

Smarter Systems for a Smarter Planet



SESSION 10724

SHARE Atlanta 2012

Tuesday, March 13, 2012: 4:30 PM-5:30 PM

zEnterprise System – Secure Networking with the zEnterprise Ensemble (Part 3)

Gwendolyn J. Dente – gdente@us.ibm.com
IBM Advanced Technical Skills (ATS)
Gaithersburg, MD, USA



IBM Advanced Technical Skills (ATS)

© 2012 IBM Corporation



zEnterprise System – Secure Networking with the zEnterprise Ensemble (Part 3)

Session number:	10724
Date and time:	Tuesday, March 13, 2012: 4:30 PM-5:30 PM
Location:	Hickory Room at the OMNI Hotel CNN
Program:	Communications Infrastructure
Project:	Network Management and Security
Track:	Capitalizing on zEnterprise and Network Support and Management
Classification:	Technical
Speaker:	Gwendolyn J. Dente, IBM Advanced Technical Skills (ATS)
Abstract:	IBM's zEnterprise System introduces new paradigms for managing a hybrid topology that integrates heterogeneous platforms into an "Ensemble." The new Ensemble paradigm requires revolutionary thinking about securing a single federated platform while allowing multiple administrators to work on the zEnterprise System "system of systems." The federation of multiple IT architectures into an Ensemble, all managed from a single HMC with Unified Resource Manager ("zManager"), provides not only simplified and centralized administration, but also centralized and simplified security within a virtualized, high-speed environment. Even the challenges of government security mandates and Payment Card Industry (PCI) compliance can be met with Ensemble networking. Come to this session to understand the security foundations of the zEnterprise System. This is the third in a series of three presentations on zEnterprise networking.

SHARE session specifications.



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:

- | | | | | |
|--|--|---|--|---|
| <ul style="list-style-type: none">Advanced Peer-to-Peer Networking®AIX®alphaWorks®AnyNet®AS/400®BladeCenter®Candle®CICS®DataPower®DB2 ConnectDB2®DRDA®e-business on demand®e-business (logo)e-business (logo)®ESCON®FICON® | <ul style="list-style-type: none">CDDM®GDPS®Geographically Dispersed Parallel SysplexHyperSocketsHPR Channel ConnectivityHyperSwapi5/OS (logo)i5/OS®IBM eServerIBM (logo)®IBM®IBM zEnterprise™ SystemIMSInfiniBand®IP PrintWayIPDSiSeriesLANDP® | <ul style="list-style-type: none">Language Environment®MCSeries®MVSNetView®OMEGAMON®Open PowerOpenPowerOperating System/2®Operating System/400®OS/2®OS/390®OS/400®Parallel Sysplex®POWER®POWER7®PowerVMPR/SMpSeries®RACF® | <ul style="list-style-type: none">Rational Suite®Rational®RedbooksRedbooks (logo)Sysplex Timer®System i5System p5System x®System z®System z9®System z10Tivoli (logo)®Tivoli®VTAM®WebSphere®xSeries®z9®z10 BCz10 EC | <ul style="list-style-type: none">zEnterprisezSeries®z/Architecturez/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.



Acknowledgments

From IBM z/OS® Communications Server Development:

Kim Bailey, Mike Fox
Jeff Haggar, Bebe Isrel
Gary McAfee, Jerry Stevens

From IBM z/VM® Development:

Alan Altmark, Angelo Macchiano
Rick Tarcza, Romney White

From IBM International Technical Support Organization (ITSO)

Mike Ebbers, Parwez Hamid
Karl-Erik Stenfors, William G. White

From IBM Advanced Technical Skills (ATS)

Mary Astley, Jean Marc Darees
Harv Emery, John Goodyear
Gregory Hutchison

From IBM S&D Group, Systems Software Development

Friedrich M Welter

From IBM Europe

Lennie J Dymoke-Bradshaw
Peter Redman

Thanks to the many people and resources I found to help me understand the information I used to write this presentation.



Frequently Asked Questions about Securing the Ensemble Network

- **How do I limit and secure access to the Ensemble Definitions?**
 - Use controls at the HMC (Unified Resource Manager or zManager functions)
 - Use existing HMC controls
- **How do I limit and secure access from within the IEDN to the Ensemble Virtual Servers and Networks?**
 - All Virtual Servers and VLANs must pass through “access points” (Hypervisors and TORs) where their authorization is confirmed.
 - Authorize the Virtual Servers to become Ensemble Members
 - Authorize the Virtual Servers to send data across the Ensemble Networks
 - Authorize the Virtual Servers to exploit only certain VLAN IDs
 - Use existing security techniques
 - Userid and passwords
 - Access controls to access storage
 - Firewalls and IP Filtering
 - Encryption
 - Etc.
- **How do I limit access to the Ensemble by the External Networks?**
 - Deploy routers and Firewalls to Permit or Deny traffic
 - Implement controls at the Top-of-Rack (TOR) switch and at the LPARs to limit access
- **With a hybrid solution that combines heterogeneous platforms into a single Enterprise System, how do I provide the network segmentation that many security mandates (like PCI) require?**
 - Centrally assign and have Hypervisors enforce VLAN IDs in the data network (IEDN)
 - Have Hypervisors and the TOR enforce Virtual Media Access Control (VMAC) addresses
- **Why are VLAN implementations in the Ensemble considered more secure than VLANs in a non-Ensemble environment?**
 - Without Ensemble, Network Interface Cards (NICs) of platforms connect to switch ports, which, if improperly configured, could allow miscoded VLAN IDs at the server level to gain access to the network infrastructure.
 - With Ensemble, the Unified Resource Manager firmware in the hypervisors and virtual switches ensures that servers are only allowed to connect to the VLAN(s) in the IEDN that they are explicitly authorized for. A misconfiguration of a virtual server (e.g. incorrect VLAN ID) results in a failure to connect to the IEDN.

There are many controls to permit or deny access to the IEDN. The Hypervisor controls (including the hypervisor component known as a VSwitch) operate within the IEDN to authorize access. The internal ports (Blade switch modules and TOR ports) are configured as trunk mode and allow all IEDN VLANs to flow BUT with the understanding that someone else – the Hypervisor and VSwitch) is enforcing the access control. For access from the external network into the IEDN, the Top-of-Rack (TOR) switches can implement access controls to the IEDN as can security technologies outside of or within an Ensemble LPAR or Virtual Server.

HMC Security White Paper: IBM System z Hardware Management Console Security White Paper. Author Kurt Schroeder (schroedk@us.ibm.com), Sept. 2008

<http://nascpok.pok.ibm.com/rsf/zHMCSecurityWhitepaper.pdf>

zEnterprise Network Security White Paper (ZSW03167-USEN-00) and Other Resources

www.ibm.com/systems/z/resources (Select “Literature” Entries)


http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfid=ZSW03167USEN&attachment=ZSW03167USEN.PDF

zEnterprise Network Security Frequently Asked Questions:

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FQ130131>

Agenda

- ❑ The New Paradigms for Intelligent Computing and for Security
- ❑ Unified Resource Manager and its Role in Security
- ❑ Security Mandates for Networking Segmentation (Isolation)
- ❑ Secure Ensemble Networking with VLANs
 - ❑ Security Enforcement Points
 - ❑ Within IEDN and between External Data Network & IEDN
- ❑ Ensemble Access and Routing Controls
- ❑ HMC Security
 - ❑ Resource Roles, Task Roles
- ❑ Security Framework: Layers of Security
- ❑ Summary of Security
- ❑ Appendices: OSM & IEDN Security Tables, OSX & "ISOLATE", References



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 – Secure Networking with the zEnterprise Ensemble (Part 3)


The New Paradigms for Intelligent Computing and for Security



Intelligent Computing for a Smarter Planet™

New possibilities:

- Breakthrough productivity
- Accelerated value creation
- Increased velocity



The zEnterprise™ System addresses all three client imperatives for achieving a Smarter Planet. Each of the initiatives is reflected in the architecture of the Hybrid Solution. Pervasive themes in the zEnterprise System design architecture are: Virtualization, Centralization, and Redundancy, which lead to Simplification, enhanced Performance, enhanced Availability, enhanced Security, and enhanced Management.



zEnterprise Value

1. Network Simplification (“Network in a Box”)

- Single physical network and IBM zEnterprise BladeCenter® Extension (zBX) “package” (physical integration)
 - “Ensemble” and Clustering of Heterogeneous Resources & Capabilities
 - Centralization of Appliances, Distributed Systems at zEnterprise
- Central point of Management for heterogeneous platforms (HMC/SE)
- Reduced network path length; reduced number of hops
- Permits sharing of the same IP network by HiperSockets and the IEDN OSA ports
- Co-location of Business Processes and Business Data
 - Reduce path length and latency issues
- “Fit for Purpose” (select the optimal platform for the workload type)

2. Secure communications

- Physical security (internal network equipment)
- Logical security (controlled access through centralized definition at zManager of the Hypervisor functions)
- Network Virtualization and Isolation within Internal Networks

3. High Availability

- Redundant Network Hardware
- Logical failover

4. Unique IBM System z® QoS and Performance

- Exploit centralized security of RACF®
- Improved throughput (see 1)

5. Increased Opportunity for Collaboration of Technical Talent

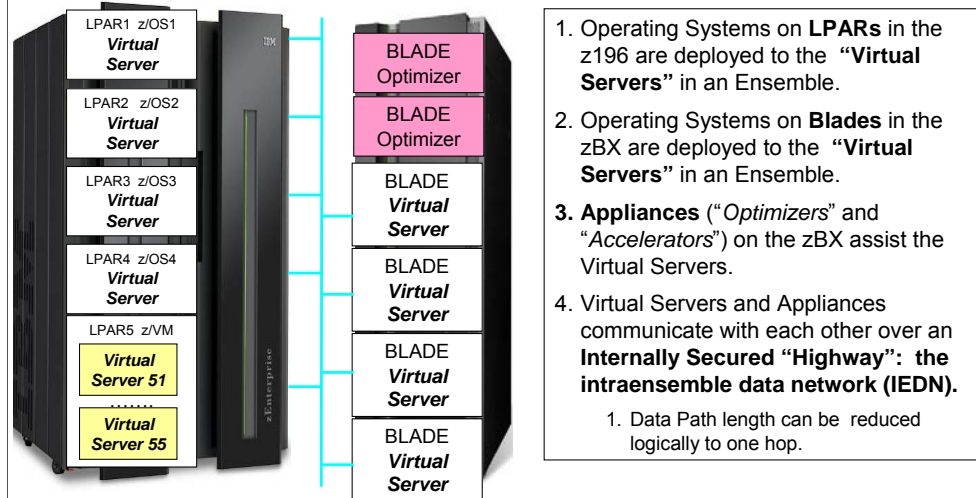
- Broadening the base of System z Skills
- Make System z relevant to a new IT Generation

The IBM zEnterprise™ System is a first-of-a-kind workload-optimized multiplatform (or multi-architecture) computing environment that spans (and tightly integrates) mainframe and distributed technologies. This system (of systems) consists of an IBM zEnterprise™ 196 (z196), the IBM® zEnterprise™ BladeCenter® Extension (zBX) Model 002, and the IBM zEnterprise™ Unified Resource Manager (“Unified Resource Manager” or “zManager”). The z196 is designed with improved scalability, performance, security, resiliency and availability. The Unified Resource Manager, working with the z196, the zBX infrastructure, and the attached blades, can help to deliver end-to-end virtualization and management providing the ability to align the technology deployment environment according to individual workload requirements.

A Hybrid Environment in the Ensemble: LPARs, Appliances, Distributed Operating Systems

LPARs and Blades as Virtual Servers
Communications Private Super Highway = IEDN

Server Consolidation and Federation on zEnterprise



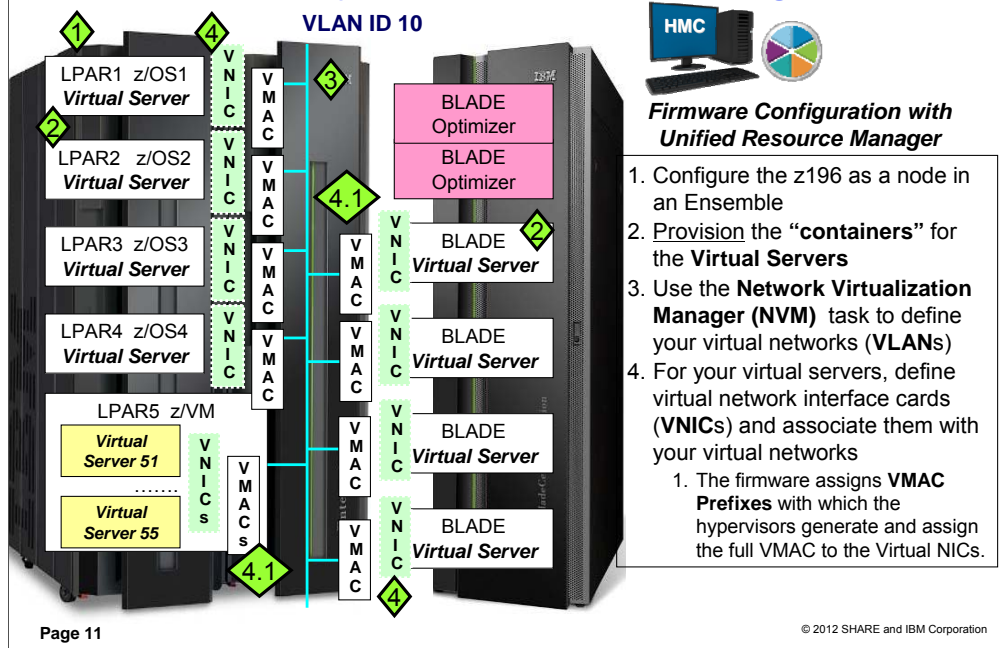
1. Operating Systems on **LPARs** in the z196 are deployed to the “**Virtual Servers**” in an Ensemble.
2. Operating Systems on **Blades** in the zBX are deployed to the “**Virtual Servers**” in an Ensemble.
3. **Appliances** (“*Optimizers*” and “*Accelerators*”) on the zBX assist the Virtual Servers.
4. Virtual Servers and Appliances communicate with each other over an **Internally Secured “Highway”**: the **intraensemble data network (IEDN)**.
 1. Data Path length can be reduced logically to one hop.

In this visual you see how formerly Distributed Servers can be consolidated at a zEnterprise site to form a hybrid environment in an Ensemble that can be managed as a federated whole with selected LPARs in the zEnterprise. Appliances like ISA Optimizer and DataPower can be integrated as well by housing them in the blades of the zBX. This type of consolidation not only can reduce the path length from multiple hops to reach a server in the external network to a single hop to reach the same server in a zBX, but it can also transform the data path into an internal path that can be inherently more secure than a path over the external network.

The “Virtual Server” is a “Virtual Entity” that is defined to the Ensemble. As a result, z/OS is not a virtual server but rather an Operating System inside the Virtual Server; you can think of a native LPAR as a Virtual Server. On the other hand, under z/VM, the Guest Machine would be the Virtual Server and Linux or z/OS would be the Operating Systems in that Virtual Server.

Recall what you heard earlier: “The IBM zEnterprise™ System is a first-of-a-kind workload-optimized multiplatform (or multi-architecture) computing environment that spans (and tightly integrates) mainframe and distributed technologies. This system (of systems) consists of an IBM zEnterprise™ 196 (z196), the IBM® zEnterprise™ BladeCenter® Extension (zBX) Model 002, and the IBM zEnterprise™ Unified Resource Manager (Unified Resource Manager). The z196 is designed with improved scalability, performance, security, resiliency and availability. The Unified Resource Manager, working with the z196, the zBX infrastructure, and the attached blades, can help to deliver end-to-end virtualization and management providing the ability to align the technology deployment environment according to individual workload requirements.”

Ensemble Membership: Firmware and Software Configuration 1



The Unified Resource Manager performs the virtual server definitions up to, but not including, the Operating System tasks. Therefore, Operating System tasks like installation of the software, applying fixes or patches, or backing up and restoring are performed as usual. For example, you might connect to the customer’s external network to perform these tasks. After configuring the z196 as a node in an Ensemble, the “containers” for the Virtual Servers must be provisioned. Remember that the Virtual Server is wherever the Operating System is deployed. For example, an LPAR is a Virtual Server and z/OS might be deployed in it. Alternatively, Virtual Machines (or “Guests”) under z/VM could be Virtual Servers and Operating Systems might be deployed inside these Virtual Guests/Servers. The Virtual Server “container” provided by Ensemble management represents the configuration required to define the Virtual Server – including LPARs, z/VM guests, Operating Systems, Blade guests -- to the hypervisor. If the operating system definitions do not match the definitions of the Ensemble’s Virtual Server provisioning in the firmware, the Ensemble membership will not succeed. In z/VM the work at the HMC to define the Ensemble and its Virtual Servers provisions many of the definitions that are required to create Guests under VM, thus performing much of the work that a VM Systems Programmer would have to do in order to build a Virtual Guest. **Some steps in the firmware configuration are explicitly executed by the administrator. Some occur transparently: MAC prefix assignment, MAC address assignment, connection of the Virtual Network Interface Card to the virtual network, etc. Some Operating Systems and Hypervisors – like z/VM and POWER -- introduce more transparent provisioning of object definitions than others.** There are four components of the network virtualization of the IEDN:

VLAN (Virtual LAN) — A logical local area network that flows across the IEDN. A name and a numeric VLAN IDentifier define a VLAN.**the TOR enforces membership in the VLAN.**

VSwitch (Virtual switch) — A virtual switch is a hypervisor component that provides virtualized network resources to a virtual server. Conceptually a VSwitch is represented for a z/OS LPAR by the OSX port. In z/VM the Hypervisor is a software VSwitch.

VNIC (Virtual network interface card) — The Virtual Network Interface Card is the network resource that a virtual server uses to access the IEDN. The Virtual Network Interface Card is defined in the hypervisor through a VSwitch. (Again, conceptually the OSX port assumes the role of the VSwitch for z/OS.)

