Revision 0515

NFV 350-WBT: AT&T Vyatta 5600 vRouter BGP ESSENTIALS

The training materials and other content provided herein for assistance in training on the Vyatta vRouter may have references to Brocade as the Vyatta vRouter was formerly a Brocade product prior to AT&T's acquisition of Vyatta. Brocade remains a separate company and is not affiliated to AT&T."

T&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under pplicable nondisclosure terms and is not for general distribution

Welcome to the AT&T Vyatta 5600 vRouter BGP Essentials course.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Legal Disclaimer

All or some of the products detailed in this presentation may still be under development and certain specifications may be subject to change.

Nothing in this presentation shall be deemed to create a warranty of any kind.

©2017 AT&T Intellectual Property. All rights reserved. AT&T, the Globe logo, Vyatta and VPlane are registered trademarks of AT&T Intellectual Property. All other marks are the property of their respective owners.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Please take a moment to read our legal disclaimer.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T Vyatta 5600 vRouter
BGP Basics

3

Let’s being the first module in the BGP Essentials course, BGP Basics.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Objectives

After completing this module, you will be able to

- Describe BGP functionality and applications
- Discuss 5600 vRouter BGP support for IPv6
- Configure eBGP connectivity to a single ISP
- Configure basic iBGP connectivity
- Verify BGP operations
- Troubleshoot common problems

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

After completing this module, you will be able to:

- Describe BGP functionality and applications
- Discuss 5600 vRouter BGP support for IPv6
- Configure eBGP connectivity to a single ISP
- Configure basic iBGP connectivity
- Verify BGP operations
- Troubleshoot common problems

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

BGP Overview

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

We begin with an overview of BGP, what it is and what it does.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

What is BGP-4?

Routing protocol between Autonomous Systems (AS)

- AS: a group of networks under one administration

Two implementations

- External BGP (eBGP) – exchanges BGP routes between two ASs
- Internal BGP (iBGP) – exchanges BGP routes within a single AS

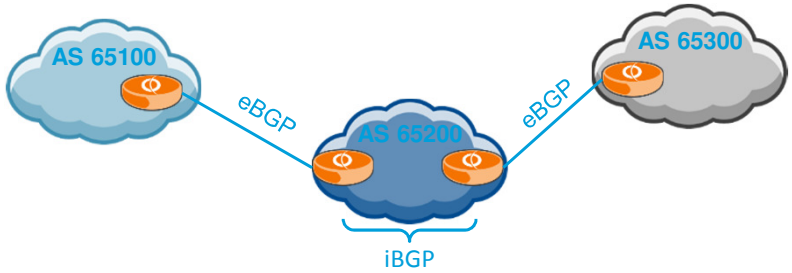

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

BGP stands for Border Gateway Protocol. The current version is BGP-4 defined in RFC 4271.
BGP is the standard for exchanging routing information between Autonomous Systems.
An autonomous system, or AS, is a group of networks under a single administration. An AS may be a corporation, or an ISP.
BGP is deployed in one of two ways, each having its own requirements.
External BGP, or eBGP, runs between routers in two different autonomous systems.
Internal BGP, or iBGP, runs between routers within the same autonomous system.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

How Does BGP Differ From an IGP?

- Connection-oriented
  - BGP peers connect via TCP session
- Change-based – no periodic updates
  - Full BGP table is ≈ 500,000 entries (Sept. 2014)
- No single metric
  - Set of route attributes used to determine routing preference
  - Attributes can be modified to adjust routing behavior

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

So, what makes BGP different from an interior gateway protocol like OSPF or RIP?
First of all, BGP is connection-oriented
Unlike interior protocols, which rely on broadcasts or multicasts for routing updates, BGP peers maintain an active TCP session. Routing updates are sent directly between peers over this connection.
Secondly, BGP updates are only sent when there is a change in the network. Even OSPF sends out “sanity check” link state advertisements every thirty minutes, but BGP peers assume their routing information is stable as long as the peer remains connected via the TCP connection.
Remember, BGP routing keeps entries for every network in the Internet. Change-based updates are essential in a routing environment where a device may be maintaining over five hundred thousand entries.
Finally, unlike RIP’s hop count or OSPF’s path cost, BGP does not have a single metric it uses to determine a best path to a particular network.
Instead, each route has several associated route attributes.
Administrators can use these attributes to determine preferred routes, or modify the attributes advertised to BGP peers to affect upstream routing choices. We discuss BGP route attributes in detail in the *AT&T Vyatta 5600 vRouter BGP Policies* module of this course.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

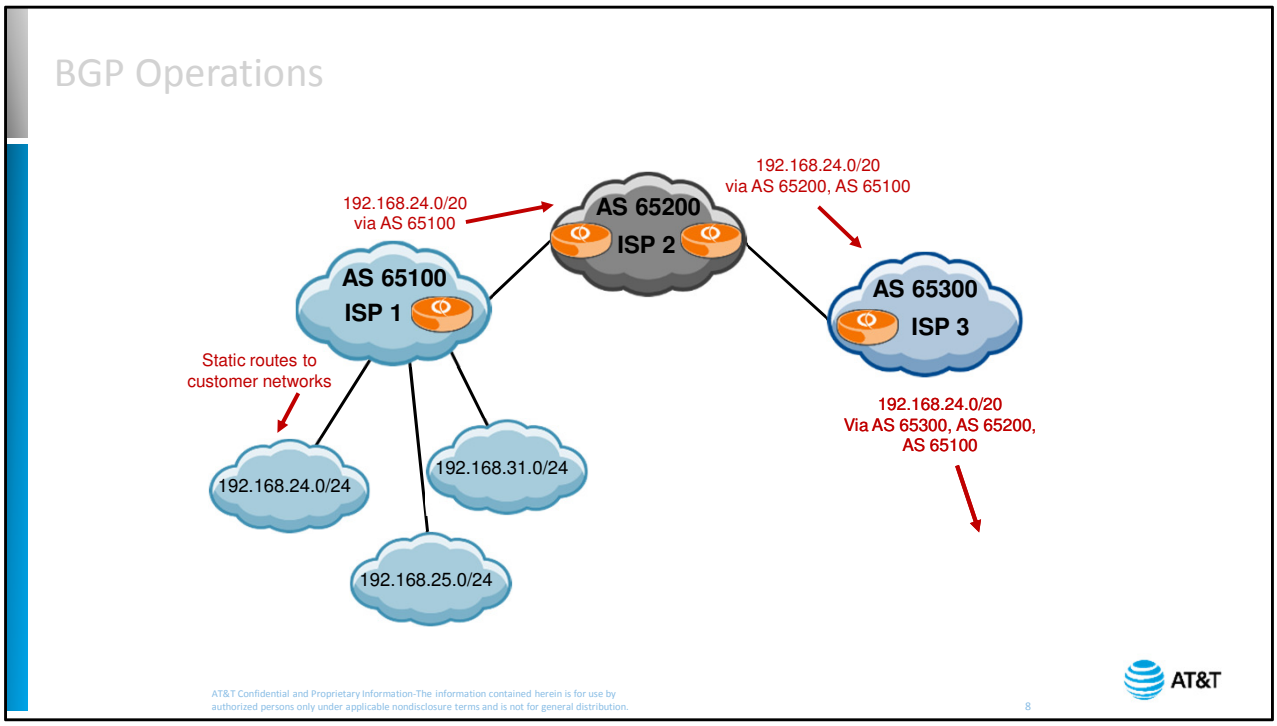

Let's take a high-level look at how BGP operates across the Internet.
ISP 1 has several customers. It has assigned each customer a block of IP addresses, and has configured static routing within its domain to those individual networks. Likewise, the end customers have a static default route directing traffic to their upstream ISP.
Upstream, ISP 1 aggregates those networks into a summary route and advertises the summary using BGP. The advertisement includes the AS used to reach the advertised network. This AS path information is one of the BGP route attributes.
ISP 2 advertises the route via BGP to ISP 3, adding its AS to the path.
ISP 3 advertises the route via BGP to its peers, adding its own AS to the path. And so on.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

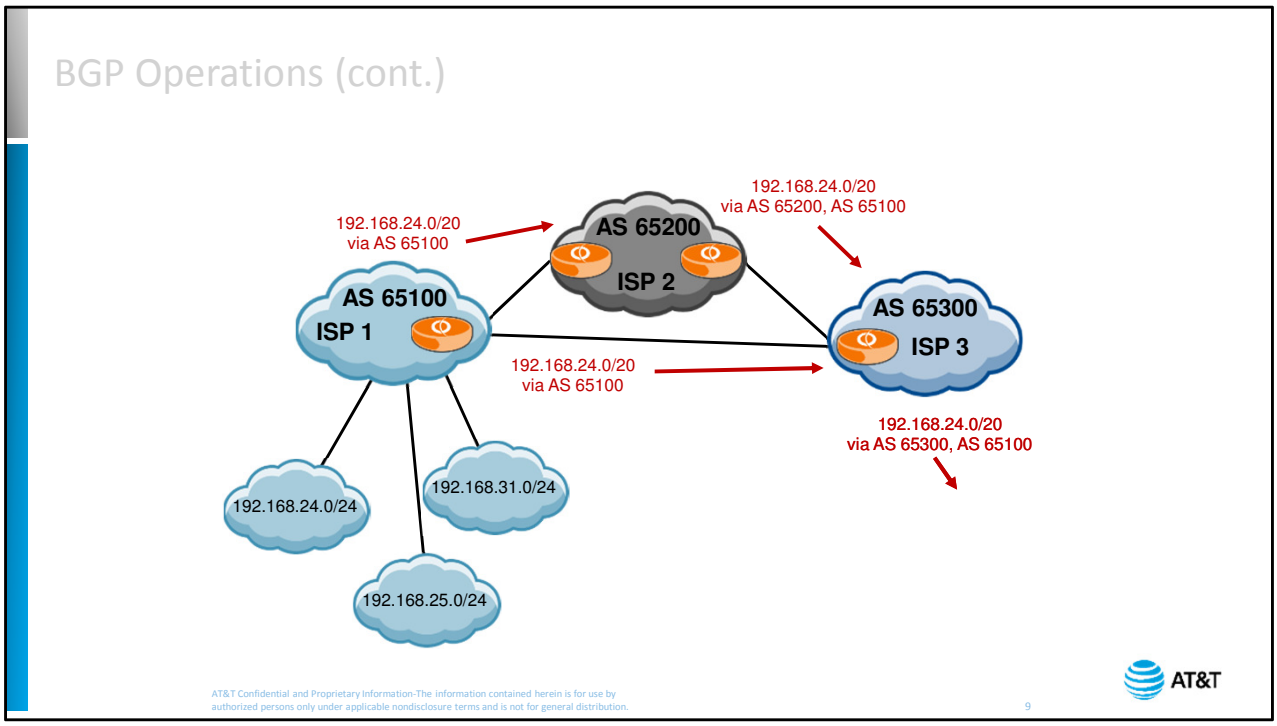

In cases where a device receives multiple paths to the same destination network, It will by default choose the shortest AS path for its upstream advertisements. In this case, the shortest path to network 192.168.24.0 is through ISP 3 directly to ISP 1. ISP 3 will continue to receive the information from ISP 2, but will not forward that path information unless the direct connection to ISP 1 fails.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When Should I Use BGP?

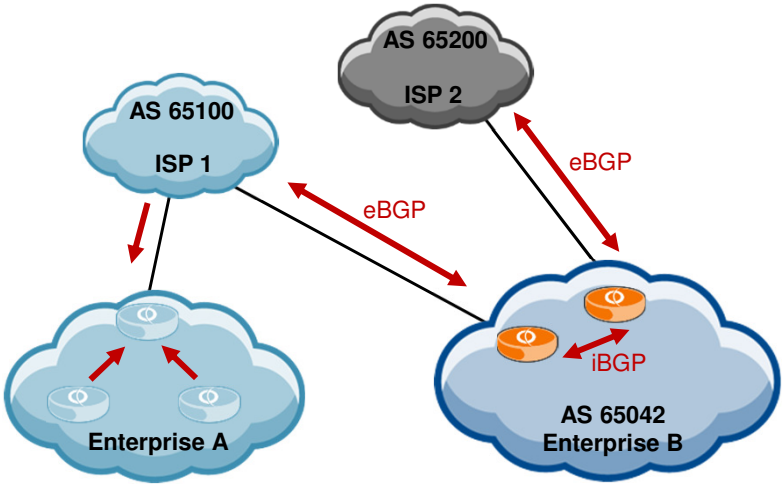

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

As you can see, BGP is essential to ISPs. But what if you are an enterprise? Do you need BGP?

Generally, if you are working in an enterprise network with a single ISP connection, you do not *need* BGP.

You can configure a static default route to your ISP and propagate that information throughout your network. All traffic for destinations outside your network will be directed to your ISP connection.

Likewise, your ISP has no need for you to run BGP and will likely have a static route in place to reach your network.

If your enterprise is multi-homed to different ISPs, you and the ISPs will benefit from running BGP. Traffic destined for your network will use the most efficient path to reach you, and you can use BGP policy controls to intelligently select the exit point from your network. You will also benefit from BGP being a dynamic protocol, allowing you to automatically route around link and node failures.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

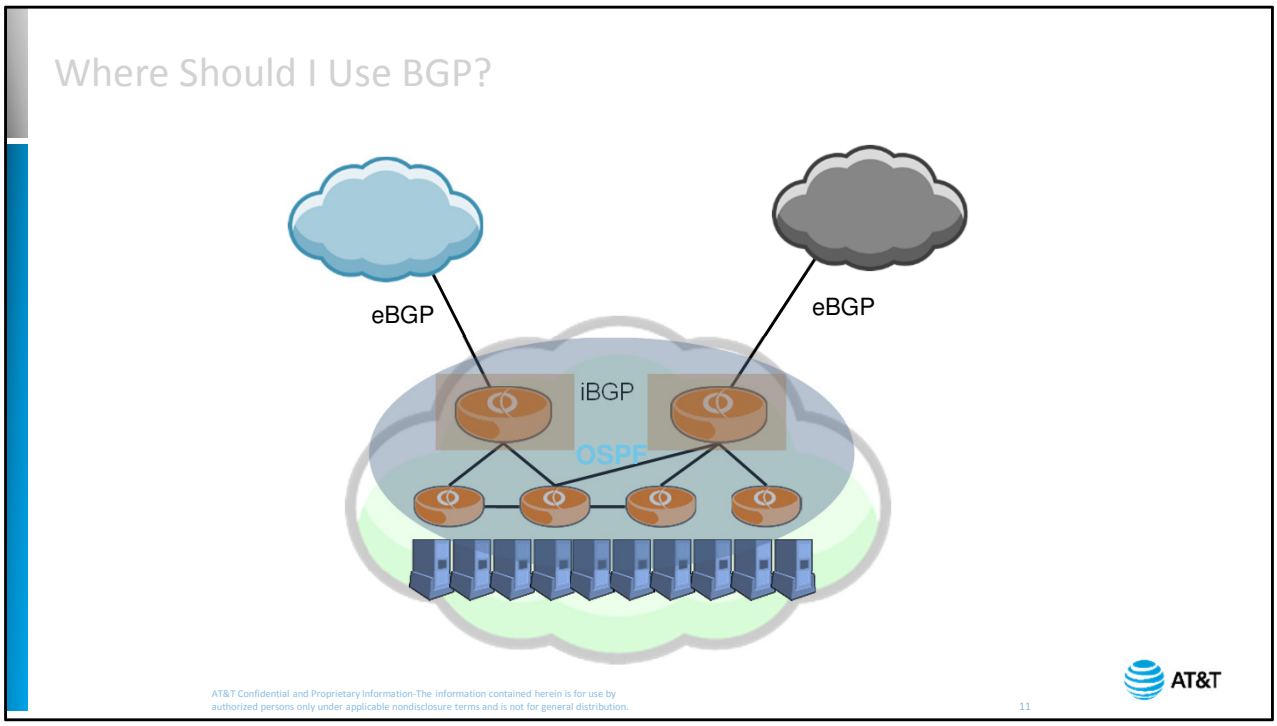

As we said before, BGP is NOT an interior routing protocol. You will only run BGP on those routers that need detailed information about the best paths to networks outside your AS. In this network, only the highlighted routers are running BGP,
eBGP to the ISPs and iBGP between them within the enterprise. Note that the iBGP peers are not directly connected; they do not need to be as long as they have IP reachability. All the routers within the network are running OSPF to advertise internal routing information, including the routes the iBGP peers will use to reach each other.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

MP-BGP for IPv6

Multiprotocol BGP (MP-BGP) supports the distribution of IPv6 routing information using extensions, also known as BGP4+

- Uses address families to support sharing of additional prefix information such as IPv6 unicast and multicast, and VPNv4 prefixes
- Address families define what types of prefixes BGP is carrying, which by default are IPv4 unicast
- To support IPv6 prefixes, the IPv6 unicast address family must be configured

Address families defined per neighbor (IPv4, IPv6, or both)

```
edit protocols bgp as-number
set neighbor ip-address address-family ipv4-unicast
set neighbor ip-address address-family ipv6-unicast
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's take a minute to discuss the 5600 and BGP support for IPv6.

BGP-4 is designed to distribute routing information for the IPv4 protocol. To support the distribution of IPv6 routing information BGP uses Multiprotocol BGP (MP-BGP) extensions, also known as BGP4+. MP-BGP uses address families to support the sharing of additional prefix information such as IPv6 unicast and multicast, and VPNv4 prefixes. Address families define what types of prefixes BGP is carrying, which by default are IPv4 unicast. To support IPv6 prefixes, the IPv6 unicast address family must be configured.

To define what type of prefixes a peering session between neighbors will carry, configure the neighbor statement with the `address-family ipv4-unicast` command, the `address-family ipv6-unicast` command, or both. We will look at the full BGP configuration in the next section of this module.

In this course we discuss the basics for configuring IPv6 for BGP on the 5600 vRouter, we do not discuss MP-BGP in detail. For the full details for MP-BGP, please refer to RFC 2545.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

IPv4 and IPv6 Peers

5600 vRouter supports the following:

- Peering between IPv4 peers
- Peering between IPv6 peers
- IPv4 routes carried over IPv4 *or* IPv6 peer session
- IPv6 routes carried over IPv4 *or* IPv6 peer session
- IPv4 *and* IPv6 routes carried over a single IPv4 *or* IPv6 peer session

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The 5600 supports peering between two IPv4 neighbors, or two IPv6 neighbors. IPv4 routes can be carried over an IPv4 peering session, or an IPv6 peering session. Likewise, IPv6 routes can be carried over an IPv4 peering session, or an IPv6 peering session. Any two IPv4 or IPv6 peers can be configured to carry both IPv4 and IPv6 routes.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuring eBGP

14

Let’s begin our configuration by setting up eBGP.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

eBGP Design Requirements

Directly-connected peers

- Exception? Add BGP multi-hop to the configuration

AS number

- AT&T vRouter training uses private AS numbers as well as private addresses in our examples

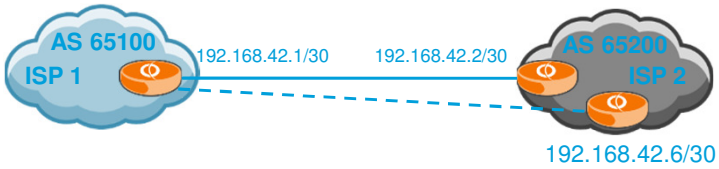

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To implement eBGP, you need to meet a couple of requirements.
As a rule, eBGP must be deployed between two directly-connected peers. This connection is typically a point-to-point WAN connection, but the only real requirement is that the peers be on the same subnet.
As networks have become more complex, eBGP allows exceptions to this rule through the use of a feature called BGP Multi-hop. However, where possible, it is a best practice to use a direct connection.
You must also apply for a unique AS number. These numbers are assigned by the same authority that assigns IP addresses.
Note that all the examples in our courseware use private AS numbers, just as we use private IP addresses. As with private IP addresses, private AS numbers cannot be used across the Internet.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuration Commands

Configure AS number

```
set protocols bgp as-number
edit protocols bgp as-number
```

Configure router ID

```
set parameters router-id ip-address
```

Configure peer information

```
set neighbor ip-address remote-as as-number
set neighbor ip-address address-family ipv4-unicast
set neighbor ip-address address-family ipv6-unicast
```

Specify networks to advertise

```
set network address/mask
```

- Network must already exist in routing table
- Best practice – create static route to blackhole

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

eBGP only has a few required configuration commands. First, you need to create the BGP instance by setting the AS number for your autonomous system. This is the AS number you received from your addressing authority.

Because you have several parameters to configure within the BGP instance, we recommend using the `edit` command to create the BGP instance and move into the appropriate location in the configuration hierarchy. This saves you some typing, and ensures that all subsequent commands are under the correct BGP instance.

The next step is to set the router ID for use with BGP. This address should be an IP address configured on your device, generally the address of the loopback interface. As a best practice, you should always explicitly set the router ID rather than allow a system to select one of the interface IP addresses.

Next, configure the BGP peer. First set the IP address and AS number of the remote peer. Second, configure the address families to be carried between the peers, IPv4, IPv6, or both. Finally, specify the networks you will advertise in BGP. This will generally be a supernet of networks and subnets within your AS. In order for BGP to advertise this network, the exact network and mask combination must already exist in the device routing table. If BGP cannot find the network and mask in the routing table, it will not include it in BGP updates. It is a best practice is to add a “blackhole” static route to the summary network and mask. All routing uses the longest match to determine the next-hop, so this practice not only ensures that BGP will always advertise the network as long as your device is up, but that

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

incoming packets to non-existent subnets within your domain are discarded immediately.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

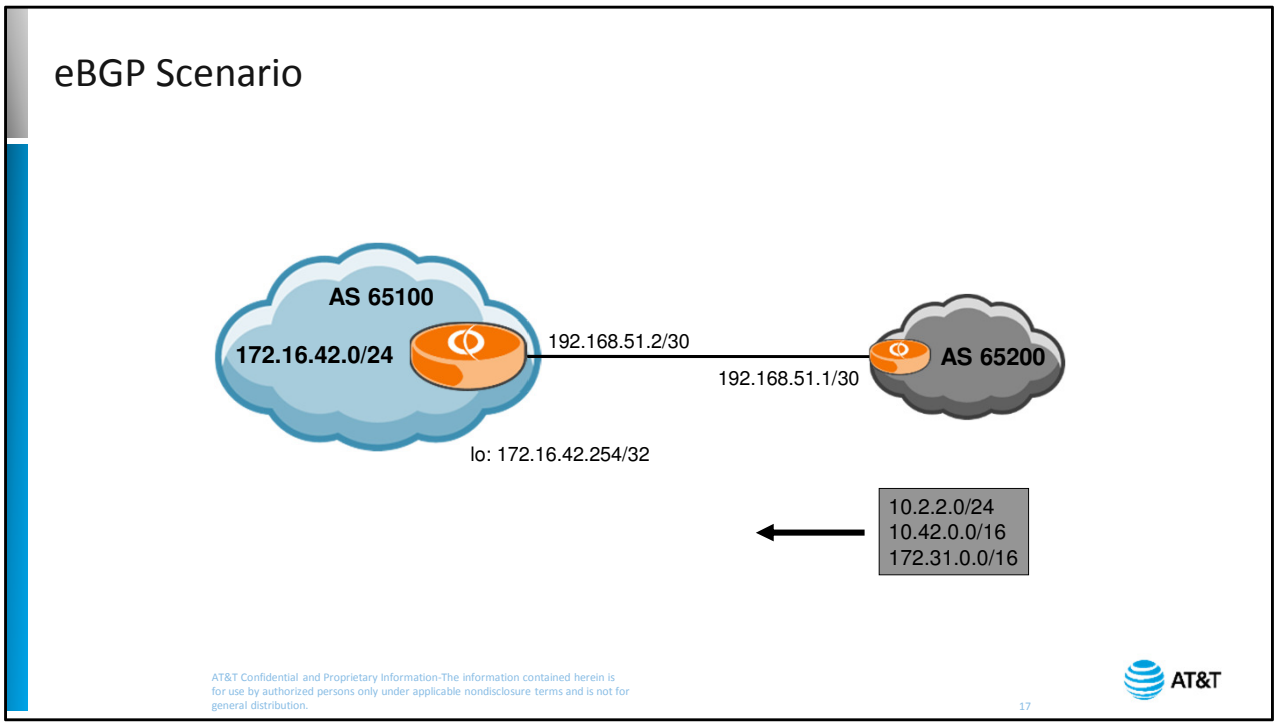

We will use this scenario to build our eBGP configuration.
Our AS number is 65100.
For the router ID, we use the address of our loopback interface. It is a best practice to use the loopback address as the router ID.
Our neighbor's IP address is 192.168.51.1 and we are directly connected.
The neighbor's AS number is 65200
The network prefix we will be advertising via BGP is 172.16.42.0/24
Our neighbor will be sending us several routes via BGP.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

eBGP Configuration

