GNS3 simulation tests of BGP delays in re-routing for IP Anycast

Enrique Alfredo Ariza Zagarra

1/20/2011

Project Supervisor: Dr. Mike MacGregor

University of Alberta

Master in Internetworking

MINT 709

February 2011

Abstract

BGP and IP Anycast are two major technologies used by Internet Service Providers to support their various services such as, Hosting, FTP, e-mail, Data Storage, Printer Server. BGP is mainly used to support routing decisions in the core of Internet and permits forwarding packets among autonomous systems. On the other hand, IP Anycast is used to allow different providers to have mirror sites and to allow for load balancing and backup/restore points in the network.

This project focuses on analyzing the convergence time when using BGP as a unique routing protocol and, in combination with IP Anycast, as a tool to re-route traffic to another point in the network. For example, when an outage occurs and one of the IP Anycast servers is down, the network has to find a new route to the destination. Using BGP's process, the network must wait until the path is withdrawn, and then a new path is established to the server. It is possible that this process could take more than 180 seconds in default conditions to be able to transmit data again to the same service in the new location.

The main finding of this project was that convergence time is close to the hold-time setting for BGP. Moreover this convergence time is higher than the convergence time for IP Anycast re-routing when using OSPF as the routing protocol. Also when the default BGP timer values were changed, the same tendency in convergence time was confirmed. Thus, I conclude that anycast convergence time is related to the value of the BGP hold-timer. In addition, I found that some sample convergence times are higher than the value of the hold-timer, and this situation occurs more often when the path to the destination is longer.

Acknowledgment

I am heartily thankful to Dr. Mike MacGregor, whose guidance and support, from the initial to the final level, enabled me to develop an understanding of the subject. Also, I give special recognition to Shahnawaz Mir for allowing me to use the MINT equipment during afterhours and for helping me to connect cables in order to allow me to access the lab remotely.

Lastly, I offer my gratitude to my wife Marcela for having patience and for giving me the encouragement to finish this project. Thank you also to my little baby Valentina, she was the most beautiful gift that I could receive while working on this project.

Table of Contents

List of Figures

List of Tables

1 Border Gateway Protocol (BGP)

1.1 Definition of BGP

BGP is one of the most widely used routing protocols in the Internet today; it is used along with Autonomous System (AS) numbers to determine which networks an AS can access. BGP was started as IETF RFC-1150 in 1998 and has been improving since; the current version 4 was adopted in 1995 as RFC-4271, and has been implemented by many vendors around the world.

RFC-4271 defines BGP's primary function as: "*The primary function of a BGP speaking system is to exchange network reachability information with other BGP systems. This network reachability information includes information on the list of Autonomous Systems (ASes) that reachability information traverses. This information is sufficient for constructing a graph of AS connectivity for this reachability from which routing loops may be pruned, and, at the AS level, some policy decisions may be enforced.*

BGP-4 provides a set of mechanisms for supporting Classless Inter-Domain Routing (CIDR). These mechanisms include support for advertising a set of destinations as an IP prefix, and eliminating the concept of network "class" within BGP. BGP-4 also introduces mechanisms that allow aggregation of routes, including aggregation of AS paths."

1.2 BGP PATH Attributes

There are ten path attributes used in BGP and, there are five mandatory fields: ORIGIN, AS_PATH, AS_SET, AS_SEQUENCE and NEXT_HOP. When using "*show ip bgp*" on Cisco routers, the result is a table with the fields shown in Table 1.

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values:
	s—The table entry is suppressed.
	*—The table entry is valid.
	>—The table entry is the best entry to use for that network.
	i—The table entry was learned via an internal BGP session.
Origin codes	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values:
	i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command.
	e—Entry originated from an Exterior Gateway Protocol (EGP).
	?—Origin of the path is not clear Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	If shown, the value of the interautonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

Table 1 BGP Fields on Cisco Implementation

Figure 1 shows results from the *"show ip bgp"* command; here, Weight is a Cisco-specific attribute.

```
AS201#sh ip bgp
BGP table version is 50, local router ID is 203.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.1/32      131.0.1.2                0             0 101 ?
*  10.1.1.1/32      132.0.1.1                              0 301 401 1001 1000 400 ?
*>                  131.0.1.2                              0 101 1001 1000 400 ?
*> 10.1.1.2/32      131.0.1.2                0             0 101 ?
*  112.0.0.1/32     132.0.1.1                              0 301 401 1001 1000 100 ?
*>                  131.0.1.2                              0 101 1001 1000 100 ?
*> 113.0.0.1/32     131.0.1.2                0             0 101 ?
*  121.0.1.0/30     132.0.1.1                              0 301 401 1001 1000 100 ?
*>                  131.0.1.2                              0 101 1001 1000 100 ?
```

Figure 1 "Show IP BGP" command result

This command shows all routes available that lead to a certain network. BGP routes the packet through the best route; the best route is shown with ">" symbol. For example. there are two possible paths to 112.0.0.1/32. These are AS 301 – AS 401 – AS 1001 – AS 1000 – AS 100, and AS 101 – AS 1001 – AS 1000 - AS 100. The latter is the best path, and it will be used by BGP.

1.3 BGP Finite State Machine

BGP can be summarized using the finite state machine (FSM) shown in Figure 2. This FSM has six states: Idle; Connect; Active; OpenSent; OpenConfirm; and Established.

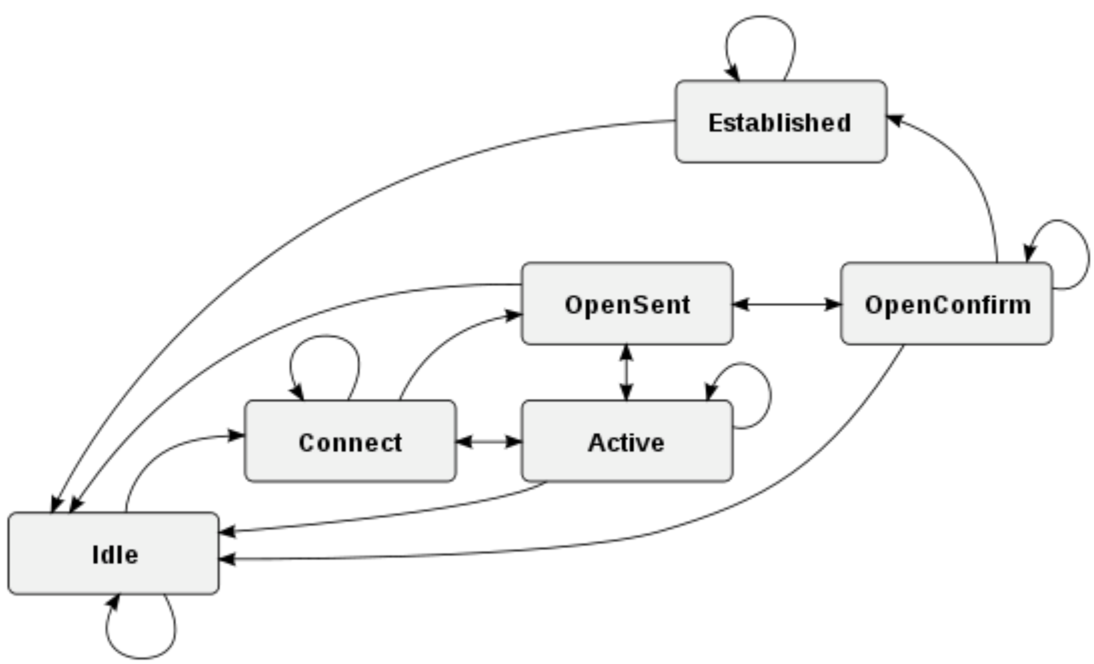

Figure 2 BGP finite state machine

The functions of the states are:

1. Idle
   - Refuses all incoming BGP connections.
   - Start event triggers the initialization of resources for the BGP process.
   - Initiates a TCP connection with its configured BGP peer.
   - Listens for a TCP connection from its peer.
   - Changes state to Connect.
   - If an error occurs at any state of the FSM process, the BGP session is terminated immediately and returned to the idle state. Some of the reasons why a router may not progress from the idle state are:
     - TCP port 179 is not open.
     - A random TCP port over 1023 is not open.
     - Peer address configured incorrectly on either router
     - AS number configured incorrectly on either router
2. Connect
   - Waits for successful TCP negotiation with peer.
   - BGP does not spend much time in this state if the TCP session has been successfully established.
   - Sends Open message to peer and changes state to OpenSent.
   - If an error occurs, BGP moves to the Active state. Possible reasons for an error are:
     - TCP port 179 is not open.
     - A random TCP port over 1023 is not open.
     - Peer address configured incorrectly on either router.
     - AS number configured incorrectly on either router.
3. Active
   - If the router is unable to establish a successful TCP session, then it moves to the Active state.
   - BGP will try again to start a TCP session with the peer and, if successful, then it will send an Open message to the peer.
   - If it is unsuccessful again, the FSM is reset to the idle state.
   - Repeated failures may result in a router cycling between Idle and Active. Possible reasons for this include:
     - TCP port 179 is not open.
     - A random TCP port over 1023 is not open.
     - BGP configuration error.
     - Network congestion.
     - Flapping network interface.
4. OpenSent
   - BGP listens for an Open message from its peer.
   - Once the Open message has been received, the router checks its validity.
   - If there is an error, it will be because one of the fields in the Open message doesn't match the value expected, e.g. BGP version mismatch, MD5 password mismatch, a different My AS is

expected. The router will then send a Notification message to the peer indicating why the error occurred.
- If there is no error, a Keepalive message is sent, various timers are set, and the state is changed to OpenConfirm.

5. OpenConfirm
   - The peer is listening for a Keepalive message from its peer.
   - If a Keepalive message is received and no timer has expired before reception of the Keepalive, BGP transitions to the Established state.
   - If a timer expires before a Keepalive message is received, or if an error condition occurs, the router transitions back to the Idle state.
6. Established
   - In this state, the peers send Update messages to exchange information about each route being advertised.
   - If there is any error in the Update message then a Notification message is sent to the peer, and BGP transitions back to the Idle state.
   - If a timer expires before a Keepalive message is received, or if an error condition occurs, the router transitions back to the Idle state.

1.4 BGP Timers

There are two BGP timers that are controllable in Cisco routers: the hold-timer and keepalive timer. These timers have default values, and the BGP process will select the smallest value between peers when establishing the connection.

Keepalive Timer: Is the number of seconds that BGP waits for a keepalive message from a peer. The default value is 60 seconds. The hold-timer must be set first, so the BGP process first negotiates the value of the hold-timer, which has a default value of 180 seconds. In order to configure the keepalive tiemr on a Cisco router under the BGP router config, use the command **"*timers bgp keepalive holdtime*"** in seconds.

Hold-timer: Is the number of seconds that BGP waits for a keepalive, update, or notification message before deciding that the connection is down. Its default value is 180 seconds.

In my project I used default values for the first set of samples, then I changed the hold-timer to 90 seconds, and left the keepalive timer at the default value of 60 seconds. During negotiation, BGP established the keepalive at 30 seconds.

1.5 BGP convergence time

Convergence is said to have occurred once all routers participating in a network have interchanged all the routing information and have achieved steady-state. On the Internet, BGP convergence never occurs because there are constant changes. However, in small networks, convergence will occur. For a single router failure, the convergence time is close to the value of the hold-timer. This is because this timer determines when the route to the failed router is released by the continuing router.

2 IP Anycast

IP Anycast allows a user to connect to the closest server in the IP Anycast cloud. This technology allows companies to optimize their services by sharing load over multiple sites. It is also used to protect them against server failures. A company like Google that has a world presence uses this technology to bring up your browser more quickly and in your language. This is due to the fact that it selects the closest server to your location. For example, someone in Mexico who types www.google.com will get redirected to www.google.com.mx, whereas someone in Canada will get redirected to www.google.ca. IP Anycast is responsible for choosing the closest available server. Figure 3 shows an IP Anycast connection. In this diagram the red dot is a user and green dots are servers with the same IP address. In this diagram, the red dot is connected to only one green dot. In case of failure, the network will re-route the traffic to the nearest surviving green dot.

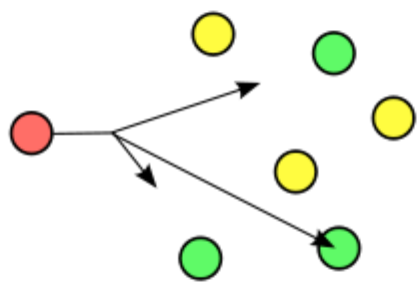

Figure 3 Anycast Diagram

IP Anycast routes traffic to a single member of a group of potential receivers that are identified by the same destination address. When building IP Anycast using Cisco routers, a Cisco Anycast RP is used to support Anycast connections. It creates MSPD peering between routers with the same IP address. Two loopback interfaces must be created in each particpating router. The first loopback interface is used to configure the unique IP address, and the second is used to create the MSPD peering. Figure 4 shows an example of the use of Anycast RP in two Cisco routers.

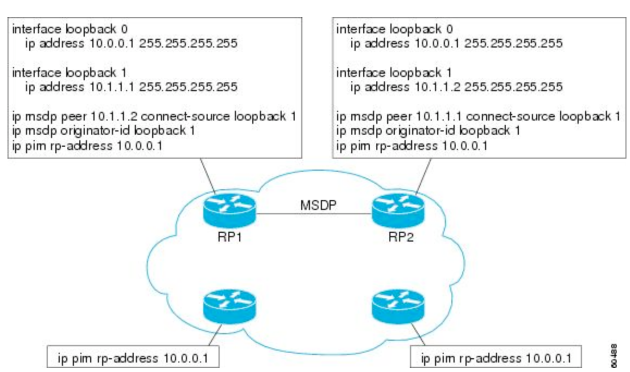

Figure 4 Example of Cisco Anycast Configuration

2.1 Benefits of IP Anycast

- ***Router and link resource reduction:*** Standard IP routing will deliver packets over the shortest path to the closest available host.
- ***Simplified configuration:*** Clients only need to configure a single Anycast address, which identifies one of a group of possible servers offering a particular service or application.
- ***Network resiliency:*** If a server in the Anycast group becomes unavailable, the network will deliver packets to the next-closest Anycast server.
- ***Load balancing:*** Anycast servers distributed over the network will effectively balance the traffic load from many clients.

3 Virtual Lab Configuration

In this chapter I describe hardware and software used to run the virtual lab for this project, as well as all the configuration details.

