NFV 363-WBT vRouter Multicast Routing

NFV 363-WBT vRouter Multicast Routing

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.

Welcome to the AT&T vRouter Dynamic Multipoint VPN course.

NFV 363-WBT vRouter Multicast Routing

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.

Before we begin the course, please take a moment to read our legal disclaimer.

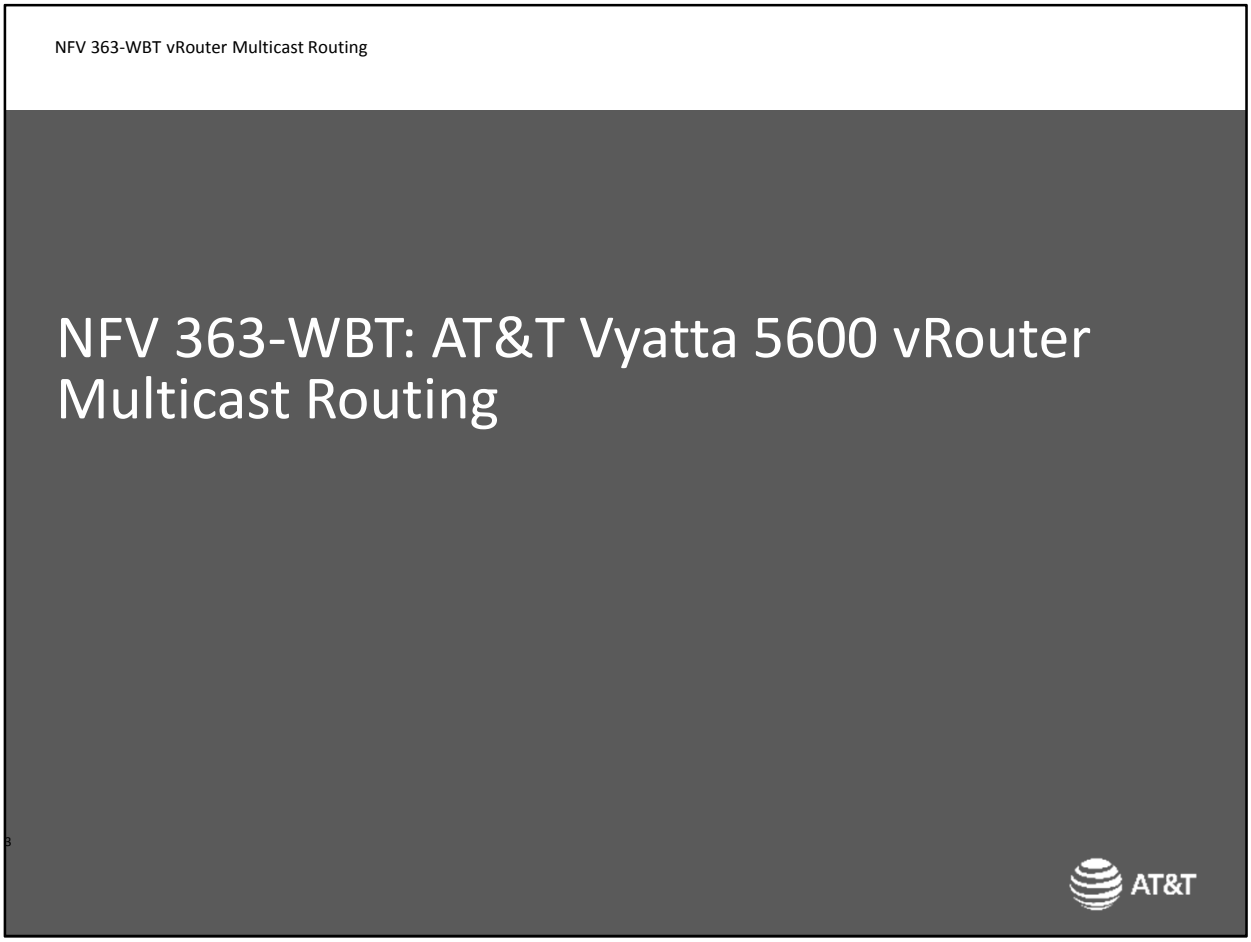

Welcome to the AT&T Vyatta 5600 vRouter Multicast Routing course.

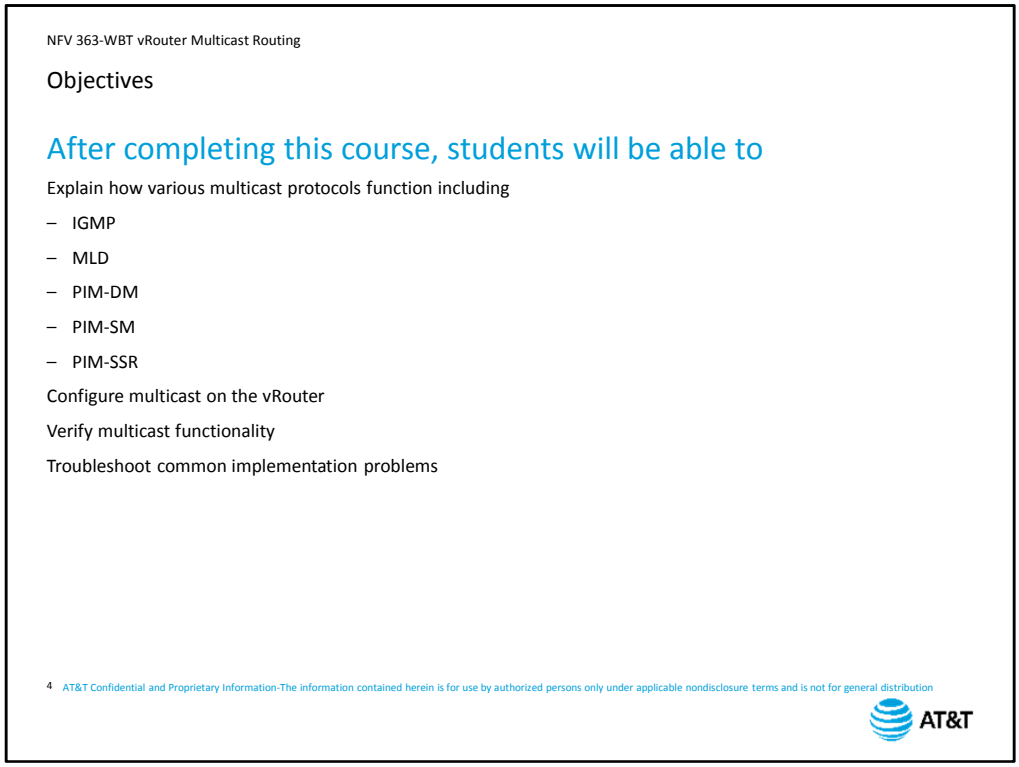

After completing this course, you will be able to:

- Explain how various multicast protocols function, including IGMP, MLD, PIM-dense mode, PIM sparse mode, and PIM source-specific multicast
- Configure multicast support on the vRouter
- Verify multicast functionality
- And troubleshoot common implementation problems

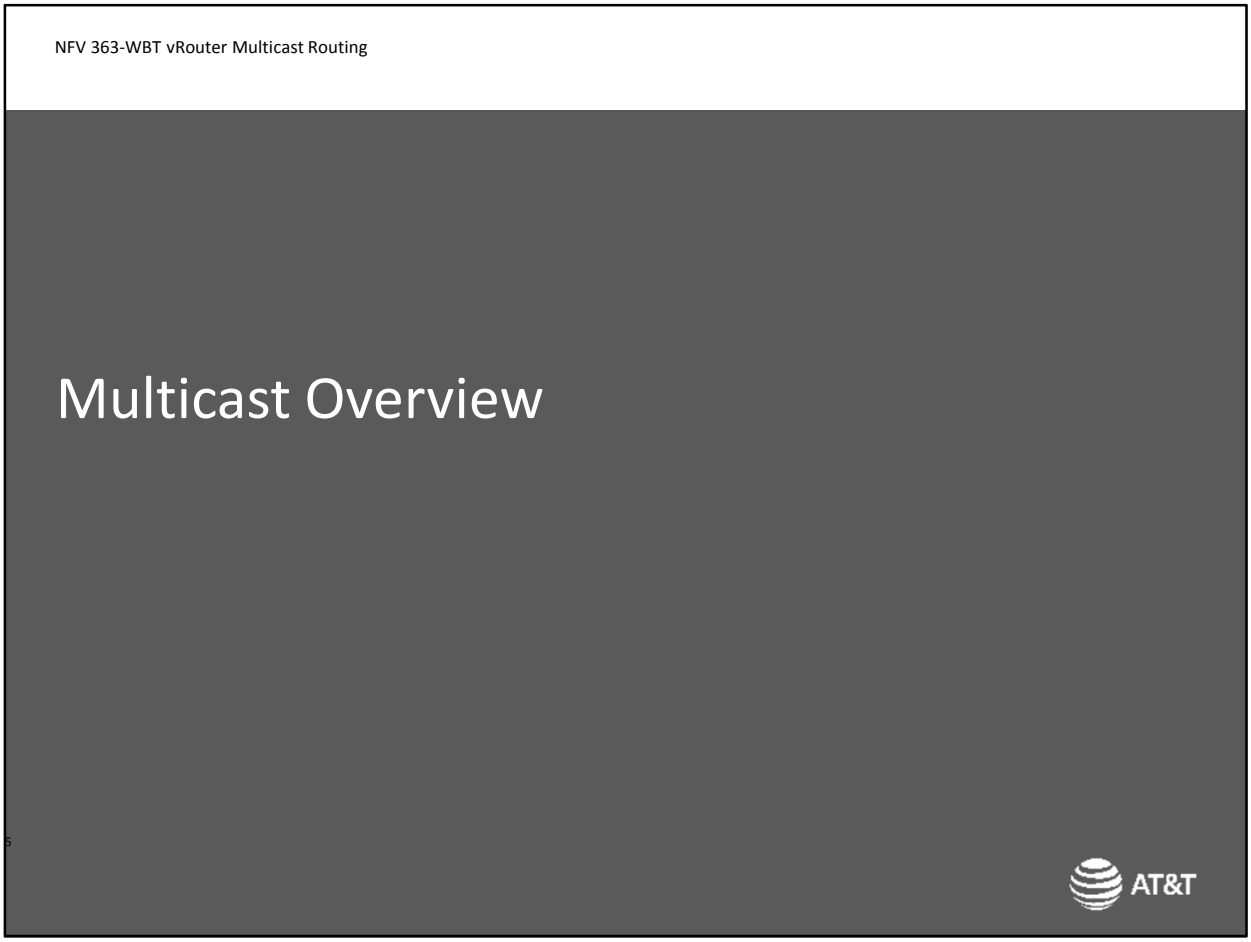

We'll begin with a quick overview of multicast at Layers 2 and 3.

NFV 363-WBT vRouter Multicast Routing

What is Multicast?

Multicast is traffic sent from a single source to a specific group of destinations

Nodes respond via unicast

Multicast uses UDP

L2 switches flood multicast by default

L3 routers do not forward by default

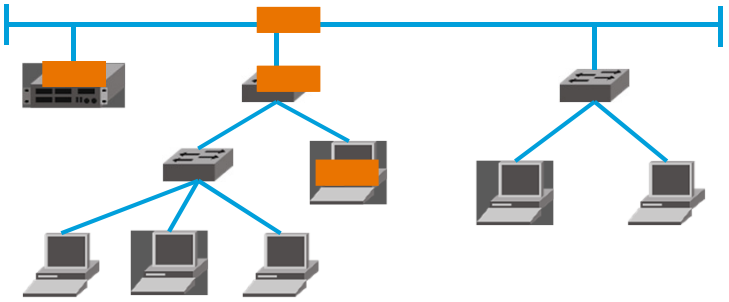

AT&T

Multicast is an efficient way of sending one-to-many communications across IP networks. Multicast applications enable end stations to listen for specific multicast addresses in addition to the configured unicast IP addressing.
Multicast is one-way traffic; individual nodes that need to send traffic in response to a multicast source will send it via unicast.
Multicast traffic is also connectionless, using UDP rather than TCP. This enables individual nodes to join or leave a multicast on demand without requiring a new connection to be established for each node.
At Layer 2, multicast traffic is handled like broadcast traffic by default; any multicast traffic it receives gets flooded out all other ports. Multicast-intelligent switches can modify this behavior.
At Layer 3, multicast traffic is dropped by default. In order for multicast traffic to cross routers, you have to enable multicast routing.

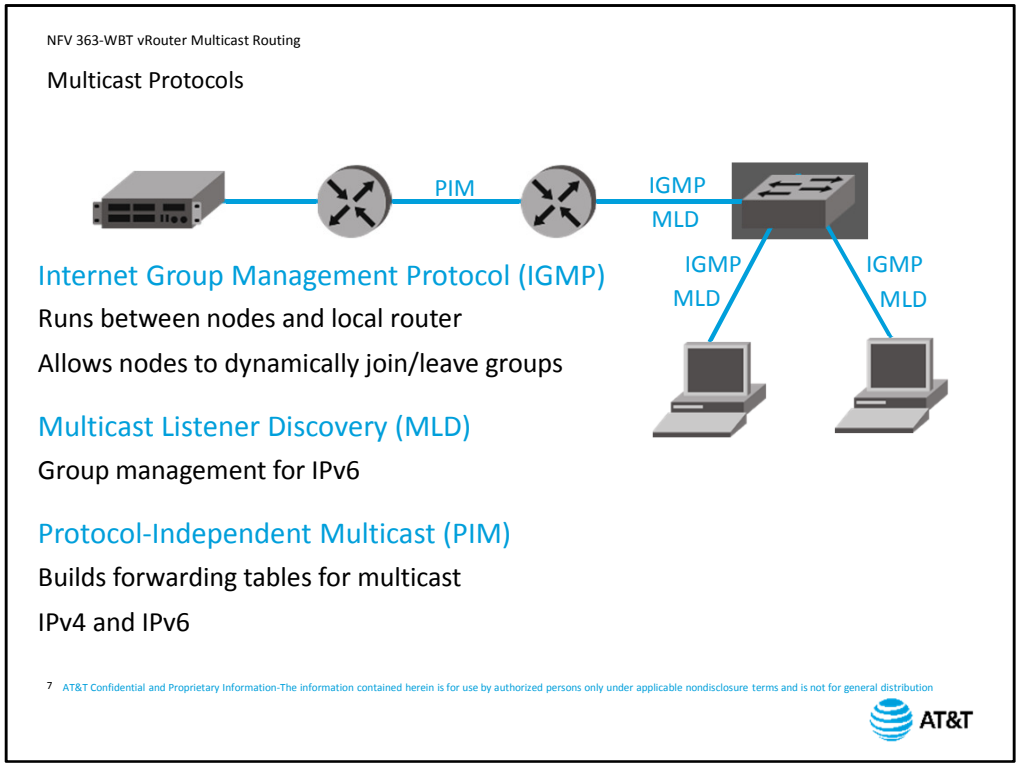

When you implement multicast, you will be working with several different protocols, and you need to understand how they work together.
The first protocol is Internet Group Management Protocol, or IGMP.
IGMP runs between nodes that are receiving multicast traffic, and the nearest multicast-capable router.
IGMP allows hosts to dynamically join and leave a multicast group. We will take a detailed look at IGMP later in this course.
As we said earlier, Layer 2 switches flood multicast traffic by default. However, if your Layer 2 switch has IGMP support, it can prune multicast to only those ports where hosts are actively participating in multicast. The switch will not participate in the IGMP exchange itself; it will learn from the IGMP messages exchanged by the node and the router.
If you are using IPv6, you will use Multicast Listener Discovery instead of IGMP in order to manage multicast group membership. MLD is part of ICMP in IPv6 and is not a separate protocol.
Between multicast routers, you will configure Protocol-Independent Multicast, or PIM. PIM builds forwarding trees for IP multicast traffic so that it is only sent across links where there are devices listening for the multicast.
PIM is used for both IPv4 and IPv6, but the processes run separately and build separate forwarding tables.

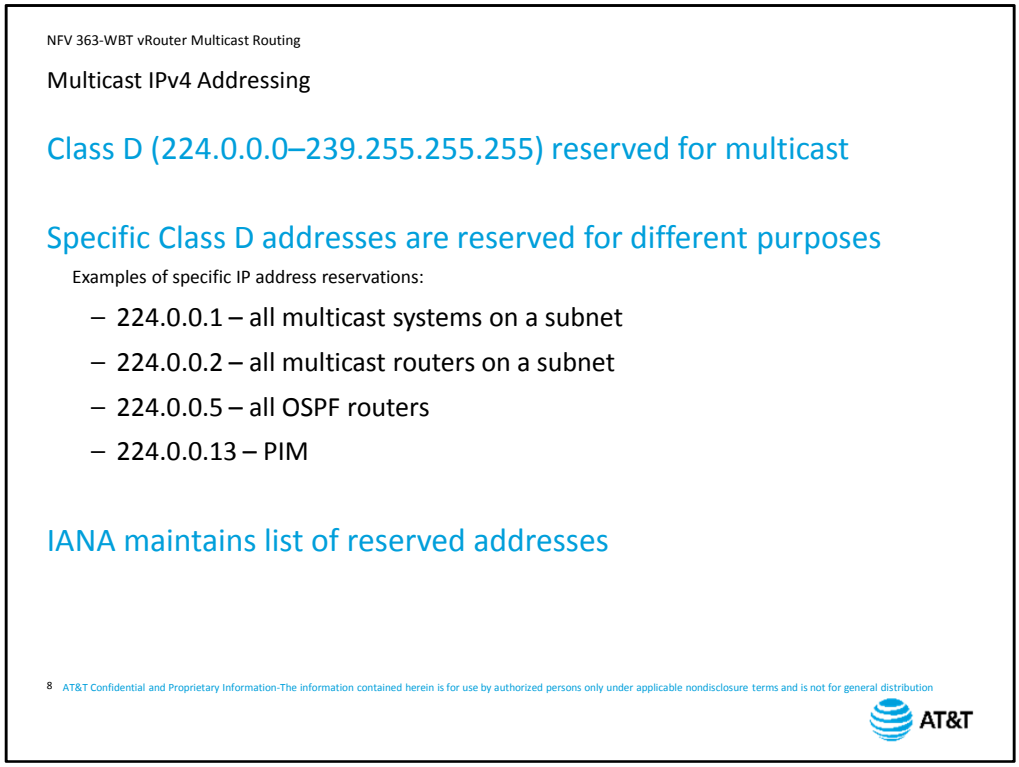

When you set up a multicast service, you configure the source to use a particular multicast Layer 3 address. The address must be from the range reserved for multicast traffic. If you are using IPv4, these addresses will come from the range reserved by Class D. The first four bits of an IPv4 multicast address will always be 1110 or 224.
You need to make sure the multicast address you select is not reserved for other uses. Your screen shows a few of the reserved address that you may see on your network.
The Internet Assigned Numbers Authority maintains the list of reserved addresses.

