



G06

# Ride the System z9 Business Class Bus

Connie K. Beuselinck

**IBM**  
**SYSTEM z9 AND zSERIES EXPO**  
**October 9 - 13, 2006**

Orlando, FL

IBM Hardware Product Planning

[conniek@us.ibm.com](mailto:conniek@us.ibm.com)

Poughkeepsie, NY U.S.A



# Trademarks

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

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

\* 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.

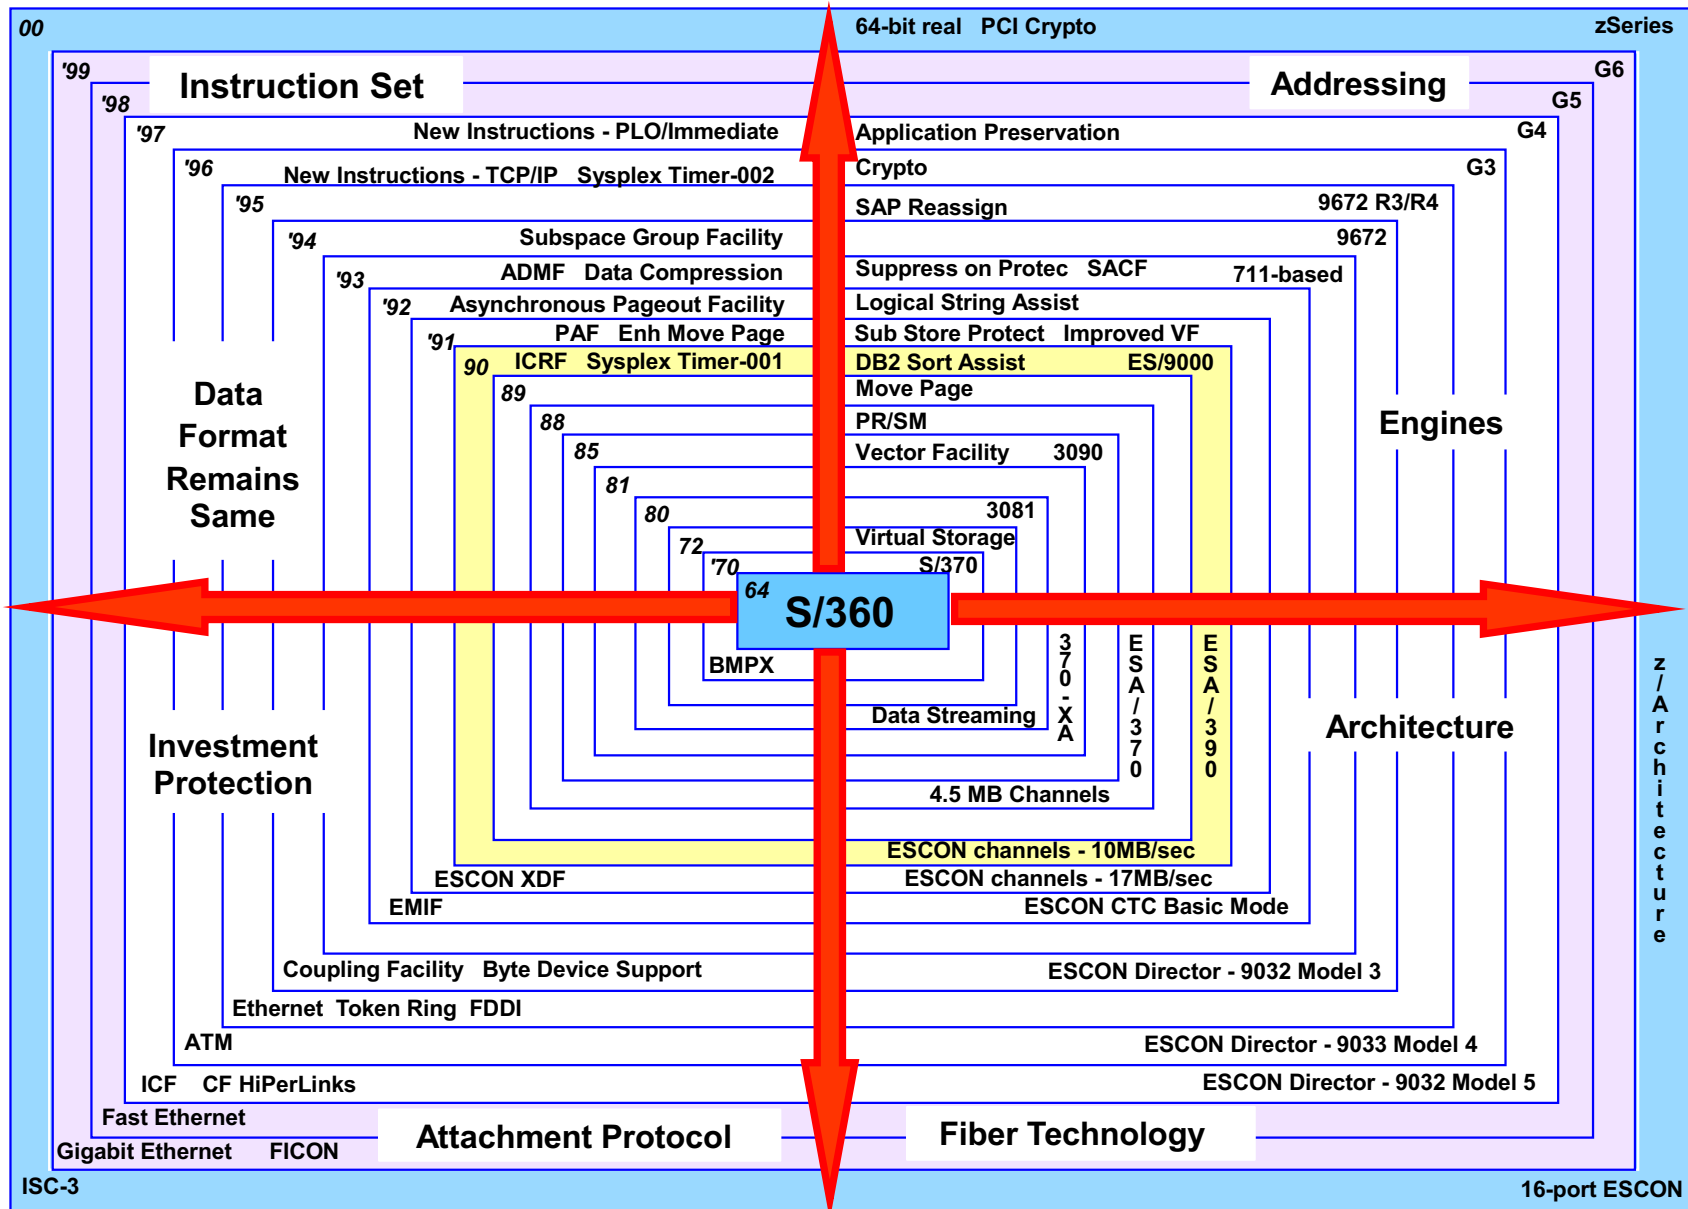


# Glossary


Acronym	Full Name	Use
CFCC	Coupling Facility Control Code	Parallel Sysplex
CP	Central Processor	General Purpose PU
CPACF	CP Assist for Cryptographic Function	Cryptography
ESCON	Enterprise Systems Connection	Storage, Printers
ETR	External Time Reference	Sysplex Timer, STP, Parallel Sysplex
FCP	Fibre Channel Protocol	SCSI devices (z/VM, z/VSE, Linux on System z9)
FICON	Fibre Connection	Storage, Printers
IC	Internal Coupling Channel	Parallel Sysplex
ICB	Integrated Cluster Bus	Parallel Sysplex
ICF	Internal Coupling Facility	PU for Coupling Facility Control Code
IFL	Integrated Facility for Linux	PU for use by Linux on System z
ISC-3	InterSystem Channel-3	Parallel Sysplex
MBA	Memory Bus Adapter	Part of Central Electronic Complex (CEC)
MCM	Multichip Module	Part of CEC
OSA	Open Systems Adapter	Local Area Network connectivity
PCI	Peripheral Component Interconnect	Local bus standard (used with OSA and Crypto)
PCICA	PCI Cryptographic Accelerator	Cryptography
PCIXCC	PCI-X Cryptographic Coprocessor	Cryptography
PU	Processor Unit	Becomes a CP, ICF, IFL, zAAP, zIIP
SCSI	Small Computer System Interface	Storage – fixed block devices
STI	Self-Timed Interconnect	Internal host bus on System z
TKE	Trusted Key Entry	Cryptography – key management system (feature)
zAAP	System z9 Application Assist Processor	PU for Java execution environment
zIIP	System z9 Integrated Information Processor	PU for use by DB2 UDB for z/OS V8 workloads




