


Systems and Technology Group

IBM System z9 Technology Update & Trends *A System z9 for Everyone*

Svend Erik Bach
Distinguished Engineer, IBM Nordic/Europe
seb@dk.ibm.com



LSU 2006, October-November, 2006
25 years of I/T education and positioning

IBM Systems

Trademarks

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

AIX*	GDPS*	Parallel Sysplex*	VisualAge*
CICS*	HiperSockets	PR/SM	VM/ESA*
DB2*	IBM*	RACF*	VSE/ESA
DB2 Connect	IBM eServer	Rational*	VTAM*
DB2 Universal Database	IBM logo*	RMF	WebSphere*
DirMaint	IMS	System i	z/Architecture
Domino	Language Environment*	System z	z/OS*
DRDA*	Lotus*	System z9	z/VM*
Enterprise Storage Server*	MQSeries*	System Storage	z/VSE
ESCON*	Multiprise*	Tivoli*	zSeries*
FICON*	OMEGAMON*	TotalStorage*	zSeries Entry License Charge
FlashCopy*	OS/390*	Virtualization Engine	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Intel is a trademark of Intel Corporation in the United States, other countries, or both.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Linux is a trademark of Linus Torvalds in the United States and other countries..

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

Microsoft is a registered trademark of Microsoft Corporation in the United States and other countries.

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

Notes:

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

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

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

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

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

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

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

Notable quotable...

“I predict that the last mainframe will be unplugged on March 15, 1996”

Stewart Alsop, former InfoWorld columnist (now at Fortune Magazine), March, 1991



- Annual growth in MIPS of over 30% since 1992
- More than \$20B in mainframe revenue since 1996 (when the last one was to have been unplugged...)
- But there have been many changes in the mainframe since 1991!
- Prior to the mid-90's, mainframes were the IBM "cash cow". Revenues are still healthy and margins high, but things have changed...

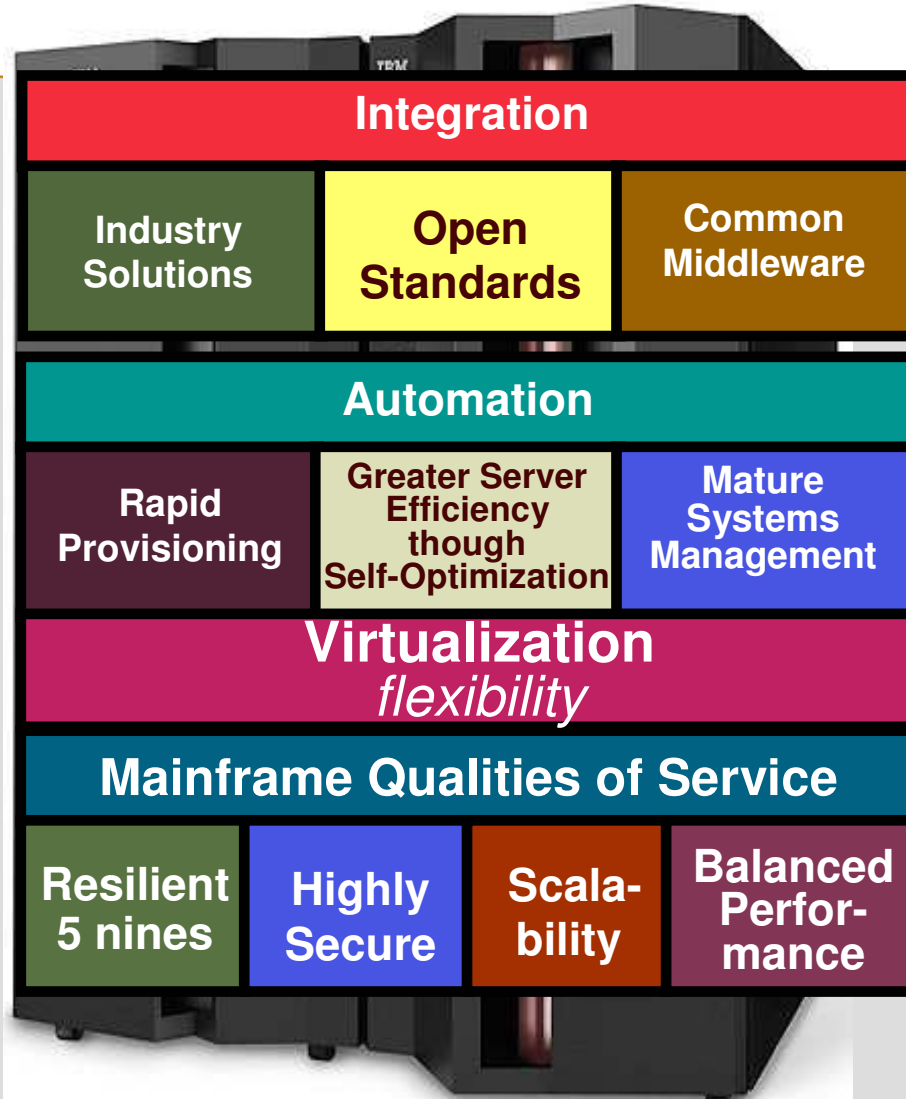
System z Core Values - 40 years investments...

still and increasingly at work

uniquely positioned for the on demand challenge

Business Challenges

- Financial Pressures
- **Increase Revenue**
- **Accelerate Time-to-Market**
- Integration
- Deploy New Capabilities
- Security and Operational Resiliency



IT Challenges

- **Respond to changing business needs**
- Meet service levels
- Help manage risk
- Increased server utilization
- Simplify Infrastructure Complexity
- **Reduce IT Costs**

Investment Flexibility

2Q Highlights - System z

- **Market Share 34%, up 3% YTY - (*up 14% over the last 5 years...*)**
- **MIPS Growth +7% YTY**
- **2% Worldwide Total revenue growth; HW Revenue up 7% YTY**
- **YTY Revenue growth in all Geo's except Asia Pacific**
- **z more “*open*” = accelerated acceptance of zAAP, zIIP, IFL engines**
- **System z Linux and Java workloads show very healthy upswing**
- **60% of total revenue driven by new workloads - SOA, Java, DB, Linux**



The Mainframe Charter – investing in the future



Innovation

- **Focus on enterprise wide roles**
 - ▶ Security – Security Hub
 - ▶ BR and Workload management
 - ▶ Data Hub and Business Integration
 - ▶ On Demand solutions
 - ▶ Continue to “Raise the Bar” on technology leadership
 - ▶ **SIMPLIFICATION**



Value

- **Attractive for new workloads**
- **Continued focus on specialty engines & accelerators**
- **Drive granularity to support broad market**
- **Generation to generation price / performance improvements**



Community

- **Drive ISV applications & strengthen partner relationships**
- **Build new skills in marketplace, including next generation**
- **Focus on emerging geographies**
- **SIMPLIFICATION**

System z9 delivers price / performance and investment flexibility for on demand computing

Generation to generation price / performance improvements:	z9 BC	z9 EC
Reduction in chargeable MSUs versus z890 / z990	10%	10%
Reduction in chargeable MSUs versus z800 / z900	19%	19%
Reduction in maintenance costs (*) (up to)	20%	20%
Hardware performance improvement for IFL (Linux), zIIP (DB2, etc.), zAAP (Java) and Internal coupling facilities (ICF) (*) up to	37%	35%
Typical charge for MES upgrades for IFLs and zAAPs	0	0
Technology-driven value	z9 BC	z9 EC
Number of capacity settings	73	78
Specialty engines (IFL, zAAP) and the new System z9 Integrated Information Processors (zIIPs) which can help reduce the cost of certain DB2® Data Serving Workloads (**)	\$95k	\$125k
IBM SW charges for zAAP & zIIP capacity	0	0

Plus

- IBM MLC SW Cost per MIP is reduced with a factor 2 since 2001
- Significant price reductions for memory etc..
- On/Off Capacity on Demand (On/Off CoD) enhancements to better manage volatile business requirements

(*) – comparisons shown are z9 BC vs. z890 and z9 EC vs. z990
 (**) Prices may vary by country

The Mainframe Charter

Investing in the future

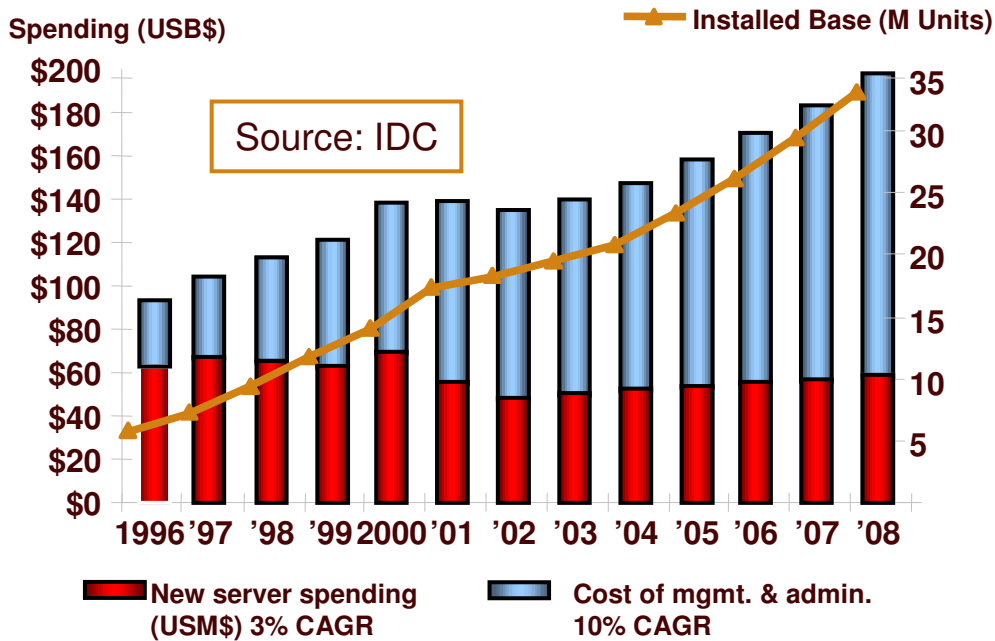
Innovation

Value

Community



Cost of Complexity – System z is unique in managing growth and cost



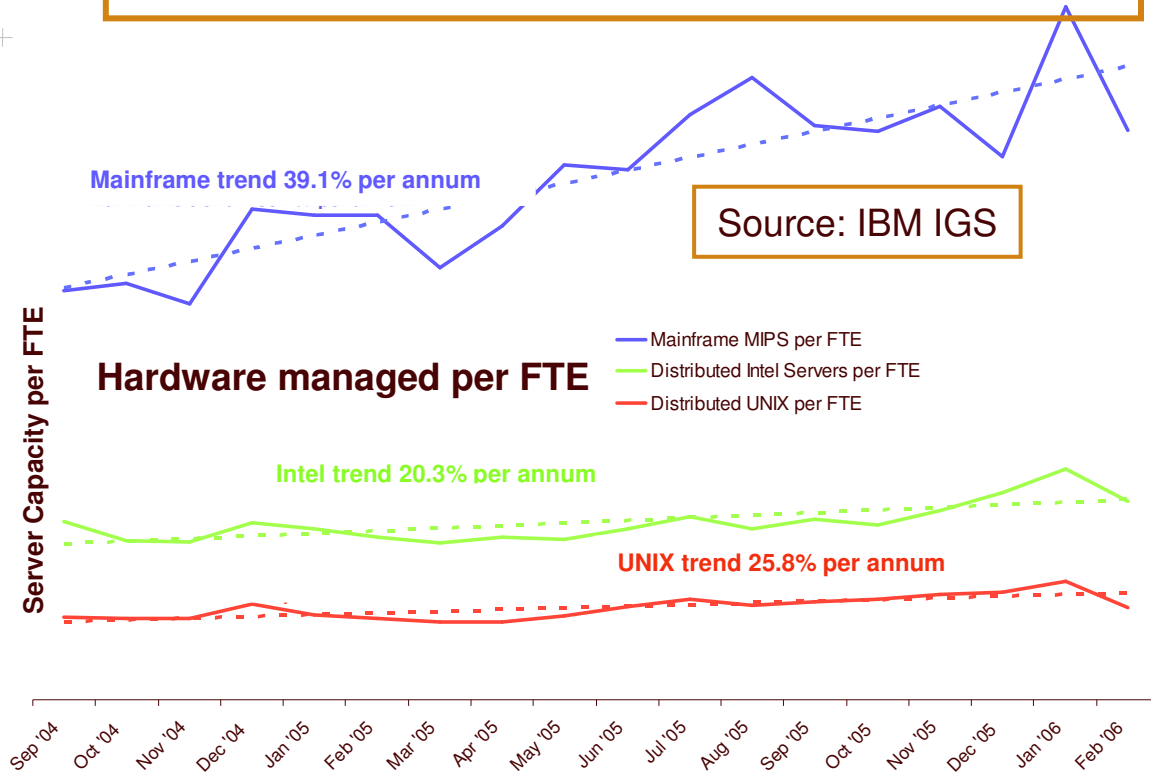
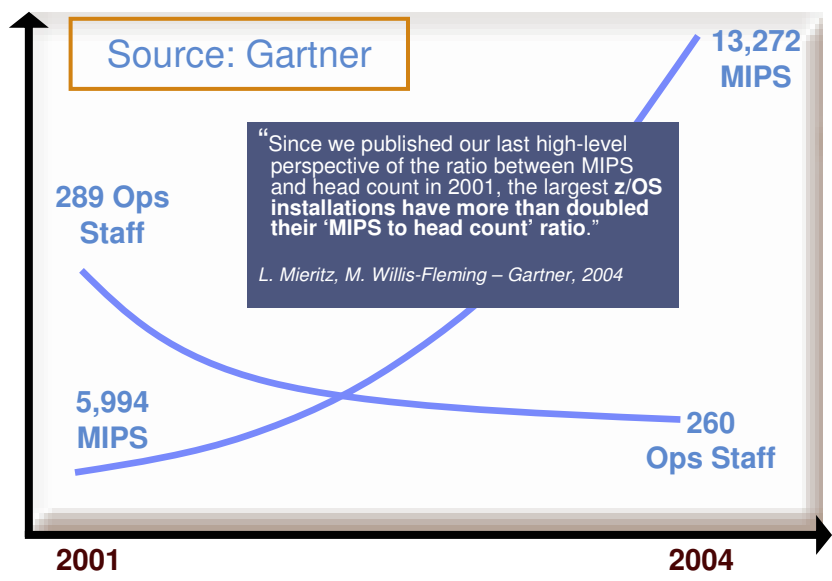
IBM IGS Observation

Total HW/SW/Operational cost / MIP has been reduced by a factor 2 over the last 4 years

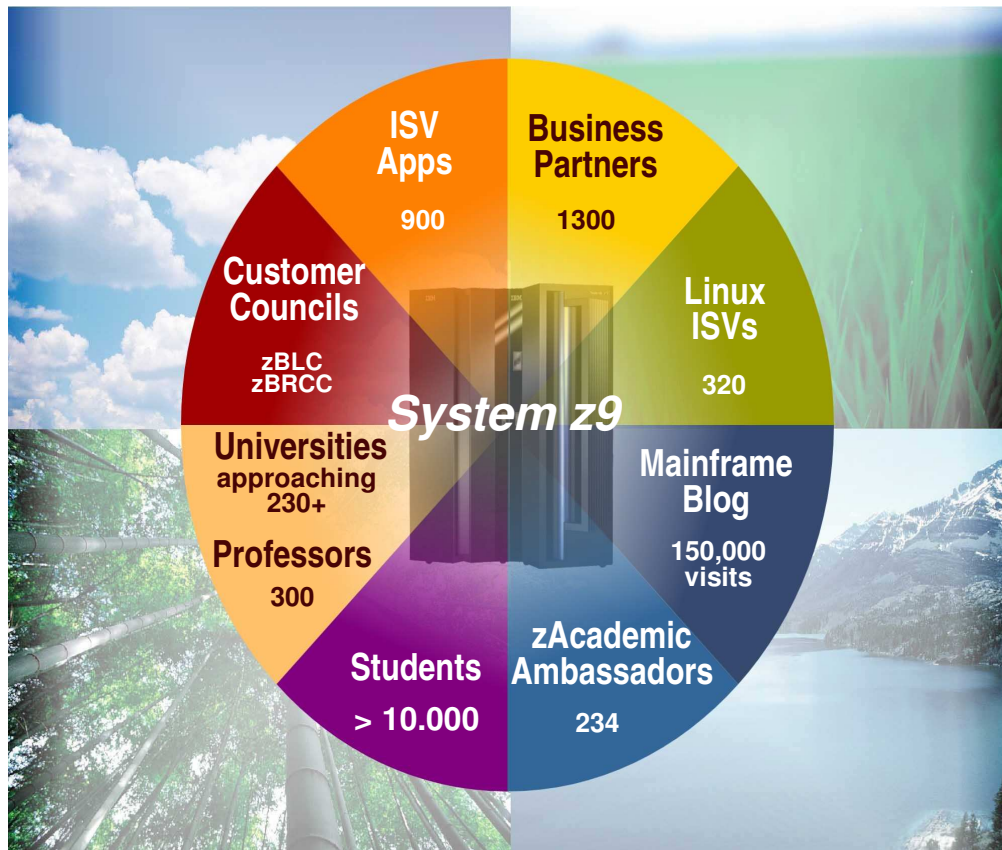
Predicted average cost per end user in 2010:

- Mainframes \$6,250
- Unix Minis \$19,000
- PC Servers \$24,000

5yr costs for hardware, software and maintenance
Arcati Research 2005 – The Dinosaur Myth 2004 Update



Mainframe Community Ecosystem



Customer councils

IBM Academic Initiative

- **GOAL - 20,000 new skills by 2010**
✓ *more than 10,000 so far*
- **230+ colleges and universities**
✓ *with 50% outside US*
- **300 professors registered**
- **200+ IBM ambassadors**
- **12 Enterprise systems courses**
- **New e-learning courses developed**

Partners – heavy investments

- **300+ mainframe Linux ISVs**
- **1,300 mainframe partners**
- **training, loaner program, discounts**
- **24/7 HW/SW/support**

Mainframe blog

- <http://mainframe.typepad.com>

Meet the Community” via the mainframe community portal

<http://www.ibm.com/servers/eserver/zseries/about/charter/community.html>

z Solution and Technology Rollouts

heavy application focus

	3Q06	4Q06	1Q07	2Q07	3Q07	4Q07
Marketing Theme	<p>Information Management</p> <ul style="list-style-type: none"> Enterprise Security Data Serving <p>WebSphere</p>	<ul style="list-style-type: none"> Virtualization Consolidation SOA 	<ul style="list-style-type: none"> Virtualization Consolidation 	<ul style="list-style-type: none"> Data Serving Platform Competitiveness Performance 	<ul style="list-style-type: none"> SOA Enterprise Security 	<ul style="list-style-type: none"> Platform Competitiveness Application Enablement
Solutions	<ul style="list-style-type: none"> System z SAP Advantage for SAP Oracle DB on Linux consolidation <p>ORACLE</p>	<p>ACI</p> <ul style="list-style-type: none"> Electronics Payments with ACI and eFunds <p>E-FUNDS</p>	<p>ORACLE</p> <ul style="list-style-type: none"> Oracle App Server on Linux for System z 	<p>SAP</p> <p>Back end Retail with SAP</p>	<p>WebSphere</p> <p>SOA Solution for Insurance</p>	<p>Information Management</p> <ul style="list-style-type: none"> Data Warehousing / Business Intelligence <p>Information Builders</p>
Technology	<ul style="list-style-type: none"> Security Tape Encryption Data Warehousing (SWG) 	<ul style="list-style-type: none"> Data Hub DB2 V9 	<ul style="list-style-type: none"> Linux virtualization 	<ul style="list-style-type: none"> Data Hub Performance SOA (SWG) 	<ul style="list-style-type: none"> Security 	<ul style="list-style-type: none"> Application Enablement
Industry Focus	<ul style="list-style-type: none"> Industrial, Govt & Financial 	<ul style="list-style-type: none"> Banking & Financial Markets 	<ul style="list-style-type: none"> Cross-industry and FSS 	<ul style="list-style-type: none"> Retail 	<ul style="list-style-type: none"> Insurance, Govt, & Financial 	<ul style="list-style-type: none"> Industrial

The mainframe as a platform for people integration

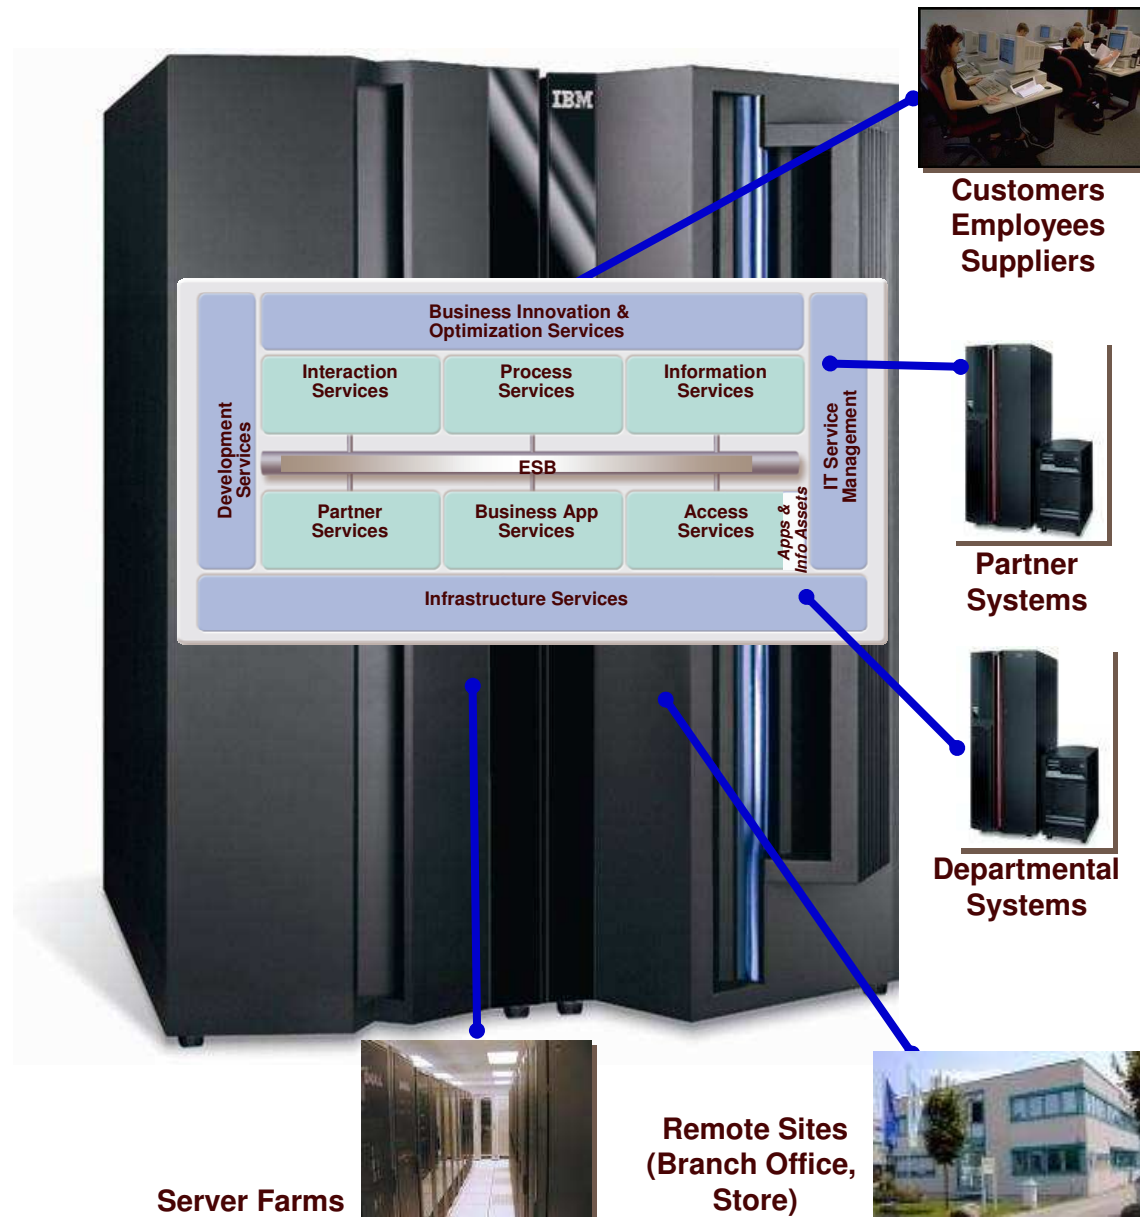
Delivers all the essential SOA qualities

Architecturally compatible:

- SOA Reference Architecture
- Architectural services
- Communications and integration
- Process control
- Reuse of core assets
- Helps address industry standards

Operationally superior:

- End-to-end **security** features in the enterprise
- **24/7 availability**
- **Massive scalability**
- **Automated recoverability**
- **Centralized operations**





System z9 - today

**Unlimited Scalability - High Flexibility
Resilience - Security
Integration**

**Continued effective price reduction
HW and SW**

IBM Announces the System z9

Investment:

- 3 years
- \$1.2 billion
- 5,000 tech professionals

System z - today

Scalability - High Flexibility
Resilience - Security
Integration

Continued effective price reduction
HW and SW

z9 EC Enterprise Class

Up to 64 engines (cores)
Multiple Capacity levels
zAAP, zIIP, IFL engines
RoHS compliant
Investment protection
Upgrade from z900, z990, z9-BC



z9 BC Business Class

Up to 8 engines (cores)
1-3 way + speciality engines
0-4 way + speciality engines
Multiple Capacity Levels
zAAP, zIIP, IFL engines
RoHS compliant
Investment protection
Upgrade from z800-4, z890



26 to 17500++ MIPS
Corresponds to:
3000 to more than 2 million TPC-C



System z9 - Scalability and Upgrade strategy

Each new range has delivered:

- New function
- Improved performance
- Improved availability and serviceability
- Better price/performance
- Field Upgrades to protect investments

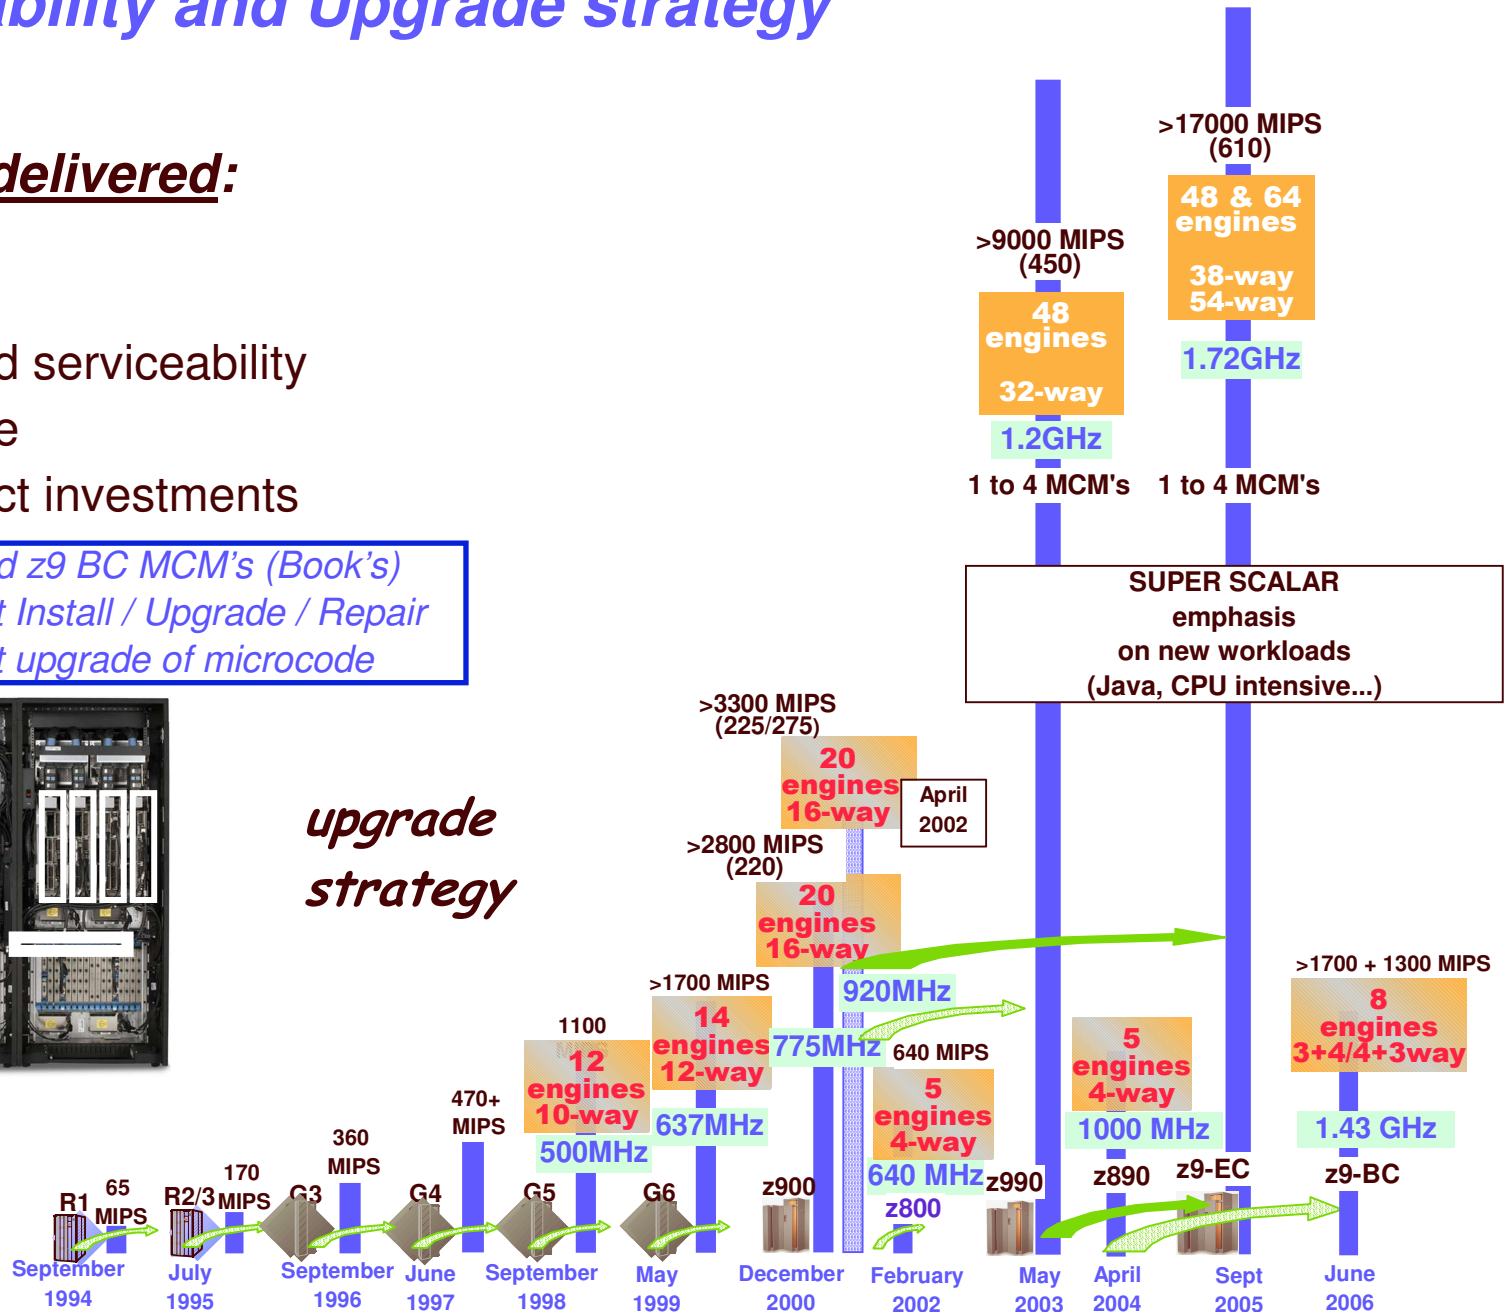
4 Advanced z9 BC MCM's (Book's)
 Concurrent Install / Upgrade / Repair
 Concurrent upgrade of microcode



System z9 EC MCM

95x95mm
 16 chipsites, 217 capacitors
3.9 billion transistors
 104 layer of glass ceramic
 476 meters of wire
more than 4000 MIPS
more than 400000++ TPC-C

upgrade strategy



Restriction of Hazardous Substances (RoHS) - Overview

What is RoHS

- European Union Restriction of Hazardous Substance Directive
- Bans placing new electrical and electronic equipment (EEE) containing more than agreed levels of lead, cadmium, mercury, hexavalent chromium, polybrominated biphenyl (PBB) and polybrominated diphenyl ether (PBDE) flame retardants on the EU market beginning July 1, 2006
- Most IT hardware is included within the scope: PCs, servers, storage, printers, options
- Products and their components must comply
- China plans to adopt this legislation or very similar legislation, date is not firm at this point. Other countries/states are looking at it.
- Impacts all products/equipment put on the market (sold) beginning July 1, 2006 -- not just the launch of a new product or product line
- Spares/FRU's for non-compliant equipment put (sold) on the market prior to 7/1/2006 are exempt from the legislation

System z9/zSeries Allowable Ships Matrix based on RoHS Legislation

MT	Box Ships	System MES	Hybrid	Loose piece feature MESs	
z900	No	No	No	Yes *	Non compliant, Worldwide withdrawal target 2Q2006
z800	No	No	n/a	Yes *	Non compliant, Worldwide withdrawal announced for 12/31/2005
z990	No	No	No	Yes *	Non compliant, RoHS country specific withdrawal target 2Q2006
z890	No	No	n/a	Yes *	Non compliant, RoHS country specific withdrawal target 2Q2006
2094	Yes	Yes	Yes #	Yes	Plan to be a fully compliant by 7/1/2006
2096	Yes	Yes	Yes #	Yes	Plan to be a compliant at GA
2074	No	No	n/a	Yes *	Non compliant, WW withdrawal tgt Feb 2006
9037	No	No	n/a	Yes *	Non compliant, RoHS country specific withdrawal target 2Q2006

* installed base for Feature MES must have been installed prior to 7/1/2006 to be allowed

MES kit being applied to the starting point base needs to be RoHS compliant on 7/1/2006 and SN can not change and must originate previously in that jurisdiction (such as the EU)



IBM System z9 Business Class (BC) systems

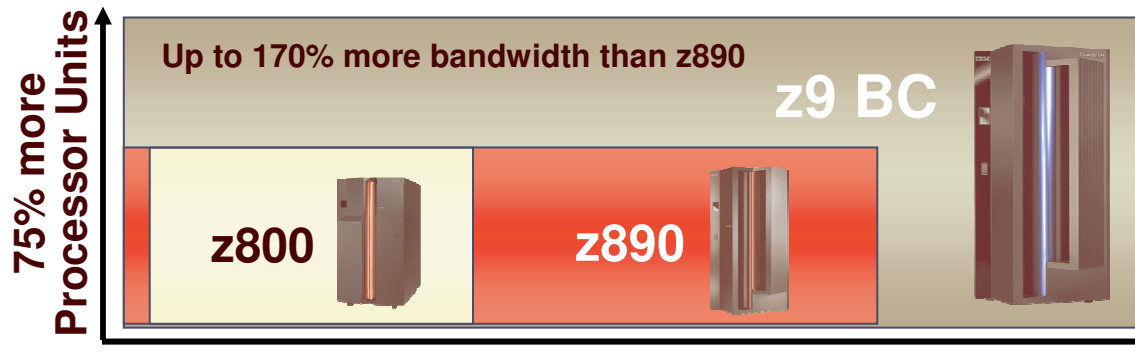
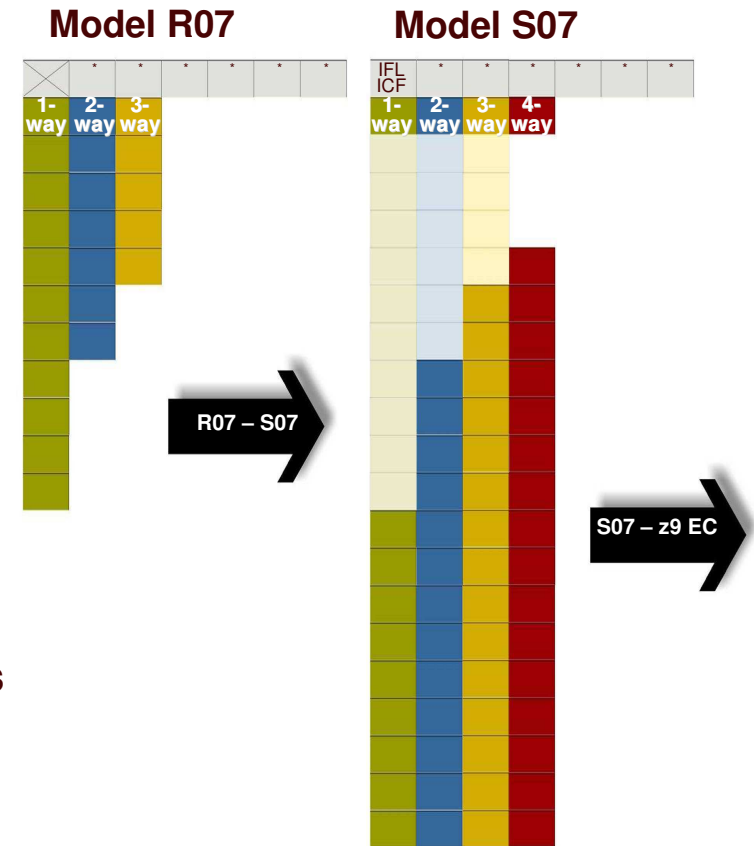


ENABLING BUSINESS.
A THROUGH Z.

z9 BC - The modern mainframe for the small to medium enterprises

Low entry point and more choices

- **Designed for flexibility in 2 new models - R07 and S07**
 - ▶ Entry model (R07) with 1 to 3 standard engines
 - ▶ 1 to 4 way (S07) model
 - ▶ **Granularity options – same entry size as z890 (26 MIPS)**
 - ▶ Up to 7 way with speciality engines
- **More speciality engines for more workloads (480 MIPS)**
 - ▶ Integrated Facility for Linux (IFL), Internal Coupling Facility (ICF)
 - ▶ System z™ Application Assist Processor (zAAP)
 - ▶ **System z Integrated Information Processor (zIIP)**
- **On demand upgrade capability**
 - ▶ Any-to-Any upgradeability including upgrade to z9 EC
 - ▶ On/Off Capacity on Demand (On/Off CoD) functions available
 - ▶ Sub-capacity CBU options
- **Double Memory (64GB) compared to z890**
- **Enhanced networking and connectivity options**
- **Built with System z9's Cryptographic & Encryption functions**
 - ▶ ATM/POS Remote Key Load
- **EWLC and Tiered EWLC Software Pricing Structure**
- **Operating system support - similar to z9 EC**
 - ▶ z/OS.e continues to be supported



+38% Engine Size

0.7 ns Cycle Time

+ 37% Capacity

Based on System z9 technology

- Availability enhancements
- More FICON channels
- MIDAW (FICON performance)
- 10 Gb Ethernet
- Multiple Subchannel Sets (MSS)

Flexibility - improved granularity and scalability

A choice that is just right

* = Specialty Engines

Model R07

	*	*	*	*	*	*
	1-way	2-way	3-way			
	A01	A02	A03			
	B01	B02	B03			
	C01	C02	C03			
	D01	D02	D03			
	E01	E02				
	F01	F02				
	G01					
	H01					
	I01					
	J01					

+ up to 6 FULL-size
Specialty Engines
(480 MIPS at "1/10 MIPS price)



Model S07

IFL/ICF	*	*	*	*	*	*
	1-way	2-way	3-way	4-way		
						K04
			L03			L04
			M03			M04
		N02	N03			N04
		O02	O03			O04
		P02	P03			P04
		Q02	Q03			Q04
R01	R02	R03				R04
S01	S02	S03				S04
T01	T02	T03				T04
U01	U02	U03				U04
V01	V02	V03				V04
W01	W02	W03				W04
X01	X02	X03				X04
Y01	Y02	Y03				Y04
Z01	Z02	Z03				Z04

+ up to 7 FULL-size
Specialty Engines
(480 MIPS at "1/10 MIPS price)



z9 BC Model S07

- ▶ Low entry point (26 MIPS)
- ▶ Full size engines and Speciality Engines = 480 MIPS
- ▶ Granularity for cost effective growth
- ▶ System z9 I/O packaging on a smaller scale
- ▶ More specialty engines compared to z890
- ▶ Any to any capacity upgradeability within the Model R07 and an upgrade path to the S07

z9 BC Model S07

- ▶ Granularity designed for flexibility and growth
- ▶ IFL/ICF only model
- ▶ Any to any capacity upgradeability within the Model S07 and upgradeable to the z9 EC
- ▶ More specialty engines including Linux only and ICF only servers

z9 BC – Delivering increased capacity and performance

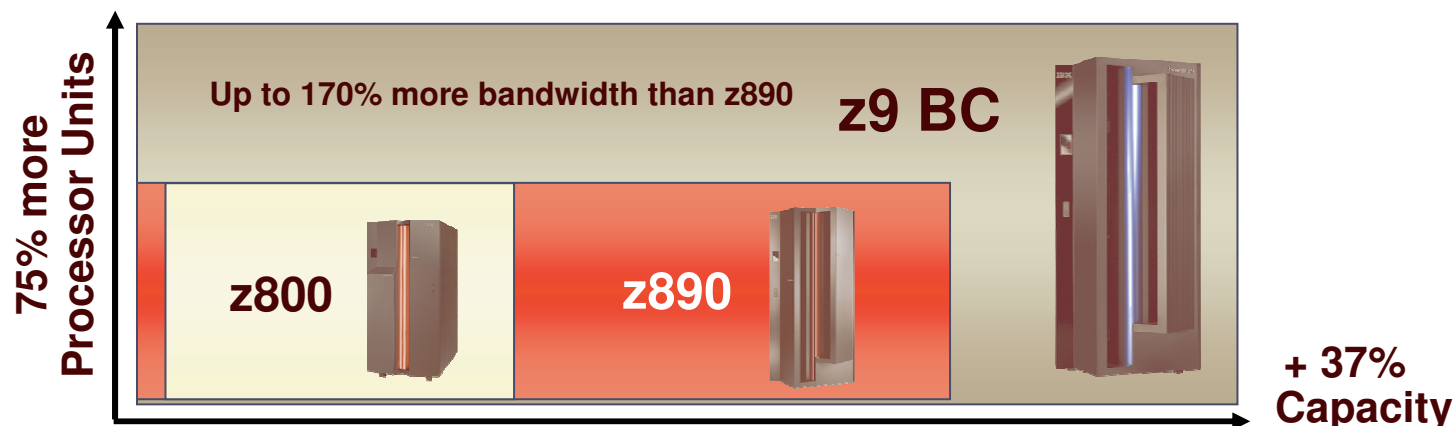
Flexibility for growth

Greater granularity and scalability

- Two models with one machine type (2096)
 - ▶ Entry model with 1 to 3-way standard engines
 - ▶ 1 to 4-way standard engines
 - ▶ Up to a 7-way with specialty engines
- 73 capacity settings for a 2.6 times increase in flexibility over IBM eServer™ zSeries® 890 (z890)
- Delivers over 36% more capacity with the same low entry point as the z890
- Up to 38% hardware performance improvement for Linux® (IFLs), Java™ (zAAPs) and coupling (ICFs)
- New zIIP for data serving workloads
- Sub-capacity CBU options
- Double the memory – up to 64 GB per server

Improved I/O Performance

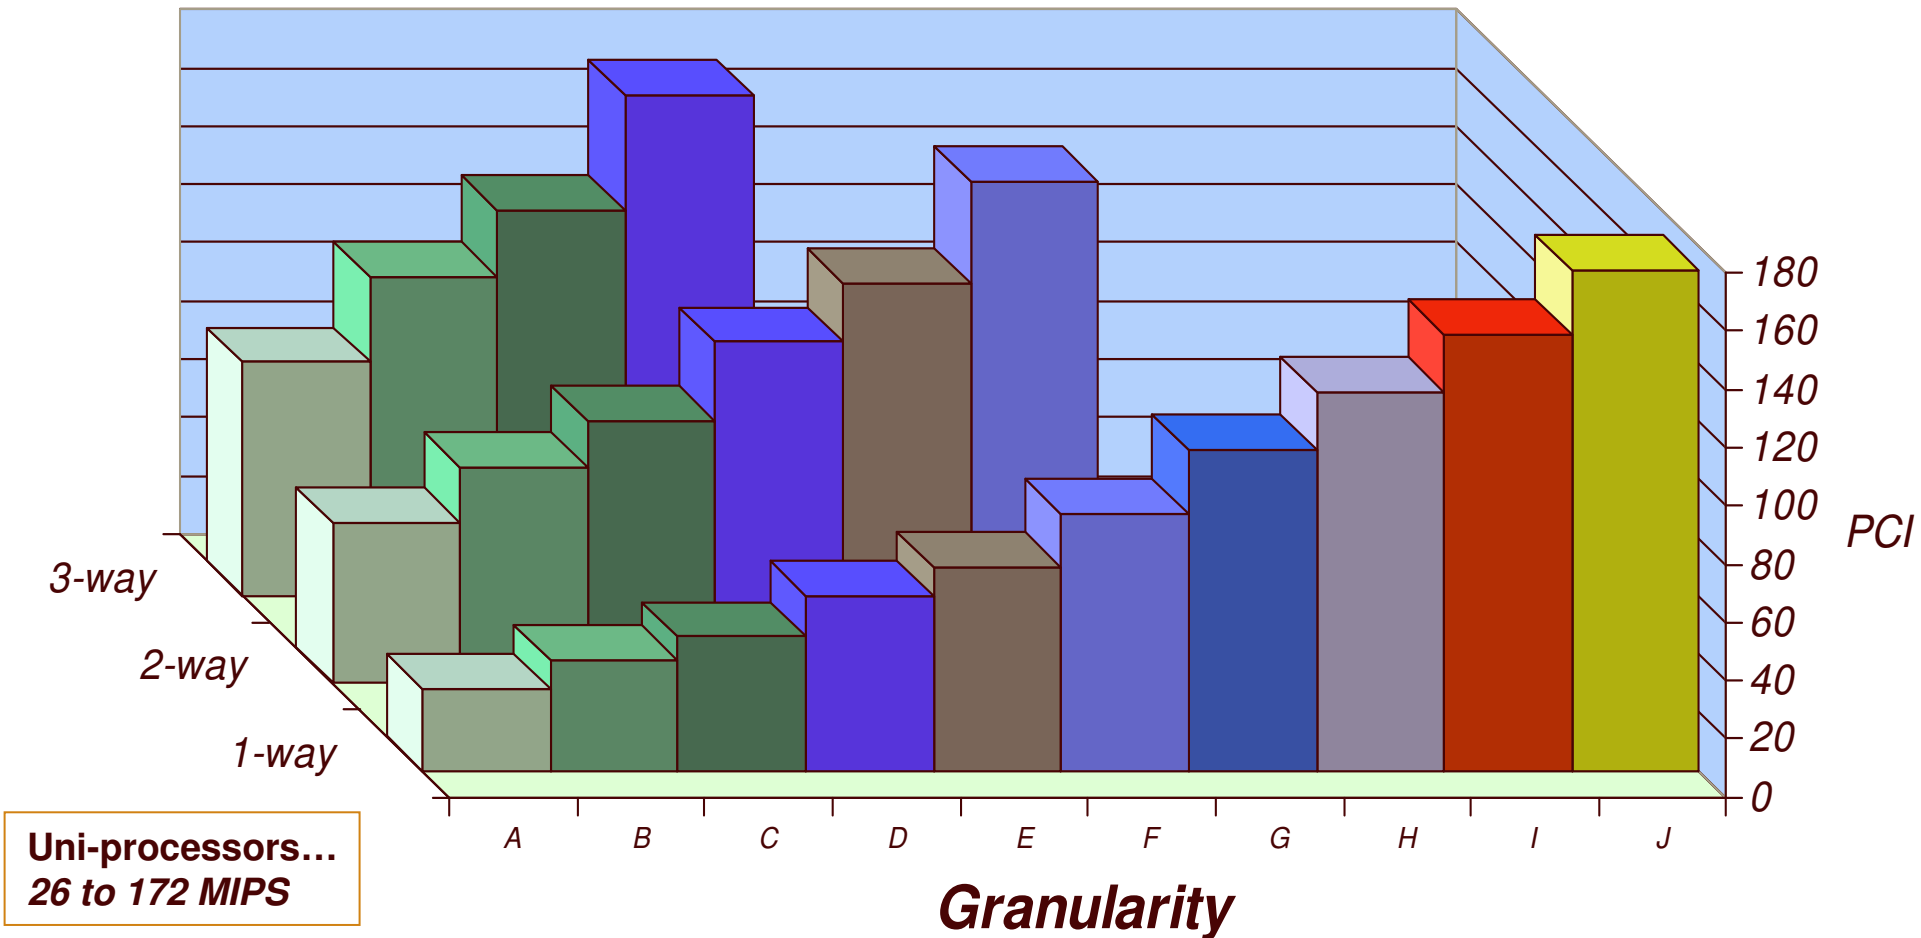
- 40% more FICON® channels – up to 112
- Up to 170% more bandwidth than z890
- Can improve FICON performance with Modified Indirect Data Address Word (MIDAW) facility
- Double the FICON concurrent I/O operations from 32 to 64 on FICON channel
- Multiple Subchannel Sets (MSS) for an increased number of logical volumes





Processor Granularity z9 BC R07 – 26 to 172 MIPS

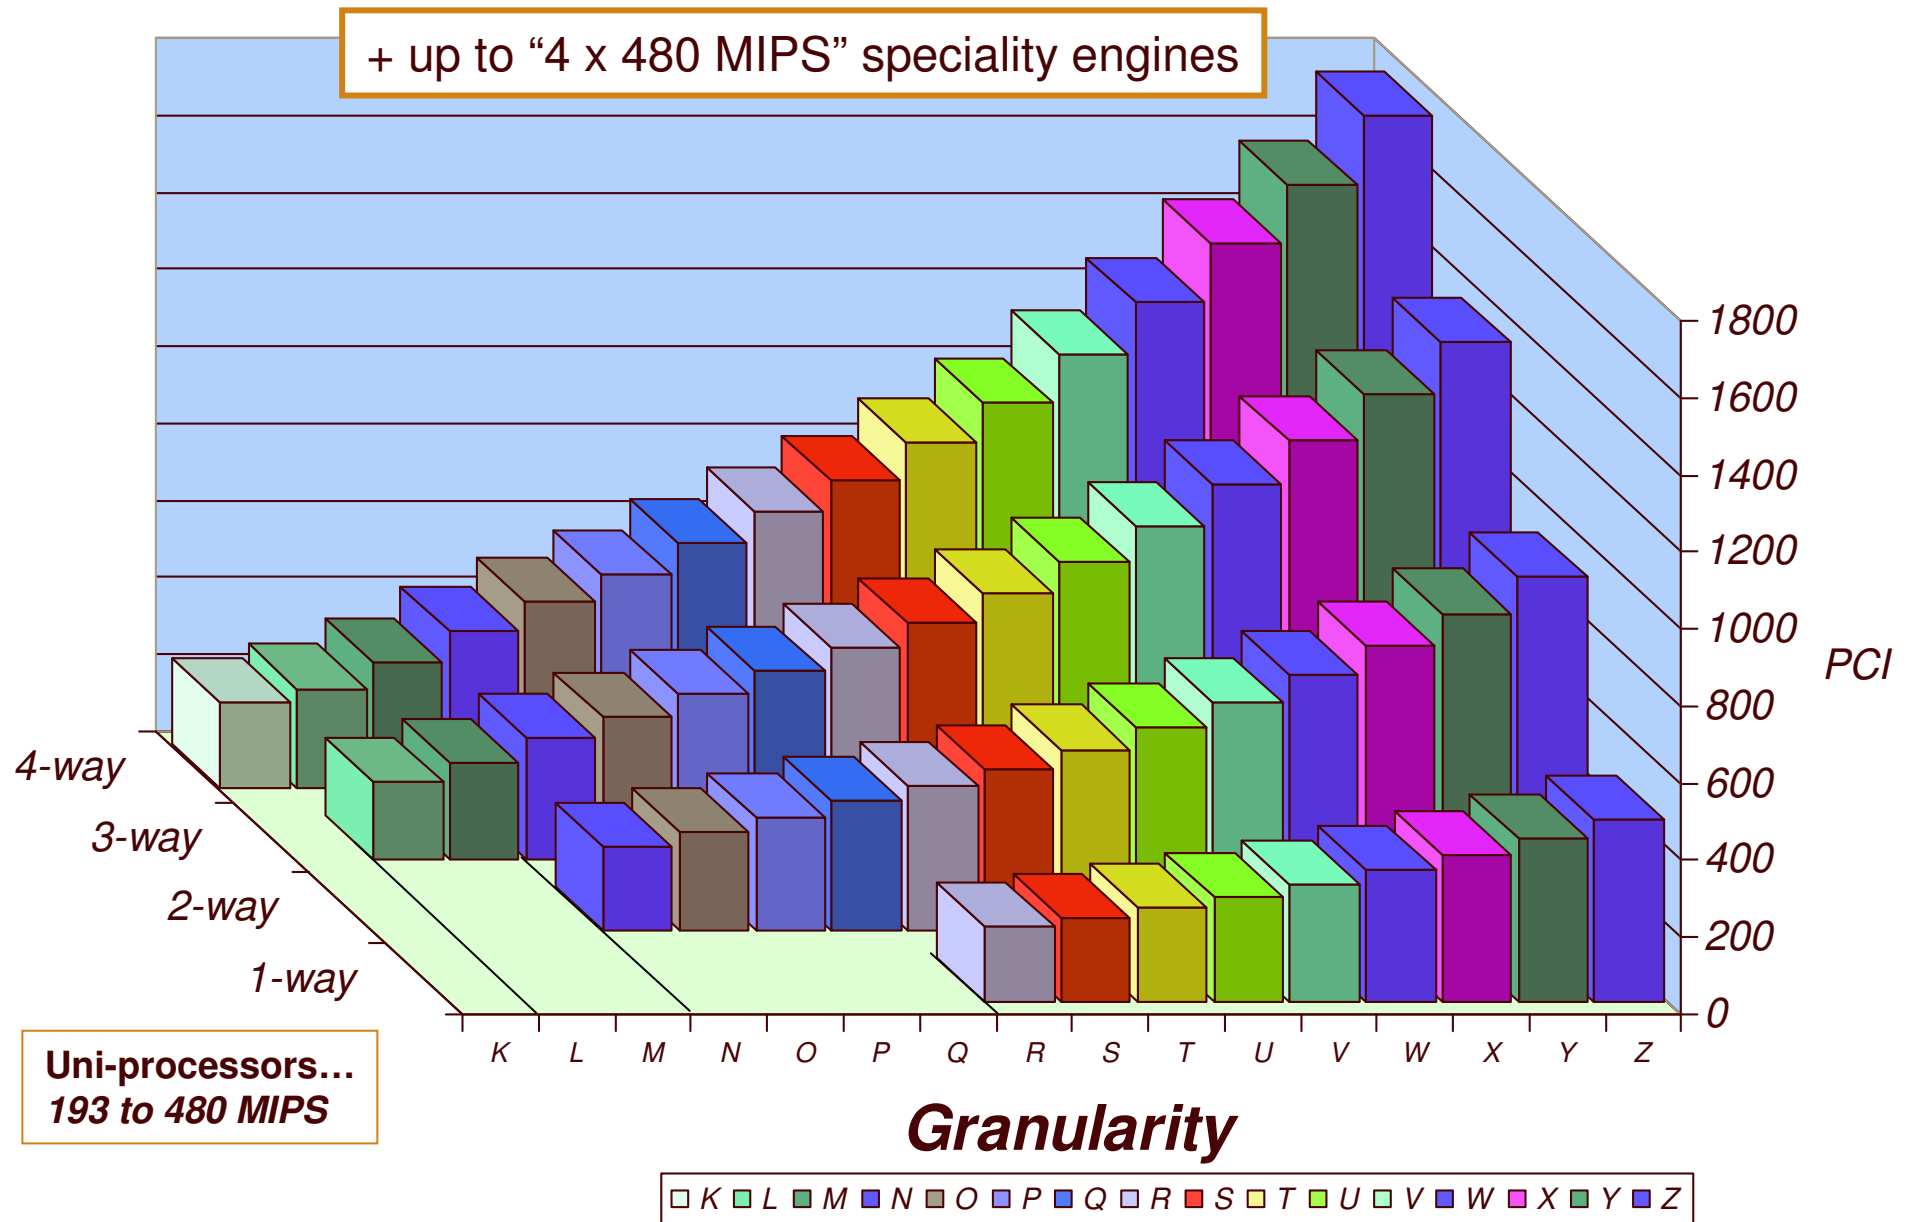
+ up to "4 x 480 MIPS" speciality engines