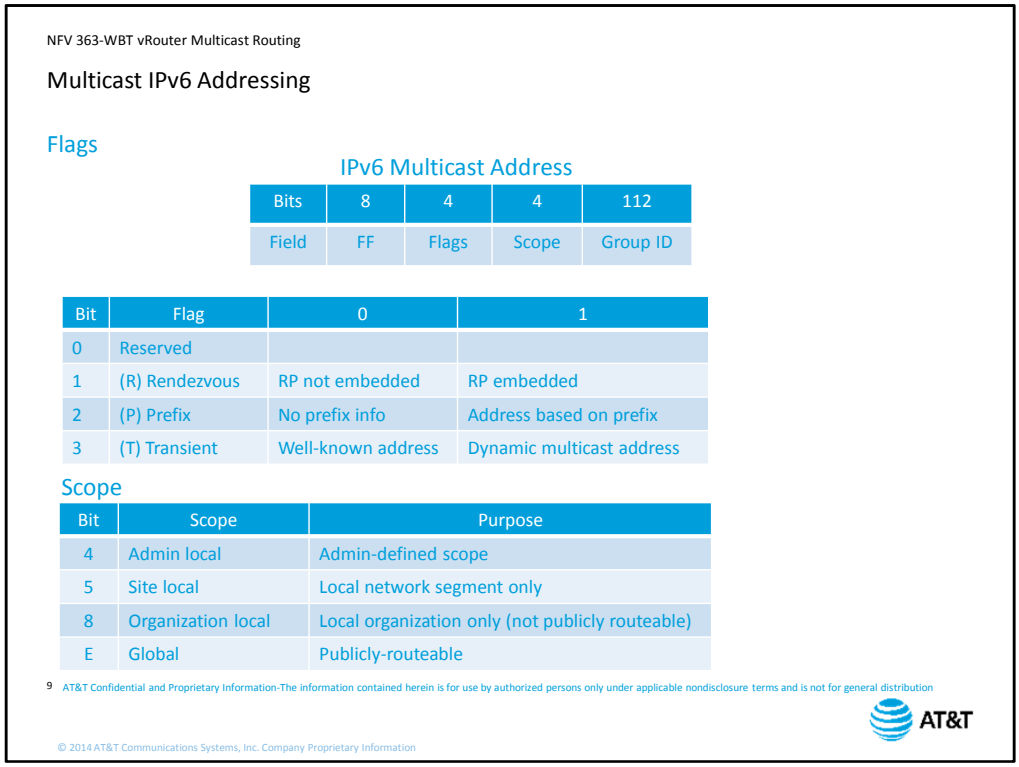

IPv6 also has reserved addresses for multicast. An IPv6 multicast address will begin with the first 8 bits set to FF. The next 8 bits are used to embed information regarding the address, and to limit the scope of the packet forwarding, and the last one hundred and twelve bits identify the multicast group.

The flags include information on the nearest multicast rendezvous point. We will discuss rendezvous points when we cover routing protocols later in this course. The flags also include information about the rest of the address.

The scope field specifies the limit in which the address is valid. Publicly-routeable addresses will have the field set to *E*. Other settings specify narrower scopes and can be used to limit the forwardability of traffic to a network or set of networks.

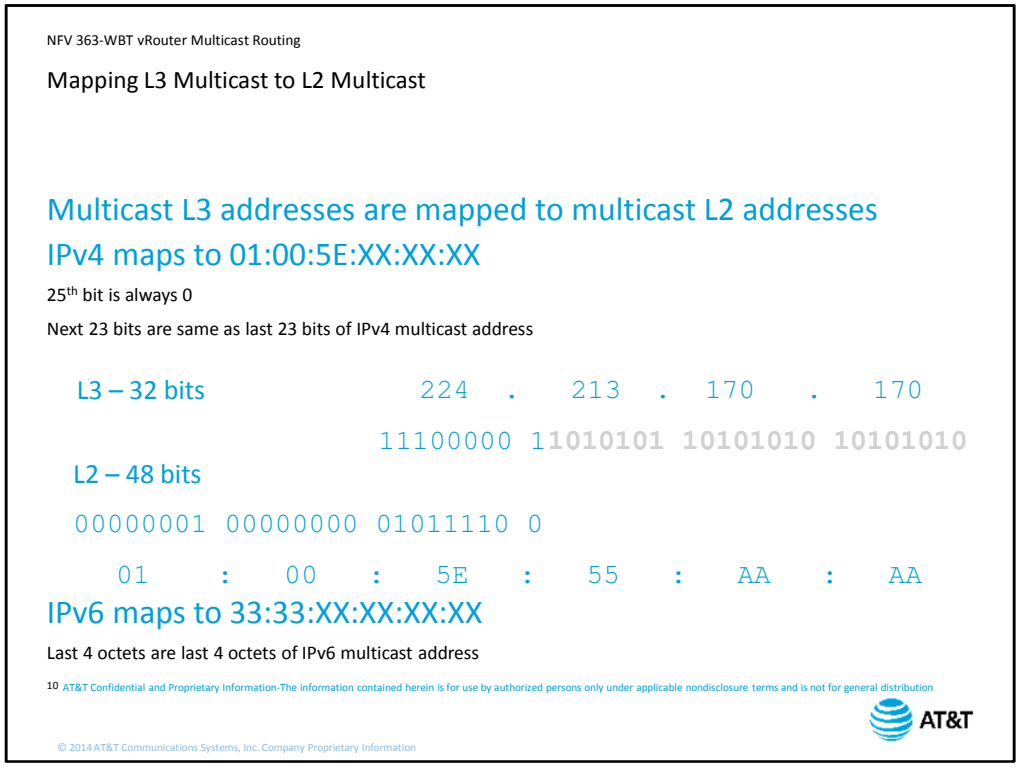

All multicast addresses map to Layer 2 multicast addresses.
If you are using IPv4, a multicast MAC address will begin with *zero one zero zero five e* (01 00 5e).
The 25th bit of the MAC address is zero, and the remaining 23 bits are the same as the last 23 bits of the IPv4 multicast address. Let's look at this more closely.
We are using multicast address 224.213.170.170.
We convert this to binary so we can see how the bits move into the MAC address.
The MAC address always begins with the same 25 bits.
The last 23 bits of the IP address are copied to the MAC address field.
Of course, MAC addresses are usually expressed in hexadecimal, so this is the address we would actually see on the network.
IPv6 is simpler. The first two octets are always 33:33, and the last 4 octets are the same as the last 4 octets of the Layer 3 address.

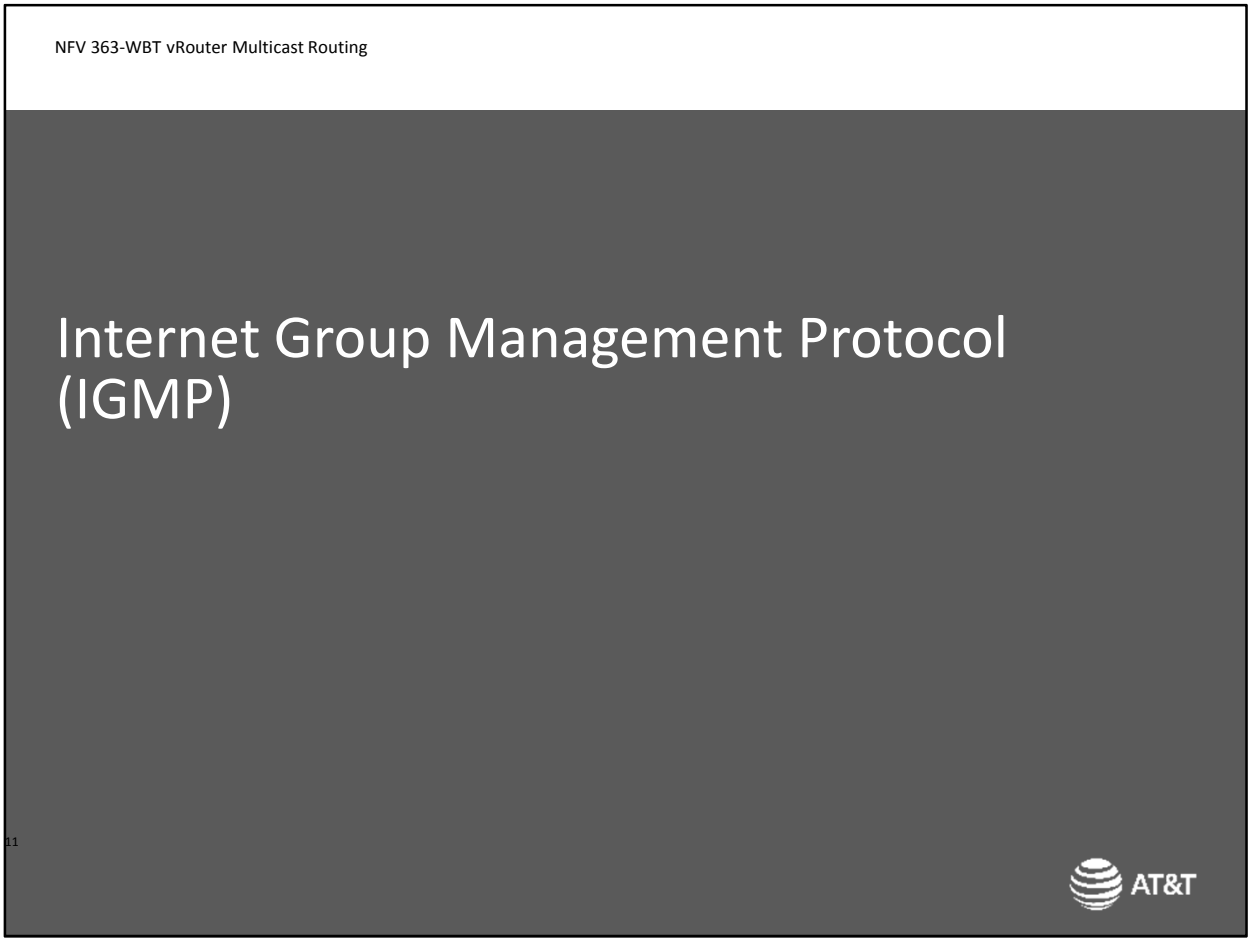

Next, we'll discuss IGMP and how multicast sources and clients use it to join and leave multicast groups.

NFV 363-WBT vRouter Multicast Routing

IGMP Overview

IGMP is a communications protocol used by hosts and adjacent routers on IP networks to establish multicast groups

IGMP specifies how groups are created, NOT how traffic is forwarded

Multi-segment networks need additional protocols (e.g. PIM)

There are currently three versions of IGMP

vRouter defaults to v3; v1 and v2 are also supported

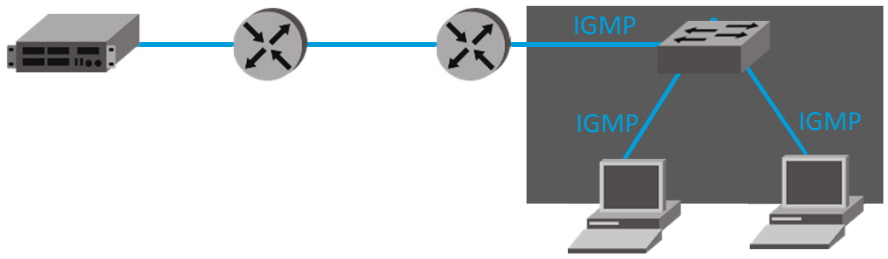

AT&T

As the name implies, Internet Group Management Protocol (IGMP) is used to manage multicast groups.
IGMP takes care of the “last hop” between listening multicast group members and their nearest multicast-capable router. It does not specify how to build paths from one multicast router to another.
Router-to-router path information is built using a multicast routing protocol such as PIM.
There are currently three versions of IGMP specified by RFCs.
The vRouter supports all versions, but defaults to version 3.

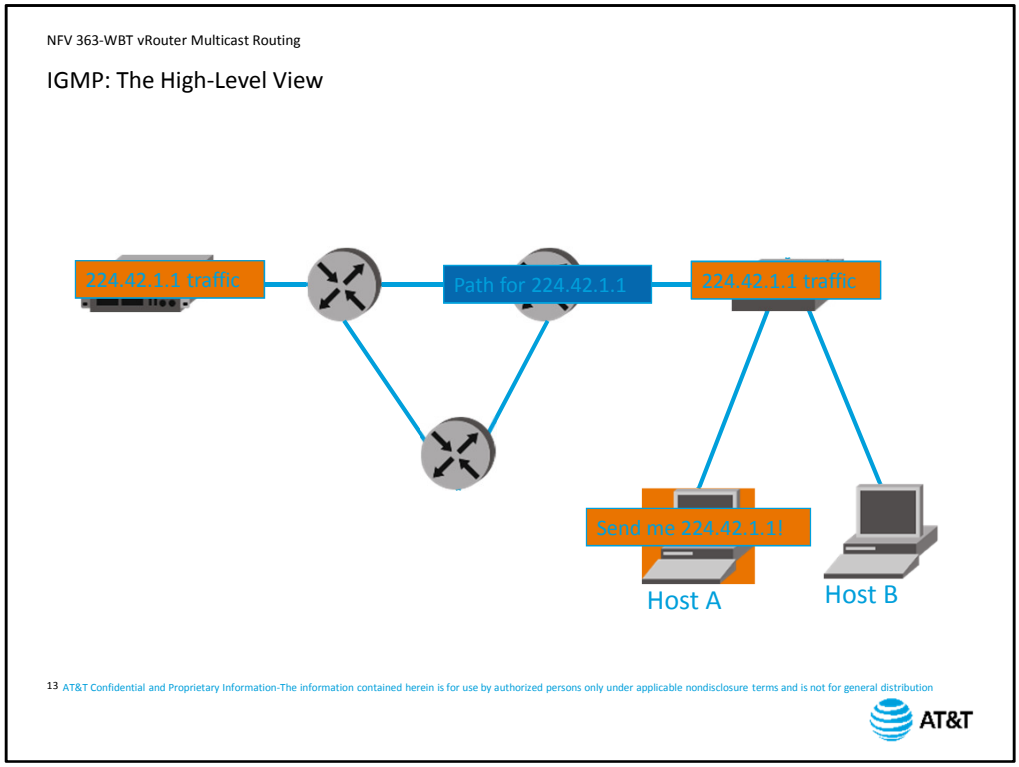

Before we get into version specifics, let's look briefly at how IGMP and a multicast routing protocol work together to build a path for multicast traffic.
Host A is running an application that needs to receive multicast traffic coming to address 224.42.1.1. It will send an IGMP message to its nearest multicast-capable router
The multicast router will communicate with other multicast routers to open a path for 224.42.1.1.
When the multicast server sends traffic, the routers will forward it along the path they have created, as well as on segments where there are group members.
By default, the switch will flood the multicast traffic out all ports, even to hosts that have not joined the group.
If your switch is capable of IGMP snooping, it will only forward to the IGMP group member.

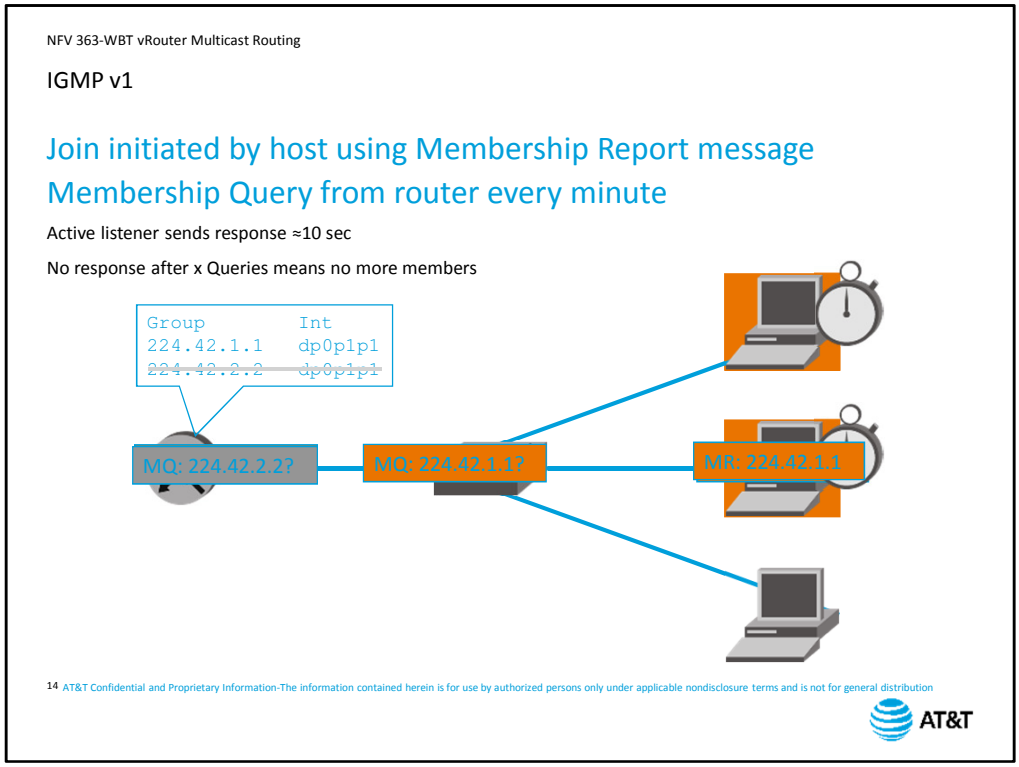

IGMP version 1 introduced the Membership Report and Membership Query message exchange.
When a host wants to join a group, it sends a Membership Report message to the All-multicast-hosts address, 224.0.0.1. The report contains the address of the multicast group the host wants to join.
The router keeps a table of multicast addresses and associated ports where hosts are active group members.
To keep this table current, the router sends out a multicast query for each group every minute.
Hosts that are active in the group set a countdown timer for a random interval between 0 and 10 seconds.
The first host whose timer expires sends a Membership Report. All others then cancel their reply.
If no host responds after a configurable number of queries, the router stops forwarding multicast traffic for that group number on the link.

