

Mainframe OSA Connectivity and Routing in a Cisco Environment


Junnie Sadler
Cisco Systems

August 02, 2010



Objective

- Give audience members exposure to the router/networking world beyond the mainframe
- Impart an appreciation of network design issues
- Present the network as a whole vs Box Centric
- Enable greater ability to troubleshoot connectivity/performance problems
- Not to make CCIE's out of anyone
- Not to sell Cisco equipment



SHARE

Technology • Connections • Results

Overview

- **Background –
IP From the Mainframe's View**
- **Mainframe IP Communications
Hypersockets , XCF**
- **IP Routing Protocols:
Static, RIP (v1 & v2), OSPF**
- **Network Design Considerations:**




SHARE

Technology • Connections • Results


Mainframe view of the Network

- The network is just a big cloud
- Stuff happens



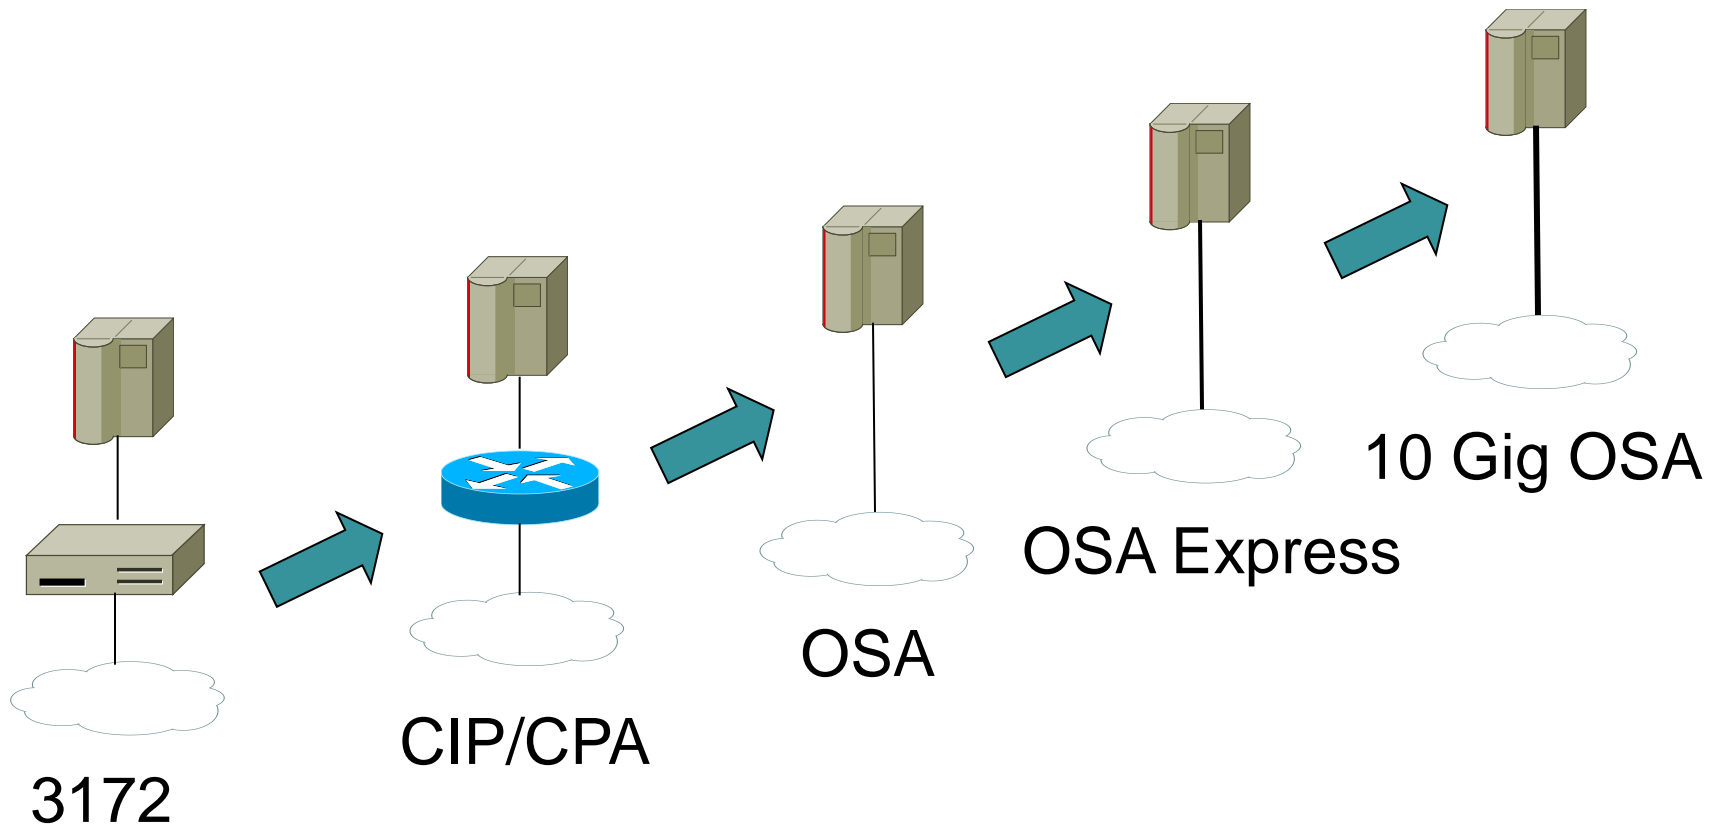
Network view of the Mainframe

- Sysplex – isn't that a big movie theater
- VIPA – that's a poisonous snake, right?



- LPAR – is that something Tiger Woods worries about?

IP Connectivity Evolution



Background

- - Mainframe as a IP Server
- - Driving Factors toward IP Routing Requirements:
 - New Applications are IP Based
 - QoS (Quality of Service) required for these Applications
 - Old SNA Applications being re-written or replaced with IP transport (EE) Enterprise Extender
 - TN3270 Application Requirements (IP to & from the MF)
 - Data Replication – IP Based High Performance Requirements
 - VIPA (Virtual IP Address)
 - IP Load Balancing between Sysplex

Mainframe IP Communications - Part 1

- **Hipersockets support:** Hipersockets (Internal Queued Direct I/O - iQDIO) is a S/390 zSeries hardware feature that provides high-speed communicating LPAR-to LPAR on the same processor (via memory). It also provides secure data flows between LPARs and high availability, in that there is no network attachment dependency or exposure to adapter failures.

- **XCF:**

Cross-System Coupling Facility (XCF) support for CS for z/OS IP allows IP traffic across the XCF connection being used by systems that are operating in the same sysplex group.

By default,VTAM will use XCF if it is available.

Since CS for z/OS IP uses the I/O facilities of VTAM, the use of XCF must be shared by APPN and TCP/IP.

The deactivation of the XCF major node (ISTLSXCF) will take down any connection (IP or APPN) that is using the XCF connection.

Mainframe IP Communications – Part 2

- **z/OS-supported routing daemons:** CS for z/OS IP ships two routing applications, **ORouteD** and OMPROUTE. OMPROUTE and **ORouteD** cannot run on the same TCP/IP stack concurrently.
 - **ORouteD:** ORouteD is the CS for z/OS IP implementation of the Routing Information Protocol (RIP) Version 1 (RFC 1058) and RIP Version 2 (RFC 1723). **(Removed in z/OS v1.7)**
 - **OMPROUTE:** In OS/390 V2R6 IP and later, OMPROUTE implements the Open Shortest Path First (OSPF) protocol described in RFC 1583 (OSPF Version 2) as well as Static & RIP V1 and RIP V2.



SHARE

Technology • Connections • Results

Mainframe IP Communications – Part 2

- **z/OS-Configuration**



SHARE

Technology • Connections • Results

Cisco Routers

- There is a wide variety of routers to serve different purposes: Branch, Aggregation, Service Provider.
- The most common device connecting OSA Express adapters to a Cisco network is the Catalyst 6500.



