IBM GLOBAL SERVICES

Session G02

IBM eServer zSeries 990 and 890 Update

Harv Emery, IBM, Washington Systems Center

zSeries® EXPO
FEATURING Z/OS, Z/VM, Z/VSE
AND LINUX ON ZSERIES

September 19 - 23, 2005 San Francisco, CA

IBM Americas ATS, Washington Systems Center

IBM eServer zSeries
990 and 890 Update

Session G02, Expo 2005 in San Francisco
Harv Emery, Washington Systems Center
System z9 and zSeries Hardware ATS

RETURN TO INDEX

Trademarks

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

APPN*	IBM logo*	Virtual Image Facility
DB2*	IBM System z9	VM/ESA*
e-business logo*	IMS	VSE/ESA
Enterprise Storage Systems	Magstar*	VTAM*
ESCON*	MVS	WebSphere
FICON	Netfinity*	z/Architecture
GDPS	OS/390*	z/OS
Geographically Dispersed Parallel Sysplex	Parallel Sysplex*	z/OS.e
HiperSockets	PR/SM	z/VM
IBM*	S/390*	zSeries
	S/390 Parallel Enterprise Server	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Lotus, Notes, and Domino are trademarks or registered trademarks of Lotus Development Corporation
LINUX is a registered trademark of Linus Torvalds
Penguin (Tux) complements of Larry Ewing
Tivoli is a trademark of Tivoli Systems Inc.
Java and all Java-related trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States and other countries
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.
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.

IBM considers a product "Year 2000 ready" if the product, when used in accordance with its associated documentation, is capable of correctly processing, providing and/or receiving date data within and between the 20th and 21st centuries, provided that all products (for example, hardware, software and firmware) used with the product properly exchange accurate date data with it. Any statements concerning the Year 2000 readiness of any IBM products contained in this presentation are Year 2000 Readiness Disclosures, subject to the Year 2000 Information and Readiness Disclosure Act of 1998.

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.

IBM eServer z990/z890 Overview

IBM eServer zSeries 990 (z990) Overview

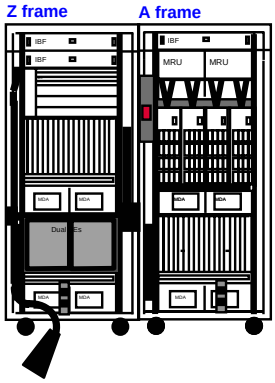

Two Frame System

New function for z990.
Updates for GA3 on 5/28/04

- Processor
  - *4 flexible models*
  - *Unified SMP structure of up to 4 books*
    - *Concurrent model (book) upgrade*
  - 64-bit z/Architecture™
  - Logical Partition Mode *Only*
    - *More than 16 logical PUs in an LPAR*
    - ***Dynamic LPAR name***
    - *Up to 30 LPARs active*
  - *Up to 48 PUs (12 per book), , SuperScalar*
  - *CMOS9S-SOI Technology*
  - *Up to 32 PUs characterizable*
    - *CP, zAAP (6/30/04), ICF, IFL, oSAP*
  - *Concurrent PU Conversion*
  - *Crypto assist for DES and SHA in every PU*
  - Capacity Upgrade on Demand
  - *On/Off Capacity on Demand*
    - *For CPs and IFLs*
    - *For ICFs (5/28/04) and zAAPs (6/30/04)*
    - *Activation coexistence with CBU*
  - Capacity Backup (CBU)
  - Customer Initiated Upgrade (CIU)
  - *Hybrid Cooling (Air/Liquid)* Optional ETR attachment
- Memory
  - *Maximum system memory 256 GB (D32)*
    - *Minimum system memory 16 GB*
  - *Card sizes 8, 16, 32 GB (2 cards per book)*
  - *Bi-directional redundant ring memory interconnect among books*

z990 12 Processor Unit (PU) MCM

- Advanced 93mm x 93mm MCM
  - 16 chip sites, 185 capacitors
  - 100 Glass Ceramic layers
  - 1.3 Volts, 713 Watts
  - 46% smaller than z900

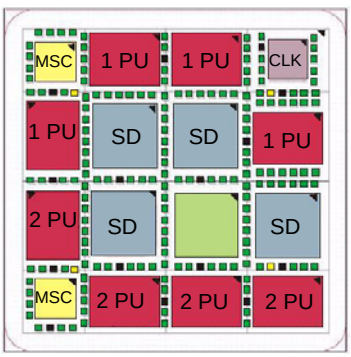

- CMOS 9S - SOI chip technology
  - .83 nsec cycle time
  - PU, SC, SD and MSC chips
  - Copper interconnections, 8 copper layers
- 4 Dual PU, 4 Single PU Chips per MCM
  - 14.1 mm x 18.9 mm
  - 122 million transistors/Chip
  - L1 cache/PU
    - 256 KB I-cache
    - 256 KB D-cache
- 4 System Data (SD) cache chips per MCM
  - 521 million transistors/chip
  - 8 MB L2 cache per chip
  - Single 32 MB L2 cache per MCM
- 1 Storage Control (SC) chip
  - L2 cache crosspoint switch
  - L2 access rings to/from other MCMs
  - L2 access to/from MBAs (off MCM)
- 2 Storage Control (MSC) chips
  - Memory cards (L3) interface to L2
- 1 Clock (CLK-ETR) chip
  - CMOS 8SF, 7 copper layers
  - Clock and ETR Receiver

z990 MSU and ITR Comparisons

z990 STSI "MCI"	CPs	Base	ITR Ratio
301	1	z900 2C1	1.54 - 1.61
308	8	z900 2C8	1.52 - 1.56
316	16	z900 216	1.51 - 1.55
332	32	z900 216	2.46 - 2.98

MCI = Model Capacity Indicator. For z990, this is "3nn", where "nn" is the number of CPs.

See: http://www.ibm.com/servers/eserver/zseries/library/swpriceinfo/pdf/swp_ibm.pdf
for exact MSU values represented by the height of each bar.

