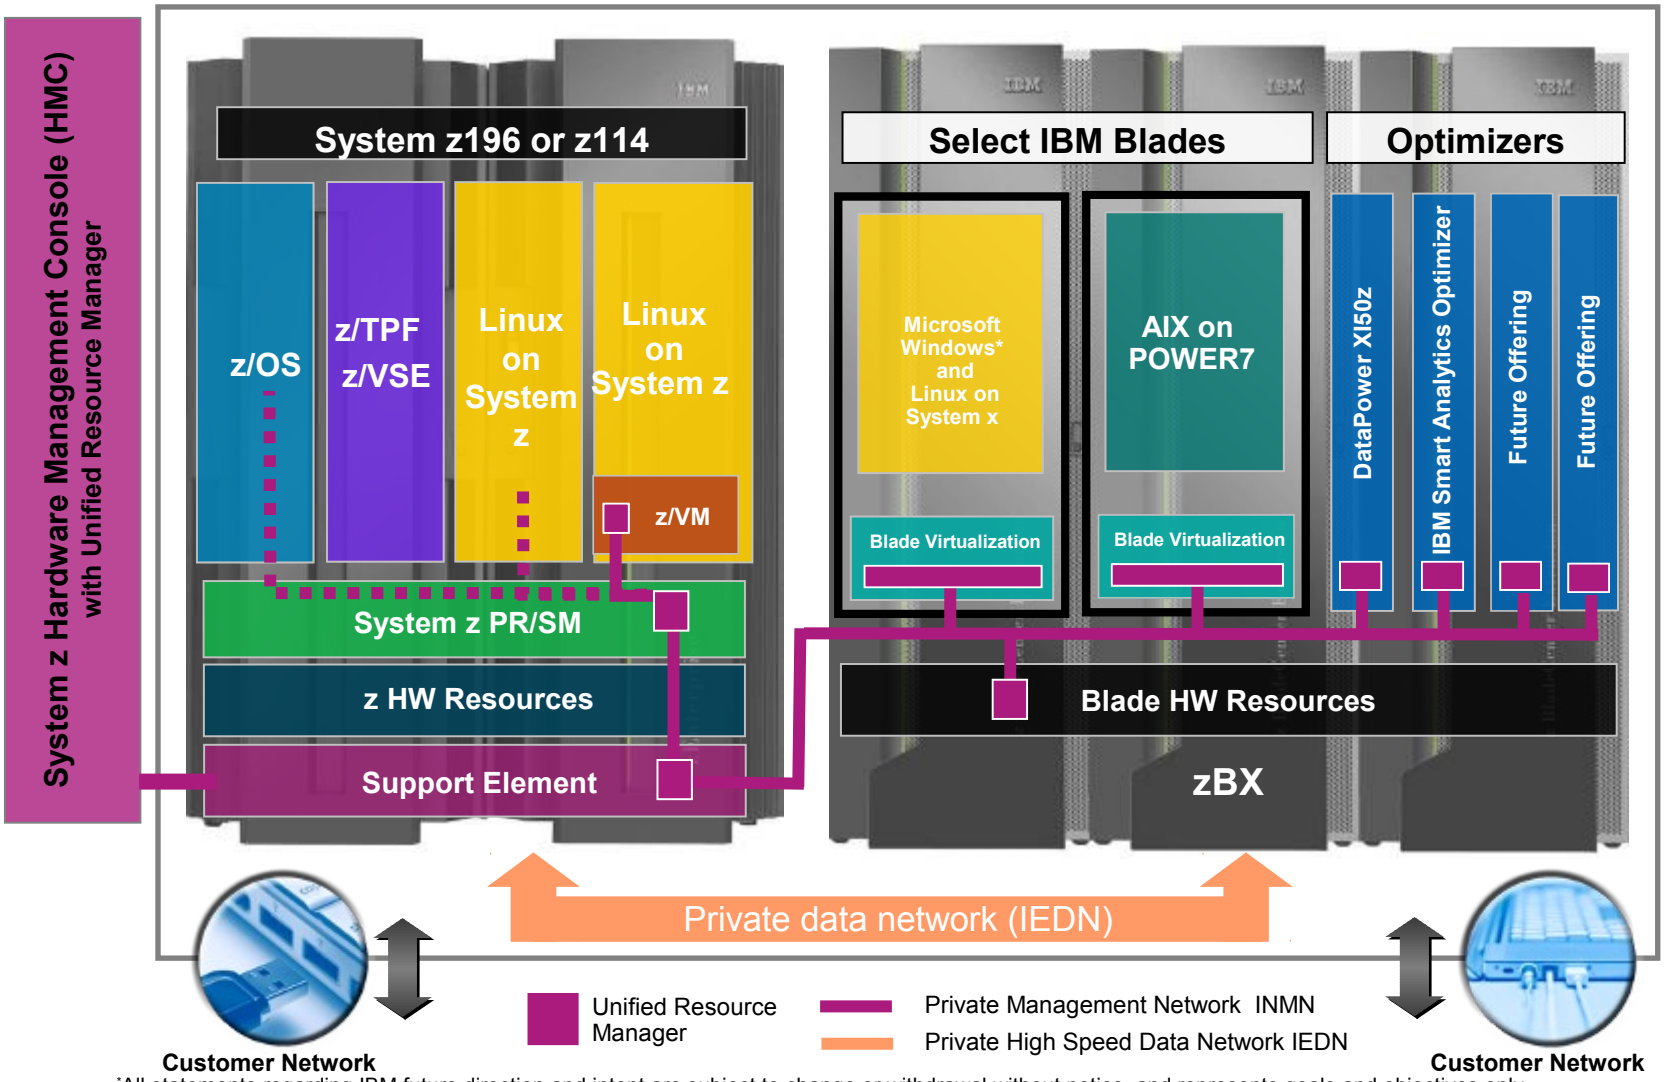

*All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

This Library of charts has been created with the assistance of various IBM staff in Poughkeepsie and the ATS organisation in Gaithersburg by:

Parwez Hamid

Executive Certified IT Specialist

IBM UK Limited

+44 (20) 8818 4711

parwez_hamid@uk.ibm.com

Any comments/corrections, please send these to Parwez.

IBM

What is IBM zEnterprise System?

Re-write the rulebook and set new standards for business-centric IT with IBM System z, to be the world's premier workload-optimized platform for enterprise applications.

Our Vision:

Deliver the best of all worlds - Mainframe, UNIX®, x86 and single function processors - integrated in a single system for ultimate flexibility and simplicity to optimize service, risk, and cost across multiple heterogeneous workloads.

z114TLLB2

So, with future generations we will breakthrough some of the barriers that limit the way we use IT today. The objective is to move away from today's Silo-ed approach to IT where each platform is managed as a separate and distinct entity, and instead focus on managing across heterogeneous platforms to give an end-to-end solution view. This will enable IT resources and services to be properly aligned according to business viewpoint, rather than limited to an IT viewpoint

This new alignment will enable resources to be optimized to give the right level of service with appropriate risk at the right cost to the business.

IBM

IBM zEnterprise: A centrally managed enterprise cloud for the flexible delivery of high value services.

- **A highly scalable heterogeneous pool of virtualized resources managed in a single system.**
- **Activate, allocate, prioritize and retire resources on demand and automate service delivery.**
- **Maximize utilization of resources for improved ROI and for low cost service delivery.**
- **Bring new levels of security, resiliency and manageability to create a cloud environment that is enterprise ready.**

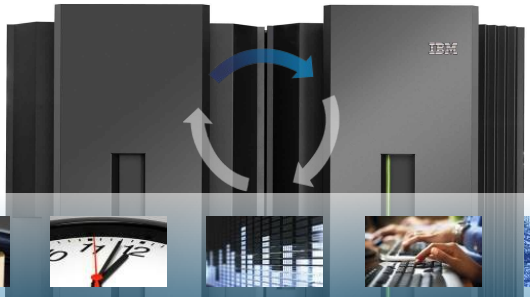

Security ***Availability*** ***Scalability*** ***Efficiency*** ***Virtualization***

