

Introducing the IBM zEnterprise 114

Gregory Hutchison
IBM

August 09, 2011
Session Number 9796

Technical Review - Agenda

- zEnterprise 114 Overview
 - z BladeCenter Extension (zBX)
 - Functions
 - Performance
 - Upgrades
 - Memory
- I/O, Security, Miscellaneous
 - I/O Drawers
 - I/O Features
 - Discontinued I/O Features
 - Cryptography
 - Server Time Protocol
 - Installation Options
 - Capacity on Demand Enhancements
 - Operating Systems
 - Hardware Management Console



Covered in Session 9797
zEnterprise 196 I/O Infrastructure Update
4:30PM
Room - Southern Hemisphere 1/2

z114 Business Value

- <http://www.youtube.com/watch?v=SXWonQEvI1Y>
- Other
 - http://www.youtube.com/results?search_query=zenterprise+114&aq=f

Introducing the zEnterprise

Bringing hybrid computing to a broader set of businesses



IBM zEnterprise 114 (z114)

The next generation midrange mainframe delivering extensive growth options, flexibility, efficiency and improved price performance.

zEnterprise Unified Resource Manager

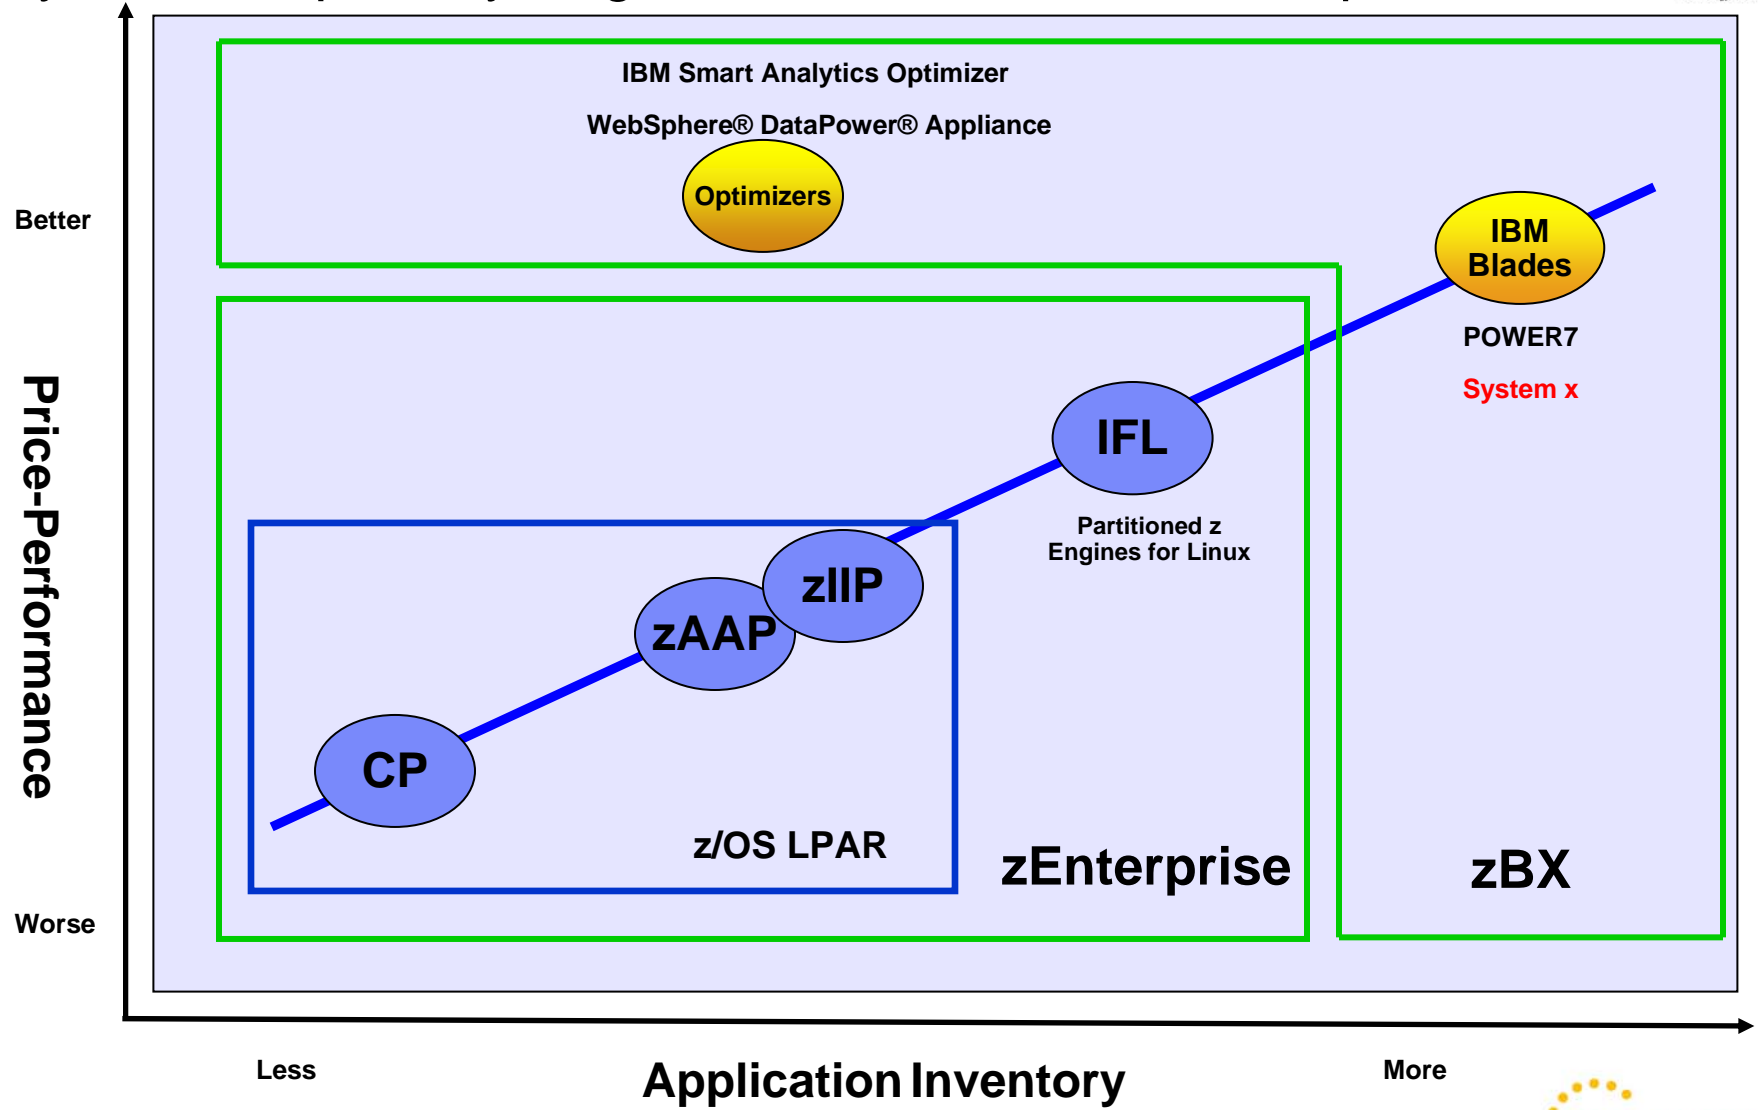
Centralized management of heterogeneous resources for simplification and resiliency

zEnterprise BladeCenter Extension (zBX)

Integrated IBM POWER7® blades, IBM System x blades, and High-performance optimizers and appliances*