Uni-processors...
26 to 172 MIPS

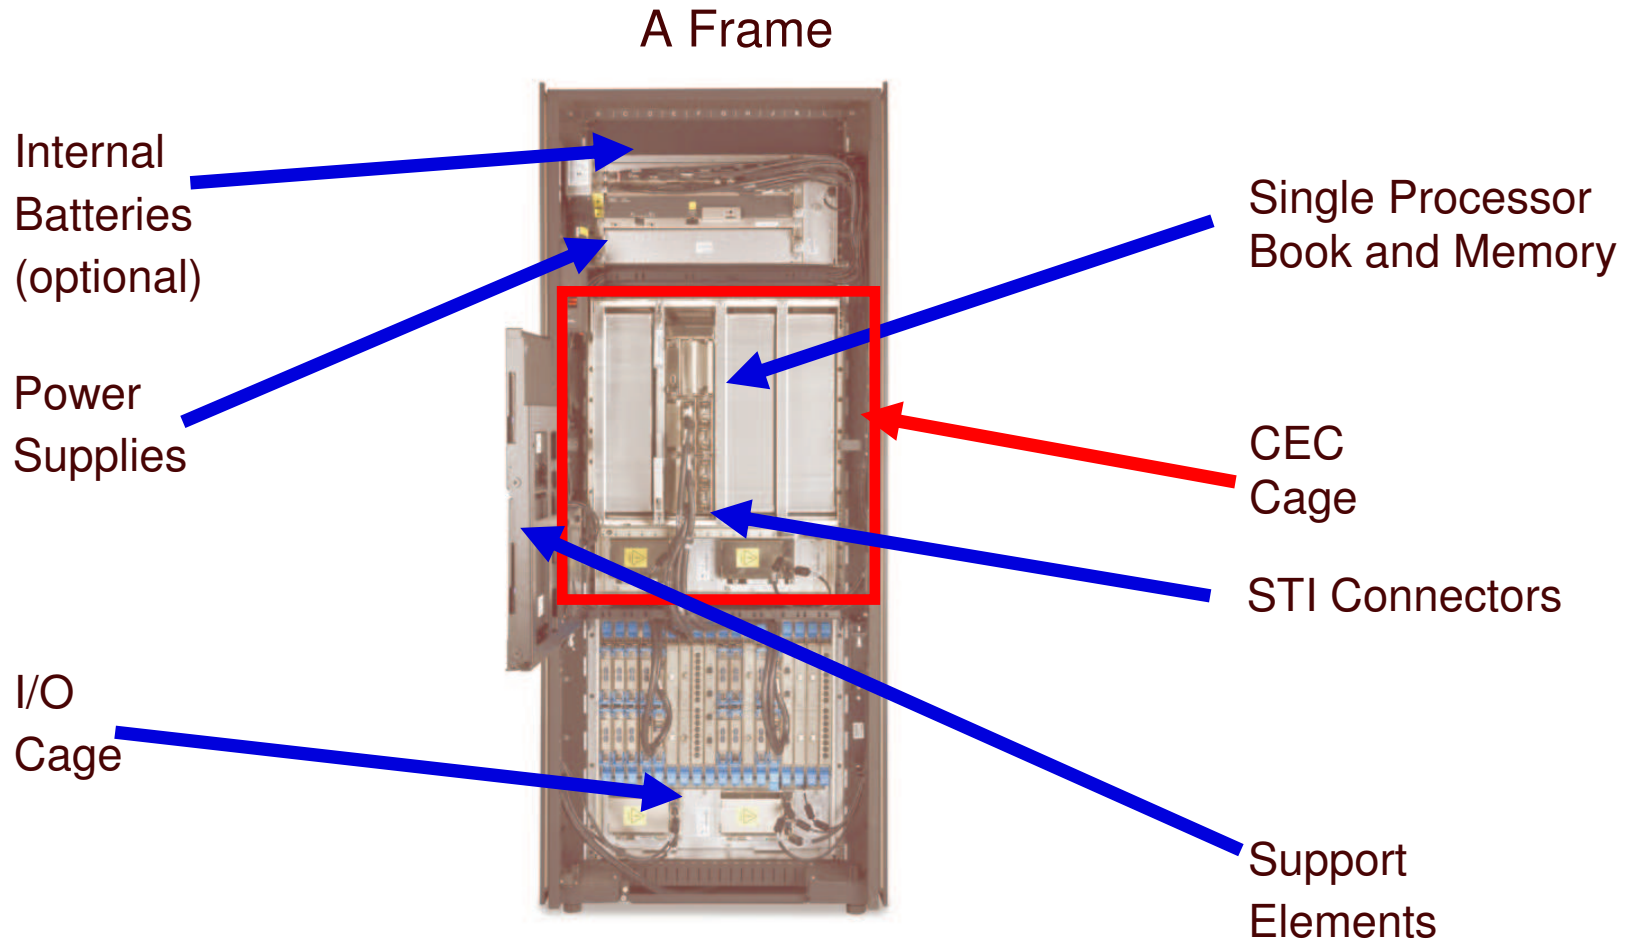
- A
- B
- C
- D
- E
- F
- G
- H
- I
- J

Processor Granularity z9 BC S07 – 193 to 1782 MIPS





z9 BC – Under the covers



Fiber Quick Connect Feature (optional)




Front View

z9 BC 8-way MCM

Advanced 95mm x 95mm MCM

- ▶ 102 Glass Ceramic layers
- ▶ 16 chip sites, 217 capacitors
- ▶ 0.545 km of internal wire

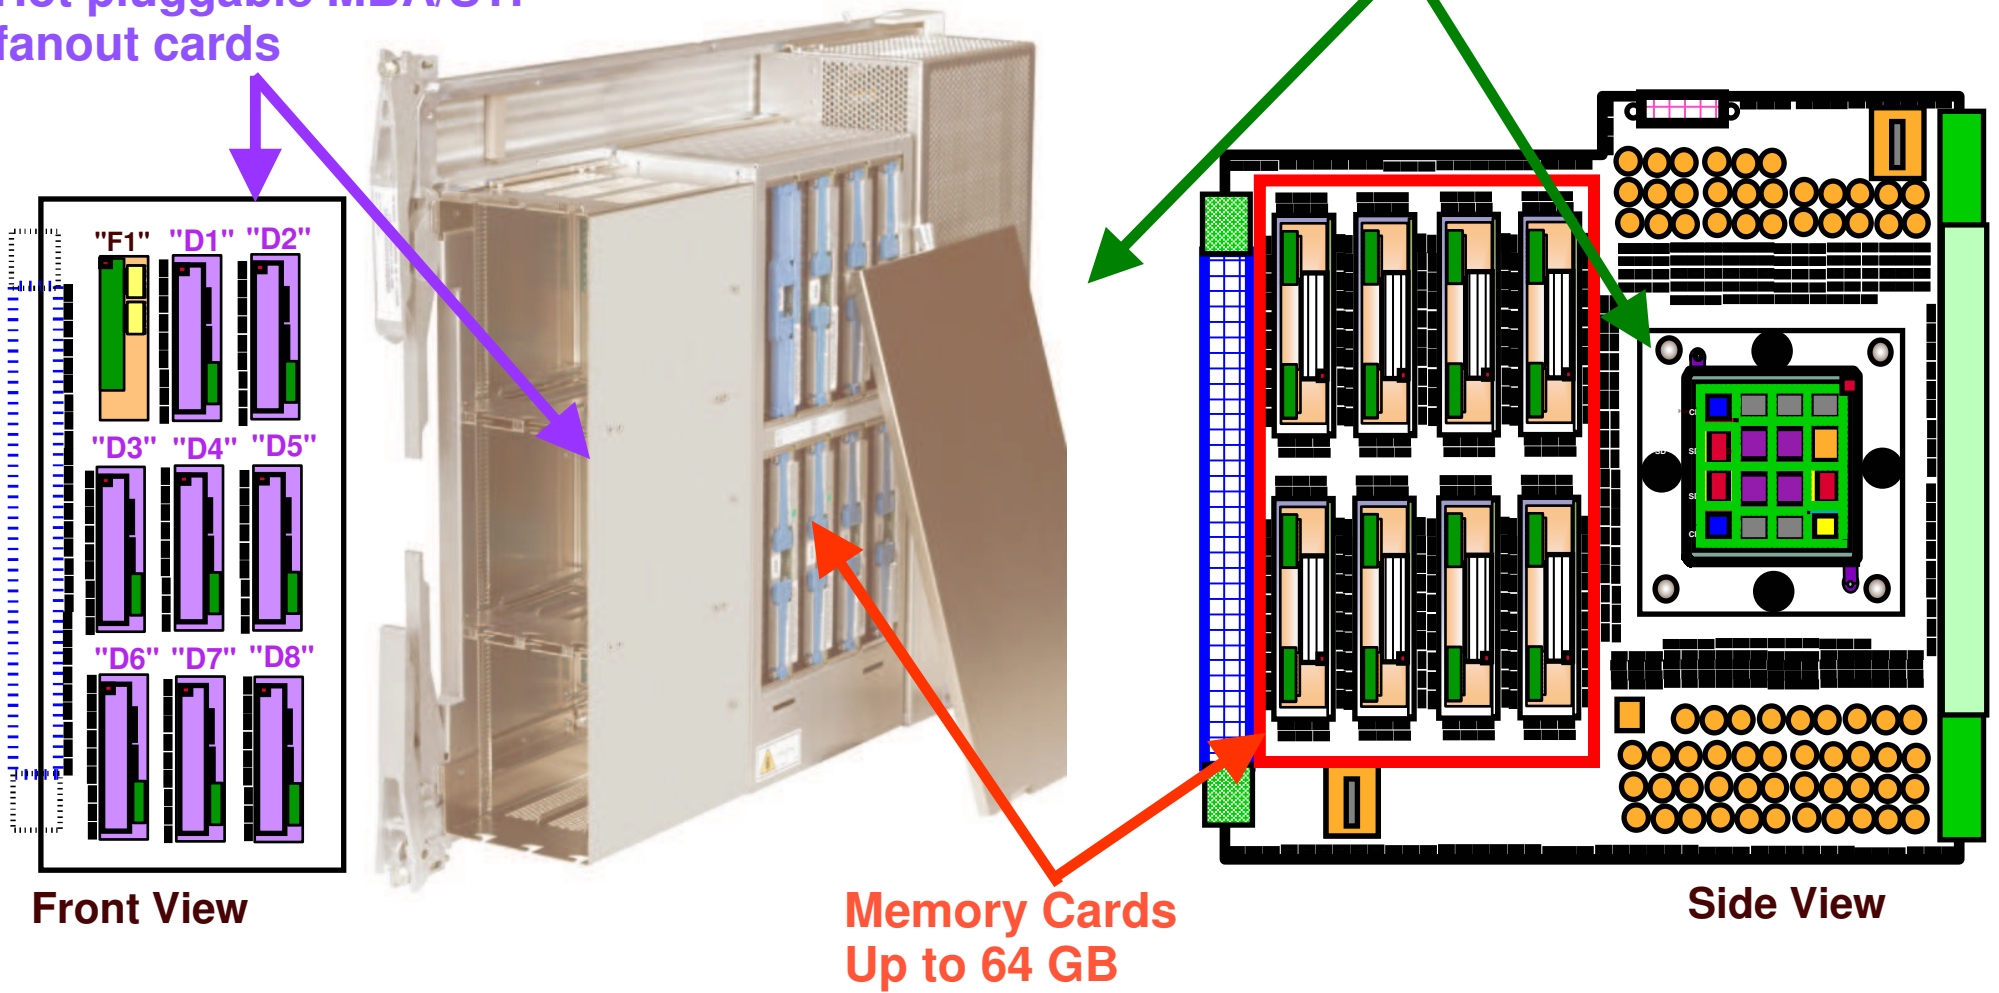


CMOS 10Ks0 chip Technology

- ▶ PU, SC, SD and MSC chips
- ▶ Copper interconnections, 10 copper layers
- ▶ 8 PU chips/MCM
 - 15.78 mm x 11.84 mm
 - 121 million transistors/chip
 - L1 cache/PU
 - 256 KB I-cache
 - 256 KB D-cache
 - **0.7 ns Cycle Time**
- ▶ 4 System Data (SD) cache chips/MCM
 - 15.66 mm x 15.40mm
 - L2 cache per Book
 - 660 million transistors/chip
 - 40 MB
- ▶ One Storage Control (SC) chip
 - 16.41mm x 16.41mm
 - 162 million transistors
 - L2 cache crosspoint switch
 - L2 access rings to/from other MCMs
- ▶ Two Memory Storage Control (MSC) chips
 - 14.31 mm x 14.31 mm
 - 24 million transistors/chip
 - Memory cards (L3) interface to L2
 - L2 access to/from MBAs (off MCM)
- ▶ One Clock (CLK) chip - CMOS 8S
 - Clock and ETR Receiver

z9 BC Processor Book Layout

Up to 8
Hot pluggable MBA/STI
fanout cards



- Note:**
1. Concept Illustration only - not to scale
 2. 4 or 8 pluggable Memory Cards
 3. Each MBA fanout card is hot-pluggable and has 2 STIs

z9 BC Memory Upgrade Options

From	To	To	To	To	To	To	To
8 GB	16 GB	24 GB	32 GB	40 GB	48 GB	56 GB	64 GB
16 GB	-	24 GB	32 GB	40 GB	48 GB	56 GB	64 GB
24 GB	-	-	32 GB	40 GB	48 GB	56 GB	64 GB
32 GB	-	-	-	40 GB	48 GB	56 GB	64 GB
40 GB	-	-	-	-	48 GB	56 GB	64 GB
48 GB	-	-	-	-	-	56 GB	64 GB
56 GB	-	-	-	-	-	-	64 GB
64 GB	-	-	-	-	-	-	-

Red - Disruptive upgrade

Green - Concurrent upgrade

Card Sizes = 2 GB, 4 GB and 8 GB

z9 BC – HSA considerations

- **HSA Estimator on Resource Link**

- ▶ Previously only available on a machine Support Element
 - Estimates relative to the machine family of the Support Element

- **HSA significantly larger than pre-z890 processors**

- **HSA larger than that experienced on z800 or z890 processors**

- ▶ Probable values in the 1.5 - 3 GB range depending on configuration
 - Dependent on number of LPARs, Multiple Channel Subsets, etc.
 - Some incremental growth for a 2086 to 2096 upgrade with no major changes to configuration

Standalone z9 BC Software Pricing

For Sub-Capacity Eligible Products *
 Entry Workload License Charges (EWLC)

For non Sub-Capacity Eligible Products
 EWLC Tiered Price Structure

EWLC Price Structure

Base	3 MSUs
Level 1	4 - 17 MSUs
Level 2	18 - 30 MSUs
Level 3	31 - 45 MSUs
Level 4	46 - 87 MSUs
Level 5	88 - 175 MSUs
Level 6	176 - 260 MSUs
Level 7	261+ MSUs

cumulative monthly pricing

EWLC Tiered Price Structure

Tier A	1-11 MSUs
Tier B	12-15 MSUs
Tier C	16-40 MSUs
Tier D	41 - 75 MSUs
Tier E	76 - 1500 MSUs
Tier F	1501+ MSUs



Flat monthly pricing.
 Select the tier based on the
MSU rating of your server

* Note: The z9 BC Model A01 is priced using zSeries Entry License Charges (zELC).

Protecting your investment in System z technology


- Full upgrades within the z9 (R07 to S07 to z9 EC)
- Any to any upgrade from the z890
- Upgrade from the z800 model 004*
- No charge MES upgrades on IFLs and zAAPs
- Capability of the System z9 servers to nondisruptively increase computing resources within the server
 - ▶ Can enable dynamic and flexible capacity growth for mainframe servers
 - ▶ Temporary capacity upgrade available through On/Off Capacity on Demand
 - ▶ Temporary, nondisruptive addition of CP processors/capacity, IFLs, ICFs, zAAPs or zIIPs
 - ▶ New options for changing On/Off CoD configurations
 - ▶ Sub-capacity CBU engines



IBM System z9 BC model comparison

Model R07

- **Processor Units (PUs)**
 - ▶ 7 PUs + 1 SAP
 - ▶ 1 - 3 CPs
 - ▶ 0 – 3 zAAPs or zIIPs
 - ▶ 0 – 6 IFLs or ICFs
 - ▶ 20 Capacity Settings
- **Memory**
 - ▶ 8 – 64GB
- **I/O**
 - ▶ 240 ESCON®
 - ▶ 64 FICON Express4
 - ▶ 32 OSA-Express2 (2-port)
 - ▶ Only 24 OSE-E on A01
 - ▶ 8 Crypto Express2
 - ▶ 16 STIs



Model S07

- **Processor Units (PUs)**
 - ▶ 7 PUs + 1 SAP
 - ▶ 0 - 4 CPs
 - ▶ 0 – 3 zAAPs or zIIPs
 - ▶ 0 – 7 IFLs or ICFs
 - ▶ 53 Capacity Settings
- **Memory**
 - ▶ 8 – 64GB
- **I/O**
 - ▶ 420 ESCON
 - ▶ 112 FICON Express4
 - ▶ 48 OSA-Express2 (2-port)
 - ▶ 16 Crypto Express2
 - ▶ 16 STIs

Both models have Sub-capacity CBU CPs and Specialty Engine CBU capabilities for more robust disaster recovery possibilities

IBM zSeries 890 to System z9 BC comparison

z890

- **Processor Units (PUs)**
 - ▶ 4 PUs + 1 SAP
 - ▶ 0 - 4 CPs
 - ▶ 0 – 2 zAAPs (*no zIIPs*)
 - ▶ 0 – 4 IFLs or ICFs
 - ▶ 28 Capacity Settings
- **Memory**
 - ▶ 8 – 32GB
- **I/O**
 - ▶ 420 ESCON
 - ▶ 80 FICON Express2
 - ▶ 40 OSA-Express2 (2-port)
 - ▶ 16 Crypto Express2
 - ▶ 8 STIs
- **Coupling Links (64 max)**
 - ▶ 32 IC
 - ▶ 48 ISC-3 (peer mode only)
 - ▶ 16 ICB-3
 - ▶ 8 ICB-4



Model S07

- **Processor Units (PUs)**
 - ▶ 7 PUs + 1 SAP
 - ▶ 0 - 4 CPs
 - ▶ 0 – 3 zAAPs or zIIPs
 - ▶ 0 – 7 IFLs or ICFs
 - ▶ 53 Capacity Settings
(73 capacity settings on z9 BC)
- **Memory**
 - ▶ 8 – 64GB
- **I/O**
 - ▶ 420 ESCON
 - ▶ 112 FICON Express4
 - ▶ 48 OSA-Express2 (2-port)
 - ▶ 16 Crypto Express2
 - ▶ 16 STIs
- **Coupling Links (64 max)**
 - ▶ 32 IC
 - ▶ 48 ISC-3
 - ▶ 16 ICB-3
 - ▶ 16 ICB-4

z9 BC models have Sub-capacity CBU CPs and Specialty Engine CBU capabilities for more robust disaster recovery possibilities

z9 BC Functions and Features

- Two hardware models with high levels of granularity available – 73 CIs
- Faster Uni Processor
- Up to 7 customer PUs
- Up to 64 GB memory
- Up to 30 LPARs
- Separate PU pool management
- PU Conversions for zAAPs and zIIPs
- CBU for IFL, ICF, zAAP and zIIP
- CBU and On/Off CoD Enhancements
- Redundant I/O interconnect
- Dynamic oscillator switchover
- 54 additional hardware Instructions
- Enhanced Driver Maintenance



Server Time Protocol

- Hot pluggable/ maintainable MBA/STI fanout cards
- Up to 16 2.7 GB STIs per system
- MIDAW facility
- Multiple Subchannel Sets per LCSS
- 63.75K Subchannels for Set-0
- Increased Number of FICON Express2 and 4 Features
- N_Port ID Virtualization
- IPv6 Support for HiperSockets
- OSA-Express2 1000BASE-T
- FICON and OSA Enhancements
- Crypto Enhancements
- Configurable Crypto Express2
- Availability in EU after implementation of RoHS

*This statement represents IBM's current intentions. IBM development plans are subject to change or withdrawal without further notice.

Note: Please refer to the latest PSP bucket for latest PTFs for new functions/features



IBM System z9 Enterprise Class (EC) systems



ENABLING BUSINESS.
A THROUGH Z.

IBM System z9 EC overview

- Machine Type
2094

- 5 Models
S08, S18, S28, S38, S54*



- Processor Units (PUs)
 - 48way & 64way
 - 580 to 17500+ MIPS
 - 0.58 ns cycle time
 - 12 PUs/book for S08/18/28/38
 - 16 PUs/book for S54
 - 2 SAPs per book, standard
 - 2 spare PUs per server
 - 1-38 or 1-54 PUs available...
 - CPs, IFLs, ICFs, zAAPs
 - optional SAPs
- Compared to z990
 - +35% more UNI-capacity
 - +95% more system capacity

- Memory
 - Minimum of 16 GB
 - up to 128 GB/book & 512 GB/system in 16 GB increments
- Bandwidth for I/O cage
 - up to 16 STIs per book
 - ▶ 2.7 GB/s for each I/O
 - ▶ 2.0 GB/s for ICBs
- Total system I/O bandwidth capability of 170+ GB/sec
- New generation of FICON/FCP
- Improved FICON performance with the MIDAW facility (Modified Indirect Data Address Word)
- Multiple subchannel sets
- 4 Channel Subsystems
- 60 Logical Partitions
- Concurrent ucode upgrade
- Concurrent book upgrade/repair
- Upgradeability
 - Upgrade paths from z900 & z990
 - Disruptive upgrade from zSeries and from other z9-BC models to model S54

Models	MCMs	Available PUs	Standard SAPs	Standard Spares	CP/ICF/IFL's and zAAP	Max Memory	Max** Channels
S08*	1	12	2	2	8	128 GB	960***
S18*	2	24	4	2	18	256 GB	1024
S28*	3	36	6	2	28	384 GB	1024
S38*	4	48	8	2	38	512 GB	1024
S54*	4	64	8	2	54	512 GB	1024

The name of the product is 'IBM System z9 109'

The SW Model Capacity Indicator field will read 700-754 depending on the number of active CPs on the machine.

z9 EC – Delivering increased capacity and performance

▪ Delivering new levels of scalability

- ▶ Built on modular book design – one to four books
- ▶ Five models with one machine type
 - 1 to 38-way high performance server (four models)
 - Up to 54-way enhanced model for high performance and maximum capacity
- ▶ The z9 EC full capacity uniprocessor is expected to deliver 35% more capacity than the z990 uniprocessor *
- ▶ The S54 offers 95% more server capacity than z990 **
- ▶ Two spare processor units per server
- ▶ Increased memory – up to 512 GB per server
- ▶ Multiple Subchannel Sets (MSS) for an increased number of logical volumes
- ▶ Up to 60 logical partitions (2X improvement)

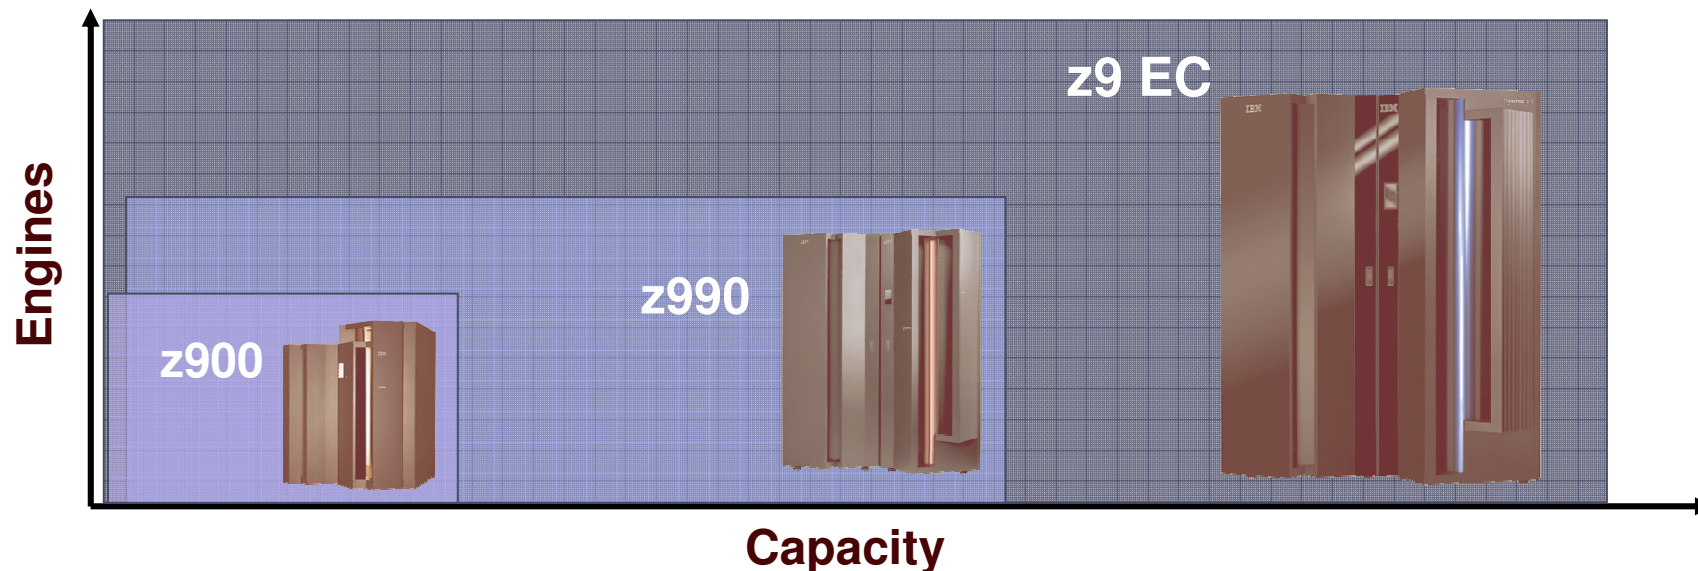
▪ Improved I/O Performance

- ▶ Up to 80%*** more bandwidth than the IBM eServer zSeries 990 (z990)
- ▶ Can improve FICON performance with Modified Indirect Data Address Word (MIDAW) facility
- ▶ New generation of FICON/FCP

* LSPR mixed workload average. z9 EC-701 Vs z990-301

** This is a comparison of the z9 EC 54-way and the z990 D32 and is based on LSPR mixed workload average.

*** When comparing a z990 Model A08 with a z9 EC Model S08

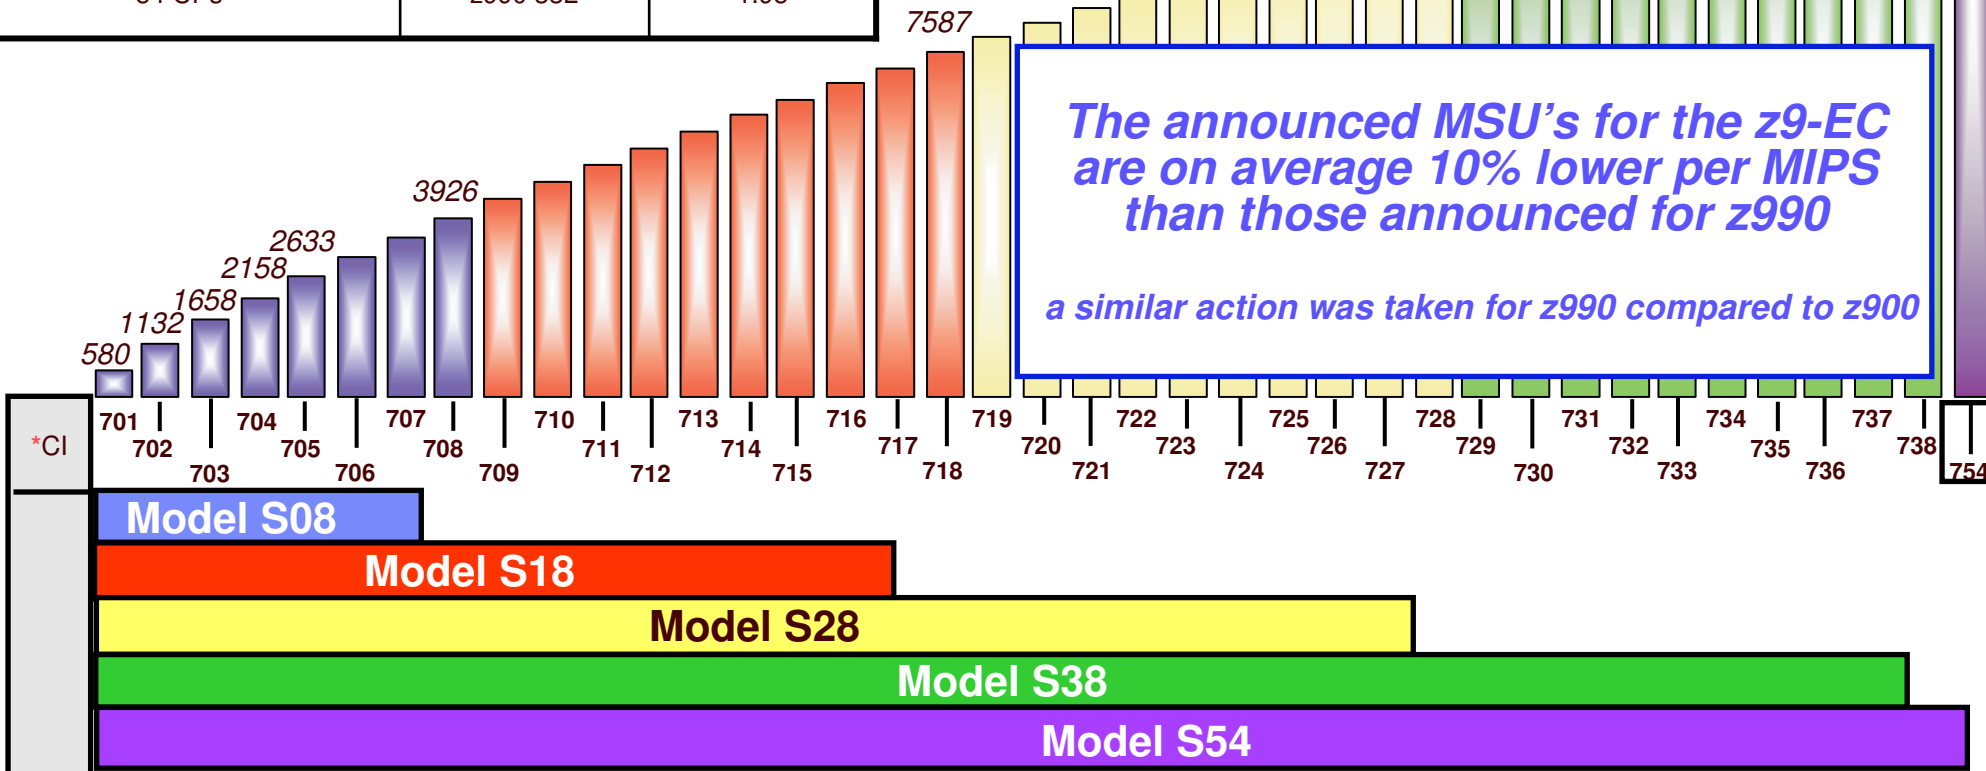




z9-EC Performance Comparison

Number of z9-109 CPs	Base	Ratio
1 CP	z990 301	1.35
8 CPs	z990 308	1.34
18 CPs	z990 318	1.34
28 CPs	z990 328	1.32
38 CPs	z990 332	1.51
54 CPs	z990 332	1.95

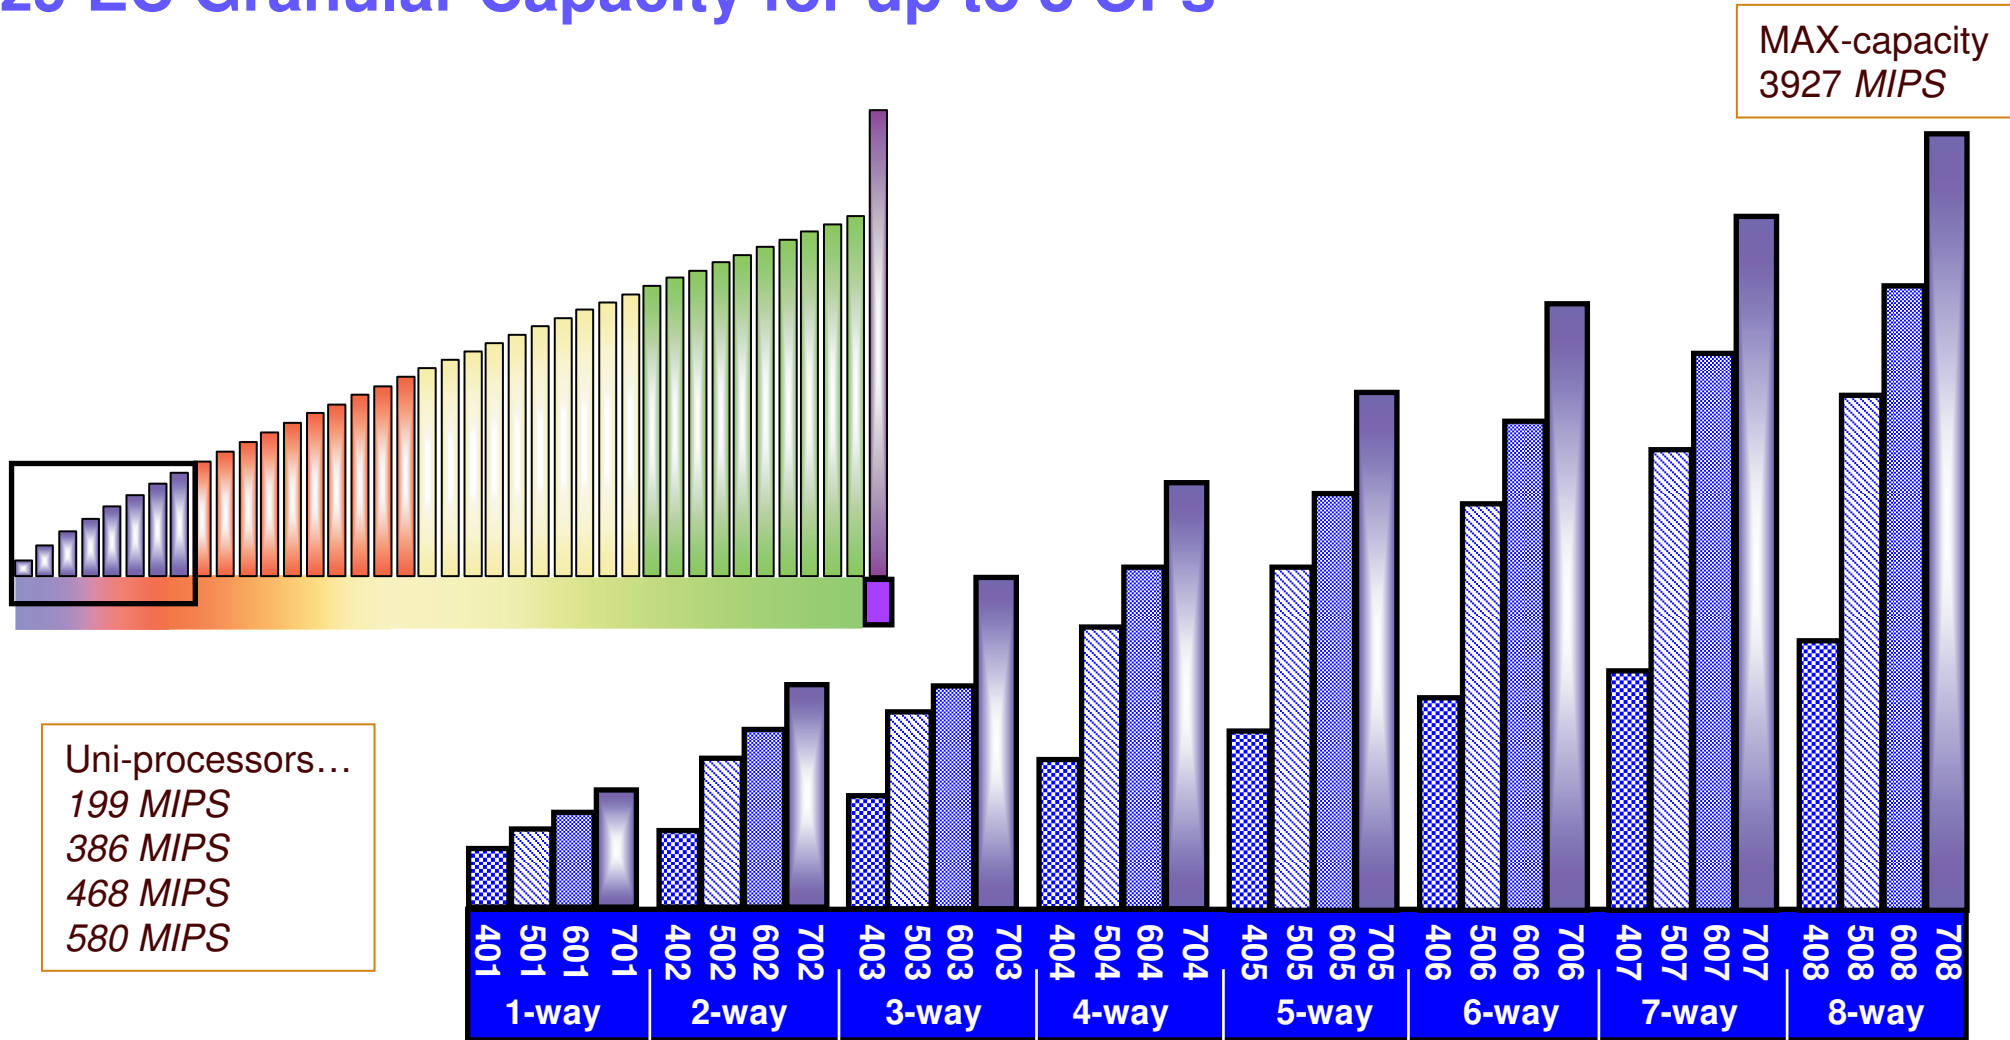
580 MIPS - 17801 MIPS
workload dependent variation
15800-21400



Note: For MSU values, refer to:
www-1.ibm.com/servers/eserver/zseries/library/swpriceinfo/
 For ITRs refer to: www-1.ibm.com/servers/eserver/zseries/lspr/zSerieszOS.html

- CI = Capacity Indicator and refers to number of installed CPs.
- CI Reported by STSI instruction.
- Model 700 does not have any CPs.

z9 EC Granular Capacity for up to 8 CPs

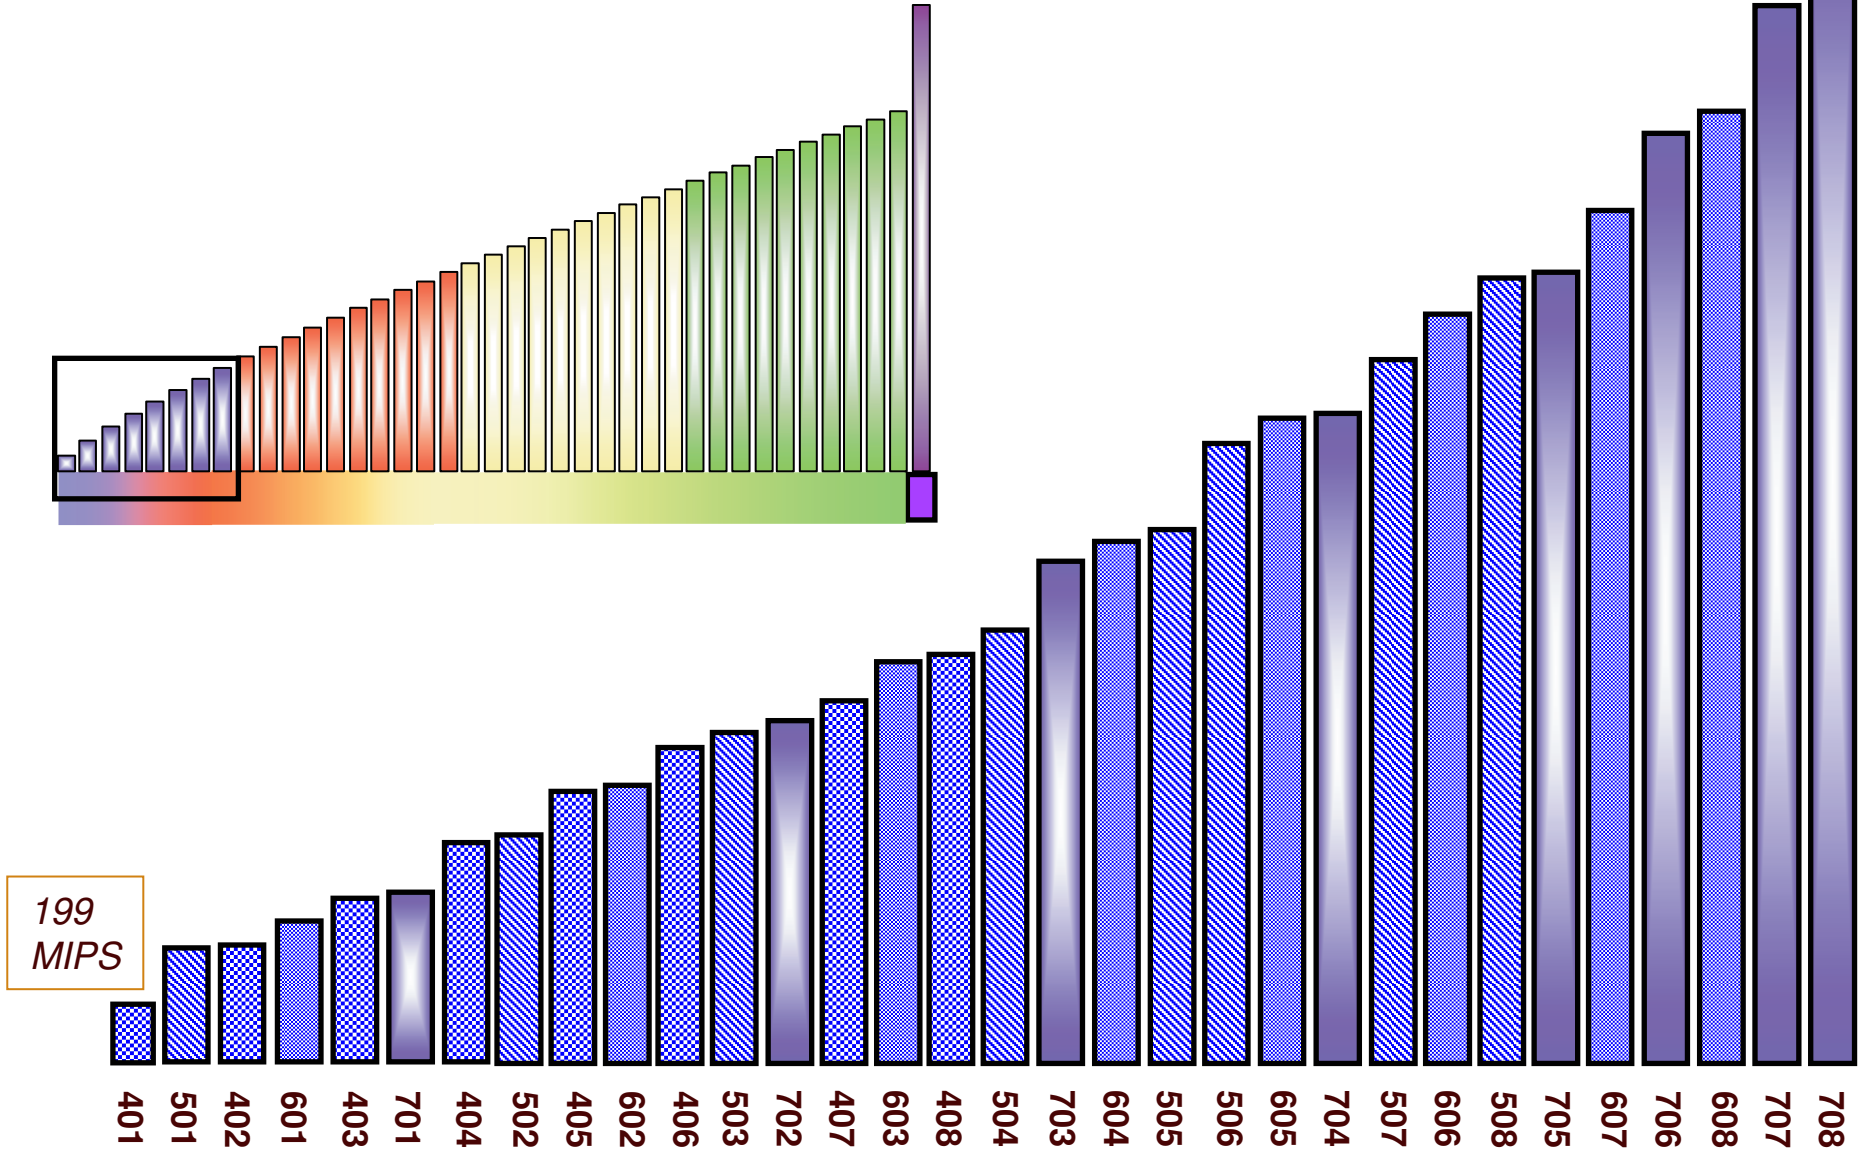


- The z9 EC now has 24 additional capacity settings at the low end of the processor (any-to-any upgrade)
- **Entry point is approximately 34% the capacity of the 701**
- All CPs must be the same capacity within one z9 EC
- **Combined zAAPs and/or zIIPs can not be more than 2x the number of CPs**
- Only 8 CPs can have granular capacity, other PUs must be CBU or characterized as specialty engines

z9 EC Granular Capacity for up to 8 CPUs

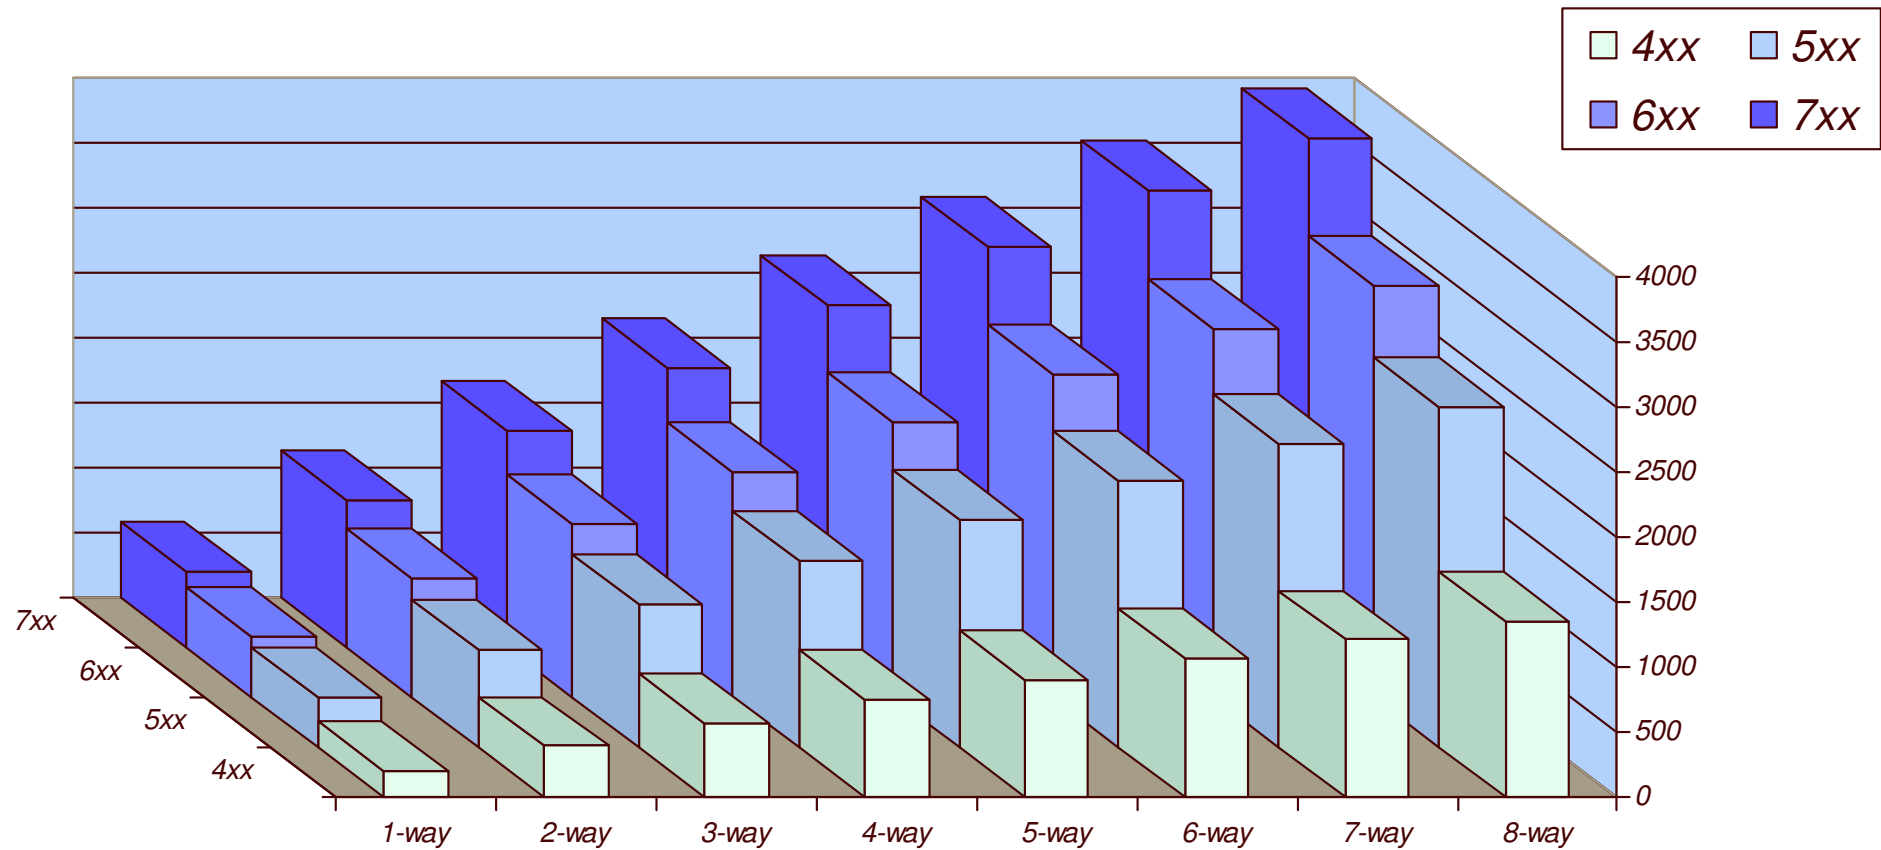
any to any upgrade – also to z9 EC models with more than 8 engines