VMAC (Virtual media access control) — Virtual MAC addresses are assigned to Virtual Network Interface Cards. The VMAC replaces the manufacturer’s “burned-in” MAC address on a physical network card. A specific MAC address for an ensemble network resource is not assigned – only a specific MAC prefix is assigned. The enforcement points (i.e. OSX and z/VM) are each assigned a 3-byte MAC prefix from which a full 6-byte MAC address is created and assigned to QDIO connections. Access to the intra node management network is restricted to authorized management applications, and is only available through Port 0 of any OSA-Express3 CHPID configured with type OSM. Port 1 is not available for these communications. Connectivity to the intra node management network is restricted to stacks that are enabled for IPv6. Connectivity to the intra node management network and to the intra ensemble data network is allowed only when zEnterprise is part of an ensemble.


Dependencies: [This function is limited to OSA-Express3 ethernet features configured with CHPID types of OSX and OSM running on a zEnterprise. See the 2817DEVICE Preventive Service Planning (PSP) bucket for more information. This function is dependent upon the z/OS LPAR participating in a zEnterprise ensemble. See *System zEnterprise Ensemble Planning Guide* for more information.



Ensemble Membership: Firmware and Software Configuration 2



Software Configuration for Operating System in the Virtual Servers



1. Enable the Ensemble in the Operating Systems of the selected Virtual Servers and Define the Interfaces to the IEDN and its VLAN or VLANs.
 1. z/OS enabled for IPv6:
 1. VTAM® Start Option
 2. Interfaces to IEDN (OSX)
 1. VTAM TRLEs
 2. TCP/IP Interfaces
 2. z/VM enabled for IPv6:
 1. Provisioning z/VM and its Guests via the Firmware supplies the System Management API with information so that the VSwitch and Guest Configurations can be built dynamically!
 3. zTPF (as z/VM Guest on VSwitch) ¹
 4. zVSE™ (as z/VM Guest on VSwitch) ¹
 5. Linux on System z (as z/VM Guest on VSwitch) ¹
 6. Operating Systems in Virtual Servers on Blades of zBX

© 2012 SHARE and IBM Corporation

Z/OS and VTAM: A VTAM Start Option enables ensemble participation by the TCP/IP stacks running in z/OS. The Start Option is “ENSEMBLE=YES|NO.” The ENSEMBLE setting is used to either permit or deny connectivity to the intraensemble data network and the intranode management network by allowing or denying activation of OSX and OSM interfaces. When the zEnterprise node is configured at the HMC as a member of an ensemble, you can change the value of the ENSEMBLE start option with the MODIFY VTAMOPTS command while VTAM is running.

Result: The ENSEMBLE setting is used to either permit or deny activation attempts for OSX and OSM interfaces. Modifying the ENSEMBLE start option from YES to NO does not cause z/OS Communications Server to take action on any active OSX or OSM interfaces.

The other operating systems all have their own definition types to configure what needs to match in the Ensemble Configuration that was performed at the HMC with Unified Resource Management.

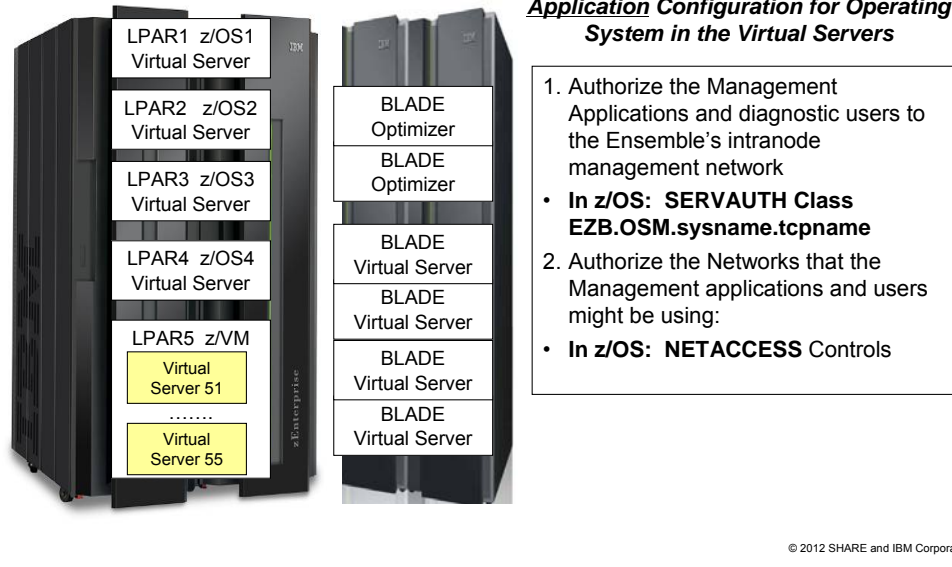
Access to the intra node management network is restricted to authorized management applications, and is only available through Port 0 of any OSA-Express3 CHPID configured with type OSM. Port 1 is not available for these communications.

- Connectivity to the intra node management network is restricted to stacks that are enabled for IPv6.
- Connectivity to the intra node management network and to the intra ensemble data network is allowed only when zEnterprise is part of an ensemble.
- This function is limited to OSA-Express3 ethernet features configured with CHPID types of OSX and OSM running on a zEnterprise. See the 2817DEVICE Preventive Service Planning (PSP) bucket for more information.
- This function is dependent upon the z/OS LPAR participating in a zEnterprise ensemble. See *System zEnterprise Ensemble Planning Guide* for more information.

Footnote 1: zTPF, zVSE, and Linux on z can all run native in an LPAR. However, in such a configuration they cannot participate in an Ensemble.



Ensemble Membership: Firmware and [Software Configuration 3](#)



Only authorized applications and users can use the INMN. We show you how to do this for z/OS. Other operating systems have their own definitions and controls to authorize access to the INMN.

The SERVAUTH class EZB.OSM.sysname.tcpname controls ability to access the intranode management network using OSM interfaces.

The NETACCESS controls are defined for z/OS with the SERVAUTH class **EZB.NETACCESS.sysname.tcpname.security_zonename**. This SERVAUTH class controls local user inbound and outbound access to network resources, and local user access to local IP address when explicitly binding to local interface (or using job-specific or destination-specific source IP addresses)

zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)


Unified Resource Manager and its Role in Security



Unified Resource Manager for Ensemble Management

Two Suites of Functionality:

1. Manage Suite
2. Automate Suite



- Integrate, monitor, and manage multi-OS resources as a single, logical virtualized system
- Consolidate multi-tier application policies to a single management interface
- Extends Mainframe Qualities of Service to all assets in a mission critical multi-tier application
- Delivered as integrated platform firmware
- Protected through standard HMC security and additional security controls

The ensemble is managed by a System z Hardware Management Console running firmware called the IBM zEnterprise™ Unified Resource Manager. The Unified Resource Manager provides energy monitoring and management, goal-oriented policy management, increased security, virtual networking, and data management for physical and logical resources of a given ensemble.

The Primary HMC for an Ensemble includes the Unified Resource Manager which provides two suites of functions:

1. Manage Suite (Manage). The first suite of functionality associated with the Unified Resource Manager. The Manage suite includes core operational controls, installation, configuration, Virtual Server provisioning and management, Virtual Network management, Hypervisor Management, Storage Virtualization Management, Energy Controls, Energy Monitoring, Monitors Dashboard, Default Workload Performance Context, Monitoring, and Reporting

2. Automate Suite (Automate). The second of two suites of functionality associated with the IBM zEnterprise™ Unified Resource Manager (Unified Resource Manager). The Automate suite includes goal-oriented Energy Management, Workload Performance Context, Monitoring, and Reporting, and Performance Management. The Manage suite must be installed in order to install the Automate suite.

Any HMC can manage up to 100 z196s. The primary HMC can perform all non-ensemble HMC functions on z196s that are not members of the ensemble (up to the 100 z196s limit which includes any ensemble members).

A primary HMC is the only HMC that can perform ensemble-related management tasks (create virtual server, manage virtual networks, create workload)

Any other HMC can only perform non-ensemble tasks on a z196 that is a member of an ensemble.

A customer can have multiple ensembles - they would need a Primary and an Alternate HMC pair for each ensemble they create.

Security in the HMC and the Unified Resource Manager:

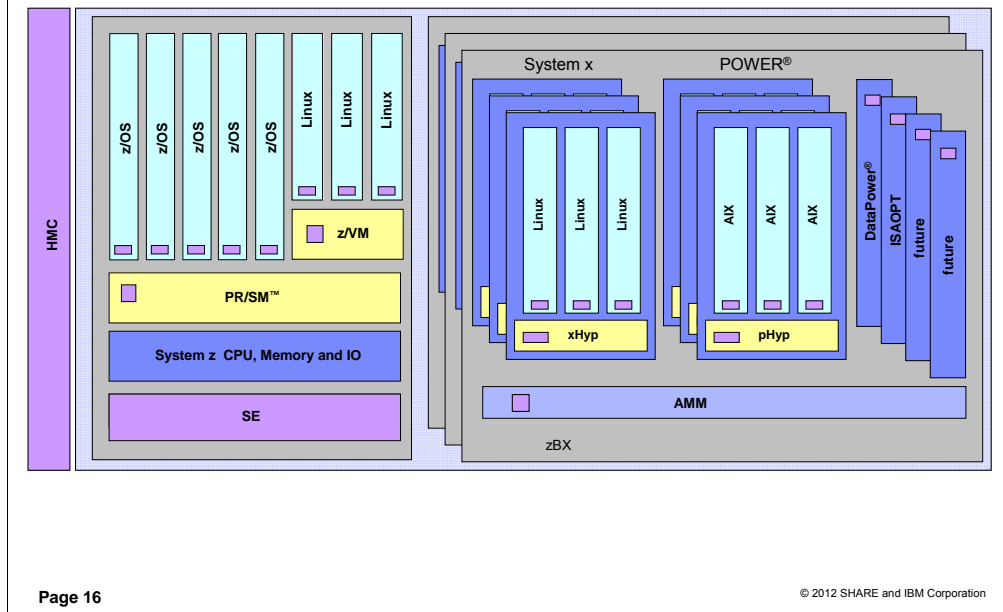
The HMC continues to enjoy the security controls inherent in the System z Console. Please see IBM System z Hardware Management Console Security White Paper by Kurt Schroeder (schroedk@us.ibm.com) published in Sept. 2008 and available at:

<http://nascpok.pok.ibm.com/rsf/zHMCSecurityWhitepaper.pdf>

This paper explains: "The HMC *Licensed Internal Code* includes a full-function firewall that is used to control network access to the HMC. As previously described, by default the HMC allows for virtually no inbound network traffic. As different features of the HMC are enabled (e.g., remote access, SNMP based automation etc.) additional inbound network traffic is allowed."

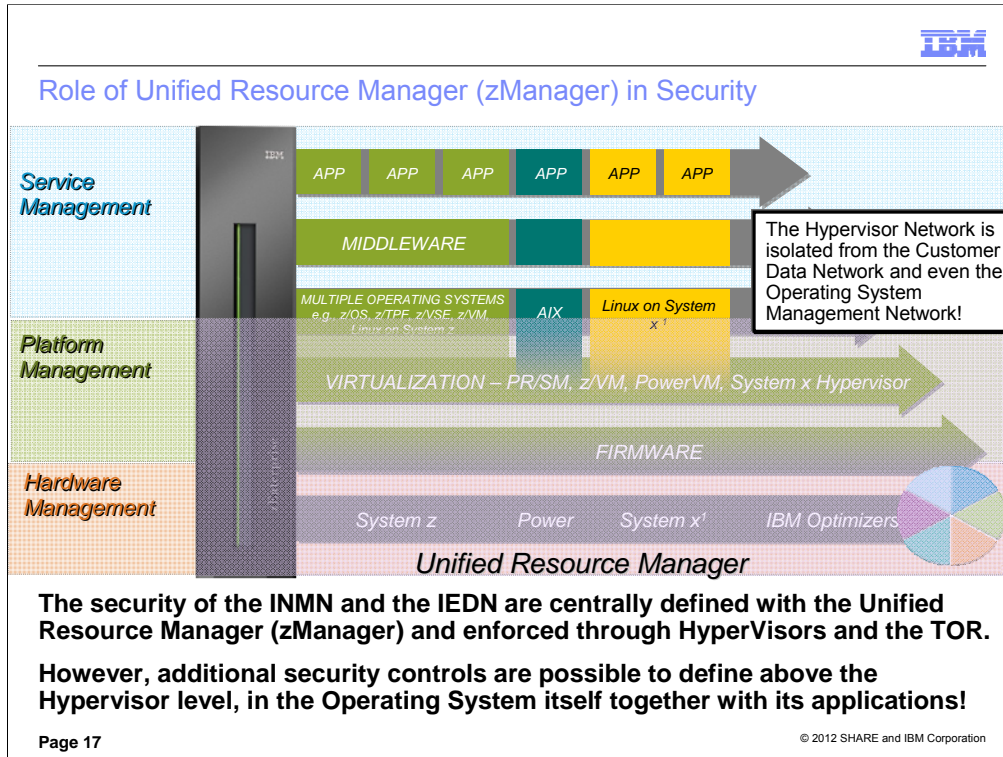
In addition, the Unified Resource Manager screens of the console enjoy additional security controls, like role-based security and operations over the closed, Intranode Management network (INMN).

Unified Resource Manager Controls: Firmware and Software



IBM-controlled Firmware and sometimes Software in the zEnterprise system work together to provide definition, access, and operational controls to reduce security exposures. The purple boxes illustrate various elements of Unified Resource Manager (aka zManager). While zManager starts at the HMC, you can see how that there are aspects of the new platform firmware management function spread across the entire zEnterprise environment. All virtual servers sit on top of their respective hypervisors (yellow boxes), and zManager controls all heterogeneous hypervisors within zEnterprise, thus controlling network access for all virtual servers.

Notice how the zManager definitions and controls assist the Advanced Management Module (AMM) in a Blade Center to autodetect Hardware and to detect errors. zManager also interfaces with the Support Element (SE), which controls the access to the OXM and OSX OSA ports. zManager influences Hypervisor definitions in System z, in z/VM, and in the blade operating system Hypervisors. Then each operating system may also software exploit agents to report on performance in the Ensemble.



SECURITY and zManager: Unified Resource Manager will orchestrate various forms of platform management and virtualization by interacting with various elements of platform firmware and hardware. **The security of the INMN and the IEDN is centrally defined with the Unified Resource Manager (zManager) and enforced through HyperVisors and the TOR.** All Virtual Servers and VLANs must pass through "access points" (Hypervisors and TORs) where their authorization is confirmed; the "access points" contain security enforcement that has been defined with Unified Resource Manager.

However, additional security controls are possible to define above the Hypervisor level, in the Operating System itself together with its applications! Understanding the level of security required and the isolation provided by the network virtualization management function of the Unified Resource Manager in collaboration with other firmware elements of the IBM zEnterprise System will help clients determine what, if any, additional security devices, like firewalls, are required in their enterprise solutions.

Built on this construct – zEnterprise – Innovation at every level

In the hardware management layer we have the System z server, and POWER based hardware, System x hardware and IBM appliances or optimizers (the specialty purpose processors versus the general purpose processors). Let's ask ourselves ... why should the hardware/firmware management be any different for each piece when they are all working together to provide the same business results? All have pieces in the same business problem.

Moving into the Platform management layer – In the System z environment today customers manage the System z as an integrated and unified system – with a single point of control for everything in the environment. But when you go into the distributed world you immediately get more risk. It's just inherently more risky because every component has its own service and change methodology – there is one website for this piece and another website for another piece. As a customer, no matter how much you have control over your system – each piece of the solution is handled on an individual basis – with individual change control policies and nobody .. except you (the customer) .. has executed them in any given sequence. Every time something is done the customer takes the risk of doing something that isn't compatible with the operating environment as a whole.

So this is what we want to help do with Unified Resource Manager. Take these many layers and extend them across the different architectures whether it's the hardware (System z, Power, System x) or the platform (the virtualization and how it's extended and managed).

IBM Systems Director and Unified Resource Manager both provide function at the hardware and platform management level. Use Unified Resource Manager for resources in zEnterprise.

Service Management – that's Tivoli's domain and this is fundamentally unchanged. And you need to understand that we are NOT talking about changes to the application environment. Anything working at the application level is unchanged.



Traditional Unified Resource Manager Security Services

▪The Unified Resource Manager security design is based upon:

– Centralized and Simplified Control at the HMC

- Limited to a Few Selected Administrators
- Access Controls for Authorized Administrators (Role-based Security)
 - Access Controls for Tasks and Resources
 - * VLANs
 - * Z196 / z114 LPARs
 - * zBX Blades
 - * Virtual Servers
 - * Policies
 - * Operations

– Security Enforcement by the Hypervisors and the Top-of-Rack Switches

- Access Controls for Ensemble Participation by Hybrid Resources:
 - Which Virtual Servers are authorized to which:
 - * VLANs
 - * Policies
 - Which Applications may exploit the INMN
- Access Controls for Ensemble Participation by External Network
 - Which External Equipment can attach to the IEDN?

The security of the Ensemble Network is provided primarily by the functions executed at the HMC, including those that reside in the Unified Resource Manager firmware. As long as a security implementation is managed through the Unified Resource Manager, it is part of the Ensemble “story.” The security is tightly controlled through a limited group of individuals with access to the definition panels of the HMC and the Unified Resource Manager.

However, there is more to a security architecture than what lies in the single point of control provided for the Ensemble by the HMC and the Unified Resource Manager. You may choose to implement additional security controls, but, in so doing, you run the risk of having to manage too many points of security control, at which point the Ensemble network can be exposed to additional security vulnerabilities.

In the next sections of this document we describe to you these Unified Resource Manager and Hypervisor controls.

