

BrainBuzz

Cramsession

Last updated October, 2000.
Click [here](#) for updates.

Click [here](#) to see additional documents related to this study guide.

Contents

Contents.....	1
Routing Protocols.....	2
Hybrid Protocols	3
Link State Protocols	10
TCP/IP.....	15
Encapsulation Protocols	17
Routed v. Routing Protocols ..	17
Network Services and Gateways	18
BGP	20

Cramsession™ for Building Scalable Cisco Networks Routing 2.0

Abstract:

This Cramsession will help you to prepare for Cisco exam 640-503, Building Scalable Cisco Networks Routing 2.0. Exam topics include: The Advanced Fundamentals for Implementing, Planning and Configuring a Scalable Cisco Network such as TCP/IP, WAN, LAN, Routing Protocols, Interface Configuration, and Networking.

Notice: While every precaution has been taken in the preparation of this material, neither the author nor BrainBuzz.com assumes any liability in the event of loss or damage directly or indirectly caused by any inaccuracies or incompleteness of the material contained in this document. The information in this document is provided and distributed "as-is", without any expressed or implied warranty. Your use of the information in this document is solely at your own risk, and Brainbuzz.com cannot be held liable for any damages incurred through the use of this material. The use of product names in this work is for information purposes only, and does not constitute an endorsement by, or affiliation with BrainBuzz.com. Product names used in this work may be registered trademarks of their manufacturers. This document is protected under US and international copyright laws and is intended for individual, personal use only. For more details, [visit our legal page](#).

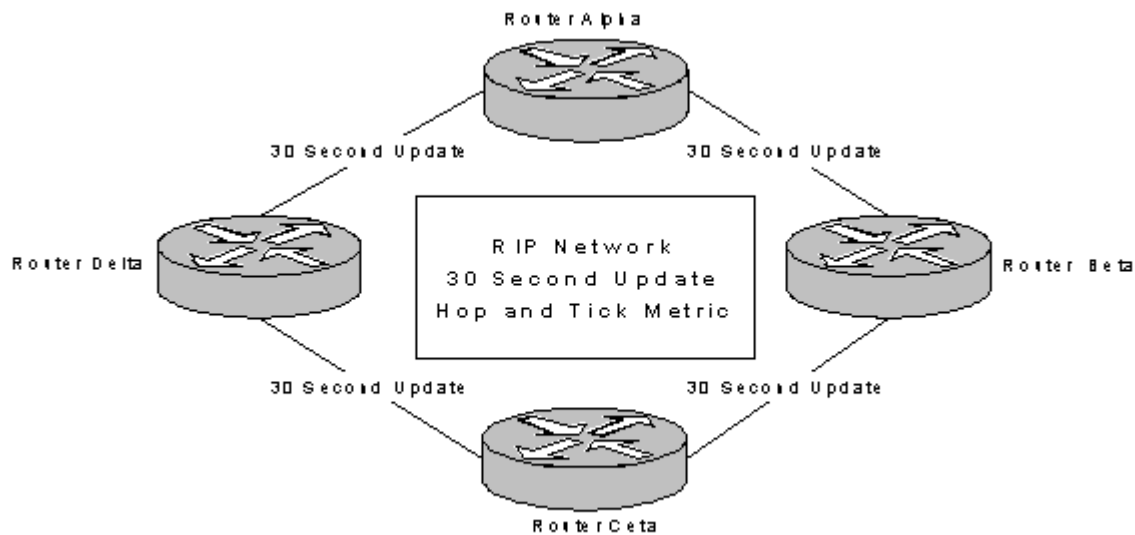
Routing Protocols

Classfull Routing Protocols RIP and IGRP

Link state protocols support classfull addressing as well, but for exam purposes, consider them primarily used in a classless infrastructure. RIP and IGRP do not pass subnet information.

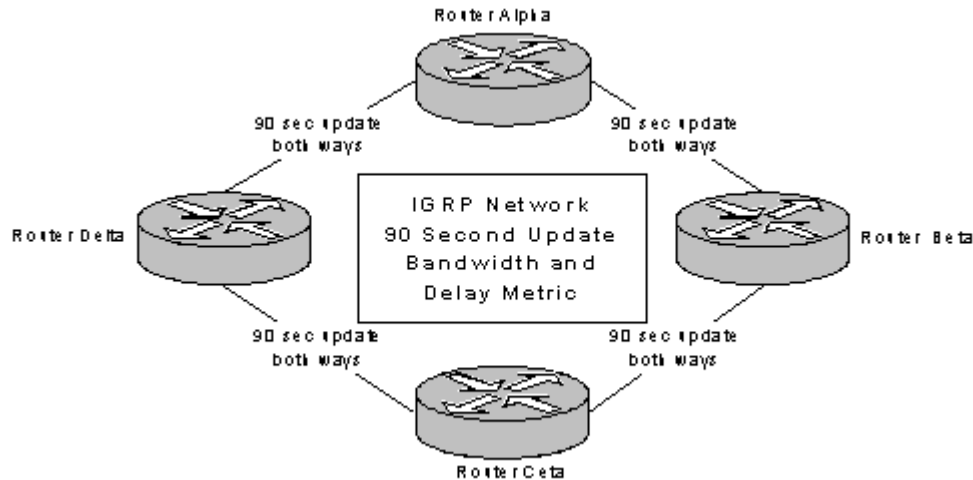
RIP

Route Metric HOPS and ticks. The max hop count is 16. Used in small diameter networks, does not scale well in a very large enterprise environment. Sends its entire routing table out to all its interfaces every 30 seconds.



