

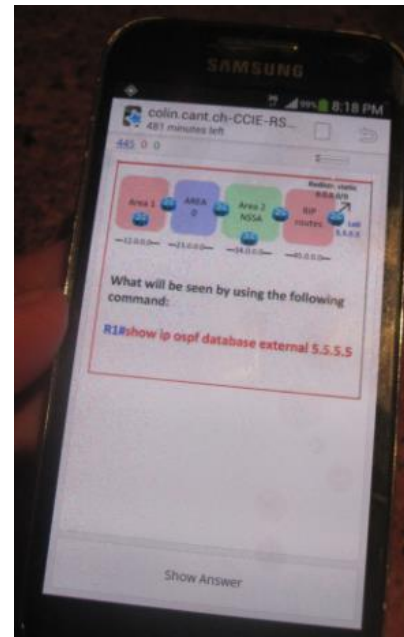
1600+ Cisco CCIE RS ver 4 / 5 / CCNP RS study flashcards

Best viewed with an iPad or similar!

As I was going through the entire CCIE training material and had found that I tend to forget some of the many details I had studied month earlier and had to come up with a solution on how to remember all the nibbly details and keep the learned fresh. This file contains over 1600 designs, config snippets, explanations of command output, and handy debug commands to keep in mind. Based on the "APP Free study cards" I have created the ANKI version, so people can go through the flashcards on their mobile phones etc while commuting to work and back again, utilizing that time as study time. I am still working on my CCIE number, therefore this document is subject to change without notice, I keep adding things I think one could easily forget etc or is just generally good to know. Keep an eye on the "revision number" on the top left to see if I had made any changes since you last visited the file. If my "APP Free" card deck had a great impact on your CCIE trail, please feel free to let me know and post me your CCIE number!

Have fun studying!

ANKI version for mobile devices:



For all the folks who rather use the "APP Free" study flashcards on a **mobile device**, please download the **ANKI version** of the cards Found here:

<http://www.flashcardguy.ch>

There are two versions, one with all the 1600 cards in one deck, and another ZIP version where I have separated each technology into a separate ANKI file which might be easier, better to use initially.

ANKI for mobile devices found here: <http://ankisrs.net/>

Please use ankisrs support / forums if you are facing problems running ANKI on your mobile devices. Thanks




"APP Free", the classic way:

Instructions "APP Free" study flashcards

Print them A3 colored, cut them in rows, pre-fold them in rows, glue them together, cut the single cards from the row.
Mark the question side with a highlighter/marker to make it easier for you to sort.

or simply use the PDF to search for a command etc. Lets say you forgot how to do OSPF authentication, then go to the first page of OSPF, CTRL-F, "authentication" and hopefully you will find something within minutes that can help you.



Topic:	Page:
Switch/Bridge	2
DMVPN	6
IP routing	8
RIP	10
EIGRP	15
Redistribution	21
OSPF	22
BGP	33
Multicast	41
IPv6	50
Security	57
VPN/MPLS	66
System	74
Services	79
QoS	85
Frame-Relay	96
Study approach	102

<h2>802.1q Tunneling</h2>		<p>show interface fa0/x pruning</p>	<p>Show interface fa0/x pruning</p> <p>Port Vlans pruned for lack of request by neighbor Fa0/16 7-8,10,22,58,67,146</p> <p>Port Vlan traffic requested of neighbor Fa0/16 1,5,7-10,22,43,58,67,79,146</p> <p>Show interface fa0/x trunk -> offers easier output</p>	<p>show interface trunk</p>	<pre> Switch# int trunk Port Mode Encapsulation Status Native vlan Fa0/13 on isl trunking 1 Fa0/13 Vlans allowed on trunk: 1-4094 Fa0/13 Vlans allowed and active in management domain: 1,5,7-10,22,35,43,58,67,79,100,146,200,300,500,600,1000 Fa0/13 Vlans in spanning tree forwarding state and not pruned: 1,3,7-10,22,35,43,58,67,79,100,146,200,300,500,600,1000 Switch# </pre>																														
<p>What is a important pre-requisit for Dot1Q Tunnel setups ?</p>	<p>Set the MTU to 1504 and reload the switch.</p>	<h2>VTP Prune-Eligible List</h2>	<p>Vlans not specified in the list will NOT be pruned, vlans within the list could be pruned:</p> <p>interface FastEthernet0/x switchport trunk pruning vlan 2-6, 8-10</p> <p>2-6, 8-10 are prune-eligible!</p> <p>Vlan 7 will never be pruned!</p>	<p>Show spanning-tree uplinkfast</p>	<pre> VLAN0010 Spanning tree enabled protocol ieee Root ID Priority 10 Address 001b.d480.7c00 Cost 4000 Root 15 (FastEthernet0/13) Hello Time 3 sec Max Age 10 sec Forward Delay 10 sec Bridge ID 49162 (priority 49152 sys-id-ext 10) Address 001b.d4d4.e000 Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec Aging Time 300 Uplinkfast enabled ----- Interface Role Sts Cost Prio.Nbr Type ----- Fa0/6 Desg FWD 3019 128.5 F2p Fa0/13 Root FWD 4000 128.15 F2p Fa0/14 Altn RSP 4000 128.16 F2p Fa0/14 11+- RSP 4000 128.17 S2p </pre>																														
<h2>EtherChannel over 802.1q Tunneling</h2>	<pre> interface FastEthernet0/[X,Y,Z] switchport access vlan [10,20,30] switchport mode dot1q-tunnel l2protocol-tunnel cdp l2protocol-tunnel stp l2protocol-tunnel point-to-point [Pagp,Lacp] </pre>	<p>Spanning-tree global vs Interface commands</p>	<p>Global command over-rides the interface command:</p> <p>Spanning-tree vlan 1 - x port-prio 16</p> <p>Interface Fa0/x Spanning-tree port-priority 16</p>	<p>debug spanning-tree backbonefast</p>	<pre> Rack15W2#debug spanning-tree backbonefast Spanning Tree backbonefast: general debugging is on Rack15W2# STP FAST: received inferior BPDU on VLAN0001 FastEthernet0/19. STP FAST: sending RLQ request FDU on VLAN0001(1) Fa0/13 Vlan1 STP FAST: sending RLQ request FDU on VLAN0001(1) Fa0/14 Vlan1 STP FAST: sending RLQ request FDU on VLAN0001(1) Fa0/15 Vlan1 STP FAST: sending RLQ request FDU on VLAN0001(1) Fa0/20 Vlan1 STP FAST: sending RLQ request FDU on VLAN0001(1) Fa0/21 Vlan1 STP FAST: received inferior BPDU on VLAN0001 FastEthernet0/20. STP FAST: sending RLQ request FDU on VLAN0001(1) Fa0/13 Vlan1 STP FAST: sending RLQ request FDU on VLAN0001(1) Fa0/14 Vlan1 </pre>																														
<h2>Switchports types</h2>	<p>Interface Fa0/x</p> <p>Switchport mode dynamic auto</p> <p>Switchport mode dynamic desirable</p> <p>Switchport mode trunk</p> <p>Switchport mode access</p> <p>switchport mode dot1q-tunnel</p> <p>switchport mode private-vlan host</p> <p>switchport mode private-vlan promiscuous</p> <p>no switchport</p>	<p>Change Spanning-tree listening / learning timers</p>	<p>Spanning-tree vlan x-y forward-delay [seconds]</p> <p>Default is 15, command impacts LISTENING and LEARNING</p> <p>There is no separate command for the two states, only one for both.</p>		<p>Bridge 1 protocol ieee</p> <p>int fa0/x bridge-group X</p> <p>int ser0/x bridge-group X</p>																														
<p>switchport mode dynamic desirable</p> <p>switchport mode dynamic desirable</p> <p>switchport mode dynamic desirable</p> <p>switchport mode dynamic auto</p>	<table border="1"> <tr> <th>Port</th> <th>Mode</th> <th>Encapsulation</th> <th>Status</th> <th>Native vlan</th> </tr> <tr> <td>Fa0/13</td> <td>desirable</td> <td>n-isl</td> <td>trunking</td> <td>1</td> </tr> <tr> <td>Fa0/13</td> <td>desirable</td> <td>n-isl</td> <td>trunking</td> <td>1</td> </tr> </table> <table border="1"> <tr> <th>Port</th> <th>Mode</th> <th>Encapsulation</th> <th>Status</th> <th>Native vlan</th> </tr> <tr> <td>Fa0/13</td> <td>desirable</td> <td>n-isl</td> <td>trunking</td> <td>1</td> </tr> <tr> <td>Fa0/13</td> <td>auto</td> <td>n-isl</td> <td>trunking</td> <td>1</td> </tr> </table>	Port	Mode	Encapsulation	Status	Native vlan	Fa0/13	desirable	n-isl	trunking	1	Fa0/13	desirable	n-isl	trunking	1	Port	Mode	Encapsulation	Status	Native vlan	Fa0/13	desirable	n-isl	trunking	1	Fa0/13	auto	n-isl	trunking	1	<p>How to identify STP portfast ports in the debug output</p> <p>Debug spanning-tree events</p>	<p>Port fast enabled ports will have a log entry such as:</p> <p>JUMP TO FORWARDING FROM BLOCKING</p>	<p>What types of bridging over Frame relay are there?</p>	<p>Bridging over Frame-Relay -non multicast</p> <p>Bridging over Frame-Relay with multicast</p> <p>Bridging over Frame-Relay via Subinterfaces</p> <p>Remote Transparent bridging with Circuit-groups (MFR like)</p>
Port	Mode	Encapsulation	Status	Native vlan																															
Fa0/13	desirable	n-isl	trunking	1																															
Fa0/13	desirable	n-isl	trunking	1																															
Port	Mode	Encapsulation	Status	Native vlan																															
Fa0/13	desirable	n-isl	trunking	1																															
Fa0/13	auto	n-isl	trunking	1																															
<p>switchport mode dynamic desirable</p> <p>Switchport trunk encapsulation dot1q</p> <p>switchport mode dynamic desirable</p> <p>Switchport trunk encapsulation isl</p>	<table border="1"> <tr> <th>Port</th> <th>Mode</th> <th>Encapsulation</th> <th>Status</th> <th>Native vlan</th> </tr> <tr> <td>Fa0/13</td> <td>desirable</td> <td>802.1q</td> <td>trunking</td> <td>1</td> </tr> <tr> <td>Fa0/13</td> <td>desirable</td> <td>n-802.1q</td> <td>trunking</td> <td>1</td> </tr> </table> <table border="1"> <tr> <th>Port</th> <th>Mode</th> <th>Encapsulation</th> <th>Status</th> <th>Native vlan</th> </tr> <tr> <td>Fa0/13</td> <td>desirable</td> <td>isl</td> <td>trunking</td> <td>1</td> </tr> <tr> <td>Fa0/13</td> <td>desirable</td> <td>n-isl</td> <td>trunking</td> <td>1</td> </tr> </table>	Port	Mode	Encapsulation	Status	Native vlan	Fa0/13	desirable	802.1q	trunking	1	Fa0/13	desirable	n-802.1q	trunking	1	Port	Mode	Encapsulation	Status	Native vlan	Fa0/13	desirable	isl	trunking	1	Fa0/13	desirable	n-isl	trunking	1	<p>switchport trunk encapsulation dot1q</p> <p>switchport mode trunk</p> <p>switchport trunk encapsulation isl</p> <p>switchport mode trunk</p>	<p>switchport trunk encapsulation dot1q</p> <p>switchport mode trunk</p> <p>switchport trunk encapsulation isl</p> <p>switchport mode trunk</p>	<p>Handy set of Spanning tree debug commands:</p>	<p>Debug spanning-tree events</p> <p>Debug spanning-tree backbone fast</p> <p>Debug spanning-tree pvst+</p>
Port	Mode	Encapsulation	Status	Native vlan																															
Fa0/13	desirable	802.1q	trunking	1																															
Fa0/13	desirable	n-802.1q	trunking	1																															
Port	Mode	Encapsulation	Status	Native vlan																															
Fa0/13	desirable	isl	trunking	1																															
Fa0/13	desirable	n-isl	trunking	1																															

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

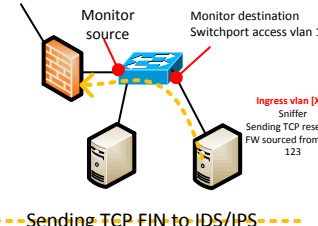
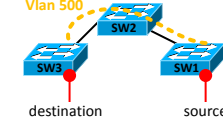
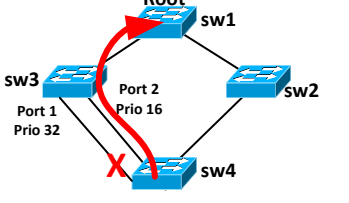
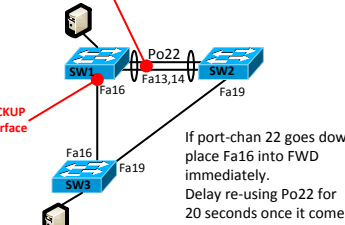
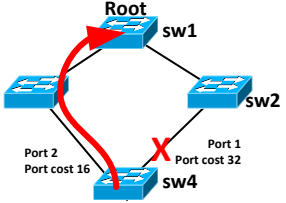
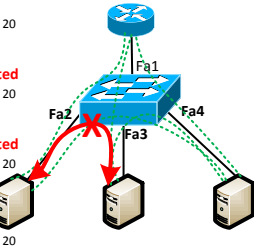
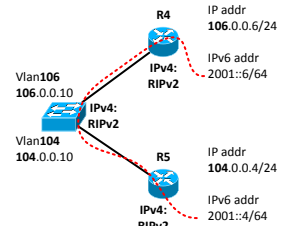
Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

Switch / Bridge

<p>UDLD port modes</p>	<p>Int fa0/x Udld port aggressive</p> <p>UDLD well-known MAC 0100.0ccc.cccc Device ID + Originator Port + Timeout Echo Value</p> <p>Normal Mode: Undetermined – continues to operate/ does not prevent loops</p> <p>Aggressive Mode: 8 UDLD frames / second No response for 1 second -> Port is error-disabled</p> <p>Reset all UDLD error disabled ports via: UDLD RESET</p>	<p>Storm control setup</p>	<pre>SW1: interface FastEthernet0/1 storm-control unicast level pps 100 SW2: interface FastEthernet0/6 storm-control broadcast level bps 10m SW4: interface FastEthernet0/4 storm-control broadcast level 10.0 Interface bandwidth 10 % SW1(config-if)#storm-control action [trap, shutdown]</pre>	<p>IP Phone Trust and CoS Extend</p>	<pre>interface FastEthernet0/2 mls qos trust cos mls qos trust device cisco-phone switchport priority extend cos 1 ! interface FastEthernet0/4 mls qos trust cos mls qos trust device cisco-phone switchport priority extend cos 1 ! interface FastEthernet0/6 mls qos trust cos mls qos trust device cisco-phone switchport priority extend cos 1</pre>												
<p>Spanning-tree MST config</p>	<p>spanning-tree mst configuration name MST1 revision 1 instance 1 vlan 1-100 instance 2 vlan 101-200 instance 3 vlan 201-4094</p> <p>spanning-tree mst 1 priority 0 spanning-tree mst 2 priority 4096 spanning-tree mst 2 priority 8192</p> <p>spanning-tree mode mst</p> <p>Show spanning-tree mst X detail</p>	<p>SPAN sessions</p>	<p>Monitor session [X] source [vlan x, int y] Monitor session [X] destination fa0/x ingress vlan [YY]</p> 	<p>Smartport Macros</p>	<p>Default interface fa0/x</p> <p>macro name MACRO-NAME switchport mode access switchport access vlan 146 spanning-tree bpdupfilter enable @</p> <p>Interface fa0/x macro apply MACRO-NAME</p> <p>show parser macro</p>												
<p>Spanning-tree MST port priority and cost</p>	<p>Interface fa0/x Spanning-tree mst X cost [COST]</p> <p>Interface fa0/x Spanning-tree mst X priority [16]</p>	<p>RSPAN sessions:</p>	<p>SW1: Vlan [500] remote-span</p> <p>monitor session [X] source interface [fa0/x] monitor session [X] destination vlan [500]</p> <p>SW2: Vlan [500] remote-span</p> <p>SW3: Vlan [500] remote-span</p> <p>monitor session [X] source vlan [500] monitor session [X] destination fa0/x ingress vlan [YY]</p> 	<p>Applying dynamic macros</p>	<p>Standard / existing macros can be verified by:</p> <p>Show parser macro</p> <p>And applied with dynamic parameters:</p> <p>Interface fa0/10 Macro apply cisco-desktop \$access_vlan 10</p>												
<p>Spanning-tree port priority and its influence?</p>	 <p>Spanning-tree port-priority influences only the direct attached SW4, but is configured on SW3!</p>	<p>Span session show output</p>	<pre>Rack1SW1#show monitor session 1 Session 1 ----- Type : Local Session Source VLANs : Both : 146 Destination Ports : Fa0/24 Encapsulation : Native Ingress : Disabled Rack1SW4#show monitor session 1 Session 1 ----- Type : Local Session Source Ports : Both : Fa0/4 Destination Ports : Fa0/24 Encapsulation : Native Ingress : Enabled, default VLAN = 146</pre>	<p>Flex Links</p> <p>Alternative to spanning-tree</p>	<p>Config applied to SW1:</p> <pre>Int Port-Channel 22 switchport backup interface Fa0/16 switchport backup interface Fa0/16 preemption mode forced switchport backup interface Fa0/16 preemption delay 20</pre> 												
<p>Spanning-tree port cost and its influence?</p>	 <p>Spanning-tree port cost affects SW4 and also other switches possible down stream switches to use port 2</p>	<p>RSPAN session show output:</p>	<pre>Rack1SW2#show monitor session 2 Session 2 ----- Type : Remote Source Session Source Ports : Both : Fa0/4 Dest RSPAN VLAN : 500 Rack1SW2#show monitor session 2 Session 2 ----- Type : Remote Source Session Source Ports : Both : Fa0/4 Dest RSPAN VLAN : 500</pre>	<p>Flex Links</p> <p>Show outputs</p>	<p>Primary link is UP/UP</p> <pre>debug backup all show interfaces po1 switchport backup</pre> <table border="1"> <thead> <tr> <th>Switch</th> <th>Backup Interface</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>Port-channel1</td> <td>FastEthernet0/16</td> <td>Active Up/Backup Standby</td> </tr> </tbody> </table> <p>Primary link is DOWN</p> <pre>show interfaces po1 switchport backup</pre> <table border="1"> <thead> <tr> <th>Switch</th> <th>Backup Interface</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>Port-channel1</td> <td>FastEthernet0/16</td> <td>Active Down/Backup Up</td> </tr> </tbody> </table>	Switch	Backup Interface	State	Port-channel1	FastEthernet0/16	Active Up/Backup Standby	Switch	Backup Interface	State	Port-channel1	FastEthernet0/16	Active Down/Backup Up
Switch	Backup Interface	State															
Port-channel1	FastEthernet0/16	Active Up/Backup Standby															
Switch	Backup Interface	State															
Port-channel1	FastEthernet0/16	Active Down/Backup Up															
<p>Protected Ports setup:</p>	<p>Int fa1 Switchport access vlan 20</p> <p>Int fa2 Switchport protected Switchport access vlan 20</p> <p>Int fa3 Switchport protected Switchport access vlan 20</p> <p>Int fa4 Switchport access vlan 20</p>  <p>Fa2 and Fa3 can't communicate even being in the same Vlan</p>	<p>What are the three different configuration methods for Voice ports?</p>	<pre>interface FastEthernet0/2 switchport access vlan 146 switchport voice vlan 600 spanning-tree portfast ! interface FastEthernet0/4 switchport trunk encapsulation dot1q switchport trunk native vlan 146 switchport trunk allowed vlan 146,600 switchport mode trunk switchport voice vlan 600 spanning-tree portfast trunk spanning-tree bpdupfilter enable ! interface FastEthernet0/6 switchport access vlan 146 switchport voice vlan dot1p</pre>	<p>Fallback Bridging</p>	<p>SW1: IP routing</p> <p>bridge 1 protocol vlan-bridge</p> <p>Int fa0/4 bridge-group [1]</p> <p>Int fa0/6 bridge-group [1]</p>  <p>show bridge [1] group</p>												

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

<h2>Fallback Bridging</h2>	<pre> R4: interface FastEthernet0/1 ip address 104.0.0.4 255.255.255.0 ipv6 address 2001::4/64 SW1# ip routing bridge 1 protocol vlan-bridge interface Vlan104 ip address 104.0.0.10 255.255.255.0 bridge-group 1 interface FastEthernet0/6 No switchport ip address 106.0.0.10 255.255.255.0 bridge-group 1 interface FastEthernet0/1 ip address 106.0.0.6 255.255.255.0 ipv6 address 2001::6/64 </pre>	<h2>PPP</h2> <h3>Bi-directional CHAP Authentication</h3>	<pre> R4# username USER-R4 pass cisco interface Serial 0/1 encapsulation ppp ppp authentication chap ppp chap hostname USER-R5 R5# username USER-R5 pass cisco interface Serial 0/1 encapsulation ppp ppp authentication chap ppp chap hostname USER-R4 </pre>	<h2>PPP AAA Authentication</h2> <h3>(Radius)</h3>	<pre> R4# (User admin password 0 cisco ! Don't get locked out) aaa new-model aaa authentication login CONSOLE none aaa authentication ppp PPP-AUTH-LIST group GRP-RADIUS local ! aaa group server radius GRP-RADIUS server-private 155.1.146.100 key CISCO ! interface Serial 0/1/0 ppp authentication ppp chap PPP-AUTH-LIST ! line console 0 login authentication CONSOLE (Auth via Radius, if not available fallback to local database) </pre>
<h2>Private VLANs</h2> <h3>Design</h3>		<h2>PPP</h2> <h3>Useful PPP debug commands</h3>	<pre> Debug ppp negotiations Debug ppp authentication Debug ppp packet Debug ppp error (debug aaa authentication) </pre>	<h2>PPP AAA Authentication</h2> <h3>(TACACS+)</h3>	<pre> R5# (User admin password 0 cisco ! Don't get locked out) aaa new-model aaa authentication login CONSOLE none aaa authentication ppp default group tacacs local ! tacacs-server host 155.1.146.200 key CISCO ! line console 0 login authentication CONSOLE (Auth via TACACS+, if not available fallback to local database) </pre>
<h2>Private VLANs</h2> <h3>config</h3>	<pre> vtp domain CCIE vtp mode transparent vlan 1000 private-vlan community vlan 3000 private-vlan isolated vlan 100 private-vlan primary private-vlan association 1000,2000,3000 (Associate Prim / Sec as last part in cfig phase 1) interface FastEthernet0/1 switchport private-vlan mapping 100 add 1000,2000,3000 switchport mode private-vlan promiscuous interface FastEthernet0/3 switchport private-vlan host-association 100 1000 switchport mode private-vlan host interface FastEthernet0/5 switchport private-vlan host-association 100 2000 switchport mode private-vlan host interface FastEthernet0/13 switchport trunk encapsulation dot1q switchport trunk </pre>	<h2>How do you filter out Vlan IDs out of SPAN sessions for Trunks ?</h2>	<pre> monitor session 2 source interface fa0/2 rx monitor session 2 filter vlan 5-10 monitor session 2 destination interface fa0/1 (SPAN session will only SPAN traffic within Vlan 5-10) </pre>	<h2>PPPoE</h2> <h3>Client / Server</h3>	<pre> R4 Client: interface Dialer1 ip address dhcp encapsulation ppp dialer pool 1 ppp chap hostname USER-1 ppp chap password CISCO ! interface FastEthernet0/1 no shutdown pppoe enable pppoe-client dial-pool-number 1 R5 Server: aaa authentication ppp PPPoE-LIST local username USER-1 password CISCO ip dhcp pool POOL-PPPoE network 155.1.35.0 255.255.255.0 interface FastEthernet 0/1 encapsulation dot1q 35 pppoe enable group PPPoE ! interface Virtual-Template 1 bba-group pppoe PPPoE virtual-template 1 sessions per-mac throttle 10 60 300 interface Virtual-Template 1 encapsulation ppp ip address 155.1.35.1 255.255.255.0 ppp authentication chap PPPoE-LIST </pre>
<h2>Private VLANs</h2> <h3>Show commands</h3>	<pre> show vlan private-vlan Primary Secondary Type Ports --- 100 1000 community Fa0/1, Fa0/3 100 2000 community Fa0/1, Fa0/5 100 3000 isolated Fa0/1 </pre> <p>check after configuring "phase 1"</p>	<h2>Whats the purpose of vlan dot1q tag native</h2> <h3>In dot1q tunneling ?</h3>	<pre> enables tagging of native VLAN frames on all IEEE 802.1Q trunk ports in dot1q tunnels. show vlan dot1q tag native -> Service-provider network mis-direction issue </pre>	<h2>PPPoE</h2> <h3>Debugs / Show CMDs</h3>	<pre> debug pppoe packets debug pppoe events debug ppp negotiation clear pppoe all show pppoe session </pre>
<h2>PPP</h2> <h3>Uni-directional PAP Authentication</h3>	<pre> R4# interface Serial 0/1 encapsulation ppp ppp pap sent-username USER-1 pass cisco R5# username USER-1 pass cisco interface Serial 0/1 encapsulation ppp ppp authentication pap Clock rate [64000] </pre>	<h2>How can one SPAN CDP, STP and other control protocols to the destination port ?</h2>	<pre> monitor session [1] destination interface Fa0/x encapsulation replicate C3750#sh monitor session 1 Session 1 ----- Type : Local Session Source Ports : Both : Gi1/0/9 Destination Ports : Gi1/0/24 Encapsulation : Replicate ← Ingress : Disabled </pre>	<h2>PPPoE Client status views:</h2>	<pre> Session on Client is UP R4#sh pppoe sess 1 Client session Uniq ID PPPoE RADIUS Port Source VA State N/A 19 0001:4098:F4C1 Fa0/1 011 V12 UP --- Transit to PPPoE SRV has just been disrupted R4#sh pppoe sess 1 Client session Uniq ID PPPoE RADIUS Port Source VA State N/A 20 0001:4098:F4C1 Fa0/1 011 N/A SHUTDOWN --- Transit disrupted, attempting connection to SRV R4#sh pppoe sess 1 Client session Uniq ID PPPoE RADIUS Port Source VA State N/A 20 0001:4098:F4C1 Fa0/1 011 N/A PENDING </pre>
<h2>PPP</h2> <h3>Uni-directional CHAP Authentication</h3>	<pre> R4# username USER-R4 pass cisco interface Serial 0/1 encapsulation ppp ppp authentication chap callin ppp chap hostname USER-R5 R5# username USER-R5 pass cisco interface Serial 0/1 encapsulation ppp ppp authentication chap </pre>	<h2>How to configure dot1Q tunnel on a routers subInterface</h2>	<pre> Interface Fa0/1 MTU 1504 Interface Fa0/1.41 encapsulation dot1Q 99 second-dot1q 200,300 </pre>	<h2>PPPoE Server's status views:</h2>	<pre> Connection from PPPoE SRV to Client established R4#sh pppoe sess 1 session in LOCALLY_TERMINATED (PTA) state 1 session total Uniq ID PPPoE RADIUS Port Source VA State 518 21 000c:ceb7:05c1 010/2/35 0001:4098:F4C1 VLAN 135 vcl1 UP Connection from PPPoE SRV to Client was lost due to problem on transit from client to server R4#sh pppoe sess Uniq ID PPPoE RADIUS Port Source VA State --- </pre>

Help me create more flashcards:

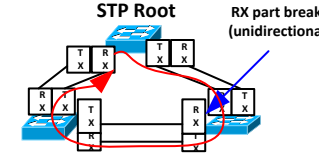
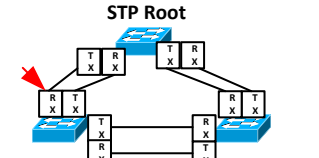
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

Switch / Bridge

<p>Advanced LACP / PAGP troubleshooting commands:</p>	<pre>Show int trunk Show etherchannel summary Rack1SW1#show pagp ? <1-48> Channel group number counters Traffic information internal Internal information neighbor Neighbor information Rack1SW1#show lacp ? <1-48> Channel group number counters Traffic information internal Internal information neighbor Neighbor information sys-id LACP System ID</pre>			<p>spanning-tree extend system-id explained</p>	<p>If "spanning-tree extend system-id" is NOT ENABLED: one MAC address per VLAN to make the bridge ID unique for each VLAN, using a lot of MAC addresses -> (chassis with only 64 MAC addresses!)</p> <p>extended system ID enabled: One MAC address used in all STP Vlans. In order to uniquely identify the root the VLAN-ID value is added to the STP priority value, with the same MAC address in all VLANs.</p>
<p>What does: define interface-range Do?</p>	<pre>conf t define interface-range VPORTS FastEthernet 0/7-8 Used for macro's</pre>	<p>VTP Version 3 And MST</p>	<ul style="list-style-type: none"> - MST domain = domain name and revision number (64 instances max) - MST domain configuration can be distributed via VTP Version 3 instance. (Instead of error prone manual config) - MST config changes only allowed on primary server. <pre>SWITCH# show spanning-tree mst configuration % Switch is not in mst mode Name [MST] Revision 1 Instances configured 3 Instance Vlans mapped ----- 0 2-400,406-4000,4006-4094 ----- config changes of MST only allowed on VTP 3 primary server.</pre>	<p>Spanning-tree loopguard</p> <p>Mainly (mainly used on fiber links, can be reproduced with bpdfilter on ethernet)</p>	 
<p>VTP Version 3 facts:</p>	<ul style="list-style-type: none"> - only the vtp primary server, is allowed to update other devices - Two instances in VTP version 3, VLAN and MST instance - Instances can be on the same or separate switches. - If there is more than one primary server a warning message will indicate conflicting devices. - Reserved VlanIDs 1000-1017, (show vlan intern usage) - VTPVer 2 compatible to Version 3, where as Ver 1 is not. - VTP 2 device will never update a Version 3 device. - After reload a primary server will take the secondary server role 	<p>VTP Version 3 With routed interfaces and Default VLAN-IDs:</p>	<pre>int fa1/24 no switchport ip address 1.2.3.4 255.255.255.0 SWITCH# show vlan internal usage VLAN Usage ----- <SNIP> 1018 FastEthernet1/24 VlanID 1018 for Fa1/24 ! default allocation policy the switch starts to allocate beginning at 1018! In order to use Vlan 1018, default interface fa1/24, shutdown the port, force a VTP revision number change, by changing vlan 1018's name to flood the change from an Internal Vlan to a regular vlan to other switches.</pre>	<p>Global: Spanning-tree loopguard default</p> <p>Per interface: Spanning-tree guard loop</p>	
<p>VTP Version 3 Configuration:</p>	<pre>spanning-tree extend system-id <- required vtp version 3 vtp domain DOMAIN-X vtp primary-server vlan [force] vtp primary-server mst [force] service password-encryption SW1# show vtp password VTP Encrypted Password: 02270A5F19030E32 Display of the password is encrypted, vlan.dat IS CLEAR TEXT</pre>	<p>VTP Version 3 modes</p>	<ul style="list-style-type: none"> - vtp mode server - vtp mode client - vtp mode transparent - vtp OFF <p>Disable vtp per port (VTP off)</p> <pre>int fa0/24 no vtp</pre>	<p>SDM Prefer on Catalyst</p>	<p>Access Used for QoS classification and security</p> <p>Routing Used for routing</p> <p>Vlan Disables routing and sets the switch to be a layer 2 switch</p> <p>Extended-match Reformats routing memory space to allow 144-bit layer 3 TCAM support needed for WCCP and/or multiple VRF instances</p> <pre>show sdm prefer conf t sdm prefer [Access Routing Vlan Extended-match]</pre>
<p>VTP Version 3 Difference between vtp password X hidden And vtp password / service password-encryption</p>	<pre>service password-encryption vtp password CISCO SWITCH# show vtp password VTP Encrypted Password: 02270A5F19030E32 Within vlan.dat password is still in clear text. vtp password CISCO hidden SW1# show vtp password VTP Password: CF94C2FF1CDCEB8DC795CEB21E305F10 vlan.dat is encrypted</pre>	<p>VTP Version 3 Cisco 3750 Stack Master is the vtp primary-server</p>	<p>The Stack Master announces its own MAC Address as vtp version 3 primary-server.</p> <p>-> THE CURRENT STACK MASTER FAILS!</p> <p>With the following command, the newly assigned Stack Master will send a VTP Version 3 take-over message after 5 minutes down time of the previous Stack Master. Announcing the new Stack Masters MAC address as primary-server</p> <pre>stack-mac persistent timer <S></pre>	<p>Catalyst VMPS config:</p>	<pre>SW1#show vmps VQP Client Status: ----- Vmps reconfirm 30 vmps retry 5 vmps server 1.1.1.1 primary vmps server 2.2.2.2 VMPS VQP Version: 1 Reconfirm Interval: 60 min Server Retry Count: 3 VMPS domain server: Reconfirmation status: ----- VMPS Action: No Dynamic Port int fa0/x switchport mode access switchport access vlan dynamic</pre>
<p>VTP Version 3 Show commands:</p>	<pre>show vtp status enhanced show vtp devices conflicts show vtp devices feature show vtp interface show vtp counters (can discover configs from different primary servers) show vlan internal usage (Vlan ID 1000-1018 issue)</pre>	<p>VTP Version 3 show vtp devices Output:</p>	<pre>SWITCH# show vtp devices Gathering information from the domain, please wait. VTP Database Conf switch ID Primary Server Revision System Name ----- VLAN No 000c.0012.3456=000c.0012.3456 1001 SWITCH MST No 000c.0012.3456=000c.0012.3456 42 SWITCH (Neighbor with two instances shown, VLAN and MST instance)</pre>	<p>Spanning-tree selection rules</p>	<p>STP rules:</p> <ol style="list-style-type: none"> 1. Lower root BID 2. Lower path cost to the root bridge 3. lower sending BID 4. lower Sending Port-ID (Priority).Port-iD

Home Lab tip:

If you have a lab on which the Lines keep miss-behaving, Clear all lines instead of individually:

```
alias exec line-clear event manager run CLEAR-LINES


event manager applet CLEAR-LINES
event none sync yes
action 1.0 cli command "clear line 65"
action 1.1 cli command "clear line 66"
action 1.2 cli command "clear line 67"
action 1.3 cli command "clear line 68"
action 1.4 cli command "clear line 69"
action 1.5 cli command "clear line 70"
action 1.6 cli command "clear line 71"
action 1.7 cli command "clear line 72"
action 1.8 cli command "clear line 73"
action 1.9 cli command "clear line 74"
action 2.0 cli command "clear line 75"
action 2.1 cli command "clear line 76"
action 2.2 cli command "clear line 77"
action 2.3 cli command "clear line 78"
action 2.4 cli command "clear line 79"
action 2.5 syslog msg "All lines cleared"
```


R3#line-clear
%HA_EM-6-LOG: CLEAR-LINES: All lines cleared

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!





Thanks for appreciating my efforts

Colin

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

DMVPN Phase 1 Static config

```
HUB
interface Tunnel1
ip address 10.1.1.1 255.255.255.0
no ip redirects
ip nhrp map 10.1.1.2 192.1.2.2
ip nhrp map 10.1.1.3 192.1.3.3
ip nhrp map 10.1.1.4 192.1.4.4
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel mode gre multipoint

R1#show ip nhrp
10.1.1.2/32 via 10.1.1.2
Tunnel1 created 00:11:18, never expires
Type: static, Flags:
NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
Tunnel1 created 00:11:18, never expires
Type: static, Flags:
NBMA address: 192.1.3.3
10.1.1.4/32 via 10.1.1.4
Tunnel1 created 00:11:18, never expires
Type: static, Flags: used
NBMA address: 192.1.4.4

Spoke
interface Tunnel1
ip address 10.1.1.2 255.255.255.0
ip nhrp map 10.1.1.1 192.1.1.1
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel destination 192.1.1.1

R2#show ip nhrp
10.1.1.1/32 via 10.1.1.1
Tunnel1 created 00:11:49, never expires
Type: static, Flags:
NBMA address: 192.1.1.1
```

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

DMVPN Phase 1 Dynamic mapping config

```
HUB
interface Tunnel1
ip address 10.1.1.1 255.255.255.0
no ip redirects
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel mode gre multipoint

R1#show ip nhrp
10.1.1.2/32 via 10.1.1.2
Tunnel1 created 00:07:29, expires 00:07:29
Type: dynamic, Flags: unique register
NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
Tunnel1 created 00:07:24, expires 00:07:24
Type: dynamic, Flags: unique register
NBMA address: 192.1.3.3

Spoke
interface Tunnel1
ip address 10.1.1.2 255.255.255.0
ip nhrp map 10.1.1.1 192.1.1.1
ip nhrp network-id 111
ip nhrp nhs 10.1.1.1
tunnel source Ethernet0/0
tunnel destination 192.1.1.1
Type: static, Flags:
NBMA address: 192.1.1.1
```

What is the difference between the two DMVPN flavors?

GRE
int tun 1
tunnel source x.x.x.x
tunnel destination y.y.y.y

mGRE
int tun 1
tunnel source x.x.x.x
tunnel mode gre multipoint

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

Trace from Spoke 2 to Spoke 3

DMVPN Phase 1 Static config

```
R2#traceroute 10.1.1.3 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.1 13 msec 6 msec 5 msec
 2 10.1.1.3 3 msec * 2 msec

2nd traceroute:

R2#traceroute 10.1.1.3 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.1 13 msec 6 msec 5 msec
 2 10.1.1.3 3 msec * 2 msec
```

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

Trace from Spoke 2 to Spoke 3

DMVPN Phase 1 Dynamic config

```
R2#traceroute 10.1.1.3 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.1 12 msec 1 msec 6 msec
 2 10.1.1.3 6 msec * 2 msec

2nd traceroute:

R2#traceroute 10.1.1.3 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.1 12 msec 1 msec 6 msec
 2 10.1.1.3 6 msec * 2 msec
```

Hub:
interface Tunnel1
ip address 10.1.1.1 255.255.255.0
no ip redirects
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel mode gre multipoint

Spoke:
interface Tunnel1
ip address 10.1.1.4 255.255.255.0
no ip redirects
ip nhrp map 10.1.1.1 192.1.1.1
ip nhrp network-id 111
ip nhrp nhs 10.1.1.1
tunnel source Ethernet0/0
tunnel mode gre multipoint

On Hub: debug nhrp cache:

debug nhrp cache:
NHRP: Tunnel1: Cache add for target 10.1.1.4/32 next-hop 10.1.1.4 192.1.4.4
NHRP: Inserted subblock node for cache: Target inserted subblock node for cache: Target 10.1.1.4/32 next-hop 10.1.1.4
NHRP: Converted internal dynamic cache entry for 10.1.1.4/32 interface Tunnel1 to external
NHRP: Updating our cache with NBMA: 192.1.1.1, NBMA_ALT: 192.1.1.1
NHRP: Setting 'used' flag on cache entry with nhop: 10.1.1.4
NHRP: NHRP successfully mapped '10.1.1.4' to NBMA 192.1.4.4
NHRP: Tunnel1: Cache update for target 10.1.1.4/32 next-hop 10.1.1.4 192.1.4.4
NHRP: Updating our cache with NBMA: 192.1.1.1, NBMA_ALT: 192.1.1.1
NHRP: Setting 'used' flag on cache entry with nhop: 10.1.1.4
NHRP: NHRP successfully mapped '10.1.1.4' to NBMA 192.1.4.4

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

DMVPN Phase 2 Static config

```
HUB
interface Tunnel1
ip address 10.1.1.1 255.255.255.0
no ip redirects
ip nhrp map 10.1.1.2 192.1.2.2
ip nhrp map 10.1.1.3 192.1.3.3
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel mode gre multipoint

R1#show ip nhrp
10.1.1.2/32 via 10.1.1.2
Tunnel1 created 00:03:41, never expires
Type: static, Flags: used
NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
Tunnel1 created 00:03:41, never expires
Type: static, Flags: used
NBMA address: 192.1.3.3

Spoke
interface Tunnel1
ip address 10.1.1.2 255.255.255.0
no ip redirects
ip nhrp map 10.1.1.1 192.1.1.1
ip nhrp map 10.1.1.3 192.1.3.3
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel mode gre multipoint

R2#show ip nhrp
10.1.1.1/32 via 10.1.1.1
Tunnel1 created 00:02:41, never expires
Type: static, Flags: used
NBMA address: 192.1.1.1
10.1.1.3/32 via 10.1.1.3
Tunnel1 created 00:02:41, never expires
Type: static, Flags: used
NBMA address: 192.1.3.3
```

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

DMVPN Phase 2 Dynamic config

```
HUB
interface Tunnel1
ip address 10.1.1.1 255.255.255.0
no ip redirects
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel mode gre multipoint

R1#show ip nhrp
10.1.1.2/32 via 10.1.1.2
Tunnel1 created 00:15:59, expires 00:15:59
Type: dynamic, Flags: unique register
NBMA address: 192.1.2.2
10.1.1.3/32 via 10.1.1.3
Tunnel1 created 00:15:51, expires 00:15:51
Type: dynamic, Flags: unique register
NBMA address: 192.1.3.3

Spoke
interface Tunnel1
ip address 10.1.1.2 255.255.255.0
no ip redirects
ip nhrp map 10.1.1.1 192.1.1.1
ip nhrp network-id 111
ip nhrp nhs 10.1.1.1
tunnel source Ethernet0/0
tunnel mode gre multipoint

R2#show ip nhrp
10.1.1.1/32 via 10.1.1.1
Tunnel1 created 00:04:29, never expires
Type: static, Flags: used
NBMA address: 192.1.1.1
10.1.1.4/32 via 10.1.1.4
Tunnel1 created 00:02:31, expires 00:02:31
Type: dynamic, Flags: router
NBMA address: 192.1.4.4
```

Hub:
interface Tunnel1
ip address 10.1.1.1 255.255.255.0
no ip redirects
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel mode gre multipoint

Spoke:
interface Tunnel1
ip address 10.1.1.4 255.255.255.0
no ip redirects
ip nhrp map 10.1.1.1 192.1.1.1
ip nhrp network-id 111
ip nhrp nhs 10.1.1.1
tunnel source Ethernet0/0
tunnel mode gre multipoint

On Hub: debug nhrp packet:

NHRP: Receive Registration Request via Tunnel1 vrf 0, packet size: 92
src NBMA: 192.1.4.4
src protocol: 10.1.1.4, dst protocol: 10.1.1.1

NHRP: Send Registration Reply via Tunnel1 vrf 0, packet size: 112
src: 10.1.1.1, dst: 10.1.1.4
(M) flags: "unique nat", reqid: 65544
src NBMA: 192.1.4.4
src protocol: 10.1.1.4, dst protocol: 10.1.1.1

NHRP: Receive Registration Request via Tunnel1 vrf 0, packet size: 92
src NBMA: 192.1.4.4
src protocol: 10.1.1.4, dst protocol: 10.1.1.1

NHRP: Send Registration Reply via Tunnel1 vrf 0, packet size: 112
src: 10.1.1.1, dst: 10.1.1.4

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

Trace from Spoke 4 to Spoke 2

DMVPN Phase 2 static config

```
R4#traceroute 10.1.1.2 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.2
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.2 7 msec * 1 msec

2nd traceroute:

R4#traceroute 10.1.1.2 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.2
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.2 7 msec * 1 msec
```

Direct access due to static mapping and Tunnel mode gre multipoint

NBMA Range 192.1.1.x.x
Tunnel Range 10.1.1.x

Trace from Spoke 4 to Spoke 2

DMVPN Phase 2 dynamic config

```
R2#traceroute 10.1.1.3 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.1 11 msec 1 msec 6 msec
 2 10.1.1.3 1 msec * 1 msec

2nd traceroute:

R2#traceroute 10.1.1.3 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.3
VRF info: (vrf in name/id, vrf out name/id)
 1 10.1.1.3 1 msec * 2 msec
```

What are main configuration differences between DMVPN Phase 1 Static and dynamic Config?

Phase 1 static
Hub:
ip nhrp map SPOKE-TUNNEL-IP SPOKE-NBMA-IP
tunnel mode gre multipoint

Spoke:
ip nhrp map HUB-TUNNEL-IP HUB-NBMA-IP
tunnel destination HUB-NBMA-IP

Phase 1 dynamic
Hub:
ip nhrp map SPOKE-TUN-IP SPOKE-NBMA-IP
tunnel mode gre multipoint

Spoke:
ip nhrp map HUB-TUNNEL-IP HUB-NBMA-IP
tunnel destination HUB-NBMA-IP
ip nhrp nhs HUB-NBMA-IP

DMVPN Phase 1 setup Using EIGRP

```
HUB
interface Tunnel123
ip address 123.1.1.1 255.255.255.0
no ip redirects
ip nhrp map 123.1.1.2 200.1.2.2
ip nhrp map 123.1.1.3 200.1.3.3
ip nhrp map multicast 200.1.2.2
ip nhrp map multicast 200.1.3.3
ip nhrp network-id 111
no ip split-horizon eigrp 10
tunnel source Ethernet0/0
tunnel mode gre multipoint

Spoke
router eigrp 10
network 123.0.0.0
network 2.2.2.2

HUB:
router eigrp 10
network 123.0.0.0
network 1.1.1.1
```

You have issued clear ip nhrp on the HUB and discover that you no longer have reachability from and to your spokes!

How do you solve this in the lab?

Force the spokes to re-register

On Spokes:
int tun x
ip nhrp registration timeout <1 sec>

What are main configuration differences between DMVPN Phase 2 Static and dynamic Config?

Phase 2 static
Hub:
ip nhrp map SPOKE-TUNNEL-IP SPOKE-NBMA-IP
tunnel mode gre multipoint

Spoke:
ip nhrp map HUB-TUNNEL-IP HUB-NBMA-IP
ip nhrp map SPOKE-x-TUNNEL-IP SPOKE-x-NBMA-IP
tunnel mode gre multipoint

Phase 2 dynamic
Hub:
ip nhrp map SPOKE-TUNNEL-IP SPOKE-NBMA-IP
tunnel mode gre multipoint

Spoke:
ip nhrp map HUB-TUNNEL-IP HUB-NBMA-IP
ip nhrp nhs HUB-TUNNEL-IP
tunnel mode gre multipoint

NBMA Range 200.1.x.x
Tunnel Range 123.1.1.x

DMVPN Phase 1 setup Using EIGRP

Do not use "no ip split-horizon eigrp X"

```
HUB
interface Tunnel123
ip address 123.1.1.1 255.255.255.0
no ip redirects
ip nhrp map 123.1.1.2 200.1.2.2
ip nhrp map 123.1.1.3 200.1.3.3
ip nhrp map multicast 200.1.2.2
ip nhrp map multicast 200.1.3.3
ip nhrp network-id 111
ip summary-address eigrp 10 0.0.0.0 0.0.0.0
tunnel source Ethernet0/0
tunnel mode gre multipoint

Spoke
router eigrp 10
network 123.0.0.0
network 2.2.2.2

HUB:
router eigrp 10
network 123.0.0.0
network 1.1.1.1
```

Sending a default route via summarization

NBMA Range 200.1.x.x
Tunnel Range 123.1.1.x

DMVPN Phase 1 setup Using RIP

```
HUB
interface Tunnel123
ip address 123.1.1.1 255.255.255.0
no ip redirects
ip nhrp map 123.1.1.2 200.1.2.2
ip nhrp map 123.1.1.3 200.1.3.3
ip nhrp map multicast 200.1.2.2
ip nhrp map multicast 200.1.3.3
ip nhrp network-id 111
ip rip advertise 2
no ip split-horizon
tunnel source Ethernet0/0
tunnel mode gre multipoint

Spoke
router rip
version 2
network 2.0.0.0
network 123.0.0.0
no auto-summary

HUB:
router rip
version 2
network 1.0.0.0
network 123.0.0.0
no auto-summary

Spoke
interface Tunnel123
ip address 123.1.1.2 255.255.255.0
ip nhrp map 123.1.1.1 200.1.1.1
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel destination 200.1.1.1
```

DMVPN Phase 1 setup Using RIP

do not use "no ip split-horizon"

```
HUB
interface Tunnel123
ip address 123.1.1.1 255.255.255.0
ip summary-address rip 0.0.0.0 0.0.0.0
ip nhrp map 123.1.1.2 200.1.2.2
ip nhrp map 123.1.1.3 200.1.3.3
ip nhrp map multicast 200.1.2.2
ip nhrp map multicast 200.1.3.3
ip nhrp network-id 111
ip rip advertise 2
no ip split-horizon
tunnel source Ethernet0/0
tunnel mode gre multipoint

Spoke
router rip
version 2
network 2.0.0.0
network 123.0.0.0
no auto-summary

HUB:
router rip
version 2
network 1.0.0.0
network 123.0.0.0
no auto-summary

Spoke
interface Tunnel123
ip address 123.1.1.2 255.255.255.0
ip nhrp map 123.1.1.1 200.1.1.1
ip nhrp network-id 111
tunnel source Ethernet0/0
tunnel destination 200.1.1.1
```

Sending a default route via summarization (NBMA must be specifically routed /32 or higher AD)

Help me create more flashcards:


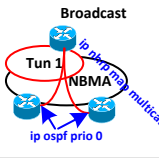




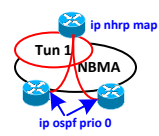
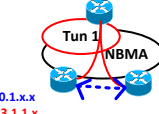
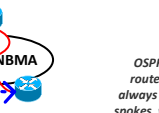
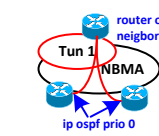
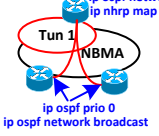
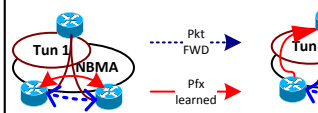
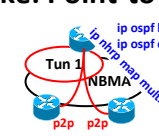
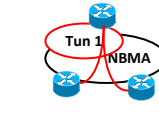
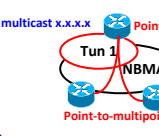
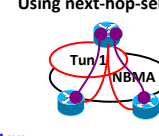
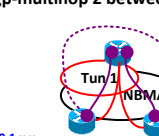

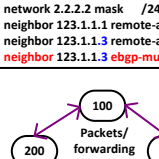
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin


<p>What three solutions are there in regards to DMVPN Phase 1 and OSPF?</p>	<p>Phase 1 – OSPF Point-to-point</p>  <p>Error messages! EXCHANGE TO DOWN ...</p> <p>Phase 1 – OSPF Broadcast</p>  <p>Phase 1 – OSPF Non Broadcast</p>  <p>Phase 1 – OSPF Point-to-Multipoint</p> 	<p>DMVPN Phase 2 and RIPv2 (direct Spoke to Spoke)</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 /24 no ip redirects ip nhrp map multicast dynamic ip nhrp network-id 111 no ip split-horizon tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 no ip redirects ip nhrp map multicast 200.1.1.1 ip nhrp network-id 111 ip nhrp nhs 123.1.1.1 tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>router rip version 2 network 1.0.0.0 network 123.0.0.0 no auto-summary</p> <p>R2#show ip route i 3.3.3.3 R 3.3.3.3 [120/2] via 123.1.1.3, 00:00:16, Tunnel123</p> <p>R2#tracroute 3.3.3.3 source 2.2.2.2 numeric VRF info: (vrf in name/id, vrf out name/id) 1 123.1.1.3 2 msec * 6 msec</p>	<p>DMVPN Phase 3 and EIGRP (direct Spoke to Spoke)</p> <p>Not using: no ip next-hop-self eigrp 10</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 /24 no ip redirects ip nhrp map multicast dynamic ip nhrp network-id 111 ip nhrp network-id 222 no ip split-horizon eigrp 10 tunnel source FastEthernet0/0 tunnel mode gre multipoint</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 no ip redirects ip nhrp map multicast 200.1.1.1 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 222 ip nhrp nhs 123.1.1.1 ip nhrp shortcut tunnel source FastEthernet0/0 tunnel mode gre multipoint</p> <p>router eigrp 10 network 1.1.1.1 0.0.0.0 network 123.1.1.1 0.0.0.0</p> <p>router eigrp 10 network 2.2.2.2 0.0.0.0 network 123.1.1.2 0.0.0.0</p> <p>Notice the route pointing to .1, traffic flowing to .3 directly as of the 2nd packet!</p> <p>R2#show ip route i 3.3.3.3 D 3.3.3.3 [90/28288000] via 123.1.1.1, 00:00:16, Tunnel123</p> <p>R2#tracroute 3.3.3.3 sou 2.2.2.2 numeric (2nd traceroute) 1 123.1.1.3 4 msec * 0 msec</p>
<p>DMVPN Phase 1 / OSPF Broadcast</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 255.255.255.0 ip nhrp map 123.1.1.2 200.1.2.2 ip nhrp map 123.1.1.3 200.1.3.3 ip nhrp map multicast 200.1.2.2 ip nhrp map multicast 200.1.3.3 ip nhrp network-id 111 ip ospf network broadcast tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>Spoke: router ospf 1 router-id 0.0.0.2 network 2.2.2.2 0.0.0.0 area 0 network 123.1.1.2 0.0.0.0 area 0</p> <p>HUB router ospf 1 router-id 0.0.0.1 network 1.1.1.1 0.0.0.0 area 0 network 123.1.1.1 0.0.0.0 area 0</p> <p>Spoke: interface Tunnel123 ip address 123.1.1.2 255.255.255.0 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 111 ip ospf network broadcast ip ospf priority 0 tunnel source Ethernet0/0 tunnel destination 200.1.1.1</p>	<p>DMVPN Phase 2 and EIGRP (direct Spoke to Spoke)</p> <p>-> change next-hop behaviour</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 /24 no ip redirects no ip next-hop-self eigrp 10 no ip split-horizon eigrp 10 ip nhrp map multicast dynamic ip nhrp network-id 111 ip nhrp network-id 111 tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 no ip redirects ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp map multicast 200.1.1.1 ip nhrp network-id 111 ip nhrp nhs 123.1.1.1 ip ospf network broadcast tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>router eigrp 10 network 1.0.0.0 network 123.0.0.0</p> <p>router eigrp 10 network 2.0.0.0 network 123.0.0.0</p> <p>R2#show ip route i 3.3.3.3 D 3.3.3.3 [90/28288000] via 123.1.1.3, 00:18:48, Tunnel123</p> <p>R2#tracroute 3.3.3.3 source 2.2.2.2 numeric VRF info: (vrf in name/id, vrf out name/id) 1 123.1.1.3 1 msec * 5 msec</p>	<p>DMVPN Phase 3 and OSPF (direct Spoke to Spoke)</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p> <p>OSPF Point2Multipoint on all routers would cause spokes to always travel via the HUB to other spokes, with Phase 3, the routes are advertised by the HUB, but travel/redirect directly from spoke to spoke!</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 255.255.255.0 no ip redirects ip nhrp map multicast dynamic ip nhrp network-id 111 ip nhrp network-id 222 ip nhrp network point-to-multipoint tunnel source FastEthernet0/0 tunnel mode gre multipoint</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 255.255.255.0 no ip redirects ip nhrp map multicast 200.1.1.1 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 222 ip nhrp network-id 222 ip nhrp nhs 123.1.1.1 ip nhrp shortcut ip ospf network point-to-multipoint tunnel source FastEthernet0/0 tunnel mode gre multipoint</p> <p>router ospf 1 network 1.1.1.1 0.0.0.0 area 0 network 123.1.1.1 0.0.0.0 area 0</p> <p>router ospf 1 network 2.2.2.2 0.0.0.0 area 0 network 123.1.1.2 0.0.0.0 area 0</p> <p>R3#show ip route i 2.2. O 2.2.2.2 [110/2001] via 123.1.1.1, 00:04:48, Tunnel123</p> <p>R3#tracroute 2.2.2.2 source 3.3.3.3 numeric 1 123.1.1.3 0 msec * 0 msec</p> <p>As of 2nd packet!</p>
<p>DMVPN Phase 1 / OSPF non-broadcast</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 255.255.255.0 ip nhrp map 123.1.1.2 200.1.2.2 ip nhrp map 123.1.1.3 200.1.3.3 ip nhrp map multicast 200.1.2.2 ip nhrp map multicast 200.1.3.3 ip nhrp network-id 111 ip ospf network non-broadcast tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>HUB router ospf 1 router-id 0.0.0.1 network 1.1.1.1 0.0.0.0 area 0 network 123.1.1.1 0.0.0.0 area 0 neighbor 123.1.1.3 neighbor 123.1.1.2</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 255.255.255.0 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 111 ip ospf network non-broadcast ip ospf priority 0 tunnel source Ethernet0/0 tunnel destination 200.1.1.1</p> <p>Spoke router ospf 1 router-id 0.0.0.2 network 2.2.2.2 0.0.0.0 area 0 network 123.1.1.2 0.0.0.0 area 0</p>	<p>DMVPN Phase 2 and OSPF (direct Spoke to Spoke)</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 /24 ip nhrp map multicast dynamic ip nhrp network-id 111 ip ospf network broadcast tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 no ip redirects ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 111 ip nhrp nhs 123.1.1.1 ip ospf network broadcast tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>router ospf 1 network 1.1.1.1 0.0.0.0 area 0 network 123.1.1.1 0.0.0.0 area 0</p> <p>router ospf 1 network 2.2.2.2 0.0.0.0 area 0 network 123.1.1.2 0.0.0.0 area 0</p> <p>R2#show ip route ospf i 3.3.3 O 3.3.3.0 [110/1001] via 123.1.1.3, 00:05:10, Tunnel123</p> <p>R2#tracroute 3.3.3.3 source 2.2.2.2 numeric 1 123.1.1.3 2 msec * 2 msec</p>	<p>What is the difference between DMVPN Phase 2 and Phase 3?</p>  <p>Phase 2: The routing protocol on a Spoke learns prefixes of other spokes directly of the other spokes through the tunnel IP address. Traffic then forwarded directly to the spokes tunnel IP. (Resolution done via IGP, not NHRP)</p> <p>Phase 3: A Spokes routing protocol points towards the HUB for a prefix from another Spoke, NHRP kicks in and redirects traffic directly between spokes. On Hub: ip nhrp redirect On Spokes: ip nhrp shortcut</p>	
<p>DMVPN Phase 1 / OSPF HUB: Point-to-Multipoint Spoke: Point-to-Point</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 255.255.255.0 ip nhrp map 123.1.1.2 200.1.2.2 ip nhrp map 123.1.1.3 200.1.3.3 ip nhrp map multicast 200.1.2.2 ip nhrp map multicast 200.1.3.3 ip nhrp network-id 111 ip ospf network point-to-multipoint ip ospf hello-interval 10 ip ospf dead-interval 40 tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>HUB router ospf 1 router-id 0.0.0.1 network 1.1.1.1 0.0.0.0 area 0 network 123.1.1.1 0.0.0.0 area 0</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 111 ip ospf network point-to-multipoint tunnel destination 200.1.1.1</p> <p>Spoke router ospf 1 router-id 0.0.0.2 network 2.2.2.2 0.0.0.0 area 0 network 123.1.1.2 0.0.0.0 area 0</p>	<p>DMVPN Phase 2 and BGP (direct Spoke to Spoke)</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 /24 ip nhrp map multicast dynamic ip nhrp network-id 111 tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp map multicast 200.1.1.1 ip nhrp network-id 111 ip nhrp nhs 123.1.1.1 tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>router bgp 100 bgp log-neighbor-changes network 1.1.1.1 mask 255.255.255.0 neighbor 123.1.1.2 remote-as 200 neighbor 123.1.1.3 remote-as 300</p> <p>router bgp 200 bgp log-neighbor-changes network 2.2.2.2 mask 255.255.255.0 neighbor 123.1.1.1 remote-as 100 neighbor 123.1.1.3 remote-as 300 neighbor 123.1.1.3 ebgp-multihop 2</p> <p>R2#show ip route i 3.3.3 B 3.3.3.0 [20/0] via 123.1.1.3, 00:13:08</p> <p>R2#tracroute 3.3.3.3 source 2.2.2.2 numeric 1 123.1.1.3 6 msec * 1 msec</p>		
<p>DMVPN Phase 1 / OSPF HUB: Point-to-Multipoint Spoke: Point-to-Multipoint</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 255.255.255.0 ip nhrp map 123.1.1.2 200.1.2.2 ip nhrp map 123.1.1.3 200.1.3.3 ip nhrp map multicast 200.1.2.2 ip nhrp map multicast 200.1.3.3 ip nhrp network-id 111 ip ospf network point-to-multipoint tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>HUB router ospf 1 router-id 0.0.0.1 network 1.1.1.1 0.0.0.0 area 0 network 123.1.1.1 0.0.0.0 area 0</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 111 no ip route-cache ip ospf network point-to-multipoint tunnel source Ethernet0/0 tunnel destination 200.1.1.1</p> <p>Spoke router ospf 1 router-id 0.0.0.2 network 2.2.2.2 0.0.0.0 area 0 network 123.1.1.2 0.0.0.0 area 0</p>				
<p>DMVPN Phase 1 BGP Using next-hop-self</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 /24 ip nhrp map 123.1.1.2 200.1.2.2 ip nhrp map 123.1.1.3 200.1.3.3 ip nhrp network-id 111 tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>HUB router bgp 100 bgp log-neighbor-changes network 1.1.1.1 mask /32 neighbor 123.1.1.2 remote-as 200 neighbor 123.1.1.2 next-hop-self neighbor 123.1.1.3 remote-as 300 neighbor 123.1.1.3 next-hop-self</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp network-id 111 tunnel source Ethernet0/0 tunnel destination 200.1.1.1</p> <p>Spoke router bgp 200 bgp log-neighbor-changes network 2.2.2.2 mask /32 neighbor 123.1.1.1 remote-as 100 neighbor 123.1.1.1 next-hop-self</p>	<p>DMVPN Phase 1 BGP ebgp-multihop 2 between spokes</p>  <p>NBMA Range 200.1.x.x Tunnel Range 123.1.1.x</p>	<p>HUB interface Tunnel123 ip address 123.1.1.1 /24 ip nhrp map 123.1.1.2 200.1.2.2 ip nhrp map 123.1.1.3 200.1.3.3 ip nhrp network-id 111 tunnel source Ethernet0/0 tunnel mode gre multipoint</p> <p>Spoke interface Tunnel123 ip address 123.1.1.2 /24 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp map multicast 200.1.1.1 ip nhrp network-id 111 tunnel source Ethernet0/0 tunnel destination 200.1.1.1</p> <p>router bgp 100 bgp log-neighbor-changes network 1.1.1.1 mask /32 neighbor 123.1.1.2 remote-as 200 neighbor 123.1.1.3 remote-as 300</p> <p>router bgp 200 bgp log-neighbor-changes network 2.2.2.2 mask /24 neighbor 123.1.1.1 remote-as 100 neighbor 123.1.1.3 remote-as 300 neighbor 123.1.1.3 ebgp-multihop 2</p> <p>BGP view</p> 	<p>What are the main differences between DMVPN Phase 1</p> <p>Static mapping</p> <p>Dynamic mapping</p> 	<p>HUB int tun 123 ip nhrp map 123.1.1.x 200.1.x.x ip nhrp map multicast 200.1.x.x</p> <p>Spoke int tun 123 ip nhrp map 123.1.1.x 200.1.1.1 tunnel destination 200.1.1.1</p> <p>HUB int tun 123 ip nhrp map multicast dynamic</p> <p>Spoke int tun 123 ip nhrp map 123.1.1.1 200.1.1.1 ip nhrp nhs 123.1.1.1 tunnel destination 200.1.1.1</p> <p>Static mapping</p> <p>Dynamic mapping</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

IP routing


<p>Enabling / disabling Proxy-Arp:</p> <p>Debugging proxy-arp:</p>	<p>Disable:</p> <pre>interface GigabitEthernet0/0.146 encapsulation dot1Q 146 ip address 155.1.146.6 255.255.255.0 no ip proxy-arp</pre> <p>Enable:</p> <pre>interface GigabitEthernet0/0.146 encapsulation dot1Q 146 ip address 155.1.146.6 255.255.255.0 no ip proxy-arp</pre> <p>Debug arp Debug ip packet</p> <p>IP ARP: rcvd rep src 150.1.6.6 0026.0b57.b960, dst 155.1.146.1 FastEthernet0/0</p>	<p>Policy routing configuration:</p>	 <pre>R1# interface FastEthernet0/0 ip policy route-map POLICY_ROUTING ip access-list extended FROM_R4 permit ip host 155.1.146.4 any ip access-list extended FROM_R6 permit ip host 155.1.146.6 any route-map POLICY_ROUTING permit 10 match ip address FROM_R4 set ip next-hop 155.1.13.3 route-map POLICY_ROUTING permit 20 match ip address FROM_R6 set ip next-hop 155.1.0.5</pre>	<h3>Local Policy Routing</h3> <p>Configuration:</p>	<pre>ip local policy route-map RMP-POL-LOCAL route-map RMP-POL-LOCAL permit 10 match ip address TO_R4 set ip next-hop 155.1.0.5 !</pre> <pre>route-map RMP-POL-LOCAL permit 20 match ip address TO_R5 set ip next-hop 155.1.146.4</pre> <p>Affects local by the router generated traffic.</p>
<h3>Routing to NBMA Interfaces</h3> <p>Possible configurations (two):</p>	<p>Use static route using Next-hop IP / LMI:</p> <pre>ip route X.X.X.X 255.255.255.255 10.0.0.1 (Frame-Relay PVC learned via LMI for 10.0.0.1)</pre> <p>Using Interface command and frame-relay map:</p> <pre>Int Serial0 frame-relay map ip 10.0.0.1 502 broadcast Exit ip route 10.0.0.1 255.255.255.255 Serial0</pre>	<p>Debugging Policy routing:</p>	<pre>R1#debug ip policy Policy routing debugging is on RIP *Jul 15 10:34:39.146: IP: s=155.1.146.6 (FastEthernet0/0), d=155.1.5.5, len 100, FIB policy match *Jul 15 10:34:39.146: IP: s=155.1.146.6 (FastEthernet0/0), d=155.1.5.5, g=155.1.0.5, len 100, FIB policy routed *Jul 15 10:39:11.902: IP: s=54.1.1.6 (FastEthernet0/0), d=54.1.2.254, len 56, FIB policy rejected(no match) - normal forwarding *Jul 15 10:39:11.906: IP: s=54.1.1.6 (FastEthernet0/0), d=54.1.2.254, len 56, policy rejected -- normal forwarding</pre>	<p>What is the pre-requisite for visible output on the following command?</p> <p>debug ip packet detail</p>	<p>Int X no ip route-cache</p> <p>debug ip packet detail</p>
<h3>Where to use Longest Match Routing?</h3>	 <p>Ip route 99.99.99.0/24 via Ser0</p> <p>Ip route 99.0.0.0/8 via Ethernet0 (longest match back route)</p> <p>If you need to create a backup path in case the primary more specific route goes down.</p>	<p>Show track brief output:</p>	<pre>R1# show track brief Track Object Parameter Value 123 rtr 1 reachability Up 124 rtr 2 reachability Up</pre>	<h3>GRE Tunneling and Recursive Routing</h3> <p>Using RIP</p>	<pre>Interface Loopback0 Ip address 150.1.9.9 255.255.255.0 (/24 Mask!) interface Tunnel34 tunnel source Loopback0 router rip distribute-list prefix PFX-DISALLOW-TUN-DST out Tunnel34 ip prefix-list PFX-DISALLOW-TUN-DST seq 5 deny 150.1.9.0/24 ip prefix-list PFX-DISALLOW-TUN-DST seq 10 permit 0.0.0.0/0 le 32</pre> <p>Dis-allow the tunnel SOURCE PREFIX out the tunnel to the other side. Need to be done on both sides!</p>
<h3>Floating Static Routes</h3>	 <p>Ip route 99.99.99.0/24 Ser0 20</p> <p>Ip route 99.99.99.0/24 Ser1 10</p> <p>Administrative distance 10 prioritized over 20, using Serial 1</p>	<h3>Reliable Policy Routing</h3> <p>(ip policy with IP SLA combined)</p>	<pre>Interface X ip policy route-map RMP-POL-R3-IN route-map RMP-POL-R3-IN permit 20 match ip address ACL_SRC_R3_DST_R5 set ip next-hop verify-availability 155.1.146.10 1 track 50 set ip next-hop verify-availability 155.1.146.20 2 track 99 set ip next-hop 155.1.146.10 set ip next-hop 155.1.146.20 set ip default next-hop 155.1.0.5</pre> <p>Sequence Nr Track Nr</p>	<h3>GRE Tunneling and Recursive Routing</h3> <p>What two solutions are there?</p>	<p>Solution A:</p> <p>Use prefix list outbound the Tunnel interface, dis-allowing the Tunnel Source to be advertise through the Tunnel.</p> <p>Solution B:</p> <p>Using static routes with lower administrative Distance.</p>
<h3>Backup Interface on switches and routers:</h3>	<p>Routers:</p> <pre>interface Serial0/0/0.10 point-to-point backup delay 3 60 backup interface Serial0/1/0</pre> <p>Switches:</p> <pre>Int Port-Channel 22 switchport backup interface Fa0/16 switchport backup interface Fa0/16 preemption mode forced switchport backup interface Fa0/16 preemption delay 20</pre>	<p>What two possible ways of tracking are available in policy routing, Unsing an IP and a non-IP protocol?</p>	<p>Tracking through IP SLA</p> <p>Tracking by the use of CDP</p>	<p>What is special about routes with Administrative Distance, but are configured as the Backup interface?</p>	<p>They can not be used, only if the primary interface goes down, are the backup interfaces/routes installed!</p>
<p>IP SLA (RTR) configuration and monitoring:</p>	<p>Config:</p> <pre>ip sla monitor 10 type echo protocol icmpEcho 150.1.1.1 source-interface Gi0/1 timeout 2000 frequency 5 ip sla monitor schedule 10 life forever start-time now</pre> <pre>ip sla 1 icmp-echo 155.1.146.1 source-interface FastEthernet0/1 timeout 2000 frequency 5 ip sla schedule 1 life forever start-time now</pre> <p>Monitoring:</p> <pre>Show rtr config Show rtr statistics Show ip sla XX</pre>	<h3>Reliable Policy Routing</h3> <p>(ip policy utilizing CDP, not IP SLA)</p>	<pre>Neighbour interface Cdp enable ip address 155.1.0.5 255.255.255.0 Interface X Cdp enable ip addr 155.1.0.1 255.255.255.0 route-map RELIABLE_POLICY_ROUTING permit 10 match ip address FROM_R3_TO_R4_LOOPBACK set ip next-hop 155.1.0.5 set ip next-hop verify-availability set ip default next-hop 155.1.146.4</pre> <p>As long as 155.1.0.5 is learned via CDP that next-hop is used as it is verified available.</p>	<h3>Reliable Backup Interface with GRE</h3> <p>Using Keepalives to detect End-To-End connectivity over Frame-Relay Networks</p>	<p>without GRE-Keepalives the router would not be able to detect an underlying path error. If the Keepalive stop, the tunnel line-protocol goes down, and the backup interface is triggered.</p> <pre>interface Tunnel0 ip address 10.0.0.5 255.255.255.0 tunnel source Serial0/0/0 tunnel destination 155.1.0.4 keepalive 1 3 backup interface Serial0/1/0</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

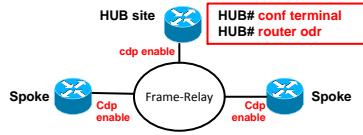
[Donate](#)



Thanks for appreciating my efforts

Colin


IP routing

<h2>On-Demand Routing (ODR)</h2>	 <pre> Spoke# sh ip route o* 0.0.0/0 [160/1] via 155.1.0.5, 00:00:28, Serial0/0/0.1 HUB#sh ip route o 150.1.4.0 [160/1] via 155.1.0.4, 00:00:16, Serial0/0/0 o 150.1.3.0 [160/1] via 155.1.0.3, 00:00:15, Serial0/0/0 </pre>	<p>What types of aggregation modes are there for OER?</p>	<p>aggregation-type {bgp non-bgp prefix-length <prefix-mask>}</p>	<p>What is measured in OER passive mode?</p>	<p>Enabled by default, uses NetFlow</p> <p>Delay - OER measures the average delay of TCP flows for a given prefix</p> <p>Packet loss - OER measures packet loss by tracking TCP sequence Numbers</p> <p>Reachability - OER measures reachability by tracking TCP SYN messages without receiving a TCP ACK packet</p> <p>throughput - OER measures throughput by measuring the total number of bytes and packets for each traffic class enabled by default</p>
<h2>On-Demand Routing (ODR)</h2>	<pre> conf t router odr timers basic <update> <invalid> <holddown> <flushtimer> distribute-list route-map in interface X Use distribute-list to filter the routes in a "complex" ODR scenario! ☺ </pre>	<p>OER:</p> <pre> Periodic interval <minute> prefixes <number> </pre>	<p>OER Prefix learning starts every time interval</p> <p>By default, 100 top flows are learned</p>	<p>What are properties of OER active mode?</p>	<p>synthetic traffic are applied to the corresponding traffic class in the MTC list</p> <p>OER activates the probes on all the Border Routers</p> <pre> oer-map OER 20 set mode monitor active oer master active-probe tcp-conn 150.1.1.1 target-port 23 active-probe tcp-conn 150.1.4.4 target-port 23 </pre>
<p>OER / Pfr Phases:</p>	<p>(1) OER Profile phase: Discovering traffic classes.</p> <p>(2) OER Measure phase: OER Border Routers measure traffic actively (SLA) or passively (Netflow)</p> <p>(3) OER Apply Policy phase: Master controller has set of thresholds</p> <p>(4) OER Control phase. Injecting routes, changing metrics</p> <p>(5) OER Verify phase. Verifies new policies, checking if in-policy</p>	<p>How do you exclude certain prefixes from OERs monitoring?</p>	<pre> ip prefix-list MONITOR deny 114.0.1.0/24 ip prefix-list MONITOR permit 114.0.0.0/8 ! oer-map OER 10 match ip address prefix-list MONITOR oer master policy-rules OER </pre>	<p>show oer master prefix</p> <p>Output explanations:</p>	<p>PasSDly – passively measured delay, short-range (5 minutes interval)</p> <p>ActSDly – actively measured delay, short-range</p> <p>PasLDly – passively measured delay, long-range (60 minute interval)</p> <p>PasSun/PasLun – passively measured short and long range unreachable metric</p> <p>ActSun/ActLun – the same unreachable, just measured actively</p> <p>PasSLoS/PasLSLoS – short and long range loss, measured passively</p> <p>EBw/Bw – egress and ingress bandwidth usage for this class in kbps</p> <p>State – any of the states from the traffic state diagram</p> <p>Time – the amount of time spent in the state</p> <p>Curr BR – current Border Router selected for this class</p> <p>Curr I/F – current exit interface selected for this class</p> <p>Protocol – protocol used to influence routing for this traffic class</p>
<p>OER traffic class</p>	<p>Can be a network prefix</p> <p>Can be a complex object consisting of network prefix and application port number.</p>	<p>IP SLA jitter emulating G.729 codec between R6 and SW2</p>	<p>R6:</p> <pre> ip sla 2 udp-jitter 150.1.8.8 16384 source-ip 150.1.6.6 codec g729a codec-numpackets 10 codec-interval 10 exit ! ip sla schedule 2 life forever start-time now SW2: ip sla responder </pre>	<p>show oer master active-probes</p>	<pre> Rack1R5#show oer master active-probes OER Master Controller active-probes Border = Border Router running this Probe State = On/Assigned to a Prefix Prefix = Probe is assigned to this Prefix Type = Probe Type Target = Target Address TPort = Target Port How = Was the probe Learned or Configured N = Not applicable The following Probes exist: State Prefix Type Target TPort How Codec Assigned 150.1.6.0/24 tcp-conn 150.1.6.6 23 Cfgd N Assigned 150.1.4.0/24 tcp-conn 150.1.4.4 23 Cfgd N Assigned 150.1.1.0/24 tcp-conn 150.1.1.1 23 Cfgd N Assigned 150.1.6.0/24 echo 150.1.6.6 N Lrnd N Assigned 150.1.1.0/24 echo 150.1.1.1 N Lrnd N Assigned 150.1.4.0/24 echo 150.1.4.4 N Lrnd N The following Probes are running: Border State Prefix Type Target TPort 150.1.3.3 ACTIVE 150.1.1.0/24 echo 150.1.1.1 N 150.1.3.3 ACTIVE 150.1.1.0/24 echo 150.1.1.1 N 150.1.5.5 ACTIVE 150.1.1.0/24 echo 150.1.1.1 N 150.1.3.3 ACTIVE 150.1.6.0/24 echo 150.1.6.6 N </pre>
<p>OER Automatic prefix traffic class learning is based on what:</p>	<p>Master Controller collects NetFlow information from the Border Routers</p>	<p>show oer border</p>	<pre> show oer border oer on 150.1.2.2 ACTIVE, MC 150.1.5.5 UP/DOWN: UP 00:20:04, Auth Failures: 0 Conn Status: SUCCESS, PORT: 3949 Exits Fa0/0 EXTERNAL Se0/0.1 INTERNAL </pre>	<h2>RIPv2 Authentication</h2> <p>Configuration:</p>	<p>R1:</p> <pre> key chain RIP key 1 key-string CCIE ! interface FastEthernet0/0 ip rip authentication mode text (ip rip authentication mode md5) ip rip authentication key-chain RIP </pre>
<p>oer master learn throughpug delay</p>	<p>highest <i>outbound</i> throughput (number of bytes transferred)</p> <p>delay (RTT time)</p>	<p>show oer master</p>	<pre> Rack1R5#show oer master OER state: ENABLED and ACTIVE Conn Status: SUCCESS, PORT: 3949 Version: 2.2 Number of Border routers: 3 Number of Exits: 4 Number of monitored prefixes: 0 (max 5000) Max prefixes: total 5000 learn 2500 Prefix count: total 0, learn 0, cfg 0 PBR Requirements not met Nbar Status: Inactive Border Status UP/DOWN AuthFail Version 150.1.5.5 ACTIVE UP 00:01:12 0 2.2 150.1.2.2 ACTIVE UP 00:11:08 0 1.0 150.1.3.3 ACTIVE UP 00:12:06 0 1.0 </pre>	<h2>How to find spaces within Passwords:</h2> <p>What can "show key chain" reveal?</p>	<p>Can identify spaces within a password at its end, which can not be seen otherwise.</p> <pre> R6# show key chain Key-chain RIP: key 1 - text "CCIE" accept lifetime (always valid) - (always valid) [valid now] send lifetime (always valid) - (always valid) [valid now] show run i CCIE\$ R2#sh run i CCIE\$ No output due to the password being "CCIE" R2#sh run i CCIE \$ key-string CCIE </pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin

<p>Clearing routing processes in RIP, EIGRP, OSPF and BGP:</p>	<p>RIP: clear ip route *</p> <p>EIGRP: clear ip eigrp 100 neighbors</p> <p>OSPF: clear ip ospf process yes</p> <p>BGP: clear ip bgp *</p>	<p>PREFIX Lists</p> <p>Allow class B networks that are or are not subnetted</p>	<pre>ip prefix-list PFX seq 5 permit 128.0.0.0/2 ge 16</pre> <pre>R 128.1.0.0/16 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 132.1.0.0/24 is subnetted, 1 subnets</pre> <pre>R 132.1.1.0 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 191.1.0.0/16 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre>	<p>What will R1 routing table show?</p> <p>What will R1 RIP and OSPF database show?</p>	<pre>R1#show ip route 1.0.0.0</pre> <pre>O IA 1.0.0.0/24 [110/11] via 2.0.0.3, 00:00:23, e0/0.13</pre> <pre>R1#show ip rip database 1.0.0.0</pre> <pre>1.0.0.0/24 auto-summary</pre> <pre>1.0.0.0/24 redistributed</pre> <pre>[1] via 10.1.13.3, from 0.0.0.3,</pre> <pre>R1#show ip ospf database</pre> <pre>Summary Net Link States (Area 0)</pre> <pre>1.0.0.0 0.0.0.3 752 0x80000001 0x00D1E6</pre> <p>OSPF has the better admin distance than RIP! Therefore the OSPF route ends in the routing table!</p>												
<p>Don't waste time looking at stars! If you forgotten how to use control, shift, 6 and X or what it was, change it to something you remember:</p> <p>R1#traceroute 2.2.2.2 numeric</p> <p>Type escape sequence to abort. Tracing the route to 2.2.2.2 VRF info: (vrf in name/id, vrf out name/id)</p> <pre>1 * * *</pre> <pre>2 * * *</pre> <pre>3 * * *</pre> <pre>4 *</pre>	<p>Change the escape-character to your easy to remember escape-character:</p> <p>Permanent: line con 0 Per session: R1#terminal escape-character 27 escape-character 27</p> <table border="1"> <thead> <tr> <th>Decimal</th> <th>Hex</th> <th>Key</th> </tr> </thead> <tbody> <tr> <td>27</td> <td>1B</td> <td>ESC</td> </tr> <tr> <td>3</td> <td>03</td> <td>Ctrl-C</td> </tr> <tr> <td>127</td> <td>7F</td> <td>Delete</td> </tr> </tbody> </table> <p>Lookup in document: "ASCII Character Set and Hex Values"</p>	Decimal	Hex	Key	27	1B	ESC	3	03	Ctrl-C	127	7F	Delete	<p>PREFIX Lists:</p> <p>Allow class C networks that are or are not subnetted:</p>	<pre>ip prefix-list PFX seq 5 permit 192.0.0.0/3 ge 24 le 32</pre> <pre>R 192.1.1.0/24 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre> <pre>R 193.1.1.0/25 is subnetted, 1 subnets</pre> <pre>R 193.1.1.0 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre> <pre>R 194.1.1.0/26 is subnetted, 1 subnets</pre>	<p>What will R1 routing table show?</p> <p>What will R1 RIP and OSPF database show?</p>	<pre>R1#show ip route</pre> <pre>R 1.0.0.0/24 [120/1] via 10.1.12.2, 00:00:01, e0/0.12</pre> <pre>R1#show ip rip database</pre> <pre>1.0.0.0/24 auto-summary</pre> <pre>1.0.0.0/24</pre> <pre>[1] via 10.1.12.2, 00:00:00, e0/0.12</pre> <pre>R1#show ip ospf database 1.0.0.0</pre> <pre>-> EMPTY</pre> <p>As soon as the OSPF route disappears, the RIP route to 1.0.0.0/24 will be visible via the RIP domain! (-> OSPF has the better Admin Distance)</p>
Decimal	Hex	Key															
27	1B	ESC															
3	03	Ctrl-C															
127	7F	Delete															
<p>PREFIX Lists:</p> <p>Match all unsubnetted /8 prefixes:</p>	<p>ALL unsubnetted /8</p> <pre>ip prefix-list BLA permit 0.0.0.0/1 ge 8 le 8</pre> <pre>R 1.0.0.0/8 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 2.0.0.0/8 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 3.0.0.0/8 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <p>0nnnnnnn hhhhhhhh hhhhhhhh hhhhhhhh First bit has to be "0", Up to the 8th bit could be 1 or 0. Initial byte: 0 - 127</p>	<p>PREFIX Lists:</p> <p>Allows networks with a prefix-length of 25 or greater in its routing table:</p>	<pre>ip prefix-list PFX seq 5 permit 0.0.0.0/0 ge 25</pre> <pre>6.0.0.0/26 is subnetted, 1 subnets</pre> <pre>R 6.6.6.0 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 193.1.1.0/25 is subnetted, 1 subnets</pre> <pre>R 193.1.1.0 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre>	<p>Will R4 see 1.0.0.0 in its routing table?</p>	<p>In this case we are redistributing RIP prefixes into EIGRP. 1.0.0.0/24 will show up in R4's routing table as the prefix is NOT learned from the OSPF database.</p> <p>If the prefix 1.0.0.0/24 "would" be learned via OSPF (no shut of Lo0) with a lower Admin distance, the route would NOT be redistributed into EIGRP as the route 1.0.0.0 is installed into the routing table based on OSPF and not RIP!</p> <p>1.0.0.0 is still visible in the RIP database but learned from OSPF!</p> <pre>Admin Distance:</pre> <pre>RIP 120</pre> <pre>OSPF 110</pre> <pre>EIGRP 90</pre> <pre>R1#show ip route 1.0.0.0</pre> <pre>Known via "rip"</pre>												
<p>PREFIX Lists:</p> <p>Match all unsubnetted /16 prefixes:</p>	<pre>ip prefix-list ROUTES seq 5 permit 128.0.0.0/2 ge 16 le 16</pre> <pre>R 128.1.0.0/16 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre> <pre>R 191.1.0.0/16 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre> <p>10nnnnnnn nnnnnnnn hhhhhhhh hhhhhhhh First two bits have to be "10", Up to the 16th bit could be 1 or 0. Initial byte: 128 - 191</p>	<p>PREFIX Lists:</p> <p>Allow networks with a prefix-length of 16 or less in its routing table:</p>	<pre>ip prefix-list PFX seq 5 permit 0.0.0.0/0 le 16</pre> <pre>R 3.0.0.0/8 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 4.0.0.0/16 is subnetted, 1 subnets</pre> <pre>R 4.4.0.0 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 125.0.0.0/8 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 128.1.0.0/16 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre>	<p>Will R4 see 1.0.0.0 in its routing table?</p>	<p>R1 learns 1.0.0.0 from 110 and 120!</p> <p>1.0.0.0 is inserted by OSPF into R1's routing table. Therefore, 1.0.0.0 will NOT be redistributed into EIGRP, due to the fact that 1.0.0.0 was inserted to R1's routing table via OSPF and not RIP!</p> <p>1.0.0.0 is found in R1's RIP database, but learned via OSPF.</p> <pre>Admin Distance:</pre> <pre>RIP 120</pre> <pre>OSPF 110</pre> <pre>EIGRP 90</pre> <pre>R1#show ip rip database</pre> <pre>1.0.0.0/24 auto-summary</pre> <pre>1.0.0.0/24 redistributed</pre> <pre>[1] via 10.1.13.3, from 0.0.0.3,</pre> <pre>R1#show ip route 1.0.0.0</pre> <pre>Known via "ospf 1"</pre>												
<p>PREFIX Lists:</p> <p>Match all unsubnetted /24 prefixes:</p>	<pre>ip prefix-list ROUTES seq 5 permit 192.0.0.0/3 ge 24 le 24</pre> <pre>R 192.1.1.0/24 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <pre>R 195.1.1.0/24 [120/1] via 10.1.12.1, 00:00:01, e0/0</pre> <p>110nnnnnnn nnnnnnnn hhhhhhhh hhhhhhhh First three bits have to be "110", Up to the 24th bit could be 1 or 0. Initial byte: 192 - 223</p>	<p>PREFIX Lists</p> <p>Allow networks with a prefix-length of 16 to 25 in its routing table:</p>	<pre>ip prefix-list PFX seq 5 permit 0.0.0.0/0 ge 16 le 25</pre> <pre>R 128.1.0.0/16 [120/1] via 10.1.12.1, 00:00:00, e0/0</pre> <pre>R 132.1.0.0/24 is subnetted, 1 subnets</pre> <pre>R 132.1.1.0 [120/1] via 10.1.12.1, 00:00:00, e0/0</pre> <pre>R 191.1.0.0/16 [120/1] via 10.1.12.1, 00:00:00, e0/0</pre> <pre>R 192.1.1.0/24 [120/1] via 10.1.12.1, 00:00:00, e0/0</pre> <pre>R 193.1.1.0/25 is subnetted, 1 subnets</pre> <pre>R 193.1.1.0 [120/1] via 10.1.12.1, 00:00:00, e0/0</pre>	<p>How can you solve this situation so that 1.0.0.0/24 ends up in R4's routing table?</p>	<p>Increasing the Admin Distance of the OSPF route, in order to trigger the RIP route on R1!</p> <pre>R1#</pre> <pre>router ospf 1</pre> <pre>distance 121 0.0.0.3 0.0.0.0 99</pre> <pre>access-list 99 permit 1.0.0.0</pre>												
<p>PREFIX Lists:</p> <p>Allow Class A networks which are NOT subnetted:</p>	<p>- class A, with a prefix-length of greater than or equal 8 and less than or equal 32 - class A plus all class A networks that are subnetted</p> <pre>ip prefix-list PFX permit 0.0.0.0/1 ge 8 le 32</pre> <pre>R 3.0.0.0/8 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre> <pre>R 4.0.0.0/16 is subnetted, 1 subnets</pre> <pre>R 4.4.0.0 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre> <pre>R 5.0.0.0/24 is subnetted, 1 subnets</pre> <pre>R 5.5.5.0 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre> <pre>R 6.0.0.0/26 is subnetted, 1 subnets</pre> <pre>R 6.6.6.0 [120/1] via 10.1.12.1, 00:00:02, e0/0</pre>	<p>PREFIX lists:</p> <p>Only permit the 10.7.0.0/22 networks, using one line:</p> <pre>10.7.1.1 255.255.255.0</pre> <pre>10.7.0.1 255.255.255.0</pre> <pre>10.7.3.1 255.255.255.0</pre> <pre>10.7.2.1 255.255.255.0</pre> <pre>10.7.0.1 255.255.255.0</pre>	<p>Checks the first 22 bits, Then the subnet mask needs to be GE 23 and LE 24</p> <pre>ip prefix-list PFX seq 5 permit 10.7.0.0/22 ge 23 le 24</pre> <pre>10.7.1.1 255.255.255.0</pre> <pre>10.7.0.1 255.255.255.0</pre> <pre>10.7.3.1 255.255.255.0</pre> <pre>10.7.2.1 255.255.255.0</pre> <pre>10.7.0.1 255.255.255.0</pre>	<p>ping ipv6</p> <p>And all its options:</p>	<pre>R6#ping ipv6</pre> <pre>Target IPv6 address: 12::4</pre> <pre>Repeat count [5]:</pre> <pre>Datagram size [100]:</pre> <pre>Timeout in seconds [2]:</pre> <pre>Extended commands? [no]: y</pre> <pre>Source address or interface: 12::6</pre> <pre>UDP protocol? [no]:</pre> <pre>Verbose? [no]:</pre> <pre>Precedence [0]: 0</pre> <pre>OSPF [0]: 20</pre> <pre>Include hop by hop option? [no]:</pre> <pre>Include destination option? [no]:</pre> <pre>Sweep range of sizes? [no]:</pre> <pre>Type escape sequence to abort.</pre> <pre>Sending 5, 100-byte ICMP Echos to 12::4, timeout is 2 seconds:</pre> <pre>Packet sent with a source address of 12::6</pre> <pre>!!!!</pre> <pre>Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms</pre>												

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

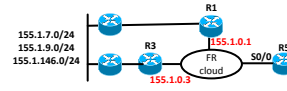
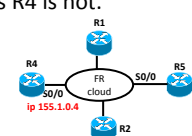
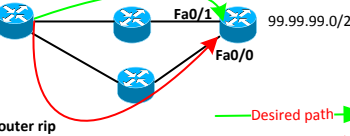
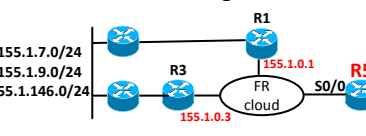
Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin



RIP

<p>debug ip rip</p> <p>Output, seeking for authentication:</p>	<pre>int fa0/0 ip rip authentication mode md5 ip rip authentication key-chain RIP debug ip rip RIP: received packet with MD5 authentication RIP: received v2 update from 192.10.1.254 on FastEthernet0/0 key chain RIP key 1 key-string CCIE int fa0/0 ip rip authentication mode text ip rip authentication key-chain RIP debug ip rip RIP: received packet with text authentication CCIE RIP: received v2 update from 155.1.146.1 on GigabitEthernet0/0</pre> <p><i>ALWAYS KEY-CHAIN FIRST -> order of operation</i></p>	<p>Prefix lists for:</p> <p>The default route:</p> <p>ALL routes (any):</p>	<p>The default route:</p> <pre>ip prefix-list DEFAULT permit 0.0.0.0/0</pre> <hr/> <p>ANY ANY, all routes:</p> <pre>ip prefix-list ANY permit 0.0.0.0/0 le 32</pre>	 <p>Explain RIP filtering with Extended ACLs:</p> <pre>R5# router rip distribute-list 100 in Serial0/0 access-list 100 deny ip host 155.1.0.3 host 155.1.7.0 access-list 100 deny ip host 155.1.0.3 host 155.1.9.0 access-list 100 permit ip any any</pre>	<pre>router rip distribute-list 100 in Serial0/0 access-list 100 deny ip host 155.1.0.3 host 155.1.7.0 Prefix/network to be filtered access-list 100 deny ip host 155.1.0.3 host 155.1.9.0 Source Next-Hop access-list 100 permit ip any any Permit all other prefixes</pre>
<p>Configuring IP rip versions (send/receive):</p>	<pre>router rip network 150.1.0.0 no auto-summary interface fa0/0 ip address 155.1.108.8 255.255.255.0 ip rip send version 1 ip rip receive version 1</pre>	<p>Filtering RIP route updates from a single router, connected to a frame-relay cloud using extended access-lists</p> <p>R1 – R3 should be allowed updates, whereas R4 is not.</p> 	<p>What differences are there with Extended access-lists between IGP and BGP?</p> <pre>router rip version 2 distribute-list prefix PERMIT-ALL-OTHER gateway DENY-R4-IN in ip prefix-list DENY-R4-IN seq 5 deny 155.1.0.4/32 ip prefix-list DENY-R4-IN seq 10 permit 0.0.0.0/0 le 32 ip prefix-list PERMIT-ALL-OTHER seq 5 permit 0.0.0.0/0 le 32</pre>	<p>Use static Null routes for summaries.</p> <p>Use distribute lists, which prevent the outbound summaries being learned back to the router.</p>	<p>Extended ACLs:</p> <p>BGP:</p> <p>Source field = Network Address Destination field = Subnet Mask</p> <pre>access-list 100 permit ip 10.1.1.0 0.0.0.0 255.255.255.0 0.0.0.0</pre> <p>IGP:</p> <p>Source field = Source Next-Hop Destination field = Network Address</p> <pre>access-list 100 deny ip host 155.1.0.3 host 155.1.7.0</pre>
<p>Configuring IP RIP summarization:</p>	<pre>interface Serial0/0/0.1 point-to-point ip summary-address rip 30.0.0.0 255.252.0.0 ip summary-address rip 31.0.0.0 255.252.0.0</pre> <p>Don't forget to turn off auto-summary!!</p>	<p>What can be used to handle route feedback with RIP ?</p>	<p>Use static Null routes for summaries.</p> <p>Use distribute lists, which prevent the outbound summaries being learned back to the router.</p>	<p>Explain the different extended ACLs used with BGP:</p> <pre>access-list 100 permit ip 10.1.1.0 0.0.0.0 255.255.255.0 0.0.0.0 access-list 100 permit ip 10.0.0.0 0.0.255.0 255.255.255.0 0.0.0.0 access-list 100 permit ip 10.0.0.0 0.0.0.255 255.255.255.240 0.0.0.0</pre>	<p>Matches 10.1.1.0/24 – Only</p> <p>Matches 10.0.X.0/24 – Any for X with /24 prefix length</p> <p>Matches 10.0.0.X/28 – Any for X with /28 prefix length</p>
<p>Changing RIP's default timers:</p>	<pre>router rip version 2 timers basic 10 60 60 80 100 timers <update-interval> <invalid> <holddown> <flush> <sleep-time in msec> Show ip protocol helps to see the original timers in case you are asked to tune them to a third of the defaults which are: timers basic 30 180 180 240 100 Interface Fa0/0 Ip rip advertise 30 <- Updates per interface config</pre> <p><i>Postpone periodic updates for 100 msec until sent</i></p>	<p>Standard ACLs, filtering all ODD IPs on the 2nd Octet:</p>	<p>ALL ODD IPs in 2nd octet:</p> <pre>access-list 50 permit 0.1.0.0 255.254.255.255</pre> <p>00000001</p> <p>.254. instructs to look at the last bit within the 2nd octet.</p>	<p>RIPv2 Filtering with Offset Lists</p> <p>Alternative to passive interface, per prefix:</p>	<pre>router rip offset-list 1 out 16 Vlan79 access-list 1 permit 155.1.5.0 Setting Hop-Count to 16 for prefix 155.1.5.0/24 outbound of Vlan79 router rip offset-list 22 out 9 Fa0/0 access-list 22 deny 1.0.0.0 access-list 22 permit any</pre> <p><i>Set offset to 9 for all routes except for 1.0.0.0/8</i></p>
<p>How to check timers of RIP protocol:</p>	<pre>R1#show ip protocols include seconds Sending updates every 30 seconds, next due in 22 seconds Invalid after 180 seconds, hold down 180, flushed after 240</pre>	<p>Standard ACL, filtering all EVEN IPs on the 2nd Octet:</p>	<p>ALL EVEN IPs in 2nd octet:</p> <pre>access-list 51 permit 0.0.0.0 255.254.255.255</pre> <p>00000000</p> <p>.254. instructs to look at the last bit within the 2nd octet.</p>	<p>RIPv2 Filtering with Administrative Distance</p>	<p>Filtering prefixes by setting the admin distance to UNKNOWN for certain prefixes: (locally significant only)</p> <pre>router rip distance 255 0.0.0.0 255.255.255.255 88 From any neighbor access-list 88 permit 150.1.4.0</pre>
<p>RIP offset list configuration:</p>	 <pre>router rip version 2 offset-list 1 out 4 fa0/0 network 150.1.0.0 no auto-summary access-list 1 permit 99.99.99.0 0.0.0.0 (dst prefix) Adding 4 Hops to the path out over Fa0/0, making sure 99.99.99.0 is reached via Fa0/1</pre> <p>Desired path →</p> <p>current →</p>	<p>RIP From R5: Allowing Vlan 7, 9 via R1.</p> <p>Allowing Vlan 146 and R1's Lo0 via R3 config:</p> 	<pre>R5# router rip distribute-list 100 in Serial0/0 access-list 100 deny ip host 155.1.0.3 host 155.1.7.0 access-list 100 deny ip host 155.1.0.3 host 155.1.9.0 access-list 100 deny ip host 155.1.0.1 host 155.1.146.0 access-list 100 deny ip host 155.1.0.1 host 150.1.1.0 access-list 100 permit ip any any Allow Vlan 7 and Vlan 9 via R1 while denying via R3. Allow Vlan 146 and R1's Loopback 155.1.1.1 via R3, while denying via R1</pre>	<p>RIPv2 Filtering with Per Neighbor AD</p>	<pre>router rip distance 255 155.1.37.3 0.0.0.0 22 Source Next-Hop Router Exact match access-list 22 permit 150.1.3.0 Filtering out prefix 150.1.3.0/24 from Neighbour 155.1.37.3 by setting the Admin Distance to UNKNOWN 255.</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<p>How can the distance command be applied with RIP ?</p>	<ul style="list-style-type: none"> - globally for the routing process - globally for the routing process per route type - on a per-prefix basis - on a per-neighbor per-prefix basis 	<h3>RIPv2 Unicast Updates</h3>	<p>Router rip passive-interface default neighbor 1.1.1.1</p> <p>If passive int not set, router can still process Multicast announcements!</p> <p>Can be verified by show ip protocol</p>	<p>Tricky RIP problem</p> <p>Find the fault:</p> <pre>R1#ping 4.4.4.4 Not successful show ip route i 4.4.4.4 R 4.4.4.4 [120/3] via x.x.x.x, 00:00:06, Fa0/1</pre>	<pre>R1#ping 4.4.4.4 Not successful show ip route i 4.4.4.4 R 4.4.4.4 [120/3] via x.x.x.x, 00:00:06, Fa0/1</pre> <p>RIP routes visible! But no connectivity</p> <p>R2# int ser0/0 ip address 10.0.0.1/24</p> <p>R3# int ser0/0 no ip address</p> <p>ALL routers have rip configured correctly</p>
<p>How can one prevent route-feedback in RIP?</p>	<p>Configure static routes for summaries pointing to the NULL0.</p> <p>Make sure that the summary is denied coming back to the router via Distribute-list or similar via one path, but is allowed via the other.</p>	<h3>RIPv2 Broadcast Updates</h3>	<pre>R6#sh run int gi0/0.146 interface GigabitEthernet0/0.146 encapsulation dot1Q 146 ip address 155.1.146.6 255.255.255.0 ip rip v2-broadcast</pre> <p>Debug ip packet: IP: s=155.1.146.6 (local), d=255.255.255.255 (GigabitEthernet0/0.146), len 272, sending broad/multicast</p> <p>Broadcast can be changed from 255.255.255.255 to: int gi0/0.146 ip broadcast-address 155.1.146.255</p>	<h3>RIP BFD</h3> <h3>RIPv2 BFD</h3>	<pre>router rip version 2 bfd all-interfaces network 10.0.0.0 neighbor 10.10.20.2 bfd</pre> <p>int fa0/1 ip address 10.10.20.1 255.255.255.0 bfd interval 50 min_rx 50 multiplier 5</p> <p>int fa0/2 ip address 10.90.90.1 255.255.255.0 bfd interval 50 min_rx 50 multiplier 5</p> <p>debug ip rip bfd events</p>
<h3>RIPv2 Default Routing</h3> <p>Setting the source</p>	<pre>router rip default-information originate route-map DEFAULT_TO_R1 route-map DEFAULT_TO_R1 permit 10 set interface FastEthernet0/0.146</pre>	<h3>RIPv2 Triggered Updates</h3>	<p>RIP sends entire routing table every 30 seconds by default:</p> <pre>interface Serial0/0/0.1 point-to-point ip address 155.1.0.4 255.255.255.0 ip rip triggered</pre> <p>Triggered updates used in order to Support Demand Circuits. With the triggered option RIP converts in an EVENT triggered protocol, only sending changes based on changes to its database.</p> <p>SHUT / NO SHUT to take action!</p>	<p>What are all the variations one can run RIPv1 and RIPv2 in?</p>	<p>Broadcast ip rip v2-broadcast</p> <p>send version [1,2]</p> <p>multicast</p> <p>Unicast Neighbor cmd and Passive interface</p> <p>Directed broadcast neighbor</p>
<h3>RIPv2 Conditional Default Routing</h3> <p>Announcing a default route, if another prefix exists utilizing a RMP:</p>	<pre>router rip version 2 default-information originate route-map RMP-DEFAULT route-map RMP-DEFAULT permit 10 match ip address prefix-list CHECK_EXISTING_ROUTE ip prefix-list CHECK_EXISTING_ROUTE seq 5 permit 204.12.1.0/24 ! IF 204.12.1.0/24 exists in the routing table, announce the default-route via RIP.</pre>	<p>How can triggered updates within RIP be identified by using:</p> <p>show ip rip database:</p>	<pre>interface Serial0/1/0 ip rip triggered</pre> <p>Show ip rip database</p> <pre>150.1.5.0/24 [1] via 155.1.45.5, 00:00:16 (permanent), Serial0/1/0 [1] via 155.1.0.5, 00:00:01, Serial0/0/0.1</pre>	<p>Filter RIP routes so that R2 will not see the announcements of R1:</p>	<p>router rip passive interface default</p> <p>int x no ip route-cache debug ip packets should not see RIP updates sourced from R1 (mcast)</p>
<h3>RIPv2 Reliable Conditional Default Routing</h3>	<pre>ip sla 1 icmp-echo 204.12.1.254 source-interface FastEthernet0/0 timeout 50 frequency 1 ip sla schedule 1 start-time now track 99 rtr 1 ! router rip default-information originate route-map RMP-TRACK ip route 169.254.0.1 255.255.255.255 Null0 track 99 ip prefix-list PFX-DUMMY seq 5 permit 169.254.0.1/32 route-map RMP-TRACK permit 10 match ip address prefix-list PFX-DUMMY Tracking Route the DUMMY route via SLA and Track 99, if successful announce the default route.</pre>	<h3>RIPv2 Source Validation</h3>	<p>R1: interface Serial0/1 ip address negotiated encapsulation ppp</p> <p>router rip no validate-update-source</p> <p>R3: interface Serial1/2 encapsulation ppp peer default ip address 155.1.13.1</p> <p><i>the two links negotiate a host route and are no longer considered on the same network, no validate needed to get RIP process talking</i></p>	<h3>Speciality about the RIP Process:</h3>	<p>The RIP process does not start unless you have entered a network statement below:</p> <pre>router rip network X.X.X.X</pre>
<h3>How to disable split-horizon on RIP?</h3>	<p>Conf t Int fa0/0 no ip split-horizon End</p>	<h3>How to inject a RIP default route</h3> <ol style="list-style-type: none"> 1) int fa0/x ip summary-address rip 0.0.0.0 0.0.0.0 (per neighbor) 2) router rip default-information originate route-map RMP (use route-map if used with conditions) 3) ip route 0.0.0.0 0.0.0.0 Null0 3a) router rip network 0.0.0.0 3b) router rip redistribute static 4) ip default network 1.0.0.0 <p>If the static route from 3) is removed, ALL routes will be advertised!</p>	<p>RIPv1 / RIPv2 default route:</p> <p>router rip network 1.0.0.0 offset-list 0 out 7</p>	<p>How will R2 routing table look in regards to 1.0.0.0/8</p> <p>router rip network 1.0.0.0 offset-list 0 out 7</p> <p>offset-list 0 will add 7 hops to ALL routes!!</p> <pre>R2# R 1.0.0.0/8 [120/8] via 10.0.0.1, 00:00:03, e0/0</pre>	

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

1.0.0.1/8 | R1 --- R2 | 2.0.0.1/8

Explain what the following does:
router rip
no validate-update-source

```
R1#
int ser0/0
ip address 10.0.0.1 255.255.255.0

router rip
no validate-update-source

R2#
int ser0/0
ip address 22.0.0.22 255.255.255.0

router rip
no validate-update-source

R2#
R 1.0.0.0/8 [120/2] via 10.0.0.1, 00:00:03
```

RIP:

How to add an offset value of 7 to ALL routes

```
router rip
offset-list 0 out 7
```

What filtering options are there in regards to RIP?

- Option 1: **router rip**
distribute-list X
- Option 2: **router rip**
offset-list X
- Option 3: **router rip**
distance 255 ...

Bunch of routes

R1 --- R2 --- R3

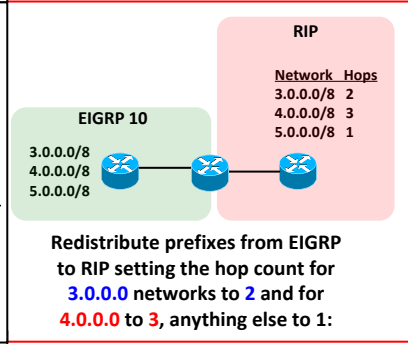
deny x.x.x.x/8 to x.x.x.x/10

router RIP

Filter out routes with a /8 to /10 prefix sourced from Router R1

```
ip prefix-list NET permit 0.0.0.0/0 ge 8 le 10
ip prefix-list R2 permit 10.0.0.2/32

router rip
distribute-list prefix NET gateway R2 in
```



```
router rip
version 2
redistribute eigrp 10 route-map EIGRP-2-RIP

ip prefix-list PFX-3 seq 5 permit 3.0.0.0/8
ip prefix-list PFX-4 seq 5 permit 4.0.0.0/8

route-map EIGRP-2-RIP permit 10
match ip address prefix-list PFX-3
set metric 2

route-map EIGRP-2-RIP permit 20
match ip address prefix-list PFX-4
set metric 3

route-map EIGRP-2-RIP permit 30
```

What filtering options are there in regards to RIP while redistribution routes from other protocols?

- 1) matching the prefixes with a ACL/PFX-list and denying them into RIP.
- 2) matching prefix via ACL/PFX-list and setting hop count (to infinity) for the different prefixes via route-map
- 3) Setting a tag value on the originating router, denying that tag to be redistributed using route-maps
- 4) using distribute lists
- 5) using an offset-list (RIP sourced PFXs)

1.0.0.1/8 | R1 --- R2 | 2.0.0.1/8

R1 and R2 # int e0/0
ip rip authentication mode md5
ip rip authentication key-chain BLA

R1: key chain BLA
key 1
key-string CISCO

R2: key chain BLA
key 2
key-string CISCO

```
R1#show ip route
NO ROUTE

R2#show ip route
R 1.1.1.1 [120/1] via 10.0.0.1, 00:00:24, e0/0
```

What methods are there to bypass RIP's sanity checks if you have the following situation:

1. host route to the other interface IP specifying the local interface (ip route x.x.x.x/32 fa0/x)
2. on the connecting interface use: **ip unnumbered LoX** in order to bypass the sanity check, using the interface rather than the IP address as next-hop
3. change the encapsulation to PPP and do **router rip** **no validate-update-source**
4. "adjust" the problem by using a correct secondary IP address (may not be a valid CCIE Lab answer)

How to bypass RIP's sanity check by using ip unnumbered LoX

```
R1#
int ser0/0
ip address unnumbered lo88

int lo88
ip address 192.168.1.1 255.255.255.0

router rip
version 2
no auto
network 192.168.1.0

R2#
int ser0/1
ip address unnumbered lo99

int lo99
ip address 10.0.0.2 255.255.255.0

router rip
version 2
no auto
network 10.0.0.0
```

What does the following RIP command do?

router rip
flash-update-threshold <seconds>
10

Something happened here

Regular update

Regular update

Suppresses triggered updates when the arrival of a regularly scheduled update matches, or is less than the number of seconds that is configured.

If configured visible via **show ip protocol | i Flash**
Flash update is suppressed when next update due within 10 seconds

If not used, show ip protocol:
Sending updates every 60 seconds, next due in 41 seconds

How to bypass RIP's sanity check by using the "host route method"

```
R1#
int fa0/0
ip address 192.168.1.1 255.255.255.0
ip route 10.0.0.2 255.255.255.0 fa0/0

router rip
version 2
no auto
no validate-update-source
Network 192.168.1.0

R2#
int fa0/1
ip address 10.0.0.2 255.255.255.0
ip route 192.168.1.1 255.255.255.0 fa0/0

router rip
version 2
no auto
no validate-update-source
Network 10.0.0.0
```

How to bypass RIP's sanity check by changing the encapsulation to PPP

```
R1#
int ser0/0
ip address 192.168.1.1 255.255.255.0
encapsulation ppp
shutdown
no shut

router rip
version 2
no auto
network 192.168.1.0
no validate-update-source

R2#
int ser0/1
ip address 10.0.0.2 255.255.255.0
encapsulation ppp
shutdown
no shutdown

router rip
version 2
no auto
network 10.0.0.0
no validate-update-source
```

How to check which key-chain is used on which interface with RIP?

```
R2#show ip protocols
....
Interface Send Recv Triggered KEY-chain
FastEthernet0/0 2 2 KEY
FastEthernet0/1 2 2 KEY23

key chain KEY
key 1
key-string cisco

interface FastEthernet0/0
ip address x.x.x.x 255.255.255.0
ip rip authentication key-chain KEY

key chain KEY23
key 1
key-string cisco23

interface FastEthernet0/1
ip address y.y.y.y 255.255.255.0
ip rip authentication mode md5
ip rip authentication key-chain KEY23
```

How to "bypass" RIP's sanity check by using a secondary address:

```
R1#
int fa0/0
ip address 192.168.1.1 255.255.255.0
ip address 10.0.0.1 255.255.255.0 secondary

router rip
version 2
no auto
network 10.0.0.0
network 192.168.1.0

R2#
int fa0/1
ip address 10.0.0.2 255.255.255.0

router rip
version 2
no auto
network 10.0.0.0
```

How can you instruct the fast RIP router to slow down its pace while sending updates to the slow one?

```
Fast Router running RIP --- Slow Router running RIP
```

FastRouter#
router rip
output-delay <50 in msec>

The fast router adds 50 msec delay between each of the update packet.