Notes: MSU values originally announced with z990 were reduced 10% by the IBM Mainframe Charter as a way of reducing software license charges. Capacity and ITR ratios did NOT change. STSI is the "Store System Information" instruction, which returns configuration information.

z890 System Overview

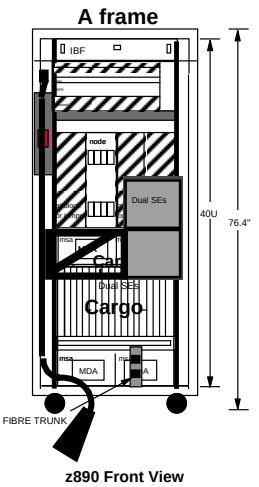

z890 Front View

- **Processor**
  - Single Hardware model
  - 28 Capacity Settings
  - New to z890 - zAAPs
  - 64-bit z/Architecture
  - 5 PUs - One standard SAP and up to 4 PUs characterizable
  - CMOS 9S-SOI Technology
  - SuperScalar
  - Crypto assist for DES and SHA in every PU
  - Capacity Upgrade on Demand
  - On/Off Capacity on Demand
    - For CPs
    - For ICFs
    - For IFLs
    - For zAAPs
  - Capacity Backup (CBU) - To full size engines only
  - Customer Initiated Upgrade (CIU)
  - Up to 30 LPs active
  - Optional ETR attachment
  - Internal Battery Feature
- **Memory**
  - Maximum system memory 32 GB
    - Minimum system memory 8 GB
  - Card sizes 8, 16, 32 GB (single card)

z890 5-way MCM

- **Advanced 93mm x 93mm MCM**
  - 101 Glass Ceramic layers
  - 16 chip sites, 13 in use
  - 0.4 km of internal wire
  - 5,184 LGA connectors vs 1,849 pins for z800

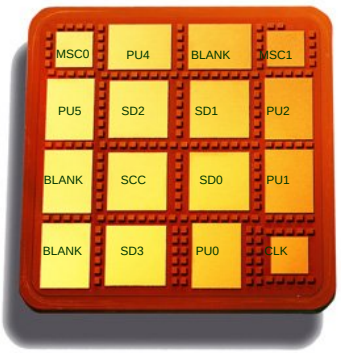

Note: MBAs not on MCM

- **CMOS 9S-SOI chip Technology**
  - PU, SC, SD and MSC chips
  - Copper interconnections, 8 copper layers
  - 5 PU chips/MCM
    - 14.1 mm x 18.9 mm
    - 122 million transistors/PU
    - L1 cache/PU
      - 256 KB I-cache
      - 256 KB D-cache
    - 1.0 ns Cycle Time
  - 4 System Data (SD) cache chips/MCM
    - 17.5 mm x 17.5mm
    - L2 cache per Book
      - 521 million transistors/chip
      - 32 MB
  - One Storage Control (SC) chip
    - 17.3mm x 17.3mm
    - 98 million transistors
    - 3692 Power Signal I/Os
    - L2 access to/from MBAs (off MCM)
  - Two Memory Storage Control (MSC) chips
    - Memory cards (L3) interface to L2
  - One Clock (CLK) chip - CMOS 8S
    - Clock and ETR Receiver

z890 MSU Comparison

* Capacity setting refers to number of installed CPs and the capacity. Reported by STSI instruction.

z890 Uni is 1.98 - 2.09 times the z800 Uni
z890 4-way is 2.19 - 2.28 times the z800 4-way.

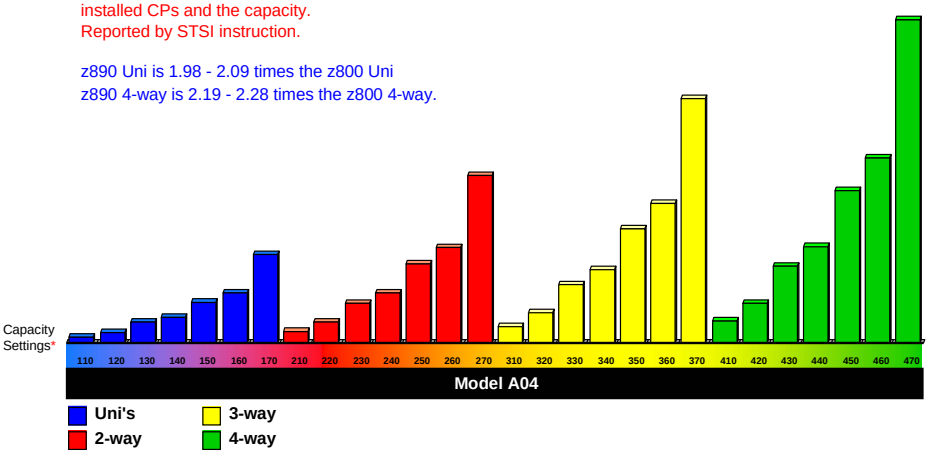

Note: For MSU values, refer to: http://ibm.com/servers/eserver/zseries/library/swpriceinfo/

zSeries Type/Model Upgrades to z890 and z990

- **z800**
  - **Model 0A2**
  - **Model 0X2**
  - Model 002
  - **Model 003**
  - Model 004
- **z890**
  - From z890 Model A04 Capacity Setting:
    170 (Uni)
    250 to 270 (2-way)
    330 to 370 (3-way)
    430 to 470 (4-way)
- **z990**
  - Model A08 (1 to 8-way)

- Selected z800 Models to any z890 capacity setting
- z890 Model A04 to z990 Model A08
- **z990 any model to z9-109 any model (not shown)**
- No upgrade to z890 from any IBM S/390® 9672 or IBM eServer zSeries 900 (z900)

z990 and z890 Logical Channel Subsystems (LCSSes)

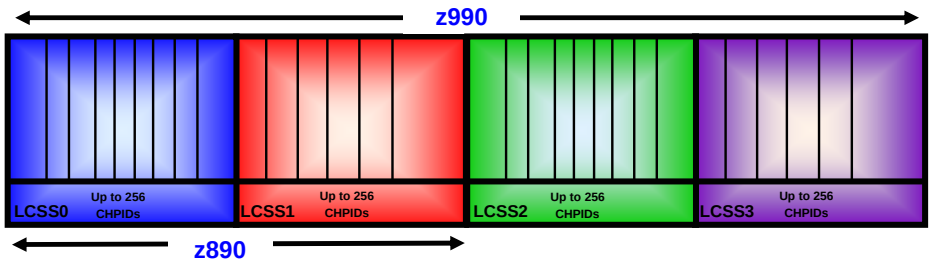

- Up to four Logical Channel SubSystems (LCSSs) z990, two LCSSes on z890
  - Up to 15 LPARs per LCSS
  - Up to 256 channels per LCSS
- Multiple LCSSes Enable
  - Up to 30 Logical Partitions per CEC (Even with three or four LCSSes on z990)
  - **Up to 1024 external channels on z990, Up to 421 external channels on z890**
- 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, CBP, CBS, CFP, CFS
  - But not ESCON, FICON Conversion, Coupling Receiver (CBR, CFR)