3.1 Hardware

I used a Sony VAIO with an Intel T4300 CPU (2.1 GHz with 2 cores), 4 GB of RAM, 320 GB hard drive, wireless LAN connection, and Gigabit LAN interface. The operating system was Ubuntu 10.04.

3.2 Software

- GNS3 Graphical Network Simulator, version 0.7.2
- WireShark Network Protocol Analyzer, version 1.4.1
- VMWare Workstation 7, Virtual Machine Manager
- IOS images for the routers used in the simulations
- Damn Small Linux (DSL) - small version of Linux using the 2.4 kernel, requires only 50MB

3.3 GNS3 Graphical Network Simulator

This open source simulator comes with all applications needed to run properly including: Dynamips, the core program that allows Cisco IOS emulation; Dynagen, a text-based front-end for Dynamips; Qemu, a generic, open source machine emulator and virtualizer; PUTTY; and WINPCAP 4.1.1. You can get these programs at http://www.gns3.net/download.

After installing the software the first step is adding the IOS to the simulted Cisco routers. In our case, images for 2600 and 3600 series were used.

Dynamips is only capable of a 1 Kpps forwarding rate, depending on the hardware used. In contrast real Cisco 2620/21 routers offer a forwarding rate of about 25 Kpps (see Figure 5).

Bundle	Router	IOS Image	WAN Interface Card	Fast Ethernet Ports	Flash (MB)	DRAM (MB)	Router Performance (kpps)
CISCO2611XM-ADSL	2611XM	IP Plus	WIC-1ADSL	2	32	128	20kpps
CISCO2621XM-ADSL	2621XM	IP Plus	WIC-1ADSL	2	32	128	30kpps
CISCO2651XM-ADSL	2651XM	IP Plus	WIC-1ADSL	2	32	128	40kpps
CISCO2611XM-SHDSL	2611XM	IP Plus	WIC-1SHDSL	2	32	128	20kpps
CISCO2621XM-SHDSL	2621XM	IP Plus	WIC-1SHDSL	2	32	128	30kpps
CISCO2651XM-SHDSL	2651XM	IP Plus	WIC-1SHDSL	2	32	128	40kpps

Figure 5 Cisco 2600 XM specs

3.3.1 Main issues with GNS3

GNS3 can overload the processor in any operating system (Windows or Ubuntu). When I started adding routers to my network, memory usage rose and the processor was running at a constant 100%. The solution is to assign an Idle PC option, which allows reducing CPU usage when routers are in standby mode. Another way to improve performance is by using unzipped Cisco images. In addition, sharing router images helps to enhance computer performance. We used eight Cisco 2600 and two Cisco 3600 images.

Notice that configuring Idle PC option for the first time results in some delay because it is necessary to wait for the calculations to be done. The options that appear with the **"*"** are the best choices. Using one of these will result in processor load dropping from 100% to 25% or less.

In Windows 7, Dynamips seems to halt several times when other programs are opened. At this time you are forced to exit from all other applications.

Figure 6 shows the reduction in processor usage when choosing the best possible Idle PC value in GNS3.

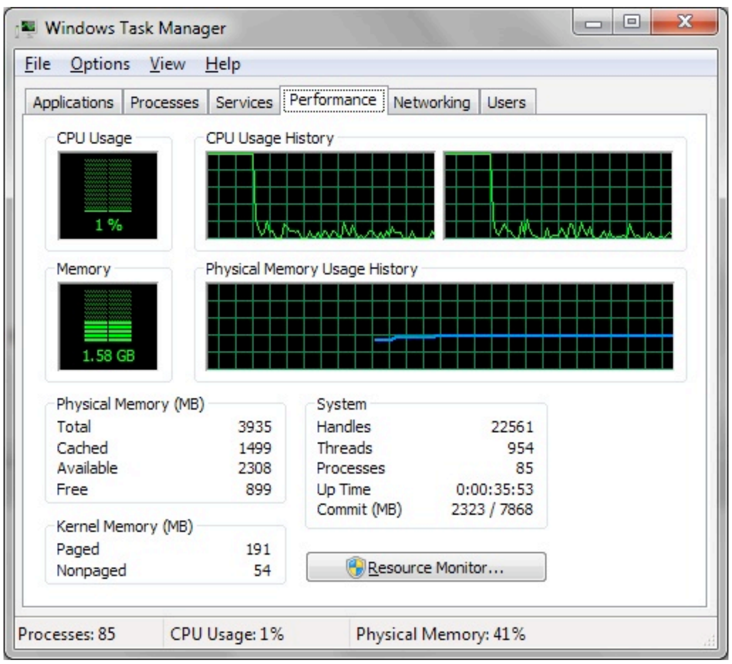

Figure 6 Idle PC in use

3.3.2 Basic configuration of GNS3

In order to begin configuring routers, run the **"wr"** command in the router, save in the GNS program, and then the current running config will appear in the selected folder, in this case C:\GNS3\newproject_configs. Then the message will appear in the GNS3 command "Exporting R "x" to the desired folder"

Once the required routers have been connected, it is possible to select the start option and run the console to be able to load the IOS, and start programming the router to your requirements.

GNS3 very easy and practical to use. It is important to remember that this only works for layer 3 and higher.

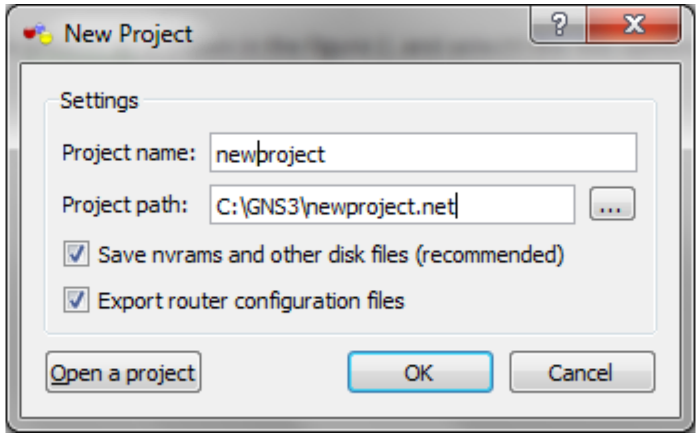

Figure 7 GNS 3 Saving Settings

3.4 WireShark

WireShark is the most widely-used network protocol analyzer on the market. I have used WireShark in many projects for other MINT courses. It allows one to analyze packets crossing a selected network interface. Easily the most important feature for our virtual lab is that it is directly embedded into GNS3. A simple right click on the interface in GNS3 will open WireShark.

3.5 VMWare Workstation 7

Workstation 7 allows you to run different operating systems using a host PC. It also comes with a virtual network manager in which we can configure all the virtual network interfaces, except for two which are used internally by Workstation 7 (VMNET 0 and 8). Figure 8 shows the virtual network editor with the interfaces that we used for the project; these interfaces were configured as host-only.

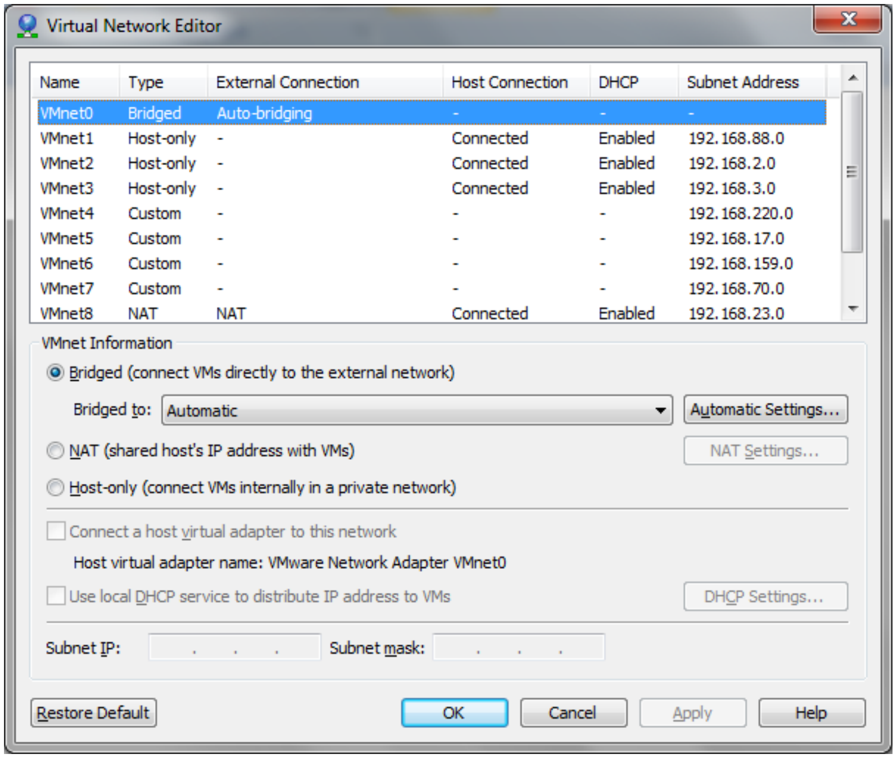

Figure 8 Virtual Network Editor

3.6 IOS Images

Cisco 2600 and Cisco 3600 images were used. The IOS version used is the mainline version 12.3 revision 26, which is the current version used in the MINT lab.

3.7 Damn Small Linux (DSL)

This version of Linux uses only 50MB to run, thus it is very portable and runs very fast in a two-core processor platform using VMWare Workstation 7 (see Figure 9).

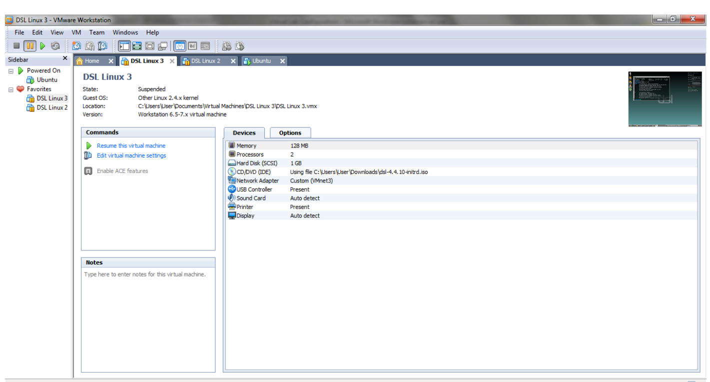

Figure 9 DSL on VMware Workstation 7

Using DSL made it possible for me to use multiple virtual workstations at the same time. I distributed them in different subnets. Processor load is very low compared with other operating systems that require more resources to operate. I created 4 DSL virtual machines within different subnets.

4 Comparison of the Real and Simulated Routers

The main purpose of this chapter is to document the performance of real and simulated routers. I implemented the same configuration on the Rack 1 routers in the MINT lab, and within GNS3 (see Figure 10).

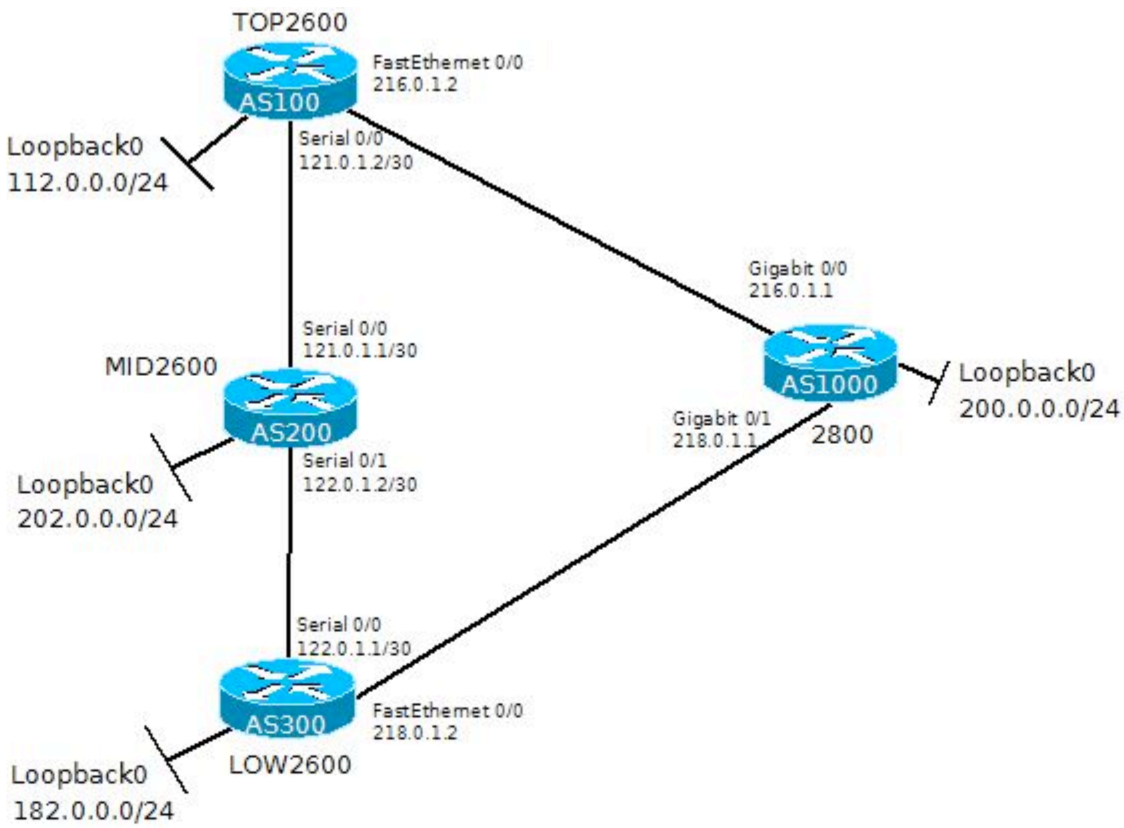

Figure 10 Rack 1 Diagram

Figure 11 shows the configuration diagram from GSN3. However, Cisco 2800 routers can't be simulated in GNS3, so we used a Cisco 3600 emulator instead.

