



# HiperSockets vs. OSA: My DB2 and Linux Connections in the Sysplex Are Taking the Wrong Routes

Linda Harrison  
IBM

14 August 2013  
Session 13634



# Trademarks



- **The following are Registered Trademarks of the International Business Machines Corporation in the United States and/or other countries.**
  - IBM
  - z/OS
- **The following are trademarks or registered trademarks of other companies.**
  - Microsoft is a registered trademark of Microsoft Corporation in the United States and other countries.
- All other products may be trademarks or registered trademarks of their respective companies.
- Refer to [www.ibm.com/legal/us](http://www.ibm.com/legal/us) for further legal information.

# Agenda

- TCP/IP Routing
- TCP/IP Session
- Example Environment and Routing Tables
- z/OS to z/OS
- Linux on System z to z/OS with Static Routing
- z/OS to Linux on System z with Static Routing
- Considerations
- Linux on System z to z/OS with Dynamic Routing
- z/OS to Linux on System z with Dynamic Routing
- Static Host Route Solutions



# TCP/IP Routing

# Physical Network



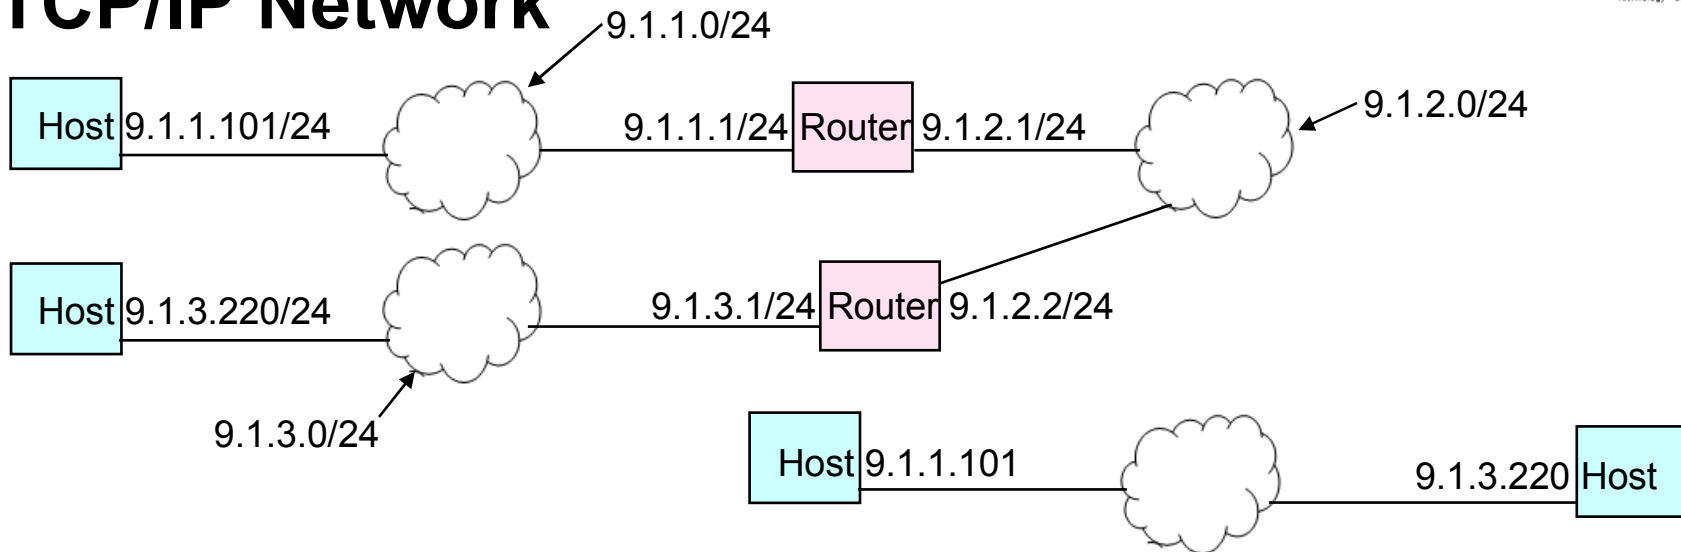
- There is some physical network.
- It can be shared between TCP/IP and SNA, and even other protocols concurrently.
- Devices are connected with cables.
- Each device network connection has a Media Access Control (MAC) address that is used for sending and receiving messages.

# Logical Network




- There are logical network views that use the underlying physical network.
  - The level of detail depends upon the discussion.
- Each device network connection has a Logical address that is used for sending and receiving messages.

# TCP/IP Network



- IPv4 Address has 32 bits 1111 1111 1111 1111 1111 1111 1111 1111
- Hex calculation 9.1.1.101 = 0000 1001 0000 0001 0000 0001 0110 0101
- Subnet Mask 24 indicates that the first 24 bits identify the subnetwork and only the last 8 bits identify the host address.
- 24 = 1111 1111 1111 1111 1111 1111 0000 0000
  - Also referred to as 255.255.255.0

# TCP/IP Routing



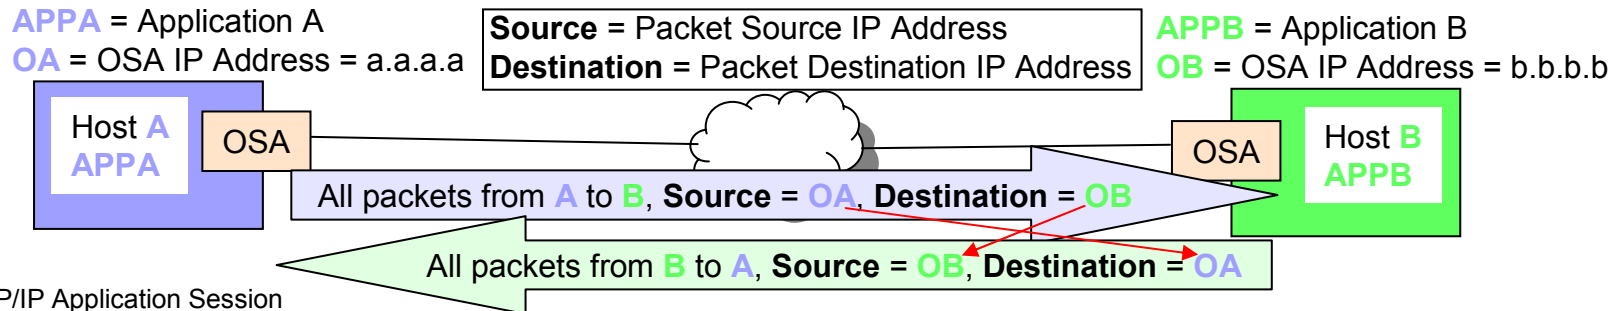
Host A:  
 Direct Route to 9.1.1.0/24  
 Indirect Route to 9.1.3.0/24 via 9.1.1.1  
 Default Route to 9.1.1.1

- Direct Routes
  - Subnetworks that this host connects to.
- Indirect Routes
  - Subnetworks that this host can reach by routing through a router.
- Default Routes
  - Where to send messages when there is no explicit route.



# TCP/IP Session

# Source and Destination IP Addresses



- TCP/IP Application Session
  - One partner (Host A) initiates the connection and sends connection request packet. For TCP connections the host that sends the connection request is the client.
    - What destination IP address to connect to?
      - *If hostname is given then IP address is resolved by DNS or Local Host Name file (IPNodes).*
        - *Result may be cached.*
    - What source IP address to use in the initiate packet?
      - *Hierarchy of source IP address is detailed in z/OS Communication Server (CS) IP Configuration Guide, SC31-8775.*
        1. *Sendmsg( ) using the IPV6\_PKTINFO ancillary option specifying a nonzero source address (RAW and UDP sockets only)*
        2. *Setsockopt( ) IPV6\_PKTINFO option specifying a nonzero source address (RAW and UDP sockets only)*
        3. *Explicit bind to a specific local IP address*
        4. *bind2addrsel socket function (AF\_INET6 sockets only)*
        5. *PORT profile statement with the BIND parameter*
        6. *SRCIP profile statement (TCP connections only)*
        7. *TCPSTACKSOURCEVIPA parameter on the IPCONFIG or IPCONFIG6 profile statement (TCP connections only)*
        8. *SOURCEVIPA: Static VIPA address from the HOME list or from the SOURCEVIPAINTERFACE parameter*
        9. *HOME IP address of the link over which the packet is sent*
  - The other partner (Host B) receives the connection request packet and responds. For TCP connections the host that receives the connection request is the server.
    - The host that receives the connection request takes the source IP address from the connection packet and uses that for the destination IP address in the packet that it sends back.
    - The host that receives the connection request takes the destination IP address from the connection packet and uses that for the source IP address in the packet that it sends back.
    - These source and destination assignments are usually used for the life of the connection.