In subsequent sections of this document, you see other types of security implementations that can be added to a running environment. You, the implementer, must then make sure that the business and IT processes you use to safeguard your applications, servers, network, and business are strong enough to mitigate the dangers of the widened span of control that you have introduced.

zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

Network Segmentation Requirements of Security Mandates



What is Network Segmentation? (Defined by Payment Card Industry)

- **Used to protect data and data flows**
 - Through physical or logical isolation or segregation of ...
 - storage areas
 - process areas
 - transmission paths
 - even of users who need to access the data and the data paths
- **In the PCI documentation, the wording relates to "Network Segmentation"**
 - "PCI DSS requirements apply to all system components. In the context of PCI DSS, "system components" are defined as any network component, server or application that is included in, or connected to, the cardholder data environment. System components" also include any virtualization components such as virtual machines, virtual switches/routers, virtual appliances, virtual applications/desktops, and hypervisors. The cardholder data environment is comprised of people, processes and technology that handle cardholder data or sensitive authentication data.
 - Network segmentation can be achieved through a number of physical or logical means, such as properly configured internal network firewalls, routers with strong access control lists, or other technologies that restrict access to a particular segment of a network.

PCI DSS2.0 refines the definition as follows:

Navigating the Intent of PCI: PCI DSS requirements apply to all system components. In the context of PCI DSS, "system components" are defined as any network component, server or application that is included in, or connected to, the cardholder data environment. System components" also include any virtualization components such as virtual machines, virtual switches/routers, virtual appliances, virtual applications/desktops, and hypervisors. The cardholder data environment is comprised of people, processes and technology that handle cardholder data or sensitive authentication data.

PCI DSS 2.0: Without adequate network segmentation (sometimes called a "flat network") the entire network is in scope of the PCI DSS assessment. Network segmentation can be achieved through a number of physical or logical means, such as properly configured internal network firewalls, routers with strong access control lists, or other technologies that restrict access to a particular segment of a network.

At a high level, adequate network segmentation isolates systems that store, process, or transmit cardholder data from those that do not. However, the adequacy of a specific implementation of network segmentation is highly variable and dependent upon a number of factors, such as a given network's configuration, the technologies deployed, and other controls that may be implemented.

Selected Examples of Network Segmentation Techniques

- **Physically Separated (Sub)nets and Address Groups** ❖ Exploited in IEDN architectural design
 - **Physically Separated Network Adapters (OSA Ports)**
 - ❖ **Logically Separated Network Adapters (OSA Ports)** ➤ Possibly unnecessary with IEDN
 - Virtualized into multiple ports with Separate VMACs and VLANs
 - **Networks that are logically (virtually) isolated by Encryption or by Filtering with Firewalls, Access Controls, Virtual LANs (VLANs), etc.**
 - ❖ **Operating Systems Separated through Hypervisor Techniques**
 - LPARs in z/OS (PR/SM) or LPARs on P-Blades (Power/VM)
 - Guests under z/VM
 - ❖ **“Network in a Box” with HiperSockets Convergence (z/OS) and zVM HiperSockets Completion Queue and VM Bridge**
 - ❖ **Concealed Networks in the IEDN and the Intranet**
 - IPv6: Link-Local Addresses (non-routable across IP nodes)
 - IPv6: Unique Local Unicast Addresses (non-routable outside the Intranet)
 - IPv4: Private Network Address Ranges (non-routable outside the Intranet)
 - **Concealed Networks with Dynamic Routing Protocols (Dynamic Routing Protocols Not recommended for the IEDN)**
 - OSPF: concealing a range behind an Area Border router
 - OSPF: "Ignore undefined interfaces" in order to restrict link state / route advertisements
- Possible, but not recommended in IEDN due to its Hypervisor Security Enforcement. Nevertheless, it can be implemented with careful planning.

Notice the different formats of the bullets. Each one indicates the statement’s applicability to the Ensemble environment.

IMPORTANT: Although it is possible to exploit dynamic routing protocols within the IEDN, IBM discourages you from using such protocols because the security of the Ensemble VLAN implementation requires very careful planning. That is, an invalid dynamic routing design can lead to inconsistent topology databases and undesirable routing tables due to the VLAN ID enforcement in the Hypervisors.

Notes: Virtualization and Security (PCI DSS 2.0)

- If virtualization is implemented, all components within the virtual environment will need to be identified and considered in scope for the review, including the individual virtual hosts or devices, guest machines, applications, management interfaces, central management consoles, hypervisors, etc.
- All intra-host communications and data flows must be identified and documented, as well as those between the virtual component and other system components.
- The implementation of a virtualized environment must meet the intent of all requirements, such that the virtualized systems can effectively be regarded as separate hardware. For example, there must be a clear segmentation of functions and segregation of networks with different security levels; segmentation should prevent the sharing of production and test/development environments; the virtual configuration must be secured such that vulnerabilities in one function cannot impact the security of other functions; and attached devices, such as USB/serial devices, should not be accessible by all virtual instances.
- Additionally, all virtual management interface protocols should be included in system documentation, and roles and permissions should be defined for managing virtual networks and virtual system components. Virtualization platforms must have the ability to enforce separation of duties and least privilege, to separate virtual network management from virtual server management.
- Special care is also needed when implementing authentication controls to ensure that users authenticate to the proper virtual system components, and distinguish between the guest VMs (virtual machines) and the hypervisor.


zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

**Network Virtualization Management:
Secure Networking with VLANs in the
Ensemble**




What is a Virtual LAN (VLAN)? Traffic Isolation on a Segment

- A single Physical LAN network segment should have only one IP Subnet assigned to it.



- A single Physical LAN network segment can be segmented into multiple virtual LAN segments by exploiting VLANs with multiple IP Subnets.



A local area network (LAN) is a broadcast domain. Nodes on a LAN can communicate with each other without a router, and nodes on different LANs need a router to communicate. A virtual LAN (VLAN) is a configured logical grouping of nodes using switches. Nodes on a VLAN can communicate with each other as if they were on the same LAN, and nodes on different VLANs need a router to communicate. (Layer 3 routers can add, remove, or validate VLAN tags.) The IBM Open Systems Adapter provides support for IEEE standards 802.1q, which describes VLAN identifier tagging. **(Note that currently the OSX implementation supports 802.1q only.)** Deploying VLAN IDs allows a physical LAN to be partitioned or subdivided into discrete virtual LANs. This support is provided by the z/OS TCP/IP stack and the OSA-Express feature in QDIO mode. When you use VLAN IDs, the z/OS TCP/IP stack can have multiple connections to the same OSA-Express feature. One connection is allowed for each unique combination of VLAN ID and IP version (IPv4 or IPv6).

Note in the top half of the visual how one network takes advantage of the physical connectivity. In the bottom half of the visual, we have split the physical LAN into two VLANs: one with VLAN ID of 10 and another with VLAN ID of 11.


Z/OS and z/VM are implemented with a VLAN technology called “Global VLAN.” With Global VLAN, z/OS and z/VM can define a VLAN ID which is then registered in the OSA port. The OSA port then performs the VLAN tagging. The implementation of Global VLAN causes the stacks to be technically unaware of the VLAN, or “vlan-unaware.” However, many people find this subtle distinction confusing and refer to z/OS and z/VM as “vlan-aware” stacks since they can define a VLAN ID for a LAN connection. For Linux on z (native), the TCP/IP stack itself performs the VLAN tagging, and, thus, Linux on z when running native is not using the Global VLAN ID but rather the standard 802.1q implementation of VLAN. Linux on z is thus technically a “vlan-aware” stack. You may read about Global VLAN IDs at: <http://publib.boulder.ibm.com/infocenter/zvm/v5r3/index.jsp?topic=/com.ibm.zvm.v53.hcpa6/hcsc9b2131.htm>

“A GLOBAL VLAN ID is OSA’s VLAN support to provide access to a virtual LAN segment for a VLAN unaware host so the host can receive and send its network traffic. This host does not tag its outbound frames nor receive tagged inbound frames. The GLOBAL VLAN ID participates on the VLAN transparently with OSA handling all the tagging work (VLAN-unaware). A host device driver can register a Global VLAN ID with the OSA-Express adapter. Typically each host defines only one Global VLAN ID per connection. Some device drivers allow configuration of one VLAN ID for IPv4 and a second VLAN ID for IPv6. The OSA-Express will use the Global VLAN ID to tag frames and send out Gratuitous ARP requests (ARP requests to check for duplicate IP addresses) on behalf of the host. The NIC simulation in z/VM also provides this support, which is separate from the virtual switch support. The Global VLAN ID processing for the virtual NIC is performed prior to any virtual switch port ingress processing and after virtual switch port egress processing.

One example of this is in z/VM. You can specify the VLAN keyword on a LINK configuration statement for a QDIOETHERNET link to register a Global VLAN ID. In the past with OSD and the ability to code VLAN IDs on a Virtual Switch, the following recommendation has been made: “To reduce complexity and host TCP/IP configuration changes when configuring a virtual switch host connection, it is recommended that you do not configure a global VLAN ID for a host that will be connected to a trunk port. Instead, connect the host to an access port and authorize it for the desired VLAN ID. This assigns a port VLAN ID (pvid) for the access port and all VLAN operations occur within the virtual switch.” With Ensemble Networking, the z/VM VSwitch is the enforcement point for VLANs that have been assigned with the HMC through the Network Virtualization Management functions; in addition the TOR is coded in TRUNK Mode and not ACCESS Mode. Therefore, the z/VM guests that are capable of doing so should assign the VLAN ID that the HMC recognizes; then the VSwitch enforces the activation of that connection to the VSwitch using the appropriate VLAN ID.


Can Traffic “Hop” from one VLAN to Another? No, unless ...

- A single VLAN Can Span Switches:



- Layer 2 Forwarding with a Switch – No Router Is Required to forward packets.

- One VLAN can be interconnected to another if a Router routes between the two VLAN IDs:



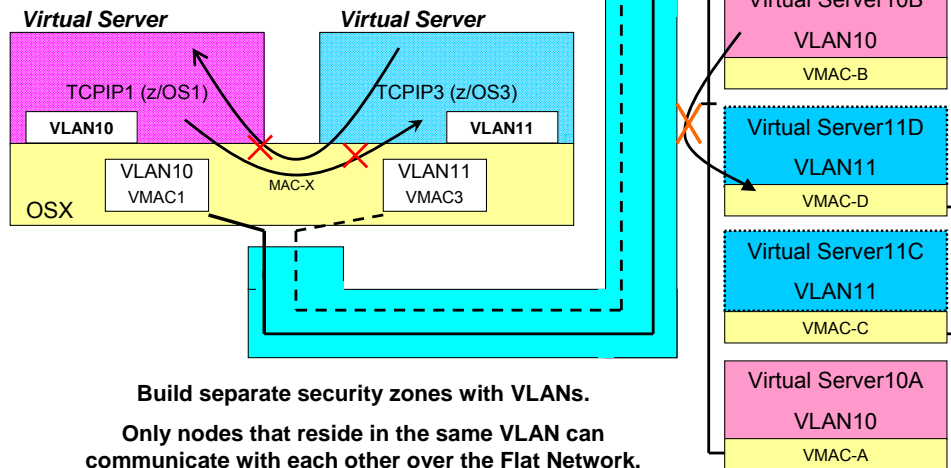
- Layer 3 Routing for forwarding packets between VLAN 10 and VLAN 11.

A VLAN may have multiple physical hops that are interconnected using a SWITCH. The multiple physical segments of the VLAN can even be associated with completely different IP addresses but the SWITCH can connect the segments to provide a single path through the SWITCH. A SWITCH operates at what is called “LAYER 2” of the IP protocol stack, and examines the contents of the Frame Header to determine where to forward packets across physical segments in the switch.

If a connection path contains more than one VLAN ID, as you see in the second diagram above, you must use a ROUTER or a combination ROUTER/SWITCH to interconnect the parts of the entire path. A ROUTER operates at what is called “LAYER3” of the IP protocol stack, and examines the contents of the IP Header to determine where to route (forward) the packets across physical segments of the path.

Different VLAN IDs over a Shared IEDN in the Ensemble: Security

Extra Security: VLAN ID Enforcement takes place at the TOR and Hypervisor: PR/SM™, z/VM, VSwitch, Blade Hypervisor, OSX (VSwitch).



VLANs are a means of building distinct security zones across the same physical network segment. Over the years there have been concerns in the industry about VLANs in terms of security, because both the MAC address and the VLAN ID can be spoofed or altered. However, in an IEDN these types of issues do not apply, because we use the Network Virtualization Management (NVM) features to control all access to the LAN via our hypervisors – including the OSX implementation as a Hypervisor VSwitch. The VLANs and the VMACs are controlled through NVM assignments. (With OSX an Operating System cannot build its own Logical Link Control – “LLC” – header.) The only exception to the security with a VLAN implementation is when an outside network connects to the IEDN via the external ports on a TOR; for this exception, the customer is responsible for securing the access.

Ensemble members that reside in a zBX may not be sharing ports at all. Even in such a situation, the VLAN ID prevents sending a message intended for one member to another. In this visual you see that Virtual Servers z/OS1 with TCPIP1 and z/OS3 with TCPIP3 cannot communicate with each other over the OSA port because they are attached to different VLANs. However, Virtual Server TCPIP1 can communicate via the OSX OSA port to the Virtual Servers in the zBX that are also attached to the same VLAN ID of 10 (Virtual Servers 10B and 10A). Likewise, Virtual Server z/OS 3 with TCPIP3 on VLAN 11 can communicate with the Virtual Servers in the zBX that are also attached to VLAN ID of 11 (Virtual Servers 11C and 11D).

Virtual Server Operating Systems should not forward from one VLAN ID to another on an OSX CHPID via Layer 3 routing. In fact, on z/OS IP Forwarding has been disabled when traffic comes in on an OSX port and wishes to exit on another (or the same) OSX port. Both z/OS and z/VM can implement the ISOLATE function for an OSD or OSX interface; although ISOLATE is not necessary to isolate traffic over a Shared OSA implementation when one has deployed different VLAN IDs (as in the diagram above), you can still code ISOLATE on the Interface definition to prevent communication over a shared OSA port.

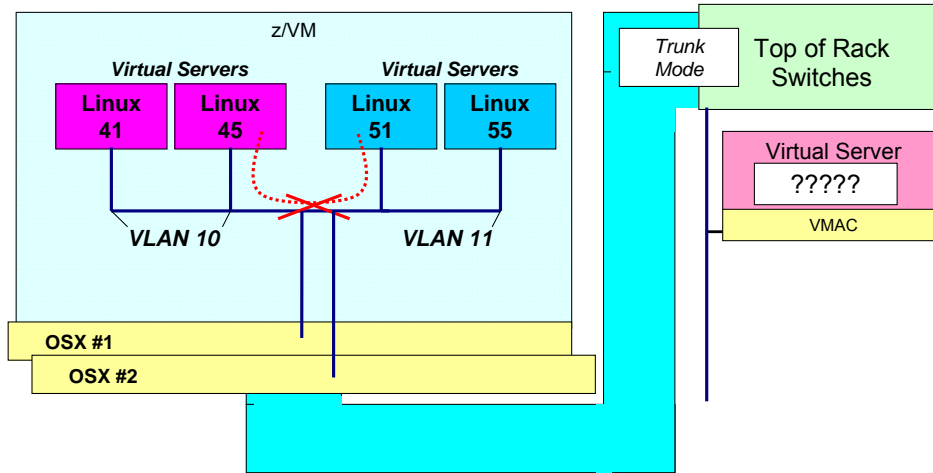
When anything is connected to the IEDN, a VLAN must be used. When the traffic hits the TOR port, the sending adapter (server) must have applied a VLAN ID tag. Therefore, for z/OS it will be tagged by the OSA at z/OS, and for an external router that wishes to connect to the TOR, the traffic must be tagged by that router.

VLAN ID enforcement adds a layer of security; it takes place at the Hypervisor.

hypervisor. A program that allows multiple instances of operating systems or virtual servers to run simultaneously on the same hardware device. A hypervisor can run directly on the hardware, can run within an operating system, or can be imbedded in platform firmware. Examples of hypervisors include PR/SM, z/VM, and PowerVM." In fact, even the OSX is considered a Hypervisor VSwitch, or a type of PR/SM Hypervisor. z/VM is also a Hypervisor for a direct attachment to the OSX, where zVM itself is considered a type of PR/SM Hypervisor.

In the scenario depicted, the TOR IEDN Ports to the LPARs would be configured in TRUNK Mode because the Operating Systems are defining their VLAN IDs (are “VLAN-aware”). If z/VM is a host for a Virtual Server whose Operating System is not VLAN-aware, the z/VM VSwitch – which is a Hypervisor – handles the VLAN ID on behalf of the Virtual Server. TOR ports to the Virtual Servers on the BLADES are also configured in TRUNK mode. The TOR ports to any ISAOPT blades are configured in ACCESS Mode. **The hypervisor -- whichever it may be (OSX for native z/OS or z/VM, VSwitch for Virtual Servers on a VM VSwitch, or the Hypervisor on the blade) -- performs the enforcement. The TOR performs the enforcement for anything that does not fall into these categories just mentioned.** For example, if a native LPAR is communicating with a virtual server on a blade, the OSX will perform VLAN enforcement; the data passes through the TOR and the Ethernet Switch Module (ESM) without checking, but it is then checked again at the blade's VSwitch (=hypervisor).


Secure Segregation of VLAN 10 from VLAN11 (IEDN): 1 VSwitch



- If there is a z/VM VSwitch as depicted, the VSwitch, a Hypervisor component, prevents communication between the VLANs.
 - Layer 2 protocols cannot route between different VLANs; a layer 3 router is required to route between different VLANs.

There is VLAN ID enforcement on the z/VM VSwitch

Secure Segregation of VLAN 10 from VLAN11 (IEDN): 2 VSwitches



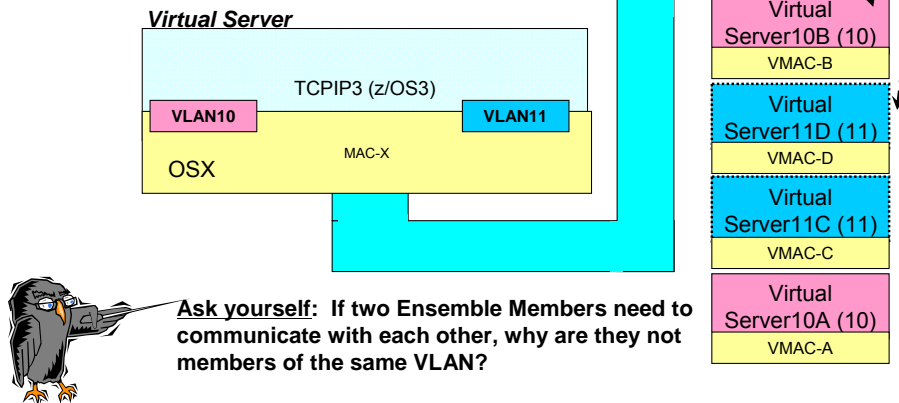
- If there are two z/VM VSwitches as depicted, the OSX VSwitch prevents communication between the VLANs.
 - Layer 2 protocols cannot route between different VLANs; a layer 3 router is required to route between different VLANs.

