



G06

OSA-Express2 – The latest offerings for the LAN

Connie K. Beuselinck, conniek@us.ibm.com



September 19 - 23, 2005

San Francisco, CA

IBM Corporation
Poughkeepsie, NY

Trademarks

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

390	FICON *	OS/390 *	Tivoli *
ACF/VTAM *	HiperSockets	Parallel Sysplex *	TotalStorage *
APPN *	HPR	PR/SM	VM/ESA *
CICS *	IBM*	RACF *	VSE/ESA
DB2 *	IBM logo*	Redbooks	VTAM *
e-business logo *	IMS	Resource Link	WebSphere *
ESCON *	Infoprint	RMF	z/Architecture
eServer	OS/2 *	RS/6000 *	z/OS *
GDPS *		S/390 *	zSeries *
Geographically Dispersed Parallel Sysplex		S/390 Parallel Enterprise Server	z/VM *
		Sysplex Timer *	

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Linux is a registered trademark of Linus Torvalds

Penguin (Tux) compliments of Larry Ewing

Java and all Java-related trademarks and logos are 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.

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
1000BASE-T	Standard term for Ethernet when	capable of operating at 10/100/1000 Mbps
CHPID	Channel Path Identifier	
CRH	Channel Request Handler	Internal host bus for I/O communication
dB	decibel	Metric used to measure signal strength in fiber
GB	GigaByte	
GbE	Gigabit Ethernet	LAN protocol
km	kilometer	1 kilometer is 0.62 miles
LAN	Local Area Network	
LCS	LAN Channel Station	
LCSS	Logical Channel Subsystem	Architecture that allows more than one physical channel subsystem
LIC	Licensed Internal Code	
LPAR	Logical Partition	
LR	Long Reach	Transceiver used with single mode fiber optic cabling
LX	Long wavelength	Transceiver used with single mode fiber optic cabling
MCP	Mode Conditioning Patch	Cable used with sm fiber to accommodate reuse of mm
MBA	Memory Bus Adapter	Part of Central Electronic Complex
MBps	Megabits per second	
MCM	Multichip Module	Part of Central Electronic Complex
MM	Multimode	50 or 62.5 micron fiber optic cabling
MTU	Maximum Transmission Unit	Largest physical packet size in bytes that a network can transmit
OSA	Open Systems Adapter	zSeries family name for LAN adapters
PCI	Peripheral Component Interconnect	Local bus standard
PCI-X	Peripheral Component Interconnect Extended	Enhanced PCI bus - for increased performance over PCI
SM	Single mode	9 micron fiber optic cabling
SNMP	Simple Network Management Protocol	TCP/IP application layer protocol to manage network devices
STI	Self-Timed Interconnect	Internal host bus for I/O communication
SX	Short wavelength	Transceiver used with multimode fiber optic cabling



| IBM System z9™

OSA on z9-109, z990, z900



zSeries Expo

San Francisco

September 10 – 23, 2005

© 2005 IBM Corporation

Agenda

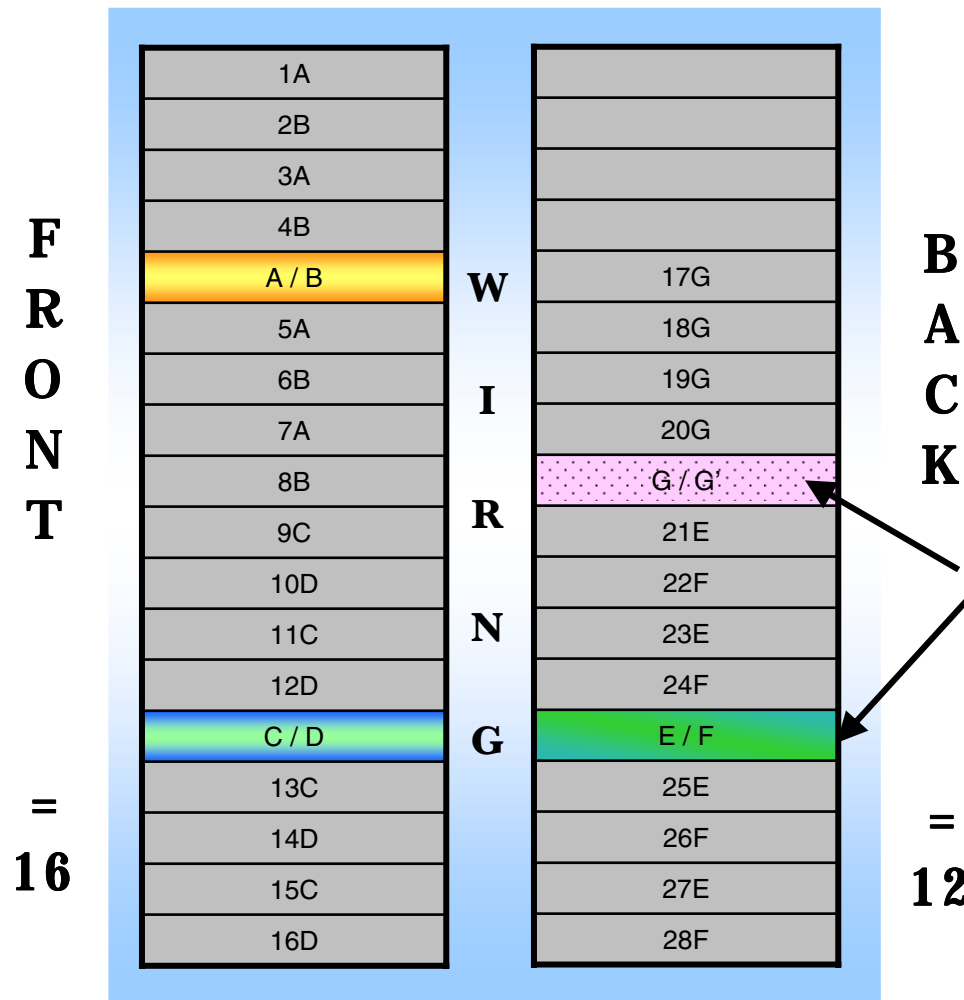


IBM System z9™

- I/O cage and plugging rules
- History and hardware
- Technology and performance
- Personalities of OSA - CHPIDs
- Newest functions
- Cabling, connectors
- Distance, Link loss budgets



Now 8 STIs per I/O cage



□ I/O cage

➤ Top View

➤ 8 STIs/STI-MPs

▪ G' is for redundancy

□ STI speeds for downstream channels/ports/links

➤ 500 MBps

OSA-Express

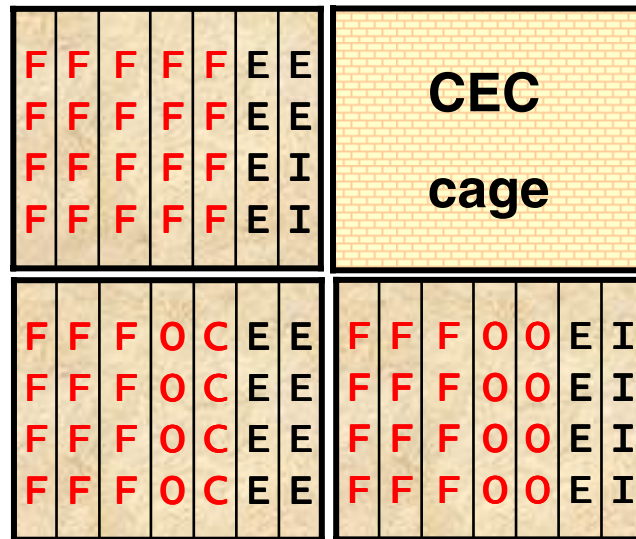
➤ 1 GBps

▪ OSA-Express2


Increased connectivity per I/O cage

- **Now can populate the 28 I/O slots in one I/O cage with any mix of:**
 - FICON Express2, FICON Express, up to 28 in one I/O cage, 84 per server
 - OSA-Express2, OSA-Express, up to 24 per I/O cage or server
 - Crypto Express2, up to 8 per I/O cage or server
- **Removed limitation of up to 20 FICON, OSAs, and Crypto in one I/O cage**
- **84 I/O slots in three I/O cages**

z900 and z990, 20 per I/O cage



z9-109, 28 per I/O cage 



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

Spanned channels, z9-109 and z990

Shared channels among LPARs across LCSSs

★ **Internal spanned channels** – **First available October 2003**

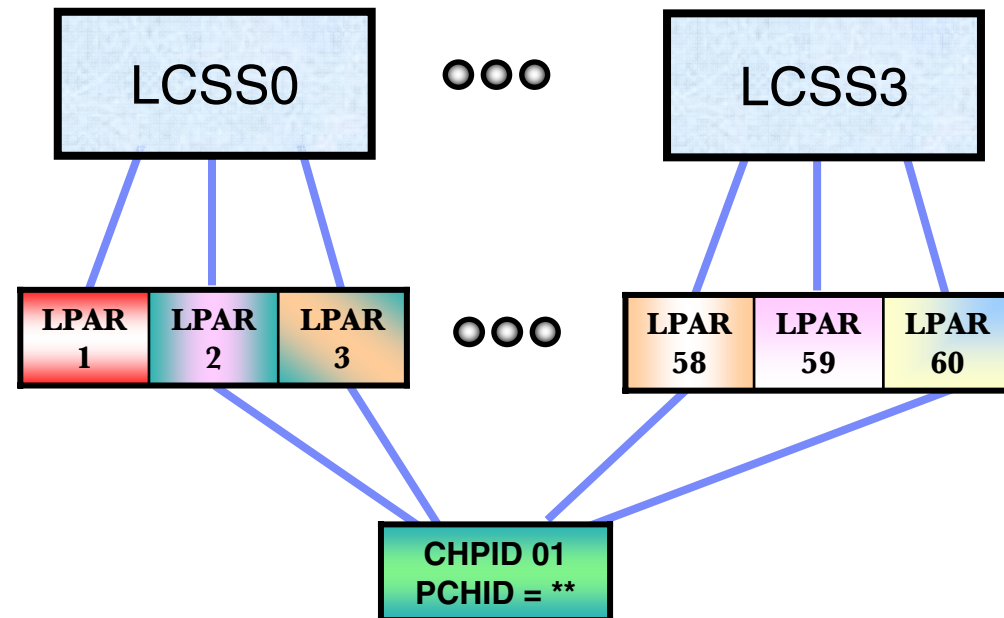
– HiperSockets and Internal Coupling links (ICs)

★ **External spanned channels** – **First available May 2004**

– OSA-Express, OSA-Express2

– FICON Express, FICON Express2

– ICB-3s, ICB-4s, ISC-3



** No PCHID for HiperSockets and Internal Coupling links. PCHID required for FICON, ICs, ICBs, ISC-3, OSA
Spanning reduces the number of channels that can be defined for all LCSSs on server
Worst case - 256 if all channels are spanned between all LCSSs

Enhanced performance assists for z/VM V5.2 guests

1. QDIO Enhanced Buffer-State Management (QEBSM)

- Two new machine instructions
 - Designed to help eliminate overhead of hypervisor interception

2. Host Page-Management Assist (HPMA)

- Interface to z/VM paging-storage management
 - Allow machine to assign, lock, unlock page frames without hypervisor assistance

