zEnterprise System – z/OS IEDN network design and implementation (part 2)

Gus Kassimis – kassimis@us.ibm.com
Alfred B Christensen – alfredch@us.ibm.com
IBM Raleigh, NC, USA

Session: 8317
Tuesday, March 1, 2011: 11:00 AM-12:00 PM

zEnterprise System – z/OS IEDN network design and implementation (Part 2)

Session number:	8317
Date and time:	Tuesday, March 1, 2011: 11:00 AM-12:00 PM
Location:	Room 212A (Anaheim Convention Center)
Program:	Communications Infrastructure
Project:	Communications Server
Track:	Capitalizing on zEnterprise and Network Support and Management
Classification:	Technical
Speaker:	Gus Kassimis, IBM Alfred B Christensen, IBM
Abstract:	There are many technologies included in the z/OS Communications Server that enable you to design and implement highly available and scalable z/OS Sysplex networking scenarios, and most of you have been using those technologies through many years. How do you extend those capabilities to the zEnterprise environment, or more specifically, to the Intra-Ensemble Data Network (IEDN)? When a z/OS system is part of a z/OS Sysplex and needs to be connected to both a remote network and to the IEDN, some special considerations must be used when designing for high availability and workload management. This session will analyze those considerations and discuss a few selected network topologies to illustrate the design and implementation considerations. The session will cover topologies that include applications deployed on AIX on Power blades and also optimizers in the zBX, such as the IBM Smart Analytics Optimizer.

Trademarks, notices, and disclaimers

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both:

- Advanced Peer-to-Peer Networking®
- AIX®
- alphaWorks®
- AnyNet®
- AS/400®
- BladeCenter®
- Candle®
- CICS®
- DataPower®
- DB2 Connect
- DB2®
- DRDA®
- e-business on demand®
- e-business (logo)
- e business(logo)®
- ESCON®
- FICON®
- GDDM®
- GDPS®
- Geographically Dispersed Parallel Sysplex
- HiperSockets
- HPR Channel Connectivity
- HyperSwap
- i5/OS (logo)
- i5/OS®
- IBM eServer
- IBM (logo)®
- IBM®
- IBM zEnterprise™ System
- IMS
- InfiniBand ®
- IP PrintWay
- IPDS
- iSeries
- LANDP®
- Language Environment®
- MQSeries®
- MVS
- NetView®
- OMEGAMON®
- Open Power
- OpenPower
- Operating System/2®
- Operating System/400®
- OS/2®
- OS/390®
- OS/400®
- Parallel Sysplex®
- POWER®
- POWER7®
- PowerVM
- PR/SM
- pSeries®
- RACF®
- Rational Suite®
- Rational®
- Redbooks
- Redbooks (logo)
- Sysplex Timer®
- System i5
- System p5
- System x®
- System z®
- System z9®
- System z10
- Tivoli (logo)®
- Tivoli®
- VTAM®
- WebSphere®
- xSeries®
- z9®
- z10 BC
- z10 EC
- zEnterprise
- zSeries®
- z/Architecture
- z/OS®
- z/VM®
- z/VSE

* All other products may be trademarks or registered trademarks of their respective companies.

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both:

- Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
- IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

Notes:

- Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
- IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
- This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
- All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
- Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Refer to www.ibm.com/legal/us for further legal information.

Agenda

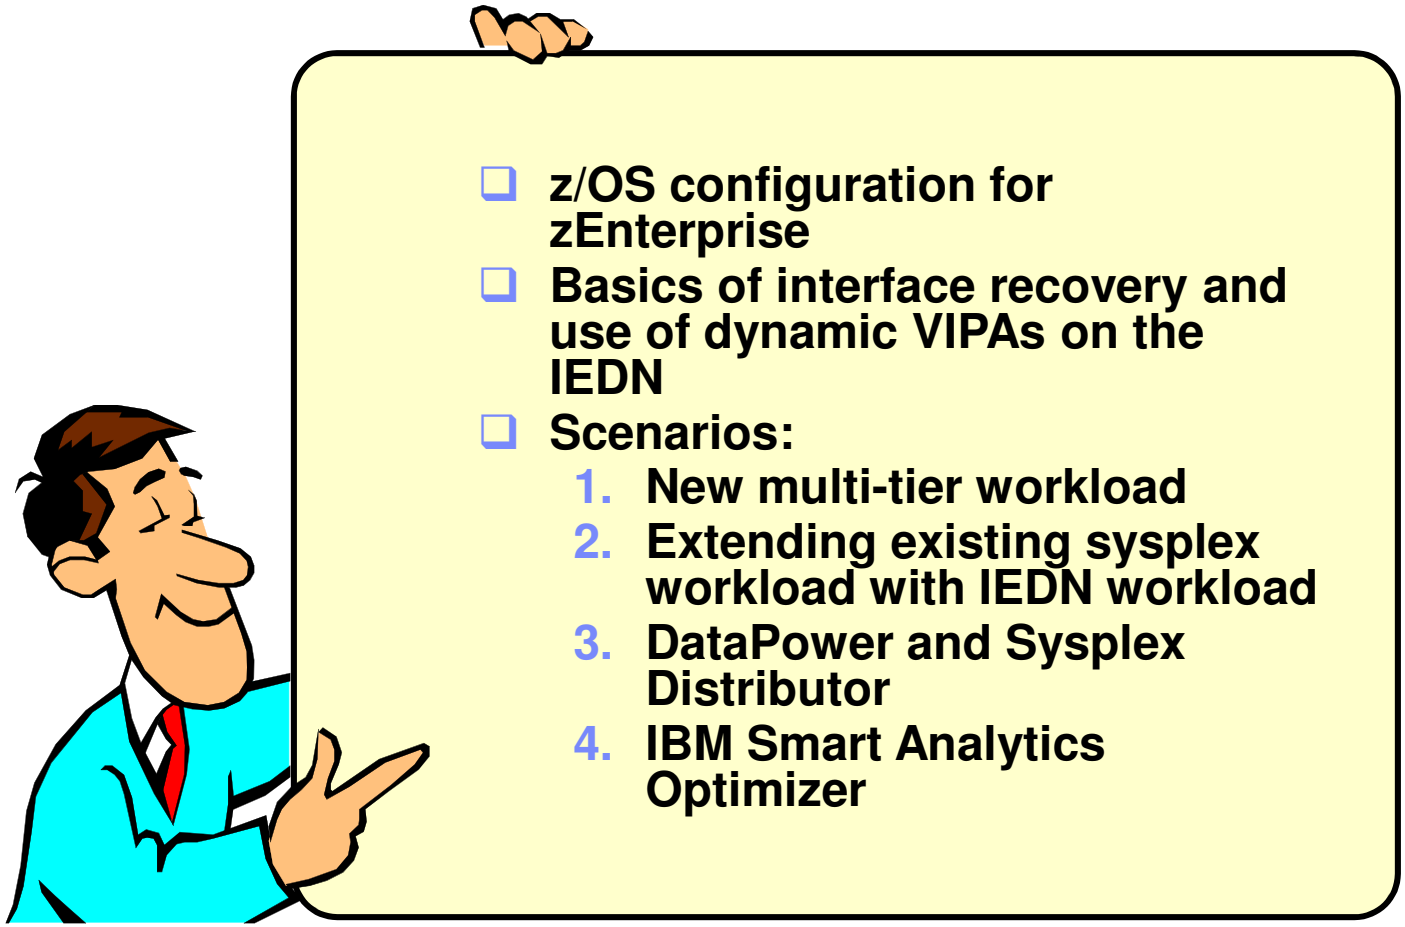

- z/OS configuration for zEnterprise
- Basics of interface recovery and use of dynamic VIPAs on the IEDN
- Scenarios:
  1. New multi-tier workload
  2. Extending existing sysplex workload with IEDN workload
  3. DataPower and Sysplex Distributor
  4. IBM Smart Analytics Optimizer

Disclaimer: All statements regarding IBM future direction or intent, including current product plans, are subject to change or withdrawal without notice and represent goals and objectives only. All information is provided for informational purposes only, on an "as is" basis, without warranty of any kind.

zEnterprise networking – z/OS IEDN network design and implementation (Part 2)

z/OS configuration for zEnterprise

The zEnterprise Management Network 1000Base-T OSA configured as CHPID type OSM

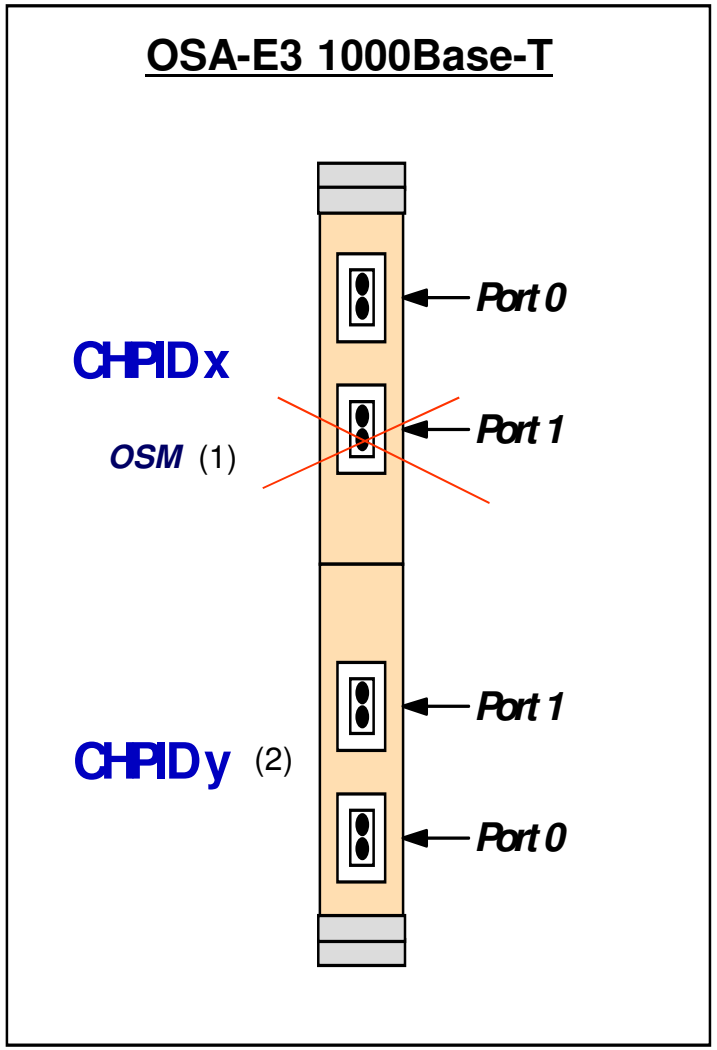

IOCDS

```
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 Connections are dynamically created when an Ensemble is defined.

z/OS Communications Server and OSM connectivity

- z/OS LPAR must participate in the ensemble
  - –New VTAM start option: ENSEMBLE=YES
    - Required for both OSM and OSX connectivity

```
ENSEMBLE=YES,                LPAR is part of zEnterprise ENSEMBLE
```

- LPAR must be IPv6-enabled for OSM connectivity
  - –BPXPRMxx PARMLIB updates
- Two IPAQENET6 interface definitions are dynamically generated and started
  - –If OSM CHPIDs are defined to the z/OS image, the two CHPIDs with the lowest device numbers are assigned to these interfaces
- TRLEs dynamically generated if connectivity allowed and CHPIDs found
  - –Only port 0 supported
- IPv6 link-local address only
  - –No globally unique IPv6 addresses needed

z/OS Communications Server and OSM connectivity (cont)

- Uses VLAN in access mode
  - Switch handles VLAN tagging, stack unaware
- Not reported to OMPROUTE
- Cannot add static or dynamic routes
- Supports stop, start, packet trace, and OSA NTA
- Only applications permitted to EZB.OSM.sysname.tcpname can communicate over OSM interfaces
  - The Guest Platform Management Provider (GPMP) is the only application that needs to

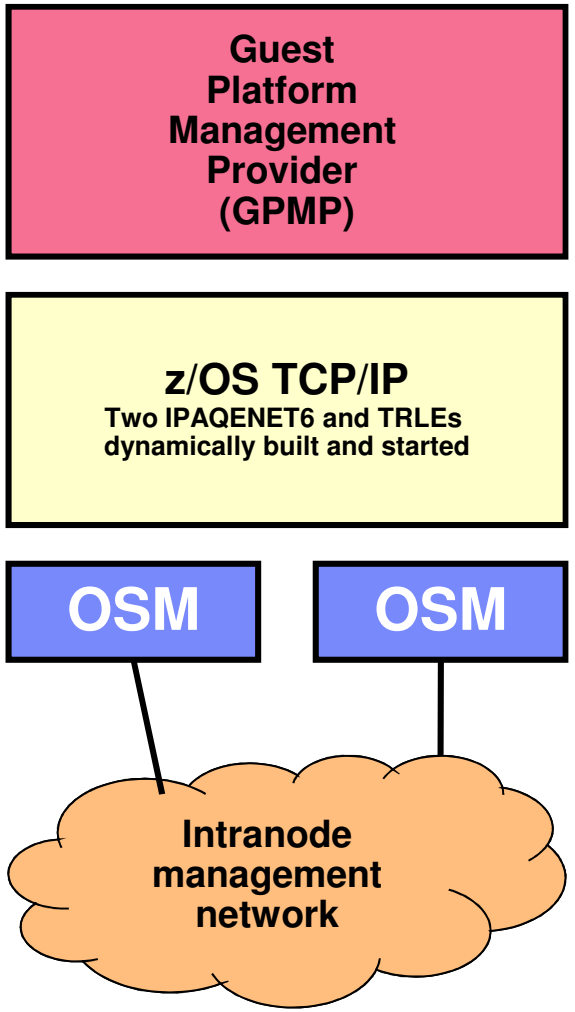

Steps to enable the intranode management network

1. Authorize the management application to the EZB.OSM.sysname.tcpname resource.
   - To send or receive data over an OSM interface, an application must have READ authorization to the EZB.OSM.sysname.tcpname resource. If used on this image, authorize the application to this resource.
2. Reserve the UDP port that the platform management application is to use to listen for multicast traffic over the intranode management network.
3. Authorize any user IDs to this resource that might issue diagnostic commands, such as Ping and Traceroute, over OSM interfaces to verify connectivity.
4. If you enable IP security for IPv6, you can configure a security class for IP filtering that applies to all OSM interfaces.
   - Use the OSMSECCLASS parameter on the IPCONFIG6 statement. This enables you to configure filter rules for traffic over the EZ6OSM01 and EZ6OSM02 interfaces.
5. If the multicast address that is used by the platform management application is configured into a network access zone, then give the user ID for this application read permission to the resource profile for that zone.

How bad is it enabling IPv6?

- Add a NETWORK stmt. for AF_INET6 to your BPXPRMxx PARMLIB member
- No changes needed to your TCP/IP Profile
  - Unless you want to exploit and test specific IPv6 features
- Testing needed in the network management area
  - All Netstat reports will use the LONG format
    - Also when reporting on IPv4-only activity
  - Home-written Netstat "scraping" logic will need to be changed
  - Network management products may fail if they are not prepared for IPv6 addresses

```
FILESYSTYPE TYPE(INET) ENTRYPOINT(EZBPFINI)
 NETWORK DOMAINNAME(AF_INET)
         DOMAINNUMBER(2)
         MAXSOCKETS(2000)
         TYPE(INET)
 NETWORK DOMAINNAME(AF_INET6)
         DOMAINNUMBER(19)
         MAXSOCKETS(3000)
         TYPE(INET)
```

AF_INET

AF_INET6

```
MVS TCP/IP NETSTAT CS V1R12        TCPIP Name: TCPCS           13:02:52
User Id  Conn     State
-------  ----     -----
MYINETD1 00000022 Listen
  Local Socket:   9.42.130.98..23
  Foreign Socket: 0.0.0.0..0
TN3270A  0000004D Establsh
  Local Socket:   ::ffff:9.42.105.45..23
  Foreign Socket: ::ffff:9.76.144.213..4211
  Application Data:  EZBTNSRV TCPABC80 TSO10001 ET B
TN3270A  0000003F Listen
  Local Socket:   ::..23
  Foreign Socket: ::..0
  Application Data:  EZBTNSRV LISTENER
```

AF_INET socket

AF_INET6 sockets

It's a piece of cake

Netstat devlinks example of an OSM interface

