

Avoiding the Pitfalls of an OSA-E3 or OSA-E4S Migration

Session 9295

(CHPID Types OSD, OSE, OSM, OSX)

Wednesday, August 10, 2011: 4:30 PM -5:30 PM
Europe 11 (Walt Disney World Dolphin Resort)



Gwendolyn J. Dente (gdente@us.ibm.com)
z/OS Communications Server and Networking Communications Specialist
IBM Advanced Technical Skills (ATS)
Gaithersburg, MD 20879

© Copyright IBM Corporation 2011

- **So you think you know all about the OSA-E3 or OSA-E4S multiport design? Did you know that you may not be able to seamlessly transfer a network design of a pair of two-port adapter cards to the four ports of a single OSA-E3 adapter or even the two ports of an OSA-E4S? Did you know that working with the OSA/SF utility to configure the OSA Port Configuration and the OAT table for CHPID type of OSE requires special consideration? Did you know that new DLC status codes have been introduced to describe OSA errors that you might run into? Wait till you hear about the very common mistakes that you need to avoid when planning and implementing your migration to multiport from the OSA-E2.**
- **This presentation also contains some information on the OSM and OSX CHPID Types for the zEnterprise™ System.**

● **Note: OSA VLANs and ICC (Console) Configuration are not subjects of this presentation.**

- **OSA Adapter Types**
- **OSA-E3 Multiport Layout**
- **OSA-E4s Multiport Layout**
- **IOCDS Samples for OSA-E3 and OSA-E4S (OSD - QDIO)**
- **QDIO Definitions: z/OS VTAM & TCP/IP**
- **IOCDS Samples for OSA-E3 (OSE - non-QDIO)**
- **OSA/SF Definitions in REXX for CHPID Type OSE**
 - Configuration Definitions: OSE (IOAFENET)
 - OSE: OSA/SF OAT Definition
 - z/OS VTAM Definitions for SNA Link Station Architecture (LSA)
 - z/OS TCP/IP Definitions for IP LAN Channel Station (LCS)
- **Appendix: Common Problems Implementing OSA-E3**
- **Appendix: Breaking zEnterprise™ News - OSA-E3 for QDIO CHPID Types of OSM and OSX**
- **Appendix: Breaking zEnterprise™ News - OSA-E4S for QDIO CHPID Types of OSD and OSX**

● **Note: OSA VLANs and ICC (Console) Configuration are not subjects of this presentation.**



OSA Adapters and Features




System z and Available OSA Adapter Types



Feature	Feature Name	Ports	z900 12/00	z990 06/03	z9 EC 09/05	z10 EC 02/08	z196 09/10
OSA-Express							
1364	GbE LX	2	09/04	X	C	N/A	N/A
1365	GbE SX	2	09/04	X	C	N/A	N/A
1366	1000BASE-T	2	N/A	X	C	N/A	N/A
OSA-Express2							
3364	GbE LX	2	N/A	01/05	X	vd Jun09	C
3365	GbE SX	2	N/A	01/05	X	vd Jun09	C
3366	1000BASE-T	2	N/A	05/06	X	vd Dec09	C
3368	10 GbE LR	1	N/A	01/05	X	vd Jun08	C
OSA-Express3							
3362	GbE LX	4	N/A	N/A	N/A	30May08	X
3363	GbE SX	4	N/A	N/A	N/A	30May08	X
3367	1000BASE-T	4	N/A	N/A	N/A	28Oct08	X
3370	10 GbE LR	2	N/A	N/A	N/A	30May08	X
3371	10 GbE SR	2	N/A	N/A	N/A	28Oct08	X
OSA-Express4S							
0404	GbE LX	2	N/A	N/A	N/A	N/A	09/11
0405	GbE SX	2	N/A	N/A	N/A	N/A	09/11
0406	10 GbE LR	1	N/A	N/A	N/A	N/A	09/11
0407	10 GbE SR	1	N/A	N/A	N/A	N/A	09/11

- Implementation Planning:
 - Particularly Important: "Ports per CHPID"
 - Notice OSA-E4S:
 - Gigabit Ethernet
 - ▶ 2 ports per CHPID
 - Notice OSA-E3:
 - Gigabit Ethernet
 - ▶ 2 ports per CHPID
 - 1000Base-T (copper)
 - ▶ 2 ports per CHPID



The configurations for the "multi-port-per-CHPID" adapters require additional planning over the traditional "single-port-per-CHPID" adapters.

© Copyright IBM Corporation 2011

1. This is a subset of a larger set of charts. The arrows indicate the OSA LAN adapters that are engineered with two ports per CHPID.
2. Historically many customers have been dealing with System z LAN adapters (OSA ports) that include only one port per CHPID. With the advent of two ports per CHPID you simply need to examine your OSA design if you are migrating from a single-port-per-CHPID model to a two-port-per-CHPID model.

Open Systems Adapter features on System z (Enterprise Class)



Feature	Feature Name	Ports	z900 12/00	z990 06/03	z9 EC 09/05	z10 EC 02/08	z196 09/10	CHPIDs	Connectors
OSA-Express									
1364	GbE LX	2	09/04	X	C	N / A	N / A	OSD, L2/L3**	LC Duplex
1365	GbE SX	2	09/04	X	C	N / A	N / A	OSD, L2/L3**	LC Duplex
1366	1000BASE-T	2	N / A	X	C	N / A	N / A	OSC, OSD L2/L3, OSE	RJ-45
OSA-Express2									
3364	GbE LX	2	N / A	01/05	X	wd Jun09	C	OSD L2/L3, OSN*	LC Duplex
3365	GbE SX	2	N / A	01/05	X	wd Jun09	C	OSD L2/L3, OSN*	LC Duplex
3366	1000BASE-T	2	N / A	05/06	X	wd Dec09	C	OSC, OSD L2/L3, OSE, OSN*	RJ-45
3368	10 GbE LR	1	N / A	01/05	X	wd Jun08	C	OSD L2/L3**	SC Duplex
OSA-Express3									
3362	GbE LX	4	N / A	N / A	N / A	30May08	X	OSD L2/L3, OSN	LC Duplex
3363	GbE SX	4	N / A	N / A	N / A	30May08	X	OSD L2/L3, OSN	LC Duplex
3367	1000BASE-T	4	N / A	N / A	N / A	28Oct08	X	OSC, OSD L2/L3, OSE, OSN*, OSM***	RJ-45
3370	10 GbE LR	2	N / A	N / A	N / A	30May08	X	OSD L2/L3, OSX***	LC Duplex
3371	10 GbE SR	2	N / A	N / A	N / A	28Oct08	X	OSD L2/L3, OSX***	LC Duplex
OSA-Express4S									
0404	GbE LX	2	N / A	N / A	N / A	N / A	09/11	OSD L2/L3	LC Duplex
0405	GbE SX	2	N / A	N / A	N / A	N / A	09/11	OSD L2/L3	LC Duplex
0406	10 GbE LR	1	N / A	N / A	N / A	N / A	09/11	OSD L2/L3, OSX***	LC Duplex
0407	10 GbE SR	1	N / A	N / A	N / A	N / A	09/11	OSD L2/L3, OSX***	LC Duplex

X = Available for ordering C = Carry forward on an upgrade

* CHPID type OSN is supported on z196, z114, z10, and z9 servers

** L2/L3 = Layer 2/Layer3 which is applicable to z990 and later servers

*** OSM and OSX are exclusive to z196 and z114

**SOD - z196 and z114 are last System z servers to support OSA-Express 2
CHPID Type OSN not offered on OSA-Express4S GbE**

© Copyright IBM Corporation 2011

1. This is the full support chart for the various OSA Types available to the various System z Enterprise Class models.
2. The arrows point to the OSA Express Feature Codes that support Two Ports per CHPID.
3. Advice about OSN on z196 and z114: An OSA-E3 can only be ordered if there are empty slots in an I/O cage or I/O drawer and the PCIe I/O drawer is fully populated. OSA-E3 can also be carried forward to zEnterprise.

Open Systems Adapter features on System z (Business Class)



Feature	Feature Name	Ports	z800 03/02	z890 05/04	z9 BC 05/06	z10 BC 10/08	z114 09/11	CHPIDs
OSA-Express								
1364	GbE LX	2	09/04	X	C	N/A	N/A	OSD, L2/L3 **
1365	GbE SX	2	09/04	X	C	N/A	N/A	OSD, L2/L3 **
1366	1000BASE-T	2	N/A	X	C	N/A	N/A	OSC, OSD L2/L3, OSE
OSA-Express2								
3364	GbE LX	2	N/A	01/05	X	wd Jun09	C	OSD L2/L3, OSN *
3365	GbE SX	2	N/A	01/05	X	wd Jun09	C	OSD L2/L3, OSN *
3366	1000BASE-T	2	N/A	05/06	X	wd Dec09	C	OSC, OSD L2/L3, OSE, OSN *
3368	10 GbE LR	1	N/A	01/05	X	C	N / A	OSD L2/L3 **
OSA-Express3								
3362	GbE LX	4	N/A	N/A	N/A	X	C, RPQ	OSD L2/L3, OSN
3363	GbE SX	4	N/A	N/A	N/A	X	C, RPQ	OSD L2/L3, OSN
3367	1000BASE-T	4	N/A	N/A	N/A	X	X	OSC, OSD L2/L3, OSE, OSN*, OSM***
3369	2P 1000BASE-T	2	N/A	N/A	N/A	X	X	OSC, OSD L2/L3, OSE, OSN*, OSM***
3370	10 GbE LR	2	N/A	N/A	N/A	X	C, RPQ	OSD L2/L3
3371	10 GbE SR	2	N/A	N/A	N/A	X	C, RPQ	OSD L2/L3
3373	2P GbE SX	2	N/A	N/A	N/A	X	C, RPQ	OSD L2/L3, OSN
OSA-Express4S								
0404	GbE LX	2	N/A	N/A	N/A	N/A	X	OSD L2/L3
0405	GbE SX	2	N/A	N/A	N/A	N/A	X	OSD L2/L3
0406	10 GbE LR	1	N/A	N/A	N/A	N/A	X	OSD L2/L3, OSX***
0407	10 GbE SR	1	N/A	N/A	N/A	N/A	X	OSD L2/L3, OSX***



X = Available for ordering C = Carry forward on an upgrade
*** CHPID type OSN is supported on z196, z114, z10, and z9 servers**
**** L2/L3 = Layer 2/Layer3 which is applicable to z990 and later servers**
***** OSM and OSX are exclusive to z196 and z114**

**SOD - z196 and z114 are last System z servers to support OSA-Express 2
 CHPID Type OSN not offered on OSA-Express4S GbE**

© Copyright IBM Corporation 2011

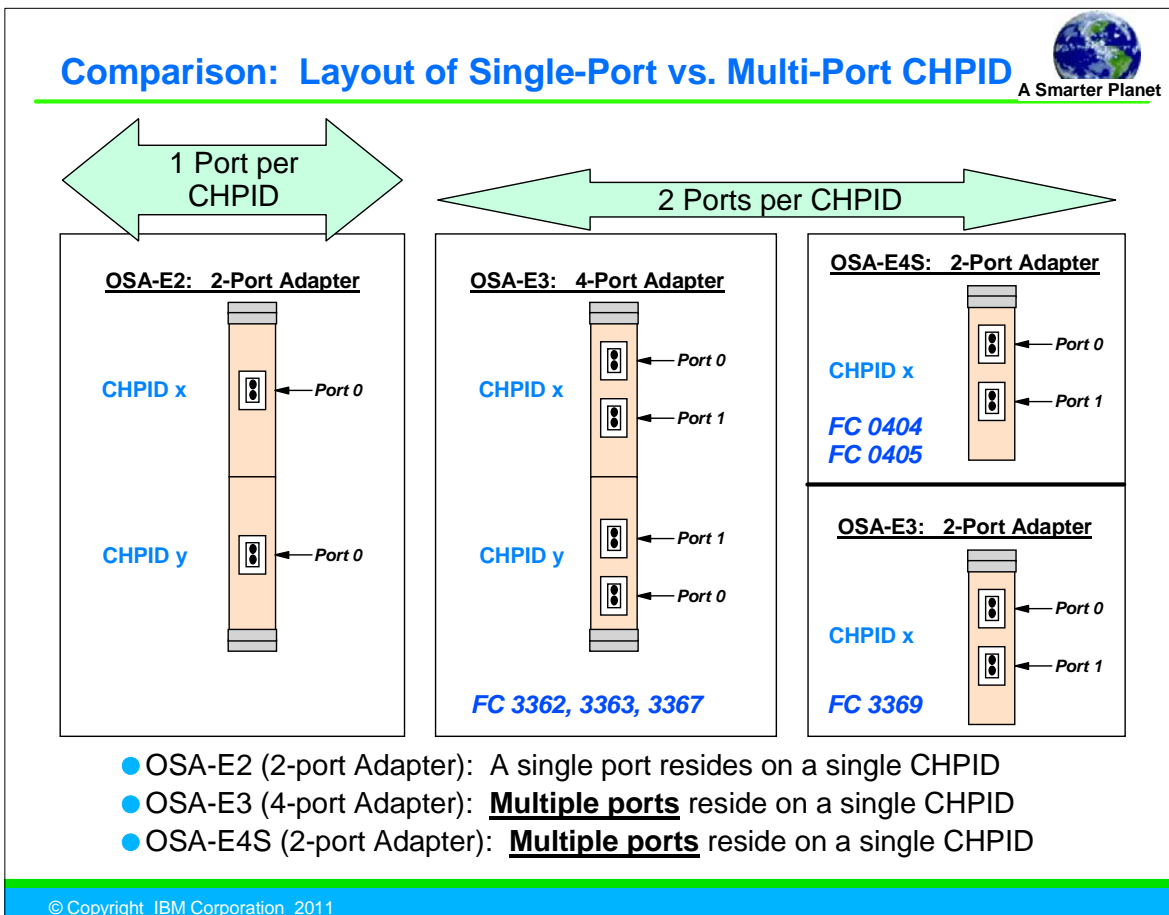
1. This is the full support chart for the various OSA Types available to the various System z Enterprise Class models.
2. The arrows point to the OSA Express Feature Codes that support Two Ports per CHPID.
3. Advice about OSN on z196 and z114: An OSA-E3 can only be ordered if there are empty slots in an I/O cage or I/O drawer and the PCIe I/O drawer is fully populated. OSA-E3 can also be carried forward to zEnterprise.