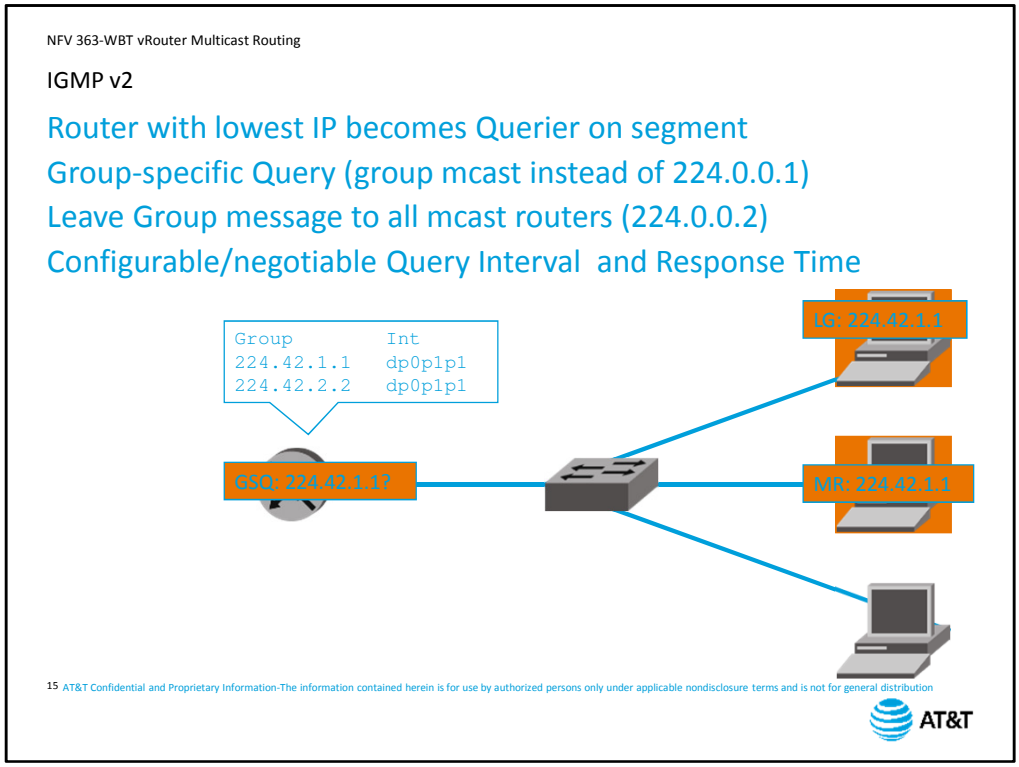

IGMP version 2 has four major changes to version 1.
First, Version 2 adds a Querier election process. All routers start out sending queries, and then listen to the queries to determine which device has the lowest IP address on the segment. That device becomes the Querier for the segment.
Second, the multicast query is now sent to the group multicast address instead of the “all multicast hosts” address of 224.0.0.1.
Third, version 2 adds a Leave Group message.
Now when a host is done listening to multicast, it sends a Leave Group message to the all multicast routers address.
The router immediately sends a Group-Specific Query to the group to see if there are any other hosts on the segment that are still active for the group.
Finally, the one minute General Query interval, and the up-to-10 second interval for response was replaced with a configurable Query Interval and Query Response Time.

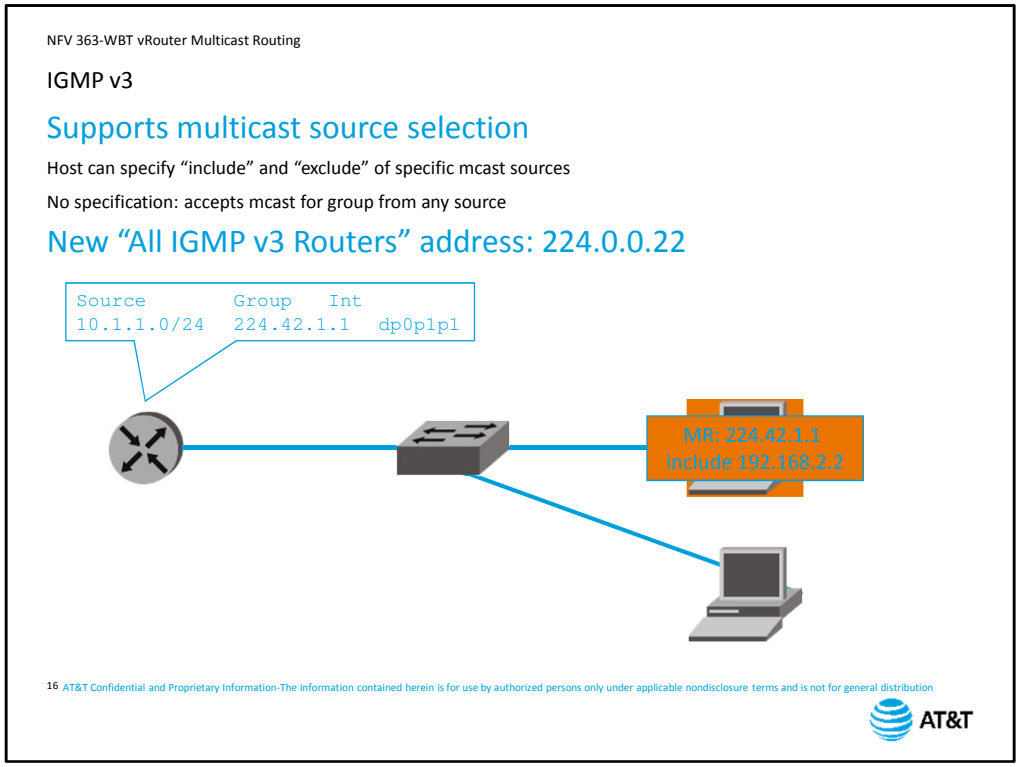

IGMP version 3 adds the ability for group members to specify sources of multicast traffic. Now when a host sends a membership report to join a group, it can specify a list of multicast sources that it will accept traffic from, or a list of hosts it will NOT accept traffic from.
This requires the multicast router to track the source of multicast traffic for each group. Groups can have multiple sources of traffic.
If the membership report does not specify an include or exclude list, it will accept multicast traffic from any source for that group.
The RFC also includes a new address for IGMP version 3 routers – 224.0.0.22. All version 3-capable routers listen on this address for version 3 membership reports.

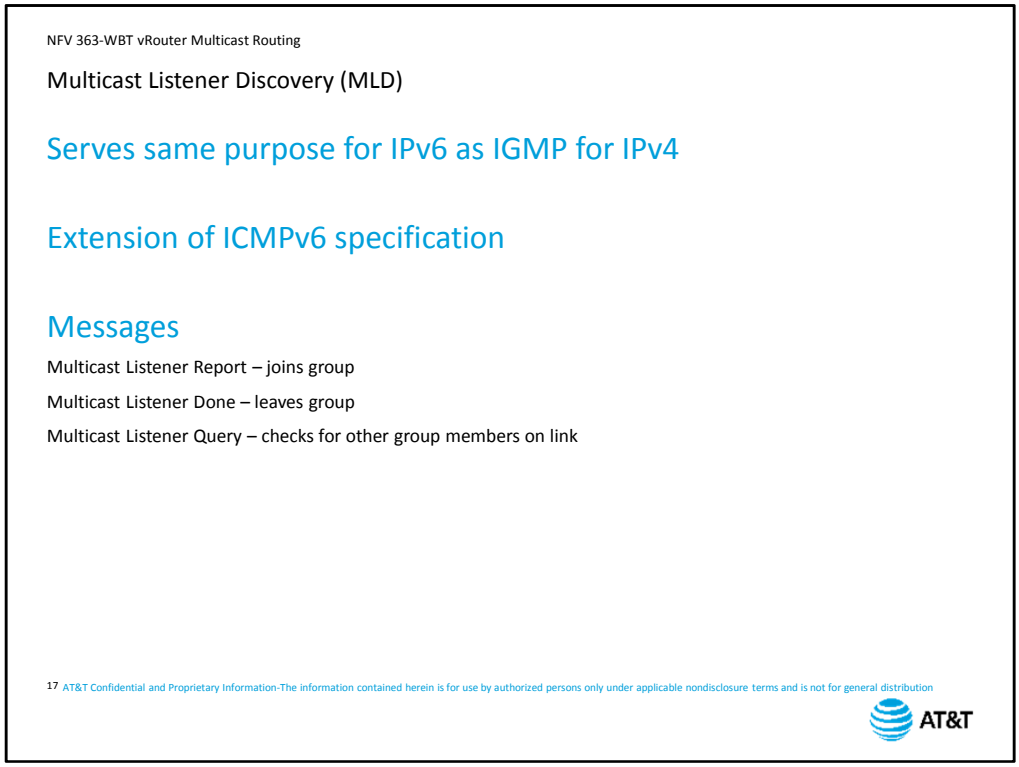

Multicast Listener Discovery (MLD) is essentially IGMP for IP version 6.
However, rather than a separate protocol, MLD is an extension of ICMP for IPv6.
MLD consists of three messages. A host sends a Multicast Listener Report when it wants to join a multicast group. This message is sent to all IPv6 multicast routers.
When the host leaves the group, it sends a Multicast Listener Done message to all multicast routers.
The router then sends a Multicast Listener Query on the link to see if there are any other members on the link. This sequence of messages should sound familiar, because it is the same basic exchange as IGMP.

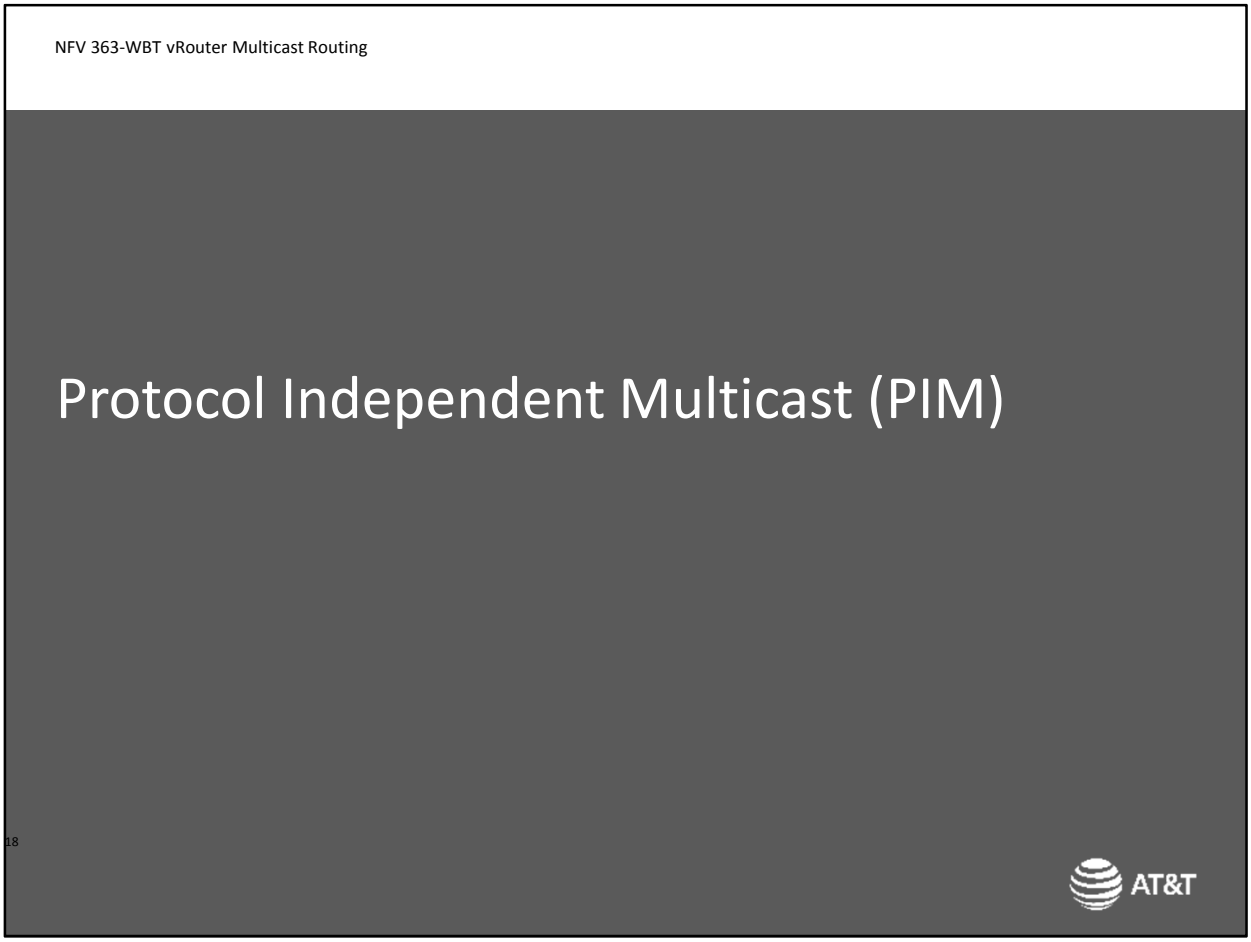

Now we'll look at Protocol-independent Multicast or PIM.

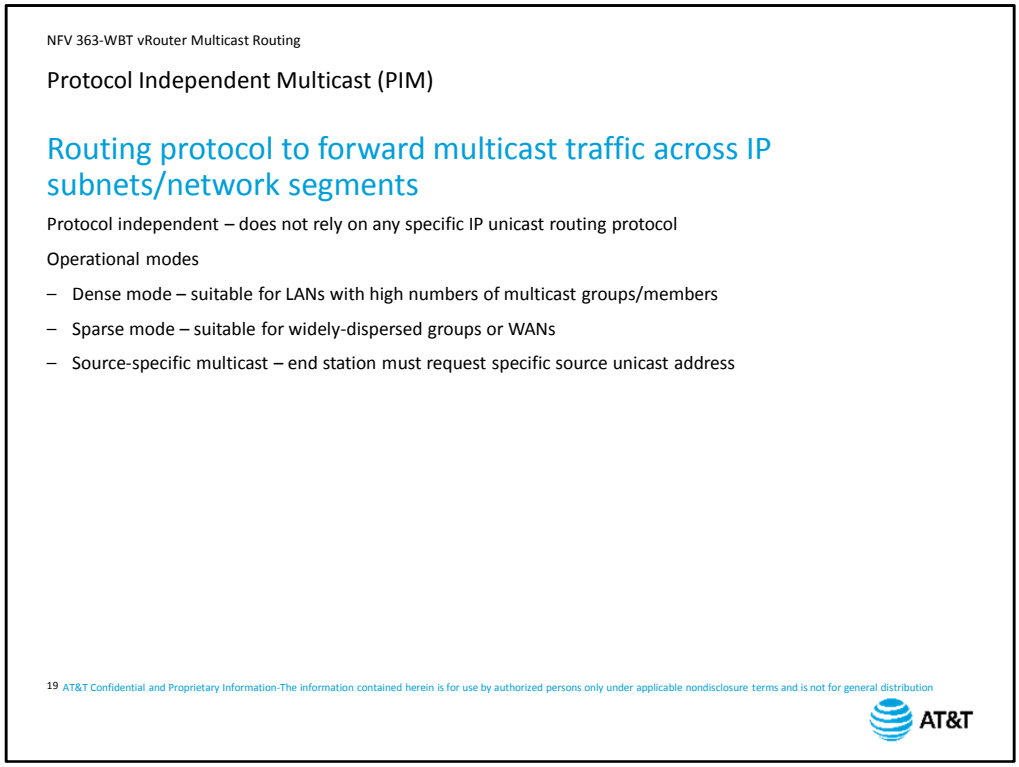

When you have multicast crossing multiple network segments, you need PIM to build the forwarding paths for those multicasts.
PIM does not rely on a specific unicast routing protocol like OSPF; instead, PIM uses the forwarding tables built by any Layer 3 unicast routing protocol to build forwarding paths for multicast traffic.
PIM has three operational modes:
Dense mode is typically used in a LAN environment where you have many multicast groups, each with many members.
Sparse mode is typically used with more widely-dispersed multicast groups, particularly over wide area networks.
Source-specific multicast requires that the end station specify which multicast source it wants to receive traffic from.