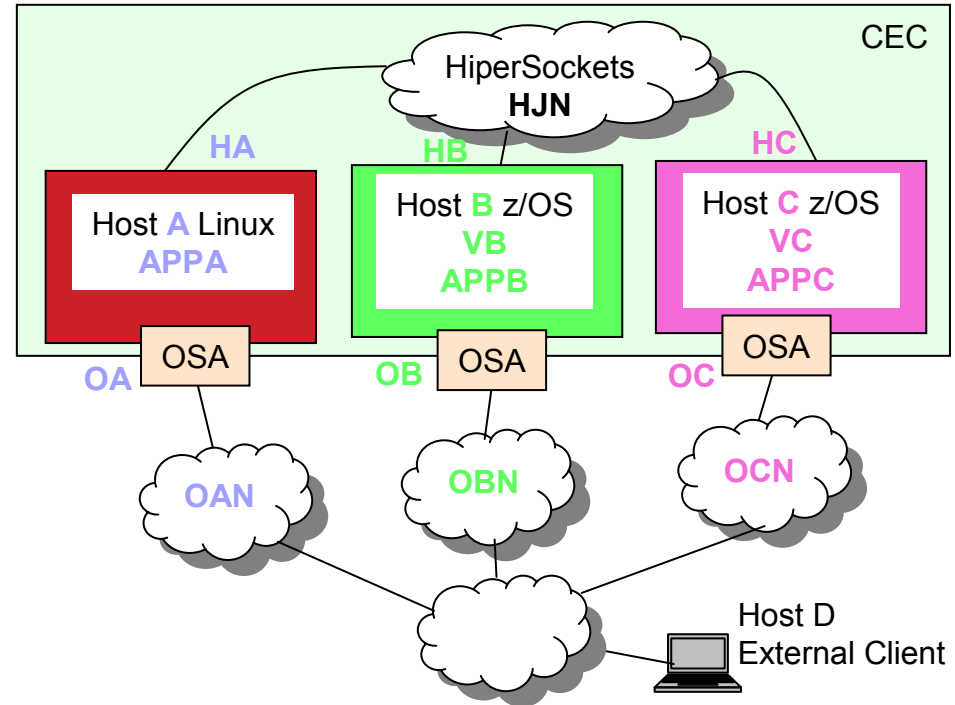
# Example Environment and Routing Tables

# Linux on System z and z/OS Hosts

**APPA** = Application A  
**OA** = OSA IP Address = 10.1.190.51  
**OAN** = OSA Network = 10.1.190.0/24  
**HA** = HiperSockets IP Address = 10.20.1.51  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24

**APPB** = Application B  
**VB** = VIPA IP Address = 10.4.1.101  
**OB** = OSA IP Address = 10.10.4.101  
**OBN** = OSA Network = 10.10.4.0/24  
**HB** = HiperSockets IP Address = 10.20.1.101  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24

**APPC** = Application C  
**VC** = VIPA IP Address = 10.3.1.202  
**OC** = OSA IP Address = 10.10.3.202  
**OCN** = OSA Network = 10.10.3.0/24  
**HC** = HiperSockets IP Address = 10.20.1.202  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24



- Linux on System z typically only has a single OSA Address.
  - z/VM is able to provide multiple OSA connectivity via guest LAN.
  - When HiperSockets is added to the Linux environment a choice must be made to route traffic over the OSA or over the HiperSockets.
- z/OS handles the choice of OSA or HiperSockets with VIPA.
  - VIPA and dynamic routing provide a single IP Address (VIPA) that may be used as source or destination IP Address.
    - HiperSockets or OSA are chosen depending on dynamic routing weights.
    - HiperSockets may be used between hosts on the same CEC concurrently while OSA is used for non-CEC partners.

# Routing Table

- z/OS B Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

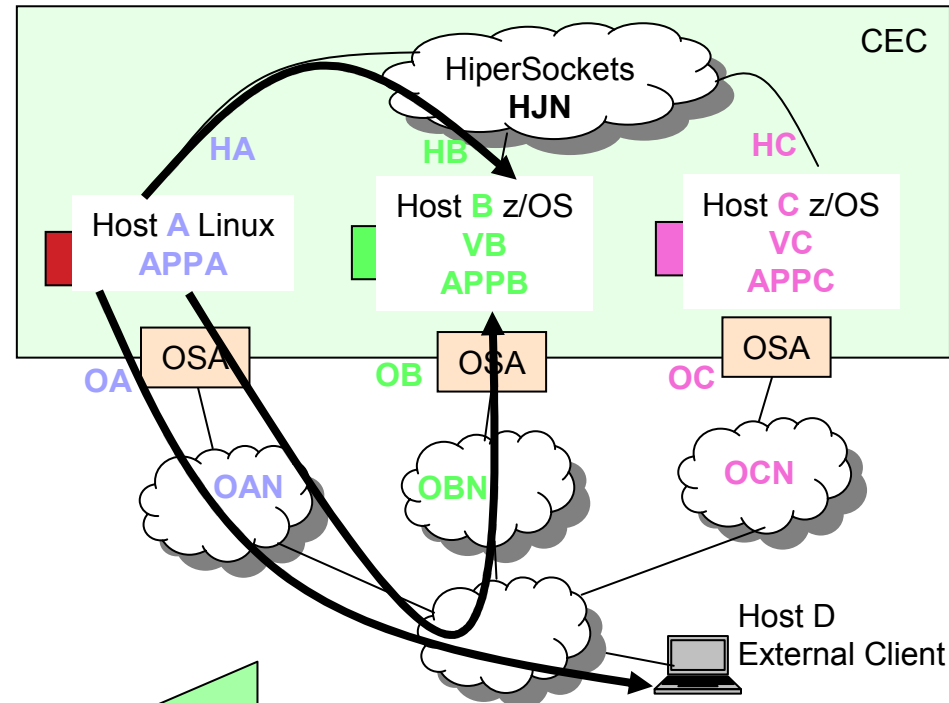
- Subnet mask of 32 indicates Host Routes, the most specific type of route possible. More specific routes always take precedence over less specific routes.
- Subnet or Network routes, in this case subnet route with mask of 24, are less specific than host routes but more specific than default routes.
- Default routes are the least specific routes possible and are therefore only used when no other route matches the destination.

- The destination IP Address, IP subnet, or IP network is in the first column above.
  - DEFAULT is a special destination keyword indicating where to send packets when the destination does not match any other route in the table.
- Gateway IP Address is in the second column above. A zero, “0”, indicates there is no next hop gateway.
  - A route without a next hop gateway is known as a direct route.
  - A route with a next hop gateway is known as an indirect route.
- The interface name or link name is in the third column above to indicate which interface the packet is to be sent over.
- Each route in the table may be read as “If this destination, then send packet to this gateway over this interface”.

# Host A Linux on System z Routing Table

**APPA** = Application A  
**OA** = OSA IP Address = 10.1.190.51  
**OAN** = OSA Network = 10.1.190.0/24  
**HA** = HiperSockets IP Address = 10.20.1.51  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24

**APPB** = Application B  
**VB** = VIPA IP Address = 10.4.1.101  
**OB** = OSA IP Address = 10.10.4.101  
**OBN** = OSA Network = 10.10.4.0/24  
**HB** = HiperSockets IP Address = 10.20.1.101  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24



- Host A Linux Routing Table
 

Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

Linux routing display might be in a little different format than what I have depicted here (like z/OS), but you should get the general idea of the routes I am trying to explain.

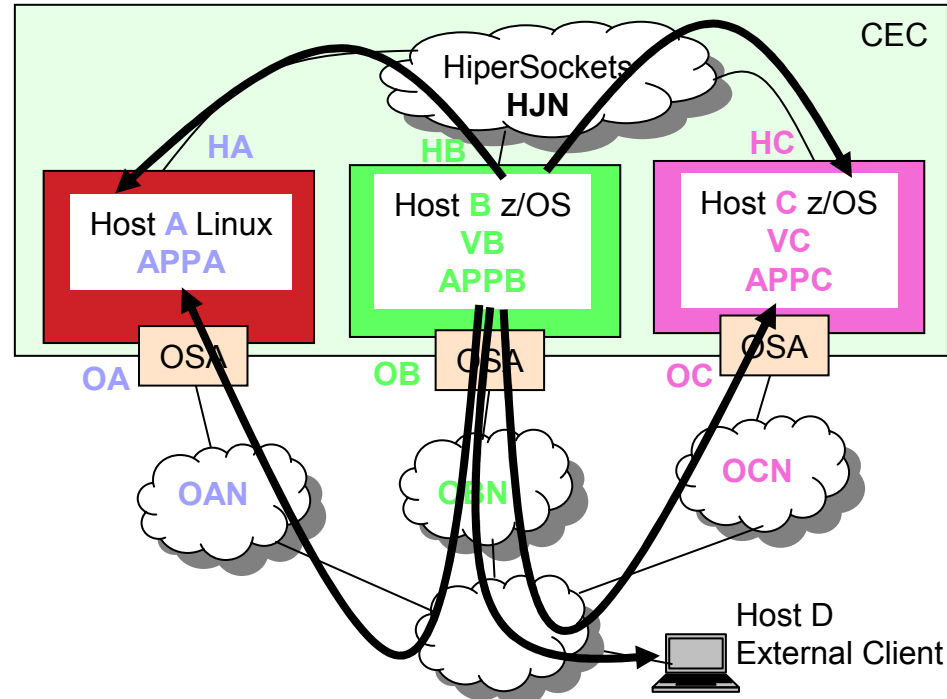
- This presentation will only discuss three session paths for Host A sending packets:
  - To Host B, Host A may route over the OSA or HiperSockets.
  - To Host D, Host A can only route over the OSA.