There is still VLAN ID enforcement on each individual z/VM VSwitch. But now you see that the OSX VSwitch will also not allow forwarding between two VLANs when the VLANs are on separate VSwitches.

What if I Must Route between VLANs of the IEDN? Example 1

Extra Security: You may route between VLANs through a non-IEDN router if necessary, but use a Firewall to maintain secure environment.

Example: non-IEDN router



Ask yourself: If two Ensemble Members need to communicate with each other, why are they not members of the same VLAN?

When anything is connected to the IEDN, a VLAN must be used. When the traffic hits the TOR port, the sending adapter (server) must have applied a VLAN ID tag. Therefore, for z/OS it will be tagged by the OSA at z/OS, and for an external router that wishes to connect to the TOR, the traffic must be tagged by that router.

Virtual Server Operating Systems should not forward from one VLAN ID to another on an OSX CHPID via Layer 3 routing. In fact, on z/OS IP Forwarding has been disabled when traffic comes in on an OSX port and wishes to exit on another OSX port.

If you find it necessary to route between VLAN IDs by exiting to the external network to allow a router to route via Layer 3 back into the IEDN and another VLAN, then you will want to implement a Firewall to keep the traffic secure. Many security mandates require Stateful Firewalls, and most external Firewalls with IP Filtering do implement stateful packet inspection.

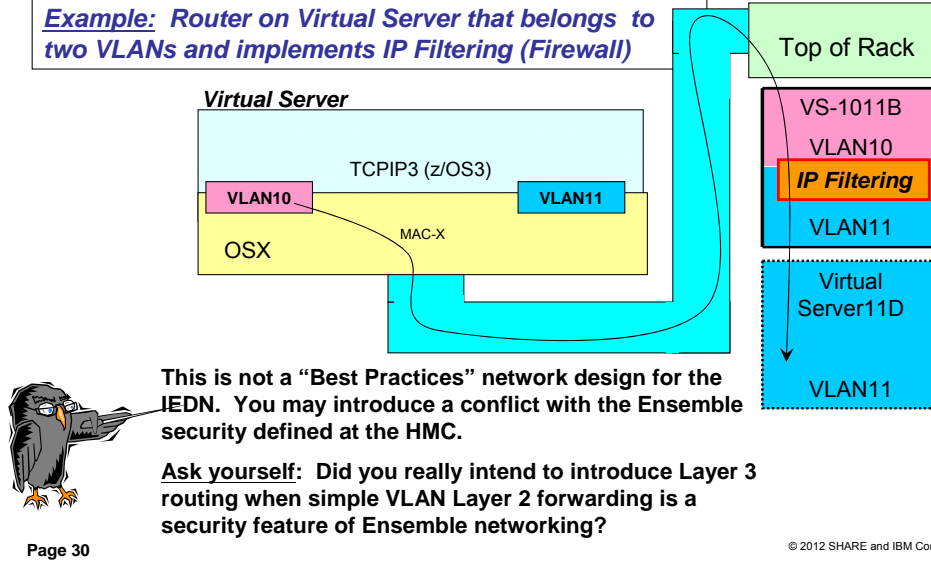
You must also implement static routes to force Virtual Server11D to route to the external router depicted; the external Layer 3 router then routes back into the zBX to create the connection to Virtual Server10B. You must use static routing to accomplish this, because dynamic routing will have built the route between the two servers unless the Virtual Server implementation has blocked this internal routing over the IEDN as z/OS has done. The router must have tagged packets with the VLAN ID that is to be used within the IEDN. (That is, in our scenario here the router will have tagged the traffic to the IEDN's VLAN10 with VLAN10 and the traffic to the IEDN's VLAN11 with VLAN11.)

IMPORTANT: Although it is possible to exploit dynamic routing protocols within the IEDN, IBM discourages you from using such protocols because the security of the Ensemble VLAN implementation requires very careful planning. That is, an invalid dynamic routing design can lead to inconsistent topology databases and undesirable routing tables due to the VLAN ID enforcement in the Hypervisors.

What if I Must Route between VLAN IDs of the IEDN? Example 2

Extra Security: You may route between VLAN IDs through an IEDN virtual server *if necessary*.

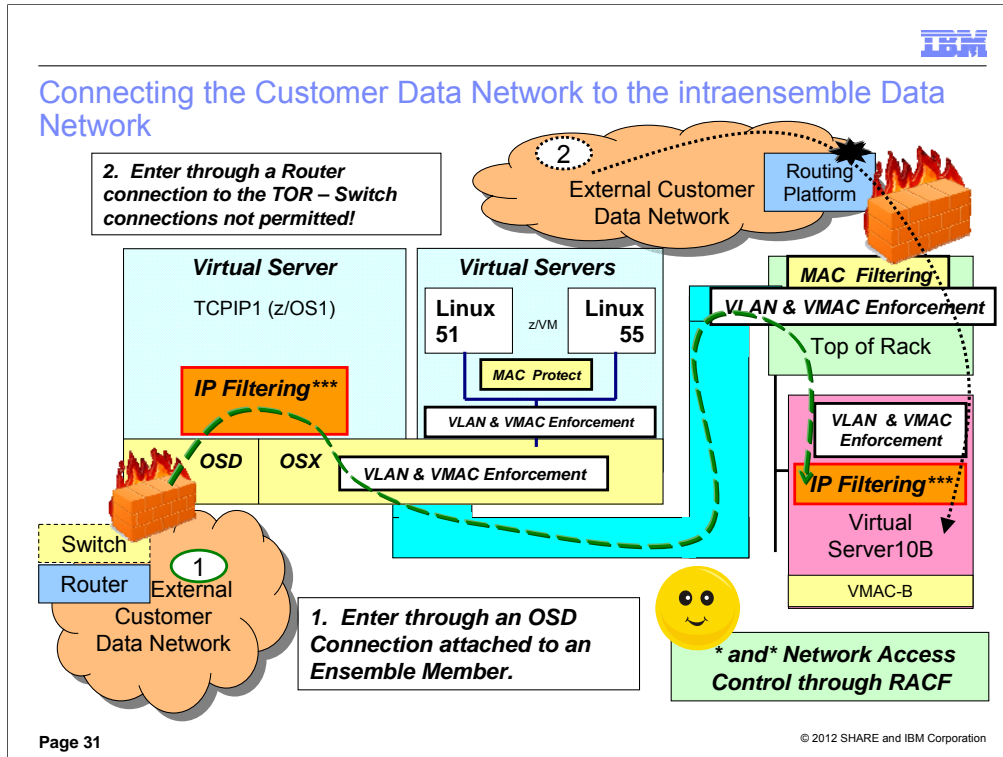
Example: Router on Virtual Server that belongs to two VLANs and implements IP Filtering (Firewall)



Virtual Server Operating Systems should not forward from one VLAN ID to another within the IEDN via Layer 3 routing since this can conflict with the security that the administrator handling the Network Virtualization Manager tasks has established. By setting up routing within a Virtual Server you now have another administrator potentially interfering with the security that was defined at the HMC. In fact, you have seen that on z/OS IP Forwarding has been disabled when traffic comes in on an OSX port and wishes to exit on the same or a different OSX port.

If the implementer of the Virtual Server inside the zBX decides to permit routing through that virtual server, then this visual illustrates how the packets would flow: from VLAN10 at TCPIP3, over the OSX port, through the IEDN to the TOR, and then to VS-1011B on VLAN10; VS-1011B inspects the packet with IP Filtering or an internal Firewall and if allowed, routes the packet to VLAN11 and out over the IEDN to the Virtual Server11D which is connected to VLAN11. Remember that this routing should be permitted ONLY IF IP FILTERING has been implemented to permit the traffic to flow as depicted. Bear in mind that currently there are no stateful IP Filtering or Firewall mechanisms that can run on a blade or an LPAR.

Connecting the Customer Data Network to the Intraensemble Data Network



If you decide to permit communication between the External Customer Data Network and the Ensemble, you can keep this path secure. **First**, determine whether you want to do this. **Second**, understand that the Ensemble contains its own enforcement points within the TOR (for external connections on the egress ports) and also within the Hypervisors. All Virtual Servers and VLANs must pass through the enforcement points -- "access points" (Hypervisors and TORs) -- where their authorization is confirmed; the "access points" contain security enforcement that has been defined with Unified Resource Manager. However, you may continue to implement other security services within the Ensemble by exploiting traditional security mechanisms: IP Filtering (a firewall function), encryption, access control lists, userid and password authentication, etc. **Note that any security implementation that falls outside the strict controls of Unified Resource Manager can be a security weakness if the implementers are not tightly controlled and if there are too many diverse implementers who are not subject to the centralized controls of Unified Resource Manager.** **Third**, determine how you will secure the external connections. Be aware of the fact that the TOR performs VLAN ID enforcement for connections to servers outside the zBX that are not attached to an OSX OSA port. If an ISAOPT appliance is on the zBX, then the TOR also performs the VLAN ID enforcement. (You configure the authorized VLAN IDs at the HMC as part of the Network Virtualization infrastructure.) Note that for security purposes a Layer 2 connection from the external network into the TOR is not supported -- the connection must be using Layer 3 protocols. (Bridge Protocol Data Units -- BPDUs -- at Layer 2 cannot be successfully exchanged between an external Layer 2 switch and the IEDN TOR. Spanning Tree Protocol (STP) messages that might be received from external switches are filtered out at the TOR. This together with an external Firewall is to protect the security of the IEDN network by avoiding VLAN ID collisions. For example, if a customer were to attach an external switch to the TOR, and BPDUs and STP messages were not being filtered out, a customer's external VLAN ID might be the same VLAN ID used within the IEDN and thus mistakenly cause the interconnection of external VLAN segments to the IEDN VLAN segments, thus impinging the security of the IEDN.)

More Background on Security Services and Mechanisms: The VLAN ID enforcement for any Virtual Server attached to an OSX works as follows: If a z/OS Native LPAR is on the OSX, then the OSX performs the VLAN ID enforcement; if the Virtual Server is under z/VM and attached to a VSwitch, then the VSwitch performs the VLAN ID enforcement. The hypervisors on the Blade of the Virtual Servers of the zBX perform VLAN ID Enforcement in the zBX.

Remember that Security protection is much more than just inserting a firewall or IP Filtering along a path. It encompasses all layers of the IP Stack: Application Security Mechanisms (Access Control Lists, Userid and Password checking, mapping mechanisms), Transport Security Mechanisms (SSL/TLS, AT-TLS), IP Layer Security Mechanisms (IPSec, IP Filtering, Intrusion Detection Services, Network Address Tables), Data Link Control Security Mechanisms (MAC Address Filtering, VLAN Segmentation or Segregation), and many more mechanisms too numerous to mention here.

With regard to MAC Filtering, the zManager/HMC can define MAC filtering for external MAC Addresses and must define permitted VLAN IDs on external connections. (These are the connections established on the TOR's "egress" ports to and from the external network.) The TOR then enforces these VLAN IDs and filtered MACs or VMACs. The MACs within the IEDN are managed by the Network Virtualization Manager. All MACs are allowed that originate from within the IEDN (they are managed by NVM). z/VM VSwitch MAC Protect function for Layer 2 is on by default for IEDN type VSwitches. This MAC Protect function enforces that a VMAC sent during guest link initialization (SETVMAC) matches with what has been assigned by the z/VM hypervisor. In addition, all SOURCE MAC addresses on egress frames from the guest are verified to insure that only the assigned VMAC for the guest is being sent on outbound data transfers. This eliminates any attempt by the guest to spoof its source MAC address.

Note on use of VMACs: When an Operating System on z is using layer 2, it performs an ARP and builds Ethernet headers with VMAC. z/OS does not use layer 2 and so ARP is handled by the OSA and the OSA builds the Ethernet header. On a VSwitch under z/VM, you might have a Linux guest using Layer 2. In this case the VSwitch builds the Frame Header with the VMAC in it. But, if the guest system is using Layer 3 on a VSwitch, then a frame header is not necessary and the packet is forwarded over the VSwitch using only the IP address.

If you are still interested in introducing firewalls consider these possibilities: Firewalls: IP Filtering in z/OS Policy Agent (a "host-based firewall" that is not stateful); IP Filtering with Proventia Intrusion Prevention Services for Linux on z (a "host-based firewall" that is not stateful); IP Filtering in Virtual Servers residing on the zBX (may or may not be stateful); Other IP Filtering mechanisms (could be stateful or not); Firewall in front of LPAR that is attached to an OSD OSA (External firewalls are usually "appliance-based" firewalls that -- unlike "host-based firewalls" -- are stateful.); Firewall in front of TOR in External network (External firewalls are usually "appliance-based" firewalls that -- unlike "host-based firewalls" -- are stateful.); SERVAUTH Classes: NETACCESS CONTROLS for IEDN; MAC ADDRESS FILTERING at the TOR; MultiLevel Security

A Closer Look at the IEDN TOR and its Secure Connections to the External Customer Data Network

2. Enter through a Router connection to the TOR – Switch connections not permitted!

❖ Selected IEDN VLANs terminate at the external, Layer 3 Routing platform.

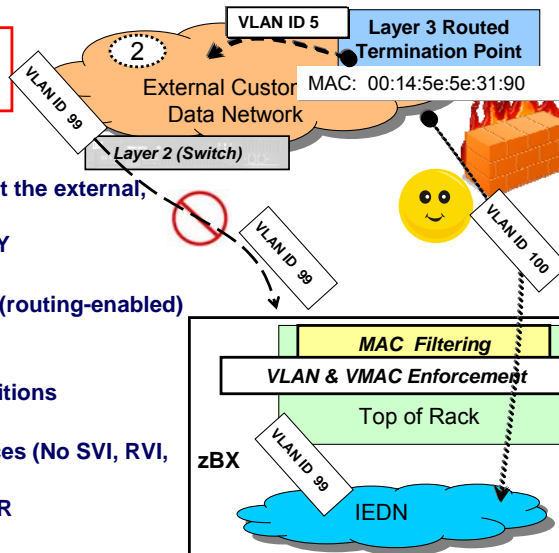
• Routed termination points ONLY

- ✓ Dedicated Router Platform
- ✓ Operating System Platform (routing-enabled)
- ✓ L2/L3 Switch with
 - Routed Interface or
 - Sub-interface definitions



• No external Layer 2 Switch!

- ✓ No Switched Virtual Interfaces (No SVI, RVI, VRI)
- ✓ No Layer 2 Messages to TOR
- ✓ No STP messages
- ✓ No BPDUs



When we talk about the IEDN, we point out that it is a FLAT network, using Layer 2 Virtual LANs to forward traffic from one end of this FLAT network to the other end. We also point out how this "one-hop" configuration reduces the complexity of the network design by eliminating network equipment (routers, cables, administration, and so on). Finally, we emphasize how the reduced complexity of the network design leads to the elimination of the typical security vulnerabilities of a multi-hop network and the reduction of network latencies.

• Integrity of the IEDN is preserved by preventing VLAN ID collisions between the external Customer Data Network and the IEDN!

Certain VLANs in the IEDN can extend one hop to the external Layer 3 routing platform. (See VLAN ID 100 in the visual.) A "routing platform" can be a dedicated router platform, or a platform running routing in the installed operating system, or even a Layer2/Layer 3 switch that has been implemented only with Router Interfaces or Subinterfaces. (The Switched Virtual Interface definition available in some platforms is still switched and therefore not allowed.) Keep in mind: IEDN VLANs cannot extend into the external customer network by connecting through an external switch. (See VLAN ID 99 in the visual.)

In the visual you see that -- for security purposes -- a Layer 2 connection from the external network into the TOR is not supported -- the connection must be using Layer 3 protocols. (Bridge Protocol Data Units -- BPDUs -- at Layer 2 cannot be successfully exchanged between an external Layer 2 switch and the IEDN TOR. Spanning Tree Protocol (STP) messages that might be received from external switches are filtered out at the TOR. This together with an external Firewall is to protect the security of the IEDN network by avoiding VLAN ID collisions. For example, if a customer were to attach an external switch to the TOR, and BPDUs and STP messages were not being filtered out, a customer's external VLAN ID might be the same VLAN ID used within the IEDN and thus mistakenly cause the interconnection of external VLAN segments to the IEDN VLAN segments, thus impinging the security of the IEDN.)

IMPORTANT Reminder: The TOR is not just any switch .. the switch ensures that possibilities for LAN collisions and for misconfiguration do not impinge on the security of the network because certain standard switch functions like the exchange of Layer 2 messages have been disabled. -- Which is another reason why the Juniper switch itself cannot be swapped out for a different switch -- Unified Resource Manager integrates with the Juniper switch so as to eliminate these Layer 2 security exposures and to provide a simplified configuration interface that is independent of the platform- and vendor-unique Graphical User Interfaces with which an administrator would normally have to deal. As a result, it is not important that an administrator be familiar with the configuration syntax of a particular switch brand. Due to this simplified GUI and its integration into the zBX, the TORs require very little configuration -- many of the functions are fixed and relieve the Ensemble administrators of typical switch tasks. The only configuration necessary is for securing the attachment to the external network through access control lists to VLAN IDs and to Virtual MACs.