```
[edit]
vyatta@AS65100# edit protocols bgp 65100
[edit protocols bgp 65100]
vyatta@AS65100# set parameters router-id 172.16.42.254
[edit protocols bgp 65100]
vyatta@AS65100# set neighbor 192.168.51.1 remote-as 65200
[edit protocols bgp 65100]
vyatta@AS65100# set neighbor 192.168.51.1 address-family ipv4-unicast
[edit protocols bgp 65100]
vyatta@AS65100# set address-family ipv4-unicast network 172.16.42.0/24
[edit protocols bgp 65100]
vyatta@AS65100# exit
[edit]
vyatta@AS65100# set protocols static route 172.16.42.0/24 blackhole
[edit]
vyatta@AS65100# commit
[edit]
vyatta@AS65100# save
[edit]
vyatta@AS65100
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's build the configuration for our scenario.
First, we enter the `edit` command to create our BGP instance with AS 65100 and move us within the configuration hierarchy. Note the prompt changes to indicate where we are within the command hierarchy.
Next, we set the router ID. This address is the address we have already configured on our loopback interface.
Next, we enter the neighbor information, IP address, remote-AS number, and IPv4 address family.
Finally, we enter the network we want to include in our BGP advertisements.
To set the blackhole static route, we enter into the protocols configuration node, then enter the `static route` command, using the parameter `blackhole` for the next-hop.
Finally, we commit and save our changes.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

eBGP Multihop

- For peers that are not directly connected
  `set neighbor ip-address ebgp-multihop number-of-hops`

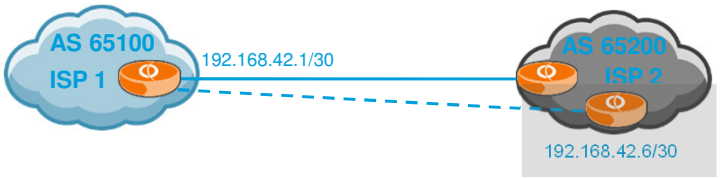

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

If your eBGP peer is not directly connected, you need to add the `ebgp-multihop` command to your neighbor configuration.
Specify the neighbor address, and the number of hops allowed between the two peers. By specifying a limit on the number of hops, you maintain a certain level of forwarding performance. If a routing change within the upstream ISP increases the hop count past the number of hops allowed, the BGP session drops.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying BGP Operations

20

Before we discuss iBGP configuration, let's make sure eBGP is working properly

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Displaying BGP Routes

Displays active routing entries

```
vyatta@AS65100:~$ show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       > - selected route, * - FIB route, p - stale info

Gateway of last resort is not set

B>* 10.2.2.0/24 [20/0] via 192.168.51.1, dp0p1p1, 00:51:42
B>* 10.42.0.0/16 [20/0] via 192.168.51.1, dp0p1p1, 00:51:42
C>* 10.224.7.0/24 is directly connected, dp0p1p2
C>* 127.0.0.0/8 is directly connected, lo
S>* 172.16.42.0/24 [1/0] is directly connected, Null0, bh
C>* 172.16.42.254/32 is directly connected, lo
B>* 172.31.0.0/16 [20/0] via 192.168.51.1, dp0p1p1, 00:51:42
C>* 192.168.51.0/30 is directly connected, dp0p1p1
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

One of the first places to look to verify BGP operations is the device routing table using the command `show ip route`. This command displays the routes the device is using to forward traffic.

Here we see that we are receiving BGP routes from the configured peer, so we know BGP is working properly. The BGP routes are represented by the *B* at the beginning of the route. However, what we cannot tell from this output is what networks we are adding to the BGP environment.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Displaying the BGP Table

Displays BGP table

- This is NOT the device routing table
- Routes may appear here that do not appear in the routing table

```
vyatta@AS65100:~$ show ip bgp
BGP table version is 0, local router ID is 172.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
l - labeled, S Stale,
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.2.2.0/24      192.168.51.1             0             0 65200 i
*> 10.42.0.0/16     192.168.51.1             0             0 65200 i
*> 172.16.42.0/24   0.0.0.0                  0     100 32768 i
*> 172.31.0.0       192.168.51.1             0             0 65200 i

Total number of prefixes 4
vyatta@AS65100:~$ show ip bgp | grep 0.0.0.0
*> 172.16.42.0/24   0.0.0.0                  0     100 32768 i
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To view the BGP table, use the command `show ip bgp`.
Note that this is not the device routing table. The vRouter maintains a separate table of BGP information that is then imported into the device routing table.
The BGP table lists all routes learned from all BGP peers, whether the routes are active or not. You may see networks with multiple next-hops in the table.
The best route is indicated by an arrow (>) next to the entry. You may also see networks with no best route indicator. This means the route is unusable and cannot be added to the device routing table. BGP will only advertise the best routes to other BGP peers.
We can find the routes we are originating by looking for entries with a next-hop of 0.0.0.0, a high weight, and no external path information.
In a real-world environment, this output may be several screens long. A quick way to reduce the output to just the routes we are originating is to use the Linux `grep` command and search for 0.0.0.0. The output will not include any headers – only the lines including the string we have searched for.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

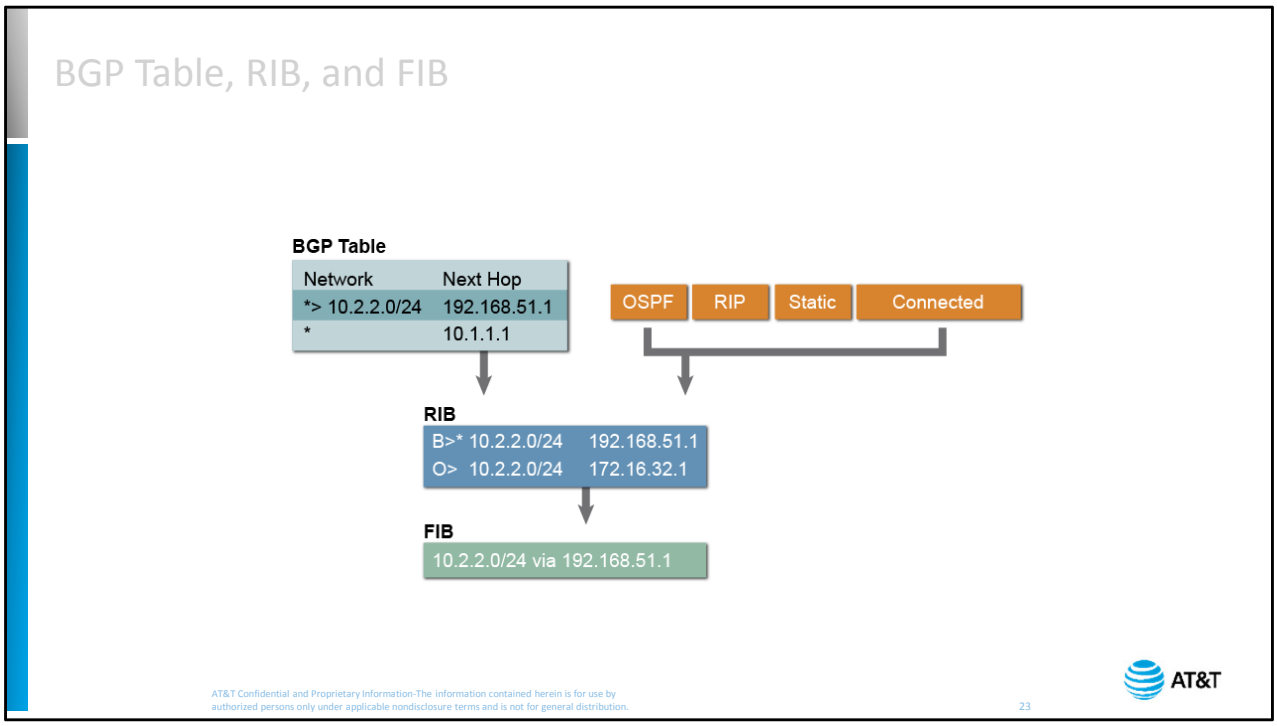

Let's pause a moment to explain the relationship between the BGP table, the routing information base, or RIB, and the forwarding information base, or FIB.
The BGP table includes all routes learned from all BGP peers. If a route is advertised by multiple peers, the BGP process flags one of the routes as the best path, as indicated here by the right arrow (>).
Information from all available route sources, including BGP, is combined in the Routing Information Base (RIB), also known as the routing table. This is the table you see with the `show ip route` command. Again, if a network is advertised by multiple sources, the vRouter selects one as the best route. Here, the BGP route is the best route, as indicated by the asterisk (*) next to the route.
The best route is added to the forwarding information base (FIB), or forwarding table. This is the table the vRouter uses to make traffic forwarding decisions.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Display BGP Status Summary

Display a summary of BGP configuration information and neighbor status

`show ip bgp summary`

```
vyatta@AS65100:~$ show ip bgp summary
BGP router identifier 192.168.101.1, local AS number 65100
BGP table version is 60
2 BGP AS-PATH entries
0 BGP community entries

IPv4 Unicast address family:
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.51.1    4 65200      13      12        6    0    0 00:05:05        3
Total number of neighbors 1

vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The `show ip bgp summary` command displays a summary of your BGP configuration, including your router ID, AS number, number of entries in the BGP routing information base, or RIB; amount of memory the RIB is using, number of peers, and counters for each peer.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Displaying BGP Neighbors

```
vyatta@AS65100:~$ show ip bgp neighbors
BGP neighbor is 192.168.51.1, remote AS 65200, local AS 65100, external link
  BGP version 4, remote router ID 10.4.5.6
  BGP state = Established, up for 02:17:13
  Last read 00:00:17, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    Address family IPv4 Unicast: advertised and received
  Received 120 messages, 0 notifications, 0 in queue
  Sent 121 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
 For address family: IPv4 Unicast
  BGP table version 61, neighbor version 61
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  1 announced prefixes
<Truncated Output>
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

For detailed information regarding your BGP peers, use the command `show ip bgp neighbor`. The output of this command can be useful when troubleshooting, as we will see later in the module.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

BGP Neighbor Options

Displays BGP neighbor information with options

`show ip bgp neighbors ip-address option`

Option	Displays
`advertised-routes`	Routes advertised to BGP neighbor
`received`	Information received from BGP neighbor
`received-routes`	Received routes from BGP neighbor
`routes`	Routes learned from BGP neighbor

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The `show ip bgp neighbors` command has several additional options you can use to limit the output to specific information. We will use some of these options later in the troubleshooting section of the module.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuring iBGP

27

We will now add iBGP to our configuration.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

iBGP Design Requirements

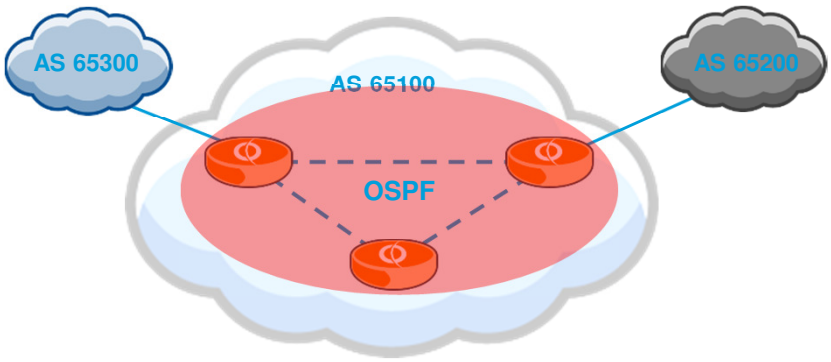

IP connectivity via IGP

- Best practice: peer to loopback addresses

Full mesh with all iBGP peers

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

As with eBGP, iBGP has some design requirements.
Unlike eBGP, iBGP peers do not have to be physically connected to each other. However, because BGP is connection-oriented, they do have to have routes to each other so they can establish a connection. Although you can use static routes to meet this requirement, you will generally use an interior gateway protocol (IGP) like OSPF to build the required routing.
A best practice is to use loopback interfaces for the peer addresses, then advertise the loopback addresses within the IGP. This ensures that iBGP connectivity does not rely on any one physical interface on your device; as long as a route exists to the peer's loopback address, the devices can maintain the iBGP session.
A second requirement is that you must fully mesh all your iBGP peers. This requirement stems from the way iBGP peers propagate routing information, which we will look at on the next screen.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

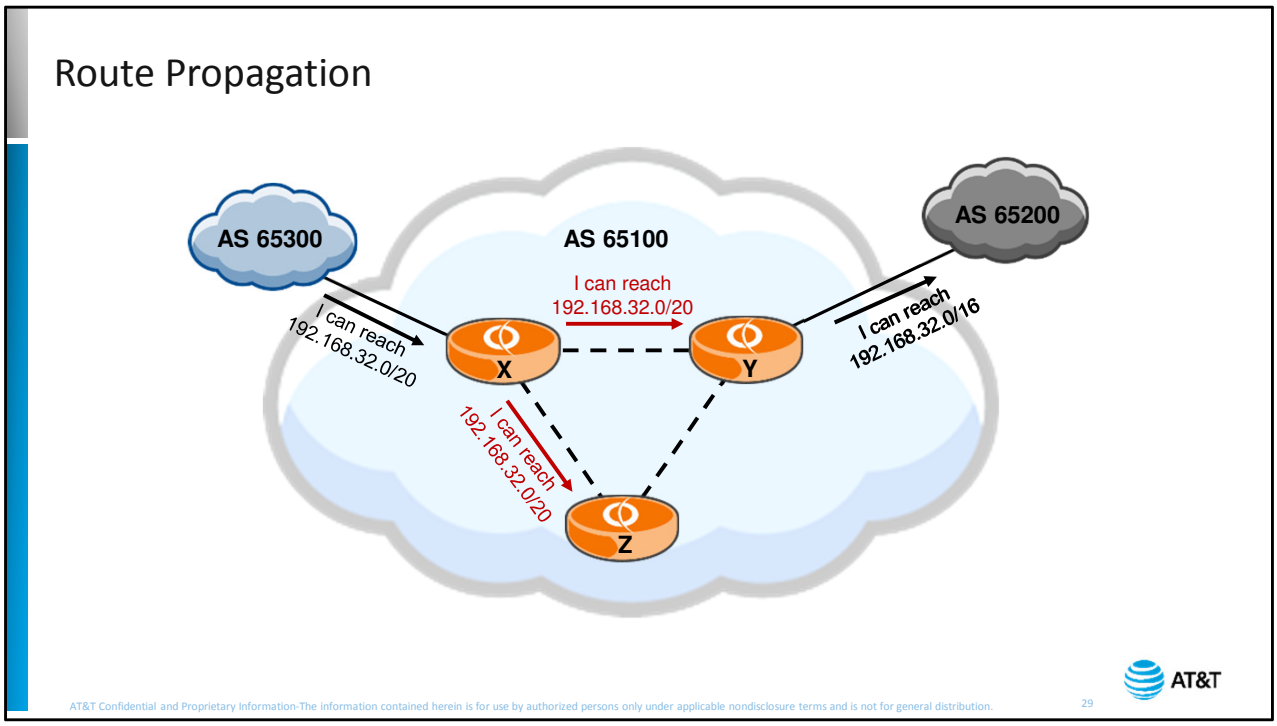

As we said on the previous screen, you must configure a full mesh of iBGP peers. The reason is to prevent routing loops, iBGP peers do not forward information learned via iBGP to other iBGP peers. Let's look at example to explain this further.

In our network, we have three routers running iBGP. Router Z has two peers – Router X and Router Y – but we have not configured X and Y to be peers with each other.

Router X receives an eBGP update from its peer in AS 65300. Because this is an eBGP route entry, router X will forward the information to all its iBGP peers – in this case, only to router Z. Because Router Z received the information via iBGP, it will not forward it to any other iBGP peer, meaning that Router Y has no information about reaching network 192.168.32.0.

By implementing a full mesh, Router X will forward the eBGP-received information to both Router Z and Router Y.

Router Y can then forward the information to its eBGP peer. Router Y and Router Z will still not propagate the information to each other, thus preventing an internal routing loop within AS 65100.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

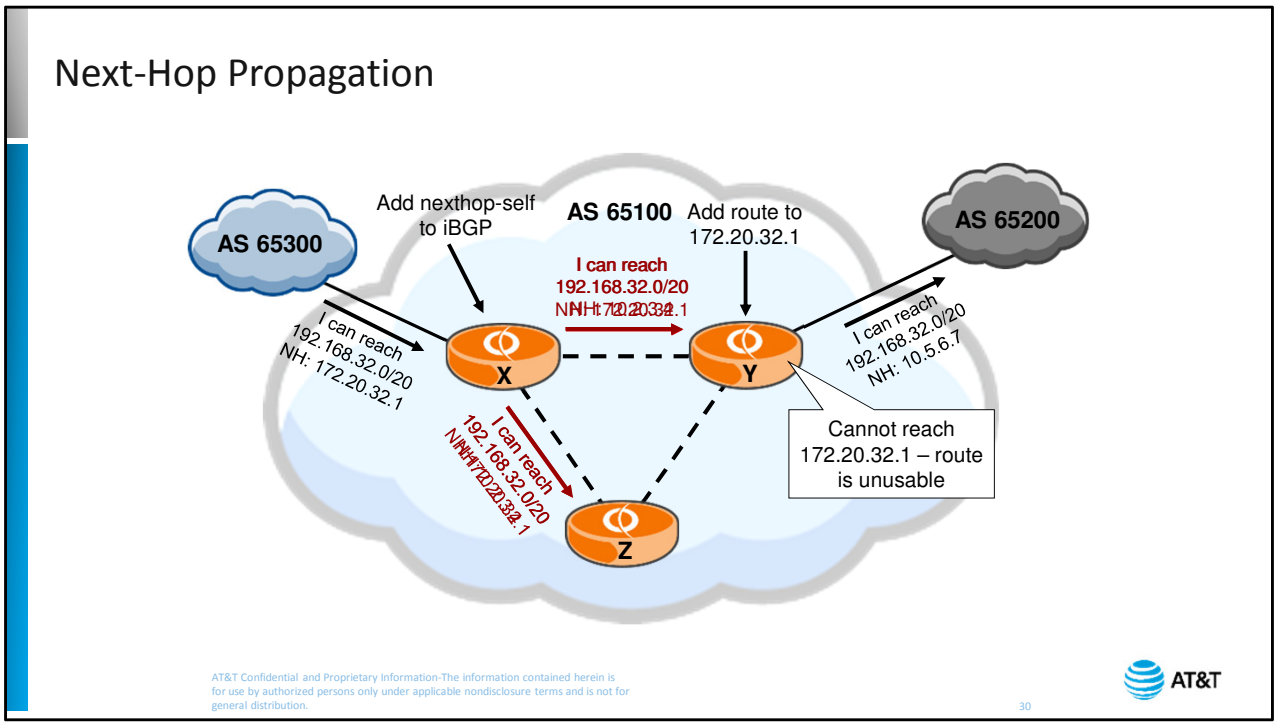

A related issue is how iBGP propagates next-hop information. By default, BGP routers do not modify next-hop information over iBGP links.

So, when router X receives an eBGP update for network 192.168.32.0, it propagates the update without changing the next-hop address. In order for routers Y and Z to mark the route as valid, they need to have a route to the next-hop address in their routing table.

If the next-hop is not reachable, then the routers mark the route as unusable, and do not add the information to their internal routing tables. It also means that router Y will not forward any advertisements for network 192.168.32.0 to AS 65200.

One solution is to add routing information for the next-hop, but this frequently is not a practical option.

A better solution is to add the `next-hop self` parameter to the iBGP configuration on router X.

This instructs router X to replace the next-hop address with the address it is using for the BGP session. This ensures that the peers always have a route, since they have to be able to reach router X in order to maintain the BGP session.

Router Y can then propagate the route to AS 65200. Because the connection is an eBGP connection, router Y automatically replaces the next-hop address with its own address.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuring iBGP

- Verify IP reachability
  - Implement IGP, ping peers, SSH between peers
  - Any firewalls/security to block BGP sessions?
- Configure iBGP
  - Router ID
  - Neighbors – use local AS number
    - Configure address to use for updates
      `set neighbor ip-address update-source local-ip`
    - Enable next-hop self for the IPv4 unicast address-family
      `set neighbor ip-address address-family ipv4-unicast nexthop-self`

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Before you begin the iBGP configuration, you should first verify that your iBGP peers can establish the required TCP connection.
Implement whatever routing you plan to use within your AS, then use Ping to verify reachability. Once you have verified reachability with Ping, verify that you can establish an SSH session between peers. This tests not only reachability, but that TCP is functioning between peers.
You should also verify that there are no firewalls or other security devices that may block the establishment of BGP sessions between your peers.
Once you have verified that the peers can successfully establish a connection, go ahead and configure iBGP.
The process is similar to configuring eBGP. First set the router ID, then configure the neighbors using the local AS number.
One additional parameter you need to add for iBGP is the address you will use as the source of your iBGP updates. A best practice is to set this address to the loopback address you are using for router ID. This address must not only be reachable by your iBGP peer, but must match the address your peer device has configured in its neighbor statement.
If needed, enable the `nexthop-self` command.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

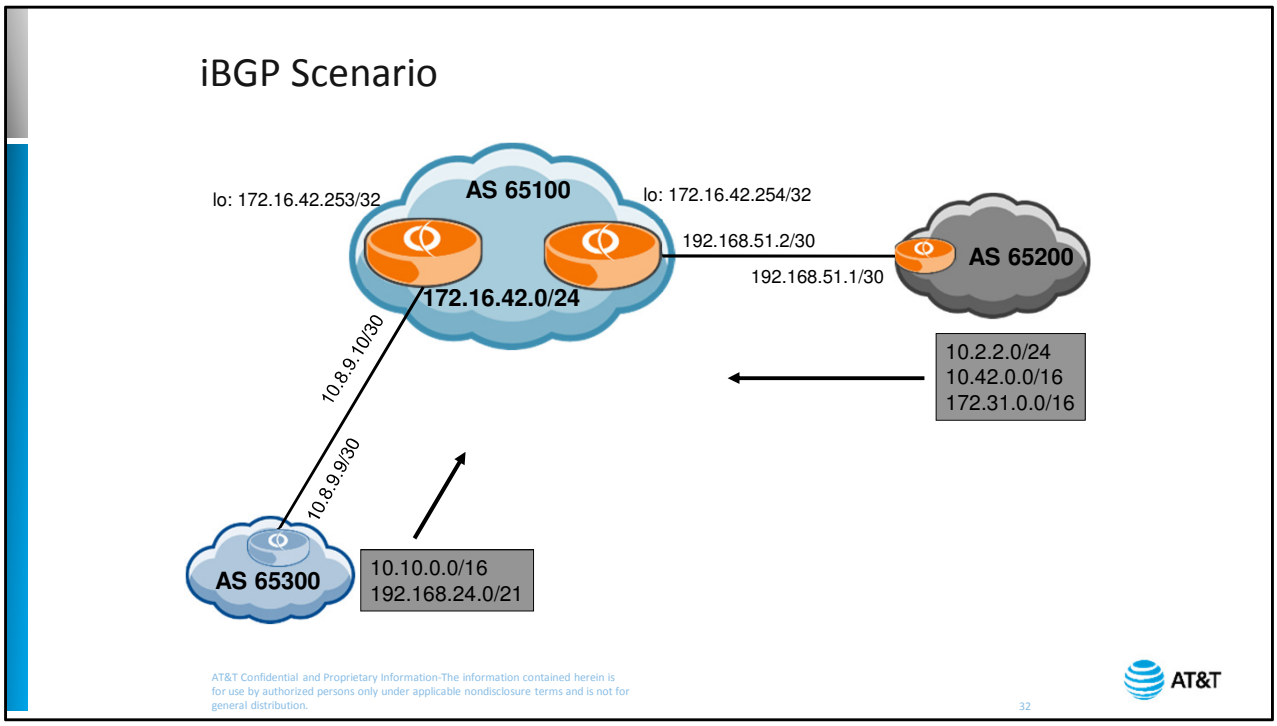

We have expanded our BGP scenario to include iBGP.
Our internal peer's loopback address is 172.16.42.253.
The peer also has an eBGP connection to AS 65300, so when our iBGP configuration is complete, we should see internal BGP routes for the networks coming from AS 65300.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying Peer Reachability

```
vyatta@AS65100:~$ show ip route ospf | grep 172.16.42.253
O>* 172.16.42.253/32 [110/20] via 172.16.42.2, dp0p1p1, 00:01:59
vyatta@AS65100:~$ ping 172.16.42.253
PING 172.16.42.253 (172.16.42.253) 56(84) bytes of data.
64 bytes from 172.16.42.253: icmp_seq=1 ttl=64 time=4.68 ms
64 bytes from 172.16.42.253: icmp_seq=2 ttl=64 time=0.608 ms
64 bytes from 172.16.42.253: icmp_seq=3 ttl=64 time=0.822 ms
^C
--- 172.16.42.253 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2005ms
rtt min/avg/max/mdev = 0.608/2.038/4.686/1.874 ms
vyatta@AS65100:~$ ssh 172.16.42.253
vyatta@172.16.42.253's password:
Welcome to Vyatta
Version:      3.2R2
Description:  AT&T Vyatta vRouter 5600 3.2 R2
Copyright:    2006-2014 Vyatta, a Brocade Company.
Last login: Mon Sep 14 19:24:39 2014 from 192.168.2.104
vyatta@AS65100-2:~$ logout
Connection to 172.16.42.253 closed.
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

33

Our first step is to verify that we can reach the peer.
We look in the routing table to see if we have a route to the peer. We are running OSPF in our network, and we know that we will be using a loopback address to reach the peer, so we can use the `grep` command to search for a routing entry specific to the address.
Next, we run the `ping` command to verify basic IP reachability. Ping is successful!
Next, we try opening an SSH connection to verify that TCP works end to end. We enter the password for the peer device, and complete the login, showing that we can connect.
We then log out and return to our device to begin configuring iBGP.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

iBGP Configuration