3927
MIPS





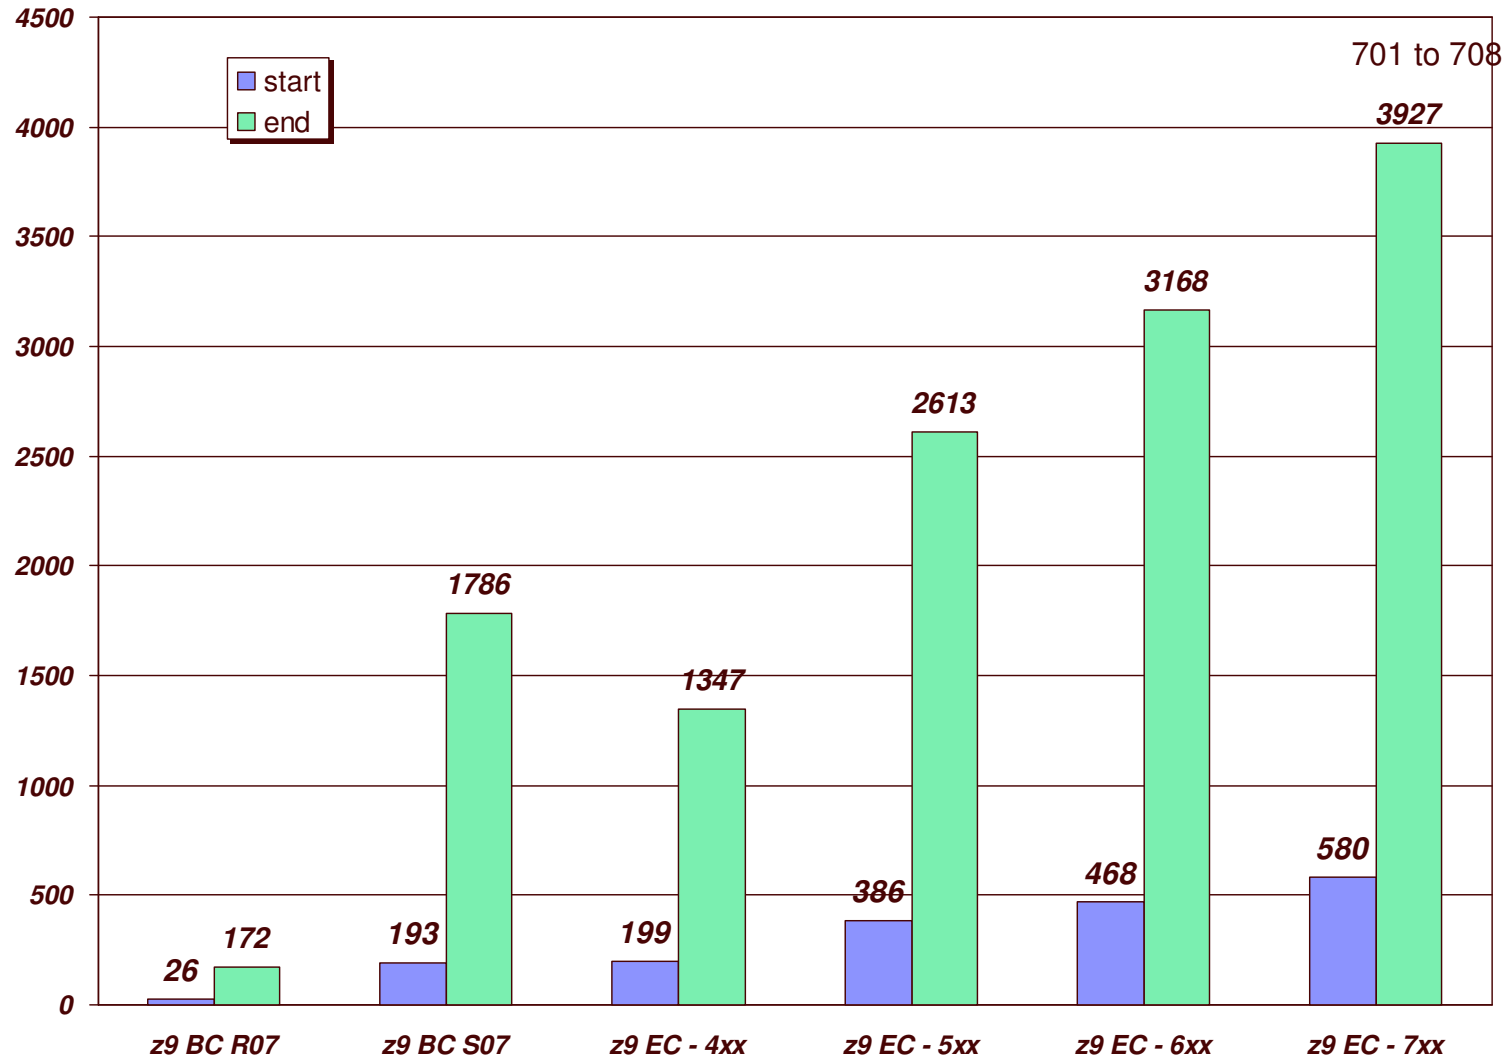
Processor Granularity z9 EC S08



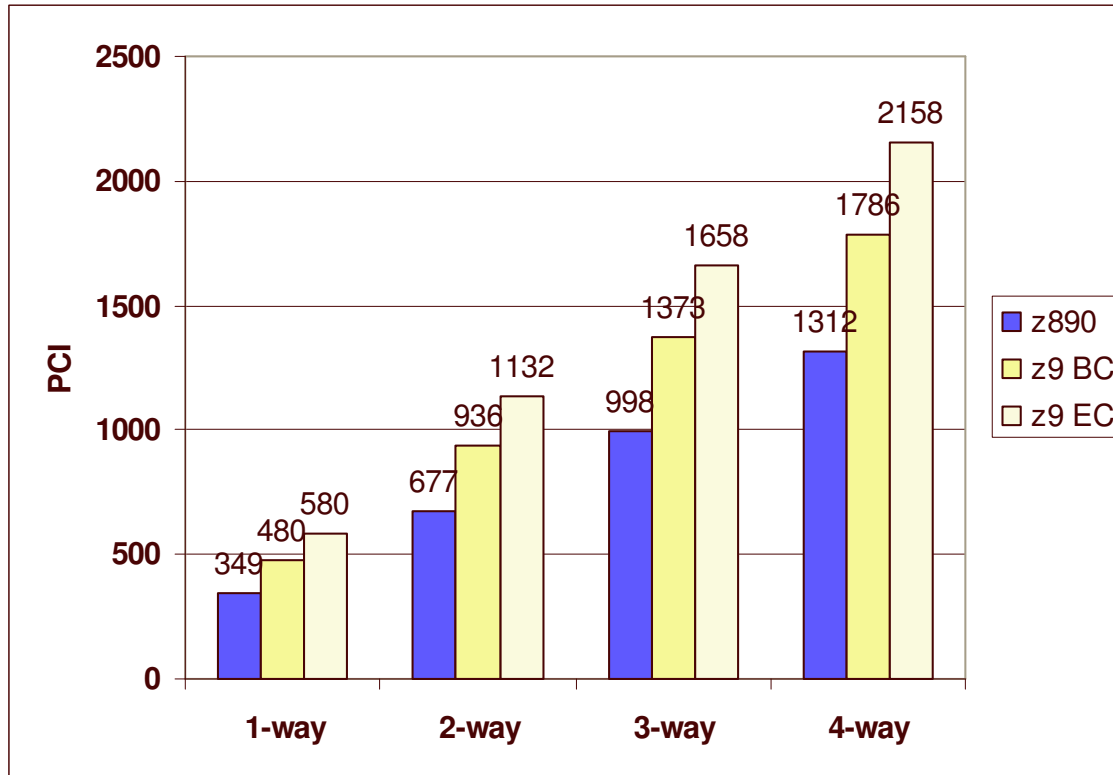
1-way	2-way	3-way	4-way	5-ways	6-ways	7-ways	8-ways	SW Model
199	388	569	740	903	1058	1206	1347	4xx
386	753	1104	1436	1752	2053	2339	2613	5xx
468	913	1338	1741	2124	2489	2836	3168	6xx
580	1131	1659	2158	2633	3086	3515	3927	7xx



Capacity Overlapping



Max Processor Capacity Index – 1 to 4way systems



Note: Estimates based on z/OS multi-image configuration.

z9 EC Model Structure

A flexible model structure that can be optimized for your business

- **One machine type – 2094 – five models – S08, S18, S28, S38, and S54**
- **Model number indicates PUs available for characterization**
 - ▶ Single serial number
 - ▶ PU characterization is identified by number of features ordered
- **2 System Assist Processors (SAPs) per book**
- **2 spares standard per server**
- **z9 EC software models**
 - ▶ 700, 401 to 408, 501 to 508, 601 to 608 and 701 to 754
 - ▶ nxx, where n = the capacity setting of the engine, and xx = the number of PU characterized as CPs in the CEC
 - ▶ Once xx exceeds 08, then all CP engines are full capacity

Models	MCMs	Available PUs	Max Available Sub-capacity CP PUs	Standard SAPs	Standard Spares	CP/IFL/ICF/zAAP/zIIP *****	Max Memory	Max Channels
S08*	1	12	8	2	2	8	128 GB	960 **
S18*	2	24	8	4	2	18	256 GB	1024 ***
S28*	3	36	8	6	2	28	384 GB	1024 ***
S38*	4	48	8	8	2	38	512 GB	1024 ***
S54*	4	64	8	8	2	54	512 GB	1024 ***

Notes:

* Must have a minimum of 1 CP, IFL or ICF

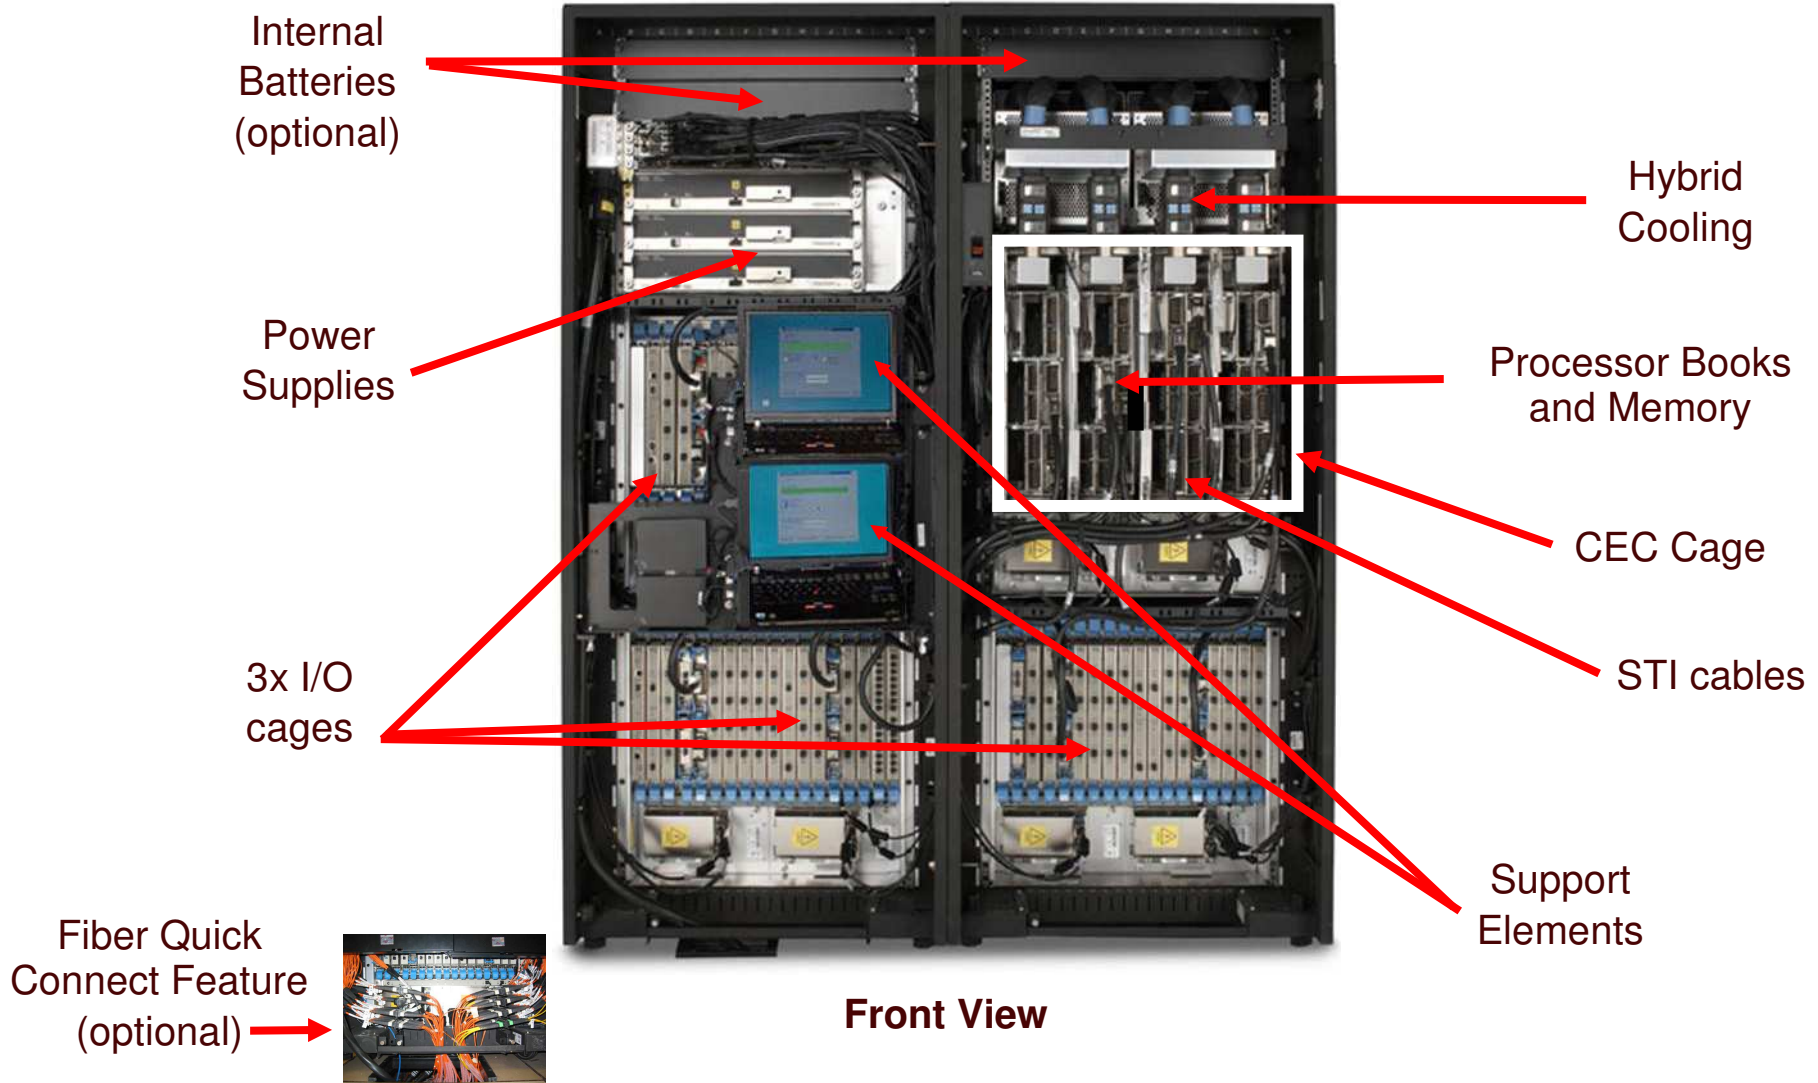
** There is a max of 64 ESCON features/960 active channels and a max of 64 FICON features/256 channels on Model S08.

*** The one for one relationship of zAAP or zIIP to CP still exists, but one CP can satisfy requirement for either or both specialty engines

**** Maximum of 16 ICFs




z9 EC – Under the covers (Model S38 or S54)



z9-EC 12-way MCM

- **Advanced 95mm x 95mm MCM**
 - ▶ 104 Glass Ceramic layers
 - ▶ 16 chip sites, 217 capacitors
 - ▶ 0.476 km of internal wire



CMOS 10K chip Technology

- **PU, SC, SD and MSC chips**
- **Copper interconnections, 10 copper layers**
- **8 PU chips/MCM**
 - ▶ 15.78 mm x 11.84 mm
 - ▶ 121 million transistors/chip
 - ▶ L1 cache/PU
 - 256 KB I-cache
 - 256 KB D-cache
 - *0.58 ns Cycle Time*
- **4 System Data (SD) cache chips/MCM**
 - ▶ 15.66 mm x 15.40mm
 - ▶ L2 cache per Book
 - 660 million transistors/chip
 - 40 MB
- **One Storage Control (SC) chip**
 - ▶ 16.41mm x 16.41mm
 - ▶ 162 million transistors
 - ▶ L2 cache crosspoint switch
 - ▶ L2 access rings to/from other MCMs
- **Two Memory Storage Control (MSC) chips**
 - ▶ 14.31 mm x 14.31 mm
 - ▶ 24 million transistors/chip
 - ▶ Memory cards (L3) interface to L2
 - ▶ L2 access to/from MBAs (off MCM)
- **One Clock (CLK) chip - CMOS 8S**
 - ▶ Clock and ETR Receiver

Notice: IBM z9 EC S54 is based on a 16way MCM module

Protecting Your Investment in System z Technology *Enhanced flexibility for upgradeability*

■ Full upgrades within the z9 EC


- ▶ Including any to any upgradeability in the 32 sub-capacity matrix

■ Any to any upgrade from the IBM eServer zSeries 990 (z990), IBM eServer zSeries 900 (z900) - except Model 100, or IBM System z9 BC Model S07

■ Capability of the System z9 servers to nondisruptively increase computing resources within the server such as processors, memory and I/O*

- ▶ Can enable dynamic and flexible capacity growth for mainframe servers
- ▶ Temporary capacity upgrade available through On/Off Capacity on Demand of CP processors, IFLs, ICFs, zAAPs or zIIPs
- ▶ New options for reconfiguring specialty engines if the business demands it
- ▶ New options for changing On/Off CoD configurations
- ▶ Sub-capacity CBU engines

* When properly configured. Also, upgrading to an S54 from other z9 EC models will require a planned outage



z9-109 (EC) New Functions / Features introduced in 2005 – expanded in 2006

Five HW models

Faster Uni Processor

50++ additional HW Instructions

Up to 54 CPs

Up to 512 GB memory

Up to 60 LPARs

CBU for IFL, ICF and zAAP

Separate PU pool mgnt.

Redundant I/O interconnect

Enhanced Driver maintenance

Enhanced Book availability

Dynamic oscillator switchover



Server Time Protocol

**Hot pluggable/ maintainable
MBA/STI fanout cards**

Up to 16 x 2.7 GB STIs per book

MIDAW facility

Multiple Subchannel Sets / LCSS

63.75K Subchannels for Set-0

**Increased Capacity
FICON Express2 & Express 4 Feat.**

N_Port ID Virtualization

IPv6 Support for HiperSockets

OSA-Express2 1000BASE-T

OSA-Express2 OSN (OSA for NCP)

**Enhanced CPACF with AES, PRNG
and SHA-256**

Configurable Crypto Express2

Now the System z9 BC does it too - but in a smaller package, etc...

New Enhanced microcode assist for Linux/zVM I/O & paging performance to be retrofitted back on to z890/z990

What's new in the LSPR for the z9-BC

- Business As Usual changes
 - ▶ uplevel z/OS V1R6, subsystem and compilers
- Improve customer representativeness of workload mix
 - ▶ drop CB-S (very short batch) workload
 - ▶ add CB-J (java-based) batch workload
 - ▶ new Mixed workload scales closer to LoIO-mix
- Improve customer representativeness of LPAR environment
 - ▶ 95% of z990s are configured with more than one z/OS image
 - ▶ provide two LSPR tables
 - single-image (SI) table
 - one z/OS image equal in size to Nway of model (z/OS V1R6 limit to 32way)
 - multi-image (MI) table
 - average complex LPAR configuration for each model based on customer profile
 - most representative for vast majority of customers
 - basis for single-number metrics

Average LPAR Configuration Profiles for the Multi-image Table

- Total number of z/OS images
 - ▶ 5 images at low-end models to 9 images at high-end
- Number of major images (>10% weight each)
 - ▶ 2 images across full range of models
- Size of images
 - ▶ low- to mid-range models have at least one image close to Nway of model
 - ▶ high-end models generally have largest image well below Nway of model
 - these models tend to be used for consolidation
- Logical to physical CP ratio
 - ▶ low-end near 5-1
 - ▶ most of the range 2-1
 - ▶ high-end near 1.5-1



IBM System z9 special processor considerations

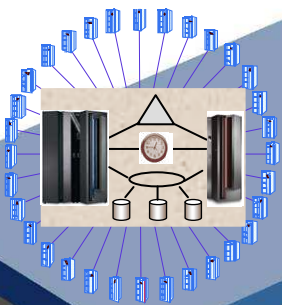


ENABLING BUSINESS.
A THROUGH Z.

More choice for your business

Evolution of specialty engines

No SW charges
\$125K per engine




Internal Coupling Facility (ICF) 1997

Centralized data sharing across mainframes




Integrated Facility for Linux (IFL) 2001

Support for new workloads and open standards



IBM System z Application Assist Processor (zAAP) 2004

Designed to help improve resource optimization for z/OS Java technology-based workloads



IBM System z9 Integrated Information Processor (IBM zIIP)

Designed to help improve resource optimization for eligible data workloads within the enterprise

*more functions in the future
Like TCP/IP Security, etc..*

IBM zSeries & System z9 - Specialty Engines

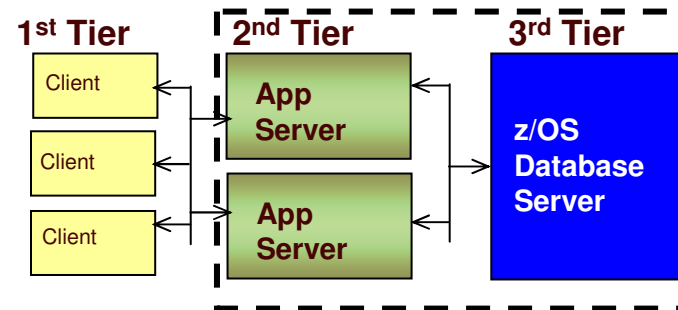
Delivering improved price/performance for new workloads

- **Application related integrated HW Specialty Engines**
 - Internal Coupling Facility (ICF) - *for clusters (sysplex)*
 - Integrated Facility for Linux (IFL) - *for Linux*
 - System z Application Assist Application (zAAP) - *for Java work*
 - System z Integrated Information Processor (zIIP) - *for some DB2 work*
- **NO SW charges related to Speciality Engines**
- **Specialty engines provide price / performance improvements**
 - Price independent of CAPACITY - increased benefit as engines grow
 - Low cost ...125K\$ per engine....independent of engine size
 - Reduced price options available on z9 Business Class systems
(95K\$ per engine for z9 S07 models, - even lower prices available on z9 R07)
- **Capacity BackUp (CBU) extended to Speciality Engines**
- **Speciality engines are managed “out of” individual PU pools**


Distributed versus a Centralized paradigm

Simplify and improve TCO by integration – z/OS

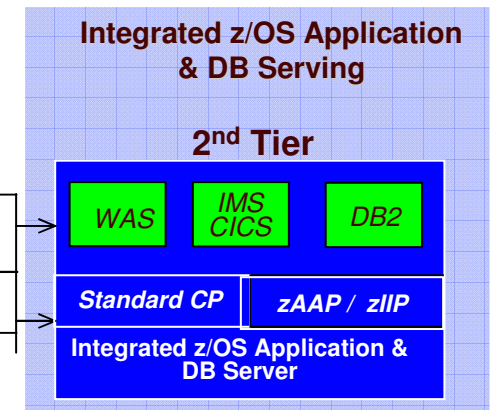
zAAP/zIIP processors



reduced HW cost
reduced SW cost for standard work
no SW charge for work on zAAP/zIIP engines



***provides the economic justification
 for
 the technical “correct”
 and
 integrated application solution***



**Potential opportunity
 DW server in future !**

*Use of zLinux is an option in between
 which provides for
 a simplified infrastructure*

Helping customers integrate data across the enterprise

The new IBM System z9 Integrated Information Processor (IBM zIIP)



Announcing
Availability

- **z/OS manages and directs work between the general purpose processor and the zIIP**
 - ▶ No changes anticipated to DB2 UDB for z/OS V8 applications
 - ▶ Number of zIIPs per z9 EC not to exceed number of general purpose processors
 - ▶ Price for each zIIP on the z9 EC is \$125K *
 - ▶ Price for each zIIP on the z9 BC is \$95K for S07, may be even lower for R07
 - ▶ No IBM SW charges on the zIIP - consistent with the zAAP specialty engine

- **DB2 UDB for z/OS V8 will be first IBM exploiter of the zIIP with:**
 - ▶ System z9 with z/OS 1.6 or later
 - ▶ *Other exploiters to be expected in the future*

- **Portions of the following DB2 UDB for z/OS V8 workloads may benefit from zIIP**:**
 - ▶ ERP, CRM, Business Intelligence and other enterprise applications – via DRDA over a TCP/IP connection
 - ▶ Data warehousing applications** – requests that utilize star schema parallel queries – plus parallel queries in general
 - ▶ DB2 UDB for z/OS V8 utilities** – select internal DB2 utility functions used to maintain index maintenance structures

* Prices may vary outside the US

** The zIIP is designed so that a program can work with z/OS to have eligible portions of its enclave Service Request Block (SRB) work directed to the zIIP. The above types of DB2 V8 work are those executing in enclave SRBs, of which portions can be sent to the zIIP.

How does the zIIP work ?

The zIIP is designed so that a program can work with z/OS to have all or a portion of its enclave Service Request Block (SRB) work directed to the zIIP.

The types of DB2 V8 work listed below are those executing in enclave SRBs, portions of which can be sent to the zIIP.

Example 1 = Distributed SQL requests (DRDA)

Queries that access DB2 for z/OS V8 via DRDA over a TCP/IP connection are dispatched within z/OS in enclave SRBs. z/OS directs a portion of this work to the zIIP.

Example 2 = Complex parallel query (BI)

Complex star schema parallel queries will now use enclave SRBs. z/OS directs a portion of this work to the zIIP.

Has been extended to long running queries in general

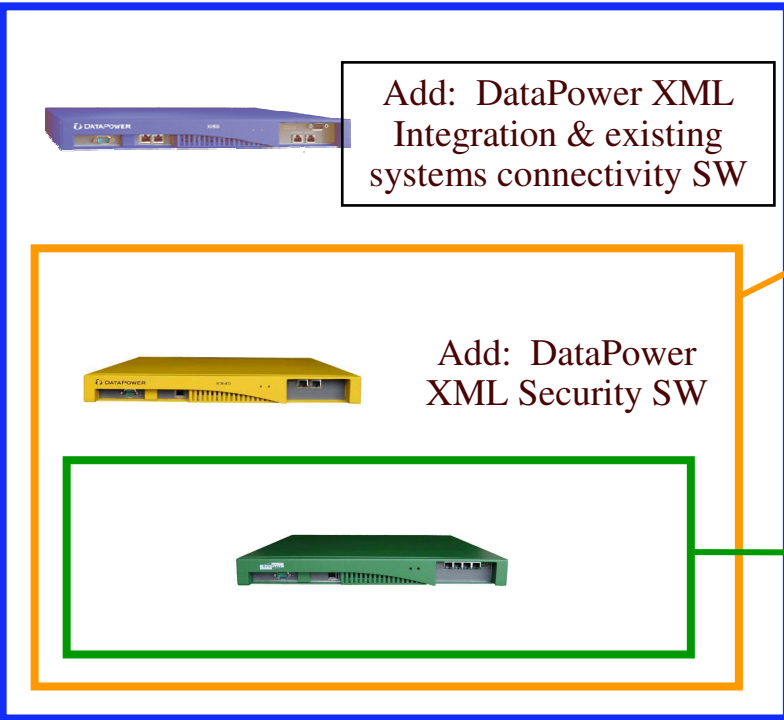
Example 3 = DB2 utilities for index maintenance

DB2 utilities LOAD, REORG, and REBUILD will now use enclave SRBs for the portion of the processing that is related to index maintenance. z/OS directs a portion of this work to the zIIP.

System z9 PR/SM Hypervisor PU Pool Rules

- **Logical PUs dispatched from supporting pool only**
 - ▶ Logical CPs from CP pool only, for example
- **Pool “width”**
 - ▶ Width equals the number of PUs in the pool
 - ▶ Limits LPAR maximum number of shared logical PUs brought ON
- **PUs placed in pools by**
 - ▶ Activate (POR)
 - ▶ Concurrent Upgrade – IBM System z9 On/Off Capacity on Demand (CoD), Capacity BackUp (CBU), Customer Initiated Upgrade (CIU), Capacity Upgrade on Demand (CUoD) MES
 - ▶ Dedicated LPAR deactivation
 - ▶ Dedicated LPAR configure logical PU OFF
- **PUs removed from pools by**
 - ▶ Concurrent Downgrade – On/Off CoD, CBU, PU Conversion MES
 - ▶ Dedicated LPAR activation (“width” permitting)
 - ▶ Dedicated LPAR configure logical PU ON (“width” permitting)
- **PUs can be CPs, IFLs, ICFs, zAAPs or zIIPs**

DataPower Products



■ XI50 Integration Appliance

- ▶ Expands support to non-XML solutions
- ▶ Advanced architecture
- ▶ Integrated message-level security



■ XS40 XML Security Gateway

- ▶ Security, agility and performance
- ▶ Device can off-load application security software
- ▶ Performs XML Web services security functions (parse, filter, validate schema, encrypt/decrypt, signatures, access control, and more)



■ XA35 XML Accelerator

- ▶ Offloads overtaxed servers by processing XML, XSD, XPath and XSLT at wire speed
- ▶ SW provides significant performance improvements over WebSphere solutions
- ▶ HW + SW provides enterprise-class performance



**XML / Security
HW/ucode support**

*DataPower Box owned by SWG today.....
External attachment.....may reduce XML cost by 5-20 times (16 at UBS)
Opportunity for further integration.....*



IBM System z9 Capacity on Demand



ENABLING BUSINESS.
A THROUGH Z.

System z9 Concurrent Upgrade – CUoD

standard functions

- **Non-disruptive addition of CPs (‘model capacity’ upgrade), ICFs, IFLs, zAAPs, zIIPs**
- **Some administrative improvements**
- **LIC enable additional memory increments**
- **Concurrent model upgrade (capacity setting) to add active PUs, PU capacity and memory**
- **Downgrades are non-disruptive**
- **Note: I/O feature adds and removes are also non-disruptive**
- ***CBU – CP’s, zAAPs, IFL’s & ICF’s***

Notes:

1. CUoD is built on a base of concurrent “hot-plug” maintenance
2. I/O feature adds and removes are also non-disruptive
3. Customer planning and operator action are required to take full advantage of CUoD. To avoid a planned outage, it may be necessary to predefine LPAR profiles with “reserved” resource specified. It may also be necessary to use z/OS or z/VM dynamic I/O capabilities. In some cases, disruption of certain LPARs is required following a concurrent hardware change.

System z9 Concurrent Upgrade – Customer Controlled

- **On/Off Capacity on Demand - Temporary upgrade**
 - ▶ Nondisruptive temporary addition of CPs, **CP capacity**, IFLs, ICFs, zAAPs and zIIPs
 - ▶ Upgrades requiring parts (e.g., for a z9-EC Model S08 to S18 upgrade) not supported
 - ▶ "Right to use" feature - Orderable as MES or with new build to initiate contract and administrative setup
 - ▶ Customer orders and installs upgrade via Resource Link and IBM RSF
 - ▶ Nondisruptive removal when capacity is no longer wanted
- **CIU – Customer Initiated Upgrade - Express - Permanent upgrade**
 - ▶ Customer capability to order and install permanent upgrade
 - ▶ Not included
 - Upgrades requiring parts (e.g., for a z9-EC Model S08 to S18 upgrade)
 - Channel upgrades by LIC enable of existing ports
 - ▶ CIU feature - MES ordered to initiate contract and administrative setup
 - ▶ Customer orders and installs upgrade via Resource Link and IBM RSF
- **CBU – Capacity BackUp**
 - ▶ Nondisruptive temporary addition of CPs, ICF, IFL, zAAPs in an emergency situation
 - ▶ CBU contract required to order CBU features and CBU LIC CC
 - ▶ Customer activates upgrade for test or temporary emergency
 - ▶ Nondisruptive downgrade after test or recovery completed

System z9 On/Off Capacity on Demand (On/Off CoD)

- **For customers who have a requirement for short-term additional capacity**

- **What Is It?**

- ▶ Temporary, nondisruptive addition of one or more CPs, IFLs, zAAPs, or ICFs
 - Memory and channels are not included
- ▶ Contract between IBM and customer
- ▶ Features required: CIU Enablement (#9898) and On/Off CoD Enablement (#9896)
- ▶ Add a quantity of temporary CPs, IFLs, and ICFs up to the quantity of permanently active plus dormant CPs, IFLs, and ICFs, respectively
- ▶ Add a quantity of temporary zAAPs up to the quantity of permanent zAAPs, with the limitation that the combined quantity of temporary zAAPs plus permanent zAAPs may not exceed the combined quantity of temporarily active CPs plus permanently active CPs plus dormant CPs on the same machine
- ▶ Billed for associated usage through one or more of the following features: On/Off CoD Active CP-Day (#9897), On/Off CoD Active IFL Day (#9888), On/Off Active ICF-Day (#9889), or On/Off CoD Active zAAP-Day (#9893)
 - HW charges calculated monthly and billed in arrears based on 24-hour usage period

- **Nondisruptive temporary upgrade**

- ▶ Execute On/Off via Resource Link
- ▶ Customer must take action to deactivate (same action as CBU undo)

- **Test**

- ▶ One no-charge test per contract.
- ▶ A maximum duration of 24 hours commencing with the download and activation of an On/Off CoD order.
- ▶ On/Off CoD tests that do exceed 24 hours in duration will be treated in their entirety as billable On/Off CoD upgrades.

System z9 Capacity Backup Upgrade

- **For customers who have a requirement for robust Disaster Recovery**
- **What Is It?**
 - ▶ Temporary, nondisruptive addition of one or more CPs, IFLs, zAAPs, or ICFs
 - Memory and channels are not included
 - ▶ Must plan ahead for memory and connectivity requirements
 - ▶ Contract between IBM and customer
 - ▶ Count of CBU Features is the number of active CPs, IFLs, zAAPs or ICFs to be added
 - FC: 7820 for CBU CP
 - FC: 7821 for CBU IFL
 - FC: 7822 for CBU ICF
 - FC: 7824 for CBU zAAP
 - ▶ Count of active PU features plus CBU features limited to available PUs
- **Nondisruptive temporary upgrade or test process**
 - ▶ Execute CBU from HMC
 - ▶ CBU features activate as CPs, IFLs, zAAPs, or ICFs
 - ▶ Configure additional logical CPs ON to active partition
 - Predefine as "Reserved" PU(s)
- **Nondisruptive downgrade process**
 - ▶ Required after recovery or test completed
 - ▶ Follow procedures to quiesce workload
 - ▶ Configure CBU PU(s) OFF or deactivate using Logical Partition
 - ▶ Execute downgrade from HMC
- **Sub-capacity CBU's supported**



System z Capacity on Demand Summary

Capacity Upgrade on Demand (CUoD) ^{***}	Capacity BackUp (CBU) ^{***}	Customer Initiated Upgrade (CIU) ^{***}	On/Off Capacity on Demand (On/Off CoD) ^{***}
Permanent capacity upgrade; a standard System z9 and zSeries feature that allows you to order extra capacity resources such as processors, memory*, and I/O	Reserve backup PU capacity (CP, ICF#, IFL#, zAAP# or zIIP#) for specified duration; original configuration must be restored after test or disaster recovery	Facility for ordering, configuring, pricing and installing capacity upgrades. It is a Web-based solution available through Resource Link	Temporary capacity upgrade (CP, ICF, IFL, zAAP, zIIP) of unlimited duration; orderable through CIU; customer activates and deactivates.
Available on LIC enabled System z9, z990, z890, z900, z800 and G5/G6	Available on System z9, z990, z890, z900, z800 and G5/G6	Available on LIC enabled System z9, z990, z890, z900 and z800	Available only on System z9, z990 and z890; orderable feature
Inherent capability of System z9 and zSeries servers; spare processors, memory and/or I/O slots must be available	A CBU contract must be in place prior to implementation and reserve PUs available for test or disaster recovery	A CIU contract must be in place prior to implementation	A CIU contract with special On/ Off CoD terms and conditions and right-to-use feature must be in place prior to implementation
Capacity upgrade Installed by customer or IBM Service representative	Capacity reserve installed by customer or IBM Service representative for predetermined period of use	CIU contract and registration required to use CIU application to order capacity	Feature ordered through IBM Sales; once enacted, customer orders temporary CP, ICF, IFL, zAAP or zIIP upgrade through CIU
Customer or IBM planning required	Customer or IBM planning required	Customer planning required	Customer planning required
Nondisruptive** capacity activation (may require deactivation or activation of LPAR partition)	Nondisruptive** capacity activation System z9, z990, z890, z900 and z800 and G5, G6	Ordering facility available with the System z9, z990, z890, z900 and z800	Nondisruptive temporary CP, ICF, IFL, zAAP or zIIP upgrade; customer deactivates; mutually exclusive with CBU enablement

* Memory cannot be upgraded on z800 with CUoD. Limited option for z890 and z9 BC
 ** CUoD and CBU may need IPL for z800/z890 "sub" model upgrades with older levels of OS

*** Additional terms and conditions apply
 # System z9 ONLY

Note: Upgrades are nondisruptive only where there is sufficient hardware resource available and provided pre-planning has been done

z9 EC Concurrent PU Conversions

- **Must order (characterize one PU as) a CP, an ICF or an IFL**
- **Concurrent model upgrade (book add) is supported – S08 to S18 to S28 to S38**
 - ▶ Upgrades to Model 54 are disruptive
- **Concurrent processor upgrade is supported if PUs are available**
 - ▶ Add CP, IFL, unassigned IFL, ICF, zAAP, zIIP or optional SAP
- **Concurrent processor type conversion is supported with limitations:**
 - ▶ Unassigned IFL must be activated as an IFL first before it can be converted to any other PU type

From/To->	CP	IFL	Unassigned IFL	ICF	zAAP	zIIP
CP	x	Yes	No	Yes	Yes	Yes
IFL	Yes	x	Yes	Yes	Yes	Yes
Unassigned IFL	No	Yes	x	No	No	No
ICF	Yes	Yes	No	x	Yes	Yes
zAAP	Yes	Yes	No	Yes	x	Yes
zIIP	Yes	Yes	No	Yes	Yes	x

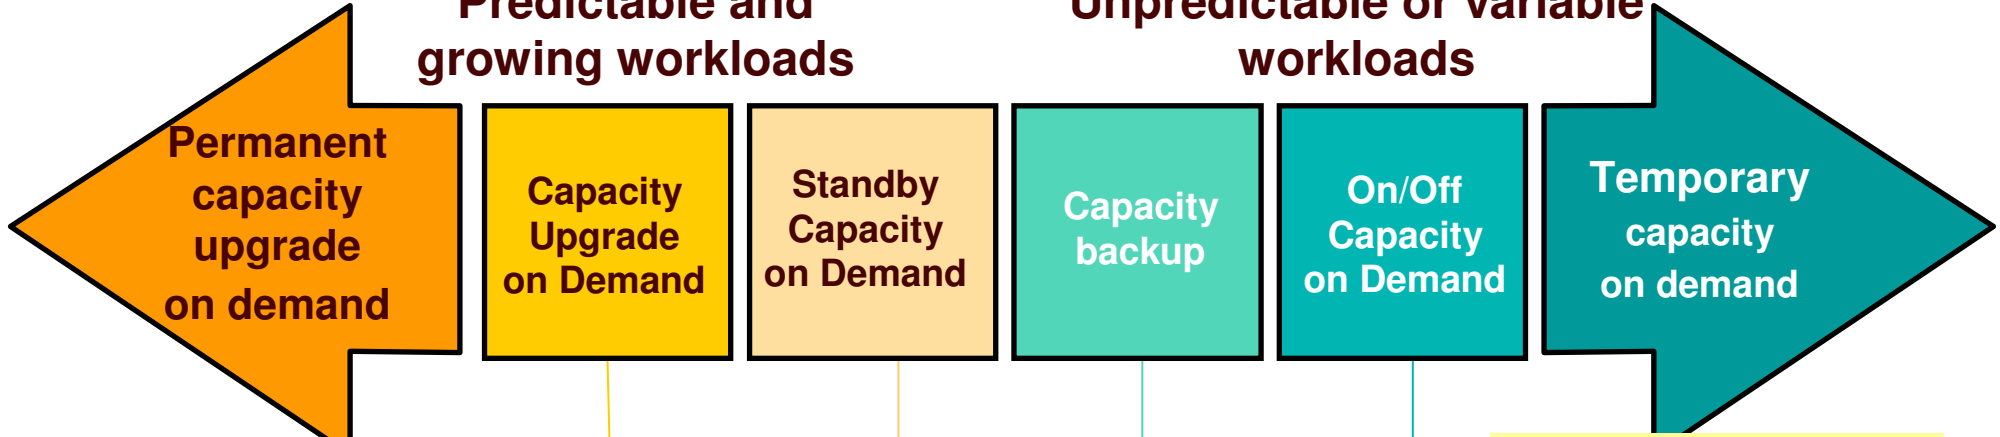
Exceptions: Disruptive if ALL current PUs are converted to different types May require individual LPAR disruption if dedicated PUs are converted.

Capacity on Demand

a must in a consolidated and simplified on demand infrastructure

Predictable and growing workloads

Unpredictable or variable workloads



- Enables scale up on same hardware
- Little impact on current operations
- Predictable pricing model

Provides non-disruptive growth, limited downtime and payment when activated

Cost-effective disaster recovery resources

Assures capacity is always available and allows for reduced purchased or financed capacity

Provides processing for planned and unplanned spikes to improve cost control

- Business continuity
- Cost efficient
- Rapidly available

Matching resource acquisition to business needs dynamically - exactly when you need it





IBM System z9
Availability

Concurrent MES upgrade

Redundant I/O Interconnect
Concurrent Book Upgrade/Replace

Concurrent Oscillator Switch

other items...



ENABLING BUSINESS.
A THROUGH Z.

z9 – Providing new levels of availability

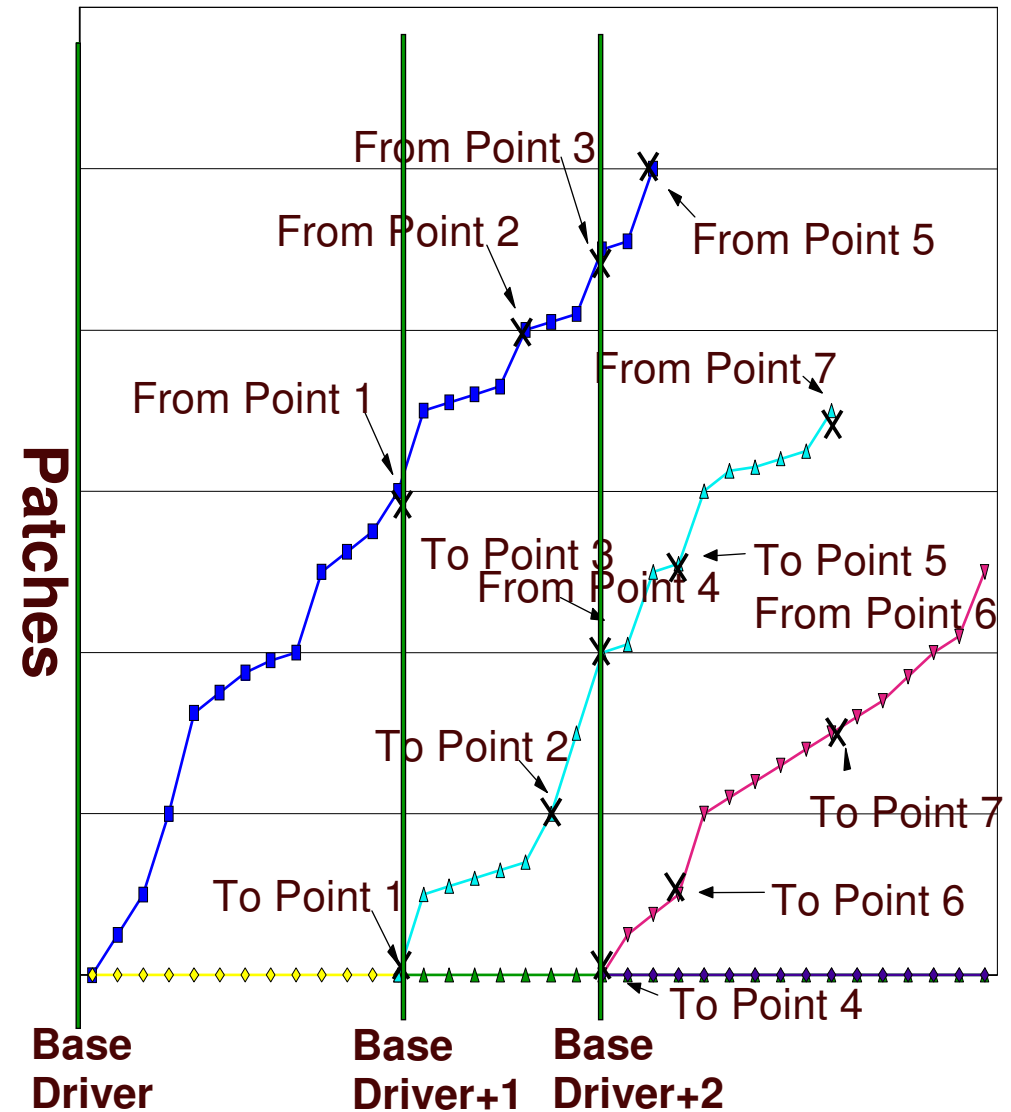
- **Improving the application of HW driver maintenance:***
 - ▶ Concurrent microcode level upgrade (at synch points)
 - ▶ reducing planned outages using enhanced driver maintenance
- **New enhanced book availability** (*EC models only*)
 - ▶ Enhancing recovery of resources
 - ▶ Improving ability to nondisruptively add/repair memory resources
- **Redundant I/O Interconnect (RII)**
 - ▶ Allows for dynamic book replacement without I/O loss (*EC models*)
 - ▶ Allows for MBA replacement without I/O loss
- **Extending capability for Capacity Backup (CBU) to include specialty engines**
 - ▶ *Sub-capacity CBU's are supported*
- **Improving memory availability with flexible memory offering** (*EC models only*)
- **Dynamic Oscillator switchover**

* Customer pre-planning is required and may require purchasing additional hardware resources



Enhanced Driver Maintenance

- The ability to concurrently move from one patch point on major driver N, to a patch point on major driver N+1
- Cannot move any to any, must move from a specific "from" patch bundle to a specific "to" patch bundle
- A Limited number of specific cross-over bundles will be defined for each driver.
- Crossover bundles (target dates) will be communicated (via resource link) in advance for planning.
- Disruptive driver upgrades are permitted at any time
- Concurrent cross-over from driver N to driver N+1, to driver N+2 must be done serially. No composite moves.
- No concurrent back off possible. Must move forward to driver N+1 once CDU is initiated. Catastrophic errors during CDU may dictate a disruptive outage



z9-EC Enhanced Driver Maintenance

- **Ability to concurrently install and activate a new driver***
 - ▶ Can eliminate common planned outage
 - ▶ Select window of opportunity within code maintenance stream
 - ▶ Like some concurrent patches, may need to vary off/on certain devices
- **The ability to concurrently move from one patch point on major driver N, to a patch point on major driver N+1.**
 - ▶ Cannot move any-to-any, must move from a specific "from" patch bundle to a specific "to" patch bundle
- **A limited number of specific crossover bundles will be defined for a driver.**
- **Disruptive driver upgrades are permitted at any time**
- **Concurrent cross-over from driver N to driver N+1, to driver N+2 must be done serially; no composite moves**
- **No concurrent back-off possible. Must move forward to driver N+1 once enhanced driver maintenance is initiated. Catastrophic errors during update may dictate a disruptive outage**

*Requires proper configuration and planning for exploitation

z9-EC Enhanced Driver Maintenance limitations

- **Specific complex code changes may dictate a disruptive driver upgrade. This will be alerted in advance for planning purposes.**
 - ▶ Design data fixes
 - ▶ CFCC level change
 - ▶ LPAR code fixes

- **Non-QDIO OSA CHPID types will still require CHPID Vary OFF/ON in order to activate new code. Not needed if these types do not exist in the IOCDS**
 - ▶ Adapter type: OSA3270 (OSC)
 - ▶ Adapter type: Ethernet (OSE)

- **Crypto code load will require a Config OFF/ON in order to activate new code. Not needed if crypto is not installed.**

- **FICON and FCP code changes involving code loads will require CHPID "reset" to activate. Base code is infrequently changed**

z9 BC Redundant I/O Interconnect (RII)

Function Provided:

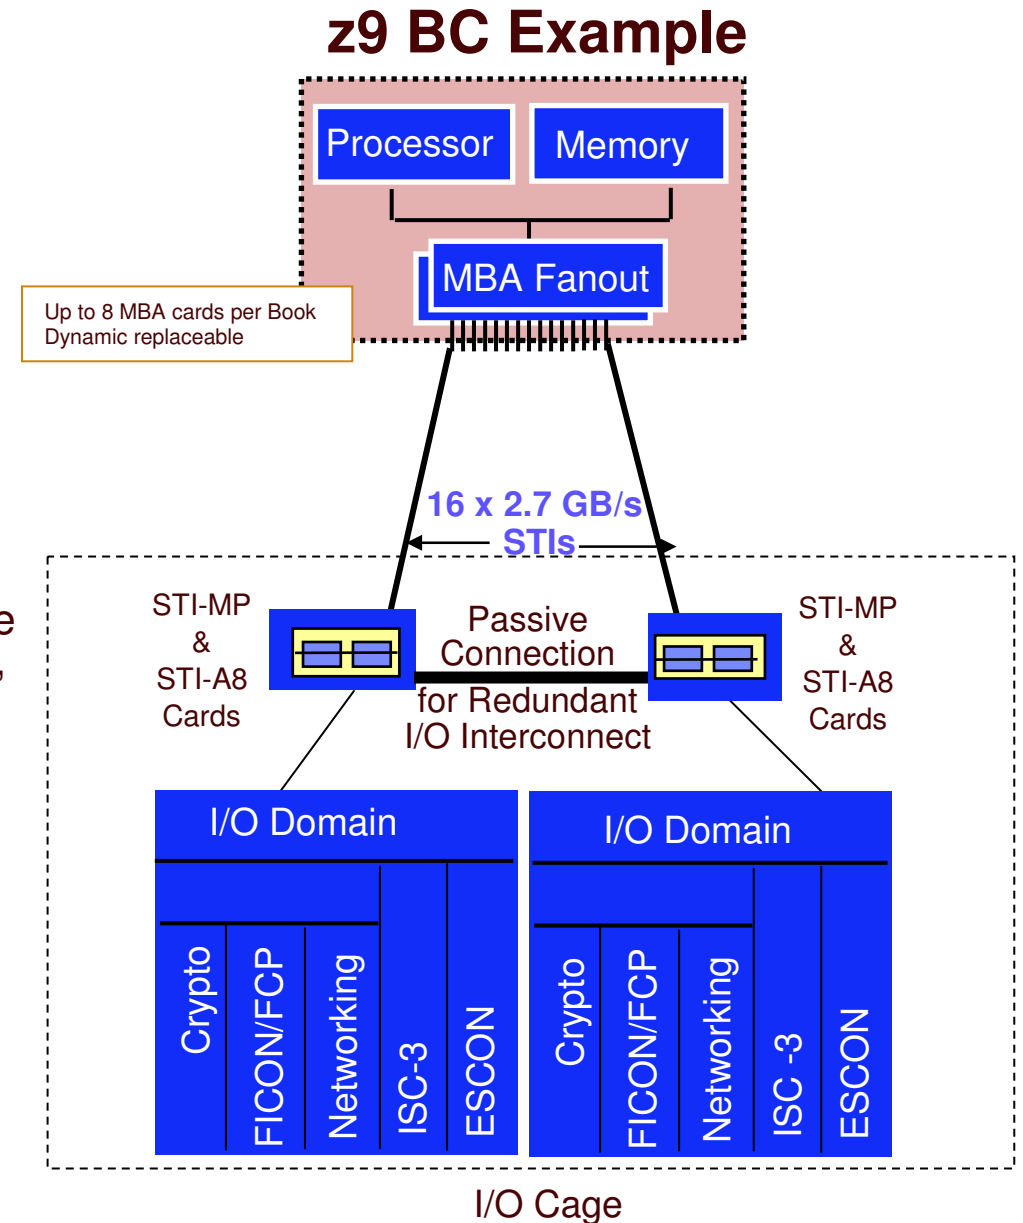
- ▶ Support MBA fanout card and/or STI cable maintenance while maintaining system connectivity to the attached I/O
- ▶ Manually initiated for configuration and maintenance actions
- ▶ System initiated for recovery support