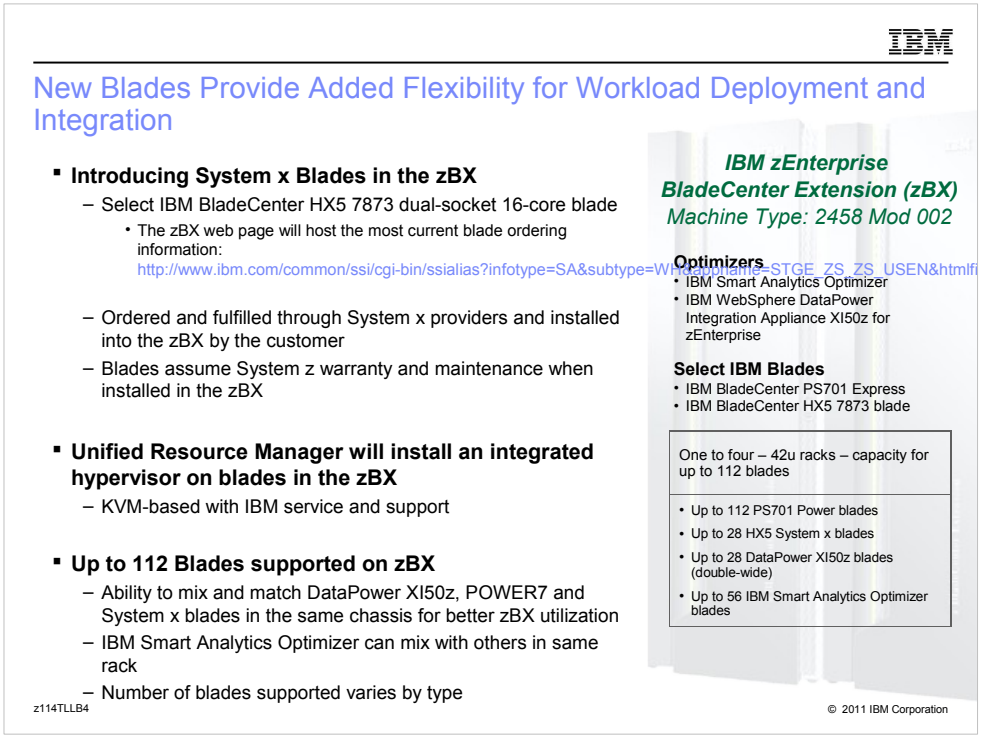

On July 12 we announced the introduction of System x Blades installed into the zBX. We will be supporting select configurations of the IBM BladeCenter HX5 7873 dual-socket 16-core blade. It is critical that only the supported HX5 7873 configurations be ordered for installation into the zBX. The zBX webpage (http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfid=ZSL03128USEN&attachment=ZSL03128USEN.PDF) will host the supported configurations (just like it has hosted the PS701 configurations). Like the POWER7 blades, these blades will be ordered and fulfilled through System x providers.) They can not be ordered using the zBX configurator. Just like the POWER7 blades, they are installed by you (or you could create a service contract with IBM to install), but once in there, support and warranty on the blades takes on z characteristics and will be provided by the IBM System z Service Support Rep.

The IBM Unified Resource Manager will auto detect the blades and then install an integrated hypervisor on the BladeCenter HX5 – InxHY. This hypervisor is KVM based. The hypervisor will have IBM service and support. We can not put this on the charts, and ask that you not share this unless asked… but in one on one discussions you can share that the INxHY will be 'RHEV-H 5.5 based on KVM'.

The zBX can host up to 112 blades. You have the ability to mix and match blades and optimizers in the zBX for better utilization. The blades for BladeCenter PS701 Express blade, BladeCenter HX5 blade and DataPower XI50z can be shared in the same BladeCenter chassis—note that DataPower XI50z blades are "doublewide" and use two slots. They can be mixed in the same zBX with IBM Smart Analytics Optimizer blades—but not in the same BladeCenter Chassis. Remember that the total zBX capacity can not exceed 112 total blades.

The number of blades supported varies by type – seen on the right side of the chart:

- Up to 112 PS701 Power blades
- Up to 28 HX5 System x blades
- Up to 28 DataPower XI50z blades (double-wide)
- Up to 56 IBM Smart Analytics Optimizer blades

Transition: And new blades offer new operating system support for more application choices.

IBM zEnterprise 114 (z114)

z114 Overview

- **Machine Type**
  - 2818
- **2 Models**
  - M05 and M10
  - Single frame, air cooled
  - Non-raised floor option available
  - Overhead Cabling and DC Power Options
- **Processor Units (PUs)**
  - 7 PU cores per processor drawer (One for M05 and two for M10)
  - Up to 2 SAPs per system, standard
  - 2 spares designated for Model M10
  - Dependant on the H/W model - up to 5 or 10 PU cores available for characterization
    - Central Processors (CPs), Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z Integrated Information Processor (zIIP), optional - additional System Assist Processors (SAPs)
  - 130 capacity settings
- **Memory**
  - Up to 256 GB for System including HSA
    - System minimum = 8 GB (Model M05), 16 GB (Model M10)
    - 8 GB HSA separately managed
    - RAIM standard
    - Maximum for customer use 248 GB (Model M10)
    - Increments of 8 or 32 GB
- **I/O**
  - Support for non-PCIe Channel Cards
  - Introduction of PCIe channel subsystem
    - Up to 64 PCIe Channel Cards
  - Up to 2 Logical Channel Subsystems (LCSSs)
- **STP - optional** (No ETR)

z114TLLB6

Continued qualities of service: The z114 continues the heritage in mainframe qualities of service with extreme granularity. structure and significant improvements in packaging, performance, and total system scalability over prior generations.

More processing capacity and performance: System resources are powered by up to 14 microprocessors running at 3.8 GHz. The z114 provides up to an 18% improvement in uniprocessor performance and up to a 12% increase in total system capacity for z/OS, z/VM, z/VSE, and Linux on System z workloads as compared to its predecessor, the z10 BC.

Its superscalar microprocessor chip has a higher-frequency design that leverages IBM technology leadership with a new out-of-order execution sequence that delivers world-class per-thread performance. With over 100 new instructions and numerous compiler related enhancements, the z114 can deliver up to 25% performance improvement, based on measurements and projections, for CPU Intensive workloads when accompanied by multiple C/C++ compiler level improvements going from XL C/C++ V1R9 to V1R12.

Scalability and flexibility for growth: The z114 will be available in two models: a single central processing drawer model, the M05, and a two drawer model, the M10, which offers the additional flexibility for I/O and coupling expansion and/or increased specialty engine capability. With up to 10 configurable cores, the model naming is indicative of how many total processor units are available for user characterization. The cores can be configured as general purpose processors (CPs), Integrated Facilities for Linux (IFLs), System z Application Assist Processors (zAAPs), System z Integrated Information Processors (zIIPs), Internal Coupling Facilities (ICFs), additional System Assist Processors (SAPs), or can be used as additional spares (M10 only).

With 130 available capacity settings and a granular cost structure offered across either model, you have the freedom to choose the right capacity setting for your needs with the flexibility to scale on demand as workload demands increase.

To help secure sensitive data and business transactions, the z114 is designed for Common Criteria Evaluation Assurance Level 5 (EAL5) certification for security of logical partitions. Support for the next generation of public key technologies is available with Elliptic Curve Cryptography (ECC), which is ideal for constrained environments such as mobile devices. The z114 also offers support for key ANSI and ISO standards for the banking and finance industry.

256 GB memory for application growth: The z114 will support up to 248 GB of real (usable) RAIM-protected memory, an industry exclusive currently available only on System z. Beyond the purchased memory, there is an additional 8 GB of memory for the hardware system area (HSA), which holds the I/O configuration data for the server.

High-speed access to your resources: High-speed connectivity is critical in achieving sufficient levels of transaction

zEnterprise 114 Models M05 and M10

- **M/T 2818 – Model M05**
  - Air cooled
  - Single Frame
  - Non-raised floor option available
  - 30 LPARs
- **Processor Units (PUs)**
  - New processor drawer design (1 CPC Drawer)
  - 7 per system
    - 2 SAPs standard
    - Up to 5 CPs
    - Up to 5 specialty engines
    - Up to 2 zIIPs/zAAPs
    - 0 spares when fully configured

- **M/T 2818 – Model M10**
  - Air cooled
  - Single Frame
  - Non-raised floor option available
  - 30 LPARs
- **Processor Units (PUs)**
  - New processor drawer design (2 CPC Drawers)
  - 14 per system
    - 2 SAPs standard
    - Up to 5 CPs
    - Up to 10 specialty engines
    - Up to 5 zIIPs/zAAPs
    - 2 dedicated spares

