



## G02

IBM System z9 Business Class  
*-Overview and Update 2007*

**John Hughes – Washington Systems Center**

**IBM System z Expo**

September 17-21, 2007  
San Antonio, TX



© IBM Corporation 2007

2007 IBM System z Expo



Washington Systems Center

**G02**  
**IBM System z9 Business Class**  
*-Overview and Update 2007*

John Hughes  
Washington Systems Center  
[jjhughes@us.ibm.com](mailto:jjhughes@us.ibm.com)

Advanced Technical Support

© 2003 IBM Corporation

## Trademarks

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

AIX*	Geographically Dispersed Parallel Sysplex	OS/390*	VSE/ESA
APPN*	HiperSockets	Parallel Sysplex*	VTAM*
CICS*	HyperSwap	PR/SM	WebSphere*
DB2*	IBM*	Processor Resource/Systems Manager	z/Architecture
DB2 Connect	IBM eServer	RACF*	z/OS*
DirMant	IBM e(logo)server*	Resource Link	z/VM*
DRDA*	IBM logo*	RMF	z/VSE
Distributed Relational Database Architecture	IMS	S/390*	zSeries*
e-business logo*	InfoPrint*	Sysplex Timer*	
ECKD	Language Environment*	System z	
Enterprise Storage Server*	MQSeries*	System z9	
ESCON*	Multiprise*	TotalStorage*	
FICON*	NetView*	Virtualization Engine	
GDS*	On demand business logo	VSE/ESA	

\* 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 registered trademark of Linus Torvalds in the United States, other countries, or both.

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

Microsoft, Windows and Windows NT are registered trademarks of Microsoft Corporation.

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

SET and Secure Electronic Transaction are trademarks owned by SET Secure Electronic Transaction LLC.

\* 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.

Review of z9 BC

Washington Systems Center IBM


## IBM System z family

### IBM eServer zSeries z990 z990 (2084) (WdFm in EMEA)




- Announced 5/03 – first zSeries Superscalar Server with up to 48 PUs
- 4 models – Up to 32-way
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 256 GB
- Channels
  - Four LCSSs
    - Up to 1024 ESCON channels
    - Up to 240 FICON Express2 channels
  - Token-Ring, GbE, 1000BASE-T Ethernet
  - Coupling Links
- Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions
- Operating Systems
  - z/OS, z/VM, VSE/ESA, z/VSE, TPF, z/TPF, Linux on zSeries

### IBM eServer zSeries 890 z890 (2086) (WdFm in EMEA)




- Announced 4/04 – zSeries Superscalar Server with 5 PUs
- 1 model – Up to 4-way
- 28 capacity settings
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 32 GB
- Channel
  - Two LCSSs
    - Up to 420 ESCON channels
    - Up to 80 FICON Express2 channels
  - Networking Adapters (OSA)
  - Coupling Links
- Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions
- Operating Systems
  - z/OS, z/OS.e, z/VM, VSE/ESA, z/VSE, TPF, z/TPF, Linux on zSeries

### IBM System z9 z9 EC (2094)



- Announced 7/05 - Superscalar Server with up to 64 PUs
- 5 models – Up to 54-way
- Granular Offerings for up to 8 CPs
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 512 GB
- Channels
  - Four LCSSs
    - Multiple Subchannel Sets
    - MIDAW facility
    - 63.75 subchannels
    - Up to 1024 ESCON channels
    - Up to 336 FICON channels
    - Enhanced FICON Express2 and 4
    - 10 GbE, GbE, 1000BASE-T
  - Coupling Links
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 60 logical partitions
- Enhanced Availability
- Operating Systems
  - z/OS, z/VM, VSE/ESA, z/VSE, TPF, z/TPF, Linux on System z9

### IBM System z9 z9 BC (2096)



- Announced 4/06 - Superscalar Server with 8 PUs
- 2 models – Up to 4-way
- High levels of Granularity available
- 73 Capacity Indicators
- PU (Engine) Characterization
  - CP, SAP, IFL, ICF, zAAP, zIIP
- On Demand Capabilities
  - CUoD, CIU, CBU, On/Off CoD
- Memory – up to 64 GB
- Channels
  - Two LCSSs
    - Multiple Subchannel Sets
    - MIDAW facility
    - 63.75 subchannels
    - Up to 420 ESCON channels
    - Up to 112 FICON channels
    - Enhanced FICON Express2 4 Gbps
    - 10 GbE, GbE, 1000BASE-T
  - Coupling Links
- Configurable Crypto Express2
- Parallel Sysplex clustering
- HiperSockets – up to 16
- Up to 30 logical partitions
- Enhanced Availability
- Operating Systems
  - z/OS, z/OS.e, z/VM, VSE/ESA, z/VSE, TPF, z/TPF, Linux on System z9

5
Advanced Technical Support
© 2003 IBM Corporation

Washington Systems Center IBM

## z9 BC – Under the covers


Internal Batteries (optional)

Power Supplies

I/O Cage

Fiber Quick Connect Feature (optional)

A Frame



Front View

Single Processor Book and Memory

CEC Cage

STI Connectors


Support Elements (gate with Laptops swung open)

6
Advanced Technical Support
© 2003 IBM Corporation

## IBM System z9 BC Processor and Memory

## 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

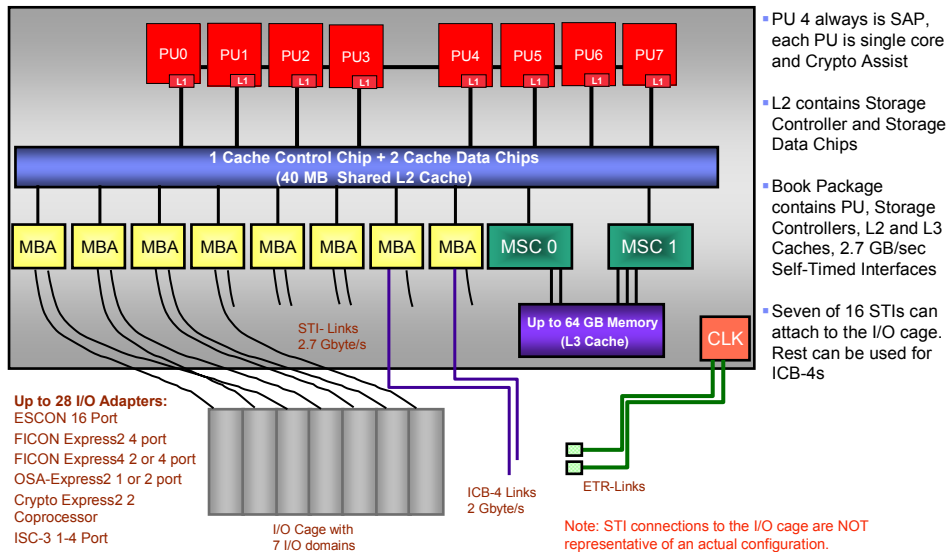
### z9 BC Processor Book Layout

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 8 PU Logical Structure



- PU 4 always is SAP, each PU is single core and Crypto Assist
- L2 contains Storage Controller and Storage Data Chips
- Book Package contains PU, Storage Controllers, L2 and L3 Caches, 2.7 GB/sec Self-Timed Interfaces
- Seven of 16 STIs can attach to the I/O cage. Rest can be used for ICB-4s

Up to 28 I/O Adapters:  
 ESCON 16 Port  
 FICON Express2 4 port  
 FICON Express4 2 or 4 port  
 OSA-Express2 1 or 2 port  
 Crypto Express2 2  
 Coprocessor  
 ISC-3 1-4 Port


I/O Cage with 7 I/O domains

Note: STI connections to the I/O cage are NOT representative of an actual configuration.

## Technology evolution with specialty engines

- Building on a strong track record of technology innovation with specialty engines, IBM introduces the System z9 Integrated Information Processor

Centralized data sharing across mainframes



**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 (zIIP) 2006**

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

## System z9 PU Characterization

- The type of Processor Units (PUs) that can be ordered on System z9:
  - Central Processor (CP)
    - Provides processing capacity for z/Architecture™ and ESA/390 instruction sets
    - Runs z/OS, z/VM, VSE/ESA, z/VSE, TPF4, z/TPF, Linux for System z and Linux under z/VM or Coupling Facility
    - z9 EC has Capacity Marker features NOT Unassigned CP features
  - IBM System z Application Assist Processor (zAAP)
    - Under z/OS, the Java Virtual Machine (JVM) assists with Java processing to a zAAP
  - IBM System z9 Integrated Information Processor (zIIP)
    - Provides processing capacity for selected workloads e.g., DB2 for z/OS V8 workloads executing in SRB mode
  - Integrated Facility for Linux (IFL)
    - Provides additional processing capacity for Linux workloads
  - Internal Coupling Facility (ICF)
    - Provides additional processing capacity for the execution of the Coupling Facility Control Code (CFCC) in a CF LPAR
  - Optional System Assist Processors (SAP)
    - SAP manages the start and ending of I/O operations for all Logical Partitions and all attached I/O

## z9 BC Memory Subsystem Offering and Assignment

- **Memory card sizes: 2 GB, 4 GB and 8 GB**
- **Memory cards per System: 4 or 8 of same card size**
- **8 GB increments**
- **Maximum memory size: 64 GB per System**
- **All memory is shared within the server**
- **Cache structure allows uniform memory access**
- **eConfig specifies LICCC**
- **LPAR allocates the amount of memory purchased by the customer**

**Remember:**

PR/SM manages all of this as a single large SMP. A single logical partition can use resources located in multiple books (multi-Book Systems e.g. z9 EC).