Possible Benefits:

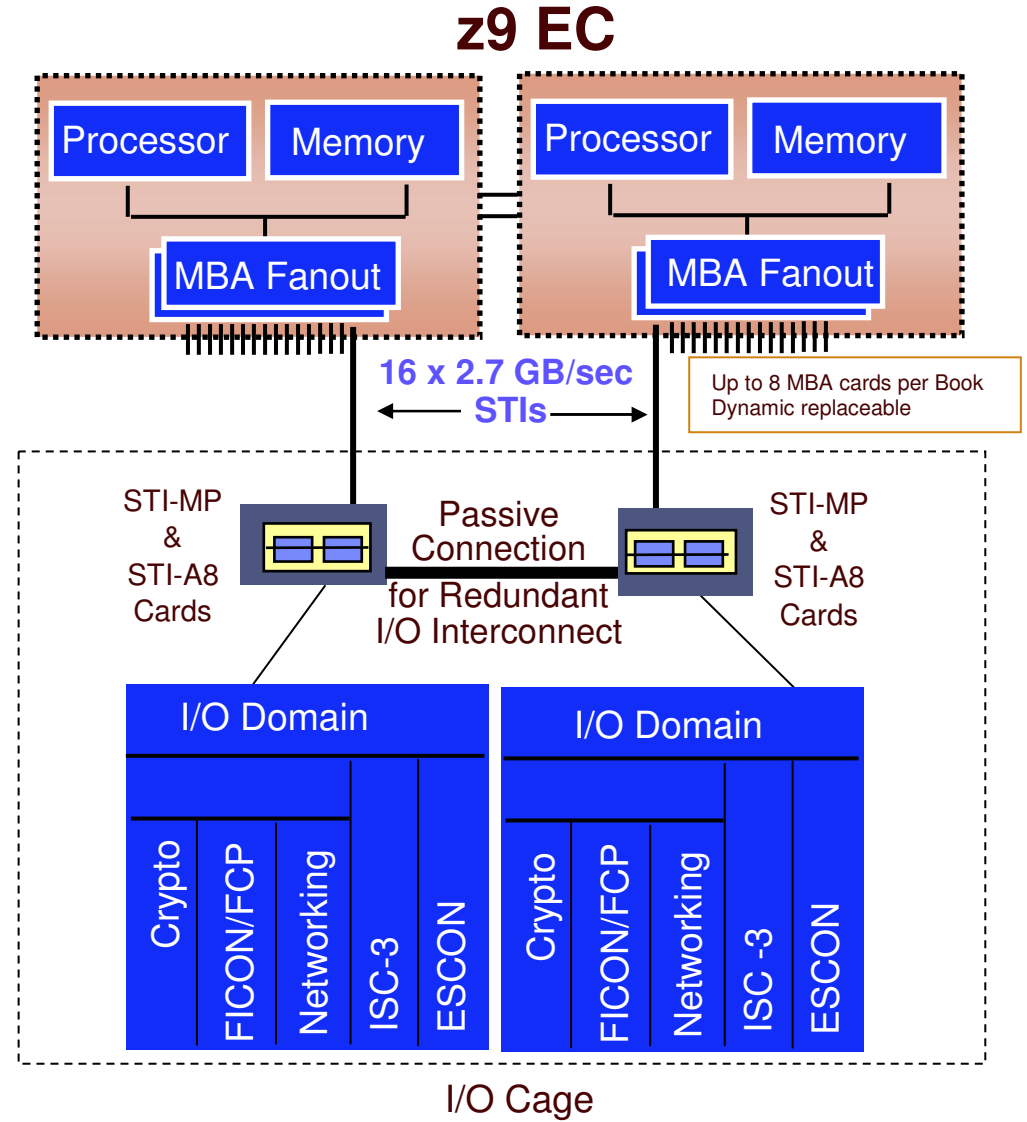
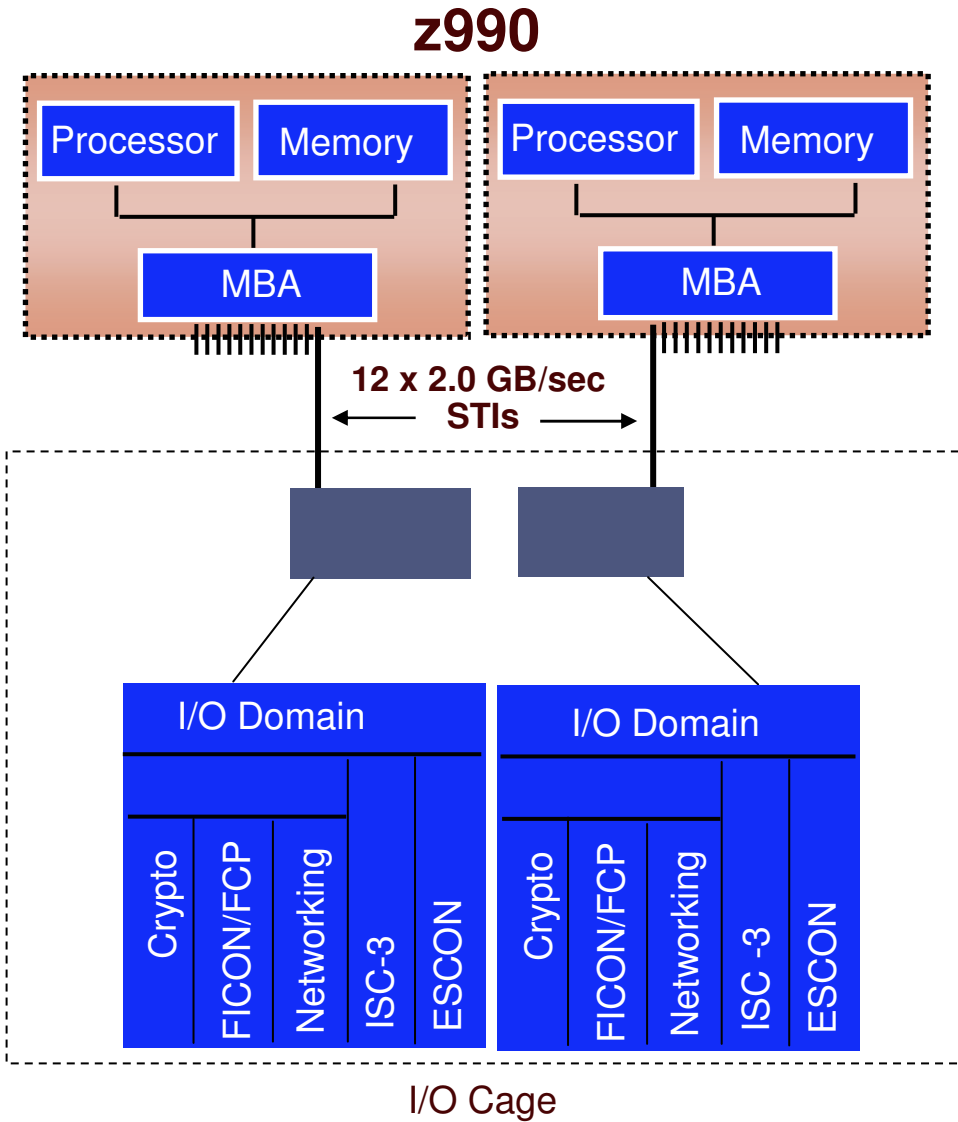
- ▶ No loss of connectivity to I/O (except for ICB-4)
- ▶ Protects up to 60 CHPIDs / STI, 16 STIs
- ▶ No planning required or dependency on alternate I/O attachment interfaces (e.g. FICON, Ethernet, ESCON, ...)
- ▶ Automatic traffic re-routing and return to original configuration upon maintenance completion

I/O Characteristics

- ▶ All PUs can access all adapters via Shared Memory and I/O Infrastructure
- ▶ Provide Dual STI Attachment for each STI-MP, each with unique MBA connectivity via MBA fanout cards within same Book



z9 EC vs. z990 System Structure



z9 EC Redundant I/O Interconnect (RII)

Function Provided:

- ▶ Support dynamic book maintenance while maintaining system connectivity to the attached I/O (requires planning)
- ▶ Support MBA fanout card and/or STI cable maintenance while maintaining system connectivity to the attached I/O
- ▶ Manually initiated for configuration and maintenance actions
- ▶ System initiated for recovery support

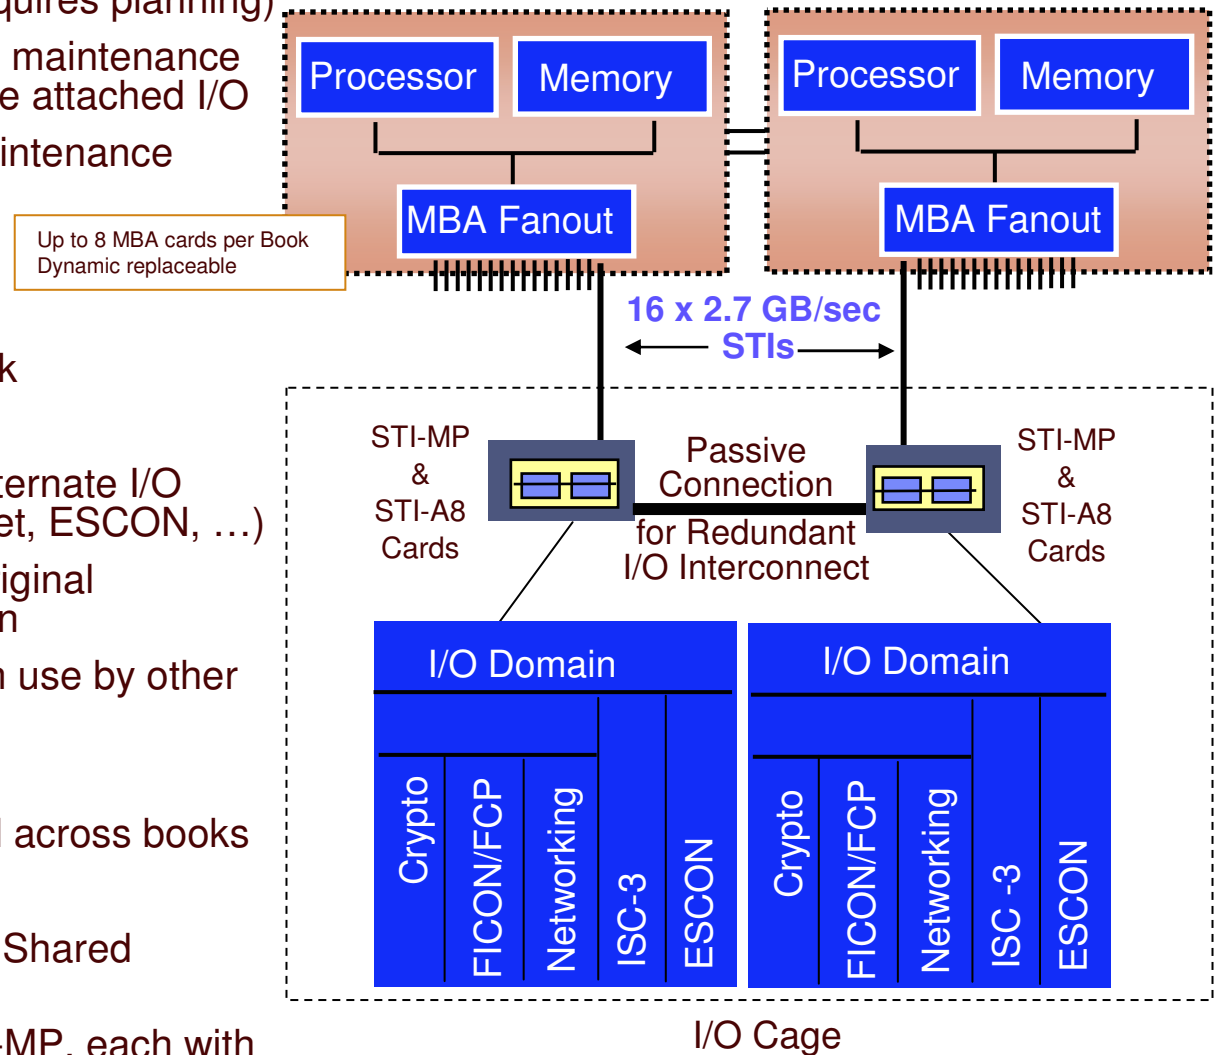
Possible Benefits:

- ▶ No loss of connectivity to I/O from any book
- ▶ Protects up to 60 CHPIDs / STI, 16 STI's
- ▶ No planning required or dependency on alternate I/O attachment interfaces (e.g. FICON, Ethernet, ESCON, ...)
- ▶ Automatic traffic re-routing and return to original configuration upon maintenance completion
- ▶ Book removal does not cause loss of I/O in use by other Books

I/O Characteristics

- ▶ On a multi-book machine, I/O is distributed across books
- ▶ I/O has no intrinsic BOOK AFFINITY
- ▶ All Books/PUs can access all adapters via Shared Memory and I/O Infrastructure
- ▶ Provide Dual STI Attachment for each STI-MP, each with unique MBA connectivity via different Books or MBA fanout cards within same Book

z9 EC Example



z9-EC Book Replacement Flow

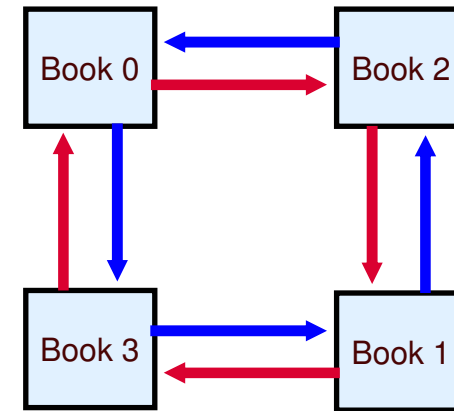
■ Prepare for Book Replacement

- ▶ Standalone SE Application
- ▶ Customer Workload Reduction

■ Perform Book Replacement

- ▶ Verification of Prepared System (*pre-defined activation profile*)
- ▶ Resource Evacuation (*pre-defined activation profile*)
- ▶ Fence Targeted Book
- ▶ Physically Repair/Replace/Upgrade Targeted Book
- ▶ Book Add Operation
 - Rebalance Processor Resources
 - Optimize Memory Resources
 - Restore I/O Connectivity (ICB's)

■ Restore Customer Workload



z9-EC: Options to Configure for Maximum Availability

- **Concept: configure enough physical memory and engines such that loss of a book does not result in any degradation**
 - ▶ During a degraded restart in the rare event of a book failure
 - ▶ During a book replacement for repair or physical memory upgrade
 - ▶ Added benefit: reasonable 'plan ahead' CUoD memory upgrade paths

- **How ?**
 - ▶ **Select PU/model configuration such that 100% of the customer-owned PUs can be activated even when one book within a model is fenced.**
 - a maximum of 8 customer PUs are configurable on the S18
 - a maximum of 18 customer PUs are configurable on the S28
 - a maximum of 28 customer PUs are configurable on the S38
 - a maximum of 40 customer PUs are configurable on the S54

 - ▶ **Requires no special feature codes for PU/model configuration**

 - ▶ **New optional features for Memory: deliver physical memory cards such that 100% of the purchased memory increment can be activated even when one book is fenced**



System z RAS Overview for high end Systems

	z900	z990	z9 BC/EC
Microcode Driver Updates	6 Hr Scheduled outage	6 Hr Scheduled outage	Concurrent*
Book Replacement**	Not Applicable	Scheduled Outage	Concurrent
Memory Replacement	Scheduled Outage	Scheduled Outage	Concurrent (Book Offline)
ECC on Memory Control Circuitry (EX: SMI)	Unscheduled Outage	Unscheduled Outage	Transparent
Memory Bus Adapter (MBA) Replacement	Scheduled Outage. Lose connectivity to I/O Domain	Scheduled Outage. Lose connectivity to I/O Domain	Concurrent. Connectivity to I/O Domain remains
STI Failure	As for MBA	As for MBA	As for MBA
Oscillator Failure	Unscheduled Outage	Unscheduled Outage	Transparent
Processor Upgrades	Concurrent	Concurrent	Concurrent
Physical Memory Upgrades	Scheduled Outage	Scheduled Outage	Concurrent ** (Book Offline)
I/O Upgrades	Concurrent	Concurrent	Concurrent
Spare PUs	1 System	2 / Book	2 / System **

*In select circumstances

**System z9 EC ONLY - Customer pre-planning required, may require acquisition of additional hardware resources



IBM System z9 Channel and I/O System Connectivity considerations



ENABLING BUSINESS.
A THROUGH Z.

IBM System z9 Channel and I/O System - Connectivity considerations

**up to 4 Logical Channel Subsystems
Multiple Subchannel Sets (63.75K + 64K)
1024 channels - 336 FICON**

FCP Sharing (zVM & zLinux)

**MIDAW Architecture
FICON Express / Express 2 (2Gbps) / Express 4 (4Gbps)**

**OSA Express 2 - up to 10Gb Ethernet
Hipersockets**

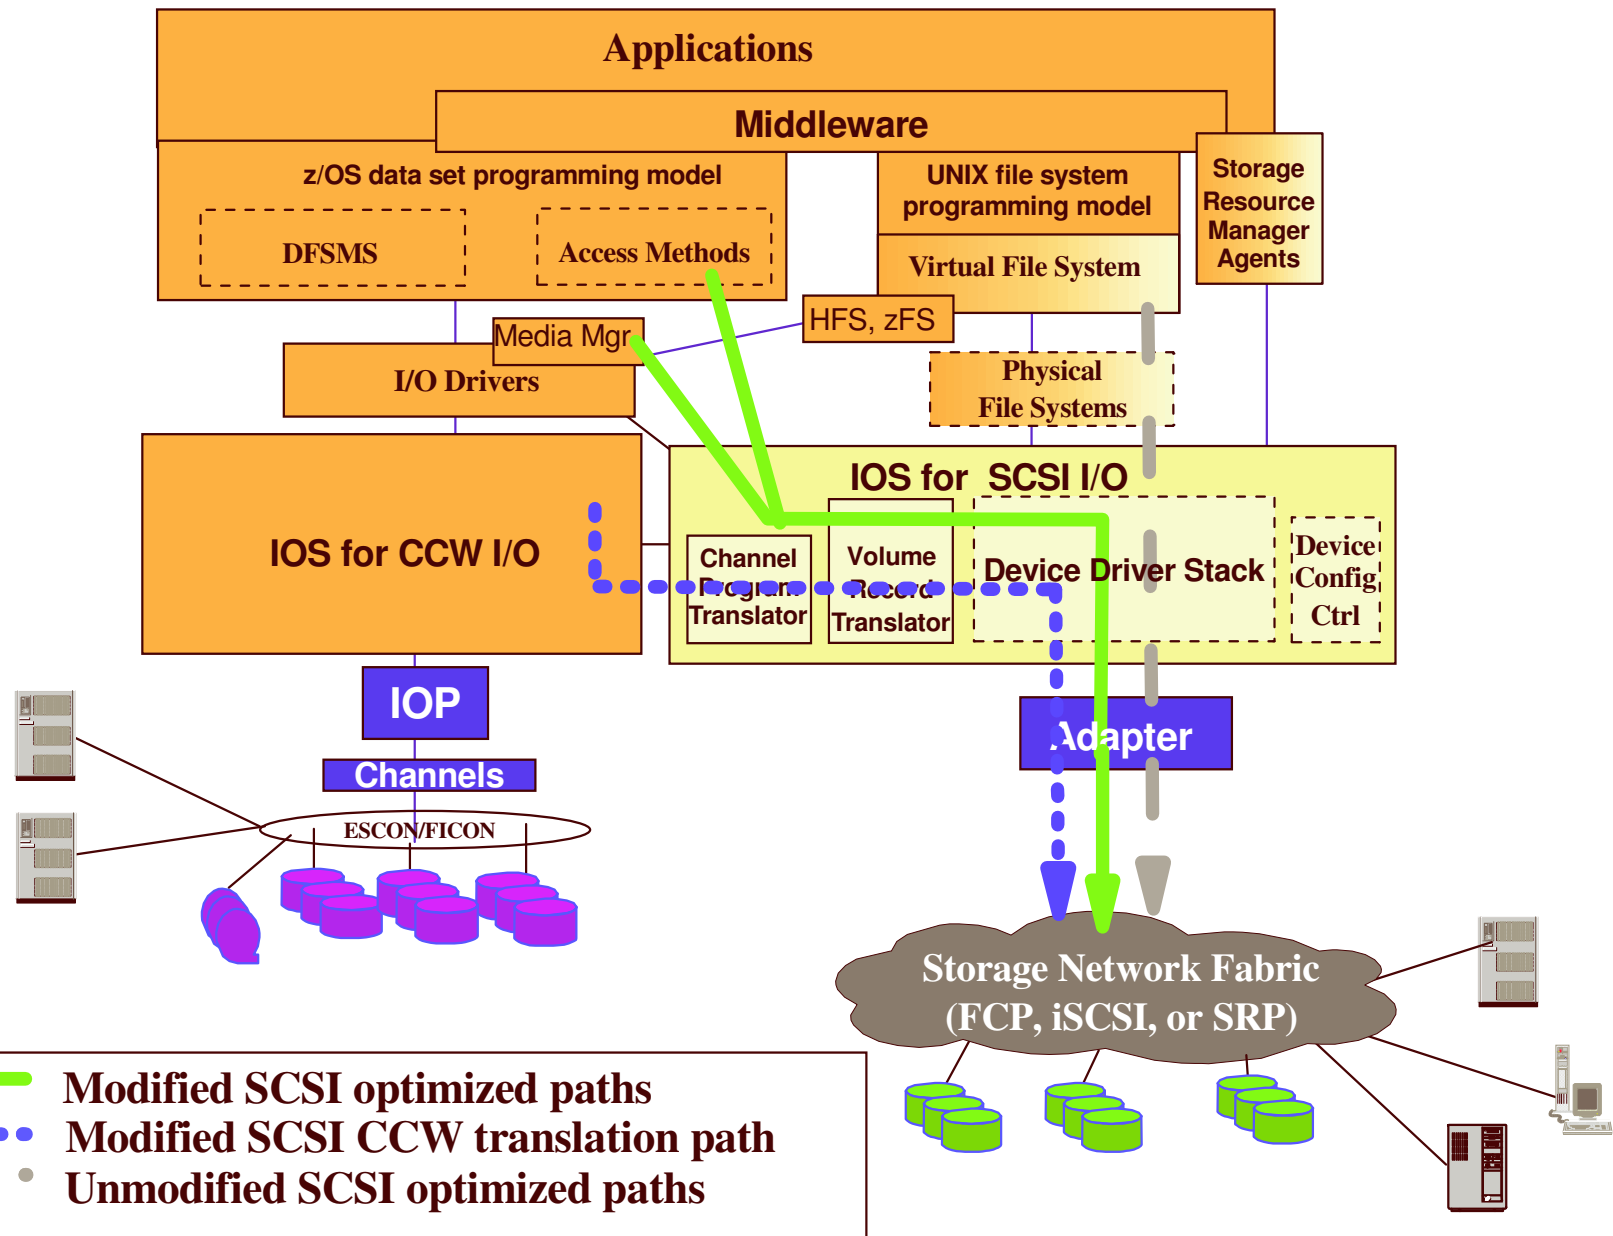
**Enhanced CP Assist
new Crypto Express 2 Adapter with combined SSL Accelerator & Coprocessor**

Peer mode Coupling links only – ICB4

Future options or possibilities

**higher bandwidth and Bus speed - Faster I/O exceeding 10Gbs
Exploitation of Infiniband - internal and external
FCP support within z/OS – sharing of I/O with Unix/Wintel**

System z - potential future I/O structure



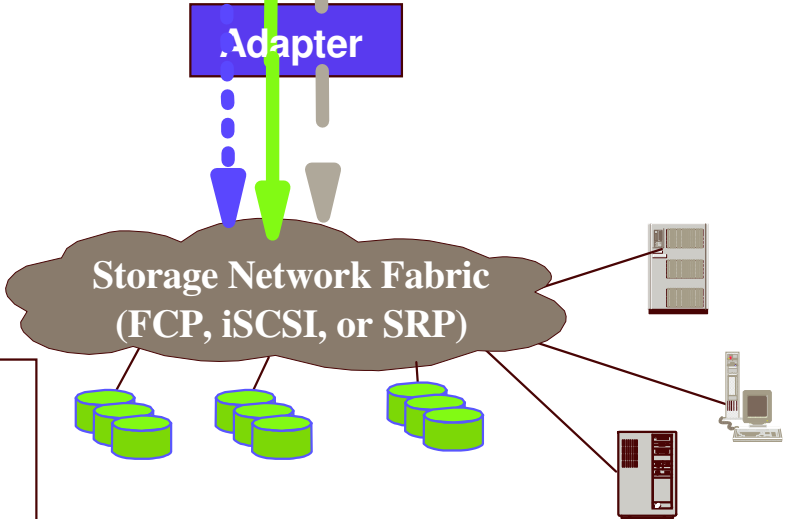
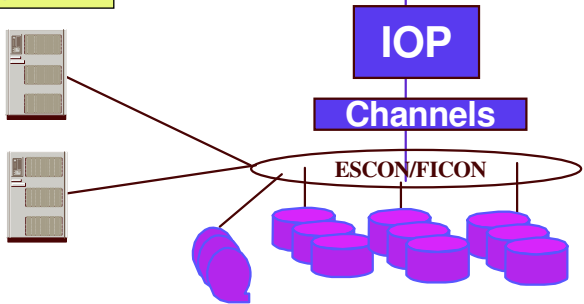
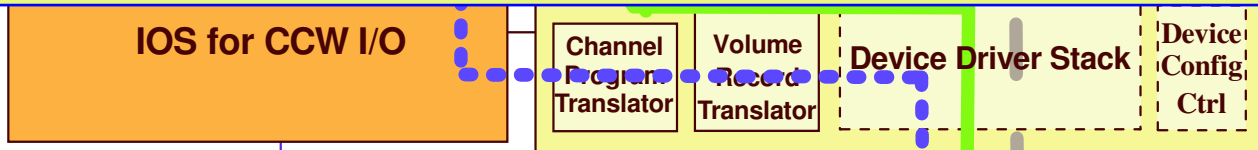


System z - potential future I/O structure

Applications

Full Sharing of the external storage infrastructure between the Mainframe (z/OS & Linux) and UNIX and Wintel

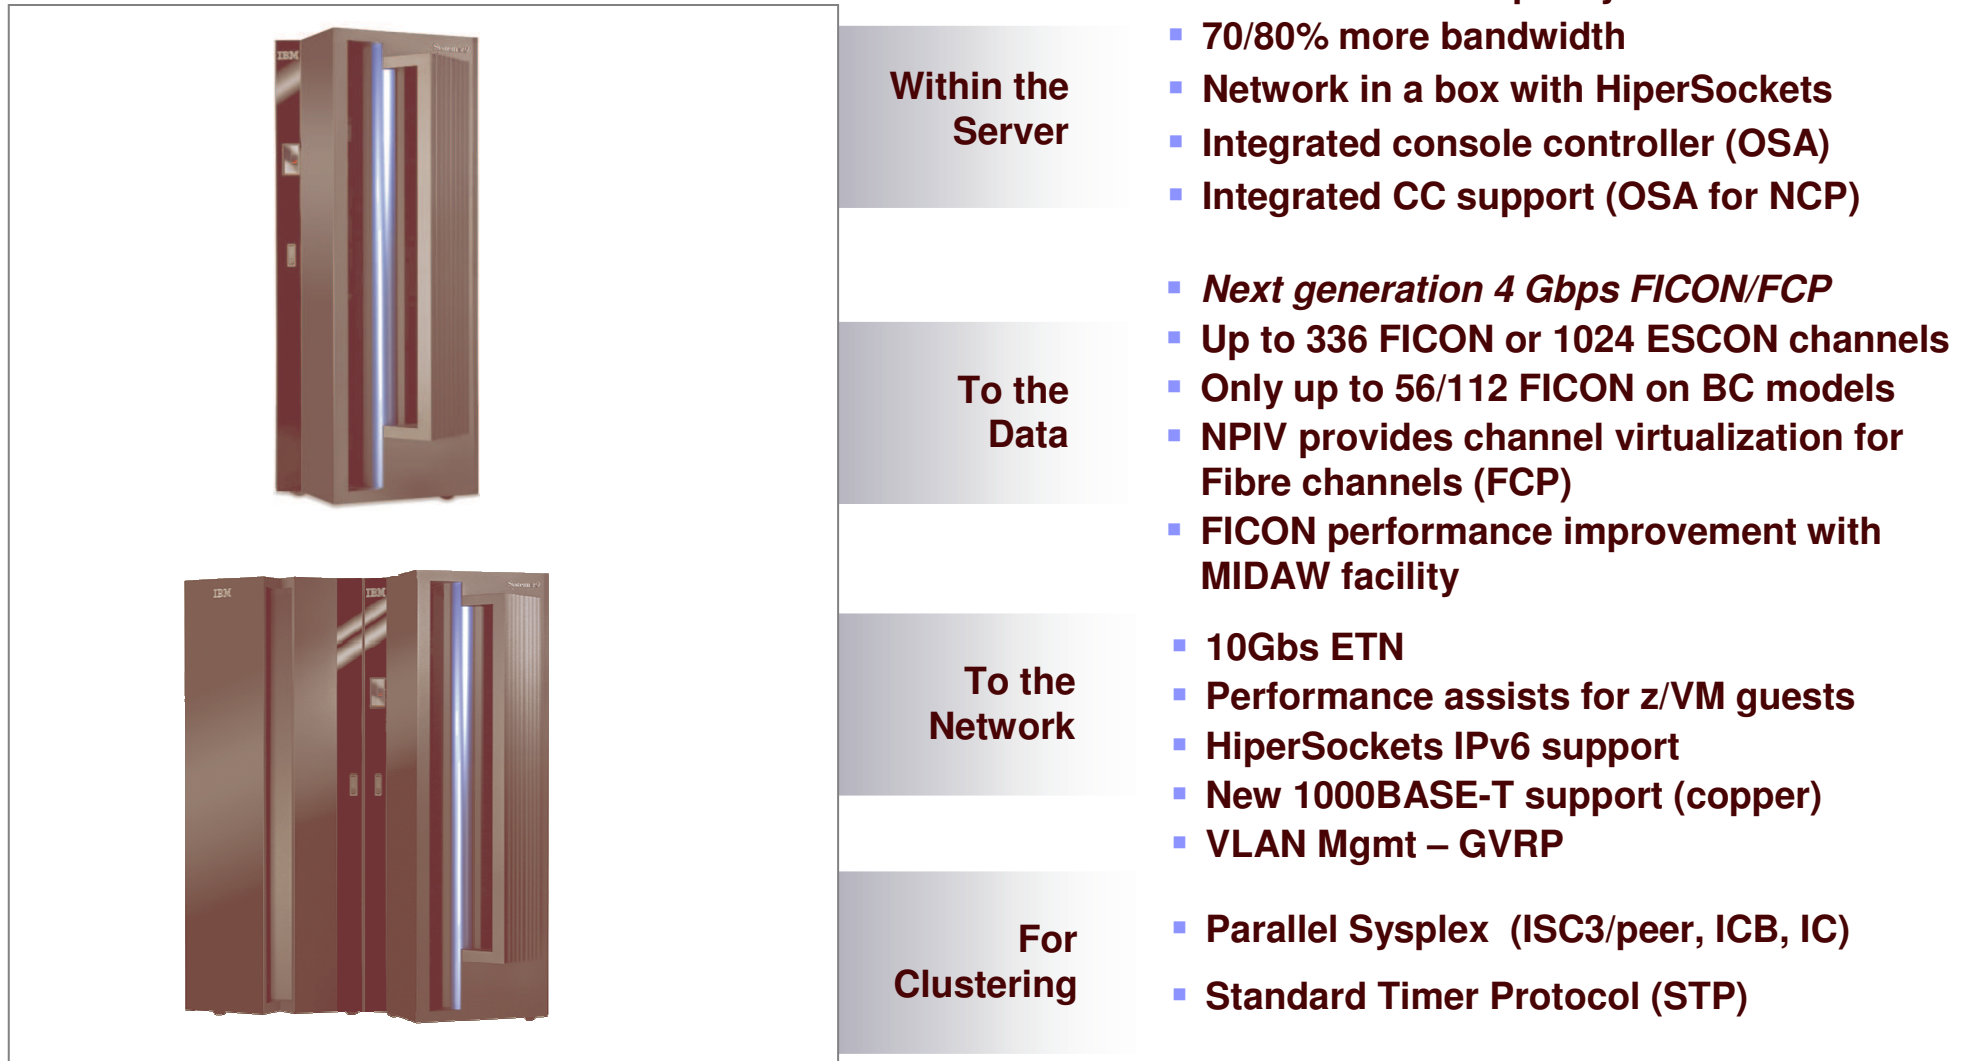
Plus
Infiniband
10Gbs++
x00GB bandwidth



 **Modified SCSI optimized paths**
 **Modified SCSI CCW translation path**
 **Unmodified SCSI optimized paths**

z9 – Delivering enhanced connectivity for the system

Within the server, between servers, to the data and to the network



System z9 I/O Overview

I/O Enhancements

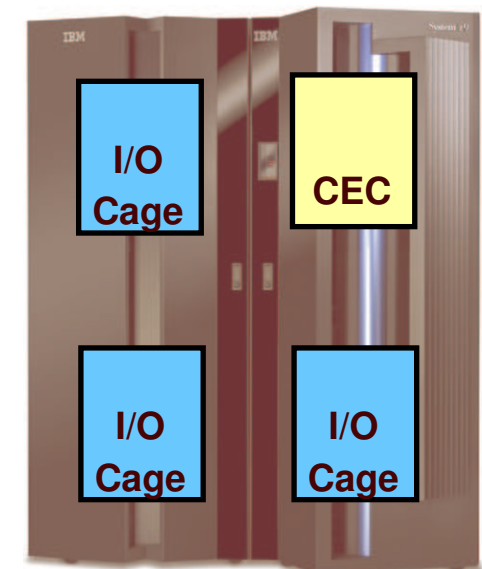

- ▶ STI Capacity = 2.7GB/sec
- ▶ Up to 70% to 80% more bandwidth than respectively z890 and z990
- ▶ Up to 28 FICON Express, FICON Express2, FICON Express4 features / book
 - 4 channels/feature FICON/FCP
 - 1, 2, 4 Gbps auto-negotiated. 4 Gbps for FICON Express4
- ▶ Redundant I/O Interconnect (RII)
- ▶ Modified Indirect Data Address Word (MIDAW) facility
- ▶ Multiple (2) Subchannel sets (MSS) and increase to 63.75K Subchannels for Set-0
- ▶ **System z9 BC:** Up to 16* x 2.7GB STI's. (7 STIs max for the single I/O cage)
- ▶ **System z9 EC:** Up to 64* x 2.7GB STI's. (21 STIs max for 3 I/O cages)

Storage Area Networks (SANs) enhancements

- ▶ N_Port ID Virtualization
- ▶ Program Directed re-IPL
- ▶ FICON Link Incident Reporting

Networking enhancements

- ▶ HiperSockets IPv6
- ▶ OSA Express 2 - up to 10Gbs
- ▶ OSA-Express2 1000BASE-T Ethernet
- ▶ OSA-Express2 OSN (OSA for NCP support)
- ▶ GARP VLAN management (GRVP)



* z9 BC & EC exploits a subset of its designed I/O capability

System z9 Channel Type and Crypto Overview

▪ FICON/FCP

- ▶ FICON Express4 (4Gbps)
- ▶ FICON Express2 (2Gbps)
- ▶ FICON Express
- ▶ up to 336, 112, 56 chpid's on z9 EC, BC S07, BC R07

▪ Networking

- ▶ OSA-Express2
 - Gigabit Ethernet LX and SX
 - 10 Gigabit Ethernet LR
 - 1000BASE-T Ethernet
- ▶ OSA-Express (*carry forward on upgrade*)
 - Gigabit Ethernet LX and SX
 - 1000BASE-T Ethernet (copper and fiber)
 - Fast Ethernet
- ▶ HiperSockets

▪ Coupling Links

- ▶ ISC-3 (Peer mode only)
- ▶ ICB-3, ICB-4
- ▶ IC

▪ ESCON


- ▶ Up to 420 channels on BC models
- ▶ Up to 1024 channels on EC models

▪ Crypto – Crypto Express 2

- ▶ Configurable Coprocessor or Accelerator

▪ Channel types not supported:

- ▶ FICON (pre-FICON Express)
- ▶ OSA-Express Token-Ring (SOD Oct 2004)
- ▶ PCIXCC
- ▶ PCICA
- ▶ ICB-2 (SOD 2003)
- ▶ ISC-3 Links in Compatibility Mode (SOD April 2004)
- ▶ Parallel (use ESCON Converter)
- ▶ OSA-Express ATM 155
- ▶ OSA-2



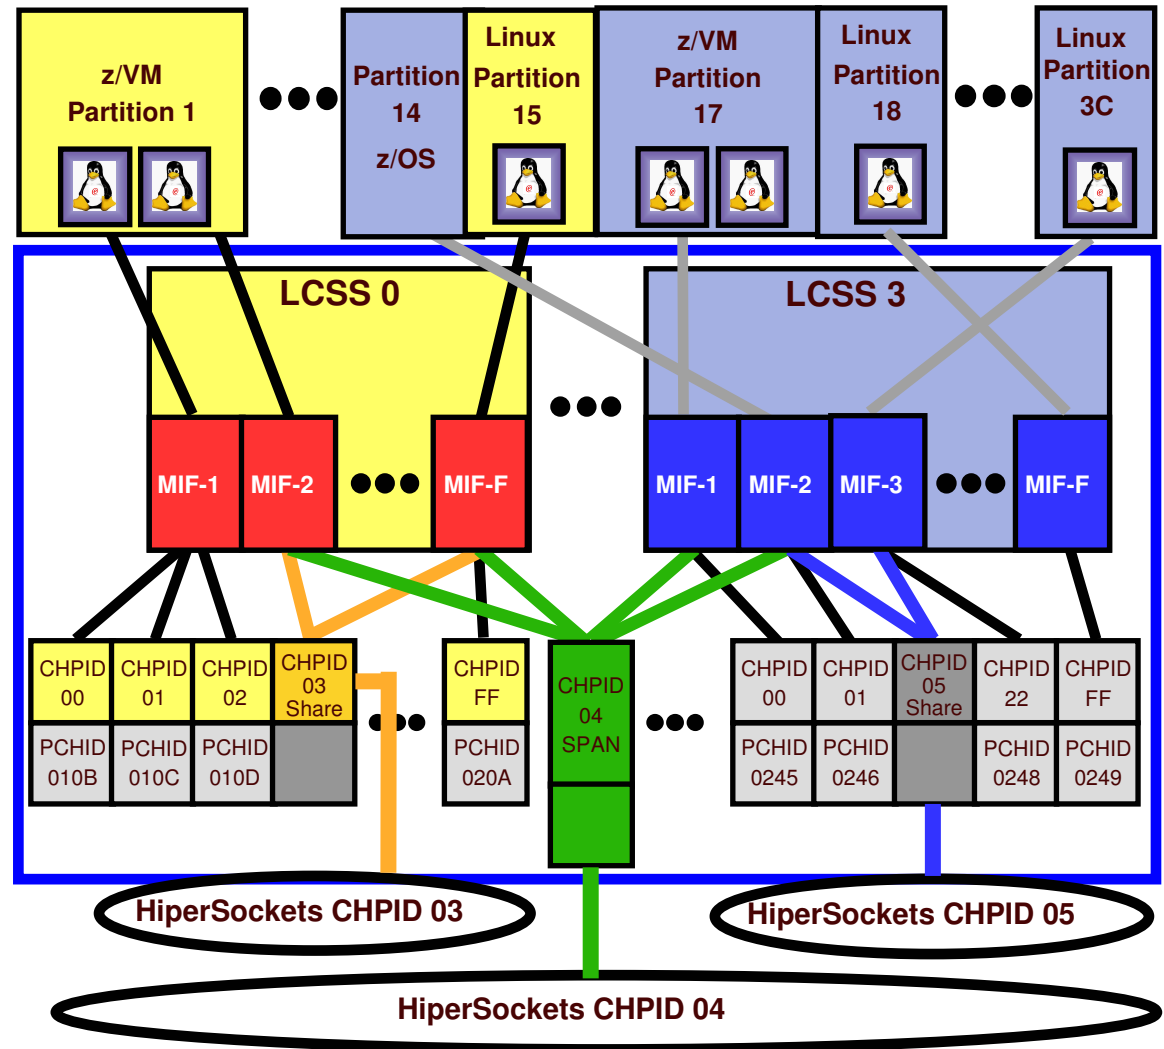
Note: Only ICB cables orderable.- All other cables have to be sourced separately.

System z9 I/O Connectivity – some details

- **HiperSockets, up to 16 (internal LAN)**
- **Crypto Express2, up to 8**
 - ▶ Now configurable, Coprocessor for secure key transactions and Accelerator for SSL acceleration
- **Coupling Links, up to 64 (EC ,models) in combination**
 - ▶ IC (up to 32), ICB-3 (up to 16), ICB-4 (up to 16), ISC-3 (up to 48 active links)
- **16-port ESCON**
 - ▶ Up to 420 channels on BC models
 - ▶ Up to 1024 channels on EC models
- **BC Models - FICON Express4 and Express2, FICON Express**
 - ▶ Up to 84 features / 336 channels (FICON Express4 and Express2)
- **EC Models - FICON Express4 and Express2, FICON Express**
 - ▶ Up to 28 features / 112 channels (FICON Express4 and Express2)
 - ▶ Up to 28 features / 56 channels (FICON Express4 2-port)
- **OSA-Express2, OSA-Express**
 - ▶ Up to 24 features
 - ▶ Fast Ethernet, 1000BASE-T Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet

System z9 HiperSockets IPv6

- **Internet Protocol Version 6 (IPv6)**
- **More unique IP addresses**
 - ▶ Expands the IP address space
 - From 32 bits to 128 bits
 - ▶ Follow-on to IPv4
- **Minimum software:**
 - ▶ z/OS V1.7
 - ▶ z/VM V5.2 with PTFs in May 2006



Very High Speed Interconnection between programs running z/OS, z/VM or Linux

System z9 OSA-Express2 GbE and 10 GbE

10 Gigabit Ethernet LR (long reach)

- ▶ One port per feature
- ▶ CHPID type OSD (QDIO)
- ▶ 9 micron single mode fiber, SC Duplex connector

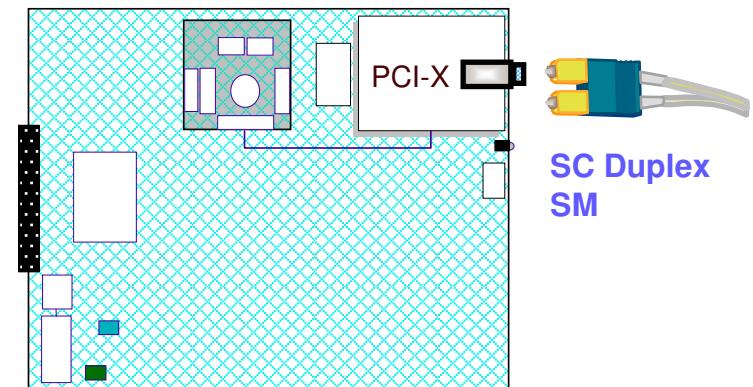
Gigabit Ethernet features, 2 ports per feature

- ▶ CHPID types OSD (QDIO), OSN (OSA for NCP)
- ▶ Designed to achieve line speed - 1 Gbps in each direction
- ▶ Gigabit Ethernet LX (Long wavelength)
 - 9 micron single mode fiber, LC Duplex connector
- ▶ Gigabit Ethernet SX (Short wavelength)
 - 50 or 62.5 micron multimode fiber, LC Duplex connector

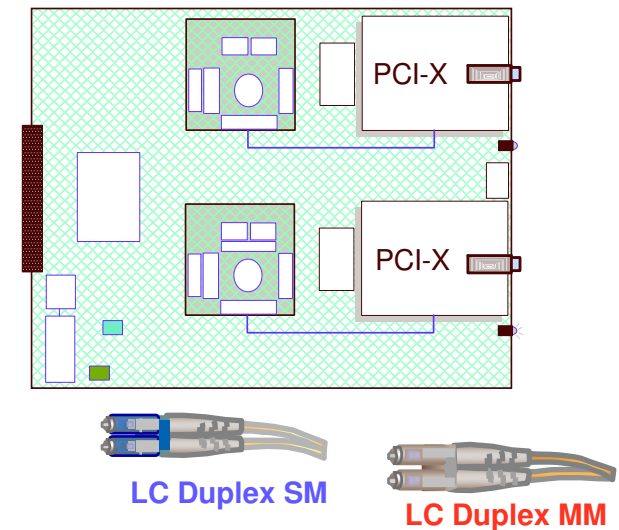
OSA-Express2 GbE and 10 GbE support

- ▶ Large send - offloading TCP segmentation
- ▶ Concurrent LIC update to minimize network traffic disruption
- ▶ 640 TCP/IP stacks - improved virtualization
- ▶ Layer 2 support - protocol-independent packet

10 Gigabit Ethernet
Feature 3368

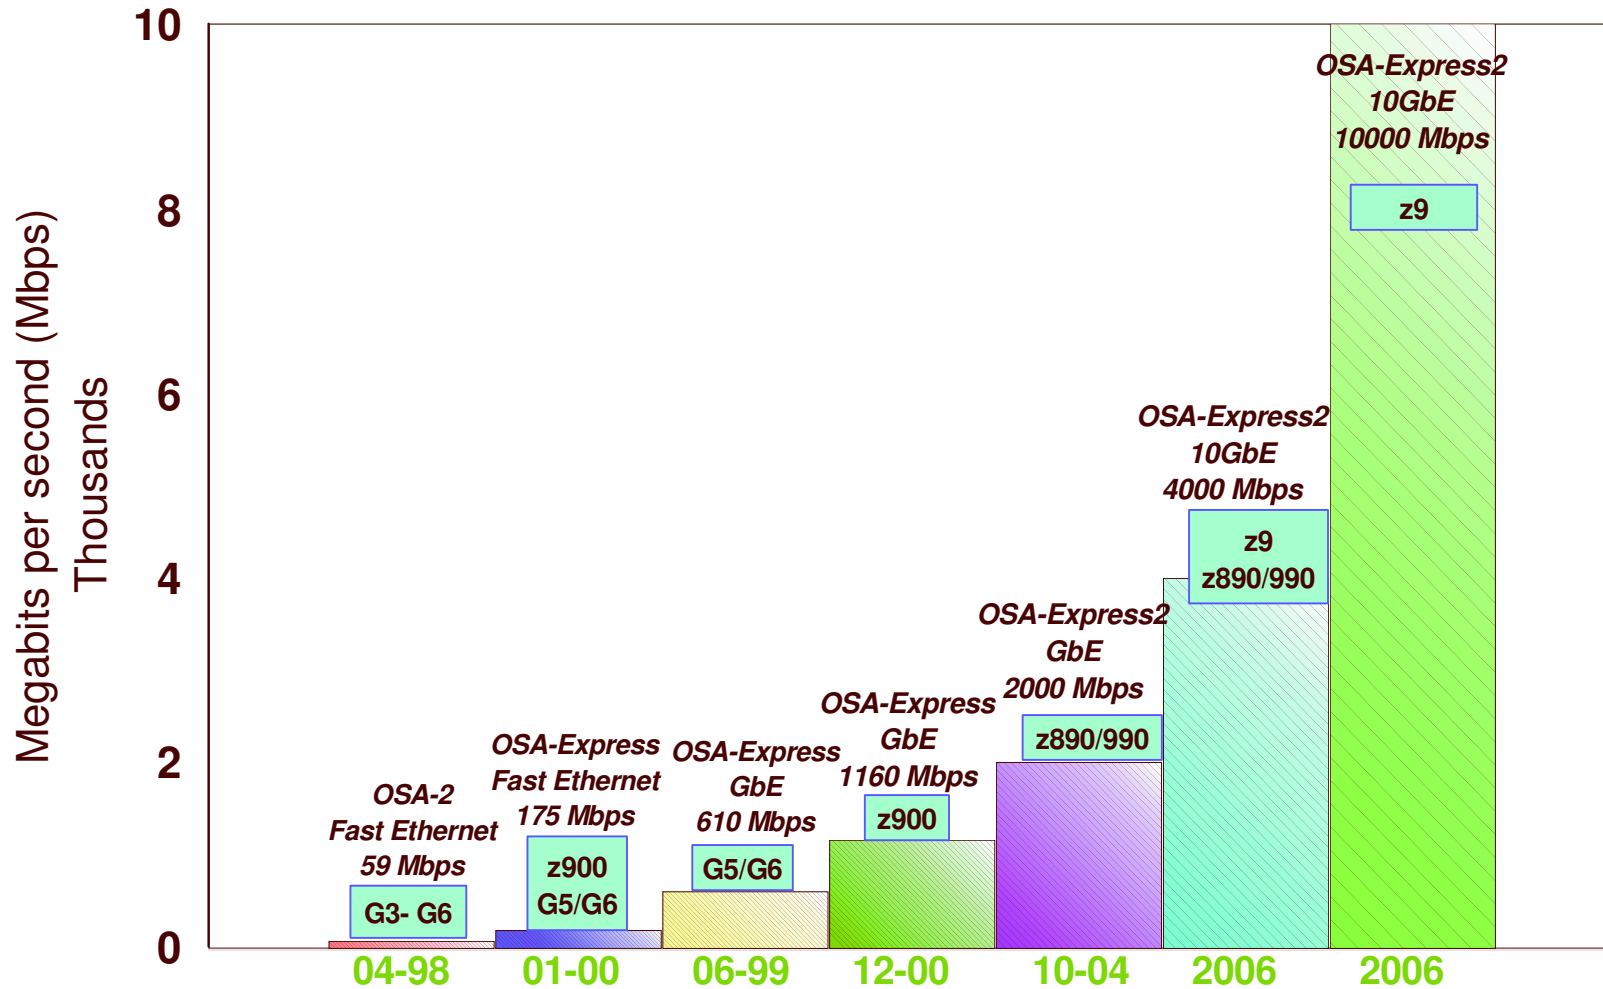


Gigabit Ethernet
Features 3364 (LX), 3365 (SX)





System z Network I/O Performance objectives



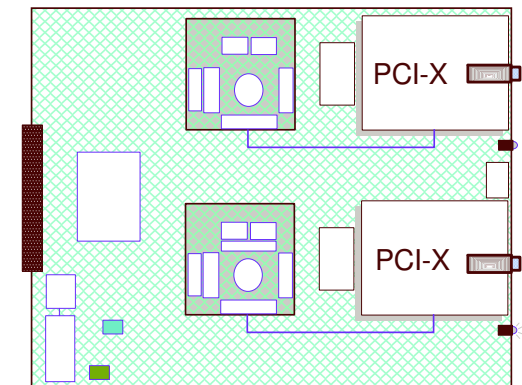
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only

System z9 OSA-Express2 1000BASE-T Ethernet

- **Planned to be available for z890 during 3Q06**
 - ▶ It intended that FC 1366 be no longer orderable
 - ▶ Channel type OSN not supported on z890
- **Supports auto-negotiation to 10, 100, 1000 Mbps over Category 5 copper**
- **Capable of achieving line speed**
 - ▶ Actual throughput is dependent upon environment
- **Supports:**
 - ▶ Large send - **offloading TCP segmentation**
 - ▶ Concurrent LIC update to **minimize network traffic disruption**
 - ▶ 640 TCP/IP stacks - **improved virtualization**
 - ▶ Layer 2 support - **protocol-independent packet**

Mode	CHPID	Description
OSA-ICC	OSC	3270 data streams (console support)
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 NCP	OSN	Channel Data Link Control for Linux NCP