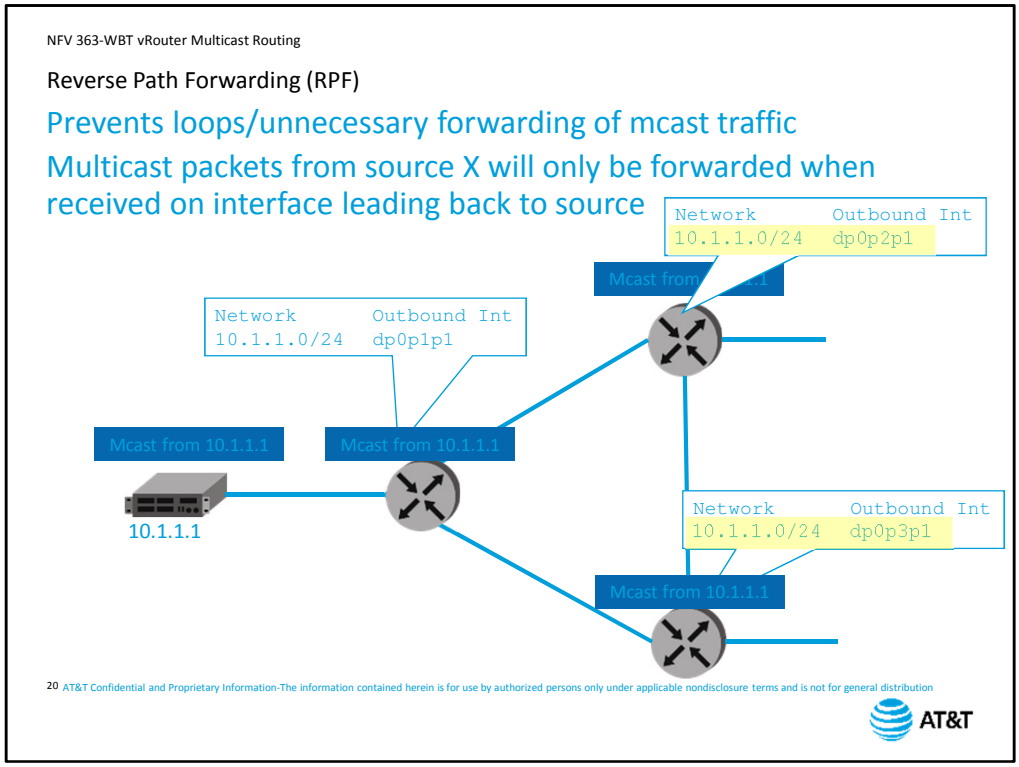

Regardless of mode, one key component of PIM is the idea of reverse-path forwarding. RPF prevents loops and unnecessary forwarding of multicast traffic by checking the source of each multicast packet.
A router will only forward a multicast packet if it was received on the interface that the unicast routing table associates with the source IP address.
Let's look at an example. Source 10.1.1.1 is sending out multicast packets.
When a router receives the packet, it looks at its routing table to find the reverse path back to the multicast source – in this case, data plane interface 1.
The packet arrived on interface 1, so the router forwards the packet out other interfaces based on the forwarding tables built by PIM. We will look at this in detail in a moment.
Note that the router will not forward the packet out interface 1, the interface where the packet arrived.
This process continues as the multicast continues to propagate throughout the network.
Routers check to see that the packet arrived via the correct interface,
Then forward out all other active interfaces as determined by PIM. Note that the routers here have received the same packet again through flooding.
However, the packet did not arrive via the interface leading back to the source, so the router knows the packets are a result of flooding and discards them.

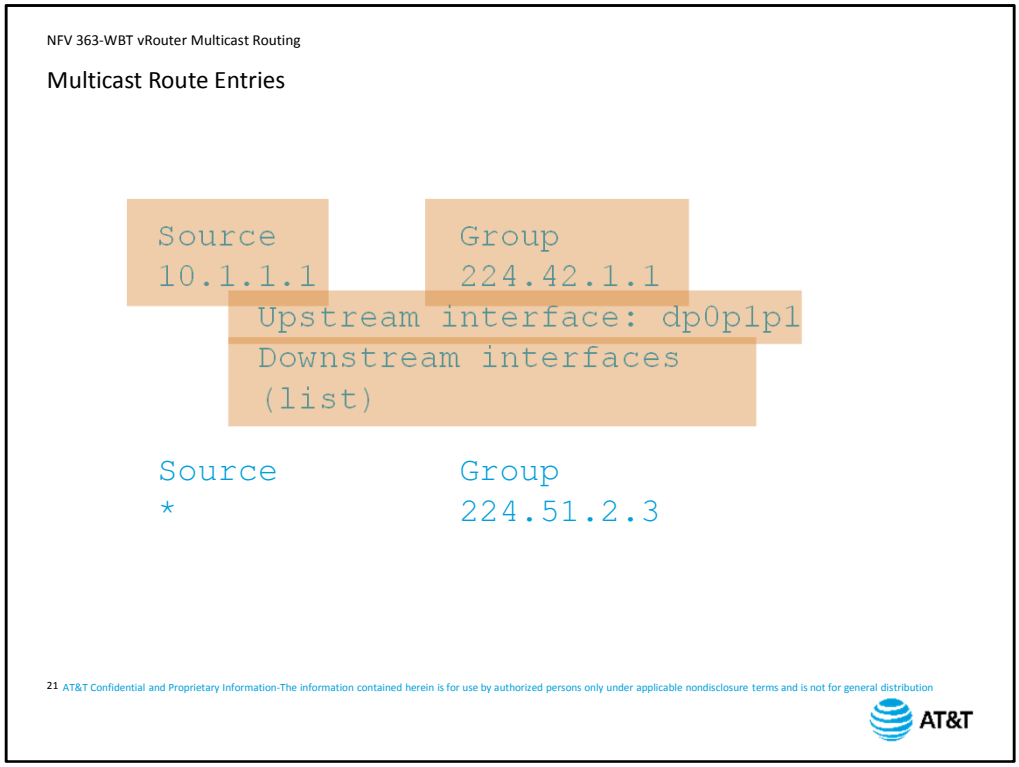

A complete multicast route entry will include the following information:
The IP address of the multicast source.
The multicast group address.
The interface that leads back to the multicast source.
And any interfaces where the multicast traffic is being flooded toward multicast group members.
In some cases, you will see a group address without a specific source address. In sparse mode, you may have downstream routers requesting membership in a group where a source has not yet been identified.

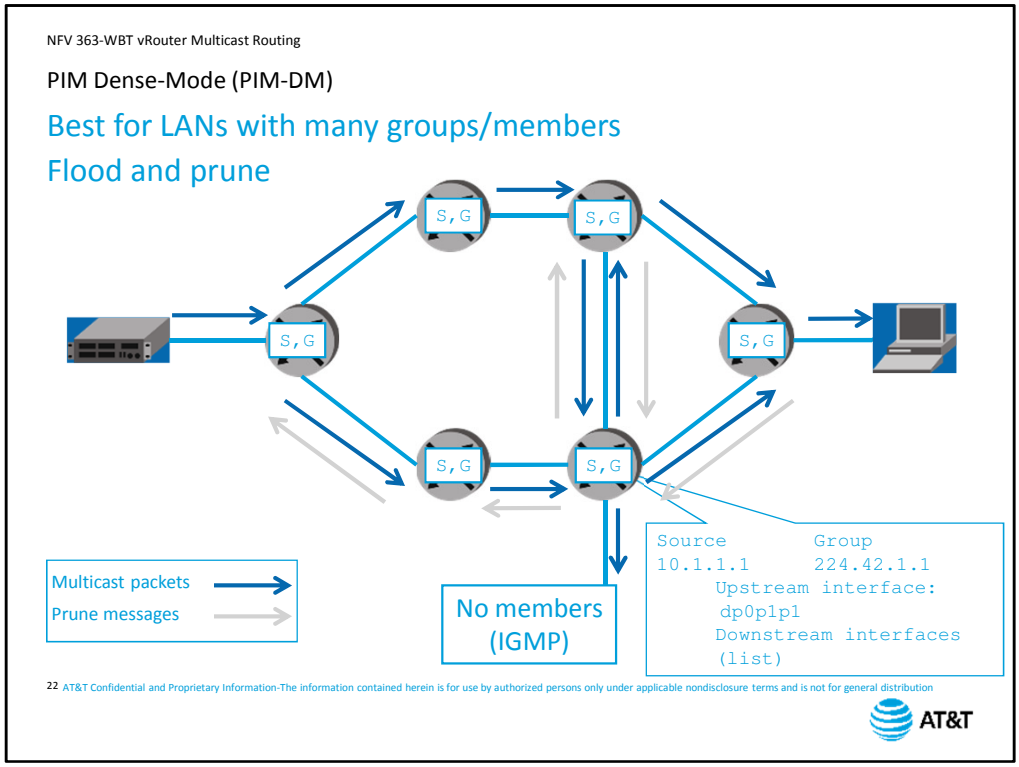

As we said earlier, PIM dense mode is designed for use in local area networks with many multicast groups and widely-distributed group members.

PIM dense mode works on a flood-and-prune basis to build a forwarding tree for multicast traffic.

When a multicast source begins transmitting, each router uses reverse-path forwarding to flood the multicast packet throughout the network.

As the packet passes through each router, it builds a Source-Group (S,G) table that lists the source IP address for each multicast group address, and the interface leading back to the source as determined by RPF. The table also lists downstream interfaces for forwarding the multicast. Those interfaces are determined based on PIM neighbors, and the local IGMP table for directly-connected group members.

The routers then begin sending prune messages to each other to trim the forwarding tree. The first prune messages are sent on all non-RPF interfaces to eliminate forwarding loops.

The routers then look at the list of members. If there are no downstream peers interested in the multicast, then the router sends a prune message back up the RPF path.

This 'no interested members' pruning continues back toward the source of the traffic until the tree is pruned to the minimum path necessary for forwarding to all interested group members. The resulting forwarding path is called a source-based distribution tree.

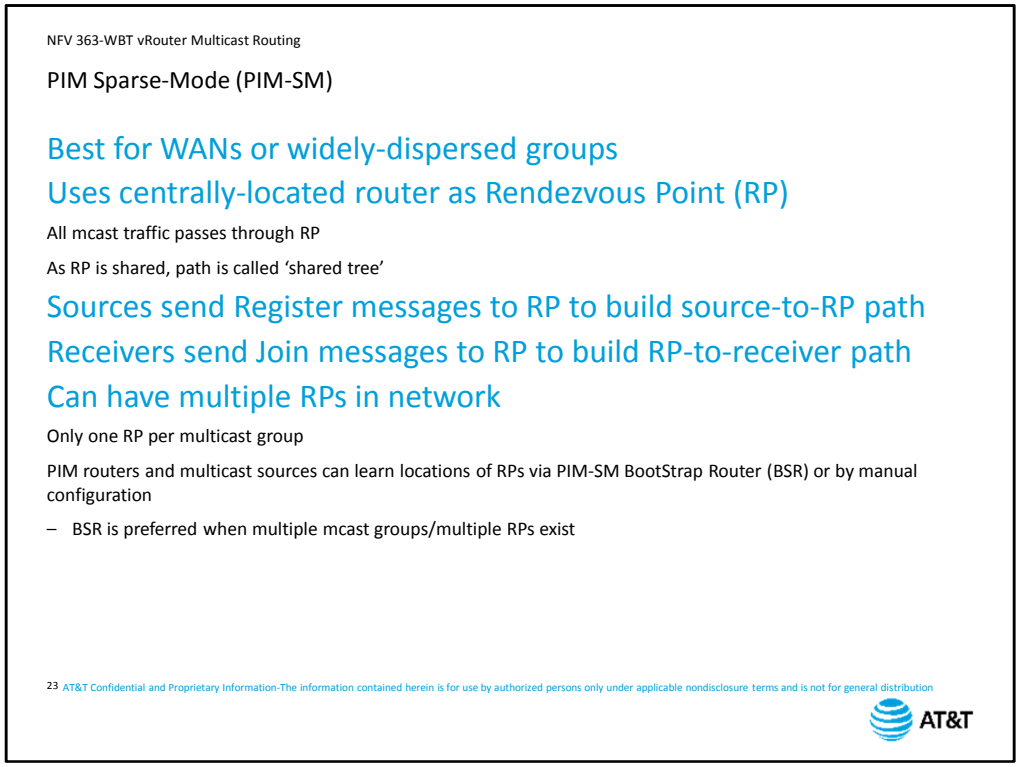

PIM sparse mode is designed for use on wide-area networks or with widely-dispersed multicast groups where dense mode path discovery would result in widespread flooding of multicast traffic.
Instead of building the forwarding tree from the multicast source outward, sparse mode uses a Rendezvous Point (RP) as the starting point of a multicast forwarding tree.
All multicast traffic will pass through the RP, so it needs to be both centrally-located and have the necessary forwarding capacity for your network's multicast requirements.
The resulting forwarding tree is called a shared tree, as opposed to the source-based tree created by dense mode.
When a multicast source comes online, it sends a Register message to the RP in order to build the source-to-RP path.
When a multicast receiver comes online, it sends a join message to the RP in order to build the RP-to-receiver forwarding path.
You can distribute the load throughout the network by designating multiple routers as RPs; however, each multicast group will only have a single active RP.
PIM-enabled routers and multicast source devices can learn the locations of the RPs in the network from a PIM Bootstrap router (BSR), or you can manually configure the addresses of each RP. If you have multiple multicast groups or multiple RPs in the network, using a BSR simplifies your overall configuration.

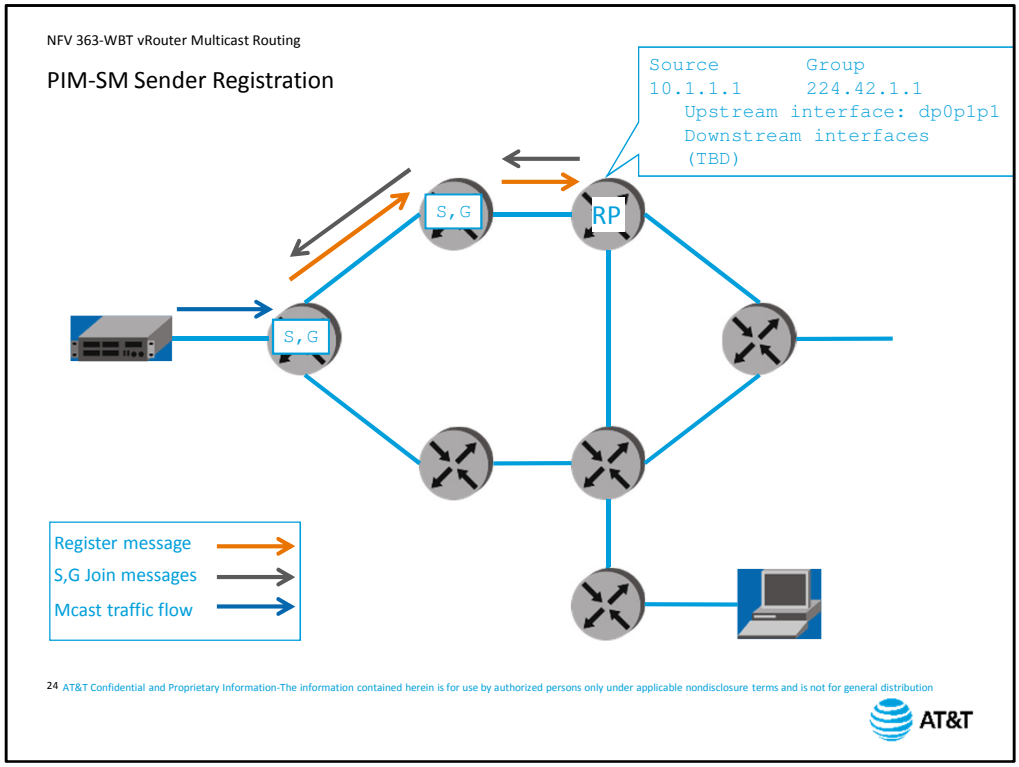

The router directly connected to the multicast source is responsible for registering the source with the RP. The Register message is a unicast message specifically for the RP, and it contains the source and group information for the multicast. The RP enters this in the Source-Group table, along with the upstream interface information.
The RP then sends a Source-Group Join message back up the reverse path to populate the Source-Group table in all multicast routers between the source and the RP.

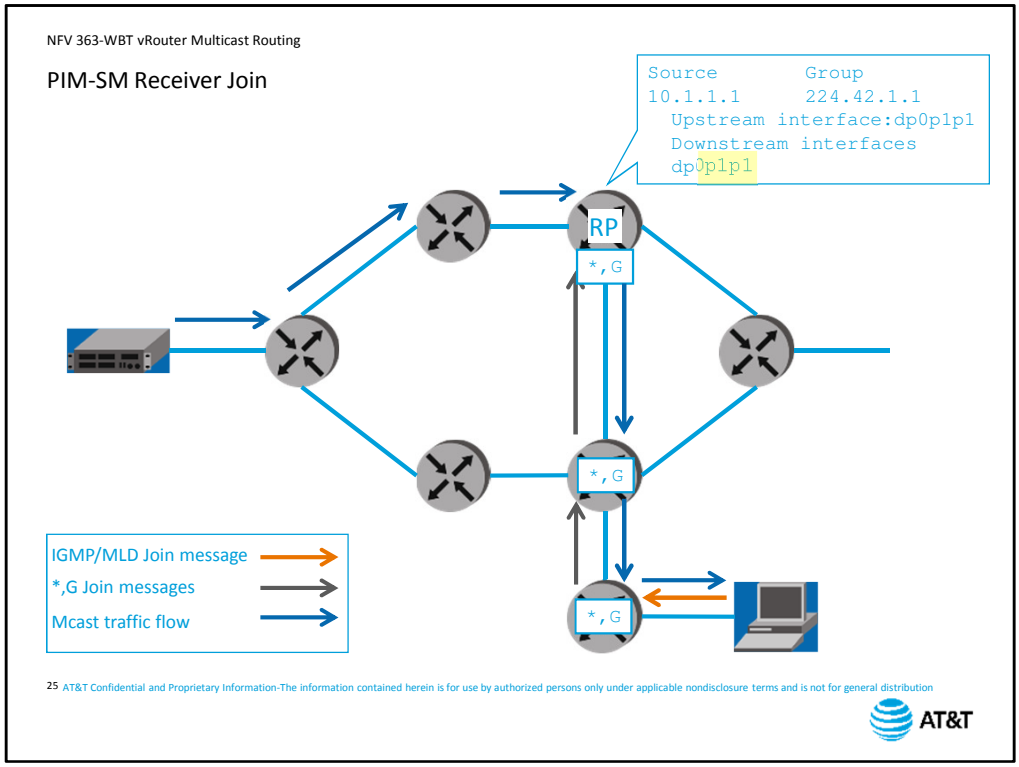

With the sender registered, we have half the forwarding path created. We'll complete the path when a receiver joins the multicast group.
When a router receives an IGMP or MLD join message for a downstream receiver, it adds the interface to its local Source-Group table. Note that the source is not specified at this point, unless the endpoint is running IGMP version 3 and has included a source specification in join message.
The router then sends a Group join message toward the RP.
All the routers in the path to the RP add the group to their own Source-Group tables. When the RP makes this addition, it begins forwarding the multicast group downstream, completing the forwarding path.

