



Connectivity mit z9

Vortrag (G11)

Jörg Härtel, IBM München



VM / VSE / Linux for System z

Berlin 2007
Frühjahrstagung

26. - 28. März



© 2007 IBM Corporation

IBM System z9 FTSS



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	zVM *
		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.

Open Systems Adapter (OSA) – First availability dates


OSA-2, 12 features

G2 (Oct '95) G3 (Sept '96), G4 (June '97)

1995 (Oct) = OSA-2 Token Ring, Ethernet, FDDI

1996 (Aug) = 155 ATM

1998 (April) = Fast Ethernet



EN/TR

FDDI

OSA-Express


G5 (Sept '98), G6 (May '99)

z900 (Dec '00), z990 (June '03)


1999 = OSA-Express GbE (June)

2000 = Fast Ethernet & ATM (Jan)

2003 = **24 features, 1000BASE-T (June)**



G5, G6



NIC

13xx: z990, z890


2005, OSA-Express2

GbE, 10 GbE (Jan)


#3366: 1000BASE-T Ethernet (Sept, z9 EC)

Not shown

also available on z9 BC, z990, z890



PCI-X



PCI-X

3368: z9 EC, z9 BC

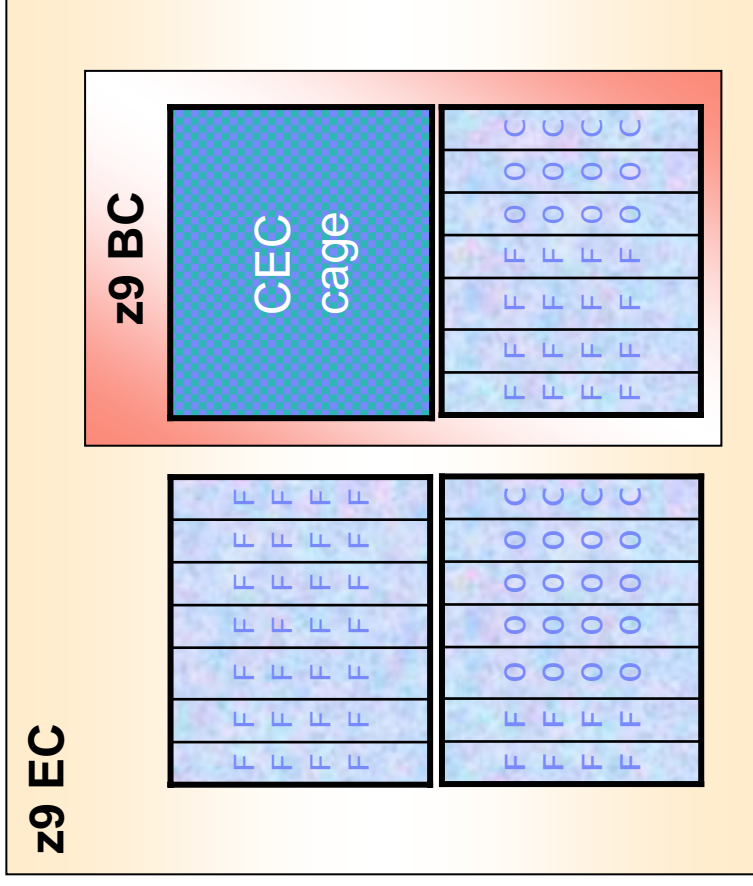
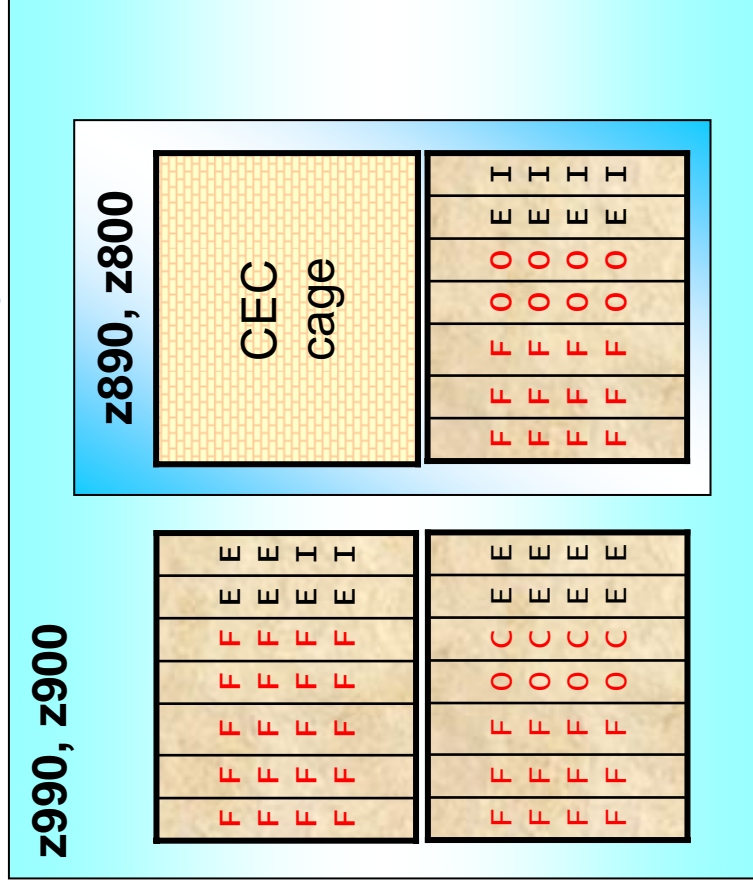
z990, z890

3364, 3365: z9 EC, z9 BC

z990, z890

Increased connectivity per I/O cage on z9 EC, z9 BC

- All supported FICON features - 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**
- z990, z890, z900, z800
20 per I/O cage
- z9 EC, z9 BC - 28 per I/O cage



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

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

Feature	Feature Name	Ports	z900	z990	z9 EC	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 z900 or z990

* 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

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

Feature	Feature Name	Ports	z800	z890	z9 BC	CHIPIDs	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	N/A	N/A	OSD, OSE	SC Duplex
2363	OSA-E 155 ATM MM	2	X	N/A	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 z900 or z990

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

640 TCP/IP stacks for improved virtualization

- **OSA-Express2 on z9 EC, z9 BC, z990, z890**
 - ▶ 640 TCP/IP stacks per OSA-Express2 port/CHPID
- **For hosting more images on server**
- **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, z/VSE V3.1
 - ▶ Linux on System z








Limits	S/390 G5/G6	z990 Dec 00 z800 Mar 02		z990 Jun 03 z890 May 04		z990, z890 Oct 04 OSA-E		z990, z890 > Jan 05 OSA-E2	
OSD									
Subchannels per stack	3	3	3	3	3	3	3	3	3
IP stacks per port/CHPID on server	15	80	160	160	160	160	160	640 \$	640 \$
Subchannels per port	240	240	480	480	480	480	480	1920 \$	1920 \$
IP stacks per LPAR	15	80	84	84	160	160	160	640 \$	640 \$
Devices per LPAR	240	240	254	254	480	480	480	1920 \$	1920 \$
Maximum control units	1	1	1	1	16	16	16	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

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)

Ethernet connectors and cabling

Feature #	Feature name	Connector	Cabling
OSA-Express on z900, z800			
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, z890			
OSA-Express on z900, z800 (only #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 EC, z9 BC, z990, z890			
3364	OSA-Express2 GbE LX	LC Duplex	9µ SM
3365	OSA-Express2 GbE SX	LC Duplex	50, 62.5µ MM
3366	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

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



OSA-Express2 10 Gigabit Ethernet LR cabling

LR= Long reach 1310 nm transceiver



- **z9 EC, z9 BC, z990, z890**

OSA-Express2, OSA-Express limits

Hardware limits		S/390 G5/G6	z900 Dec 00	z900 Oct 01 z800 Mar 02	zSeries May 02	z990 Jun 03	z990,z890 Oct 04 OSA-E	z990,z890 Jan 05 OSA-E2	z9 EC Sept 05 z9 BC May 06
Ports per feature (G5/G6 #5201 (EN/TR) has 2 ports)		1	2	2	2	2	2	GbE = 2 10 GbE = 1	100BASE-T 2 ports
Features per server		12	12	12	12	24	24	24	24
CHPID per port		1	1	1	1	1	1	1	1
IP									
Home IP addresses (IPv4 + IPv6 + DVIPA) per port		512	512	2048	2048	2048	2048	4096	4096
Multicast addresses (IPv4 + IPv6)		64	64	64	1024	1024	1024	2048	2048
ARP table size *		512	2048	8192	8192	8192	8192	16384	16384

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

The dates are general availability dates for hardware.