# Evolution of the mainframe




# Data center – 1964. There was one protocol!



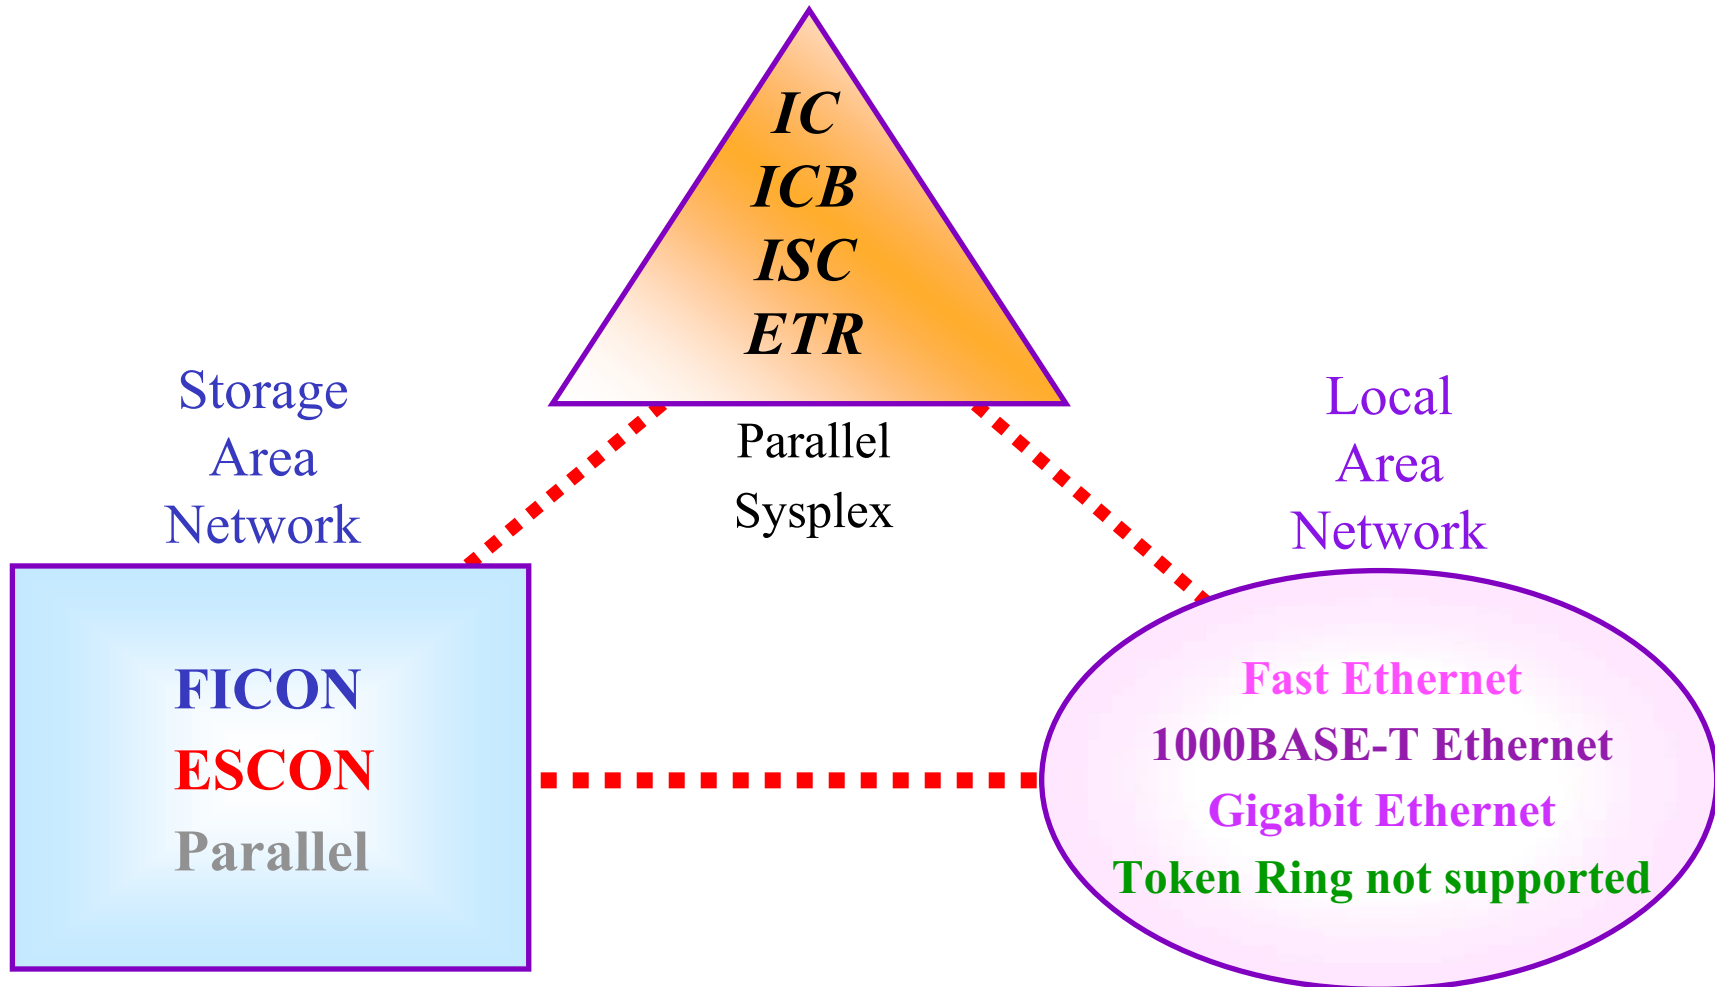


# Getting connected





# ..... multiple architectures and protocols





# What's different with System z9 BC?

**STI speed**

**STI granularity**

**I/O domains**

**Availability**

**I/O cage capacity  
for Crypto, OSA, FICON**

**FICON Express4**

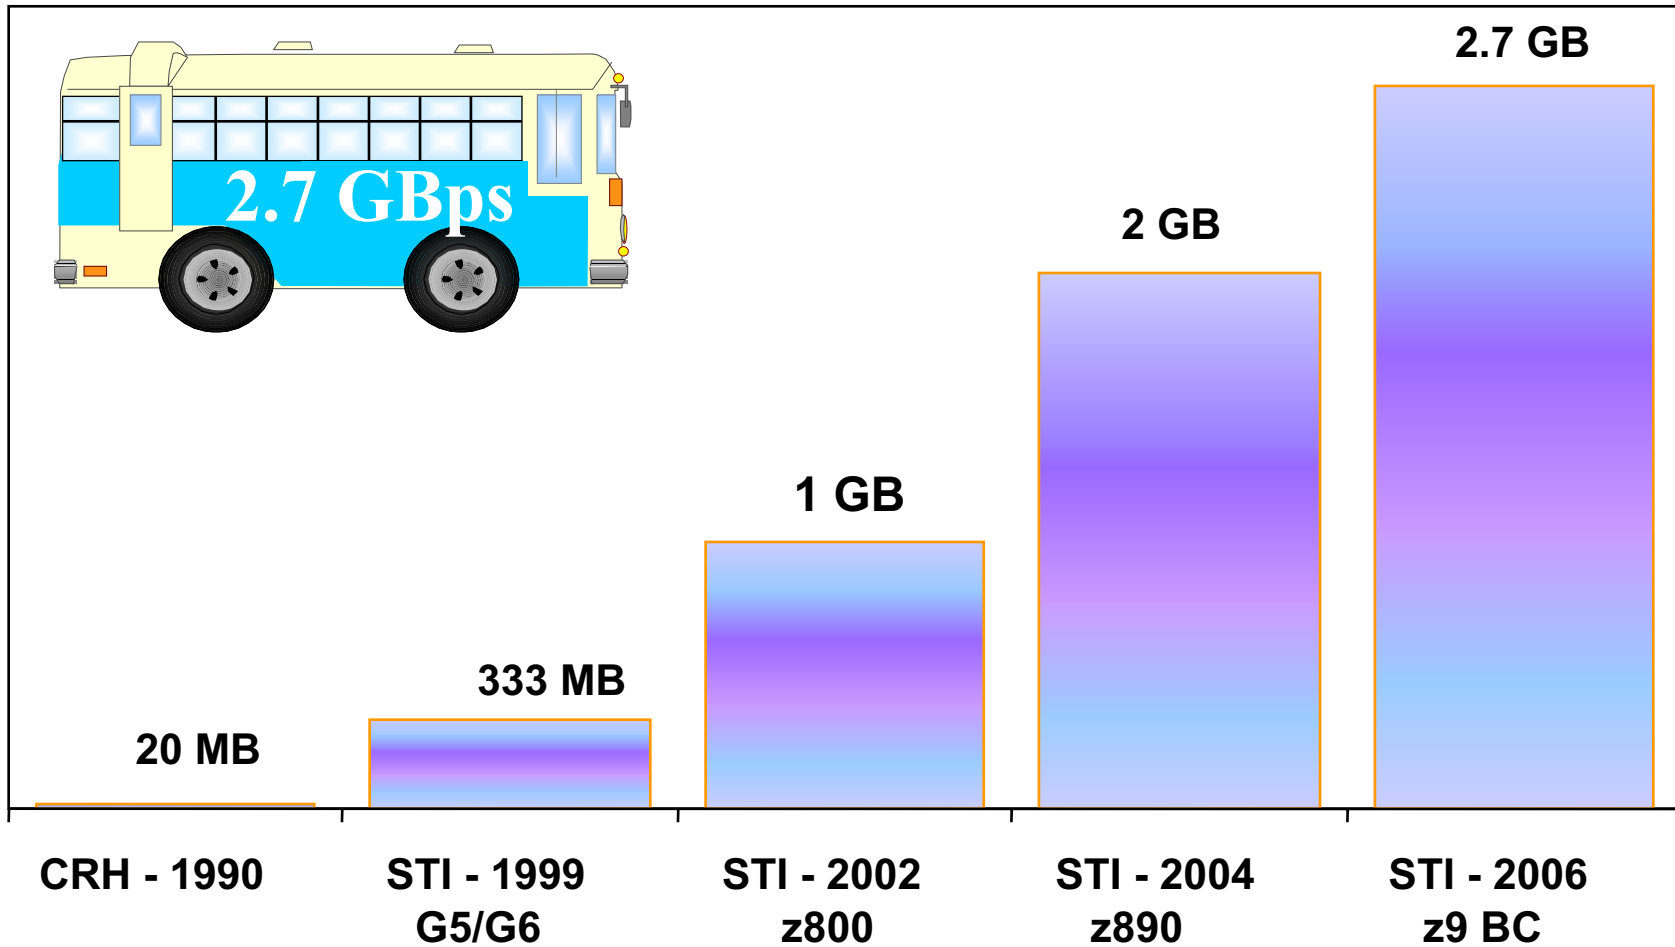
**Configurable  
Crypto Express2**







# I/O subsystem bus bandwidth



CRH – Channel Request Handler bus  
STI – Self-Timed Interconnect bus



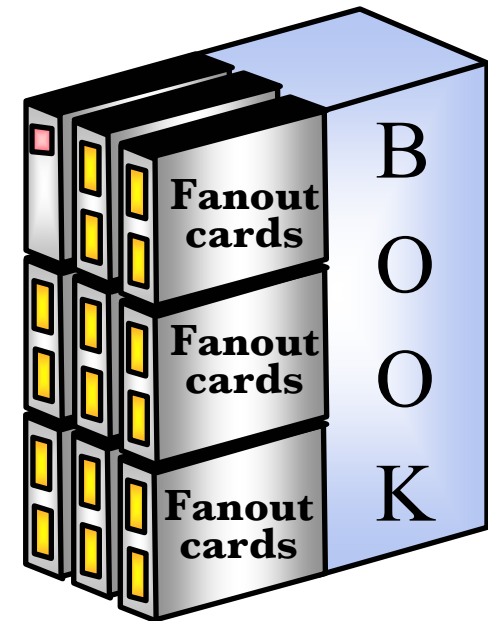
# I/O subsystem bus bandwidth

★ 1990 - CRH bus = 20 MegaByte  
(Channel Request Handler)

★ 1999 - STI bus = 333 MegaByte  
(Self-Timed Interconnect)

One STI  
Shared  
by up to  
8 features

Up to 16 STIs  
on book



★ 2000 - STI bus = 1 GigaByte  
– zSeries 800 - Up to 4 STIs for I/O

★ 2003 - STI bus = 2 GigaByte  
– zSeries 890 - Up to 8 STIs for I/O

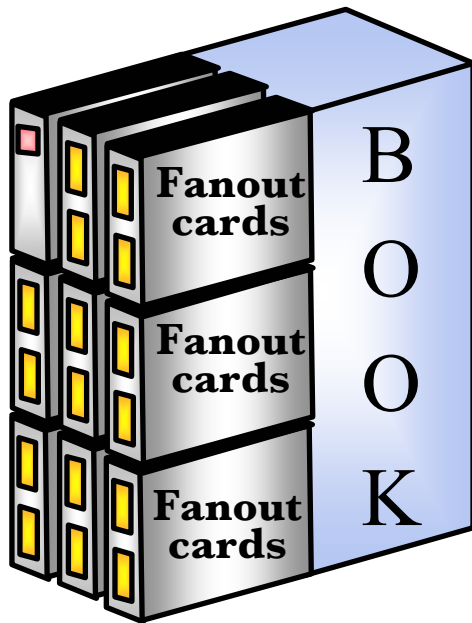
★ 2006 - STI bus = 2.7 GigaByte  
– System z9 BC – Up to 16 STIs for I/O

One STI  
Shared  
by up to  
4 features

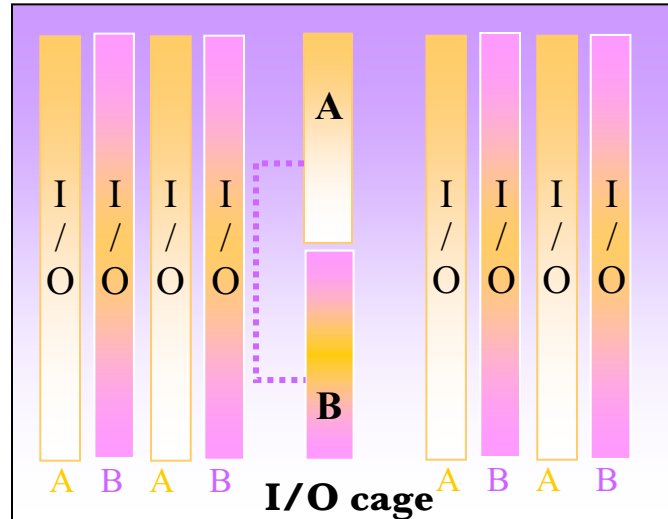
–An additional STI in I/O cage for redundancy (now 8)



# z9 BC new infrastructure




+



=


Availability  
Scalability






# z9 BC – Under the covers

A Frame




# z9 BC Processor Book Layout

Up to 8  
Hot pluggable MBA/STI  
fanout cards




Front View




Memory Cards  
Up to 64 GB

MCM




Side View

- 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 STIs in I/O cage – Top view