Layout of Multiport OSA-E3 and OSA-E4S



Comparison: Layout of Single-Port vs. Multi-Port CHPID



1. Each port of the OSA-Express2 adapter depicted is on a separate CHPID: CHPID x and CHPID y. Each port on each CHPID is defined with a separate portname and resides at port number 0.
2. This OSA-Express3 (either a Gigabit Ethernet or a 1000Base-T) is engineered with two ports on each CHPID: CHPID x and CHPID y. The two ports on each CHPID are numbered port 0 and port 1. But note how the top half of the OSA-E3 4-port card is the mirror image of the bottom half with regard to the port number assignments; reading from top to bottom you see Port 0, Port 1, Port 1, Port 0. As with any OSA port, the portnames on the multi-port OSA-E3 must be unique to a CHPID; an explanation of this portname assignment is provided in a later paragraph of this section.
3. The OSA-Express4S (GbE) is engineered with two ports on each CHPID: CHPID x and CHPID y. As with any OSA port, the portnames on the multi-port OSA-E4S must be unique to a CHPID; an explanation of this portname assignment is provided in a later paragraph of this section.

Portnames: A Pitfall (Must Be Unique Names!)



OSA-E2: 2-Port Adapter

- On an OSA CHPID, the Portname value must be unique to the CHPID.
- This example depicts a single port per CHPID, as in the design of an OSA-E2.
 - The Portnames are not only unique to the CHPID but also different from each other (GIGx and GIGy)..
 - However, **certain configurations would permit the Portnames to be the same as in "GIG0."**
 - Example: If different VTAMs control the OSA TRLE definitions, the Portnames could be the same (e.g., GIG0) across the two CHPIDs.

OSA-E3: 4-Port Adapter

- On an OSA CHPID, the Portname value must be unique to the CHPID.
- This example depicts multiple ports per CHPID, as in the design of an OSA-E3 or OSA-E4S.
 - The Portnames in the top left of the visual are not only unique to the CHPID but also different from each other: "GIG0x" and "GIG1x."
 - No configuration can allow two OSA ports on the same CHPID to be assigned the same Portname.
 - Example: The Portnames in the bottom half of the visual must bear unique portnames. Otherwise, one port will fail to activate.

OSA-E4S: 2-Port Adapter

OSA-E3: 2-Port Adapter

© Copyright IBM Corporation 2011


1. Port Names must be unique per CHPID and also unique per controlling VTAM.

Portname Assignments: A Solution




RECOMMENDATION for PORTNAME and QDIO: Use unique name for each OSA "PORTNAME" on a CHPID, even though some exceptions to this recommendation are valid..


OSA-E2



Unique names on each CHPID across all operating systems on platform




A shared OSA port must have the same name in all sharing LPARs (Exception: Portname independence for some operating systems)




Unique names on each CHPID but same name because different operating systems own each CHPID.


OSA-E3, E4S



Unique names on each CHPID across all operating systems on platform



Unique name on each CHPID but same name because different operating systems own each CHPID.



Portnames must be unique on a CHPID regardless of controlling operating system.

If you use the same portname for separate ports on same CHPID of an OSA-E3, you can get in trouble:

Message EZZ4310I,
Code 8010302C,
Diagnostic Code
8010311B

Compare diagrams 3 and 6

1. Note how each of the configurations (1 through 6) assigns a unique QDIO Portname to each OSA Port on a CHPID, However, with the multiport design of the OSA-E3, you must now ensure that each of two ports on a CHPID is assigned a unique name on that CHPID.
2. In migrating from OSA-E2 QDIO Configuration3 TO OSA-E3 QDIO Configuration 6, you can no longer use the same portname. Each portname on a CHPID must be unique!

OSA-E3 multiport activation failure: EZZ4310I, 8010302C or 8010311B

Item: RTA000190883 Last Updated: 01/26/2010

Source: US CSYS Created Date: 01/22/2010

Question & Answer Text +Topic Thread:

- COMMUNICATION SERVER FOR OS/390 AND z/OS IP

Q:The second port of a multiport OSA-E3 fails to activate and I receive EZZ4310I with error of 8010302C and a diagnostic code of 03. LPAR1 and VTAM1 own the OSA Portnum 0 on CHPID F7; LPAR2 and VTAM2 own the OSA Portnum 1 on CHPID F7. Although the CHPID is shared across the two LPARS, the OSA ports are not shared, and I therefore named them the same name: PORTNAME GIGBIT. My microcode is uplevel; my operating system code is uplevel. What could be going on here?

A:Because of the shared CHPID, and even though the OSA ports themselves are unique and attached to different partitions and VTAMs, you must provide a unique name to each PORT on the same CHPID. You might choose GIGBIT0 and GIGBIT1 or something similar. Note that with an OSA-E2, where each OSA port is attached to a different CHPID, your configuration could have worked given the attachment to two separate LPARs and VTAMs. With the OSA-E3 you must ensure that ports on the same CHPID have different names from each other even when being activated by separate VTAMs and IP Stacks.

At V1R9 of z/OS, the status code that is produced is 8010302C, but at **z/OS V1R10** the status code is more granular: 8010311B. See the wording of status code 8010311B:

"Duplicate port name

Explanation: An attempt was made to activate an OSA-Express3 port in QDIO mode. The port name that was used for this activation attempt was already in use on the other port that belongs to that channel path ID (CHPID). Two ports on the same CHPID cannot have the same port name."

The z/OS V1R10 SNA IP Sense Codes manual (SC31-8791) was amended with APAR OA25064.

The new status code is in the base of the V1R11 manual.

Search - keywords:

OSA-E3 multiport OSA EZZ4310I 8010302C 8010311B

- http://www-01.ibm.com/support/docview.wss?rs=852&context=SSSN3L&q1=311B&uid=swg21376527&loc=en_US&cs=utf-8&lang=en from March 2009

- =====

- The following codes were added to Chapter 3 Data link control (DLC) status codes in the Bytes 2 and 3 (completion code) of the DLC status code table:

- **Hexadecimal Code Meaning**

- **X'31nn'** OSA-Express rejection of an attempt to activate a port

- Explanation: Codes that begin with X'31' are issued only when there is an attempt to activate a port in OSA-Express QDIO Mode.
- These codes indicate that the OSA adapter has rejected an activation attempt. The value nn indicates the reason for the rejection. Specific nn codes are listed in this table. If you receive a code that is not listed in this table, contact IBM Service.

- **X'311B'** Duplicate port name

- Explanation: An attempt was made to activate an OSA-Express3 port in QDIO mode. The port name that was used for this activation attempt was already in use on the other port that belongs to that channel path ID (CHPID). Two ports on the same CHPID cannot have the same port name.

- **X'3150'** Incorrect port name

- Explanation: An attempt was made to activate an OSA-Express port in QDIO mode. The port name that was used for this activation attempt did not match the port name that was already assigned to this port by a previous user. All z/OS users sharing that port must activate that port with the same port name.


1. This is a z/OS Communications Server Technote that you can access at URL:
http://www-01.ibm.com/support/docview.wss?rs=852&context=SSSN3L&q1=311B&uid=swg21376527&loc=en_US&cs=utf-8&lang=en

Notes on "Port Name Relief" for z/VM and Linux on z



- From SA22-7935 OSA Cust Guide and Reference

Port Name on z/VM and Linux on z no longer verified at OSA Port activation. (See specifics below.)



● Port name relief

- For OSA-Express features running in QDIO mode, the port name identifies the OSA port for sharing by other operating system instances.
- **OLD RULE:** When the port name is defined, all operating system instances that share the port must use the same port name.
- **NEW RULE:** Beginning with a particular level of OSA-Express Licensed Internal Code (LIC) and some operating systems, this restriction has been lifted. A port name is not required at the following OSA-Express LIC levels **and** operating system levels:
 - z800 and z900 with OSA-Express (LIC) level September 2003 or higher
 - z890 and z990 with all levels of OSA-Express LIC levels
 - z9 EC
 - z10
 - z/VM 4.3 with APAR PQ73878 and z/VM 4.4
 - Linux – V2.4 kernel June 2003 stream and above – V2.6 kernel April 2004 stream and above
- z/OS, VSE/ESA, and TPF require a port name, which must be the same when the OSA port is shared with the same operating system in different images or in a mixed OS environment. In addition, the port name must match the device name in the z/OS TCP/IP profile [or must match the PORTNAME parameter of the INTERFACE statement].

© Copyright IBM Corporation 2011

1. PORTNAME relief is applicable to all of the OSA-Express features in QDIO mode (OSD CHPID type).
2. For PORTNAME relief, z/VM V4.3 with APARs VM63308 (PORTNAME relief for Guest LAN) and PQ73878 (VM TCP/IP exploitation for the PORTNAME relief) are required; PORTNAME relief is in the z/VM V4.4 base.
3. Linux on zSeries support was delivered as part of the "June 2003 stream" at:
<http://www10.software.ibm.com/developerworks/opensource/linux390>
4. Existing Linux distributions and z/VM releases without support for PORTNAME relief continue to run on this MCL as before.
5. Treat the last sentence of the visual with an open mind ... the INTERFACE Definition on z/OS actually defines the PORTNAME parameter and uses a different INTERFACE Name.

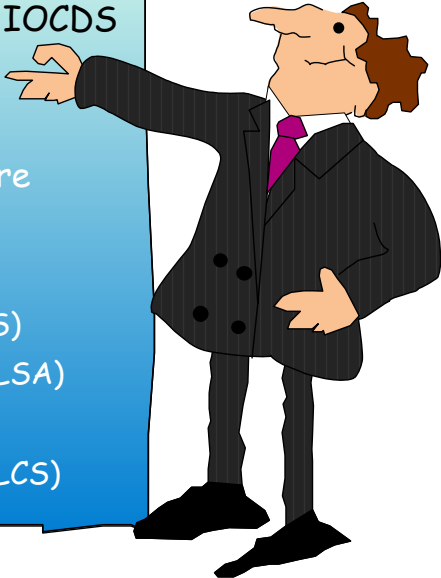


IOCDs Samples for Multiport OSA-E3 or OSA-E4S (OSD - QDIO)



For QDIO (IP): 3 devices per connection
VTAM requires an even-numbered device address for the READ device and an
odd-numbered device address for the WRITE device.
OSA microcode itself no longer imposes this restriction.

1. The OSA-E microcode no longer requires that the READ device be assigned to an Even-numbered device address. The only requirement is that the WRITE device be assigned to the device address that represents $\langle \text{read_device_address} + 1 \rangle$.
 1. z/VM and zLinux have no problems assigning an odd-numbered address to the READ device.
2. However, VTAM for z/OS does require that the READ device be assigned to an even-numbered device address and that the WRITE device be assigned to an odd-numbered device address.



- **No!**
- Port Numbers are not assigned in the IOCDs of System z.
- Port numbers are assigned in
 - VTAM
 - QDIO (z/OS)
 - non-QDIO (LSA)
 - TCP/IP
 - non-QDIO (LCS)

© Copyright IBM Corporation 2011

1. IOCDs DOES NOT CARE ABOUT THE PORT ASSIGNMENTS. ONLY VTAM or TCP/IP or SNA networking definitions care about the port assignments.
 1. In z/OS QDIO for TCP/IP, the TRLE definition in VTAM provides a PORTNUM parameter associated with a PORTNAME parameter.
 1. The PORTNAME parameter is then used as the DEVICE name on an IPv4 interface definition in z/OS.
 2. If deploying a QDIO interface with the INTERFACE definition of z/OS TCP/IP, you will find the same VTAM PORTNAME on the INTERFACE statement.
 3. In z/OS non-QDIO for SNA, you must define an XCA Major Node definition that includes a port number.
 2. In z/OS non-QDIO for TCP/IP, you must define a DEVICE and LINK statement for Lan Channel Station (LCS). The port number of the OSA-Express3 is assigned in this LCS definition.
2. NOTE about other Operating Systems:
 1. z/VSE, Linux on z, TPF, z/VM also have interface definitions with parameters to assign device addresses to port number 0 or 1; please consult the appropriate manuals or device driver descriptions for these operating systems.

Assigning Device Addresses to Port 1: Alternative 1



Alt. 1: Keep the IOCDS the same.