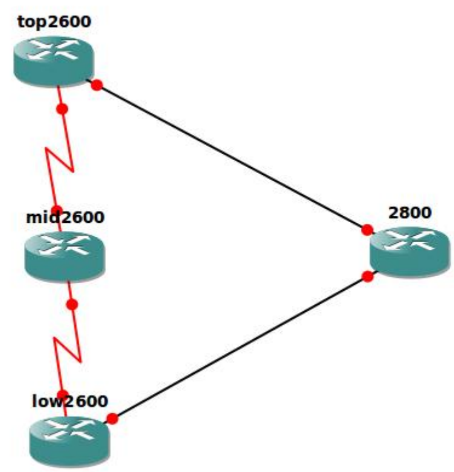

Figure 11 GNS3 Diagram

First, all BGP connections were removed in the MID2600. Next, there was a wait time for messages to be sent to all its neighbours. Times were gathered from the routers labelled LOW2600 and TOP 2600. In addition, the 2800 router was configured as a time (ntp) server to keep the routers synchronized. The results shown in Table 2 were gathered from MINT Rack 1.

REAL RACK 1 MINT LAB			TEST 1		TEST 2		TEST 3		TEST 4		TEST 5	
Router Narr	Neighbor	Event	Time Stamp	Difference	Time Stamp	Difference	Time Stamp	Difference	Time Stamp	Difference	Time Stamp	Difference
mid2600	121.0.1.2	Down	17:58:09.751	REF	18:02:20.623	REF	18:05:41.613	REF	18:10:04.636	REF	18:13:11.622	REF
mid2600	122.0.1.1	Down	17:58:09.751	REF	18:02:20.623	REF	18:05:41.613	REF	18:10:04.636	REF	18:13:11.622	REF
mid2600	121.0.1.2	Up	17:58:47.415	37.664	18:02:57.084	36.461	18:06:21.326	39.713	18:10:42.272	37.636	18:13:52.595	40.973
mid2600	122.0.1.1	Up	17:58:47.875	38.124	18:03:06.392	45.769	18:06:18.410	36.797	18:10:49.088	44.452	18:13:51.459	39.837
low2600	122.0.1.2	Down	17:58:11.258	1.507	18:02:22.534	1.911	18:05:43.539	1.926	18:10:06.529	1.893	18:13:13.528	1.906
low2600	122.0.1.2	Up	17:58:47.942	38.191	18:03:06.399	45.776	18:06:18.415	36.802	18:10:49.086	44.450	18:13:51.465	39.843
top2600	121.0.1.1	Down	17:58:10.878	1.127	18:02:22.543	1.920	18:05:43.536	1.923	18:10:06.538	1.902	18:13:12.689	1.067
top2600	121.0.1.1	Up	17:58:47.110	37.359	18:02:57.102	36.479	18:06:21.343	39.730	18:10:42.294	37.658	18:13:52.620	40.998

Table 2 MINT Rack Results

In the same configuration, GNS3 using Ubuntu achieved the results shown in Table 3.

SIMULATED GNS3 USING UBUN			TEST 1		TEST 2		TEST 3		TEST 4		TEST 5	
Router Narr	Neighbor	Event	Time Stamp	Difference	Time Stamp	Difference	Time Stamp	Difference	Time Stamp	Difference	Time Stamp	Difference
mid2600	121.0.1.2	Down	13:22:33.644	REF	13:26:22.424	REF	13:29:35.513	REF	13:31:46.521	REF	13:34:53.294	REF
mid2600	122.0.1.1	Down	13:22:33.648	REF	13:26:22.428	REF	13:29:35.513	REF	13:31:46.521	REF	13:34:53.298	REF
mid2600	121.0.1.2	Up	13:23:12.316	38.672	13:27:08.121	45.697	13:30:14.125	38.612	13:32:30.101	43.580	13:35:31.286	37.992
mid2600	122.0.1.1	Up	13:23:16.748	43.100	13:27:06.053	43.625	13:30:12.945	37.432	13:32:26.529	40.008	13:35:39.458	46.16
low2600	122.0.1.2	Down	13:22:35.548	1.900	13:26:24.180	1.752	13:29:36.156	0.643	13:31:47.181	0.660	13:34:55.152	1.854
low2600	122.0.1.2	Up	13:23:16.560	42.872	13:27:05.684	43.256	13:30:12.532	37.019	13:32:26.073	39.552	13:35:39.440	46.142
top2600	121.0.1.1	Down	13:22:35.336	1.692	13:26:23.772	1.348	13:29:36.021	0.508	13:31:47.041	0.520	13:34:55.201	1.907
top2600	121.0.1.1	Up	13:23:12.008	38.364	13:27:07.625	45.201	13:30:13.593	38.080	13:32:29.521	43.000	13:35:31.079	37.785

Table 3 GNS3 Results

In order to analyze the results, I compared the average values from the tests. These averages seem to be very close (see Figure 12).

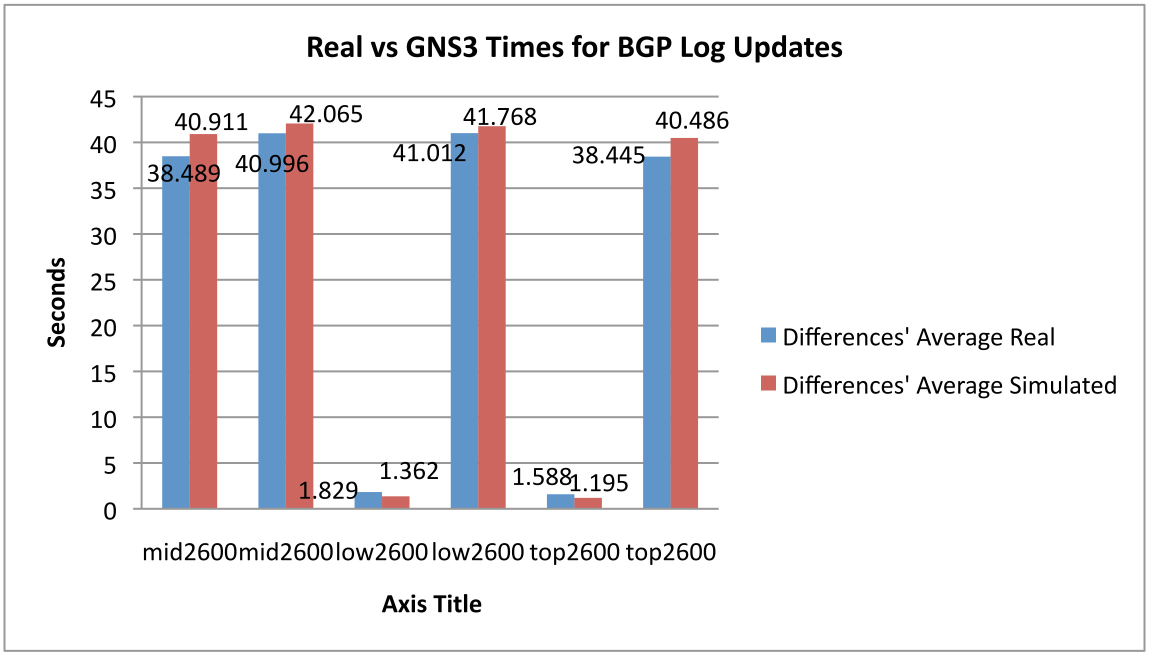

Figure 12 Comparison of BGP Log Updates

Note: It seems that, in GNS3, the interfaces in the MID2600 router are not released simultaneously, although when using the real MID2600 router, the interfaces are released at the same time.

5 Configuring the Virtual Lab

The objective of this chapter is to outline the step-by-step process of the router configuration using GNS3, and to highlight other issues that I had to solve in order to use the simulation environment.

5.1 Network Configuration

Figure 13 shows the desired configuration for this project.

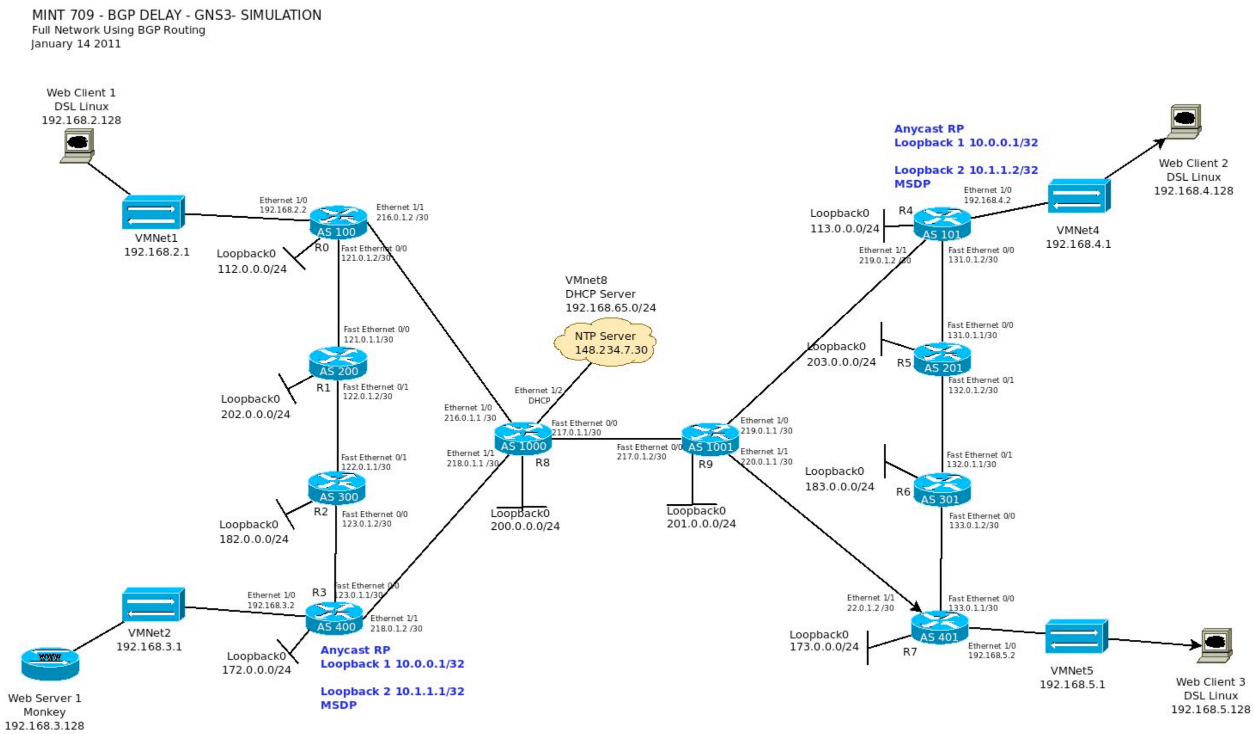

Figure 13 Network Diagram

5.2 Configuring Routers in GNS3

In this case there are ten routers (R8 and R9, using Cisco 3600 emulation and the rest using Cisco 2600 emulation). There are four virtual network clouds that allow the connection with the virtual network interfaces in Ubuntu, using the VMware network editor to set them up. All connections are Ethernet or FastEthernet.

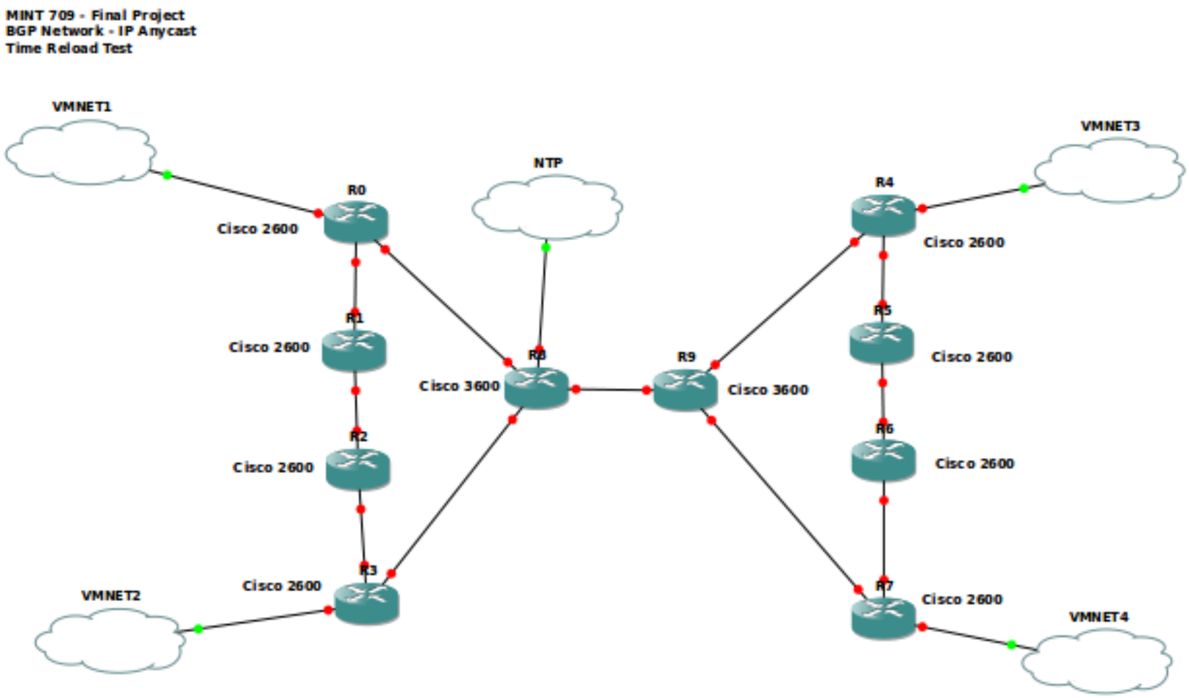

Figure 14 Network Architecture in GNS3

5.2.1 Router configurations

Table 4 gives a summary of the addresses of the interfaces in each router. The first part of the table corresponds to the Cisco 2600 series routers (R0 to R7) and the last part corresponds to the Cisco 3600 series routers (R8 and R9) .

R#	AS	ETH 1/0	ETH 1/1	FastETH 0/0	FastETH 0/1	Loopback 0	Loopback 1	Loopback 2
Cisco 2600 Series								
R0	100	192.168.2.2	216.0.1.2	121.0.1.2		112.0.0.0		
R1	200			121.0.1.1	122.0.1.2	202.0.0.0		
R2	300			123.0.1.2	122.0.1.1	182.0.0.0		
R3	400	192.168.3.2	218.0.1.2	123.0.1.1		172.0.0.0	10.0.0.1	10.1.1.1
R4	101	192.168.4.2	219.0.1.2	131.0.1.2		113.0.0.0	10.0.0.1	10.1.1.2
R5	201			131.0.1.1	132.0.1.2	203.0.0.0		
R6	301			133.0.1.2	132.0.1.1	183.0.0.0		
R7	401	192.168.5.2	220.0.1.2	133.0.1.1		173.0.0.0		
R#	**AS**	**ETH 1/0**	**ETH 1/1**	**ETH ½**	**FastETH 0/0**	**Loopback 0**	**Loopback 1**	**Loopback 2**
Cisco 3600 Series								
R8	1000	216.0.1.1	218.0.1.1	DHCP (NTP)	217.0.1.1	200.0.0.0		
R9	1001	219.0.1.1	220.0.1.1		217.0.1.2	201.0.0.0		

Table 4 Routers Interface IP addresses

In this network, there is one IP Anycast configuration involving R3 and R4 using the Anycast RP configuration in the Cisco routers. The routers are configured with only BGP routing processes; the BGP keepalive timer was set at 60 seconds and the hold-timer was set at 180 seconds.