On which router would you configure:
ip default-network x.x.x.x

In this RIP network?

```
R2#conf t
ip default-network 2.0.0.0

R3#show ip route
Gateway of last resort is 2.0.0.2 to network 0.0.0.0
R* 0.0.0.0/0 [120/1] via 2.0.0.2, 00:00:01, Serial1/2
```

Describe **ip default-network x.x.x.x** in a RIP domain:



increase the capable number of NOT yet processed RIP update packets to 120 on a RIP router:

```
router rip
input-queue 120
```

Help me create more flashcards:

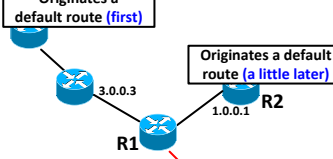
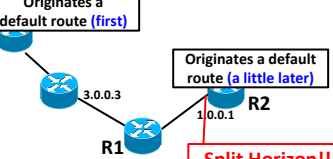
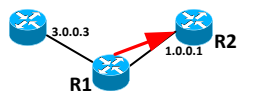
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts


Colin

<p>Injecting a default route in RIP</p>	<p>OR case: <pre>router rip default-information originate route-map DEFAULT ip prefix-list PFX-1 permit 1.0.0.0/8 ip prefix-list PFX-1 permit 2.0.0.0/8 route-map RMP-DEFAULT match ip address prefix PFX-1</pre> </p>				
<p>If there are two networks which either OR have to be up</p>	<p>track 9 list boolean and object 15 object 16</p>				
<p>If there are two networks which both (AND) have to be up</p>	<p>AND case: track 15 interface loopback 15 track 16 interface loopback 16</p>				
<p>In order to generate the default route:</p>	<p>ip route 169.254.0.0 255.255.255.0 null 0 track 90 router rip default-information originate route-map DEFAULT ip prefix-list DUMMY permit 169.254.0.0/24 route-map RMP-DEFAULT match ip address prefix DUMMY</p>				
<p>Originates a default route (first)</p>  <p>show ip route i 0.0.0.0 R* 0.0.0.0/0 [120/2] via 3.0.0.3 fa0/0</p> <p>Which default route will R1 use?</p>	<p>Originates a default route (first)</p>  <p>R2 will NOT send its default route, because its already receiving a default route through the same interface its about to send its default route out! -> split horizon prevents this</p> <p>In order for R2 to send a default route, deny the incoming default route via distribute-list, so R1 will use R2 due to shorter hop-count!</p>				
<p>Advertise a default route in RIP if loopback 9 is up and loopback 99 is down:</p>	<pre>int loopback 9 ip address 9.9.9.9 255.255.255.0 int loopback 99 ip address 99.99.99.99 255.255.255.0 track 9 interface Loopback9 line-protocol track 99 interface Loopback99 line-protocol track 20 list boolean and object 9 object 99 not router rip network 0.0.0.0 ip route 0.0.0.0 0.0.0.0 Null0 track 20</pre>				
<p>How to match a default route with standard ACLs, extended ACLs, and a prefix list:</p>	<p>Deny the default route, allow everything else</p> <p>Standard access-list: access-list 22 deny host 0.0.0.0 access-list 22 permit any</p> <p>Extended access-list: access-list 102 deny host 0.0.0.0 host 0.0.0.0 access-list 102 permit any any</p> <p>Prefix-list: ip prefix-list DEFAULT deny 0.0.0.0/0 ip prefix-list DEFAULT permit 0.0.0.0/0 le 32</p>				
<p>Configuring RIP authentication before configuring the key chain can cause problems that no RIP routes are being passed on. Can these problems "survive" a reload?</p>	<p>The answer is YES!!</p> <p>The only way on fixing this is:</p> <pre>Int ser0/x No ip rip authentication mode text No ip rip authentication key-chain XXX Key-chain BLA Key 1 Key-string blabla Int ser0/x Ip rip authentication mode text Ip rip authentication key-chain BLA</pre>				
 <p>Explain two options on how to send a default route only on the link towards R2 in RIP</p>	<pre>int fa0/x ip summary-address rip 0.0.0.0 0.0.0.0 router rip default-information originate route-map rMP-DEF-fa0/x route-map RMP-DEF-fa0/x permit 10 set interface fa0/x</pre>				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin



EIGRP

<p>show ip eigrp neighbors:</p> <p>What does one have to watch out for?</p>	 <pre>R3#show ip eigrp neighbors IP-EIGRP neighbors for process 100 H Address Interface Hold Uptime SRTT RTO Q Seq (sec) (ms) 2 155.1.37.7 Fa0/0 12 00:05:22 6 200 0 27 1 155.1.0.5 Se1/0/1 139 00:07:12 39 1140 0 124 0 155.1.13.1 Se1/2 11 00:07:51 20 1140 0 104</pre> <p>Q Cnt Num should be 0, meaning there are no packets in the queue. No congestion present.</p>	<h2>EIGRP Unicast Updates / Neighbor</h2>	<pre>router eigrp 100 neighbor 155.1.58.8 Fa0/0</pre>	<h2>EIGRP Filtering with Prefix-Lists</h2> <p>(not from gateway)</p>	<pre>router eigrp 100 distribute-list prefix PERMIT_ALL gateway NOT_FROM_R4 in ip prefix-list NOT_FROM_R4 seq 5 deny 155.1.146.4/32 ip prefix-list NOT_FROM_R4 seq 10 permit 0.0.0.0/0 le 32 ip prefix-list PERMIT_ALL seq 5 permit 0.0.0.0/0 le 32</pre>
<h2>How to disable EIGRP Split Horizon</h2>	<pre>interface Serial0/0/0 no ip split-horizon eigrp 100</pre> <p>Verify using:</p> <pre>show ip eigrp interface</pre>	<p>Debugging EIGRP traffic using "debug ip packet detail":</p>	<pre>EIGRP Using static neighborship statements: R5#debug ip packet detail IP packet debugging is on (detailed) IP: s=155.1.58.5 (local), d=155.1.58.8 (Fa0/0), len 60, sending, proto=88 EIGRP Using multicast: R5#debug ip packet detail IP packet debugging is on (detailed) IP: s=155.1.58.5 (local), d=224.0.0.10 (Fa0/0), len 60, sending broad/multicast, proto=88</pre>	<h2>EIGRP Filtering with Standard Access-Lists</h2> <p>(allowing all EVEN prefixes, disallowing all ODD prefixes)</p>	<pre>router eigrp 10 distribute-list 1 in Serial0/0/0 access-list 1 permit 0.0.0.0 255.255.254.255</pre>
<p>Explain the EIGRP feasibility condition?</p>	<ul style="list-style-type: none"> - end-to-end composite metric is compared between routes. - If the Advertised Distance is equal to the Feasible Distance, it is considered an alternative path. - if the Advertised Distance is smaller than the Feasible Distance, it is considered as a new Feasible successor. 	<h2>EIGRP Default Network</h2>	<pre>router eigrp 10 network 200.0.0.0 ! ip default-network 200.0.0.0 (Prefix 200.0.0.0 was received from a EIGRP neighbor) R5# show ip route eigrp i EX D*EX 200.0.0.0/24 [170/2812416] via 155.1.0.1, 00:00:52, Ser0/0/0 R5# sh ip route ia - IS-IS inter area, * - candidate default, U - per-user static route o - ODR, P - periodic downloaded static route Gateway of last resort is 155.1.0.1 to network 200.0.0.0</pre>	<h2>EIGRP Filtering with Extended Access-Lists</h2>	<pre>access-list 100 deny ip host 155.1.0.2 host 150.1.2.0 access-list 100 deny ip host 155.1.0.4 host 150.1.2.0 access-list 100 permit ip any any router eigrp 100 distribute-list 100 in Serial0/0/0 access-list 100 deny ip host 155.1.0.2 host 150.1.2.0 Where 155.1.0.2 is the next hop 150.1.2.0 is the prefix.</pre>
<h2>EIGRP MD5 Authentication</h2>	<pre>key chain MD5_KEYS key 1 key-string CISCO interface Serial0/0/0 ip authentication mode eigrp 10 md5 ip authentication key-chain eigrp 10 MD5_KEYS</pre>	<p>How to identify the default network with a show ip route output:</p>	<pre>R1# ip default-network 200.0.0.0 R3#sh ip route 200.0.0.0 Routing entry for 200.0.0.0/24 Known via "eigrp 10", distance 170, metric 2300, candidate default path, type external Redistributing via eigrp 100 Last update from 155.1.37.7 on FastEthernet0/0, 00:07:47 ago Routing Descriptor Blocks: * 155.1.37.7, from 155.1.37.7, 00:07:47 ago, via FastEthernet0/0 Route metric is 2300/2, traffic share count is 1 Total delay is 25110 microseconds, minimum bandwidth is 1544 Kbit Reliability 255/255, minimum MTU 1500 bytes Loading 1/255, Hops 3</pre>	<h2>EIGRP Filtering with Offset Lists</h2>	<pre>router eigrp 100 offset-list 99 in 2147483647 Fa0/3 access-list 99 permit 150.1.3.0 modify the metric on a per route basis or a per-interface basis, setting the maximum value inbound for that prefix.</pre>
<p>What are usefull EIGRP Debug commands:</p>	<pre>debug eigrp packets hello debug eigrp fsm</pre>	<p>Configuring EIGRP summary addresses</p>	<pre>Interface fa0/0 ip summary-address eigrp 100 30.0.0.0 0.0.0.0 EIGRP AS: 100</pre>	<h2>EIGRP Filtering with Administrative Distance</h2> <p>Per prefix, any neighbour:</p>	<pre>access-list 4 permit 150.1.4.0 router eigrp 100 distance 255 0.0.0.0 255.255.255.255 4 0.0.0.0 255.255.255.255 specify from any neighbour. Setting the Admin Distance to UNKNOWN 255 from any neighbour for prefix 150.1.4.0/xx</pre>
<h2>Key Chain Rotation</h2>	<pre>key chain KEY_ROTATION key 10 key-string CISCO10 accept-lifetime 00:00:00 Jan 1 1993 00:15:00 Jan 1 2030 send-lifetime 00:00:00 Jan 1 1993 00:05:00 Jan 1 2030 key 20 key-string CISCO20 accept-lifetime 00:00:00 Jan 1 2030 infinite send-lifetime 00:00:00 Jan 1 2030 infinite ! Make sure the clocks are synchronized either by "clock set" or NTP !</pre>	<p>Show commands with special includes:</p>	<pre>show ip route include via 155.1.(0 45).4 show ip route include 30\.\ 31\. show ip route include 3(0 1).[0-3].0.0</pre>	<h2>EIGRP Filtering with Per Neighbor AD</h2> <p>Per prefix, per neighbour:</p>	<pre>access-list 7 permit 150.1.7.0 router eigrp 100 distance 255 155.1.37.7 0.0.0.0 7 Route: 150.1.7.0/24 Next-Hop: 155.1.37.7 Making the route unreachable outbound to Host 155.1.37.7.</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin


EIGRP

<p>EIGRP Summarization sending only a Default Route:</p>	<pre>interface Ser0/0/0.1 point-to-point ip summary-address eigrp 100 0.0.0.0 0.0.0.0 5</pre> <p>Everything else will be suppressed out this interface, only the Default route will be advertised out Ser0/0/0.1.</p>	<p>Checking the usage of the EIGRP Metric Weights / K-Values using:</p> <p>show ip protocols</p>	<pre>Router# show ip protocols Routing Protocol is "eigrp 100" Outgoing update filter list for all interfaces is not set Incoming update filter list for all interfaces is not set Default networks flagged in outgoing updates EIGRP metric weights K1=1, K2=0, K3=1, K4=0, K5=0 EIGRP maximum hopcount 100 EIGRP maximum metric variance 1 Redistributing: eigrp 100 EIGRP NSF-aware route hold timer is 240s Automatic network summarization is not in effect Maximum path: 4 Routing for Networks: 150.1.0.0 155.1.0.0 Routing Information Sources: Gateway Distance Last Update 155.1.0.2 90 00:15:49 155.1.45.4 90 00:15:06 Distance: internal 90 external 170</pre>	<p>Configure EIGRP so that if a query to a missing route is stuck in Active will be declared dead after one minute:</p>	<pre>router eigrp 100 timers active-time 1</pre> <p>(Stuck in Active SIA) Declare route down if no answer was received after one minute.</p>
<p>EIGRP Summarization with Leak Map</p>	<pre>interface Serial0/1/0 ip summary-address eigrp 100 0.0.0.0 0.0.0.0 5 leak-map LEAK_LOOPBACK0 route-map LEAK_LOOPBACK0 permit 10 match ip address prefix-list LOOPBACK0 ip prefix-list LOOPBACK0 seq 5 permit 150.1.4.0/24</pre> <p>This will send out only a default route and the leaked out Loopback0 prefix of 150.1.4.0/24</p> <p>(subnets will be advertised / leaked in addition to the summary)</p>	<p>Checking the values of EIGRP Metric Weights / K-Values using the:</p> <p>show ip eigrp topology PREFIX MASK</p>	<pre>SW3#show ip eigrp topology 150.1.9.0 255.255.255.0 IP-EIGRP (AS 100): Topology entry for 150.1.9.0/24 State is Passive, Query origin flag is 1, 1 Successor(s), FD is 128000 Routing Descriptor Blocks: 0.0.0.0 (Loopback0), from Connected, Send flag is 0x0 Composite metric is (128000/0), Route is Internal Vector metric: Minimum bandwidth is 10000000 Kbit Total delay is 5000 microseconds Reliability is 255/255 Load is 1/255 Minimum MTU is 1514 Hop count is 0</pre> <p>RD FD</p>	<p>Configure EIGRP so that SIA condition can remain indefinitely</p>	<p>router eigrp 100 timers active-time disable</p> <p>Disables the timers and permits the routing wait time to remain active indefinitely.</p>
<p>EIGRP Floating Summarization</p> <p>Generates summary with route to Null0, having more specific route:</p>	<pre>interface FastEthernet0/0 ip summary-address eigrp 100 150.1.4.0 255.255.254.0 5 ! ip route 150.1.4.0 255.255.254.0 155.1.0.4</pre>	<p>Using solely EIGRP's Delay K-Value, after performed changes, what needs to be done?</p>	<pre>conf t interface fa0/0 delay 999999 end clear ip eigrp neighbours</pre>	<p>Which debug command would you use to troubleshoot EIGRP Stuck In Active situations?</p>	<p>debug eigrp packet terse</p> <pre>EIGRP: Enqueueing QUERY on Port-channel1 EIGRP: Received QUERY on Vlan79 nbr 155.1.79.7 EIGRP: Enqueueing ACK on Vlan79 nbr 155.1.79.7</pre>
<p>EIGRP Poisoned Floating Summarization</p> <p>Propagating a summary out, but locally deny of using it via AD:</p>	<pre>interface GigabitEthernet0/0 ip summary-address eigrp 100 150.1.4.0 255.255.254.0 255</pre> <p>Poisoning the route via setting the Admin Distance to UNKNOWN, prevents the summary to be announced out Gi0/0</p>	<p>Feasible Distance:</p> <p>Reported Distance:</p> <p>Advertised Distance:</p>	<pre>155.1.67.6 (Vlan67), from 155.1.67.6, Send flag is 0x0 Composite metric is (25728000/128000), Route is Internal (Feasible Distance FD / Advertised Distance AD) Advertised Distance is the distance received from the neighbour. Feasible Distance is the routers distance to the network, once added the local interface "cost" / metric to the path.</pre>	<p>What for does one use EIGRP stub routing?</p>	<p>EIGRP stub is used to limit the query messages, limiting SIA conditions.</p> <pre>router eigrp 100 eigrp stub connected</pre> <pre>R5#show ip eigrp neighbors detail IP-EIGRP neighbors for process 100 H Address Interface Hold Uptime SRTT RTO Q Seq (sec) (ms) Cnt Num S 155.1.58.8 Gi0/0 29 00:00:28 8 200 0 150 Version 12.2/3.0, Retrans: 0, Retries: 0, Prefixes: 3 Stub Peer Advertising (CONNECTED) Routes Suppressing queries</pre>
<p>Changing EIGRP Metric Weights</p> <p>(use of K-Values)</p>	<pre>conf t router eigrp 100 metric weights 0 0 1 0 0 metric weights 0 K1 K2 K3 K4 K5 K1 bandwidth, K2 load, K3 delay, K4 reliability, K5 MTU</pre> <p>VERIFY using used K Values via show ip protocols</p>	<p>How can one identify the effect of EIGRP variance of an IP route?</p>	<pre>R6#show ip route 155.1.9.9 Routing entry for 155.1.9.0/24 Known via "eigrp 100", distance 90, metric 3072, type internal Redistributing via eigrp 10, eigrp 100 Advertised by eigrp 10 Last update from 155.1.146.1 on GigabitEthernet0/0.146, 00:00:43 ago Routing Descriptor Blocks: 155.1.146.1, from 155.1.146.1, 00:00:43 ago, via GigabitEthernet0/0.146 Route metric is 15360, traffic share count is 1 Total delay is 600 microseconds, minimum bandwidth is 1544 Kbit Reliability 255/255, minimum MTU 1500 bytes Loading 1/255, Hops 4 * 155.1.67.7, from 155.1.67.7, 00:00:43 ago, via GigabitEthernet0/0.67 Route metric is 3072, traffic share count is 5 Total delay is 120 microseconds, minimum bandwidth is 100000 Kbit Reliability 255/255, minimum MTU 1500 bytes Loading 1/255, Hops 2</pre> <pre>conf t int fa0/x, fa0/z no route-cache ip load-sharing per-packet USE ACLs to count ICMPs on the opposite side, to count packets to analyse traffic share.</pre>	<p>EIGRP Stub Routing with Leak Map</p>	<pre>ip prefix-list SW2_LOOPBACK seq 5 permit 150.1.8.0/24 ! route-map STUB_LEAK_MAP deny 10 match ip address prefix-list SW2_LOOPBACK ! route-map STUB_LEAK_MAP permit 20 ! router eigrp 100 eigrp stub connected leak-map STUB_LEAK_MAP</pre> <p>Leaks out all routes except of 150.1.8.0/24.</p>
<p>EIGRP composite metric calculation formula:</p>	$\text{metric} = [k1 * \text{bandwidth} + (k2 * \text{bandwidth}) / (256 - \text{load}) + k3 * \text{delay}] * [k5 / (\text{reliability} + k4)]$ <p>Breaks down to (metric weights)</p> <p>256 (BW + DLY)</p>	<p>EIGRP Convergence Timers:</p>	<pre>(eigrp AS 100) interface FastEthernet0/0 ip hello-interval eigrp 100 1 ip hold-time eigrp 100 3</pre> <p>in seconds</p>	<p>Explain the passive-interface command in EIGRP:</p>	<p>will stop the forming of an adjacency on the interface, and hence the learning of any updates on the link.</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

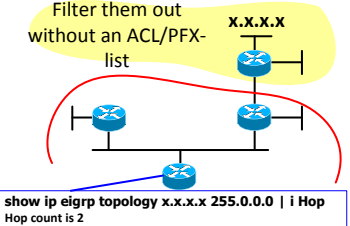
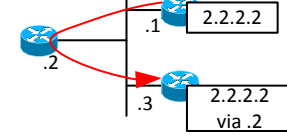
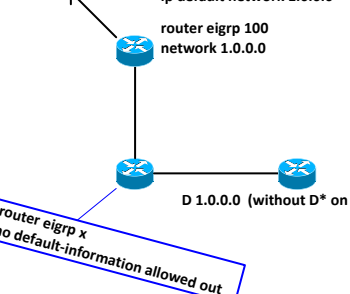
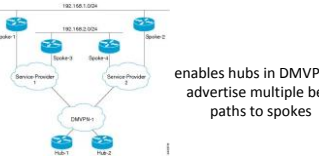
Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin



EIGRP

<p>EIGRP Filtering with Per Neighbor AD</p> <p>Differences between Internal EIGRP routes and External ones:</p>	<p>The administrative distance for EIGRP internal routes can be changed on a per prefix basis, but external EIGRP routes cannot.</p>	<p>EIGRP Router-ID</p>	<pre>router eigrp 100 eigrp router-id 150.1.2.2</pre> <p>Can be used to filter out external EIGRP routes of a different router, by using his Router-ID, the local Router will ignore those routing updates.</p>	<p>EIGRP BFD</p>	<pre>Interface e0/0 bfd interval 50 min_rx 50 multiplier 3 router eigrp 100 bfd interface e0/0 bfd all-interfaces</pre>										
<p>EIGRP Filtering with Route Maps</p>	<pre>router eigrp 100 distribute-list route-map FILTER_ON in ! route-map FILTER_ON deny 10 match tag 4 ! route-map FILTER_ON permit 20</pre>	<p>EIGRP Maximum Hops</p>	<pre>router eigrp 100 metric maximum-hops 1</pre> 	<p>Interface X</p> <pre>no ipv6 next-hop-self eigrp 55</pre>	<p>by default, the IPv6 next-hop value is set to be itself for routes that it is advertising, even when advertising those routes back out the same interface where it learned them</p> <pre>interface type number no ipv6 next-hop-self eigrp as-number</pre> 										
<p>Configure a route-map filter for EIGRP that filters routes with a metric that ranges within 500000 – 700000 from entering the routing table:</p> <p><i>metric +/-</i></p>	<pre>router eigrp 100 distribute-list route-map METRIC_RANGE in route-map METRIC_RANGE deny 10 match metric 625000 +- 125000 route-map METRIC_RANGE permit 20 (METRIC 625K +/- 125K = 500K or 750K)</pre>	<p>EIGRP Default routes</p> <p>3 options</p>	<pre>ip route 0.0.0.0 0.0.0.0 Null0 A) network 0.0.0.0 B) redistribute static into EIGRP ip summary-address eigrp 100 0.0.0.0 0.0.0.0 ip default network 99.99.0.0 router eigrp 100 network 99.99.0.0</pre>	<p>Route-map METRIC</p> <p>match different metrics</p> <p>metric +/-</p>	<pre>match routes with a metric of 110, or 200, or a metric within a range of 700 to 800 (750 plus / minus 50 = low 700 high 800)</pre> <pre>route-map METRIC match metric 110 200 750 +/- 50 router eigrp 1 redistribute ospf route-map METRIC</pre>										
<p>EIGRP Bandwidth Pacing</p> <p>Make sure EIGRP does not use more than 154Kbit of Bandwidth on the link:</p> <p>(10% of 1544 is 154Kbit)</p>	<pre>interface Serial0/1 bandwidth 1544 ip bandwidth-percent eigrp 100 10</pre>	<p>EIGRP</p> <p>no default-information allowed out</p>	<pre>ip default network 1.0.0.0 router eigrp 100 network 1.0.0.0</pre> 	<p>EIGRP named mode</p> <p>Add path support</p> 	<p>Add Path Support on a Hub:</p> <pre>router eigrp name address-family ipv4 autonomous-system 10 af-interface tunnel 0 no next-hop-self no-ecmp-mode add-paths 4</pre> <p>Add Path Support on Spoke:</p> <pre>router eigrp name address-family ipv6 autonomous-system 10 af-interface tunnel 0 no next-hop-self no-ecmp-mode add-paths 4</pre> <p><i>Add Path only with EIGRP named mode!</i></p>										
<p>EIGRP Default Metric</p> <p>Redistributing a static route, changing default metrics:</p>	<pre>ip route 222.22.2.2 255.255.255.255 192.10.1.254 router eigrp 100 redistribute static default-metric 100000 10 255 1 1500</pre>	<p>EIGRP</p> <p>default-information</p>	<pre>default-information allowed { in out } default-information { in out } [acl] no default-information ...</pre> <table border="1"> <tr> <td>allowed</td> <td>Configures EIGRP to accept default routing information.</td> </tr> <tr> <td>in</td> <td>Configures EIGRP to accept exterior or default routing information.</td> </tr> <tr> <td>out</td> <td>Configures EIGRP to advertise external routing information.</td> </tr> <tr> <td>ac-number</td> <td>(Optional) Standard access list number from 1 to 99 or an expanded standard access list from 1300 to 1999.</td> </tr> <tr> <td>ac-name</td> <td>(Optional) Named standard access list.</td> </tr> </table> <p>ACL specifies for which prefix the default D* prefix is allowed to be allowed in/out</p>	allowed	Configures EIGRP to accept default routing information.	in	Configures EIGRP to accept exterior or default routing information.	out	Configures EIGRP to advertise external routing information.	ac-number	(Optional) Standard access list number from 1 to 99 or an expanded standard access list from 1300 to 1999.	ac-name	(Optional) Named standard access list.	<p>EIGRP Stub possibilities</p>	<pre>router eigrp 100 eigrp stub eigrp stub receive-only eigrp stub leak-map RMP-NAME eigrp stub connected eigrp stub static eigrp stub summary eigrp stub redistributed</pre> <p>connected and summary Nothing sent identifies suppressed PFXs Connected routes static routes, no summaries Allows summaries redistributes other proto's</p> <pre>router eigrp NAME address-family ipv4 vrf X eigrp stub connected</pre>
allowed	Configures EIGRP to accept default routing information.														
in	Configures EIGRP to accept exterior or default routing information.														
out	Configures EIGRP to advertise external routing information.														
ac-number	(Optional) Standard access list number from 1 to 99 or an expanded standard access list from 1300 to 1999.														
ac-name	(Optional) Named standard access list.														
<p>EIGRP Neighbor Logging</p>	<pre>router eigrp 100 no eigrp log-neighbor-changes eigrp log-neighbor-warnings 20</pre> <p>Only create a warning message every 20 seconds.</p>	<p>EIGRP</p> <p>neighbor maximum prefix</p> <p>per vrf / per neighbor</p>	<pre>router eigrp 55 address-family ipv4 vrf X autonomous-system 66 neighbor 1.2.3.4 description MY-DESCRIPTION neighbor 1.2.3.4 maximum-prefix 600 <threshold> [warning-only] neighbor 1.2.3.4 maximum-prefix 600 <threshold> reset-time <min> restart <min> ...</pre> <pre>router eigrp INSTANCE-NAME address-family ipv4 vrf X autonomous-system 66 neighbor 1.2.3.4 description MY-DESCRIPTION neighbor 1.2.3.4 maximum-prefix 600 <threshold> [warning-only]</pre>	<p>EIGRP</p> <p>Redistribute maximum-prefix</p> <p>Autonomous-System</p> <p>Named-Mode</p>	<pre>router eigrp 100 address-family ipv4 [unicast] vrf vrf-name redistribute maximum-prefix maximum</pre> <pre>router eigrp INSTANCE-NAME address-family ipv4 vrf X autonomous-system 66 topology (base topology-name tid number) redistribute maximum-prefix maximum</pre>										

Help me create more flashcards:

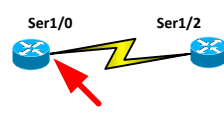



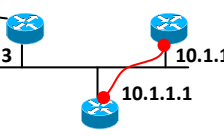
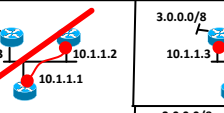

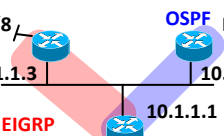
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin


<p>EIGRP Route Tag Enhancements</p>	<pre>conf t route-tag notation dotted-decimal range: 1 to 4294967295 range: 0.0.0.0 to 255.255.255.255 eigrp default-route-tag <123 or 0.0.1.2> for internal routes only route-map X permit match tag 1235 match tag 0.0.12.99 match tag list TAG-LIST1 conf t route-tag list NAME permit 1235 8879 show route-tag list</pre> <p><i>route-tag list TAG-UST1 permit 2.2.2.1 0.0.0.1 tag nullcard</i></p>	<p>EIGRP</p> <p>Nonstop Forwarding (NSF) Awareness</p> <p>Graceful-restart</p>	<p>Sets the route-hold timer to determine how long an NSF-aware router that is running EIGRP will hold routes for an inactive peer.</p> <p>timers graceful-restart purge-time <seconds></p> <p>Verify using: <code>show ip protocols</code></p> <p>debug ip eigrp notifications</p> <p>EIGRP:NSF:AS2. Rec RS update from x.x. Wait for EOT. UAL-5-NBRCHANGE:IP-EIGRP(0) 2:Neighbor x.x (fa3/0) is up:peer NSF restarted EIGRP-IPV4 100: Neighbor x.x (fa1/0) is resync: peer graceful-restart</p>	<p>What will R1 have in its routing table in regards to the 1.0.0.0 networks?</p> <p>R2# router eigrp 100 network 1.1.0.0 0.0.255.255 interface Serial1/2 ip summary-address eigrp 100 1.1.0.0 255.255.252.0 router eigrp 100 eigrp stub connected</p> <p>Loopbacks: 1.1.0.0/24 Ser1/1 1.1.1.0/24 Ser1/2 1.1.2.0/24 Ser1/2 1.1.3.0/24 R1</p>	<p>Loopbacks: 1.1.0.0/24 Ser1/1 1.1.1.0/24 Ser1/2 1.1.2.0/24 Ser1/2 1.1.3.0/24 R1</p> <p>R1#show ip route eigrp 1.0.0.0/24 is subnetted, 4 subnets D 1.1.0.0 [90/206400] via 12.1.1.1, 00:02:08, Serial1/1 D 1.1.1.0 [90/206400] via 12.1.1.1, 00:02:08, Serial1/1 D 1.1.2.0 [90/206400] via 12.1.1.1, 00:02:08, Serial1/1 D 1.1.3.0 [90/206400] via 12.1.1.1, 00:02:08, Serial1/1</p> <p>THE INDIVIDUAL ROUTES SEEN!! NO SUMMARY, even it is configured!</p>
<p>EIGRP Wide Metrics</p>	<p>- supports 64-bit metric calculations - (classic mode configurations use 32-bit calculations) - > 1 Gbit and up to 4.2 terabits EIGRP path selections</p> <p>metric rib-scale</p> <p>router eigrp name1 address-family ipv4 autonomous-system 4533 metric weights 0 2 0 1 0 0 1</p> <p>K1 Bandwidth set to 2</p>	<p>EIGRP Loop-Free Alternate Fast Reroute</p> <p>IP FRR/fast reroute (single hop)</p>	<p>router eigrp NAME address-family ipv4 autonomous-system 88 topology base fast-reroute per-prefix all fast-reroute per prefix route-map RMP-NAME</p> <p>Disabling Load sharing among prefixes: fast-reroute load-sharing disable</p> <p>Enabling Tie breaking rules: fast-reroute tie-break [interface-disjoint linecard-disjoint lowest-backup-path-metric srlg-disjoint] priority-number</p> <p>show ip eigrp topology fr</p>	<p>Configuring EIGRP stub connected In Classic Mode</p>	<p>router eigrp 100 eigrp stub connected</p>
<p>Named EIGRP Authentication md5</p>	<p>router eigrp NAME address-family ipv4 autonomous-system 258 af-interface [default] fa0/x af-interface fa0/0 authentication key-chain KEY-CHAIN authentication mode md5</p>	<p>Redistributing static and connected routes into router EIGRP:</p>	<p>router eigrp 100 redistribute static</p> <p>ip route 3.0.0.0 255.0.0.0 Null0</p> <p>When redistributing connected or static routes into EIGRP NO seed metric is required!</p>	<p>What will R1 have in its routing table in regards to the 1.0.0.0 networks?</p> <p>R2# router eigrp 100 network 1.1.0.0 0.0.255.255 interface Serial1/2 ip summary-address eigrp 100 1.1.0.0 255.255.252.0 router eigrp 100 eigrp stub summary</p> <p>Loopbacks: 1.1.0.0/24 Ser1/1 1.1.1.0/24 Ser1/2 1.1.2.0/24 R2 1.1.3.0/24 R1</p>	<p>Loopbacks: 1.1.0.0/24 Ser1/1 1.1.1.0/24 Ser1/2 1.1.2.0/24 R2 1.1.3.0/24 R1</p> <p>R1#show ip route eigrp 1.0.0.0/22 is subnetted, 1 subnets D 1.1.0.0 [90/206400] via 12.1.1.1, 00:00:12, Serial1/1</p>
<p>Named EIGRP Authentication HMAC-SHA-256</p>	<p>router eigrp NAME address-family ipv4 autonomous-system 258 af-interface [default] fa0/x af-interface default authentication mode hmac-sha-256 0 PASSWORD</p>			<p>Loopbacks: 1.1.0.0/24 Ser1/1 1.1.1.0/24 Ser1/2 1.1.2.0/24 R2 1.1.3.0/24 R1</p>	<p>R2#show ip route eigrp 1.0.0.0/8 is variably subnetted, 9 subnets, 3 masks D 1.1.0.0/22 is a summary, 00:56:55, Null0</p>
<p>EIGRP over the top Background info:</p>	<p>EIGRP on the control plane and Locator ID Separation Protocol (LISP) encapsulation</p> <p>- must configure the neighbor command with LISP encapsulation on every CE</p> <p>- having many CE, could use EIGRP Route Reflectors (E-RRs)</p> <p>E-RR: remote-neighbors source to listen to unicast messages from peer Ces</p>	<p>Make sure EIGRP uses max 520 kbit of bandwidth, do NOT use ip bandwidth-percent eigrp x 520 or any service-policy applied to the interface:</p> 	<p>R1# conf t interface serial 1/0 bandwidth 1040</p> <p>By default EIGRP will use up to 50% of the links bandwidth for EIGRP.</p> <p>bandwidth 1040 / 2 = 520 kbit/s</p>	<p>What will R1 have in its routing table ?</p> <p>R2# router eigrp 100 network 1.1.0.0 0.0.255.255 interface Serial1/2 ip summary-address eigrp 100 1.1.0.0 255.255.252.0 router eigrp 100 redistribute static redistribute rip metric 1 1 1 1 eigrp stub redistributed</p> <p>Static route 11.0.0.0/8 Loopbacks: 1.1.0.0/24 Ser1/1 1.1.1.0/24 Ser1/2 1.1.3.0/24 R2</p> 	<p>Static route 11.0.0.0/8 RIP prefix 200.1.1.0/24 Loopbacks: 1.1.0.0/24 Ser1/1 1.1.1.0/24 Ser1/2 1.1.3.0/24 R2</p> <p>R2# router eigrp 100 eigrp stub redistributed</p> <p>R1#show ip route eigrp D EX 11.0.0.0/8 [170/2051] via 12.1.1.1, 00:00:36, Ser1/1 D EX 200.1.1.0/24 [170/2560] via 12.1.1.1, 00:00:36, Ser1/1</p>
<p>EIGRP over the top E-RR</p>	<p>EIGRP over the top on CE</p> <pre>router eigrp virtual-name address-family ipv4 autonomous-system 65 neighbor 1.2.3.4 fa0/x remote <max-hops> lisp-encap <lisp-id></pre> <p>EIGRP over the top Route Reflectors E-RR</p> <pre>router eigrp NAME address-family ipv4 unicast autonomous-system 55 af-interface fa0/0 no next-hop-self no split-horizon ! remote-neighbors source fa0/x unicast-listen lisp-encap 1</pre>	<p>EIGRP Summary-address (classic, old fashioned)</p> <hr/> <p>EIGRP Named Mode Summary-address</p> 	<p>router eigrp 100 network 1.1.0.0 0.0.255.255 interface Serial1/2 ip summary-address eigrp 100 1.1.0.0 255.255.252.0</p> <p>OLD</p>  <p>NEW</p> <pre>router eigrp NAME address-family ipv4 unicast autonomous-system 100 af-interface Serial1/1 summary-address 2.2.0.0 255.255.252.0 exit-af-interface topology base exit-af-topology network 2.2.0.0 0.0.255.255</pre>	<p>Will 3.0.0.0/8 be seen by the other two routers in this EIGRP network?</p>  <p>R1# router eigrp NAME address-family ipv4 autonomous-system 100 neighbor 10.1.1.2 fa0/x</p>	<p>EIGRP a neighbor command is used on an Ethernet Broadcast medium, EIGRP will no longer use broadcast for other still broadcast enabled routers on the same segment!!</p>  <p>Will not work</p>  <p>works</p> <p>mcast</p>
<p>EIGRP Classic to Named Mode Conversion</p> <p>EIGRP conversion</p>	<p>Convert from EIGRP classic mode to Named mode, BUT NOT backwards from Named to Classic!!</p> <p>eigrp upgrade-cli</p> <p>write memory</p> <p>From:</p> <pre>router eigrp 6524</pre> <p>to</p> <pre>router eigrp NAME</pre>	<p>How to redistribute static routes</p> <p>In EIGRP "classic" mode:</p> <pre>ip route X.0.0.0 255.0.0.0 FastEthernet0/0</pre> <p>Classic</p> <pre>router eigrp 100 redistribute static</pre> <p>In EIGRP Named Mode:</p> <pre>router eigrp NAME address-family ipv4 unicast autonomous-system 100 topology base redistribute static</pre> <p>Named</p>	<p>ip route X.0.0.0 255.0.0.0 FastEthernet0/0</p> <p>router eigrp 100 redistribute static</p> <p>ip route X.0.0.0 255.0.0.0 FastEthernet0/0</p> <p>router eigrp NAME address-family ipv4 unicast autonomous-system 100 topology base redistribute static</p>	<p>Redistribute OSPF into EIGRP so that R2 see's 3.0.0.0/8 with a next hop of 10.1.1.3</p>  <p>Perform EIGRP "classic config" and EIGRP named mode</p>	<p>R1# router ospf 1 network 10.1.1.1 0.0.0.0 area 0 router eigrp 100 network 10.1.1.1 0.0.0.0 redistribute ospf 1 metric 1 1 1 1 int fa0/0 no ip next-hop-self eigrp 100</p> <p>R2#show ip route 3.0.0.0 D 3.0.0.0 via 10.1.1.3</p> <p>R3#show ip route 3.0.0.0 D 3.0.0.0 via 10.1.1.3</p> <p>router ospf 1 network 10.1.1.1 0.0.0.0 area 0 router eigrp TST address-family ipv4 unicast autonomous-system 200 af-interface FastEthernet0/1 no next-hop-self exit-af-interface topology base redistribute ospf 1 metric 1 1 1 1</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

How to configure / change hold-down and hello interval in EIGRP

“classic mode”

```
router eigrp 100
network 10.1.1.2 0.0.0.0

interface FastEthernet0/1
ip address 10.1.1.2 255.255.255.0
ip hello-interval eigrp 200 10
ip hold-time eigrp 200 30
```

Verify:
show ip eigrp 100 interfaces detail fa0/1 | time

And

Named Mode:

```
router eigrp TST
address-family ipv4 unicast autonomous-system 200
af-interface FastEthernet0/1
hello-interval 10
hold-time 30
```

How to configure the router to never exceed more than 10% of any links bandwidth for EIGRP in named mode:

```
router eigrp NAME
address-family ipv4 unicast autonomous-system 100
af-interface default
bandwidth-percent 10
```

EIGRP offset-lists Configuration:

“classic mode”

```
access-list 99 permit 1.1.0.0 0.0.255.255

router eigrp A-100
address-family ipv4 unicast autonomous-system 100
topology base
offset-list 99 out 8888 Serial1/3
```

Named mode:

```
access-list 99 permit 1.1.0.0 0.0.255.255

router eigrp 10
offset-list 99 out 8888 FastEthernet0/17
```

What is the difference between the following:

RIP:
router rip
passive-interface fa0/x
neighbor 10.0.0.1 fa0/x

EIGRP
router eigrp 100
passive-interface fa0/x
neighbor 10.0.0.1 fa0/x

RIP:
router rip
passive-interface fa0/x
neighbor 10.0.0.1 fa0/x

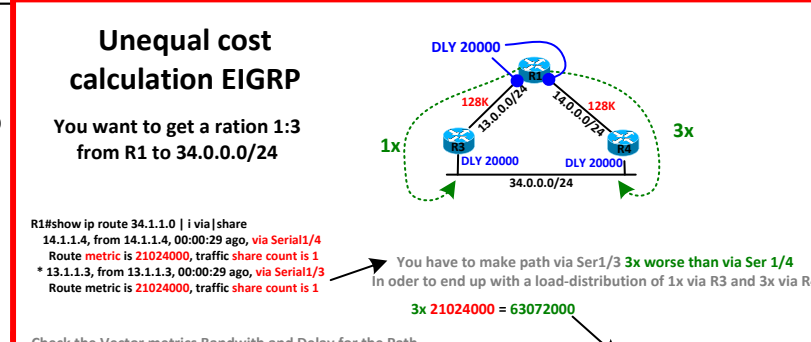
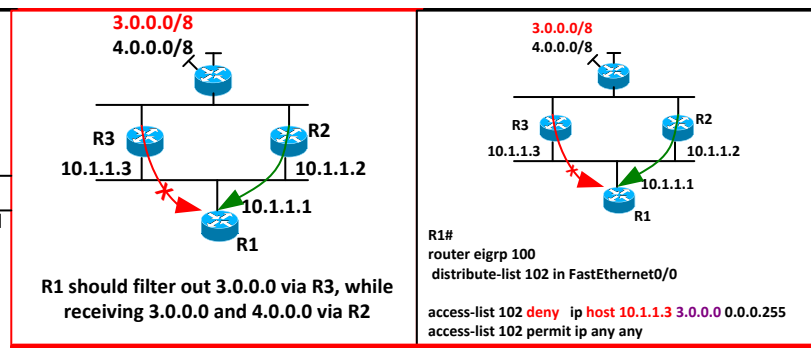
EIGRP
router eigrp 100
passive-interface fa0/x
neighbor 10.0.0.1 fa0/x

Works

WILL NOT WORK
Session is teared down

What parameter have to match in order to have EIGRP establish an adjacency?

- K-Values
- Autonomous System Number
- Authentication



Check the Vector metrics Bandwidth and Delay for the Path from R1 to 34.1.1.0/24 and extract the two values per path

R1#show ip eigrp 300 topology 34.1.1.0 255.255.255.0

EIGRP-IPv4 Topology Entry for AS(300)/ID(0.0.0.1) for 34.1.1.0/24

State is Passive, Query origin flag is 1, 2 Successor(s), FD is 1024000

Descriptor Blocks:

13.1.1.3 (Serial1/3), from 13.1.1.3, Send flag is 0x0

Composite metric is (1024000/512000), route is Internal

Vector metric:

Minimum bandwidth is 128 Kbit

Total delay is 40000 microseconds

Hop count is 1

Originating router is 0.0.0.1

14.1.1.4 (Serial1/4), from 14.1.1.4, Send flag is 0x0

Composite metric is (1024000/512000), route is Internal

Vector metric:

Minimum bandwidth is 128 Kbit

Total delay is 40000 microseconds

Hop count is 1

Originating router is 0.0.0.4

Background info regarding the calculation:

Bandwidth: 10'000'000 is a fixed EIGRP value.

Delay: under the interface delay is configured in 10th of micro-seconds

Calculate Path via Serial 1/3

Bandwidth = 10'000'000 / 128 = 78125

Delay 40000 / 10 = 4000

You have two choices, either manipulate the Delay value or the Bandwidth Value. Its recommended to use Delay, as routing protocols and QoS uses the Bandwidth statement under the interface for their calculations, whereas only EIGRP considers delay.

EIGRP Metric formula:

256 (BW + DLY) = COMPOSITE-Metric

In order to have 3x the metric (63072000) received from R4 (21024000) we need to calculate the following:

63072000 = 256 (BW + DLY)

63072000 = 256 (78125 + 4000)

63072000 / 256 = 82125

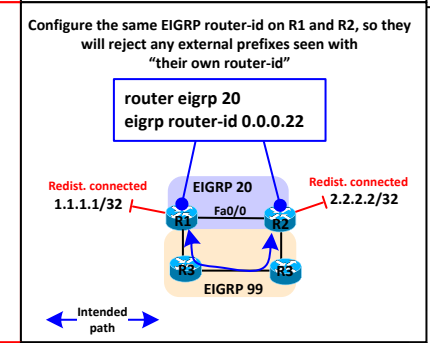
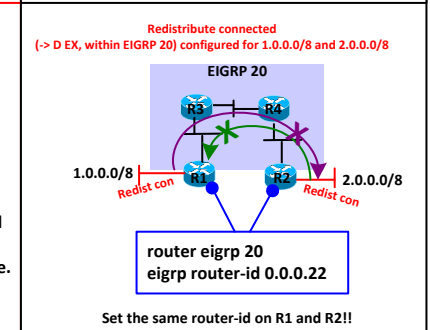
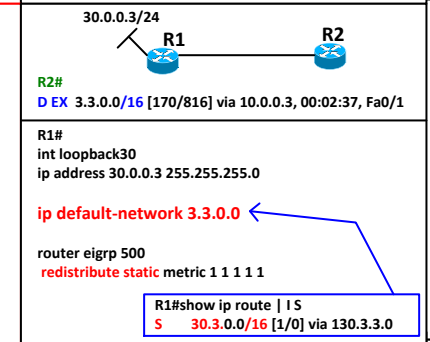
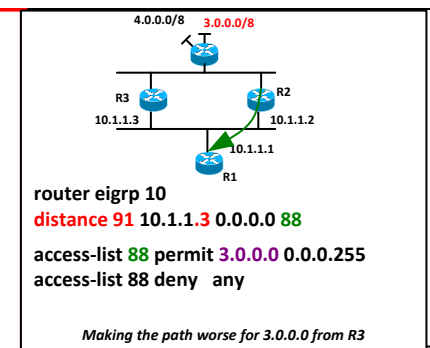
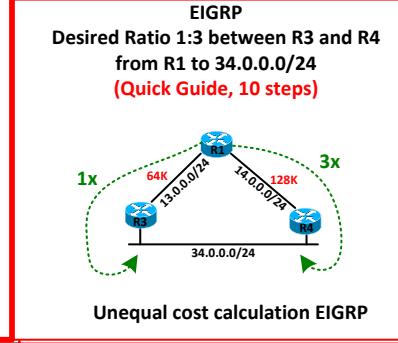

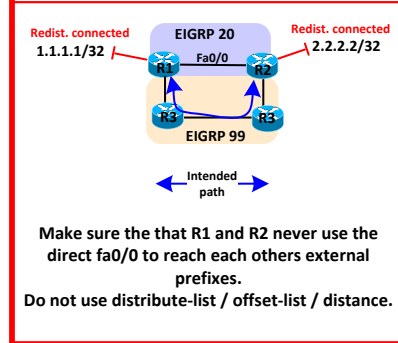
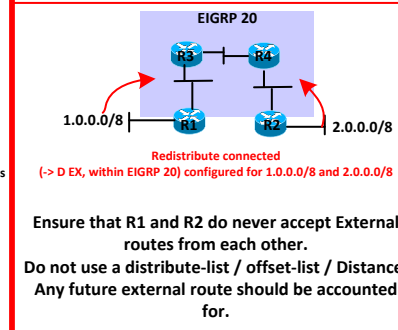
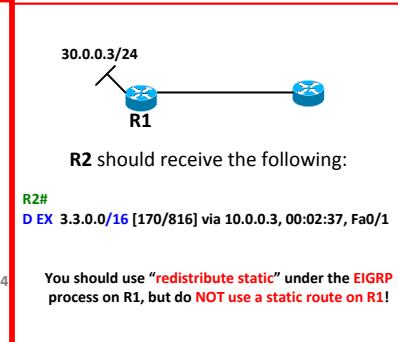
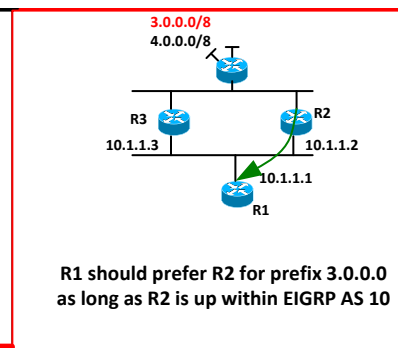
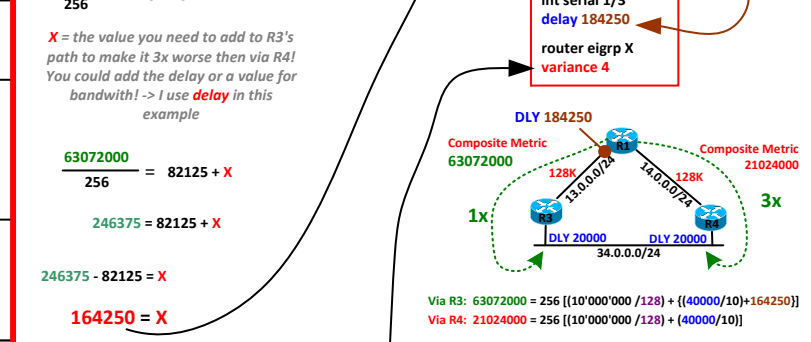
164250 + 20000 = 184250

The solution:

```
R1#conf t
int serial 1/3
delay 184250
router eigrp X
variance 4
```

R1#show interface ser 1/3

```
MTU 1500 bytes, BW 128 Kbit/sec, DLY 20000 usec,
```



R1#show ip route 34.1.1.0 | i share

```
Route metric is 21024000, traffic share count is 1 (Serial 1/3)
Route metric is 21024000, traffic share count is 1 (Serial 1/4)
```

Currently the share is 1:1, use the metric multiply by 3x

3x 21024000 = 63072000 (used to make R3's path 3x worse than R4's)

R1#show ip eigrp 300 topology 34.1.1.0 255.255.255.0

Composite metric is (1024000/512000), route is Internal

Vector metric:

Minimum bandwidth is 128 Kbit

Total delay is 40000 microseconds

Bandwidth: 10'000'000 / 128 = 78125

Delay: 40'000 / 10 = 4000

Calculation!

4000 + 78125 = 82125

R1#show int ser 1/3

```
..., BW 128 Kbit/sec, DLY 20000
```

63072000 = 82125 + X

256

246375 = 82125 + X

246375 - 82125 = X

164250 = X

164250 + 20000 = 184250

router eigrp X

```
variance 3
```

R1#conf t

```
int serial 1/3
delay 184250
```

- Show ip route 34.0.0.0 | i metric
- Path 1 via R3 metric = 21024000
- Path 2 via R4 metric = 21024000
- Calculate the needed metric, to make the second path worse enough to fit variance (3x 21024000 = 63072000)
- Identify Path Total Delay and minimum bandwidth using show ip eigrp X topology 34.0.0.0 255.255.255.0
- Calculate Total Delay / 10 = DLY (4000)
- 63072000 = 256 (BW + DLY)
- 63072000 / 256 = 82125
- add 164250 to 20000 = 184250
- Check interface delay R1 to R3 show ip int ser 1/3 | i Delay DLY 20000 usec
- add 164250 to 20000 = 184250
- R1#conf t
- int serial 1/3
- delay 184250
- router eigrp X
- variance 3

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

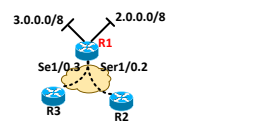
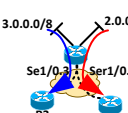
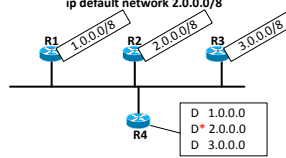
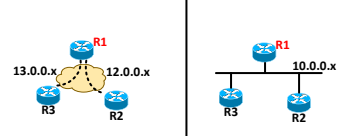
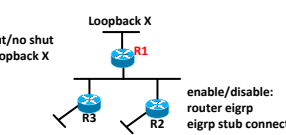
Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

EIGRP

<p>Have R1 leak 2.0.0.0/8 to R2 while you leak 3.0.0.0/8 to R3. Do not remove "eigrp stub connected" on R1:</p> 	<pre>R1# router eigrp 100 network 2.0.0.0 network 3.0.0.0 eigrp stub connected leak-map RMP-LEAK-EIGRP ip prefix-list PFX-LEAK-2 seq 5 permit 2.0.0.0/8 ip prefix-list PFX-LEAK-3 seq 10 permit 3.0.0.0/8 route-map RMP-LEAK-EIGRP permit 10 match ip address prefix-list PFX-LEAK-2 match interface ser0/0.2 route-map RMP-LEAK-EIGRP permit 20 match ip address prefix-list PFX-LEAK-3 match interface Ser0/0.3</pre> 				
<pre>R1# router eigrp 100 network 1.0.0.0 ip default network 1.0.0.0/8 R2# router eigrp 100 network 2.0.0.0 ip default network 2.0.0.0/8 R3# router eigrp 100 network 3.0.0.0 ip default network 3.0.0.0/8</pre>  <p>How can you instruct R4 to have all 3 routes in its routing table, but only accepts 2.0.0.0 as the default route?</p> <pre>R4# D 1.0.0.0 D* 2.0.0.0 D 3.0.0.0</pre>	<p>R4# router eigrp 100 default-information in 2 access-list 2 permit 2.0.0.0</p> <pre>R1#show ip route D 1.0.0.0/8 [90/156160] via 10.0.0.2, 00:09:53, Fa0/0 D* 2.0.0.0/8 [90/156160] via 10.0.0.2, 00:09:53, Fa0/0 D 3.0.0.0/8 [90/156160] via 10.0.0.2, 00:09:53, Fa0/0</pre>				
<p>Configure R1 into an EIGRP stub in these Situations</p> 	<pre>R1# router eigrp 100 network 12.0.0.1 0.0.0.0 network 13.0.0.1 0.0.0.0 eigrp stub</pre>	<pre>R1# router eigrp 100 network 10.0.0.1 0.0.0.0 neighbor 10.0.0.2 fa0/0 neighbor 10.0.0.3 fa0/0 eigrp stub R2 and R3# router eigrp 100 network 10.0.0.[2,3] 0.0.0.0 neighbor 10.0.0.1 fa0/0</pre>			
<p>How to test EIGRP Stub behaviour:</p>  <p>What will debug eigrp packets terse display in the case R2 and R3 have eigrp stub connected enabled and once disabled?</p>	<p>NO EIGRP STUB:</p> <pre>debug ip packet terse (on R1 and then shutdown loopback X) EIGRP: Enqueueing QUERY on Fa0/21 tid 0 iidbQ un/rely 0/1 serno 43-43 EIGRP: Enqueueing QUERY on Fa0/19 tid 0 iidbQ un/rely 0/1 serno 43-43 EIGRP: Sending QUERY on Fa0/21 tid 0 EIGRP: Sending QUERY on Fa0/19 tid 0 EIGRP is enqueueing and sending a query to R2 and R3. EIGRP STUB enabled on R2 debug ip packet terse (on R1 and then shutdown loopback X) EIGRP: Enqueueing QUERY on Fa0/21 tid 0 iidbQ un/rely 0/1 serno 43-43 EIGRP: Enqueueing QUERY on Fa0/19 tid 0 iidbQ un/rely 0/1 serno 43-43 EIGRP: Sending QUERY on Fa0/19 tid 0 EIGRP is enqueueing for both but only sending the query to R3! </pre>				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin

Redistribution

<p>Narbiks Redistribution Route-Map:</p> <p>2 to 1 route-map</p> <pre> route-map O->E deny 10 match tag 90 route-map O->E permit 20 set tag 110 route-map E->O deny 10 match tag 110 route-map E->O permit 20 set tag 90 router ospf 1 redistribute eigrp 35 subnets route-map E->O metric-type 2 router eigrp 35 redistribute ospf 1 metric 100000 100 255 1 1500 route-map O->E </pre>	<p>Show ip alias output explained:</p> <pre> SW4#show ip alias Address Type IP Address Port Interface 150.1.10.10 Interface 155.1.10.10 Interface 155.1.108.10 Shows all connected interfaces: Similar to show ip int brie e una </pre>	<p>Redistribution solution 1</p> <p><i>Use distribute-lists</i></p> <p>R1 and R4 router eigrp 20 redistribute rip metric 1 1 1 1 1 distance eigrp 90 110</p> <p>router rip redistribute eigrp 20 metric 1</p> <p>R X via R2 [3 hop] R X via R5 [2 hop] R X via R3 [6 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	<p>Redistribution solution 2</p> <p><i>Use EIGRP duplicated router-id</i></p> <p>R1 and R4 router eigrp 20 redistribute rip metric 1 1 1 1 1 distance eigrp 90 110</p> <p>router rip redistribute eigrp 20 metric 1</p> <p>R X via R2 [3 hop] R X via R5 [2 hop] R X via R3 [6 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	<p>Redistribution solution 3</p> <p><i>Using a route-map and prefix-lists</i></p> <p>R1 and R4 router eigrp 20 redistribute rip metric 1 1 1 1 1 distance eigrp 90 110</p> <p>router rip redistribute eigrp 20 metric 1</p> <p>R X via R2 [3 hop] R X via R5 [2 hop] R X via R3 [6 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	<p>Redistribution solution 4</p> <p><i>Use the 2-to-1 route-map</i></p> <p>R1 and R4 router eigrp 100 redistribute rip metric 1 1 1 1 1 route-map RMP-R-2-E</p> <p>router rip redistribute eigrp 100 metric 1 route-map RMP-E-2-R</p> <p>route-map RMP-E-2-R deny 10 match tag 120</p> <p>route-map RMP-E-2-R permit 20 set tag 90</p> <p>route-map RMP-R-2-E deny 10 match tag 90</p> <p>route-map RMP-R-2-E permit 20 set tag 120</p> <p>R X via R2 [7 hop] R X via R3 [6 hop] R X via R5 [8 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>
<p>Handy VTP debug commands:</p> <pre> debug sw-vlan vtp events debug sw-vlan vtp packets </pre>	<p>Narbiks 8 redistribution methods</p> <p>Briefly mentioned:</p> <ol style="list-style-type: none"> 1. RIP distribute list, deny advertised routes inbound 2. Same Router ID on EIGRP 3. Redistribute appropriate routes using RMPs and PFX 4. 2 to 1 route-map 5. Filter based on route summarization (longest match) 6. Set lower distance on neighbor (RIP / OSPF) 7. Distance? 8. EIGRP races OSPF condition (EIGRP is super quick) <pre> acl into-eigrp deny 5.0.0.0 acl into-eigrp permit any router ospf 1 distribute-list into-eigrp out ospf 1 router eigrp 100 distribute-list into-ospf out eigrp 100 </pre>	<p>Redistribution solution 5</p> <p><i>Using Summaries</i></p> <p>R1 and R4 router eigrp 20 redistribute rip metric 1 1 1 1 1 distance eigrp 90 110</p> <p>router rip redistribute eigrp 20 metric 1</p> <p>R X via R2 [3 hop] R X via R5 [2 hop] R X via R3 [6 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	<p>Redistribution solution 6</p> <p><i>Using Summaries</i></p> <p>R1 and R4 int fa0/0 (int pointing to R2, R5) ip summary-address eigrp 20 3.0.0.0 255.0.0.0</p> <p>R1 to R2 advertises: 3.0.0.0/8 not /24! R1 has 3.3.3.0/24 via R2 in RT</p> <p>R4 to R5 advertises: 3.0.0.0/8 not /24! R4 has 3.3.3.0/24 via R5 in RT</p> <p>R X via R2 [7 hop] R X via R3 [6 hop] R X via R5 [8 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	<p>Redistribution solution 7</p> <p><i>Using Summaries</i></p> <p>R1 and R4 int fa0/0 (int pointing to R2, R5) ip summary-address eigrp 20 3.0.0.0 255.0.0.0</p> <p>R1 to R2 advertises: 3.0.0.0/8 not /24! R1 has 3.3.3.0/24 via R2 in RT</p> <p>R4 to R5 advertises: 3.0.0.0/8 not /24! R4 has 3.3.3.0/24 via R5 in RT</p> <p>R X via R2 [7 hop] R X via R3 [6 hop] R X via R5 [8 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	
<p>Disable error messages for MOSPF LSA Type 6 messages</p> <pre> Conf t Router ospf X Ignore LSA MOSPF </pre>	<p>How to troubleshoot redistribution issues?</p> <p>Use the following command on all redistributing routers:</p> <p>debug ip routing</p> <p>Wait a few moments and see what routes are flapping due to wrong redistribution</p> <pre> SW4# debug ip routing RT: add 54.1.1.0/24 via 183.1.105.5, eigrp metric [170/2560002816] </pre>	<p>Redistribution solution 8</p> <p><i>Using Summaries</i></p> <p>R1 and R4 int fa0/0 (int pointing to R2, R5) ip summary-address eigrp 20 3.0.0.0 255.0.0.0</p> <p>R1 to R2 advertises: 3.0.0.0/8 not /24! R1 has 3.3.3.0/24 via R2 in RT</p> <p>R4 to R5 advertises: 3.0.0.0/8 not /24! R4 has 3.3.3.0/24 via R5 in RT</p> <p>R X via R2 [7 hop] R X via R3 [6 hop] R X via R5 [8 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	<p>Redistribution solution 9</p> <p><i>Using Summaries</i></p> <p>R1 and R4 int fa0/0 (int pointing to R2, R5) ip summary-address eigrp 20 3.0.0.0 255.0.0.0</p> <p>R1 to R2 advertises: 3.0.0.0/8 not /24! R1 has 3.3.3.0/24 via R2 in RT</p> <p>R4 to R5 advertises: 3.0.0.0/8 not /24! R4 has 3.3.3.0/24 via R5 in RT</p> <p>R X via R2 [7 hop] R X via R3 [6 hop] R X via R5 [8 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	<p>Redistribution solution 10</p> <p><i>Using Summaries</i></p> <p>R1 and R4 int fa0/0 (int pointing to R2, R5) ip summary-address eigrp 20 3.0.0.0 255.0.0.0</p> <p>R1 to R2 advertises: 3.0.0.0/8 not /24! R1 has 3.3.3.0/24 via R2 in RT</p> <p>R4 to R5 advertises: 3.0.0.0/8 not /24! R4 has 3.3.3.0/24 via R5 in RT</p> <p>R X via R2 [7 hop] R X via R3 [6 hop] R X via R5 [8 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	
<p>What could be forgotten to redistribute in this situation?</p> <p>Redistributing EIGRP 10 into OSPF</p> <pre> router ospf 1 redist eigrp 10 subnets route-map RMP-E-2-O route-map RMP-E-2-O deny 10 match tag 110 route-map RMP-E-2-O permit 20 set tag 90 route-map RMP-E-2-O permit 20 match interface Fa0/1 </pre> <p>Make sure to redistribute the connected interface on the EIGRP side into OSPF too!!!!</p>	<p>What could be the problem here:</p> <pre> R1#show run i route ip route 200.1.12.0 255.255.255.0 200.1.12.11 R1#show run i route ip route 200.1.12.0 255.255.255.0 200.1.12.11 %No matching route to delete R1(config)#no ip route 200.1.12.0 255.255.255.0 200.1.12.11 %No matching route to delete R1#conf t R1(config)#no ip default-network 200.1.12.11 </pre> <p>Someone configured a default network, but not to a classfull network! You will not be able to NO out the static route other than doing this:</p>	<p>How can you solve this that R1 and R2 will be able to ping the 3.3.3.0/24 network?</p> <p>Fix: router rip distance 109</p>	<p>Will R2 be able to ping 3.3.3.3 ?</p> <p>What happens here in detail?</p> <ol style="list-style-type: none"> 1. R2 received route 3.x via RIP from R3 2. R2 advertises 3.0.0.0 via RIP to R1 3. R1 redistributes 3.0.0.0 into OSPF advertises towards R2 4. R2 starts using OSPF route to 3.x. flushes the RIP route and announces it as inaccessible within RIP. 5. R1 receives inaccessible RIP route from R2 for 3.0.0.0 and 6. R1 advertises OSPF max-age for 3.0.0.0 route, R2 flushes route 7. R2 receives RIP update from R3, and it starts at 1. again 	<p>Redistribution solution 11</p> <p><i>Using Summaries</i></p> <p>R1 and R4 int fa0/0 (int pointing to R2, R5) ip summary-address eigrp 20 3.0.0.0 255.0.0.0</p> <p>R1 to R2 advertises: 3.0.0.0/8 not /24! R1 has 3.3.3.0/24 via R2 in RT</p> <p>R4 to R5 advertises: 3.0.0.0/8 not /24! R4 has 3.3.3.0/24 via R5 in RT</p> <p>R X via R2 [7 hop] R X via R3 [6 hop] R X via R5 [8 hop] R X via R4 [1 hop] D EX X via R1 [1 hop]</p>	

Help me create more flashcards:

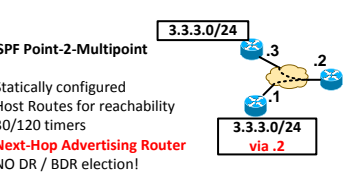
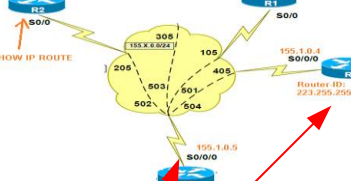
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts



Colin

<p>Configure OSPFv2 for a prefix without using the network statement:</p>	<p>OSPFv2</p> <p>interface FastEthernet0/0</p> <p>ip ospf 1 area 1</p>	<p>show ip ospf interface Serial10/0:</p>	<pre>Serial10/0 is up, line protocol is up Internet Address 155.1.0.5/24, Area 0 Process ID 1, Router ID 150.1.5.5, Network Type NON_BROADCAST, Cost: 64 Enabled by interface config, including secondary ip addresses Transmit Delay is 1 sec, State DR, Priority 1 Designated Router (ID) 150.1.5.5, Interface address 155.1.0.5 Backup Designated router (ID) 223.255.255.255, Interface address 155.1.0.4 Timer intervals configured, Hello 30, Dead 120, Wait 120, Retransmit 5 oob-resync timeout 120 Hello due in 00:00:16 Supports Link-local Signaling (LLS) Index 1/2, flood queue length 0 Next 0x0(0)/0x0(0) Last flood scan length is 1, maximum is 4 Last flood scan time is 0 msec, maximum is 4 msec Neighbor Count is 4, Adjacent neighbor count is 4 Adjacent with neighbor 150.1.3.3 Adjacent with neighbor 150.1.1.1 Adjacent with neighbor 150.1.2.2 Adjacent with neighbor 223.255.255.255 (Backup Designated Router)</pre>	<p>Serial interfaces P-2-P output of OSPF</p>	<pre>R3#show ip ospf neighbor Neighbor ID Pri State Dead Time Address Interface 150.1.2.2 0 FULL/- 00:00:37 155.1.23.2 Serial1/3 R3#show ip ospf interface brief Interface PID Area IP Address/Mask Cost State Nbrs F/C Ser1/3 1 5 155.1.23.3/24 781 P2P 1/1</pre>
<p>Configure all attached interfaces into OSPF area 2 with one line:</p>	<p>router ospf 1</p> <p>network 0.0.0.0 255.255.255.255 area 2</p>	<p>Show ip ospf neighbor</p> <p>Output:</p>	<pre>sh ip ospf neighbor Neighbor ID Pri State Dead Time Address Interface 150.1.1.1 1 FULL/DR 00:00:34 155.1.146.1 Gi0/1 150.1.6.6 1 FULL/DR 00:00:38 155.1.146.6 Gi0/1</pre>	<p>OSPF Point-2-Multipoint:</p>	 <p>OSPF Point-2-Multipoint</p> <ul style="list-style-type: none"> - Statically configured - Host Routes for reachability - 30/120 timers - Next-Hop Advertising Router - NO DR / BDR election! <pre>interface Serial0/0 ip ospf network point-to-multipoint frame-relay map ip 155.1.0.5 105 broadcast</pre>
<p>What is special about OSPF network statement and ip unnumbered interfaces?</p>	<p>If there is an IP unnumbered command configured on Fa0/0 and the network statement covers only the IP unnumbered address space, Fa0/0 will also be enabled for that OSPF instance.</p> <p>Interface fa0/0</p> <p>ip unnumbered</p>	<p>What is important to keep in mind when it comes to OSPF configuration with Loopbacks and Router-ID?</p>	<p>Within lab, double check for highest Loopback, in case you do not configure a router-id for OSPF, they may deliberately want to break it like that.</p>	<p>Show IP route output explained in combination with OSPF</p>	 <pre>R2#show ip route 155.1.0.4 Routing entry for 155.1.0.4/32 Known via "ospf 1", distance 110, metric 128, type intra area Last update from 155.1.0.5 on Serial0/0, 00:23:39 ago Routing Descriptor Blocks: * 155.1.0.5, from 223.255.255.255, 00:23:39 ago, via Serial0/0 Route metric is 128, traffic share count is 1</pre>
<p>OSPF network statement described:</p>	<p>it simply enables the OSPF process on the interface.</p> <p>If multiple network statements overlap the same interface, the most specific match based on the wildcard mask wins.</p>	<p>How to detect a duplicated router-id in ospf?</p>	<pre>Router1: int Loopback222 ip address 222.255.255.255 /32 Router3: int Loopback222 ip address 222.255.255.255 /32 %OSPF-4-DUP_RTRID_AREA: Detected router with duplicate router ID 222.255.255.255 in area 2 %OSPF-4-FLOOD_WAR: Process 1 re-originates LSA ID 155.1.79.9 type-2 adv-rtr 222.255.255.255 in area 2</pre>	<p>OSPF Broadcast:</p>	<p>OSPF Broadcast</p> <ul style="list-style-type: none"> - Ethernet - DR / BDR Election - Multicast 224.0.0.5 / 224.0.0.6 - 10/40 Timers - Next-hop does not change
<p>Matching specifically one address into OSPF area 3:</p>	<p>router ospf</p> <p>network 155.1.10.10 0.0.0.0 area 3</p> <p>IP 155.1.10.10</p> <p>Total match via 0.0.0.0</p>	<p>What to check if networks are visible in the OSPF database but not in the routing table:</p>	<pre>R2#show ip ospf database router OSPF Router with ID (150.1.2.2) (Process ID 1) Router Link States (Area 0) Adv Router is not-reachable LS age: 1419 Options: (No TOS-capability, DC) LS Type: Router Links Link State ID: 150.1.1.1 Advertising Router: 150.1.1.1 LS Seq Number: 80000007 Checksum: 0x4B33 Length: 36</pre>	<p>OSPF Non-Broadcast:</p> <p>(Frame-Relay Multipoint)</p>	<p>OSPF Non-Broadcast</p> <ul style="list-style-type: none"> - Frame-Relay Multipoint - DR / BDR - Unicast (Neighbor command) - 30/120 Timers - Next-hop does not change
<p>show ip ospf interface brief:</p>	<p>show ip ospf interface brief</p> <pre>R2#show ip ospf interface brief Interface PID Area IP Address/Mask Cost State Nbrs F/C Fa0/0 1 51 192.10.1.2/24 1 BDR 1/1</pre> <p>Use this command after initial ospf config, to verify</p>	<p>Show ip ospf database (what Net Link status'es reveal):</p>	<pre>R3#show ip ospf database OSPF Router with ID (150.1.3.3) (Process ID 1) NET Link States (Area 2) L Link ID ADV Router Age Seq# Checksum 155.1.37.7 150.1.7.7 428 0x80000005 0x0098E8 155.1.97.6 150.1.5.6 1314 0x80000007 0x0089A7 155.1.79.7 150.1.7.7 1403 0x80000004 0x0060E1</pre> <p>DR or advertising router</p> <p>network segment</p> <p>If no "Net Link States (Area X)" is visible in the OSPF database, there was no DR elected!</p>	<p>OSPF Point-to-Point:</p> <p>HDLC / FR P-2-P / PPP</p>	<p>OSPF Point-to-Point</p> <ul style="list-style-type: none"> - No DR / BDR Election - Traffic to 224.0.0.5 - 10/40 Timers - Next-hop own address

Help me create more flashcards:

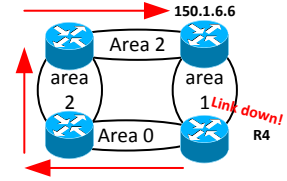
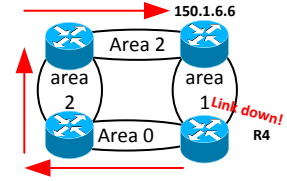
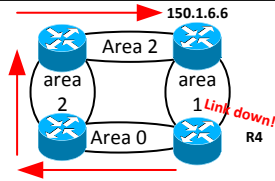
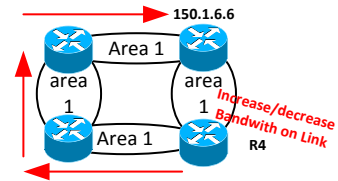
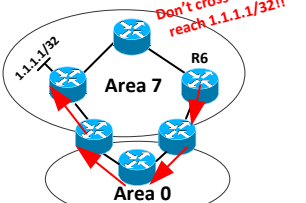
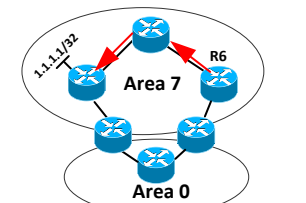
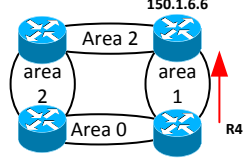
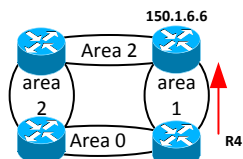
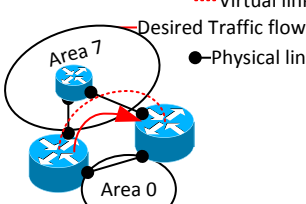
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts



Colin

<p>OSPF Point-Multipoint Non-Broadcast</p> <ul style="list-style-type: none"> - Unicast Neighbor command! - Host Routes for reachability - 30/120 timers - Next-Hop Advertising Router - NO DR / BDR election! <p>interface Serial0/0 ip ospf network point-to-multipoint non-broadcast frame-relay map ip 155.1.0.5 105</p> <p>router ospf X neighbour 155.1.0.5</p>	<p>OSPF Point-2-Multipoint NON-Broadcast</p> <p>Show ip route output of R4 using an inter-area route to R6:</p> 		 <pre>R4#sh ip route 150.1.6.6 Routing entry for 150.1.6.6/32 Known via "ospf 1", distance 110, metric 40431, type inter area Last update from 155.1.45.5 on Serial0/1/0, 00:00:00 ago Routing Descriptor Blocks: 155.1.45.5, from 150.1.1.1, 00:00:00 ago, via Serial0/1/0 Route metric is 40431, traffic share count is 1</pre>	<p>What does (DNA) within a "show ip ospf database" stand for?</p>	<pre>R6# show ip ospf database OSPF Router with ID (150.1.6.6) (Process ID 1) Router Link States (Area 0) Link ID ADV Router Age Seq# Checksum Link count 150.1.1.1 150.1.1.1 1 (DNA) 0x8023 0x002C1E 4 150.1.2.2 150.1.2.2 793 (DNA) 0x8028 0x00CB61 3 150.1.3.3 150.1.3.3 760 (DNA) 0x8025 0x00C76B 3</pre> <p>Do Not Age bit, or learned via Virtual Link most likely</p>
<p>Different OSPF network types:</p> <ul style="list-style-type: none"> ip ospf network point-to-point ip ospf network broadcast ip ospf network non-broadcast ip ospf network point-to-multipoint ip ospf network point-to-multipoint non-broadcast 	<p>Useful command troubleshooting IP OSPF Cost and paths:</p>	<pre>R#show ip ospf interface i cost Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_POINT, Cost: 15 Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_MULTIPoint, Cost: 15 Process ID 1, Router ID 150.1.5.5, Network Type POINT_TO_POINT, Cost: 1 Process ID 1, Router ID 150.1.5.5, Network Type BROADCAST, Cost: 10 Process ID 1, Router ID 150.1.5.5, Network Type BROADCAST, Cost: 10</pre>	<p>What is special in terms of Virtual Link and Router-IDs, associated to the highest interface?</p>	<p>If a Virtual link is setup, pointing to the current Router-ID which happens to be the highest Loopback interface.</p> <p>A new, even higher Loopback number is configured, the Virtual Link will be broken, in the event that the OSPF process is restarted as the new even higher Loopback will take its place as OSPF Router-ID, whereas the other, configured Router is pointing at the wrong Router-ID and will fail to establish the virtual-link.</p> <p>Virtual-link is only a control-plane solution, and not a data-plane!</p>	
<p>Debug IP packet of OSPF:</p> <ul style="list-style-type: none"> Point-2-Point Broadcast Point-to-multipoint non-broadcast 	<pre>interface Serial0/1/0 ip ospf 1 area 0 IP: s=155.1.45.4 (Serial0/1/0), d=224.0.0.5, len 80, rcvd 0, proto=89 interface GigabitEthernet0/0 ip address 155.1.58.5 255.255.255.0 ip ospf 1 area 3 s=155.1.58.8 (GigabitEthernet0/0), d=224.0.0.5, len 80, rcvd 0, proto=89 interface Serial0/0/0 ip address 155.1.0.5 255.255.255.0 encapsulation frame-relay ip ospf network point-to-multipoint non-broadcast frame-relay map ip 155.1.0.1 501 no frame-relay inverse-arp IP: s=155.1.0.5 (local), d=155.1.0.3 (Serial0/0/0), len 92, sending, proto=89</pre>	<p>How to display the local OSPF database which all local networks</p>	<pre>show ip ospf database router 150.1.8.8 self-originate (150.1.8.8 is its own local Loopback IP address)</pre>	<h3>Repairing Discontiguous OSPF Areas with Virtual-Links</h3>	<p>R1: router ospf 1 area 1 virtual-link 150.1.6.6</p> <p>R6: router ospf 1 area 1 virtual-link 150.1.1.1</p> <p>Make sure Router-ID has been manually configured!</p>
<p>How do you advertise the same prefix into two OSPF areas using one network statement and one interface command?</p>	<pre>int loopback99 ip address 99.99.99.99 255.255.255.255 ip ospf 1 area 88 router ospf 1 network 99.99.99.99 255.255.255.255 area 99</pre>	<h3>OSPF Path Selection with Bandwidth</h3>	 <pre>interface Serial0/0 bandwidth 10000</pre>	<p>OSPF</p> <p>router ospf 1 no capability transit</p> <p>What does this command do?</p>	<p>Within Area 1: router ospf 1 no capability transit</p> 
<p>Auto-cost reference bandwidth could potentially harm the network, explain why:</p>	<p>A routing loop could occur due to mismatched auto-cost reference bandwidth.</p> <p>Set the auto-cost value consistently throughout the entire OSPF domain</p>	<h3>Calculating OSPF path costs</h3> <p>Calculate cost via ABR:</p> <p>Calculate cost to ABR:</p> <p>Summarize both = route metric:</p>	<pre>show ip ospf database summary 150.1.8.0 Link State ID: 150.1.8.0 (summary Network Number) Advertising Router: 150.1.1.1 (IP of the ABR) Metric: 19731 (shows metric as of the ABR) show ip ospf database router 150.1.6.6 self-originate Link connected to: a Transit Network (Link ID) Designated Router address: 155.1.146.6 (Link Data) Router Interface address: 155.1.146.6 Number of TOS metrics: 0 TOS 0 Metrics: 300 (shows metric to the ABR) R6#show ip route 150.1.8.8 via FastEthernet0/146 Route metric is 20031, traffic share count is 1</pre>	<p>What does the following OSPF command do?</p> <p>router ospf 1 capability transit</p>	<p>capability transit enabled by default!</p> 
<p>Show ip route output of R4 using an intra-area route to R6:</p> 	 <pre>R4#sh ip route 150.1.6.6 Routing entry for 150.1.6.6/32 Known via "ospf 1", distance 110, metric 301, type intra area Last update from 155.1.146.6 on GigabitEthernet0/1, 00:14:19 ago Routing Descriptor Blocks: * 155.1.146.6, from 150.1.6.6, 00:14:19 ago, via GigabitEthernet0/1 Route metric is 301, traffic share count is 1</pre>	<h3>OSPF Path Selection with Per-Neighbor Cost</h3>	<pre>router ospf 1 neighbor 155.1.0.1 cost 1000 neighbor 155.1.0.4 cost 9999</pre>	<h3>OSPF Path Selection with Virtual-Links</h3>	 <p>Use Virtual links to adjust traffic flow over other areas. Cost for Vlink calculated based on physical interfaces!</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

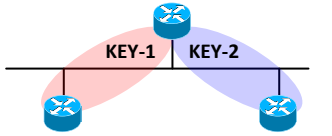
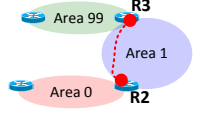
Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin



OSPF

<p>OSPF Demand Circuit</p> <p>reduce periodic OSPF hello transmission</p>	<pre>interface Serial0/1/0 ip ospf demand-circuit R4# show ip ospf interface Serial0/1/0 Serial0/1/0 is up, line protocol is up Enabled by interface config, including secondary ip addresses Configured as demand circuit. Run as demand circuit. DoNotAge LSA allowed. oob-resync timeout 40 Hello due in 00:00:07 Supports Link-local Signaling (LLS) Last flood scan time is 0 msec, maximum is 0 msec Neighbor Count is 1, Adjacent neighbor count is 1 Adjacent with neighbor 150.1.5.5 (Hello suppressed) Suppress hello for 1 neighbor(s)</pre> <p><i>Suppresses Hello msg. while maintaining the ADJ</i></p>	<p>Authentication and Area 0 or Virtual links:</p>	<p>a virtual-link is an interface in area 0 !!</p> <p>Be carefull not to forget this while enabling:</p> <p>area 0 authentication</p>	<p>OSPF Path Selection with Summarization</p>	<p>Deliberately make routes a longer prefix to force traffic another, more specific way.</p> <p>Use</p> <p>Router ospf X Area X range x.x.x.x 255.255.254.0 to make a /24 look like a /23 to force a different route.</p>
<p>OSPF Flooding Reduction</p> <p>Reducing "paranoid update"</p>	<p>feature stops unnecessary LSA flooding by setting the DoNotAge (DNA) bit in the LSA, removing the requirement for the periodic refresh.</p> <pre>interface Vlan10 ip ospf flood-reduction</pre>	<p>OSPF Null Authentication</p>	<pre>interface Vlan7 ip ospf authentication null</pre>	<p>OSPF metric and forward metric on E2 routes:</p>	<pre>SW3#show ip route 51.51.51.51 Routing entry for 51.51.51.51/32 Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 50 Last update from 155.1.79.7 on Vlan79, 00:06:49 ago Routing Descriptor Blocks: * 155.1.79.7, from 192.10.1.254, 00:06:49 ago, via Vlan79 Route metric is 20, traffic share count is 1</pre> <p>Forward metric, gives information about which paths is taken effectively to the E2 destination.</p>
<p>OSPF Type 1 (clear text) Authentication</p> <p>Per interface</p>	<pre>interface FastEthernet0/0 ip ospf authentication ip ospf authentication-key SECRET</pre>	<p>OSPF authentication, the difference between interface / process config:</p>	<p>The authentication type configured at the interface level overrides the authentication type configured at the process level</p>	<p>OSPF Stub config and output:</p>	<pre>router ospf 1 area 3 stub Show ip route O 155.1.5.0/24 [110/2] via 155.1.58.5, 00:38:42, Vlan58 O IA 150.1.1.1/32 [110/66] via 155.1.58.5, 00:00:21, Vlan58 O IA 0.0.0.0/0 [110/2] via 155.1.58.5, 00:00:21, Vlan58</pre> <pre>SW4#sh ip ospf database OSPF Router with ID (150.1.10.10) (Process ID 1) Router Link States (Area 3) Net Link States (Area 3) Summary Net Link States (Area 3) Several entries due to IA routes</pre>
<p>OSPF Type 1 (clear text) Authentication</p> <p>Per area / interface</p>	<pre>interface Vlan67 ip ospf authentication-key SECRET router ospf 1 area 2 authentication</pre>	<p>OSPF MD5 Authentication with Multiple Keys</p>	<pre>interface FastEthernet0/0 ip ospf authentication message-digest ip ospf message-digest-key 16 md5 KEY-1 ip ospf message-digest-key 46 md5 KEY-2</pre> 	<p>What four OSPF stub types are there?</p>	<p>stub area</p> <p>totally stubby area</p> <p>not-so-stubby area (NSSA)</p> <p>not-so-totally stubby area</p>
<p>How to troubleshoot OSPF Authentication mis-match:</p>	<pre>debug ip ospf adj</pre> <p>OSPF: Send with youngest Key 0 OSPF: Rcv pkt from 155.1.79.7, Vlan79 : Mismatch Authentication type. Input packet specified type 1, we use type 2</p> <p>(Mis-match can be either authentication type or the password.)</p>	<p>Debugging different received OSPF MD5 Keys on the same interface:</p>	<pre>debug ip ospf adj</pre> <p>OSPF: Send with youngest Key 0 OSPF: Rcv pkt from 155.1.146.4, Fa0/0 : Mismatch Authentication Key - No message digest key 46 on interface OSPF: Rcv pkt from 155.1.146.1, Fa0/0: Mismatch Authentication Key - No message digest key 16 on interface</p>	<p>Default route within the OSPF Database:</p> <pre>show ip ospf database summary 0.0.0.0</pre> <p>Output:</p>	<pre>Router ospf 1 Area 3 stub LSA Type3 SW4#show ip ospf database summary 0.0.0.0 OSPF Router with ID (150.1.10.10) (Process ID 1) Summary Net Link States (Area 3) Routing Bit Set on this LSA LS age: 1415 Options: (No TOS-capability, DC, Upward) LS Type: Summary Links(Network) Link State ID: 0.0.0.0 (summary Network Number) Advertising Router: 150.1.5.5 LS Seq Number: 80000001 Checksum: 0x1D7F Length: 28 Network Mask: /0 TOS: 0 Metric: 1</pre>
<p>OSPF MD5 Authentication</p>	<pre>interface Serial0/0 ip ospf 1 area 0 ip ospf message-digest-key 1 md5 KEY</pre> <pre>router ospf 1 area 0 authentication message-digest area 1 virtual-link 150.1.6.6 message-digest-key 1 md5 KEY</pre> 	<p>OSPF Internal Area Summarization</p>	<pre>router ospf 1 area 3 range 155.1.8.0 255.255.252.0</pre>	<p>OSPF Totally Stubby Areas</p> <p>Config:</p>	<p>ONLY ON ABR:</p> <pre>router ospf 1 area 3 stub no-summary</pre> <p>Internal Area routers have:</p> <pre>router ospf 1 area 3 stub</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts



Colin

<h2>OSPF Totally Stubby Areas</h2> <p>show ip ospf database</p> <p>Output:</p>	<p>ABR: Router ospf 1 Area 3 stub no-summary Internal Area Router: Router ospf 1 Area 3 stub</p> <p>Show ip route</p> <p>O 150.1.8.8/32 [110/2] via 155.1.108.8, 00:31:57, Po1 O*IA 0.0.0.0/0 [110/3] via 155.1.108.8, 00:03:46, Po1</p> <p>SW4#show ip ospf database</p> <p>OSPF Router with ID (150.1.10.10) (Process ID 1)</p> <p>Router Link States (Area 3)</p> <p>Net Link States (Area 3)</p> <p>Summary Net Link States (Area 3)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.0</td> <td>150.1.5.5</td> <td>49</td> <td>0x80000002</td> <td>0x001880</td> </tr> </tbody> </table> <p>Only 0.0.0.0 entry in Summary Net Link States! LSA Type 3</p>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.0	150.1.5.5	49	0x80000002	0x001880	<h2>OSPF Not-So-Totally-Stubby Areas</h2>	<p>ABR: router ospf 1 area 2 nssa no-summary area 2 default-cost 500 Internal Area Router: Router ospf 1 area 2 nssa</p> <p>O 155.1.67.0 [110/2] via 155.1.79.7, 00:01:19, Vlan79 O N2 200.0.3.0/24 [110/20] via 155.1.79.7, 00:01:20, Vlan79 O*IA 0.0.0.0/0 [110/502] via 155.1.79.7, 00:01:20, Vlan79</p> <p>SW3#sh ip ospf database</p> <p>OSPF Router with ID (150.1.9.9) (Process ID 1)</p> <p>Router Link States (Area 2)</p> <p>Net Link States (Area 2)</p> <p>Summary Net Link States (Area 2)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.0</td> <td>150.1.3.3</td> <td>382</td> <td>0x80000001</td> <td>0x004F54</td> </tr> </tbody> </table> <p>Type-7 AS External Link States (Area 2)</p>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.0	150.1.3.3	382	0x80000001	0x004F54	<h2>OSPF Forwarding Address Suppression</h2> <p>OSPF Forwarding Address Suppression in Translated Type-5 LSAs</p> <p>with forwarding address suppression enabled the traffic will always flow through the Type-7 to 5 translator.</p> <p>(NSSA with multiple exits, Highest IP is translator)</p> <p>router ospf 1 area 3 nssa no-redistribution no-summary translate type7 suppress-fa</p> <p>Instructs the ABR not to send the original forward address but to set it to 0.0.0.0.</p>	
Link ID	ADV Router	Age	Seq#	Checksum																					
0.0.0.0	150.1.5.5	49	0x80000002	0x001880																					
Link ID	ADV Router	Age	Seq#	Checksum																					
0.0.0.0	150.1.3.3	382	0x80000001	0x004F54																					
<p>Show ip ospf database ?</p> <p>Output:</p>	<p>SW4# show ip ospf database ?</p> <table border="1"> <thead> <tr> <th>router</th> <th>Type 1</th> </tr> </thead> <tbody> <tr> <td>network</td> <td>Type 2 (ADV Router is the DR)</td> </tr> <tr> <td>summary</td> <td>Type 3</td> </tr> <tr> <td>asbr-summary</td> <td>Type 4</td> </tr> <tr> <td>external</td> <td>Type 5, translated Type 7's</td> </tr> <tr> <td>nssa-external</td> <td>NSSA External link states</td> </tr> <tr> <td>self-originate</td> <td>Self-originated link states</td> </tr> </tbody> </table>	router	Type 1	network	Type 2 (ADV Router is the DR)	summary	Type 3	asbr-summary	Type 4	external	Type 5, translated Type 7's	nssa-external	NSSA External link states	self-originate	Self-originated link states	<h2>OSPF Stub Areas with Multiple Exit Points (O*IA / O*N2)</h2>	<p>router ospf 1 area 2 nssa default-information-originate area 2 default-cost 500</p> <p>O*N2 ABR-1 O*IA ABR-2 RTR-1 O*IA 0.0.0.0/0 (Type 3 Sum) O*N2 0.0.0.0/0 (Type 7)</p>	<h2>OSPF Default Routing</h2> <p>2 config options:</p>	<p>router ospf 1 default-information originate always</p> <hr/> <p>ip route 0.0.0.0 0.0.0.0 54.1.1.254</p> <p>router ospf 1 default-information originate metric 60</p>						
router	Type 1																								
network	Type 2 (ADV Router is the DR)																								
summary	Type 3																								
asbr-summary	Type 4																								
external	Type 5, translated Type 7's																								
nssa-external	NSSA External link states																								
self-originate	Self-originated link states																								
<h2>OSPF Not-So-Stubby Areas</h2> <p>Speciality about the Type 7 to Type 5 conversion:</p>	<p>Only one ABR is sending the translated Type 5 into Area 0!</p>	<p>List ospf route preference:</p> <ol style="list-style-type: none"> intra-area (to disable no capability transit) inter-area external nssa-external 	<h2>OSPF Conditional Default Routing</h2> <p>Config:</p>	<p>router ospf 1 default-information originate always route-map RMP-TRACK</p> <p>ip prefix-list PFX-BB1 seq 5 permit 54.1.1.0/24</p> <p>ip prefix-list PFX-R3 seq 5 permit 150.3.3.0/24</p> <p>route-map RMP-TRACK permit 10 match ip address prefix-list PFX-BB1 PFX-R3</p> <p>(Advertise 0.0.0.0/0 if 54.1.1.0/24 or 150.3.3.0/24 is in the routing table)</p>	<p>ip sla monitor 10 type echo protocol icmpEcho 204.12.1.254 timeout 2000 frequency 5 ip sla monitor schedule 10 life forever start-time now</p> <p>track 1 rtr 10</p> <p>ip route 169.254.0.1 255.255.255.255 Null0 track 1 name is.up.always.as.is.route.to.Null0</p> <p>ip prefix-list PLACEHOLDER seq 5 permit 169.254.0.1/32</p> <p>route-map TRACK_PLACEHOLDER permit 10 match ip address prefix-list PLACEHOLDER</p> <p>router ospf 1 default-information originate always route-map TRACK_PLACEHOLDER</p> <p>Logic uses Track on a dummy route which is always UP/UP due to pointed at Null0. Therefore only the IP SLA is important.</p>																				
<p>show ip ospf database nssa-external x.x.x.x</p> <p>Output:</p>	<p>Rack1R3#show ip ospf database nssa-external 200.0.0.0</p> <p>OSPF Router with ID (150.1.3.3) (Process ID 1)</p> <p>Type-7 AS External Link States (Area 2)</p> <p>Routing Bit Set on this LSA</p> <p>LS age: 312</p> <p>Options: (No TOS-capability, Type 7/5 translation, DC)</p> <p>LS Type: AS External Link</p> <p>Link State ID: 200.0.0.0 (External Network Number)</p> <p>Advertising Router: 150.1.6.6</p> <p>LS Seq Number: 80000001</p> <p>Checksum: 0xF94A</p> <p>Length: 36</p> <p>Network Mask: /24</p> <p>Metric Type: 2 (Larger than any link state path)</p> <p>TOS: 0</p> <p>Metric: 20</p> <p>Forward Address: 155.1.67.6</p> <p>External Route Tag: 0</p>	<h2>OSPF NSSA Type-7 to Type-5 Translator Election</h2>	<p>Multiple ABRs connect the NSSA to area 0</p> <p>ABR with the highest router-id is elected as the Type-7 to 5 translator and is responsible for re-originating the Type-5 LSA into area 0.</p> <p>Only one translator</p>	<h2>OSPF Reliable Conditional Default Routing</h2> <p>Config:</p>	<p>ip sla monitor 10 type echo protocol icmpEcho 204.12.1.254 timeout 2000 frequency 5 ip sla monitor schedule 10 life forever start-time now</p> <p>track 1 rtr 10</p> <p>ip route 169.254.0.1 255.255.255.255 Null0 track 1 name is.up.always.as.is.route.to.Null0</p> <p>ip prefix-list PLACEHOLDER seq 5 permit 169.254.0.1/32</p> <p>route-map TRACK_PLACEHOLDER permit 10 match ip address prefix-list PLACEHOLDER</p> <p>router ospf 1 default-information originate always route-map TRACK_PLACEHOLDER</p> <p>Logic uses Track on a dummy route which is always UP/UP due to pointed at Null0. Therefore only the IP SLA is important.</p>																				
<h2>OSPF Not-So-Stubby Areas</h2>	<p>router ospf 1 area 2 nssa</p> <p>Show ip route</p> <p>O 150.1.7.7/32 [110/2] via 155.1.79.7, 00:31:00, Vlan79 O IA 150.1.5.5/32 [110/784] via 155.1.79.7, 00:31:00, Vlan79 O N1 200.0.0.0/24 [110/20] via 155.1.79.7, 00:30:59, Vlan79 O N2 200.0.1.0/24 [110/20] via 155.1.79.7, 00:30:59, Vlan79</p> <p>SW3#show ip ospf database</p> <p>OSPF Router with ID (150.1.9.9) (Process ID 1)</p> <p>Router Link States (Area 2)</p> <p>Net Link States (Area 2)</p> <p>Summary Net Link States (Area 2)</p> <p>Type-7 AS External Link States (Area 2)</p> <p><i>No OSPF domain external routes visible! Default route not automatically generated on ABR!</i></p>	<h2>OSPF NSSA Redistribution Filtering</h2>	<p>SW2#show ip route ospf</p> <p>O N2 5.5.5.5 [110/20] via 155.1.58.5, 00:05:32, Vlan58 O*IA 0.0.0.0/0 [110/2] via 155.1.58.5, 00:01:58, Vlan58</p> <p>Additional N2 is not needed due to the existing default route pointing to the same ABR. The additional N2's can be disabled via:</p> <p>ABR# router ospf 1 area 3 nssa no-redistribution no-summary</p> <p>Area 3 SW2 5.5.5.5/32 9.9.9.9/xx</p>	<h2>OSPF Filtering with Distribute-Lists</h2> <p>Intra-area filtering</p> <p>Config:</p>	<p>Intra-area filtering can be accomplished in OSPF with an inbound distribute-list:</p> <p>All routers have to have the same config, otherwise there is a danger of blackholing networks!</p> <p>router ospf 1 distribute-list 1 in</p> <p>access-list 1 deny 150.1.1.1 access-list 1 deny 150.1.2.2 access-list 1 permit any</p> <p>Area 2 ABR</p>																				
<h2>OSPF Not-So-Stubby Areas and Default Routing</h2>	<p>router ospf 1 area 2 nssa default-information-originate area 2 default-cost 500</p> <p>Internal Area Router output:</p> <p>O 150.1.7.7/32 [110/2] via 155.1.79.7, 00:46:01, Vlan79 O IA 150.1.1.1/32 [110/848] via 155.1.79.7, 00:46:01, Vlan79 O*N2 0.0.0.0 [110/500] via 155.1.79.7, 00:00:38, Vlan79</p>	<h2>OSPF LSA Type-3 Filtering</h2> <p>Config:</p> <p>—Filter out the pfx going into Area 0—</p> <p>Area 3 ABR area 0 155.1.108.0/24</p>	<p>ABR# ip prefix-list AREA_3_ROUTES deny 155.1.108.0/24 ip prefix-list AREA_3_ROUTES permit 0.0.0.0/0 le 32</p> <p>router ospf 1 area 3 filter-list prefix AREA_3_ROUTES out</p> <p>—Filter out the pfx going into Area 0—</p> <p>Area 3 ABR area 0 155.1.108.0/24</p>	<h2>OSPF LSA Type-3 Filtering</h2> <p>Inter-area filtering</p> <p>Config:</p> <p>←Deny pfx from entering Area 0—</p> <p>Area 0 ABR Area 2 155.1.108.0/24</p>	<p>Inter-area filtering</p> <p>coming from Area 0 going out to Area 2</p> <p>router ospf 1 area 0 filter-list prefix FILTER in</p> <p>ip prefix-list FILTER seq 5 deny 155.1.23.0/24 ip prefix-list FILTER seq 10 permit 0.0.0.0/0 le 32</p> <p>←Deny pfx from entering Area 0—</p> <p>Area 0 ABR Area 2 155.1.108.0/24</p>																				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<h3>OSPF Summarization and Discard Routes</h3>	<pre>router ospf 1 no discard-route internal area 0 range 150.1.0.0 255.255.252.0</pre> <p>Discard-route prevent the forwarding of traffic towards a shorter match. automatic origination of the discard route can be disabled via "no discard-route" Internal = inter-area External = summary-address Discard route a ospf generated Null0 route for area range or ip summary address command. If this automatic Null route should be disabled, use NO discard-route (internal/external)</p>	<h3>OSPF Stub Router Advertisement</h3> <p>Config:</p>	<pre>router ospf 1 max-metric router-lsa</pre> <ul style="list-style-type: none"> - prevent traffic black holes - advertise a maximum metric for non-stub destinations - initializing the OSPF process, transit traffic will not flow through the stub router - once the ospf domain is converged, the max metric is withdrawn. <p>max-metric router-lsa on-startup wait-for-bgp (waits for BGP keepalives) max-metric router-lsa on-startup announce-time (how long to wait after a reload)</p>	<p>Ignore OSPF LSA type 6 error messages:</p>	<pre>router ospf 1 ignore lsa mospf</pre>																																																																														
<h3>OSPF Filtering with Administrative Distance</h3>	<pre>access-list 99 permit 155.1.67.0</pre> <pre>router ospf 1 distance 255 150.1.6.6 0.0.0.0 99</pre>	<h3>OSPF Interface Timers</h3>	<pre>interface Serial0/0 ip ospf hello-interval 5 ip ospf dead-interval (seconds)</pre> <pre>interface Serial0/1/0 ip ospf dead-interval minimal hello-multiplier 4</pre> <p>minimal hello-multiplier 4 = 250 msec</p> <p>Timers need to match on all Routers, otherwise the Adjacency will not form!</p>	<h3>OSPF cost calculation</h3> <h3>Changing network types</h3>	<table border="1"> <tr> <th colspan="2">Cost via R4</th> <th colspan="2">Cost via R3</th> </tr> <tr> <td>Fa0/1</td> <td>1</td> <td>Ser0/1</td> <td>20</td> </tr> <tr> <td>Redir static</td> <td>20</td> <td>Redir static</td> <td>20</td> </tr> <tr> <td>Fa0/0</td> <td>1</td> <td>Fa0/1</td> <td>1</td> </tr> <tr> <td>Cost sum:</td> <td>22</td> <td>Cost sum:</td> <td>21</td> </tr> </table>	Cost via R4		Cost via R3		Fa0/1	1	Ser0/1	20	Redir static	20	Redir static	20	Fa0/0	1	Fa0/1	1	Cost sum:	22	Cost sum:	21																																																										
Cost via R4		Cost via R3																																																																																	
Fa0/1	1	Ser0/1	20																																																																																
Redir static	20	Redir static	20																																																																																
Fa0/0	1	Fa0/1	1																																																																																
Cost sum:	22	Cost sum:	21																																																																																
<h3>OSPF Filtering with Route-Maps</h3> <p>Config:</p>	<pre>router ospf 1 distribute-list route-map DENY_R3_LOOPB_FROM_R4 in</pre> <pre>access-list 3 permit 150.1.3.3</pre> <pre>access-list 4 permit 155.1.146.4</pre> <pre>route-map DENY_R3_LOOPB_FROM_R4 deny 10 match ip address 3 match ip next-hop 4</pre> <pre>route-map DENY_R3_LOOPB_FROM_R4 permit 20</pre> <p>Loopback IP is 150.1.3.3/32 Next-Hop of R4 is 155.1.146.4</p>	<h3>OSPF Global Timers</h3>	<pre>router ospf 1 timers throttle spf 100 1000 10000 timers pacing flood 50 timers pacing retransmission 75 timers throttle lsa all 10 4000 6000 timers lsa arrival 2000</pre> <pre>interface Serial0/1 ip ospf transmit-delay 2 ip ospf retransmit-interval 10</pre> <p>SPF pacing/throttling timers control how fast OSPF responds to convergence events.</p> <p>Check via: show ip ospf</p>	<h3>OSPF cost calculation</h3> <h3>Changing network types</h3> <p>Next</p>	<p>R1 will see R4's IP address as forward-address and will calculate the cost to it! → Cost of 1+20+1 = 22</p> <p>R1 will see R3's IP address as forward-address due to the Serial Link and calculate the cost to it: → Cost of 1+20 = 21</p>																																																																														
<h3>OSPF LSA Type-3 Filtering</h3> <p>Config:</p>	<pre>router ospf 1 area 1 filter-list prefix FILTER out</pre> <pre>ip prefix-list FILTER seq 5 deny 155.1.23.0/24</pre> <pre>ip prefix-list FILTER seq 10 permit 0.0.0.0/0 le 32</pre>	<h3>OSPF Resource Limiting</h3> <p>Config:</p>	<pre>router ospf 1 max-lsa 5000 redistribute maximum-prefix 500 process-min-time percent 20</pre> <p>max-lsa: defines a max of 5000 LSA in the OSPF database</p> <p>Redistribute maximum-pref 500: No more than 500 pfx show be originated through redistribution</p> <p>Process-min-time percent: OSPF should not use more than 20% of the CPU resources</p>	<h3>OSPF</h3> <h3>Auto-cost reference-bandwidth</h3> <h3>Calculation:</h3>	<p>cost = Reference-Bandwidth / Bandwidth_of_the_link</p> <p>Desired metrics:</p> <table border="1"> <tr> <th>Bandwidth</th> <th>OSPF Cost</th> </tr> <tr> <td>10'000</td> <td>2</td> </tr> <tr> <td>10</td> <td>2000</td> </tr> <tr> <td>1.544</td> <td>12953</td> </tr> <tr> <td>0.768</td> <td>26041</td> </tr> </table> <p>Used command: auto-cost reference-bandwidth 20000</p> <p>20000/10 = 2000</p>	Bandwidth	OSPF Cost	10'000	2	10	2000	1.544	12953	0.768	26041																																																																				
Bandwidth	OSPF Cost																																																																																		
10'000	2																																																																																		
10	2000																																																																																		
1.544	12953																																																																																		
0.768	26041																																																																																		
<h3>OSPF NSSA ABR External Prefix Filtering</h3> <p>Config:</p>	<p>On ABR:</p> <pre>router ospf 1 area 2 nssa summary-address 200.0.0.0 255.255.255.0 not-advertise</pre> <p>On internal Area router:</p> <pre>router ospf 1 area 2 nssa</pre> <p>ABR translates Type-7 NSSA External LSA to a Type-5 External LSA. Now creating a summary address and using "not-advertise" in order to suppress that LSA Type 5.</p>	<p>Use DNS resolution on the OSPF router-id value in show commands:</p> <pre>R5#sh ip ospf neighbor</pre> <table border="1"> <thead> <tr> <th>Neighbor ID</th> <th>Pri</th> <th>State</th> <th>Dead Time</th> <th>Address</th> <th>Interface</th> </tr> </thead> <tbody> <tr> <td>150.1.4.4</td> <td>0</td> <td>FULL/-</td> <td>00:00:33</td> <td>155.1.45.4</td> <td>Serial0</td> </tr> <tr> <td>Globi-R1</td> <td>0</td> <td>FULL/-</td> <td>00:00:39</td> <td>155.1.0.1</td> <td>Serial0</td> </tr> <tr> <td>150.1.4.4</td> <td>0</td> <td>FULL/-</td> <td>00:00:38</td> <td>155.1.0.4</td> <td>Serial0</td> </tr> <tr> <td>HOBBIT-R2</td> <td>0</td> <td>FULL/-</td> <td>00:00:39</td> <td>155.1.0.2</td> <td>Serial0</td> </tr> <tr> <td>150.1.3.3</td> <td>0</td> <td>FULL/-</td> <td>00:00:38</td> <td>155.1.0.3</td> <td>Serial0</td> </tr> </tbody> </table>	Neighbor ID	Pri	State	Dead Time	Address	Interface	150.1.4.4	0	FULL/-	00:00:33	155.1.45.4	Serial0	Globi-R1	0	FULL/-	00:00:39	155.1.0.1	Serial0	150.1.4.4	0	FULL/-	00:00:38	155.1.0.4	Serial0	HOBBIT-R2	0	FULL/-	00:00:39	155.1.0.2	Serial0	150.1.3.3	0	FULL/-	00:00:38	155.1.0.3	Serial0	<pre>ip host Globi-R1 150.1.1.1</pre> <pre>ip host HOBBIT-R2 150.1.2.2</pre> <pre>ip ospf name-lookup</pre> <pre>R5#sh ip ospf neighbor</pre> <table border="1"> <thead> <tr> <th>Neighbor ID</th> <th>Pri</th> <th>State</th> <th>Dead Time</th> <th>Address</th> <th>Interface</th> </tr> </thead> <tbody> <tr> <td>150.1.4.4</td> <td>0</td> <td>FULL/-</td> <td>00:00:33</td> <td>155.1.45.4</td> <td>Serial0</td> </tr> <tr> <td>Globi-R1</td> <td>0</td> <td>FULL/-</td> <td>00:00:39</td> <td>155.1.0.1</td> <td>Serial0</td> </tr> <tr> <td>150.1.4.4</td> <td>0</td> <td>FULL/-</td> <td>00:00:38</td> <td>155.1.0.4</td> <td>Serial0</td> </tr> <tr> <td>HOBBIT-R2</td> <td>0</td> <td>FULL/-</td> <td>00:00:39</td> <td>155.1.0.2</td> <td>Serial0</td> </tr> <tr> <td>150.1.3.3</td> <td>0</td> <td>FULL/-</td> <td>00:00:38</td> <td>155.1.0.3</td> <td>Serial0</td> </tr> <tr> <td>150.1.8.8</td> <td>1</td> <td>FULL/BDR</td> <td>00:00:31</td> <td>155.1.58.8</td> <td>Gi0/0</td> </tr> </tbody> </table>	Neighbor ID	Pri	State	Dead Time	Address	Interface	150.1.4.4	0	FULL/-	00:00:33	155.1.45.4	Serial0	Globi-R1	0	FULL/-	00:00:39	155.1.0.1	Serial0	150.1.4.4	0	FULL/-	00:00:38	155.1.0.4	Serial0	HOBBIT-R2	0	FULL/-	00:00:39	155.1.0.2	Serial0	150.1.3.3	0	FULL/-	00:00:38	155.1.0.3	Serial0	150.1.8.8	1	FULL/BDR	00:00:31	155.1.58.8	Gi0/0		
Neighbor ID	Pri	State	Dead Time	Address	Interface																																																																														
150.1.4.4	0	FULL/-	00:00:33	155.1.45.4	Serial0																																																																														
Globi-R1	0	FULL/-	00:00:39	155.1.0.1	Serial0																																																																														
150.1.4.4	0	FULL/-	00:00:38	155.1.0.4	Serial0																																																																														
HOBBIT-R2	0	FULL/-	00:00:39	155.1.0.2	Serial0																																																																														
150.1.3.3	0	FULL/-	00:00:38	155.1.0.3	Serial0																																																																														
Neighbor ID	Pri	State	Dead Time	Address	Interface																																																																														
150.1.4.4	0	FULL/-	00:00:33	155.1.45.4	Serial0																																																																														
Globi-R1	0	FULL/-	00:00:39	155.1.0.1	Serial0																																																																														
150.1.4.4	0	FULL/-	00:00:38	155.1.0.4	Serial0																																																																														
HOBBIT-R2	0	FULL/-	00:00:39	155.1.0.2	Serial0																																																																														
150.1.3.3	0	FULL/-	00:00:38	155.1.0.3	Serial0																																																																														
150.1.8.8	1	FULL/BDR	00:00:31	155.1.58.8	Gi0/0																																																																														
<h3>OSPF Database Filtering</h3> <p>Config per interface:</p> <p>Config per neighbor:</p>	<p>Per interface:</p> <pre>interface Vlan79 ip ospf database-filter all out</pre> <p><i>clear ip ospf process!</i></p> <p>Per Neighbor:</p> <pre>router ospf 1 neighbor 155.1.0.2 database-filter all out</pre>	<p>OSPF, ignore MTU while creating OSPF Adjacency:</p>	<pre>interface Port-channel1 ip ospf mtu-ignore</pre>	<p>Parameters that have to match for a OSPF adjacency to come up:</p>	<p>OSPF parameters that have to match for OSPF Adjacency:</p> <ul style="list-style-type: none"> - Timers - Area ID - Authentication - Area Stub Flag - MTU 																																																																														

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

Ensure other routers of segment cannot intercept the OSPF communication:

Using unicast, instead of multicast via the LAN segment so Router 3 "can" not intercept the ospf conversation between R1 and R2

ip ospf network non-broadcast
 router ospf 1
 neighbor 1.1.1.1

ip ospf network non-broadcast
 router ospf 1
 neighbor 1.1.1.2

What does one have to keep in mind when dealing with

ip ospf message-digest-key X

In this situation?

