

# *OSAs: What you know you know, What you know you don't know, and What you don't know you don't know...*

Gordon Webber  
William Data Systems

August 2010



# Agenda ...

***To identify and access the information available  
to manage an OSA:***

- **What is an OSA?**
  - **Overview**
  - **Functionality**
  - **Construction – the required environment**
  
- **MIBs ...**
- **What Information is available**
  
- **Accessing the OSA data**
- **Making sense of the data**
- **OSAENTA?**
  
- **Summary**

# What is an OSA?



**SHARE** in Boston

## What is an OSA? - Overview ...

- High-speed connection hardware that provides all of the required networking functionality (internal and external) to the Sysplex : IBM's premier device for connecting the mainframe with the outside world of IP!
- However, it is not a “black-box” in the same sense as a router. There are many components that comprise the path from application to WAN.
- The OSA device(s) are connected via the IO sub-system (CHPIDS).
- OSAs can connect to multiple LPARS ; up to 640 stacks...  
*(an important point for later...)*
- An LPAR can connect to multiple OSAs.
- The OSA is “operated” by Comm. Server using 4 Priority Queues.
- OSAs are managed by assorted software on the host.

## What is an OSA? - Overview ...

- The LPARS an OSA connects to may not all be z/OS
- OSAs hold management and performance data : This data is accessed by an Agent (IOBSNMP on z/OS) which presents it to the requester as an industry-standard MIB II.
- The range of OSAs available can present different “types” of device, for different functionality:
  - type **OSD** – for **QDIO** (Queued Direct I/O)
  - type **OSE** – for **non-QDIO**
  - type **OSC** – for the Integrated Console Controller function
  - type **OSN** – for NCP support under CCL

*(OSC and OSN are specialized functions, and omitted from this presentation.)*

# What is an OSA? - Overview ...

**Wait !** , why do we need to know this? ...

- OSA management requires data collection from multiple locations, using multiple tools/techniques, and therefore has a pre-requisite that the operator fully understands the OSA configuration, and the resource names in detail.
- Without this knowledge it is impossible to locate/collect the data that is available.
- This information is required to then explain/interpret the OSA data returned and the resource relationships in z/OS, as follows ...

# What is an OSA? - Functionality ...

## QDIO

- QDIO is a mode/protocol created by IBM to empower these high-speed (gigabit) OSAs.
- It encompasses such techniques as “fast-path” I/O, direct memory addressing (latest h/ware) which is reminiscent of the 3745 direct buffer addressing, LPAR-to-LPAR comms, and **configuration-from-host**.
- QDIO mode is required for many configurations (e.g. “Enterprise Extender” and Hipersockets {technically “Internal QDIO”}).  
*(Hipersockets: any-to-any connx between z/OS, z/VM, and zLinux)*
- Non-QDIO mode only supports the earlier methods of SNA communications as an XCA, runs channel programs ...  
and requires **OSA/SF** ...

# What is an OSA? - Functionality ...

## OSA/SF (Support Facility)

- OSAs require an OAT (OSA Address Table): this assigns/defines the path between an LPAR and the OSA port, including the IP addresses and MAC address, as required.
- When using QDIO mode, this OAT is created/updated dynamically (esp. from TCPIP profile data), therefore OSA/SF is not a requirement : If **not** using QDIO then in most cases (SNA), OSA/SF is a requirement.
- Since all OSAs have an OAT, OSA/SF can be used to retrieve information from that table. ( *how common is this?* )



# What is an OSA? - Functionality ...

- **OSAs also ...**

Perform router-like functions:

- o True routing
- o IP forwarding
- o ARP caching / responding (inc VIPAs)
- o VLAN support
- o MAC level support (& VMACs)
- o Checksum offloading
- o And many other performance-related and high-availability (aggregation) assists.
- o **And** support MIF (Multi-Image Facility – *more later ...*)

# What is an OSA? - Functionality ...

## IOBSNMP

- IOBSNMP is the “Subagent” found in z/OS (*other platforms handle this differently*).
- This is the agent that sits between the OS and the OSA, and intercepts the request made by the user/user application of the OSA data.
- IOBSNMP services the MIB responses back to the SNMP agent (OSNMPD).
- SNMP requests raised by the user/application, are directed at the OSNMP daemon, which directs the request via TCP/IP to the OSA Subagent (IOBSNMP) which, in turn, gets the data from the OSA.

**Before we can  
use an OSA ...**




**SHARE** in Boston

# What is an OSA? - Construction ...

## Components:

- Addresses – presentation to the Ethernet.
- Physical Connections
- **The “path”** ... 



# What is an OSA? - Construction ...

To use an OSA we need the “PATH” from OS to the Network:

- TCP/IP actually uses a VTAM TRL definition to connect to the OSA:

```
TRL400      VBUILD TYPE=TRL  
            TRLE   LNCTL=MPC,  
                MPCLEVEL=QDIO,  
                READ=0400,  
                WRITE=0401,  
                DATAPATH=(0402,0403),  
                PORTNAME=OSA24,  
                PORTNUM=0
```

- READ/WRITE must be an EVEN/ODD pair.
- DATAPATH is the list of devices available for data transfer/OSAENTA
- PORTNAME *must* match the name in the TCP/IP Profile!!
- PORTNUM is the physical port number on the OSA card.

# What is an OSA? - Construction ...

... In the TCP/IP Profile we define:

- The Device & Link (*as in this extract*):

```
DEVICE OSA24 MPCIPA NONROUTER
LINK   LNKOSA40 IPAQGNET OSA24
. . .
DEVICE VIPA VIRTUAL 0
LINK   VIPALINK VIRTUAL 0 VIPA
. . .
HOME
  10.5.1.10      LNKOSA40
  10.5.1.160    VIPALINK
. . .
BEGINROUTES
  ROUTE 10.5.1.0  HOST    =  LNKOSA40  MTU  4092
  ROUTE DEFAULT  10.5.1.1  LNKOSA40  MTU  1500
ENDROUTES
. . .
START OSA24
```

# What is an OSA? - Construction ...

To z/OS we will have previously defined the Control Unit and Devices via the HCD and an IOCPgen...

This information can be retrieved from an HCD report:

```
z/OS V1.10 HCD
Command ==> _____

Hardware Configuration

Select one of the following.

1. Define, modify, or view configuration data
2. Activate or process configuration data
3. Print or compare configuration data
4. Create or view graphical configuration report
5. Migrate configuration data
6. Maintain I/O definition files
7. Query supported hardware and installed UIMs
8. Getting started with this dialog
9. What's new in this release

For options 1 to 5, specify the name of the IODF to be used.
I/O definition file . . . 'SYS1.IODF10' +
```





# What is an OSA? - Construction ...

... And via z/OS Commands:

That tell use the status of the device(s) and who is using them :

```
-----
SDSF SYSLOG 7911.101 LW10 LW10 10/29/2009 0w 287006 COMMAND ISSUED
```

```
COMMAND INPUT ==> /D U,,,0400,8 SCROLL ==> CSR
RESPONSE=ADCD
```

```
IEE457I 11.08.07 UNIT STATUS 435
UNIT TYPE STATUS VOLSER VOLSTATE
0400 OSA A-BSY
0401 OSA A
0402 OSA A-BSY
0403 OSA O
0404 OSA OFFLINE
0405 OSA OFFLINE
0406 OSA OFFLINE
0407 OSA OFFLINE
```

```
COMMAND INPUT ==> /D U,,ALLOC,0400,4 SCROLL ==> CSR
```

```
RESPONSE=ADCD
IEE106I 11.27.24 UNITS ALLOCATED 697
UNIT JOBNAME ASID JOBNAME ASID
0400 VTAM 001D
0401 VTAM 001D
0402 VTAM 001D
0580 BACKTHR 01B6
```

# What is an OSA? - Construction ...

... Finally: some “Missing Links” in OSA terminology:

**CSS** - Channel Sub System. There can be more, but generally we find one CSS per LPAR (?).  
This defaults to an id number of “0”, and will appear as part of the “image” id...