z990 Four LCSSes and External Spanned Channels

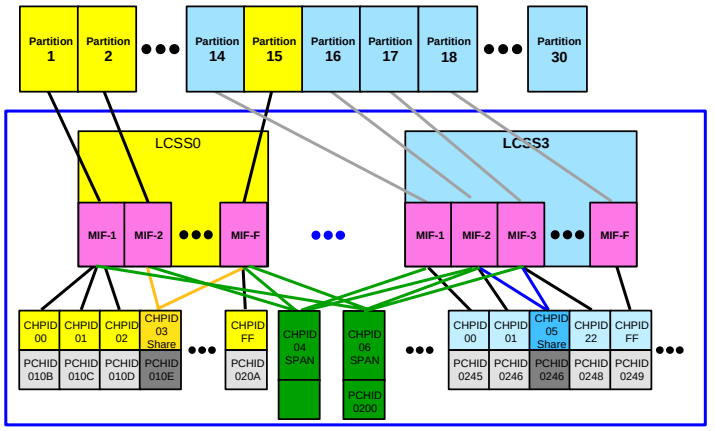

- **CHPID 04 Spanned Internal HiperSockets (IQD) or Internal Coupling Link (ICP)**
- **CHPID 06 Spanned external channel (FICON, OSA, or External Coupling Link) – New**
  - **Supported: FC, FCP, OSE, OSD, OSC, CBP, CBS, CFP, CFS**
  - **Not supported: ESCON, FICON Conversion, or Coupling Receiver (CBR, CFR)**

z990 and z890 Physical Channel IDs (PCHIDs)

- **CHPID numbers are no longer pre-assigned**
  - i.e. there are NO default CHPID numbers
- **Customer MUST assign CHPIDs to channels**
  - HCD/IOCP 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 exception
    - IC and HiperSockets - no PCHID
    - ICB-4 - assigned to CEC cage

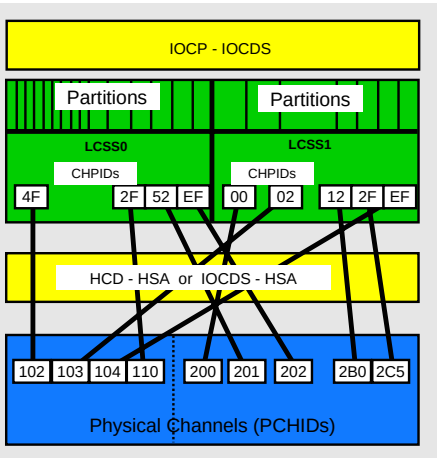

IBM eServer z990/z890
Update for 2005

IBM z890/z990 Highlights for October 29, 2004

- On Demand
  - On/Off Capacity on Demand (On/Off CoD) test
  - Extended order staging for CIU-Express and On/Off CoD
  - z/OS® and z/VM® support for IBM eServer zSeries 890 (z890) CP speed change for On/Off CoD, Customer Initiated Upgrade (CIU) and Capacity BackUp (CBU)
  - Additional IBM eServer zSeries 800 (z800) – z890 upgrade paths
- LAN
  - New OSA-Express2 Gb Ethernet and 10 Gb Ethernet LR **(January 2005)**
    - Concurrent LIC updates*
    - Up to 640 TCP/IP Stacks
    - Large send for TCP/IP traffic
    - Layer 2 support
  - OSA-Express functional improvements
    - Layer 2 Support
    - Improved TCP/IP stack utilization
- SAN
  - Preview of FCP LUN access control **(Date to be announced)**
  - FICON purge path extended
- Security
  - New Crypto Express2 **(January 2005)**
  - New Cryptographic Support
    - 19-digit Personal Account Numbers
    - 2048-bit clear and secure key RSA operations
    - Less than 512-bit clear-key RSA operations
  - TKE 4.2 workstation with Smart Card Reader support
  - IBM eServer zSeries 990 (z990) PR/SM™ EAL5 Certification
- Availability, Clustering, and Virtualization
  - GDPS/PPRC Multiplatform Resiliency for zSeries
  - Coupling Facility dispatcher improvements – CFCC level 14
  - New Hardware Management Console (HMC), HMC flat panel, and HMC LAN connectivity features **(Nov 2004)**
  - GA2/4 function designed for delivery by concurrent LIC updates

IBM z890 and z990 Highlights for January 2005

- **On Demand support for Argentina and Venezuela – Planned June 30, 2005**
  - CIU and CIU Express on IBM eServer™ zSeries® 990, 890, 900 and 800 (z990, z890, z900 and z800)
  - On/Off Capacity on Demand (On/Odd CoD) on z990 and z890
    - On/Off CoD test
  - Extended order staging for CIU-Express and On/Off CoD
- LAN
  - OSA-Express2 Gb Ethernet and 10 Gb Ethernet LR **(January 28, 2005)**
    - **Preview - Large send support planned for z/OS and z/OS.e 1.7 (1.6 with PTF)**
  - OSA-Express Gb Ethernet withdrawn from marketing on z990 and z890
- **SAN**
  - **FICON® Express2**
    - Four LX or four SX ports per feature
    - Up to 240/80 FICON Express2 ports on z990/z890
    - Supports improved channel performance
    - Supports FC and FCP definition (not FCV)
  - **FICON Express withdrawn from marketing on z990 and z890**
    - **Orderable by RPQ 8P2295 if additional FCV is required**
  - Preview of FCP LUN access control
- Security
  - New Crypto Express2 **(January 28, 2005)**
  - PCICA and PCIXCC withdrawn from marketing on z990 and z890
  - **SOD: A Cryptographic Coprocessor for the IBM eServer xSeries® designed to meet FIPS 140-2 Level 4 is planned for 2H 2005**
- Availability, Clustering, and Virtualization
  - **GDPS® 3.2 Enhancements (Planned March 31, 2005)**
    - **GDPS/PPRC HyperSwap™ function extensions**
    - **GDPS/PPRC and XRC FlashCopy® exploitation**
  - **GDPS/PPRC support for IBM TotalStorage® Enterprise Storage Server® (ESS) Global Mirror-Preview**
  - GA2/4 function designed for delivery by concurrent LIC updates

z890 Upgrades – Concurrent support for z/OS and z/VM

- Single Machine: 2086 and a single Model: A04
- A dramatic new way to consider upgrading
- One MCM per model with 5 Processor Units (PUs)
  - Four PUs available for characterization
    - CPs, IFLs, ICFs, or zAAPs
  - One PU standard as a SAP
- Standard CPs –
  - Four full capacity processors each with 7 capacity settings
    - Entry point is 65% less capacity than z800-0E1 and largest capacity setting is 110% more than z800-004
  - Upgrades can be horizontal, vertical, or diagonal or whatever way best fits your needs. All are designed to be concurrent to hardware (no POR)
  - Horizontal upgrades add same speed CP and are designed to be OS concurrent
  - **Vertical or diagonal upgrades change CP speed (Formerly – IPL required)**
    - **Designed to be concurrent: z/OS or z/OS.e 1.4 and higher with PTF for APAR OA07510**
    - **Designed to be concurrent: z/VM 5.1 for Linux guests and z/OS guests with PTF for OA07510**