```
IntfName: EZ6OSM01           IntfType: IPAQENET6  IntfStatus: Ready
    PortName: IUTMP0F2  Datapath: BE22     DatapathStatus: Ready
    CHPIDType: OSM
    QueSize: 0    Speed: 0000001000
    VMACAddr: 02006FEB8363   VMACOrigin: OSA    VMACRouter: All
    DupAddrDet: 1
    CfgMtu: None                       ActMtu: 1500
    VLANid: None                       VLANpriority: Disabled
    ReadStorage: GLOBAL (4096K)
    InbPerf: Dynamic
      WorkloadQueueing: No
    ChecksumOffload: No                SegmentationOffload: No
    SecClass: 255                      MonSysplex: No
    Isolate: Yes                       OptLatencyMode: No
    TempPrefix: None
 Multicast Specific:
    Multicast Capability: Yes
    Group:     ff02::1:ffeb:8363
      RefCnt:  0000000001  SrcFltMd: Exclude
      SrcAddr: None
(more multicast group memeberships)
 Interface Statistics:
    BytesIn                            = 4256
    Inbound Packets                    = 29
    Inbound Packets In Error           = 9
    Inbound Packets Discarded          = 0
    Inbound Packets With No Protocol   = 0
    BytesOut                           = 958
    Outbound Packets                   = 9
    Outbound Packets In Error          = 0
    Outbound Packets Discarded         = 0
```

CHPIDType OSM

VMAC

INBPERF is Dynamic

The zEnterprise Internal Data Network 10 Gigabit OSA configured as CHPID type OSX (z/OS)

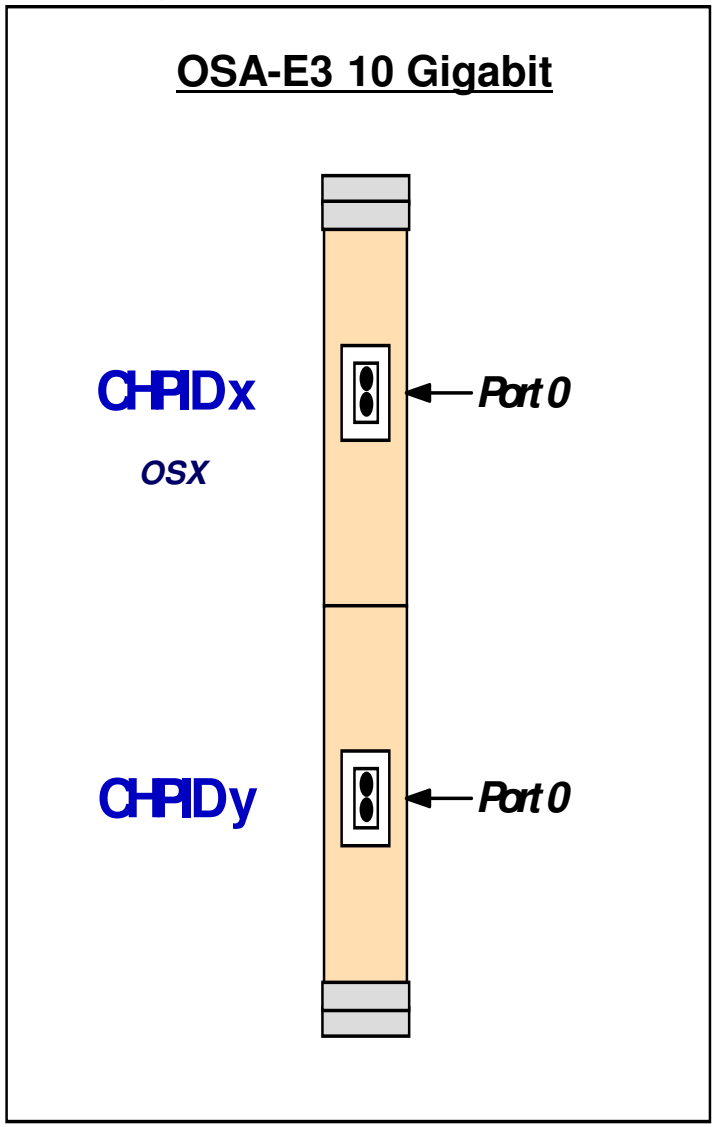

IOCDS

```
CHPID PCHID=5E1,PATH=(CSS(0,1,2,3),2F),          *
TYPE=OSX,SHARED,   …

CNTLUNIT CUNUMBR=09F0,PATH=((CSS(0),2F)), *
UNIT=OSX

IODEVICE ADDRESS=(09F0,15),CUNUMBR=(091F),*
UNIT=OSA,UNITADD=00,                              *
MODEL=X,DYNAMIC=YES,LOCANY=YES
```

Operating System Definitions (z/OS)

VTAM Definitions

- Dynamic TRLEs
- Manually Defined TRLEs

TCP/IP Definitions

- INTERFACE IPAQENET
- INTERFACE IPAQENET6

z/OS Communications Server and OSX connectivity

- Configure with INTERFACE statement
  - IPAQENET and IPAQENET6
- Either specify CHPID
  - Dynamically created TRLE similar to HiperSockets
- Or configure TRLE and point to it
  - Useful in VM guest LAN environment where CHPID is unpredictable
- Always uses VLAN in trunk mode
  - VLANID required and must be authorized at HMC
    - If not authorized, OSA activation fails
- Prevents IP forwarding from OSX ⇨ OSX
  - Sysplex distributor forwarding is allowed when using VIPAROUTE

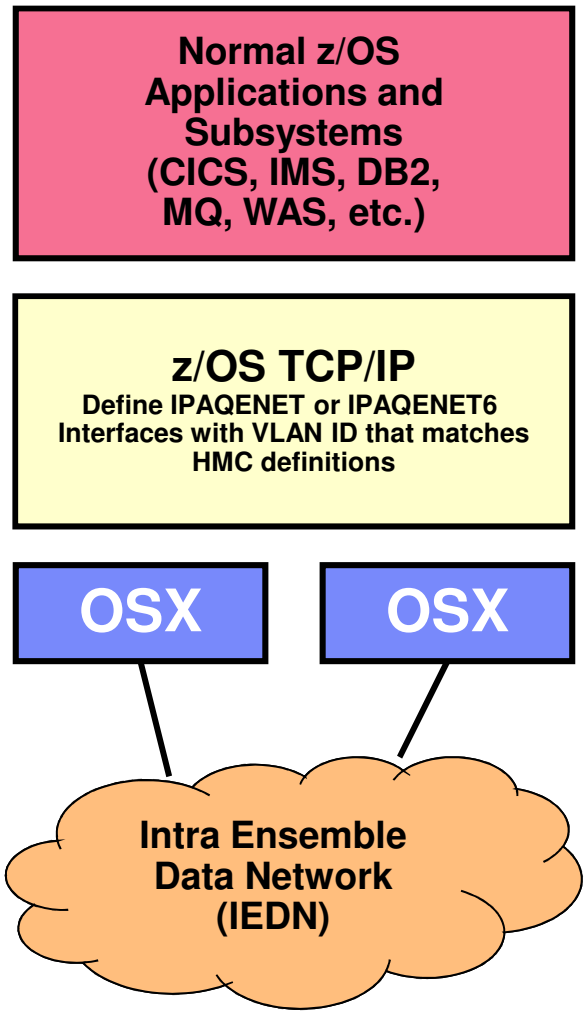

z/OS Communications Server and OSX connectivity

- Supports stop, start, packet trace, and OSA NTA
- To prevent external traffic from being routed to/from the OSX VLAN
  - Define OSX as INTERFACE or IPV6_INTERFACE
  - Do not enable IMPORT_DIRECT_ROUTES function
  - Alternatively, do not define the OSX interfaces to OMPROUTE and tell OMPROUTE to ignore undefined interfaces
- To allow external traffic to be routed to/from the OSX VLAN
  - Define OSX as OSPF_INTERFACE or IPV6_ISPF_INTERFACE
  - Configure a non-0 value for ROUTER_PRIORITY
- If you are already using OMPROUTE and OSPF, there seem to be no general reason not to define the OSX interfaces as normal OSPF interfaces
  - Routing to/from OSX interfaces can still be controlled via IPSec filter rules

IEDN OSX IOCP and TCP/IP sample definitions

IOCP

```
CHPID PATH=(CSS(0,1),D3),SHARED,                                    *
      PCHID=131,TYPE=OSX
CNTLUNIT CUNUMBR=BF30,PATH=((CSS(0),D3),(CSS(1),D3)),UNIT=OSA
IODEVICE  ADDRESS=(BF30,14),CUNUMBR=BF30,UNIT=OSA,                  *
      NOTPART=((CSS(1),RALHCD)),                                    *
      UNITADD=00
IODEVICE  ADDRESS=(BF3E,1),CUNUMBR=BF30,UNIT=OSAD,                  *
      NOTPART=((CSS(1),RALHCD)),                                    *
      UNITADD=FE
```

CHPIDTYPE OSX

CHPID D3

VLANID 85

TCP/IP

```
INTERFACE O3OSXA0 DEFINE IPAQENET
 CHPIDTYPE OSX
 CHPID D3
 IPADDR 16.11.160.108/21
 SOURCEVIPAINT LFRVIPA1
 MTU 8992
 VLANID 85
 READSTORAGE GLOBAL
 INBPERF DYNAMIC WORKLOADQ
 NOISOLATE
 IPBCAST
 MONSYSPLEX
 DYNVLANREG
 NOOLM
 VMAC
```

Virtual MAC address

OMPROUTE

```
OSPF_Interface
  Name                      = O3OSXA0
  IP_Address                = 16.11.160.108
  Subnet_Mask               = 255.255.248.0
  MTU                       = 9000
  Hello_Interval            = 10
  Dead_Router_Interval      = 40
  Retransmission_Interval   = 20
  Cost0                     = 100
  Attaches_To_Area          = 11.11.11.11
```

Netstat devlinks example of an OSX interface