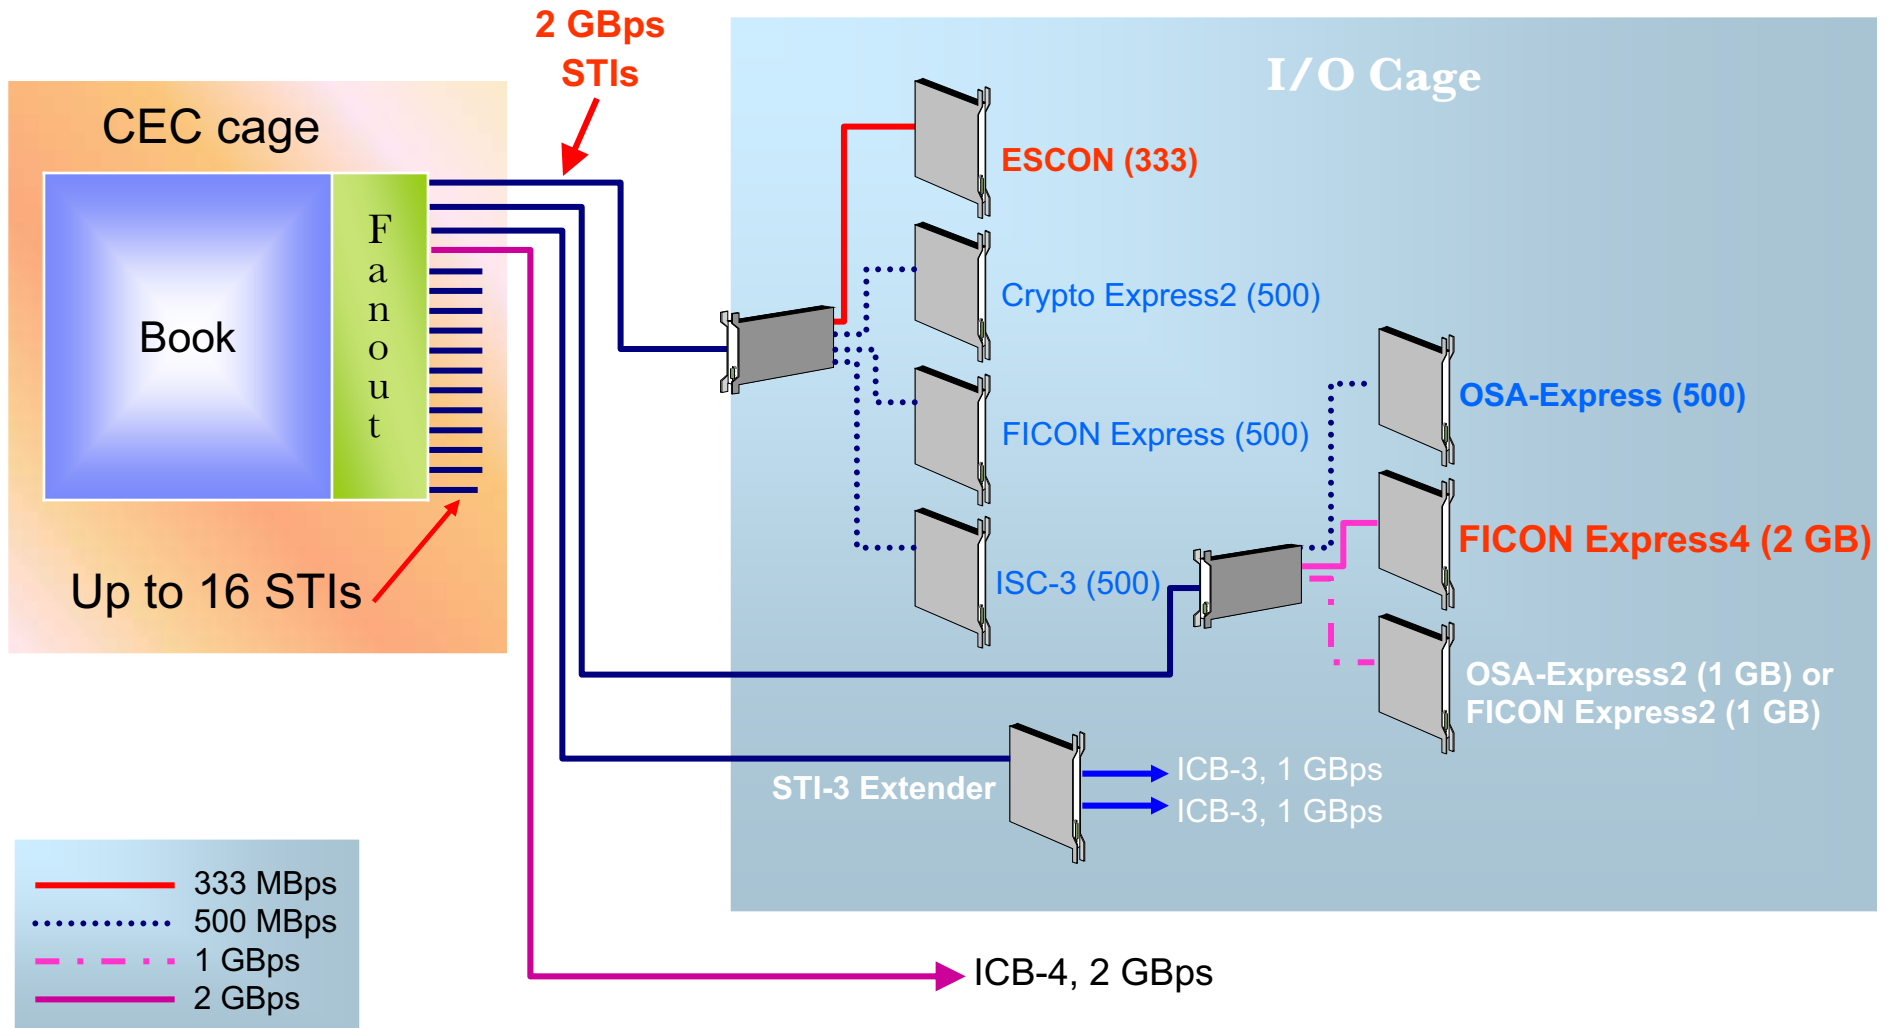
- I/O cage
  - 8 STIs
    - G' is for redundancy
- STI speeds for downstream channels/ports/links
  - 333 MBps – ESCON
  - 500 MBps
    - Crypto Express2
    - FICON Express
    - ISC-3
    - OSA-Express
  - 1 GBps
    - FICON Express2
    - OSA-Express2
  - 2 GBps
    - FICON Express4

STI-MPs



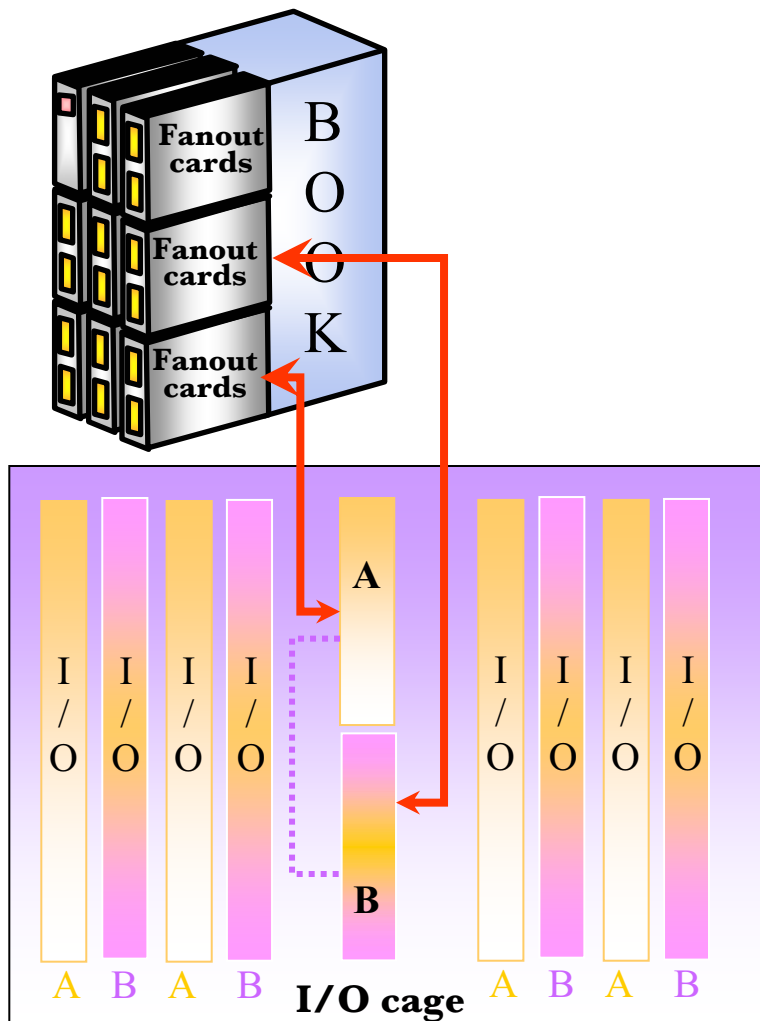


# z9 BC STIs servicing I/O cage and features





# z9 BC STI hot-plug; Redundant I/O interconnect



- **STI-MP (multiplexer) card, hot-plug**

- Each attaches to four I/O features
  - I/O domains A – G
  - Alternate path to **second STI-MP**

- **Concurrency / automatic failover**

- Manual or system-initiated
  - ▶ Add, remove, move of MBA fanout cards
  - ▶ Repair of MBA fanout card, STI cable

- **Redundant I/O interconnect** .....







# z9 BC STI granularity



- **MBA fanouts hot-plugged into book**
  - **Concurrent upgrade, repair**
- **Up to 8 MBA Fanouts – up to 16 STIs**
  - **Two STIs per MBA Fanout**
- **Compare to z890 book package**
  - Two MBAs - 8 STIs



No MBA card hot-plugging





## MIDAW facility on z9 BC

- **Modified Indirect Data Address Word (MIDAW) facility**
  - ▶ Designed to increase throughput and reduce link overhead
- **Alternative to using CCW data chaining in channel programs**
- **Breaks the 2k or 4k boundary restriction**
  - ▶ Channel Control Word (CCW) architecture is enhanced
    - Can handle larger blocks of data
    - Reduces the number of data moves
    - Reduced chained data
- **Design "in line" with Fibre Channel architecture**
- **Designed to reduce**
  - ▶ ESCON and FICON I/O processing
  - ▶ Control unit overhead
- ▶ Supported by
  - ▶ ESCON (CHPID type CNC)
  - ▶ FICON features supported on z9 BC (CHPID types FC, FCV)
- **Potential exploiters**
  - ▶ Applications that use: DB2, VSAM, PDSE, HFS, zFS
  - ▶ z/OS V1.6 or V1.7 with PTFs





# MIDAW facility 4k DB2 page comparisons

## Non-EF datasets

- CCW READ 4K
- CCW READ 4K
- CCW READ 4K
- CCW READ 4K
- CCW READ 4K
- CCW READ 4K
- CCW READ 4K
- CCW READ 4K
- CCW READ 4K

...32 CCW's in total

## EF datasets

- CCW READ 4K
- CCW READ 32 byte suffix
- CCW READ 4K
- CCW READ 32 byte suffix
- CCW READ 4K
- CCW READ 32 byte suffix
- CCW READ 4K
- CCW READ 32 byte suffix
- CCW READ 4K
- CCW READ 32 byte suffix