# Host B z/OS Routing Table

**APPA** = Application A  
**OA** = OSA IP Address = 10.1.190.51  
**OAN** = OSA Network = 10.1.190.0/24  
**HA** = HiperSockets IP Address = 10.20.1.51  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24

**APPB** = Application B  
**VB** = VIPA IP Address = 10.4.1.101  
**OB** = OSA IP Address = 10.10.4.101  
**OBN** = OSA Network = 10.10.4.0/24  
**HB** = HiperSockets IP Address = 10.20.1.101  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24

**APPC** = Application C  
**VC** = VIPA IP Address = 10.3.1.202  
**OC** = OSA IP Address = 10.10.3.202  
**OCN** = OSA Network = 10.10.3.0/24  
**HC** = HiperSockets IP Address = 10.20.1.202  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24



- Host B z/OS Routing Table


Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

- This presentation will discuss five session paths for Host B sending packets:
  - To send to Host A, Host B may route over the OSA or HiperSockets.
  - To send to Host C, Host B may route over the OSA or HiperSockets.
  - To send to Host D, Host A can only route over the OSA.

# Host C z/OS Routing Table

**APPB** = Application B  
**VB** = VIPA IP Address = 10.4.1.101  
**OB** = OSA IP Address = 10.10.4.101  
**OBN** = OSA Network = 10.10.4.0/24  
**HB** = HiperSockets IP Address = 10.20.1.101  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24

**APPC** = Application C  
**VC** = VIPA IP Address = 10.3.1.202  
**OC** = OSA IP Address = 10.10.3.202  
**OCN** = OSA Network = 10.10.3.0/24  
**HC** = HiperSockets IP Address = 10.20.1.202  
**HJN** = HiperSockets Joint Network = 10.20.1.0/24



- Host C z/OS Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
10.20.1.0/24	0	HC
DEFAULT	10.10.3.1	OC

- This presentation will only discuss two session paths for Host C sending packets:
  - To Host B, Host C may route over the OSA or HiperSockets.



# z/OS to z/OS

# Destination and Source are OSA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

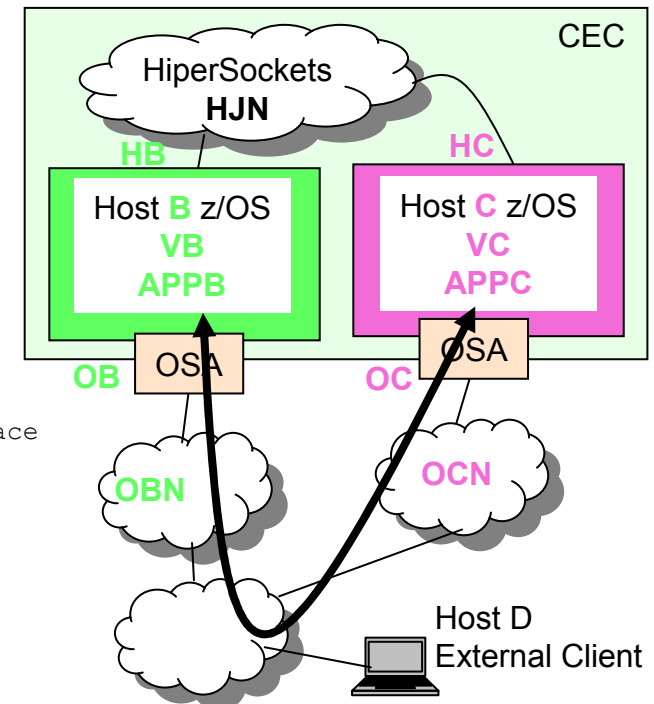
Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
10.20.1.0/24	0	HC
→ <b>DEFAULT</b>	<b>10.10.3.1</b>	<b>OC</b>



Source = OB, Destination = OC

Source = OC, Destination = OB

- If APPB sends a connection request to APPC at OC, and source IP Address of OB is selected, then packets will flow over OSAs.

# Destination is OSA and Source is HiperSockets

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table



Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
→ <b>10.20.1.0/24</b>	<b>0</b>	<b>HC</b>
DEFAULT	10.10.3.1	OC



- If APPB sends a connection request to APPC at OC, and source IP Address of HB is selected, then packets from B to C will flow over OSAs, and packets from C to B will flow over HiperSockets.

# Destination is OSA and Source is VIPA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table



Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
→ <b>10.4.1.101/32</b>	<b>10.20.1.101</b>	<b>HC</b>
10.10.3.0/24	0	OC
10.20.1.0/24	0	HC
DEFAULT	10.10.3.1	OC



- If APPB sends a connection request to APPC at OC, and source IP Address of VB is selected, then packets from B to C will flow over OSAs, and packets from C to B will flow over HiperSockets.

# Destination is HiperSockets and Source is OSA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

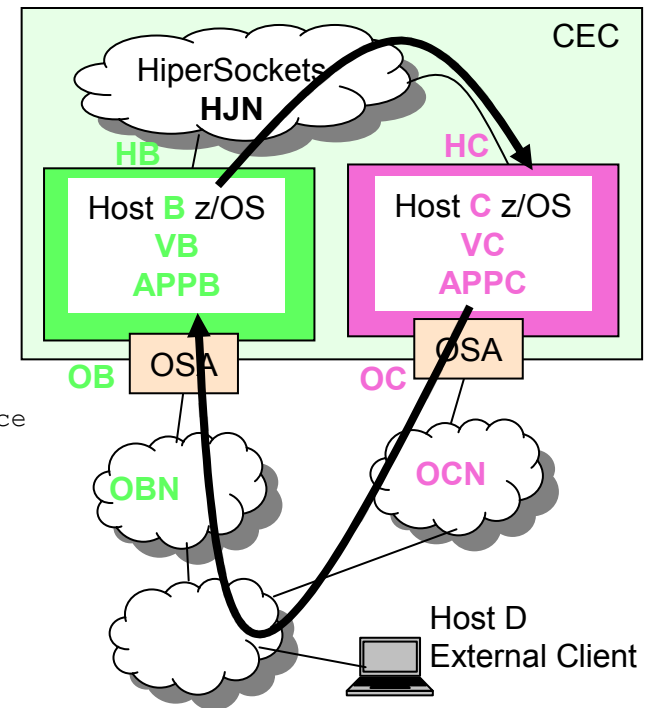
Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ 10.20.1.0/24	0	<b>HB</b>
DEFAULT	10.10.4.1	OB

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
10.20.1.0/24	0	HC
→ DEFAULT	10.10.3.1	<b>OC</b>



Source = **OB**, Destination = **HC**

Source = **HC**, Destination = **OB**

- If APPB sends a connection request to APPC at HC, and source IP Address of OB is selected, then packets from B to C will flow over OSAs, and packets from C to B will flow over HiperSockets.

# Destination and Source are HiperSockets

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table



Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ 10.20.1.0/24	0	<b>HB</b>
DEFAULT	10.10.4.1	OB

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
→ 10.20.1.0/24	0	<b>HC</b>
DEFAULT	10.10.3.1	OC



- If APPB sends a connection request to APPC at HC, and source IP Address of HB is selected, then packets between B and C will flow over HiperSockets.

# Destination is HiperSockets and Source is VIPA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table



Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ 10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
→ 10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
10.20.1.0/24	0	HC
DEFAULT	10.10.3.1	OC



- If APPB sends a connection request to APPC at HC, and source IP Address of VB is selected, then packets between B and C will flow over HiperSockets.

# Destination is VIPA and Source is OSA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table



Destination	Gateway	Interface
→ 10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
10.20.1.0/24	0	HC
→ DEFAULT	10.10.3.1	OC



- If APPB sends a connection request to APPC at VC, and source IP Address of OB is selected, then packets from B to C will flow over HiperSockets, and packets from C to B will flow over OSAs.