Router R8 is also the NTP server for the entire network, using 217.0.1.1 as the default ntp server. R8 is connected to a real NTP server using VMnet8 (VMware host connection to the internet); this connection is to Ethernet 1/2 of R8. The NTP server 148.234.7.30 was used in order to get time updates.

5.2.2 BGP Routing Tables

5.2.2.1 R0 BGP "SHOW IP BGP" Result

```
AS100#sh ip bgp
BGP table version is 45, local router ID is 112.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  10.0.0.1/32      121.0.1.1                              0 200 300 400 ?
*>                  216.0.1.1                              0 1000 400 ?
*  10.1.1.1/32      121.0.1.1                              0 200 300 400 ?
*>                  216.0.1.1                              0 1000 400 ?
*> 10.1.1.2/32      216.0.1.1                              0 1000 1001 101 ?
*> 112.0.0.1/32     0.0.0.0                  0         32768 ?
*> 113.0.0.1/32     216.0.1.1                              0 1000 1001 101 ?
*  121.0.1.0/30     121.0.1.1                0             0 200 ?
*>                  0.0.0.0                  0         32768 ?
*> 122.0.1.0/30     121.0.1.1                0             0 200 ?
*  123.0.1.0/30     216.0.1.1                              0 1000 400 ?
*>                  121.0.1.1                              0 200 300 ?
*> 131.0.1.0/30     216.0.1.1                              0 1000 1001 101 ?
*> 132.0.1.0/30     216.0.1.1                              0 1000 1001 401 301 ?
*> 133.0.1.0/30     216.0.1.1                              0 1000 1001 401 ?
*  172.0.0.1/32     121.0.1.1                              0 200 300 400 ?
*>                  216.0.1.1                              0 1000 400 ?
*> 173.0.0.1/32     216.0.1.1                              0 1000 1001 401 ?
*  182.0.0.1/32     216.0.1.1                              0 1000 400 300 ?
*>                  121.0.1.1                              0 200 300 ?
*> 183.0.0.1/32     216.0.1.1                              0 1000 1001 401 301 ?
*> 192.168.2.0      0.0.0.0                  0         32768 ?
*  192.168.3.0      121.0.1.1                              0 200 300 400 ?
*>                  216.0.1.1                              0 1000 400 ?
*> 192.168.4.0      216.0.1.1                              0 1000 1001 101 ?
*> 192.168.5.0      216.0.1.1                              0 1000 1001 401 ?
*> 192.168.65.0     216.0.1.1                0             0 1000 ?
*> 200.0.0.1/32     216.0.1.1                0             0 1000 ?
*> 201.0.0.1/32     216.0.1.1                              0 1000 1001 ?
*> 202.0.0.1/32     121.0.1.1                0             0 200 ?
*> 203.0.0.1/32     216.0.1.1                              0 1000 1001 101 201 ?
*> 216.0.1.0/30     0.0.0.0                  0         32768 ?
*                   216.0.1.1                0             0 1000 ?
*> 217.0.1.0/30     216.0.1.1                0             0 1000 ?
*> 218.0.1.0/30     216.0.1.1                0             0 1000 ?
*> 219.0.1.0/30     216.0.1.1                              0 1000 1001 ?
*> 220.0.1.0/30     216.0.1.1                              0 1000 1001 ?
```

5.2.2.2 R1 BGP "SHOW IP BGP" Result

```
AS200#sh ip bgp
BGP table version is 31, local router ID is 202.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  10.0.0.1/32      121.0.1.2                              0 100 1000 400 ?
*>                  122.0.1.1                              0 300 400 ?
*  10.1.1.1/32      121.0.1.2                              0 100 1000 400 ?
*>                  122.0.1.1                              0 300 400 ?
*  10.1.1.2/32      122.0.1.1                              0 300 400 1000 1001 101 ?
*>                  121.0.1.2                              0 100 1000 1001 101 ?
*> 112.0.0.1/32     121.0.1.2                0             0 100 ?
*  113.0.0.1/32     122.0.1.1                              0 300 400 1000 1001 101 ?
*>                  121.0.1.2                              0 100 1000 1001 101 ?
*  121.0.1.0/30     121.0.1.2                0             0 100 ?
*>                  0.0.0.0                  0         32768 ?
*  122.0.1.0/30     122.0.1.1                0             0 300 ?
*>                  0.0.0.0                  0         32768 ?
*  123.0.1.0/30     121.0.1.2                              0 100 1000 400 ?
*>                  122.0.1.1                0             0 300 ?
*  131.0.1.0/30     122.0.1.1                              0 300 400 1000 1001 101 ?
*>                  121.0.1.2                              0 100 1000 1001 101 ?
*  132.0.1.0/30     122.0.1.1                              0 300 400 1000 1001 401 301 ?
*>                  121.0.1.2                              0 100 1000 1001 401 301 ?
*  133.0.1.0/30     122.0.1.1                              0 300 400 1000 1001 401 ?
*>                  121.0.1.2                              0 100 1000 1001 401 ?
*  172.0.0.1/32     121.0.1.2                              0 100 1000 400 ?
*>                  122.0.1.1                              0 300 400 ?
*  173.0.0.1/32     122.0.1.1                              0 300 400 1000 1001 401 ?
*>                  121.0.1.2                              0 100 1000 1001 401 ?
*> 182.0.0.1/32     122.0.1.1                0             0 300 ?
*  183.0.0.1/32     122.0.1.1                              0 300 400 1000 1001 401 301 ?
*>                  121.0.1.2                              0 100 1000 1001 401 301 ?
*> 192.168.2.0      121.0.1.2                0             0 100 ?
*  192.168.3.0      121.0.1.2                              0 100 1000 400 ?
*>                  122.0.1.1                              0 300 400 ?
*  192.168.4.0      122.0.1.1                              0 300 400 1000 1001 101 ?
*>                  121.0.1.2                              0 100 1000 1001 101 ?
*  192.168.5.0      122.0.1.1                              0 300 400 1000 1001 401 ?
*>                  121.0.1.2                              0 100 1000 1001 401 ?
*  192.168.65.0     122.0.1.1                              0 300 400 1000 ?
*>                  121.0.1.2                              0 100 1000 ?
*  200.0.0.1/32     122.0.1.1                              0 300 400 1000 ?
*>                  121.0.1.2                              0 100 1000 ?
*  201.0.0.1/32     122.0.1.1                              0 300 400 1000 1001 ?
*>                  121.0.1.2                              0 100 1000 1001 ?
*> 202.0.0.1/32     0.0.0.0                  0         32768 ?
*  203.0.0.1/32     122.0.1.1                              0 300 400 1000 1001 101 201 ?
*>                  121.0.1.2                              0 100 1000 1001 101 201 ?
*  216.0.1.0/30     122.0.1.1                              0 300 400 1000 ?
*>                  121.0.1.2                0             0 100 ?
*  217.0.1.0/30     122.0.1.1                              0 300 400 1000 ?
*>                  121.0.1.2                              0 100 1000 ?
*  218.0.1.0/30     122.0.1.1                              0 300 400 ?
*>                  121.0.1.2                              0 100 1000 ?
*  219.0.1.0/30     122.0.1.1                              0 300 400 1000 1001 ?
*>                  121.0.1.2                              0 100 1000 1001 ?
*  220.0.1.0/30     122.0.1.1                              0 300 400 1000 1001 ?
*>                  121.0.1.2                              0 100 1000 1001 ?
```

5.2.2.3 R2 BGP "SHOW IP BGP" Result

```
AS300#sh ip bgp
BGP table version is 75, local router ID is 182.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.1/32      123.0.1.1                0             0 400 ?
*> 10.1.1.1/32      123.0.1.1                0             0 400 ?
*> 10.1.1.2/32      123.0.1.1                              0 400 1000 1001 101 ?
*                   122.0.1.2                              0 200 100 1000 1001 101 ?
*  112.0.0.1/32     123.0.1.1                              0 400 1000 100 ?
*>                  122.0.1.2                              0 200 100 ?
*> 113.0.0.1/32     123.0.1.1                              0 400 1000 1001 101 ?
*                   122.0.1.2                              0 200 100 1000 1001 101 ?
*> 121.0.1.0/30     122.0.1.2                0             0 200 ?
*  122.0.1.0/30     122.0.1.2                0             0 200 ?
*>                  0.0.0.0                  0         32768 ?
*  123.0.1.0/30     123.0.1.1                0             0 400 ?
*>                  0.0.0.0                  0         32768 ?
*> 131.0.1.0/30     123.0.1.1                              0 400 1000 1001 101 ?
*                   122.0.1.2                              0 200 100 1000 1001 101 ?
*> 132.0.1.0/30     123.0.1.1                              0 400 1000 1001 401 301 ?
*                   122.0.1.2                              0 200 100 1000 1001 401 301 ?
*> 133.0.1.0/30     123.0.1.1                              0 400 1000 1001 401 ?
*                   122.0.1.2                              0 200 100 1000 1001 401 ?
*> 172.0.0.1/32     123.0.1.1                0             0 400 ?
*> 173.0.0.1/32     123.0.1.1                              0 400 1000 1001 401 ?
*                   122.0.1.2                              0 200 100 1000 1001 401 ?
*> 182.0.0.1/32     0.0.0.0                  0         32768 ?
*> 183.0.0.1/32     123.0.1.1                              0 400 1000 1001 401 301 ?
*                   122.0.1.2                              0 200 100 1000 1001 401 301 ?
*  192.168.2.0      123.0.1.1                              0 400 1000 100 ?
*>                  122.0.1.2                              0 200 100 ?
*> 192.168.3.0      123.0.1.1                0             0 400 ?
*> 192.168.4.0      123.0.1.1                              0 400 1000 1001 101 ?
*                   122.0.1.2                              0 200 100 1000 1001 101 ?
*> 192.168.5.0      123.0.1.1                              0 400 1000 1001 401 ?
*                   122.0.1.2                              0 200 100 1000 1001 401 ?
*> 192.168.65.0     123.0.1.1                              0 400 1000 ?
*                   122.0.1.2                              0 200 100 1000 ?
*> 200.0.0.1/32     123.0.1.1                              0 400 1000 ?
*                   122.0.1.2                              0 200 100 1000 ?
*> 201.0.0.1/32     123.0.1.1                              0 400 1000 1001 ?
*                   122.0.1.2                              0 200 100 1000 1001 ?
*> 202.0.0.1/32     122.0.1.2                0             0 200 ?
*> 203.0.0.1/32     123.0.1.1                              0 400 1000 1001 101 201 ?
*                   122.0.1.2                              0 200 100 1000 1001 101 201 ?
*  216.0.1.0/30     123.0.1.1                              0 400 1000 ?
*>                  122.0.1.2                              0 200 100 ?
*> 217.0.1.0/30     123.0.1.1                              0 400 1000 ?
*                   122.0.1.2                              0 200 100 1000 ?
*> 218.0.1.0/30     123.0.1.1                0             0 400 ?
*                   122.0.1.2                              0 200 100 1000 ?
*> 219.0.1.0/30     123.0.1.1                              0 400 1000 1001 ?
*                   122.0.1.2                              0 200 100 1000 1001 ?
*> 220.0.1.0/30     123.0.1.1                              0 400 1000 1001 ?
*                   122.0.1.2                              0 200 100 1000 1001 ?
```

5.2.2.4 R3 BGP "SHOW IP BGP" Result

```
AS400#sh ip bgp
BGP table version is 30, local router ID is 172.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.1/32      0.0.0.0                  0         32768 ?
*> 10.1.1.1/32      0.0.0.0                  0         32768 ?
*> 10.1.1.2/32      218.0.1.1                              0 1000 1001 101 ?
*  112.0.0.1/32     123.0.1.2                              0 300 200 100 ?
*>                  218.0.1.1                              0 1000 100 ?
*> 113.0.0.1/32     218.0.1.1                              0 1000 1001 101 ?
*> 121.0.1.0/30     123.0.1.2                              0 300 200 ?
*                   218.0.1.1                              0 1000 100 ?
*> 122.0.1.0/30     123.0.1.2                0             0 300 ?
*                   218.0.1.1                              0 1000 100 200 ?
*  123.0.1.0/30     123.0.1.2                0             0 300 ?
*>                  0.0.0.0                  0         32768 ?
*> 131.0.1.0/30     218.0.1.1                              0 1000 1001 101 ?
*> 132.0.1.0/30     218.0.1.1                              0 1000 1001 401 301 ?
*> 133.0.1.0/30     218.0.1.1                              0 1000 1001 401 ?
*> 172.0.0.1/32     0.0.0.0                  0         32768 ?
*> 173.0.0.1/32     218.0.1.1                              0 1000 1001 401 ?
*> 182.0.0.1/32     123.0.1.2                0             0 300 ?
*> 183.0.0.1/32     218.0.1.1                              0 1000 1001 401 301 ?
*  192.168.2.0      123.0.1.2                              0 300 200 100 ?
*>                  218.0.1.1                              0 1000 100 ?
*> 192.168.3.0      0.0.0.0                  0         32768 ?
*> 192.168.4.0      218.0.1.1                              0 1000 1001 101 ?
*> 192.168.5.0      218.0.1.1                              0 1000 1001 401 ?
*> 192.168.65.0     218.0.1.1                0             0 1000 ?
*> 200.0.0.1/32     218.0.1.1                0             0 1000 ?
*> 201.0.0.1/32     218.0.1.1                              0 1000 1001 ?
*> 202.0.0.1/32     123.0.1.2                              0 300 200 ?
*                   218.0.1.1                              0 1000 100 200 ?
*> 203.0.0.1/32     218.0.1.1                              0 1000 1001 101 201 ?
*  216.0.1.0/30     123.0.1.2                              0 300 200 100 ?
*>                  218.0.1.1                0             0 1000 ?
*> 217.0.1.0/30     218.0.1.1                0             0 1000 ?
*> 218.0.1.0/30     0.0.0.0                  0         32768 ?
*                   218.0.1.1                0             0 1000 ?
*> 219.0.1.0/30     218.0.1.1                              0 1000 1001 ?
*> 220.0.1.0/30     218.0.1.1                              0 1000 1001 ?
```