```
IntfName: O3OSXA0           IntfType: IPAQENET   IntfStatus: Ready
    PortName: IUTXP0D3  Datapath: BF32     DatapathStatus: Ready
    CHPIDType: OSX      CHPID: D3
    Speed: 0000010000
    IpBroadcastCapability: Yes
    VMACAddr: 029BA400024F   VMACOrigin: OSA    VMACRouter: All
    SrcVipaIntf: LFRVIPA1
    ArpOffload: Yes                  ArpOffloadInfo: Yes
    CfgMtu: 8992                     ActMtu: 8992
    IpAddr: 16.11.160.108/21
    VLANid: 85                       VLANpriority: Enabled
    DynVLANRegCfg: Yes               DynVLANRegCap: Yes
    ReadStorage: GLOBAL (4096K)
    InbPerf: Dynamic
      WorkloadQueueing: Yes
    ChecksumOffload: Yes             SegmentationOffload: Yes
    SecClass: 255                    MonSysplex: Yes
    Isolate: No                      OptLatencyMode: No
 Multicast Specific:
    Multicast Capability: Yes
    Group             RefCnt         SrcFltMd
    -----             ------         --------
    224.0.0.6         0000000001     Exclude
      SrcAddr: None
    224.0.0.5         0000000001     Exclude
      SrcAddr: None
    224.0.0.1         0000000001     Exclude
      SrcAddr: None
  Interface Statistics:
    BytesIn                          = 32182373
    Inbound Packets                  = 194050
    Inbound Packets In Error         = 0
    Inbound Packets Discarded        = 0
    Inbound Packets With No Protocol = 0
    BytesOut                         = 21913806
    Outbound Packets                 = 183213
    Outbound Packets In Error        = 0
    Outbound Packets Discarded       = 0
```

CHPIDType OSX

VMAC

VLANID 85

INBPERF Dynamic

Basics of interface recovery and use of dynamic VIPAs on the IEDN

Some basic LAN technology overview

- The LAN infrastructure transports “Frames” between Network Interface Cards (NICs) that are attached to the LAN media (Copper or fiber optic)
- Each NIC has a physical hardware address
  - A Media Access Control (MAC) address
    - Burned in (world-wide unique by vendors) or alternatively locally administered
- Every frame comes from a MAC and goes to a MAC
  - There are special MAC values for broadcast and multicast frames
- Every frame belongs to the physical LAN or to one of multiple Virtual LANs (VLAN) on the physical LAN
  - A VLAN ID is in the IEEE801.Q header if VLAN technologies are in use
- A frame carries a payload of a specified protocol type, such as ARP, IPv4, IPv6, SNA LLC2, etc.

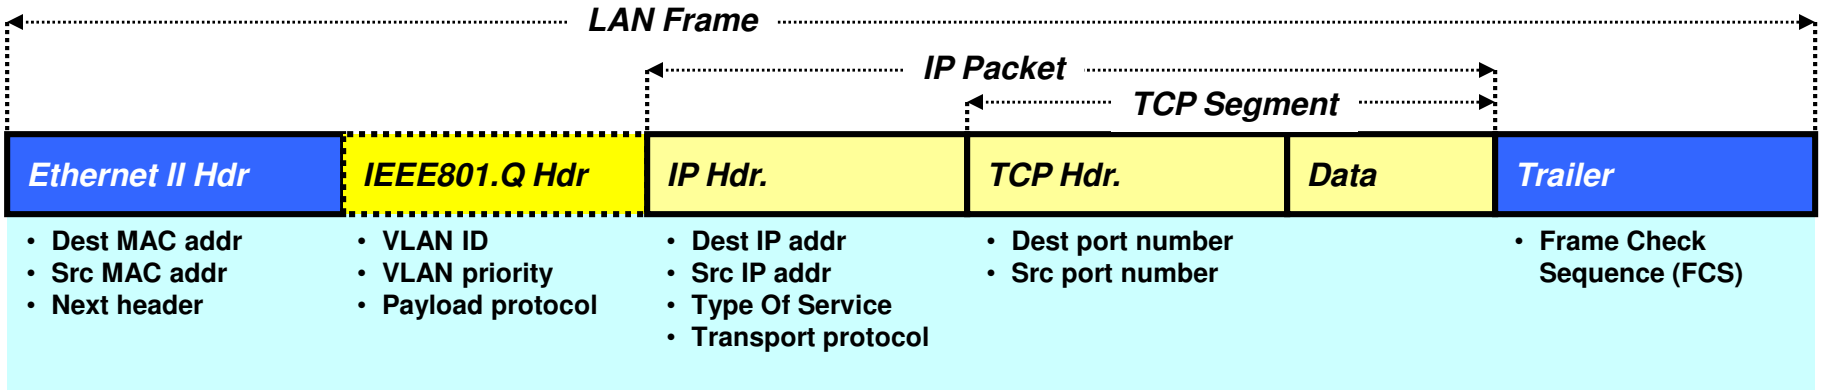

Correlation of IPv4 addresses and MAC addresses on a LAN – Address Resolution Protocol (ARP)

- An IPv4 node uses the ARP protocol to discover the MAC address of another IPv4 address that belongs to the same IPv4 subnet as it does itself.
- ARP requests are broadcasted to all NICs on the LAN
- The one NIC that has a TCP/IP stack with the requested IPv4 address responds directly back to the IPv4 node that sent out the broadcast
- Each IPv4 node maintains a cache of IPv4 addresses and associated MAC addresses on their directly connected LANs

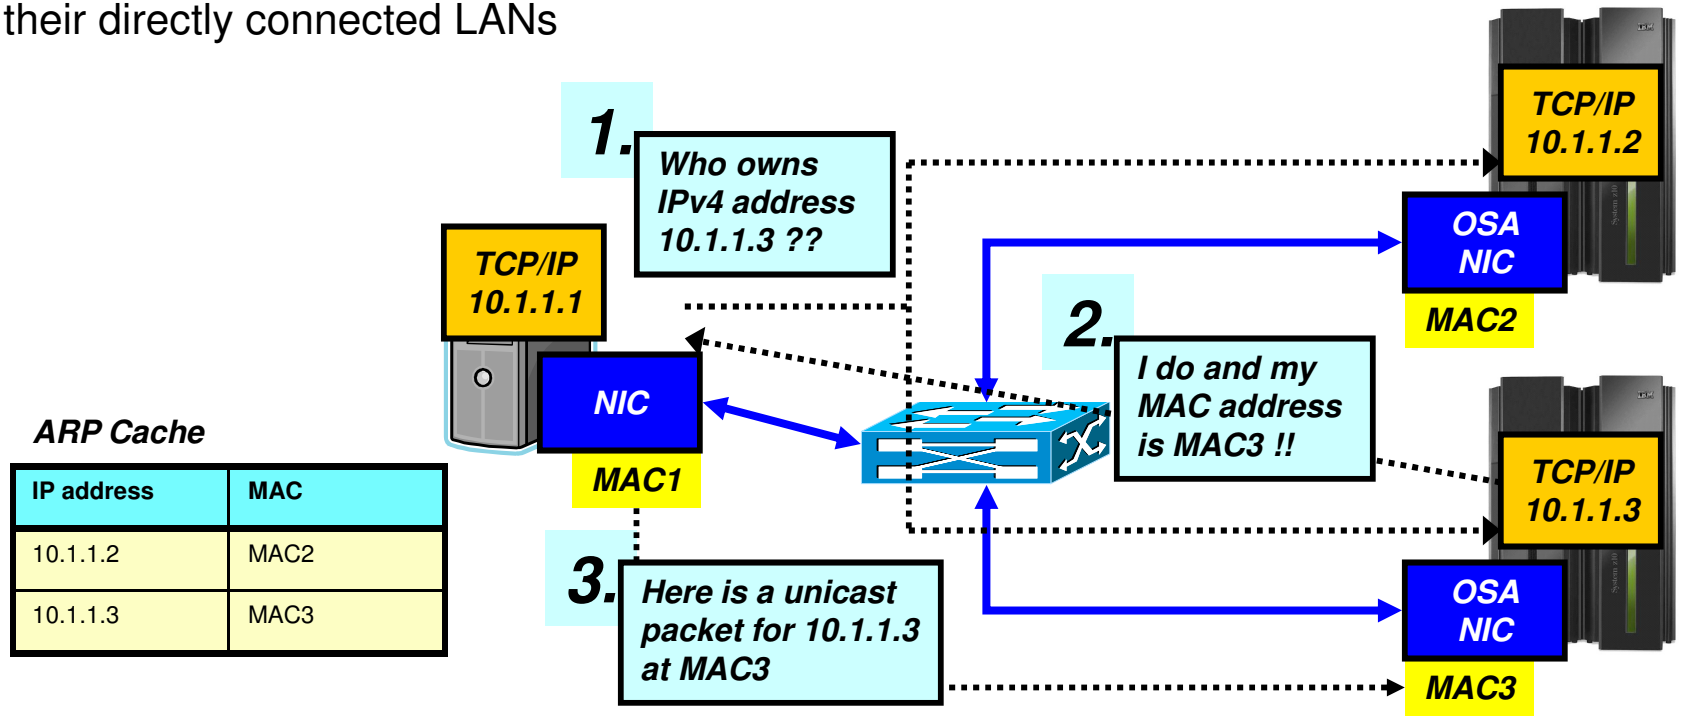

z/OS VIPA address visibility on the IEDN

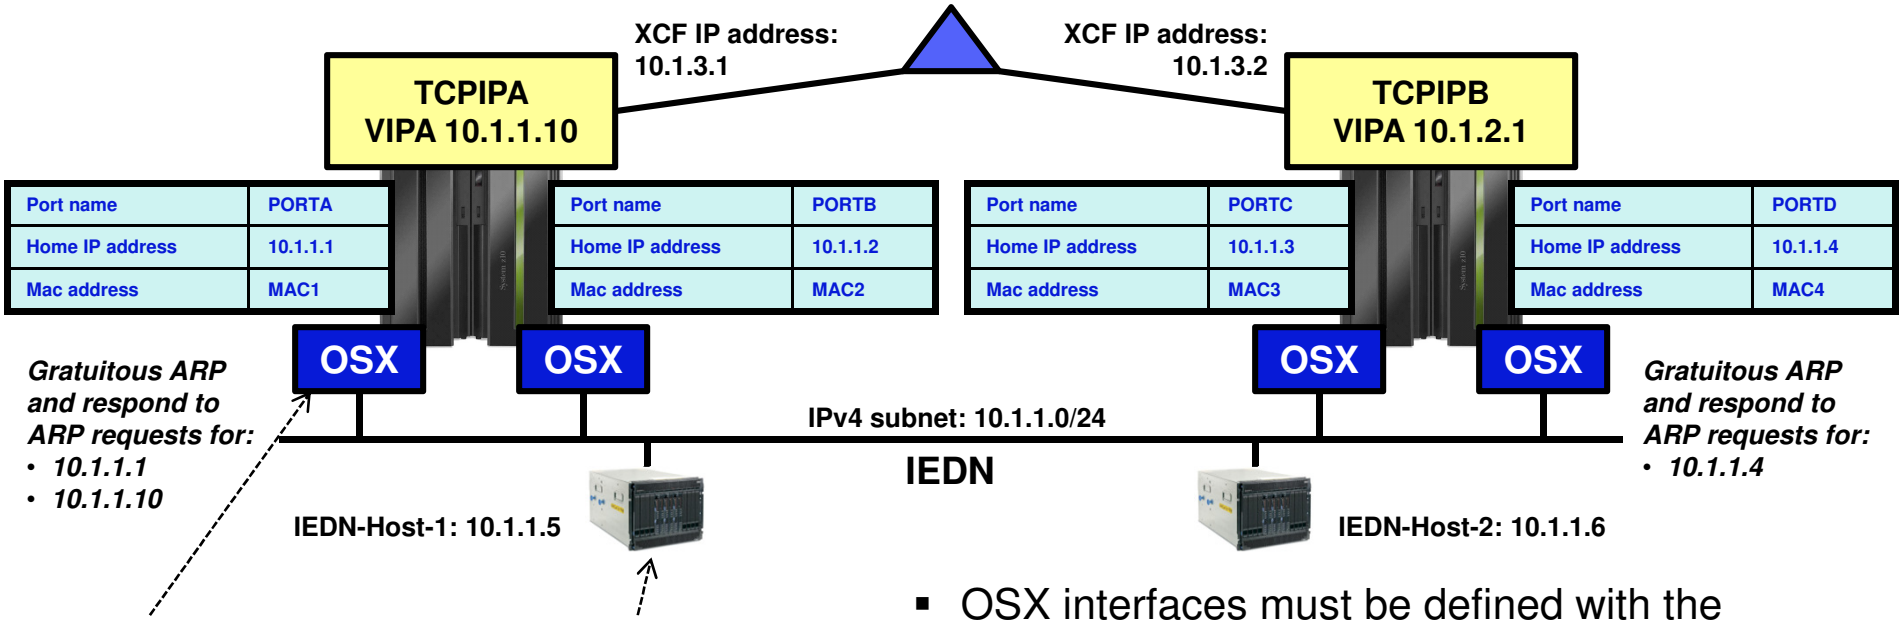

OSA PORTA's OAT

IP Address	ARP Owner
10.1.1.1	Yes
10.1.1.10	Yes
10.1.1.2	No
10.1.3.1	No

IEDN-Host-1's ARP cache

IP Address	MAC Address
10.1.1.1	MAC1
10.1.1.2	MAC2
10.1.1.3	MAC3
10.1.1.4	MAC4
10.1.1.10	MAC1

- OSX interfaces must be defined with the INTERFACE statement
- With VMAC and ROUTEALL, only addresses for which OSA has to perform ARP are registered in the OAT
- In all other cases, all HOME IP addresses will be registered in the OAT and the OAT content will be changed as the HOME lists change due to (dynamic) movement of IP addresses.
- OSX interfaces will do gratuitous ARP for the OSA interface IP address and for VIPA addresses that belong to the ***same*** subnet as the OSA interface.

Network connectivity resilience on the IEDN

z/OS TCP/IP supports interface recovery if multiple network interfaces to the same subnet exist.
In this example, both OSA PORTA and PORTB are connected to the IEDN (10.1.1.0/24 subnet).

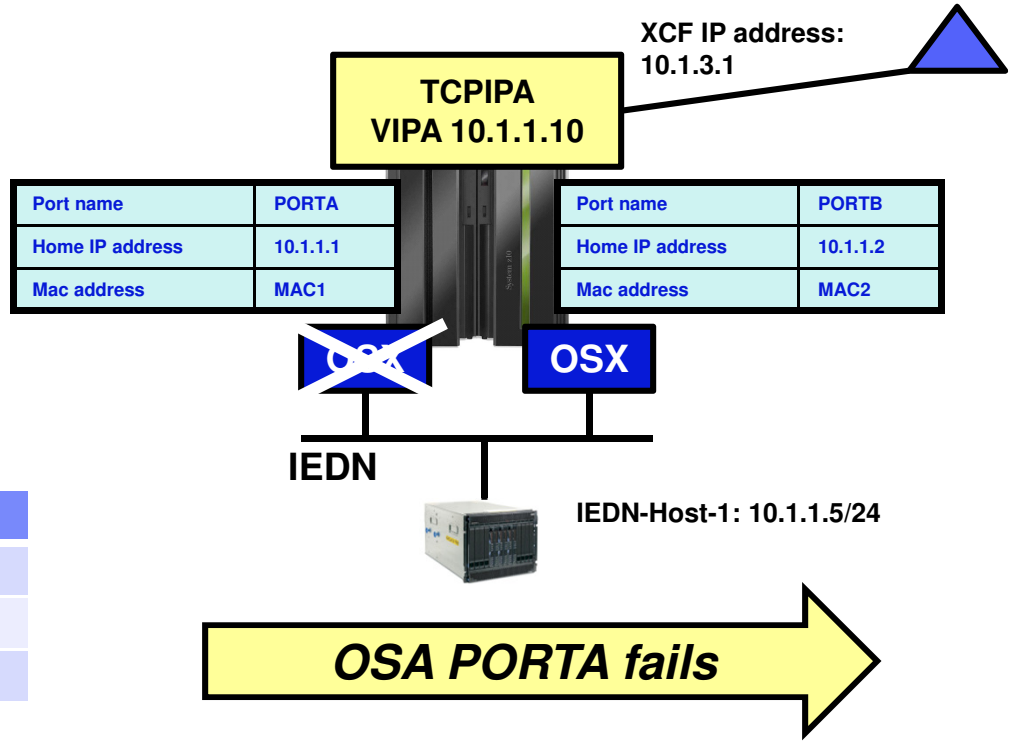

When PORTA fails, PORTB is given ARP ownership of the addresses PORTA previously had. PORTB sends gratuitous ARPs to enable downstream hosts to update their ARP cache.

IEDN-Host-1's ARP cache

IP Address	MAC Address
10.1.1.1	MAC1
10.1.1.2	MAC2
10.1.1.10	MAC1

IEDN-Host-1's ARP cache

IP Address	MAC Address
10.1.1.1	MAC2
10.1.1.2	MAC2
10.1.1.10	MAC2

OSA PORTA's OAT

IP Address	ARP Owner
10.1.1.1	Yes
10.1.1.10	Yes
10.1.1.2	No
10.1.3.1	No

OSA PORTB's OAT

IP Address	ARP Owner
10.1.1.1	No
10.1.1.10	No
10.1.1.2	Yes
10.1.3.1	No

OSA PORTA's OAT

IP Address	ARP Owner
10.1.1.1	Yes
10.1.1.10	Yes
10.1.1.2	No
10.1.3.1	No

OSA PORTB's OAT

IP Address	ARP Owner
10.1.1.1	Yes
10.1.1.10	Yes
10.1.1.2	Yes
10.1.3.1	No

How do you know which OSA interfaces are on the same subnet and which OSA interface currently handles ARP for your VIPA addresses?

- Messages are issued when an interface takes over ARP responsibility
  - EZD0040I INTERFACE OSAQDIO2 HAS TAKEN OVER ARP RESPONSIBILITY FOR INACTIVE INTERFACE OSAQDIO1
- Messages are issued whenever a previously taken over link or interface recovers and takes back the ARP responsibility.
  - EZD0041I INTERFACE OSAQDIO1 HAS TAKEN BACK ARP RESPONSIBILITY FROM INTERFACE OSAQDIO2
- Use Netstat DEvlinks/-d report to tracks the state of takeover:
  - Displays ARP/ND information.
  - LAN group membership is determined dynamically per interface during interface initialization
  - LAN group numbers are determined dynamically, they are not configured