## 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

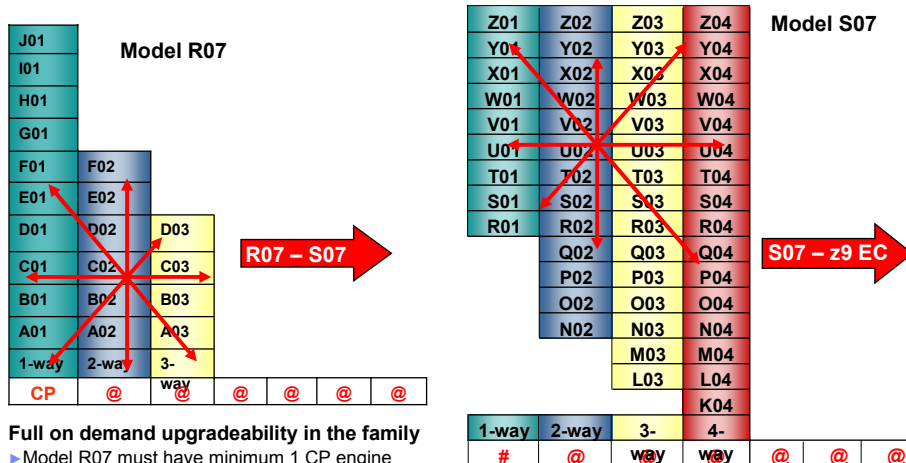
**Green** - Concurrent upgrade

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

## z9 BC Improved granularity and scalability





## System z9 Capacity Planning in a nutshell



**Don't use "one number" capacity comparisons!**  
 Work with IBM technical support for capacity planning!  
 Customers can now use zPCR

## Upgrade Paths

- 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, IFLs, ICFs, zAAPs or zIIPs
  - New options for reconfiguring specialty engines if the business demands it
  - New options for changing On/Off CoD configurations
  - Subcapacity CBU engines



## IBM System z9 BC Channel Subsystem


### z9 BC I/O Overview

- **I/O Enhancements**
  - Up to 28 FICON Express, FICON Express2, FICON Express4 features
    - 4 channels/feature FICON/FCP
    - 1, 2, 4 Gbps auto-negotiated. 4 Gbps for FICON Express4
  - Modified Indirect Data Address Word (MIDAW) facility
  - Multiple (2) Subchannel sets (MSS)
    - Increase to 63.75K Subchannels for Set-0
  - Up to 16 x 2.7GB STI's (7 STIs max for the single I/O cage. Possible to use remaining STIs for ICB-4s)
- **Storage Area Networks (SANs) enhancements**
  - N\_Port ID Virtualization
  - Program Directed re-IPL
  - FICON Link Incident Reporting
- **Networking enhancements**
  - HiperSockets IPv6
  - OSA-Express2 1000BASE-T Ethernet
  - OSA-Express2 OSN (OSA for NCP support)
  - GARP VLAN management (GRVP)



## z9 BC Compatibility or Exploitation ready OS

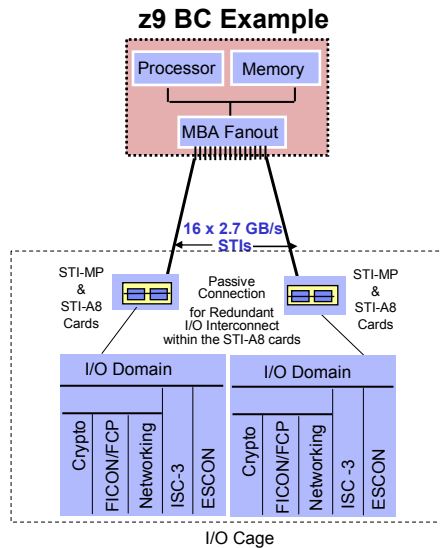
- Up to 2 Logical Channel Subsystems (LCSS)
- Up to 15 Logical Partitions (LPARs) per Logical Channel Subsystem
- Up to 256 CHPIDs per LCSS
- A Logical Partition uses I/O from a single LCSS
- 2 digit LPAR ID
- 1 digit MFID



256 Channel limitation per OS image remains unchanged!

## 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

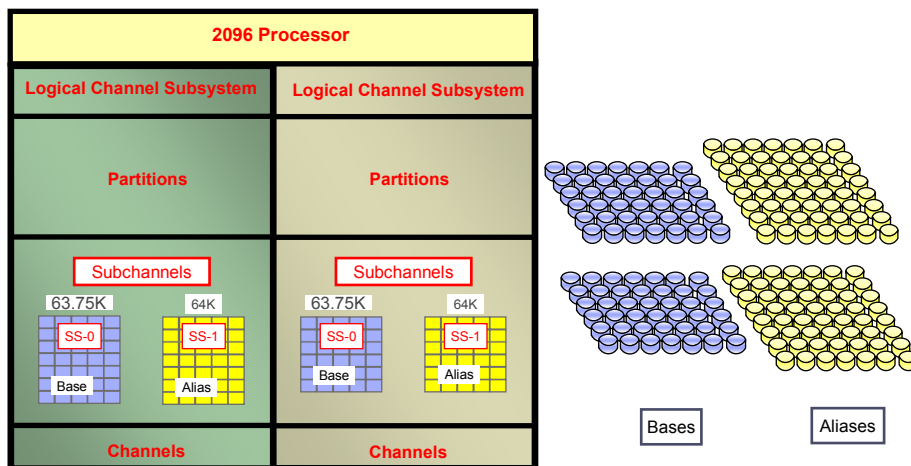


## Modified Indirect Data Address Word (MIDAW) Facility

- **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

## System z9 – Multiple Subchannel Sets per LCSS

Each 2096 Logical Channel Subsystem can have up to two sets of Subchannels



## IBM System z9 BC Connectivity

## System z9 Channel Type and Crypto Overview

- FICON/FCP
  - FICON Express4
  - FICON Express2 (carry forward on upgrade)
  - FICON Express (carry forward on upgrade)
- 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
    - Fast Ethernet
  - HiperSockets
- Coupling Links
  - ISC-3 (Peer mode only)
  - ICB-3, ICB-4
  - IC
- ESCON
- Crypto
  - Crypto Express2
    - 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.

New feature codes for ICB cables when connecting  
to a z9 BC System with 'non-raised floor' option.

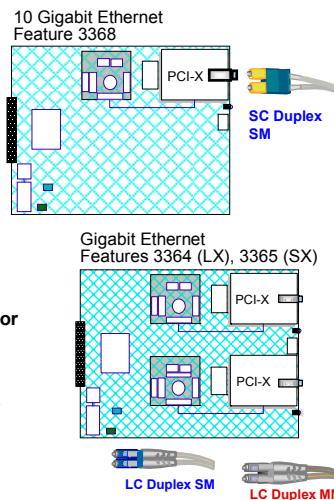


## z9 BC I/O Connectivity

- 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 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
- 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 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




## System z9 OSA-Express2 1000BASE-T Ethernet

- 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
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



## z9 BC 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
  - Two or Four 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

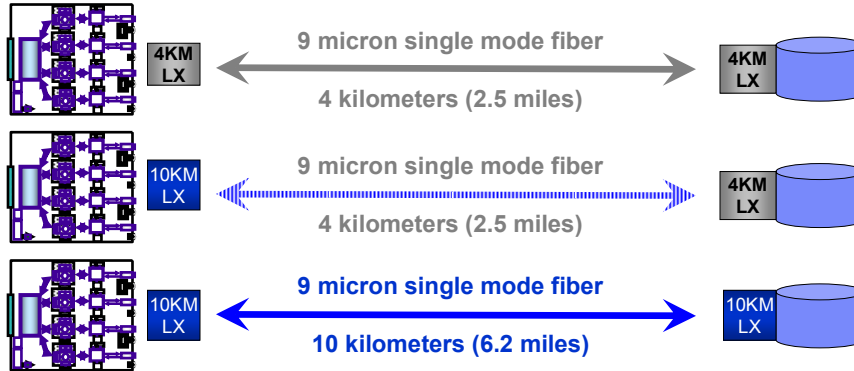


FC 3321	FICON Express4 10 KM LX
FC 3322	FICON Express4 SX
FC 3323	FICON Express4-2C 4KM LX (2 ports)
FC 3324	FICON Express4 4KM LX