5.2.2.5 R4 BGP "SHOW IP BGP" Result

```
AS101#sh ip bgp
BGP table version is 54, local router ID is 113.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.1/32      0.0.0.0                  0         32768 ?
*> 10.1.1.1/32      219.0.1.1                              0 1001 1000 400 ?
*> 10.1.1.2/32      0.0.0.0                  0         32768 ?
*> 112.0.0.1/32     219.0.1.1                              0 1001 1000 100 ?
*> 113.0.0.1/32     0.0.0.0                  0         32768 ?
*> 121.0.1.0/30     219.0.1.1                              0 1001 1000 100 ?
*> 122.0.1.0/30     219.0.1.1                              0 1001 1000 100 200 ?
*> 123.0.1.0/30     219.0.1.1                              0 1001 1000 400 ?
*  131.0.1.0/30     131.0.1.1                0             0 201 ?
*>                  0.0.0.0                  0         32768 ?
*> 132.0.1.0/30     131.0.1.1                0             0 201 ?
*                   219.0.1.1                              0 1001 401 301 ?
*  133.0.1.0/30     131.0.1.1                              0 201 301 ?
*>                  219.0.1.1                              0 1001 401 ?
*> 172.0.0.1/32     219.0.1.1                              0 1001 1000 400 ?
*  173.0.0.1/32     131.0.1.1                              0 201 301 401 ?
*>                  219.0.1.1                              0 1001 401 ?
*> 182.0.0.1/32     219.0.1.1                              0 1001 1000 400 300 ?
*> 183.0.0.1/32     131.0.1.1                              0 201 301 ?
*                   219.0.1.1                              0 1001 401 301 ?
*> 192.168.2.0      219.0.1.1                              0 1001 1000 100 ?
*> 192.168.3.0      219.0.1.1                              0 1001 1000 400 ?
*> 192.168.4.0      0.0.0.0                  0         32768 ?
*  192.168.5.0      131.0.1.1                              0 201 301 401 ?
*>                  219.0.1.1                              0 1001 401 ?
*> 192.168.65.0     219.0.1.1                              0 1001 1000 ?
*> 200.0.0.1/32     219.0.1.1                              0 1001 1000 ?
*> 201.0.0.1/32     219.0.1.1                0             0 1001 ?
*> 202.0.0.1/32     219.0.1.1                              0 1001 1000 100 200 ?
*> 203.0.0.1/32     131.0.1.1                0             0 201 ?
*> 216.0.1.0/30     219.0.1.1                              0 1001 1000 ?
*> 217.0.1.0/30     219.0.1.1                0             0 1001 ?
*> 218.0.1.0/30     219.0.1.1                              0 1001 1000 ?
*> 219.0.1.0/30     0.0.0.0                  0         32768 ?
*                   219.0.1.1                0             0 1001 ?
*> 220.0.1.0/30     219.0.1.1                0             0 1001 ?
```

5.2.2.6 R5 BGP "SHOW IP BGP" Result

```
AS201#sh ip bgp
BGP table version is 50, local router ID is 203.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.1/32      131.0.1.2                0             0 101 ?
*  10.1.1.1/32      132.0.1.1                              0 301 401 1001 1000 400 ?
*>                  131.0.1.2                              0 101 1001 1000 400 ?
*> 10.1.1.2/32      131.0.1.2                0             0 101 ?
*  112.0.0.1/32     132.0.1.1                              0 301 401 1001 1000 100 ?
*>                  131.0.1.2                              0 101 1001 1000 100 ?
*> 113.0.0.1/32     131.0.1.2                0             0 101 ?
*  121.0.1.0/30     132.0.1.1                              0 301 401 1001 1000 100 ?
*>                  131.0.1.2                              0 101 1001 1000 100 ?
*  122.0.1.0/30     132.0.1.1                              0 301 401 1001 1000 100 200 ?
*>                  131.0.1.2                              0 101 1001 1000 100 200 ?
*  123.0.1.0/30     132.0.1.1                              0 301 401 1001 1000 400 ?
*>                  131.0.1.2                              0 101 1001 1000 400 ?
*  131.0.1.0/30     131.0.1.2                0             0 101 ?
*>                  0.0.0.0                  0         32768 ?
*  132.0.1.0/30     132.0.1.1                0             0 301 ?
*>                  0.0.0.0                  0         32768 ?
*  133.0.1.0/30     131.0.1.2                              0 101 1001 401 ?
*>                  132.0.1.1                0             0 301 ?
*  172.0.0.1/32     132.0.1.1                              0 301 401 1001 1000 400 ?
*>                  131.0.1.2                              0 101 1001 1000 400 ?
*> 173.0.0.1/32     132.0.1.1                              0 301 401 ?
*                   131.0.1.2                              0 101 1001 401 ?
*  182.0.0.1/32     132.0.1.1                              0 301 401 1001 1000 400 300 ?
*>                  131.0.1.2                              0 101 1001 1000 400 300 ?
*> 183.0.0.1/32     132.0.1.1                0             0 301 ?
*  192.168.2.0      132.0.1.1                              0 301 401 1001 1000 100 ?
*>                  131.0.1.2                              0 101 1001 1000 100 ?
*  192.168.3.0      132.0.1.1                              0 301 401 1001 1000 400 ?
*>                  131.0.1.2                              0 101 1001 1000 400 ?
*> 192.168.4.0      131.0.1.2                0             0 101 ?
*> 192.168.5.0      132.0.1.1                              0 301 401 ?
*                   131.0.1.2                              0 101 1001 401 ?
*  192.168.65.0     132.0.1.1                              0 301 401 1001 1000 ?
*>                  131.0.1.2                              0 101 1001 1000 ?
*  200.0.0.1/32     132.0.1.1                              0 301 401 1001 1000 ?
*>                  131.0.1.2                              0 101 1001 1000 ?
*  201.0.0.1/32     132.0.1.1                              0 301 401 1001 ?
*>                  131.0.1.2                              0 101 1001 ?
*  202.0.0.1/32     132.0.1.1                              0 301 401 1001 1000 100 200 ?
*>                  131.0.1.2                              0 101 1001 1000 100 200 ?
*> 203.0.0.1/32     0.0.0.0                  0         32768 ?
*  216.0.1.0/30     132.0.1.1                              0 301 401 1001 1000 ?
*>                  131.0.1.2                              0 101 1001 1000 ?
*  217.0.1.0/30     132.0.1.1                              0 301 401 1001 ?
*>                  131.0.1.2                              0 101 1001 ?
*  218.0.1.0/30     132.0.1.1                              0 301 401 1001 1000 ?
*>                  131.0.1.2                              0 101 1001 1000 ?
*  219.0.1.0/30     132.0.1.1                              0 301 401 1001 ?
*>                  131.0.1.2                0             0 101 ?
*  220.0.1.0/30     132.0.1.1                              0 301 401 ?
*>                  131.0.1.2                              0 101 1001 ?
```

5.2.2.7 R6 BGP “SHOW IP BGP” Result

```
AS301#sh ip bgp
BGP table version is 52, local router ID is 183.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.1/32      132.0.1.2                              0 201 101 ?
*                   133.0.1.1                              0 401 1001 101 ?
*  10.1.1.1/32      132.0.1.2                              0 201 101 1001 1000 400 ?
*>                  133.0.1.1                              0 401 1001 1000 400 ?
*> 10.1.1.2/32      132.0.1.2                              0 201 101 ?
*                   133.0.1.1                              0 401 1001 101 ?
*  112.0.0.1/32     132.0.1.2                              0 201 101 1001 1000 100 ?
*>                  133.0.1.1                              0 401 1001 1000 100 ?
*> 113.0.0.1/32     132.0.1.2                              0 201 101 ?
*                   133.0.1.1                              0 401 1001 101 ?
*  121.0.1.0/30     132.0.1.2                              0 201 101 1001 1000 100 ?
*>                  133.0.1.1                              0 401 1001 1000 100 ?
*  122.0.1.0/30     132.0.1.2                              0 201 101 1001 1000 100 200 ?
*>                  133.0.1.1                              0 401 1001 1000 100 200 ?
*  123.0.1.0/30     132.0.1.2                              0 201 101 1001 1000 400 ?
*>                  133.0.1.1                              0 401 1001 1000 400 ?
*  131.0.1.0/30     133.0.1.1                              0 401 1001 101 ?
*>                  132.0.1.2                0             0 201 ?
*  132.0.1.0/30     132.0.1.2                0             0 201 ?
*>                  0.0.0.0                  0         32768 ?
*  133.0.1.0/30     133.0.1.1                0             0 401 ?
*>                  0.0.0.0                  0         32768 ?
*  172.0.0.1/32     132.0.1.2                              0 201 101 1001 1000 400 ?
*>                  133.0.1.1                              0 401 1001 1000 400 ?
*> 173.0.0.1/32     133.0.1.1                0             0 401 ?
*  182.0.0.1/32     132.0.1.2                              0 201 101 1001 1000 400 300 ?
*>                  133.0.1.1                              0 401 1001 1000 400 300 ?
*> 183.0.0.1/32     0.0.0.0                  0         32768 ?
*  192.168.2.0      132.0.1.2                              0 201 101 1001 1000 100 ?
*>                  133.0.1.1                              0 401 1001 1000 100 ?
*  192.168.3.0      132.0.1.2                              0 201 101 1001 1000 400 ?
*>                  133.0.1.1                              0 401 1001 1000 400 ?
*> 192.168.4.0      132.0.1.2                              0 201 101 ?
*                   133.0.1.1                              0 401 1001 101 ?
*> 192.168.5.0      133.0.1.1                0             0 401 ?
*  192.168.65.0     132.0.1.2                              0 201 101 1001 1000 ?
*>                  133.0.1.1                              0 401 1001 1000 ?
*  200.0.0.1/32     132.0.1.2                              0 201 101 1001 1000 ?
*>                  133.0.1.1                              0 401 1001 1000 ?
*  201.0.0.1/32     132.0.1.2                              0 201 101 1001 ?
*>                  133.0.1.1                              0 401 1001 ?
*  202.0.0.1/32     132.0.1.2                              0 201 101 1001 1000 100 200 ?
*>                  133.0.1.1                              0 401 1001 1000 100 200 ?
*> 203.0.0.1/32     132.0.1.2                0             0 201 ?
*  216.0.1.0/30     132.0.1.2                              0 201 101 1001 1000 ?
*>                  133.0.1.1                              0 401 1001 1000 ?
*  217.0.1.0/30     132.0.1.2                              0 201 101 1001 ?
*>                  133.0.1.1                              0 401 1001 ?
*  218.0.1.0/30     132.0.1.2                              0 201 101 1001 1000 ?
*>                  133.0.1.1                              0 401 1001 1000 ?
*  219.0.1.0/30     132.0.1.2                              0 201 101 ?
*>                  133.0.1.1                              0 401 1001 ?
*  220.0.1.0/30     132.0.1.2                              0 201 101 1001 ?
*>                  133.0.1.1                0             0 401 ?
```

5.2.2.8 R7 BGP "SHOW IP BGP" Result

```
AS401#sh ip bgp
BGP table version is 52, local router ID is 173.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  10.0.0.1/32      133.0.1.2                              0 301 201 101 ?
*>                  220.0.1.1                              0 1001 101 ?
*> 10.1.1.1/32      220.0.1.1                              0 1001 1000 400 ?
*  10.1.1.2/32      133.0.1.2                              0 301 201 101 ?
*>                  220.0.1.1                              0 1001 101 ?
*> 112.0.0.1/32     220.0.1.1                              0 1001 1000 100 ?
*  113.0.0.1/32     133.0.1.2                              0 301 201 101 ?
*>                  220.0.1.1                              0 1001 101 ?
*> 121.0.1.0/30     220.0.1.1                              0 1001 1000 100 ?
*> 122.0.1.0/30     220.0.1.1                              0 1001 1000 100 200 ?
*> 123.0.1.0/30     220.0.1.1                              0 1001 1000 400 ?
*  131.0.1.0/30     133.0.1.2                              0 301 201 ?
*>                  220.0.1.1                              0 1001 101 ?
*> 132.0.1.0/30     133.0.1.2                0             0 301 ?
*  133.0.1.0/30     133.0.1.2                0             0 301 ?
*>                  0.0.0.0                  0         32768 ?
*> 172.0.0.1/32     220.0.1.1                              0 1001 1000 400 ?
*> 173.0.0.1/32     0.0.0.0                  0         32768 ?
*> 182.0.0.1/32     220.0.1.1                              0 1001 1000 400 300 ?
*> 183.0.0.1/32     133.0.1.2                0             0 301 ?
*> 192.168.2.0      220.0.1.1                              0 1001 1000 100 ?
*> 192.168.3.0      220.0.1.1                              0 1001 1000 400 ?
*  192.168.4.0      133.0.1.2                              0 301 201 101 ?
*>                  220.0.1.1                              0 1001 101 ?
*> 192.168.5.0      0.0.0.0                  0         32768 ?
*> 192.168.65.0     220.0.1.1                              0 1001 1000 ?
*> 200.0.0.1/32     220.0.1.1                              0 1001 1000 ?
*> 201.0.0.1/32     220.0.1.1                0             0 1001 ?
*> 202.0.0.1/32     220.0.1.1                              0 1001 1000 100 200 ?
*> 203.0.0.1/32     133.0.1.2                              0 301 201 ?
*                   220.0.1.1                              0 1001 101 201 ?
*> 216.0.1.0/30     220.0.1.1                              0 1001 1000 ?
*> 217.0.1.0/30     220.0.1.1                0             0 1001 ?
*> 218.0.1.0/30     220.0.1.1                              0 1001 1000 ?
*> 219.0.1.0/30     220.0.1.1                0             0 1001 ?
*> 220.0.1.0/30     0.0.0.0                  0         32768 ?
*                   220.0.1.1                0             0 1001 ?
```