Cisco Router Basics – OS & Configuration

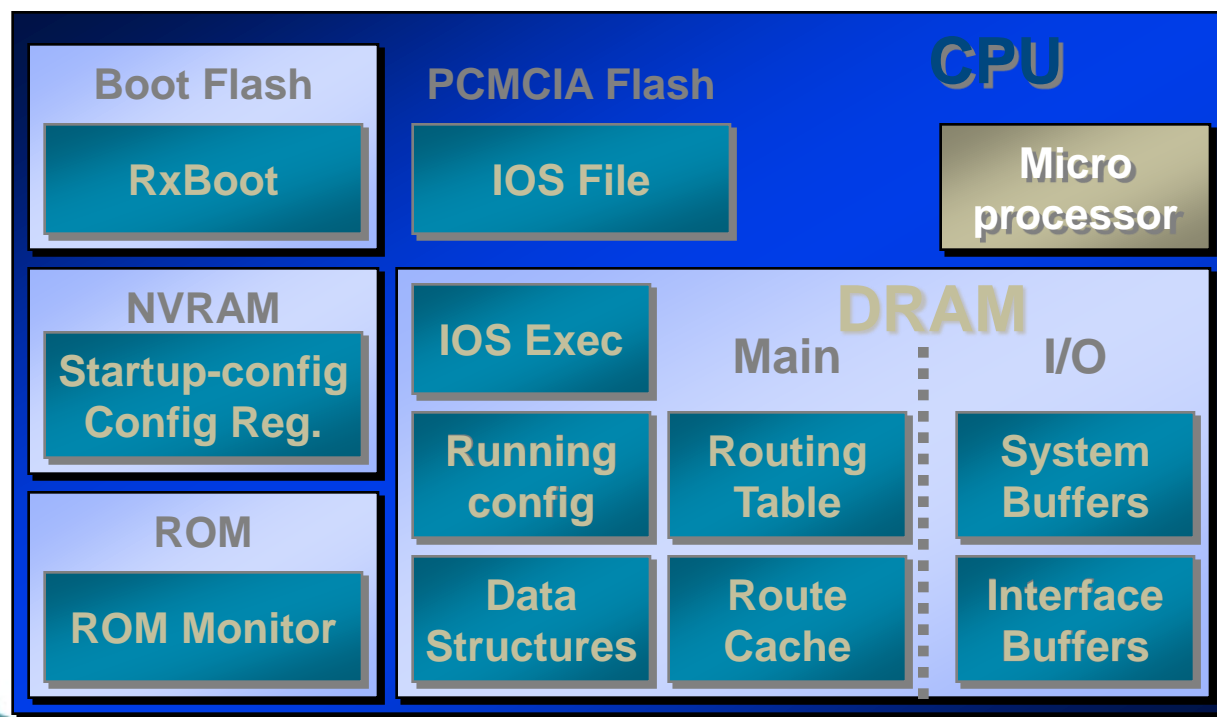


- Runs Cisco IOS – Internetwork Operating System –
- Configuration is stored as single ASCII flat file
- Configuration commands can be global, interface, or protocol commands
- Configuration is dynamic – commands take effect immediately
- CLI interface for configuration and display commands
- Exec and Configuration modes
- Username/Password login – usually via TACACS to an ACS server with different privilege levels



Router Basics – File Systems

- Flash drives
- NVRAM (non-volatile RAM) stores running configuration
- Code and configurations can be transferred using ftp and tftp.






SHARE

Technology • Connections • Results

Catalyst 6500

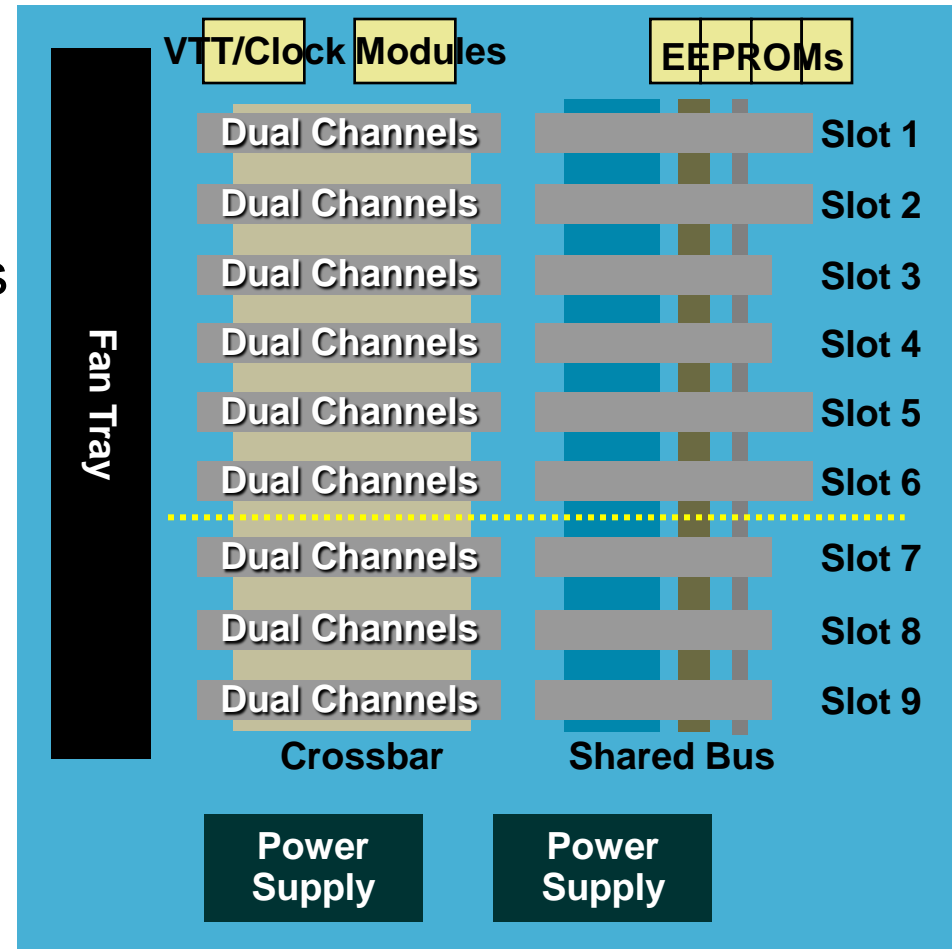
- Enterprise Class Switch
- Multilayer Switching (MLS) = ASIC forwarding
- Multiple chassis types and line cards
- Shared backplane architecture
- Up to 400 Mpps performance





Catalyst 6500 - Chassis

- Shared bus architecture
- Cross-bar switching matrix
- Dual power supplies and Fans



Catalyst 6500 – Supervisor Card

- “Brains” of the box
- Performs packet forwarding
- Out of band management of line cards
- Can be deployed redundantly in a chassis
- Console connection
- Flash drives





SHARE

Technology • Connections • Results

Catalyst 6500 – Line Cards

- Ethernet ports –
10/100/1000 Mbps/10 Gbps
- Optical and WAN cards
- Can be Fabric enabled
- DFC – distributed forwarding card






SHARE

Technology • Connections • Results

Catalyst 6500 – Service Modules

- Firewalls
- Load Balancing
- Intrusion Detection
- Wireless controllers
- SSL acceleration





SHARE

Technology • Connections • Results

IP Routing

Static

RIP

OSPF

Route Manipulation



IP Routing: IP Address Breakdown (Network, Subnetwork, Host)

IP Address Breakdown: using binary math:

00000000	>	0 0 0 0 0 0 0 0	= 0
11111111	>	1, 2, 4, 8, 16, 32, 64, 128	= 255

1.1.1.1 255.255.255.0

00000001.00000001.00000001.00000001

11111111.11111111.11111111.00000000

1 . 1 . 1 . 0

Network . Sub -----Net . Host

Where there is a 1 in the subnet mask bring down what is in the IP address, where there is a 0 that is the host address

IPv4 Network Class Layout:

- Network **Class A** = 0-127 ----- 0, 10 & 127 reserved and will not be advertised into the Internet.
- Network **Class B** = 128 -191 ----- 172.16.0.0 thru 172.31.255.255 reserved and will not be advertised into the Internet.
- Network **Class C** = 192– 223 -----192.168.0.0 thru 192.168.255.255 reserved and will not be advertised into the Internet
- **IPv4 private addresses**
- Computers not connected to the Internet (e.g., factory machines that communicate with each other via TCP/IP) need not have globally-unique IP addresses.
- Three ranges of IPv4 addresses for [private networks](#), one range for each [class](#), were reserved in [RFC 1918](#); these addresses are not routed on the Internet, and thus need not be coordinated with an IP address registry.
- 0, 10 & 127 reserved and will not be advertised into the Internet.
- 172.16.0.0 thru 172.31.255.255 reserved and will not be advertised into the Internet.
- 192.168.0.0 thru 192.168.255.255 reserved and will not be advertised into the Internet
- **IP version 6 addresses**
- IPv6 is the next generation of the Internet Protocol, designed to replace IPv4 on the Internet.
- Example of an IPv6 address: 2001:0db8:85a3:08d3:1319:8a2e:0370:7334

Just as there are addresses for private, or internal networks in IPv4, there are blocks of addresses set aside in IPv6 for private addresses. Referred to as unique local addresses (ULA). [RFC 4193](#) sets aside the routing prefix fc00::/7 for this block. The addresses include a 40-bit pseudorandom number that minimize the risk of address collisions if sites merge or packets are misrouted.

IP Routing: Static Routing MF



- **Static routing is normally used in simple network topologies:**
 - Static routes are configured manually for each destination network or device.
Note: If a destination (sub)network becomes unreachable, the static routes for that (sub)network remain in the routing table, and packets will still be forwarded to the destination. We call this (**BlackholedTraffic**).
- **There are other circumstances when static routing is useful:**
 - To define a default route that is used to forward traffic when the routing table does not contain a more specific route to a destination
 - To define a route that is not automatically advertised/ known within a network.
 - When line utilization or tariffs make it undesirable to send routing advertisements through lower capacity connections.
 - To ensure traffic for a specific destination takes a specific route through the network.
 - To provide a secure network by defining routes to authorized (sub)networks only.