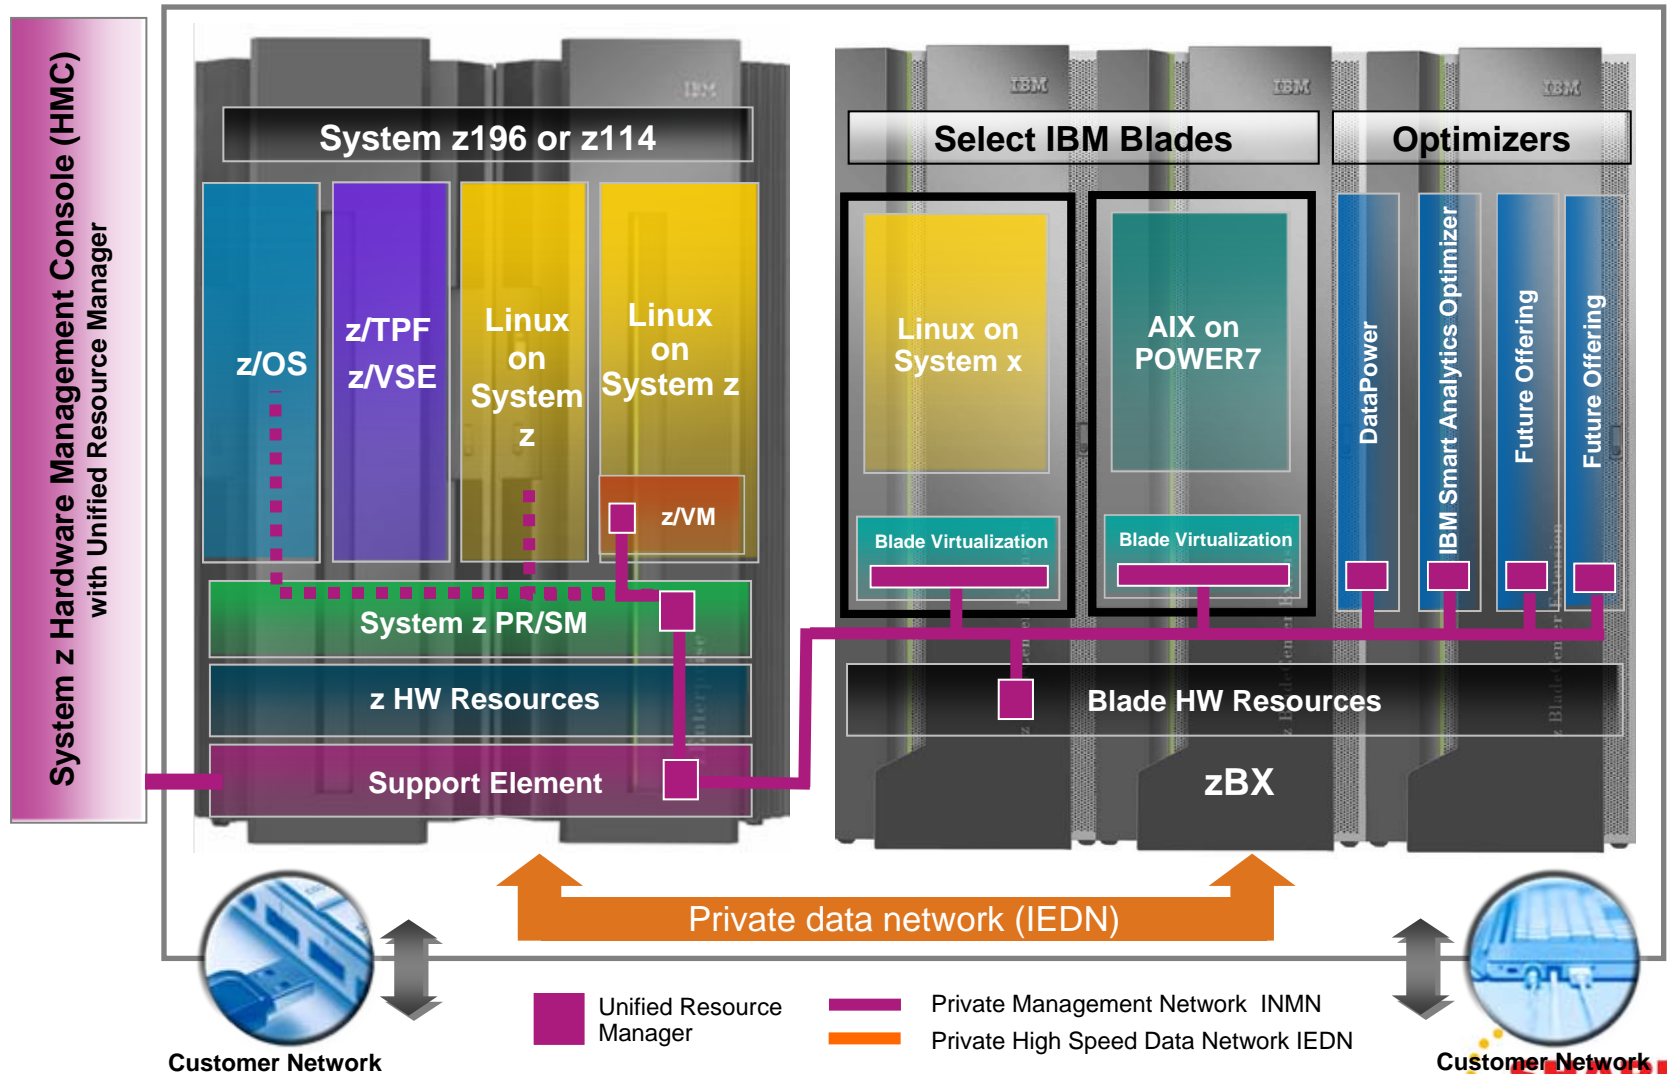
* Statement of Direction

System z “Specialty Engine” Evolution to the zEnterprise Ensemble



Putting zEnterprise System to the task

Use the smarter solution to improve your application design



2458-002 - IBM Smart Analytics Optimizer



- Pre-packaged and pre-tested
- zBX components are a logical extension to System z as a new System z Machine Type/Model.
 - Machine Type 2458
 - Model 002
- Used for specialized workload processing which can be handled more economically than if those workloads were processed directly in the System z server
- zBX processing components are provided using standard BladeCenter[®] components.
- Specific disk requirement, DS5020s, use for backing up data on the blades.
- Impressive Performance
 - Compressed DB2 data
 - Parallel file system
 - In memory execution

IBM Smart Analytics Optimizer - Sizing

- **How do I size the right machine?**
 - Watch this space, things may change
 - Initially, go here
- **For requests outside of North America**
 - dwhz@de.ibm.com
- **For requests in North America**
 - Forward the sizing request to the BI Swat team under Beth Hamel
 - [DW on System z/Silicon Valley/Contr/IBM](#)
- <https://w3.tap.ibm.com/w3ki08/display/isao/Home>
<https://w3.tap.ibm.com/w3ki08/display/isao/Process>
 - Download an off-line version of the questionnaire ([ISAO_Assessment_Questionnaire.doc](#)) from <https://w3.tap.ibm.com/w3ki08/display/isao/Process>
 - Complete Questionnaire
 - System Environment and Data Warehouse workload (to make sure that the customer meets the requirements)
- Send the completed Questionnaire to the User ID dwhz@de.ibm.com or to BI Swat team under Beth Hamel in North America [DW on System z/Silicon Valley/Contr/IBM](#) or use dwonz@us.ibm.com.
 - It is not recommended that you approach the customer until you have had feed back on the ISAO Assessment
 - a quick analysis of real workload should be performed (Quick Workload test)
- Down load the [ISAO Assessment Description.zip](#) from the <https://w3.tap.ibm.com/w3ki08/display/isao/Process>


IBM POWER7

General purpose processors under one management umbrella



What is it?

The zBX infrastructure can host select IBM POWER7 blades. Each blade comes with an installed hypervisor that offers the possibility of running an application that spans z/OS, Linux on System z, AIX on POWER, but have it under a single management umbrella.



IBM System X blade - Not offered with this Early Support Program

How is it different?

- **Complete management:** Advanced management brings operational control and cost benefits, improved security, workload management based on goals and policies.
- **Virtualized and Optimized:** Virtualization means fewer resources are required to meet peak demands with optimized interconnection.
- **Integrated:** Integration with System z brings heterogeneous resources together that can be managed as one.
- **Transparency:** Applications certified to run on AIX 5.3 or 6.1 will also be certified and run on the POWER7 blade. No changes to deployed guest images.
- **More applications:** Brings larger application portfolio to System z.

IBM Blade based on Power7



- MT 8406-71Y (PS701)
 - Power7 8 Core Processor
 - 8 Processor Cores activated
 - 1 Processor socket
 - Single wide Blade only
 - 3.0GHz
 - 16 dimm slots (4, 8, & 16 GB/core)
 - 300GB HDD Internal Disk
- 3 Configurations are supported.
- IBM POWER7 supports the 10GbE IEDN.
- IBM Blade Chassis attach to the INMN TOR via 1 GbE.

- Blades acquired by the customer through existing channels or through IBM (not from System z).
- A PowerVM Enterprise Edition licence and Software Maintenance Agreement is required for all 8 Cores, and must be maintained for the duration of use.
- AIX 5.3+, 6.1+

Customer procured
With AIX and PowerVM EE Licenses!

Hardware Warranty and Maintenance

24x7 on-site support for parts and service during the 1 year System z warranty and subsequent post warranty maintenance terms. Do not purchase a separate blade warranty. Provided as part of the zBX warranty and terms.

Power ASB	Feature Code	Config 1	Config 2	Config 3
Processor 3.0GHz@150W		1	1	1
Processor Activations (8)	8411 8412	4 4	4 4	4 4
Memory kits 8 GB (2 x 4 GB) 16 GB (2 x 8 GB)	8208 8209	32GB 4 0	64GB 8 0	128GB 0 8
HDD 300GB	8274	1	1	1
CFFh 10GbE	8275	1	1	1
CIOv 8Gb FC	8242	1	1	1
PowerVM EE	5228	8	8	8
Required SW	PID			
SW License PID 5765-PVE	0001	8	8	8
1 YR SWMA PID 5771-PVE	1191	Choose Qty 8 of 1 YR or 3 YR		
3 YR SWMA PID 5773-PVE	0999			




IBM System x blade

General purpose processors under one management umbrella



What is it?

The zBX infrastructure can host select IBM System x blades. Each blade comes with an installed hypervisor that offers the possibility of running an application that spans z/OS, Linux on System z, Linux on System x but have it under a single management umbrella.



How is it different?

- **Complete management:** Advanced management brings operational control and cost benefits, improved security, workload management based on goals and policies.
- **Virtualized and Optimized:** Virtualization means fewer resources are required to meet peak demands with optimized interconnection.
- **Integrated:** Integration with System z brings heterogeneous resources together that can be managed as one.
- **Transparency:** Applications certified to run on RHEL 5.5 or SLES 11 SP1 will also be certified and run on the HX5 7873 blade. No changes to deployed guest images.
- **More applications:** Brings larger application portfolio to System z.

New Blades Provide Added Flexibility for Workload Deployment and Integration



Introducing System x Blades in the zBX

- IBM BladeCenter HX5 7873 dual-socket 16-core blades
- Complements existing portfolio of POWER7, DataPower XI50z and IBM Smart Analytic Optimizer blades.
- Ordered and fulfilled through System x providers
- Blades assume System x 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

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

Extending support to New Operating System Environments

NEW

- **Support for Linux and in the future Windows¹ environments on select System x blades**

- 64-bit version support only
- Linux: RHEL 5.5, SLES 11 SP1
- Additional versions to follow¹
- The zBX web page hosts the most current blade ordering information:
http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfid=ZSL03128USEN&attachment=ZSL03128USEN.PDF
- In the future we are planning to support Microsoft[®] Windows[®] Server 2008 - Datacenter Edition¹

- **Operating Systems are customer acquired and installed**

Manage your mainframe and distributed environment with the same tools, same techniques, same practices

¹ All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

IBM zEnterprise™ BladeCenter® Extension (zBX) IBM System x® Blades



After August 30th, new models will be preconfigured for you in SSCT.
This table is useful for pricing today.

MT 7873 (HX5)
July 12th Announce
GA September 26th

Customer Configuration

- Intel 8 core Processor
- 2 Processor sockets
- 2.13 GHz 105W
- Max 14 A16M's per BC-H
- Memory 1066 Mhz with 6.4 GTs
- 16 DIMM slots
- 100GB SSD Internal Disk

- **Blades acquired by the customer through existing channels or through IBM.**

- **Virtualization: Integrated Hypervisor supplied by Unified Resource Manager**

Description	Part Number	Option Part Number	Feature Code	Config 0	Config 1
Blade Base	69Y3056	69Y3056	A16M	1	1
Initial Processor 2.13 GHz 105W (E7-2830 8C)	69Y3071	69Y3071	A16S	1	1
Additional Processor 2.13 GHz 105W (E7-2830 8C)	69Y3072	69Y3074	A179	1	1
# Intel Processors (Sockets)				2	2
Blade Width				Single	Single
Total Cores				16	16
Memory 8GB 1333 MHz	46C0558	46C0570	A17Q	64GB 8	128GB 16
GB/core				4	8
Speed Burst	46M6843	59Y5889	1741	1	1
SSD Expansion Card	46M6906	46M6908	5765	1	1
50 GB SSD MLC	46W7727	43W7726	5428	2	2
No Internal RAID			9012	1	1
CFFh 10 GbE	46M6170	46M6168	0099	1	1
CIOv 8Gb FC	44X1946	44X1945	1462	1	1

<http://www.ibm.com/systems/z/hardware/zenterprise/zbx.html>



System x Blade Orderings

- Use The IBM Standalone Solutions Configuration Tool (SSCT)
 - <https://www-947.ibm.com/support/entry/myportal/docdisplay?brand=5000008&Indocid=MIGR-62168>
- Will release four hardware configurations with Operating System choices.
 - Only two configurations are available on the zBX initially.
- The supported System x blades (new model numbers) will be available using the IBM SSCT configuration tool on August 30, 2011.


IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise helps extend the value of zEnterprise



Purpose-built hardware for simplified deployment and hardened security helps businesses quickly react to change and reduce time to market

What is it?

The IBM WebSphere DataPower Integration Appliance XI50 for zEnterprise can help simplify, govern, secure and integrate XML and IT services by providing connectivity, gateway functions, data transformation, protocol bridging, and intelligent load distribution.



How is it different?

- **Security:** VLAN support provides enforced isolation of network traffic with secure private networks.
- **Improved support:** Monitoring of hardware with “call home” for current/expected problems and support by System z Service Support Representative.
- **System z packaging:** Increased quality with pre-testing of blade and zBX. Upgrade history available to ease growth.
- **Operational controls:** Monitoring rolled into System z environment from single console. Consistent change management with Unified Resource Manager.

zEnterprise provides the foundation for the “smart” infrastructure on which we can build the workloads of today and tomorrow

They are workloads that ...

- Rely on data serving and application components on IBM System z®
- Solutions that need to leverage strengths of System z... Security, Reliability, Availability
- Have application components on UNIX (HP, Sun, Power) or Linux (x86, System z) but require a higher level of integration capabilities and efficiency



... and / or ...

- Reside in low utilization / development environments
- Can be made more efficient through consolidation
- Can be optimized by using the newest virtualization technology


... but also may ...