IGRP

Is a more robust than RIP. IGRP is a Cisco-proprietary classfull routing protocol. IGRP does not support VSLM or discontinuous subnets. IGRP will send its entire routing table every 90 seconds. Over slow links this can be a problem with the amount of bandwidth that would be consumed in updates.



IGRP supports load balancing over paths with unequal bandwidth. The *variance* command is used to load balance IGRP.

```
Router iGRP 90
```

```
Variance 3
```

The default variance is 1 for equal balancing. The variance multiplier can go 1-128.

IGRP uses a "keep alive" timer to determine if links are still up. Triggered updates cause the IGRP router to send a routing table update out of its interfaces. Periodic updates still get broadcasts sent to connected routers every 90 seconds.

Routing loops

Distance vector protocols use the following to prevent routing loops:

Poison Reverse	Split Horizon	Holddown Timers
----------------	---------------	-----------------

Hybrid Protocols

EIGRP

Considered a hybrid protocol. It combines the best features of both link state and distance vector. It can detect a link failure within one second. It converges rapidly and scales well into large networks. EIGRP sends routing updates to directly connected neighbors; only changes are sent rather than the entire routing table.

Bandwidth is the primary metric. That is why it is important to specify the bandwidth on an interface. EIRGP has a default for serial interfaces of 1.5Mbps. So you must set the correct bandwidth with the *bandwidth* command.

Advantages to EIGRP:


- VSLM Support
- Supports Discontiguous Subnets
- Multi-Protocol Support
- Automatic Route Summarization
- Keeps copy Neighbors Routes
- Load Balances up to 6 Paths

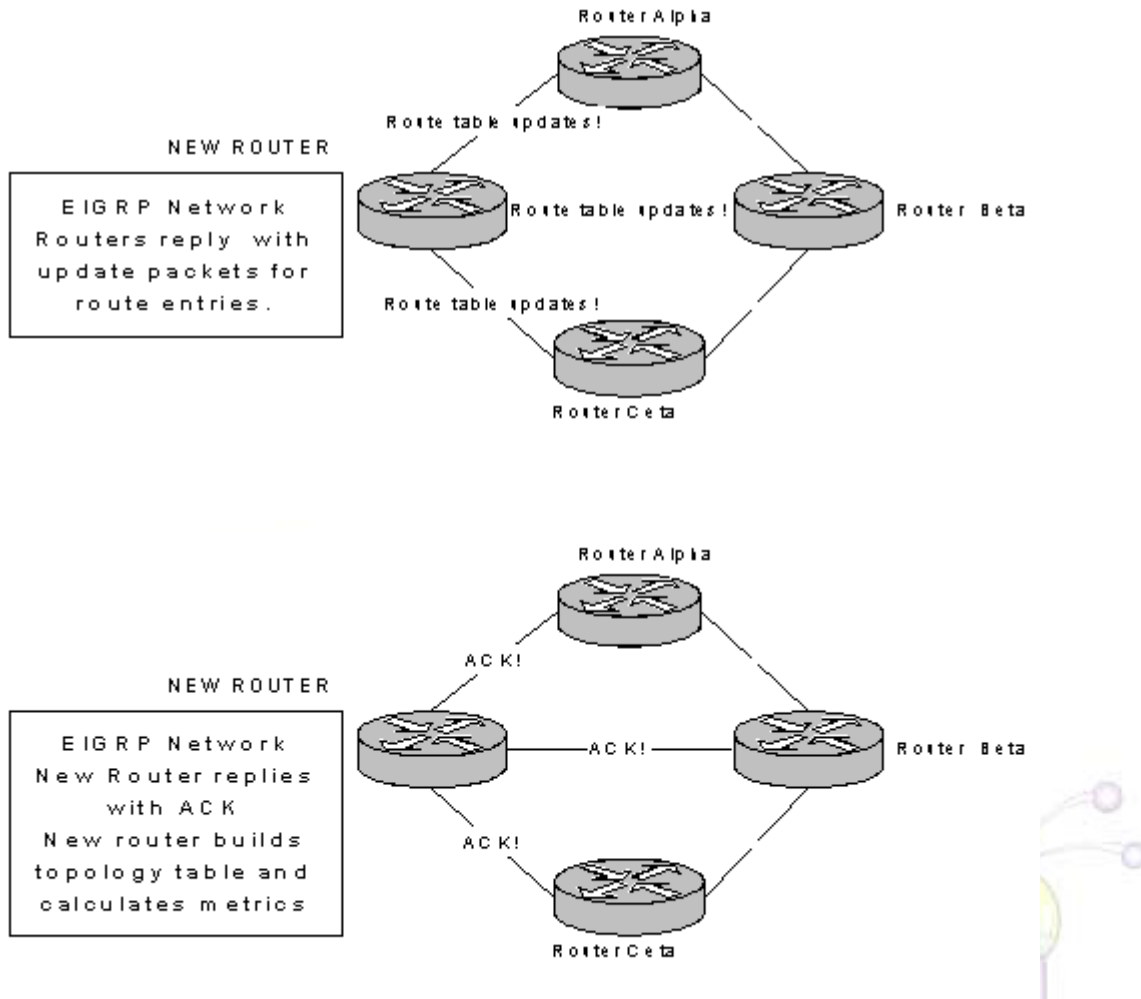
DUAL

Diffusing Update ALgorithm – The routing engine behind EIGRP. It allows for routers to update at the same time and allows for multi-protocol routing. It tracks route updates sent by neighbors and ensures against black holes.