<p>OSA-E2: 2-Port Adapter</p> <p>CHPID x ← Port 0 :IODEVICE 2080 - 208E:</p> <p>CHPID y ← Port 0</p>	<p>OSA-E3: 4-Port Adapter</p> <p>CHPID x ← Port 0 :IODEVICE 2080 - 2089: ← Port 1 :IODEVICE 208A - 208E:</p> <p>CHPID y ← Port 1 ← Port 0</p> <p>FC 3362, 3363, 3367</p>	<p>OSA-E4S: 2-Port Adapter</p> <p>CHPID x ← Port 0 :2080 - 2089: ← Port 1 :208A - 208E: FC 0404 FC 0405</p> <p>OSA-E3: 2-Port Adapter</p> <p>CHPID x ← Port 0 :2080 - 2089: ← Port 1 :208A - 208E: FC 3369</p>
---	---	--

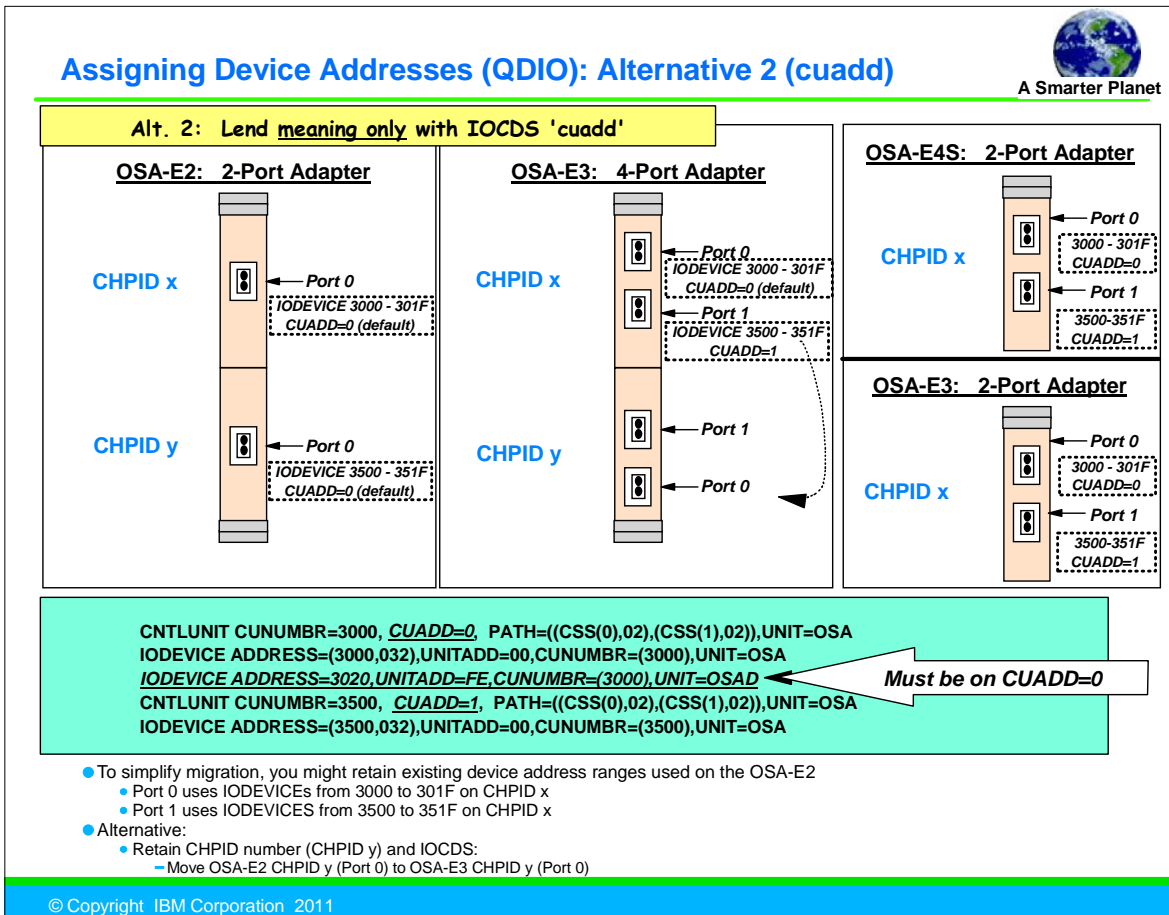
**CNTLUNIT CUNUMBR=2080, PATH=((CSS(2),02)), UNIT=OSA
IODEVICE ADDRESS=(2080,015), CUNUMBR=(2080), UNIT=OSA**

- Migration from OSA-E2 (2-port Adapter) to an OSA with 2-Port-per-CHPID (IOCDS Considerations):
 - You might leave all ports on a single CHPID within the same overall device range
 - Port 0 uses IODEVICEs from 2080 to 2089 on CHPID x
 - ▶ Assigned in the Operating System or Communications Access Method Configuration
 - Port 1 uses IODEVICEs from 208A to 208E on CHPID x
 - ▶ Assigned in the Operating System or Communications Access Method Configuration

© Copyright IBM Corporation 2011

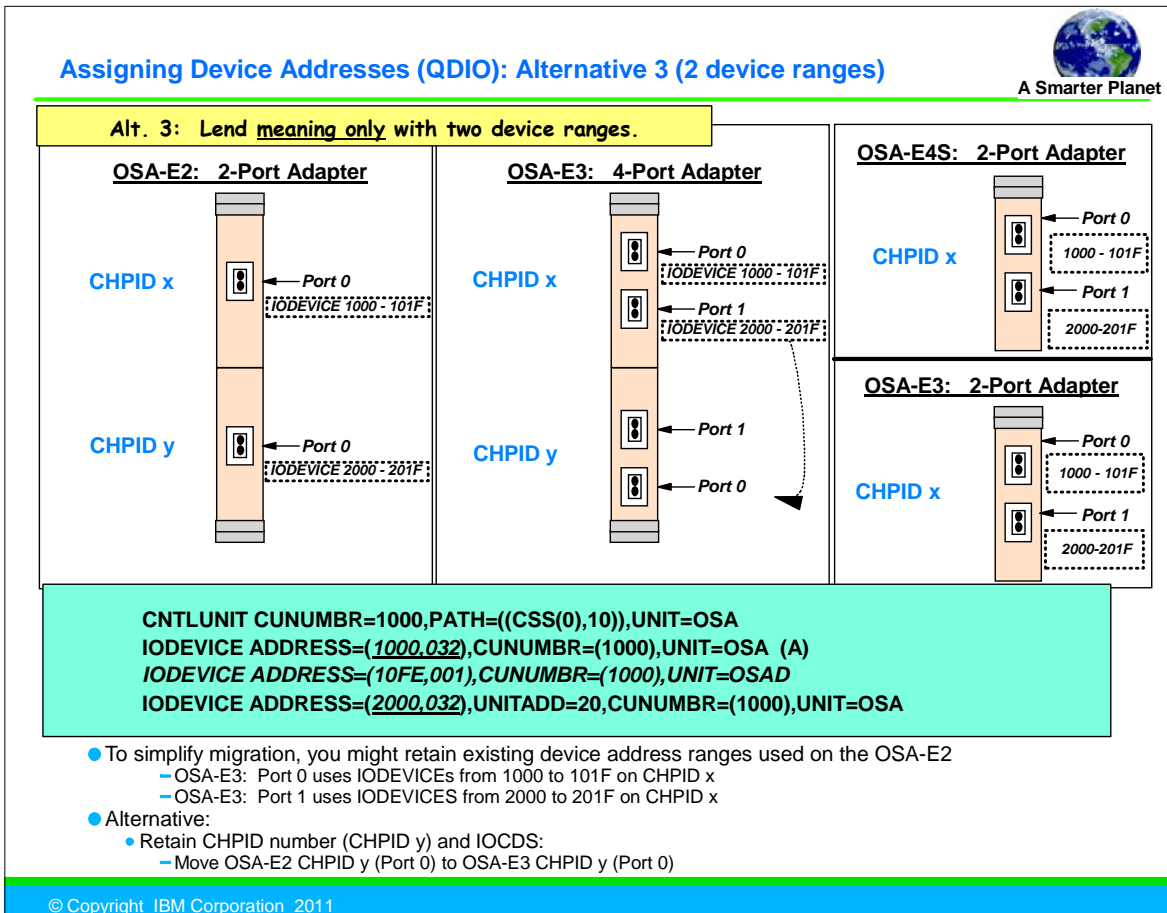
1. In migrating to a multiport-per-chpid OSA-Express3 or OSA-Express4S you can choose to use an IOCDS similar to what you already have for an OSA-E2, and simply spread the device addresses a single address range across two ports of the same CHPID that originally connected to only one port.
2. This is part of the IOCDS that was already in place for a Gigabit OSA-E2. We have allocated fifteen addresses (2080-208E) to QDIO connections starting with device address 2080.
3. We use the same IOCDS for the multiport models on the right, but now we split the assigned range of 15 addresses into two ranges
4. IOCP definitions have no awareness of the OSA adapter's two ports and simply assign device addresses; the Operating System or VTAM definition for z/OS do care about the port numbers and maps the number to the addresses to a Port Number of 0 or 1.
5. Or you can choose to change your IOCDS to reflect separate address ranges or, for QDIO, even separate logical control units, despite the presence of only a single physical control unit on the CHPID. We now want to show you a couple of different ways to implement an IOCDS for an OSA-Express3 implementation.

Assigning Device Addresses (QDIO): Alternative 2 (cuadd)

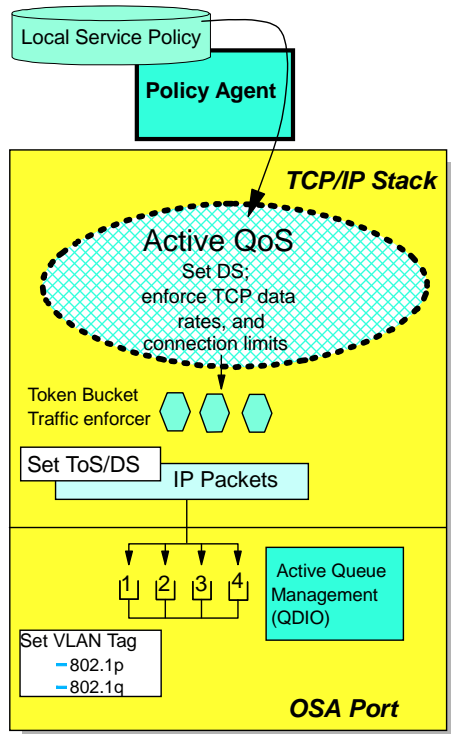


- The examples show you how to make a logical distinction within the IOCP between ports 0 and 1 of a multiport-per-CHPID OSA. (This is a logical distinction only, since the IOCDs itself does not assign the device numbers to the individual OSA port numbers. Ultimately, it is still VTAM that makes that assignment and not the IOCDs.)
 - You can specify separate logical control units with the CUADD parameter. While in the past defining multiple Control Units had value only if defining many devices, it appears customers migrating from OSA-Express2 channels to multi-port OSA-Express3s are finding it is an easy way to combine two OSA-Express CHPIDs into the two ports of an OSA-Express3 CHPID.
- Migration from OSA-E2 (2-port Adapter) to a multiport-per-CHPID OSA.:
 - To simplify migration, you might retain existing device address ranges used on the OSA-E2
 - OSA-E3 or OSA-E4S: Port 0 uses IODEVICES from 3000 to 301F on CHPID x
 - That is, OSA-E2 CHPID x (Port 0) moves to OSA-E3 CHPID x (Port 0, CUADD=0)
 - Assign Logical Control Unit Address (only possible with QDIO Implementations)
 - OSA-E3 or OSA-E4S: Port 1 uses IODEVICES from 3500 to 351F on CHPID x
 - That is, OSA-E2 CHPID y (Port 0) moves to OSA-E3 CHPID x (Port 1, CUADD=1)
 - Assign Logical Control Unit Address (only possible with QDIO Implementations)
 - Alternative:
 - Retain CHPID number (CHPID y) and IOCDs:
 - Move OSA-E2 CHPID y (Port 0) to OSA-E3 or OSA-E4S CHPID y (Port 0)
 - If you are exploiting OSA/SF, the OSA/SF definition (UNIT=OSAD) must be defined during OSA/SF setup in the CUADD=0 range. See SA22-7935-11 .. Customer Guide and Reference Chapter 2, Configuring an OSA using HCD
 - "For a QDIO CHPID (OSD) you must define the OSAD device on logical control unit X'00' when multiple control units are defined or OSA/SF will not recognize the OSAD device."
 - IMPORTANT: OSA/SF device must be associated with CUADD=0 and is not valid for an OSX CHPID.
- Another Advantage of CUADD approach:
 - If you need to define more than 254 devices for an unshared OSD channel path, multiple control units must be defined.
 - Specify a unique logical address for each control unit using the CUADD keyword.

Assigning Device Addresses (QDIO): Alternative 3 (2 device ranges)



- An alternative to the coding scheme with 2 CUADDs, you can also use a different address for each of the two ports. Such a scheme might make problem determination and operator procedures easier for you, as message displays very clearly show the distinction between the two ports, although they reside on the same CHPID.
 - Again, the IOCDS does NOT CARE about port numbers ... it is the VTAM or the TCPIP or SNA definition that decides what is on port 0 or port 1. So if you want to be able to look at an IOCDS and determine which device numbers are for Port 0 and Port 1, you can play some tricks with the IOCDS to create a convention that lets someone realize: Oh, this is now this maps to the VTAM and the TCPIP/SNA
- Migration from OSA-E2 (2-port Adapter) to OSA-E3 or OSA-E4S:
- To simplify migration, you might retain existing device address ranges used on the OSA-E2
 - OSA-E3 or OSA-E4S: Port 0 uses IODEVICEs from 1000 to 101F on CHPID x
 - That is, OSA-E2 CHPID x (Port 0) moves to OSA-E3/OSA-E4S CHPID x (Port 0)
 - OSA-E3 or OSA-E4S: Port 1 uses IODEVICEs from 2000 to 201F on CHPID x
 - That is, OSA-E2 CHPID y (Port 0) moves to OSA-E3/OSA-E4S CHPID x (Port 1)
- Alternative:
 - Retain CHPID number (CHPID y) and IOCDS:
 - Move OSA-E2 CHPID y (Port 0) to OSA-E3/OSA-E4S CHPID y (Port 0)
- If you are exploiting OSA/SF, the OSA/SF definition (UNIT=OSAD) must be defined during OSA/SF setup in the CUADD=0 range.