**IMAGE** - This is the identifier used by the Multiple Image Facility (“MIF”).  
The image ID is made up as follows:

The image ID is “**x.y**” where “x” is the CSS number and “y” is the MIF id,

So, commonly “**0.1**”

The OSA MIB prefixes this with index value (*more later*) and the result is a value such as “**3.0.1**” (*remember this for later!*)

# What is an OSA? - Construction ...

## ...To recap:

- TRL/TRLE - defined by VTAM major node (also links OSA Port)
  - o verified/interrogated by VTAM display commands
- Device /Link - defined in TCP/IP profile (also assigns IP addresses)
  - o verified/interrogated by TCPIP display commands
- CUU/Device - defined in HCD/IOCP
  - o verified/interrogated by MVS display commands
- CHPID/PCHID - defined by HCD/IOCP
- **“OSA view”** - overall status and performance data is in the **MIB** ...

# MIB Data *(a brief diversion)*



**SHARE** in Boston

# Management Information Base ...

## MIB Objects ...

- **OBJECT-TYPE**
  - String that describes the MIB object.
  - Object Identifier (**OID**).
- **SYNTAX**
  - Defines what kind of data is stored in the MIB object.
- **ACCESS**
  - READ-ONLY, READ-WRITE.
- **STATUS**
  - State of object with regard to the SNMP community.
- **DESCRIPTION**
  - Reason why the MIB object exists.

Example of a Standard (**RFC1213**) MIB Object:

```
sysContact OBJECT-TYPE
    SYNTAX DisplayString (SIZE (0..255))
    ACCESS read-write
    STATUS mandatory
    DESCRIPTION "The textual identification of
the                contact person for this
                managed node,
                together with information on how to contact
                this person."
    ::= { system 4 }
```

"RIM" (Blackberry) MIB extract:

```
besConfigVersionString OBJECT-TYPE
    SYNTAX DisplayString
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "BlackBerry Server version information."
    ::= { besConfigEntry 10 }
```

## MIB Objects are Tables, Entries, and Indexes.

Imagine tables to be spreadsheets... consider the “**interface**” object:

Each interface type requires a “**row**” in the table : lets say 5 types (*5 indices*)

Each “**column**” represents a MIB object, as defined by the **entry** node

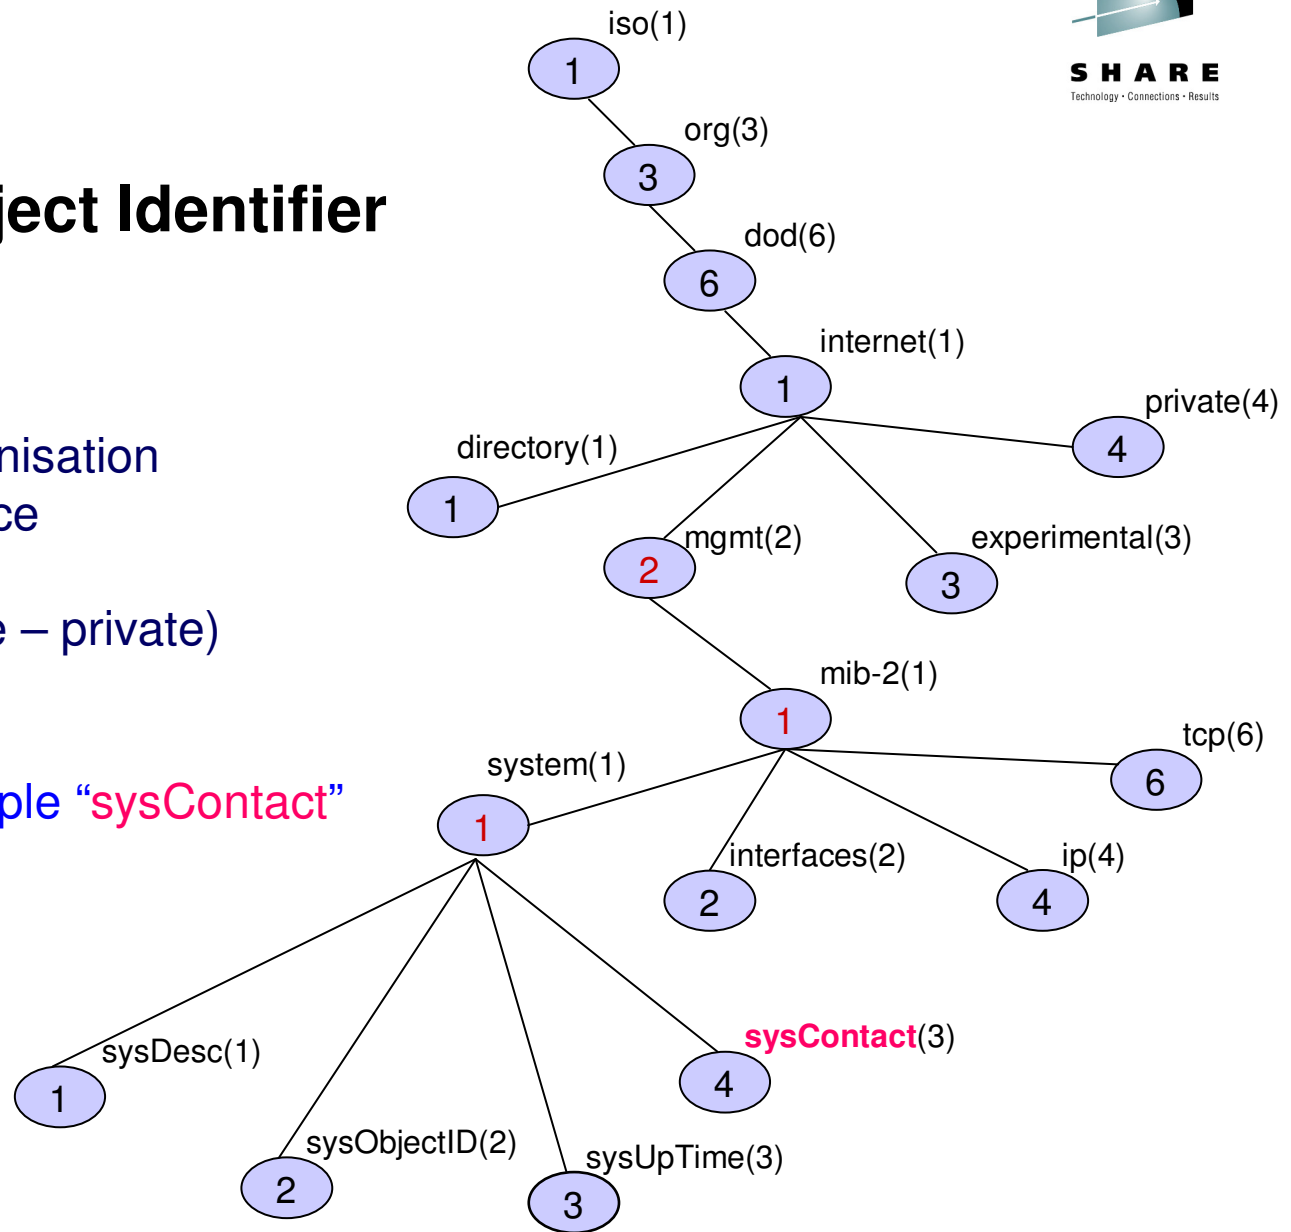
The “**spreadsheet**” might look like this:

ifEntry + ifIndex is used to locate the “Instance”	ifIndex	ifDescr	ifType	ifMtu	ifSpeed	ifPhysAddress	.etc.
	1	other	1	1500	0	-	-
	2	ddn-x25	4	-	256	-	-
	3	ethernet-csmacd	6	1500	100	Eth0	-
	4	iso88025-tokenRing	9	-	16	Tr0	-
	5	software Loopback	24	-	0	Lo	-
	.etc.	-	-	-	-	-	-