- Reside in complex multi-platform IT environments
- Require flexible development and test infrastructure
- Require simplified, integrated policy and management

Hardware Withdrawal: IBM System z10 EC and IBM System z10 BC

- Effective June 30, 2012, IBM is withdrawing the following selected products from marketing worldwide
 - All models of the IBM System z10 Enterprise Class (z10 EC) and all upgrades to the z10 EC from the IBM eServer zSeries 990 (z990), IBM System z9 EC (z9 EC), or IBM System z10 BC (z10 BC).
 - All models of the IBM System z10 Business Class (z10 BC) and all upgrades to the z10 BC from the IBM eServer zSeries 890 (z890) or IBM System z9 BC (z9 BC).
 - Model conversions and hardware MES features applied to an existing z10 EC or z10 BC server.
- Field installed features and conversions that are delivered solely through a modification to the machine's Licensed Internal Code (LIC) will continue to be available until June 30, 2013. After June 30, 2013, features and conversions that are delivered solely through a modification to the LIC will be withdrawn.
- The Capacity on Demand offerings that are configured prior to withdrawal are usable until the offering expiration date or termination date, as applicable.

zEnterprise 114 Product Positioning



- zEnterprise provides increased capacity in a single footprint
 - Designed for up to a 18% performance improvement per core and up to 12% improvement in total system capacity for z/OS, z/VM, and Linux workloads on System z compared to the z10 BC.
 - 12s0 technology
 - higher clock frequency 3.8 Ghz
 - out-of-order instruction processing
 - larger caches
 - compiler enhancements
- Connectivity improvements include bandwidth and throughput

System zEnterprise 114 Functions and Features



- Two hardware models
- Up to 10 processors configurable as CPs, zAAPs, zIIPs, IFLs, ICFs, or optional SAPs
- Up to 26 subcapacity settings across a maximum of 5 CPs
- Increased capacity processors
- Out of order instruction execution
- Improved processor cache design
- New and additional instructions
- Dedicated Spares on the Model M10
- Up to 248 GB of Redundant Array of Independent Memory (RAIM)
- Memory power save
- Cryptographic enhancements
- On Demand enhancements
- 6.0 GB/sec InfiniBand I/O interconnect



- 2 New OSA CHPIDs – OSX and OSM
- New 32 slot PCIe Based I/O Drawer
- Concurrent I/O drawer add, remove, replace
- Doubled HiperSockets to 32
- Physical Coupling Links increased to 72
- Doubled Coupling CHPIDs to 128
- CFCC Level 17 enhancements
- Optional High Voltage DC power
- Optional overhead I/O cable exit
- NRF Support with either top exit or bottom exit I/O and power.
- STP enhancements
- zBX Model 002 with ISAOPT, POWER7, DataPower and IBM System x Blades
- Platform Management from HMC

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 CEC Drawer design (1 processor 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 CEC Drawer design (2 processor 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

Model Structure and Upgrades

	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 processing drawers
- The number of processing drawers based upon the number of CPs and specialty engines.

Capacity Matrix – 130 Capacity Settings



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

zEnterprise 114

- 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

Processor / Memory Subsystem Drawers

(Model M05 and M10)

One z10 BC Drawer

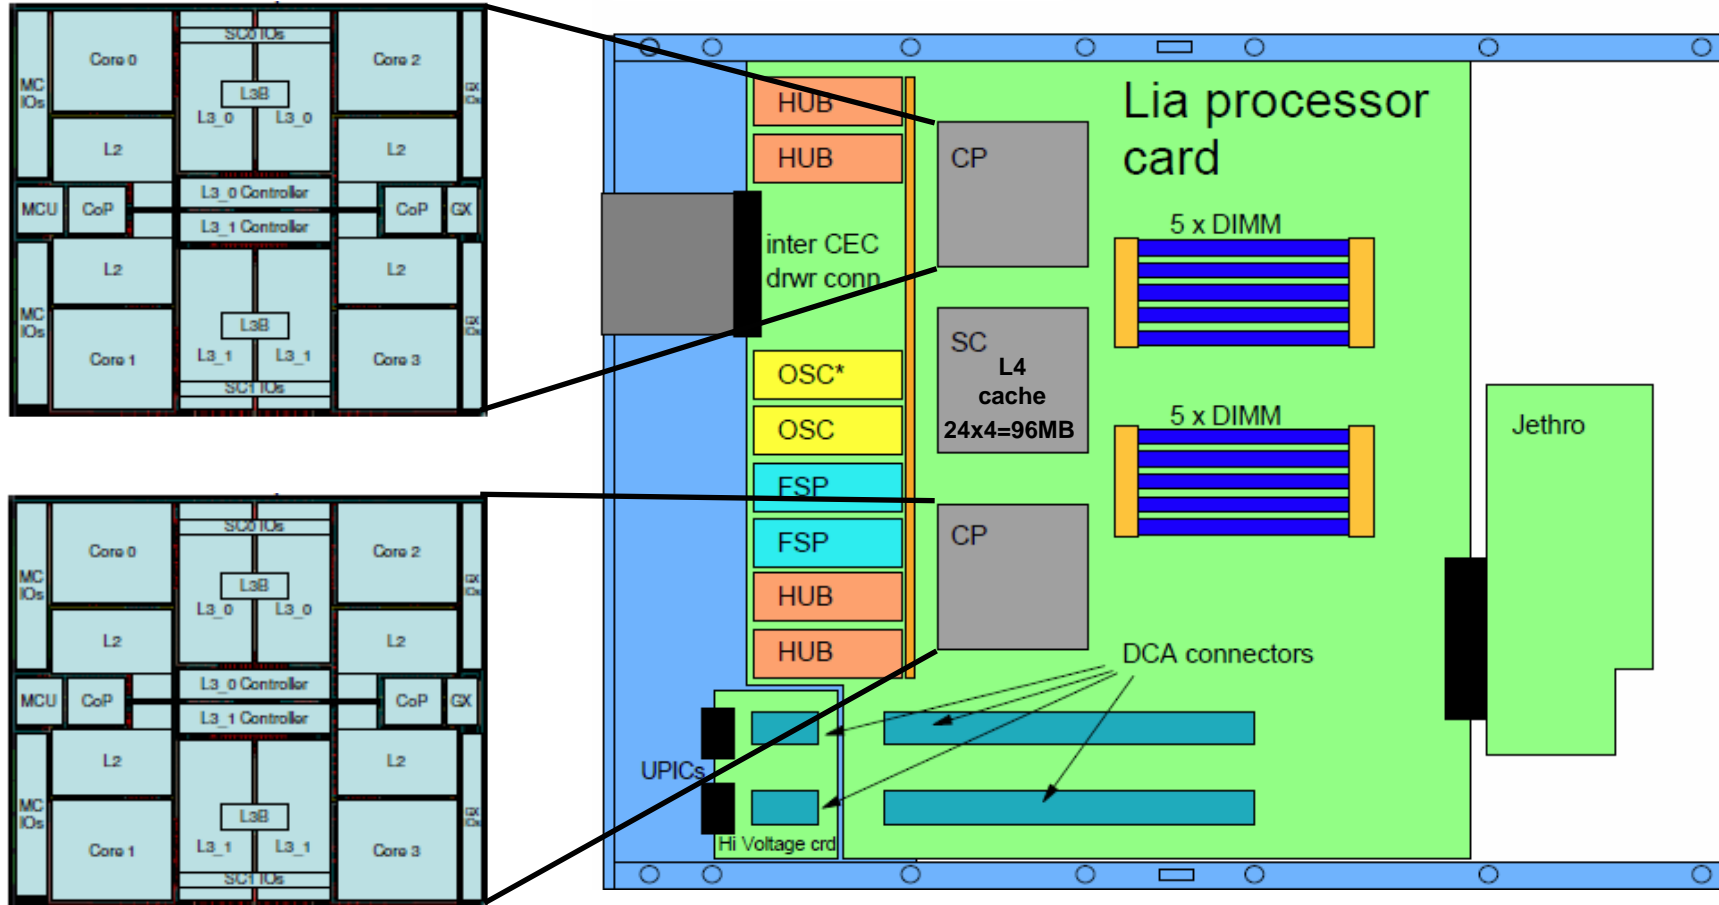


Two z114 Drawers (Model M10)




- **System resources split between 2 drawers (Model M10)**
- **Second CEC drawer (Model 10) for:**
 - Increased specialty engine capability
 - Increased memory capability
 - Increased I/O capability
 - More coupling links than z10 BC
 - More I/O features than z10 BC
- **Planning Note: Unlike the z196 Books, add/remove/repair of the CEC drawer is disruptive**

Single-Chip Module (SCM) in processing drawer(s)



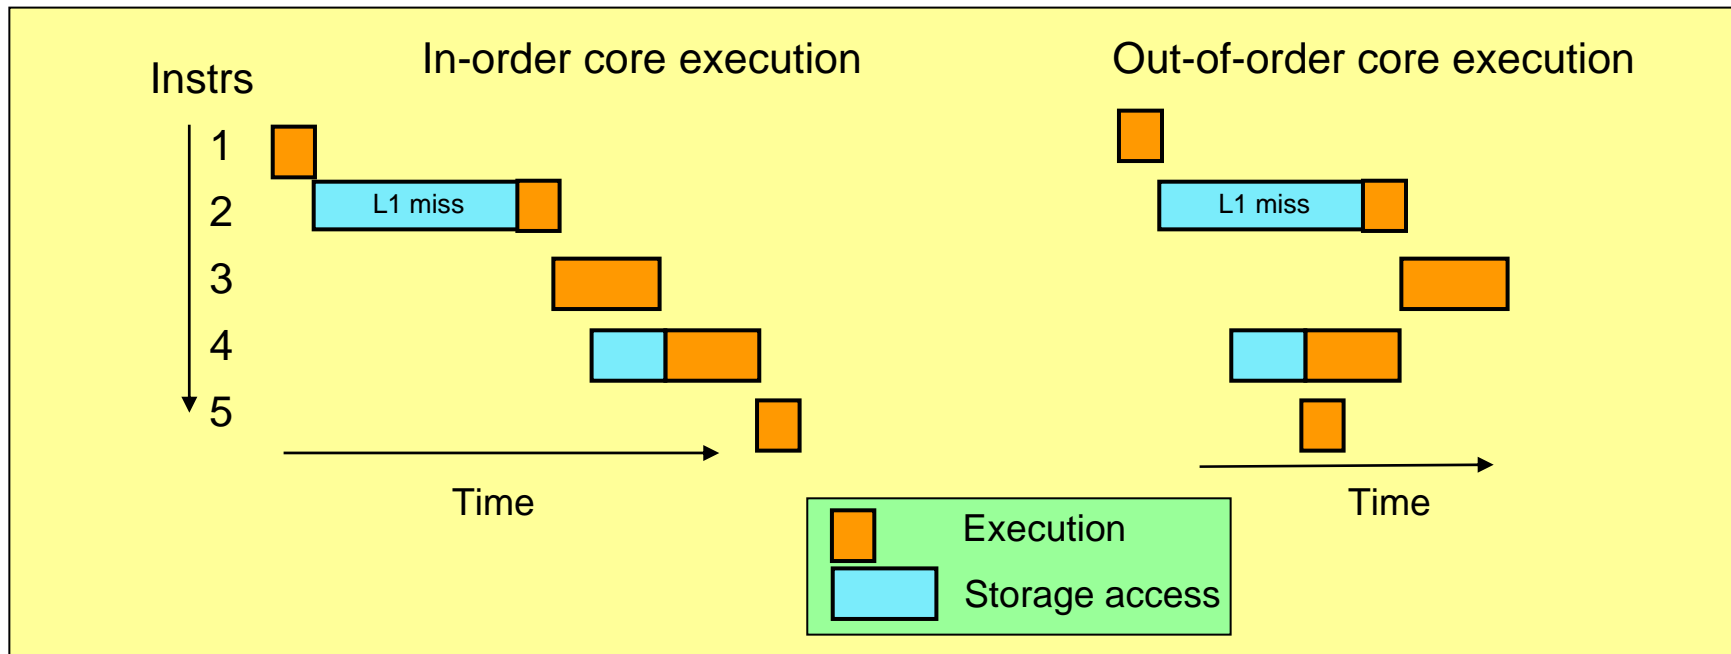
Single-Chip Module (SCM) in processing drawer(s)

