SHARE
User-Driven Training
Event & Expo
Long Beach, CA February 22 - 27, 2004 Long Beach Convention Center

zSeries OSA-Express Update

Session # 3803

Connie K. Beuselinck
IBM Corporation
Poughkeepsie, N.Y.
conniek@us.ibm.com

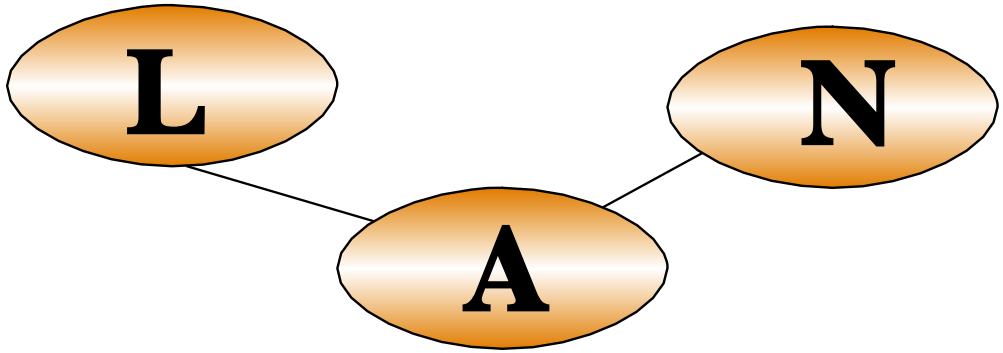

February 23, 2004

Trademarks

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

390
ACF/VTAM*
AIX*
APPN*
CICS*
DB2*
e-business logo*
ESCON*
GDPS*
Geographically Dispersed Parallel Sysplex*
FICON
HiperSockets
HPR
IBM *
IBM logo *
IMS
Magstar *
MVS/ESA
Net.Data*
Netfinity
OS/2 *
OS/390 *
Parallel Sysplex *
pSeries
RACF *
RMF
RS/6000 *
S/390 *
S/390 Parallel Enterprise Server
Sysplex Timer *
Virtual Image Facility *
VM/ESA *
VSE/ESA
VTAM *
WebSphere *
xSeries
z/OS
zSeries
z/VM

* 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.

Family of LAN Adapters

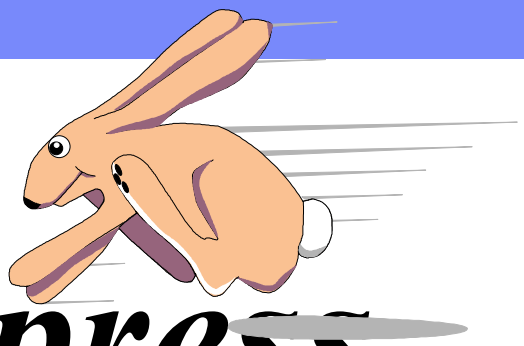

OSA-Express

Open, Industry-standard
Local Area Network Interfaces

Gigabit Ethernet
(June 1999)

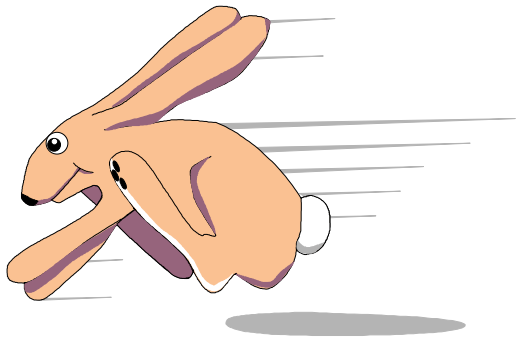

Fast Ethernet
(January 2000)

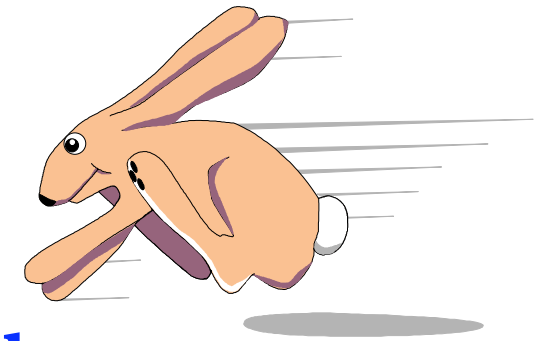

Token Ring
(October 2001)

1000BASE-T Ethernet
(June, 2003)

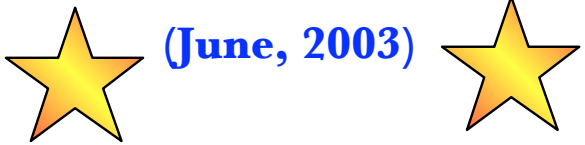

Any-to-Any Connectivity

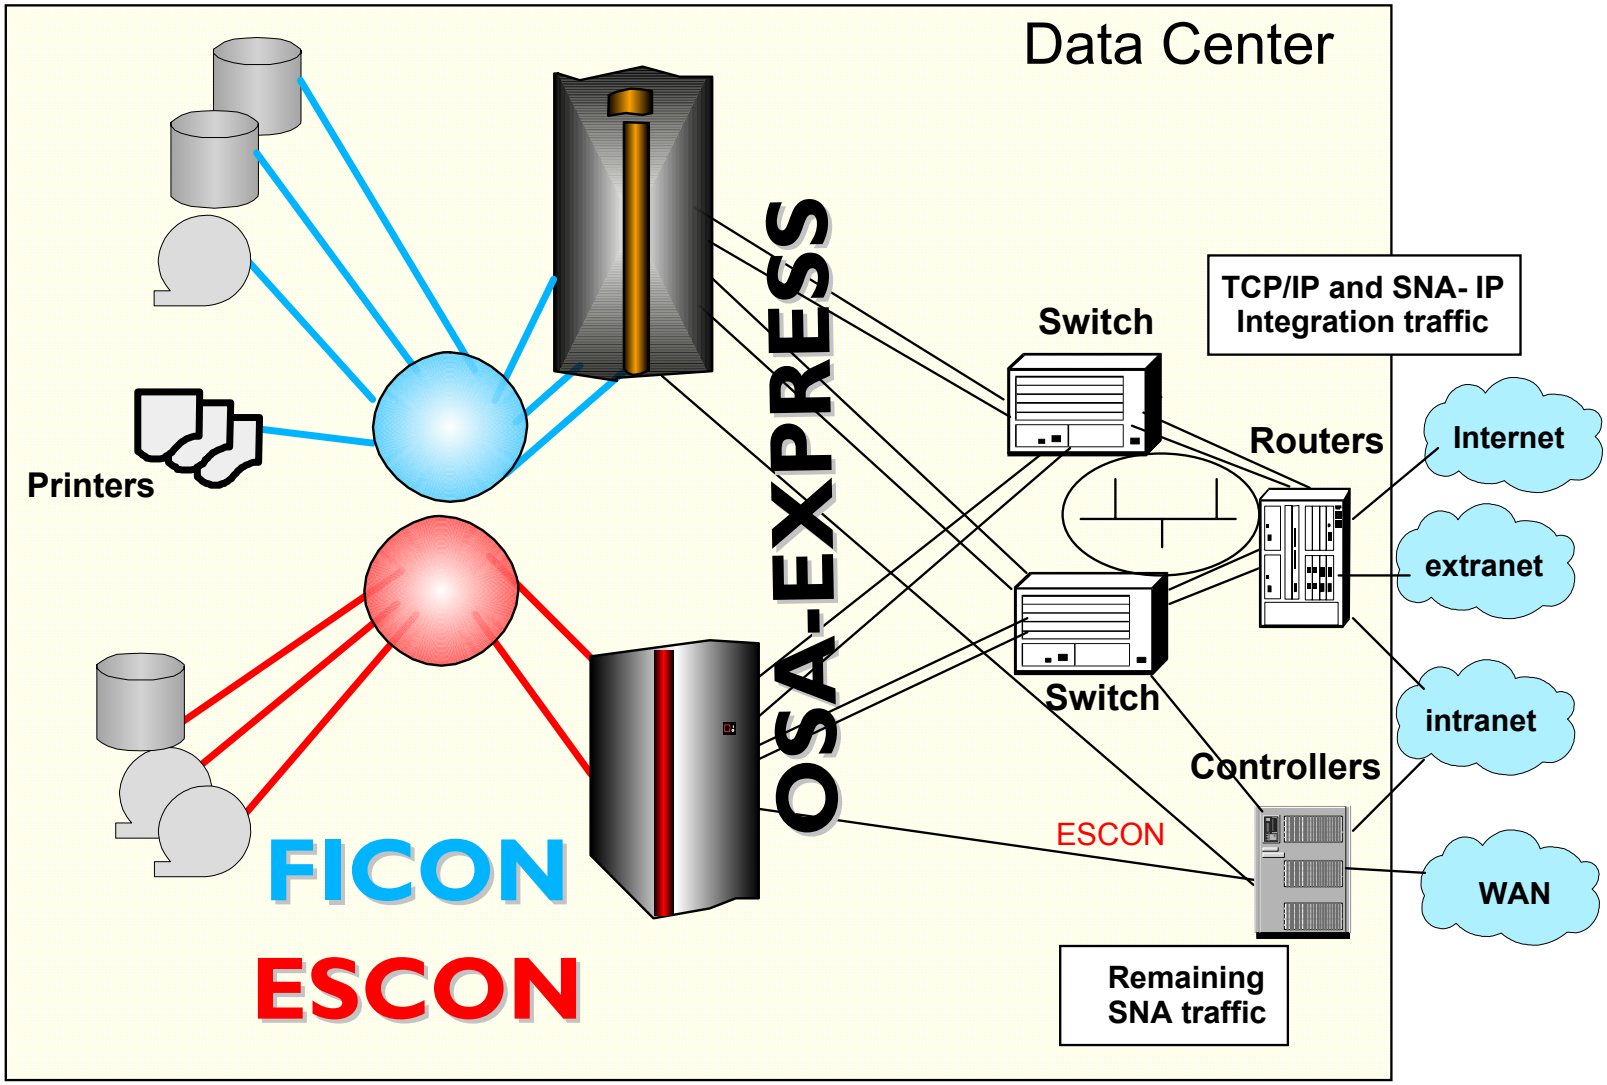

Speed Zones on the Information Highway

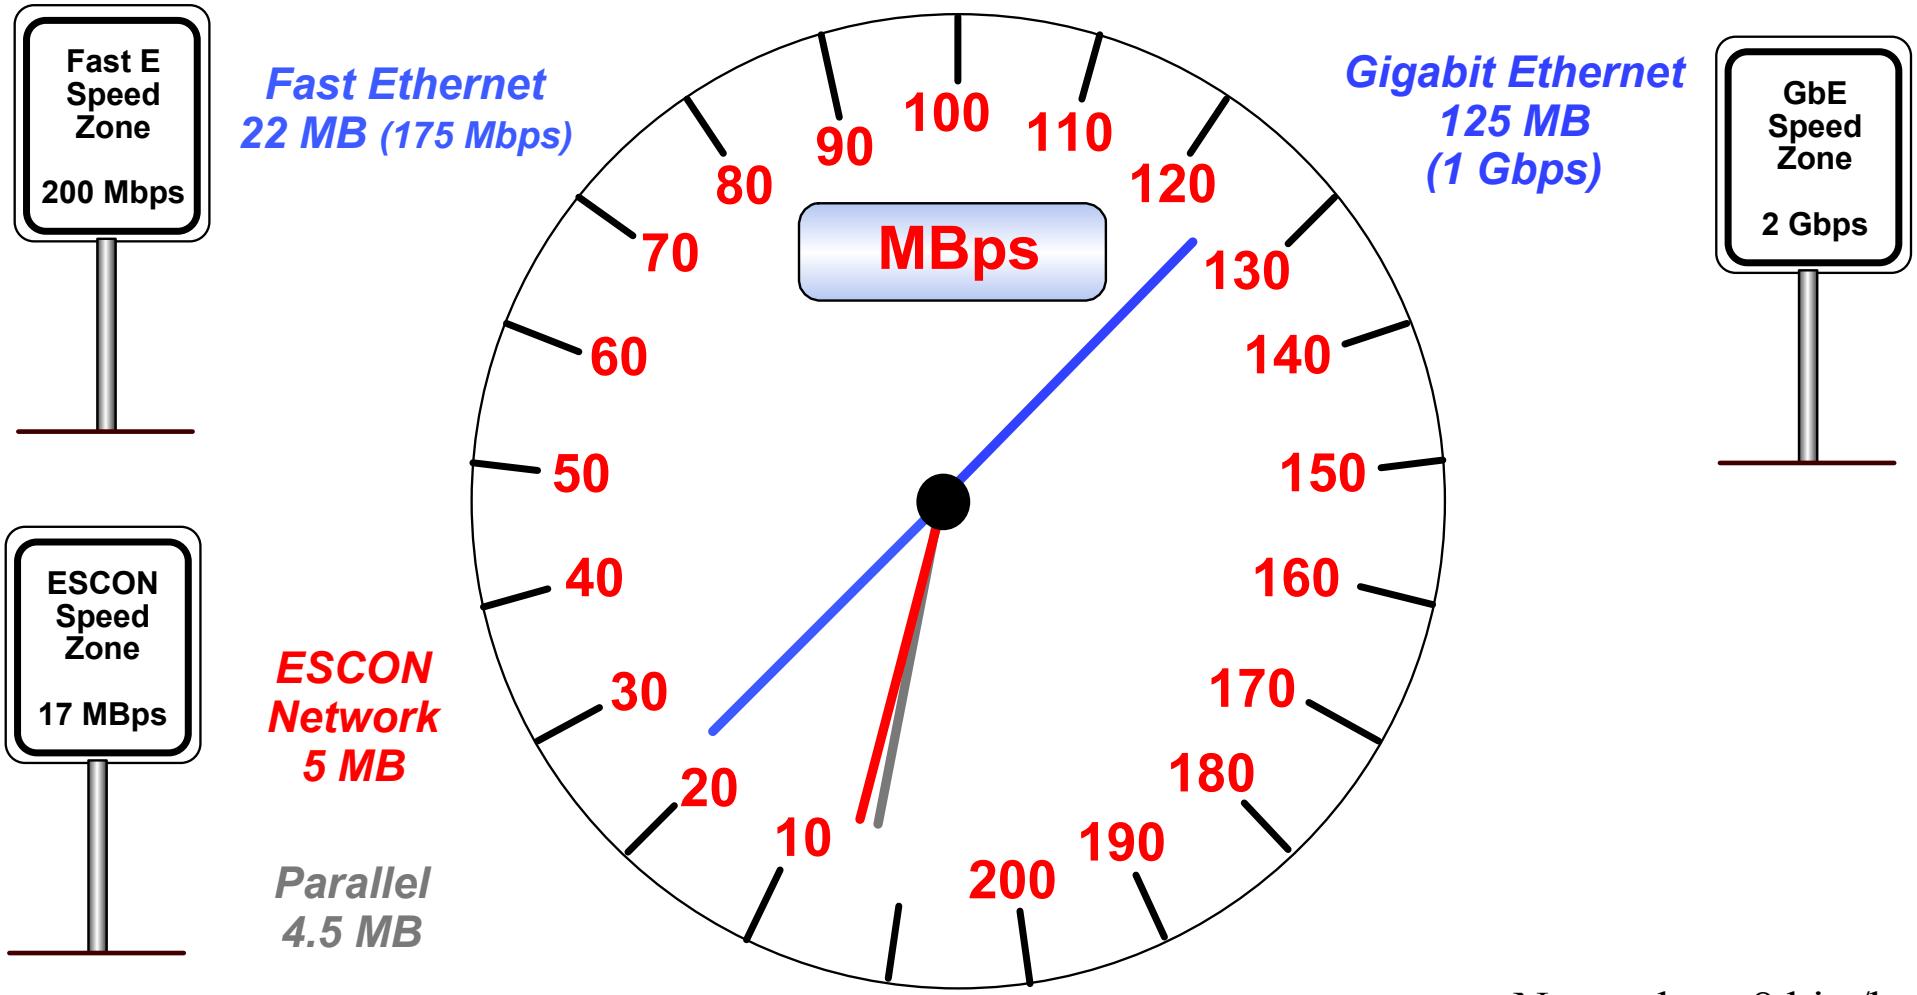

MBps = MegaBytes per second
Gbps = Gigabits per second
Mbps = Megabits per second