# The “OID”, or Object Identifier

E.g. 1.3.6.1.2.1.1.4

- 1 - ISO
- 3 - ISO recognised Organisation
- 6 - Department of Defence
- 1 - Internet type entry
- 2 - mgmt (alternative – private)
- 1 - MIB-2
- 1 - system mib
- 4 – is our previous example “sysContact”
- ..... etc .....



Host Name 192.168.1.231Community public MaxRequest 128

Object	Value
system	
interfaces	
ifNumber	3
ifTable	
ifEntry	
ifIndex	
.1	1
.2	2
.3	3
ifDescr	(I)
ifType	(I)
ifMtu	(I)
ifSpeed	(I)
ifPhysAddress	(I)
ifAdminStatus	(I)
ifOperStatus	(I)
ifLastChange	(I)
ifInOctets	(I)
ifInUcastPkts	(I)
ifInNUcastPkts	(I)
ifInDiscards	(I)
ifInErrors	(I)
ifInUnknownProtos	(I)
ifOutOctets	(I)
ifOutUcastPkts	(I)
ifOutNUcastPkts	(I)
ifOutDiscards	(I)
ifOutErrors	(I)
ifOutQLen	(I)
ifSpecific	(I)
at	
ip	
icmp	
tcp	

## Value Selection

g Get  
n GetNext  
b GetBulk  
w Walk/Show  
m Monitor  
i Index View

This display is the result of a "GetBulk" and the expansion of the **interfaces** MIB.

The "(I)" descriptor signifies that these objects can be "indexed" for each of the 3 interfaces.

Interface 2 is the "ethernet", and will be selected using the "i" command...

Objects 265

7671



Host 192.168.1.231  
Index .2

Object	Value
ifIndex	2
ifDescr	eth0
ifType	ethernet-csmacd
ifMtu	1500
ifSpeed	95m
ifPhysAddress	No Data
ifAdminStatus	1
ifOperStatus	1
ifLastChange	---
ifInOctets	926m
ifInUcastPkts	7004k
ifInNUcastPkts	---
ifInDiscards	0
ifInErrors	0
ifInUnknownProtos	---
ifOutOctets	1421m
ifOutUcastPkts	54m

Update MIB Monitor

ADCDPL

P390 TCPIP

14:54:38

### MIB Details

Host 192.168.1.231 Community public  
ObjectID 1.3.6.1.2.1.2.2.1.8.2  
Name ifOperStatus.2

### Monitor Details

Interval	0	Frequency (minutes) object value will be monitored
Low Value	0	Alert if object value is less than this
High Value	0	Alert if object value is more than this
Monitor Id		Displayed in alert messages

Objects 270

7791

## Mapping the “OIDs” ...

' system	'	OID='1.3.6.1.2.1.1.'
' interfaces	'	OID='1.3.6.1.2.1.2.'
' at	'	OID='1.3.6.1.2.1.3.'
' ip	'	OID='1.3.6.1.2.1.4.'
' icmp	'	OID='1.3.6.1.2.1.5.'
' tcp	'	OID='1.3.6.1.2.1.6.'
' udp	'	OID='1.3.6.1.2.1.7.'
' egp	'	OID='1.3.6.1.2.1.8.'
' ipoaList	'	OID='1.3.6.1.2.1.10.'
' snmp	'	OID='1.3.6.1.2.1.11.'
' ospf	'	OID='1.3.6.1.2.1.14.'
' ifXTable	'	OID='1.3.6.1.2.1.31.1.1.'
' atmInterfaceConfTable	'	OID='1.3.6.1.2.1.37.1.2.'
' ibmResearch	'	OID='1.3.6.1.4.1.2.2.'
' sa	'	OID='1.3.6.1.4.1.2.4.12.'
' IBM3172	'	OID='1.3.6.1.4.1.2.6.1.'
' IBMTcpip	'	
OID='1.3.6.1.4.1.2.6.19.2.'		
' OSAExpress-OSA/SF	'	OID='1.3.6.1.4.1.2.6.19.2.2.10.'
' OSAExpress-Direct	'	OID='1.3.6.1.4.1.2.6.188.1.'
' usecAgent	'	OID='1.3.6.1.6.3.6.1.'

**Entries can be “scalar” (arithmetic value – cf ifOperStatus )**  
**or “columnar” (tabular data).**

# Available Data



**SHARE** in Boston

# MIBs – General Data ...

OID:  Operations:

Name/OID	Value
sysDescr.0	Sysname: OS/390 Nodename: ADCD Release: 20.00 Version: 03 Machine: 2096
sysObjectID.0	.1.3.6.1.4.1.2.3.13
sysUpTime.0	532 hours 52 minutes 17 seconds
sysContact.0	SNMPBASE - Unspecified
sysName.0	SNMPBASE - Unspecified
sysLocation.0	SNMPBASE - Unspecified
sysServices.0	0
.1.3.6.1.2.1.1.8.0	532 hours 48 minutes 13 seconds
.1.3.6.1.2.1.1.9.1.2.1	.1.3.6.1.4.1.2.11.7.1
.1.3.6.1.2.1.1.9.1.2.2	.1.3.6.1.4.1.2.11.26.1
.1.3.6.1.2.1.1.9.1.2.3	.1.3.6.1.4.1.2.11.7.2
.1.3.6.1.2.1.1.9.1.3.1	z/OS SNMP Agent
.1.3.6.1.2.1.1.9.1.3.2	OSA subagent
.1.3.6.1.2.1.1.9.1.3.3	z/OS TCP/IP SNMP Subagent
.1.3.6.1.2.1.1.9.1.4.1	0 millisecond
.1.3.6.1.2.1.1.9.1.4.2	532 hours 48 minutes 13 seconds
.1.3.6.1.2.1.1.9.1.4.3	23 seconds

# MIBs – Interface Information ...

ifNumber.0	11
ifIndex.1	1
ifIndex.2	2
ifIndex.3	3
ifIndex.4	4
ifIndex.5	5
ifIndex.6	6
ifIndex.7	7
ifIndex.8	8
ifIndex.9	9
ifIndex.10	10
ifIndex.11	11
ifDescr.1	Loopback Device
ifDescr.2	Loopback
ifDescr.3	Loopback IPv6
ifDescr.4	Multipath Channel IP Assist Device
ifDescr.5	IP Assist QDIO Ethernet
ifDescr.6	Multipath Channel Point-to-Point Device
ifDescr.7	Multipath Channel Point-to-Point
ifDescr.8	Virtual IP Address Device
ifDescr.9	Virtual IP Address Link
ifDescr.10	Virtual IP Address Device
ifDescr.11	Virtual IP Address Link
ifType.1	propVirtual
ifType.2	softwareLoopback
ifType.3	softwareLoopback
ifType.4	propVirtual
ifType.5	ethernetCsmacd
ifType.6	propVirtual
ifType.7	mpc
ifType.8	propVirtual
ifType.9	virtualIpAddress
ifType.10	propVirtual
ifType.11	virtualIpAddress
ifMtu.1	0
ifMtu.2	65535
ifMtu.3	65535
...	-

z/OS Interfaces ...