## 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)



## System z9 FICON Express2

- Four channels per feature of LX or SX
  - Maximum 28, 112 channels
- 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](http://www.ibm.com/servers/eserver/zseries/connectivity)





## IBM System z9 BC Availability

## z9 BC Availability Enhancements

- **Enhanced driver maintenance**
- **Concurrent MBA fanout add/replacement**
- **Redundant I/O Interconnect**
- **Enhanced ECC memory**
- **Dynamic oscillator switchover**
- **Pluggable optics for FICON Express4**

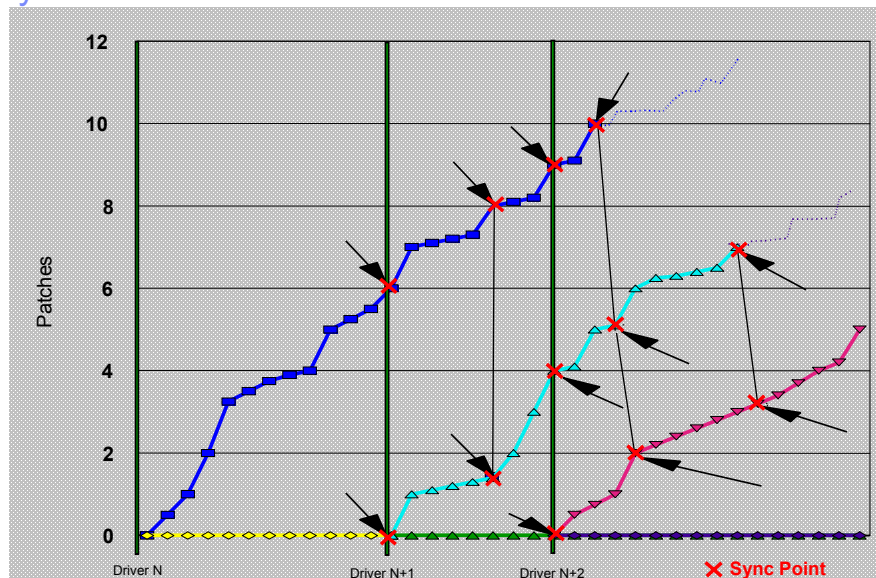


## System z9 Enhanced Driver Maintenance

- **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 crossover 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 cause a disruptive outage**

\*Requires proper configuration and planning for exploitation

## System z9 Enhanced Driver Maintenance

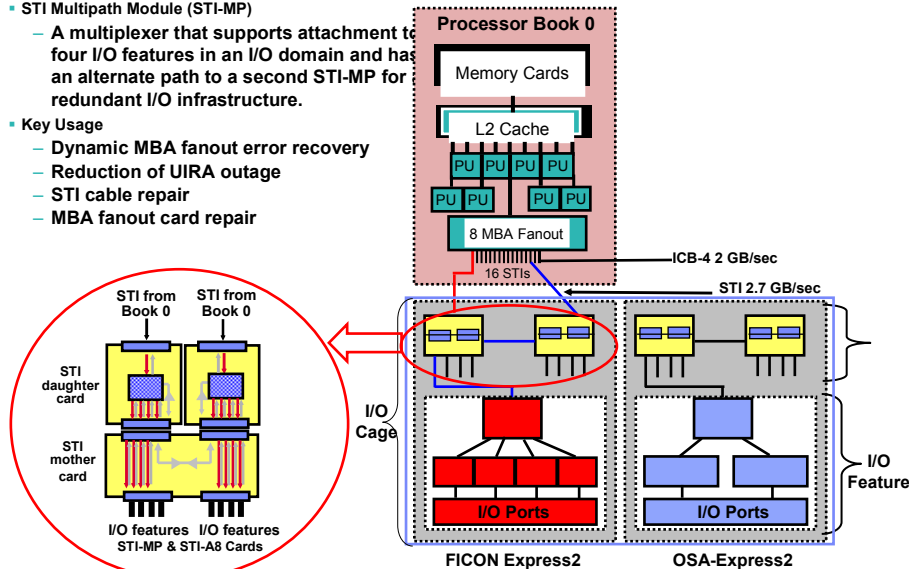


## System z9 Enhanced Driver Maintenance limitations

- **Specific complex code changes may cause 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) Capability

- **STI Multipath Module (STI-MP)**
  - A multiplexer that supports attachment to four I/O features in an I/O domain and has an alternate path to a second STI-MP for redundant I/O infrastructure.
- **Key Usage**
  - Dynamic MBA fanout error recovery
  - Reduction of UIRA outage
  - STI cable repair
  - MBA fanout card repair



## System z9 Concurrent Upgrade – CUoD

- **CUoD – Capacity Upgrade on Demand - Standard machine function**
  - Concurrent addition of CPs ('model capacity' upgrade), IFLs, ICFs, zAAPs, zIIPs, and memory
  - LIC enable additional 8 GB memory increments (for certain circumstances only)
  - Concurrent model upgrade (capacity setting) to add active PUs and memory
  - Downgrades are nondisruptive
  - **Note: I/O feature adds and removes are also nondisruptive but not really "CUoD"**

**Notes:**

1. CUoD is built on a base of concurrent "hot-plug" maintenance
2. I/O feature adds and removes are also nondisruptive but not really "CUoD"
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.

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

## System z9 Concurrent Upgrade – Customer Controlled

- **On/Off Capacity on Demand - Temporary upgrade**
  - Nondisruptive temporary addition of CPs, IFLs, ICFs zAAPs and zIIPs
  - "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., memory)
    - 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 - Temporary emergency capacity upgrade**
  - Nondisruptive temporary addition of CPs, IFL, ICF, zAAPs, zIIPs 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

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

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

- For customers who have a requirement for short-term additional capacity, temporary, nondisruptive addition of CPs, IFLs, ICFs, ZAAPs and/or zIIPs
  - Requires a contract between IBM and the customer and features on the machine
    - CIU Enablement (#9898) and On/Off CoD Enablement (#9896)
  - Add temporary CPs, IFLs and/or ICFs up to the number of base active CPs plus base unassigned CP capacity, base active and unassigned IFLs, and base ICFs, respectively
  - Add temporary ZAAPs up to the quantity of base ZAAPs
    - Temporary ZAAPs plus base ZAAPs may not exceed the quantity of temporary CPs plus base active CPs plus base unassigned CP capacity.
  - Add temporary zIIPs up to the quantity of base zIIPs.
    - Temporary zIIPs plus base zIIPs may not exceed the quantity of temporary CPs plus base active CPs plus base unassigned CP capacity.
  - Billing through one or more of the following features:
    - On/Off CoD Active CP-Day for base CP speed, On/Off CoD Active IFL Day, On/Off Active ICF-Day, or On/Off CoD Active zAAP-Day, or On/Off CoD Active zIIP-Day
    - HW charges calculated monthly and billed in arrears based on 24-hour usage period
- Nondisruptive temporary upgrade and downgrade
  - Order On/Off CoD upgrade on Resource Link
  - Download and install upgrade On/Off CoD upgrade LIC from RSF using an HMC
  - Deactivate On/Off CoD upgrade using an HMC (same action as CBU undo)

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

## System z9 On/Off CoD Enhancements – 1

- Full Function Test
  - One no-charge test per Server contract.
  - Enables customer to use On/Off CoD function and install/remove additional capacity
  - 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.
- Introducing special Administrative On/Off CoD Test
  - Enables customers to order 'zero' quantity features via Resource Link for:
    - Pre-staging On/Off CoD order
    - Activating and deactivating 'zero' quantity On/Off CoD
  - To allow customer staff to order/test/rehearse/document whole On/Off CoD process without incurring any cost. Zero quantity features = zero cost
  - Unlimited number of tests
  - No time period restrictions

## System z9 On/Off CoD Enhancements – 2

- **Additional flexibility for OOCoD function**
  - With OOCoD already activated, customer can add additional OOCoD capacity without having to restore system to 'purchased' capacity – as today
  - Customer can keep adding or removing capacity using OOCoD without have to go back to 'purchased' capacity using OOCoD
    - Linked to purchased capacity rule i.e. maximum capacity customer can have with OOCoD is 2 x of purchased capacity
    - Limit controlled by Capacity Marker feature
  - Customer charged for additional capacity on 24 hour basis
  - If customer increases capacity multiple times during a 24 hour period, charge applies to the highest amount of capacity activated
  - Total of 60 minutes grace period allowed within a 24 hour OOCoD "day"
- **Enhancement for Automation Code**
  - New API being provided
  - Currently Operations have to select OOCoD from the HMC
  - New API is designed to enable customers to modify automation code to use SNMP to send command with order number to HMC to enable activation
  - Can provide improved flexibility of operation of OOCoD function

## 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, ICFs, zAAPs, or zIIPs
    - 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 IFLs, ICFs, zAAPs or zIIPs to be added. For CPs, if within the same 'speed' it's the number of CPs to be added. If going to a different 'speed' CP, it's the TOTAL number of CPs which will be active.
    - FC: 7880 – 7895 options for CBU CP
    - FC: 7821 for CBU IFL
    - FC: 7822 for CBU ICF
    - FC: 7824 for CBU zAAP
    - FC: 7825 for CBU zIIP
  - 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, ICFs, zAAPs, or zIIPs
  - 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