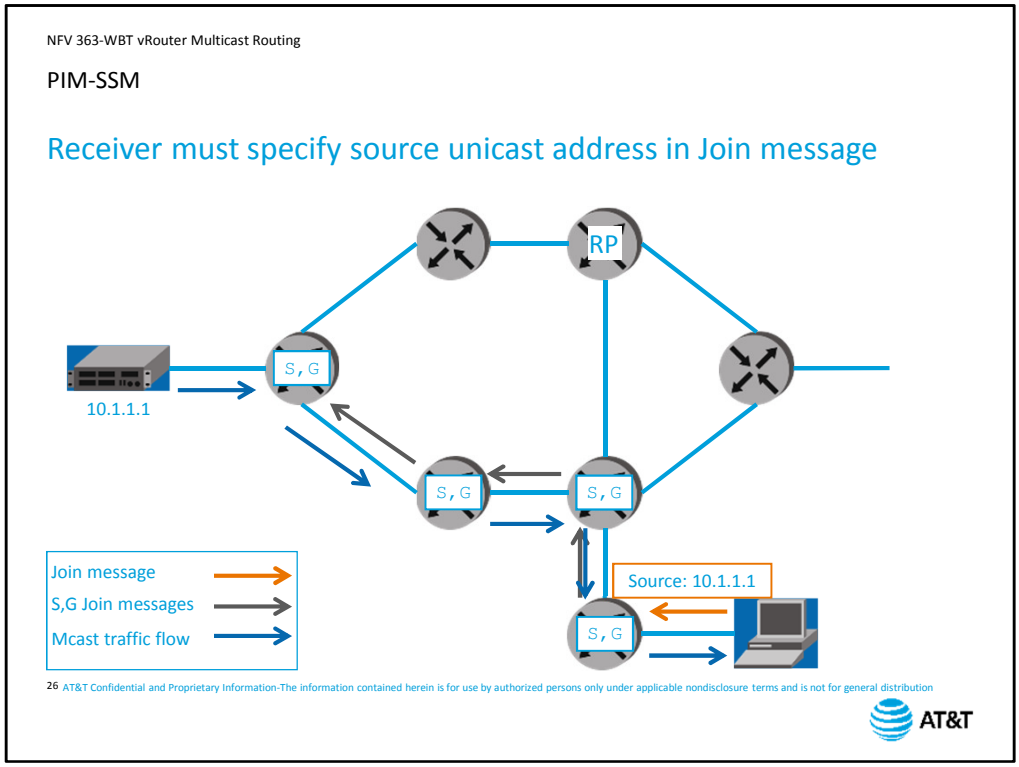

PIM source-specific multicast works in conjunction with either IGMPv3 or MLDv2. In both cases, the multicast receiver device specifies the unicast address that it wants to receive multicast traffic from.

The multicast router generates a Source-Group join message. It uses the unicast routing table to direct the *S,G* messages upstream toward the source. Each multicast router along the path adds the source, group data to its own table.

Multicast traffic then follows the path built by the *S,G* join messages.

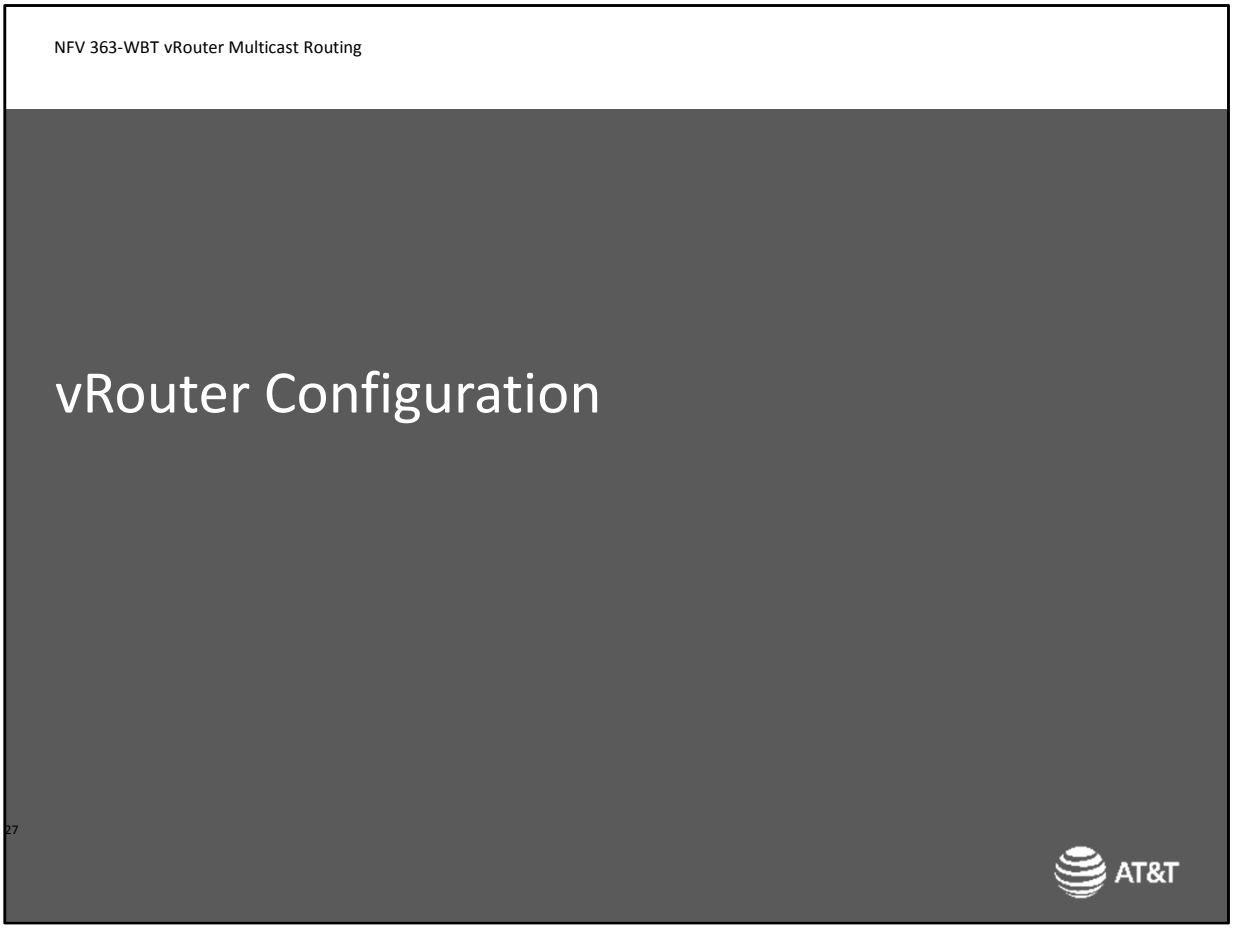

Next, we'll look at configuring multicast support on the vRouter.

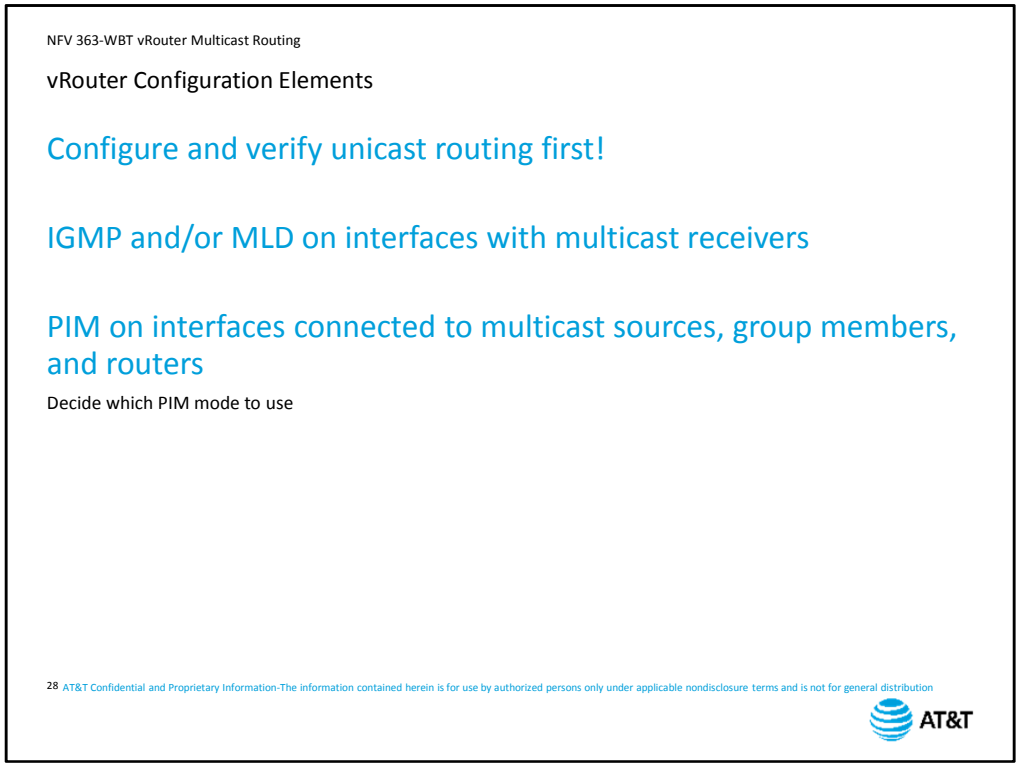

Before you can configure multicast routing, you need to make sure you have full reachability across your network for unicast traffic. Remember, multicast relies on the underlying unicast network.
Next, enable multicast group management on the interfaces connected to multicast receivers. You can enable IGMP, MLD, or both, depending on which version of IP is deployed in your network.
Next, enable PIM on the interfaces connected to multicast sources, receivers, and other multicast routers.
You will need to decide which version of PIM to use before you begin configuration – all routers within the network need to use the same version.

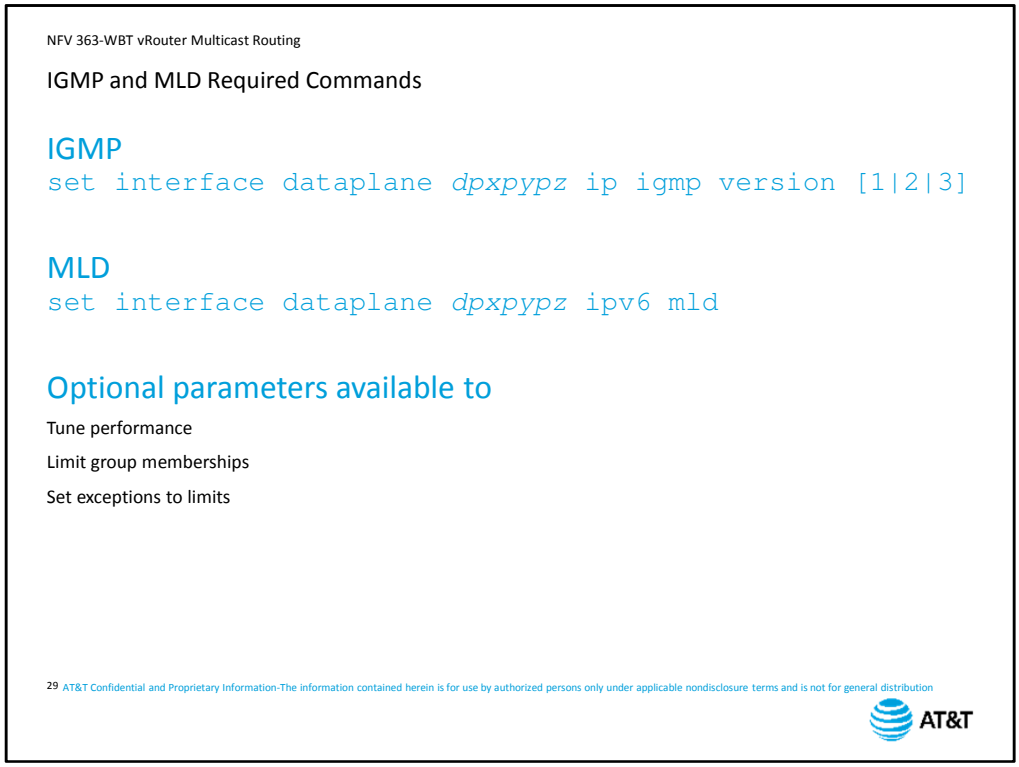

To enable IGMP, specify the interface name and the IGMP version.
To enable MLD, specify the interface name and MLD. These are the only required commands to enable the protocols.
For both protocols, there are many optional parameters, including parameters to tune performance, limit the size of multicast groups or which groups can be accessed on an interface.
There are also commands to set exceptions to those limits. We will not cover these commands in this course, but they are covered in the vRouter documentation, available online.

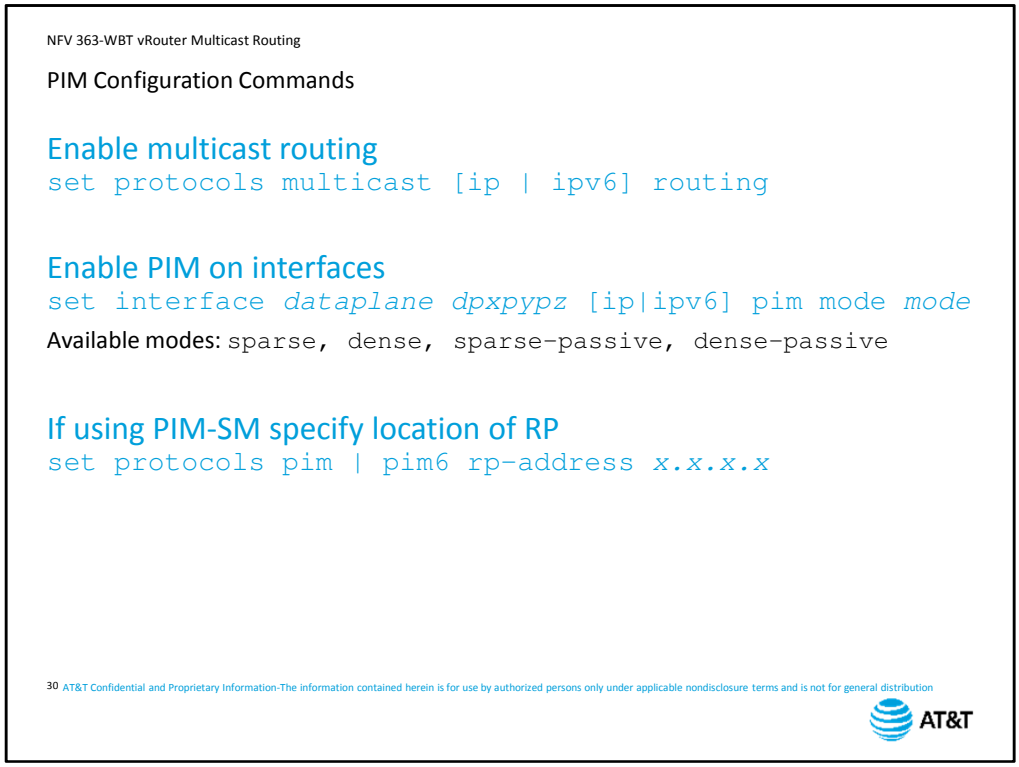

Configuring PIM requires two commands.
The first command is to enable multicast routing. Enable it for the version of IP running in your network.
Second, enable PIM on the appropriate interfaces, selecting either dense mode or sparse mode. If the interface is connected to multicast group members, but no other PIM routers, you can enable PIM in passive mode.
If you are using sparse mode, you also need to specify the address of the Rendezvous Point in your network.

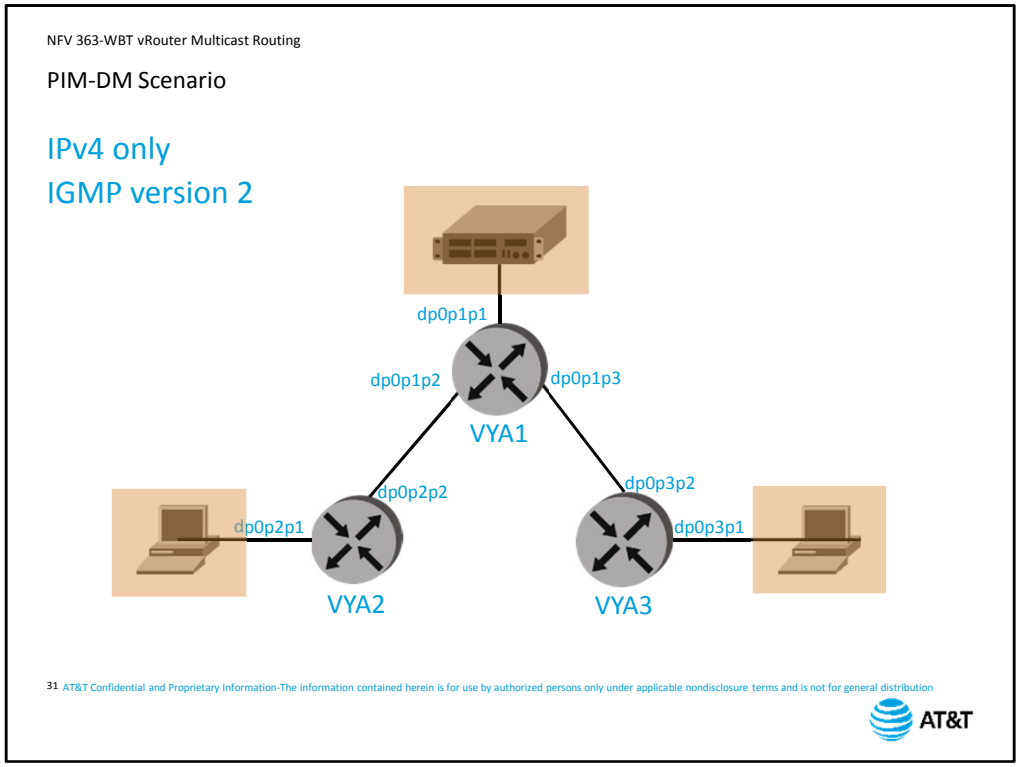

We will use this simple network to demonstrate the configuration of PIM dense mode.
The multicast source is connected to VYA1.
Multicast receivers are connected to VYA2 and VYA3.
We are only running IPv4 and we are configuring IGMP version 2.
You can open this scenario in a separate window by clicking on the *Attachments* tab and selecting *PIM-DM Scenario* from the list of attachments.