IP Routing: Static Routing > MainFrame

- Defining Static Routes (MF):

There are two ways to define static routes:

Using the BEGINROUTES and or **GATEWAY (Deprecated)** statements in PROFILE.TCPIP..

- Type of Static Routes:

replaceable static routes – Used in conjunction with OMPRoute-- if a static route is defined as being replaceable (**Keyword {REPL}**), OMPROUTE will replace it with a found dynamic route to the same destination.

Note: The **GATEWAY** statement and routing daemon ORouted do not support replaceable static routes

nonreplaceable static routes -- **(Default Gateway Statements) as an example.**

- Components of a Static Route:

IP Destination Address / Subnet Mask

Next Hop address – (first gateway along the route to the destination)

- Route Selection:

A **specific host route** takes precedence over a **subnetwork route**, followed by a **network route**, followed by a **supernetwork route**, and finally, a **default route**.



Static Routing – Cisco Router

```

router-4#conf t
router-4(config)#
ip route 100.1.0.0 255.255.255.0 10.1.10.12
ip route 100.2.0.0 255.255.255.0 10.1.10.12
ip route 100.3.0.0 255.255.255.0 10.1.10.12
ip route 100.4.0.0 255.255.255.0 10.1.10.13
ip route 100.5.0.0 255.255.255.0 10.1.10.13
ip route 100.6.0.0 255.255.255.0 10.1.10.13
ip route 100.7.0.0 255.255.255.0 10.1.10.14
ip route 100.8.0.0 255.255.255.0 10.1.10.14
ip route 100.9.0.0 255.255.255.0 10.1.10.14

```

```

router-4# sh ip route static
100.0.0.0/24 is subnetted, 9 subnets
S    100.4.0.0 [1/0] via 10.1.10.13
S    100.5.0.0 [1/0] via 10.1.10.13
S    100.6.0.0 [1/0] via 10.1.10.13
S    100.7.0.0 [1/0] via 10.1.10.14
S    100.1.0.0 [1/0] via 10.1.10.12
S    100.2.0.0 [1/0] via 10.1.10.12
S    100.3.0.0 [1/0] via 10.1.10.12
S    100.8.0.0 [1/0] via 10.1.10.14
S    100.9.0.0 [1/0] via 10.1.10.14

```

```

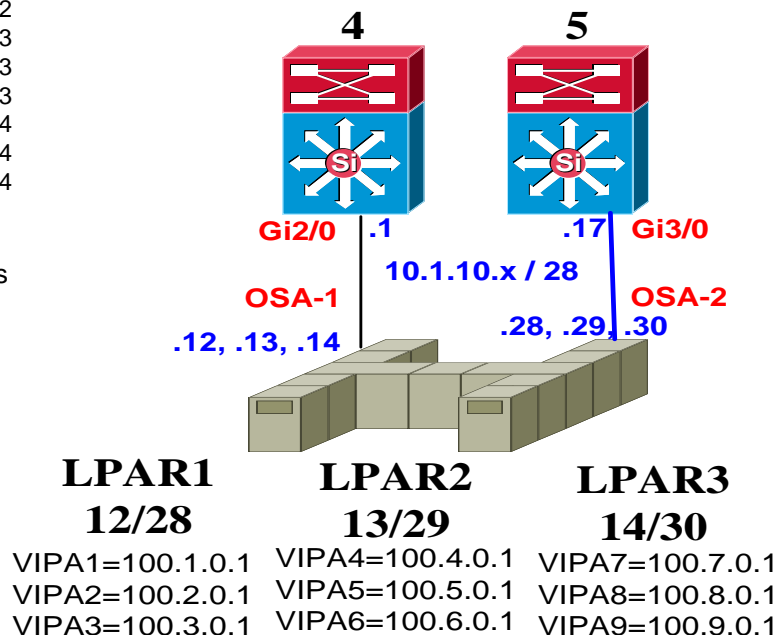
router-5#conf t
router-5(config)#
ip route 100.1.0.0 255.255.255.0 10.1.10.28
ip route 100.2.0.0 255.255.255.0 10.1.10.28
ip route 100.3.0.0 255.255.255.0 10.1.10.28
ip route 100.4.0.0 255.255.255.0 10.1.10.29
ip route 100.5.0.0 255.255.255.0 10.1.10.29
ip route 100.6.0.0 255.255.255.0 10.1.10.29
ip route 100.7.0.0 255.255.255.0 10.1.10.30
ip route 100.8.0.0 255.255.255.0 10.1.10.30
ip route 100.9.0.0 255.255.255.0 10.1.10.30

```

```

router-5#sh ip route static
100.0.0.0/24 is subnetted, 9 subnets
S    100.4.0.0 [1/0] via 10.1.10.29
S    100.5.0.0 [1/0] via 10.1.10.29
S    100.6.0.0 [1/0] via 10.1.10.29
S    100.7.0.0 [1/0] via 10.1.10.30
S    100.1.0.0 [1/0] via 10.1.10.28
S    100.2.0.0 [1/0] via 10.1.10.28
S    100.3.0.0 [1/0] via 10.1.10.28
S    100.8.0.0 [1/0] via 10.1.10.30
S    100.9.0.0 [1/0] via 10.1.10.30

```



MF Static Default Route – pointing to

10.1.10.1

10.1.10.17

Route Filtering and Manipulation

- A number of mechanisms exist to filter routing information
- Distribute lists
- Route maps
- Prefix lists
- Route summarization
- Routing Protocol Redistribution



SHARE

Technology • Connections • Results

Static Routing – Cisco Router

```
router-6#sh ip route 100.1.0.0
```

% Network not in table

```
router-6#sh ip route 10.1.10.0
```

% Subnet not in table

```
router-7#sh ip route 100.1.0.0
```

% Network not in table

```
router-7#sh ip route 10.1.10.0
```

% Subnet not in table

WHY aren't they we ask ????

How do we fix it ???

```
router-4 & 5 (config-router)#
```

```
router ospf 100
```

```
log-adjacency-changes
```

```
redistribute static metric 100 subnets
```

```
router-6#sh ip route 10.1.10.0
```

% Subnet not in table

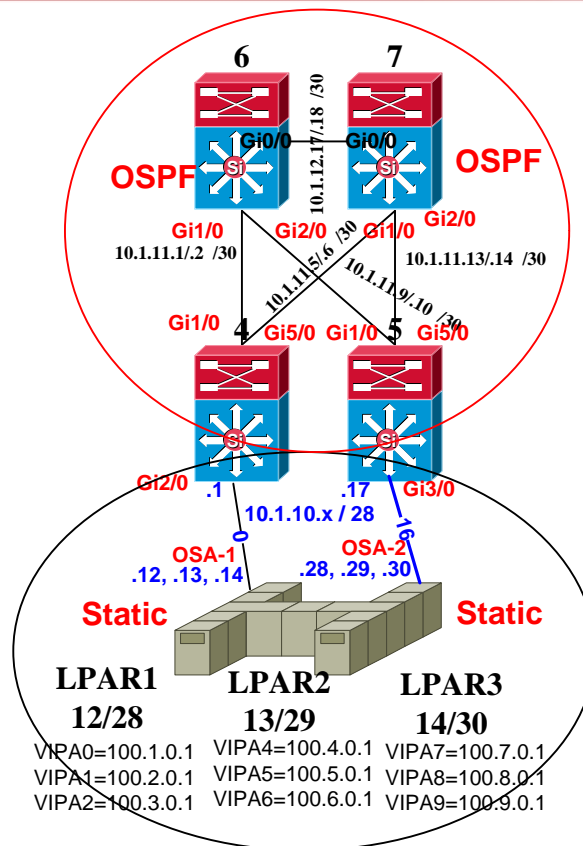
```
router-6#sh ip route 100.1.0.0
```

Routing entry for 100.1.0.0/24 Known via "ospf 100", distance 110, metric 100, type extern 2,

Routing Descriptor Blocks:

* 10.1.11.9, from 10.1.11.18, 00:48:35 ago, via gi2/0

WHY is 100.x.x.x in the table but not 10.1.10.0 we ask ????



```
router-4 & 5 (config-router)#
```

```
router ospf 100
```

```
log-adjacency-changes
```

```
redistribute static metric 100 subnets
```

```
redistribute connected metric 150 subnets
```

```
router-6 & 7 #sh ip route 100.1.0.0
```

Routing entry for 100.1.0.0/24

Known via "ospf 100", distance 110, metric 100, type extern 2, forward metric 10

Last update from 10.1.11.13 on gi2/0, 00:16:02 ago

Routing Descriptor Blocks:

* 10.1.11.13, from 10.1.11.18, 00:16:02 ago, via gi2/0

Route metric is 100, traffic share count is 1

```
router-6# & sh ip route 10.1.10.16
```

Routing entry for 10.1.10.16/28

Known via "ospf 100", distance 110, metric 150, type extern 2, forward metric 10

Last update from 10.1.11.13 on gi2/0, 00:04:48 ago

Routing Descriptor Blocks:

* 10.1.11.13, from 10.1.11.18, 00:04:48 ago, via gi2/0

Route metric is 150, traffic share count is 1

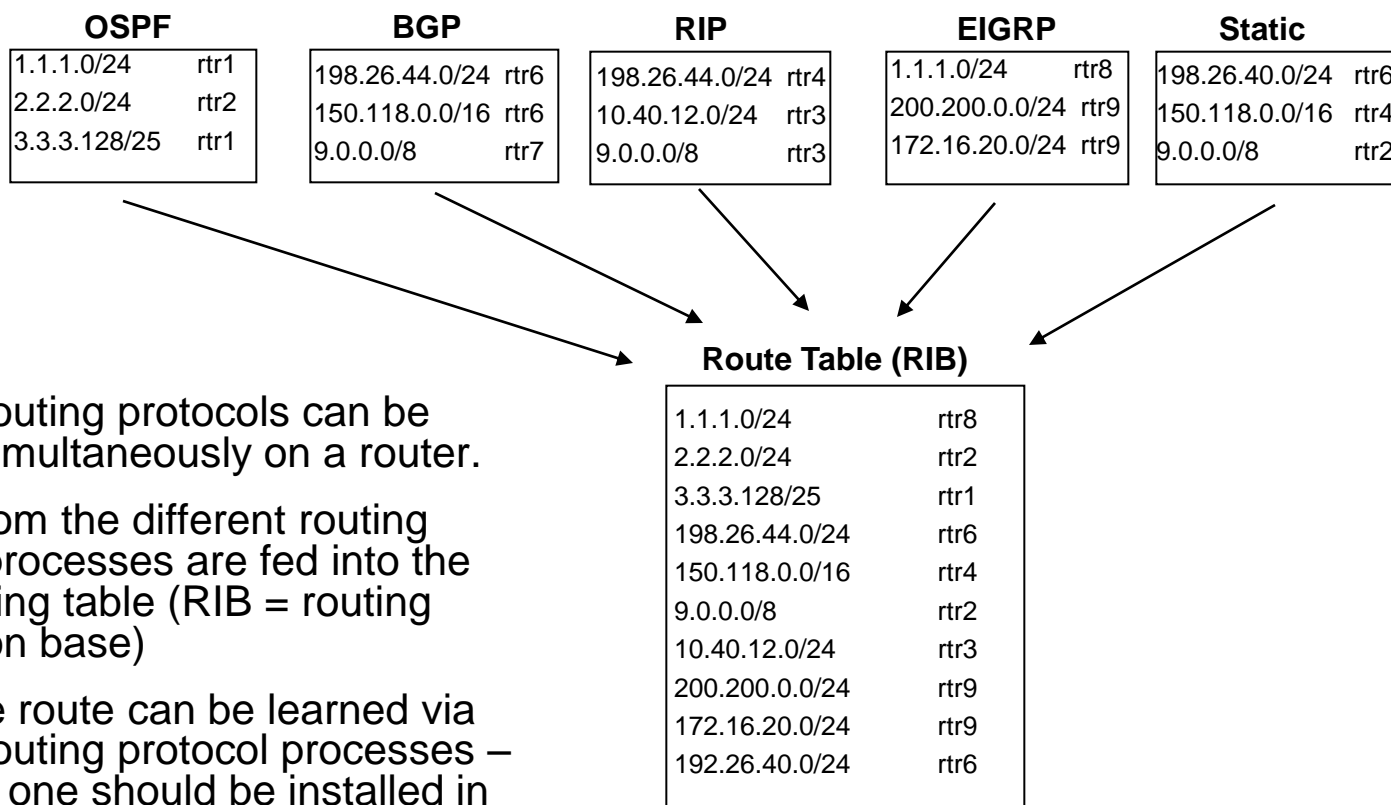
MF Static Default Route – pointing to

10.1.10.1

10.1.10.17



Multiple Routing Protocols



- Multiple routing protocols can be running simultaneously on a router.
- Routes from the different routing protocol processes are fed into the main routing table (RIB = routing information base)
- The same route can be learned via multiple routing protocol processes – but which one should be installed in the RIB? → the one with the best Administrative Distance



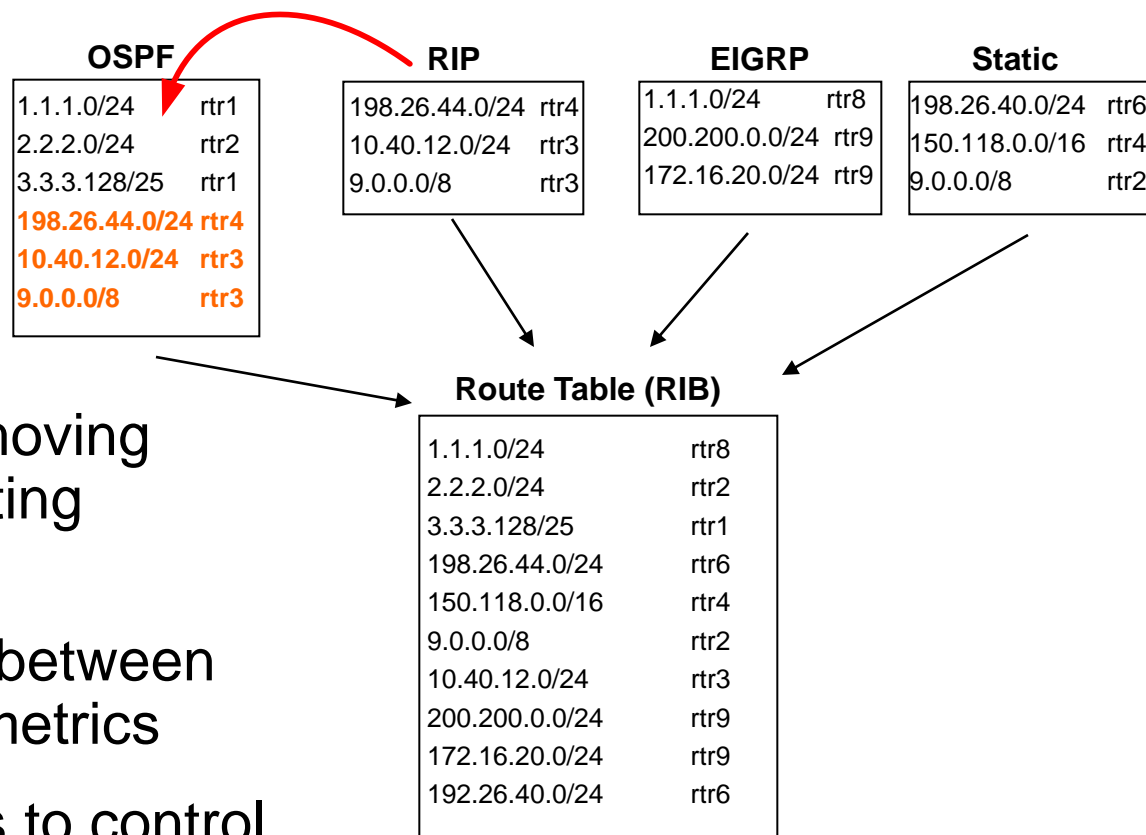
Administrative Distance

- Same concept as Route Precedence on the mainframe
- The lower the Admin Distance, the more desirable the route
- Can be changed from the default values listed
- Note differences between internal and external EIGRP and BGP routes
- OSPF internal and external routes given same Admin Distance by default
- Most specific route is still chosen during route table look-up. eg, a more specific RIP route will be used over a summarized OSPF route

Connected	0
Static	1
EIGRP Summary	5
eBGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
ODR	160
External EIGRP	170
iBGP	200
Unknown	255



Redistribution



- Method of sharing/moving routes from one routing protocol to another
- Involves converting between different protocols' metrics
- Can use route-maps to control which routes are redistributed