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

## z9 BC CBU Enhancements

- CBU for Specialty Engines
  - CBU is available for CPs, IFLs, ICFs, zAAPs and zIIPs
  - FULL size specialty engines for CBU
  - During CBU can't reduce engine count or convert engine types of the base machine
- CBU for CPs
  - No change for FULL size CPs
  - CBU for sub-capacity CPs
    - In CBU mode CP engine count must be equal or greater than purchased CPs
  - Example below (xx = number of CPs) shows valid directions of CBU 'paths' for CI U02

	1-way	2-way	3-way	4-way
CI Txx	34	66	95	124
CI Uxx	38	73	106	138
CI Vxx	42	82	119	155
CI Wxx	47	92	134	174

- CBU Pricing
  - Cost is per CBU feature MSU values shown in above table
  - Its not always cost effective to have larger number of sub-capacity CPs
    - In the above example, if a CI U02 is CBUed to CI U04, its 2 CBU features, however if CBU goes to CI V04, its 4 CBU features

## z9 BC Concurrent PU Conversions

- Must order (characterize one PU as) a CP, an ICF or an IFL
- Concurrent processor upgrade is supported if PUs are available
  - Add CP, IFL, unassigned IFL, ICF, zAAP, zIIP or optional SAP
- Conversion of unassigned IFL to any other PU type direct is supported on the z9 BC (This is different than the z9 EC). Conversion to unassigned IFL is via a IFL

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	Yes	Yes	x	Yes	Yes	Yes
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.

## IBM System z9 BC Hardware Management Console (HMC) Trusted Key Entry (TKE) Workstation

## System z9 Hardware Management Console (HMC)

### ▪ Objective

- **Provide full function remote access via standard browsers**
- **Utilize user interface technologies which are common with other IBM products**
- **Minimize changes to the user interface**
- **Enhance security**
- **Support as many levels of machines as possible**



## Similarities to previous HMC

- **Similar User Interface**
  - Main UI as much the same as possible
  - Most tasks present the same panels, but in a Web-based form
- **Same capabilities**
  - Change management / phone server, automation access point, etc.
  - Most configuration data will carry forward in an upgrade
- **Same PC hardware as current HMCs/SEs**
- **Networks: Ethernet, Token-Ring (for server upgrades to System z9 and where already installed on HMCs from prior zSeries servers only)**

## Differences from previous HMC

- **Closed system**
  - No desktop access
  - No installation of (customer) applications
  - ESCON Director / Sysplex Timer® WinOS2 HMC application features are no longer supported. (Director/Timer console definition/single object operations still supported.)
  - Hardware Management Console Application Programming Interface libraries and sample programs can no longer be obtained from the HMC hard disk. They will be available from the IBM Resource Link Web site.
- **TCP/IP is the only supported communications protocol**
  - Auto-discovery uses IP multicast instead of NetBIOS
  - SNA protocol no longer supported
  - Asynch/LAN only supported call home method. SDLC no longer supported

## Differences (continued)

- **The following systems will be supported, but they must be upgraded to a new AROM level, and support will only be provided for the specified driver.**
  - 9672 G5, G6 AND Multiprise® 3000 (Driver 26)
  - z800/z900 (Driver 3G)
  - z890/z990 (Driver 55)
- **The following systems will not be supported:**
  - 9221
  - 967x G1-G4
  - Multiprise 2000 (Driver 24)
  - 7060 (Multiprise 3000)

## System z9 HMC Security aspects

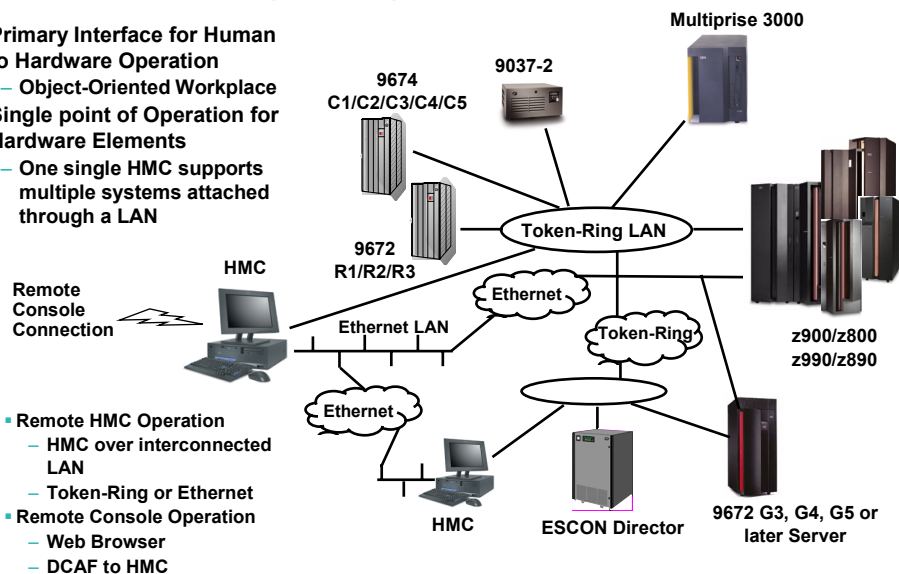
- **Remote support connections for service are initiated by the customer machine (HMC) to IBM Retain. With customer consent, IBM may connect to the system**
- **An RSF connection can be made via a phone connection or an Internet connection.**
  - **If it is via a phone connection, the protocol used is PPP and TCPIP. Data transfer is done using SSL encryption.**
    - Phone connections are made using ATT as the global service provider.
    - Provision is made for the use of local phone connections where available.
    - ATT system is designed to limit access for these connections to a restricted set of destination IP addresses. ATT system is designed so that general Internet access is not available through these connections.
  - **If it is via an Internet connection, the protocol used is TCPIP and data transfer is done using SSL encryption.**
  - **Internet connections are assumed to go through a customer firewall system before entering the global Internet.**
- **On either style connection, the IBM Retain system is designed to validate that the incoming requesting system is known and authorized**

## System z9 HMC Security aspects – continued

- Data transferred to IBM is for fixing the reported issue
- Data transferred from IBM to the customer system is encrypted and checked prior to use
- The HMC has an integrated firewall designed to disable all access by default opening only specific configured ports needed by the HMC functions.
- The HMC is designed to be a closed platform and it is designed so that there is no customer access to the underlying system.
- The HMC, by default, will not answer an incoming call via the modem. This can be changed by the customer if desired to allow remote backup operations
- All remote access is now encrypted using Secure Socket Layer (SSL) certificate based technology

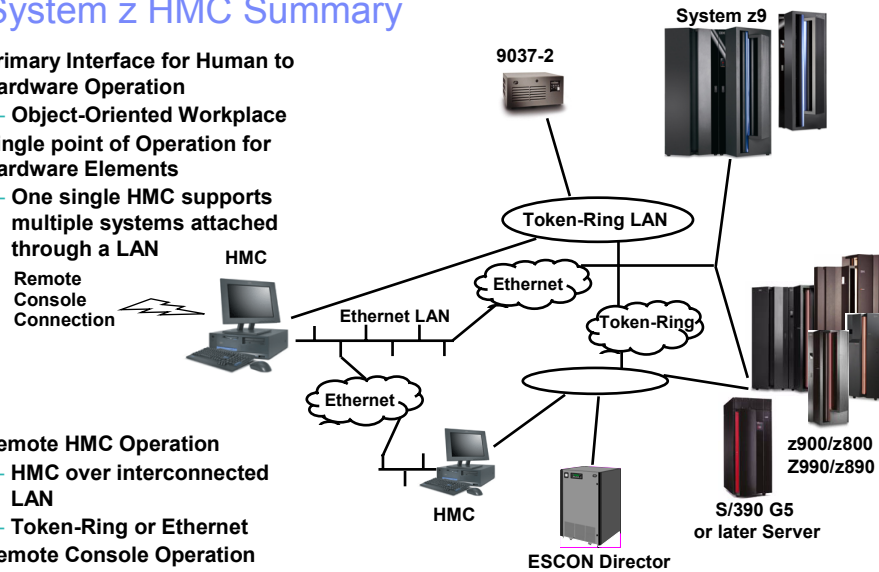
## HMC Summary pre-System z9

- Primary Interface for Human to Hardware Operation
  - Object-Oriented Workplace
- Single point of Operation for Hardware Elements
  - One single HMC supports multiple systems attached through a LAN



## System z HMC Summary

- **Primary Interface for Human to Hardware Operation**
  - Object-Oriented Workplace
- **Single point of Operation for Hardware Elements**
  - One single HMC supports multiple systems attached through a LAN
- **Remote HMC Operation**
  - HMC over interconnected LAN
  - Token-Ring or Ethernet
- **Remote Console Operation**
  - Web Browser