- **When Model M10 (requires the 2nd processor drawer)?**
  - > 5 Customer PUs
  - > 120 GB memory
  - > 4 Fanouts for additional I/O connectivity – especially PSIFB links
    - Depends - numbers vary for drawers, I/O features and PSIFB links

z114 Model M10 – Under the covers

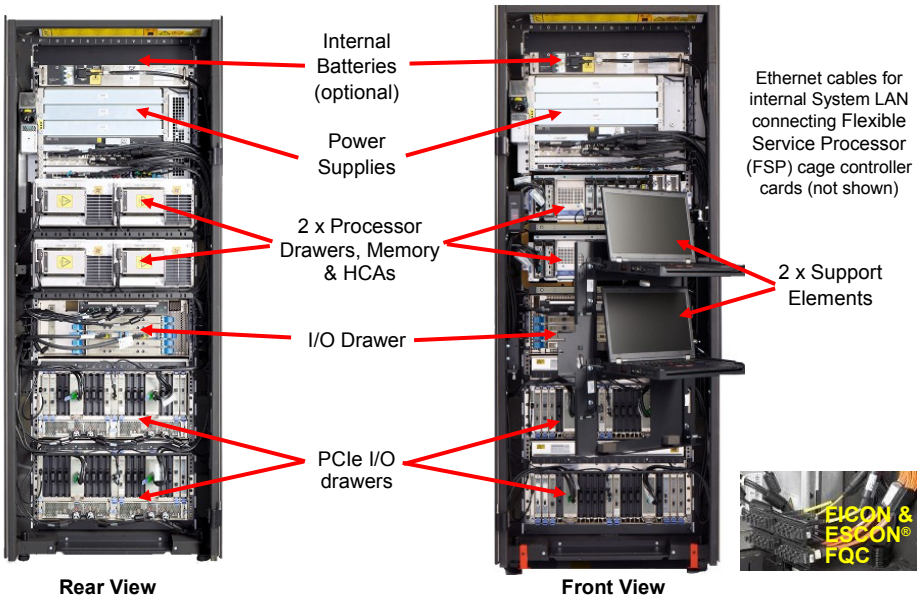

z114TLLB8

z114 Model Structure and Upgrades

Model	CPs	IFLs Unassigned IFLs	zAAPs	zIIPs	ICFs	Std. SAPs	Add'l SAPs	Spares
M05	0-5	0-5	0-2	0-2	0-5	2	0-2	0
M10	0-5	0-10	0-5	0-5	0-10	2	0-2	2

- **Model structure based on number of drawers**
- **M05 sparing based on prior Business Class (BC) offerings – no dedicated spares**
- **M10 sparing based on Enterprise Class (EC) offerings – dedicated spares**
  - SAP and PU Allocation/Sparing in the M10
    - Default assignment is one SAP per drawer; one Spare per drawer. Spill and fill CP low to high; spill and fill specialty engines high to low
    - Two defective PUs may cause the default assignment to spill and fill into the second processor drawer. LPAR has the capability to request PU of a specified type to be grouped together in a book/drawer (i.e. LPAR may change the default assignment)
- ***Disruptive* upgrade from M05 to M10**
  - No model downgrades
- **Upgrades from z9 BC and z10 BC into either model M05 or M10**
- **Only the M10 will upgrade to z196 Model M15 (Air cooled only)**

z114 Upgrade paths

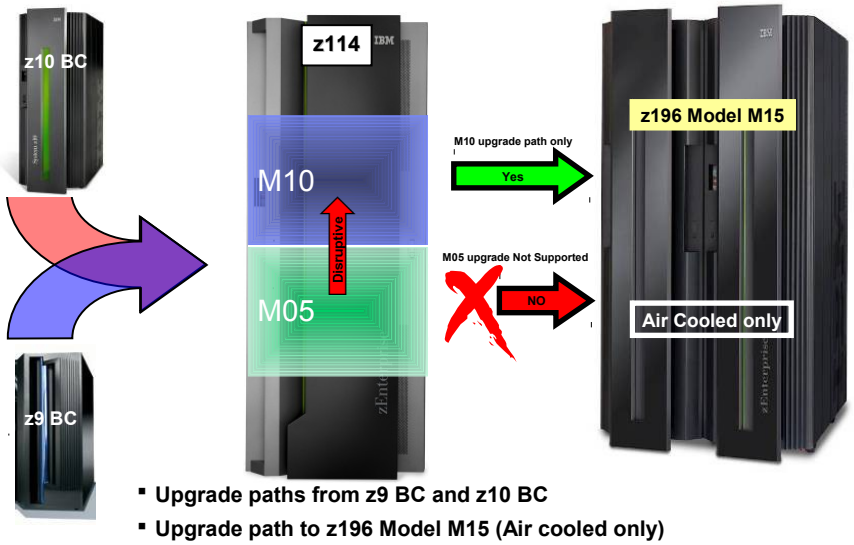

- Upgrade paths from z9 BC and z10 BC
- Upgrade path to z196 Model M15 (Air cooled only)
- Disruptive upgrade M05 to M10 and from M10 to z196 M15

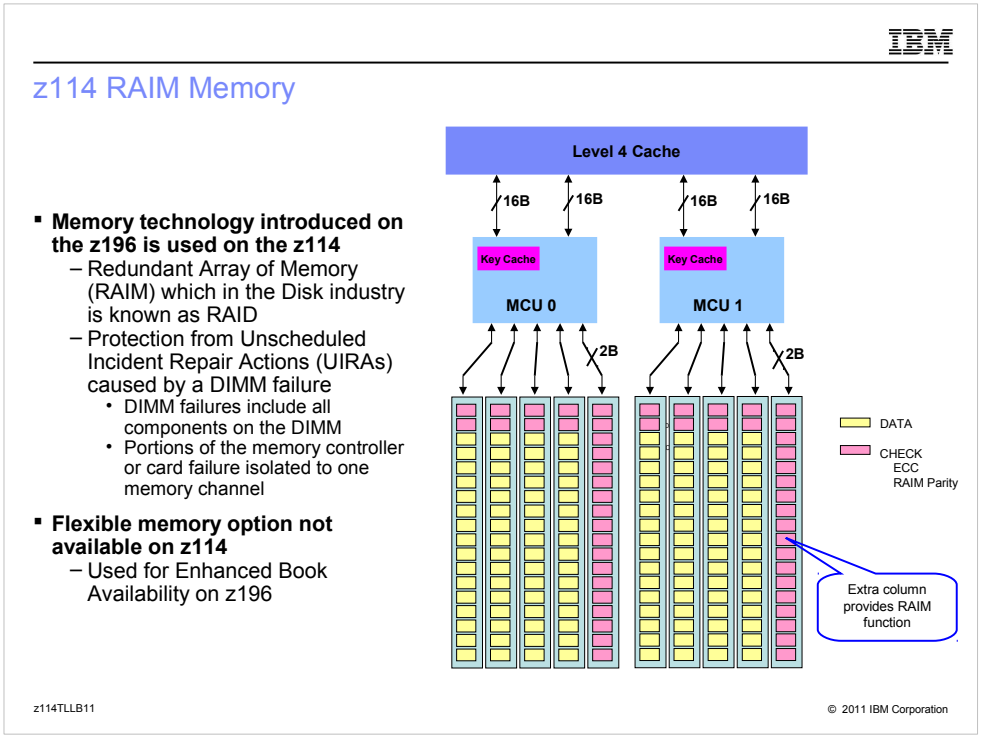

Major redesign of memory subsystem for improved availability:

IBM's most robust error correction to date can be found in the memory subsystem. A new redundant array of independent memory (RAIM) technology is being introduced to provide protection at the dynamic random access memory (DRAM), dual inline memory module (DIMM), and memory channel level. Three full DRAM failures per rank can be corrected. DIMM-level failures, including components such as the controller application specific integrated circuit (ASIC), the power regulators, the clocks, and the board, can be corrected. Memory channel failures such as signal lines and control lines can be corrected. Upstream and downstream data signals can be spared using two spare wires on both the upstream and downstream paths. One of these signals can be used to spare a clock signal line (one up stream and one down stream). Together these improvements are designed to deliver System z's most resilient memory subsystem to date.

System z HSA Comparisons

- **HSA significantly larger than pre-z10 Servers**
- **Fixed HSA, does not affect customer purchased memory**
- **Size of HSA on prior Servers (dependant on defined configuration)**
  - Multiprise® 2000 From 12 MB up to 40 MB
  - 9672 G4 From 48 up to 64 MB
  - Multiprise 3000 From 38 MB up to 136 MB
  - 9672 G5/G6 From 64 MB up to 192 MB
  - z800 From 160 MB up to 256 MB
  - z900 From 288 MB up to 512 MB
  - z890 From 768 MB up to 1.9 GB
  - z990 From 1 GB MB up to 2 GB
  - z9 BC From 896 MB up to 2.7 GB
  - z9 EC From 1.2 GB up to 4.2 GB
  - z10 BC 8 GB – Fixed
  - z10 EC 16 GB – Fixed
  - z196 16 GB – Fixed
  - z114 8 GB - Fixed
- **HSA Estimator on Resource Link not relevant**