5.2.2.9 R8 BGP "SHOW IP BGP" Result

```
AS1000#sh ip bgp
BGP table version is 51, local router ID is 200.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.0.0.1/32      218.0.1.2                0             0 400 ?
*                   217.0.1.2                              0 1001 101 ?
*> 10.1.1.1/32      218.0.1.2                0             0 400 ?
*> 10.1.1.2/32      217.0.1.2                              0 1001 101 ?
*> 112.0.0.1/32     216.0.1.2                0             0 100 ?
*> 113.0.0.1/32     217.0.1.2                              0 1001 101 ?
*  121.0.1.0/30     218.0.1.2                              0 400 300 200 ?
*>                  216.0.1.2                0             0 100 ?
*  122.0.1.0/30     218.0.1.2                              0 400 300 ?
*>                  216.0.1.2                              0 100 200 ?
*> 123.0.1.0/30     218.0.1.2                0             0 400 ?
*                   216.0.1.2                              0 100 200 300 ?
*> 131.0.1.0/30     217.0.1.2                              0 1001 101 ?
*> 132.0.1.0/30     217.0.1.2                              0 1001 401 301 ?
*> 133.0.1.0/30     217.0.1.2                              0 1001 401 ?
*> 172.0.0.1/32     218.0.1.2                0             0 400 ?
*> 173.0.0.1/32     217.0.1.2                              0 1001 401 ?
*> 182.0.0.1/32     218.0.1.2                              0 400 300 ?
*                   216.0.1.2                              0 100 200 300 ?
*> 183.0.0.1/32     217.0.1.2                              0 1001 401 301 ?
*> 192.168.2.0      216.0.1.2                0             0 100 ?
*> 192.168.3.0      218.0.1.2                0             0 400 ?
*> 192.168.4.0      217.0.1.2                              0 1001 101 ?
*> 192.168.5.0      217.0.1.2                              0 1001 401 ?
*> 192.168.65.0     0.0.0.0                  0         32768 ?
*> 200.0.0.1/32     0.0.0.0                  0         32768 ?
*> 201.0.0.1/32     217.0.1.2                0             0 1001 ?
*  202.0.0.1/32     218.0.1.2                              0 400 300 200 ?
*>                  216.0.1.2                              0 100 200 ?
*> 203.0.0.1/32     217.0.1.2                              0 1001 101 201 ?
*  216.0.1.0/30     216.0.1.2                0             0 100 ?
*>                  0.0.0.0                  0         32768 ?
*  217.0.1.0/30     217.0.1.2                0             0 1001 ?
*>                  0.0.0.0                  0         32768 ?
*  218.0.1.0/30     218.0.1.2                0             0 400 ?
*>                  0.0.0.0                  0         32768 ?
*> 219.0.1.0/30     217.0.1.2                0             0 1001 ?
*> 220.0.1.0/30     217.0.1.2                0             0 1001 ?
```

5.2.2.10 R9 BGP “SHOW IP BGP” Result

```
AS1001#sh ip bgp
BGP table version is 42, local router ID is 201.0.0.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  10.0.0.1/32      217.0.1.1                              0 1000 400 ?
*>                  219.0.1.2                0             0 101 ?
*> 10.1.1.1/32      217.0.1.1                              0 1000 400 ?
*> 10.1.1.2/32      219.0.1.2                0             0 101 ?
*> 112.0.0.1/32     217.0.1.1                              0 1000 100 ?
*> 113.0.0.1/32     219.0.1.2                0             0 101 ?
*> 121.0.1.0/30     217.0.1.1                              0 1000 100 ?
*> 122.0.1.0/30     217.0.1.1                              0 1000 100 200 ?
*> 123.0.1.0/30     217.0.1.1                              0 1000 400 ?
*> 131.0.1.0/30     219.0.1.2                0             0 101 ?
*  132.0.1.0/30     219.0.1.2                              0 101 201 ?
*>                  220.0.1.2                              0 401 301 ?
*> 133.0.1.0/30     220.0.1.2                0             0 401 ?
*> 172.0.0.1/32     217.0.1.1                              0 1000 400 ?
*> 173.0.0.1/32     220.0.1.2                0             0 401 ?
*> 182.0.0.1/32     217.0.1.1                              0 1000 400 300 ?
*  183.0.0.1/32     219.0.1.2                              0 101 201 301 ?
*>                  220.0.1.2                              0 401 301 ?
*> 192.168.2.0      217.0.1.1                              0 1000 100 ?
*> 192.168.3.0      217.0.1.1                              0 1000 400 ?
*> 192.168.4.0      219.0.1.2                0             0 101 ?
*> 192.168.5.0      220.0.1.2                0             0 401 ?
*> 192.168.65.0     217.0.1.1                0             0 1000 ?
*> 200.0.0.1/32     217.0.1.1                0             0 1000 ?
*> 201.0.0.1/32     0.0.0.0                  0         32768 ?
*> 202.0.0.1/32     217.0.1.1                              0 1000 100 200 ?
*  203.0.0.1/32     220.0.1.2                              0 401 301 201 ?
*>                  219.0.1.2                              0 101 201 ?
*> 216.0.1.0/30     217.0.1.1                0             0 1000 ?
*  217.0.1.0/30     217.0.1.1                0             0 1000 ?
*>                  0.0.0.0                  0         32768 ?
*> 218.0.1.0/30     217.0.1.1                0             0 1000 ?
*  219.0.1.0/30     219.0.1.2                0             0 101 ?
*>                  0.0.0.0                  0         32768 ?
*  220.0.1.0/30     220.0.1.2                0             0 401 ?
*>                  0.0.0.0                  0         32768 ?
```

5.2.3 Empty Network Behaviour

Using WireShark, It was possible to verify that the BGP updates were occurring every 60 seconds, and that the holding time was 180 seconds. Figure 15 shows the BGP open message from router R9. and Figure 16 shows R8 keepalive message every 60 seconds to R9.

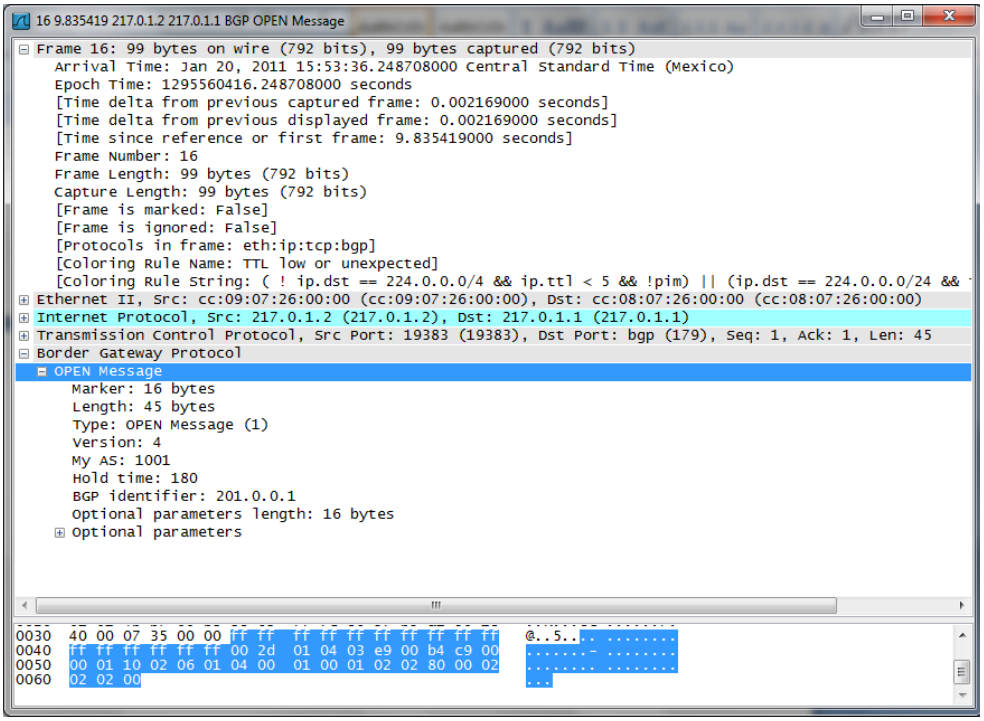

Figure 15 R9 Open Message to R8

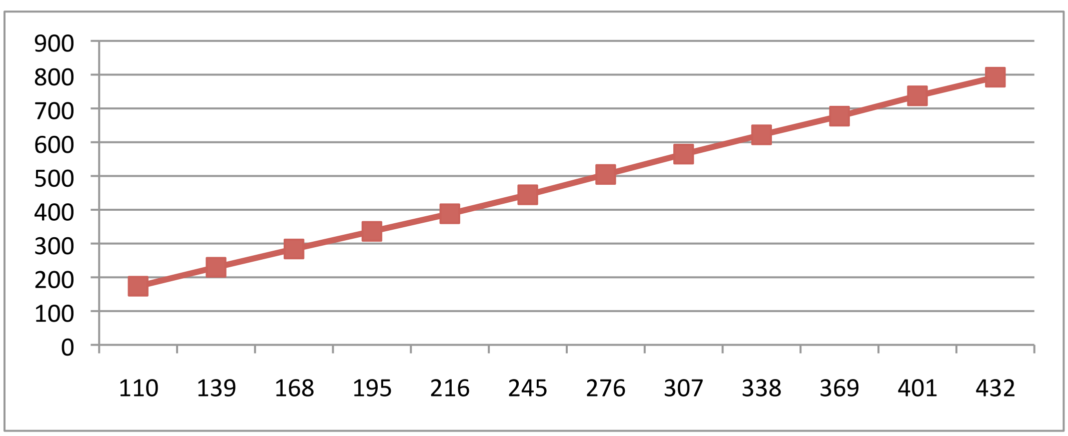

Figure 16 Keepalive Time R8 to R9

5.2.4 IP Anycast Configuration

In order to configure two routers with the same IP address it is necessary to program two loopback addresses with same IP address (10.0.0.1). This was implemented in the following router configurations for R3 and R4.

R3 anycast configuration commands

Loopback1
IP address 10.0.0.1 255.255.255.255
Loopback2
IP address 10.1.1.1 255.255.255.255
IP PIM rp-address 10.0.0.1
IP MSDP peer 10.1.1.2 connect-source Loopback2
IP MSDP originator-id Loopback2

R4 anycast configuration commands

Loopback1
IP address 10.0.0.1 255.255.255.255
Loopback2
IP address 10.1.1.2 255.255.255.255
IP PIM rp-address 10.0.0.1
IP MSDP peer 10.1.1.1 connect-source Loopback2
IP MSDP originator-id Loopback2

After this configuration is applied, the routers closest to R3 will choose it to get to 10.0.0.1, while the routers closest to R4 will choose it instead.

When using the PIM and MSDP features of Cisco IOS, PIM is used to forward traffic between multicast domains, while MSDP sets up a multicast peering between R3 and R4. Figure 17 shows one of the traces from R8 to R9 which captures some of the messages interchanged in this peering. Figure 18 shows the times of the keepalive messages for MSDP.

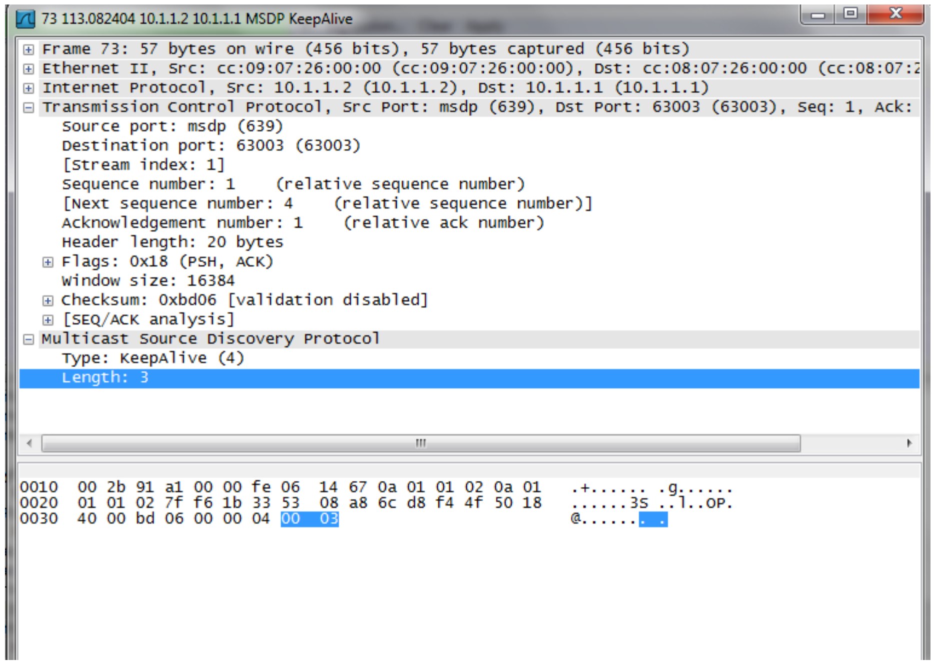

Figure 17 MSDP Keepalive Message

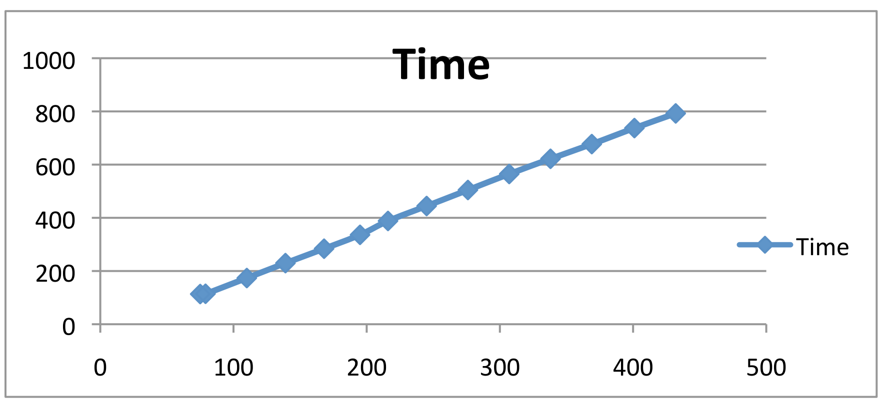

Figure 18 Keepalive MSDP (60 seconds default)

5.3 IP Anycast testing

The following procedure was used to test IP Anycast:

1. Start the network and wait until all routers reach steady-state.
2. From any router, run a traceroute command to 10.0.0.1, which is the IP Anycast address.
3. From the same router, ping 10.0.0.1, using the following ping options: Protocol = IP, Target IP Address= 10.0.0.1, Repeat count=1000, Datagram size = 100, Timeout in seconds = 1, extended commands =n, and Sweep Range of Sizes = n
4. Turn off the router through which the connection to 10.0.0.1 has been established.