## System z Optional TKE

- **TKE 5.0 workstation (up to 3 per server)**
  - Only TKE feature that supports System z9
  - Supports z990, z890, z900, z800 also
  - May carry forward optional Smart Card Reader
  - Feature includes system unit, keyboard, mouse, 17" flat panel display
    - LAN cables are a customer responsibility
- **TKE 4.x workstation**
  - Customers may use to control z990, z890 and prior servers
  - MES available for optional Smart Card Reader support
- **TKE 3.x workstation**
  - Customers may use to control z900, z800 and prior servers
  - No optional Smart Card Reader support



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

## Smart Card Reader feature

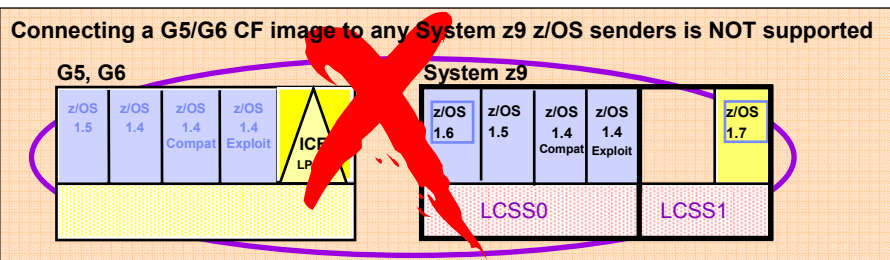
- Each optional Smart Card Reader feature consists of two Smart Card Readers, two cables, and 20 smart cards
- Optional Smart Card Reader feature, attached to TKE 5.0 workstation allows:
  - Use of smart cards, which contain an embedded microprocessor and memory for data storage
  - Storage of master and operational key parts
- Smart cards are protected by a user-defined PIN
- TKE additional smart cards are available via MES

## IBM System z9 BC Parallel Sysplex


## CF Links – System z9 and connectivity to G5/G6

- If you have a System z9 CF image,
  - You cannot connect that CF image to any G5/G6 z/OS senders (or for duplexing, to a G5/G6 CF)
  - Having a System z9 CF, therefore, introduces coexistence issues if there are any G5/G6 CECs participating in that sysplex
  - Need to start thinking about removing any existing G5/G6 images in the sysplex
- If you have a G5/G6 CF image,
  - You cannot connect that CF to any System z9 z/OS senders (or for duplexing to a System z9).
  - Having a G5/G6 CF, therefore, introduces coexistence issues if there are any System z9 CECs participating in that sysplex.
  - Need to start thinking about migrating off of the G5/G6 CFs to something more current
- If you have G5/G6 images AND System z9 images,
  - They can coexist in the same sysplex as long as their CFs are not on G5/G6 or System z9 (e.g., on a z990/z890/z900/z800),
    - These "intermediate" CFs can connect to both the old (G5/G6) and the System z9 systems, so they can provide a "bridge" that allows them to coexist in a sysplex as described above.

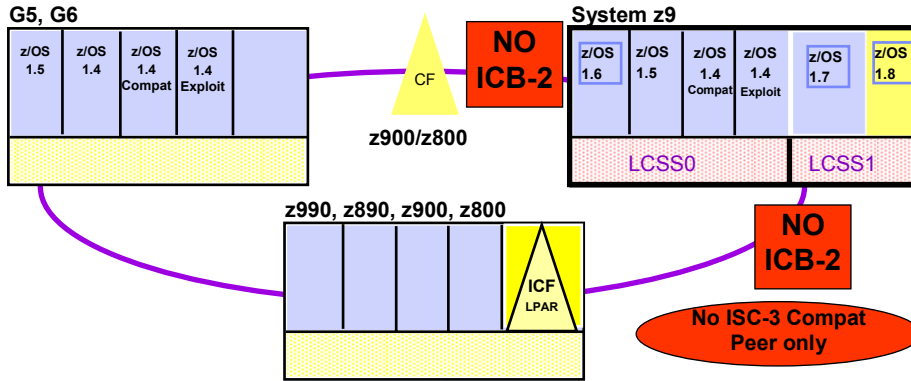
## System z9 – Parallel Sysplex coexistence with a G5/G6 CF



Connecting a G6/G6 CF image for duplexing to a System z9 CF is NOT supported



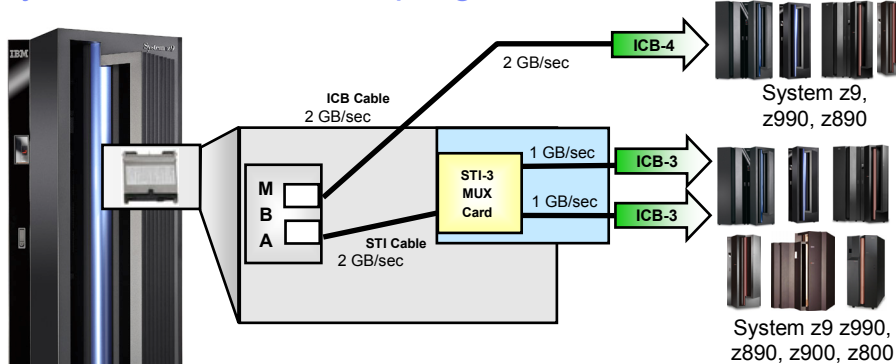
## Parallel Sysplex coexistence with G5/G6 and System z9 Senders



**G5/G6 z/OS senders CAN be in the same sysplex with System z9 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 System z9
- ▶ 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 System z9
- ICB Cables – different types for ICB-3 and ICB-4
  - ▶ 10 meter copper wire cables

## System z 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**



## System z 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
System 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
System z9/z990/z890 ICB-3	N/A	1 GB/s Peer Mode Recommendation: Use ICB-4	N/A
System 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

**New feature codes for ICB cables when connecting to a z9 BC System with 'non-raised floor' option.**



## System z – 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

## Additional Information

- **Parallel Sysplex Web Site:**

- [www.ibm.com/servers/eserver/zseries/psa](http://www.ibm.com/servers/eserver/zseries/psa)

- **GDPS Web Site:**

- [www.ibm.com/servers/eserver/zseries/gdps](http://www.ibm.com/servers/eserver/zseries/gdps)

- **GDPS Executive White Paper**

- [www.ibm.com/servers/eserver/zseries/library/whitpapers/gf225114.html](http://www.ibm.com/servers/eserver/zseries/library/whitpapers/gf225114.html)

## IBM System z9 BC Summary

## z9 BC Model R07 (2006)



- **Machine Type**
  - 2096
- **Model**
  - R07
- **Processor Units (PUs)**
  - 8 PUs per System
  - 1 SAP per book, standard
  - No dedicated spares
  - 7 PUs available for characterization
    - 1 to 3 Central Processors (CPs)
    - IFLs, ICFs, zAAPs, zIIPs, optional SAPs
  - Up to 15 LPARs
- **Memory**
  - Minimum of 8 GB
  - Up to 64 GB per System
    - 8 GB increments
- **I/O**
  - Up to 16 STIs per System @ 2.7 GB/s each
  - Total system I/O bandwidth capability of 43.2 GB
  - Up to 2 Logical Channel Subsystem (LCSS)
  - Up to a maximum of 4 I/O Domains
    - Up to 240 channels – dependent on Channel types

## z9 BC Model S07 (2006)




- Machine Type
  - 2096
- Model
  - S07
- Processor Units (PUs)
  - 8 PUs per System
  - 1 SAP per book, standard
  - No dedicated spares
  - 7 PUs available for characterization
    - 0 to 4 Central Processors (CPs)
    - Integrated Facility for Linux (IFLs), Internal Coupling Facility (ICFs), System z Application Assist Processors (zAAPs), System z9 Integrated Information Processor (zIIP), optional System Assist Processors (SAPs)
    - Can have an IFL-only System
  - Up to 30 LPARs
- Memory
  - Minimum of 8 GB
  - Up to 64 GB per System
    - 8 GB increments
- I/O
  - Up to 16 STIs per System @ 2.7 GB/s each
  - Total system I/O bandwidth capability of 43.2 GB
  - Up to 2 Logical Channel Subsystems (LCSSs)
  - Up to a maximum of 7 I/O Domains
    - Up to 420 channels – dependent on Channel types

## System z9 Channel Type and Crypto Overview

- FICON/FCP
  - FICON Express4
  - FICON Express2 (carry forward on upgrade)
  - FICON Express (carry forward on upgrade)
- 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
    - Fast Ethernet
  - HiperSockets
- Coupling Links
  - ISC-3 (Peer mode only)
  - ICB-3, ICB-4
  - IC
- ESCON
- Crypto
  - Crypto Express2
    - 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.


New feature codes for ICB cables when connecting  
to a z9 BC System with 'non-raised floor' option.