# Destination is VIPA and Source is HiperSockets

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table



Destination	Gateway	Interface
→ 10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
→ 10.20.1.0/24	0	HC
DEFAULT	10.10.3.1	OC



- If APPB sends a connection request to APPC at VC, and source IP Address of HB is selected, then packets between B and C will flow over HiperSockets.

# Destination and Source are VIPA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

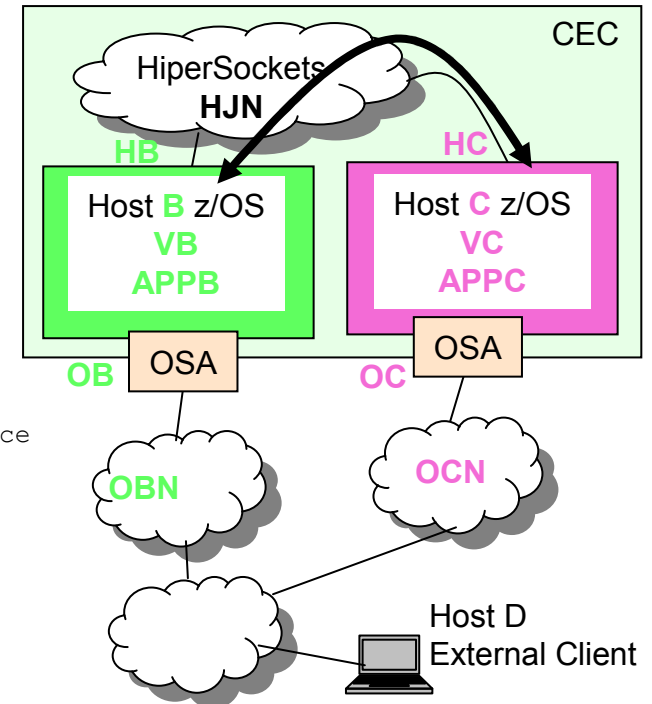

Destination	Gateway	Interface
→ 10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

## Host C

**VC** = 10.3.1.202  
**OC** = 10.10.3.202  
**OCN** = 10.10.3.0/24  
**HC** = 10.20.1.202  
**HJN** = 10.20.1.0/24

### Host C z/OS Routing Table

Destination	Gateway	Interface
→ 10.4.1.101/32	10.20.1.101	HC
10.10.3.0/24	0	OC
10.20.1.0/24	0	HC
DEFAULT	10.10.3.1	OC



- If APPB sends a connection request to APPC at VC, and source IP Address of VB is selected, then packets between B and C will flow over HiperSockets.

# z/OS to External Host

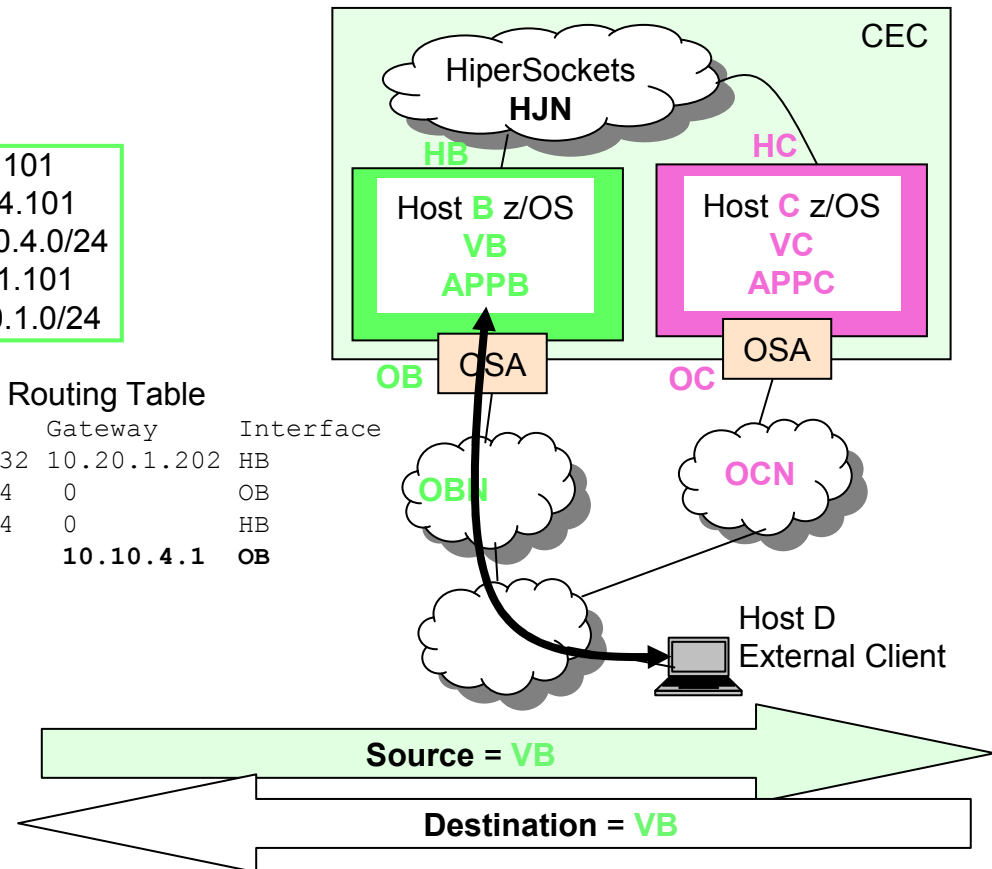
- If APPB sends a connection request to an application on Host D, what source IP Address will be chosen?
- Typically an application picks the same source IP Address no matter what partner it is connecting to.
- You must keep this in mind if a z/OS application is going to initiate connections to same CEC partners as well as external partners.
- If APPB sends a connection request and source IP Address of VB is selected then packets from B to D will flow over OSA (using the default route) and packets from D to B will flow over OSA (assuming the routers between B and D have static or dynamic routes that indicate that packets to VB should be routed to OB.)

## Host B

**VB = 10.4.1.101**  
**OB = 10.10.4.101**  
**OBN = 10.10.4.0/24**  
**HB = 10.20.1.101**  
**HJN = 10.20.1.0/24**

## Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>



# z/OS Static Routing

- Static Routing wasn't even covered between z/OS's but before we summarize the Dynamic Routing results here are the Static Routing results:

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
z/OS	z/OS	OSA	OSA	No	No	Yes
z/OS	z/OS	HiperSockets	OSA	No	No	No
z/OS	z/OS	VIPA	OSA	No	No	Yes <sup>1</sup>
z/OS	z/OS	OSA	HiperSockets	No	No	Yes
z/OS	z/OS	HiperSockets	HiperSockets	No	Yes	No
z/OS	z/OS	VIPA	HiperSockets	No	No <sup>2</sup>	Yes <sup>1</sup>
z/OS	z/OS	OSA	VIPA	No	No	Yes
z/OS	z/OS	HiperSockets	VIPA	No	No	No
z/OS	z/OS	VIPA	VIPA	No	No <sup>2</sup>	Yes <sup>1</sup>

<sup>1</sup> Requires route on network routers to reach VIPA over OSA.

<sup>2</sup> Assumes no static host route to reach VIPA.

- Note:
  - There is no way to have same CEC traffic over HiperSockets and external traffic over OSA using a single Source IP Address for same CEC partners as well as external partners.
  - There is no way to have same CEC traffic over HiperSockets and external traffic over OSA using a single Destination IP Address from same CEC partners as well as external partners.


# z/OS Dynamic Routing

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
z/OS	z/OS	OSA	OSA	Yes	No	Yes
z/OS	z/OS	HiperSockets	OSA	Yes	No <sup>3</sup>	No
z/OS	z/OS	VIPA	OSA	Yes	No <sup>3&amp;4</sup>	Yes
z/OS	z/OS	OSA	HiperSockets	Yes	No <sup>3</sup>	Yes
z/OS	z/OS	HiperSockets	HiperSockets	Yes	Yes	No
z/OS	z/OS	VIPA	HiperSockets	Yes	Yes <sup>4</sup>	Yes
z/OS	z/OS	OSA	VIPA	Yes	No <sup>3&amp;4</sup>	Yes
z/OS	z/OS	HiperSockets	VIPA	Yes	Yes <sup>4</sup>	No
<b>z/OS</b>	<b>z/OS</b>	<b>VIPA</b>	<b>VIPA</b>	<b>Yes</b>	<b>Yes<sup>4</sup></b>	<b>Yes</b>

<sup>3</sup> Not included here, but if OSA is advertised over HiperSockets and the cost is less to send over HiperSockets, then it is possible that packets destined for the OSA IP Address will be sent over HiperSockets.

<sup>4</sup> If the cost is less to send over HiperSockets then it is possible that packets destined for VIPA will be sent over HiperSockets.