Networks = 8 bits/byte
I/O = 10 bits/byte

Breaking the Barrier

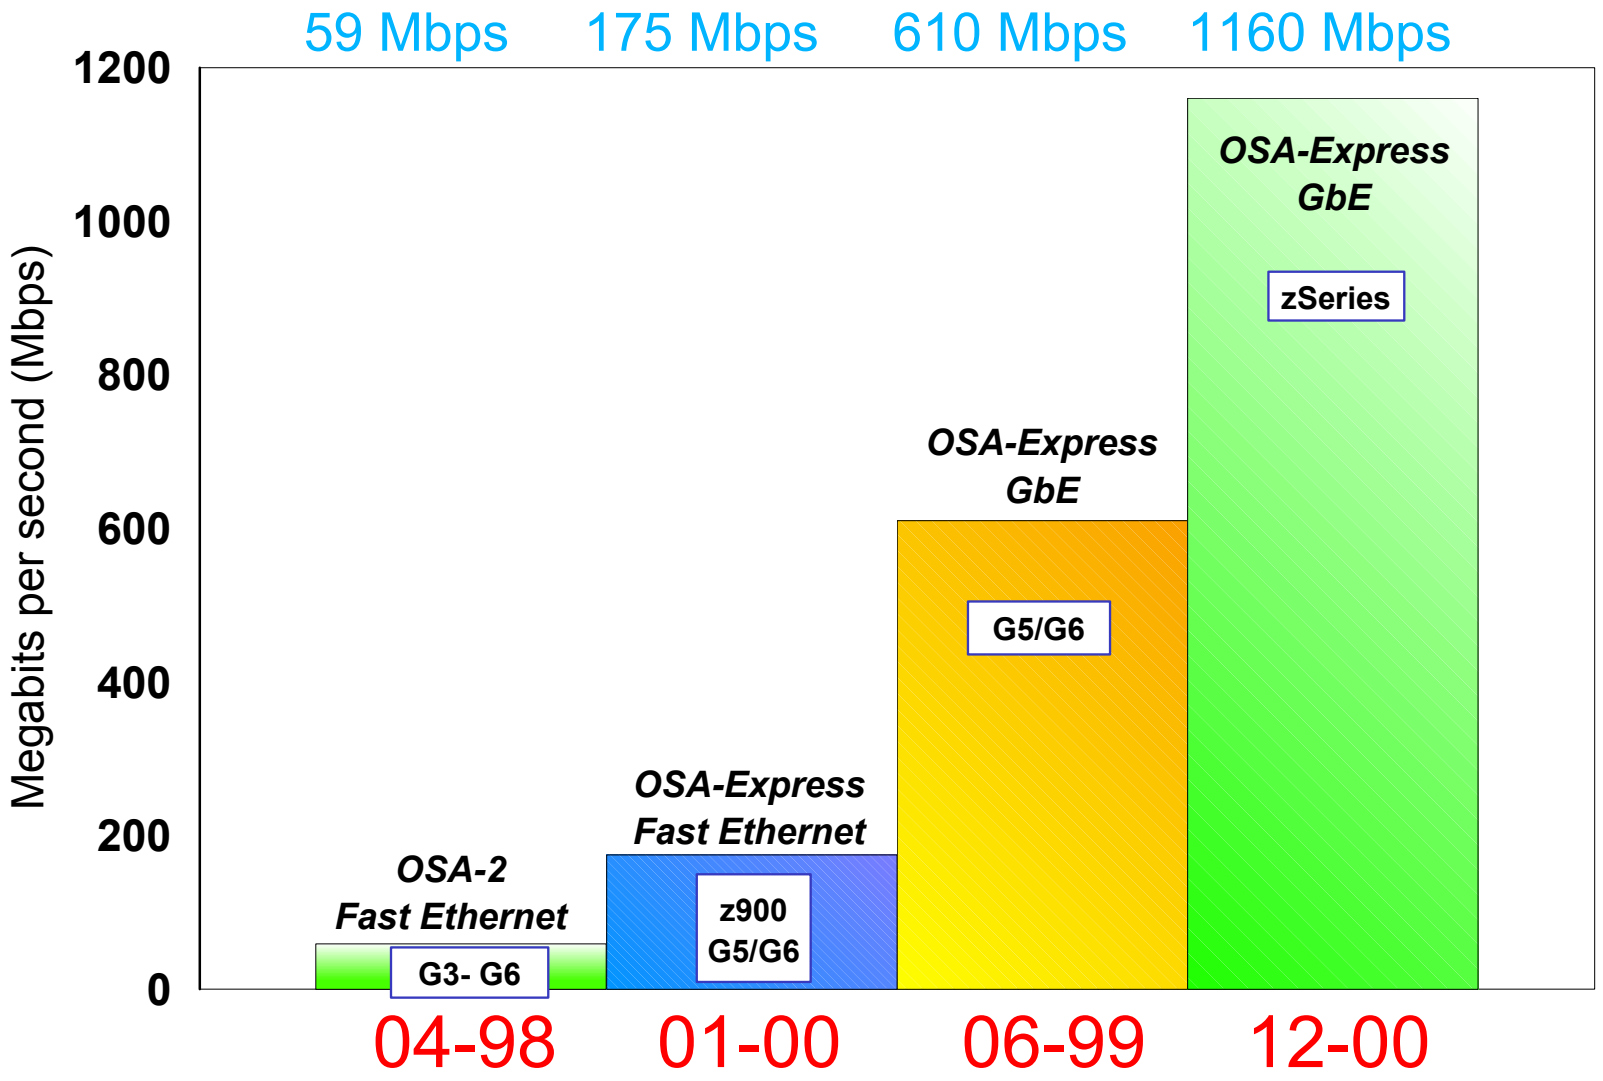

I/O Subsystem Bus

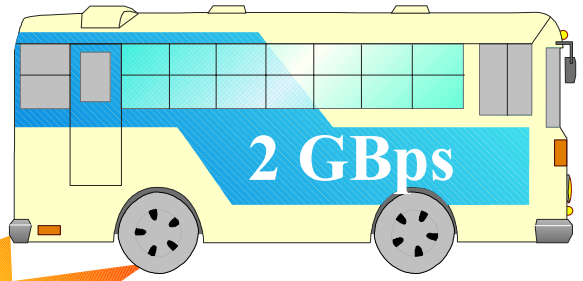

- **CRH bus = 20 MegaByte (Channel Request Handler)**
  - Used by I/O since 1990

Shared by up to 8 features

- **STI bus = 333 MegaByte (Self-Timed Interconnect)**
  - Used by FICON and OSA-Express beginning in 1999 on G5/G6 Servers

Shared by 4 features

- **STI bus = 1 GigaByte**
  - Introduced with the zSeries 900

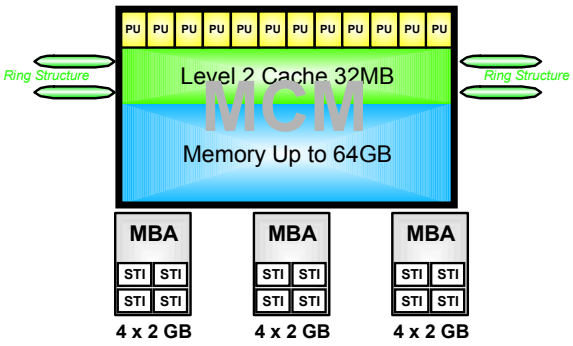

- **STI bus = 2 GigaByte**
  - Introduced with z990

Shared by 4 features

Plugging Rules: Total cannot exceed 20 features per I/O cage for:

FICON Express, OSA-Express, PCIXCC, PCICA

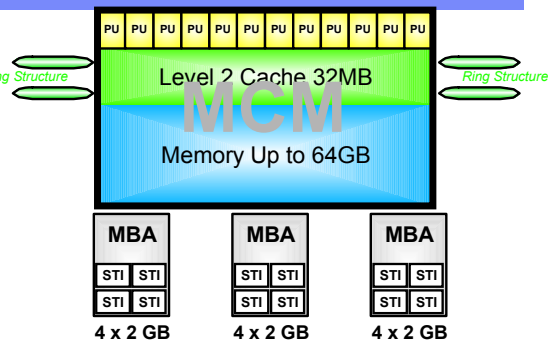

Model A08

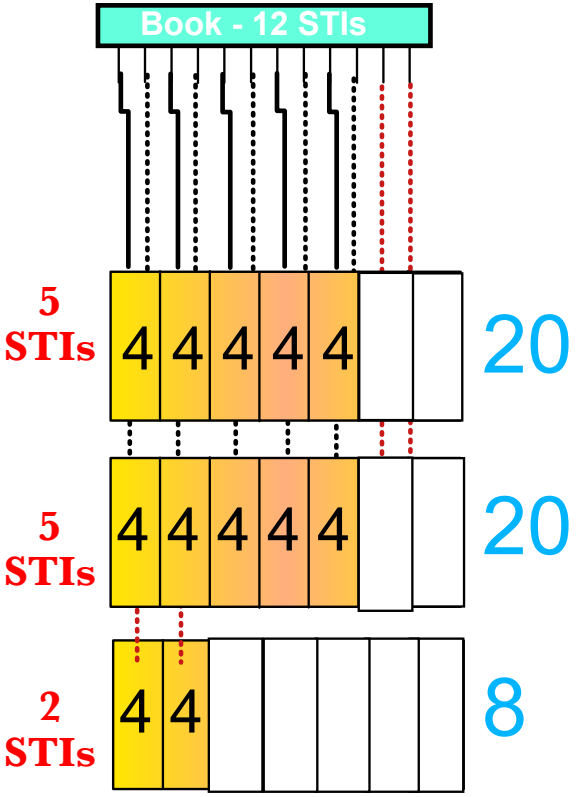

30 GBytes for I/O if desired

Up to 12 STIs for I/O

Models B16, C24, D32

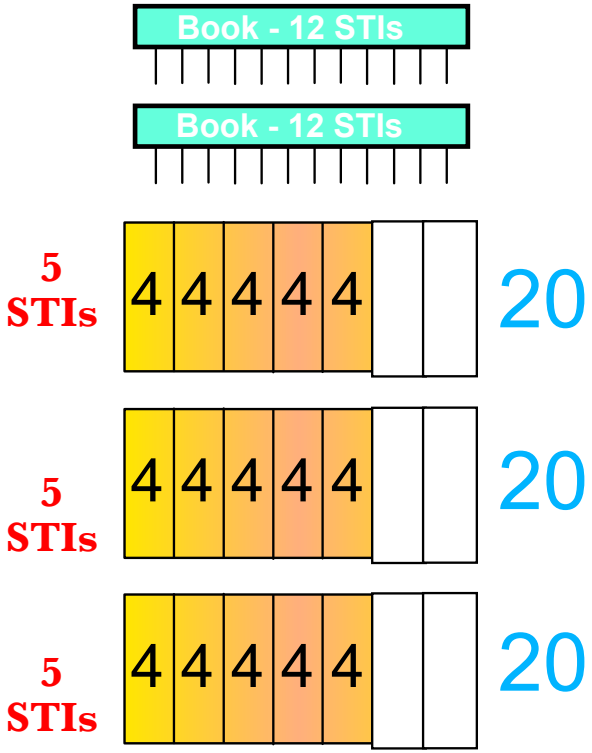

52.5 GBytes for I/O if desired

7 Domains in each I/O cage

Up to 21 STIs for I/O

Model A08
12 STIs
30 GBytes
Up to 48 features

Model B16
24 STIs
60 GBytes
Up to 84 features

Model C24
36 STIs
90 GBytes
Up to 84 features

Model D32
48 STIs
120 GBytes
Up to 84 features

Down to basics - One I/O cage

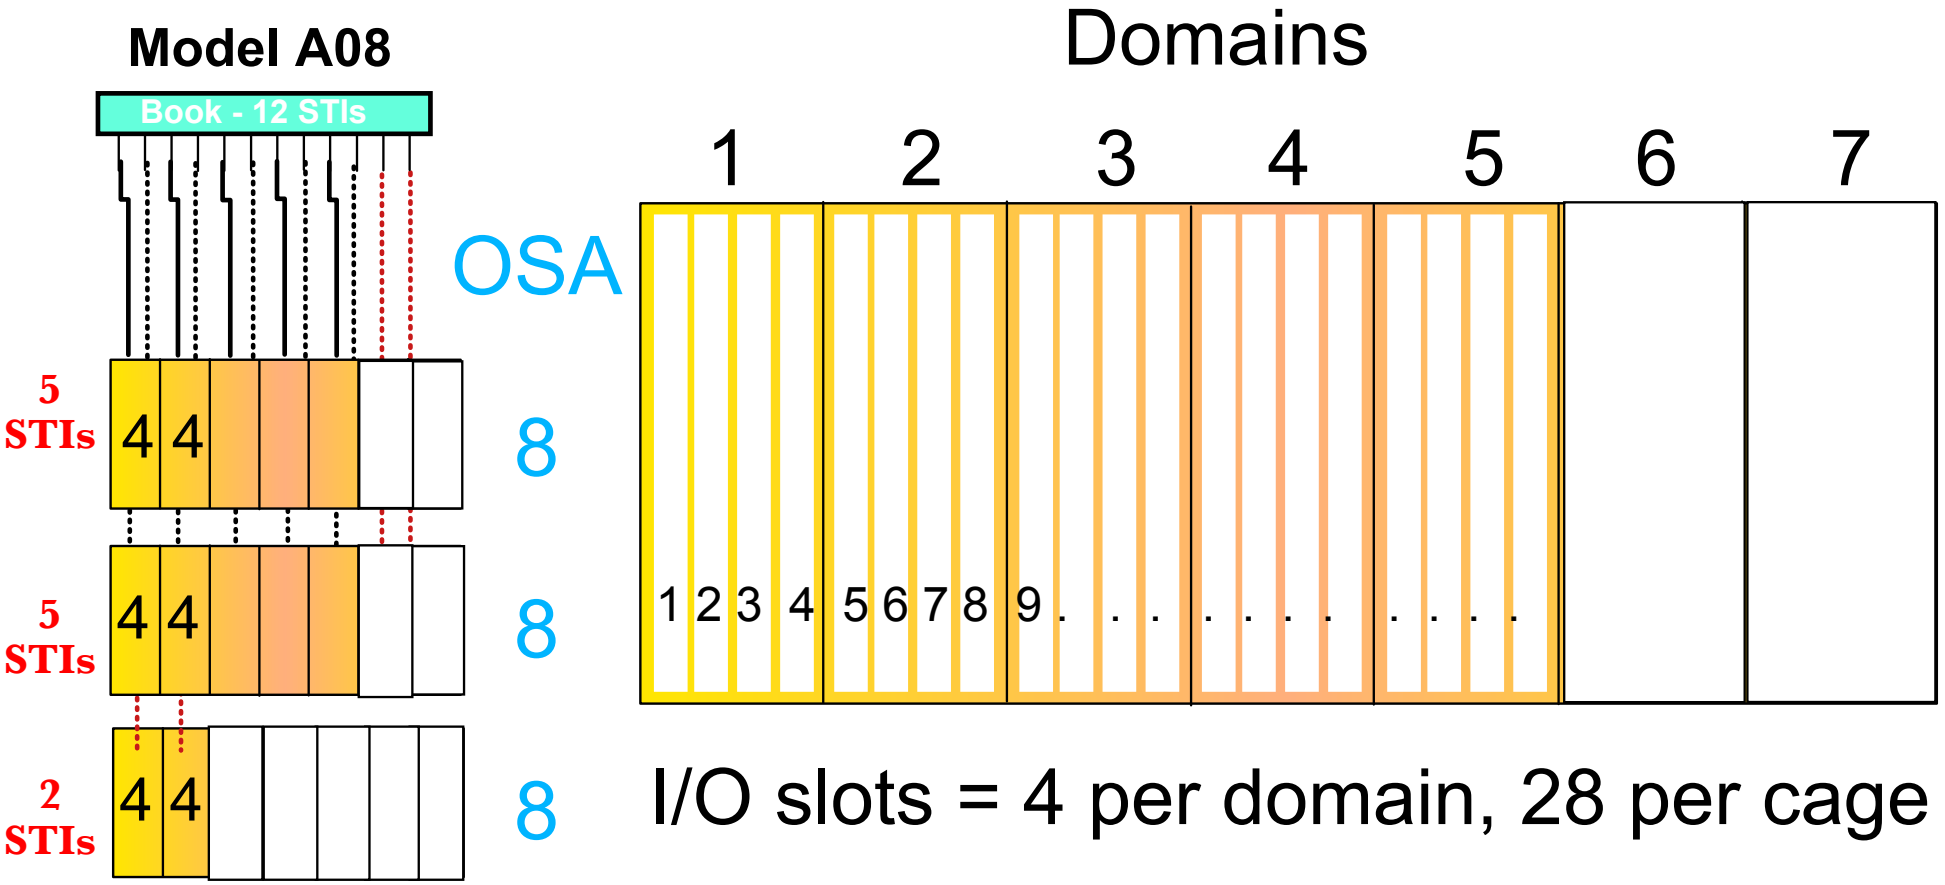

End-to-End Connectivity

Server → Feature → Transceiver → Connector → Fiber

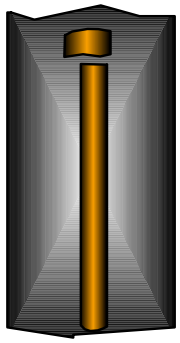

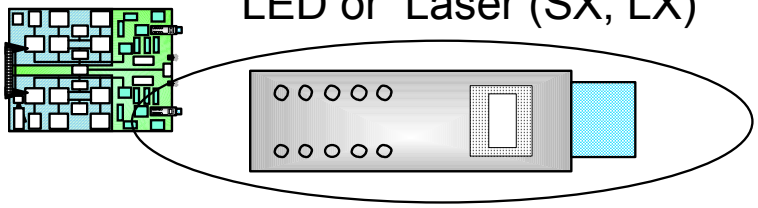

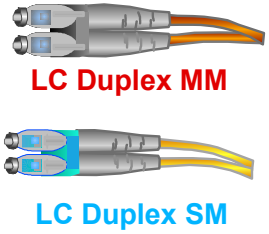

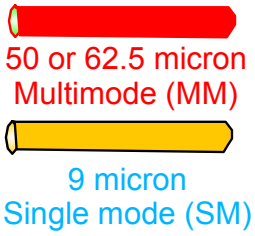

Multimode (MM) fiber

"Multiple paths" for the light to travel

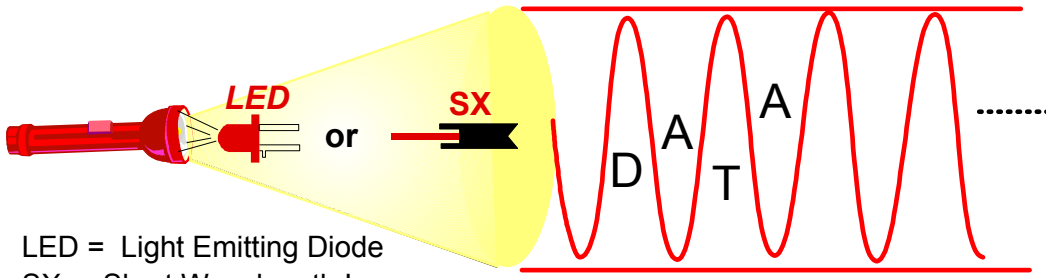

LED = Light Emitting Diode
SX = Short Wavelength Laser

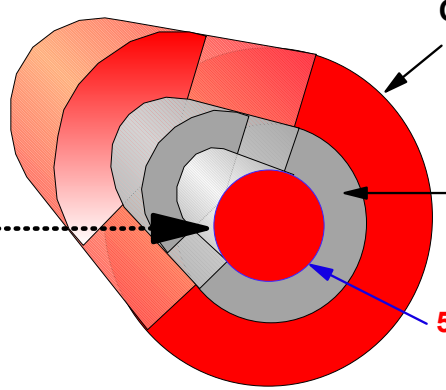

Single Mode (SM) Fiber

"Single path" for light to travel

LX = Long Wavelength Laser
Single mode support on 9032-005
is referred to as XDF - Extended Distance Feature

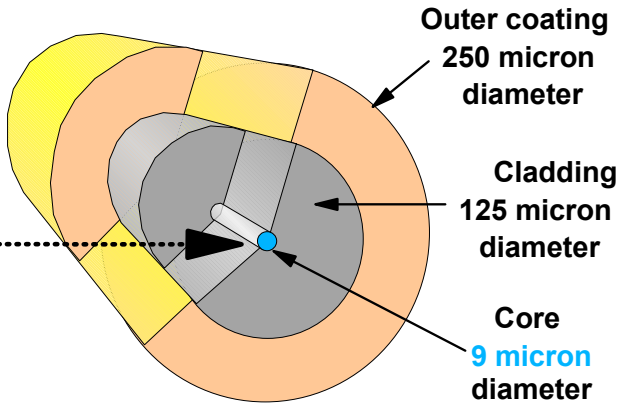

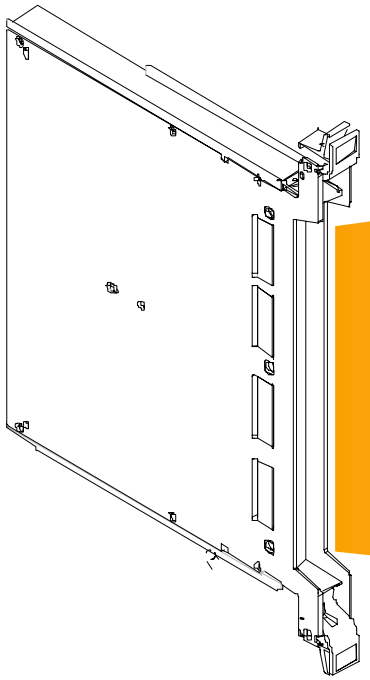

The Network Connections

z990 OSA-Express New Build Family

- New maximum of 24 features / 48 ports per system
- Each feature has two identical ports, each capable of achieving line speed *
- New Checksum offload - QDIO mode only (z/OS V1.5, Linux for zSeries)
  - GbE, 1000BASE-T Ethernet
- Jumbo frames: GbE, 1000BASE-T Ethernet
- New Gigabit Ethernet features
  - New connector type, LC Duplex
  - Gigabit Ethernet LX (Long wavelength)
    - 9 single mode **fiber**
  - Gigabit Ethernet SX (Short wavelength)
    - 50, 62.5 multimode **fiber**
- New 1000BASE-T Ethernet (10/100/1000 Mbps)
    - Same Category 5 **copper** as Fast Ethernet
- Token Ring (4/16/100 Mbps)
    - Category 5 **copper**

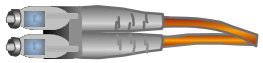

LC Duplex MM

LC Duplex SM

Modes of Operation
- QDIO for all ..
  - CHPID type OSD
  - TCP/IP only (TCP/IP traffic; TN3270 or Enterprise Extender for SNA traffic)
- Non-QDIO
  - CHPID type OSE - TCP/IP (LCS) and SNA/APPN/HPR (LSA)
    - 1000BASE-T Ethernet
    - Token Ring

Gigabit Ethernet Features 1364, 1365

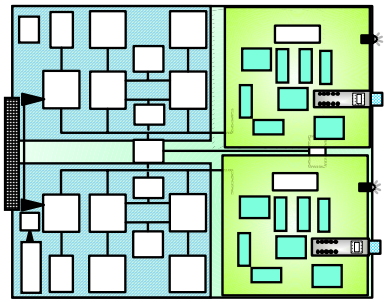

Ethernet (1000BASE-T) Feature 1366

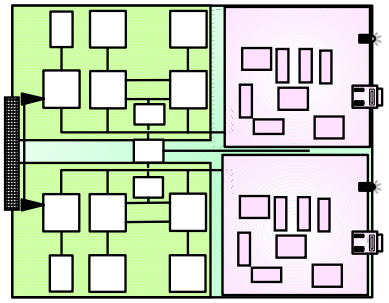

Token Ring Feature 2367

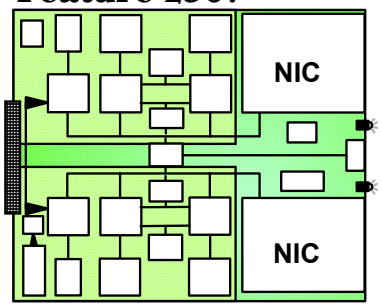

* Actual throughput is dependent upon customer environment

z990 OSA-Express 1000BASE-T Ethernet

Ethernet (1000BASE-T)
Feature 1366

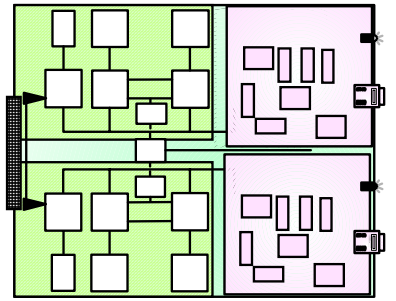

- **For z990 new builds**
- **Supports auto-negotiation: 10, 100, 1000 Mbps**
- **QDIO and a non-QDIO**
- **TCP/IP and SNA/APPN/HPR environments at up to gigabit speeds**
- **Checksum Offload when in QDIO mode (OSD CHPID type)**
  - Supported by z/OS V1.5, Linux

- **When configured at 1 Gbps**
  - Operates in full-duplex mode only
  - Operates in QDIO mode or non-QDIO mode
  - Can carry SNA/APPN/HPR traffic (non-QDIO mode)
  - Can carry TCP/IP packets (QDIO or non-QDIO mode)
  - Supports jumbo frames in QDIO mode

NIC = Network Interface Card

* Actual throughput is dependent upon customer environment

z990 OSA-Express Upgrade Family

- Each feature has two identical ports, each capable of achieving line speed *
- All can be carried forward on an upgrade from z900
- Gigabit Ethernet LX
  - 9 micron single mode **fiber**
  - Continues to use SC Duplex connector
- Gigabit Ethernet SX
  - 50 or 62.5 micron multimode **fiber**
  - Continues to use SC Duplex connector
- Fast Ethernet (10/100 Mbps)
  - Category 5 **copper**
- Token-Ring (4/16/100 Mbps)
  - Carry forward on an upgrade from z900
  - Category 5 **copper**
- Modes of Operation
  - QDIO for all ..
    - CHPID type OSD
    - TCP/IP only (TCP/IP traffic; TN3270 or Enterprise Extender for SNA traffic)
  - Non-QDIO
    - CHPID type OSE - TCP/IP (LCS) and SNA/APPN/HPR (LSA)
      - 1000BASE-T Ethernet
      - Token Ring

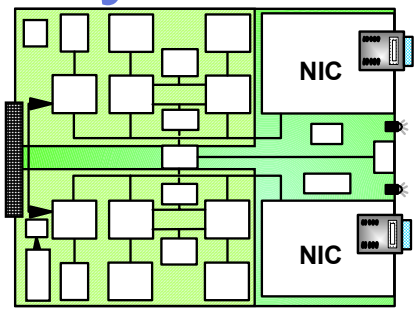

Gigabit Ethernet LX
Feature 2364

Gigabit Ethernet SX
Feature 2365

Fast Ethernet
Feature 2366

Token Ring
Feature 2367

NIC = Network Interface Card

* Actual throughput is dependent upon customer environment

Refresh of OSA-Express Technology

	1st Generation	2nd Generation	3rd Generation	3rd + Generation
	G5/G6	z800, z900	z800, z900	z990 New NIC
	GbE - 06/99 Fast E - 01/00 ATM - 01/00	Introduced on z900 12/00 z800 03/02	Introduced on z900 12/00 z800 03/02	
	GbE Fast Ethernet ATM	GbE Fast Ethernet ATM Token Ring	GbE	GbE, 1000BASE-T Ethernet (FC 1364, 1365, 1366)
Number of Ports	One port per feature	Two ports per feature	Two ports per feature	Two ports per feature
Microprocessor	166 MHz	333 MHz	333 MHz	333 MHz
PCI Bus	32-bit 33 MHz	32-bit 33 MHz	64-bit 66 MHz	64-bit 66 MHz
Maximum Features/Ports	12 / 12	12 / 24	12 / 24	24 / 48

OSA New Builds and Upgrades

Features on z900 (12/00) and z800 (03/02)	z990 Upgrade	z990 New Build
OSA-Express	***OSA-Express***	***OSA-Express***
All carried forward from z900 - - - >		
2367 Token Ring	- - > Carried forward	- - >
2366 Fast Ethernet	- - > Carried forward	**1366** 1000BASE-T Ethernet
2365 Gigabit Ethernet SX	- - > Carried forward	**1365** Gigabit Ethernet SX
2364 Gigabit Ethernet LX	- - > Carried forward	**1364** Gigabit Ethernet LX
2363 155 ATM MM	Offered as RPQ 8P2258	Not offered
2362 155 ATM SM	Offered as RPQ 8P2258	Not offered
OSA-2	**OSA-2**	**OSA-2**
5201 Converted to either **2366** or **2367**	**5201** Converted to either **2367** or **1366**	**5201** Converted to either **2367** or **1366**
5202 FDDI (not offered on z800)	Not offered	Not offered

z990 New GbE Features - Connector Change

May require a Conversion Kit to attach to switches

Feature	New Connector	Previous Connector
OSA Express GbE LX Single mode (SM) fiber	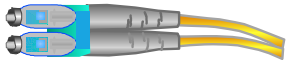 **LC Duplex SM**	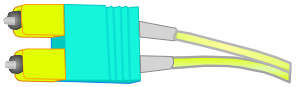 **SC Duplex SM**
OSA Express GbE SX Multimode (MM) fiber	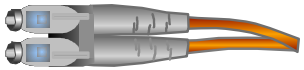 **LC Duplex MM**	 **SC Duplex MM**

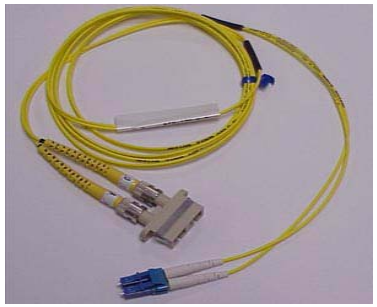

Conversion Kit example
2 meters (6.5 feet)

Mode Conditioning Patch Cables

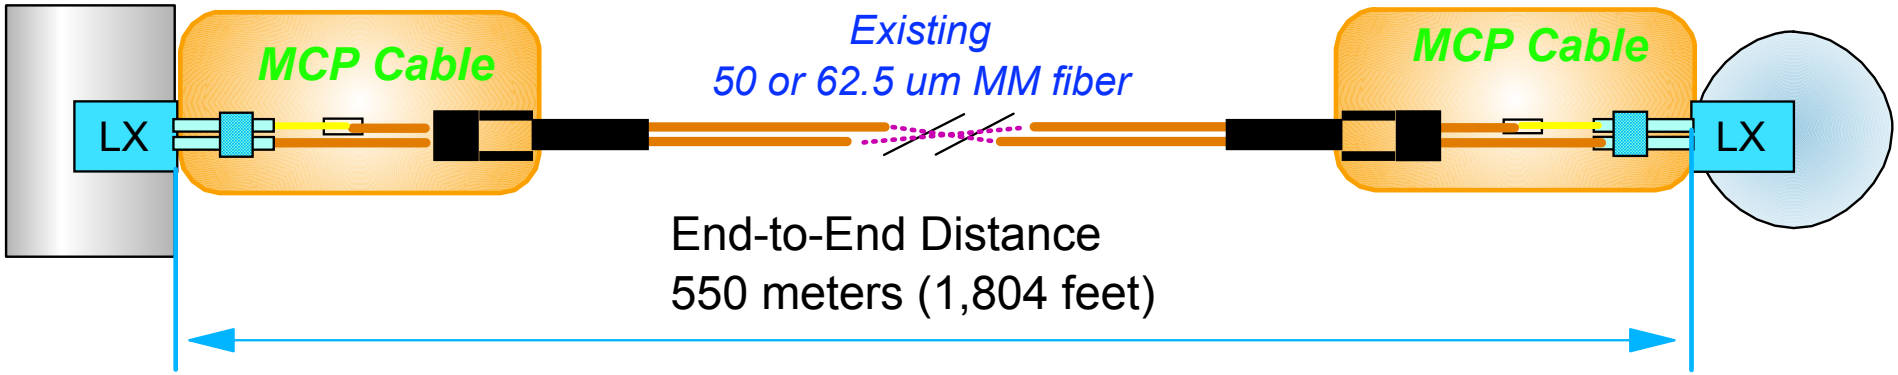

MCP Cable example
2 meters (6.5 feet)