```
.........
IPv4 LAN Group Summary
 LanGroup: 00010
  Name               Status      ArpOwner          VipaOwner
  ----               ------      --------          ---------
  O3OSXA0            Active      O3OSXA0           Yes
IPv6 LAN Group Summary
LanGroup: 00001
  Name               Status      NDOwner           VipaOwner
  ----               ------      -------           ---------
  EZ6OSM02           Active      EZ6OSM02          Yes
  EZ6OSM01           Active      EZ6OSM01          No
```

Dynamic VIPA movement on the IEDN (without dynamic routing)

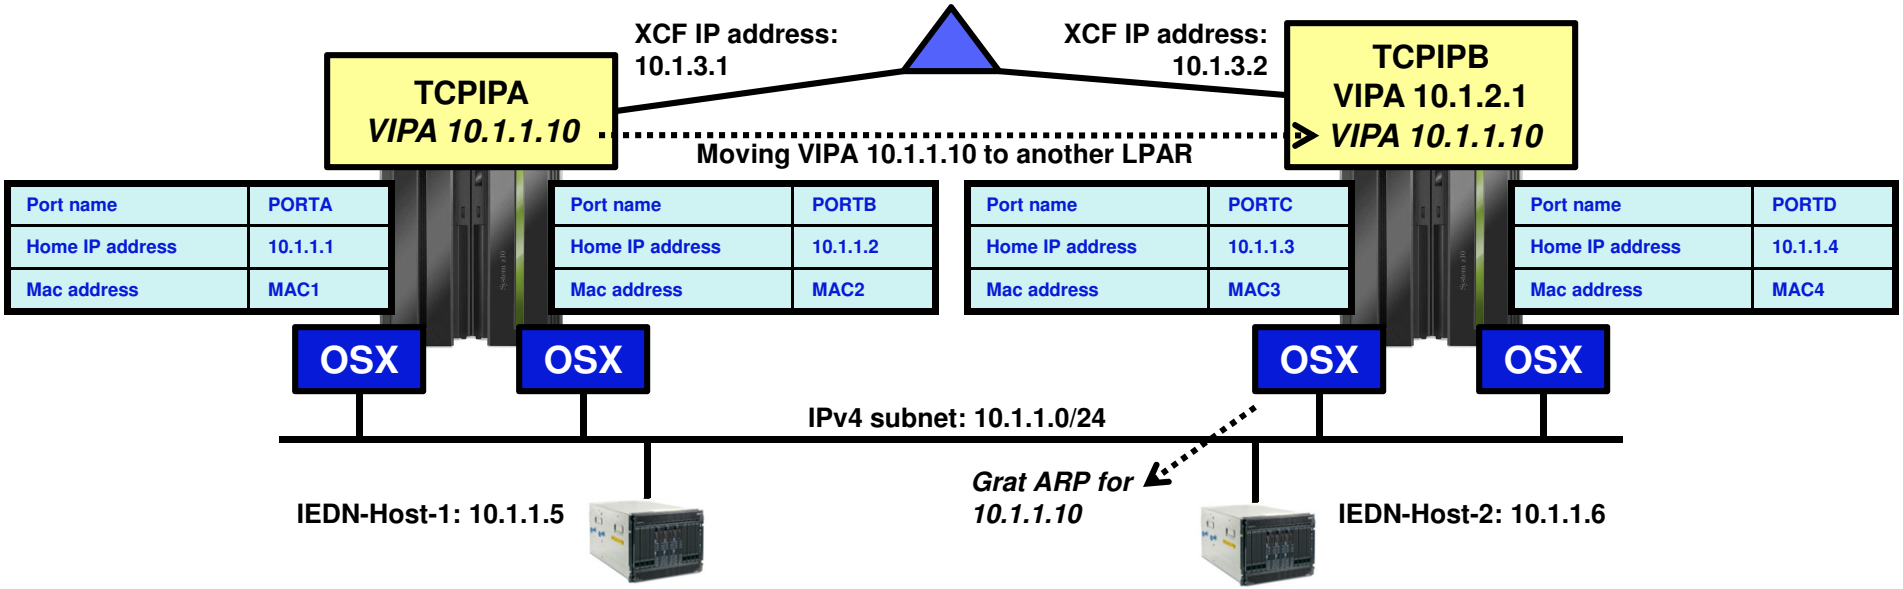

When 10.1.1.10 is added to OSA PORTC's OAT with ARP ownership, PORTC sends out a gratuitous ARP that forces nodes on the LAN with that IP address in their ARP cache to update their ARP cache entry with the new MAC address.

OSA PORTA's OAT

IP Address	ARP Owner
10.1.1.1	Yes
10.1.1.2	No
10.1.3.1	No

IEDN-Host-1's ARP cache

IP Address	MAC Address
10.1.1.1	MAC1
10.1.1.2	MAC2
10.1.1.10	***MAC1***
10.1.1.3	MAC3
10.1.1.4	MAC4

OSA PORTC's OAT

IP Address	ARP Owner
10.1.1.3	Yes
10.1.1.10	***Yes***
10.1.1.4	No
10.1.3.2	No
10.1.2.1	No

IEDN-Host-1's ARP cache

IP Address	MAC Address
10.1.1.1	MAC1
10.1.1.2	MAC2
10.1.1.10	***MAC3***
10.1.1.3	MAC3
10.1.1.4	MAC4

Some of the rules for availability and DVIPA movement in flat networks without dynamic routing – such as on the IEDN

- z/OS VIPA addresses in a flat network configuration without dynamic routing must be allocated out of the same subnet as the directly attached network - in this example, the 10.1.1.0/24 subnet.
  - If multiple VLANs are used on the IEDN, VIPA addresses belonging to one VLAN (one subnet) can in general only be accessed and recovered via that VLAN
  - A default router that does know about all DVIPAs, can be used to reach VIPAs on other VLANs (subnets)
- All LPARS in the Sysplex must be attached to one and the same IP subnet via OSA ports.
- Network interfaces belonging to other IP subnets cannot be used for automatic re-routing around failed OSA ports.
  - That includes MPC links, XCF links, or other OSA-attached subnets
- Overall physical availability of the network to which the OSA ports are attached becomes of outmost importance and must generally be based on what is known as Layer2-and-below availability functions in the switches and the physical links (cables).
  - Redundant switches with trunk links – part of the zEnterprise infrastructure
  - Redundant OSA adapters in each LPAR – always use at least two OSX ports
  - OSA port cabling to multiple switches – part of the zEnterprise infrastructure

Scenario 1:
New multi-tier workload

Single IEDN VLAN, new zEnterprise application workload, z/OS as the gateway to the external network

Hosts/routers on the external network can access the DVIPA addresses as long as they listen to dynamic route updates (as normal)

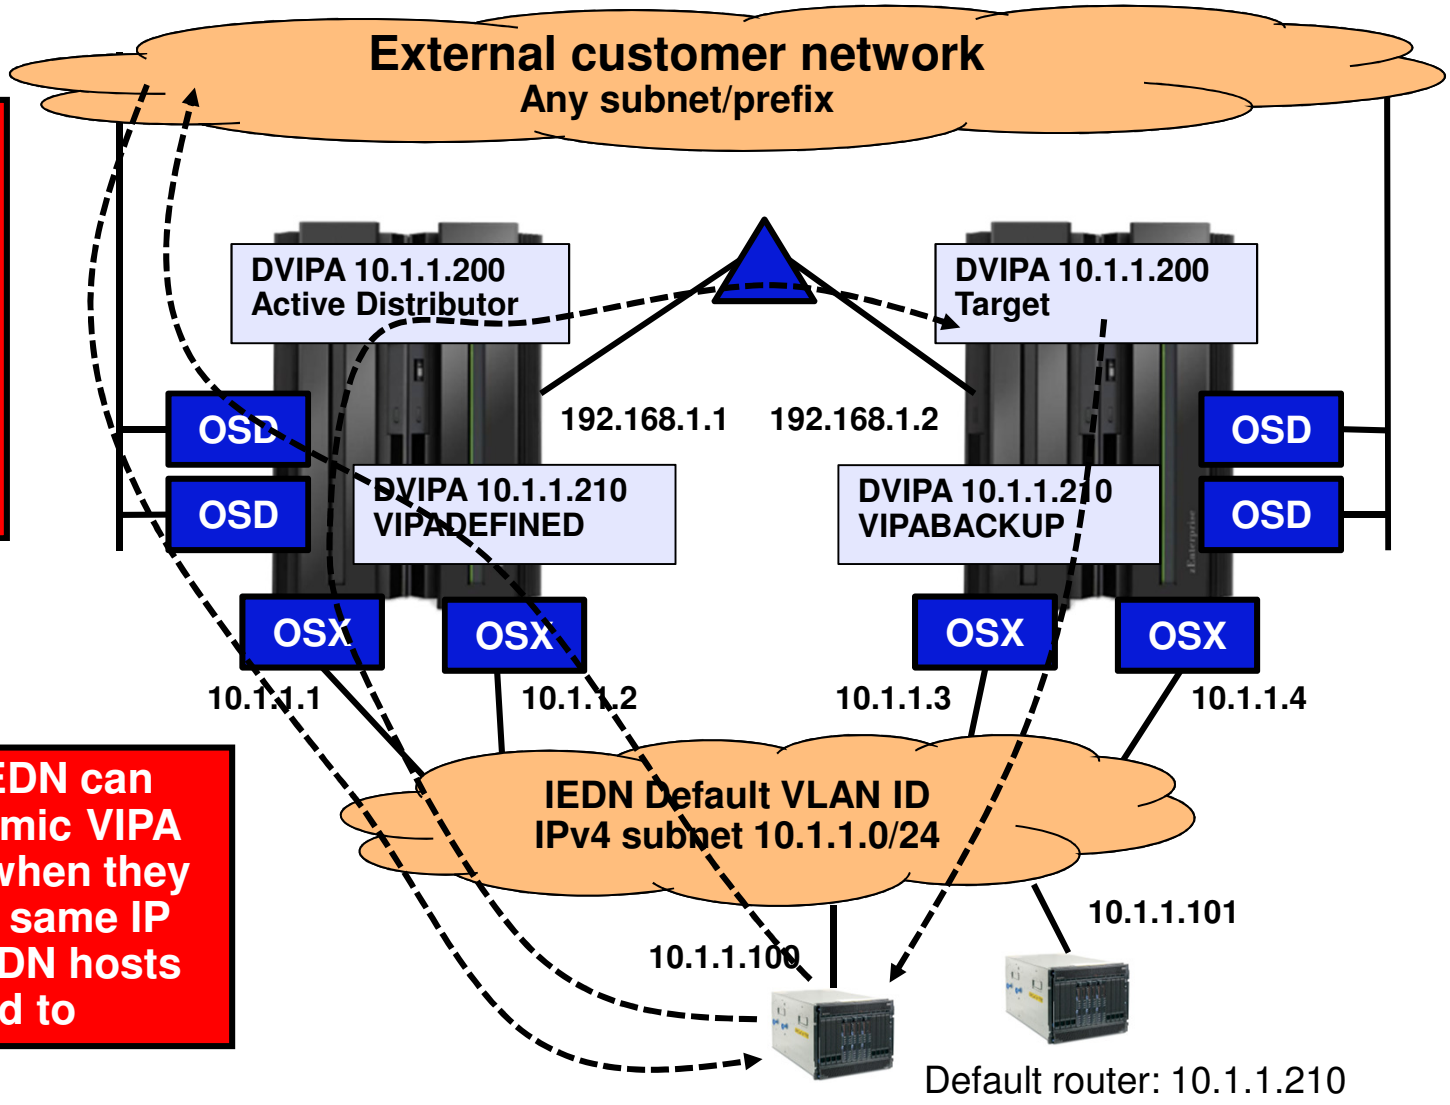

Hosts on the IEDN can access the dynamic VIPA addresses only when they come out of the same IP subnet as the IEDN hosts are attached to

New IEDN workload considerations

- Assume in this scenario, a single VLAN on the IEDN
- Assume also, the IPv4 subnet assigned to that IEDN VLAN is 10.1.1.0/24