- Quad core chips with 3 or 4 active cores
 - Same as the zEnterprise 196
- 3.8 GHz
- **L1: 64K I / 128K D private/core**
- **L2: 1.5M I+D private/core**
- **L3: 12MB**
 - Same chip as z196, but enabled half of the available 24MB
- L4: 96MB per processing drawer
 - On the SC Chip
 - 24MB assigned to each core
 - $24 \times 4 = 96$
 - Half of that on the z196



zEnterprise Out-of-Order (OOO) Value

- **OOO yields significant performance benefit for applications through**
 - Re-ordering instruction execution
 - Later (younger) instructions can execute ahead of an older stalled instruction



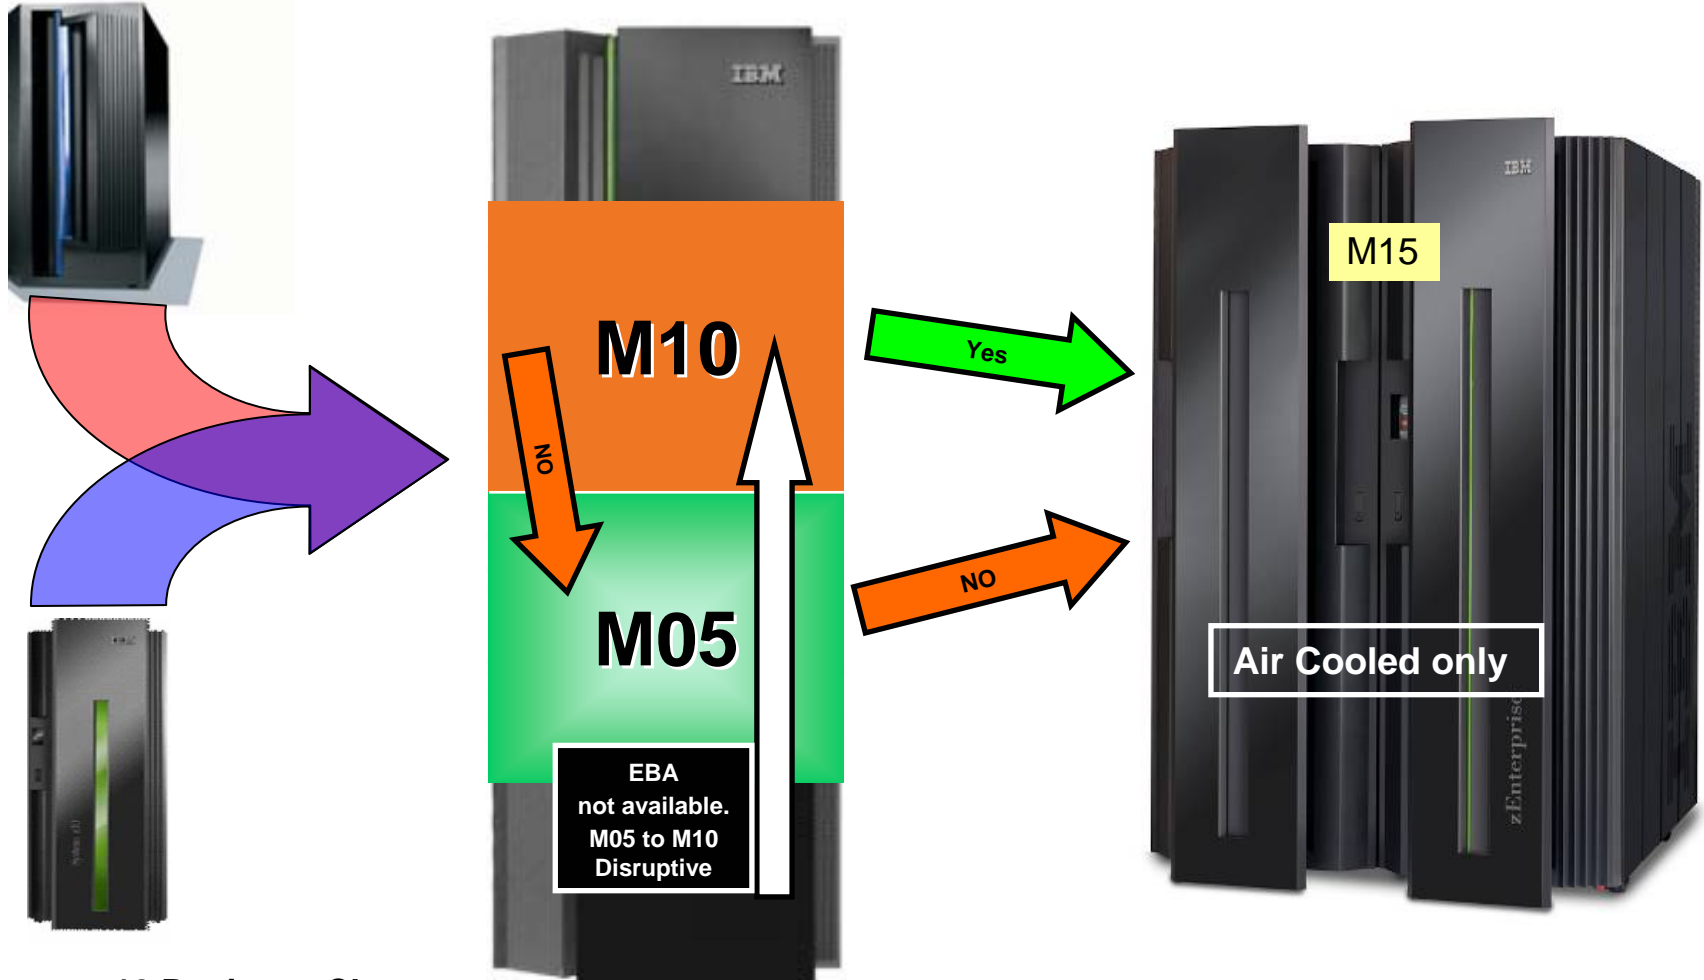
Family Upgrades

Cards being moved from the z9 or z10 will not maintain the existing PCHIDs on the z114. Upgrades will be configured as New Builds. See the "Moved Report" in eConfig.

IBM System z9 Business Class

IBM z114

IBM zEnterprise 196



IBM System z10 Business Class

EBA=Enhanced Book Availability

zEnterprise 114 Concurrent Conversions



- Must order (characterize one PU as) a CP, an ICF or an IFL
- Conversions within zEnterprise 114 family
- zEnterprise 114 to z196 Model M15 (401-715)
- Concurrent processor upgrade is supported if PUs are available
 - **Add CP, IFL, unassigned IFL, ICF, zAAP, zIIP or optional SAP**

From/To->	CP	IFL	Unassigned IFL	ICF	zAAP	zIIP	Additional SAP
CP	x	Yes	Yes	Yes	Yes	Yes	Yes
IFL	Yes	x	Yes	Yes	Yes	Yes	Yes
Unassigned IFL	Yes	Yes	x	Yes	Yes	Yes	Yes
ICF	Yes	Yes	Yes	x	Yes	Yes	Yes
zAAP	Yes	Yes	Yes	Yes	x	Yes	Yes
zIIP	Yes	Yes	Yes	Yes	Yes	x	Yes
Additional SAP	Yes	Yes	Yes	Yes	Yes	Yes	x

Exceptions: Disruptive if ALL current PUs are converted to different types may require individual LPAR disruption if dedicated PUs are converted. Note: Conversion of a CP to a Specialty Engine is a delete/add.

zEnterprise 114 Conversions (continued)

- Any to Any is allowed (upgrades or downgrades)
 - CP Capacities are handled separately from specialty engines.
 - When CP capacity on the target machine is ordered using fewer CPs.....
 - *the remaining CPs are not available to be converted to specialty engines without a fee.*

Examples:

2-way to 1-way capacity downgrade

The PU in the green column must maintain its high water mark and can not be purchased as a specialty engine.

2-way to 1-way capacity upgrade

The high water mark is now in the red column and the PU in the green column can be used during the purchase of a specialty engine.

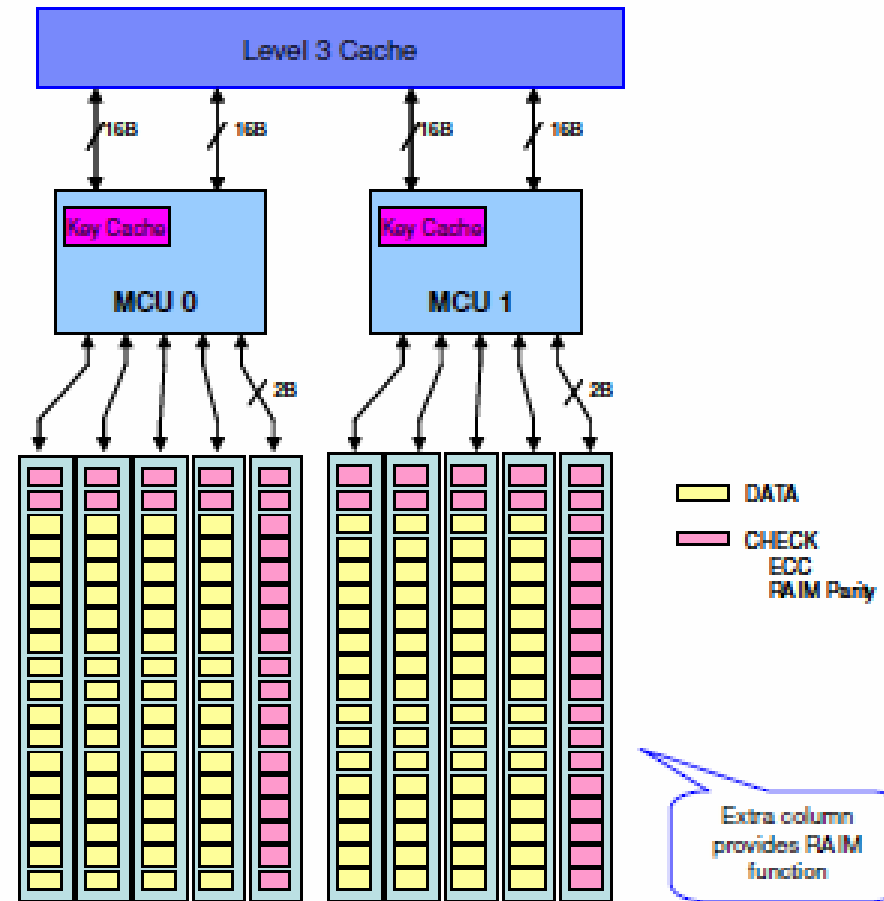
F01	F02	F03	F04	F05
E01	E02	E03	E04	E05
D01	D02	D03	D04	D05
C01	C02	C03	C04	C05
B01	B02	B03	B04	B05
A01	A02	A03	A04	A05
1-way	2-way	3-way	4-way	5-way



Memory

8 GB to 120 GB (M05)
16 GB to 248 GB (M10)

- Memory technology introduced on the z196 is used on the zEnterprise 114.
 - Redundant Array of Memory (RAIM) which in the Disk industry is known as RAID.
 - Protection from UIRAs (outages) caused by a DIMM failure.
 - DIMM failures include all components on the DIMM including
 - *Supernova*
 - *Drams*
 - *Connectors*
 - Portions of the memory controller or card failure isolated to 1 memory channel.



HSA history

- HSA significantly larger than pre-z10 Servers
- Fixed 8 GB HSA and 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 EC 16 GB – Fixed
 - z10 BC 8 GB – Fixed
 - z196 16 GB – Fixed
 - **z114 8 GB - Fixed**
- HSA Estimator on Resource Link not relevant

Memory Offerings

Memory upgrades within the same color (except white) are concurrent without the need for Memory Plan Ahead.