● Outbound Priority Queuing

- OSA maintains 4 outbound priority queues
 - Applications assigned to queues based on **Policy Agent QoS Policy**, or
 - By **application definitions** (as with Enterprise Extender), or
 - By **WLM and IPCONFIG WLM PRIORITYQ** (in V1R11)
- OSA supports Outbound Priority Queuing (multiple Outbound Queues) as long as no more than 480 valid subchannels are defined for all LPARs sharing a CHPID.
- If your definition requires more than 480 valid subchannels (with a maximum of 1920), then the user must explicitly turn off Outbound Priority Queuing on the CHPID definition by specifying CHPARM=02 in the IOCP or by specifying it in HCD.
- **Inbound Priority Queuing**
 - At z/OS V1R12 and V1R13.

© Copyright IBM Corporation 2011

1. The diagram above depicts some general facts about Policy Agent.
2. The following is an overall description of Policy Agent components.
3. The policy agent (pagent) is responsible for reading policy definitions from a configuration file (pagent.conf) and/or from an LDAP server and installing them into one or more TCP/IP stacks.
4. The flat policy file can be located in any MVS dataset or HFS file. A search order is supported, but for simplicity, the file will be referred to as pagent.conf in this presentation.
5. The easiest way to implement PAGENT is with a set of flat files that define policies. The easiest way to code the policies is to use the z/OS IBM Configuration Assistant GUI.
6. Note, if multiple TCP/IP stacks are running on one MVS image, only one policy agent is used. The configuration for different stacks can be in different files, with a main file that uses the Tcplmage statement to identify the other files.
7. Policy Agent in CS for z/OS can define policies for:
 8. Automatic QOS enforcement
 9. Quality of Service for setting precedence bits in the IP Packet header
10. OSA supports Outbound Priority Queuing (multiple Outbound Queues) as long as no more than 480 valid subchannels are defined for all LPARs sharing a CHPID.
11. Each logical partition sharing a CHPID gets a subchannel for every device defined on that CHPID.
12. Therefore, if you define a CHPID shared by 15 logical partitions and define 32 devices (either on one port or across two ports), you have used 480 valid subchannels ($15 * 32 = 480$).
13. Default is to enable Outbound Priority Queuing
14. If your definition requires more than 480 valid subchannels (with a maximum of 1920), then the user must explicitly turn off Outbound Priority Queuing on the CHPID definition by specifying CHPARM=02 in the IOCP or by specifying it in HCD.
15. HCD will prevent a device definition that will cause the 480 subchannel limit to be broken. IOCP will issue an error message and not create an IOCDs if the limit is broken.
16. Inbound Priority Queuing
 1. At z/OS V1R12: Bulk/Streaming, Sysplex Distributor, Other
 2. At z/OS V1R13: All above plus Enterprise Extender



QDIO Definitions: z/OS VTAM & TCP/IP

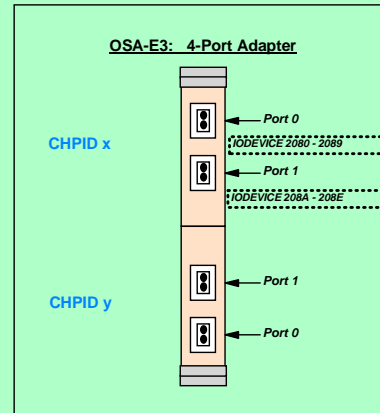


CHPID Type=OSD: VTAM TRLE



```
OSA2080 VBUILD TYPE=TRL
OSA2080P TRLE LNCTL=MPC,
              READ=2080,
              WRITE=2081,
              DATAPATH=(2082-2089),
              PORTNAME=OSA2080,
              PORTNUM=0,
              MPCLEVEL=QDIO
```

```
OSA208A VBUILD TYPE=TRL
OSA208AP TRLE LNCTL=MPC,
              READ=208A,
              WRITE=208B,
              DATAPATH=(208C-208E),
              PORTNAME=OSA208A,
              PORTNUM=1,
              MPCLEVEL=QDIO
```



© Copyright IBM Corporation 2011

1. The IOCDs knows nothing of the actual OSA Port number.
2. But the VTAM assigns the PORT Number to the range of device addresses.
3. Here you see how we have specified that PORTNUM 0 is for the IODEVICE range indicated, whereas PORTNUM 1 is for the other range.

QDIO: TCP/IP Profile QDIO Device / Link or INTERFACE



A Smarter Planet

```
DEVICE OSA2080 MPCIPA
LINK  OSA2080L IPAQENET OSA2080 VLANID 10
;
DEVICE OSA208A MPCIPA
LINK  OSA208A IPAQENET OSA208A VLANID 11
;
```

```
OSA2080 VBUILD TYPE=TRL
OSA2080P TRLE LNCTL=MPC,
          READ=2080,
          WRITE=2081,
          DATAPATH=(2082-2089),
          PORTNAME=OSA2080,
          PORTNUM=0,
          MPCLEVEL=QDIO
```

```
OSA208A VBUILD TYPE=TRL
OSA208AP TRLE LNCTL=MPC,
          READ=208A,
          WRITE=208B,
          DATAPATH=(208C-208E),
          PORTNAME=OSA208A,
          PORTNUM=1,
          MPCLEVEL=QDIO
```

```
INTERFACE OSA2080I
DEFINE IPAQENET
PORTNAME OSA2080
IPADDR 10.1.2.12/24
MTU <value>
VLANID 10
```

```
Advanced Parameters
[INBPERF DYNAMIC WORKLOADQ]
VMAC [ <value> ] [ROUTEALL]
[SOURCEVIPAIN T VIPA2L]
```

```
INTERFACE OSA208AI
DEFINE IPAQENET
PORTNAME OSA208A
IPADDR 10.1.2.14/24
MTU <value>
VLANID 11
```

```
Advanced Parameters
[INBPERF DYNAMIC WORKLOADQ]
VMAC [ <value> ] [ROUTEALL]
[SOURCEVIPAIN T VIPA2L]
```

© Copyright IBM Corporation 2011

1. VTAM assigns the PORT Number to the range of device addresses.
2. The TCP/IP Profile in z/OS then defines the IP devices and links (or the INTERFACES) by pointing in some fashion to the VTAM PORTNAME.
 1. The DEVICE NAME must be the same as the TRL PORTNAME, or
 2. If using INTERFACE, then the PORTNAME parameter is specified in the INTERFACE coding.
3. In the Interface definitions we also show you some "Advanced Parameters" that you might consider coding as well:
 1. Determine how the INBPERF values can benefit performance in your environment.
 2. The VMAC setting is particularly of value if you want to virtualize the OSA port so that each connection assumes a different physical identity. You can allow the system to generate the VMAC value or you can establish a new value that does not change with each cycle of TCP/IP. If you use VMAC, you can eliminate the confusion that arises with trying to determine whether you should code PRIROUTER, SECROUTER, etc. But, if you are sharing the same Interface and/or VLANs with multiple LPARs, you can still code PRIROUTER, SECROUTER, etc. if you desire unknown destination IP addresses to be sent to one of the LPARs for further routing. Please see the IP Configuration Guide and the IP Configuration Reference for more detail about routing packets inbound through the OSA port.
 3. If you want to point to a SOURCEVIPA on the Interface Definition, you must code SOURCEVIPAIN. (Interface does not take its SOURCEVIPA from the Home List sequence since there is no Home List entry for the IP Addresses of an Interface statement.)

QDIO: Where's My Source VIPA Now?



Source IP Address Selection

- **IPv6:** Sendmsg() using the IPV6_PKTINFO ancillary option specifying a nonzero source address (RAW and UDP sockets only)
- **IPv6:** Setsockopt() IPV6_PKTINFO option specifying a nonzero source address (RAW and UDP sockets only)
- **IPv4 and IPv6:** Explicit bind to a specific local IP address
- **IPv6:** bind2addrsel socket function (AF_INET6 sockets only)
- **IPv4 and IPv6:** PORT profile statement with the BIND parameter
- **IPv4 and IPv6:** SRCIP profile statement (TCP connections only)
 - a. JOBNAME entries, other than JOBNAME *
 - b. DESTINATION entries
 - c. JOBNAME * entries
- **IPv4 and IPv6:** TCPSTACKSOURCEVIPA parameter on the IPCONFIG or IPCONFIG6 profile statement (TCP connections only)
 - Recommend use of SRCIP block over TCPSTACKSOURCEVIPA
- **IPv4 and IPv6:** SOURCEVIPA -- Static VIPA address from the HOME list or from the SOURCEVIPAINTERFACE parameter
- **IPv4 and IPv6:** HOME IP address of the link over which the packet is sent

```
;
INTERFACE OSA2080I
DEFINE IPAQENET
PORTNAME OSA2080
IPADDR 10.1.2.12/24
MTU <value>
VLANID 10
[INBPERF DYNAMIC WORKLOADQ]
VMAC [ <value> ] [ROUTEALL]
SOURCEVIPAINTE VIPA2L
```

```
;
INTERFACE OSA208AI
DEFINE IPAQENET
PORTNAME OSA208A
IPADDR 10.1.2.14/24
MTU <value>
VLANID 11
[INBPERF DYNAMIC WORKLOADQ]
VMAC [ <value> ] [ROUTEALL]
SOURCEVIPAINTE VIPA2L
```

© Copyright IBM Corporation 2011

1. From V1R12 IP Configuration Guide:
 1. For a TCP connection, the source address is selected for the initial outbound packet, and the same source IP address is used for the life of the connection.
 2. For the UDP and RAW protocols, a source IP address selection is made for each outbound packet. A more detailed description of this sequence follows.
2. There are many different ways to choose a SOURCE IP address for outbound-initiated connections.
3. One common way is to assign a SOURCEVIPA address. In moving from a DEVICE/LINK definition to an INTERFACE Definition, the "traditional" sourcevipa based upon the HOME LIST sequence is no longer valid. And so please use the SOURCEVIPAINTE parameter on the INTERFACE statement.
4. IPCONFIG SOURCEVIPA
 1. From Home List for interfaces defined with DEVICE/LINK
 2. From SOURCEVIPAINTE for interfaces defined with INTERFACE
5. IPCONFIG6 SOURCEVIPA
 1. From SOURCEVIPAINTE for interfaces defined with INTERFACE

OAT Registration for QDIO: Different Behaviors



Single Connection per Stack per Subnet (2)

Connection Definition Type	Subnet coded?	Basic definition (no VMACs)	VMAC ROUTEALL configured? <i>For routing: IP addresses are ignored and routing is based solely on VMAC</i>	VMAC ROUTELCL configured? <i>For routing: IP addresses used for routing once packet has reached the VMAC</i>
DEVICE and LINK (2)	N/A	Most active IP addresses in HOME List -- OSA (3) -- CTC -- XCF -- HiperSockets, etc. -- All VIPAs	V1R11 & Prior: Most active IP addrs. in HOME List: (1) -- OSA (3) -- CTC -- XCF -- HiperSockets, etc. -- All VIPAs	V1R12 & Later (6): --OSA IP address (3) --All VIPAs (4)
INTERFACE (2)	No	Most active IP addresses in HOME List -- OSA (3) -- CTC -- XCF -- HiperSockets, etc. -- All VIPAs	-- OSA IP address (3) -- All VIPAs (4)	Most active IP addresses in HOME List -- OSA (3) -- CTC -- XCF -- HiperSockets, etc. -- All VIPAs
INTERFACE (2)	Yes	Most active IP addresses in HOME List -- OSA (3) -- CTC -- XCF -- HiperSockets, etc. -- VIPAs in same subnet as OSA (4) -- All other VIPAs (5)	-- OSA IP address (3) -- VIPAs in same subnet as OSA (4)	Most active IP addresses in HOME List -- OSA (3) -- CTC -- XCF -- HiperSockets, etc. -- VIPAs in same subnet as OSA (4) -- All other VIPAs (5)

General NOTES: z/OS registers IP addresses for two purposes: for inbound routing and for ARP offload. The OSA/SF "GET OAT" function displays only addresses that have been registered for ARP offload purposes. The coding of VLANID does not influence ARP registration; it does, however, influence routing; if VLAN is coded on the z/OS or z/VM stack, only correctly tagged packets are considered for routing, after which the IP address will be considered for VMAC ROUTELCL and ignored for VMAC ROUTEALL.

Note (1): Since routing in this case is based solely on VMAC, the registration of all IP addresses for routing purposes is superfluous. If ROUTEALL is coded, the OSA will route any received packets to this stack if the destination MAC address matches the coded VMAC, regardless of which addresses are registered.

Note (2): If a stack has multiple connections to the same IP subnet or VLAN, the VIPAs are registered for ARP offload purposes with only one of the multiple OSA connections.

Note (3): If ARP takeover is in effect, the OSA address that we have taken over is also registered for ARP offload purposes

Note (4): Registered for ARP offload purposes (sends gratuitous ARPs and responds to ARP requests)

Note (5): Registered for inbound routing

Note (6): The change in V1R12 avoids registrations for inbound routing which serve no purpose

Powerful QDIO Displays (D TCPIP,,OSAINFO)



```
DISPLAY TCPIP,,OSAINFO,INTFNAME=LNK29D,MAX=500.
EZD0031I TCP/IP CS V1R12 TCPIP Name: TCPCS 15:14:15
Display OSAINFO results for IntfName: LNK29D
PortName: DEV29D PortNum: 01 Datapath: 3902 RealAddr: 0002
PCHID: 0451 CHPID: 29 CHPID Type: OSD OSA code level: 6760
Gen: OSA-E3 Active speed/mode: 1000 mb/sec full duplex
Media: Singlemode Fiber Jumbo frames: Yes Isolate: No
PhysicalMACAddr: 643B88F30000 LocallyCfgMACAddr: 000000000000
Queues defined Out: 4 In: 3 Ancillary queues in use: 2
Connection Mode: Layer 3 IPv4: Yes IPv4
SAPSup: 00010293 SAPEna: 00010293
IPv4 attributes:
VLAN ID: N/A VMAC Active: No
Defined Router: Non Active Router: No
AsstParmsEna: 00215C66 OutCkSumEna: 00
```

INFORMATION ON:

- IOCDs
- OSA Config.