- As long as the z/OS DVIPA addresses are assigned out of the same IPv4 subnet, the IEDN hosts can access them
  - OSX interface recovery and DVIPA movement addressed via ARP updates – as discussed earlier
- If the IEDN hosts need a route to the external network, a VIPADEFINED DVIPA can be used as the default router from the IEDN hosts
  - The z/OS default router responsibility will move to backup a z/OS if primary z/OS fails – retaining access to he external network for the IEDN hosts
- z/OS should be defined with MULTIPATH PerConnection to load-balance outbound traffic from z/OS

OMPROUTE and the IEDN

- If you want to prevent external traffic from being routed to the IEDN VLAN, then do one of the following so that OMPROUTE does not advertise the intraensemble subnet:
  - Define the OSX interface to OMPROUTE using an INTERFACE statement or IPV6_INTERFACE statement, and do not enable the IMPORT_DIRECT_ROUTES function of AS boundary routing.
  - Do not define the OSX interface to OMPROUTE, and ensure that GLOBAL_OPTIONS IGNORE_UNDEFINED_INTERFACES is configured to OMPROUTE..
- If you want to allow external traffic to be routed to the IEDN VLAN, then define the OSX interface to OMPROUTE as an OSPF_INTERFACE or IPV6_OSPF_INTERFACE, and code a nonzero value for the ROUTER_PRIORITY parameter on the interface.
  - As long as no other hosts on that OSX VLAN have coded their interfaces as OSPF interfaces, then OMPROUTE advertises the subnet (or IPv6 prefixes) of the intraensemble data network into the OSPF network.
  - This advertisement makes all addresses that fall into the intraensemble subnet (or IPv6 prefixes) reachable using OSPF.
- **Tip**: These definitions apply per interface, so you could implement advertising on one VLAN while not advertising on a different VLAN attached to the same z/OS router.
- **Remember**: z/OS will never route traffic between IEDN VLANs

Scenario 2: Extending existing Sysplex workload with IEDN workload

Considerations for Extending existing Sysplex workload with IEDN workload

- DVIPA addresses that must be accessed from an IEDN VLAN, must be assigned out of the same subnet as the IEDN VLAN itself.
- If you are adding IEDN access to an existing z/OS Sysplex that already use VIPA addresses for application access and you need to access those same applications from the IEDN, you have a few choices:
  1. Renumber the affected DVIPAs to use IEDN VLAN addresses
     - Everything looks like it did in the previous scenario
     - Update your name servers with the new addresses
     - Hope (!) that everyone in the external network does use a name server when accessing those applications
  2. If the applications use generic IP address binding (0.0.0.0), you can assign new additional VIPA addresses to be used when accessing the applications from the IEDN
     - The TN3270 server may be accessed through both an existing VIPA address and a new VIPA address
  3. If the application uses a bind-specific DVIPA address, and the application supports being started on multiple addresses, define a new “instance” of the listener running on that new DVIPA address
     - This will work for many bind-specific applications, such as CICS Sockets Domain, CICS Sockets, IMS Connect, FTP servers, TN3270 servers, etc.