What does ip ospf prefix-suppression do?

Without group pacing, LSAs need to be refreshed frequently and at random intervals. Individual LSA timers require many refresh packets that contain few LSAs.

What does ip ospf prefix-suppression do?

Will hide the transit networks between the routers

Prefixes in routing table:
 1.1.1.0 / 2.2.2.0 and 3.3.3.0

What options are there if you would like to use the FR connection as a backup in case the ethernet segment goes down in respect to OSPF?

If this should be the backup path, make sure its in AREA 0, If not, use a virtual-link. Otherwise no traffic will use the serial link as transit interface.

Virtual-link / or place serial in Area 0

OSPF LSAs on Individual Timers with Group Pacing

LSA Type 3 filtering via area 1 range x.x.x.x y.y.y.y not-advertise

Similar to **area X filter-list [in/out]**

Configuring OSPF Distance

For intra-area / inter-area / external:

```

router ospf 1
distance ospf intra-area 200
distance ospf inter-area 100
distance ospf external 120
  
```

OSPF show ip ospf virtual-links

```

R4#show ip ospf virtual-links
Virtual Link OSPF_VL0 to router 150.1.5.5 is up
Run as demand circuit
DoNotAge LSA allowed.
Transit area 45, via interface GigabitEthernet0/0, Cost of using 10000
Transmit Delay is 1 sec, State POINT_TO_POINT,
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:02
Adjacency State FULL (Hello suppressed)
Index 2/4, retransmission queue length 0, number of retransmission 0
First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
Last retransmission scan length is 0, maximum is 0
Last retransmission scan time is 0 msec, maximum is 0 msec
  
```

OSPF Verifying OSPF fast hello's:

```

R5#show ip ospf interface S0/0/0 | include Timer
Timer intervals configured, Hello 333 msec, Dead 1, Wait 1, Retransmit 5
Interface ser0/0/0
ip ospf dead-interval minimal hello-multiplier 3
  
```

OSPF no discard-route

no discard-route [internal] [external]

If not specified, will remove the route for both.

OSPF and how many routes visible in each area:

Conditional default route

TRACK NOT

(if SLA X down / network x is NOT in the table, advertise the default route)

```

ip sla 53
icmp-echo 172.16.0.22
timeout 500
frequency 3
ip sla schedule 53 life forever start-time now
track 13 rtr 53 reachability
! Define another object that is the negation of the previous object
track 77 list boolean and object 13 not
! Insert a static route if the second object is UP (thus the ! IP SLA probe failed)
ip route 1.0.0.0 255.0.0.0 Null0 track 77
  
```

OSPF Filtering with Distribute-Lists

Config:

Inter-area filtering using distribute lists

```

acl 1 deny 1.0.0.0
acl 1 permit any
router ospf X
distribute-list 1 in
  
```

OSPF Area stub flags and config:

How do you connect AREA 25 to Area 0 ?

OSPF LSA Type 3 from Area 1 to Area 0 will be redistributed (usual rules apply)

Area 0 to 2

Filtering has been applied ->

OSPF

area 1 stub
 area 2 stub no-sum
 area 3 stub nssa
 area 4 nssa default inf
 area 5 nssa no-sum

OSPF

area 0
 area 1 stub
 area 2 stub no-sum
 area 3 stub nssa
 area 4 nssa default inf
 area 5 nssa no-sum

OSPF

area 0
 area 1 stub
 area 2 stub no-sum
 area 3 stub nssa
 area 4 nssa default inf
 area 5 nssa no-sum

OSPF

area 0
 area 1 stub
 area 2 stub no-sum
 area 3 stub nssa
 area 4 nssa default inf
 area 5 nssa no-sum

AREA x nssa translate type7 suppress-fa TO BE DONE

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

<p>IP OSPF Authentication</p> <p>Enable per Interface / Area</p> <p>Apply authentication per Interface</p> <p>Disable authentication per interface</p>	<p>Enable OSPF Authentication:</p> <p>Per Interface Per Area</p> <pre>int ser0/x router ospf 1 ip ospf authentication area 0 authentication message ... ip ospf authentication message ... area 0 authentication message ..</pre> <p>Apply authentication</p> <pre>Int ser0/x ip ospf authentication-key TEXT-TYPE-1 ip ospf message-digest key 1 MD5 PASSWORD</pre> <p>Disable authentication per link</p> <pre>Int ser0/x ip ospf authentication null</pre>	<p>OSPFv2 Cryptographic Authentication</p> <p>key chain TEST</p> <pre>key 1 key-string PASSWORD cryptographic-algorithm hmac-sha-256 send-lifetime {start-time} [infinite end-time] duration <seconds></pre> <p>Interface fa0/x</p> <pre>ip ospf authentication key-chain TEST</pre>	<p>What LSA Types will be announced in each topology, and why?</p>		
<p>OSPF virtual-link Authentication</p> <p>Problem with Area 0 authentication and VLinks</p>	<p>R3</p> <pre>router ospf 1 NO area 0 authentication area 99 virtual-link 2.2.2.2</pre> <p>If the authentication would not explicitly be set to null on the Virtual link, it would try to Auth via Type 2, MD5 to R3 on R2, failing to establish the Vlink.</p> <p>R2#</p> <pre>area 0 authentication area 0 authentication message-digest area 99 virtual-link 3.3.3.3 authentication null</pre>	<p>OSPFv3 IPsec ESP Encryption and Authentication</p> <p>Interface fa0/x</p> <pre>ospfv3 encryption {ipsec spi spi esp encryption-algorithm key-encryption-type key authentication-algorithm key- encryption-type key null} ipv6 ospf encryption {ipsec spi spi esp {encryption- algorithm [[key-encryption-type] key] null} authentication-algorithm [key-encryption-type] key null}</pre> <pre>ospfv3 encryption ipsec spi 1001 esp null md5 0 2757... ipv6 ospf encryption ipsec spi 1001 esp null sha1 1234...</pre>	<p>Explain why it is essential that OSPF router-id's should be kept unique within the entire OSPF domain?</p>		<p>R9: Show ip ospf database external LS-ID: 4.4.4.0/24 Advertising Router 0.0.0.3</p> <p>If R9 had RID of 0.0.0.3 in Area 2, 4.4.4.0/24 would be filtered out!</p>
<p>OSPF Default route</p> <p>3 different methods</p>	<pre>1. ip route 0.0.0.0 0.0.0.0 192.168.1.1 router ospf 1 default-information originate 2. router ospf 1 default-information originate always 3. router ospf 1 default-information originate route-map RMP-X route-map RMP-X match ip address 77 access-list 77 permit 10.0.0.1</pre>	<p>OSPF multiarea</p> <p>authentication mode {strict deployment normal}</p>	<p>OSPF multiarea</p>	<p>What segment announces which LSA Type?</p>	<p>What segment announces which LSA Type?</p> <p>Type 1 LSA</p> <p>DR, Type 2 LSA</p>
<p>OSPF BFD</p> <p>BFD is configured globally for all interfaces associated with the OSPF process.</p> <pre>interface Fast Ethernet 0/1 ip address x.x.x.x bfd interval 50 min_rx 50 multiplier 3 interface Fast Ethernet 0/2 ip address u.u.u.u</pre> <p>router ospf 123</p> <pre>log-adjacency-changes detail network x.x.x.x 0.0.0.0 area 0 network u.u.u.u 0.0.0.0 area 0 bfd all-interfaces</pre> <p>show bfd neighbors details</p>	<p>How can you make R1 and R2 form a working OSPF adjacency without changing the Subnet mask on the Ethernet interface?</p> <pre>int fa0/0 ip address 10.0.0.1 255.255.255.0 int fa0/0 ip address 10.0.0.2 255.255.254.0 router ospf 1 network 10.0.0.1 0.0.0.0 area 0 network 10.0.0.2 0.0.0.0 area 0</pre>	<p>Change the ospf network type to P-2-P</p> <p>-> will ignore the mis-matched subnet mask!</p> <pre>int fa0/0 ip ospf network point-to-point</pre> <pre>router ospf 1 network 10.0.0.1 0.0.0.0 area 0 network 10.0.0.2 0.0.0.0 area 0</pre>	<p>What could be a potential problem here?</p> <p>Something changed here</p> <p>OSPF priority is the same on all routers</p>	<p>Something changed here</p> <p>R1 sends update to DR of the Segment!</p> <p>R3 has the highest IP on the common segment and is the DR. R2 will never hear the update unless there is a static mapping from R3 to R2!</p>	
<p>IPsec on OSPFv3</p> <pre>ospfv3 authentication {ipsec spi} {md5 sha1} {key- encryption-type key} null ipv6 ospf authentication {null ipsec spi spi authentication-algorithm [key-encryption-type] [key]}</pre> <pre>interface fa0/x ospfv3 authentication md5 0 2757613409...09727 Or ipv6 ospf authentication ipsec spi 500 md5 123456789</pre>	<p>What if DB filter / distribute list / prefix-suppression in operation?</p> <p>This router needs to see this subnet!</p>	<p>Check the ospf database for the network and the forward address!</p> <p>This router needs to see this subnet!</p>	<p>Who is going to be the DR ?</p> <p>ip ospf priority all are set to the same</p>	<p>If the OSPF priority is the same, the higher IP wins!</p> <p>ip ospf priority all are set to the same</p> <p>(R1 should be the DR in this situation)</p>	
<p>IPv6 OSPF Null authentication</p> <pre>conf t interface fa0/x ospfv3 authentication null ipv6 ospf authentication null</pre>	<p>What could be the problem here?</p> <p>All routers have:</p> <pre>Interface e0/0 ip ospf priority 0</pre>	<p>What could be the problem here?</p>	<p>Explain how OSPF LSA Type 1 flooding works?</p> <p>Something changed here</p>	<p>OSPF LSA Type 1 flooding</p> <p>Something changed here</p>	

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

Explain output of:

R1#show ip ospf border-routers

```
R1#show ip ospf border-routers
OSPF Router with ID (0.0.0.1) (Process ID 1)

Base Topology (MTID 0)

Internal Router Routing Table
Codes: I - Intra-area route, I - Inter-area route

I 0.0.0.2 [10] via 12.1.1.2, e0/0, ABR, Area 1, SPF 12
I 0.0.0.4 [75] via 12.1.1.2, e0/0, ASBR, Area 1, SPF 12
```

Cost to the ASBR

Explain output of:

R1#show ip ospf border-routers

```
R1#show ip ospf border-routers
OSPF Router with ID (0.0.0.1) (Process ID 1)

Base Topology (MTID 0)

Internal Router Routing Table
Codes: I - Intra-area route, I - Inter-area route

I 0.0.0.2 [10] via 12.1.1.2, e0/0, ABR, Area 1, SPF 12
I 0.0.0.3 [20] via 12.1.1.2, e0/0, ASBR, Area 1, SPF 12
```

Notice R3 is the ASBR once Area 2 converted to NSSA!

Explain output of:

R2#show ip ospf database (focus on LSA type 1)

```
R2#show ip ospf database
OSPF Router with ID (0.0.0.2) (Process ID 1)

Router Link States (Area 0)
Link ID ADV Router Age Seq# Checksum Link count
0.0.0.2 0.0.0.2 800 0x8000000E 0x005190 1
0.0.0.3 0.0.0.3 1198 0x8000000F 0x004D90 1

Router Link States (Area 1)
Link ID ADV Router Age Seq# Checksum Link count
0.0.0.1 0.0.0.1 533 0x80000015 0x007B79 1
0.0.0.2 0.0.0.2 795 0x8000001B 0x00707A 1
```

Explain output of:

R1#show ip ospf database adv-router 0.0.0.4

```
R1#show ip ospf database adv-router 0.0.0.4
OSPF Router with ID (0.0.0.1) (Process ID 1)

Type-5 AS External Link States
Link ID ADV Router Age Seq# Checksum Tag
9.9.9.0 0.0.0.4 1614 0x80000002 0x003D52 0
```

Redistributed prefix Router-ID of R4

Explain output of:

R1#show ip ospf database adv-router 0.0.0.4
R1#show ip ospf database adv-router 0.0.0.3

```
R1#show ip ospf database adv-router 0.0.0.4
OSPF Router with ID (0.0.0.1) (Process ID 1)

Type-5 AS External Link States
Link ID ADV Router Age Seq# Checksum Tag
9.9.9.0 0.0.0.3 265 0x80000001 0x000E5B 0
```

Explain output of:

R2#show ip ospf database (focus on LSA type 4,5)

```
R2#show ip ospf database (partial output, looking at LSA 4, 5)

Summary ASB Link States (Area 0)
Link ID ADV Router Age Seq# Checksum
0.0.0.4 0.0.0.3 693 0x80000002 0x006795

Summary ASB Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
0.0.0.4 0.0.0.2 296 0x80000001 0x00D321

Type-5 AS External Link States
Link ID ADV Router Age Seq# Checksum Tag
9.9.9.0 0.0.0.4 331 0x80000002 0x0088F7 0
```

Explain output of:

R1#show ip ospf database external

```
R1#show ip ospf database external
OSPF Router with ID (0.0.0.1) (Process ID 1)

Type-5 AS External Link States
Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 1979
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 9.9.9.0 (External Network Number)
Advertising Router: 0.0.0.4
LS Seq Number: 80000002
Checksum: 0x3D52
Length: 36
Network Mask: /24
Metric Type: 2 (Larger than any link state path)
MTID: 0
Metric: 20
Forward Address: 0.0.0.0
External Route Tag: 0
```

What is the effective forward metric to 9.9.9.0? Its NOT 20!

Explain output of:

R1#show ip ospf database external

```
R1#show ip ospf database external
OSPF Router with ID (0.0.0.1) (Process ID 1)

Type-5 AS External Link States
Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 431
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 9.9.9.0 (External Network Number)
Advertising Router: 0.0.0.3
LS Seq Number: 80000001
Checksum: 0x5B
Length: 36
Network Mask: /24
Metric Type: 2 (Larger than any link state path)
MTID: 0
Metric: 20
Forward Address: 34.1.1.4
External Route Tag: 0
```

Explain output of:

R1#show ip ospf database external

Explain output of:

R1#show ip ospf database external

router ospf x summary-address 4.0.0.0 255.0.0.0

Either on R4 or R3

How can you calculate the forward metric from R1 to 9.9.9.0/24?

What show commands are needed?

```
R1#show ip route 9.9.9.0
Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 75

R1#show ip ospf database external
Link State ID: 9.9.9.0 (External Network Number)
Advertising Router: 0.0.0.4
Forward Address: 0.0.0.0 (0.0.0.0 = set to self)

R1#show ip ospf database asbr-summary
Link State ID: 0.0.0.4 (AS Boundary Router address)
Advertising Router: 0.0.0.2
MTID: 0 Metric: 65 From ABR Area1 (R2) to ASBR

R1#sh ip ospf database router 0.0.0.2
Link State ID: 0.0.0.2
Advertising Router: 0.0.0.2
TOS 0 Metrics: 10 From R1 to ABR Area 1
```

How can you calculate the forward metric from R1 to 9.9.9.0/24?

What show commands are needed?

```
R1#show ip route 9.9.9.0
Known via "ospf 1", distance 110, metric 20, type extern 2, forward metric 75

R1#show ip ospf database external
Link State ID: 9.9.9.0 (External Network Number)
Advertising Router: 0.0.0.3
Forward Address: 34.1.1.4

R1#show ip ospf database asbr-summary
Link State ID: 0.0.0.3 (AS Boundary Router address)
Advertising Router: 0.0.0.2
MTID: 0 Metric: 10

R1#show ip ospf database summary 34.1.1.0
Link State ID: 34.1.1.0 (summary Network Number)
Advertising Router: 0.0.0.2
Metric: 65
```

Where can you summarize these internal AREA 2 routes?

```
R1#show ip route ospf | 1.4
O IA 4.0.0.0/8 [110/76] via 12.1.1.2, 00:05:54, e0/0

R3#
router ospf 1
area 2 range 4.0.0.0 255.0.0.0 [not-advertise]
```

What info will the following provide:

R1#show ip ospf database adv-router 0.0.0.2

```
R1#show ip ospf database adv-router 0.0.0.2
OSPF Router with ID (0.0.0.1) (Process ID 1)

Router Link States (Area 1)
Link ID ADV Router Age Seq# Checksum Link count
0.0.0.2 0.0.0.2 1774 0x8000000E 0x008A6D 1

Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
12.1.1.2 0.0.0.2 1774 0x80000003 0x00E939

Summary Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
23.1.1.0 0.0.0.2 536 0x80000003 0x009A7B
34.0.0.0 0.0.0.2 1774 0x80000002 0x004C8A

Summary ASB Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
0.0.0.4 0.0.0.2 43 0x80000005 0x00C825
```

What info will the following provide:

R1#show ip ospf database adv-router 0.0.0.2

```
R1#show ip ospf database adv-router 0.0.0.2
OSPF Router with ID (0.0.0.1) (Process ID 1)

Router Link States (Area 1)
Link ID ADV Router Age Seq# Checksum Link count
0.0.0.2 0.0.0.2 222 0x80000010 0x00866F 1

Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
12.1.1.2 0.0.0.2 222 0x80000005 0x00E53B

Summary Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
23.1.1.0 0.0.0.2 989 0x80000004 0x00987C
34.0.0.0 0.0.0.2 222 0x80000004 0x00488C

Summary ASB Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
0.0.0.3 0.0.0.2 1147 0x80000001 0x00B577
```

What is the potential problem of using OSPF Point-to-Multipoint in this scenario with DMVPN?

OSPF Point-to-Multipoint

Underlying NBMA access rates: 10, 100

10.0.0.0/8

OSPF Point-to-Multipoint

Underlying NBMA access rates: 10, 100

10.0.0.0/8

Use P-2-M non-broadcast to solve!

What info will the following provide:

R1# show ip ospf database

```
R1#show ip ospf database
OSPF Router with ID (0.0.0.1) (Process ID 1)

Router Link States (Area 1)
Link ID ADV Router Age Seq# Checksum Link count
0.0.0.1 0.0.0.1 1952 0x80000007 0x009F6B 1
0.0.0.2 0.0.0.2 6 0x8000000F 0x00886E 1

Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
12.1.1.2 0.0.0.2 6 0x80000004 0x00E73A

Summary Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
23.1.1.0 0.0.0.2 748 0x80000003 0x009A7B
34.0.0.0 0.0.0.2 6 0x80000003 0x004A8B

Summary ASB Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
0.0.0.4 0.0.0.2 254 0x80000005 0x00C825

Type-5 AS External Link States
Link ID ADV Router Age Seq# Checksum Tag
9.9.9.0 0.0.0.4 1474 0x80000003 0x003B53 0
```

What info will the following provide:

R1#show ip ospf database

```
R1#show ip ospf database
OSPF Router with ID (0.0.0.1) (Process ID 1)

Router Link States (Area 1)
Link ID ADV Router Age Seq# Checksum Link count
0.0.0.1 0.0.0.1 481 0x80000009 0x00936D 1
0.0.0.2 0.0.0.2 548 0x80000010 0x00866F 1

Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
12.1.1.2 0.0.0.2 548 0x80000005 0x00E53B

Summary Net Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
23.1.1.0 0.0.0.2 1315 0x80000004 0x00987C
34.0.0.0 0.0.0.2 548 0x80000004 0x00488C

Summary ASB Link States (Area 1)
Link ID ADV Router Age Seq# Checksum
0.0.0.3 0.0.0.2 1474 0x80000001 0x00B577

Type-5 AS External Link States
Link ID ADV Router Age Seq# Checksum Tag
9.9.9.0 0.0.0.3 1469 0x80000001 0x000E5B 0
```

How can you solve this situation, accounting for the unequal, physical access rates using OSPF?

OSPF Point-to-Multipoint

Underlying NBMA access rates: 10, 100

10.0.0.0/8

OSPF Point-to-Multipoint

Underlying NBMA access rates: 10, 100

10.0.0.0/8

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Colin

<p>What info will the following provide:</p> <p>R1#show ip ospf database external</p>	<p>R1#show ip ospf database external OSPF Router with ID (0.0.0.1) (Process ID 1) Type-5 AS External Link States</p> <p>Routing Bit Set on this LSA in topology Base with MTID 0 LS age: 28 Options: (No TOS-capability, DC) LS Type: AS External Link Link State ID: 9.9.9.0 (External Network Number) Advertising Router: 0.0.0.4 LS Seq Number: 80000001 Checksum: 0x8AF6 Length: 36 Network Mask: /24 Metric Type: 2 (Larger than any link state path) MTID: 0 Metric: 20 Forward Address: 0.0.0.0 External Route Tag: 0</p>	<p>What info will the following provide:</p> <p>R1#show ip ospf database external</p>	<p>R1#show ip ospf database external OSPF Router with ID (0.0.0.1) (Process ID 1) Type-5 AS External Link States</p> <p>Routing Bit Set on this LSA in topology Base with MTID 0 LS age: 28 Options: (No TOS-capability, DC) LS Type: AS External Link Link State ID: 9.9.9.0 (External Network Number) Advertising Router: 0.0.0.3 LS Seq Number: 80000001 Checksum: 0x5901 Length: 36 Network Mask: /24 Metric Type: 2 (Larger than any link state path) MTID: 0 Metric: 20 Forward Address: 34.1.1.4 External Route Tag: 0</p>	<p>R1#show ip ospf database external</p> <p>connections = e0/x ip ospf network broadcast</p>	<p>R1#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.4 Network Mask: /24 Metric: 20 Forward Address: 0.0.0.0</p>																																																																																																																																							
<p>What info will the following provide:</p> <p>R1#show ip ospf events</p>	<p>show ip ospf events: clear ip ospf events</p> <p>Network 9.9.9.0 from UP to DOWN status: Timer Exp: if_ack_delayed 0xAB5314C0 RIB Delete, Topo Base, dest 9.9.9.0, mask 255.255.255.0, gw 12.1.1.2, via Ethernet0/0, source 0.0.0.4, type Ext2 Insert MAXAGE lsa: 0xAAA423E8 9.9.9.0 Rcv Changed Type-5 LSA, LSID 9.9.9.0, Adv-Rtr 0.0.0.4, Seq# 80000002, Age 3600</p> <p>Network 9.9.9.0 from DOWN to UP status: Timer Exp: if_ack_delayed 0xAB5314C0 RIB Replace, Topo Base, dest 9.9.9.0, mask 255.255.255.0, gw 12.1.1.2, via Ethernet0/0, source 0.0.0.4, type Ext2 Rcv New Type-5 LSA, LSID 9.9.9.0, Adv-Rtr 0.0.0.4, Seq# 80000001, Age 3 DB add: 9.9.9.0 0x4223E8 175</p>	<p>Output of</p> <p>show ip ospf topology-info</p>		<p>R2#show ip ospf database external</p> <p>connections = e0/x ip ospf network broadcast</p>	<p>R2#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.3 Network Mask: /24 Metric: 20 Forward Address: 0.0.0.0</p>																																																																																																																																							
<p>R1-4#show ip ospf database external Link State ID: 4.4.4.0 (External Network Number) connections = ethernet / broadcast!</p>	<p>R1-4#show ip ospf database external Link State ID: 4.4.4.0 (External Network Number)</p> <p>R1# Advertising Router: 0.0.0.4 Metric: 20 Forward Address: 0.0.0.0</p> <p>R2# Advertising Router: 0.0.0.4 Metric: 20 Forward Address: 24.1.1.4</p> <p>R3# Advertising Router: 0.0.0.4 Metric: 20 Forward Address: 0.0.0.0</p>	<p>R1-4#show ip ospf database external Link State ID: 4.4.4.0 (External Network Number) connections = ethernet / point-to-point</p>	<p>R1-4#show ip ospf database external Link State ID: 4.4.4.0 (External Network Number)</p> <p>R1# Advertising Router: 0.0.0.4 Metric: 20 Forward Address: 0.0.0.0</p> <p>R2# Advertising Router: 0.0.0.4 Metric: 20 Forward Address: 24.1.1.4</p> <p>R3# Advertising Router: 0.0.0.4 Metric: 20 Forward Address: 0.0.0.0</p>	<p>R2#show ip ospf database external</p> <p>connections = e0/x ip ospf network broadcast</p>	<p>R2#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.3 Network Mask: /24 Metric: 20 Forward Address: 24.1.1.4</p>																																																																																																																																							
<p>R1-4#show ip ospf database external Link State ID: 4.4.4.0 (External Network Number) connections = ethernet / broadcast!</p>	<p>R1#show ip route ospf b 4.4.4.0 O E2 4.4.4.0 [110/20] via 13.1.1.3, 00:00:02, e0/1 [110/20] via 12.1.1.2, 00:00:02, e0/0</p> <p>R1#show ip ospf database Summary ASB Link States (Area 0)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.4</td> <td>0.0.0.2</td> <td>61</td> <td>0x80000001</td> <td>0x00A8B0</td> </tr> <tr> <td>0.0.0.4</td> <td>0.0.0.3</td> <td>49</td> <td>0x80000003</td> <td>0x00A187</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>57</td> <td>0x80000001</td> <td>0x003F51 0</td> </tr> </tbody> </table> <p>R1#show ip route ospf b 4.4.4.0 O E2 4.4.4.0 [110/20] via 12.1.1.2, 00:00:03, e0/0</p> <p>R1#show ip ospf database Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>33</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.4	0.0.0.2	61	0x80000001	0x00A8B0	0.0.0.4	0.0.0.3	49	0x80000003	0x00A187	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	57	0x80000001	0x003F51 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	33	0x80000001	0x009FD3 0	<p>R1-4#show ip ospf database external Link State ID: 4.4.4.0 (External Network Number) connections = e0/x ip ospf network point-to-point</p>	<p>R1#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 13.1.1.3, 00:23:45, Ethernet0/1 [110/20] via 12.1.1.2, 00:23:45, Ethernet0/0</p> <p>R1#show ip ospf database Summary ASB Link States (Area 0)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.4</td> <td>0.0.0.2</td> <td>1446</td> <td>0x80000001</td> <td>0x00A8B0</td> </tr> <tr> <td>0.0.0.4</td> <td>0.0.0.3</td> <td>1441</td> <td>0x80000001</td> <td>0x00A585</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>1451</td> <td>0x80000001</td> <td>0x003F51 0</td> </tr> </tbody> </table> <p>R1#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 12.1.1.2, 00:01:04, Ethernet0/0</p> <p>R1#show ip ospf database Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>108</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.4	0.0.0.2	1446	0x80000001	0x00A8B0	0.0.0.4	0.0.0.3	1441	0x80000001	0x00A585	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	1451	0x80000001	0x003F51 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	108	0x80000001	0x009FD3 0	<p>R2#show ip ospf database external</p> <p>connections = e0/x ip ospf network point-to-point</p>	<p>R1#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.4 Network Mask: /24 Metric: 20 Forward Address: 0.0.0.0</p> <p>R1#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.3 Network Mask: /24 Metric: 20 Forward Address: 24.1.1.4</p>																																																																	
Link ID	ADV Router	Age	Seq#	Checksum																																																																																																																																								
0.0.0.4	0.0.0.2	61	0x80000001	0x00A8B0																																																																																																																																								
0.0.0.4	0.0.0.3	49	0x80000003	0x00A187																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	57	0x80000001	0x003F51 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	33	0x80000001	0x009FD3 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum																																																																																																																																								
0.0.0.4	0.0.0.2	1446	0x80000001	0x00A8B0																																																																																																																																								
0.0.0.4	0.0.0.3	1441	0x80000001	0x00A585																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	1451	0x80000001	0x003F51 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	108	0x80000001	0x009FD3 0																																																																																																																																								
<p>R2#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 24.1.1.4, 00:00:00, Ethernet0/1</p> <p>R2#show ip ospf database Summary ASB Link States (Area 0)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.4</td> <td>0.0.0.2</td> <td>86</td> <td>0x80000001</td> <td>0x00A8B0</td> </tr> <tr> <td>0.0.0.4</td> <td>0.0.0.3</td> <td>76</td> <td>0x80000001</td> <td>0x00A585</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>83</td> <td>0x80000001</td> <td>0x003F51 0</td> </tr> </tbody> </table> <p>R2#show ip route b 4.4.4.0 O N2 4.4.4.0 [110/20] via 24.1.1.4, 00:00:03, Ethernet0/1</p> <p>R2#show ip ospf database Type-7 AS External Link States (Area 1)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>39</td> <td>0x80000001</td> <td>0x000563 0</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>22</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.4	0.0.0.2	86	0x80000001	0x00A8B0	0.0.0.4	0.0.0.3	76	0x80000001	0x00A585	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	83	0x80000001	0x003F51 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	39	0x80000001	0x000563 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	22	0x80000001	0x009FD3 0	<p>R2#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 24.1.1.4, 00:00:09, Ethernet0/1</p> <p>R2#show ip ospf database Type-7 AS External Link States (Area 1)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>48</td> <td>0x80000001</td> <td>0x000563 0</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>30</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	48	0x80000001	0x000563 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	30	0x80000001	0x009FD3 0	<p>R2#show ip ospf database external</p> <p>connections = e0/x ip ospf network point-to-point</p>	<p>R2#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 24.1.1.4, 00:00:08, Ethernet0/1</p> <p>R2#show ip ospf database Summary ASB Link States (Area 0)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.4</td> <td>0.0.0.2</td> <td>42</td> <td>0x80000001</td> <td>0x00A8B0</td> </tr> <tr> <td>0.0.0.4</td> <td>0.0.0.3</td> <td>37</td> <td>0x80000001</td> <td>0x00A585</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>51</td> <td>0x80000001</td> <td>0x003F51 0</td> </tr> </tbody> </table> <p>R2#show ip route b 4.4.4.0 O N2 4.4.4.0 [110/20] via 24.1.1.4, 00:00:09, Ethernet0/1</p> <p>R2#show ip ospf database Type-7 AS External Link States (Area 1)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>48</td> <td>0x80000001</td> <td>0x000563 0</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>30</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.4	0.0.0.2	42	0x80000001	0x00A8B0	0.0.0.4	0.0.0.3	37	0x80000001	0x00A585	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	51	0x80000001	0x003F51 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	48	0x80000001	0x000563 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	30	0x80000001	0x009FD3 0	<p>R2#show ip ospf database external</p> <p>connections = e0/x ip ospf network point-to-point</p>	<p>R2#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.4 Network Mask: /24 Metric: 20 Forward Address: 0.0.0.0</p> <p>R2#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.3 Network Mask: /24 Metric: 20 Forward Address: 24.1.1.4</p>																									
Link ID	ADV Router	Age	Seq#	Checksum																																																																																																																																								
0.0.0.4	0.0.0.2	86	0x80000001	0x00A8B0																																																																																																																																								
0.0.0.4	0.0.0.3	76	0x80000001	0x00A585																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	83	0x80000001	0x003F51 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	39	0x80000001	0x000563 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	22	0x80000001	0x009FD3 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	48	0x80000001	0x000563 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	30	0x80000001	0x009FD3 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum																																																																																																																																								
0.0.0.4	0.0.0.2	42	0x80000001	0x00A8B0																																																																																																																																								
0.0.0.4	0.0.0.3	37	0x80000001	0x00A585																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	51	0x80000001	0x003F51 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	48	0x80000001	0x000563 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	30	0x80000001	0x009FD3 0																																																																																																																																								
<p>R3#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 34.1.1.4, 00:00:01, Ethernet0/0</p> <p>R3#show ip ospf database Summary ASB Link States (Area 0)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.4</td> <td>0.0.0.2</td> <td>34</td> <td>0x80000003</td> <td>0x00A782</td> </tr> <tr> <td>0.0.0.4</td> <td>0.0.0.3</td> <td>35</td> <td>0x80000001</td> <td>0x00A585</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>43</td> <td>0x80000001</td> <td>0x003F51 0</td> </tr> </tbody> </table> <p>R3#show ip route b 4.4.4.0 O N2 4.4.4.0 [110/20] via 34.1.1.4, 00:00:10, Ethernet0/0</p> <p>R3#show ip ospf database Type-7 AS External Link States (Area 1)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>36</td> <td>0x80000001</td> <td>0x000563 0</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>28</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.4	0.0.0.2	34	0x80000003	0x00A782	0.0.0.4	0.0.0.3	35	0x80000001	0x00A585	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	43	0x80000001	0x003F51 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	36	0x80000001	0x000563 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	28	0x80000001	0x009FD3 0	<p>R3#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 34.1.1.4, 00:00:07, Ethernet0/0</p> <p>R3#show ip ospf database Summary ASB Link States (Area 0)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.4</td> <td>0.0.0.2</td> <td>53</td> <td>0x80000001</td> <td>0x00A8B0</td> </tr> <tr> <td>0.0.0.4</td> <td>0.0.0.3</td> <td>43</td> <td>0x80000001</td> <td>0x00A585</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>58</td> <td>0x80000001</td> <td>0x003F51 0</td> </tr> </tbody> </table> <p>R3#show ip route b 4.4.4.0 O N2 4.4.4.0 [110/20] via 34.1.1.4, 00:00:09, Ethernet0/0</p> <p>R3#show ip ospf database Type-7 AS External Link States (Area 1)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>50</td> <td>0x80000001</td> <td>0x000563 0</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>35</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.4	0.0.0.2	53	0x80000001	0x00A8B0	0.0.0.4	0.0.0.3	43	0x80000001	0x00A585	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	58	0x80000001	0x003F51 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	50	0x80000001	0x000563 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	35	0x80000001	0x009FD3 0	<p>R3#show ip ospf database external</p> <p>connections = e0/x ip ospf network point-to-point</p>	<p>R3#show ip route b 4.4.4.0 O E2 4.4.4.0 [110/20] via 34.1.1.4, 00:00:07, Ethernet0/0</p> <p>R3#show ip ospf database Summary ASB Link States (Area 0)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum</th> </tr> </thead> <tbody> <tr> <td>0.0.0.4</td> <td>0.0.0.2</td> <td>53</td> <td>0x80000001</td> <td>0x00A8B0</td> </tr> <tr> <td>0.0.0.4</td> <td>0.0.0.3</td> <td>43</td> <td>0x80000001</td> <td>0x00A585</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>58</td> <td>0x80000001</td> <td>0x003F51 0</td> </tr> </tbody> </table> <p>R3#show ip route b 4.4.4.0 O N2 4.4.4.0 [110/20] via 34.1.1.4, 00:00:09, Ethernet0/0</p> <p>R3#show ip ospf database Type-7 AS External Link States (Area 1)</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.4</td> <td>50</td> <td>0x80000001</td> <td>0x000563 0</td> </tr> </tbody> </table> <p>Type-5 AS External Link States</p> <table border="1"> <thead> <tr> <th>Link ID</th> <th>ADV Router</th> <th>Age</th> <th>Seq#</th> <th>Checksum Tag</th> </tr> </thead> <tbody> <tr> <td>4.4.4.0</td> <td>0.0.0.3</td> <td>35</td> <td>0x80000001</td> <td>0x009FD3 0</td> </tr> </tbody> </table>	Link ID	ADV Router	Age	Seq#	Checksum	0.0.0.4	0.0.0.2	53	0x80000001	0x00A8B0	0.0.0.4	0.0.0.3	43	0x80000001	0x00A585	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	58	0x80000001	0x003F51 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.4	50	0x80000001	0x000563 0	Link ID	ADV Router	Age	Seq#	Checksum Tag	4.4.4.0	0.0.0.3	35	0x80000001	0x009FD3 0	<p>R3#show ip ospf database external</p> <p>connections = e0/x ip ospf network point-to-point</p>	<p>R3#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.4 Network Mask: /24 Metric: 20 Forward Address: 0.0.0.0</p> <p>R3#show ip ospf database external LS Type: AS External Link Link State ID: 4.4.4.0 (External Network Number) Advertising Router: 0.0.0.3 Network Mask: /24 Metric: 20 Forward Address: 24.1.1.4</p>
Link ID	ADV Router	Age	Seq#	Checksum																																																																																																																																								
0.0.0.4	0.0.0.2	34	0x80000003	0x00A782																																																																																																																																								
0.0.0.4	0.0.0.3	35	0x80000001	0x00A585																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	43	0x80000001	0x003F51 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	36	0x80000001	0x000563 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	28	0x80000001	0x009FD3 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum																																																																																																																																								
0.0.0.4	0.0.0.2	53	0x80000001	0x00A8B0																																																																																																																																								
0.0.0.4	0.0.0.3	43	0x80000001	0x00A585																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	58	0x80000001	0x003F51 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	50	0x80000001	0x000563 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	35	0x80000001	0x009FD3 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum																																																																																																																																								
0.0.0.4	0.0.0.2	53	0x80000001	0x00A8B0																																																																																																																																								
0.0.0.4	0.0.0.3	43	0x80000001	0x00A585																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	58	0x80000001	0x003F51 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.4	50	0x80000001	0x000563 0																																																																																																																																								
Link ID	ADV Router	Age	Seq#	Checksum Tag																																																																																																																																								
4.4.4.0	0.0.0.3	35	0x80000001	0x009FD3 0																																																																																																																																								

I used ethernet connections instead of serial connections!
Int e0/0
ip ospf network broadcast

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

What options are available to filter out 1.0.0.0/24 on R4 ?

R4 = ASBR
Redistr. static
1.0.0.0/24
11.0.0.0/24
100.0.0.0/24
RID: 0.0.0.4

Option 1:
R4#
router ospf 1
summary-address 1.0.0.0 255.255.255.0 not-advertise

Option 2:
ip prefix-list PFX-1 deny 1.0.0.0/24
ip prefix-list PFX-1 permit 0.0.0.0/0 le 32
router ospf 1
distribute-list prefix PFX-1 out

Redistr. static
0.0.0.0/0
1.0.0.0/24
11.0.0.0/24
100.0.0.0/24
RID: 0.0.0.4

What should be the expected output in regards to the Virtual Link in this example of the following show command?

show ip ospf database

R2#show ip ospf database
OSPF Router with ID (0.0.0.2) (Process ID 1)

Router Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum	Link count
0.0.0.1	0.0.0.1	1	(DNA)	0x800002	0x00C646 1
0.0.0.2	0.0.0.2	300		0x800004	0x00088E 3
0.0.0.3	0.0.0.3	559		0x800003	0x0015A4 2

Summary Net Link States (Area 0)

Link ID	ADV Router	Age	Seq#	Checksum
1.1.1.1	0.0.0.1	6	(DNA)	0x800001
1.1.1.1	0.0.0.2	547		0x800002
10.1.1.1	0.0.0.1	6	(DNA)	0x800001
12.1.1.0	0.0.0.1	6	(DNA)	0x800001
12.1.1.0	0.0.0.2	547		0x800002

What does one have to remember about Virtual-links and OSPF Authentication?

If you change authentication type here, the vlinks will not "updating" this info due to their demand circuit behaviour. Bounce an interface to force a Area 0 topology change in order to verify the VLINKs status. If necessary add authentication or set it to Null.

What will be seen by using the following command:
R1#show ip ospf database external 5.5.5.5

Redistr. static
0.0.0.0/0
5.5.5.5

R1#show ip ospf database external 5.5.5.5

OSPF Router with ID (0.0.0.1) (Process ID 1)

Type-5 AS External Link States

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 121
Options: (No TOS-capability, DC)
LS Type: AS External Link
Link State ID: 5.5.5.5 (External Network Number)
Advertising Router: 0.0.0.3
LS Seq Number: 80000003
Checksum: 0x9AC6
Length: 36
Network Mask: /32
Metric Type: 2 (Larger than any link state path)
MTID: 0
Metric: 20
Forward Address: 34.0.0.4
External Route Tag: 0

How to verify OSPF authentication? Type 1

Config

```
router ospf 1
router-id 0.0.0.2
area 0 authentication

interface Serial1/1
ip address 12.1.1.2 255.255.255.0
ip ospf authentication-key Cisco
```

Verification:

R2#show ip ospf interface | i Serial | authentication

Serial1/3 is up, line protocol is up
Simple password authentication enabled
Serial1/1 is up, line protocol is up
Simple password authentication enabled

How do you configure a new MD5 password/key without interrupting traffic? How can you verify it?

interface Serial1/1
ip ospf message-digest-key 1 md5 ccie

interface Serial1/2
ip ospf message-digest-key 1 md5 ccie

R2#show ip ospf interface ser 1/1
Message digest authentication enabled
Youngest key id is 2
Rollover in progress, 1 neighbor(s) using the old key(s):
key id 1

Once you see this:
R2#show ip ospf interface ser 1/1
Message digest authentication enabled
Youngest key id is 2

R1 and R2(config)#int ser 1/1
R1 and R2(config-if)#no ip ospf message-digest-key 1 md5 ccie

Where would you place the following command in order for R1 to see its effect?
area 2 nssa translate type7 suppress-fa

Redistr. static
0.0.0.0/0
5.5.5.5

R3#
router ospf 1
area 2 nssa translate type7 suppress-fa

show ip ospf database external 5.5.5.5

Without suppress-fa:
R1#
Advertising Router: 0.0.0.3
Forward Address: 34.0.0.4
Metric: 20

Using suppress-fa on R3:
R1#
Advertising Router: 0.0.0.3
Forward Address: 0.0.0.0
Metric: 20

How to verify OSPF authentication? Type 2

Config

```
router ospf 1
router-id 0.0.0.2
area 0 authentication message-digest

interface Serial1/3
ip address 23.1.1.2 255.255.255.0
ip ospf message-digest-key 1 md5 Cisco
```

Verification:

R2#show ip osp interface | i Serial | authentication | key

Serial1/3 is up, line protocol is up
Message digest authentication enabled
Youngest key id is 1

You are here: And want to know all prefixes R4 is attached to. What command would you use?

R1#show ip ospf database router adv-router 0.0.0.4

OSPF Router with ID (0.0.0.1) (Process ID 1)

Router Link States (Area 1)

LS age: 590
Options: (No TOS-capability, DC)
LS Type: Router Links
Link State ID: 0.0.0.4
Advertising Router: 0.0.0.4
LS Seq Number: 80000005
Checksum: 0xD919
Length: 120
Number of Links: 8
Link connected to: a Stub Network (Link ID) Network/subnet number: 1.0.0.0 (Link Data) Network Mask: 255.255.255.0
Number of MTID metrics: 0
TOS 0 Metrics: 1
....

Stub Network, No other OSPF neighbor attached!

What will show up in R6 ospf database, routing table for 0.0.0.0 ?

Redistr. static
0.0.0.0/0
5.5.5.5

R6#show ip route
O*IA 0.0.0.0/0 [110/11] via 34.0.0.3, 00:04:05, e0/0

R6#show ip ospf database

Summary Net Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum
0.0.0.0	0.0.0.3	280		0x80000001
				0x00270C

Explain the output of the following command with Type 2 Auth: debug ip ospf packet

TTL set to 1
Length of message 48 bytes
Router-id of the received packet.

OSPF Version 2

OSPF-1 PAK : rcv. v:2 t:1 l:48 rid:0.0.0.1 aid:0.0.0.0 chk:0 aut:2
keyid:1 seq:0x539A6A65 from Serial1/1

Authentication key
Authentication type 1

How can you filter out 1.0.0.0/24 on R1 located on R4 based on distance within the same OSPF area?

R1#
router ospf 1
distance 255 0.0.0.4 0.0.0.0 99

access-list 99 permit 1.0.0.0 0.0.0.255

You can specify the Router-ID here of R4!!!

What is mandatory for the following command to work correctly?
router ospf 1
area 2 nssa no-summary

Redistr. static
0.0.0.0/0
5.5.5.5

The router needs to have a minimum of one interface within Area 0 for "no-summary" to work!

router ospf 1
area 2 nssa no-summary

Explain the output of the following command with Type 1 Auth: debug ip ospf packet

TTL set to 1
Length of message 48 bytes
Router-id of the received packet.

OSPF Version 2

OSPF-1 PAK : rcv. v:2 t:1 l:48 rid:0.0.0.3 aid:0.0.0.0 chk:EC94 aut:1
auk: from Serial1/3

Authentication key
Authentication type 1

What will be expected output of show ip ospf database network On the DR of this network?

R4#show ip ospf database network

OSPF Router with ID (0.0.0.4) (Process ID 1)

Net Link States (Area 0)

Routing Bit Set on this LSA in topology Base with MTID 0
LS age: 185
Options: (No TOS-capability, DC)
LS Type: Network Links
Link State ID: 10.1.1.4 (Address of Designated Router)
Advertising Router: 0.0.0.4
LS Seq Number: 80000002
Checksum: 0x70A2
Length: 40
Network Mask: /24
Attached Router: 0.0.0.4
Attached Router: 0.0.0.1
Attached Router: 0.0.0.2
Attached Router: 0.0.0.3

All OSPF routers on that subnet

What will the output of the following command show on all Routers in regards to 0.0.0.0 ?
show ip ospf database

Redistr. static
0.0.0.0/0
5.5.5.5

R1#
Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	0.0.0.3	182		0x80000001	0x00C6C3 0

R2#
Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	0.0.0.3	206		0x80000001	0x00C6C3 0

R3#
Type-7 AS External Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	0.0.0.4	222		0x80000001	0x002C53 0

Type-5 AS External Link States

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	0.0.0.3	221		0x80000001	0x00C6C3 0

R6#
Type-7 AS External Link States (Area 2)

Link ID	ADV Router	Age	Seq#	Checksum	Tag
0.0.0.0	0.0.0.4	250		0x80000001	0x002C53 0

What to expect of: show ip ospf route

R1#show ip ospf route

OSPF Router with ID (0.0.0.1) (Process ID 1)

Base Topology (MTID 0)

Area BACKBONE(0)

Intra-area Route List

- 10.1.1.0/24, Intra, cost 1, area 0, Connected via 10.1.1.1, FastEthernet0/0
- 1.0.0.0/8, Intra, cost 1, area 0, Connected via 1.1.1.1, Loopback0
- 2.0.0.0/8, Intra, cost 2, area 0 via 10.1.1.2, FastEthernet0/0
- 3.0.0.0/8, Intra, cost 2, area 0 via 10.1.1.3, FastEthernet0/0
- 4.0.0.0/8, Intra, cost 2, area 0 via 10.1.1.4, FastEthernet0/0

* connected, do not show a >
* OSPF learned route

Do not advertise the secondary address:

interface Loopback0
ip address 10.2.2.2 255.255.255.0 secondary
ip address 2.2.2.2 255.0.0.0

router ospf 1
network 10.2.2.2 0.0.0.0 area 0
network 2.2.2.2 0.0.0.0 area 0

interface Loopback0
ip address 10.2.2.2 255.255.255.0 secondary
ip address 2.2.2.2 255.0.0.0
ip ospf 1 area 0 secondaries none

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Donate

Thanks for appreciating my efforts



Colin

<p>What two solutions are available in this OSPF network to connect all Areas?</p>	<p>Option 1:</p> <pre>R2# router ospf 1 area 2 virtual-link 0.0.0.3 R3# router ospf 1 area 2 virtual-link 0.0.0.2</pre> <p>Option 2:</p> <pre>R2# interface Tunnel1 ip unnumbered Loopback0 ip ospf 1 area 0 tunnel source Serial1/x tunnel destination 23.0.0.3 R3# interface Tunnel1 ip unnumbered Loopback0 ip ospf 1 area 0 tunnel source Serial1/x tunnel destination 23.0.0.2</pre>	<p>What will show ip ospf database display in regards to 0.0.0.0?</p> <p>What will show ip ospf database display in regards to 0.0.0.0?</p>	<pre>Summary Net Link States (Area 1) (type 3) Link ID ADV Router Age Seq# Checksum 0.0.0.0 0.0.0.2 760 0x80000001 0x002D07</pre> <pre>Type-7 AS External Link States (Area 2) Link ID ADV Router Age Seq# Checksum Tag 0.0.0.0 0.0.0.4 408 0x80000001 0x00F488 0</pre>		
<p>Important facts regarding OSPF stub areas:</p>	<p>OSPF stub area:</p> <ul style="list-style-type: none"> - can NOT be a transit area (use GRE Tunnels) - can NOT have an ASBR - the backbone area can NOT be a stub - External routes are not allowed into a stub - A stub area can not have LSA Type 4's - An ABR of a stub injects a default route via summary with a default cost of 1, which can be changed. area x default-cost <37> 	<p>What will be seen in the routing table in this situation?</p> <p>What will be seen in the routing table in this situation?</p>	<pre>O*IA 0.0.0.0/0 [110/65] via 12.1.1.2, 00:26:43, Serial1/2 O 222.22.2.0/24 [110/65] via 12.1.1.2, 00:04:49, Ser1/2 no type 3!</pre> <pre>O*N2 0.0.0.0/0 [110/1] via 45.1.1.4, 00:21:10, Ethernet0/0 O 3.0.0.0/24 is subnetted, 1 subnets O IA 3.3.3.0 [110/75] via 45.1.1.4, 00:21:37, Ethernet0/0 O 4.0.0.0/24 is subnetted, 1 subnets Type 3's!</pre>		
<p>Important facts regarding OSPF totally stub areas:</p>	<p>OSPF totally stub area</p> <ul style="list-style-type: none"> - can NOT be used as transit area (use GRE tunnels) - can NOT have an ASBR - backbone area can not be a totally stub area. - external routes not allowed in totally stub area. - default route (summary 3) injected, cost of default route can be changed. area x default-cost <z> - do not get IA routes of other areas. 	<p>Configure R1 to limit the number of non-self generated LSA's in its database to 10. R1 should generate a warning message if 50 % of this threshold is reached.</p>	<p>Verify current amounts of LSA in the entire domain:</p> <pre>show ip os database database-summary Non-self Process 1 database summary LSA Type Count Delete Maxage Router 3 0 0 ... Non-self 3 Total 4 0 0</pre> <pre>router ospf 1 max-lsa 10 50 warning-only Max 10 Non-self LSA's Warning thresh. 50%</pre> <pre>%OSPF-4-OSPF_MAX_LSA_THR: Threshold for maximum number of non self-generated LSA has been reached "ospf 1" - 8 LSAs %OSPF-4-OSPF_MAX_LSA: Maximum number of non self-generated LSA has been exceeded "ospf 1" - 11 LSAs</pre> <pre>show ip os database database-summary Non Non-self 6</pre>		
<p>What alternative is there for the following config snip:</p> <pre>router ospf 1 no discard-route external</pre>	<p>Alternative:</p> <pre>router ospf 1 discard-route external 255</pre> <p><i>Sets the admin distance of the external discard route to 255!</i></p>	<p>How can you bypass OSPF's sanity check in this situation -> to form an adjacency</p> <p>show ip ospf neighbor LIST IS EMPTY</p>	<pre>R1# int Lo1 ip address 10.0.0.1 255.255.255.0 int ser1/1 ip unnumbered Lo1 router ospf 1 network 10.0.0.1 0.0.0.0 area 0</pre> <pre>R2# int Lo1 ip address 22.0.0.22 255.255.255.0 int ser1/1 ip unnumbered Lo1 router ospf 1 network 22.0.0.22 0.0.0.0 area 0</pre> <pre>%OSPF-5-ADJCHG: Process 1, Nbr 1.1.1.1 on Serial1/1 from LOADING to FULL, Loading Done</pre>		
<p>Configure R1 such that it retransmits an LSA if no ACK is heard after 10 seconds:</p>	<p>Verify current settings:</p> <pre>R1#show ip ospf interface fa0/0 i Retransmit Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 5</pre> <p>Configure:</p> <pre>conf t int fa0/0 ip ospf retransmit-interval 10</pre> <p>Verify configuration</p> <pre>R1#show ip ospf interface fa0/0 i Retransmit Timer intervals configured, Hello 250 msec, Dead 1, Wait 1, Retransmit 10</pre>	<p>OSPF filtering route-source</p> <p>Configure R1 to filter out 22.22.22.0/24</p>	<pre>R1# distribute-list route-map RMP-NO-22-NET in route-map RMP-NO-22-NET deny 10 match ip route-source 11 Route-map RMP-NO-22-NET permit 20 access-list 11 permit 0.0.0.2 OSPF router-id of R2</pre>		
<p>Configure R1 such that in case there is a topology change in the domain R1 only re-calculates the affected LSA type 1 and 2's in its database.</p>	<p>Have R1 only calculate the affected LSA 1 / LSA 2 recalculated. Speeding up convergence and lowering the amount of needed CPU cycles!</p> <pre>R1# conf t router ospf 1 ispf</pre>				

Help me create more flashcards:

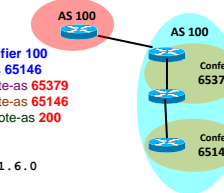

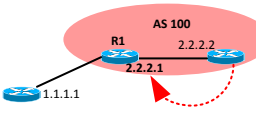
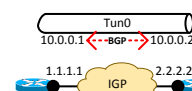

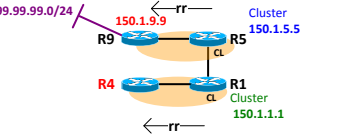
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin


<p>show ip bgp x.x.x.x y.y.y.y explained:</p>	<pre>R1# R1#show ip bgp 112.0.0.0 255.0.0.0 BGP routing table entry for 112.0.0.0/8, version 22 Paths: (2 available, best #2, table Default-IP-Routing-Table) Not advertised to any peer 54 50 60 54.1.1.254 (metric 2172416) from 155.1.146.6 (150.1.6.6) Origin IGP, metric 0, localpref 100, valid, internal 54 50 60 204.12.1.254 (metric 30720) from 155.1.0.4 (150.1.4.4) Origin IGP, metric 0, localpref 100, valid, internal, best</pre> <p>204.12.1.254 eBGP next-hop 155.1.0.4 Peering-IP (150.1.4.4) router-id</p>	<h2>iBGP Confederation</h2>	 <pre>router bgp 65379 bgp confederation identifier 100 bgp confederation peers 65146 neighbor 155.1.37.3 remote-as 65379 neighbor 155.1.37.7 remote-as 65146 neighbor 155.1.44.44 remote-as 200</pre> <pre>R1#show ip bgp 150.1.6.0 Local 155.1.108.10 from 155.1.58.5 (150.1.5.5) Origin IGP, metric 0, localpref 100, valid, confed-internal (65379 65146) 155.1.146.4 (metric 2172416) from 155.1.0.3 (150.1.3.3) Origin IGP, metric 0, localpref 100, valid, confed-external 200 192.10.1.254 from 192.10.1.254 (222.22.2.1) Origin incomplete, metric 0, localpref 100, valid, external</pre>	<h2>iBGP Synchronization</h2> <p>Disable RIBfailure marked routes from being automatically propagated:</p>	<p>By default BGP routes that have RIB-failure are advertised to neighbors.</p> <p>This can be disabled by using:</p> <p>bgp suppress-inactive</p>
<h2>BGP Update Source Modification</h2>	<pre>router bgp 100 neighbor 150.1.4.4 update-source Loopback0</pre>	<h2>iBGP Confederation default path decision:</h2>	<pre>show ip bgp 150.1.7.0 ... (65379) ... 155.1.37.7 .. from 155.1.0.1 (150.1.3.3) Origin ..., valid, confed-external (65146 65379) ... 155.1.37.7 .. from 155.1.0.1 (150.1.1.1) Origin ..., valid, confed-external, best</pre> <p>Lower Router-ID wins over confed-as path length by default!</p>		
<h2>BGP neighbor 172.16.103.3 disable-connected-check</h2> <p>Command:</p>	 <pre>R1# ip route 172.16.103.0 255.255.255.0 10.1.0.3 ip route 172.16.103.0 255.255.255.0 10.1.13.3 ! router bgp 10 no synchronization bgp router-id 1.1.1.1 bgp log-neighbor-changes neighbor 172.16.103.3 remote-as 30 neighbor 172.16.103.3 disable-connected-check neighbor 172.16.103.3 update-source Loopback0 no auto-summary</pre>	<h2>BGP Next-Hop Processing – Next-Hop-Self:</h2>	<pre>R1# router bgp 100 neighbor 2.2.2.2 next-hop-self</pre>  <p>If next-hop is not set and 1.1.1.1 is not know the following could happen:</p> <pre>R2#show ip bgp 100.100.100.0 ... 54 50 60, (Received from a RR-client) 1.1.1.1 (inaccessible) from 155.1.146.4 (150.1.4.4) ...</pre>	<h2>BGP over GRE</h2> <p>Run BGP Session across non-BGP capable router cloud</p>	 <pre>interface Tunnel0 ip address 10.0.0.1 255.255.255.0 tunnel source 1.1.1.1 tunnel destination 2.2.2.2 router bgp 200 neighbor 10.0.0.2 remote-as 100</pre> <pre>interface Tunnel0 ip address 10.0.0.2 255.255.255.0 tunnel source 2.2.2.2 tunnel destination 1.1.1.1 router bgp 200 neighbor 10.0.0.1 remote-as 100</pre> <p>Transporting / Hiding a BGP next-hop via using a GRE tunnel. IGP does not need to know BGP next-hops.</p>
<h2>Authenticating BGP Peerings</h2>	<pre>Router bgp 200 neighbor 192.10.1.254 password CISCO</pre> <p>Verify the established, autenticated BGP session:</p> <pre>show ip bgp neighbors 192.10.1.254 include state Flags BGP state = Established, up for 00:01:19 Connection state is ESTAB, I/O status: 1, unread input bytes: 0 Flags: active open, nagle, md5</pre>	<h2>BGP Next-Hop Processing - Manual Modification</h2>	<pre>Set NHS for your coming from R6</pre> <pre>router bgp 100 neighbor 1.1.1.1 route-map SET_NHS_FROM_R6 in route-map SET_NHS_FROM_R6 permit 10 set ip next-hop 2.2.2.1</pre> <pre>router bgp 100 neighbor 1.1.1.1 route-map SET_NHS_TO_R6 out route-map SET_NHS_TO_R6 permit 10 set ip next-hop 2.2.2.1</pre> <p>Set NHS for routes being sent to R6</p>	<h2>BGP Redistribute Internal</h2>	<pre>Router [ospf, eigrp, rip] redistribute bgp</pre> <p>only eBGP prefixes will be redistributed by default to the other routing protocols!</p> <p>bgp redistribute Internal</p> <p>Will enable iBGP prefix to be redistributed into other routing protocols. -> Dangerous, as iBGP has no internal Loop prevention mechanism.</p>
<h2>iBGP Route Reflection</h2>	 <pre>R1# Router bgp 200 neighbor 155.1.0.3 route-reflector-client neighbor 155.1.0.4 route-reflector-client</pre> <pre>R1#show ip bgp 150.1.2.0 255.255.255.0 ... Local, (Received from a RR-client) ... R2#show ip bgp 150.1.10.0 255.255.255.0 ... Originator: 204.12.1.10, Cluster list: 150.1.1.1 ...</pre>	<h2>iBGP Synchronization</h2> <p>Route is not in IGP:</p>	<pre>router bgp 100 Synchronization</pre> <pre>R1#show ip bgp 112.0.0.0 ... Origin IGP, ..., valid, internal, not synchronized ...</pre> <p>Route not in IGP</p> <p>Route needs to be visible in IGP before BGP is allowed to propagate the prefix.</p>	<h2>BGP Peer Groups</h2>	<pre>router bgp 100 neighbor IBGP_PEERS peer-group neighbor IBGP_PEERS <commands> neighbor 150.1.5.5 peer-group IBGP_PEERS</pre> <pre>show ip bgp peer-group ... BGP neighbor is IBGP_PEERS, peer-group internal, members: 150.1.3.3 150.1.4.4 150.1.5.5 150.1.6.6 Index 0, Offset 0, Mask 0x0 Route-Reflector Client Update messages formatted 0, replicated 0 Number of NLRI in the update sent: max 0, min 0</pre>
<h2>Route Reflection with Clusters</h2>	 <pre>R4#show ip bgp 99.99.99.0 BGP routing table entry for 99.99.99.0/24, version 16 Paths: (2 available, best #2, table Default-IP-Routing-Table) Not advertised to any peer Local 155.1.79.9 (metric 2175488) from 155.1.58.5 (150.1.5.5) Origin IGP, metric 0, localpref 100, valid, internal Originator: 150.1.9.9, Cluster list: 150.1.1.1, 150.1.5.5</pre>	<h2>iBGP Synchronization</h2> <p>Route is in IGP:</p>	<pre>router bgp 100 Synchronization</pre> <pre>Rack1R1#show ip bgp 222.22.2.0 ... Paths: (1 available, best #1, table Default-IP-Routing-Table, RIBfailure(17)) ... Origin incomplete, metric 0, localpref 100, valid, internal, synchronized, best</pre> <p>RIB-failure output is an informational message to let us know that although the BGP route is valid, it is not being installed in the routing table -> route via an IGP route with a lower administrative distance</p>	<h2>BGP Network Statement</h2>	<pre>Interface loop0 Ip address 99.9.9.9 255.255.255.0</pre> <pre>ip route 150.0.0.0 255.252.0.0 Null0</pre> <pre>router bgp 100 network 150.0.0.0 mask 255.252.0.0 network 99.9.9.0 mask 255.255.255.0</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

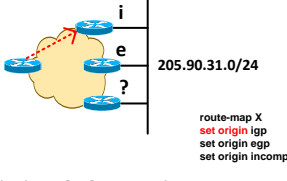
Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

<p>BGP Auto-Summary</p> <p>Create a summary, but do not use aggregate-address:</p>	<p>Router bgp 100 auto-summary</p>	<p>How to make BGP ignore the AS-Path length:</p>	<p>Router bgp x bgp bestpath as-path ignore</p>	<p>BGP DMZ Link Bandwidth</p> <p>unequal cost load-balancing</p> <p>Described:</p>	<p>Router bgp 200 maximum-paths ibgp bgp dmzlink-bw (enable feature) neighbor <IP> dmzlink-bw neighbor <IP> send-community extended</p> <p>load-balance based on the bandwidth of the links used to connect to the external BGP peers</p> <p>bandwidth value is copied into a new extended community attribute</p> <p>Internal routers need maximum-path activated and need to exchange ext-communities</p>
<p>Describe BGP path selection:</p>	<ol style="list-style-type: none"> 1) Ignore invalid paths (no valid next hop, not synchronized, looped). 2) Prefer path with the highest locally assigned weight value. 3) Prefer path with the highest Local Preference attribute value. 4) Prefer locally originated prefixes (i.e. originated via the network, aggregate-address or redistribution commands). 5) Prefer path with the shortest AS_PATH attribute length 6) Prefer path with the lowest numerical value of the Origin code (IGP < EGP < Incomplete) 7) Prefer path with the lowest MED attribute value (provided that the first AS in the list is the same). 8) Prefer external BGP paths over Internal 9) Prefer path with the smallest IGP metric to reach the NEXT_HOP IP address 10) Prefer path originated from the router with the lowest BGP Router ID 	<p>BGP - Origin</p>	<p>BGP origin type can be used for path selection:</p>  <pre> route-map X set origin igp set origin egp set origin incomplete </pre> <p>BGP table version is 18, local router ID is 150.1.1.1 Status codes: s suppressed, d damped, h history, * valid, > best, i internal</p> <pre> R#show ip bgp 205.90.31.0 RIP-Failure, S Stale Origin codes: i - IGP, e - EGP, ? - Incomplete Network Next Hop Metric LocPrf Weight Path * 205.90.31.0 155.1.45.3 0 100 0 200 254 e * 155.1.38.7 0 100 0 200 254 e * 155.1.13.3 0 100 0 200 254 ? </pre>	<p>BGP DMZ Link Bandwidth</p>	

<h3>BGP Filtering Standard Access-Lists</h3>	<p>AS-100 config: Lo10 10.0.0.1/24 Lo10 10.0.1.1/24 Lo10 10.0.2.1/24 Lo10 10.0.3.1/24</p> <p>AS-600 in BGP table: 10.0.0.0/24 10.0.1.0/24 10.0.2.0/24 10.0.3.0/24</p> <p>access-list 100 permit ip host 10.0.1.0 host 255.255.255.0 access-list 100 permit ip host 10.0.3.0 host 255.255.255.0</p> <p>route-map RMP-100-IN deny 10 match ip address 100</p> <p>route-map RMP-100-IN permit 20</p> <p>router bgp 400 neighbor 155.1.146.1 route-map RMP-100-IN in</p>	<h3>BGP Local AS</h3> <p>No prepend replace-as</p>		<h3>BGP Dampening with Route-Map</h3>	<pre>router bgp 200 bgp dampening route-map DAMPENING route-map DAMPENING permit 10 match as-path 100 set dampening 4 750 2000 16 route-map DAMPENING permit 90 catch-all others!</pre> <p>show ip bgp dampening parameters debug ip bgp dampening</p>																		
<h3>BGP Filtering Extended Access-Lists</h3>	<p>BGP filtering: ACL extended: variable part: permit ip 192.168.0.0 0.0.0.255 255.255.255.0 0.0.0.255</p> <p>AND</p> <p>matching the subnet Range 192.168.0.0 to 192.168.0.255</p> <p>prefix length of /24 or greater</p> <p>10.0.0.X with /28 mask permit ip 10.0.0.0 0.0.0.255 255.255.255.240 0.0.0.0</p> <p>10.0.X.0 with /24 mask permit ip 10.0.0.0 0.0.0.255.0 255.255.255.0 0.0.0.0</p>	<h3>BGP Local AS Replace-AS/Dual-AS</h3>		<h3>BGP Timers Tuning</h3>	<pre>Router bgp 100 timers bgp <keepalive> <holdtime> timers bgp 60 180</pre> <p>time it takes for the IGP update to be redistributed into BGP For VPNv4 pfxs</p> <p>bgp scan-interval <5-60> bgp scan-time import <5-60> debug ip bgp keepalive</p> <p>batches all new prefixes and delays the sending of an update packet to the peer until the next advertisement interval timer expires: set to 0 advertises immediately</p> <p>neighbor <IP> advertisement-interval <seconds></p>																		
<h3>BGP Regular Expressions</h3>	<p>^\$ locally originated prefix _254\$ pfx originated in AS 254 ^254_ pfx received from adjacent neighbor _254_ pfx was transiting AS 254 ^{[0-9]+}_254 pfx from AS-254 one AS hop away ^254_{[0-9]+} pfx of an AS connected to AS 254 ^{254_}{[0-9]+} same as above but allows 254 to prepend ^254_{[0-9]+}_ An AS connected to 254 can prepend ^{65100}\$ pfx from confederation peer 65100</p> <p>show ip bgp quote-regexp _101\$_100\$ ip as-path access-list 1 permit _54\$</p>	<h3>BGP Remove Private AS</h3>		<h3>BGP Fast Fallover</h3>	<p>if line protocol to eBGP speaker goes down, dont wait for holddown timer, tear session down. Enabled by default.</p> <p>bgp fast-external-fallover</p> <p>Informs BGP about line-proto down</p> <p>same feature but for shared ip segments enabled by neighbor:</p> <p>neighbor <IP> fall-over</p>																		
<h3>BGP Filtering with Maximum Prefix</h3>	<p>neighbor <IP> maximum-prefix <Number> [<Threshold%>] [warning-only] [restart <minutes>]</p> <p>router bgp 100 neighbor 54.1.1.254 maximum-prefix 20 80 restart 3</p> <p>router bgp 300 neighbor 155.1.37.3 maximum-prefix 20 warning-only</p> <p>SW1#show ip bgp neighbors 155.1.37.3 include Maximum Thresh Maximum prefixes allowed 20 (warning-only) Threshold for warning message 75%</p>	<h3>BGP Dampening</h3>	<p>Have a route being dampened if it flaps 2 for 5 minutes:</p> <pre>router bgp 200 bgp dampening 4 750 2000 16</pre> <p>Bgp dampening HL RL SL MS</p> <p>bgp dampening Half_Life ReuseLimit SuppressLimit MaximumSuppressTime</p> <p>show ip bg dampening flap-statistics show ip bg dampening dampened-path</p>	<p>Good BGP route troubleshooting debug command:</p> <p>debug ip bgp rib-filter</p>	<p>debug ip bgp rib-filter</p> <pre>R1(config)#no ip route 99.99.99.99 255.255.255.255 Null0</pre> <p>*Jul 16 16:46:19.538: BGP- ATF: EVENT 99.99.99.99/32 RIB update DOWN *Jul 16 16:46:19.538: BGP- ATF: EVENT 99.0.0.0/8 RIB update DOWN</p> <pre>R1(config)#no ip route 99.99.99.99 255.255.255.255 Null0</pre> <p>*Jul 16 16:46:30.470: BGP- ATF: EVENT 99.99.99.99/32 RIB update UP</p>																		
<h3>BGP Default Routing</h3>	<p>Router bgp 100 network 0.0.0.0 mask 0.0.0.0</p> <p>default route is advertised if the route-map match conditions are satisfied</p> <p>ip prefix-list PFX-CHECK permit 10.10.10.0/24</p> <p>route-map RMP-DEFAULT permit 10 match ip address prefix-list PFX-CHECK</p> <p>neighbor <IP> default-originate route-map RMP-DEFAULT</p>	<h3>BGP Dampening</h3> <p>Show outputs:</p>	<pre>sh ip bgp 1.1.1.1/32 ... 400 (suppressed due to dampening) (history entry) ... Dampinfo: penalty 3556, flapped 4 times in 00:04:41, reuse in 00:00:53</pre> <p>Flapped 1x, but not yet dampened</p> <table border="1"> <thead> <tr> <th>Network</th> <th>Next Hop</th> <th>Metric</th> <th>LocPrf</th> <th>Weight</th> <th>Path</th> </tr> </thead> <tbody> <tr> <td>h 1.1.1.1/32</td> <td>155.1.146.1</td> <td>0</td> <td></td> <td></td> <td>400 i</td> </tr> <tr> <td>d 1.1.1.1/32</td> <td>155.1.146.1</td> <td>0</td> <td></td> <td></td> <td>400 i</td> </tr> </tbody> </table> <p>Flapped several times, is dampened</p>	Network	Next Hop	Metric	LocPrf	Weight	Path	h 1.1.1.1/32	155.1.146.1	0			400 i	d 1.1.1.1/32	155.1.146.1	0			400 i	<h3>BGP Outbound Route Filtering (ORF)</h3>	<p>neighbor <IP> capability orf prefix-list (send receive both)</p> <pre>Received: 1.1.1.1/32 6.6.6.6/32</pre> <p>R1# ip prefix-list ORF seq 3 deny 6.6.6.6/32 ip prefix-list ORF seq 10 permit 0.0.0.0/0 le 32</p> <p>router bgp 400 neighbor 155.1.146.6 capability orf prefix-list both neighbor 155.1.146.6 prefix-list ORF in</p> <p>R2# sh ip bgp neighbors 155.1.146.4 advertised-routes *> 1.1.1.1/32 0.0.0.0 0 32768 i</p>
Network	Next Hop	Metric	LocPrf	Weight	Path																		
h 1.1.1.1/32	155.1.146.1	0			400 i																		
d 1.1.1.1/32	155.1.146.1	0			400 i																		
<h3>BGP Local AS</h3>		<h3>BGP dampening calculation:</h3> <p>Routes flaps twice in a row, they should only advertise after 10 minutes of stability:</p> <p>Maximum penalty = reuse-limit * 2 ^ (maximum suppress time/half-life tim)</p> <p>2000 = 750 * 2 ^ (10 minutes / half-life time) 2000/750 = 2 ^ (10 / half-life) 8/3 = 2 ^ (10 / half-life) using basic exponential equation, take log of both side</p> <p>log (8/3) = log [2 ^ (10 / half-life)] 0.4259 = 10 log 2 / half-life 0.4295 = 3.0102 / half-life half-life = 7.067 (approximately 7)</p>	<h3>BGP Outbound Route Filtering (ORF)</h3> <p>Show commands:</p>	<pre>R6#show ip bgp neighbors 155.1.146.4 received prefix-filter Address family: IPv4 Unicast ip prefix-list 155.1.146.4: 3 entries seq 3 deny 6.6.6.6/32 seq 10 permit 0.0.0.0/0 le 32</pre> <p>Config -----</p> <p>R1# ip prefix-list ORF seq 3 deny 6.6.6.6/32 ip prefix-list ORF seq 10 permit 0.0.0.0/0 le 32</p> <p>router bgp 400 neighbor 155.1.146.6 capability orf prefix-list both neighbor 155.1.146.6 prefix-list ORF in</p> <p>Force neighbor to update filter: clear ip bgp * x.x.x.x out</p>																			

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

BGP Soft Reconfiguration

Router bgp 100
neighbor <IP> soft-reconfiguration inbound

```

Adj In | RIB | Adj out
-----|---|-----
Prefix-list deny 6.6.6.6/32
sh ip bgp neighbors 155.1.146.6 received-routes
* 6.6.6.6/32 155.1.146.6 0 0 600 i
*> 1.1.1.1/32 155.1.146.6 0 0 600 i

sh ip bgp neighbors 155.1.146.6 routes
*> 1.1.1.1/32 155.1.146.6 0 0 600 i
    
```

BGP multipath

BGP Multipath	BGP address family command
eBGP	maximum-path n
iBGP	maximum-path ibgp n
eIBGP	maximum-path eibgp n

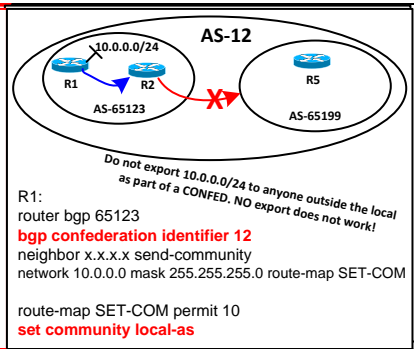
Weight, LP, AS-Path, Origin, MED need to be the same.
Default multipath = 1

BGP multipath

BGP Multipath	BGP address family command
eBGP	maximum-path n
iBGP	maximum-path ibgp n
eIBGP	maximum-path eibgp n

Weight, LP, AS-Path, Origin, MED need to be the same.
Default multipath = 1

BGP community explained



BGP Next-Hop Trigger

bgp nexthop trigger enable

As soon as any change that affects an existing NEXT_HOP occurs, the watch process notifies the BGP router Process

bgp nexthop trigger delay <seconds>
Schedules the next BGP table scan after a change happened to a next-hop after <seconds> of it being detected.

BGP SOO Configuration:

```

neighbor 1.1.1.1 soo 10:12

neighbor 1.1.1.1 route-map RMP-X in
route-map RMP-X permit 10
set extcommunity 10:12
    
```

BGP 4-Byte Autonomous System Numbers

```

Router# show ip bgp summary
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down Statd
192.168.1.2 4 65536 7 7 1 0 000:03:04 0
192.168.3.2 4 65550 4 4 1 0 000:00:15 0

router bgp X
bgp asnotation dot asplain: 65536 to 4294967295

Router# show ip bgp summary
BGP table version is 1, main routing table version 1
Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down Statd
192.168.1.2 4 1.0 9 9 1 0 000:04:13 0
192.168.3.2 4 1.14 6 6 1 0 000:01:24 0

show ip bgp regexp ^1\.0$ (ASN 1.0 / 65536)
    
```

BGP TTL Security

neighbor <IP> ttl-security hops <hop-count>

BGP packets will be tolerated with a TTL no lower than (255 - 2 hops) = 253

R5#show ip bgp neighbors 1.1.1.1| inc TTL
Minimum incoming TTL 253, Outgoing TTL 255

BGP scan interval VPN performance tuning

```

bgp scan-interval
general scanning interval

advertisement-interval
periodic "batching" of events:
BGP waits for the timer to expire and accumulates the updates

scan-time import 15
process to import the MP-BGP VPNv4 prefixes into the local VRF table

router bgp 100
address-family vpnv4 unicast
scan-time import 15
neighbor 150.1.5.5 advertisement-interval 0
bgp scan import 5
    
```

BGP BFD for IPv4 and IPv6

```

interface X
bfd interval <msec> min_rx <msec> multiplier <interval-multiplier>

bfd interval 50 min_rx 50 multiplier 3 ! = 150 msec holddown

router bgp 65000
neighbor 2001:DB8:5::2 fall-over bfd
neighbor 1.2.3.4 fall-over bfd

show bfd neighbors
    
```

BGP TTL Security Debugging:

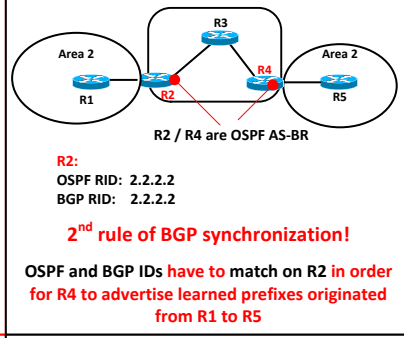
R1/R3# neighbor <IP> ttl-security hops 2

Debug ip packet detail DUMP //do not use in production!

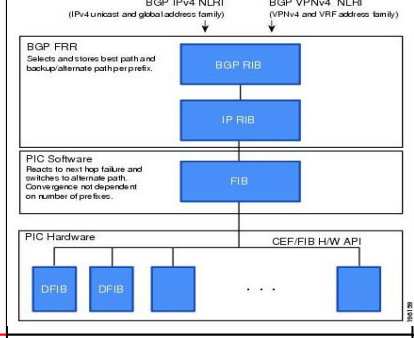
```

*Feb 10 12:15:37.485: IP: s=155.1.146.4 (GigabitEthernet0/0.146), d=155.1.146.6 (GigabitEthernet0/0.146), len 61, rcvd 3
*Feb 10 12:15:37.485: TCP src=47218, dst=179, seq=1061966907, ack=369930400, win=16384 ACK PSH
3F200C00: 0026 0857B960 .&.W9*
3F200C10: 00260857 BA610800 45C0003D 00150000 .&.W.a.E@.=...
3F200C20: FE0661D8 98019204 98019206 B87200B3 ~.aX.....8r.3
3F200C30: 3F4C543B 160CB0AD 50184000 FF360000 ?L.T..O P@..6..
3F200C40: FFFFFFFF FFFFFFFF FFFFFFFF FFFFFFFF .....
3F200C50:
FE Hex = 254 packet has arrived with a TTL of 254 from 155.1.146.4
    
```

2nd Rule of IGP / BGP synchronization



BGP PIC Edge and BGP FRR Explained:



BGP AllowAS in

```

AS-6500 config: AS-6500 AS-40 AS-6500 Received: 10.0.0.1/24 10.0.1.1/24
AS-6500 config: AS-6500 AS-40 AS-6500 Received: 10.0.0.1/24 10.0.1.1/24

debug ip bgp updates
rcv UPDATE about 7.7.7.7/32 -- DENIED due to: AS-PATH contains our own AS:

neighbor <IP> allowas-in
    
```

BGP as-override

```

router bgp X
address-family ipv4 unicast vrf x
neighbor 1.1.1.1 as-override
neighbor 2.2.2.2 as-override

AS-6500 config: CE1 AS-6500 CE2 AS-6500
Lo10 10.0.0.1/24 1.1.1.1
Lo10 10.0.1.1/24 2.2.2.2

CE1: AS-path for 20.0.0.0/24: 40
CE2: AS-path for 10.0.0.0/24: 40

PE hides CE AS number behind AS-40
    
```

BGP info:

- idle
- connect
- active (3 way handshake)
- open sent (capabilities)
- open confirm
- open

router bgp X
neigh 1.1.1.1 remote-as 200 (reachability to peer)

neighbor transport-mode

OpenSent:
- ASN Number
- Holddown
- Router-ID
- Options (AFI/SAFI)
- Open

Keep alives

BGP can NOT peer using a default 0/0 route!

BGP Neighbor capabilities:

AFI:	Reserved
0	Reserved
1	IPv4
2	IPv6
11	IPX
12	AppleTalk

debug ip bgp afi/safi

SAFI:	Reserved
1	unicast forwarding
2	multicast forwarding
3	both_unicast and multicast forwarding
4	IPv4 label forwarding
128	labeled VPN forwarding

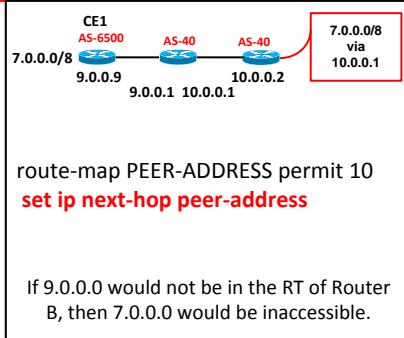
Alternative to neighbor x.x.x.x set-next-hop-self and neighbor x.x.x.x route-map OUT Route-map OUT permit 10 Set ip next-hop-self

```

neighbor x.x.x.x set-next-hop-self

neighbor x.x.x.x route-map OUT
Route-map OUT permit 10
Set ip next-hop-self

If 9.0.0.0 would not be in the RT of Router B, then 7.0.0.0 would be inaccessible.
    
```



BGP neighbor transport connection-mode

```

bgp neighbor transport connection-mode [active / passive]

TCP SYN to 1.1.1.1:179

1.1.1.1
Server Side | Client Side
outside | inside

Watch out for stateful firewalls or Router ACLs with keyword established
    
```

BGP

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Donate

VISA, giro pay, etc.

Thanks for appreciating my efforts

Colin

BGP

<p>BGP next-hop default</p>		<p>bgp bestpath med missing-as-worst</p> <p>bgp bestpath med confed</p>	<p>bgp bestpath med missing-as-worst</p> <p>consider a missing MED as having a value of infinity, making the path without a MED value the least desirable path.</p> <p>bgp bestpath med confed</p> <p>Enable Path/MED comparison between different Sub-Ases within a confederation.</p> <ul style="list-style-type: none"> - comparison between MEDs is made only if there are no external autonomous systems in the path - external autonomous system in the path, then the external MED is passed transparently through the confederation 	<p>bgp route-map priority</p>	<pre>router bgp 50000 bgp route-map priority address-family ipv4 unicast vrf inside bgp route-map priority</pre>																								
<p>BGP 4 Byte ASN calculation explained</p>	<p>2 Byte ASN Range: 0 to 65535</p> <p>4 Byte ASN Range: 0 to 4294967295</p> <p>2 Byte max ASN 65535 = 0.65535 4 Byte ASN Format</p> <p>65536 = 1.0 65537 = 1.1</p> <p>Example calculation using ASN 140000</p> <p>140000 / 65535 = 2.1 (gives the factor)</p> <p>2x 65535 = 131070</p> <p>140000 - 131070 = 8930</p> <p>Final 4 Byte ASN Nr: 2.8930</p>	<p>bgp bestpath med confed</p> <p>path= 65000 65004, med=2 path= 65001 65004, med=3 path= 65002 65004, med=4 path= 65003 22, med=1</p> <p>Which path will be used?</p>	<p>Used due to lowest MED</p> <p>path= 65000 65004, med=2 path= 65001 65004, med=3 path= 65002 65004, med=4 path= 65003 22, med=1</p> <p>Has lower MED, but it is not involved in the MED comparison because there is an external autonomous system in this path</p> <p>bgp bestpath med confed</p>	<p>BGP router-id</p>	<pre>bgp router-id { ip-address vrf auto-assign }</pre> <pre>Router bgp 45000 bgp router-id 1.1.1.1</pre> <pre>router bgp 45000 bgp router-id vrf auto-assign</pre> <pre>router bgp 45000 address-family ipv4 vrf VRF2 bgp router-id auto-assign</pre>																								
<p>BGP 4 Byte ASN: Data entities that carry ASNs</p>	<ul style="list-style-type: none"> - The AS_PATH attribute; (AS4_PATH) - The AGGREGATOR attribute; (AS4_AGGREGATOR) - The COMMUNITES attribute; (4 Byte EXT_COMM) - The Open message (New capability) <ul style="list-style-type: none"> - Neighbor is either New_BGP or Old_BGP implementation (Capability). - New to old BGP speaker uses reserved 2-byte ASN, 23456, called AS_TRANS and no OLD ASN SHOULD USE THIS ASN! - AS4_PATH new optional path attribute, unlike "historic" AS_PATH attribute which is mandatory 	<p>BGP bgp deterministic med</p> <p>(ensures MED gets compared where AS paths are the same)</p>	<p>bgp deterministic-med</p> <p>Groups the same AS path together, checks the oldest path within the group. Then compares the oldest path towards AS 22 with the oldest path of the group going to AS 44. The oldest path between the two group wins. The oldest path is always displayed at the bottom of show bgp!</p> <p>Compare Oldest between groups (red group contains the oldest, flag best)</p>	<p>BGP RPKI</p> <p>bgp rpki server tcp 192.168.1.1 port 1033 refresh 600</p> <p>(Origin AS Validation)</p>	<pre>router bgp X bgp rpki server tcp 192.168.1.1 port 1033 refresh 600</pre> <p>show ip bgp rpki servers</p> <p>Displays the current state of communication with the RPKI servers.</p> <p>show ip bgp rpki table</p> <p>cached list of prefix/AS pairs.</p> <p>bgp bestpath prefix-validate</p> <p>Invalid prefixes are allowed to be used as the best path, even if valid prefixes are available, or disables the checking of prefixes.</p> <p>clear ip bgp rpki server</p> <p>Purges SOVC records downloaded from the specified server</p> <p>debug ip bgp event rpki</p> <p>neighbor announce rpki state</p>																								
<p>4 Byte – 2 Byte BGP ASN operation</p>		<p>BGP bgp bestpath compare-router-id</p>	<p>Do not use the oldest path as the best</p> <p>NEWEST PATH -> Compare 1st OLDER PATH -> Compare 2nd OLDEST PATH</p> <p>bgp bestpath compare-router-id (will select lowest RID)</p> <p>Making bgp NOT compare down to the oldest path.</p>	<p>bgp update-delay</p>	<p>Delay the first update with prefixes for X seconds</p> <p>bgp update-delay <seconds></p> <p>Could be used in to initialize learned prefixes learned older or newer depending how much you delay on which router.</p> <p>Command can be used with:</p> <p>bgp graceful-restart</p>																								
<p>BGP Dynamic Neighbors</p> <p>(Up to 5 optional AS numbers)</p>	<p>2-byte AS Transit</p> <p>(Using lowest RID)</p> <p>Even path has higher MED</p>	<p>BGP Dynamic Neighbors</p> <p>(Up to 5 optional AS numbers)</p>	<pre>router bgp bgp listen limit <max-number> bgp listen range 10.0.0.0/24 peer-group GROUP neighbor GROUP ebgp-multihop <ttl> neighbor GROUP remote-as 22 alternative-as 44</pre> <p>debug ip bgp range show ip bgp peer-group X</p> <p>bgp slow-peer detection [threshold seconds]</p> <p>oldest message in a peers queue can be lagging behind the current time before the peer is determined to be a slow peer</p> <p>neighbor slow-peer detection</p> <p>clear ip bgp slow</p> <p>bgp slow-peer split-update-group dynamic [permanent]</p>	<p>BGP table-map</p>	<pre>router bgp 500 table-map buckets</pre> <pre>ip as-path access-list 99 permit _10_ ip as-path access-list 77 permit _11_</pre> <pre>route-map buckets permit 10 match as-path 99 set traffic-index 1</pre> <pre>route-map buckets permit 20 match as-path 77 set traffic-index 2</pre> <pre>route-map buckets permit 80 set traffic-index 7</pre> <p>Interface fa0/x</p> <p>bgp-policy accounting input source</p>																								
<p>BGP BMP Server</p>	<p>4-byte AS Transit</p> <p>AS seen as "23456" ...</p> <p>Best BGP Path</p>	<p>BGP BMP Server</p>	<p>Both neighbors seen as 4 Byte, best path lower MED</p>	<p>BGP BMP Server</p>	<pre>router bgp 65 bmp initial-refresh delay 30</pre> <p>show ip bgp bmp server neighbors</p> <p>neighbor bmp-activate</p> <p>Device# show ip bgp bmp server neighbors</p> <p>Number of BMP neighbors configured: 10 BMP Refresh not in progress, refresh not scheduled Initial Refresh delay configured, refresh value 30s BMP buffer size configured, buffer size 2048 MB, buffer size bytes used 0 MB</p> <table border="1"> <thead> <tr> <th>Neighbor</th> <th>PeerID</th> <th>MsgQ</th> <th>CfgBvrf#</th> <th>ActBvrf#</th> <th>RW Bvrf</th> </tr> </thead> <tbody> <tr> <td>30.1.1.1</td> <td>0</td> <td>0</td> <td>1</td> <td>2</td> <td>16</td> </tr> <tr> <td>2001:1001:1001</td> <td>0</td> <td>0</td> <td>1</td> <td>2</td> <td>15</td> </tr> <tr> <td>40.1.1.1</td> <td>0</td> <td>0</td> <td>1</td> <td>2</td> <td>26</td> </tr> </tbody> </table>	Neighbor	PeerID	MsgQ	CfgBvrf#	ActBvrf#	RW Bvrf	30.1.1.1	0	0	1	2	16	2001:1001:1001	0	0	1	2	15	40.1.1.1	0	0	1	2	26
Neighbor	PeerID	MsgQ	CfgBvrf#	ActBvrf#	RW Bvrf																								
30.1.1.1	0	0	1	2	16																								
2001:1001:1001	0	0	1	2	15																								
40.1.1.1	0	0	1	2	26																								

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

BGP additional path (BGP Add path)

```

router bgp 100
address-family ipv4 unicast
  bgp additional-paths select all
  neighbor 192.168.1.2 additional-paths send receive
  neighbor 192.168.1.2 advertise additional-path all

address-family ipv4 unicast
  bgp additional-path select all best 3 group-best
  neighbor 1.1.1.1 route-map ADD-PATH out
  neighbor 1.1.1.1 advertise additional-paths best 2

route-map ADD-PATH permit 10
match additional-paths advertise-set best 2
set metric 888
    
```

How do you configure eBGP multipath?

```

router bgp 50
  bgp bestpath as-path multipath-relax
  maximum-paths 2
  neighbor 2.2.13.3 remote-as 100
  neighbor 2.2.16.6 remote-as 200

Router#show bgp
Network Next Hop Metric LocPrf Weight Path
*> 1.2.3.4/32 2.2.16.6 0 200 888 i
*m 2.2.13.3 2.2.13.3 0 100 888 i

Router#show ip route bgp
1.0.0.0/32 is subnetted, 1 subnets
B 1.2.3.4 [20/0] via 2.2.16.6, 00:01:46 [20/0] via 2.2.13.3, 00:01:46
    
```

Weight, LP, AS-Path, Origin, MED need to be the same.

Which of the following prefixes have been originated locally?

Where these networks learned by eBGP or iBGP?

```

R3#show ip bgp
...
Network Next Hop Metric LocPrf Weight Path
*> 1.0.0.0 10.1.1.1 0 100 0 i
*> 2.0.0.0 10.1.1.2 0 100 0 i
*> 3.0.0.0 0.0.0.0 0 32768
*> 4.0.0.0 10.1.1.4 0 100 0 i
    
```

R3#show ip bgp BGP table version is 17, local router ID is 192.168.3.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter, x best-external, a additional-path, c RIB-compressed, Origin codes: i - IGP, e - EGP, ? - incomplete RPKI validation codes: V valid, I invalid, N Not found

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 1.0.0.0	10.1.1.1	0	100	0	i
*> 2.0.0.0	10.1.1.2	0	100	0	i
*> 3.0.0.0	0.0.0.0	0		32768	
*> 4.0.0.0	10.1.1.4	0	100	0	i

Prefixes learned by iBGP, due to Local Pref visible in the output!

Where these networks learned by eBGP or iBGP?

```

R2#show ip bgp
...
Network Next Hop Metric LocPrf Weight Path
* 1.0.0.0 10.1.1.1 0 300 100 i
* 10.1.1.1 10.1.1.1 0 400 100 i
* 10.1.1.1 0 0 100 i
*> 2.0.0.0 0.0.0.0 0 32768
*> 3.0.0.0 10.1.1.3 0 400 300 i
* 10.1.1.3 0 0 100 300 i
* 10.1.1.3 0 0 300 i
* 4.0.0.0 10.1.1.4 0 100 400 i
* 10.1.1.4 0 0 300 400 i
*> 10.1.1.4 0 0 400 i
    
```

Routes where learned from eBGP identified by a zero Local Preference field!
The Metric of 0 only shows up from Peers where that peer originates that prefix!

```

R2#show ip bgp
...
Network Next Hop Metric LocPrf Weight Path
* 1.0.0.0 10.1.1.1 0 300 100 i
* 10.1.1.1 10.1.1.1 0 400 100 i
* 10.1.1.1 0 0 100 i
*> 2.0.0.0 0.0.0.0 0 32768
*> 3.0.0.0 10.1.1.3 0 400 300 i
* 10.1.1.3 0 0 100 300 i
* 10.1.1.3 0 0 300 i
* 4.0.0.0 10.1.1.4 0 100 400 i
* 10.1.1.4 0 0 300 400 i
*> 10.1.1.4 0 0 400 i
    
```

Which route will be seen with *> in the BGP table of R2 in regards to 3.3.3.3?

BGP route-reflector

```

R2#show ip bgp 33.33.33.0
BGP routing table entry for 33.33.33.0/24, version 10
Paths: (1 available, best #1, table default)
Not advertised to any peer
Refresh Epoch 1
Local
13.1.1.3 (metric 128) from 12.1.1.1 (1.1.1.1)
Origin IGP, metric 0, localpref 100, valid, internal, best
Originator: 3.3.3.3, Cluster list: 1.1.1.1
Refresh Epoch 2
Local
13.1.1.3 (metric 128) from 13.1.1.3 (3.3.3.3)
Origin IGP, metric 0, localpref 100, valid, internal, best
    
```

Output of show ip bgp 33.33.33.0 On R2 and the RR:

How can you filter out any originated prefixes originated by R5 and R6 without using Route-maps, ACLs / Prefix lists on RR-1?

```

RR-1#show ip bgp
Network Next Hop Metric LocPrf Weight Path
*> 1.0.0.0 0.0.0.0 0 32768 i
*> 2.0.0.0 12.1.1.2 0 100 0 i
*> 3.0.0.0 13.1.1.3 0 100 0 i
*> 4.0.0.0 14.1.1.4 0 100 0 i
*> 5.0.0.0 45.1.1.5 0 100 0 i
*> 6.0.0.0 46.1.1.6 0 100 0 i
    
```

Filter them out

Keyword BGP Backdoor:

Configure R3 and R2 so that traffic from Loopback to loopback takes the ethernet interfaces. Use two different configuration on both sides!

```

R3#
router bgp 300
network 2.0.0.0 mask 255.255.255.0 backdoor

R2#
router bgp 200
distance 91 12.1.1.1 0.0.0.0 88
access-list 88 permit 3.0.0.0
    
```

What are the values for the following well-known BGP communities?

Internet no-export no-advertise local-as

Wellknown BGP community range: 4294901760 - 4294967295

Community	Value
Internet	does not have a value
no-export	4294967041
no-advertise	4294967042
local-as	4294967043

What to keep in mind with BGP in the CCIE lab / real world?

```

ip prefix-list PFX-2 permit 2.2.2.0/24
ip as-path access-list 500 permit *200$

route-map RMP-2-IN permit 10
match ip address prefix-list PFX-2
match as-path 500
SET / DO what ever

route-map RMP-2-OUT permit 20
    
```

How can you influence AS-100 to use the path with the lowest metric in this situation?

AS-400 seen from AS-100: (metric 100) 200 300 400 (metric 150) 400

clear ip bgp * (hard clear the session!!)

What are 3 different examples of filtering methods in BGP?

```

router bgp 100
neighbor 1.1.1.1 distribute-list 88 in
access-list 88 deny 2.0.0.0 0.0.0.0
access-list 88 permit any

router bgp 100
neighbor 1.1.1.1 prefix-list PFX-1 in
ip prefix-list PFX-1 deny 2.0.0.0/24
ip prefix-list PFX-1 permit 0.0.0.0/0 le 32

router bgp 100
neighbor 1.1.1.1 filter-list 77 in
ip as-path access-list 77 deny _300$
ip as-path access-list 77 permit *
    
```

BGP regex Utilizing (\1):

```

AS-200
R2
  bgp confederation identifier 900
  neighbor R1.R1.R1.R1 remote-as 65511

AS-100
R1
  bgp confederation identifier 900
  neighbor R2.R2.R2.R2 remote-as 65511
    
```

R1#show ip bgp regex ^([0-9]+)(\1)*\$

(\1) stores, what ever was discovered as directly attached neighbor AS!

* allows the prepend of one more instance of ANY prepended number.

show ip bgp regex ^([0-9]+)(\1)*\$

+ only allows 200 200, the prepended number needs to be the same!

How will the config look in regards to the BGP confederation?

```

R1#
router bgp 65511
  bgp confederation identifier 900
  neighbor R2.R2.R2.R2 remote-as 65511

R2#
router bgp 65511
  bgp confederation identifier 900
  bgp confederation peers 65522
  neighbor R1.R1.R1.R1 remote-as 65511
  neighbor R3.R3.R3.R3 remote-as 65522

R3#
router bgp 65522
  bgp confederation identifier 900
  bgp confederation peers 65511
  neighbor R2.R2.R2.R2 remote-as 65511
  neighbor R4.R4.R4.R4 remote-as 65522
    
```

Confed peer not configured!

Confederation peers only configured where necessary!

How will a show ip bgp look from router R1 in this BGP confederation?

```

R1#show ip bgp
Network Next Hop Metric LocPrf Weight Path
*> 1.0.0.0 0.0.0.0 0 32768 i
*> 2.0.0.0 12.1.1.2 0 100 0 (65511) i
*> 3.0.0.0 23.1.1.3 0 100 0 (65522) i
*> 4.0.0.0 34.1.1.4 0 100 0 (65522) i
*> 5.0.0.0 45.1.1.5 0 100 0 (65522 65533) i
*> 6.0.0.0 45.1.1.6 0 100 0 (65522 65533) 123 i
    
```

show ip bgp template peer-session

```

Template:COMMON, index:1
Local policies:0x12, Inherited policies:0x0
*Inherited by Template iBGP, index: 2
Locally configured session commands:
  version 4
  password is configured
Inherited session commands:
  Template:iBGP, index:2
Local policies:0x81, Inherited policies:0x12
This template inherits:
COMMON index:1 flags:0x0
Locally configured session commands:
  update-source Loopback0
Inherited session commands:
  version 4
  password is configured
    
```

router bgp 100 network 10.0.0.0 mask 255.0.0.0 template peer-session ONE version 4 password cisco template peer-session TWO update-source loopback 0 inherit peer-session ONE

BGP Session templates

BGP Policy templates

advertisement-interval allows-in as-override capability default-originate distribute-list dmzlink-bw filter-list maximum-prefix next-hop-self next-hop-unchanged prefix-list remove-private-as route-map route-reflector-client send-community send-label soft-reconfiguration soo unsuppress-map weight

Session templates allows-in description disable-connected-check ebgp-multipath fall-over local-as password remote-as shutdown timers translate-update transport ttl-security update-source version

BGP Route-reflector cluster:

All IGP metrics are the same. Which path will R2 choose to reach 3.0.0.0/8?