EIGRP Discovery Process

EIRGP sends hello packets out of all of its interfaces to find its neighbors.





The new router then exchanges update packets with the adjacent routers. The adjacent routers then send ACKs to the new router. When this step is completed the new router will then choose its routes (successors).

Route Summarization for EIGRP

This is on by default, but only at the network or class boundaries. Manual configuration for route summarization is primarily done at the core or distribution layers. An example of the syntax for configuration is as follows:

```
ip summary-address eigrp 100 10.98.0.0 255.255.0.0
```

Important! By default, EIRGP does not support VSLM. You must use the *no auto-summary* command. Summaries are done at the interface level rather than at the router level. Use *no auto-summary* command to enable support of discontinuous subnets.

Example:

```
router eigrp 100
network 10.0.0.0
network 192.64.0.0
no auto summary
interface serial 1
ip address 10.98.98.24 255.255.255.0
bandwidth 128
ip summary-address eigrp 100 192.64.0.0 255.255.0.0
```

EIGRP Tables

EIGRP keeps three active tables in its database.


- **Topology Table** – Is kept for each protocol. It is a table of all the route entries the router has learned. Syntax: *show ip eigrp topology* shows the number of successors, active or passive, and distance to destination.
- **Neighbor Table** – EIGRP keeps a table for each adjacent router and one for each protocol. Syntax: *show ip eigrp neighbors*
- **Routing Table** – Each protocol has its own routing table. EIGRP calculates the best route, or successor, from the topology table and puts the entry in the routing table. Syntax: *show route eigrp*

EIGRP Traffic

You can use the *show ip eigrp traffic* to view traffic statistics. This includes hello packets, updates, ACKs and replies.

EIGRP Configuration

```
EIGRP Network Configuration
interface ethernet 1
ip address 192.32.2.1 255.255.255.0
interface ethernet 2
ip address 192.32.64.1 255.255.255.0
router eigrp 500
network 192.32.0.0
```



Default routes

You can also configure the route to forward to a default network if a route is not found in the routers routing table. Use the following syntax:

```
router(config)#ip default-network 10.98.98.1
```

Static routes

Can be used between two AS's or in a DDR situation. You define specific routes between the AS's so you do not have to spend bandwidth on routing updates.

```
router(config)# ip route 10.98.8.0 255.255.255.0 172.16.32.0 e0 permanent
```

10.98.8.0 255.255.255.0 is the destination

172.16.32.0 is the next hop address.

Passive interface

You can configure EIGRP with a passive interface, which will not allow it send any routing updates including hello packets. Use the following syntax:

```
router(config-router)#passive-interface e0
```

Route filters

It is also possible to filter out inbound and outbound routing updates.

Outgoing Syntax:

```
router(config-router)# distribute-list 101 out e0 static
```

101= the access-list-number

out = filters for outbound

e0 = the interface

static = name of route process

Incoming Syntax:

```
router(config-router)# distribute-list 101 in e0
```

in = filters for inbound

e0 = the interface

Route Redistribution with EIGRP and IGRP

Routes can be redistributed within routing protocols. With EIGRP and IGRP this process is automatic if the AS systems numbers are the same. If the numbers are not the same then the redistribution will have to be done manually with the default-metric command:

router(config-router)# redistribute protocol X metric Y

router(config-router)# default-metric bandwidth delay reliability loading mtu

Route Redistribution with EIGRP and OSPF

Determine the core routing protocol.

Locate the router or ASBR where the route distribution needs to take be configured.

Make the decision of which is your short term or edge protocol.

Determine the route process where you want the route distributed.

Example: ***router(config)# router ospf X***

X= the process id

Use the redistribute command to populate from the short term or edge protocol to the core routing protocol.

Example: ***router(config-router)# redistribute protocol X metric Y***

X= the process id and Y = the metric value

router(config-router)# redistribute ospf 213 metric 110 subnets

router(config-router)# default-metric X

X= the metric value wanted

Define the seed metric to be used in the calculation of the value of the route before distribution.

For EIGRP

Router(config-router)# default-metric bandwidth delay reliability loading mtu

Administrative Distance

Advertised Distance

Advertised Distance is the distance a neighbor router says, or advertises, is the distance to a destination. This is key in electing a feasible successor or backup route. The lower the distance, the better. The lower the value the better the route is believable. These values can be changed with the *distance* command.

Administrative Distances

Protocols	Distance Value
Connected Interface	0
Static routes	1
EIGRP Summary routes	5
EBGP	20
Internal BGP	90
IGRP	100
OSPF	110
RIP	120
EGP	140
External IEGRP	170
Internal BGP	200

EIGRP and Dropped Links

When EIGRP discovers a down link the router does a lookup in its topology table for a successor. If a suitable one is found, it changes to the new route. The router then does a recalculation for the next successor. If one is not found, the router begins a new process.

The failed route or link is moved into active status and sends query packets to its adjacent neighbors. It multicasts out to every interface except the one it was learned from.

The router then waits for a response. If the adjacent router does not have a route to the destination, nor has a feasible successor, it sends a unicast packet back to the troubled querying router. If the neighbor router is using the troubled query router as a successor, it sends its own query packet to its neighbors. This query will propagate the entire AS.

If the troubled query router receives a successor, or feasible successor, the data is changed in the topology table and the router waits to receive more responses. The routing table is recalculated and the route is returned to passive state.

If no route updates for a successor or feasible successor are found, then the router deletes the entry from the topology database and routing table. If routers do not respond within 180 seconds their routes are also put into active state, and the query router begins to look for the additional routes it lost through the lost router.

