



IBM System z9

IBM System z9 109

I/O Infrastructure and Connectivity

zSTSU, August 2, 2005
Harv Emery
Washington Systems Center



© 2005 IBM Corporation

Technical Sales Support, Americas, ATS, Washington Systems Center



Trademarks

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

DB2*	IBM logo*	z/VM*
ESCON*	IBM System z9*	z/VSE
FICON*	Parallel Sysplex*	zSeries*
FICON Express	S/390*	
HiperSockets	z/Architecture	
IBM*	z/OS*	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

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, 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 in the United States, other countries or both.

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.

Agenda

- I/O Introduction and New Function
- I/O Infrastructure
- LCSSes and PCHIDs
- Channels
- Networking
- Parallel Sysplex® Support
- Statements of Direction
- Backup Charts



IBM System z9™ 109 (z9-109) I/O Introduction and New Function



ENABLING BUSINESS.
A THROUGH Z.

z9-109 I/O Enhancements

I/O Enhancements

- Up to 16 x 2.7GB STIs per book* with Redundant I/O Interconnect
- Up to 80% increase in I/O bandwidth per book
- Multiple Subchannel Sets (MSS)
 - Second subchannel set definable per LCSS
 - Up to 63.75K subchannels for Set-0
- Modified Indirect-Data-Address Word (MIDAW) facility
- FICON® Express2 – 64 Open Exchanges
- Up to 28 FICON Express2/Express, Crypto Express2, OSA-Express2/Express features per I/O cage
 - Up to 336 FICON Express2 channels per machine (Up to 256 for the Model S08)
- Storage Area Network (SAN) enhancements:
 - N_Port ID Virtualization (NPIV) for switched FCP
 - Program Directed re-IPL for FCP
 - Remote Node ID Reporting for FICON FC



*z9-109 exploits only a subset of its designed I/O capability


Networking enhancements

- HiperSockets™ IPv6
- OSA-Express2 1000BASE-T Ethernet
- OSA-Express2 OSN (OSA for NCP support)
- GARP VLAN management (GRVP)


Single Subchannel Set per LCSS (per Partition)

z990 2084 Processor	
LCSS 0	LCSS 1
Partitions	Partitions
Subchannels 64512 (63 K) 	Subchannels 64512 (63 K)
Channels	Channels

Subchannel per device




Base Devices




Alias Devices

System z9™ Multiple Subchannel Sets per LCSS (MSS)



- **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**


System z9 MSS Capabilities

z9-109 2094 Processor			
LCSS 0		LCSS 1	
Partitions		Partitions	
Subchannels		Subchannels	
65280 (63.75 K)	65535 (64 K -1)	65280 (63.75 K)	65535 (64 K -1)
Channels		Channels	

Subchannel per device



Base Devices



Alias Devices

- More base volumes, more data in an LPAR or sysplex
- More alias volumes, better performance
- Synergy with z/OS® large volume capability



System z9 Modified Indirect-Data-Address Words (MIDAWs)

- **Modified-Indirect-Data-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



- **Designed to:**

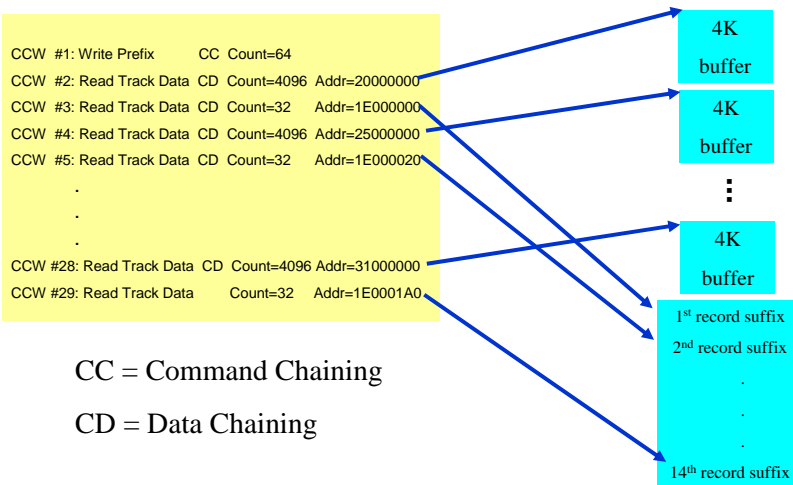
- Be compatible with existing disk control units
- Decrease response time for exploiting I/O.
- Increase the number of I/O operations per second that can be processed and thus move more data per second, especially on faster FICON channels

- **Applications that may benefit include:**

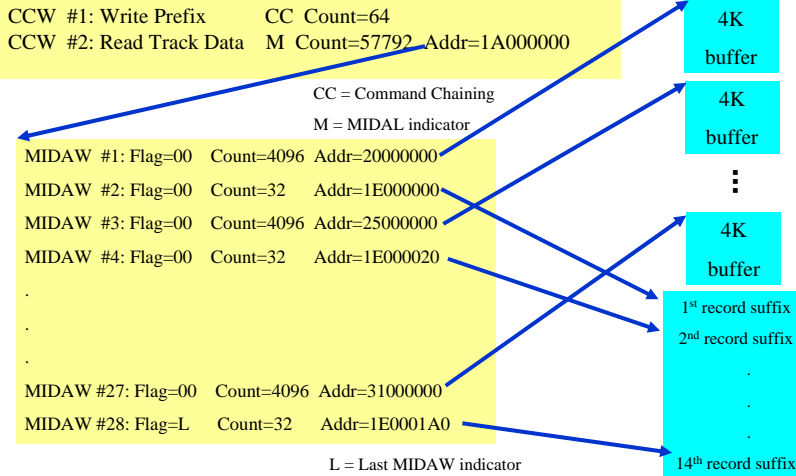
- DB2®, VSAM, Partitioned Data Set Extended (PDSE), Hierarchical File System (HFS), z/OS File System (ZFS), and other datasets exploiting striping and compression




Without the MIDAW facility – 29 CCWs with data chaining Reading 14 disk records with 4K data, 32 byte suffix



System z9 using the MIDAW facility – 2 CCWs Reading 14 disk records with 4K data, 32 byte suffix



Program-directed Re-IPL for Linux on System z9

- **Program-directed Re-IPL is designed to allow an operating system running natively in a logical partition to IPL itself.**
- **This function is designed to support both SCSI and ECKD devices** 
- **To initiate an IPL, an operating system can determine**
 - **How** it was loaded (i.e., via Channel Control Words (CCWs) or via SCSI IPL) and
 - **From where** it was loaded in case of SCSI IPL (World Wide Port Name (WWPN) and Logical Unit Number (LUN) of the load device).
- **Using these parameters, the operating system subsequently calls a firmware function to request that it be reloaded from the same load device, using the same load parameters.**
- **Operating System Support**
 - Linux® on System z9 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.