FC	GB	Increment	M05			M10 (2 processing drawers)		
			Dial Max	Dimm (GB)	# plugged	Dial Max	Dimm (GB)	# plugged
3609	8	8	24	4	10	N/A	N/A	N/A
3610	16	8		4	10	56	4/4	10/10
3611	24	8		4	10		4/4	10/10
3612	32	8	56	8	10		4/4	10/10
3613	40	8		8	10		4/4	10/10
3614	48	8		8	10	4/4	10/10	
3615	56	8		8	10	4/4	10/10	
3616	64	8	120	16	10	88	4/8	10/10
3617	72	8		16	10		4/8	10/10
3618	80	8		16	10		4/8	10/10
3619	88	8		16	10		4/8	10/10
3620	96	8		16	10	120	8/8	10/10
3621	104	8		16	10		8/8	10/10
3622	112	8		16	10		8/8	10/10
3623	120	8	16	10	8/8		10/10	
3624	152	32	N/A	N/A	N/A	152	4/16	10/10
3625	184	32	N/A	N/A	N/A	184	8/16	10/10
3626	216	32	N/A	N/A	N/A	248	16/16	10/10
3627	248	32	N/A	N/A	N/A		16/16	10/10

M05 Memory Features

- Plan Ahead Memory
 - Pre-plugged memory based on target capacity specified by the customer.
 - Enabled by LICCC, concurrently.
 - **FC1993** tracks the quantity of 8GB physical increments.
 - Charged (half price) when physical memory is installed
 - **FC1903** generally indicates 8GB (or 32 GB in larger configurations) LICC'd increments of Memory Capacity.
 - Charged by increments when Plan Ahead memory is enabled
 - Subsequent memory upgrade orders will use up the Plan Ahead memory first.

Plan Ahead Plan Ahead



10 x 4 GB DIMMs	10 x 8 GB DIMMs	10 x 16 GB DIMMs
Feature Size	Feature Size	Feature Size
8	32	64
16	40	72
24	48	80
	56	88
		96
		104
		112
		120

Physical memory upgrades are *DISRUPTIVE*

M10 Memory Features

2 Drawers

32 GB

Plan Ahead

Plan Ahead

Plan Ahead

Plan Ahead


Plan Ahead

4 GB/4GB	4GB/8GB 8GB/4GB	8GB/8GB	4GB/16GB 16GB/4GB	8GB/16GB 16GB/8GB	16GB/16GB
Feature Size	Feature Size	Feature Size	Feature Size	Feature Size	Feature Size
16	64	96	152	184	216
24	72	104			248
32	80	112			
40	88	120			
48					
56					

Physical memory upgrades are *DISRUPTIVE*

Agenda

- zEnterprise 114 Overview
 - z BladeCenter Extension (zBX)
 - Functions
 - Performance
 - Upgrades
 - Memory
- I/O, Security, Miscellaneous
 - I/O Drawers
 - I/O Features
 - Discontinued I/O Features
 - Cryptography
 - Server Time Protocol
 - Installation Options
- Capacity on Demand Enhancements
- Operating Systems
- Hardware Management Console



Covered in Session 9797
zEnterprise 196 I/O Infrastructure Update
4:30PM
Room - Southern Hemisphere 1/2

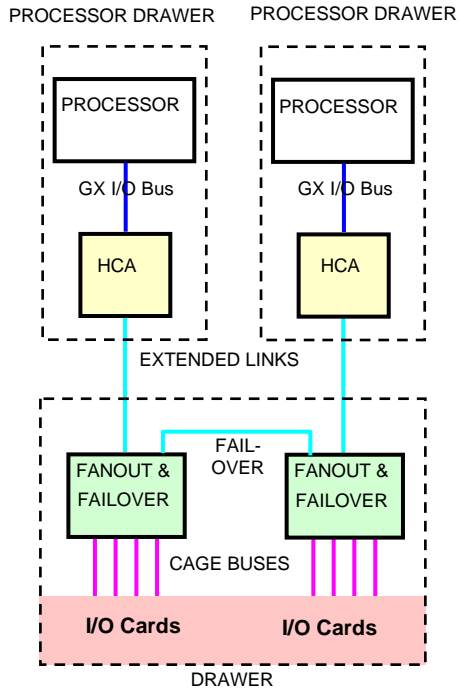
PCIe I/O Drawer (FC4003)



- Designed to
 - Support concurrent add and repair
 - Support for an industry standard
 - Potential attachment of select industry standard PCIe cards
 - Physically carry forward in an MES upgrade to z196
 - Provide improved performance for new and traditional workloads
 - Provide lower power requirements while increasing connectivity for all workloads
- Higher bandwidth
 - I/O bus infrastructure data rate up to 8GB/s
- Power Save Mode
 - Power can be reduced on unused ports
 - Clock gating of unused functions on the ASIC
 - Core voltage operation reduction
- Fewer ports per I/O card, but support for four times as many slots.
 - I/O card density – 14% more capacity (more slots, fewer ports per feature)
 - 32 I/O card slots
 - Maximum of two new I/O drawers per z114
 - Intended for new I/O features (FICON Express8S, OSA Express4S)
 - Legacy I/O technology (ESCON, ISC-3, PSC, Crypto Express3 and carry forward) still supported on I/O Drawer (FC4000)

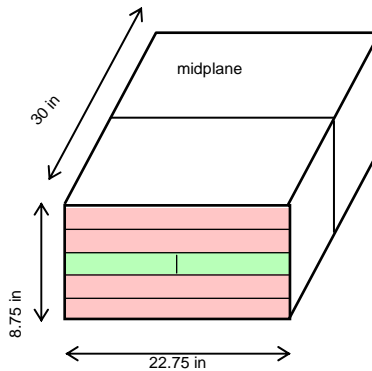
zEnterprise 114 I/O Drawers

I/O Infrastructure




Redundant I/O Reconnect

I/O Drawer FC4000 (same as z10 BC)




- 8 horizontal slots
- 4 Front/4 Back
- 32 FICON Ports
- 2 DOMAINS
- 2 & 4 port cards

PCIe I/O Drawer FC4003






- 32 vertical slots
- 16 Front/16 Back
- 64 FICON Ports
- 4 DOMAINS
- 1 & 2 port cards

FANOUT Cards




Concurrent add/delete

Description	F/C	Ports	Comments
PCIe copper fanout	0169	2 	To PCIe I/O Drawers (FC4003)
HCA3-O LR 1x IFB	0170	4 	PSIFB coupling 10 KM (100 KM repeated)
HCA3-O 12x IFB	0171	2 	PSIFB coupling 150 meters
HCA2-C copper fanout	0162	2	To I/O Drawers (FC4000)
HCA2-O 12x IB-DDR	0163	2	Coupling (150 meters)
HCA1-O 12x IB-SDR	0167	2	z9 feature for PSIFB to zEnterprise/z10
HCA2-O LR 1x IB-DDR Carry Forward only	0168	2	Coupling -10 KM (100 KM repeated) Withdraw from marketing on Dec 31, 2011 Carry Forward only

Each HCA3-O can communicate with the z10 via HCA2-O.
HCA3-O can not communicate with the z9.
HCA2 on zEnterprise can communicate with z9 HCA1.

Legacy Coupling Links



Description	F/C	Ports	Available	Comments
ISC-D	0217	N/A	New/Carry Forward	RoHS compliant – Mother Card
ISC-D	0218	1 to 2	New/Carry Forward	ISC-D (Daughter Card)
ISC-3 Link	0219	1 to 4	New/Carry Forward	Port(s) Enabled
RPQ	8P2197	2	New/Carry Forward	ISC-3 20 KM
ICB-3	0993		Not Available	Will be deleted
ICB-4	3393		Not Available	Will be deleted

**The z114 is the last server to offer ordering of new ISC-3 features.
ISC-3 requires an I/O Drawer.**

Server Time Protocol (STP) Required

Parallel Sysplex Coupling Connectivity


z9 EC and z9 BC S07
 IFB 12x SDR, ISC-3
 z9 to z9 IFB is **NOT** supported




ISC-3, 2 Gbps
 10/100 km

12x IFB, 3 GBps
 Up to 150 m

HCA1



z10 EC and z10 BC
 IFB 12x and 1x, ISC-3,



1x IFB, 5 Gbps
 10/100 km

ISC-3, 2 Gbps
 10/100 km

12x IFB, 6 GBps
 150 m

HCA2-O LR

HCA2-O

HCA2-O LR

ISC-3
 10/100 km

12x IFB3
 or IFB
 6 GBps
 150 m

1x IFB, 5 Gbps
 10/100 km

1x IFB, 5 Gbps, 10/100 km

ISC-3, 2 Gbps, 10/100 km



z800, z900
z890 and z990
 Not supported!

Note: ICB-4 and ETR are NOT supported on z196 or z114

*HCA2-O LR carry forward only on z196 and z114

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

System z – Maximum Coupling Links and CHPIDs



Server	1x IFB (HCA3-O LR)	12x IFB & 12x IFB3 (HCA3-O)	1x IFB (HCA2-O LR)	12x IFB (HCA2-O)	IC	ICB-4	ICB-3	ISC-3	Max External Links	Max Coupling CHPIDs
z196	48 M15 – 32*	32 M15 – 16*	32 M15 – 16* CF only	32 M15 – 16*	32	N/A	N/A	48	104 ⁽¹⁾	128
z114	M10 – 32* M05 – 16*	M10 – 16* M05 – 8*	M10 – 12 M05 – 8* CF only	M10 – 16* M05 – 8*	32	N/A	N/A	48	M10 ⁽²⁾ M05 ⁽³⁾	128
z10 EC	N/A	N/A	32 E12 – 16*	32 E12 – 16*	32	16 (32/RPQ)	N/A	48	64	64
z10 BC	N/A	N/A	12	12	32	12	N/A	48	64	64
z9 EC	N/A	N/A	N/A	HCA1-O 16 S08 - 12	32	16	16	48	64	64
z9 BC	N/A	N/A	N/A	HCA1-O 12	32	16	16	48	64	64

1. A z196 M49, M66 or M80 supports a maximum 104 extended distance links (48 1x IFB and 48 ISC-3) plus 8 12x IFB links.
A z196 M32 supports a maximum 96 extended distance links (48 1x IFB and 48 ISC-3) plus 4 12x IFB links*.
A z196 M15 supports a maximum 72 extended distance links (24 1x IFB and 48 ISC-3) with no 12x IFB links*.
2. z114 M10 supports a maximum of 72 extended distance links (24 1x IFB and 48 ISC-3) with no 12x IFB links*.
3. z114 M05 supports a maximum of 56 extended distance links (8 1x IFB and 48 ISC-3) with no 12x IFB links*.

* Uses all available fanout slots. Allows no other I/O or coupling.

SAN Connectivity



Description	F/C	Ports	Available	Comments
FICON Express8S 10Km LX	0409	2	New	Initial orders
FICON Express8S SX	0410	2	New	Initial orders
FICON Express4-2C SX	3318	2	Carry Forward	New during upgrade. RPQ 8P2534 if FC4000 slots are open and PCIe drawer is full.
FICON Express4 10KM LX	3321	4	Carry Forward	
FICON Express4 SX	3322	4	Carry forward	
FICON Express4-2C 4KM LX	3323	2	Carry Forward	New during upgrade. RPQ 8P2534 if FC4000 slots are open and PCIe drawer is full.
FICON Express4 4KM LX	3324	4	Carry Forward	
FICON Express8 10KM LX	3325	2	Carry Forward	New during upgrade. RPQ 8P2534 if FC4000 slots are open and PCIe drawer is full.
FICON Express8 SX	3326	2	Carry Forward	New during upgrade. RPQ 8P2534 if FC4000 slots are open and PCIe drawer is full.