- Always remember to check your networking policy rules when adding new IP addresses.

Adding IEDN access to existing z/OS servers

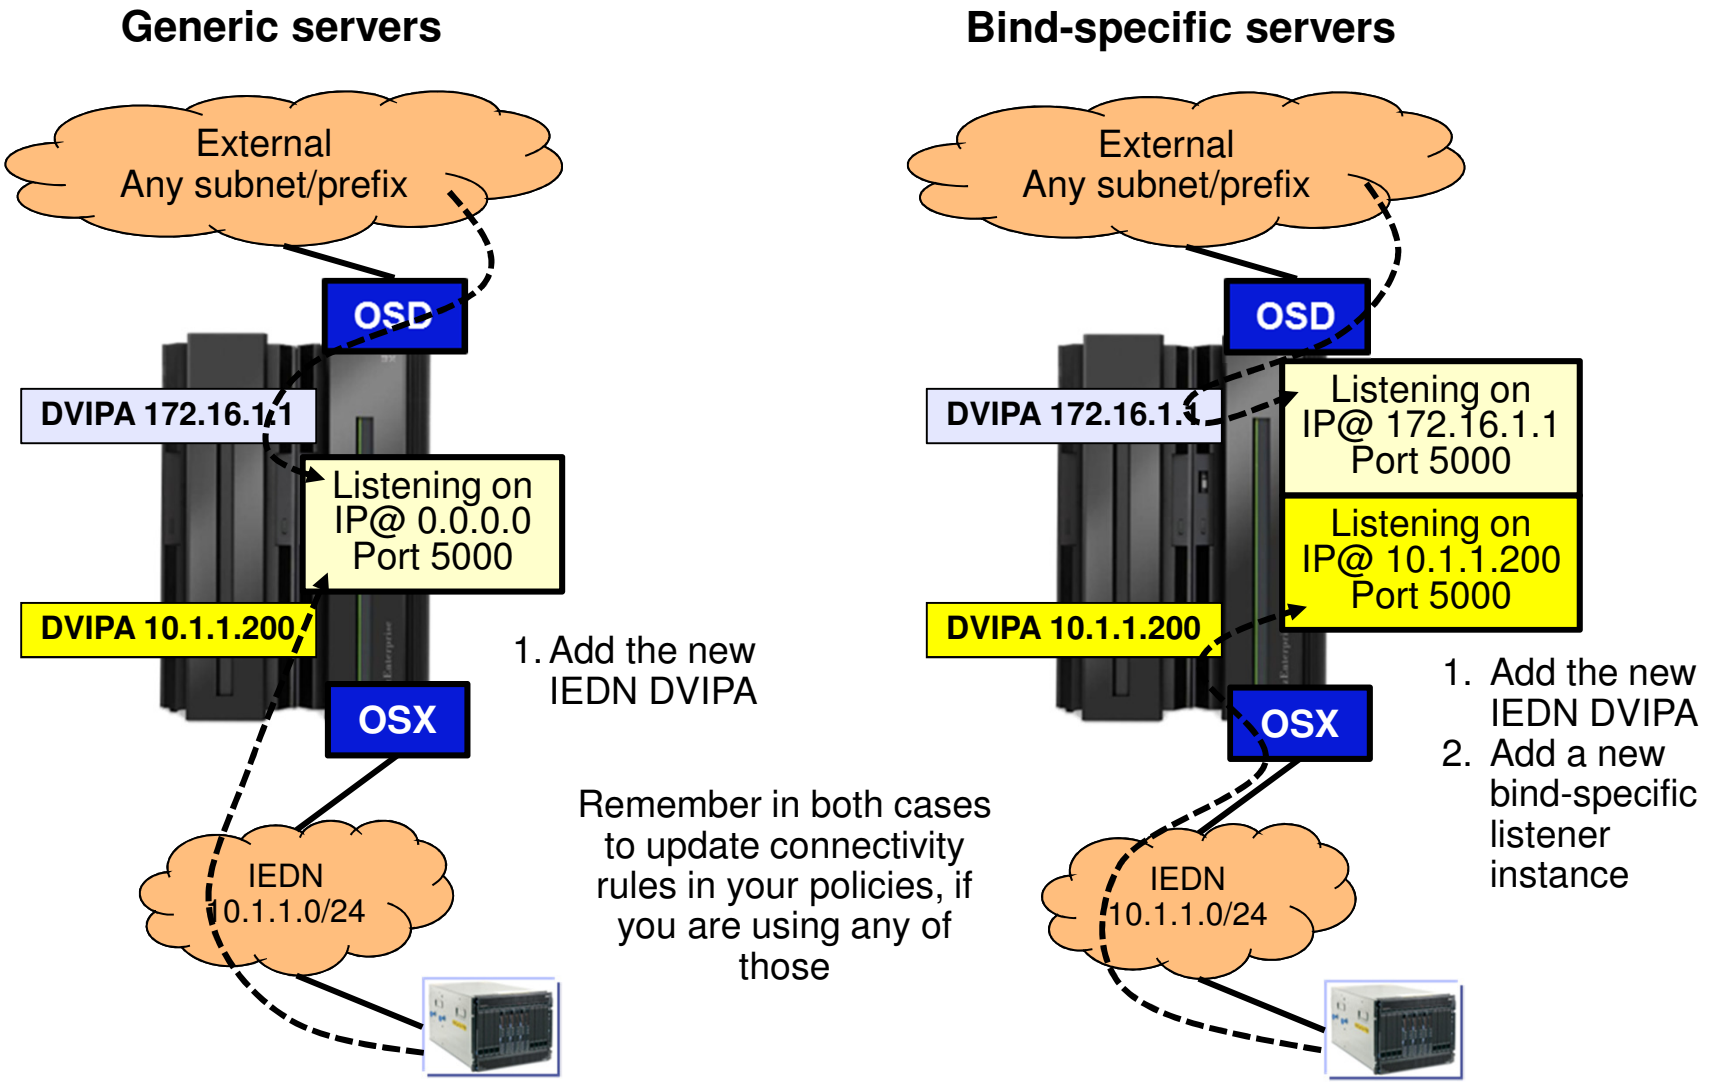

IEDN with multiple VLANs

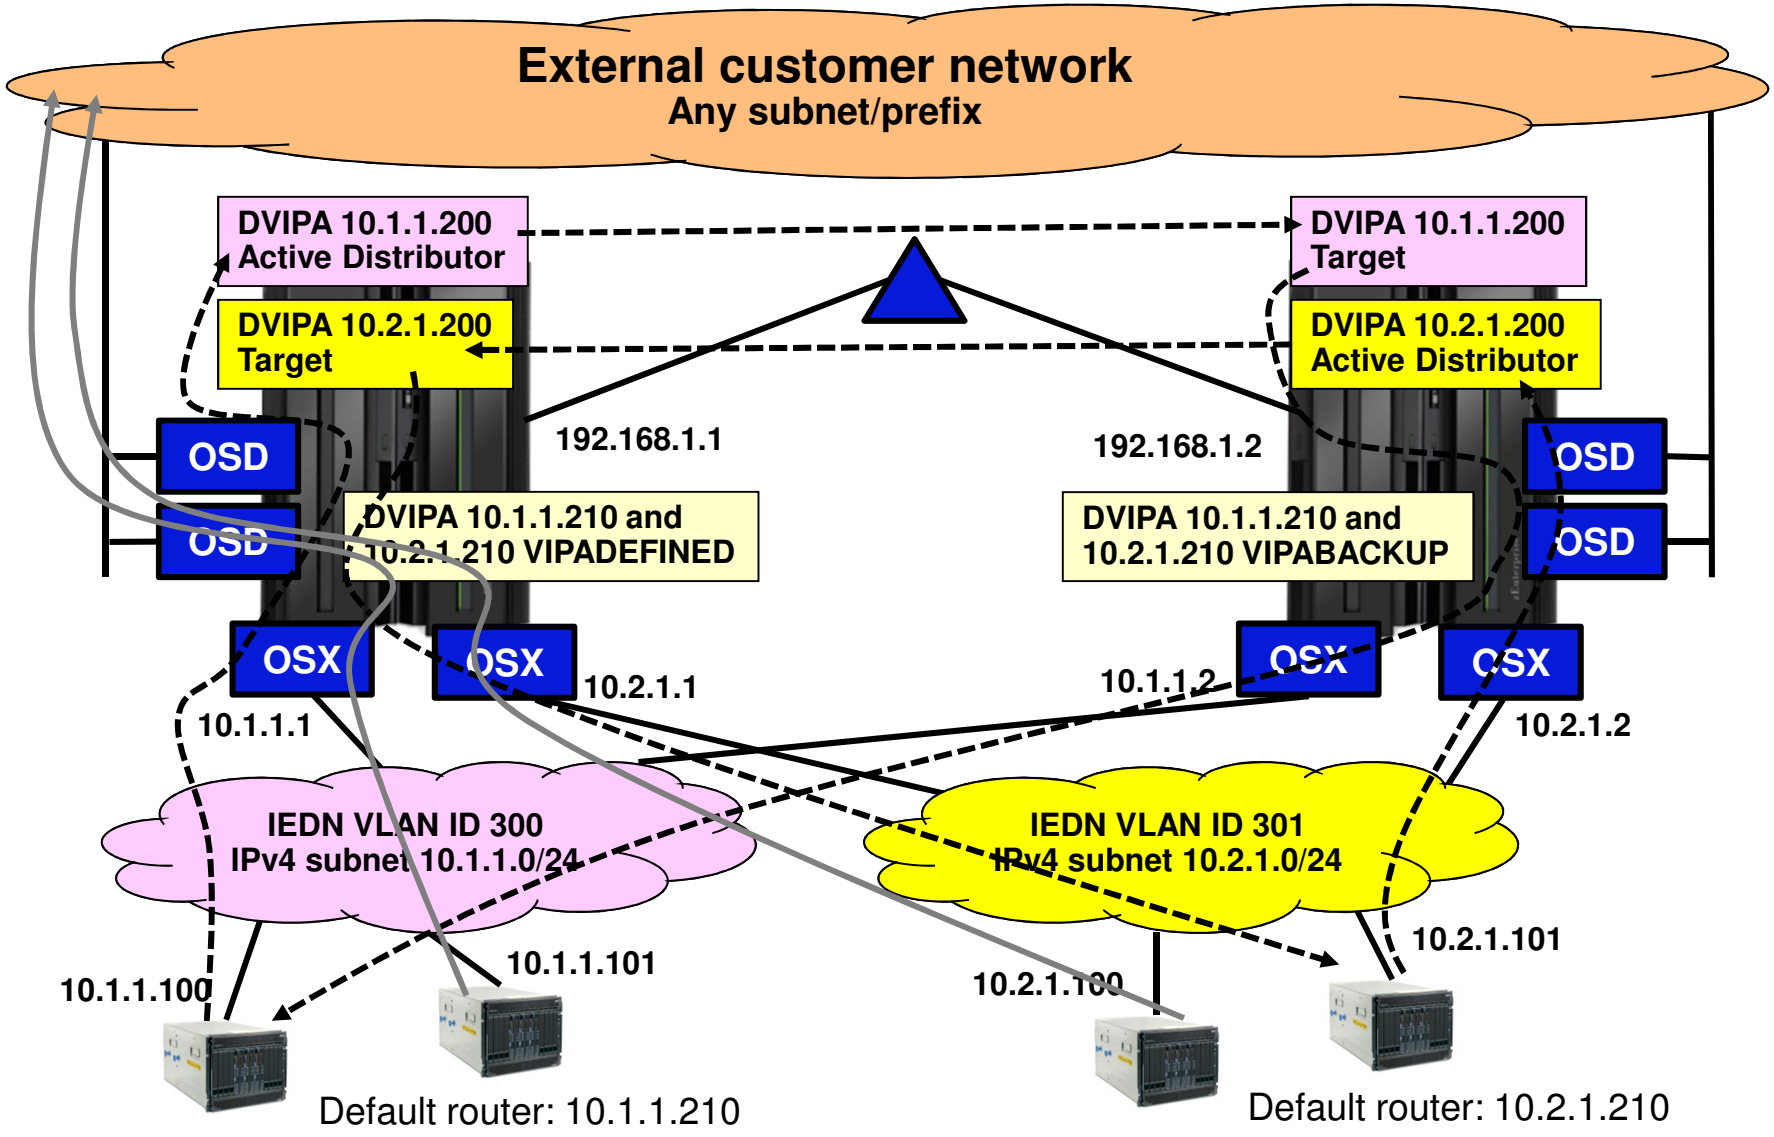

Considerations for multiple VLAN IEDN access

- Basic rule is that IEDN hosts can access DVIPAs in their own IEDN VLAN subnet/prefix
  - Exception is if a z/OS system is used as the default router, in which case the default router z/OS system may know how to reach the DVIPA from another IEDN VLAN
    - This may result in non optimal routing: route to the default router, forward to the distributing z/OS system, connection forwarding to the target z/OS system, which may or may not have direct access to the source IEDN VLAN (in which case another routing hop via another z/OS system will be needed on the way back)
    - You can control at a very detailed level what amount of routing z/OS is to perform in these cases through IPSec filter rules on z/OS
- Each IEDN VLAN must have a default router DVIPA defined on z/OS - if z/OS is used as the gateway to the external network
- **Remember**: z/OS will never route traffic between IEDN VLANs

Scenario 3: DataPower and Sysplex Distributor

Introducing the WebSphere DataPower XI50z for zEnterprise

Announced 2/15/2011

- XI50 features optimized in a dense, high compute IBM zEnterprise BladeCenter Extension (zBX) form-factor
- Supports all ESB, Security, and Integration capabilities of DataPower XI50 v3.8.1
- **Purpose-built** Integration Appliance
  - *Sysplex*, CICS, IMS, DB2, SAF, RACF integration
- **Highest capacity** DataPower appliance for SOA workloads optimized for zEnterprise environments
- Tightly **integrated** with zEnterprise
  - Unified hardware and firmware management through the Hardware Management Console (HMC)
  - Inherits serviceability, monitoring and reporting capabilities of zEnterprise

Sysplex Distributor support for DataPower

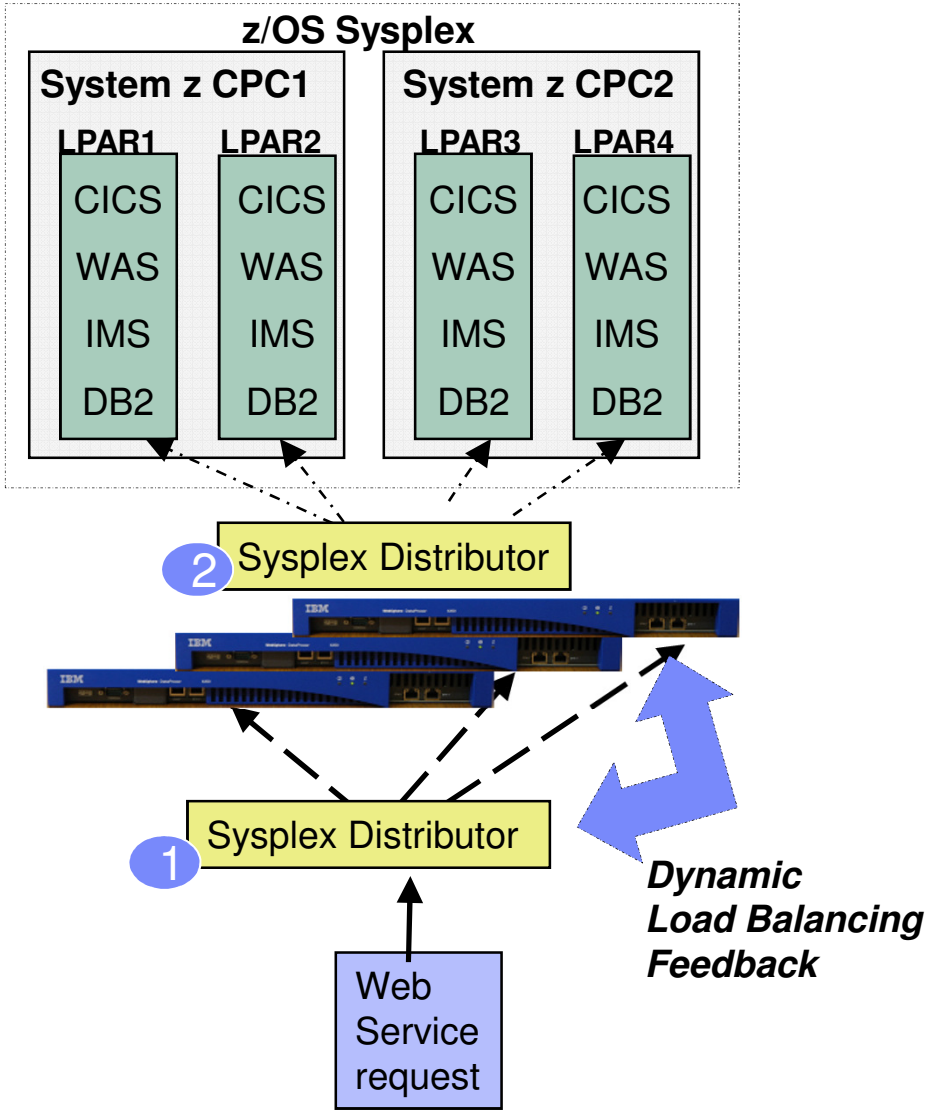

- Introduced in z/OS V1R11 Communications Server
  - DataPower Support in Firmware 3.8.1
- Allows Sysplex Distributor to load balance connections to a cluster of DataPower appliances that "front-end" a z/OS Sysplex environment (Tier 1)
  - Complements Sysplex Distributor support for back-end workflows (DataPower to z/OS – Tier 2)
- Sysplex Distributor and DataPower communicate over a control connection
  - Allows SD to have awareness of state and utilization levels of each DataPower instance
  - Facilitates TCP connection management and use of GRE to preserve client's IP address visibility to DataPower

WebSphere DataPower XI50z – Sysplex Distributor use case

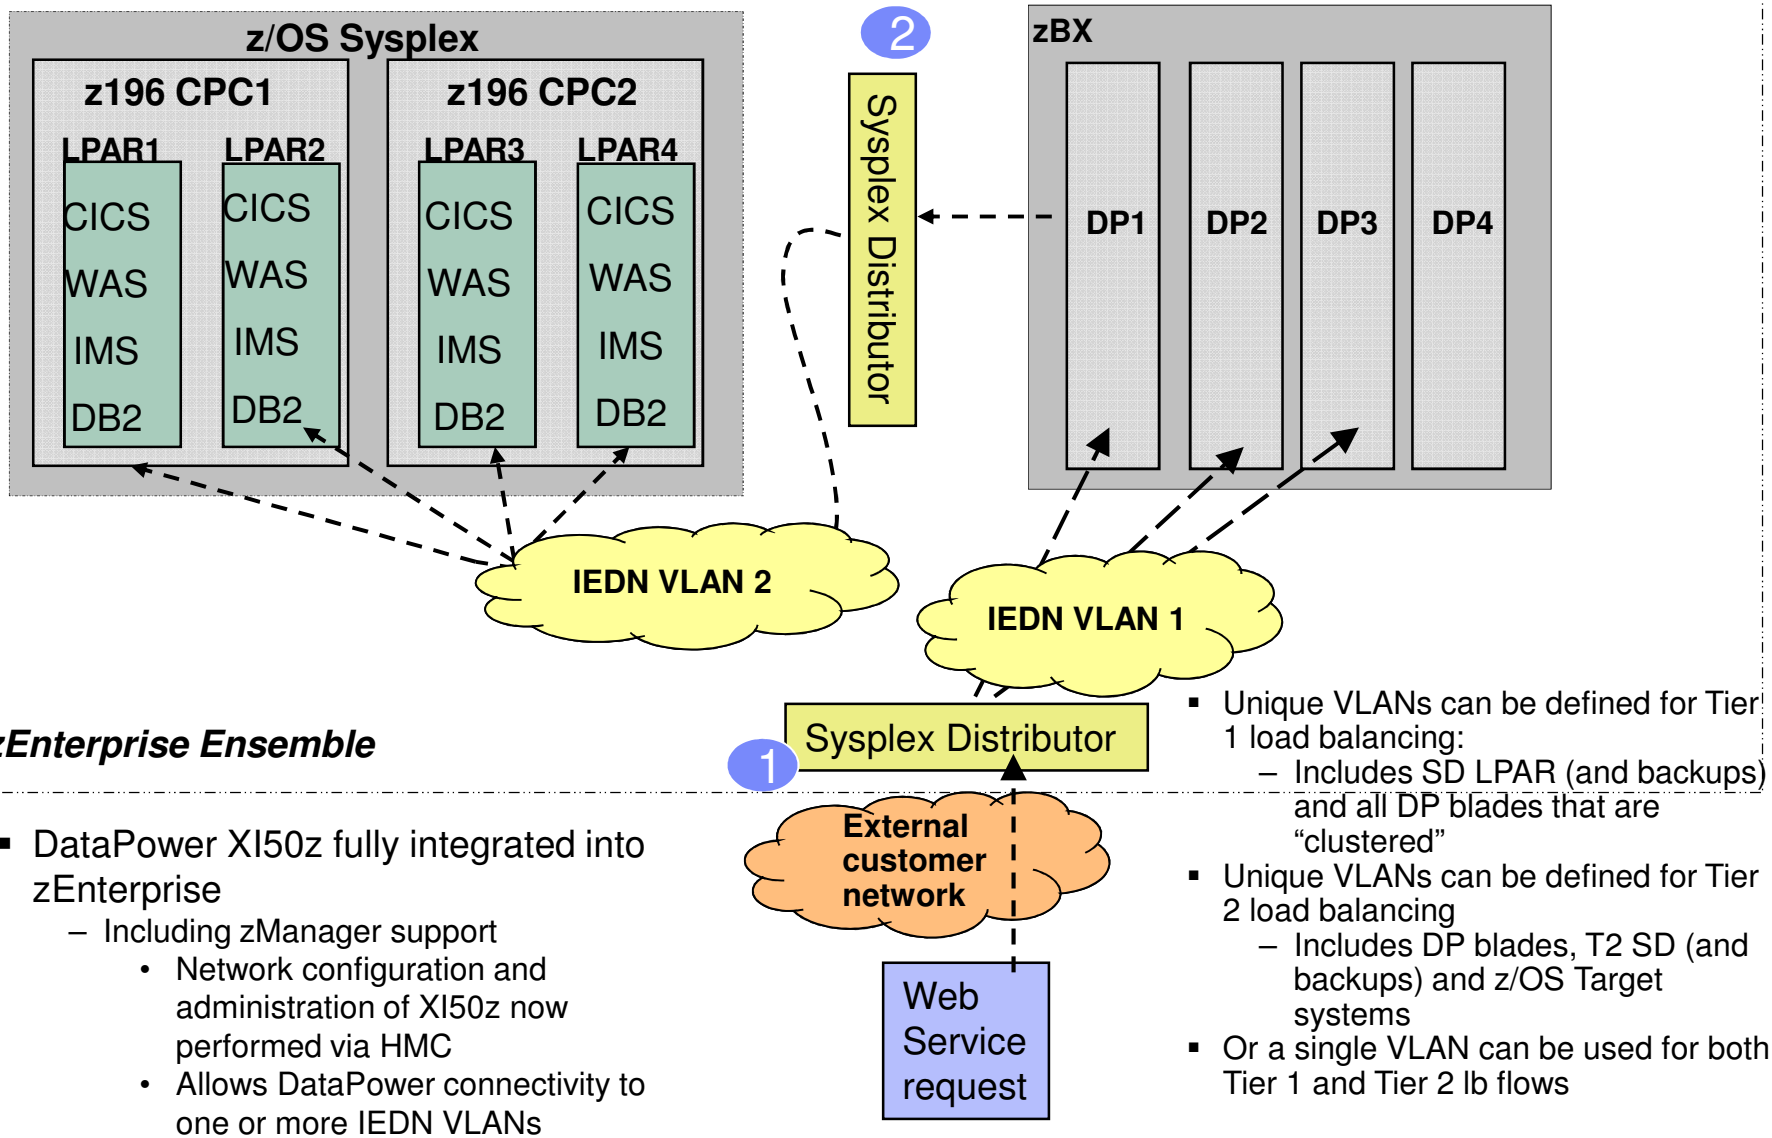

- DataPower XI50z fully integrated into zEnterprise
  - Including zManager support
    - Network configuration and administration of XI50z now performed via HMC
    - Allows DataPower connectivity to one or more IEDN VLANs

- Unique VLANs can be defined for Tier 1 load balancing:
  - Includes SD LPAR (and backups) and all DP blades that are "clustered"
- Unique VLANs can be defined for Tier 2 load balancing
  - Includes DP blades, T2 SD (and backups) and z/OS Target systems