MAC Filtering for External Connections into the TOR

SCZHMCA: Configure Top-of-rack (TOR) Switch - Mozilla Firef...

https://sczhmca.itso.ibm.com/hmc/content?taskId=88&refresh=211

Configure Top-of-rack (TOR) Switch - SCZP301

Switch Port:

Select	Port	Type	VLAN Mode	Allowed Virtual Networks
<input checked="" type="radio"/>	38	External	Trunk	199
<input type="radio"/>	39	External	Trunk	
<input type="radio"/>	8	Internal	Trunk	

VLAN Settings:

Allow all VLAN IDs

VLAN Mode:

Allowed Virtual Networks:

Select	Virtual Network
<input type="checkbox"/>	1034 - External-VSAdmin1034-172.30.10.0
<input type="checkbox"/>	10 - Default
<input type="checkbox"/>	110 - Internal-VLAN110-172.30.110.0

MAC Address Filtering:

Allow all MAC addresses

MAC Address:

Allowed MAC Addresses:

Example:
00:11:22:33:44:55

Done

Page 33

© 2012 SHARE and IBM Corporation

On this zManager screen, which shows the configuration for the TOR Switch, we have selected External TOR Port #38. Currently we have allowed only VLAN ID of 199 to operate on this Port. We are going to add other VLAN IDs later, and so we have set the VLAN Mode to TRUNK. We can begin other VLAN IDs that are acceptable from the outside network to this port. And we can also choose to apply MAC Filtering to ensure that the external device that is connecting is the one whose MAC address matches what we have established on this panel. We do not entirely trust the external network that is connected to this port so we will not select "Allow all MAC addresses." But we will filter the MAC addresses that are permitted by "Add"ing a specific MAC address as shown in the figure.

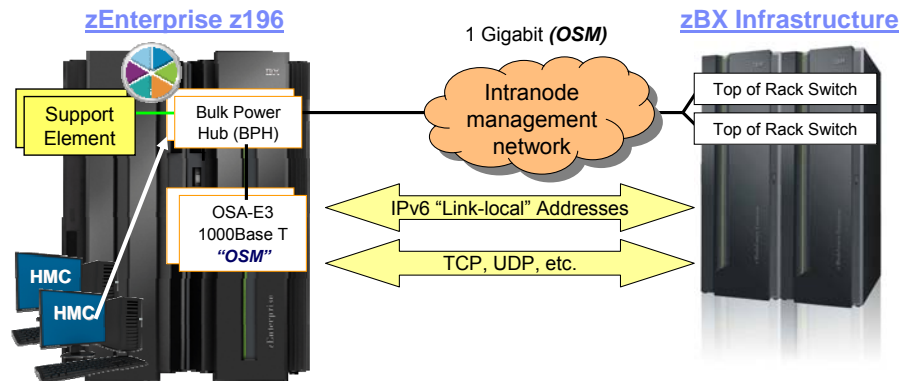
zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

Security: Access and Routing Controls in the Ensemble Networks



As you know, zManager controls stop at the operating system level. This means that there are all the other layers of security from the security framework that software at the operating system level can enforce. These are additional security implementations that reside in RACF, in networking protocols, and in software that can be added to an Ensemble implementation to protect the Ensemble network.

The intranode Management Network: A Closed, Secure Network



- **Intranode management network (INMN)**
 - 1000Base-T OSA-Express3 (copper) — QDIO (*CHPID Type OSM*) – Cables from OSM to BPH are 3.2 meters long; from BPH to 1Gig TOR 26 meters long
 - 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.

For reliability in an Ensemble, you must order redundant switches and redundant OSA-E3 cards to attach to the switches in order to interconnect the members of an Ensemble. You must also provide definitions that secure the access to the two networks depicted: the intranode management network (INMN) and the intraensemble data network (IEDN).

The INMN is used for firmware management (platform management). It is not an application network at all and no application software uses this network. Once it is defined, it is essentially invisible to the users and the Virtual Servers; the Unified Resource Manager performs all the definition that is imbedded in the firmware; although there is an underlying VLAN, it is invisible and is not even defined through Unified Resource Manager panels. Virtual Servers communicate only to the SE over this network. One Virtual Server cannot even ping another Virtual Server over this network since each Virtual Server is isolated from the others. Besides the SE, only management applications (existing management applications and future ones when they become available) will be able to communicate with the Virtual Servers. (Management applications are called Guest Performance Agents and there are several others types of management applications as well.)

For each z196 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.

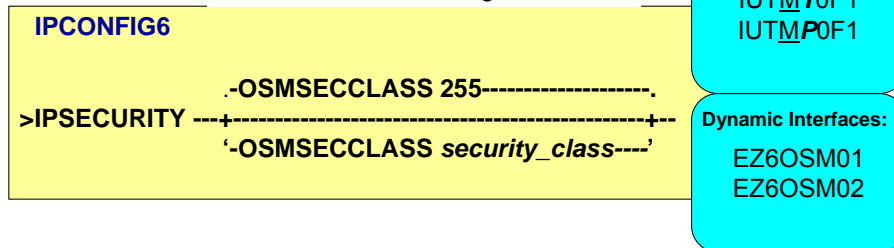
If the z/OS stack is enabled for IPv6, the stack defines two OSM interfaces. If the connectivity requirements on the previous chart are met, then Comm Server automatically starts these interfaces and dynamically creates TRLEs for them. These are IPv6 interfaces which only have a link-local address. These interfaces are always on a VLAN which is handled at the switch so the stack is unaware of the VLAN ID.

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. Generally speaking, z/OS as a member of an ensemble does not require a connection to the INMN. Z/VM with Virtual Machines that are Virtual Servers MUST be connected to the INMN. (The z/VM connection to the INMN is required even if any of the Virtual Guest are loaded with z/OS.) All members of the Ensemble MUST be connected to the intraensemble data network (IEDN).

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.

OSM Definitions on z/OS

TCP/IP Profile: IPConfig6 Statement



- If desired to apply a Security Class to OSM, add above value to the **IPCONFIG6** statement
 - The value is meaningful ONLY if the stack has been configured with **IPSECURITY**

SAF Controls

- RDEFINE SERVAUTH EZB.OSM.sysname.tcpname UACC(NONE)
- PERMIT EZB.OSM.sysname.tcpname CLASS(SERVAUTH) ID(userid) ACCESS(READ)

For OSM, VTAM dynamically creates TRLE with up to nine DATAPATH devices. You can see the dynamically created TRLE names and PORTNAMES on this chart: . These TRLEs always use PORTNUM 0. VTAM looks for the smallest consecutive even/odd pair of devices to use for READ and WRITE. 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**

The IPCONFIG6 statement has a parameter which defines a security class for IP filtering for the automatically created OSM interfaces. This parameter is only meaningful if IPSECURITY is enabled for IPv6. This allows you to create a filter rule to permit traffic for OSM interfaces. This is similar to the SECCLASS parameter under DYNAMICXCF which covers dynamic XCF interfaces.

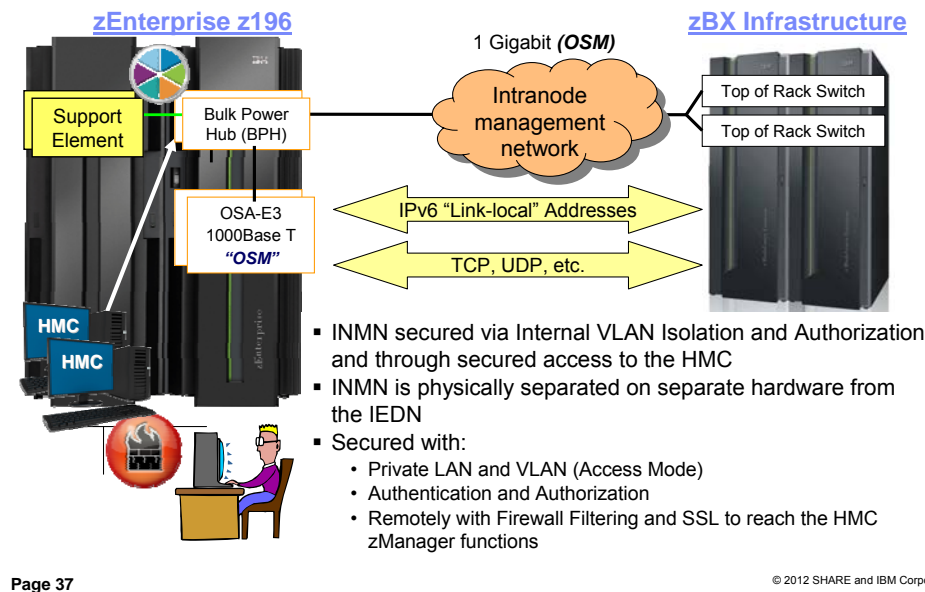
Interface names for the dynamically generated OSM Interfaces are: EZ6OSM01 or EZ6OSM02.

Several of the attributes for an OSM CHPID and Interface are different from those associated with the OSD Defaults: these attributes that are different are listed here:

- INBPERF DYNAMIC
- ISOLATE
- SECCLASS (value from IPCONFIG6 OSMSECCLASS)
- VMAC (OSA-generated VMAC with ROUTEALL)

Security Controls: Besides the closed, Layer 2 (VLAN) implementation of this network, and the enforcement performed by the hypervisors for use of this network, some operating systems exercises additional security controls. In z/OS OSM usage is restricted to management applications with proper authorization. In order to establish a TCP connection over OSM, send or receive non-TCP traffic over OSM, or join a multicast group over OSM, an application must have READ access to the SAF resource named "**EZB.OSM.sysname.tcpname.**" On z/OS, OSM access control is mutually exclusive with network access control. In other words, if packets are sent or received over an OSM interface, network access control is not enforced, but OSM access control is enforced instead. Stack-generated ICMPv6 traffic (such as Neighbor Discovery) is exempt from the need for OSM access control. There are additional MLS considerations for OSM interfaces and authorized platform management applications using them. See the MLS chapter in the IP Configuration Guide for details on these.

Can the External Customer Data Network Connect to the intranode Management Network? No!



The INMN connections use the VLAN in access mode because only one VLAN ID is recognized on this internal management network.

The TOR switch handles VLAN tagging and the stack remains unaware of VLAN IDs for these interfaces.

The External Customer Data Network cannot connect directly to the intranode management network. The INMN is accessible only through the HMC. And the HMC is locally secured on a Private LAN with Authentication and Authorization. If being accessed remotely, the HMC is secured with Firewall Filtering, with a connection secured with Secure Sockets Layer (SSL), and further discrete authorizations. The Unified Resource Management functions are also secured with further discrete authorizations to its special functions.

The INMN uses hardware that is entirely different from the hardware being used by the external network or even the IEDN. There is physical separation. So there is no physical connection to the IEDN except to management ports which cannot forward traffic or receive traffic from data ports. (The 10-Gig TORs are connected to a management port in the 1-Gig TORs.)

An administrator can make use of the INMN if he uses ping, traceroute, diagnostics from an Operating System attached to the INMN, but such an administrator is allowed to issue such diagnostic commands only if authorized. Furthermore, such diagnostic commands can reach only as far as the Support Element, since each Virtual Server (including VSs in LPARs) is isolated unto itself over this INMN.

A Customer Requirement: To access the INMN through a Network Management Application under secured conditions.

Summary: 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. (You do not have to exploit IPv6 at all to create an Ensemble; the use of IPv6 is transparent to the user.) 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.

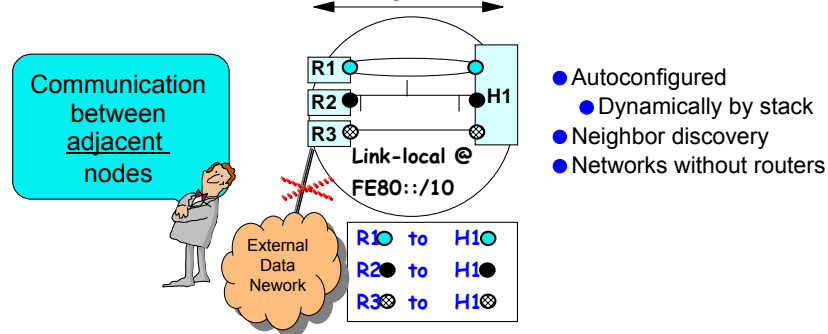
The INMN interfaces and their addresses are not reported to OMPROUTE nor can the operating system systems administrator add static or dynamic routes for these interfaces. Optionally you can specify a Security Class on the IPCONFIG6 statement of the z/OS TCP/IP profile (SecClass) to perform IPsec filtering. If desirable, you may allow stops, starts for these INMN interfaces or you may take packet traces and use OSAENTA on them.

Security Through Non-Routable IPv6 Addresses

Special Unicast Address: Link-local

- **Link-local Scope / Address - unique on a link only**
 - "Format prefix" or "Scope Prefix" is FE80::/10
 - Used for addressing between stations on the same link or LAN
 - **Cannot be routed**
 - One link-local required for each interface (excluding VIPA and loopback)
 - z/OS CS only allows a single link-local address per interface.

Communication using Link-Local Addresses



The scope of a packet's source and destination addresses controls where in the network the packet will be routed.

Every IPv6 interface will have a link-local address, which is automatically generated by the stack (considered one type of autoconfiguration)..

The link-local "format prefix" -- now called a "link-local routing prefix" or "link-local scope prefix" -- is FE80::.

Site-local (now deprecated) and global address can either be manually configured or dynamically generated.

A packet with a link-local source/destination address will not leave its originating LAN. A router receiving the packet will not forward it onto another physical LAN. Notice in the diagram how you cannot route between R3 and the External Data Network.

Used for any kind of temporary network

Dynamically configured by stack.

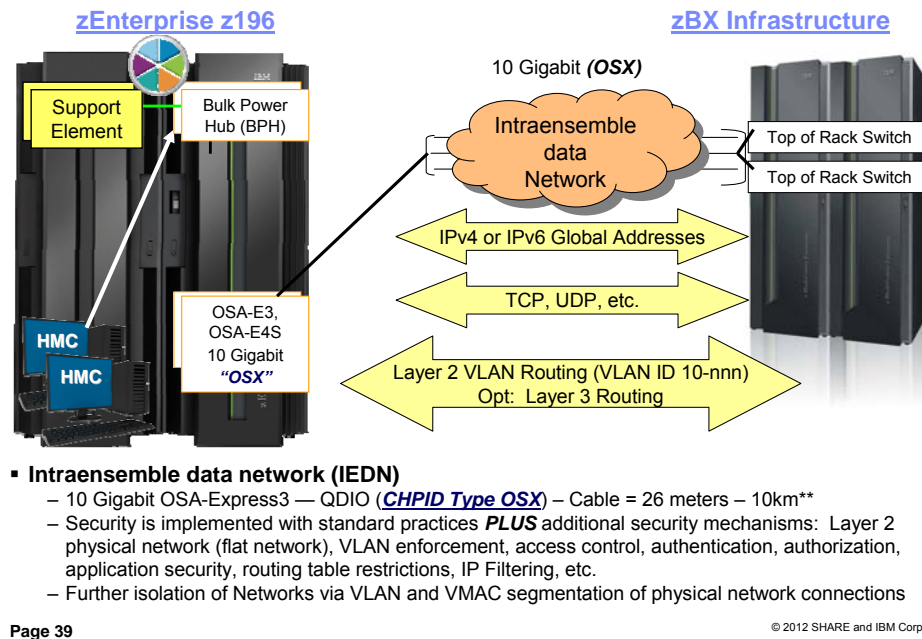
Neighbor discovery

Networks without routers

So remember: A packet with a link-local source/destination address will not leave its originating LAN. A router receiving the packet will not forward it onto another physical LAN.

Used for any kind of temporary network: Autoconfiguration, Neighbor discovery, Networks without routers

The intraensemble Data Network: Closed or Open Network



** In a long-reach environment, the IEDN connection can be up to 10km long. However the actual limit on the distance is the practical distance that would allow the Unified Resource Manager firmware at the HMC to continue to manage a second Ensemble Node. The intraensemble data network (IEDN) carries the data traffic among members of the ensemble. This network may be defined in the Operating System as an IPv4 or an IPv6 network. If you have not yet migrated your existing Customer Network to IPv6 you will probably for the short term continue to use IPv4 addressing in the IEDN. If you choose to implement IPv6 on the IEDN, you will want to understand the basics of networking with IPv6: IPv6 addressing, IPv6 protocol headers, IPv6 routing, IPv6 security, and so forth. Whichever type of IP addressing protocol you choose to implement, you will discover that the intraensemble data network relies on Layer 2 VLAN routing. All members of an ensemble with a requirement to communicate with each other must belong to the same VLAN. If they are on the same VLAN, then they also all belong to the same IP Subnet. All network connections in an IEDN require that a Virtual Medium Access Control address (VMAC) be assigned to the connection together with a VLAN ID. The default VLAN ID is 10. Any VLAN ID(s) used on the IEDN must be authorized at the HMC. The TOR for the IEDN ports to the LPARs are defined in [Trunk Mode](#).

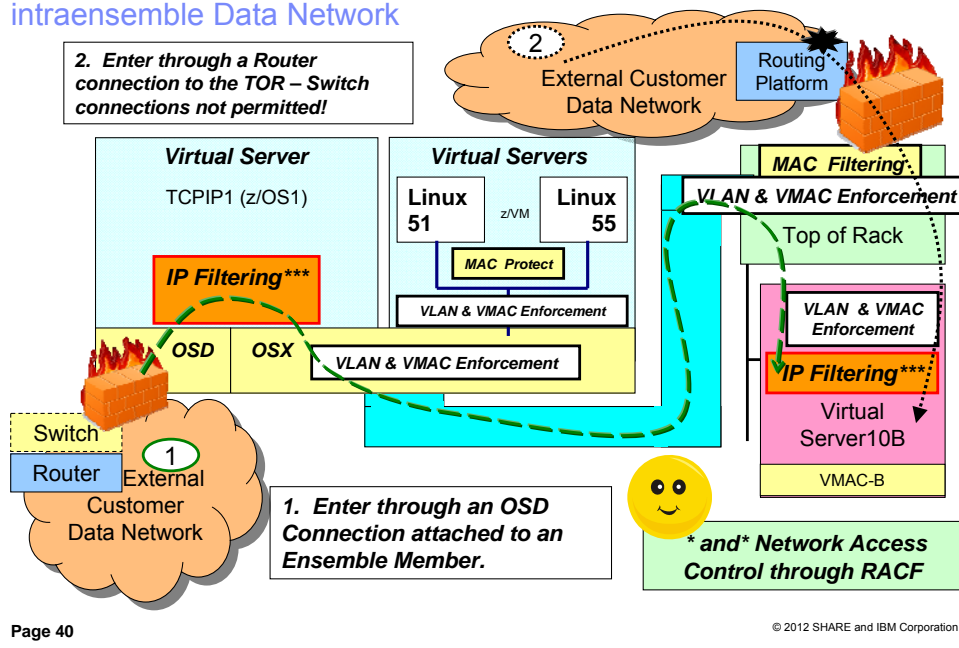
All Operating System TCP/IP stacks require a Layer 3 routing table, but this table can be very simple, because the basic implementation of the IEDN uses a flat network, that is, one with all addresses in the same IP subnet. If you choose to implement some addresses in the members that are outside this IP subnet, or, if you elect to allow communication between the IEDN and your External Customer Network, you may need to implement more complex routing tables – preferably with static route definitions. For reliability in an Ensemble, you must order redundant switches and redundant OSA-E3 cards to attach to the switches in order to interconnect the members of an Ensemble. You must also provide definitions that secure the access to the intraensemble data network (IEDN).