System z9 Availability Enhancements for FICON

▪ FICON Remote Node ID (RNID) Reporting

- Each FICON (FC mode) channel requests Remote Node ID (RNID) data from each control unit defined to it and stores that data in HSA
- This RNID data can be:
 - Displayed on the SE "Analyze Control Unit Header" IOPD panel.
 - Provided to z/OS so it can be displayed on the "D M=DEV" (Display Matrix Device) command to help debug configuration/cabling problems.

▪ FICON Link Incident Reporting

- Designed to allow z/OS 1.7 operating system images (without operator intervention) to register to receive link incident reports from FICON (FC mode) channels. FICON link incident reporting improves the ability to capture data for link error analysis. z/OS can display this information in console messages and save it in the system log and in the LOGREC error recording data set.




z9-109 I/O Infrastructure



ENABLING BUSINESS.
A THROUGH Z.

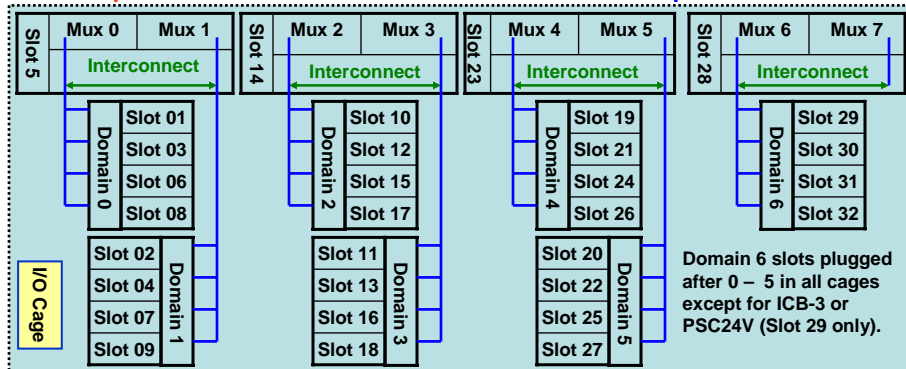


z9-109 Model S08 with three I/O cages

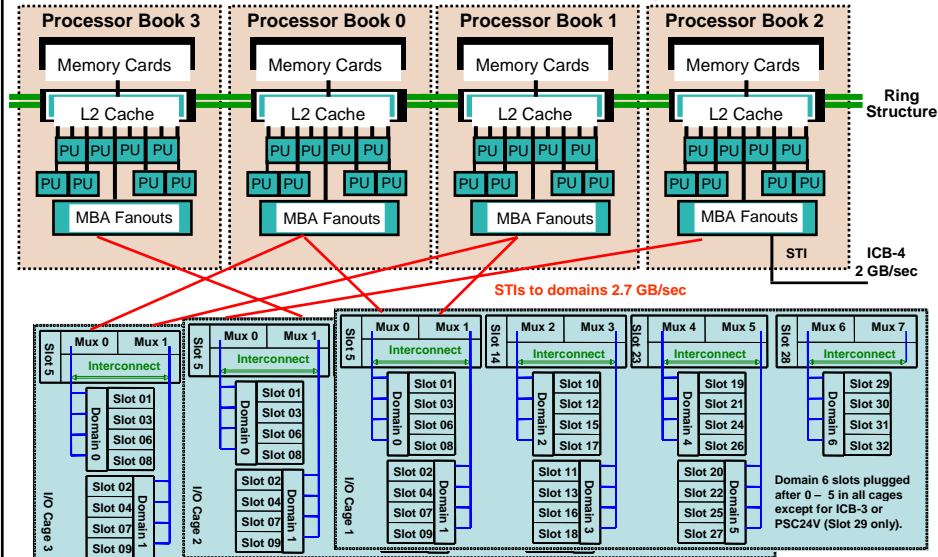


System z9 I/O cage

- **Increased power** – Adequate for any combination of 28 I/O cards
- **New STI multiplexers with Redundant I/O Interconnect**
 - Up to four pairs supporting Domains 0/1, 2/3, 4/5, and 6
 - Up to 1 GB/sec bandwidth per I/O card in each domain
 - Hot pluggable (Airflow cards used if domains not populated)
- **Disruptive MES add** – Use **“Plan Ahead”** to avoid disruption




z9-109 Book, STI and I/O Cage Configuration



Note: Each MBA Fanout card has 2 STI ports. STI connectivity is normally balanced across all installed Books. MBA supports 2 GB/sec for ICB3 and ICB-4 and 2.7 GB/sec for I/O channels. ICB-3 actually runs at 1GB/sec

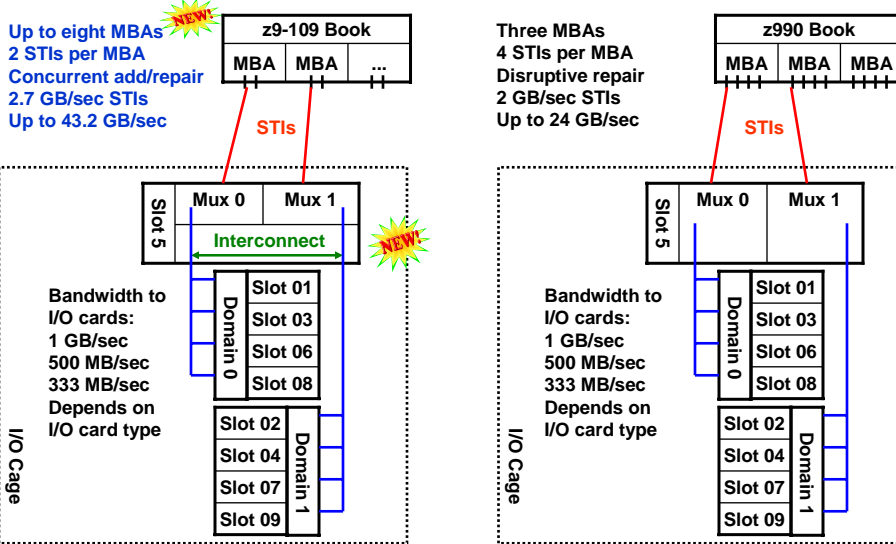
System z9 MBA Fanouts and STIs

- **Up to 8 MBA fanout cards per book**
 - Hot pluggable and maintainable
- **Two STIs per fanout**
 - 2.7 GB/sec for I/O domains
 - 2 GB/sec for ICB coupling links
- **Up to 80% greater I/O bandwidth per book than z990***