Open Systems Adapter - in PCIe I/O Drawer



Description	F/C	Ports	Available
OSA-Express4S GbE LX	0404	2	New Build
OSA-Express4S GbE SX	0405	2	New Build
OSA-Express4S 10 GbE Long Reach	0406	1	New Build
OSA-Express4S 10 GbE Short Reach	0407	1	New Build

Note: OSA-Express3 1000Base-T requires an I/O Drawer (FC4000)

Open Systems Adapter --- FC4000 Drawer



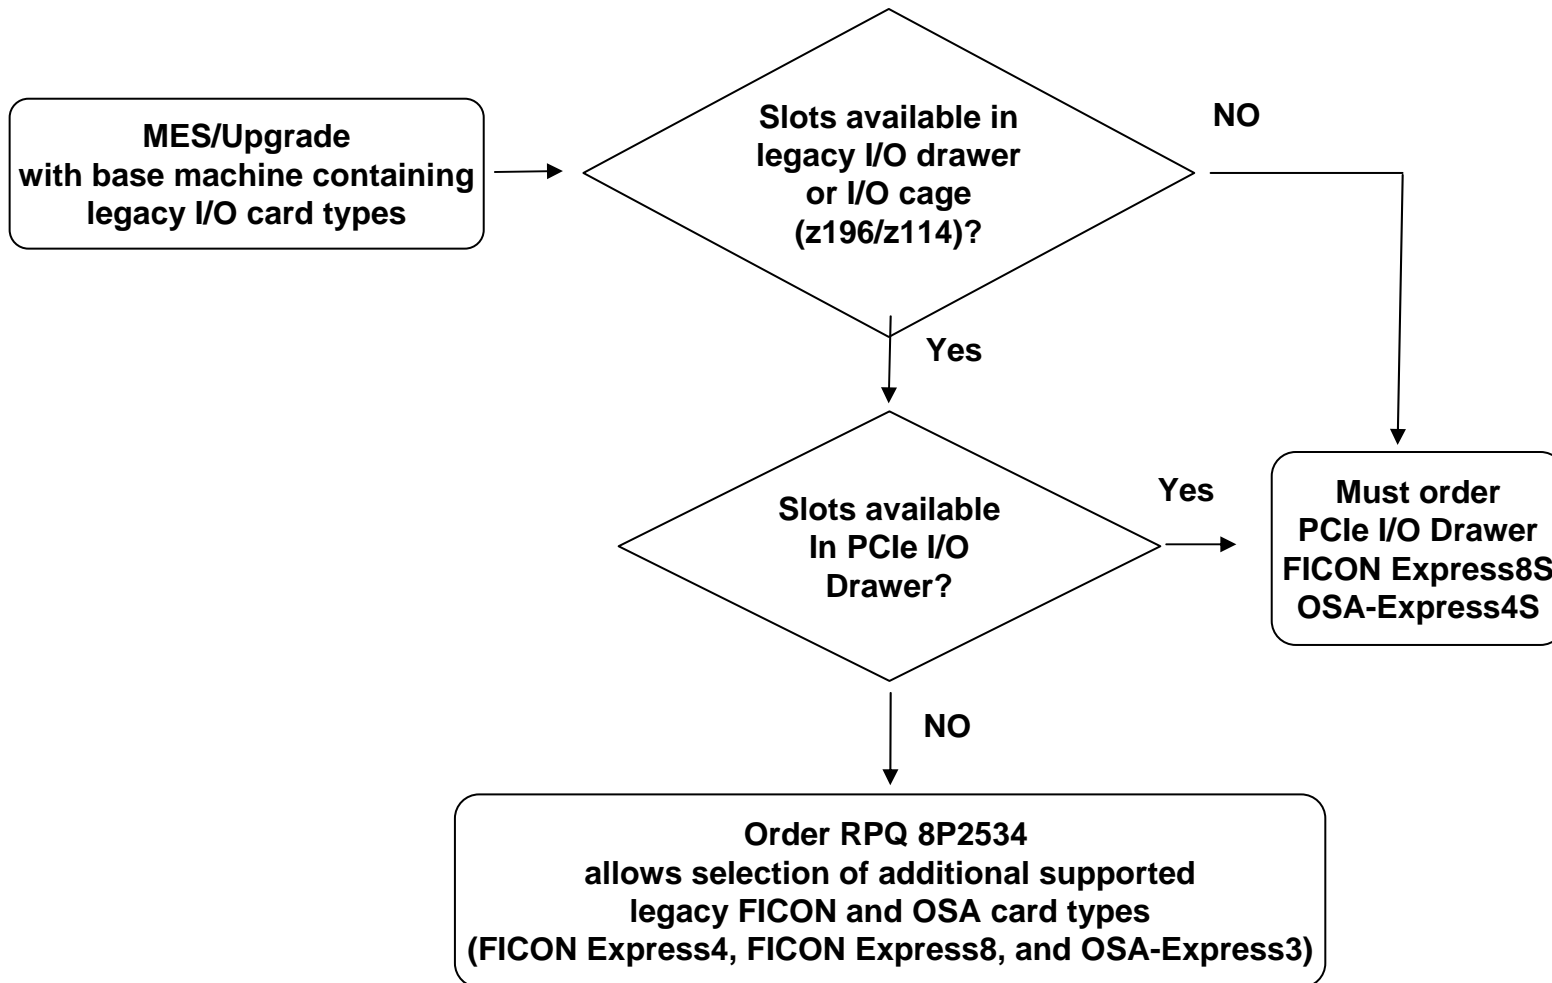
SHARE
Technology - Connections - Results

Description	F/C	Ports	Available	Comments
OSA-Express	23xx/13xx	2	NO	
OSA-Express3 GbE LX	3362	4 ¹	Carry Forward	Note 2
OSA-Express3 GbE SX	3363	4 ¹	Carry Forward	Note 2
OSA-Express2 GbE LX	3364	2	Carry Forward	
OSA-Express2 GbE SX	3365	2	Carry Forward	
OSA-Express2 1000BASE-T	3366	2	Carry Forward	
OSA-Express3 1000BASE-T	3367	4 ¹	New/Carry Forward	
OSA-Express2 10 GbE Long Reach	3368	1	NO	
OSA-Express3-2P 1000BASE-T	3369	2 ¹	New/Carry Forward	
OSA-Express3 10 GbE Long Reach	3370	2	Carry Forward	Note 2
OSA-Express3 10 GbE Short Reach	3371	2	Carry Forward	Note 2
OSA-Express3-2P GbE SX	3373	2 ¹	Carry Forward	Note 2

¹ two ports per CHPID

² New during upgrade. RPQ 8P2534 if FC4000 slots are open and PCIe drawer is full.

RPQ 8P2534 – Access to Legacy IO



RPQ 8P2534 notes:

RPQ is only available for MES and not New Build, Migration Offerings, or System z Exchange Program (formerly known as a hybrid offering)
 Does not allow for the addition of I/O drawers, only allows for existing drawers to be filled with current legacy I/O cards.

Statements of Direction



- **Application Program Interfaces (APIs) for Unified Resource Manager:**
IBM intends to offer Application Program Interfaces (APIs) for IBM zEnterprise Unified Resource Manager. These APIs provide access to the same underlying functions that support the Unified Resource Manager user interface.

IBM plans to enhance Tivoli's Integrated Service Management for System z portfolio of products to provide integrated end-to-end monitoring, alerting, discovery, automation, storage, and security solutions to take advantage of the zEnterprise ensemble monitoring and management capabilities provided by the API support.
- **Dynamic discovery of storage resources by Unified Resource Manager:**
IBM intends to offer dynamic discovery of storage resources by Unified Resource Manager. A server administrator will be able to trigger discovery of additional storage resources through the user interface of Unified Resource Manager.
- **Microsoft Windows support:**
In the fourth quarter of 2011, IBM intends to support running the Microsoft Windows operating system on select IBM BladeCenter HX5 blades installed in the IBM zEnterprise BladeCenter Extension (zBX) Model 002.



Statements of Direction

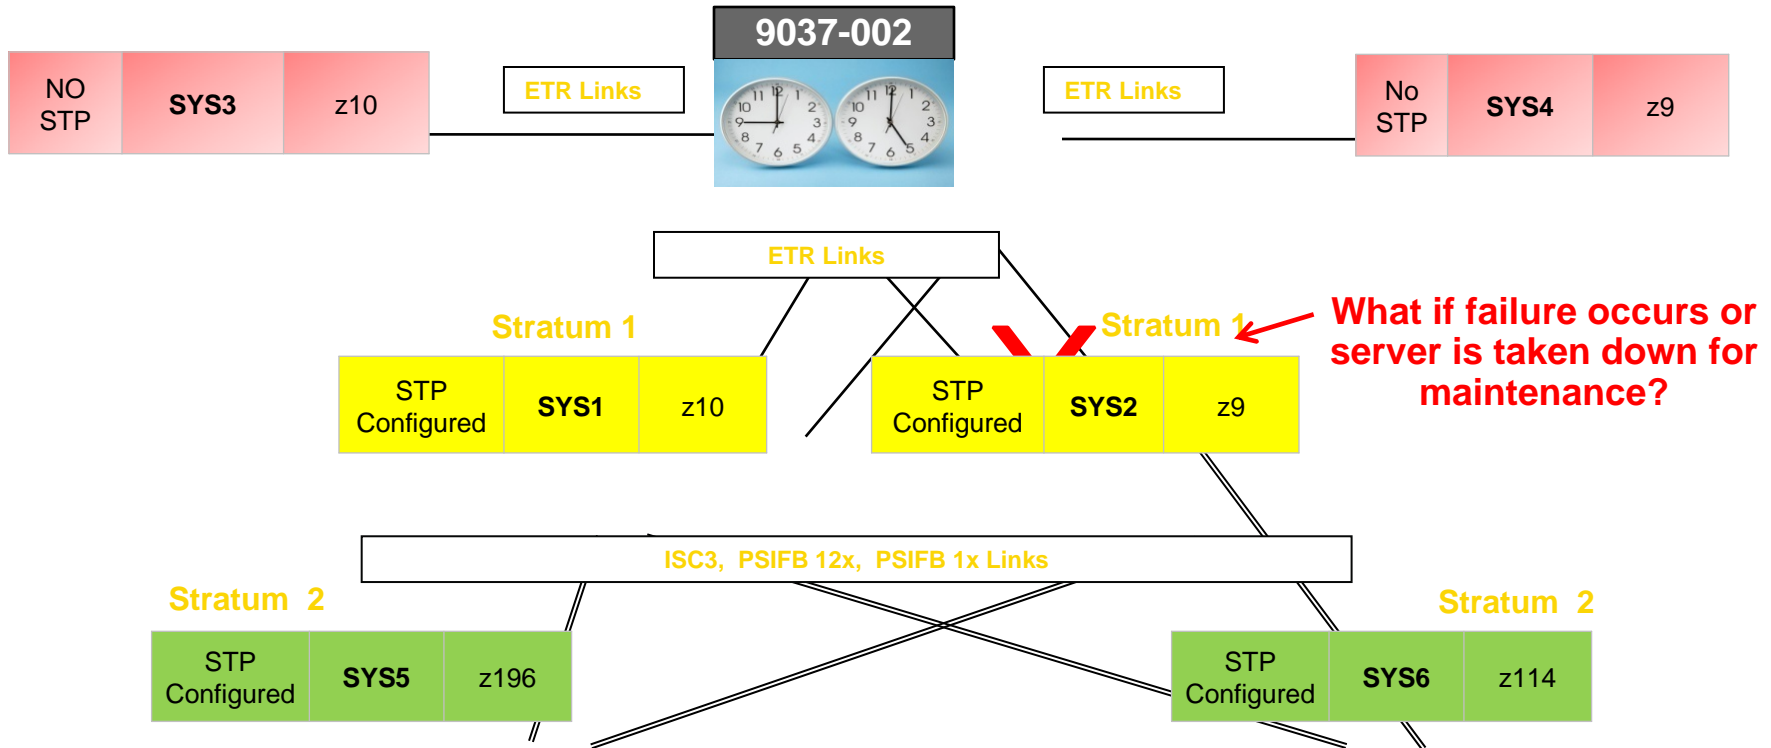


- **HiperSockets integration with the IEDN:**
Within a zEnterprise environment, it is planned for HiperSockets to be integrated with the intraensemble data network (IEDN), extending the reach of the HiperSockets network outside of the central processor complex (CPC) to the entire ensemble, appearing as a single Layer 2 network. HiperSockets integration with the IEDN is planned to be supported in z/OS V1.13 and z/VM in a future deliverable.
- **HiperSockets Completion Queue:**
IBM plans to support transferring HiperSockets messages asynchronously, in addition to the current synchronous manner on z196 and z114. This could be especially helpful in burst situations. The Completion Queue function is designed to allow HiperSockets to transfer data synchronously if possible and asynchronously if necessary, thus combining ultra-low latency with more tolerance for traffic peaks. HiperSockets Completion Queue is planned to be supported in the z/VM and z/VSE environments.
 - **z/VSE support of HiperSockets Completion Queue:**
z/VSE plans to exploit HiperSockets Completion Queue in a future deliverable.
- **z114 will be the last server to support ESCON channels:**
System z customers should continue to eliminate ESCON channels from the mainframe wherever possible. Alternate solutions are available for connectivity to ESCON devices. IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to FICON. This offering is designed to help customers to simplify and manage a single physical and operational environment - FICON channels on the mainframe with continued connectivity to ESCON devices.