## z/OS to z/OS Conclusion

- VIPA and Dynamic routing
  - A single z/OS source IP Address, VIPA, may be used for both same CEC partners as well as external partners.
  - A single z/OS destination IP Address, VIPA, may be used on both same CEC partners as well as external partners.
-  It is possible to use HiperSockets between same CEC partners and OSA with non-CEC partners concurrently.
  - OSA between same CEC partners may hypothetically be used as backup in the event that the HiperSockets fails. “Hypothetically” since a HiperSockets failure may indicate serious issues on the CEC.

# Linux on System z to z/OS with Static Routing

# Destination and Source are OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table


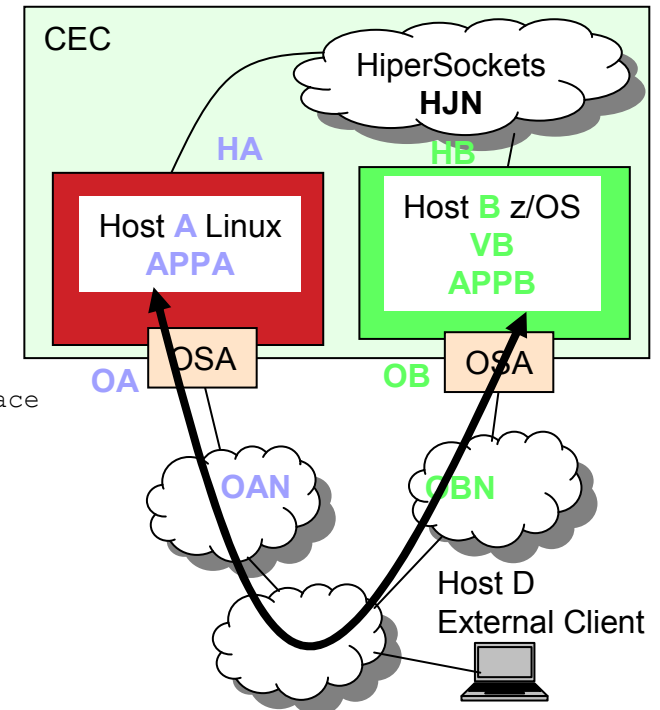
Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>



- If APPA sends a connection request to APPB at OB, and source IP Address of OA is selected, then packets will flow over OSAs.



# Destination is OSA and Source is HiperSockets

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ <b>10.20.1.0/24</b>	<b>0</b>	<b>HB</b>
DEFAULT	10.10.4.1	OB



- If APPA sends a connection request to APPB at OB, and source IP Address of HA is selected, then packets from A to B will flow over OSAs, and packets from B to A will flow over HiperSockets.

# Destination is HiperSockets and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
→ 10.20.1.0/24	0	<b>HA</b>
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ DEFAULT	10.10.4.1	OB



- If APPA sends a connection request to APPB at HB, and source IP Address of OA is selected, then packets from A to B will flow over HiperSockets, and packets from B to A will flow over OSAs.

# Destination and Source are HiperSockets

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
→ 10.20.1.0/24	0	<b>HA</b>
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ 10.20.1.0/24	0	<b>HB</b>
DEFAULT	10.10.4.1	OB



- If APPA sends a connection request to APPB at HB, and source IP Address of HA is selected, then packets between A to B will flow over HiperSockets.

# Destination is VIPA and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>



- If APPA sends a connection request to APPB at VB, and source IP Address of OA is selected, then packets between A and B will flow over OSAs.

# Destination is VIPA and Source is HiperSockets

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

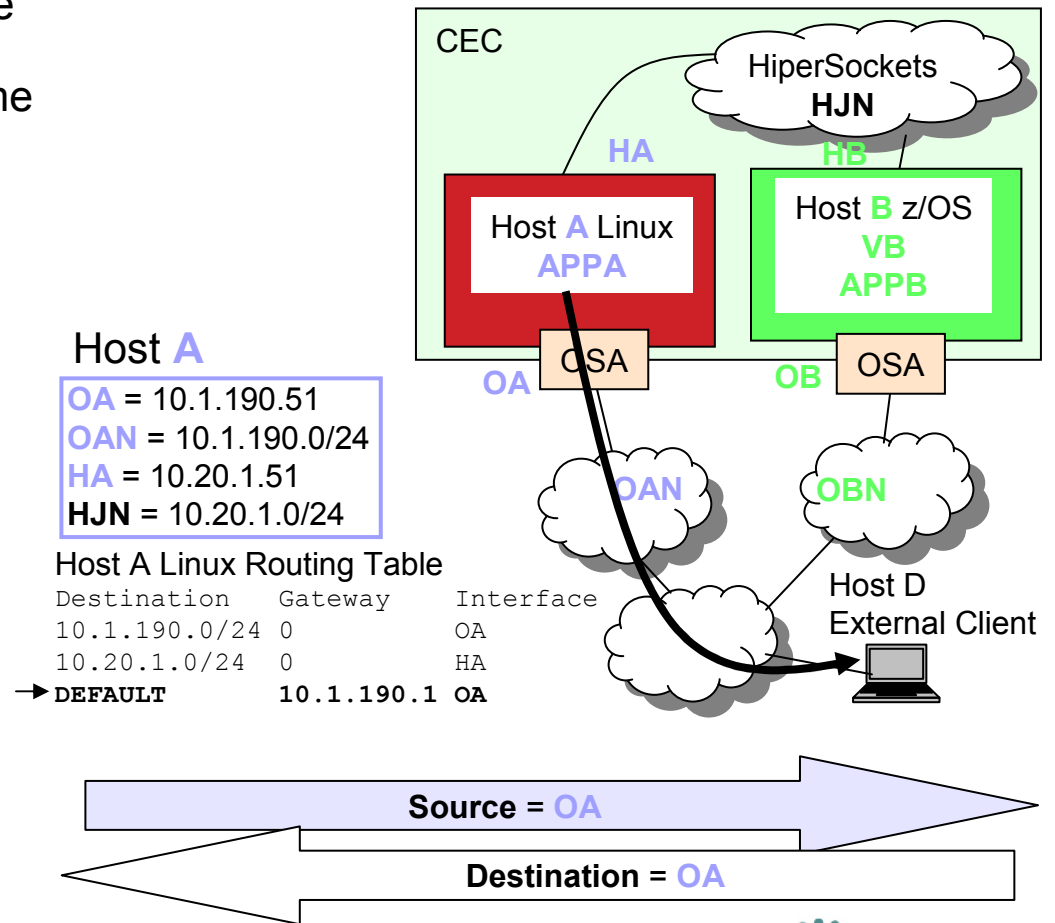
Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ <b>10.20.1.0/24</b>	<b>0</b>	<b>HB</b>
DEFAULT	10.10.4.1	OB



- If APPA sends a connection request to APPB at VB, and source IP Address of HA is selected, then packets from A to B will flow over OSAs, and packets from B to A will flow over HiperSockets.

# Linux on System z to External Host

- If APPA sends a connection request to an application on Host D, what source IP Address will be chosen?
- Typically an application picks the same source IP Address no matter what partner it is connecting to.
- You must keep this in mind if a Linux on System z application is going to initiate connections to same CEC partners as well as external partners.
- The OSA had better be the source IP Address that is chosen because if HiperSockets is chosen the external host might not be able to send a response back.



# Static Routing from Linux on System z

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
Linux	z/OS	OSA	OSA	No	No	Yes
Linux	z/OS	HiperSockets	OSA	No	No	No
Linux	z/OS	OSA	HiperSockets	No	No	Yes
Linux	z/OS	HiperSockets	HiperSockets	No	Yes	No
Linux	z/OS	OSA	VIPA	No	No	Yes
Linux	z/OS	HiperSockets	VIPA	No	No <sup>2</sup>	No

<sup>2</sup> Assumes no static host route to reach VIPA.

- **Problem:**

- When using z/OS Client and Linux Server, without VIPA on Linux, and with Static Routing,
- How to have same CEC traffic over HiperSockets and external traffic over OSA using a single Source IP Address for same CEC partners as well as external partners?

# z/OS to Linux on System z with Static Routing



# Destination and Source are OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>



- If APPB sends a connection request to APPA at OA, and source IP Address of OB is selected, then packets will flow over OSAs.

# Destination is OSA and Source is HiperSockets

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
→ 10.20.1.0/24	0	<b>HA</b>
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at OA, and source IP Address of HB is selected, then packets from B to A will flow over OSAs, and packets from A to B will flow over HiperSockets.

# Destination is OSA and Source is VIPA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
→ <b>DEFAULT</b>	<b>10.10.4.1</b>	<b>OB</b>



- If APPB sends a connection request to APPA at OA, and source IP Address of VB is selected, then packets between A and B will flow over OSAs.