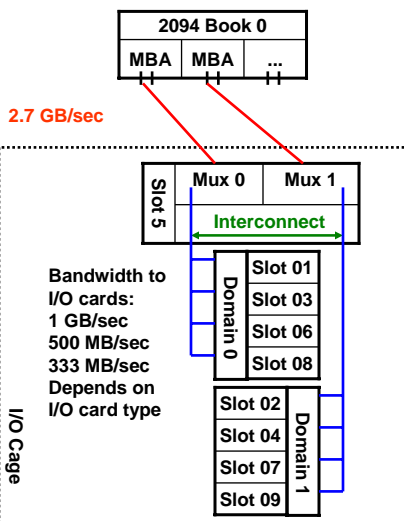


*z9-109 exploits only a subset of its designed I/O capability

System z9 I/O structure compared to z990 (one book)

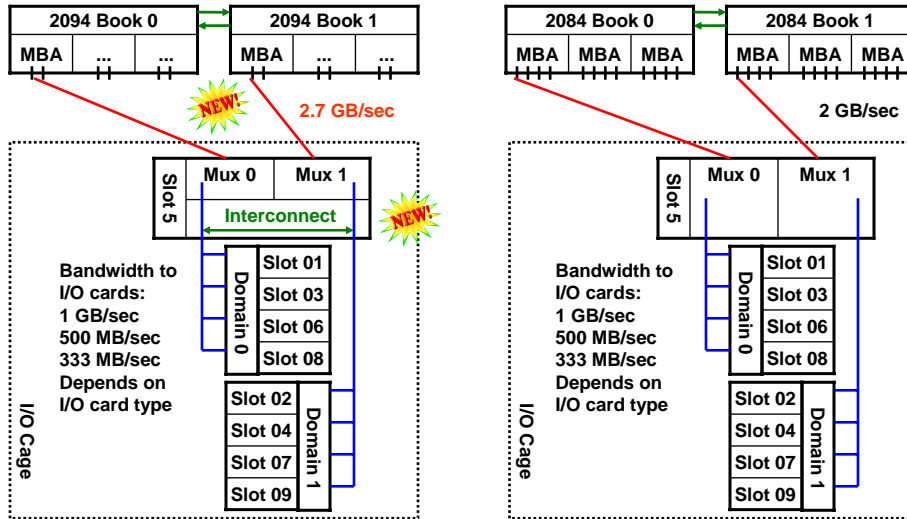


z9-109 Redundant I/O Interconnect – single book



- **Normal Operation**
 - Mux 0 STI supports only Domain 0
 - Mux 1 STI supports only Domain 1
 - STIs from different MBAs
- **Redundant operation – one book**
 - Using the interconnect Mux 0 STI **OR** Mux 1 STI supports **both** domains
 - Supports new function on z9-109:
 - Concurrent MBA repair
 - Concurrent STI cable repair

z9-109 Redundant I/O Interconnect compared to z990

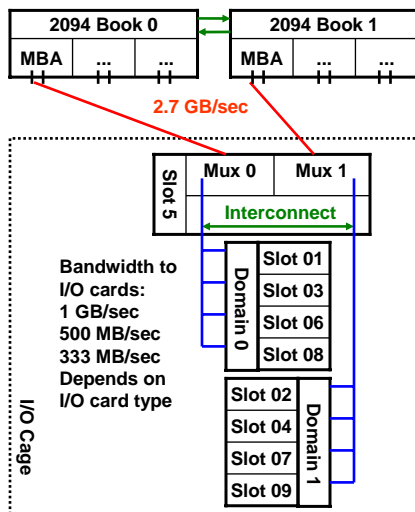


21

zSTSU: August, 2005

© 2005 IBM Corporation

z9-109 Redundant I/O Interconnect - multiple book



Normal Operation

- Mux 0 STI supports only Domain 0
- Mux 1 STI supports only Domain 1
- MBAs from different books

Redundant operation – multi-book

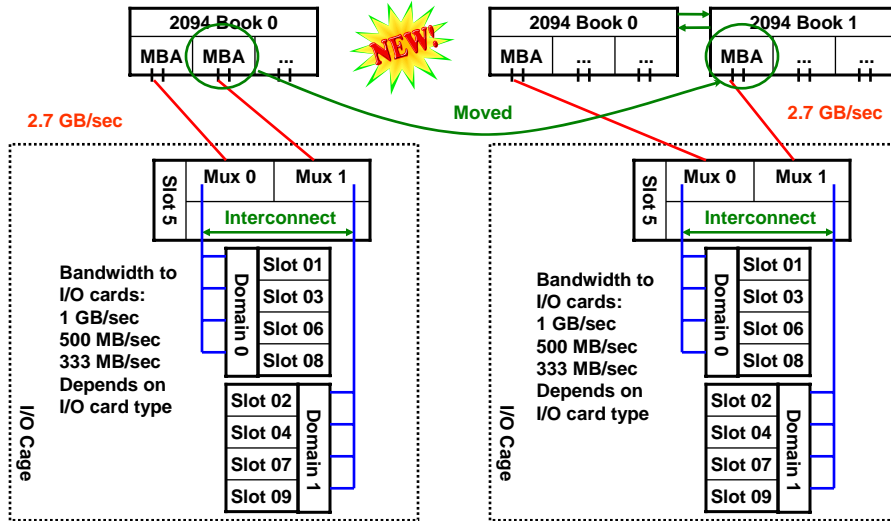
- Using the interconnect Mux 0 STI **OR** Mux 1 STI supports both domains **NEW!**
- Supports: STI/MBA repair and
 - Concurrent MBA/STI reconfiguration for **Concurrent Book Add**
 - **Enhanced Book Availability**
Concurrent book shutdown without loss of I/O (except ICB) and Restart with fenced book