...64 CCW's in total

## EF or non-EF datasets with MIDAWs

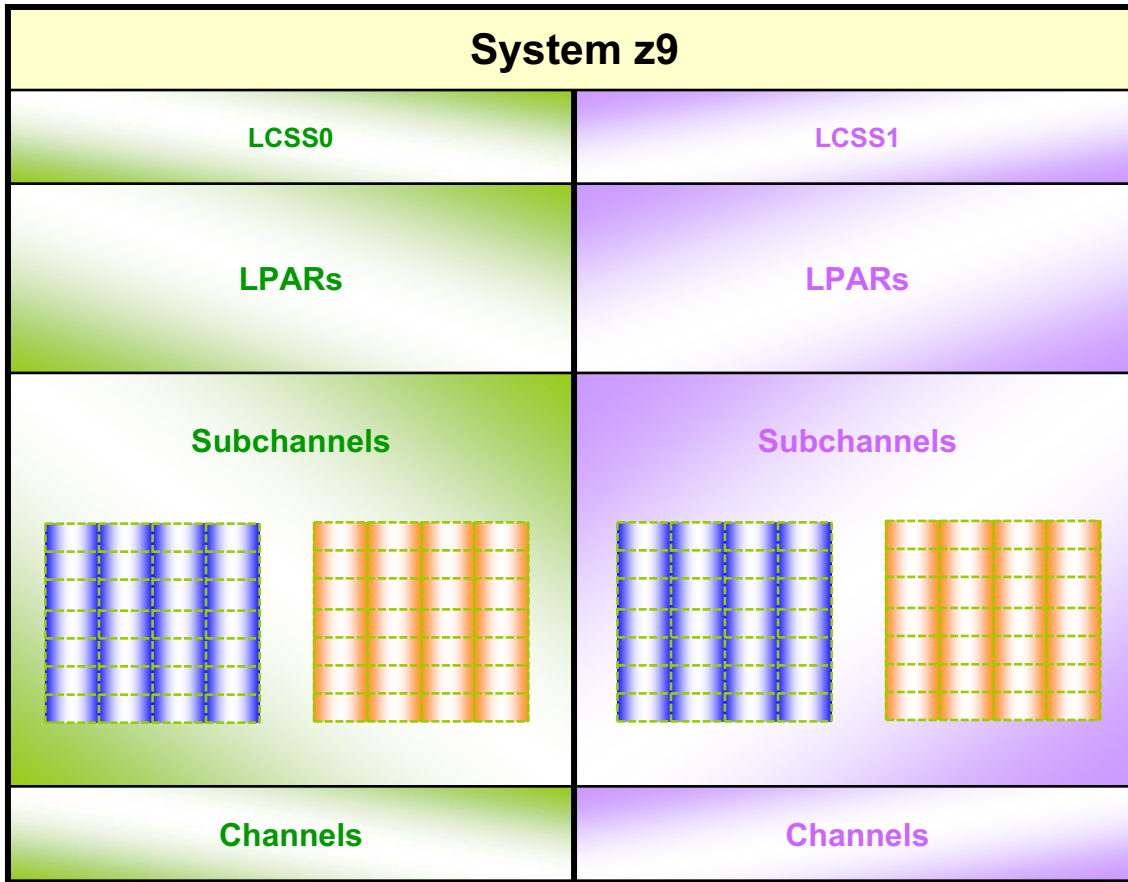
- Track level CCW
- Track level CCW
- Track level CCW
- Track level CCW

3 or 4 CCW's in total




# Multiple subchannel sets – two sets on z9 BC


- Subchannel – I/O device to hardware
  - Used by System Control Program (SCP)
    - Passes I/O requests from SCP to CSS
  - One or two subchannel sets per CSS
- PAV alias devices only in SS-1
    - Designed to be compatible with existing storage CUs that support PAV



**Logical Channel Subsystem (LCSS)**



**Base, SS-0  
63.75k**



**Aliases, SS-1  
64k**



## 63.75k subchannels on z9 BC

- **64k-1 subchannels in set 0**
  - ▶ **Previously 1024 (1k) reserved for system use**
  - ▶ **Returning 768 of 1k**
  - ▶ **For increased addressable storage**
  - ▶ **Now symmetry – server and storage**
    - **63.75k, System z9**
    - **63.75k, IBM TotalStorage DS8000 Series**
  - ▶ **All channel types**
  - ▶ **z/OS V1.4 + PTFs**
  - ▶ **z/VM V4.4 (HCD PTF)**
  - ▶ **Linux on System z9**
    - **SUSE SLES9, Red Hat RHEL4 distributions**
  
  - ▶ *Example: 3390 volume sizes*
    - *768 volumes of 54 GB/volume = 41 terabytes of increased storage*
      - *54 GB/volume \* 768 volumes = 41 TB*





## Program-directed re-IPL on z9 BC



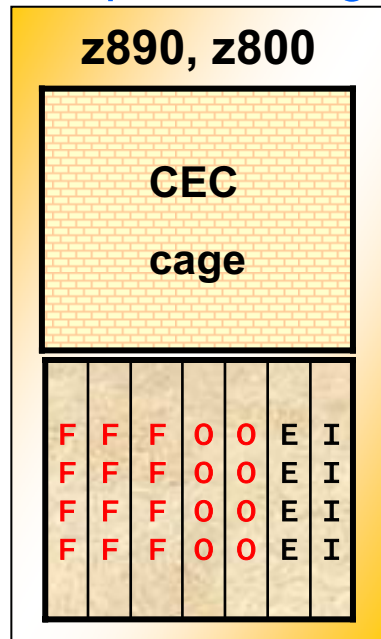
- **Linux running natively in a logical partition can re-IPL itself**
- **Supported for SCSI and ECKD devices**
- **To initiate a Re-IPL, Linux can determine**
  - ▶ **How it was loaded (i.e., via Channel Control Words (CCWs)) or via SCSI-type IPL, and**
  - ▶ **From where it was loaded in case of SCSI-type IPL (World Wide Port Name (WWPN) and Logical Unit Number (LUN) of the load device).**
- **Linux can request**
  - ▶ **That it be reloaded from the same load device**
  - ▶ **Using the same load parameters.**
- **Linux on System z9 – watch for releases from distribution partners**
- **Note: z/VM already supports an interface that allows a program running as a guest under z/VM to re-IPL itself.**



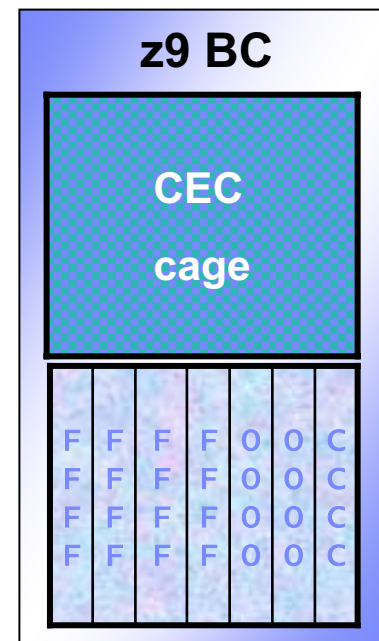
# Increased connectivity per I/O cage on z9 BC

- All supported FICON features - up to 28 in I/O cage
- OSA-Express2, OSA-Express, up to 24 in I/O cage
- Crypto Express2, up to 8 in I/O cage
- **Removed limitation of up to 20 FICON, OSAs, and Crypto in I/O cage**

z890, z800  
20 per I/O cage



z9 BC  
28 per I/O cage





E = ESCON, F = FICON, I = ISC-3, O = OSA, C = Crypto





# Speed zones on the information highway ("Best can do" throughput)

**Fast Ethernet**  
22 MBps  
(176 Mbps)




**FICON**  
270 MBps

**Gigabit Ethernet**  
231 MBps  
(1850 Mbps FDX)



**ESCON**  
Tape  
17 MBps  
Disk  
12 MBps  
Network  
5 MBps



**10 GbE**  
423 MBps  
(3380 Mbps FDX)

FDX = Full duplex






# Parallel Sysplex

---

## Coupling connectivity



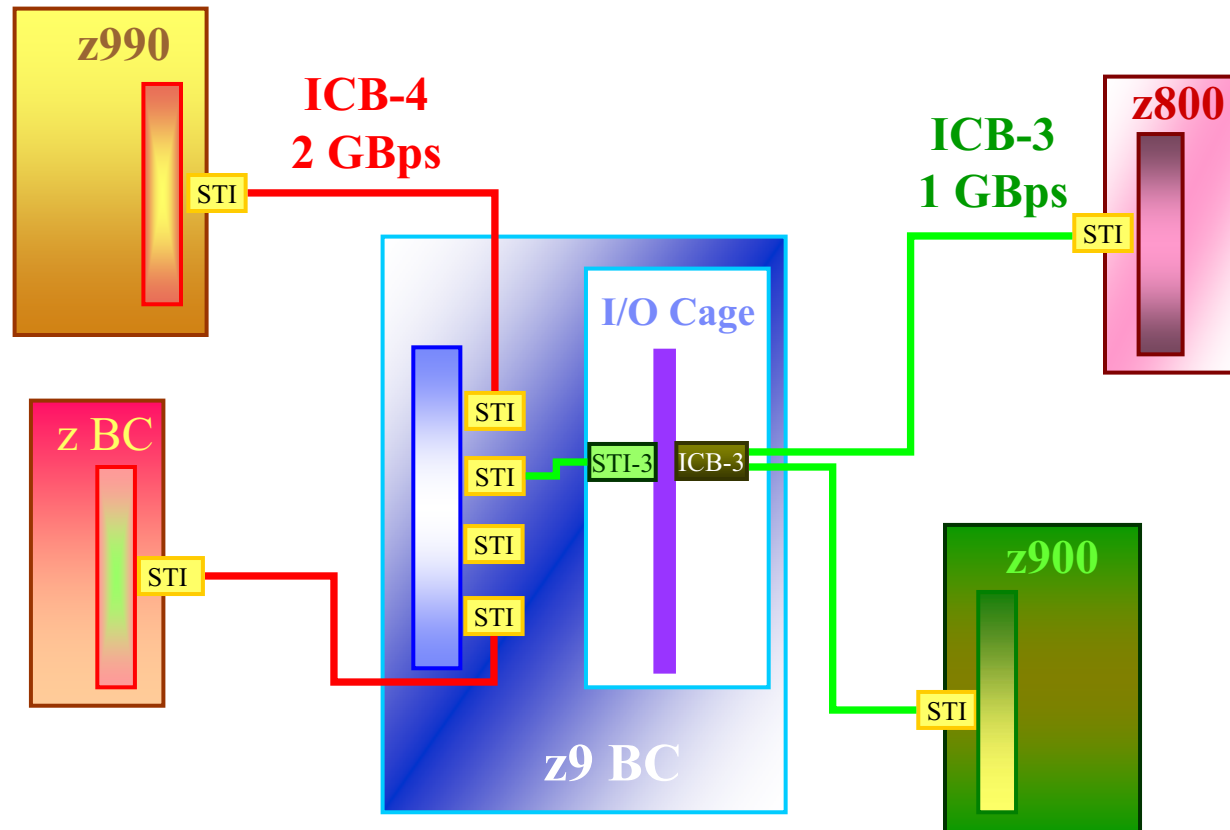


## Coupling link connectivity on z9 BC, z890, z800

Link type	Description	Use	Link data rate	Distance	z800 Max.	z890 Max.	z9 BC S07 Max.
IC	Internal Coupling channel	Internal	Internal speeds	N / A	32	32	32
ICB-2	Integrated Cluster Bus-2	z990, z900, z800 to G5 and G6	333 MBps	7 meters (23 ft)	8	N/A	N/A
ICB-3	Integrated Cluster Bus-3	z9 EC, z9 BC to z900, z800	1 GBps	7 meters (23 ft)	5	16	16
ICB-4	Integrated Cluster Bus-4	z9 EC, z9 BC z990, z890	2 GBps	7 meters (23 ft)	N/A	8	16
ISC-3	InterSystem Channel-3	z9 EC, z9 BC z990, z890 z900, z800	2 Gbps	10 km (6.2 m)	24	48	48

- The maximum number of Coupling Links combined cannot exceed 64 per server
  - (ICs, ICB-3s, ICB-4s, and active ISC-3 links)
- ISC-3 Peer (2 Gbps) and Compatibility (1 Gbps) mode on z890 and z800
- ISC-3 Peer mode only on z9 BC


## ICB-4, ICB-3 coexistence



- ICB-4 links support a link data rate of 2 GBps: z9 BC, z890
- ICB-3 links support a link data rate of 1 GBps: z9 BC to z900, z800
- ICB-2 links support a link data rate of 333 MBps: z890 to G5/G6 (not supported on z9 BC)