5. Wait for the ping command to finish.
6. Take the values obtained from the measurement.
7. Run traceroute to 10.0.0.1 again to confirm the new path has been established.

The first test was run using R0 and turning R3 down a few seconds after ping command to 10.0.0.1 started. Figure 19 shows the test result. In this case, the new path to 10.0.0.1 is established after 136 seconds. The first traceroute to 10.0.0.1 shows the path to be R0-R8-R3 (216.0.1.1 and 218.0.1.2). The second traceroute shows the path to be R0-R8-R9-R4 (216.0.1.1, 217.0.1.2 and 219.0.1.2).

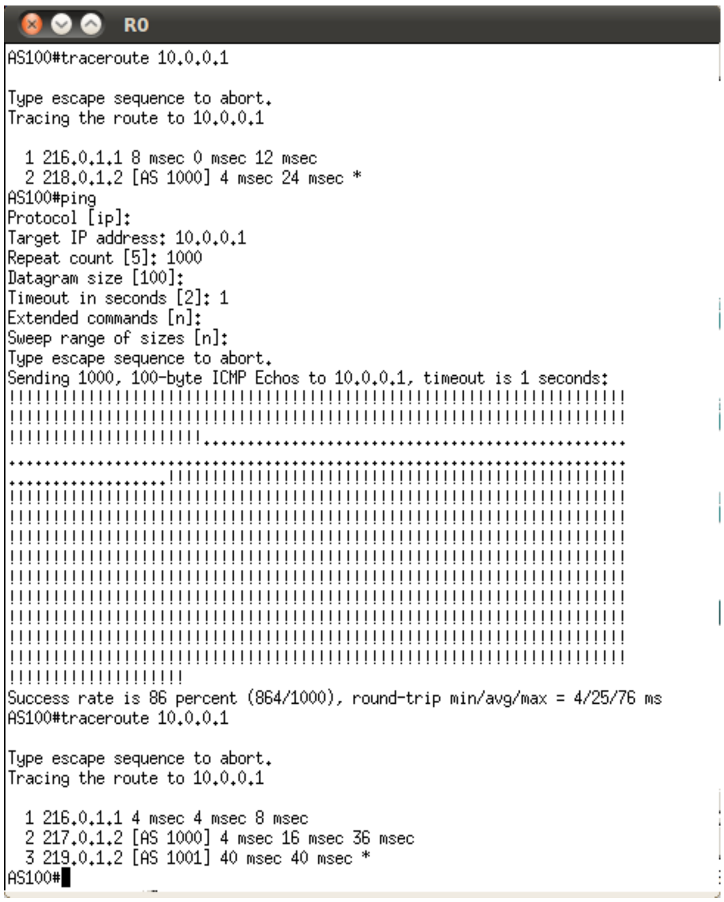

```
R0
AS100#traceroute 10.0.0.1

Type escape sequence to abort.
Tracing the route to 10.0.0.1

  1 216.0.1.1 8 msec 0 msec 12 msec
  2 218.0.1.2 [AS 1000] 4 msec 24 msec *
AS100#ping
Protocol [ip]:
Target IP address: 10.0.0.1
Repeat count [5]: 1000
Datagram size [100]:
Timeout in seconds [2]: 1
Extended commands [n]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 1000, 100-byte ICMP Echos to 10.0.0.1, timeout is 1 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!................................................
......................................................................
..................!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!
Success rate is 86 percent (864/1000), round-trip min/avg/max = 4/25/76 ms
AS100#traceroute 10.0.0.1

Type escape sequence to abort.
Tracing the route to 10.0.0.1

  1 216.0.1.1 4 msec 4 msec 8 msec
  2 217.0.1.2 [AS 1000] 4 msec 16 msec 36 msec
  3 219.0.1.2 [AS 1001] 40 msec 40 msec *
AS100#
```

Figure 19 Ping and Traceroute from R0 to 10.0.0.1

From each of four routers - R0, R1, R7 and R6 - ten sample runs were conducted to see the behaviour of the convergence time for IP Anycast using only BGP as routing protocol. Table 5 shows the summary of the results for the tests from R0.

Test #	Convergence Time (s)
1	181
2	136
3	180
4	169
5	172
6	178
7	145
8	162
9	143
10	160
Average	162.6

Table 5 R0 Convergence Time Samples

The second test was run using R1 and turning down R3 a few seconds after the ping to 10.0.0.1 was started. Figure 20 shows the results. In this case, 165 seconds were needed to establish the new path to 10.0.0.1. The first traceroute to 10.0.0.1 shows the path to be R1-R2-R3 (122.0.1.1 and 123.0.1.2). The second traceroute shows the path to be R1-R0-R8-R9-R4 (121.0.1.2, 216.0.1.1, 217.0.1.2 and 219.0.1.2). Table 6 shows the summary of the results for this set of tests

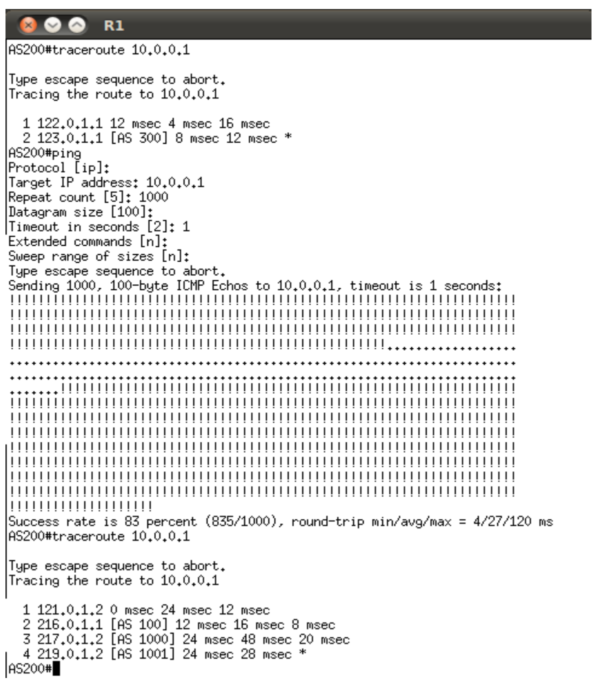

Figure 20 Ping and traceroute from R1 to 10.0.0.1

Test #	Convergence Time (s)
1	165
2	179
3	181
4	152
5	182
6	180
7	182
8	169
9	176
10	150
Average	**171.6**

Table 6 R1 Convergence Time Samples

The third test was run using R7 and turning down R4 a few seconds after the ping to 10.0.0.1 was started. Figure 21 shows the first sample result. In this case, 153 seconds were needed to establish the connection to the new path to 10.0.0.1. The first traceroute to 10.0.0.1 shows the path to be R7-R9-R4 (220.0.1.1 and 219.0.1.2). The second traceroute shows the path to be R7-R9-R8-R3 (220.0.1.1, 217.0.1.1 and 218.0.1.2). Table 7 shows the results for this set.

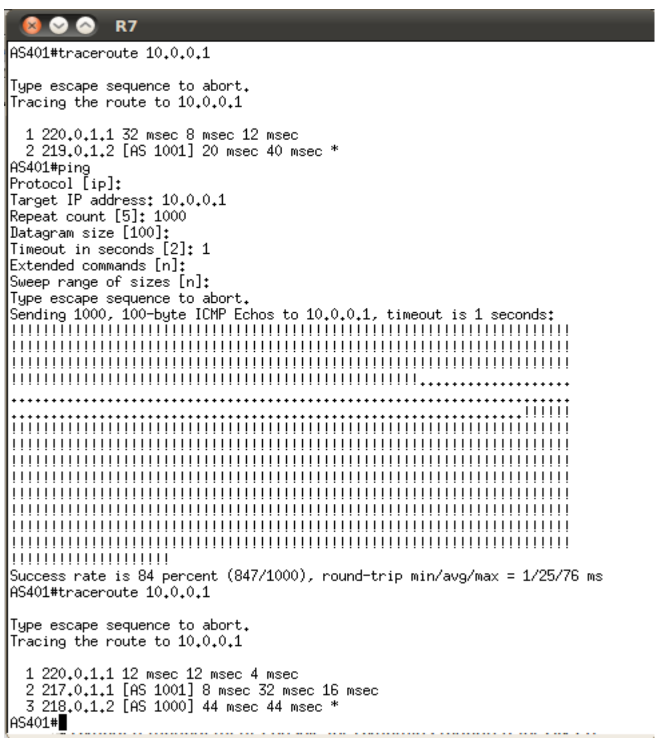

Figure 21 Ping and traceroute from R7 to 10.0.0.1

Test #	Convergence Time (s)
1	153
2	136
3	139
4	163
5	146
6	133
7	159
8	160
9	150
10	171
Average	**151**

Table 7 R7 Convergence Time Samples

This fourth test was run using R6, and turning down R4 a few seconds after the ping to 10.0.0.1 has started. Figure 22 shows the first sample performed. In this case, 153 seconds were needed to re-establish the connection to the new path to 10.0.0.1. The first traceroute to 10.0.0.1 shows the path to be is R6-R5-R4 (132.0.1.2 and 131.0.1.2). The second traceroute shows the path to be R6-R7-R9-R8-R3 (133.0.1.1, 220.0.1.1, 217.0.1.1 and 218.0.1.2). Table 8 shows the summary of results for this set.

Figure 22 Ping and traceroute from R6 to 10.0.0.1

Test #	Convergence Time (s)
1	176
2	162
3	183
4	186
5	156
6	142
7	176
8	164
9	171
10	161
Average	167.7

Table 8 R6 Convergence Time Samples

Convergence time in all cases is around 10% less than the BGP hold-timer. This trend was maintained when we changed the hold-timer to 90 seconds in all routers (see Table 9). The only additional observation is that the average time from R1 and R6 seems to be closer to the hold-timer. Also, the average for R6 is higher than the hold-timer, at an average value of 95.6 seconds.

IP ANYCAST BGP HOLD TIME: 90 s

Test #	R3 Down Convergence Time from R0 (90 s)	R3 Down Convergence Time from R1 (90 s)	R4 Down Convergence Time from R7 (90 s)	Convergence Time from R6 (90 s)
1	82	73	87	100
2	78	90	85	91
3	78	66	83	93
4	86	88	88	95
5	81	88	80	109
6	89	101	74	95
7	69	98	92	95
8	70	111	81	102
9	90	89	86	88
10	79	82	74	88
Average	**80.2**	**88.6**	**83**	**95.6**
STDEV	**7.084**	**13.100**	**5.869**	**6.535**

Table 9 Convergence times when hold-time is 90 seconds

In addition to the previous tests, some data collection was done using Wireshark . The results are very similar. For example, a trace from R2 to R3 was analyzed in order to get the time when no more messages were received from R3. There were seven TCP re-transmissions since the last successful keepalive message from R3 (123.0.0.1) so the next message from R2 to R3 was 60 seconds after that. TCP then tried to send the keepalive message to R3 but there was no successful response after

44.396365 seconds. In summary the time elapsed since the last successful transmission from R3 was a total of 104.866046 seconds.

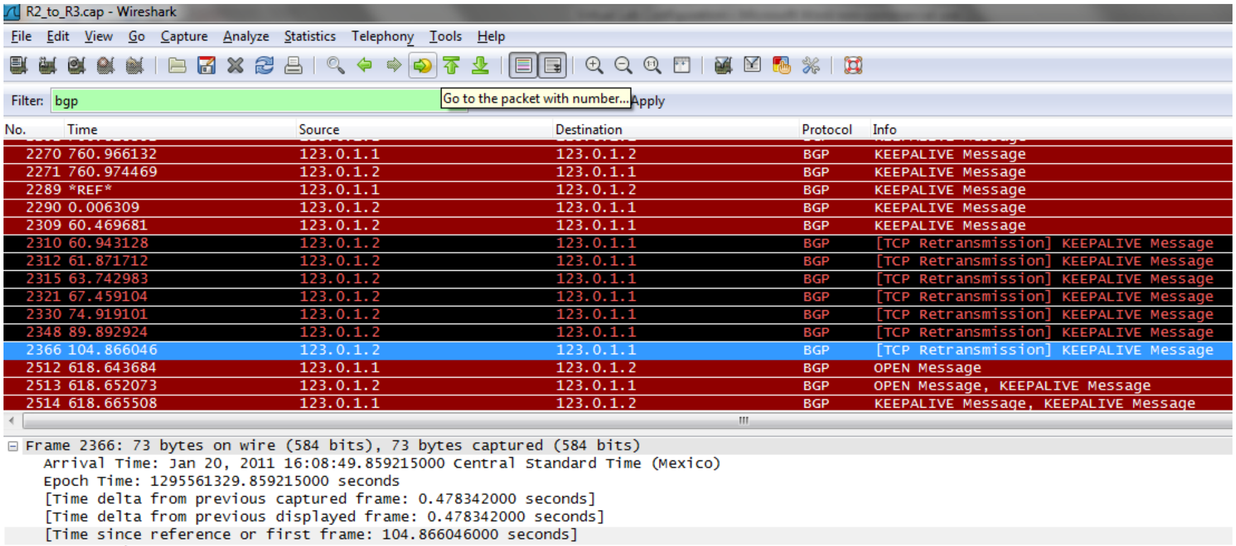

Figure 23 WireShark R2 to R3 trace when R3 goes down

In the R4-to-R9 trace we find an update message where the routes from R3 were withdrawn. After that the pings from R1 to 10.0.0.1 were successful (see Figure 24).