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

Gigabit Ethernet Cabling Options

LX = Long wavelength 1300 nm transceiver SX - Short wavelength 850 nm transceiver

- LX transceiver/feature on each end
- Requires 9 micron single mode (SM) fiber

1 Gbps	4.6 dB	5 km (3.1 miles)

- 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

1 Gbps	2.4 dB	550 meters (1804 feet)

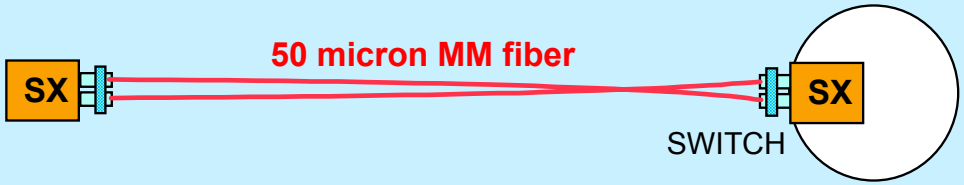

- SX transceiver/feature on each end
- Requires 50 micron multimode fiber
- **Reduced distance** and link budget

1 Gbps	3.6 dB	550 meters (1804 feet)

- SX transceiver/feature on each end
- Requires 62.5 micron multimode fiber
- **Reduced distance** and link budget

1 Gbps	2.6 dB	220 meters (722 feet)

Link Rate ***Link Budget*** ***Unrepeated Distance***

Two CHPID Types Control Operation

Mutually exclusive

- CHPID type of OSE: native SNA/APPN/HPR and TCP/IP traffic can flow
- CHPID type of OSD: TCP/IP traffic can flow; use Enterprise Extender or TN3270 for SNA applications

CHPID Type	MEDIA	SNA/APPN/HPR Traffic	TCP/IP Traffic	OSA/SF Required
OSE (non-QDIO)	1000BASE-T Ethernet Fast Ethernet	YES	YES	YES
OSE (non-QDIO)	Token Ring	YES	YES	YES
OSD (QDIO)	Gigabit Ethernet	NO (Use EE or TN3270)	YES	NO
OSD (QDIO)	1000BASE-T Ethernet Fast Ethernet	NO (Use EE or TN3270)	YES	NO
OSD (QDIO)	Token Ring	NO (Use EE or TN3270)	YES	NO

Logical View: Interconnect Controller

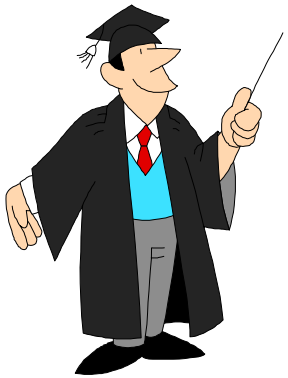

- Used for all LAN communications
  - ✓ ***Start I/O's***
  - ✓ ***Channel Control Words (CCWs)***
  - ✓ ***I/O interrupts***
- Involved in all data transfers
  - ✓ ***I/O Processor (IOP)***
  - ✓ ***Channel***
- Control unit layer
  - ✓ ***Interprets CCWs (Read/Write)***
  - ✓ ***Blocks/deblocks data***
- Device runs the LAN driver

Controller, Router, OSA-2, OSA-Express Non-QDIO

Queued Direct Input/Output (QDIO)

- For TCP/IP traffic only
  - For SNA/APPN/HPR traffic with QDIO use TN3270, Enterprise Extender
- Design for high speed communication
  - ***Reduced TCP/IP path length***
    - Internal testing reflects host utilization drop of 15% for FTP with 1500 byte datagrams

1. ***QDIO IP Processing Assist (see next slide)***
2. ***LPAR-to-LPAR Communication with port sharing***
3. ***Direct Memory Access (DMA) Protocol***
   - Memory-to-memory communication
     - I/O interrupts minimized
     - Continuous direct data exchanges
4. ***Dynamic customization***

- z/OS, OS/390, z/VM, VM/ESA, VSE/ESA, LINUX, TPF
- Gigabit Ethernet, 1000BASE-T Ethernet, Fast Ethernet, Token Ring

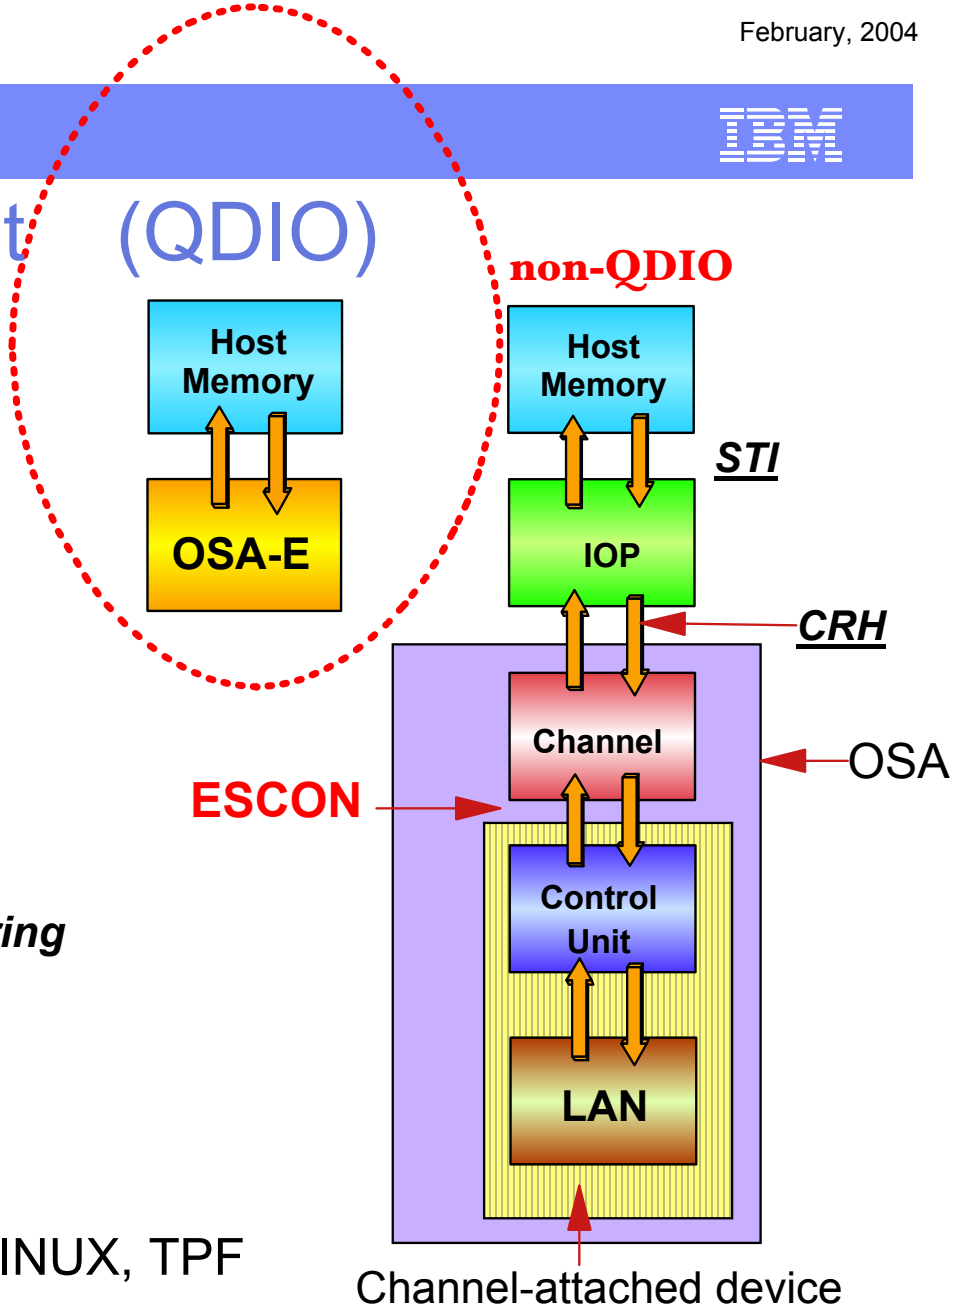

z990 OSA-Express New Functions - QDIO only

z/OS V1.5, Linux for zSeries

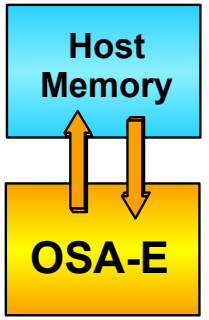

- **QDIO IP Processing Assists**
  - Performs all ARP processing
  - Provides Multicast support
  - Builds MAC and LLC headers
  - Performs filtering - TCP/IP stack sees only IP datagrams
  - New Checksum offload for IPv4 packets
    - Calculates the TCP/UDP and IP header checksums
      - Verifies the correctness of files
    - Reduces host CPU cycles
      - Checksum offload support is offered for z/OS V1.5, Linux for zSeries on:
        - GbE features (new features 1364, 1365)
        - 1000BASE-T Ethernet (new feature 1366)
  - All the OSA-Express features support z/OS Communications Server Intrusion Detection Services (IDS). IDS helps detect attacks on the TCP/IP stack that can potentially harm its ability to function normally and cause a misuse of system resources. This was introduced in z/OS V1.2.
    - Enhancements are introduced in z/OS V1.5

Connection is Dynamically Defined (OSD)

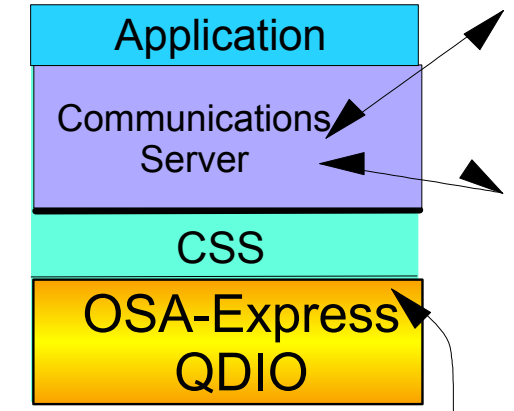

TCP/IP Profile

```
DEVICE OSAPORT1 MPCIPA NONRouter
LINK OSALINK1 1PAQGNET OSAPORT1
HOME OSALINK1 192.60.41.1
```

TRL Macro Definition - Transport Resource List Element (TRLE)

```
OS11 TRLE LNCTL=MPC, READ=(100),
WRITE=(101), DATAPATH=(102),
PORTNAME=OSAPORT1, MPCLEVEL=QDIO
```

IOCP statements (can be done using HCD)

```
CHPID PATH=C8, TYPE=OSD
CNTLUNIT CUNUMBR=100, UNIT=OSA, PATH=C8
IODEVICE ADDRESS=(100,3), CUNUMBR=100, UNIT=OSA
IODEVICE ADDRESS=(1FE), CUNUMBER=100, UNIT=OSAD,
         UNITADD=FE
```

Client

OSA-Express

OSA Address Table (**built dynamically; data received from host**)
- LP=0, UA=00, Type=MPC(QDIO)-control,OSAName=OSA11, Port=0
- LP=0, UA=02, Type=MPC(QDIO)-data,OSAName=OSA11, Port=0

IP Address Table
- IP Address=192.60.40.1, Def=None, OAT entry=(LP=0,UA=02)

Port Parm
- Portname=OSAPORT1

Virtual IP Address – VIPA

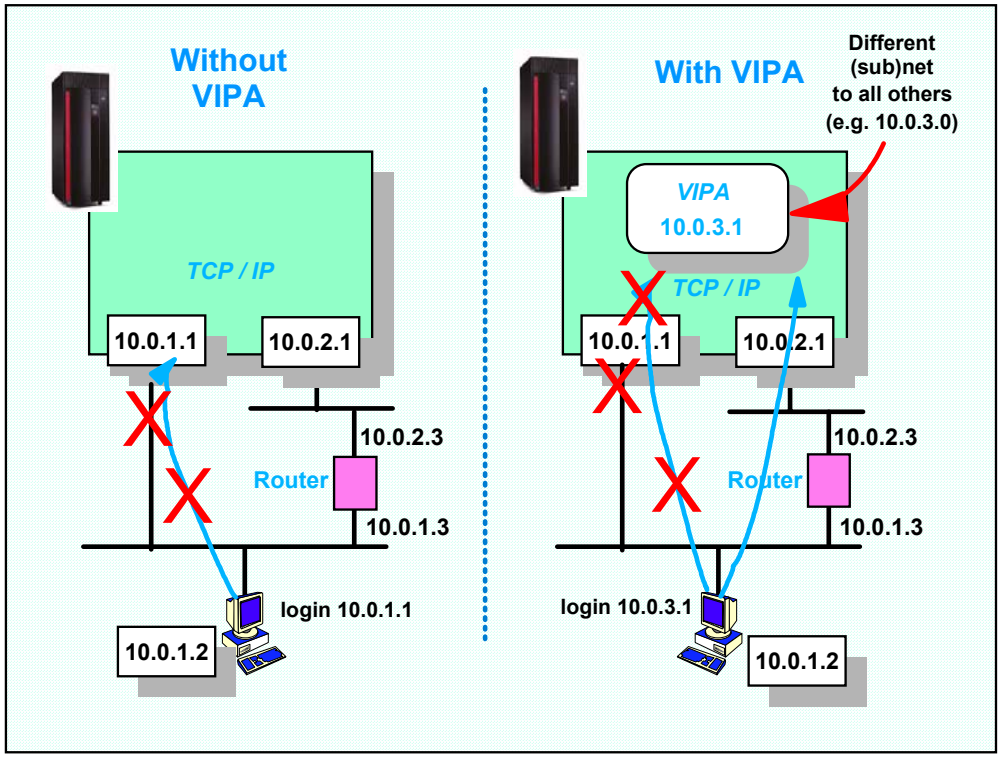

Use:
RIP (Routing Information Protocol) or
OSPF (Open Shortest Path First)

- Frees hosts from dependence on particular network attachments
  - Provides IP address of TCP/IP stack without a specific network attachment
- With VIPA, OMPROUTE provides automatic and transparent recovery from
  - Interface failures
  - Network failures
- VIPA can represent stack or specific TCP application
  - Multiple VIPAs can be defined for a TCP/IP stack
- Other hosts can use VIPA to connect to TCP/IP applications