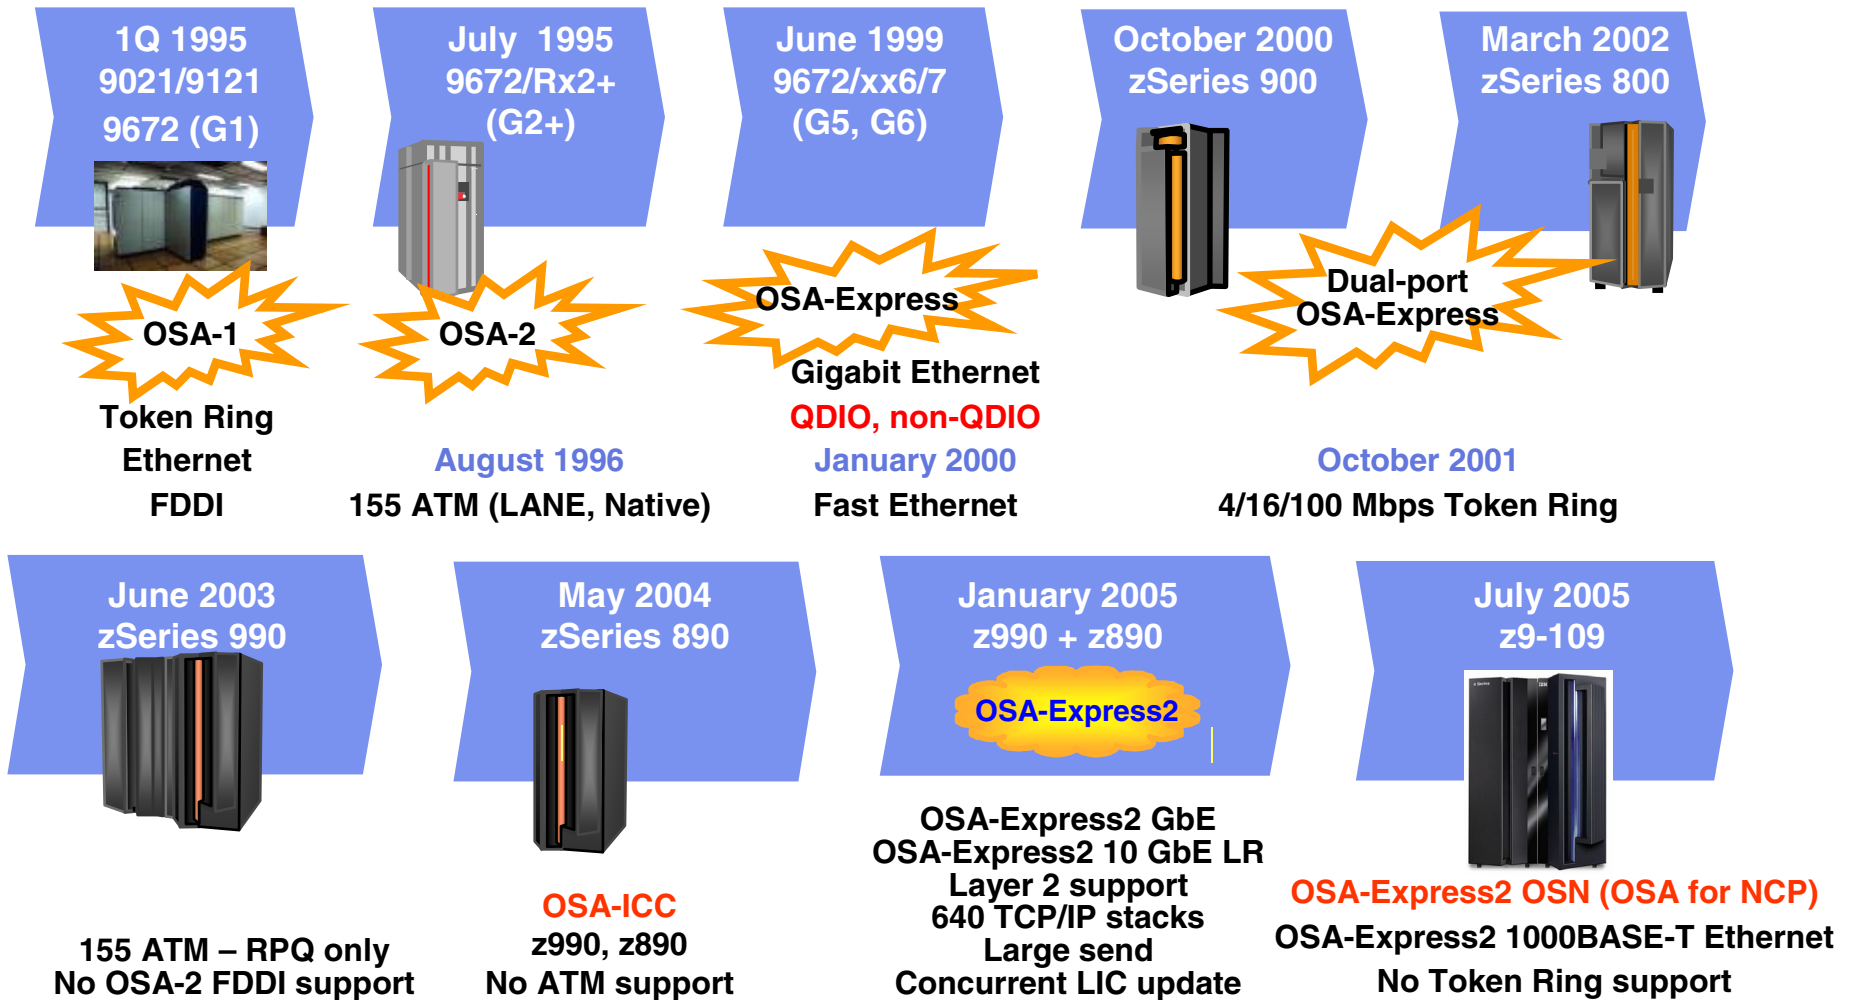
▪ **Applicable to the following on z9-109, z990**

- First-level guests of z/VM V5.2
- All OSA features (CHPID type OSD)
- All FICON features (CHPID type FCP)
- HiperSockets (CHPID type IQD)


▪ **Complements performance assists introduced in z/VM V4.4**

Guest operating systems can initiate QDIO operations directly to channel
No interception by z/VM

History of the Open Systems Adapter (OSA)



OSA features over time

Feature	Feature Name	Ports	z900	z990	z9-109	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	SC Duplex
2365	OSA-E GbE SX	2	X	C	C	OSD	SC Duplex
2366	OSA-E Fast Ethernet	2	X	C	C	OSD, 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	06/03	C	OSD	LC Duplex
1365	OSA-E GbE SX	2	09/04	06/03	C	OSD	LC Duplex
1366	OSA-E 1000BASE-T Ethernet	2	N / A	06/03	C	OSC, OSD, OSE	RJ-45
3364	OSA-E2 GbE LX	2	N / A	01/05	X	OSD, OSN *	LC Duplex
3365	OSA-E2 GbE SX	2	N / A	01/05	X	OSD, OSN *	LC Duplex
 3366	OSA-E2 1000BASE-T Ethernet	2	N / A	N / A	X	OSC, OSD, OSE, OSN *	RJ-45
3368	OSA-E2 10 GbE LR	1	N / A	01/05	X	OSD	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 z900 or z990

* OSN is exclusive to z9-109. Hardware availability is 09/16/05



IBM System z9™

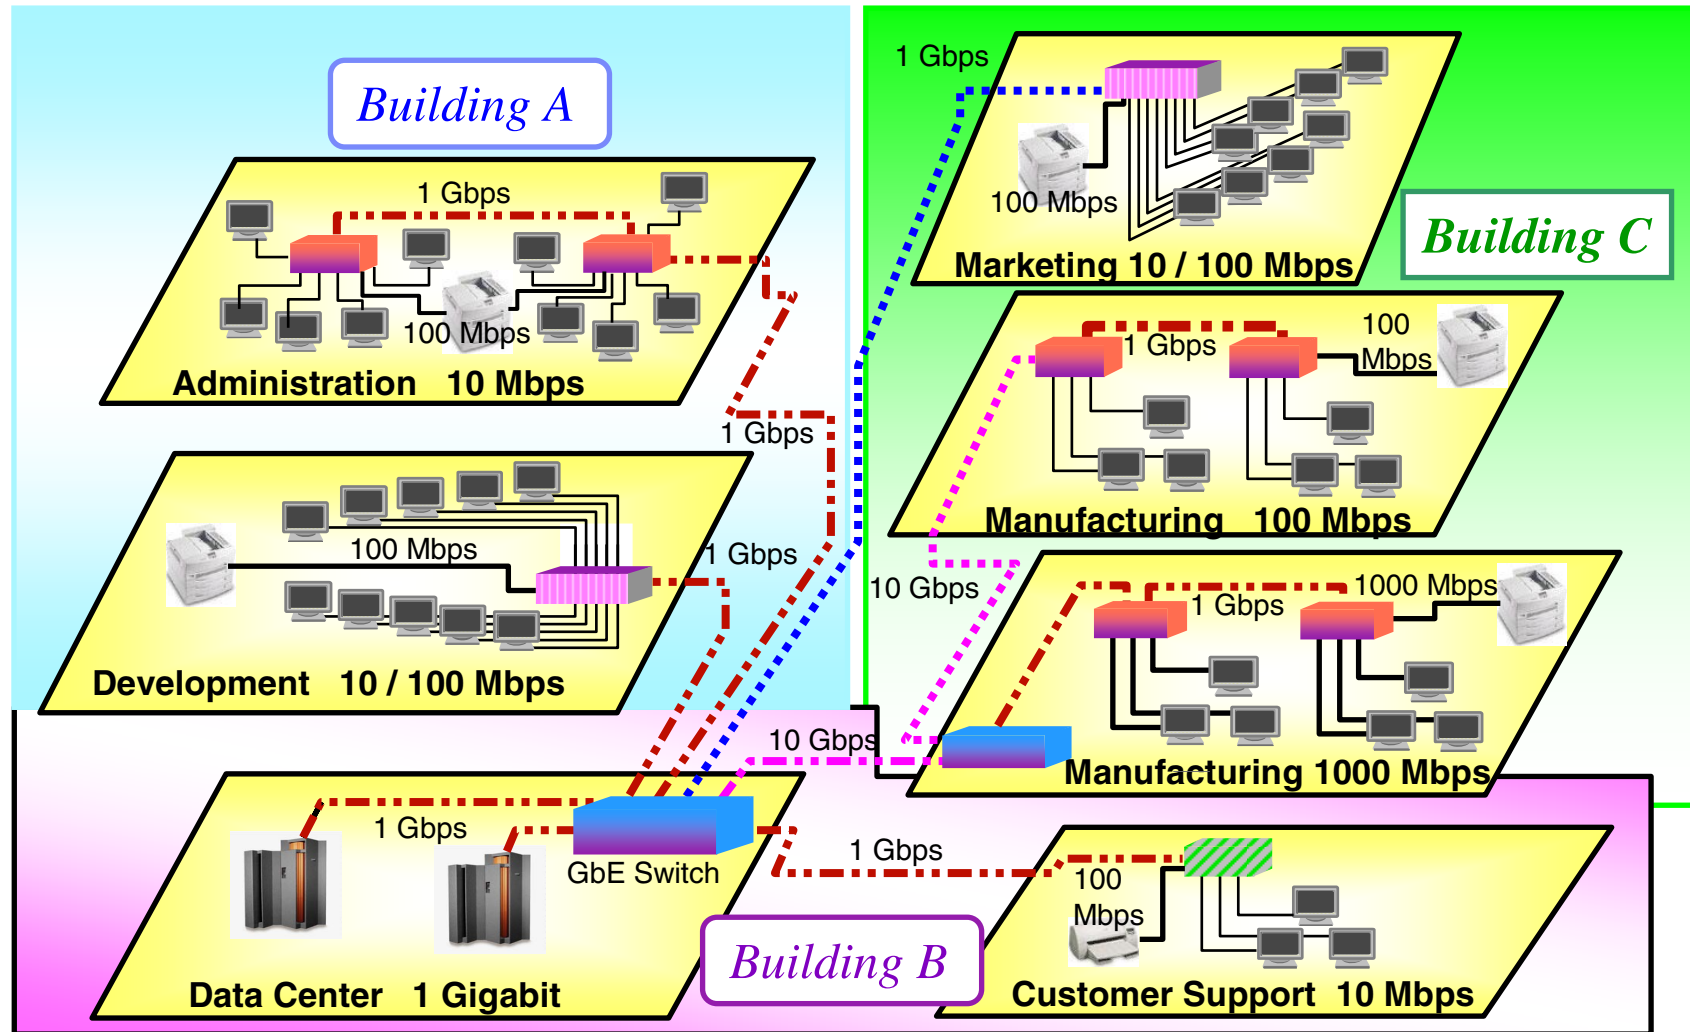
The features on
z9-109 ...
An Ethernet server



ON DEMAND BUSINESS™

Ethernet Coexistence

- ⋯ = SM fiber, 10 Gbps
- ⋯ = SM fiber
- ⋯ = MM fiber
- = 1000 or 100 Mbps, Category 5 copper
- = 10 Mbps, Category 3 copper