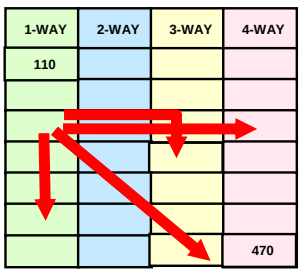

* Note: No mixing of standard CP capacity sizes in multi-engine machines

Think of the possibilities:
Define the processor the way your business requires!

z890/z990 Crypto Roadmap to Crypto Express2

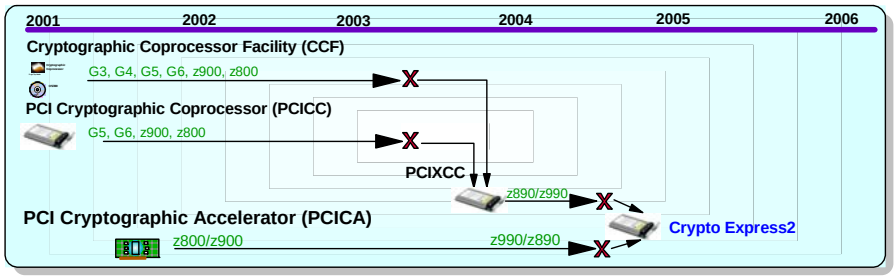

- **z990/890 includes NO standard cryptographic function**
- CP Assist for Cryptographic Function (message security assist) Feature #3863
  - Allows access to crypto functions from any CP (Limited to CP0 and CP1 on z900)
  - Supports limited clear key processing **running on the CP** – Compute intensive!
  - **NOT equivalent to CCF on older machines in function or offload**
- PCIXCC Feature – Supports "Secure key" cryptographic processing
- PCICA Feature – Supports "Public key" SSL cryptographic processing
- Crypto Express2 – Combines function and performance of PCICA and PCICC
- **Migration to z990 when CCF, PCICC or PCICA is in use on an older machine usually requires Crypto Express2, PCIXCC and/or PCICA on z890/990.**

z890/z990 CP Assist for Cryptographic Function (CPACF)

- High performance crypto engine in every CP
  - Clear key DES and hashing
  - Optimized for low-latency SSL transactions
- Two algorithms:
  - DES (single, double, and triple)
    - Up to 2**64 byte message, interruptible execution
    - Requires FC #3863 to enable (Export control)
  - SHA - Defined in FIPS PUB 180-1 publication
    - Always enabled
- 5 new problem state instructions:
  - KMAC - Compute Message Authentication Code
  - KM – Cipher Message
  - KMC – Cipher Message with Chaining
  - KIMD - Compute Intermediate Message Digest
  - KLMD - Compute Last Message Digest

Crypto Express2

- Available – January 28, 2005
- Dual Integrated Cryptographic Coprocessors
  - Provides PCIXCC and PCICA functionality
- Improved throughput over the PCIXCC
  - Multitasking required to use both coprocessors
- Scalable (no CP affinity) - 0 to 8 features
  - The total number of Crypto Express2, PCICA and PCIXCC features cannot exceed 8 features per server
  - All 8 Crypto Express2 features can plug in a single I/O cage without restrictions
  - Minimum purchase increment is two (Crypto Express2 and/or PCIXCC)
- Current applications expected to run without change
- Connection to STI interface; no external cables
- Fully programmable, User Defined Extensions (UDX) support
- Designed for FIPS 140-2 Level 4 Certification
- Trusted Key Entry (TKE) 4.X support
  - Secure operational and master key loading
  - Smart Card Reader support (TKE 4.2 – October 2004)

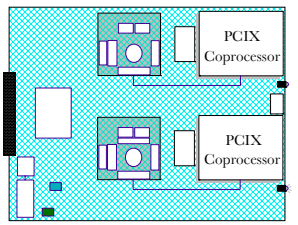

All z990/z890 cryptographic features are designed to be seamlessly managed by ICSF for optimum performance!

Crypto Express2 Support Requirements

- z890 or z990 hardware LIC support for GA2/4 (January 2005)
- z/OS 1.3 or z/OS.e 1.3 or later with Web Deliverable:
  - z990 and z890 Enhancements to Cryptographic Support
  - **Service for ICSF APAR OA09157** and RMF APAR OA07347
- z/VM 5.1 or later with service (January 2005)
  - Dedicated queue support for clear-key and secure-key functions for z/OS guests
  - Shared and dedicated queue support for clear-key functions for Linux on zSeries guests, with up to 256 dedicated queues
- z/VSE 3.1 (planned March 4, 2005), VSE/ESA 2.7 and IBM TCP/IP for VSE/ESA 1.5
  - Clear-key functions only
- Linux on zSeries with IBM Open Source code:
  - Delivered in October 2004 for kernel 2.4
  - Planned for delivery early in 2005 for kernel 2.6
- **See the 2084DEVICE or 2086DEVICE PSP for additional service.**

System z9™ Trusted Key Entry (TKE) Workstation 5.0

- Optional TKE Workstation:
  - **The only TKE feature that supports z9-109**
  - Orderable on z9-109, **z990, z890**, z900 and z800
  - TKE 5.0 LIC: FC 0855
    - **Requires TKE 5.0 hardware**
  - TKE 5.0 hardware: FC 0859
    - **Requires TKE 5.0 LIC**
    - xSeries-based system unit, keyboard, flat panel, mouse
    - PCI-X Crypto Coprocessor
    - **Ethernet connectivity only**
  - Optional Smart Card Reader: FC 0887
  - Optional Additional Smart Cards: FC 0888
- TKE 5.0 Hardware and LIC support to enter secure cryptographic keys for:
  - z9-109: Crypto Express2
  - **z990 and z890: PCIXCC and Crypto Express2**
  - z900 and z800: CCF and PCICC

PCI-X Crypto Coprocessor

FICON Express2

- Four channels per feature
  - Up to 240 channels (60 features, 20 per I/O cage) on z990 Compared to 120 channels with FICON Express
- Performance improvement compared to FICON Express on z990 and z890
  - Up to 50% MB/sec for duplex large sequential R/Ws*
  - Up to 40% small block I/Os per second*
- Two operating modes (**no FCV support**)
  - Defined on a port basis
  - FC (Fibre Channel): Native FICON and FICON CTC
  - FCP (Fibre Channel Protocol): SCSI LUN access for Linux, z/VM and z/VSE environments
- Connectivity options
  - 1 or 2 Gbps, auto-negotiated
  - Can be shared among LPARs, and defined as a spanned channel
  - FICON point-to-point or switched
  - FICON cascade - two directors
  - FCP SAN fabric - one or more directors
- Available – January 28, 2005

FICON Express2
Fourth Generation

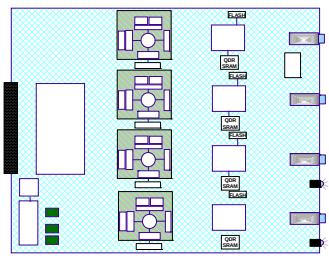

#3319 – LX , #3320 - SX

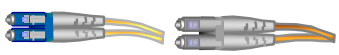

LC Duplex Connectors

*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.