EIGRP and NBMA

Take particular care when configuring EIGRP in a NBMA environment.

Three NBMA Scenarios:

- Pure point-to-point configuration – each PVC on subinterfaces
- Multi-point configuration – no subinterfaces
- Hybrid (mixed) multi-point and point-to-point

In these scenarios it is important that:

- The traffic bandwidth given to EIGRP on each VC must be the same in both directions;
- The total EIGRP traffic for the sum of all the VC's must not exceed the line speed of the interface; and
- The EIGRP traffic allowed on one VC must not exceed the capacity of the VC.

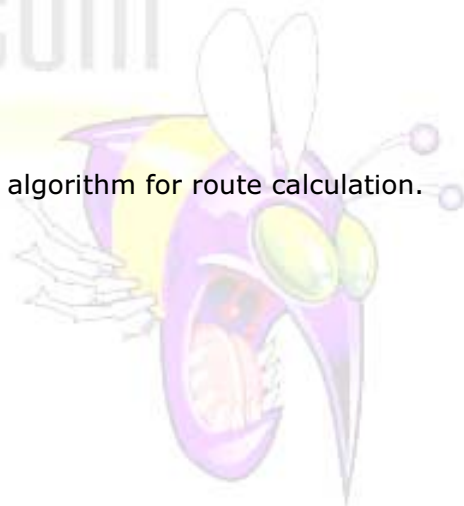
Link State Protocols

OSPF (Open Shortest Path First)

Is a link state routing protocol that uses Dijkstra's algorithm for route calculation.

OSPF has several advantages:

- VSLM Support
- Fast Convergence
- No Hop Count Limit
- Route Selection Based on Cost
- Low Bandwidth Usage



OSPF Commands (Single Area) Setup

router (config)#router ospf # (#=process ID)

- Enables ospf on the router

router (config-router)#network address wildcard-mask area # (#=area id)

- Address can be a subnet, network or the address of the interface
- Selects the networks that will be in the OSPF network

router (config-if)#interface loopback #(#=the loopback address)

- OSPF has been proven more reliable with a loopback number
- Loopback address can override the highest ip address for the router id


router (config-if)#ip ospf priority 0-255

router (config-if)#ip ospf cost # (# = the cost value 1-65535)


Cost values Ethernet=10, T1=128, 56k serial=1785

OSPF in a Single Area

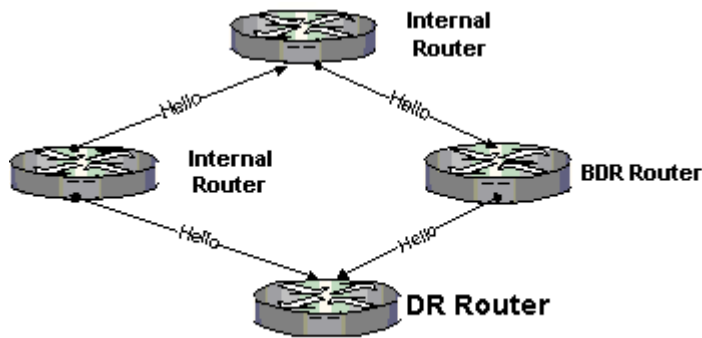
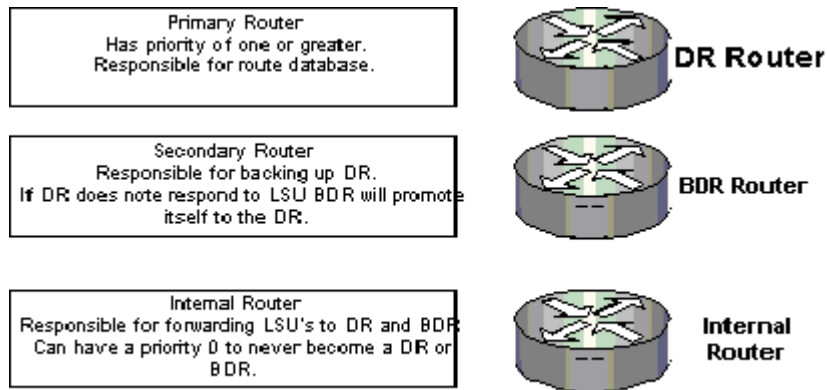
OSPF Exchange Phase
 1. DR sends the new router its routing information
 2. New router sends an ACK to confirm delivery
 3. New router begins compiling route table. This is the Loading Phase



OSPF New Router Ex Start
 1. Router sends hello packets out its interfaces develops adjacencies
 2. Develops adjacencies
 3. Router begins the Exchange Phase



Once the loading phase is completed and the router is ready to join the OSPF network, the router is in Full State.



Routers exchange hello packets every 10 seconds to verify links are up. The Dead interval is 4X the hello interval.

SPF hold time - Is the time the router waits before doing a SPF route calculation. The default hold time is 10 seconds.

Neighbor relationships will vary with the connection types. Different OSPF configurations are required for each.

NBMA - Non-Broadcast Manual Configuration needed. Static lists for routes must be created.

Router must poll for the routing information since broadcasts are not allowed. Use the neighbor command to accomplish this.

neighbor ip-address X Y [X=priority #] [Y=poll-interval]

OSPF Multiple-Areas Configuration Commands Enable OSPF on the Router

router(config)# router ospf X (X=process id)

- Next step is to tell the router which networks are on the OSPF network

router(config-router)#network address wildcard-mask area (area id)

LSA1 – Router Links LSA – Sends information about the routers links.