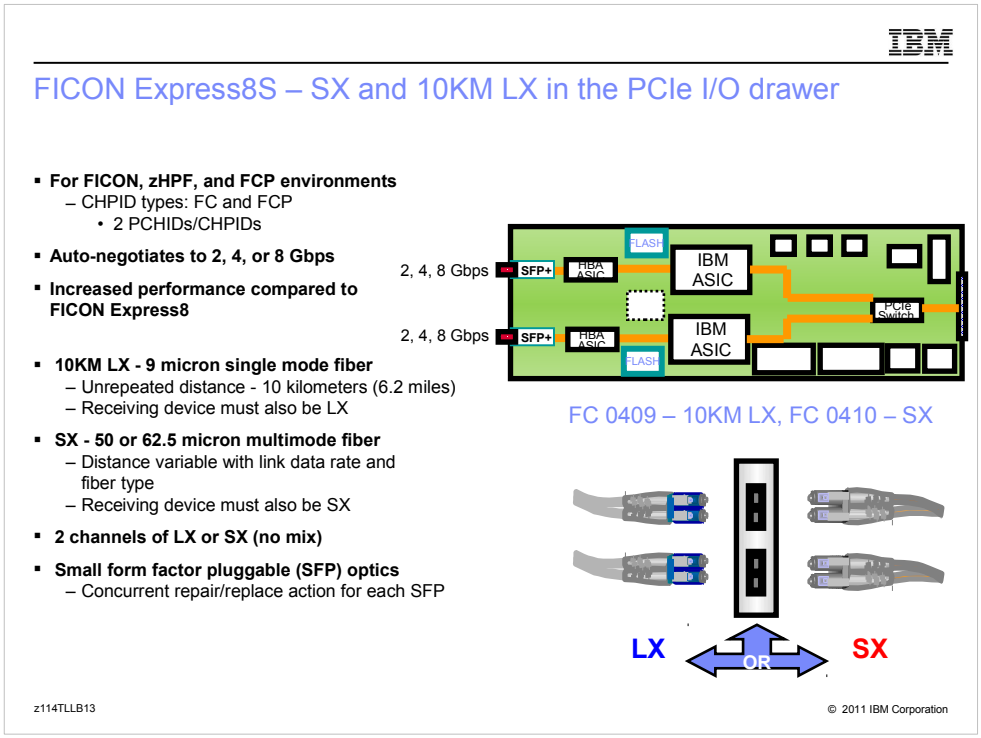

FICON Express8S - a new generation for FICON, zHPF, and FCP:

A new generation of features for the storage area network (SAN) is being introduced in support of the PCIe 8 GBps host bus and the PCIe I/O drawer. The new features for the multimode and single mode fiber optic cabling environments have path length reductions for High Performance FICON for System z (zHPF) and Fibre Channel Protocol (FCP), increased start I/Os, improved throughout for zHPF and FCP with the introduction of a hardware data router, and increased port granularity - two channels/ports per feature.

New design for increased performance for zHPF and FCP: FICON Express8S contains a new IBM ASIC which is designed to support the 8 GBps PCIe interface to the PCIe I/O drawer and increased start I/Os. In addition, a hardware data router has been added in support of the zHPF and FCP protocols for path length reduction and increased throughput. FICON Express8S supports a link data rate of 2, 4, or 8 Gbps autonegotiated.

With these changes FICON Express8S, when supporting the zHPF or FCP protocols, has been designed to achieve full duplex line speed - 8 Gbps - in each direction.

Increased port granularity: The FICON Express8S 10KM LX and SX features for single mode and multimode fiber optic cabling environments each now have two channels/ports per feature versus the four channels per feature for the FICON Express8 features. This design helps facilitate purchasing the right number of ports to help satisfy your application requirements and to better optimize for redundancy.

The FICON Express8S features, supporting CHPID types FC (zHPF, FICON, channel-to-channel) and FCP, are exclusive to the z196 and z114. They are for use exclusively in the PCIe I/O drawer and are supported by z/OS, z/VM, z/VSE, z/TPF, and Linux on System z

OSA-Express4S GbE and 10 GbE fiber for the PCIe I/O drawer

- **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

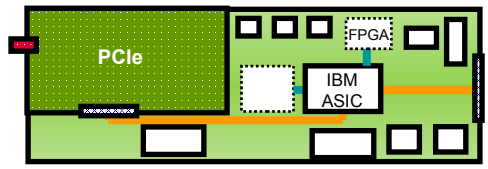

FC 0406 – 10 GbE LR, FC 0407 – 10 GbE SR

- **Gigabit Ethernet (GbE)**
  - CHPID types: OSD (no CHPID=OSN)
  - Single mode (LX) or multimode (SX) fiber
  - Two ports of LX or two ports of SX
    - 1 PCHID/CHPID

- **Small form factor optics – LC Duplex**

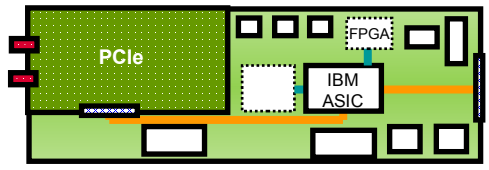

FC 0404 – GbE LX, FC 0405 – GbE SX

OSA-Express3 1000BASE-T 'copper' for I/O Drawer

- **Auto-negotiation to 10, 100, 1000 Mbps**
- **Double the port density of OSA-Express2**
- **Reduced latency & improved throughput**
  - Ethernet hardware data router
- **Improved throughput – standard & jumbo frames**
  - New microprocessor
  - New PCIe adapter
- **Usage of ports**
  - OSC, OSD, OSE can exploit four ports
  - OSM, Port 0 only
  - CHPID type OSN is for LPAR-to-LPAR communication. Does not use ports

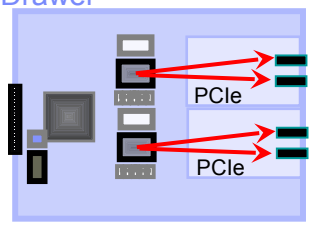

CHPID shared by two ports

4 ports - FC 3367 1000BASE-T

2 ports - FC 3369 for z10 BC and z114 only

	OSA-Express2	OSA-Express3
Microprocessor	448 MHz	667 MHz
PCI bus	PCI-X	PCIe G1

Mode	CHPID	Description
OSA-ICC	OSC	TN3270E, non-SNA DFT, IPL CPCs, and LPARs, 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
OSA for NCP (LP-to-LP)	OSN	NCPs running under IBM Communication Controller for Linux (CDLC)
Unified Resource Manager	OSM	Connectivity to intranode management network (INMN) from z196/z114 to zBX

z114TLLB15

The new OSA-Express3 10 GbE LR with double the port density designed to deliver improved connectivity to web applications.

Open Systems Adapter CHPIDs types: OSM (Express3 1000BASE-T)) and OSX (Express3 and 4S 10 GbE)

- **Two new OSA CHPID types to support new types of z196/z114 networks**
- **A z196/z114 System can have up to 6 types of OSA CHPID's**
  - External (customer managed) networks
    - Defined as OSC,OSD,OSE, & OSN
    - Existing customer provided and managed OSA ports used for access to the current customer external networks - no changes
  - Intranode management network (INMN)
    - Defined as CHPID type OSM, OSA-Express for Unified Resource Manager
      - When the PCIe adaptor on 1000BASE-T is defined as CHPID type OSM, the second port cannot be used for anything else
    - OSA-Express3 1000BASE-T configured as CHPID type OSM for connectivity to INMN from z196/z114 to Unified Resource Manager functions
    - OSA connection via the Bulk Power Hub (BPH) on the z196/z114 to the Top of the Rack (TORs) switches on zBX
  - Intraensemble data networks (IEDN)
    - Defined as CHPID OSX, OSA-Express for zBX
    - OSA-Express3 or 4S 10 GbE configured as CHPID type OSX for connectivity and access control to IEDN from z196/z114 to zBX
- **Functions Supported:**
  - Dynamic I/O support
  - HCD
  - CP Query capabilities
  - Ensemble Management for these new channel paths and their related subchannels.
- **z/VM 5.4 – CHPID types OSX and OSM cannot be varied online**

z114 continues the CMOS Mainframe heritage

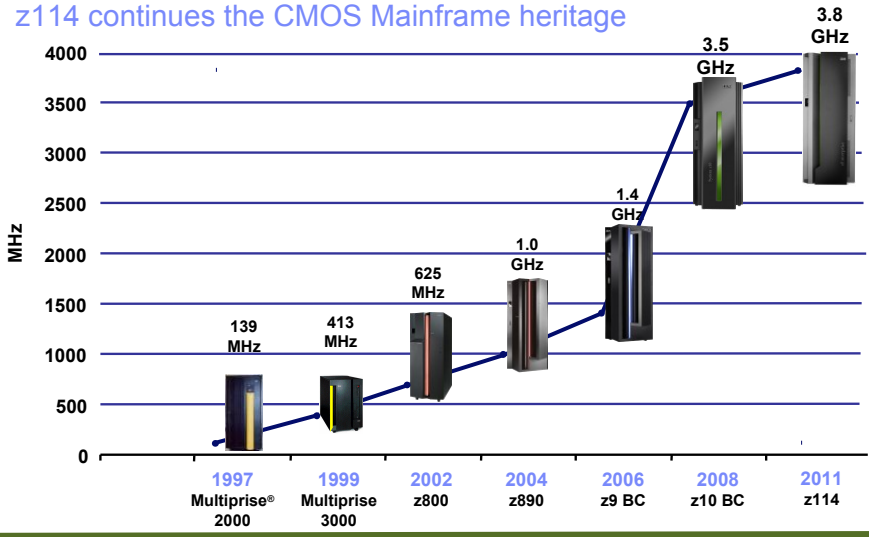

- Multiprise 2000 – 1st full-custom Mid-range CMOS S/390
- Multiprise 3000 – Internal disk, IFL introduced on midrange