1000BASE-T Ethernet #3366



System z OSA Technology Refresh over time

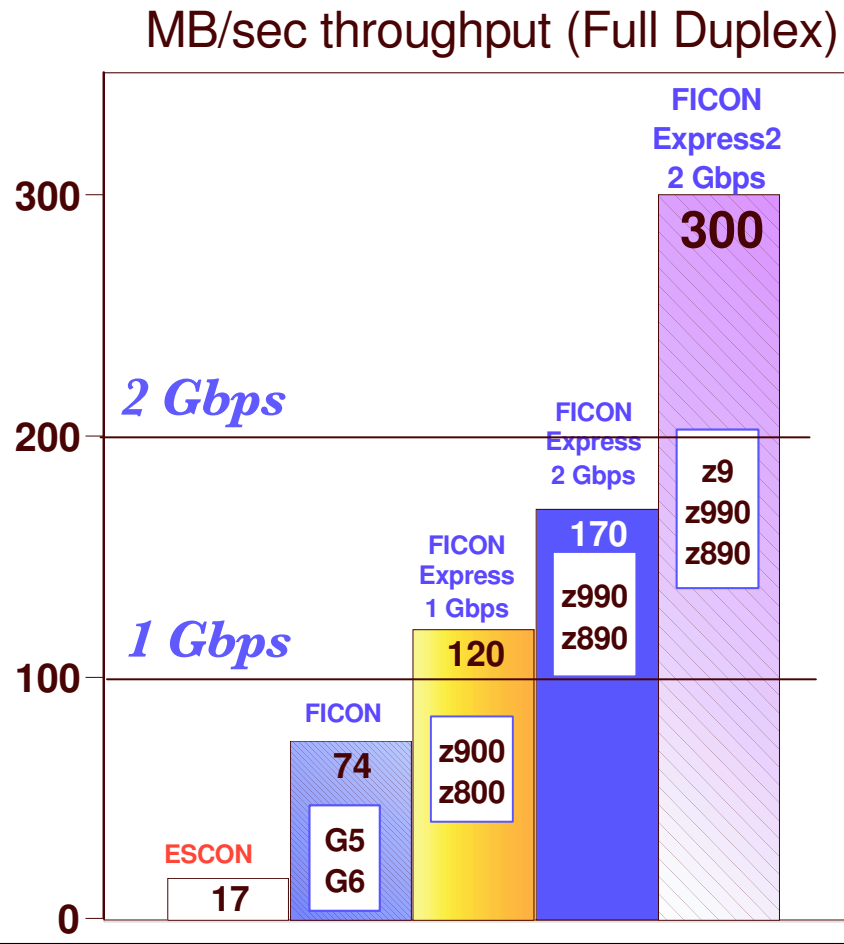
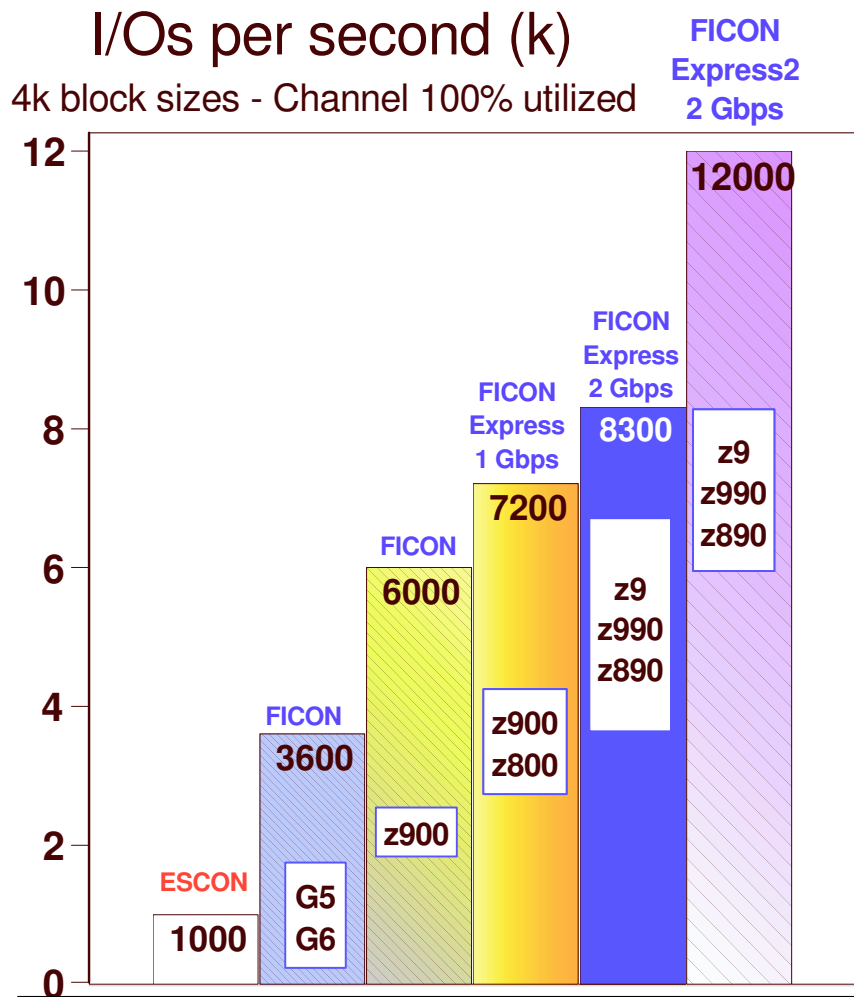
Feature	Feature Name	z900	z800	z990	z890	System z9	CHPIDs
3364	OSA-Express2 GbE LX	N/A	N/A	01/05	01/05	X	OSD, OSN (z9 only)
3365	OSA-Express2 GbE SX	N/A	N/A	01/05	01/05	X	OSD, OSN (z9 only)
3366	OSA-Express2 1000BASE-T Ethernet	N/A	N/A	05/06	05/06	X	OSC, OSD, OSE, OSN (z9 only)
3368	OSA-Epress2 10 GbE LR	N/A	N/A	01/05	01/05	X	OSD
1364	OSA-Express GbE LX	09/04	09/04	06/03	05/04	C	OSD
1365	OSA-Express GbE SX	09/04	09/04	06/03	05/04	C	OSD
1366	OSA-Express 1000BASE-T Ethernet	N/A	N/A	06/03	05/04	C	OSC, OSD, OSE
2362	OSA-Express 155 ATM SM	Y	Y	RPQ	N/A	N/A	OSD, OSE
2363	OSA-Express 155 ATM MM	Y	Y	RPQ	N/A	N/A	OSD, OSE
2364	OSA-Express GbE LX	Y	Y	C	C	C	OSD
2365	OSA-Express GbE SX	Y	Y	C	C	C	OSD
2366	OSA-Express Fast Ethernet	Y	Y	C	C	C	OSD, OSE
2367	OSA-Express Token-Ring	Y	Y	Y	Y	N/A	OSD, OSE
5201	OSA-2 Token-Ring	Y	N/A	N/A	N/A	N/A	OSA
5202	OSA-2 FDDI	Y	N/A	N/A	N/A	N/A	OSA

LX = Long wavelength transceiver, SX = Short wavelength transceiver, LR - Long Reach transceiver
 C = Carry forward on an upgrade from z900 to z990, z800 to z890. Replacements available.
 Y = Supported



FICON Express2 - Potential Performance increase

10Gb/sec feature in the future



FICON Express = SX (2320), LX (2319)
FICON numbers represent native mode (FC mode)

Introducing FICON Express4 for System z9

- **Designed to improve capacity and performance with next generation 4 Gbps FICON/FCP**
 - ▶ Up to 25% improvement in FICON channel throughput when processing a mix of read and write data transfers¹
 - ▶ Up to 65% improvement in FICON channel throughput when processing all read or all write data transfers¹
 - ▶ 220% cumulative MB/sec throughput improvement in DB2 table scan tests for extended format data sets with FICON Express4 on z9 EC with the MIDAW facility compared to FICON Express2 with the IDAW facility on z9-109².
 - **46% without the MIDAW facility on z9EC**
- **Benefits**
 - ▶ Helps to support reduced cost of storage operations and shorter backup windows with faster channel link data rates
 - ▶ Enables migration to higher performance with 1/2/4 Gbps auto-negotiating links
- **FCP channel performance improvements for z/VM and Linux environments¹**
 - ▶ *Up to 50% in FCP channel throughput when processing a mix of read & write large data transfers*
 - ▶ *Up to 100% in FCP channel throughput when processing all read or all write large data transfers*
- **2-port/4-port cards for z9 BC and 4-port for z9 EC**

1. Large sequential data transfers on z9 EC with FICON Express4 operating at 4 Gbps (running z/OS V1.7) when compared to FICON Express2 on z9-109 (running z/OS V1.6)
2. Results of internal DB2 table scan tests with the z9 EC, the MIDAW facility, FICON Express4 operating at 4 Gbps (running z/OS V1.7), and the DS8000 compared to z9-109, and FICON Express2 operating at 2 Gbps (running z/OS V1.6)

Next generation 4 Gbps FICON/FCP ... helping to improve capacity and performance

z9 FICON Express4

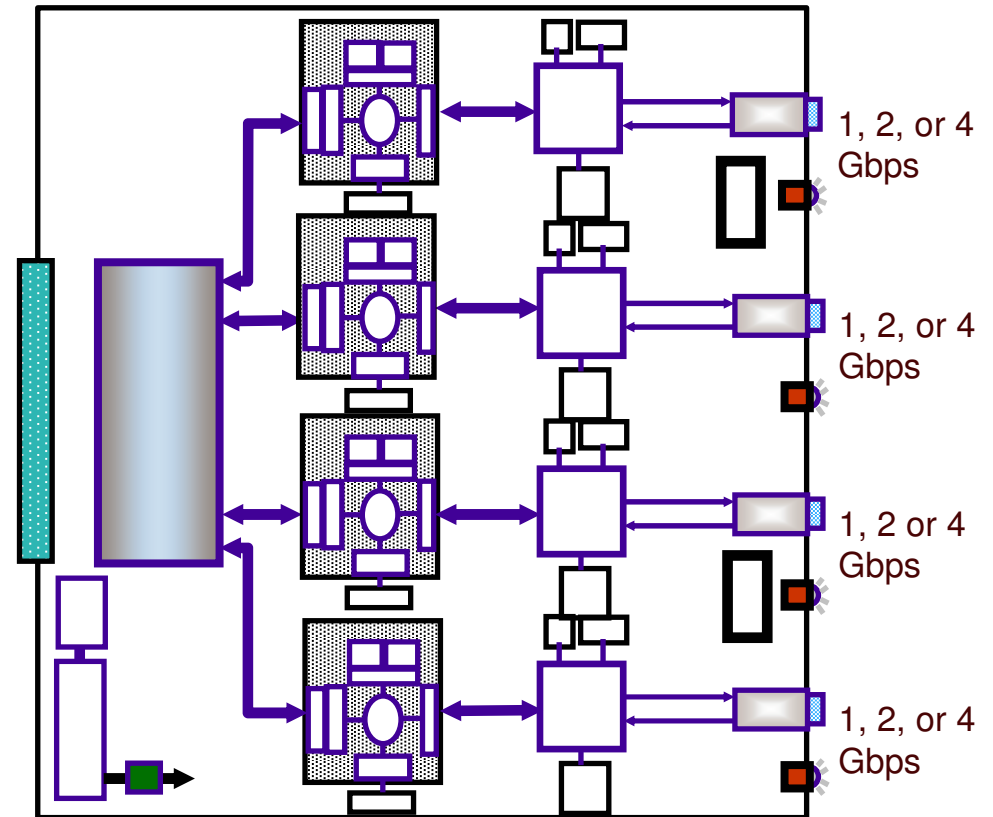
Supports all the function of the FICON Express2 feature plus:

- ▶ 4 Gbps with Auto-negotiate capability (1, 2, or 4 Gbps)
- ▶ Can be shared among LPARs, and defined as spanned
- ▶ Small Form Factor Pluggable (SFP) optics for Service / Repair
 - Concurrent repair/replace action for each SFP

Ordering

- ▶ 2 (*on BC models only*) or 4 port increments
- ▶ Intermix is not supported on a single card
- ▶ All ports must be of the same type, either LX or SX
 - LX Feature Code 3321/24 - gives 4 SFP
 - SX Feature Code 3322 - gives 4 SFP
 - LX Feature Code 3323 - gives 2 SFP

FICON Express4 4 Port card shown

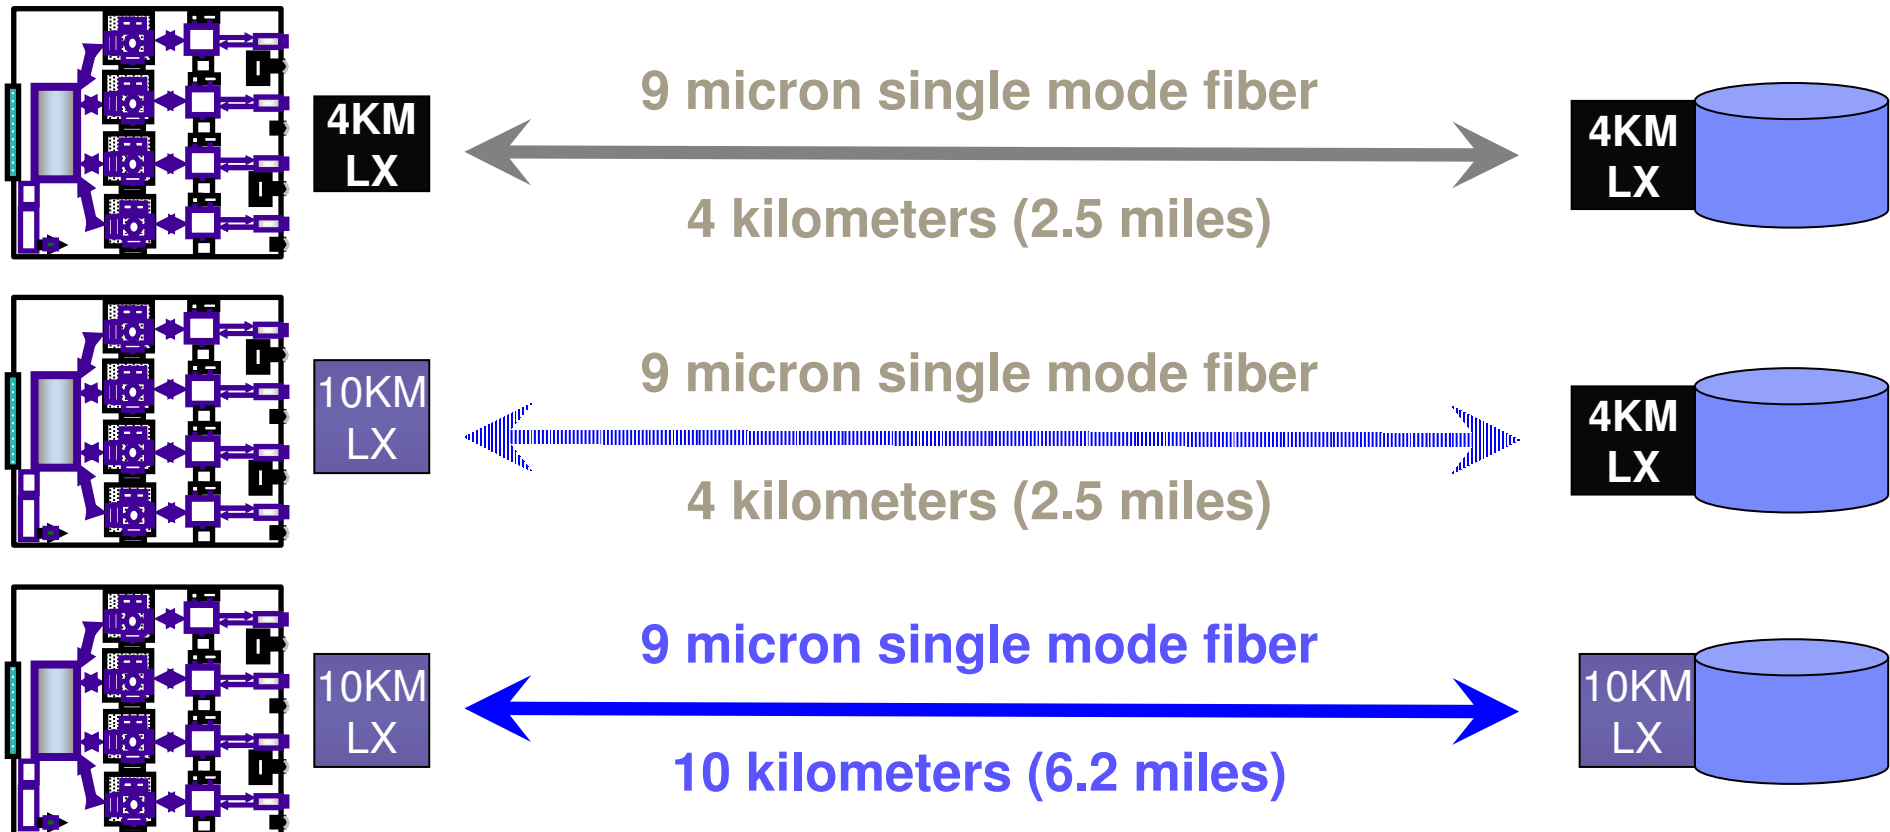


FC 3321 FICON Express4 10 KM LX
 FC 3322 FICON Express4 SX
 FC 3323 FICON Express4-2C 4KM LX (2 ports)
 FC 3324 FICON Express4 4KM LX

System z9 FICON Express4 Unrepeated Distances

IBM System z9 4Gb FICON/FCP connectivity:

- The ANSI Fibre Channel Physical Interface (FC-PI-2) standard 10 kilometer (km) and 4 kilometer transceivers when using 9 micron single mode fiber optic cabling
- IBM supports FC-PI-2 variants
- IBM supports interoperability of 10 km transceivers with 4 km transceivers provided the unrepeated distance between a 10 km transceiver and a 4 km transceiver does not exceed 4 kilometers (2.5 miles)



IBM Storage Ready for System z9 and FICON Express4

IBM System z9 and IBM storage 4 Gb FICON/FCP connectivity may help to:

- Support faster link speeds and shorter backup windows
- Enable channel and link consolidation to help simplify management and reduce the cost of the storage infrastructure
- Support easier migration to 4 Gbps bandwidth with auto-sensing links



*IBM has a full range of
Disk, SAN, Tape, Software,
& Services for System z9*

Disk

DS8000 – 4 Gbps FICON/FCP Planned 2Q06
DS6000 – 2 Gbps FICON/FCP

SAN

IBM SAN256B and SAN32B-2 FICON/FCP
IBM SAN256M (Planned for 2006) and SAN32M,
and SAN140M 4 Gbps FICON/FCP
Cisco MDS 9500 and 9216 4Gbps FICON/FCP Planned 2006

Virtualization

IBM SVC 4 Gb FCP for Linux on System z Planned 2Q06
VTS 2 Gbps FICON/FCP
TS7510 Virtualization Engine™ – 2 Gbps FCP for Linux
on System z Planned 2Q06

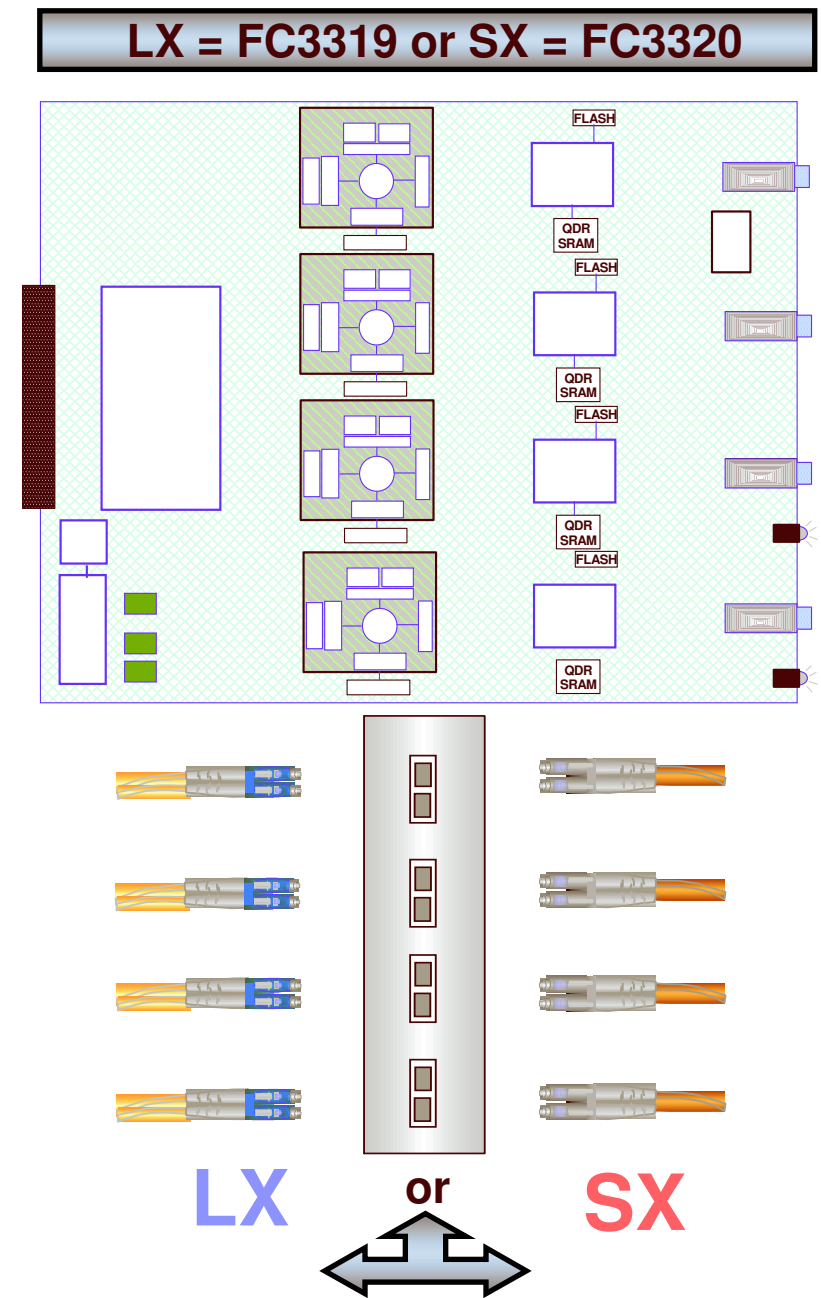
Tape

IBM TS1120 4 Gbps FCP Tape Drive
IBM TS1120 Tape Controller 4 Gbps FICON Planned 2Q06
IBM LTO Gen 3 – 4 Gbps FCP for Linux on
System z Planned 2006
IBM 3494 and 3584 Tape Libraries
IBM TS3310 Tape Library – 4 Gbps FCP for Linux
on System z Planned 2Q06

All statements regarding IBM's future direction and plans are subject to change or withdrawal without notice

System z9 FICON Express2

- **Four channels per feature of LX or SX**
- **Two CHPID types**
 - ▶ FC – native FICON and CTC
 - ▶ FCP – communication with SCSI devices
- **Connectivity options for each channel**
 - ▶ Two operating modes (no FCV support)
 - Defined on a port basis
 - FC (Fibre Channel): Native FICON and FICON CTC
 - Up to 64 Open Exchanges
 - Point to point or two director cascade
 - FCP (Fibre Channel Protocol)
 - SCSI LUN access for Linux on System z9 environments
 - Point to point or multiple director fabric
 - ▶ 1 or 2 Gbps, auto-negotiated
 - ▶ Can be shared among LPARs, and defined as a spanned channel
- **Connector – LC Duplex**
- **LX - 9 micron single mode fiber**
 - ▶ Unrepeated distance – up to 10 kilometers (6.2 miles)
 - ▶ Receiving device must also be LX
- **SX - 50 or 62.5 micron multimode fiber**
 - ▶ Variable distance with speed and fiber type
 - ▶ Receiving device must also be SX
- **Supported connectivity devices**
 - ▶ Refer to:
www.ibm.com/servers/eserver/zseries/connectivity



System z FICON features over time

Feature	Feature name	z900 12/00	z990 06/03	z9 EC 09/05	z9 BC 05/06	CHPIDs	Connector
3321	FICON Express4 10 KM LX	N/A	N/A	05/06	05/06	FC, FCP	LC Duplex
3322	FICON Express4 SX	N/A	N/A	05/06	05/06	FC, FCP	LC Duplex
3323	FICON Express4-2C 4 KM LX	N/A	N/A	N/A	05/06	FC, FCP	LC Duplex
3324	FICON Express4 4KM LX	N/A	N/A	05/06	05/06	FC, FCP	LC Duplex
3319	FICON Express2 LX	N/A	01/05	Y	05/06	FC, FCP	LC Duplex
3320	FICON Express2 SX	N/A	01/05	Y	05/06	FC, FCP	LC Duplex
2319	FICON Express LX	10/01	Y	C	C	FCV, FC, FCP	LC Duplex
2320	FICON Express SX	10/01	Y	C	C	FC, FCP	LC Duplex
2315	FICON LX	Y	N/A	N/A	N/A	FCV, FC, FCP	SC Duplex
2318	FICON SX	Y	N/A	N/A	N/A	FC, FCP	SC Duplex

LX = Long wavelength transceiver, used with LX senders and receivers and single mode fiber

SX = Short wavelength transceiver, used with SX senders and receivers and multimode fiber

C = Carry forward on an upgrade from z900 or z990 to z9 EC. z800 or z890 to z9 BC

Y = Available for ordering

System z9 Connectors and fiber optic cabling

Feature #	Feature name	Connector	Cable type
0219	ISC-3	LC Duplex	9 u SM
6155	External Time Reference (ETR)	MT-RJ	62.5 u MM
2324	ESCON channel	MT-RJ	62.5 u MM
3319/21/23/24	FICON Express2 and Express4 LX	LC Duplex	9 u SM
3320/22	FICON Express2 and Express4 SX	LC Duplex	50, 62.5 u MM
2319	FICON Express LX	LC Duplex	9 u SM
2320	FICON Express SX	LC Duplex	50, 62.5 u MM
3364	OSA-Express2 GbE LX	LC Duplex	9 u SM
3365	OSA-Express2 GbE SX	LC Duplex	50, 62.5 u MM
3366	OSA-Express2 1000BASE-T Ethernet	RJ-45	Category 5 UPT
3368	OSA-Express2 10 GbE LR	SC Duplex	9 u SM
2364	OSA-Express GbE LX	SC Duplex	9 u SM
2365	OSA-Express GbE SX	SC Duplex	50, 62.5 u MM
2366	OSA-Express Fast Ethernet	RJ-45	Category 5 UTP
1364	OSA-Express GbE LX	LC Duplex	9 u SM
1365	OSA-Express GbE SX	LC Duplex	50, 62.5 u MM
1366	OSA-Express 1000BASE-T Ethernet	RJ-45	Category 5 UTP

U = micron SM = Single mode fiber, MM = Multimode fiber, LX = Long wavelength transceiver, SX = Short wavelength transceiver, UTP = Unshielded Twisted Pair, STP = Shielded Twisted Pair, LR = Long Reach transceiver

z9 BC Model R07 I/O Features

Maximum of 240 CHPIDs, One I/O cage, 16 I/O slots

Features	Minimum # of features	Maximum # of features	Maximum connections	Increments per feature	Purchase increments
16-port ESCON	0 ⁽¹⁾	16	240 channels	16 channels 1 reserved as a spare	4 channels
FICON Express4**	0 ⁽¹⁾	16	64 channels**	4 channels**	4 channels**
FICON Express2*	0 ⁽¹⁾	16	64 channels	4 channels	4 channels
FICON Express*	0 ⁽¹⁾	16	32 channels	2 channels	2 channels
STI-3 ⁽²⁾	0	8	N/A	2 outputs	N/A
ICB-3 link	0 ⁽¹⁾	N/A	16 links ⁽³⁾	N/A	1 link
ICB-4	0 ⁽¹⁾	N/A	16 links ^{(3) (4) (6)}	N/A	1 link
ISC-3	0 ⁽¹⁾	12	48 links ⁽³⁾	4 links	1 link
OSA-Express2	0	16	32 ports ⁽⁶⁾	2 or 1 (10 GbE has 1)	2 ports/1 port
OSA-Express*	0	16	32 ports ⁽⁶⁾	2 ports	2 ports
Crypto Express2	0	4	8 PCI-X adapters	2 PCI-X adapters	2 PCI-X adapters ⁽⁵⁾

1. Minimum of one I/O feature (ESCON, FICON) or one Coupling Link (ICB, ISC-3) required.
2. Each STI-3 distribution card occupies one I/O slot (supports ICB-3s).
3. Maximum number of Coupling Links combined (ICs, ICB-3s, ICB-4s, and active ISC-3 links) cannot exceed 64 per server.
4. ICB-4s are not included in the maximum feature count for I/O slots but are included in the CHPID count.
5. Initial order of Crypto Express2 is 4 PCI-X adapters (two features). Each PCI-X adapter can be configured as either a coprocessor or an accelerator.
6. A01 has up to 8 ICB-4 links. Up to 12 OSA-Express2/ OSA-Express features.

* Carry forward on an upgrade only.

** FICON Express4-2C 4KM LX has two channels per feature

z9 BC Model S07 I/O Features

Maximum of 420 CHPIDs, One I/O cage, 28 I/O slots

Features	Minimum # of features	Maximum # of features	Maximum connections	Increments per feature	Purchase increments
16-port ESCON	0 ⁽¹⁾	28	420 channels	16 channels 1 reserved as a spare	4 channels
FICON Express4**	0 ⁽¹⁾	28	112 channels**	4 channels**	4 channels**
FICON Express2*	0 ⁽¹⁾	20	80 channels	4 channels	4 channels
FICON Express*	0 ⁽¹⁾	20	40 channels	2 channels	2 channels
STI-3 ⁽²⁾	0	8	N/A	2 outputs	N/A
ICB-3 link	0 ⁽¹⁾	N/A	16 links ⁽³⁾	N/A	1 link
ICB-4	0 ⁽¹⁾	N/A	16 links ^{(3) (4)}	N/A	1 link
ISC-3	0 ⁽¹⁾	12	48 links ⁽³⁾	4 links	1 link
OSA-Express2	0	24	48 ports	2 or 1 (10 GbE has 1)	2 ports/1 port
OSA-Express*	0	20	40 ports	2 ports	2 ports
Crypto Express2	0	8	16 PCI-X adapters	2 PCI-X adapters	2 PCI-X adapters ⁽⁵⁾

1. Minimum of one I/O feature (ESCON, FICON) or one Coupling Link (ICB, ISC-3) required.
 2. Each STI-3 distribution card occupies one I/O slot (supports ICB-3s).
 3. Maximum number of Coupling Links combined (ICs, ICB-3s, ICB-4s, and active ISC-3 links) cannot exceed 64 per server.
 4. ICB-4s are not included in the maximum feature count for I/O slots but are included in the CHPID count.
 5. Initial order of Crypto Express2 is 4 PCI-X adapters (two features). Each PCI-X adapter can be configured as either a coprocessor or an accelerator.
- * Carry forward on an upgrade only.
- ** FICON Express4-2C 4KM LX has two channels per feature

z9 EC I/O Features Summary

Maximum of 1024 CHPIDs, three I/O cages, 84 I/O slots (28 I/O slots per I/O cage)

Features	Minimum # of features	Maximum # of features	Maximum connections	Increments per feature	Purchase increments
16-port ESCON	0 ⁽¹⁾	69	1024 channels	16 channels 1 reserved as a spare	4 channels
FICON Express4	0 ⁽¹⁾	84	336 channels	4 channels	4 channels
FICON Express2*	0 ⁽¹⁾	84	336 channels	4 channels	4 channels
FICON Express*	0 ⁽¹⁾	60	120 channels	2 channels	2 channels
STI-3 ⁽²⁾	0	8	N/A	2 outputs	N/A
ICB-3 link	0 ⁽¹⁾	N/A	16 links ⁽³⁾	N/A	1 link
ICB-4	0 ⁽¹⁾	N/A	16 links ^{(3) (4)}	N/A	1 link
ISC-3	0 ⁽¹⁾	12	48 links ⁽³⁾	4 links	1 link
OSA-Express2	0	24	48 ports	2 or 1 (10 GbE has 1)	2 ports/1 port
OSA-Express*	0	24	48 ports	2 ports	2 ports
Crypto Express2	0	8	16 PCI-X adapters	2 PCI-X adapters	2 PCI-X adapters ⁽⁵⁾

1. Minimum of one I/O feature (ESCON, FICON) or one Coupling Link (ICB, ISC-3) required.

2. Each STI-3 distribution card occupies one I/O slot (supports ICB-3s).

3. Maximum number of Coupling Links combined (ICs, ICB-3s, ICB-4s, and active ISC-3 links) cannot exceed 64 per server.

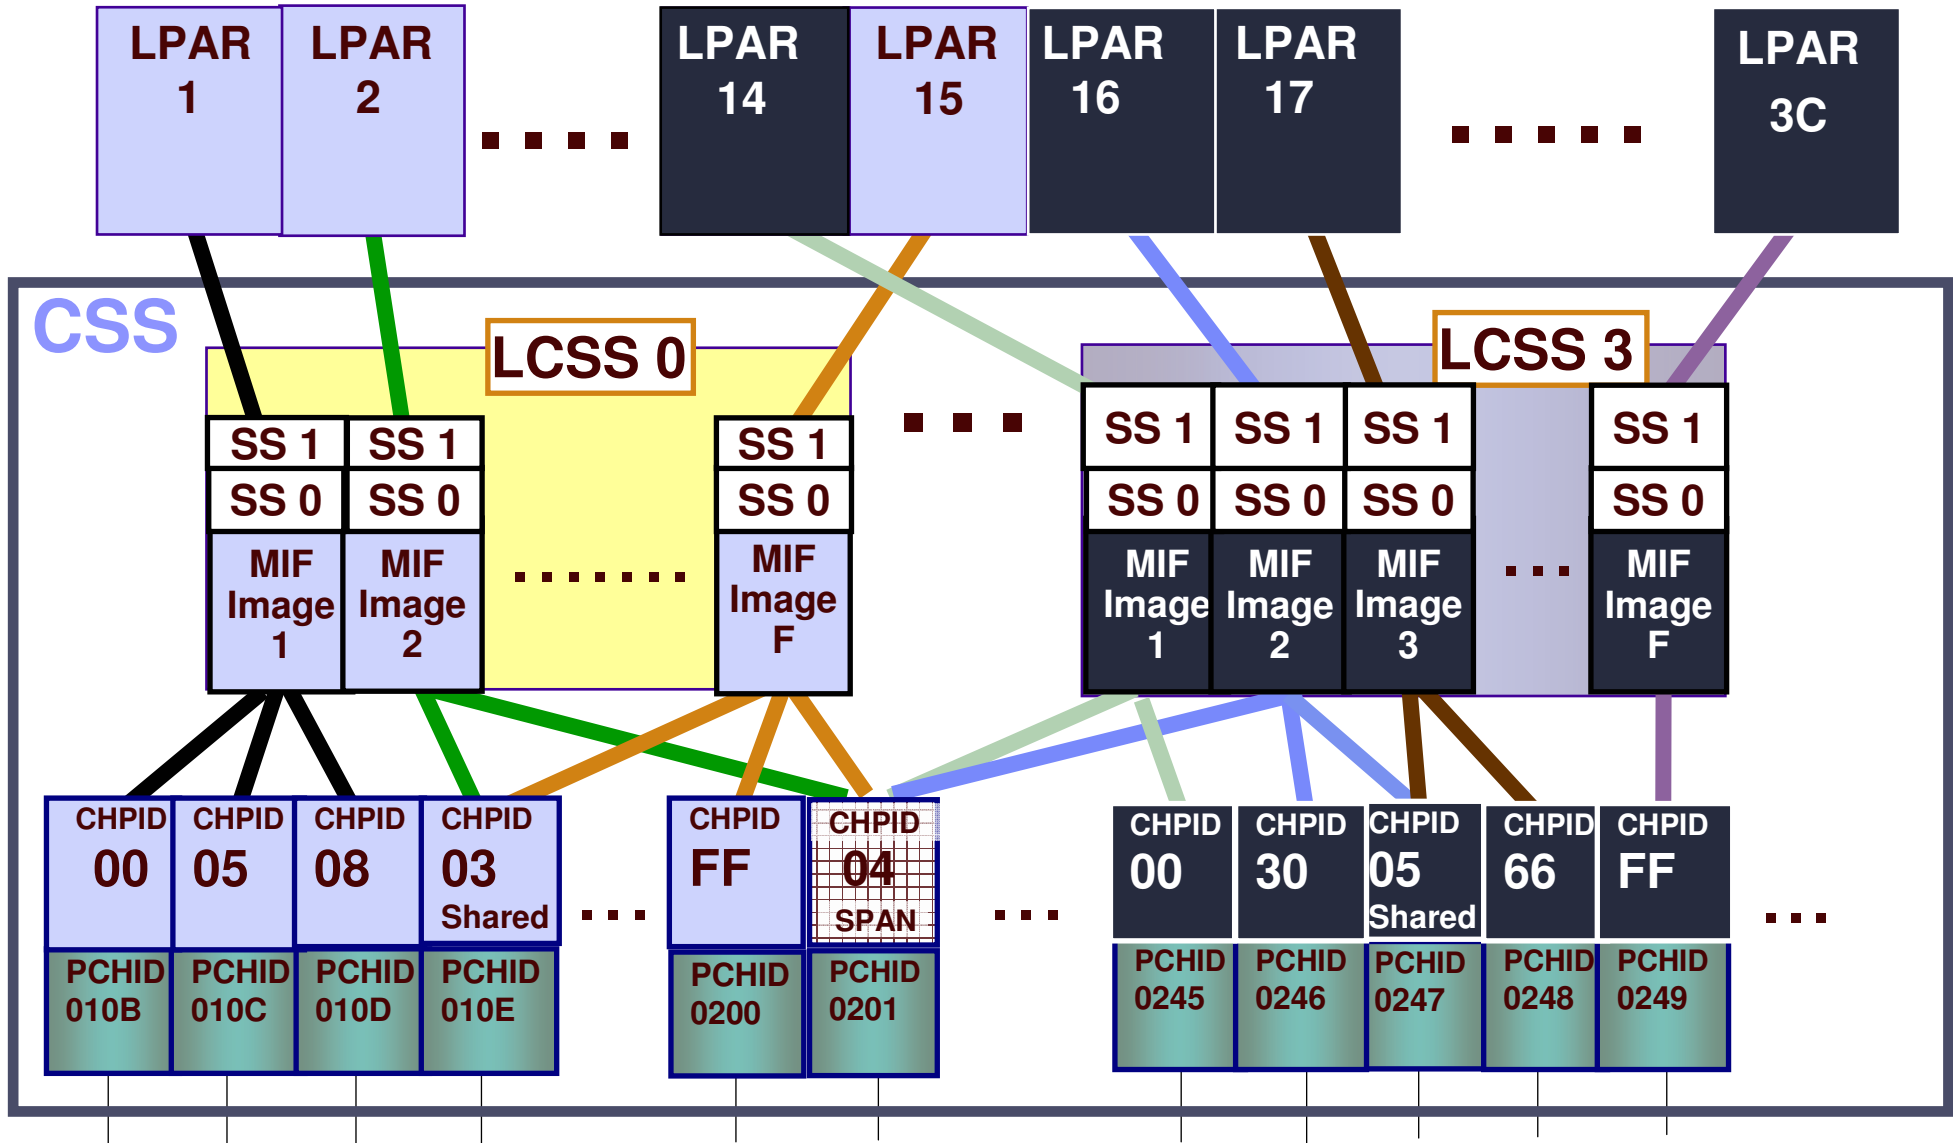
4. ICB-4s are not included in the maximum feature count for I/O slots but are included in the CHPID count.

5. Initial order of Crypto Express2 is 4 PCI-X adapters (two features). Each PCI-X adapter can be configured as a coprocessor or an accelerator.

* Available only when carried forward on an upgrade.

Note: There is a maximum of 64 ESCON features (960 active channels) and a maximum of 64 FICON feature (256 channels) Model S08.

System z9 Connectivity Summary



System z9 MIDAW

Designed to enhance performance for FICON Express 2 and 4

- **Can improve FICON performance for**
 - Extended Format (EF) data sets – including DB2 and VSAM
- **Can improve channel utilization and can significantly improve I/O response times**
 - Internal IBM DB2 Table Scan tests(*) with the z9 EC, FICON Express2 and the IBM TotalStorage® DS8000 control unit comparing MIDAW facility configurations to pre-MIDAW configurations showed:
 - 36% to 58% reduction in response times
 - 35% to 56% reduction in channel busy
 - 56% to 126% improvement in I/O throughput
- **Supported on z/OS 1.6 and above and corresponding supporting devices**

Modified Indirect Data Address Word (MIDAW) Facility

■ **Modified Indirect Address Word facility:**

- ▶ The MIDAW facility is a new CCW-indirect-data-address word facility being added to z/Architecture to coexist with the current IDAW facility
- ▶ Both MIDAW and IDAW facilities offer, for FICON and ESCON channels, alternatives to using CCW data chaining in channel programs
- ▶ Both facilities are designed to reduce channel, director, and control unit overhead by reducing the number of CCWs and frames processed
- ▶ The MIDAW facility is usable in certain case where the IDAW facility is not because it does not have IDAW boundary and data length restrictions
- ▶ Supported on z/OS 1.6 and higher

■ **Designed to:**

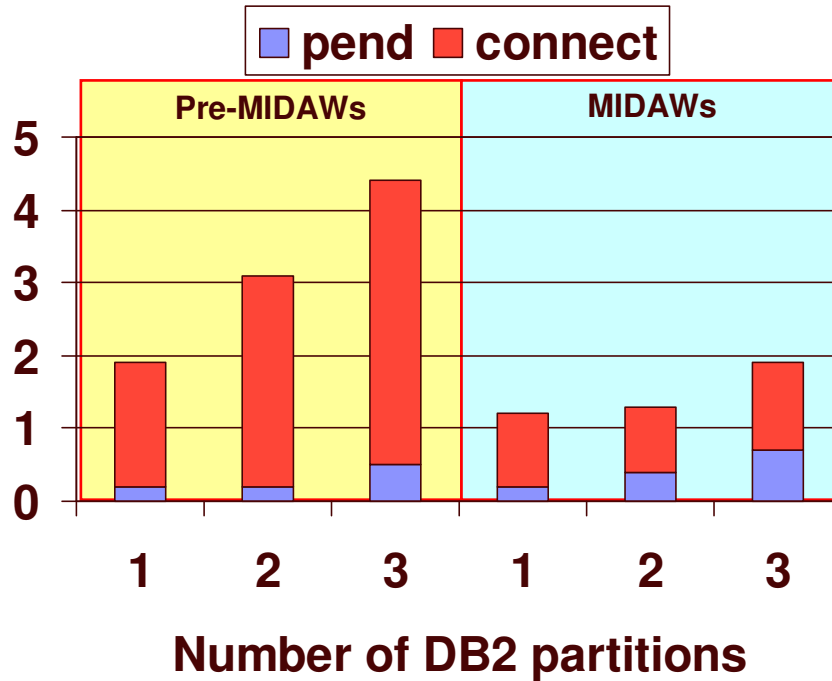
- ▶ Be compatible with existing IBM and non IBM disk control units (Note: non IBM storage devices will require support from their vendors and should be contacted as part of the installation systems assurance process)
- ▶ Decrease response time for exploiting I/O
- ▶ Increase the number of I/O operations per second that can be processed and thus move more data per second, especially on faster FICON channels

■ **Applications that may benefit include:**

- ▶ DB2, VSAM, Partitioned Data Set Extended (PDSE), Hierarchical File System (HFS), z/OS File System (zFS), and other datasets exploiting striping and compression

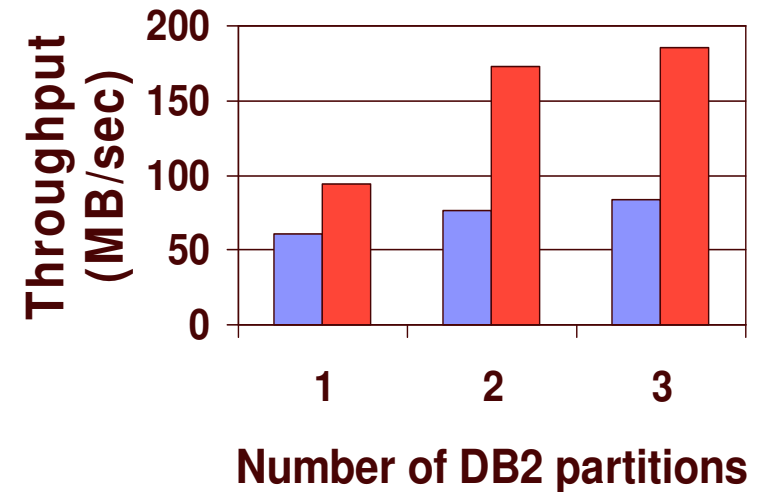
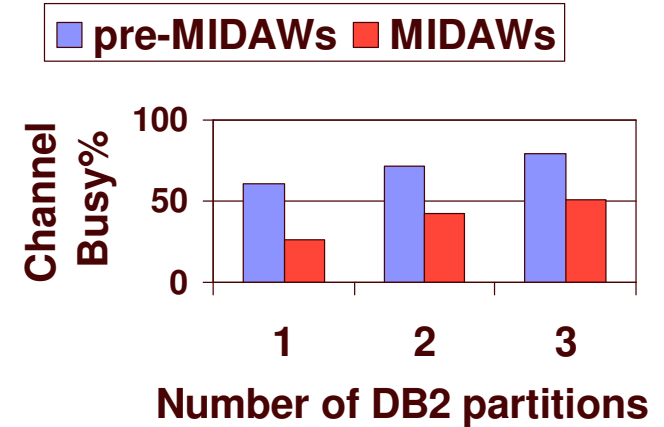
Parallel DB2 Table Scan, EF 4K (single channel)

I/O Response Time (sec)



Configuration:
 MIDAW : z/OS 1.7
 Pre-MIDAW: z/OS 1.4

 DB2 for z/OS Version 8
 4000 byte row size
 System z9
 FICON Express2
 2 Gbit/sec link
 DS8000 control unit



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

System z9 Multiple Subchannel Sets

■ Description

- ▶ Only one set of 63K subchannels are available with z990, z890, z900, z800
- ▶ For System z9, two subchannel sets are now available per LCSS, enabling a total of 63.75K subchannels in set-0 and adding 64K-1 subchannels in set-1
- ▶ Multiple subchannel sets provides growth for I/O device configuration

■ Operating System Requirements

- ▶ z/OS 1.7 and later
- ▶ z/OS will only allow Parallel Access Volume Alias (PAV-alias) devices in the second subchannel set.
- ▶ Subchannels for any other devices not allowed in subchannel Set-1

System z9 – Increased Subchannels to 63.75K

- **zSeries can address a maximum of 64K subchannels**
- **1024 (1K) of these previously reserved for system use**
- **IBM has made available 768 of these 1K reserved subchannels for customer use**
- **Can increase the storage attachment capability of the System z9 processor**
- **For example, in the largest case, using 3390 volumes with 54 GB/volume and 768 additional volumes, you could have 41 Terabytes of additional disk storage addressability" (i.e., 54 GB/volume * 768 volumes = 41 TB).**
- **The IBM TotalStorage DS8000 can be defined to attach 63.75K unit addresses**
- **z/OS 1.4 (with applicable PTFs) and later**
- **z/VM V4.4 and V5.1 (with HCD PTF) and V5.2**



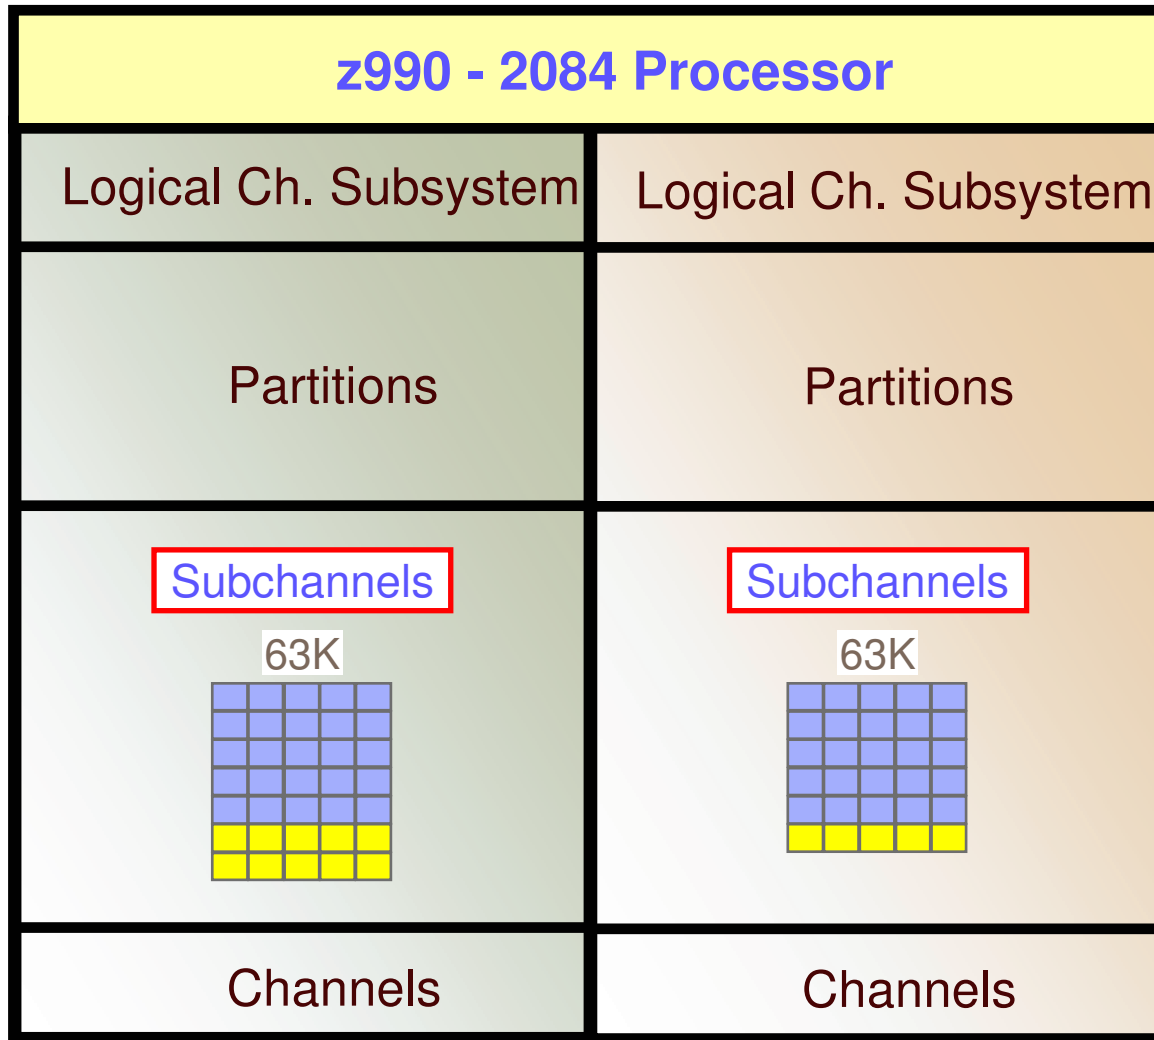
System z9 vs. z990, z890, z900, z800

zSeries Processor Type	MCSS (LCSSs)	Logical Partitions	MSS	Subchannels Per LCSS (Per LPAR)	Base Devices (per CSS)	SPANNED Channels
2064 z900	No (1 CSS)	15	No	63 K	64512	N/A
2066 z800	No (1 CSS)	15	No	63 K	64512	N/A
2084 z990	Yes 4 x LCSS	30 (Max 15 for a CSS)	No	63 K	64512	Yes
2086 z890	Yes 2 x LCSS	30 (Max 15 for a CSS)	No	63 K	64512	Yes
2094 z9 EC	Yes 4 x LCSS	60 (Max 15 for a LCSS)	1 or 2 per LCSS (Up to 8 total)	127.75 K	65280 (in MSS0)	Yes
2096 z9 BC	Yes 2 x LCSS	30 (Max 15 for a LCSS)	1 or 2 per LCSS (Up to 4 total)	127.75 K	65280 (in MSS0)	Yes