22

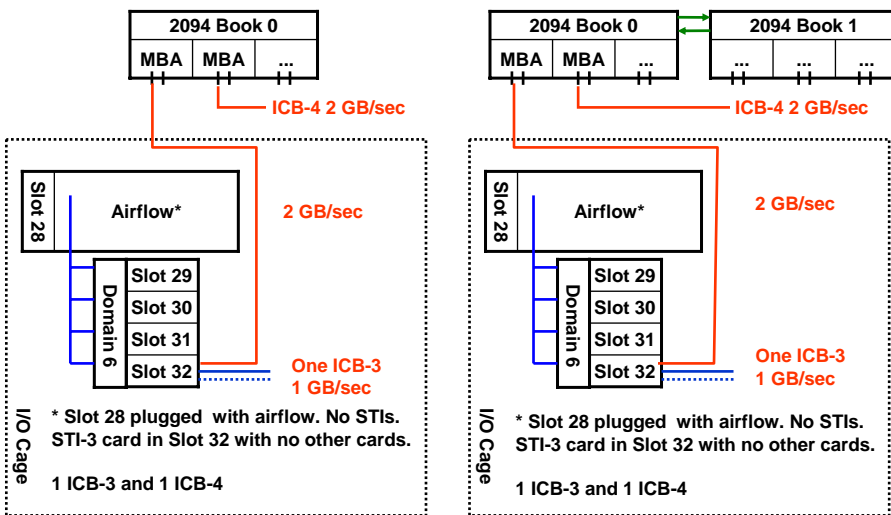
zSTSU: August, 2005

© 2005 IBM Corporation

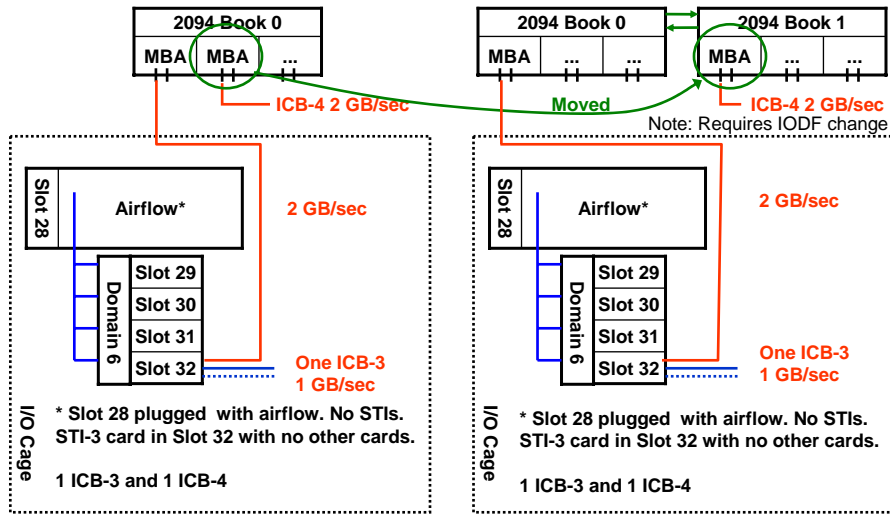
z9-109 S08 concurrent book add upgrade to S18 - Concurrent move of MBAs supporting I/O domains



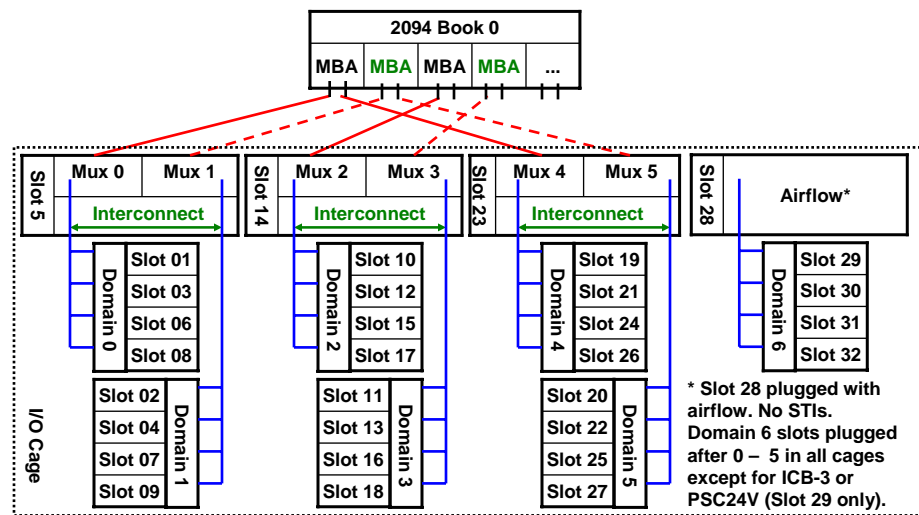
z9-109 S08 concurrent upgrade to S18 (with ICBs) - No Concurrent move of MBAs supporting ICB connectivity



z9-109 S08 disruptive upgrade to S18 (with ICBs) - Disruptive move of MBAs with FC #2400 (optional)

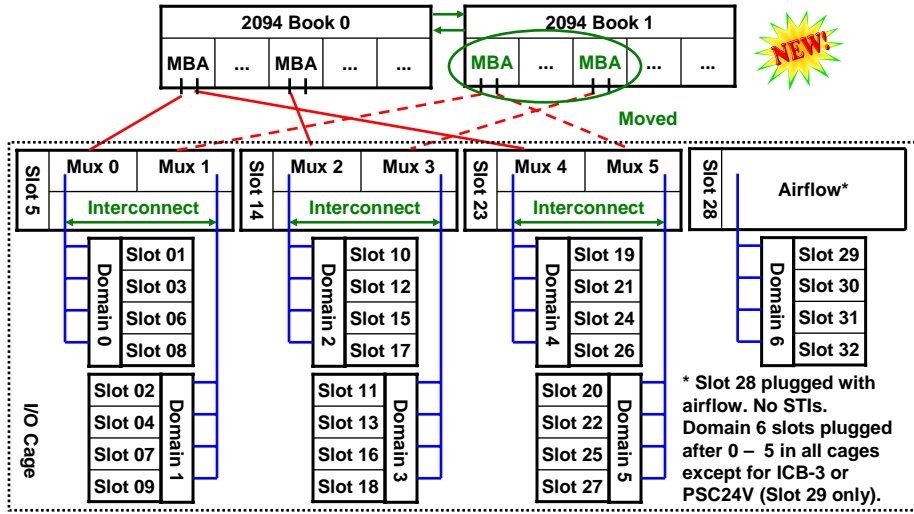


z9-109 S08 with 6 domains plugged, one I/O cage

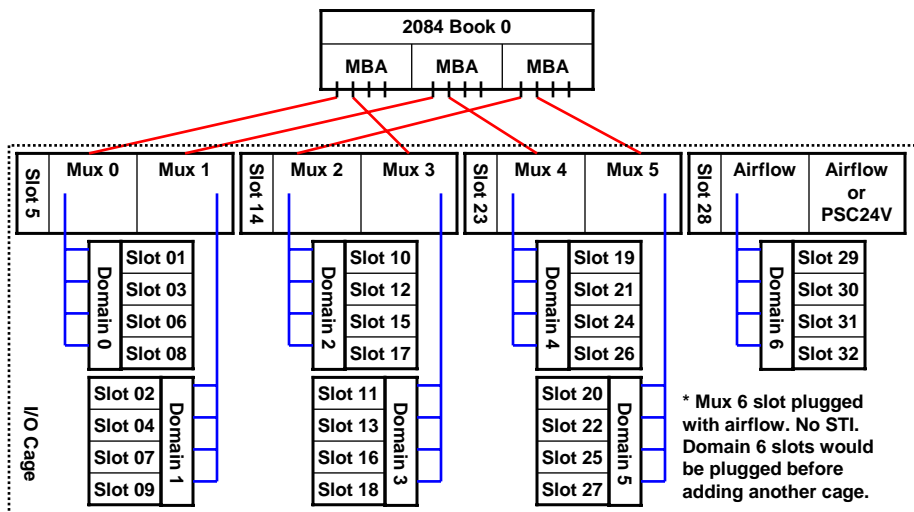




z9-109 S08 to S18 concurrent upgrade, 6 I/O domains - Concurrent move of MBAs supporting I/O domains