```
[edit]
vyatta@AS65100# edit protocols bgp 65100
[edit protocols bgp 65100]
vyatta@AS65100# edit neighbor 172.16.42.253
[edit protocols bgp 65100 neighbor 172.16.42.253]
vyatta@AS65100# set remote-as 65100
[edit protocols bgp 65100 neighbor 172.16.42.253]
vyatta@AS65100# set update-source 172.16.42.254
[edit protocols bgp 65100 neighbor 172.16.42.253]
vyatta@AS65100# set address-family ipv4-unicast nexthop-self
[edit protocols bgp 65100 neighbor 172.16.42.253]
vyatta@AS65100# commit
[edit protocols bgp 65100 neighbor 172.16.42.253]
vyatta@AS65100# save
[edit]
vyatta@AS65100#
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

We enter Configuration mode, then begin editing our BGP instance.
We have already set the router ID, so we can begin by adding our iBGP neighbor. We use the `edit` command since we have a few parameters to set for this neighbor.
We set the remote AS to be our AS number. Using the same AS number tells the vRouter that this is an iBGP peer.
We set the update source to be the address of our loopback interface. This is also our router ID.
Because we are connected to another AS, we set the `nexthop-self` command to ensure our learned routes propagate correctly through the iBGP connection.
Finally, we commit and save our changes.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Side-by-Side Configurations for AS65100

```
AS65100# show protocols bgp 65100
address-family {
     ipv4-unicast {
          network 172.16.42.0/24
     }
 }
          neighbor 192.168.51.1 {
              address-family {
                  ipv4-unicast {
                  }
              }
              remote-as 65200
}
          neighbor 172.16.42.253 {
              address-family {
                  ipv4-unicast {
                  nexthop-self
              }
     }
     remote-as 65100
     update-source 172.16.42.254
 }
 parameters {
     router-id 172.16.42.254
```

```
AS65100-2# show protocols bgp 65100
address-family {
     ipv4-unicast {
          network 172.16.42.0/24 {
     }
 }
          neighbor 10.8.8.9 {
              address-family {
                  ipv4-unicast {
                  }
              }
              remote-as 65300
}
          neighbor 172.16.42.254 {
              address-family {
                  ipv4-unicast {
                  nexthop-self
              }
     }
     remote-as 65100
     update-source 172.16.42.253
 }
 parameters {
     router-id 172.16.42.253
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look at the configurations for the two vRouter devices inside AS 65100.
We can see our iBGP peer configuration, including the `nexthop-self` parameter.
Remember, this is only necessary if the peers do not have routes to the eBGP peers.
We also see the eBGP peer configuration.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

eBGP and iBGP routes

```
vyatta@AS65100:~$ show ip bgp
BGP table version is 0, local router ID is 172.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
l - labeled, S Stale,
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.2.2.0/24      192.168.51.1             0             0 65200 i
*>i10.10.0.0/16     172.16.42.253            0    100      0 65300 i
*> 10.42.0.0/16     192.168.51.1             0             0 65200 i
* i172.16.42.0/24   172.16.42.253            0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 172.31.0.0       192.168.51.1             0             0 65200 i
*>i192.168.24.0/21  172.16.42.253            0    100      0 65300 i

Total number of prefixes 6
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look closely at the BGP table.
Routes that we have learned through iBGP are marked with an i before the network information.
We can also see that our iBGP neighbor is advertising the same network we are – 172.16.42.0. However, since we are originating a route to the same network, our originated route is considered the best path.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Troubleshooting Basic BGP Operations

37

We end this module with a look at some common BGP configuration problems and how to identify and resolve them.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

No Connection to Peer

What displays in the summary table

- No routes from peer
- Peer shows Active state

```
vyatta@AS65100:~$ show ip bgp summary
BGP router identifier 172.16.42.254, local AS number 65100
BGP table version is 4
1 BGP AS-PATH entries
0 BGP community entries

IPv4 Unicast address family:
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.42.253   4 65100       0       0        0    0    0 never    Active
192.168.51.1    4 65200       1       4        0    0    0 00:20:13        3
Total number of neighbors 2

vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When you first configure BGP, a common problem is that you are not able to establish a BGP connection with a peer router.
The first obvious symptom is that you have no routes from the peer.
The second symptom is that your device shows the peer state as *Active*. You may think “Active” is good, but in the case of BGP, it indicates that your device is actively trying to establish a connection with the peer and has not been successful.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

No Connection to Peer - Solutions

Check for basic IP connectivity

Check for misconfigured neighbor addressing

```
vyatta@AS65100-2# show protocols bgp
65100
<Truncated Output>
        neighbor 172.16.42.254 {
            address-family {
                ipv4-unicast {
                nexthop-self
            }
        }
        remote-as 65100
        update-source 172.16.42.253
<Truncated Output>
```

```
vyatta@AS65100# show protocols bgp
65100
<Truncated Output>
        neighbor 172.16.42.253 {
            address-family {
                ipv4-unicast {
                nexthop-self
            }
        }
        remote-as 65100
        update-source 172.16.42.254
<Truncated Output>
```

Check for intervening firewalls

- Make sure firewalls permit BGP (TCP port 179)

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To resolve this problem,
first check for basic IP connectivity. Can you ping the peer device? If not, then you have a routing or addressing problem to solve.
If you can ping, make sure that the neighbor address is configured correctly on both peers. A device will not accept a BGP connection from an IP address it does not have defined as a neighbor.
If you are using iBGP, you also need to make sure the neighbor address on your device matches the update source on the peer and vice-versa.
If the configuration is correct, then check for any firewalls in the path between your device and the peer.
If there are firewalls, make sure those firewalls are permitting BGP traffic. BGP uses TCP port 179.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

No Active Routes From iBGP Peer

What displays in the BGP table

- Peer session is established
- No routes in the IP routing table from peer (`show ip route`)
  - `show ip route bgp` limits output to BGP routes
- Routes in BGP table are inactive

```
vyatta@AS65100:~$ show ip bgp
BGP table version is 0, local router ID is 172.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
l - labeled, S Stale,
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.2.2.0/24      192.168.51.1             0             0 65200 i
* i10.10.0.0/16     10.8.9.9                 0    100      0 65300 i
*> 10.42.0.0/16     192.168.51.1             0             0 65200 i
* i172.16.42.0/24   172.16.42.253            0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 172.31.0.0       192.168.51.1             0             0 65200 i
* i192.168.24.0/21  10.8.9.9                 0    100      0 65300 i
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Another common problem is that you are not seeing iBGP routes in your routing table.
After you have eliminated a peer connectivity problem,
you still do not see any routes from the peer in the routing table.
Note that you can use `show ip route bgp` to limit the show command output to just BGP routes.
When you look in the BGP table, you see that you have routes from your iBGP peer, but they are inactive; they do not have a best route arrow (>) next to them.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

No Active Routes From iBGP Peer - Solutions

```
vyatta@AS65100:~$ show ip bgp
BGP table version is 0, local router ID is 172.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
l - labeled, S Stale,
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.2.2.0/24      192.168.51.1             0             0 65200 i
* i10.10.0.0/16     10.8.9.9                 0    100      0 65300 i
*> 10.42.0.0/16     192.168.51.1             0             0 65200 i
* i172.16.42.0/24   172.16.42.253            0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 172.31.0.0       192.168.51.1             0             0 65200 i
* i192.168.24.0/21  10.8.9.9                 0    100      0 65300 i
```

- Configure `nexthop-self` on peer, or
- Configure routing to next-hop address

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When we look more closely at the BGP table, we see that the next-hop address shown for these routes is outside our network. Unless we have a route to this destination, these routes will remain inactive.

The solution is to either configure `nexthop-self` on the peer so that it uses its own address as the next-hop address, or configure a route entry for the address outside our network. Configuring `nexthop-self` is generally the best practice.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Upstream Devices Can't Reach Your Network

Symptom – complaints from people trying to reach your network

Am I advertising the network to my eBGP peer?

`show ip bgp neighbors ip-address advertised-routes`

```
vyatta@AS65100:~$ show ip bgp neighbors 192.168.51.1 advertised-routes
BGP table version is 0, local router ID is 172.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

IPv4 Unicast address family:
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.0.0/16     192.168.51.2                 100      0 65300 i
*> 172.16.42.0/24   192.168.51.2             0          32768 i
*> 192.168.24.0/21  192.168.51.2                 100      0 65300 i

Total number of prefixes 3
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

One final problem you may encounter is that upstream devices are unable to reach your network.

The first symptom is going to be complaints from users outside your network who are trying to reach your public resources, but are unable to do so.

The first thing to check is whether you are advertising your network to your eBGP peer. You can do this by using the `show ip bgp neighbor` command, checking for the advertised routes for a specific peer.

You want to see the route for your network, as shown here. Since no AS number is displayed in the path column this is an indication that this route is originating from your own AS.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Upstream Devices Can't Reach Your Network (cont.)

If you are not advertising

- Check network statement in BGP configuration
- Check for active route in IP routing table
- Check for routing policies applied to the peer configuration

If you are advertising

- Clear the eBGP advertised routes for the specified address family
  `reset ip bgp external ipv4 | ipv6 | soft`
- Contact your ISP to verify their routing policies

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

If your network is not listed in the advertised routes table,
first, check to see that you have a network statement in your BGP configuration for the network.
If the network statement is there, then check your routing table using `show ip route` to make sure you have an active route to the network.
If the route is there, then check for any routing policies that may be blocking advertisements to your peer.
If you have resolved one of the above problems and the route now appears in your BGP table, but is still not visible by your peer,
You need to clear the externally-advertised routes. By default, BGP will not change its externally advertised routes due to administrative changes until the administrator clears the connection. Remember, a BGP table can contain hundreds of thousands of routes, so you want the administrative control over when a change triggers a full table exchange.
If that does not resolve the problem, you need to ask your ISP to verify their routing policies. They may have a policy blocking your network advertisement upstream.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

You are Carrying Transit Traffic

If you are not an ISP, you do not want this to happen!

Solution: Filter out all non-local networks from eBGP advertisements with a routing policy

For more information on routing policies, please refer to the *AT&T Vyatta 5600 vRouter Software Documentation*

```
vyatta@AS65100# show policy
route {
  prefix-list MyNetwork {
     rule 1 {
        action permit
        prefix 172.16.42.0/24
     }
 }
  route-map Export2eBGP {
      rule 10 {
        action permit
        match {
             ip {
                 address {
                     prefix-list MyNetwork
                 }
[edit]
vyatta@AS65100# set protocols bgp 65100
neighbor 192.168.51.1 address-family ipv4-
unicast route-map export Export2eBGP
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

A final problem to consider is whether you want to allow transit traffic; that is, do you want traffic from other companies crossing your network.
Unless you are an ISP, chances are the answer to this is NO, you do not want to allow outside traffic to cross your network.
To prevent this from happening, you need to use a routing policy to block all networks except for yours in your eBGP updates.
The example shows a sample policy, as well as the command to apply it to an eBGP neighbor.
For more information on routing policies, please refer to the *AT&T Vyatta 5600 vRouter Software Documentation* on http://businesscenter.att.com

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Summary

You should now be able to

- Describe BGP functionality and applications
- Describe 5600 vRouter BGP support for IPv6
- Configure eBGP connectivity to a single ISP
- Configure basic iBGP connectivity
- Verify BGP operations
- Troubleshoot common problems

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

This concludes the AT&T Vyatta 5600 vRouter BGP Basics module.
You should now be able to:

- Describe BGP functionality and applications
- Describe 5600 vRouter BGP support for IPv6
- Configure eBGP connectivity to a single ISP
- Configure basic iBGP connectivity
- Verify BGP operations
- Troubleshoot common problems

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T Vyatta 5600 vRouter BGP Policies

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Objectives

After completing this module, you will be able to

- Describe BGP routing policy operations and applications
- Configure BGP address-based and AS-path-based route filters
- Verify filter configuration and operations

Note: This module assumes that you have completed the *AT&T Vyatta 5600 vRouter BGP Basics* module or have equivalent experience

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

After completing this module, you will be able to
- Describe BGP routing policy operations and applications
- Configure BGP address-based and AS-path-based route filters
- Verify filter configuration and operations

Note: This module assumes that you have completed the AT&T Vyatta 5600 vRouter BGP Basics module or have equivalent experience.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Policy Overview

48

We begin with an overview of routing policies.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

What is a Routing Policy?

Allows you to accept, reject, or modify attributes for

- Routes received via dynamic routing protocols
- Routes sent via dynamic routing protocols

Some BGP-specific applications

- Prevent transit traffic by blocking all non-originated routes from eBGP neighbors
- Block private networks from being advertised in BGP
- Modify BGP route attributes to affect path selection

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Routing policies modify the default behavior of routing protocols. In short, a routing policy allows you to accept, reject, or modify attributes of a routing table entry.
You can use policies to affect routes you receive from other routers, or affect routes before you advertise them to other routers.
Some typical applications of routing policies for BGP include:
Blocking all non-originated routes from being advertised to BGP neighbors. This prevents your network from being used as a transit domain.
Another application is to ensure that private networks are not included in BGP updates. You can apply this to both your advertisements and to advertisements you receive from your BGP peers.
Another common application is to modify BGP route attributes to change the preferred route that is added to the routing table.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Routing Policy Components

Filters

- Define what route entries to act upon
- Can filter based on route information, route source/destination, AS-path, and other criteria

Route Maps

- Applies filters and defines what actions to take
- Actions can include modifying routing attributes, permitting routes, or dropping routes

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

A routing policy consists of two components.
The first is the filter. A filter defines which routing table entries you want to act upon. The filter does not specify the action itself – it only identifies the target of the action.
You can filter based on the route information, the source or destination of the routing update, AS-path information, and certain BGP route attributes.
The second policy component is a route map.
The route map references the defined filter, then specifies the action to take on whatever routes match the filter.
The action can be as simple as accepting or dropping the route from the table, or can modify one or more of the route attributes.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

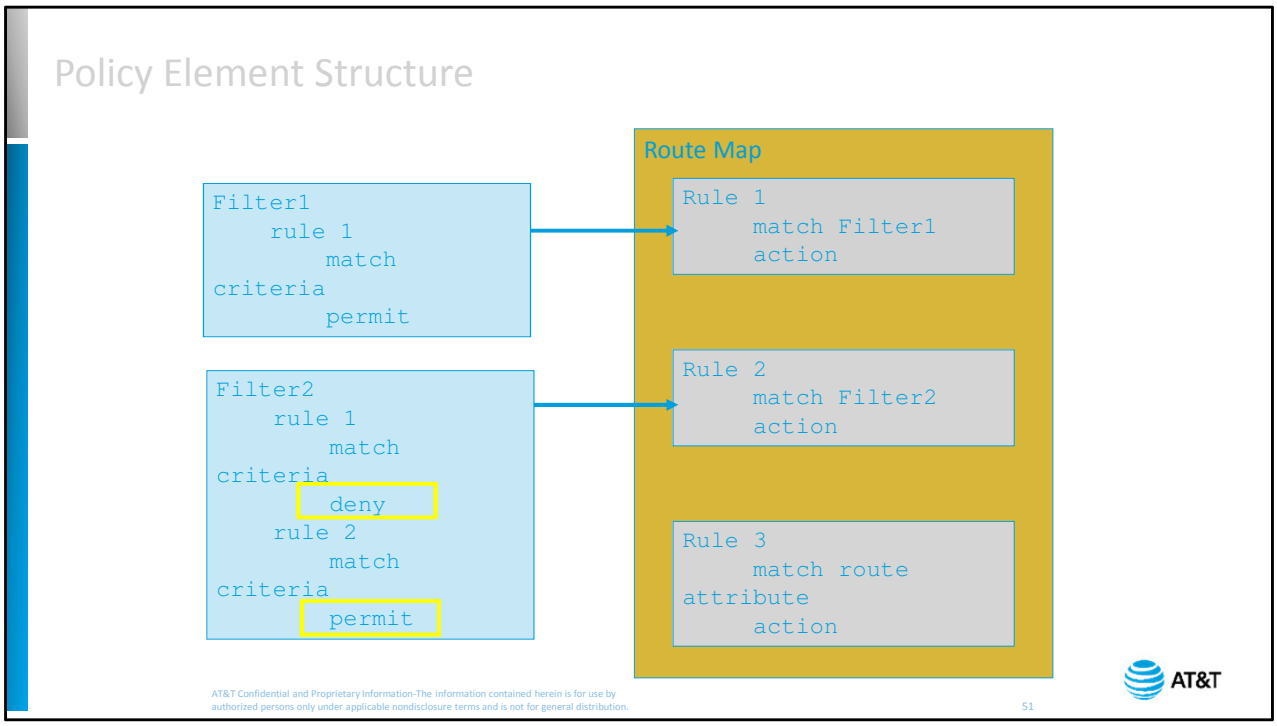

This slide shows the structure of the policy elements and how they relate to each other.
If you are using address-based or AS-path filters, you create those filters separately. You can have a simple filter,
or you may need a more complex filter to specifically identify the routes you want to modify. The default action for a filter is to deny everything, so you must include at least one permit statement for the filter to be valid.
In the route map, you create rules that invoke the filters you created. If a route matches the filter criteria, the vRouter will apply the configured action to that route
If the route does not match, the vRouter will go to the next rule in the list.
It is important to note that the permit or deny statements in the filter do not drop or allow traffic from the routing table itself. They permit or deny a routing entry from being affected by the action in the route map. A permitted route will be modified based on the rule. A denied route will not be affected and will pass to the next rule in the list.
If you are filtering based on route attributes, you do not need to configure an external filter. Attribute filters are configured directly within the route map itself.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

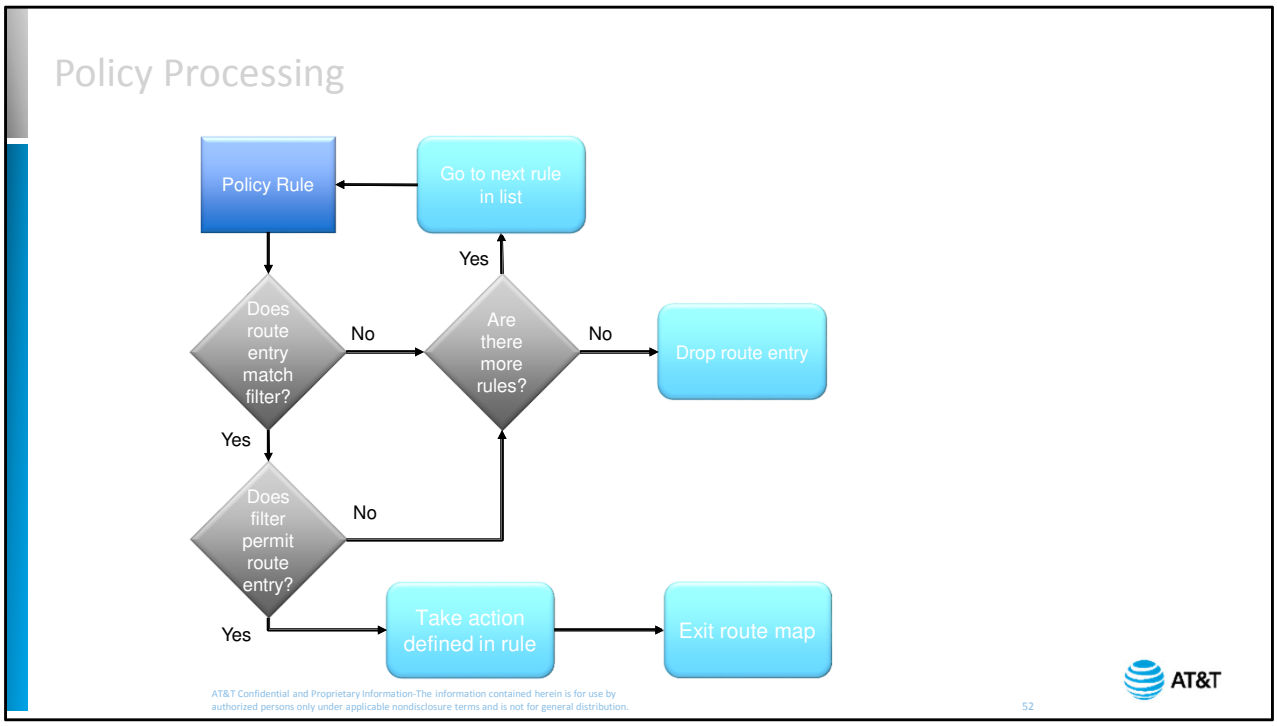

Now let's look at how the vRouter evaluates the routing policy.
We start with the first rule in the policy and evaluate whether the route entry matches the defined filter criteria.
If the answer is Yes, then the next question is whether the filter action is permit or deny.
If the filter action is to permit the route entry, then the vRouter applies the action defined in the rule and exits the list. It will not evaluate any other rules in the list by default.
If the route entry does not match the filter criteria, the device checks to see if there are more rules in the list.
If there are, the device goes to the next rule in the list.
Going to the next rule, we could have a route entry that matches the filter criteria, but is denied by the filter.
In that case, it is treated as not matching the filter, and the device attempts to go to the next rule.
If there are no more rules in the list, the device drops the route entry from the routing table. It is important to remember that the default action is to drop the routing entry; if you want to override this, you configure your last rule with a filter that matches all route entries and an action to permit the routes into the routing table.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuring Filters

Next, we will look at the commands used to configure several different filter types.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Filter Options for BGP

- Prefix list
  - Defines networks (routing table entries) to filter
- AS-path list
  - Uses regular expressions to match information in BGP AS-path
- Access list
  - Defines individual IP address or range of addresses to filter
    - Use to modify route entries based on next-hop address, source or destination address of routing update, peer address, or route source
    - Out of scope for this course
- Community list
  - Uses regular expressions to match BGP community information
    - Out of scope for this course; please refer to the *AT&T vRouter BGP Tuning* module for more information on BGP communities