Statements of Direction

- **z114 will be the last server to support FICON Express4 channels:**
 - Enterprises should migrate to FICON Express8s channels.
- **z114 will be the last server to support OSA-Express2 features:**
 - Enterprises should migrate to OSA-Express4S features.
- **z114 will be the last server to offer ordering of ISC-3:**
 - Enterprises should migrate to 12x InfiniBand or 1x InfiniBand LR coupling links.
- **z114 will be the last server to offer ordering of the PSC feature.**
- **z114 will be the last server to support dial-up modems** for use with the Remote Support Facility (RSF), and the External Time Source (ETS) option of Server Time Protocol (STP).

No Support for the 9037 Sysplex Timer





- It is possible to have a z114 server as a Stratum 2 or Stratum 3 server in a Mixed CTN linked to z10s or z9s (STP configured) attached to the Sysplex Timer operating as Stratum 1 servers
- Two Stratum 1 servers are highly recommended to provide redundancy and avoid a single point of failure
- Suitable for a customer planning to migrate to an STP-only CTN.
- The z114 can not be in the same Mixed CTN as a z990 or z890 (n-2)


z114 Power and I/O Cabling Options




z114 Raised Floor



z114 Non-Raised Floor



Floor Independent



Reclassification from “general business” environment to “data center”

Agenda

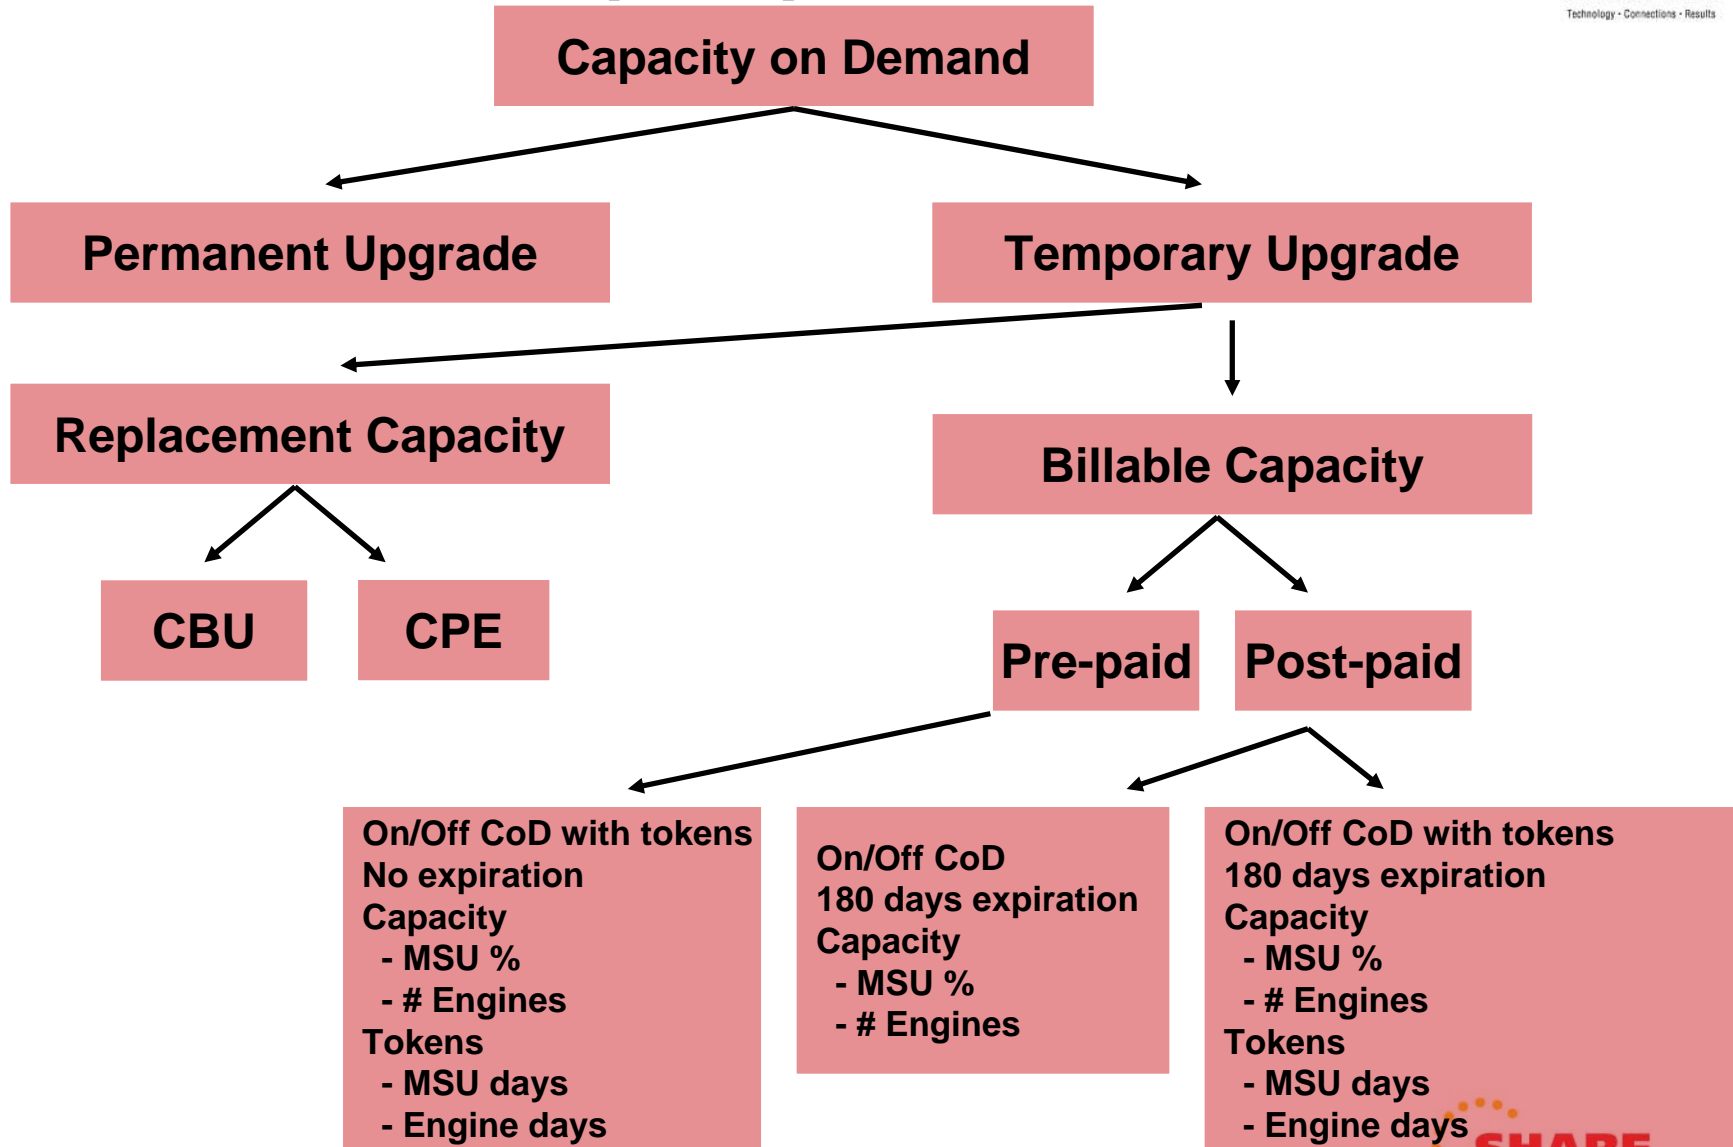
- zEnterprise 114 Overview
 - z BladeCenter Extension (zBX)
 - Functions
 - Performance
 - Upgrades
 - Memory
- I/O, Security, Miscellaneous
 - I/O Drawers
 - I/O Features
 - Discontinued I/O Features
 - Cryptography
 - Server Time Protocol
 - Installation Options
- Capacity on Demand Enhancements
- Operating Systems
- Hardware Management Console



The Basics – Temporary Upgrades

- Capacity Backup (CBU)
 - Predefined capacity for disasters on a other “lost” server(s)
 - Concurrently add CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs
 - Pre-paid
- Capacity for Planned Events (CPE)
 - CBU-like offering, when a disaster is not declared
 - Example: System migration (push/pull) or relocation (data center move)
 - Predefined capacity for a fixed period of time (3 days)
 - Pre-paid
- On/Off Capacity on Demand (On/Off CoD)
 - Satisfy periods of peak demand for computing resources
 - Concurrent 24 hour rental of CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs
 - Supported through a new software offering – Capacity Provisioning Manager (CPM)
 - Post-paid & Pre-Paid

z114 – Basics of Capacity on Demand



z10 to z114 Capacity on Demand Enhancements

System z10


Separate orders for purchase of unassigned engines

On/Off CoD records must be replenished manually

CoD records staged on machine deliver

No On/Off CoD administrative test


z114



Unassigned engine purchase via CIU



Auto replenishment of On/Off CoD records



Manufacturing install of up to 4 CoD records with system ship on a new build.



On/Off CoD Administrative tests

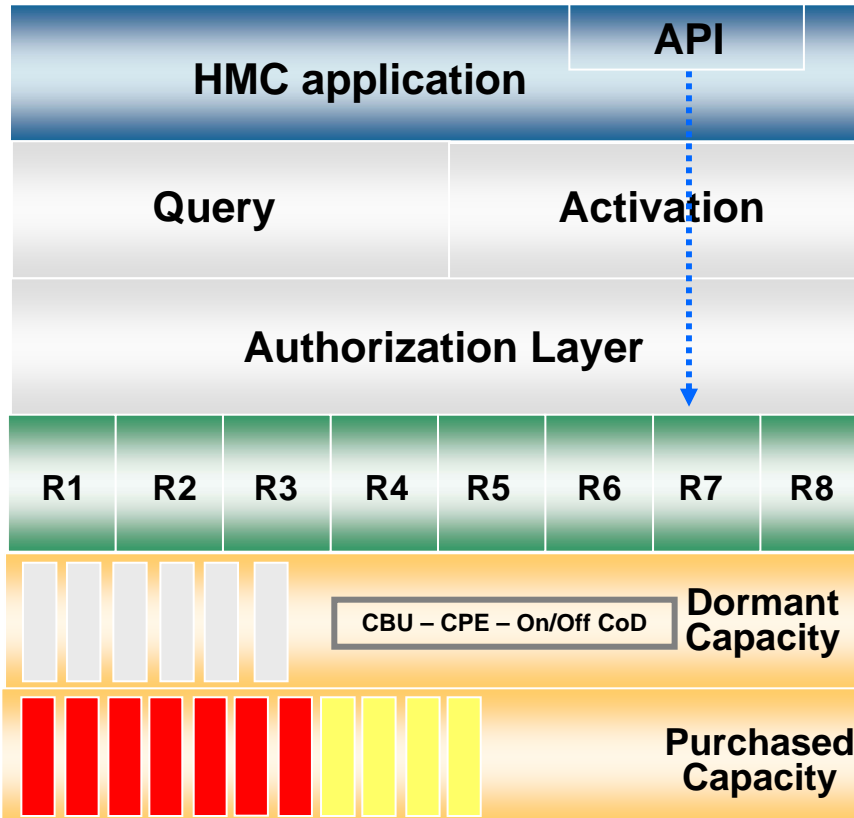
CoD Provisioning Architecture

Manual operations

Customer defined policy or user commands

CPM (z/OS 1.9 or higher)