Service Policy Server in CS for OS/390 V2.10 >

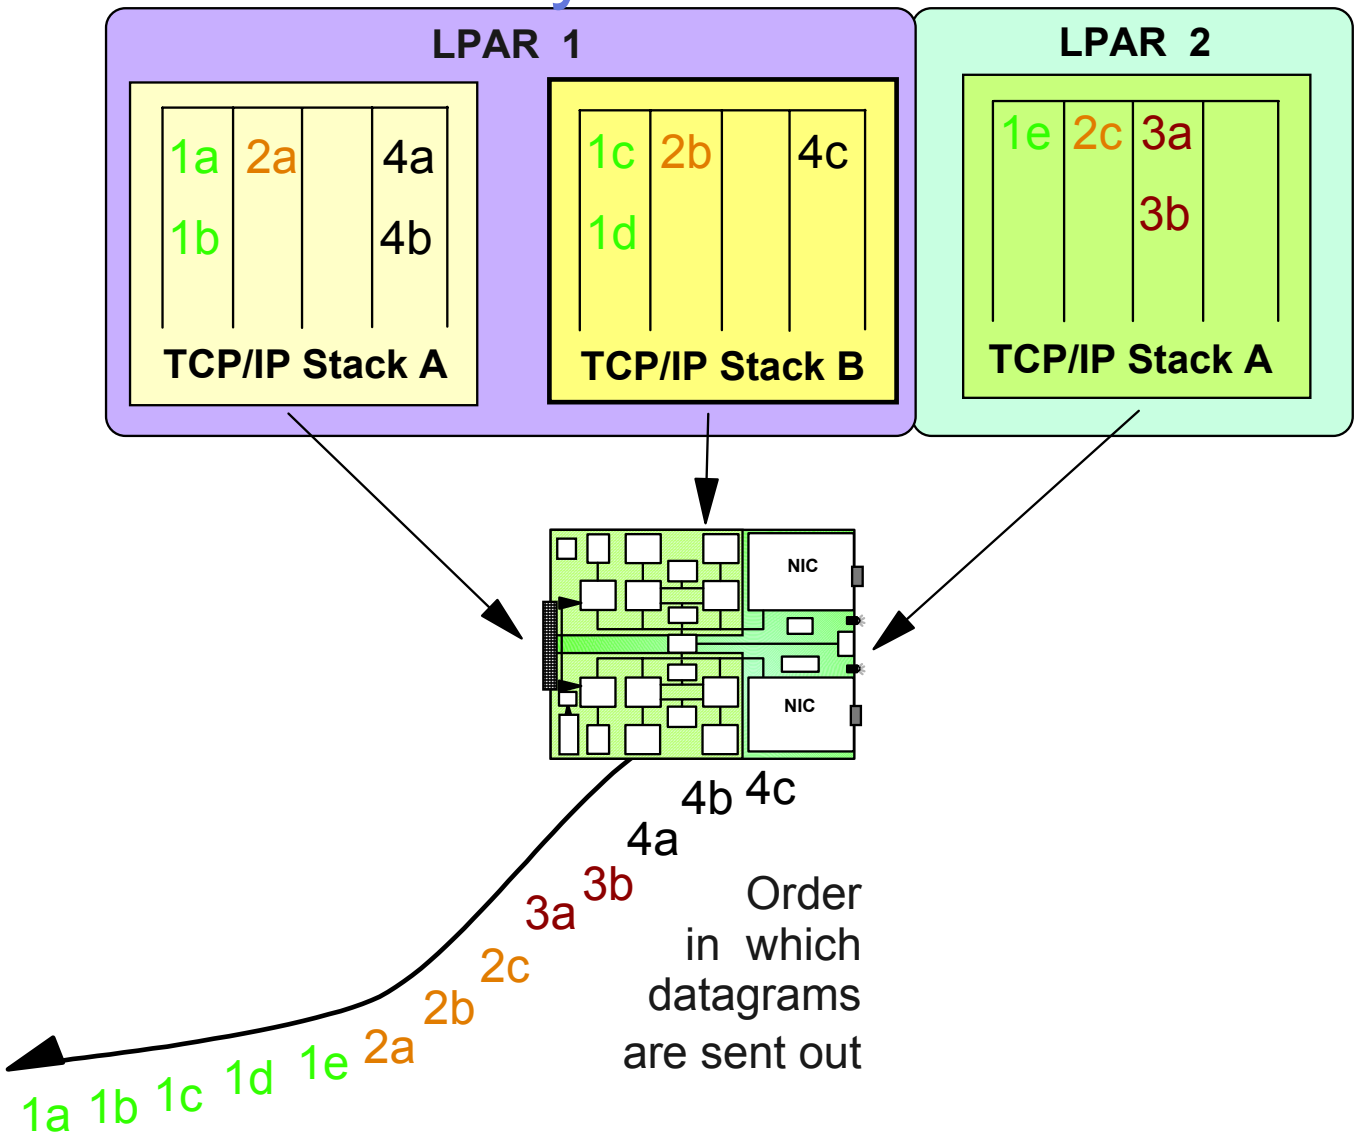

QDIO (OSD)

Extends Business Priority Across the Network

4 Priority Queues per TCP/IP Stack

Queues correspond to TOS priorities

- Service?
  - ***Premium***
  - ***Shared***
  - ***Best Effort***

OSA-Express Functions - QDIO only

Virtual Local Area Network Identifications (VLAN IDs) - IEEE 802.1q

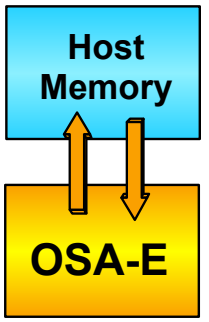

- **Full VLAN support for z/OS V1.5, z800, z900, z990**
  - Now z/OS supports VLAN priority tagging as well as VLAN IDs.
    - One global VLAN ID per IP version per IP stack (IPv4 and IPv6)
- **VLAN support for z/VM V4.4, z990**
  - One global VLAN ID for IPv4
    - One VLAN ID per OSA-Express port
    - Each port can be configured with a different VLAN ID
- **VLAN support for Linux on zSeries support, z800, z900, z990**
  - Delivered as part of the "May, 2003, stream" at:
    http://www10.software.ibm.com/developerworks/opensource/linux390
  - VLAN support is integrated in SUSE SLES 8, Turbolinux TLES8, Conectiva CLEE.

- On October 1, 2003 the following MCL was released for z800 and z900 in support of VLANs:
  Driver 3G - EC stream J11204, MCL032 (OSA level 3.33)

z/OS V1.5 - Full VLAN support

- Organize traffic by logical patterns, independent of location, over common physical network
  - Save networking equipment and bandwidth, minimize impact of broadcast traffic
  - Increase security among user groups on shared network and shared OSA-Epxress
- Full VLAN (IEEE 802.1Q) for OSA-Express Ethernet QDIO
  - Extends z/OS 1.2 VLAN priority tagged frames to now include non-null VLAN ID, called full VLAN
    - VLAN ID configured via TCP/IP
    - Defines additional information in the MAC frame (tag carries VLAN ID)
  - One Global VLAN ID each for IPv4 and IPv6 per TCP/IP stack
  - Multiple TCP/IP stacks, each with IPv4 and IPv6 VLANs, over single OSA-Express port
  - LPAR-to-LPAR Communication supported

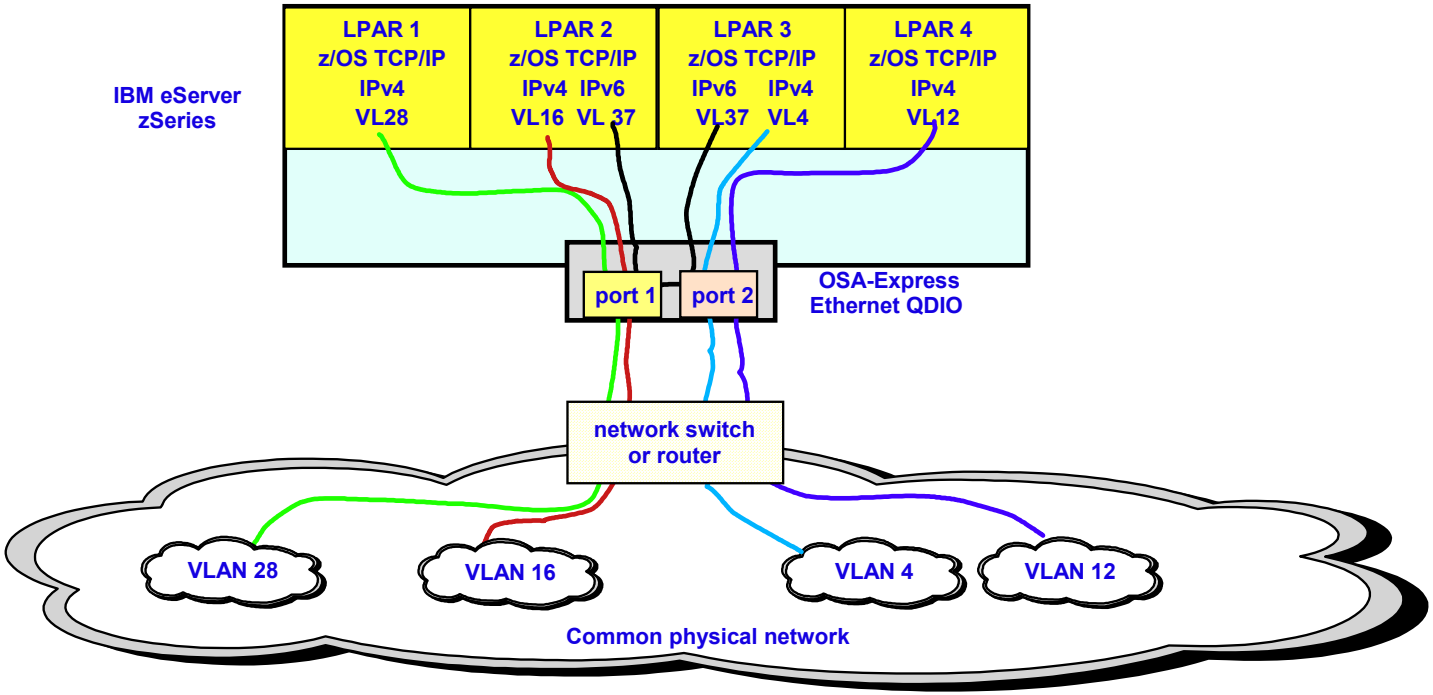

z/VM V4.4 - VLAN Support

- z/VM may now participate in an IPv4 VLAN (IEEE 802.1q)
  - Logically separates group traffic patterns independent of location across a common physical network
  - Can lower broadcast traffic and provide secure traffic flows among groups of users
- z/VM TCP/IP stack over HiperSockets and OSA-Express Ethernet in QDIO mode
  - Supports one VLAN ID per OSA-Express connection (LINK)
  - IP stack can have connections (LINKS) to multiple OSA-Express ports
  - Each connection (LINK) could be configured with a different VLAN ID
- Guest LAN support for guests using either QDIO or HiperSockets virtual adapter connections
- Virtual Switch connection between guest and OSA-Express connected external network
- z/VM Virtual Switch provides dynamic, centralized network configuration and control of VLANs

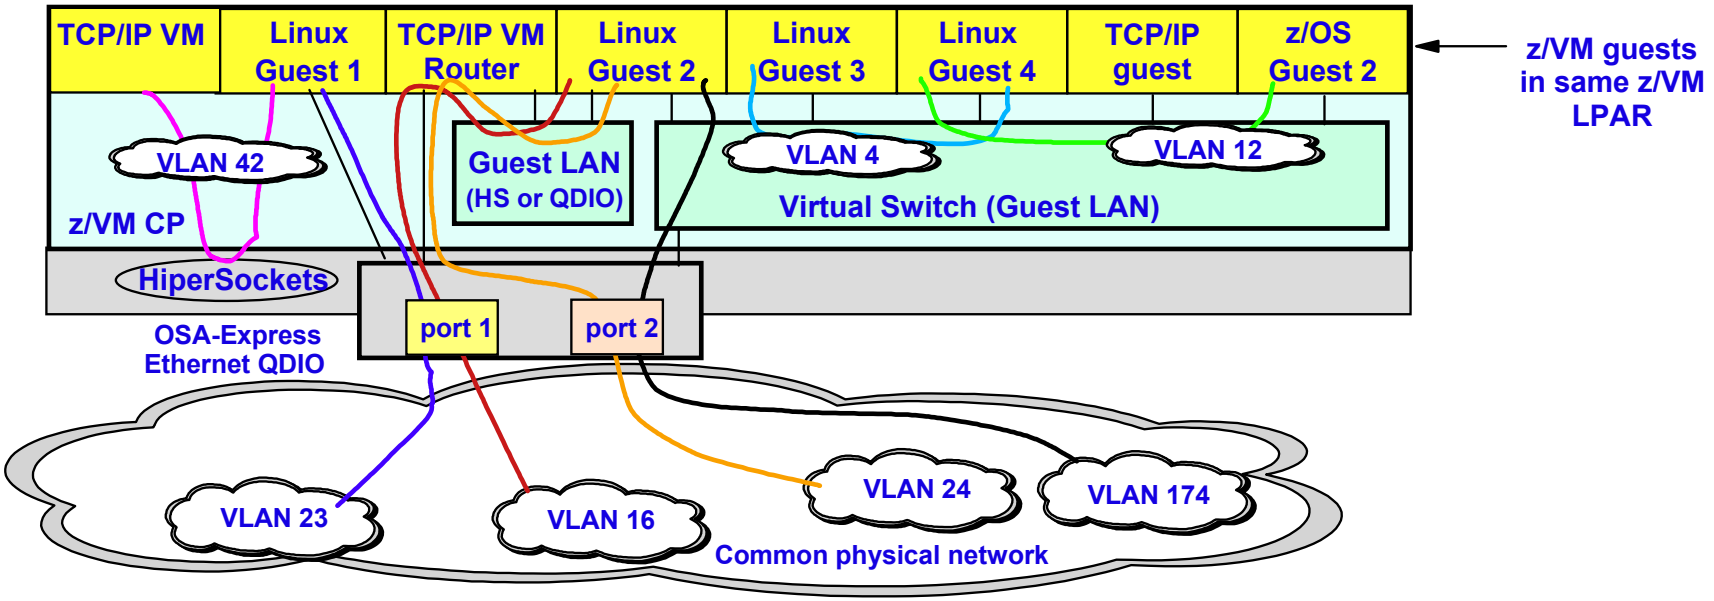

z990 Performance Assists - QDIO only

- **Adapter Interruptions - z990 + z/VM V4.4**
  - OSA-Express (OSD CHPID type only)
  - Extension to QDIO architecture
  - New technique for I/O interruptions
    - Reduces overhead to process an I/O interruption designed to enhance Queued-Direct-I/O (QDIO) Performance
      - Operating system and OSA-Express

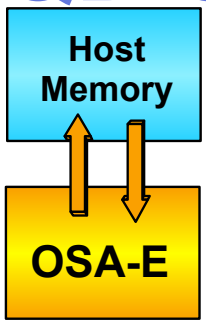

Support for adapter interruptions is designed to improve performance of VSE/ESA V2.7 guest operating systems using TCP/IP for VSE/ESA when VSE/ESA V2.7 is running as guest system under z/VM V4.4.