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

OSA-Express2, OSA-Express limits


Hardware limits		S/390 G5/G6	z900 Dec 00	z900 Oct 01 z800 Mar 02	zSeries May 02	z990 Jun 03	z990,z890 Oct 04 OSA-E	z990,z890 Jan 05 OSA-E2	z9 EC Sept 05 z9 BC May 06
OSE (non-QDIO) Applicable only to copper OSA									
Subchannels per stack		2	2	2	2	2	2	NA	2
Devices per port		240	240	240	240	240	240	NA	240
IP stacks per port/CHIPID		120	120	120	120	120	120	NA	120
PUs per port/CHIPID/MAC address		2048	4096	4096	4096	4096	4096	4096	4096
Control units per port/CHIPID		1	1	1	1	1	1	1	1
OSD (QDIO)									
Devices/subchannels per stack, 2 control & 1 data		3	3	3	3	3	3	3	3
Devices/subchannels per LPAR		240	240	240	240	254	480	1920 \$	1920
IP stacks per port / CHIPID		15	80	80	80	160	160	640 \$	640
IP stacks per LPAR		15	80	80	80	84	160	640 \$	640
Devices / subchannels per port		240	240	240	240	480	480	1920	1920
Control units per port/CHIPID		1	1	1	1	1	16	16	16
OSN (OSA for NCP) GbE, 1000BASE-T Ethernet									
Connections per CHIPID (374x subchannels)		NA	NA	NA	NA	NA	NA	NA	180
Number of QDIO devices (to communicate with CCL)		NA	NA	NA	NA	NA	NA	NA	480
MAC addresses per port		NA	NA	NA	NA	NA	NA	NA	2048

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

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

CHPIDs 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

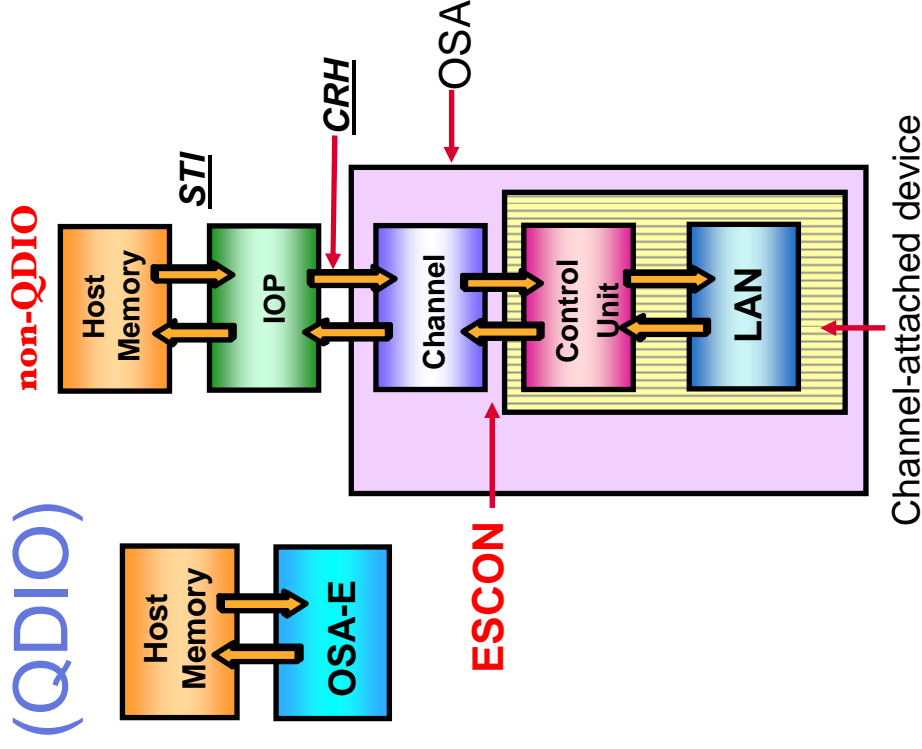
CHPID type OSD Queued Direct Input/Output (QDIO) since June 1999

- 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 System z, TPF

- 10 Gigabit Ethernet, Gigabit Ethernet, 1000BASE-T Ethernet, Fast Ethernet



Functions - QDIO only

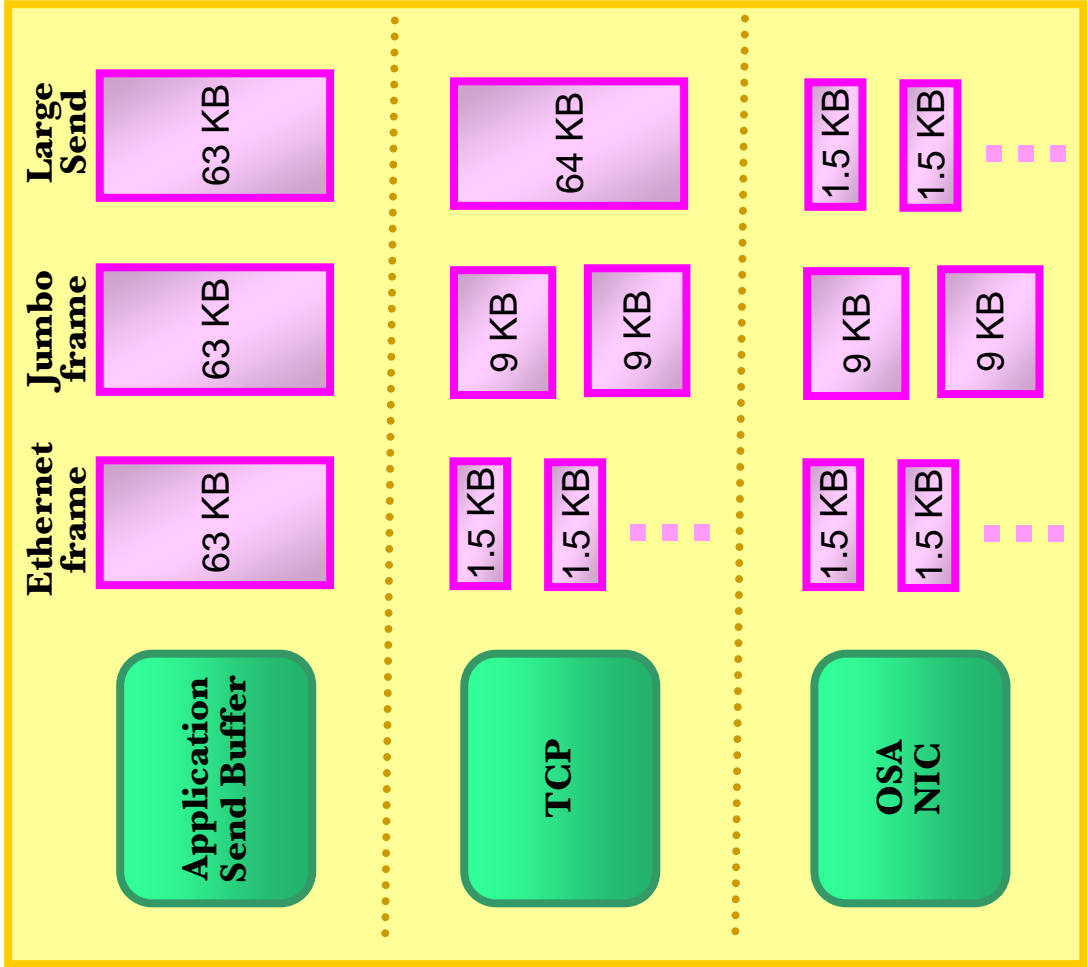
- **QDIO IP Processing Assists**
 - Provides Multicast support
 - Builds MAC and LLC headers
 - Performs filtering - TCP/IP stack sees only IP datagrams
 - Performs all ARP processing