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

1000BASE-T Ethernet #3366



▪ Modes of operation



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 GbE, 10 GbE

10 Gigabit Ethernet LR (long reach)

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

10 Gigabit Ethernet Feature 3368




SC Duplex SM

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

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



LC Duplex SM



LC Duplex MM


OSA-Express2 GbE and 10 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



IBM System z9™

Technology and performance



ON DEMAND BUSINESS™

zSeries Expo

San Francisco

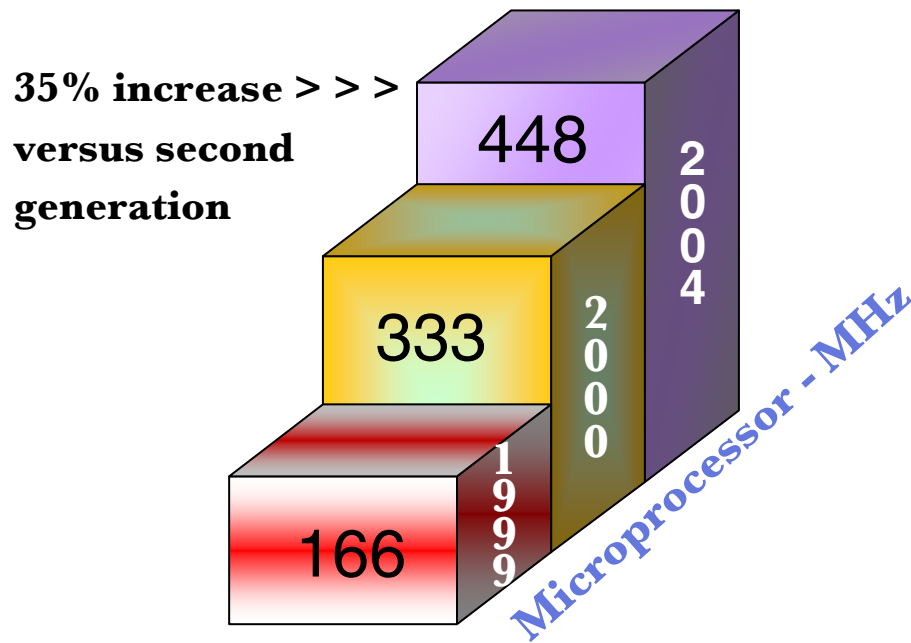
September 10 – 23, 2005

© 2005 IBM Corporation

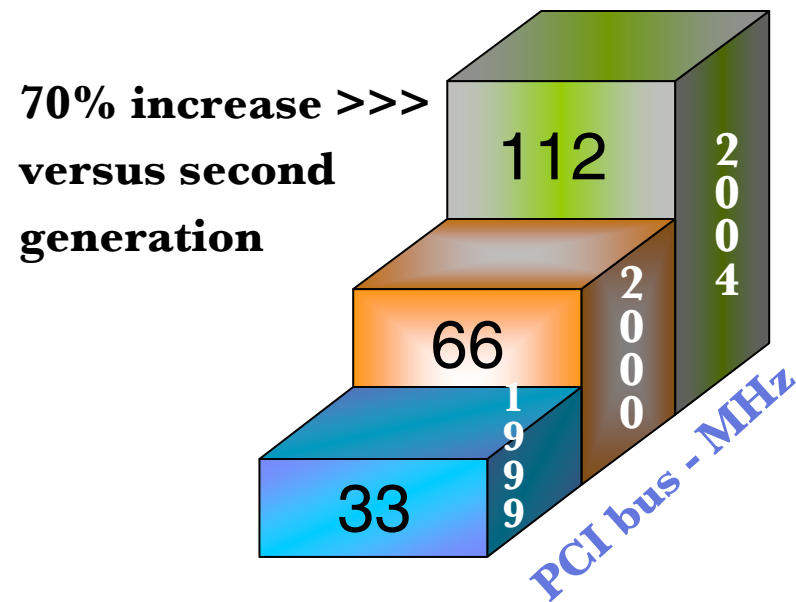
History of refresh of OSA technology

	1st Generation	2nd Generation	2nd Generation	2+ Generation	3rd Generation OSA-Express2
Server family	G5 / G6	z900 - 12/00	z900 – 12/00	z990 - 06/03 New NIC z900 (#1364, #1365)	z990, z9-109
Features	GbE – 06/99 Fast E – 01/00 ATM – 01/00	ATM Token Ring	Fast Ethernet GbE	1000BASE-T GbE (1364, 1365, 1366)	GbE, 10 GbE – 01/05 1000BASE-T – 09/05 (z9-109 only)
Number of ports	One port per feature	Two ports per feature	Two ports per feature	Two ports per feature	GbE – 2 ports 1000BASE-T – 2 ports 10 GbE – 1 port
Microprocessor	166 MHz	333 MHz	333 MHz	333 MHz	448 MHz GbE 500 MHz 10 GbE
PCI bus	32-bit 33 MHz	32-bit 33 MHz	64-bit 66 MHz	64-bit 66 MHz	64-bit 112 MHz GbE 125 MHz 10 GbE
Maximum Features / ports	12 / 12	12 / 24	12 / 24	24 / 48	24 / 48 z990, z9-109

OSA-Express2 Gigabit Ethernet A refresh of technology



Generation
Third
Second
First

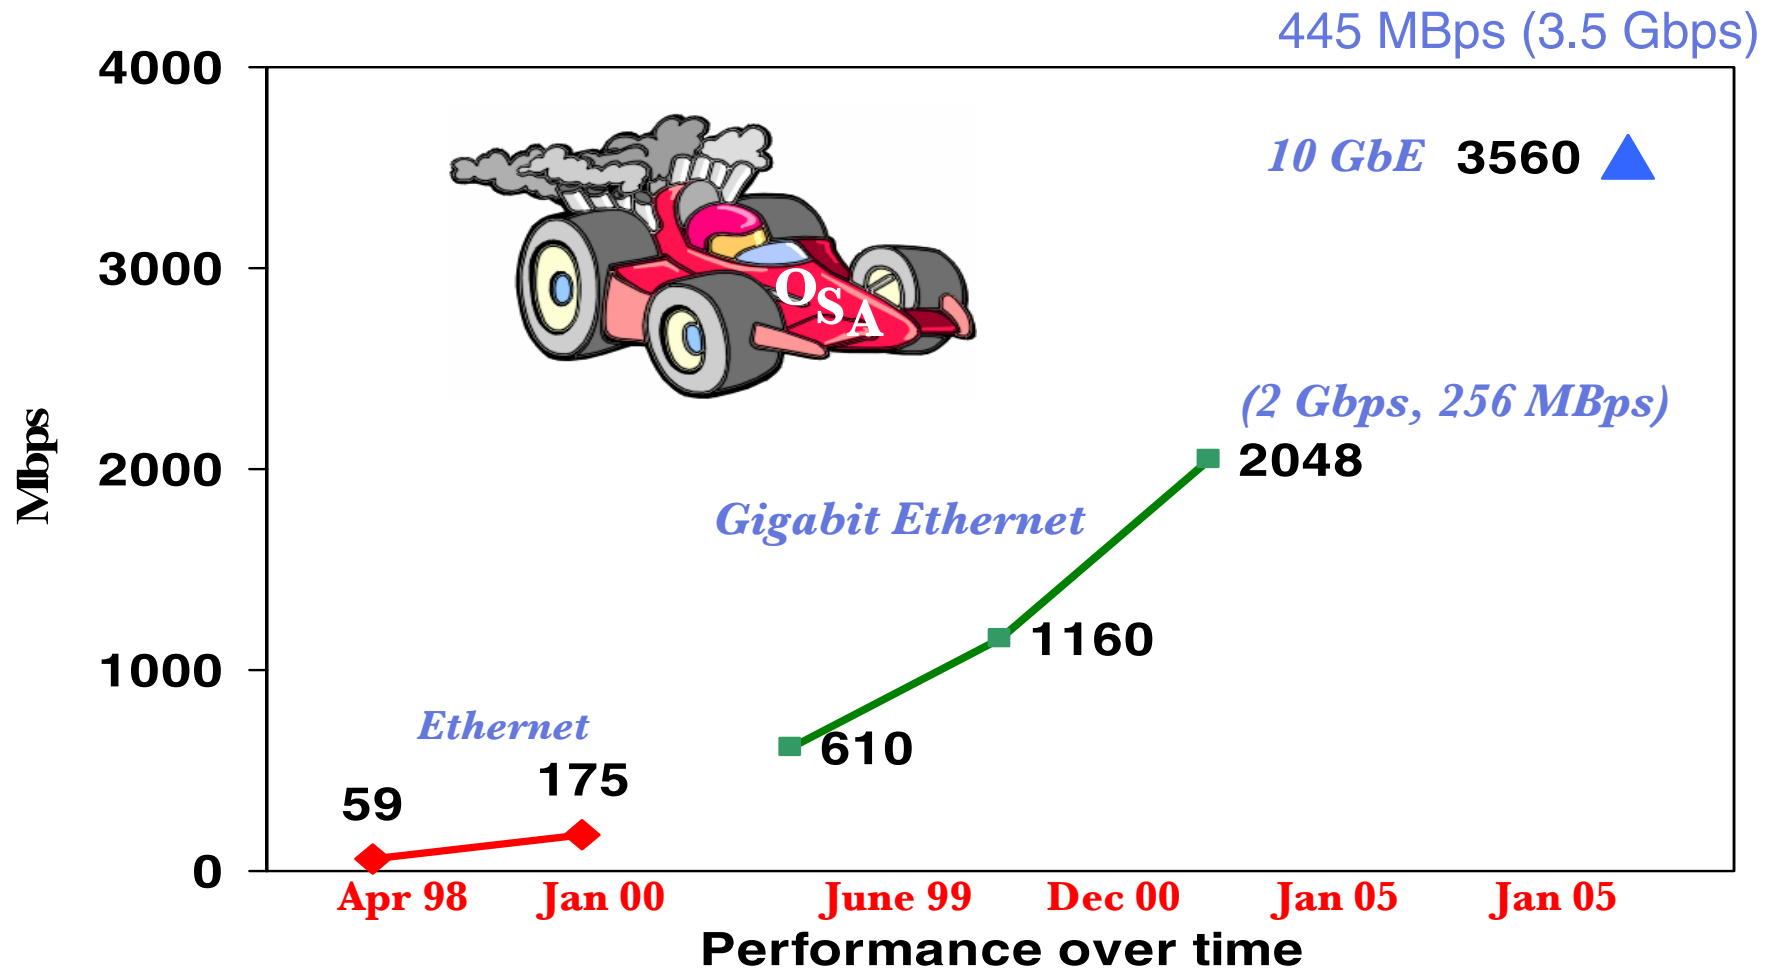


- OSA-Express2 is the 3rd generation of Ethernet technology

OSA-Express2 GbE is designed to achieve line speed - 1 Gbps in each direction

Ethernet on the LAN

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



OSA-Express performance

- ★ **OSA processor becomes more efficient as throughput increases**
- ★ **Window size**
 - **TCP window determines amount of data that the sender can transmit to receiver without needing an acknowledgment from the receiver**
 - **Faster and longer networks require larger windows to keep data flowing smoothly**
- ★ **Blocking**
 - **Performance is affected by the amount of data blocked together for transfer between OSA and TCP.**
- ★ **Frame size**
 - **Larger frames perform better**
 - **Larger frames reduce host and OSA processing costs**
 - **Size of frame depends on LAN type, MTU setting, size of data sent**
- ★ **Measure: throughput, transaction response time, server utilization**
 - **Bulk data transfer and interactive transactions**
- ★ **QDIO and jumbo frames (8992 byte MTUs) yield the highest streams**

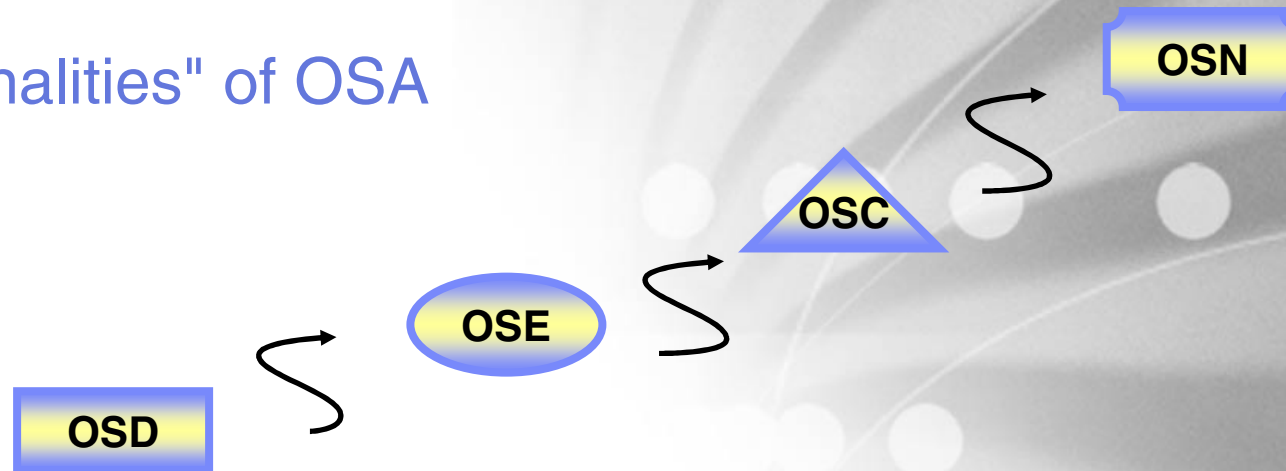
OSA solutions - how do they measure up?