INFORMATION ON:

- Registered OSA Addresses
- INBOUND Workload Queuing

```
Registered Addresses:
IPv4 Unicast Addresses:
ARP: Yes Addr: 10.10.10.10
Total number of IPv4 addresses: 1
IPv4 Multicast Addresses:
MAC: 01005E000001 Addr: 224.0.0.1
Total number of IPv4 addresses: 1
Ancillary Input Queue Routing Variables:
Queue Type: BULKDATA Queue ID: 2 Protocol: TCP
Src: 11.1.1.11..100
Dst: 12.12.12.12..100
Src: 13.3.3.13..101
Dst: 14.14.14.14..101
Total number of IPv4 connections: 2
Queue Type: SYSDIST Queue ID: 3 Protocol: TCP
Addr: 10.10.10.10
Total number of IPv4 addresses: 1
33 OF 33 Lines Displayed
End of report
```

1. Please consult the IP System Administrator's Guide for more Information.

Powerful QDIO Displays (D TCPIP,,N,DEV)



```
D TCPIP,,N,DEV
IntfName: OSAQDIOINTF IntfType: IPAQENET IntfStatus: Ready
PortName: OSAQDIO2 Datapath: 0E2A DatapathStatus: Ready
ChpidType: OSD
Speed: 0000000100
IpBroadcastCapability: No
VMacAddr: 020629DC21BD VMacOrigin: Cfg VMacRouter: All
SrcVipaIntf: VIPAV4
CfgRouter: Non ActRouter: Non
ArpOffload: Yes ArpOffloadInfo: Yes
CfgMtu: 1492 ActMtu: 1492
IpAddr: 100.1.1.1/24
VLANid: 1261 VLANpriority: Enabled
DynVLANRegCfg: Yes DynVLANRegCap: No
ReadStorage: GLOBAL (8064K)
InbPerf: Dynamic
WorkloadQueueing: Yes
ChecksumOffload: Yes SegmentationOffload: Yes
SecClass: 9 MonSysplex: Yes
Isolate: Yes OptLatencyMode: Yes
Multicast Specific:
Multicast C
Group RefCn
-----
224.0.0.1 0
SrcAddr: No
```

INFORMATION ON:
• IOCDs
• VMACs & Sourcevipa
• Performance Tuning

INFORMATION ON:
• Inbound & Outbound Statistics
• Are VIPAs for ARP Ownership Registered?

```
Interface Statistics:
BytesIn = 12834
Inbound Packets = 16
Inbound Packets In Error = 0
Inbound Packets Discarded = 0
Inbound Packets With No Protocol = 0
BytesOut = 5132
Outbound Packets = 10
Outbound Packets In Error = 0
Outbound Packets Discarded = 0
IPv4 LAN Group Summary
LanGroup: 001
Name Status ArpOwner VipaOwner
-----
TR1 Active TR1 No
LanGroup: 002
Name Status ArpOwner VipaOwner
-----
OSAQDIOLINK Active OSAQDIOLINK Yes
OSAQDIOINTF Active OSAQDIOINTF No
```

1. Please consult the IP System Administrator's Guide for more Information.



IOCDS Samples for OSA-E3 (OSE -- non-QDIO)



For non-QDIO LSA (SNA): 1 device per connection
For non-QDIO LCS (IP): 2 devices per connection
(READ=even-numbered device address; WRITE=odd-numbered device address)

Sample IOCDs for OSE



```
ID MSG1='IODF29',MSG2='SYS6.IODF29 - 2008-10-08 12:04', *
SYSTEM=(2098,1),LSYSTEM=SCZP202, *
TOK=('SCZP202',00800006991E2094120450420108282F00000000,*
00000000,'08-10-08','12:04:50','SYS6','IODF29')

RESOURCE PARTITION=((CSS(0),(A01,1),(*,2),(*,3),(*,4),(*,5),(**
,6),(*,7),(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F*
)),(CSS(1),(A11,1),(A12,2),(*,3),(*,4),(*,5),(*,6),(*,7)*
,(*,8),(*,9),(*,A),(*,B),(*,C),(*,D),(*,E),(*,F)))

CHPID PATH=(CSS(0),0C),SHARED,PARTITION=((A01),(=)),PCHID=230,*
TYPE=OSE
CNILUNIT CUNUMBR=2E40,PATH=((CSS(0),0C)),UNIT=OSA

IODEVICE ADDRESS=(2E40,031),UNITADD=00,CUNUMBR=(2E40),UNIT=OSA

IODEVICE ADDRESS=2E5F,UNITADD=FE,CUNUMBR=(2E40),UNIT=OSAD
```

For IP (LCS): 2 devices per connection (Read=even; Write=odd)
For SNA (LSA): 1 device per connection



OSA/SF Definitions in REXX for CHPID Type OSE





Configuration Definitions: OSE (IOAFENET)



Differences: OSA-E, OSA-E2 vs. OSA-E3 multiport

Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2 (Part 1)



A Smarter Planet

Instructions

```
/******  
/*          Input for configuring an OSA-Express  
/*          Fast Ethernet or 1000Base-T Ethernet CHPID  
/*  
/* This file contains the required input parameters to customize an  
/* OSA-Express fast Ethernet or 1000Base-T Ethernet CHPID.  
/* Follow the instructions to modify the data and then run IOACMD,  
/* specifying the 'Configure OSA CHPID (CONFIG_OSA)' command,  
/* to put the parameters on the OSA (CHPID).  
/*  
/* Notes: 1) This should be a copy of the sample file (IOAFENET)  
/*        2) Lines that start with a slash asterisk (/*) are comments.  
/*        3) The file is not case sensitive.  
/*  
/* Instructions:  
/*  
/* 1) All parameters have 2 indices. The first specifies the port  
/*    the parameter is associated with. The second indicates which  
/*    parameter. A brief description of the parameter is in the  
/*    comment field at the end of the line. The following sample line  
/*    fenet.0.4 = 1234567890AB /* Local MAC address (12 hex digits)  
/*    is for parameter 4, which corresponds to the local MAC address  
/*    you can use to override the universal MAC address.  
/*  
/* 2) You must have the base information that starts with 'fenet'.  
/*    The SNA information is optional. If you choose to configure SNA  
/*    The SNA information is optional. If you choose to configure SNA  
/*    parameters, they must start with 'sna.p', where 'p' is the port  
/*    number. Otherwise remove all the SNA entries from the input.  
/*    SNA is only valid for OSE (non-QDIO) CHPIDs.  
/*
```


Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2 (Part 2)



A Smarter Planet

```
/* 3) Change the sample values to match your installation.
/*
/* 4) The following parameters are for your use only. These fields
/* are not used by the OSA but are good places to keep data you
/* can associate to the specified fields.
/* a) Configuration name
/* b) User data
/*
/* 5) MAC address - You can have the OSA use a local MAC address by
/* changing parameter fenet.0.4 to match your installation.
/* To use the universal, burned in MAC address, set the parm to 0.
/* The format of a MAC address is 12 hex digits, e.g.-1234567890AB
/* with the following guideline:
/*
/* A fast Ethernet MAC address must have the following format:
/* (Ethernet MAC addresses use canonical notation)
/* Bits 0 through 5 can be anything
/* Bit 6 must be 1
/* Bit 7 must be 0
/* Bits 8 through 47 can be anything
/*
/* 6) For OSD (QDIO) CHPIDs, no OAT file is required or asked for by
/* IOACMD.
/*
/* 7) Port name - For OSE CHPIDs it must match the name you define in
/* the TCP/IP profile DEVICE and LINK statements for SNMP. If you
/* are not using SNMP, this field is ignored and can be used for
/* additional information for your installation.
/*
```

Instructions



```
/*=====
/* Fast Ethernet parameters
/*=====
fenet.0.1 = config file name      /* Configuration name (32-char max)
fenet.0.2 = user data             /* User data (32-char max)
fenet.0.3 = portname              /* Port name (8-char max)
                                  /* Data ignored for OSD CHPIDs
fenet.0.4 = 000000000000          /* Local MAC address (12 hex digits)
fenet.0.5 = auto                  /* Speed/mode
                                  /* Auto - auto negotiate
                                  /* 10H - 10 Mb, half duplex
                                  /* 10F - 10 Mb, full duplex
                                  /* 100H - 100 Mb, half duplex
                                  /* 100F - 100 Mb, full duplex
                                  /* 1000F - 1000 Mb, full duplex
                                  /* (only valid for 1000Base-T)
/*=====
/* Parameters fenet.0.6.1 to fenet.0.6.32 are used to specify group
/* addresses. You can have up to 32 assigned per OSA.
/* The last digit specifies which group address you are setting.
/* As an example, fenet.0.6.4 is used to set group address 4.
/* The format of this parameter is 12 hex digits - 1234567890AB.
/*
/* To specify a group address, modify the proper fenet.0.6.y entry.
/* To omit a group address, set fenet.0.6.y to 0 or delete the entry
/* from this file. If you do not want any group addresses, you can
/* omit all these parameters.
/*=====
fenet.0.6.1 = 000000000000        /* Group address 1 (12 hex digits)
fenet.0.6.5 = 000000000000        /* Group address 5
```

IP (LCS)
Parameters &
Instructions



```
/*=====
/* SNA parameters - Valid only for OSE (non-QDIO) CHPIDs
/*
/* The ti, t1 and t2 timers can be set as indicated below. The values
/* shown are in seconds. Any values entered that are not multiples of
/* the indicated increment will be rounded to the next highest
/* incremental value.
/* For the inactivity timer (ti), a value of 0 will disable the timer.
/* The recommended settings for large system environments are:
/* ti-90, t1-10, t2-1.04, and N3-4
/* The recommended settings for small system environments are:
/* ti-90, t1-2, t2-0.08, and N3-1
/*=====
sna.0.1 = Configuration name      /* Configuration name (32-char max)
sna.0.2 = 90.00                   /* Inactivity timer (ti)
                                   /* .24-90 in increments of .12
                                   /* 0 disables the inactivity timer
sna.0.3 = 10.00                   /* Response timer (t1)
                                   /* .20-51 in increments of .20
sna.0.4 = 1.04                    /* Acknowledgement timer (t2)
                                   /* .08-20.4 in increments of .08
sna.0.5 = 4                       /* N3 (1-4)
sna.0.6 = 8                       /* TW (1-16)
```

SNA (LSA)
Parameters &
Instructions

Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2, E3 (1)



```
/* *****
/* Input for configuring an OSA-Express Fast Ethernet or OSA-Express,
/* OSA-Express2, or OSA-Express3 1000Base-T Ethernet CHPID
/*
/* This file contains the required input parameters to customize an
/* OSA-Express fast Ethernet or any 1000Base-T Ethernet CHPID.
/* Follow the instructions to modify the data and then run IOACMD,
/* specifying the 'Configure OSA CHPID (CONFIG_OSA)' command,
/* to put the parameters on the OSA (CHPID).
/* Notes: 1) This should be a copy of the sample file (IOAFENET)
/* -OR- the output from GET_CONFIG which will give you
/* the current configuration on the OSA
/* 2) Lines that start with a slash asterisk (/*) are comments.
/* 3) The file is not case sensitive.
/*
/* Instructions:
/*
/* 1) All parameters have 2 indices. The first specifies the port
/* the parameter is associated with. The second indicates which
/* parameter. A brief description of the parameter is in the
/* comment field at the end of the line. The following sample line
/* fenet.0.4 = 1234567890AB /* Local MAC address (12 hex digits)
/* is for parameter 4 on port 0, which corresponds to the local
/* MAC address you can use to override the universal MAC address
/*
/* 2) You must have the base information that starts with 'fenet' for
/* all valid ports on the OSA.
/* The SNA information is optional. If you choose to configure SNA
/* parameters, they must start with 'sna.p', where 'p' is the port
/* number. Otherwise remove the SNA entries from the input.
/* The SNA entries for each port must follow the FENET entries
/* for the same port if they are to be used.
/* SNA is only valid for OSE (non-QDIO) CHPIDs.
/*
```


Subtle Changes
for OSA-E3:
Instructions
See APAR on IP
and SNA
definition
sequence

Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2, E3 (2)



```
/* 3) Change the sample values to match your installation.
/*
/* 4) The following parameters are for your use only.  These fields
/* are not used by the OSA but are good places to keep data you
/* can associate to the specified fields.
/* a) Configuration name
/* b) User data
/*
/* 5) MAC address - You can have the OSA use a local MAC address by
/* changing parameter fenet.0.4 to match your installation.
/* To use the universal, burned in MAC address, set the parm to 0.
/* The format of a MAC address is 12 hex digits, e.g.-1234567890AB
/* with the following guideline:
/*
/* A fast Ethernet MAC address must have the following format:
/* (Ethernet MAC addresses use canonical notation)
/* Bits 0 through 5 can be anything
/* Bit 6 must be 1
/* Bit 7 must be 0
/* Bits 8 through 47 can be anything
/*
/* 6) For OSD (QDIO) CHIPIDs, no OAT file is required or asked for by
/* IOACMD.
/*
/* 7) Port name - For OSE CHIPIDs it must match the name you define in
/* the TCP/IP profile DEVICE and LINK statements for SNMP if you
/* are using IOASNMP. If you are not using SNMP, or you are using
/* OSA Direct SNMP, IOBSNMP, this field is ignored and can be used
/* for additional information for your installation.
/*
/* 8) Group addresses are no longer used on OSA-Express3 features
/* and should be removed from the input file
```

Subtle Changes
for OSA-E3:
Instructions



Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2, E3 (3)