13th Prefer the path that comes from the lowest neighbor address

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<h3>BGP Inject-map</h3> <h4>Inject ANY prefix into BGP</h4> <p>R1 sends a default route to R2</p> <p>Inject network 99.99.99.0/24 which is NOT configured on R2, into R2s BGP table. R3 will see the route but will not be able to reach it.</p>	<p>R2 has ALL routes in table due the default route, R2 can inject ANY route specified by the INJECT route-map into BGP!</p> <pre> R2# router bgp 200 bgp inject-map INJECT exist-map EXIST neighbor 12.0.0.1 remote-as 100 neighbor 13.0.0.3 remote-as 300 route-map INJECT permit 10 SET ip address prefix INJECT-THIS-NETWORK route-map EXIST match ip address prefix DEFAULT match ip route-source prefix R1-NEXT-HOP ip prefix-list DEFAULT permit 0.0.0.0/0 ip prefix-list INJECT-THIS-NETWORK 99.99.99.0/24 ip prefix-list R1-NEXT-HOP permit 12.0.0.1/32 </pre>	<pre> R1# neighbor 12.1.1.2 send-community neighbor 12.1.1.2 route-map RMP-SET-COM route-map RMP-SET-COM match ip address 1 set community no-export access-list 1 permit 1.0.0.0 0.0.0.255 </pre> <p>Ensure that 1.0.0.0/8 is in R4 routing table, do NOT create/modify existing route-maps:</p>	<p>R2# Router bgp 200 no neighbor 23.1.1.3 send-community</p> <p>By disabling communities from R2 to R3, the community "no-export" set by R1 will be removed. Therefore R2 will advertise 1.0.0.0/8 to R4</p>	<h3>BGP communities / tags</h3> <p>Ensure R3 takes takes path via R1 towards 4.0.0.0/8 And R2 towards 8.0.0.0/8</p> <p>Use communities to accomplish this task:</p>	<pre> router bgp 100 neighbor 2.2.2.2 remote-as 100 neighbor 2.2.2.2 route-map RMP-R2-IN in route-map RMP-R2-IN permit 10 match community 4 set ip next-hop 2.2.2.2 route-map RMP-R2-IN permit 10 R1 match community 4 set ip next-hop to R1 router bgp 100 neighbor 4.4.4.4 remote-as 400 neighbor 4.4.4.4 send-community neighbor 4.4.4.4 route-map RMP-R4-IN in route-map RMP-R4-IN permit 10 match ip address 4 set community 4 route-map RMP-R4-IN permit 20 access-list 4 permit 4.0.0.0 0.255.255.255 </pre>
<h3>BGP Aggregation</h3> <p>Configure two solutions on how to aggregate the R1 networks on R2 towards R3.</p>	<p>Solution 1</p> <pre> R2# router bgp 200 aggregate-address 10.0.0.0 255.255.0.0 summary-only </pre> <p>Solution 2</p> <pre> R2# ip route 10.0.0.0 255.255.0.0 null 0 router bgp 200 network 10.0.0.0 mask 255.255.0.0 </pre>	<p>Ensure that 1.0.0.0/8 is in R4 routing table. You are only allowed to change R2 config:</p> <pre> R1# neighbor 12.1.1.2 send-community neighbor 12.1.2 route-map RMP-SET-COM route-map RMP-SET-COM match ip address 1 set community no-export access-list 1 permit 1.0.0.0 0.255.255.255 </pre>	<pre> R1 sets community no-export R2# router bgp 200 neighbor 12.1.1.1 send-community neighbor 12.1.1.1 route-map RMP-R1-IN in route-map RMP-R1-IN match ip address 1 set community internet access-list 1 permit 1.0.0.0 0.255.255.255 </pre>		
<h3>R2#</h3> <h4>show ip bgp community no-advertise</h4> <pre> R2# router bgp 100 aggregate-address 10.0.0.0 255.255.252.0 summary-only as-set </pre> <p>show ip bgp ?</p>	<p>Easy way of finding why the summary is not advertised from R2 to R1 when using aggregate-address and as-set!</p> <pre> R3#show ip bgp community no-advertise ... Network Next Hop Metric LocPrf Weight Path *> 10.0.2.0/24 35.1.1.5 0 0 19 i *-> 10.0.0.0/22 0.0.0.0 100 32768 {22,57,19} i </pre>				
<h3>Give two solutions on how to fix the missing summary on R1</h3> <p>show ip bgp ?</p>	<p>Solution 1</p> <p>Change outbound route-map of AS-19 NOT to set community no-advertise.</p> <p>Solution 2</p> <pre> R2# router bgp 35 aggregate-address 10.0.0.0 255.255.252.0 summary-only as-set attribute-map ATTR-MAP route-map ATTR-MAP permit 10 set community internet </pre>				
<h3>What do the following commands do?</h3> <pre> router bgp 300 bgp log-neighbor-changes aggregate-address 10.1.0.0 255.255.252.0 as-set summary-only attribute-map ATTR-MAP advertise-map ADVERTISE-MAP </pre>	<p>Includes all AS numbers of all specific routes of the summary.</p> <p>Will only send the summary address and suppress the more specific routes</p> <p>Sets attributes such as communities etc</p> <pre> route-map ATTR-MAP permit 10 set community internet </pre> <p>Can be used to filter out either single prefixes and their attached attributes (no-export with AS-SET!) or entire AS numbers out of the summary (with or without no-advertise / no-export attributes)</p> <pre> route-map ADVERTISE-MAP permit 10 match as-path 400 (if you only want to have AS 400 in AS-SET) or match ip address prefix-list PFX-MATCH-ONLY-THIS-INTO-AS-SET </pre>				
<p>Unsuppress 10.0.1.0/24 to R7 but not to R1</p>	<pre> R2# router bgp 100 aggregate-address 10.0.0.0 255.255.252.0 summary-only as-set R2# router bgp 35 neighbor 7.7.7.7 unsuppress-map DONT-SUPPRESS route-map DONT-SUPPRESS match ip address 2 access-list 2 permit 10.0.1.0 0.0.0.255 </pre>				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

Multicast

<p>show ip pim neighbor:</p>	<pre>R5#show ip pim neighbor PIM Neighbor Table Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority, 5 - State Refresh Capable Neighbor Interface Uptime/Expires Ver DR Address Interface Uptime/Expires Pri/Mode 155.1.45.4 Serial0/1/0 00:01:38/00:01:36 v2 1/S 155.1.58.8 GigabitEthernet0/0 00:01:11/00:01:32 v2 1/DR S</pre>	<p>ip igmp join-group:</p> <p>ip igmp static-group:</p>	<p>ip igmp join-group</p> <p>Simulates a router acting as a client which is attached/joining a group. Will answer to pings to the group. (*, 224.1.1.1), 03:50:11/stoppped, RP 0.0.0.0, flags: DC</p> <p>ip igmp static-group</p> <p>Will make sure the router is accepting the group and is forwarding traffic for that group but it will NOT answer ICMPs for that group. Can be rate-limited only on Serial interfaces.</p>	<p>DVMRP explained:</p>	<p>DVMRP:</p> <ul style="list-style-type: none"> - Dense-Mode (S,G) - Infinity 32 hops 												
<p>Show ip pim interface</p>	<pre>R5#show ip pim interface Address Interface Ver/ Mode Nbr Query DR DR 155.1.45.5 Serial0/1/0 v2/D 1 30 1 0.0.0.0 155.1.58.5 GigabitEthernet0/0 v2/D 1 30 1 155.1.58.8 155.1.62.3 GigabitEthernet0/1 v2/S 1 30 1 155.1.62.1 155.1.88.2 GigabitEthernet0/2 v2/SD 1 30 1 155.1.88.4</pre> <p>Dense-Mode Sparse-Mode Sparse-Dense-Mode</p>	<p>Show ip igmp group:</p>	<p>Group is known for:</p> <p>155.1.58.5 was the last "client" requesting that group</p> <pre>R5#show ip igmp group IGMP Connected Group Membership Group Address Interface Uptime Expires Last Reporter 224.99.99.99 Gi0/0 03:04:17 00:02:49 155.1.58.5 224.0.1.40 Gi0/0 04:10:06 00:02:46 155.1.58.8</pre>	<p>IGMPv1:</p>	<p>Hosts leave a group quietly Querier is eleted via DR Time-out = 3x query interval 3x 60 = 180 seconds</p> <p>The difference from IGMPv1 and v2 can be seen in the Maximum-Response-Time field which is always set to Zero in Version 1.</p> <p>Where as Version 2 it's a non-zero value.</p>												
<p>Multicast Test setup:</p>	<p>Multicast Source Receiver</p> <pre>R4# int fa0/x ip igmp join-group 224.11.22.33 R3# ping 224.11.22.33 rep 100</pre> <p>Or use a IP SLA for a constant multicast stream. Only Group Members should respond to the sent ping</p>	<p>Multicast</p> <p>Source Tree:</p> <p>Shared Tree:</p>	<p>Source Tree (S,G) known as SPT</p> <p>Shared Tree (*,G) via RP known as RTP</p>	<p>IGMPv2:</p>	<p>Non Zero Maximum-Response-Time field. Group specific joins and leaves Querier election when router starts by sending 224.0.0.1 general-query. Lowest IP is querier.</p>												
<p>Show ip mroute described:</p> <p>DENSE-MODE</p>	<p>Multicast Source Receiver</p> <p>DENSE-MODE example</p> <pre>SW4#sh ip mroute (*, 224.10.10.10), 00:11:44/00:02:49, RP 0.0.0.0, flags: DCL Incoming interface: Null, RPF nbr 0.0.0.0 Outgoing interface list: Port-channel1, Forward/Dense, 00:11:33/00:00:00 Vlan10, Forward/Dense, 00:11:44/00:00:00</pre> <p>D: Dense Mode C: Directly connected L: router itself is a member</p>	<p>IGMPv3</p>	<p>Join / Prunes include the source of the group:</p> <p>Join (192.168.77.22, 224.1.1.1)</p> <p>Prune (192.168.59.32, 224.1.1.1)</p>	<p>PIM Dense-Mode:</p>	<p>Depends on unicast routing table (RPF check)</p> <p>RPF check towards source, checks AdminDistance/Metric If several path exist, interface with highest IP is used.</p> <p>Neighbor discovery via 224.0.0.13, Hello Period 30 sec</p> <p>DR election in PIMv1 Dense-Mode, Highest IP wins.</p> <p>Join/Prune Prune override Graft Assert-Message</p> <p>No flags set WC,RP flags set</p>												
<p>Show ip igmp interface X/X:</p>	<pre>R5#sh ip igmp interface gi0/0 GigabitEthernet0/0 is up, line protocol is up Internet address is 155.1.58.5/24 IGMP is enabled on interface Current IGMP host version is 2 Current IGMP router version is 2 IGMP query interval is 60 seconds IGMP querier timeout is 120 seconds IGMP max query response time is 10 seconds Last member query count is 2 Last member query response interval is 1000 ms Inbound IGMP access group is not set IGMP activity: 2 joins, 0 leaves Multicast routing is enabled on interface Multicast TTL threshold is 0 Multicast designated router (DR) is 155.1.58.8 IGMP querying router is 155.1.58.5 (this system) Multicast groups joined by this system (number of users): 224.0.1.40(1) 224.99.99.99(1)</pre>	<p>PIM Spare Mode Register Messages</p>	<p>PIM Sparse-Mode Register</p> <p>DR Encapsulates Register-Message in a Unicast to the RP. RP then sends (S,G) join to the source. -> forces a RPT to SPT switchover</p>	<p>PIM Assert message Explained:</p> <p>Higher IP wins, if both have the same metric</p>	<p>PIM Assert message</p> <p>Higher IP wins, if both have the same metric</p>												
<p>Multicast flags DENSE-MODE</p> <table border="1"> <tr> <td>(*,G) ? (S,G) ?</td> <td>(*,G) ? (S,G) ?</td> <td>nothing</td> </tr> <tr> <td>(*,G) RPF (S,G) RPF int</td> <td>(*,G) RPF (S,G) RPF int</td> <td>nothing</td> </tr> </table>	(*,G) ? (S,G) ?	(*,G) ? (S,G) ?	nothing	(*,G) RPF (S,G) RPF int	(*,G) RPF (S,G) RPF int	nothing	<p>Multicast flags DENSE-MODE</p> <table border="1"> <tr> <td>(*,G) DCL (S,G) LT</td> <td>(*,G) D (S,G) T</td> <td>nothing</td> </tr> <tr> <td>(*,G) RPF 0.0 (S,G) RPF int -></td> <td>(*,G) RPF 0.0 (S,G) RPF int -></td> <td>nothing</td> </tr> </table>	(*,G) DCL (S,G) LT	(*,G) D (S,G) T	nothing	(*,G) RPF 0.0 (S,G) RPF int ->	(*,G) RPF 0.0 (S,G) RPF int ->	nothing	<p>PIM Dense-Mode</p> <p>Which router should be checked for the RPF failure in this PIM DM scenario?</p>	<p>PIM Dense-Mode</p> <p>The unicast table is using E0 to reach 2.2.2.2 which is a Non-PIM enabled interface/route to the source! -> Results in a RPF failure</p>	<p>MAC multicast address conversion:</p> <p>Convert the mac to the corresponding MultiCast IP address:</p> <p>0100.5e07.1925</p>	<p>0100.5e07.1925 = 238.7.25.37</p> <p>0100 0001 0000 0000 1001 1110 0000 0111 0001 1001 0010 0101 00000001.00000000.10011110.00000111.00011001.00100101 Dez: 1 0 .158 .7 .25 .37</p> <p>25 bits ← 23 bits</p> <p>1110 = Multicast Prefix</p> <p>11101110.00000111.00011001.00100101 = 238.7.25.37</p> <p>32 bits</p> <p>00000 - 11110 11100000.0xxxx111.x = 224.7.xx.xx 11101111.0xxxx111.x = 229.7.xx.xx Range 00001 - 11111 11100000.1xxxx111.x = 224.135.xx.xx 11101111.1xxxx111.x = 239.135.xx.xx</p>
(*,G) ? (S,G) ?	(*,G) ? (S,G) ?	nothing															
(*,G) RPF (S,G) RPF int	(*,G) RPF (S,G) RPF int	nothing															
(*,G) DCL (S,G) LT	(*,G) D (S,G) T	nothing															
(*,G) RPF 0.0 (S,G) RPF int ->	(*,G) RPF 0.0 (S,G) RPF int ->	nothing															

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

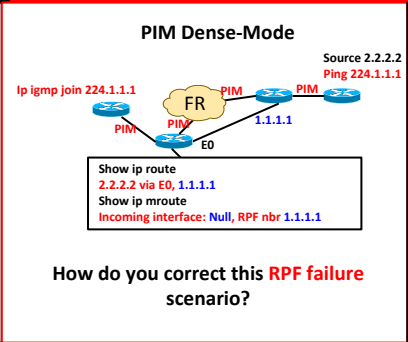
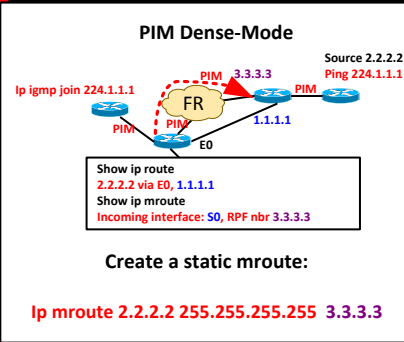
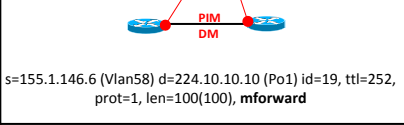
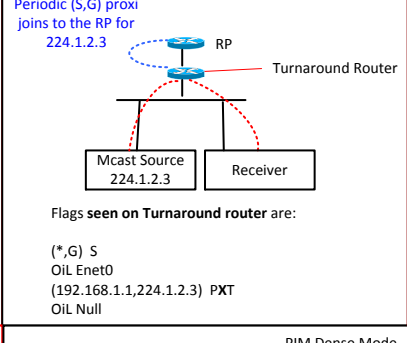
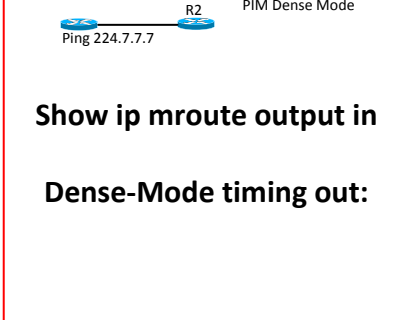
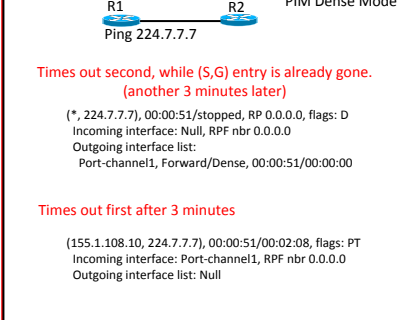
Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin



Multicast

<p>PIMv2 Messages:</p> <ol style="list-style-type: none"> 0 Hello 1 Register 2 Register-Stop 3 Join/Prune <small>(Join WC/RP flag set. Prune has none set)</small> 4 Bootstrap 5 Assert 6 Graft (PIM DM only) 7 Graft-Ack (PIM only) 8 Candidate-RP-Advertisement 	<p>Explain the command</p> <p>Ip pim sparse-dense-mode:</p>	<p>ip pim sparse-dense-mode</p> <p>Will perform in Sparse mode for all groups where there is an Group/RP binding known / existing.</p> <p>For all others it will flood and prune using IP PIM Dense-Mode.</p>	<p>PIM Dense-Mode</p>  <p>How do you correct this RPF failure scenario?</p>	<p>PIM Dense-Mode</p>  <p>Create a static mroute:</p> <p>Ip mroute 2.2.2.2 255.255.255.255 3.3.3.3</p>
<p>Show ip mroute Dense-Mode (*,G) Explained:</p>	<p>Difference between (*,G) in PIM Dense-Mode And (*,G) in PIM Sparse-Mode</p>	<p>Dense-Mode: Traffic is never forwarded according to (*,G), a separate (S,G) needs to be created to forward traffic. OIL is copied from (*,G) to (S,G) which triggers flood and prune behavior.</p> <p>Sparse-Mode: (*,G) is used for forwarding multicast traffic. (*,G) is a result of an explicit join of a PIM neighbor or a directly connected Host. The incoming interface of (*,G) always points to the RP.</p>	<p>How to troubleshoot Multicast problems using Debug ip mpacket:</p>	<pre>Conf t Int x no ip mroute-cache End debug ip mpacket</pre>  <p>s=155.1.146.6 (Vlan58) d=224.10.10.10 (Po1) id=19, ttl=252, prot=1, len=100(100), mforward</p>
<p>MOSPF explained:</p>	<p>IP PIM Sparse Mode Flags:</p>	<p>DR originates group membership LSA's (S,G) whereas S refers to the source subnet not the source ip!</p> <p>Each time there is a topology change MOSPF must flush the entire forwarding cache!</p>	<p>IP PIM Sparse Mode Flags:</p> <p>S indicates sparse-mode C directly connected member for this group attached L router itself is a member of the group P causes to prune to be sent to the upstream RPF neighbor T (S,G) only: traffic is being forwarded X Proxy-Join Timer is running J (*,G) traffic rate is exceeding the SPT-threshold, will switch over to (S,G). J (S,G) too little traffic, will switch back to (*,G) R Prune (*,G) traffic to the RP and start direct (S,G) traffic</p>	<p>delay their RPF interface re-computation after topology change: specifies the amount of milliseconds to wait after the topology change to re-calculate RPF interfaces</p> <p>multicast rpf backoff <min-delay ms> <max-delay ms></p> <p><min-delay ms> <max-delay ms></p> <p><min-delay> gets doubled after every consequent topology change but never gets higher than specified in the <max-delay ms></p>
<p>PIM Dense-Mode (*,G) OIL explained On show ip mroute:</p>	<p>IP PIM Dense Mode Flags</p>	<p>IP PIM Dense-Mode flags</p> <p>D Dense Mode C directly connected Member attached L router itself is a member to this group P OIL is Null and a Prune is sent upstream towards the RPF interface T Traffic is forwarding via (S,G) J Used internally, tells (*,G) to create a (S,G)</p>	<p>How can you tune the Multicast RPF computation?</p>	<p>ip multicast rpf interval <in seconds></p>
<p>PIM Dense-Mode (*,G) and (S,G) Show ip mroute explained:</p>	<p>IP PIM Sparse-Mode Turnaround Router explained:</p>	<p>Periodic (S,G) proxy joins to the RP for 224.1.2.3</p>  <p>Flags seen on Turnaround router are:</p> <p>(*,G) S OIL Enet0 (192.168.1.1,224.1.2.3) PXT OIL Null</p>	<p>Debugging IP multicast traffic</p> <p>Using: no ip mroute-cache debug ip mpackets</p> <p>Finding RPF failures:</p>	<p>s=155.1.146.6 (Serial0/0/0) d=224.10.10.10 id=21, ttl=253, prot=1, len=104(100), RPF lookup failed for source</p> <p>s=155.1.146.6 (Serial0/0/0) d=224.10.10.10 id=21, ttl=253, prot=1, len=104(100), not RPF interface</p>
<p>Configuring IP PIM Sparse-Mode Statically:</p> <p>ip multicast-routing</p> <p>interface X ip pim sparse-mode</p> <p>ip pim rp-address X.X.X.X</p> <p>RP address tells the router where to send the (*,G) joins to.</p>	<p>Show ip mroute output in Dense-Mode timing out:</p>	<p>PIM Dense Mode</p>  <p>Ping 224.7.7.7</p> <p>Times out second, while (S,G) entry is already gone. (another 3 minutes later)</p> <p>(*, 224.7.7.7), 00:00:51/stopped, RP 0.0.0.0, flags: D Incoming interface: Null, RPF nbr 0.0.0.0 Outgoing interface list: Port-channel1, Forward/Dense, 00:00:51/00:00:00</p> <p>Times out first after 3 minutes</p> <p>(155.1.108.10, 224.7.7.7), 00:00:51/00:02:08, flags: PT Incoming interface: Port-channel1, RPF nbr 0.0.0.0 Outgoing interface list: Null</p>	<p>IP PIM Sparse-Mode Router manually joined</p>  <p>NO SOURCE</p> <p>R4# int fa0/x ip igmp join-group 224.10.10.10</p>	<p>R4# int fa0/x ip igmp join-group 224.10.10.10</p> <p>NO SOURCE</p> <p>R4# (*, 224.10.10.10), 00:09:33/00:02:11, RP 0.0.0.0, flags: SJCL Incoming interface: Null, RPF nbr 0.0.0.0 Outgoing interface list: Vlan10, Forward/Sparse, 00:09:33/00:02:11</p> <p>Sparse-Mode J reached threshold, will switchover to S,G Connected L Router is member of group</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

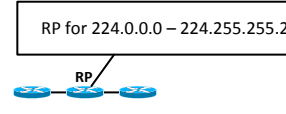
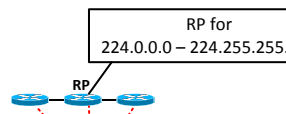
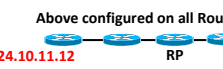


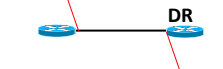
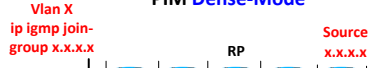

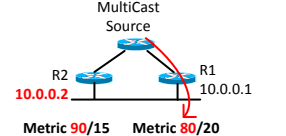
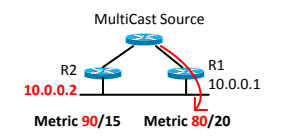
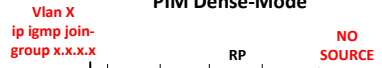

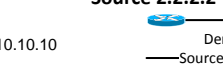
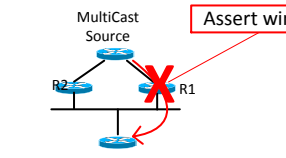
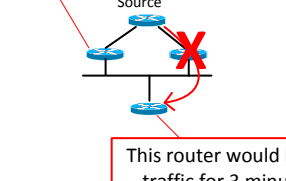
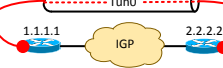
Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin


Multicast

<p>Explain the following command:</p> <pre>ip pim rp-address <IP> <ACL> <override></pre>	<p>ip pim rp-address x.x.x.x ACL override</p> <p>Sets the RP address for groups specified in the Access-Lists. Keyword override will prioritise static configuration versus dynamically learned Group/RP mappings!</p>	<p>PIM Sparse-Dense-Mode</p>  <p>Setup config so that Routers MultiCast in Sparse-Mode for any groups within 224.0.0.0 - 224.255.255.255,</p> <p>All others should be using dense-mode</p>	<p>PIM Sparse-Dense-Mode</p>  <pre>ip pim rp-address 150.1.5.5 SPARSE_GROUPS ip access-list standard SPARSE_GROUPS permit 224.0.0.0 0.255.255.255</pre>	<h2>PIM Accept RP</h2> <p>Allows only groups specified from RP</p>	<pre>ip pim accept-rp 150.1.5.5 ALLOWED_GRP ip access-list standard ALLOWED_GRP permit 224.10.10.10 permit 224.110.110.110</pre>  <p>%PIM-6-INVALID_RP_JOIN: Received (*, 224.10.11.12) Join from 0.0.0.0 for invalid RP 150.1.5.5</p>																																																																																	
<p>Explain:</p> <h2>ip pim spt-threshold</h2>	<p>ip pim spt-threshold <bandwidth in kilobits/s></p> <p>Will switchover from (*,G) to (S,G) entry if the bandwidth specified is reached.</p> <p>ip pim spt-threshold infinity</p> <p>Will never switch to from a (*,G) to (S,G)</p>	<h2>PIM Assert messages over NBMA networks:</h2> <h3>Dangerous!</h3>	<p>xxxxxx</p>	<h2>Forcing a group to revert to Dense-Mode Operation in a Sparse-Dense-Mode setup for specific groups:</h2>	<pre>ip pim sparse-dense-mode igmp join 224.10.11.12</pre> <p>Setting the range for Sparse Mode:</p> <pre>ip pim rp-address 150.1.5.5 SPARSE_GROUPS ip access-list standard SPARSE_GROUPS permit 224.0.0.0 0.255.255.255</pre> <p>ip pim accept-rp x.x.x.x ALLOWED_GRP</p> <p>ip access-list standard ALLOWED_GRP</p> <pre>DENY 224.10.11.12</pre> <p>Disabling the allowed Group to RP mapping results in Dense-Mode operation for Group 224.10.11.12</p>																																																																																	
<p>PIM Sparse-Mode</p>  <table border="1"> <tr><td>(*,G) flags</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) i-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) o-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) flags</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) i-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) o-list</td><td></td><td></td><td></td><td></td></tr> </table>	(*,G) flags					(*,G) i-list					(*,G) o-list					(S,G) flags					(S,G) i-list					(S,G) o-list					<p>PIM Sparse-Mode</p>  <table border="1"> <tr><td>(*,G) flags</td><td>SJCL</td><td>S</td><td>S</td><td>SP</td><td>SPF</td></tr> <tr><td>(*,G) i-list</td><td>→</td><td>→</td><td>Null</td><td>←</td><td>←</td></tr> <tr><td>(*,G) o-list</td><td>←</td><td>←</td><td>←</td><td>Null</td><td>Null</td></tr> <tr><td>(S,G) flags</td><td>LJT</td><td>T</td><td>T</td><td>T</td><td>PFT</td></tr> <tr><td>(S,G) i-list</td><td>→</td><td>→</td><td>→</td><td>→</td><td>SRC int</td></tr> <tr><td>(S,G) o-list</td><td>←</td><td>←</td><td>←</td><td>←</td><td>Null</td></tr> </table>	(*,G) flags	SJCL	S	S	SP	SPF	(*,G) i-list	→	→	Null	←	←	(*,G) o-list	←	←	←	Null	Null	(S,G) flags	LJT	T	T	T	PFT	(S,G) i-list	→	→	→	→	SRC int	(S,G) o-list	←	←	←	←	Null	<p>Multicast with different routing domains within the same MultiCast Domain!</p>	<p>xxxxxx</p>	<h2>PIM DR Election</h2>	<pre>interface FastEthernet0/0 ip pim dr-priority 0</pre>  <pre>interface FastEthernet0/0 ip pim dr-priority 100</pre>															
(*,G) flags																																																																																						
(*,G) i-list																																																																																						
(*,G) o-list																																																																																						
(S,G) flags																																																																																						
(S,G) i-list																																																																																						
(S,G) o-list																																																																																						
(*,G) flags	SJCL	S	S	SP	SPF																																																																																	
(*,G) i-list	→	→	Null	←	←																																																																																	
(*,G) o-list	←	←	←	Null	Null																																																																																	
(S,G) flags	LJT	T	T	T	PFT																																																																																	
(S,G) i-list	→	→	→	→	SRC int																																																																																	
(S,G) o-list	←	←	←	←	Null																																																																																	
<p>PIM Dense-Mode</p>  <table border="1"> <tr><td>(*,G) flags</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) i-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) o-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) flags</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) i-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) o-list</td><td></td><td></td><td></td><td></td></tr> </table>	(*,G) flags					(*,G) i-list					(*,G) o-list					(S,G) flags					(S,G) i-list					(S,G) o-list					<p>PIM Dense-Mode</p>  <table border="1"> <tr><td>(*,G) flags</td><td>DCL</td><td>D</td><td>D</td><td>D</td><td></td></tr> <tr><td>(*,G) i-list</td><td>Null</td><td>Null</td><td>Null</td><td>Null</td><td></td></tr> <tr><td>(*,G) o-list</td><td>←</td><td>←</td><td>←</td><td>←</td><td></td></tr> <tr><td>(S,G) flags</td><td>LT</td><td>T</td><td>T</td><td>T</td><td></td></tr> <tr><td>(S,G) i-list</td><td>→</td><td>→</td><td>→</td><td>→</td><td></td></tr> <tr><td>(S,G) o-list</td><td>←</td><td>←</td><td>←</td><td>←</td><td></td></tr> </table>	(*,G) flags	DCL	D	D	D		(*,G) i-list	Null	Null	Null	Null		(*,G) o-list	←	←	←	←		(S,G) flags	LT	T	T	T		(S,G) i-list	→	→	→	→		(S,G) o-list	←	←	←	←		<p>Troubleshooting PIM assert winners:</p> 	<p>debug ip pim</p> <pre>PIM(0): Received v2 Assert on FastEthernet0/0 from 155.1.146.4 PIM(0): Assert metric to source 155.1.108.10 is [90/2174976] PIM(0): We win, our metric [80/20]</pre> 	<h2>PIM DR Election</h2> <p>Verification:</p>	<pre>show ip pim interface fastEthernet 0/0 detail</pre> <pre>... PIM: enabled PIM version: 2, mode: sparse-dense PIM DR: 155.1.146.1 (this system) PIM neighbor count: 2 ...</pre> <pre>R6#show ip pim neighbor</pre> <table border="1"> <thead> <tr><th>Neighbor Address</th><th>Interface</th><th>Uptime/Expires</th><th>Ver</th><th>DR Prio/Mode</th></tr> </thead> <tbody> <tr><td>1.1.1.1</td><td>Gi0/0.146</td><td>00:27:04/00:01:19</td><td>v2</td><td>255/ DR S</td></tr> <tr><td>2.2.2.2</td><td>Gi0/0.146</td><td>02:44:56/00:01:43</td><td>v2</td><td>1/ S</td></tr> </tbody> </table>	Neighbor Address	Interface	Uptime/Expires	Ver	DR Prio/Mode	1.1.1.1	Gi0/0.146	00:27:04/00:01:19	v2	255/ DR S	2.2.2.2	Gi0/0.146	02:44:56/00:01:43	v2	1/ S
(*,G) flags																																																																																						
(*,G) i-list																																																																																						
(*,G) o-list																																																																																						
(S,G) flags																																																																																						
(S,G) i-list																																																																																						
(S,G) o-list																																																																																						
(*,G) flags	DCL	D	D	D																																																																																		
(*,G) i-list	Null	Null	Null	Null																																																																																		
(*,G) o-list	←	←	←	←																																																																																		
(S,G) flags	LT	T	T	T																																																																																		
(S,G) i-list	→	→	→	→																																																																																		
(S,G) o-list	←	←	←	←																																																																																		
Neighbor Address	Interface	Uptime/Expires	Ver	DR Prio/Mode																																																																																		
1.1.1.1	Gi0/0.146	00:27:04/00:01:19	v2	255/ DR S																																																																																		
2.2.2.2	Gi0/0.146	02:44:56/00:01:43	v2	1/ S																																																																																		
<p>PIM Dense-Mode</p>  <table border="1"> <tr><td>(*,G) flags</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) i-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) o-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) flags</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) i-list</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) o-list</td><td></td><td></td><td></td><td></td></tr> </table>	(*,G) flags					(*,G) i-list					(*,G) o-list					(S,G) flags					(S,G) i-list					(S,G) o-list					<p>PIM Dense-Mode</p>  <table border="1"> <tr><td>(*,G) flags</td><td>DCL</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) i-list</td><td>Null</td><td></td><td></td><td></td><td></td></tr> <tr><td>(*,G) o-list</td><td>←</td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) flags</td><td></td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) i-list</td><td></td><td></td><td></td><td></td><td></td></tr> <tr><td>(S,G) o-list</td><td></td><td></td><td></td><td></td><td></td></tr> </table>	(*,G) flags	DCL					(*,G) i-list	Null					(*,G) o-list	←					(S,G) flags						(S,G) i-list						(S,G) o-list						<p>How to identify the PIM Assert winner in a show ip mroute output:</p>	<pre>(*, 239.6.6.6),00:02:05/stopped,RP 0.0.0.0, flags: DC Incoming interface: Null, RPF nbr 0.0.0.0 Outgoing interface list: FastEthernet0/0, Forward/Sparse-Dense, 00:02:05/00:00:00 Serial0/0.1, Forward/Sparse-Dense, 00:02:05/00:00:00 (155.1.108.10, 239.6.6.6),00:00:22/00:02:46, flags: T Incoming interface: Serial0/0.1, RPF nbr 155.1.0.5 Outgoing interface list: FastEthernet0/0, Forward/Sparse-Dense, 00:00:24/00:00:00, A</pre>	<h2>PIM Accept Register</h2> <p>(registering the source)</p>	<p>Source: Source 2.2.2.2 RP</p>  <pre>ping 224.10.10.10</pre> <p>Denies Source 2.2.2.2 For (S,G)</p> <p>RP Config:</p> <pre>ip pim accept-register route-map ACCEPT_REGISTER ACCEPT_REGISTER route-map ACCEPT_REGISTER deny 10 match ip address SOURCES</pre> <pre>ip access-list extended SOURCES permit ip host 2.2.2.2 any deny ip any any</pre> <p>%PIM-4-INVALID_SRC_REG: Received Register from 155.1.0.1 for (2.2.2.2, 224.10.10.10), not willing to be RP</p>															
(*,G) flags																																																																																						
(*,G) i-list																																																																																						
(*,G) o-list																																																																																						
(S,G) flags																																																																																						
(S,G) i-list																																																																																						
(S,G) o-list																																																																																						
(*,G) flags	DCL																																																																																					
(*,G) i-list	Null																																																																																					
(*,G) o-list	←																																																																																					
(S,G) flags																																																																																						
(S,G) i-list																																																																																						
(S,G) o-list																																																																																						
<p>How to figure out if PIM Sparse or Dense mode is used on an interface configured with IP PIM SPARSE-DENSE-MODE</p> <pre>(*, 239.0.0.1), 00:00:10/stopped, RP 0.0.0.0, flags: D Incoming interface: Null, RPF nbr 0.0.0.0 Outgoing interface list: Serial0/1/0, Forward/Sparse-Dense, 00:00:10/00:00:00 FastEthernet0/1, Forward/Sparse-Dense, 00:00:10/00:00:00 (155.1.146.6, 239.0.0.1), 00:00:10/00:02:55, flags: T Incoming interface: FastEthernet0/1, RPF nbr 0.0.0.0 Outgoing interface list: Serial0/1/0, Forward/Sparse-Dense, 00:00:12/00:00:00</pre>	<pre>(*, 239.0.0.1), 00:00:10/stopped, RP 0.0.0.0, flags: D Incoming interface: Null, RPF nbr 0.0.0.0 Outgoing interface list: Serial0/1/0, Forward/Sparse-Dense, 00:00:10/00:00:00 FastEthernet0/1, Forward/Sparse-Dense, 00:00:10/00:00:00 (155.1.146.6, 239.0.0.1), 00:00:10/00:02:55, flags: T Incoming interface: FastEthernet0/1, RPF nbr 0.0.0.0 Outgoing interface list: Serial0/1/0, Forward/Sparse-Dense, 00:00:12/00:00:00</pre>	<p>What happens when an Assert winner stops operation?</p> 	<p>The Assert loser needs to time-out the Assert winner!</p>  <p>This router would loose traffic for 3 minutes</p>	<h2>Multicast Tunneling</h2>	 <pre>interface Tunnel 0 ip unnumbered Loopback0 tunnel source Loopback0 tunnel destination 2.2.2.2 ip pim sparse-dense-mode</pre> <pre>interface Tunnel 0 ip unnumbered Loopback0 tunnel source Loopback0 tunnel destination 1.1.1.1 ip pim sparse-dense-mode</pre> <pre>ip mroute 0.0.0.0 0.0.0.0 tunnel 0</pre> <p>Use different IGP for tunnel, or make sure tunnel SRC/DST not learned through the tunnel!</p>																																																																																	

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin

Multicast

<p>PIM NBMA Mode (sparse-mode only)</p>		<p>Configuring a Auto-RP Candidate RP And Mapping Agent:</p>	<pre>interface Loopback0 ip pim sparse-dense-mode cRP: ip pim send-rp-announce Loopback 0 scope 10 ip pim send-rp-discovery loopback 0 scope 10 R5#show ip pim rp mapping PIM Group-to-RP Mappings This system is an RP (Auto-RP) This system is an RP-mapping agent (Loopback0) Group(s) 224.0.0.0/4 RP 150.1.5.5 (?), v2v1 Info source: 150.1.5.5 (?), elected via Auto-RP Uptime: 00:02:33, expires: 00:02:26</pre>	<p>Auto-RP - Filtering Candidate RPs</p>	<pre>ip access-list standard RP_LIST permit 150.1.10.10 ip access-list standard GROUP_LIST deny 224.110.110.110 permit any ip pim rp-announce-filter RP_LIST group-list GROUP_LIST</pre>
<p>PIM NBMA Mode (sparse-mode only) Show ip mroute output:</p>		<p>How to change the ? From the output below?</p> <pre>R5#show ip pim rp mapping ... Group(s) 224.0.0.0/4 RP 150.1.5.5 (?), v2v1 Info source: 150.1.5.5 (?), elected via Auto-RP Uptime: 00:02:33, expires: 00:02:26</pre>	<pre>conf t ip host Router5 150.1.5.5 end R5#show ip pim rp mapping ... Group(s) 224.0.0.0/4 RP 150.1.5.5 (Router5), v2v1 Info source: 150.1.5.5 (Router5), elected via Auto-RP Uptime: 00:07:34, expires: 00:02:21</pre>	<p>Debug ip pim auto-rp output Disallowing group 224.110.110.110:</p>	<pre>debug ip pim auto-rp Auto-RP(0): Received RP-announce, from 150.1.10.10, RP_cnt 1, ht 181 Auto-RP(0): Filtered -224.110.110.110/32 for RP 150.1.10.10 Auto-RP(0): Update (232.0.0.0/5, RP:150.1.10.10), PIMv2 v1</pre>
<p>How do you configure a MultiCast cRP, Candidate RP?</p>	<p>MultiCast Candidate RP: ip pim send-rp-announce <Interface> scope <TTL> group-list <Std-ACL> interval <seconds> Router will start sending traffic destined to: 224.0.1.39 UDP 496 Served to groups specified in the group-list. Denied group-list entries will be served in Dense-Mode. TTL can be used for admin scoping Requires: pim sparse-dense-mode</p>			<p>Auto-RP Listener</p>	<pre>interface x/x ip pim sparse-mode ip pim autorp listener Only 224.0.1.39 and 224.0.1.40 are flooded in dense-mode All other possible denied groups are never flooded in dense-mode</pre>
<p>How to configure a MultiCast MA Mapping Agent?</p>	<p>MultiCast Mapping Agent ip pim send-rp-discovery <Interface> scope <TTL> interval <Seconds> Listens to 224.0.1.39 udp 496 Sends to 224.0.1.40 udp 496 Requires: pim sparse-dense-mode</p>	<p>What are reasons why the RP does not have a S,G entry in its mroute table?</p>	<p>S,G not seen on RP! Source IP not in RPs routing table!</p>	<p>Show ip pim autorp Once with ip pim sparse-dense-mode: Once with ip pim sparse-mode and ip pim autorp listener:</p>	<pre>Transit interfaces: ip pim sparse-dense-mode SW4#show ip pim autorp AutoRP Information: AutoRP is enabled. PIM AutoRP Statistics: Sent/Received RP Announce: 810/334, RP Discovery: 342/417 Transit interfaces: ip pim sparse-mode SW4#sh ip pim autorp AutoRP Information: AutoRP is enabled. AutoRP groups over sparse mode interface is enabled PIM AutoRP Statistics: Sent/Received RP Announce: 816/340, RP Discovery: 345/421</pre>
<p>Auto RP Explained as picture:</p>		<p>How can one verify a group is running in dense-mode in combination with Auto-RP:</p>	<pre>SW4#show ip pim rp mapping PIM Group-to-RP Mappings This system is an RP (Auto-RP) Group(s) 224.0.0.0/5 RP 150.1.8.8 (?), v2v1 Info source: 150.1.5.5 (?), elected via Auto-RP Uptime: 00:51:39, expires: 00:02:36 Group(s) (-) 224.14.14.14/32 RP 150.1.10.10 (?), v2v1 Info source: 150.1.5.5 (?), elected via Auto-RP Uptime: 00:00:22, expires: 00:02:37 ip pim send-rp-announce Loopback0 scope 10 group-list GROUPS ip access-list standard GROUP deny 224.110.110.110 permit 224.0.0.0 7.255.255.255</pre> <p>Shows that group 224.14.14.14 runs in Dense-Mode due to denied group-list</p>	<p>Auto-RP and RP/MA Placement</p>	
<p>Auto RP and group-lists and more specifics: Diagram:</p>		<p>Explain: ip pim rp-announce-filter rp-list <access-list> group-list <accesslist></p>	<p>List of allowed RPs ip pim rp-announce-filter rp-list <access-list> group-list <accesslist> List of allowed groups, per allowed RP</p>	<p>Filtering Auto-RP Messages</p>	<pre>cRP# ip pim send-rp-discovery Loopback0 scope 2</pre>

Help me create more flashcards:

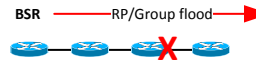
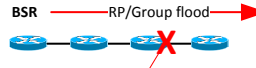
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

Multicast


<p>Multicast Boundary</p> <p>Acl:</p>	<p>filtering control plane traffic (IGMP,PIM,AutoRP) and data plane not relying on TTL.</p> <p>ip multicast boundary <access-list> [filter-autorp]</p> <p>standard ACL: ingress IGMP/PIM, forwards group traffic if it has a match permit 224.10.10.10</p> <p>extended ACL: specifies multicast source and group interface FastEthernet 0/0 ip multicast boundary PERMITTED_GROUPS filter-autorp</p> <p>ip access-list standard PERMITTED_GROUPS deny 232.0.0.0 7.255.255.255 permit any</p>	<p>BSR Multiple RP Candidates</p> <p>Distribute all odd / even multicast groups between two RPs using 31 bits:</p>	<p>BSR: ip pim bsr-candidate Loopback0 31</p> <p>Router1: ip pim rp-candidate Loopback0</p> <p>Router2: ip pim rp-candidate Loopback0</p> <p>Router2#show ip pim rp-hash 239.1.1.1 ... PIMv2 Hash Value (mask 255.255.255.254) RP 150.1.10.10, via bootstrap, priority 0, hash value 989207280 ...</p> <p>Router2#show ip pim rp-hash 239.1.1.2 PIMv2 Hash Value (mask 255.255.255.254) RP 150.1.8.8, via bootstrap, priority 0, hash value 1364246456</p>	<p>IGMP Timers</p>	<p>Interface X ip igmp query-interval <seconds> ip igmp querier-timeout <seconds> ip igmp query-max-response-time <seconds> ip igmp last-member-query-count <2 sec> ip igmp last-member-query-interval <msec> ip igmp immediate-leave group-list <ACL></p>
<p>PIM Bootstrap Router</p> <p>PIMv2 RP</p>	<p>ip pim rp-candidate <PIM-Enabled-Interface> group-list <Standard-ACL> priority <0-255></p> <p>Lower priority is preferred</p>  <p>BSR floods RP/group info via PIM messages hop by hop, NOT Dense-Mode</p>	<p>Filtering BSR Messages</p>  <p>Stop BSR flooding here</p>	 <p>Interface fa0/x ip pim bsr-border</p>	<p>Checking IGMP timers:</p>	<p>show ip igmp interface Fa0/0</p> <p>... IGMP query interval is 20 seconds IGMP querier timeout is 40 seconds IGMP max query response time is 4 seconds Last member query count is 2 Last member query response interval is 1000 ms ...</p>
<p>PIM Bootstrap Router</p> <p>PIMv2 BSR</p>	<p>If both priors are the same, higher IP is used</p> <p>ip pim bsr-candidate <Interface-Name> hash-mask-length priority</p> <p>Hash is used to loadbalance to different RPs</p> <p>Higher priority is preferred</p> <p>ip pim rp-candidate Loopback0 ip pim bsr-candidate Loopback0</p>	<p>Stub Multicast Routing & IGMP Helper</p>	<p>ip igmp helper-address 1.1.1.1</p>  <p>Forwards join/prune messages to 1.1.1.1 without creating (*,G)(S,G) entries locally</p> <p>Keeps track for Stub Router behind low bandwidth link, only sends 2.2.2.2 groups requested by the clients</p> <p>access-list 33 deny 2.2.2.2 access-list 33 permit any</p> <p>Int ser0/0 ip pim sparse-mode ip pim neighbor-filter 33</p>	<p>Multicast Helper Map</p>	 <p>ip forward-protocol udp 5000 ip access-list extended TRAFFIC permit udp any any eq 5000 interface FastEthernet 0/0 ip multicast helper-map broadcast 224.1.2.3 TRAFFIC</p> <p>ip forward-protocol udp 5000 ip access-list extended TRAFFIC permit udp any any eq 5000 interface FastEthernet 0/0 ip directed-broadcast ip broadcast-address 155.1.37.255</p> <p>interface Serial 0/0 ip multicast helper-map 224.1.2.3 155.1.37.255 TRAFFIC</p>
<p>debug ip pim bsr:</p>	<p>debug ip pim bsr</p> <p>PIM-BSR(0): RP-set for 224.0.0.0/4 PIM-BSR(0): RP(1) 150.1.5.5, holdtime 150 sec priority 0 PIM-BSR(0): Bootstrap message for 150.1.5.5 originated PIM-BSR(0): Build v2 Candidate-RP advertisement for 150.1.5.5 priority 0, holdtime 150 PIM-BSR(0): Candidate RP's group prefix 224.0.0.0/4 PIM-BSR(0): Send Candidate RP Advertisement to 150.1.5.5</p> <p>When RPF check fails: PIM-BSR(0): bootstrap from non-RPF neighbor 155.1.146.6</p>	<p>IGMP Filtering</p>	<p>Receivers wanting to IGMP join/report</p> <p>ip igmp access-group <ACL></p> <p>Standard ACL: permit 239.1.1.0 0.0.0.255 allow all groups within 239.1.1.0/24 to be joined.</p> <p>Extended ACL: permit ip <srcip> <src-mask> <group-ip> <group-mask> Allows to specify source and group</p>	<p>Multicast Rate Limiting</p>	<p>ip multicast rate-limit {in out} group-list <acl> source-list <acl> <limit></p> <p>limit is specified in Kilobits per second</p> <p>ip multicast rate-limit out group-list 100 128 ip multicast rate-limit out 512 Aggregate max flow Over all is 512</p> <p>If limit is "forgotten" it results in ALL multicast traffic being dropped!</p>
<p>show ip pim bsr-router</p>	<p>Rack1R1#show ip pim bsr-router</p> <p>PIMv2 Bootstrap information BSR address: 150.1.5.5 (?) Uptime: 00:21:04, BSR Priority: 0, Hash mask length: 32 Expires: 00:01:35</p>	<p>ip igmp limit <N></p> <p>Explained:</p>	<p>Used globally:</p> <p>Allows only the configured count of multicast groups over all multicast enabled interfaces.</p> <p>Per Interface:</p> <p>Limits the number of multicast groups joined on this interface.</p>	<p>Multicast Rate Limiting</p> <p>Show commands:</p>	<p>Show ip mroute (155.1.146.6, 239.1.1.100), 00:00:04/00:02:57, flags: LJT Incoming interface: Serial1/0/1, RPF nbr 155.1.0.5 Outgoing interface list: FastEthernet0/0, Forward/Dense, 00:00:04/00:00:00, limit 256 kbps</p>
<p>How to check that the RP/Group mapping has been learned by BSR?</p>	<p>R1#show ip pim rp mapping</p> <p>PIM Group-to-RP Mappings Group(s) 224.0.0.0/4 RP 150.1.5.5 (?), v2 Info source: 150.1.5.5 (?), via bootstrap, priority 0, holdtime 150 Uptime: 00:05:23, expires: 00:02:06</p>	<p>Check for:</p> <p>ip igmp limit</p> <p>And</p> <p>ip igmp access-group</p> <p>With a show command:</p>	<p>Interface fa0/0 Ip igmp limit 10 Ip igmp filter IGMP_FILTER</p> <p>show ip igmp interface fastEthernet 0/0 ... Inbound IGMP access group is IGMP_FILTER ... Interface IGMP State Limit : 1 active out of 10 max ...</p>	<p>Bidirectional PIM</p> <p>Shared tree (*,G):</p>	<p>Enable PIM Bidir in all routers: ip pim bidir-enable</p> <p>Statically assign the RP: ip pim rp-address <IP> <ACL> bidir</p> <p>Using Auto-RP: ip pim send-rp-announce <interface> scope <TTL> group-list <ACL> bidir</p> <p>Using BSR: ip pim rp-candidate <interface> group-list <ACL> bidir</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

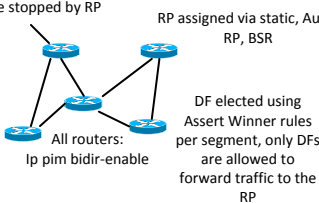
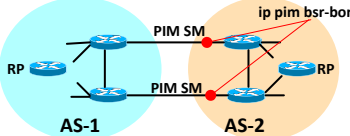


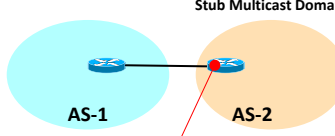
[Donate](#)



Thanks for appreciating my efforts

Colin

Multicast


<h3>Bidirectional PIM</h3>	<p>Source always sends to RP, can't be stopped by RP</p>  <p>RP assigned via static, Auto-RP, BSR</p> <p>DF elected using Assert Winner rules per segment, only DFs are allowed to forward traffic to the RP</p> <p>All routers: ip pim bidir-enable</p>	<p>Establish a DVMRP connection:</p>	<pre>ip dvmrp interoperability access-list 40 permit 155.1.0.0 0.0.255.255 interface FastEthernet0/1 ip dvmrp metric 1 list 40 eigrp 100 interface tunnel 0 ip unnumbered Loopback0 ip pim dense-mode tunnel source Loopback0 tunnel destination 204.12.1.100 tunnel mode dvmrp</pre>	<p>Show ip msdp peer</p> <p>Output:</p>	<pre>SW1#show ip msdp peer MSDP Peer 150.1.5.5 (?), AS 200 (configured AS) Connection status: State: Listen, Resets: 0, Connection source: Loopback0 (150.1.7.7) Uptime(Downtime): 00:00:12, Messages sent/received: 0/0 Output messages discarded: 0 Connection and counters cleared 00:00:12 ago SA Filtering: Input (S,G) filter: none, route-map: none Input RP filter: none, route-map: none Output (S,G) filter: none, route-map: none Output RP filter: none, route-map: none SA-Requests: Input filter: none Peer ttl threshold: 0 SAs learned from this peer: 0 Input queue size: 0, Output queue size: 0</pre>
<h3>Bidirectional PIM</h3> <p>Identify DFs and the RP with show ip mroute</p>	<p>DF routers will show: (* 238.1.1.1), 00:08:15/00:02:59, RP 150.1.5.5, flags: BCL Bidir-Upstream: Serial0/0.1, RPF nbr 155.1.0.5 Outgoing interface list: FastEthernet0/0, Forward/Sparse, 00:08:15/00:02:40 Serial0/0.1, Bidir-Upstream/Sparse, 00:08:15/00:00:00</p> <p>The RP will NOT show Bidir-Upstream !</p>	<p>Verify DVMRP packet generation:</p>	<pre>R2#debug ip dvmrp detail DVMRP(0): Building Report for FastEthernet0/1 DVMRP(0): Report 155.1.146.0/24, metric 32 DVMRP(0): Report 155.1.10.0/24, metric 1 DVMRP(0): Report 155.1.8.0/24, metric 1 DVMRP(0): Report 150.1.5.0/24, metric 1 DVMRP(0): Report 150.1.10.0/24, metric 1 DVMRP(0): Report 150.1.8.0/24, metric 1 DVMRP(0): Delay Report on FastEthernet0/1 DVMRP(0): 12 unicast, 0 MBGP, 0 DVMRP routes advertised DVMRP(0): Send Report on FastEthernet0/1 to 224.0.0.4</pre>	<p>show ip msdp summary</p> <p>Output:</p>	<pre>SW1#show ip msdp summary MSDP Peer Status Summary Peer Address AS State Uptime/ Reset SA Peer Name Downtime Count Count 150.1.5.5 200 Up 00:00:57 0 0 ? 150.1.8.8 200 Up 00:01:35 0 0 ?</pre>
<h3>Source Specific Multicast</h3>	<p>Specify groups that should use SSM:</p> <pre>ip pim ssm range default range <Standard-ACL></pre> <p>Default = 232.0.0.0/8</p> <p>Interface X ip igmp version 3</p> <p>Interface Y ip igmp version 3 ip igmp join 232.2.2.2 source 2.2.2.2</p>	<h3>Multicast BGP Extension</h3>	<pre>router bgp 2 address-family ipv4 multicast neighbor 1.1.1.1 activate</pre>  <pre>router bgp 1 address-family ipv4 multicast neighbor 2.2.2.2 activate</pre>	<p>debug ip msdp detail</p> <p>While pinging 239.1.1.1:</p>	<pre>start_index = 0, mroute_cache_index = 0, Qlen = 0 Sent entire mroute table, mroute_cache_index = 0, Qlen = 0 start_index = 0, sa_cache_index = 0, Qlen = 0 Sent entire sa-cache, sa_cache_index = 0, Qlen = 0 Received 120-byte TCP segment from 150.1.5.5 Append 120 bytes to 0-byte msg 26 from 150.1.5.5, qs 1 WAVL Insert SA Source 155.1.10.10 Group 239.1.1.1 RP 150.1.5.5 Successful Forward decapsulated SA data for (155.1.10.10, 239.1.1.1) on Vlan79 Received 120-byte TCP segment from 150.1.5.5 Append 120 bytes to 0-byte msg 27 from 150.1.5.5, qs 1 WAVL Insert SA Source 155.1.108.10 Group 239.1.1.1 RP 150.1.5.5 Successful Forward decapsulated SA data for (155.1.108.10, 239.1.1.1) on Vlan79</pre>
<h3>Source Specific Multicast</h3> <p>Show ip igmp groups x.x.x.x detail</p>	<pre>show ip igmp groups 232.6.6.6 detail Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group, SS - Static Source, VS - Virtual Source, Ac - Group accounted towards access control limit Interface: FastEthernet0/0 Group: 232.6.6.6 Flags: SSM Uptime: 00:26:06 Group mode: INCLUDE Last reporter: 155.1.146.6 Group source list: (C - Cisco Src Report, U - URD, R - Remote, S - Static, V - Virtual, M - SSM Mapping, L - Local, Ac - Channel accounted towards access control limit) Source Address Uptime v3 Exp CSR Exp Fwd Flags 150.1.10.10 00:26:06 00:02:54 stopped Yes R</pre>	<p>Show ip bgp ipv4 multicast summary</p> <p>Output:</p>	<pre>R6#show ip bgp ipv4 multicast summary Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd 150.1.7.7 4 100 57 43 21 0 00:00:29 16 155.1.146.4 4 200 29 49 21 0 00:01:18 2</pre>	<p>show ip pim rp-hash 239.1.1.1</p> <p>Output:</p>	<pre>R4#show ip pim rp-hash 239.1.1.1 RP 150.1.5.5 (?), v2 Info source: 150.1.10.10 (?), via bootstrap, priority 0, holdtime 150 Uptime: 02:35:01, expires: 00:02:24 PIMv2 Hash Value (mask 0.0.0.0) RP 150.1.5.5, via bootstrap, priority 0, hash value 623125189 RP 150.1.8.8, via bootstrap, priority 0, hash value 613026582 Lowest hash is selected, if both are the same, highest RP IP wins.</pre>
<h3>Source Specific Multicast</h3> <p>Show ip igmp groups x.x.x.x detail</p>	<pre>Rack1R1#show ip mroute 232.6.6.6 150.1.10.10 IP Multicast Routing Table (150.1.10.10, 232.6.6.6), 00:26:09/00:02:26, flags: sTI Incoming interface: Serial0/0.1, RPF nbr 155.1.0.5 Outgoing interface list: FastEthernet0/0, Forward/Sparse, 00:26:09/00:02:26</pre>	<h3>MultiCast MSDP</h3> <p>ip msdp peer</p> <p>ip msdp sa-limit x.x.x.x <nr></p>	<pre>ip msdp peer ip msdp originator-id ip msdp peer 2.2.2.2 connect-source Loop0 remote-as 2 ip msdp [vrf X] sa-limit (peer-address) sa-limit %MSDP-4-SA_LIMIT: RP <RP address> for <mroute> exceeded sa-limit of</pre>	<p>mtrace source group:</p> <p>Once successful</p> <p>Once failing, due to the group not joined</p>  <p>Show mtrace Ping 224.44.44.44</p>	<pre>Successful: SW3#mtrace 150.1.10.10 224.44.44.44 ... -2 155.1.79.7 PIM/MBGP [150.1.10.0/24] -3 155.1.37.3 PIM/MBGP [150.1.10.0/24] -4 155.1.0.5 [AS 200] PIM Reached RP/Core [150.1.10.0/24] Failed: SW3#mtrace 150.1.10.10 239.1.1.4 ... -6 155.1.108.10 [AS 200] PIM [150.1.10.0/24] -7 150.1.10.10 PIM Reached RP/Core is missing!</pre>
<h3>DVMRP Interoperability</h3>	<p>Globally enable: ip dvmrp interoperability <i>PIM needs to be enabled!</i></p> <p>Per interface: ip dvmrp unicast-routing</p> <p>DVMRP summarizes groups by default, to disable summary: no ip dvmrp auto-summary</p> <p>By default will only advertise directly connecteds</p> <p>Redistribute static subnets: Interface X ip dvmrp metric <hops> list <access-list> protocol <process-id></p>	<h3>MultiCast MSDP</h3> <p>On a stub multicast domain:</p> 	 <pre>ip msdp default-peer ip msdp default-peer 10.1.1.1 prefix-list site-a</pre>	<p>mtrace 150.1.10.10 239.1.1.1</p> <p>Output:</p>	<pre>SW3#mtrace 150.1.10.10 239.1.1.1 Type escape sequence to abort. Mtrace from 150.1.10.10 to 155.1.79.9 via group 239.1.1.1 From source (?) to destination (?) Querying full reverse path... 0 155.1.79.9 -1 155.1.79.9 PIM/MBGP [150.1.10.0/24] -2 155.1.79.7 PIM/MBGP Reached RP/Core [150.1.10.0/24] -3 155.1.37.3 PIM/MBGP [150.1.10.0/24] -4 155.1.0.5 [AS 200] PIM Reached RP/Core [150.1.10.0/24] -5 155.1.58.8 [AS 200] PIM [150.1.10.0/24] -6 155.1.108.10 [AS 200] PIM [150.1.10.0/24]</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

Multicast

<p>Anycast RP intra-domain solution</p>	<p>- PIM Joins are being sent to the closest RP - Groups of RPs use the same IP address. - To maintain consistent source information configure MSDP sessions.</p> <p>1. Use the same IP address on all routers as the candidate RP IP address. (Propagate via Auto-RP or BSR)</p> <p>OR</p> <p>2. Using different IP addresses on every router. source MSDP sessions and link all candidate RPs in a mesh. Manually specify the MSDP originator ID to be different on every RP</p>	<p>show ip igmp snooping vlan 146</p> <p>Output:</p>	<p>SW1#show ip igmp snooping vlan 146 Global IGMP Snooping configuration:</p> <pre>IGMP snooping : Enabled IGMPv3 snooping (minimal) : Enabled Report suppression : Enabled TCN solicit query : Disabled TCN flood query count : 2 Last Member Query Interval : 1000 Vlan 146: ----- IGMP snooping : Enabled IGMPv2 immediate leave : Enabled Explicit host tracking : Enabled Multicast router learning mode : pim-dvmrp Last Member Query Interval : 1000 CGMP interoperability mode : IGMP_ONLY</pre>	<p>CGMP messages:</p> <table border="1"> <thead> <tr> <th>Type</th> <th>GDA</th> <th>USA</th> <th>Function</th> </tr> </thead> <tbody> <tr><td>Join</td><td></td><td></td><td></td></tr> <tr><td>Join</td><td></td><td></td><td></td></tr> <tr><td>Join</td><td></td><td></td><td></td></tr> <tr><td>Leave</td><td></td><td></td><td></td></tr> <tr><td>Leave</td><td></td><td></td><td></td></tr> <tr><td>Leave</td><td></td><td></td><td></td></tr> <tr><td>Leave</td><td></td><td></td><td></td></tr> <tr><td>Leave</td><td></td><td></td><td></td></tr> </tbody> </table>	Type	GDA	USA	Function	Join				Join				Join				Leave				Leave				Leave				Leave				Leave				<table border="1"> <thead> <tr> <th>Type</th> <th>GDA</th> <th>USA</th> <th>Function</th> </tr> </thead> <tbody> <tr><td>Join</td><td>0</td><td>Router MAC</td><td>Identify Router Port</td></tr> <tr><td>Join</td><td>Group MAC</td><td>Member MAC</td><td>Adds member</td></tr> <tr><td>Leave</td><td>Group MAC</td><td>Member MAC</td><td>Removes member</td></tr> <tr><td>Leave</td><td>Group MAC</td><td>0</td><td>Removes Group</td></tr> <tr><td>Leave</td><td>0</td><td>Router MAC</td><td>Removes all groups Affected switch</td></tr> <tr><td>Leave</td><td>0</td><td>0</td><td>Removes all groups all switches</td></tr> </tbody> </table>	Type	GDA	USA	Function	Join	0	Router MAC	Identify Router Port	Join	Group MAC	Member MAC	Adds member	Leave	Group MAC	Member MAC	Removes member	Leave	Group MAC	0	Removes Group	Leave	0	Router MAC	Removes all groups Affected switch	Leave	0	0	Removes all groups all switches
Type	GDA	USA	Function																																																																		
Join																																																																					
Join																																																																					
Join																																																																					
Leave																																																																					
Leave																																																																					
Leave																																																																					
Leave																																																																					
Leave																																																																					
Type	GDA	USA	Function																																																																		
Join	0	Router MAC	Identify Router Port																																																																		
Join	Group MAC	Member MAC	Adds member																																																																		
Leave	Group MAC	Member MAC	Removes member																																																																		
Leave	Group MAC	0	Removes Group																																																																		
Leave	0	Router MAC	Removes all groups Affected switch																																																																		
Leave	0	0	Removes all groups all switches																																																																		
<p>Anycast RP intra-domain solution</p> <p>Config:</p>	<p>R1# interface loopback 0 ip address 9.9.9.9 255.255.255.255 interface loopback 1 ip address 1.1.1.1 255.255.255.255 ip msdp peer 2.2.2.2 connect-source loopback 1 ip msdp originator-id loopback1 ip pim rp-address 9.9.9.9</p> <p>R2# interface loopback 0 ip address 9.9.9.9 255.255.255.255 interface loopback 1 ip address 2.2.2.2 255.255.255.255 ip msdp peer 1.1.1.1 connect-source loopback 1 ip msdp originator-id loopback1 ip pim rp-address 9.9.9.9</p>	<p>show ip igmp snooping mrouter vlan <nr></p> <p>show ip igmp snooping groups vlan <nr></p> <p>Output:</p>	<p>Rack1SW1#show ip igmp snooping mrouter vlan 146 Vlan ports ----- 146 Fa0/1(dynamic), Fa0/19(dynamic)</p> <p>SW4#show ip igmp snooping groups vlan 146 Vlan Group Version Port List ----- 146 239.1.1.100 v2 Fa0/4, Fa0/13, Fa0/16</p>	<table border="1"> <thead> <tr> <th>Protocol</th> <th>Implicit Join</th> <th>Explicit Join</th> </tr> </thead> <tbody> <tr><td>DVMRP</td><td></td><td></td></tr> <tr><td>MOSPF</td><td>X</td><td>X</td></tr> <tr><td>PIM-DM</td><td></td><td></td></tr> <tr><td>PIM-SM</td><td>X</td><td>X</td></tr> <tr><td>CBT</td><td></td><td>X</td></tr> </tbody> </table>	Protocol	Implicit Join	Explicit Join	DVMRP			MOSPF	X	X	PIM-DM			PIM-SM	X	X	CBT		X	<table border="1"> <thead> <tr> <th>Protocol</th> <th>Implicit Join</th> <th>Explicit Join</th> </tr> </thead> <tbody> <tr><td>DVMRP</td><td></td><td></td></tr> <tr><td>MOSPF</td><td>X</td><td>X</td></tr> <tr><td>PIM-DM</td><td></td><td></td></tr> <tr><td>PIM-SM</td><td>X</td><td>X</td></tr> <tr><td>CBT</td><td></td><td>X</td></tr> </tbody> </table> <p>Sender originates Hosts join via IGMP</p>	Protocol	Implicit Join	Explicit Join	DVMRP			MOSPF	X	X	PIM-DM			PIM-SM	X	X	CBT		X																												
Protocol	Implicit Join	Explicit Join																																																																			
DVMRP																																																																					
MOSPF	X	X																																																																			
PIM-DM																																																																					
PIM-SM	X	X																																																																			
CBT		X																																																																			
Protocol	Implicit Join	Explicit Join																																																																			
DVMRP																																																																					
MOSPF	X	X																																																																			
PIM-DM																																																																					
PIM-SM	X	X																																																																			
CBT		X																																																																			
<p>Anycast RP intra-domain solution</p> <p>Diagram:</p>		<p>Catalyst Multicast VLAN Registration (MVR)</p>	<p>Int Gi0/2 mvr type receiver Switchport access vlan 20</p> <p>Int Gi0/1 mvr type source Switchport access vlan 10</p> <p>mvr vlan 10</p> <p>mvr mvr group 224.9.9.9</p> <p>Int vlan [10,20] Ip address x.x.x.x ip pim sparse-dense-mode</p> <p>no ip multicast-routing distributed !!</p>	<table border="1"> <thead> <tr> <th>Protocol</th> <th>Source-Based-Tree (S,G) RPT</th> <th>Shared-Tree (*,G) SPT</th> </tr> </thead> <tbody> <tr><td>DVMRP</td><td></td><td></td></tr> <tr><td>MOSPF</td><td>X</td><td></td></tr> <tr><td>PIM-DM</td><td>X</td><td></td></tr> <tr><td>PIM-SM</td><td></td><td>X</td></tr> <tr><td>CBT</td><td></td><td>X</td></tr> </tbody> </table>	Protocol	Source-Based-Tree (S,G) RPT	Shared-Tree (*,G) SPT	DVMRP			MOSPF	X		PIM-DM	X		PIM-SM		X	CBT		X	<table border="1"> <thead> <tr> <th>Protocol</th> <th>Source-Based-Tree (S,G) RPT</th> <th>Shared-Tree (*,G) SPT</th> </tr> </thead> <tbody> <tr><td>DVMRP</td><td>X</td><td></td></tr> <tr><td>MOSPF</td><td>X</td><td></td></tr> <tr><td>PIM-DM</td><td>X</td><td></td></tr> <tr><td>PIM-SM</td><td></td><td>X</td></tr> <tr><td>CBT</td><td></td><td>X</td></tr> </tbody> </table>	Protocol	Source-Based-Tree (S,G) RPT	Shared-Tree (*,G) SPT	DVMRP	X		MOSPF	X		PIM-DM	X		PIM-SM		X	CBT		X																												
Protocol	Source-Based-Tree (S,G) RPT	Shared-Tree (*,G) SPT																																																																			
DVMRP																																																																					
MOSPF	X																																																																				
PIM-DM	X																																																																				
PIM-SM		X																																																																			
CBT		X																																																																			
Protocol	Source-Based-Tree (S,G) RPT	Shared-Tree (*,G) SPT																																																																			
DVMRP	X																																																																				
MOSPF	X																																																																				
PIM-DM	X																																																																				
PIM-SM		X																																																																			
CBT		X																																																																			
<p>show ip msdp sa-cache</p> <p>Output:</p>	<p>R5#show ip msdp sa-cache MSDP Source-Active Cache - 3 entries (150.1.10.10, 239.1.1.1), RP 150.1.8.8, MBGP/AS 200, 00:00:26:00:05:34, Peer 150.1.8.8</p> <p>(155.1.10.10, 239.1.1.1), RP 150.1.8.8, MBGP/AS 200, 00:00:26:00:05:34, Peer 150.1.8.8</p> <p>(155.1.108.10, 239.1.1.1), RP 150.1.8.8, MBGP/AS 200, 00:00:25:00:05:34, Peer 150.1.8.8</p>	<p>Catalyst IGMP Profiles</p>	<p>Switches allow filtering of IGMP messages sent by directly connected hosts to multicast routers similar to the ip igmp accessgroup</p> <p>applies ingress to layer 2 ports only</p> <p>ip igmp profile 1 permit range 232.0.0.0 232.255.255.255</p> <p>interface FastEthernet 0/4 ip igmp filter 1</p>	<p>PIMv2 DM messages:</p>	<p>0 Hello</p> <p>3 Join/Prune</p> <p>6 Graft</p> <p>7 Graft-Ack</p> <p>5 Assert</p>																																																																
<p>Catalyst IGMP Snooping</p>	<p>IGMP snooping is enabled by default on Catalyst multi-layer switches</p> <p>to disable IGMP snooping globally: no ip igmp snooping</p> <p>or disable per vlan no ip igmp snooping vlan <VLAN-ID></p> <p>Statically configure a port to a router: ip igmp snooping vlan <vlan-id> mrouter interface <interface-id></p> <p>switchports with only one host attached, can immediately leave the group if a leave is heard: ip igmp snooping vlan <vlan-id> immediate-leave</p>	<p>MSDP default-peer</p>	<p>STUB# ip msdp peer 1.2.3.4 connect-source Lo0 ip msdp peer 9.8.7.6 connect-source Lo0 ip msdp default-peer 1.2.3.4 ip msdp default-peer 9.8.7.6</p> <p>Active default peer is the first in the config. SAs are not accepted from 9.8.7.6 unless 1.2.3.4 fails. (accepts only a single peer)</p>	<p>PIMv2 SM messages:</p>	<p>0 Hello</p> <p>4 Bootstrap</p> <p>8 Candidate-RP-Advertisement</p> <p>3 Join/Prune</p> <p>5 Assert</p> <p>1 Register</p> <p>2 Register-Stop</p>																																																																
<p>Ethernet MAC address range for multicast</p> <p>IP Address range for MultiCast</p>	<p>01:00:5E:00:00:00</p> <p>To</p> <p>01:00:5E:7F:FF:FF</p> <p>224.0.0.0</p> <p>To</p> <p>239.255.255.255</p>	<p>Limiting multicasts IPv4 boundaries:</p>	<p>ip multicast threshold <TTL-value></p> <p>-----</p> <p>ip multicast boundary 10</p> <p>accesslist 10 deny 224.x.x.x 0.0.0.0</p>	<p>What dangers are there using the following command:</p> <p>RTR(config)#ip pim send-rp-discovery Loopback0 scope 5</p>	<p>RTR(config)#ip pim send-rp-discovery Loopback0 scope 5 No PIM interface ignored in accepted command. RTR(config)#</p> <p>Results in:</p> <p>Ip pim send-rp-discovery scope 5</p> <p>MAPPING AGENT, WILL NOT WORK, Loopback 0 has to be enabled with PIM first!!!!</p>																																																																

Help me create more flashcards:

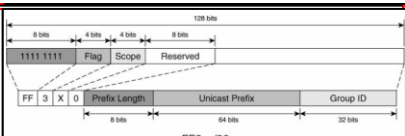
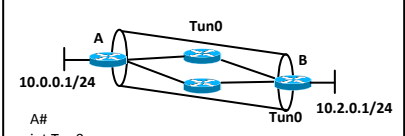
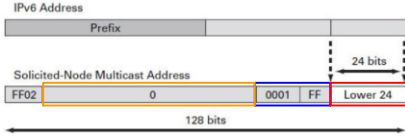
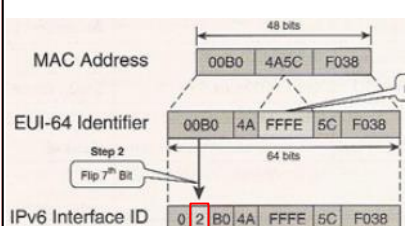
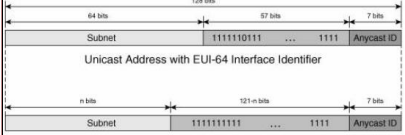
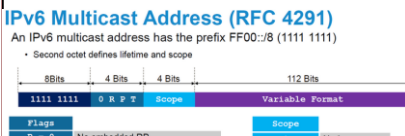
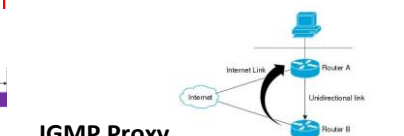
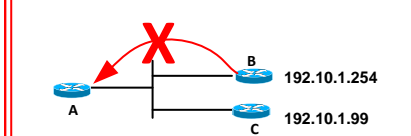
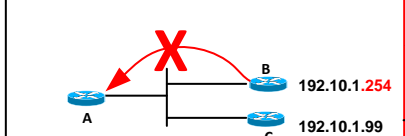
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

Multicast


<p>Explain two ways on how to perform RP load balancing for multicast groups?</p>	<p>1. use group lists: 2. use hash values:</p>	<p>Unicast-Prefix based Multicast Address</p> <p>Similar to GLOB under IPv4:</p>	 <p>IPv6 Unicast Address: 2001:100:abc:1::/64</p> <p>Global Multicast Address: FF3E:0040:2001:100:abc:1::11FE:11EE</p>	<p>What types of Multicast RP assignments are there?</p>	<ul style="list-style-type: none"> • Static RP • Bootstrap Router (BSR) • Auto-RP • Anycast-RP • Phantom RP • Embedded RP 																								
<p>Explain how to Load Balance MultiCast traffic over an Equal cost path?</p>	<p>MultiCast Load balance over equal cost Path</p>  <pre> A# int Tun0 ip addr 1.1.1.1 255.255.255.0 ip pim sparse-dense-mode tunnel source lo0 tunnel destination x.x.x.x ip mroute 2.2.2.2 255.255.255.255 tunnel0 </pre>	<p>IPv6 Solicited-Node Multicast Address:</p>	 <p>2001:100:abc:1:0:0:aabb:ccdd/64</p> <p>FF02::1:FFbb:ccdd</p>	<p>Multicast Vlan Registration</p> <p>MVR</p> <p>Show commands:</p>	<p>show commands:</p> <p>show mvr</p> <p>show mvr interface</p> <p>show mvr members</p> <p>show ip igmp groups</p> <p><i>mvr mode dynamic Only forwards traffic if receivers are attached on receiver side</i></p>																								
<p>Explain mroute's output:</p> <pre> Router# mroute 192.1.7.37 (b.x.com) [version cisco 11.1] [flags: PMSA] 192.1.7.37 -> 192.1.7.34 (s.x.com) [1/0/pim] 192.1.7.37 -> 192.1.7.47 (d.x.com) [1/0/pim/querier/leaf] 192.1.7.37 -> 192.1.7.44 (d2.x.com) [1/0/pim] 131.9.26.10 -> 131.9.26.9 (su.bbnplanet.net) [1/32/pim] </pre>	<p>Router# mroute</p> <pre> 192.1.7.37 (b.x.com) [version cisco 11.1] [flags: PMSA] 192.1.7.37 -> 192.1.7.34 (s.x.com) [1/0/pim] 192.1.7.37 -> 192.1.7.47 (d.x.com) [1/0/pim/querier/leaf] 192.1.7.37 -> 192.1.7.44 (d2.x.com) [1/0/pim] 131.9.26.10 -> 131.9.26.9 (su.bbnplanet.net) [1/32/pim] </pre> <p>Metric/no TTL threshold set(0)/Protocol/Role/Type</p> <p>P = prune-capable M = mtrace-capable S = SNMP-capable A = Auto-RP-capable</p>	<p>IPv6</p> <p>EUI-64 Interface ID generation:</p>	 <p>IPv6 Interface ID: 02 B0 4A FFFE 5C F038</p>	<p>Show ip mroute active</p> <p>Output:</p>	<pre> Router# show ip mroute active Active IP Multicast Sources - sending >= 4 kbps Group: 239.1.2.3, (?) Source: 10.0.0.1 (?) Rate: 20 pps/4 kbps[1sec], 4 kbps[last 30 secs], 4 Kbps(life avg) </pre>																								
<p>Explain the following output of mtrace:</p> <pre> R6# mtrace 10.0.12.1 Type escape sequence to abort. Mtrace from 10.0.12.1 to 10.0.56.6 via RPF From source (?) to destination (?) Querying full reverse path... 0 10.0.56.6 -1 10.0.56.6 PIM [10.0.12.0/24] -2 10.0.56.5 PIM [10.0.12.0/24] -3 10.0.45.4 PIM/Static [10.0.12.1/32] static mroute -4 10.0.45.4 None No route Intentionally disabled PIM, -> RPF fails -5 10.0.12.1 </pre>	<pre> R6# mtrace 10.0.12.1 Type escape sequence to abort. Mtrace from 10.0.12.1 to 10.0.56.6 via RPF From source (?) to destination (?) Querying full reverse path... 0 10.0.56.6 -1 10.0.56.6 PIM [10.0.12.0/24] -2 10.0.56.5 PIM [10.0.12.0/24] -3 10.0.45.4 PIM/Static [10.0.12.1/32] static mroute -4 10.0.45.4 None No route Intentionally disabled PIM, -> RPF fails -5 10.0.12.1 </pre>	<p>IPv6 AnyCast Address format:</p>	 <p>AnyCast ID field can take following values: 00 - 7D, 7F</p> <p>7E is reserved for MIPv6</p>	<p>show ip pim interface count</p> <p>output</p>	<pre> Router# show ip pim interface count State: * - Fast Switched, D - Distributed Fast Switched H - Hardware Switching Enabled Address Interface FS Mpackets In/Out 172.31.100.2 GigabitEthernet0/0/0 * 4122/0 10.1.0.1 GigabitEthernet1/0/0 * 0/3193 </pre>																								
<p>Explain the mstat syntax:</p> <pre> mstat 1.1.1.1 2.2.2.2 224.9.9.9 </pre>	<p>source destination group</p> <pre> mstat 1.1.1.1 2.2.2.2 224.9.9.9 </pre>	<p>IPv6 Multicast Address format:</p>	<p>IPv6 Multicast Address (RFC 4291)</p> <p>An IPv6 multicast address has the prefix FF00::8 (1111 1111)</p> <p>Second octet defines lifetime and scope</p>  <table border="1"> <thead> <tr> <th>Flags</th> <th>Scope</th> <th>Variable Format</th> </tr> </thead> <tbody> <tr> <td>R = 0</td> <td>1</td> <td>Node</td> </tr> <tr> <td>R = 1</td> <td>2</td> <td>Link</td> </tr> <tr> <td>P = 0</td> <td>3</td> <td>Subnet</td> </tr> <tr> <td>P = 1</td> <td>4</td> <td>Admin</td> </tr> <tr> <td>T = 0</td> <td>5</td> <td>Site</td> </tr> <tr> <td>T = 1</td> <td>8</td> <td>Organization</td> </tr> <tr> <td></td> <td>E</td> <td>Global</td> </tr> </tbody> </table> <p>FF00::/8</p>	Flags	Scope	Variable Format	R = 0	1	Node	R = 1	2	Link	P = 0	3	Subnet	P = 1	4	Admin	T = 0	5	Site	T = 1	8	Organization		E	Global	<p>IGMP Proxy</p> 	<p>Upstream UDL Device for IGMP UDRL</p> <p>Interface fa0/x</p> <p>ip pim dense-mode</p> <p>ip igmp unidirectional-link</p> <p>Downstream UDL Device for IGMP UDRL</p> <p>Interface fa0/1</p> <p>ip pim dense-mode</p> <p>ip igmp unidirectional-link</p> <p>Interface fa0/2</p> <p>ip pim dense-mode</p> <p>ip igmp mroute-proxy loopback 0</p> <p>Request IGMP reports sent back to Lo0 for all groups in mroute table forwarded to fa0/2</p> <p>interface loopback 0</p> <p>ip pim dense-mode</p> <p>ip igmp helper-address udl fa0/1</p> <p>ip igmp proxy-service</p>
Flags	Scope	Variable Format																											
R = 0	1	Node																											
R = 1	2	Link																											
P = 0	3	Subnet																											
P = 1	4	Admin																											
T = 0	5	Site																											
T = 1	8	Organization																											
	E	Global																											
<p>Do not allow router B to build a PIM adjacency with Router A: (config on RTR A)</p> 	 <pre> interface Fa0/0 ip pim neighbor-filter 75 access-list 75 deny 192.10.1.254 access-list 75 permit any </pre>	<table border="1"> <thead> <tr> <th>Scope Hex</th> <th>Scope Binary</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>0001</td> <td>Interface-local</td> </tr> <tr> <td>2</td> <td>0010</td> <td>Link-local</td> </tr> <tr> <td>3</td> <td>0011</td> <td>Subnet-local</td> </tr> <tr> <td>4</td> <td>0100</td> <td>Admin-local</td> </tr> <tr> <td>5</td> <td>0101</td> <td>Site-local</td> </tr> <tr> <td>8</td> <td>1000</td> <td>Organizational-local</td> </tr> <tr> <td>E</td> <td>1110</td> <td>Global-local</td> </tr> </tbody> </table> <p>Fill in the blanks</p> <p>IPv6 Multicast Scope and values:</p>	Scope Hex	Scope Binary	Description	1	0001	Interface-local	2	0010	Link-local	3	0011	Subnet-local	4	0100	Admin-local	5	0101	Site-local	8	1000	Organizational-local	E	1110	Global-local	<p>Multicast Service Reflection</p> <p>"Multicast NAT"</p>	<p>To be continued....</p> <p>Found under Implementing Multicast Service Reflection</p>	
Scope Hex	Scope Binary	Description																											
1	0001	Interface-local																											
2	0010	Link-local																											
3	0011	Subnet-local																											
4	0100	Admin-local																											
5	0101	Site-local																											
8	1000	Organizational-local																											
E	1110	Global-local																											

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

Multicast

<p>ECMP Multicast Load Splitting based on:</p> <p>S,G</p> <p>S,G next-hop</p>	<p>ip multicast multipath Based on source address</p> <p>ip multicast multipath s-g-hash basic Source and group address S-G-Hash algorithm.</p> <p>ip multicast multipath s-g-hash next-hop-based source, group, and next-hop address using the next-hop-based algorithm</p> <p>Alternative - Tunnel interface (static mroutes)</p>				
<p>Multicast CAC</p> <p>Multicast Limit</p>	<p>ip multicast limit out acl-basic 75 ip multicast limit out acl-premium 25 ip multicast limit out acl-gold 25</p> <p>ip multicast limit out ACL <kbps permitted></p> <p>debug ip mrouting limits [group-address] show ip multicast limit type number clear ip multicast limit</p>				
<p>What info does the following command provide:</p> <p>show ip multicast:</p>	<p>R2#show ip multicast Multicast Routing: disabled Multicast Multipath: disabled Multicast Route limit: No limit Multicast Fallback group mode: Sparse Number of multicast boundaries configured with filter-autorp option: 0</p>				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin


<p>IPv6 interface serial 0/0/0</p> <p>Frame map / show frame-relay map output:</p>	<pre>interface Serial0/0/0 encapsulation frame-relay ipv6 address FE80::5 link-local frame-relay map ipv6 FE80::1 501 broadcast R5#show frame-relay map Serial0/0/0 (up): ipv6 FE80::2 dlci 502(0x1F6,0x7C60), static, broadcast, CISCO, status defined, active</pre>	<p>IPv6 Auto-Configuration</p>	<p>Announce prefix via ND RA, but hosts are not allowed to use it for autoconfig:</p> <pre>ipv6 nd prefix fc00:1:0:58::/64 14400 14400 no-autoconfig Lifetime set to 4 hours</pre> <hr/> <p>Announce prefix via ND RA, hosts are allowed to use it:</p> <pre>ipv6 nd prefix fc00:1:0:85::/64 14400 14400</pre>	<p>ipv6 nd prefix:</p>	<p>manipulates the IPv6 network prefixes included into RA. By default, all prefixes are included.</p>
<p>Ping ipv6 xx.x.x</p> <p>On 12.2-24T:</p>	<pre>R4#ping ipv6 FE80::5 Output interface: serial0/0/0 Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to FE80::5, timeout is 2 seconds: Packet sent with a source address of FE80::4 !!!! Success rate is 100 percent (5/5), round-trip min/avg/max = 56/56/56 ms R4#</pre>	<p>IPv6 ULA address config:</p>	<pre>interface FastEthernet 0/0 ipv6 address fc00:1:0:58::5/64 ipv6 address fc00:1:0:85::5/64</pre>	<p>IPv6 Auto-Configuration:</p> <p>Explain its function:</p>	<p>With auto-configuration, an IPv6 host may automatically learn the IPv6 prefixes assigned to the local segment, as well as determine the default routers on that segment.</p> <p>Client: ipv6 address autoconfig default</p> <p>Router: ipv6 nd ...</p>
<p>IPv6 Unique Local Address</p> <p>ULA format:</p>	<p>FC00 (7 bits) Unique ID (41 bits) Link ID (16 bits) Interface ID (64 bits).</p>	<p>IPv6 ND RA:</p> <p>advertise itself as the default router every 40 seconds Lifetime interval 60 seconds:</p>	<pre>ipv6 address fc00:1:0:85::5/64 ipv6 nd prefix fc00:1:0:85::/64 14400 14400 ipv6 nd ra-interval 40 ipv6 nd ra-lifetime 60</pre>	 <p>show ipv6 int gi0/0 prefix:</p>	<p>Shows R2 announced prefixes via ND RA's:</p> <pre>R2#show ipv6 int gi0/0 prefix IPv6 Prefix Advertisements GigabitEthernet0/0 Codes: A - Address, P - Prefix-Advertisement, O - Pool U - Per-user prefix, D - Default N - Not advertised, C - Calendar default [LA] Valid lifetime 2592000, preferred lifetime 604800 AP FC00:1:0:58::/64 [L] Valid lifetime 14400, preferred lifetime 14400 AP FC00:1:0:85::/64 [LA] Valid lifetime 14400, preferred lifetime 14400</pre>
<p>IPv6 Global Aggregatable Addressing</p>	<p>1/8th of the total IPv6 address space is currently allocated: 2001::/16</p> <p>binary prefix 001 (2000::/3)</p> <p>2000:: - 3FFF::</p> 	<p>R1 learns its IPv6 address automatically and use R2 as its default gateway</p>  <p>Fill in the blanks:</p>	<p>R1 learns its IPv6 address automatically and use R2 as its default gateway</p>  <pre>R1-Client# interface fa0/1 ipv6 address autoconfig default R2# interface FastEthernet 0/0 ipv6 address fc00:1:0:85::5/64 ipv6 nd prefix fc00:1:0:85::/64 14400 14400 ipv6 nd ra-interval 40 ipv6 nd ra-lifetime 60 no ipv6 nd suppress-ra</pre>	<p>RIPng</p> <p>Config:</p>	<p>ipv6 unicast-routing</p> <pre>interface FastEthernet 0/0 ipv6 address fc00:1:0:1::1/64 ipv6 rip RIPNG enable</pre>
<p>IPv6 EUI-64 Addressing</p>	<pre>interface FastEthernet 0/0.146 ipv6 address 2001:1:0:146::/64 eui-64</pre> <p>Results in:</p> <p>2001:1:0:146:213:7FFF:FE7F:62A0</p>  <p>MAC: 0013 7F7F 62A0</p> <p>Use: show ipv6 interface</p>	<p>ipv6 nd ra-interval:</p>	<p>specifies the periodic interval to send RAs.</p>	<p>show ipv6 route rip</p> <p>Output:</p>	<pre>R6#show ipv6 route rip IPv6 Routing Table - 10 entries Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP U - Per-user Static route I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2 R FC00:1:0:1::/64 [120/2] via FE80::213:7FFF:FE7F:62A0, Gi0/0.146 R FC00:1:0:4::/64 [120/2] via FE80::226:BFF:FE57:BA61, Gi0/0.146</pre>
<p>show ipv6 interface of EUI-64 address:</p>	<pre>R1#show ipv6 interface FastEthernet0/0 is up, line protocol is up IPv6 is enabled, link-local address is FE80::213:7FFF:FE7F:62A0 Global unicast address(es): 2001:1:0:146:213:7FFF:FE7F:62A0, subnet is 2001:1:0:146::/64 [EUI] Joined group address(es): FF02::1 FF02::2 FF02::1:FF7F:62A0 MTU is 1500 bytes ICMP error messages limited to one every 100 milliseconds ICMP redirects are enabled ND DAD is enabled, number of DAD attempts: 1 ND reachable time is 30000 milliseconds ND advertised reachable time is 0 milliseconds ND advertised retransmit interval is 0 milliseconds ND router advertisements are sent every 200 seconds ND router advertisements live for 1800 seconds Hosts use stateless autoconfig for addresses.</pre>	<p>ipv6 nd ra-lifetime:</p>	<p>specifies the validity interval of the router's IPv6 address</p>	<p>Output of debug ipv6 rip</p>	<pre>R6#debug ipv6 rip RIP Routing Protocol debugging is on Rack1R6# RIPng: response received from FE80::20D:65FF:FE84:6560 on FastEthernet0/0.146 for RIPNG src=FE80::20D:65FF:FE84:6560 (Fa0/0.146) dst=FF02::9 sport=521, dport=521, length=52 command=2, version=1, mbz=0, #rte=2 tag=0, metric=1, prefix=2001:1:0:146::/64 tag=0, metric=1, prefix=FC00:1:0:1::/64 Sending multicast update on Loopback100 for RIPNG src=FE80::20C:85FF:FEC1:FC60 dst=FF02::9 (Loopback100) sport=521, dport=521, length=92 command=2, version=1, mbz=0, #rte=4 tag=0, metric=1, prefix=2001:1:0:146::/64</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

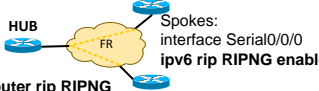
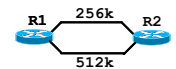

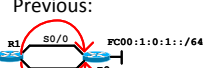
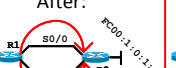
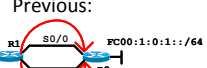
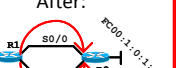

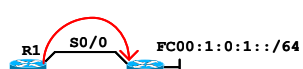


Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts


Colin


<h3>RIPng over NBMA</h3>	 <p>HUB: ipv6 router rip RIPNG no split-horizon</p> <p>interface Serial 0/0/0 ipv6 rip RIPNG enable</p> <p>Do not forget to use the broadcast keyword with the IPv6 mapping statements on DLCIs</p> <p>NBMA networks the next-hop is always going to be a link-local address</p>	<h3>RIPng Default Routing</h3>	<p>interface g0/0.146 ipv6 rip RIPNG default-information originate metric 5</p>	<h3>IPv6 Summarization</h3> <p>FC00:1:0:5::/64 FC00:1:0:8::/64</p>	<p>FC00:1:0:5::/64 FC00:1:0:8::/64</p> <p>0000 0000 0000 0 0101 = 5 0000 0000 0000 0 1000 = 8</p> <p>FC00:1::/60</p> <p>Shifting 4 bits, to the left ending up with a /60 as summary</p>
<h3>Show ipv6 rip database</h3> <p>Output:</p>	<p>R5#show ipv6 rip database RIP process "RIPNG", local RIB 2001:1:0:146::/64, metric 2, installed Serial0/0/0/FE80::1, expires in 179 secs Serial0/0/0/FE80::4, expires in 157 secs 2001:1:0:1234::/64, metric 2 Serial0/0/0/FE80::1, expires in 179 secs Serial0/0/0/FE80::4, expires in 157 secs FC00:1:0:1::/64, metric 2, installed Serial0/0/0/FE80::1, expires in 179 secs</p> <p><i>In routing table</i></p>	<h3>EIGRPv6 Enable process:</h3>	<p>ipv6 router eigrp <AS-Nr> no shutdown</p>	<h3>EIGRPv6 Summarization</h3>	<p>With IPv6 EIGRP it is not possible to configure a leak-map to leak more specifics!</p> <p>interface Serial 0/0/0 ipv6 summary-address eigrp 100 FC00:1::/60</p>
<h3>RIPng Summarization</h3>	<p>Interface fa0/x ipv6 rip RIPNG enable ipv6 rip RIPNG summary-address FC00:1::/61</p>	<h3>EIGRPv6 authentication config:</h3>	<p>key chain EIGRPV6 key 1 key-string CISCO</p> <p>interface FastEthernet 0/1 ipv6 eigrp 100 ipv6 authentication mode eigrp 100 md5 ipv6 authentication key-chain eigrp 100 EIGRPV6</p> <p>ipv6 router eigrp 100 No shutdown</p>	<h3>EIGRPv6 Prefix Filtering</h3> <h3>Distribute-Lists</h3> <p>Blocking prefix to enter routing table</p>	<p>ipv6 prefix-list PFX-1 seq 10 deny FC00:1:0:6::/64 ipv6 prefix-list PFX-1 seq 20 permit ::0 le 128</p> <p>ipv6 router eigrp 100 distribute-list prefix-list PFX-1 in</p>
<h3>RIPng Prefix Filtering</h3>	<p>Denying a prefix in, allowing all others: ipv6 prefix-list PFX deny fc00:1:0:6::/64 ipv6 prefix-list PFX permit ::0 le 128</p> <p>ipv6 router rip RIPNG distribute-list prefix-list PFX in</p>	<h3>Explain:</h3> <p>ipv6 split-horizon eigrp 100</p> <p>no ipv6 next-hop-self eigrp</p>	<p>ipv6 split-horizon eigrp 100</p> <p>disable the split-horizon rule on a particular interface</p> <p>no ipv6 next-hop-self eigrp</p> <p>used on the hub router, explicitly sets the next-hop field in the relayed EIGRPv6 updates to the spoke router's IP address → used in DMVPN setups</p>	<h3>EIGRPv6 Metric Manipulation</h3>	<p>EIGRPv6 can only implement equal cost loadbalancing due to CEF IPv6 limitations.</p>  <p>ipv6 router eigrp 100 metric weight 0 0 0 1 0 0 variance 3</p> <p>interface Serial 0/0 delay 2000</p> <p>interface Serial 0/1 delay 1000</p> <p>Manipulate path metrics to adjust for unequal load balancing settings</p>
<h3>RIPng Metric Manipulation</h3> <p>Offset-list</p>	<p>Interface serial 0/0/0 ipv6 rip RIPNG enable ipv6 rip RIPNG metric-offset 2</p> <p>You only need link local addresses for RIP to work!</p>	<h3>EIGRPv6</h3> <p>Usefull show commands to troubleshoot</p>	<p>show ipv6 eigrp 100 interfaces</p> <p>show ipv6 protocols</p> <p>show ipv6 eigrp neighbors</p> <p>show ipv6 eigrp interfaces detail fastEthernet X</p> <p>(gives clues about authentication type / used key-chain)</p>	<h3>EIGRPv6 checking unequal load balancing</h3> 	<p>ipv6 router eigrp 100 metric weight 0 0 0 1 0 0 variance 3</p> <p>interface Serial 0/0 delay 2000</p> <p>interface Serial 0/1 delay 1000</p> <p>show ipv6 route FC00:1::/60</p> <p>show ipv6 eigrp topology FC00:1::/60 Topology should have 2 entries</p> <p>show ipv6 eigrp topology FC00:1::/60 Check that CEF equally allocates the 16 buckets</p>
<p>IPv6 router rip outputs after applying offset-lists:</p> <p>Previous:  show ipv6 route FC00:1:0:1::/64</p> <p>After:  show ipv6 route FC00:1:0:1::/64</p> <p>?</p>	<p>IPv6 router rip outputs after applying offset-lists:</p> <p>Previous:  show ipv6 route FC00:1:0:1::/64</p> <p>After:  show ipv6 route FC00:1:0:1::/64</p> <p>R FC00:1:0:1::/64 [120/2] via FE80::1, Serial0/0</p> <p>R FC00:1:0:1::/64 [120/2] via FE80::22, Serial0/1</p> <p>show ipv6 route FC00:1:0:1::/64</p> <p>R FC00:1:0:1::/64 [120/2] via FE80::1, Serial0/0</p>	<p>Make sure that RIP only uses one path, while serial0/1 is used as a backup connection, if serial0/0 fails.</p>  	 <p>Interface serial 0/1 ipv6 rip RIPNG enable ipv6 rip RIPNG metric-offset 2</p>	<h3>Show ipv6 cef <prefix> internal</h3> <p>Output:</p> 	<p>Rack1#show ipv6 cef FC00:1::/60 internal</p> <p>FC00:1::/60, epoch 0, RIB[0], refcount 4, per-destination sharing</p> <p>sources: RIB</p> <p>feature space:</p> <p>IPRM: 0x00038000</p> <p>ifname:</p> <p>Serial0/0(0/5): FE80::5 Serial0/1(0/6): FE80::213:1AFF:FE68:B04C path 65E78FC8, path list 65F9332C, share 1/1, type attached nexthop, for IPv6 nexthop FE80::5 Serial0/0(0), adjacency IPv6 adj out of Serial0/0(0), addr FE80::5 65E7AAE0 path 65E78C28, path list 65F9332C, share 1/1, type attached nexthop, for IPv6 nexthop FE80::213:1AFF:FE68:B04C Serial0/1(0), adjacency IPv6 adj out of Serial0/1(0) 65E79E60</p> <p>output chain:</p> <p>loadinfo 665FD68, per-session, 2 choices, flags 0005, 8 locks</p> <p>Flags: Per-session, for-rx IPv6</p> <p>16 hash buckets</p> <p>< 0 > IPv6 adj out of Serial0/0(0), addr FE80::5 65E7AAE0</p> <p>< 1 > IPv6 adj out of Serial0/1(0) 65E79E60</p> <p>Subblocks:</p> <p>None</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

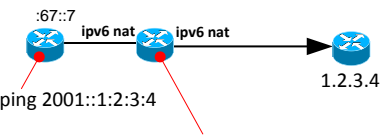
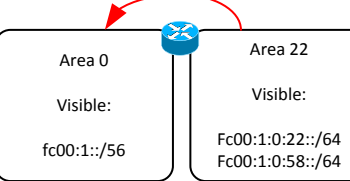

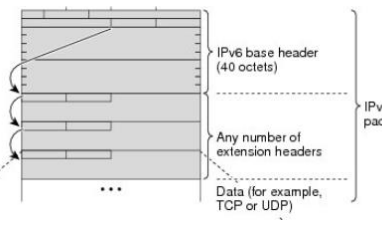
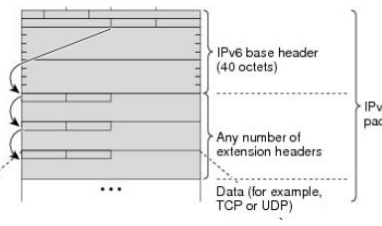
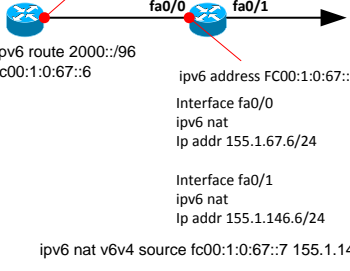
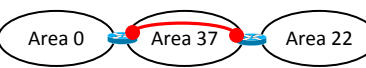
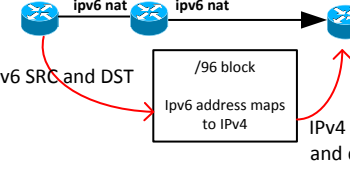
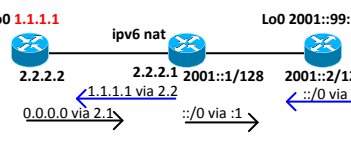
Ranging 5 bucks to unlimited!





Thanks for appreciating my efforts

Colin


<h3>EIGRPv6 Default Routing</h3> <p>using redistribution</p> <p>using summarization</p>	<p>Redistribution: EIGRP an external route with an AD of 170. Allows metric manipulations!</p> <hr/> <p>Summarization: summarize all routes, no leak-map! Metric is calculated, not changeable</p> <pre>interface Fa0/1 ipv6 summary-address eigrp 100 ::/0 5</pre>	<p>Show ipv6 route ospf</p> <p>C / L / S / R / B :</p>	<pre>show ipv6 route codes: C - Connected L - Local S - Static R - RIP B - BGP U - Per-user Static route I1 - ISIS L1 I2 - ISIS L2 IA - ISIS interarea IS - ISIS summary O - OSPF intra OI - OSPF inter, OE1 - OSPF ext 1 OE2 - OSPF ext 2 ON1 - OSPF NSSA ext 1 ON2 - OSPF NSSA ext 2</pre>	<p>debug ipv6 nat detailed</p>	 <p>debug ipv6 nat detailed</p> <p>NOT GOOD: IPv6 NAT: Dropping v6tov4 packet</p> <p>GOOD: ipv6nat_find_entry_v4tov6:</p>																																																																																				
<p>IPv6 autoconfiguration</p> <p>Administrative Distance</p> <p>For default routes:</p>	<p>IPv6 autoconfiguration has a Admin Distance of 1 and could overwrite any routing protocol speaking for an injected default route!</p> <pre>SW2#show ipv6 route IPv6 Routing Table - Default - 12 entries S ::/0 [1/0] via FE80::213:1AFF:FE68:B04C, Vlan58</pre> <p>EIGRPv6 was configured, but autoconfigs default route was used instead!</p>	<h3>OSPFv3 Summarization</h3>	<p>ipv6 router ospf 1 area 22 range fc00:1::/56</p> 	<p>IPv6 address types:</p>	<p>Unspecified Address: 0:0:0:0:0:0:0:0</p> <p>Loopback: 0:0:0:0:0:0:0:1</p> <p>IPv4-compatible-IPv6 addr: 0:0:0:0:0:IPV4</p> <p>IPv4-mapped IPv6 addr: 0:0:0:0:0:FFFF:IPV4</p>																																																																																				
<h3>OSPFv3</h3> <p>Basic config:</p>	<p>ipv6 unicast-routing</p> <p>ipv6 router ospf 1 router-id 150.1.7.7</p> <p>interface Vlan 67 ipv6 ospf hello-interval 1 ipv6 ospf 1 area 0</p>	<h3>IPv6 Redistribution</h3>	<p>redistribute <protocol> will not redistribute the locally connected</p> <pre>ipv6 router ospf 1 redistribute rip RIPNG metric 8 redistribute eigrp 100 metric 8 redistribute connected metric 8 ipv6 router rip RIPNG redistribute eigrp 100 include-connected metric 8 redistribute ospf 1 include-connected metric 8 ipv6 router eigrp 100 redistribute rip RIPNG include-connected metric 1000 0 255 1 1500 redistribute ospf 1 include-connected metric 1000 0 255 1 1500</pre>	<h3>IPv6 Packet header</h3> <p>Image:</p> 	<table border="1"> <tr> <td>Version</td> <td>Traffic Class</td> <td>Flow Label</td> </tr> <tr> <td>Payload Length</td> <td>Next Header</td> <td>Hop Limit</td> </tr> <tr> <td colspan="3">Source Address</td> </tr> <tr> <td colspan="3">Destination Address</td> </tr> <tr> <td colspan="3">Next Header</td> </tr> <tr> <td colspan="3">Extension Header information</td> </tr> <tr> <td colspan="3">Data Portion</td> </tr> </table>	Version	Traffic Class	Flow Label	Payload Length	Next Header	Hop Limit	Source Address			Destination Address			Next Header			Extension Header information			Data Portion																																																																	
Version	Traffic Class	Flow Label																																																																																							
Payload Length	Next Header	Hop Limit																																																																																							
Source Address																																																																																									
Destination Address																																																																																									
Next Header																																																																																									
Extension Header information																																																																																									
Data Portion																																																																																									
<p>show ipv6 ospf interface fa0/3</p> <p>Output:</p>	<pre>SW1#show ipv6 ospf interface fastEthernet 0/3 FastEthernet0/3 is up, line protocol is up (connected) Link Local Address FE80::212:1FF:FE31:41, Interface ID 1021 Area 37, Process ID 1, Instance ID 0, Router ID 150.1.7.7 Network Type POINT_TO_POINT, Cost: 1 Transmit Delay is 1 sec, State POINT_TO_POINT, Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5 Hello due in 00:00:01 Index 1/1/2, flood queue length 0 Next 0x0(0)/0x0(0)/0x0(0) Last flood scan length is 1, maximum is 1 Last flood scan time is 0 msec, maximum is 0 msec Neighbor Count is 1, Adjacent neighbor count is 1 Adjacent with neighbor 150.1.3.3 Suppress hello for 0 neighbor(s)</pre>	<h3>IPv6 Filtering</h3>	<p>ipv6 access-list FILTER_OUT permit tcp fc00:1:0:67::/64 any eq 80 permit tcp fc00:1:0:67::/64 any range 20 21 permit udp fc00:1:0:67::/64 any eq 43</p> <p>interface Serial 1/0 ipv6 traffic-filter FILTER_OUT out</p>	<h3>IPv6 packet extension header</h3> <p>Image:</p> 	<p>IPv6 packet extension header</p> 																																																																																				
<h3>OSPFv3 over NBMA</h3> <p>No DR election</p> <p>Not using broadcasts</p>	<p>ipv6 unicast-routing</p> <p>ipv6 router ospf 1 router-id 150.1.2.2</p> <p>interface Serial 0/0 ipv6 ospf 1 area 0 ipv6 ospf network point-to-multipoint non-broadcast ipv6 ospf neighbor fe80::5</p> <p>OSPFv3 neighbours configured under the interface, not the process!</p>	<h3>IPv6 NAT-PT</h3>	<p>ipv6 address FC00:1:0:67::7/64</p>  <p>ipv6 route 2000::96 fc00:1:0:67::6</p> <p>Interface fa0/0 ipv6 nat ip addr 155.1.67.6/24</p> <p>Interface fa0/1 ipv6 nat ip addr 155.1.146.6/24</p> <p>ipv6 nat v6v4 source fc00:1:0:67::7 155.1.146.7</p>	<table border="1"> <thead> <tr> <th>Order</th> <th>Header Type</th> <th>Next Header Code</th> </tr> </thead> <tbody> <tr><td>1</td><td>-</td><td>-</td></tr> <tr><td>2</td><td>0</td><td>0</td></tr> <tr><td>3</td><td>60</td><td>60</td></tr> <tr><td>4</td><td>43</td><td>43</td></tr> <tr><td>5</td><td>44</td><td>44</td></tr> <tr><td>6</td><td>51</td><td>51</td></tr> <tr><td>7</td><td>50</td><td>50</td></tr> <tr><td>8</td><td>60</td><td>60</td></tr> <tr><td>9</td><td>135</td><td>135</td></tr> <tr><td></td><td>59</td><td>59</td></tr> <tr><td>Upr L</td><td>6</td><td>6</td></tr> <tr><td>Upr L</td><td>17</td><td>17</td></tr> <tr><td>Upr L</td><td>58</td><td>58</td></tr> </tbody> </table>	Order	Header Type	Next Header Code	1	-	-	2	0	0	3	60	60	4	43	43	5	44	44	6	51	51	7	50	50	8	60	60	9	135	135		59	59	Upr L	6	6	Upr L	17	17	Upr L	58	58	<table border="1"> <thead> <tr> <th>Order</th> <th>Header Type</th> <th>Next Header Code</th> </tr> </thead> <tbody> <tr><td>1</td><td>Basic IPv6 header</td><td>-</td></tr> <tr><td>2</td><td>Hop-by-Hop Options</td><td>0</td></tr> <tr><td>3</td><td>Destination options (with routing options)</td><td>60</td></tr> <tr><td>4</td><td>Routing header</td><td>43</td></tr> <tr><td>5</td><td>Fragment header</td><td>44</td></tr> <tr><td>6</td><td>Authentication header</td><td>51</td></tr> <tr><td>7</td><td>Encapsulation Security Payload header</td><td>50</td></tr> <tr><td>8</td><td>Destination Options</td><td>60</td></tr> <tr><td>9</td><td>Mobility header</td><td>135</td></tr> <tr><td></td><td>No next header</td><td>59</td></tr> <tr><td>Upr L</td><td>TCP</td><td>6</td></tr> <tr><td>Upr L</td><td>UDP</td><td>17</td></tr> <tr><td>Upr L</td><td>ICMPv6</td><td>58</td></tr> </tbody> </table>	Order	Header Type	Next Header Code	1	Basic IPv6 header	-	2	Hop-by-Hop Options	0	3	Destination options (with routing options)	60	4	Routing header	43	5	Fragment header	44	6	Authentication header	51	7	Encapsulation Security Payload header	50	8	Destination Options	60	9	Mobility header	135		No next header	59	Upr L	TCP	6	Upr L	UDP	17	Upr L	ICMPv6	58
Order	Header Type	Next Header Code																																																																																							
1	-	-																																																																																							
2	0	0																																																																																							
3	60	60																																																																																							
4	43	43																																																																																							
5	44	44																																																																																							
6	51	51																																																																																							
7	50	50																																																																																							
8	60	60																																																																																							
9	135	135																																																																																							
	59	59																																																																																							
Upr L	6	6																																																																																							
Upr L	17	17																																																																																							
Upr L	58	58																																																																																							
Order	Header Type	Next Header Code																																																																																							
1	Basic IPv6 header	-																																																																																							
2	Hop-by-Hop Options	0																																																																																							
3	Destination options (with routing options)	60																																																																																							
4	Routing header	43																																																																																							
5	Fragment header	44																																																																																							
6	Authentication header	51																																																																																							
7	Encapsulation Security Payload header	50																																																																																							
8	Destination Options	60																																																																																							
9	Mobility header	135																																																																																							
	No next header	59																																																																																							
Upr L	TCP	6																																																																																							
Upr L	UDP	17																																																																																							
Upr L	ICMPv6	58																																																																																							
<h3>OSPFv3 Virtual Links</h3>	 <p>ipv6 router ospf 1 area 37 virtual-link 150.1.7.7</p>	<h3>IPv6 NAT-PT</h3> <p>Drawing / rules</p>	 <p>IPv6 SRC and DST</p> <p>/96 block</p> <p>IPv6 address maps to IPv4</p> <p>IPv4 src and dst</p> <ol style="list-style-type: none"> Rules to translate IPv4 source addrs to IPv6 addrs Rules to translate IPv6 source addrs to IPv4 addrs The /96 prefix to map the IPv4 address space to 	<p>Explain IPv6 NAT-PT</p> <p>NAT statements:</p> 	<p>IPv6 nat prefix 2000::/96</p> <p>ip nat v6v4 source 2001::99:99 1.1.1.1</p> <p>ip nat v4v6 source 2.2.2.2 2000::960B:202</p> <p>CHECK ANSWER MIGHT BE WRONG!</p>																																																																																				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

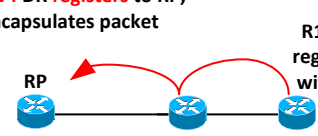
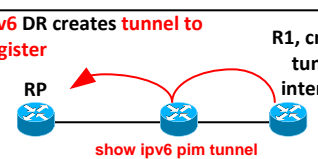
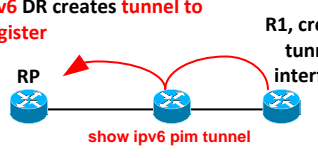
Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin





<p>Loopback1: prefix/64 Loopback2: prefix /61</p> <p>Advertising /59 summary to peer</p> <h2>IPv6 MP-BGP:</h2>	<p>Loopback1: prefix/64 Loopback2: prefix /61</p> <p>Advertising /59 summary to peer</p> <pre>interface Loopback 101 ipv6 address 2003:1:0:1::1/64 interface Loopback 102 ipv6 address 2003:1:0:11::11/61 router bgp 100 address-family ipv6 unicast neighbor 2001:1:0:1234::5 remote-as 500 neighbor 2001:1:0:1234::5 activate network 2003:1:0:1::/64 network 2003:1:0:10::/61 aggregate-address 2003:1::/59 summary-only</pre>	<h3>Show ipv6 pim interface</h3> <p>Output:</p>	<pre>R5#show ipv6 pim interface Interface PIM Nbr Hello DR Count Intvl Prior VoIP-Null0 off 0 30 1 Address: :: DR : not elected GigabitEthernet0/0 on 0 30 1 Address: FE80::221:A0FF:FE9A:F4C0 DR : this system GigabitEthernet0/1 off 0 30 1 Address: :: DR : not elected</pre>	<p>What is the difference between a DR registering with an RP in IPv4 and IPv6?</p>	<p>IPv4 DR registers to RP, encapsulates packet</p>  <p>RP R1, DR registers with RP</p> <p>IPv6 DR creates tunnel to register</p>  <p>R1, creates tunnel interface</p> <p>show ipv6 pim tunnel</p>
<p>Router bgp XXX</p> <pre>bgp default ipv4-unicast:</pre>	<p>bgp default ipv4-unicast</p> <p>means that all unicast IPv4 peers are automatically "activated"</p>	<h3>Show ipv6 pim neighbor</h3> <p>Output:</p>	<pre>R5#show ipv6 pim neighbor PIM Neighbor Table Mode: B - Bidir Capable, G - GenID Capable NeighborAddress Interface Uptime Expires Mode DR pri FE80::1 Serial0/0/0 00:06:21 00:01:19 B G 1 FE80::3 Serial0/0/0 00:06:21 00:01:20 B G 1 FE80::4 Serial0/1/0 00:06:21 00:01:17 B G 1</pre>	<h3>Show ipv6 pim tunnel:</h3> <p>Output:</p>	<p>IPv6 DR creates tunnel to register</p>  <p>R1, creates tunnel interface</p> <p>show ipv6 pim tunnel</p> <pre>R5#show ipv6 pim tunnel Tunnel0* Type : PIM Encap RP : Embedded RP Tunnel Source: 2001::5</pre>
<h2>IPv6 PIM and MLD</h2> <p>Flooding scope details:</p>	<p>PIM supports only Sparse-Mode operation</p> <p>flooding scope enforcement must be configured administratively using multicast filtering</p> <p>As soon as ipv6 multicast-routing is entered PIM becomes active</p> <p>To disable IPv6 pim use: no ipv6 pim</p>	<h2>IPv6 PIM BSR</h2> <p>Announcing itself as RP candidate:</p> <p>Announcing itself as BSR:</p> <p>BSR only announcing a static list of cRPs:</p>	<pre>Router announce itself as a RP via the BSR: ipv6 pim bsr candidate rp <IPv6 Address> ----- Announce itself as BSR: ipv6 pim bsr candidate bsr <IPv6 Address> BSR only announcing a static list of cRPs: ipv6 pim bsr announced rp <IPv6 Address></pre>	<p>What is the difference between a IPv6 unicast and a Multicast route:</p> <p>Using:</p> <p>Network 2001::1/64 Next-Hop FE80::1</p>	<pre>ipv6 route 2001::1/64 Gi0/0 FE80::1 ipv6 route 2001::1/64 Gi0/0 FE80::1 multicast</pre> <p>Route only usable by multicast</p>
<h2>IPv6 MLD details</h2> <p>IPv6 IPv4</p> <p>MLDv1 = IGMP... MLDv2 = IGMP...</p>	<p>Multicast Listener Discovery protocol, based on ICMPv6, replaced IGMP.</p> <p>Message types are: Query, Report, Done</p> <p>IPv6 IPv4</p> <p>MLDv1 = IGMPv2 MLDv2 = IGMPv3</p>	<h3>Usefull MLD commands:</h3>	<pre>show ipv6 pim group show ipv6 pim topology show ipv6 pim interface show ipv6 mfib</pre>	<h3>Show ipv6 pim bsr election</h3> <p>Output:</p> <p>To find the successful BSR candidate:</p>	<pre>R4#show ipv6 pim bsr election PIMv2 BSR information BSR Election Information Scope Range List: ff00::/8 This system is the Bootstrap Router (BSR) BSR Address: FC00:1:0:4::4 Uptime: 00:00:02, BSR Priority: 100, Hash mask length: 126 RPF: FE80::226:BFF:FE57:BA60, Loopback100 BS Timer: 00:00:57 This system is candidate BSR Candidate BSR address: FC00:1:0:4::4, priority: 100, hash mask length: 126</pre>
<h3>MLD:</h3> <p>Limit maximum of groups a member can join:</p> <p>Change mld query interval</p> <p>Mld time-out</p> <p>Response time</p>	<pre>ipv6 mld limit ipv6 mld queryinterval ipv6 mld query-timeout ipv6 mld query-max-reponsetime</pre>	<p>Joining a group via MLD:</p>	<pre>interface FastEthernet 0/0 ipv6 mld join-group ff76:0640:2001:CC1E::8</pre>	<p>How can you display the mappings of RPs to the multicast group ranges in IPv6?</p>	<pre>R4#show ipv6 pim range-list Static SSM Exp: never Learnt from :: FF33::/32 Up: 00:51:08 FF34::/32 Up: 00:51:08 FF35::/32 Up: 00:51:08 FF36::/32 Up: 00:51:08 FF37::/32 Up: 00:51:08 FF38::/32 Up: 00:51:08 FF39::/32 Up: 00:51:08 FF3A::/32 Up: 00:51:08 FF3B::/32 Up: 00:51:08 FF3C::/32 Up: 00:51:08 FF3D::/32 Up: 00:51:08 FF3E::/32 Up: 00:51:08 FF3F::/32 Up: 00:51:08 BSR SM RP: FC00:1:0:6::6 Exp: 00:02:13 Learnt from : FC00:1:0:4::4 FF00::/8 Up: 00:01:16</pre>
<h3>MLDv1 filtering:</h3> <p>Group based filtering:</p> <p>SSM source/group based filtering:</p>	<pre>ipv6 access-list ACL-X (permit deny) ipv6 <part1> <part2> Filter MLDv1, group based, set <part1> to any: MLDv1: ipv6 access-list MLDv1_FILER permit ipv6 any ff08::/64 source / group Filter MLDv1, SSM based, source/group: MLDv2: ipv6 access-list MLDv1_FILER permit ipv6 2000::25 ff08::/64</pre>	<h2>IPv6 SSM</h2> <p>Join a group ff36::8 Only sourced from 2001::25</p> <p>config</p>	<pre>IPv6 SSM, join group ff36::8 but only sourced from 2001::25 interface FastEthernet 0/0 ipv6 mld join-group ff36::8 2001::25</pre>	<h2>IPv6 Embedded RP</h2>	<p>Using 2002:6666::6 as the RP address</p> <pre>FF7x:y40:2002:6666::1 or FF7e:640:2002:6666::1 Only the RP needs to know that he is serving as RP, all other routers will figure out based on the Embedded IPv6 address!</pre> <pre>FF7E:0640:2002:6666::1 On the RP: conf t: ipv6 pim rp-address 2002:6666::6</pre> <p>FF Multicast FF7 Embedded RP IPv6 FF7E Global scope FF7E:0 must be 0 FF7E:06 Last digit of "rp-address"</p> <p>FF7E:0640 Prefix length of IPv6 addr 40 Hex = 64 decimal</p> <p>FF7E:0640:2002:6666::1111:1111 Multicast Group</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