← Note the “index” values!  
 (“4” is for the OSA...)

Loopback Device  
 Loopback  
 Loopback IPv6  
 Multipath Channel IP Assist Device  
 IP Assist QDIO Ethernet  
 Multipath Channel Point-to-Point Device  
 Multipath Channel Point-to-Point  
 Virtual IP Address Device  
 Virtual IP Address Link  
 Virtual IP Address Device  
 Virtual IP Address Link

# MIBs - Mapping the OIDs ...

Loopback Device	
Loopback	
Loopback IPv6	
Multipath Channel IP Assist Device	DEVICE OSA24 MPCIPA NONROUTER
IP Assist QDIO Ethernet	LINK LNKOSA40 IPAQGNET OSA24
Multipath Channel Point-to-Point Device	DEVICE IUTSAMEH MPCPTP AUTORESTART
Multipath Channel Point-to-Point	LINK IUTSAMEH MPCPTP IUTSAMEH
Virtual IP Address Device	DEVICE EEVIPAD VIRTUAL 0
Virtual IP Address Link	LINK EEVIPAL VIRTUAL 0 EEVIPAD
Virtual IP Address Device	DEVICE VIPA VIRTUAL 0
Virtual IP Address Link	LINK VIPALINK VIRTUAL 0 VIPA

Extract from TCP/IP Profile :

# MIBs – The OSA Mapping ...

OID:  Operations:

Name/OID	Value
ibmOSAEExpChannelType.4	osaDirectExpress
ibmOSAEExpChannelHdwLevel.4	4
ibmOSAEExpChannelSubType.4	97
ibmOSAEExpChannelShared.4	shared
ibmOSAEExpChannelNodeDesc.4	0x20 0x00 0x06 0x24 0xF0 0xF0 0xF1 0xF7 0xF3 0xF0 0xF0 0xF0 0xF4 0xC9 ...
ibmOSAEExpChannelProcCodeLevel.4	0x08 0x92
ibmOSAEExpChannelPCIBusUtil1Min.4	0
ibmOSAEExpChannelProcUtil1Min.4	
ibmOSAEExpChannelPCIBusUtil5Min.4	
ibmOSAEExpChannelProcUtil5Min.4	
ibmOSAEExpChannelPCIBusUtilHour.4	0
ibmOSAEExpChannelProcUtilHour.4	0
ibmOSAEExpPEMaxSizeArpCache.4	8192
ibmOSAEExpPEArpPendingEntries.4	0
ibmOSAEExpPEArpActiveEntries.4	4
ibmOSAEExpPEIPEntries.4	15
ibmOSAEExpPEMulticastEntries.4	3
ibmOSAEExpChannelNumber.4	0x00 0x24

E.G. note this OID - 1.3.6.1.4.1.2.6.188.1.10.1.46.1.1.1  
... the **“ibmOsaExp3ReceiveLenErrorCount”** !

## MIBs – “Catch 22” ...

**For OSA MIBS, you only see YOUR LPAR data!**

- Note the “indices” in the general RFC1213 MIB.
- The MIB data will be filtered to reflect the point of view of the system from which the request comes.
- To see a complete picture, you must collect the data from all paths to the OSA/MIB!



# Accessing the Data



**SHARE** in Boston

# Accessing the OSA data ...

- **Console Commands!**

COMMAND INPUT ==> /D M=DEV(0400)

```

RESPONSE=ADCD
IEE174I 11.36.44 DISPLAY M 79
DEVICE 0400 STATUS=ONLINE
CHP 24
ENTRY LINK ADDRESS ..
DEST LINK ADDRESS 0D
PATH ONLINE Y
CHP PHYSICALLY ONLINE Y
PATH OPERATIONAL Y
MANAGED N
CU NUMBER 0080
MAXIMUM MANAGED CHPID(S) ALLO
DESTINATION CU LOGICAL ADDRESS
SCP CU ND = 001730.00
SCP TOKEN NED = 001730.00
SCP DEVICE NED = 001732.00
    
```

/D M=CHP

IEE174I 11.39.57 DISPLAY M 800

```

CHANNEL PATH STATUS
      0 1 2 3 4 5 6 7 8 9 A B C D E F
0 + + + + + + + + + + + + + +
1 + + + + + + + + + + + + + +
2 + + + + + + + + + + + + + +
3 + + + + + + + + + + + + + +
. . .
    
```

\*\*\*\*\* SYMBOL EXPLANATIONS \*\*\*\*\*  
+ ONLINE @ PATH NOT VALIDATED - OFFLINE . DOES NOT EXIST  
• MANAGED AND ONLINE # MANAGED AND OFFLINE

```

CHANNEL PATH TYPE STATUS
      0 1 2 3 4 5 6 7 8 9 A B C D E F
0 00 00 00 00 00 00 1A 1A 1A 1A 00 00 00 00 06 00
1 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
2 00 14 11 11 11 00 00 00 00 00 00 00 00 00 00 00
3 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
. . .
    
```

```

10 OSA EXPRESS OSE
11 OSA DIRECT EXPRESS OSD
12 OPEN SYSTEMS ADAPTER OSA
13 INTERNAL SYSTEM DEVICE ISD
14 OSA CONSOLE OSC
15 OSA NCP OSN
    
```

# Accessing the OSA data ...

...And more commands:

- **NETSTAT ...**

“onetstat”  
“TSO NETSTAT xxx”  
“D TCPIP,,NETSTAT,xxx”

Where xxx can be...

ARP - ARP cache  
DEV - Device details  
HOME - Home IP Addresses  
ROUTE - Routing table  
etc...

```
/D TCPIP,,NETSTAT,DEV  
  
RESPONSE=ADCD  
EZD0101I NETSTAT CS V1R10 TCPIP 399  
HOME ADDRESS LIST:  
LINKNAME: LNKOSA40  
ADDRESS: 10.5.1.10  
FLAGS: PRIMARY  
LINKNAME: EEVIPAL  
ADDRESS: 10.5.1.150  
FLAGS:  
LINKNAME: VIPALINK  
ADDRESS: 10.5.1.160  
FLAGS:  
.  
.  
.
```

# Accessing the OSA data ...

**/D TCPIP, ,NETSTAT,DEV**

```

DEVNAME: OSA24                DEVTYPE: MPCIPA
DEVSTATUS: READY
LNKNAME: LNKOSA40            LNKTYPE: IPAQENET   LNKSTATUS: READY
SPEED: 0000000100
IPBROADCASTCAPABILITY: NO
CFGROUTER: NON                ACTROUTER: NON
ARPOFFLOAD: YES               ARPOFFLOADINFO: YES
ACTMTU: 1492
VLANID: NONE                  VLANPRIORITY: DISABLED
READSTORAGE: GLOBAL (4096K)   INBPERF: BALANCED
CHECKSUMOFFLOAD: YES
SECCLASS: 255                 MONSYSPLEX: NO

BSD ROUTING PARAMETERS:
MTU SIZE: N/A                 METRIC: 00
DESTADDR: 0.0.0.0            SUBNETMASK: 255.255.255.0
PACKET TRACE SETTING:
PROTOCOL: 253                 TRRECCNT: 000
DISCARD: NONE                 DESTPORT: *
SRCPORT: *                     SUBNET: *
IPADDR: *
    
```

**MULTICAST SPECIFIC:**

GROUP	REFCNT	SRCFLTMD
224.0.0.1	0000000001	EXCLUDE

SRCADDR: NONE

**LINK STATISTICS:**

```

BYTESIN = 8635841755
INBOUND PACKETS = 21945924
INBOUND PACKETS IN ERROR = 273215
INBOUND PACKETS DISCARDED = 0
INBOUND PACKETS WITH NO PROTOCOL = 0
BYTESOUT = 6091250164
OUTBOUND PACKETS = 17393059
OUTBOUND PACKETS IN ERROR = 0
OUTBOUND PACKETS DISCARDED = 0
    
```

# Accessing the OSA data ...

...And more commands: VTAM

- **D NET,xxxx**

Where xxxx can be...