SUMMARY: Intraensemble data network (IEDN)

10 Gigabit OSA-Express3 or OSA-Express4S --- QDIO (CHPID Type OSX); Connected to authorized members of the Ensemble via the 10Gig TOR Switch; note that all physical switches are managed by the Support Element. (The 10-Gig TORs are connected to a management port in the 1-Gig management TORs.) 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.

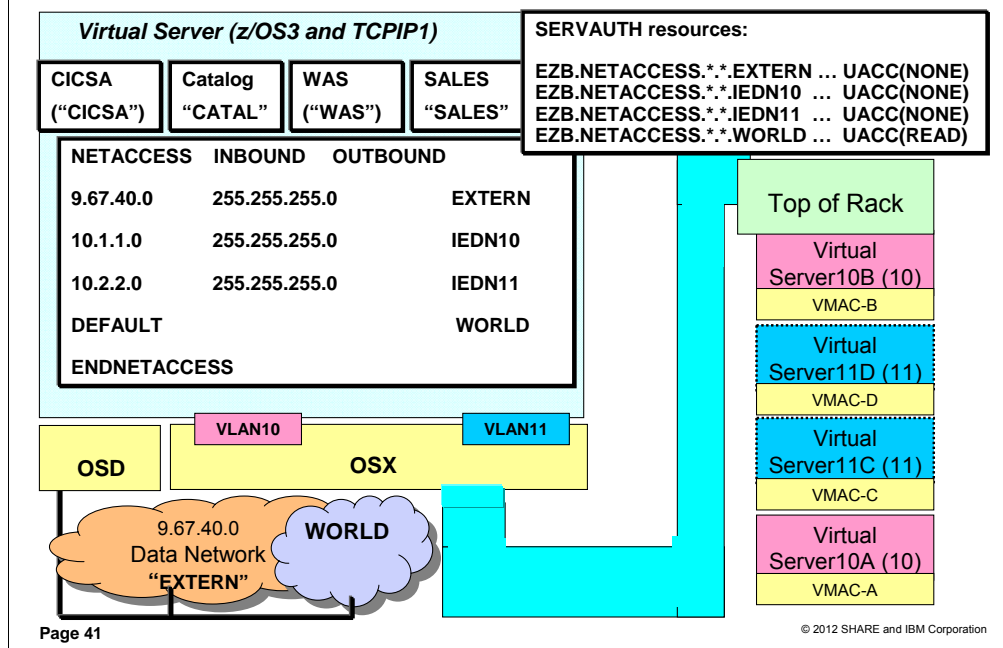
Security is implemented with standard practices **PLUS** additional security mechanisms: access control, authentication, authorization, application security, routing table restrictions, IP Filtering, etc. VLAN and VMAC segmentation of the network connections; Can assign Private network addresses (non-routable 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. We recommend implementing such a design with static routing.

Review: Connecting the Customer Data Network to the Intraensemble Data Network



This review is to remind you that RACF on z/OS can define a SERVAUTH class to permit or deny Network Access from outside networks. The next page gives you an example.

Network Access Controls to Limit Access between z/OS & IEDN



Notice in the diagram above how CICSA, Catalog, WAS and SALES are running under the OMVS Segment User ID of CICSA, CATAL, WAS, and SALES. Notice how we have defined separate security zones in the NETACCESS statements of the TCP/IP Profile. Notice how these security zones are also defined with SERVAUTH Resources in RACF.

Network access control gives system administrators the ability to assign permission for z/OS users to access certain networks and hosts. With this function, the ability of users to send or receive data between z/OS and certain networks can be controlled through z/OS. Network access control provides an additional layer of security to any authentication and authorization security that is used in the network or at the peer system by disallowing the unauthorized user to communicate with the peer network resource.

The NETACCESS statement in the TCP/IP profile is used to configure portions of your IP network into named security zones. Each defined security zone must have a SERVAUTH profile for the resource named EZB.NETACCESS.sysname.tcpname.zonename.

Network Access Control (NAC) limits user access to certain IP security zones defined by the NETACCESS statement. A security product, such as RACF, is used to check the permission of user IDs to send data to or receive data from these security zones.

Use of the zEnterprise System intra node management network (INMN) is protected by OSM access control and is exempt from network access control.

IP addresses are classified into *security zones*, in which each zone has a certain level of security sensitivity. A default security zone exists for interfaces that are not explicitly associated with a specific security zone. Security zones consist of one or more, perhaps discontinuous, IP address ranges that have the same security sensitivity and are identified by a specific zone name.

SAF is used to check whether users or groups of users have permission to access the security zone.

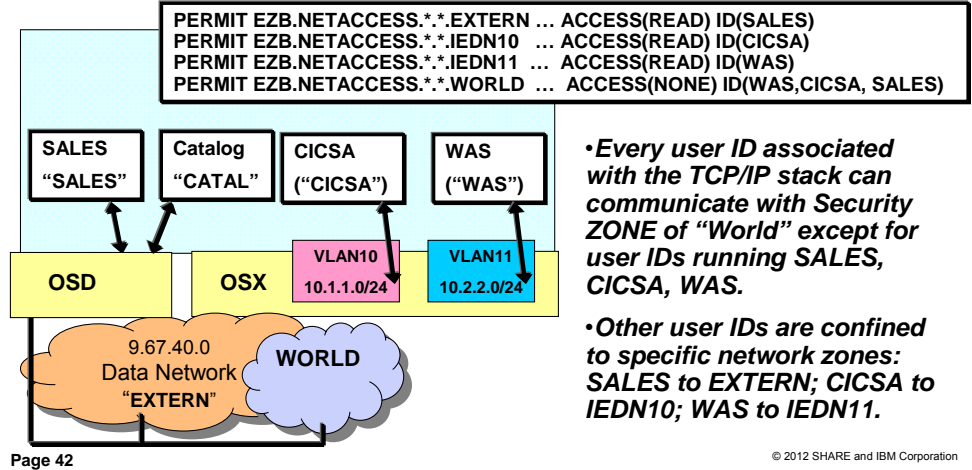
The installation defines a network access resource for each security zone and permits users or groups of users access to the resource. The security zone is represented by an SAF SERVAUTH profile name of EZB.NETACCESS.sysname.tcpname.zonename.

TCP/IP keeps a mapping of network resources by IP address to security zones. This mapping is consulted on certain inbound and outbound operations to determine the corresponding resource zone name for the most specific network defined. Then the current user's access to that resource is queried using SAF, and the operation is allowed or denied completion accordingly. This mapping is also consulted when the security ioctl is issued to extract the port of entry zone name of a socket's current peer.

Network access control is used to control z/OS user access to a peer address in an IP network through a sockets application. Resource access checks occur at connection setup or acceptance time for TCP, peer identification time for UDP and RAW, and on the first and potentially subsequent sends or receives (TCP, UDP, or RAW) to a particular destination in a socket's lifetime.

Network Access Controls to Limit Access between z/OS & IEDN ...

NETACCESS	INBOUND	OUTBOUND	EXTERN	SERVAUTH resources:
9.67.40.0	255.255.255.0		IEDN10	EZB.NETACCESS.*.EXTERN ... UACC(NONE)
10.1.1.0	255.255.255.0		IEDN11	EZB.NETACCESS.*.IEDN10 ... UACC(NONE)
10.2.2.0	255.255.255.0		WORLD	EZB.NETACCESS.*.IEDN11 ... UACC(NONE)
DEFAULT				EZB.NETACCESS.*.WORLD ... UACC(READ)
ENDNETACCESS				



Definitions for the scenario in the visual:

```

RDEFINE EZB.NETACCESS.*.EXTERN CLASS(SERVAUTH) UACC(NONE)
RDEFINE EZB.NETACCESS.*.IEDN10 CLASS(SERVAUTH) UACC(NONE)
RDEFINE EZB.NETACCESS.*.IEDN11 CLASS(SERVAUTH) UACC(NONE)
RDEFINE EZB.NETACCESS.*.WORLD CLASS(SERVAUTH) UACC(READ)
    
```

```

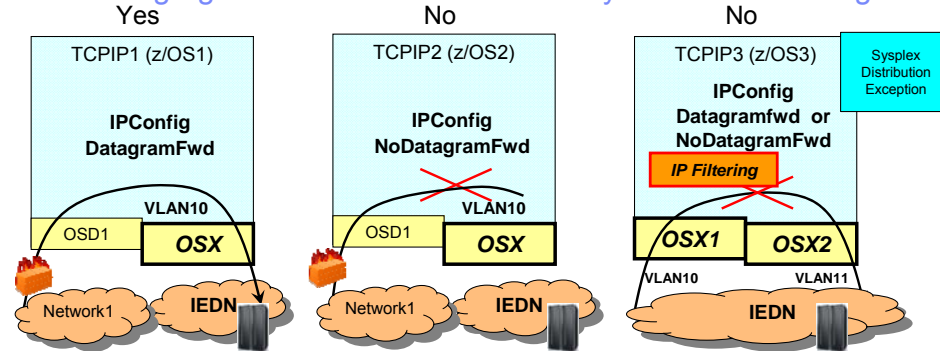
PERMIT EZB.NETACCESS.*.EXTERN CLASS(SERVAUTH) ACCESS(READ) ID(SALES)
PERMIT EZB.NETACCESS.*.IEDN10 CLASS(SERVAUTH) ACCESS(READ) ID(CICSA)
PERMIT EZB.NETACCESS.*.IEDN11 CLASS(SERVAUTH) ACCESS(READ) ID(WAS)
PERMIT EZB.NETACCESS.*.WORLD CLASS(SERVAUTH) ACCESS(NONE) ID(WAS,CICSA, SALES)
    
```

RESULT of these Definitions: **Every user ID associated with the TCP/IP stack can communicate with Security ZONE of "World" except for user IDs running SALES, CICSA, WAS.** The user ID running SALES is the only one permitted to communicate with the network zone EXTERN. Likewise CICSA with IEDN10 and WAS with IEDN11.

WARNING: Avoid using the same OMVS user ID for an end user as the one you use for Started Procs. You could inadvertently prevent some users from signing onto applications because of the way you may have defined NETACCESS. Always use discrete USERIDS for procedures from those you use for end users.

Understand that NETACCESS has its place, but it is not the perfect tool for all situations. You might implement something similar with IP Filtering in front of the System z, or within the z/OS LPAR, or elsewhere. It uses very few cycles, and certainly fewer than application access control lists or perhaps even IP filtering in the System z Stack, but an external device (firewall) would not use any cycles on System z.

Secure Segregation within an IEDN: No Layer 3 IP Forwarding



1. DatagramFwd and NoDatagramFwd operate as usual when routing is to occur between a non-IEDN interface and an IEDN interface. (*Visuals 1 and 2*)
 1. In addition, you may also deploy IP Filtering or Firewalls to permit or deny traffic.
2. Within an IEDN, only Layer 2 forwarding of traffic is permitted.
 1. Forwarding occurs only between Interfaces with identical VLAN IDs.
3. If forwarding between different VLAN IDs is desired, then, as usual in the TCP/IP architecture, Layer 3 routing is required. (Packets must be sent to a router for routing to a different VLAN.)
4. However, z/OS does not permit Layer 3 routing between OSX or IEDN interfaces at all.
 1. IEDN security has disabled Layer 3 Routing between OSX interfaces. (*Visual 3*)
5. **Other operating systems (z196 or zBX) may allow Layer 3 routing between VLANs in the IEDN.**

A feature of TCP/IP routing is to enable IP forwarding or to disable it. IP forwarding permits the transfer of data between networks. This IP forwarding capability is enabled or disabled in the z/OS TCP/IP stack with the IPConfig statement "DatagramFwd" or "NoDatagramFwd."

As you see in Visual #1, as long as routed traffic is permitted with IPConfig DatagramFwd and with any potential IP Filtering, then traffic entering an Ensemble Virtual Server by means of a non-IEDN path may be routed over the IEDN OSX OSA port into the IEDN. Visual #2 shows you that IP forwarding is generally disabled in TCPIP2 and so traffic cannot be routed between the external network and the IEDN.

However, as the third visual shows, traffic may not be routed between separate Ensemble VLAN IDs by the z/OS Ensemble Member. This is true regardless of the coding of DATAGRAMFWD or NODATAGRAMFWD. In the IEDN only Layer 2 forwarding is permitted and this can occur only if the VLAN IDs are the same. If it is necessary to route between two separate VLAN IDs in the IEDN, then the layer 3 routing table must be invoked to route outside the IEDN, through a router there, and then back into the IEDN over another VLAN. (See examples later in this presentation.)

As described IP Forwarding has been disabled in z/OS, but other operating systems (z/VM, for example) in the z196 or in the zBX may indeed still permit IP Forwarding to be able to route between separate VLANs within the IEDN in the same stack.

EXCEPTION: Regardless of the coding of NODATAGRAMFWD or DATAGRAMFWD, a Sysplex Distribution stack may forward received packets to a Sysplex Distribution target node over an XCF path or a VIPAROUTE path even within the IEDN.

zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

**Security: Identity, Authentication,
and Access Controls for Ensemble
Resources and Tasks**



Creating the Ensemble with the HMC and zManager



https://9.60.92.193 -

Ensemble Management Guide

Use this guide to assist you with setting up an ensemble. Click any of the links to take you directly to the tasks. Click the notes link to add notes about your ensemble, such as steps completed or number of members added.


Task	Allows you to...
Customize User Controls	(Optional) View an
User Profiles	(Optional) View an
View documentation	(Optional) Read or
Task	Allows you to...
Define alternate HMC	Choose another H
Create Ensemble	Create an ensemb
Add Member	Add a member to
Entitle zBX blades	Use the Perform M the Single Object
Manage Storage Resources	Add or remove sto
Manage Virtual Networks	Add or remove virt
New Virtual Server	Create a virtual se
Mount Virtual Media	Install your operati install the guest pl
Activate	Activate a virtual s
Open Text Console	Open a console w
New Workload	Create a workload business applicati
Add Performance Policies	Define the rules as
View Performance Metrics	View performance
View Workload Reports	View workload rep

- **Customer User Controls (Roles)**
- **User Profiles**
- **Create Ensemble**
- **Add Member**
- **Entitle zBX Blades**
- **Manage Storage Resources**
- **Manage Virtual Networks**
- **Create Virtual Server**
 - Install Operating System & Applications
- **Create Workload**
 - Performance Policies
 - View Performance Metrics
 - View Workload Reports

Done Internet

The HMC functions for working with the Unified Resource Manager perform many tasks that otherwise a system administrator or system programmer would have to perform. Some operating systems exploit the virtualization functions more than others. Note how the first task is to customize the User Controls and Roles for specific User Profiles. With the Ensemble, we now have more User Profiles and Roles to consider.

Ensemble Roles & User IDs – HMC



- You can create customized user profiles which would allow you to have unique user IDs and multiple user roles.
- The management of these user roles is performed by using the **Customize User Controls** task.
- Roles are assigned to users with the **User Profiles** or **Manage Users Wizard** tasks.

Using the Hardware Management Console or the Support Element you can assign and or design new Ensemble task roles, resource roles, and new User IDs, although templates exist for default user IDs and predefined Resource and Task Roles for the Ensemble. You can further customize these roles or define new ones.



Role-Based Access Controls with the Unified Resource Manager

Role	Description
Ensemble Administrator	Responsible for creating and managing the zEnterprise ensemble Create Ensemble, Add Member...
Virtual Network Administrator	Responsible for Managing Virtual Networks, Hosts, and MAC Prefixes Manage Virtual Networks, Add Hosts to Virtual Networks, Create VLAN IDs...
Virtual Server Administrator	Responsible for managing virtual servers New /Modify Virtual Server, Add Virtual Disk, Migrate...
Virtual Server Operator	Responsible for performing and scheduling virtual server activation/deactivation, mounting virtual media Activate, Deactivate, Mount Virtual Media, Console session...
Storage Resource Administrator	Responsible for managing storage resources – Storage Access Lists, WWPNs, z/VM Storage Groups Export WWPN, Import SAL, Add Storage Resources...
Workload Administrator	Responsible for managing workloads New /Modify workload, Add / Remove Virtual Servers..
Performance Management Administrator	Responsible for managing performance policies New /Modify performance policy, Import policy
Performance Management Operator	Responsible for performing and scheduling policy activations and creating threshold notifications Activate, Export Policy, Monitor System Events
Energy Management Administrator	Responsible for managing power settings including power capping and power savings Set Power Cap, Set Power Savings Mode, Set zBX Power Policy

- New task and resource roles enable isolation across management disciplines
- New predefined users EnsOperator and EnsAdmin

The role-based access control in the zManager provides granularity for the different functions that are used to manage an Ensemble.

Since these panels are available only from the HMC, the security for the Ensemble functions is another security layer already available in the HMC itself.