## System z9 FICON Express4 Unrepeated Distances

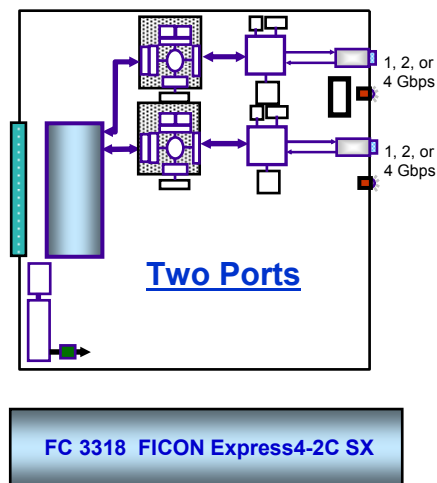
IBM System z9 4Gb FICON/FCP connectivity:

- 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)



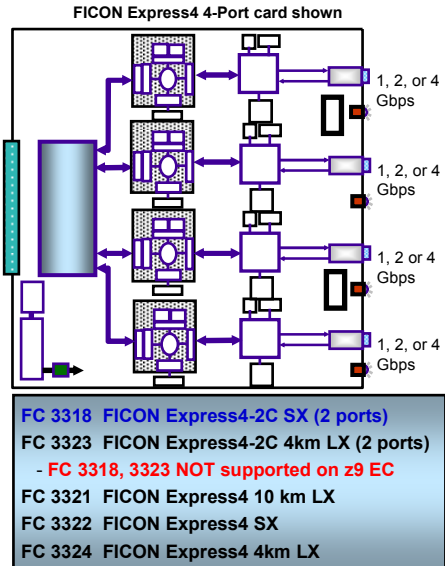
## System z9 BC FICON Express4-2C SX

- FICON Express4 supports:
  - 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
  - One feature = two port increment
  - Intermix is not supported on a single card
  - Both ports must be SX
  - 2 port card can not be 'upgraded' to a 4 port card – it's a physical card replacement
- z9 BC FC 3318 SX (2 ports) is NOT available on z9 EC and cannot be carried forward from z9 BC S07 to z9 EC



## System z9 BC FICON Express4

- **FICON Express4 supports:**
  - 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**
  - Two port or four depending on feature
  - Intermix is not supported on a single card
  - All ports must be of the same type, LX or SX
    - SX Feature Code 3318 has 2 SX SFPs
    - SX Feature Code 3322 has 4 SX SFPs
    - LX Feature Code 3324 has 4 LX 4 km SFPs
    - LX Feature Code 3323 has 2 LX 4 km SFPs
    - LX Feature Code 3321 has 4 LX 10 km SFPs



Enhancements to  
 z9 BC GA2  
 April 2007 Announcement

## System z9 BC GA2 Enhancements (GA May 11, 2007)




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

## Using System z9 Enhanced Driver Maintenance for updating to GA2/GA3 Driver levels

- Need to apply GA MCL(s) to establish Sync Points in order to use EDM to upgrade from driver level 63 to 67
- All HMCs that have the CEC defined MUST be upgraded to Driver 67 (2.9.2 level) BEFORE the EDM upgrade
  - By design, the HMC must be equal too or higher then the highest SE on the network
- Planning required for Systems with the following:
  - CFCC code for CF LPARs. Will require partition activate for CF LPAR(s).
  - OSA CHPID types OSC, OSE, OSD and OSN will require Vary OFF/ON in order to activate new code. However Vary OFF/ON not required with OSA-Express2 defined as CHPID types OSD and OSN.
  - Crypto Express2 code load will require a Config OFF/ON in order to activate new code. Not needed if Crypto Express2 is not installed.
- Option to install new level of Driver code with a planned outage remains

## Using z9 EDM for non-disruptive\* Driver update for GA2/GA3 i.e. Driver level 63 to level 67



\*Requires proper configuration and planning for exploitation

## System z9 LPAR Group Capacity Limit

- Adds capability to define a z/OS LPAR as a member of a group of LPARs
  - Group can cross sysplex boundaries
  - Group can include LPARs not participating in a sysplex
- Adds capability to specify capacity of the group of LPARs in MSUs per hour
  - Synergy with LPAR defined capacity
- PR/SM™ and WLM work together to help:
  - Enforce the capacity defined for the group
  - Enforce the capacity optionally defined for each individual LPAR
- May provide better control of CP resource consumed for WLC pricing
- Exclusive to System z9
- Requires at a minimum:
  - z/OS or z/OS.e Version 1 Release 8 (1.8)

## Change LPAR Group Controls – Assign to Group

**Change LPAR Group Controls**

To remove a partition from a group, enter the word "NONE" in the corresponding "New Group" field.  
Assign partitions to a group, then click "OK" to save ...

Partition Name	Partition Active	Current Group	New Group
LP01	No	LENORA	<input type="text"/>
LP02	No	LENORA	<input type="text"/>
LP03	No	LENORA	<input type="text"/>
LP05	No	GROUP1	<input type="text"/>
LP06	Yes	LENORA	<input type="text"/>
LP07	No		<input type="text"/>
LP08	Yes	GROUP1	<input type="text"/>
LP4	No		<input type="text"/>
LP09	No		<input type="text"/>
LP10	No	LENORA	<input type="text"/>
LP11	No	GROUP1	<input type="text"/>
LP12	No		<input type="text"/>

OK Cancel Help

## Change LPAR Group Controls – Group Capacity

**Change LPAR Group Controls**

Edit group capacities, then click "OK" to save ...

Group Name	Capacity
DEFAULT	<input type="text" value="33"/>
LENORA	<input type="text" value="123"/>
GROUP1	<input type="text" value="33"/>


OK Cancel Help



## Enhanced sysplex management – GLC

- The ability to define a Group Capacity Limit for not only a single logical partition (LPAR), but for a group of LPARs as well
  - System z to manage the group of LPARs in such a way that the sum of the LPAR capacity limits will not be exceeded.
  - Capacity limits based on four-hour rolling average
  - Available on System z9 EC and z9 BC only
  - Available with z/OS 1.8 and higher

- ! May help reduce the amount of 'capping'
- ! For more productive use of 'white space' and higher utilization



## Hardware Decimal Floating Point facility

- Provides a new z/Architecture® Hardware Decimal Floating Point facility providing 4, 8, and 16 byte data formats, an encoded decimal (base 10) representation for data, instructions for performing decimal floating point computations, and an instruction which performs data conversions to and from the decimal floating point representation.
- Exclusive to System z9 EC and BC
  - No emulation or simulation support is provided for earlier machines.
- Requires at a minimum:
  - z/OS or z/OS.e 1.6 with PTF UA26136
  - z/VM 5.2 for guest support
  - IBM High Level Assembler V1.4 or higher plus applicable service for PTF PK18170
- Additional Information
 

<http://www.research.ibm.com/journal/rd/511/duale.html>

## Coupling Facility Control Code Level 15

- All CFs in the configuration have to be at CFCC level 15 to get benefit
- Increasing the allowable tasks in the CF from 48 to 112
- Designed to provide a basis for future\* CF enhancements when CFs are at CFCC level 15
- Exclusive to System z9
- Requires at a minimum:
  - z/OS or z/OS.e 1.6 with APAR OA17055 for function and APAR OA17070 for RMF™ support
  - z/VM V5.1 for z/OS guests requiring virtual coupling facility support
  - Exploitation of future\* enhancements will require z/OS 1.9 (Planned)

Note: When migrating CF levels, lock, list and cache structure sizes may need to be increased to support the new function. This adjustment can impact the system when it allocates structures or copies structures from one coupling facility to another at different CF levels. The coupling facility structure size tool is designed to size structures for you, and takes into account the amount of space needed for the current CFCC levels.

\* 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.

## System-initiated CHPID Reconfiguration

- Adds function to the System z9 to signal each LPAR sharing a CHPID that is to be configured physically Offline or Online from the HMC.
  - New operating system function then configures the LPAR channel image.
  - Designed to eliminate the requirement to enter an operator command manually for each LPAR sharing the channel.
- Designed to reduce operator effort required to configure the CHPID from each LPAR, which can be substantial with up to 60 sharing LPARs.
- May reduce the duration of a service action when an ESCON or FICON channel, and OSA port or a coupling link must be taken physically offline.
- Exclusive to System z9 and applicable to channel types when configured in modes that can be shared.
- Requires at a minimum:
  - z/OS or z/OS.e 1.6 with PTF for APAR OA16166
    - Messages: IOS288A and IOS289I report initiation and completion
  - Linux on System z9 – IBM is working with its distribution partners to provide support for this function in future distributions or service updates.