A Smarter Planet

```
/*=====
/* Fast Ethernet parameters for port 0
/*=====
fenet.0.1 = config file name      /* Configuration name (32-char max)
fenet.0.2 = user data             /* User data (32-char max)
fenet.0.3 = portname             /* Port name (8-char max)
                                /* Data ignored for OSD CHPIDs
fenet.0.4 = 000000000000         /* Local MAC address (12 hex digits)
fenet.0.5 = auto                 /* Speed/mode
                                /* Auto - auto negotiate
                                /* 10H - 10 Mb, half duplex
                                /* 10F - 10 Mb, full duplex
                                /* 100H - 100 Mb, half duplex
                                /* 100F - 100 Mb, full duplex
                                /* 1000F - 1000 Mb, full duplex
                                /* (only valid for 1000Base-T)

/*=====
/* Parameters fenet.0.6.1 to fenet.0.6.32 are used to specify group
/* addresses. You can have up to 32 assigned per OSA.
/* The last digit specifies which group address you are setting.
/* As an example, fenet.0.6.4 is used to set group address 4.
/* The format of this parameter is 12 hex digits - 1234567890AB.
/*
/* To specify a group address, modify the proper fenet.0.6.y entry.
/* To omit a group address, set fenet.0.6.y to 0 or delete the entry
/* from this file. If you do not want any group addresses, you can
/* omit all these parameters.
/*
/* For OSA-Express3, these fields are not used and should be removed
/*=====
fenet.0.6.1 = 000000000000      /* Group address 1 (12 hex digits)
fenet.0.6.5 = 000000000000      /* Group address 5
```

Subtle Changes for OSA-E3:

- See APAR on "Speed"
- Syntax for Port 0: "fenet.0.n"

Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2, E3 (4)



A Smarter Planet

```
/*=====
/* SNA parameters for port 0 - Valid only for OSE (non-QDIO)
/*
/* The ti, t1 and t2 timers can be set as indicated below. The values
/* the indicated increment will be rounded to the next highest
/* incremental value.
/* For the inactivity timer (ti), a value of 0 will disable the timer.
/* The recommended settings for large system environments are:
/* ti-90, t1-10, t2-1.04, and N3-4
/* The recommended settings for small system environments are:
/* ti-90, t1-2, t2-0.08, and N3-1
/*=====
sna.0.1 = Configuration name      /* Configuration name (32-char max)
sna.0.2 = 90.00                  /* Inactivity timer (ti)
                                  /* .24-90 in increments of .12
                                  /* 0 disables the inactivity timer
sna.0.3 = 10.00                 /* Response timer (t1)
                                  /* .20-51 in increments of .20
sna.0.4 = 1.04                  /* Acknowledgement timer (t2)
                                  /* .08-20.4 in increments of .08
sna.0.5 = 4                     /* N3 (1-4)
sna.0.6 = 8                     /* TW (1-16)
```

Subtle Changes for
OSA-E3:

- Syntax for Port 0:
"sna.0.n"

Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2, E3 (5)



```
/*=====
/* Fast Ethernet parameters for port 1
/*=====
fenet.1.1 = config file name      /* Configuration name (32-char max)
fenet.1.2 = user data             /* User data (32-char max)
fenet.1.3 = portname              /* Port name (8-char max)
                                  /* Data ignored for OSD CHPIDs
fenet.1.4 = 000000000000          /* Local MAC address (12 hex digits)
fenet.1.5 = auto                  /* Speed/mode
                                  /* Auto - auto negotiat
                                  /* 10H - 10 Mb, half duplex
                                  /* 10F - 10 Mb, full duplex
                                  /* 100H - 100 Mb, half duplex
                                  /* 100F - 100 Mb, full duplex
                                  /* 1000F - 1000 Mb, full duplex
                                  /* (only valid for 1000Base-T)
/*=====
/* Parameters fenet.0.6.1 to fenet.0.6.32 are used to specify group
/* addresses. You can have up to 32 assigned per OSA.
/* The last digit specifies which group address you are setting.
/* As an example, fenet.0.6.4 is used to set group address 4.
/* The format of this parameter is 12 hex digits - 1234567890AB.
/*
/* To specify a group address, modify the proper fenet.0.6.y entry.
/* To omit a group address, set fenet.0.6.y to 0 or delete the entry
/* from this file. If you do not want any group addresses, you can
/* omit all these parameters.
/*
/* For OSA-Express3, these fields are not used and should be removed
/*=====
fenet.1.6.1 = 000000000000        /* Group address 1 (12 hex digits)
fenet.1.6.5 = 000000000000        /* Group address 5
/*=====
```

Subtle Changes for OSA-E3:

- See APAR on "Speed"
- Syntax for Port 1: "fenet.1.n"

Ex.: SYS1.OSA.IOASAMP(IOAFENET) -- OSA-E2, E3 (6)



A Smarter Planet

```
/*=====
/* SNA parameters for port 1 - Valid only for OSE (non-QDIO)
/*
/* The ti, t1 and t2 timers can be set as indicated below. The values
/* shown are in seconds. Any values entered that are not multiples of
/* the indicated increment will be rounded to the next highest
/* incremental value.
/* For the inactivity timer (ti), a value of 0 will disable the timer.
/* The recommended settings for large system environments are:
/* ti-90, t1-10, t2-1.04, and N3-4
/* The recommended settings for small system environments are:
/* ti-90, t1-2, t2-0.08, and N3-1
/*=====
sna.1.1 = Config name port 1      /* Configuration name (32-char max)
sna.1.2 = 90.00                  /* Inactivity timer (ti)
                                /* .24-90 in increments of .12
                                /* 0 disables the inactivity timer
sna.1.3 = 10.00                 /* Response timer (t1)
                                /* .20-51 in increments of .20
sna.1.4 = 1.04                  /* Acknowledgement timer (t2)
                                /* .08-20.4 in increments of .08
sna.1.5 = 4                     /* N3 (1-4)
sna.1.6 = 8                     /* TW (1-16)
```

Subtle Changes for
OSA-E3:

- Syntax for Port 1:
"sna.1.n"



OSE: OSA/SF OAT Definition



Sample was not changed for OSA-E3 ... still says port is always 00

```
***** Top of Data *****
* This OAT template is a sample for setting up TCP/IP and SNA modes
* with port sharing between 2 Images running on a z990 which supports
* logical channel subsystems.
* Image 0.5 and Image 0.7 are sharing port 0.
*
* It can also be used for other combinations of modes by following
* the instructions below. There are four cases where you MUST
* have an IP address on your passthru entries to allow all the
* defined modes to operate properly.
* 1) There is TCP/IP traffic to different Images
* 2) TCP/IP and SNA traffic are sharing a port
* 3) TCP/IP and MPC traffic are in use on the CHPID
* 4) There is more than 1 TCP/IP stack using a port
*
* In this example, Image 0.5 and Image 0.7 are sharing port 0.
* Port 0 is also being shared by 2 TCP/IP stacks on the same image.
* To use this template, do the following:
*
* 1) Change the Image numbers to match your installation.
* The Image statement must precede all entries for that Image.
* 2) Change the unit addresses. UAs must be even numbers for passthru.
* The odd entries are automatically added by the CHPID.
* 3) The mode must be passthru or sna (for this example).
* 4) The port number must be 00 for all OSA-Express CHPIDs.
```

Ex.: SYS1.OSA.IOASAMP(IOAOSHRA) -- OSA-E2, E3 (2)



A Smarter Planet

```

* 5) If you are not using TCP/IP, delete all passthru OAT entries and
* continue at (6) below.
* Otherwise, change the IP addresses. They are required for TCP/IP
* when sharing a port.
* You can have up to 8 IP addresses per OAT entry.
* You can have up to 512 IP addresses per port spread over multiple
* OAT entries for OSA-Express CHPIDs.
* An IP address of 0.0.0.0 indicates no port sharing.
* You CANNOT have a mixture of zero and non-zero IP addresses in
* the same OAT.
* Specify if the Image is the default entry (No, PRI, or SEC).
* Only one entry per port can be the PRImary default entry.
* Only one entry per port can be the SECondary default entry.
* 6) If you are not using SNA, delete all SNA OAT entries and continue
* at (7) below.
* 7) Add additional entries as needed for each mode you are configuring
*****
* UA(Dev) Mode  Port  Entry specific information  Entry Valid
*****
                        Image 0.5
00(2E40)* passthru 00 Pri 105.001.005.005          SIU  ALL
                        105.001.005.015
                        105.001.005.025
                        105.001.005.035
02(2E42)* passthru 00 No 100.100.100.100          SIU  ALL
0A(2E4A) SNA      00                          SIU  ALL
*****
                        Image 0.7
00(2E50)* passthru 01 No 107.001.075.075          SIU  ALL
                        107.100.075.085
02(2E52)* passthru 01 Sec 107.005.035.035        SIU  ALL
0A(2E5A) SNA      01                          SIU  ALL
    
```

Subtle Changes for OSA-E3:

- Must change to reflect Port 01, as if this were an old 3172.
- Documented in the OSA Configuration Guides



z/OS VTAM Definitions for SNA Link Station Architecture (LSA)



For non-QDIO LSA (SNA): 1 device per connection

CHPID Type=OSE: SNA XCA Major Node (LSA - Ex. 1)



```
XCAOSA VBUILD TYPE=XCA
OSAX30  PORT  MEDIUM=CSMACD,                X
          ADAPNO=0,                          X
          CUADDR=2E4A,                       X
          TIMER=60,                          X
          SAPADDR=04
*****
OSAX3G0  GROUP DIAL=YES,                    X
          DYNPU=YES,                         X
          ANSWER=ON,                         X
          AUTOGEN=(3,L0,P0),                 X
          CALL=INOUT,                       X
          ISTATUS=ACTIVE
*****
CCL1AG0  GROUP DIAL=NO,                     X
          ISTATUS=INACTIVE
CCL1AL0  LINE  USER=SNA
CCL1AP0  PU    MACADDR=02000000B000,PUTYPE=4,SUBAREA=2,TGN=1,X
          SAPADDR=4,ALLOWACT=YES
```

© Copyright IBM Corporation 2011

1. Here you see how VTAM specifies that the address is available on Port 0 of a multiport CHPID.
2. The DIAL=YES group is for peripheral SNA devices.
3. The DIAL=NO group for "leased lines" is for Subarea Node attachments, like other VTAMs or NCPs.

CHPID Type=OSE: SNA XCA Major Node (LSA - Ex. 2)



```
XCAOSA VBUILD TYPE=XCA
OSAX31  PORT  MEDIUM=CSMACD,                X
        ADAPNO=1,                            X
        CUADDR=2E5A,                          X
        TIMER=60,                             X
        SAPADDR=04
*****
OSAX3G1  GROUP DIAL=YES,                    X
        DYNPU=YES,                            X
        ANSWER=ON,                            X
        AUTOGEN=(3,L1,P1),                    X
        CALL=INOUT,                           X
        ISTATUS=ACTIVE
*****
CCL1AG1  GROUP DIAL=NO,                      X
        ISTATUS=INACTIVE
CCL1AL1  LINE  USER=SNA
CCL1AP1  PU    MACADDR=02000000B000,PUTYPE=4,SUBAREA=2,TGN=1,X
        SAPADDR=4,ALLOWACT=YES
```

© Copyright IBM Corporation 2011

1. Here you see how VTAM specifies that the address is available on Port 1 of a multiport CHPID.
2. The DIAL=YES group is for peripheral SNA devices.
3. The DIAL=NO group for "leased lines" is for Subarea Node attachments, like other VTAMs or NCPs.



z/OS TCP/IP Definitions for IP LAN Channel Station (LCS)



For non-QDIO LCS (IP): 2 devices per connection
(READ=even-numbered device address; WRITE=odd-numbered device address)

CHPID Type=OSE: TCP/IP Profile for LCS



```
;
;-----Name-----Protocol-----Device Number
DEVICE OSA2E40   LCS                2E40
;-----Name-----Protocol-----Portnum-----Device Name
LINK   OSA2E40L  ETHEROR802.3    0        OSA2E40
;
```

```
;
;-----Name-----Protocol-----Device Number
DEVICE OSA2E50   LCS                2E50
;-----Name-----Protocol-----Portnum-----Device Name
LINK   OSA2E50L  ETHEROR802.3    1        OSA2E50
;
```

© Copyright IBM Corporation 2011

1. The Lan Channel Station definition also requires that you specify the OSA Port number.



APPENDIX: Summary of Common Problems Implementing OSA-E3 and OSA-E4S





Common Problems with OSA Implementations

- **Neglecting to update Microcode, VTAM and TCP/IP PTF levels**
- **Converting to INTERFACE statement and forgetting that one must specify SOURCEVIPAIN for the Source VIPA**
- **If assigning OSA/SF device, trying to allocate it on any CUADD but 0.**
 - OSA/SF device must be assigned to CUADD=0.
- **QDIO**
 - Neglecting to set Missing Interrupt Handler (MIH)
 - Portnames not unique on the CHPID and on the VTAM
- **Failing to exploit the power of commands like:**
 - D TCPIP,,OSAINFO,INTFNAME=xxxx
 - D TCPIP,,N,DEV,INTFNAME=xxxx
- **non-QDIO with IOACMD (REXX) Definitions**
 - OAT Definition:
 - Wrong Port number in OAT if multiport CHPID
 - Neglecting to update OSA/SF PTFs
 - Including Group Addresses for non-QDIO FENET (not valid for OSA-E3)
 - Configuration Definition for 1000Base-T "FENET" (UA46430 available 04/01/2009):
 - SNA must immediately follow FENET definitions
 - Both PORTNUMS must be included even if one is not in use
 - Use speed of AUTO unless APAR is applied to allow 1000F
 - AUTO-NEGOTIATE
 - 1000Base-T normally supports auto-negotiate for
 - LAN Speed and DUPLEX Mode
 - Ethernet switch and OSA can negotiate between themselves, or
 - Ethernet switch may be configured with a specific speed and duplex setting
 - OSA accommodates itself to the signal if the OSA is set up for auto-negotiate
 - If OSA port is configured with a value that is incompatible with the switch, the OSA will not connect.
- **Loss of Connectivity through Spanning Tree Protocol Problems**

© Copyright IBM Corporation 2011

1. The following information is quoted from the Red Book "OSA Express Implementation Guide" (SG24-5948-05).
2. The OSA-Express3 1000BASE-T Ethernet feature supports auto-negotiation when attached to an Ethernet router or switch. If you allow the LAN speed and duplex mode to default to auto-negotiation, the OSA port and the attached router or switch auto-negotiate the LAN speed and duplex mode settings between them, and connect at the highest common performance speed and duplex mode of interoperation. If the attached Ethernet router or switch does not support auto-negotiation, the OSA port examines the signal it is receiving, and connects at the speed and duplex mode of the device at the other end of the cable.
3. You can choose any of the following settings for the OSA-Express3 1000BASE-T Ethernet feature port:
 1. 7 Auto-negotiate
 2. 7 10 Mbps half-duplex
 3. 7 10 Mbps full-duplex
 4. 7 100 Mbps half-duplex
 5. 7 100 Mbps full-duplex
 6. 7 1000 Mbps full-duplex
4. If you are not using auto-negotiate, the OSA port will attempt to join the LAN at the specified speed and duplex mode. If this does not match the speed and duplex mode of the signal on the cable, the OSA port will not connect. LAN speed and duplex mode can be set explicitly, using OSA/SF or the OSA Advanced Facilities function of the hardware management console (HMC). The explicit settings will override the OSA feature port's ability to auto-negotiate with its attached Ethernet switch.