Orders downloaded from Retain/media to SE hard drive




* Only one On/Off CoD record can be active

Enforce Terms and Conditions and physical model limitations

Up to 8 records installed and active on the CEC and up to 200 records staged on the SE

Change permanent capacity via CIU or MES order


Base Model



SHARE
olutions - Results

Agenda

- zEnterprise 114 Overview
 - z BladeCenter Extension (zBX)
 - Functions
 - Performance
 - Upgrades
 - Memory
- I/O, Security, Miscellaneous
 - I/O Drawers
 - I/O Features
 - Discontinued I/O Features
 - Cryptography
 - Server Time Protocol
 - Installation Options
- Capacity on Demand Enhancements
- Operating Systems
- Hardware Management Console



Operating System Support for z114

- **Currency is key to operating system support and exploitation of future servers**
- **The following are the minimum operating systems planned to run on z114**

<i>Operating System</i>	<i>Supported levels</i>
z/OS	<ul style="list-style-type: none">• V1.11, 1.12, 1.13 or higher• V1.10* (requires Lifecycle Extension after Sept. 30, 2011)• V1.8 and 1.9, in Lifecycle Extension• zBX Ensemble support: z/OS V1.10* or higher
Linux	<ul style="list-style-type: none">• Red Hat RHEL 5• Novell SUSE SLES 11
z/VM	<ul style="list-style-type: none">• V5.4• zBX Ensemble support: V6.1
z/VSE	<ul style="list-style-type: none">• V4.2• zBX Ensemble support V4.3 or higher
z/TPF	<ul style="list-style-type: none">• V1.1 or higher



* z/OS 1.10 support ends Sept. 30, 2011, and Lifecycle Extension is required after that date.

FCP end-to-end data checking



- Enhanced RAS for business critical applications
- Supporting ANSI T10 DIF standard and its extensions
- SCSI device support for cyclical redundancy check (CRC) protection
 - For end-to-end data integrity
- Extension to standard allows for
 - Checksum protection
 - CRC protection
 - Requires support by control unit
- z/VM V5.4, V6.1 guest exploitation
- Linux on System z


Enhances the RAS characteristics for FCP links through the continued adoption of industry standard protocols.

IPL for alternate subchannel set



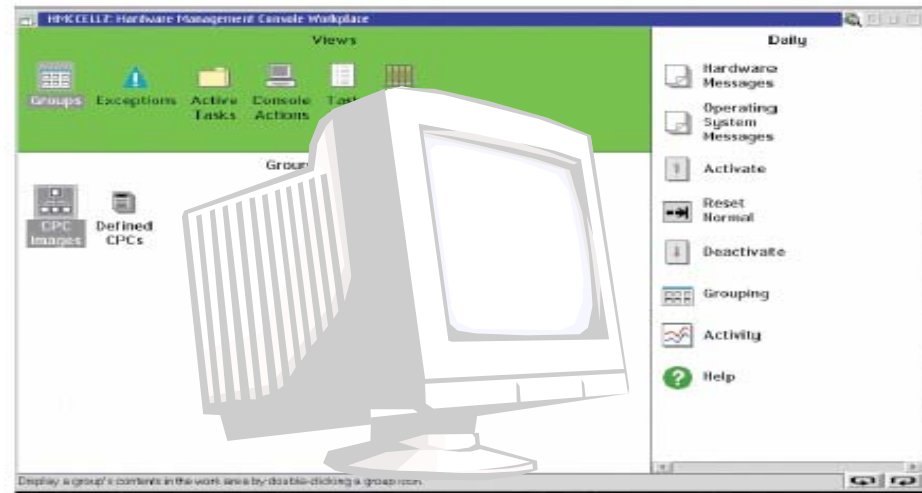
- Ability to IPL from an alternate subchannel set
 - No longer require IPL devices to be in subchannel set 0
 - Utilize additional device addressability with reduced complexity
- z/OS V1.13
- z/OS V1.11 and V1.12 with PTFs

Reduces the complexity for customers utilizing GDPS HyperSwap and exploiting the alternate subchannel set for PPRC secondary devices.



Agenda

- zEnterprise xxx Overview
 - z BladeCenter Extension (zBX)
 - Functions
 - Performance
 - Upgrades
 - Memory
- I/O, Security, Miscellaneous
 - I/O Drawers
 - I/O Features
 - Discontinued I/O Features
 - Cryptography
 - Server Time Protocol
 - Installation Options
- Capacity on Demand Enhancements
- Operating Systems
- Hardware Management Console



Hardware Management Console Features

Description	F/C	Available	Comments
HMC w/Dual Ethernet	0091	New	
22" Flat Panel Display	6096	New	
HMC w/Dual Ethernet	0084	Carry Forward	Can be decremented Can not be used with Unified Resource Manager
HMC w/Dual Ethernet	0090	Carry Forward	
17" Flat Panel Display	6094	Carry Forward only	
20" Flat Panel Display	6095	Carry Forward only	
Ethernet Switch	0089	Carry Forward	10/100 mbps
Ethernet Switch	0070	Yes, Carry Forward	10/100/1000 mbps

Primary and Alternate Hardware Management Consoles




- **Any V2.11.1 HMC can become the Primary HMC that controls the ensemble**
 - The Primary HMC can perform all non-ensemble HMC functions on CPCs that aren't members of the ensemble
- **The HMC that creates an ensemble (the HMC that performed the "Create Ensemble" wizard) becomes the Primary HMC**
- **The Alternate HMC is specified when executing the "Create Ensemble" wizard**
 - Any V2.11.1 HMC is eligible to be an Alternate HMC after running the "Manage Alternate Hardware Management Console task"
- **The title of Primary Hardware Management Console and Alternate Hardware Management Console will appear on the Login HMC panel and the title line once you are logged in**
 - The default HMC titles will change to these titles when the ensemble is created
 - The titles will revert back to the default if the ensemble is deleted
- **A Primary HMC is the only HMC that can perform ensemble related management tasks (create virtual server, manage virtual networks, create workload)**

zEnterprise 114 Summary

- Integration with the zEnterprise BladeExtension
- Two Models
- Increased capacity in a single footprint
 - Designed for up to 1.2 times the z10 BC in total system capacity
 - 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





SHARE

Technology - Connections - Results

Thank you

- Backup


AES KEKs and typed AES keys

- Previous AES support was only for DATA keys
 - For data encrypt/decrypt only
 - No associated key usage attribute
- Enhanced to support AES IMPORTER and EXPORTER KEKs as well as CIPHER keys
- Key_Generate2 generates keys in pairs, similar to Key_Generate
 - Can be wrapped under AES master key or an AES KEK
- Symmetric Key_Import & Symmetric_Key_Export enhanced to transport keys wrapped by AES KEKs
- Other updated verbs for key part import, key test, prohibit export, key translate

KEK = Key Encryption Key

TR-31


- TR-31 is a method for wrapping TDES keys
 - Not a “standard”, but an example method conforming to X9.24-1
 - Wraps key material + key attributes
- Key types and usage attributes are not standardized
 - TR-31 and CCA differ quite a bit
 - Different implementations of TR-31 interpret attributes differently!
- CCA will provide functions to convert between CCA and TR-31 key formats
 - Convert either internal or external CCA tokens
 - Restrictions to prevent security attacks, such as key type changes
 - CCA will have access control to let you choose which to allow
- TR-31 key blocks cannot be directly used for CCA crypto functions
 - Must first be converted to CCA key tokens
 - TR-31 is provided as an interchange format



• Failure to comply with this guideline may be cause for any one of the following payment networks to refuse connection: Star, Pulse, and NYCE.

• The Payment Card Industry has mandated that PIN Pad manufacturers provide for TR-31 or equivalent methodology.

EC-DH key establishment



- The problems:
 - RSA is not strong enough to transport long AES keys
 - Some organizations (such as NSA) have standardized on EC key management
- CCA is adding EC-DH for key agreement
 - CCA operational keys can be established using EC-DH techniques
 - Same curves (NIST and Brainpool) that are supported for ECDSA
- Note that this is a key agreement protocol, not key transport
 - The EC algorithm does not support key transport

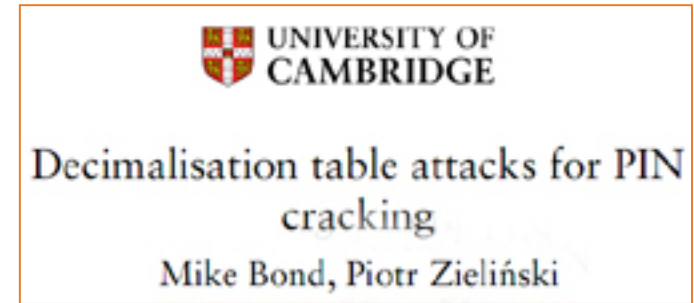
en.wikipedia.org

Elliptic curve Diffie–Hellman (ECDH) is a key agreement protocol that allows two parties, each having an elliptic curve public-private key pair, to establish a shared secret over an insecure channel.

This shared secret may be directly used as a key, or better yet, to derive another key which can then be used to encrypt subsequent communications using a symmetric key cipher. It is a variant of the Diffie–Hellman protocol using elliptic curve cryptography.

Protection of decimalization tables for financial security

- Researchers identified attacks based on unprotected PIN decimalization tables
- Blocking the attacks requires use of only approved decimalization tables.
 - This is different than using encrypted tables!
- CCA solution is to allow using decimalization tables stored inside the HSM
HSM = Hardware Security Module
 - Tables are managed using secure techniques
 - Potential attacker cannot tell HSM to use a table of his design

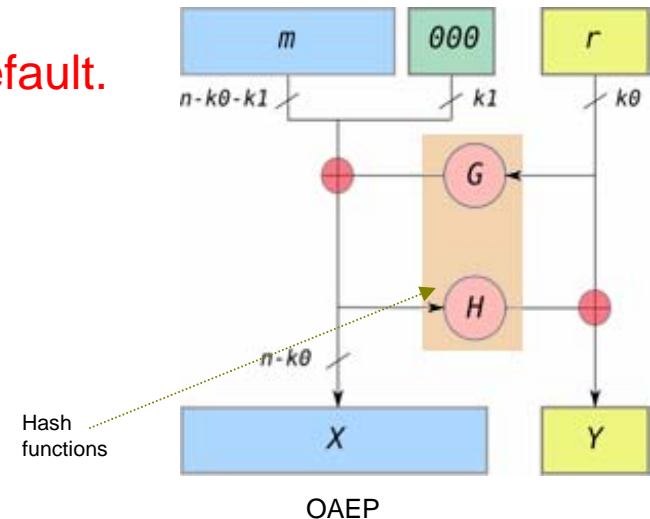


en.wikipedia.org

A decimalization table attack is a technique that may allow a corrupt insider at a bank to discover Personal Identification Numbers (PINs) by exploiting a design flaw in the Hardware Security Module used to protect the PIN.

RSA-OAEP with SHA-256

- Extends RSA key management to **add SHA-256** hash method for RSA-Optimal Asymmetric Encryption Padding (OAEP) method
 - CSNDSYX (Symmetric key export under RSA)
 - CSNDSYI (Symmetric key import under RSA)
 - CSNDSYG (Generate symmetric key wrapped by RSA)
- **Previous implementation used only SHA-1.**
 - **Hash method is now an option. SHA-1 is the default.**
- Can be used for AES or DES/TDES DATA keys.



Optimal asymmetric encryption padding