# Destination is HiperSockets and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ <b>10.20.1.0/24</b>	<b>0</b>	<b>HB</b>
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at HA, and source IP Address of OB is selected, then packets from B to A will flow over HiperSockets, and packets from A to B will flow over OSAs.

# Destination and Source are HiperSockets

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table

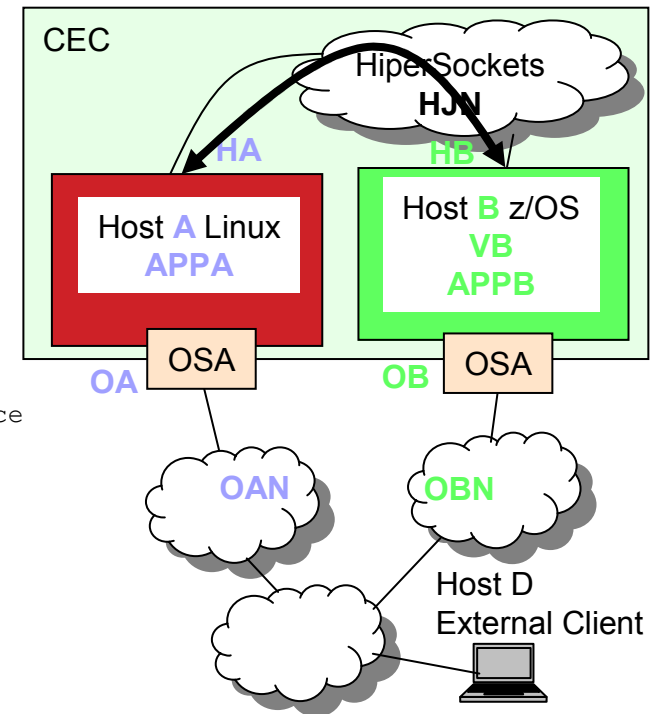

Destination	Gateway	Interface
10.1.190.0/24	0	OA
→ 10.20.1.0/24	0	<b>HA</b>
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ 10.20.1.0/24	0	<b>HB</b>
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at HA, and source IP Address of HB is selected, then packets between A and B will flow over HiperSockets.

# Destination is HiperSockets and Source is VIPA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
→ <b>DEFAULT</b>	<b>10.1.190.1</b>	<b>OA</b>

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
10.3.1.202/32	10.20.1.202	HB
10.10.4.0/24	0	OB
→ <b>10.20.1.0/24</b>	<b>0</b>	<b>HB</b>
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at HA, and source IP Address of VB is selected, then packets from B to A will flow over HiperSockets, and packets from A to B will flow over OSAs.

# Static Routing to Linux on System z

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
z/OS	Linux	OSA	OSA	No	No	Yes
z/OS	Linux	HiperSockets	OSA	No	No	No
z/OS	Linux	VIPA	OSA	No	No	Yes <sup>1</sup>
z/OS	Linux	OSA	HiperSockets	No	Yes	Yes
z/OS	Linux	HiperSockets	HiperSockets	No	No	No
z/OS	Linux	VIPA	HiperSockets	No	No <sup>2</sup>	Yes <sup>1</sup>

<sup>1</sup> Requires route on network routers to reach VIPA over OSA.

<sup>2</sup> Assumes no static host route to reach VIPA.

- Note:
  - There is no way to have same CEC traffic over HiperSockets and external traffic over OSA using a single Destination IP Address from same CEC partners as well as external partners.

## Linux on System z Static Routing Conclusion

- If a single source IP Address is desired for both same CEC partners and external partners then OSA must be used.
- If a single destination IP Address is desired by same CEC partners and external partners then OSA must be used.
- In order to use HiperSockets between same CEC partners and OSA with non-CEC partners:
  - Use HiperSockets IP Address as source IP Address to same CEC partners.
  - Use HiperSockets IP Address as destination IP Address by same CEC partners.
  - Use OSA IP Address as source IP Address to external partners.
  - Use OSA IP Address as destination IP Address by external partners.

### • **Problem:**

- When using Linux Client and z/OS Server, without VIPA on Linux, and with Static Routing,
- How to have same CEC traffic over HiperSockets and external traffic over OSA using a single Source IP Address for same CEC partners as well as external partners?



# Considerations

## Source and Destination IP Addresses

- Does the application use the same source IP Address for same CEC partners as well as non-CEC partners?
  - Does the application bind to an IP Address?
  - Does the application use the stack host name for source?
- For z/OS initiated sessions, looking back at the Hierarchy of source IP address, items 6 and 9 might be of interest.
  6. SRCIP profile statement (TCP connections only)
    - Does the application match any items 1 through 5?
    - If not then using SRCIP you may identify a different source IP Address depending upon what destination IP Address is used.
  9. HOME IP address of the link over which the packet is sent
    - Does the application match any items 1 through 9?
    - If not then the IP address of the interface that is chosen to send the packet over will cause a different source IP to be chosen depending upon what the destination IP Address is (the HiperSockets IP Address or the non-CEC partner IP Address).
- Is this useful?
  - I don't think this is very useful since using VIPA as a source IP Address enables the same thing.

Which partner initiates the connection and what source and destination IP Address the application chooses is key to creating a solution!

# Linux on System z to z/OS with Dynamic Routing

## Dynamic Routing Costs

- With Dynamic Routing it is possible that the OSA networks are advertised over HiperSockets. If the HiperSockets path cost is less than the OSA path then the packets will flow over HiperSockets.
- This routing is not always possible or desirable. What if the Linux on System z OSA and z/OS OSA are in the same subnet? What if there are other destinations in the same subnet as the partner OSA? You probably would not want traffic to initiate from Linux on System z and traverse z/OS to reach a destination host out on the z/OS OSA network.
- However for the purpose of this example, let us assume that you are able to design an OSPF routing network so that the HiperSockets are the least cost route for the OSA destination between LPARs.

# Destination and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HA
→ 10.10.4.0/24	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
→ 10.1.190.0/24	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB



- If APP A sends a connection request to APP B at OB, and source IP Address of OA is selected, then packets between A and B will flow over HiperSockets.

# Destination is HiperSockets and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HA
10.10.4.0/24	10.20.1.101	HA
10.1.190.0/24	0	OA
→ 10.20.1.0/24	0	<b>HA</b>
DEFAULT	10.1.190.1	OA

## Host B

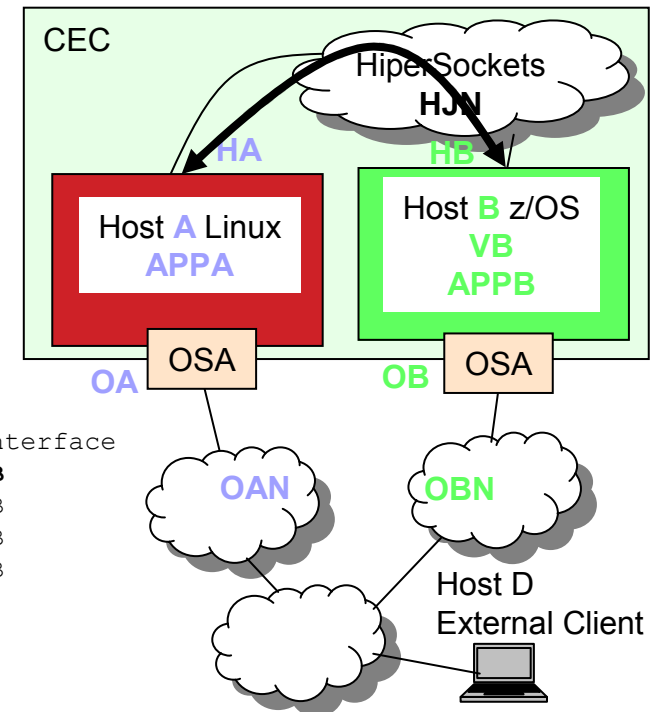
**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
→ 10.1.190.0/24	10.20.1.51	<b>HB</b>
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

Source = **OA**, Destination = **HB**

Source = **HB**, Destination = **OA**



- If APP A sends a connection request to APP B at HB, and source IP Address of OA is selected, then packets between A and B will flow over HiperSockets.