- **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 z
 - 1000BASE-T Ethernet (#1366, #3366)
 - Gigabit Ethernet (#1364, 1365, #3364, #3365)
 - 10 Gigabit Ethernet (#3368)

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

- z9 EC, z9 BC, z990, z890
- OSA-Express2
- 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
 - ▶ z/OS V1.6 with PTF, z/OS V1.7
 - ▶ Linux on System z9
 - ▶ Linux SUSE SLES 9, SP2




OSA VLAN support

VLAN support			
	Linux on System z	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 EC, z9 BC: QDIO	No	Yes. V1.7 with PTF	Yes, V5.1 with PTF 2Q06

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 supports 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




Also OSA-Express
1000BASE-T Ethernet #1366

OSA-Express2 OSN (OSA for NCP)

Hardware available September 2005

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



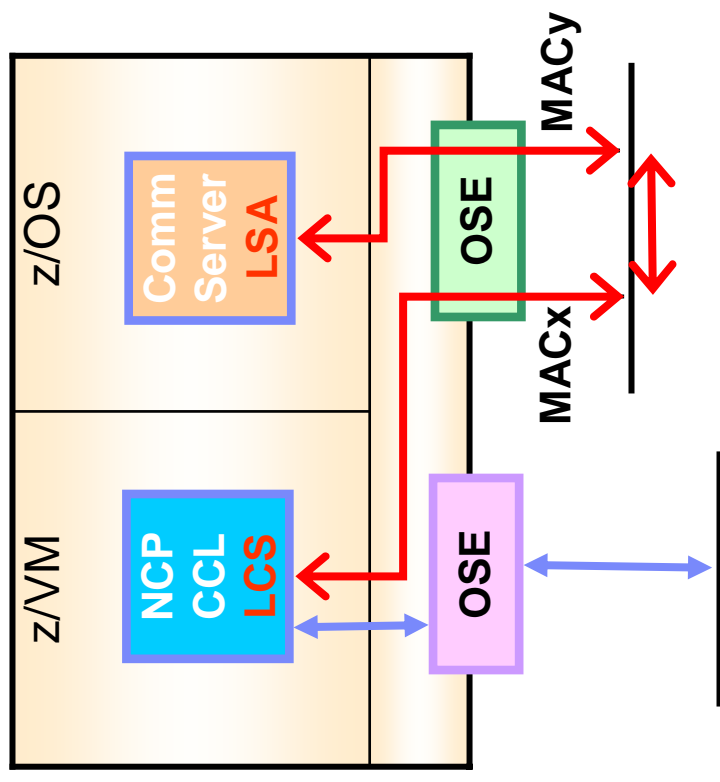
No longer limited to
“copper OSA”

374X Network Control Program

zSeries



IBM Communication Controller for Linux V1.1 zSeries

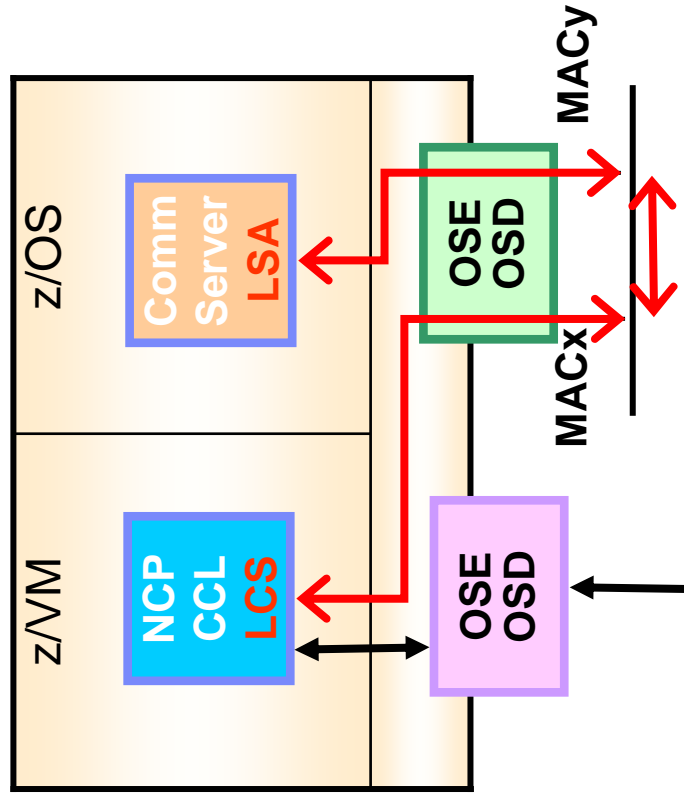


1. z/OS defined as LSA, own MAC/port
2. CCL defined as LCS, own MAC/port
3. 3rd port for SNA devices on LAN
4. All CHPID type OSE

IBM Communication Controller
for Linux V1.1

CCL traffic 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

CDLCL 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/CDLCL to CCL

z/VM Connectivity Options

Virtual Lan

- ▶ Hiper
- ▶ ODIO
- ▶ virtual NIC


Virtual Switch

- ▶ IP (Layer 3) | Ethernet (Layer 2)
- ▶ VLAN unaware
- ▶ VLAN aware
 - Trunk
 - Access
- ▶ virtual NIC

z/VM Connectivity

- z/VM virtual LAN Concept
 - ▶ Virtual network adapters connect IP stacks in virtual machines.
 - ▶ No hardware is required.
 - ▶ It's all done by CP commands, directory statements, configuration file statements, etc.
 - ▶ High speed and high volume networks.
 - ▶ One z/VM system can have multiple guest lans.
 - ▶ Guest lans can connect to other guest lans ...
 - or be isolated from other guest lans
 - ▶ One IP stack can belong to multiple guest lans.
 - ▶ Supports multicast, unicast, broadcast networks.
 - ▶ Supports all protocols.
 - ▶ VM TCPIP and linux support guest lan


z/VM virtual LAN




z/VM Connectivity

- z/VM VSWITCH Concept
 - ▶ Special kind of Guest LAN
 - ▶ Like a Guest LAN Provides network of virtual network interfaces
 - ▶ Connects directly to an OSA-Express QDIO Interface
 - Or can run disconnected from real devices.
 - ▶ Connects to external LAN segments without need for Routing on z/VM.
 - ▶ Operates as layer 2 or layer 3.
 - ▶ Can have multiple Vswitches on one z/VM LPAR.