VTAMOPTS - ARP cache

ID=nnnnn - Device details

TRL - Transport Resource List

TRLE - TRLE Entry

```
D NET,TRL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TRL 632
IST924I -----
. . .
. . .
IST1954I TRL MAJOR NODE = P10TRL
IST1314I TRLE = TRL400 STATUS = ACTIV CONTROL = MPC
IST1454I 1 TRLE(S) DISPLAYED
IST314I END
```

## Accessing the OSA data ...

**D NET,TRL,TRLE=TRL400**

```
IST097I DISPLAY ACCEPTED
IST075I NAME = TRL400, TYPE = TRLE 652
IST1954I TRL MAJOR NODE = P10TRL
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST087I TYPE = LEASED, CONTROL = MPC , HPDT = YES
IST1715I MPCLEVEL = QDIO MPCUSAGE = SHARE
IST1716I PORTNAME = OSA24 LINKNUM = 0 OSA CODE LEVEL = 0892
IST1577I HEADER SIZE = 4096 DATA SIZE = 0 STORAGE = ***NA***
IST1221I WRITE DEV = 0401 STATUS = ACTIVE STATE = ONLINE
IST1577I HEADER SIZE = 4092 DATA SIZE = 0 STORAGE = ***NA***
IST1221I READ DEV = 0400 STATUS = ACTIVE STATE = ONLINE
IST1221I DATA DEV = 0402 STATUS = ACTIVE STATE = N/A
IST1724I I/O TRACE = OFF TRACE LENGTH = *NA*
IST1717I ULPID = TCPIP
IST1815I IQDIO ROUTING DISABLED
IST1918I READ STORAGE = 4.0M(64 SBALS)
IST1757I PRIORITY1: UNCONGESTED PRIORITY2: UNCONGESTED
IST1757I PRIORITY3: UNCONGESTED PRIORITY4: UNCONGESTED
IST2190I DEVICEID PARAMETER FOR OSAENTA TRACE COMMAND = 00-01-00-05
IST1801I UNITS OF WORK FOR NCB AT ADDRESS X'0FEAF010'
```

“Congestion State (congstate) specifies the state of this priority level. It will be CONGESTED when, at least once within the last congestion reporting window, all 128 writes for the priority level were unavailable...otherwise congstate will be UNCONGESTED.”

# Why You Need This! ...

## So why go to all this trouble to monitor your OSA?

- **ALL of your network traffic comes through the OSA!**
- What you see from any one position may not be the whole story!
- A 5% increase in bandwidth use per path, may be 20%+ overall!
- You should be aware of :
  - state changes (*in time to react!*)
  - performance thresholds
  - error counter (no errors... are you **sure** !?)
- What about the Queue (remember, p1 - p4?) depths? Are the performing correctly??
- Congestion – a fluctuating target! (remember that a congested state relates to a Priority Level ... *what exactly does this mean??*)

# Why You Need This! ...

## Managing the data ...

OSA/SF: Recommended by IBM as a tool for interrogating OSAs ;  
*(even if not required for setting up the OSA)*

- OSA/SF runs a an STC
- OSA/SF is operated via APPC, therefore a VTAM LU is required.
- Access is via a workstation GUI, or REXX (TSO)
- Provides such data as:
  - o Managing message logs
  - o Access to MIB data
  - o Debugging aids



# Why You Need This! ...

## Managing the data ...

- Clear the error message cache! (*note OSA/SF functions*)
- Monitor the Congestion state:— “**D NET,TRL**”, then...
  - Search the returned lists for the TRLE name you require
  - Then repeat the display using “...TRLE=*name*”
- Browse the MIB!
- TSO “SNMP” command ... but you need to know which “OID” you require  
*(how frequently is this required?)*
- RMF also collects some performance and availability data.  
*(see: Monitor I & II Channel Path Activity Report)*

*... And remember that z/OS only shows the view of the MIB (remember the index value) pertinent to that LPAR/Stack/or/Link !*

# Making Sense of the Data



**SHARE** in Boston

OSA Dashboard

### OSA Summary

OSA	LPARs	Channels	Ethernet Ports	PE Table	Processor Busy % (1m)	PCI Busy % (5m)	Processor Busy % (5m)	PCI Busy % (1h)	Processor Busy % (1h)
OSA2	0	1	0	1	0	0	0	0	0
OSA1	1	1	1	1	0	0	0	0	0
Total									

OSA name

### Channel Summary for OSA: OSA2

Port	CHPID	PCI Busy %(1m)	Processor Busy %(1m)	PCI Busy %(5m)	Processor Busy %(5m)	PCI Busy %(1h)	Processor Busy %(1h)
x'0024'	0	0	0	0	0	0	0
Total	0	0	0	0	0	0	0

CHPID

### Channel Summary for OSA: OSA1

Port	CHPID	PCI Busy %(1m)	Processor Busy %(1m)	PCI Busy %(5m)	Processor Busy %(5m)	PCI Busy %(1h)	Processor Busy %(1h)
OSA408	x'0021'	0	0	0	0	0	0
Total		0	0	0	0	0	0

Port name

### LPAR Summary

Image Id	Processor Util 1 Min	In Kbytes Rate 1 Min	Out Kbytes Rate 1 Min	Processor Util 5 Min	In Kbytes Rate 5 Min	Out Kbytes Rate 5 Min	Processor Util Hour	In Kbytes Rate Hour	Out Kbytes Rate Hour
3.0.1	0	2	6	0	2	6	0	2	5

Image ID

14:25:20

Link name

Links Dashboard

LinkName	OSA	IP Address	Port	CHPID	Stack Name	MIB Index	Speed/Sec In/Out	% Util In/Out	VTAM State	Congestion P1 P2 P3 P4				Alerts	Online
LNKOSA40	OSA2	10.5.1.10	OSA24		TCPIP		1KB 922	0/0	ACTIV	●	●	●	●	0	●
LNKOSA40	OSA2	10.5.1.19	OSA24				0 0	0/0	***	●	●	●	●	0	●
LNKOSA44	OSA2	192.168.1.5					0 0	0/0	***	●	●	●	●	0	●
LNKOSA48	OSA1	10.5.1.20					0 0	0/0	***	●	●	●	●	0	●


Links by Stack: TCPIP

LinkName	OSA	IP Address	Port	CHPID	MIB Index	Speed/Sec In/Out	% Util In/Out	VTAM State	Congestion P1 P2 P3 P4				Alerts	Online
LNKOSA40	OSA2	10.5.1.10	OSA24			317 206	0/0	ACTIV	●	●	●	●	0	●

Or drawing on the data relative to the device ...