# Destination is VIPA and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table

Destination	Gateway	Interface
→ 10.4.1.101/32	10.20.1.101	HA
10.10.4.0/24	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

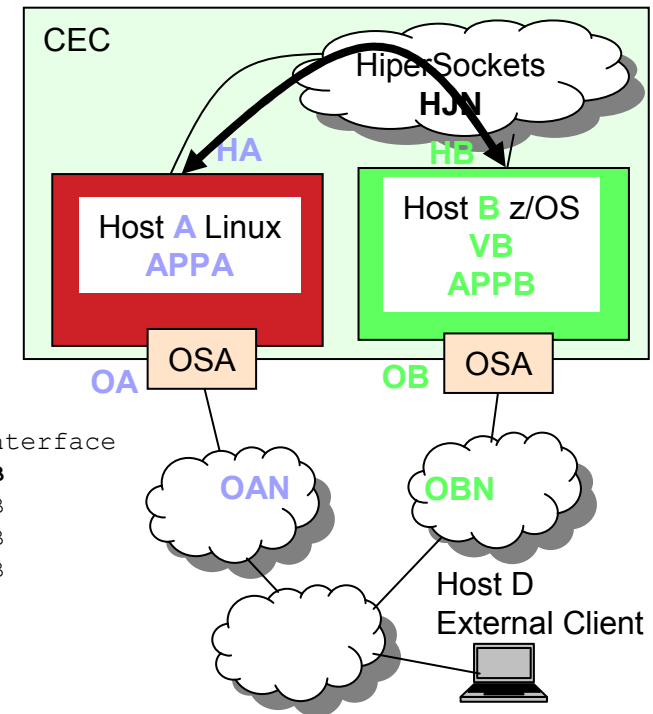
**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
→ 10.1.190.0/24	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

Source = OA, Destination = VB

Source = VB, Destination = OA



- If APP A sends a connection request to APP B at VB, and source IP Address of OA is selected, then packets between A and B will flow over HiperSockets.

# Dynamic Routing from Linux on System z

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
Linux	z/OS	OSA	OSA	Yes	Yes <sup>3</sup>	Yes
Linux	z/OS	HiperSockets	OSA	Yes	Yes <sup>3</sup>	No
Linux	z/OS	OSA	HiperSockets	Yes	Yes <sup>3</sup>	Yes
Linux	z/OS	HiperSockets	HiperSockets	Yes	Yes	No
Linux	z/OS	OSA	VIPA	Yes	Yes <sup>3&amp;4</sup>	Yes
Linux	z/OS	HiperSockets	VIPA	Yes	Yes <sup>4</sup>	No

<sup>3</sup> If OSA is advertised over HiperSockets and the cost is less to send over HiperSockets then it is possible that packets destined for OSA IP Address will be sent over HiperSockets.

<sup>4</sup> If the cost is less to send over HiperSockets then it is possible that packets destined for VIPA will be sent over HiperSockets.

• Note:

- It is possible to have same CEC traffic over HiperSockets and external traffic over OSA using a single Source IP Address for same CEC partners as well as external partners.
- For same CEC partners it is possible to route packets over HiperSockets even though the destination IP Address in the packet is the OSA or VIPA.

## • Little Used Solution:

- As stated in the beginning of this section, this routing is not always possible or desirable.
- When using Linux on System z Client and z/OS Server, with Dynamic Routing, it is possible to have same CEC traffic over HiperSockets and external traffic over OSA.



# z/OS to Linux on System z with Dynamic Routing

# Destination and Source is OSA

## Host A

OA = 10.1.190.51  
 OAN = 10.1.190.0/24  
 HA = 10.20.1.51  
 HJN = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HA
→ 10.10.4.0/24	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

VB = 10.4.1.101  
 OB = 10.10.4.101  
 OBN = 10.10.4.0/24  
 HB = 10.20.1.101  
 HJN = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
→ 10.1.190.0/24	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at OA, and source IP Address of OB is selected, then packets between A and B will flow over HiperSockets.

# Destination is OSA and Source is VIPA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
→ 10.4.1.101/32	10.20.1.101	HA
10.10.4.0/24	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
→ 10.1.190.0/24	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at OA, and source IP Address of VB is selected, then packets between A and B will flow over HiperSockets.

# Dynamic Routing to Linux on System z

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
z/OS	Linux	OSA	OSA	Yes	Yes <sup>3</sup>	Yes
z/OS	Linux	HiperSockets	OSA	Yes	Yes <sup>3</sup>	No
z/OS	Linux	VIPA	OSA	Yes	Yes <sup>3&amp;4</sup>	Yes
z/OS	Linux	OSA	HiperSockets	Yes	Yes <sup>3</sup>	Yes
z/OS	Linux	HiperSockets	HiperSockets	Yes	Yes	No
z/OS	Linux	VIPA	HiperSockets	Yes	Yes <sup>4</sup>	Yes

<sup>3</sup> If OSA is advertised over HiperSockets and the cost is less to send over HiperSockets then it is possible that packets destined for OSA IP Address will be sent over HiperSockets.

<sup>4</sup> If the cost is less to send over HiperSockets then it is possible that packets destined for VIPA will be sent over HiperSockets.

- Note:
  - It is possible to have same CEC traffic over HiperSockets and external traffic over OSA using a single Destination IP Address from same CEC partners as well as external partners.
  - For some CEC partners it is possible to route packets over HiperSockets even though the destination IP Address in the packet is the OSA or VIPA.

## • Little Used Solution:

- Again, this routing is not always possible or desirable.
- When using z/OS Client and Linux on System z Server, with Dynamic Routing, it is possible to have same CEC traffic over HiperSockets and external traffic over OSA.

## OSPF Stub Area

- When using OSPF on z/OS it is recommended that z/OS be configured as part of an OSPF Stub Area.
  - Depending upon the size of the OSPF network, making z/OS part of an OSPF Stub Area can significantly reduce the amount of cycles used for OSPF routing.
- When using HiperSockets and OSA between z/OS and Linux on System z it is recommended that z/OS and Linux on System z be in the same Stub Area.
  - Since a Stub Area only uses default routing outside of the Stub Area, it is not possible to use HiperSockets between same CEC partners concurrently with OSA between non-same CEC partners, unless same CEC partners are in the same Stub Area.

# Static Host Route Solutions

## Reasons for Static Host Routes

- Do you want to use Static Routing instead of Dynamic Routing between Linux on System z and z/OS?
- Do you want to use a single Linux on System z source IP Address for same CEC partners as well as external partners?
- Do want to use a single Linux on System z destination IP Address on same CEC partners as well as external partners?
- For the purpose of this example, let us assume that you have created static host routes for the OSA and VIPA Addresses to route them over the HiperSockets.

# Destination and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table

Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HA
★ → 10.10.4.101/32	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B


**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
★ → 10.1.190.51/32	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

Source = OA, Destination = OB

Source = OB, Destination = OA



- If APP A sends a connection request to APP B at OB, and source IP Address of OA is selected, then packets between A and B will flow over HiperSockets.



# Destination is HiperSockets and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table


Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HA
10.10.4.101/32	10.20.1.101	HA
10.1.190.0/24	0	OA
★ → 10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
★ → 10.1.190.51/32	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB



Source = OA, Destination = HB

Source = HB, Destination = OA

- If APP A sends a connection request to APP B at HB, and source IP Address of OA is selected, then packets between A and B will flow over HiperSockets.

# Destination is VIPA and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table

Destination	Gateway	Interface
★ → 10.4.1.101/32	10.20.1.101	HA
10.10.4.101/32	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B


**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
★ → 10.1.190.51/32	10.20.1.51	HB
10.10.4.101/32	0	OB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB

Source = OA, Destination = VB

Source = VB, Destination = OA



- If APP A sends a connection request to APP B at VB, and source IP Address of OA is selected, then packets between A and B will flow over HiperSockets.

# Static Host Routes from Linux on System z

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
Linux	z/OS	OSA	OSA	No	Yes <sup>5</sup>	Yes
Linux	z/OS	HiperSockets	OSA	No	Yes <sup>5</sup>	No
Linux	z/OS	OSA	HiperSockets	No	Yes <sup>5</sup>	Yes
Linux	z/OS	HiperSockets	HiperSockets	No	Yes	No
Linux	z/OS	OSA	VIPA	No	Yes <sup>5&amp;6</sup>	Yes
Linux	z/OS	HiperSockets	VIPA	No	Yes <sup>5</sup>	No

<sup>5</sup> Requires static host route to reach OSA over HiperSockets.

<sup>6</sup> Requires static host route to reach VIPA over HiperSockets.

• **Conclusion:**

- It is possible to have same CEC traffic over HiperSockets and external traffic over OSA using a single Source IP Address, OSA, to same CEC partners as well as external partners.
- For same CEC partners it is possible to route packets over HiperSockets even though the destination IP Address in the packet is the OSA or VIPA.



