

* The Hardware Management Console (HMC) is attached to the same LAN as the server's support element (SE) - See #65 zTidBits(Unfiled Resource Manager) - a.k.a. zManager
 This LAN is referred to as the **Customer Managed Management Network** and the HMC communicates with each Central Processor Complex (CPC), and optionally to one or more zEnterprise BladeCenter Extensions (zBXs), through the Support Element (SE).

- If the zEnterprise System server is **not** a member of an ensemble, it is operated and managed from one or more HMCs (just as any previous generation System z server).

NOTE: Previously, the role of HMCs were *stateless* (they did not keep any system status) and therefore not affecting system operations when, if necessary, they were disconnected from the system. The system can (however not recommended) be managed from either SE.

* If the **zEnterprise System node** is defined as a member of an **ensemble**, the **primary HMC** is the authoritative owning (*stateful*) component for zManager configuration and policies that have a scope that spans all of the managed CPCs/SEs in the ensemble.

* It will no longer simply be a console/access point for configuration and policies that is owned by each of the managed CPC's SEs. The managing HMC has an **active role in ongoing system monitoring and adjustment**.

- This requires the HMC to be configured in an primary/alternate configuration and **cannot** be disconnected from the managed ensemble members.

NOTE: The **primary HMC** and its **alternate** must be connected to the **same subnetwork** to allow the alternate HMC to take over the IP address of the primary HMC during failover processing.

- Customers often would deploy multiple HMC instances to manage an overlapping collection of systems.

- Until the **zEnterprise**, all of the HMCs **were** peer consoles to the managed systems and all management actions are performed by each z196 with its optional zBX possible to any of the reachable systems while logged into a session on any of the HMCs (subject to access control)-makes up a node of an ensemble.

* With the **zEnterprise zManager**, this paradigm **has now changed** where **only one primary alternate pair** of HMCs can manage ensembles.

- In this environment, if a **zEnterprise System node** has been added to an ensemble, management actions targeting that system can **only be done from the managing (primary) HMC** for that ensemble.

NOTE: A **remote HMC browser session** to the HMC that is the ensemble-managing HMC for an ensemble allows a user currently logged onto another HMC or a workstation to perform related actions.

Network Extensions: Each node in an ensemble includes as many as **five distinct networks**.

* Figure directly on right shows a **high-level summary** of the connectivity required for the zBX environment.

- There are **three types of LANs** (each with redundant connections) that attach to the zBX: the **INMN**, the **IEDN**, and the **customer managed data network**. The **INMN** is **fully isolated** and only established between the owning z196 server and the zBX. **NOTE:** The **IEDN** connects the zBX to a maximum of **eight z196 servers**. Each z196 server must have a **minimum of two connections** to the zBX.

- The **IEDN** is used to connect a zBX to a **maximum of seven other zBXs**.

- The **IEDN** is a **VLAN-capable network** that allows enhanced security by **isolating data traffic** between virtual servers (both test and production can run on same VLAN).

- The right figure shows the z196 connections through **two OSA-Express3 1000BASE-T features** (CHPID type OSM) to the **INMN TOR switches**.

- The **OSA-Express3 10 GbE features** (CHPID type OSX) connect to the **two IEDN TOR switches**. Depending on workload requirements, any OSA-Express2 or OSA-Express features (CHPID type OSD) can connect to the customer managed data network.

- The **Fibre Channel (FC)** connections are **only required** between the zBX and the attached Fibre Channel disk or storage area network (SAN).

NOTE: It is the **client's responsibility** to supply the cables for the IEDN, the customer managed network, and the connection between the zBX and the FC disk.

* The **IEDN** provides **private and secure** 10 GbE high speed data paths between **all elements** of a zEnterprise ensemble (up to eight z196s with optional zBXs).

- The zBX is managed by the HMC through the **physically isolated INMN**, which interconnects all resources of the zEnterprise System (z196 and zBX components).

- The **scope** of the intranode management network (INMN) is **within an ensemble node**.

- INMNs in different nodes are **not** connected to each other.

- The **INMN connects** the Support Element (SE) of the z196 to the hypervisor, optimizer, and guest management agents within the node. **Communication across the INMN** is exclusively for the purpose of enabling the zManager of the HMC to perform its various management disciplines (for example, performance management, network virtualization management, or energy management) for the node.

- The z196 connection to the INMN is achieved through the **definition of a CHPID type OSM**, which can be defined over an OSA-Express3 1000BASE-T Ethernet feature. **NOTE:** There is also a 1 GbE (OSX) infrastructure within the zBX.

- The **key points to consider** for an **INMN** are:

> Each z196 server must have **two OSA-Express3 1000BASE-T ports** connected to the Bulk Power Hub in the same z196

(see #66 zTidBits(z196&IOS) middle-diagram);

>> The two ports provide a redundant configuration for failover purposes in case one link fails.