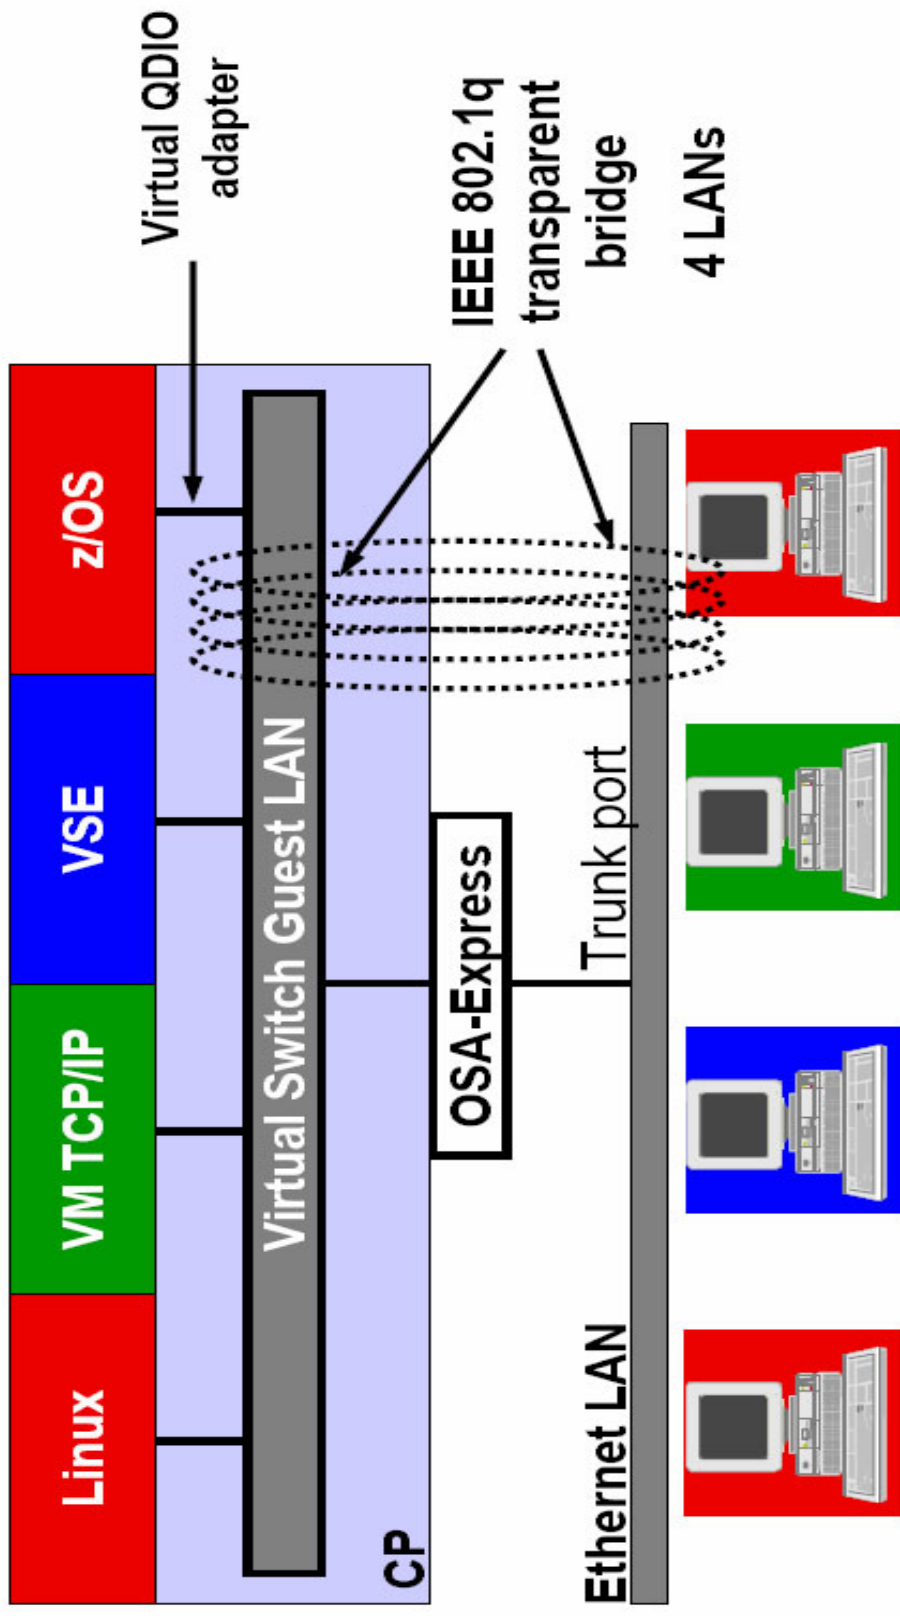
z/VM VSWITCH



z/VM Virtual Switch – VLAN unaware



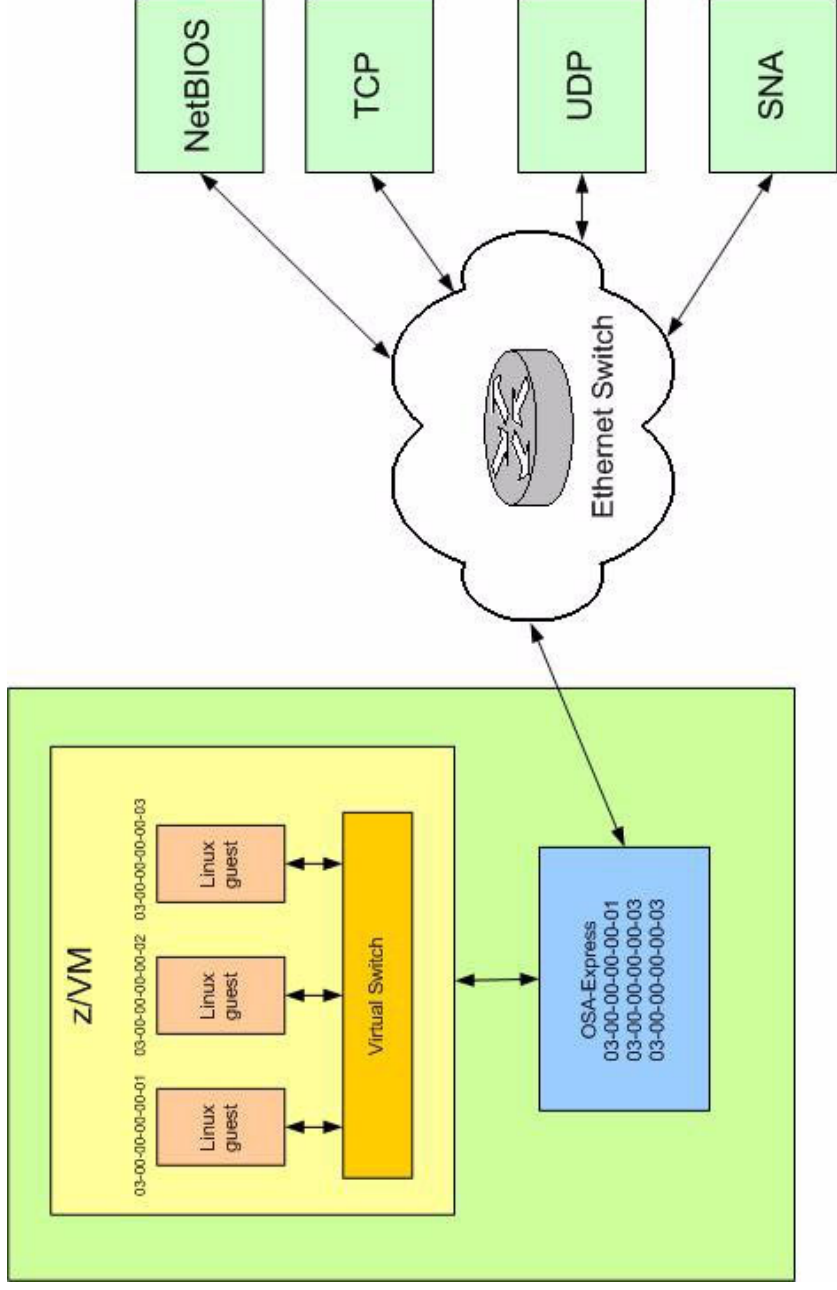
z/VM Virtual Switch – VLAN aware