# Destination and Source is OSA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table


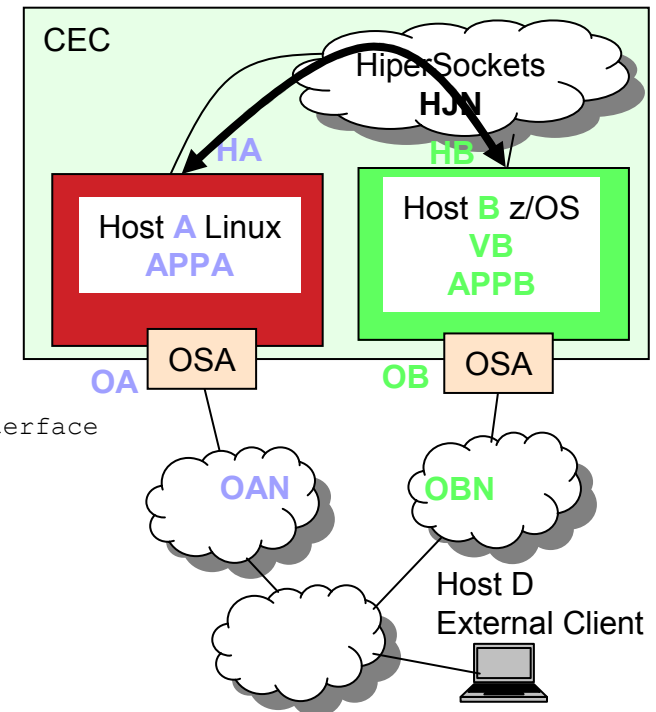
Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HA
★ → 10.10.4.101/32	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
★ → 10.1.190.51/32	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at OB, and source IP Address of OA is selected, then packets between A and B will flow over HiperSockets.

# Destination is OSA and Source is HiperSockets

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
10.4.1.101/32	10.20.1.101	HA
10.10.4.101/32	10.20.1.101	HA
10.1.190.0/24	0	OA
★ → 10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
★ → 10.1.190.51/32	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at OB, and source IP Address of HA is selected, then packets between A and B will flow over HiperSockets.

# Destination is OSA and Source is VIPA

## Host A

**OA** = 10.1.190.51  
**OAN** = 10.1.190.0/24  
**HA** = 10.20.1.51  
**HJN** = 10.20.1.0/24

### Host A Linux Routing Table



Destination	Gateway	Interface
★ → 10.4.1.101/32	10.20.1.101	HA
10.10.4.101/32	10.20.1.101	HA
10.1.190.0/24	0	OA
10.20.1.0/24	0	HA
DEFAULT	10.1.190.1	OA

## Host B

**VB** = 10.4.1.101  
**OB** = 10.10.4.101  
**OBN** = 10.10.4.0/24  
**HB** = 10.20.1.101  
**HJN** = 10.20.1.0/24

### Host B z/OS Routing Table

Destination	Gateway	Interface
★ → 10.1.190.51/32	10.20.1.51	HB
10.10.4.0/24	0	OB
10.20.1.0/24	0	HB
DEFAULT	10.10.4.1	OB



- If APPB sends a connection request to APPA at OB, and source IP Address of VA is selected, then packets between A and B will flow over HiperSockets.

# Static Host Routes to Linux on System z

Sending connection request	Receiving connection request	Source IP Address	Destination IP Address	Dynamic Routing	Same CEC over HiperSockets	External partner with same Source IP Address
z/OS	Linux	OSA	OSA	No	Yes <sup>5</sup>	Yes
z/OS	Linux	HiperSockets	OSA	No	Yes <sup>5</sup>	No
z/OS	Linux	VIPA	OSA	No	Yes <sup>5&amp;6</sup>	Yes
z/OS	Linux	OSA	HiperSockets	No	Yes <sup>5</sup>	Yes
z/OS	Linux	HiperSockets	HiperSockets	No	Yes	No
z/OS	Linux	VIPA	HiperSockets	No	Yes <sup>6</sup>	Yes

<sup>5</sup> Requires static host route to reach OSA over HiperSockets, like 10.10.1.101/32 10.20.1.101 HA

<sup>6</sup> Requires static host route to reach VIPA over HiperSockets, like 10.30.1.101/32 10.20.1.101 HA

- **Conclusion:**

- It is possible to have same CEC traffic over HiperSockets and external traffic over OSA using a single Destination IP Address, OSA, from same CEC partners as well as external partners.

For some CEC partners it is possible to route packets over HiperSockets even though the destination IP Address in the packet is the OSA.




## Static Routes

- Do you need all the static host routes in the previous example?
  - Two for Linux on System z to route to the z/OS VIPA and OSA Addresses.
  - One on z/OS to route to the Linux on System z OSA Address.
- The answer is that it depends.
  - Is all the traffic to and from z/OS is to the VIPA rather than both the VIPA and OSA address? If so then the static host route to the z/OS OSA Address is not needed. However it might be easier just to add all three just to be absolutely sure all the sessions are sent over HiperSockets.
- Static Routes are manual routes that do not offer dynamic backup, however since HiperSockets is internal to the CEC the lack of backup may be acceptable.
- Static Routes may be administratively onerous, depending upon the number of host partners.



## Bind InAddrAny and IQDX

- Does the Linux on System z application have to bind to a single IP Address? If the Linux application binds in InAddrAny (does not bind to a single IP address):
  - Clients on z/OS may use the HiperSockets address to identify the Linux Server at the same time that non-CEC clients use the OSA address to identify the Linux Server.
  - Client on Linux may use the HiperSockets addresses to identify same CEC Servers at the same time that traffic to non-CEC Servers is routed over the OSA.



If the non-CEC z/OS partner is running in zBX then HiperSockets Integration with OSA for IEDN (IQDX) is the solution!

OSA Supported Features	z/OS	z/VM	Linux on System z	z/VSE	OSE	OSD	OSX	OSM
IPv4 Support	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
IPv6 Support	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
Outbound Priority Routing Queues	Yes	No	No	No	No	Yes	Yes	
IPv4 ARP Offload	Yes	No	Yes	No	No	Yes	Yes	
Checksum Offload	Yes	No	Yes	No	No	Yes	Yes	
Segmentation Offload (Large Send)	Yes	No	Yes	Yes	No	Yes	Yes	
Hardware Assist (QEBSM & HPMA)	No	Yes	Yes	No	No	Yes	Yes	
Diagnostic Synchronization	Yes	No	No	No	No	Yes	Yes	
OSA NTA (Network Traffic Analyser)	Yes	No	No	No	No	Yes	Yes	
Native SNA flows	Yes	Yes	Yes**	No	Yes	No	No	
Layer 2 Support	No	Yes	Yes	Yes	No	Yes	Yes	
Layer 3 Virtual MAC	Yes	No	No	No	No	Yes	Yes	
Jumbo Frames (MTU >=8000)	Yes	No	Yes	Yes	No	Yes	Yes	
VLAN Support	Yes	Yes	Yes	Yes	No	Yes	Yes	Yes
VLAN Support of GVRP	Yes	Yes	No	No	No	Yes	Yes	
Gratuitous ARP Takeover	Yes	No	Yes	No	No	Yes	Yes	
Primary/Secondary Router Function	Yes	Yes	Yes	No	Yes	Yes	No	
SNMP Support	Yes	Yes	Yes	No	No	Yes	Yes	
Dynamic LAN Idle	Yes	No	No	No	No	Yes	Yes	
Optimized Latency Mode	Yes	No	No	No	No	Yes	Yes	
Data Connection Isolation	No	Yes	Yes	No	No	Yes	Yes	
Interface Isolation	Yes	Yes	No	No	No	Yes	Yes	
Adapter Interruptions	No	Yes	Yes	Yes	No	Yes	Yes	
Inbound Workload Queuing	Yes	No	No	No	No	Yes	Yes	
OSAINFO	Yes	Yes	Not Yet	No	No	Yes	Yes	Yes
HiperSockets Integration with IEDN	Yes	No*	No	No	No	No	Yes	
HiperSockets Virtual Switch Bridge Support	No	Yes	Not Yet	No	No	Yes	Yes	
* Depends upon the z/VM release								
** Requires CCL (Communication Controller for Linux) or Communications Server for Linux								

<b>HiperSockets Supported Features</b>	<b>z/OS</b>	<b>z/VM</b>	<b>Linux on System z</b>	<b>z/VSE</b>
IPv4 Support	Yes	Yes	Yes	Yes
IPv6 Support	Yes	Yes	Yes	Yes
VLAN Support	Yes	Yes	Yes	Yes
Network Concentrator	No	No	Yes	No
Layer 2 Support	No	Yes	Yes	No
Multiple Write Facility	Yes	No	No	No
zIIP Assisted Multiple Write Facility	Yes	No	No	No
HiperSockets NTA (Network Traffic Analyzer)	No	No	Yes	No
Integration with IEDN (IQDX)	Yes	No*	Yes	No
Virtual Switch Bridge Support	No	Yes	No	No
Fast Path to Linux (LFP) Support / IUCV over HiperSockets	No	No	Yes	Yes
Completion Queue	No	Yes	No	Yes
* Depends upon the z/VM release				

The End