Security in the HMC and the Unifed Resource Manager:

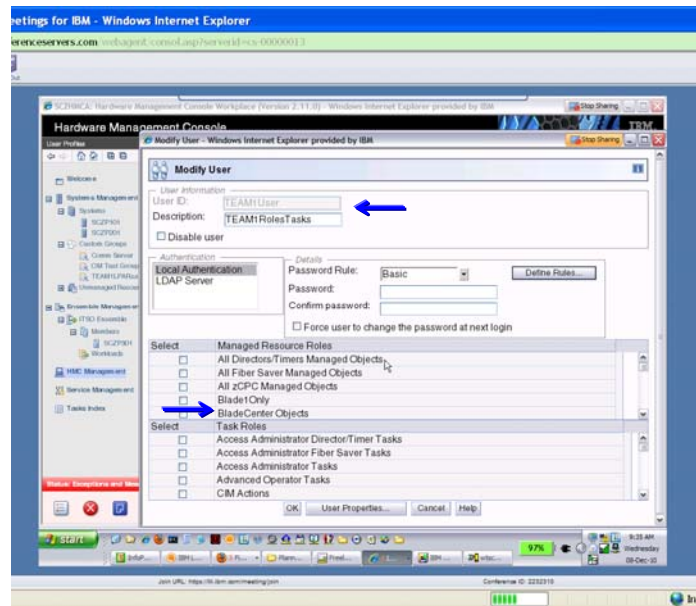
The HMC continues to enjoy the security controls inherent in the System z Console. Please see IBM System z Hardware Management Console Security White Paper by Kurt Schroeder (schroedk@us.ibm.com) published in Sept. 2008 and available at:

<http://nascpok.pok.ibm.com/rsf/zHMCSecurityWhitepaper.pdf>

This paper explains: “The HMC *Licensed Internal Code* includes a full-function firewall that is used to control network access to the HMC. As previously described, by default the HMC allows for virtually no inbound network traffic. As different features of the HMC are enabled (e.g.. remote access, SNMP based automation etc.) additional inbound network traffic is allowed.”

In addition, the Unified Resource Manager screens of the console enjoy additional security controls, like role-based security and operations over the closed, Intranode Management network (INMN).

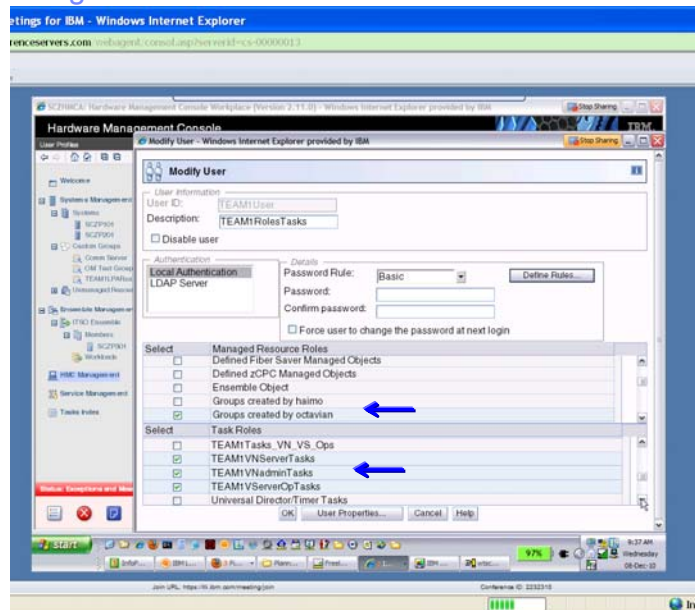
Restricting z196 Resources to a Particular User



This screen shot from the zManager panels shows that we have defined a new userid called “TEAM1User.” It also shows that we have defined a separate set of Resource Roles just for BLADE 1. (Blade Center Objects is for a role that comes automatically with the firmware of zManager.) On the next visual you will see that we have also defined new Task Roles (not part of the default definitions) that will be assigned to this same User.

In this manner, we can secure access to the HMC to ensure that a particular userid is only allowed to execute functions against specific members of the z196 and the Ensemble. That is, a Power Blade administrator might not be allowed to execute any functions against certain LPARs of the z196 and would be allowed access only to the Virtual Servers on Blade 1.

Restricting Certain Tasks to a Particular User



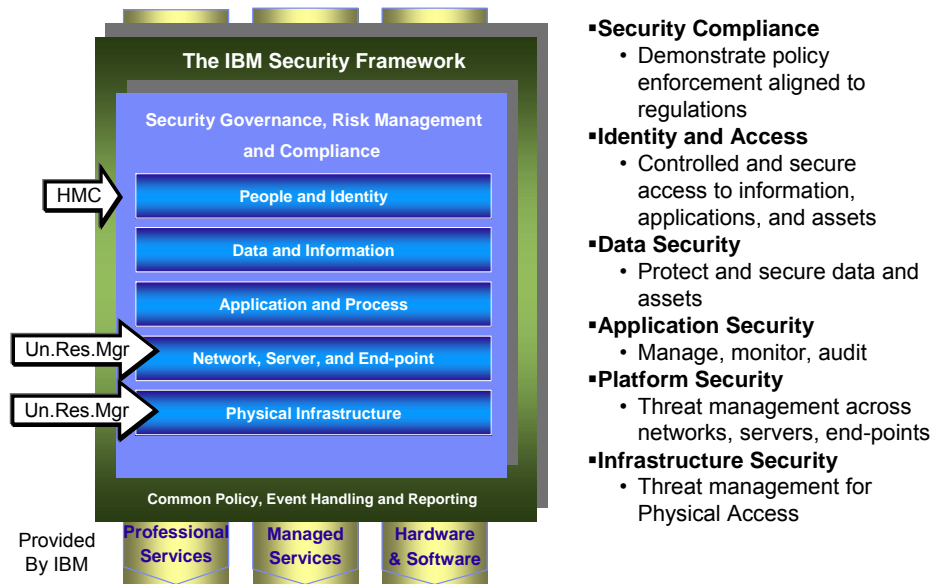
This screen shows that we have defined server Task Roles just for this userid of TEAM1User. In this manner, we can secure access to the HMC to ensure that a particular userid is only allowed to execute functions against specific members of the z196 and the Ensemble. That is, a Power Blade administrator might not be allowed to execute any functions against certain LPARs of the z196 and would be allowed access only to the Virtual Servers on Blade 1.

zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

IBM Security Framework – Traditional Security Services



The Security Story for Information Technology



The IBM Security Framework provides a model – a structured approach -- for selecting, designing, and monitoring technologies to protect all aspects of an IT organization.

IBM provides the professional services to assess an organization's needs for security with regard to compliance mandates and general security requirements. These services can design, implement, and manage security technologies and can recommend hardware and software solutions for an organization.

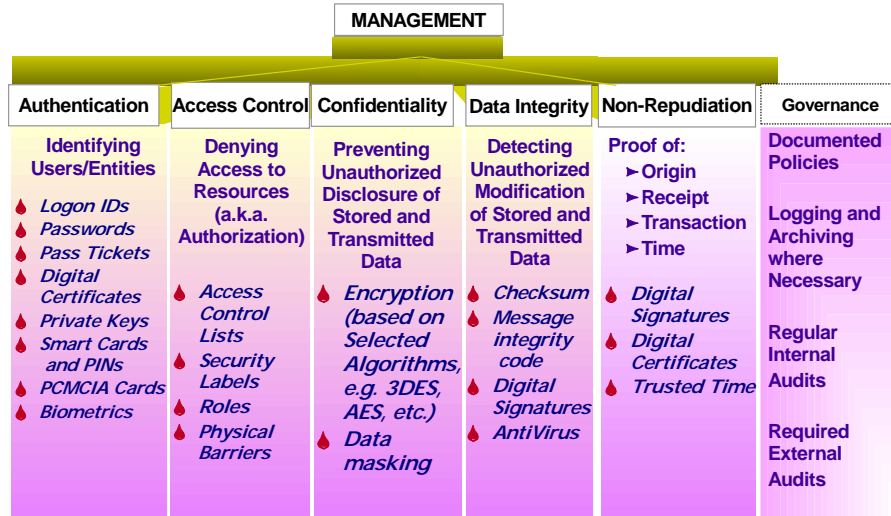
More information: Introducing the IBM Security Framework and IBM Security Blueprint to Realize Business-Driven Security; IBM RedGuide REDP-4528-00, July 2009 at www.redbooks.ibm.com

This model applies to any Information Technology implementation. You may apply it to the z196 or System z alone, or you may also apply it to the new zEnterprise System (z196 + zBX + zUnified Resource Manager).

Particularly the security layers for “Physical Infrastructure and Network,” and “Server, and Endpoint” are enhanced with features defined with the zEnterprise Unified Resource Manager and with the controlled access to the physical components of the zBX.

The Security Story for Information Technology

Security Services and Mechanisms



This is an older version of the ISO security model. Note the entry for "Governance" and "Logging." This is not part of the ISO model, but it is nevertheless integral for any security implementation. We have added it here to show its importance.

You have now seen a couple of architectures for security. What is the difference?

The IBM Security Framework describes WHAT needs to be protected in an IP installation.

This ISO Security architecture describes:

HOW to protect the data (i.e., which Services should be used to protect data and resources)

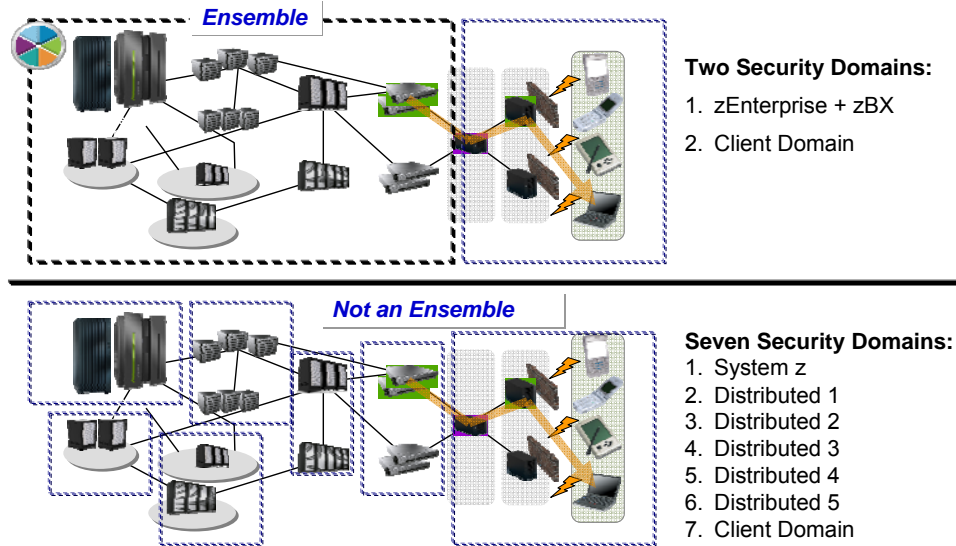
WHICH ways to protect the data (i.e., which security mechanisms could be used to protect the data and resources)

These Services and Mechanisms apply to the z196 or System z alone or even to the zEnterprise System. In fact, the zEnterprise System architectures provides innovative ways to strengthen these services when they are implemented. Most of these ways revolve around the use of the z Unified Resource Manager (zManager).

As you will see, Ensemble networking with the zEnterprise System provides many layers of security to ensure that data is not compromised. The rest of this presentation outlines for you many of these methods.

NOTE: We believe there is no reason to implement IPsec for encryption with the zEnterprise System because we believe that encryption should not be required in the IEDN. However, if the customer wants to encrypt, we would recommend SSL/TLS/AT-TLS.

Scope of Security Vulnerability: Ensemble vs. Traditional Network



Where fewer administrators are responsible for security implementations, the risks of human error and intentional security breaches are lower than where many security domains exist. Centralization, Simplification, and Enforcement of Security through the zEnterprise System with its Unified Resource Manager strengthen security at many levels by restricting the number of security domains required and by enforcing certain rules on specific definitions made by administrators of the Virtual Server Operating Systems.

zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

Summary: Security in the Ensemble



As you know, zManager controls stop at the operating system level. This means that there are all the other layers of security from the security framework that software at the operating system level can enforce. These are additional security implementations that reside in RACF, in networking protocols, and in software that can be added to an Ensemble implementation to protect the Ensemble network.



Frequently Asked Questions about Securing the Ensemble Network

- **How do I limit and secure access to the Ensemble Definitions?**
 - Use controls at the HMC (Unified Resource Manager or zManager functions)
 - Use existing HMC controls
- **How do I limit and secure access from within the IEDN to the Ensemble Virtual Servers and Networks?**
 - All Virtual Servers and VLANs must pass through “access points” (Hypervisors and TORs) where their authorization is confirmed.
 - Authorize the Virtual Servers to become Ensemble Members
 - Authorize the Virtual Servers to send data across the Ensemble Networks
 - Authorize the Virtual Servers to exploit only certain VLAN IDs
 - Use existing security techniques
 - Userid and passwords
 - Access controls to access storage
 - Firewalls and IP Filtering
 - Encryption
 - Etc.
- **How do I limit access to the Ensemble by the External Networks?**
 - Deploy routers and Firewalls to Permit or Deny traffic
 - Implement controls at the Top-of-Rack (TOR) switch and at the LPARs to limit access
- **With a hybrid solution that combines heterogeneous platforms into a single Enterprise System, how do I provide the network segmentation that many security mandates (like PCI) require?**
 - Centrally assign and have Hypervisors enforce VLAN IDs in the data network (IEDN)
 - Have Hypervisors and the TOR enforce Virtual Media Access Control (VMAC) addresses
- **Why are VLAN implementations in the Ensemble considered more secure than VLANs in a non-Ensemble environment?**
 - Without Ensemble, Network Interface Cards (NICs) of platforms connect to switch ports, which, if improperly configured, could allow miscoded VLAN IDs at the server level to gain access to the network infrastructure.
 - With Ensemble, the Unified Resource Manager firmware in the hypervisors and virtual switches ensures that servers are only allowed to connect to the VLAN(s) in the IEDN that they are explicitly authorized for. A misconfiguration of a virtual server (e.g. incorrect VLAN ID) results in a failure to connect to the IEDN.

There are many controls to permit or deny access to the IEDN. The Hypervisor controls (including the hypervisor component known as a VSwitch) operate within the IEDN to authorize access. The internal ports (Blade switch modules and TOR ports) are configured as trunk mode and allow all IEDN VLANs to flow BUT with the understanding that someone else – the Hypervisor and VSwitch) is enforcing the access control. For access from the external network into the IEDN, the Top-of-Rack (TOR) switches can implement access controls to the IEDN as can security technologies outside of or within an Ensemble LPAR or Virtual Server.

HMC Security White Paper: IBM System z Hardware Management Console Security White Paper. Author Kurt Schroeder (schroedk@us.ibm.com), Sept. 2008

<http://nascpok.pok.ibm.com/rsf/zHMCSecurityWhitepaper.pdf>

zEnterprise Network Security White Paper (ZSW03167-USEN-00) and Other Resources

www.ibm.com/systems/z/resources (Select “Literature” Entries)

http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfid=ZSW03167USEN&attachment=ZSW03167USEN.PDF

zEnterprise Network Security Frequently Asked Questions:

<http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FQ130131>

Enforcement and Filtering in the Ensemble

- **Centralized Control of the Enforcement Definitions in the IEDN**
 - VMACs (Physical Adapter Addresses)
 - VLANs (Virtual Local Area Network IDs)
 - IP Addresses
- **Hypervisors in the Ensemble ENFORCE the Definitions**
 - VMACs
 - VLANs
 - IP Addresses
- **The Top of Rack Switch FILTERs for ...**
 - Correct external MAC Addresses (Physical Addresses)
 - Correct external VLAN IDs
- **An Ensemble Member in the z196 LPARs PERMITS or DENIES**
 - Routing into the IEDN
 - ACCESS Controls into the IEDN
- **Deploy traditional security measures although many – like Encryption and Decryption -- may be unnecessary in an Ensemble**



Comparison of Today's Network Security with zEnterprise

Function	Today's Networks	zEnterprise
IP filtering	External Firewall	Possible Opportunity for elimination. zManaged VLANs restrict IP addresses.
Port filtering	External Firewall	Possible Opportunity for elimination. zManaged VLANs control Authorized Virtual Servers.
Application Gateway	External Firewall	Possible Opportunity for elimination. zManaged VLANs restrict IP addresses.
Stateful Packet Inspection	External Firewall	Possible Opportunity for elimination. zManaged Network leads to highest internal trust.
Partner Authentication	SSL/TLS, AT-TLS, IPSec, SSH	Requirement typically will not change. Exposures reduced (= application layer requirement).
Message Authentication	SSL/TLS, AT-TLS, IPSec, SSH	Requirements typically will not change. Exposures reduced (= application layer requirement).
Data confidentiality (Encryption)	SSL/TLS, AT-TLS, IPSec, SSH	Possible Opportunity for elimination. Physical closed network and zManaged VLANs restrict potential interference.
Message Integrity	SSL/TLS, AT-TLS, IPSec, SSH	Possible Opportunity for elimination. Physical closed network and zManaged VLANs restrict potential interference.

Traditional security implementations may no longer be necessary in the Ensemble. If you choose to use them anyway, you are responsible for their controlled and audited implementation. Some requirements may not change or may become an application layer requirement.

zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

Appendix: Tables on INMN and IEDN Security



Protecting the Ensemble Infrastructure – Part 1

INMN:	<p>1. Authentication and Access Control:</p> <ol style="list-style-type: none">1. Isolated and non-configurable VLANs in the INMN<ul style="list-style-type: none">• Exploited only by Management Applications communicating between the Support Element and the Virtual Servers2. Virtual Servers cannot communicate with each other over the INMN3. ISOLATE over a shared OSM OSA is enforced.4. Restrict access to OSM with SERVAUTH Class EZB.OSM.sysname.tcpname (z/OS only)5. IP Filtering is optional<ul style="list-style-type: none">• Can implement Security Classes for IP Filtering <p>2. External Network Access impossible; the INMN ends at the HMC and Support Element.</p>
--------------	---

Protecting the Ensemble Infrastructure – Part 2

IEDN:	<ol style="list-style-type: none"> 1. Reduce the Scope of Security Vulnerability through Elimination of Routing Hops and Network Administration 2. Authentication and Access Control: <ol style="list-style-type: none"> 1. Only authorized servers can access the IEDN <ul style="list-style-type: none"> • OSDSIM support in z/VM grants access to certain Guests on the VSwitch under VM 2. Isolated VLANs in the IEDN with VLAN Enforcement for nodes on the VLAN. 3. By default: No routing within the IEDN between different VLAN IDs <ul style="list-style-type: none"> • VMAC enforcement is required within the IEDN; VMAC(MAC) Filtering is recommended between External Customer Network and the IEDN • ISOLATE over a shared OSX OSA is optional, but plan carefully! • If necessary, implement Firewall for routing between VLANs in the IEDN 4. Restrict Access with z/OS TCP/IP NETACCESS Controls 5. IP Filtering (IPSec) is optional <ul style="list-style-type: none"> • z/OS can implement Security Classes for IP Filtering (IPSec) 3. Confidentiality (Encryption) is optional. 4. External Network Access permitted only if connection is authorized in the LPAR or the TOR <ol style="list-style-type: none"> 1. Recommendation: Implement Customer Firewalls and/or IP Filtering 2. Implement Multi-Level security where desirable.
--------------	---

Some topics mentioned here have not been part of the main body of this presentation. Please review the materials in the Appendix and the References at the back of this document for more information.