<p>IPv6 Embedded RP address format:</p>		<p>IPv6 ISATAP Tunneling</p> <p>Show interface tun X</p> <p>Output:</p>	<pre>R4#show interface tunnel 345 Tunnel345 is up, line protocol is up Hardware is Tunnel MTU 1514 bytes, BW 9 Kbit/sec, DLY 500000 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation TUNNEL, loopback not set Keepalive not set Tunnel source 150.1.4.4 (Loopback0), destination UNKNOWN Tunnel protocol/transport IPv6 ISATAP Fast tunneling enabled Tunnel transmit bandwidth 8000 (kbps) Tunnel receive bandwidth 8000 (kbps)</pre>	<p>OSPFv3</p> <p>Clearing process/redistr./spf:</p>	<pre>clear ospfv3 <process-ID> force-spf clear ospfv3 <process-ID> process clear ospfv3 <process-ID> redistribution clear ipv6 ospf <process-ID> [process, force-spf, redistr]</pre>
<p>What is the difference between the following two tunnel configs:</p> <p>Interface Tunnel 88 Tunnel mode ipv6ip</p> <p>Interface Tunnel 26 (no mode specified)</p>	<p>Protocol 41</p> <pre>interface tunnel 88 tunnel source Loopback0 tunnel destination 1.1.1.1 ipv6 address 2001:2::2/64 tunnel mode ipv6ip</pre> <p>Protocol 47</p> <pre>interface tunnel 26 tunnel source Loopback0 tunnel destination 1.1.1.1 ipv6 address 2001:1::2/64 tunnel mode ipv6ip</pre>	<p>IPv6 ISATAP Tunneling</p>	<pre>R5# interface Tunnel345 ipv6 address 2001:1:0:345::/64 eui-64 tunnel source Loopback0 tunnel mode ipv6ip isatap ! interface Loopback100 ipv6 address 2001:1:0:5::/64 ! ipv6 route 2001:1:0:4::/64 2001:1:0:345:0:SeFe:9601:404 ipv6 route 2001:1:0:3::/64 2001:1:0:345:0:SeFe:9601:303</pre>	<p>BFD Support for EIGRP IPv6</p>	<pre>interface fa0/x bfd interval 50 min_rx 50 multiplier 3 router eigrp NAME address-family ipv6 autonomous-system 88 af-interface default bfd</pre>
<p>Based on what can you filter the establishment of IP tunnels utilizing Access-Lists?</p>	<p>Protocol 41: (ipv6ip)</p> <pre>access-list 100 [permit deny] 41 any any interface tunnel 88 ... ipv6 address 2001:2::2/64 tunnel mode ipv6ip</pre> <p>Protocol 47: (GRE)</p> <pre>access-list 100 [permit deny] 47 any any interface tunnel 26 ... ipv6 address 2001:1::2/64</pre>	<p>IPv6 ISATAP Tunneling</p> <p>Addressing format:</p>	<p>ISATAP Addressing:</p> <p>EUI-64 = 0000 (16 bits) + 5EFE (16 bits) + IPv4 Address (32 bits):</p> <p>IPv6 Prefix 2001:1:0:345::/64 will have:</p> <p>2001:1:0:345:0:5efe:9601:0303/64 150.1.3.3</p>	<p>Assigning FE80::1 the same Link-local addresses to several Interfaces:</p> <p>How do you ping other FE80 addresses?</p>	
<p>Automatic 6to4 Tunneling</p>	<p>Automatic 6to4 Tunneling are Multipoint by design:</p> <pre>interface Tunnel 345 tunnel source Loopback0 tunnel mode ipv6ip 6to4 ipv6 address 2002:9601:303::3/64 ipv6 route 2002::/16 Tunnel 345 (route destination through tun)</pre> <p><i>No Tunnel destination configured</i></p>	<p>Automatic 6to4 Tunneling:</p> <p>General info about 6to4:</p>	<p>- 6to4 tunnels are multipoint by design</p> <p>- router extracts the IPv4 address embedded in the IPv6 address</p> <p>- need to use the 16-bit prefix 2002,</p> <p>- only static routing is possible, (usually 2002::/16)</p> <p>2002 (16 bits): IPv4 address (32 bits): Subnet ID(16 bits): Interface ID (64 bits)</p> <p>150.1.3.3 = 2002:9601:303::/48 150.1.3.3</p>	<p>Will R1 be able to ping FE80::3 ?</p>	<p>R1 can NOT ping R3's FE80::3 Link local address!</p>
<p>IPv6 Subnet Reference Prefix Lengths:</p>		<p>Automatic 6to4 Tunneling:</p> <p>Show interface tunnel:</p>	<pre>R4#show interfaces tunnel 345 Tunnel345 is up, line protocol is up Hardware is Tunnel MTU 1514 bytes, BW 9 Kbit/sec, DLY 500000 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation TUNNEL, loopback not set Keepalive not set Tunnel source 150.1.4.4 (Loopback0), destination UNKNOWN Tunnel protocol/transport IPv6 6to4 Fast tunneling enabled Tunnel transmit bandwidth 8000 (kbps) Tunnel receive bandwidth 8000 (kbps) Queueing strategy: fifo Output queue: 0/0 (size/max) 0 output buffer failures, 0 output buffers swapped out</pre>	<p>What will the outputs show in regards to joined multicast groups?</p> <pre>interface Ethernet0/0 mac-address 0000.7799.8888 ipv6 enable show ipv6 interface Ethernet0/0 is up, ... IPv6 is enabled, link-local addr.. FE80::200:77FF:FE99:8888 Joined group address(es): FF02::1 FF02::1:FF99:8888</pre>	<p>Config 1:</p> <pre>interface Ethernet0/0 mac-address 0000.7799.8888 ipv6 enable</pre> <p>Config 2:</p> <pre>ipv6 unicast-routing interface Ethernet0/0 mac-address 0000.7799.8888 ipv6 enable</pre> <p>Output 1</p> <pre>show ipv6 interface Ethernet0/0 is up, ... IPv6 is enabled, link-local addr.. FE80::200:77FF:FE99:8888 Joined group address(es): FF02::1 FF02::1:FF99:8888</pre> <p>Output 2</p> <pre>show ipv6 interface IPv6 is enabled, link-local addr.. FE80::200:77FF:FE99:8888 Joined group address(es): FF02::1 FF02::2 ← ipv6 unicast-routing enabled, adds group FF02::2 ! FF02::1:FF99:8888</pre>
<p>Example of /48 allocations:</p>	<p>2402:9400:10::/48</p> <p>2402:9400:11::/48</p> <p>.....</p> <p>2402:9400:1F::/48</p> <p>2402:9400:20::/48</p> <p>2402:9400:21::/48</p> <p>2402:9400:22::/48</p> <p>.....</p> <p>2402:9400:2F::/48</p> <p>2402:9400:30::/48</p>	<p>Leave your IPv6 calculator at home:</p> <p>2402:9400:1234:1234::/?</p> <p>2402:9400:1234:123X::/?</p> <p>2402:9400:1234:12XX::/?</p> <p>2402:9400:1234:1XXX::/?</p> <p>2402:9400:1234:XXXX::/?</p> <p>2402:9400:123X:XXXX::/?</p> <p>2402:9400:12XX:XXXX::/?</p>	<p>2402:9400:1234:1234::/64</p> <p>2402:9400:1234:123X::/60</p> <p>2402:9400:1234:12XX::/56</p> <p>2402:9400:1234:1XXX::/52</p> <p>2402:9400:1234:XXXX::/48</p> <p>2402:9400:123X:XXXX::/44</p> <p>2402:9400:12XX:XXXX::/40</p>	<p>IPv6 Multicast addresses:</p> <p>IPv6 RIP: FF02::9 (224.0.0.9)</p> <p>IPv6 EIGRP: FF02::10 (224.0.0.10)</p> <p>IPv6 OSPF: FF02::5 (224.0.0.5)</p> <p>After DR election FF02::6 appears (224.0.0.6)</p> <p>mDNS: FF02::B</p> <p>send to FF02::B, if there is something it will reply</p>	<p>IPv6 Multicast addresses:</p> <p>IPv6 RIP: FF02::9 (224.0.0.9)</p> <p>IPv6 EIGRP: FF02::10 (224.0.0.10)</p> <p>IPv6 OSPF: FF02::5 (224.0.0.5)</p> <p>After DR election FF02::6 appears (224.0.0.6)</p> <p>mDNS: FF02::B</p> <p>send to FF02::B, if there is something it will reply</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

<p>IPv6 SLAAC Configuration:</p>	<pre> R1# ipv6 unicast-routing int ser1/2 ipv6 address 2001::1/64 ipv6 enable clock rate 64000 </pre>	<p>Configure int Loopback 0 with IPv6 addr: 1:1:1:1::/64 using the mac address of the first LAN interface:</p>	<pre> conf t interface Loopback 0 ipv6 addr 2:2:2:2::/64 eui-64 </pre> <p>Will instruct the router to use the first LAN interface in regards to the MAC address!</p>	<p>What will the following command display:</p> <pre> show ipv6 interface e0/x </pre> <pre> R1# interface Ethernet0/0 ipv6 address FE80::9 link-local interface Ethernet0/1 ipv6 address FE80::9 link-local SW1# interface e0/0 switchport access vlan 10 interface e0/1 switchport access vlan 10 </pre>	<pre> R1# interface Ethernet0/0 ipv6 address FE80::9 link-local interface Ethernet0/1 ipv6 address FE80::9 link-local R5#show ipv6 eigrp 100 interfaces detail serial 2/1 EIGRP-IPv6 Interfaces for AS(100) </pre>
<p>IPv6 SLAAC And options via DHCP</p>	<pre> ipv6 unicast-routing int e0/0 ipv6 address 2001::1/64 ipv6 enable ipv6 nd other-config-flag ipv6 dhcp pool TEST address prefix 99::/64 lease server 2000::1 domain name its-a-bit-buggy.com </pre>	<p>What is the correct FE80::x address ?</p> <p>interface loopback0 ipv6 address 2:2:2:2:/64</p> <p>MAC of first LAN interface: aabb.cc00.0200</p>	<pre> interface Loopback 0 ipv6 address x:x:x:x:/64 eui-64 </pre> <p>MAC of first LAN interface: AABB.CC00.0200</p> <p>AA = 1010 1010 AA flipped 7th bit = 1010 1000 HEX = A 8 Finally: FE80::A8BB:CCFF:FE00:0200 FFFE inserted</p>	<p>What to expect of</p> <pre> show ipv6 eigrp 100 interfaces detail serial 2/x </pre>	<pre> R5#show ipv6 eigrp 100 interfaces detail serial 2/1 EIGRP-IPv6 Interfaces for AS(100) Multicast Pending Xmit Queue PeerQ Mean Pacing Time Interface Peers Un/Reliable Un/Reliable SRTT Un/Reliable Flow Timer Routes Se2/1 1 0/0 0/0 14 0/15 71 0 Hello-interval is 5, Hold-time is 15 Split-horizon is enabled Next xmit serial <none> Packetized sent/expedited: 12/0 Hello's sent/expedited: 159/3 Un/reliable mcasts: 0/0 Un/reliable ucasts: 12/15 Mcast exceptions: 0 CR packets: 0 ACKs suppressed: 1 Retransmissions sent: 1 Out-of-sequence rcvd: 1 Topology-ids on interface - 0 Authentication mode is not set </pre>
<p>IPv6 over DMVPN (IPv4)</p>		<p>Which address do you use for static routing? Using public IP's</p> <pre> R1# interface Loopback0 ipv6 address FEED:BEEF::1/128 interface Ethernet0/0 mac-address 0000.0000.1111 ipv6 address FE80::1 link-local ipv6 address 2001::1/64 ipv6 route FEED:BEEF::2/128 ??????? R2# interface Loopback0 ipv6 address FEED:BEEF::2/128 interface Ethernet0/0 ipv6 address FE80::2 link-local ipv6 address 2001::2/64 ipv6 route FEED:BEEF::1/128 ????? </pre>	<pre> R1# interface Loopback0 ipv6 address FEED:BEEF::1/128 interface Ethernet0/0 mac-address 0000.0000.1111 ipv6 address FE80::1 link-local ipv6 address 2001::1/64 ipv6 route FEED:BEEF::2/128 2001::2 R2# interface Loopback0 ipv6 address FEED:BEEF::2/128 interface Ethernet0/0 ipv6 address FE80::2 link-local ipv6 address 2001::2/64 ipv6 route FEED:BEEF::1/128 2001::1 </pre>	<p>Full IPv6 EIGRP authentication</p> <p>Config:</p> <p>Use a key of "Cisco" EIGRP 100</p>	<pre> R3# key chain KEY key 1 key-string Cisco interface Serial1/x ipv6 address FE80::3 link-local ipv6 eigrp 100 ipv6 authentication mode eigrp 100 md5 ipv6 authentication key-chain eigrp 100 KEY ipv6 router eigrp 100 eigrp router-id 0.0.0.3 no shut Verify: show ipv6 eigrp 100 interfaces detail serial 1/x Authentication mode is md5, key-chain is "KEY" </pre>
<p>IPv6 over DMVPN (IPv4) Using link-local address space</p>	<p>If using Link-Local Addresses configure them fixed. It makes your live way easier</p> <p>Looking at show ipv6 nhrp</p>	<p>What does the static route look like using link local addresses?</p> <pre> R1# interface Loopback0 ipv6 address FEED:BEEF::1/128 interface Ethernet0/0 mac-address 0000.0000.1111 ipv6 address FE80::1 link-local ipv6 address 2001::1/64 ipv6 route FEED:BEEF::2/128 ??????? R2# interface Loopback0 ipv6 address FEED:BEEF::2/128 interface Ethernet0/0 ipv6 address FE80::2 link-local ipv6 address 2001::2/64 ipv6 route FEED:BEEF::1/128 ????? </pre>	<pre> R1# interface Loopback0 ipv6 address FEED:BEEF::1/128 interface Ethernet0/0 mac-address 0000.0000.1111 ipv6 address FE80::1 link-local ipv6 address 2001::1/64 ipv6 route FEED:BEEF::2/128 e0/0 FE80::2 R2# interface Loopback0 ipv6 address FEED:BEEF::2/128 interface Ethernet0/0 ipv6 address FE80::2 link-local ipv6 address 2001::2/64 ipv6 route FEED:BEEF::1/128 e0/0 FE80::2 </pre>	<p>IPv6 EIGRP filtering:</p> <p>Filter 1::1/64 on R6:</p>	<pre> R6# ipv6 router eigrp 100 distribute-list prefix-list PFX-1 in eigrp router-id 0.0.0.6 ipv6 prefix-list PFX-1 seq 5 deny 1::/64 ipv6 prefix-list PFX-1 seq 10 permit ::/0 le 128 </pre>
<p>Output of: debug ipv6 nd and following config:</p> <pre> ipv6 unicast-routing int e0/0 mac-address 0000.0000.1111 ipv6 enable </pre>	<pre> ICMPv6-ND: IPv6 Opr Enabled on Ethernet0/0 ICMPv6-ND: L2 came up on Ethernet0/0 (Layer 2 UP) [performing DAD] IPv6-Addrmgr-ND: DAD request for FE80::200:FF:FE00:1111 on Ethernet0/0 [checking if unique] ICMPv6-ND: Sending NS for FE80::200:FF:FE00:1111 on Ethernet0/0 (no response back, address seems unique) IPv6-Addrmgr-ND: DAD: FE80::200:FF:FE00:1111 is unique. [sending last "warning" here I come] ICMPv6-ND: Sending NA for FE80::200:FF:FE00:1111 on Ethernet0/0 ICMPv6-ND: L3 came up on Ethernet0/0 (Layer 3 UP) ICMPv6-ND: Linklocal FE80::200:FF:FE00:1111 on Ethernet0/0, Up [due to ipv6 unicast routing enabled on the router] ICMPv6-ND: Created RA context for FE80::200:FF:FE00:1111/Ethernet0/0 [Router sending RA announcements to FF02::1] ICMPv6-ND: Request to send RA for FE80::200:FF:FE00:1111 ICMPv6-ND: Setup RA from FE80::200:FF:FE00:1111 to FF02::1 on Ethernet0/0 ICMPv6-ND: MTU = 1500 ICMPv6-ND: ND output feature SEND executed on 3 - rc=0 </pre>	<p>How do you configure a static IPv6 default route?</p> <pre> ipv6 route ::/0 e0/0 fe80::1 ipv6 route ::/0 2001::1 </pre>	<p>What to expect of:</p> <pre> R1#show ipv6 ospf database </pre> <pre> interface Loopback0 ipv6 address 2::2/64 interface Ethernet0/0 ipv6 address 12::2/64 ipv6 ospf 1 area 1 ipv6 router ospf 1 router-id 0.0.0.2 interface Loopback0 ipv6 address 1::1/64 interface Ethernet0/0 ipv6 address 12::1/64 ipv6 ospf 1 area 1 ipv6 router ospf 1 router-id 0.0.0.1 </pre>	<pre> R1#show ipv6 ospf database OSPFv3 Router with ID (0.0.0.1) (Process ID 1) Router Link States (Area 1) ADV Router Age Seq# Fragment ID Link count Bits 0.0.0.1 409 0x80000002 0 1 None 0.0.0.2 410 0x80000002 0 1 None Net Link States (Area 1) ADV Router Age Seq# Link ID Rtr count 0.0.0.2 410 0x80000001 3 2 Link (Type-8) Link States (Area 1) ADV Router Age Seq# Link ID Interface 0.0.0.1 471 0x80000001 3 Et0/0 0.0.0.2 451 0x80000001 3 Et0/0 Intra Area Prefix Link States (Area 1) ADV Router Age Seq# Link ID Ref-Istype Ref-LSID 0.0.0.1 409 0x80000003 0 0x2001 0 0.0.0.2 410 0x80000002 0 0x2001 0 0.0.0.2 410 0x80000001 3072 0x2002 3 </pre>	
<p>What will be the IPv6 link local address of the following config look like:</p> <pre> interface Ethernet0/2 mac-address 0012.3456.8910 ipv6 enable </pre>	<pre> interface Ethernet0/2 mac-address 0012.3456.8910 ipv6 enable FE80::212:34FF:FE56:8910 or FE80::0212:34FF:FE56:8910 </pre>	<p>Host entries for IPv6 Ips on routers:</p> <p>Make your life easier pinging:</p> <pre> R1# ping R1 Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms </pre>	<p>Configure IPv6 EIGRP to gain reachability between the loopbacks:</p> <pre> R1# interface Loopback0 ipv6 address 12::1/64 int e0/0 ipv6 address fe80::1 link-local R2# interface Loopback0 ipv6 address 2::2/64 int e0/0 ipv6 address 12::2/64 ipv6 address fe80::2 link-local </pre>	<pre> R1# ipv6 unicast-routing interface Loopback0 ipv6 address 1::1/64 int e0/0 ipv6 address 12::1/64 ipv6 address fe80::1 link-local ipv6 eigrp 100 ipv6 router eigrp 100 NO SHUTDOWN eigrp router-id 0.0.0.1 R2# ipv6 unicast-routing interface Loopback0 ipv6 address 2::2/64 int e0/0 ipv6 address 12::2/64 ipv6 address fe80::2 link-local ipv6 eigrp 100 ipv6 router eigrp 100 NO SHUTDOWN eigrp router-id 0.0.0.2 </pre>	

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

Why are the following OSPFv6 routers not forming an adjacency?
Or seen with a flapping adjacency?

```

R1 interface Serial1/0
  ipv6 address FE80::3 link-local
  ipv6 address 23::1/64
  clock rate 64000
R2 interface Serial1/0
  ipv6 address FE80::3 link-local
  ipv6 address 23::2/64
  clock rate 64000
  
```

```

R1#show ipv6 interface e0/0
Ethernet0/0 is up, line protocol is up
IPv6 is stalled, link-local address is FE80::3 [DUP]
R2#show ipv6 interface e0/0
Ethernet0/0 is up, line protocol is up
IPv6 is stalled, link-local address is FE80::3
  
```

How to calculate the path from here
To there in OSPFv6:

```

R1#show ipv6 ospf database router
Routing Bit Set on this LSA
Advertising Router: 0.0.0.2
Area Border Router
Link connected to: a Transit Network
Link Metric: 10
Local Interface ID: 3
Neighbor (DR) Interface ID: 3
Neighbor (DR) Router ID: 0.0.0.2
  
```

```

R1#show ipv6 ospf database inter-area prefix adv-router 0.0.0.2
...
Advertising Router: 0.0.0.2
Metric: 64
Prefix Address: 3::3
  
```

```

R1#show ipv6 route ospf
O1 3::3/128 [110/74]
via FE80::2, Ethernet0/0
  
```

What to expect of a
R1# show ipv6 ospf database router

Output:

```

R1#show ipv6 ospf database router
OSPFv3 Router with ID (0.0.0.1) (Process ID 1)
Router Link States (Area 1)
LS age: 446
Options: (V6-Bit, E-Bit, R-bit, DC-Bit)
LS Type: Router Links
Link State ID: 0
Advertising Router: 0.0.0.1
LS Seq Number: 8000000A
Checksum: 0xF5ED
Length: 40
Number of Links: 1
Link connected to: a Transit Network
Link Metric: 10
Local Interface ID: 3
Neighbor (DR) Interface ID: 3
Neighbor (DR) Router ID: 0.0.0.2
  
```

What can happen to an existing OSPFv3 session if you do the following:

```

R1#conf t
interface e0/0
ipv6 address fe80::1 link-local
  
```

The OSPFv3 session gets dropped!!!
→ keep in mind in regards to the daily operation!!

```

R1#
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial2/1 from FULL to DOWN, Neighbor Down: Interface down or detached
%OSPFv3-5-ADJCHG: Process 1, Nbr 0.0.0.3 on Serial2/1 from LOADING to FULL, Loading Done
  
```

```

R1#conf t
interface e0/0
ipv6 address fe80::1 link-local
  
```

What to expect of a
R1# show ipv6 ospf database link

Output:

```

R1#show ipv6 ospf database link
OSPFv3 Router with ID (0.0.0.1) (Process ID 1)
Link (Type-8) Link States (Area 1)
LS age: 1147
Options: (V6-Bit, E-Bit, R-bit, DC-Bit)
LS Type: Link-LSA (Interface: e0/0)
Link State ID: 3 (Interface ID)
Advertising Router: 0.0.0.1
LS Seq Number: 80000006
Checksum: 0xB973
Length: 56
Router Priority: 1
Link Local Address: FE80::1
Number of Prefixes: 1
Prefix Address: 12::
Prefix Length: 64, Options: None
  
```

How can you find the cost to the ASBR in OSPFv3 from R1?

```

R1#show ipv6 ospf database
Type-5 AS External Link States
ADV Router Age Seq# Prefix
0.0.0.4 953 0x80000001 4::4/64
  
```

```

R1#show ipv6 ospf database inter-area router 0.0.0.4
OSPFv3 Router with ID (0.0.0.1) (Process ID 1)
Inter Area Router Link States (Area 1)
Advertising Router: 0.0.0.2
LS Seq Number: 80000001
Checksum: 0xDD9E
Length: 32
Metric: 128
Destination Router ID: 0.0.0.4
  
```

```

R1#show ipv6 route ospf
OE2 4::4/64 [110/20]
via FE80::2, e0/0
  
```

Forward metric 128+10 = 138

What to expect of a
R1# show ipv6 ospf database prefix

Output:

```

R1#show ipv6 ospf database prefix
OSPFv3 Router with ID (0.0.0.1) (Process ID 1)
Intra Area Prefix Link States (Area 1)
Routing Bit Set on this LSA
LS age: 729
LS Type: Intra-Area-Prefix-LSA
Link State ID: 0
Advertising Router: 0.0.0.1
LS Seq Number: 80000004
Checksum: 0x9C3E
Length: 72
Referenced LSA Type: 2001
Referenced Link State ID: 0
Referenced Advertising Router: 0.0.0.1
Number of Prefixes: 2
Prefix Address: 1::1
Prefix Length: 128, Options: LA, Metric: 0
Prefix Address: 1::1
Prefix Length: 128, Options: LA, Metric: 0
  
```

How will 4::4/64 be seen on R1 in case of it being redistributed into OSPFv3 as Type-1 or Type-2 ?

```

Type-1:
R1#show ipv6 route ospf
OE1 4::4/64 [110/158]
via FE80::2, Ethernet0/0
  
```

```

R4#
ipv6 router ospf 1
redistribute connected route-map RMP-CON metric-type 1
  
```

```

Type-2:
R1#show ipv6 route ospf
OE2 4::4/64 [110/20]
via FE80::2, Ethernet0/0
  
```

```

R4#
ipv6 router ospf 1
redistribute connected route-map RMP-CON metric-type 2
  
```

Total forward metric + 20 = 158

What to expect of a
R1#show ipv6 ospf database prefix 0

Output:

```

R2#show ipv6 ospf database prefix 0
OSPFv3 Router with ID (0.0.0.2) (Process ID 1)
Intra Area Prefix Link States (Area 1)
Routing Bit Set on this LSA
LS age: 1159
LS Type: Intra-Area-Prefix-LSA
Link State ID: 0
Advertising Router: 0.0.0.1
LS Seq Number: 80000004
Checksum: 0x9C3E
Length: 72
Referenced LSA Type: 2001
Referenced Link State ID: 0
Referenced Advertising Router: 0.0.0.1
Number of Prefixes: 2
Prefix Address: 1::1
Prefix Length: 128, Options: LA, Metric: 0
Prefix Address: 1::1
Prefix Length: 128, Options: LA, Metric: 0
  
```

prefix 0 within an area shows all attached prefixes, or look at the specific LSA-ID

How to check prefixes within the Area, and for prefixes learned from other Areas in OSPFv6 ?

For prefixes within the Area:

```

R1#show ipv6 ospf database prefix 0
Referenced Advertising Router: 0.0.0.1
Number of Prefixes: 2
Prefix Address: 1::1
...
Referenced Advertising Router: 0.0.0.2
Number of Prefixes: 1
Prefix Address: 2::2
  
```

For Prefixes learned from other areas

```

R1#show ipv6 ospf database
Inter Area Prefix Link States (Area 1)
ADV Router Age Seq# Prefix
0.0.0.2 739 0x80000001 23::/64
0.0.0.2 575 0x80000001 3::3/128
  
```

What to expect of:
R1# show ipv6 route ospf

```

R1#show ipv6 route ospf
IPv6 Routing Table - default - 11 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
ND - ND Default, NDp - ND Prefix, DCE - Destination, NDR - Redirect
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, I - ISIP
  
```

```

O 2::128 [110/10]
via FE80::2, Ethernet0/0
O 3::3/128 [110/74]
via FE80::2, Ethernet0/0
O 23::/64 [110/74]
via FE80::2, Ethernet0/0
  
```

How can you look at those in detail?

```

R1#show ipv6 ospf database
...
Inter Area Prefix Link States (Area 1)
ADV Router Age Seq# Prefix
0.0.0.2 1925 0x80000001 2::2/128
0.0.0.2 1925 0x80000001 23::/64
0.0.0.2 1761 0x80000001 3::3/128
  
```

```

R1#show ipv6 ospf database inter-area prefix
OSPFv3 Router with ID (0.0.0.1) (Process ID 1)
Inter Area Prefix Link States (Area 1)
Routing Bit Set on this LSA
LS age: 139
LS Type: Inter Area Prefix Links
Link State ID: 0
Advertising Router: 0.0.0.2
LS Seq Number: 80000002
Checksum: 0x891E
Length: 44
Metric: 0
Prefix Address: 2::2
Prefix Length: 128, Options: None
  
```

```

Routing Bit Set on this LSA
LS age: 139
LS Type: Inter Area Prefix Links
Link State ID: 1
Advertising Router: 0.0.0.2
LS Seq Number: 80000002
Checksum: 0x8807
Length: 36
Metric: 64
Prefix Address: 23::
Prefix Length: 64, Options: None
  
```

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<p>Describe all available types of ACLs:</p>	<p>Standard ACL Extended ACL IP named ACL Lock and Key ACL (Dynamic ACLs) Reflexive ACL Established ACL Time-based ACL (time-range) Distributed time-based ACL Turbo ACL Receive ACL Infrastructure protection ACL Transit ACL Classification ACL Debugging traffic ACL</p>	<p>Turbo ACLs:</p>	<p>Compiles access-list for acceleration:</p> <p>Router-config# access-list compiled</p> <p>show access-list compiled</p>	<p>Radius Packet header:</p>	<p>Code 1 Access-Request 2 Access-Accept 3 Access-Reject 4 Accounting-Request 5 Accounting-Response 11 Access-Challenge</p> <p>Identifier: Message Sequence Number, allows Radius client to match a Radius response Authenticator: Used to authenticate the reply from the Radius server, MD5 hash</p> <p>UDP 1812 (1645) for authentication / authorization UDP 1813 (1646) for accounting requests</p>
<p>Standard Access-list parameters:</p>	<pre>access-list 10 permit 10.1.1.0 0.0.0.255 access-list 10 permit 10.1.1.0 0.0.0.255 log-input (Logs Layer 2 interface information) Interface X ip access-group 10 [in,out]</pre>	<p>Classification ACL:</p>	<p>Usually holds only permit statements for various protocols, ports or flags. Used to identify traffic, or DoS attacks:</p> <pre>access-list 101 permit icmp any any eq echo access-list 101 permit tcp any any eq syn access-list 101 permit tcp any any eq fragment access-list 101 permit udp any any eq fragment access-list 101 permit ip any any eq fragment access-list 101 permit tcp any any access-list 101 permit udp any any access-list 101 permit icmp any any access-list 101 permit ip any any</pre>	<p>VACL Vlan access-list config:</p>	<pre>vlan access-map VACL-1 10 match ip address ACL-IPv4 match ipv6 address ACL-IPv6 match mac address ACL-MAC-ADDR action [forward,drop] vlan filter VACL-1 vlan-list 25-37 (vlans 25-37)</pre>
<p>Extended Access-list Parameters:</p>	<pre>access-list 101 permit [tcp,udp,icmp,ip,PROTOCOL] access-list 101 permit tcp any any eq 23 Interface X ip access-group 101 [in,out] TCP options: Timeout in minutes, port, established, precedence, tos, log, log-input, time-range, fragments UDP options: TCP options: Timeout in minutes, port, precedence, tos, log, log-input, time-range, fragments</pre>	<p>Tokens used for Banner messages:</p>	<pre>\$(hostname) \$(domain) \$(peer-ip) \$(gate-ip) \$(encap) \$(encap-alt) \$(mtu) \$(line) \$(line-desc)</pre> <p>SLIP/PPP gateway SLIP/PPP SL/IP</p>	<p>VACL on a routed port Order of processing:</p>	<ol style="list-style-type: none"> VACL for input Vlan Input IOS ACL Output IOS ACL VACL for output Vlan
<p>Lock and Key ACL (Dynamic ACL) Config:</p>	<pre>username test password cisco123 line vty 0 4 login local ! Time-out per uses basis after 10 minutes username test autocommand access-enable host timeout 10 ! Time-out any user globally after 10 minutes: line vty 0 4 autocommand access-enable host timeout 10 access-list 102 permit tcp any host <router-ip> eq telnet access-list 102 dynamic ACCESS timeout 15 permit tcp any any eq 80 deny ip any any log interface X ip access-group 102 in</pre>	<p>Explanations:</p> <ul style="list-style-type: none"> No ip source-route No ip proxy-arp No ip gratuitous-arps No ip directed-broadcast No ip redirects No ip unreachable 	<p>Disables strict / loose source routing</p> <p>Router replies with own MAC for an IP attached to its other interface.</p> <p>Turns of unsolicited ARP broadcasts IP of host but Routers MAC addr.</p> <p>Can be used to map a network, also to use as amplifier in a DoS attack.</p> <p>Disables ICMP redirects, is enabled with HSRP by default</p> <p>Will not send unreachable messages for traffic pointing to Null0 interface.</p>	<p>Switchport Port-security Config:</p>	<pre>switchport port-security switchport port-security max 5 / broken down to 3 and 2 switchport port-security 3 vlan access switchport port-security 2 vlan voice switchport port-security agging -timeout - static - type inactivity (idle) - type absolute (ttl 5min clears cam tbl!) switchport port-security agging time switchport port-security violation shutdown Discards traffic of exceeding MACs SNMP msg! switchport port-security violation restrict Discards silently, no SNMP msg switchport port-security violation protect</pre>
<p>Reflexive ACLs Config:</p>	<pre>ip access-group ACL-IN in ip access-group ACL-OUT out ip access-list extended ACL-IN evaluate tcp_reflect ip access-group extended ACL-OUT permit tcp 10.0.0.0/24 192.x.x/16 reflect tcp_reflect ip reflexive-list timeout <TIMEOUT> ! timing out of old sessions</pre>	<p>Auto-Secure:</p>	<pre>conf t auto-secure ! Disables common services auto-secure no-interact ! User not prompted for interaction show auto secure config</pre> <p>Shows all config that has been added part of the secure process.</p>	<p>Port Access-List PACL config:</p>	<p>Does not support – log option Does not support MPLS / ARP filtering Filtering is based on the fields of the Ethernet datagr see how much TCAM space is available: show tcam counts</p> <pre>interface X [ip, mac] access-group XXX [in, out]</pre>
<p>Time-based ACLs:</p>	<pre>time-range MY-RANGE absolute [start time date] [end time date] periodic days-of-week hh:m to days-of-week hh:mm Apply via: access-list <NR> permit tcp <SRC> <DST> time-range MY-RANGE time-range WEEKDAYS_EVES periodic weekdays 18:00 to 23:59 periodic weekdays 0:00 to 8:59 - interval from 00:01 to 02:00 will last from 00:01:00 to 02:00:59 - So if you want to end at exactly 04:00pm use to 15:59</pre>	<p>How to block unicast/multicast frames with a unknown destination on switches?</p>	<pre>conf t interface X switchport block unicast switchport block multicast show interface fa0/1 switchport Unknown unicast blocked: enabled Unknown multicast blocked: enabled</pre> <p>Bad guy flooding the switch with random SRC MAC addrs -> forcing the switch to go into open fail mode</p> <p>Flood traffic hits B, if switchport block not used</p>	<p>PACL, VACL, IOS ACL Diagram:</p>	

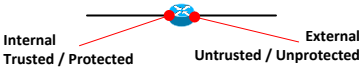
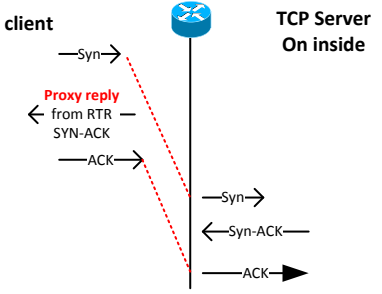
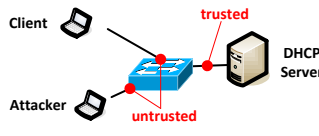
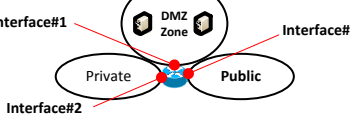
Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts



Colin

<p>IP Source Guard</p> <p>Config:</p>	<p>DHCP Snooping (prerequisite) - IP Source Guard</p> <p>ip dhcp snooping ip dhcp snooping vlan X</p> <p>L3 check only: interface fa0/x ip verify source</p> <p>Any traffic incoming with a source other than assigned via DHCP or static will be filtered out.</p> <p>L2+L3 combination check interface fa0/x switchport port-security ip verify source port-security</p> <p>ip source binding 1234.1234.abcd vlan 22 1.1.1.1 int fa0/x Static entry for MAC 1234.x in vlan 22.</p>	<p>Configuring CBAC</p> <p>Steps:</p> 	<p>1. select interface internal, external 2. configure ACL 3. define inspection rule 4. configure timeouts and thresholds 5. apply ACL and inspection rule to an interface 6. verify CBAC</p> <p>Fragmented packets can't be checked by CBAC!</p>	<p>IP source tracker</p>	<p>Collect traffic flow statistics from host 1.1.1.1:</p> <pre>Router-config# ip source-track 1.1.1.1 ip source-track syslog-interval 2 ip source-track export-interval 30 Syslog every 2 sec / exports flow info every 30 sec from route processor. show ip source-track <IP-ADDR> Address SRC IF Bytes Pkts Bytes/s Pkts/s 1.1.1.1 Fa0/0 123G345M 562342 134535 show ip source-track summary</pre>
<p>IP Source Guard</p> <p>Config / troubleshooting:</p>	<p>ip dhcp snooping ip dhcp snooping vlan X</p> <p>OR</p> <p>L3 SRC and SRC MAC check interface x ip verify source port-security switchport port-security</p> <p>L3 src check only: interface x ip verify source</p> <p>Configure static binding: ip source binding xxxx.xxxx.xxxx vlan X 10.1.1.1 int fa0/x</p> <p>show ip verify source show ip source binding</p>	<p>Configuring CBAC:</p>	<p>Inspection Rule: ip inspect name CBAC http ip inspect name CBAC ftp</p> <p>Inspection timeout / thresholds ip inspect tcp synwait-time <30 sec> ip inspect udp</p> <p>Interface fa0/1 ip inspect CBAC in Interface fa0/0 ip inspect CBAC out</p> <p>Internal Trusted / Protected External Untrusted / Unprotected</p> <p>Show ip inspect [config- interface]</p>	<p>TCP intercept</p> <p>Diagram:</p>	 <p>Router glues ACK together once the "outside" three way hand-shake has been confirmed.</p>
<p>DAI</p> <p>Dynamic ARP inspection</p> <p>Config:</p>	<p>DAI relies on IP DHCP Snooping. Checks if ARP message is correct, claiming to be the right IP -> mitigates ARP poisoning</p> <p>ip dhcp snooping ip dhcp snooping vlan X</p> <p>ip arp inspection vlan x ip arp inspection filter ACL vlan X</p> <p>arp access-list ACL permit ip host 1.1.1.1 mac host xxxx.xxxx.xxxx</p> <p>NOT relying on / using IP DHCP Snooping (static keyword):</p> <p>ip arp inspection vlan x ip arp inspection filter ACL vlan 100 static arp access-list ACL permit ip host 1.1.1.1 mac host xxxx.xxxx.xxxx</p>	<p>Virtual Fragmentation Reassembly (VRF)</p> <p>Configuration:</p>	<p>Interface fa0/x ip virtual-reassembly max-reassemblies 100 max-fragments 20 timeout 5</p> <p>Max 100 IP datagrams to be reassembled</p> <p>Max of 20 fragments per packet, at any given time.</p> <p>Waits 5 seconds, to reassemble, if not received in 5 secs, received packets and future fragments will be dropped.</p>	<p>TCP intercept</p> <p>Config:</p>	<p>Client (any) TCP Server Inside 1.1.1.1</p> <p>access-list 101 permit tcp any 1.1.1.1 0.0.0.0</p> <p>ip tcp intercept list 101</p> <p>ip tcp intercept max-incomplete low 400 high 500</p> <p>ip tcp intercept one-minute low 30 high 60</p> <p>Watch Mode: is passive, terminates connection after timeout.</p> <p>Intercept Mode: actively intercepts SYNs, responds on behalf of the internal Server.</p> <p>Drop incoming SYN lower than 30 session If 60 open sessions are reached.</p> <p>show tcp intercept connections show tcp intercept statistics</p>
<p>IP DHCP Snooping:</p> 	<p>ip dhcp snooping ip dhcp snooping vlan X</p> <p>Puts ports in vlan X in untrusted mode</p> <p>DHCP Server: Interface X ip dhcp snooping trust</p> <p>Clients / untrusted ports: interface X ip dhcp snooping limit rate 100</p> <p>Errdisable recovery cause arp-insepection interval <seconds></p> <p>show ip dhcp snooping show ip dhcp snooping binding</p>	<p>Zone-Based Policy Firewall (ZFW)</p> <p>Comparison ZFW and CBAC:</p>	<p>Built to overcome limitations of CBAC: Traffic passing through the interface was subject to the same inspection policy.</p> <p>ZFW: Interfaces are assigned to zones, policy inspections are applied to traffic moving between zones. Uses class-maps via CPL!</p> 	<p>Unicast Reverse Path Forwarding</p> <p>uRPF</p> <p>Strict mode:</p>	<p>uRPF Strict Mode:</p> <p>Source Address must match the FIB Adjacency Info in the CEF table</p> <p>Packets sourced from 10.0.0.0/8 arriving at ser0/0 failing the uRPF check will be logged and dropped:</p> <p>access-list 101 deny ip 10.0.0.0/8 any log-input</p> <p>Packets sourced from 172.x arriving at ser0/0 failing the uRPF will be logged and forwarded:</p> <p>access-list 101 permit ip 172.x.x any log-input</p> <p>interface Serial0/0 ip verify unicast reverse-path 101</p>
<p>IP DHCP snooping config:</p>	<p>ip dhcp snooping ip dhcp snooping vlan X no ip dhcp snooping information option</p> <p>ip dhcp snooping vlan X ip dhcp snooping database flash://snoop.db ip dhcp snooping database write-delay X</p> <p>Write-delay: Default time from local entry to remote DB writing time.</p> <p>Option 82: ip dhcp snooping information option RemoteID + CircuitID within DHCP Discover message, seen on the servers trusted port outbound. (SW# within exec mode, not config!) ip dhcp snooping binding xxxx.xxxx.xxxx vlan X 1.2.3.4 int X expiry</p>	<p>Zone-Based Policy Firewall:</p> <p>Steps:</p>	<p>1 Define zones 2 Define zone-pairs 3 Define class-map(s) which identify traffic (traverses zone-pair) 4 Define policy-map / action to the traffic in class-map 5 Apply Policy-Map to zone-pair 6 Assign Interfaces to Zones</p>	<p>Unicast Reverse Path Forwarding</p> <p>uRPF</p> <p>Loose mode</p>	<p>Strict Mode: Resolves source IP address / source interface Loose Mode: Resolves source network</p> <p>In case Loose Mode is configured with a default route, Loose mode is useless!</p> <p>ip cef</p> <p>interface X ip verify unicast source reachable-via any</p>
<p>Control Plane Policing CoPP</p>	<p>class map CMAP-COPP match [ACL, protocol, ip prec, ip dhcp, vlan]</p> <p>policy-map PMAP-COPP class CMAP-COPP police <rate> conform-action <action> exceed-action <action></p> <p>control-plane service-policy [input, output] PMAP-COPP</p> <p>(IGP/BGP packets from an to Router)</p>	<p>Zone-Based Policy Firewall</p> <p>Config:</p>	<p>class-map type inspect match-any CLASS match protocol tcp match protocol udp</p> <p>policy-map type inspect MYPOL class type inspect CLASS inspect</p> <p>zone-pair security PAIR source private destination public service-policy type inspect MYPOL</p> <p>interface ethernet0 zone-member security private</p> <p>interface ethernet1 zone-member security public</p> <p>zone security public zone security private</p>	<p>Difference of</p> <p>uRPF Strict Mode</p> <p>uRPF Loose Mode:</p>	<p>uRPF Strict Mode: Routing entry for Source IP? Packet received on Interface where the IP should source?</p> <p>Interface X ip verify unicast source reachable-via rx</p> <p>uRPF loose mode: Is the source network in the routing table?</p> <p>Interface X ip verify unicast source reachable-via any</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts



Colin

<p>uRPF</p> <p>Strict Mode</p> <p>Loose Mode</p> <p>Show outputs:</p>	<p>Strict Mode:</p> <pre>show ip interface fastEthernet 0/1 include drops 5 verification drops 0 suppressed verification drops</pre> <p>Packet came from wrong interface, been dropped.</p> <hr/> <p>Loose Mode:</p> <pre>show ip interface fastEthernet 0/1 include drops 0 verification drops 7 suppressed verification drops</pre> <p>suppressed = Packet came in wrong interface, but there is a source network in the FIB.</p>	<p>TACACS+</p> <p>Response types:</p>	<pre>Accept: Successfully authenticated Reject: Incorrect credentials Error: communication error between NAS and Server Continue: Server is expecting additional info, user Prompt</pre> <p>TACACS encrypts entire body of the packet.</p> <p>TACACS uses TCP port 49</p>	<p>Configuring Login authentication using TACACS+:</p>	<p>Using local user database as fallback</p> <pre>aaa new-model aaa authentication login default group tacacs+ local</pre> <p>tacacs-server host 1.2.3.4</p> <p>tacacs-server key secret123</p>																																							
<p>Verify uRPF show commands:</p>	<pre>show ip interface <interface> ... IP verify source reachable-via RX, allow default, ACL 101 56 verification drops 192 suppressed verification drops</pre> <pre>show cef interface <interface> ... IP unicast RPF check is enabled</pre> <pre>show ip traffic ... 0 no route, 0 unicast RPF, 0 forced drop</pre>	<p>Implementing AAA</p> <p>Three types:</p>	<ol style="list-style-type: none"> self-contained AAA local security database Cisco Secure ACS Server Cisco Secure ACS Solutions Engine appliance 	<p>Authorization Services:</p>	<table border="1"> <thead> <tr> <th>Keyword</th> <th>Description of Use</th> </tr> </thead> <tbody> <tr> <td>Network Exec</td> <td>authorizes connections PPP SLIP ARAP attributes associated with EXEC term sessions</td> </tr> <tr> <td>Command</td> <td>Auths EXEC mode commands associated with privilege levels</td> </tr> <tr> <td>Auth-proxy Configuration</td> <td>applying specific sec policies per user downloads configs from AAA server</td> </tr> </tbody> </table>	Keyword	Description of Use	Network Exec	authorizes connections PPP SLIP ARAP attributes associated with EXEC term sessions	Command	Auths EXEC mode commands associated with privilege levels	Auth-proxy Configuration	applying specific sec policies per user downloads configs from AAA server																															
Keyword	Description of Use																																											
Network Exec	authorizes connections PPP SLIP ARAP attributes associated with EXEC term sessions																																											
Command	Auths EXEC mode commands associated with privilege levels																																											
Auth-proxy Configuration	applying specific sec policies per user downloads configs from AAA server																																											
<p>Compromising Private Vlan implementation:</p> <p>PVLAN</p>	<p>Attacker sends packet with Destination IP of the Victim while setting the Destination MAC to the Router!</p> <p>Router receives bridged packet, performs the routing lookup and forwards the packet to the Victim.</p> <pre>access-list 101 deny ip 10.0.0.0 0.0.0.255 10.0.0.0 0.0.0.255 access-list 101 permit ip any any Interface X ip access-group 101 in</pre> <p><i>mitigation</i></p>	<p>AAA overall configuration:</p>	<ol style="list-style-type: none"> enable aaa new-model configure IP of Radius or TACACS with a shared key define authentication service aaa authentication cmd apply authentication login authentication on line/int. Define authorization method list service Apply authorization method under line/interface Define accounting method list Apply accounting under line/interface <p>Named Method: applied to specific interfaces Default Method: globally / automatically applied to all interfaces if no other list is specified.</p>	<table border="1"> <thead> <tr> <th></th> <th>Radius</th> <th>TACACS+</th> </tr> </thead> <tbody> <tr> <td>Developer</td> <td>industry standard</td> <td>Cisco proprietary</td> </tr> <tr> <td>Transport Protocol</td> <td>UDP 1645,1646 UDP 1812,1813</td> <td>TCP 49</td> </tr> <tr> <td>AAA Support</td> <td>combines authen and authorization separate accounting</td> <td>uses AAA architecture separates the three services.</td> </tr> <tr> <td>Challenge Response</td> <td>Unidirectional single challenge</td> <td>Bidirectional Multiple challenges</td> </tr> <tr> <td>Protocol Support</td> <td>No NetBEUI</td> <td>Full support</td> </tr> <tr> <td>Security</td> <td>Encrypts only the password</td> <td>Encrypts the entire packet</td> </tr> </tbody> </table>		Radius	TACACS+	Developer	industry standard	Cisco proprietary	Transport Protocol	UDP 1645,1646 UDP 1812,1813	TCP 49	AAA Support	combines authen and authorization separate accounting	uses AAA architecture separates the three services.	Challenge Response	Unidirectional single challenge	Bidirectional Multiple challenges	Protocol Support	No NetBEUI	Full support	Security	Encrypts only the password	Encrypts the entire packet	<table border="1"> <thead> <tr> <th>Keyword</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>arap</td> <td>authentication list for Appletalk Remote Access Protocol ARAP</td> </tr> <tr> <td>login</td> <td>enable authentication for ASCII based logins such as Telnet / SSH</td> </tr> <tr> <td>enable</td> <td>Authentication list for enabling access to the router</td> </tr> <tr> <td>ppp</td> <td>Auth list for any PPP-based protocol such as ISDN, remote dial-in...</td> </tr> </tbody> </table>	Keyword	Description	arap	authentication list for Appletalk Remote Access Protocol ARAP	login	enable authentication for ASCII based logins such as Telnet / SSH	enable	Authentication list for enabling access to the router	ppp	Auth list for any PPP-based protocol such as ISDN, remote dial-in...								
	Radius	TACACS+																																										
Developer	industry standard	Cisco proprietary																																										
Transport Protocol	UDP 1645,1646 UDP 1812,1813	TCP 49																																										
AAA Support	combines authen and authorization separate accounting	uses AAA architecture separates the three services.																																										
Challenge Response	Unidirectional single challenge	Bidirectional Multiple challenges																																										
Protocol Support	No NetBEUI	Full support																																										
Security	Encrypts only the password	Encrypts the entire packet																																										
Keyword	Description																																											
arap	authentication list for Appletalk Remote Access Protocol ARAP																																											
login	enable authentication for ASCII based logins such as Telnet / SSH																																											
enable	Authentication list for enabling access to the router																																											
ppp	Auth list for any PPP-based protocol such as ISDN, remote dial-in...																																											
<p>802.1x Authentication process diagram:</p>	<p>802.1x Authentication process</p>	<p>AAA Authentication Methods</p>	<p>AAA authentication login</p> <p>Applied by login authentication</p> <table border="1"> <thead> <tr> <th>keyword</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>enable</td> <td>enable pass authentication</td> </tr> <tr> <td>group radius</td> <td>list all RADIUS servers</td> </tr> <tr> <td>group tacacs+</td> <td>list all TACACS server</td> </tr> <tr> <td>krb5</td> <td>Use Kerberos 5 authentication</td> </tr> <tr> <td>krb4-telnet</td> <td>Use Kerberos 5 using Telnet</td> </tr> <tr> <td>line</td> <td>use the line password for authentication</td> </tr> <tr> <td>local</td> <td>use local username database</td> </tr> <tr> <td>local-case</td> <td>uses case-sensitive local database</td> </tr> <tr> <td>none</td> <td>no authentication used</td> </tr> </tbody> </table>	keyword	Description	enable	enable pass authentication	group radius	list all RADIUS servers	group tacacs+	list all TACACS server	krb5	Use Kerberos 5 authentication	krb4-telnet	Use Kerberos 5 using Telnet	line	use the line password for authentication	local	use local username database	local-case	uses case-sensitive local database	none	no authentication used	<p>Authentication Services</p>	<table border="1"> <thead> <tr> <th>Keyword</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Network</td> <td>accounting info for all network related services such as PPP, SLIP, ARAP includes packet/byte count</td> </tr> <tr> <td>Connection</td> <td>all outbound connections from NAS telnet, local-area transport LAT PAD, rlogin</td> </tr> <tr> <td>Exec</td> <td>terminal sessions, username start, stop telephone number call origination</td> </tr> <tr> <td>System</td> <td>all system level events, reboots accounting is disabled or enabled.</td> </tr> <tr> <td>Command</td> <td>command accounting, exec commands which are executed on NAS</td> </tr> </tbody> </table>	Keyword	Description	Network	accounting info for all network related services such as PPP, SLIP, ARAP includes packet/byte count	Connection	all outbound connections from NAS telnet, local-area transport LAT PAD, rlogin	Exec	terminal sessions, username start, stop telephone number call origination	System	all system level events, reboots accounting is disabled or enabled.	Command	command accounting, exec commands which are executed on NAS							
keyword	Description																																											
enable	enable pass authentication																																											
group radius	list all RADIUS servers																																											
group tacacs+	list all TACACS server																																											
krb5	Use Kerberos 5 authentication																																											
krb4-telnet	Use Kerberos 5 using Telnet																																											
line	use the line password for authentication																																											
local	use local username database																																											
local-case	uses case-sensitive local database																																											
none	no authentication used																																											
Keyword	Description																																											
Network	accounting info for all network related services such as PPP, SLIP, ARAP includes packet/byte count																																											
Connection	all outbound connections from NAS telnet, local-area transport LAT PAD, rlogin																																											
Exec	terminal sessions, username start, stop telephone number call origination																																											
System	all system level events, reboots accounting is disabled or enabled.																																											
Command	command accounting, exec commands which are executed on NAS																																											
<p>TACACS+ packet header:</p>	<p>TACACS+ packet header</p> <table border="1"> <thead> <tr> <th>Major</th> <th>Minor</th> <th>Packet Type</th> <th>Seq-Num</th> <th>Flags</th> </tr> </thead> <tbody> <tr> <td colspan="5">Session ID (4 bytes)</td> </tr> <tr> <td colspan="5">Length (4 bytes)</td> </tr> </tbody> </table> <p>Major Version: TACACS+ version number Minor Version: maintaining backward compatibility</p> <p>Packet Type: TAC_PLUS_AUTHEN = 0x01 authentication TAC_PLUS_AUTHOR: = 0x02 authorization TAC_PLUS_ACCT: = 0x03 Accounting</p> <p>Flags: Signify if packet is encrypted</p>	Major	Minor	Packet Type	Seq-Num	Flags	Session ID (4 bytes)					Length (4 bytes)					<p>AAA Authorization Methods:</p>	<p>AAA authorization login</p> <p>Applied by authorization xxxx</p> <table border="1"> <thead> <tr> <th>Keyword</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>group radius</td> <td>NAS requests from server, defines rights for users by associated attribute-value pairs</td> </tr> <tr> <td>group tacacs+</td> <td>NAS requests, user rights utilizing attribute-value pairs</td> </tr> <tr> <td>if-authenticated</td> <td>allowed access if successfully authenticated.</td> </tr> <tr> <td>local</td> <td>using local user database for authentication</td> </tr> <tr> <td>none</td> <td>Does not request authorization.</td> </tr> </tbody> </table>	Keyword	Description	group radius	NAS requests from server, defines rights for users by associated attribute-value pairs	group tacacs+	NAS requests, user rights utilizing attribute-value pairs	if-authenticated	allowed access if successfully authenticated.	local	using local user database for authentication	none	Does not request authorization.	<p>Accounting Services</p>	<table border="1"> <thead> <tr> <th>Keyword</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Network</td> <td>accounting info for all network related services such as PPP, SLIP, ARAP includes packet/byte count</td> </tr> <tr> <td>Connection</td> <td>all outbound connections from NAS telnet, local-area transport LAT PAD, rlogin</td> </tr> <tr> <td>Exec</td> <td>terminal sessions, username start, stop telephone number call origination</td> </tr> <tr> <td>System</td> <td>all system level events, reboots accounting is disabled or enabled.</td> </tr> <tr> <td>Command</td> <td>command accounting, exec commands which are executed on NAS</td> </tr> </tbody> </table>	Keyword	Description	Network	accounting info for all network related services such as PPP, SLIP, ARAP includes packet/byte count	Connection	all outbound connections from NAS telnet, local-area transport LAT PAD, rlogin	Exec	terminal sessions, username start, stop telephone number call origination	System	all system level events, reboots accounting is disabled or enabled.	Command	command accounting, exec commands which are executed on NAS
Major	Minor	Packet Type	Seq-Num	Flags																																								
Session ID (4 bytes)																																												
Length (4 bytes)																																												
Keyword	Description																																											
group radius	NAS requests from server, defines rights for users by associated attribute-value pairs																																											
group tacacs+	NAS requests, user rights utilizing attribute-value pairs																																											
if-authenticated	allowed access if successfully authenticated.																																											
local	using local user database for authentication																																											
none	Does not request authorization.																																											
Keyword	Description																																											
Network	accounting info for all network related services such as PPP, SLIP, ARAP includes packet/byte count																																											
Connection	all outbound connections from NAS telnet, local-area transport LAT PAD, rlogin																																											
Exec	terminal sessions, username start, stop telephone number call origination																																											
System	all system level events, reboots accounting is disabled or enabled.																																											
Command	command accounting, exec commands which are executed on NAS																																											
<p>TACACS+ Communication Authentication:</p>	<p>TACACS+ Communication Authentication:</p>	<p>Configuring AAA RADIUS Server Groups</p> <p>"globally and privately"</p>	<pre>aaa group server radius TEST-1 server 1.2.3.4 radius-server key super-secret123</pre> <hr/> <pre>aaa group server radius TEST-2 server-private 2.2.2.2 key secret123</pre> <hr/> <pre>aaa authentication login default group TEST-1 aaa authentication ppp default group TEST-2</pre>	<p>PPP authentication</p> <p>if-needed:</p> <p>if-authenticated:</p>	<pre>aaa new-model aaa authentication ppp default if-needed group radius aaa authorization network default group radius if-authenticated</pre> <p>if-needed: If the user has already authenticated by going through the ASCII login procedure, PPP authentication is not necessary and skipped.</p> <p>if-authenticated: Indicates that users can be given access to requested services only if they have been authenticated first.</p>																																							

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<p>AAA login authentication</p> <p>Password retry / lockout:</p>	<p>Aaa local authentication attempts max-fail 3</p> <p>Show aaa local user locked</p> <p>Clear aaa local user lockout</p> <p>Users with privilege level 15 can not be locked out with ASCII based logins.</p>	<p>IKE Phase 2:</p> <p>Quick mode:</p>	<p>IKE phase 2 protects user data and establishes SA for Ipsec. <i>3 messages</i></p> <p>IKE phase 2 negotiates:</p> <ul style="list-style-type: none"> - Protection suite (using ESP, AH) - Algorithms in the protection suite (DES,3DES,AES,SHA) - IP traffic that is being protected, proxy identities - Optional keying material for negotiated protocols. <p>- At the end of phase 2 negotiations, two unidirectional IPsec SAs are established. One for sending, and one for receiving encrypted traffic</p> <p>- Only one mode: Quick mode.</p> <p>- Multiple phase 2 SA can be established over the same phase 1 SA.</p>	<p>Easy VPN</p> <p>Head-End Router:</p> <p>Part 1</p>	<pre>aaa new-model aaa authentication login vpnauthen local aaa authorization network vpnauthen local username cisco password cisco1 crypto isakmp policy 1 encryption 3des authentication pre-share group 2 crypto isakmp client configuration address-pool local POOL crypto isakmp xauth timeout 60 crypto isakmp client configuration group EASYvpn key cisco2 dns 1.2.3.4 domain cisco.com pool POOL access-list 101</pre> <p><i>aaa for Xauth</i> <i>Xauth credentials</i> <i>Phase 1 params</i> <i>Easy VPN group Parameters & split tunnel</i></p>																																																						
<p>Types of VPN technologies:</p>	<p>Secure VPN (cryptographic VPN)</p> <ul style="list-style-type: none"> - IPsec - L2TP over IPsec - SSL encryption <p>Trusted VPN (non-cryptographic VPN)</p> <ul style="list-style-type: none"> - MPLS VPN (Layer 3) - BGP VPN (Layer 3) - Multicast VPN (Layer 3) - Transport of Layer 2 frames over MPLS (AtoM) (Layer 2) - Virtual Private LAN Services VPLS (Layer 2) <p>Hybrid VPN combination of trusted and secure tunnel, tunnel within tunnel setup</p>	<p>IKEv2 features:</p>	<ul style="list-style-type: none"> -IKE dead peer detection / Initial contact - NAT traversal support - identities are always protected - Certs can be referenced through URL + hash to avoid fragmentation - EAP (MD-5, OTP, GTC) support - Remote address acquisition, New Config Payload CP - Two kinds of SA: IKE_SA used by IKE self and CHILD_SA used by IPsec - Four message types: <ul style="list-style-type: none"> IKE_SA_INIT IKE_AUTH CREATE_CHILD_SA INFORMATIONAL 	<p>Easy VPN</p> <p>Head-End Router:</p> <p>Part 2</p>	<pre>crypto ipsec transform-set TRANS esp-3des esp-sha-hmac crypto dynamic-map DYNMAP 10 set transform-set TRANS reverse-route crypto map DYNMAP client authentication list vpnauthen crypto map DYNMAP isakmp authorization list vpnauthen crypto map DYNMAP client configuration address respond crypto map CISCO ipsec-isakmp dynamic DYNMAP interface X (outside) crypto map CISCO ip local pool POOL 10.0.0.1 10.0.0.100 access-list 101 permit 1.1.1.0 0.0.0.255 any</pre> <p><i>Phase 2 params</i></p>																																																						
<p>IPSec Modes:</p> <p>Diagram:</p>	<p>Transport Mode Protects the payload of the original datagram.</p> <p>Tunnel Mode Encrypts traffic, encapsulates the entire datagram.</p>	<table border="1"> <thead> <tr> <th></th> <th>IKEv1</th> <th>IKEv2</th> </tr> </thead> <tbody> <tr> <td>UDP Port</td> <td></td> <td></td> </tr> <tr> <td>Phases</td> <td></td> <td></td> </tr> <tr> <td>Keepalives</td> <td></td> <td></td> </tr> <tr> <td>Identity Hiding</td> <td></td> <td></td> </tr> <tr> <td>UDP/NAT</td> <td></td> <td></td> </tr> <tr> <td>SA Negotiation</td> <td></td> <td></td> </tr> <tr> <td>Number of MSGs</td> <td></td> <td></td> </tr> <tr> <td>EAP/CP</td> <td></td> <td></td> </tr> </tbody> </table>		IKEv1	IKEv2	UDP Port			Phases			Keepalives			Identity Hiding			UDP/NAT			SA Negotiation			Number of MSGs			EAP/CP			<table border="1"> <thead> <tr> <th></th> <th>IKEv1</th> <th>IKEv2</th> </tr> </thead> <tbody> <tr> <td>UDP Port</td> <td>500</td> <td>500,4500</td> </tr> <tr> <td>Phases</td> <td>Phase 1 6 / 3 messages Phase 2 3 messages</td> <td>Phase 1 4 messages Phase 2 2 messages</td> </tr> <tr> <td>Keepalives</td> <td>No</td> <td>Yes</td> </tr> <tr> <td>Identity Hiding</td> <td>in main-mode</td> <td>Yes</td> </tr> <tr> <td>UDP/NAT</td> <td>no</td> <td>Yes</td> </tr> <tr> <td>SA Negotiation</td> <td>Responder selects initiators proposal</td> <td>Same as IKEv1, simpler</td> </tr> <tr> <td>Number of MSGs</td> <td>6-9</td> <td>4-8</td> </tr> <tr> <td>EAP/CP</td> <td>No</td> <td>Yes</td> </tr> </tbody> </table>		IKEv1	IKEv2	UDP Port	500	500,4500	Phases	Phase 1 6 / 3 messages Phase 2 3 messages	Phase 1 4 messages Phase 2 2 messages	Keepalives	No	Yes	Identity Hiding	in main-mode	Yes	UDP/NAT	no	Yes	SA Negotiation	Responder selects initiators proposal	Same as IKEv1, simpler	Number of MSGs	6-9	4-8	EAP/CP	No	Yes	<p>Easy VPN</p> <p>Client Router config:</p> <p>(Network extension mode)</p>	<pre>crypto ipsec client ezvpn MYVPN connect auto group easyvpn key cisco123 mode network-extension peer 2.2.2.2 username cisco password cisco xauth userid mode local interface gi0/0 description outside crypto ipsec client ezvpn MYVPN outside interface gi0/1 description inside crypto ipsec client ezvpn MYVPN inside</pre>
	IKEv1	IKEv2																																																									
UDP Port																																																											
Phases																																																											
Keepalives																																																											
Identity Hiding																																																											
UDP/NAT																																																											
SA Negotiation																																																											
Number of MSGs																																																											
EAP/CP																																																											
	IKEv1	IKEv2																																																									
UDP Port	500	500,4500																																																									
Phases	Phase 1 6 / 3 messages Phase 2 3 messages	Phase 1 4 messages Phase 2 2 messages																																																									
Keepalives	No	Yes																																																									
Identity Hiding	in main-mode	Yes																																																									
UDP/NAT	no	Yes																																																									
SA Negotiation	Responder selects initiators proposal	Same as IKEv1, simpler																																																									
Number of MSGs	6-9	4-8																																																									
EAP/CP	No	Yes																																																									
<p>IPSec Protocol Headers:</p> <p>ESP:</p> <p>AH:</p>	<p>Encapsulating Security Payload (ESP):</p> <ul style="list-style-type: none"> - UDP port 50 - offers anti-replay protection - does not provide protection to the outer IP header <p>Authentication Header (AH):</p> <ul style="list-style-type: none"> - UDP port 51 - offers anti-replay protection - AH provides protection to the IP header. - No confidentiality protection. <table border="1"> <thead> <tr> <th>Mode</th> <th>Transport</th> <th>Tunnel</th> </tr> </thead> <tbody> <tr> <td>AH</td> <td>IP AH Data</td> <td>IP AH IP Data</td> </tr> <tr> <td>ESP</td> <td>IP ESP Data ESP-T</td> <td>IP ESP IP Data ESP-T</td> </tr> <tr> <td>AH-ESP</td> <td>IP AH ESP Data ESP-T</td> <td>IP AH ESP IP Data ESP-T</td> </tr> </tbody> </table>	Mode	Transport	Tunnel	AH	IP AH Data	IP AH IP Data	ESP	IP ESP Data ESP-T	IP ESP IP Data ESP-T	AH-ESP	IP AH ESP Data ESP-T	IP AH ESP IP Data ESP-T	<p>IPSec VTI Configuration:</p>	<pre>crypto isakmp policy 10 encryption 3des authentication pre-share group 2 crypto isakmp key cisco address 2.2.2.2 255.255.255.255 crypto transform-set MYTRANS esp-aes esp-sha mode transport crypto ipsec profile vti_profile set transform-set MYTRANS interface tunnel 0 ip address 172.16.1.1 255.255.255.252 tunnel source serial0 tunnel destination 3.3.3.3 tunnel mode ipsec ipv4 tunnel protection ipsec profile vti_profile ip route X.X.X.X 255.255.255.0 Tunnel0</pre>	<p>Cisco Easy VPN with Dynamic VTI</p> <p>Server-Router</p> <p>Part 1</p>	<pre>aaa new-model aaa authentication login default local aaa authorization network default local ip cef username cisco privilege 15 password 0 cisco1 policy-map TEST class class-default shape average 128000 crypto isakmp policy 10 encryption 3des authentication pre-share group 2 crypto isakmp key cisco12 address 0.0.0.0 0.0.0.0 crypto isakmp keepalive 10</pre> <p><i>aaa for Xauth</i> <i>Phase 1</i></p>																																										
Mode	Transport	Tunnel																																																									
AH	IP AH Data	IP AH IP Data																																																									
ESP	IP ESP Data ESP-T	IP ESP IP Data ESP-T																																																									
AH-ESP	IP AH ESP Data ESP-T	IP AH ESP IP Data ESP-T																																																									
<p>IPSEC IKE phase 1:</p> <p>Main mode:</p> <p>Aggressive mode:</p>	<p>Modes of IKE</p> <p>Main mode: 6 messages Default method</p> <p>Aggressive mode: 3 messages Faster but less secure, does not provide identity protection.</p>	<p>IPSec VPN Site-to-Site configuration:</p>	<pre>crypto keyring SPOKES pre-shared-key address 1.2.3.4 0.0.0.0 key Secret1 crypto isakmp policy 10 encryption 3des authentication pre-share group 2 crypto isakmp profile ISAKMP-PROF keyring SPOKES match identity address 1.2.3.4 crypto map CISCO 10 ipsec-isakmp set peer 1.2.3.4 set transform-set MYTRANS set isakmp-profile ISAKMP-PROF match address 101 Interface Gi0/1 crypto map CISCO access-list 101 permit 4.0.0.0 0.0.0.0 5.0.0.0 0.0.0.255 ip route 5.0.0.0 255.0.0.0 1.2.3.4</pre>	<p>Cisco Easy VPN with Dynamic VTI</p> <p>Server-Router</p> <p>Part 2</p>	<pre>crypto isakmp client configuration group cisco key cisco dns 1.2.3.4 pool POOL access-list 101 crypto isakmp profile ISAKMP match identity group cisco isakmp authorization list default client configuration address respond virtual-template 1 crypto ipsec transform-set TRANS esp-3des esp-sha-hmac crypto ipsec profile MYIPSEC set transform-set TRANS set isakmp-profile ISAKMP</pre> <p><i>Easy VPN group</i> <i>ISAKMP and bind parameters</i> <i>IPSec bind parameters</i></p>																																																						
<p>Troubleshooting ISAKMP or Phase 1 VPN connections:</p>	<p>Verify ISAKMP parameters match exactly.</p> <p>Verify pre-shared-keys match exactly.</p> <p>Check that each side has a route to the peer address that you are trying to form a tunnel with.</p> <p>Verify ISAKMP is enabled on the outside interfaces.</p> <p>Is ESP traffic permitted in through the outside interface?</p> <p>Is UDP port 500 open on the outside ACL?</p> <p>Some situations require that UDP port 4500 is open for the outside.</p>	<p>Remote Access IPsec VPN Overview:</p>	<ul style="list-style-type: none"> - RA IPsec VPN - RA Secure Sockets Layer SSL VPN - Cisco Easy VPN - Dynamic VTI (DVTI) <p>Easy VPN Solution: Have centralized security policies at the head-end VPN server, pushed to the remote sites upon connection.</p> <p>Easy VPN Software Client (laptop) Easy VPN Hardware Client (ASA, PIX for LAN-to-LAN setups)</p> <ul style="list-style-type: none"> - Client Mode / PAT mode: single source ip fully routable over tunnel - Network Extension Mode: able to request IP address via mode configuration. - Network Extension Plus Mode: 	<p>Cisco Easy VPN with Dynamic VTI</p> <p>Server-Router</p> <p>Part 3</p>	<pre>int gi0/0 description outside ip address x.x.x.x 255.255.255.0 int gi0/1 description inside ip address y.y.y 255.255.255.0 interface virtual-template 1 type tunnel ip unnumbered Gi0/0 tunnel source gi0/0 tunnel mode ipsec ipv4 tunnel protection ipsec profile MYIPSEC service-policy output TEST ip local pool POOL 172.16.1.1 172.16.1.100 ip route 0.0.0.0 0.0.0.0 gi0/0 access-list 101 permit 1.1.1.0 0.0.0.255 any</pre> <p><i>For DVTI cloning & apply IPsec profile</i> <i>IP pool for VPN users</i></p>																																																						

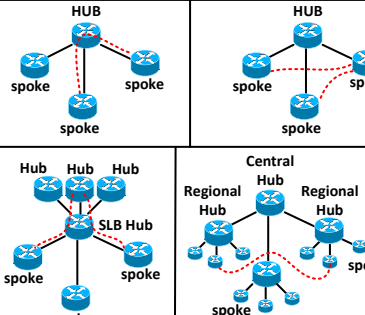
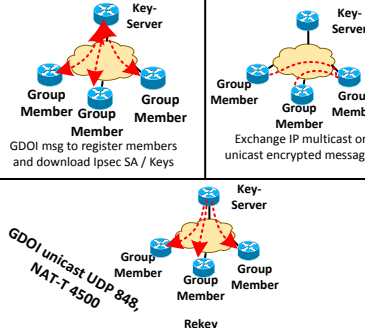
Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<p>Cisco Easy VPN with Dynamic VTI Client-Router Part 1</p>	<pre>ip cef Username cisco privilege 15 password 0 cisco1 Policy-map TEST Class class-default Shape average 128000 Crypto isakmp policy 10 Encr 3des Authentication pre-share Group 2 Crypto isakmp key cisco2 address 0.0.0.0 0.0.0.0 Crypto isakmp keepalive 10 Crypto ipsec client ezvpn MYVPN Connect auto Group cisco key cisco3 Peer 2.2.2.2</pre>	<p>DMVPN Single Hub DMVPN HUB config:</p>	<pre>crypto isakmp policy 10 encr 3des authentication pre-share group 2 crypto isakmp key SECRET address 0.0.0.0 0.0.0.0 crypto ipsec transform-set TRANS esp-3des esp-sha-hmac crypto ipsec profile VPNPROF set transform-set TRANS interface Tunnel0 ip address 172.16.1.11 ip mtu 1400 ip nhrp authentication cisco ip nhrp map multicast dynamic ip nhrp network-id 123 ip nhrp holdtime 360 delay 1000 ip tcp adjust-mss 1360 tunnel source gi0/0 tunnel mode gre multipoint tunnel protection ipsec profile VPNPROF</pre>	<p>AAA Local Command Authorization</p>	<pre>username ADMIN privilege 7 password 0 CISCO privilege exec level 7 configure terminal privilege exec level 7 undebbug all privilege exec level 7 show running-config privilege exec level 7 debug ip rip privilege configure level 7 interface privilege interface level 7 shutdown privilege interface level 7 no shutdown privilege interface all level 7 ip</pre>				
<p>Cisco Easy VPN with Dynamic VTI Client-Router Part 2</p>	<pre>crypto ipsec client ezvpn MYVPN Username cisco password ciscoX Xauth userid mode local Local-address Gi0/0 Mode client Interface X Ip virtual-assembly Crypto ipsec client ezvpn MYVPN Interface Y Ip virtual-assembly Crypto ipsec client ezvpn MYPN inside Interface Virtual-Template1 type tunnel No ip address Ip virtual-reassembly Tunnel mode ipsec ipv4 Service-policy output TEST</pre>	<p>DMVPN Single Hub DMVPN Spoke config:</p>	<pre>crypto isakmp policy 10 encr 3des authentication pre-share group 2 crypto isakmp key SECRET address 0.0.0.0 0.0.0.0 crypto ipsec transform-set TRANS interface Tunnel0 ip mtu 1400 ip nhrp authentication cisco ip nhrp map 172.16.1.11 11.11.11.11 ip nhrp network-id 123 ip nhrp holdtime 360 ip nhrp nhs 12.16.1.11 delay 1000 ip tcp adjust-mss 1360 tunnel source fa0/0 tunnel destination 11.11.11.11 tunnel protection ipsec profile VPNPROF</pre>	<p>Extended ACLs which permits UDP traceroute:</p>	<p>traceroute UDP probes</p> <p>permit udp any any range 33434 33474</p>				
<p>DMVPN Diagrams</p> <table border="1" data-bbox="118 819 341 966"> <tr> <td>Phase 1</td> <td>Phase 2</td> </tr> <tr> <td>Phase 1 Server load balancing SLB</td> <td>Phase 3</td> </tr> </table>	Phase 1	Phase 2	Phase 1 Server load balancing SLB	Phase 3		<p>SSL VPN Access modes:</p>	<p>Clientless Mode (Layer 7):</p> <ul style="list-style-type: none"> - Browser-based - Web-enabled applications (Outlook webaccess) <p>Thin-Client Mode (Layer 7):</p> <ul style="list-style-type: none"> - Delivered via Java Applet - TCP-forwarding - Extension of application support - Telnet, pop3, SMTP, SSH, static port based applications <p>Thick-Client Mode (Layer 3):</p> <ul style="list-style-type: none"> - Traditional SSL VPN Client through Java, ActiveX - AnyConnect VPN Client software - Support of all IP-based applications 	<p>Filtering Fragmented Packets</p>	<p>Non-fragmented packets or initial fragments have a fragment offset of zero and hold upper level protocol information.</p> <p>Non-initial fragments have a non-zero fragment offset, hold no upper layer info, but would still slip through an ACL like the following, if the IP addresses match:</p> <pre>permit tcp host 1.1.1.1 host 2.2.2.2 eq 80 ip access-list extended NO_FRAGMENTS deny ip any any fragments permit tcp any any eq 80</pre>
Phase 1	Phase 2								
Phase 1 Server load balancing SLB	Phase 3								
<p>DMVPN deployment topologies</p>	<p>Hub and Spoke Designs</p> <ul style="list-style-type: none"> - Single Hub Single DMVPN (SHSD) - Dual Hub Dual DMVPN (DHDD) - Server Load Balancing (SLB) <p>Dynamic Mesh Designs</p> <ul style="list-style-type: none"> - Dual Hub Single DMVPN (DHSD) - Multihub Single DMVPN (MHSD) - Hierarchical (Tree-Based) 	<p>AAA Authentication Lists</p>	<pre>aaa new-model aaa authentication login CONSOLE local aaa authentication login VTY group tacacs+ line aaa authentication enable default group tacacs+ enable aaa authentication password-prompt "Please Enter Your Password:" aaa authentication username-prompt "Please Enter Your ID:" aaa authentication banner # This system requires you to identify yourself. # aaa authentication fail-message # Authentication Failed, Sorry. # tacacs-server host 155.1.146.100 tacacs-server directed-request tacacs-server key CISCO</pre>	<p>Lock N Key Dynamic Access-lists:</p>	<ul style="list-style-type: none"> - only one dynamic entry per access-list - using dynamic ACLs with AAA enabled, make sure you are using local AAA. <pre>username ENABLE password CISCO username ENABLE autocommand access-enable host timeout 5 ip access-list extended DYNAMIC-1 dynamic PERMIT permit icmp any any deny ip any any interface FastEthernet0/0 ip access-group DYNAMIC-1 in line vty 4 rotary 1 password CISCO login autocommand access-enable timeout 5 clear access-template DYNAMIC-1 PERMIT any 10.0.0.0 0.0.0.255</pre> <p><i>access-enable command unlocks the dynamic entries.</i></p>				
<p>DMVPN show commands:</p>	<pre>show ip nhrp reveals NBMA address show crypto socket Local/Remote network show crypto ipsec sa remote crypto endpoints show crypto map peer, ACLs</pre>	<p>Debugging AAA for a telnet line:</p>	<pre>debug aaa authentication R5#telnet 150.1.1.1 Trying 150.1.1.1 ... Open AAA/AUTHEN/START (649338587): using "default" list AAA/AUTHEN/START (649338587): Method=tacacs+ (tacacs+) AAA/AUTHEN(649338587): Status=ERROR AAA/AUTHEN/START (649338587): Method=ENABLE AAA/AUTHEN(649338587): Status=GETPASS AAA/AUTHEN/CONT (649338587): continue_login (user=(undef)) AAA/AUTHEN(649338587): Status=GETPASS AAA/AUTHEN/CONT (649338587): Method=ENABLE AAA/AUTHEN(649338587): Status=PASS</pre>	<p>Lock N Key Additional features:</p>	<pre>line vty 4 autocommand access-enable timeout 5 autocommand access-enable host timeout 5 access-enable host replaces the source IP specification (any) in the access-list entry with the IP of the authenticated user access-list dynamic-extended re-login to the router to extend the absolute timeout manual clear feature, with numbered extended ACL only clear access-template <ACL-NUMBER> <DYNAMIC-ENTRY-NAME> <SRC-IP> <SRC-MASK> <DST-IP> <DST-MASK></pre>				
<p>Get VPN Diagram:</p> <p>GDOI msg / traffic / re-keying</p>		<p>AAA Exec Authorization example</p>	<pre>aaa authorization console aaa authorization exec default none aaa authorization exec CONSOLE group tacacs+ local aaa authorization exec VTY group tacacs+ if-authenticated username ADMIN privilege 7 password 0 CISCO line con 0 authorization exec CONSOLE line vty 0 4 privilege level 15 password cisco authorization exec VTY login authentication VTY</pre>	<p>Traffic Filtering with Policy Based Routing</p>	<p>Define Traffic:</p> <pre>ip access-list extended ICMP permit icmp any any route-map DROP match ip address ICMP match interface gi 0/0.67 match length 100 100 (min 100 bytes, max 100 bytes) set interface Null0 (send to Null0, drop) interface gi 0/0.146 ip policy route-map DROP interface Null0 no ip unreachable Drop packet, but don't send ip unreachable</pre>				


```
username TELNET password 0 CISCO
username TELNET autocommand access-enable timeout 5
or
username TELNET autocommand access-enable host timeout 5
!
ip access-list extended DYNAMIC1
dynamic PERMIT_TELNET permit tcp any any eq telnet
deny tcp any host 191.1.27.7 eq telnet
deny tcp any host 191.1.7.7 eq telnet
deny tcp any host 191.1.77.7 eq telnet
deny tcp any host 191.1.177.7 eq telnet
deny tcp any host 150.1.7.7 eq telnet
permit ip any any
!
interface Serial0/0
ip access-group DYNAMIC1 in
!
interface Serial0/1
ip access-group DYNAMIC1 in
!
line vty 0 4
login local
```

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

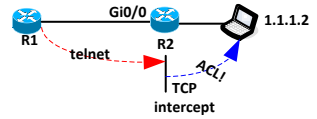
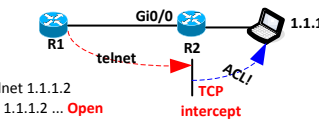
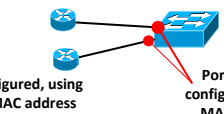

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin


<h3>NBAR for Content-Based Filtering</h3>	<pre>class-map match-all EXTENSION match protocol http url "*.bin .exe .com" policy-map DROP class EXTENSION drop interface Gi 0/0.146 service-policy output DROP</pre> <p><i>Tricky to find, while troubleshooting</i></p>	<h3>CBAC Filtering</h3> <p>Show commands:</p>	<pre>show ip inspect config show ip inspect interfaces show ip inspect sessions show ip inspect all show ip port-map inc http</pre> <p>TFTP is the perfect protocol for testing, due to it using dynamic ports</p>	<table border="1"> <thead> <tr> <th>Ethertype</th> <th>Protocol</th> <th>Ethertype</th> <th>Protocol</th> </tr> </thead> <tbody> <tr><td>IPv4</td><td>0x800</td><td>IPv4</td><td>0x800</td></tr> <tr><td>ARP</td><td>0x0806</td><td>ARP</td><td>0x0806</td></tr> <tr><td>Reverse-ARP</td><td>0x0835</td><td>Reverse-ARP</td><td>0x0835</td></tr> <tr><td>802.1Q tagged frame</td><td>0x8100</td><td>802.1Q tagged frame</td><td>0x8100</td></tr> <tr><td>IPX</td><td>0x8137</td><td>IPX</td><td>0x8137</td></tr> <tr><td>IPX</td><td>0x8138</td><td>IPX</td><td>0x8138</td></tr> <tr><td>?</td><td>IPv6</td><td>0x86DD</td><td>IPv6</td></tr> <tr><td></td><td>MPLS unicast</td><td>0x8847</td><td>MPLS unicast</td></tr> <tr><td></td><td>MPLS multicast</td><td>0x8848</td><td>MPLS multicast</td></tr> <tr><td></td><td>PPPoE Discovery stage</td><td>0x8863</td><td>PPPoE Discovery stage</td></tr> <tr><td></td><td>PPPoE Session Stage</td><td>0x8864</td><td>PPPoE Session Stage</td></tr> <tr><td></td><td>LLDP</td><td>0x88CC</td><td>LLDP</td></tr> <tr><td></td><td>EAP over LAN</td><td>0x888E</td><td>EAP over LAN</td></tr> <tr><td></td><td>TRILL</td><td>0x22F3</td><td>TRILL</td></tr> </tbody> </table>	Ethertype	Protocol	Ethertype	Protocol	IPv4	0x800	IPv4	0x800	ARP	0x0806	ARP	0x0806	Reverse-ARP	0x0835	Reverse-ARP	0x0835	802.1Q tagged frame	0x8100	802.1Q tagged frame	0x8100	IPX	0x8137	IPX	0x8137	IPX	0x8138	IPX	0x8138	?	IPv6	0x86DD	IPv6		MPLS unicast	0x8847	MPLS unicast		MPLS multicast	0x8848	MPLS multicast		PPPoE Discovery stage	0x8863	PPPoE Discovery stage		PPPoE Session Stage	0x8864	PPPoE Session Stage		LLDP	0x88CC	LLDP		EAP over LAN	0x888E	EAP over LAN		TRILL	0x22F3	TRILL	
Ethertype	Protocol	Ethertype	Protocol																																																														
IPv4	0x800	IPv4	0x800																																																														
ARP	0x0806	ARP	0x0806																																																														
Reverse-ARP	0x0835	Reverse-ARP	0x0835																																																														
802.1Q tagged frame	0x8100	802.1Q tagged frame	0x8100																																																														
IPX	0x8137	IPX	0x8137																																																														
IPX	0x8138	IPX	0x8138																																																														
?	IPv6	0x86DD	IPv6																																																														
	MPLS unicast	0x8847	MPLS unicast																																																														
	MPLS multicast	0x8848	MPLS multicast																																																														
	PPPoE Discovery stage	0x8863	PPPoE Discovery stage																																																														
	PPPoE Session Stage	0x8864	PPPoE Session Stage																																																														
	LLDP	0x88CC	LLDP																																																														
	EAP over LAN	0x888E	EAP over LAN																																																														
	TRILL	0x22F3	TRILL																																																														
<h3>debug ip tcp intercept:</h3>	<pre>INTERCEPT: new connection (155.1.67.7:11008 SYN -> 155.1.146.4:23) INTERCEPT(*): (155.1.67.7:11008 <- ACK+SYN 155.1.146.4:23) INTERCEPT: 1st half of connection is established (155.1.67.7:11008 ACK -> 155.1.146.4:23) INTERCEPT(*): (155.1.67.7:11008 SYN -> 155.1.146.4:23) INTERCEPT: 2nd half of connection established (155.1.67.7:11008 <- ACK+SYN 155.1.146.4:23)</pre>	<h3>Advanced CBAC Features</h3>	<p>Inspects router generated traffic, incl. RIP ip inspect name INSPECT udp router-traffic</p> <p>Use FTP inspect on port 80 for host 1.1.1.1 access-list 55 permit 1.1.1.1 ip port-map ftp port 80 list 55</p>	<h3>Show port-security interface X:</h3>	<pre>SW3#show port-security interface fa0/6 Port Security : Enabled Port Status : Secure-up Violation Mode : Protect Aging Time : 10 mins Aging Type : Inactivity SecureStatic Address Aging : Disabled Maximum MAC Addresses : 2 Total MAC Addresses : 2 Configured MAC Addresses : 0 Sticky MAC Addresses : 0 Last Source Address:Vlan : 0011.200a.f000:146 Security Violation Count : 0</pre> <p>Vlan ID</p>																																																												
<h3>Tricky TCP intercept scenario:</h3> 	 <p>R1#telnet 1.1.1.2 Trying 1.1.1.2 ... Open</p> <p>R1 is receiving a three Way handshake even Tcp established is Denied by the inbound ACL TCP intercept Answers in this case!</p> <p>R2# ip access-list extended NO_ACK deny tcp any any established permit ip any any</p> <p>interface Gi0/0 ip access-group NO_ACK in</p>	<h3>CBAC TCP/UDP Intercept Feature:</h3> <ul style="list-style-type: none"> - Maximum of 100 half-open connections, keep dropping the sessions until their number reaches 80. - Limit the new connections rate to 30 per minute, damp it to 15 if exceeded. - Limit half-open tcp per server to 10, block for 60 sec if exceeded. 	<pre>ip inspect max-incomplete low 80 ip inspect max-incomplete high 100 ip inspect one-minute low 15 ip inspect one-minute high 30 ip inspect tcp max-incomplete host 10 block-time 1 ip inspect tcp synwait-time 10</pre> <p>ip inspect name DEFEND tcp ip inspect name DEFEND udp</p> <p>interface FastEthernet 0/0 ip inspect DEFEND out</p>	<h3>HSRP and Port-Security</h3>	 <p>HSRP configured, using the BIA MAC address instead of an additional HSRP virtual MAC addr</p> <p>Port-security configured, only 1 MAC allowed</p> <pre>interface FastEthernet 0/0 ip address 155.1.146.4 255.255.255.0 standby 2 ip standby use-bia</pre> <p>Arp cache of clients need to be flushed!</p>																																																												
<h3>TCP intercept watch mode:</h3>	<pre>ip tcp intercept mode watch ip tcp intercept watch-timeout <TIMEOUT></pre> <p>router forcefully sends RST if the client/server connection did not establish after <timeout> seconds. also after 20 seconds without any interaction.</p>	<h3>VLAN Filters for Non-IP Traffic</h3>	<p>Allows STP, ARP on the Vlan, every thing else is denied by implicit deny any any:</p> <pre>mac access-list extended ALLOWED_L2_TRAFFIC permit any any Isap 0x4242 0x0 permit any any 0x010B 0x0 permit any any 0x806 0x0</pre> <p>The list is then applied using the vlan filter</p> <pre>vlan access-map VLAN22_FILTER 20 match mac address ALLOWED_L2_TRAFFIC action forward</pre> <p>Apply filter to VLAN:</p> <pre>vlan filter VLAN22_FILTER vlan-list 22 vlan filter TEST vlan-list all (1-4096)</pre>	<h3>What does the following command do:</h3> <h3>ip dhcp relay information trust-all:</h3>	<p>Instruct the IOS DHCP Server to accept DHCP messages with a zero "giaddr" using the global command:</p> <pre>ip dhcp relay information trust-all</pre> <p>A DHCP Relay is supposed to set the "giaddr" field to its own IP address</p>																																																												
<h3>Aggregating logging entries:</h3>	<p>Aggregate logging messages so that a log entry is generated after five access-list entry hits, also generates the five minutely logging message</p> <pre>ip access-list log-update threshold 5</pre> <p>reduce CPU by process switching only one packet per second</p> <pre>ip access-list logging interval 1000</pre> <p>Rate-limiting logging -> CPU process switching if log / log-input is used on ACLs</p>	<h3>MAC address-lists</h3> <h3>Filtering Layer 2 protocols:</h3> <h3>PVST+, ARP, CDP, VTP, DTP, UDLD, STP</h3>	<pre>mac access-list extended ALLOWED_L2_TRAFFIC permit any any Isap 0x4242 0x0 permit any any 0x010B 0x0 permit any any 0x806 0x0</pre> <pre>PVST+ 0x010B ARP 0x806 CDP 0x2000 VTP 0x2003 DTP 0x2004 UDLD 0x0111 IEEE STP BPDUs 0x4242</pre> <p>SNAP-encapsulated packets can be matched using an LSAP value of 0xAAAA</p> <p>To check filtering, make opposite switch the root for one Vlan and filter it on the port, check if the local switch can not see the root on that port.</p>	<h3>Show ip dhcp snooping</h3> <h3>Output:</h3>	<pre>SW1#show ip dhcp snooping Switch DHCP snooping is enabled DHCP snooping is configured on following VLANs: 5 Insertion of option 82 is enabled circuit-id format: vlan-mod-port remote-id format: MAC Option 82 on untrusted port is not allowed Verification of hwaddr field is enabled Interface Trusted Rate limit (pps) ----- FastEthernet0/1 no 10 FastEthernet0/13 yes unlimited</pre> <p>Option 82: Remote-id = BIA of switch Port-ID = Port where the Client is attached</p>																																																												
<h3>Stateful Filtering with CBAC</h3>	 <pre>interface Serial 0/1/0 ip inspect INSPECT out ip access-group INBOUND in ip access-list extended INBOUND deny ip any any ip inspect name INSPECT ftp alert on / alerts only ip inspect name INSPECT http audit-trail on ip inspect name INSPECT udp timeout 30 /inactivity 30sec ip inspect name INSPECT tftp</pre>	<h3>Use IPX for troubleshooting Layer2 MAC address lists filtering Non-IP protocols:</h3>	<pre>Router1# ipx routing interface FastEthernet 0/1 ipx network 146 encapsulation snap (Vlan146) R1#show ipx interface fastEthernet 0/0 FastEthernet0/0 is up, line protocol is up IPX address is 146.000d.edc8.4f60, SNAP [up] On router 2 test reachability via: ping 146.000d.edc8.4f60</pre>	<h3>Useful DHCP Snooping show commands:</h3>	<pre>show ip dhcp snooping database detail more flash:/dhcp-bindings.txt SW1#show ip dhcp snooping binding MacAddress IpAddress Lease(sec) Type VLAN Interface ----- 00:0C:31:EF:4E:60 155.1.146.5 85188 dhcp-snooping 5 FastEthernet0/1 Total number of bindings: 1</pre>																																																												

Help me create more flashcards:

Simply press this button and send me your credit cards regards!


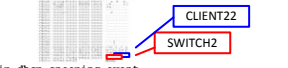
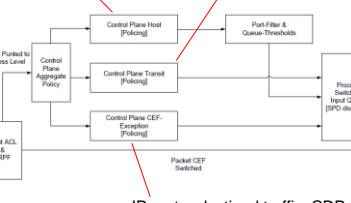
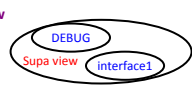
Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts



Colin

<h2>DHCP Snooping and the Information Option</h2> <p>Setting the remote-id and circuit-id strings:</p>	<p>SW2: </p> <pre>ip dhcp snooping ip dhcp snooping vlan 5 ip dhcp snooping information option format remote-id string SWITCH2 ip dhcp snooping information option allow-untrusted interface FastEthernet0/6 ip dhcp snooping vlan 5 information option format-type circuit-id string CLIENT22</pre> <p>Interface fa0/1 Ip dhcp snooping trust</p> <p>Remote-ID = dhcp relay (sw2) Circuit-id =point of client's attachment</p>	<h3>show ip arp inspection vlan X</h3> <p>Output:</p>	<pre>SW2#show ip arp inspection vlan 146 Source Mac Validation : Enabled Destination Mac Validation : Enabled IP Address Validation : Enabled Vlan Configuration Operation ACL Match Static ACL ----- 146 Enabled Active ARP_VLAN146 No Vlan ACL Logging DHCP Logging Probe Logging ----- 146 Acl-Match Deny Off</pre>	<h2>Control Plane Protection (CPPr)</h2> <p>Explaining:</p> <p>control-plane host subinterface</p> <p>control-plane transit subinterface</p>	<p>CPPr treats the Route Processor (RP) as a virtual interface attached to the router classified into three categories or sub-interfaces:</p> <p>control-plane host subinterface receives all control plane TCP/UDP traffic that is directly destined for router interfaces.non TCP/UDP control traffic will end up on the CEF exception sub-interface. Has most policing, port-filtering, per protocol queue thresholds</p> <p>control-plane transit subinterface. Handles <i>transit</i> IP packets not handled via CEF Ethernet interface and no ARP lookup has been made yet for the next-hop, making the CEF adjacency incomplete</p>
<h2>What happens when a partially trusted port receives a DHCP packet with a zero GiAddr content?</h2>	<p>ip dhcp snooping information option allow-untrusted & Interface X & ip dhcp snooping trust</p> <p>Need to be set that the switch will NOT reject packets with a Giaddr of zero value.</p> <p>Troubleshoot via: access-list 100 permit udp any any eq bootps debug ip packet 100 dump</p>  <pre>debug ip dhcp snooping event \DHCP_SNOOPING-5-DHCP_SNOOPING_NONZERO_GIADDR: DHCP_SNOOPING_drop</pre>	<h3>show ip verify source</h3> <p>Output:</p>	<pre>show ip verify source Interface Filter-type Filter-mode IP-address Mac-address Via ----- Fa0/6 ip-mac active 155.1.67.6 00:0C:CE:65:50:A0 67 Fa0/6 ip-mac active 155.1.146.3 00:0C:CE:65:50:A0 25</pre>	<h2>Control Plane Protection (CPPr)</h2> <p>Explaining:</p> <p>control-plane CEF exception subinterface</p>	<p>control-plane CEF exception subinterface</p> <p>This is where packets causing an exception in the CEF switching path land. include non- IP router-destined traffic, such as CDP, L2 keepalive messages, ARP packets and IP packets with options or TTL <=1. Non-UDP local multicast traffic destined to the router, OSPF updates fall under this category.</p> <p>NBAR can NOT be used for control-plane traffic classification !!</p>
<h2>Dynamic ARP Inspection</h2> <p>maximum, 16 entries in the ARP logging buffer. Enable all additional sanity checks for ARP packets</p>	<pre>arp access-list ARP_VLAN146 permit ip host 155.1.146.1 mac host 000d.edc8.4f60 log permit ip host 155.1.146.4 mac host 0015.c634.6c61 log ip arp inspection vlan 146 ip arp inspection vlan 146 logging acl-match matchlog ip arp inspection log-buffer entries 16 ip arp inspection log-buffer logs 4 interval 10 ! Checks content of the ARP packets ip arp inspection validate src-mac dst-mac ip ! Apply the ARP ACL ip arp inspection filter ARP_VLAN146 vlan 146</pre>	<h2>Controlling Terminal Line Access</h2>	<p>! Global ACL 99 in place for VTY access:</p> <pre>line vty 0 4 access-class 99 in login local</pre> <p>User TELNET has a second ACL, permitting this user only from ACL 100 allowed networks.</p> <pre>username TELNET password CISCO username TELNET access-class 100</pre> <p>Access-class 100 could be useful in combination of restricting rotary access to ports like 80, 161 etc..</p>	<h2>Control Plane Protection (CPPr)</h2> <p>Diagram:</p>	<p>no egress policing for any of the host subinterfaces</p> <p>All TCP/UDP to directly destined on router</p> <p>no ARP lookup, CEF adjacency still incomplete</p>  <p>non- IP router-destined traffic, CDP, L2 keepalives, multicast OSPF updates</p>
<h2>What speciality is to keep in mind using Reflexive Access-lists?</h2>	<p>account for local traffic by either statically permitting the traffic in the inbound ACL or use local policy routing to divert the local traffic across the loopback interface and make it re-enter the router</p>	<h2>IOS Login Enhancements</h2>	<pre>username TEST password TEST access-list 99 permit 150.1.5.5 ! 3 unsuccessful attempts in 30sec block for 40 sec login block-for 40 attempts 3 within 30 ! Exempt traffic sourced in 99 from block restriction login quiet-mode access-class 99 Log every 3rd unsuccessful attempt: login on-failure log every 3 Log every successful attempt: login on-success log login delay 2 line vty 0 4 login local</pre>	<h2>Control Plane Protection (CPPr)</h2> <p>Control-plane host:</p>	<p>policy-map CPP_HOST_POLICY class BGP police rate 256000</p> <p>control-plane host service-policy input CPP_HOST_POLICY</p>
<h2>Dynamic arp inspection command by command:</h2>	<p>Used to trust/disable inspection on trunks: ip arp inspection trust</p> <p>Check correctness of contents of ARP packets for src and DST ip arp inspection validate src-mac dst-mac</p> <p>IP address consistency checks for ARP packets Ensures no host binds 0.0.0.0 or 255.255.255.255 ip arp inspection validate ip</p> <p>For host NOT using DHCP create static entries: ip arp inspection filter <ARP_ACL> vlan <vlan_ID> static</p> <p>Starts logging allowed or denied packets ip arp inspection vlan <VLAN_ID> logging acl-match</p>	<h2>Role Based CLi</h2>	<pre>aaa new-model enable view parser view DEBUG secret CISCO commands exec include show running-config commands exec include all debug commands exec include all undebug parser view INTERFACE1 secret CISCO commands interface include all ip commands configure include interface commands exec include configure terminal commands configure include interface FastEthernet0/0 parser view SUPER superview secret CISCO view DEBUG view INTERFACE1</pre> 	<h2>How can one check for OPEN ports in the control-plane:</h2>	<h3>show control-plane host open ports</h3>
<h2>Dynamic arp inspection Commands:</h2>	<p>Logging packets with a source IP of 0.0.0.0 ip arp inspection vlan <VLAN_ID> logging dhcp-bindings {all permit none}</p> <p>ip arp inspection vlan <VLAN_ID> logging arp-probe</p> <p>regulate the size of this buffer ip arp inspection log-buffer entries <N></p> <p>switch empties this buffer using a rate-limited procedure ip arp inspection log-buffer logs <number> <interval></p> <p>estricit the rate of received ARP packets to <pps> per <seconds> interval ip arp inspection limit rate {<pps> burst interval <seconds> none }</p>	<h2>Controlling the ICMP Messages Rate</h2>	<pre>Interface X no ip unreachable limits the total rate of all router-generated unreachable messages (100 per second = <10>) ip icmp rate-limit unreachable <once per this ms> Control the rate of "packet-too-big" messages (<1> = 1000 times per second) ip icmp rate-limit unreachable DF <once per ms></pre>	<h2>Control Plane Protection (CPPr)</h2> <p>Control-plane host / port-filter:</p>	<p>Control-plane host, port-filtering feature: will effectively drop all packets destined to closed ports before affecting router's CPU</p> <pre>class-map type port-filter match-all CLOSED_PORTS match closed-ports match not port 3020 match not port 520</pre> <p>Close all except of 3020 (rotary-group) or RIP routing protocol</p> <pre>policy-map type port-filter FILTR_CLOSED_PORTS class CLOSED_PORTS drop</pre> <p>control-plane host service-policy type port-filter input FILTR_CLOSED_PORTS</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts


Colin

<p>Control-Plane Protection (CPPr)</p> <p>Control-plane host / queue-threshold</p>	<p>separate limit enforced for this particular protocol in the common input queue, only one threshold policy map possible. Map different threshold class maps to route map.</p> <pre>class-map type queue-threshold CMAP_BGP match protocol bgp policy-map type queue-threshold BGP-Tresh class CMAP_BGP queue-limit 50 ----- In packets control-plane host service-policy type queue-threshold input BGP-Tresh</pre>	<p>Flexible Packet Matching</p>	<p>-Load a PHDF (optional). PHDF Packet Header Definition File written in XML</p> <pre>load protocol system:fpm/phdf/ip.phdf load protocol system:fpm/phdf/icmp.phdf system:fpm/phdfs Define a protocol stack (optional). class-map type-stack match-all <NAME></pre>	<p>ZFW Application Inspection</p> <p>AIC</p>	<p>Config too big to display here!</p> <p>Look up on page 211 INE, section-11-security.pdf Lab Volume 1, task 11.41</p>																																						
<p>Rotary / Base TCP ports explained</p>	<table border="1"> <thead> <tr> <th>Services</th> <th>Base TCP port</th> <th>Base for Lines</th> </tr> </thead> <tbody> <tr> <td>Telnet protol</td> <td>3000</td> <td>2000</td> </tr> <tr> <td>Raw TCP no telnet</td> <td>5000</td> <td>4000</td> </tr> <tr> <td>Telnet binary mode</td> <td>7000</td> <td>6000</td> </tr> <tr> <td>XRemote protol</td> <td>10000</td> <td>9000</td> </tr> </tbody> </table> <p>line vty 0 password cisco login rotary 20</p> <p>using rotary 20, one could telnet to 3020 to the router rotary 20</p>	Services	Base TCP port	Base for Lines	Telnet protol	3000	2000	Raw TCP no telnet	5000	4000	Telnet binary mode	7000	6000	XRemote protol	10000	9000	<p>Flexible Packet Matching</p> <p>Packet matching</p> <p>(order of match statements is important):</p>	<pre>class-map type stack match-all TCP_IP_IP_ETHER stack-start l2-start match field layer 1 ETHER type eq 0x800 next IP match field layer 2 IP protocol eq 4 next IP match field layer 3 IP protocol eq 6 next TCP</pre>	<p>Bridge IRB</p> <p>Input-type-list</p>	<p>Define Bridging either (IRB or CRB)</p> <pre>access-list 201 deny 0x86dd deny IPv6 access-list 201 permit 0x0 ffffff interface X bridge-group 1 bridge-group 1 input-type-list 201</pre>																							
Services	Base TCP port	Base for Lines																																									
Telnet protol	3000	2000																																									
Raw TCP no telnet	5000	4000																																									
Telnet binary mode	7000	6000																																									
XRemote protol	10000	9000																																									
<p>How to check CPPr</p> <p>Queue-thresholds:</p> <p>CEF-exceptions subinterface statistics:</p>	<pre>show policy-map type queue-threshold control-plane host show policy-map control-plane cef-exception show policy-map control-plane transit</pre>	<p>Flexible Packet Matching</p> <p>Filter ICMPs with string AAA in payload, look no deeper than 256 bytes in packet:</p>	<pre>class-map type stack ICMP_IN_IP_IN_ETHER stack-start l2-start match field ether type eq 0x800 next ip match field layer 2 ip protocol eq 1 next icmp class-map type access-control match-all ICMP_ECHO_STRING match field icmp type eq 8 match start icmp payload offset 0 size 256 regex "AAAA" policy-map type access ACCESS_CONTROL_POLICY class ICMP_ECHO_STRING log policy-map type access-control STACK_POLICY class ICMP_IN_IP_IN_ETHER service-policy ACCESS_CONTROL_POLICY</pre>	<p>Control Plane Protection</p> <p>Show commands:</p>	<p>show control-plane features</p> <pre>show control-plane counters Feature Path Packets processed/dropped/errors Aggregate 164103/0/0 Host 13840/0/0 Transit 9736/0/0 Cef-exception 140558/0/0 class-map type port-filter match-all CLOSED_PORTS match closed-ports match not port TCP 2020 match not port TCP 2040</pre>																																						
<p>IOS ACL Selective IP Option Drop</p> <p>IP source route [strict / loose]:</p>	<p>silently discard all packets with IP options using the command:</p> <p>ip options drop</p> <p>Or use an access-list:</p> <p>IP source route loose:</p> <p>deny ip any any option lsr</p> <p>IP source route strict:</p> <p>deny ip any any option ssr</p>	<p>ASCII table:</p> <table border="1"> <thead> <tr> <th>Decimal</th> <th>Hex</th> <th>Character</th> </tr> </thead> <tbody> <tr><td>65</td><td>41</td><td>A</td></tr> <tr><td>..</td><td>..</td><td>..</td></tr> <tr><td>90</td><td>5B</td><td>Z</td></tr> <tr><td>..</td><td>..</td><td>..</td></tr> <tr><td>97</td><td>61</td><td>a</td></tr> <tr><td>..</td><td>..</td><td>..</td></tr> <tr><td>122</td><td>7A</td><td>z</td></tr> <tr><td>..</td><td>..</td><td>..</td></tr> <tr><td>48</td><td>30</td><td>0</td></tr> <tr><td>49</td><td>31</td><td>1</td></tr> <tr><td>..</td><td>..</td><td>..</td></tr> <tr><td>57</td><td>39</td><td>9</td></tr> </tbody> </table>	Decimal	Hex	Character	65	41	A	90	5B	Z	97	61	a	122	7A	z	48	30	0	49	31	1	57	39	9	<p>Control Plane Protection</p> <p>RMPs:</p>	<pre>control-plane host service-policy input HOST_RATE_LIMIT service-policy type port-filter input HOST_PORT_FILTER control-plane transit service-policy input TRANSIT_RATE_LIMIT control-plane cef-exception service-policy input CEF_EXCEPTION_RATE_LIMIT</pre>
Decimal	Hex	Character																																									
65	41	A																																									
..																																									
90	5B	Z																																									
..																																									
97	61	a																																									
..																																									
122	7A	z																																									
..																																									
48	30	0																																									
49	31	1																																									
..																																									
57	39	9																																									
<p>Troubleshooting IP source route [loose / strict]</p> <p>using ping</p> <p>deny ip any any option lsr is set in an ACL on the other side:</p>	<pre>Rack1SW1#ping Protocol [ip]: Target IP address: 155.1.37.3 Loose, Strict, Record, Timestamp, Verbose[none]: Loose Source route: 155.1.37.3 ... Sending 5, 100-byte ICMP Echos to 155.1.37.3, timeout is 2 seconds: Packet has IP options: Total option bytes= 7, padded length=8 Loose source route: <*> (155.1.37.3) Unreachable from 155.1.37.3. Received packet has options Total option bytes= 7, padded length=8 Loose source route: <*> (155.1.37.3)</pre>	<p>Show ip port-map</p> <p>Output:</p>	<pre>R4#show ip port-map Default mapping: snmp udp port 161 system defined Default mapping: echo tcp port 7 system defined Default mapping: echo udp port 7 system defined Default mapping: telnet tcp port 23 system defined Default mapping: wins tcp port 1512 system defined Default mapping: n2h2server tcp port 9285 system defined Default mapping: n2h2server udp port 9285 system defined</pre>	<p>Zone Based Firewall</p> <p>show parameter-map type inspect</p> <p>Output:</p>	<pre>R5#show parameter-map type inspect parameter-map type inspect PMAP_PARAMS audit-trail on alert off max-incomplete low 1000 max-incomplete high 2000 one-minute low 10 one-minute high 100 udp idle-time 10 icmp idle-time 5 dns-timeout 15 tcp idle-time 3600 tcp finwait-time 5 tcp synwait-time 30 tcp max-incomplete host 200 block-time 1 sessions maximum 5000</pre>																																						
<p>Flexible Packet Matching</p> <p>FPM overview:</p>	<ul style="list-style-type: none"> - Using FPM you can match any string, byte or even bit at any position in an IP (or theoretically non-IP) packet - FPM is completely stateless and can not discover dynamic protocol ports - Only the initial fragment can be inspected. IP packets with IP options are not matched by FPM -> sent to CPU - Inspects only unicast packets, no MPLS packets <p>Configuring FPM filter:</p> <ol style="list-style-type: none"> (1) Loading protocol headers. (2) Defining a protocol stack. (3) Defining a traffic filter. (4) Applying the policy & Verification 	<p>Zone Based Firewall</p> <p>ZFW Rate Limiting flows</p>	<p>Supports two types of rate-limiting:</p> <ol style="list-style-type: none"> 1) Limiting aggregate packet rate for the flows between security zones. 2) Limiting the maximum number and/or rate of the half-open connections for TCP/UDP sessions. <pre>policy-map type inspect OUTSIDE_TO_INSIDE class ICMP inspect police 256000 burst 8000</pre>	<p>Zone Based Firewall</p> <p>Limiting max number of half open connections / paramaters:</p>	<pre>parameter-map type inspect PMAP_PARAMS max-incomplete low 1000 max-incomplete high 2000 one-minute low 10 one-minute high 100 tcp max-incomplete host 200 block-time 1 sessions maximum 5000 dns-timeout 15 policy-map type inspect PMAP_OUTSIDE_DMZ class CMAP_OUTSIDE_TO_DMZ_ACCESS police rate 512000 burst 32000 inspect PMAP_PARAMS</pre>																																						

Help me create more flashcards:



Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts



Colin

<p>The difference between bridge crb and bridge irb</p>	<p>bridge crb (router has no SVI in L2 domain, not reachable)</p> <p>bridge x bridge ip (bridge IP packets, no ip route mode)</p> <p>bridge x route ip (ignore IP packets, follow IP routing rules)</p> <hr/> <p>bridge irb</p> <p>If BVI is in bridge group, enable either IP bridging or IP routing, but not both!</p> <p>bridge irb (bvi represents L3 in L2 domain) bridge X route ip (similar as a SVI) bridge X protocol IEEE (probably necessary) int Z, bridge group X int BVIx</p>	<p>802.1x Authentication</p> <p>Config using Radius:</p>	<pre>aaa new-model aaa authentication login default none aaa authentication dot1x default group radius dot1x system-auth-control interface FastEthernet0/9 switchport mode access dot1x port-control auto interface FastEthernet0/10 switchport mode access dot1x port-control auto ip radius source-interface Loopback0 radius-server host 204.12.1.100 radius-server key CISCO</pre>	 <p>Securing sub-interfaces</p> <p>For traffic transiting by using a policy-map:</p>	 <pre>class-map match-all FROM_ETHERNET match input-interface FastEthernet0/1 policy-map SAME_INTERFACE class FROM_ETHERNET drop interface FastEthernet0/1.52 service-policy output SAME_INTERFACE interface FastEthernet0/1.53 service-policy output SAME_INTERFACE</pre>
<p>Classic IOS Transparent Firewall</p> <p>Using bridge IRB</p>	<pre>interface BVI1 ip address 10.0.0.6 255.255.255.0 ip inspect name INSIDE_PROTOCOLS http ip inspect name OUTSIDE_PROTOCOLS http ip access-list extended OUTSIDE_IN permit tcp any any eq 80 deny ip any any log access-list 201 deny 0x86dd ! block IPv6 access-list 201 permit 0x0 0xFFFF interface FastEthernet 0/0 ip inspect INSIDE_PROTOCOLS IN bridge-group 1 bridge-group 1 input-type-list 201 interface FastEthernet 0/1 ip inspect OUTSIDE_PROTOCOLS OUT bridge-group 1 bridge-group 1 input-type-list 201</pre>	<p>802.1x</p> <p>show dot1x all</p> <p>Output:</p>	<pre>Dot1x Info for FastEthernet0/9 ----- PAE = AUTHENTICATOR PortControl = ALTO ControlDirection = Both HostMode = SINGLE_HOST ReAuthentication = Disabled QuietPeriod = 60 ServerTimeout = 30 SuppTimeout = 30 ReAuthPerio = 3600 (Locally configured) ReAuthMax = 2 MaxReq = 2 TxPeriod = 30 RateLimitPeriod = 0 SW1#show dot1x all Sysauthcontrol Enabled Dot1x Protocol Version 2 Critical Recovery Delay 100 Critical EAPOL Disabled</pre>	<p>ERSPAN Source Session:</p>	<pre>ERSPAN Source Session: monitor session <session-nr> type erspan-source description SNIFF of bla source interface fa0/x (optional filter vlan_range ..) destination ip address <IP.IP.IP> erspan-id <flow-ID> origin ip address <IP.IP.IP> [force] ! SRC of flow ip ttl <ttl> ip prec <IPP-Value> no shutdown</pre>
<p>Troubleshooting:</p> <p>Classic IOS Transparent Firewall</p> <p>Using bridge IRB</p> <p>(L2 firewall mode)</p>	<pre>show bridge Total of 300 station blocks, 297 free Codes: P - permanent, S - self Bridge Group 1: Address Action Interface Age RX count TX count 0013.c451.f240 forward Fa0/0.146 0 123 23 0013.c440.3980 forward Fa0/0.67 0 297 23 show ip inspect config show ip inspect sessions debug ip inspect l2-transparent packets</pre>	<p>Troubleshooting AAA Servers:</p> <p>show aaa servers:</p>	<pre>SW1#show aaa servers RADIUS: id 1, priority 1, host 204.12.1.100, auth-port 1645, acct-port 1646 State: current UP, duration 1636s, previous duration 0s Dead: total time 0s, count 0 Quarantined: No Authen: request 0, timeouts 0 Response: unexpected 0, server error 0, incorrect 0, time 0ms Transaction: success 0, failure 0 Author: request 0, timeouts 0 Response: unexpected 0, server error 0, incorrect 0, time 0ms Transaction: success 0, failure 0 Account: request 0, timeouts 0 Response: unexpected 0, server error 0, incorrect 0, time 0ms Transaction: success 0, failure 0 Elapsed time since counters last cleared: 14m</pre>	<p>ERSPAN Destination Session:</p>	<pre>ERSPAN Destination monitor session <session-nr> type erspan-destination description bla destination interface fa0/x source ip address <ip.ip.ip> [force] ! must match source erspan-id <flow-ID> no shutdown</pre>
<p>ZFW-Based IOS Transparent Firewall</p> <p>CPL (L2 firewall mode)</p> <p>Part-1</p>	<pre>class-map type inspect match-any CMAP_PROTOCOLS_FROM_INSIDE match protocol http class-map type inspect match-any CMAP_PROTOCOLS_TO_INSIDE match protocol ftp class-map type inspect CMAP_RIP_TRAFFIC match access-group name ACL_RIP_TRAFFIC policy-map type inspect PMAP_INSIDE_TO_OUTSIDE class CMAP_PROTOCOLS_FROM_INSIDE inspect policy-map type inspect PMAP_OUTSIDE_TO_INSIDE class CMAP_PROTOCOLS_TO_INSIDE inspect</pre>				
<p>ZFW-Based IOS Transparent Firewall</p> <p>CPL (L2 firewall mode)</p> <p>Part-2</p>	<pre>zone security INSIDE zone security OUTSIDE zone-pair security ZP_INSIDE_TO_OUTSIDE source INSIDE destination OUTSIDE service-policy type inspect PMAP_INSIDE_TO_OUTSIDE zone-pair security ZP_OUTSIDE_TO_INSIDE source OUTSIDE destination INSIDE service-policy type inspect PMAP_OUTSIDE_TO_INSIDE interface FastEthernet 0/0 bridge-group 1 zone-member security INSIDE interface FastEthernet 0/2 bridge-group 1 zone-member security OUTSIDE</pre>				
<p>IOS IPS</p> <p>Integration steps</p>	<p>Step 1. Downloading IOS IPS files Step 2. Creating IOS IPS configuration directory on flash Step 3. Configuring IOS IPS crypto key Step 4. Enabling IOS IPS Step 5. Loading IOS IPS signature package to router Step 6: Tuning Signatures</p>	<p>IOS IPS</p> <p>(config reference only)</p> <p>Copy paste only, as no longer in version 5 lab</p>	<pre>show ip ips all ip ips name MYIPS ip ips config location flash:ips ip ips notify sds ip ips notify log interface FastEthernet 0/1 ip ips MYIPS in interface FastEthernet 0/0 ip ips MYIPS in ip ips signature-category category all retired true exit exit ip ips signature-definition signature 2004 0 status enabled true retired false exit exit signature 1004 0 status enabled true retired false exit exit engine event-action deny-attacker-inline exit exit</pre>		

Help me create more flashcards:


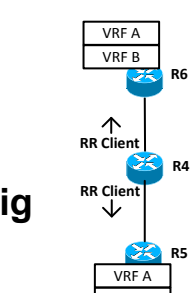
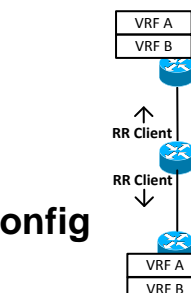
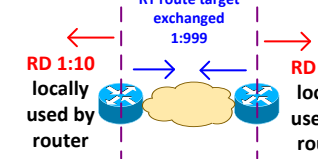
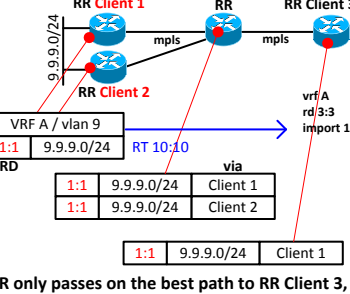
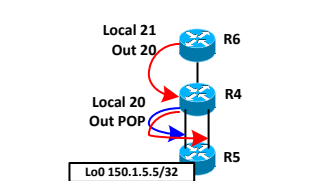
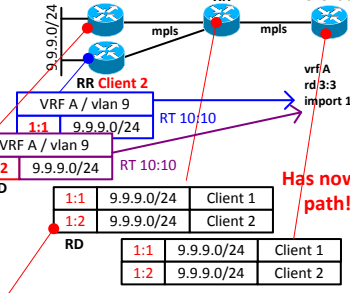
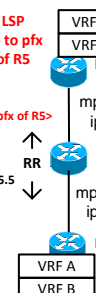
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!