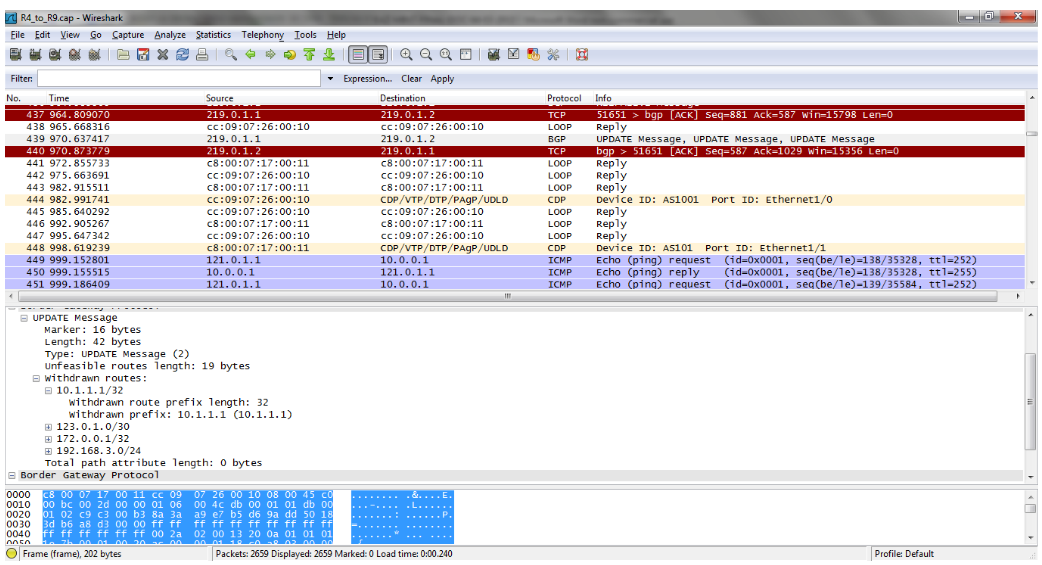

Figure 24 WireShark R4 to R9 trace when R3 goes down

Traces were taken from every link in our network. These traces were compared against the time reference. Table 10 shows all times when R3 went down, and when it was restored. When R3 was down the configuration of the network required more time than when R3 was re-started. Some of this delay is due to the holding time (180 seconds). Table 10 shows the results for all network connections.

Traces Analysis from WireShark
Initial Time when R3 when Down: 15:18:03

Test 1

Link	Message (WireShark)	Time(hh:mm:ss.ssssss)	Difference(s)
R2 to R3	Last Keepalive	15:18:50.763418	47.763418
R2 to R3	Last TCP Retransmission	15:19:21.156953	78.156953
R2 to R3	First ARP for 123.0.0.1	15:19:31.395004	88.395004
R1 to R2	First BGP Update after R3 Down	15:20:52.898145	169.898145
R8 to R9	First BGP Update after R3 Down	15:20:39.940560	169.940560
R0 to R8	First BGP Update after R3 Down	15:20:39.940645	169.940645
R9 to R7	First BGP Update after R3 Down	15:20:39.947234	169.947234
R4 to R9	First BGP Update after R3 Down	15:20:39.947366	169.947366
R4 to R5	First BGP Update after R3 Down	15:20:39.954117	169.954117
R6 to R7	First BGP Update after R3 Down	15:20:39.954213	169.954213
R0 to R1	First BGP Update after R3 Down	15:20:39.965055	169.965055
R5 to R6	First BGP Update after R3 Down	15:20:39.97689	169.976890

Table 10 Time differences when R3 is down

As noted in Table 10, the last update for the routers takes around 170 seconds to be received. This is close to 10% less than the hold-timer.

6 Conclusion

Border Gateway Protocol is one of the most widely used routing protocols in the Internet today. In small environments it is possible for BGP to converge, but this never occurs in the global Internet due to the number and frequency of changes. A small network of ten routers was used In the tests documented in this report. The main finding is that when a failure occurs in one router, re-routing to an alternative IP Anycast address converges after almost the same time as the hold-timer value. In our first series of tests, the hold-timer value had the default value of 180 seconds, and the average convergence time was 163.2 seconds. In our second series of tests, the hold-timer was set to 90 seconds, and the average convergence time was 86.9 seconds.

A second observation is that convergence time increases as the number of hops increases. For example, convergence times measured from a router one hop further away from the failed router are higher. In our experiments, recovering the path from R0 required 162.6 seconds, while recovering the path from R1 took 171.6 seconds. We can also see the same trend in the tests with R7 and R6. The R7 average is 16% less and the R6 is 7% less.

In this project, I used the graphic network simulator GNS3, rather than real Cisco routers. GNS3 was very useful in testing a small, quiet network. The simulation environment worked better using Ubuntu rather than Windows 7, which wasn't able to keep GNS3 running for more than 10 minutes. Previous work (Ref. 10) used a BGP network simulator called SSFNET. With this tool, researchers were able to simulate and analyze the performance of large BGP networks. Our conclusions using GNS3 agree with the findings from similar tests with SSFNET.

The performance of the GNS3 simulation is very close to that of a real network using the same set of four routers in a quiet traffic situation. The conclusion is that GNS3 is a good tool to build your own home lab without using real routers, but it only works with layer 3 or higher. Further, GNS3 allows you to practice in order to get your Cisco Certifications. It also allows you to connect your home network or internet to your lab environment.

A possible future project could use the tool SSFNET to simulate larger networks and to analyze the BGP convergence time for different types of networks.

7 Appendix 1: Router Configurations

R0

```
!
Version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS100
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 112.0.0.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 121.0.1.2 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
```

```
 no ip address
 shutdown
 duplex auto
 speed auto
!
interface Ethernet1/0
 ip address 192.168.2.2 255.255.255.0
 half-duplex
!
interface Ethernet1/1
 ip address 216.0.1.2 255.255.255.252
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 100
 no synchronization
 bgp log-neighbor-changes
 network 112.0.0.0 mask 255.255.255.0
 redistribute connected
 neighbor 121.0.1.1 remote-as 200
 neighbor 216.0.1.1 remote-as 1000
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
line con 0
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

R1

```
!
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS200
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 202.0.0.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 121.0.1.1 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
 ip address 122.0.1.2 255.255.255.252
 duplex auto
 speed auto
```

```
!
interface Ethernet1/0
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/1
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 200
 no synchronization
 bgp log-neighbor-changes
 network 202.0.0.0
 redistribute connected
 neighbor 121.0.1.2 remote-as 100
 neighbor 122.0.1.1 remote-as 300
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
line con 0
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

R2

```
! Last configuration change at 14:58:22 MST Wed Jan 19 2011
! NVRAM config last updated at 14:58:30 MST Wed Jan 19 2011
!
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS300
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 182.0.0.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 123.0.1.2 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
```

```
 ip address 122.0.1.1 255.255.255.252
 duplex auto
 speed auto
!
interface Ethernet1/0
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/1
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 300
 no synchronization
 bgp log-neighbor-changes
 network 182.0.0.0 mask 255.255.255.0
 redistribute connected
 neighbor 122.0.1.2 remote-as 200
 neighbor 123.0.1.1 remote-as 400
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
line con 0
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

R3

```
!
! No configuration change since last restart
! NVRAM config last updated at 15:11:48 MST Wed Jan 19 2011
!
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS400
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 172.0.0.1 255.255.255.255
!
interface Loopback1
 ip address 10.0.0.1 255.255.255.255
!
interface Loopback2
 ip address 10.1.1.1 255.255.255.255
!
```

```
interface FastEthernet0/0
 ip address 123.0.1.1 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
 no ip address
 shutdown
 duplex auto
 speed auto
!
interface Ethernet1/0
 ip address 192.168.3.2 255.255.255.0
 half-duplex
!
interface Ethernet1/1
 ip address 218.0.1.2 255.255.255.252
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 400
 no synchronization
 bgp log-neighbor-changes
 network 172.0.0.0 mask 255.255.255.0
 redistribute connected
 neighbor 123.0.1.2 remote-as 300
 neighbor 218.0.1.1 remote-as 1000
 no auto-summary
!
ip http server
ip classless
!
ip pim rp-address 10.0.0.1
ip msdp peer 10.1.1.2 connect-source Loopback2
ip msdp originator-id Loopback2
!
!
!
!
!
!
!
line con 0
```

```
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

R4

```
!
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS101
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 113.0.0.1 255.255.255.255
!
interface Loopback1
```

```
 ip address 10.0.0.1 255.255.255.255
!
interface Loopback2
 ip address 10.1.1.2 255.255.255.255
!
interface FastEthernet0/0
 ip address 131.0.1.2 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
 no ip address
 shutdown
 duplex auto
 speed auto
!
interface Ethernet1/0
 ip address 192.168.4.2 255.255.255.0
 half-duplex
!
interface Ethernet1/1
 ip address 219.0.1.2 255.255.255.252
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 101
 no synchronization
 bgp log-neighbor-changes
 network 113.0.0.0 mask 255.255.255.0
 redistribute connected
 neighbor 131.0.1.1 remote-as 201
 neighbor 219.0.1.1 remote-as 1001
 no auto-summary
!
ip http server
ip classless
!
ip pim rp-address 10.0.0.1
ip msdp peer 10.1.1.1 connect-source Loopback2
ip msdp originator-id Loopback2
!
!
!
```

```
!
!
!
!
line con 0
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

R5

```
!
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS201
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
```

```
interface Loopback0
 ip address 203.0.0.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 131.0.1.1 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
 ip address 132.0.1.2 255.255.255.252
 duplex auto
 speed auto
!
interface Ethernet1/0
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/1
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 201
 no synchronization
 bgp log-neighbor-changes
 network 203.0.0.0
 redistribute connected
 neighbor 131.0.1.2 remote-as 101
 neighbor 132.0.1.1 remote-as 301
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
```

```
line con 0
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

R6

```
!
! No configuration change since last restart
! NVRAM config last updated at 15:11:53 MST Wed Jan 19 2011
!
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS301
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
```

```
interface Loopback0
 ip address 183.0.0.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 133.0.1.2 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
 ip address 132.0.1.1 255.255.255.252
 duplex auto
 speed auto
!
interface Ethernet1/0
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/1
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 301
 no synchronization
 bgp log-neighbor-changes
 network 183.0.0.0 mask 255.255.255.0
 redistribute connected
 neighbor 132.0.1.2 remote-as 201
 neighbor 133.0.1.1 remote-as 401
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
```

```
line con 0
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

R7

```
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS401
!
boot-start-marker
boot-end-marker
!
!
memory-size iomem 15
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 173.0.0.1 255.255.255.255
!
interface FastEthernet0/0
```

```
 ip address 133.0.1.1 255.255.255.252
 duplex auto
 speed auto
!
interface FastEthernet0/1
 no ip address
 shutdown
 duplex auto
 speed auto
!
interface Ethernet1/0
 ip address 192.168.5.2 255.255.255.0
 half-duplex
!
interface Ethernet1/1
 ip address 220.0.1.2 255.255.255.252
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 401
 no synchronization
 bgp log-neighbor-changes
 network 173.0.0.0 mask 255.255.255.0
 redistribute connected
 neighbor 133.0.1.2 remote-as 301
 neighbor 220.0.1.1 remote-as 1001
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
line con 0
line aux 0
line vty 0 4
 login
!
```

```
ntp server 217.0.1.1
!
end
```

R8

```
! Last configuration change at 14:57:42 MST Wed Jan 19 2011
! NVRAM config last updated at 14:58:35 MST Wed Jan 19 2011
!
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS1000
!
boot-start-marker
boot-end-marker
!
!
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 200.0.0.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 217.0.1.1 255.255.255.252
 duplex auto
```

```
 speed auto
!
interface Ethernet1/0
 ip address 216.0.1.1 255.255.255.252
 half-duplex
!
interface Ethernet1/1
 ip address 218.0.1.1 255.255.255.252
 half-duplex
!
interface Ethernet1/2
 ip address dhcp
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 1000
 no synchronization
 bgp log-neighbor-changes
 network 251.0.0.0 mask 255.255.255.0
 redistribute connected
 neighbor 216.0.1.2 remote-as 100
 neighbor 217.0.1.2 remote-as 1001
 neighbor 218.0.1.2 remote-as 400
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
!
!
line con 0
line aux 0
line vty 0 4
 login
!
ntp clock-period 17179872
ntp server 148.234.7.30
!
end
```

R9

```
version 12.3
service timestamps debug datetime msec
service timestamps log datetime msec
no service password-encryption
!
hostname AS1001
!
boot-start-marker
boot-end-marker
!
!
clock timezone MST -7
no aaa new-model
ip subnet-zero
!
!
!
ip cef
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 201.0.0.1 255.255.255.255
!
interface FastEthernet0/0
 ip address 217.0.1.2 255.255.255.252
 duplex auto
 speed auto
!
interface Ethernet1/0
 ip address 219.0.1.1 255.255.255.252
 half-duplex
!
interface Ethernet1/1
 ip address 220.0.1.1 255.255.255.252
```

```
 half-duplex
!
interface Ethernet1/2
 no ip address
 shutdown
 half-duplex
!
interface Ethernet1/3
 no ip address
 shutdown
 half-duplex
!
router bgp 1001
 no synchronization
 bgp log-neighbor-changes
 network 201.0.0.0
 redistribute connected
 neighbor 217.0.1.1 remote-as 1000
 neighbor 219.0.1.2 remote-as 101
 neighbor 220.0.1.2 remote-as 401
 no auto-summary
!
ip http server
ip classless
!
!
!
!
!
!
!
!
!
!
line con 0
line aux 0
line vty 0 4
 login
!
ntp server 217.0.1.1
!
end
```

8 References

1. William R. Parkhurst, *Cisco BGP-4 Command and Configuration Handbook* (Cisco Press), May 7 2001, 205.
2. Scott Empson and Hans Roth, CCNP *ROUTE Portable Command Guide,* (Cisco Press), March 12, 2010.
3. Jeff Doyle and Jennifer DeHaven Carroll, *Routing TCP/IP, Volume II (CCIE Professional Development)* (Cisco Press), April 21, 2001.
4. www.gns3.org, GNS3 Web Site, and repository of examples using this simulator
5. www.vmware.com
6. www.cisco.com
7. www.ubuntu.com
8. www.wireshark.com
9. http://www.rfc-archive.org/getrfc.php?rfc=4271&tag=A-Border-Gateway-Protocol-4-%28BGP-4%29
10. Griffin, TG, Premore, BJ. *An experimental analysis of BGP convergence time.* 9th International Conference on Network Protocols (ICNP 2001), IEEE COMPUTER SOCIETY, 2001
11. www.cisco.com/en/US/docs/ios/solutions_docs/ip_multicast/White_papers/anycast.html
12. Metz, Chris. *IP Anycast Point-to-(Any) Point Communication. IEEE INTERNET COMPUTING. MARCH - APRIL 2002.*
13. Han, Kevin*. MINT 709 Capstone Project Report, Anycast network solutions for DataGardens' Syntropy implementation and their performance investigations, Nov 23 2009.*
14. Patel, Ashishkumar. Implement virtual Data sharing application on WAN using DDNS and its performance investigation. MINT Capstone Project