OSA Device Information

OSA	Stack Name	Link Name	Device Name	IP Address	Start Time	State	Speed	MTU Size	Inbound Packets	Outbound Packets	Inbound Bytes	Outbound Bytes
OSA2	TCPIP	LNKOSA40	OSA24	10.5.1.10	Sep 19 20:03:13	Online	100	1492	21700295	17440572	8GB	6GB



SHARE

Priority Queues  
-  
Congestion indicators

Link: LNKOSA40 Information					
<b>Device Information</b>					
Link Name	LNKOSA40				
Device Name	OSA24				
Stack Name	TCPIP				
IP Address	10.5.1.10				
Port	OSA24				
Type	OSA Express				
Started At	Sep 19 20:03:13				
State	Online				
Alerts	0				
<b>VTAM TRLE Information</b>					
TRLE Name	TRL400				
Major Node	P10TRL				
DeviceId parm for OSAENTA Trace Cmd	00-01-00-05				
Microcode Level	0892				
VTAM Status	ACTIV				
Read Device and State	0400	ACTIVE			
Write Device and State	0401	ACTIVE			
Data Device and State	0402	ACTIVE			
Trace Device and State	0403	RESET			
PR1 Congestion / Queue lengths Curr/Avg/Max	●	0	2	18	
PR2 Congestion / Queue lengths Curr/Avg/Max	●	0	0	0	
PR3 Congestion / Queue lengths Curr/Avg/Max	●	0	1	4	
PR4 Congestion / Queue lengths Curr/Avg/Max	●	0	1	18	
<b>Throughput Information for LPAR: LW10, Stack: TCPIP</b>					
Max Speed	100Mbps				
MTU	1492				
Packets In	21697979				
Packets Out	17437883				
Bytes In	8GB				
Bytes Out	6GB				
Current speed/sec In	957				
Current speed/sec Out	718				
Current %Util In	0				
Current %Util Out	0				
Average speed/sec In	35				
Average speed/sec Out	194				
Average %Util In	0				
Average %Util Out	0				
<b>OSA Component Information</b>					
OSA					OSA2
Channels					1
14:31:29					

Detail MIB Display for 10.5.1.20

Type	Data	Description
Sequence		
Integer	1	SNMP Version
CharacterString	public	Password
GetResponse PDU		
Integer	296142560	Request Id
Integer	0	Error Count
Integer	0	Error Index
Sequence		
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.1.3	ibmOSAEExpChannel
CharacterString		
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.2.3	ibmOSAEExpChannel
Integer	17	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.3.3	ibmOSAEExpChannelHdwLevel.3
Integer	4	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.4.3	ibmOSAEExpChannelSubType.3
Integer	97	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.5.3	ibmOSAEExpChannelShared.3
Integer	1	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.6.3	ibmOSAEExpChannelNodeDesc.3
CharacterString		
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.8.3	ibmOSAEExpChannelProcCodeLevel.3
CharacterString	□k	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.9.3	ibmOSAEExpChannelPCIBusUtil1Min.3
Integer	0	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.10.3	ibmOSAEExpChannelProcUtil1Min.3
Integer	0	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.11.3	ibmOSAEExpChannelPCIBusUtil5Min.3
Integer	0	
Sequence		
OID	1.3.6.1.4.1.2.6.188.1.1.1.12.3	ibmOSAEExpChannelProcUtil5Min.3

Basic MIB Display...  
With data values and applied MIB mapping

OSA:OSA1 Channel: LNKOSA48

**OSA:OSA1 Channel: LNKOSA48**

Field	Value	Description
Channel Number	x'0022'	The CHPID corresponding to this ifIndex.
Type	osa Direct Express	The type of channel for this interface. OSA Express (OSE) has a value of 16. OSA Direct Express (OSD) has a value of 17.
Hdw Level	osaExp300	Hardware model of the channel. The value osaExp150(2) indicates a feature as OSA-Express. The value osaExp175(3) indicates a hardware as OSA-Express. The value osaExp300(4) indicates a hardware level Express2. The value osaExp400(5) indicates a hardware level of 4.00 Express3.
Sub Type	One Thousand BaseT Ethernet	Indicates the type of OSA feature present. Each value further indicates a value of gigabitEthernet, fastEthernet or oneThousandBaseTEthernet. A value of tokenRing returns the ibmOSAExpTRPortTable. A value of tenGigabitEthernet returns the ibmOSAExp10GigEthPortTable. A value of osaexp3gigabitEthernet, osaexp3oneThousandBaseTEthernet returns the ibmOSAExp3PortTable. A value of atmEmulatedEthernet returns the ibmOSAExpATMPortTable.
Shared Channel	Being Shared	All OSA features can be shared across multiple images. This object indicates if this feature is currently being shared.
Node Desc	x'20000622' x'F0F0F1F7' x'F3F0F0F0' x'F4C9C2D4' x'F8F3F0F0' x'F0F0F0F0' x'F0C6F8F6' x'F4F32200'	This is the Node Descriptor of the OSA feature. It represents the ND obtained from the Channel Subsystem.
Processor Code Level	x'0892'	This is the Licensed Internal Code (micro code level) of the OSA feature. For example, OSA code level 05.6A would be represented as 0x056A.
PCI Busy % (1m)	0	The average, over a 1 minute interval, of the percentage of time that the PCI bus was utilized to transfer data. It does not include idle time or time used by routine maintenance tasks. The range for this value is from 0 to 100%.
Processor Busy % (1m)	0	The average, over a 1 minute interval, of the percentage of time that the CHPID Processor was utilized to transfer data. It does not include idle time or time used by routine maintenance tasks. The range for this value is from 0 to 100%.
PCI Busy % (5m)	0	The average, over a 5 minute interval, of the percentage of time that the PCI bus was utilized to transfer data. It does not include idle time or time used by routine maintenance tasks. The range for this value is from 0 to 100%.
Processor Busy % (5m)	0	The average, over a 5 minute interval, of the percentage of time that the CHPID Processor was utilized to transfer data. It does not include idle time or time used by routine maintenance tasks. The range for this value is from 0 to 100%.
PCI Busy % (1h)	0	The average, over an hour interval, of the percentage of time that the PCI bus was utilized to transfer data. It does not include idle time or time used by routine maintenance tasks.
Processor Busy % (1h)	0	The average, over an hour interval, of the percentage of time that the CHPID Processor was utilized to transfer data. It does not include idle time or time used by routine maintenance tasks. The range for this value is from 0 to 100%.