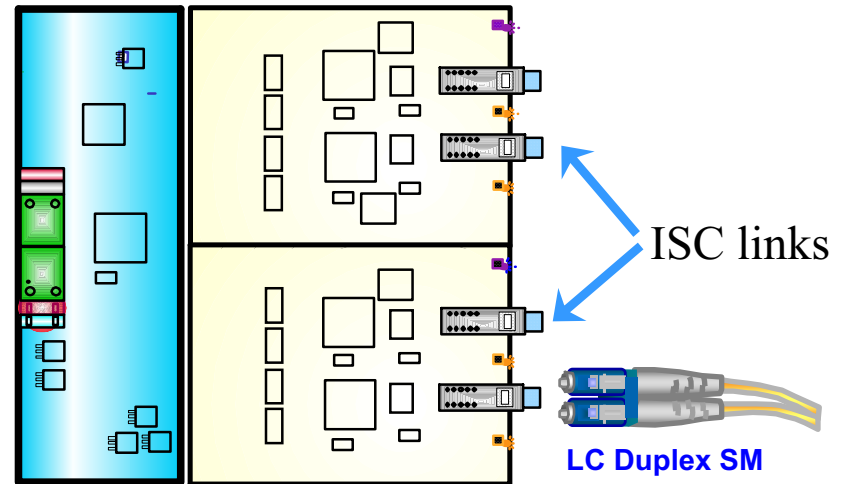
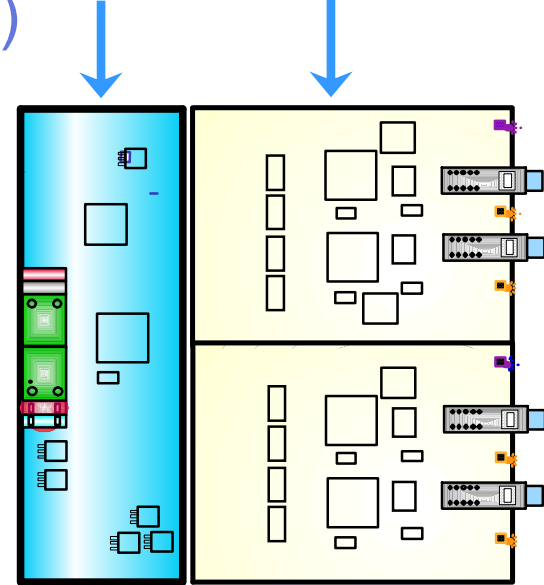


# System z InterSystem Channel-3 (ISC-3)



ISC-M = Mother card  
 ISC-D = Daughter card


ISC-M    ISC-D (hot-plug)



- Two modes of operation
  - Peer Mode (2 Gbps)
  - Compatibility Mode (1 Gbps) not supported on z9 BC




# External Time Reference (ETR) feature for attachment to Sysplex Timer



## Connector



MTRJ MM





## Coupling connectors and fiber optic cabling

Feature	Description	z800	z890	z9 BC	Connector	Cable type
<b>0219</b>	<b>ISC-3 link</b>	<b>X</b>	<b>X</b>	<b>X</b>	<b>LC Duplex</b>	<b>9 μ SM</b>
<b>6154</b>	<b>External Time Reference (ETR)</b>	<b>X</b>	<b>X</b>	<b>N / A</b>	<b>MT-RJ</b>	<b>62.5 μ MM</b>
<b>6155</b>	<b>External Time Reference (ETR)</b>	<b>N / A</b>	<b>N / A</b>	<b>X</b>	<b>MT-RJ</b>	<b>62.5 μ MM</b>

**X** = Available for ordering  
**μ** = micron  
**SM** = Single mode fiber  
**MM** = Multimode fiber



## Server Time Protocol (STP)

- **Feature of z9 EC, z9 BC, z990, z890**
- **Maintain time synchronization with each other**
  - ▶ **Multiple z9 EC, z9 BC, z990, and z890 servers**
  - ▶ **Does not require Sysplex Timer if all servers STP-capable**
- **ISC-3 links (Peer mode), ICB-3 and ICB-4 links**
  - ▶ **Timing information transmitted**
- **Allows a Parallel Sysplex cluster to span up to 100 km (62 miles)**
- **Can coexist with an External Time Reference (ETR) network**
  - ▶ **Sysplex Timer-based network**
- **Prerequisite**
  - ▶ **z/OS V1.7**
  - ▶ **z9 BC - HMC Licensed Internal Code**

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





# Cryptography

---








# Cryptographic support on z9 BC

Integrated Cryptographic Service Facility (ICSF)



## ■ CP Assist for Cryptographic Function (CPACF)

- Available on every Processor Unit (PU)
  - Defined as a Central Processor (CP) or Integrated Facility for Linux (IFL)
- No-charge enablement feature (SHA-1, SHA-256 shipped enabled)
- High performance clear key symmetric encryption
  - Data Encryption Standard (DES)
  - Triple Data Encryption Standard (TDES)
  - Secure Hash Algorithm (SHA-1, [SHA-256](#))
- [Enhancements on z9 BC](#)
  - [Advanced Encryption Standard \(AES\) for 128-bit keys](#)
  - [Pseudo Random Number Generation \(PRNG\)](#)
  - [SHA-256](#)
- ICSF, component of z/OS
  - Uses available crypto functions
  - Balances the workload





# Crypto Express2 configurable on z9 BC

## □ Crypto Express2

❖ Two configuration modes for the PCI-X adapters

- Secure Coprocessor (default)

- ✓ Certified for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification (Certification # 661)

- ✓ Secure key transactions

- ✓ User Defined Extensions (UDX)

- Accelerator

- For SSL acceleration

- Clear key RSA operations

❖ Three configuration options

## □ Trusted Key Entry (TKE) 5.0 workstation

- ❖ For Crypto Express2 secure coprocessor




- ❖ New Graphical User Interface (GUI)

- ❖ Smart Card Reader

1

2

3





## System z cryptographic functions

Feature	CCF	PCICC	PCICA	CPACF	PCIXCC	Crypto Express2	CPACF	Crypto Express2
System	z800	z800	z800/z890	z890	z890	z890	z9 BC	z9 BC
Certification	140-1 L4	140-1 L4	None	None	140-2 L4	140-2 L4	None	140-2 L4
Secure key	Yes	Yes	No	No	Yes	Yes	No	Yes
Clear key	RSA only	RSA only	RSA only	Yes	RSA only	RSA only	Yes	RSA only
SSL acceleration	Yes	Yes	Yes	No	Yes	Yes	No	Yes
Configurable	No	No	No	No	No	No	No	Yes
DES/TDES	Yes	Yes	No	Yes	Yes	Yes	Yes	Yes
AES-128	No	No	No	No	No	No	Yes	No
SHA-1	Yes	No	No	Yes	No	No	Yes	No
SHA-256	No	No	No	No	No	No	Yes	No
PRNG	Yes	Yes	No	No	Yes	Yes	Yes	Yes
RSA-1024	Yes	Yes	Yes	No	Yes	Yes	No	Yes
RSA-2048	No	Yes	No	Yes	Yes	Yes	Yes	Yes
DUKPT	No	DES	No	No	DES/TDES	DES/TDES	No	DES/TDES
EMV-2000	No	No	No	No	Yes	Yes	No	Yes
19-digit PAN	No	No	No	No	Yes	Yes	No	Yes



## Cryptographic features over time

Feature	Feature Name	G5/G6	z800, 03/02	z890, 05/04	z9 BC
0860	PCICC	06/99	N/A	N/A	N/A
0861	PCICC replaces 0860	N/A	X	N/A	N/A
0862	PCICA	N/A	X	X	N/A
0868	PCIXCC replaces 0861	N/A	N/A	X	N/A
0863	Crypto Express2	N/A	N/A	01/05	X

X = Available for ordering



# Cryptographic features over time

Generation	1 <sup>st</sup>	2 <sup>nd</sup>	2 <sup>nd</sup>	3 <sup>rd</sup>	4 <sup>th</sup>	4 <sup>th</sup>
<b>Servers</b>	G5/G6 06/99	z800, 03/02	z890	z890	z890	z9 BC
<b>Name</b>	PCICC	PCICC PCICA	PCICA	PCIXCC	Crypto Express2	Crypto Express2
<b>Features numbers</b>	0860	PCICC, 0861 PCICA, 0862	0862	0868	0863	0863
<b>Number of PCI adapters</b>	1	2	2	1	2	2 *
<b>Maximum SSLs per second per Coprocessor</b>	125 SSLs	250 SSLs	N/A	1000 SSLs	2000 SSLs	2000 SSLs
<b>Maximum SSLs per second per Accelerator</b>	N/A	2000 SSLs	2000 SSLs	N/A	N/A	6000 SSLs
<b>Maximum Features / PCIs</b>	8 / 8	PCICA – 6 / 12 PCICC – 8 / 16	2 / 4 – z890	4 / 4	8 / 16	8 / 16

PCICA = PCI Cryptographic Accelerator

PCICC = PCI Cryptographic Coprocessor

\* On z9 BC each PCI-X adapter is configurable as either a coprocessor or an accelerator

The Secure Sockets Layer (SSL) measurements are examples of the maximum handshakes per second per feature achieved in a laboratory environment with no other processing occurring and do not represent actual field measurements.

The feature may have 1 or 2 PCI adapters. Details are available upon request.




## Glossary of Cryptographic terms

Acronym	Full name
AES	Advanced Encryption Standard
CCF	CMOS Cryptographic Coprocessor Facility
CP	Central Processor - a Processor Unit defined as a general purpose processor
CPACF	CP Assist for Cryptographic Function
DES	Data Encryption Standard
DUKPT	Derived Unique Key Per Transaction (algorithm) – key word management method ANSI X9.24 standard
EMV 2000	Europay MasterCard and Visa – standard; financial applications; heterogeneous hardware and software
FIPS	Federal Information Processing Standard
PCI	Peripheral Component Interconnect; An internal local bus for communication (industry standard)
PCICA	PCI Cryptographic Accelerator
PCICC	PCI Cryptographic Coprocessor
PCIXCC	PCIX Cryptographic Coprocessor
PKD	Public Key Decrypt – a service supporting Zero-Pad option for clear RSA private keys
PKE	Public Key Encrypt – a service supporting Mod_Raised_to Power (MRP) function used to offload compute-intensive portion of Diffie-Hellman protocol
PRNG	Pseudo Random Number Generation
SHA	Secure Hash Algorithm
SSL	Secure Sockets Layer
TDES	Triple Data Encryption Standard
TKE	Trusted Key Entry
UDX	User-Defined Extensions