SHARE

Technology • Connections • Results

RIP Routing – Cisco Router

Customer's Network

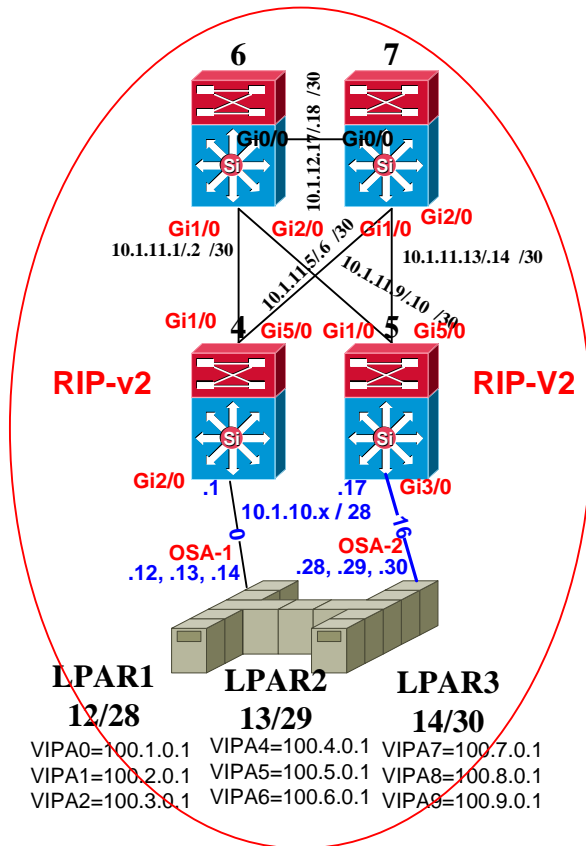
```
router-4#config-router
router rip
version 2
network 10.0.0.0
```

```
router-4#sh ip route
Gateway of last resort is not set
100.0.0.0/24 is subnetted, 9 subnets
```

```

R 100.4.0.0 [120/1] via 10.1.10.13, E2/0
R 100.5.0.0 [120/1] via 10.1.10.13, E2/0
R 100.6.0.0 [120/1] via 10.1.10.13, E2/0
R 100.7.0.0 [120/1] via 10.1.10.14, E2/0
R 100.1.0.0 [120/1] via 10.1.10.12, E2/0
R 100.2.0.0 [120/1] via 10.1.10.12, E2/0
R 100.3.0.0 [120/1] via 10.1.10.12, E2/0
R 100.8.0.0 [120/1] via 10.1.10.14, E2/0
R 100.9.0.0 [120/1] via 10.1.10.30, E3/0

```



```
router-5#config-router
router rip
version 2
network 10.0.0.0
```

```
router-5#sh ip route
Gateway of last resort is not set
100.0.0.0/24 is subnetted, 9 subnets
```

```

R 100.4.0.0 [120/1] via 10.1.10.29, E3/0
R 100.5.0.0 [120/1] via 10.1.10.29, E3/0
R 100.6.0.0 [120/1] via 10.1.10.29, E3/0
R 100.7.0.0 [120/1] via 10.1.10.30, E3/0
R 100.1.0.0 [120/1] via 10.1.10.28, E3/0
R 100.2.0.0 [120/1] via 10.1.10.28, E3/0
R 100.3.0.0 [120/1] via 10.1.10.28, E3/0
R 100.8.0.0 [120/1] via 10.1.10.30, E3/0
R 100.9.0.0 [120/1] via 10.1.10.30, E3/0

```

What is required for the MF running RIPv2 to route out to the internet or to a different routing AS network with this configuration ?



SHARE

Technology • Connections • Results

RIP Routing – Cisco Router

```


router-4# sh run
router rip
version 2
passive-interface default
no passive-interface gi2/0
network 10.0.0.0
default-information originate

```

```

Router-5# sh run
router rip
version 2
passive-interface default
no passive-interface gi3/0
network 10.0.0.0
default-information originate


```



```

MF#sh ip route
Gateway of last resort is 10.1.10.1 to network 0.0.0.0
 100.0.0.0/24 is subnetted, 9 subnets
R*  0.0.0.0/0 [120/1] via 10.1.10.1, gi2/0
    [120/1] via 10.1.10.17, gi3/0
 10.0.0.0/28 is subnetted, 2 subnets
C   10.1.10.0 is directly connected, gi2/0
C   10.1.10.16 is directly connected, gi3/0

```



SHARE

Technology • Connections • Results

RIP Routing – Cisco Router

router-6#sh ip route 100.9.0.1

Routing entry for 100.9.0.0/24

Known via "ospf 100", distance 110, Metric 100,
type extern 2, on gi1/0,

Routing Descriptor Blocks:

* 10.1.11.1, from 10.1.11.1, via gi1/0

router-6#sh ip ospf nei

Neighbor ID	Pri	State	Address	Interface
10.1.11.18	1	FULL/BDR	10.1.11.9	gi0/0
10.1.11.5	1	FULL/BDR	10.1.11.1	gi1/0
10.1.11.1	1	FULL/BDR	10.1.11.1	gi2/0

router-6#sh ip route

Gateway of last resort is not set

100.0.0.0/24 is subnetted, 9 subnets

- O E2 100.4.0.0 [110/100] via 10.1.11.1, gi1/0
- O E2 100.7.0.0 [110/100] via 10.1.11.1, gi1/0
- O E2 100.1.0.0 [110/100] via 10.1.11.1, gi1/0

And all of the others ****

router-4# sh run

router ospf 100

redistribute connected

redistribute rip metric 100 subnets

network 10.1.11.0 0.0.0.255 area 0.0.0.0

!

router rip

version 2

passive-interface default

no passive-interface gi2/0

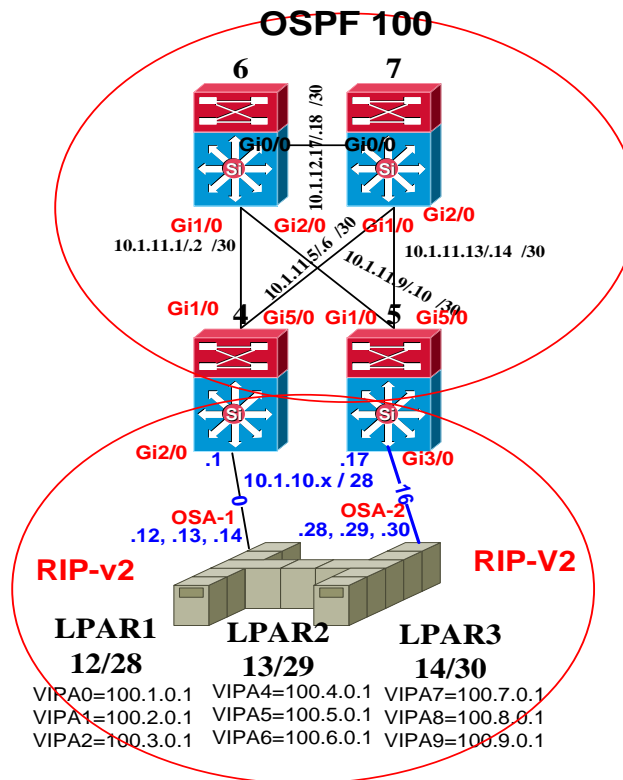
network 10.0.0.0

default-information originate

distribute-list 51 out gi2/0

!

access-list 51 permit 0.0.0.0



MF#sh ip route

Gateway of last resort is 10.1.10.1 to network 0.0.0.0

100.0.0.0/24 is subnetted, 9 subnets

R* 0.0.0.0/0 [120/1] via 10.1.10.1, gi2/0

[120/1] via 10.1.10.17, gi3/0

10.0.0.0/28 is subnetted, 2 subnets

C 10.1.10.0 is directly connected, gi2/0

C 10.1.10.16 is directly connected, gi3/0

router-7#sh ip route 100.9.0.1

Routing entry for 100.9.0.0/24

Known via "ospf 100", distance 110, Metric 100,
type extern 2, on gi1/0,

Routing Descriptor Blocks:

* 10.1.11.5, from 10.1.11.5, via gi1/0

router-7#sh ip ospf nei

Neighbor ID	Pri	State	Address	Interface
10.1.11.18	1	FULL/BDR	10.1.11.13	gi2/0
10.1.11.5	1	FULL/BDR	10.1.11.5	gi1/0
10.1.11.14	1	FULL/BDR	10.1.11.1	gi2/0

router-7# ip route

Gateway of last resort is not set

100.0.0.0/24 is subnetted, 9 subnets

- O E2 100.6.0.0 [110/100] via 10.1.11.5, gi1/0
- O E2 100.3.0.0 [110/100] via 10.1.11.5, gi1/0
- O E2 100.9.0.0 [110/100] via 10.1.11.5, gi1/0

And all of the others ****

Router-5# sh run

router ospf 100

redistribute connected

redistribute rip metric 100 subnets

network 10.1.11.0 0.0.0.255 area 0.0.0.0

!

router rip

version 2

passive-interface default

no passive-interface gi3/0

network 10.0.0.0

default-information originate

distribute-list 51 out gi3/0


!

access-list 51 permit 0.0.0.0



Mutual Redistribution

- Occurs with Bi-directional redistribution of routes between two or more routers
- Can be problematic with OSPF - Internal & External routes have the same admin distance
- In certain scenarios a pair of routers can end up pointing to each other for a group of routes
- Timing related – can depend on which router learns the routes first.