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When working with BGP, you have four stand-alone filters you can configure in addition to the route attribute filters embedded in the route map.
The first type is the prefix list.
The prefix list identifies the routing table entries you want to filter. You can configure an exact match, or specify an address as a summary and then specify a mask length to match.
The second type of BGP policy filter is the AS-path list. You use regular expressions to create filters that match AS numbers in the BGP route AS path.
The third type of filter is an IP access list. This list defines an individual IP address or range of addresses to filter. Although access lists are generally used to block traffic flowing through a device, they can be used in route maps to filter route entries based on the next-hop address in the routing advertisement, or the source or destination of the routing update, or the BGP peer address, or the ultimate source of the route.
Because this is rarely used in production networks, we do not cover access lists in this module.
The last type of filter is the community list. Again, you use regular expressions, in this case to match information in the BGP community attribute.
We will cover BGP communities in the *AT&T Vyatta 5600 vRouter BGP Tuning* module.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Prefix List Commands

Create Prefix List

```
set policy route prefix-list name
edit policy route prefix-list name
```

- Create rule

```
set rule number <1-65535>
edit rule number <1-65535>
```

  - Set action

```
set action [permit | deny]
```

  - Specify prefix

```
set prefix address/mask
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

We start with the command syntax for a prefix list.
First, you create the list. A prefix list is part of the policy configuration hierarchy, and is a named entity. A best practice is to use a name that is related to the list function.
Because you will be configuring several parameters within the prefix list, we recommend using the `edit` command to ensure that all the parameters are set within the same prefix list, as well as to save typing and prevent typing mistakes.
Within the prefix list, create a rule. Rules are numbered, and it is a best practice to leave space between the rule numbers in case you need to insert rules later on. In our examples, we begin with rule 10, then rule 20, and so on. Also note that each rule has both actions and filter criteria, so we again recommend using the `edit` command.
Within the rule, set the action to either permit or deny. Remember, this does not mean that the route information that matches the rule will be permitted or denied into the routing table. A permit means the route map action will apply to the route; a deny action means the route will be passed to the next rule in the list.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Prefix List Options

Optional configuration

- Match prefix lengths greater than or equal to the specified prefix length
  `set ge prefix-length <0-32>`

- Match prefix lengths less than or equal to the specified prefix length
  `set le prefix-length <0-32>`

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Still within the rule, specify the routing prefix you want to match. If you type in an address and mask combination, that exact route entry will match the rule. Subnets or supernets of the route entry will not match.
You can expand the range of routing prefixes by adding the following options:
The greater than option (*ge*) converts the specified prefix into a summary address and mask. Any subnets with the specified prefix length or greater, that fall within the specified prefix, will now match the rule.
The less-than option (*le*) also converts the specified prefix into a summary address and mask. All prefixes with the specified length or less, up to the length of the prefix itself, will now match the rule.
The prefix length you set for the greater than value must be larger than the mask in the prefix statement. The less than value must be greater than or equal to the greater than value if used together.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Using Greater Than and Less Than

Prefixes in table

10.0.0.0/8
10.24.0.0/12
10.99.0.0/16
10.100.32.0/20
10.101.48.0/24

```
prefix-list BlockNet10 {
   rule 10 {
        action permit
        prefix 10.0.0.0/8
    }
}
```

```
prefix-list BlockNet10 {
   rule 10 {
        action permit
        ge 9
        prefix 10.0.0.0/8
    }
}
```

```
prefix-list BlockNet10 {
   rule 10 {
        action permit
        ge 16
        le 23
        prefix 10.0.0.0/8
    }
}
```

```
prefix-list BlockNet10 {
   rule 10 {
        action permit
        le 32
        prefix 10.0.0.0/8
    }
}
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look at an example to see how adding the greater than and less than parameters affect which routing entries match the filter.
When we do not add the greater than or less than parameters, we have specified an exact match list. The only prefix in the table that matches the list is 10.0.0.0 with an 8 bit mask.
When we add the greater than setting, we match all routes that fall into the specified prefix that have the specified mask length or longer. Note that the *ge* option does not allow you to specify a prefix length that is the same as the prefix itself. If we want the prefix AND all summarizable routes to match the filter
We use the *le* option with a mask length of 32. This matches all routes within the specified prefix range, including the prefix itself.
We can use both settings in combination to set a range of matches. Now, only the routes with masks between 16 and 23 bits long will match the filter.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AS-Path List Commands

Create list

```
set policy route as-path-list name
edit policy route as-path-list name
```

- Create rule

```
set rule number
edit rule number
```

  - Set action

```
set action [permit | deny]
```

  - Set filter criteria

```
set regex regular-expression
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

An AS-path list has a similar structure to the prefix list.
First, you create the list, using either the `set` or `edit` command. Again, we recommend using the `edit` command so that all subsequent commands are associated with the right list.
Next, create a rule within the list. Again, these are numbered rules.
For each rule, configure an action.
The filter criteria for an AS-path list is a regular expression that matches information in the AS-path attribute.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Regular Expressions

- AS-path is a character string
  - Example: `65100 65298 65430 65123`
  - Use regular expression to match part of the string
- Regular expression special characters

^	Matches the beginning of the input string
$	Matches the end of the input string
_	Matches a comma, braces, parentheses, space, beginning or end of input string
.	Matches any single character, including white space
*	Matches 0 or more sequences of the pattern
*$	Matches everything

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The AS-path route attribute is simply a string of numbers that lists the autonomous systems the route has passed through. Follow the AS-path back and you will eventually reach the origin of the route.
Each AS number is separated by a space. The first number in the list is the closest AS. The last number is the originating AS.
To filter on this information, you need to configure a regular expression that matches the relevant information in the string.
Regular expressions use some special characters that you need to be familiar with.
The ones most often used in AS-path filters are listed both on the screen and in the configuration commands.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Regular Expression Examples

- The most common mistake is forgetting the spaces before and after the AS number

String	Description
65300	Matches any instance of 65300, including 765300, 65300x, etc
_65300_	Matches 65300 anywhere in the AS-path
^65300	Matches 65300 at the beginning of the AS-path, including 65300, 653001122, etc.
^65300_	Matches 65300(space) at the beginning of the AS-path
_65300$	Matches 65300 at the end of the AS-path

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The most common mistake made when using regular expressions is forgetting to include the spaces in your expression.
For example, if you enter 65300 as your expression, it will match any instance of the pattern 65300, including 765300, 65300 followed by other integers, and so on.
To include the spaces, add the underscore (_) characters to the expression. The expression now matches only 65300. Because the underscore also matches the beginning and the end of the string, this expression will match 65300 no matter where its found in the AS path.
If you want to match 65300 just at the beginning of the string, use the caret (^) character. However, you still have to remember the space after 65300. If you do not specify something, ^65300 will match 65300 followed by anything, as long as its at the beginning of the string.
To match just 65300, use the caret (^) at the beginning and the underscore (_) at the end.
If you want to match a pattern at the end of the string, enter the pattern followed by the dollar sign ($).

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Testing Filters

Operational mode commands

- Displays all BGP table entries that match the specified prefix list
  `show ip bgp prefix-list name`

- Displays all BGP table entries that match the specified AS-path list
  `show ip bgp filter-list name`

```
vyatta@AS65100:~$ show ip bgp prefix-list Match172.16.0.0 | grep 65200
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Before using your filter in a route map, you can check it against the active BGP table. These commands are Operational mode commands and do not affect the active configuration or operation of your device.
To check your prefix list, enter the command `show ip bgp prefix-list`, followed by the list name. The output will display all the routes in the BGP table that match your configured list.
To check the AS-path list, enter the command `show ip bgp filter-list`, followed by the list name. The output will display all the routes in the BGP table that match the specified AS-path list.
These commands are a great way to preview the effect your filter would have on the BGP table, by showing you all the routes that match the configured list. Keep in mind, that these commands check the entire BGP table, and not routes sent to you by a particular neighbor. If you want to check which routes from a particular neighbor would match, you can pipe (|) the output through a `grep` search for the neighbor AS number. Only the lines that include the neighbor AS will display in the output. We will look at an example of this when we get to our configuration scenarios.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuring Route Maps

Next, we will apply our filters using route maps. This section includes several example configurations.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Route Map Commands

Create route map

```
set policy route-map name
edit policy route-map name
```

- Create rule

```
set rule number
edit rule number
```

- Set match criteria

```
set match ip | ipv6 address prefix-list name
set match as-path as-path-list-name
```

  - Other match criteria include community, interface, metric, origin, and tag

- Specify action for rule

```
set action [permit | deny]
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The structure of a route map is similar to the structure of a filter list.
First, you create the route map, and then create a rule.
Within the rule, you set the match criteria. This can reference an existing prefix list, an existing AS-path name, or can be a routing attribute, such as OSPF cost or BGP origin.
Next, specify the action for the rule. The action is either *permit* or *deny*.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Route Map Options

Optionally use `set set` *`attribute`* to modify route attributes

Attribute	Description
aggregator	BGP aggregator attribute
as-path-prepend	Prepend string for a BGP AS-path attribute
atomic-aggregate	BGP atomic aggregate attribute
community	BGP community attribute
delete-community	BGP communities to delete
ip-next-hop	Next-hop IP address
ipv6-next-hop	Next-hop IPv6 address
local-preference	BGP local preference attribute
metric	Destination routing protocol metric
metric-type	OSPF external metric-type
origin	BGP origin code
originator-id	BGP originator ID attribute
tag	Tag value for routing protocol
weight	BGP weight attribute

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

You can optionally modify one or more routing attributes for the routes that match the filter. Notice the command syntax; you will type `set` two times. The attribute options depend on the routing protocol you are modifying; BGP has several attributes you can set.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Applying Route Map to BGP Peers

- Test before applying
  `show ip bgp route-map name`
  - Tests effects of route map on BGP table without interfering with BGP operations
  - Operational mode command
- Apply to BGP peer
  `set protocol bgp as-num neighbor ip-addr route-map [import | export] name`
  - Configuration mode command
- Manually reset BGP peer connections
  `reset ip bgp external [ipv4 | ipv6 | soft] unicast soft`
  - Operational mode command

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

You can test the effect of your route map on the BGP table without interfering with BGP operations by using the Operational mode command `show ip bgp route-map` followed by the name of the route map. As with the filter tests, this command applies the route map to the entire BGP table rather than to routes coming from a particular neighbor, so unless you plan to apply the route map to all your neighbors, the output is not an accurate reflection of the route map.

Once you have configured the route map, you apply it to the desired BGP peer. You can apply it in either direction; import affects routes you are receiving, export affects routes you are advertising.

Unlike other configuration changes, which take effect immediately, changes to eBGP peers do not take effect until an administrator manually resets the peer connections. This allows you to make changes to the configuration without affecting the stability of your BGP network; you can wait to apply the changes until a scheduled maintenance window.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

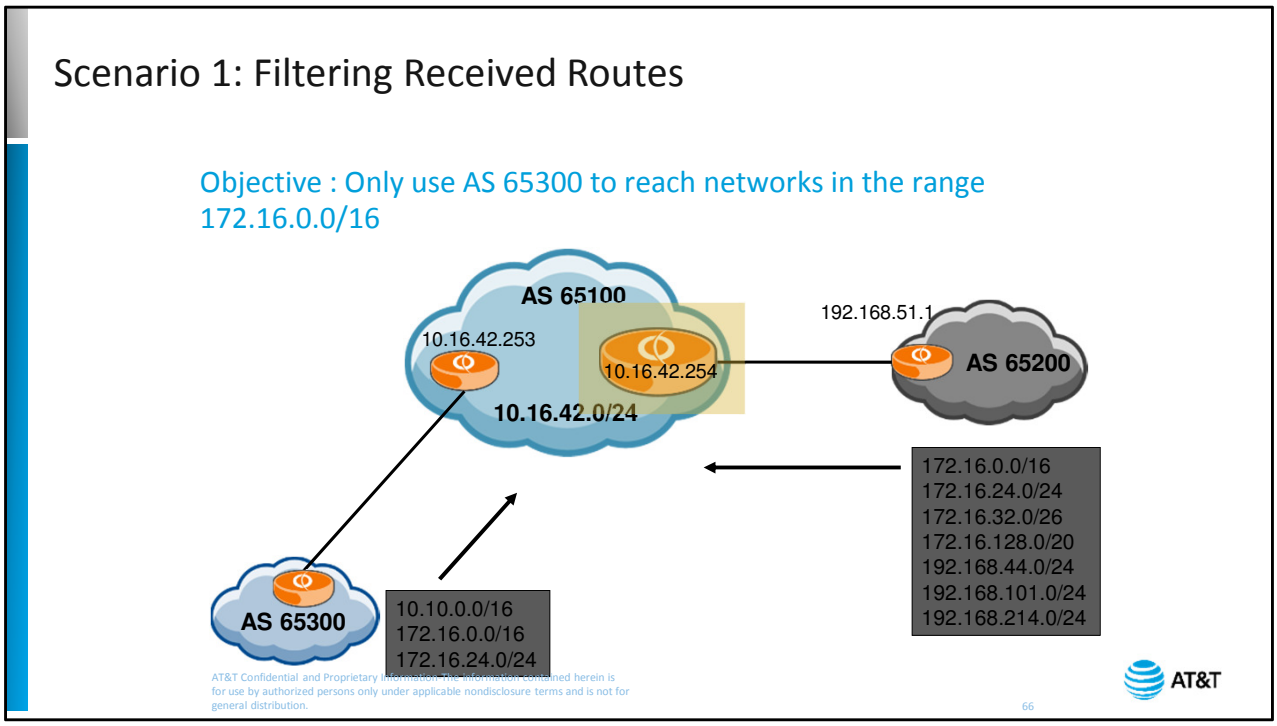

We will create the configuration for this scenario on the next screen. Our policy is to use AS 65300 to reach networks in the prefix range 172.16.0.0/16. We are receiving advertisements within this network range from both AS 65200 and AS 65300, so we want to filter out these networks from AS 65200.

We will place our configuration on the vRouter directly connected to AS 65200.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 1: Configuration

```
[edit]
vyatta@AS65100# edit policy route route-map Deny172.16.0.0
[edit policy route route-map Deny172.16.0.0]
vyatta@AS65100# edit rule 10
[edit policy route route-map Deny172.16.0.0 rule 10]
vyatta@AS65100# set match ip address prefix-list Match172.16.0.0
[edit policy route route-map Deny172.16.0.0 rule 10]
vyatta@AS65100# set action deny
[edit policy route route-map Deny172.16.0.0 rule 10]
vyatta@AS65100# up
[edit policy route route-map Deny172.16.0.0 rule]
vyatta@AS65100# up
[edit policy route route-map Deny172.16.0.0]
vyatta@AS65100# set rule 20 action permit
[edit policy route route-map Deny172.16.0.0]
vyatta@AS65100# top
[edit]
vyatta@AS65100# commit
[edit]
vyatta@AS65100# exit
Warning: configuration changes have not been saved.
exit
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's create our configuration. We begin by creating our prefix list and giving it a meaningful name. The name in the example is *Match172.16.0.0*.
Next, we edit the rule in the prefix list.
We set the prefix we want to match, then set the length to less than or equal to 32. This will match *all* prefixes within the specified range.
Finally, we set the action to *permit* and use the `top` command to exit out of the prefix list.
Next, we create our route map, again giving it a meaningful name. We edit rule 10 and set the match criteria. We want to match IP addresses as defined in the prefix list we just created.
For the action, we *deny* routes that match the filter.
We use the *up* command twice to go up to the *route-map* level of the hierarchy.
Now we create rule 20 and set the action to *permit*. If we do not set a filter for rule 20, all routes will match the rule. The end result is that we permit all routes except those that match rule 10.
We use the `top` command to exit the route map.
Next, we commit our changes and exit Configuration mode so we can test our filters before applying and activating them.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Testing and Completing the Configuration

```
vyatta@AS65100:~$ show ip bgp prefix-list Match172.16.0.0 | grep 65200
*> 172.16.0.0       192.168.51.1            0             0 65200 65478 65398
65219 i
*> 172.16.24.0/24   192.168.51.1            0             0 65200 65478 65398
65219 i
*> 172.16.32.0/26   192.168.51.1            0             0 65200 65478 65398
65219 i
*> 172.16.128.0/20  192.168.51.1            0             0 65200 65478 65398
65219 i
vyatta@AS65100:~$
```

```
[edit]
vyatta@AS65100# set protocols bgp 65100 neighbor 192.168.51.1
address-family ipv4-unicast route-map import Deny172.16.0.0
[edit]
vyatta@AS65100# commit
[edit]
vyatta@AS65100# exit
Warning: configuration changes have not been saved.
exit
vyatta@AS65100:~$ reset ip bgp external
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Before we apply the route map to our neighbor,
let's check to see which routes in the BGP table match our filter. We use the `show ip bgp prefix-list` command, checking specifically for routes coming from AS 65200. The output shows the routes that match. These are the routes that will be removed from the routing table once our route map is activated.
We have verified that these are the routes we want to affect, so we go back into Configuration mode and apply the filter to routes we are importing from our eBGP neighbor at 192.168.51.1. This is the neighbor in AS 65200.
We commit our changes and exit Configuration mode.
In Operational mode, we activate the filter by clearing the external BGP connection. Note that we did not save our changes. Until we know that the filter is working as desired, we do not want to save the changes permanently. This allows us to go back to a known-working configuration by loading the unchanged *config.boot* file.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 1: Before and After

Before Applying the Filter

```
vyatta@AS65100:~$ show ip bgp
BGP table version is 0, local router ID is 10.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop         Metric LocPrf Weight Path
*>i10.10.0.0/16     10.16.42.253     0      100    0      65300 i
* i10.16.42.0/24    10.16.42.253     0      100    0      i
*>                  0.0.0.0          0                     32768 i
*>i172.16.0.0       10.16.42.253     0      100    0      65300 65432 65321 65219 i
*>i172.16.24.0/24   10.16.42.253     0      100    0      65300 65432 65321 65219 i
*> 192.168.44.0     192.168.51.1     0             0      65200 65477 i
*> 192.168.101.0    192.168.51.1     0             0      65200 65477 i
*> 192.168.214.0    192.168.51.1     0             0      65200 65477 i

Total number of prefixes 7
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Before we applied the filter, our BGP table included several routes in the prefix range 172.16.0.0/16. For some of them, we have multiple entries.
For example, 172.16.0.0/16 has two entries – one from our external BGP neighbor at 192.168.51.1, and one from our internal BGP neighbor at 10.16.42.253. The active route is the route received via eBGP.
After we apply the filter, we do not have any 172.16 routes from our neighbor in AS 65200. Looking again at the route to 172.16.0.0/16, we see that the only route is coming from our iBGP neighbor. The 172.16 routes from AS 65200 are blocked.
However, we are still receiving some routes from AS 65200, so we know that the neighbor is up and working properly.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 2: Blocking Private Networks

Objective: Remove all private networks from BGP outgoing and incoming advertisements

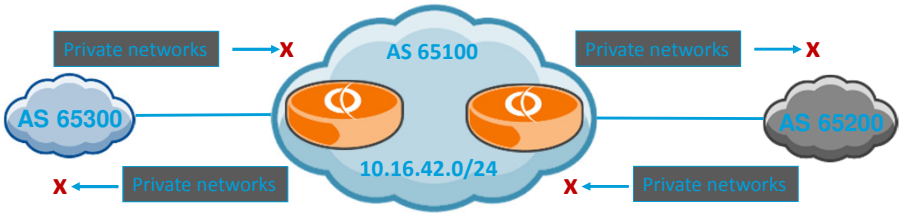

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In this scenario, our AS has two eBGP connections to two different ISPs. Our goal is to ensure that we are neither advertising, nor accepting advertisements, for private networks. Configuring filters to block these networks is a real-world best practice and should be implemented on all eBGP connections.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 2: Configuration

```
policy {
  route {
    prefix-list PrivateNets {
      rule 10 {
        action permit
        le 32
        prefix 10.0.0.0/8
      }
      rule 20
        action permit
        le 32
        prefix 172.16.0.0/12
      }
      rule 30
        action permit
        le 32
        prefix 192.168.0.0/16
    }
}
```

```
route-map BlockPrivateNets {
    rule 10 {
       action deny
       match {
            ip {
               address {
                   prefix-list PrivateNets
               }
    }
    rule 20 {
        action permit
```

```
Protocols
  bgp 65100{
    address-family {
        ipv4-unicast {
}
    neighbor 192.168.51.1 {
       address-family {
           ipv4-unicast {
               route-map {
                 import BlockPrivateNets
                 export BlockPrivateNets
       }
       remote-as 65200
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The trick to this filter is to make sure you have configured the prefixes correctly.
Prefixes in network 10 are simple – we want to filter anything that begins with 10. So we set the prefix itself to 10.0.0.0/8, then add the `le 32` parameter to match all prefixes less than or equal to 32 bits in length. Remember, if we did not include the `le 32` parameter, then the only thing we would match is 10.0.0.0/8 exactly.
To capture the range between 172.16. and 172.31, we configure a prefix with a 12 bit mask, then match all prefix lengths within the range.
To capture all networks beginning with 192.168, we configure a prefix mask of 16, then match all prefixes in the range.
Next, we configure the route map to remove all the private networks but allow all other networks.
Finally, we apply this to our eBGP neighbor, both on imported routes and exported routes.
The example shows the configuration for one of our routers. We then have to repeat these steps on all routers running eBGP.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 3: No Transit Routing

.

Objective: Only advertise local network in eBGP updates; do not advertise any other BGP routes

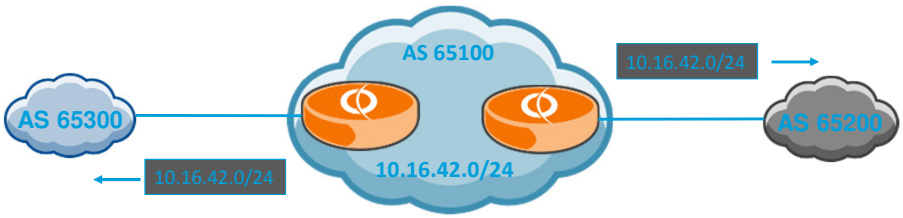

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In this scenario, our network again has two eBGP connections to two different ISPs. Because we are running iBGP to exchange information across our network, we are by default relaying advertisements for networks outside our AS. This means that our network could be used as a transit domain; traffic from other networks destined for other networks could cross ours, just as if we were an ISP. If we are an enterprise, we want to prevent this.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 3: Configuration Using Prefix Lists

```
Policy {
   route {
     prefix-list MyNetwork {
       rule 10 {
           action permit
           prefix 10.16.42.0/24
       }
     }
     route-map MyNetOnly {
         rule 10 {
             action permit
             match {
                 ip {
                     address {
                        prefix-list MyN
                     }
```

```
protocols
    bgp 65100 {
        address-family {
            ipv4-unicast {
             }
         }
        neighbor 192.168.51.1 {
            address-family {
                ipv4-unicast {
                    route-map {
                        export MyNetOnly
                    }
                }
            }
            remote-as 65200
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

You can address this problem with a prefix list.
Configure a prefix list that matches your network or networks,
Then create a route map that only permits the networks that match your prefix list. The default *deny* action will drop all other routes.
Then apply that route map to routes exported to your eBGP peers. Again, the example shows the configuration on one of the routers, but you need to replicate this configuration on all your eBGP devices.
There is one downside to this configuration. If you expect to be adding prefixes to your network over time, you will have to remember to add the prefixes to your prefix list.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 3: Configuration Using AS-Path

```
policy {
    route {
      as-path-list MyASOnly {
          rule 10 {
             action permit
             regex ^$
            }
        }
      route-map MyASOnly {
          rule 10 {
             action permit
             match {
                    as-path MyASOnly
                 }
            }
```

```
protocols
    bgp 65100 {
      address-family {
           ipv4-unicast {
           }
      }
      neighbor 192.168.51.1 {
         address-family {
             ipv4-unicast {
                 route-map {
                    export MyASOnly
                 }
             }
         }
         remote-as 65200
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Another solution is to use an as-path list.
We set the AS path list to match an empty string. This matches any route we are originating, regardless of the prefix. Any routes we have received via BGP will have at least one AS number in the path, so they will not match this expression.
Again, we create a route map with a single permit rule. In this case, it permits routes that match our AS path list.
Finally, we apply this to routes exported to our eBGP peers.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 4: Set Maximum AS-Path Length

Goal: Only accept routes from AS 65200 that have an AS-path length of 2 or less

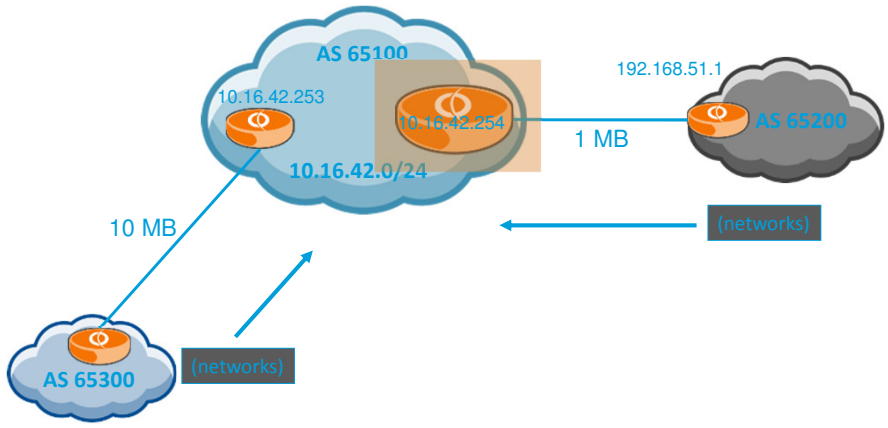

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In this scenario, our goal is to use AS 65200 only for “nearby” destinations; that is, for networks that are either being sourced by AS 65200, or by a directly-connected neighbor of AS 65200. We want to do this because our link to AS 65200 is relatively low bandwidth, so we want to limit the amount of traffic crossing that link. We cannot feasibly use a prefix list for this objective; we will use a regular expression to “measure” the length of the AS path list and only allow routes with a path that has two AS numbers or less.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 4: Configuration

```
policy {
  route {
     as-path-list ShortPaths {
        rule 10 {
            action permit
            regex "^65200(_[0-9]+){0,1}$"
        }
     }
     route-map AS65200_in {
        rule 10 {
            action permit
            match {
                as-path ShortPaths
            }
        }
```

```
protocols
    bgp 65100 {
      address-family {
          ipv4-unicast {
          }
      }
      neighbor 192.168.51.1 {
         address-family {
             ipv4-unicast {
                 route-map {
                    import AS65200_in
                 }
             }
         }
         remote-as 65200
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The regular expression we need to meet our requirements is a bit complex, so let's take a moment to look at it closely.