- **Performance Assist for V=V guests - z990 + z/VM V4.4**
  - OSA-Express (OSD CHPID type only)
  - Direct presentation of adapter interruptions by server to pageable guests
  - No need to stop guest processing

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/

z990 OSA-Express - QDIO only

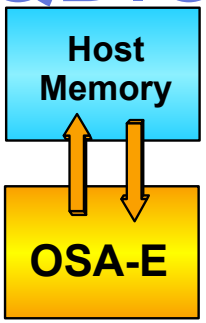

- **Increased TCP/IP connectivity (transparent to operating systems)**
  - All OSA-Express features in QDIO mode (OSD CHPID Type only)
  - Now allow up to 160 IP stacks per OSA port and 480 devices
    - Previously was 80 IP stacks and 240 devices
  - Maximum of 84 IP stacks per Logical Partition (LPAR) and 255 devices per LPAR

OSA Support Facility

- **Open Systems Adapter Support Facility (OSA/SF) - new version**
  - Complete replacement for OSA/SF V2.1 (5655-B57)
  - Applicable to all OSA-Express and OSA-2 features on all supported servers
- **New OSA/SF Java-based Graphical User Interface (GUI)**
- **Operates wherever the Java 1.4 runtimes are available**
  - GUI is supported via TCP/IP only
    - REXX command interface still available
  - Interoperability testing performed for
    - Windows 2000, Windows XP, Linux for zSeries
  - OSA/SF not required when using QDIO mode

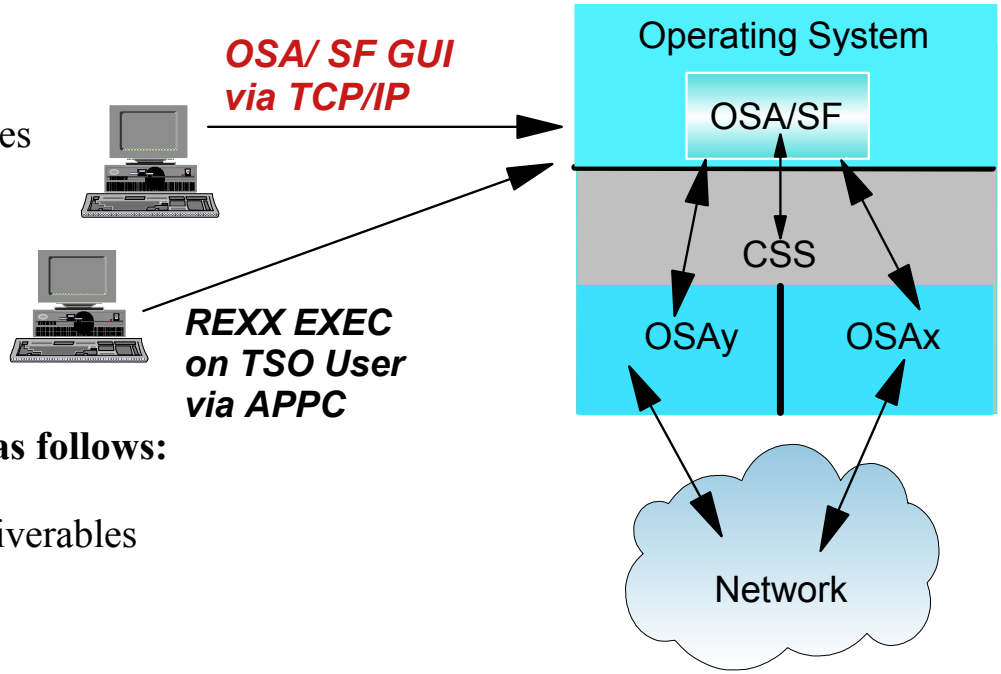

- **Integrated in the operating systems and available as follows:**
  - z/OS V1.4 z990 Compatibility Support feature
  - z990 Compatibility for Selected Releases web deliverables
    - OS/390 V2.10, z/OS V1.2, z/OS V1.3
    - ~~New version can coexist with OSA/SF V2.1~~
  - Integrated in z/VM V4.4
    - Overlays OSA/SF V2.1
  - z/VM V3.1, z/VM V4.2, z/VM V4.3 delivered as a PTF
    - Overlays OSA/SF V2.1
  - VSE/ESA V2.6 delivered as a PTF
    - Overlays OSA/SF V2.1

Direct SNMP Subagent - QDIO only

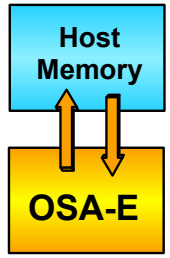

- **Direct SNMP subagent support is applicable to all of the OSA-Express features when configured in QDIO mode (CHPID type OSD).**
- **Query function supported using the SNMP "get" command, on z800, z900, z990**
  - First introduced April, 2002
  - Minimum software requirements:
    - z/OS V1.4 with z/OS CS, APAR PQ73478,
    - Linux on Series support: SuSE SLES 8, Turbo Linux TLES, and Conectiva CLEE.
- **New on z990 (May, 2003), z800, z900 (October, 2003)**
  - Simple Network Management Protocol (SNMP) subagent supports dot3StatsTable
    - From the SNMP EtherLike Management Information Base (MIB) module in RFC 2665 for all of the Ethernet features supported on z800, z900, and z990
  - Now, performance data for each image, reflecting OSA-Express utilization, is returned in individual objects without the need to "decode" a larger object that previously contained all of the information. The information returned includes the percentage of time the microprocessor was utilized to transfer data, as well as the number of inbound and outbound packets
- **Direct SNMP subagent MIB changes and dot3StatsTable**
  - Minimum software requirement s:
    - z/OS V1.4 and z/OS CS, APAR PQ73478 is required
    - Linux on zSeries support was delivered as part of the "June 2003 stream" at: http://www10.software.ibm.com/developerworks/opensource/linux390
      Note that earlier Linux on zSeries OSA SNMP support (from the "May 2002 stream", integrated in SuSE SLES 8, Turbo Linux TLES8, or Conectiva CLEE) is not supported with the updated MCLs.

Direct SNMP Subagent

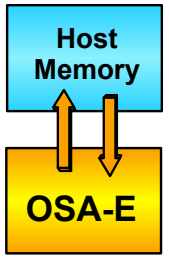

- **Operating Systems notes - prior to z800, z900, z990 MCLs:**
  - z/OS: Support is offered in z/OS V1.4 with z/OS CS, APAR PQ73478.
  - Linux on zSeries: Support was delivered as part of the "May 2002 stream", and is integrated in SuSE SLES 8, Turbo Linux TLES8, and Conectiva CLEE.
- **Operating Systems notes - with the z800, z900, z990 MCLs:**
  - General: Operating systems without OSA SNMP support continue to run with this MCL as before.
  - z/OS: Support is offered in z/OS V1.4 with z/OS CS, APAR PQ73478 (unchanged from prior MCL).
  - Linux on zSeries: Support was delivered as part of the "June 2003 stream" at:
    http://www10.software.ibm.com/developerworks/opensource/linux390
    - WARNING: If the MCL is installed and the co-requisite "June 2003 stream" is not present, Direct SNMP subagent support is no longer functional in a Linux environment. The earlier Linux on zSeries OSA SNMP support (from the "May 2002 stream", integrated in SuSE SLES 8, Turbo Linux TLES8, and Conectiva CLEE) does not support this change

- On October 1, 2003 the following MCL was released for z800 and z900 in support of the Direct SNMP subagent MIB changes and the dot3StatsTable:
  Driver 3G - EC stream J11204, MCL032 (OSA level 3.33).

OSA-Express Direct SNMP with z/OS V1.4

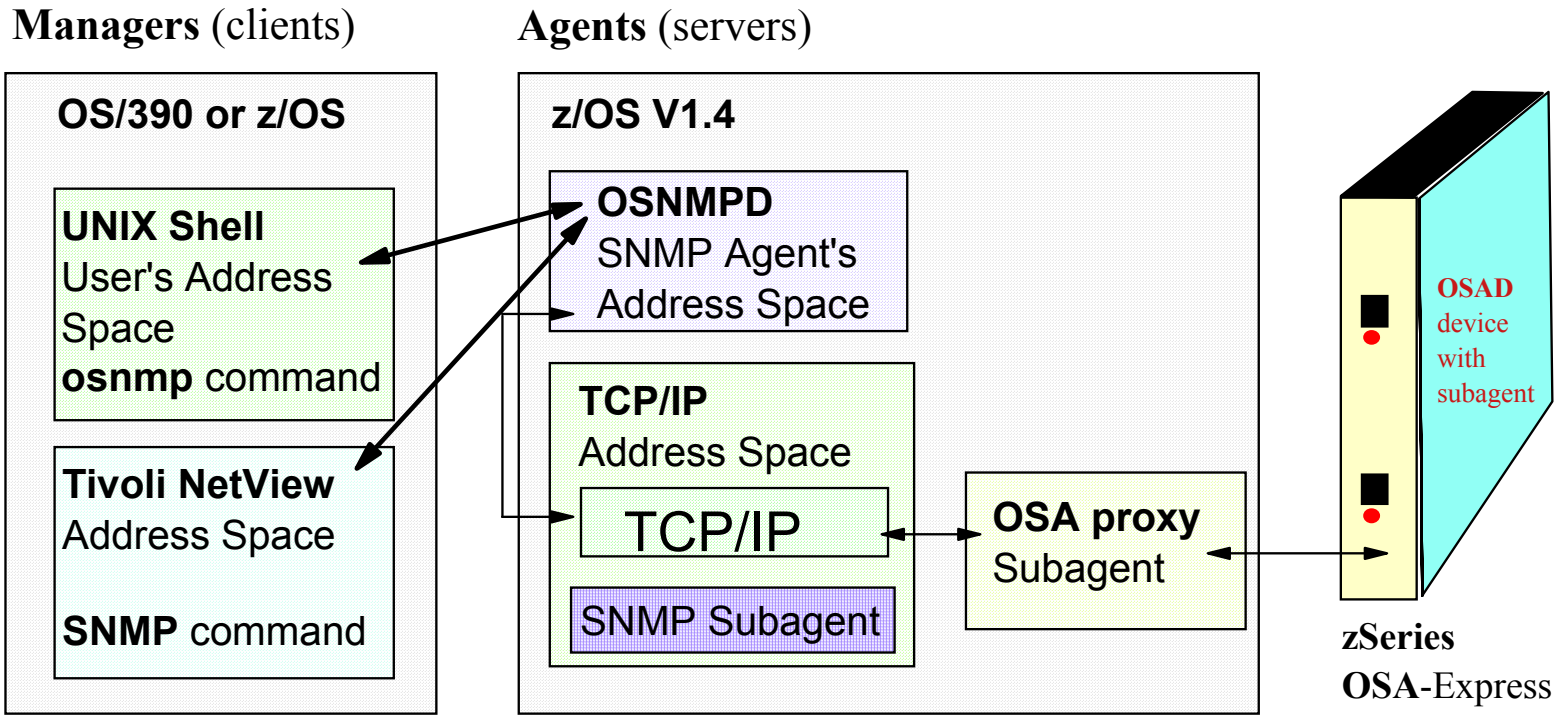

- z/OS CS can request network management information from OSA/SF
  - Then requests the information from OSA-Express
- z/OS V1.4 introduces a new structure
  - SNMP subagent on the OSA-Express feature
  - Proxy subagent which communicates with the subagent via a special z/OS CS-provided socket interface
  - New OSA-Express Enterprise Specific MIB on the adapter

Direct SNMP Support - Updated module

- **Updated Direct SNMP MIB module supports z800, z900, z990 as well as:**
  - New Gigabit Ethernet features
  - New 1000BASE-T Ethernet feature
  - Logical Channel Subsystems (LCSS's)
  - Updated performance table with more detailed information
- **OSA-Express 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

PORTNAME relief - QDIO only

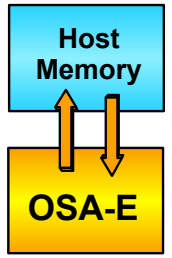

- PORTNAME relief, z800, z900, z990
- Specification of a PORTNAME is no longer required as part of the operating system configuration and can be omitted if not needed by the operating system.
- For compatibility purposes, OSA-Express will now permit activation with or without a PORTNAME.
- Minimum software requirements:
  - z/VM V4.3 with APARs VM63308 (PORTNAME relief for Guest LAN) and PQ73878 (VM TCP/IP exploitation for the PORTNAME relief).
    - PORTNAME relief is in the z/VM V4.4 base.
  - Linux on zSeries support was delivered as part of the "June 2003 stream" at: http://www10.software.ibm.com/developerworks/opensource/linux390
    - Existing Linux distributions and z/VM releases without support for PORTNAME relief continue to run on this MCL as before.
    - Users should ensure that OSA PORTNAME relief is available and installed on their system before activating the interface without specifying a PORTNAME
  - z/OS Communications Server (z/OS CS) still requires the use of the OSA-Express PORTNAME. The PORTNAME must be the same among all z/OS systems (or any operating system that still uses PORTNAME) sharing the OSA-Express CHPID

    - On October 1, 2003 the following MCL was released for z800 and z900 in support of PORTNAME relief:

      Driver 3G - EC stream J11204, MCL032 (OSA level 3.33)

OSA-Express Token Ring outlook

- Demand for Token Ring is declining
- Increasing shortage of Token Ring components and suppliers
- Are you using OSA-Express Token Ring in your enterprise?

  Yes or No ?
- When do you plan to migrate off Token Ring? Estimate year.