- **Want to compare?**
- **Use AWM (Application Workload Modeler).**
 - ▶ Replaces Netmarks, and is now externally available.
- **Can model your current network configuration in your environment**
- **Rerun the same workload against a proposed alternative (e.g. OSA gigabit).**
- **Access information at**
 - ▶ <http://www-3.ibm.com/software/network/awm/>



IBM System z9™

The "personalities" of OSA



ON DEMAND BUSINESS™

zSeries Expo

San Francisco

September 10 – 23, 2005


© 2005 IBM Corporation

CHPID type OSD Queued Direct Input/Output since June 1999

- For TCP/IP traffic only - Layer 3
 - For SNA/APPN/HPR traffic with QDIO use TN3270, Enterprise Extender
- Protocol-independent when using Layer 2
 - ▶ **z/VM V5.1 with PTFs, Linux on System z9 and zSeries**
- Design for high speed communication

- ▶ **Reduced TCP/IP path length**
- ▶ **QDIO IP Processing Assist**
- ▶ **LPAR-to-LPAR Communication with port sharing**
- ▶ **Direct Memory Access (DMA) Protocol**
 - Memory-to-memory communication
 - I/O interrupts minimized
 - Continuous direct data exchanges
- ▶ **Dynamic customization**


- QDIO: z/OS, z/OS.e, z/VM, z/VSE, VSE/ESA, Linux on zSeries, TPF
- 10 Gigabit Ethernet, Gigabit Ethernet, 1000BASE-T Ethernet, Fast Ethernet



Functions - QDIO only

■ QDIO IP Processing Assists

- ▶ Performs all ARP processing
- ▶ Provides Multicast support
- ▶ Builds MAC and LLC headers
- ▶ Performs filtering - TCP/IP stack sees only IP datagrams




■ Checksum offload for IPv4 packets

- ▶ Calculates the TCP/UDP and IP header checksums
- ▶ Verifies the correctness of files
- ▶ Reduces host CPU cycles
- ▶ Checksum offload for IPv4 packets
 - ▶ z/OS V1.5, Linux on System z9 and zSeries
 - 1000BASE-T Ethernet (#1366, #3366)
 - Gigabit Ethernet (#1364, 1365, #3364, #3365)
 - 10 Gigabit Ethernet (#3368)

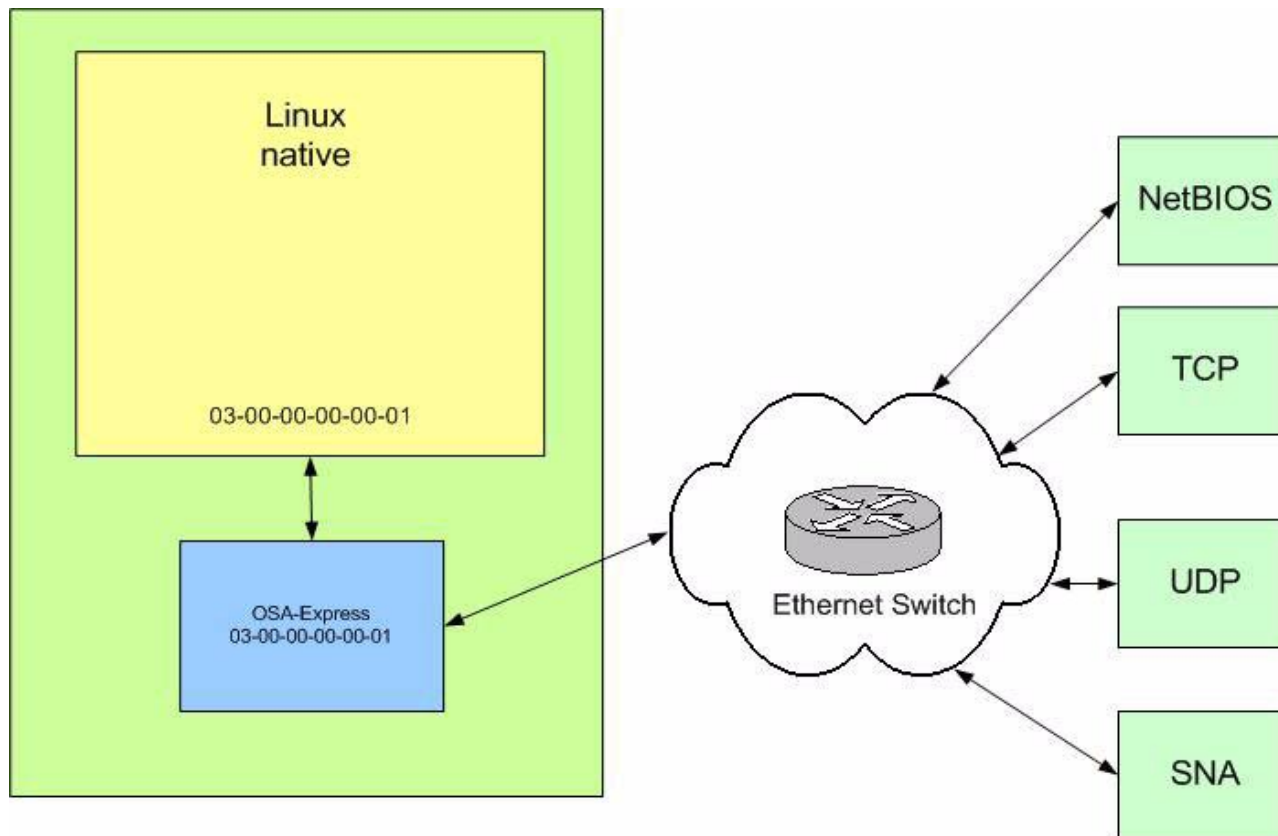
Layer 3 versus Layer 2

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

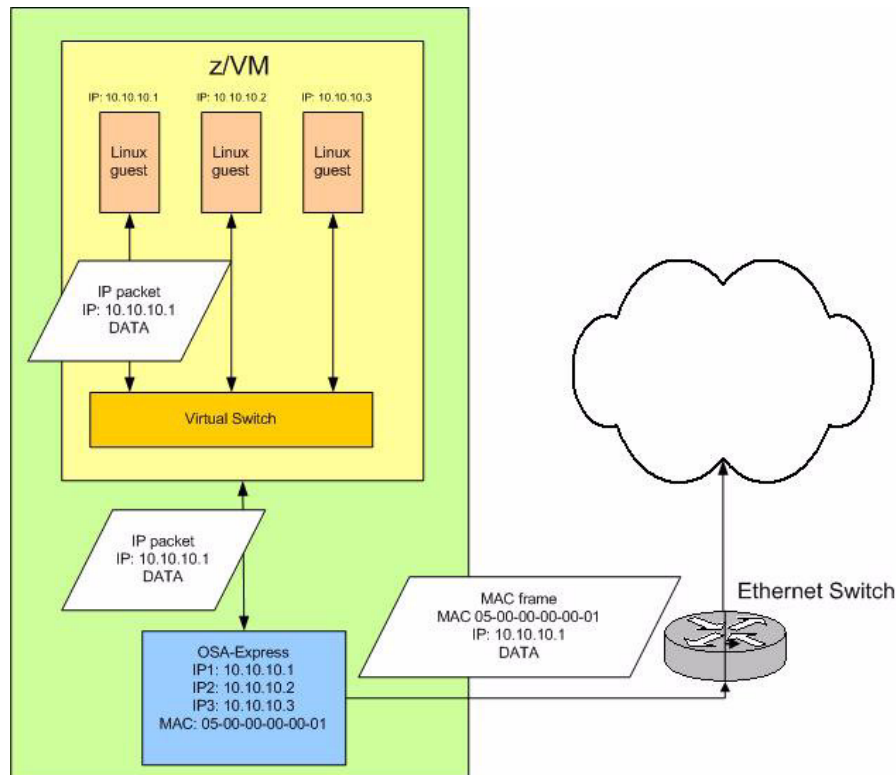
Layer 2 - protocol independent z/VM - Linux