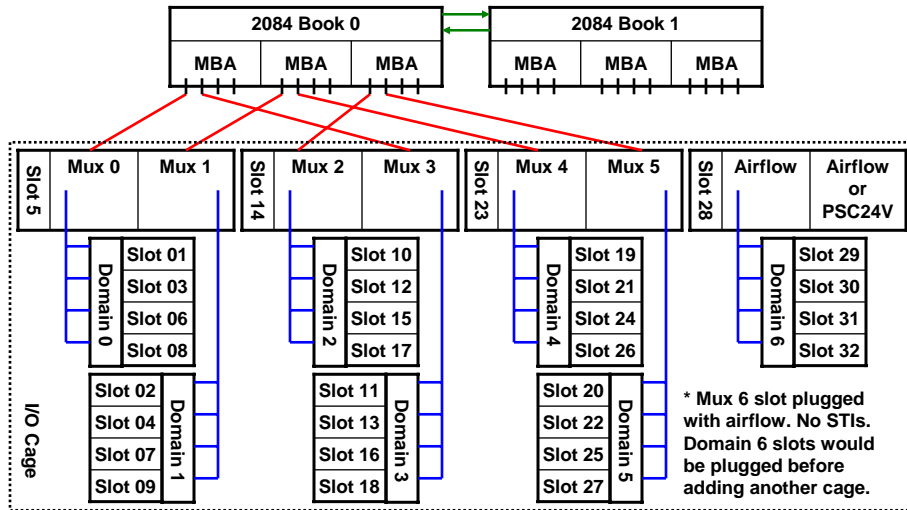


z990 A08 with 6 domains plugged, one cage

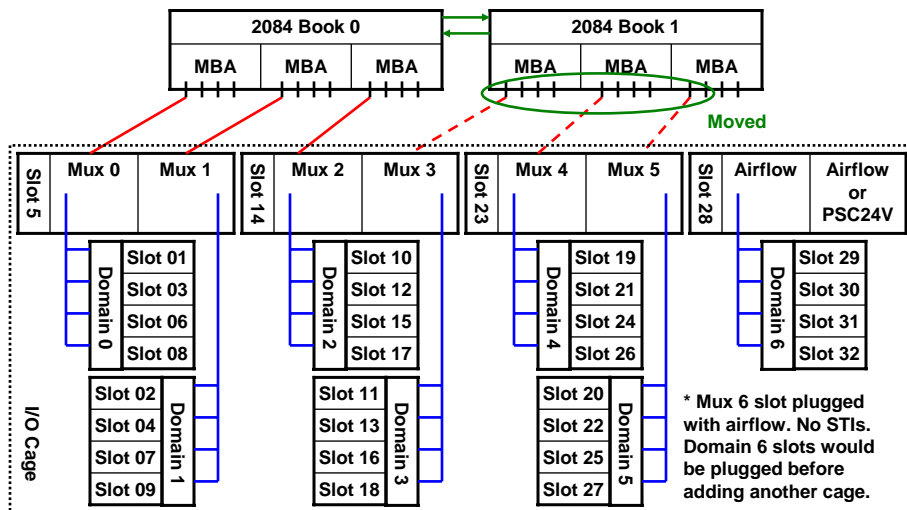




z990 A08 concurrent upgrade to B16 without FC #2400 - Concurrent STI reconfiguration not possible



z990 A08 upgrade to B16 - Disruptive STI reconfiguration with FC #2400 (optional)




System z9 LCSSes and PCHIDs



ENABLING BUSINESS.
A THROUGH Z.

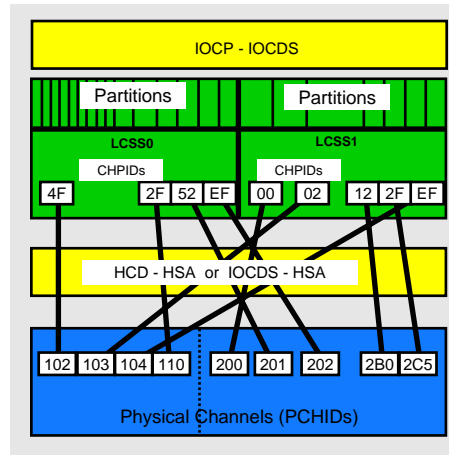
z9-109 Logical Channel SubSystems (LCSSes)



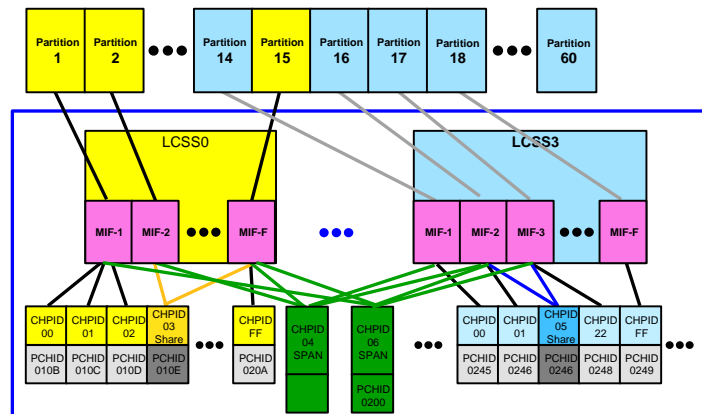
- Up to four Logical Channel SubSystems (LCSSs) on z9-109
 - ▶ Up to 15 LPARs per LCSS
 - ▶ Up to 256 channels per LCSS
- Multiple LCSSes Enable
 - ▶ **Up to 60 Logical Partitions per CEC (Requires four LCSSes on z9-109)**
 - ▶ Up to 1024 external channels on z9-109
- An LPAR can access channels ONLY in its assigned LCSS
- Some channels may be assigned to multiple LCSSes - "Spanned Channels"
 - ▶ ICP, IQD, FC, FCP, OSE, OSD, OSC, **OSN**, CBP, CFP
 - ▶ But not ESCON or FICON Conversion

z9-109 Physical Channel IDs (PCHIDs)