Ethernet Coexistence

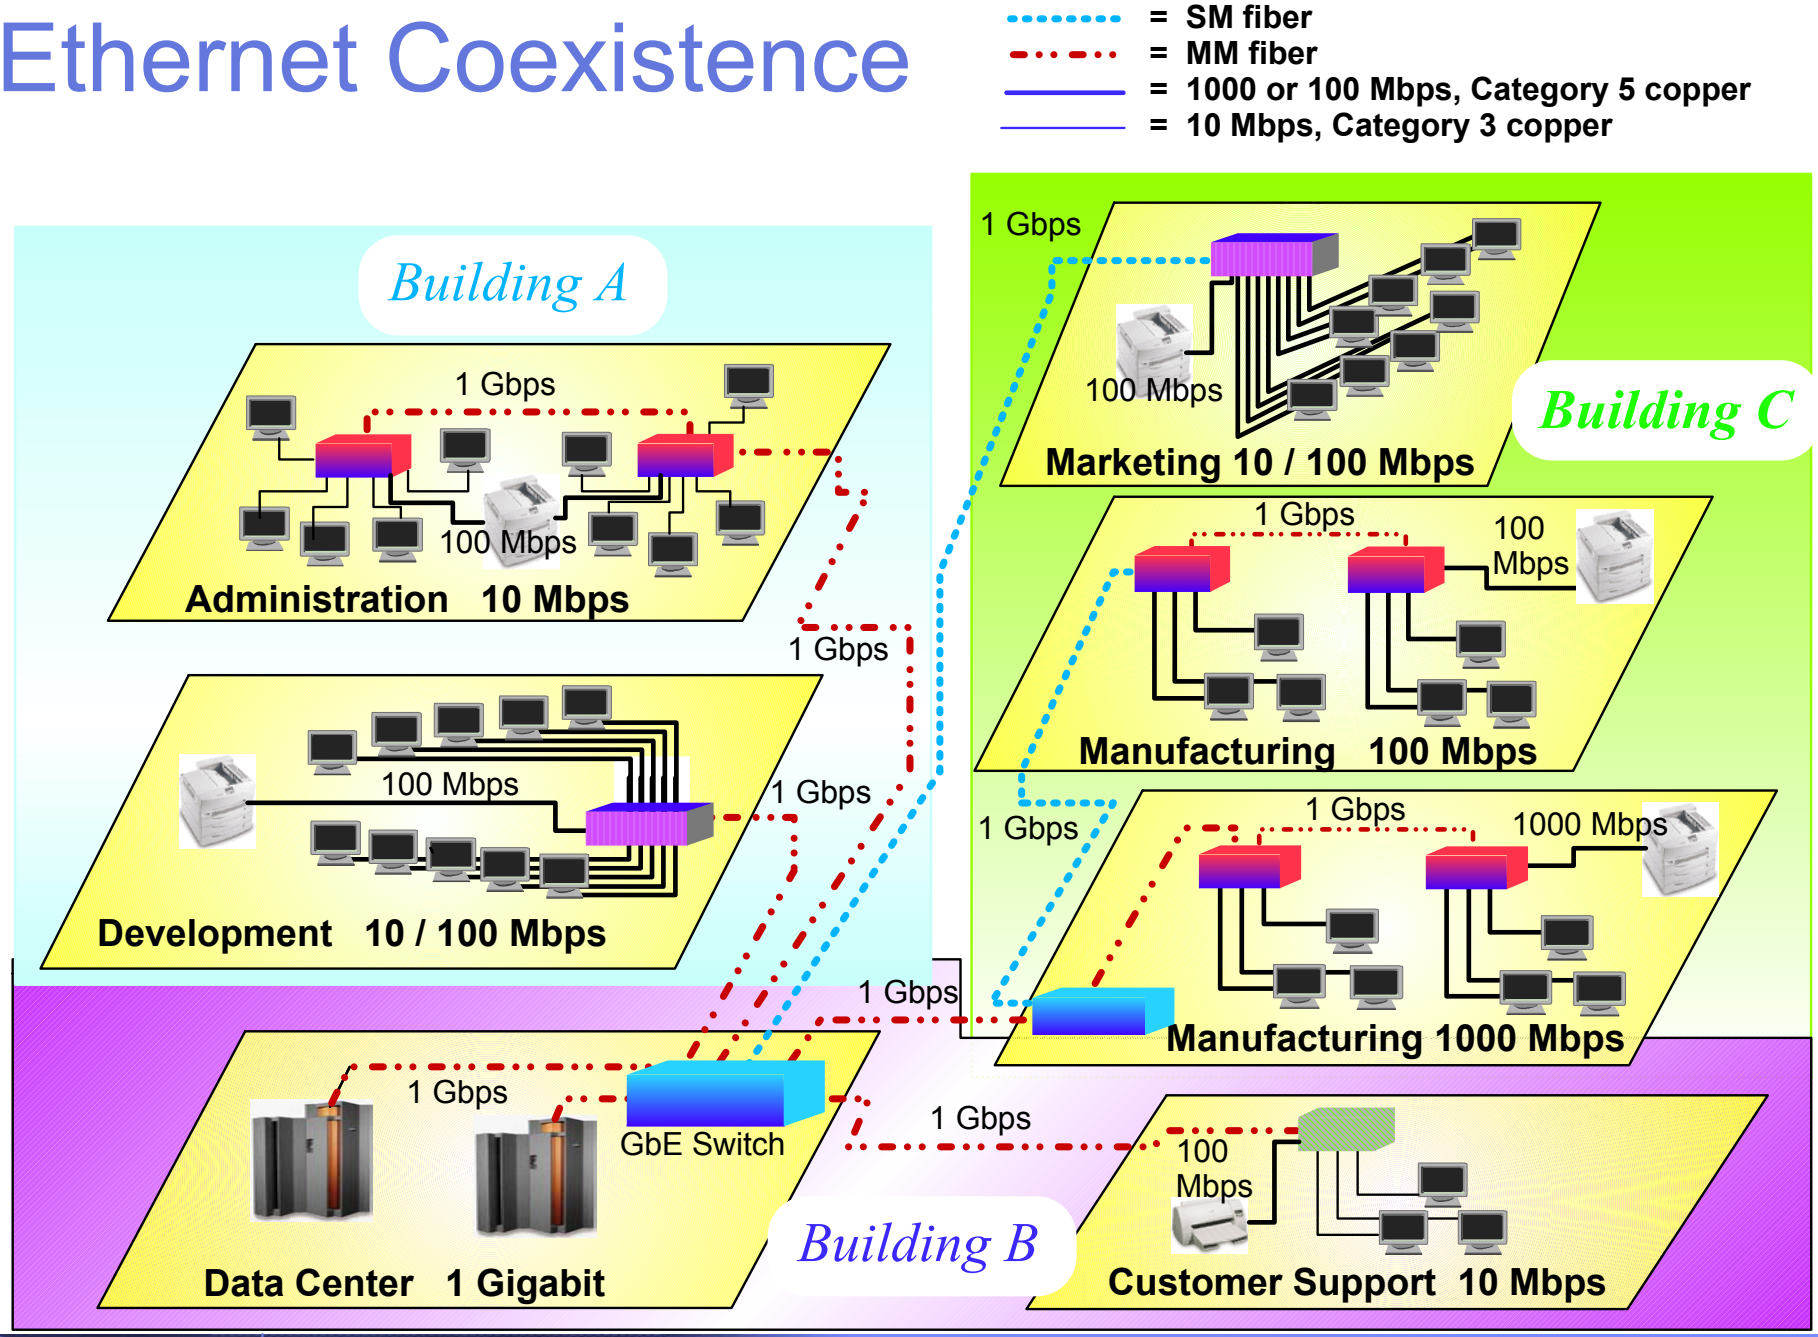

Fibre channel and Ethernet specifications

- The following charts summarize 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.
  - Fibre channel standard physical interface (FC-0)
  - Ethernet physical layer specification (IEEE 802.3)
- For 10 Gigabit speeds, the draft Fibre Channel specification (April 9, 2003) is closely aligned with the 10 Gigabit Ethernet physical layer specification (IEEE 802.3ae).
- 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 802.3)

		1 Gigabit Ethernet	/sec	10 Gigabit Ethernet	/sec
Fiber Core Microns (u) Light Source	Fiber Bandwidth @wavelength	Unrepeated distance	* Link loss budget	Unrepeated distance	* Link loss budget
9 u SM LX laser	@ 1310 nm	5 km 3.1 miles	4.6 dB	10 km 6.2 miles	6.0 dB
9u SM LX laser with MCP cable 50 or 62.5 u	500 MHz km	550 meters 1804 feet	2.4 dB	N / A	N / A
NEW 50 u MM SX laser	2000 MHz km @850 nm	Manufacturers guarantee to be equal or better than standard 50 um fiber (not yet qualified by IBM)		300 meters 984 feet	2.6 dB
50 u MM SX laser	500 MHz km @850 nm	550 meters 1804 feet	3.6 dB	82 meters 269 feet	2.3 dB
62.5 u MM SX laser	200 MHz km @ 850 nm	275 meters 902 feet	2.6 dB	33 meters 108 feet	2.5 dB
62.5 u MM SX laser	160 MHz * km @850nm	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 new aqua 2000 MHz km multimode fiber became available September, 2003.

*** Most often applicable to currently installed ESCON environments

Gigabit Ethernet link loss budget at high data rates

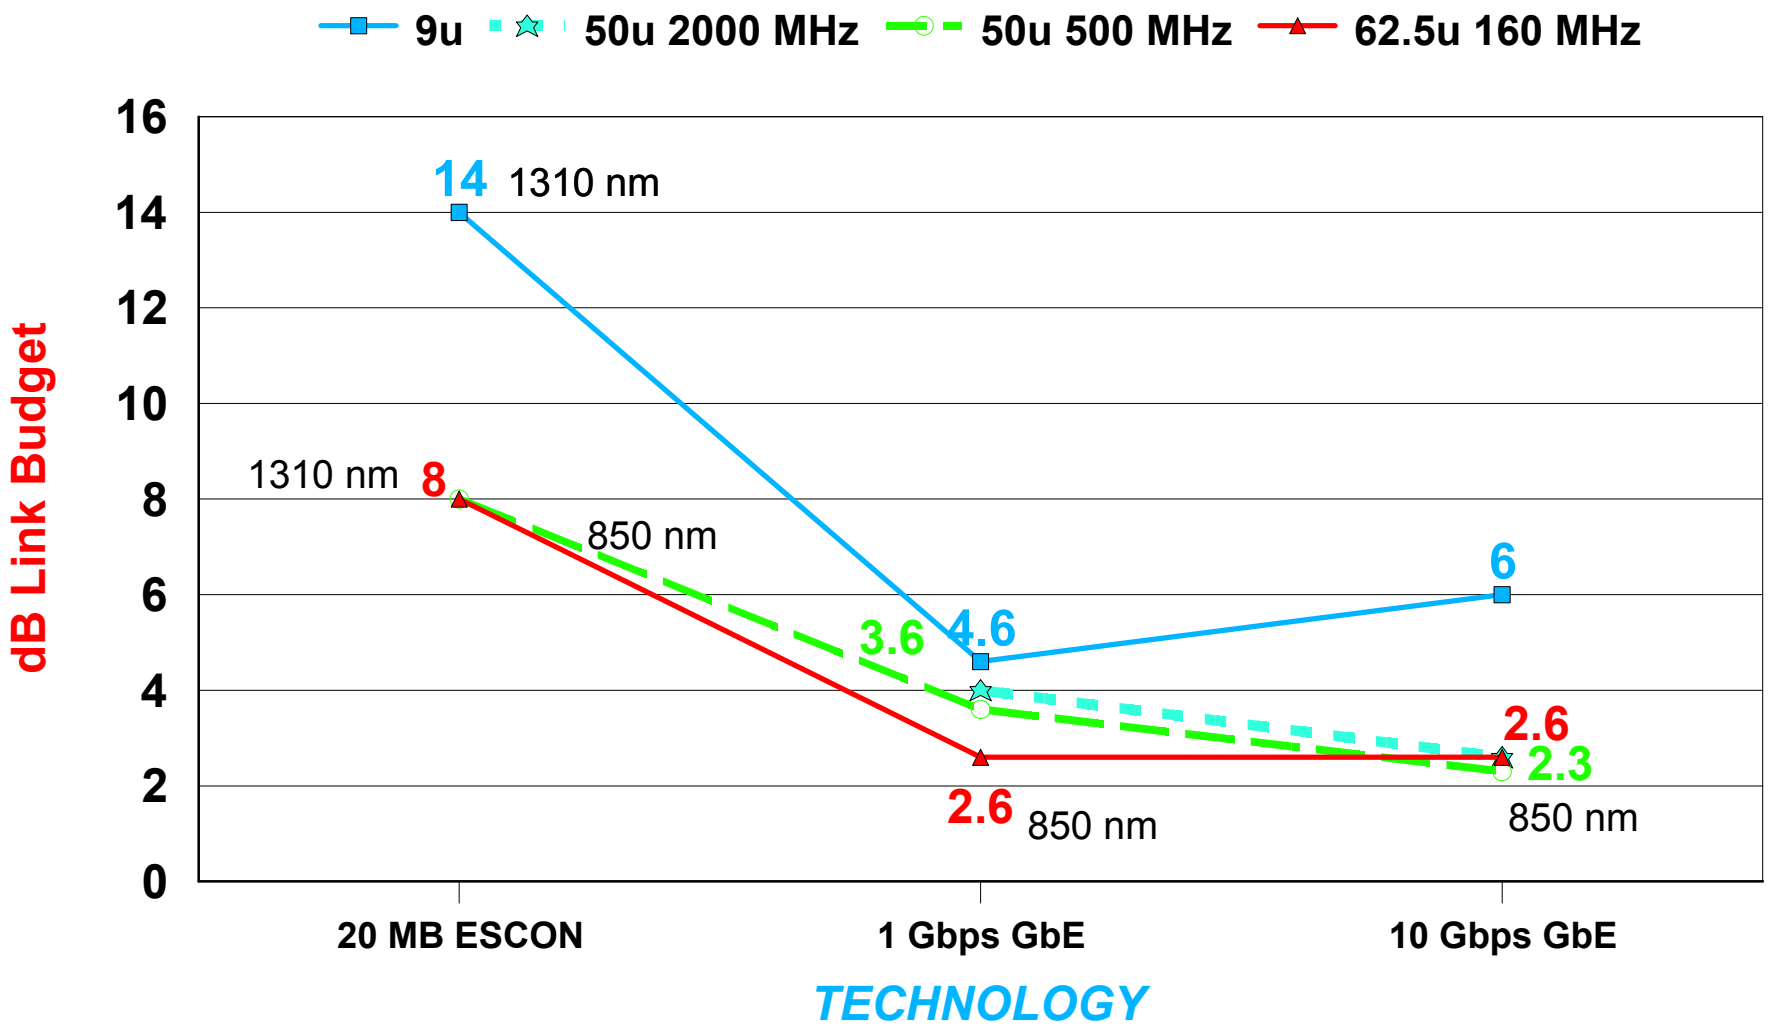

Gigabit Ethernet distances when using multimode fiber optic cabling

1 and 10 Gbps = 9u @1310nm - - - 10,000 meters (10 km)

Compare using 9u single mode fiber = 10,000 meters

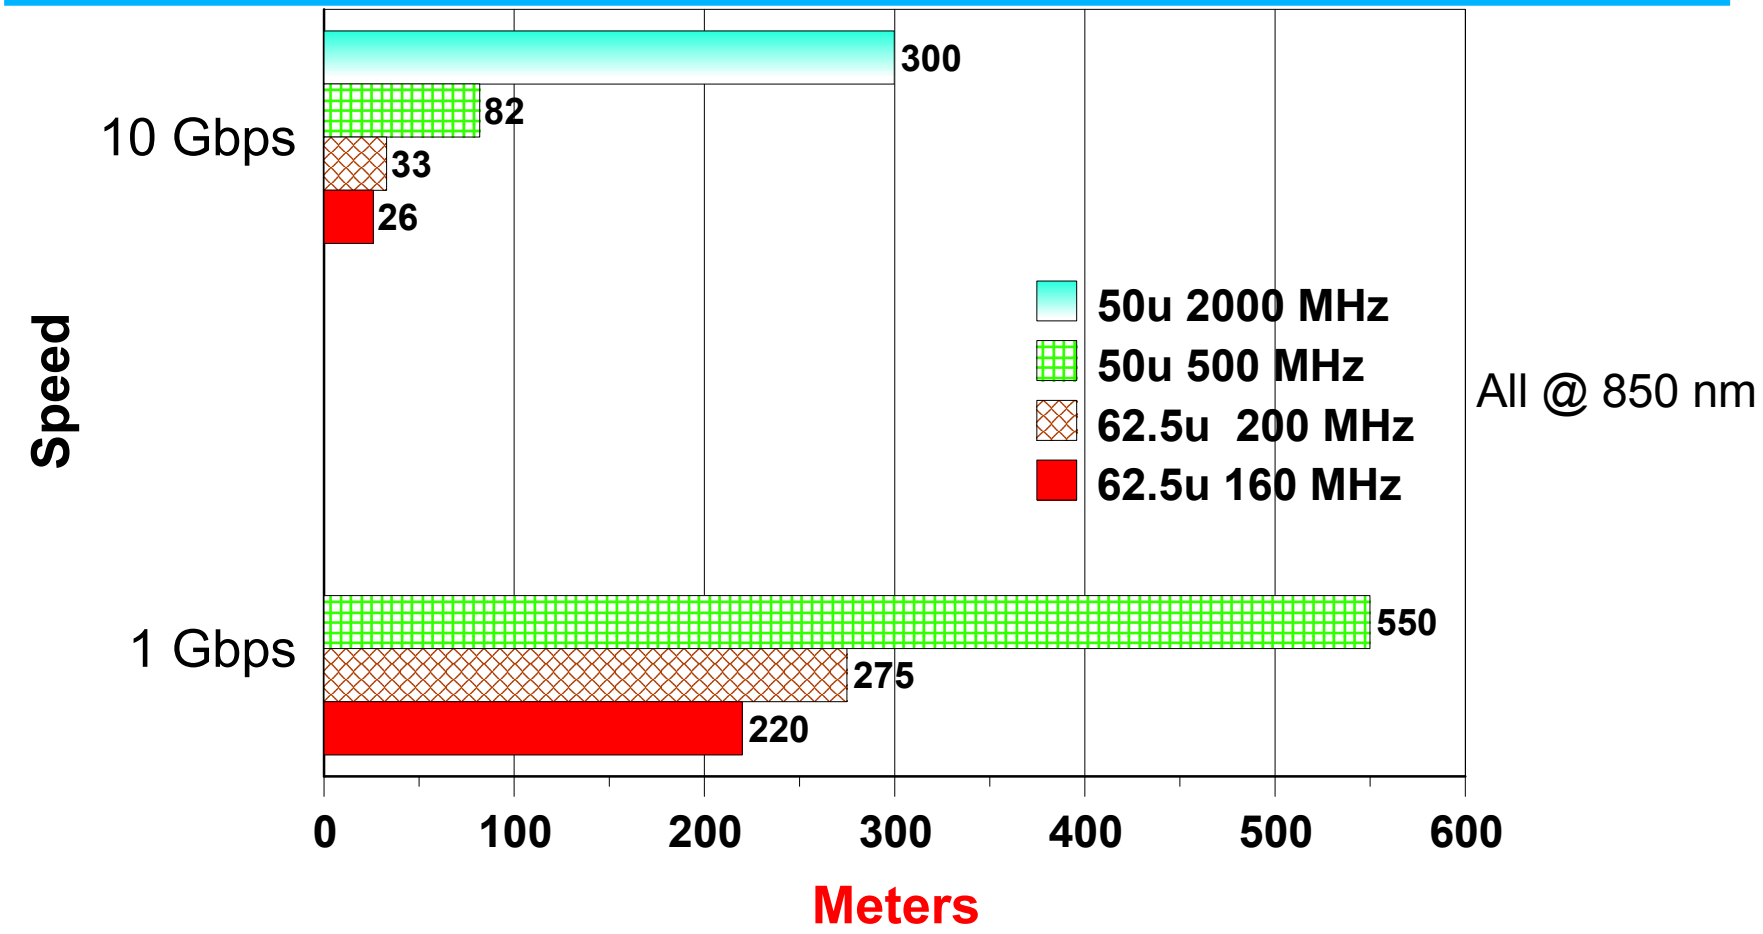

OSA -Express == Building on what we have . . .