## Multi-path Initial Program Load

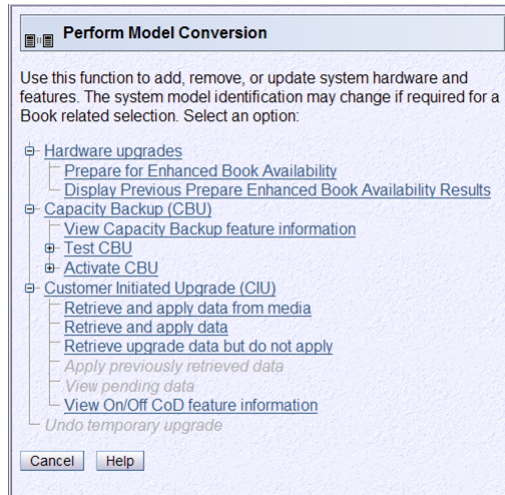
- Adds function to permit use of alternate paths to the IPL device if errors occur during IPL.
- Designed to help increase availability by allowing IPL to complete if possible using alternate paths and to reduce the requirement for manual problem determination when errors occur during IPL.
- Exclusive to System z9 and applicable to FICON channels in FC mode and ESCON channels in CNC mode.
- Requires at a minimum:
  - z/OS or z/OS.e 1.6 with PTF for APAR 0A09050

## Channel Subsystem Enhancements

- Designed to increase availability of ESCON, FICON, OSA and ISC-3 Coupling Links to facilitate recovery from errors – should they occur.
- Exclusive to System z9.
- Transparent to Operating System i.e. no service required.

## On/Off Capacity on Demand Enhancement

- Download one or more On/Off CoD LIC CC records to the System z9 Support Element without application of the capacity
  - Multiple records can have same or different capacities
  - Each record can have customer assigned description name
- May provide more responsive activation of temporary capacity
- Billing begins at activation
- Similar to current function for Customer Initiated Upgrade (CIU) LIC CC records
- Exclusive to System z9



## HMC Logon LDAP Authentication

- Provides optional Lightweight Directory Access Protocol (LDAP) support for HMC logon USERID and password authentication.
  - HMC USERID settings and roles will still be defined on the HMC.
    - These continue to reside on the HMC
  - LDAP “bind mode” support delegates password authentication to a customer furnished LDAP server.
  - SSL and non-SSL connection support is provided between the HMC and the LDAP server.
- Designed to improve management and audit functions for HMC USERIDs and passwords.
  - Authentication is on a per user basis. User1 can have LDAP Authentication while User2 can have Local Server authentication.
- Requires a suitable LDAP Server with LAN connectivity to the HMC.
- Transparent to operating systems supported on System z9.

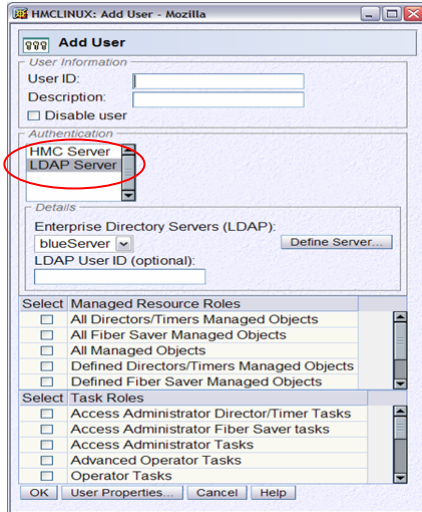
## HMC LDAP Setup Sample – 1

The screenshot shows the 'Add User' dialog box in the HMC Linux environment. The 'Authentication' dropdown menu is open, and 'HMC Server' is selected and circled in red. The dialog includes the following fields and options:

- User Information:** User ID, Description, and a checkbox for 'Disable user'.
- Authentication:** A dropdown menu with 'HMC Server' and 'LDAP Server' options.
- Details:** Password, Confirm password, Password Rule (Basic), and a checkbox for 'Force user to change the password at next login'.
- Select Managed Resource Roles:** A list of roles including 'All Directors/Timers Managed Objects', 'All Fiber Saver Managed Objects', 'All Managed Objects', 'Defined Directors/Timers Managed Objects', and 'Defined Fiber Saver Managed Objects'.
- Select Task Roles:** A list of tasks including 'Access Administrator Director/Timer Tasks', 'Access Administrator Fiber Saver tasks', 'Access Administrator Tasks', 'Advanced Operator Tasks', and 'Operator Tasks'.

This panel illustrates what an administrator will see when they are creating a new user and selecting the HMC to handle the userid and password authentication

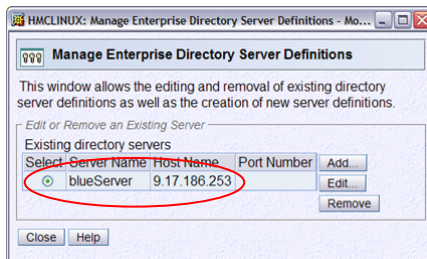
## HMC LDAP Setup Sample – 2



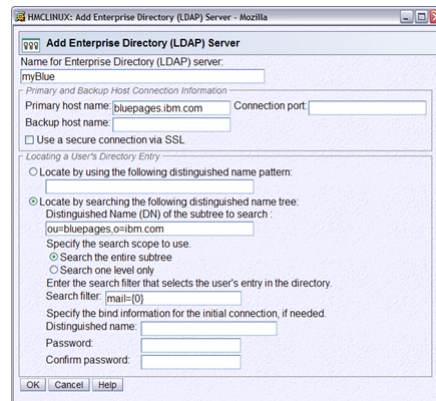
This panel illustrates what an administrator will see when they are creating a new user and selecting a defined LDAP Server to handle the userid and password authentication

## HMC LDAP Setup Sample – 3

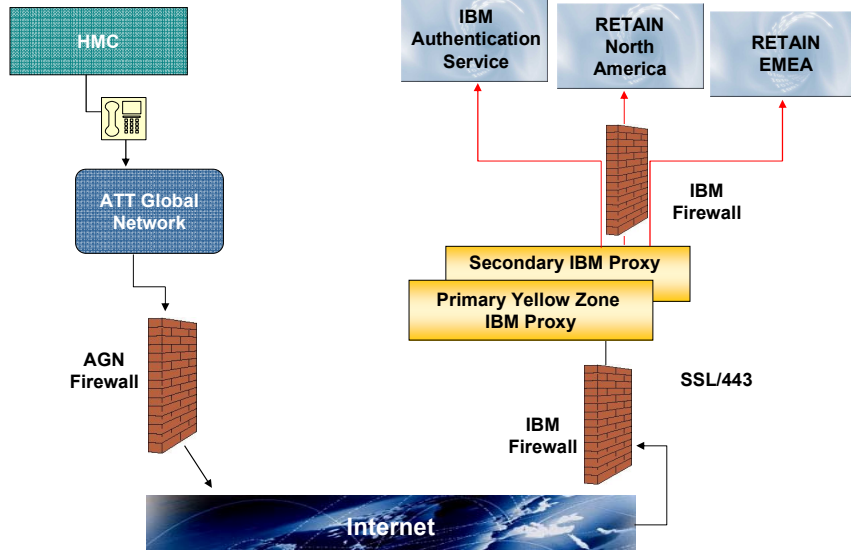
This is the panel that an administrator will see when they either select the Manage Enterprise Directory Server Definitions task, or if they select the Define Server button on the Add User Panel. This panel lists the currently defined LDAP Servers



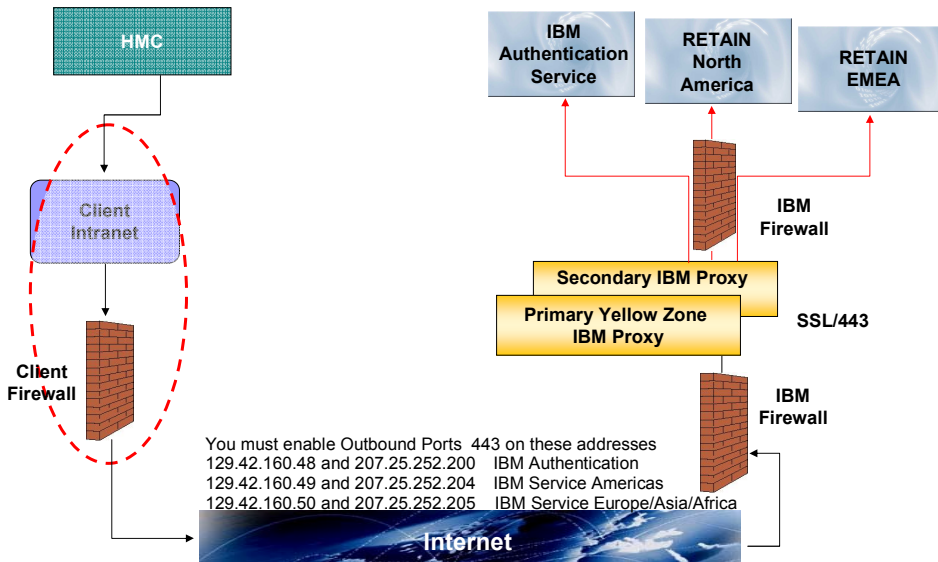
This is the panel which allows the administrator to create or modify an LDAP Enterprise Directory Server definition.