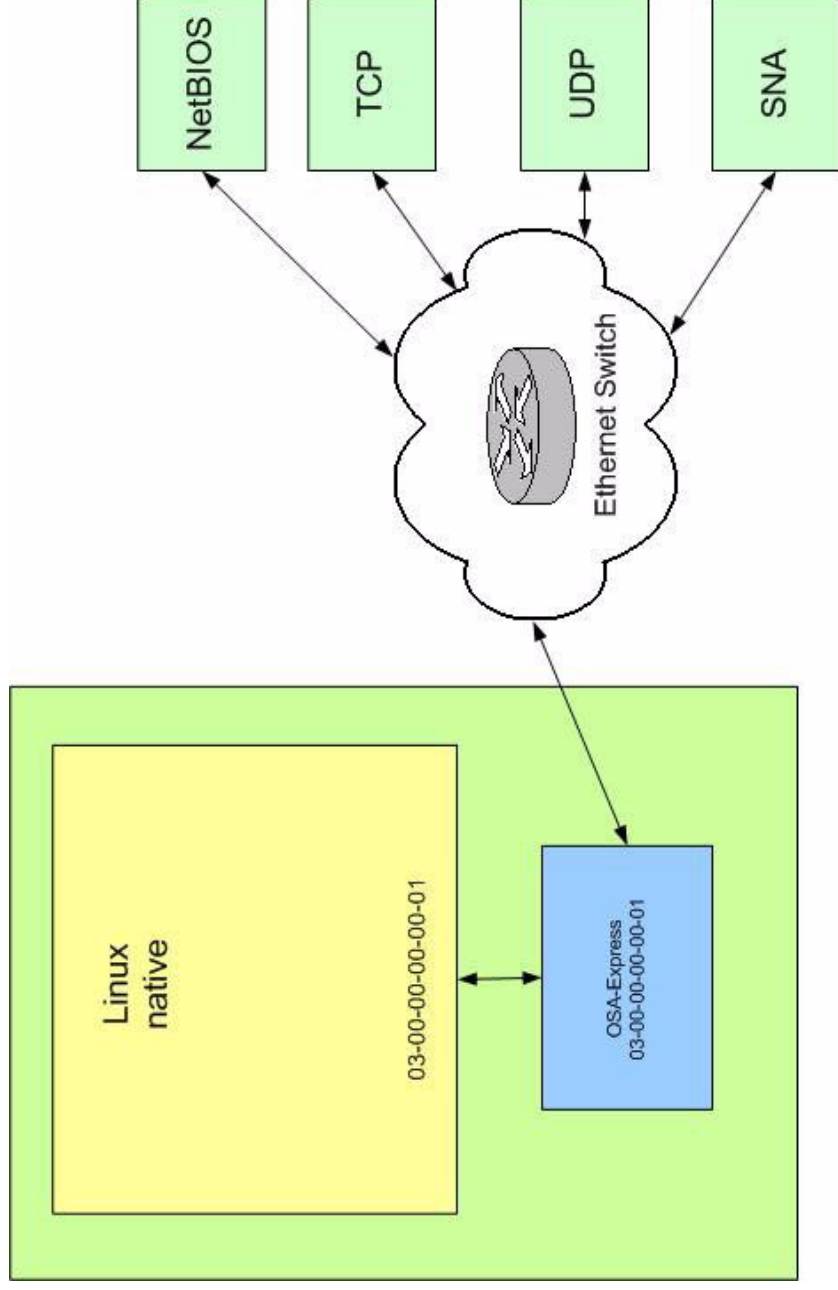
Layer 3 versus Layer 2

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

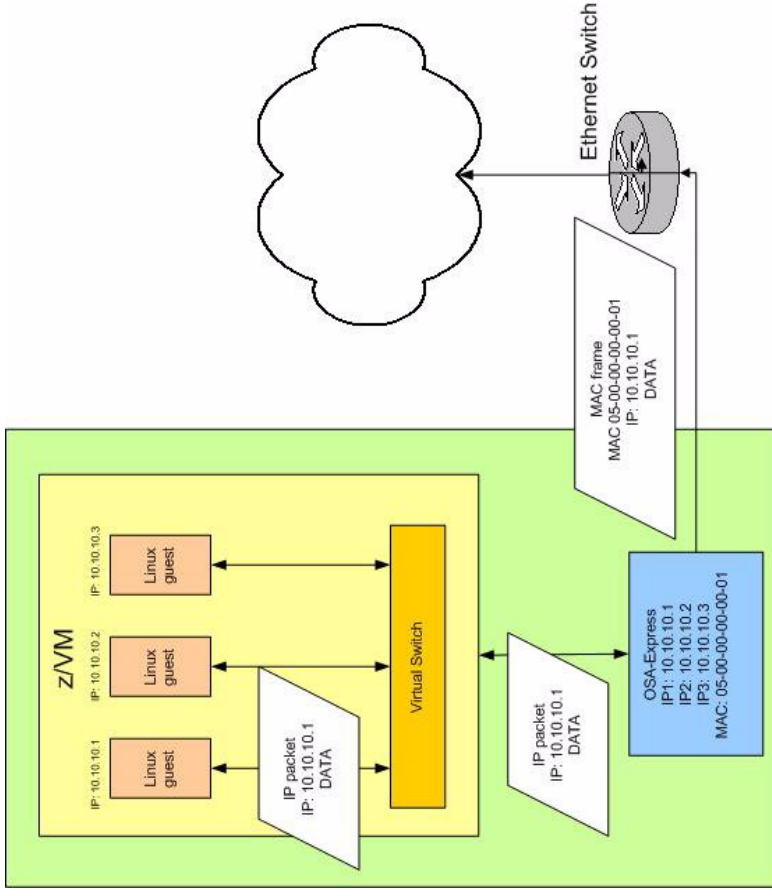
Layer 2 - protocol independent z/VM - Linux