LSA2 – Network Link LSA – Sent by the DR to all routers in the AS. A list of routers in the segment.

LSA3 – Summary Link LSA – Sent by ASBR's list of networks available outside the area.

LSA4 – Summary Link LSA – Sent by ASBR's list of networks available outside the area.

LSA5 – External Link LSA – Sent by ASBR's list of external network routes.

OSPF recalculates a new table when a route goes down. So, if you have a link flapping you may want to increase the amount of time to wait. Use *spf holdtime* command. If not, it could overload CPU and cause performance issues.



LSA's are sent and an ACK is expected as a response. If an ACK is not received it will resend in 5 seconds. This can be changed with the command: *ip ospf retransmit-interval 9* (changes retransmit time to 9 seconds)

Area Types – The area types determine what kind of LSA's the area will receive.

Stub Area – Does not accept External LSA's. LSA Type 5's are rejected. Can except route summaries.


Totally Stubby Areas – Will not except any LSA's with external or summaries.

Internal Routers – Exchange LSA's 1 and LSA's 2. They share the same routing database and all interfaces are within the same area.

Backbone Routers BBR – Exchange LSA's 1 and LSA's 2. Share at least on interface in the area 0.

Area Border Router ABR– Share an interface with another OSPF area. This router keeps a database for each area.

Autonomous System Border Router ASBR – Have at least one interface in a non-OSPF network. Its uses LSA 5's to distribute this routing information into the OSPF network.



Route Table Updates

Routers take Type 1 and Type 2 LSA's and make their own route calculations. These entries are used to modify the routing table.

The routers will also take Type 3 and Type 4 LSA's and use them to calculate a route for areas in the internetwork. When a router has both an inter-area route and intra-area route are always kept. An easy way to remember this is, "Why go around the block to do next door?"

Only stub area routers do not expect type 5 LSA's. All other routers perform path calculations to external AS's.

Commands for Stub Area Configuration

router(config-router)#area area-id stub

- Configures Regular Stub area

router(config-router)#area area-id stub no-summary

- Configures A Total Stub Area

Commands for Route Summarization on OSPF

For ASBR's:

router(config-router)#summary-address address mask

- Condenses inter-area routes into summary

For ABR's:

router(config-router)#area area-id range address mask

- Condenses inter-area routes into summary

Commands for Troubleshooting OSPF

router#show ip route

- Gives the route information learned by the router.

router#show ip protocol

- Router information along with metrics and networks. Used to verify how OSPF is configured.

router#show ip ospf

- Displays how many times the SPF algorithm was calculated and update interval time.

router#show ip ospf interface

- Displays hello interval, adjacencies, and the ospf area id.

router#show ip ospf neighbor detail

- Shows the list of neighbors, DR and BDR info priorities and states.

router#show ip ospf database

- Displays the database topology, link state database, router id and ospf process id.

router#show ip ospf border-routers

- List the ABR's in the AS.

router#show ip ospf virtual-links

- Shows the status on all the virtual links.

TCP/IP

TCP/IP is a widely used routable protocol, and its biggest challenge is proper management of addressing, security and broadcast management.

Private Addressing

The usual address prefixes are 10, 172 and 192. Used for private networks not openly exposed to the Internet (inside a firewall).

Public Addressing

Assigned by an ISP, not recommended for private networks. Private to public network communication can be accomplished by NAT through a PIX or other firewalls. Options also include VPN (Virtual Private Networks) or extranets secured through PPTP (Point-To-Point-Tunneling-Protocol) and/or L2TP (Layer 2 Tunneling Protocol).

Hierarchical Addressing

Using an address scheme where the different network numbers determine whether a destination is local or remote. Longer subnet masks are used at the access layers. Network prefix gets smaller as you move up the network hierarchy.

Prefix Routing

This is how a router forwards packets. Router uses the network number to make routing determination.

Classfull Addressing

This addressing scheme is commonly used where the subnet mask reflects the number of bits used to calculate the default gateway (e.g. Class A 10.0.0.0 mask 255.X.0.0, Class B 172.0.0.0 255.255.0.0, Class C 192.0.0.0 255.255.255.0). RIP and IGRP can only be used with a classfull addressing scheme.

Classless Addressing

CIDR - Classless Inter-Domain Routing - is used to conserve and use address space effectively (see VLSM). This is required for route summarization to work correctly. Careful planning and implementation are required. An easy way to identify a classless address is to look at the subnet mask. You will commonly see a Class A address with a Class B or C subnet mask. Some protocols require additional configuration to support discontinuous subnets. Link state protocols support classless addressing. RIP and IGRP do not, because they do not pass subnet information.

VLSM

Variable Length Subnet Mask - classless addressing allows using, for example, a Class B address with a Class C subnet mask. Usually summarized in this fashion 172.98.98.24/30. "30", or 255.255.255.252, specifies the number of bits used to calculate the network portion. This allows effective use of your IP addresses and should only be used with routing protocols that support VLSM, like IEGRP and OSPF.

Secondary Addressing

Is assigning a second IP gateway address for the same interface on a router. This is not recommended as a good practice and should be used only when you have to.

Encapsulation Protocols

Authentication

CHAP or PAP - CHAP is encrypted; with PAP login and password information are sent in plain text.

GRE

Generic Routing Encapsulation – Used primarily in the backbone. Can be used to tunnel IPX or AppleTalk. Fast switching supported.

NWLINK


Used to encapsulate NetBIOS over IPX. Requires type 20 packets to operate properly. Use the *ipx type-20-propagation* commands on the interface.