What May Catch You by Surprise: Simplicity of IPv6 & Security



Page 61

- **IPv6 for INMN need not be a “Big Deal”**
 - You need not learn IPv6!
 - Be aware of Display Output Changes when a z/OS stack is IPv6-enabled
 - IPv6 is completely transparent in z/VM
- **INMN Implementation is Simple and Secure**
 - Customer needs no control over the IPv6 link-local addresses, which are dynamically assigned
 - A CLOSED, FLAT network
 - Authorizations and OSM Access Control (z/OS) are required to reach it
 - IP Forwarding is disabled
- **IEDN enablement adds only a few parameters**
 - In z/OS: Dynamic VTAM Definitions or Manual
 - In z/VM the Systems Management API builds the z/VM Directories and VSwitch definitions
- **The IEDN network is inherently more secure than the External Customer Network**
 - Customer decides if External Connectivity is needed
 - It is an Internal Network secured by
 - Additional Security Definitions
 - VLAN Routing with VLAN Enforcement
 - VMAC Enforcement
 - VMAC Protect Function in z/VM VSwitch and in Network Virtualization Mgr.
 - VMAC Filtering via HMC
 - IEDN IP Forwarding is disabled on z/OS
 - Secured by traditional means to secure all security layers
 - Secured by traditional means to secure the networking layer:
 - Firewall Filtering (not stateful) that can be loaded into the Operating Systems
 - External stateful firewall filtering
 - Network Access Controls (z/OS)
 - Technologies to protect Data in Transit
 - Multilevel Security (z/OS)

© 2012 SHARE and IBM Corporation

Visual: Edvard Munch “The Scream” (“The Cry”), 1893-1910

IPv6 for INMN need not be a “Big Deal” --- The Virtual Servers, including z/OS, z/VM, and Linux on z become Dual-Mode Stacks; **BPXPRMxx** In z/OS there are changes for UNIX initialization

IPv4 and IPv6 interfaces and routing

In z/OS IPv6 is “INTERFACE” only (not DEVICE/LINK)

On z/OS: Long Format only of NETSTAT

NETSTAT ROUTE for IPv6 and not NETSTAT GATE

INMN Implementation is Simple and Secure

The INMN enablement results in “Stateless Autoconfiguration” of IPv6 addresses

Customer needs no control over the IPv6 link-local addresses, which are dynamically assigned

The INMN network need not be secured beyond what you would normally do with access to the HMC

A CLOSED, FLAT network

Authorizations and OSM Access Control (z/OS) are required to reach it

IEDN enablement adds only a few parameters to definitions

allows customer to define IPv4 addresses as usual, or ...

IPv6 addresses

Dynamically, through stateless autoconfiguration (advertisements from routers build the addresses)

Manually, through customer definition.

The IEDN network is inherently more secure than the External Customer Network

It is an Internal Network secured by Additional definitions in the HMC authorized only to specific administrators; VLAN Routing with VLAN Enforcement; z/VM VMAC Protect Function enforces VMACs assigned and generated with Ensemble Management; VMAC Filtering via HMC; Additional definitions in RACF to limit access; Secured by all other means already available to the Operating Systems loaded into the Virtual Servers; Multilevel Security (z/OS); Secured by traditional networking security:

Firewall Filtering (not stateful) that can be loaded into the Operating Systems

External stateful firewall filtering

Network Access Controls (z/OS)

Technologies to protect Data in Transit

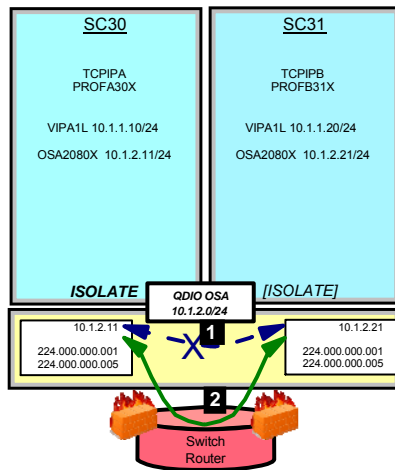
zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

Appendix: ISOLATE on the OSX OSA Port for Security



As you know, zManager controls stop at the operating system level. This means that there are all the other layers of security from the security framework that software at the operating system level can enforce. These are additional security implementations that reside in RACF, in networking protocols, and in software that can be added to an Ensemble implementation to protect the Ensemble network.

Secure Segregation of Traffic over Shared OSD: ISOLATE



1 = Direct route between TCPIPA and TCPIPB
2 = Indirect route between TCPIPA and TCPIPB

- Code *ISOLATE* on the OSD definition or the z/VM VSwitch to eliminate direct routing over the shared OSA to or from the image coded with *ISOLATE*. A z/OS Example:

```
INTERFACE OSADI
CHPIDTYPE OSD
DEFINE IPAQENET
PORTNAME OSADPORT
IPADDR 10.1.2.11/24
MTU 1492
VLANID 10
VMAC ROUTEALL
ISOLATE
```

- For further isolation you may even want to implement *NOTPART* in the **IOCDS** to eliminate the use of certain device addresses in partitions sharing the OSA OSD Port.



BACKGROUND INFORMATION on ISOLATE:

ISOLATE was introduced in z/VM 5.3 with APAR VM64281 as “Port Isolation” on a VSwitch and then later in z/OS as “Connection Isolation” with the V1R11 release. Since VLAN enforcement did not exist on an OSD device the way it does on an OSX device, many customers were asking for a means to prohibit traffic forwarding across a shared OSA port.

VLANs, when properly implemented, can isolate traffic over a shared network and shared OSA port. The isolation is complete if all TCP/IP stacks that share an OSA port implement VLAN ID tagging and assign separate VLAN IDs.

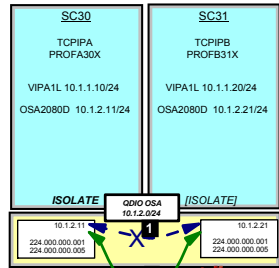
Another method that is available to isolate traffic across a shared OSA port is by using “Connection Isolation” in either z/OS or z/VM. This method can be deployed for OSD CHPIDs with or without out assigning a VLAN ID or a VMAC to the OSA port and simply requires that the parameter *ISOLATE* be coded on the QDIO Interface definition. If it is deemed necessary, you can also implement Connection Isolation over an OSX port, although the security implementations of Ensemble networking already perform isolation over the shared port when separate VLAN IDs are coded. Introducing *ISOLATE* without thought into the OSX environment can have unintended consequences as subsequent charts show.

WHY YOU SHOULD NOT NEED “ISOLATE” WITH OSX DEVICES:

With sharing over an OSX port, the VLAN Isolation is complete because the Hypervisors enforce VLAN IDs. Unlike with QDIO, an OSX implementation requires a VLAN ID and a VMAC as assigned by the firmware definitions in the Unified Resource Manager. Remember that OSX is a z/Managed resource and the isolation method that is deployed and required is VLANs. OSX is a flat Layer 2 broadcast domain using VLANs; technically no Layer 3 routing is required. Deploying *ISOLATE* can cause unintended networking issues in the Ensemble environment, as you can examine on subsequent charts.

A note about SHARED OSX ports: If you are sharing an OSX OSA port among LPARs and wish to exclude certain devices on the OSX CHPID from being shared, take advantage of the *NOTPART* keyword 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, improperly deploying *ISOLATE* can cause unintended networking issues.

Secure Segregation of Traffic over Shared OSD: ISOLATE & Layer 3 Routing



10.1.2.240
 1 = Direct route between TCPIPA and TCPIPB
 2 = Indirect route between TCPIPA and TCPIPB

- Combine OMPROUTE with Static Routes to bypass direct routing through OSA port.

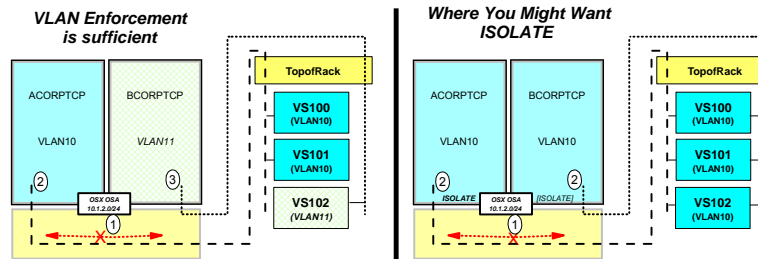
```

;TCPIPA.TCPPARMS(ROUTA30X)
;AUTOLOG LIST: INITIALIZE OMPROUTE
...
BEGINRoutes
; Direct Routes - Routes directly connected to my interfaces
; Destination Subnet Mask First Hop Link Name Packet Size
ROUTE 10.1.2.0/24 10.1.2.240 OSA2080D mtu 1492
ROUTE 10.1.1.0/24 10.1.2.240 OSA2080D mtu 1492
ROUTE 10.1.1.20/32 10.1.2.240 OSA2080D mtu 1492
ENDRoutes
  
```

```

;TCPIPB.TCPPARMS(ROUTB31X)
;AUTOLOG LIST: INITIALIZE OMPROUTE
...
BEGINRoutes
; Direct Routes - Routes directly connected to my interfaces
; Destination Subnet Mask First Hop Link Name Packet Size
ROUTE 10.1.2.0/24 10.1.2.240 OSA2080D mtu 1492
ROUTE 10.1.1.0/24 10.1.2.240 OSA2080D mtu 1492
ROUTE 10.1.1.10/32 10.1.2.240 OSA2080D mtu 1492
ENDRoutes
  
```


Effect of Introducing ISOLATE with OSX into the Ensemble: 2 Examples



1 = Two TCP images on LPARs in the Ensemble are isolated from each other in the IEDN because of VLAN Enforcement.
 2 = The ACORPTCP image on VLAN10 can communicate over the IEDN only with Virtual Servers in the zBX that belong to the appropriate VLAN10.
 3 = The BCORPTCP image on VLAN11 can communicate over the IEDN only with Virtual Servers in the zBX that belong to the appropriate VLAN11.
 4 = If you introduce an external router solution, ACORP and BCORP can communicate with each other -- just not over the IEDN!

1 = Two TCP images on LPARs in the Ensemble are isolated from each other on the shared OSX port because of ISOLATE coding.
 2 = Each LPAR can still communicate with all the Virtual Servers in the zBX that belong to the appropriate VLAN.
 3 = If you introduce an external router solution, ACORP and BCORP can communicate with each other -- just not over the IEDN!

Best Practices:

- Use separate VLAN IDs to segregate Virtual Servers sharing the same OSX port from each other. The VLAN enforcement of the Ensemble makes this port sharing secure and you should not need ISOLATE.
- If you must introduce ISOLATE for further segregation or segmentation on the OSA port, you may – and without negative consequences *in certain scenarios*. See PRS4160 for more detail at www.ibm.com/support/techdocs.

NOTE: ISOLATE is enforced on the INMN network; it is optional on the IEDN network. In z/OS the ISOLATE parameter is coded on the INTERFACE definition; in z/VM ISOLATE is coded on the VSwitch.

In the first visual above you see that the z196 has two LPARs, with each LPAR dedicated to a different company's production TCP/IP stack: ACORPTCP and BCORPTCP. We are managing the ensemble and keeping the processing secure by ensuring that the traffic from ACORPTCP flows over VLAN10 in the IEDN, and that the traffic from BCORPTCP flows over VLAN11 in the IEDN. ACORPTCP needs to communicate only with VS100 and VS101 in the zBX; BCORPTCP needs to communicate only with VS102 in the zBX. The VLAN enforcement required for Ensemble Networking ensures that the two TCP/IP stacks on the z196 cannot communicate over the IEDN with each other.

In the second visual we show you a different design: ACORPTCP and BCORPTCP are both allowed to send traffic over VLAN10. They are both using the same Virtual Servers in VLAN10 in the zBX, but the TCP stacks themselves are not allowed to communicate directly with each other. For this reason the system administrators have introduced Connection Isolation into the picture by coding ISOLATE on at least one (and possibly both) of the OSX Interface definitions of the two z/OS TCP/IP stacks. Therefore, while permitting communication freely over VLAN10 to Virtual Servers in the zBX, they are not permitting communication over the shared OSX port between the two TCP/IP stacks in the LPARs. This is an unusual network design and one would have to have very good reasons for introducing it. You could eventually see a negative consequence of this design, for example, if you needed to assign a second TCP/IP stack to each company. In such a case, the two TCP/IP stacks of a single company would not be able to communicate with each other at all over the IEDN because of the ISOLATE coding that had already been introduced in the first TCP/IP stack.

WHY YOU SHOULD NOT NEED ISOLATE WITH OSX DEVICES:

With sharing over an OSX port, the VLAN Isolation is complete because the Hypervisors enforce VLAN IDs. Unlike with QDIO, an OSX implementation requires a VLAN ID and a VMAC as assigned by the firmware definitions in the Unified Resource Manager. Remember that OSX is a z/Managed resource and the isolation method that is deployed and required is VLANs. OSX is a flat Layer 2 broadcast domain using VLANs; technically no Layer 3 routing is required. Deploying ISOLATE can cause two stacks sharing an OSX port and using the same VLAN not to be able to communicate with each other over the internally shared OSA path; if the two stacks need to communicate with each other despite the ISOLATED internal path, then you must deploy a router or a switch solution with IP Firewall Filtering rules to maintain the security and integrity of the Ensemble design.

A note about SHARED OSX ports: If you are sharing an OSX OSA port among LPARs and wish to exclude certain devices on the OSX CHPID from being shared, take advantage of the NOTPART keyword 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, improperly deploying ISOLATE can cause unintended networking issues as you may explore further on the subsequent charts.

zEnterprise networking – Secure Networking with the zEnterprise Ensemble (Part 3)

REFERENCES



As you know, zManager controls stop at the operating system level. This means that there are all the other layers of security from the security framework that software at the operating system level can enforce. These are additional security implementations that reside in RACF, in networking protocols, and in software that can be added to an Ensemble implementation to protect the Ensemble network.

References (White Papers, FAQs, Presentations)

- zEnterprise System Frequently Asked Questions (FAQs)
 - www.ibm.com/systems/z/faq

- zEnterprise Network Security White Paper (ZSW03167-USEN-00) and Other Resources
 - www.ibm.com/systems/z/resources (Select “Literature” Entries)
 - http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&appname=STGE_ZS_ZS_USEN&htmlfid=ZSW03167USEN&attachment=ZSW03167USEN.PDF

- zEnterprise Network Security Frequently Asked Questions:
 - <http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FQ130131>

- IBM zEnterprise System Network Virtualization, Management, and Security (Parts 1 and 2: Overview and Detail)
 - w3.ibm.com/support/techdocs

- IBM System z Hardware Management Console Security White Paper
 - Author Kurt Schroeder (schroedk@us.ibm.com), Sept. 2008
 - <http://nascpok.pok.ibm.com/rsf/zHMCSecurityWhitepaper.pdf>

References (Hardware)

▪ zBX Publications

- zBX Service Guide GC28-6884-01
- zBX Installation Manual (2458-002) GC27-2610-00
- zBX IMPP (2458-002) GC27-2611-00
- zBX Service Education SE245800
- zBX Safety Inspection (for mod 1 and 2) GC28-6889-00
- IBM License Agreement for Machine Code SC28-6872-00
- Systems Environmental Notices and User Guide Z125-5823-02
- Systems Safety Notices G229-9054-02

▪ Redbooks (www.redbooks.ibm.com)

- IBM zEnterprise Technical Introduction, SG24-7832
- IBM zEnterprise Technical Guide, SG24-7833
- IBM zEnterprise Configuration Setup, SG24-7834
- IBM zEnterprise Platform Management, SG24-7835
- IBM System p® Advanced POWER Virtualization Best Practices, redp4194
- IBM BladeCenter JS12 and JS22 Implementation Guide, SG24-7655)


▪ zBX 2458-002 SAPR Guide



- SA10-006
 - Note: Different SAPR Guide than 2458-001
 - 2458 TDA Confirmation Form

References (Software and Security)

- **z/OS Ensemble Implementation**
 - z/OS Communications Server V1R12 SNA Network Implementation Guide (SC31-8777)
 - z/OS Communications Server V1R12 SNA Network Definition Reference (SC31-8778)
 - z/OS Communications Server V1R12 IP Configuration Guide (SC31-8775)
- **IPv6 Information**
 - IPv6 Network and Application Design Guide Version 1 Release 12 (SC31-8885)
- **z/VM Ensemble Implementation**
 - *z/VM: CP Planning and Configuration (SC24-6083)*
- **Introducing the IBM Security Framework and IBM Security Blueprint to Realize Business-Driven Security; IBM RedGuide REDP-4528-00, July 2009**
 - www.redbooks.ibm.com
- **[Security on the IBM Mainframe](#), SG24-7803-00 Redbooks®, published 30 April 2010**
 - www.redbooks.ibm.com
- **[Introduction to the New Mainframe: Security](#), SG24-6776-00 Redbooks, published 3 April 2007, last updated 26 April 2007**
 - www.redbooks.ibm.com

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/atmastr.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



Questions ?



Thank You!

gdente@us.ibm.com



Questions ?



Thank You!

gdente@us.ibm.com