- Or a single VLAN can be used for both Tier 1 and Tier 2 lb flows

Scenario 4:
IBM Smart Analytics Optimizer

IBM Smart Analytics Optimizer

Capitalizing on breakthrough technologies to accelerate business analytics

What is it?

The IBM Smart Analytics Optimizer is a workload optimized, appliance-like, add-on, that enables the integration of business insights into operational processes to drive winning strategies. It accelerates select queries, with unprecedented response times.

How is it different?

- **Performance:** Unprecedented response times to enable 'train of thought' analyses frequently blocked by poor query performance.
- **Integration:** Connects to DB2® through deep integration providing transparency to all applications.
- **Self-managed workloads:** Queries are executed in the most efficient way.
- **Transparency:** Applications connected to DB2, are entirely unaware of IBM Smart Analytics Optimizer.
- **Simplified administration:** Appliance-like hands-free operations, eliminating many database tuning tasks.

Faster insights for enabling new opportunities

Optimizing to the Workload
Marrying the best of each system

Total solution remains centrally managed by System z…

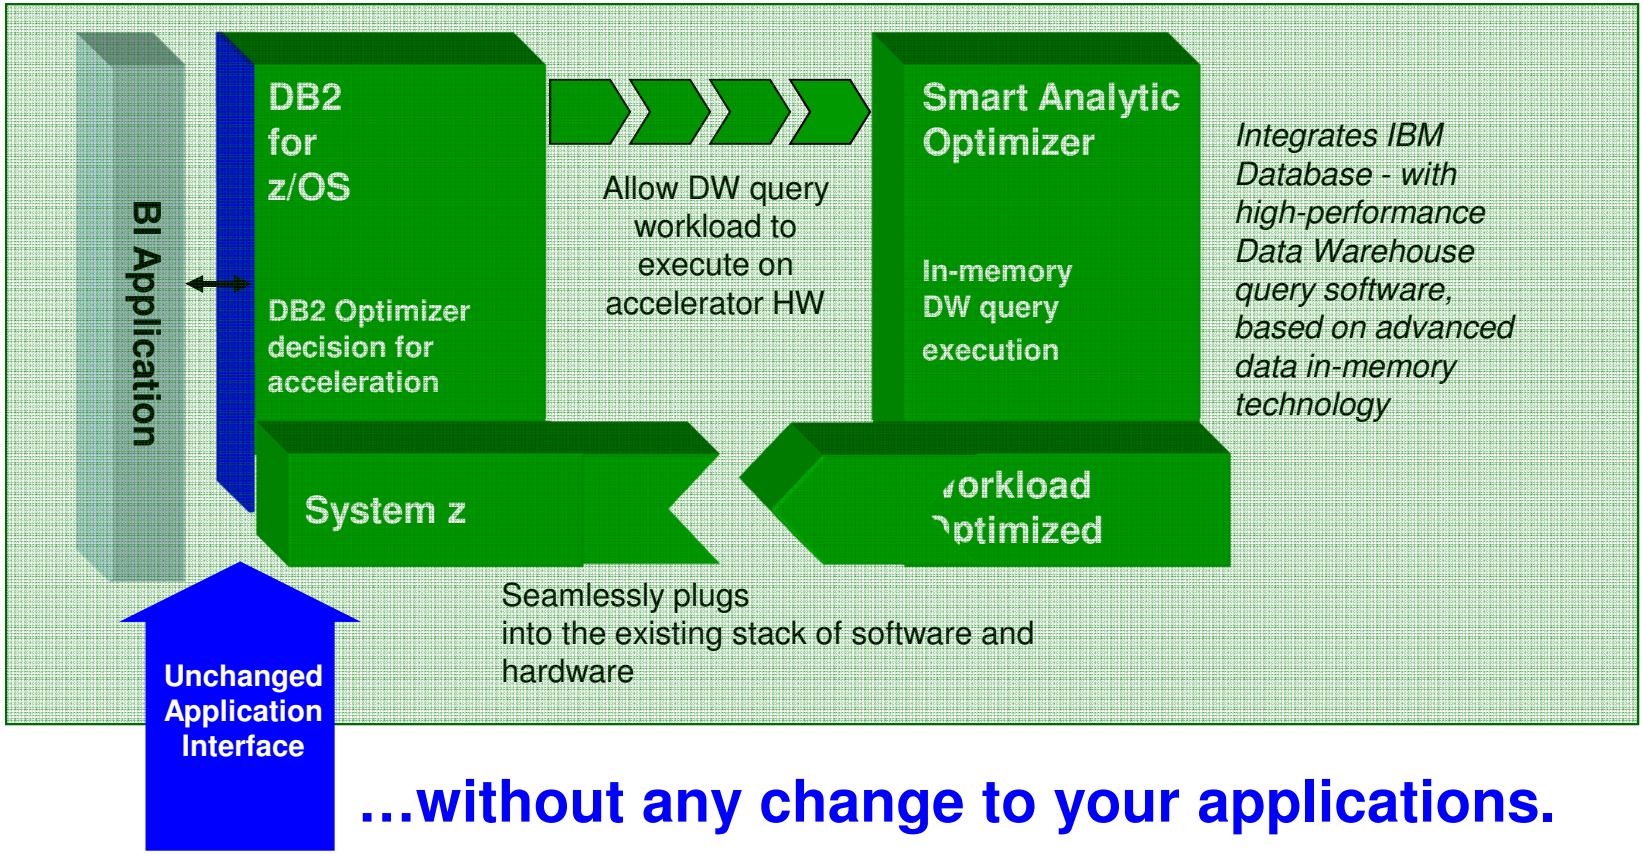

…without any change to your applications.

zBX / Smart Analytics Optimizer configuration example

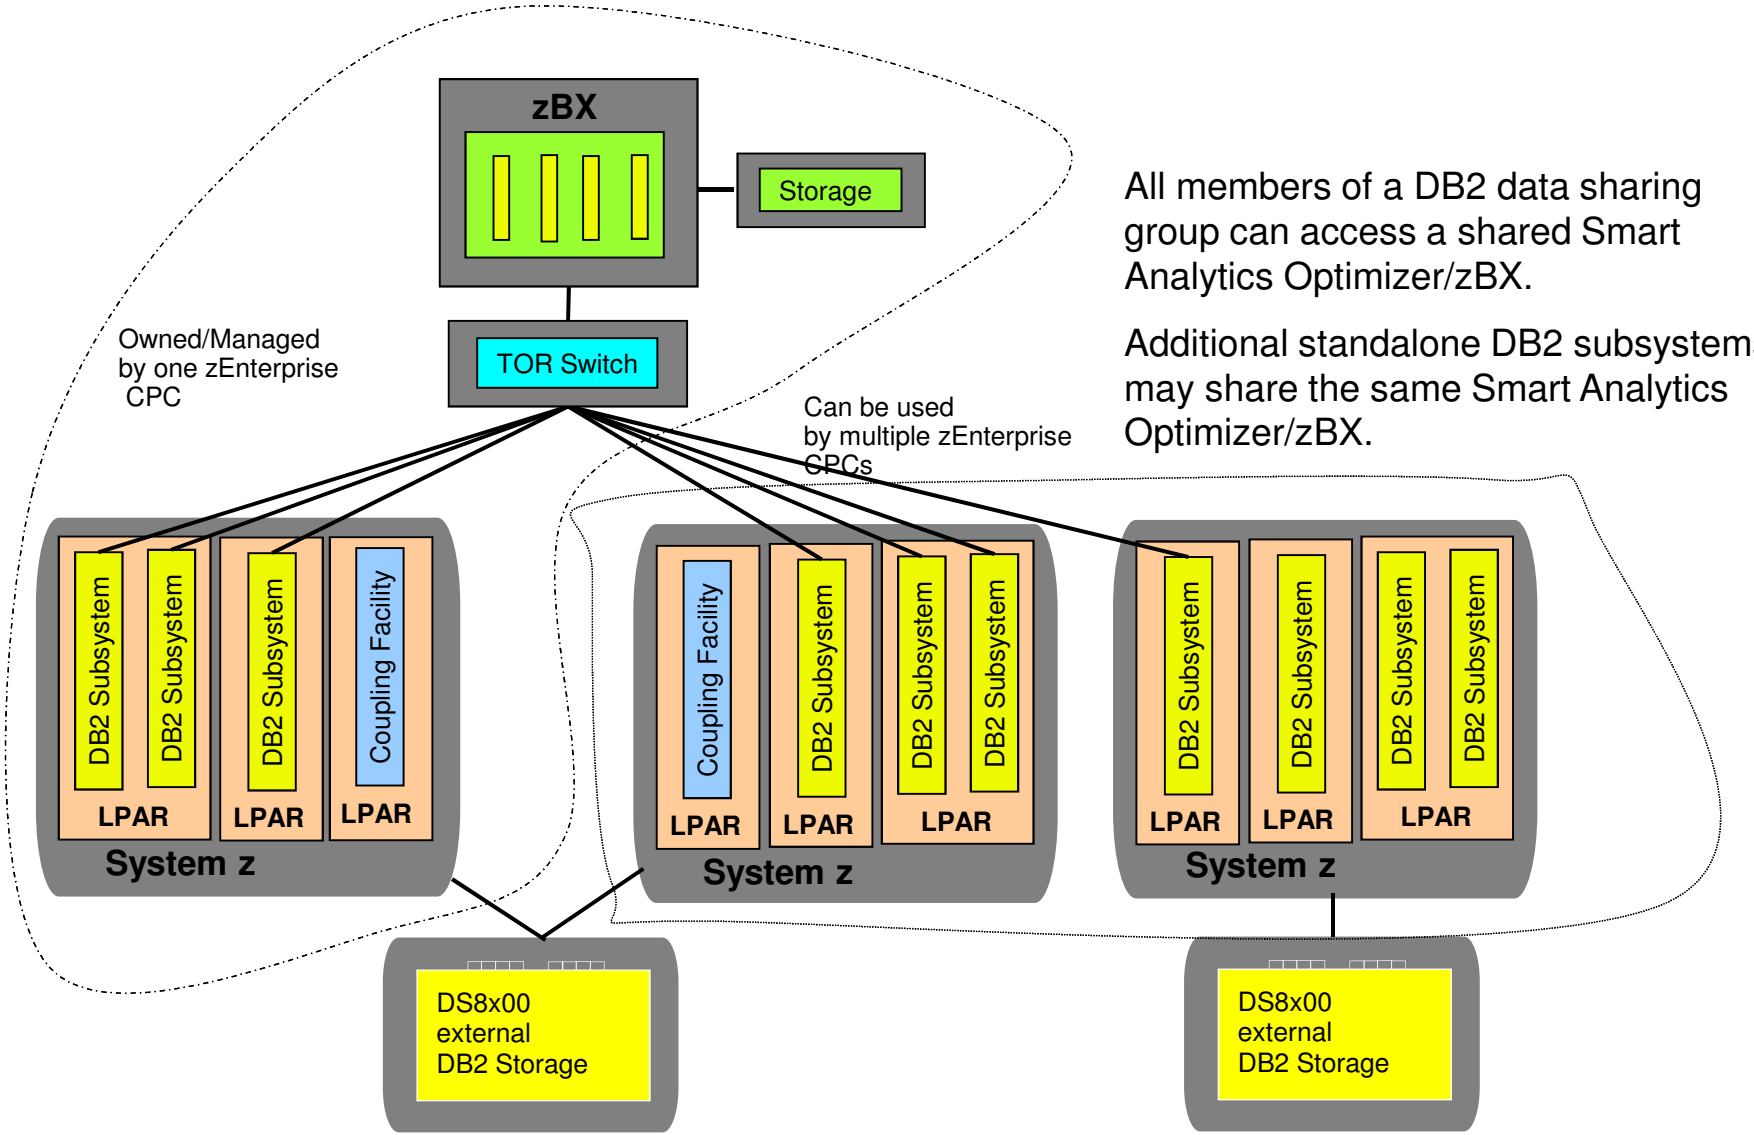

All members of a DB2 data sharing group can access a shared Smart Analytics Optimizer/zBX.

Additional standalone DB2 subsystems may share the same Smart Analytics Optimizer/zBX.

IBM Smart Analytics Optimizer Overview

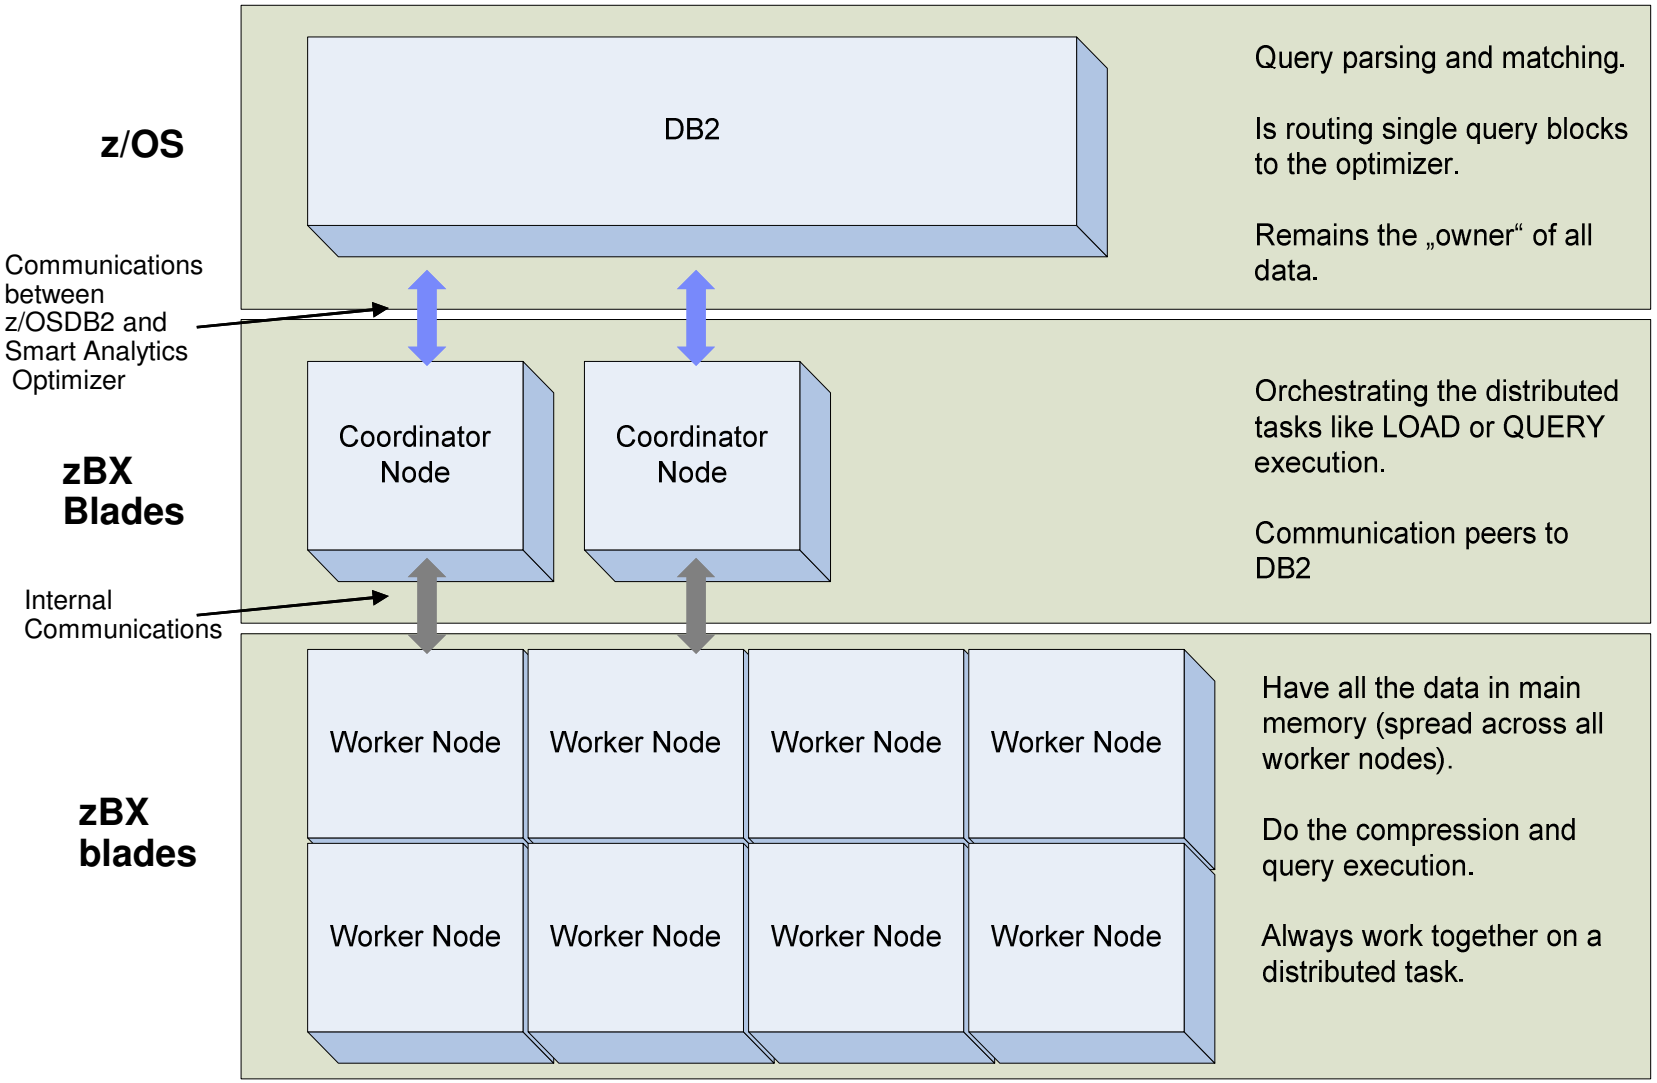

IBM Smart Analytics Optimizer – Network Considerations

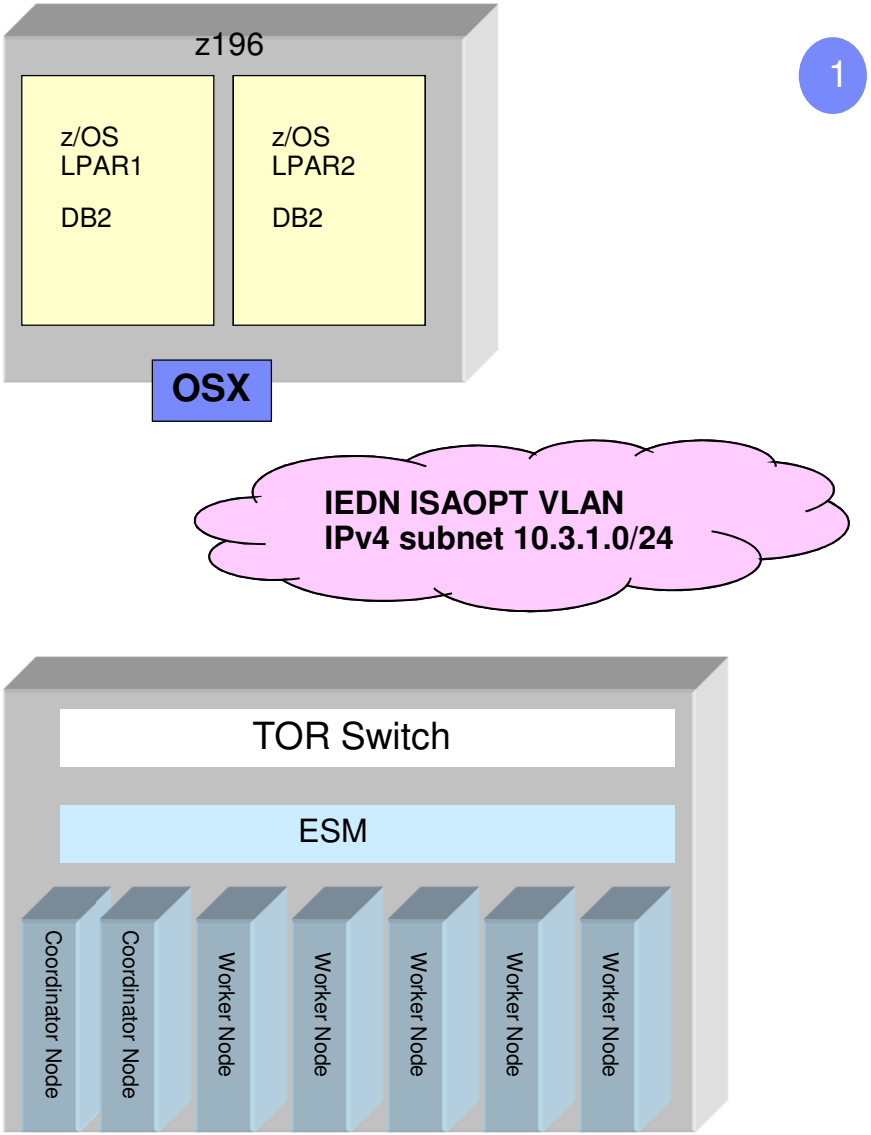

1. Define a subnet for use between z/OS and IBM Smart Analytics

   Need to ensure that subnet is large enough to accommodate all members of the subnet:
   - Each z/OS System accessing the Smart Analytics Optimizer will require a minimum of 2 IP addresses (one per OSA)
     - May require more if VIPAs are also used on the subnet
   - Each Coordinator Node will require an IP address
     - The number of Coordinator nodes varies based on the size of the Smart Analytics Optimizer system
     - The Coordinator nodes need a ***contiguous range*** of IP addresses
     - Worker nodes ***do not*** need IP addresses from this subnet

IBM Smart Analytics Optimizer – Network Considerations...

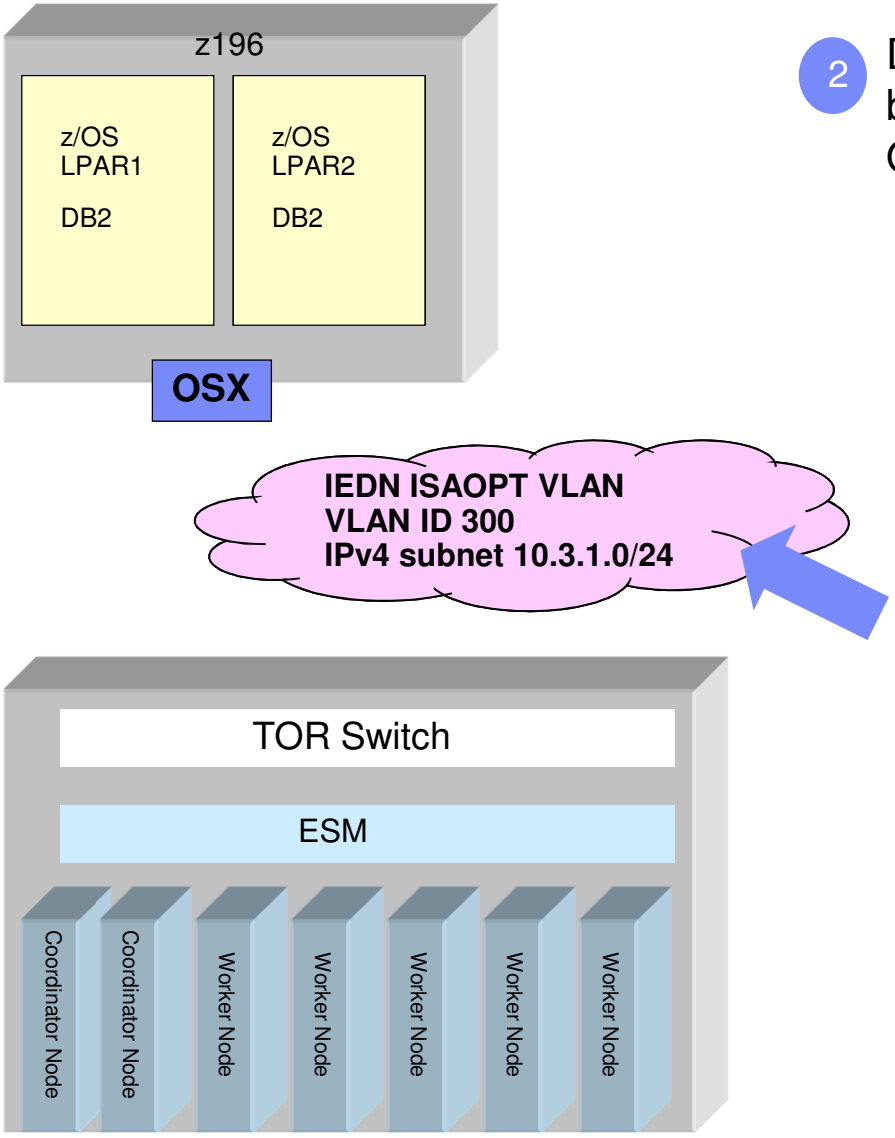

2 Define VLAN for connectivity between IBM Smart Analytics Optimizer and z/OS systems
- Using zManager NVM
- An existing IEDN VLAN can be used if it meets the security/isolation requirements of your environment (this presentation assumes a new VLAN will be set up for use by the Smart Analytics Optimizer workload)

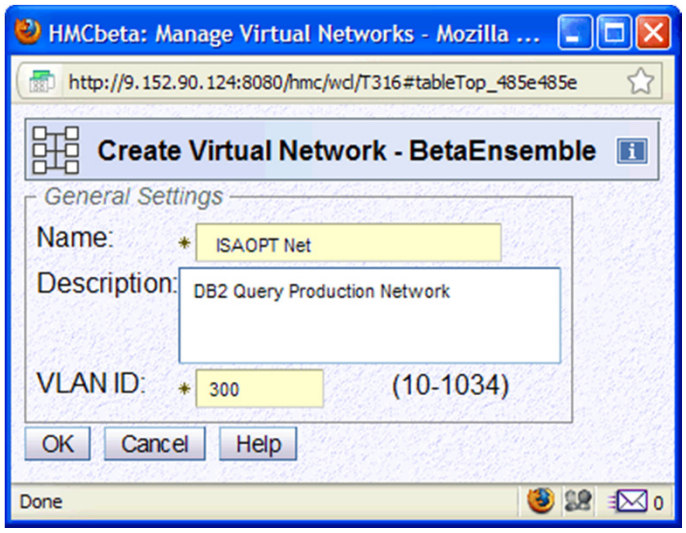

IBM Smart Analytics Optimizer – Network Considerations…

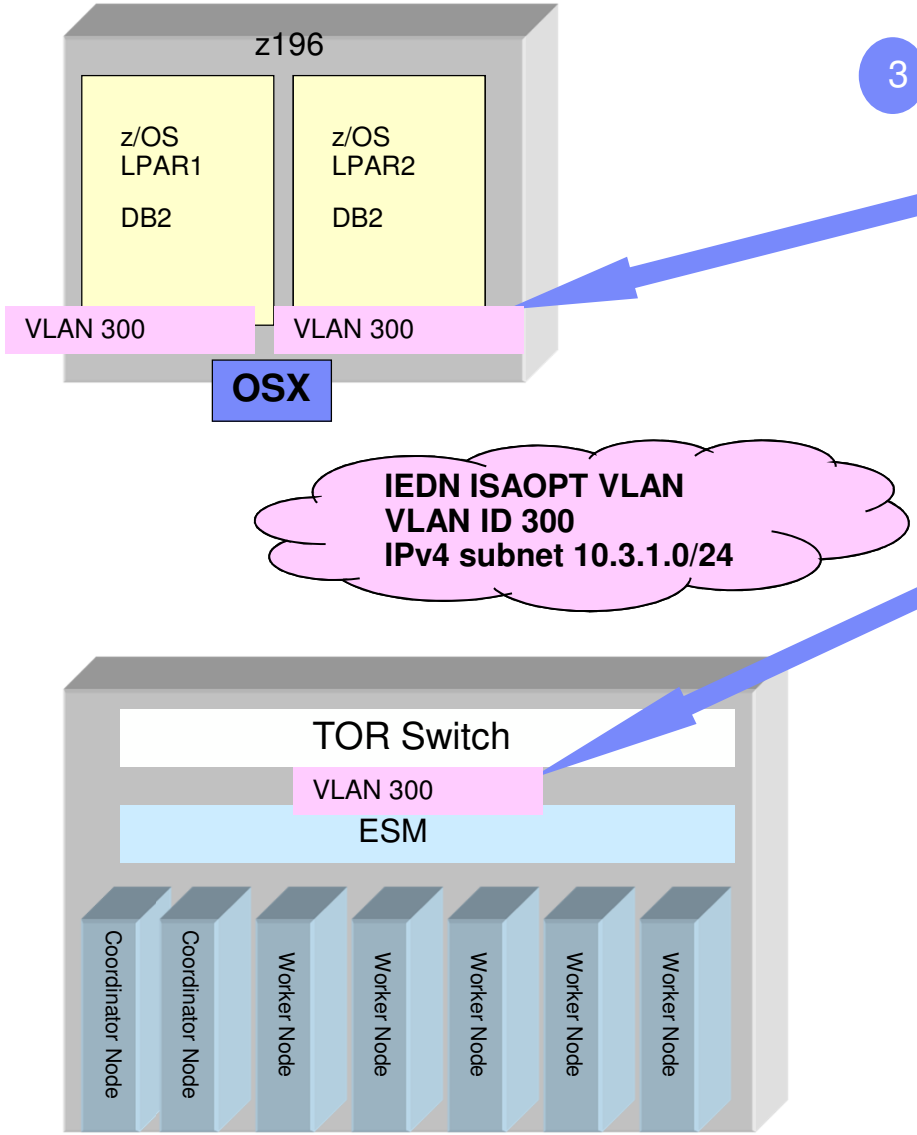

3 Using zManager associate all members of the subnet with the ISAOPT VLAN

1. Associate all z/OS Systems with the ISAOPT VLAN (Using "Add Hosts to Virtual Networks" task)
2. Associate the Smart Analytics Optimizer with the ISAOPT VLAN