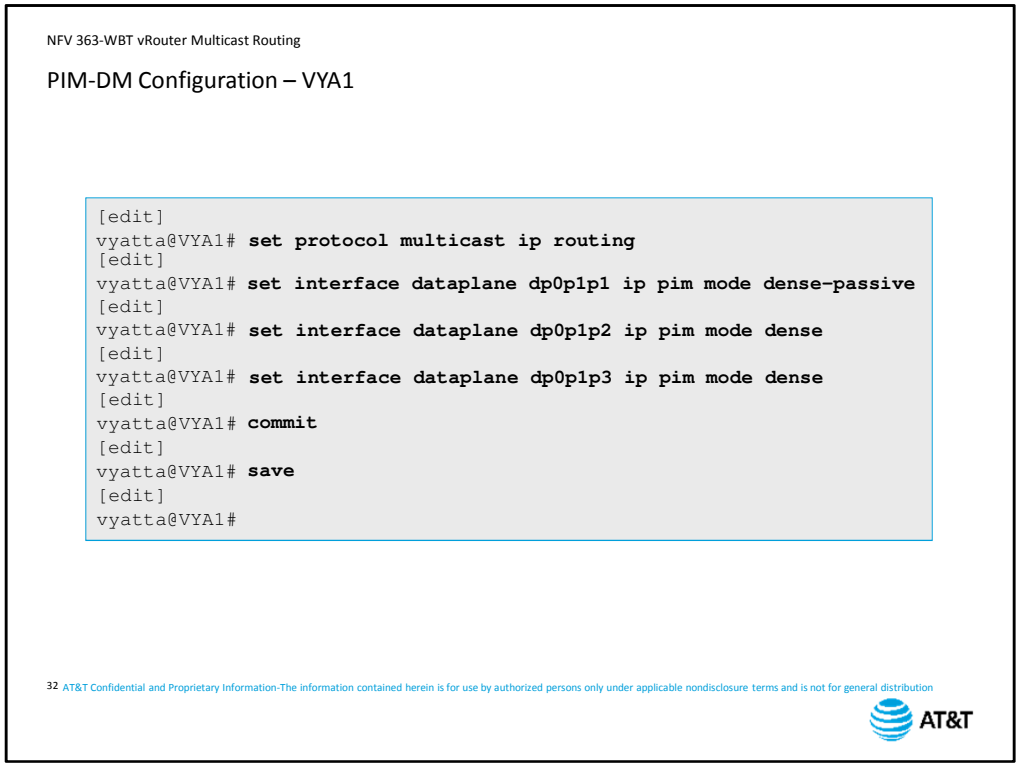

We will configure VYA1 first.
We enable multicast routing, then enable PIM on all three interfaces. On data plane 1, which connects to the multicast source, we enable PIM in passive mode.
On data plane 2, we enable dense mode and do the same on data plane 3.
We commit our configuration to make it active, then save it to make it permanent. This is all that is required to enable multicast routing.

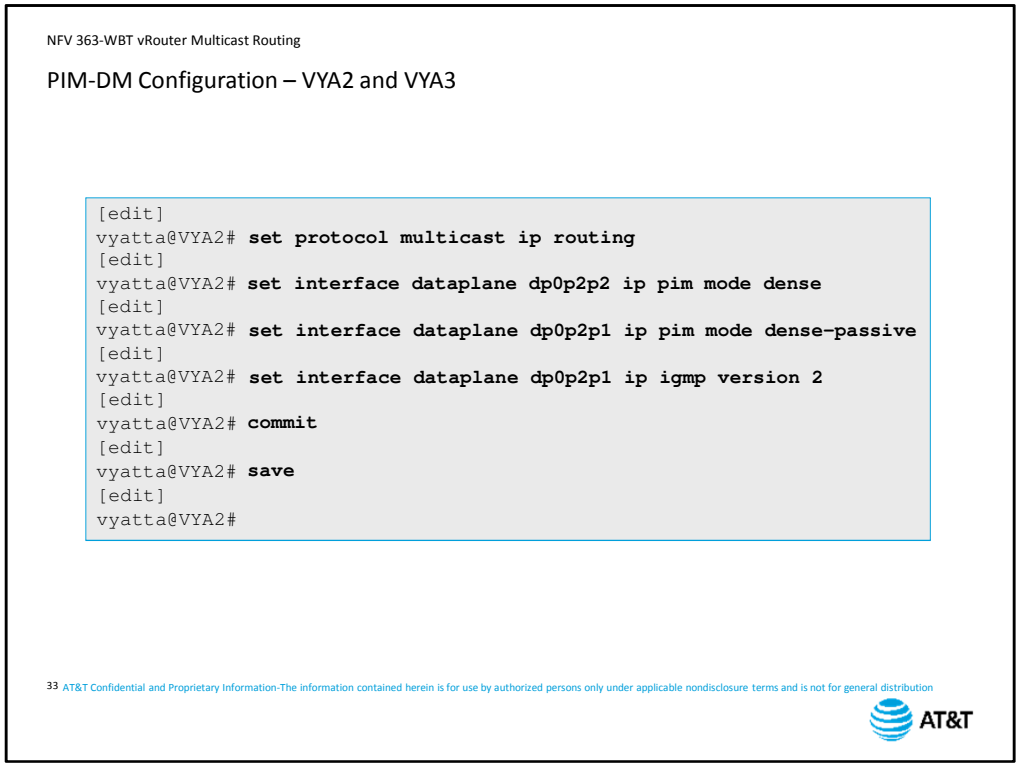

On VYA2 and VYA3, the steps are similar. The example displays the configuration for VYA2.
Enable multicast routing, enable PIM on the interface connected to the other multicast routers, and enable PIM in passive mode on the interface connected to the multicast group members.
We also need to enable IGMP on the interface connected to the group members.
Again, we commit, and save the configuration.

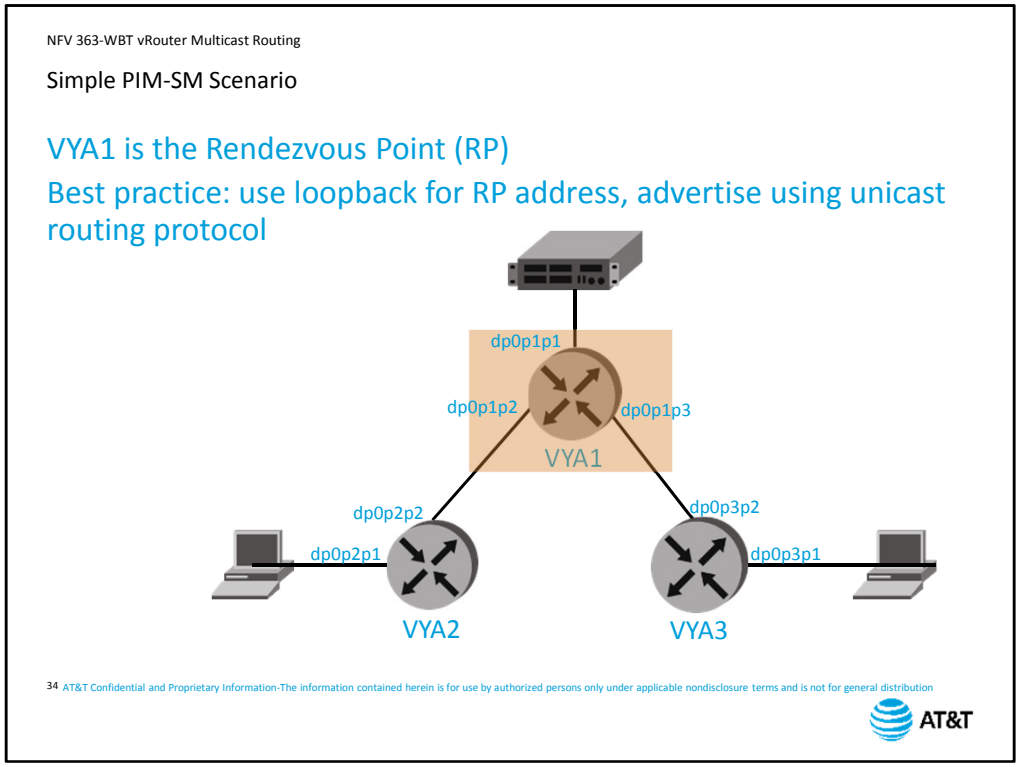

For sparse mode, you need to consider the location of the Rendezvous Point. In our network, VYA1 is the RP.

A best practice is to use the loopback address of VYA1 as the RP address rather than a physical interface address, then include the loopback address in the unicast routing protocol. This ensures that, as long as VYA1 is reachable via any routing path, the RP will be available.

You can open this scenario in a separate window by clicking on the *Attachments* tab and selecting *PIM-SM Scenario* from the list of attachments.

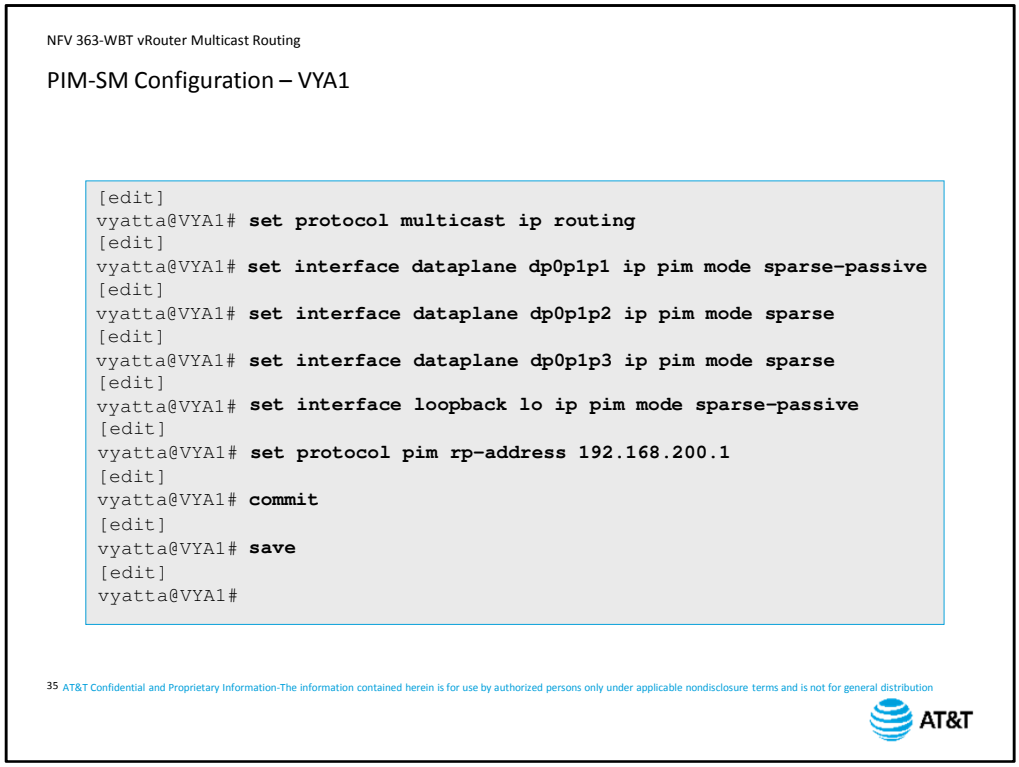

To start your sparse mode configuration, enable multicast routing, then configure the interfaces for PIM sparse mode. Note that we are using passive on the interface connected to the source, because there are no other PIM routers on that interface.
Because we are using the loopback address for the Rendezvous Point address, we need to activate PIM on the loopback interface.
Next, we specify the IP address of the Rendezvous Point – in this case, the address of the local loopback interface.
Then commit and save the configuration.

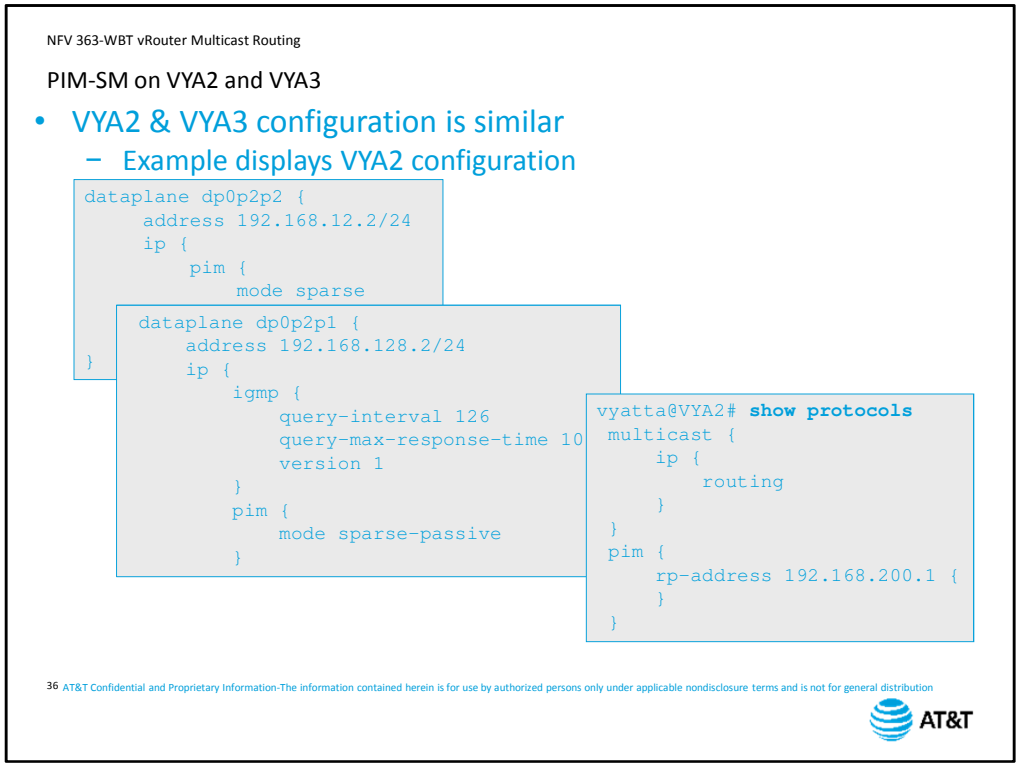

The PIM-SM configuration on VYA2 and VYA3 are similar. The example displays the configuration for VYA2.
On the other routers, we enabled PIM sparse mode on the interfaces connected to other PIM routers,
IGMP and PIM sparse in passive mode on the interfaces connected to multicast group members, and configure the Rendezvous Point address under the PIM protocol.

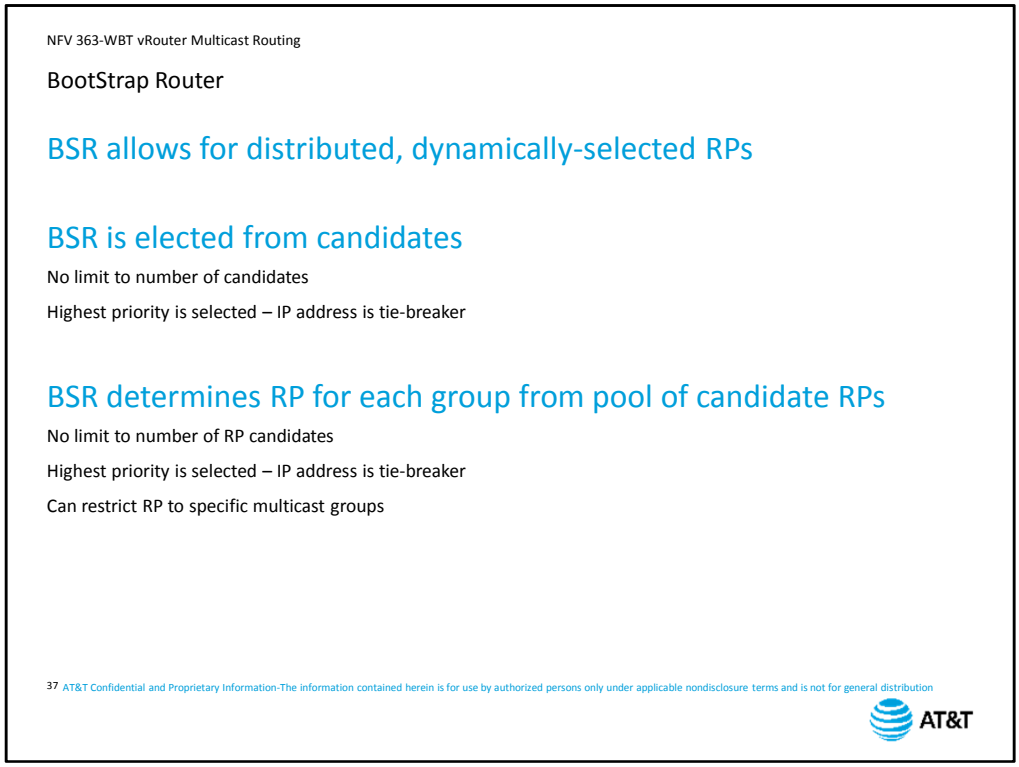

In larger networks, you do not use a static Rendezvous Point. Instead, you use a BootStrap Router, or BSR.
Instead of being statically configured, the bootstrap router is selected by the PIM protocol from a pool of candidate devices. There is only one active BSR in a network.
The BSR then determines the Rendezvous Point for each multicast group, again from a list of candidate devices. You configure the candidate devices for both the BSR and the RP.
In our scenario, we will configure all devices to be BSR and RP candidates. In a production network, you would determine which devices are best suited to maintain these functions based on traffic load, and only make those devices candidate BSRs or RPs.