>> For availability, each connection should be from two different OSA-Express3 1000BASE-T features within the same z196 server.

- **OSA-Express3 1000BASE-T ports** can be defined in the IOCDs as SPANNED, SHARED, or DEDICATED. See #24 zTidBits (An I/O White Paper) on channel types

- z/OS Communication server **TCP/IP stack must be enabled for IPv6** ; The CHPID type OSM related definitions will be dynamically created.

NOTE: No IPv4 address is needed.

- z/VM virtual switch types provide **INMN access**.

- Two 1000BASE-T top of rack switches in the zBX (Rack B) are used for the **INMN**; **no additional** 1000BASE-T Ethernet switches are required.

* **1000BASE-T supported cable:**

- 3.2 meter Category 6 Ethernet cables are shipped with the z196 ensemble management flag feature (FC 0025). Those cables connect the OSA-Express3 1000BASE-T ports to the Bulk Power Hubs.

- 26 meter (85') Category 5 Ethernet cables are shipped with the zBX. Those cables are used to connect the z196 Bulk Power Hubs and the zBX top of rack switches.


* Right figure displays the **primary and alternate HMC** configuration connecting into the **two bulk power hubs** (BPHs) in the z196. The 1000BASE-T TOR switches in the zBX is also connected to the BPHs in the z196.

* A **zBX rack can support a maximum of two BladeCenter chassis**. Each rack is designed for enhanced air flow and is shipped loaded with the initial configuration. It can be upgraded on-site to a larger configuration.

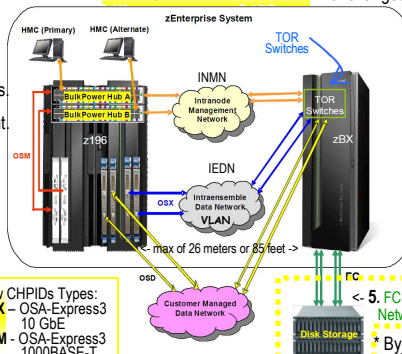
There are 14 blade server slots (BS01 to BS14) available in a zBX BladeCenter chassis.

* **Each slot** is capable of housing **any zBX supported blades**.

NOTE: In the future, some supported blade types will be double-wide.



Networking is a pervasive component of an ensemble.





New CHPIDs Types:
 OSX - OSA-Express3 10 GbE
 OSM - OSA-Express3 1000BASE-T

The IBM OSA-Express is an integrated hardware feature that provides direct connection to clients on local area networks (LANs). The OSA-Express feature plugs into an I/O slot as a channel card. The OSA-Express is identified in the hardware I/O config by its channel path identifier (CHPID).

The networks (INMN and IEDN) that connect the z196 to the zBX are constructed with extreme security in mind.

The first rack (Rack B) in the zBX is the **primary rack** where one or two BladeCenter chassis and four top of rack (TOR) switches reside.

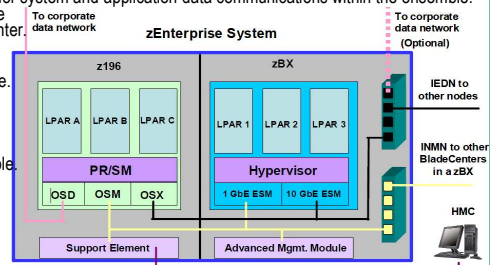


* Each node in an ensemble includes as many as **five distinct networks**, four of which are depicted in the illustration below.

1. **Intranode management network (INMN):** This private internal network provides the connections necessary to monitor and control components of the node, such as virtual servers or physical switches. The INMN connects to OSA adapters with CHPID type OSM in the z196 CPC, and to the 1Gb Ethernet Electronic Switch Module (ESM) in the zBX BladeCenters. **Note:** This network requires user definition only for its connection to z/OS or z/VM LPARs.

2. **Intraensemble data network (IEDN):** This is the network for system and application data communications within the ensemble. It connects to OSA adapters with CHPID type OSX in the z196 CPC, and to the 10 GbE ESM in the zBX BladeCenter. This network connects all nodes, including z196 and zBX frames, together. This is the network that will be virtualized for the use of the virtual servers in the ensemble.

3. **Customer managed network:** Also known as the HMC LAN, this network provides the communication link between Hardware Management Consoles (HMCs) and the nodes of the ensemble. It may connect to other System z machines that are not members of the ensemble. It may connect to support servers such as a Network Time Protocol (NTP) server that provides accurate time to the Server Time Protocol (STP) server in a System z processor. This network will be familiar to previous users of System z, and definitions on this network are unchanged from past System z implementations.