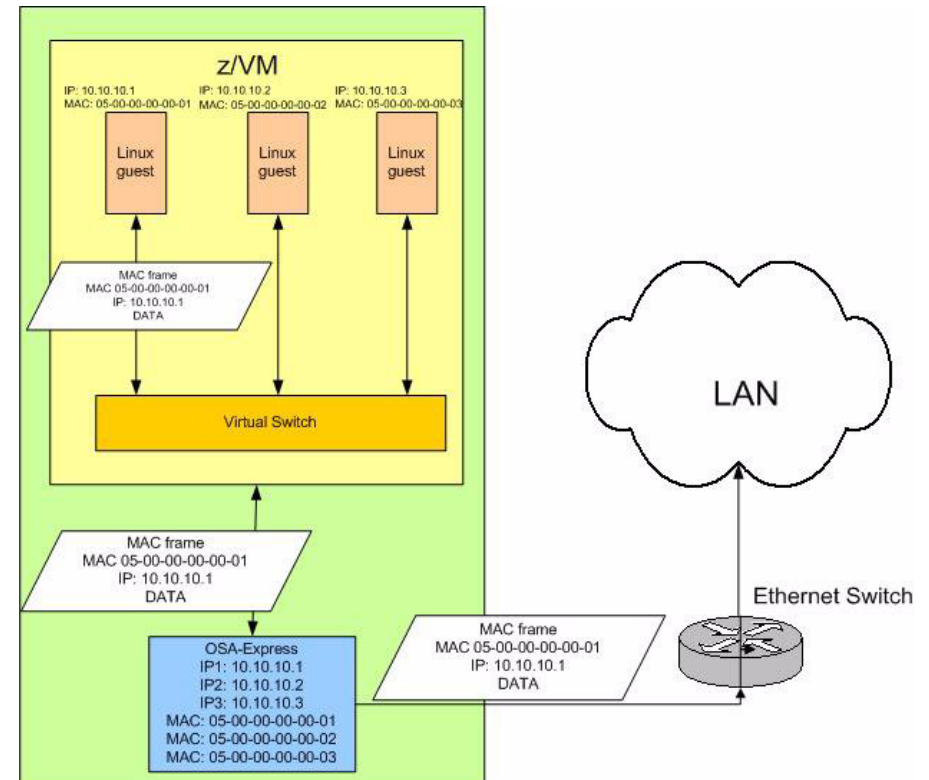
Layer 2 - protocol independent - Linux



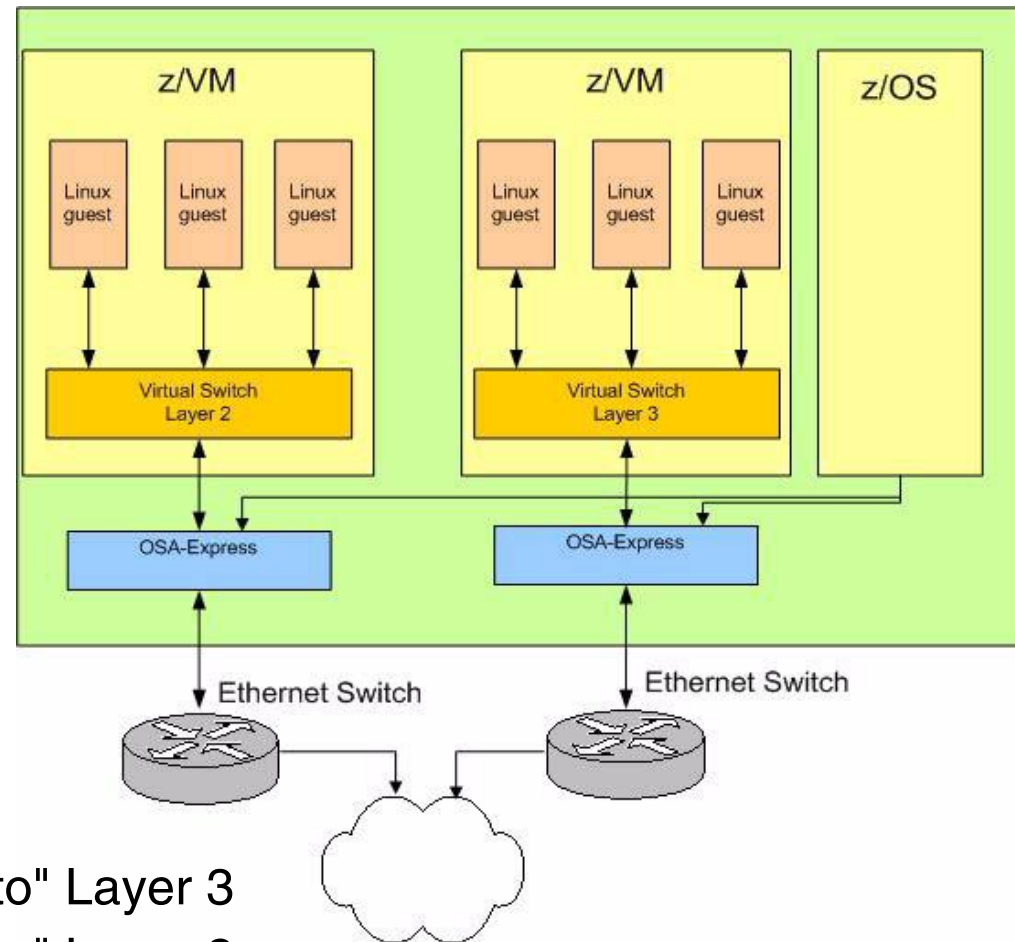
Layer 3 mode



Layer 2 mode



Layer 2 and Layer 3 traffic passing through same OSA



Layer 3 can only "talk to" Layer 3
Layer 2 can only "talk to" Layer 2


CHPID type OSE for SNA and/or TCP/IP since June 1999

- Used for all LAN communications
 - ✓ *Start I/O's*
 - ✓ *Channel Command Words (CCWs)*
 - ✓ *I/O interrupts*

- Involved in all data transfers
 - ✓ *I/O Processor (IOP)*
 - ✓ *Channel*

- Control unit layer
 - ✓ *Interprets CCWs (Read/Write)*
 - ✓ *Blocks/deblocks data*

- Device runs the LAN driver




Controller, Router, OSA-2, OSA-Express, OSA-Express2 non-QDIO

CHPID type OSC OSA-Integrated Console Controller (OSA-ICC) since May 2004

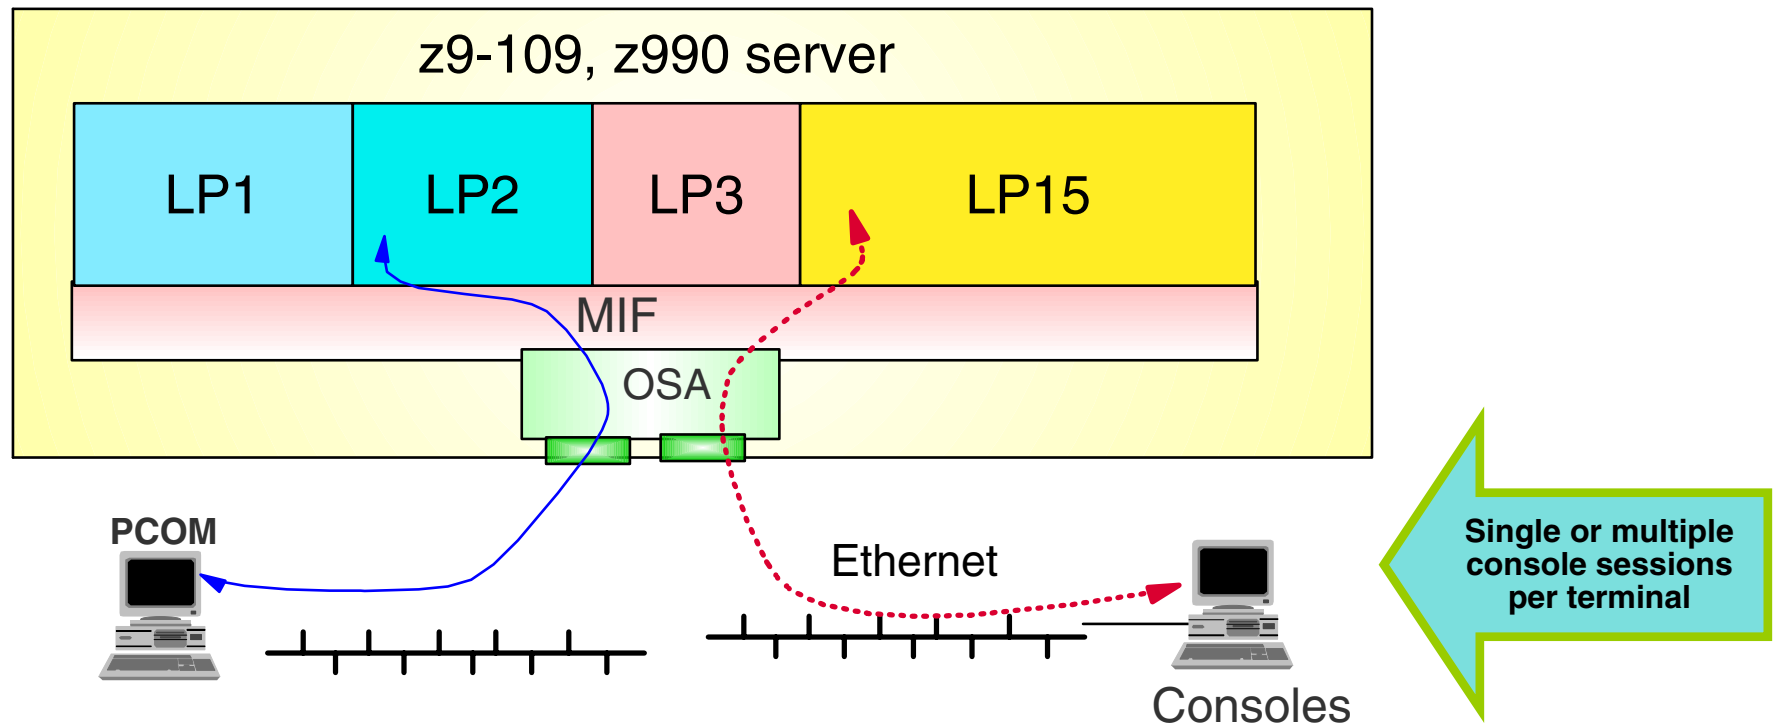
- Exclusive to 1000BASE-T Ethernet
 - ▶ CHPID type **OSC**
 - ▶ OSA-ICC configured on a port-by-port basis
 - Each port can be CHPID type OSC, OSD, OSE, or OSN
- LAN Connectivity
 - ▶ LAN-attached consoles running TN3270E clients
 - Non-SNA DFT data streams to TN3270E
 - Capable of operating at 10, 100, or 1000 Mbps (1 Gbps)
 - ▶ Uses RJ-45 connection to Category 5 Unshielded Twisted Pair (UTP) copper cabling
- Configuration support via Support Element (SE) and Hardware Management Console (HMC)
- Operating systems - z/OS, z/VM, VSE/ESA, TPF
- Can coexist in configurations using prior IBM 2074 models and older 3174 controllers
- No coaxial cable support or Token Ring support

OSA-Express2
1000BASE-T Ethernet #3366



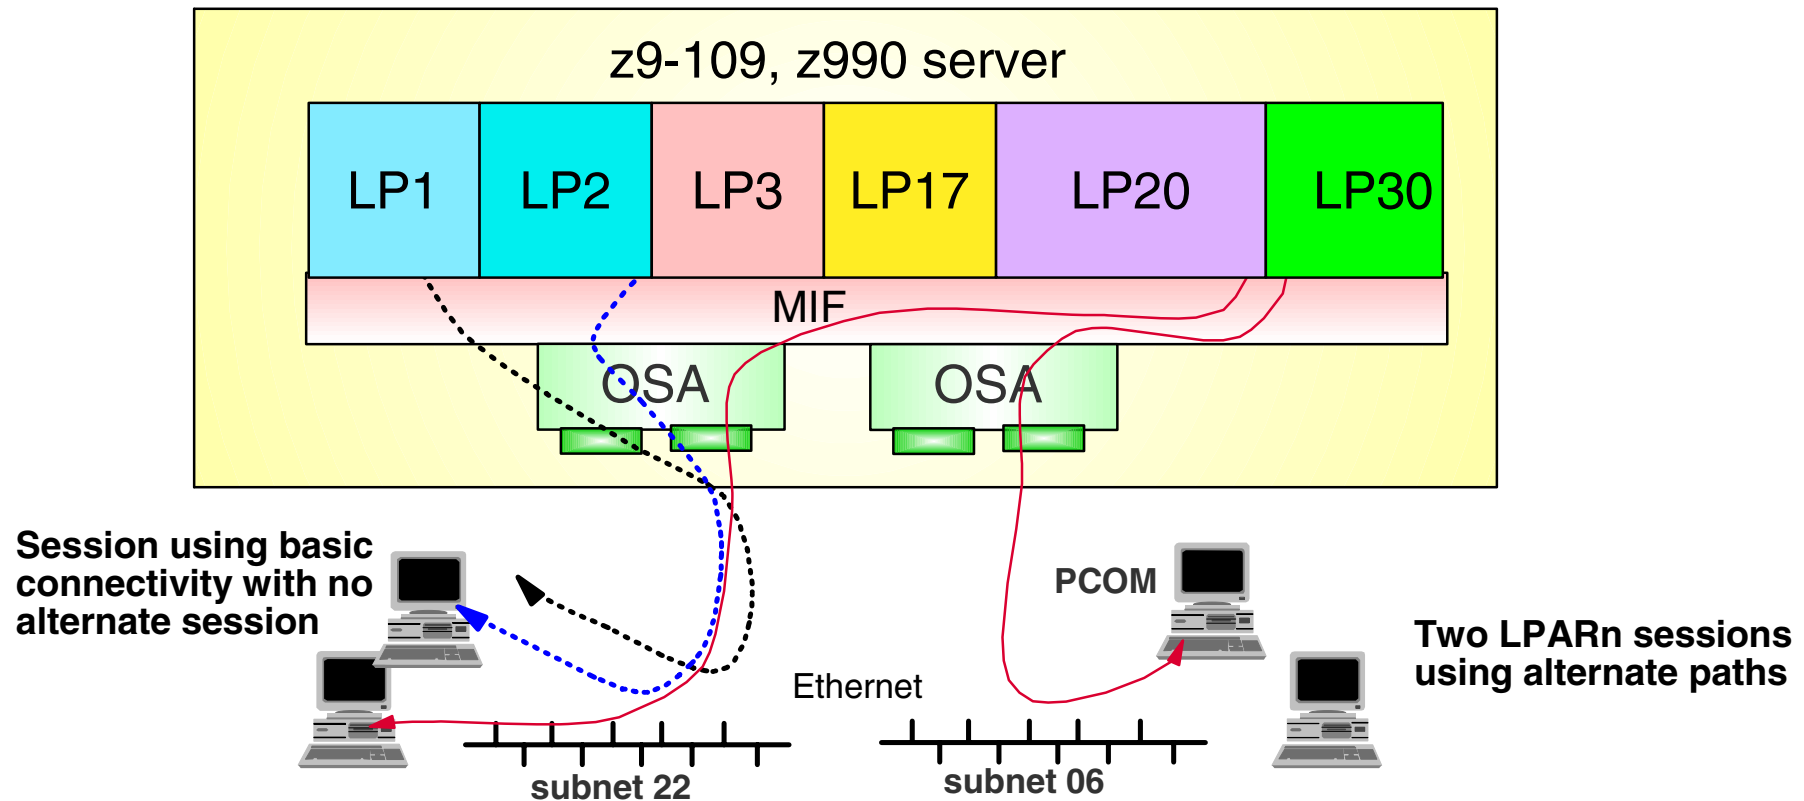
Also OSA-Express
1000BASE-T Ethernet #1366

OSA-ICC Single System with Alternate Sessions



- To IPL CEC and LPARs, and operating system console operations
- Up to 120 sessions per port across multiple LPARs .. MIF capable
 - ▶ Can run multiple sessions on a single workstation
- For increased console session availability, plan for separate consoles, on separate LANs, on separate ports, on separate OSA-Express features
 - ▶ Manual, disruptive console session switch possible