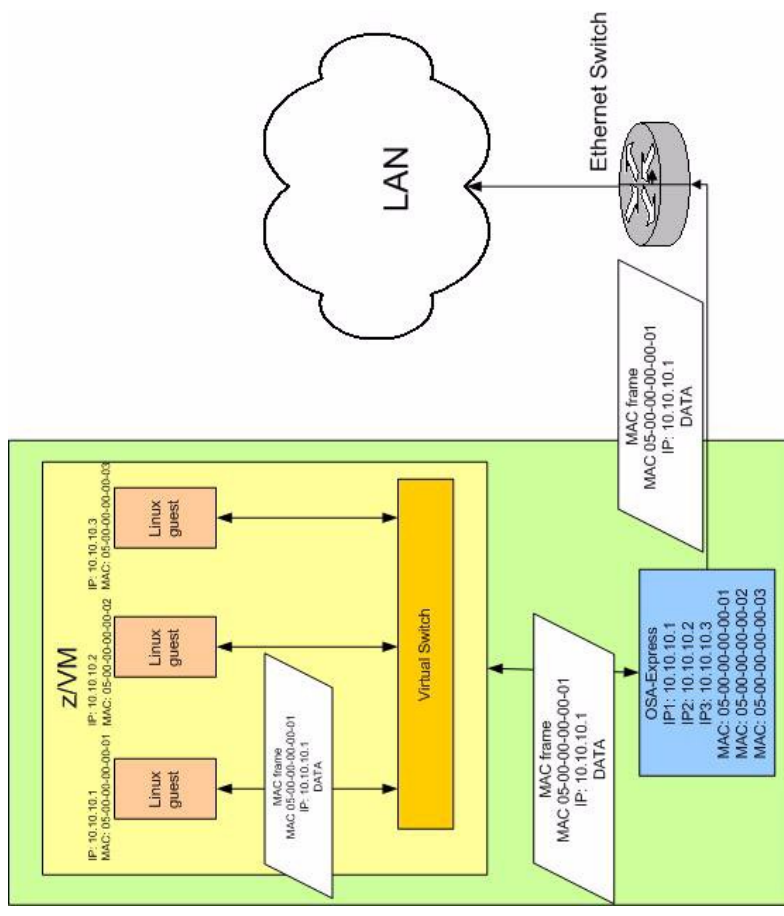
Layer 2 - protocol independent - Linux



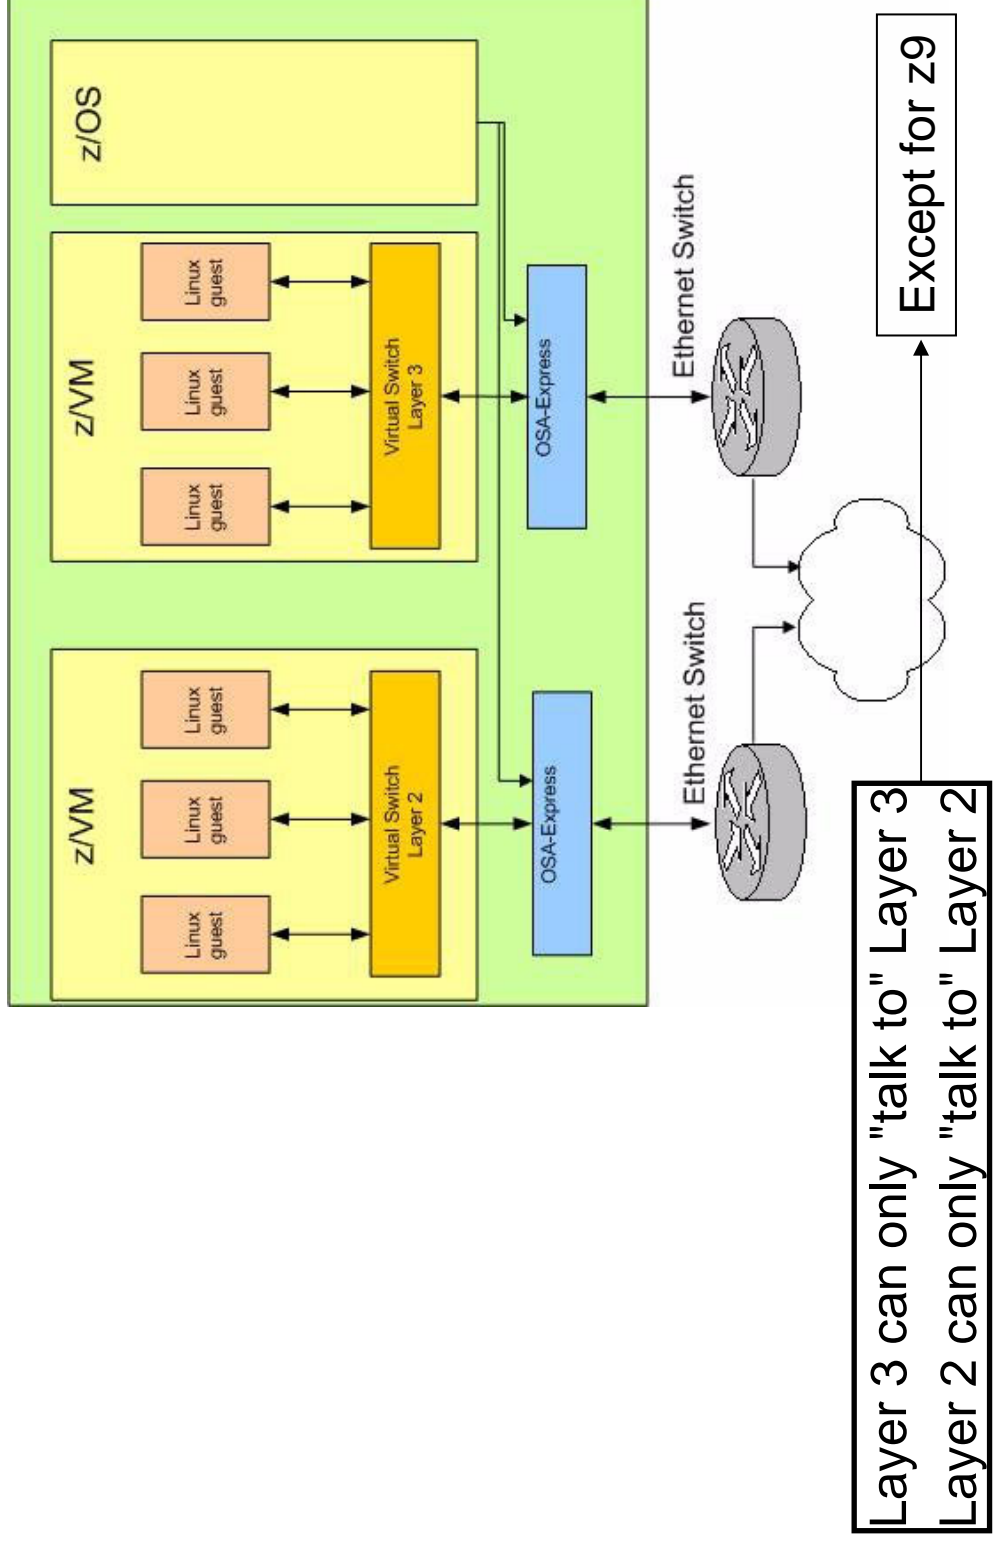
Layer 3 mode




Layer 2 mode



Layer 2 and Layer 3 traffic passing through same OSA



Performance



Problem Determination

- Inside z/VM
 - Q VMLAN
 - Service Level
 - MAC Details
 - Q VSWITCH DET
 - Type of Vswitch
 - registered MAC and IP Addresses
 - Q VSWITCH ACC
 - Q NIC DET

Problem Determination

- Inside z/VM
 - ▶ z/VM CP
 - TRSOURCE
 - LAN
 - TRSAVE
 - IPFORMAT
 - ▶ Linux for zSeries (z/VM 5.2, SLES 10)
 - Trace virtual LAN and VSWITCH Traffic
 - NIC Device in PROMiscuous Mode
 - Take a Trace using tcpdump (Linux Tool)
 - Analyse Trace using Ethereal

PROMiscuous Mode captures everything

Problem Determination

- Inside z/VM
 - ▶ z/VM TCP/IP
 - Netstat
 - HOME
 - DEV
 - GATEWAY
 - CONN
 - ARP
 - CP Q NIC DET
 - Trace
 - Process Names
 - ◆ IPUP, IPDOWN, PACKET, ...
 - Group Process Names
 - ◆ ALL, RAWIP, TCP,

Problem Determination

- Inside z/VSE
 - ▶ QUERY
 - SET
 - VERSION
 - ARP
 - ROUTES
 - LINKS
 - ▶ PING *host ipaddress*
 - ▶ TRACERT *host ipaddress*
 - ▶ Take a Trace

Problem Determination

- Inside z/VSE
 - ▶ How to take a trace
 - Enter the following commands on your VSE console:
 - MSG xx (xx = partition ID of target TCP/IP partition)
 - DEFINE TRACE,ID=xxxx,IPADDR=ipaddr-of-target-system
 - > recreate the problem
 - DUMP TRACES
 - SEGMENT NEW
 - DELETE TRACE,ID=xxxx
 - Download the SYSLSST output containing the trace data to your PC in ASCII format.
 - Now you can use the IP Trace Tool to convert and view this trace in Ethereal (now Wireshark).

www-03.ibm.com/servers/eserver/zseries/zvse/downloads/tools.html

IP Trace Tool

Updated: 01/2007

This Java tool is able to read IP packet traces captured with TCP/IP for VSE 1.5D as well as 1.5E and convert it into the CAP trace format that can be opened and viewed with Wireshark (formerly Ethereal) or Packetalyzer. Download and extract the ZIP file into a new directory (e.g. C:\IPTraceTool). A HTML documentation is contained in the 'doc' directory. Please refer to this documentation for further setup and usage information.