NBT

Used to encapsulate NEBIOS over TCP/IP.

AURP

AppleTalk Update Routing Protocol – Encapsulated in TCP/IP over WAN links. Sends updates only like EIGRP.



Routed v. Routing Protocols

It is important to distinguish the difference between *routed* and *routing* protocols. Routing protocols use metrics, hop counts, ticks, etc. to make a routing decision. Since routers do not forward broadcasts, routers separate networks into different broadcast domains. Switches and bridges separate media into separate collision domains. Routers are responsible for:

- Switching and/or relaying packets
- Path determination

IPX on the WAN

Use NLSP for faster convergence over IPX/RIP and reduced routing traffic. It uses cost as calculation metric and is more CPU intensive. NLSP redistributes RIP but retains 15-hop limit. NLSP supports up to 1023 hops.

EIGRP for IPX

Increases bandwidth by only sending updates over the WAN and full updates over the LAN. When a route goes from IPX/RIP to EIGRP it increases the hop count by two. From EIRGP to IPX/RIP, the route tick count is unchanged.

AppleTalk

RTMP

AppleTalk's version of a routing protocol. Very similar to RIP broadcasts entire table in 10 Seconds. Max Hop count is still 15, uses split horizon.

Design Rule

Uses EIRGP for routing AppleTalk.

EIGRP

Saves bandwidth because only updates are sent. Fast convergence.

Network Services and Gateways

DHCP

BOOTP server used to assign IP addresses to requesting clients. Can be configured to specify node type, WINS, DNS and other information.

There are several options for DHCP configuration. Cisco offers IOS features to forward DHCP packets. **ip helper-address* command forwards broadcasts to DHCP servers like an NT server.

WINS

Windows Internet Name Service – is a static addressed server that performs NetBIOS-name-to-IP-address resolution. It resolves NetBIOS names to IP addresses, which takes away the need to ARP (broadcast) to resolve network names. After booting and obtaining a DHCP IP address, the client sends a unicast packet to the WINS server requesting to register its NetBIOS name. DNS servers and WINS servers (sometimes on the same server) work together to resolve name lookup. Acts as a register for windows machines.

DNS

Application server that provides Internet-name-to-IP-address conversion. Windows DNS servers can be directed to query a WINS server for NetBIOS names.

Helper Addressing Syntax

ip helper-address X.X.X.X (where X.X.X.X is the destination IP)

no ip helper-address X.X.X.X (where X.X.X.X is the destination IP) to remove.

Multiple ports and servers can be configured:

int e1

ip address 10.98.98.24 255.255.255.0

ip helper-address 172.98.65.2

ip forward protocol 69


ip helper-address 192.54.12.8

ip forward protocol 2020

NAT

Network Address Translation

Can be used to merge two large networks without having to re-address the whole network. Another function of NAT is overloading inside global addresses. This process allows several inside addresses to use a single IP address. NAT can also use a pool of addresses or multiple interfaces.



BGP

Border Gateway Protocol

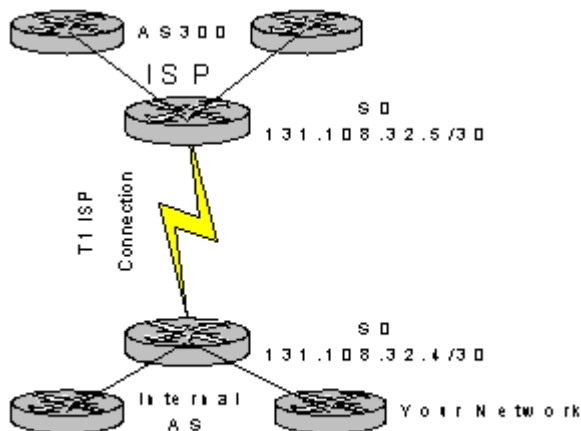
BGP is an EGP (exterior gateway protocol) it is used to pass routing information between autonomous systems. It is still a routing protocol and, like the other routing protocols, it passes routing information and uses a metrics for route determination. It functions to advertise which networks can be reached. BGP can act as an IBGP or EBGP. Which means it can be configured to advertise networks within an AS or between different AS's. The trick to BGP is you have to tell BGP which networks to advertise whether they are directly connected or not. BGP uses TCP (port 179) to communicate with other routers.

To avoid having to manually update several routers with a single BGP change you can use the peer-group command to group router together. This enables you to implement a common policy, distribute-list, route-maps, and update policies. Use the following syntax:

```
#neighbor peer-group-name peer-group
#neighbor ip-address peer-group peer-group-name
```

Common BGP Implementation

Establishing a connection to ISP. Most of the configuration choices for you will be made for you by your ISP.



Like OSPF, a BGP-configured router will first exchange the entire routing table. Once a peer has been established only changes are sent. BGP routers also send keepalives to make sure connections are still active. Like OSPF and EIGRP, BGP uses AS numbers. BGP uses a single metric for path selections. The network administrator manipulates the metric value to achieve the desired result.

BGP Metric Criteria

Speed Delay Cost Stability # of AS's passed through

Traits of BGP Path Selections

- Use routes with a higher BGP administrative value first
- When routes have identical weights use the route with the higher local preference
- When routes have identical local preference use the route that the server originated
- When the next hop is not accessible DO NOT consider it
- When AS paths are the identical length choose the external path over an internal path
- Use the shorter AS path if no route was originated
- When all routes are external use the route with the lowest origin code
- Use the path with the lowest MULTI_EXIT_DISC metric if the origin codes are the same and the paths came from the same AS
- Use the route with the lowest ip address value for the BGP router ID
- Use the path through the closest neighbor when IGP is disabled and only internal paths are left

BGP Commands

Network command - used to tell BGP which network to advertise.