SHARE

Technology • Connections • Results

OSPF Routing – Cisco Router & MF

```

Router-24# sh run
router ospf 200
router-id 10.2.10.1
log-adjacency-changes
area 0.0.0.1 stub no-summary
passive-interface default
no passive-interface gi2/0
network 1.1.1.0 0.0.0.255 area 0
network 10.2.10.0 0.0.0.255 area 0.0.0.1

```

distribute-list 50 in gi2/0

!

```

access-list 50 deny 0.0.0.0
access-list 50 permit any

```

router-24# sh ip ospf nei 2/0

Neighbor ID	Pri	State	Address	Int
10.2.10.12	0	FULL/DROTHER	10.2.10.12	gi2/0
10.2.10.13	0	FULL/DROTHER	10.2.10.13	gi2/0
10.2.10.14	0	FULL/DROTHER	10.2.10.14	gi2/0


LPAR-2X# ip ospf nei gi2/0

Neighbor ID	Pri	State	Address	Int
10.2.10.1	1	FULL/DR	10.2.10.1	gi2/0
10.2.10.13	0	2WAY/DROTHER	10.2.10.13	gi2/0
10.2.10.14	0	2WAY/DROTHER	10.2.10.14	gi2/0

Router_Priority=0 prevents z/OS being a DR

BEST PRACTICE: Code Router_Priority=0 for the z/OS platform unless you are dealing with HiperSockets; allow Router to be the DR.

Per - Gwen Dente – S3921



LPAR-2X>sh ip route

Gateway of last resort is 10.2.10.1

to network 0.0.0.0

```

O*IA 0.0.0.0/0 [110/11] via 10.2.10.1, gi2/0
[110/11] via 10.2.10.17, gi3/0

```

```

router-25#sh run
router ospf 200
router-id 10.2.10.17
log-adjacency-changes
area 0.0.0.1 stub no-summary
passive-interface default
no passive-interface gi3/0
network 2.2.2.0 0.0.0.255 area 0.0.0.0
network 10.2.10.0 0.0.0.255 area 0.0.0.1
distribute-list 50 in gi3/0

```

!

```

access-list 50 deny 0.0.0.0
access-list 50 permit any

```

router-25# sh ip ospf nei gi3/0

Neighbor ID	Pri	State	Address	Int
10.2.10.12	0	FULL/DROTHER	10.2.10.28	gi3/0
10.2.10.13	0	FULL/DROTHER	10.2.10.29	gi3/0
10.2.10.14	0	FULL/DROTHER	10.2.10.30	gi3/0

LPAR-2X# ip ospf nei gi3/0

Neighbor ID	Pri	State	Address	Int
10.2.10.13	0	2WAY/DROTHER	10.2.10.29	gi3/0
10.2.10.14	0	2WAY/DROTHER	10.2.10.30	gi3/0
10.2.10.17	1	FULL/DR	10.2.10.17	gi3/0

Router_Priority=0 prevents z/OS being a DR

BEST PRACTICE: Code Router_Priority=0 for the z/OS platform unless you are dealing with HiperSockets; allow Router to be the DR.

Per - Gwen Dente – S3921



SHARE

Technology • Connections • Results

OSPF Routing – Cisco Router & MF

```

Router-24# sh run
router ospf 200
router-id 10.2.10.1
log-adjacency-changes
area 0.0.0.1 stub no-summary
passive-interface default
no passive-interface gi2/0
network 1.1.1.0 0.0.0.255 area 0
network 10.2.10.0 0.0.0.255 area 0.0.0.1

```

distribute-list 50 in gi2/0

!

access-list 50 deny 0.0.0.0

access-list 50 permit any

```
router-24# sh ip ospf nei 2/0
```

Neighbor ID	Pri	State	Address	Int
10.2.10.12	0	FULL/DROTHER	10.2.10.12	gi2/0
10.2.10.13	0	FULL/DROTHER	10.2.10.13	gi2/0
10.2.10.14	0	FULL/DROTHER	10.2.10.14	gi2/0

```
LPAR-2X# ip ospf nei gi2/0
```

```

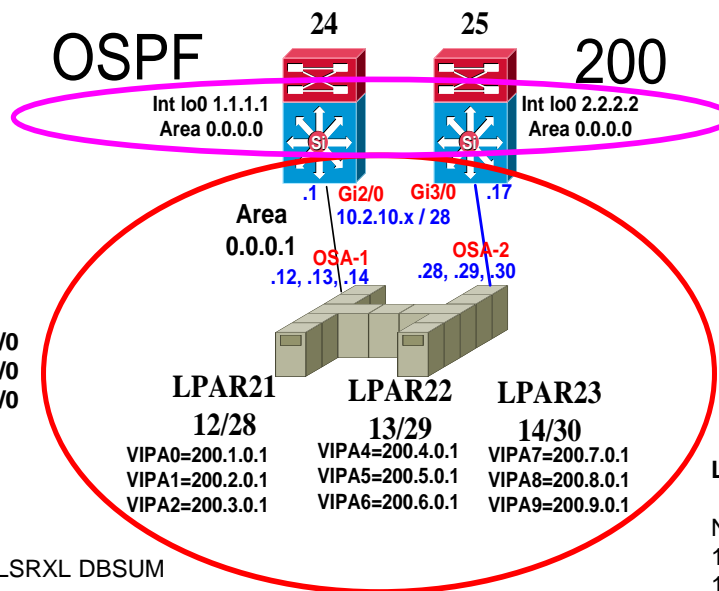
d tcpip,,omproute,ospf,nbrs
EZZ78511 NEIGHBOR SUMMARY 121
NEIGHBOR ADDR NEIGHBOR ID STATE LSRXL DBSUM
LSREQ HSUP IFC
10.2.10.1 1.1.1.1 128 0 0 0 OFF GIG
10.2.10.13 1.1.1.1 128 0 0 0 OFF GIG

```

Router_Priority=0 prevents z/OS being a DR

BEST PRACTICE: Code Router_Priority=0 for the z/OS platform unless you are dealing with HiperSockets; allow Router to be the DR.

Per - Gwen Dente – S3921



```
LPAR-2X>sh ip route
```

Gateway of last resort is 10.2.10.1

to network 0.0.0.0

```

O*IA 0.0.0.0/0 [110/11] via 10.2.10.1, gi2/0
[110/11] via 10.2.10.17, gi3/0

```

```

router-25#sh run
router ospf 200
router-id 10.2.10.17
log-adjacency-changes
area 0.0.0.1 stub no-summary
passive-interface default
no passive-interface gi3/0
network 2.2.2.0 0.0.0.255 area 0.0.0.0
network 10.2.10.0 0.0.0.255 area 0.0.0.1
distribute-list 50 in gi3/0

```

!

access-list 50 deny 0.0.0.0

access-list 50 permit any

```
router-25# sh ip ospf nei gi3/0
```

Neighbor ID	Pri	State	Address	Int
10.2.10.12	0	FULL/DROTHER	10.2.10.28	gi3/0
10.2.10.13	0	FULL/DROTHER	10.2.10.29	gi3/0
10.2.10.14	0	FULL/DROTHER	10.2.10.30	gi3/0

```
LPAR-2X# ip ospf nei gi3/0
```

Neighbor ID	Pri	State	Address	Int
10.2.10.13	0	2WAY/DROTHER	10.2.10.29	gi3/0
10.2.10.14	0	2WAY/DROTHER	10.2.10.30	gi3/0
10.2.10.17	1	FULL/DR	10.2.10.17	gi3/0

Router_Priority=0 prevents z/OS being a DR

BEST PRACTICE: Code Router_Priority=0 for the z/OS platform unless you are dealing with HiperSockets; allow Router to be the DR.

Per - Gwen Dente – S3921



SHARE

Technology • Connections • Results

Cisco Routing - OSPF & EIGRP

```

router-26#sh ip route 200.1.0.1
% Network not in table
router-26#sh ip route 200.9.0.1
% Network not in table

```

Why???

```

Router-24# sh run
router eigrp 300
passive-interface default
no passive-interface gi0/0
no passive-interface gi1/0
no passive-interface gi2/0
no passive-interface gi3/0
no passive-interface gi4/0
network 10.2.11.0 0.0.0.255
network 10.2.12.0 0.0.0.255
auto-summary