In order to install and use the IP Trace Tool, you need a **Java runtime environment** (JRE) or a **Java developers kit** (JDK) 1.4 or higher. You can download JDKs/JREs for various platforms from IBM developerWorks or Sun Microsystems, Inc.

Note: This tool is distributed 'as-is', no support, no warranty.

 **Download now** (zip, 68KB)

Problem Determination

- Inside Linux for zSeries
 - ▶ Check /var/log/message
 - ▶ Ifconfig
 - ▶ route
 - ▶ Netstat -s
 - ▶ traceroute
- ▶ Take a Trace using tcpdump (Linux Tool)
- ▶ Analyse Trace using Ethereal

Ethereal the Network Sniffer




Ethereal is Open Source software

- Ethereal is free
- released under the GNU General Public License
- <http://www.ethereal.com/download.html>

OFFICIAL RELEASES

Installable packages for Windows, Red Hat Linux/Fedora, and Solaris can be found on the main Ethereal website. Source code can be found on the main Ethereal website. [Ethereal](#) can be found on the main Ethereal website. [Ethereal](#) can be found on the main Ethereal website.

Windows:  or or [SourceForge](#)

Red Hat / Fedora:  or or [SourceForge](#)

Solaris:  or or [SourceForge](#)

Source code:  or or [SourceForge](#)

IBM: z/VM Operating System - Mozilla Firefox

Datei Bearbeiten Ansicht Gehe Lesezeichen Extras Hilfe

Privat Google Log on to IBM Messa... Erste Schritte Eingang MyOffice@Home

http://www.ibm.com/

z/VM[®]
the newest VM hypervisor based on 64-bit z/Architecture.

Currently supported releases of z/VM

Announced: z/VM V5.3
Also supported: z/VM V5.2
z/VM V5.1

The z/VM hypervisor is designed to help clients extend the business value of mainframe technology across the enterprise by integrating applications and data while providing exceptional levels of availability, security, and operational ease. z/VM virtualization technology is designed to allow the capability for clients to run hundreds to thousands of Linux servers on a single mainframe running with other System z operating systems, such as z/OS, or as a large-scale Linux-only enterprise server solution. z/VM V5.3 can also help to improve productivity by hosting non-Linux workloads such as z/OS, z/VSE, and z/TPF.

Summary of News and Updates

View 21 Mar. 2007 updates.
Read the z/VM and VM Site News and Changes for a summary of VM-related news, announcements, pointers, new classes, and places to hear about z/VM virtualization technology.

Worldwide announcement letters (US letters / product links below)

- Feb. 06, 2007 IBM z/VM V5.3 - Improving scalability, security, and virtualization technology
- Apr. 27, 2006 z/VM V5.2 New Function Added in Support of System z9
- Apr. 27, 2006 Introducing new members to the System z9™ family
- Aug. 15, 2005 V1.2.0 IBM Backup and Restore Manager for z/VM
- Feb. 21, 2006 z/VM V5.2 Generally Available
- Dec. 16, 2005 Archive Manager and Operations Manager for z/VM
- Aug. 23, 2005

z/VM platform resources

VM Technical resources

z/VM V5.3 Resources	Developer pages	Pipelines
z/VM V5.2 Resources	Downloads	Printing resources
z/VM V5.1 Resources	e-business	Security
Library and Literature	Networking	Storage management
Technical resources	OpenExtensions	System management
Database	Performance	Virtual Networking

Mainframe history

40 years and counting
Explore IBM mainframe innovation

Is your VM current ?
Thinking about migration?

Technical Conference

Munich
April 16-20, 2007

San Antonio
Sept. 17-21, 2007

New to Linux on VM?
This 4.5-day class, Installing, Configuring and Servicing z/VM for Linux Guests (ZV062) may be just what you need

New z/VM Class!
Enroll in this 4.5-day class to learn about z/VM RACF and DirMaint Implementation (ZV200)

Take back control

IBM z/VM Virtual Networking resources - Mozilla Firefox

Datei Bearbeiten Ansicht Gehe Lesezeichen Extras Hilfe

http://www.ibm.com/virtualnetworkk/

Privat Google Log on to IBM Messa... DasTelefonbuch Erste Schritte Eingang MyOffice@Home

Country/region [select] Search

Home Products Services & solutions Support & downloads My account

IBM System z

→ More about System z
Optimized to help meet the requirements of today's on demand solutions.

→ Linux on System z

Who is using Linux on System z? What applications are there? Have you heard about the Community Development System for Linux ?

→ Check out upcoming events where you can hear about z/VM Virtual Networking and more

IBM System z > System z > z/VM >

z/VM Virtual Networking

Did you know that you can create virtual LAN segments that connect your z/VM guests together without the need for all those messy point-to-point connections? And that you can do that without creating new subnets? Did you know you can do away with using a virtual router for your Linux farm and use VSWITCH to provide direct connection to physical LAN segments for all your guests? And you can also design, configure and operate your network using VSWITCH to provide High Availability and Automatic Network failover and/or the failover of an OSA?

Intigued? Introduce yourself to these powerful technologies provided in z/VM Guest LANs and the z/VM Virtual Switch.

The links on this page are intended to provide helpful information to customers who are using the following types of z/VM virtual networks:

- Virtual Switches
- QDIO Guest LANs
- HiperSockets Guest LANs.

For an overview of this support, see the [z/VM V5.2 library](#):

- "Planning for Guest LANs and Virtual Switches" in the **Connectivity** book
- **Getting Started with Linux on zSeries** book
- Virtual Switch and Guest LAN CP Maintenance Levels
- Virtual Switch TCP/IP Maintenance Levels
- Virtual Networking Release Enhancements
- Virtual Networking Hints and Tips

A new Redpaper:

- Networking Overview for Linux on zSeries **New !**

z/VM

- News
- About z/VM
- Events calendar
- Products and features
- Downloads
- Technical resources
- Library
- How to buy
- Service
- Education
- Site map
- Site search
- Printer-friendly
- Notify me
- Contact z/VM

SC24-6080-03



z/VM

Connectivity

version 5 release 2

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	Planning for Fiber Optic Links
SC33-7991	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

Haben Sie noch Fragen, sprechen Sie mich bitte an,
oder schreiben mir eine Mail an:

Haertel@de.ibm.com

Auch telefonisch bin ich für Sie da
089 4504 - 3240
oder 0171 3059653

Danke!



OSA-Express2/OSA-Express . . . Building for tomorrow

Features/functions	G5/G6 June 99	G5/G6 Jan 00	z900 Oct 01 z800 Mar 02	z990 June 03	z890 z990 May 04	z890 z990 Jan 05	z9 EC Sept 05 z9 BC May 06
OSA-Express features							
Gigabit Ethernet LX and SX (fiber)	X	>	>	>	>	>	>
Fast Ethernet (10/100 Mbps) (copper)		X	>	>	>	>	>
Token Ring (4/16/100 Mbps) (copper)			X	>	>	>	NA
1000BASE-T Ethernet (10/100/1000 Mbps) (copper)				X	>	>	>
Functions							
EMIF / MIF	X	>	>	>	>	>	>
OSA/SF Java GUI integrated in all operating systems				X	>	>	>
Spanned channels					X	>	>
OSA-Express2 features							
10 Gigabit Ethernet LR (fiber)						X	>
Gigabit Ethernet LX and SX (fiber)						X	>
1000BASE-T Ethernet (copper)							X

The dates are general availability dates for hardware.