FICON Express2

- Connectors - LC Duplex (same as FICON Express)
- LX - 9 micron single mode fiber
  - Maximum unrepeated distance 10 km (6.2 miles)
  - Maximum repeated distance 100 km (62 miles)
  - Supports MCP cables to reuse multimode infrastructure
    - –Only at 1Gbps, not at 2 Gbps
    - –Maximum unrepeated distance up to 550 meters
  - Receiving port must also be LX
- SX - 50 or 62.5 micron multimode fiber
  - Maximum unrepeated distance 120 to 500 meters depending on multimode fiber specification
  - MCP cables not applicable
  - Maximum repeated distance 100 km (62 miles)
  - Receiving port must also be SX
- Maximum number of features supported
  - z990 – 60 features, 20 per cage, 240 channels
    - –48 features, Model A08
  - z890 – 20 features, 80 channels
    - –16 features, 64 channels on z890 capacity setting 110

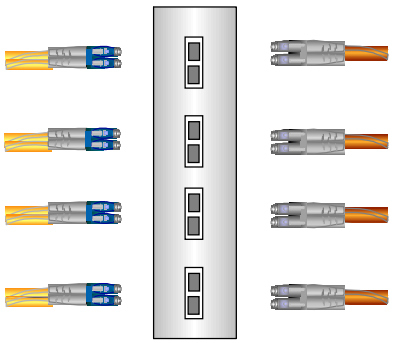

#3319 – LX or #3320 - SX

All LX **All SX**

FICON Express2 Performance (FC channel)

I/Os per second*
4k block size, Channel 100% utilized

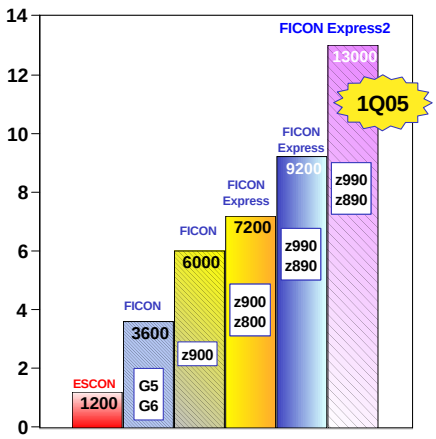

MB/sec throughput (Full Duplex)*
Large Sequential R/W mix

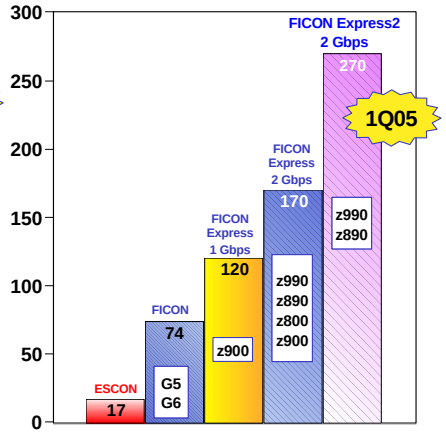

*This performance data was measured in a controlled environment on a z990 running an I/O driver program under z/OS 1.6. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

z990 and z890 FCP Attachment Options

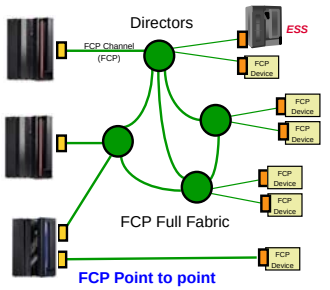

- **FCP Full Fabric Connectivity**
  - Homogeneous, single vendor fabric
  - Fibre channel directors, switches

NEW!

- FCP point to point – Designed to support all FICON features on z9-109, z990 and z890
  - **Direct attachment to FCP CU port**
  - **MCLs will be needed**

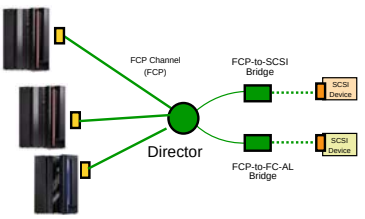

- **FCP switched to SCSI Bridge**
  - FCP-to-SCSI Bridges
  - FCP-to-FC-AL bridge

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

FICON Express2 – minimum software

- FICON Express2 (CHPID Type FC), on z890 and z990 including Channel-To-Channel (CTC), requires at a minimum:
  - z/OS 1.3 and z/OS.e 1.3
    - HCD APAR OA09114 and HCM (optional feature) APAR IR54497
    - For FICON purge path extended: z/OS and z/OS.e 1.4, and later, with PTFs for APAR 0A06846 and EREP APAR IR51695.
  - z/VM 3.1, and 4.3, 4.4, 5.1 and later.
    - HCD APAR VM63610 is required
  - z/VSE 3.1 (planned March 4, 2005) and VSE/ESA 2.6 and later
  - TPF 4.1 at PUT 16 and later
  - Linux on zSeries
    - The currently available distributions: SUSE SLES 8 and SLES 9, Red Hat RHEL 3
- FICON Express2 (CHPID Type FCP) on z890 and z990 for support of SCSI disks requires at a minimum (refer to 2084 and 2086 PSP buckets for any required service):
  - z/VSE 3.1 (Planned March 4, 2005)
  - Linux on zSeries
    - The currently available distributions: SUSE SLES 8 and SLES 9, Red Hat RHEL 3
  - z/VM 5.1 (for z/VM install, IPL, and operation from SCSI disks)
  - z/VM 4.4, 5.1 and later, for
    - Performance Assist for Adapter Interruptions
    - Performance Assist for V=V Guests
    - Guest IPL from SCSI devices
    - FCP LUN Access Control - APAR VM63328 is required
  - z/VM 4.3, 4.4, 5.1 and later for Linux or z/VSE as a guest under z/VM.
- For CHPID Mapping (optional), updated CHPID Mapping Tool from Resource Link

zSeries 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
  - FCP (Fibre Channel Protocol): SCSI devices in Linux environments
- Replaced by FICON Express2
  - Still fully supported on z990
  - Carried forward in upgrades to z990
  - **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
- Maximum Numbers for zSeries 990 and 890
  - z990 - 60 features, 20 per cage for 120 channels
    - –Model A08 - 48 features for 96 channels
  - z890 - 20 for 40 channels
    - –Model 110 - 16 for 32 channels