- **Missing Interrupt Handler for QDIO**
- **The WRITE devices (as defined in the TRLE) should have a**
 - Missing Interrupt Handler (MIH) value of at least 15 seconds
 - (or 30 seconds, if running as a guest system on z/VM).
- **To determine the current MIH value for the device (E201, in our example), enter the command:**
 - D IOS,MIH,DEV=E201
- **To dynamically change the MIH value, enter the command:**
 - SETIOS MIH,DEV=E201,TIME=00:15
- **QDIO mode for z/OS**
 - To set these MIH values at IPL time, update the IECIOSxx member in PARMLIB.
 - Important:
 - On a multisubchannel device, the MIH is automatically configured OFF by VTAM on the READ subchannel or subchannels.
 - Setting an MIH value of ZERO for a TCP/IP or VTAM WRITE device disables MIH on those devices.

- Original Info APAR on MIH:
 - <http://www-01.ibm.com/support/docview.wss?uid=isg1112320>
- Newer Technote on MIH:
 - <http://www-01.ibm.com/support/docview.wss?rs=852&uid=swg21326865>

OSA/SF and IOACMD Errors (OSE): APAR 0A27643



APAR - OA27643

APAR - OA27643

OSA/SF FOR Z/OS OSA-EXPRESS3 1000BASE-T FEATURE FOR 1000F
MODEIOACMD

OSA/SF FOR Z/OS OSA-EXPRESS3 1000BASE-T FEATURE FOR 1000F
MODEIOACMD

Summary Page

APAR Identifier: OA27643 Last Changed: 04/02/2009
Symptom: IN INCORROUT Status: CLOSED PER
Severity: 3 Date Closed: 03/18/2009
Component: 565510400 Duplicate of:
Reported Release: 400 Fixed Release: 999
Component Name: OSA/SF Special Notice:
Customer Name: IBM INTERNAL - Customer #: 9999999 00
Current Target Date: 03/31/2009 Flags:
SCP:
Platform:
Status Detail: SHIPMENT
' Fix Page
PE PTF List:

PTF List: Release 400 : UA46430 available 04/01/2009 (F903)

OSE OSA/SF Config file (FENET) Errors with IOACMD:

1. Speed must be AUTO
2. SNA must follow each set of IP Definitions
3. Must Configure Ports 0 and 1

' Submitter Page Error.

Description: Using the following definition in the configuration file doesn't work when configuring an OSA Express3 1000base-T Ethernet feature using IOACMD EXEC:

fenet.0.5 = 1000F

Messages received:

IOACMD: The speed in the config file is 1000F. This is not
IOACMD: valid for a 1000Base-T Ethernet CHPID. EXEC ending.
--> Error during IOACMD CONFIG_OSA 81 - RC=8
Using 'fenet.0.5 = AUTO' works fine.

Additional symptoms:

Trying to configure an OSA-Express3 1000Base-T with
SNA and passthru has found another bug in the IOACMD EXEC. The
problem is twofold:

First, if someone is using SNA also, they must code the
SNA.port.entry parameters right after the corresponding
FENET.port.entry parms. They can't put both SNA sets of parms
after both sets of FENET entries. This is minor.

Second problem is even if they fix this, unless they code BOTH
ports or port 1 only, it won't work due to a
programming bug. In fact, if they just use port 0, IOACMD will
hang on z/OS.

If customer wants to make the change, here is what needs to be
changed: ...

- **VIPA Address Moves (Static or Dynamic VIPAs)**

- **Implement PortFast, a Spanning Tree Protocol (STP) feature on Cisco Routers**

- Shortens the time needed to detect availability of a new link in a switch while continuing to implement Spanning Tree Protocol

- **Disable Spanning Tree Protocol on the Router connection to the OSAs**

- Spanning Tree detects what it thinks is a loop during takeover or move processing of the VIPAs and can block the transfer of packets
 - ▶ Can cause loss of connectivity to the VIPAs that have moved if STP left enabled

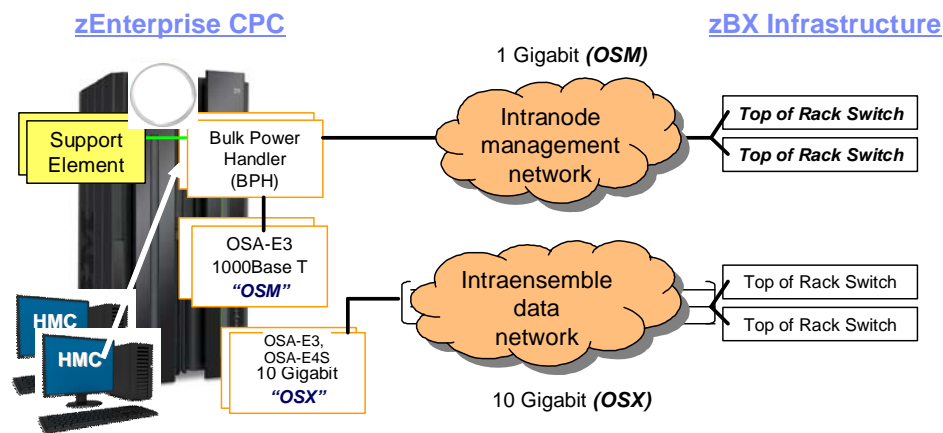
1. The following paragraphs are taken from the z/OS Communications Server IP Configuration Guide.
 1. Spanning Tree Protocol (STP) can potentially impact environments where both OSA-Express QDIO and VIPA are deployed.
 2. With dynamic VIPA, a TCP/IP address takeover occurs when a TCP/IP stack fails, or a static or dynamic VIPA is manually moved by operator intervention. In either case, when a static or dynamic VIPA moves, the IP address and respective workloads are taken over by another TCP/IP stack through other OSA-Express Ethernet devices running on a different server. When the original TCP/IP stack and respective OSA-Express QDIO devices are returned to operation, both the IP address and respective workload traffic are taken back by the recovered TCP/IP stack.
 3. If the network bridge or switch is not configured properly, packets can get lost in the network or be blocked by the networking equipment. This is a result of a physical looping condition identified by the STP, or expired OSA-Express QDIO timers due to the increased latencies associated with blocked ports or delayed packets. In these cases, static or dynamic VIPA connectivity can fail.
 4. When configuring STP, use care in the bridge or switch configuration to avoid or minimize potential loop conditions. For example, if the STP is not an integral component of the overall network, disabling the STP on all of the Virtual LANs (VLANs) that connect to OSA-Express QDIO devices will help avoid the problem. Also, some networking switches provide a mechanism for suppressing the STP's listening and learning states on specific switch ports. For example, Cisco Systems provides an STP configuration feature called PortFast that can place specific switch ports into forwarding state as soon as a link is detected. Without PortFast enabled, a switch port has to transition through the listening and learning stages (30 seconds total) of Spanning Tree reconvergence before the switch port can actually pass valid traffic. For PortFast capable Cisco switches, enabling PortFast on all switch ports that OSA-Express QDIO is connected into allows network administrators to both preserve their original Spanning Tree configuration by not having to change it, while also providing a viable mechanism to avoid potential static or dynamic VIPA problems.
2. For more information on configuring STP, see the OSA Red Book "OSA Express Implementation Guide" (SG24-5948-05).
 1. Warning: Spanntree port fast start should only be enabled on ports connected to a single host. Connecting hubs, concentrators, switches, bridges, etc. to a fast start port can cause temporary spanning tree loops. Use with caution.



Appendix: News about New CHPID Types (OSM, OSX) for zEnterprise™



The Ensemble Networks in zEnterprise & zBX Model 002



Intranode management network (INMN)

- 1000Base-T OSA-Express3 (copper) --- QDIO (CHPID Type OSM) – Cables are 3.2 meters long from OSM to BPH in CEC and 26 meters from BPH to TOR
- HMC security is implemented with standard practices PLUS additional security mechanisms:
 - Isolated IPv6 network with "link-local" addresses only; authentication and authorization and access control, etc.


Intraensemble data network (IEDN)

- 10 Gigabit OSA-Express3 --- QDIO (CHPID Type OSX) – Cables are maximum of 26 meters long to TOR & 10km long-range (**OSA-Express4S 10 Gigabit as OSD or OSC announced July 12, 2011**)
- Security is implemented with standard practices PLUS additional security mechanisms: VLAN ID enforcement, access control, authentication, authorization, application security, routing table restrictions, IP Filtering, etc.
- Networks can be further isolated using VLAN and VMAC segmentation of the network connections

© Copyright IBM Corporation 2011

- For each central processor complex (CPC) that participates in an Ensemble, define GbE ports of an OSA-Express3 1000Base-T card as CHPID Type of OSM in the IOCDs; the OSM ports connect to the Intranode Management Network (INMN) over which the Unified Resource Manager defines, accesses, and manages the members of the ensemble. You can define ports that are shared among multiple logical partitions (LPARs) or ports that are dedicated to a single LPAR. A dedicated port is not required. It is recommended that you define ports that are shared just between the LPARs that work with your IBM BladeCenter Extension.
 - NOTE: You do NOT define devices, links, or interfaces to this INMN from any of the Virtual Servers; this wholly self-contained private network dynamically builds the connections to the INMN when the server becomes a member of the Ensemble.
 - Note how the Support Element is still connected to the BPH switch as with the z10; however, now the OSM CHPID is also attached to the BPH Switch.
 - HMC security is implemented with standard practices, but there are also additional safeguards for security, because the IPv6 network is automatically created without a chance of human error during device definition. This is an isolated network that uses link-local addresses only; further authentication and authorization are implemented through the Firmware and through Operating System enablement to restrict access to the INMN.
- Intraensemble data network (IEDN)
 - 10 Gigabit OSA-Express3 --- QDIO (CHPID Type OSX)
 - Connected to authorized members of the Ensemble and to the Support Element with which the Hardware Management Console communicates
 - Maximum of 16 data paths (8 pairs of redundant paths)
 - There are eight OSA adapters (16 OSA ports) needed for maximum configuration in a node. Only the first pair of OSA cables is required to be connected to the managing zEnterprise.
 - Security is implemented with standard practices PLUS additional security mechanisms: VLAN ID enforcement, access control, authentication, authorization, application security, routing table restrictions, IP Filtering, etc.
 - VLAN and VMAC segmentation of the network connections
 - Each connection to the IEDN must be assigned a VLAN. Without a VLAN a Virtual Guest / Virtual Server cannot communicate on the IEDN. Untagged VLAN traffic within the IEDN is not supported. In addition Virtual Guests / Virtual Servers are assigned a specific vMAC by z/VM and OSX using the VMAC prefix defined with Unified Resource manager. This is not an option and is enforced.
 - Can assign Private network addresses (non-routable addresses) to the resources in the zBX, or can assign Public network addresses (routable) addresses to the resources, or can assign a mixture, especially if desiring to provide network reachability to a VIPA in the zBX.
 - Can implement with static routing or with dynamic routing.