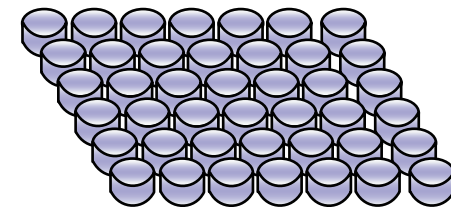


z990 – I/O Configuration Support

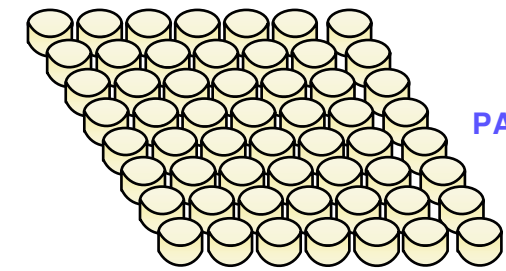
Each 2084 Logical Channel Subsystem has a set of Subchannels



zSeries 2084
DASD Configurations
63K Device



Bases

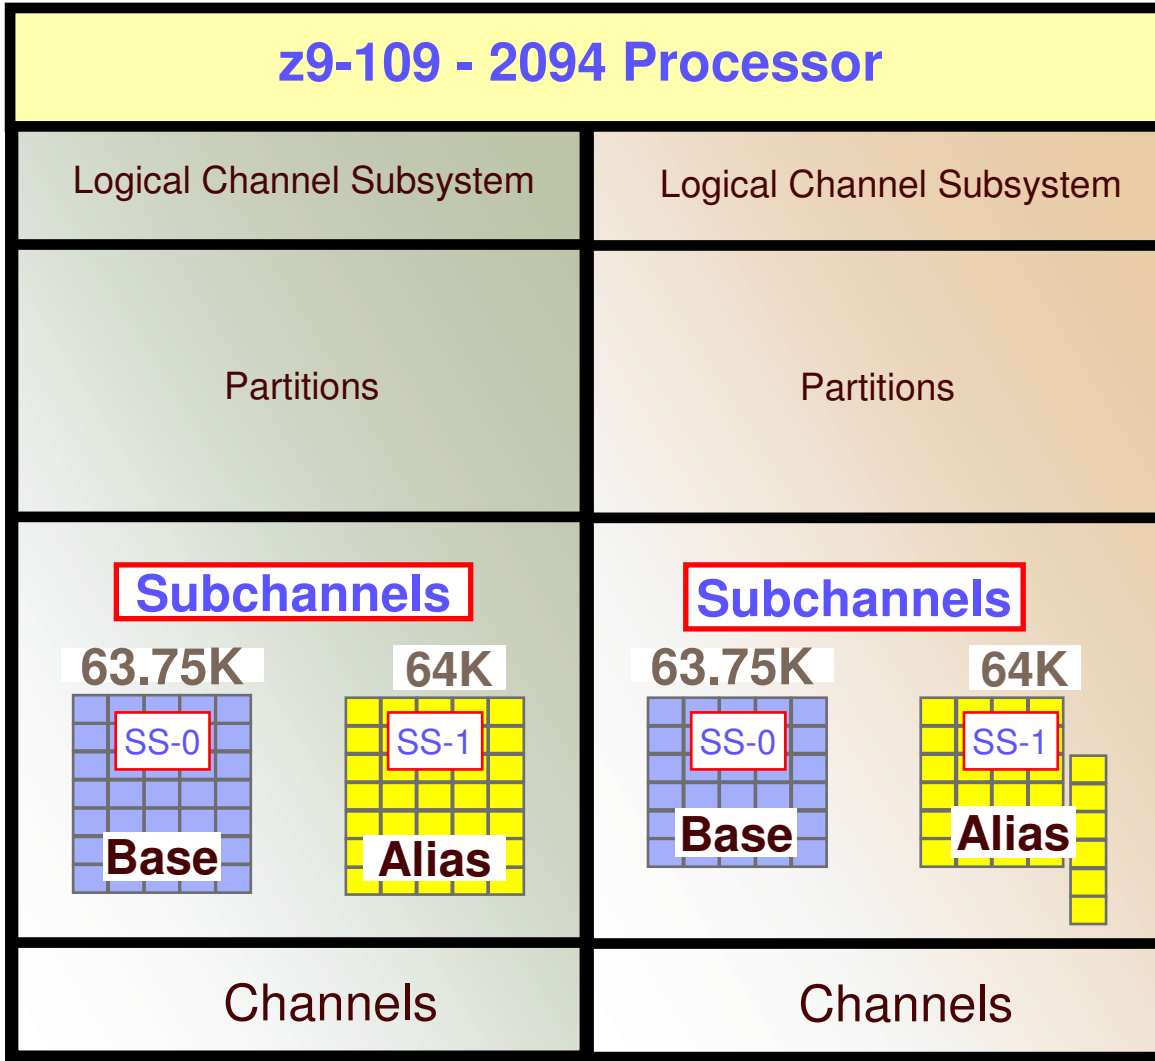


PAV

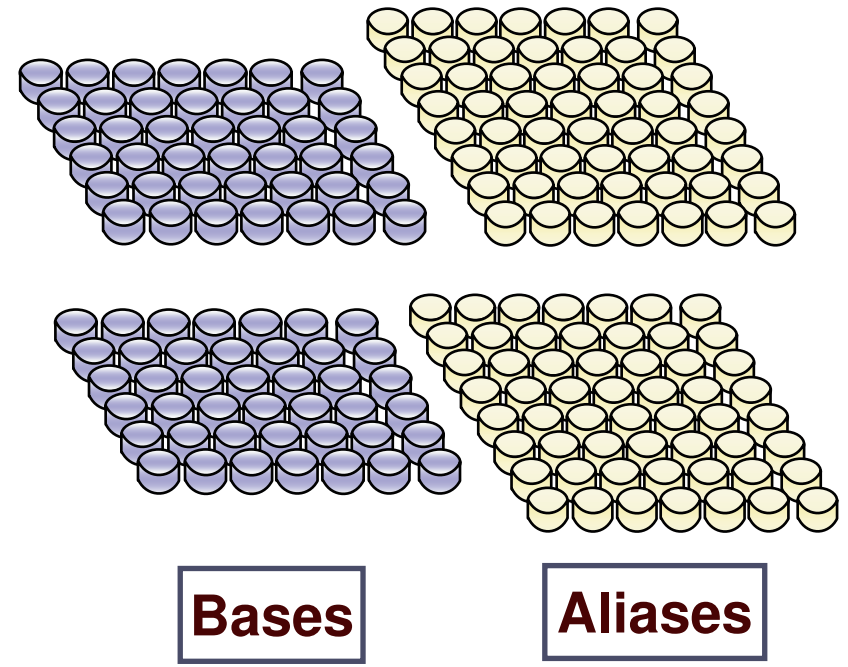
Aliases

z9 – Multiple Subchannel Sets per LCSS

Each 2094 Logical Channel Subsystem can have up to two sets of Subchannels




System z9 2094
DASD Configurations
Increased to 128K Devices




- More base volumes, more data in an LPAR or Sysplex
- More alias volumes, better performance

System z9 Multiple Subchannel Sets Summary



Set 0 – Up to 65,280 subchannels

- Always present if the LCSS is defined
- Any device for allocation
- Reference in JCL
- UCB Services
- Messages
- Commands



Set 1 – Up to 65,535 subchannels

- Optionally defined
- PAV alias devices only
- Display Commands
- Limited messages

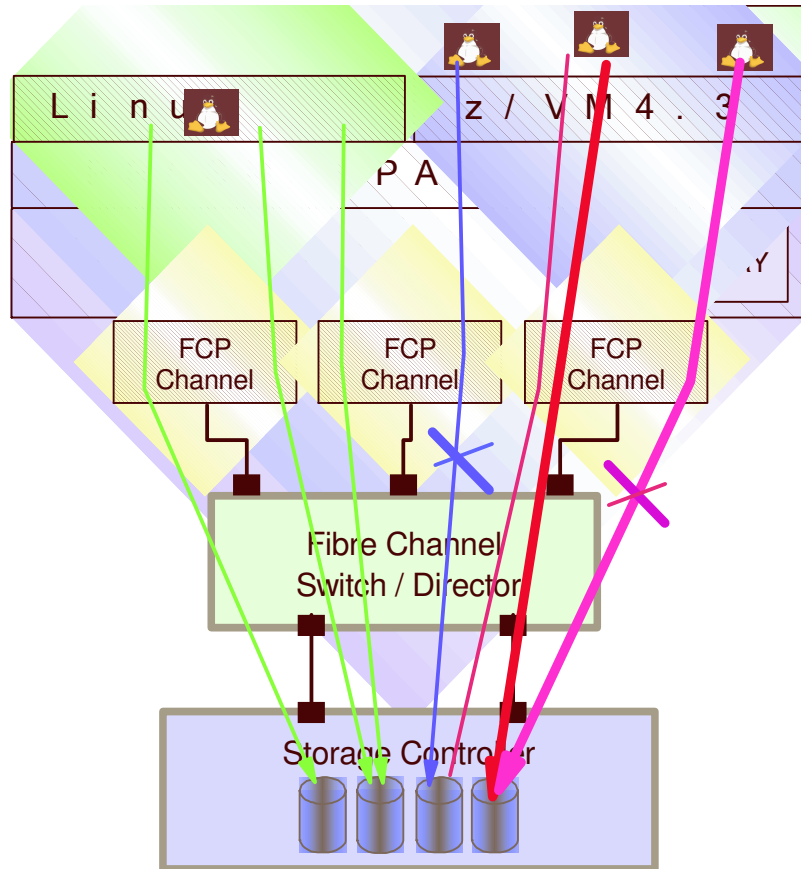
Designed to be compatible with existing storage CUs that support PAV

N-Port ID Virtualization (NPIV)

- **Multiple operating system images can now concurrently access the same or different SAN attached devices (LUNs) via a single, shared FCP Channel.**
 - ▶ Can improve channel utilization
 - ▶ Less hardware required
 - ▶ Helps reduce complexity of physical I/O connectivity.
- **NPIV provides each OS instance a set of virtual FCP Channel ports.**
 - ▶ Each virtual FCP Channel port has its own unique SAN identity (WWPN)
 - ▶ A virtual FCP Channel port behaves like a physical FCP Channel port
- **Each virtual FCP Channel port identifier (WWPN) can be used in device-level access control in storage controllers (LUN masking), as well as switch-level zoning.**

Effects of FCP Channel Virtualization on Access Control

1 2



LUN 1

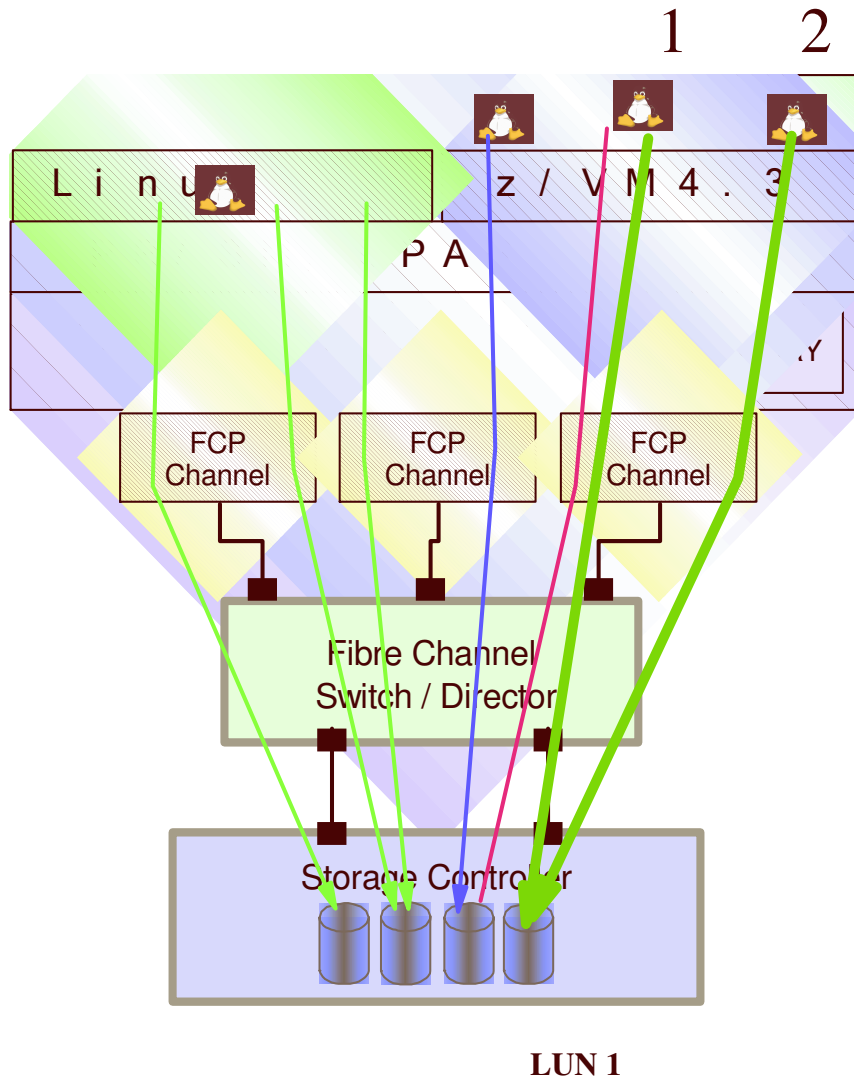
Without FCP Channel Virtualization (prior to z9-109)

FCP Channel ports can be **shared** among multiple operating system instances, with the following **restrictions**:

- SAN attached devices only recognize the single port name (WWPN) of the FCP Channel port.
- Devices (eg, **LUN 1**) can not distinguish between OS instances sharing a single FCP Channel port.
- As a consequence, several SCSI commands and SAN functions may not work properly, such as
 - SCSI reservations
 - SAN-based Access Control (LUN masking, zoning)

-FCP Channel thus prohibits shared access to the same SCSI LUN

Effects of FCP Channel Virtualization on Access Control



With FCP Channel Virtualization :

FCP Channel Virtualization provides virtual FCP Channel ports to operating systems

- Each **virtual** FCP Channel port has its own SAN identity (WWPN)
- Devices attached (eg, **LUN 1**) to the Fibre Channel fabric **can now distinguish between operating system instances** that share a physical FCP Channel.
- FCP Channel no longer prohibits shared access to the same SCSI LUN.**

Announced z/VM Support of N_Port ID Virtualization (NPIV)

- **NPIV is available with Linux on System z9 in a logical partition or as a guest of z/VM V4.4, V5.1 and later for SCSI disks accessed via dedicated subchannels and for guest IPL**
- **For guest use of NPIV, z/VM V4.4 and z/VM V5.1 provide support transparently; i.e., no PTF is required. z/VM V5.2 continues this support when released**
- **z/VM V5.1 and later provide NPIV support for VM-system use of SCSI disks (including emulated-FBA minidisks for guests). z/VM V5.1 requires a PTF to properly handle the case of a Fibre Channel switch not being able to assign a new N_Port ID when one is requested (due to the switch's capacity being exceeded).**
- **A z/VM V5.2 PTF is planned to be provided in May 2006 to add support allowing VM users and VM guest operating systems to obtain the worldwide port name(s) (WWPNs) being used in a virtual machine. The QUERY command will be enhanced for VM users, and virtualization of a machine function will be enhanced for VM guests.**

Notes:

- 1) IBM is working with its distribution partners to provide this function in future Linux for System z9 and zSeries distribution releases, or service updates
- 2) z/VM V5.1 cannot be installed from DVD to SCSI disks when NPIV is enabled.

Program-Directed Re-IPL

- **Program-Directed Re-IPL enables an operating system running natively in a logical partition to re-IPL itself.**
- **This function is supported both for**
 - ▶ SCSI and ECKD™ devices
- **In order to initiate a Re-IPL, an operating system can determine**
 - ▶ *How* it was loaded (i.e., via Channel Control Words (CCWs)) or via SCSI-type IPL, and
 - ▶ *From where* it was loaded in case of SCSI-type IPL
 - ▶ (World Wide Port Name (WWPN) and Logical Unit Number (LUN) of the load device).
- **Using these parameters, the program can subsequently call a firmware function (SCLP call) to request that it be reloaded from the same load device, using the same load parameters.**
- **Software Support**
 - ▶ Linux for System z9 and zSeries support for Program-Directed Re-IPL is expected to become available from IBM's Linux distribution partners at a later time.

Note:

z/VM already supports an interface that allows a program running as a guest under z/VM to re-IPL itself.



IBM System z9 109 Parallel Sysplex

Additional Information:

Parallel Sysplex Web Site:

www.ibm.com/servers/eserver/zseries/pso



**ENABLING BUSINESS.
A THROUGH Z.**

Coupling Facility Control Code (CFCC) for System z9

- **Structure size (and CF storage) increase when moving from**
 - ▶ CF Level 11 to CF Level 14 (or 12/13)
 - ▶ CF Level 12 or CF Level 13 to CF Level 14

- **CFCC level 14 does not change on a System z9**

- **Recommendations have not changed for System z9 for CP designations**
 - ▶ Still Workload volume dependent
 - ▶ Recommendation still to use dedicated CPs for production CF/ICFs

System z9 CFCC Level 14

■ Function and Potential Benefit

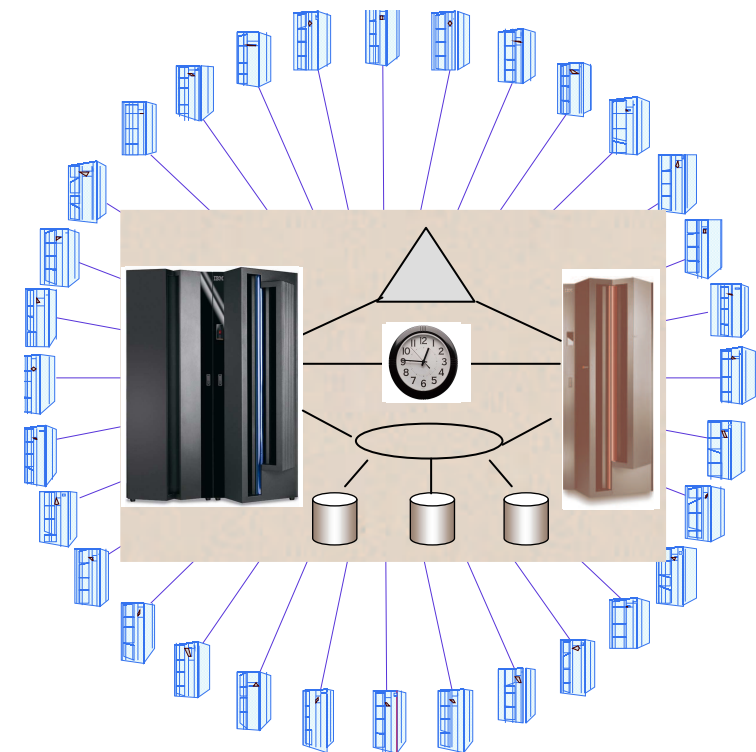
- ▶ Contains improvements to the CF dispatcher and internal serialization mechanisms designed to better manage coupled workloads

■ Requirements and Support

- ▶ z/OS 1.4 and higher
 - Optional APAR fix OA08742 to allow sysplex connectors to request structure allocation in a Level 14 Coupling Facility
- ▶ z/VM 4.4 and higher for virtual CF support

■ CF Storage Sizing with CFCC level 14

- ▶ May increase storage requirements
- ▶ Use CFSIZER tool to determine:
www.ibm.com/servers/eserver/zseries/cfsizer/



1 to 32 Systems

zSeries continues to meet the requirements for advanced clustering

CF Links for System z9 and connectivity to G5/G6

Having a System z9 CF - introduces coexistence issues, if there are any G5/G6 CECs participating in that sysplex....and vice versa....

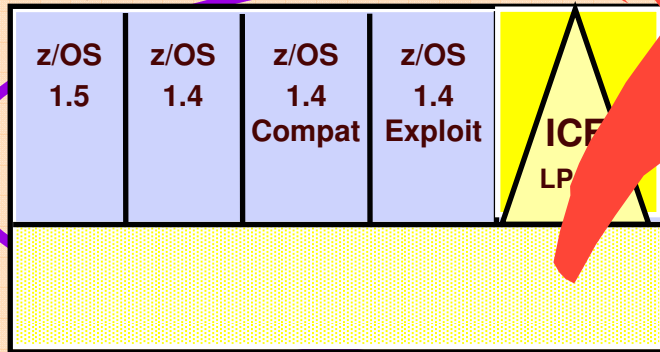
Need to start thinking about removing any existing G5/G6 CEC or CF images in the sysplex

- **If you have a System z9 CF image,**
 - ▶ You cannot connect that CF image to any G5/G6 z/OS senders
 - ▶ You cannot connect that CF image to any G5/G6 CF when using duplexing
- **If you have a G5/G6 CF image,**
 - ▶ You cannot connect that CF to any z9-EC/BC z/OS senders (or for duplexing, to a z9 CF)
- **If you have G5/G6 images AND z9-109 images**
 - ▶ They can coexist in the same sysplex as long as their CFs are not on G5/G6 or z9 EC/BC (e.g., on a z990/z890/z900/z800),
 - ▶ These "intermediate" CFs can connect to both the old (G5/G6) and the new z9 EC/BC systems, so they can provide a "bridge" that allows them to coexist in a sysplex as described above.

System z9 & Parallel Sysplex coexistence with a G5/G6 CF

Connecting a G5/G6 CF image to any z9-109 z/OS senders is NOT supported

G5, G6

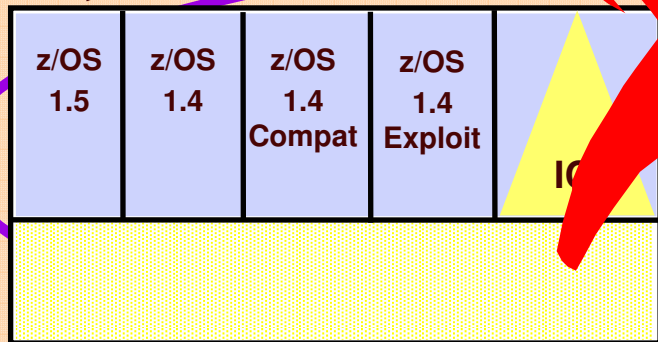


z9-109

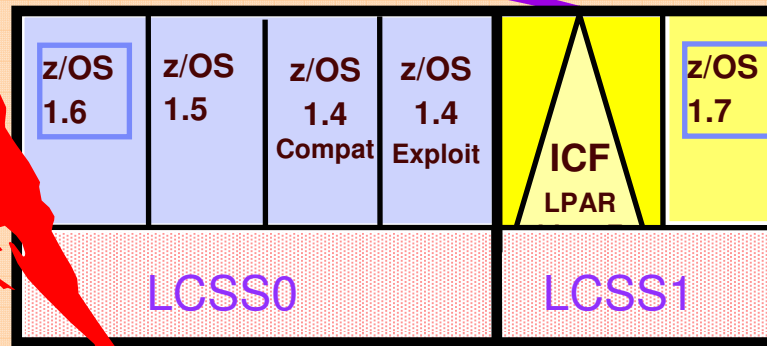


Connecting a G6/G6 CF image for duplicating to a z9-109 CF is NOT supported

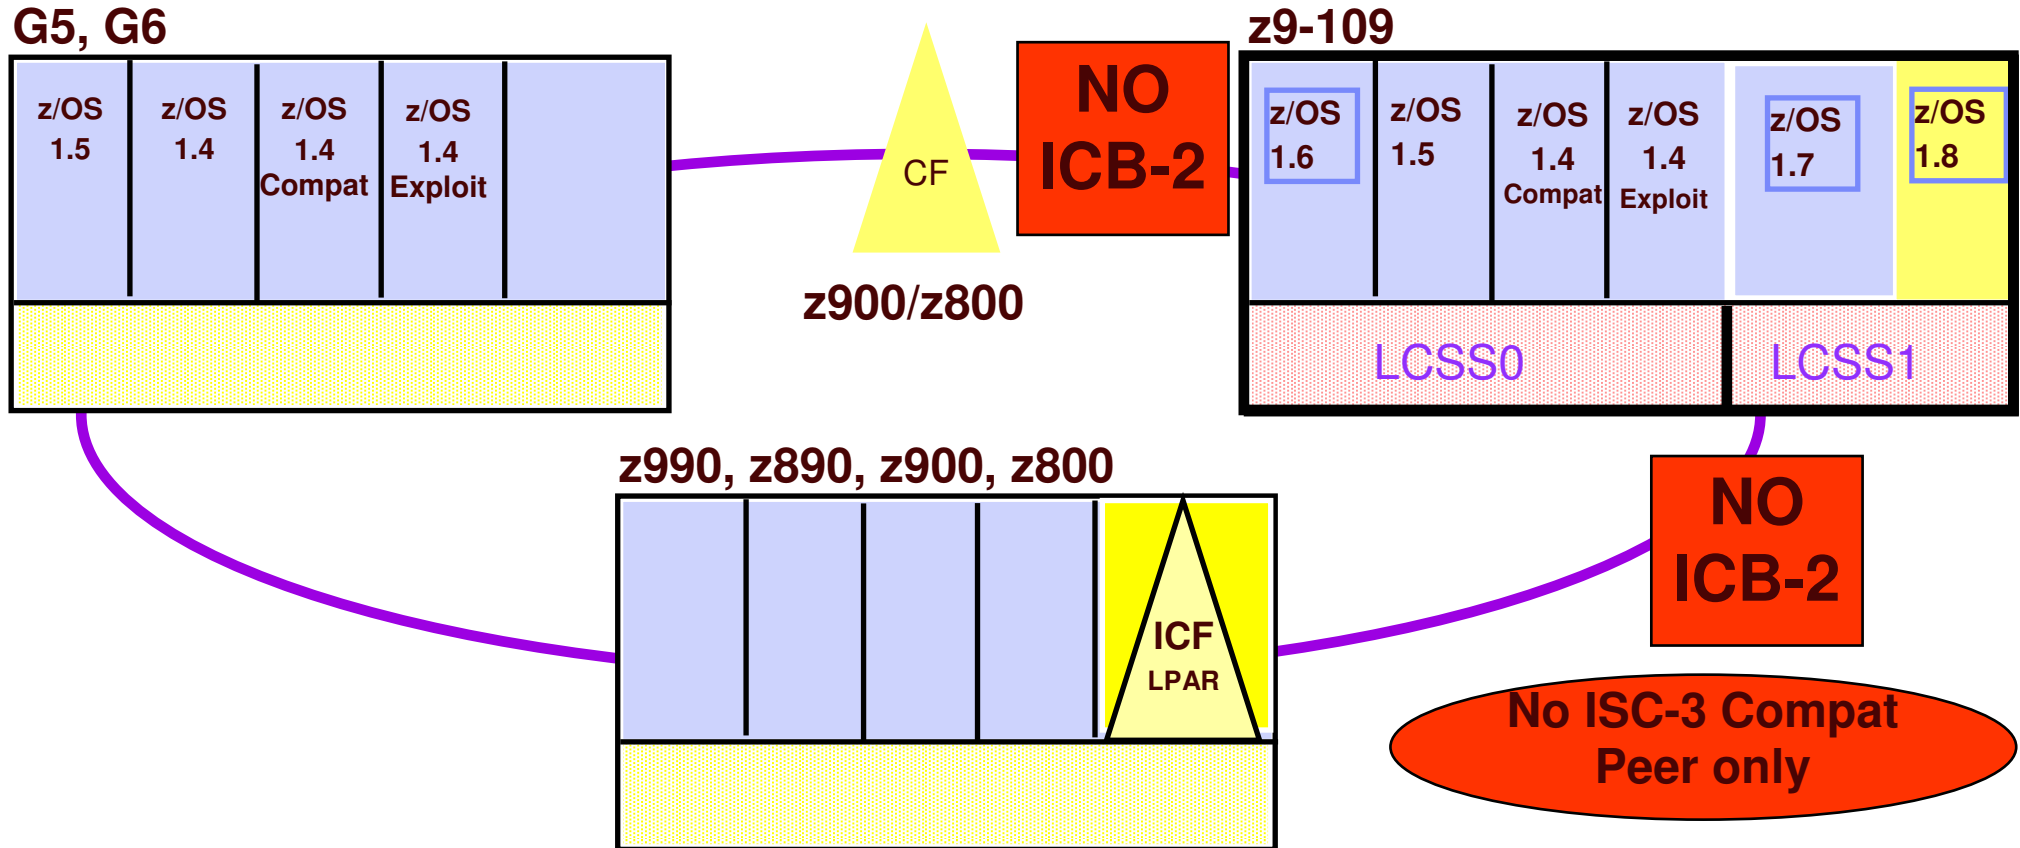
G5, G6



z9-109



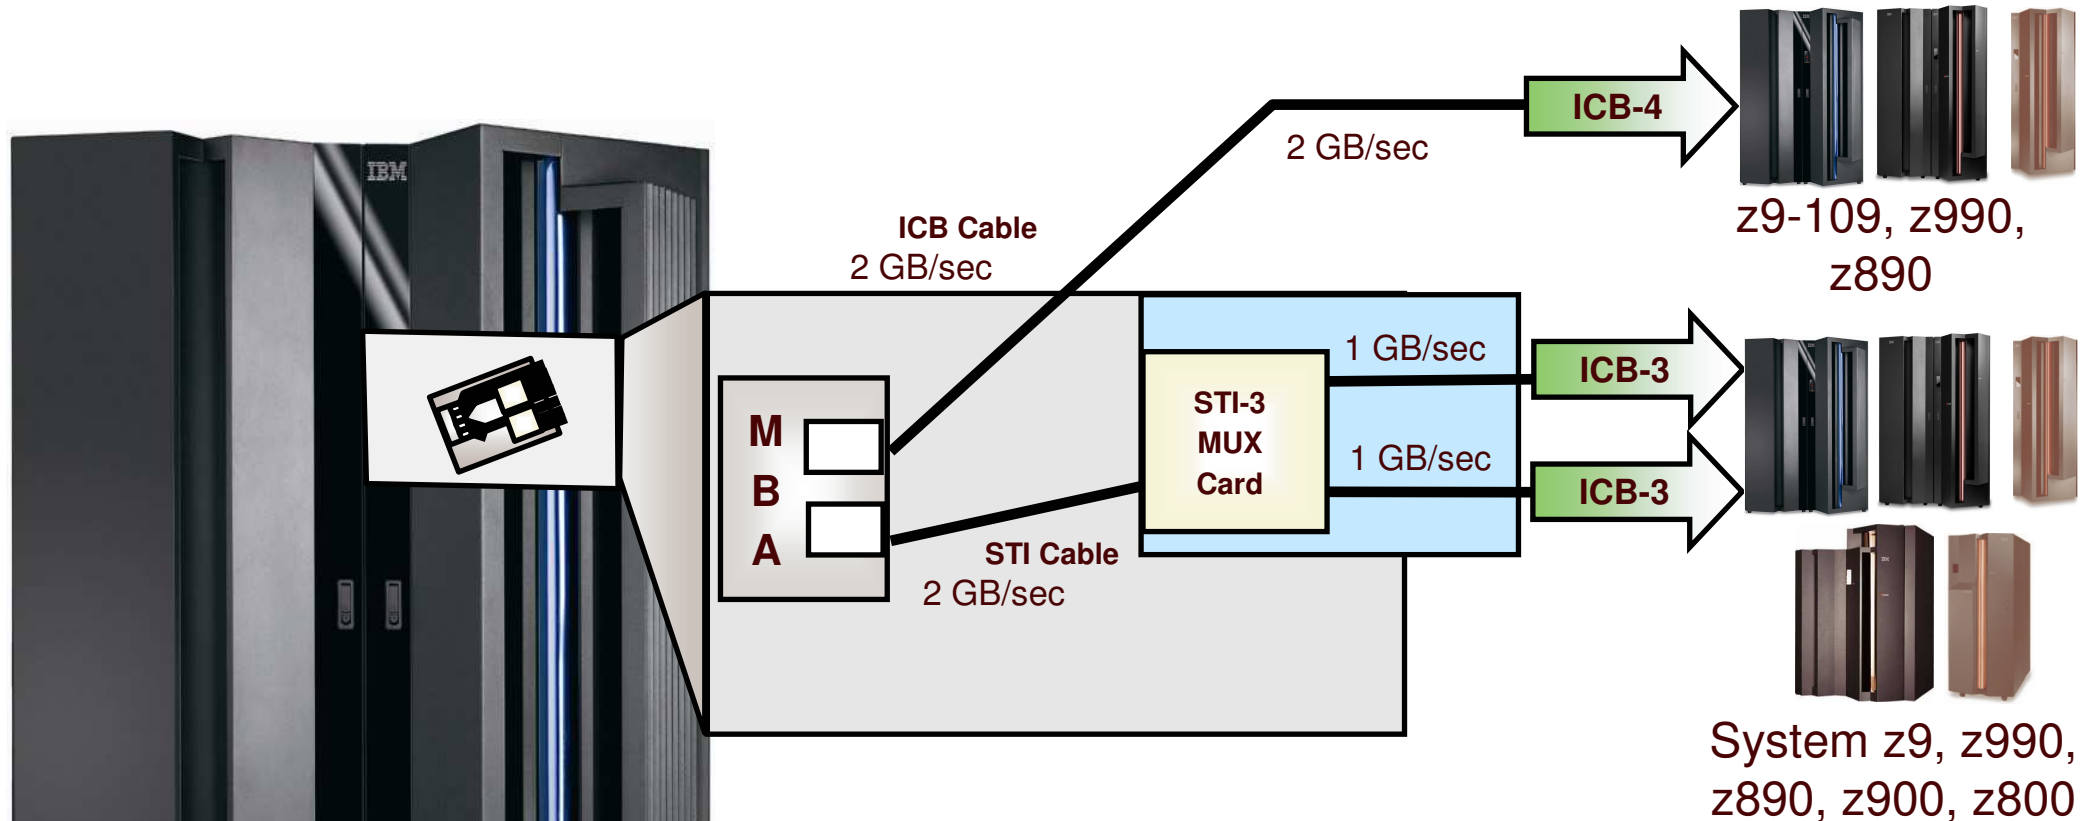
Parallel Sysplex coexistence with G5/G6 & System z9 Senders



G5/G6 z/OS senders CAN be in the same sysplex with z9-109 senders if they connect with CFs on z990, z890, z900 or z800 servers!

- These "intermediate" CFs can connect to both the old (G5/G6) and the z9-109.
- They can provide a "bridge" that allows them to coexist in a sysplex


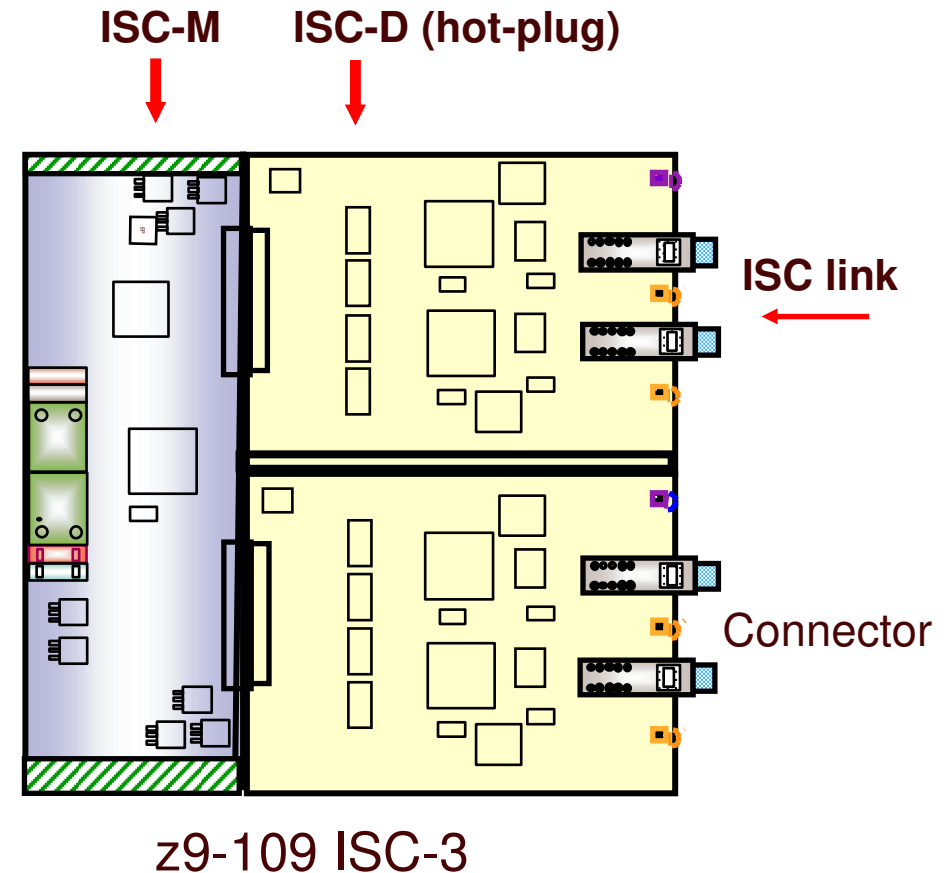
System z9 and ICB Coupling Links



- STI multiplexor cards take I/O card slots
- No I/O slot needed for ICB-4 on z9-109
- ICB Cables – different types for ICB-3 and ICB-4
 - 10 meter copper wire cables

System z9 Fiber Optic Coupling Links

- **InterSystem Channel-3**
 - ▶ 3rd generation Coupling Link
 - ▶ ISC-3 links ordered in increments of one
 - ▶ Activated links balanced across features
- **Peer mode only**
 - ▶ Peer Mode (2 Gigabits per second - Gbps)
 - ▶ Connects to zSeries only
 - ▶ No connectivity to any 9672 or 9674
- **ISC-3 FCs: 0217(ISC-M), 0218 (ISC-D / ISC link)**
 - ▶ Activate link - FC 0219
 - ▶ Four ports per ISC-M (two ports per ISC-D)
 - ▶ Supports 9 micron single mode fiber
- **Up to 48 links**



LC Duplex SM



System z9 CF Link Connectivity

Connectivity Options	System z9 ISC-3	System z9 ICB-3	System z9 ICB-4
z900/z800 ISC-3	2 Gb/s Peer Mode	N/A	N/A
z9/z990/z890 ISC-3	2 Gb/s Peer Mode	N/A	N/A
z900/z800 ICB-3	N/A	1 GB/s Peer Mode	N/A
z9/z990/z890 ICB-3	N/A	1 GB/s Peer Mode Recommendation: Use ICB-4	N/A
z9/z990/z890 ICB-4	N/A	N/A	2 GB/s Peer Mode

RPQ 8P2197- 1 Gb/sec in Peer mode and supports distance of up to 20 km unrepeated instead of 10 km
 Coupling Connectivity to 9672s and any system with ICB-2 is not supported
 Compatibility mode (sender/receiver) ISC-3 not supported

Configuring CF Links

Server	IC	ICB-4	ICB-3	ICB	ISC-3	Max # Links
z800	32	-	5 6 (OCF)	-	24	26 + 32
z900-100 CF	32	-	16	16	32 42 w/ RPQ	64
z900	32	-	16	8 16 w/ RPQ	32	64
z890	32	8	16	-	48	64
z990	32	16	16	8	48	64
System z9	32	16	16	-	48 Peer Mode Only	64



IBM System z9 Security & Cryptography



ENABLING BUSINESS.
A THROUGH Z.

System z9 EC Security Summary

- **z9 EC PR/SM – Common Criteria EAL5 Certification, March 2006**
- **Crypto Express2 feature**
 - ▶ Support high levels of security for demanding applications – fully programmable
 - ▶ Designed to meet FIPS 140-2 level 4 certification
 - ▶ Offers high-scale performance for SSL transactions
 - ▶ Integration using ICSF
- **Trusted Key Entry optional smart card reader**
 - ▶ Smart Cards – Certified to meet FIPS 140-2 Level 2
- **Common Criteria (ISO/IEC 15408) Evaluation Assurance Levels Reference: <http://niap.nist.gov/cc-scheme/>**
 - ▶ z/OS V1.7 – Controlled Access Protection Profile (CAPP) EAL4+ and Labeled Security Protection Profile (LSPP) EAL4+, February 2006
 - ▶ z/VM V5.1 with RACF® for z/VM – Controlled Access Protection Profile (CAPP) EAL3+ and the Labeled Security Protection Profile (LSPP) EAL3+, October 2005
 - ▶ Novell SUSE SLES 9 and Red Hat RHEL 4 – Controlled Access Protection Profile (CAPP) EAL4+



EAL5: Evaluation Level 5
CAPP: Controlled Access Protection Profile (CAPP)
LSPPS: Labeled Security Protection Profile (LSPP)

System z Cryptographic Technology

- Continue to provide flexible Secure Sockets Layer (SSL) acceleration
- Continue to provide competitive symmetric performance in a security-rich environment
- Provides integration of Crypto features via ICSF
- Focus on required certifications and open standards
- Continue to improve performance
 - ▶ Each Crypto Express2 feature on a System z9, with both adapters configured as accelerators, is designed to provide up to 6000* SSL handshakes per second

z900/z800 – Dec. 2000/ May 2002
 2 Chips on CEC Board -
 CMOS7s+ PCICC/PCICA (10/01)

G6 – June 1999
 2 Chips on Processor
 MCM - CMOS5x +
 PCICC (6/99)



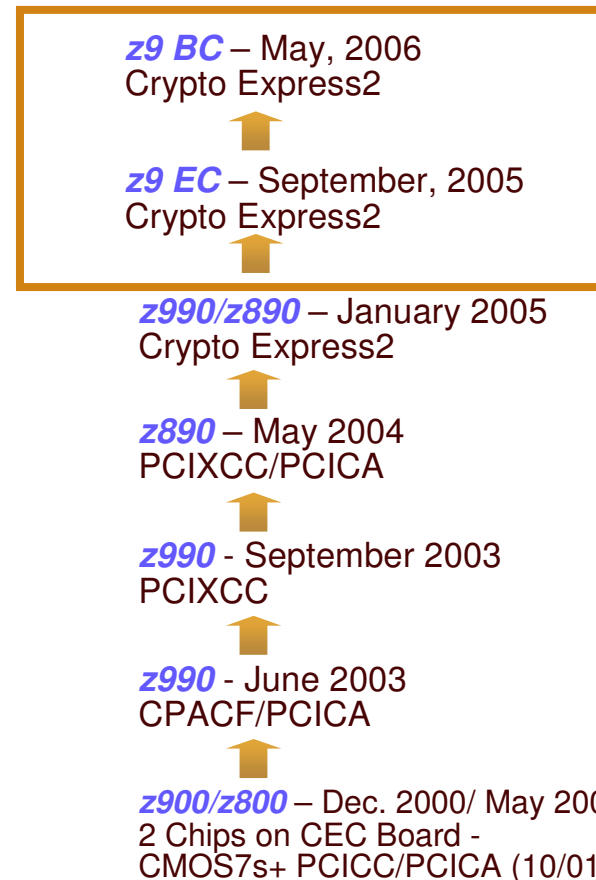
G5 – Sept. 1998
 2 Chips on Processor
 MCM - CMOS5x +
 PCICC (6/99)



G4 – Sept. 1997
 SCMs on Planar
 Board - CMOS5x

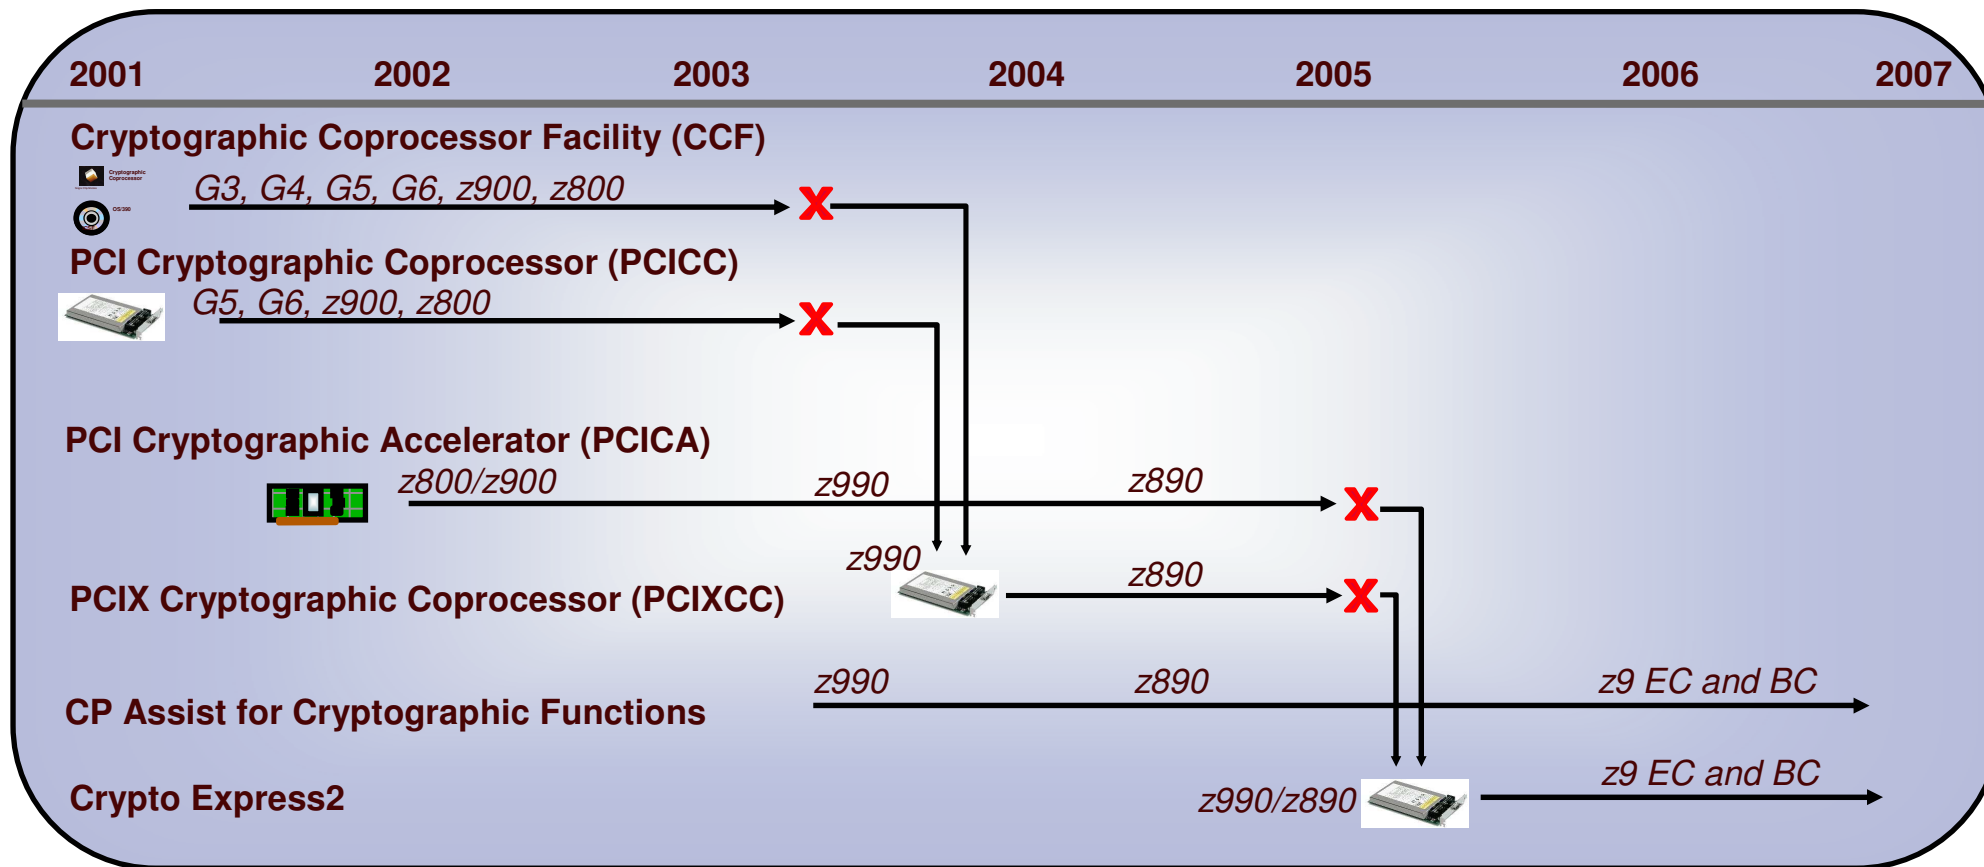


G3 – June, 1997
 SCMs on Planar
 Board - CMOS5x



*These measurements are examples of the maximum handshakes per second achieved in a laboratory environment with no other processing occurring and do not represent actual field measurements. Details are available upon request.

System z9 and zSeries Crypto Roadmap



- 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 “Clear key” crypto functions from any CP/IFL
- Crypto Express2 – Combines function and performance of PCICA and PCICC

z9 BC/EC – Enhancing security

- **Integrated cryptography features offer more security options on z9 EC**
 - ▶ Advanced Encryption Standard (AES) support in z9 EC hardware
 - ▶ Stronger hash algorithm with SHA-256
 - ▶ Pseudo Random Number Generator
 - ▶ ATM/POS remote key loading support
- **Crypto Express2 improved flexibility and speed**
 - ▶ Configurability options, two coprocessors, two accelerators or one of each
 - ▶ With both adapters configured as accelerators each Crypto Express2 card is designed to provide up to 6000 SSL handshakes per second *
- **Secure encryption facility for z/OS to help protect data shared with partners, suppliers, and customers**
 - ▶ Designed to leverage z/OS key management and high performance hardware encryption
- **Can help to achieve higher levels of certifications and compliance**
- **Virtualized cryptographic capabilities for card sharing by Linux virtual servers**
- **Complementary IBM technology and vendors' advanced security solutions**
 - ▶ Can enable a cross-platform model that can extend RACF® capabilities to the enterprise
 - ▶ Expansion of ISV community ensures application availability



System z9 Cryptographic Support

Summary

■ CP Assist for Cryptographic Function (CPACF) 1

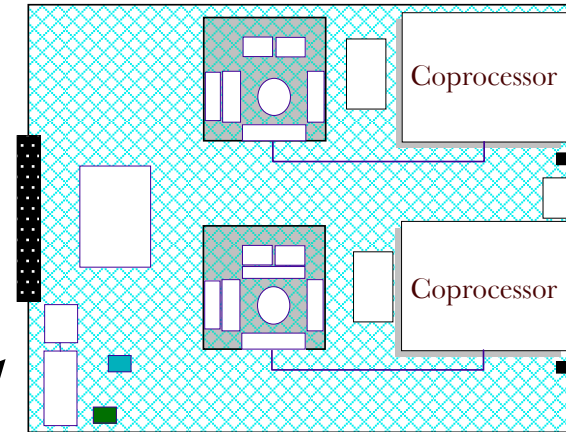
- ▶ Standard on every CP and IFL
- ▶ Supports DES, TDES and SHA-1
- ▶ New to z9 EC
 - Advanced Encryption Standard (AES)
 - Secure Hash Algorithm – 256 (SHA-256)
 - Pseudo Random Number Generation (PRNG)

■ Crypto Express2

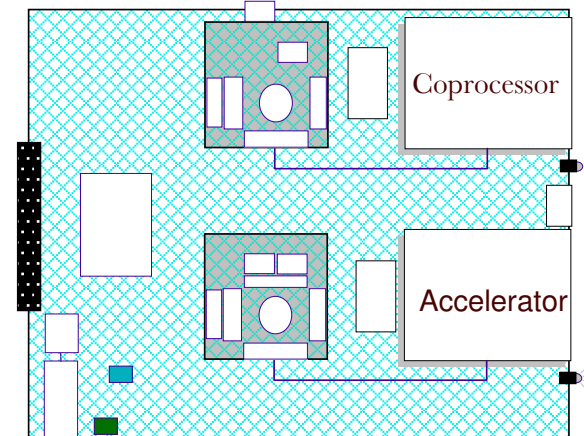
- ▶ Two configuration modes
 - Coprocessor (default)
 - Designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification
 - Accelerator (configured from the HMC)
- ▶ Three configuration options
 - Default set to Coprocessor

■ TKE workstation with 5.0 level of LIC

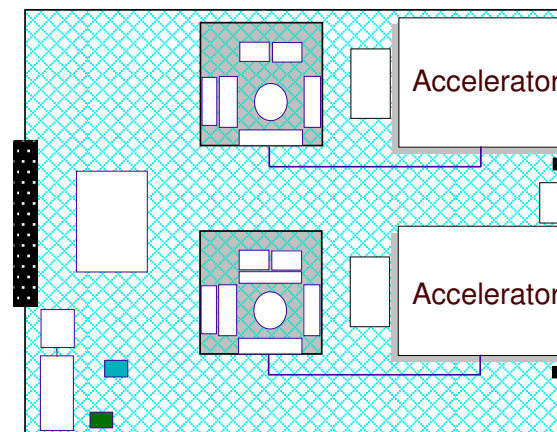
- ▶ Supports configurable Crypto Express2 feature
- ▶ New Graphical User Interface (GUI)
- ▶ Smart Card Reader



2



3



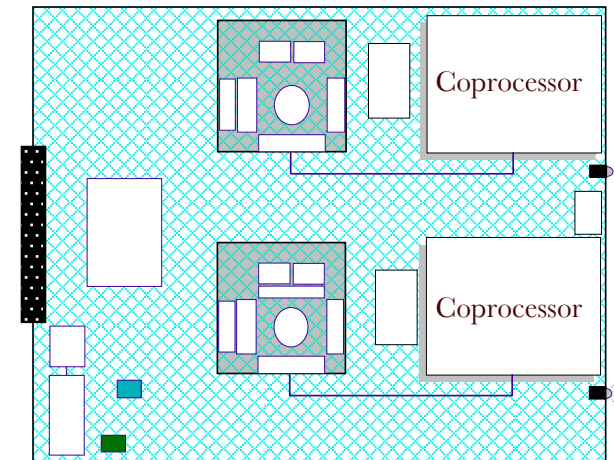
System z9 CPACF Support

- **CP Assist for Cryptographic Function (CPACF)**
 - ▶ Available on every CP & IFL
 - ▶ High performance clear key symmetric encryption/decryption
 - Advanced Encryption Standard (AES) – 128-bit
 - Triple DES / DES
 - Requires no charge enablement feature
 - ▶ High performance clear key hashing
 - Secure Hash Algorithm (SHA)-256
 - SHA-1
 - Shipped enabled on all systems
 - ▶ High performance Pseudo Random Number Generator (PRNG)
 - Requires no charge enablement feature
 - ▶ Called via ICSF API or Problem State Instructions
- **CPACF Enabler Feature**
 - ▶ No additional charge export control feature
 - ▶ Required to enable AES, DES/DES, and PRNG (SHA-1 and SHA-256 are always enabled)
 - ▶ Required to order Crypto Express2

System z9 Cryptographic Coprocessor

▪ Crypto Express2 Coprocessor

- ▶ Default configuration for Crypto Express2 feature
 - Provides 'secure-key' and 'public key' functionality
- ▶ Scalable - 0 to 8 features
 - Minimum purchase increment is two
- ▶ Configurable
 - 0, 1, or 2 coprocessors per feature
 - Individually by PCIX adapter
- ▶ Current applications expected to run without change
- ▶ Connection to STI interface; no external cables
- ▶ Fully programmable, User Defined Extensions (UDX) support
- ▶ Designed for FIPS 140-2 Level 4 Certification
- ▶ Trusted Key Entry (TKE) 5.0 support
 - Supports Crypto Express2 coprocessor
 - Smart Card Reader support
- ▶ PCIXCC cannot be carried forward to z9 EC
 - Replaced by Crypto Express2 Coprocessor

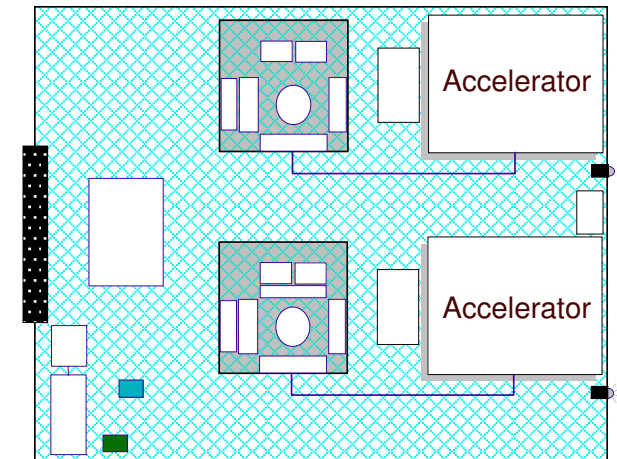


Note: A TKE workstation is required to manage Crypto Express2 features WHEN configured as a coprocessor

System z9 Cryptographic Accelerator

■ Crypto Express2 Accelerator

- ▶ Non-default configuration for Crypto Express2 feature
 - Configured from the HMC
 - Provides SSL acceleration functions
- ▶ Scalable - 0 to 8 features
 - Minimum purchase increment is two
- ▶ Configurable
 - 0, 1, or 2 accelerators per feature
 - Individually by PCI-X adapter
- ▶ Hardware acceleration for Secure Sockets Layer (SSL transactions)
- ▶ High performance public key (RSA) acceleration
- ▶ Connection to STI interface; no external cables
- ▶ PCICA cannot be carried forward to z9 EC
 - Replaced by Crypto Express2 Accelerator



System z9 Cryptographic Enhancements

- **Remote loading of initial Automated Teller Machines (ATM) keys**
 - ▶ Typically, a new ATM has none of the financial institutions keys installed.
 - ▶ Remote Key Loading refers to the process of loading Data Encryption Standard (DES) keys to Automated Teller Machines (ATMs) from a central administrative site without the need for personnel to visit each machine to manually load DES keys.
 - ▶ This has been done by manually loading each of the two cleartext key parts individually and separately into ATMs.
 - ▶ Manual entry of keys is one of the most error-prone and labor-intensive activities that occur during an installation, making it expensive for the banks and financial institutions.
- **Remote Key Loading possible Benefits**
 - ▶ Provides a mechanism to load initial ATM keys without the need to send technical staff to ATMs.
 - ▶ Reduces downtime due to key entry errors.
 - ▶ Reduces service call and key management costs.
 - ▶ Improves the ability to manage ATM conversions and upgrades.
- **ATMs which can support the IBM System z9 Remote Key Load function:**
 - ▶ ATMs supporting the mechanisms, outlined in following standards, should be compatible with IBM System z9 Remote Key Load solution:
 - ISO/IEC 11770-3: Information Technology, Security Techniques, Key Management, Part 3: Mechanisms Using Asymmetric Techniques.
 - ANS X9.24-2 (Draft): Retail Financial Services, Symmetric Key Management, Part 2: Using Asymmetric Techniques for the Distribution of Symmetric Keys.