OSA-ICC Single System with Redundancy Configuration



- Up to 240 sessions to multiple LPARs on the CEC
- Session-level redundancy using different paths on ports, LANs, and consoles
 - ▶ Different console sessions from different LAN to same LPAR for dual connectivity
 - ▶ Manual, disruptive console session switch possible for failed session
 - ▶ Individual session cannot be shared between LPARs
 - ▶ LAN ports must attach to different subnets
- Loss of one port does not nondisruptively switch sessions to second port on same or different OSA

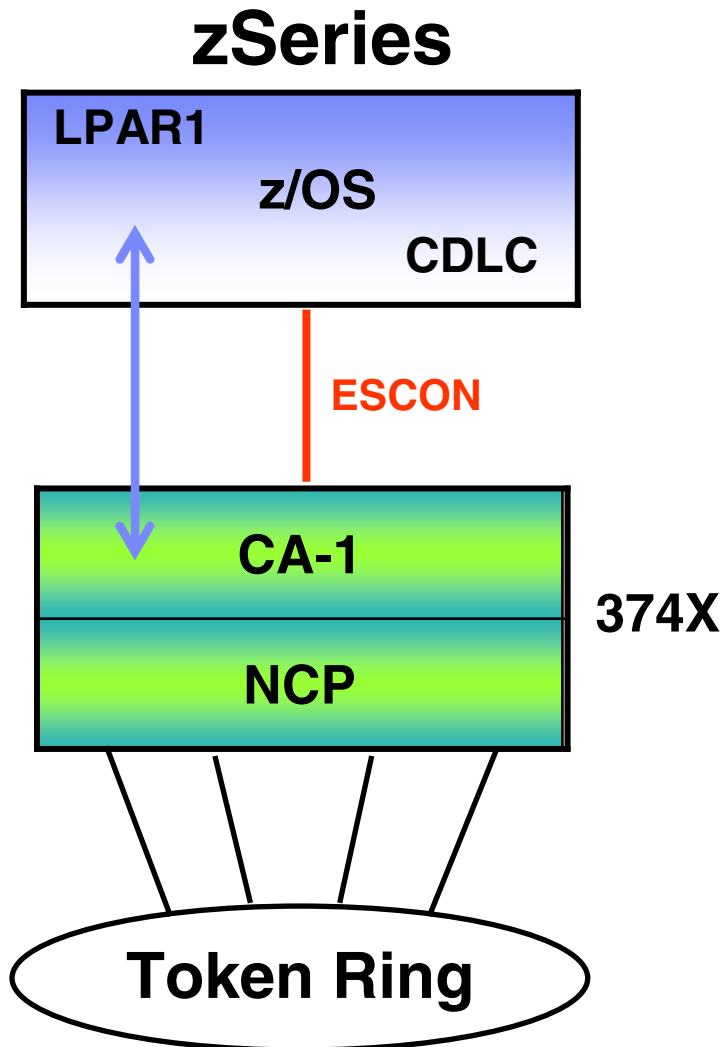
OSA-Express2 OSN (OSA for NCP)

September 2005

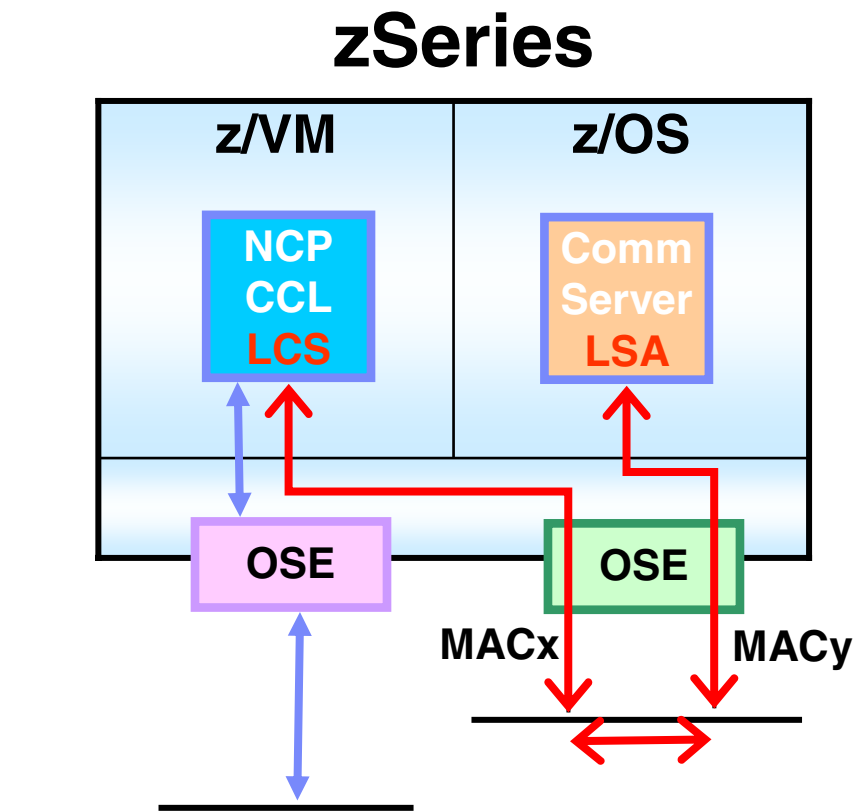
and IBM Communication Controller for Linux (CCL)

CCL V1.1	CCL V1.2 (SOD)
OSA – Fast Ethernet or 1000BASE-T Ethernet	OSA-Express2 GbE or 1000BASE-T Ethernet
Non-QDIO (CHPID type OSE)	OSA-Express2 OSN - OSA for NCP (CHPID type OSN)
External communication	Internal communication
Traffic flows on the LAN	Traffic flows LPAR-to LPAR

374X Network Control Program



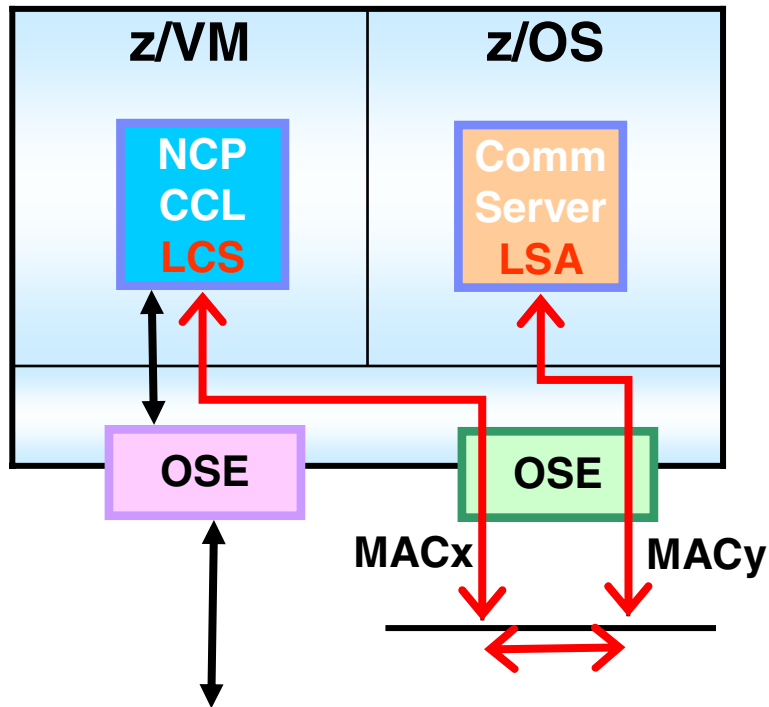
IBM Communication Controller for Linux V1.1



1. z/OS defined as LSA, own MAC/port
2. CCL defined as LCS, own MAC/port
3. 3rd port for SNA devices on LAN
4. All CHPID type OSE

IBM Communication Controller
for Linux V1.1
CCL traffi flowing on LAN

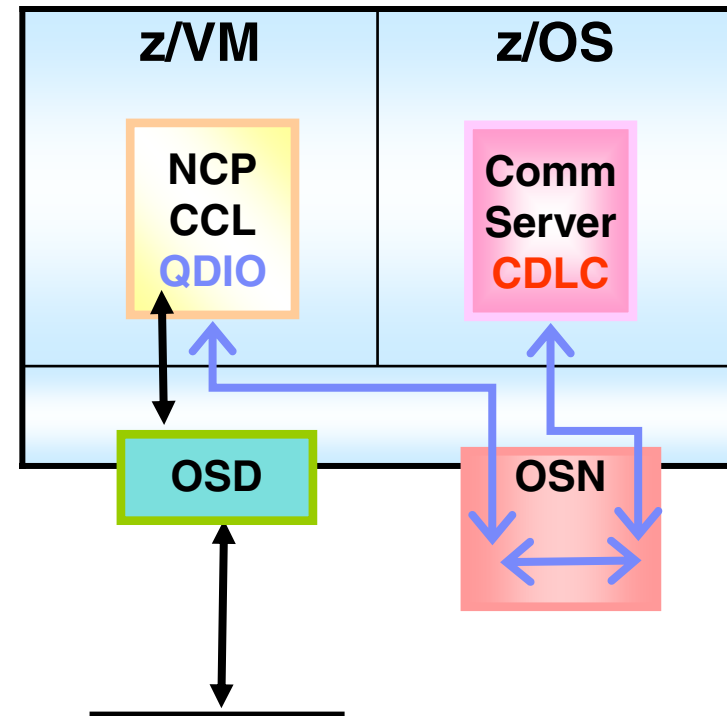
zSeries



1. z/OS defined as LSA, own MAC/port
2. CCL defined as LCS, own MAC/port
3. 3rd port for SNA devices on LAN
4. All CHPID type OSE

IBM Communication Controller
for Linux V1.2
CDLC traffic flowing LPAR-to-LPAR
Using OSA for NCP (OSN)

System z9



- ✓ Connectivity from OS to CCL
- ✓ SNA PUs Type 2.1 and 5
- ✓ VTAM/CDLC to CCL


CHPID type controls operation

CHPID type	Feature	SNA/APPN/HPR	TCP/IP	3270	NCP	OSA/SF required
OSD zSeries System z9	GbE, 10 GbE 1000BASE-T Ethernet Fast Ethernet	No (L3), Yes (L2) Use EE or TN3270E	Yes	No	No	No
OSE zSeries System z9	1000BASE-T Ethernet Fast Ethernet	Yes	Yes	No	No	Yes
OSC z990, z890 z9-109	1000BASE-T Ethernet	No	No	Yes	No	No
OSN z9-109 exclusive	1000BASE-T Ethernet GbE	No	No	No	Yes	No



IBM System z9™

The
newest
functions



ON DEMAND BUSINESS™

zSeries Expo

San Francisco

September 10 – 23, 2005

© 2005 IBM Corporation

VLAN management enhancement – support of GVRP exclusive to z9-109 and OSA-Express2

■ IEEE 802.1

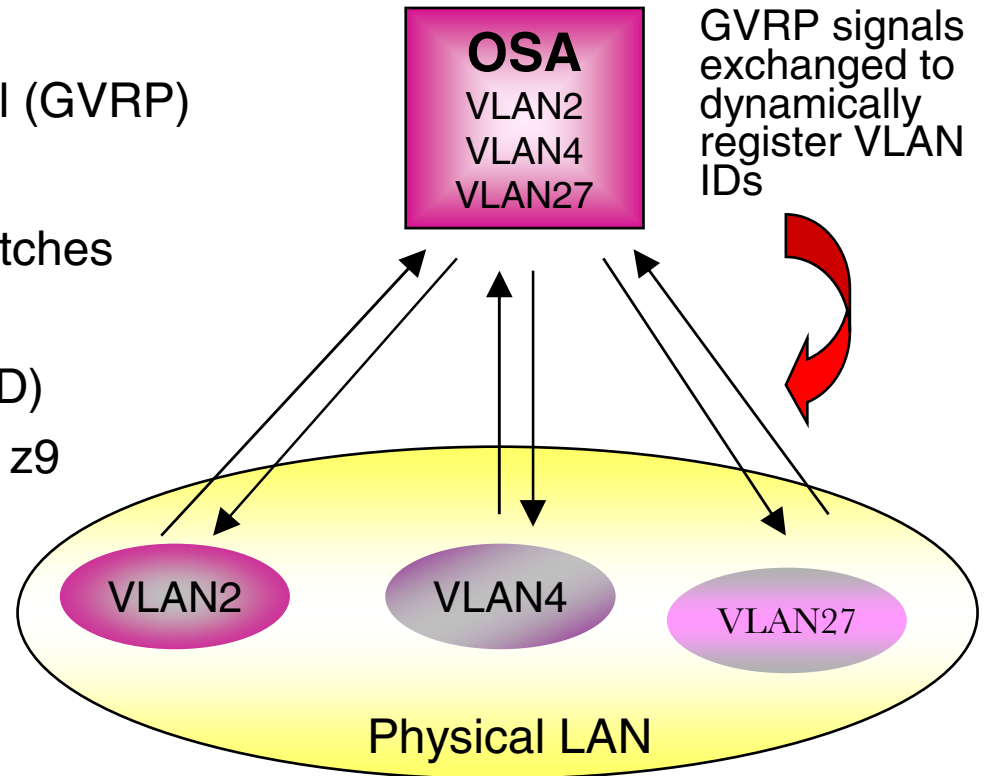
- GARP VLAN Registration Protocol (GVRP)