```

```

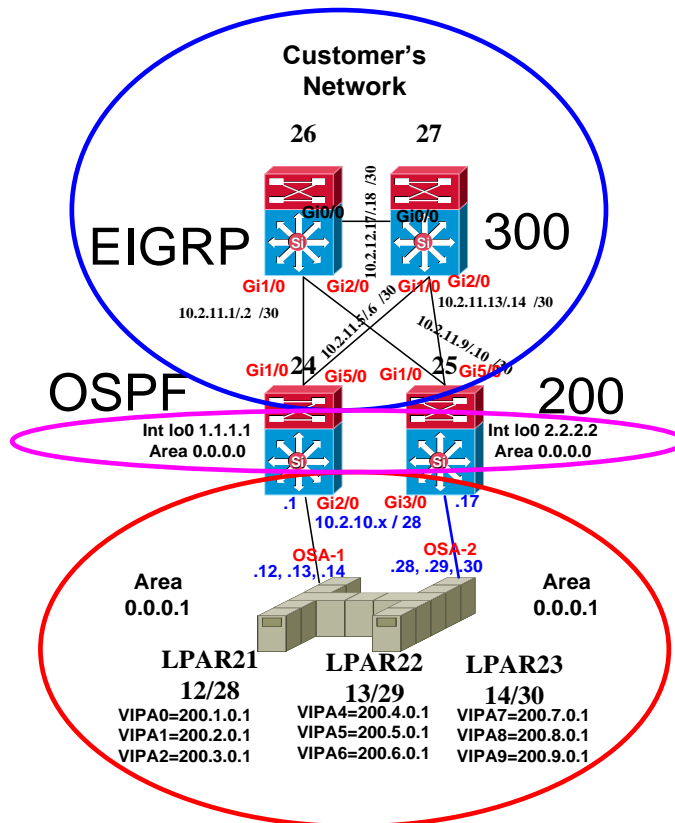
Router-24# sh run
router ospf 200
router-id 10.2.10.1
log-adjacency-changes
area 0.0.0.1 stub no-summary
passive-interface default
no passive-interface gi2/0
network 1.1.1.0 0.0.0.255 area 0
network 10.2.10.0 0.0.0.255 area 0.0.0.1
distribute-list 50 in gi2/0
!
access-list 50 deny 0.0.0.0
access-list 50 permit any

```

```

router-24#sh ip route 200.1.0.1
Routing entry for 200.1.0.1/32
Known via "ospf 200", distance 110,
router-24#sh ip route 200.9.0.1
Routing entry for 200.9.0.1/32
Known via "ospf 200", distance 110,

```



```

LPAR-2X>sh ip route
Gateway of last resort is 10.2.10.1
to network 0.0.0.0
O*IA 0.0.0.0/0 [110/11] via 10.2.10.1, gi2/0
[110/11] via 10.2.10.17, gi3/0

```

```

router-27#sh ip route 200.1.0.1
% Network not in table
router-27#sh ip route 200.9.0.1
% Network not in table

```

Why???

```

Router-25# sh run
router eigrp 300
passive-interface default
no passive-interface gi0/0
no passive-interface gi1/0
no passive-interface gi2/0
no passive-interface gi3/0
no passive-interface gi4/0
network 10.2.11.0 0.0.0.255
network 10.2.12.0 0.0.0.255
auto-summary

```

```

router-25#sh run
router ospf 200
router-id 10.2.10.17
log-adjacency-changes
area 0.0.0.1 stub no-summary
passive-interface default
no passive-interface gi3/0
network 2.2.2.0 0.0.0.255 area 0.0.0.0
network 10.2.10.0 0.0.0.255 area 0.0.0.1
distribute-list 50 in gi2/0
!
access-list 50 deny 0.0.0.0
access-list 50 permit any

```

```

router-25#sh ip route 200.1.0.1
Routing entry for 200.1.0.1/32
Known via "ospf 200", distance 110,
router-25#sh ip route 200.9.0.1
Routing entry for 200.9.0.1/32
Known via "ospf 200", distance 110,

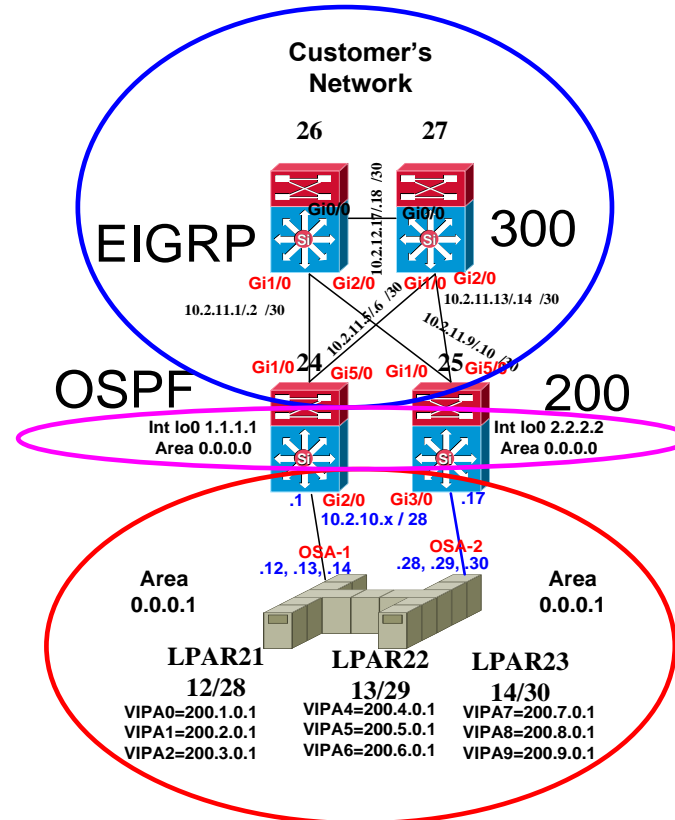
```

Cisco Routing - OSPF / EIGRP Redistribution

```
router-26#sh ip route 200.1.0.1
Routing entry for 200.1.0.1/32
Known via "eigrp 300", distance 170,type external
router-26#sh ip route 200.9.0.1
Routing entry for 200.9.0.1/32
Known via "eigrp 300", distance 170,type external
```

```
router-26#sh ip route
Gateway of last resort is not set
200.1.0.0/32 is subnetted, 1 subnets
D EX 200.1.0.1 [170/25651200] via 10.2.11.1, gi1/0
[170/25651200] via 10.2.11.9, gi2/0
Thru
D EX 200.9.0.1 [170/25651200] via 10.2.11.1, gi1/0
[170/25651200] via 10.2.11.9, gi2/0
```

```
Router-24#
router eigrp 300
redistribute ospf 200 metric 100 100 255 255 1500
passive-interface default
no passive-interface gi1/0
no passive-interface gi5/0
network 10.2.11.0 0.0.0.255
auto-summary;
```



```
LPAR-2X>sh ip route
Gateway of last resort is 10.2.10.1
to network 0.0.0.0
O*IA 0.0.0.0/0 [110/11] via 10.2.10.1, gi2/0
[110/11] via 10.2.10.17,gi3/0
```

```
router-27#sh ip route 200.1.0.1
Routing entry for 200.1.0.1/32
Known via "eigrp 300", distance 170, type external
router-27#sh ip route 200.9.0.1
Routing entry for 200.9.0.1/32
Known via "eigrp 300", distance 170, type external
```

```
router-27#sh ip route
Gateway of last resort is not set
200.1.0.0/32 is subnetted, 1 subnets
D EX 200.1.0.1 [170/25651200] via 10.2.11.5, gi1/0
[170/25651200] via 10.2.11.13,gi2/0
Thru
D EX 200.9.0.1 [170/25651200] via 10.2.11.5, gi1/0
[170/25651200] via 10.2.11.13,gi2/0
```

```
Router-25#
router eigrp 300
redistribute ospf 200 metric 100 100 255 255 1500
passive-interface default
no passive-interface gi1/0
no passive-interface gi5/0
network 10.2.11.0 0.0.0.255
auto-summary
```




SHARE

Technology • Connections • Results



Routing in the Bigger Picture






SHARE


Technology • Connections • Results

Network Model

- Ideally, the network between data centers would look like this:




- Unfortunately, nothing is ever quite that straightforward....