Functions	G5 G6	z900 March 2001	z900 October 2001	zSeries May 2002	z990 June 2003
OSE (non-QDIO, SNA and / or TCP/IP)					
IPv4, Broadcast, Multicast, VIPA (manual)	X	>	>	>	>
SNMP (via OSA/SF), SNA (LSA), TCP/IP (LCS)	X	>	>	>	>
SNA at 1 Gbps over copper (1000BASE-T Ethernet)					X
OSD (QDIO, TCP/IP Only)					
IPv4, Multicast, VIPA (Dynamic)	X	>	>	>	>
SNMP (via OSA/SF) QDIO and non-QDIO	X	>	>	>	>
Primary/Secondary Routers	X	>	>	>	>
VLAN (Priority Tagging); 10/100/1000 Ethernet		X	>	>	>
ARP Query		X	>	>	>
IPv6 for 10/100/1000 Ethernet				X	>
VLAN for 10/100/1000 Ethernet (IEEE 802.1q), Linux				X	>
ARP Query and ARP Purge entire cache for IPv4				X	>
Broadcast support for RIP V1, Direct SNMP subagent				X	>
Multiple Secondary Routers for 10/100/1000 Ethernet				X (08/02)	>
Virtualization of adapter interruptions					X
Checksum offload for IPv4 packets					X
VLAN support (IEEE 802.1q), z/OS 1.5, z/VM 4.4					X
Intrusion Detection Services enhancements, z/OS 1.5					X
SNMP subagent support for dot3StatsTable					X

OSA-Express Limits

Limits	S/390 G5/G6	z900 (12/2000)	z900 (10/2001)	zSeries (5/2002)	z990 (6/2003)
Ports per feature (FC 5201 had 2 ports)	1	2	2	2	2
Features per system	12	12	12	12	24
CHPIDs required per port	1	1	1	1	1
CHPIDs required per feature	1	2	2	2	2
IP					
Home IP Addresses (IPv4+IPv6+DVIPA)	512	512	2048	2048	2048
Maximum # Dynamic VIPA @ is based on TCP/IP stack: OS/390 V2R8-R9: 64, V2R10, z/OS V1R1-V1R4 : 255					
Maximum IPv6 addresses				1024	1024
Maximum Multicast addresses (IPv4+ IPv6)	64	64	64	1024	1024
ARP Table Size *	512	2048	8192	8192	8192
SNA					
PUs per port	2,048	4096	4096	4096	4096
OSE					
Subchannels per stack	2	2	2	2	2
Devices	240	240	240	240	240
IP Stacks per port/CHPID on system	120	120	120	120	120
OSD					
Subchannels per stack	3	3	3	3	3
IP Stacks per port/CHPID on system	15	80	80	80	160
Subchannels per port on system	240	240	240	240	480
IP stacks per LPAR	15	80	80	80	84
Devices per LPAR	240	240	240	240	255

* Note: The ARP Table's capacity limits equals the sum of the IPv4 Home Addresses, plus the IPv6 Home Addresses, plus the IPv4 Multicast Addresses, plus the IPv6 Multicast Addresses, and plus the IPv4 Remote Addresses stored in the table.

OSA-Express Limits

Limits	S/390 G5/G6	z900 GA1 (12/2000)	z900 GA2 (10/2001)	z900 GA3 (5/2002)	z990 (6/2003)
VLANs per stack					
Linux				4,096	4,096 IPv4 + IPv6
z/OS			Null tagging	Null tagging	1 per IPv4, 1 per IPv6 stack on z/OS V1.5
z/VM					1 per IPv4 stack, z/VM 4.4

Sum it Up!

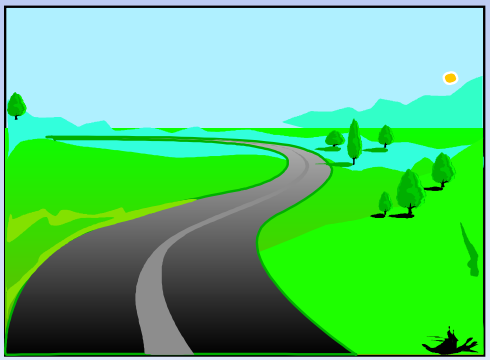

Speed Zone

? MBps

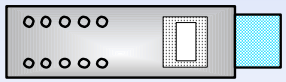

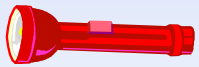

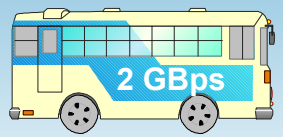

Direct SNMP

OSA/SF GUI

VLAN 802.1q

OSD

OSE

1000BASE-T Ethernet

48 NICs

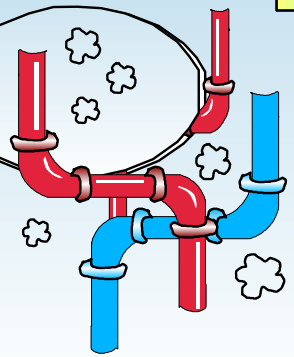

160 stacks

Host Memory

OSA-E

Checksum Offload

New GbEs

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: zSeries courses (z800, z900)
      - General Information for Planning a Physical Site (GIM)
- http://www.ibm.com/services/networking/
  - Product and Enterprise cabling offerings
- http://www.redbooks.ibm.com
  - IBM Redbooks

Check these out!

On the Internet

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

OSA-Express LAN Transmission Matrix

OSA-Express Feature	Auto-sense	Auto-negotiate	Operating Mode	Frame Size
Gigabit Ethernet	YES	YES	Full Duplex	802.3: 1492 bytes DIX: 1492 bytes Jumbo: 8992 bytes
1000BASE-T Ethernet Fast Ethernet	NO	YES	Full Duplex Half Duplex	802.3: 1492 bytes DIX: 1492 bytes
Token Ring	YES	N / A	4 Mbps: Half/full duplex 16 Mbps: Half/full duplex 100 Mbps: Full duplex	4 Mbps: 4550 bytes 16/100 Mbps: 18200

zSeries Features: Connectors / Cables

Feature Code	Feature Name	Connector Type	Cable Type	Comments
2362	OSA-Express 155 ATM SM	SC Duplex	9 micron SM	z800, z900
2363	OSA-Expresss 155 ATM MM	SC Duplex	50, 62.5 micron MM	z800, z900
1364	OSA-Express GbE LX	LC Duplex	9 micron SM	z990
2364	OSA-Express GbE LX	SC Duplex	9 micron SM	z900 to z990 upgrade only *
1365	OSA-Express GbE SX	LC Duplex	50, 62.5 micron MM	z990
2365	OSA-Express GbE SX	SC Duplex	50, 62.5 micron MM	z900 to z990 upgrade only *
1366	OSA-Express 1000BASE-T Ethernet	RJ-45	Category 5 UTP	z990
2366	OSA-Express Fast Ethernet	RJ-45	Category 5 UTP	z900 to z990 upgrade only *
2367	OSA-Express Token Ring	RJ-45	STP or UTP	z990

MM = Multimode fiber SM = Single mode fiber

* 2364, 2365, 2366 can be brought forward on an upgrade from z900.
They have been replaced by 1364, 1365, 1366 for new builds.

OSA reference materials

SA22-7935	z990 Open Systems Adapter-Express Customer's Guide and Reference
SA22-7476	z800, z900 Open Systems Adapter-Express Customer's Guide and Reference
SA22-7403	S/390 OSA-Express Customer's Guide and Reference (G5/G6)
GA22-7477	Planning for the Open Systems Adapter-2 for zSeries
SG24-5948	S/390 OSA-Express Implementation Guide (Redbook)
SG24-5444	IBM eServer zSeries Connectivity Handbook (Redbook)
SG24-5443	S/390 OSA-Express Gigabit Ethernet Implementation Guide (Redbook)
GX28-8002-10	Network and e-business Products Reference booklet (Redbook)
SG24-4770	Open Systems Adapter 2 Implementation Guide (Redbook)
SC28-1950-04	OS/390 Resource Measurement Facility Report Analysis
G221-0110	OSA-Express for IBM eServer zSeries 900 and S/390 Specification Sheet
GA23-0367-07b	Planning for Fiber Optic Links (ESCON, FICON, Coupling Links, and Open Systems Adapters)

OSA-Express Support Summary

VLAN support	Linux for zSeries	z/OS V1.5	z/VM V4.4
OSA-Express	yes, IPv4 & IPv6	yes, IPv4 & IPv6	yes, IPv4
HiperSockets	yes, z990, v4 and v6	no	yes, IPv4
z/VM Guest LAN	yes, IPv4	yes, IPv4	yes, IPv4

Broadcast support	Linux for zSeries	z/OS	z/VM V4.4
OSA-Express	yes, IPv4	yes, IPv4	yes, IPv4
HiperSockets	yes, IPv4, z990	yes, IPv4, z990	yes, IPv4, z990
z/VM Guest LAN	yes, IPv4	yes, IPv4	yes, IPv4

IPv6 support	Linux for zSeries	z/OS	z/VM
OSA-Express	yes	yes, V1.4 minimum	no
HiperSockets	yes, z990 only	no	no
z/VM Guest LAN	yes, z/VM 4.4	yes, z/VM 4.4	yes, V4.4

Minimum Software: z/OS, OS/390

Support inn service releases

- **OSA-Express Gigabit Ethernet**
  - OS/390 V2.10 and Communications Server (CS)
  - z/OS.e V1.3
- **OSA-Express 1000BASE-T Ethernet, Fast Ethernet**
  - OS/390 V2.10 and CS with PTFs if using the QDIO mode
  - OS/390 V2.10 if using the non-QDIO mode
    - Appropriate release of CS
    - Appropriate release of ACF/VTAM
  - z/OS.e V1.3
- **OSA-Express Token Ring**
  - OS/390 V2.10 and CS with PTFs if using the QDIO mode
  - OS/390 V2.10 with PTFs if using the non-QDIO mode
    - Appropriate release of CS
    - Appropriate release of ACF/VTAM
  - z/OS.e V1.3

Minimum Software: z/VM
Support in service releases

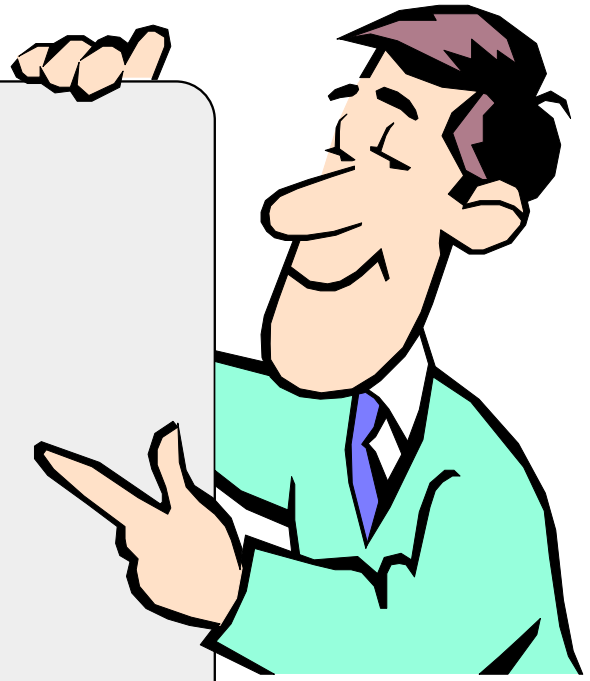

- **OSA-Express Gigabit Ethernet**
  - z/VM V3.1 with TCP/IP feature 330
  - z/VM V4.2, and later, with TCP/IP feature 330
- **OSA-Express 1000BASE-T Ethernet, Fast Ethernet**
  - z/VM V3.1, z/VM V4.2 , and later
    - With TCP/IP feature 330 if using QDIO mode
  - z/VM V3.1, z/VM V4.2, and later, if using non-QDIO mode
    - Appropriate releases of TCP/IP
    - Appropriate releases of ACF/VTAM
- **OSA-Express Token Ring**
  - z/VM V4.2, and later
    - With TCP/IP feature 330 if using QDIO mode
  - z/VM V3.1, z/VM V4.2, and later if using non-QDIO mode
    - Appropriate releases of TCP/IP
    - Appropriate releases of ACF/VTAM

Minimum Software: VSE/ESA

Support in service releases

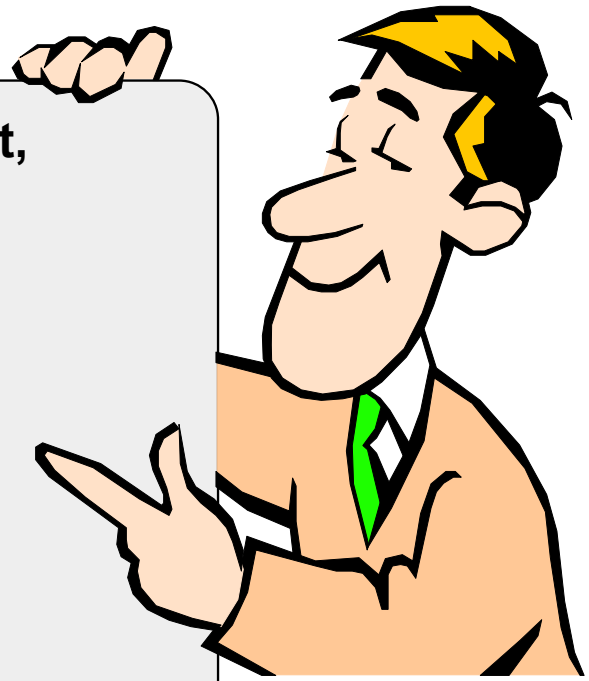

- **OSA-Express Gigabit Ethernet, 1000BASE-T Ethernet, Fast Ethernet, Token Ring**
  - VSE/ESA V2.6 for QDIO mode
    - Approrpriate release of TCP/IP
  - VSE/ESA V2.6 for non-QDIO mode
    - Appropriate release of TCP/IP
    - Appropriate release of ACF/VTAM
  - OSA/SF Java GUI requires a PTF
    - APAR is PQ73567
    - PTFs are
      - UQ76882 (VSE/ESA 2.5)
      - UQ76883 (VSE/ESA 2.6)
      - UQ76884 (VSE/ESA 2.7)

Minimum Software: TPF

- **OSA-Express Gigabit Ethernet only**
  - TPF 4.1 at PUT 13
    - Requires TPF APAR J27333
    - Includes VIPA (Virtual Internet Protocol Address) support, which allows TCP/IP connections to be balanced in a loosely coupled environment

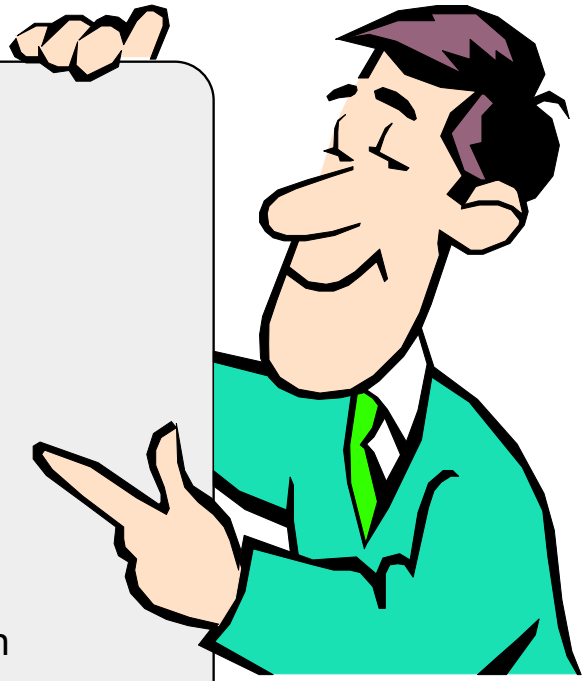

Minimum Software: Linux on zSeries

- **OSA-Express Gigabit Ethernet**
  - Linux kernel 2.2.16
- **OSA-Express 1000BASE-T Ethernet, Fast Ethernet**
  - Linux kernel 2.2.16
- **OSA-Express Token Ring**
  - Linux kernel 2.4

Thank You!