NOTE: The primary HMC and its alternate must be connected to the **same subnetwork** to allow the alternate HMC to take over the IP address of the primary HMC during failover processing.

4. **Customer managed data network:** This network represents the existing enterprise data communication network. This network is attached to Open Systems Adapters (OSAs) such as OSD, in the z196 node, just as it has been attached to previous System z machines. In addition, this network may **optionally** be connected directly to the IEDN, depending on your configuration requirements.

* The **IEDN requires user customization** before it can be used by using the **Network Virtualization Manager (NVM)**.

- The **IEDN** is the network used for **application communications** within an ensemble.

> It exists **only** within an ensemble, although it **might also** have a connection to the customer data network **outside** the ensemble (see bottom right).

- It is implemented as a **flat layer-2 network**, which means that **all the network interfaces** can communicate **directly with each other** as if they were all connected to a **single network switch**.

- **No routers** are necessary to communicate across the IEDN.

- While there are physical network switches that are part of the IEDN, the appearance of a single network is maintained through **virtualization**.

* The physical construction of the IEDN contributes to the security and reliability of the ensemble.

- All the network switches are inside the frames of the z196 and zBX frames and all network cables are point-to-point between the frames.

- With no intervening switches or routers the opportunity to compromise network integrity is greatly reduced.

- The switches are defined and configured only from the zEnterprise System firmware.

* By **virtualizing the network definitions** it is possible to **isolate** the virtual servers from the **physical definitions** of the network interfaces and devices.

- This allows the virtual servers to be **placed anywhere within the ensemble** without changing the network definitions **inside** the virtual server.

- It **isolates** the virtual servers from **"burned-in"** addresses on physical network interface cards which allows **failed** cards to be replaced without changing definitions.

* Finally, the **network provisioning** is based on the concept of virtual LANs (VLANs) which provides for multiple logical networks to be defined over the **same physical infrastructure**.

- **VLANs** are a proven method for **separating data traffic for multiple applications**, as might be required for privacy rules, regulatory requirements, and even separation of production and test communications, all flowing over the same physical network.

- This **virtualization** helps fully utilize the **physical network capacity** while still meeting your organization's security requirements.

* 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 are required to define a VLAN.

VSWITCH (Virtual switch): A virtual switch is a hypervisor component providing virtualized network resources to a virtual server.

VNIC (Virtual network interface card): The VNIC is the network resource that a virtual server uses to access the IEDN. The VNIC is defined in the hypervisor through a VSWITCH.

VMAC (Virtual media access control): Virtual MAC addresses are assigned to VNICs. The VMAC **replaces** the manufacturer's **"burned-in"** MAC address on a physical network card.

* The **physical IEDN** is connected to **all the zBX Blade Centers** in the ensemble, as well as **all the z196 CPCs**.

- Thus the physical network is **shared** by all members of the ensemble.

- A **VLAN** then provides a **logical network on top of the physical IEDN** where the virtual servers can connect using a virtual NIC which has a virtual MAC address.

- All this **virtualization** is maintained by the ensemble management **firmware** cooperating with the **hypervisors**.

- The operating systems running in the virtual servers see the **VNIC** as a real network interface into a real network. They **don't need** to be aware of the virtualization, but are able to utilize the virtualized resources.

- The many parts of the **virtualized network environment** are connection points for various resources of the ensemble. **Hypervisors** contain the **VSWITCH** definitions and **VNICs** that are contained in the **VSWITCH**.

- The virtual server and VNICs **must** be associated to the **VLAN** where they will connect.

Connection to the existing customer data network to the IEDN.

- The network configuration tasks allow specific ports on the **TOR switches** to be configured for attachment to your existing data network, **external from the ensemble**, and can impose restrictions on the attaching network.

- Configuring the **switches** for external connections you **must** consider whether to extend the **VLANs** out into the existing data network or keep them internal to the ensemble. If the IEDN switch port chosen for **external connection** is defined in **trunk mode**, then the VLAN-tagged data is **passed through** to the **external network**.

- By choosing **access mode** for the IEDN switch port it can be **restricted** to a single **VLAN** from the IEDN.

- These **decisions** will depend on your network implementation and **your network engineers should be consulted** and involved before defining the **external connection** to the IEDN.

* The **virtual servers** need IP addresses assigned just as they would as **real servers** on a real network.

- The IP addressing scheme is **not defined** in the network **virtualization** because it is a **layer-3 function**, which is built on top of the IEDN's **layer-2 structure** and it is **client's responsibility** to choose an IP addressing scheme appropriate for each of the VLANs on the IEDN.

- Either IPv4 or IPv6 (preferred) addressing can be used, **depending on the OS capability of the virtual servers**.