# Storage Area Network connectivity






# 16-port ESCON



MT-RJ connector

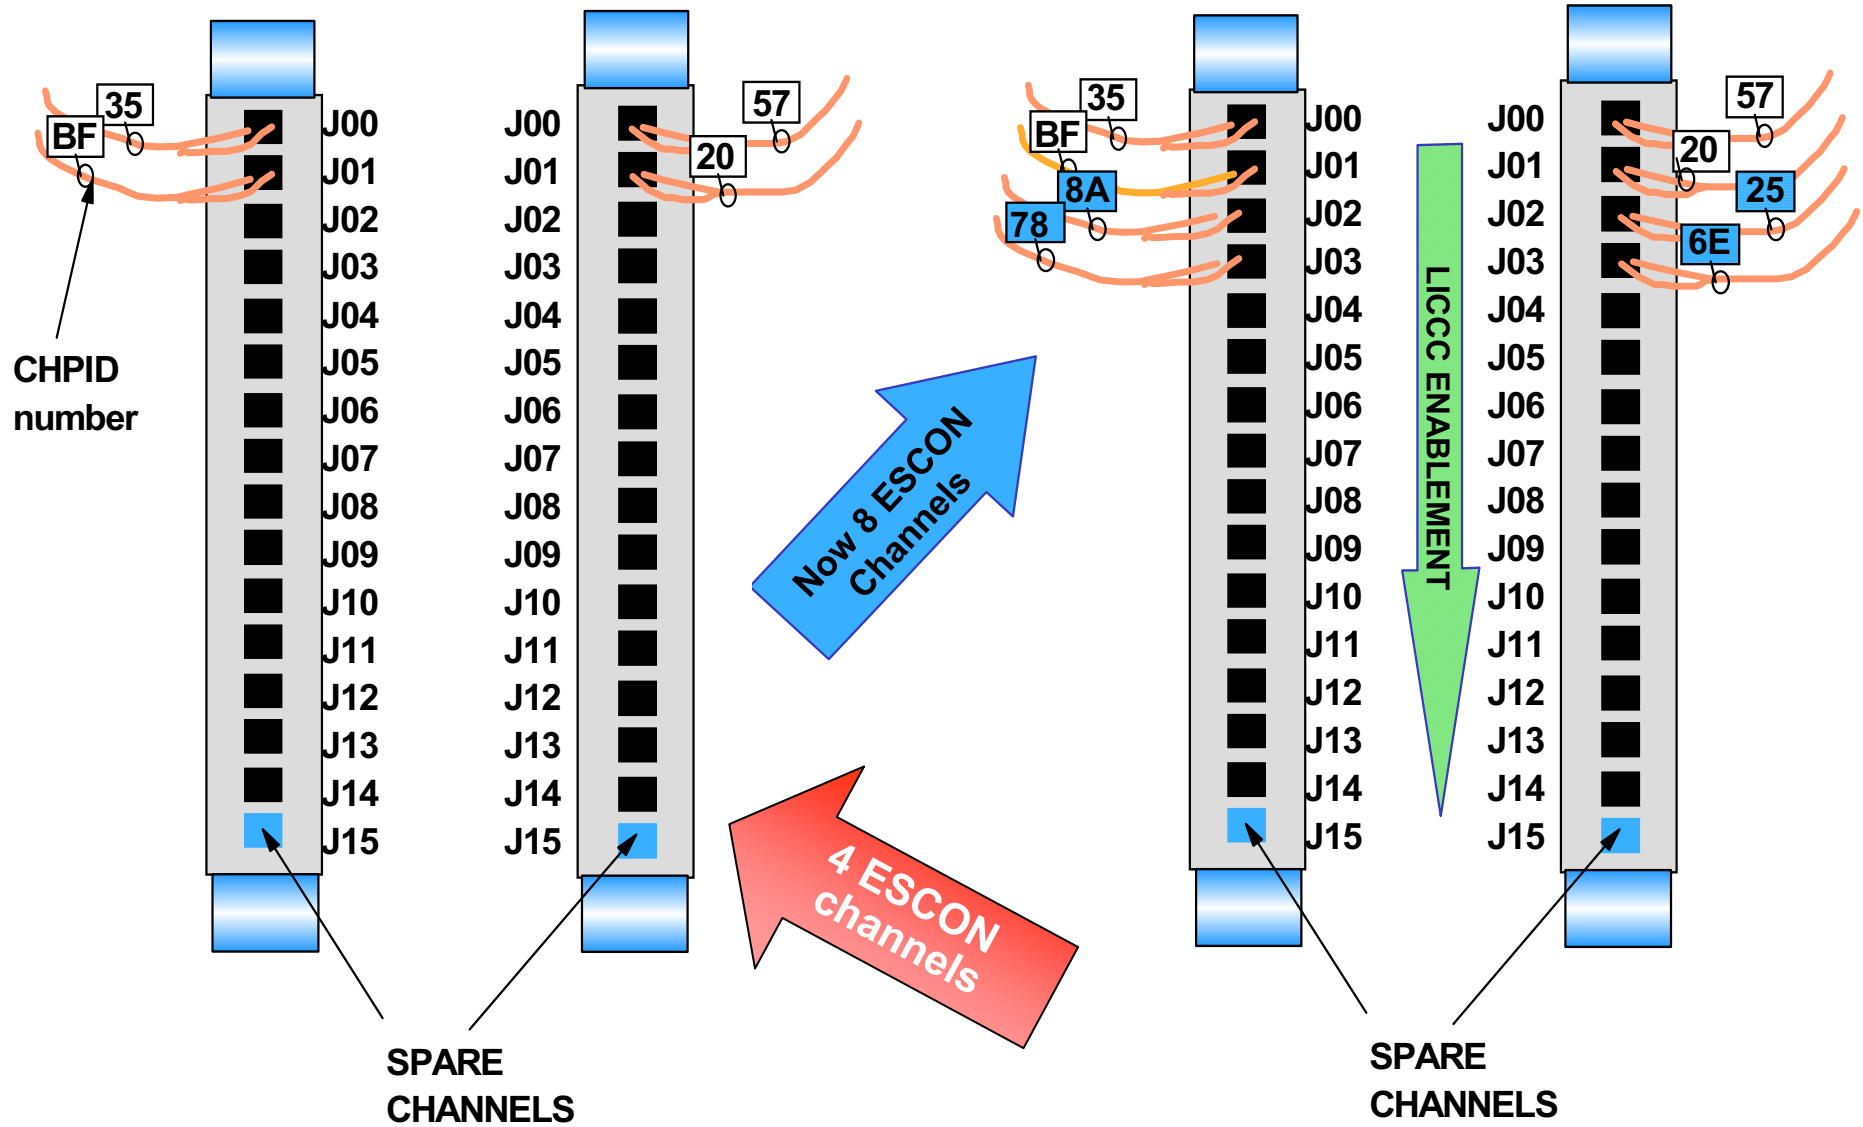
Machine Type	Maximum channels	Maximum features
z9 BC S07/R07	420/240	28/16
z890	420/240	28/16
z800	240	16








# ESCON channel plugging





# FICON Express4 on z9 BC



▪ **Personalize as:**

- ▶ FC (Fibre Channel)
  - Native FICON
  - Channel-To-Channel (CTC)
- ▶ FCP (Fibre Channel Protocol)
  - Support of SCSI devices
    - z/VM, z/VSE, Linux on System z9
- ▶ FCV not offered on
  - ▶ FICON Express4
  - ▶ FICON Express2

Machine Type	Maximum channels	Maximum features
z9 BC S07/R07	112/64	28/16
z890	112/64	28/16
z800	32	16




## Fibre Channel distances

9u single mode fiber 10 km - 10,000 meters, 6.2 miles

✕ 50u 2000 MHz   
 ▲ 50u 500 MHz   
 ✕ 62.5u 200 MHz   
 ● 62.5u 160 MHz

## Multimode fiber distances




# FICON Express4 on z9 BC

**Small Form Factor Pluggable (SFP) optics**  
**Concurrent repair/replace action for each SFP**

■ **Ordering**

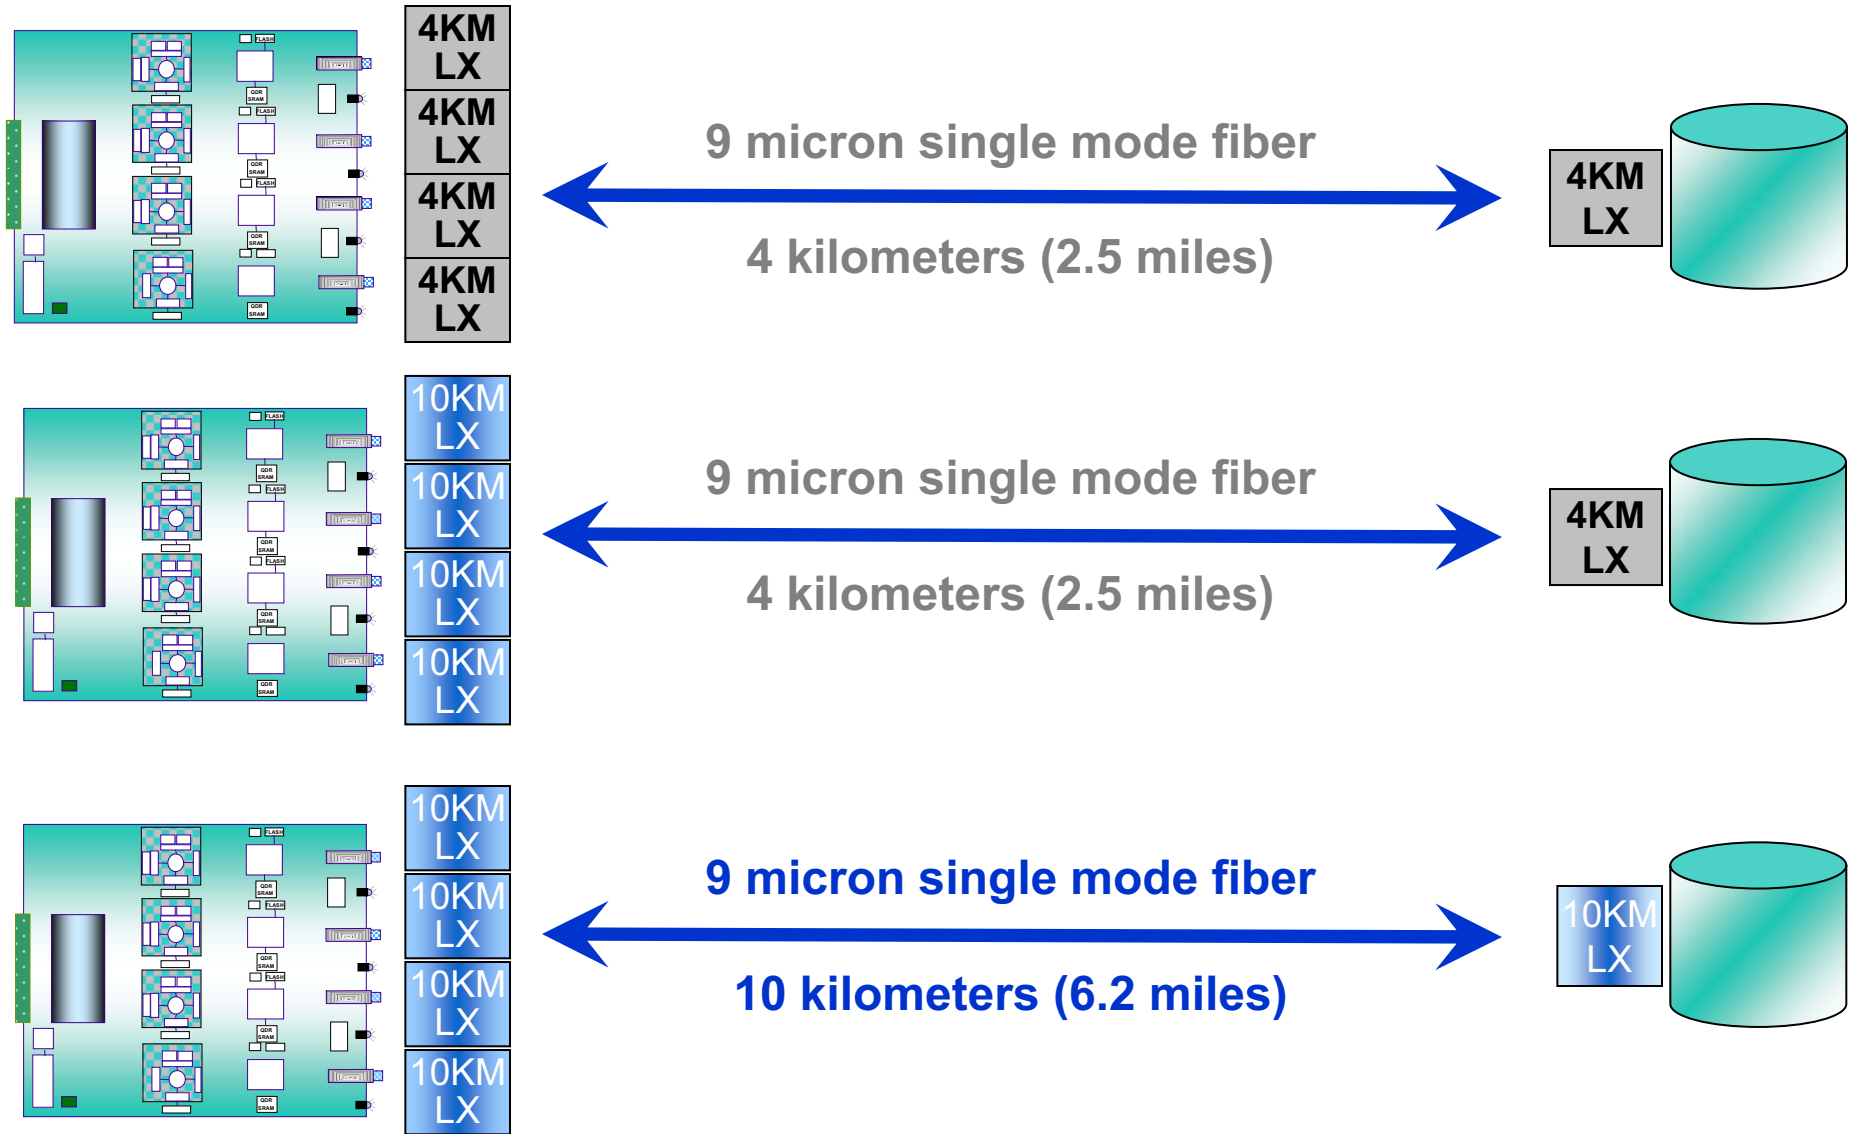
- ▶ Four-channel increments
- ▶ All channels - same type
  - LX (single mode fiber) or
  - SX (multimode fiber)