Setting up BGP on a router

1. router bgp AS#
2. network ip-network# mask network-mask

#router bgp 300

#network 131.108.0.0 mask 255.255.0.0

Neighbor command - used to establish the BGP peer. This is used for IBGP and EBGP. (**Note the difference in the AS number to tell if it is an EBGP configuration or an IBGP configuration.)

External BGP Configuration:


Establishing EBGP Peer:

1. router bgp AS#
2. network ip-network# mask network-mask
3. neighbor ip-address remote-as AS#

Router Dallas

#router bgp 200**#network 131.108.0.0 mask 255.255.0.0****#neighbor 131.108.32.5 remote as 300**

Note AS numbers are different



```
External BGP
Router Dallas
#router bgp 200
#network 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.5 remote-as 300
*Note AS numbers are different*
```

Internal BGP Configuration:

Establishing IBGP Peer:

router bgp AS#**network ip-network# mask network-mask****neighbor ip-address remote-as AS#**

Router Austin

```
#router bgp 200
#network 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.4 remote-as 200
```


Note AS numbers are the same

Router Dallas

```
#router bgp 200
#network 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.5 remote-as 200
```

Note AS numbers are the same

The similar AS number on the last line indicates to the BGP router that the neighbor is an internal BGP neighbor, as opposed to the external configuration, where the BGP network number and the remote AS were different.



Internal BGP

```
Router Austin
#router bgp 200
#network 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.4 remote-as 200
*Note AS numbers are the same*
```

Internal BGP

```
Router Dallas
#router bgp 200
#network 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.5 remote-as 200
*Note AS numbers are the same*
```

Shutdown command – used to disable a peer.

To stop (disable) the peer relationship to Router Austin:

```
#neighbor ip-address|peer-group-name shutdown
```

```
# neighbor 131.108.32.5 shutdown
```

No command - to reverse the effect of the command use the **no** command

```
#no neighbor ip-address|peer-group-name shutdown
```

```
#no neighbor 131.108.32.5 shutdown
```

Reset Connections Commands

To clear the BGP database, cache or table use the following commands.

****Note they all start with clear ip bgp****

```
#clear ip bgp ip-address
```

- To clear a single bgp connection

```
#clear ip bgp peer-group tag
```

- To clear all the members of the bgp peer group

```
#clear ip bgp
```

- To clear all the bgp connections

```
#bgp fast-external-failover
```

- To clear sessions for external peers directly connected use

BGP Control Commands

Route flapping can cause havoc on your BGP configurations. Use the following commands to control updates sent and enable dampening.

```
#bgp dampening
```

- Enables route dampening for BGP

```
#clear ip bgp dampening address mask
```

- Use the clear command to reverse dampening

```
#show ip bgp flap-statistics
```

- Use the **flap-statistic** command to show flapping routes

```
#clear ip bgp flap-statistics
```

- Use the **clear** command to clear the statistics

BGP Statistic Commands

#show ip bgp summary

- Use the **summary** command to display the status of all BGP connections

#show ip bgp paths

- Use the path command to view the BGP database

#show ip bgp neighbors address

- Use the **show neighbors** commands to a detailed list of the bgp neighbors and the TCP information

BGP Path Filtering by Neighbor

You can path filter by configuring an access list and applying that access list using the **as-path** command. BGP filtering is done with the **filter-list** command. **** Note that filtering by neighbor uses the as-path command with expressions (and filter-list) and route filtering uses the distribute list command. Both use access list****

Use the following syntax to filter by neighbor.

*Note: an access list must be configured first!

#ip as-path access-list access-list-number# permit|deny as-expression

Example:

```
#router bgp 200
#network 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.5 remote- as 300
#neighbor 131.108.32.5 filter-list 1 out
#ip as-path access-list 1 permit ^$
```

^\$= AS-expression



BGP Route Filtering

This syntax will filter by route filter by neighbor. This configuration uses access list and the distribute list command.

*Note: an access list must be configured first!

#access-list access-list# permit|deny source source-mask

Use the following syntax:

```
#neighbor ip-address distribute list assess-list# in|out
```

For Example:

Router Dallas

```
#router bgp 200
#network 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.5 remote-as 300
#neighbor 131.108.32.4 remote-as 300 distribute-list 1 out
```

Backdoor Routes

Routes that you do not want advertised but used by a border router are called backdoor routes. To enable a backdoor route use the following command.

```
#network address backdoor
```

BGP Summary Routes

Use the aggregate-address command to summarize network routes.

Syntax:

```
#aggregate-address address mask
```

or

```
#aggregate-address address mask summary-only
```

There are more variations of this command.

Distributing Network 0.0.0.0 into BGP

To allow the distribution of network 0.0.0.0. Which is not allowed by default. Use the following syntax:

```
#default-information originate
```

BGP Route Distribution

Static Routes BGP – Is one method of dealing with flapping networks and preventing BGP instability. The drawback with static routes is BGP will show the route to be active even if the route is down. Static routing with BGP enables the route to always be advertised and always in the routing database. Use the **redistribute static** command. Use the following syntax to distribute static routes.

```
#router bgp 100
#neighbor 131.108.0.0 mask 255.255.0.0
#neighbor 131.108.32.5 remote-as 300
#redistribute static
ip route 10.0.0.0 0.255.255.255 null 0
```

null 0 is a null interface. With the configuration above it will cause any packet destined for the 10.x.x.x to be discarded.