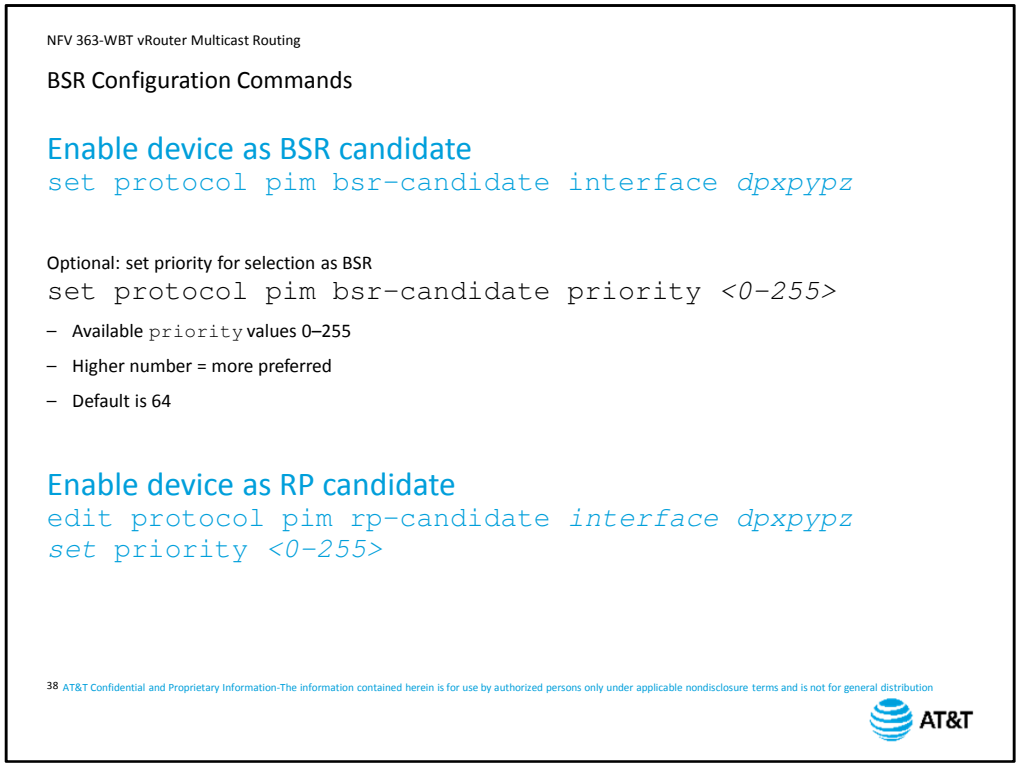

To configure a device as a bootstrap router, first, enable it as a BSR candidate. If you have PIM running on multiple interfaces, you should enable this on all the interfaces.
Optionally, if you have multiple BSR candidates in your network, you can set a priority value of 0 to 255. The higher the priority, the more likely the device will be selected as the BSR. The default value is 64.
To enable devices to be potential Rendezvous Points, set the `rp-candidate` command. Again, you can specify a priority value of 0 to 255.

Looking at the configuration on VYA1, we start with the same interface configuration, enabling PIM sparse mode on data plane and loopback interfaces.
Next, under the PIM protocol, we configure the loopback interface as a BSR candidate, and as an RP candidate.
As always, we commit and save our configuration.

Let’s now look at the commands available to verify multicast operations

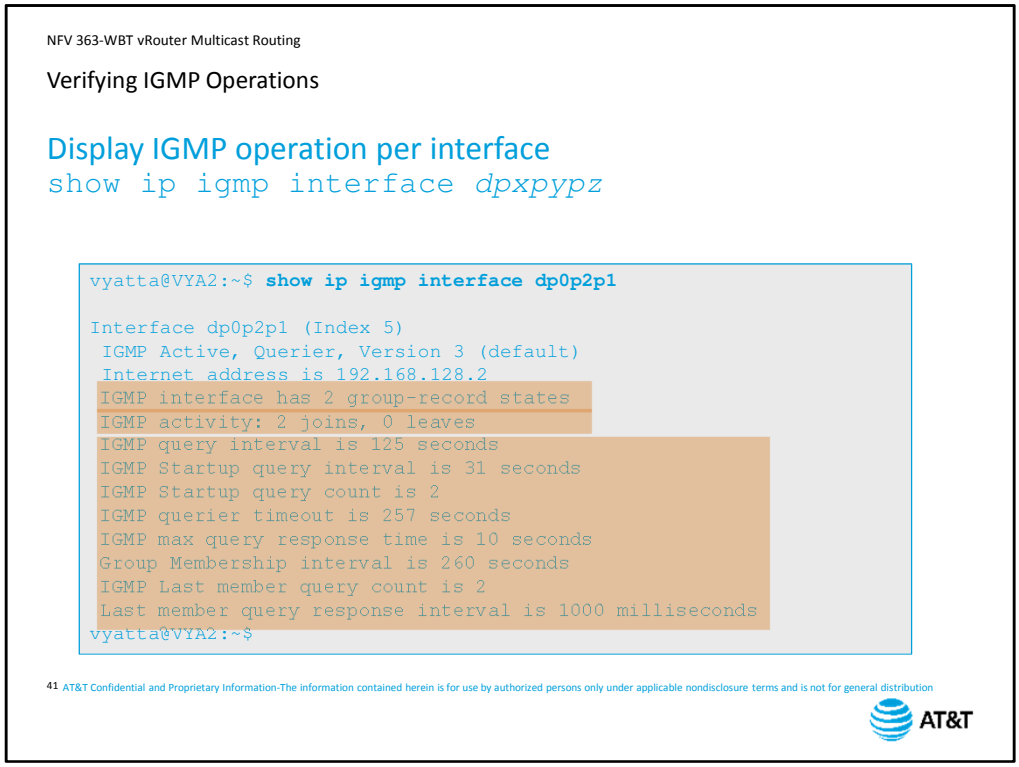

To verify that IGMP is working on an interface, use the command `show ip igmp interface` followed by the interface name. The output includes information on:
Number of current IGMP groups active on the interface.
Total end station IGMP activity.
And current timer and limit configuration.
As we mentioned earlier in the configuration section, you can adjust many of these settings as needed by your network.

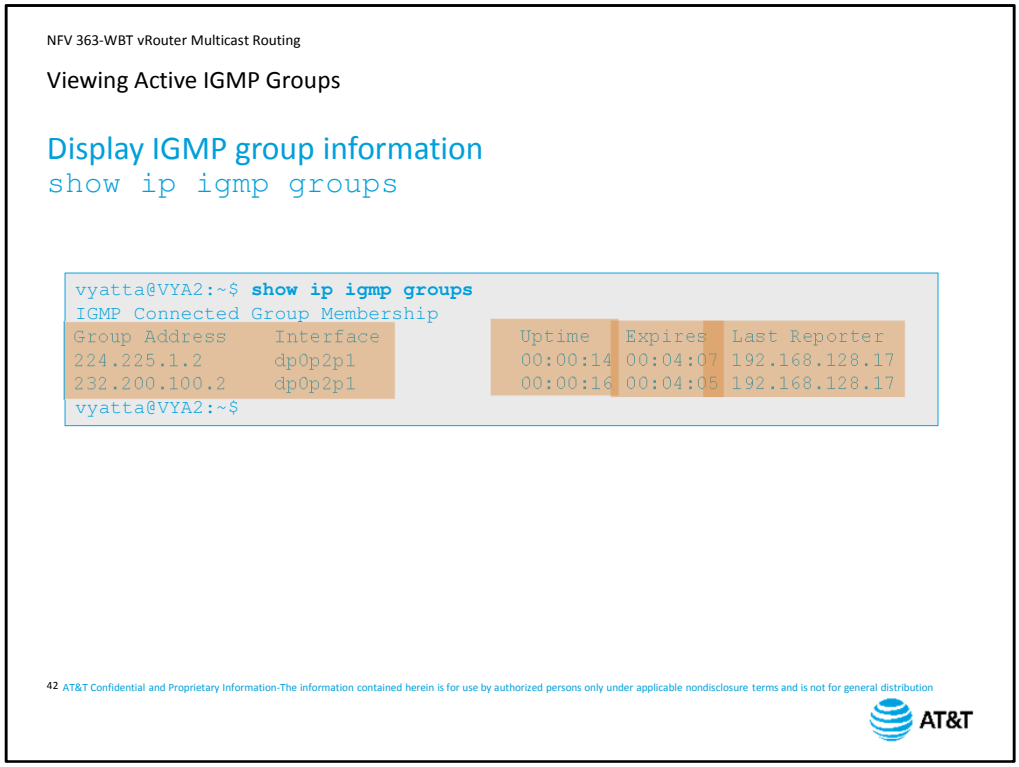

To view all active IGMP groups on a vRouter, use the command `show ip igmp groups`. The output will list all multicast group addresses currently active on the device, and the interface(s) with group members.
Also displayed is how long the group has been active, when the group will expire, and the IP address of the last device to send a membership report.

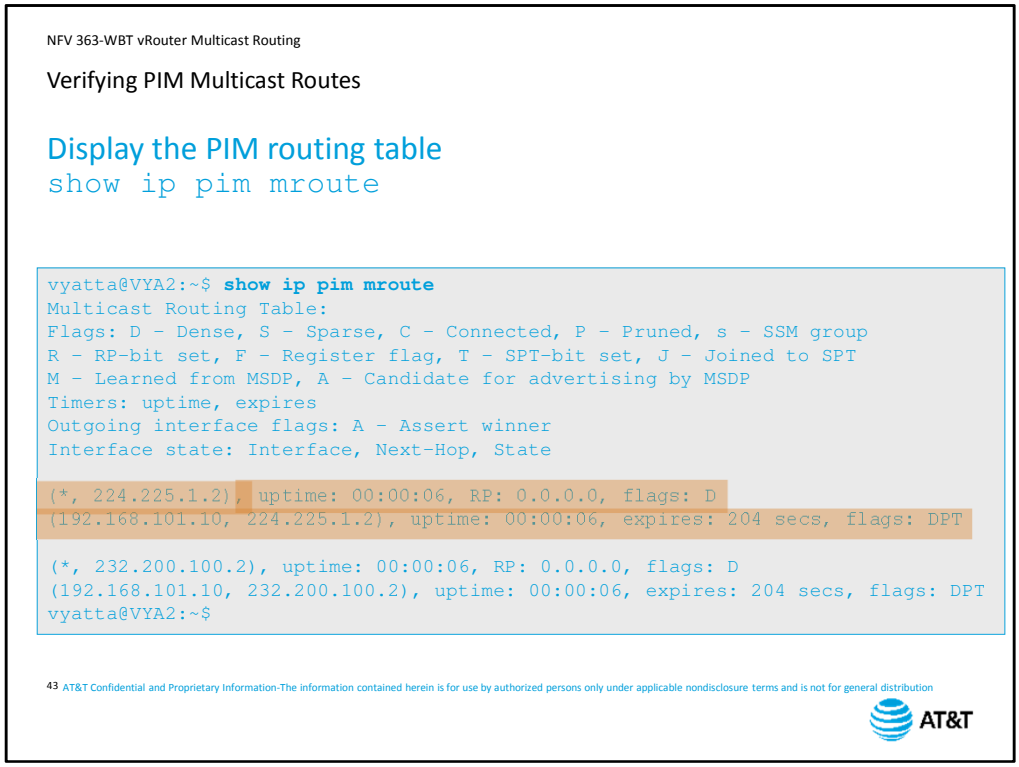

To verify that your device is learning multicast routes, use the command `show ip pim mroute`.

The route entry begins with the multicast group address. Because a group can have multiple sources, the route entry header does not include a specific source.

For the multicast group, you can see the entry uptime, the Rendezvous Point if applicable, and flags that describe the route entry – in this case, a *D* flag because this route was learned via PIM dense mode.

Under the multicast group, you will have entries for an individual host or hosts that are the source for the group address, with uptime, route expiration, and flags for each specific multicast source. The flags for this specific multicast source indicate that it was learned via dense mode, the route has been pruned, and that the SPT-bit has been set – in other words, the packets have been received by the shortest path tree.

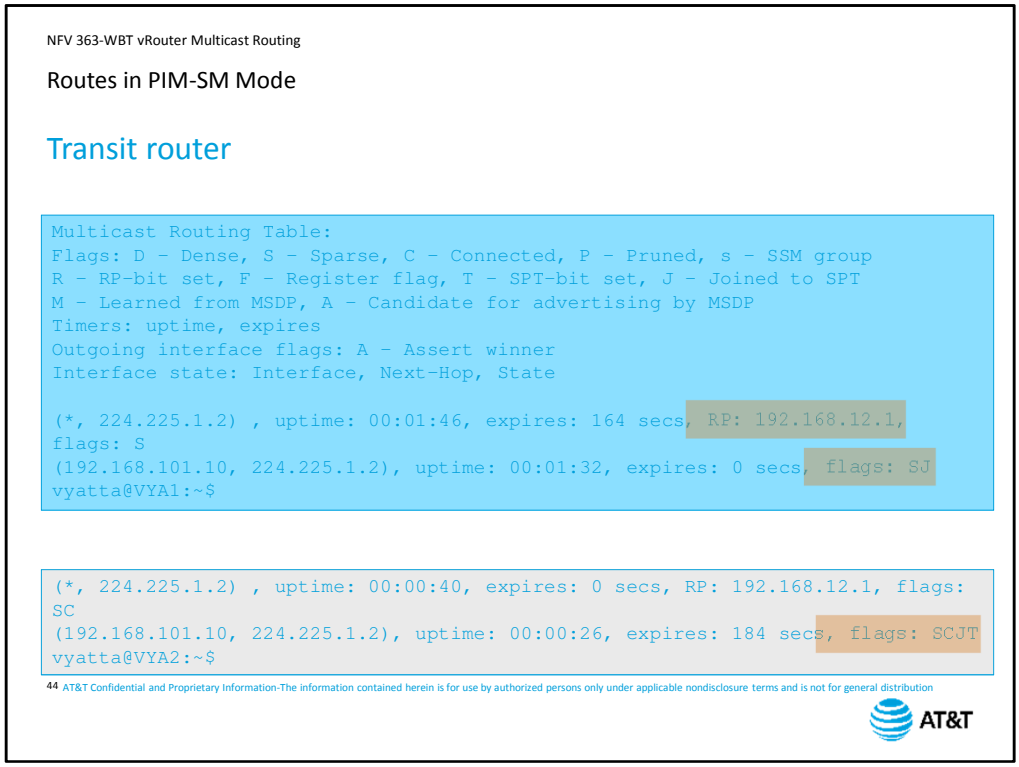

In a sparse mode network, you will see the Rendezvous Point for each multicast group and a join flag that indicates the device has explicitly joined the shortest path tree.
On devices with multicast group members, the flags will also include the connected bit, indicating that there are directly-connected receivers, and the SPT-bit flag.

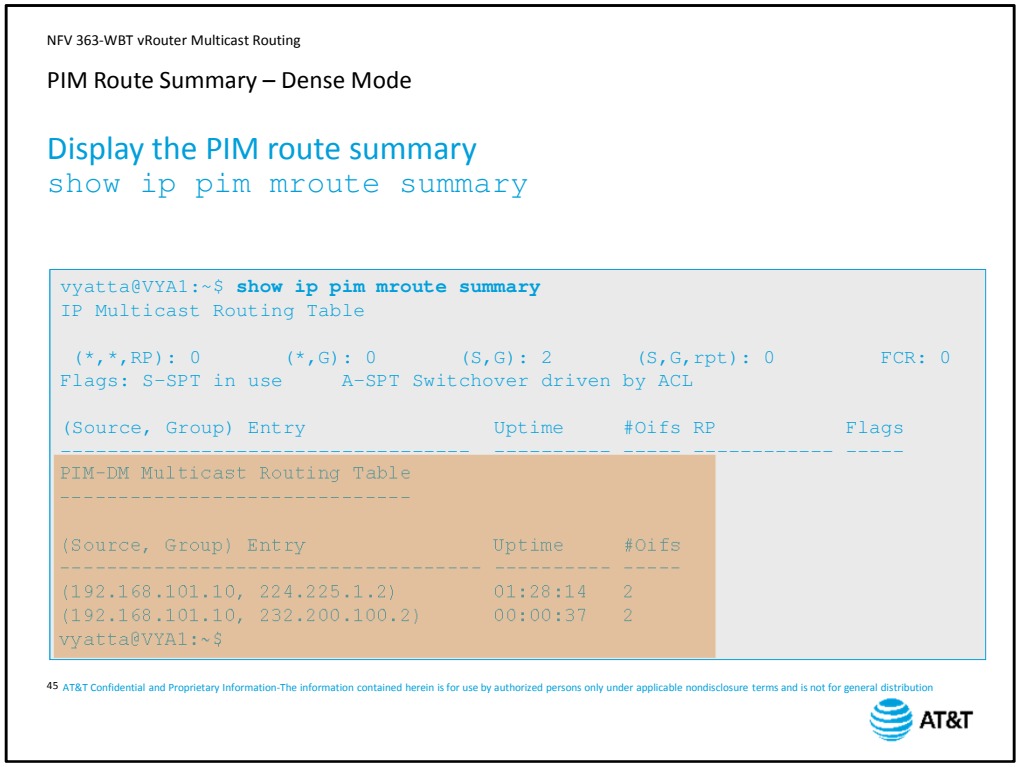

You can see a summary of this information by adding the keyword `summary` to the same show command. The output on the screen is for the dense mode configuration. Note that the table headers are for a sparse mode network, but because we are in dense mode, there is no data in those fields.
Instead, the output adds a header indicating dense mode routes (PIM-DM), then lists the known multicast sources and groups, plus the number of interfaces where the route is active.