<b>FC 3321</b>	<b>FICON Express4 10KM LX</b>
<b>FC 3324</b>	<b>FICON Express4 4KM LX</b>
<b>FC 3322</b>	<b>FICON Express4 SX</b>
<b>FC 3323</b>	<b>FICON Express4-2C 4KM LX</b> (exclusive to z9 BC)





# FICON Express4 LX optics - unrepeated Distances





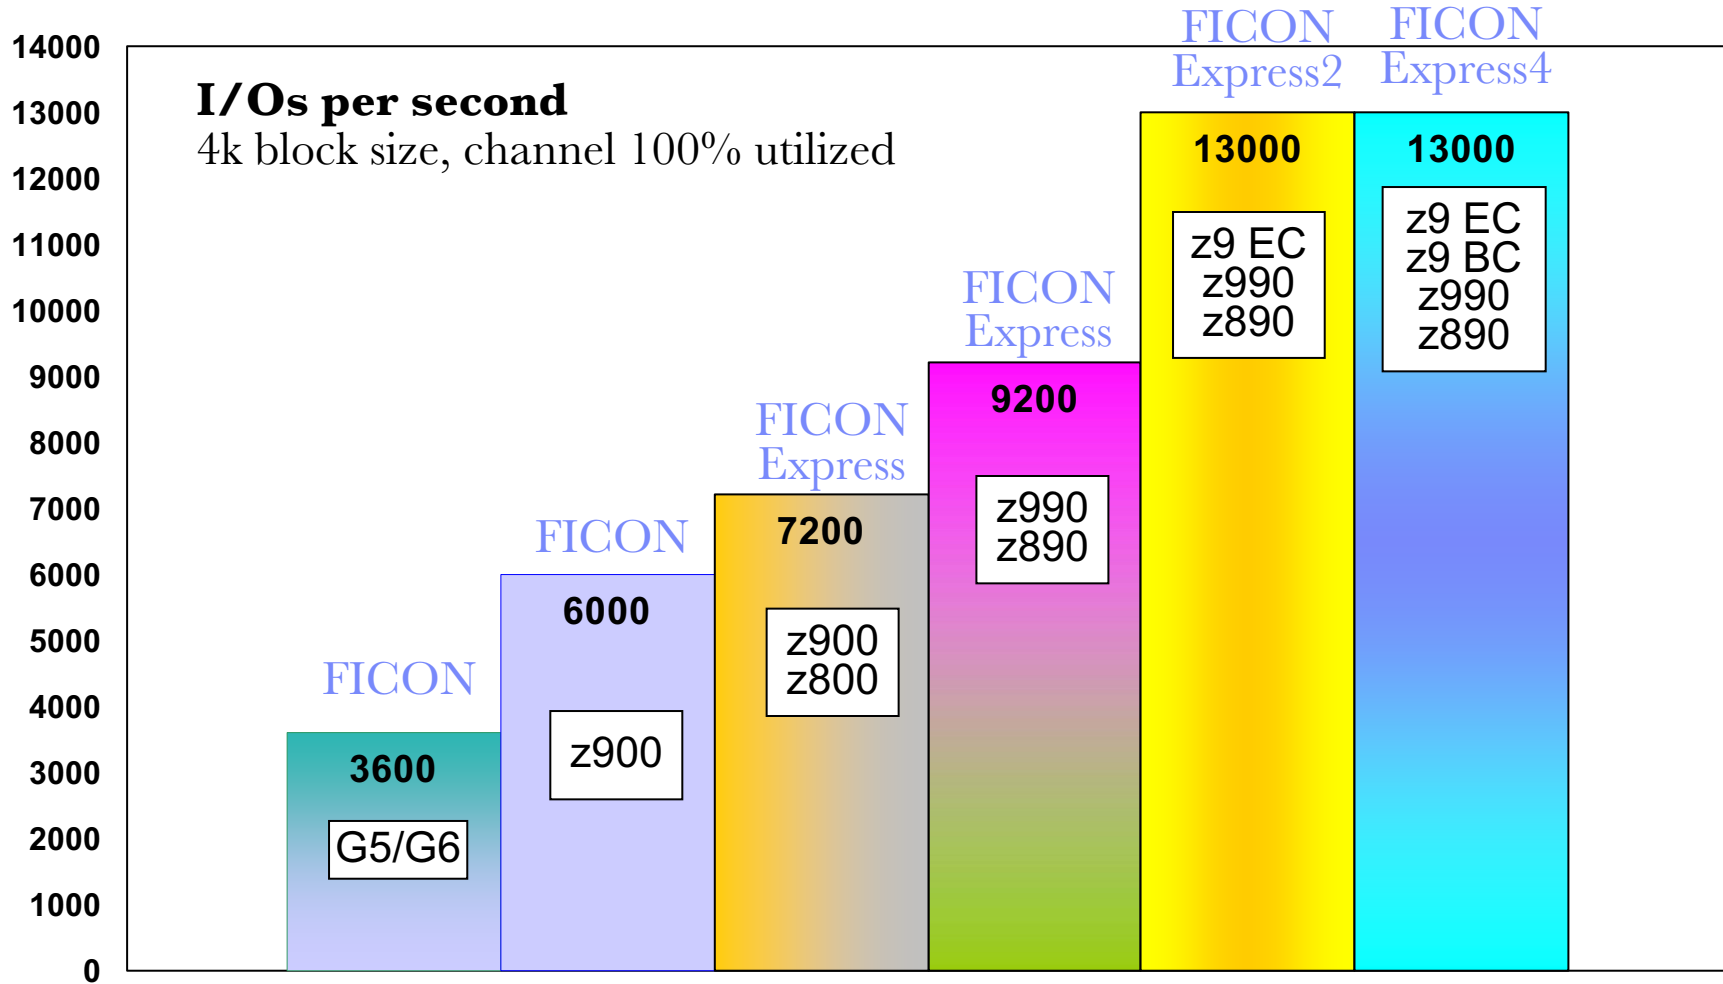
# System z9 BC FICON/FCP 4 Gbps tested products

Product	Type	MT	Models	Mode	Gbps	Intermix ISL	Transceiver	Connector
IBM TotalStorage SAN256B Director	b-type	2109	M48	FICON FCP	1, 2, 4	Yes Yes	SX, LX	LC Duplex
IBM TotalStorage SAN32B-2	b-type	2005	B32	FICON FCP	1, 2, 4	Yes Yes	SX, LX	LC Duplex
Cisco MDS 9000 – 9506, 9509	c-type	2062	D04 , D07	FICON FCP	1, 2, 4	Yes, VSAN Yes	SX, LX CWDM *	LC Duplex
Cisco MDS 9000 – 9216a, 9216i	c-type	2062	D1A, D1H	FICON FCP	1, 2, 4	Yes, VSAN Yes	SX, LX CWDM *	LC Duplex
Cisco MDS 9513 Multilayer Director	c-type	2062	E11	FICON FCP	1, 2, 4	Yes, VSAN Yes		
IBM TotalStorage SAN32M-2	m-type	2026	432	FICON FCP	1, 2, 4	Yes Yes	SX, LX	LC Duplex
IBM TotalStorage SAN140M	m-type	2027	140	FICON FCP	1, 2, 4	Yes Yes	SX, LX	LC Duplex

April / May 2006 announce



# FICON performance – Start I/Os



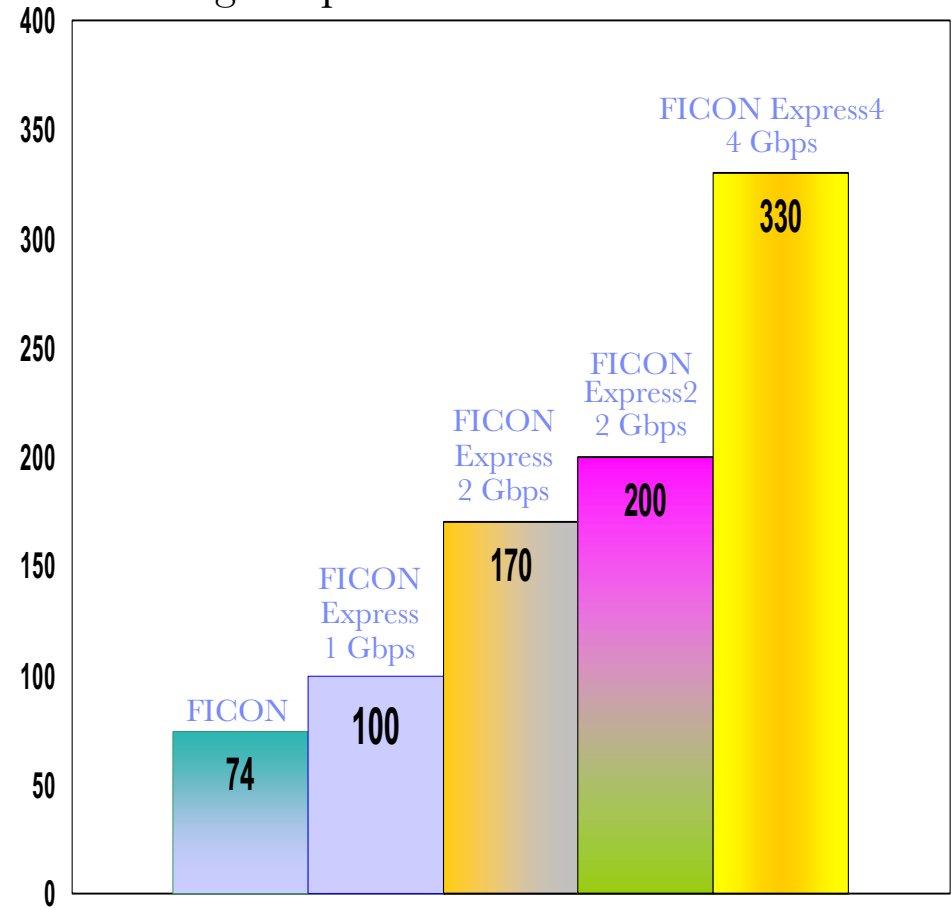

\*This performance data was measured in a controlled environment running an I/O driver program under z/OS 1.7. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.