SHARE

Technology • Connections • Results

The Bigger Picture



3 Layer Hierarchical Model

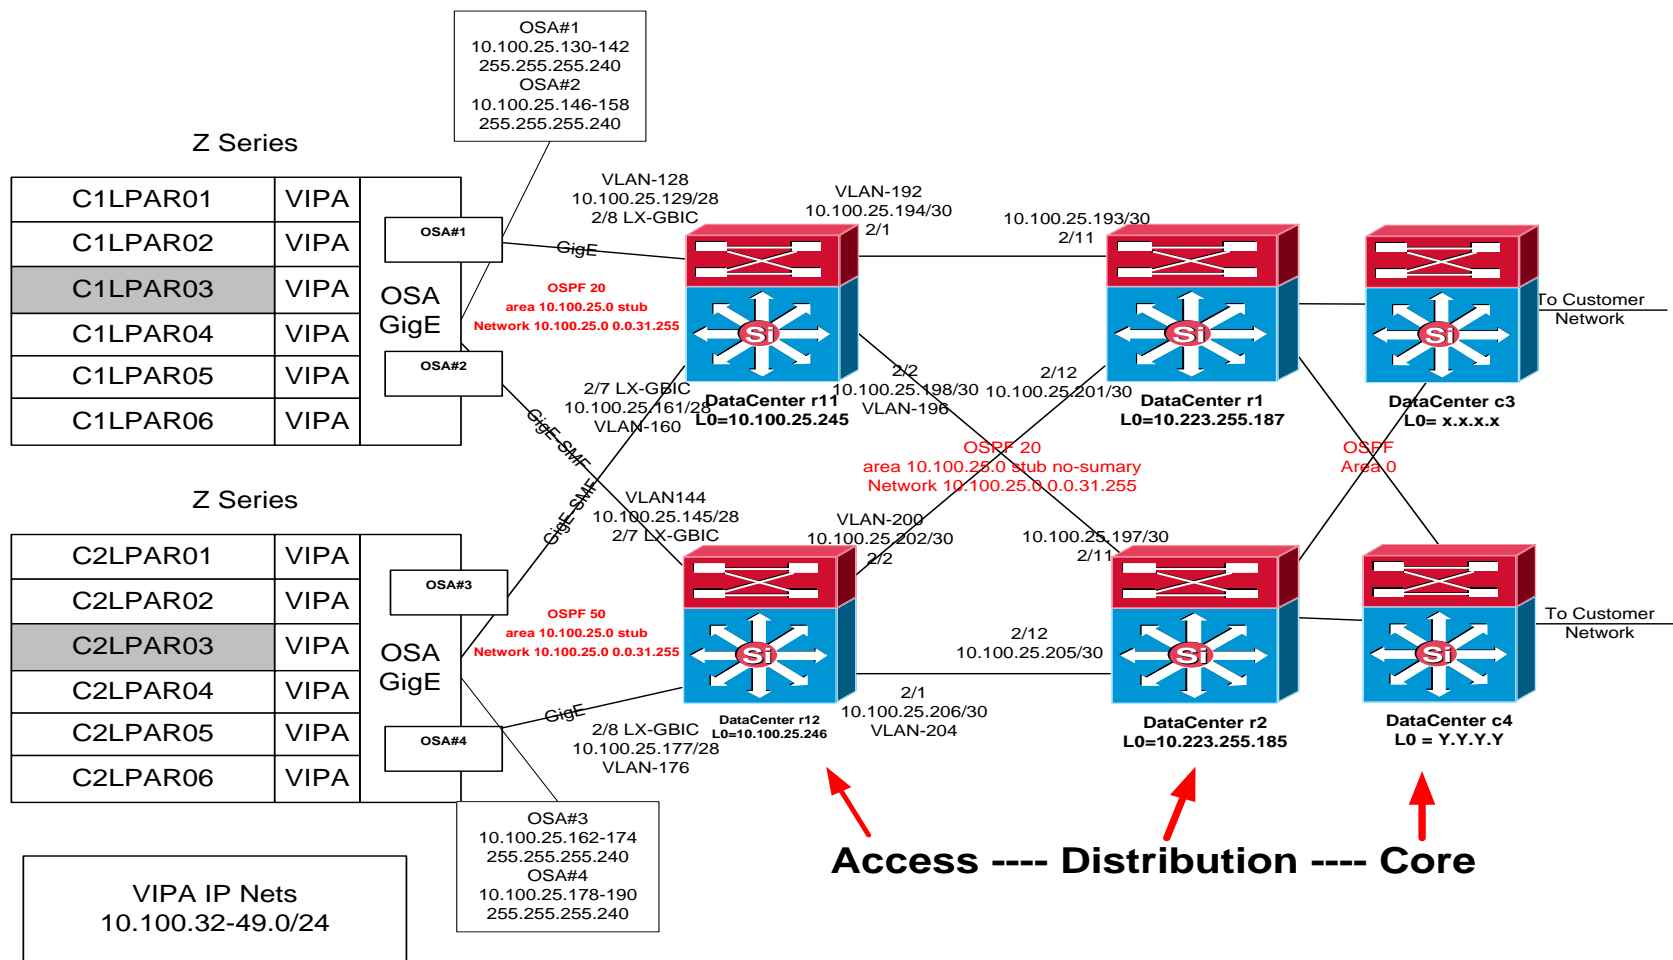




SHARE

Technology • Connections • Results

Mainframe OSA Express Example






SHARE

Technology • Connections • Results



Auxiliary Slides





SHARE

Technology • Connections • Results

MTU Considerations

- Jumbo Frame (> 1500 bytes) use in the data center
- Fragmentation can cause performance issues
- Max packet sizes can vary by platform
- Extra headers imposed with IPSec and GRE tunnels can lead to MTU problems
- 'show ip traffic' command displays the number of packets a router has had to fragment

QoS – Quality of Service

- Consists of classifying, marking, queuing, policing, shaping
- Converged network of voice/video/data depends heavily on QoS
- Most concerned about WAN QoS but LAN based QoS important too – microbursts, worms, etc.

Application	L3 Classification		IETF RFC
	PHB	DSCP	
Network Control	CS6	48	RFC 2474
VoIP Telephony	EF	46	RFC 3246
Call Signaling	CS5	40	RFC 2474
Multimedia Conferencing	AF41	34	RFC 2597
Real-Time Interactive	CS4	32	RFC 2474
Multimedia Streaming	AF31	26	RFC 2597
Broadcast Video	CS3	24	RFC 2474
Low-Latency Data	AF21	18	RFC 2597
OAM	CS2	16	RFC 2474
High-Throughput Data	AF11	10	RFC 2597
Best Effort	DF	0	RFC 2474
Low-Priority Data	CS1	8	RFC 3662



SHARE

Technology • Connections • Results

Nexus 7000

- Data Center switch
- Designed to deliver high density 10 GigE and unified fabric
- Runs Cisco NX-OS – different from IOS but relatively close CLI and configuration commands




MLS – Multilayer Switching

- ASIC assisted routing
- Goes beyond standard CPU interrupt packet processing
- Scales to wire-speed forwarding rates
- RP = Route Processor
- SP = Switch Processor
- RP handles control traffic (ie, routing protocols)
- SP handles data plane traffic
- RP pushes routing information down to SP



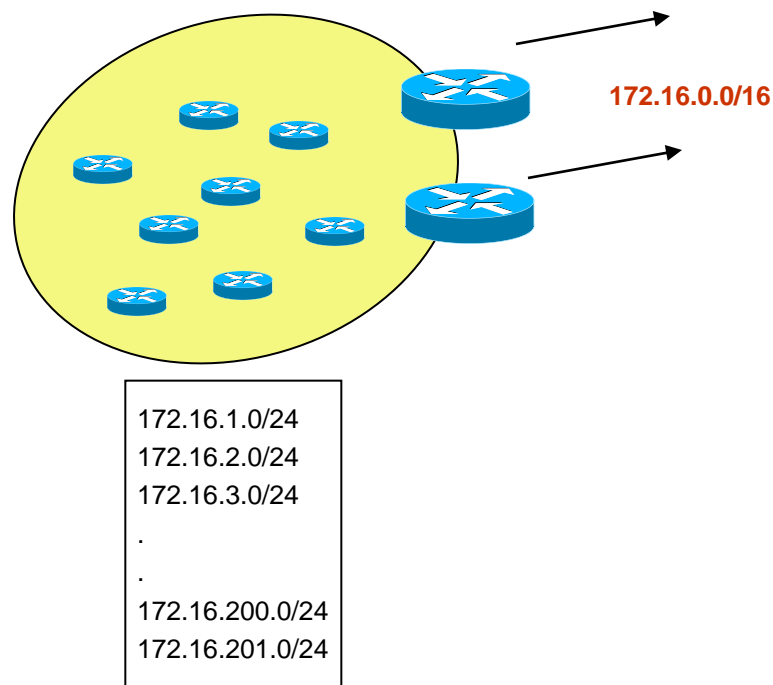
CEF – Cisco Express Forwarding

- IP routing mechanism
- Consolidates IP route table lookup, ARP resolution, and outgoing interface identification
- Performs load-balancing based on Source-Destination address pair and TCP/UDP port numbers.
- Source-Destination pair path is deterministic on a per router basis




Route Summarization


- Summarization can reduce and simplify the routing table
- Useful when the network addressing scheme is well organized



MPLS



- Multi-Protocol Label Switching
- Optimized forwarding methodology
- Scalable – core routers do not have to carry full routing table – just need routes to P & PE routers
- MPLS-VPN keeps separate route tables for each VPN
- P = Provider router (Core)
- PE = Provider Edge router
- CE = Customer Edge router
- C = Customer router



SHARE

Technology • Connections • Results

Contact Information

- Kevin Manweiler
 - kmanweil@cisco.com
 - 919 392-3683
- Junnie Sadler
 - jsadler@cisco.com
 - 919 392-5765

References

- Cisco IBM Alliance
http://www.cisco.com/web/partners/pr67/pr30/partners_strategic_alliance.html
- IBM Redbook: Communications Server for z/OS V1R2 TCP/IP Implementation Guide Volume 4: Connectivity and Routing Oct. 2002
- Route-Maps for IP Protocol Redistribution
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a008047915d.shtml
- Nexus 7000
<http://www.cisco.com/en/US/products/ps9402/index.html>
- IP address
http://en.wikipedia.org/wiki/IP_address
- SHARE S3921 Configuring OSPF for z/OS Communications Server: Part 1 (An OSPF Tutorial)