System z9 Further Cryptographic Enhancements

- **Improved key Exchange with Non-CCA Cryptographic systems:**
 - ▶ IBM Common Cryptographic Architecture (CCA) employs Control Vectors to control usage of cryptographic keys.
 - ▶ Non-CCA systems use other mechanisms, or may use keys that have no associated control information.
- **Possible benefits of improved key exchange:**
 - ▶ Enhances the ability to exchange keys between CCA systems, and systems that do not use Control Vectors.
 - ▶ Allows the CCA system owner to define permitted types of key import and export.
 - Can help to prevent uncontrolled key exchange that can open the system to an increased threat of attack.
- **These enhancements are exclusive to System z9 and supported by z/OS 1.6 or higher**

System z9 Further Cryptographic Enhancements

- **ISO 16609 CBC Mode T-DES MAC enhancement**
- **Supports the requirements for Message Authentication, using symmetric techniques**
- **The Integrated Cryptographic Service Facility (ICSF) will use the following callable services to access the ISO 16609 CBC Mode T-DES MAC enhancement in the Cryptographic coprocessor:**
 - ▶ MAC Generate (CSNBMGN)
 - ▶ MAC Verify (CSNVMVR)
 - ▶ Digital Signature Verify (CSNDDSV)
- **ISO 16609 CBC mode T-DES MAC is accessible through ICSF function calls made in the Cryptographic Adapter Segment 3 Common Cryptographic Architecture (CCA) code**
- **This enhancement is exclusive to System z9 and supported by z/OS 1.6 or higher**

System z9 New and Modified ICSF Callable Services

- **Integrated Cryptographic Service Facility (ICSF), together with Crypto Express2, support the basic mechanisms in Remote Key Loading**
- **The implementation offers a bridge between the highly secure CCA environment and the various formats and encryption schemes offered by the ATM vendors**
- **The following are new ICSF services are offered for Remote Key loading:**
 - ▶ **Trusted Block Create (CSNDTBC)**
 - This callable service is used to create a trusted block containing a public key and some processing rules. The rules define the ways and formats in which keys are generated and exported.
 - ▶ **Remote Key Export (CSNDRKX)**
 - This callable service uses the trusted block to generate or export DES keys for local use and for distribution to an ATM or other remote device.
- **Modified ICSF Callable Services for ISO 16609 CBC Mode T-DES MAC enhancement:**
 - ▶ **Digital Signature Verify (CSNDDSV)**
 - ▶ **MAC Generate (CSNBMGN)**
 - ▶ **MAC Verify (CSNVMVR)**
- **Refer to Application Programmers Guide, SA22-7522, for additional details.**

System z9 z/OS Crypto Support

z/OS Release	Crypto Web Download	FMID	APAR	Comments
z/OS 1.4/1.5 (1.4 with z990 compatibility feature or z990 exploitation feature)	09/2003, z990 Cryptographic Support (Web download no longer available) OR	HCR770A	OA09157 OA11946	Co-Processor - okay No Accelerator
	05/2004, z990 and z890 Enhancements to Cryptographic Support	HCR770B	OA09157 OA11946	Co-Processor - okay No Accelerator
z/OS 1.6	[05/2004] z990 and z890 Enhancements to Cryptographic Support OR [12/2004] ICSF 64-bit Virtual Support for z/OS 1.6 and z/OS.e 1.6	HCR770B	OA09157 OA11946	Co-Processor - okay No Accelerator
		HCR7720	OA11946	
z/OS 1.7	[12/2004] ICSF 64-bit Virtual Support for z/OS 1.6 and z/OS.e 1.6 Replaced by Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7 09/2005	HCR7720	OA11946	Co-Processor - okay No Accelerator
z/OS V1.6 or V1.7	[09/2005] Cryptographic Support for z/OS V1R6/R7 and z/OS.e V1R6/R7	HCR7730	All Included	Co-Processor - okay Accelerator - okay

- **OA09157 - Co-Processor**

- Permits the use of the z9-109 Crypto Express2 Co-Processor as a z990/z890 Crypto Express2

- **OA11946 - Accelerator toleration**

- When an Accelerator is defined, ICSF will abend without PTF for OA11946

Secure Enterprise Data - Solutions for a Heterogeneous Environment

Helping to reduce risk end-to-end across your value chain



zSeries is becoming the HUB for enterprise security

Leverage zSeries Encryption Facilities – Integrated/Crypto Express2
Leverage z/OS key management capabilities across the enterprise



Protect Data Leaving Your Enterprise

Protect Archived Data

Businesses are proactively focusing on securing customer and business data

- Increasing regulatory requirements driving need for security of data for audit compliance
- Recent events highlight impacts caused by loss/theft of removable data
- Requirements for tighter security driving need for encryption of data

Enterprise-wide Key Management



Protect Enterprise-wide Data

IBM future directions*

- Encryption for IBM TotalStorage
- Continued enhancements of mainframe key mgnt services

- Encryption with key management for removable media
- Provide highly secure data transfer with partners, suppliers, and customers
- Enable high volume encryption of data for remote site archival
- Can create trusted exchange with heterogeneous systems with support for open standards for IP encryption

Directions for Data Encryption in IBM Storage

Statement of Direction

- ▶ Support of encryption capabilities within storage environments such that the capability does not require the use of host server resources (so called "outboard" encryption capabilities)
- ▶ Includes the intent to support outboard encryption and to leverage the centralized key management provided by z/OS ICSF
- ▶ First implementation of outboard encryption is planned for the IBM System Storage TS1120 Tape Drive in the second half of 2006. Plan is to include:
 - A chargeable microcode and hardware upgrade option for older TS1120 Tape Drives shipped before availability of the encryption feature
 - A new software program for management of encryption keys for tape drives across the enterprise – the new software may use standard key repositories, including ICSF on z/OS



Improve data protection and resiliency across the IT infrastructure

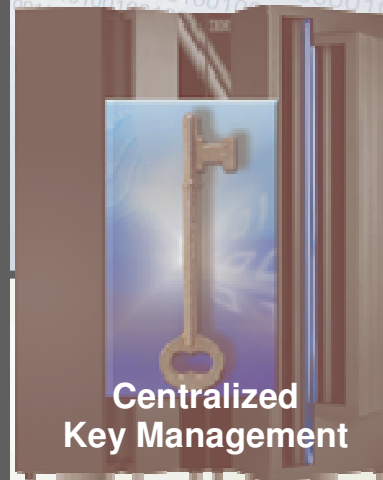
Tape Encryption with Key Management on System z

Why z/OS centralized key management?

- Can help to protect and manage keys
 - Highly secure and available key data store
 - Long term key management
 - Disaster recovery capabilities
- Single point of control
- Over a decade of production use

Encryption Facility for z/OS, V1.1

Data Encryption in the Server



- Flexible options for business partner exchange
- Partners can encrypt and decrypt using no-charge Java client
- Supports public key or password based exchange
- Plans to support OpenPGP standard*

Protected Encryption Keys

Data Encryption in TS1120*



Plans for encryption in IBM System Storage™ (*)



Enterprise scope

- Highly secure tape library
- High performance archive encryption
- Transparent to existing processes and applications
- Can help provide audit compliance



Server Time Protocol

Time Synchronization – The Next Generation



**ENABLING BUSINESS.
A THROUGH Z.**

What is Server Time Protocol ?

- **Designed to provide the capability for multiple servers to maintain time synchronization with each other and form a Coordinated Timing Network (CTN)**
 - ▶ CTN: a collection of servers that are time synchronized to a time value called Coordinated Server Time (CST)

- **IBM Server-wide facility implemented in IBM System z9™, IBM eServer™ zSeries® 990 (z990), zSeries 890 (z890) Licensed Internal Code (LIC)**
 - ▶ FC1021 on each server and/or Coupling Facility
 - ▶ Single view of “time” to PR/SM™
 - ▶ PR/SM can virtualize this view of time to the individual logical partitions (LPARs)
 - ▶ STP is not available on z900, z800 or 9672 Gx servers

- **Message based time synchronization protocol**
 - ▶ Similar to Network Time Protocol (NTP) – an industry standard
 - ▶ Timekeeping info transmitted over Coupling Links (*ISC-3 links in Peer mode, ICB-3 and ICB-4*)
 - ▶ **NOT** standard NTP (*see SOD*)

- **Designed to support a multi-site Sysplex of up to 100 km (62 miles) with repeaters**

- **Allows use of dial-out time services to set time to international time standard (UTC) to within +/- 100 ms as well as adjust to UTC on a periodic basis**

Terminology

- **STP-capable server/CF**

- ▶ IBM System z9 Enterprise Class (z9 EC), System z9 Business Class (z9 BC), z990, z890 server/CF with STP LIC installed

- **STP-enabled server/CF**

- ▶ STP-capable server/CF with STP Feature Code 1021 installed
- ▶ STP panels at the HMC/SE can now be used

- **STP-configured server/CF**

- ▶ STP-enabled server/CF with a CTN ID assigned
- ▶ STP message exchanges can take place

- **CTN**

- ▶ Collection of servers that are time synchronized to a time value called Coordinated Server Time (CST)

- **CTN ID**

- ▶ Servers / Coupling Facilities (CFs) that make up a CTN are all configured with a common identifier - CTN ID

Terminology (continued)

Two types of CTN configurations possible:

■ Mixed CTN

- ▶ Allows servers/CFs that can only be synchronized to a Sysplex Timer (ETR network) to coexist with servers/CFs that can be synchronized with CST in the “same” timing network
- ▶ Sysplex Timer[®] provides timekeeping information
- ▶ CTN ID format - STP network ID concatenated with ETR network ID

■ STP-only CTN

- ▶ All servers/CFs synchronized with CST
- ▶ Sysplex Timer is NOT required and not part of CTN
- ▶ CTN ID format - STP network ID only (ETR network ID field has to be null)

Terminology (continued)

- **STP transmits timekeeping information in layers or Stratum**

- **Stratum 1**

- ▶ Highest level in the hierarchy of timing network that uses STP to synchronize to CST

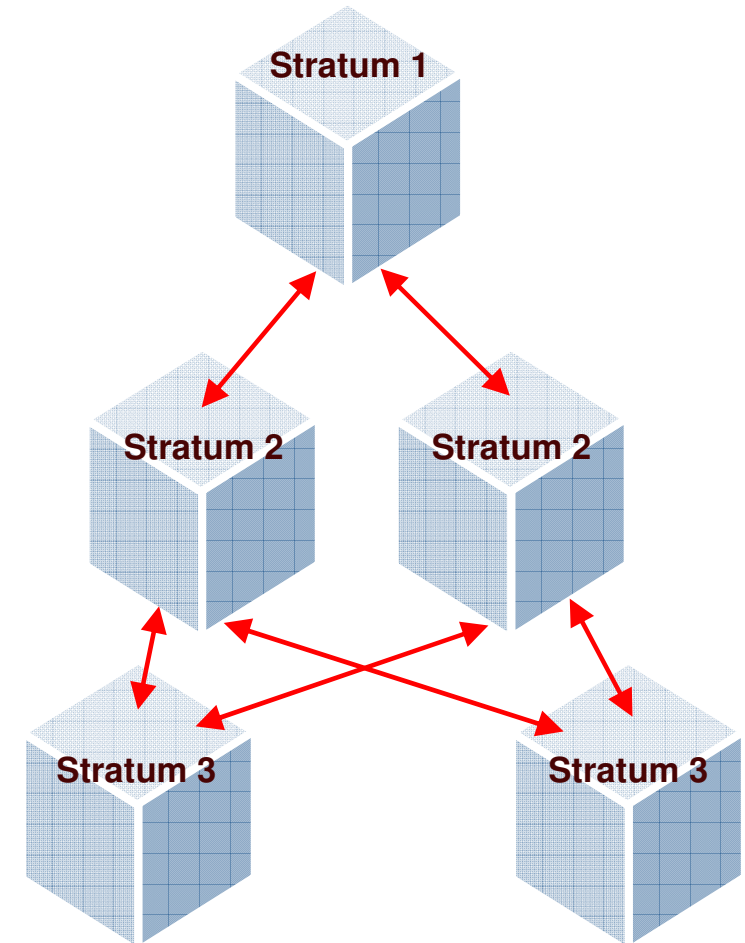
- **Stratum 2**

- ▶ Server/Coupling Facility (CF) that uses STP messages to synchronize to Stratum 1

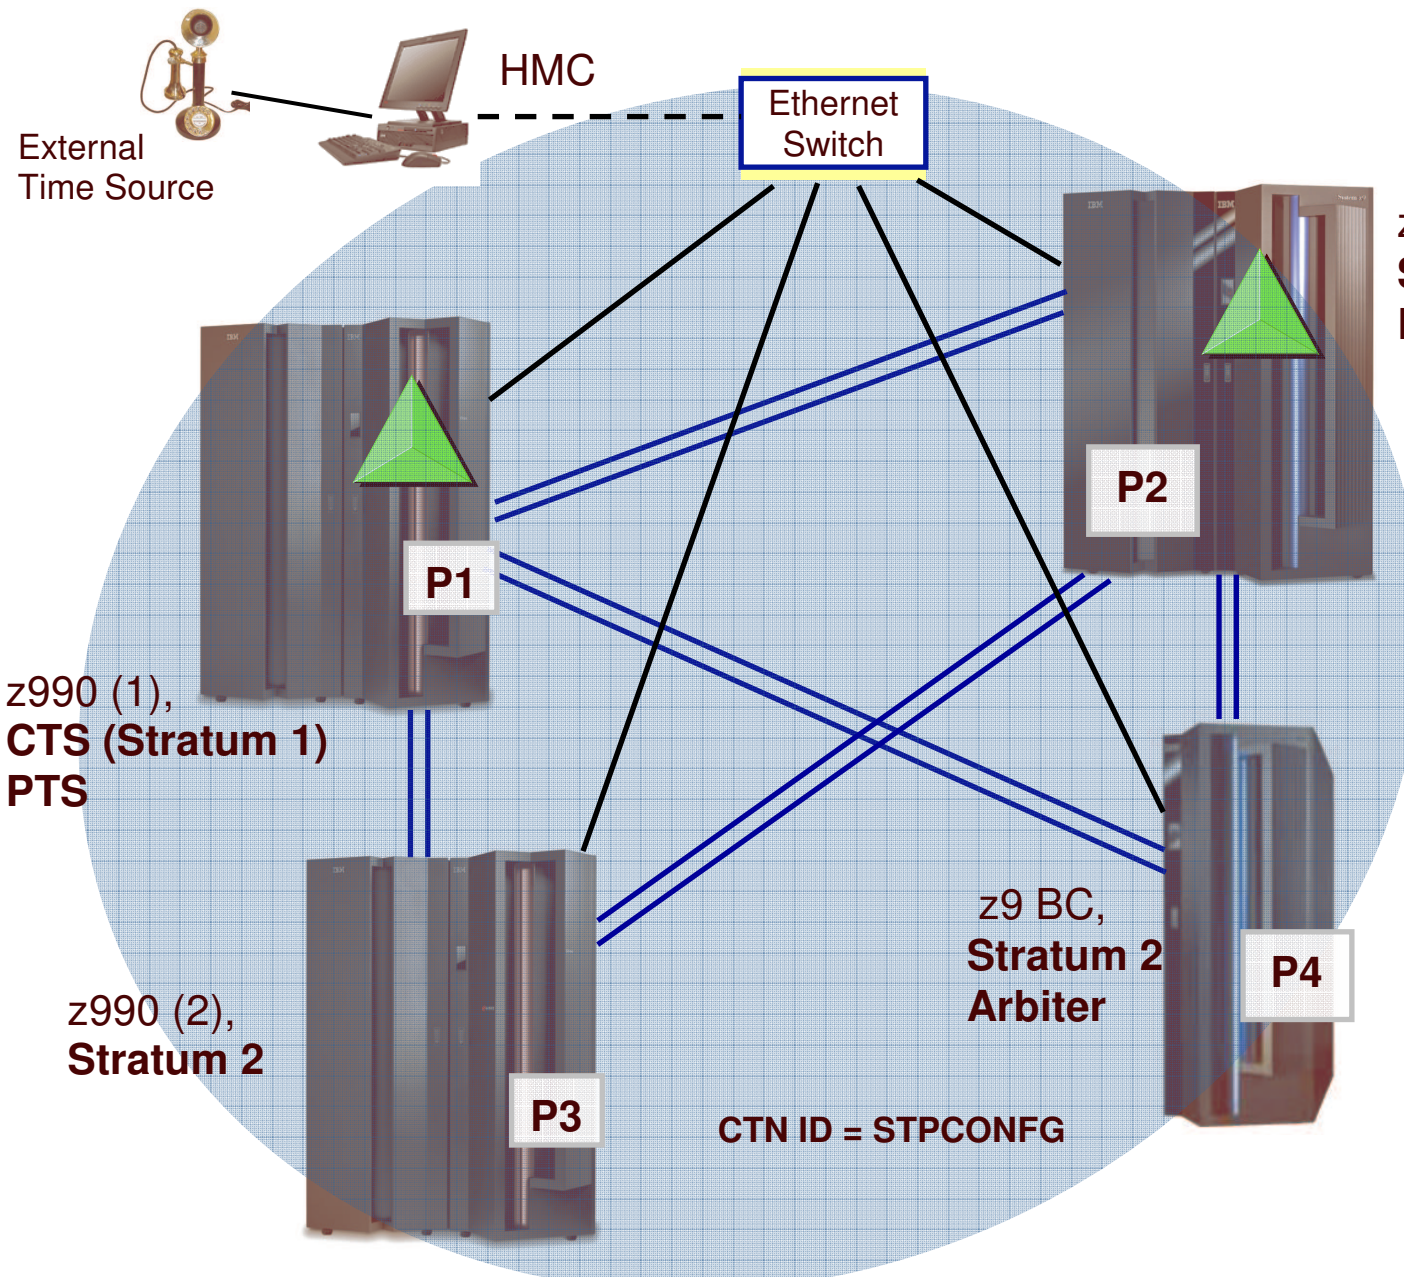
- **Stratum 3**

- ▶ Server/Coupling Facility (CF) that uses STP messages to synchronize to Stratum 2

- **STP supports configurations up to Stratum 3**



Example of STP-only CTN



STP can use existing Coupling links

- STP Configuration**
- **Must do: Assign PTS/CTS**
 - ▶ Preferred Time Server
 - ▶ Current Time Server
 - **Option: Assign Backup Time Server (BTS)**
 - ▶ Strongly recommended to allow for continuous availability
 - **Option: Assign Arbiter**
 - ▶ Recommended for configurations of 3 or more servers/CFs
 - ▶ Can help improve recovery

STP-only CTN

All servers in STP-only CTN have to be STP capable - 9037s no longer required

Server roles

- **Preferred Time Server (PTS)**
 - ▶ Server that is preferred to be the “active” Stratum 1 server

- **Backup Time Server (BTS)**
 - ▶ Role is to take over as the Stratum 1 under planned or unplanned outages, without disrupting synchronization capability of STP-only CTN

- **Current Time Server (CTS) - “Active” Stratum 1 server**
 - ▶ Only one “active” S1 allowed
 - ▶ Only the PTS or BTS can be assigned as the CTS
 - ▶ Normally the PTS is assigned the role of CTS

- **Arbiter**
 - ▶ Provides additional means to determine if BTS should take over as the CTS under unplanned outages

STP-only CTN (continued)

- **HMC must be used to provide the following functions:**
 - ▶ Initialize Coordinated Server Time (CST) manually
 - ▶ Initialize CST to an international time standard (UTC)
 - Dial-out from HMC to set CST to within +/- 100 ms of UTC
 - ▶ Schedule periodic dial outs to maintain accurate time
 - ▶ Set Time Zone Offset, Daylight Saving Time Offset, Leap seconds Offset
 - ▶ Schedule and change Offsets (Daylight Saving, Leap seconds)
 - Automatic scheduling of Daylight Savings Time based on algorithm
 - ▶ Adjust time by up to +/- 60 seconds (currently 9037 allows 4.999 seconds)
 - ▶ Define, modify, view the STP-only CTN ID

- **Concurrent migration (with proper planning)**
 - ▶ Concurrently migrate from Mixed CTN to STP-only CTN or
 - ▶ Concurrently migrate from existing ETR network to STP-only CTN
 - ▶ Concurrently migrate from STP-only CTN to Mixed CTN
 - Allows testing STP-only CTN during change window and go back concurrently to Mixed CTN or ETR network at start of production.

Role Recommendations

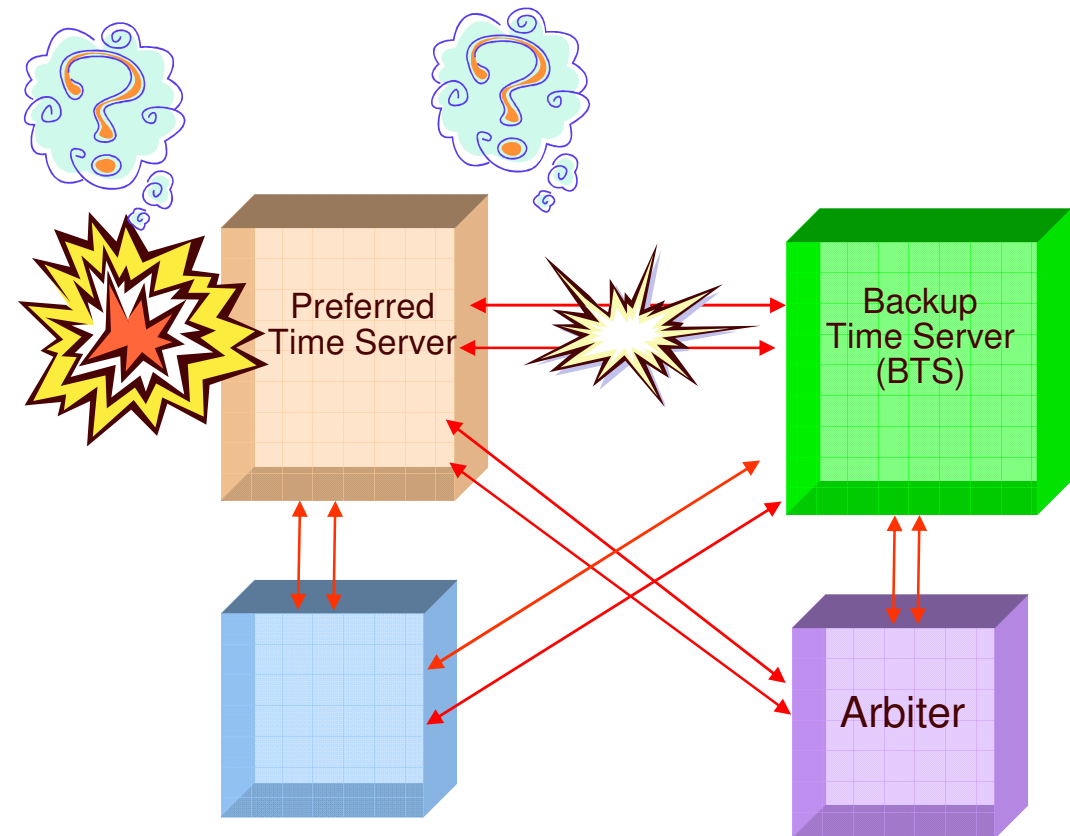
Your choice of Preferred and Current Time Server, should not be made lightly.

- **This is called the preferred server for good reason. Some consideration should be given to:**
 - ▶ The location of the server. Is it in the most important site (if relevant)?
 - ▶ Is it of the latest technology in the complex? New STP function will be introduced to the latest technology, not necessarily older server generations.
 - ▶ Maintenance considerations. A server that is regularly unavailable due to scheduled maintenance is not a good choice.
 - ▶ Available connectivity. An established coupling facility server will typically have coupling link connectivity to existing servers. This existing connectivity can be exploited for timing links without having to deploy additional hardware.
 - ▶ z/OS messaging. A stand-alone coupling facility will produce no z/OS messages for interception by automation routines. Also, no display commands are possible from this server. However, other servers in the complex will produce STP related z/OS messages and report on the status of the timing networks which may meet your requirements.
- **Once these points have been considered, a similar review should be made for the Backup Time Server and Arbiter.**

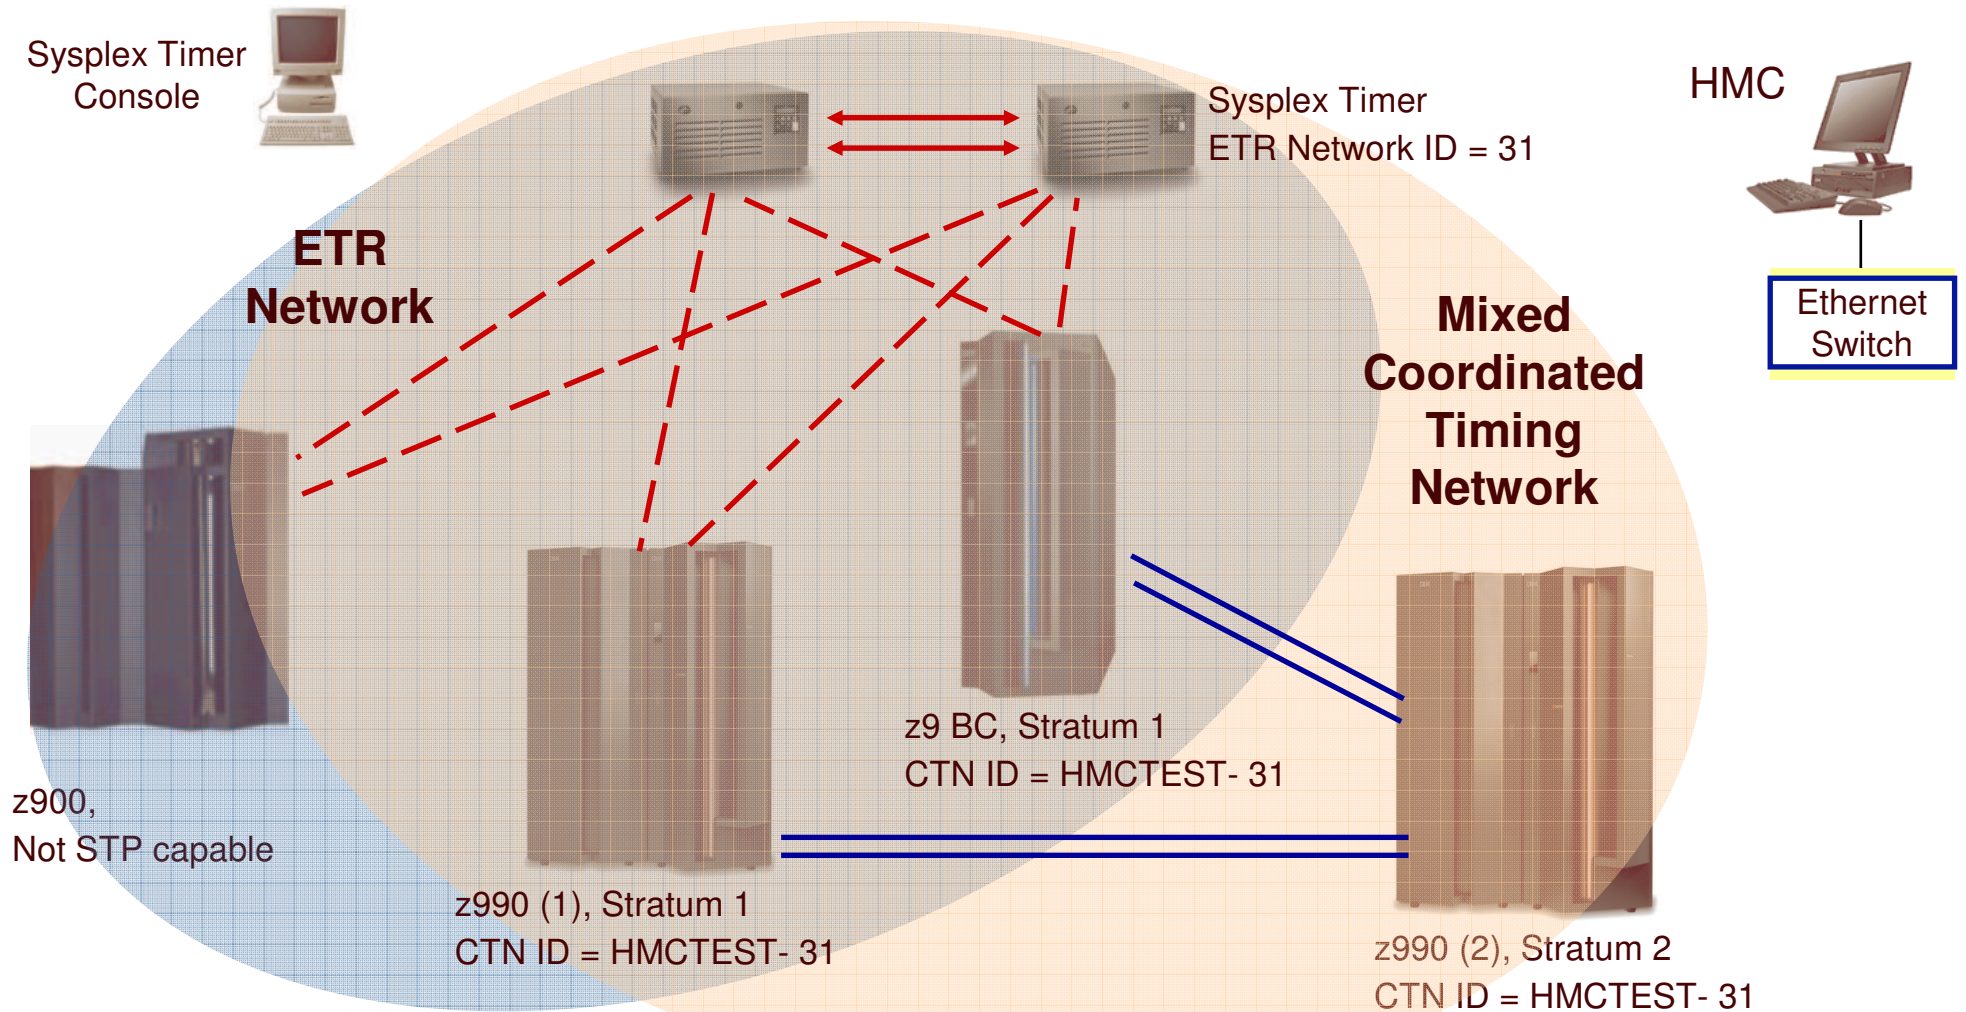
Example of Backup Time Server / Arbiter

If the Backup Time Server loses communication with PTS on all established paths, what does it do ?

- It asks the Arbiter for its status
- If Arbiter still getting signals (just links are lost), BTS becomes Stratum 3 (get signals from a level 2 Stratum)
- If Arbiter also lost communication, (assume PTS failed), BTS takes over as CTS (Stratum 1)



Mixed Coordinated Timing Network



■ z9 BC, z990(1), z900 synchronized to Sysplex Timer

■ z9 BC, z990(1) are Stratum 1 servers

▶ Two Stratum 1 servers recommended to help avoid single point of failure

■ z990(2) synchronized to either z9 BC or z990(1) via STP is a Stratum 2 server

▶ z990(2) does not need ETR link connections

▶ z990(2) can be located up to 100 km away from z9 BC, 990(1)

Mixed Coordinated Timing Network (CTN)

- **Need at least one STP-enabled server to configure Mixed CTN**
 - ▶ Selected STP-enabled server **MUST** also be synchronized to the Sysplex Timer
 - Automatically becomes a Stratum 1 server for the Mixed CTN
 - ▶ Stratum 2 server/CF uses Stratum 1 as clock source
 - ▶ Stratum 3 server/CF uses Stratum 2 as clock source

- **Sysplex Timer provides timekeeping information for Mixed CTN**

- **Sysplex Timer console continues to be used for all timing related functions of the Mixed CTN**
 - ▶ Initialize time; Set Time Zone, Daylight Saving Time (DST), Leap seconds offsets
 - ▶ Schedule DST and Leap seconds offset changes
 - ▶ Adjust time up to +/- 4.999 seconds

- **Hardware Management Console (HMC) must be used for Mixed CTN ID initialization and modification**

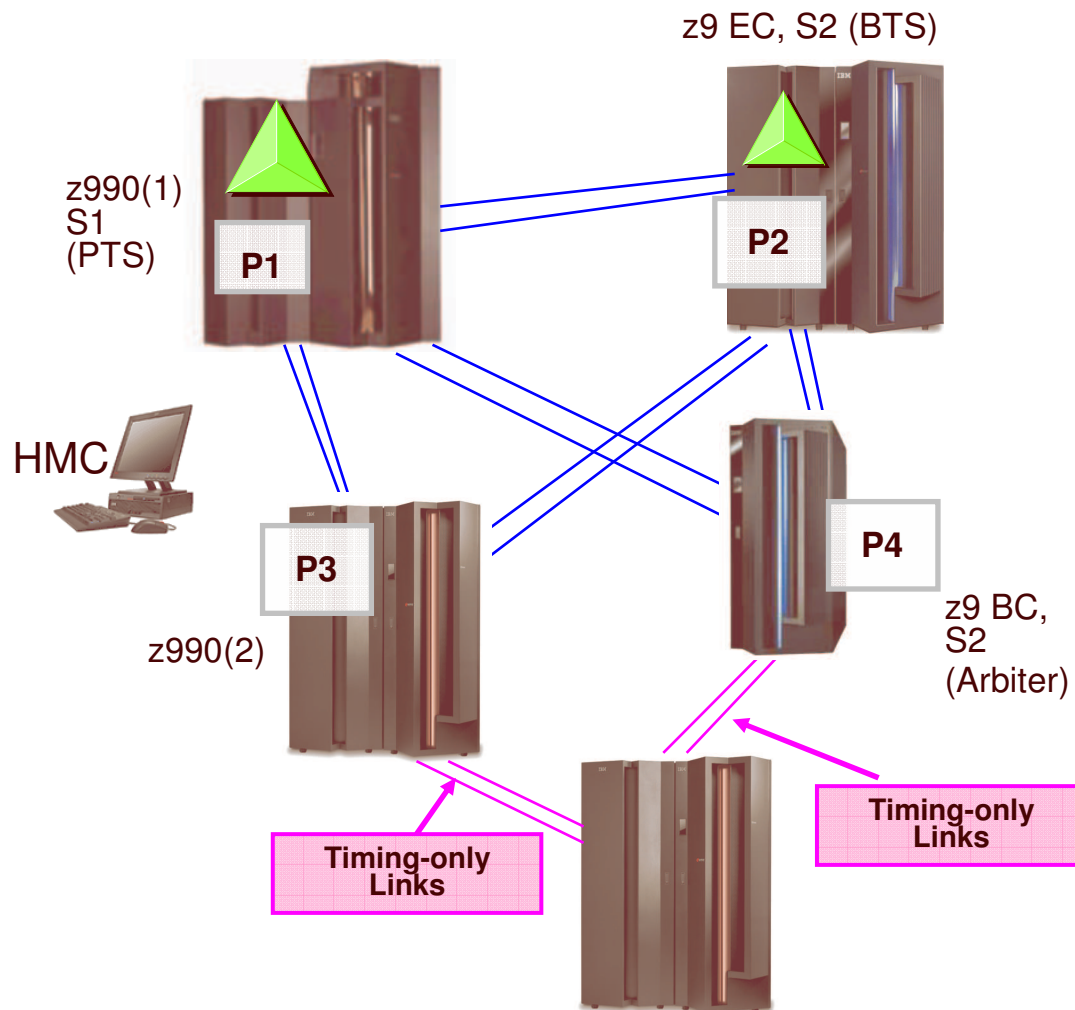
Mixed CTN (continued)

Mixed CTN allows:

- **Concurrent migration from ETR network (with proper planning)**
 - ▶ Concurrently migrate from existing ETR network to Mixed CTN
 - ▶ Concurrently migrate from Mixed CTN to ETR Network
 - Allows testing Mixed CTN during change window and go back concurrently to ETR network at start of production

- **Coexistence in the “same” timing network of**
 - ▶ Servers/CFs that can only be synchronized to a Sysplex Timer (ETR network) and
 - ▶ Servers/CFs that can be synchronized with CST
 - ▶ In a Parallel Sysplex configuration, the only non STP-capable server and CF that can coexist are the z900 and z800 server and CF
 - Non STP-capable Server/CF MUST support Message Time Ordering Facility (MTOF)
 - Non STP-capable server/CF MUST be attached to Sysplex Timer


Timing-only Links – *what is it...?*



- **Coupling links that allow 2 servers to be synchronized when a CF does not exist at either end of link**
- These timing-only links are used to transmit STP messages only
- Typically required when synchronization needed in a Non-Parallel Sysplex configuration *(for example XRC)*
- HCD enhanced to define Timing-only links
- Can be defined in either Mixed CTN or STP-only CTN

STP – key attributes

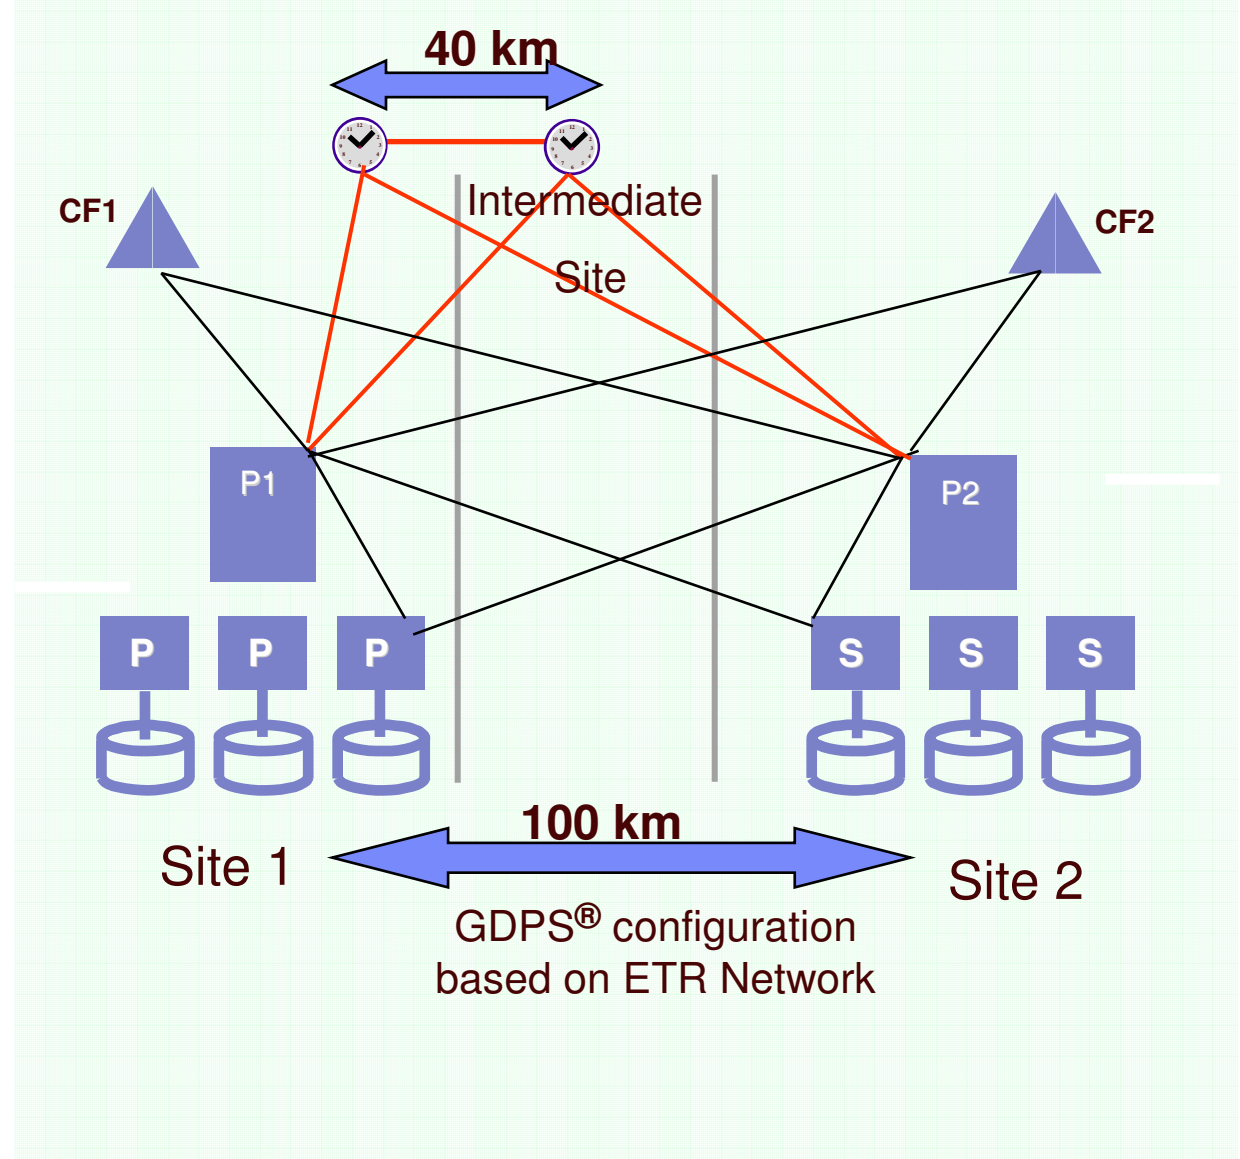
- Allows Parallel Sysplex distances to extend beyond the current 40 km limit
 - ▶ Limits set by coupling protocol and links
- Can help meet more stringent (precise) time synchronization requirements
- Expected to Scale with technology as processors and messaging technology improve
- Does not require dedicated Timer links
 - ▶ Uses same hardware and protocols as data
 - ▶ Reduced requirements to infrastructure
- Allows concurrent migration from an ETR network with proper planning
- Allows coexistence with ETR network
- Allows time to be set to a dialup service to within 100 ms of UTC



STP Enhancements over ETR Network

STP design can allow longer distances with less complex and less costly infrastructure

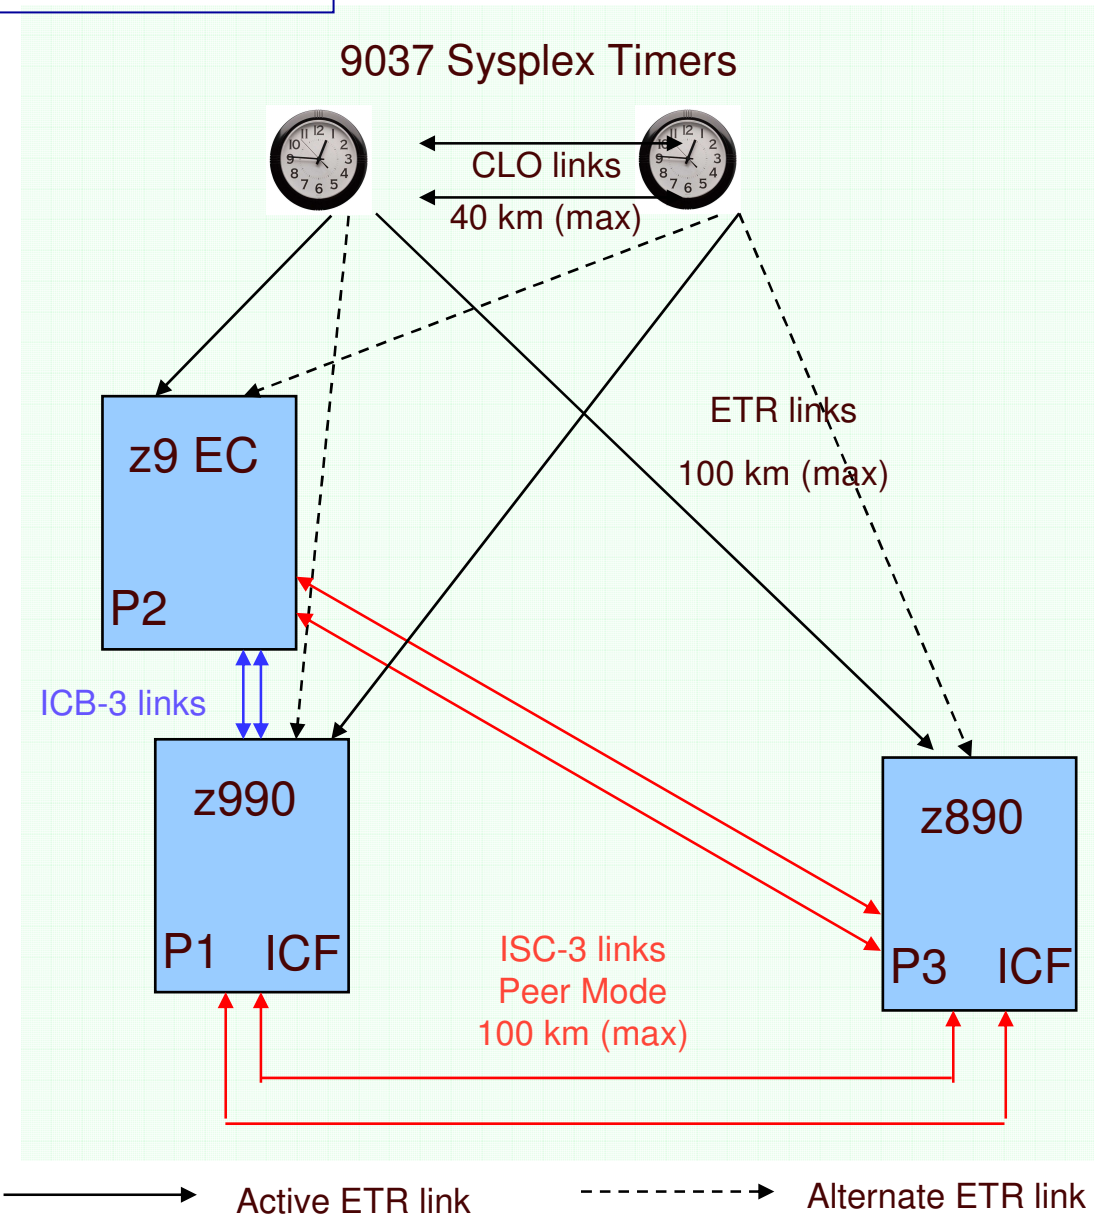
- STP supports a multi-site timing network of up to 100 km without requiring an intermediate site
- Fiber distance between Sysplex Timers cannot exceed 40 km
- Intermediate site to locate second timer recommended to avoid a single point of failure, if data centers more than 40 km apart



STP Enhancements over ETR Network (cont) ...

STP design can allow more stringent synchronization

- Time of Day (TOD) clocks of servers must be synchronized within the fastest messaging time between servers
- In this diagram, between IBM System z9 Enterprise Class (z9 EC) and z990 (shortest messaging time in example) compared to between z990 and z890 (longer messaging time in example)
- “Best case” messaging times over ICB links in Parallel Sysplex (8 us approximately) approaching “Worst case” TOD synchronization between CECs stepping to 9037s 40 km apart (4 us approx)



STP - key attributes *(cont)* ...

Allows

- **Use of dial-out time services to initialize Coordinated Server Time (CST) to within +/- 100 ms of international time standard (UTC)**
 - ▶ NIST Automated Computer Time Service (ACTS)
 - ▶ NRC Canadian Time Service (CTS)
 - ▶ IEN Telephone Date Code (CTD)
- **Scheduling of periodic dial-outs to time services so that CST can be gradually steered to an international time standard (UTC)**
- **Setting of local time parameters,**
 - ▶ Time zone offset
 - ▶ Daylight Saving Time (DST) offset
 - ▶ Leap Seconds offset
- **Automatic updates of DST offset based on time zone algorithm**
- **Adjustment of CST up to +/- 60 seconds**

Other potential STP values

■ Helps Eliminate

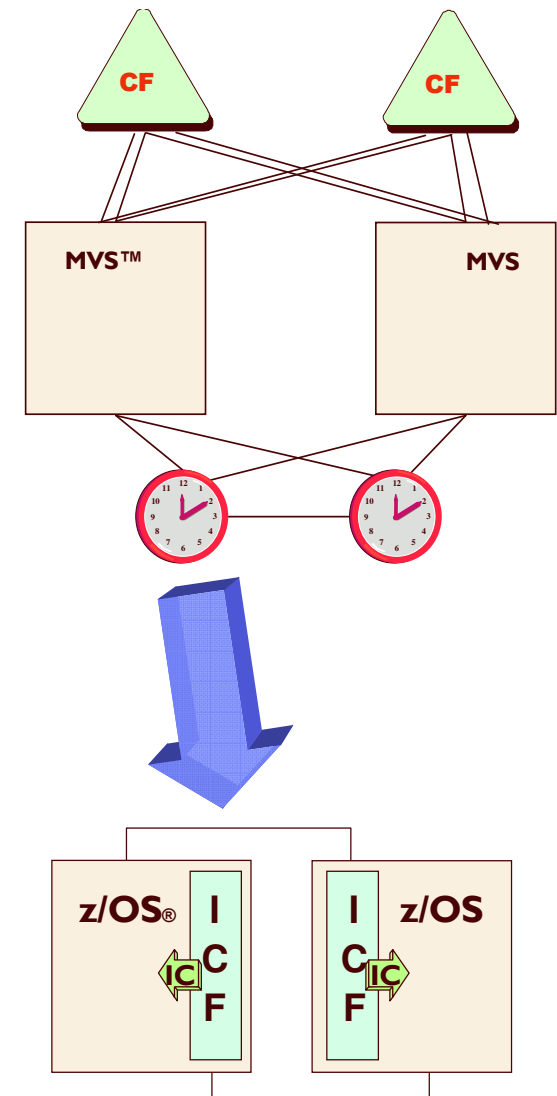
- ▶ Infrastructure requirements (space, power, etc.) to support Sysplex Timers
- ▶ Sysplex Timer maintenance costs.
- ▶ Dark fiber between sites for ETR and CLO links

■ Helps Reduce

- ▶ Fiber optic infrastructure requirements for DWDM ports, patch/trunk cables


■ Helps improve System Management

- ▶ Allows automatic adjustment of Daylight Saving Time offset based on time zone algorithm
 - With ETR network, you need to schedule DST offsets at least twice a year manually at the Sysplex Timer console
- ▶ Allows gradual time adjustment of up to +/- 60 seconds
 - Sysplex Timer allows time adjustments of up to +/- 4.999 seconds



Statements of Direction

- IBM intends to enhance the accuracy of initializing and maintaining Coordinated Server Time to an international time standard such as Coordinated Universal Time (UTC). The then current server is planned to have the capability of attaching to an external time source, such as a Global Positioning System (GPS) receiver.
- Network Time Protocol (NTP) client support: IBM intends to enhance the STP design to provide Network Time Protocol (NTP) client capability, so that Coordinated Server Time may be initialized and maintained to time provided by an NTP server. The purpose of this function is to allow the same time across an enterprise comprised of heterogeneous platforms.



* 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.