FICON Express
3rd Generation

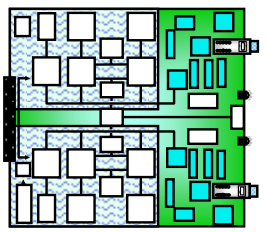

FC 2319 (LX), 2320 (SX)

z990 and z890 OSA-Express2

- Newest member - 10 Gigabit Ethernet LR (long reach)
  - One port per feature
  - 9 micron single mode fiber, SC Duplex connector
- New - Gigabit Ethernet features
  - 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
  - Designed to achieve line speed - 1 Gbps in each direction
- Support offered by both 10 GbE and GbE:
  - Queued Direct Input/Output (QDIO) for TCP/IP traffic only
    - –Use TN3270 or Enterprise Extender for SNA traffic
  - **Layer 2 support** for flexible and efficient data transfer
  - **Up to 640 TCP/IP stacks** for improved virtualization
  - **Large send** for CPU efficiency
  - **Concurrent LIC update** to help minimize network traffic disruption
- CHPID type for all features and functions listed is OSD
- **Availability – January 28, 2005**

10 Gigabit Ethernet
Feature 3368

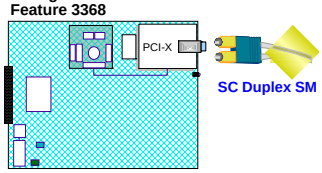

Gigabit Ethernet
Features 3364 (LX), 3365 (SX)

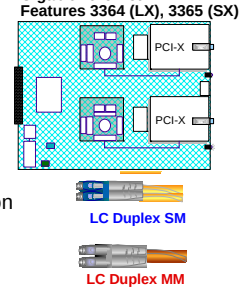

OSA-Express2 Large Send Support

- OSA-Express2 (GbE and 10 GbE) (Available January 28, 2005)
- Segmentation of IP packets done by OSA-Express2, not TCP/IP stack
  - Offloads the TCP segmentation processing from CPs or IFLs
  - Host code path length reduced
  - Sends 64 KB blocks to OSA-Express2
- Processing performed by OSA-Express2
  - TCP/IP checksum processing
  - TCP packet processing
  - Sends out 1.5 KB packets (1492 byte)
- QDIO mode only (CHPID type OSD)
  - For outbound traffic only
  - For IPv4, IPv6
  - For unicast datagrams
- Support planned:
  - **z/OS Communications Server with z/OS or z/OS.e 1.7** (1.6 with PTF) for TCP/IPv4 traffic only
  - **Linux® on zSeries** with code IBM intends to deliver Open Source in early 2005

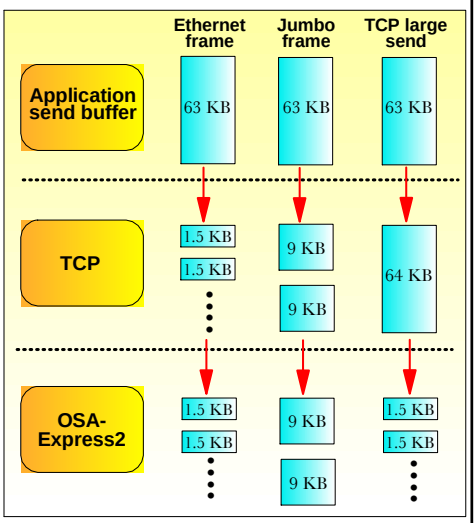

640 TCP/IP stacks for improved virtualization

- Exclusive to OSA-Express2 (Planned January 28, 2005)
  - 640 TCP/IP stacks per OSA-Express2 port/CHPID
- Doubled the number of home IP addresses
- For hosting more images on zSeries
- Reduces the number of OSA features required to host multiple images
- Exclusive to OSA-Express2 (GbE, 10 GbE)
  - QDIO mode only (CHPID type OSD)
- Support planned by
  - z/OS and z/OS.e 1.6 with HCD APAR OA09114, HCM (optional feature) APAR IR54497, and CommServer APARs (**Check PSP**!)
  - z/VM V5.1 with APARs VM63524 and PQ91421
  - Linux on zSeries with code IBM intends to deliver through Open Source in early 2005
  - z/VSE 3.1 (March 4, 2005) and VSE/ESA 2.6 and later with APAR DY46170

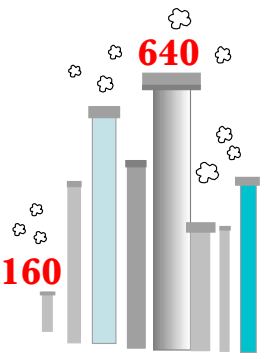

Limits	S/390 G5/G6	z900	z900	zSeries	z990	z990, z890 OSA-E	z990, z890 OSA-E2
OSD		Dec 00	Oct 01	May 02	Jun 03	Oct 04	Jan 05
Subchannels per stack	3	3	3	3	3	3	3
IP Stacks per port/CHPID on server	15	80	80	80	160	160	640 $
Subchannels per port	240	240	240	240	480	480	1920 $
IP stacks per LPAR	15	80	80	80	84	160	640 $
Devices per LPAR	240	240	240	240	254	480	1920 $
Maximum Control Units Supported	1	1	1	1	1	16	16

$ If multiple priorities for queues is enabled (one to four QDDIO priorities) the maximum remains at 160 stacks/480 devices

z890/990 OSA-Express2/OSA-Express new build features

- Up to 48 network connections
- Up to 40 network connections - z890
  - 24 on z890 capacity setting 110
- Choose from 5 features
  - **OSA-Express2 10 GbE, GbE LX and SX**
  - OSA-Express 1000BASE-T Ethernet
    - Same Cat 5 cable as Fast Ethernet
  - Token-Ring (4/16/100 Mbps)
    - Cat 5 copper cable
- Modes of Operation for 1000BASE-T Ethernet, Token-Ring
  - QDIO = TCP/IP traffic only
    - TN3270 or Enterprise Extender for SNA traffic
  - Non-QDIO = TCP/IP and/or SNA/APPN®/HPR
  - OSA-ICC for 1000BASE-T only
- **SOD - z990/z890 are the last zSeries servers to support Token-Ring OSA - new build, upgrade, MES, or carry forward**