OSA-Express2 / OSA-Express Building for tomorrow

CHPID types and functions	G5/G6 Jun 99	G5/G6 Jan 00	z990 Jun 03	z890 z990 May 04	z9 EC Sept 05 z9 BC May 06
OSE (non-QDIO, SNA and/or TCP/IP)	X				
IPv4, Broadcast, Multicast (OSPF), Primary/secondary routers/OAT entry, all media	X	>	>	>	>
VIPA (manual using OSA/SF), all media types		OS/390 VM/ESA	>	>	>
SNA at 1 Gbps over copper (1000BASE-T Ethernet)			SW - T	>	>
OSC (OSA Integrated Console Controller, OSA-ICC)				X	
TN3270E Non-SNA DFT 3270 emulation 1000BASE-T Ethernet only				z/OS V1.3 z/VM V4.4 VSE/ESA V2.6 TPF 4.1	>
120 console sessions per port				SW - T	>
OSN (OSA for NCP) OSA-Express2 exclusive					X
OSA-Express2 OSN supporting Communication Controller for Linux on System z9 Operating system PTFs required					z/OS V1.4 z/VM V5.1 VSE/ESA 2.7 TPF 4.1 Linux

The dates are general availability dates for hardware.

SW-T = Transparent to software.

OSA-Express2 / OSA-Express Building for tomorrow

CHPID type OSD functions	G5/G6 Jun 99	G5/G6 Jan 00	z900 Oct 01 z800 Mar 02	zSeries May 02	z990 Jun 03 z890 May 04	z9 EC Sept 05 z9 BC May 06
OSD (QDIO)	X					
Jumbo frames, Ethernet only at 1 / 10 Gbps	SW - T	>	>	>	>	>
Dynamic OAT, RMF – all media types	X	>	>	>	>	>
IPv4, Broadcast, Multicast (OSPF) Primary/secondary routers/OAT entry, all	X	>	>	>	>	>
VIPA (dynamic), all media		OS/390 VM/ESA	>	>	>	>
ARP query – all media Collect statistics regarding ARP cache			OS/390 V2.10	>	>	>
IPv6, Ethernet only				z/OS V1.4 Linux	z/VM V4.4 Aug 03	>
ARP cache management – all media types Purge IPv4 cache				z/OS V1.4	>	>
ARP cache management – all media For IP v4 - Query ARP table, Purge entries				Linux	>	>
TCP/IP broadcast support for RIP V1, all				z/OS V1.4 z/VM V4.3 Linux	>	>
Multiple secondary routers (Ethernet)				Aug 02 SW - T	>	>

The dates are general availability dates for hardware.
SW-T = Transparent to software.

OSA-Express2 / OSA-Express Building for tomorrow

CHPID type OSD functions		z990 Jun 03	Oct 2003	z890 z990 Oct 04	z890 z990 Jan 05	z9 EC Sept 05 z9 BC May 06
OSD (QDIO)						
Adapter interruptions, all media types		z/VM Linux	>	>	>	>
Performance assist for V=V guests, all media types		z/VM	>	>	>	>
Checksum offload for IPv4 packets, Ethernet only		z/OS V1.5 Linux *	>	>	>	>
160 TCP/IP stacks per port, all media types		SW - T	>	>	>	>
Intrusion Detection Services enhancements, all media		z/OS V1.5	>	>	>	>
Port name relief, zSeries, all media types			z/OS - NA z/VM V4.3 Linux	>	>	>
160 stacks per LPAR for all OSA-Express features (was 84)				z/OS V1.6 Linux *	z/VM V3.1 z/VM V4.3	>
Layer 2 for IP and non-IP workloads, Ethernet only				OSA-E z/VM V5.1 Linux **	OSA-E2 z/VM V5.1 Linux **	>

The dates are general availability dates for hardware.

SW-T = Transparent to software.

* SUSE SLES 8 and Red Hat RHEL 3 distributions

** SUSE SLES 9 SP2

OSA-Express2 / OSA-Express . . . Building for tomorrow

CHIPID type OSD functions	z990 Jun 03	Oct 2003	z890 z990 May 04	z890 z990 Oct 04	z890 z990 Jan 05	z9 EC Sept 05 z9 BC May 06
OSD (QDIO) – all exclusive to OSA-Express2						
640 TCP/IP stacks					z/OS V1.6 z/VM V5.1 VSE/ESA V2.6 Linux SUSE SLES 9 SP2	>
Large send					z/OS V1.6 Linux SUSE SLES 9 SP2	>
Concurrent LIC update					X	>
GARP VLAN Registration Protocol (GVRP)						z/OS V1.7 with PTF z/VM V5.1 *

GVRP = Generic Attribute Registration Protocol (GARP) Virtual Local Area Network (VLAN) Reigstration Protocol
 The dates are general availability dates for hardware.

* z/VM support planned to be available in second quarter 2006.

OSA-Express2 / OSA-Express Building for tomorrow

SNMP support	z900 Mar 01	zSeries May 02	z990 Jun 03	Oct 03	z990 z890 May 04	z9 EC Sept 05 z9 BC May 06
OSE (non-QDIO, LCS, LSA)						
SNMP MIB extension ATM Ethernet LANE	OS/390 V2.10	>	>	>	>	>
Direct SNMP for LCS – Get, GetNext, Trap, Set, all media					z/OS V1.6 Sept 04	>
OSD (QDIO)						
SNMP MIB extension; Ethernet and ATM Ethernet LANE	OS/390 V2.10	>	>	>	>	>
Direct SNMP query - Get, GetNext, all media types		z/OS V1.4 Linux *	>	>	>	>
dot3StastsTable (RFC 2665) Ethernet only, copper & fiber			z990 z/OS V1.4 Linux *	z900 z800	>	>
Direct SNMP – performance data all media types, z990, z900, z890				z/OS V1.4 Linux *	>	>
Direct SNMP – traps and set all media types					z/OS V1.5	>

The dates are general availability dates for hardware.
 LAN Channel Station (LCS) supporting TCP/IP traffic. Link Services Architecture (LSA) supporting SNA traffic.
 Direct SNMP for QDIO (OSD) removed the prerequisite of OSA/SF for SNMP management – first available May 2002.
 Direct SNMP for LAN Channel Station (LCS) - TCP/IP traffic.
 * Linux distribution SUSE SLES 8

OSA-Express2 / OSA-Express Building for tomorrow

CHPID type OSD functions for VLANs (IEEE 802.1q)	z900 Oct 01	zSeries May 02	z990 Jun 03	Oct 03	z990 z890 May 04	z9 EC Sept 05 z9 BC May 06
OSD (QDIO) – Ethernet only, IEEE 802.1q						
Null tagging Ethernet, ATM Ethernet LANE	z/OS V1.2	>	>	>	>	>
Full VLAN support		Linux	>	>	>	>
Full VLAN support One global VLAN ID for IPv4 One global VLAN ID for IPv6			z/OS V1.5	>	>	>
VLAN support One global VLAN ID for IPv4			z/VM V4.4	>	>	>
Full VLAN for z/OS on z800, z900				z/OS V1.5	>	>
VLAN support One global VLAN ID for IPv6					z/VM V5.1 Sep 04	>

The dates are general availability dates for hardware.

Enhanced performance assists for z/VM V5.2 guests


December 2005

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 EC, z9 BC, z990, z890
 - ▶ 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

Direct 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



Offerings and dates first available:

- Get and getNext requests - zSeries (April 2002)
- dot3StatsTable (RFC 2665) - 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 - z990 (May 2004)
 - ▶ Same support as listed above
 - ▶ Non-QDIO CHPID type OSE
 - ▶ z/OS V1.6

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