- z800 - Full 64-bit z/Architecture®
- z890 - Superscalar CISC pipeline
- z9 BC - System level scaling

- z10 BC - Architectural extensions
  - Higher frequency CPU
- z114 – Additional Architectural extensions and new cache structure

z114TLLB17

zEnterprise System Out of Order (OOO) Value

- **OOO yields significant performance benefit for compute intensive apps through**
  - –Re-ordering instruction execution
    - Later (younger) instructions can execute ahead of an older stalled instruction
  - –Re-ordering storage accesses and parallel storage accesses
- **OOO maintains good performance growth for traditional apps**

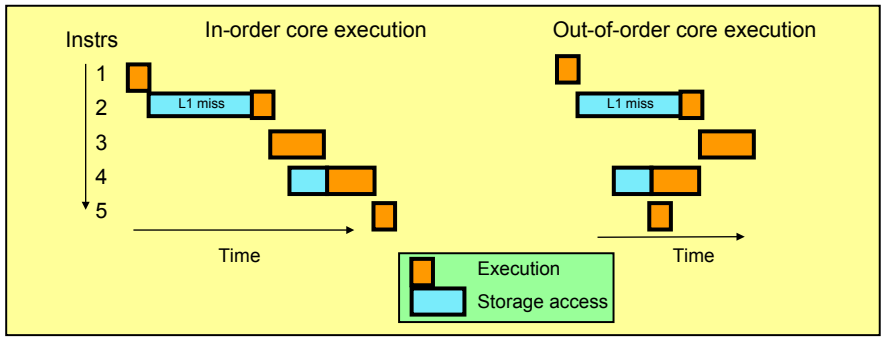

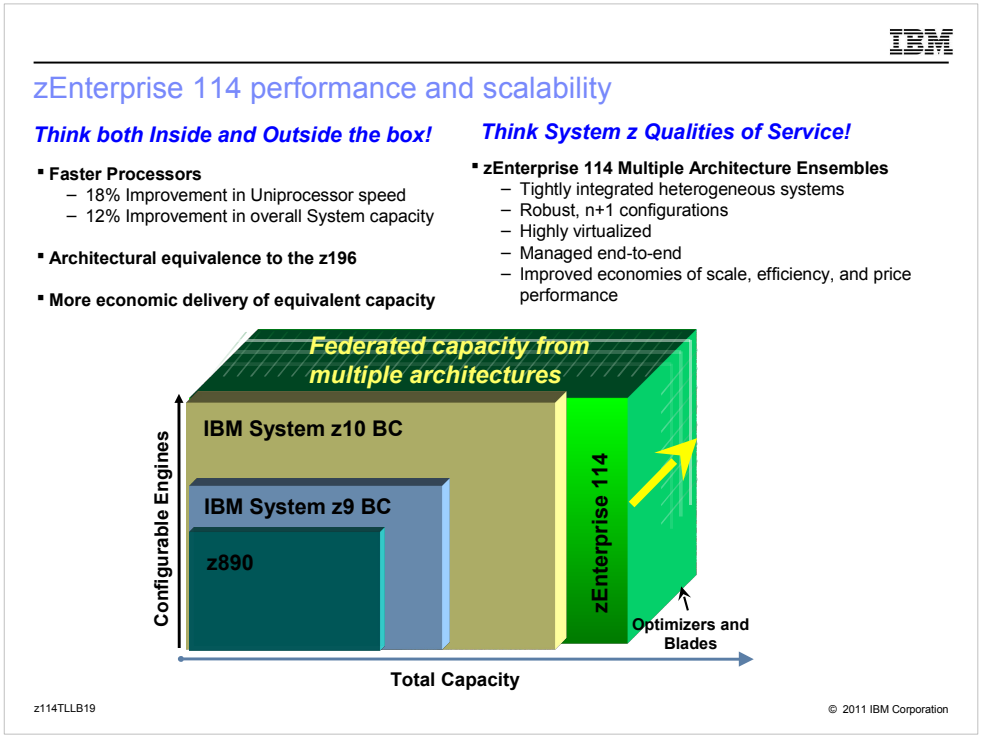

So – let 's move close to what the future for System z looks like. Given the trends we discussed and the customer requirements and pain points, the System z strategy is to acknowledge and address the challenges in the industry through the traditional System z balanced approach. The goal for zEnterprise, the next generation system, is to continue to increase the core speed while also increasing the number of cores available. The goal for zEnterprise is to increase the core speed by 1.3-1.7x depending on the type of workloads. The more CPU intensive workloads will benefit more from the faster clock speed. Although the growth is modest compared to the z10, we believe it will be very competitive industry wide given the aforementioned trends. We will couple this with larger SMP capabilities. We will proved more customer useable cores and also enhance z/OS to take advantage of the additional cores and provide larger and more efficient n-way capability. Together, they will meet the largest customer requirements as well as maintain our leadership in Single System Image scaleability.

Complementary to this increased traditional capacity, we will also add heterogeneous processing capacity. With the combination of specialized optimizers and other heterogeneous processing capabilities that we'll discuss later, we will increase the overall processing footprint of zEnterprise. This is a fundamental extension to our strategy. We will not only focus on the traditional growth of z architecture, but extend our sights into satisfying a larger part of the customer's end to end processing.

System Comparisons

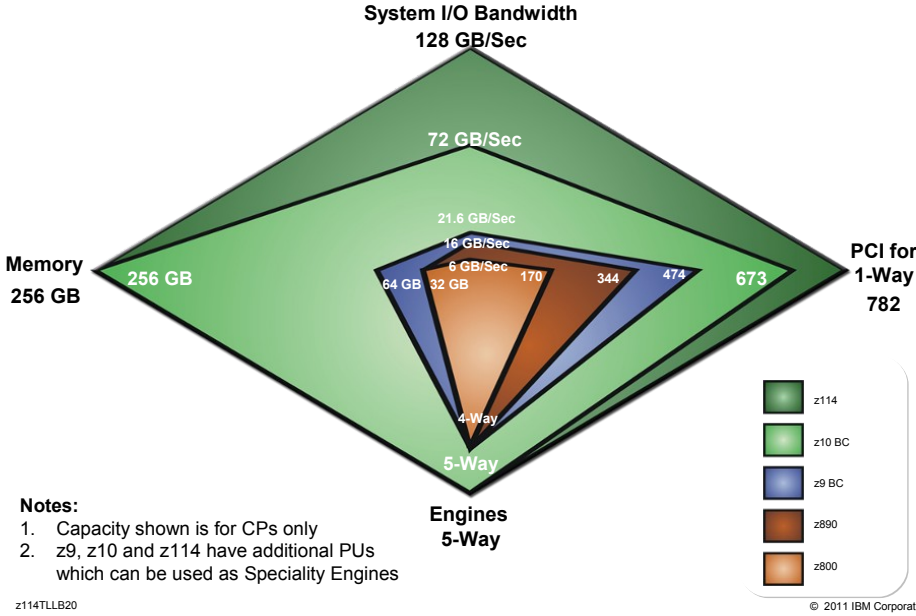

Notes:

1. Capacity shown is for CPs only
2. z9, z10 and z114 have additional PUs which can be used as Speciality Engines

z114 Capacity Matrix

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

z114

- Granularity levels similar to z10 BC to facilitate upgrades and incremental growth
- Nomenclature: **XYY**
  - **X = Capacity level**
  - **YY= Number of processors**
  - **A00 = ICF or IFL only**
- 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
- Linux only and ICF only servers
- Model M10 provides specialty engine scale out capabilities

Additional engines available on the M10

Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine	Specialty Engine

z114 Sub-capacity Processor Granularity

- **The z114 has 26 CP capacity levels (26 x 5 = 130)**
  - Up to 5 CPs at any capacity level
    - All CPs must be the same capacity level
- **The one for one entitlement to purchase one zAAP and/or one zIIP for each CP purchased is the same for CPs of any speed.**
  - All specialty engines run at full speed
  - Processor Unit Value for IFL = 100

Number of z114 CPs	Base Ratio	Ratio z10 BC to z114
1 CP	z10 BC Z01	1.18
2 CPs	z10 BC Z02	1.16
3 CPs	z10 BC Z03	1.14
4 CPs	z10 BC Z04	1.13
5 CPs	z10 BC Z05	1.12