- Channel ports have no CHPID numbers pre-assigned as on z900/800 and earlier machines
- CHPID numbers must be assigned to channel ports (PCHIDs) as part of the definition process
 - HCD/IOCP manual process, and/or
 - [CHPID Mapping Tool](#)
- CHPID assignment
 - Define the channel to an LCSS(s)
 - Associate the CHPID number to a physical channel port location (PCHID)
 - CHPID numbers are still 00 - FF and must be unique within an LCSS
- Physical channel location, known as the PCHID, is assigned by manufacturing and reported by eConfig in the PCHID report
- Except for ESCON sparing, a PCHID relates directly to a jack location on a channel card in a specific I/O slot, in a specific I/O cage
 - Other exceptions
 - ICs and HiperSockets - no PCHID
 - ICB-4 - assigned to CEC cage



z9-109 Four LCSSes and External Spanned Channels



- CHPID 04 Spanned Internal HiperSockets (IQD) or Internal Coupling Link (ICP)
- CHPID 06 Spanned external channel (FICON, OSA, or External Coupling Link)
 - Supported: FC, FCP, OSE, OSD, OSC, **OSN**, CBP, CFP
 - **Not supported: ESCON, FICON Conversion (FCV)**

z9-109 Channels



ENABLING BUSINESS.
A THROUGH Z.

z9-109 Channel and Cryptographic Feature Overview

- **FICON**
 - FC and FCP:
 - FICON Express2
 - 64 Open Exchanges
 - FICON Express
 - FCV: FICON Express
- **Networking**
 - OSA-Express2
 - Gigabit Ethernet LX and SX
 - 10 Gigabit Ethernet LR
 - 1000BASE-T Ethernet
 - OSA-Express
 - Gigabit Ethernet LX and SX
 - 1000BASE-T Ethernet
 - Fast Ethernet
 - HiperSockets
- **Coupling Links**
 - ISC-3 (Peer mode only)
 - ICB-3, ICB-4
 - IC (Peer mode only)
- **ESCON**
- **Crypto**
 - Crypto Express2
 - Secure Coprocessor
 - Default
 - Public Key Accelerator
 - Configurable
- **Channel and crypto types not supported:**
 - OSA-Express Token-Ring (SOD Oct 2004)
 - ICB-2 (SOD 2003)
 - ISC-3 Links in Compatibility Mode (SOD April 2004)
 - Older crypto: PCICA and PCIXCC





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

z9-109 I/O connectivity overview

- **FICON Express2, FICON Express**
 - Up to 84 features
 - Up to 336 channels (256 on Model S08) if all FICON Express2
 - **FICON Express required for FCV support**
- **ESCON Channels, 16-port cards: up to 15 channels per card with spare**
 - Up to 1024 channels, 69 cards (960 channels, 64 cards on Model S08)
- **OSA-Express2, OSA-Express**
 - Up to 24 features, 48 ports
 - Gigabit Ethernet, 10 Gigabit Ethernet, 1000BASE-T Ethernet, Fast Ethernet
- **HiperSockets internal LANs, up to 16**
 - Nothing to order created by I/O definition
- **Coupling Links, up to 64 total**
 - IC (up to 32), ICB-3 (up to 16), ICB-4 (up to 16), ISC-3 (up to 48 links)
- **Crypto Express2**
 - Up to 8 features, 16 coprocessors (If ordered, minimum order is 2 features)
 - Default Secure Coprocessor or **configurable as "Public key" Accelerator**

z9-109 FICON Express2

- **Four channels per feature of LX or SX**
 - Carry forward and new build
- **Connectivity options for each channel**
 - 1 or 2 Gbps, auto-negotiated
 - Can be shared among LPARs, and defined as spanned
- **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
- **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

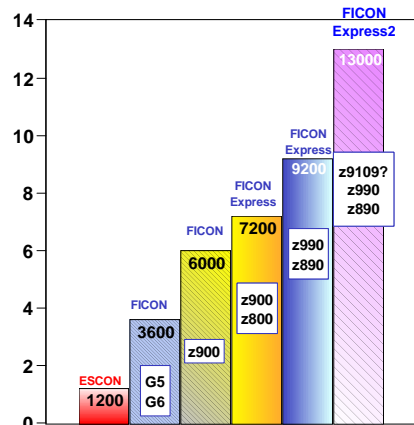


LX = FC3319 or SX = FC3320

z990 FICON Express2 Performance (FC channel)

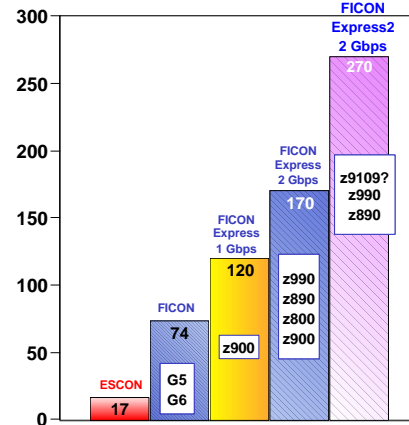
I/Os per second*

4k block size, Channel 100% utilized



MB/sec throughput (Full Duplex)*

Large Sequential R/W mix




*This performance data was measured in a controlled environment on a z990 running an I/O driver program under z/OS 1.6. 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.

z9-109 FICON Express – for FCV Support

- Modes of Operation: defined on a port basis
 - **FCV (FICON Bridge Converted) LX feature only**
 - FC (Fibre Channel): Native FICON and FICON CTC
 - Up to 32 Open Exchanges
 - FCP (Fibre Channel Protocol): SCSI devices in Linux on System z9 environments
- Replaced by FICON Express Express2
 - Still fully supported on z9-109
 - Carried forward in upgrades to z9-109
 - **Not orderable by feature code**
 - **Orderable by RPQ 8P2295 if required to add additional FCV channels ONLY (As Available basis)**
- **Note:** IBM 9032-5 ESCON Director FICON Bridge cards were withdrawn as of December 31, 2004
- Bandwidth
 - 1 or 2 Gbps link data rate- Auto-negotiated
- 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

FICON Express 3rd Generation

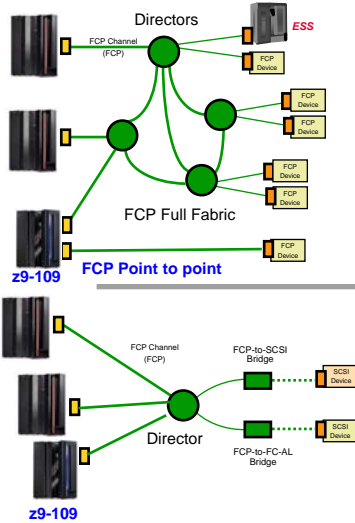


FC 2319 (LX), 2320 (SX)