The 1000Base-T OSA Configured as CHPID Type OSM



IOCDs

```

CHPID PCHID=191,PATH=(CSS(0,1,2,3),23), *
TYPE=OSM,CHPARM=01,SHARED, ...
CNTLUNIT CUNUMBR=0910,PATH=((CSS(0),23)), *
UNIT=OSM
IODEVICE ADDRESS=(0910,15),CUNUMBR=(0910), *
UNIT=OSA,UNITADD=00, *
MODEL=M,DYNAMIC=YES,LOCANY=YES
    
```

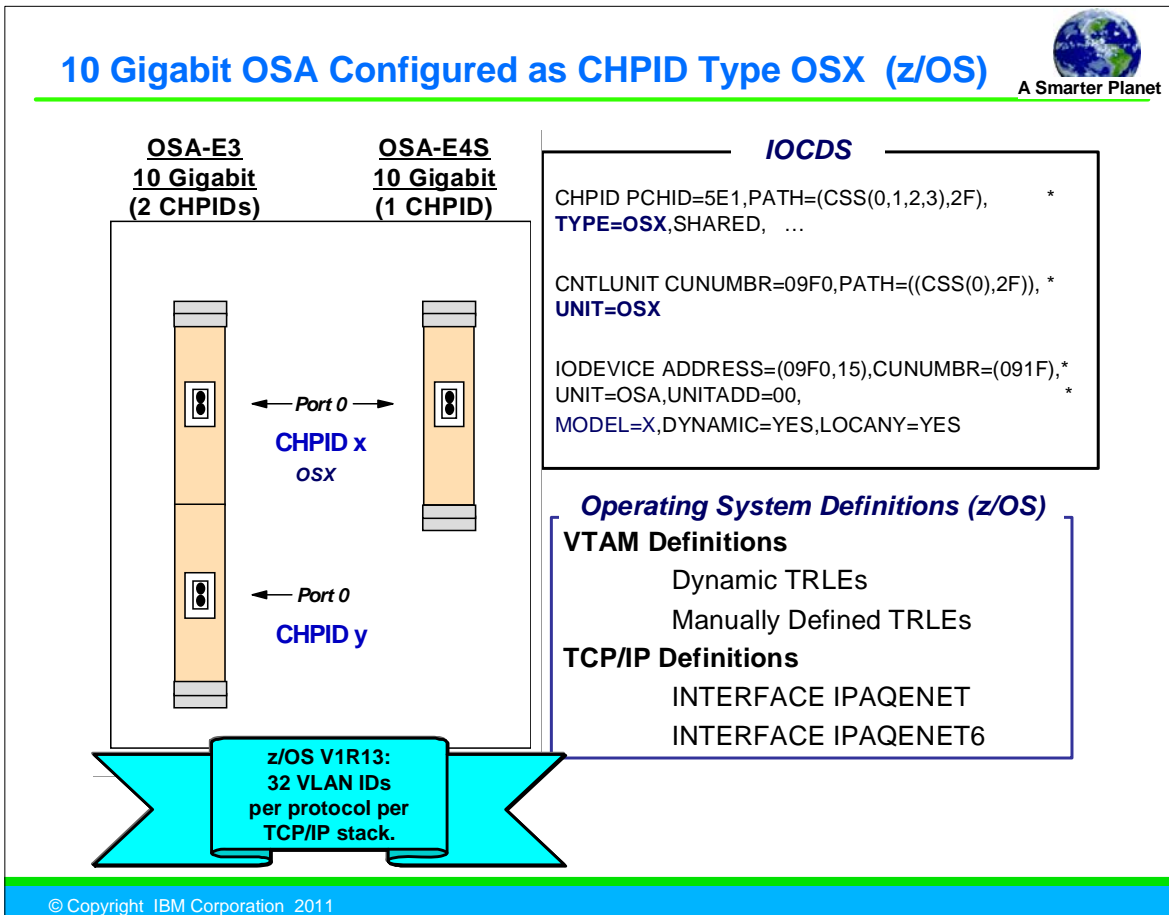
z/OS Operating System Definitions

Not Applicable
 These OSM Connections are dynamically created when an Ensemble is defined.

© Copyright IBM Corporation 2011

1. Notice first of all that this adapter is a four-port adapter and is also called a “multiport adapter.” This OSA-Express3 (1000Base-T) is engineered with two ports on each CHPID: CHPID x and CHPID y. The two ports on each CHPID are numbered port 0 and port 1. But note how the top half of the OSA-E3 is the mirror image of the bottom half with regard to the port number assignments; reading from top to bottom you see Port 0, Port 1, Port 1, Port 0. As with any OSA port, the portnames on the multi-port OSA-E3 must be unique to a CHPID in z/OS. Portname relief is implemented in z/VM and Linux on z, meaning that the portname is ignored by the OSA firmware whether specified in the Operating System definitions or not.
2. You may want to read more about this OSA type in technical document PRS3950 named “Migrating to a Multi-port CHPID OSA-E3: Avoiding Common Problems (CHPID Types OSD and OSE)” It is available at: www.ibm.com/support/techdocs.
3. You may also want to review Flash number FLASH10706 at the same web site; it is named “OSA-E3 Multiport and Portname Conflicts.”
4. Note (1): When an OSA CHPID is configured with a CHPID type of “OSM,” only PORT 0 is available; PORT 1 is not used at all in this case.
5. Note (2): The second CHPID on an OSA-E3 1000Base-T may be defined in the IOCDs as OSC, OSE, OSD, or even OSM. As per Note (1), if configured as OSM, PORT 1 is blocked from usage.
6. Operating System Definitions: QDIO TRLE definitions are always dynamically generated for connectivity to the
7. intra node management network (CHPID type OSM). For each dynamic TRLE for OSM, the TRLE name is IUTMT0xx and the corresponding portname is IUTMP0xx (where xx is the value of the OSM CHPID).

10 Gigabit OSA Configured as CHPID Type OSX (z/OS)



1. Notice first of all that the 10 Gigabit adapters have one port per CHPID.
2. NOTE: You may want to read more about portnames that you can assign to these OSA ports in VTAM. Please see Flash number FLASH10706 (named "OSA-E3 and OSA-E4S Multiport and Portname Conflicts") at www.ibm.com/support/techdocs.
3. z/OS Operating System Definitions:
4. Dynamic Definition: When z/OS Communications Server participates in an ensemble environment, QDIO TRLE definitions are dynamically generated for connectivity to the intraensemble data network (CHPID type OSX) if the QDIO interface is defined with CHPIDTYPE OSX and the CHPID parameter. For each dynamic TRLE for OSX, the TRLE name is IUTXT0xx and the corresponding portname is IUTXP0xx (where xx is the configured CHPID parameter).
5. Manual Definition:
6. Tip: If the values that are chosen for the dynamically created TRLE definitions do not suit your needs (for example, you need more than 17 DATAPATH devices), you can define your own QDIO TRLE and configure the QDIO interface with the PORTNAME parameter. See INTERFACE -- IPAQENET OSA-Express QDIO
7. interfaces statement and INTERFACE -- IPAQENET6 OSA-Express QDIO interfaces statement in z/OS Communications Server: IP Configuration Reference for more information.
8. Restrictions:
 1. If the QDIO interface that is used for connectivity to the intra ensemble data network is defined with the PORTNAME parameter, you must define a QDIO TRLE definition. The PORTNAME parameter value on the TRLE definition must match the PORTNAME parameter value. For information about restrictions on the TRLE name, see Transport resource list major node in z/OS Communications Server: SNA Resource Definition Reference.
 2. If z/OS Communications Server runs as a guest on z/VM and guest LAN definitions is used for connectivity to the intra ensemble data network, you must define the QDIO interface with the PORTNAME parameter.
9. When one CHPID is configured as Type OSX, the other CHPID can be configured as either OSD or as OSD.
10. Specify the CHPARM keyword for the OSD, OSM, or OSX channel path with bit 6 of the CHPARM value on (CHPARM=02) to disable priority specification and support a maximum of 1920 OSA valid subchannels (640 TCP/IP stacks). Otherwise, an OSD, OSM, or OSX channel path has priority specification enabled and supports a maximum of 480 OSA valid subchannels (160 TCP/IP stacks).
11. A note about MODEL on the IODEVICE: MODEL=model specifies the model number, if any, for the device. You can specify a maximum of four alphanumeric or special characters (#, @, or \$) that represent the model number of the device. IOCP checks only the syntax for alphanumeric or special characters. IOCP does not validate the model value. For VM the MODEL=model parameter of the IODEVICE statement in the IOCP input file has the same value as the MODEL=model parameter of the RDEVICE statement in the SYSTEM CONFIG file.
12. For example, the following RDEVICE statement is in the SYSTEM CONFIG file: RDEVICE ... MODEL 11 The corresponding IODEVICE statement is: IODEVICE ...,MODEL=11,...
13. A note about SHARED OSX ports: If you are sharing an OSX OSA port among LPARs and wish to prevent certain devices on the OSX CHPID from being shared, take advantage of the NOTPART keyword on the IODEVICE definition in the IOCDs. An LP cannot access a device if the LP is not specified in the device candidate list for the device, even if the LP can access a channel path assigned to the device. The PART or PARTITION keyword specifies the LPs that are in the device candidate list. The NOTPART keyword specifies the LPs that are not in the device candidate list. For example, if a CSS has three LPs (LP1, LP2, and LP3) and you specify NOTPART=(LP2), LP1 and LP3 can access the device, but LP2 cannot. This capability may provide assurances that certain Operating System and TCP/IP stack images cannot define access to the OSX OSA over certain addresses despite the SHARED CHPID. However, if this level of assurance does not satisfy an auditor, you may still exploit the z/VM or z/OS ISOLATE keyword to enhance the segregation of the Operating Systems sharing the OSX OSA. ISOLATE by itself may satisfy a security auditor or you may use ISOLATE in the interface definition together with the NOTPART definition in the IOCDs. However, only in very rare circumstances should you ever have to use ISOLATE in an ensemble environment, because the VLAN and VMAC enforcement of the ENSEMBLE already segregate traffic flows across the shared OSA. Implementing ISOLATE could, in fact, introduce network problems if not properly planned for.

```

VTAM Start Options (ATSCTR00)
  _ENSEMBLE=NO
> _____ >
  |_____|
  |_ENSEMBLE=_____|_NO_____|
  |_____|_YES_____|
    
```

```

Dynamic TRLE:
IUTXTOF1
IUTXPOF1
    
```

```

TCP/IP Profile INTERFACE Definitions
  _____ _CHPIDTYPE OSD_____ | OSD Interface Definition |
INTERFACE ...DEFINE ... -+-----+
  |_____ _CHPIDTYPE OSX_____| | OSX Interface Definition |
    
```

Dynamic Creation of VTAM TRLE for OSX, or Predefined TRLE in VTAM?	Example TCP/IP INTERFACE Definition for OSX (adjust values as required: MTU, etc.)
1. DYNAMIC TRLE (up to 17 Datapath Devices): Specify 'CHPID <number>' on INTERFACE for Dynamically generated PORTNAME (See 2.)	<pre> INTERFACE QDIOOSX1 DEFINE IPAQENET CHPIDTYPE OSX _CHPID F1 VLANID 10 MTU 8992 IPADDR 172.16.1.1/24 VMAC ROUTEALL INBPERF DYNAMIC WORKLOADQ READSTORAGE GLOBAL </pre>
2. Predefined TRLE (> 17 Datapath Devices): Specify 'PORTNAME <TRLEportname>' on INTERFACE	<pre> INTERFACE QDIOOSX1 DEFINE IPAQENET CHPIDTYPE OSX _PORTNAME OSXNY1 VLANID 10 MTU 8992 IPADDR 172.16.1.1/24 VMAC ROUTEALL INBPERF DYNAMIC WORKLOADQ READSTORAGE GLOBAL </pre>

© Copyright IBM Corporation 2011

- This shows the syntax for the VTAM start option. This option indicates whether or not this z/OS LPAR will participate in the ensemble. If the CPC is not configured as a member of an ensemble, then VTAM ignores this start option and issues existing message IST448I.
- For OSX configured with CHPID, VTAM dynamically creates a TRLE with up to 17 DATAPATH devices. The dynamically assigned names for the TRLE and for the OSX Portname are: TRLE name = IUTXT0xx (where xx = CHPID) and PORTNAME = IUTXP0xx. If you need more than 17 DATAPATH devices, you can predefine your TRLEs. All other defaults are the same as for DynamicXCF HiperSockets: LNCTL = MPC MPCLEVEL = QDIO MPCUSAGE = SHARE MAXBFRU = 2 MAXREADS = 2 PACKING = OFF REPLYTO = 30 STORAGE = DS LASTRW = DISALLOW
- Alternatively you may choose to create your own TRLE definitions and assign your own portnames. You may wish to predefine the TRLE definitions if you are operating as a GUEST under z/VM.
- The visual also shows the syntax for an INTERFACE definition for OSX (CHPIDTYPE OSX). (CHPIDTYPE defaults to OSD.) The OSX syntax diagram is similar to that of OSD, but has some different defaults from OSD:
- 1) For OSX, you may specify either CHPID or PORTNAME depending on whether the underlying TRLE has been dynamically or statically defined.
- 2) VLANID is required
- 3) Defaults to INBPERF DYNAMIC
- 4) Requires an OSA-generated VMAC; you may not specify your own VMAC value.
- 5) You must specify a subnet mask for IPADDR if you are defining an IPv4 interface.
- 6) PRIROUTER/SECROUTER/NONROUTER are not applicable to an OSX Interface definition.
- 7) All other parameters – including ISOLATE|NOISOLATE -- default to OSD defaults
- This chart also contains an example statement which defines an IPv4 OSX interface for CHPID F1 because the OSX TRLE for this CHPID was dynamically created by VTAM when the ENSEMBLE Start Option was enabled. (If the VTAM TRLE was Dynamically Created, use CHPID=<number> on the INTERFACE statement in TCP/IP. Alternatively you may instead use the PORTNAME assigned dynamically by VTAM.)
- You also see an example for an INTERFACE statement when the TRLE was predefined; this is a case when you must specify PORTNAME in the INTERFACE definition. (If the VTAM TRLE was defined in a VBUILD TYPE=TRL, use PORTNAME=<TRLEportname> on the INTERFACE statement.)



References



- IP System Administrator's Commands (z/OS V1R12) - SC31-8781-10
- IP Configuration Reference (z/OS V1R12) - SC31-8776-18
- Converting to INTERFACE Statement for OSA Port Interfaces in z/OS
 - <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FLASH10744>
- OSA-E3 and OSA-E4S Multiport and Portname Conflicts
 - <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FLASH10706>
- Migrating to a Multi-port CHPID OSA-E3 or OSA-E4S: Avoiding Common Problems (CHPID Types OSD and OSE)
 - <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS3950>
- SA22-7990 OSA-ICC User's Guide
- SG24-6364 OSA-ICC Implementation Guide Redbook
- SC23-2266 OSA-Express3 ICC Dual-Port User's Guide
- OSA-E Integrated Console Controller Presentation
 - <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS3591>



End of Topic