PCI – Processor Capacity Index

IBM

System z I/O Subsystem Internal Bus Interconnect Speeds (GBps)

PCIe
z196/z114
Sep 2011

8 GBps

InfiniBand
z10/z196/z114

6 GBps

STI
z9

2.7 GBps

STI
z990/z890

2 GBps

STI
z900/z800

1 GBps

STI: Self-Timed Interconnect

z114TLLB23

z114 and z196 GA2 Parallel Sysplex Coupling Connectivity

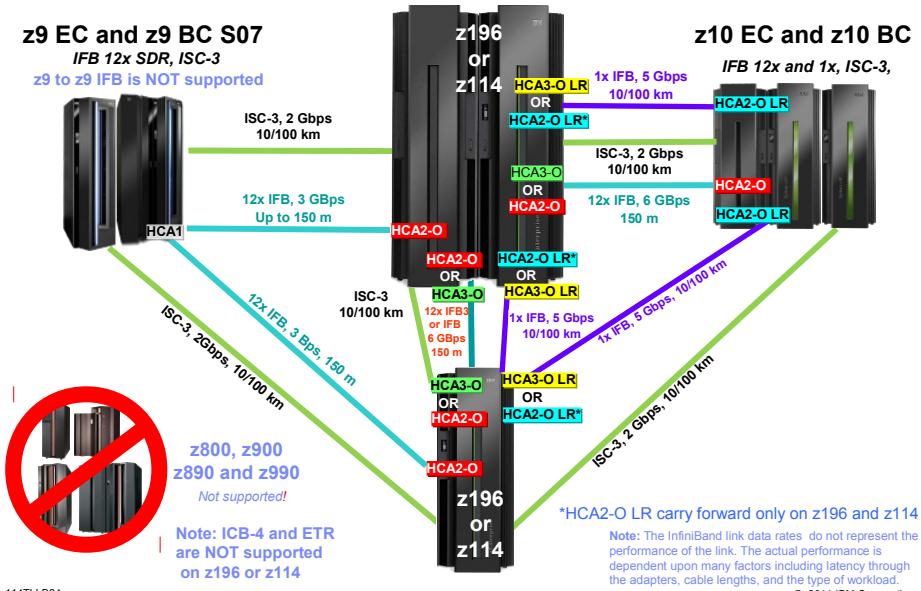

z114TLLB24

System z Crypto History

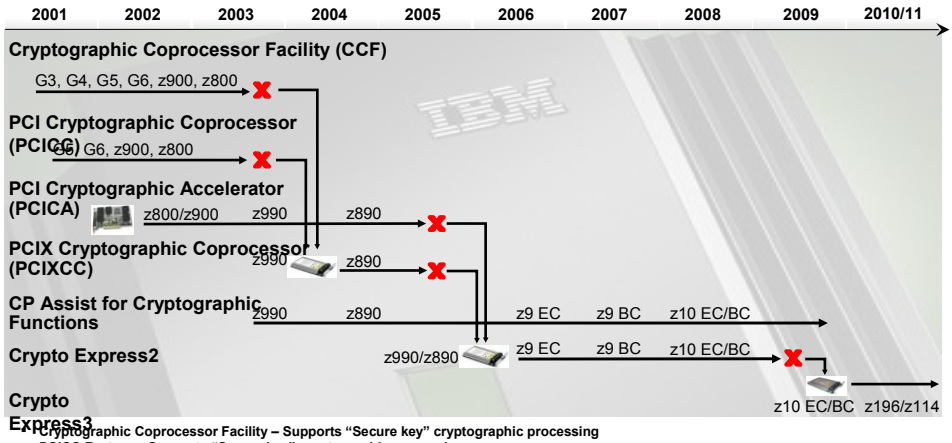

- **Cryptographic Coprocessor Facility – Supports "Secure key" cryptographic processing**
- **PCICC Feature – Supports "Secure key" cryptographic processing**
- **PCICA Feature – Supports "Clear key" SSL acceleration**
- **PCIXCC Feature – Supports "Secure key" cryptographic processing**
- **CP Assist for Cryptographic Function allows limited "Clear key" crypto functions from any CP/IFL**
  - NOT equivalent to CCF on older machines in function or Crypto Express2 capability
- **Crypto Express2 – Combines function and performance of PCICA and PCICC**
- **Crypto Express3 – PCI-e Interface, additional processing capacity with improved RAS**

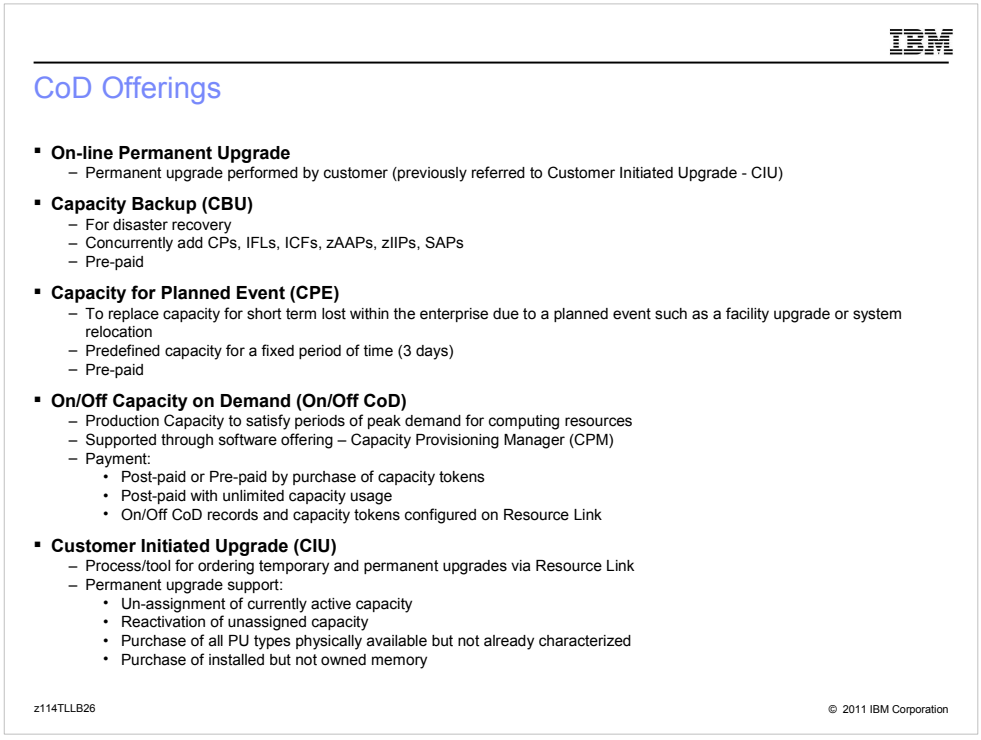

General

The existing set of contract documents which support the various Capacity on Demand offerings also support z114, and are described below.

Base Capacity on Demand (CoD) Terms (for all CoD offerings)

The base terms for all z114 CoD offerings are contained in this contract: IBM Customer Agreement Attachment for IBM System z Capacity on Demand Offerings (in the US this is form number Z125-7879). Each enterprise is required to sign this contract one time within a given country before IBM will accept an order for its first-ever instance of the enablement feature code for Capacity for Planned Events (CPE) (feature #9912), Capacity Back Up (CBU)(feature #9910), On/Off CoD (feature #9896), Permanent Upgrade authorization (feature #9898), or On-Line Ordering (feature #9900). Further, this contract document is required if an enterprise's CBU contract expires and it wishes to renew CBU.

Replacement Capacity Offering Terms

In addition to the base terms, our replacement capacity offerings (CBU and CPE) require this contract, which contains terms common to each replacement capacity offering: Customer Agreement Attachment for IBM System z Replacement Capacity Offerings (in the US this is form number Z125-7880). Each enterprise is required to sign this contract one time within a given country before IBM will accept an order for its first-ever instance of the enablement feature code for Capacity for Planned Events (CPE) (feature #9912) or Capacity Back Up (CBU)(feature #9910). Further, this contract document is required if an enterprise's CBU contract expires and it wishes to renew CBU.

Capacity For Planned Events (CPE)

In addition to the base terms and the replacement capacity terms, CPE requires this contract, which contains terms specific to this offering: IBM Customer Agreement Attachment for IBM System z Capacity for Planned Events (in the US this is form number Z125-7882). Each enterprise is required to sign this contract one time within a given country before IBM will accept an order for its first-ever instance of the enablement feature code for CPE (feature #9912).