- Propagate VLAN information
 - Register host VLANs with switches
- OSA-Express2
 - QDIO mode (CHPID type OSD)
 - z/OS, z/VM, Linux on System z9

■ Current support:

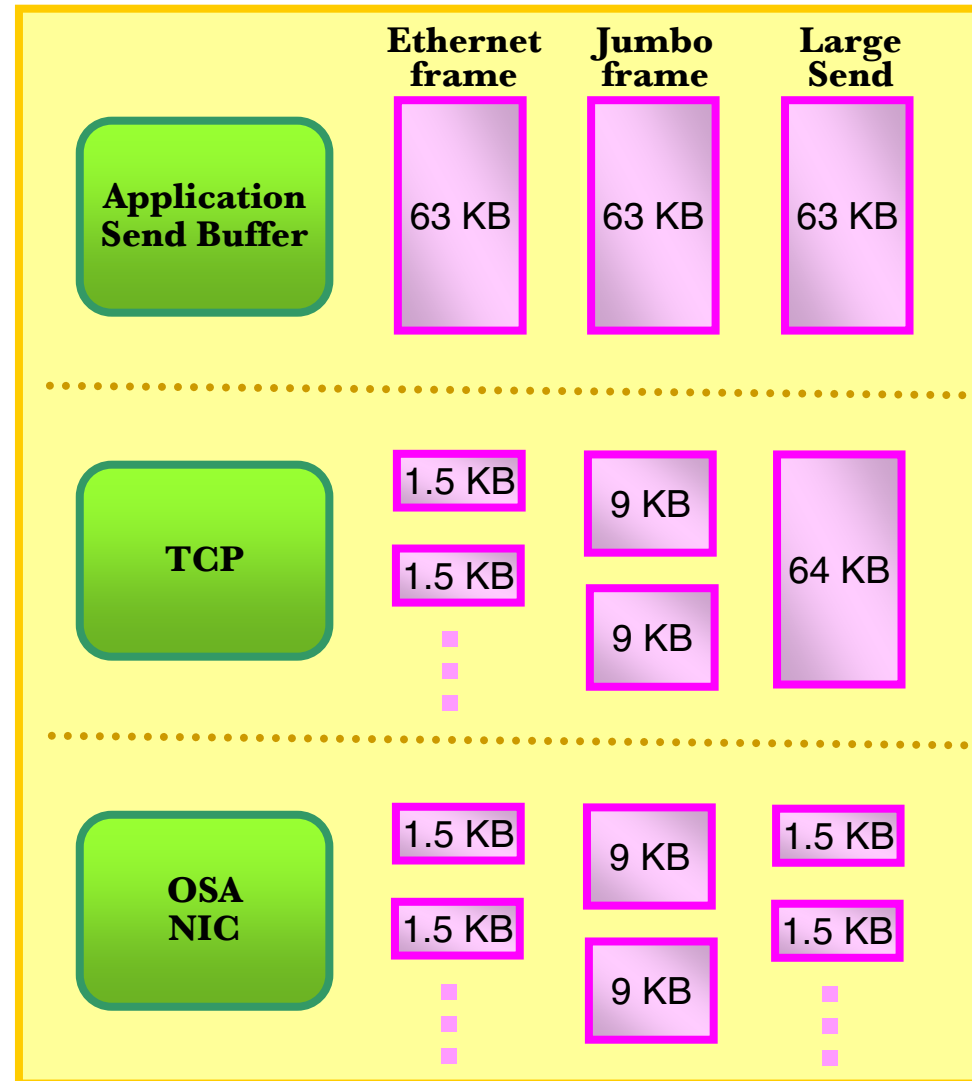
- VLAN tagging
 - z/OS, z/VM, Linux
- VLAN priority queueing
 - z/OS, Linux



Simplified network administration and management


Large Send for TCP/IP traffic (TCP segmentation offload) for CPU efficiency

- z990 and z9-109
- OSA-Express2 (GbE and 10 GbE)
- Segmentation of IP packets done by OSA NIC, not IP stack
- Offloads the TCP segmentation processing from host TCP/IP stack
- Host code path length reduced
- **Sends 64 kilobyte (KB) blocks to OSA**
- Processing performed by OSA NIC
 - ▶ IP/TCP checksum processing
 - ▶ TCP packet processing
 - ▶ Sends out 1.5 KB packets (1492 byte)
- For outbound traffic only
- For IPv4
- For unicast datagrams
- QDIO mode only (CHPID type OSD)
- Supported by
 - ▶ Linux on zSeries
 - ▶ z/OS V1.6 with PTF (2nd qtr 2005)
 - ▶ z/OS V1.7



640 TCP/IP stacks for improved virtualization

- **OSA-Express2 on z990, z9-109**
 - ▶ 640 TCP/IP stacks per OSA-Express2 port/CHPID
- **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)
- **Supported by**
 - ▶ z/OS and z/OS.e V1.6 with PTF
 - ▶ z/VM V5.1, Linux on zSeries, z/VSE V3.1



Limits	S/390 G5/G6	z990 Dec 00	z990 June 03	z990,z890 Oct 04 OSA-E	z990, z890 > Jan 05 OSA-E2
OSD					
Subchannels per stack	3	3	3	3	3
IP stacks per port/CHPID on server	15	80	160	160	640 \$
Subchannels per port	240	240	480	480	1920 \$
IP stacks per LPAR	15	80	84	160	640 \$
Devices per LPAR	240	240	254	480	1920 \$
Maximum control units	1	1	1	16	16

OSA-E = OSA-Express, OSA-E2 is OSA-Express2


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

SNMP - QDIO only

- **SNMP - Monitor network devices/functions, identify problems**
- **Direct SNMP subagent MIB for all of the OSA-Express features when configured in QDIO mode (CHPID type OSD)**
- **Do not require OSA/SF to obtain data**

- **Get and GetNext requests - zSeries (April 2002)**
- **dot3StatsTable - zSeries (May 2003)**
- **Performance data - zSeries (May 2003)**
- **Traps and Set - z990 (May 2004)**
 - ▶ Trap - Asynchronously receive an alert
 - ▶ Set - Change value (currently limited to set traps on/off)


- **Direct SNMP support for LCS - 990 (May 2004)**
 - ▶ Same support as listed above
 - ▶ Non-QDIO CHPID type OSE
 - ▶ z/OS V1.6





IBM System z9™

Cabling and connectors



ON DEMAND BUSINESS™

zSeries Expo


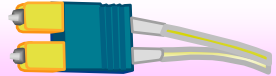


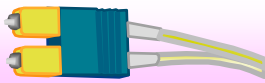
San Francisco

September 10 – 23, 2005

© 2005 IBM Corporation

GbE Features - Connector Change


May require a Conversion Kit to attach to switches

Feature	Connector	Previous Connector
Gigabit Ethernet LX Single mode (SM) fiber	 LC Duplex SM # 1364, # 3364	 SC Duplex SM #2364
Gigabit Ethernet SX Multimode (MM) fiber	 LC Duplex MM # 1365, # 3365	 SC Duplex MM # 2365
10 GbE Ethernet LR Single mode (SM) fiber	 SC Duplex SM # 3368	N / A



Conversion Kit example
2 meters (6.5 feet)

Mode Conditioning Patch Cables



I have a multimode fiber infrastructure.
Now what?

**MCP Cables can be used
(for 1 Gigabit links only)**

*A pair is required for each link
Offered as features on z900 only*

MCP Cable example
2 meters (6.5 feet)



Gigabit Ethernet Cabling Options


LX = Long wavelength 1300 nm transceiver

SX - Short wavelength 850 nm transceiver


<p>9 micron SM fiber</p>	<ul style="list-style-type: none"> LX transceiver/feature on each end Requires 9μ single mode (SM) fiber <table border="1"> <tr> <td>1 Gbps</td> <td>4.6 dB</td> <td>5 km (3.1 miles)</td> </tr> </table>	1 Gbps	4.6 dB	5 km (3.1 miles)
1 Gbps	4.6 dB	5 km (3.1 miles)		
<p>50 or 62.5 micron MM fiber</p>	<ul style="list-style-type: none"> LX transceiver/feature on each end combined with a pair of MCP cables Uses current 50 or 62.5 multimode (MM) fiber infrastructure Reduced distance and link budget <table border="1"> <tr> <td>1 Gbps</td> <td>2.4 dB</td> <td>550 meters (1804 feet)</td> </tr> </table>	1 Gbps	2.4 dB	550 meters (1804 feet)
1 Gbps	2.4 dB	550 meters (1804 feet)		
<p>50 micron MM fiber</p>	<ul style="list-style-type: none"> SX transceiver/feature on each end Requires 50μ multimode fiber Reduced distance and link budget <table border="1"> <tr> <td>1 Gbps</td> <td>3.6 dB</td> <td>550 meters (1804 feet)</td> </tr> </table>	1 Gbps	3.6 dB	550 meters (1804 feet)
1 Gbps	3.6 dB	550 meters (1804 feet)		
<p>62.5 micron MM fiber</p>	<ul style="list-style-type: none"> SX transceiver/feature on each end Requires 62.5μ multimode fiber Reduced distance and link budget <table border="1"> <tr> <td>1 Gbps</td> <td>2.6 dB</td> <td>220 meters (722 feet)</td> </tr> </table> <p><i>Link Rate</i> <i>Link Budget</i> <i>Unrepeated Distance</i></p>	1 Gbps	2.6 dB	220 meters (722 feet)
1 Gbps	2.6 dB	220 meters (722 feet)		

OSA-Express2 10 Gigabit Ethernet LR cabling

LR= Long reach 1310 nm transceiver



Connector



SC Duplex SM

- **z990 and z9-109**

Ethernet connectors and cabling

Feature #	Feature name	Connector	Cabling
OSA-Express on z900			
2364	OSA-Express GbE LX	SC Duplex	9μ SM
2365	OSA-Express GbE SX	SC Duplex	50, 62.5μ MM
2366	OSA-Express Fast Ethernet	RJ-45	Category 5 UTP
OSA-Express on z990, z900 (#1364, #1365)			
1364	OSA-Express GbE LX	LC Duplex	9μ SM
1365	OSA-Express GbE SX	LC Duplex	50, 62.5μ MM
1366	OSA-Express 1000BASE-T Ethernet	RJ-45	Category 5 UTP
OSA-Express2 on z9-109 and z990			
3364	OSA-Express2 GbE LX	LC Duplex	9μ SM
3365	OSA-Express2 GbE SX	LC Duplex	50, 62.5μ MM
3366 z9-109 only	OSA-Express2 1000BASE-T Ethernet	RJ-45	Category 5 UTP
3368	OSA-Express2 10 GbE LR	SC Duplex	9μ SM

SM = Single mode fiber, MM = Multimode fiber

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

UTP = Unshielded Twisted Pair, STP - Shielded Twisted Pair

IEEE 802.3ae Ethernet specification

- The following chart summarizes the unrepeated distances and link loss budgets supported by the standards. The link loss budget is the channel insertion loss plus the unallocated link margin as identified by the standard.
 - ▶ Ethernet physical layer specification (IEEE 802.3)

- As a light signal traverses a fiber optic cable, the light signal loses strength
 - ▶ dB (decibel) is the metric used to measure the signal strength (loss or gain)
 - ▶ The Link loss budget identified in the following slides is represented in dBs.

- Factors that contribute to the loss of signal strength
 - ▶ Number of connections (conversion kits, MCP cables, jumpers, trunks, patch panels)
 - ▶ Length of the fiber optic cable

- All industry standard links (FICON, FCP, Ethernet) follow published standards.