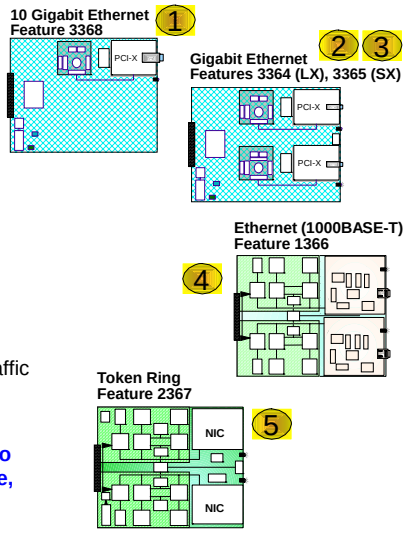

OSA-Express2 Support Requirements

- OSA-Express2 Gigabit Ethernet requires:
  - z890 or z990 hardware LIC support for GA2/4 (January 2005)
  - z/OS 1.3 or z/OS.e 1.3 or later with HCD APAR OA09114, HCM (optional feature) APAR IR54497, and CommServer APARs (**Check PSP**!)
  - z/VM 3.1 or z/VM 4.3 or later with APARs VM63610, VM63524, and PQ91421
  - z/VSE 3.1 (planned March 4, 2004) and VSE/ESA™ 2.6 or later plus service for APAR DY46170
  - TPF 4.1 PUT13 with service for APAR PJ27333
  - Linux on zSeries with Gigabit Ethernet support:
    - SUSE LINUX SLES 8 or 9, Red Hat RHEL 3, Turbolinux TLES 8 or Conectiva CLEE
  - **See the 2084DEVICE or 2086DEVICE PSP for additional service required**
- OSA-Express2 10 Gigabit Ethernet requires:
  - z890 or z990 hardware LIC support for GA2/4 (January 2005)
  - z/OS 1.3 or z/OS.e 1.3 with APARs listed above
  - For Checksum Offload, z/OS or z/OS.e 1.5 or later
  - z/VM 3.1 or z/VM 4.3 or later with service listed above
  - z/VSE 3.1 (planned March 4, 2004) and VSE/ESA 2.6 or later service for APAR DY46170
  - TPF 4.1 PUT13 with service for APARs PJ27333 and PJ29930
  - Linux on zSeries with code IBM plans to deliver as Open Source in early 2005
  - **See the 2084DEVICE or 2086DEVICE PSP for additional service required**
- **For CHPID Mapping (optional), updated CHPID Mapping Tool from Resource Link**

z990 CFCC Level 14 – z890/990 GA2/4

- Function and Potential Benefit
  - Contains improvements to the CF dispatcher and internal serialization mechanisms designed to better manage coupled workloads
- Requirements and Support
  - Requires z890/990 hardware LIC support for GA2/4 (October 29, 2004)
  - z/OS 1.3 or z/OS.e 1.3 and higher
    - Optional APAR fix OA08742 to allow sysplex connectors to request structure allocation in a Level 14 Coupling Facility
  - z/VM 3.1 and z/VM 4.3 and higher for virtual CF support
- CF Storage Sizing with CFCC level 14
  - May increase storage requirements
  - Use CFSIZER tool to determine: www.ibm.com/servers/eserver/zseries/cfsizer/

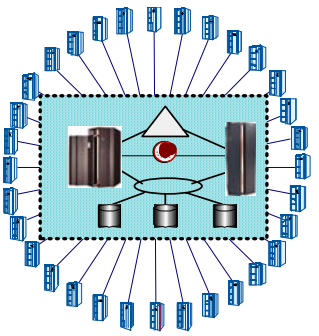

1 to 32 Systems

zSeries continues to meet the requirements for advanced clustering

z990 GA3 availability enhancement for CFCC changes

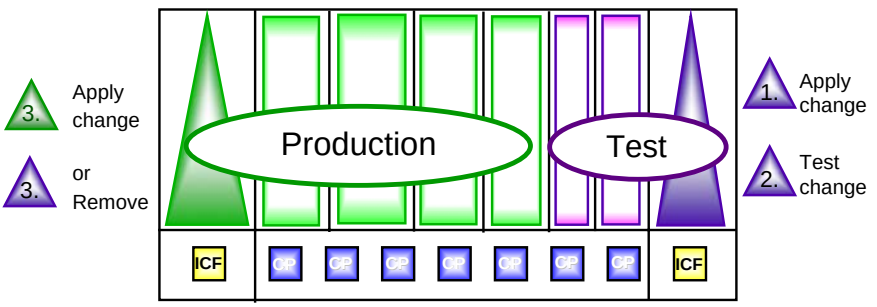

- **Apply previously disruptive CFCC changes with little disruption to z990**
  - Disruption occurs one CFCC LPAR at a time to activate or remove a change
  - Allows rolling CFCC maintenance across CF LPARs
    - Similar to rolling z/OS maintenance across OS images
  - Helps reduce the requirement to isolate test CFs from production OS/CF images

Server Time Protocol (STP) Overview

- Designed to provide capability for multiple IBM System z9, zSeries 990 and 890 and servers to maintain time synchronization with each other
  - Does not require the 9037 Sysplex Timer if all servers STP capable
- Timing information transmitted over peer mode links: ISC-3, ICB-3 and ICB-4
- Supports a multi-site timing network of up to 100 km (62 miles)
  - Allows a Parallel Sysplex to span up to 100 km
- May reduce the cross-site connectivity required for a multi-site Parallel Sysplex
- Can coexist with an External Time Reference (ETR) network (9037 based)
  - Mixed Timing Network
- Designed to allow use of dial-out time services to set the time to international time standard (UTC) as well as adjust to UTC
- Planned to be available as an optional feature (RPQ?) on 2094, 2084 and 2086
- Prerequisites
  - z9-109 HMC LIC
  - z/OS V1.7

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

STP Implementation Assistance Program (IAP)

- Objective
  - Accelerate the adoption of STP with IBM assistance
- Assistance planned
  - Consultation
  - Review of migration plans
  - Technical support
- Planned Availability dates
  - October 2005 (z990, z890)
  - 1Q06 (2094)

z990 and z890 HMC configurations and recommendations

Currently orderable HMCs (November 2004):
HMC #0079 – Dual Ethernet
HMC #0080 – Ethernet/Token-Ring
Note: Both LANs on HMC cannot connect to same CEC

Old HMCs OK:
HMC #0077, #0078
HMC #0076 conversion of #0075
HMC #0074, #0075
HMC #0073 + #0047 DVD
Note: Only #0076 and #0077 have dual Ethernet support