To match an AS path that contains only AS 65200, we use the carat (^) which signifies the beginning of the string, then the AS number 65200, then the dollar sign ($) which signifies the end of the string.

But we want to match AS 65200 by itself *and* AS 65200 followed by one other AS number. We do not know what that second AS number might be, so we configure the regular expression to accept a range of options.

First, we know we need a space between 65200 and any subsequent AS number, so we add the underscore (_) character to represent the space.

The AS number could be any string of digits from one to five characters long. We use [0-9] in brackets to represent this allowable range of characters – in other words, any number.

We add the plus sign (+) which means that the expression will match one or more occurrences of the preceding pattern. In other words, match one or more occurrences of a number. Now, remember that we want this regular expression to also match 65200 by itself.

So, we enclose the part of the expression that represents the second AS number in parenthesis (),

Then add a limiter that specifies how many times the pattern inside the parenthesis can occur. The pattern can occur zero times – in other words, just as 65200 – or it can occur

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

one time – in other words, AS 65200 followed by some other AS number.
We can now use that regular expression in our as path list.
We use the AS path list in our route map, permitting only routes that match the AS path list, then apply it to routes we are importing from our neighbor in AS 65200.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Scenario 4: Before and After

Before the Filter

```
vyatta@AS65100:~$ show ip bgp | grep 65200
*> 172.16.0.0       192.168.51.1             0     0 65200 65478 65398 65219 i
*> 172.16.24.0/24   192.168.51.1             0     0 65200 65478 65398 65219 i
*> 172.16.32.0/26   192.168.51.1             0     0 65200 65478 65398 65219 i
*> 172.16.128.0/20  192.168.51.1             0     0 65200 65478 65398 65219 i
*> 192.168.44.0     192.168.51.1             0     0 65200 65477 i
*> 192.168.101.0    192.168.51.1             0     0 65200 65477 i
*> 192.168.214.0    192.168.51.1             0     0 65200 65477 i
vyatta@AS65100:~$
```

After the Filter

```
vyatta@AS65100:~$ show ip bgp | grep 65200
*> 192.168.44.0     192.168.51.1             0             0 65200 65477 i
*> 192.168.101.0    192.168.51.1             0             0 65200 65477 i
*> 192.168.214.0    192.168.51.1             0             0 65200 65477 i
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Before we apply the filter, our BGP table has several routes from our peer in AS 65200. We used the Linux `grep` command to limit the output to only those lines that include 65200. Some of these paths are longer than we want.
Once we apply the filter, the BGP table only includes the shorter routes.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Combining Policies

Each peer can have only one policy in each direction

Order is important!

Rule 10: allow only short routes (2 AS numbers)

Rule 20: block network 172.16.0.0/16

Route: 172.16.0.0/16 – 65200 65432

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In a production environment, chances are your policies are going to be more complicated than the examples we have presented here.
Remember that each BGP peer can only have one policy applied in each direction. This means that your route maps will most likely have multiple rules, and the order in which those rules are processed are important.
Let's combine two of our scenarios to illustrate this point. We use rule 10 from Scenario 1, which blocks all routes from prefix 172.16.0.0/16, with rule 20 from Scenario 4, which allows routes with two AS numbers in its AS-path. If we place both rules in a single route map in the order shown, what will happen to the route 172.16.0.0/16 with two AS numbers in the AS path?
The route will match rule 10, and get dropped. Even though it also matches rule 20, we never check that rule.
If we swap the order of the rules, then look at the same route,
It matches the first rule, and this time, the route is allowed into the routing table.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying Operations

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Displaying Prefix Lists

Display configured prefix lists

`show ip prefix-list [details | summary]`

– Operational mode command

```
vyatta@AS65100:~$ show ip prefix-list
ip prefix-list Match172.16.0.0: 1 entries
     seq 10 permit 172.16.0.0/16 le 32
ip prefix-list Match21: 1 entries
     seq 10 permit 172.16.21.0/24
vyatta@AS65100:~$ show ip prefix-list detail
Prefix-list with the last deletion/insertion: Match172.16.0.0
ip prefix-list Match172.16.0.0:
   count: 1, range entries: 0, sequences: 10 - 10
    ripd:
     seq 10 permit 172.16.0.0/16 le 32 (hit count: 0, refcount: 0)
    ripngd:
     seq 10 permit 172.16.0.0/16 le 32 (hit count: 0, refcount: 0)
    ospfd:
     seq 10 permit 172.16.0.0/16 le 32 (hit count: 0, refcount: 0)
    ospf6d:
     seq 10 permit 172.16.0.0/16 le 32 (hit count: 0, refcount: 0)
<Truncated Output>
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In addition to the BGP process commands we have been using to verify our lists, we have additional `show` commands specific to the routing policies themselves. These commands are useful if you do not have access to Configuration mode.

The first is `show ip prefix list`. The output shows all prefix lists configured on the device.

When you add the `detail` parameter to the command, the prefix lists display with information per routing protocol.

You can see the *hit count* (how many routing entries matched each rule in the list), and *refcount* (how many routes were evaluated against the rule).

There is also a summary version of the command that does not display the rule details, just the number of rules in each named list.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Displaying the AS-path List

- Display AS-path list
  `show ip as-path-access-list`
  - Operational mode command

```
vyatta@AS65100:~$ show ip as-path-access-list
AS path access list AS65200
    permit 65200(_[0-9]+){0,1}$
vyatta@AS65100:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The output for checking the AS path list is simpler; it displays the list name and rules. The summary and detail options are not available for this command.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Displaying the Route Map

Display all route maps configured on the system
`show route-map`

- Operational mode command

```
vyatta@AS65100:~$ show route-map
route-map Deny172.16.0.0, deny, sequence 10
  Match clauses:
    ip address prefix-list Match172.16.0.0
  Set clauses:
route-map Deny172.16.0.0, permit, sequence 20
  Match clauses:
  Set clauses:
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The `show route-map` command displays the configured route maps on the system. The output displays the configuration for each rule, the overall action for the rule – *permit* or *deny,* the name of the match filter list, and any `set` commands associated with the rule.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Reviewing BGP `show` commands

`show ip bgp`	Displays BGP table
`show ip bgp summary`	Displays summary information for BGP neighbors
`show ip bgp neighbors address`	Output includes information on applied route maps
`show ip bgp neighbors address [advertised-routes | received-routes]`	Shows just the routes that have been received or sent to the specified neighbor

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When applying filters to BGP operations, you will most often check what is happening with the BGP process rather than the list configuration itself. We have used these commands throughout this course, but include them here for a brief review.

`show ip bgp` shows the contents of the BGP table. You can use this command to check the effects of your policies once they have been activated.

`show ip bgp summary` displays the state of your BGP neighbors. Things to check here are the status of your neighbor connection, as well as the number of prefixes you have learned from a neighbor.

If you do not have access to Configuration mode and need to see whether a filter's been applied to a BGP neighbor, use the `show ip bgp neighbors` command, followed by the IP address of the neighbor. The output includes information on applied filters.

To check the specific routes you are advertising or receiving from a neighbor, add the keywords `advertised-routes` or `received-route` to the `show ip bgp neighbors` command.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Summary

You should now be able to

- Describe BGP routing policy operations and applications
- Configure BGP address-based and AS-path-based route filters
- Verify filter configuration and operations

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

This concludes the AT&T Vyatta 5600 vRouter BGP Policies module.
You should now be able to:

- Describe BGP routing policy operations and applications
- Configure BGP address-based and AS-path-based route filters
- Verify filter configuration and operations

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

BGP Tuning

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Objectives

- After completing this module, you will be able to
  - Explain how BGP uses route attributes to select preferred routes
  - Modify BGP route attributes to affect route selection
  - Explain the function of the BGP community attribute
  - Use the BGP community attribute to filter BGP routes
  - Describe two methods for adding routes to BGP and select the best method for your network
  - Explain when to use a soft reset

This module assumes that you have completed the *AT&T Vyatta 5600 vRouter BGP Basics* and *BGP Policies* modules or have equivalent experience

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

After completing this module, you will be able to:

- Explain how BGP uses route attributes to select preferred routes
- Modify BGP route attributes to affect route selection
- Explain the function of the BGP community attribute
- Use the BGP community attribute to filter BGP routes
- Describe two methods for adding routes to BGP and select the best method for your network
- And Explain when to use a soft reset

Please note that this module assumes that you have completed the *AT&T Vyatta 5600 vRouter BGP Basics* and *AT&T Vyatta 5600 vRouter BGP Policies* modules, or have equivalent experience configuring BGP on a AT&T vRouter. If you do not meet this prerequisite, you may not get the full benefits of this training module.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

BGP Route Selection and Route Attributes

87

We begin with an examination of how BGP selects routes and how you can modify route selection by altering route attributes.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

BGP Route Attributes

Purpose of BGP is to determine the best path

- Shortest is not necessarily best!
- Multiple attributes for each route can be weighed to select the best path for your domain

Attributes include

- AS_Path
- Origin
- Multi-Exit Discriminator (MED)
- Local Preference
- Community

Not all attributes are forwarded

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The primary purpose of BGP is not to find the shortest path to a given network – it is to find the best path. The shortest path is not always the best path; because we are dealing with the Internet rather than just our own network, “best” can depend on several other factors. To accommodate these several factors, BGP routes include multiple attributes rather than just a single metric. An AS administrator can use these different attributes to select the path that is best for that particular AS.

BGP attributes include the AS-path, which lists the AS numbers in order which trace the path back to the route’s origin.

The Origin attribute, which describes how the route was added to BGP. Origin can be an interior gateway protocol, eBGP, or unknown.

Multi-exit discriminator (MED), indicates the preferred path into an AS to external neighbors when multiple paths exist. This is typically used by service providers to specify the preferred connection to their service.

Local preference defines the preferred exit point from the local AS for a specific route. This is typically set by an enterprise when they want to specify paths for given networks.

The community attribute is used to group routes that share common properties so that policies can be applied at the group level.

Not all attributes are forwarded to other BGP routers. For example, AS-path is always forwarded, but local preference is only forwarded within an AS, and is removed from advertisements to eBGP neighbors.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

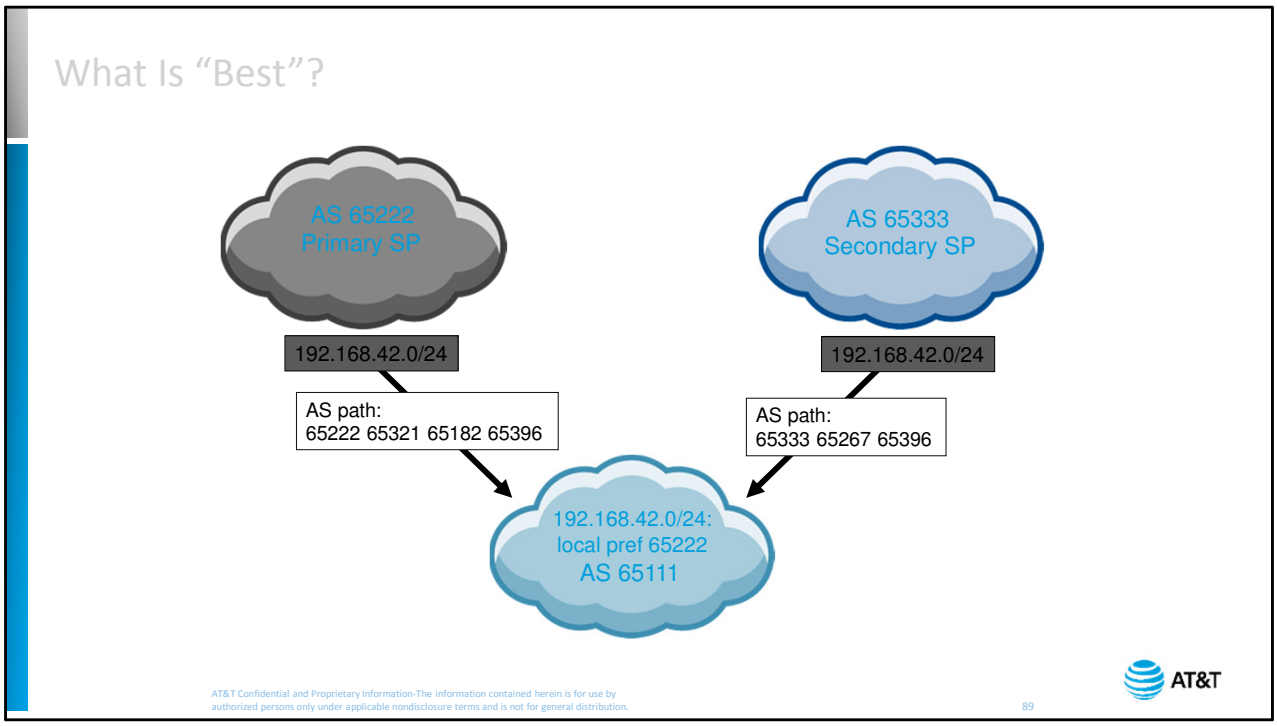

So, if the purpose of BGP is to determine the best path, how do we determine what is best?
In this example, AS 65111 has two peer autonomous systems.
Both peers are advertising a route to network 192.168.42.0. Which path is the best path for AS 65111 to that network?
If we look at the AS path, we see that the shorter path is through AS 65333, so based on just that information, that is the best path.
However, AS 65222 may be our primary service provider, with a faster connection. AS 65333 is our secondary service provider and should only be used if paths are not available via the primary.
We can use the local preference attribute to tell our network that the path via 65222 is the preferred path. We set this attribute using BGP routing policies. We look at specific examples later in this module, but in general, you can modify any BGP attribute using policies.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Choosing the Best Path

1. Prefer the highest local preference
2. Prefer the shortest AS path
3. Prefer the lowest origin
4. Prefer the lowest MED value
5. Prefer routes learned via eBGP over iBGP
   - If all from iBGP, prefer routes with lowest IGP metric
     - Prefer shortest cluster length
   - If all from eBGP, go to step 6
6. Prefer routes from peer with lowest peer ID

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Before you begin modifying attributes, you need to know how BGP uses the attributes to determine a best path to add to the routing table. Once you understand this, you can determine which attributes to modify to produce the desired result.

If a BGP router receives the same route from multiple sources, it first considers local preference. The route with the highest local preference is added to the routing table.

If local preference is equal, then the router will prefer the shortest AS path.

If the AS paths are equal, then the router will consider the Origin value. IGP is preferred over eBGP, which is preferred over an unknown origin.

If the origin values are equal, then BGP considers the Multi-Exit Discriminator, or MED, value. This value is set by service providers to indicate the path they prefer you to use to enter their network. Note that local preference is higher than MED; you can override your service provider's preferences.

If the MED is equal, then the router will prefer eBGP routes over iBGP routes. This is not the same as origin; this is the protocol used by the router to receive the route from a BGP peer. Did it come from an eBGP peer in another AS, or from an iBGP peer within the same AS?

If both advertisements arrived via iBGP, then the BGP router considers the internal metric to the iBGP peer.

If the internal metrics are the same, then the BGP router looks to the Cluster Length attribute. You can set the Cluster Length attribute when configuring route reflectors, which we discuss in the *AT&T Vyatta 5600 vRouter iBGP Design* module.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

If all the routes are learned via eBGP, then the router moves on to step 6.
Finally, if all other attributes are equal, the BGP router looks at the peer router ID and chooses the path from the peer with the lowest ID.
If you are an enterprise network engineer, you may wish to modify one or more of the first three attributes to achieve the desired routing policies.
If you are working for a service provider, you may also be modifying the MED value for your customers.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Outbound and Inbound Policies

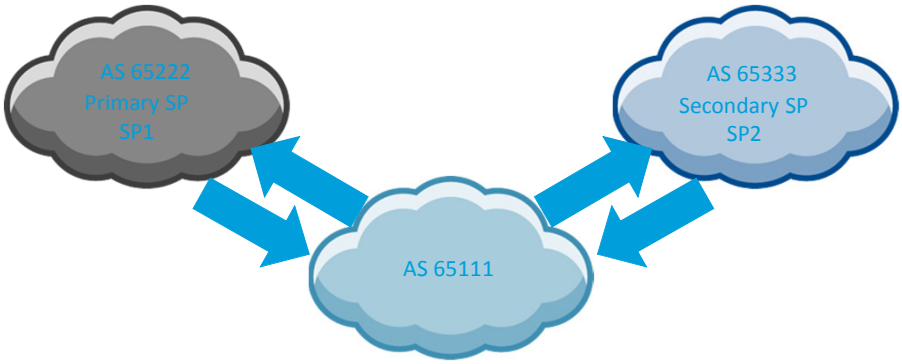

- Control outbound routing
- Influence inbound routing

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When planning your routing policies, you need to consider both directions of traffic.
For traffic that originates inside your autonomous system, you can directly control the path it takes to exit your network, through adjustments to your interior gateway protocol and setting local preference values within your iBGP network.
However, for traffic originating outside your network and destined for your autonomous system, the attributes you set can only influence path choice; other autonomous system administrators can choose to ignore or override your settings once the BGP route information arrives inside their network. We will look at a specific example of outbound and inbound policies on the next screens.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

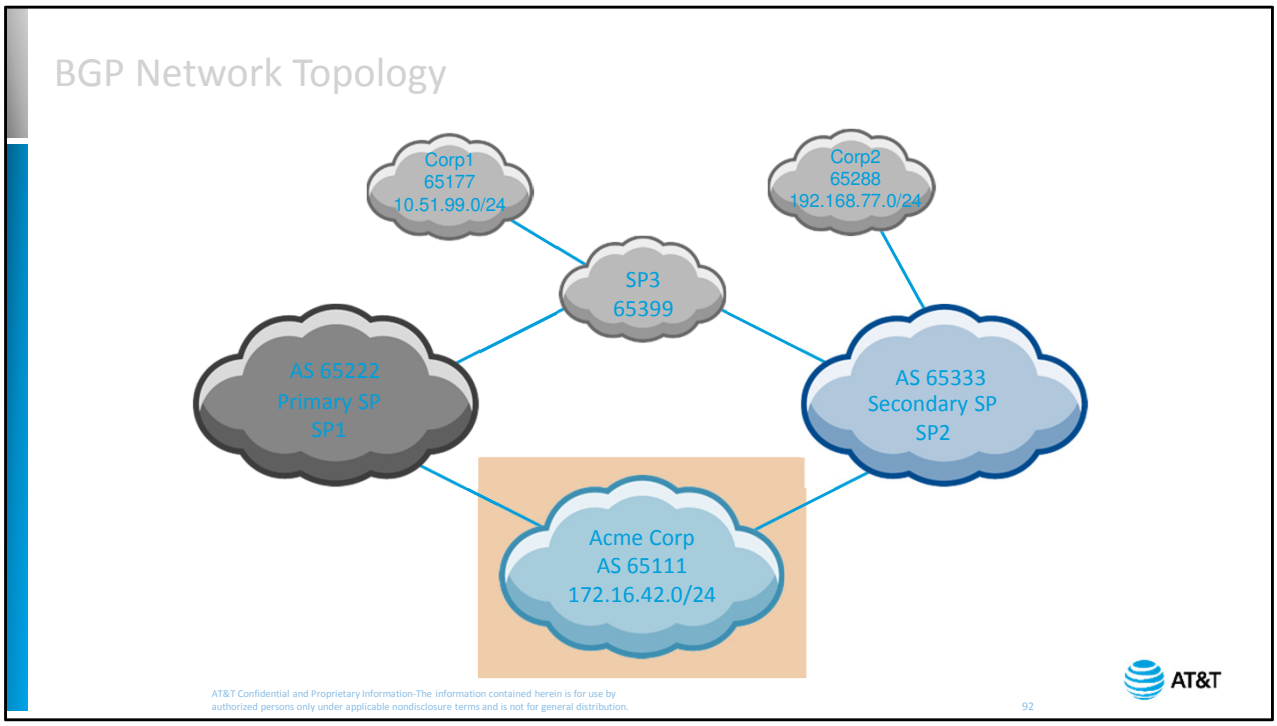

We will use this topology for all the examples in this module.
We are configuring BGP for AS 65111, Acme Corp.
AS 65222 is our primary service provider, and AS 65333 is our secondary service provider.
Upstream, we have two partner companies, Corp1 and Corp2. Corp2 is directly connected to our secondary service provider.
Corp1 is connected to service provider 3 (SP3), which has a connection to each of our service providers.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Setting Outbound Policy

Example 1: Strict primary/secondary

- All traffic exits via SP1 unless connection to SP1 goes down

Policy: set high local preference for all routes from SP1

- Default local preference is 100; higher is better

```
vyatta@AcmeCorp# show policy
route
  route-map 65222Pref {
     rule 10 {
         action permit
         set {
             local-preference 200
         }
     }
 }
```

```
vyatta@AcmeCorp# show protocols bgp 65111
  address-family {
     ipv4-unicast {
     }
}
neighbor 10.1.0.1 {
     address-family {
         ipv4-unicast {
         route-map {
             import 65222Pref
     }
 }
     remote-as 65222
<Truncated Output>
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In our first example,
We look at the configuration for a strict primary/secondary routing policy.
We want all traffic to exit our network using service provider 1 unless the connection to service provider 1 goes down.
To accomplish this, we create a routing policy that sets the local preference for all routes from SP1. The default local preference is 100. A higher value is more preferred, so we set the preference to 200.
First we create the route map for our policy. In the first rule, we match all routes, allow them into the routing table, with an adjusted local preference of 200.
We then apply the rule to our neighbor in AS 65222. Now all routes coming from that neighbor will have a of 200.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Example 1: Verifying Outbound Policy

```
vyatta@AcmeCorp:~$ show ip bgp
BGP table version is 2, local router ID is 172.16.42.251
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.51.99.0/24    10.1.0.1                 0    200      0 65222 65399 65177 i
*                   192.168.42.6             0             0 65333 65399 65177 i
*> 172.16.42.0/24   0.0.0.0                  0               32768 i
*> 192.168.77.0/24  10.1.0.1                 0    200      0 65222 65399 65333 65288 i
*                   192.168.42.6             0             0 65333 65288 i