Optional High Voltage DC Power

- **Using high voltage DC power can save power by eliminating DC to AC and AC to DC conversion losses**
- **The bulk power in z196/z114 is modified to support HVDC so the only differences in shipped HW to implement the option are the DC line cords:**
  - This adds DC line cord feature codes
  - Expect that nominal DC supply voltages supported will include:
    - 570V (500 - 600V)
    - 380V (330- 420V)

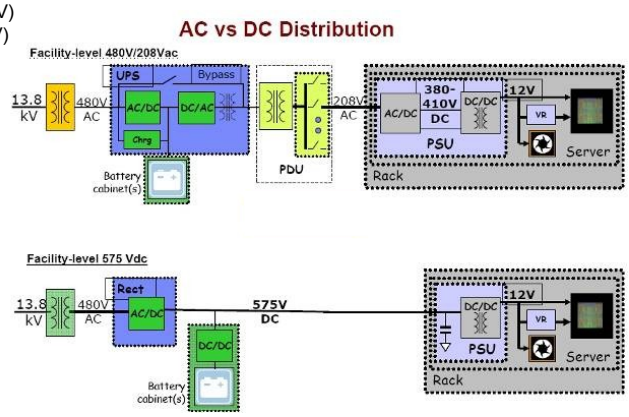

In data centers today, many businesses are paying increasing electric bills and are also running out of power. High-voltage DC power, an optional feature on z114, increases the voltage directly into the system. This can save System z users 1% - 3%, on average, on their power bills without having to go through a step-down.

High-voltage DC power is added to the universal power input that is available on the z114. The voltage level is in the range of 380V - 570V DC, and requires two or four 60 A line cords, depending on configuration. These line cords (#8965, #8963) are offered on new build as well as MES orders.

z/OS Support Summary

Release	z9 EC WdfM	z9 BC WdfM	z10 EC	z10 BC	z196	z114	End of Service	Coexists with z/OS
z/OS V1.8	X	X	X	X	X	X	9/09[1]	V1.10
z/OS V1.9	X	X	X	X	X	X	9/10[1]	V1.11
z/OS V1.10	X	X	X	X	X	X	9/11[2]	V1.12
z/OS V1.11	X	X	X	X	X	X	9/12*	V1.13*
z/OS V1.12	X	X	X	X	X	X	9/13*	V1.14*
z/OS V1.13	X	X	X	X	X	X	9/14*	V1.15*

Notes:
[1]The IBM Lifecycle Extension for z/OS provides the ability for customers to purchase extended defect support for that release of z/OS for up to 24 months after the z/OS release's end of service date.
[2] z/OS 1.10 EOS is September 29, 2011. Requires Lifecycle Extension after this date
* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

System z z/VM & z/VSE Support Summary

		z9 EC WdfM	z9 BC WdfM	z10 EC	z10 BC	z196	z114	Ship Date	End of Market	End of Service
z/VSE*	4.2	x	x	x	x	x	X	10/08	TBD	TBD
	4.3	x	x	x	x	x	X	4Q10	TBD	TBD
	5.1*	x	x	x	x	x	X	4Q11	TBD	TBD
z/VM	5.4	x	x	x	x	x	X	09/08	TBD	9/13*(1)
	6.1	No	No	x	x	x	x	10/09	TBD	4/13*

Note: z/VSE V4 is designed to exploit 64-bit real memory addressing, but does not support 64-bit virtual addressing. z/VSE V5.1 has been previewed.
Crypto Express3 requires z/VSE V4.2 or later.

1) End of Service date for z/VM 5.4 extended
* Planned. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice. Any reliance on these Statements of General Direction is at the relying party's sole risk and will not create liability or obligation for IBM.

Note: z/VM 6.1 Implements a new Architecture Level Set (ALS) available only on IBM System z10 servers or later

z114TLLB29

Self explanatory

System z Linux Support






	z9 EC WdfM	z9 BC WdfM	z10 EC	z10 BC	z196	z114	Availability Date
RHEL 5	x	x	x	x	x	x	03/2007
RHEL 6	x	x	x	x	x	x	11/2010
SLES 10	x	x	x	x	x	x	08/2006
SLES 11	x	x	x	x	x	x	03/2009

	End of Production Ph 1	**End of Production Ph 2**	**End of Production Ph 3**
RHEL 5 support*	4Q 2011	4Q 2012	03/312014
RHEL 6 support*	4Q 2014	4Q 2015	11/30/3017

	General support	**Extended support**	**Self support**
SLES 10 support*	07/31/2013	07/31/2016	07/31/2016
SLES 11 support*	03/31/2016	03/31/2019	03/31/2019

- **For latest information and details contact your Linux distributor**
- Recommendation: use RHEL 6 or SLES 11 for new projects
- For latest information about supported Linux distributions on System z refer to: http://www.ibm.com/systems/z/os/linux/resources/testedplatforms.html

* SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux
Support dates may be changed by Linux distributors

zEnterprise 114 in a nutshell:

- **Integration with the zEnteprise BladeExtension**
- **Two Models**
- **Increased capacity in a single footprint**
  - 12s0 technology
  - out-of-order instruction processing
  - higher clock frequency
  - larger cache
- **Robust Memory**
- **Upgrades**
  - Investment protection with upgrades from two previous families
    - z10 BC
    - z9 BC
  - Upgradeability to z196 (M15)
- **I/O Improvements**
  - new I/O features and New I/O Drawer

IBM zEnterprise 114 (z114) and z196 GA2 Key Dates for Hardware

- **Planned Availability Dates:**
  - **September 09, 2011 – Driver 93**
    - IBM zEnterprise 114 Models M05 and M10 new builds
    - Field installed features and conversions that are delivered solely through a modification to the machine's Licensed Internal Code (LIC):
      - LIC only upgrades for any PU type and memory
      - LIC for CBU and On/Off Capacity upgrade for installed z114s
      - z114 Capacity Setting downgrades
    - IBM System z9® Business Class (z9® BC) upgrades to z114
    - IBM System z10™ Business Class (z10 BC™) upgrades to z114
    - NEW z196s at GA2 level
      - MES for new function/features for installed z196s
    - TKE Support for LIC 7.1
  - **September 26, 2011**
    - Unified Resource manager functions:
      - Manage suite (#0019) enhancements
      - Automate/Advanced Management Firmware Suite (#0020) enhancements
      - Manage Firmware System x Blade (#0042)
      - Advanced Management Firmware System x Blade (#0046)
  - **October 21, 2011**
    - Add a zBX as an MES to the installed z114
  - **December 31, 2011**
    - z114 Model M05 to M10 upgrades
    - Feature adds (memory, I/O, RPQs and zBX features)
  - **Future***
    - Microsoft Windows (64-bit only) for select IBM System x Blades in a zBX

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

Thank You!

IBM

zEnterprise Backups

What is a zEnterprise Ensemble?

- A zEnterprise ensemble is a collection of 1 to 8 z196/z114 CPCs with/without zBX managed collectively by the Unified Resource Manager as a single logical virtualized system using the HMC
- A zEnterprise node is a z196/z114 CPC with 0 to 4 racks up to 2 BladeCenters per rack
- Blade based fit-for-purpose Solutions
- Integrated Advanced Virtualization Management
- Implements well-defined external interface to Data Center Service Management functions
- Virtual Resource Management and Automation

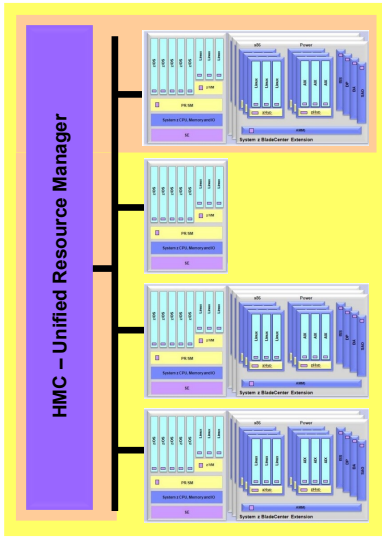

zEnterprise – What are the INMN, IEDN and Customer networks

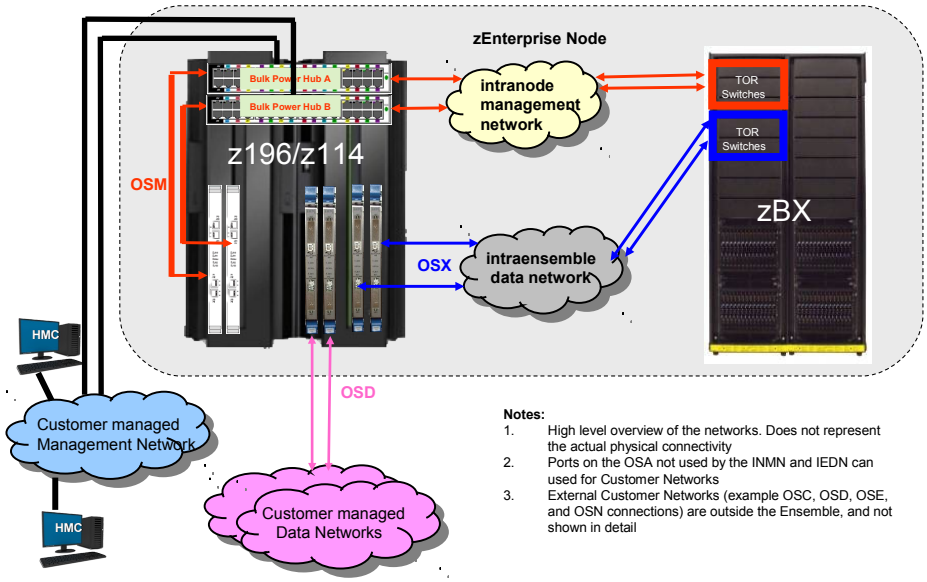

Notes:

1. High level overview of the networks. Does not represent the actual physical connectivity
2. Ports on the OSA not used by the INMN and IEDN can used for Customer Networks
3. External Customer Networks (example OSC, OSD, OSE, and OSN connections) are outside the Ensemble, and not shown in detail

z114TLLB36

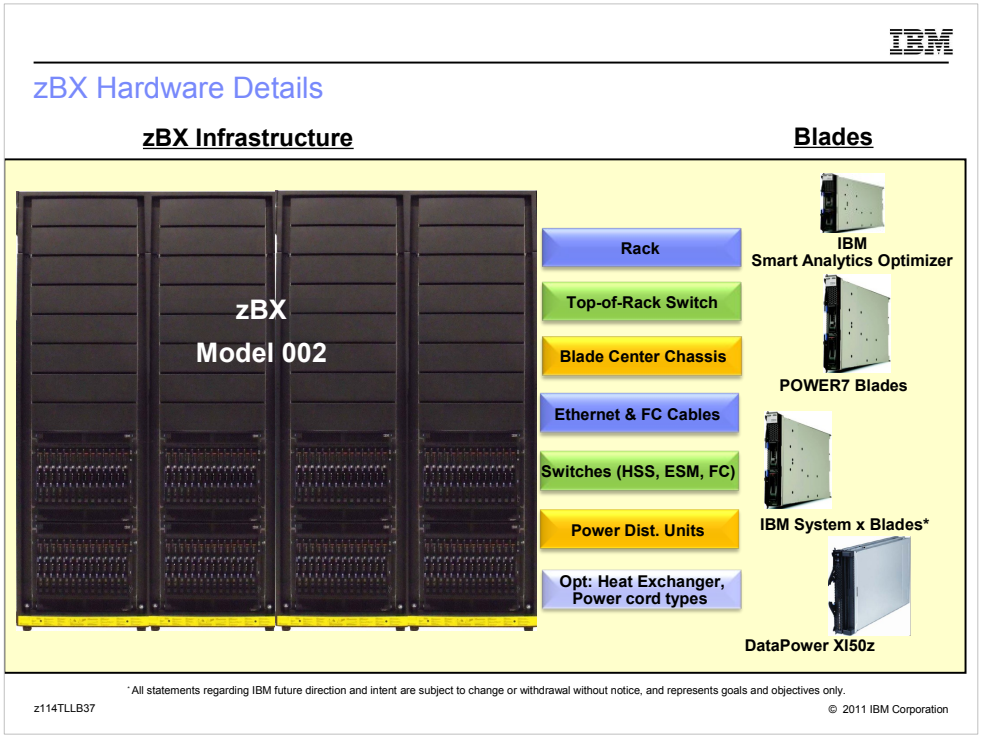

The zBX can support the IBM Smart Analytics Optimizer for DB2 for z/OS, V1.1 (5697-AQT), the IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z), and select POWER7 and IBM System x blades.

Smart Analytics Optimizer - a specialized accelerator for queries: The IBM Smart Analytics Optimizer is a high-performance, integrated hardware and software specialized accelerator that delivers dramatically faster analytic response times to select queries, complementing traditional query processing. It offers business value at lower cost through the simplicity and rapid deployment characteristics of an appliance, while retaining the benefits of having the data managed and secured by DB2 for z/OS. This workload-optimized, appliance-like add-on is installed in the zBX, connects to DB2 transparently to the user and applications, and requires no changes to an existing application to take immediate advantage of the analyses.

The IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z) is a multifunctional appliance that can help provide multiple levels of XML optimization, streamline and secure valuable service-oriented architecture (SOA) applications, and provide drop-in integration for heterogeneous environments by enabling core enterprise service bus (ESB) functionality, including routing, bridging, transformation, and event handling. It can help to simplify, govern, and enhance the network security for XML and web services.

POWER7 blades (supporting AIX) and/or System x blades (supporting Linux) installed in the zBX can enable application integration with System z transaction processing, messaging, and data serving capabilities. The blades are managed as part of a single logical virtualized environment by the IBM zEnterprise Unified Resource Manager.

The IBM zEnterprise Unified Resource Manager manages System z ensembles, collections of one or more zEnterprise System nodes in which each node is comprised of a zEnterprise (z114 or z196) and its optionally attached IBM zEnterprise BladeCenter Extension (zBX). An ensemble can consist of a single zEnterprise with no zBX attached, or up to eight CPCs where at least one of the CPCs has a zBX attached. The resources of a zEnterprise System ensemble are managed and virtualized as a single pool of resources, integrating system and workload management across the multisystem, multitier, multiarchitecture environment.

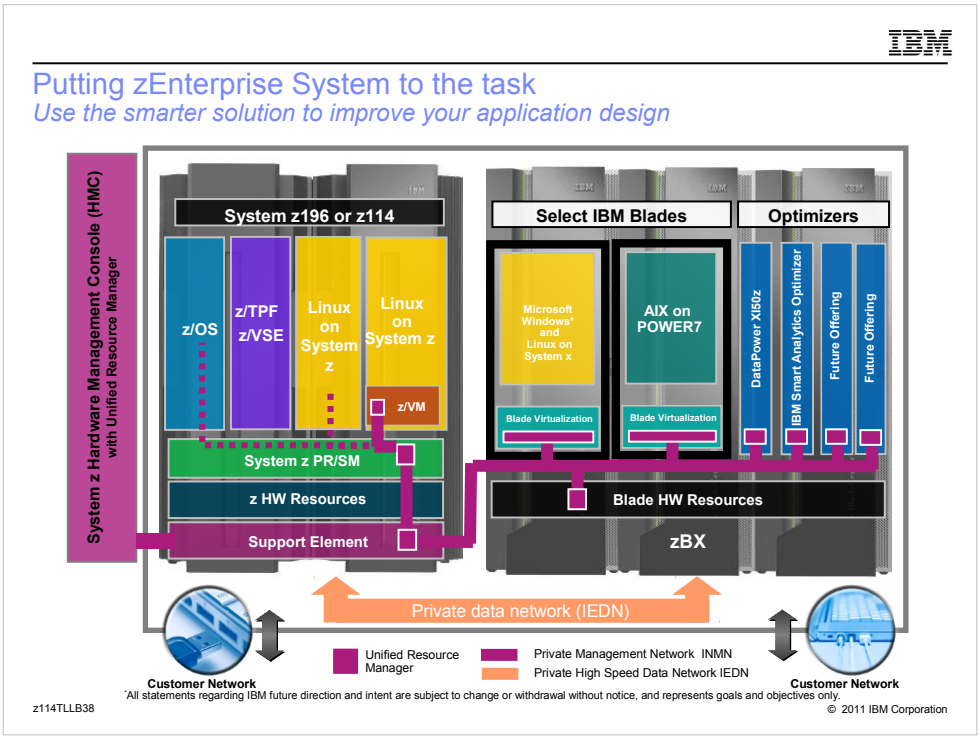

Putting zEnterprise System to the task – Use the smarter solution to improve your application design

Note to a presenter.. this chart can wrap up what you've told the customer about the new solution in 1 picture. The Purple boxes and lines represent the Unified Resource Manager, it's agents and the private support network (intra-node network management). The dotted line shows that z/OS and Linux on System z will offer monitor capabilities but not turn over management (provide agents) to Unified Resource Manager. The orange arrow at the bottom is the private data network (called Intra-ensemble data network – IEDN). You can point out the zEnterprise xxx, the zBX, the blades and optimizers.

Note to the presenter:

Depending on your comfort level, and time, you can talk about:

•How the hypervisors are treated as system z firmware

•That the general purpose x and p blades are virtualized unconditionally

•The SE has been extended to orchestrate hardware operational control for the blades as well as for System z (config, problem management, call home, etc)

•The blades are optional, and you get the Unified Resource Manager capabilities with the system zEnterprise xxx (don't need a zBX to get value)

•The role of the HMC is being extended, and is evolving to be a full function platform management appliance

•The HMC represents the "management access point" for all Unified Resource Manager management functions