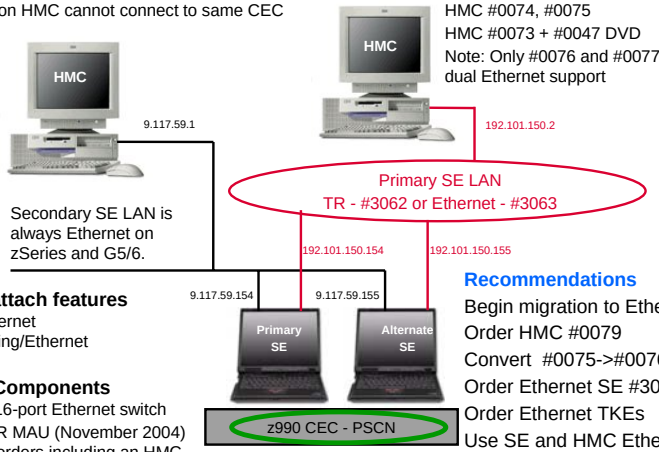

z990 SE LAN attach features
#3063 – Dual Ethernet
#3062 – Token-Ring/Ethernet

HMC/SE LAN Components
#0089 – 10/100, 16-port Ethernet switch
#0088 – 8 port T/R MAU (November 2004)
Note: Offered on orders including an HMC

Recommendations
Begin migration to Ethernet
Order HMC #0079
Convert #0075->#0076
Order Ethernet SE #3063
Order Ethernet TKEs
Use SE and HMC Ethernet available on older equipment

z9-109 HMC Support for z990 and z890

HMC orderable on 2094:
HMC #0079 – Dual Ethernet
Note: Both LANs on HMC cannot connect to same CEC

Other HMCs OK:
HMC #0077, #0078, #0080
HMC #0076 conversion of #0075
HMC #0075
Note: **Ethernet attach standard.**
Only HMCs #0076 and #0077 have dual Ethernet support

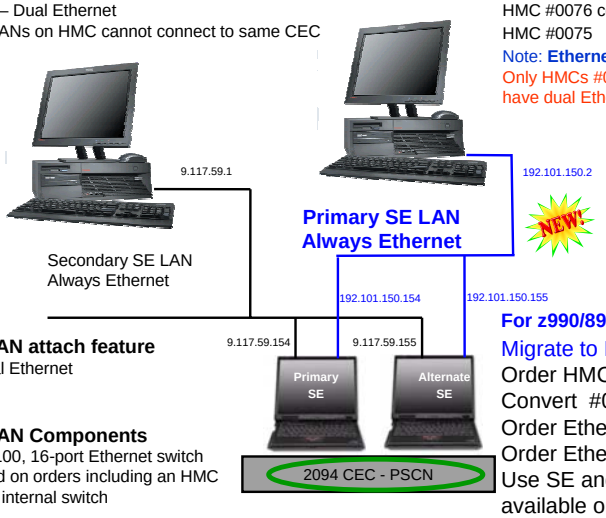

2094 SE LAN attach feature
#3065 – Dual Ethernet

HMC/SE LAN Components
#0089 – 10/100, 16-port Ethernet switch
Note: Offered on orders including an HMC
2094 has no internal switch

For z990/890, z900/800:
Migrate to Ethernet
Order HMC #0079
Convert #0075->#0076
Order Ethernet SE #3063
Order Ethernet TKE
Use SE and HMC Ethernet available on older equipment

Non-OS2 Operating System "View Licenses"

This product contains certain code packages that are licensed pursuant to the terms of the GNU General Public License ("GPL") and/or the GNU Lesser General Public Licnese "LGPL"). Those terms are reproduced below for your reference. The code packages that are licensed under the GPL or LGPL include:

aaa_base, aaa_skel, apmd,at, bash, binutils, cloop, diffutils, dosfstools, e2fsprogs, eject, ethtool, fbset, fileutils, fillup, findutils, ftplib, gawk, gdbm, glib, glibc, gpm, grep, grub, gzip, hotplug, hwinfo, imlib, iproute2, iptables, iptraf, iputils, kbd, kernel, ksymoops, l2tpd, less, libelf, libstdc++, libxml , lkcdutils, logrotate, m4, makedev, mkinitrd, mktemp, modutils, mtools, net-tools, netcfg, nfsutils, openswan, pam, pciutils, pcmcia, permissions, procmail, ps, rpm, sed, setserial, sh-utils, sysconfig, syslinux, syslogd, sysvinit, tar, textutils, udftools, udf, utempter, util-linux, vsftpd, and xfntscl2.

Note: Source code to any of the above-listed packages is available upon written request to the following address: IBM Corporation, Linux Technology Center, Dep't 7UDA, 11501 Burnett Road, Austin, TX 78758.

The terms of the GPL and LGPL follow below: ….

z9-109 HMC Compatibility Driver and MCL Levels

- zSeries 990 and 890
  - Driver 55k with MCL 132 to EC J13486
- zSeries 900 and 800
  - Driver 3Gf with MCL 194 to EC J11213
- S/390 9672 G6 and G5
  - Driver 26 with:
    - MCL 174 to EC F99918

As of 8/14/05!

These MCLs are concurrent.

Earlier machines are not supported.

Statements of Direction for z990 and z890

- ✓ **April 2004: Hardware Management Consoles (HMCs):** Beginning with the next zSeries server, after the IBM zSeries 890 and 990, all new HMCs on all currently marketed zSeries servers are intended to become closed platforms. They will support only the HMC application and not the installation of other applications such as the IBM ESCON® Director and the IBM Sysplex Timer® console applications.
  When available, the next-generation HMC is expected to communicate only with G5 Servers, and above (Multiprise® 3000, G5/G6, z800, z900, z890, z990). TCP/IP is intended to be the only communications protocol supported.

- ✓ **April 2004: ISC-3s in compatibility mode:** IBM intends z890 and z990 to be the last family of zSeries servers to support:
  - **Any** sysplex coupling link connectivity to S/390 G5/G6 or earlier
    - Attachment of ISC-3 links to HiPerLinks (ISC-2) on G5/G6 servers
  - Compatibility mode (CHPID types CFS and CFR) definitions for ISC-3 links

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Any reliance on this Statement of Direction is at the relying party's sole risk and will not create any liability or obligation for IBM.

Statements of Direction for z990 and z890

- ✓ **October 2004: OSA-Express Token-Ring not to be offered:** The zSeries 890 and 990 are expected to be the last zSeries servers to offer a Token-Ring feature. It is intended that the OSA-Express Token-Ring feature will not be available for ordering on a new build or upgraded server, or for carrying forward on an upgrade.
  A migration from a Token-Ring to an Ethernet environment should be a part of all strategic LAN planning.
- ✓ **April 2004: Token-Ring on HMC, SE, TKE workstation, IBM 2074:** The z890 and z990 will be the last zSeries servers to offer Token-Ring adapter features on the Hardware Management Consoles (HMCs), Support Element (SEs), and Trusted Key Entry (TKE) workstations. The IBM 2074 Model 3 Console Support Controller will be the last controller to offer Token-Ring adapter features.
  IBM zSeries is making these statements to allow enterprises sufficient opportunity to prepare for a migration to Ethernet environments.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Any reliance on this Statement of Direction is at the relying party's sole risk and will not create any liability or obligation for IBM.

End of Presentation

RETURN TO INDEX