PIM Route Summary – Sparse Mode

```
vyatta@VYA1:~$ show ip pim mroute summary
IP Multicast Routing Table

 (*,*,RP): 0       (*,G): 2       (S,G): 2       (S,G,rpt): 2         FCR: 0
Flags: S-SPT in use     A-SPT Switchover driven by ACL

(Source, Group) Entry                Uptime     #Oifs RP            Flags
-------------------------------------------------------------------------
(*, 224.225.1.2)                     00:00:31   1     192.168.12.1     S
(192.168.101.10, 224.225.1.2)        00:00:22   0
(192.168.101.10,224.225.1.2,rpt)     00:00:22   0     192.168.12.1
(*, 232.200.100.2)                   00:00:30   1     192.168.12.1     S
(192.168.101.10, 232.200.100.2)      00:00:29   1
(192.168.101.10,232.200.100.2,rpt)   00:00:29   0     192.168.12.1
vvyatta@VYA1:~$
```


In sparse mode, the output is quite different.
We see the group, then the specific Source,Group (S,G) combination, then a separate entry indicating which Rendezvous Point is reporting the S,G combination.

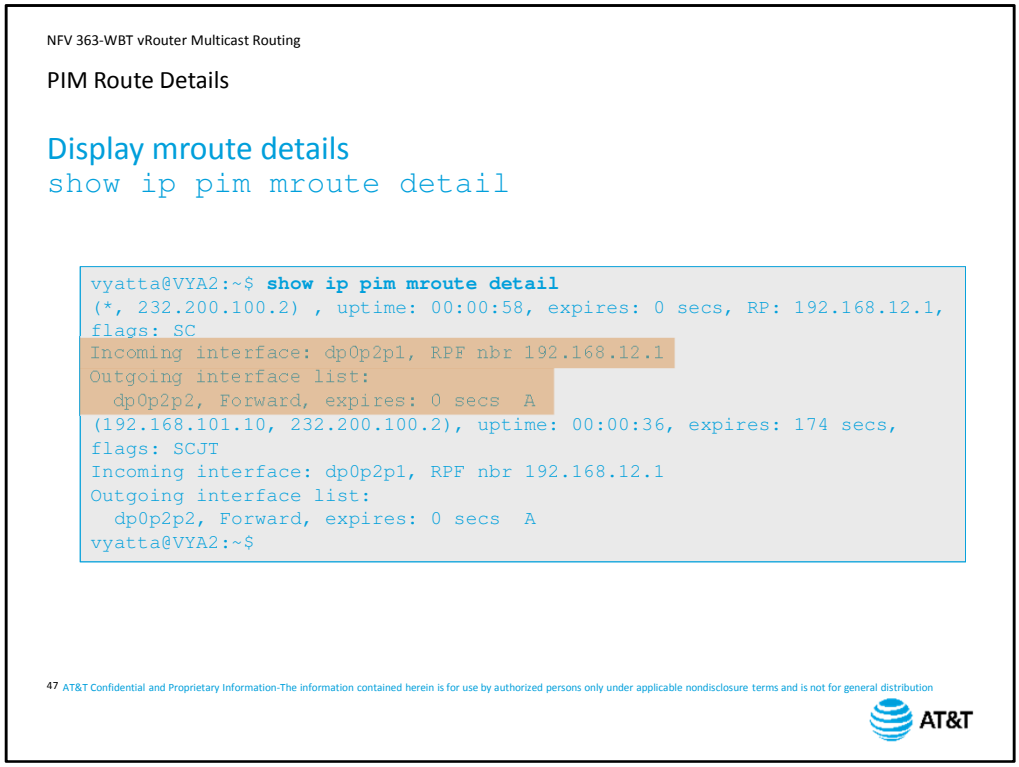

When you look at the detailed output for a route, you will see information about the upstream path back to the multicast source, and the downstream path to the multicast group members.

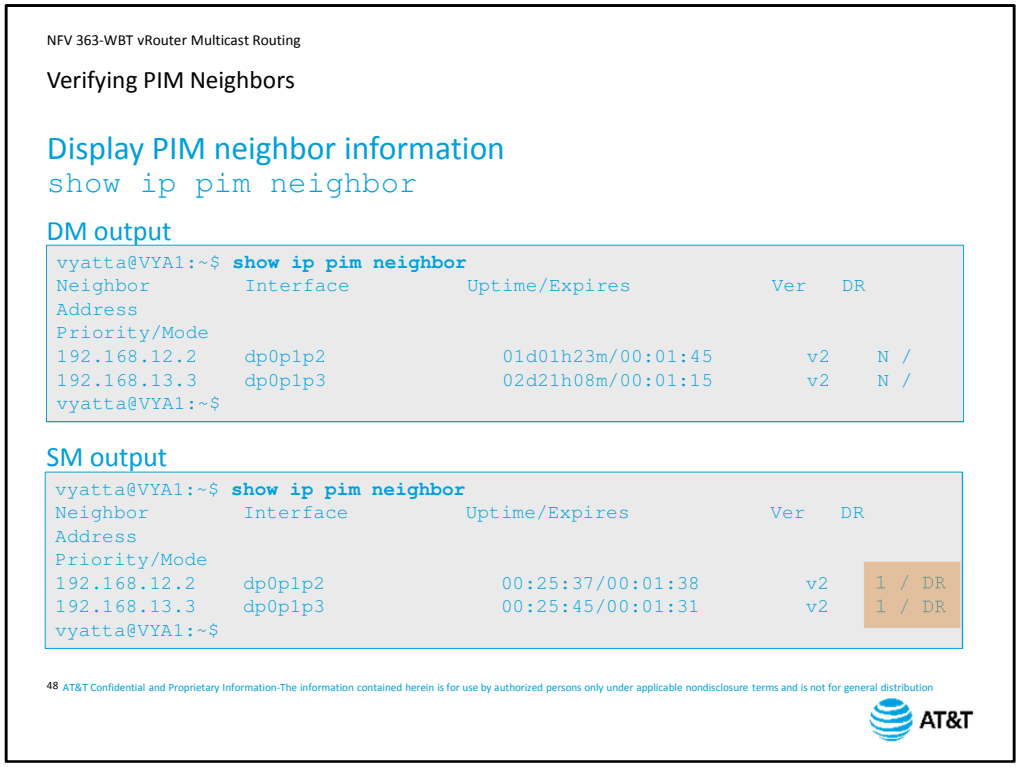

You can verify connections to PIM neighbors with the command `show ip pim neighbor`. The output is essentially the same for both dense mode and sparse mode. The one difference in sparse mode is that, for each group, the router will report whether it is functioning as the Designated Router for that group. In the context of PIM sparse mode, a Designated Router is a router on a segment that includes the group address in its periodic Join/Prune messages to the Rendezvous Point. This is only meaningful if there is more than one router on a segment with multicast group members, and ensures that the rendezvous point does not receive multiple join/prune messages for a given network segment.

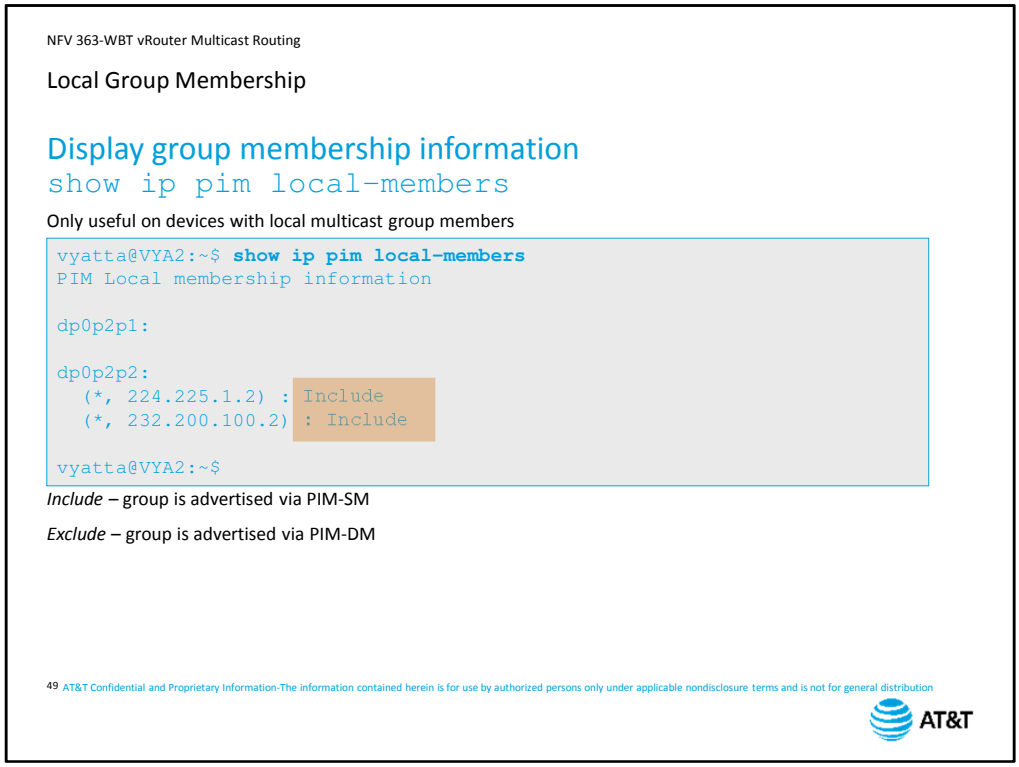

You can view a list of which multicast groups have local members using the command `show ip pim local-members`. The output here is essentially a duplicate of the output of `show ip igmp groups`.
However, there is one thing to note: each group will have a label indicating “include” or “exclude”.
Include means that the path to the group source has been learned via PIM sparse mode, so the group needs to be included in the periodic PIM Join/Prune messages to maintain the path.
“Exclude” means that the path has been learned via PIM dense mode because dense mode paths are already pruned to a minimum forwarding tree, there is no need to include information about local group membership in PIM Join/Prune messages.

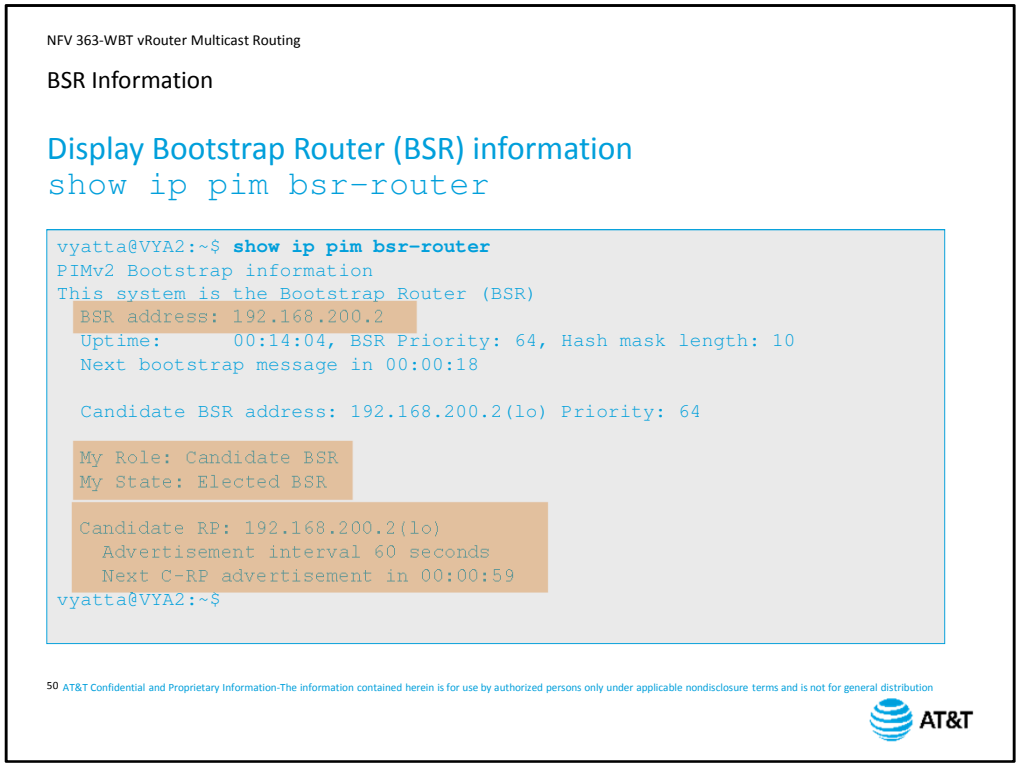

You can display information about the Bootstrap router on any BSR candidate device. The output tells you the address of the bootstrap router, the local device's role (candidate or elected BSR), and any locally-configured candidate Rendezvous Point addresses.

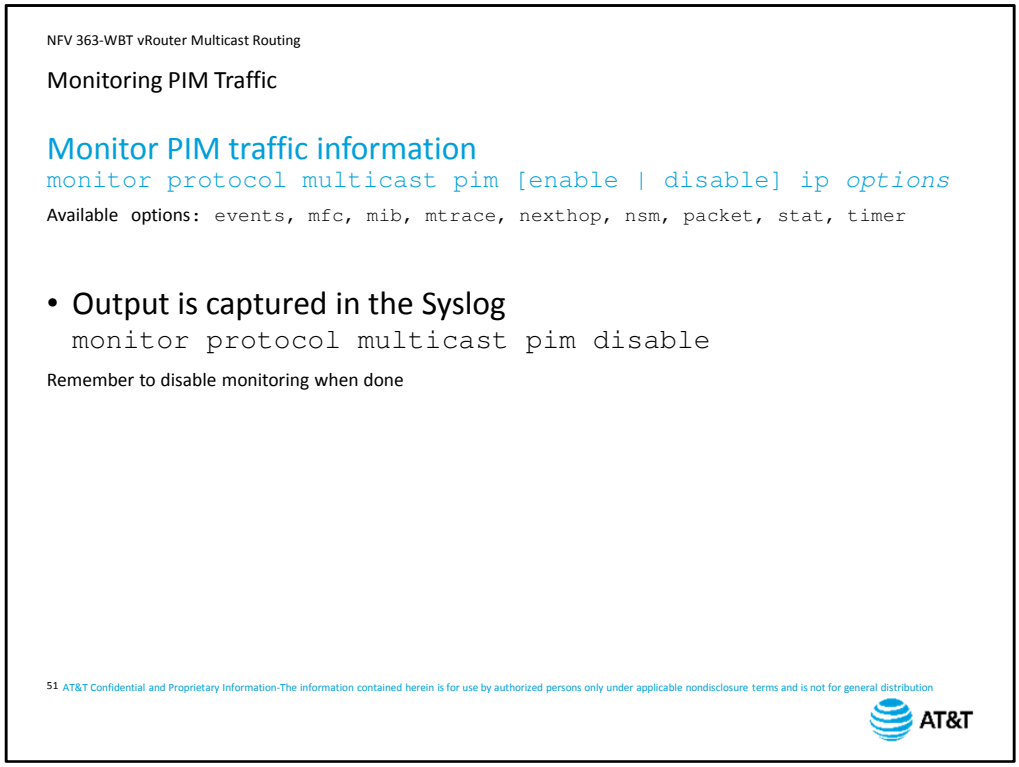

You can monitor PIM traffic using the `monitor` command. Options allow you to select the specific PIM packets you wish to monitor.
As with all monitoring, the output is automatically captured in the Syslog.
If you want to view the output in real time on your console, repeat the `monitor ip pim` command to place your console into viewing mode.
Remember that exiting the on-screen monitoring does not disable the background logging process. You must manually disable monitoring. For more information on using vRouter system logs and monitor commands, please refer to the *AT&T Vyatta 5600 vRouter Software Documentation* on www.AT&T.com.

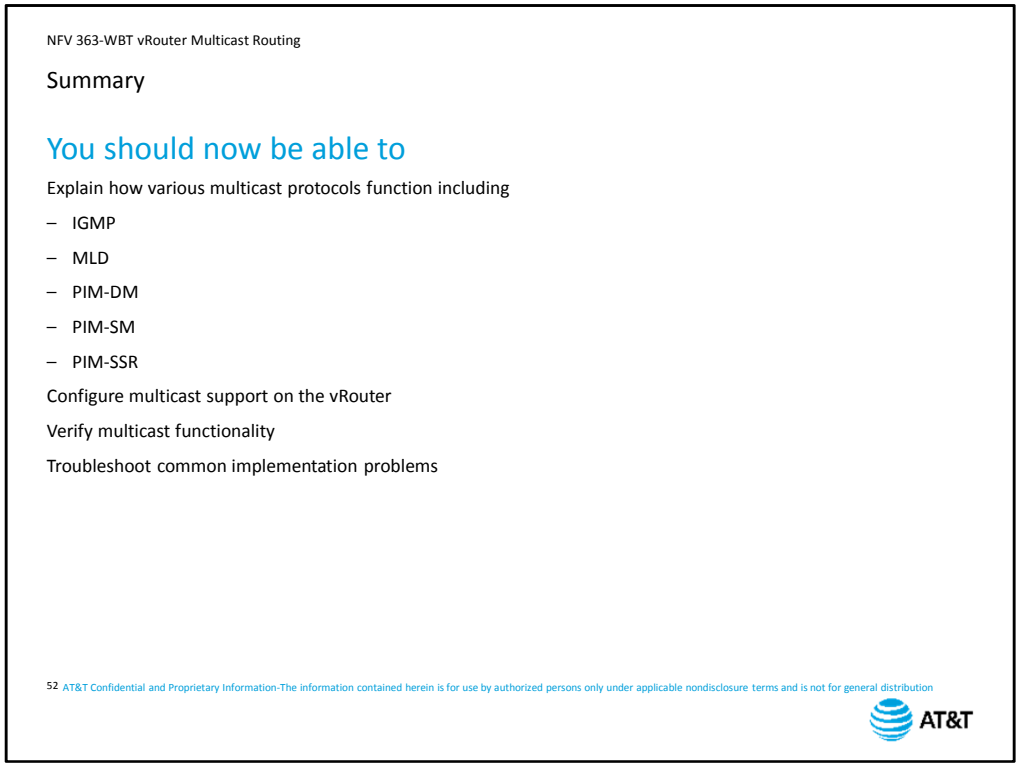

This concludes the AT&T Vyatta 5600 vRouter Multicast Routing course.
You should now be able to:

- Explain how various multicast protocols function, including IGMP, MLD, PIM-dense mode, PIM sparse mode, and PIM source-specific multicast
- Configure multicast support on the vRouter
- Verify multicast functionality
- Troubleshoot common implementation problems

We hope that this information has been useful to you, and that you will take additional AT&T University courses in the future.
Thank you.

NFV 363-WBT vRouter Multicast Routing

End of Course – vRouter Multicast Routing