# FICON performance

**Full duplex data transfers**  
**MegaBytes per second (full duplex)**  
 Large sequential read/write mix

**Half duplex data transfers**  
**MegaBytes per second (half duplex)**  
 Large sequential all reads or all writes



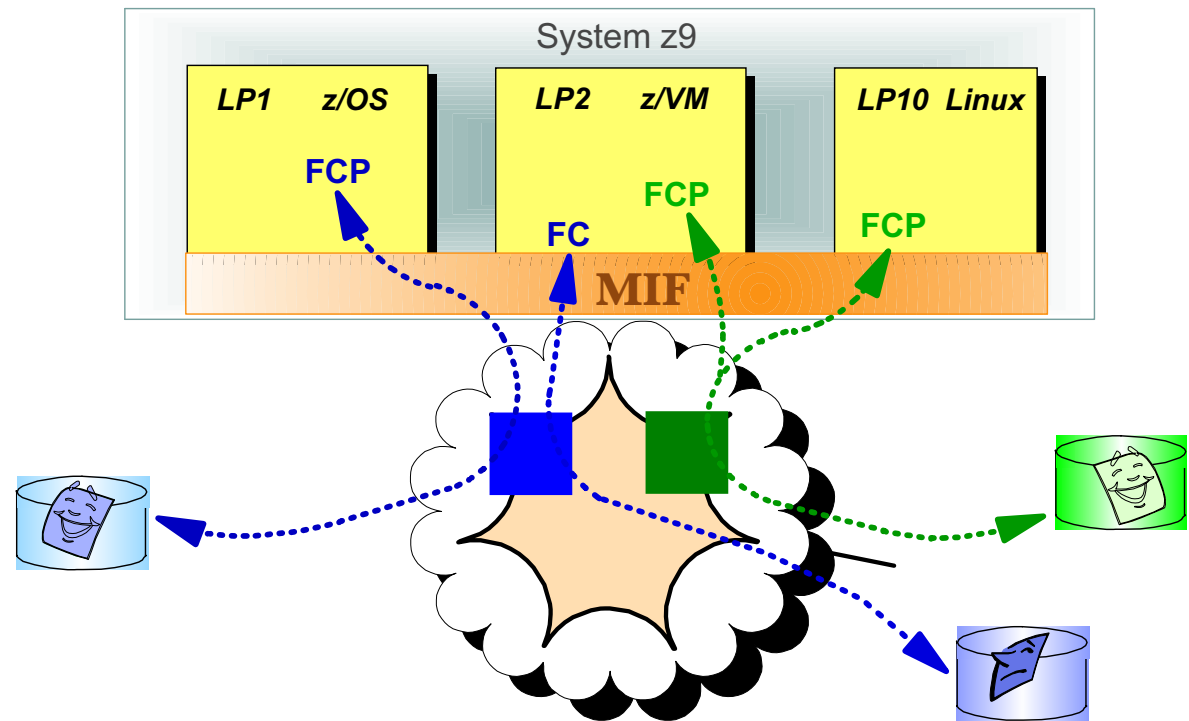
\*This performance data was measured in a controlled environment running an I/O driver program under z/OS 1.7. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.





## N\_Port ID Virtualization (NPIV) for all FICON features on z9 BC

- ❑ Multiple N\_Port\_Names (WWPNs) assigned to N\_Port of FCP channel
  - ❖ Each OS uses its own unique N\_Port\_Name to log into fabric
  - ❖ Unique FCP identifier (N\_Port\_ID) assigned to each OS
  - ❖ OS can access fabric with one or multiple N\_Port\_names
- ❑ Extension to the Fibre Channel standard
- ❑ Entry point into fabric must also support NPIV
- ❑ All FICON features supported on z9 BC (CHPID type FCP)






## FICON availability enhancements


### All FICON features supported on z9 BC

- **Request Node Identification Data (RNID)**
  - ▶ RNID data is stored in HSA for each control unit attached to FICON channels
  - ▶ Can now request RNID data for native FICON channels (CHPID type FC)
    - For each device or control unit attached to channel
    - Formatted and displayed on the SE "Analyze Control Unit Header" IOPD panel.
    - Provided to z/OS so it can display it on the "D M=DEV" (Display Device Matrix) command to help debug configuration/cabling problems.
    - z/OS V1.4 and above with PTFs
- **FICON link incident reporting**
  - ▶ Operating system image can register for link incident reports
    - Without operator intervention
  - ▶ Displayed on operator console
  - ▶ Saved in system log and LOGREC
  - ▶ z/OS V1.7





# FICON features over time

Feature	Feature name	Ports	z800 03/02	z890 05/04	z9 BC 05/06	CHPIDs	Connector
2315	FICON LX	2	N/A	N / A	N / A	FCV, FC, FCP	SC Duplex
2318	FICON SX	2	N/A	N / A	N / A	FC, FCP	SC Duplex
2319	FICON Express LX	2	X	X	C	FCV, FC, FCP	LC Duplex
2320	FICON Express SX	2	X	X	C	FC, FCP	LC Duplex
3319	FICON Express2 LX	4	N / A	01/05	X	FC, FCP	LC Duplex
3320	FICON Express2 SX	4	N / A	01/05	X	FC, FCP	LC Duplex
3321	FICON Express4 10 KM LX	4	N/A	N/A	05/06	FC, FCP	LC Duplex
3322	FICON Express4 SX	4	N/A	N/A	05/06	FC, FCP	LC Duplex
 3323	FICON Express4-2C 4KM LX	2	N/A	N/A	05/06	FC, FCP	LC Duplex
3324	FICON Express4 4KM LX	4	N/A	N/A	05/06	FC, FCP	LC Duplex

LX = Long wavelength transceiver, used with LX senders and receivers and single mode fiber  
 SX = Short wavelength transceiver, used with SX senders and receivers and multimode fiber  
 C = Carry forward on an upgrade from z800 or z890 to z9 BC  
 X = Available for ordering


# Local Area Network connectivity

## Ethernet server

**1000BASE-T**

**Gigabit**

**10 Gigabit**






# OSA-Express2 1000BASE-T Ethernet

- **Supports auto-negotiation to 10, 100, 1000 Mbps over Category 5 copper**
- **Capable of achieving line speed**
  - ▶ Actual throughput is dependent upon environment
- **Supports:**
  - ▶ Layer 2 for protocol-independent packet forwarding
  - ▶ Large send for offloading TCP segmentation processing
  - ▶ 640 TCP/IP stacks for improved virtualization
  - ▶ Concurrent LIC update to minimize network traffic disruption
- **Modes of operation**

1000BASE-T Ethernet #3366



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 for NCP	OSN	NCPs running under IBM Communication Controller for Linux

# OSA-Express2 10 GbE and 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 type OSD (QDIO)
- CHPID type OSN (OSA for NCP) – exclusive to z9 EC, z9 BC
- 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 10 GbE and GbE support


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

10 Gigabit Ethernet  
Feature 3368



SC Duplex SM

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



LC Duplex SM




LC Duplex MM



# Ethernet on the LAN

- ✓ OSA-Express2 GbE is designed to achieve line speed - 1 Gbps in each direction
- ✓ Greatest throughput achieved with jumbo frames





## Personalities of OSA

CHPID type	OSA-Express2 features	Traffic
OSC	1000BASE-T Ethernet	OSA-Integrated Console Controller (OSA-ICC) TN3270E, non-SNA DFT IPL CECs and LPARs, OS console operations
OSD	1000BASE-T Ethernet Gigabit Ethernet 10 Gigabit Ethernet	Queued Direct Input/Output (QDIO) TCP/IP traffic when Layer 3 Protocol-independent when Layer 2
OSE	1000BASE-T Ethernet	Non-QDIO SNA, APPN, HPR and/or TCP/IP
OSN	1000BASE-T Ethernet Gigabit Ethernet Exclusive to z9 EC, z9 BC	OSA-Express2 OSN (OSA for NCP) Supporting NCPs running under IBM Communication Controller for Linux on System z9 and zSeries V1.2






## Layer 3 versus Layer 2

Layer 3	Layer 2
Guests share the same MAC address	Guests have own unique MAC address
Uses MAC address of shared OSA	z/VM virtual switch assigns MAC addresses
Guests are TCP/IP only	Guests can use any network layer protocol
No emulation of physical LAN segment for guests	Guests appear to have own physical LAN segment
Guests are not known to the hosts on the physical side of the LAN segment. ARP always resolves to a single OSA feature	ARP allows guests to become known to the host residing on the physical side of the LAN segment
Data sent by guest encapsulated within IP packet	Data sent by guest encapsulated in Ethernet frame




## OSA-Express2 OSN (OSA for NCP) and IBM Communication Controller for Linux (CCL)

CCL V1.1	CCL V1.2, available October 2005 (5724-J38)
OSA – Fast Ethernet, 1000BASE-T Ethernet, Gigabit Ethernet zSeries	OSA-Express2 GbE or 1000BASE-T Ethernet System z9 exclusive
Non-QDIO (CHPID type OSE) QDIO (CHPID type OSD) Layer 2	OSA-Express2 OSN – OSA for NCP (CHPID type OSN)
External communication	Internal communication
Traffic flows on the LAN	Traffic flows LPAR-to LPAR



No longer limited to  
“copper OSA”



**OSA – available on new build and what you can carry forward®**

Feature	Feature Name	Ports	z800	z890	z9 BC	CHPIDs	Connectors
5201	OSA-2 Token Ring	2	X	N / A	N / A	OSA	RJ-45
5202	OSA-2 FDDI	1	X	N / A	N / A	OSA	SC Duplex
2362	OSA-E 155 ATM SM	2	X	RPQ	N / A	OSD, OSE	SC Duplex
2363	OSA-E 155 ATM MM	2	X	RPQ	N / A	OSD, OSE	SC Duplex
2364	OSA-E GbE LX	2	X	C	C	OSD, L2/L3 **	SC Duplex
2365	OSA-E GbE SX	2	X	C	C	OSD, L2/L3 **	SC Duplex
2366	OSA-E Fast Ethernet	2	X	C	C	OSD L2/L3 **, OSE	RJ-45
2367	OSA-E Token Ring	2	X	X	N / A	OSD, OSE	RJ-45
1364	OSA-E GbE LX	2	09/04	X	C	OSD, L2/L3 **	LC Duplex
1365	OSA-E GbE SX	2	09/04	X	C	OSD, L2/L3 **	LC Duplex
1366	OSA-E 1000BASE-T Ethernet	2	N / A	X	C	OSC, OSD L2,L3, OSE	RJ-45
3364	OSA-E2 GbE LX	2	N / A	01/05	X	OSD L2/L3, OSN *	LC Duplex
3365	OSA-E2 GbE SX	2	N / A	01/05	X	OSD L2/L3, OSN *	LC Duplex
3366	OSA-E2 1000BASE-T Ethernet	2	N / A	05/06	X	OSC, OSD L2/L3, OSE, OSN *	RJ-45
3368	OSA-E2 10 GbE LR	1	N / A	01/05	X	OSD L2/L3 **	SC Duplex

LX = Long wavelength transceiver, SX = Short wavelength transceiver, LR - Long Reach transceiver

X = Available for ordering C = Carry forward on an upgrade from z800 or z890

\* OSN is exclusive to z9 EC and z9 BC \*\* L2/L3 = Layer 2/Layer 3 which is applicable to z9 EC, z9 BC, z990, z890



# z9 BC S07 I/O connectivity summary

Up to a maximum of 420 CHPIDs, one I/O cage, 28 I/O slots

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

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




# z9 BC R07 I/O connectivity summary

Up to a maximum of 240 CHPIDs, one I/O cage, 16 I/O slots

Features	Minimum # of features	Maximum # of features	Maximum connections	Increments per feature	Purchase increments
16-port ESCON	0 (1)	16	240 channels	16 channels 1 reserved as spare	4 channels
FICON Express4	0 (1)	16	64 / 32 channels	4 / 2 channels **	4 / 2 channels **
FICON Express2 *	0 (1)	16	64 channels	4 channels	4 channels
FICON Express *	0 (1)	16	32 channels	2 channels	2 channels
STI-3 (2)	0	8	N/A	2 outputs	N/A
ICB-3 link	0 (1)	N/A	16 links (3)	N/A	1 link
ICB-4 (4)	0 (1)	N/A	16 links (3) (5)	N/A	1 link
ISC-3	0 (1)	12	48 links (3)	4 links	1 link
OSA-Express2	0	15 (6)	30 ports (6)	2 or 1 (10 GbE has 1)	2 ports/1 port
OSA-Express *	0	15 (6)	30 ports (6)	2	2 ports
Crypto Express2	0	4	8 PCI-X adapters	2 PCI-X adapters	2 PCI-X adapters (7)

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. Model Capacity ID A01 ICB-4 links is maximum of 8
  6. Model Capacity ID A01 OSA-Express2, OSA-Express ix maximum of 12 features (24 ports)
  7. Initial order of Crypto Express2 is 4 PCI-X adapters (two features). Each PCI-X adapter can be configured as either a coprocessor or an accelerator.
- Available only when carried forward on an upgrade.
- \*\* The FICON Express4-2C 4KM LX feature has two channels per feature



# Ride the System z9 Business Class Bus

*Coupling*

*Cryptography*

*SAN*

*LAN*

**Thank you!**  
**Session G06**