Ethernet physical layer specification (IEEE 802.3ae)


μ = microns	Nanometer (nm)	1 Gbps Ethernet		10 Gbps Ethernet	
Fiber core (μ) Light source	Fiber Bandwidth @ wavelength	Unrepeat ed distance	Link loss budget *	Unrepeated distance	Link loss budget *
9 μ SM LX	@ 1310 nm	5 km 3.1 miles	4.6 dB	10 km 6.2 miles	6.0 dB
9 μ SM LX with MCP cable 50 or 62.5 μ	500 MHz km	550 meters 1804 feet	2.4 dB	N / A	N / A
Newest 50 μ MM SX	2000 MHz km @ 850 nm	Not qualified		300 meters 984 feet	2.6 dB
50 μ MM SX	500 MHz @ 850 nm	550 meters 1804 feet	3.6 dB	82 meters 269 feet	2.3 dB
62.5 μ MM SX	200 MHz km @ 850 nm	275 meters 902 feet	2.6 dB	33 meters 108 feet	2.5 dB
62.5 μ MM SX	160 MHz * km @ 850 nm	220 meters 722 feet	2.6 dB	26 meters 85 feet	2.6 dB

* The **link loss budget** is the channel insertion loss + unallocated link margin as defined by the standard.

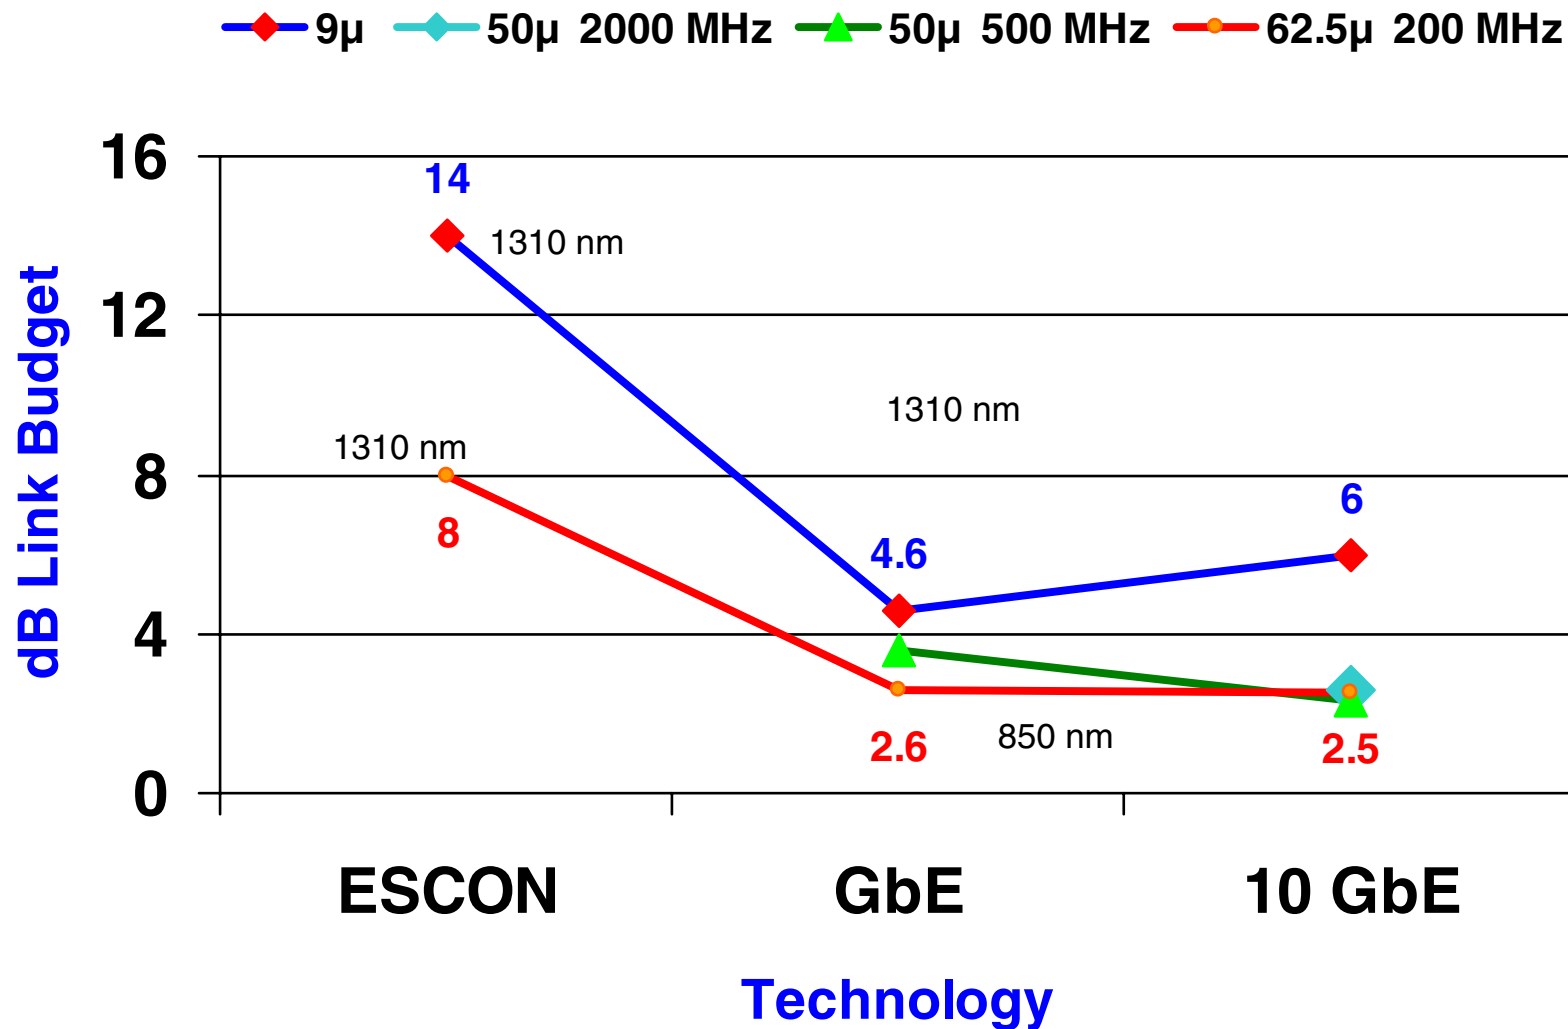
** The aqua 2000 MHz km multimode fiber became available September, 2003.

*** Most often applicable to currently installed ESCON environments

Gigabit Ethernet distances when using multimode fiber optic cabling

Compare using 9 μ single mode fiber 10,000 meters (10 km)◆ 50 μ 2000 MHz ▲ 50 μ 500 MHz ● 62.5 μ 200 MHz ▲ 62.5 μ 160 MHz

Gigabit Ethernet link loss budget at high data rates



OSA-Express2

OSA-Express

OSD, OSE, OSC, OSN

Connectivity

Distance

Thank You!

More information on the www

- Visit <http://www-306.ibm.com/common/ssi/OIX.wss>
 - ▶ Announcement letters
 - ▶ Sales manuals
- Refer to announcement letters 104-346, 105-241
- The *sales manual* includes:
 - ▶ All of the software requirements
 - For OSA-Express2 and OSA-Express and all of the functions
 - Feature descriptions including minimums, maximums, limitations, etc.

Direct SNMP Support - module

▪ **Direct SNMP MIB module**

- ▶ OSA-Express2 and OSA-Express features
- ▶ Logical Channel Subsystems (LCSSs)
- ▶ Updated performance table with more detailed information
- ▶ Get, GetNext, Traps, Set

▪ **OSA Direct SNMP MIB module is available via Resource Link**

- ▶ Resource Link is set up to house multiple copies of the MIB in anticipation of version or release changes/updates.


▪ **To retrieve the MIB follow these steps:**

- ▶ Locate Resource Link at: www.ibm.com/servers/resourcelink
- ▶ Log in
- ▶ Click on Library (on the bottom of the page or the left in the navigation bar)
- ▶ Under "Library shortcuts" on the right side of the screen, click on "Open System Adapter (OSA) Library"
- ▶ Choose "OSA-Express SNMP Direct MIB Module" for a description, or click on "TXT" for the module itself

▪ **Another technique:**

- ▶ Once logged on, select "site search" from the navigation bar on the left of the screen
- ▶ Type in "MIB" as the search argument, and hit submit at the bottom of the page

On the Internet

- IBM Resource Link, Web-based tool
 - ▶ www.ibm.com/servers/resourcelink/
 - Services section: zSeries Fiber Cabling Service
 - Planning section/Physical Planning
 - Physical Planning manuals, GIM
 - Education section: System z9 and zSeries courses
 - General Information for Planning a Physical Site (GIM)
- <http://www.ibm.com/services/networking/>
 - ▶ IT services - Product and Enterprise cabling offerings
- <http://www.redbooks.ibm.com>
 - ▶ IBM Redbooks 
- <http://www.ibm.com/servers/eserver/zseries/networking>
 - ▶ The network connectivity home page
- <http://www.ibm.com/servers/eserver/zseries/connectivity>
 - ▶ The I/O connectivity home page
 - ▶ Go to this location for a list of FICON/FCP supported devices
- <http://www-306.ibm.com/common/ssi/OIX.wss>
 - ▶ Announcement Letters, sales manual

On the Internet

**IBM Resource Link
www.ibm.com/servers/resourcelink/**

**A fiber optic cabling presentation with
narrative is available**

**Covers ISCs, ETR, ESCON, FICON/FCP, OSA
Overview of each feature with fiber optic
cabling requirements**

**FQC, Conversion kits, MCP cables
Extended distance implications**

**IBM Networking Services
zSeries fiber cabling services
Enterprise fiber cabling services**

**You can locate the Fiber Optic Cabling
presentation on Resource Link and subscribe
to receive updates.**

**After logging in click on Education in the blue
on the left.**

**Under zSeries Courses click on z990.
Click on Fiber Optic Cabling.**

On the Internet – current as of August 12, 2005

URL	Content
http://www.ibm.com/eserver/zseries/networking/technology.html	Networking technologies
http://www.ibm.com/software/network	Software - Networking & communications
http://www.ibm.com/software/network/commserver	Communications Server
http://www.ibm.com/software/network/commserver/library	CS white papers, product documentations
http://www.ibm.com/support/techdocs/	Advanced technical support (flashes, presentations, white papers)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)
http://www.ibm.com/servers/eserver/zseries/networking/dsnmp.html	Direct SNMP Subagent MIB

OSA reference materials

SA22-7935	OSA-Express Customer's Guide and Reference
SA22-7990	OSA-Express Integrated Console Controller User's Guide
SA22-7476	z800, z900 Open Systems Adapter-Express Customer's Guide and Reference
SG24-5444	System z9 and zSeries Connectivity Handbook (Redbook)
SG24-5948	OSA-Express Implementation Guide (Redbook)
SG24-6364	OSA-ICC Implementation Guide (Redbook)
GA23-0367-09	Planning for Fiber Optic Links
SC33-7991-10	z/OS V1R6.0 Resource Measurement Facility (RMF) Report Analysis (-10 has System z9 updates)
Look in CS library	For a complete description of the SNMP management data supported by OSA-Express features: z/OS Communications Server: IP Configuration Guide and the SNMP chapter of IP System Administrator's Commands

OSA-Express VLANs

VLAN support			
	Linux on System z9 and zSeries	z/OS	z/VM
OSA-Express OSA-Express2	Yes, IPv4, IPv6	Yes, V1.5 IPv4, IPv6	Yes, V4.4 - - IPv4 Yes, V5.1 - - IPv6
z/VM guest LAN	Yes, IPv4	Yes, IPv4	Yes, IPv4
IPv6 support			
OSA-Express OSA-Express2	Yes	Yes, V1.4	? NO
z/VM guest LAN	Yes, z/VM V4.4	Yes, z/VM V1.4	Yes, V4.4
Broadcast support			
OSA-Express OSA-Express2	Yes, IPv4	Yes, IPv4	Yes, IPv4
z/VM guest LAN	Yes, IPv4	Yes, IPv4	Yes, IPv4
Generic Attribute Registration Protocol VLAN Registration Protocol (GVRP)			
OSA-Express2 z9-109, QDIO	No	Yes. V1.7 with PTF	Yes, V5.1 with PTF 2Q06