Default Route (Gateway)

0.0.0.0 is the default gateway. It is also the gateway of last resort. It is easier to implement this to an interface on your border router going to your ISP. Use the following syntax:

```
#ip route 0.0.0.0 0.0.0.0 s1
```

Dynamic Route Distribution

To distribute IGP routing information use the **redistribute protocol process-id** command. See the earlier notes in the EIRGP route distribution section.

To prevent any routes from being advertised to BGP use the **passive-interface** command.

Route Filters

Route Maps – used by BGP to control which routes are advertised between routing domains. They typically used in a community and are used for filtering. The map-tag command refers to the name of the route map.

Syntax:

```
#route-map map-tag permit|deny sequence#
```

BGP Communities

BGP routers can be grouped into logical groups known as communities.

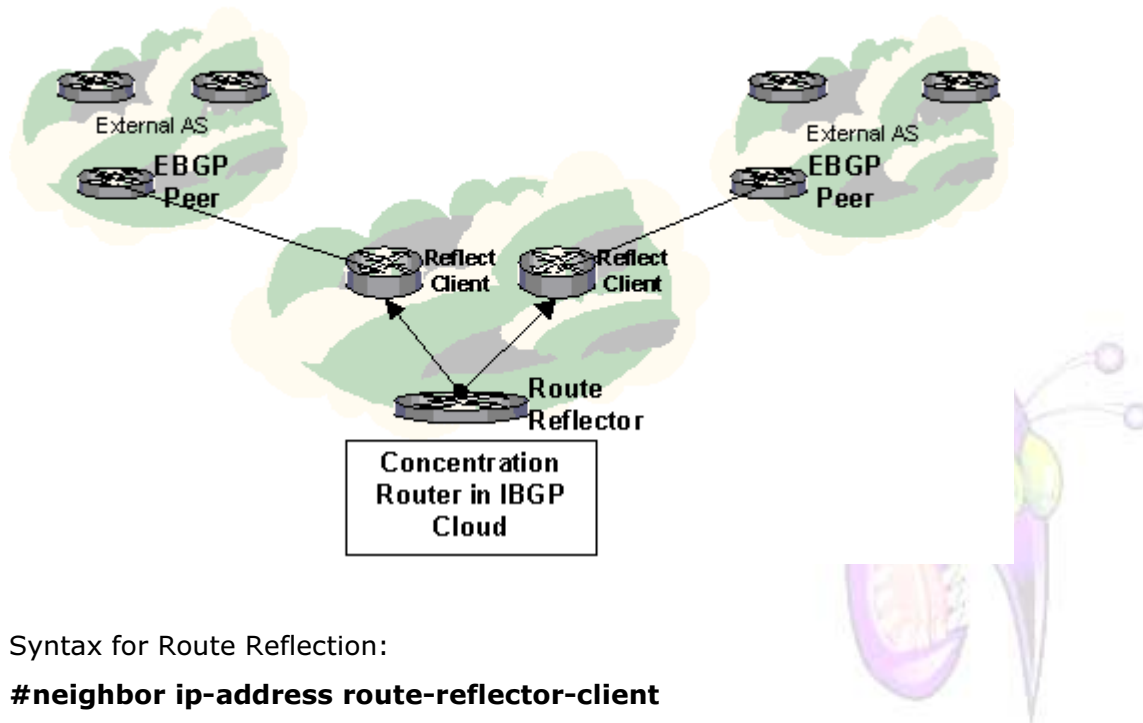


Route Reflectors

As you can tell by the multiple BGP commands and options, configuration for BGP can get quite cumbersome (not to mention making sure all the routes have been configured correctly). External AS's make this problem even more cumbersome. To help reduce the amount of peers a router will have in an AS, there are route reflectors. The routers peer with another router, or concentration router. These

peers are then known as Clients. They receive updates from and send updates to the route reflector.

In an BGP AS, there can be several routers 40+ exchanging route information. To reduce the amount of peers a router has to establish, a concentration and or route reflector router can be configured. Reflect clients are configured to exchange information with a central concentration router (route reflector). So reflect clients only peer with the central concentration route instead of every router in the AS. The central router (route reflector) is responsible for sending updates to the reflector clients. The reflector clients are responsible for sending updates to the central concentration router instead of every router in the AS.



Syntax for Route Reflection:

#neighbor ip-address route-reflector-client

disables client to client reflection

#no bgp client-to-client reflection

Confederations (route reflection on steroids)

A confederation is a large single autonomous system split into smaller sub-autonomous systems. The other networks see the confederation as a single AS. The sub AS's are transparent to the outside world. The advantage to this is the same as route reflection: it reduced the number of peers IBGP.


Syntax for BGP Confederation for Router Z:

#router bgp 275

#bgp confederation identifier 300

```
#bgp confederation peers 250 300 350
```

```
#neighbor 131.108.32.4 remote-as 300
```



Modifying Parameters and Administrative Distances for BGP

How to apply route maps to the routing table:

```
#table-map route-map name
```

Changing administrative distances:

BGP uses three administrative distances local, external and internal.

```
#distance bgp external-distance internal-distance local-distance
```

The default timer is 60 seconds for keep-alives. BGP sends out periodic keepalives to make sure routes are still up. If the route is down then the holddown timer is used before declaring the route dead.

To adjust the time, use the following syntax: **#timers bgp keepalive holdtime**

- This changes the times for all bgp peers.

Special Thanks to [Julian Laredo](#) for contributing material for this Cramsession.