**Interpreted MIB Data ...**

(relevant to the selected resource – here, the OSA name)

14:38:43

OSA: OSA1 LPAR Performance Table: ⏪ ⏩ ⏹

**OSA: OSA1 LPAR Performance Table:**

Field	Value	Description
CSS Id	3.0.1	The Logical Channel Subsystem the performance data refers to. For a z900 or z800 system, there is only 1 Logical Channel subsystem and will be indicated by a CSSid of 0. For z990 systems, numbered 0 to (n-1).
Processor Util 1 Min	0	The average, over a 1 minute interval, of the percentage of time that the OSA Processor was utilized to transfer data for a specific Image Id. It does not include idle time or time used by routine maintenance tasks. The range of valid values for this object is 0 to 100%. A value of -1 indicates that the value could not be retrieved from the adapter.
In Kbytes Rate 1 Min	2	The average, over a 1 minute interval, of the number of inbound kilobytes processed for a specific Image. The value is expressed in kilobytes per second. When the ibmOSAExpV2PerfProcUtil1Min object for a specific Image has a value of -1, then the interval data could not be retrieved from the OSA and this object will have a value of zero.
Out Kbytes Rate 1 Min	6	The average, over a 1 minute interval, of the number of outbound kilobytes processed for a specific Image. The value is expressed in kilobytes per second. When the ibmOSAExpV2PerfProcUtil1Min object for a specific Image has a value of -1, then the interval data could not be retrieved from the OSA and this object will have a value of zero.
Processor Util 5 Min	0	The average, over a 5 minute interval, of the percentage of time that the OSA Processor was utilized to transfer data for a specific Image Id. It does not include idle time or time used by routine maintenance tasks. The range of valid values for this object is 0 to 100%. A value of -1 indicates that the value could not be retrieved from the adapter.
In Kbytes Rate 5 Min	2	The average, over a 5 minute interval, of the number of inbound kilobytes processed for a specific Image. The value is expressed in kilobytes per second. When the ibmOSAExpV2PerfProcUtil5Min object for a specific Image has a value of -1, then the interval data could not be retrieved from the OSA and this object will have a value of zero.
Out Kbytes Rate 5 Min	5	The average, over a 5 minute interval, of the number of outbound kilobytes processed for a specific Image. The value is expressed in kilobytes per second. When the ibmOSAExpV2PerfProcUtil5Min object for a specific Image has a value of -1, then the interval data could not be retrieved from the OSA and this object will have a value of zero.
Processor Util Hour	0	The average, over a 60 minute interval, of the percentage of time that the OSA Processor was utilized to transfer data for a specific Image Id. It does not include idle time or time used by routine maintenance tasks. The range of valid values for this object is 0 to 100%. A value of -1 indicates that the value could not be retrieved from the adapter.
In Kbytes Rate Hour	2	The average, over a 60 minute interval, of the number of inbound kilobytes processed for a specific Image. The value is expressed in kilobytes per second. When the ibmOSAExpV2PerfProcUtilHour object for a specific Image has a value of -1, then the interval data could not be retrieved from the OSA and this object will have a value of zero.
Out Kbytes Rate Hour	5	The average, over a 60 minute interval, of the number of outbound kilobytes processed for a specific Image. The value is expressed in kilobytes per second. When the ibmOSAExpV2PerfProcUtilHour object for a specific Image has a value of -1, then the interval data could not be retrieved from the OSA and this object will have a value of zero.

Interpreted MIB Data ...

(relevant to the selected resource – here, the IMAGE ID)

14:39:44



OSA History Capture Admin

History File Status	
History Capture is:	Active
Sampling Interval is:	5 minutes

History File Contents	
Starting Date/Time:	Wed Aug 19 10:32:17 2009
Ending Date/Time:	Fri Oct 30 12:40:18 2009
Reference Date/Time:	Fri Oct 30 11:56:59 2009

History File Capacity	
History File Capacity is:	1260000 records, 315000 samples
History File Capacity is:	43 days, 18 hours

Suspend   Soft Reset   Cancel

12:44:21

Resetting the  
OSA error  
stats...

## Beyond z/OS ...

So far we have looked at the “common” usage:

... z/OS calling MIB data via the local agent, where the OSA is in QDIO mode ...

**BUT** there are other considerations:

- Not all z/OS systems run QDIO  
*(OSC is common, OSE is still in use?)*
- Not all OSAs run in QDIO mode
- Not all systems are z/OS !!  
*(What about z/VM & zLinux?)*

***We need to centralize the collection & monitoring ...***

OSA Configuration

Config file: /wds/support/zenprod/user/data/configuration.zom

✗	OSA	LinkName	IP Address	Port	Community Name	Alert	Status	Error
✗	OSA2	LNKOSA40	10.5.1.10	OSA24	●●●●●●	<input checked="" type="checkbox"/>	●	
✗	OSA2	LNKOSA40	10.5.1.19	OSA24	●●●●●●	<input checked="" type="checkbox"/>	●	IP is not a known OSA
✗	OSA2	LNKOSA44	192.168.1.5		●●●●●●	<input checked="" type="checkbox"/>	●	IP is not a known OSA
✗	OSA1	LNKOSA48	10.5.1.20		●●●●●●	<input checked="" type="checkbox"/>	●	IP is not a known OSA

New Apply Save Discover Cancel

**OSAs are accessible via IP !**

... we can address OSAs *not known* by this LPAR ... all we lose is the local data (TCPIP stats, etc)

# Then there's OSAENTA ...

## The OSA-Express Network Traffic Analyzer

- Requires a z9 or z10 with OSA-Express2/3 in QDIO mode (CHPID type OSD)
- Provides support to trace inbound and outbound traffic to the OSA
- Trace functions controlled and formatted by CS; collected at the OSA port.

### Set-up requires:

- The correct microcode level for the OSA
- a CHPID
- VTAM TRLE definitions (*as noted*)
- TCPIP definitions (*as noted*)
- Customizing OSAENTA itself
- Defining a resource profile in RACF (MVS.VARY.TCPIP.OSAENTA)
- Allocating a VSAM linear data set
- Create/start an OSAENTA trace

## Then there's OSAENTA ...

- Determining the microcode level for OSA/OSAExpress3:  
HMC console or D NET,TRL,TRLE=*trlname*

```
VBUILD TYPE=TRL
TRL400 TRLE LNCTL=MPC,
MPCLEVEL=QDIO,
READ=0400,
WRITE=0401,
DATAPATH=(0402,0403),
PORTNAME=OSA24
```

- ADD TRLE definitions

```
DATAPATH=(0402,0404), (check with NETSTAT DEVLINKS)
```