Total number of prefixes 3

vyatta@AcmeCorp:~$ show ip route bgp
B  *> 10.51.99.0/24 [20/0] via 10.1.0.1, dp0p1p1, 00:13:32
B  *> 192.168.77.0/24 [20/0] via 10.1.0.1, dp0p1p1, 00:13:32
vyatta@AcmeCorp:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When we look at the BGP table, we see the local preference (LocPrf) has been set for routes from our neighbor in AS 65222.
Looking at the best route indicator, or angle bracket (>), we see that the routes from AS 65222 are the best routes, even though the AS path for network 192.168.77.0 is longer.
We can further verify that the routes from that neighbor are the active routes by looking at the routing table. The routes added to the routing table use 10.1.0.1 as the next-hop address, and that is our neighbor in AS 65222.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Example 2: Setting Outbound Policy

Example 2: Loose primary/secondary

- Use SP2 for networks directly connected to SP2
- All other traffic exits via SP1 unless connection is down

Policy

- Set local preference for SP2 directly-connected networks to 250
  - Use as-path filter to match “65333 <any>” in path
- Set Local Preference for all routes from SP1 to 200

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In this second example, we want to implement a loose primary/secondary policy.
That is, we want to use the path through service provider 2 for all networks directly connected to service provider 2.
All other traffic for all other destinations uses the path through service provider 1 unless the connection is down.
To enforce this policy, We set the local preference for networks directly connected to AS 65333 to 250.
To accomplish this, we examine the AS path and filter out routes that have a path with only one AS number after as 65333.
We continue to use the filter we set previously to alter preference for all the routes coming from service provider 1. Because we have set specific paths from service provider 2 to a higher value, those paths will now be preferred.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Loose Primary/Secondary Outbound Configuration

```
vyatta@AcmeCorp# show policy
 route
   as-path-list 65333Peers {
     rule 10 {
         action permit
         regex "^65333(_[0-9]+){0,1}$"
     }
 }
   route-map 65222P
     rule 10 {
         action perm
         set {
             local-p
         }
     }
 }
```

```
route
  route-map 65333Pref {
     rule 10 {
         action permit
         match {
             as-path 653
         }
         set {
             local-prefe
         }
     }
     rule 20 {
         action permit
     }
  }
```

```
vyatta@AcmeCorp# show protocols bgp 65111
 address-family {
     ipv4-unicast {
     }
}
neighbor 10.1.0.1 {
   address-family {
       ipv4-unicast {
           route-map {
               import 65222Pref
           }
       }
   }
   remote-as 65222
}
neighbor 192.168.42.6 {
 address-family {
       ipv4-unicast {
           route-map {
               import 65333Pref
           }
       }
   }
   remote-as 65333
}
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To match the short paths from AS 65333, we create an as-path-list with the appropriate regular expressions.
We still have the route map from the previous example, and have added a second route map for SP2. The first rule invokes the as-path-list and sets the preference for those matching routes to 250.
The second rule accepts all other routes without altering them.
We then apply the appropriate route maps to each BGP neighbor.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Example 2: Verifying Outbound

```
vyatta@AcmeCorp:~$ show ip bgp
BGP table version is 2, local router ID is 172.16.42.251
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.51.99.0/24    10.1.0.1                 0    200      0 65222 65399 65177 i
*                   192.168.42.6             0             0 65333 65399 65177 i
*> 172.16.42.0/24   0.0.0.0                  0               32768 i
*  192.168.77.0     10.1.0.1                 0    200      0 65222 65399 65333 65288 i
*>                  192.168.42.6             0    250      0 65333 65288 i

Total number of prefixes 3

vyatta@AcmeCorp:~$  show ip route bgp
B  *> 10.51.99.0/24 [20/0] via 10.1.0.1, dp0p1p1, 00:18:32
B  *> 192.168.77.0/24 [20/0] via 192.168.42.6, dp0p1p2, 00:5:48
vyatta@AcmeCorp:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Now when we look at the BGP table, we see that the path through AS 65333 has a local preference (LocPrf) of 250, and is the active route.
The routing table also reflects this change.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Influencing Inbound Routing

Prefer traffic to arrive via SP1

Policy: increase AS path length for advertisement to SP2

- Prepend our AS number 3 times
- Policy also blocks transit routes – permits only AcmeCorp network

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Remembering that we cannot completely control the routing of traffic arriving at our network, we would prefer to have traffic arrive at our network over the connection to SP1.
To implement this policy, we modify the AS path for advertisements to SP2.
This practice is called prepending, and inserts information at the beginning of the AS path.
When you do this, the best practice is to insert your autonomous system number multiple times to increase the path size.
In addition, we need to block transit route advertisements, as we are not a service provider and do not want traffic transiting our network.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Influencing Inbound Routing Configuration

```
vyatta@AcmeCorp# show policy
 route
   as-path-list
       rule 10 {
         action
         regex "
     }
   }
   prefix-list M
       rule 10 {
         action
         prefix
     }
   }
   route-map 652
       rule 10 {
         action
         set {
             loc
         }
     }
```

```
route-map 65333Pref {
    rule 10 {
       action permit
       match {
           as-path 65333P
       }
       set {
           local-preferen
       }
    }
    rule 20 {
       action per
    }
}
route-map MyNetOn
    rule 10 {
       action per
       match {
           ip {
               ad
    }
}
```

```
route-
    r
}
```

```
vyatta@AcmeCorp# show protocols bgp 65111
 address-family {
     ipv4-unicast {
     }
}
neighbor 10.1.0.1 {
   address-family {
       ipv4-unicast {
           route-map {
               import 65222Pref
               export MyNetOnly
           }
       }
   }
   remote-as 65222
}
neighbor 192.168.42.6 {
 address-family {
       ipv4-unicast {
           route-map {
               import 65333Pref
               export PrependAS
           }
       }
   }
   remote-as 65333
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

We are still adding to our existing configuration.
First, we add a prefix list to match our network. We use this to filter out all other routing advertisements when we send BGP updates to our peers.
Next, we add a new route map that invokes our prefix list and permits the routes without modification. We use this route map for exporting routes to our neighbor in SP1.
The next new route map uses the same match parameter, but adds the AS path information to increase the AS path.
We can see the prepending route map applied to our neighbor in the SP2 network.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying Inbound Routing Configuration

```
vyatta@AcmeCorp:~$ show ip bgp neighbors 192.168.42.6 advertised-routes
BGP table version is 2, local router ID is 172.16.42.251
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 172.16.42.0/24   192.168.42.5             0         32768 65111 65111 i

Total number of prefixes 1
vyatta@AcmeCorp:~$  show ip bgp neighbors 10.1.0.1 advertised-routes
BGP table version is 2, local router ID is 172.16.42.251
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 172.16.42.0/24   10.1.0.2                 0         32768 i

Total number of prefixes 1
vyatta@AcmeCorp:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To verify that the AS path has been altered according to your policy, use the `show ip bgp neighbor` command.

The path we are advertising to our neighbor in AS 65333 has the prepended AS numbers already in the path.

The path we are advertising to our neighbor in AS 65222 does not.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The Community Attribute

101

Next we take a look at the community attribute and how we can use it to tune BGP route choices.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The Community Attribute

Community: a group of prefixes with a common property

- Community attribute identifies the group
- Can use the community attribute in place of prefix-filter or AS-filter to identify routes

Example: SP A uses community attribute to identify routes from peer service providers

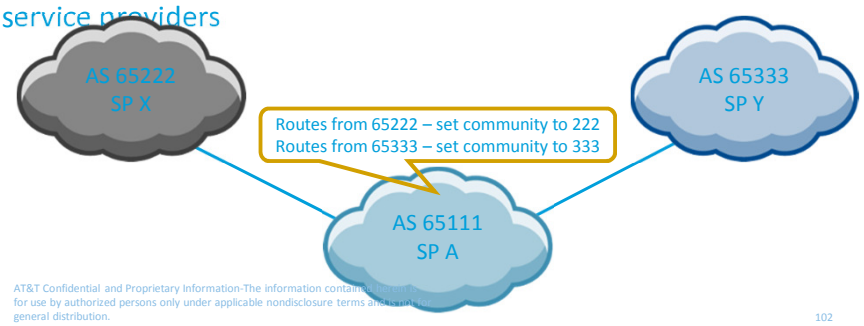

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

A community is a group of prefixes that share some common property and can be configured with the BGP community attribute.
While communities themselves do not alter the BGP decision making process, communities can be used as flags in order to mark a set of routes.
The community attribute can then be used to apply routing policies, just as you can use prefix filters or as-path-lists to identify routes you want to modify.
In this example, service provider A (SP A) is using the community attribute to identify the peer AS from which a route arrives. The service provider can then use the community attribute to apply an attribute, such as local preference.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Community Attribute Options

Community attribute options

`no-export`	Do not advertise to eBGP peers
`no-advertise`	Do not advertise this route to any peer
`Internet`	Advertise this route to the internet community
`local-AS`	Use in confederation scenarios to prevent sending the route outside the local sub-AS

Set attribute to a number string using the *aa:nn* format

- where *aa* is the AS number and *nn* is the community number

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

The community attribute can be set to one of four well-known options. Remember that these attributes are simply flags; for example, if you want to prevent routes flagged as *no-export* from being exported, you have to configure a routing policy to match the attribute then *deny* the route.

You can also configure the community attribute to a number string. The format of the string is the AS number followed by a colon, then the community number.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Using the Community Attribute

Setting the community attribute

```
set policy route route-map name rule num set community string
```

Using the community attribute

- Create community-list

```
edit policy route community-list num rule num
```

  - Set match criteria

```
set regex regular-expression
```

  - Set action

```
set action [permit | deny]
```

- Invoke community list in route-map

```
set policy route route-map name rule num match
community community-list num
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

You set the community attribute as part of a routing policy.
In your route map, after you configure the `match` parameter and the `permit` action, use the `set` command to set the community string, either to one of the well-known values, or to the *AS number-colon-communitynumber* format.
To use the community attribute, you first have to create a filter list called a community list. This is similar in function and format to an as-path-list, although the lists are numbered instead of named.
The match criteria is defined in a regular expression. The action is either *permit* or *deny*. Remember that this action does not permit or deny a route into the routing table. Permitting the route means it matches and will be affected by the settings in the route map; denying a route means it will be passed to the next rule in the route map list.
To invoke the community list filter in a route map statement, use the `match` command, matching the number of the community list you configured.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Using the Community Attribute

Service providers have default local preferences

- Directly-connected customers have preference of 100
- Routes from other SPs have preference of 80

Service providers use the community attribute to set local preference

- For example, if the SP receives a route with the community attribute set to 110, they set the local preference to 110

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look at an application of the community attribute as it applies to our topology.
We first have to look at some of our service provider policies.
By default, directly-connected customer routes are set to a local preference of 100.
Routes learned from other service providers are set to a local preference of 80.
In order to allow customers to adjust these values, the service providers have also set up policies that use the community attribute.
If they receive a route with the community attribute set to 110, they set the local preference to 110.
If the community attribute is 100, they set the local preference to 100,
and so on for communities 80 and 90.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Excerpt from Service Provider Policy

```
vyatta@SP2# show policy
 community-list 100 {
     rule 10 {
         action permit
         description "Match Community 80"
     }
 }
 communit
     rule
         description "Match Community 90"
         regex :90_
     }
<Truncated Output>
```

65111:80
65222:90

- Order is important! If routes have multiple community attributes, lowest will match first

```
route-map CustomerPref {
     rule 10 {
         action permit
         match {
             community {
                 community-list 100
             }
         }
         set {
             local-preference 80
         }
     }
     rule 20 {
         action permit
         match {
             community {
                 community-list 101
             }
         }
         set {
             local-preference 90
<Truncated Output>
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look at the actual service provider policy configuration.
We see here that the service provider has created a community list to match any community attribute ending in colon 80. Remember, the format of the community attribute is an AS number followed by a community number, so we want to match any AS that is associated with the number 80.
We have another community list to match the community attribute 90. Not shown are similar community lists for community attribute 100 and 110.
We invoke these community lists in a route map that actually sets the local preference.
Here, rule 10 invokes the list that matches community attribute 80,
Then sets the local preference to 80.
Rule 20 matches community attribute 90 and sets the local preference to 90. Again, not shown are rules 30 and 40, which sets local preference to 100 and 110, respectively. Then rule 50 permits all remaining routes and sets the preference to 100. The service provider has a similar route map for service provider peers that has a default rule setting the preference to 80.
As with any policy, the rule order is important! A route can have multiple community attributes. We have configured this policy so that routes are set to the lowest matching community string. Remember, a list processes until it finds the first match, then executes the action and exits.
So, if a route arrived with community attributes of 65111:80 and 65222:90,

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

It would match rule 10, have the local preference set to 80, then exit the policy without looking at any other rules. If we wanted to set higher preferences first, we would reverse the order of our route map rules.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

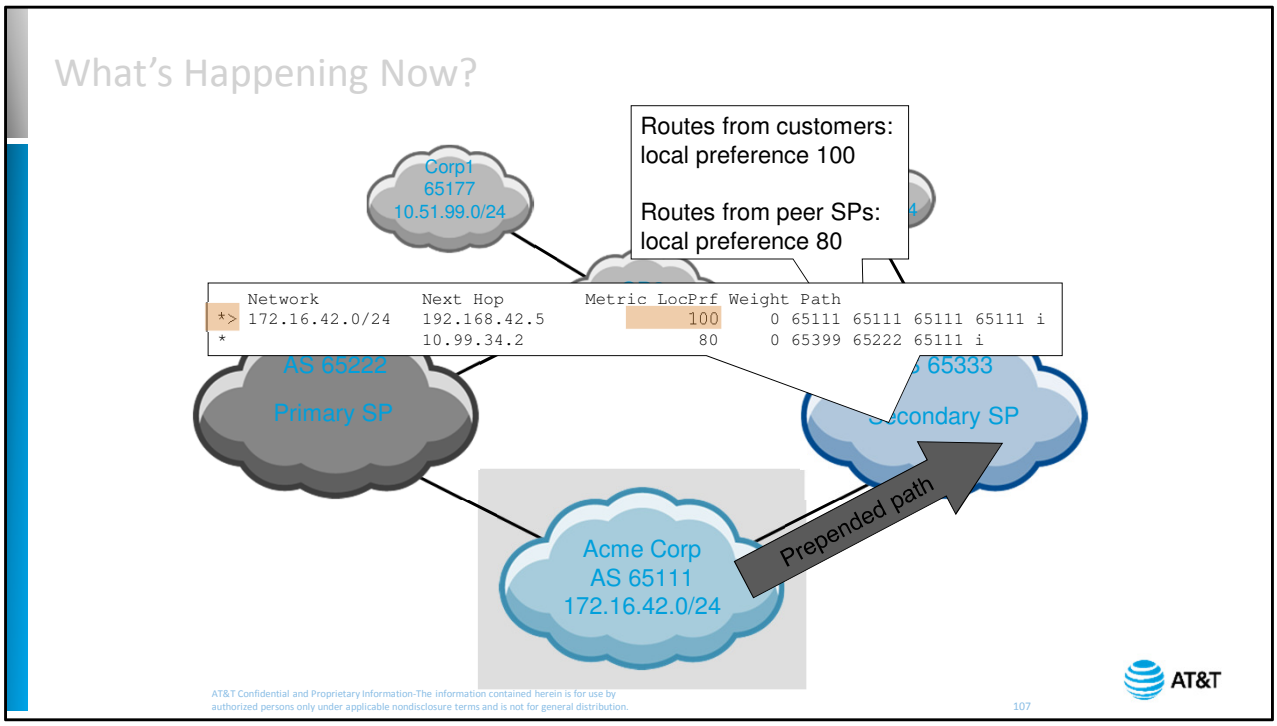

At the end of our last example, we were sending a prepended AS-path to our secondary service provider in order to make that path less preferable than the one through our primary service provider. However, when we analyze link usage we see that we are still receiving a great deal of traffic over that link.

When we consider our service provider's policy, we can see why.

The secondary service provider receives announcements for network 172.16.42.0/24 directly from us, and through service provider 3.

Even though we have prepended the AS path to make it appear longer, and less preferred, the service provider's local preference takes precedence and makes the direct connection the active route. To overcome this, we can take advantage of our service providers' policies and set the community string.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Setting Our Community Attribute

```
vyatta@AcmeCorp# show policy
 route-map AppendAS {
     rule 10 {
         action permit
         match {
             ip {
                 address {
                     prefix-list MyNet
                 }
             }
         }
         set {
             as-path-prepend "65111 65111 65111"
             community 65111:80
         }
<Truncated Output>
vyatta@AcmeCorp#
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To set the community attribute for routes sent to our secondary service provider, we add to the existing route map that prepends the AS path, setting the community attribute to our AS number followed by 80.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying Community

Use `show ip bgp community-info` to display all communities known by the device

- Operational mode command

```
vyatta@AcmeCorp:~$ show ip bgp community-info
Address Refcnt Community
[0xb8a92080] (1) 65111:80
vyatta@AcmeCorp:~$
```

```
vyatta@SP2:~$ show ip bgp community-info
Address Refcnt Community
[0xb90df570] (1) 65111:80
vyatta@SP2:~$
```

Check policy configuration to see what you are setting

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

You can show all communities known by your device, whether they are ones you set, or ones you learned from your BGP peers, with the command `show ip bgp community-info`.

On the AcmeCorp device, we see the community attribute we set – 65111 colon 80 (65111:80).

When we issue the same command on the device in service provider 2's network, we see the same community string.

The only way to check what community strings you are actually setting is to look at your routing policy configuration.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Service Provider 2 – BGP Table

```
vyatta@rtr3:~$ show ip bgp
BGP table version is 2, local router ID is 192.168.51.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.51.99.0/24    10.99.34.2               0     80      0 65399 65177 i
*  172.16.42.0/24   192.168.42.5                 80      0 65111 65111 65111 65111 i
*>                  10.99.34.2                   80      0 65399 65222 65111 i
*> 192.168.77.0     10.99.34.6               0    100      0 65288 i

Total number of prefixes 3
vyatta@rtr3:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Now when we look at the BGP table in service provider 2's network, we see that the preferred path to our network is not the directly connected path. The local preference values are equal, so the router goes to the next criteria in the BGP decision process, which is AS path length. The path through the service provider is shorter than our prepended AS path, so it is preferred.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Adding Routes to BGP

111

We have talked about modifying routes, but we have not yet discussed how to add routes to BGP in the first place, so we will discuss that now.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Option 1: Adding Routes to BGP

- Network statement and blackhole route
  - Fulfills requirement that BGP-advertised route must exist in device routing table

```
vyatta@AcmeCorp# show protocols bgp 65111
  <Truncated Output>
  network 172.16.42.0/24 {
  }
vyatta@AcmeCorp# show protocols static
  route 172.16.42.0/24 {
      blackhole {
      }
```

- Advantages
  - Simple to configure – one network statement, one static route
  - Simple security measure – traffic to unreachable hosts discarded at entry point
  - BGP route not dependent on IGP advertisements

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

One method of adding routes to your BGP network, is using a single statement for your public network address, and a blackhole static route to that network-level address on the edge router.
Using this configuration fulfills the BGP requirement that the route must exist in the device routing table before it can be included in BGP advertisements.
Here we have a single network statement in BGP, and a corresponding static blackhole route.
This approach has several advantages. It is simple to configure. Instead of creating multiple network statements for every possible subnet in your network, you configure a single network statement and a single static route.
It also provides a basic level of network security. Traffic that arrives at your network for addresses that do not exist, or are simply unreachable at the time, are discarded automatically by the vRouter.
Finally, using a static route separates BGP completely from your interior routing protocol operations. No matter what changes occur within your internal network, the BGP advertisement remains stable.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Option 1: Redundancy

- Option 1 breaks if you use VRRP for redundancy

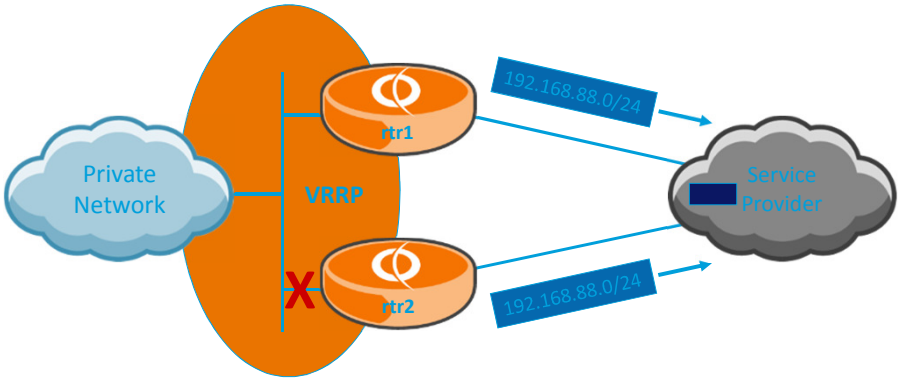

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

However, Option 1 has one shortcoming if used in conjunction with VRRP.
In this network, we are using VRRP on the Ethernet segment connected to our private network.
Each router has its own BGP peer in the service provider network, and is advertising the network based on the static blackhole route in the routing table.
If router 2's (rtr2) link to the private network fails, it is unable to forward traffic to any destination in the private network.
However, it will continue to advertise reachability because the BGP advertisement is not dependent on the state of the internal network.
Service provider can forward traffic to rtr2,
But because no internal hosts are reachable, rtr2 drops the traffic, effectively black-holing the entire network.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Option 2 - Adding Routes to BGP

Use IGP-learned subnets in BGP network statements

Configure BGP aggregate address

```
set protocols bgp as-num address-family ipv4-unicast | ipv6-
unicast aggregate-address prefix/mask [summary-only]
```

- `summary-only` parameter filters out subnets

```
vyatta@AcmeCorp# show protocols bgp 65111
  address-family {
      ipv4-unicast {
          aggregate-address 172.16.42.0/24 {
              summary-only
          }
          network 172.16.42.32/27 {
          }
          network 172.16.42.48/27 {
          }
          network 172.16.42.251/32{
      }
    }
}
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

In a redundant environment, you want to make your BGP advertisements dependent on IGP routes. To do this, you first use subnets in your BGP network statements. These should be subnets that your device is learning via the IGP. It does not have to be every subnet in your network, but should be more than one subnet for stability purposes.

Next, configure an aggregate address for your network. The aggregate is the summary network you want advertised via BGP. As long as the routing table includes a route within the range of the aggregate route, the vRouter will advertise the aggregate in BGP.

If you add the `summary-only` parameter, ONLY the aggregate is advertised in BGP. If you leave off the `summary-only` parameter, then the subnets and the aggregate are all included in BGP advertisements.

So, in our example, we have an aggregate address for our network. We have included the `summary-only` parameter so we do not advertise any of the individual subnets.

We have three subnets in our network. As long as we are learning about one of these three subnets in our routing table, we will advertise the 172.16.42.0 network in BGP.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Option 2: Confirmation

```
vyatta@AcmeCorp:~$ show ip route ospf
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

O>* 172.16.42.32/27 [110/10] via 172.16.42.25, dp0p1p2, 00:18:32
O>* 172.16.42.48/27 [110/30] via 172.16.42.25, dp0p1p2, 00:18:48
<Truncated Output>
vyatta@AcmeCorp:~$  show ip bgp
BGP table version is 0, local router ID is 172.16.42.251
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 172.16.42.0/24   0.0.0.0                  0         32768 i
s> 172.16.42.32/27  0.0.0.0                  1         32768 i
s> 172.16.42.48/27  0.0.0.0                  1         32768 i

Total number of prefixes 3
vyatta@AcmeCorp:~$  show ip bgp neighbors 10.1.0.1 advertised-routes
<Truncated Output>
   Network          Next Hop            Metric LocPrf Weight Path
*> 172.16.42.0/24   10.1.0.2                 0         32768 i

Total number of prefixes 1
vyatta@AcmeCorp:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To verify that our configuration is working, we first look at our interior routing protocol. Of the three network statements we configured, two are currently active and learned via OSPF.

Looking at our own BGP table, we see the aggregate route and the individual subnets are all in the BGP table. The subnets are there due to the network statements, and the aggregate because there is an active route that falls into the summary range.

Note that the subnet routes have the code *s* in front of them. This means that these routes are being suppressed in outgoing advertisements, which we would expect with the `summary-only` parameter enabled.