Thanks for appreciating my efforts

Colin

<h3>VRF Lite</h3>	<p>Router# ip vrf VPN_A rd 100:1 ip vrf VPN_B rd 100:2 Int Gi0/0.67 ip vrf forwarding VPN_A ip addr ... Int Gi0/0.76 ip vrf forwarding VPN_B ip addr ... ip route vrf VPN_A 192.168.7.0 255.255.255.0 Gi0/0.76 1.1.1.1 ip route vrf VPN_B 172.16.7.0 255.255.255.0 Gi0/0.67 2.2.2.2</p> <p>Don't forget Switch# ip routing ip vrf VPN_A rd 100:1 ip vrf VPN_B rd 100:2 Interface / routes...</p> 	<h3>Show mpls ldp neighbour</h3> <p>Output:</p>	<pre>R6#show mpls ldp neighbor Peer LDP Ident: 150.1.4.4:0; Local LDP Ident 150.1.6.6:0 TCP connection: 155.1.146.4.646 - 155.1.146.6.33662 State: Oper; Msgs sent/rcvd: 17/17; Downstream Up time: 00:03:58 LDP discovery sources: GigabitEthernet0/0.146, Src IP addr: 155.1.146.4 Addresses bound to peer LDP Ident: 204.12.1.4 155.1.146.4 155.1.0.4 155.1.45.4 150.1.4.4</pre>	<h3>MP-BGP VPNv4</h3> <h4>R4's config</h4> <p>Route-reflector</p> 	<pre>router ospf 1 mpls ldp autoconfig router-id 150.1.4.4 router bgp 100 no bgp default ipv4-unicast bgp log-neighbor-changes neighbor 150.1.5.5 remote-as 100 neighbor 150.1.5.5 update-source Loopback0 neighbor 150.1.6.6 remote-as 100 neighbor 150.1.6.6 update-source Loopback0 address-family vpnv4 neighbor 150.1.5.5 activate neighbor 150.1.5.5 send-community extended neighbor 150.1.5.5 route-reflector-client neighbor 150.1.6.6 activate neighbor 150.1.6.6 send-community extended neighbor 150.1.6.6 route-reflector-client exit-address-family</pre>
<h3>Show ip vrf output:</h3>	<pre>SW1#show ip vrf Name Default RD Interfaces VPN_A 100:1 VI67 Lo101 VPN_B 100:2 VI76 Lo102</pre>	<h3>MPLS LDP</h3> <h4>Configuring MPLS LDP password on only one side:</h4> <h4>Configuring a wrong MPLS LDP password:</h4>	<pre>conf t no mpls ldp neighbor 2.2.2.2 password CISCO TCP-6-BADAUTH: No MD5 digest from 1.1.1.1(47531) to 2.2.2.2(646) conf t mpls ldp neighbor 2.2.2.2 password WRONG-pass TCP-6-BADAUTH: Invalid MD5 digest from 1.1.1.1(43524) to 2.2.2.2(646)</pre>	<h3>MP-BGP VPNv4</h3> <h4>R5/R6's config</h4> 	<pre>interface FastEthernet 0/0 ip vrf forwarding VPN_A rd 100:1 ip address x.x.x.x 255.255.255.0 route-target both 100:1 interface FastEthernet 0/1 ip vrf forwarding VPN_B rd 100:2 ip address x.x.x.x 255.255.255.0 route-target both 100:2 router bgp 100 no bgp default ipv4-unicast neighbor 150.1.4.4 remote-as 100 neighbor 150.1.4.4 update-source Loopback0 address-family vpnv4 unicast neighbor 150.1.4.4 activate neighbor 150.1.4.4 send-community extended address-family ipv4 vrf VPN_A redistribute connected redistribute static address-family ipv4 vrf VPN_B redistribute connected redistribute static</pre>
<h3>Significants of RD route-distinguisher and RT route target</h3>	 <pre>ip vrf A rd 1:10 route target export 1:999 route target import 1:999 ip vrf A rd 1:10 route target export 1:999 route target import 1:999</pre>	<h3>Show mpls neighbor password</h3> <p>Output:</p>	<p>Necessary config:</p> <pre>mpls ldp password required mpls ldp neighbor 150.1.5.5 password CISCO</pre> <pre>R1#show mpls ldp neighbor password Peer LDP Ident: 150.1.5.5:0; Local LDP Ident 150.1.4.4:0 TCP connection: 150.1.5.5.42010 - 150.1.4.4.646 Password: required, neighbor, in use State: Oper; Msgs sent/rcvd: 28/28</pre>	<h3>MP-BGP VPNv4</h3> <h4>And Route-Reflector</h4> <p>Using the same RD on PE's:</p>	 <p>RR only passes on the best path to RR Client 3, to overcome this use different route targets</p>
<h3>Enable mpls in two ways:</h3>	<p>Interface X mpls ip</p> <p>or</p> <p>if OSPF is configured use following command to enable all OSPF enabled interfaces for MPLS:</p> <pre>mpls ldp autoconfig [area 2]</pre>	<h3>Tracing from R6 to R5's loopback address:</h3> <h4>Following the unidirectional label path:</h4> 	<p>On R5 no info can be found in:</p> <pre>show mpls forwarding-table</pre> <pre>R6#traceroute 150.1.5.5 1 155.1.146.4 [MPLS: Label 20 Exp 0] 92 msec 92 msec 88 msec 2 155.1.0.5 12 msec * 8 msec R4# 1 155.1.45.5 0 msec 155.1.0.5 12 msec * R5# 1 150.1.5.5 0 msec * 0 msec</pre>	<h3>MP-BGP VPNv4</h3> <h4>and Route-Reflector</h4> <p>Using the different RD on PE's:</p> <p>RR only passes best path to Client 3:</p>	 <p>Has now 2 path!</p> <p>Routes have community set to 10:10</p>
<h3>MPLS LDP discovery:</h3>	<ol style="list-style-type: none"> initially LDP sends multicast discovery messages using 224.0.0.2 UDP 646 Hears other LDP router, (highest IP or mpls ldp router-id <interface>) Will start to establish a TCP session to other LDP router-ID (normally from loopback) to use the physical address instead: mpls ldp discovery transport-address <interface> Using LDP-Router-IDs, rtrs establish TCP session to port 646. Can be authenticated using: interface X, mpls ldp neighbor <IP> password <password> Or globally: mpls ldp password required LFIB should be getting populated 	<h3>MPLS Label Filtering</h3>	<p>Normally labels are generated for all entries in the routing table outbound. In order to generate labels only for specific prefixes use an access-list and specify the prefixes you want to have labels assigned for. Generate labels only for Loopbacks for example:</p> <pre>access-list 99 permit 150.1.6.6</pre> <pre>no mpls ldp advertise-labels mpls ldp advertise-labels for 99</pre> <pre>R6#show mpls forwarding-table Local Outgoing Prefix Bytes tag Outgoing Next Hop tag tag or VC or Tunnel Id switched interface 16 Untagged 204.12.1.0/24 0 Gi0/0.146 1.2.3.4 17 Untagged 155.1.0.0/24 0 Gi0/0.146 1.2.3.4</pre>	<h3>Prevent BGP to automatically activate a configured session:</h3>	<pre>router bgp 100 no bgp default ipv4-unicast</pre> <p>Prevents bgp from automatically activating the session:</p>
<h3>MPLS IP</h3> <h4>Configuration using Local address for LDP</h4> <p>Enable all OSPF interfaces for MPLS</p> <p>Make passwords mandatory:</p>	<pre>mpls ip Enable MPLS mpls ldp router-id Loopback0 force set mpls router-id interface FastEthernet 0/1 mpls ldp discovery transport-address interface Using local address to establish LDP session router ospf 1 mpls ldp autoconfig Enable MPLS on all OSPF enabled interfaces mpls ldp password required Make password mandatory mpls ldp neighbor 150.1.5.5 password CISCO</pre>	<h3>MPLS IP</h3> <h4>Configuration using The loopback for the LDP session</h4> <p>Enable MPLS per interface</p> <p>Make passwords mandatory</p>	<pre>mpls ip Using Loopback0 for LDP session mpls ldp router-id Loopback0 force interface Serial 0/1/0 mpls ip Enable MPLS per interface Make password mandatory mpls ldp password required mpls ldp neighbor 150.1.4.4 password CISCO</pre>	<h3>MPLS label switch path</h3> <h4>Explained:</h4> <p>LSP from R6 to a prefix announced from R5</p> <p>PE-P-PE</p>	<pre>R6# show ip bgp vrf A <prfx of R5> Originator R5 (5.5.5.5) In label none Out label 19 show mpls forwarding table 5.5.5.5 Local 17 Out 16 R4# show ip bgp vpnv4 rd 100:1 <prfx of R5> Originator R5 (5.5.5.5) In label none Out label 19 show mpls forwarding table 5.5.5.5 Local 16 Out POP-tag R5# show ip bgp vrf A <prfx of R5> Originator R5 (5.5.5.5) locally originated In label 19 Out tableAggregate(VPN_A) show mpls forwarding table 5.5.5.5 No entry Label 19 is VPN label for the prefix on R5, distributed via BGP</pre> 

Help me create more flashcards:

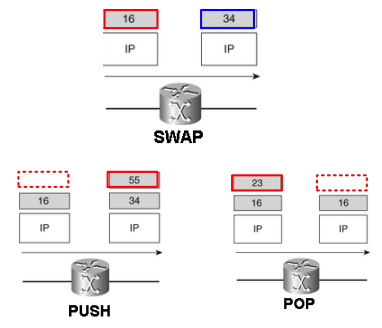
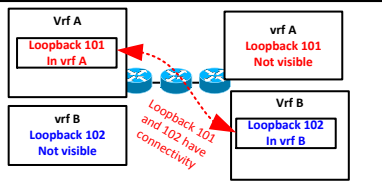


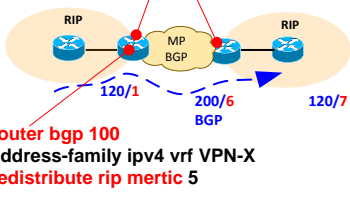

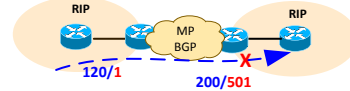
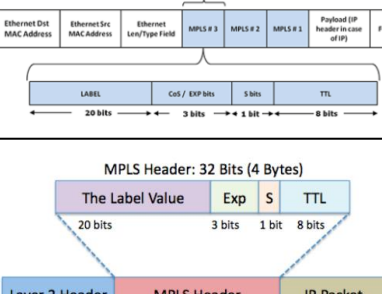
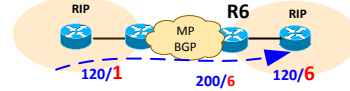
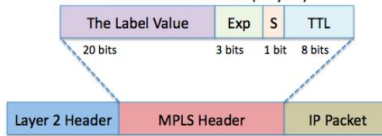
Simply press this button and send me your credit cards regards!



Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<h3>MP-BGP Prefix Filtering</h3> <p>Import/export maps</p>	<pre>import map <ROUTE_MAP_NAME> export map <ROUTE_MAP_NAME> can match based on: - ACLS - prefix-lists - extended-communities Can be used to set extended-communities via: set extcommunity rt Import-map has implicit deny as default! If not permitted by export maps the prefix will not be exported into the BGP process</pre>	<h3>Troubleshooting MP-BGP / MPLS</h3> <p>Commands:</p>	<h3>MPLS</h3> <p>SWAP, PUSH, POP</p> 
<h3>MP-BGP Prefix Filtering</h3> <p>Example diagram:</p> <p>Allowing single prefixes from VPN A to VPN B of another PE</p>	 <pre>route-map VPN_A_EXPORT permit 10 match ip address prefix-list LO101 set extcommunity rt 100:55 route-map VPN_A_EXPORT permit 20 set extcommunity rt 100:1 ip vrf VPN_A export map VPN_A_EXPORT route-target both 100:1 route-target import 100:66 VPN B Exports Lo102 to 100:66 Everything else to 100:2 imports 100:2 imports 100:55</pre>	<h3>PE-CE Routing with RIPv2</h3> <p>RIPv1 not supported as PE-CE protocol</p> 	<pre>router rip version 2 address-family ipv4 vrf VPN-X redistribute bgp <AS-Nr> metric <value> Sets the RIP metrics via BGP MED attribute redistribute bgp <AS-Nr> metric transparent Is utilizing BGP MED attribute, RIP metrics learned from the remote site (transparent) router bgp 100 address-family ipv4 vrf VPN-X redistribute rip metric <1-16> or <16+></pre>
<h3>MP-BGP Prefix Filtering</h3> <p>Importing a prefix into VPN A</p> <p>VPN A RT 100:1 PFX was set to RT 100:66</p>	<pre>R5#sh ip bgp vpnv4 all 192.168.6.6 BGP routing table entry for 100:1:192.168.6.0/24, version 29 Paths: (1 available, best #1 table VPN_A) ... Extended Community: RT:100:66 ... BGP routing table entry for 100:2:192.168.6.0/24, version 28 Paths: (1 available, best #1, no table) ... Extended Community: RT:100:66 ... VPN a has imported this prefix from RT 100:66</pre>	<h3>PE-CE Routing with RIP</h3> <p>Filtering the rip route out using a bgp redistributed metric within 0-16</p> 	<pre>router rip version 2 address-family ipv4 vrf VPN-X redistribute bgp 100 router bgp 100 address-family ipv4 vrf VPN-X redistribute rip metric 5</pre> 
<h3>MP-BGP Prefix Filtering:</h3> <p>Config:</p>  <pre>RD 100:1 both RT 100:1 Import RT 100:66 Export-map PFX 100:55 RD 100:2 both RT 100:2 Import RT 100:55 Export-map PFX 100:66</pre>	<pre>ip vrf VPN_A rd 100:1 route-target both 100:1 route-target import 100:66 export map VPN_A_EXPORT interface Loopback101 ip vrf forwarding VPN_A ip address 172.16.5.5 255.255.255.0 ip prefix-list LO101 permit 172.16.5.0/24 route-map VPN_A_EXPORT permit 10 match ip address prefix-list LO101 set extcommunity rt 100:55 route-map VPN_A_EXPORT permit 20 set extcommunity rt 100:1</pre>	<h3>PE-CE Routing with RIP</h3> <p>Filtering the rip route out using a bgp redistributed metric of over 16</p> 	<pre>router rip version 2 address-family ipv4 vrf VPN-X Redistribute bgp 100 router bgp 100 address-family ipv4 vrf VPN-X redistribute rip metric 500 (anything over 16)</pre>
<h3>MPLS Header:</h3> <p>Picture:</p> 	<pre>R6#show ip route vrf VPN_B 31.3.0.0 Routing entry for 31.3.0.0/16 Known via "bgp 100", distance 200, metric 12345, type internal Redistributing via rip Advertised by rip metric transparent Last update from 150.1.4.4 00:21:04 ago Routing Descriptor Blocks: * 150.1.4.4 (Default-IP-Routing-Table), from 150.1.4.4, 00:21:04 ago Route metric is 12345, traffic share count is 1 AS Hops 0</pre>	<h3>PE-CE Routing with RIPv2</h3> <p>Router rip redistribute bgp <AS-Nr> metric 5</p> <pre>router bgp redistribute rip</pre> 	<h3>MPLS label operations:</h3> <p>Pop Swap Push Untagged/No Label Aggregate</p>
<h3>MPLS Header:</h3> <p>And Label distribution modes:</p>  <p>20 bits equivalent to 1'048'575 label IDs</p> <p>Label Distribution mode: Unsolicited downstream UC mode</p> <p>Label retention mode: Liberal Label Retention LLR mode</p> <p>LSP control mode: Independent LSP control mode</p>	<pre>show ip ospf 100</pre> <p>Connected to MPLS VPN Superbackbone, VRF VPN_A It is an area border and autonomous system boundary router</p>	<h3>PE-CE Routing with OSPF</h3>	<h3>MPLS labels</h3> <p>Default range</p> <p>Labels within 0 - 15</p>

RIB

RIB commands:
show ip route vrf x

LIB commands:
show adjacency detail
show ip cef vrf x

FIB commands:
show mpls bindings
show mpls ip bindings

LFIB commands:
Show mpls forwarding

Incomming label 23 is swapped with label 20, and label 16 is pushed onto label 20

```
R1#show mpls forwarding-table 10.200.254.4 detail
Local  Outgoing  Prefix  Bytes tag  Outgoing  Next Hop
Tag    tag or VC or Tunnel Id  switched  interface
23     16           10.200.254.4/32  0      Tu1       point2point
MAC/Encaps=14/22, MRU=1496, Tag Stack(20 16), via
E10/0/0
00604700881D00024A400800847 00014000000000
No output feature configured
```

Pushed 16 on top


23 —swap→ 20

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

<p>MPLS labels</p> <p>0, 2, 3, 1, 14 explained:</p>	<p>Label 0 explicit 0 for IPv4 Label 2 explicit 0 for IPv6</p> <p>Copy EXP bits of outer label to inner label, preserving QoS bits</p> <p>Label 3 implicit 0 PE requests POP operation from P router. Only for directly connected or summary routes</p> <p>Label 1 router alert label Label check in Software!</p> <p>Label 14 OAM alert label Not used by Cisco</p>	<p>Basic MPLS LDP config:</p>	<p>ip cef</p> <p>mpls ldp router-id Loopback0 force mpls label protocol LDP</p> <p>int loopback0 ip address 10.10.10.10 255.255.255.255</p> <p>interface X ip address x.x.x.x/yy mpls ip</p>	<p>mpls ldp discovery transport address</p> <p>Explained:</p>	<p>Int lo1 Ip address 1.1.1.1/32</p> <p>interface e1 ip addr x.x.x.x / 24 mpls ldp discovery transport address 1.1.1.1 mpls ip</p> <p>interface e2 ip addr x.x.x.x / 24 mpls ldp discovery transport address 1.1.1.1 mpls ip</p> <p>When a router has multiple links to the same LDP router, the same transport addr must be advertised on all parallel links.</p>
<p>Configuring the MPLS label range:</p> <p>And</p> <p>show mpls label range:</p>	<p>RTR(config)#mpls label range 16 1048575</p> <p>event#show mpls label range Downstream Generic label region: Min/Max label: 16/1048575</p>	<p>show mpls ldp discovery detail:</p>	<p>RTR#show mpls ldp discovery detail</p> <p>Local LDP Identifier: 10.200.254.2:0</p> <p>Discovery Sources: Interfaces: Ethernet0/1/2 (ldp): xmit/recv Enabled: Interface config Hello interval: 5000 ms; Transport IP addr: 10.200.254.2 LDP Id: 10.200.254.5:0 Src IP addr: 10.200.215.2; Transport IP addr: 10.200.254.5 Hold time: 15 sec; Proposed local/peer: 15/15 sec Reachable via 10.200.254.5/32</p> <p>mpls ldp discovery [hello / hold-time <15> / interval <5>]</p>	<p>Number of MPLS LDP session:</p> <p>Frame based transport / Mixed / LC-ATM:</p>	
<p>TTL expired in MPLS network</p> <p>Picture:</p>		<p>MPLS ldp discovery</p> <p>When there is no route to the LDP neighbor:</p> <p>Discovered via multicast, but unable to connect via TCP:</p> <p>show mpls ldp discovery [detail]:</p>	<p>RTR#show mpls ldp discovery</p> <p>Local LDP Identifier: 10.200.254.2:0</p> <p>Discovery Sources: Interfaces: Ethernet0/1/4 (ldp): xmit/recv LDP Id: 10.200.254.1:0 POS5/0/0 (ldp): xmit/recv LDP Id: 10.200.254.3:0; no route</p> <p>london#show mpls ldp discovery detail</p> <p>POS5/0/0 (ldp): xmit/recv Enabled: Interface config Hello interval: 5000 ms; Transport IP addr: 10.200.254.2 LDP Id: 10.200.254.3:0; no route to transport addr</p> <p><i>Check show ip route for LSP!</i></p>	<p>MPLS LDP bound IP addresses:</p>	<p>rtr#show mpls ldp neighbor detail</p> <p>...</p> <p>Addresses bound to peer LDP Ident: 10.200.254.2 10.200.210.2 10.200.218.2 10.200.211.1 10.200.215.1</p> <p>Peer holdtime: 180000 ms; KA interval: 60000 ms; Peer state: estab</p> <p>Remote LSR binds each local IGP prefix in the routing table and advertise them through LDP</p> <p>R1#show mpls ldp bindings</p> <p>lib entry: 10.200.211.0/24, rev 12 local binding: label: imp-null remote binding: lsr: 10.200.254.5:0, label: 18 remote binding: lsr: 10.200.254.3:0, label: imp-null</p>
<p>show mpls interfaces fastEthernet 2/6 detail:</p>	<p>RTR#show mpls interfaces fastEthernet 2/6 detail</p> <p>Interface FastEthernet2/6: IP labeling enabled LSP Tunnel labeling not enabled BGP labeling not enabled MPLS not operational MTU = 1500</p> <p>RTR(config)#interface FastEthernet2/6 RTR(config-if)#mpls mtu 1508</p> <p>RTR#show mpls interfaces fastEthernet 2/6 detail</p> <p>... MTU = 1508</p> <p><i>(system jumbo mtu)</i></p>	<p>Identifying the remote LSR and label space being used:</p>	<p>RTR#show mpls ldp discovery</p> <p>Local LDP Identifier: 10.200.254.2:0</p> <p>Discovery Sources: Interfaces: Ethernet0/1/4 (ldp): xmit/recv LDP Id: 10.200.254.1:0</p> <p>4 bytes identifies LSR uniquely</p> <p>2 Bytes identifies the label space, - if set to 0 platform-wide label space, - if set to 1, per interface label space is used.</p>	<p>show mpls ldp bindings:</p> <p>Check the local actions for a PFX:</p>	<p>RTR#show mpls ldp bindings</p> <p>lib entry: 10.200.210.0/24, rev 4</p> <p>lib entry: 10.200.211.0/24, rev 12 local binding: label: imp-null POP action remote binding: lsr: 10.200.254.5:0, label: 18 remote binding: lsr: 10.200.254.1:0, label: 32 remote binding: lsr: 10.200.254.3:0, label: imp-null</p> <p>lib entry: 10.200.254.1/32, rev 31 local binding: label: 24 SWAP action remote binding: lsr: 10.200.254.5:0, label: 22 remote binding: lsr: 10.200.254.1:0, label: imp-null remote binding: lsr: 10.200.254.3:0, label: 26</p>
<p>MPLS Maximum Receive Unit</p> <p>MRU:</p>	<p>- informs the LSR how big a received labeled packet of a certain FEC can be that can still be forwarded out of this LSR without fragmenting it.</p> <p>- value per FEC (or prefix) and not per interface!</p> <p>show mpls forwarding-table x.x.x.x detail MAC/Encaps=14/14, MRU=1512, Tag Stack{}</p>	<p>MPLS LDP session establishment:</p>	<p>1. Discover via LDP Hello's Multicast 224.0.0.2 UDP 646</p> <p>2. Attempt TCP connection port 646 LDP initialization messages containing: - Timer values - Label distribution method - VPI / VCI ranges (LC-ATM) - DLCI ranges for Frame-Relay</p> <p>If peers disagree they will retry. To throttle: mpls ldp backoff <initial-backoff> <max></p> <p>Attempts will increase exponential from 5 sec 120 sec</p>	<p>Show mpls ip binding</p> <p>Check for binding inuse:</p>	<p>london#show mpls ip binding</p> <p>Local binding</p> <p>10.200.254.1/32 in label: 24 out label: 22 lsr: 10.200.254.5:0 out label: imp-null lsr: 10.200.254.1:0 inuse out label: 26 lsr: 10.200.254.3:0</p> <p>Traffic being forwarded to the ones marked as inuse</p>
<p>MPLS LDP functions:</p>	<p>MPLS LDP functions are:</p> <ul style="list-style-type: none"> - Discovery of LSRs running LDP - Session establishment and maintenance - Advertising of label mappings - Notifications (Malformed PDU, Unknown TLV, keepalive timer expired, unilateral session shutdown, initialization messages, internal errors, loop detection, other events) 	<p>Show mpls ldp neighbor x.x.x.x detail:</p> <p>Check which end is the TCP server of the LDP session:</p>	<p>show mpls ldp neighbor 10.200.254.5 detail</p> <p>Peer LDP Ident: 10.200.254.5:0; Local LDP Ident 10.200.254.2:0</p> <p>TCP connection: 10.200.254.5.11537 - 10.200.254.2.646</p> <p>State: Oper; Msgs sent/rcvd: 16/19; Down: Last TIB rev sent 50</p> <p>Up time: 00:00:36; UID: 9; Peer Id 1; LDP discovery sources: Ethernet0/1/2; Src IP addr: 10.200.215.2 holdtime: 15000 ms, hello interval: 5000 ms</p> <p>Addresses bound to peer LDP Ident: 10.200.254.5 10.200.215.2 10.200.216.1</p> <p>Peer holdtime: 180000 ms; KA interval: 60000 ms; Peer state: estab</p> <p><i>Remote side is the TCP server</i></p>	<p>MPLS LIB</p> <p>RIB</p> <p>LFIB</p> <p>LDP peers</p> <p>Output:</p>	<p>LIB</p> <p>show mpls ldp bindings 1.1.1.1/32</p> <p>Local bind Label 20</p> <p>Remote Bind: lsr 3.3.3.3 label 18</p> <p>RIB</p> <p>show ip route 1.1.1.1/32</p> <p>Next hop IP via Interface 2.2.2.2 [POS1]</p> <p>LDP peers</p> <p>show mpls ldp neighbor pos1</p> <p>Peer LDP Ident 3.3.3.3, Local 2.2.2.2</p> <p>Address bound to peer LDP Ident: 2.2.2.2</p> <p>LFIB</p> <p>show mpls forwarding-table 1.1.1.1</p> <p>Local Out PFX interface</p> <p>20 18 1.1.1.1 [POS1]</p>

Help me create more flashcards:

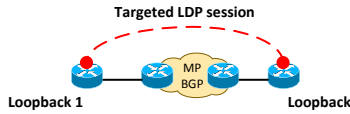
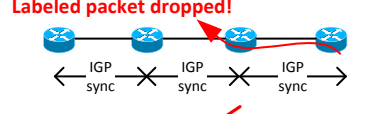
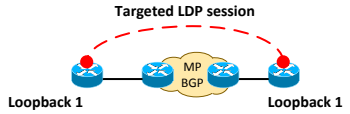
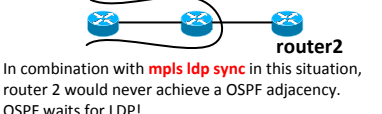
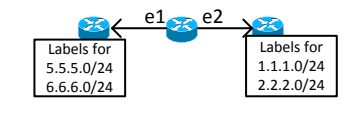
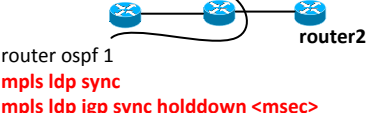
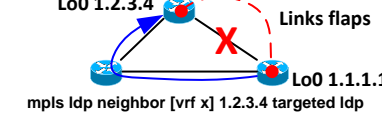
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Donate

Thanks for appreciating my efforts


Colin

<p>MPLS LDP label withdraw behaviour:</p>	<p>OLD DEFAULT: MPLS LDP does not withdraw labels from a neighbor it has learned the prefix/label. (no split-horizon) To configure this behaviour:</p> <pre>mpls ldp neighbor x.x.x.x implicit-withdraw</pre> <pre>debug mpls messages received</pre> <pre>debug mpls ldp bindings</pre>	<p>MPLS LDP autoconfig:</p>	<pre>router ospf 1 mpls ldp autoconfig area 0</pre> <p>R1#show mpls interfaces detail Interface Ethernet3/1: IP Labeling enabled (ldp): Interface config IGP config LSP Tunnel labeling enabled BGP labeling not enabled MPLS operational</p> <p>R1#show mpls ldp discovery detail Ethernet3/1 (ldp): xmit/rcv Enabled: Interface config, IGP config;</p>	<p>Verifying CEF switching:</p>	<pre>Enable cef: ip cef</pre> <pre>Clear ip cache (fast switched cache)</pre> <pre>show ip cache (cef switched packets will not be displayed)</pre> <pre>Show interface stats includes fast-switched and cef switched</pre> <pre>no ip route-cache cef</pre> <pre>show interface stats</pre> <pre>show ip cache (will show entries, switch failed back to fast-switching)</pre> <pre>no ip route-cache (packets will no be process-switched)</pre>															
<p>MPLS Targeted LDP Sessions:</p>	 <pre>mpls ldp neighbor [vrf x] 1.2.3.4 targeted ldp</pre> <p>For LDP neighbors that are NOT directly connected</p> <p>Targeted LDP can improve the label convergence time, in situations with flapping links.</p>	<p>MPLS LDP-IGP Synchronization Concept:</p>	<p>Labeled packet dropped!</p>  <p>LDP session is down, while IGP is up!</p> <p>OSPF will not form an adjacency until a link if the LDP session is not established first across that link. (No hello's on link)</p> <p>OSPF will announce the link with a max metric of 65536 or 0xFFFF until synchronization is achieved.</p> <pre>mpls ldp sync</pre>	<p>Switching methods:</p> <p>- CEF - fast-switching - process switching:</p>	<pre>ip cef</pre> <p>↓</p> <pre>no ip route-cache cef</pre> <p>Packets are now fast-switched</p> <p>↓</p> <pre>no ip route-cache</pre> <p>Packets are now process switched CPU</p>															
<p>MPLS LDP Targeted Hello Accept:</p>	 <pre>access-list standard ACCEPT-LDP permit 2.2.2.2</pre> <pre>mpls label protocol ldp mpls ldp router-id Loopback0 force</pre> <pre>mpls ldp discovery targeted-hello accept from ACCEPT-LDP</pre> <p>(Allow hello's only from 2.2.2.2)</p>	<p>MPLS LDP-IGP Synchronization Spoke problem:</p>	 <p>In combination with mpls ldp sync in this situation, router 2 would never achieve a OSPF adjacency. OSPF waits for LDP!</p> <p>Solution: Configure hold-down timer for the synchronization.</p> <pre>Router2 mpls ldp igp sync holddown <msec></pre> <p>Disable sync per interface no mpls ldp igp sync</p> <p>If the hold-down timer expires before the LDP session is established, the OSPF adjacency is built anyway.</p>	<p>IOS switching path that a packet takes</p> <table border="1"> <thead> <tr> <th>Incoming Interface</th> <th>Outgoing Interface</th> <th>Switching Method</th> </tr> </thead> <tbody> <tr> <td>CEF</td> <td>Process</td> <td>CEF</td> </tr> <tr> <td>Process</td> <td>CEF</td> <td>Fast</td> </tr> <tr> <td>Process</td> <td>Fast Switching (IP route cache)</td> <td>Fast Switching</td> </tr> <tr> <td>CEF</td> <td>Fast Switching</td> <td>CEF</td> </tr> </tbody> </table>	Incoming Interface	Outgoing Interface	Switching Method	CEF	Process	CEF	Process	CEF	Fast	Process	Fast Switching (IP route cache)	Fast Switching	CEF	Fast Switching	CEF	
Incoming Interface	Outgoing Interface	Switching Method																		
CEF	Process	CEF																		
Process	CEF	Fast																		
Process	Fast Switching (IP route cache)	Fast Switching																		
CEF	Fast Switching	CEF																		
<p>MPLS Controlling the advertisement of Labels via LDP Advertise-labels</p>	<pre>mpls ldp advertise-labels [vrf x] interface <fa0/x> for <prefix-acl> to <peer-list></pre> <div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid black; padding: 2px;"> <p>Prefix-acl</p> <pre>Permit 1.1.1.0/24</pre> <pre>Permit 192.168.20.2</pre> </div> <div style="border: 1px solid black; padding: 2px;"> <p>Peer-list</p> <pre>Permit 1.1.1.0/24</pre> <pre>Permit 192.168.20.2</pre> </div> </div>  <pre>show mpls ldp bindings advertisement-acls</pre>	<p>MPLS LDP-IGP Synchronization Show commands:</p>	 <pre>router ospf 1 mpls ldp sync mpls ldp igp sync holddown <msec></pre> <pre>show ip ospf mpls ldp interface serial 4/0</pre> <p>....</p> <p>LDP is not configured though LDP autoconfig LDP-IGP Synchronization : Required Holddown timer is not configured</p> <pre>show mpls ldp igp sync <interface></pre>	<p>BGP / IGP / CEF recursion:</p> <pre>BGP Prefix 10.99.1.1/32</pre> <pre>BGP nexthop 10.200.254.4</pre> <pre>IGP next-hop 10.200.200.2</pre>	<pre>RTR#show ip bgp 10.99.1.1</pre> <p>BGP routing table entry for 10.99.1.1/32, version 13 Paths: (1 available, best #1, table Default-IP-Routing-Table) Not advertised to any peer Local 10.200.254.4 (metric 85) from 10.200.254.4 (10.200.254.4) Origin IGP, metric 0, localpref 100, valid, internal, best</p> <pre>RTR#show ip cef 10.200.254.4</pre> <p>10.200.254.4/32 nexthop 10.200.200.2 Ethernet0/0/0 label 23</p> <hr/> <pre>RTR#show ip cef 10.99.1.1</pre> <p>10.99.1.1/32 nexthop 10.200.200.2 Ethernet0/0/0 label 23</p>															
<p>Controlling MPLS LDP label advertisement Configuration: Advertise-labels</p>	<pre>no mpls ldp advertise-labels</pre> <pre>mpls ldp advertise-labels for PFX to PEER-List1</pre> <pre>mpls ldp advertise-labels for PFX-2 to PEER-List2</pre> <pre>access-list PFX permit 10.200.254.4</pre> <pre>access-list PFX permit 10.200.254.3</pre> <pre>access-list PFX deny any</pre> <pre>access-list PEER permit 10.200.254.5</pre> <pre>access-list PEER deny any</pre> <p>Only prefixes 10.200.254.3/32 and 10.200.254.4/32 are advertised to LDP peer 10.200.254.5</p> <pre>R1#show mpls ldp bindings advertisement-acls</pre> <p>Advertisement spec: Prefix acl = PFX; Peer acl = PEER</p> <pre>lib entry: 10.200.211.0/24, rev 15</pre> <pre>lib entry: 10.200.254.3/32, rev 21</pre> <pre>Advert acl(s): Prefix acl 1; Peer acl 2</pre> <pre>lib entry: 10.200.254.4/32, rev 2</pre> <pre>Advert acl(s): Prefix acl 1; Peer acl 2</pre>	<p>MPLS LDP-IGP synchronization Show and debugs:</p>	<pre>mpls ldp igp sync holddown 3000</pre> <pre>debug mpls ldp sync [int x] peer-acl ACL</pre> <p>IGP forms an adjacency anyway to give LDP the opportunity to build an LDP session across that link</p> <pre>%LINK-3-UPDOWN: Interface Serial4/0, changed state to up</pre> <pre>LDP-SYNC: Se4/0: queue swit_down, set INTFADDR_PENDING.</pre> <pre>LDP-SYNC: Se4/0: process swit_down, clear INTFADDR_PENDING.</pre> <pre>%LINEPROTO-5-UPDOWN: Line protocol on Interface Serial4/0, changed state to up</pre> <pre>%OSPF-5-ADJCHG: Process 1, Nbr 10.200.254.4 on Serial4/0 from LOADING to FULL, Loading Done</pre> <pre>LDP-SYNC: Se4/0: No session or session has not send initial update, ignore adjoining event.</pre> <pre>%LDP-5-NBRCHG: LDP Neighbor 10.200.254.4:0 is UP</pre> <pre>LDP-SYNC: Se4/0: session 10.200.254.4:0 came up, sync_achieved up</pre> <pre>LDP-SYNC: Se4/0: OSPF 1: notify status (required, achieved, no delay, holddown 3000)</pre> <pre>OSPF: schedule to build router LSA after notification from LDP</pre>	<p>CEF load-sharing</p>	<pre>RTR#show cef interface ethernet 1/2</pre> <p>Per packet load-sharing is disabled</p> <pre>RTR(config)#int et 1/2</pre> <pre>RTR(config-if)#ip load-sharing per-packet</pre> <pre>RTR#show cef interface ethernet 1/2</pre> <p>Per packet load-sharing is enabled</p> <pre>restore CEF default:</pre> <p>ip load-sharing per-destination</p>															
<p>MPLS LDP Inbound Label Binding Filtering Accept ACLs</p>	<pre>mpls ldp neighbor x.x.x.x accept ACCEPT-LBLS</pre> <pre>mpls ldp neighbor x.x.x.x [vrf x] accept ACCEPT-LBLS</pre> <pre>Access-list standard ACCEPT-LBLS permit 1.2.3.4</pre> <pre>R1#show mpls ldp bindings</pre> <pre>lib entry: 10.200.254.2/32, rev 69</pre> <pre>local binding: label: 27</pre> <pre>lib entry: 10.200.254.3/32, rev 71</pre> <pre>local binding: label: 19</pre> <pre>remote binding: lsr: 1.2.3.4:0, label: 21</pre> <p>No label visible for PFX which are not permitted by the ACL:</p>	<p>MPLS LDP Session protection</p>	 <pre>mpls ldp neighbor [vrf x] 1.2.3.4 targeted ldp</pre> <pre>mpls ldp session protection [vrf x] for ACL</pre> <p>LDP session is kept as long as there is an alternative path between the LSRs</p> <pre>show mpls ldp neighbor serial 4/0 detail</pre> <pre>...LDP Session Protection enabled, state: Ready</pre> <pre>Duration infinite</pre> <pre>...LDP Session Protection enabled, state: Protecting</pre> <pre>Duration infinite</pre> <p>standby Link down, triggered</p>	<p>show ip cef 10.200.254.4 internal</p> <p>Output:</p>	<pre>paris#show ip cef 10.200.254.4 internal</pre> <p>10.200.254.4/32, version 26, epoch 0, RIB, refcount 5, per-destination sharing 16 hash buckets <0> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <1> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640 <2> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <3> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640 <4> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <5> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640 <6> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <7> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640 <8> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <9> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640 <10> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <11> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640 <12> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <13> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640 <14> IP adj out of Ethernet1/2, addr 10.200.201.2 6346B8C0 <15> IP adj out of Ethernet1/3, addr 10.200.203.2 6346C640</p>															

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts



Colin

<p>show ip cef exact-route <SRC-IP> <DST-IP></p> <p>Output:</p>	<pre>show ip cef exact-route SOURCE-IP DESTINATION-IP RTR#show ip cef exact-route 10.200.254.1 10.200.254.4 10.200.254.1 -> 10.200.254.4 => IP adj out of Ethernet1/2, addr 10.200.201.2 RTR#show ip cef exact-route 10.200.1.2 10.200.254.4 10.200.1.2 -> 10.200.254.4 => IP adj out of Ethernet1/3, addr 10.200.203.2</pre>	<p>VPNv4</p> <p>RR Group config:</p>	<p>-subdivide the vpnv4 routes into groups - increases scalability</p> <p>On Route Reflector 1: router bgp X address-family vpnv4 bgp rr-group <1-500> {ext-com-list}</p> <p>ip extcommunity-list 1 permit rt 1:1 ip extcommunity-list 1 deny rt 1:2 ip extcommunity-list 1 permit rt 1:3 ip extcommunity-list 1 deny rt 1:2</p> <p>On Route Reflector 2: ip extcommunity-list 1 deny rt 1:1 ip extcommunity-list 1 permit rt 1:2</p>	<p>BGP vpnv4</p> <p>OSPF metric propagation</p>	<pre>RTR#show ip ospf 42 Routing Process "ospf 42" with ID 10.99.1.1 Domain ID type 0x0005, value 0.0.0.42 Connected to MPLS VPN Superbackbone, VRF cust-one RTR#show ip bgp vpnv4 rd 1:1 10.200.200.1 BGP routing table entry for 1:1:10.200.200.1/32, version 5649 Paths: (1 available, best #1, table cust-one) Not advertised to any peer Local 10.200.254.2 (metric 3) from 10.200.254.2 (10.200.254.2) Origin incomplete, metric 10, localpref 100, valid, internal, best Extended Community: RT:1:1 OSPF DOMAIN-ID:0x0005:0x0000002A0200 OSPF RT:0.0.0.5:1 OSPF ROUTER ID:10.99.1.1:1281, mpls labels in/out nolabel/18 Area:route-type:option</pre>																					
<p>show ip cef vrf x <prefix> [detail]</p> <p>Top and bottom MPLS labels:</p> <p>Output:</p>	<pre>RTR#show ip cef vrf cust-one 10.100.103.2 10.100.103.2/32 nexthop 10.200.200.2 Ethernet0/0/0 label 23 21 RTR#show ip cef vrf cust-one 10.100.103.2 detail 10.100.103.2/32, epoch 5 recursive via 10.200.254.4 label 21 nexthop 10.200.200.2 Ethernet0/0/0 label 23 Top Label: 23 Bottom Label: 23</pre>	<p>VPNv4</p> <p>RR Group config:</p>	<pre>ip extcommunity-list 1 permit rt 1:1 ip extcommunity-list 1 deny rt 1:2 ip extcommunity-list 1 permit rt 1:3 ip extcommunity-list 1 deny rt 1:2 ip extcommunity-list 1 deny rt 1:1 ip extcommunity-list 1 permit rt 1:2 ip extcommunity-list 1 deny rt 1:3 ip extcommunity-list 1 permit rt 1:2</pre>	<p>BGP vpnv4</p> <p>Extended communities for OSPF</p>	<p>Domain-ID NOT = OSPF process ID Result: External LSA Type 5</p> <p>Domain-ID = OSPF process ID Result: internal route</p> <pre>router bgp address-family xy domain-id 77</pre>																					
<p>show mpls forwarding-table labels label exact-path ipv4 source-address destination-address</p> <p>Output:</p>	<pre>horizon#show mpls forwarding-table Local Outgoing Prefix Bytes tag Outgoing Next Hop tag tag or VC or Tunnel Id switched interface 17 Pop tag 10.200.254.3/32 252 Et1/3 10.200.203.2 Pop tag 10.200.254.3/32 0 Et1/2 10.200.201.2 18 16 10.200.254.4/32 10431273 Et1/2 10.200.201.2 16 10.200.254.4/32 238 Et1/3 10.200.203.2</pre> <p>show mpls forwarding-table labels 18 exact-path ipv4 <SRC> <DST></p> <p>show mpls forwarding-table labels 18 exact-path ipv4 1.1.1.1 2.2.2.2 1.1.1.1 -> 2.2.2.2: Eth1/3 (next-hop x.x.x.x) Label Stack: 16</p>	<p>Life of a IPv4 Packet across a MPLS VPN Backbone:</p>		<p>VPNv4</p> <p>OSPF sham Links explained:</p>	<p>By default only Type 3</p> <p>Sham-link Type 1,2 By default only Type 3</p> <p>Lo0 1.1.1.1 Lo0 2.2.2.2</p> <pre>router ospf X vrf BLA area 0 sham-link <1.1.1.1> <2.2.2.2></pre>																					
<p>show cef table</p> <p>Output:</p>	<pre>london#show cef table Global information: MTRIE information: TAL: node pools: pool[C/8 bits]: 45 allocated (0 failed), 46800 bytes (3 recount) 3 active IPv4 tables out of a maximum of 2048 VRF Prefixes Memory Flags Default 38 23620 LCS cust-one 11 14680 LCS cust-one-ipv4 11 14680 LCS 1 active IPv6 table out of a maximum of 1 VRF Prefixes Memory Flags Default 0 60</pre>	<p>show ip cef vrf x <PFX></p> <p>show ip bgp vpnv4 rd 1:1 <PFX></p> <p>Explained:</p>	<pre>inPE#show ip cef vrf cust-one 10.10.100.1/32 detail 10.10.100.1/32, epoch 0 recursive via 10.200.254.2 label 30 nexthop 10.200.214.1 POS0/1/0 label 16 VPN label 30 IGP label to loopback 16</pre> <pre>inPE#show ip bgp vpnv4 rd 1:1 10.10.100.1 BGP routing table entry for 1:1:10.10.100.1/32, version 81 Paths: (1 available, best #1, table cust-one) Not advertised to any peer Local 10.200.254.2 (metric 3) from 10.200.254.2 (10.200.254.2) Origin incomplete, metric 1, localpref 100, valid, internal, best Extended Community: RT:1:1, mpls labels in/out nolabel/30 BGP loopback IGP next hop</pre>	<p>Verifying OSPF sham links</p>	<pre>RTR#show ip ospf 42 neighbor Neighbor ID Pri State Dead Time Address Interface 10.200.200.1 1 FULL/DR 00:00:35 10.10.2.1 Ethernet0/1/2 10.99.1.2 0 FULL/- - 10.99.1.2 OSPF_SL2</pre> <p>RTR#show ip ospf 42 sham-links Sham Link OSPF_SL2 to address 10.99.1.2 is up Area 0 source address 10.99.1.1 Run as demand circuit DoNotAge LSA allowed. Cost of using 10 State POINT_TO_POINT. Timer intervals configured, Hello 10, Dead 40, Wait 40, Hello due in 00:00:03 Adjacency State FULL (Hello suppressed)</p>																					
<p>Show ip cef table IPv4 [customer-x]</p> <p>Output:</p>	<pre>RTR#show cef table IPv4 cust-one Table: IPv4:cust-one (id 1) ref count: 3 reset count: 1 flags (0x01): LCS smp allowed: yes default network: none route count: 11 route count (fwd): 11 route count (non-fwd): 0 Database epoch: 8 (11 entries at this epoch) Subblocks: None RTR#show cef table IPv4 Default Table: IPv4:Default (id 0) ref count: 17 reset count: 1</pre>	<p>Possible OSPF MPLS VPN Scenarios</p> <p>Superbackbone diagrams:</p>		<p>Down-Bit and Domain Tag:</p> <p>MP-BGP --> OSPF = down bit is set</p> <p>OSPF --> MP-BGP = domain-id is set</p>	<p>area:type:option [3,5] [E1,E2]</p>																					
<p>Show ip bgp vpnv4 rd 1:1 labels</p> <p>Output explained:</p>	<pre>RTR#show ip bgp vpnv4 rd 1:1 labels Network Next Hop In label/Out label Route Distinguisher: 1:1 (cust-one) 10.10.2.0/24 10.200.254.2 29/36 10.10.4.0/24 0.0.0.0 26/nolabel 10.10.4.2/32 0.0.0.0 37/nolabel 10.99.1.2/32 0.0.0.0 27/nolabel 10.10.100.1/32 10.200.254.2 32/35 10.10.100.3/32 10.10.4.2 38/exp-null 10.88.1.1/32 10.200.254.2 34/34 10.99.1.1/32 10.200.254.2 28/33 10.200.200.1/32 10.200.254.2 30/32 Prefixes with next hop 0.0.0.0, have no outgoing label, learned from VRF interface, should be forwarded unlabeled towards CE Labeled traffic</pre>	<p>Internal OSPF Routes Across MPLS VPN Backbone</p>	<p>OSPF Process ID on PE's the same or domain-id X used == type 3</p> <p>Router bgp X address-family X domain-id xx</p>	<p>BGP extended communities for EIGRP</p>	<table border="1"> <thead> <tr> <th>Type</th> <th>Usage</th> <th>Value</th> </tr> </thead> <tbody> <tr> <td>0x8800</td> <td>General route info</td> <td>Flags + Tag</td> </tr> <tr> <td>0x8801</td> <td>Route metric info and Autonomous System</td> <td>Autonomous System + Delay</td> </tr> <tr> <td>0x8802</td> <td>Route metric info</td> <td>Reliability, Hop Count, Bandwidth</td> </tr> <tr> <td>0x8803</td> <td>route metric info</td> <td>Reserved field, Load, MTU</td> </tr> <tr> <td>0x8804</td> <td>External route info</td> <td>Remote Autonomous System, Remote ID</td> </tr> <tr> <td>0x8805</td> <td>External route info</td> <td>Remote Protocol, Remote metric</td> </tr> </tbody> </table>	Type	Usage	Value	0x8800	General route info	Flags + Tag	0x8801	Route metric info and Autonomous System	Autonomous System + Delay	0x8802	Route metric info	Reliability, Hop Count, Bandwidth	0x8803	route metric info	Reserved field, Load, MTU	0x8804	External route info	Remote Autonomous System, Remote ID	0x8805	External route info	Remote Protocol, Remote metric
Type	Usage	Value																								
0x8800	General route info	Flags + Tag																								
0x8801	Route metric info and Autonomous System	Autonomous System + Delay																								
0x8802	Route metric info	Reliability, Hop Count, Bandwidth																								
0x8803	route metric info	Reserved field, Load, MTU																								
0x8804	External route info	Remote Autonomous System, Remote ID																								
0x8805	External route info	Remote Protocol, Remote metric																								

Help me create more flashcards:

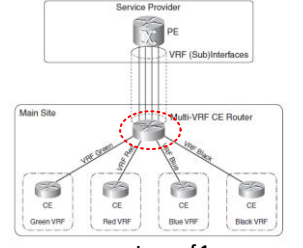
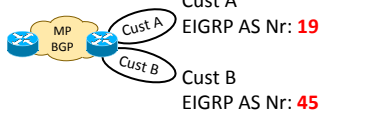
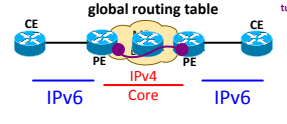
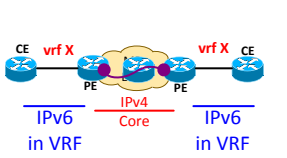
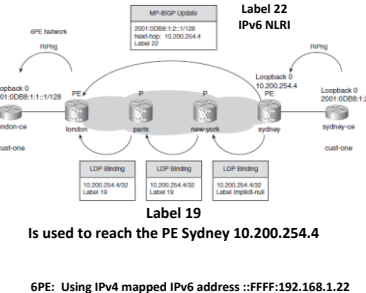
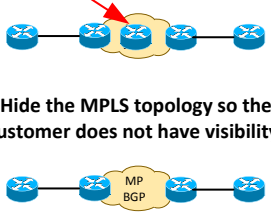
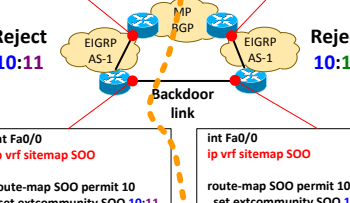
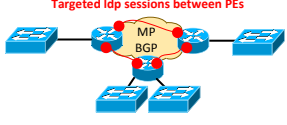
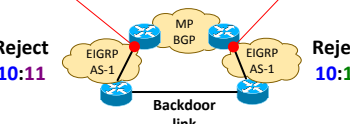
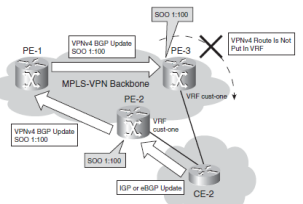
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts



Colin

<p>BGP Extended communities for EIGRP</p>	<p>Internal</p> <pre>RTR#show ip bgp vpnv4 all 10.10.100.1 BGP routing table entry for 1:1:10.10.100.1/32, version 28 ... Extended Community: RT:1:1 Cost:pre-bestpath:128:409600 0x8800:32768:0 0x8801:42:153600 0x8802:65281:256000 0x8803:65281:1500, mpls labels in/out 22/nolabel</pre> <p>BGP extcommunities for EIGRP</p> <p>External</p> <pre>RTR#show ip bgp vpnv4 all 10.200.200.1 BGP routing table entry for 1:1:10.200.200.1/32, version 91 Extended Community: RT:1:1 Cost:pre-bestpath:129:409600 0x8800:0:0 0x8801:42:153600 0x8802:65281:256000 0x8803:65281:1500 0x8804:0:168453121 0x8805:11:0, mpls labels in/out nolabel31</pre>	<p>Capability vrf-lite on OSPF enabled CE</p> <p>Explained:</p>	 <p>router ospf 1 capability vrf-lite</p> <p>CE performs down-bit / domain-tag check, disabled by capability vrf-lite</p>	<p>Upgrading ip vrf x</p> <p>To ip vrf definition x:</p>	<p>Previous:</p> <pre>ip vrf ABC rd 1:1 route-target export 1:1 route-target import 1:1</pre> <p>vrf upgrade-cli multi-af-mode common-policies</p> <p>Are you sure ? [yes]: yes Number of VRFs upgraded: 1</p> <p>vrf definition ABC</p> <pre>rd 1:1 route-target export 1:1 route-target import 1:1 address-family ipv4 exit-address-family</pre>
<p>EIGRP VRF configuration:</p>	 <pre>router eigrp X address-family ipv4 vrf cust-A autonomous-system 19 address-family ipv4 vrf cust-B autonomous-system 45</pre>	<p>Differences between 6PE and 6VPE</p> <p>Picture:</p>	<p>6PE</p>  <p>6VPE</p> 	<p>show bgp vpnv6 unicast vrf x <PFX></p> <p>Output</p> <p>RD explained:</p>	<pre>R1#show bgp vpnv6 unicast vrf A 2001:DB8:1:2::1/128 BGP routing table entry for [1:1]2001:DB8:1:2::1/128, version 5 RD for IPv6 in [1:1]</pre>
<p>Pre-Bestpath POI</p>	<p>Cost community ID: 0 - 255 Cost value: 0 - 4294967295</p> <p>EIGRP internal community ID: 128 EIGRP external community ID: 129</p> <p>The lower cost ID is more preferred.</p> <p>The lower cost ID is more preferred.</p> <p>Format: Cost-POI-com-ID:value</p> <pre>route-map X permit 10 set extcommunity cost 1 100</pre> <p>BGP considers POI over all other regular BGP comparison steps.</p> <pre>RTR#show ip bgp vpnv4 all 10.10.100.1 BGP routing table entry for 1:1:10.10.100.1/32, version 28 ... Extended Community: RT:1:1 Cost:pre-bestpath:128:409600 0x8800:32768:0</pre>	<p>MPLS 6PE</p> <p>Routing and Label Distribution:</p>	 <p>Is used to reach the PE Sydney 10.200.254.4</p> <p>6PE: Using IPv4 mapped IPv6 address ::FFFF:192.168.1.22</p>	<p>Hide the MPLS topology so the customer does not have visibility:</p> 	<p>no mpls ip propagate-ttl</p> <p>no mpls ip propagate-ttl forwarded</p> <p>Disabling TTL propagation for customers only, ISP internal still has visibility</p>
<p>EIGRP PE-CE Backdoor links</p> <p>Extcommunity SOO</p>	<pre>int Ser0/0 ip vrf forwarding AA ip vrf sitemap SOO route-map SOO permit 10 set extcommunity SOO 10:11</pre>  <pre>int Fa0/0 ip vrf sitemap SOO route-map SOO permit 10 set extcommunity SOO 10:11</pre>	<p>6PE / MPLS PE configuration:</p>	<pre>int s0/0 ip address 2001::1/64 ipv6 rip customer-1 enable ipv6 enable int fa0/0 ip address 10.0.0.1 255.255.255.0 mpls ip router bgp 1 address-family ipv6 neighbor 1.1.1.1 activate neighbor 1.1.1.1 send-community both neighbor 1.1.1.1 send-label redistribute rip customer-1</pre>	<p>Basic VPLS</p> <p>PE config:</p>	 <p>l2 vfi <name> manual</p> <p>vpn id <number></p> <p>neighbor 1.2.3.4 encapsulation mpls</p> <p>interface X xconnect vfi <name> L2tunnel-protocol [cdp,stp,vtp,...]</p>
<p>SOO set for an EIGRP route:</p>	 <pre>PE-1#show ip eigrp vrf cust-one topology 10.10.100.3 255.255.255.255 IP-EIGRP (AS 42): Topology entry for 10.10.100.3/32 Extended Community: SoO:10:10</pre> <pre>R6#show bgp vpnv4 unicast vrf VPN_A 150.1.8.0 Extended Community: SoO:10:10 RT:100:1 Cost:prebestpath: 128:156160...</pre>	<p>Verifying 6PE operation:</p>	<pre>RTR#show bgp ipv6 unicast neighbors ... Neighbor capabilities: ... Address family IPv6 Unicast: advertised and received ipv6 MPLS Label capability: advertised and received</pre> <pre>RTR#show bgp ipv6 unicast 2001:DB8:1:2::1/128 ... Local ::FFFF:10.200.254.4 (metric 4) from 10.200.254.4 (10.200.254.4) Origin incomplete, metric 2, localpref 100, valid, internal, best, mpls labels in/out nolabel/22</pre> <pre>RTR#show bgp ipv6 unicast labels Network Next Hop In label/Out label 2001:DB8:1:1::1/128 :: 29/nolabel pop 2001:DB8:1:2::1/128 ::FFFF:10.200.254.4 nolabel/22 push</pre>	<p>VPLS</p> <p>Troubleshooting commands:</p>	<p>VPLS-PE-1#show vfi cust-one</p> <p>VFI name: cust-one, state: up Local attachment circuits: Vlan111 Neighbors connected via pseudowires: 10.100.100.2 10.100.100.3</p> <p>show mpls l2transport summary</p> <p>VPLS-PE-1#show mpls l2transport vc 1 detail</p> <p>Local interface: VFI cust-one up Destination address: 10.100.100.2, VC ID: 1, VC status: up Tunnel label: 17, next hop point2point Output interface: POS5/1, imposed label stack (17 18) Create time: 2d17h, last status change time: 01:04:54 Signaling protocol: LDP, peer 10.100.100.2:0 up MPLS VC labels: local 16, remote 18 Group ID: local 0, remote 0 ...</p>
<p>SOO config explained:</p> 	<pre>SOO route-map route-map CUST-A permit 10 set extcommunity soo 1:100</pre> <p>Applying SOO route-map for BGP</p> <pre>router bgp 1 address-family ipv4 vrf CUST-A neighbor 1.2.3.4 route-map CUST-A in</pre> <p>Applying SOO on the VRF interface</p> <pre>interface fa0/0 ip vrf sitemap CUST-A</pre> <p>Applying SOO route-map for static routes</p> <pre>router bgp 1 address-family ipv4 vrf CUST-A Redistribute static route-map CUST-A</pre>	<p>6VPE PE configuration:</p>	<pre>ipv6 unicast-routing ipv6 cef vrf definition CUST-1 rd 1:1 address-family ipv6 route-target export 1:1 route-target import 1:1 exit-address-family mpls ldp route-id loopback0 force int fa0/0 ip address 10.0.0.1/24 mpls ip int ser0/0 vrf forwarding CUST-1 ipv6 address 2001::1/64 ipv6 enable</pre> <p>Facing P router</p> <p>Facing vrf customer</p> <pre>router bgp 1 neighbor 1.1.1.1 update-source Loopback 0 address-family vpnv6 neighbor 1.1.1.1 activate neighbor 1.1.1.1 send-community both address family ipv6 vrf CUST-1 neighbor 2001::2 remote-as XXX</pre>	<p>EoMPLS Carrying Simple Ethernet</p>	<p>pseudowire-class one encapsulation mpls</p> <pre>interface FastEthernet9/0/0 no ip address xconnect 10.200.254.4 2000 pw-class one</pre> <pre>PE1#show mpls l2transport vc 2000 Local intf Local circuit Dest address VC ID Status Fa9/0/0 Ethernet 10.200.254.4 2000 UP</pre>

Help me create more flashcards:

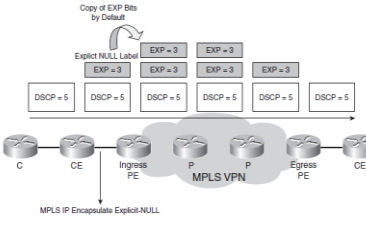
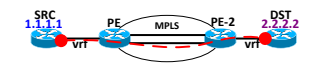
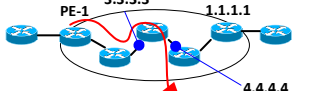


Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts



Colin

<p>MPLS Diffserver Tunneling Models:</p> <p>Uniform Model</p> <p>Short Pipe Model</p> <p>Pipe Model</p>	<p>Uniform Model Customer sets DSCP, copied into EXP field of label</p> <p>Short Pipe Model EXP bits set according to Service Provider's policy Scheduling / Discarding based on DSCP</p> <p>Pipe Model EXP bits set according to Service Provider's policy Scheduling / Discarding based on EXP</p>	<p>MPLS LSP Ping:</p>	<p>- tests one particular FEC - uses UDP port 3503 - LSR never forwards such a packet if LSP is broken</p> <p>Reply Modes: Meaning: 1 Do not reply 2 Reply via IPv4/IPv6 UDP packet 3 Reply via IPv4/IPv6 UDP packet Router Alert 4 Reply via an application-level control channel</p> <p>ping mpls ipv4 1.2.3.4/32 [verbose] ping mpls pseudowire ping mpls traffic-eng</p> <p>ping mpls ipv4 1.2.3.4/32 destination 1.1.1.1 1.1.1.20 repeat 1 destination range</p>	<p>L2TPv2</p>	<p>- Uses IP protocol 115 or UDP packets</p>																																																						
<p>Explicit Null label on CE Router</p>	 <p>Explicit null which means penultimate hop router does not pop the label. Sends with label value of 0 but with other fields including EXP bits intact, QoS is preserved.</p> <p>mpls ip encapsulate explicit-null</p>	<p>MPLS echo packet format:</p>	<table border="1"> <tr> <td colspan="2">Version Number</td> <td colspan="2">Global Flags</td> </tr> <tr> <td>Message Type</td> <td>Reply Mode</td> <td>Return Code</td> <td>Return Subcode</td> </tr> <tr> <td colspan="4">Sender's Handle</td> </tr> <tr> <td colspan="4">Sequence Number</td> </tr> <tr> <td colspan="4">Timestamp Sent (seconds)</td> </tr> <tr> <td colspan="4">Timestamp Sent (microseconds)</td> </tr> <tr> <td colspan="4">Timestamp received (seconds)</td> </tr> <tr> <td colspan="4">Timestamp received (microseconds)</td> </tr> <tr> <td colspan="4">TLVs ...</td> </tr> </table>	Version Number		Global Flags		Message Type	Reply Mode	Return Code	Return Subcode	Sender's Handle				Sequence Number				Timestamp Sent (seconds)				Timestamp Sent (microseconds)				Timestamp received (seconds)				Timestamp received (microseconds)				TLVs ...				<p>L2TPv3</p> <p>Background Info:</p>	<p>- mandatory to specify a pseudowire class - specify the source IP address used for L2TPv3 session via the pw-command ip local interface <IFNAME>.</p> <p>- ip dfbit set, which avoids in-core fragmentation and performance degradation</p> <p>- copy the TOS byte from encapsulated packets to the tunnel, ip tos reflect or statically via ip tos <VALUE></p> <p>- ip local pmtu, allows PE to send CE icmp unreachable</p> <p>- ip dfbit set, drop if packets can not be fragmented</p>																		
Version Number		Global Flags																																																									
Message Type	Reply Mode	Return Code	Return Subcode																																																								
Sender's Handle																																																											
Sequence Number																																																											
Timestamp Sent (seconds)																																																											
Timestamp Sent (microseconds)																																																											
Timestamp received (seconds)																																																											
Timestamp received (microseconds)																																																											
TLVs ...																																																											
<p>Categories of QoS Information for Table-Map</p>	<table border="1"> <thead> <tr> <th>Packet Marking Category</th> <th>Value Range</th> </tr> </thead> <tbody> <tr> <td>Cos</td> <td>0 - 7</td> </tr> <tr> <td>IP Precedence</td> <td>0 - 7</td> </tr> <tr> <td>DSCP</td> <td>0 - 63</td> </tr> <tr> <td>QoS-Group</td> <td>0 - 99</td> </tr> <tr> <td>MPLS EXP imposition</td> <td>0 - 7</td> </tr> <tr> <td>MPLS EXP topmost</td> <td>0 - 7</td> </tr> </tbody> </table>	Packet Marking Category	Value Range	Cos	0 - 7	IP Precedence	0 - 7	DSCP	0 - 63	QoS-Group	0 - 99	MPLS EXP imposition	0 - 7	MPLS EXP topmost	0 - 7	<p>Troubleshooting Load balancing in MPLS LSPs with traceroute mpls:</p>	<pre>RTR#traceroute mpls ipv4 2.2.2.2/32 destination 127.0.0.1 127.0.0.10 Type escape sequence to abort. Destination address 127.0.0.1 0 10.200.210.1 MRU 1500 [Labels: 44 Exp: 0] 1 1 10.200.210.2 MRU 1500 [Labels: 37 Exp: 0] 68 ms 2 10.200.211.2 MRU 1504 [Labels: implicit-null Exp: 0] 100 ms 3 10.200.214.2 80 ms ... Destination address 127.0.0.4 0 10.200.210.1 MRU 1500 [Labels: 44 Exp: 0] 1 1 10.200.210.2 MRU 1500 [Labels: 41 Exp: 0] 80 ms 2 10.200.215.2 MRU 1504 [Labels: implicit-null Exp: 0] 60 ms 3 10.200.216.2 68 ms</pre> 	<p>L2TPv3</p> <p>Config:</p>	<p>pseudowire-class L2TPV3 encapsulation l2tpv3</p> <p>ip local interface Loopback0 ip pmtu ip dfbit set ip tos reflect</p> <p>default interface FastEthernet 0/1 interface FastEthernet 0/1 xconnect 1.2.3.4 100 encapsulation l2tpv3 pw-class L2TPV3</p>																																								
Packet Marking Category	Value Range																																																										
Cos	0 - 7																																																										
IP Precedence	0 - 7																																																										
DSCP	0 - 63																																																										
QoS-Group	0 - 99																																																										
MPLS EXP imposition	0 - 7																																																										
MPLS EXP topmost	0 - 7																																																										
<table border="1"> <thead> <tr> <th>To Packet-Marking Type</th> <th>From Packet-Marking Type</th> <th>To Packet-Marking Type</th> <th>From Packet-Marking Type</th> </tr> </thead> <tbody> <tr> <td>Precedence</td> <td></td> <td>Precedence</td> <td>CoS QoS group</td> </tr> <tr> <td>DSCP</td> <td></td> <td>DSCP</td> <td>CoS QoS group</td> </tr> <tr> <td>CoS</td> <td></td> <td>CoS</td> <td>Precedence DSCP</td> </tr> <tr> <td>MPLS EXP topmost</td> <td></td> <td>MPLS EXP topmost</td> <td>QoS group</td> </tr> <tr> <td>MPLS EXP imposition</td> <td></td> <td>MPLS EXP imposition</td> <td>Precedence DSCP</td> </tr> </tbody> </table>	To Packet-Marking Type	From Packet-Marking Type	To Packet-Marking Type	From Packet-Marking Type	Precedence		Precedence	CoS QoS group	DSCP		DSCP	CoS QoS group	CoS		CoS	Precedence DSCP	MPLS EXP topmost		MPLS EXP topmost	QoS group	MPLS EXP imposition		MPLS EXP imposition	Precedence DSCP		<p>MPLS LSP Ping</p> <p>return codes</p>	<table border="1"> <thead> <tr> <th>Value</th> <th>Meaning</th> </tr> </thead> <tbody> <tr><td>0</td><td>no return code</td></tr> <tr><td>1</td><td>malformed echo request received</td></tr> <tr><td>2</td><td>one or more TLVs misunderstood</td></tr> <tr><td>3</td><td>Replying router is egress for the FEC</td></tr> <tr><td>4</td><td>Replying router has no mapping for the FEC</td></tr> <tr><td>5</td><td>Downstream mapping mismatch</td></tr> <tr><td>6</td><td>Upstream interface index unknown</td></tr> <tr><td>7</td><td>Reserved</td></tr> <tr><td>8</td><td>Label-Switched at stack depth RSC</td></tr> <tr><td>9</td><td>Label-switched but no MPLS forwarding at stack</td></tr> <tr><td>10</td><td>Mapping for this FEC is not given label at stack</td></tr> <tr><td>11</td><td>No label entry at stack depth</td></tr> <tr><td>12</td><td>Protocol not associated with interface at FEC</td></tr> <tr><td>13</td><td>Premature termination of ping due to label stack shrinking to a single label</td></tr> </tbody> </table>	Value	Meaning	0	no return code	1	malformed echo request received	2	one or more TLVs misunderstood	3	Replying router is egress for the FEC	4	Replying router has no mapping for the FEC	5	Downstream mapping mismatch	6	Upstream interface index unknown	7	Reserved	8	Label-Switched at stack depth RSC	9	Label-switched but no MPLS forwarding at stack	10	Mapping for this FEC is not given label at stack	11	No label entry at stack depth	12	Protocol not associated with interface at FEC	13	Premature termination of ping due to label stack shrinking to a single label	<p>Verifying L2TPv3</p> <p>Config:</p>	<pre>R5#show l2tp session all L2TP Session Information Total tunnels 1 sessions 1 Session id 19547 is up, tunnel id 38503 Remote session id is 55660, remote tunnel id 34507 Locally initiated session Call serial number is 1694600001 Remote tunnel name is R6 Internet address is 150.1.6.6 Local tunnel name is R5 Internet address is 150.1.5.5 IP protocol 115 Session is L2TP signaled Session state is established, time since change 00:00:35 Session PMTU enabled, path MTU is 1496 bytes</pre>
To Packet-Marking Type	From Packet-Marking Type	To Packet-Marking Type	From Packet-Marking Type																																																								
Precedence		Precedence	CoS QoS group																																																								
DSCP		DSCP	CoS QoS group																																																								
CoS		CoS	Precedence DSCP																																																								
MPLS EXP topmost		MPLS EXP topmost	QoS group																																																								
MPLS EXP imposition		MPLS EXP imposition	Precedence DSCP																																																								
Value	Meaning																																																										
0	no return code																																																										
1	malformed echo request received																																																										
2	one or more TLVs misunderstood																																																										
3	Replying router is egress for the FEC																																																										
4	Replying router has no mapping for the FEC																																																										
5	Downstream mapping mismatch																																																										
6	Upstream interface index unknown																																																										
7	Reserved																																																										
8	Label-Switched at stack depth RSC																																																										
9	Label-switched but no MPLS forwarding at stack																																																										
10	Mapping for this FEC is not given label at stack																																																										
11	No label entry at stack depth																																																										
12	Protocol not associated with interface at FEC																																																										
13	Premature termination of ping due to label stack shrinking to a single label																																																										
<p>What does mpls ip ttl-expiration pop 1 Do?</p>	<p>Unlike the default, where if a TTL expires along the path, the packet is sent to the Egress PE.</p> <p>With mpls ip ttl-expiration pop 1</p> <p>If a ttl expires on a P router, the P router sends back a ICMP unreachable to the source router of the packet.</p>	<p>MPLS ping of a failed LSP</p> <p>MPLS traceroute of a failed LSP</p>	 <pre>PE-1#ping mpls ipv4 1.1.1.1/32 verbose B - unlabelled output interface Type escape sequence to abort. B size 100, reply addr 3.3.3.3, return code 9 B size 100, reply addr 3.3.3.3, return code 9</pre> <pre>PE-1#traceroute mpls ipv4 1.1.1.1/32 verbose Type escape sequence to abort. 0 3.3.3.1 3.3.3 MRU 1500 [Labels: 44 Exp: 0] 1 3.3.3.3 4.4.4 MRU 1504 [No Label] 80 ms, ret code 9 2 3.3.3.3 4.4.4 MRU 1504 [No Label] 72 ms, ret code 9 ... 3 3.3.3.3 4.4.4 MRU 1504 [No Label] 88 ms, ret code 9</pre>	<p>MPLS</p> <p>Basic AToM config of two PE's</p>	<pre>PE1 mpls ldp router-id Loopback0 force mpls label protocol ldp pseudowire-class one encapsulation mpls int serial 0/0 encapsulation [hdlic, ppp] xconnect 2.2.2.2 100 pw-class one PE2 mpls ldp router-id Loopback0 force mpls label protocol ldp pseudowire-class one encapsulation mpls int serial 0/0 encapsulation [hdlic, ppp] xconnect 3.3.3.3 100 pw-class one</pre>																																																						
<p>Debugging MPLS packets</p> <p>Using ACLS:</p>	<p>debug mpls packets 2700</p> <p>access-list 2700 permit [mpls label table or any] [mpls label number] [mpls exp value] [mpls End of Stack BoS]</p> <p>access-list 2700 permit any 16 any any</p>  <p>MPLS turbo: Et3/1: rx: Len 122 Stack (16 0 253) (24 0 254) - ipv4 data MPLS turbo: Se4/0: tx: Len 108 Stack (24 0 252) - ipv4 data</p>	<p>MPLS ATOM</p> <p>Config:</p>	<ol style="list-style-type: none"> VC identifiers have to match VC Type either port mode or vlan-mode <p>port mode: int fa0/x xconnect int fa0/x.1</p> <p>Vlan mode: int fa0/x.1 encap dot1q 10 xconnect</p> <ol style="list-style-type: none"> MTU Authentication <p>- Topmost label is the transport label PE Loopback - Second label identifies the remote AC.</p> <pre>interface FastEthernet 0/1 xconnect 1.1.1.1 100 encapsulation mpls mpls ldp neighbor 1.1.1.1 password CISCO</pre>	<p>MPLS</p> <p>Troubleshooting Atom:</p>	<pre>R1#show mpls l2transport vc 100 detail Local interface: Se0/0/0 up, line protocol up, HDLC up Destination address: 10.200254.4, VC ID: 100, VC status: up Output int: E0/0, imposed label stack {19 23}</pre>  <p>show mpls l2transport hw-capability int X</p>																																																						

Help me create more flashcards:

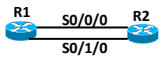
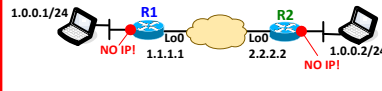
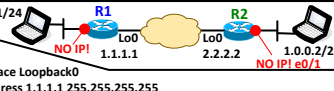
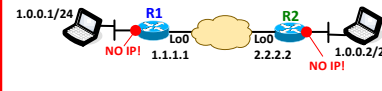
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin


<p>MPLS</p> <p>Using TDP on parralel links using the local interfaces as transport address</p>	 <pre> mpls ip mpls ldp router-id loopback 0 force interface Serial 0/1/0 mpls ip mpls ldp discovery transport-address interface mpls label protocol tdp interface Serial 0/0/0 mpls ip mpls ldp discovery transport-address interface mpls label protocol tdp </pre>				
<p>L2TPv3 configuration example:</p> 	 <pre> R1# interface Loopback0 ip address 1.1.1.1 255.255.255.255 pseudowire-class THE-WIRE encapsulation l2tpv3 ip local interface Loopback0 interface e0/1 xconnect 2.2.2.2 100 encapsulation l2tpv3 pw-class THE-WIRE R2# interface Loopback0 ip address 2.2.2.2 255.255.255.255 pseudowire-class THE-WIRE encapsulation l2tpv3 ip local interface Loopback0 interface e0/1 xconnect 1.1.1.1 100 encapsulation l2tpv3 pw-class THE-WIRE </pre>				
<p>show xconnect all</p> <p>show l2tun session [all]</p> <p>Output:</p> 	<pre> R1#show xconnect all Legend: XC ST=Xconnect State S1=Segment1 State S2=Segment2 State UP=Up DN=Down AD=Admin Down IA=Inactive SB=Standby RV=Recovering NH=No Hardware XC ST Segment 1 S1 Segment 2 S2 UP ac Et0/1(Ethernet) UP l2tp 2.2.2.2:100 UP R1#show l2tun session all Session vcid is 100 Circuit state is UP Remote tunnel name is R2 Internet address is 2.2.2.2 Local tunnel name is R1 Internet address is 1.1.1.1 IP protocol 115 </pre>				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

<p>Exec Aliases</p>	<p>alias interface si service-policy input</p> <p>alias configure iae ip access-list extended</p> <p>alias exec ri show ip route</p> <p>alias exec rb show ip bgp</p>	<p>Configuring Change Notification and Logging</p> <p>archive</p>	<p>track local config changes</p> <p>archive</p> <p>log config</p> <p>logging enable</p> <p>logging size 1000</p> <p>logging queue size 1000 entries max.</p> <p>notify syslog</p> <p>log via syslog that changes occurred</p> <p>hidekeys</p> <p>don't send passwords via syslog</p>	<p>Generating Exception Core Dumps</p>	<p>exception core-file r3-core</p> <p>Create core-dump named r3-core</p> <p>exception protocol ftp</p> <p>exception dump 155.1.146.100</p> <p>exception memory fragment 64000 reboot</p> <p>Reload case that memory fragmentation prohibits a process from allocating more than 64Kbytes of memory</p> <p>exception memory minimum 1000000 reboot</p> <p>reload as soon as free memory falls below 1Mbyte</p> <p>no ip ftp passive</p> <p>ip ftp username cisco</p> <p>ip ftp password cisco</p> <p>no exception crashinfo</p> <p>Disable local crash information collection</p> <p><i>show exception write core</i></p>																																																													
<p>System Message Logging</p>	<p>logging on</p> <p>logging buffered 8192 debugging</p> <p>save debugging up to 8192 bytes to buffers</p> <p>logging console debugging</p> <p>send debugging to console</p> <p>logging rate-limit console all 1</p> <p>limit console messages to one per second</p> <p>logging monitor informational</p> <p>users via telnet should see only informational and above messages</p> <p>line console 0</p> <p>logging synchronous</p> <p>log messages should not interrupt other command output</p> <p><i>show logging</i></p>	<p>Config / change notification And logging</p> <p>Show archive log config all</p> <p>output:</p>	<p>R4#show archive log config all</p> <table border="1"> <tr><th>idx</th><th>sess</th><th>user@line</th><th>Logged command</th></tr> <tr><td>1</td><td>1</td><td>console@console</td><td>logging enable</td></tr> <tr><td>2</td><td>1</td><td>console@console</td><td>logging size 1000</td></tr> <tr><td>3</td><td>1</td><td>console@console</td><td>notify syslog</td></tr> <tr><td>4</td><td>1</td><td>console@console</td><td>hidekeys</td></tr> <tr><td>5</td><td>2</td><td>console@console</td><td>interface Gi0/0</td></tr> <tr><td>6</td><td>2</td><td>console@console</td><td>description i am testing</td></tr> </table> <p>R4#show archive log config all provisioning</p> <p>hidekeys</p> <p>interface GigabitEthernet0/0</p> <p>description i am testing</p>	idx	sess	user@line	Logged command	1	1	console@console	logging enable	2	1	console@console	logging size 1000	3	1	console@console	notify syslog	4	1	console@console	hidekeys	5	2	console@console	interface Gi0/0	6	2	console@console	description i am testing	<p>Conditional Debugging</p>	<p>debug condition interface Gi0/0.67</p> <p>debug ip rip</p> <p>Conditions could be:</p> <ul style="list-style-type: none"> - Interfaces - usernames - calling lines <p>Undebug all does not remove the conditions!</p> <p>undebug condition interface Gi0/0.67</p> <p>Proceed with removal? [yes/no]: yes</p> <p>Condition 1 has been removed</p> <p><i>Extremely helpful on busy routers</i></p>																																	
idx	sess	user@line	Logged command																																																															
1	1	console@console	logging enable																																																															
2	1	console@console	logging size 1000																																																															
3	1	console@console	notify syslog																																																															
4	1	console@console	hidekeys																																																															
5	2	console@console	interface Gi0/0																																																															
6	2	console@console	description i am testing																																																															
<p>Syslog Logging</p>	<p>logging queue-limit trap 256</p> <p>set message queue depth to 256</p> <p>logging trap notifications</p> <p>log all messages starting at notifications to syslog</p> <p>logging origin-id string ROUTER4</p> <p>logging facility local1</p> <p>logging source-interface Loopback0</p> <p>logging host 155.1.146.100 (default, UDP 514)</p> <p>logging host 155.1.146.100 transport tcp port 5000</p>	<p>Configuration Archive & Rollback</p>	<p>- use "sw1-config" as the prefix</p> <p>- save local copy if switches saves to NVRAM (wr)</p> <p>- save a copy to IP and locally every 24 hours</p> <p>archive</p> <p>path tftp://155.1.58.100/sw1-config</p> <p>write-memory</p> <p>time-period 1440</p> <p>show archive</p> <p>show archive config differences flash:/saved-config system:/running-config</p>	<p>Telnet Service Options</p>	<p>service telnet-zeroidle</p> <p>idle outgoing telnet sessions, send remote host to pause output.</p> <p>ip telnet source-interface Loopback0</p> <p>ip telnet tos 60</p> <p>set IP precedence 3 for outgoing telnet packets</p> <p>ip telnet quiet</p> <p>Trying R4 address #1 ... Open</p> <p>ip telnet hidden addresses</p> <p>Trying R4 (155.1.146.4)... Open</p> <p>hide R4's IP when telnetting to it</p> <p>no ip domain-lookup</p> <p>ip host R4 155.1.146.4</p> <p>local host entry for R4</p> <p>busy-message R4 # Sorry, your connection failed #</p> <p>display this message if telnet to R4 fails.</p>																																																													
<p>Logging Counting and Timestamps</p>	<p>service timestamps debug uptime</p> <p>service timestamps debug datetime msec</p> <p>service timestamps log uptime</p> <p>service timestamps log datetime year</p> <p>service sequence-numbers</p> <p>revent against tampering with stored syslog information.</p> <p>logging count</p> <p>count syslog messages</p>	<p>Logging with Access-Lists</p>	<p>ip access-list extended LOGGING</p> <p>permit udp any any eq rip log-input Log L2 info MAC addr</p> <p>permit ip any any</p> <p>interface FastEthernet 0/1</p> <p>ip access-group LOGGING in</p> <p>ip access-list logging interval 10</p> <p>Send a logging message no more than once per 10 seconds</p> <p>ip access-list log-update threshold 2</p> <p>Generate a cumulative log entry for every 2 matched packets</p>	<p>Tuning Packet Buffers</p>	<p>Automatic buffer tuning:</p> <p>buffers tune automatic</p> <p>or use static assignments:</p> <p>buffers small permanent 100</p> <p>buffers middle permanent 50</p> <p>buffers big permanent 100</p> <p>buffers verybig permanent 20</p> <p>buffers large permanent 10</p> <p>buffers huge permanent 10</p> <p>Interfaces have own private buffer pool, ints have access to public buffer pools, vary in size.</p> <ul style="list-style-type: none"> - Dynamically sized buffers inefficient. - Manual sized buffers more efficient. 																																																													
<p>Syslog:</p> <p>Logging count</p> <p>Show logging count:</p>	<p>logging count</p> <p>count syslog messages</p> <p>R6#show logging count</p> <table border="1"> <thead> <tr><th>Facility</th><th>Message Name</th><th>Sev</th><th>Occur</th><th>Last Time</th></tr> </thead> <tbody> <tr><td>SYS</td><td>CONFIG_I</td><td>5</td><td>4</td><td>Jul XX 2008 23:39:00</td></tr> <tr><td colspan="5">-----</td></tr> <tr><td>SYS TOTAL</td><td></td><td></td><td>4</td><td></td></tr> <tr><td>LINEPROTO UPDOWN</td><td></td><td></td><td>5</td><td>2 Jul XX 2008 23:38:37</td></tr> <tr><td colspan="5">-----</td></tr> <tr><td>LINEPROTO TOTAL</td><td></td><td></td><td>2</td><td></td></tr> <tr><td>LINK UPDOWN</td><td></td><td></td><td>3</td><td>1 Jul XX 2008 23:38:36</td></tr> <tr><td>LINK CHANGED</td><td></td><td></td><td>5</td><td>1 Jul XX 2008 23:38:28</td></tr> <tr><td colspan="5">-----</td></tr> <tr><td>LINK TOTAL</td><td></td><td></td><td></td><td></td></tr> </tbody> </table>	Facility	Message Name	Sev	Occur	Last Time	SYS	CONFIG_I	5	4	Jul XX 2008 23:39:00	-----					SYS TOTAL			4		LINEPROTO UPDOWN			5	2 Jul XX 2008 23:38:37	-----					LINEPROTO TOTAL			2		LINK UPDOWN			3	1 Jul XX 2008 23:38:36	LINK CHANGED			5	1 Jul XX 2008 23:38:28	-----					LINK TOTAL					<p>TCP Keepalives</p>	<p>-is useful for probing idle connection to see if they are still active</p> <p>service tcp-keepalives-out</p> <p>service tcp-keepalive-in</p> <p>To test telnet from R1 to R2, perform on R2:</p> <p>show tcp brief all</p> <table border="1"> <tr><th>TCB</th><th>Local Address</th><th>Foreign Address(state)</th></tr> <tr><td>849905B4</td><td>155.1.146.1.23</td><td>155.1.146.6.17316 ESTAB</td></tr> </table> <p>R1#show tcp tcb 849905B4</p> <p>Timer Starts Wakeups Next</p> <p>KeepAlive 7 0 0x1BECF6D</p> <p>...</p> <p>Flags: passive open, active open, retransmission timeout,</p> <p>keepalive running</p>	TCB	Local Address	Foreign Address(state)	849905B4	155.1.146.1.23	155.1.146.6.17316 ESTAB	<p>Tuning Packet Buffers</p> <p>Buffer hit / misses:</p>	<p>A buffer "hit" mean a buffer was available for use when a packet arrived.</p> <p>a buffer "miss" means the IOS had to allocate a new buffer on demand for the packet.</p> <p>R4#show buffer</p> <p>Buffer elements:</p> <p>1119 in free list (1119 max allowed)</p> <p>20848 hits, 0 misses, 619 created</p> <p>Public buffer pools:</p> <p>Small buffers, 104 bytes (total 68, permanent 50, peak 68 @ 02:15:59):</p> <p>67 in free list (20 min, 150 max allowed)</p> <p>22081 hits, 6 misses, 0 trims, 18 created</p> <p>0 failures (0 no memory)</p> <p>Middle buffers, 600 bytes (total 49, permanent 25, peak 49 @ ...</p>
Facility	Message Name	Sev	Occur	Last Time																																																														
SYS	CONFIG_I	5	4	Jul XX 2008 23:39:00																																																														

SYS TOTAL			4																																																															
LINEPROTO UPDOWN			5	2 Jul XX 2008 23:38:37																																																														


LINEPROTO TOTAL			2																																																															
LINK UPDOWN			3	1 Jul XX 2008 23:38:36																																																														
LINK CHANGED			5	1 Jul XX 2008 23:38:28																																																														

LINK TOTAL																																																																		
TCB	Local Address	Foreign Address(state)																																																																
849905B4	155.1.146.1.23	155.1.146.6.17316 ESTAB																																																																
<p>Logging to Flash Memory</p>	<p>mkdir flash:/var</p> <p>mkdir flash:/var/log</p> <p>logging file flash:/var/log/syslog 32768 notifications</p> <p>logging on</p> <p>show logging:</p> <p>... File logging: file flash:/var/log/syslog, max size 32768, min size 0, level notifications, 2 messages logged</p> <p>....</p>	<p>Debugging</p> <p>TCP keepalives:</p>	<p>Established TCP session</p> <p>Interface was shutdown</p> <p>debug ip tcp transactions</p> <p>TCB: keepalive timeout (0/4)</p> <p>RackR1#</p> <p>TCB: keepalive timeout (1/4)</p> <p>RackR1#</p> <p>TCB: keepalive timeout (2/4)</p> <p>RackR1#</p> <p>TCB: keepalive timeout (3/4)</p> <p>RackR1#</p> <p>TCB: keepalive timeout (4/4)</p> <p>RackR1#</p> <p>TCB: state was ESTAB -> CLOSED [23 -> 155.1.146.6(17316)]</p> <p>TCB 0x849905B4 destroyed</p> <p>service tcp-keepalives-out</p> <p>service tcp-keepalive-in</p> <p><i>After 4 timeouts, TCP connection is forcefully closed.</i></p>	<p>Terminal Line Settings</p>	<p>vacant-message #</p> <p>refuse-message # Sorry, the line is already in use#</p> <p>line vty 0 4</p> <p>session-timeout 1</p> <p>exec-timeout 2 0</p> <p>Timeout after 2 mins of inactivity</p> <p>lockable</p> <p>absolute-timeout 5</p> <p>Will disconnect in any case in 5 min!</p> <p>ip netmask-format hexadecimal</p> <p>Display netmasks in HEX</p> <p>length 20</p> <p>Terminal length no more than 20 lines</p> <p>transport input telnet</p> <p>line vty 0</p> <p>rotary 1</p> <p>Makes vty0 listen to port 3001</p> <p>line console 0</p> <p>session-limit 1</p> <p>Allows only one user</p> <p><i>show line vty 0</i></p>																																																													

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts


Colin


<h3>SNMPv2 Server</h3>	<pre>snmp-server community CISCO RW snmp-server location Default Location snmp-server contact Default Contact snmp-server ifindex persist Ensure that interface index numbers persist between reloads. snmp-server system-shutdown access-list 98 permit 155.1.146.100 snmp-server tftp-server-list 98 Only allow configuration transfers via TFTP to/from the host 155.X.146.100</pre>	<h3>CPU and Memory Thresholds</h3>	<pre>memory free low-watermark processor 1000 set up the free memory low threshold to 1000Kbytes memory reserve critical 512 Reserve 512Kbytes of memory for the notification process using the memory command. process cpu threshold type total rising 50 interval 5 monitor CPU usage every 5 seconds using the process cpu command, and to generate a rising threshold event every time the CPU usage hits 50% snmp-server enable traps cpu threshold send snmp CPU traps</pre>	<h3>SNMPv3</h3> <p>View 2:</p> <h3>traps</h3>	<pre>Enable SNMP traps for LinkUp and LinkDown events only, and send them to the destination host 155.X.146.100 using the security model "priv" and the username TRAP. snmp-server group TRAP v3 priv snmp-server user TRAP TRAP v3 auth sha CISCO priv des56 CISCO snmp-server host 155.1.146.100 traps version 3 priv TRAP snmp-server enable traps snmp linkup linkdown</pre>
<h3>show snmp commands:</h3>	<pre>show snmp Chassis: 26388555 Contact: Default Contact Location: Default Location show snmp community Community name: ILM1 Community Index: cisco0 Community SecurityName: ILM1 storage-type: read-only active show snmp mib ifmib ifindex Ethernet0/0: ifindex = 1 Loopback0: ifindex = 7 Null0: ifindex = 6 Serial0/0: ifindex = 3 Ethernet0/1: ifindex = 2</pre>	<h3>Testing Memory thresholds:</h3>	<pre>show memory summary Head Total (b) Used (b) Free (b) Lowest (b) Largest (b) Processor 6420A860 59725728 14638656 45087072 43403156 43632028 memory free low-watermark processor 50000 %SYS-4-FREEMEMLOW: Free Memory has dropped below low watermark Pool: Processor Free: 45087396 Threshold: 51200000 memory free low-watermark processor 5000 %SYS-5-FREEMEMRECOVER: Free Memory has recovered above low watermark Pool: Processor Free: 45087660 Threshold: 5120000</pre>	<h3>SNMPv3</h3> <p>show snmp user:</p>	<pre>show snmp user User name: TRAP Engine ID: 80000009030000119221DA80 storage-type: nonvolatile active Authentication Protocol: SHA Privacy Protocol: DES Group-name: TRAP User name: NORMAL Engine ID: 80000009030000119221DA80 storage-type: nonvolatile active Authentication Protocol: SHA Privacy Protocol: DES Group-name: NORMAL</pre>
<h3>SNMPv2c Access Control</h3>	<pre>Restrict RW access from one subnet, log all other attempts for community cisco: Expose other hosts that attempt to access the router via SNMP. access-list 99 permit 155.1.146.0 0.0.0.255 access-list 99 deny any log snmp-server community CISCO RW 99 limit community PUBLIC to read-only mode, Restrict access MIB access only to the "cisco" subtree: snmp-server community PUBLIC view ROVIEW ro snmp-server view ROVIEW cisco included</pre>	<h3>Testing CPU thresholds:</h3>	<pre>process cpu threshold type total rising 5 interval 5 5 sec 5% %SYS-1-CPURISINGTHRESHOLD: Threshold: Total CPU Utilization(Total/Intr): 32%/0%, Top 3 processes(Pid/Util): 3/31%, 91/0%, 2/0%</pre>	<h3>SNMPv3</h3> <p>show snmp group</p>	<pre>show snmp group groupname: NORMAL security model:v3 priv readview : NORMAL writeview: NORMAL notifyview: <no notifyview specified> row status: active groupname: RESTRICTED security model:v3 auth readview : RESTRICTED writeview: <no writeview specified> notifyview: <no notifyview specified> row status: active access-list: 99</pre>
<h3>show snmp community</h3>	<pre>R4#show snmp community Community name: ILM1 Community Index: cisco0 Community SecurityName: ILM1 storage-type: read-only active Community name: CISCO Community Index: cisco4 Community SecurityName: CISCO storage-type: nonvolatile active access-list: 99 Community name: PUBLIC Community Index: cisco5 Community SecurityName: PUBLIC storage-type: nonvolatile active</pre>	<h3>SNMPv3</h3> <p>General info:</p>	<ul style="list-style-type: none"> - group defines what access rights a set of users have and controls which SNMP objects (MIBs) can be accessed for reading and writing - group defines which SNMP objects can generate notifications to the members of a group - security model (SNMP version) - security level (authentication and/or encryption) - read view has implicit permit, if no write or notify is defined. <pre>security levels are defined as noauth noAuthNoPriv, (no auth, no encrypt) auth AuthNoPriv (auth, no encrypt) priv AuthPriv (auth and encrypt) Group security model, but password and encryption key) are set per-user</pre>	<h3>SNMPv3</h3> <p>show snmp view</p>	<pre>show snmp view *ilmi system - included permanent active *ilmi atmForumUni - included permanent active NORMAL iso - included nonvolatile active v1default iso - included permanent active v1default internet.6.3.15 - excluded permanent active v1default internet.6.3.16 - excluded permanent active v1default internet.6.3.18 - excluded permanent active v1default ciscoMgmt.394 - excluded permanent active v1default ciscoMgmt.395 - excluded permanent active v1default ciscoMgmt.399 - excluded permanent active v1default ciscoMgmt.400 - excluded permanent active RESTRICTED ifEntry.0.3 FF:FF included nonvolatile active *tv.FFFFFFFFF.FFFFFFFFF.FFFFFFFFF0F iso.2.840.10036 - included volatile active *%FFFFFFFF.FFFFFFFFF.FFFFFFFFF0F internet - included volatile active</pre>
<h3>show snmp view</h3>	<pre>R4#show snmp view *ilmi system - included permanent active *ilmi atmForumUni - included permanent active ROVIEW cisco - included nonvolatile active v1default iso - included permanent active v1default internet.6.3.15 - excluded permanent active v1default internet.6.3.16 - excluded permanent active v1default internet.6.3.18 - excluded permanent active v1default ciscoMgmt.394 - excluded permanent active v1default ciscoMgmt.395 - excluded permanent active v1default ciscoMgmt.399 - excluded permanent active v1default ciscoMgmt.400 - excluded permanent active</pre>	<h3>SNMPv3</h3> <p>Views 1:</p> <p>Normal</p>	<pre>snmp-server ifindex persist I create view NORMAL to include iso branch. snmp-server view NORMAL iso included I create group with read, write view NORMAL snmp-server group NORMAL v3 priv read NORMAL write NORMAL I assign user NORMAL to group, set security model to priv I set auth password CISCO and encryption key to CISCO snmp-server user NORMAL NORMAL v3 auth sha CISCO priv des56 CISCO</pre>	<h3>SNMPv1</h3> <h3>SNMPv2</h3> <h3>SNMPv3</h3> <p>Ports:</p>	<pre>UDP 161 for polling, UDP 162 for notifications debug snmp packet</pre>
<h3>SNMP Traps and Informs</h3>	<pre>snmp-server enable traps snmp linkdown linkup snmp-server host 155.1.146.101 inform version 2c CISCO snmp-server host 155.1.146.100 CISCO interface Serial0/0 no snmp trap link-status show snmp host Notification host: 155.1.146.101 udp-port: 162 type: inform user: CISCO security model: v2c Notification host: 155.1.146.100 udp-port: 162 type: trap user: CISCO security model: v1</pre>	<h3>SNMPv3</h3> <p>View 2:</p> <p>restricted</p>	<pre>snmp-server ifindex persist I create view RESTRICTED to include ifEntry 3 branch. snmp-server view RESTRICTED ifEntry.*.3 included I create group with read view restricted, use security model auth. snmp-server group RESTRICTED v3 auth read RESTRICTED access 99 I Assign user RESTRICTED to group Restricted, only use auth with a key of CISCO snmp-server user RESTRICTED RESTRICTED v3 auth sha CISCO</pre>	<h3>SNMP MAC Address Notifications</h3>	<pre>interface FastEthernet0/5 snmp trap mac-notification added snmp trap mac-notification removed snmp-server enable traps mac-notification I rate-limit notifications 1 per second mac-address-table notification interval 1 I 100 notification events in history buffer mac-address-table notification history-size 100 mac-address-table notification</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!