LC Duplex Connectors

z9-109 FCP Attachment Options for z/VM and Linux



- FCP Full Fabric Connectivity
 - Homogeneous, single vendor fabric
 - Fibre channel directors, switches
 - **NPIV support**
- **NEW!** FCP point to point – Designed to support all FICON features supported on z9-109
 - Direct attachment to FCP CU port
 - NPIV support not applicable
- FCP switched to SCSI Bridge
 - FCP-to-SCSI Bridges
 - FCP-to-FC-AL bridge

Supported devices: www.ibm.com/servers/eserver/zseries/connectivity/#fcp

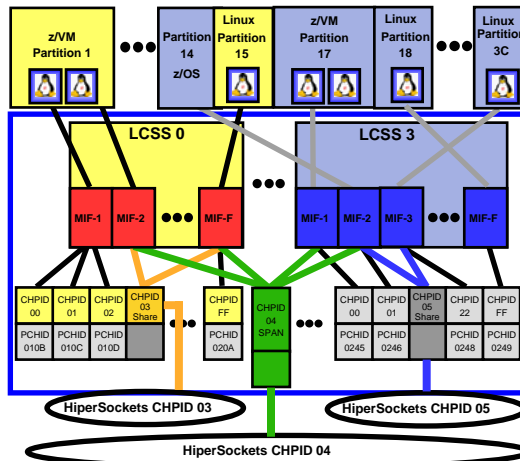
System z9 N_Port ID Virtualization (NPIV) **NEW!**

- NPIV provides the capability to define an FCP channel as multiple virtual FCP channels, each with its own unique virtual Fibre Channel port name and FC N_Port ID.
- By assigning different virtual port names to different operating system images (running in logical partitions or as z/VM guests in virtual machines), operating systems can use the virtual FCP channels as if they were using dedicated physical FCP Channels
- The I/O transactions of each operating system image are separately identified, managed, and transmitted, and are processed by the fibre channel and SAN infrastructure as if each image had its own unique physical N_Port. Each virtual FCP Channel port identifier (WWPN) can be used in device-level access control (LUN masking) in storage controllers and in switch-level zoning.
- **Designed to permit multiple operating system images concurrently to access SAN attached devices (LUNs) via a single, shared FCP Channel.**
 - Designed to work with all FICON features supported on z9-109
 - May reduce number of FCP channels needed
 - May simplify infrastructure by reducing ports and cabling
 - May improve channel utilization
- **NPIV is an industry standard solution** which supersedes FCP LUN Access Control
 - Requires Fibre Channel Director support in the FCP channel's entry director only
 - IBM is working with Linux on System z9 distribution partners to provide NPIV support

z9-109 Networking

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
- **Operating system support:**
 - z/OS V1.7
 - z/VM V5.2 with PTFs in 2Q2006



Very High Speed Interconnection between programs running z/OS, z/VM or Linux on System z9

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 switching




System z9 OSA-Express2 1000BASE-T Ethernet

- **New to OSA-Express2 family**
- **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 switching

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	Open System Adapter for NCP


1000BASE-T Ethernet #3366




OSA-Express2 OSN (Open System Adapter for NCP)

Supporting Channel Data Link Control (CDLC)

- Channel protocol for communication from a host operating system to 374x
- Communications Controller for Linux on System z9 (CCL)
 - Emulates 374x control unit



- ✓ CCL
- ✓ LAN connectivity – external communications
- ✓ Traditional connectivity
- ✓ OSA and non-QDIO (CHPID type OSE)
- ✓ Traffic flows on the LAN



- ✓ Traffic flows LPAR-to-LPAR
- ✓ OSA-Express2 (1000BASE-T, GbE) is the bridge
 - Internal communications – CCL to CDLC
 - SNA PU type 5 and PU type 2.1
- ✓ QDIO extensions for CHPID type OSN
- ✓ Unlike-to-unlike with "glue logic"

VLAN management enhancement – support of GVRP


IEEE 802.1

GARP VLAN Registration Protocol (GVRP)

- Propagates VLAN information
- Registers System z9 VLANs with switches
- Support:
 - OSA-Express2 QDIO mode (CHPID type OSD)
 - Requires z/OS1.7 with PTF
 - Planned for z/VM 5.1 2Q2006

Current support:

- VLAN tagging
 - z/OS, z/VM, Linux on System z9
- VLAN priority queuing
 - z/OS, Linux on System z9




Simplified network administration and management

z9-109 Parallel Sysplex Support

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

System z9 continues to meet the requirements for advanced clustering

System z9 concurrent CFCC changes



- **Apply previously disruptive CFCC changes with little disruption**

- Disruption occurs one CFCC LPAR at a time to activate or remove a change
- Allows rolling CFCC maintenance across CF LPARs
 - Similar to rolling z/OS maintenance across OS images
- Helps reduce the requirement to isolate test CFs from production OS/CF images

Note: Some very infrequent CFCC related changes may still require power-on-reset

z9-109 Fiber Optic Coupling Links

- **InterSystem Channel-3**

- 3rd generation fiber optic coupling link
- ISC-3 links ordered in increments of one (two for 8P2197 "long distance" links)
- Activated links balanced across features

- **Peer mode only**

- Connects to System z9 and zSeries® only
- **No connectivity to any 9672 or 9674**


- **Distance supported**

- Unrepeated up to 10 km (20 km RPQ 8P2197)
- Repeated up to 100 km

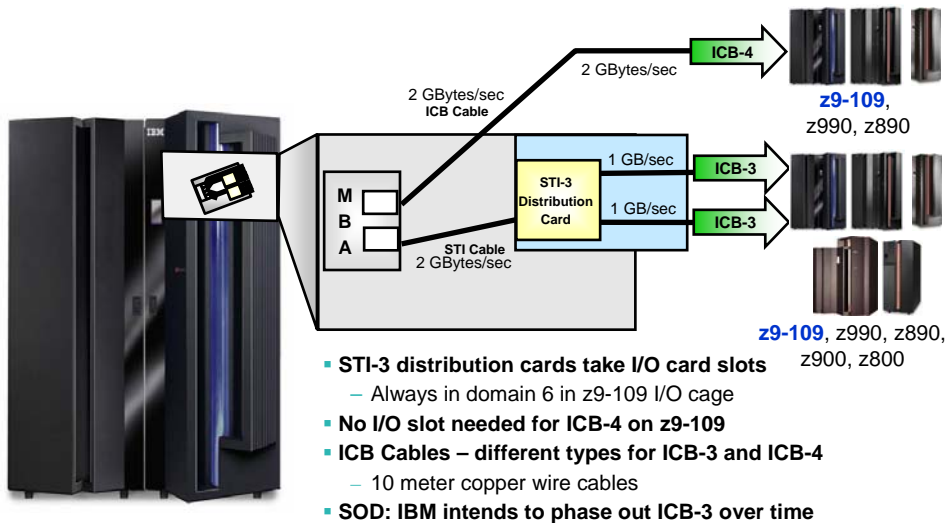
- **ISC-3 feature codes: 0217(ISC-Mother card), 0218 (ISC-D / ISC link)**