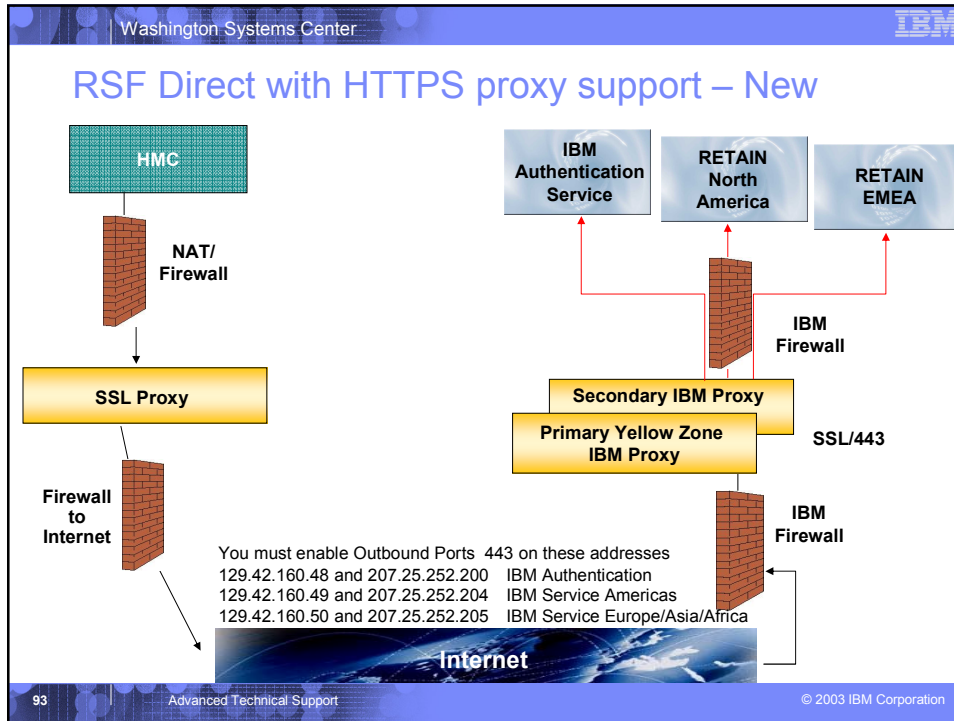


# RSF Dialup Support for zSeries and System z9

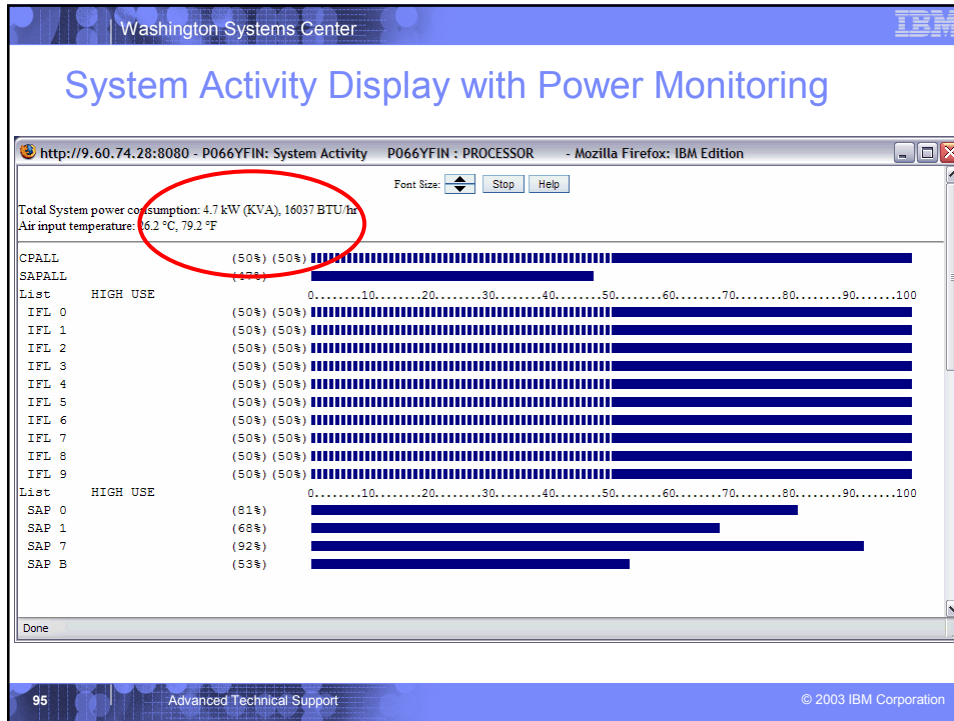


# RSF Direct HMC Connection – System z9





- Washington Systems Center IBM
- ## System z9 Power Planning and Monitoring
- **Power Planning Tool**
    - **Resource Link™** Web site tool designed to estimate power consumption of a specific System z9 machine configuration
    - **Required Input:** Model, memory size, number of I/O cages, and quantity of each type of I/O feature card.
    - Designed to be more accurate than estimates provided in existing physical planning documentation
    - For installed Systems, Resource Link can obtain input from VPD file
  - **System Activity Display (SAD) Power Monitor**
    - Additional function for SAD on the HMC
    - Designed to display Watts and BTUs per hour of power consumption
    - Designed to display cooling air input temperature
  - **Exclusive to System z9**
  - **Designed to help verify power consumption for currently installed System z9 servers and to help power and cooling planning for new System z9 servers or for upgrades to currently installed System z9 servers.**
- 94 Advanced Technical Support © 2003 IBM Corporation



- Washington Systems Center IBM
- ## Frame Bolt Down Features
- Provides hardware to make System z9 frames more rugged and to tie them down to a concrete floor beneath a raised floor or in a non-raised floor installation.
  - System z9 EC (requires two kits) and z9 BC (requires one kit)
    - #7995 Bolt Down Kit for 11.75 to 16.0 inch (298 to 405 mm) raised floor
    - #7996 Bolt Down Kit for 9.25 to 11.75 inch (235 to 298mm) raised floor
  - System z9 BC only
    - #7997 Bolt Down Kit for non-raised floor installation
  - Designed to help secure the machine frames and their contents from damage when exposed to shock and vibration such as those generated by a seismic event.
- 96 Advanced Technical Support © 2003 IBM Corporation

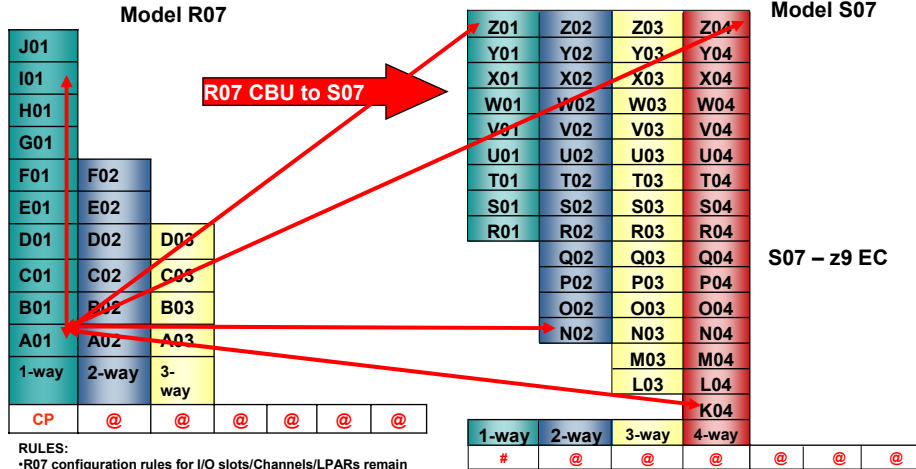


## z9 BC Only Enhancements

## z9 BC Model R07 CBU Enhancement

- **Allows z9 BC Model R07 CPs to CBU to any S07 capacity setting**
  - Adds Concurrent CBU Model Upgrade/Downgrade to/from Model S07
  - Currently available CBU upgrades within Model R07 remain available
  - Adds z9 BC Model S07 CBU CP features to Model R07
    - CBU CP K to CBU CP Z (16 different features)
    - Valid only if CBU target is Model S07
    - Up to four CBU CP features if CBU target is Model S07
  - **Rules:**
    - CBU target must be a valid Model R07 or S07 new-build configuration
      - R07 configuration rules for I/O slots/Channels/LPARs remain
    - Number of active specialty engines may not be decreased
    - Number of active CPs may not be decreased
    - CP capacity may not be decreased
    - One CBU CP feature required at base speed for each target CP added
    - One CBU CP feature required at a different speed for every target CP
- **z9 BC Model R07 CBU Specialty Engine Configuration – No change**
- **z9 BC Model S07 CBU – No change**
- **Does not apply to On/Off Capacity on Demand**

## z9 BC Model R07 – CBU Enhancement



- RULES:**
- R07 configuration rules for I/O slots/Channels/LPARs remain
  - Number of active specialty engines may not be decreased
  - Number of active CPs may not be decreased
  - CP capacity may not be decreased
  - One CBU CP feature required at base speed for each target CP added
  - One CBU CP feature required at a different speed for every target CP

# = CP or IFL or ICF  
 @ = Any Specialty Engines, zAAPs and zIIPs have T & Cs

End of Presentation