Thanks for appreciating my efforts

Colin

<p>SNMP MAC Address Notifications:</p> <p>show mac-address-table notification</p>	<pre>clear mac-address-table dynamic show mac-address-table notification MAC Notification Feature is Enabled on the switch Interval between Notification Traps : 2 secs Number of MAC Addresses Added : 7 Number of MAC Addresses Removed : 7 Number of Notifications sent to NMS : 9 Maximum Number of entries configured in History MAC Notification Traps are Enabled History Index 0, Entry Timestamp 3441456, Despatch Timestamp 3441456 MAC Changed Message : Operation: Deleted Vlan: 58 MAC Addr: 0004.9a0b.62c1 Dot1dBasePort: 7</pre>	<p>show ip http server status</p>	<pre>show ip http server status HTTP server status: Enabled HTTP server port: 8080 HTTP server authentication method: local HTTP server access class: 0 HTTP server base path: flash: Maximum number of concurrent server connections allowed: 2 Server idle time-out: 180 seconds Server life time-out: 180 seconds Maximum number of requests allowed on a connection: 1 HTTP server active session modules: ALL HTTP secure server capability: Present HTTP secure server status: Enabled HTTP secure server port: 4043 HTTP secure server ciphersuite: des-cbc-sha HTTP secure server client authentication: Disabled HTTP secure server trustpoint: HTTP secure server active session modules: ALL</pre>	<p>NTP</p> <p>ntp peer ntp broadcast ntp multicast</p>	
<p>SNMP Notifications of Syslog Messages</p>	<pre>send debugging and higher prio messages via SNMP to x.x.x.x snmp-server enable traps syslog snmp-server host 155.1.146.100 CISCO Set syslog SNMP buffer size to 100 msgs logging history debugging logging history size 100</pre>	<p>show ip http client all</p>	<pre>show ip http client all HTTP client status: Enabled HTTP client application session modules: Id : 1 Application Name : HTTP CFS Version : HTTP/1.0 Persistent : persistent Response-timeout : 0 Retries : 0 Proxy : HTTP client current connections: Persistent connection = enabled (default) Connection establishment timeout = 10s (default) Connection idle timeout = 30s (default) Maximum number of connection establishment retries = 1 (default) Maximum http client connections per host : 2 HTTP secure client capability: Present HTTP secure client ciphersuite: des-cbc-sha HTTP secure client trustpoint:</pre>	<p>Show ntp associations details</p> <p>Output:</p>	<pre>SW2#show ntp associations detail 155.1.58.5 dynamic, our_master, sane, valid, stratum 6 ref ID 150.1.4.4, time D6DFC145.83F9908B (09:37:09.515 UTC Fri Mar 28 2014) our mode bdcast client, peer mode bdcast, our poll intvl 64, peer poll intvl 64 root delay 45.29 msec, root disp 17.17, reach 177, sync dist 420.959 delay 2.18 msec, offset -2.8375 msec, dispersion 380.07 precision 2**18, version 3</pre>
<p>CDP</p>	<pre>no cdp log mismatch duplex cdp source-interface Loopback0 ! send CDP announcement every 10 seconds cdp timer 10 !instruct other devices to hold the updates for 40 seconds cdp holdtime 40 interface FastEthernet 0/0 no cdp enable</pre>	<p>FTP Server and Client</p>	<pre>R4# no ip ftp passive ip ftp source-interface Loopback0 ip ftp username CISCO ip ftp password CISCO</pre>	<p>NTP in Lab</p>	<p>Adjust NTP clocks in order to synchronize quicker. If clocks are to far apart from each other, this could take ages to converge.</p> <pre>R4#clock set 00:00:01 Jan 1 2012 R5#clock set 00:00:01 Jan 1 2012 R6#clock set 00:00:01 Jan 1 2012 SW1#clock set 00:00:01 Jan 1 2012 SW2#clock set 00:00:01 Jan 1 2012 SW3#clock set 00:00:01 Jan 1 2012</pre> <p>NTP Key ID's have to match on both ends!</p>
<p>CDP</p> <p>show commands:</p>	<pre>R4#show cdp Global CDP information: Sending CDP packets every 10 seconds Sending a holdtime value of 40 seconds Sending CDPv2 advertisements is enabled Source interface is Loopback0 R4#show cdp interface FastEthernet0/1 is up, line protocol is up Encapsulation ARPA Sending CDP packets every 10 seconds Holdtime is 40 seconds R4#show cdp traffic CDP counters : Total packets output: 160, Input: 148 Hdr syntax: 0, Chksum error: 0, Encaps failed: 0 No memory: 0, Invalid packet: 0, Fragmented: 0 CDP version 1 advertisements output: 0, Input: 0 CDP version 2 advertisements output: 160, Input: 148</pre>	<p>TFTP Server and Client</p>	<pre>R6: tftp-server flash:XXX.bin alias R6-IOS 10 access-list 10 permit 150.1.1.1 R1: ip tftp source-interface Loopback0 debug tftp events</pre>	<p>NTP Authentication</p> <p>ntp broadcast:</p>	<pre>Router# ntp server 1.1.1.1 prefer ntp authenticate ntp authentication-key 58 md5 CISCO58 interface Gi 0/0 ntp broadcast ntp broadcast key 58 Switch# ntp authenticate ntp authentication-key 58 md5 CISCO58 ntp trusted-key 58 int vlan 58 ntp broadcast client</pre>
<p>IOS HTTP Server and Client</p>	<pre>ip http client source-interface Loopback0 ip http client username CISCO ip http client password CISCO ip http client secure-ciphersuite des-cbc-sha username CISCO password 0 CISCO ip http server ip http max-connections 2 ip http path flash: ip http port 8080 access-list 80 permit 150.1.0.0 0.0.255.255 ip http access-class 80 ip http authentication local ip http secure-server ip http secure-port 4043 ip http secure-ciphersuite des-cbc-sha</pre>	<p>Remote Shell</p>	<pre>Config: R1: ip rcmd remote-username RCP ip rcmd source-interface Loopback0 R6: ip rcmd rcp-enable ip rcmd rsh-enable ip rcmd remote-host R6 150.1.1.1 Rack1R1 enable ip rcmd remote-host RCP 150.1.1.1 Rack1R1 enable Run commands on from R1 on R6: rsh 150.1.6.6 /user R6 show run int gi0/0</pre>	<p>NTP authentication</p> <p>ntp peer x.x.x.x</p>	<pre>Loopback 0 (150.1.4.4) ntp master 5 ntp peer 150.1.6.6 ntp source Loopback0 Loopback0 (150.1.6.6) ntp master 5 ntp peer 150.1.4.4 ntp source Loopback0 ntp authenticate ntp authentication-key 46 md5 CISCO46 ntp trusted-key 46 ntp peer 150.1.6.6 key 46 ntp authenticate ntp authentication-key 46 md5 CISCO46 ntp trusted-key 46 ntp peer 150.1.4.4 key 46</pre>
<p>A good way to test http / tcp oriented sessions:</p>	<p>copy http://1.2.3.4:80/IOS-X.bin null:</p>	<p>Remote Shell</p> <p>Performing remote commands:</p>	<pre>#R4 rsh 150.1.6.6 /user R6 show run int Ser0/0 R4# copy rcp://150.1.6.6/saved-config null: Source username [RCP]? Accessing rcp://RCP@150.1.6.6/saved-config...! 1941 bytes copied in 0.056 secs (34661 bytes/sec)</pre> <p><i>debug ip tcp rcmd</i></p>	<p>The difference between</p> <p>NTP peer NTP server</p>	<p>NTP Peer:</p> <p>Both routers can update their clocks vice-versa, Like a cluster of NTPs</p> <p>NTP Server:</p> <p>The local "ntp client" can only get the time, but will not update his local time to the other NTP server.</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

<p>NTP Authentication</p> <p>NTP Master ↔ NTP Server</p>	<pre> Loopback 0 (150.1.4.4) ntp master 5 ntp source Loopback0 ntp authenticate ntp authentication-key 4 md5 CISCO4 ntp server 150.1.6.6 ntp server 150.1.4.4 prefer ntp source Loopback0 ntp authenticate ntp trusted-key 4 ntp server 150.1.4.4 key 4 ntp authentication-key 4 md5 CISCO4 ntp trusted-key 6 ntp server 150.1.6.6 key 6 ntp authentication-key 6 md5 CISCO6 Loopback0 (150.1.6.6) ntp master 5 ntp source Loopback0 ntp authenticate ntp authentication-key 6 md5 CISCO6 </pre>	<p>How to find the client identifier on a Cisco Router / Switch:</p>	<pre> R4 R4(config)#interface FastEthernet0/1 DHCP SRV R4(config-if)#ip address dhcp R4#show dhcp lease Temp IP addr: 0.0.0.0 for peer on Interface: FastEthernet0/1 Temp sub net mask: 0.0.0.0 DHCP Lease server: 0.0.0.0, state: 1 Selecting DHCP transaction id: 14F4 Lease: 0 secs, Renewal: 0 secs, Rebind: 0 secs Next timer fires after: 00:00:01 Retry count: 0 Client-ID: cisco-0007.ebde.5622-E10/1 Client-ID hex dump: 636973636F2D303030372E656264652E 353632322D4574302F31 R5# ip dhcp pool HOST_R4 client-identifier 00636973636F2D303030372E656264652E 353632322D4574302F31 </pre>	<p>Reverse Telnet On AUX lines:</p>	<p>line aux 0</p> <p>no modem inOut</p> <p>transport input telnet</p> <p>Clear line <nr-of-aux-line></p> <p>Telnet to the AUX-line via:</p> <p>port number by adding 2000 + TTY = 20xx</p> <p>Telnet <IP> 20xx</p> <p>alias exec auxup telnet 127.0.0.11 2097</p> <p>alias exec auxdown clear line 97</p>
<p>NTP</p> <p>Check if ntp is authenticated:</p>	<pre> R4#show ntp associations detail inc auth 150.1.6.6 configured, authenticated, selected, sane, valid, stratum 5 </pre>	<p>Auto-Install over LAN Interfaces using DHCP</p> <p>Boot files / sequence</p>	<p>If the DHCP/BOOTP server did not provide a boot-file name, the router looks for:</p> <p>network-confg or cisco.net.cfg</p> <p>This to allow the router to determine its hostname. If the host name is found, the router then checks for</p> <p><hostname>-confg</p> <p>If that lookup fails, the router uses router-confg or outer.cfg</p>	<p>KRON Command Schedule</p>	<p>Configure policy first:</p> <pre> kron policy-list SAVE_CONFIG cli show int X redirect tftp://1.2.3.4/int-x.txt </pre> <p>Apply policy to occurrence:</p> <pre> kron occurrence SAVE_DAILY at 8:00 recurring policy-list SAVE_CONFIG </pre> <p>occurrence either one-shot or recurring</p> <p>show kron schedule</p>
<p>NTP</p> <p>Access Control</p>	<p>Server only hosts listed in ACL5</p> <pre> ntp access-group serve-only 5 access-list 5 permit 5.5.5.5 Only allow IPs out of ACL 6 to update the local clock: ntp access-group peer 6 access-list 6 permit 150.1.6.6 access-list 6 permit 127.127.7.1 </pre> <p><i>NTP locally uses 127.127.7.1 to update its clock</i></p>	<p>Auto-Install over LAN Interfaces using RARP</p>	<p>router sends out RARP requests for an IP address after it fails to obtain an address via DHCP/BOOTP</p> <pre> ip dns server ip host R4 155.1.146.4 arp 155.1.146.4 0007.ebde.5622 arpa interface FastEthernet 0/0 no shutdown ip rarp-server 155.1.146.1 tftp-server flash:r4-confg tftp-server flash:network-confg </pre>	<p>EEM Scripting: Interface Events</p> <p>Background info:</p>	<p>Event detectors:</p> <ul style="list-style-type: none"> - CLI event detector - Syslog event detector - Interface Counter thresholds - Counter (generic) - SNMP - None (event manager run) - Watchdog periodic timer events <p>Event subscribers</p> <ul style="list-style-type: none"> - start with an action keyword -TCL scripts <p>access global variables via event manager environment</p>
<p>NTP</p> <p>message types:</p>	<ul style="list-style-type: none"> - control messages <ul style="list-style-type: none"> peer status set a management parameter - update/request messages <ul style="list-style-type: none"> time synchronization 	<p>Troubleshooting</p> <p>Auto-Install over LAN Interfaces using RARP</p>	<pre> debug tftp events debug tftp packets debug arp RARP: Rcvd RARP req for 0007.ebde.5622 TFTP: Opened flash:network-confg, fd 0, size 1440 TFTP: Opened flash:R4-confg, fd 0, size 1494 </pre>	<p>EEM Scripting: Interface Events</p>	<p>Run script every 30 seconds, write to flash:</p> <pre> event manager applet TEST event timer watchdog name timer time 30 action 0.5 cli command "enable" action 1.0 cli command "show clock append flash:file.txt" action 2.0 cli command "show ip cache flow append flash:file.txt" action 3.0 cli command "show voip rtp conn" append flash:file.txt </pre> <p>If you type IDI on the cli, file.bla will be deleted:</p> <pre> event manager applet TEST event cli pattern "IDI" sync no skip yes action 1.0 cli command "enable" action 2.0 cli command "delete /force flash:file.bla" show event manager policy registered </pre>
<p>NTP</p> <p>access control levels:</p>	<p>NTP access control defines four levels:</p> <ul style="list-style-type: none"> - Peer (permits NTP updates/requests to the host as well as control queries) - Serve (permits NTP requests, but rejects NTP updates) - serve-only (permits NTP requests only, does not accept control Queries) - query-only (accepts NTP control queries, no local system time synchronization with a remote system is allowed.) 	<p>IOS Menus</p> <p>Once you login to a router:</p> <pre> Operator Menu 1 Display IP Routing Table 2 Display Running Config 3 Escape to Shell 4 Disconnect </pre>	<pre> menu OPERATOR title # Operator Menu # menu OPERATOR text 1 Display IP Routing Table menu OPERATOR command 1 show ip route menu OPERATOR text 2 Display Running Config menu OPERATOR command 2 show run menu OPERATOR text 3 Escape to Shell menu OPERATOR command 3 menu-exit menu OPERATOR text 4 Disconnect menu OPERATOR command 4 exit menu OPERATOR clear-screen username OPERATOR autocommand menu OPERATOR username OPERATOR password CISCO username OPERATOR privilege 15 line vty 0 4 login local </pre>	<p>EEM Scripting: Syslog Events</p> <p>No shutting serial interface if shut, Sending email of connected users:</p>	<pre> event manager applet INTERFACE_SHUTDOWN event tag 1.0 syslog pattern "Interface Serial0/0/0."changed."down" action 1.0 cli command "enable" action 2.0 cli command "conf t" action 3.0 cli command "interface Serial 0/0/0" action 4.0 cli command "no shutdown" action 5.0 cli command "end" action 6.0 cli command "show users" action 7.0 mail server "155.1.146.100" to "admin@INE.com" from "r5@INE.com" subject "Interface Shutdown Alert" body "Interface Serial 0/0/0 unshut, current users \$ _cli_result" show event manager policy registered </pre>
<p>Auto-Install over LAN Interfaces using DHCP</p>	<pre> R5# ip dns server ip host R4 155.1.146.4 ip host R5 150.1.5.5 ip dhcp pool HOST_R4 host 155.1.146.4 255.255.255.0 client-identifier 0063.6973.636f.2d30.30.30..... default-router 155.1.146.1 dns-server 155.1.146.6 option 66 ascii "R5" tftp-server flash:r4-confg flash:/r4-confg order of preference is: sname option 66 option 150 siaddr </pre>	<p>IOS Banners</p>	<pre> banner motd # Welcome to IOS Router # banner login # Please authenticate yourself # banner exec # Hi, you are on the line \$(line), have a nice time at \$(hostname) # banner incoming # This is a reverse telnet connection # line console 0 no motd-banner no exec-banner </pre>	<p>Troubleshooting</p> <p>EEM</p> <p>Event Manager:</p>	<p>show event manager policy registered</p> <pre> debug event manager detector syslog debug event manager action cli Run Event applet manually: event manager run TEST-THIS </pre>


no event manager applet LOOPBACK_SHUTDOWN
event manager applet LOOPBACK_SHUTDOWN
event cli pattern ".*interface Loopback.*" sync yes
action 1.0 cli command "enable"
action 1.1 cli command "configure terminal"
action 1.2 cli command "\$ _cli_msg"
action 1.3 cli command "shutdown"


```

ip sla 1
icmp-echo 115.0.0.1 Account for timed out
request-data-size 1250 packets after 25 msec
timeout 25
threshold 20 Max latency 20 msec
frequency 30 Ping every 30 sec
!
ip sla schedule 1 start-time now life forever
                    
```

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

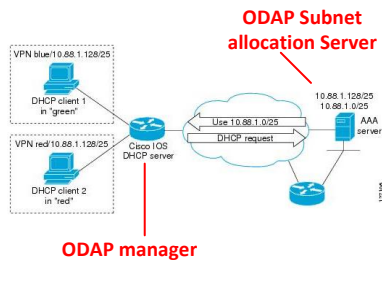




Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts


Colin

<p>DHCP address pool with secondary subnets:</p>	<p>ip dhcp pool POOL network 172.16.0.0 /16 default-router 172.16.0.1 domain-name cisco.com dns-server 172.16.1.102 172.16.2.102 lease 30</p> <p>network 172.16.1.0 /24 secondary override default-router 172.16.1.1 end</p> <p>network 172.16.2.0 /24 secondary override default-router 172.16.2.1</p>	<p>Disabling BFD Echo Mode Without Asymmetry</p>	<p>BFD Version 0 and therefore does not use the echo mode</p> <p>disable BFD echo mode without asymmetry—no echo packets will be sent by the router, and the router will not forward BFD echo packets that are received from any neighbor routers.</p> <pre>conf t no bfd echo</pre>	<p>When ever you see this while doing a write memory Or copy running-config startup-config:</p> <p>SW4#write memory startup-config file open failed (Not enough space)</p>	<p>Don't waste time!</p> <ol style="list-style-type: none"> 1. Copy the entire config into Notepad and reload the box. 2. once reloaded, paste the config back on. <p>----- Your proctor has been doing this to you:</p> <pre>write memory conf t boot config-file flash:/ end</pre> <p><i>first to saved config Then pointed the config file into nowhere</i></p>
<p>ODAP Subnet allocation Server And ODAP manager Config</p>	<p>Subnet allocation Server:</p> <pre>ip dhcp pool VRF-POOL vrf RED network 172.16.0.0 /16 subnet prefix-length 26</pre> <p>ODAP manager:</p> <pre>ip dhcp pool usergroup1 origin dhcp subnet size initial /26 autogrow /26 lease 0 1</pre> <p><i>show ip dhcp pool</i></p>	<p>Configuring BFD Templates</p>	<pre>conf t bfd-template single-hop <template-name> interval min-tx 50 min-rx 50 multiplier 5</pre>	<p>The task says something like, configure "clock timezone GMT +8"</p> <p>And this happens:</p> <p>-----</p> <pre>R2(config)#clock ? <cr> R2(config)#exit</pre> <p>R2(config)#exit R2#</p>	<p>The task says, configure clock timezone GMT +8</p> <p>The correct command is:</p> <pre>conf t clock timezone GMT +8</pre> <p><i>Because something is strange you type clock ? Then out of nothing.. the "exit" on the next line appears.</i></p> <p>R2(config)#clock ? <cr> R2(config)#exit</p> <p>The proctor or Narbik is picking on you!!</p> <pre>R2#sh run i alias alias configure clock exit</pre>
<p>ODAP address pool management Terminology:</p>		<p>Troubleshooting BFD</p>	<pre>show bfd neighbors [details] debug bfd packet debug bfd event</pre>	<p>How does HSRP elect the active Node?</p>	<ol style="list-style-type: none"> 1. Higher priority: <pre>int fa0/x standby <nr> priority <1-255></pre> <ol style="list-style-type: none"> 2. If Priorities are even, highest IP address. <p>Default priority 100</p>
<p>Bidirectional forwarding detection</p> <p>BFD per interface:</p> <p>BFD for static routes:</p>	<p>BFD per interface:</p> <pre>bfd interval <msec> min_rx <msec> multiplier <interval-multiplier></pre> <pre>interface fa0/x bfd interval 50 min_rx 50 multiplier 5</pre> <p>BFD for static routes:</p> <pre>int fa0/x ip address 10.0.0.2 255.255.255.0 bfd interval 50 min_rx 50 multiplier 5</pre> <pre>ip route static bfd fa0/0 10.0.0.1 group GRP-1 {passive} ip route 0.0.0.0 0.0.0.0 10.0.0.1</pre> <pre>show ip static route show ip static route bfd</pre>	<p>BFD in an HSRP Network</p>	<pre>standby bfd / standby bfd all-interfaces needed only if BFD has been manually disabled on a router or interface</pre> <pre>ip cef interface Fa0/x bfd interval 200 min_rx 200 multiplier 3 standby 1 ip 10.0.0.11 standby 1 preempt standby 1 priority 110</pre>	<p>What options are there in regards to HSRP authentication?</p>	<pre>standby 1 authentication cisco (default password, not visible) standby 1 authentication text Cisco standby 2 authentication md5 key-string HSRP standby 3 authentication md5 key-chain CHAIN</pre>
<p>BFD limitations with IP redirects</p>	<p>When using BFD echo mode BFD version 0, you must disable the sending of Internet Control Message Protocol (ICMP) redirect messages</p> <p>no ip redirects</p> <pre>bfd slow-timer <msec></pre>	<p>Write a EEM script which creates Interface Loopback 0 with ip 1.1.1.1/32</p>	<pre>event manager applet BLA event none sync yes action 1.0 cli command "enable" action 1.1 cli command "conf t" action 1.2 cli command "int lo0" action 1.3 cli command "ip address 1.1.1.1 255.255.255.255" action 1.4 cli command "end"</pre> <p>Manually run it: event manager run BLA</p> <pre>%SYS-5-CONFIG_I: Configured from console by vty0 (EEM:BLA) %LINEPROTO-5-UPDOWN: Line protocol on Interface Loopback0, changed state to up</pre> <p><i>none = manual trigger Sync yes = run commands one by one, prevents command over-run situations.</i></p>	<p>How can you put a "question mark" into a Password using Cisco IOS ?</p> <pre>username test password 0 I.have.a.lot.of.?.for.you</pre>	<p>As you write ...lot.of. Hit control-V on the key board, then write the ? And continue</p> <pre>R1#conf t username test password 0 I.have.a.lot.of.?.for.you</pre> <p><i>Copy and pasting the config, will end in chaos and pain due to the ? Wrongly interpreted due to the missing Control-V interaction!!!</i></p> <pre>key chain EVIL key 1 key-string ???is.this.how.you.looked.like</pre>
<p>BFD and uRPF</p> <p>Does not work!</p>	<p>BFD echo mode does NOT work in conjunction with Unicast Reverse Path Forwarding (uRPF)</p>	<p>Write an EEM script with prints "THIS WILL NOT WORK DUDE"</p> <p>If one enters this on the Cli:</p> <pre>ping 1.2.3.4</pre>	<p>Skip yes = will skip the entered ping command!</p> <pre>event manager applet PING event cli pattern "ping 1.2.3.4" skip yes sync yes action 1.0 syslog msg "THIS WILL NOT WORK DUDE"</pre> <pre>R1#ping 1.2.3.4 R1# %HA_EM-6-LOG: PING: THIS WILL NOT WORK DUDE</pre>		

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin

<p>Proxy ARP</p>		<p>DHCP Client</p> <p>Setting client-identifier</p>	<p>interface FastEthernet 0/0 ip address dhcp client-id FastEthernet 0/0</p>	<p>DHCP Proxy</p>	<pre> R2# interface Serial 0/1 encapsulation ppp ip address negotiated no peer neighbor-route R3# interface Serial 1/3 encapsulation ppp ip address 155.1.23.3 255.255.255.0 peer default ip address dhcp no peer neighbor-route ip address-pool dhcp-proxy-client ip dhcp-server 155.1.146.6 </pre>																																												
<p>Bootp Packet header:</p>	<table border="1"> <tr><th>Operation code</th><th>HW type</th><th>HW addr length</th><th>Hops</th></tr> <tr><td colspan="4">Transaction Identifier</td></tr> <tr><td colspan="2">Seconds</td><td colspan="2">Flags</td></tr> <tr><td colspan="4">Client IP address (CIAddr)</td></tr> <tr><td colspan="4">Your IP address (YIAddr)</td></tr> <tr><td colspan="4">Server IP address (SIAddr)</td></tr> <tr><td colspan="4">Gateway IP address (GIAddr)</td></tr> <tr><td colspan="4">Client Hardware Address (CHAddr)</td></tr> <tr><td colspan="4">Server Name (Sname)</td></tr> <tr><td colspan="4">Boot File Name (128 bytes)</td></tr> <tr><td colspan="4">Vendor-Specific Area (64 bytes) fixed option length!</td></tr> </table>	Operation code	HW type	HW addr length	Hops	Transaction Identifier				Seconds		Flags		Client IP address (CIAddr)				Your IP address (YIAddr)				Server IP address (SIAddr)				Gateway IP address (GIAddr)				Client Hardware Address (CHAddr)				Server Name (Sname)				Boot File Name (128 bytes)				Vendor-Specific Area (64 bytes) fixed option length!				<p>Troubleshooting DHCP</p>	<p>Sequence: D O R A debug dhcp detail debug dhcp detail Int x no ip address dhcp ip address dhcp</p> <p>debug ip dhcp server event debug ip dhcp server packets show ip dhcp pool show ip dhcp database show ip dhcp binding</p> <p>release dhcp</p>	<p>DHCP Information Option</p>	<pre> Relay: ip dhcp relay information option interface FastEthernet0/0 ip dhcp relay information option subscriber-id VLAN58 Server: ip dhcp class TEST relay agent information relay-information hex 020c020a00009b013a05000000000606564c414e3538 ip dhcp pool VLAN58 class TEST address range 155.1.58.8 155.1.58.8 </pre>
Operation code	HW type	HW addr length	Hops																																														
Transaction Identifier																																																	
Seconds		Flags																																															
Client IP address (CIAddr)																																																	
Your IP address (YIAddr)																																																	
Server IP address (SIAddr)																																																	
Gateway IP address (GIAddr)																																																	
Client Hardware Address (CHAddr)																																																	
Server Name (Sname)																																																	
Boot File Name (128 bytes)																																																	
Vendor-Specific Area (64 bytes) fixed option length!																																																	
<p>DHCP Packet header:</p>	<table border="1"> <tr><th>Operation code</th><th>HW type</th><th>HW addr length</th><th>Hops</th></tr> <tr><td colspan="4">Transaction Identifier</td></tr> <tr><td colspan="2">Seconds</td><td colspan="2">Flags</td></tr> <tr><td colspan="4">Client IP address (CIAddr)</td></tr> <tr><td colspan="4">Your IP address (YIAddr)</td></tr> <tr><td colspan="4">Server IP address (SIAddr)</td></tr> <tr><td colspan="4">Gateway IP address (GIAddr)</td></tr> <tr><td colspan="4">Client Hardware Address (CHAddr)</td></tr> <tr><td colspan="4">Server Name (Sname)</td></tr> <tr><td colspan="4">Boot File Name (128 bytes)</td></tr> <tr><td colspan="4">Vendor-Specific Area (variable size) / allows more options</td></tr> </table>	Operation code	HW type	HW addr length	Hops	Transaction Identifier				Seconds		Flags		Client IP address (CIAddr)				Your IP address (YIAddr)				Server IP address (SIAddr)				Gateway IP address (GIAddr)				Client Hardware Address (CHAddr)				Server Name (Sname)				Boot File Name (128 bytes)				Vendor-Specific Area (variable size) / allows more options				<p>DHCP Relay</p>	<p>interface FastEthernet 0/0 ip helper-address 1.2.3.4</p>	<p>DHCP option 82</p>	<pre> DHCP option 82 [5216 020C 020A 00009B013A0500000000 0606 564C414E35] option 82 (0x52) of total length 22 (0x16) First suboption remote-id (0x02) total length 12 (0xC) Remote-id TLV (type 0x02 length 0x0A) Remote-id value 10 bytes, Next suboption is 0x06 length 0x06 ASCII Value </pre>
Operation code	HW type	HW addr length	Hops																																														
Transaction Identifier																																																	
Seconds		Flags																																															
Client IP address (CIAddr)																																																	
Your IP address (YIAddr)																																																	
Server IP address (SIAddr)																																																	
Gateway IP address (GIAddr)																																																	
Client Hardware Address (CHAddr)																																																	
Server Name (Sname)																																																	
Boot File Name (128 bytes)																																																	
Vendor-Specific Area (variable size) / allows more options																																																	
<p>DHCP Server</p>	<pre> service dhcp ip dhcp excluded-address 155.1.146.100 155.1.146.254 ip dhcp pool VLAN146 network 155.1.146.0 /24 default-router 155.1.146.6 155.1.146.4 dns-server 155.1.146.6 155.1.146.4 lease 0 12 no ip bootp server ignore BOOTP ip dhcp database flash:/bindings </pre>	<p>DHCP Host Pools</p>	<pre> R6#sh ip dhcp binding Bindings from all pools not associated with VRF: IP address Client-ID/ Lease expiration Type Hardware address/ User name 155.1.146.1 0063.6973.636f.2d30. Infinite Manual 3031.332e.3766.3766. 2e36.3261.302d.4661. 302f.30 ip dhcp pool R1_HOST host 155.1.146.1 255.255.255.0 client-identifier 0063.6973.636f.2d30.3031.3..... ip address / no ip address dhcp </pre>	<p>DHCP Authorized ARP</p>	<ul style="list-style-type: none"> - Allows the DHCP process to populate the ARP cache with the DHCP-based entries - Authorized ARP enabled at interface level will disable any dynamic ARP learning on that interface. <pre> ip dhcp pool VLAN146 update arp ip dhcp pool R1_HOST update arp interface fa0/0.146 arp authorize arp timeout <seconds> arp probe interval <sec> count 15 </pre>																																												
<p>Show ip dhcp pool Output:</p>	<pre> R6#show ip dhcp pool Pool VLAN146 : Utilization mark (high/low) : 100 / 0 Subnet size (first/next) : 0 / 0 Total addresses : 254 Leased addresses : 0 Pending event : none 1 subnet is currently in the pool : Current index IP address range Leased addresses 155.1.146.1 155.1.146.1 - 155.1.146.254 0 </pre>	<p>DHCP on-Demand Pool</p>	<pre> R1: interface Serial 0/1 encapsulation ppp ip address negotiated peer default ip address 155.1.13.1 ppp ipcp mask request ppp ipcp dns request no peer neighbor-route ip dhcp pool ODAP_POOL import all origin ipcp router rip no validate-update-source R3: interface Serial 1/2 encapsulation ppp ip address 155.1.13.3 255.255.255.0 peer default ip address 155.1.13.1 ppp ipcp mask 255.255.255.0 ppp ipcp dns 155.1.146.4 155.1.146.6 no peer neighbor-route R1 imports DHCP infos via IPCP from R3 (RIP functional via DHCP received IP) </pre>	<p>IP SLA</p>	<pre> ip sla 2 tcp-connect 54.1.1.254 23 control disable timeout 5000 ip sla schedule 2 life forever start-time now ip sla monitor 2 type tcpConnect dest-ipaddr 54.1.1.254 dest-port 23 control disable timeout 5000 ip sla monitor schedule 2 life forever start-time now </pre>																																												
<p>show ip dhcp database</p>	<pre> R6#show ip dhcp database URL : flash:/bindings Read : Never Written : Never Status : Database has changed. A file transfer to the agent is pending. Delay : 300 seconds Timeout : 300 seconds Failures : 0 Successes : 0 </pre>	<p>DHCP on-Demand Pool</p> <p>Verifying the import:</p>	<pre> R1#show ip dhcp import Address Pool Name: ODAP_POOL Domain Name Server(s): 155.1.146.4 155.1.146.6 R1: show ip dhcp pool </pre>	<p>Object Tracking</p>	<pre> track 1 rtr 2 delay down 15 up 10 track 2 rtr 3 track 3 list boolean and object 1 object 2 track 4 list boolean or object 1 object 2 show track 3 Track 3 List boolean and Boolean AND is Down 3 changes, last change 00:01:45 object 1 Up object 2 Down show track 1 Track 1 Response Time Reporter 2 state State is Up 1 change, last change 00:05:10 Latest operation return code: OK Latest RTT (milliseconds) 16 Tracked by: Track-list 3 Track-list 4 show track 4 Track 4 List boolean or Boolean OR is Up 2 changes, last change 00:00:30 object 1 Up object 2 Down </pre>																																												

Help me create more flashcards:


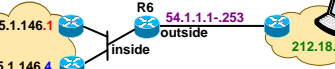
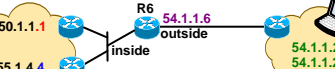
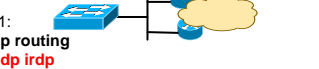
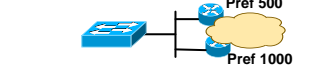
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

<h2>HSRP</h2>	<p>R1# interface FastEthernet 0/0.146 standby 146 ip 155.1.146.254 standby 146 timers 1 3 standby 146 preempt standby 146 authentication md5 key-string CISCO standby 146 name VLAN146 standby 146 priority 110</p> <p>R2# interface FastEthernet 0/1 standby 146 ip 155.1.146.254 standby 146 timers 1 3 standby 146 preempt standby 146 authentication md5 key-string CISCO standby 146 name VLAN146 standby 146 priority 105</p> <p><i>Version 1 224.0.0.2 UDP port 1985</i></p> <p><i>standby version 2 224.0.0.102 UDP port 1985</i></p>	<h2>Router Redundancy and Object Tracking</h2>	<pre>track 16 rtr 99 ip sla monitor 99 type tcpConnect dest-ipaddr 54.1.1.254 dest-port 23 control disable timeout 5000 ip sla monitor schedule 99 life forever start-time now interface Gi 0/0 standby 146 track 16 decrement 20 track 12 ip route 30.0.0.0/16 reachability interface FastEthernet 0/1 vrrp 146 track 12 decrement 20</pre>	<h2>Basic NAT</h2>	 <pre>ip nat pool VLAN43 204.12.1.1 204.12.1.253 prefix-length 24 ip access-list extended NAT_TRAFFIC permit ip 155.1.0.0 0.0.255.255 any ip nat inside source list NAT_TRAFFIC pool VLAN43 interface g0/0 ip nat outside interface g0/1 ip nat inside</pre>																									
<p>show standby brief</p> <p>show standby all</p> <p>Output:</p>	<p>R4#show standby all GigabitEthernet0/1 - Group 146 State is Standby 1 state change, last state change 00:00:38 Virtual IP address is 155.1.146.254 Active virtual MAC address is 0000.0c07.ac92 Local virtual MAC address is 0000.0c07.ac92 (v1 default) Hello time 1 sec, hold time 3 sec Next hello sent in 0.984 secs Authentication MD5, key-string Preemption enabled Active router is 155.1.146.6, priority 110 (expires in 2.164 sec) Standby router is local Priority 100 (default 100) IP redundancy name is "VLAN146" (crgd)</p> <p><i>92 hex = 146 decimal</i></p>	<p>track 10 ip route 10.2.21.128/25 metric threshold threshold metric up 20 down 50</p> <p>described</p>	<pre>track 10 ip route 10.2.21.128/25 metric threshold threshold metric up 20 down 50</pre> <p>If the metric goes beyond 50 the tracking object 10 is going invalid. The metric has then to go below 20 to become active again. (Hysteresis)</p>	<h2>Basic NAT verification</h2> 	<p>R6#show ip nat translations</p> <table border="1"> <thead> <tr> <th>Pro</th> <th>Inside global</th> <th>Inside local</th> <th>Outside local</th> <th>Outside global</th> </tr> </thead> <tbody> <tr> <td>icmp</td> <td>54.1.1.7:3</td> <td>155.1.146.1:3</td> <td>212.18.1.1:3</td> <td>212.18.1.1:3</td> </tr> <tr> <td>--</td> <td>54.1.1.7</td> <td>155.1.146.1</td> <td>--</td> <td>--</td> </tr> <tr> <td>icmp</td> <td>54.1.1.6:3</td> <td>155.1.146.4:3</td> <td>212.18.1.1:3</td> <td>212.18.1.1:3</td> </tr> <tr> <td>--</td> <td>54.1.1.6</td> <td>155.1.146.4</td> <td>--</td> <td>--</td> </tr> </tbody> </table>	Pro	Inside global	Inside local	Outside local	Outside global	icmp	54.1.1.7:3	155.1.146.1:3	212.18.1.1:3	212.18.1.1:3	--	54.1.1.7	155.1.146.1	--	--	icmp	54.1.1.6:3	155.1.146.4:3	212.18.1.1:3	212.18.1.1:3	--	54.1.1.6	155.1.146.4	--	--
Pro	Inside global	Inside local	Outside local	Outside global																										
icmp	54.1.1.7:3	155.1.146.1:3	212.18.1.1:3	212.18.1.1:3																										
--	54.1.1.7	155.1.146.1	--	--																										
icmp	54.1.1.6:3	155.1.146.4:3	212.18.1.1:3	212.18.1.1:3																										
--	54.1.1.6	155.1.146.4	--	--																										
<h2>VRRP</h2>	<p>R6: interface FastEthernet 0/0.146 vrrp 146 ip 155.1.146.253 vrrp 146 timers advertise 3 vrrp 146 authentication CISCO</p> <p>R4: interface FastEthernet 0/1 vrrp 146 ip 155.1.146.253 vrrp 146 timers advertise 3 vrrp 146 authentication CISCO vrrp 146 priority 110</p> <p><i>SRC interface IP DST 224.0.0.18 protocol 112</i></p>	<h2>IRDP</h2>	<p>ICMP Router Discovery Protocol (IRDP) allows IPv4 hosts to locate routers that provide IPv4 connectivity to other (nonlocal) IP networks</p> <p>R5 interface FastEthernet 0/0 ip irdp ip irdp address 155.1.58.5 1000 ip irdp maxadvertinterval 20 ip irdp minadvertinterval 10</p> <p>R2: interface Vlan 58 ip irdp ip irdp address 155.1.58.8 500 ip irdp maxadvertinterval 20 ip irdp minadvertinterval 10</p> <p><i>Instead using HSRP, VRRP, GLBP</i></p> <p><i>ICMP type 9 code 0 router advertisement</i></p>	<h2>NAT Overload (PAT, port address translation)</h2> 	<p>R6#show ip nat translations</p> <table border="1"> <thead> <tr> <th>Pro</th> <th>Inside global</th> <th>Inside local</th> <th>Outside local</th> <th>Outside global</th> </tr> </thead> <tbody> <tr> <td>icmp</td> <td>54.1.1.6:10</td> <td>150.1.1.1:10</td> <td>54.1.1.254:10</td> <td>54.1.1.254:10</td> </tr> <tr> <td>icmp</td> <td>54.1.1.6:11</td> <td>150.1.1.1:11</td> <td>54.1.1.222:11</td> <td>54.1.1.222:11</td> </tr> <tr> <td>icmp</td> <td>54.1.1.6:1</td> <td>150.1.4.4:11</td> <td>54.1.1.254:11</td> <td>54.1.1.254:11</td> </tr> <tr> <td>icmp</td> <td>54.1.1.6:12</td> <td>150.1.4.4:12</td> <td>54.1.1.222:12</td> <td>54.1.1.222:12</td> </tr> </tbody> </table> <p>NAT overload PAT addr 54.1.1.6</p>	Pro	Inside global	Inside local	Outside local	Outside global	icmp	54.1.1.6:10	150.1.1.1:10	54.1.1.254:10	54.1.1.254:10	icmp	54.1.1.6:11	150.1.1.1:11	54.1.1.222:11	54.1.1.222:11	icmp	54.1.1.6:1	150.1.4.4:11	54.1.1.254:11	54.1.1.254:11	icmp	54.1.1.6:12	150.1.4.4:12	54.1.1.222:12	54.1.1.222:12
Pro	Inside global	Inside local	Outside local	Outside global																										
icmp	54.1.1.6:10	150.1.1.1:10	54.1.1.254:10	54.1.1.254:10																										
icmp	54.1.1.6:11	150.1.1.1:11	54.1.1.222:11	54.1.1.222:11																										
icmp	54.1.1.6:1	150.1.4.4:11	54.1.1.254:11	54.1.1.254:11																										
icmp	54.1.1.6:12	150.1.4.4:12	54.1.1.222:12	54.1.1.222:12																										
<p>show vrrp all</p> <p>Output:</p>	<p>R6#show vrrp all GigabitEthernet0/0.146 - Group 146 State is Backup Virtual IP address is 155.1.146.253 Virtual MAC address is 0000.5e00.0192 Advertisement interval is 3.000 sec Preemption enabled Priority is 100 Authentication text, string "CISCO" Master Router is 155.1.146.4, priority is 110 Master Advertisement interval is 3.000 sec Master Down interval is 9.609 sec (expires in 7.877 sec)</p>	<h2>IRDP config</h2>	<p>SW1: no ip routing ip gdp irdp</p> <p>interface Vlan 58 ip address 155.1.58.7 255.255.255.0</p> <p>R1 interface FastEthernet 0/0 ip irdp ip irdp address 155.1.58.5 1000 ip irdp maxadvertinterval 20 ip irdp minadvertinterval 10</p> <p>R2 interface Vlan 58 ip irdp ip irdp address 155.1.58.8 500 ip irdp maxadvertinterval 20 ip irdp minadvertinterval 10</p> 	<h2>NAT with Route Maps</h2> <p>- If from Loopback PAT to 155.1.23.200 - If NOT from Loopback PAT to Serial0/1 addr</p>	<pre>ip access-list standard FROM_LOOPBACK permit 150.1.2.0 0.0.0.255 ip nat pool PPP_LOOPBACK_POOL 155.1.23.200 155.1.23.200 prefix-length 24 route-map NAT_OUT_PPP_FROM_LOOPBACK permit 10 match ip address FROM_LOOPBACK match interface Serial0/1 route-map NAT_OUT_PPP_NOT_FROM_LOOPBACK deny 10 match ip address FROM_LOOPBACK match interface Serial0/1 route-map NAT_OUT_PPP_NOT_FROM_LOOPBACK permit 20 match interface Serial0/1 ip nat inside source route-map NAT_OUT_PPP_FROM_LOOPBACK pool PPP_LOOPBACK_POOL overload ip nat inside source route-map NAT_OUT_PPP_NOT_FROM_LOOPBACK interface Serial0/1 overload</pre>																									
<h2>GLBP</h2> <p><i>SRC interface IP DST 224.0.0.102 UDP 3222</i></p>	<p>R6: interface g0/0.146 glbp 146 ip 155.1.146.252 glbp 146 timers 1 3 glbp 146 priority 110 glbp 146 preempt glbp 146 weighting 10 1x glbp 146 name VLAN146 glbp 146 authentication md5 key-string CISCO glbp 146 load-balancing weighted host-dependent</p> <p>R4: interface g0/1 glbp 146 ip 155.1.146.252 glbp 146 timers 1 3 glbp 146 preempt glbp 146 weighting 20 2x glbp 146 name VLAN146 glbp 146 authentication md5 key-string CISCO glbp 146 load-balancing weighted round-robin</p> <p><i>AVG responds to ARP requests R6 is elected AVG, (R4 and R6 AVF) Load distribution R4 and R6 in 2:1 ratio</i></p>	<h2>IRDP Show commands</h2>	<p>SW1#show ip route</p> <table border="1"> <thead> <tr> <th>Gateway</th> <th>Using</th> <th>Interval</th> <th>Priority</th> <th>Interface</th> </tr> </thead> <tbody> <tr> <td>155.1.58.5</td> <td>IRDP</td> <td>30</td> <td>1000</td> <td>Vlan58</td> </tr> </tbody> </table> <p>SW2#show ip irdp vlan 58 Vlan58 has router discovery enabled</p> <p>Advertisements will occur between every 10 and 20 seconds. Advertisements are sent with broadcasts. Advertisements are valid for 60 seconds. Default preference will be 0. Proxy for 155.1.58.8 with preference 500.</p> 	Gateway	Using	Interval	Priority	Interface	155.1.58.5	IRDP	30	1000	Vlan58	<h2>show ip interface part 1:</h2>	<p>R1#show ip interface FastEthernet0/0 is up, line protocol is up Internet address is 155.1.146.1/24 Broadcast address is 255.255.255.255 Address determined by setup command MTU is 1500 bytes Helper address is not set Directed broadcast forwarding is disabled Multicast reserved groups joined: 224.0.0.9 Outgoing access list is not set Inbound access list is not set Proxy ARP is enabled Local Proxy ARP is disabled Security level is default Split horizon is enabled ICMP redirects are never sent ICMP unreachable are never sent ICMP mask replies are never sent</p>															
Gateway	Using	Interval	Priority	Interface																										
155.1.58.5	IRDP	30	1000	Vlan58																										
<p>R4#show glbp GigabitEthernet0/1 - Group 146 State is Standby 1 state change, last state change 00:08:44 Virtual IP address is 155.1.146.252 Hello time 1 sec, hold time 3 sec Next hello sent in 0.824 secs Redirect time 600 sec, forwarder time-out 14400 sec Authentication MD5, key-string Preemption enabled, min delay 0 sec Active is 155.1.146.6, priority 110 (expires in 2.580 sec) Standby is local Priority 100 (default) Weighting 20 (configured 20), thresholds: lower 1, upper 20 Load balancing: weighted IP redundancy name is "VLAN146" Group members: 0026.0b57.b960 (155.1.146.6) authenticated 0026.0b57.ba61 (155.1.146.4) local There are 2 forwarders (1 active)</p>	<p>Forwarder 1 State is Listen MAC address is 0007.b400.9201 (learned) Owner ID is 0026.0b57.b960 Time to live: 14399.340 sec (maximum 14400 sec) Preemption enabled, min delay 30 sec Active is 155.1.146.6 (primary), weighting 10 (expires in 2.920 sec)</p> <p>Forwarder 2 State is Active 1 state change, last state change 00:08:50 MAC address is 0007.b400.9202 (default) Owner ID is 0026.0b57.ba61 Preemption enabled, min delay 30 sec Active is local, weighting 20</p>	<h2>Router ICMP Settings</h2>	<p>stop sending ICMP messages about:</p> <ul style="list-style-type: none"> - Discarded packets - ICMP messages to select a better next-hop, - ICMP messages reporting subnets mask <p>interface FastEthernet 0/0 no ip redirects <i>R1#show ip interface</i> no ip unreachable <i>FastEthernet0/0 is up, line protocol is up</i> no ip mask-reply <i>ICMP redirects are never sent</i> <i>ICMP unreachable are never sent</i> <i>ICMP mask replies are never sent</i> ...</p> <p><i>Rate-limit ICMP unreachable to 2 per-second globally</i></p> <p>ip icmp rate-limit unreachable 500 ! <msec></p>	<h2>show ip interface part 2:</h2>	<p>IP fast switching is enabled IP fast switching on the same interface is disabled IP Flow switching is disabled IP CEF switching is enabled IP CEF Fast switching turbo vector IP multicast fast switching is enabled IP multicast distributed fast switching is disabled IP route-cache flags are Fast, CEF Router Discovery is disabled IP output packet accounting is disabled IP access violation accounting is disabled TCP/IP header compression is disabled RTP/IP header compression is disabled Policy routing is disabled Network address translation is disabled BGP Policy Mapping is disabled WCCP Redirect outbound is disabled WCCP Redirect inbound is disabled WCCP Redirect exclude is disabled</p>																									

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

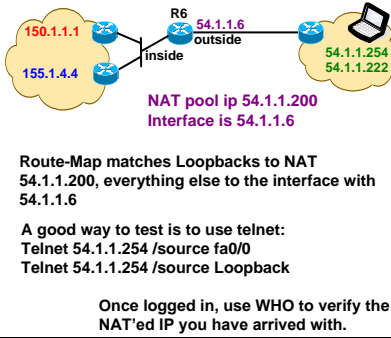
[Donate](#)

Thanks for appreciating my efforts

Colin

How to troubleshoot NAT

With different pools / interface PATs



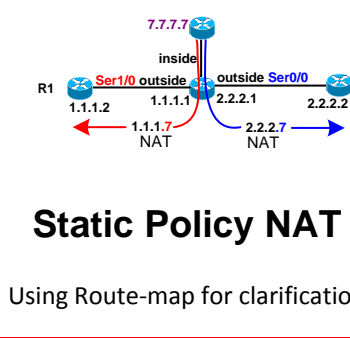
Route-Map matches Loopbacks to NAT 54.1.1.200, everything else to the interface with 54.1.1.6

A good way to test is to use telnet: Telnet 54.1.1.254 /source fa0/0 Telnet 54.1.1.254 /source Loopback

Once logged in, use WHO to verify the NAT'ed IP you have arrived with.

Static Policy NAT

Using Route-map for clarification



route-map TO_R1
match interface serial1/0

route-map TO_R2
match interface serial0/0

ip nat inside source static 7.7.7.7 1.1.1.7 route-map TO_R1

ip nat inside source static 7.7.7.7 2.2.2.7 route-map TO_R2

NAT Virtual Interface

NAT direction is always "inside" for NVI based NAT routing lookup is always performed before the translation after routing, packet source is translated

interface Serial 0/1/0
ip nat enable

interface FastEthernet 0/0
ip nat enable

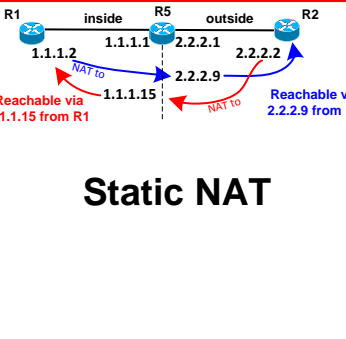
ip access-list standard VLAN8
permit 155.1.8.0 0.0.0.255

ip nat pool NVI_POOL 155.1.188.1 155.1.188.254 prefix 24 add-route

ip nat source list VLAN8 pool NVI_POOL

R5#show ip route static
S 155.1.188.0/24 [0/0] via 0.0.0.0, NVI0

Static NAT



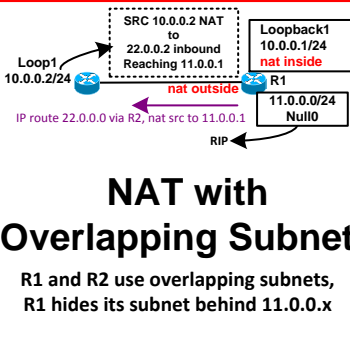
R5:
interface Serial0/1/0
ip nat outside
interface g0/0
ip nat inside

ip nat inside source static 1.1.1.2 2.2.2.9
ip nat outside source static 2.2.2.2 1.1.1.15

ip route 1.1.1.15 255.255.255.255 2.2.2.2

NAT with Overlapping Subnets

R1 and R2 use overlapping subnets, R1 hides its subnet behind 11.0.0.x



R1:
ip nat pool R2_MASQ 22.0.0.1 22.0.0.254 prefix-length 24
ip access-list extended R2_LOOPBACK1
permit ip 10.0.0.0 0.0.0.255 any

ip nat outside source list R2_LOOPBACK1 pool R2_MASQ

ip nat inside source static network 10.0.0.0 11.0.0.0/24

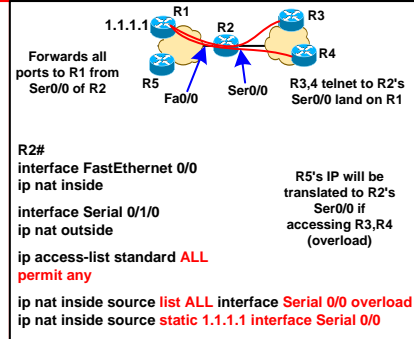
ip route 11.0.0.0 255.255.255.0 Null 0

ip route 22.0.0.0 255.255.255.0 Serial 0/1

router rip
redistribute static

NAT Default Interface

(total portforwarding / PAT to single IP)



R2#
interface FastEthernet 0/0
ip nat inside


interface Serial 0/1/0
ip nat outside

ip access-list standard ALL
permit any

ip nat inside source list ALL interface Serial 0/0 overload

ip nat inside source static 1.1.1.1 interface Serial 0/0

Troubleshooting Static NAT scenario



debug ip nat detail

NAT: I: icmp (1.1.1.2, 4) -> (1.1.1.15, 4) [9]

NAT: s=1.1.1.2->2.2.2.9, d=1.1.1.15 [9]

NAT: s=2.2.2.9, d=1.1.1.15->2.2.2.2 [9]

NAT*: O: icmp (2.2.2.2, 4) -> (2.2.2.9, 4) [9]

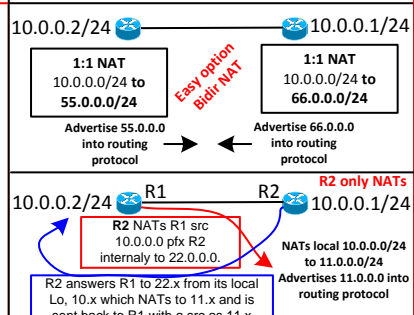
NAT*: s=2.2.2.2->1.1.1.15, d=2.2.2.9 [9]

NAT*: s=1.1.1.15, d=2.2.2.9->1.1.1.2 [9]

NAT with Overlapping Subnets

Solution 1: Bidirection NAT

Solution 2: NAT only on R2 performed, NO config on R1.



1:1 NAT
10.0.0.0/24 to 55.0.0.0/24

Easy option Bidir NAT

1:1 NAT
10.0.0.1/24 to 66.0.0.0/24

Advertise 55.0.0.0 into routing protocol

Advertise 66.0.0.0 into routing protocol

R2 NATs R1 src 10.0.0.0 pps R2 internally to 22.0.0.0

R2 answers R1 to 22.x from its local Lo, 10.x which NATs to 11.x and is sent back to R1 with a src as 11.x (R1 pings 11.x)

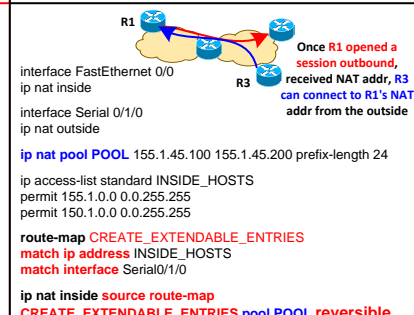
R2#
ip nat pool ROTARY prefix-length 24 type rotary
address 155.1.0.1 155.1.0.1
address 155.1.0.2 155.1.0.2
address <start-ip> <end-ip>

ip access-list extended LOAD_BALANCE
permit tcp any host 155.1.58.55 eq telnet

ip nat inside destination list LOAD_BALANCE pool ROTARY

ip alias 155.1.58.55 23

Reversible NAT



interface FastEthernet 0/0
ip nat inside

interface Serial 0/1/0
ip nat outside

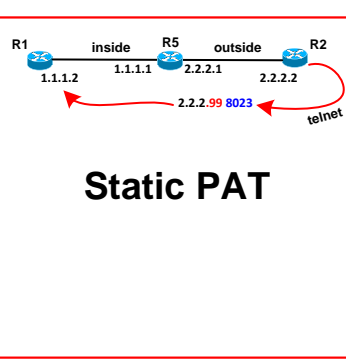
ip nat pool POOL 155.1.45.100 155.1.45.200 prefix-length 24

ip access-list standard INSIDE_HOSTS
permit 155.1.0.0 0.0.255.255
permit 150.1.0.0 0.0.255.255

route-map CREATE_EXTENDABLE_ENTRIES
match ip address INSIDE_HOSTS
match interface Serial0/1/0

ip nat inside source route-map CREATE_EXTENDABLE_ENTRIES pool POOL reversible

Static PAT



R5#
ip nat inside source static tcp 1.1.1.2 23 2.2.2.99 8023

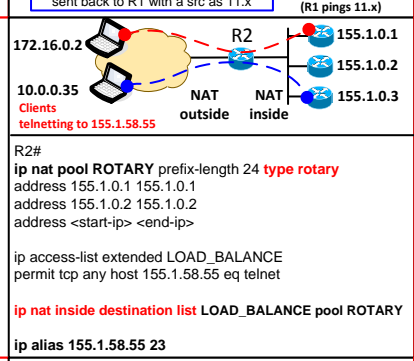
R2# telnet 2.2.2.99 8023

Login Prompt of R1

Telnet R2 to 2.2.2.99 8023 -> forwards to 1.1.1.2 23 (R1) through R5

TCP Load Distribution with NAT

(Rotary NAT)



R2#
ip nat pool ROTARY prefix-length 24 type rotary
address 155.1.0.1 155.1.0.1
address 155.1.0.2 155.1.0.2
address <start-ip> <end-ip>

ip access-list extended LOAD_BALANCE
permit tcp any host 155.1.58.55 eq telnet

ip nat inside destination list LOAD_BALANCE pool ROTARY

ip alias 155.1.58.55 23

Stateful NAT with Primary/Backup

Verification:

show ip snat distributed verbose

Stateful NAT Connected Peers

SNAT: Mode BACKUP

: State READY

: Local Address 155.1.146.4

: Local NAT id 2

: Peer Address 155.1.146.6

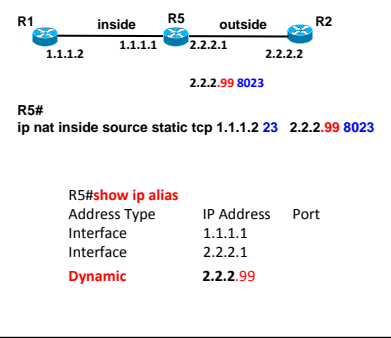
: Peer NAT id 1

: Mapping List 1

: InMsgs 15, OutMsgs 0, tcb 0x650FE6E0, listener 0x650FE22C

How to identify NAT addresses

Using show ip alias



R5#
ip nat inside source static tcp 1.1.1.2 23 2.2.2.99 8023

R5#show ip alias

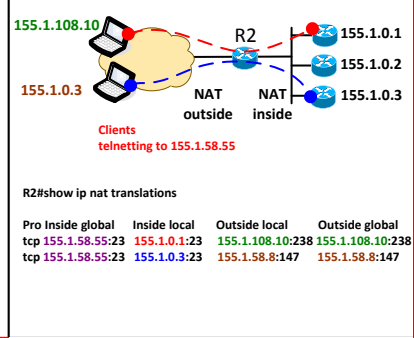
Address Type	IP Address	Port
Interface	1.1.1.1	
Interface	2.2.2.1	
Dynamic	2.2.2.99	

TCP Load Distribution with NAT

(Rotary NAT)

show ip nat translations

Output:



R2#show ip nat translations

Pro	Inside global	Inside local	Outside local	Outside global
tcp	155.1.58.55:23	155.1.0.1:23	155.1.108.10:238	155.1.108.10:238
tcp	155.1.58.55:23	155.1.0.3:23	155.1.58.8:147	155.1.58.8:147

R2#show ip nat translations

Pro Inside global Inside local Outside local Outside global

tcp 155.1.58.55:23 155.1.0.1:23 155.1.108.10:238 155.1.108.10:238

tcp 155.1.58.55:23 155.1.0.3:23 155.1.58.8:147 155.1.58.8:147

interface FastEthernet 0/0.146
standby 146 ip 155.1.146.254
standby 146 timers 1 3
standby 146 name VLAN146

ip nat stateful id 1 redundancy VLAN146 mapping-id 1

ip access-list standard NAT_LIST
permit 155.1.0.0 0.0.255.255

ip nat pool SHARED_POOL 155.1.254.1 155.1.254.254 prefix-length 24

ip nat inside source list NAT_LIST pool SHARED_POOL mapping 1

ip route 155.1.254.0 255.255.255.0 Null 0

Stateful NAT with Primary/Backup

SNAT

Config:

R4 & R6

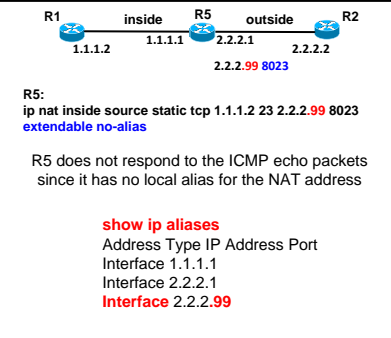
ip access-list standard NAT_LIST
permit 155.1.0.0 0.0.255.255

ip nat pool SHARED_POOL 155.1.254.1 155.1.254.254 prefix-length 24

ip nat inside source list NAT_LIST pool SHARED_POOL mapping 1

ip route 155.1.254.0 255.255.255.0 Null 0

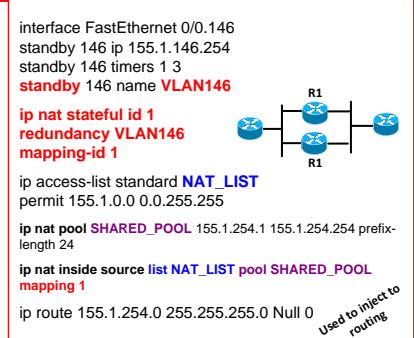
Static NAT and IP Aliasing



Stateful NAT with HSRP

(SNAT)

config



Stateful NAT with HSRP

(SNAT)

verification

R6#show ip snat distributed verbose

Stateful NAT Connected Peers

SNAT: Mode IP-REDUNDANCY :: ACTIVE

: State READY

: Local Address 155.1.146.6

: Local NAT id 1

: Peer Address 155.1.146.4

: Peer NAT id 2

: Mapping List 1

: InMsgs 4, OutMsgs 0, tcb 0x44C99F8, listener 0x0

R2

Things that did not work in the lab:

Stateful NAT prim backup or Redundancy

Overlapping

Stateful with HSRP

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Donate

colin

Thanks for appreciating my efforts

Colin

<p>Static Extendable NAT</p>	<pre> R5# ip nat inside source static 1.1.1.1 192.168.0.20 extendable ip nat inside source static 1.1.1.1 192.168.0.99 extendable R5#show ip nat translations Pro Inside global Inside local Outside local Outside global --- 192.168.0.20 1.1.1.1 --- --- --- 192.168.0.99 1.1.1.1 --- --- R2 or R3 can either Access R1 via 192.168.0.20 or 192.168.0.99 </pre>	<p>TCP Optimization</p>	<p>???????</p>	<p>WCCPv2 Services</p>	<pre> access-list 10 permit 155.1.58.100 access-list 20 permit 155.1.58.0 0.0.0.255 ip wccp version 2 ip wccp 50 group-address 224.0.1.100 redirect-list 20 group-list 10 password CISCO interface FastEthernet 0/0 ip wccp 50 redirect in ip wccp 50 group-listen </pre>
<p>IP Precedence Accounting</p>	<pre> 01100000 = 96 First 3 bits of TOS field = IP precedence interface Serial 0/0/0 ip accounting precedence input ip accounting precedence output R6#show interfaces serial 0/0/0 precedence Serial0/0/0 Input Precedence 0: 5 packets, 520 bytes Precedence 3: 5 packets, 520 bytes Precedence 6: 53 packets, 4304 bytes Output Precedence 0: 7 packets, 608 bytes Precedence 3: 5 packets, 520 bytes Precedence 6: 24 packets, 7936 bytes </pre>	<p>Verify TCP operation:</p>	<pre> show tcp brief all TCB Local Address Foreign Address (state) 856D54A0 54.1.1.6.60615 54.1.1.254.179 ESTAB R6#show tcp tcb 856BE248 Connection state is ESTAB, I/O status: 1, unread input bytes: 0 Connection is ECN Enabled, Minimum incoming TTL 0, Outgoing TTL 255 Local host: 155.1.146.6, Local port: 33209 Foreign host: 150.1.1.1, Foreign port: 23 Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 Event Timers (current time is 0x1080595): Timer Starts Wakeups Next Retrans 6 0 0x0 TimeWait 0 0 0x0 AckHold 3 1 0x0 SendWnd 0 0 0x0 KeepAlive 0 0 0x0 GiveUp 0 0 0x0 PmtuAger 0 0 0x0 DeadWait 0 0 0x0 </pre>	<p>NBAR Protocol Discovery</p>	<pre> interface Serial 0/0/0 ip nbar protocol-discovery ip nbar custom HTTP_PROXY destination tcp 3128 8080 classify connections to HTTP proxy ports 3128 and 8080 ip nbar port-map ftp tcp 2121 Change ftp protocol mapping to port 2121 ip nbar custom TEST 0 ascii A destination tcp 3001 Match ASCII "A" in the beginning of a TCP segment flowing to the destination port 3001. show ip nbar port-map show ip nbar protocol-discovery show ip nbar protocol-discovery protocol TEST </pre>
<p>IP Output Packet Accounting</p> <p>Config:</p>	<pre> ip accounting-threshold 4096 Limit database to 4096 entries ip accounting-transits 1 Set number of packets to 1 for any packets, not matching the accounting list (first packet of flow only) ip accounting-list 155.1.0.0 0.0.255.255 Only account for packets going to 155.1.0.0/16 interface FastEthernet 0/0 ip accounting output-packets </pre>	<p>IOS Small Services & Finger</p>	<pre> service udp-small-servers service tcp-small-servers ip finger TCP to port 7: echo's all characters sent back to prompt TCP to port 9: blackhole, nothing comes back TCP to port 13: receive Date and Time, exists session TCP to port 19: CHARGEN, does not stop! TCP to port 79: list of currently logged in Users </pre>	<p>Netflow Ingress & Egress</p> <p>Version 5</p>	<pre> ip flow-capture vlan-id (collect VlanIDs) ip flow-capture icmp (collect ICMP too) ip flow-export version 5 ip flow-export destination 155.1.146.100 9999 ip flow-export version 5 origin-as ip flow-cache entries 4096 (DB max entries) interface Serial 0/0/0 ip flow ingress Ingress: including the packets destined to the router itself, applied before ACLs and rate-limiting Egress: does not include packets originated from the router itself (used with MPLS, monitor leaving untagged traffic) </pre>
<p>IP Output Packet Accounting</p> <p>Show output:</p>	<pre> R1#show ip accounting Source Destination Packets Bytes 155.1.146.4 -> 155.1.13.3 2 200 155.1.13.3 <- 155.1.146.4 2 200 155.1.146.6 155.1.13.3 20 2000 155.1.13.3 155.1.146.6 20 2000 Accounting data age is 0 R1# clear ip accounting (stores the old accounting database into a checkpoint) Checkpoint can be retrieved by: show ip accounting checkpoint </pre>	<p>Directed Broadcasts & UDP Forwarding</p>	<pre> R2# no ip forward-protocol udp bootps no ip forward-protocol udp tftp no ip forward-protocol udp time no ip forward-protocol udp netbios-ns no ip forward-protocol udp netbios-dgm no ip forward-protocol udp tacacs no ip forward-protocol udp taccacs interface Serial 0/0 ip helper-address 155.1.58.255 (overrides the default of generating a 255.255.255.255, sending to 155.1.58.255 instead) interface FastEthernet 0/0 ip broadcast-address 155.1.58.255 ip directed-broadcast </pre>	<p>Netflow Top Talkers</p>	<pre> ip flow-top-talkers top 10 sort-by packets match source address 155.1.0.0 255.255.0.0 match protocol 1 Showing top talkers for ICMP traffic of hosts in 155.1.0.0/16 R4#show ip flow top-talkers SrcIf SrcIPAddress DstIf DstIPAddress Pr SrcP DstP Pkts Et0/1 155.1.146.1 Et0/0 30.0.0.1 01 0000 0800 50 </pre>
<p>IP Access Violation Accounting</p>	<pre> - deny packets with an IP precedence of 4 coming from X coming from Ser0/0/0 - Account for the packets denied by this policy access-list 101 deny ip any any precedence 4 access-list 101 permit ip any any interface Serial 0/0/0 ip access-group 101 in ip accounting access-violations R6#show ip accounting access-violations Source Destination Packets Bytes ACL 112.0.0.1 54.1.1.6 5 500 101 </pre> <p><i>Does not work with named ACLs!</i></p>	<p>DRP Server Agent</p>	<pre> - allow connections from Directors from VLAN 14 - authenticate with a password of CISCO ip drp server ip drp access-group 99 ip drp authentication key-chain DRP key chain DRP key 1 key-string CISCO access-list 99 permit 155.1.146.0 0.0.0.255 show ip drp </pre>	<p>Netflow Aggregation Cache</p> <p>Version 8</p>	<pre> ip flow-aggregation cache destination-prefix cache entries 1024 export destination 155.1.146.100 9998 mask destination minimum 8 Enabled R4#show ip cache flow aggregation destination-prefix Minimum destination mask is configured to /8 Dst If Dst Prefix Msk AS Flows Pkts B/Pk Active Null 204.0.0.0 /8 0 1 1 44 0.0 Gi0/0 204.12.1.0 /24 0 1 28 41 3.7 Gi0/1 155.1.146.0 /24 0 1 25 91 3.7 </pre> <p><i>Don't summarize more than /8 in length</i></p>
<p>MAC Address Accounting</p>	<pre> interface FastEthernet0/0 ip accounting mac-address input ip accounting mac-address output R1#show interfaces Fa0/0 mac-accounting FastEthernet0/0 Input (510 free) 0026.0b57.ba61(161): 2 packets, 1092 bytes, last: 3564ms ago 0026.0b57.b960(163): 2 packets, 612 bytes, last: 6796ms ago Total: 4 packets, 1704 bytes Output (511 free) 0100.5e00.0009(86): 2 packets, 852 bytes, last: 9444ms ago Total: 2 packets, 852 bytes </pre>	<p>WCCPv1 Web-Cache</p>	<pre> access-list 199 deny ip 155.1.146.0 0.0.0.255 any access-list 199 permit ip any any ip wccp version 1 ip wccp web-cache redirect-list 199 ip wccp outbound-acl-check interface FastEthernet 0/0.146 ip wccp web-cache redirect in show ip wccp show ip wccp interface </pre>	<p>Netflow Random Sampling</p>	<pre> Only samples every 10th packet in each flow flow-sampler-map SAMPLER mode random one-out-of 10 policy-map NETFLOW_MAP class class-default netflow-sampler SAMPLER interface Serial 0/0/0 service-policy output NETFLOW_MAP </pre>


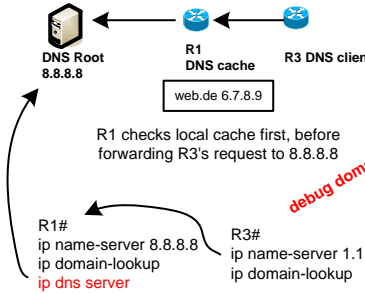
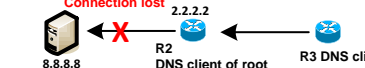
Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts



Colin

<h3>Netflow Input Filters</h3> <p>- Every packet is sampled for sources on VLAN 146</p> <p>- Other packets should still be randomly sampled</p>	<pre>ip access-list extended VLAN146 permit ip 155.1.146.0 0.0.0.255 any permit ip any 155.1.146.0 0.0.0.255 class-map VLAN146 match access-group name VLAN146 policy-map NETFLOW_MAP class VLAN146 netflow-sampler NORMAL class class-default netflow-sampler SAMPLER flow-sampler-map NORMAL mode random one-out-of 1 flow-sampler-map SAMPLER mode random one-out-of 10 interface Serial 0/0/0 service-policy output NETFLOW_MAP</pre>	<h3>IP Event Dampening</h3> <pre>dampening <half-life time> <start using> <start suppressing> <max suppress duration t> <enable restart suppression> <penalty value at restart></pre> <pre>interface Serial 0/0/0 dampening 30 1000 2000 60 restart 2000</pre> <p>- after a reload IP is not advertised into IGP for 30 seconds</p> <p>- if connection flaps: it does not disappear for more than 60 seconds from the routing table no matter how much penalty it accumulates</p> <p>To find default values: int X dampening Do show dampening</p>	<h3>DHCP Client options:</h3> <pre>interface GigabitEthernet 0/0/1 ip dhcp client client-id ascii my-test1 ip dhcp client class-id my-class-id ip dhcp client lease 0 1 0 ip dhcp client hostname host1 no ip dhcp client request tftp-server-address ip address dhcp</pre> <p>int X no ip address dhcp ip address dhcp</p> <p>Or release dhcp</p>	<h3>IOS Authoritative DNS Server</h3> <pre>ip dns server ip dns primary cisco.com soa ns.cisco.com ccie.com ip dns primary <domain> <primary-srv> <DNS mailbox> <refresh t> <retry t> <Auth expire t> <Min TTL zone t> ip host cisco.com ns 155.1.146.4 (DNS srv local own entry) ip host R4.cisco.com 150.1.4.4 155.1.146.4 155.1.45.4 All IPs which should be resolved to R4.cisco.com</pre>	<h3>NAT terminology:</h3> <p>Inside Local An actual address assigned to an inside host</p> <p>Inside Global An inside address seen from the outside</p> <p>Outside Global An actual address assigned to an outside host</p> <p>Outside Local An outside address seen from the inside</p>																																			
<h3>Router as DNS client</h3> <p>Using two DNS server in a round-robin fashion.</p>	<pre>ip name-server 155.1.146.4 155.1.146.6 ip domain-lookup ip domain round-robin ip domain name cisco.com complete all unqualified domain-names with the name "cisco.com".</pre>	<h3>Using CHARGEN to simulate TCP streams:</h3>  <pre>service tcp-small-servers CHARGEN TCP Port 19</pre> <pre>R1#show tcp brief TCB Local Address Foreign Address (state) 847557E4 1.1.1.1.19 2.2.2.2.27318 ESTAB 848D87A0 1.1.1.1.23 2.2.2.2.44491 ESTAB</pre> <pre>R1#clear tcp tcb 847557E4 (confirm) [OK] R2# HJKLMNOPQRSTUVWXYZ[!*_abcdefghijklmnopqrstuvwxyz()-!@#%& JKLMNOPQRSTUVWXYZ[!*_abcdefghijklmnopqrstuvwxyz()-!@#%&</pre>	<h3>IPv4 Address Pool for DMVPN Spokes</h3> <pre>ip dhcp pool pool1 origin dhcp number 3 odap client client-id id1 interface gi0/0 target-server 192.168.10.1 origin dhcp subnet size initial /24 autogrow /24</pre> <p>First IP of first pool assigned to interface. If second pool is requested, a secondary IP will be created with the first IP address of the second pool on the same interface.</p>	<h3>First hop redundancy protocols</h3> <table border="1"> <thead> <tr> <th rowspan="2"></th> <th colspan="3">Attributes</th> </tr> <tr> <th>HSRP</th> <th>VRRP</th> <th>GLBP</th> </tr> </thead> <tbody> <tr> <td>Standard</td> <td></td> <td></td> <td></td> </tr> <tr> <td>Load Balancing</td> <td>No</td> <td>No</td> <td>Yes</td> </tr> <tr> <td>IPv6 Support</td> <td>Yes</td> <td>No</td> <td>Yes</td> </tr> <tr> <td>Transport</td> <td>UDP/1985</td> <td>IP/112</td> <td>UDP/3222</td> </tr> <tr> <td>Default Priority</td> <td>100</td> <td>100</td> <td>100</td> </tr> <tr> <td>Default Hello</td> <td>3 sec</td> <td>1 sec</td> <td>3 sec</td> </tr> <tr> <td>Multicast Group</td> <td>224.0.0.2</td> <td>224.0.0.18</td> <td>224.0.0.102</td> </tr> </tbody> </table>		Attributes			HSRP	VRRP	GLBP	Standard				Load Balancing	No	No	Yes	IPv6 Support	Yes	No	Yes	Transport	UDP/1985	IP/112	UDP/3222	Default Priority	100	100	100	Default Hello	3 sec	1 sec	3 sec	Multicast Group	224.0.0.2	224.0.0.18	224.0.0.102	<h3>HSRPv2 224.0.0.102</h3>
	Attributes																																							
	HSRP	VRRP	GLBP																																					
Standard																																								
Load Balancing	No	No	Yes																																					
IPv6 Support	Yes	No	Yes																																					
Transport	UDP/1985	IP/112	UDP/3222																																					
Default Priority	100	100	100																																					
Default Hello	3 sec	1 sec	3 sec																																					
Multicast Group	224.0.0.2	224.0.0.18	224.0.0.102																																					
<h3>DNS related commands:</h3> <p>show ip dns primary:</p> <p>show hosts:</p>	<pre>R4#show ip dns primary Primary for zone cisco.com: SOA information: Zone primary (MNAME): ns.cisco.com Zone contact (RNAME): ccie.com Refresh (seconds): 21600 Retry (seconds): 900 Expire (seconds): 7776000 Minimum (seconds): 86400 R4#show hosts Host Port Flags Age Type Address(es) cisco.com NA (perm,OK) 0 NS 155.1.146.4 0 21600 900 7776000 86400 SOA ns.cisco.com ccie.com R4.cisco.com None (perm,OK) 0 IP 150.1.4.4 155.1.146.4 155.1.45.4</pre>	<h3>Debug ip tcp transactions</h3> <p>Output:</p> <pre>R6#debug ip tcp transactions Reserved port 0 in Transport Port Agent for TCP IP type 0 TCP0: state was LISTEN -> SYNRCVD [23 -> 155.1.146.1(64022)] TCP: tcb 454C3FEC connection to 155.1.146.1:64022, peer MSS 536, MSS is 516 TCP: sending SYN, seq 62577531, ack 199749595 TCP0: Connection to 155.1.146.1:64022, advertising MSS 536 TCP0: state was SYNRCVD -> ESTAB [23 -> 155.1.146.1(64022)] TCB454C3FEC setting property TCP_TOS (11) 448CB564 TCP322: state was ESTAB -> FINWAIT1 [23 -> 155.1.146.1(64022)] TCP322: sending FIN TCP322: state was FINWAIT1 -> FINWAIT2 [23 -> 155.1.146.1(64022)] TCP322: state was FINWAIT2 -> TIMEWAIT [23 -> 155.1.146.1(64022)]</pre>	<h3>Enhanced Object Tracking</h3> <p>track 4 list threshold weight object 1 weight 15 object 2 weight 20 object 3 weight 30 threshold weight up 30 down 10</p> <p>Combination of Object 1 + 2 will cause Track 4 to be up. (15+20 = 35, up weight = 30)</p> <p>Or if Object 3 (30 and up weight = 30)</p>	<h3>Track using weight</h3>																																				
<h3>IOS Caching DNS Server</h3>  <pre>R1# ip name-server 8.8.8.8 ip domain-lookup ip dns server R3# ip name-server 1.1.1.1 ip domain-lookup</pre>	<h3>How to monitor interface based rate-limiting:</h3>	<h3>Debug ip tcp transactions</h3> <p>Output:</p> <pre>Config: access-list 100 permit icmp any any interface FastEthernet0/0 ip address 183.1.17.1 255.255.255.0 rate-limit output access-group 100 128000 12000 12000 conform-action transmit exceed-action drop show interface fa0/0 rate-limit FastEthernet0/0 Output matches: access-group 100 params: 128000 bps, 12000 limit, 12000 extended limit conformed 0 packets, 0 bytes; action: transmit exceeded 0 packets, 0 bytes; action: drop last packet: 27522152ms ago, current burst: 0 bytes last cleared 00:05:18 ago, conformed 0 bps, exceeded 0 bps</pre>	<h3>Enhanced Object Tracking</h3> <p>track timer {interface ip route} <15 seconds> (poll interval is default 15 seconds)</p> <p>track interface ip routing Tracks whether IP routing is enabled, tracking of an IP address on an interface that was acquired through DHCP or PPP IPCP</p> <p>track line-protocol Tracks the state of a line protocol</p> <p>ip dhcp client route track <number> Assign received IP address via DHCP to tracking object <number> To be dynamically added to IP SLA</p> <p>show track timers show track brief debug track</p>	<h3>track timer</h3> <h3>track interface</h3> <h3>track line-protocol</h3> <h3>show track CMDs</h3>																																				
<h3>IOS DNS Spoofing</h3>  <pre>R2# ip dns spoofing 2.2.2.2 ip name-server 8.8.8.8 ip domain lookup ip dns server R3# ip name-server 2.2.2.2 ip domain lookup</pre> <p>DNS: No name-servers are accessible DNS: Spoofing reply to query (id#7)</p>	<h3>How to find a default IOS name, in case no local IOS is found, Routers tries to TFTP for an image:</h3>	<h3>How to find a default IOS name, in case no local IOS is found, Routers tries to TFTP for an image:</h3> <pre>Rack1R1#reload <Sent BREAK Sequence> rommon 1 > confreg Configuration Summary enabled are: load rom after netboot fails console baud: 9600 boot: image specified by the boot system commands or default to: cisco2-C2600</pre>	<h3>Enhanced Object Tracking</h3> <p>Using percentage</p> <pre>track 4 list threshold percentage object 1 object 2 object 3 threshold percentage up 51 down 10 exit At least 51% of the 3 objects need to be up in order to make Track object 4 stay up. R1#show ip route track-table ip route 0.0.0.0 0.0.0.0 Null0 track 4 state is [up]</pre>	<h3>Enhanced Object tracking</h3> <p>Using percentage</p>																																				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

<p>R2# int e0/0 ip address 20.0.0.2 255.255.255.0 vrrp ip 20.0.0.1 vrrp prio 125 mac-address 0000.1111.1111</p> <p>R3# int e0/0 ip address 20.0.0.3 255.255.255.0 vrrp ip 20.0.0.1 vrrp prio 120 mac-address 0000.2222.2222</p> <p>1.1.1.1/32</p> <p>ip route 0/0 via 20.0.0.1</p> <p>How will the traceroute and show arp look like?</p>	<p>R5#traceroute 1.1.1.1 numeric Type escape sequence to abort. Tracing the route to 1.1.1.1 VRF info: (vrf in name/id, vrf out name/id) 1 20.0.0.2 1 msec 0 msec 0 msec 2 12.1.1.1 12 msec * 13 msec</p> <p>No ARP for 20.0.0.2 but traceroute shows 20.0.0.2 in the path!</p> <p>HOST#show arp</p> <table border="1"> <thead> <tr> <th>Protocol</th> <th>Address</th> <th>Age (min)</th> <th>Hardware Addr</th> <th>Type</th> <th>Interface</th> </tr> </thead> <tbody> <tr> <td>Internet</td> <td>20.0.0.1</td> <td>0</td> <td>0000.5e00.0101</td> <td>ARPA</td> <td>Ethernet0/0</td> </tr> <tr> <td>Internet</td> <td>20.0.0.3</td> <td>-</td> <td>0000.3333.3333</td> <td>ARPA</td> <td>Ethernet0/0</td> </tr> </tbody> </table>	Protocol	Address	Age (min)	Hardware Addr	Type	Interface	Internet	20.0.0.1	0	0000.5e00.0101	ARPA	Ethernet0/0	Internet	20.0.0.3	-	0000.3333.3333	ARPA	Ethernet0/0																																																		
Protocol	Address	Age (min)	Hardware Addr	Type	Interface																																																																
Internet	20.0.0.1	0	0000.5e00.0101	ARPA	Ethernet0/0																																																																
Internet	20.0.0.3	-	0000.3333.3333	ARPA	Ethernet0/0																																																																
<p>R2# int e0/0 ip address 20.0.0.2 255.255.255.0 vrrp ip 20.0.0.1 vrrp prio 125 mac-address 0000.1111.1111</p> <p>R3# int e0/0 ip address 20.0.0.3 255.255.255.0 vrrp ip 20.0.0.1 vrrp prio 120 mac-address 0000.2222.2222</p> <p>1.1.1.1/32</p> <p>ip route 0/0 via 20.0.0.1</p> <p>How will the following output look like? show vrrp brief</p>	<p>R2#show vrrp brief</p> <table border="1"> <thead> <tr> <th>Interface</th> <th>Grp</th> <th>Pri</th> <th>Time</th> <th>Own</th> <th>Pre</th> <th>State</th> <th>Master addr</th> <th>Group addr</th> </tr> </thead> <tbody> <tr> <td>Eto/0</td> <td>1</td> <td>125</td> <td>3531</td> <td>Y</td> <td>Master</td> <td>20.0.0.2</td> <td>20.0.0.1</td> <td></td> </tr> </tbody> </table> <p>R3#show vrrp brief</p> <table border="1"> <thead> <tr> <th>Interface</th> <th>Grp</th> <th>Pri</th> <th>Time</th> <th>Own</th> <th>Pre</th> <th>State</th> <th>Master addr</th> <th>Group addr</th> </tr> </thead> <tbody> <tr> <td>Eto/0</td> <td>1</td> <td>120</td> <td>3570</td> <td>Y</td> <td>Backup</td> <td>20.0.0.2</td> <td>20.0.0.1</td> <td></td> </tr> </tbody> </table>	Interface	Grp	Pri	Time	Own	Pre	State	Master addr	Group addr	Eto/0	1	125	3531	Y	Master	20.0.0.2	20.0.0.1		Interface	Grp	Pri	Time	Own	Pre	State	Master addr	Group addr	Eto/0	1	120	3570	Y	Backup	20.0.0.2	20.0.0.1																																	
Interface	Grp	Pri	Time	Own	Pre	State	Master addr	Group addr																																																													
Eto/0	1	125	3531	Y	Master	20.0.0.2	20.0.0.1																																																														
Interface	Grp	Pri	Time	Own	Pre	State	Master addr	Group addr																																																													
Eto/0	1	120	3570	Y	Backup	20.0.0.2	20.0.0.1																																																														
<p>R2# interface Ethernet0/0 mac-address 0000.2222.2222 ip address 10.1.1.2 255.255.255.0 glbp 1 ip 10.1.1.100 glbp 1 priority 150</p> <p>R3# interface Ethernet0/0 mac-address 0000.3333.3333 ip address 10.1.1.3 255.255.255.0 glbp 1 ip 10.1.1.100 glbp 1 priority 102</p> <p>1.1.1.1/32</p> <p>ip route 0/0 via 10.1.1.100</p> <p>How will "show glbp brief" look on R2/R3?</p>	<p>R2#show glbp brief</p> <table border="1"> <thead> <tr> <th>Interface</th> <th>Grp</th> <th>Fwd</th> <th>Pri</th> <th>State</th> <th>Address</th> <th>Active router</th> <th>Standby router</th> </tr> </thead> <tbody> <tr> <td>Eto/0</td> <td>1</td> <td>-</td> <td>150</td> <td>Active</td> <td>10.1.1.100</td> <td>local</td> <td>10.1.1.3</td> </tr> <tr> <td>Eto/0</td> <td>1</td> <td>1</td> <td>-</td> <td>Active</td> <td>0007.b400.0101</td> <td>local</td> <td>-</td> </tr> <tr> <td>Eto/0</td> <td>1</td> <td>2</td> <td>-</td> <td>Listen</td> <td>0007.b400.0102</td> <td>10.1.1.3</td> <td>-</td> </tr> </tbody> </table> <p>R3#show glbp brief</p> <table border="1"> <thead> <tr> <th>Interface</th> <th>Grp</th> <th>Fwd</th> <th>Pri</th> <th>State</th> <th>Address</th> <th>Active router</th> <th>Standby router</th> </tr> </thead> <tbody> <tr> <td>Eto/0</td> <td>1</td> <td>-</td> <td>102</td> <td>Standby</td> <td>10.1.1.100</td> <td>10.1.1.2</td> <td>local</td> </tr> <tr> <td>Eto/0</td> <td>1</td> <td>1</td> <td>-</td> <td>Listen</td> <td>0007.b400.0101</td> <td>10.1.1.2</td> <td>-</td> </tr> <tr> <td>Eto/0</td> <td>1</td> <td>2</td> <td>-</td> <td>Active</td> <td>0007.b400.0102</td> <td>local</td> <td>-</td> </tr> </tbody> </table>	Interface	Grp	Fwd	Pri	State	Address	Active router	Standby router	Eto/0	1	-	150	Active	10.1.1.100	local	10.1.1.3	Eto/0	1	1	-	Active	0007.b400.0101	local	-	Eto/0	1	2	-	Listen	0007.b400.0102	10.1.1.3	-	Interface	Grp	Fwd	Pri	State	Address	Active router	Standby router	Eto/0	1	-	102	Standby	10.1.1.100	10.1.1.2	local	Eto/0	1	1	-	Listen	0007.b400.0101	10.1.1.2	-	Eto/0	1	2	-	Active	0007.b400.0102	local	-				
Interface	Grp	Fwd	Pri	State	Address	Active router	Standby router																																																														
Eto/0	1	-	150	Active	10.1.1.100	local	10.1.1.3																																																														
Eto/0	1	1	-	Active	0007.b400.0101	local	-																																																														
Eto/0	1	2	-	Listen	0007.b400.0102	10.1.1.3	-																																																														
Interface	Grp	Fwd	Pri	State	Address	Active router	Standby router																																																														
Eto/0	1	-	102	Standby	10.1.1.100	10.1.1.2	local																																																														
Eto/0	1	1	-	Listen	0007.b400.0101	10.1.1.2	-																																																														
Eto/0	1	2	-	Active	0007.b400.0102	local	-																																																														
<p>Which GLBP load-balancing method MUST be used in combination with SNAT ?</p>	<p>GLBP and SNAT requires</p> <p>int fa0/x glbp <1> load-balancing host-dependent</p> <p>In this case a Host always receives the same ARP for its Gateway which is consistent with the NAT entry.</p>																																																																				
<p>Explain GLBP</p> <p>glbp <nr> weighting 110 lower 95 upper 105</p>	<p>track 22 interface fa0/x line-protocol</p> <p>int fa0/2 glbp 1 weighting track 22 decrement 20 glbp 1 weighting 110 lower 95 upper 105 glbp 1 prio 119 glbp 1 preempt</p> <p>- sets the weight to 110 - if the weight value due to a decrement goes lower than 95 the router will NO LONGER be an active forwarder. - if the value goes over 105, it will start forwarding again.</p> <p>110 - 20 = 90, stops forwarding</p>																																																																				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

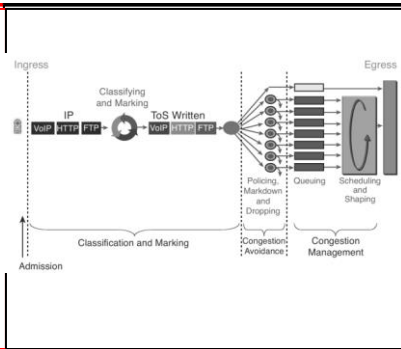
Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

QoS Explained in one picture:



Tool	Max Number of Queue	Classification possibilities
Priority Queuing (PQ)		
Custom Queuing (CQ)		
Weighted Fair Queueing (WFQ)		
Class-Based WFQ (CBWFQ)		
Low Latency Queueing LLQ		
Modified Deficit Round-Robin		

QoS action Sub-commands Inside a policy map:

Set	Mark several fields inside headers
Bandwidth	Reserve BW for a class CBWFQ
Priority	Reserve BW for LLW in CBWFQ
Shape	traffic in class to defined BW, bursts
Police	traffic in class to defined BW, bursts
Compress	performs TCP/RTP header compression/class

QoS Hold-Queue and Tx-Ring

- input software queue length of 10 packets
 - output software queue length of 30 packets
 - output hardware queue size to 15 packets

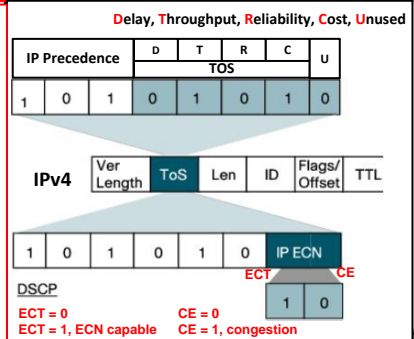
interface FastEthernet0/0
hold-queue 10 in
hold-queue 30 out
tx-ring-limit 15

R1#show interfaces fastEthernet 0/0 | include queue
 Input queue: 0/10/0/0 (size/max/drops/flushes); Total output drops: 0
 Output queue: 0/30 (size/max)

show controllers fastEthernet 0/0 | include tx
 tx_limited=0(15)

Field/Value	Binary	Name
Precedence 0	000	Routine
Precedence 1	001	Priority
Precedence 2	010	Immediate
Precedence 3	011	Flash
Precedence 4	100	Flash Override
Precedence 5	101	Critic/ECN
Precedence 6	110	Internetwork control
Precedence 7	111	Network Control

QoS ToS field explained:
 Using IP precedence Or DSCP



Input / Output Queueing explained:

- input queueing there is just one queue per interface, size of 75 packets by default (FIFO)

Two output queues by default.

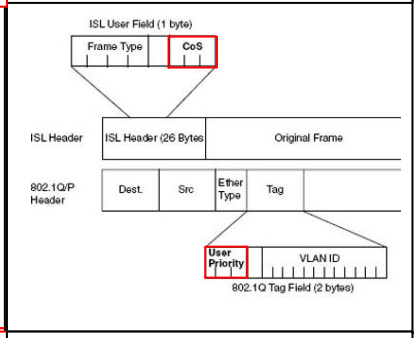
- first output queue is the "software queue" queuing strategy can be defined. (FIFO, WFQ, CBWFQ)

- Second output queue is the transmit ring tx-ring is always FIFO.

software output queue only starts to fill up when the tx-ring is full

Name of DSCP Class Selector	Binary	Precedence Value
Default	000000	0
CS1	001000	1
CS 2	010000	2
CS 3	011000	3
CS 4	100000	4
CS 5	101000	5
CS 6	110000	6
CS 7	111000	7

QoS L2 markable fields:
 ISL 802.1Q/P CoS User Priority



What are the different sources of Delay:

- Serialization Delay (fixed) = $\frac{\text{bits sent}}{\text{Link speed}}$ = msec per packet

- Propagation delay (fixed) = $\frac{\text{Length of Link (meters)}}{2.1 \times 10^8 \text{ meters/second}}$

- Queueing delay (variable) = time spent inside queues inside the device

- Forwarding / processing delay (variable) = ignored

- Shaping delay (variable)

- Network delay (variable)

- Codec delay (fixed)

- Compression delay (variable)

Range of DSCP	Binary	Precedence Value
0-7	000xxx	0
8-15	001xxx	1
16-23	010xxx	2
24-31	011xxx	3
32-39	100xxx	4
40-47	101xxx	5
48-55	110xxx	6
56-63	111xxx	7

QoS Pre-classify Type of VPN

QoS pre-qualify configured on	Type of VPN
Interface tunnel	GRE and IP/IP
Interface virtual-template	L2F and L2TP
Crypto map	IPSEC

QoS What fields can be marked?

- IP precedence [0 - 7]
 - DSCP [0 - 63]
 - 802.1P CoS [0 - 7]
 - ISL priority
 - ATM CLP
 - Frame-Relay DE
 - MPLS Experimental [0 - 7]
 - QoS Group [1 - 99]

Drop probability within class		
Low drop Name/Dec/Bin	Medium Drop Name/Dec/Bin	High Drop Name/Dec/Bin
Class 1 AF11 / / AF12 / / AF13 / /	Class 2 AF21 / 18 / 010010 AF22 / 20 / 010100 AF23 / 22 / 010110	Class 3 AF31 / 26 / 011010 AF32 / 28 / 011100 AF33 / 30 / 011110
Class 4 AF41 / / AF42 / / AF43 / /	Class 4 AF41 / 34 / 100010 AF42 / 36 / 100100 AF43 / 38 / 100110	

Calculation from AF numbers to decimal:
 AFxy
 8x + 2y = decimal value (AF41 = 34)

QoS pre-classify Crypto map

```
crypto map XXXX 10 ipsec-isakmp
...
qos pre-classify

show crypto map
...
QoS pre-classification

interface tunnel 0
qos pre-classify

interface virtual-template1
qos pre-classify
```

QoS What tools can be used to classify packets?

- IP ACLs
 - Any markable fields PREC, DSCP, ...
 - Input interface
 - MAC address (SRC or DST)
 - NBAR-enabled fields

QoS class-maps Match-all / match-any While matching DSCP / or IP Prec

class-map match-all WILL-NOT-WORK
 match dscp 0
 match dscp 1

class-map match-all WILL-WORK
 match dscp 1

class-map match-any WILL-WORK
 match dscp 0
 match dscp 1

class-map match-any WILL-WORK
 match dscp 0 1

QoS Comparing QoS Queueing tools:

Feature	Definition	QoS Characteristic Affected
Classification	determination of queue	None
Drop policy	Tail drop, modified tail drop, WRED, ECN	Loss
Scheduling Inside queue	Inside queue mostly FIFO scheduling	Bandwidth, delay jitter, loss
Scheduling Between queues	how queueing chooses queue for next packet	None
Max Queue Length	Max number of pkts in a single queue	Loss, delay

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Donate

Thanks for appreciating my efforts

Colin

<p>QoS</p> <p>Tx_limited=0(x) explained</p> <pre>Show controllers serial 0/0 i tx_ Tx_limited=0(16) Int ser0/0 Priority-group 1 Show controllers serial 0/0 i tx_ Tx_limited=1(2) Int ser0/0 tx-ring-limit 1 show controllers serial 0/0 i tx_ No queuing, but tx limited to tx_limited=0(1)</pre> <p><i>Hardware queue holds 16 packets 0 = queue size not limited due to a queuing tool being enabled on interface</i></p> <p><i>TX ring is 2 (1) -length is automatically limited as result of queuing configured.</i></p>	<p>show queue serial 0/0</p> <p>Explained:</p>	<pre>R1#show queue serial 0/0 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0 Queueing strategy: weighted fair Output queue: 0/1000/64/0 (size/max total/threshold/drops) Conversations 0/1/256 (active/max active/max total) Reserved Conversations 0/0 (allocated/max allocated) Available Bandwidth 1158 kilobits/sec Current WFQ Size / max total = hold queue, global limit of buffers / threshold ??? / count of tail drops If active = 0 no conversations / max currently active / absolute max Reserved conversation == RSVP flow reservation</pre>	<table border="1"> <tr> <td></td> <td>WFQ</td> <td>CBWFQ</td> <td>LLQ</td> </tr> <tr> <td>Requires complex classification</td> <td>No</td> <td>Yes</td> <td>Yes</td> </tr> <tr> <td>Uses MQC</td> <td>No</td> <td>Yes</td> <td>Yes</td> </tr> <tr> <td>Prefers low volume, high IPP flows</td> <td>Yes</td> <td>not flow based</td> <td></td> </tr> <tr> <td>Experiences problems with Large numbers of flows</td> <td>Yes</td> <td>No*</td> <td>No</td> </tr> <tr> <td>Can reserve bandwidth per queue</td> <td>No</td> <td>Yes</td> <td>Yes</td> </tr> <tr> <td>Provide low delay, low jitter queing</td> <td>No</td> <td>No</td> <td>Yes</td> </tr> </table> <p>*WFQ inside CBWFQ class-default, can have problems</p>		WFQ	CBWFQ	LLQ	Requires complex classification	No	Yes	Yes	Uses MQC	No	Yes	Yes	Prefers low volume, high IPP flows	Yes	not flow based		Experiences problems with Large numbers of flows	Yes	No*	No	Can reserve bandwidth per queue	No	Yes	Yes	Provide low delay, low jitter queing	No	No	Yes	<table border="1"> <tr> <td></td> <td>WFQ</td> <td>CBWFQ</td> <td>LLQ</td> </tr> <tr> <td>Requires complex classification</td> <td>No</td> <td>Yes</td> <td>Yes</td> </tr> <tr> <td>Uses MQC</td> <td>No</td> <td>Yes</td> <td>Yes</td> </tr> <tr> <td>Prefers low volume, high IPP flows</td> <td>Yes</td> <td>not flow based</td> <td></td> </tr> <tr> <td>Experiences problems with Large numbers of flows</td> <td>Yes</td> <td>No*</td> <td>No</td> </tr> <tr> <td>Can reserve bandwidth per queue</td> <td>No</td> <td>Yes</td> <td>Yes</td> </tr> <tr> <td>Provide low delay, low jitter queing</td> <td>No</td> <td>No</td> <td>Yes</td> </tr> </table> <p>*WFQ inside CBWFQ class-default, can have problems</p>		WFQ	CBWFQ	LLQ	Requires complex classification	No	Yes	Yes	Uses MQC	No	Yes	Yes	Prefers low volume, high IPP flows	Yes	not flow based		Experiences problems with Large numbers of flows	Yes	No*	No	Can reserve bandwidth per queue	No	Yes	Yes	Provide low delay, low jitter queing	No	No	Yes
	WFQ	CBWFQ	LLQ																																																									
Requires complex classification	No	Yes	Yes																																																									
Uses MQC	No	Yes	Yes																																																									
Prefers low volume, high IPP flows	Yes	not flow based																																																										
Experiences problems with Large numbers of flows	Yes	No*	No																																																									
Can reserve bandwidth per queue	No	Yes	Yes																																																									
Provide low delay, low jitter queing	No	No	Yes																																																									
	WFQ	CBWFQ	LLQ																																																									
Requires complex classification	No	Yes	Yes																																																									
Uses MQC	No	Yes	Yes																																																									
Prefers low volume, high IPP flows	Yes	not flow based																																																										
Experiences problems with Large numbers of flows	Yes	No*	No																																																									
Can reserve bandwidth per queue	No	Yes	Yes																																																									
Provide low delay, low jitter queing	No	No	Yes																																																									
<p>QoS</p> <p>show interface ser0/0 i (Que que)</p> <p>With Fair-queue</p> <p>With FIFO (no fair-queue)</p> <pre>sh int ser0/0 i (Que que) Default: fair-que R1#sh int ser0/0 i (Que que) Broadcast queue 0/64, broadcasts sent/dropped 261/0, interface broadcasts 267 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0 Queueing strategy: weighted fair Output queue: 0/1000/64/0 (size/max total/threshold/drops) FIFO (no fair-queue) R1#sh int ser0/0 i (Que que) Broadcast queue 0/64, broadcasts sent/dropped 257/0, interface broadcasts 263 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0 Queueing strategy: fifo Output queue: 0/40 (size/max)</pre>	<p>QoS</p> <p>WFQ:</p> <p>show queueing fair</p>	<pre>R1#show queueing fair Current fair queue configuration: Interface Discard Dynamic Reserved Link Priority threshold queues queues queues queues Serial0/0 64 256 0 8 1 Serial0/1 64 256 0 8 1 Alternatives: show queueing custom show queueing fair show queueing priority show queueing random-detect</pre>	<p>QoS</p> <p>When to use traffic shaping:</p>	<p>Shaping is always a egress function</p> <ul style="list-style-type: none"> - shaping slows down its sending rate, so that packets are not discarded - use shaping if the other device opposite is policing - can be used in case of speed mis-matches - Can help solve Egress blocking: 																																																								
<p>QoS</p> <p>(PQ) Scheduling Logic</p> <p>Priority Queue</p>	<p>QoS</p> <p>CBWFQ</p> <p>Details:</p>	<ul style="list-style-type: none"> - For Queues with less drop sensitive traffic WRED is a good option - CBWFQ supports 64 queues, queue lengths depends on router model - Class-Default configured automatically - "queue-limit 30" sets maximum queue size to 30 / queue <table border="1"> <tr> <td>Drop policy:</td> <td>Tail drop or WRED</td> </tr> <tr> <td>Number queues:</td> <td>64</td> </tr> <tr> <td>Scheduling inside Single queue</td> <td>FIFO on 63 queues, FIFO or WFQ on class-default queue</td> </tr> <tr> <td>Max. Queue length</td> <td>depends on router</td> </tr> </table>	Drop policy:	Tail drop or WRED	Number queues:	64	Scheduling inside Single queue	FIFO on 63 queues, FIFO or WFQ on class-default queue	Max. Queue length	depends on router	<p>QoS</p> <p>Shaping terminology</p>	<table border="1"> <tr> <td>Tc</td> <td>Time interval in msec, over which the Bc can be sent. Tc = Bc/CIR</td> </tr> <tr> <td>Bc</td> <td>Committed burst size in bits. Amount of traffic that can be sent over one Tc</td> </tr> <tr> <td>CIR</td> <td>Committed information rate, in bits/sec</td> </tr> <tr> <td>Shaped Rate</td> <td>in bits per second</td> </tr> <tr> <td>Be</td> <td>Excess burst size in Bits. Number of bits beyond Bc can be sent after a period of inactivity</td> </tr> </table>	Tc	Time interval in msec, over which the Bc can be sent. Tc = Bc/CIR	Bc	Committed burst size in bits. Amount of traffic that can be sent over one Tc	CIR	Committed information rate, in bits/sec	Shaped Rate	in bits per second	Be	Excess burst size in Bits. Number of bits beyond Bc can be sent after a period of inactivity																																						
Drop policy:	Tail drop or WRED																																																											
Number queues:	64																																																											
Scheduling inside Single queue	FIFO on 63 queues, FIFO or WFQ on class-default queue																																																											
Max. Queue length	depends on router																																																											
Tc	Time interval in msec, over which the Bc can be sent. Tc = Bc/CIR																																																											
Bc	Committed burst size in bits. Amount of traffic that can be sent over one Tc																																																											
CIR	Committed information rate, in bits/sec																																																											
Shaped Rate	in bits per second																																																											
Be	Excess burst size in Bits. Number of bits beyond Bc can be sent after a period of inactivity																																																											
<p>QoS</p> <p>Custom Queuing Logic</p> <p>(CQ)</p> <pre>- 16 queues available - guarantees minimum Bw for each queue - Bandwidth % for Queue X = (byte count for Queue X)/Sum of Bytes for all queues</pre> <p>CQ does not provide great service for delay/jitter sensitive traffic!</p>	<p>QoS</p> <p>WRED configuration:</p>	<p>WRED config</p> <p>policy-map X class class-default random-detect dscp-based</p> <pre>R1(config)#map-cj#random-detect ? discard-class parameters for each discard-class value discard-class-based Enable discard-class-based WRED as drop policy dscp parameters for each dscp value dscp-based Enable dscp-based WRED as drop policy ecn explicit congestion notification exponential-weighting-constant weight for mean queue depth calculation prec-based Enable precedence-based WRED as drop policy precedence parameters for each precedence value</pre>	<p>QoS</p> <p>Shaping formulas:</p>	<p>Bursts only in the first Tc!</p> <p>128 kbps</p> <p>Bc = Tc * CIR</p> <p>Tc = Bc/CIR</p> <p>Tc = Bc/Shaped rate</p> <p>1 sec = 1000 msec / 8 Tc = 125 msec per Tc</p> <p>Bc = Tc * Shaped rate</p> <p>Access Rate 128 kbps Shaping down to 64 kbps Default Tc is 125 msec (1 sec = 1000 msec / 8 = 125 msec)</p> <p>Bc = Tc * Shaped rate Bc = 0.125 sec * 64000bit/s = 8000 bits</p> <p><i>AR 128kbps, Shaped to 64k In each Tc, 8000 bits are sent (Bc)</i></p> <p><i>8000/128000 = 62.5ms</i></p>																																																								
<p>QoS</p> <p>Weighted Fair Queuing (WFQ)</p> <p>Logic explained</p> <p><i>Poor for Voice</i></p> <ul style="list-style-type: none"> - classifies packets based on flows - Flow consists of all packets have same SRC/DST IP addr and port numbers - weighted based on precedence - Favors low-volume, higher-precedence flows - each flow uses a different queue up to 4096 queues per interface <p>If WFQ empties a flow's queue, it removes the queue.</p> <p>In WFQ number of queues changes rapidly. show queue</p> <p>Precedence 7 traffic gets 8 times more bandwidth than 0: $(7 + 1) / (0 + 1) = 8$</p> <p>Precedence 3 traffic gets 4 times more bandwidth than 0: $(3 + 1) / (0 + 1) = 4$</p>	<p>QoS</p> <p>CBWFQ default values:</p>	<p>QoS</p> <p>- class-default receives per default 25% of bandwidth</p> <p>- max-reserved-bandwidth of 75%, meaning a policy-map can not define more than 75% of bandwidth on that interface.</p>	<p>QoS</p> <p>Fancy Shaping:</p>																																																									
<p>QoS</p> <p>Weighted Fair Queue WFQ</p> <p>Sequence Numbers explained:</p> <table border="1"> <thead> <tr> <th>Precedence</th> <th>ToS Value</th> <th>Weight</th> </tr> </thead> <tbody> <tr><td>0</td><td>0 (0x00)</td><td>32768</td></tr> <tr><td>1</td><td>32 (0x20)</td><td>16384</td></tr> <tr><td>2</td><td>64 (0x40)</td><td>10920</td></tr> <tr><td>3</td><td>96 (0x60)</td><td>8192</td></tr> <tr><td>4</td><td>128 (0x80)</td><td>6552</td></tr> <tr><td>5</td><td>160 (0xA0)</td><td>5456</td></tr> <tr><td>6</td><td>192 (0xCx)</td><td>4680</td></tr> <tr><td>7</td><td>224 (0xE0)</td><td>4096</td></tr> </tbody> </table> <p>Calculating WFQ Sequence Number (Finish Time FT): Weight = 32384 / (IP Precedence + 1) Previous_SN + (Weight * new_packet_length) = SN</p>	Precedence	ToS Value	Weight	0	0 (0x00)	32768	1	32 (0x20)	16384	2	64 (0x40)	10920	3	96 (0x60)	8192	4	128 (0x80)	6552	5	160 (0xA0)	5456	6	192 (0xCx)	4680	7	224 (0xE0)	4096	<p>QoS</p> <p>CBWFQ with LLQ logic</p>	<p>LLQ config in bandwidth: priority X in kbps</p>	<p>QoS</p> <p>Policing explained:</p>	<p>Policing, each token = 1 byte, therefore if police rate is 128000bps = 16000 bytes per second.</p> <p>Every second policing replenish 16000 tokens back into the bucket.</p> <p>Therefore, in 0.1 second, policing replenish 1600 tokens into the buckets.</p> <p>If the number of bytes of the packets is less or equal then the number of tokens, packet conforms.</p> <p>If the number of bytes of the packet is greater then the number of tokens, the packet exceeds the contract. Performs configured actions, without removing tokens!</p>																													
Precedence	ToS Value	Weight																																																										
0	0 (0x00)	32768																																																										
1	32 (0x20)	16384																																																										
2	64 (0x40)	10920																																																										
3	96 (0x60)	8192																																																										
4	128 (0x80)	6552																																																										
5	160 (0xA0)	5456																																																										
6	192 (0xCx)	4680																																																										
7	224 (0xE0)	4096																																																										

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin

QoS Dual Token Bucket (single rate)

Bc tokens placed in token bucket
1 byte = 1 token
If 128Kbps policed, Every second 16000 bytes are refilled
After 0.1 sec 1600 bytes/tokens refilled into Bc bucket

Spilled over from Bc to Be
conform
exceed
violate

If bytes of incoming packet <= Bc = conforms
If bytes of incoming packet <= Be = exceeds
If bytes of packet greater than Bc+Be = violates

Overview of Single-rate / Dual-rate Two/three color policing

Single-rate	two-color policing	single token bucket	conform/exceed
Single-rate	three-color policing	two token buckets	conform/exceed/violate
Dual-rate	three-color policing	two token buckets	PIR, CIR rates

Which TCP Flags can be used for congestion avoidance:

TCP Flags:

```

000 Reserved: not set
0 Nonce: not set
0 Congestion Window Reduced (CWR): not set
1 ECN-Echo: set
0 Urgent: not set
1 Acknowledgement: Set
0 Push: not set
0 Reset: not set
0 Syn: not set
0 Fin: not set
    
```

R1-config#ip tcp ecn
show tcp tcb 123456A | i ECN
Connection is ECN Enabled
debug ip tcp ecn

WRED: random-detect-ecn

QoS Dual Token Bucket (dual rate)

- two sending rates, but only guarantee the smaller one Tc, Bc, as with Single Rate
- Peak Information Rate (PIR): maximum sending rate. Bursts that exceed CIR but remain under PIR are allowed. May be marked for more aggressive discarding.
- Be: maximum size of the packet burst, accepted to sustain in the PIR rate

No modifications Change markings Discard, violates PIR

Type of Policing Configuration	signs in the police command	defaults	Type of Policing Configuration	signs in the police command	defaults
Single Rate Single bucket Two color			Single Rate Single bucket Two color	No violate action No Be, No PIR configured	Bc = CIR/32 Be = 0
Single Rate Dual bucket Three color			Single Rate Dual bucket Three color	No PIR configured violate-action and or Be configured	Bc = CIR/32 Be = Bc
Dual Rate Dual bucket Three color			Dual Rate Dual bucket Three color	PIR configured	Bc = CIR/32 Be = PIR/32

QoS TCP ECN / CWR Flag Explained:

TCP Flags:

```

000 Reserved: not set
0 Nonce: not set
0 Congestion Window Reduced (CWR): not set
1 ECN-Echo: set
...
    
```

- initial TCP SYN handshake includes the addition of ECN-echo capability and Congestion Window Reduced (CWR) capability flags to negotiate capabilities.
- When the TCP sender receives a packet with the ECN-echo flag set in the TCP header, the sender will adjust its congestion window as if it had undergone fast recovery from a single lost packet.
- Next sent packet will set the TCP CWR flag, to indicate to the receiver that it has reacted to the congestion

QoS Shaping example Token Bucket / Bc / Be / Tc / tokens per Tc

Pointed out using show policy-map interface X:

```

R1#sh policy-map interface ser0/0
  interface Serial0/0
    bandwidth 64
    fair-queue
    service-policy output SHAPE
  policy-map SHAPE
    class class-default
      shape average 64000
  class-map: class-default (match-any)
    8 packets, 1636 bytes
    30 second offered rate 0 bps, drop rate 0 bps
  Match: any
  Traffic Shaping
  Target/Average Byte Sustain Excess Interval Increment
  Rate Limit bits/int bits/int (ms) (bytes)
  64000/64000 2000 8000 8000 125 1000
  Adapt Queue Packets Bytes Packets Bytes Shaping
  Active Depth - 0 4 1584 0 0 no
    
```

TCP windowing:

- Receiver window or Advertised window Grants sender the right to send x bytes, before requiring an acknowledgement.
- Congestion window CWND Field never sent, is calculated by the TCP sender. Varies in size much more quickly, designed to react to congestion. If a TCP sender does not receive a ACK in time, the CWND is set to a single packet.
- SSTRESH is set to 50% of the CWND value before the last segment. CWND grows at exponential rate during slow start.
- If a packet is lost, the TCP sender decides to use the receiver window or CWND, which ever is smaller at the time.

QoS WRED configuration: Class-based Interface based:

Per interface:
interface Ser0/0
random-detect dscp-based
random-detect dscp af21 50 60
random-detect dscp af31 20 30


Per Class:
policy-map X
class class-default
random-detect dscp-based <value> <min-thres> <max-thres> <mark-prob-denominator>

Random-detect exponential-weighting-constant X
Default is 9, the smaller the number the more quickly WRED will react to changes in the Q

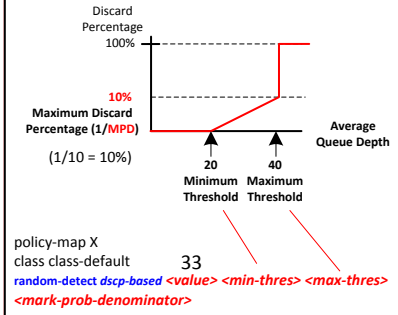
show queueing int X
show policy-map X
show queue

Shaping and latency-sensitive traffic:

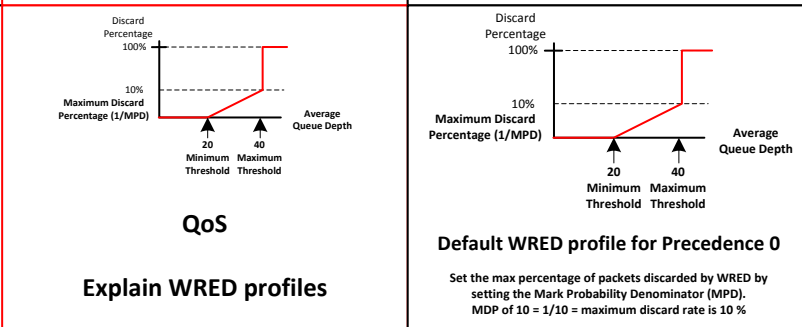
If you are sending latency-sensitive traffic, you should set Bc to drive the calculation of Tc down to 10 msec!