- Active link - FC 0219 (2 Gbit/second)
- Two links - RPQ 8P2197 (1 Gbit/second)
- Four ports per ISC-M (two ports per ISC-D)
- Supports 9 micron single mode fiber

- **Up to 48 ISC-3 links per z9-109 server**



z9-109 ICB Coupling Links



- **STI-3 distribution cards take I/O card slots**
 - Always in domain 6 in z9-109 I/O cage
- **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
- **SOD: IBM intends to phase out ICB-3 over time**

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

z9-109 Statements of Direction and Connectivity Backup Charts



**ENABLING BUSINESS.
A THROUGH Z.**

Fulfilled Statements of Direction

- **Support for 60 logical partitions – SOD of May 2003**
- **HMC/SE – SOD of April 2004**
 - Implements a closed Desktop, which discontinues support for:
 - Customer applications
 - Console applications for ESCON director
 - Console applications for Sysplex Timer
 - Withdraws support for control of G4 and older servers
 - Communicates with TCP/IP only
 - Connects to z9-109 with Ethernet only
- **Coupling Support Withdrawn – SOD of May 2003, April 2004**
 - ICB-2
 - ISC-3 configuration in compatibility (Send/Receive) mode
 - All coupling connectivity to S/390® G6 and earlier servers
- **Networking – SOD of October 2004**
 - OSA-Express Token-Ring support withdrawn

System z9 and zSeries Statements of Direction

- **Withdrawal from Marketing – IBM Sysplex Timer 9037 Model 002**
 - IBM intends to withdraw the Sysplex Timer Model 2 (9037-002) from marketing in 2006.
 - A Sysplex Timer or equivalent (refer to Server Time Protocol preview) is a mandatory hardware requirement for a sysplex consisting of more than one server.
 - All strategic planning should include a plan to migrate from a Sysplex Timer Model 2 network to an STP network, when it becomes available
- **IBM intends to phase out ICB-3 coupling links over time**
- **IBM Communication Controller for Linux on System z9**
 - IBM intends to provide CCL enhancements in network connectivity such as CDLC using OSA-Express2 OSN (OSA for NCP), DLSw, and an open interface for X.25 (NPSI).
- **Statements of Direction for z/VM 5.2**
 - IBM intends to provide the following support in z/VM V5.2 :
 - System and guest exploitation of HiperSockets supporting the IPv6 protocol
 - Improved memory management between z/VM and Linux on System z9 and zSeries
 - Simplified networking administration and management of VLANs with support for GARP VLAN Registration Protocol (GVRP) using OSA-Express2
 - IBM intends to provide a future enhancement to allow guest operating systems and z/VM users to query virtual port names when using NPIV
- **Statement of Direction for z/VSE™**
 - IBM intends to provide a software sub-capacity measurement tool for z/VSE.

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



System z9 and zSeries OSA-Express features

Feature	Feature Name	z900	z800	z990	z890	z9-109	CHPIDs
3364	OSA-E2 GbE LX	N / A	N / A	01/05	01/05	Yes, C	OSD, OSN (09/05)
3365	OSA-E2 GbE SX	N / A	N / A	01/05	01/05	Yes, C	OSD, OSN (09/05)
3366	OSA-E2 1000BASE-T Ethernet	N / A	N / A	N / A	N / A	Yes	OSC, OSD, OSE, OSN
3368	OSA-E2 10 GbE LR	N / A	N / A	01/05	01/05	Yes, C	OSD
1364	OSA-E GbE LX	09/04	09/04	WD, C	WD, C	C	OSD
1365	OSA-E GbE SX	09/04	09/04	WD, C	WD, C	C	OSD
1366	OSA-E 1000BASE-T Ethernet	N / A	N / A	Yes	Yes	C	OSC, OSD, OSE
2362	OSA-E 155 ATM SM	Yes	Yes	RPQ	N / A	N / A	OSD, OSE
2363	OSA-E 155 ATM MM	Yes	Yes	RPQ	N / A	N / A	OSD, OSE
2364	OSA-E GbE LX	WD	WD	C	C	C	OSD
2365	OSA-E GbE SX	WD	WD	C	C	C	OSD
2366	OSA-E Fast Ethernet	Yes	Yes	C	C	C	OSD, OSE
2367	OSA-E Token-Ring	Yes	Yes	Yes	Yes	N / A	OSD, OSE

LX = Long wavelength transceiver, SX = Short wavelength transceiver, LR - Long Reach transceiver
 C = Carry forward on an upgrade. Replacements available.
 WD = Formerly available, now withdrawn. Replacements available
 Yes = Currently orderable.



System z9 and zSeries Coupling Link Connectivity

To: \ From:	z9-109 ISC-3	z9-109 ICB-3	z9-109 ICB-4
z900/z800 ISC-3	2 Gb/sec* Peer Mode	No	No
z9-109/z990/z800 ISC-3	2 Gb/sec* Peer Mode	No	No
z900/z800 ICB-3	No	1 GB/sec Peer Mode	No
z9-109/z990/z890 ICB-3	No	1 GB/sec Peer Mode Recommendation: Use ICB-4	No
z9-109/z990/z890 ICB-4	No	No	2 GB/sec Peer Mode
z900/z990 ICB-2	Impossible!		

*RPQ 8P2197 links – Run at 1 Gb/sec in Peer mode and support 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



z9-109 I/O connectivity 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 capability	Capability per I/O slot	Purchase increments
16-port ESCON	0 (1)	69 64 (S08)	1024 channels 960 (S08)	15 channel ports 1 spare port	4 channels
FICON Express2	0 (1)	84 64 (S08)	336 channels 256 (S08)	4 channels	4 channels
FICON Express *	0 (1)	60	120 channels	2 channels	2 channels
STI-3 (2)	0	8	N/A	2 ICB-3 links	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 ISC-3 links	1 link (2 if 8P2197)
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 coprocessors	2 coprocessors	2 coprocessors (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 two features (4 PCI-X adapters). Each PCI-X adapter can be configured as a coprocessor (default) or an accelerator.
- * Available only when carried forward on an upgrade or, for FICON Express, by RPQ for FCV support.
- Note: There are maximums of 64 ESCON features/960 active channels and 64 FICON features/256 channels on the Model S08.



zEnd