   - Note this is done via special zManager Configuration task "Configure Top-of-Rack (TOR) switch" task
     - The internal TOR ports associated with the ESMs of the Smart Analytics Optimizer blade chassis need to be associated with the ISAOPT VLAN (in access mode)

IBM Smart Analytics Optimizer – Network Considerations…

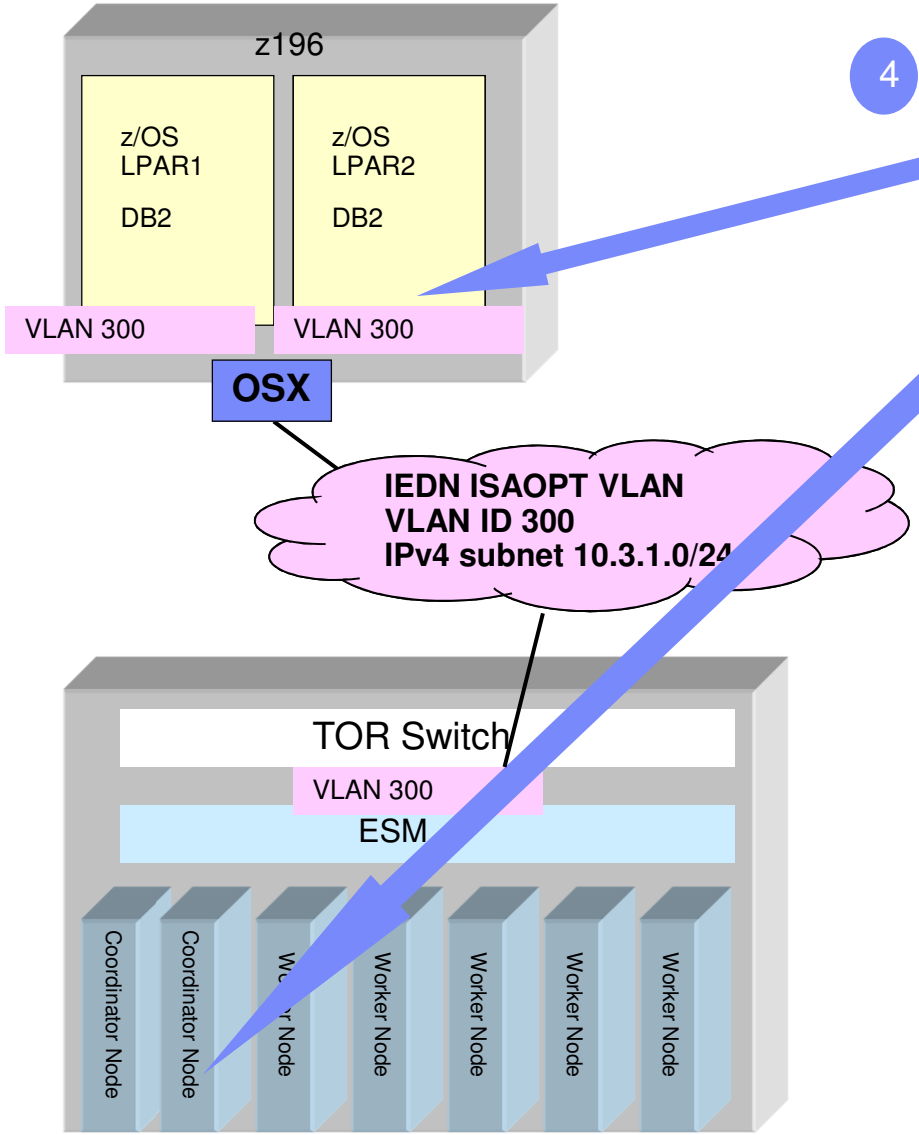

4 Specify IP addresses on ISAOPT VLAN on all subnet members:

1. On z/OS systems using INTERFACE statement
2. Using the Support Element define a range of IP addresses from the subnet for use by the Coordinator Nodes
   - For more details refer to the “IBM Smart Analytics Optimizer for DB2 for z/OS Installation Guide” (SH12-6916)

The source IP address used by DB2 for connection to the ISAO coordinator nodes must belong to the IEDN subnet.

It can be either a DVIPA address per LPAR, or it could be one of the IEDN OSX interface IP addresses.

IBM Smart Analytics Optimizer – Network Considerations….

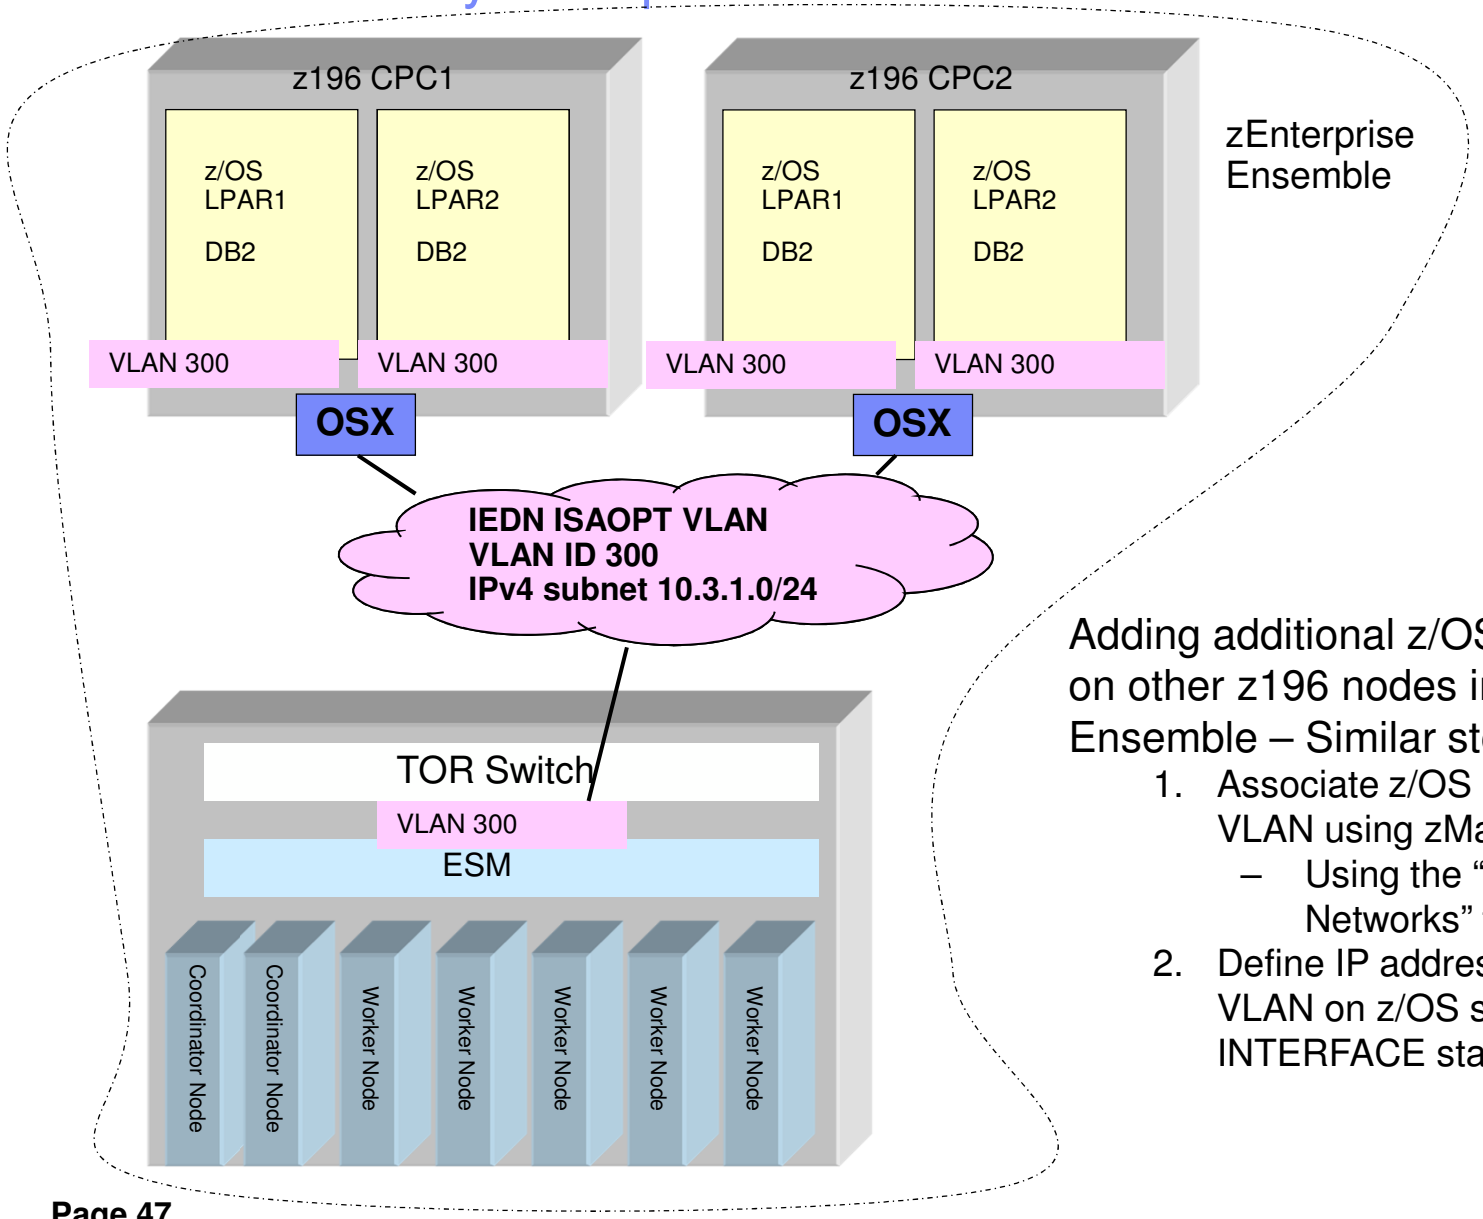

Adding additional z/OS LPARs hosted on other z196 nodes in the same Ensemble – Similar steps as earlier:

1. Associate z/OS LPARs with ISAOPT VLAN using zManager
   - Using the “Add Hosts to Virtual Networks” task
2. Define IP addresses and ISAOPT VLAN on z/OS systems using INTERFACE statement

IBM Smart Analytics Optimizer – Network Considerations…

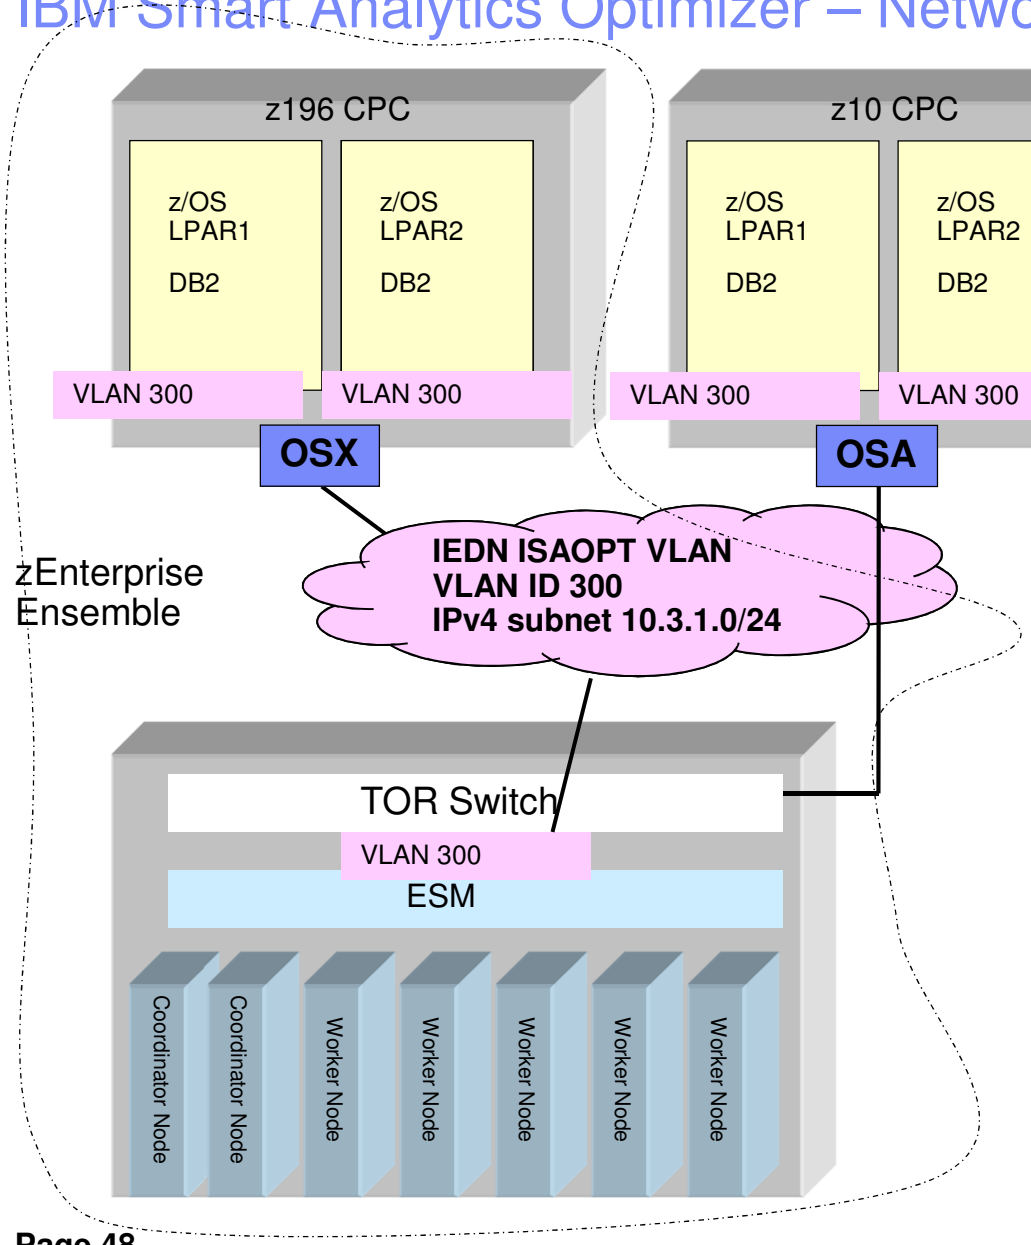

Allowing additional z/OS LPARs hosted on other z10 nodes to use the Smart Analytics Optimizer:

1. OSA Express 3 (OSD) connected directly to external TOR switch ports
   - Access to external TOR ports must be authorized via zManager Configuration task “Configure Top-of-Rack (TOR) switch” task
     - Permit the ISAOPT VLAN on the port and MAC associated with the OSA (OSD)
2. Define IP addresses and ISAOPT VLAN on z/OS systems using INTERFACE statement

For more information

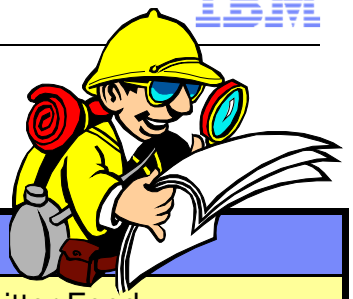

URL	Content
http://www.twitter.com/IBM_Commserver	IBM Communications Server Twitter Feed
http://www.facebook.com/IBMCommserver	IBM Communications Server Facebook Fan Page
http://www.ibm.com/systems/z/	IBM System z in general
http://www.ibm.com/systems/z/hardware/networking/	IBM Mainframe System z networking
http://www.ibm.com/software/network/commserver/	IBM Software Communications Server products
http://www.ibm.com/software/network/commserver/zos/	IBM z/OS Communications Server
http://www.ibm.com/software/network/commserver/z_lin/	IBM Communications Server for Linux on System z
http://www.ibm.com/software/network/ccl/	IBM Communication Controller for Linux on System z
http://www.ibm.com/software/network/commserver/library/	IBM Communications Server library
http://www.redbooks.ibm.com	ITSO Redbooks
http://www.ibm.com/software/network/commserver/zos/support/	IBM z/OS Communications Server technical Support – including TechNotes from service
http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs	Technical support documentation from Washington Systems Center (techdocs, flashes, presentations, white papers, etc.)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)
http://www.ibm.com/systems/z/os/zos/bkserv/	IBM z/OS Internet library – PDF files of all z/OS manuals including Communications Server

For pleasant reading ….