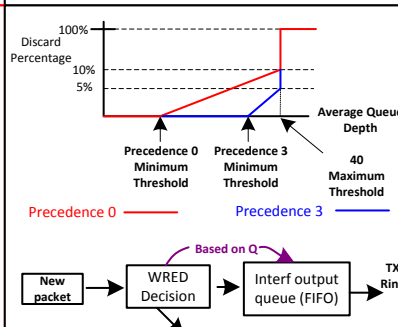
QoS WRED configuration: Class based with graph:



QoS Dual Token Bucket (dual rate)



Explain WRED profiles for Precedence 0 and Precedence 3



QoS Configuring a switch with only mls qos:


```

SW1#conf t
SW1#mls qos
    
```

Term	Meaning	Term	Meaning
Actual queue depth	number of packets in a queue	Actual queue depth	number of packets in a queue
Average queue depth	calculated measurement	Average queue depth	calculated measurement
Minimum Threshold	No pkts discarded if blow minimum	Minimum Threshold	No pkts discarded if blow minimum
Maximum Threshold	Pkts discarded if average above threshold	Maximum Threshold	Pkts discarded if average above threshold
Mark Probability Denominator	max % of packets discarded when average queue depth falls in min/max	Mark Probability Denominator	max % of packets discarded when average queue depth falls in min/max
Exponential Weighting constant	the larger the number, the slower the change in the av. Queue depth	Exponential Weighting constant	the larger the number, the slower the change in the av. Queue depth
No Drop	average below minimum threshold	No Drop	average below minimum threshold
Random Drop	between min and max threshold	Random Drop	between min and max threshold
Full Drop	Q depth exceeds max threshold, All Packets are dropped	Full Drop	Q depth exceeds max threshold, All Packets are dropped

RED Terminology

QoS The difference between TCP Slow start and Slow start and congestion avoidance



Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Donate

Thanks for appreciating my efforts



Colin

<p>QoS MLP LFI Queuing</p>		<p>QoS show mls qos map cos-dscp</p> <p>Modification and Output:</p>	<pre>SW2#show mls qos map cos-dscp Cos-dscp map: cos: 0 1 2 3 4 5 6 7 ----- dscp: 0 8 16 24 32 40 48 56 SW2(config)#mls qos map cos-dscp 0 8 16 26 32 46 48 56 SW2#show mls qos map cos-dscp Cos-dscp map: cos: 0 1 2 3 4 5 6 7 ----- dscp: 0 8 16 26 32 46 48 56</pre>	<table border="1"> <thead> <tr> <th>Codec</th> <th>speech samples per packet</th> <th>Voice Payload</th> <th>Packets/ Total Bandw. Second Per Call</th> </tr> </thead> <tbody> <tr> <td>G.711</td> <td></td> <td></td> <td></td> </tr> <tr> <td>G.711</td> <td>20 ms</td> <td>160 bytes</td> <td>50pps 80 kbps</td> </tr> <tr> <td>G.711</td> <td>30 ms</td> <td>240 bytes</td> <td>33pps 74 kbps</td> </tr> <tr> <td>G.729</td> <td>20 ms</td> <td>20 bytes</td> <td>50pps 24 kbps</td> </tr> <tr> <td>G.729</td> <td>30 ms</td> <td>30 bytes</td> <td>33pps 19 kbps</td> </tr> </tbody> </table>	Codec	speech samples per packet	Voice Payload	Packets/ Total Bandw. Second Per Call	G.711				G.711	20 ms	160 bytes	50pps 80 kbps	G.711	30 ms	240 bytes	33pps 74 kbps	G.729	20 ms	20 bytes	50pps 24 kbps	G.729	30 ms	30 bytes	33pps 19 kbps	
Codec	speech samples per packet	Voice Payload	Packets/ Total Bandw. Second Per Call																										
G.711																													
G.711	20 ms	160 bytes	50pps 80 kbps																										
G.711	30 ms	240 bytes	33pps 74 kbps																										
G.729	20 ms	20 bytes	50pps 24 kbps																										
G.729	30 ms	30 bytes	33pps 19 kbps																										
<p>QoS MLP LFI Calculation fragment size / maximum delay</p>	<p>You want a maximum serialization delay of 10 ms: Max-delay in 0.x seconds * bandwidth in bits = frag size Bandwidth = configured value with the bandwidth command. ppp multilink fragment delay X Example: 56Kbps with 10 ms delay makes fragment sizes of 70 bytes: 56000 bit * 0.1 sec = 560 bits/10 msec or 70 bytes/10msec</p>	<p>QoS show mls qos interface X</p> <p>Output:</p>	<pre>SW2#show mls qos interface fa0/1 FastEthernet0/1 QoS is disabled. When QoS is enabled, following settings will be applied trust state: not trusted trust mode: not trusted trust enabled flag: ena COS override: dis default COS: 0 DSCP Mutation Map: Default DSCP Mutation Map Trust device: none qos mode: port-based</pre>	<p>QoS Weighted Fair Queuing (WFQ)</p>	<pre>fair-queue 16 128 8 Queue length: 16 congestive discard threshold 128 number of conversations (flows) 8 reserved conversations for RSVP Tx ring as small as possible to trigger WFQ more quickly to kick in. tx-ring-limit 1 interface Serial0/1/0 clock rate 128000 bandwidth 128 tx-ring-limit 1 fair-queue 16 128 8 hold-queue 256 Hold-queue (Q length default 75, up to 4096)</pre> <p><i>A conversation is a unidirectional sequence of packets matching the same SPC/DST Ports</i> <i>Reserved Queues used for control-plane and L2 keepalive traffic</i></p>																								
<p>QoS PPP MLP LFI / multilink interleave Configuration:</p>	<pre>Interface [dialer0, virtual-template] ! Enables multilink ppp multilink ! Enables interleaving of unfragmented frames ppp multilink interleave ! Defines fragment size based on bandwidth/time formula ppp multilink fragment delay ! Disables MLP fragmentation ppp multilink fragment disable ! Configure on each interface of MLP bundle ppp multilink group <GRP-number></pre>	<p>QoS mls qos trust [x, y]</p> <p>switchport priority extended cos X</p> <p>switchport priority extended trust</p> <p>Explained:</p>	<pre>mls qos trust CoS [pass-through] Pass-through: Prevents switch from overwriting the original DSCP values sourced from the CoS-to-DSCP map. mls qos trust [device cisco-phone, dscp] mls qos trust device cisco-phone used with switchport priority extended cos <value> Overwrites original CoS value received from ethernet port of the phone. mls qos trust device cisco-phone used with switchport priority extend trust Trusts the markings sent on the phone's ethernet port of the attached PC</pre>	<p>QoS Normalize the packet flows</p>	<pre>interface Serial0/1/0 clock rate 128000 bandwidth 128 ip mtu 156 Normalize packet flows, so that each IP packet takes no more than 10ms to be sent on 128 Kbps link Bitrate of link * time-limit in msec = bit/sec 128000 bits / sec * 0.010 = 1248 bits/sec 1248 / 8 = 156 bytes = MTU size</pre>																								
<p>QoS PPP MLP LFI / multilink interleave Configuration:</p>	<pre>Interface Multilink 9 bandwidth 128 ip address 1.2.3.4 255.255.255.0 fair-queue ppp multilink multilink-group 9 int serial0/1 bandwidth 128 no ip address encapsulation ppp ppp multilink multilink-group 9 interface multilink 9 ppp multilink fragment-delay 10 ppp multilink interleave</pre>	<p>QoS mls qos cos</p> <p>mls qos cos override</p> <p>explained</p>	<pre>mls qos cos <value> Attaches specified CoS value to all untagged frames received. mls qos cos override Overwrites the original CoS value received interface FastEthernet0/1 mls qos cos 3 mls qos cos override Alternative: Trust incoming cos, set untagged to 4 int x mls qos cos 4 mls qos trust cos SW2#show mls qos int fa0/1 FastEthernet0/1 trust state: cos override trust mode: cos override trust enabled flag: ena COS override: ena default COS: 3 DSCP Mutation Map: Default DSCP Mutation Map Trust device: none qos mode: port-based</pre> <p><i>Will set untagged frames to 3, and will override tagged frames to 3!</i></p>	<p>QoS MTU's explained:</p> <p>Mtu</p> <p>ip mtu</p> <p>mpls mtu</p>	<pre>mtu maximum packet length the interface can support, oversized packets may not be interpreted correctly on the other end ip mtu fragments an IP packet if the packet arriving exceeds the value configured mpls mtu fragments the MPLS labeled packet if the labeled packet arriving exceeds the value configured</pre>																								
<p>QoS MLP Multilink Difference of Fragment-delay and interleave</p>		<p>QoS Congestion Management</p> <p>On a 2950 with 1 ingress and 4 egress transmit queues</p>		<p>QoS show queueing fair</p> <p>Output:</p> <p>WFQ</p>	<pre>R5#show queueing fair Current fair queue configuration: Interface Discard Dynamic Reserved Link Priority threshold queues queues queues queues Q Serial0/0/0 64 256 0 8 1 Serial0/1/0 16 128 8 8 1</pre>																								
<p>QoS Configuring class based header compression: Header compression on interface:</p>	<pre>class-map match-all COMPRESS-IP match ip address 101 class-map match-all COMPRESS-TCP match ip address 199 policy-map COMPRESSING class COMPRESS-IP compression header ip class COMPRESS-TCP compression header ip tcp interface Serial Encapsulation ip tcp header-compression [passive iphc-format ietf-format] ip rtp header-compression [passive, ...]</pre>	<p>QoS LAN policing</p> <p>Explained:</p>	<pre>police rate-bps burst-byte exceed-action [drop, dscp <value>] Rate-bps: average receive rate the policer will accept Burst-byte: acceptable values 4096, 8192, ... in kbps Policy-map POLICE-ME Class FTP Set ip dscp 18 police 500000 8192 exceed-action dscp 0</pre>	<p>Legacy RTP Reserved Queue</p>	<p>100% of the link bandwidth is reserved for RTP traffic in the UDP port range 16384 – 32767.</p> <p>making WFQ aware of voice bearer traffic</p> <pre>interface Serial 0/1/0 no hold-queue out fair-queue max-reserved-bandwidth 100 ip rtp reserve 16384 16383 128 ip rtp reserve <Starting UDP Port> <Port-Range> <Bandwidth> (UDP 16384 to 32767 for 128kbps)</pre>																								

Help me create more flashcards:


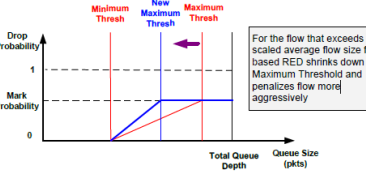
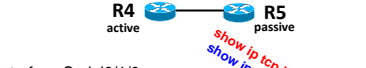
Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Thanks for appreciating my efforts

Colin


<p>Legacy RTP Reserved Queue</p> <p>Verification:</p>	<pre>interface Serial0/1/0 bandwidth 128 max-reserved-bandwidth 100 ip rtp reserve 16384 16383 128 R4#show queueing fair Current fair queue configuration: Interface Discard Dynamic Reserved Link Priority threshold queues queues queues q Serial0/0/0 64 256 0 8 1 Serial0/1/0 64 32 1 8 1</pre>	<p>Legacy Custom Queueing with Prioritization</p> <p>(not recommended, as to no way to limit the priority queues)</p>	<pre>access-list 100 permit tcp 155.1.146.0 0.0.0.255 eq www any access-list 101 permit icmp any any queue-list 1 protocol ip 0 udp 520 queue-list 1 protocol ip 1 lt 65 queue-list 1 protocol ip 2 list 100 queue-list 1 protocol ip 3 list 101 queue-list 1 default 3 queue-list 1 queue 3 limit 10 queue-list 1 queue 1 byte-count 320 queue-list 1 queue 2 byte-count 640 queue-list 1 queue 3 byte-count 104 interface Serial 0/1/0 custom-queue-list 1 queue-list 1 lowest-custom 2 lowest-custom option tells the round robin scheduler with witch Queue to start</pre> <p>Priorize RTP traffic over RIP packets</p>	<p>Legacy Flow-Based Random Early Detection</p> <p>config</p>	<p>flow-based RED:</p> <ul style="list-style-type: none"> - maximum number of flows to 16 - average flow depth scale factor to 2 - FIFO queue depth to 10 packets <p>average queue size per flow, $Avg=Q/N$. (queue size Q divided by number of currently active flows)</p> <pre>interface Serial0/1/0 random-detect random-detect flow random-detect flow count 16 random-detect flow average-depth-factor 2 hold-queue 10 out show queueing random-detect Mean queue depth: 0 Max flow count: 16 Average depth factor: 2 Flows (active/max active/max): 0/0/16</pre>
<p>To verify IP RTP</p> <p>Configure IP SLA sourcing G.729 like packet stream.</p>	<pre>R1# ip sla monitor 1 type jitter dest-ipaddr 155.1.45.5 dest-port 16384 codec g729a timeout 1000 frequency 1 ip sla monitor schedule 1 life forever start-time now R2# rtr responder ip sla responder</pre> 	<p>Legacy Priority Queueing</p> <p>Configuration:</p> <p>not very helpful for VoIP deployments</p>	<pre>- RIP packets first, second RTP VoIP - If no RIP or VoIP traffic, than Web traffic first - Last ICMP traffic - Queue-Size High 5, Medium 40, Normal 60, Low 80 pkts access-list 102 permit udp any any range 16384 32767 access-list 103 permit icmp any any priority-list 1 protocol ip high udp rip priority-list 1 protocol http normal priority-list 1 protocol ip medium list 102 priority-list 1 protocol ip low list 103 priority-list 1 queue-limit 5 40 60 80 interface Serial 0/1/0 priority-group 1</pre> <p>sends the most important traffic first, before it moves to the next important flow</p>	<p>Legacy Flow-Based Random Early Detection</p> <p>diagram</p>	<p>NewMaxThreshold = MinThreshold + (MaxThreshold-MinThreshold)/2.</p> 
<p>WFQ</p> <p>show queueing int serial X</p> <p>Output:</p>	<pre>R4#show queueing interface serial 0/1/0 Interface Serial0/1/0 queueing strategy: fair Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 18748 Queueing strategy: weighted fair Output queue: 4/1000/64/18748 (size/max total/threshold/drops) Conversations 1/3/32 (active/max active/max total) Reserved Conversations 0/1 (allocated/max allocated) Available Bandwidth 128 kilobits/sec (depth/weight/total drops/no-buffer drops/interleaves) 4/16192/0/0/0 Conversation 20, linktype: ip, length: 118 source: 150.1.6.6, destination: 155.1.108.10, id: 0x4A59, ttl: 254, TOS: 32 prot: 6, source port 19, destination port 40966</pre>	<p>Legacy Priority Queueing</p> <p>show interface X</p> <p>Output:</p>	<pre>Rack1R4#show interfaces s0/1/0 Serial0/1/0 is up, line protocol is up Queueing strategy: priority-list 1 Output queue (queue priority: size/max/drops): high: 0/5/0, medium: 0/40/0, normal: 0/60/18754, low: 0/80/0 access-list 102 permit udp any any range 16384 32767 access-list 103 permit icmp any any priority-list 1 protocol ip high udp rip priority-list 1 protocol http normal priority-list 1 protocol ip medium list 102 priority-list 1 protocol ip low list 103 priority-list 1 queue-limit 5 40 60 80 interface Serial 0/1/0 priority-group 1</pre> <p>show queueing serial 0/1/0 = high show queue serial 0/1/0 = med show queue serial 0/1/0 2 = norm show queue serial 0/1/0 3 = low</p>	<p>Selective Packet Discard</p>	<ul style="list-style-type: none"> - enable selective packet discard in aggressive mode - increase input queue - set memory headroom for IGP packets to 150 Buffers - set headroom for BGP packets should be set to 120 packets - Start dropping low-priority packets randomly when the input queue is 50% full <pre>spd extended-headroom 150 IGP's initial land here spd headroom 120 BGP initial lands here ip spd mode aggressive ip spd queue max-threshold 150 ip spd queue min-threshold 75 interface FastEthernet 0/0 hold-queue 150 in</pre> <p>SPD Extended Headroom Q is emptied before the spd headroom Queue</p>
<p>show ip sla monitor statistics X</p> <p>output</p>	<pre>R6#show ip sla monitor statistics 1 Round trip time (RTT) Index 1 Latest RTT: 12 ms Latest operation start time: *03:02:53.955 UTC Tue Apr 8 2014 Latest operation return code: OK RTT Values Number Of RTT: 1000 RTT Min/Avg/Max: 9/12/85 ms Source to Destination Jitter Min/Avg/Max: 0/2/61 ms Destination to Source Jitter Min/Avg/Max: 0/1/12 ms Packet Loss Values Loss Source to Destination: 0 Loss Destination to Source: 0 Out Of Sequence: 0 Tail Drop: 0 Packet Late Arrival: 0 Voice Score Values Calculated Planning Impairment Factor (ICPIF): 11 MOS score: 4.06 Number of successes: 13</pre>	<p>Legacy Random Early Detection</p>	<pre>- avoid tail drop, by randomly dropping packets before output queue overflows (hold-queue size set to 10 packets) (NOT weighted, no different profiles) interface Serial0/1/0 random-detect random-detect exponential-weighting-constant 4 random-detect precedence 6 11 12 hold-queue 10 out R4#show interfaces serial 0/1/0 Queueing strategy: random early detection (RED)</pre> <p>Max threshold pkts Min threshold pkts Precedence</p>	<p>Selective Packet Discard</p> <p>Normal / Aggressive mode</p> <p>show ip spd</p> <p>Output:</p>	<p>Normal mode</p> <p>treats malformed packets as it would treat regular IP packets. -> hold queue, random drop.</p> <p>Aggressive mode:</p> <p>malformed packets dropped as soon as hold queue grows above min threshold. -> unconditional drop</p> <pre>R1#show interface fastEthernet 0/0 Input queue: 0/150/0/0 (size/max/drops/flushes); Total output drops: 0 R1#show ip spd Current mode: normal. Queue min/max thresholds: 75/150, Headroom: 120, Extended Headroom: 150 IP normal queue: 0, priority queue: 0. SPD special drop mode: aggressively drop bad packets</pre>
<p>Legacy RTP Prioritization</p>	<pre>interface Serial0/1/0 bandwidth 128 no ip address fair-queue max-reserved-bandwidth 100 ip rtp priority 16384 16383 128 IP RTP Priority feature differs from the IP RTP Reserve in that the priority queue has a WFQ weight of zero, meaning that the WFQ always services it first. ip rtp priority <Starting UDP Port> <Port Range> <Bandwidth> ip rtp priority is policing <bandwidth></pre>	<p>QoS Legacy Custom Queueing</p> <p>using show interface X</p> <p>Output:</p>	<pre>R4#show interfaces serial 0/1/0 Last input 00:00:05, output 00:00:03, output hang never Last clearing of "show interface" counters 02:15:06 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 18763 Queueing strategy: custom-list 1 Output queues: (queue #: size/max/drops) 0: 0/20/0 1: 0/20/0 2: 0/20/18754 3: 0/10/0 4: 0/20/0 5: 0/20/0 6: 0/20/0 7: 0/20/0 8: 0/20/0 9: 0/20/0 10: 0/20/0 11: 0/20/0 12: 0/20/0 13: 0/20/0 14: 0/20/0 15: 0/20/0 16: 0/20/0 5 minute input rate 0 bits/sec, 0 packets/sec 5 minute output rate 0 bits/sec, 0 packets/sec Queue Number</pre>	<p>Payload Compression on Serial Links</p>	<p>-Predictor" compression algorithm uses a minimum of CPU cycles on the router, but needs more memory</p> <p>Predictor only on PPP Stacker only on HDLC</p> <pre>interface Serial0/1 encapsulation ppp compress predictor interface Serial0/1 encapsulation ppp compress predictor Serial0/1 is up, line protocol is up Encapsulation PPP, LCP Open Open: IPCP, CCP, CDPCP, loopback not set interface Serial1/3 encapsulation hdlc compress stac interface Serial1/3 encapsulation hdlc compress stac show compress</pre> <p>Check compression by pinging over the link, and watch the RTT timer</p>
<p>Legacy Custom Queueing</p>	<pre>- VoIP traffic should be guaranteed 30% - File transfers from V1146 60% - remaining 10% for ICMP, should not exceed 10 pkts in any queue in any time. - RIP packets in system prio Q access-list 100 permit tcp 155.1.146.0 0.0.0.255 eq www any access-list 101 permit icmp any any queue-list 1 protocol ip 0 udp 520 queue-list 1 protocol ip 1 lt 65 queue-list 1 protocol ip 2 list 100 queue-list 1 protocol ip 3 list 101 queue-list 1 default 3 queue-list 1 queue 3 limit 10 queue-list 1 queue 1 byte-count 320 queue-list 1 queue 2 byte-count 640 queue-list 1 queue 3 byte-count 104 interface Serial 0/1/0 custom-queue-list 1</pre> <p>packets with a size less than 60 bytes are matched in queue 1</p> <p>Ratio Calculation is a bit tricky, normalization etc..</p>	<p>QoS Legacy Custom Queueing</p> <p>Show queueing custom</p> <p>Output:</p>	<pre>- VoIP traffic should be guaranteed 30% - File transfers from V1146 60% - remaining 10% for ICMP, should not exceed 10 pkts in any queue in any time. - RIP packets in system prio Q R4#show queueing custom Current custom queue configuration: List Queue Args 1 3 default 1 0 protocol ip udp port rip 1 1 protocol ip lt 65 1 2 protocol ip list 100 1 3 protocol ip list 101 1 1 byte-count 320 1 2 byte-count 640 1 3 byte-count 104 limit 10</pre>	<p>Generic TCP/UDP Header Compression</p>	<ul style="list-style-type: none"> - maximum of 16 concurrent RTP and TCP sessions - R5 only optimizes if it detects optimized traffic from R4  <pre>interface Serial0/1/0 ip tcp header-compression ip tcp compression-connections 32 ip rtp header-compression ip rtp compression-connections 32 interface Serial0/1/0 ip tcp header-compression passive ip tcp compression-connections 32 ip rtp header-compression passive ip rtp compression-connections 32</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts


Colin


<p>Generic TCP/UDP Header Compression</p> <p>show ip tcp header-compression output</p>	<p>R5#show ip tcp header-compression TCP/IP header compression statistics: Interface Serial0/1/0 (compression on, V1, passive) Rcvd: 94 total, 93 compressed, 0 errors, 0 status msgs 0 dropped, 3 buffer copies, 0 buffer failures Sent: 96 total, 93 compressed, 0 status msgs, 0 not predicted 3243 bytes saved, 677 bytes sent 5.79 efficiency improvement factor Connect: 32 rx slots, 32 tx slots, 1 misses, 0 collisions, 0 negative cache hits, 32 free contexts 98% hit ratio, five minute miss rate 0 misses/sec, 0 max</p> <p>show ip tcp header-compression serial 0/1/0 detail IP source: 150.1.4.4, IP destination: 155.1.45.5 TCP source: 23, TCP destination: 37587 Last packet received is V1compressed-tcp and last sequence received is 0</p>	<p>QoS</p> <p>Legacy CAR for Admission Control</p> <p>(Designed for packet remarking / policing)</p>	<p>- Traffic up to 256Kbps should be marked with an IP precedence of 1 - Exceeding traffic should be marked with an IP precedence of 0 - use average traffic burst size of 4000 bytes</p> <p>access-list 111 permit ip host 155.1.146.1 any</p> <p>interface FastEthernet0/1 rate-limit input access-group 111 256000 4000 4000 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0</p> <p>(continue to next rate-limit statement..)</p>	<p>QoS</p> <p>MQC Bandwidth Reservations and CBWFQ</p>	<p>- set total size of MQC buffer to 512 - http from vlan 146 with IP Prec 0 should be guaranteed 32Kbps - Limit FIFO Q for http to 16 pkts, IP Prec 0 trail to 24 pkts - all other traffic run WFQ, dynamic flows start dropping if they reach 32 packets</p> <p>policy-map SERIAL_LINK class VOICE interface Serial 0/1/0 class HTTP bandwidth 128 bandwidth 32 max-reserved-bandwidth 100 queue-limit 16 no fair-queue class SCAVENGER hold-queue 512 out bandwidth 32 queue-limit 24 class class-default fair-queue queue-limit 32 class class-default bandwidth 96</p> <p>To use FIFO instead WFQ:</p>																																								
<p>MLP Link Fragmentation and Interleaving</p>	<p>- maximum serialization delay is 10ms - bandwidth of the link is 128Kbps. - Encapsulation ppp - PPP multilink with interleaving</p> <p>interface Virtual-Template1 bandwidth 128 ip address 155.1.45.4 255.255.255.0 fair-queue</p> <p>ppp multilink fragment delay 10 ppp multilink interleave</p> <p>multilink virtual-template 1</p> <p>interface Serial0/1/0 bandwidth 128 no ip address encapsulation ppp load-interval 30 ppp multilink</p> <p>fragment size should be based upon the physical rate of the link (clock rate) not the logical bandwidth</p>	<p>QoS</p> <p>Legacy CAR for Admission Control</p> <p>Config / show commands</p>	<p>access-list 111 permit ip host 155.1.146.1 any</p> <p>interface FastEthernet0/1 rate-limit input access-group 111 256000 4000 4000 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0</p> <p>show interfaces fastEthernet 0/1 rate-limit FastEthernet0/1 Input matches: access-group 111 params: 256000 bps, 4000 limit, 4000 extended limit conformed 3841 packets, 3885702 bytes; action: set-prec-transmit 1 exceeded 24 packets, 24336 bytes; action: set-prec-transmit 0 last packet: 20ms ago, current burst: 3928 bytes last cleared 00:02:21 ago, conformed 219000 bps, exceeded 1000 bps</p>	<p>CBWFQ weight / conversation numbers:</p>	<table border="1"> <thead> <tr> <th>Conversation Numbers</th> <th>CBWFQ Weight</th> <th>Description</th> </tr> </thead> <tbody> <tr> <td>Below 2^N</td> <td>32384(PP+1)</td> <td>Used for automatic classification of dynamic WFQ Queues. Configurable via fair-queue command under "class-default".</td> </tr> <tr> <td>2^N, 2^N+7</td> <td>1024</td> <td>Link Queues. There are 8 conversations used to queue routing updates (marked as PAK_PRIORITY internally) and Layer 2 keepalives.</td> </tr> <tr> <td>2^N+8</td> <td>0</td> <td>Priority queue which maps directly to legacy WFQ IP RTP Priority. Typically used for VoIP "beaver" traffic. CBWFQ polices this flow using a configurable token bucket procedure to ensure it does not starve other queues.</td> </tr> <tr> <td>Above 2^N+8</td> <td>Const*diffBW/ClassBW OR RSVP flow weight</td> <td>These conversations are used for user-defined traffic classes. Each class has its own FIFO queue and configurable queue depth. In addition, RSVP flows use the same range of conversation numbers.</td> </tr> </tbody> </table>	Conversation Numbers	CBWFQ Weight	Description	Below 2^N	32384(PP+1)	Used for automatic classification of dynamic WFQ Queues. Configurable via fair-queue command under "class-default".	2^N, 2^N+7	1024	Link Queues. There are 8 conversations used to queue routing updates (marked as PAK_PRIORITY internally) and Layer 2 keepalives.	2^N+8	0	Priority queue which maps directly to legacy WFQ IP RTP Priority. Typically used for VoIP "beaver" traffic. CBWFQ polices this flow using a configurable token bucket procedure to ensure it does not starve other queues.	Above 2^N+8	Const*diffBW/ClassBW OR RSVP flow weight	These conversations are used for user-defined traffic classes. Each class has its own FIFO queue and configurable queue depth. In addition, RSVP flows use the same range of conversation numbers.																									
Conversation Numbers	CBWFQ Weight	Description																																											
Below 2^N	32384(PP+1)	Used for automatic classification of dynamic WFQ Queues. Configurable via fair-queue command under "class-default".																																											
2^N, 2^N+7	1024	Link Queues. There are 8 conversations used to queue routing updates (marked as PAK_PRIORITY internally) and Layer 2 keepalives.																																											
2^N+8	0	Priority queue which maps directly to legacy WFQ IP RTP Priority. Typically used for VoIP "beaver" traffic. CBWFQ polices this flow using a configurable token bucket procedure to ensure it does not starve other queues.																																											
Above 2^N+8	Const*diffBW/ClassBW OR RSVP flow weight	These conversations are used for user-defined traffic classes. Each class has its own FIFO queue and configurable queue depth. In addition, RSVP flows use the same range of conversation numbers.																																											
<p>QoS</p> <p>show ppp multilink</p> <p>Output:</p>	<p>Rack1R4#show ppp multilink</p> <p>Virtual-Access2, bundle name is R5 Endpoint discriminator is R5 Bundle up for 00:18:36, total bandwidth 128, load 190/255 Receive buffer limit 12192 bytes, frag timeout 1000 ms</p> <p>Interleaving enabled 0/0 fragments/bytes in reassembly list 0 lost fragments, 0 reordered 0/0 discarded fragments/bytes, 0 lost received 0x7EC3 received sequence, 0xFC9C sent sequence Member links: 1 (max not set, min not set) Se0/1/0, since 00:18:34, 160 weight, 152 frag size No inactive multilink interfaces</p>	<p>QoS</p> <p>Oversubscription with Legacy CAR and WFQ</p> <p>- 64Kbps that R4 receives from R1 and R6 - Traffic from R1 and R6 should be allowed up to 128Kbps each total - Traffic above 128Kbps should be dropped - Averaging time interval of 200ms.</p>	<p>access-list 111 permit ip host 155.1.146.1 any access-list 116 permit ip host 155.1.146.6 any</p> <p>interface FastEthernet0/1 rate-limit input access-group 111 64000 3200 3200 conform-action set-prec-transmit 1 exceed-action continue rate-limit input access-group 111 128000 3200 3200 conform-action set-prec-transmit 0 exceed-action drop</p> <p>rate-limit input access-group 116 64000 3200 3200 conform-action set-prec-transmit 1 exceed-action continue rate-limit input access-group 116 128000 3200 3200 conform-action set-prec-transmit 0 exceed-action drop</p> <p>interface Serial 0/1/0 fair-queue <i>if R1 and R6 transmit, share 128K at 64/64</i> clock rate 128000</p>	<p>CBWFQ weight / conversation numbers:</p> <p>Constant:</p>	<table border="1"> <thead> <tr> <th>The "N" constant</th> <th>Number of dynamic Flows</th> <th>CBWFQ constant "Const"</th> </tr> </thead> <tbody> <tr><td>4</td><td>16</td><td>64</td></tr> <tr><td>5</td><td>32</td><td>64</td></tr> <tr><td>6</td><td>64</td><td>57</td></tr> <tr><td>7</td><td>128</td><td>30</td></tr> <tr><td>8</td><td>256</td><td>16</td></tr> <tr><td>9</td><td>512</td><td>8</td></tr> <tr><td>10</td><td>1024</td><td>4</td></tr> <tr><td>11</td><td>2048</td><td>2</td></tr> <tr><td>12</td><td>4096</td><td>1</td></tr> </tbody> </table>	The "N" constant	Number of dynamic Flows	CBWFQ constant "Const"	4	16	64	5	32	64	6	64	57	7	128	30	8	256	16	9	512	8	10	1024	4	11	2048	2	12	4096	1										
The "N" constant	Number of dynamic Flows	CBWFQ constant "Const"																																											
4	16	64																																											
5	32	64																																											
6	64	57																																											
7	128	30																																											
8	256	16																																											
9	512	8																																											
10	1024	4																																											
11	2048	2																																											
12	4096	1																																											
<p>QoS</p> <p>PPP Multilink</p> <p>show ppp multilink</p> <p>show interface virtual-access x</p> <p>show interface serial x</p> <p>Output:</p>	<p>R4#show ppp multilink Interleaving enabled Se0/1, since 00:13:44, 160 weight, 152 frag size</p> <p>R4#show interface virtual-access 2 MLP Bundle vaccess, cloned from Virtual-Template1 Output queue: 6/1000/64/0/8445 (size/max total/threshold/drops/interleaves) Conversations 2/3/256 (active/max active/max total) Reserved Conversations 0/0 (allocated/max allocated)</p> <p>R4#show interface serial 0/1/0 Link is a member of Multilink bundle Virtual-Access2 Queueing strategy: weighted fair [suspended, using FIFO] FIFO output queue 0/10, 3 drops</p> <p>show queueing interface virtual-access 2</p>	<p>QoS</p> <p>Legacy CAR for Rate Limiting</p>	<p>- drop traffic in excess traffic of 256Kbps - committed burst of 384Kbps - excess burst of 768Kbps</p> <p>interface FastEthernet 0/0 rate-limit input 256000 48000 96000 conform-action transmit exceed-action drop</p> <p>show interfaces fastEthernet 0/0 rate-limit</p> <p>(The TCP receiver's window size for TCP flows tends to be around $Traffic_Rate * RTT$)</p>	<p>QoS</p> <p>MQC Bandwidth Percent</p>	<p>policy-map SERIAL_LINK class HTTP bandwidth percent 25 class SCAVENGER bandwidth percent 25</p> <p>configured bandwidth percent values in all classes of a policy-map cannot exceed the max-reserved-bandwidth</p> <p>It is not possible to mix bandwidth with bandwidth percent commands!</p>																																								
<p>QoS</p> <p>Legacy Generic Traffic Shaping</p> <p>config</p> <p>Scheduler using WFQ!</p>	<p>- limit the rate of packets going towards R4's 99.99.99.0/24 to 128Kbps - Shaping interval of 10msec, disable extended burst - Limit shapers queue size to 1024 packets</p> <p>access-list 199 permit ip any 99.99.99.0 0.0.0.255</p> <p>interface FastEthernet 0/0 traffic-shape group 199 128000 1280 0 1024</p> <p>(1280*8)/128000 = 0.08 sec</p> <p>traffic-shape rate <CIR> <BC> <Be> <QueueLimit> traffic-shape group <ACL> <CIR> <BC> <Be> <QueueLimit></p> <p>show traffic-shape</p>	<p>QoS</p> <p>Legacy CAR Access-Lists</p>	<p>- rate limit packets going out to MAC address X of MAC-ACL - rate-limite packets towards X having IP Prec set to 1,2,4 to 256Kbps</p> <p>access-list rate-limit 100 000d.2846.8f21</p> <p>interface FastEthernet 0/0 rate-limit output access-group rate-limit 100 128000 8000 8000 conform-action transmit exceed-action drop</p> <p>access-list rate-limit 0 mask 16 binary 1 means to check, binary 0 ignore 0x16 (00010110)</p> <p>interface FastEthernet 0/0.67 rate-limit output access-group rate-limit 0 128000 8000 8000 conformaction transmit exceed-action drop</p>	<p>QoS</p> <p>MQC LLQ and Remaining Bandwidth Reservations</p>	<p>policy-map SERIAL_LINK class VOICE priority 27 class HTTP bandwidth remaining percent 33 class SCAVENGER bandwidth remaining percent 33</p> <p>(if there are 32 (25) dynamic queues, then the LLQ conversation is number 40, weight value of 0, is serviced first)</p> <p>HDLC overhead of 7 bytes, 60 bytes of layer 3 VoIP Payload = 27 Kbps (67 bytes/packet * 50 packets/second * 8 bits/byte = 26800bps)</p>																																								
<p>QoS</p> <p>Legacy Generic Traffic Shaping</p> <p>config / show cmd</p>	<p>access-list 199 permit ip any 99.99.99.0 0.0.0.255</p> <p>interface FastEthernet 0/0 traffic-shape group 199 128000 1280 0 1024</p> <p>R6#show traffic-shape</p> <p>Interface Gi0/0.146 Access Target Byte Sustain Excess Interval Increment Adapt VC List Rate Limit bits/int bits/int (ms) (bytes) Active - 104 128000 160 1280 0 10 160 -</p> <p>R6#show traffic-shape statistics</p> <table border="1"> <thead> <tr> <th>I/F</th> <th>Acc. List</th> <th>Queue Depth</th> <th>Packets</th> <th>Bytes</th> <th>Delayed</th> <th>Bytes</th> <th>Shaping</th> </tr> </thead> <tbody> <tr> <td>Fa0/0</td> <td>104</td> <td>2</td> <td>15365</td> <td>21690</td> <td>14962</td> <td>2156</td> <td>yes</td> </tr> </tbody> </table>	I/F	Acc. List	Queue Depth	Packets	Bytes	Delayed	Bytes	Shaping	Fa0/0	104	2	15365	21690	14962	2156	yes	<p>QoS</p> <p>MQC Classification and Marking</p>	<p>- Http from vlan 146 marked with IP Prec 2 - VoIP range 16384-32767 packet size <60 byte marked EF - ICMP packets >1000 bytes should be dropped - All other incoming with Prec 0 to be marked with 1</p> <p>ip access-list extended HTTP permit tcp 155.1.146.0 0.0.0.255 eq www any</p> <p>ip access-list extended VOICE permit udp any any range 16384 32767</p> <p>class-map HTTP match access-group name HTTP</p> <p>class-map match-all LARGE_ICMP match protocol icmp match packet length min 1001</p> <p>class-map match-all VOICE match access-group name VOICE match packet length min 60 max 60</p> <p>policy-map SERIAL_LINK class VOICE set ip dscp ef class HTTP set ip precedence 2 class LARGE_ICMP Drop class SCAVENGER set ip precedence 1</p>	<p>QoS</p> <p>MQC WRED</p>	<p>start dropping packets when average queue size is between 4-16 drop 1 of 4 packets when queue size reaches the max</p> <p>policy-map SERIAL_LINK class HTTP random-detect random-detect precedence 2 4 16 4</p> <p>Precedence minimum-threshold in pkts max-threshold in pkts Mark probability denominator</p> <p>show policy-map int x</p> <table border="1"> <thead> <tr> <th>class</th> <th>Transmitted</th> <th>pkts/bytes</th> <th>pkts/bytes</th> <th>pkts/bytes</th> <th>thresh</th> <th>thresh</th> <th>prob</th> </tr> </thead> <tbody> <tr> <td>2</td> <td>713/4089</td> <td>16/9280</td> <td>0/0</td> <td>4</td> <td>16</td> <td>1/4</td> <td></td> </tr> <tr> <td>3</td> <td>0/0</td> <td>0/0</td> <td>0/0</td> <td>26</td> <td>40</td> <td>1/10</td> <td></td> </tr> </tbody> </table> <p>ability to use random drop per flow, not per whole queue</p>	class	Transmitted	pkts/bytes	pkts/bytes	pkts/bytes	thresh	thresh	prob	2	713/4089	16/9280	0/0	4	16	1/4		3	0/0	0/0	0/0	26	40	1/10	
I/F	Acc. List	Queue Depth	Packets	Bytes	Delayed	Bytes	Shaping																																						
Fa0/0	104	2	15365	21690	14962	2156	yes																																						
class	Transmitted	pkts/bytes	pkts/bytes	pkts/bytes	thresh	thresh	prob																																						
2	713/4089	16/9280	0/0	4	16	1/4																																							
3	0/0	0/0	0/0	26	40	1/10																																							

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!





Thanks for appreciating my efforts

Colin

<p>QoS</p> <h3>MQC Dynamic Flows and WRED</h3>	<p>- activate random drops for unclassified traffic's dynamic flows</p> <p>- set for IP Prec 1 thres min 1 max 40</p> <p>- probability of 25% packet discard</p> <pre> policy-map SERIAL_LINK class class-default fair-queue no queue-limit random-detect random-detect precedence 1 1 40 4 </pre> <p>two ways to enable WRED within class-default:</p> <ol style="list-style-type: none"> bandwidth reservation statement - turning the class's queue into a FIFO queue - and then enabling RED enable RED with WFQ 	<p>QoS</p> <h3>MQC Single-Rate Three-Color Policer</h3> <p>Config:</p>	<p>- meter incoming HTTP traffic:</p> <ul style="list-style-type: none"> if less than 128K set IP Prec 1 if exceeds 128K then set IP Prec of 0 (using inactivity) drop violating traffic <p>Ensure burst durations of 200ms and 300ms</p> <pre> ip access-list extended HTTP permit tcp any eq 80 any class-map HTTP match access-group name HTTP policy-map POLICE_VLAN146 class HTTP police 128000 3200 4800 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0 violate-action drop </pre> <p>interface FastEthernet 0/1 service-policy input POLICE_VLAN146</p> <p><i>Be (exceed) is used as extra credit for periods of long inactivity</i></p>	<p>QoS</p> <h3>MQC Percent-Based Policing</h3>	<p>specifying the policer rate as a percentage of the interface speed, while specifying the burst rate in msec</p> <p>- set speed to 10</p> <p>- Limit the traffic entering this link to 10% of this rate with a burst value of 125ms.</p> <pre> policy-map POLICE_VLAN146 class class-default police rate percent 10 burst 125 ms </pre> <p>interface FastEthernet 0/0 speed 10 service-policy input POLICE_VLAN146</p>
<p>QoS</p> <h3>MQC WRED with ECN</h3>	<p>- overall effect of TCP ECN is better performance, as compared to simple packet drops and slow start.</p> <p>- changing the exceed action from random drop to ECN marking</p> <pre> policy-map SERIAL_LINK class HTTP random-detect ecn random-detect precedence 2 4 16 4 </pre>	<p>QoS</p> <h3>MQC Single-Rate Three-Color Policer</h3> <p>Config / show output:</p>	<pre> class HTTP police 128000 3200 4800 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0 violate-action drop </pre> <p>Class-map: HTTP (match-all) show policy-map interface 1245 packets, 1786990 bytes 30 second offered rate 127000 bps, drop rate 0 bps Match: access-group name HTTP</p> <p>police: cir 128000 bps, bc 3200 bytes, be 4800 bytes conformed 1241 packets, 1780934 bytes; actions: set-prec-transmit 1 exceeded 4 packets, 6056 bytes; actions: set-prec-transmit 0 violated 0 packets, 0 bytes; actions: drop conformed 127000 bps, exceed 0 bps, violate 0 bps</p>	<p>QoS</p> <h3>MQC Header Compression</h3>	<pre> policy-map SERIAL_LINK_OUT class VOICE_BEARER priority 24 compression header ip rtp compression header ip tcp </pre> <p>show policy-map interface serial1/2</p> <pre> ... compress: header ip rtp UDP/RTP (compression on, IPHC, RTP) ... </pre> <p>In combination with priority queues: base your calculations on compressed packet sizes</p> <p><i>show ip header-compression</i></p>
<p>QoS</p> <h3>MQC Class-Based Generic Traffic Shaping (GTS)</h3> <p>config</p>	<p>- shape rate on vian 146 to 384 Kbps</p> <p>- shape rate on vian 67 to 512 Kbps</p> <p>- burst interval Be of 20ms</p> <pre> policy-map SHAPE_VLAN146 class class-default shape average 384000 7680 policy-map SHAPE_VLAN67 class class-default shape average 512000 10240 </pre> <p>interface FastEthernet 0/1.146 service-policy output SHAPE_VLAN146</p> <p>interface FastEthernet 0/1.67 service-policy output SHAPE_VLAN67</p> <p><i>Bc = CIR * Tc / 1000 384000 * 20msec / 1000 = 10240</i></p>	<p>QoS</p> <h3>MQC Hierarchical Policers</h3> <p>(nested service-policies)</p>	<pre> policy-map POLICE_VLAN146 class HTTP police 128000 3200 4800 conform-action transmit exceed-action set-prec-transmit 0 violate-action drop service-policy SUBRATE_POLICER </pre> <pre> policy-map SUBRATE_POLICER class FROM_R1 police 64000 3200 4800 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0 violate-action set-prec-transmit 0 class FROM_R6 police 64000 3200 4800 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0 violate-action set-prec-transmit 0 </pre> <p><i>128K</i> <i>64K</i> <i>64K</i></p>	<p>QoS</p> <h3>Catalyst QoS Port-Based Classification</h3> <p>(3560 models treats IPv6 as "IP" traffic, 3550 as "non-IP" traffic)</p>	<p>mls qos</p> <pre> interface FastEthernet 0/1 mls qos trust dscp </pre> <p>Trust CoS bits in 802.1p header for encapsulated vlans. For native vlan apply the CoS value of 1 (CDP, etc.)</p> <pre> interface FastEthernet 0/6 mls qos trust cos mls qos cos 1 </pre> <pre> interface FastEthernet 0/4 mls qos trust ip-precedence mls qos cos 2 </pre> <p>Force CoS 4 on all packets</p> <pre> interface FastEthernet 0/5 mls qos cos 4 mls qos cos override </pre> <p>Will map CoS 4 to DSCP 44 mls qos map cos-dscp 0 8 16 26 44 46 48 56</p> <p><i>If mls qos configured, but no trust, incoming marked packets will be reset to 0</i></p>
<p>QoS</p> <h3>MQC Class-Based Generic Traffic Shaping (GTS)</h3> <p>Config and output:</p>	<pre> policy-map SHAPE_VLAN146 class class-default shape average 384000 7680 </pre> <p>interface FastEthernet 0/1.146 service-policy output SHAPE_VLAN146</p> <p>show policy-map interface fastEthernet 0/0.146 Service-policy output: SHAPE_VLAN146 Class-map: class-default (match-any) 65844 packets, 28003483 bytes 5 minute offered rate 334000 bps, drop rate 0 bps Match: any Traffic Shaping Target/Average Byte Sustain Excess Interval Increment Rate Limit bits/int bits/int (ms) (bytes) 384000/384000 1920 7680 7680 20 960 Adapt Queue Packets Bytes Packets Bytes Shaping Active Depth Delayed Delayed Active - 6 65838 279907 56930 270858 yes</p> <p><i>Bc = CIR * Tc / 1000 384000 * 20msec / 1000 = 7680</i></p>	<p>QoS</p> <h3>MQC Two-Rate Three-Color Policer</h3> <p>- CIR 64Kbps PIR 128Kbps</p> <p>- CIR*400ms for Bc/Be PIR*200ms for Bc/Be</p> <p>- conform action set IP Prec 1, transmit</p> <p>- exceed action set IP Prec 0, transmit</p> <p>- violate action drop</p>	<pre> policy-map POLICE_VLAN146 class HTTP service-policy SUBRATE_POLICER </pre> <pre> policy-map SUBRATE_POLICER class FROM_R1 police cir 64000 bc 3200 pir 128000 be 6400 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0 violate-action drop </pre> <p>class FROM_R6 police cir 64000 bc 3200 pir 128000 be 6400 conform-action set-prec-transmit 1 exceed-action set-prec-transmit 0 violate-action drop</p> <p><i>Four configurable parameters: CIR, PIR, Bc, Be. Bc/Be have independent fill-rate</i></p> <p><i>PIR</i> <i>CIR</i></p>	<p>QoS</p> <h3>Catalyst QoS CoS port-based classification</h3>	<p>If the incoming packet is IP, the switch will if configured trust fist DSCP or IP Precedence then CoS.</p> <p>If no CoS value had been set, it will set the default CoS value if not specified differently via</p> <pre> mls qos cos 4 </pre> <p>(the CoS-to-DSCP map will map the value accordingly into the DSCP field)</p> <p>Mls qos trust CoS: if NO 802.1q header, default CoS value of interface is applied</p> <p>mls qos cos override will override any incoming packet.</p>
<p>QoS</p> <h3>MQC Class-Based GTS and CBWFQ</h3> <p>(configuring the shaper's queues)</p> <p>No 75% rule, you may want to define a separate class for control plane traffic</p>	<pre> - 32Kbps of priority for IP pkts size of 60 bytes with a Bc of 4000 byte - guarantee 256 Kbps of shaped http traffic - default class uses WFQ </pre> <pre> class-map VOICE match packet length min 60 max 60 ip access-list extended HTTP permit tcp any eq 80 any class-map HTTP match access-group name HTTP policy-map CBWFQ class VOICE priority 32 4000 class HTTP bandwidth 256 class class-default fair-queue </pre> <pre> policy-map SHAPE_VLAN146 class class-default shape average 384000 7680 service-policy CBWFQ </pre> <p>Shaping to 384K</p> <p>voice 32K http 256K default WFQ</p>	<p>QoS</p> <h3>MQC Class-Based GTS and CBWFQ</h3> <p>Configuring the shaper's queues</p>	<pre> show policy-map int X </pre> <pre> ClassMap: POLICE_VLAN146 ... ClassMap: HTTP ... ClassMap: CBWFQ ... ClassMap: CBWFQ ... </pre> <p>Shaping to 384K</p> <p>voice 32K http 256K default WFQ</p> <p><i>"Fancy QoS"</i></p>	<p>QoS</p> <h3>Catalyst</h3> <p>show mls qos interface fastEthernet 0/6 statistics</p> <p>output</p>	<pre> show mls qos interface fastEthernet 0/6 statistics </pre> <pre> FastEthernet0/6 dscp: incoming 0 - 4: 0 0 0 0 0 5 - 9: 0 0 0 0 0 10 - 14: 35 0 0 0 0 15 - 19: 0 0 0 0 0 20 - 24: 0 0 0 0 0 25 - 29: 0 0 0 0 0 30 - 34: 0 0 0 0 0 35 - 39: 0 0 0 0 0 40 - 44: 0 0 22 0 45 - 49: 0 0 0 0 0 50 - 54: 0 0 0 0 0 55 - 59: 0 0 0 0 0 60 - 64: 0 0 0 0 0 </pre>
<p>QoS</p> <h3>Basic calculation for a G.729 packet stream for Traffic-Shaping purposes:</h3> <p>Including voice payload, ethernet header, dot1q tag</p>	<p>60 bytes VoIP packet size (G.729)</p> <p>18 bytes overhead for an Ethernet header</p> <p>4 bytes of VLAN tag</p> <p>at 50 packets per second =</p> <p>(60+18)*50*8 = 31200bps or 312Kbps</p>	<p>QoS</p> <h3>MQC Peak Shaping</h3> <p>(ios bug for show cmd)</p>	<p>(ISP allows customers to send traffic at rates up to PIR, but only guarantees CIR rate in case of network congestion)</p> <p>- shape HTTP traffic to a peak rate of 128Kbps</p> <p>- Bc and Be bursts based on a 10ms interval.</p> <pre> class-map HTTP match access-group 180 </pre> <pre> policy-map POLICY class HTTP shape peak 64000 6400 6400 </pre> <p>interface FastEthernet 0/0 service-policy output POLICY</p> <p><i>CIR/10 msec = 6400</i> <i>PIR=2*CIR</i></p>	<p>QoS</p> <h3>Catalyst 3550 monitor mode:</h3>	<pre> interface FastEthernet 0/1 mls qos monitor packets interface FastEthernet 0/4 mls qos monitor dscp 0 16 26 46 48 clear mls qos interface fastEthernet 0/4 statistics </pre> <p>Set EF</p> <pre> SW#show mls qos int fa0/4 statistics FastEthernet0/4 Ingress dscp: incoming no_change classified policed dropped (in pkts) 0 : 2 2 0 0 0 ... 48: 402 32 0 0 0 Others: 0 0 0 0 0 Egress dscp: incoming no_change classified policed dropped (in pkts) 0 : 404 n/a n/a 0 0 ... 48: 0 n/a n/a 0 0 Others: 11 n/a n/a 0 0 </pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

Donate

VISA giro pay

Thanks for appreciating my efforts

Colin

Catalyst QoS Marking Pass-Through

Used to "tunnel" one type of QoS marking through your network, while using the other type.

Classify CoS values, but do not change the DSCP value in IP packets.

no mls qos rewrite ip dscp

Trust DSCP but do not change previous set CoS values:

Interface fa0/1
mls qos trust dscp pass-through cos
mls qos trust cos pass-through dscp

Have DSCP set to 46 But have CoS at 2

Catalyst QoS Port-Based Policing and Marking

Rate-Limit to 128Kbps allow 10msec of burst generated by a full 100Mbps interface rate.

```

policy-map POLICE
class class-default
police 128000 125000
    
```

100Mbps*10ms/8 = 125000

interface FastEthernet 0/1
service-policy input POLICE

Police ICMP and remark exceeding traffic to 8

```

policy-map POLICE_INBOUND
class ICMP
trust dscp
set cos 2
police 64000 16000 exceed-action policed-dscp-transmit
    
```

mls qos map policed-dscp 0 16 24 26 to 8

Per-Tunnel QoS for DMVPN QoS Profile

Configuring NHRP Group on a Spoke:

```

int tunnel X
ip nhrp group GROUP
    
```

Config NHRP mapping on HUB:

```

int tunnel X
ip nhrp map group GROUP-1 service-policy output SERVICE-POLICY-1
ip nhrp map group GROUP-2 service-policy output SERVICE-POLICY-2
    
```

show policy-map multipoint [tunnel <nr>]
show ip nhrp group-map GROUP
Show tunnel endpoints
show dmvpn detail | (NHRP group) [service-policy]

How to troubleshooting QoS Markings / Classification

SW1#ping Protocol [ip]:
Target IP address: 150.1.6.6
Repeat count [5]: 10
Datagram size [100]: 1000
Timeout in seconds [2]: 0
Extended commands [n]: y
Source address or interface: 150.1.7.7
Type of service [0]: 160
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 10, 1000-byte ICMP Echos to 150.1.6.6, timeout is 0 seconds:
Packet sent with a source address of 150.1.7.7
.....
Success rate is 10 percent (1/10), round-trip min/avg/max = 9/9/9 ms

access-list 100 permit icmp any any dscp 40
clear access-list counters

Show mls qos int fa0/x statistics

SW4#show mls qos interface fastEthernet 0/4 statistics

```

FastEthernet0/4
Ingress
dscp: incoming no_change classified policed dropped (in pkts)
0: 1142 1 1 0 0
8: 0 0 0 0 0
16: 0 0 1141 731 0
Others: 1 0 0 0 0
Egress
dscp: incoming no_change classified policed dropped (in pkts)
0: 1746 n/a n/a 0 0
8: 0 n/a n/a 0 0
16: 0 n/a n/a 0 0
Others: 84 n/a n/a 0 0
    
```

all incoming packets are DSCP 0, switch classifies them to CS2 (16), some are policed due to exceeding rates

How to setup the easiest possible QoS lab:

Have R1 send all packets with a CoS of 4

R2 should monitor all CoS classifications

R1#show policy-map interface FastEthernet0/0

```

Service-policy input: PMAP-MONITOR
Class-map: COS0 (match-all)
  6 packets, 708 bytes
  30 second offered rate 0 bps
  Match: cos 0
Class-map: DSCP-0-7 (match-any)
  0 packets, 0 bytes
  30 second offered rate 0 bps
  Match: dscp default [0] 1 2 3 4 5 6 7
  0 packets, 0 bytes
  30 second rate 0 bps
    
```

R2#show policy-map interface FastEthernet0/0

```

Service-policy input: PMAP-MONITOR
Class-map: COS0 (match-all)
  6 packets, 708 bytes
  30 second offered rate 0 bps
  Match: cos 0
Class-map: DSCP-0-7 (match-any)
  0 packets, 0 bytes
  30 second offered rate 0 bps
  Match: dscp default [0] 1 2 3 4 5 6 7
  0 packets, 0 bytes
  30 second rate 0 bps
    
```

On switches either "no mls qos" or trust cos!

ping R2 from R1, show policy-map interface, use clear counters, after tests

Fill in the blanks (Dezimal)

TOS	DSCP	IP-PREC
0	0	0
4	1	0
8	2	0
12	3	0
16	4	0
32	8	1
40	10	1
48	12	1
56	14	1
64	16	2
72	18	2
80	20	2
88	22	2

Catalyst 3560 Per-Port Per-VLAN Policing

```

interface FastEthernet 0/13
mls qos vlan-based
class-map TRUNKS
match input-interface FastEthernet 0/13
policy-map INTERFACE_POLICY
class TRUNKS
police 128000 16000 exceed policed-dscp-transmit
policy-map VLAN_POLICY
class IP_ANY
set dscp af21
service-policy INTERFACE_POLICY
interface Vlan 146
service-policy input VLAN_POLICY
    
```

Why does R2 not see any packets with DSCP 1 via show policy-map interface ?

R1#ping 10.0.0.2

```

int fa0/0
service-policy outbound SET-DSCP-1
policy-map SET-DSCP-1
class class-default
set dscp 1
    
```

R2#show policy-map interface FastEthernet0/0

```

Service-policy input: PMAP-MONITOR
Class-map: COS0 (match-all)
  6 packets, 708 bytes
  30 second offered rate 0 bps
  Match: cos 0
Class-map: DSCP-0-7 (match-any)
  0 packets, 0 bytes
  30 second offered rate 0 bps
  Match: dscp default [0] 1 2 3 4 5 6 7
  0 packets, 0 bytes
  30 second rate 0 bps
    
```


Entries stay 0!!

R2#show policy-map interface FastEthernet0/0

```

Service-policy input: PMAP-MONITOR
Class-map: COS0 (match-all)
  6 packets, 708 bytes
  30 second offered rate 0 bps
  Match: cos 0
Class-map: DSCP-0-7 (match-any)
  0 packets, 0 bytes
  30 second offered rate 0 bps
  Match: dscp default [0] 1 2 3 4 5 6 7
  0 packets, 0 bytes
  30 second rate 0 bps
    
```

First match here:



Catalyst QoS ACL Based Classification & Marking

set DSCP under non-IP traffic. The switch computes the respective CoS value using DSCP-to-CoS table.

mls qos rewrite ip dscp

mac access-list extended IPX
permit any any 0x8137 0x0

class-map match-all IPX
match access-group name IPX

policy-map CLASSIFY
class IPX
set dscp ef
class-default
trust CoS

Set DSCP, trust CoS on 3550

(Trust cos on class default)

Catalyst QoS Aggregate Policers

Aggregate ICMP and IPX to 128Kbps, drop exceeding

```

mls qos
mls qos aggregate-policer AGG128 128000 16000 exceed-action drop
policy-map POLICE_AGGREGATE
class ICMP
police aggregate AGG128
class IPX
police aggregate AGG128
class class-default
set dscp cs1
interface FastEthernet 0/1
service-policy input POLICE_AGGREGATE
    
```

Create a DSCP mutation map which changes inbound DSCP 1 to 60.

Verify using

show mls qos maps dscp-mutation:

d1 d2 Most specific outgoing DSCP value

Default DSCP Mutation Map:

```

d1: d2 0 1 2 3 4 5 6 7 8 9
0: 00 01 02 03 04 05 06 07 08 09
1: 10 11 12 13 14 15 16 17 18 19
2: 20 21 22 23 24 25 26 27 28 29
3: 30 31 32 33 34 35 36 37 38 39
4: 40 41 42 43 44 45 46 47 48 49
5: 50 51 52 53 54 55 56 57 58 59
6: 60 61 62 63
    
```

SW2#sh mls qos maps dscp-mutation 6

```

Dscp-dscp mutation map:
TST:
d1: d2 0 1 2 3 4 5 6 7 8 9
0: 00 00 02 03 04 05 06 07 08 09
1: 10 11 12 13 14 15 16 17 18 19
2: 20 21 22 23 24 25 26 27 28 29
3: 30 31 32 33 34 35 36 37 38 39
4: 40 41 42 43 44 45 46 47 48 49
5: 50 51 52 53 54 55 56 57 58 59
6: 60 61 62 63
    
```

show mls qos !!

mls qos rewrite ip dscp

interface Fa0/16

mls qos trust dscp

mls qos dscp-mutation TST

Catalyst 3550 Per-Port Per-VLAN Classification

Make sure that classification only affects Vlan 146 of a trunk port.

ip access-list extended ICMP
permit icmp any any

class-map ICMP
match access-group name ICMP

class-map VLAN_146_ICMP
match vlan 146

match class-map ICMP

policy-map PER_PORT_PER_VLAN
class VLAN_146_ICMP
set dscp CS3

Must be the first statement in class-map!

Second layer class-map ICMP is only allowed to have 1 match criteria

Catalyst QoS DSCP Mutation

QoS DSCP mutation maps come in handy on the border of two QoS domains that use different markings

```

mls qos
mls qos map dscp-mutation MUTATION 0 to 8
mls qos map dscp-mutation MUTATION 26 to 24
mls qos map dscp-mutation MUTATION 40 to 46
interface FastEthernet 0/4
mls qos trust dscp
mls qos dscp-mutation MUTATION
    
```

Process in order to configure DSCP mutation map:

- 1) mls qos
- 2) mls qos rewrite ip dscp
- 3) mls qos map dscp-mutation BLA 28 to 44
- 4) int fa0/x
mls qos dscp-mutation BLA
- 5) int fa0/x
mls qos trust dscp

Catalyst 3560 Per-VLAN Classification

```

mls qos
interface FastEthernet 0/1
mls qos vlan-based
switchport-mode trunk
interface Fa0/22
no mls qos vlan-based
policy-map PER_VLAN
class TCP
set dscp ef
interface Vlan 146
service-policy input PER_VLAN
    
```

Inherit QoS settings from SVI per port

Disable Vlan based QoS per port

show mls qos interface fa0/4 statistics

Advanced HTTP Classification with NBAR

HTTP transfers of files with extensions ".bin", ".text" and ".txt" are limited to 256Kbps:

```

class-map match-all EXTENSION
match protocol http url ".bin|\.t[ea]xt"
policy-map SHAPE
class EXTENSION
shape average 256000
interface FastEthernet 0/0.146
service-policy output SHAPE
    
```

map DSCP 5 to CoS 6

R1#show policy-map interface FastEthernet0/0

```

Service-policy input: PMAP-MONITOR
Class-map: COS0 (match-all)
  6 packets, 708 bytes
  30 second offered rate 0 bps
  Match: cos 0
Class-map: DSCP-0-7 (match-any)
  0 packets, 0 bytes
  30 second offered rate 0 bps
  Match: dscp default [0] 1 2 3 4 5 6 7
  0 packets, 0 bytes
  30 second rate 0 bps
    
```

SW2: DSCP remains 5, but L2 CoS modified from 0 to 6 !!

mls qos

int fa0/x
mls qos trust dscp

mls qos map dscp-cos <DSCP-value> to <CoS-Value>

mls qos map dscp-cos 5 to 6

Help me create more flashcards:

Simply press this button and send me your credit cards regards!


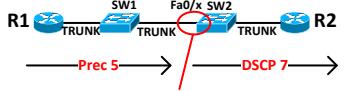
Ranging 5 bucks to unlimited!

Donate

VISA gipay

Thanks for appreciating my efforts

Colin


<p>show mls qos maps dscp-cos</p> <p>Explained:</p>	<pre>SW2#show mls qos maps dscp-cos Dscp-cos map: d1: d2 0 1 2 3 4 5 6 7 8 9 ----- 0: 00 00 00 00 00 00 00 01 01 1: 01 01 01 01 01 01 02 02 02 02 2: 02 02 02 02 03 03 03 03 03 03 3: 03 03 04 04 04 04 04 04 04 04 4: 05 05 05 05 05 05 05 05 06 06 5: 06 06 06 06 06 06 07 07 07 07 6: 07 07 07 07 SW2#show mls qos maps dscp-cos Dscp-cos map: d1: d2 0 1 2 3 4 5 6 7 8 9 ----- 0: DSCP 0-7 8,9 1: DSCP 10-15 DSCP 16-19 2: 02 02 02 02 03 03 03 03 03 03 3: 03 03 04 04 04 04 04 04 04 04 4: 05 05 05 05 05 05 05 05 06 06 5: 06 06 06 06 06 06 07 07 07 07 6: 07 07 07 07 DSCP 0-7 = Cos0 DSCP 8-15 = Cos1 DSCP 16-23 = Cos2 DSCP 24-31 = Cos3 DSCP 32-39 = Cos4 DSCP 40-47 = Cos5 DSCP 48-55 = Cos6 DSCP 56-63 = Cos7</pre> <p>8 DSCP values map to 1 CoS Value</p>				
<p>CoS to DSCP mapping:</p>  <p>map CoS 7 to DSCP 22</p>	<pre>SW1#show mls qos maps cos-dscp Cos-dscp map: cos: 0 1 2 3 4 5 6 7 ----- dscp: 0 8 16 24 32 40 48 56 conf terminal mls qos mls qos map cos-dscp 0 8 16 24 32 40 48 22 int fa0/x mls qos trust cos Mapped incoming CoS 7 to DSCP 22: SW1#show mls qos maps cos-dscp Cos-dscp map: cos: 0 1 2 3 4 5 6 7 ----- dscp: 0 8 16 24 32 40 48 22</pre> <p>Incoming CoS value</p> <p>DSCP of Outgoing ptk</p>				
<p>IP Precedence to DSCP mapping:</p>  <p>map IP Precedence 5 to DSCP 7</p>	<pre>SW1# show mls qos map ip-prec-dscp IpPrecedence-dscp map: ipprec: 0 1 2 3 4 5 6 7 ----- dscp: 0 8 16 24 32 40 48 56 conf t mls qos mls qos map ip-prec-dscp 0 8 16 24 32 7 48 56 int fa0/x mls qos trust ip-prec SW1#show mls qos maps ip-prec-dscp IpPrecedence-dscp map: ipprec: 0 1 2 3 4 5 6 7 ----- dscp: 0 8 16 24 32 7 48 56</pre> <p>Incoming IP Prec</p> <p>DSCP of Outgoing ptk</p>				
<p>Outbound policing and match source-mac address!</p> <p>Two solutions:</p>	<pre>access-list 701 permit 0000.1111.1111 0000.0000.0000 class-map match-all SRV1 match access-group 701 policy-map OUT class SRV1 police 1000 conform-action transmit exceed-action drop int fa0/x service-policy output OUT class-map match-all SRV2 match source-address mac 0000.2222.2222 class-map match-all QOS-GRP-2 match qos-group 2 policy-map IN policy-map OUT class SRV2 class QOS-GRP-2 set qos-group 2 police cir 2000000 conform-action transmit exceed-action drop int fa0/1 (inbound) int fa0/2 (outbound) service-policy input IN service-policy input OUT</pre>				
<p>Time-based QoS:</p> <p>Policing http traffic out fa0/0 on weekdays from 11:00 to 15:00</p>	<pre>access-list 102 permit tcp any eq www any time-range QOS access-list 102 permit tcp any any eq www time-range QOS time-range QOS periodic weekdays 11:00 to 15:00 class-map match-all TIME match access-group 102 policy-map OUT class TIME police 10000 conform-action transmit exceed-action drop interface FastEthernet0/0 service-policy output OUT</pre>	<p>What is the difference between shape average</p> <p>And shape peak ?</p>	<p>Shape average config:</p> <pre>policy-map OUT class class-default shape average 16000 64000 0</pre> <p>Only Bc used per Tc</p> <p>Shape peak config:</p> <pre>policy-map OUT class class-default shape average 16000 64000 64000</pre> <p>Will set Bc = Be Each Tc Bc+Be will be used!</p>		
<p>How to configure Shape average:</p> <p>Shape to 16000 bps 8000 byte Bc</p>	<pre>policy-map OUT class class-default shape average 16000 64000 0 int fa0/x service-policy output OUT R1#show policy-map in FastEthernet0/0 Service-policy output: OUT Class-map: class-default (match-any) 65 packets, 18223 bytes 5 minute offered rate 0 bps, drop rate 0 bps Match: any Traffic Shaping Target/Average Byte Sustain Excess Interval Increment Rate Limit bits/int bits/int (ms) (bytes) 16000/16000 8000 64000 0 4000 8000 Adapt Queue Packets Bytes Packets Bytes Shaping Active Depth Delayed Delayed Active - 0 1 60 0 0 no</pre> <p>"Be" has to be set to 0, if not it will be set to the same value as Bc!</p> <p>Be, must be set to 0</p> <p>Bc</p> <p>CIR</p> <p>Tc = Bc/CIR 4 sec = 64000 / 16000</p> <p>Bc = 8000 byte or 64000 bit/s</p> <p>CIR = 16000 bit/s</p>	<p>Shape average explained:</p> <p>←-4 seconds→ ←-4 seconds→</p> <p>Sent: 8000 bytes 8000 bytes</p> <p>In Shape average, no Be is sent!</p> <pre>policy-map OUT class class-default shape average 16000 64000 0</pre>			

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

Embedded packet capture

<p>Embedded Packet capture</p> <p>EPC</p>	<p>monitor capture MYCAP buffer circular packets 1000 monitor capture MYCAP buffer size 10 monitor capture MYCAP interface Gig0/0/1 in monitor capture MYCAP access-list MYACL monitor capture MYCAP start monitor capture MYCAP stop monitor capture MYCAP export bootflash:EPC1.pcap</p> <p>show monitor capture CAP parameter show monitor cap CAP buffer show monitor capture CAP buffer dump show monitor capture CAP buffer brief show monitor capture CAP buffer detail</p>				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin


<p>What do you do if you see the following:</p> <pre>R1(config)#int ser0/1 R1(config-if)#ip address 130.4.0.1 255.255.255.0 Bad mask /24 for address 130.4.0.1</pre>	<pre>R1(config)#int ser0/1 R1(config-if)#ip address 130.4.0.1 255.255.255.0 Bad mask /24 for address 130.4.0.1 R1(config-if)#do sh run i subnet-zero no ip subnet-zero R1(config-if)#ip subnet-zero <- Solution R1(config)#int ser0/1 R1(config-if)#ip address 130.4.0.1 255.255.255.0</pre>				
<p>How to detect a MFR MultiLink bundle member is not active:</p>	<pre>show ppp multilink interface Multilink1314 Multilink1314 Bundle name: R14 Remote Endpoint Discriminator: [1] R14 Local Endpoint Discriminator: [1] R13 Bundle up for 03:17:26, total bandwidth 4632, load 1/255 Receive buffer limit 36000 bytes, frag timeout 1000 ms 0/0 fragments/bytes in reassembly list 0 lost fragments, 387 reordered 0/0 discarded fragments/bytes, 0 lost received 0x1472 received sequence, 0x1487 sent sequence Member links: 3 active, 1 inactive (max not set, min not set) Se1/3, since 03:17:26 Se1/1, since 03:17:26 Se1/0, since 03:17:26 Se1/2 (inactive) Ser1/2 ppp multilink group 1413</pre>				
	<pre>show spanning-tree blockedports Show ALL frame-relay map</pre>				
<p>TCL test script</p>	<pre>show ip alias copy output of all IPs in the lab in a txt file tclsh foreach ip { 150.1.1.1 150.1.2.2 } { ping \$ip source loopback 0 }</pre>	<p>Some IOS do not support TCL scripts, use a macro instead:</p>	<pre>conf t macro name PINGS do ping 150.1.1.1 source lo0 do ping 150.1.2.2 source lo0 @ SW(config)#macro global apply PINGS</pre>		
<p>IRDP</p>	<pre>advertise themselves as default gateways for hosts on VLAN58 using ICMP messages R5: interface FastEthernet 0/0 ip irdp ip irdp address 155.1.58.5 1000 ip irdp maxadvertinterval 20 ip irdp minadvertinterval 10 SW2: interface Vlan 58 ip irdp ip irdp address 155.1.58.8 500 ip irdp maxadvertinterval 20 ip irdp minadvertinterval 10</pre> <p><i>R5 prio higher</i></p> <p><i>Not using HSRP/ VRRP/GLBP For first hop redundancy</i></p>				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!


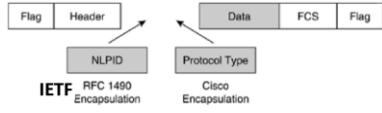
[Donate](#)



Thanks for appreciating my efforts

Colin

Frame-Relay


<p>Inverse-ARP and LMI</p> <p>Pinging the broadcast addr problem:</p>	<pre>R4# interface Serial0/0 encapsulation frame-relay ip address 54.1.1.6 255.255.255.0 R5# interface Serial6 encapsulation frame-relay ip address 54.1.1.6 255.255.255.0 interface Serial6.51 point-to-point ip address 54.1.3.254 255.255.255.0 frame-relay interface-dlci 51 interface Serial6.100 point-to-point ip address 54.1.2.254 255.255.255.0 frame-relay interface-dlci 100 interface Serial6.101 point-to-point ip address 54.1.1.254 255.255.255.0 frame-relay interface-dlci 101</pre> <p>R4# ping 255.255.255.255</p> <p>Reply to request 0 from 54.1.2.254, 76 ms Reply to request 0 from 54.1.1.254, 108 ms Reply to request 0 from 54.1.3.254, 92 ms</p> 	<p>Configuration of the different Frame-Relay encapsulation types:</p>	<pre>interface serial4/2 encapsulation frame-relay [cisco / ietf] Default is Cisco</pre>	<p>What does one have to keep in mind while changing a Frame-Relay sub-interface type from Ser0/1.22 point-to-point To a Ser0/1.22 point</p>	<pre>Conf t Int ser0/1.22 point-to-point Int ser0/1.22 multipoint RELOAD</pre>
<p>Frame-Relay</p> <p>Troubleshooting commands</p>	<pre>show frame-relay pvc [101] show frame-relay map show frame-relay lmi debug frame-relay lmi</pre>	<p>Output of Show frame-relay lmi</p>	<pre>R1#show frame-relay lmi LMI Statistics for interface Serial0/0 (Frame Relay DTE) LMI TYPE = CISCO Invalid Unnumbered Info 0 Invalid Prot Disc 0 Invalid Dummy Call Ref 0 Invalid Map Type 0 Invalid Status Message 0 Invalid Lock Shift 0 Invalid Information ID 0 Invalid Report IE Len 0 Invalid Report Request 0 Invalid Keep IE Len 0 Num Status Enq. Recv 144 Num Status Maps Rcvd 145 Num Update Status Rcvd 0 Num Status Timeouts 0</pre>	<p>Where does one specify sub-interface Frame-Relay DLCI information?</p>	<pre>interface Serial3/1.304 point-to-point frame-relay interface-dlci 304 interface Serial3/1.304 multipoint frame-relay interface-dlci 304</pre>
<p>Frame-Relay</p> <p>Configuration using LMI and Inverse arp</p>	<p>Customer:</p> <pre>interface Serial0/0 encapsulation frame-relay ip address 54.1.1.6 255.255.255.0</pre> <p>Service Provider side:</p> <pre>interface Serial1/6 no ip address encapsulation frame-relay clock rate 64000 frame-relay intf-type dce interface Serial1/6.101 point-to-point ip address 54.1.1.254 255.255.255.0 frame-relay interface-dlci 101</pre>	<p>Turning Frame-Relay LMI autosense off by:</p>	<pre>Int ser0/x frame-relay lmi-type lmi-type cisco ansi Q933a Also configure a keepalive of 10 seconds when fixing the LMI type via: keepalive 10</pre>	<p>Why does one don't create a static ip map entries for frame-relay point-to-point interfaces</p>	<p>it is always assumed that the end point of the point-to-point connection automatically resides on the same subnet as the start point.</p> <p>(There is only one destination.)</p> <p>Inverse-Arp is not necessary too.</p>
<p>Show frame-relay map</p>	<p>All is good, DLCI is active</p> <pre>R1#show frame-relay map Serial0/0 (up): ip 155.1.0.5 dlci 105(0x69,0x1890), static, broadcast, CISCO, status defined, active</pre> <p>Status deleted, re-do mapping and reload the router.</p> <pre>R1#show frame-relay map Serial0/1 (down): ip 155.1.0.5 dlci 105(0x69,0x1890), static, broadcast, CISCO, status deleted</pre>	<p>show frame-relay route</p>	<pre>R1#show frame-relay route Input Intf Input Dlci Output Intf Output Dlci Status Serial0/0 102 Serial1/1 201 active Serial1/1 201 Serial0/0 102 inactive</pre>	<p>Turning a router into a Frame-Relay Switch</p>	<pre>interface Serial4/1 encapsulation frame-relay clockrate 64000 frame-relay intf-type dce frame-relay route 304 interface Serial4/3 403 304 is the local Ser4/1 DLCI</pre>
<p>frame-relay map ip 155.1.0.5 105 broadcast</p> <p>And pings to 255.255.255.255</p>	<pre>R1#ping 255.255.255.255 Type escape sequence to abort. Sending 5, 100-byte ICMP Echos to 255.255.255.255, timeout is 2 seconds: Reply to request 0 from 155.1.0.5, 88 ms Reply to request 1 from 155.1.0.5, 88 ms Reply to request 2 from 155.1.0.5, 88 ms Reply to request 3 from 155.1.0.5, 88 ms Reply to request 4 from 155.1.0.5, 124 ms</pre> <p>This is how it should be, no other host. from another DLCI answers the ping to the broadcast.</p>	<p>Define Inverse ARP</p>	<p>resolve a next hop network protocol address to a local DLCI value</p>	<p>Show frame-relay pvc</p>	<pre>Router#show frame-relay pvc PVC Statistics for interface Serial3/0 (Frame Relay DTE) Local Active Inactive Deleted Static Switched 0 0 0 0 Unused 0 0 0 0</pre>
<p>What kind of Frame-relay encapsulations are there?</p>		<p>Enable / disable Frame-Relay inverse-arp</p>	<pre>Disable: Int ser0/x no frame-relay inverse-arp ip [DLCI 100] Enable: Int ser0/x frame-relay inverse-arp ip [DLCI 100] End clear frame-relay inarp interface Ser0/x</pre>	<p>What types of DLCI assignments are there on Frame-Relay networks?</p>	<p>Global and local DLCI assignments</p> <p>Local is usually used.</p>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin


Frame-Relay

<p>What does Frame-Relay NNI mode do generally?</p>	<p>End-to-End PVC keepalive to signal the DLCI status</p> <p>Cisco proprietary protocol</p>	<p>How to monitor the status of a Frame-Relay End-to-End keepalive session?</p>	<pre>show frame-relay end-to-end keepalive R4#show frame-relay end-to-end keepalive End-to-end Keepalive Statistics for Interface Serial1 (Frame Relay DTE) DLCI = 300, DLCI USAGE = LOCAL, VC STATUS = ACTIVE (EEK UP) SEND SIDE STATISTICS Send Sequence Number: 228, Receive Sequence Number: 221 Configured Event Window: 3, Configured Error Threshold: 2 Total Observed Events: 229, Total Observed Errors: 6 Monitored Events: 3, Monitored Errors: 0 Successful Successes: 3, End-to-end VC Status: UP RECEIVE SIDE STATISTICS Send Sequence Number: 221, Receive Sequence Number: 228 Configured Event Window: 3, Configured Error Threshold: 2 Total Observed Events: 227, Total Observed Errors: 3 Monitored Events: 3, Monitored Errors: 0 Successful Successes: 3, End-to-end VC Status: UP</pre>	<p>How many virtual templates can exist on a Router typically?</p> <p>How many virtual interfaces can exist on a Router typically?</p>	<p>-25 virtual template interfaces</p> <p>-300 virtual interfaces</p>																				
<p>Frame-Relay NNI send and receive side Status's</p>	<table border="1"> <thead> <tr> <th>Keepalive Receive Side Status</th> <th>Keepalive Send Side Status</th> <th>LMI Status</th> <th>Overall VC Status</th> </tr> </thead> <tbody> <tr> <td>UP</td> <td>UP</td> <td>UP</td> <td>UP</td> </tr> <tr> <td>DOWN</td> <td>X</td> <td>X</td> <td>DOWN</td> </tr> <tr> <td>X</td> <td>DOWN</td> <td>X</td> <td>DOWN</td> </tr> <tr> <td>UP</td> <td>UP</td> <td>DOWN</td> <td>DOWN</td> </tr> </tbody> </table>	Keepalive Receive Side Status	Keepalive Send Side Status	LMI Status	Overall VC Status	UP	UP	UP	UP	DOWN	X	X	DOWN	X	DOWN	X	DOWN	UP	UP	DOWN	DOWN	<p>Enabling Frame-Relay End-to-End Keepalive:</p> <p>(Customer Side)</p> <p>All options described</p>	<pre>interface Serial0/1.100 point-to-point frame-relay interface-dlci 100 class FRAME_EEK map-class frame-relay FRAME_EEK frame-relay end-to-end keepalive mode bidirectional frame-relay end-to-end keepalive event-window send 10 frame-relay end-to-end keepalive success-events send 5 frame-relay end-to-end keepalive error-threshold {send receive} count frame-relay end-to-end keepalive timer {send receive} interval Event-Window, number of latest events to use the check routine on, last 10. Success-window, consecutive success events required to change from DOWN to UP status. (5 OKs) Error-threshold, number of errors needed to change from UP to DOWN status</pre>	<p>PPP over Frame Relay Applied to DLCI 101</p> <p>Config:</p>	<pre>interface serial 3/0 no ip address encapsulation frame-relay ! interface serial 3/0.1 point-to-point frame-relay interface-dlci 101 ppp virtual-template1 ! interface Virtual-Template1 ip unnumbered loopback0 ppp authentication chap ! interface loopback 0 ip address 172.16.1.1 255.255.255.252</pre>
Keepalive Receive Side Status	Keepalive Send Side Status	LMI Status	Overall VC Status																						
UP	UP	UP	UP																						
DOWN	X	X	DOWN																						
X	DOWN	X	DOWN																						
UP	UP	DOWN	DOWN																						
<p>What types of Frame Relay End-to-End Keepalive Modes are there?</p>	<p>Bidirectional Mode</p> <p>Request Mode</p> <p>Reply Mode</p> <p>Passive-Reply Mode</p>	<pre>debug frame-relay end-to-end keepalive event</pre> <p>output</p>	<pre>debug frame-relay end-to-end keepalive event 1w5d: EEK SUCCESS (request, Serial2/1.100 DLCI 100) 1w5d: EEK SUCCESS (reply, Serial2/1.100 DLCI 100) 1w5d: EEK SUCCESS (request, Serial2/1.100 DLCI 100) 1w5d: EEK receiver timeout (Serial2/1.100 DLCI 100) 1w5d: EEK stopped (Serial2/1.100 DLCI 100)</pre>	<p>PPP over Frame Relay: Output of show frame-relay pvc Command of DLCI 101</p>	<pre>R4#show frame-relay pvc 101 PVC Statistics for Interface Serial1 (Frame Relay DTE) DLCI = 101, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE, INTERFACE = Serial1.1 input pkts 50 output pkts 46 in bytes 2067 out bytes 1357 dropped pkts 0 in pkts dropped 0 out pkts dropped 0 out bytes dropped 0 in FECN pkts 0 in BECN pkts 0 out FECN pkts 0 out BECN pkts 0 in DE pkts 0 out DE pkts 0 out board pkts 1 out board bytes 500 PVC create time 00:01:37, last time pvc status changed 00:01:37Bound to Virtual-Access1 [up, cloned from Virtual-Template1]</pre>																				
<p>Frame-Relay End-to-End types and their detailed description:</p>	<p>Frame-Relay End-to-End keepalive types:</p> <p>Bidirectional: Send side sends, then goes into receive mode, then back to send mode. → Ping-Pong between the two Nodes.</p> <p>Request Mode: Only send enabled</p> <p>Reply Mode: Only replies to received messages</p> <p>Passive Reply-Mode Only responds, does not track errors or adjusts timers.</p>	<p>FRF.5 Frame Relay to ATM Network Interworking</p> <p>Picture</p>		<p>debug ppp negotiation</p> <p>How to know if PPP LCP and NCP have been successful?</p>	<p>LCP successful:</p> <p>Vi1 LCP: State is Open</p> <p>Vi1 PPP: Phase is UP [0 sess, 0 load]</p> <p>Now each NCP can start negotiate, like IPCP:</p> <p>Vi1 IPCP: State is Open</p>																				
<p>Enabling Frame-Relay End-to-End Keepalive:</p> <p>(Customer Side)</p>	<pre>map-class frame-relay FRM-RLY frame-relay end-to-end keepalive mode [bi-directional] interface Serial0/1.100 point-to-point frame-relay class FRM-RLY</pre>	<p>FRF.8.1 Translation Mode</p> <p>ATM to Frame-relay configuration</p>	<pre>FRF.8.1 Translation Mode Frame-Relay to ATM Cloud Frame-Relay ATM enabled Frame-Relay-RTR# interface Serial0/0 no ip address encapsulation frame-relay IETF interface Serial0/0.16 point-to-point ip address 172.16.1.1 /24 frame-relay interface-dlci 16 ATM-RTR# interface ATM0 no ip address no atm ilmi-keepalive interface ATM0.1 multipoint ip address 172.16.1.2 /24 pvc 1/32 protocol ip 172.16.1.1 broadcast encapsulation aal5snap</pre>	<p>debug frame-relay ppp</p> <p>debug frame-relay ppp once working and once faulty</p>	<pre>debug frame-relay ppp PPPoFR working FR-PPP: process on Virtual-Access1, #out-pkts=497 FR-PPP: process on Virtual-Access1, #out-pkts=498 FR-PPP: process on Virtual-Access1, #out-pkts=499 FR-PPP: process on Virtual-Access1, #out-pkts=500 PPPoFR faulty FR-PPP: encaps failed for FR VC 101 on Serial1 down FR-PPP: input- Serial1 vc or va down, pak dropped</pre>																				
<p>Enabling Frame-Relay End-to-End Keepalive:</p> <p>(ISP Side)</p>	<pre>interface Serial1/1 encapsulation frame-relay frame-relay lmi-type ansi frame-relay intf-type nni frame-relay route 200 interface Serial4/1 100</pre>	<p>How to check ATM side between a FRF.8.1 to Frame relay connection</p>	<pre>show atm vc inter atm4/0/0 1 32 VPI = 1 VCI = 32 Interface: ATM4/0/0, Type: t1suni VPI = 1 VCI = 32 Status: UP Time-since-last-status-change: 04:40:13 Connection-type: PVC Cast-type: point-to-point Packet-discard-option: disabled Usage-Parameter-Control (UPC): pass Wrr weight: 2 Number of OAM-configured connections: 0</pre>	<p>PPP over Frame Relay</p> <p>Client negotiating for IP</p> <p>Server offering IP config</p>	<pre>PPPoFR "Client" interface Virtual-Template1 ip address negotiated interface Serial1.1 point-to-point frame-relay interface-dlci 101 ppp Virtual-Template1 PPPoFR "DHCP Server" interface Virtual-Template22 ip unnumbered Loopback0 peer default ip address pool LOCAL-DHCP-POOL-NAME interface Serial1.1 point-to-point frame-relay interface-dlci 101 ppp Virtual-Template22</pre>																				

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

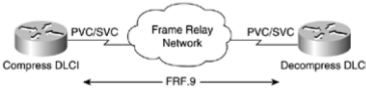
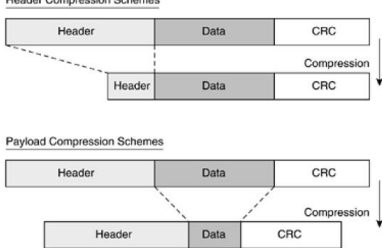
Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin


Frame-Relay

<p>debug ppp authentication</p> <p>Successful session</p>	<p>debug ppp authentication</p> <pre> Vi1 CHAP: O CHALLENGE id 24 len 28 from "R2" Vi1 CHAP: I CHALLENGE id 24 len 28 from "R3" Vi1 CHAP: O RESPONSE id 24 len 28 from "R2" Vi1 CHAP: I RESPONSE id 24 len 28 from "R3" Vi1 CHAP: O SUCCESS id 24 len 4 Vi1 CHAP: I SUCCESS id 24 len 4 </pre>	<p>Show frame-relay multilink detailed</p>	<pre> R1#show frame-relay multilink detailed Bundle: MFR0, State = up, class = A, fragmentation disabled BID = MFR0 No. of bundle links = 2, Peer's bundle-id = MFR0 Bundle links: Serial0, HW state = up, link state = Up, LID = Serial0 Cause code = none, Ack timer = 4, Hello timer = 10, Max retry count = 2, Current count = 0, Peer LID = Serial3/3, RTT = 4 ms Statistics: Add_link sent = 2, Add_link rcv'd = 1, Add_link ack sent = 1, Add_link ack rcv'd = 2, Add_link rej sent = 0, Add_link rej rcv'd = 0, Remove_link sent = 0, Remove_link rcv'd = 0, Remove_link ack sent = 0, Remove_link ack rcv'd = 0, Hello sent = 1105, Hello rcv'd = 1106, Hello_ack sent = 1106, Hello_ack rcv'd = 1105, outgoing pak dropped = 0, incoming pak dropped = 0 Serial3, HW state = up, link state = Up, LID = Serial3 Cause code = none, Ack timer = 4, Hello timer = 10, Max retry count = 2, Current count = 0, Peer LID = Serial0/0, RTT = 4 ms Statistics: ----- </pre>	<p>What is important when using frame-relay compression?</p>	<p>It needs to be configured End-to-End!</p> 
<p>debug ppp authentication</p> <p>Failed session</p>	<p>Debug ppp authentication (failed session)</p> <pre> %LINK-3-UPDOWN: Interface Virtual-Access1, changed state to up Vi1 PPP: Treating connection as a dedicated line Vi1 CHAP: O CHALLENGE id 25 len 28 from "R2" Vi1 CHAP: I CHALLENGE id 25 len 28 from "R3" Vi1 CHAP: O RESPONSE id 25 len 28 from "R2" Vi1 CHAP: I RESPONSE id 25 len 28 from "R3" Vi1 CHAP: O FAILURE id 25 len 25 msg is "MD/DES compare failed" </pre>	<p>MFR BID and LID configuration of MFR / FRF.16 configurations</p>	<pre> R1(config)#interface mfr0 R1(config-if)#frame-relay multilink bid M-LINK-BID-ID (Shut / no shut) R1(config)#interface Serial0 R1(config-if)#frame-relay multilink lid REMOTE_LINK_1 R1(config-if)#interface Serial3 R1(config-if)#frame-relay multilink lid REMOTE_LINK_2 </pre>	<p>Configuring Frame-Relay FRF.9 payload compression:</p>	<p>Physical or multipoint interface:</p> <pre> frame-relay map protocol protocol-address dlci payload-compress FRF9 stac [hardware-options] </pre> <p>Interface Serial0/1 (or multipoint)</p> <pre> frame-relay map ip x.x.x.x dlci payload-compression frf9 stac </pre> <p>Logical Sub-Interface</p> <pre> frame-relay payloadcompress FRF9 stac [hardware-options] </pre> <p>interface Serial3/0.200 point-to-point</p> <pre> frame-relay payload-compression frf9 stac software </pre>
<p>Enabling Frame Relay SVC on the Physical interface</p>	<p>Frame-Relay SVC on physical interface:</p> <pre> interface serial4/2 encapsulation frame-relay map-group svc_group frame-relay svc </pre>	<p>MFR / FRF.16</p> <p>Sub-interface config details</p>	<pre> Interface Serial 0/1 frame-relay multilink hello 10 [in seconds] frame-relay multilink retry 4 [in seconds] frame-relay multilink ack 2 default values are 10 seconds, 4 seconds, and 2 tries </pre>	<p>TCP/IP Header Compression over Frame Relay</p>	<pre> interface Serial2/3 ip address x.x.x.1 255.255.255.0 encapsulation frame-relay frame-relay ip tcp header-compression frame-relay map ip z.z.z.2 100 broadcast frame-relay map ip x.x.x.3 101 broadcast nocompress interface Serial2/3.200 multipoint ip address x.x.x.x.255.255.255.0 frame-relay map ip x.x.x.x 300 CISCO tcp header-compression passive explicitly configured for DLCI 100 active compression mode, 101 inherited compression disabled by nocompress, passive compression on DLCI 200 </pre>
<p>Enabling Frame Relay SVC on a Point-to-Point Subinterface</p>	<pre> interface serial1/0 encapsulation frame-relay frame-relay svc interface serial1/0.99 point-to-point map-group svc_group </pre>	<p>MFR / FRF.16 debug commands:</p>	<p>debug frame-relay multilink</p> <p>debug frame-relay multilink control</p>	<p>Details about TCP/IP Header Compression</p>	<p>Frame-Relay Cisco encapsulation needs to be used.</p> <pre> Serial0/1 encapsulation frame-relay cisco Active and Passive modes, active is default. </pre>
<p>show idb</p>	<p>interface descriptor block (IDB), which consists of hardware IDB and software IDB</p> <p>maximum number of IDBs that a platform can support</p> <p>hardware IDB controls the physical interface, whereas the software IDB controls the Layer 2 encapsulation.</p>	<p>Frame Relay Compression</p> <p>header compression versus payload compression</p>	 <p>Can be configured per physical or per logical sub-interface</p>	<p>RTP header compression</p> <p>CRTP</p>	<pre> interface Serial2/3 ip address 172.16.1.1 255.255.255.0 encapsulation frame-relay frame-relay map ip 172.16.1.2 100 broadcast frame-relay ip rtp header-compression ! interface Serial2/3.200 multipoint ip address 192.168.1.1 255.255.255.0 frame-relay map ip 192.168.1.2 300 CISCO rtp header-compression passive Default Mode is active </pre>
<p>FRF.16 Multilink configuration</p>	<pre> R1# interface MFR0 ! interface MFR0.1 point-to-point ip address 172.16.1.1 255.255.255.0 frame-relay interface-dlci 103 interface Serial0 no ip address encapsulation frame-relay MFR0 interface Serial3 no ip address encapsulation frame-relay MFR0 </pre> <pre> R2# interface MFR0 frame-relay intf-type dce ! interface Serial3/0 encapsulation frame-relay frame-relay intf-type dce interface Serial3/3 encapsulation frame-relay MFR0 interface Serial4/3 encapsulation frame-relay MFR0 ! connect MFR MFR0 103 Serial3/0 301 </pre>	<p>Frame Relay Cisco Proprietary Payload Compression configuration</p>	<pre> Frame-Relay compression: Using DLCI 100 interface Serial3/0 (or multipoint) encapsulation frame-relay frame-relay map ip X.X.X.X 100 payload-compression packet-by-packet interface Serial3/0.200 point-to-point frame-relay interface-dlci 200 frame-relay payload-compress packet-by-packet </pre>	<p>Verifying and Monitoring Frame Relay Compressions</p>	<p>show compress</p> <pre> R1#show compress Serial3/2 Software compression enabled uncompressed bytes xmt/rcv 1089240/1200 1 min avg ratio xmt/rcv 5.564/0.049 5 min avg ratio xmt/rcv 5.563/0.051 10 min avg ratio xmt/rcv 5.563/0.051 no bufs xmt 0 no bufs rcv 0 resyncs 0 Additional Stacker Stats: Transmit bytes: Uncompressed = 0 Compressed = 189045 Received bytes: Compressed = 920 Uncompressed = 0 show interface (reduced traffic rate 30sec) </pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

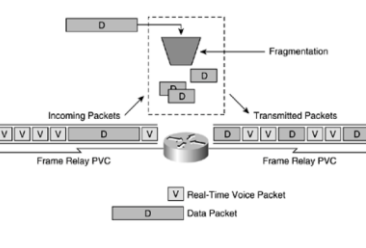
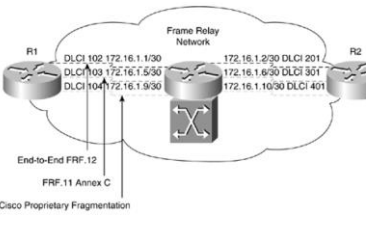
Ranging 5 bucks to unlimited!



Thanks for appreciating my efforts

Colin

Frame-Relay


<p>What types of header compressions are there for Frame-Relay?</p>	<p>Cisco Proprietary Payload Compression (per-packet)</p> <p>FRF.9 Payload Compression</p> <p>TCP/IP Header Compression</p> <p>CRTP, RTP compression</p>	<p>Frame-Relay fragmentation and interleaving</p> <p>Explained as a picture</p>		<p>What is an easy method for removing broadcast traffic for unused DLCI's which are not avoided by disabling Inverse-arp for those DLCI's?</p>	<p>Configure them on an unused sub-interface, for example Ser0/0.999</p> <pre>interface Serial0/0 encapsulation frame-relay ip address 155.1.200.3 255.255.255.0 ! interface Serial0/0.999 multipoint no ip address frame-relay interface-dlci 401 frame-relay interface-dlci 402</pre> <p>DLCI's 401, 402 are not used, and are placed into a unused sub-int ser0/0.999, Now one does not receive any Broadcast traffic for those DLCI's once pinging 255.255.255.255 while not disabling Inverse-Arp for those DLCI's.</p>
<p>Show frame relay-map, checking for header-compression:</p>	<pre>R1#show frame-relay map Serial3/2.100 (up): point-to-point dlci, dlci 100(0x64,0x1840), broadcast status defined, active, RTP Header Compression (inherited), connections: 256</pre>	<p>How is Frame-Relay FRF.12 fragmentation configured?</p>	<p>End-to-End FRF.12 Fragmentation is configured on a per-PVC basis using a Frame Relay map class</p> <pre>Map-class X Frame-relay fragment fragment_size</pre>	<p>How to disable LMI on Frame-Relay (for Back-to-Back configs)</p>	<pre>interface Serial0/1 no keepalive</pre>
<p>show frame-relay ip rtp header-compression interface Ser0/0</p>	<pre>R1#show frame-relay ip rtp header-compression interface Serial3/2 DLCI 100 Link/Destination info: point-to-point dlci Interface Serial3/2: (passive, compression on) Rcvd: 703 total, 699 compressed, 2 errors 2 dropped, 0 buffer copies, 0 buffer failures Sent: 716 total, 713 compressed, 27073 bytes saved, 115527 bytes sent 1.23 efficiency improvement factor Connect: 101 rx slots, 101 tx slots,</pre>	<p>show frame-relay fragment</p> <p>Output:</p>	<pre>R2#show frame-relay fragment interface dlci frag-type frag-size in-frag out-frag dropped-frag Serial3/3 200 end-to-end 100 0 0 0</pre>	<p>Frame-Relay Back-to-Back Configuration.</p>	<pre>R1# interface Serial0/1 ip address 155.1.45.4 255.255.255.0 encapsulation frame-relay no keepalive frame-relay map ip 155.1.45.5 514 broadcast Disabling LMI via "no keepalive"</pre> <pre>R2# interface Serial0/1 ip address 155.1.45.5 255.255.255.0 encapsulation frame-relay no keepalive clock rate 64000 frame-relay map ip 155.1.45.4 514 broadcast</pre>
<p>Frame Relay IP RTP Priority configuration</p>	<pre>interface Serial3/2 no ip address encapsulation frame-relay frame-relay traffic-shaping ! interface Serial3/2.100 point-to-point ip address 172.16.1.1 255.255.255.252 frame-relay interface-dlci 100 class SHAPING frame-relay ip rtp header-compression ! map-class frame-relay SHAPING frame-relay cir 38400 frame-relay bc 4800 frame-relay be 0 no frame-relay adaptive-shaping frame-relay fragment 250 frame-relay ip rtp priority 16384 16383 1024</pre>	<p>show frame-relay fragment interface ser0/x DLCI-NR</p> <p>Output:</p>	<pre>R2#show frame-relay fragment interface Serial3/3 200 fragment size 100 fragment type end-to-end in fragmented pkts 20 out fragmented pkts 20 in fragmented bytes 1140 out fragmented bytes 1140 in un-fragmented pkts 0 out un-fragmented pkts 0 in un-fragmented bytes 0 out un-fragmented bytes 0 in assembled pkts 10 out pre-fragmented pkts 10 in assembled bytes 1040 out pre-fragmented bytes 1040 in dropped reassembling pkts 0 out dropped fragmenting pkts 0 in timeouts 0 in out-of-sequence fragments 0 in fragments with unexpected 0 bit set 0 in fragments with skipped sequence numbers 0 out interleaved packets 0</pre>	<p>Back to back indication on</p> <p>Show frame-relay PVC output:</p>	<pre>DLCI = 514, DLCI USAGE = LOCAL, PVC STATUS = STATIC, INTERFACE = Serial0/1/0 input pkts 228 output pkts 233 in bytes 23712 out bytes 24232 dropped pkts 0 in pkts dropped 0 out pkts dropped 0 out bytes dropped 0 in FECN pkts 0 in BECN pkts 0 out FECN pkts 0 out BECN pkts 0 in DE pkts 0 out DE pkts 0 out bcst pkts 0 out bcst bytes 0 5 minute input rate 0 bits/sec, 0 packets/sec 5 minute output rate 0 bits/sec, 0 packets/sec pvc create time 00:14:29, last time pvc status changed 00:12:48 PVC STATUS = STATIC = Back-to-Back PVC STATUS = ACTIVE = learned via LMI or via frame-relay intf-dlci command.</pre>
<p>What types of Frame Relay Fragmentation are available?</p>	<p>Cisco proprietary</p> <p>FRF.11 Annex C</p> <p>FRF.12 Frame Relay Fragmentation</p>	<p>Frame Relay fragmentation types per DLCI, picture</p>		<p>Frame Relay Broadcast Queue</p>	<pre>interface Serial0/0 frame-relay broadcast-queue 100 256000 36 <100> Queue size for broadcasts 100 packets <256000> Byte rate per sec. <36> Max. packets/S broadcasts</pre>
<p>serialization delay formula:</p>	<p>serialization delay =</p> <p>frame size (in bits) / link bandwidth (in bits/sec)</p> <p>(a 1500-byte frame takes approximately 214 ms to leave the router on a 56-kbps line.)</p>	<p>LMI and Inverse-arp</p> <p>Describe their tasks</p>	<p>LMI dynamically discovers the DLCI / PVCs per interface.</p> <p>Inverse-Arp discovers the mapped IP associated to the learned DLCI number.</p>	<p>show frame-relay ip [tcp/rtp] header-compression:</p>	<pre>R5#show frame-relay ip tcp header-compression DLCI 501 Link/Destination info: ip 155.1.100.1 Interface Serial0/0.100 DLCI 501 (compression on, VJ) Rcvd: 67 total, 65 compressed, 0 errors, 0 status msgs 0 dropped, 0 buffer copies, 0 buffer failures Sent: 46 total, 43 compressed, 0 status msgs, 0 not predicted 1473 bytes saved, 435 bytes sent 4.38 efficiency improvement factor Connect: 256 rx slots, 256 tx slots, 2 misses, 0 collisions, 0 negative cache hits, 255 free contexts 95% hit ratio, five minute miss rate 0 misses/sec, 0 max</pre>

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)



Thanks for appreciating my efforts

Colin

Frame-Relay

<p>PPP over Frame Relay (PPPoFR)</p> <p>Order of implementation:</p>	<p>First create the Virtual-Template:</p> <pre>interface Virtual-Template99 ip address 155.1.0.5 255.255.255.0</pre> <p>Second, bind the Virtual-Template to the DLCI</p> <pre>interface Serial0/0/0 no ip address encapsulation frame-relay frame-relay interface-dlci 501 ppp Virtual-Template99</pre> <p>Log message appears: Virtual-Access1, changed state UP</p>	<h2>Frame-Relay Switching</h2> <p>Show connection out:</p>	<pre>R3#show connection all</pre> <table border="1"> <thead> <tr> <th>ID</th> <th>Name</th> <th>Segment 1</th> <th>Segment 2</th> <th>State</th> </tr> </thead> <tbody> <tr> <td>1</td> <td>R1_R2</td> <td>Se1/2 132</td> <td>Se1/3 231</td> <td>UP</td> </tr> </tbody> </table>	ID	Name	Segment 1	Segment 2	State	1	R1_R2	Se1/2 132	Se1/3 231	UP		
ID	Name	Segment 1	Segment 2	State											
1	R1_R2	Se1/2 132	Se1/3 231	UP											
<p>Show interface virtual-access</p> <p>Output:</p>	<pre>Rack1R1#show interface virtual-access2 Virtual-Access2 is up, line protocol is up Hardware is Virtual Access interface Internet address is 155.1.0.1/24 MTU 1500 bytes, BW 1000000 Kbit, DLY 1000000 usec, reliability 255/255, txload 1/255, rxload 1/255 Encapsulation PPP, LCP Open Open: IPCP PPPoFR vaccess, cloned from Virtual-Template1 Vaccess status 0x44 Bound to Serial0/0 DLCI 105, Cloned from Virtual-Template1, loopback not set Keepalive set (10 sec) DTR is pulsed for 5 seconds on reset Last input 00:02:00, output never, output hang never Last clearing of "show interface" counters 00:02:17 Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0 Queueing strategy: fifo Output queue: 0/40 (size/max) 5 minute input rate 0 bits/sec, 0 packets/sec 5 minute output rate 0 bits/sec, 0 packets/sec</pre>	<p>Enable CDP over Frame-Relay</p>	<p>Interface Serial 0/1</p> <p>Encapsulation frame-relay</p> <p>cdp enable</p> <p>Make sure Broadcast is enabled</p>												
<p>PPP over Frame-Relay</p> <p>Virtual-access interfaces can be seen via connected routes:</p>	<pre>R5# interface Virtual-Access1 ip address 155.1.0.5 255.255.255.0</pre> <pre>R5#show ip route connected 155.1.0.0/16 is variably subnetted, 2 subnets, 2 masks C 155.1.0.0/24 is directly connected, Virtual-Access1 C 155.1.0.1/32 is directly connected, Virtual-Access1</pre> <p>R1# (the neighbour)</p> <pre>interface Virtual-Access1 ip address 155.1.0.1 255.255.255.0 end</pre>	<h2>PPP Multi Link via Frame-Relay</h2> <p>Using CHAP</p>	<pre>username Rack1R3 password CISCO interface Multilink1 ip address 174.1.23.2 255.255.255.0 ppp multilink ppp multilink group 1 interface Serial0/0 encapsulation frame-relay no frame-relay inverse-arp interface Serial0/0.203 point-to-point frame-relay interface-dlci 203 ppp Virtual-Template1 interface Serial0/0.213 point-to-point frame-relay interface-dlci 213 ppp Virtual-Template1</pre>												
<p>Bridging over Frame Relay</p> <p>Config:</p>	<pre>bridge 1 protocol ieee interface FastEthernet0/0 bridge-group 1 ! interface Serial0/0 encapsulation frame-relay frame-relay map bridge 205 broadcast bridge-group 1 !</pre> <p>Check if you can see the spanning tree root over the Frame-Relay Bridge!</p>	<h2>PPP Multi Link via Frame-Relay</h2> <p>Using CHAP</p>	<pre>username Rack1R2 password CISCO interface Multilink1 ip address 174.1.23.3 255.255.255.0 ppp multilink ppp multilink group 1 interface Serial1/0 encapsulation frame-relay no frame-relay inverse-arp frame-relay interface-dlci 302 ppp Virtual-Template1 interface Serial1/1 encapsulation frame-relay no frame-relay inverse-arp frame-relay interface-dlci 312 ppp Virtual-Template1</pre>												
<p>Frame-Relay Switching</p> <p>config</p>	<pre>R3# frame-relay switching</pre> <pre>interface Serial 1/2 clock rate 64000 encapsulation frame-relay frame-relay intf-type doe no shutdown interface Serial 1/3 clock rate 64000 encapsulation frame-relay frame-relay intf-type doe no shutdown connect R1_R2 Serial 1/2 132 Serial 1/3 231</pre> <pre>R1: interface Serial 0/1 encapsulation frame-relay ip address 155.1.12.1 255.255.255.0 frame-relay map ip 155.1.12.1 231 R2# interface Serial 0/1 encapsulation frame-relay ip address 155.1.12.2 255.255.255.0 frame-relay map ip 155.1.12.2 231</pre>														
<p>Frame-Relay Switching</p> <p>Show frame-relay pvc</p> <p>Checking SWITCHING USAGE</p>	<pre>R3#show frame-relay pvc inc SWI</pre> <pre>DLCI = 132, DLCI USAGE = SWITCHED, PVC STATUS = ACTIVE, INTERFACE = Serial1/2 DLCI = 231, DLCI USAGE = SWITCHED, PVC STATUS = ACTIVE, INTERFACE = Serial1/3</pre>														

Help me create more flashcards:

Simply press this button and send me your credit cards regards!

Ranging 5 bucks to unlimited!

[Donate](#)

Thanks for appreciating my efforts

Colin

CCIE RS Version 5 study approach:

1. Identify the scope / estimate your know how

Use a training partner to identify all necessary topics for CCIE RS version 5 such as INE, IPEXpert etc.. In my example I have used INE's detailed expanded study blueprint.

Copy it into an Excel spreadsheet and start going through it, this will take a while to go through but its worth it. Be very honest with yourself, if you don't know the command by hart or have an idea, you do NOT know it and all the sub-commands associated with it...

<http://blog.ine.com/2009/05/12/ccie-rs-4x-expanded-study-blueprint/>

	15.3MT topic	unknown	read about it	configured 1-2	intermediate	confident
1. LAN Switching						
1.1. VLANs & Trunking						
1.1.1. Standard VLANs						
1.1.2. Extended VLANs						
1.1.3. VLAN Database						
1.1.4. Access Ports						
1.1.5. 802.1q Trunk Ports						
1.1.6. 802.1q Native VLAN						
1.1.7. Dynamic Trunking Protocol (DTP)						
1.1.8. Trunking Allowed List						
VTP 1						
VTP 2						
VTP 3						
1.2.2. VTP Authentication						

2. Estimate your efforts

	Hours, still to do:	Calculated hours in total
Reading CCIE books		
INE Vol 1 Tasks	76	581
INE Vol 2 Tasks	0	321
INE Vol 3 Tasks	60	60
INE Vol 4 Tasks	40	40
INE Vol 4 Tasks	40	40
INE 3x MOCK LABS	24	24
Nabriks Bootcamp	60	60

	Hours, still to do:	Calculated hours in total
Total hours	240	1125

calculated at 1.7 Labs per hour

Work day	8 hours
Day	24 hours
Working week	40 hours
Week	168 hours
Work-month	160 hours
Month	672 hours

Study time:

3 days during the week 12 hours
1 day weekend 8 hours

20 hours / week

4x 20 = 80 hours per month
12x 80 = 960 hours per year

5. start learning using the Cisco docCD

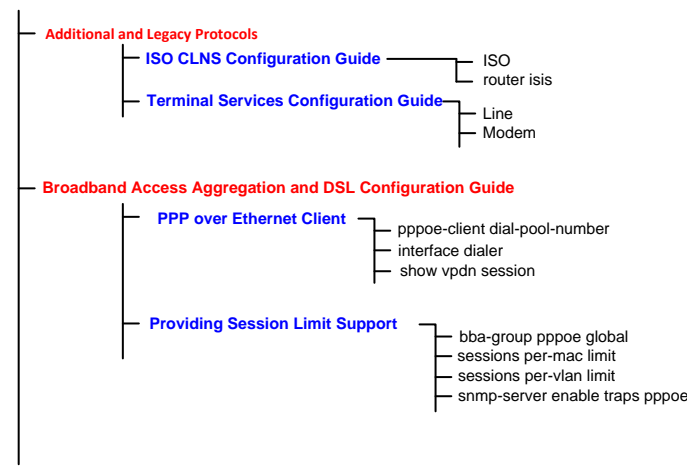
<http://www.cisco.com/c/en/us/support/ios-nx-os-software/ios-software-release-15-3-3-m/model.html>

I recommend going through the entire DocCD of version 15.3MT (configuration guides, not command reference) or what ever your future version for the CCIE is.

Browse through it and take notes where things are, can be commands, can be technologies, basically hints for yourself. Then instead of using Google to find the right page on the docCD search within your own created docCD overview and then directly jump into the section you are interested in. This will speed up your lookup time for commands/technologies initially, plus it makes you aware where to find things over time without Google. The more lab hours you spend the more you will remember this tip written here.. Plus during the lab if you know the feature name/technology but forgot the details, you find the right page within 30 seconds. Giving you another 5 minutes to answer that question.

This document just serves the purpose to CTRL-F:

<http://colin.cant.ch/projects/IOS-15.3T-docCD-overview-v3.xls>



Cisco IOS Software Release 15.3(3)M

Product Overview	
Series:	Cisco IOS 15.3M&T
Latest Version:	This Version
Status:	Orderable
End-of-Sale Date:	None Announced
End-of-Support Date:	None Announced

Documentation	Downloads	Community Content
Configuration Guides Command References Configuration Examples and TechNotes		End-of-Life and End-of-Sale Notices Release Notes Troubleshooting TechNotes

Configuration Guides
Additional and Legacy Protocols DECnet Configuration Guide, Cisco IOS Release 15M&T ISO CLNS Configuration Guide, Cisco IOS Release 15M&T Novell IPX Configuration Guide Cisco IOS Release 15M&T Terminal Services Configuration Guide, Cisco IOS Release 15M&T

6. Lab, lab, lab, and more lab....

Use a training partners workbook, there are plenty out there find one, stick to it, and don't jump from one vendor to another during your studies.

7. Bootcamps

I have visited the 5 day bootcamp performed by Narbik, the most incredible Cisco instructor I have ever met.

I had re-visited his 10 day bootcamp for a minor upgrade fee, his re-takes of the same bootcamp / version are for free and he encourages people to do so.

I highly recommend re-taking his class, the first time you go, you will have a serious buffer overflow in terms of content and quality of information received.

The second time you attend his class, you will be able to sip up details you could not process / digest the first time.

I have attended many Cisco courses in my past, but this is unlike anything you have ever seen before! His memory is sensational and he knows every command / formula / ethertype etc... by hart and never needs to look up a thing in a PDF or similar. In addition, I find his way of teaching and personality very entertaining in terms of how he delivers his class etc. Be prepared for long hours, frustrating GOOD labs which make you remember what you did wrong or what alternatives you have in each situation. It will open your eyes in terms of how protocols really work...

Narbiks training: <http://micronicstraining.com/>

8. create a process for each technology etc, follow it each time you configure that technology

Once you have gained the "entire" picture in terms of technologies, go through the the initial Scope list from Step 1 and make sure you have for each item / technology a process in place which you follow each time you configure it.

Or best, create you own set of cards!

Follow your process EVERY TIME from the beginning to the end and you will have a consistent failure resolution time!

You do NOT want to place a hip-shot assumption and waist time on upper layer stuff, when the problem after 30 minutes turns out to be an unnoticed speed mismatch or a "mac-address static drop" or similar. Keep following your created processes, and build up speed and accuracy.

You will know when it is time to go!

3. Start learning the involved technology (CCIE book list)

Book Title	Pages	already read
CCIE Routing and Switching - Official Certification Guide	1080	1080
Cisco LAN Switching	960	960
Frame Relay Solutions Guide	412	412
CCIE Routing TCP/IP Vol I	936	936
CCIE Routing TCP/IP Vol II	976	976
Cisco OSPF Command and Configuration Handbook	528	-
Cisco BGP.4 Command and Configuration Handbook	400	-
Internet Routing Architectures	528	528
Troubleshooting IP Routing Protocols	912	912
MPLS Fundamentals	672	672
MPLS and VPN Architectures Vol II	504	504
Cisco QoS	768	768
End-to-End QoS Network Design: Quality of Service in LANs, WANs, and VPNs	760	162
CCIE Practical Studies Vol I	1366	-
CCIE Practical Studies Vol II	1032	-
Developing IP Multicast Networks, Vol I	592	592
Deploying IPv6 Networks	672	672
Network Security Technologies and Solutions	912	912
CCIE Version 5 Certification guide Vol 1 and Vol 2	not available	-
Pages in total	11612	10084



Calculated based on an average of 20 pages per hour. (depends on your speed)

Read through the books, use a highlighter/marker and take personal notes of the books.

Initially the book may take you around 40 hours to go through, but later you can refresh the entire book in 4 hours. I had to have physical books as I get tired too fast reading on a screen, plus sitting in the sun staring in a screen is not much fun..

Required effort rated at reading 20 pages per hour: 580 hours

All books, 6 hours of reading each day: 97 days

Reading all, 6 hours / day in month: 3.2 month

If you are really, really sure you know all technologies, skip the books and start going through the entire IOS configuration guides.

4. Repeat going through your notes

Create you own set of notes / study repetition method.

Or use my set of APP FREE ccie rs version 5 learning study flash cards to repeat what you learned so far and keep the information fresh.

http://colin.cant.ch/projects/Visio-CCIE_Lernkarten_v11.pdf

BGP Filtering Standard Access-Lists	BGP Local AS
BGP Filtering Extended Access-Lists	BGP Local AS Replace-AS/Dual-AS
BGP Regular Expressions	BGP Remove Private AS

On your lab date X

Good luck, I wish!

and may the force be with you!