### Running the OSAENTA trace:

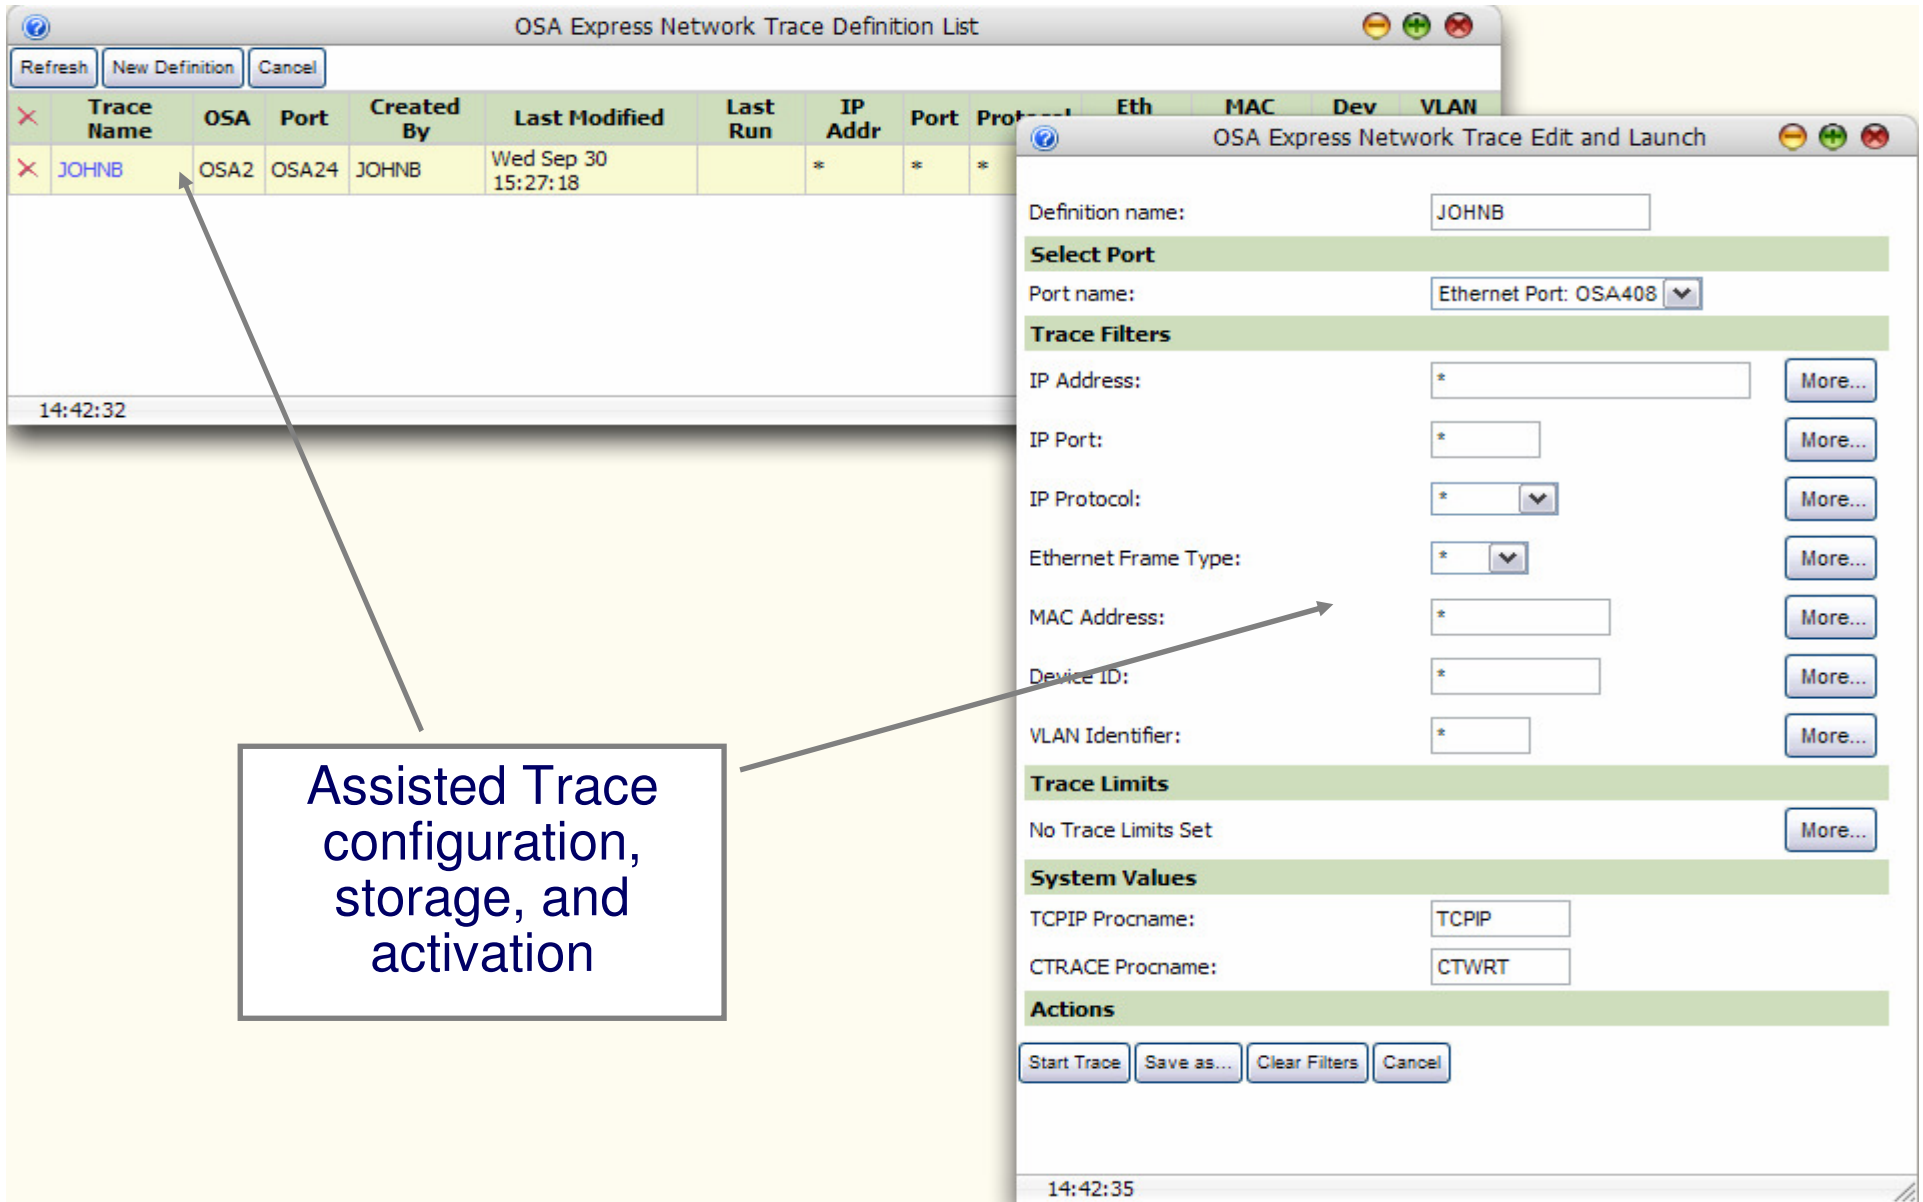
```
Start : TRACE CT,WTRSTART=WDSWRT
TRACE CT,ON,COMP=SYSTCPOT,SUB=(TCPIP) (reply WTR=WDSWRT)
(check it's active with NETSTAT DEVLINKS)

Show : D TRACE,COMP=SYSTCPOT,SUB=(TCPIPA)

Stop : TRACE CT,OFF,COMP=SYSTCPOT,SUB=(TCPIP)
TRACE CT,WTRSTOP=WDSWRT,FLUSH
```

- Analyze with IPCS, or a batch job

*(nb. OSAENTA can also be setup IN the TCPIP Profile)*



OSA Express Network Trace Definition List

Trace Name	OSA	Port	Created By	Last Modified	Last Run	IP Addr	Port	Prot	Eth	MAC	Dev	VLAN
JOHNB	OSA2	OSA24	JOHNB	Wed Sep 30 15:27:18		*	*	*				

14:42:32

OSA Express Network Trace Edit and Launch

Definition name: JOHNB

**Select Port**

Port name: Ethernet Port: OSA408

**Trace Filters**

IP Address: \* More...

IP Port: \* More...

IP Protocol: \* More...

Ethernet Frame Type: \* More...

MAC Address: \* More...

Device ID: \* More...

VLAN Identifier: \* More...

**Trace Limits**

No Trace Limits Set More...

**System Values**

TCP/IP Procname: TCP/IP

CTRACE Procname: CTWRT

**Actions**

Start Trace Save as... Clear Filters Cancel

14:42:35

Assisted Trace configuration, storage, and activation

## Summary ...

- OSAs are massively important today and deserve to be treated as such.
- **ALL** of your network traffic passes through an OSA and that OSA is probably shared, perhaps with non-z/OS partitions.
- Is the true state of the OSA (usage and error counts) really known?
- While it is complicated **collecting** and **consolidating** the available information, it is very worthwhile ...

**Get the complete picture!**

**Thank You**  
**For listening**

[www.willdata.com](http://www.willdata.com)  
The home of "ZEN"