Publications

- **Redbook™**
 - ▶ Server Time Protocol Planning Guide, SG24-7280 (available now)
 - ▶ Server Time Protocol Implementation Guide, SG24-7281 (available at GA)
 - ▶ S/390 Time Management and IBM 9037 Sysplex Timer, SG24-2070
- **Resource Link**
 - ▶ Introduction to STP Education Module (available at GA)
 - ▶ SAPR Guide, SA06-012 – Confirmation Form required prior to shipment of FC1021
- **STP WEB site** (www.ibm.com/systems/z/pso/stp.html)

- **Server Installation Planning**
 - ▶ System z9 EC Installation Manual for Physical Planning, GC28-6844
 - ▶ System z9 BC Installation Manual for Physical Planning, GC28-6855
 - ▶ zSeries 990 Installation Manual for Physical Planning, GC28-6824
 - ▶ zSeries 890 Installation Manual for Physical Planning, GC28-6828
- **Hardware Management Console User's Guide**
- **Support Element User's Guide**
- **Server Installation Manual**
- **System Overview manuals**
- **PR/SM Planning Guide**
- **Server Service Guide**
- **Online Help**



Hardware & Software Planning

- **Servers and Coupling Facilities**
 - ▶ EC Drivers and STP Feature Code
- **Hardware Management Console**
- **Coupling Links and Timing-only Links**
- **Multi-Site**
 - ▶ Dense Wave Division Multiplexer (DWDM)
- **Operating Systems**
 - ▶ OS levels
 - ▶ Preventative Service Planning (PSP)
 - ▶ IPLs
 - ▶ CLOCKxx

Prerequisites

- **Hardware - *System z9, z990 or z890 server***
 - ▶ Concurrent MCLs on Driver 55 (z990, z890) to install STP Licensed Internal Code (LIC)
 - ▶ Concurrent MCL to enable STP
 - ▶ System z9 HMC
 - New HMC
 - Upgrade of existing HMC

- **Software - *z/OS 1.7 or higher***
 - ▶ STP code in z/OS 1.7 (default=disabled)
 - ▶ PTFs for STP support
 - ▶ PTF to enable STP
 - ▶ IPL required

- **Toleration PTFs for z/OS 1.4, 1.5, 1.6 in Mixed Timing Network**
 - ▶ Mixed CTN Configurations
 - ▶ Ensures server has Message Time Ordering Facility (MTOF)
 - ▶ Various prerequisite APARs
 - ▶ IPL Required

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



IBM System z9 Software



**ENABLING BUSINESS.
A THROUGH Z.**

z/OS Support for System z9

■ Improved FICON performance

Modified Indirect Data Address Word (MIDAW) (z/OS 1.6)

- ▶ New system architecture designed to improve FICON performance for extended format data sets

including:

- DB2 queries, utilities and logs
- VSAM, HFS, zFS, PDSE, IMS Fast Path, SAM-E

■ Relief for 64K device limit

Multiple Subchannel Sets (z/OS 1.7)

- ▶ Almost two-fold increase in the number of logical volumes for typical z/OS images
- ▶ Each z/OS image can use a second set of subchannels for defining Parallel Access Volumes (PAV) aliases
- ▶ Provide an additional 64K subchannels

■ Support for more real memory (z/OS 1.8)

- ▶ z/OS 1.8 limit will be 4 TB (was 128 GB)
- ▶ Up to 512 GB supported on System z9
- ▶ Up to 256 GB supported on IBM z990

Exploitation of hardware:

■ z/OS 1.6

- ▶ Up to 32 engines¹ in single image
- ▶ Up to 60 logical partitions
- ▶ Up to 63.75K subchannels
- ▶ Modified Indirect Data Address Word (MIDAWs)
- ▶ HiperSockets support of IPv6
- ▶ OSA large send for IPv4 traffic²
- ▶ zAAP², zIIP
- ▶ CPACF enhancements
- ▶ Crypto Express2

■ z/OS 1.7

- ▶ Multiple subchannel sets
- ▶ z/OS 1.7 is planned to support Server Timer Protocol*
- ▶ FICON link incident reporting
- ▶ Wild branch diagnosis

■ z/OS 1.8

- ▶ More real memory
- ▶ WLM enhancements for zAAP
- ▶ Coupling Facility performance enhancements

¹ Also available on z990

² Also available on z990 and z890

Extending Scale and Flexibility

IMAGES

- Up to 60 LPARS on a single server (z/OS 1.7 and IBM System z9™ Enterprise Class (z9 EC), up to 30 LPARS with z/OS 1.7 and System z9 Business Class (z9 BC))
- Up to 32 processors per logical partition (z/OS 1.6 with System z9 and IBM eServer™ zSeries® 990 [z990])
- Up to 32 z/OS logical partitions can be configured in a single-image Parallel Sysplex® cluster, with shared data (up to 1,024 engines total)
- Support for up to 4 TB of real memory on a single z/OS image (z/OS 1.8). (Up to 512 GB on System z9, up to 256 GB on z990)

I/O CONFIGURATIONS

- **MIDAW** - Designed to improve FICON performance (z/OS 1.6,1.7)
- Support for 63K subchannels (z/OS 1.4)
- 63.75K subchannels (in sub-channel set =0) (z/OS 1.7 and System z9)
- Support for multiple subchannel sets – 127.75K subchannels (z/OS 1.7 and System z9)
- Support for more than 32K device groups (z/OS 1.8)
- Parallel VARY OFFLINE (z/OS 1.7) and VARY ONLINE (z/OS 1.8) processing – vary up to 32 devices in parallel
- IBM TotalStorage DS6000 & DS8000 enhancements

FILES

- Support for larger sequential and EXCP data sets (>64K tracks)
 - ▶ Larger JES spool, and DFSMSHsm™ and DFSMSrmm™ journal data sets (z/OS 1.7)
- More than 255 extents per VSAM component (z/OS 1.7)
- Support for more GRS concurrent ENQs - new maximum is 2,147,483,647 (z/OS 1.8)
- Language Environment (C/C++) support for sequential data sets larger than 64K tracks and for VSAM extended addressability data sets (z/OS 1.8)
- z/OS UNIX limit on file descriptors per process increased from 128K to 512K

z/OS Support Summary



		G5/G6 Multiprise® 3000	z900 z800	z990 z890	z9 EC z9 BC	DS8000 DS6000	End of Service	Coexists with z/OS	Planned Ship Date
z/OS	1.2	x	x	x ^c			10/05	1.5	
	1.3	x	x	x ^c			3/05	1.6	
	1.4	x	x	x ¹	x ¹	x	3/07	1.7	
	1.5	x	x	x	x	x	3/07	1.8	
	1.6		x	x	x	x	9/07	1.8*	
	1.7		x	x	x	x	9/08*	1.9*	
	1.8		x	x	x	x	9/09*	1.10*	9/06
	1.9*		x	x	x	x	9/10*	1.11*	9/07*

- z/OS.e supported on z800, z890, and z9 BC only
- z/OS 1.5 , z/OS 1.6, and z/OS 1.7 are planned to coexist with z/OS 1.8
- z/OS 1.7 is planned to coexist with z/OS 1.8 and 1.9

x^c - Compatibility support only

x¹ – z990 compatibility or exploitation feature required (orderable until December 2006)

There is no IBM Bimodal Accommodation Offering available for z/OS 1.5 or higher releases.

z9 EC and z9 BC operating system software

Operating System	ESA/390 (31-bit)	z/Arch (64-bit)
z/OS Version 1 Release 4*, 5*, 6, 7, 8	No	Yes
z/OS.e# Version 1 Release 4*, 5*, 6, 7, 8	No	Yes
Linux, 64-bit distribution	No	Yes
Linux, 31-bit distribution	Yes	No
z/VM® Version 5 Release 1, 2	No	Yes
z/VM Version 4 Release 4 **	Yes	Yes
z/VSE*** 3.1, VSE/ESA™ 2.7 ****	Yes	No
z/VSE V4 ***** (Preview – no GA announced)	No	Yes
z/TPF Version 1	No	Yes
TPF Version 4 Release 1 (ESA mode only)	Yes	No

z/OS.e - z800, z890 and z9 BC only

* Support for z/OS 1.4 and 1.5 will end on March 31, 2007

** Support for z/VM V 4.4 will end September 30, 2006

*** z/VSE V3 can execute in 31-bit mode only. It does not implement z/Architecture™ and specifically does not implement 64-bit mode capabilities. z/VSE V3 is designed to exploit select features of IBM System z hardware.

**** Support for VSE 2.7 will end February 28, 2007

***** z/VSE V4 is designed to exploit 64-bit real memory addressing, but will not support 64-bit virtual memory addressing

Note: Please refer to the latest PSP bucket for latest PTFs for new functions/features.

System z9 Minimum Operating System Support for functions – 1

	z/OS.e z/OS	z/VM	Linux on System z	z/VSE VSE/ESA⁽¹⁾	z/TPF TPF⁽²⁾
Basic System z9 support	1.4 ⁽⁴⁾	4.4	SLES 9 RHEL 4	3.1 2.7 ⁽¹⁾	1.1 4.1 ⁽²⁾
60 Logical Partitions (30 for z9 BC)	1.4 ⁽⁴⁾	4.4	SLES 9 RHEL 4	3.1	1.1 4.1 ⁽²⁾
63.75K Subchannels	1.4 ⁽⁴⁾	4.4	SLES 9 RHEL 4		
OSA-Express2 1000BASE-T Ethernet	1.4 ⁽⁴⁾	4.4	SLES 9 RHEL 4	3.1 2.7 ⁽¹⁾	1.1 4.1 PUT 13 ⁽²⁾
MIDAW Facility	1.6	Not supported	N/A		
CPACF Enhancements	1.6 ⁽⁴⁾	4.4	SLES 9 SP3 ⁽⁵⁾ RHEL 4 U3 ⁽⁵⁾	3.1	
Crypto Express2 exploitation	1.6 ⁽⁴⁾	5.1	SLES 9	3.1 2.7 ⁽¹⁾	
HiperSockets IPv6	1.7	5.2	N/A		
OSA-Express2 Large send	1.6	Not supported	SLES 9 SP2 IBM work with LDPs ⁽³⁾		
OSA-Express2 CDLC support	1.4 ⁽⁴⁾	5.1	SLES 9 SP3 RHEL 4 U3	3.1 2.7 ⁽¹⁾	1.1
Multiple Subchannel Sets (MSS)	1.7	Not supported	IBM work with LDPs ⁽³⁾		
FICON Link Incident Report	1.7	4.4	IBM work with LDPs ⁽³⁾		
Single System Image	1.6 up to 32	5.1 up to 24	SLES 9 up to 32 RHEL 4 up to 32		1.1 up to 54
Enhanced Perf Assists for z/VM Guests	N/A	5.2	IBM work with LDPs ⁽³⁾		

1. indicates VSE/ESA

2. indicates TPF

3. IBM is working with its Linux Distribution Partners (LDPs) that this function will be provided in future Linux on System z distribution releases/service updates

4. Additional features or Web downloads required

5. IBM is working with LDPs on Kernel space exploitation⁽³⁾

Note: Please refer to the latest PSP bucket for latest PTFs for new functions/features z/OS.e - z800, z890 and z9 BC only

SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux

System z9 Minimum Operating System Support for functions – 2

	z/OS.e z/OS	z/VM	Linux on System z	z/VSE VSE/ESA⁽¹⁾	z/TPF TPF⁽²⁾
N_Port ID Virtualization	N/A	4.4 (guest)	SLES 9 SP3 IBM work with LDPs ⁽³⁾	3.1	
FCP Program Directed re-IPL	N/A	Not supported	SLES 9 SP3 IBM work with LDPs ⁽³⁾		
sub-capacity	1.4 ⁽⁴⁾	4.4	IBM Software Group products are enabled ⁽⁶⁾	3.1 2.7 ⁽¹⁾	1.1 4.1 16 ⁽²⁾
zIIP Support	1.6	Not supported	N/A	Not Supported	Not supported
Crypto Remote Key Loading	1.6 ⁽⁴⁾	5.1	N/A		
Crypto ISO 16609	1.6 ⁽⁴⁾	5.1	N/A		
FICON Express4 (CHIPD type FC)	1.4 ⁽⁴⁾	4.4	SLES 9 RHEL 4	3.1 2.7 ⁽¹⁾	1.1 4.1PUT 16 ⁽²⁾
FICON Express4 (CHIPD type FCP)	N/A	4.4	SLES 9 RHEL 4	3.1	

1. indicates VSE/ESA

2. indicates TPF

3. IBM is working with its Linux Distribution Partners (LDPs) that this function will be provided in future Linux on System z distribution releases/service updates

4. Additional features or Web downloads required

5. IBM is working with LDPs on Kernel space exploitation⁽³⁾

6. Linux and z/VM do not support it, the IBM Software Group products are enabled for it on all distributions

Note: Please refer to the latest PSP bucket for latest PTFs for new functions/features z/OS.e - z800, z890 and z9 BC only

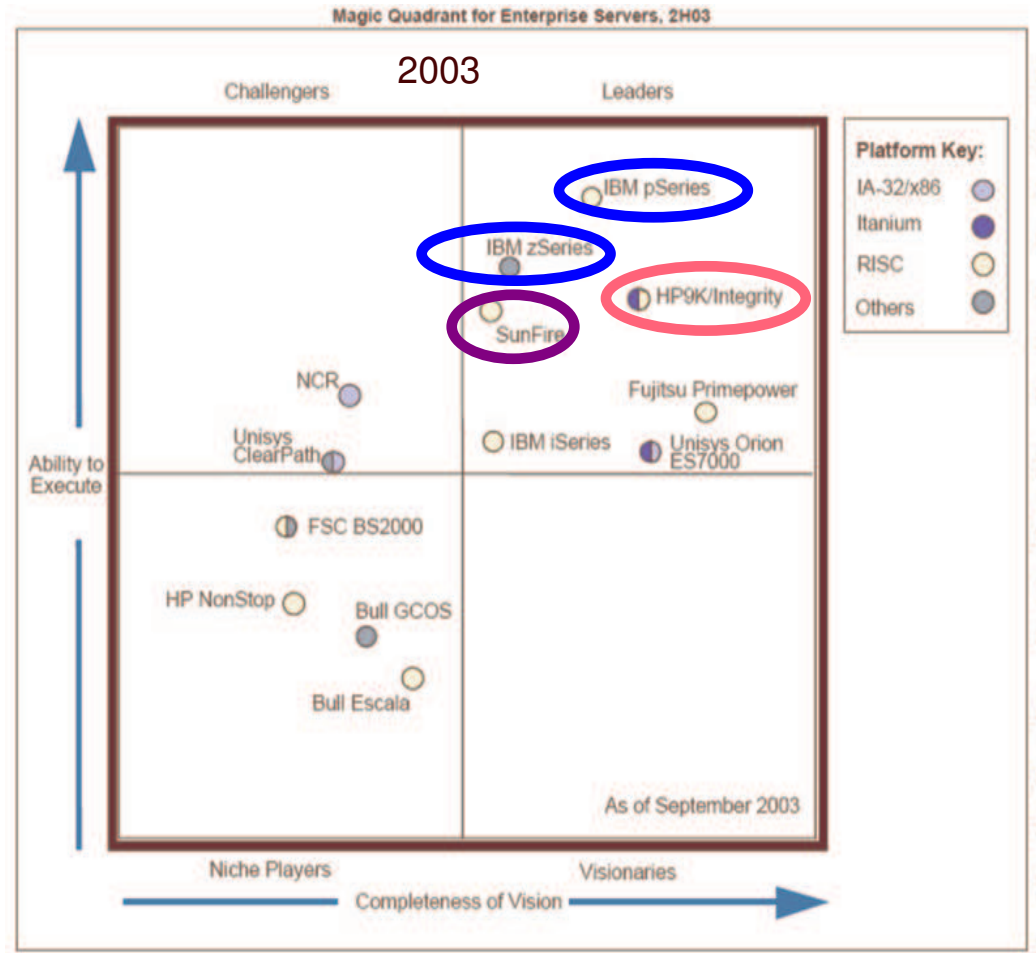
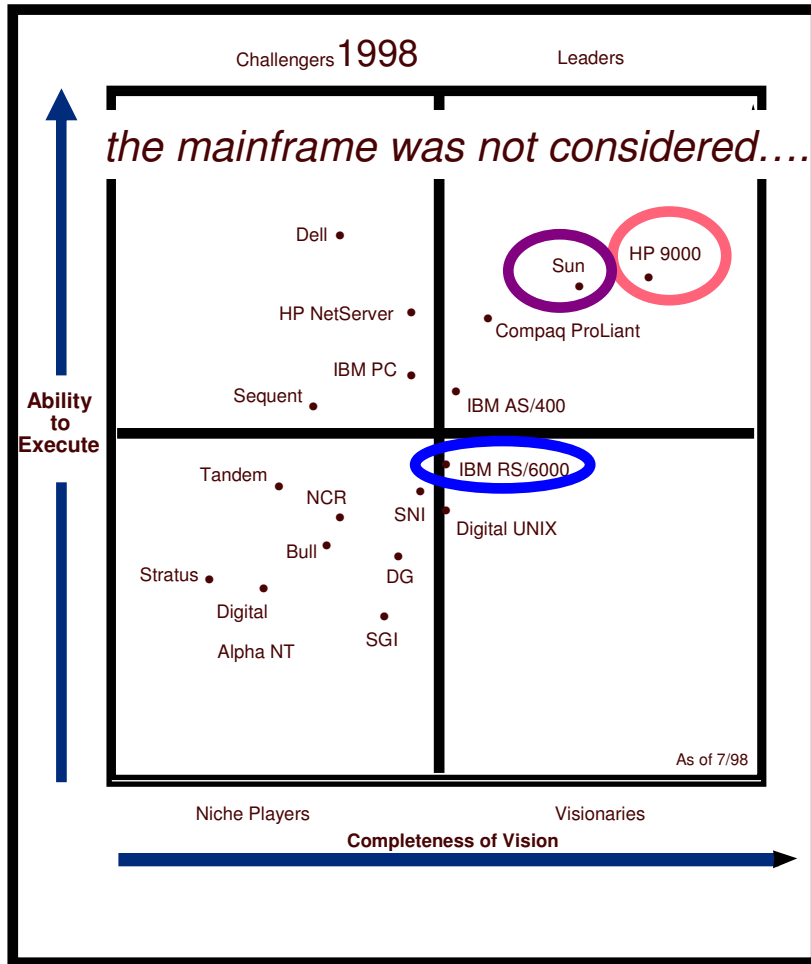
SLES = SUSE Linux Enterprise Server
RHEL = Red Hat Enterprise Linux



System z

Consultants view

Gartner's view on Server Vendor positioning – 1998 to 2003

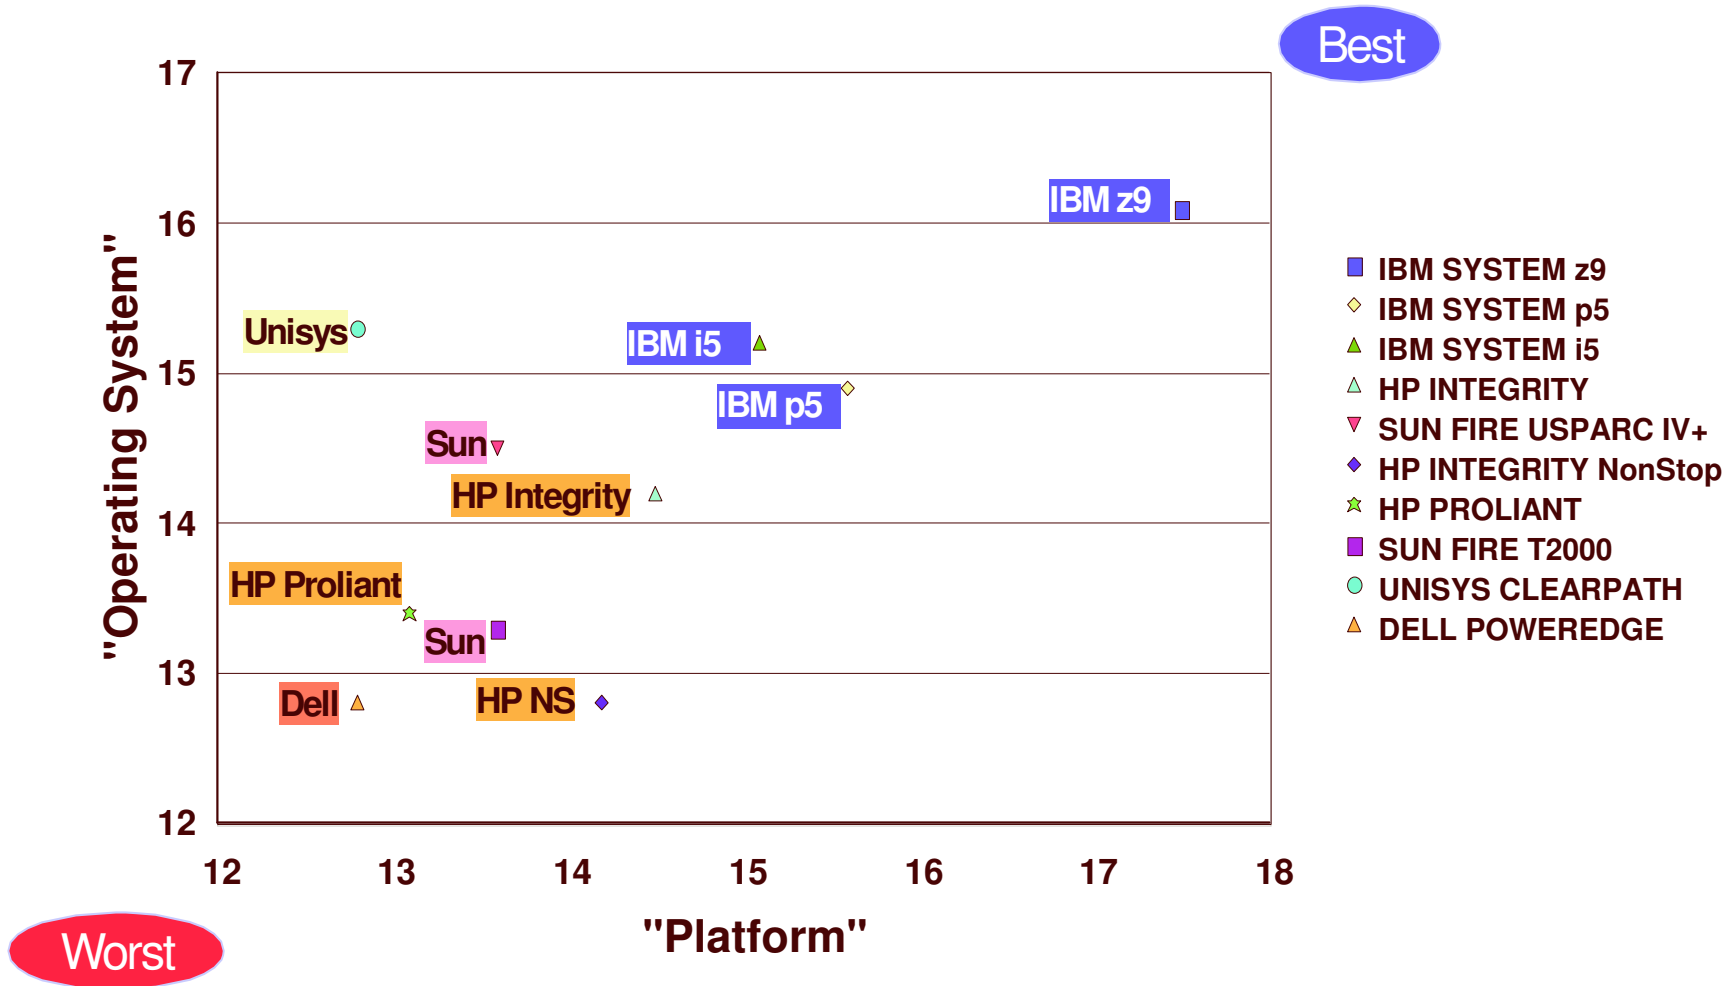


"Increasing ISV enthusiasm and market demand justifies an increase in ability to execute that now puts pSeries into the overall leading position on the Magic Quadrant, a position that Sun Microsystems has consistently held for more than three years"
Source: Gartner Group 10/02

Source: Gartner's Magic Quadrant reports July 1998, October 2002 and September 2003

Gartner's platform positioning 2q06

Top 10 Unweighted - Frame and Rack



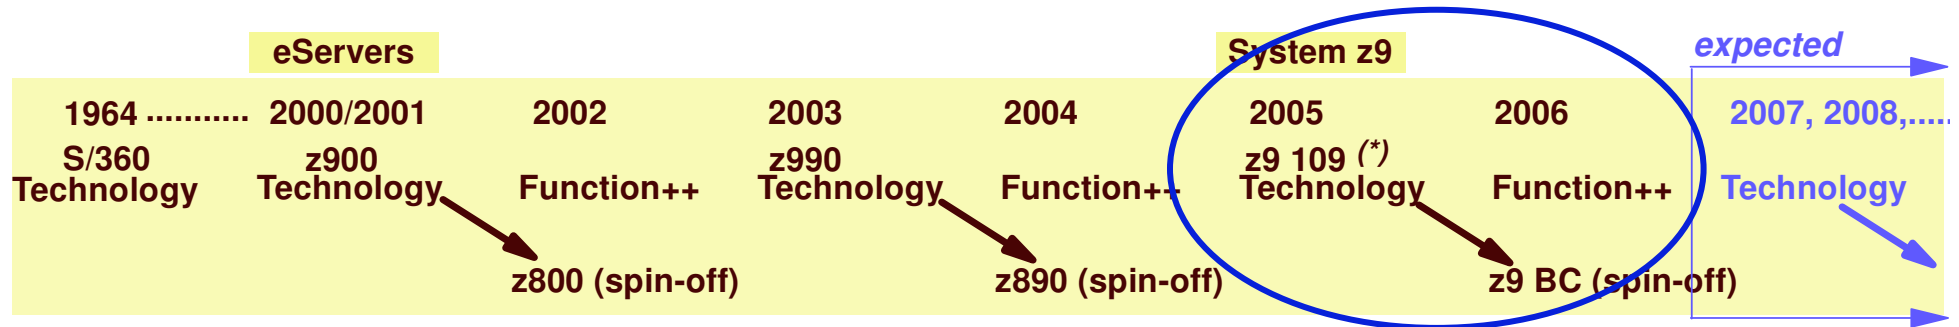
Source: Gartner, Server Scorecard Evaluation Model version 2, May 2006



System z

Future roadmap

System z Directions – Development and Delivery roadmap

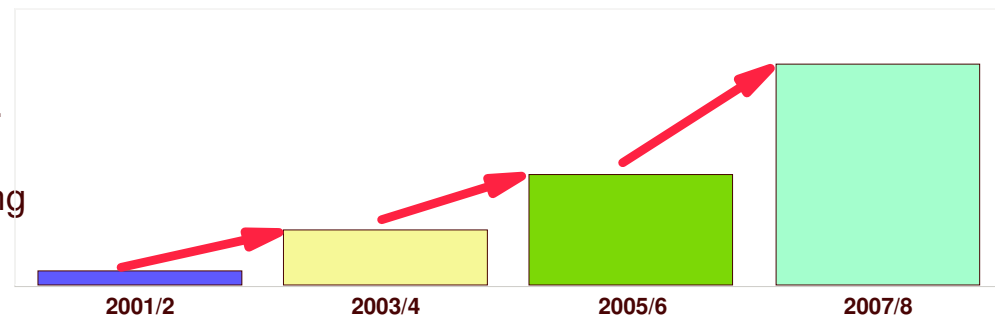


(*) has been renamed to z9 Enterprise Class (EC)

eServer Convergence, - moving mainframe like capabilities to other platforms, - but continued functional differentiation in respect to implementation of all aspects of "quality of service"

Addressing business requirements

- ▶ Continuous availability without complexity
- ▶ Improved business continuity
- ▶ Enhanced virtualization and resource mgnt.
- ▶ Increased capacity and performance
- ▶ Improved self-configuring and self-optimizing
- ▶ Reduced total cost of ownership
- ▶ Investment protection

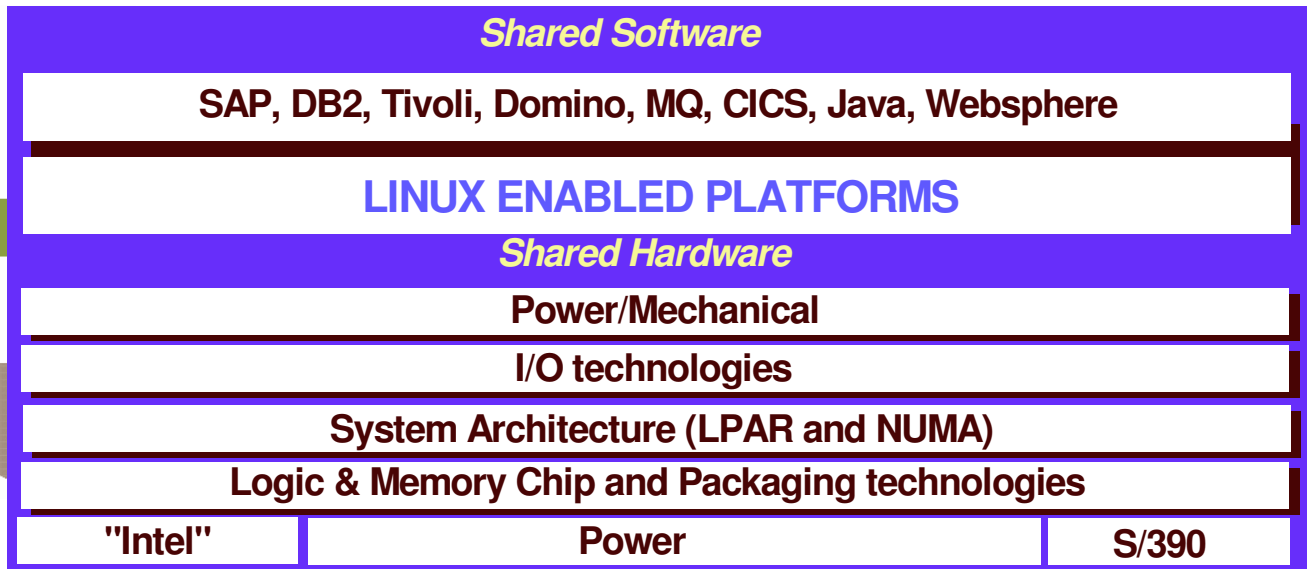
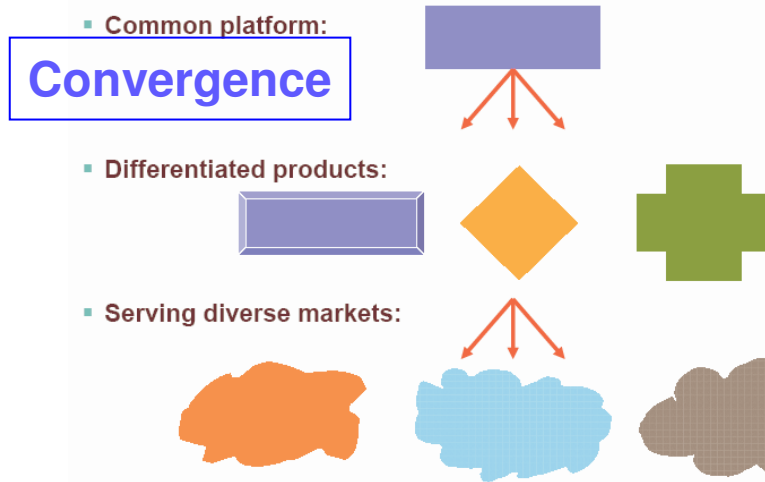


Scalability from smallest system to largest system of *refreshed* HW = 700 times
Within a System = 30 times
With Clustering = 30 times more capacity than a single system

IBM eServer convergence


to ensure

- building of high quality platform solutions to a very competitive hardware cost
- transfer of mainframe inspired technologies to enhance QoS and Flexibility
- system management convergence - "identical" functional interfaces to "attack" people cost issues
- and increase the ability to integrate end-to-end



while maintaining

- platform values in...
- support of different needs
- protecting investments

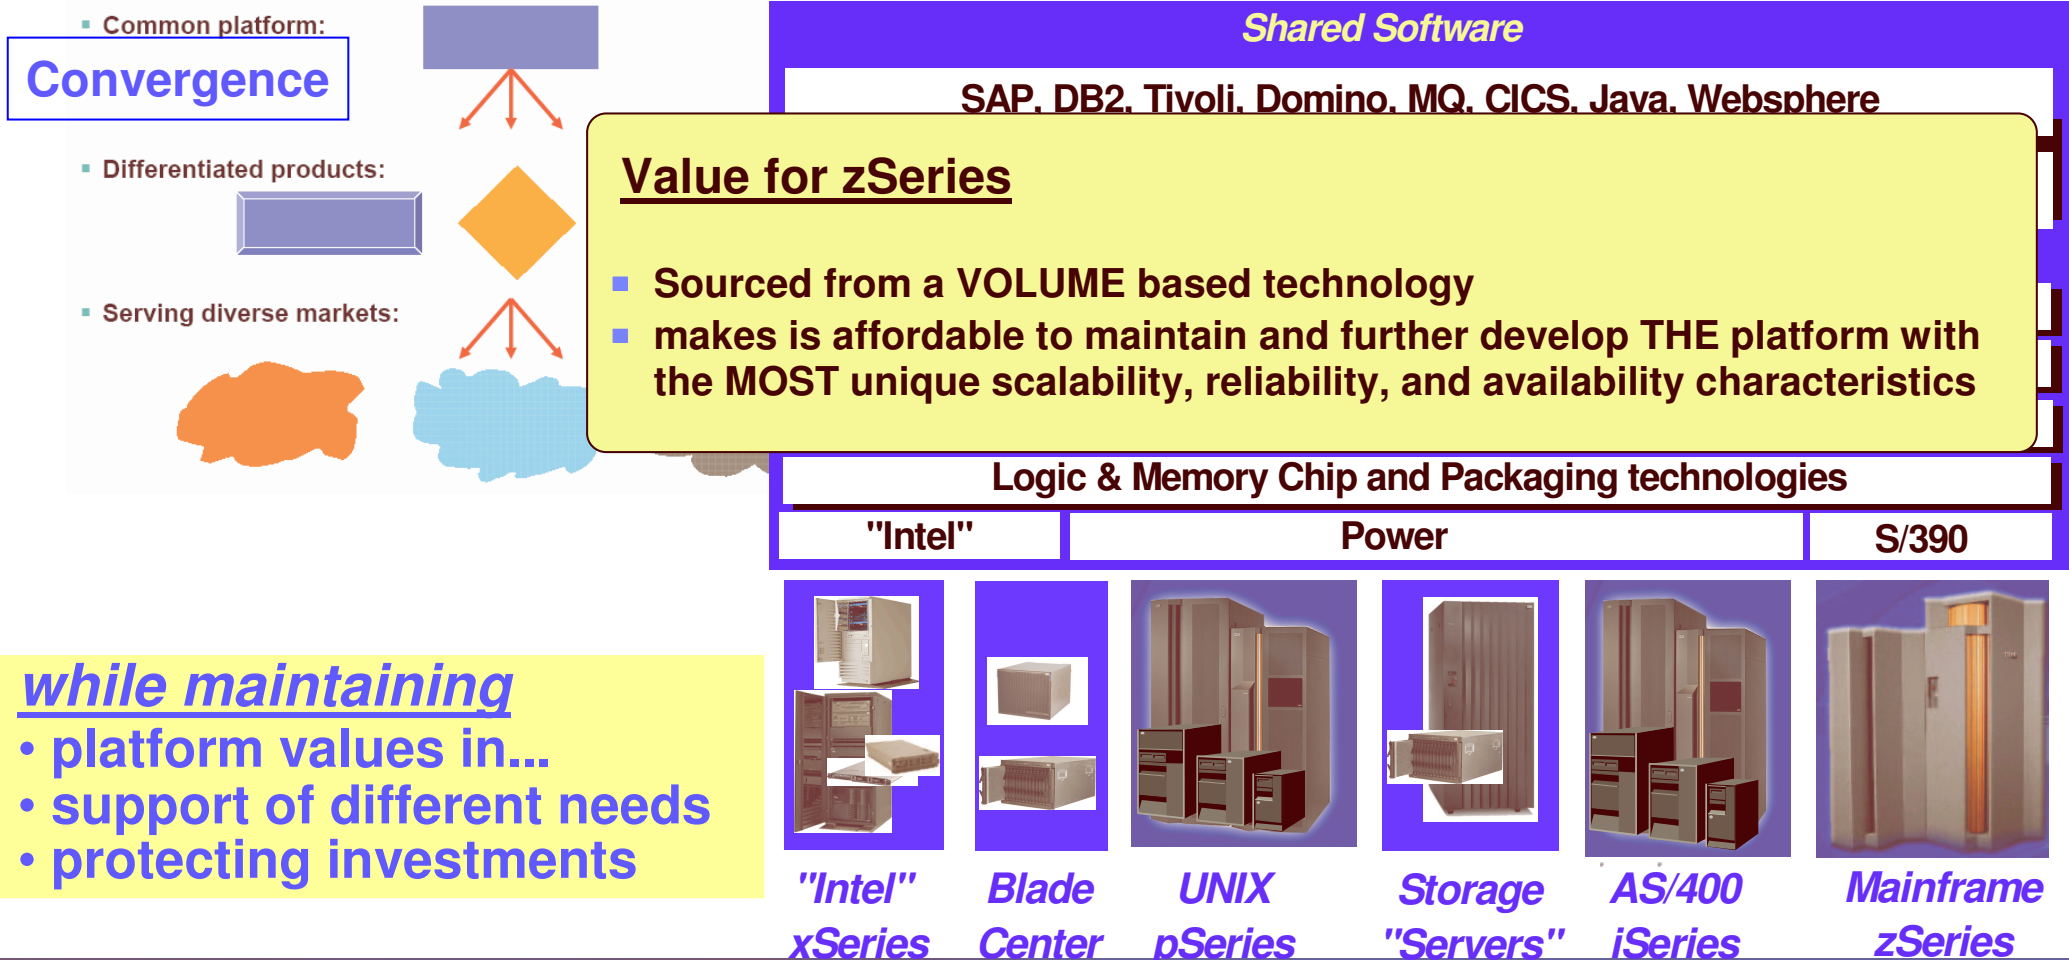


"Intel" xSeries Blade Center UNIX pSeries Storage "Servers" AS/400 iSeries Mainframe zSeries

IBM eServer convergence

to ensure

- building of high quality platform solutions to a very competitive hardware cost
- transfer of mainframe inspired technologies to enhance QoS and Flexibility
- system management convergence - "identical" functional interfaces to "attack" people cost issues
- and increase the ability to integrate end-to-end



while maintaining

- platform values in...
- support of different needs
- protecting investments



System z9 - Scalability and Upgrade strategy

Each new range has delivered:

- New function
- Improved performance
- Improved availability and serviceability
- Better price/performance
- Field Upgrades to protect investments

Solid roadmap
 2006/7/8 timeframe....
 1.72 Ghz to "x.x GHz"
 More engines per chip

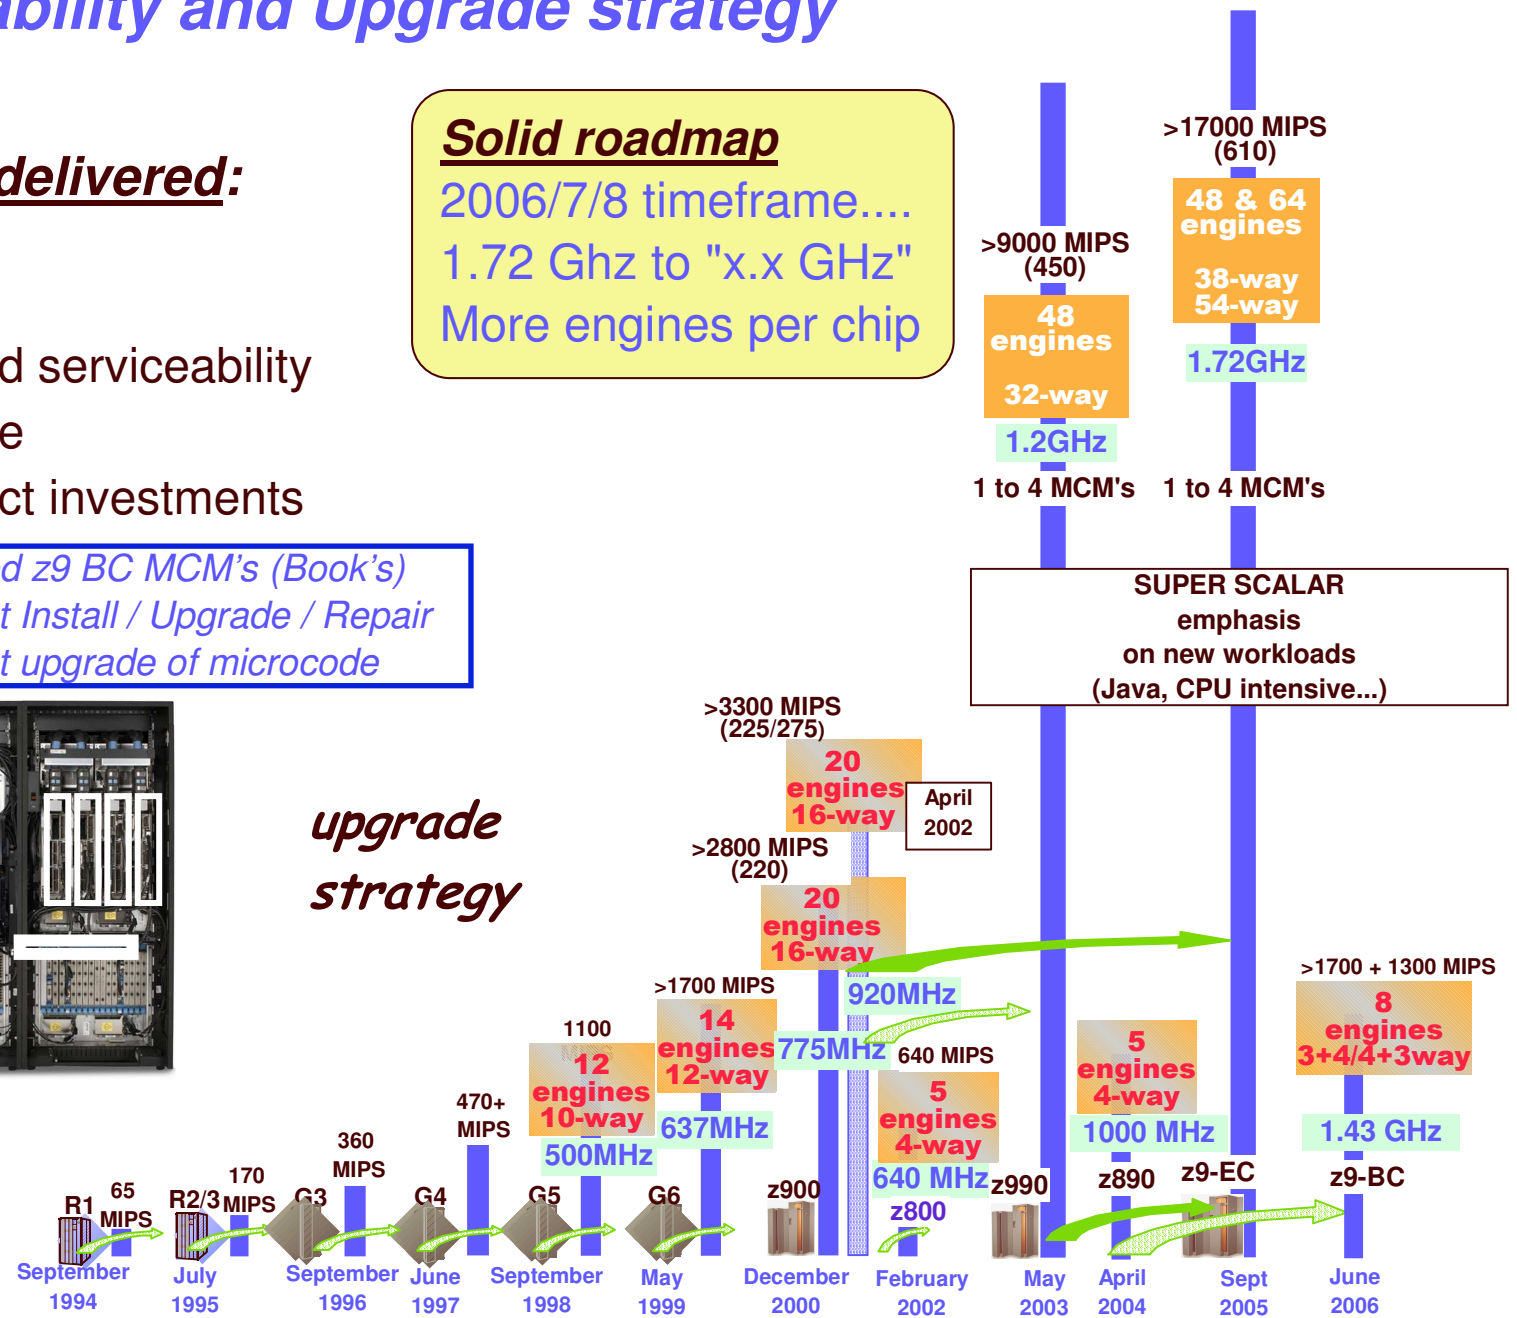
4 Advanced z9 BC MCM's (Book's)
 Concurrent Install / Upgrade / Repair
 Concurrent upgrade of microcode



System z9 EC MCM

95x95mm
 16 chipsites, 217 capacitors
3.9 billion transistors
 104 layer of glass ceramic
 476 meters of wire
more than 4000 MIPS
more than 400000++ TPC-C

upgrade strategy



SUPER SCALAR
 emphasis
 on new workloads
 (Java, CPU intensive...)

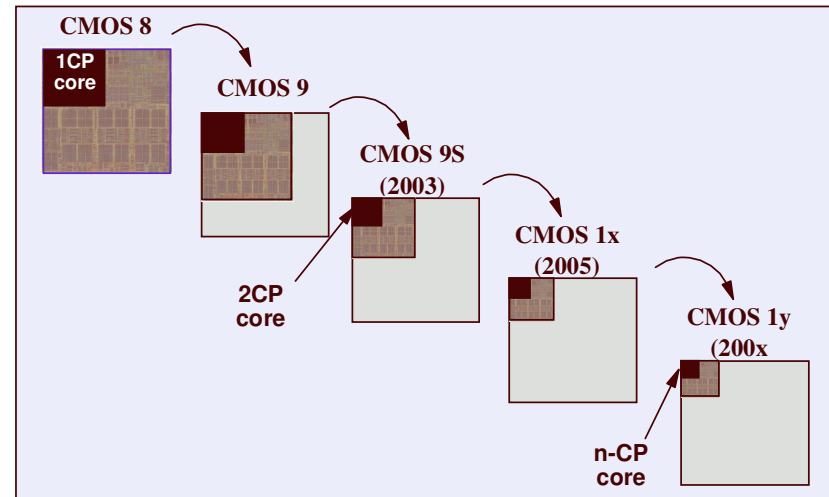
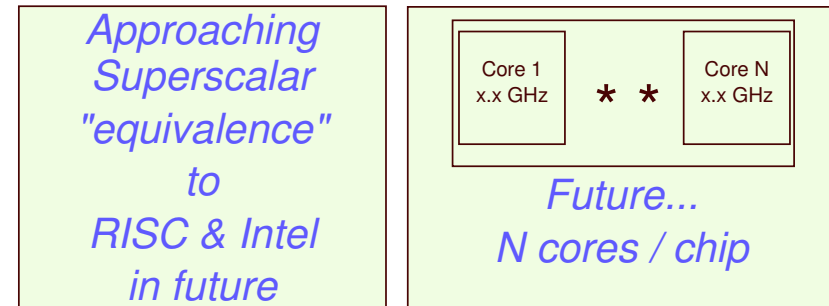
Chip technology evolution

Increasing "area" is available for SoC design

More cores (engines per chip) to be followed by more threads per core

16way systems on a chip are within reach !!!

- **Superscalar Capabilities**
- **Larger caches**
- **Memory controller functions**
- **Redundancy (processors, arrays, etc.)**
- **More concurrent threads**
- **Hardware accelerators**
 - more special assist Engines (accelerators/offload)
 - like System Assist Processors = SAPs
 - Crypto assists - special on System z
 - Compression & Translation - special on System z
 - Decimal Floating Instruction Hardware - opportunity
 - ...



Trends	2002	2005	2006/8	->>
transistors	174M	500M	1-1.3B	+3B
Mtr/cm2	65	142	350	863
Core	1-2	1-2	1-4	1 to +1K* Cores
GHz	0.7-2.1	1-3.5	1.2- 6	1.5-10

* (16w SMP x 64 chips, in 20cm x 20cm package)

Chip technology evolution


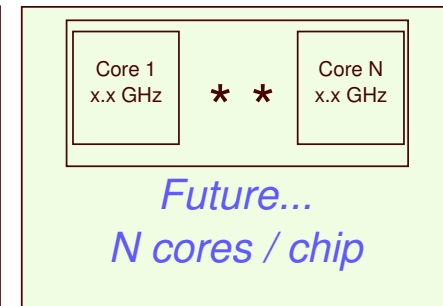
Increasing "area" is available for SoC design

More cores (engines per chip) to be followed by more threads per core

16way systems on a chip are within reach !!!


- **Superscalar Capabilities**
- **Larger caches**
- **Memory controller functions**
- **Redundancy (processors, arrays, etc.)**
- **More concurrent threads**

Approaching Superscalar "equivalence" to RISC & Intel in future



NOTICE:
 The cpu speed of z will approach the capacity of unix/intel for compute intensive work ...while System z will maintain its advantage in Bus and IO infrastructure

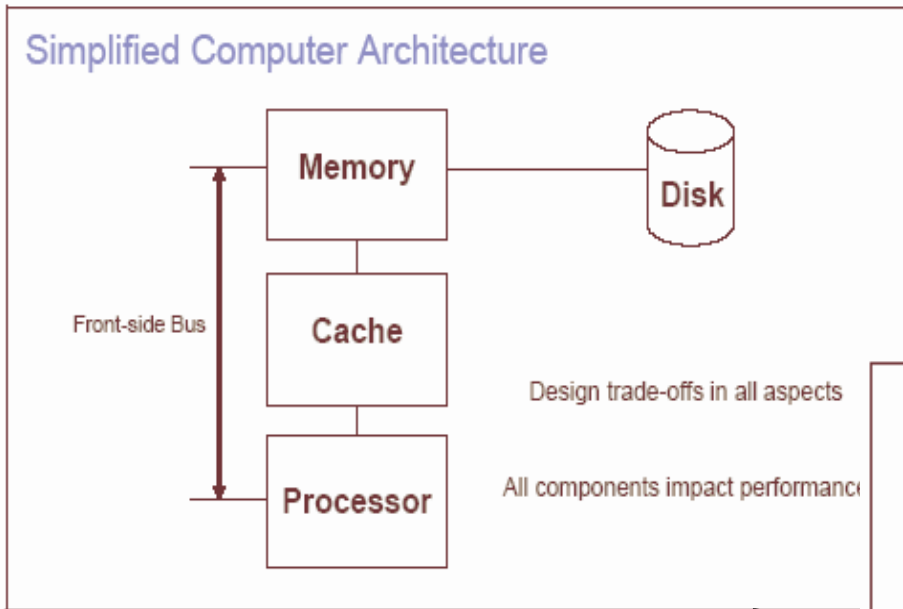
- **Crypto assists - special on System z**
- **Compression & Translation - special on System z**
- **Decimal Floating Instruction Hardware - opportunity**
- ...



Trends	2002	2005	2006/8	->>
transistors	174M	500M	1-1.3B	+3B
Mtr/cm2	65	142	350	863
Core	1-2	1-2	1-4	1 to +1K* Cores
GHz	0.7-2.1	1-3.5	1.2- 6	1.5-10

* (16w SMP x 64 chips, in 20cm x 20cm package)

Architecture - Design choices - Consequences



Speed

Multiple & Large workloads

RAS

Design choices affect "benchmark" performance
Mixed/WLM/Virtualization

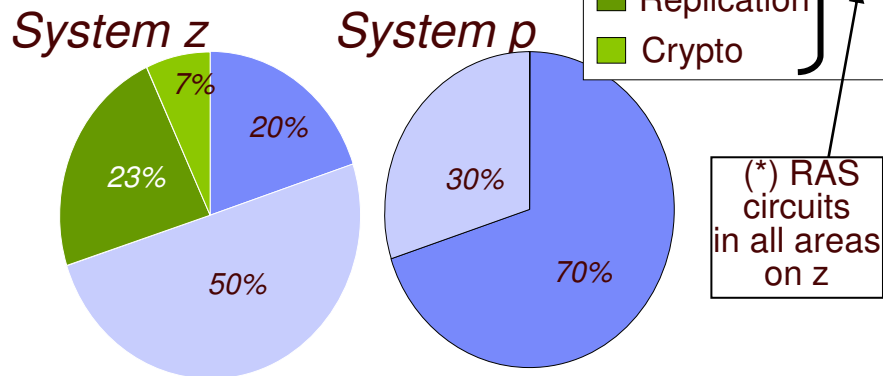
Replicated

- Working set(s) too large for cache
- Maximized Processor speed penalized
- RAS is valued in mixed workloads
- Minimized RAS is "penalized"
- "Fast" processor is under-utilized

Consolidated

- Cache contains multiple working sets
- Processor speed optimized by cache
- RAS Space is "valued"
- All of "slow" processor is used

Silicon Area of Core by Function



(*) RAS circuits in all areas on z

Comparison of processor core and L1 cache area usage

pSeries is the most "mainframe like" of the non-zSeries platforms



IBM System z9 Summary



**ENABLING BUSINESS.
A THROUGH Z.**

System z9 EC and BC – delivering new functions and features



- New IBM zIIP
- Granularity with entry one third the size of the 701
- Up to 54 configurable CPs
- Premier Availability server – with Enhanced Book Availability, RII and Enhanced Driver Maintenance
- MIDAW Facility
- FICON Express4
- Enhanced CPACF and Crypto Express2
- ATM/POS remote key loading
- Administrative On/Off CoD test
- Sub-capacity CBUs



- New low entry model
- New IBM zIIP
- Extreme Granularity
- Up to 7 PUs
- 37% more uni processor, up to 64 GB memory, 170% more bandwidth
- Sub-capacity CBUs and Administrative On/Off CoD Test
- Enhanced Driver Maintenance and RII
- MIDAW Facility and MSS
- NPIV and IPV6 Support for HyperSockets OSA-Express2 OSN (OSA for NCP)
- Enhanced CPACF with AES, PRNG and SHA-256 and Configurable Crypto Express2
- Temporary state changes allowed and new test/training option for On/Off CoD

Now there is a System z9 for everyone . . .



. . . which one is right for you?