When we check what we are advertising to our BGP neighbor, we see that only the aggregate is going out.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Soft vs. Hard Reset

We conclude with a discussion on the differences between a soft and a hard reset.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Activating Changes in BGP

Step 1: Commit your changes

- iBGP peers will be updated immediately

Step 2: Reset connection to eBGP peers

```
reset ip bgp [ip-address | all | external | peer-group]
soft [in | out]
```

- `soft` reset retransmits BGP table information
- If the `soft` parameter is not used, a hard reset occurs which drops and re-initiates BGP connection
- Because BGP table can be large, perform resets during maintenance windows

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Activating changes to BGP is a two-step process.
The first step, as with any other configuration change, is to *commit* the changes. This activates them in local memory, and initiates updates to iBGP peers immediately. However, eBGP peers are NOT updated when changes are committed.
The second step is to update your eBGP peers. To do this you need to reset the BGP connection. You can reset a connection to an individual peer, all peers, all external peers, or a peer group.
Optionally add the `soft` parameter, and a direction; i*n* requests a table update from the peers, o*ut* sends a table update to the peers. If you do not specify a direction, the command clears both inbound and outbound connections.
If you add the `soft` parameter, the connected peers only exchange tables.
If you omit the `soft` parameter, then you perform a hard reset; the vRouter drops its existing BGP connection and initiates a new one.
Because BGP tables can be multiple gigabytes of data, exchanging tables between eBGP peers should be done during scheduled maintenance windows.
For other reset options, like resetting just IPv4 or IPv6 BGP connections, please refer to the *AT&T Vyatta 5600 vRouter Software Documentation* on http://businesscenter.att.com.com.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Summary

- You should now be able to
  - Explain how BGP uses route attributes to select preferred routes
  - Modify BGP route attributes to affect route selection
  - Explain the function of the BGP community attribute
  - Use the BGP community attribute to filter BGP routes
  - Describe two methods for adding routes to BGP
  - Explain the different methods to reset BGP

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

This concludes the AT&T Vyatta 5600 vRouter BGP Tuning module.
You should now be able to:

- Explain how BGP uses route attributes to select preferred routes
- Modify BGP route attributes to affect route selection
- Explain the function of the BGP community attribute
- Use the BGP community attribute to filter BGP routes
- Describe two methods for adding routes to BGP
- Explain the different methods to reset BGP

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

iBGP Design

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Objectives

After completing this module, you will be able to

- List the design requirements for implementing iBGP
- Configure route reflectors
- Configure peer groups
- Configure confederations
- Verify correct iBGP functionality

This module assumes that you have completed the *AT&T Vyatta 5600 vRouter BGP Basics* module or have equivalent experience

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

After completing this module, you will be able to:

- List the design requirements for implementing iBGP
- Configure a route reflector
- Configure a peer group
- Configure a confederation
- And Verify correct iBGP functionality

Please note that this module assumes that you have completed the BGP Basics module, or have equivalent experience in configuring BGP on a vRouter.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

iBGP Configuration Review

121

We begin with a quick review of iBGP requirements and basic configuration.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

iBGP Design Requirements

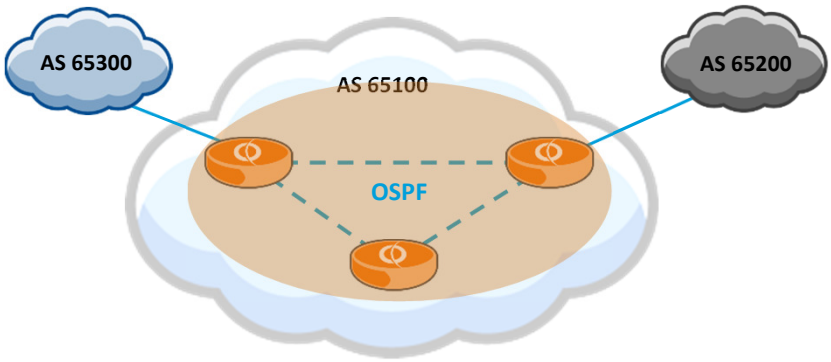

IP connectivity via IGP

- Best practice: peer to loopback addresses

Full mesh with all iBGP peers

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

As we discussed in the BGP Basics module, iBGP implementations have certain design requirements and best practices you should follow.

Unlike eBGP, iBGP peers do not have to be physically connected to each other. However, because BGP is connection-oriented, they do have to have routes to each other so they can establish a connection. Although you can use static routes to meet this requirement, you will generally use an interior gateway protocol like OSPF to build the required routing.

A best practice is to use loopback interfaces for the peer addresses, then advertise the loopback addresses within the IGP. This ensures that iBGP connectivity does not rely on any one physical interface on your device; as long as a route exists to the peer's loopback address, the devices can maintain the iBGP session.

A second requirement is that you must fully mesh all your iBGP peers. This requirement stems from the way iBGP peers propagate routing information, which we will look at on the next screen.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

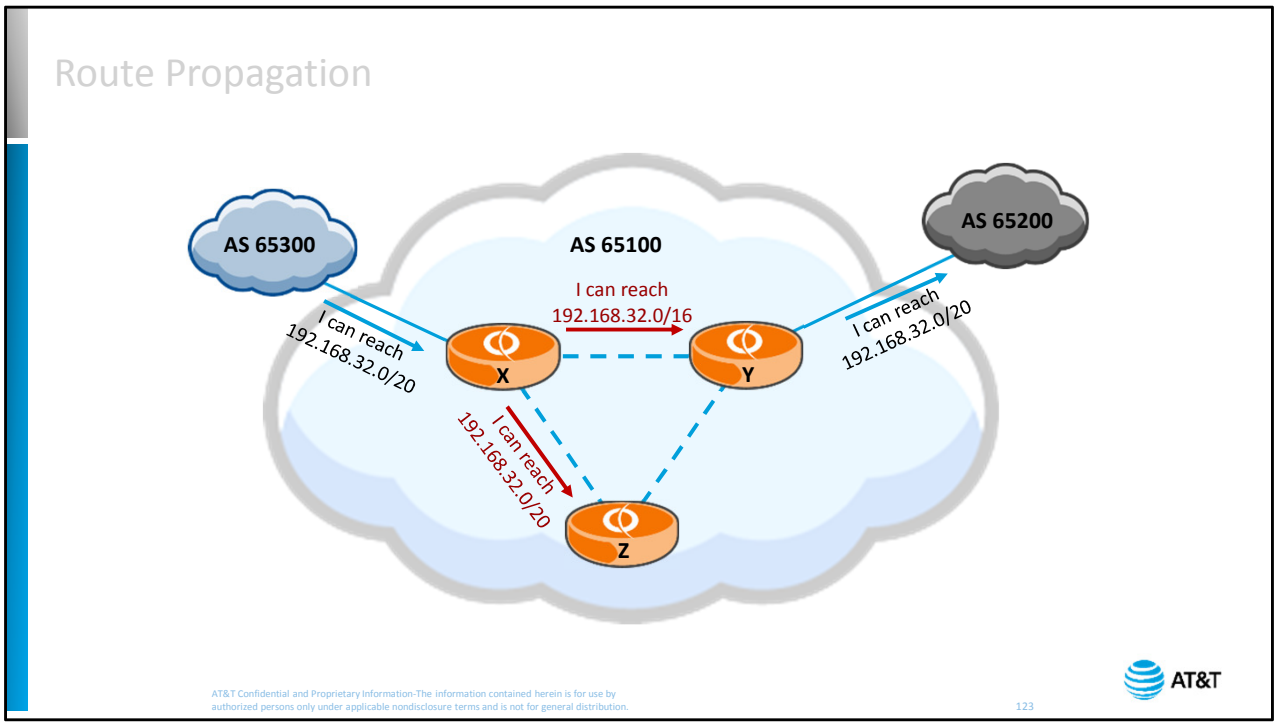

The full mesh requirement has to do with how iBGP propagates routing updates. In order to prevent routing loops, iBGP peers do not forward information learned via iBGP to other iBGP peers. Let's look at example to explain this further.

In our network, we have three routers running iBGP. Router Z has two peers – Router X and Router Y – but we have not configured X and Y to be peers with each other.

Router X receives an eBGP update from its peer in AS 65300. Because this is an eBGP route entry,

Router X will forward the information to all its iBGP peers – in this case, only to Router Z.

Because Router Z received the information via iBGP, it will not forward it to any other iBGP peer, meaning that Router Y has no information about reaching network 192.168.32.0.

By implementing a full mesh, Router X will forward it's eBGP-received information to both Router Z and Router Y.

Router Y can then forward the information to its eBGP peer. Router Y and Router Z will still not propagate the information to each other, thus preventing an internal routing loop within AS 65100.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

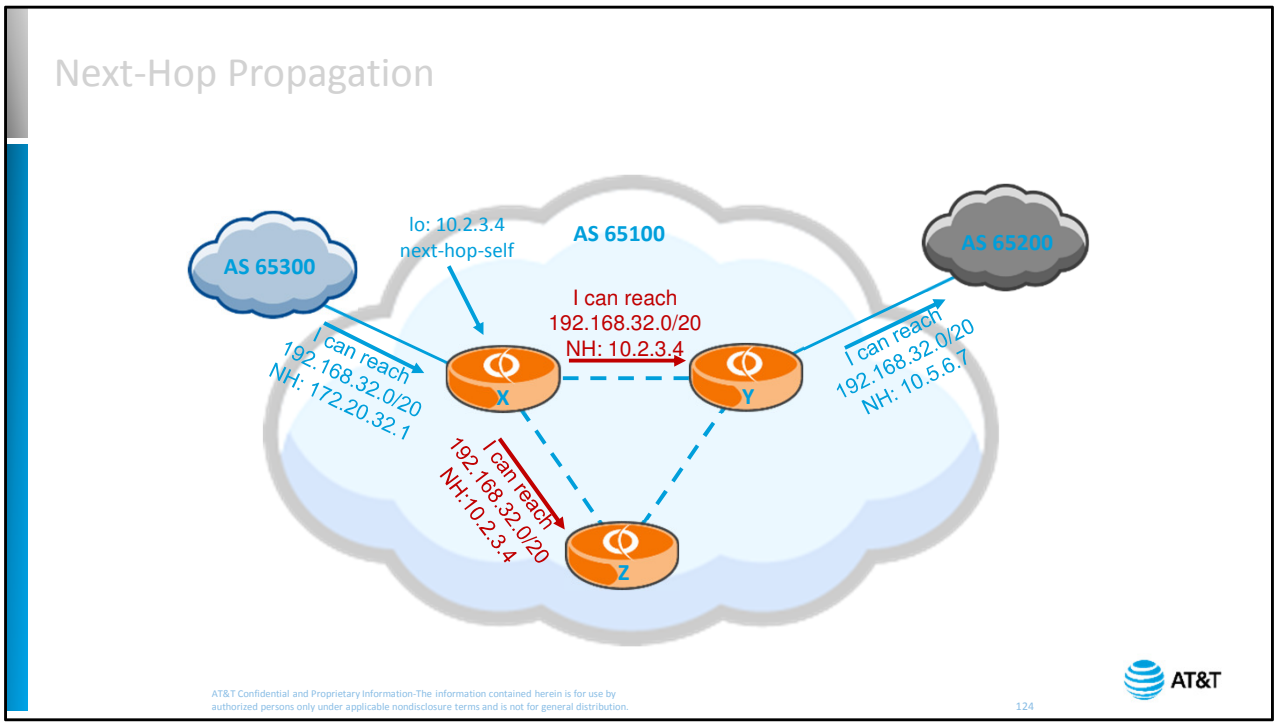

A related issue is how iBGP propagates next-hop information. By default, BGP routers do not modify next-hop information over iBGP links.

When Router X receives a route from its eBGP peer, the next-hop address is the address of the eBGP peer – 172.20.32.1. Because this address is outside the AS , it is unlikely that Routers Y and Z will have a route to this next-hop, so they would mark the route to 192.168.32.0 as unreachable.

If Router X enables the `next-hop-self` parameter for its iBGP peers, it will modify the next-hop address to its own address. Routers Y and Z must have a route to this address, or iBGP would not work at all.

Now that Router Y has an active route, it can forward the route to its eBGP peer. Because the connection is an eBGP connection, Router Y updates the next-hop address field with its own address.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

iBGP Configuration Commands

Enter BGP configuration
```
edit protocols bgp as-number
```

Set router ID
```
set parameters router-id ip-address
```

Add neighbors
```
edit neighbor ip-address
```

- Set local AS number
  ```
  set remote-as as-number
  ```
- Set address to use for updates
  ```
  set update-source local-ip
  ```
- Enable nexthop-self
  ```
  set address-family ipv4-unicast | ipv6-unicast nexthop-self
  ```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

At a minimum, you need to configure the following to enable iBGP.
First, enter the BGP configuration, specifying the local AS number.
Set the router ID if it has not already been set. This is the router ID used for BGP, and will default to the loopback IP address if not configured.
Next, enter the information for the iBGP neighbor.
Set the remote AS number for the neighbor. For an iBGP neighbor, this is the same AS the local device AS number.
Next, set the update source address. This is the address of the loopback address in a well-designed iBGP network.
If this iBGP peer also has eBGP peers, set the `nexthop-self` parameter under the address family in order to allow routes to be properly propagated through the iBGP mesh.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

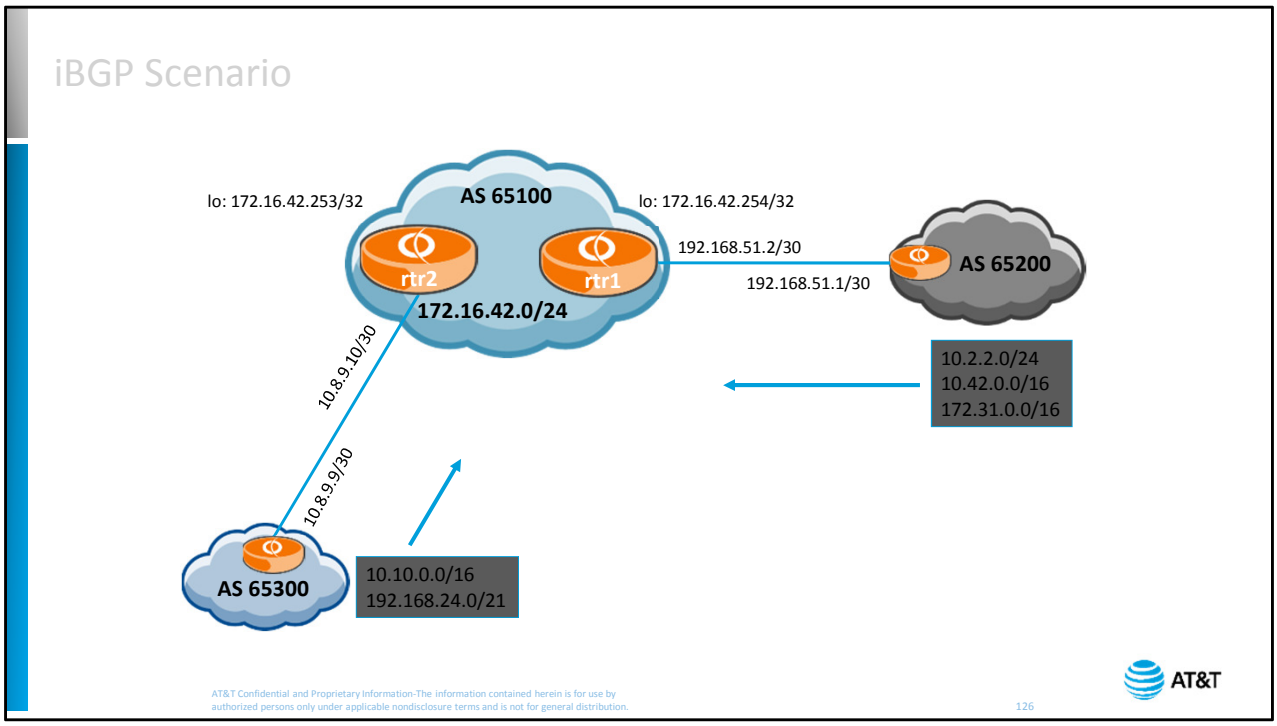

We'll look at the configuration for this scenario on the next screen.

Router 1 (rtr1) and 2 (rtr2) are iBGP peers, using their loopback addresses for iBGP peer addresses.

Router 1 has an eBGP connection to AS 65200, and Router 2 has an eBGP connection to AS 65300. When our iBGP configuration is complete, we should see both internal and external BGP routes – external from the AS directly connected to the device, and internal for the routes from our iBGP peer.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Side-by-Side Configurations for AS65100

```
vyatta@AS65100-2# show protocols bgp 65100
address-family {
     ipv4-unicast {
         network 172.16.42.0/24 {
         }
     }
 }
 neighbor 10.8.9.9 {
     remote-as 65300
 }
 neighbor 172.16.42.254 {
     address-family {
         ipv4-unicast {
               nexthop-self
         }
     }
     remote-as 65100
     update-source 172.16.42.253
 }
 parameters {
     router-id 172.16.42.253
```

```
vyatta@AS65100# show protocols bgp 65100
 address-family {
     ipv4-unicast {
         network 172.16.42.0/24 {
         }
     }
 }
neighbor 172.16.42.253 {
     address-family {
         ipv4-unicast {
               nexthop-self
         }
     }
     remote-as 65100
     update-source 172.16.42.254
 }
  neighbor 192.168.51.1 {
     remote-as 65200
 }
 parameters {
     router-id 172.16.42.254
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look at the configurations for the two vRouters inside AS 65100.
We can see our iBGP peer configuration, including the `nexthop-self` parameter. We included this because the peers also have eBGP peers, and we do not have any internal routes directly to those eBGP peer addresses.
We also see the eBGP peer configuration.
Note that this only shows the BGP configuration. Remember the first requirement for iBGP – that the peers have IP connectivity to each other. We are running OSPF inside our network in order to provide routing between the loopback IP addresses.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

eBGP and iBGP routes

```
vyatta@AS65100:~$ show ip bgp
BGP table version is 1, local router ID is 172.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.2.2.0/24      192.168.51.1             0             0 65200 i
*>i10.10.0.0/16     172.16.42.253            0    100      0 65300 i
*> 10.42.0.0/16     192.168.51.1             0             0 65200 i
* i172.16.42.0/24   172.16.42.253            0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 172.31.0.0       192.168.51.1             0             0 65200 i
*>i192.168.24.0/21  172.16.42.253            0    100      0 65300 i

Total number of prefixes 6
vyatta@AS65100:~$0
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look closely at the BGP table now.
Routes that we have learned through iBGP are marked with an "i" before the network information.
We can also see that our iBGP neighbor is advertising the same network we are – 172.16.42.0. However, since we are originating a route to the same network, our originated route is considered the best path as designated by the angle bracket (>).

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Route Reflectors

Now we'll take a look at route reflectors and how to configure them.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

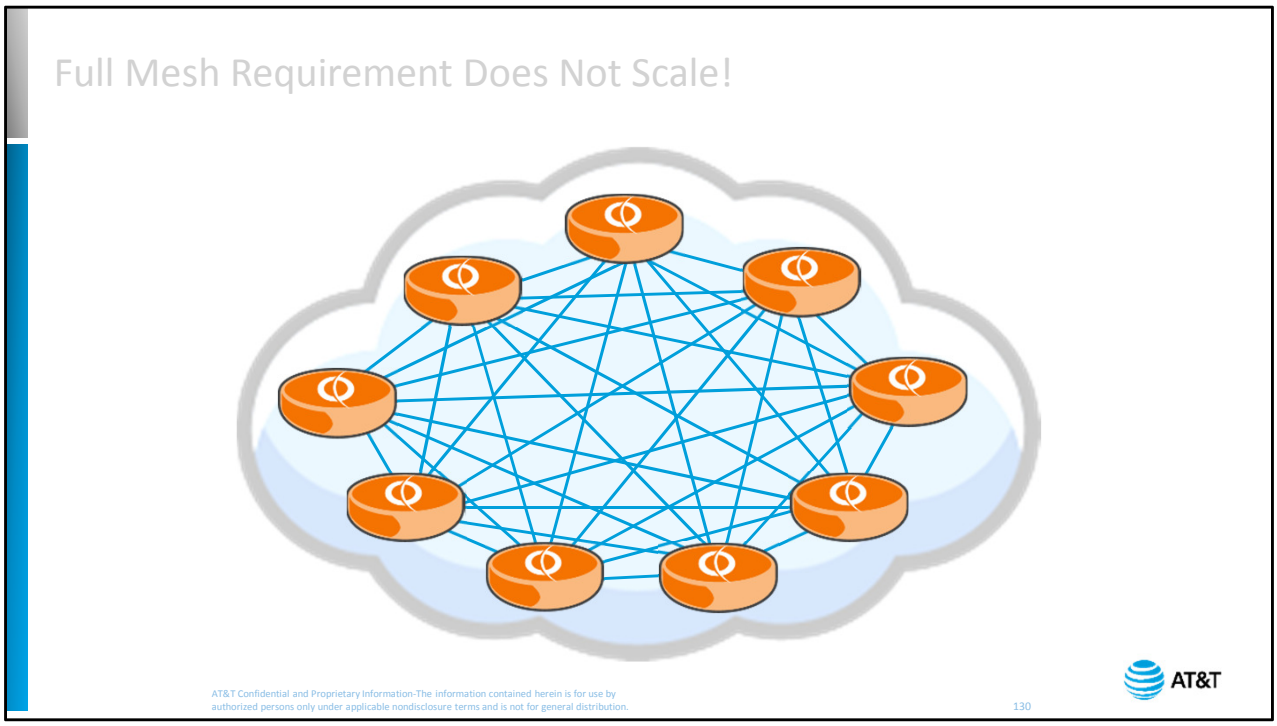

If you only have two or three iBGP peers, configuring a full mesh is not problematic, either from an administrative standpoint, or from a device resource standpoint.
However, if you are a service provider, chances are you have more than three iBGP routers – many more. And you still need to configure a full mesh of devices in order for iBGP to function properly.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Solution 1: Route Reflectors

Use a “concentrator” router to relay iBGP updates

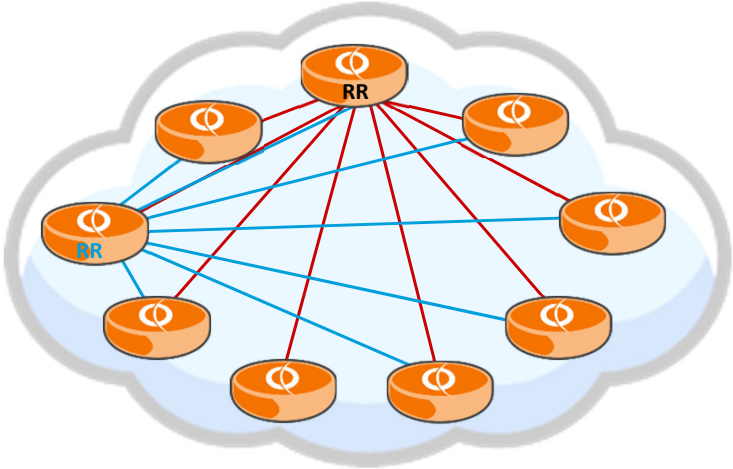

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

BGP provides two solutions to the full mesh problem.
The first is to use a route reflector.
A route reflector acts as a concentrator router for all iBGP updates. Each iBGP device peers with the route reflector and only the route reflector. A group of routers peered with a route reflector is called a cluster. The reflector forwards all iBGP routes from one cluster member to all other cluster members. The cluster members only need to maintain a single session, but you need to make sure the route reflector device has sufficient memory and processor resources in order to maintain all the peer sessions.
Or, you can use multiple route reflectors to provide redundancy while still avoiding a full mesh configuration.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuring Route Reflectors

Add client statements to route reflector device

```
set protocols bgp as-num neighbor ip-address address-family
ipv4-unicast | ipv6-unicast route-reflector-client
```

On clients, configure route reflector as neighbor

- No additional commands needed

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To configure route reflection, you only need to modify the configuration on the route reflector device. Add the `route-reflector-client` parameter to the neighbor configuration for each client device peering with the route reflector.
On the client routers, all you have to do is create a neighbor configuration for the route reflector device. You do not have to add any additional parameters to the configuration.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

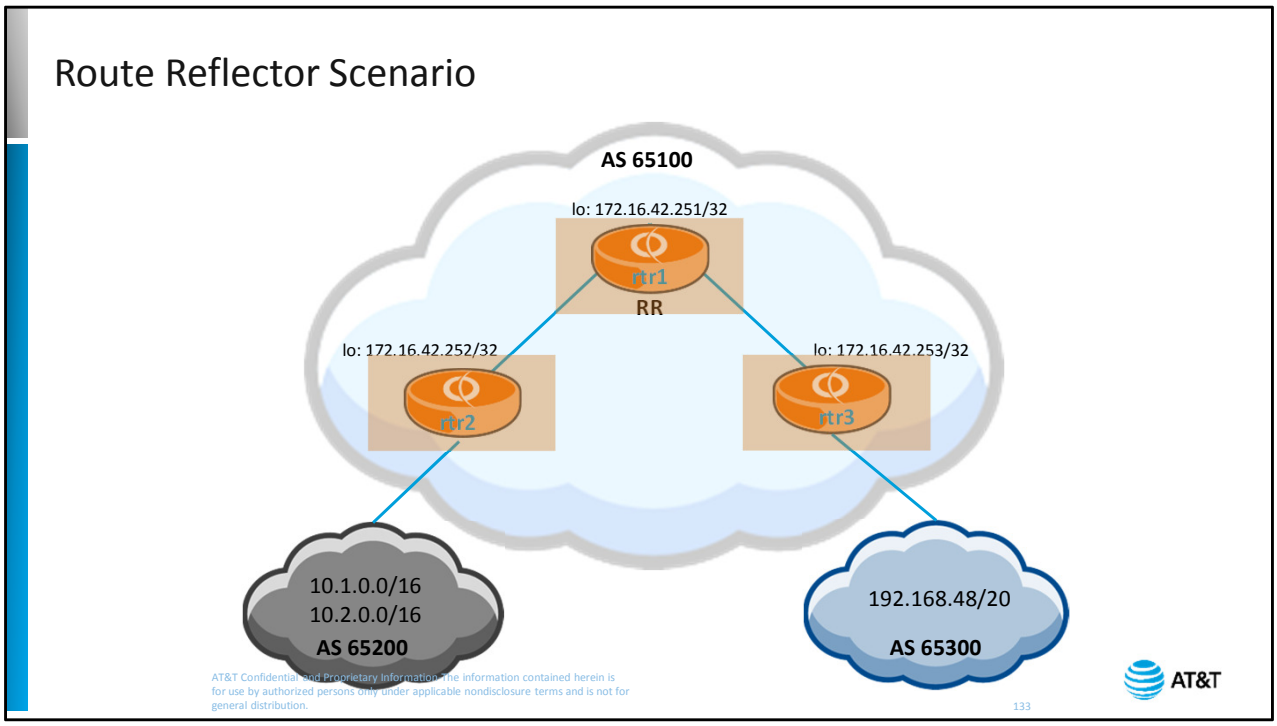

In our scenario, we will configure router 1 (rtr1) as the route reflector for the AS. In this particular configuration, router 1 does not have any eBGP connections, although that is not a requirement for route reflection.
Routers 2 (rtr2) and 3 (rtr3) are the route reflector clients.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Route Reflector Configuration – rtr1

```
vyatta@rtr1# show protocols bgp 65100
 address-family {
     ipv4-unicast {
         }
     }
 }
neighbor 172.16.42.252 {
     address-family {
         ipv4-unicast {
             route-reflector-client
         }
     }
     remote-as 65100
     update-source 172.16.42.251
 }
neighbor 172.16.42.253 {
     address-family {
         ipv4-unicast {
             route-reflector-client
         }
     }
     remote-as 65100
     update-source 172.16.42.251
<Truncated Output>
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

On router 1, we configured our two neighbors, adding the `route-reflector-client` parameter to each neighbor.

The rest of the neighbor configuration is typical for iBGP – the remote AS is the same as the local AS, and the update-source IP address is the address of loopback interface on router.

Note that the configuration for the iBGP neighbors is identical other than the neighbor IP address. With only two clients, this is not too bad, but if we had more route reflector clients, configuring the same parameters over and over gets tedious and creates more opportunities for typing errors.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Adding a Peer Group

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Peer Groups

Configure peer groups under AS number

```
set protocol bgp as-num peer-group groupname
```

- Configure peer group parameters per address family
  - Peer groups can be part of both IPv4 and IPv6 address families
  - All neighbors in the peer group inherit the group configuration

Add neighbor to peer group under neighbor address family

- Address family depends on the address family configured above
  - If peer group is a member of IPv4 and IPv6 address families, enter both neighbor statements

```
set protocol bgp as-num neighbor ipaddress address-family
ipv4-unicast peer-group groupname
```

```
set protocol bgp as-num neighbor ipaddress address-family
ipv6-unicast peer-group groupname
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

An alternative to configuring each client individually is to use a peer group. Peer groups simplify the configuration for neighbors with similar update policies. Once a peer group is created, it can be configured using the same commands for configuring individual BGP neighbors. Each member of the peer group inherits the peer group's update policies unless overridden by explicit configuration of the individual peer. This also makes update calculation more efficient.

Begin by creating a peer group under the desired AS number. Configure parameters such `nexthop-self` and `allowas-in` per address family.

Next, we configure the individual neighbors. We configure the neighbor statement with the peer group name under the desired address family. You must configure the neighbor with the same address family configured under for peer group. A peer group can be configured for both IPv4 and IPv6 address families. If that is the case, you must enter neighbor statements for both address families.

Remember, all peer group parameters will automatically apply to the neighbor.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Peer Group Configuration – rtr1

```
vyatta@rtr1# show protocols bgp 65100
 address-family {
             ipv4-unicast {
                 network 172.24.42.0/24 {
       }
       neighbor 172.16.42.251 {
           address-family {
               ipv4-unicast {
                   peer-group rrgroup
}
       neighbor 172.16.42.252 {
           address-family {
               ipv4-unicast {
                   peer-group rrgroup
       }
       parameters {
          router-id 172.16.42.251
<Truncated Output>
```

```
peer-group rrgroup {
    address-family {
        ipv4-unicast {
            route-reflector-client
        }
    }
    remote-as 65100
    update-source 172.16.42.251
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look at a peer group configuration.
We have a peer-group named *rrgroup*. The peer group configuration has the remote AS number, the `route-reflector-client` parameter, and the `update-source` parameter.
The individual neighbors are members of the peer group under the IPv4 address family. If we want to add another client router, all we have to do is add a new neighbor statement assigning the client to the peer group.
We used groups here to simplify route reflector configuration, but you can use peer groups anywhere you need to apply the same parameters to multiple peers. For example, if you want to apply route filters that block all private addresses to all your eBGP peers, you can create a peer group, assign the filters to the group, then make each eBGP neighbor a member of that group.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Client Configuration – rtr2 & rtr3

Router 2 (rtr2)

```
vyatta@rtr2# show protocols bgp 65100
 address-family {
     ipv4-unicast {
         }
neighbor 10.1.0.2 {
     address-family {
         ipv4-unicast {
         }
     }
     remote-as 65200
 }
neighbor 172.16.42.251 {
     address-family {
         ipv4-unicast {
                     nexthop-self
         }
     }
     remote-as 65100
     update-source 172.16.42.252
 }
 parameters {
     router-id 172.16.42.252
```

Router 3 (rtr3)

```
vyatta@rtr3# show protocols bgp 65100
 address-family {
     ipv4-unicast {
         }
neighbor 172.16.42.251 {
     address-family {
         ipv4-unicast {
                     nexthop-self
         }
     }
     remote-as 65100
     update-source 172.16.42.253
 }
neighbor 192.168.51.2 {
     address-family {
         ipv4-unicast {
         }
     }
     remote-as 65300
 }
parameters {
     router-id 172.16.42.253
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution. 138

On the client routers, we configured router 1 as an iBGP neighbor. We included the `nexthop-self` parameter because these two routers also have eBGP peers.
We see the eBGP configuration here. The client router configuration is the same whether you use peer groups on the route reflector or not.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying Operations

```
vyatta@rtr1:~$ show ip bgp summary
BGP router identifier 172.16.42.251, local AS number 65100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.42.252   4 65100       6       7        0    0    0 00:03:42        2
172.16.42.253   4 65100       7       8        0    0    0 00:04:20        1
Total number of neighbors 2

vyatta@rtr1:~$ show ip bgp
BGP table version is 1, local router ID is 172.16.42.251
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i10.1.0.0/16      172.16.42.252            0    100      0 65200 i
*>i10.2.0.0/16      172.16.42.252            0    100      0 65200 i
*>i192.168.48.0/20  172.16.42.253            0    100      0 65300 i
Total number of prefixes 3
vyatta@rtr1:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying route reflector operations is the same as verifying other BGP operations. First, verify that we are fully connected to our BGP peers.
Looking at the *state* column, we can see that we have received prefixes from each of our iBGP neighbors, so we know we have established a full BGP connection.
We can also look at the BGP table to see which routes we have received.
When we look at the next-hop addresses, they are the addresses of our iBGP peers, so again we know we have established a full connection.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying Operations (cont.)

```
vyatta@rtr2:~$ show ip bgp
BGP table version is 2, local router ID is 172.16.42.252
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.0.0/16      10.1.0.2                 0             0 65200 i
*> 10.2.0.0/16      10.1.0.2                 0             0 65200 i
*>i192.168.48.0/20  172.16.42.253            0    100      0 65300 i
Total number of prefixes 3
```

```
vyatta@rtr3:~$ show ip bgp
BGP table version is 2, local router ID is 172.16.42.253
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i10.1.0.0/16      172.16.42.252            0    100      0 65200 i
*>i10.2.0.0/16      172.16.42.252            0    100      0 65200 i
*> 192.168.48.0/20  192.168.51.2             0             0 65300 i
Total number of prefixes 3
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When we look at the route reflector clients, we can see that the routes are being propagated by the route reflector.
Router 2 has the route coming in to our AS via Router 3, and Router 3 has the routes coming in to our AS via Router 2.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Multiple Route Reflectors

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Multiple Reflectors for Redundancy

Create multiple route reflectors (rtrA and rtrB)

Route reflectors must be fully meshed

- `route-reflector-client` statement not needed

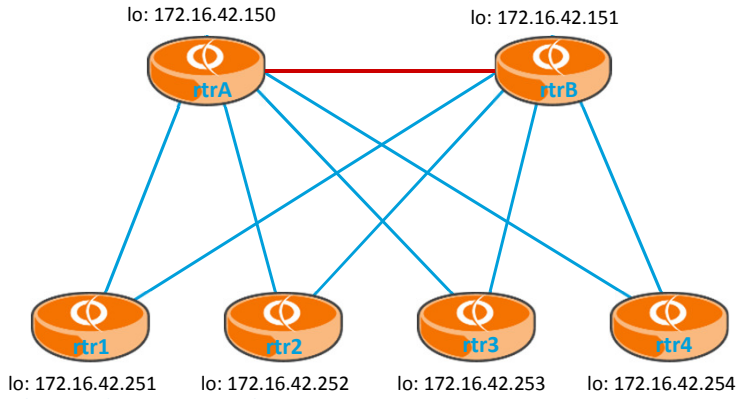

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

One concern when deploying route reflectors is redundancy. Should the reflector device fail, you will lose your entire iBGP network.
One way to address this is to create multiple route reflectors. This still avoids the full-mesh requirement; all clients peer only to the route reflectors.
However, the route reflectors themselves must be fully meshed.
You do this by not including the `route-reflector-client` parameter in the neighbor statement for the other route reflectors.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Multiple Reflectors Configuration – rtrA

```
vyatta@rtrA# show protocols bgp 65100
 address-family {
     ipv4-unicast {
         network 172.24.42.0/24 {
}
 neighbor 172.16.42.151 {
    address-family {
         ipv4-unicast {
         }
    remote-as 65100
    update-source 172.16.42.150
}
 neighbor 172.16.42.251 {
    address-family {
         ipv4-unicast {
           peer-group rrgroup
}
 neighbor 172.16.42.252 {
    address-family {
         ipv4-unicast {
           peer-group rrgroup
}
<Truncated Output>
```

```
parameters {
     router-id 172.16.42.150
}
peer-group rrgroup {
     address-family {
         ipv4-unicast {
             route-reflector-client
         }
     }
     remote-as 65100
     update-source 172.16.42.150
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Looking at the configuration on Router A (rtrA), we have created the peer group called *rrgroup* for all the route reflector clients.
Note the peer group is configured with the update source address of the rtrA router ID.
Router A also has a neighbor statement for the other route reflector router B (rtrB). RtrB is not a part of the rrgroup peer group, it is configured as a stand-alone peer. Also note that it does not have the `route-reflector-client` parameter.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Client Configuration – rtr2

```
vyatta@rtr2# show protocols bgp 65100
address-family {
     ipv4-unicast {
         }
neighbor 172.16.42.150 {
     address-family {
         ipv4-unicast {
             nexthop-self
         }
     }
     remote-as 65100
     update-source 172.16.42.252
}
neighbor 172.16.42.151 {
     address-family {
         ipv4-unicast {
              nexthop-self
         }
     }
     remote-as 65100
     update-source 172.16.42.252
}
parameters {
     router-id 172.16.42.252
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Looking at the configuration on one of the route reflector clients, router 2 (rtr2), we see that there are two neighbor statements, one to each of the route reflectors.
Each route reflector client would be configured this way.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Confederations

Next we will take a look at BGP confederations.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Solution 2: Confederations

Break AS into multiple private AS's

- Full mesh within private AS
- eBGP between private AS

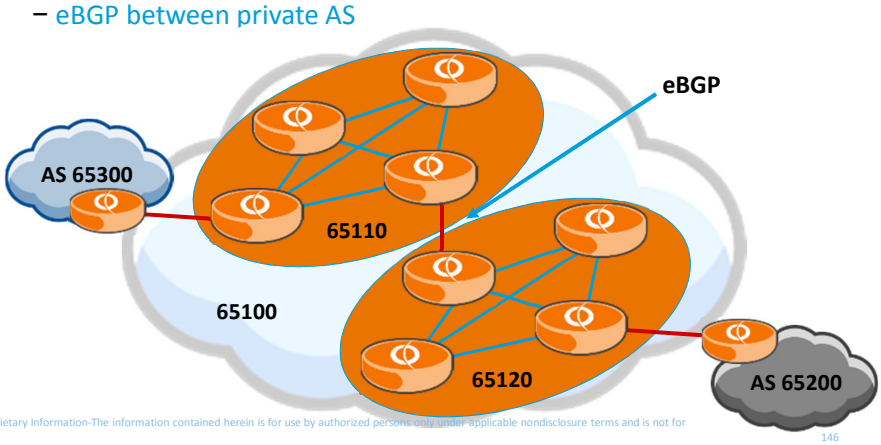

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

A second option for very large networks is to use BGP confederations.
Confederations are used to break up a large iBGP network into smaller private autonomous systems. From a BGP operational standpoint, all the routers within each private AS run iBGP in a full mesh configuration – or use route reflectors.
Between each private AS the routers run eBGP, just as if they were connected to another public autonomous system.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuration AS Numbering

Create BGP instance using private AS number

- Private AS range: 64512-65535
- In our example, 65110 and 65120

Create confederation using public AS number

- In our example, 65100

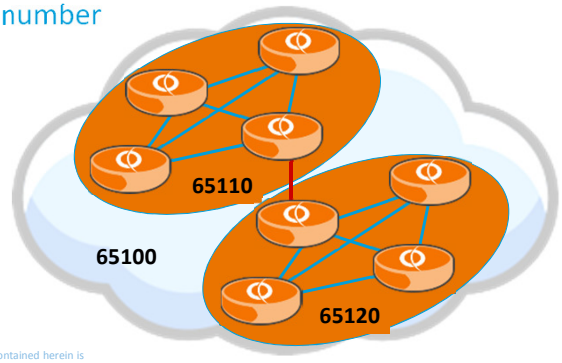

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

When configuring confederations, you use AS numbers in two different places.
When you create the BGP instance for your device, you configure it using the private AS number you selected for your confederation.
The range of private AS numbers is 64512 through 65535.
In our example, the private AS numbers – or sub-AS numbers – are 65110 and 65120.
The confederation itself uses the public AS number, the number assigned to you by the addressing authority. In our example, this number is 65100. Remember, all vRouter courses use private numbers in order to avoid using a company's real address range or AS number.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Configuring Confederations

Configure BGP peering using private AS number

```
set protocol bgp private_as-num neighbor ipaddress
remote-as private_as-num
```

Create confederation under BGP parameters

- Assign public AS number

```
edit protocol bgp private_as-num
set parameters confederation identifier public_as-num
```

- Configure peer private AS numbers

```
set confederation peers private_as-num
```

  - This is the private IP address of the neighboring confederation

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

To configure a confederation, you begin by creating the BGP peering between the routers in the private AS. Note that the neighbor statement uses the private AS number as the remote AS since these are iBGP peers.
Next add the confederation configuration to the parameters level of the BGP instance.
Use the `set parameters confederation identifier` command with the public IP address,
then list the private AS numbers, one peer per statement. This is the AS number of neighboring confederations. You can list as many private numbers as you have configured within your confederation.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

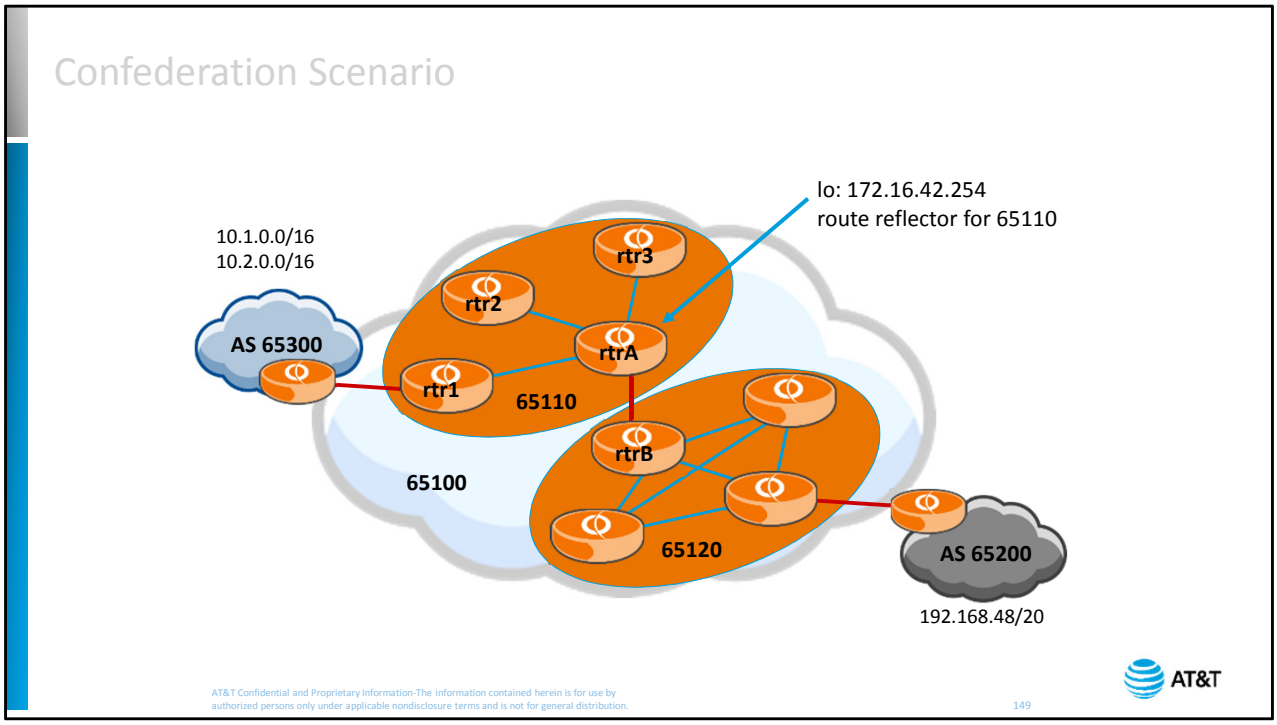

In this scenario, we have broken our iBGP network into two smaller private autonomous systems. We will focus on the configuration for Router A (rtrA), but of course for this entire scenario to be complete, we have to configure the appropriate parameters on all the iBGP routers. Note that rtrA, in addition to having an eBGP connection to Router B (rtrB), is the route reflector for all other routers in private AS 65110.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Confederation Configuration

```
vyatta@rtrA# show protocols bgp 65110
<Truncated Output>
 neighbor 172.16.42.128 {
    address-family {
         ipv4-unicast {
         }
    description rtrB
    remote-as 65120
}
 neighbor 172.16.42.251 {
    address-family {
         ipv4-unicast {
           peer-group rrgroup
    }
    description rtr1
}
 neighbor 172.16.42.252 {
    address-family {
         ipv4-unicast {
           peer-group rrgroup
    }
    description rtr2
<Truncated Output>
```

```
parameters {
     confederation {
        identifier 65100
        peers 65120
     }
     router-id 172.16.42.254
}
peer-group rrgroup {
     address-family {
         ipv4-unicast {
            nexthop-self
            route-reflector-client
         }
     }
     remote-as 65110
     update-source 172.16.42.254
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Let's look at the configuration for our scenario.
Our eBGP peer is in AS 65120. We have included a description for each of our BGP peers to make reading the configuration file a little easier, so we know this is the connection to Router B (rtrB) in our diagram.
The peers within our private AS are all part of the *rrgroup* peer group. Again, we have used a description under each neighbor so we know this group defines the connections to our route reflector clients. Router 1 (rtr1) and router 2 (rtr2) are shown here, router 3 (rtr3) is not shown, but it is configured in the same way.
In the parameters section, we see the confederation configuration. The identifier is the public AS number 65100, and the peer AS is number 65120. This is the peer confederation AS number.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Verifying Confederation Operations

```
vyatta@rtrA:~$ show ip bgp
BGP table version is 1, local router ID is 172.16.42.254
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i10.1.0.0/16      172.16.42.251            0    100      0 65200 i
*>i10.2.0.0/16      172.16.42.251            0    100      0 65200 i
*> 172.16.42.0/24   0.0.0.0                  0               32768 i
*>i172.16.51.0/24   172.16.42.128            0    100      0 (65120) i
*>i192.168.48.0/20  172.16.42.128            0    100      0 (65120) 65300 i

Total number of prefixes 5
vyatta@rtrA:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

You can verify overall confederation operations just as you would any other BGP network. To verify that the confederation is using AS numbers correctly, look at the output of the `show ip bgp` command.

Routes arriving from other private autonomous systems include the private AS in the path. If the AS number is part of the confederation, it will be called out in parentheses in the path. Any information within parentheses will be replaced with the real AS number before being advertised outside the real AS.

Note that these routes are flagged as internal routes with the “i” in front of the route, even though they are learned over an eBGP connection. The device recognizes the private AS number of the peer as part of the confederation and flags the route appropriately.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Displaying AS Numbers

```
vyatta@rtrA:~$ show ip bgp summary
BGP router identifier 172.16.42.254, local AS number 65110
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

IPv4 Unicast address family:
Neighbor        V     AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.42.128   4 65120       7       8        0    0    0 00:05:19        2
172.16.42.251   4 65110       3       4        0    0    0 00:02:31        1
172.16.42.252   4 65110       6       7        0    0    0 00:03:42        2
172.16.42.253   4 65110       7       8        0    0    0 00:04:20        1
Total number of neighbors 4

vyatta@rtrA:~$
```

```
parameters {
     confederation {
        identifier 65100
        peers 65120
     }
     router-id 172.16.42.254
vyatta@rtrA#
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Note that nowhere in the `show` commands, either in the `show ip bgp` output on the previous screen, or the `show ip bgp summary` output shown here, will you find the public AS number used by the confederation (in this case 65100).
You can see the private AS number used by the local device, and the AS numbers of the peers.
The only place that you can find the public AS number is the configuration itself.

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

Summary

You should now be able to

- List the design requirements for implementing iBGP
- Configure route reflectors
- Configure peer groups
- Configure confederations
- Verify correct iBGP functionality

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

This concludes the AT&T Vyatta 5600 vRouter iBGP Design module.
You should now be able to:

- List the design requirements for implementing iBGP
- Configure route reflectors
- Configure confederations
- Verify correct iBGP functionality

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

End of Course: AT&T Vyatta 5600 vRouter BGP Essentials

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.

AT&T

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution.