Revision 0515

NFV 340-WBT: AT&T Vyatta 5600 vRouter OSPF 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."

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

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

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

Objectives

After completing this module, you will be able to

- Explain basic OSPF operations
- Configure the vRouter for single-area OSPF networks
- Verify and troubleshoot OSPF 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

3 AT&T

After completing this module, you will be able to:

- Explain basic OSPF operations.
- Configure vRouter software for a single-area OSPF network.
- Verify and troubleshoot OSPF 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

Let's begin with an overview of OSPF and how it works.

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

OSPFv2 Overview

- Open Shortest Path First version 2 (OSPFv2) for IPv4 is a link-state protocol
  - Routers advertise local links only to all other routers using Link State Advertisements (LSA)
  - Each router builds a Link State Database (LSDB)
  - Each router runs Shortest Path first (SPF) algorithm on LSDB
- Metric is cost
  - 100/link bandwidth
  - Any link > 100 Mbps has cost of 1
  - Effectively hop count in today's high-speed networks
- Hierarchical routing protocol
  - Routes can be summarized at area boundaries

AT&T Confidential 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

OSPF stands for Open Shortest Path First. Open, because it is an open standard, not proprietary; and shortest path first because the routing algorithm selects the shortest path between two nodes.

OSPF is a Link State protocol. This means that each router advertises the networks or subnets it is directly connected to, and only those networks. The advertisements go to all other routers running OSPF.

Each router collects these link state advertisements into a database, and runs the Shortest Path First algorithm to determine the best path to each learned network or subnet.

The metric used to calculate the best path is called cost.

Cost is calculated by dividing 100 by the link bandwidth. When OSPF was designed in the early '90s, the fastest links available were 100 Mbps links, which had a cost of 1. Standard 10 Mbps links had a cost of 10, and so on.

Any link with a bandwidth greater than 100 Mbps has a cost of 1; there is no concept of a fractional link cost in OSPF.

In today's high speed networks, using this default cost calculation results in most links having a cost of 1, effectively reducing cost to hop count. You can modify the default calculation to provide a more accurate picture of bandwidth on your network, but you must modify every OSPF router in your network.

Finally, OSPF is a hierarchical routing protocol. This means that you can summarize the amount of routing information each router must maintain.

AT&T Confidential 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 OSPF, routes can only be summarized at area boundaries. We will look at the concept of areas later in this course, and look at summarization in detail in the *AT&T Vyatta 5600 vRouter Multi-area OSPF* 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

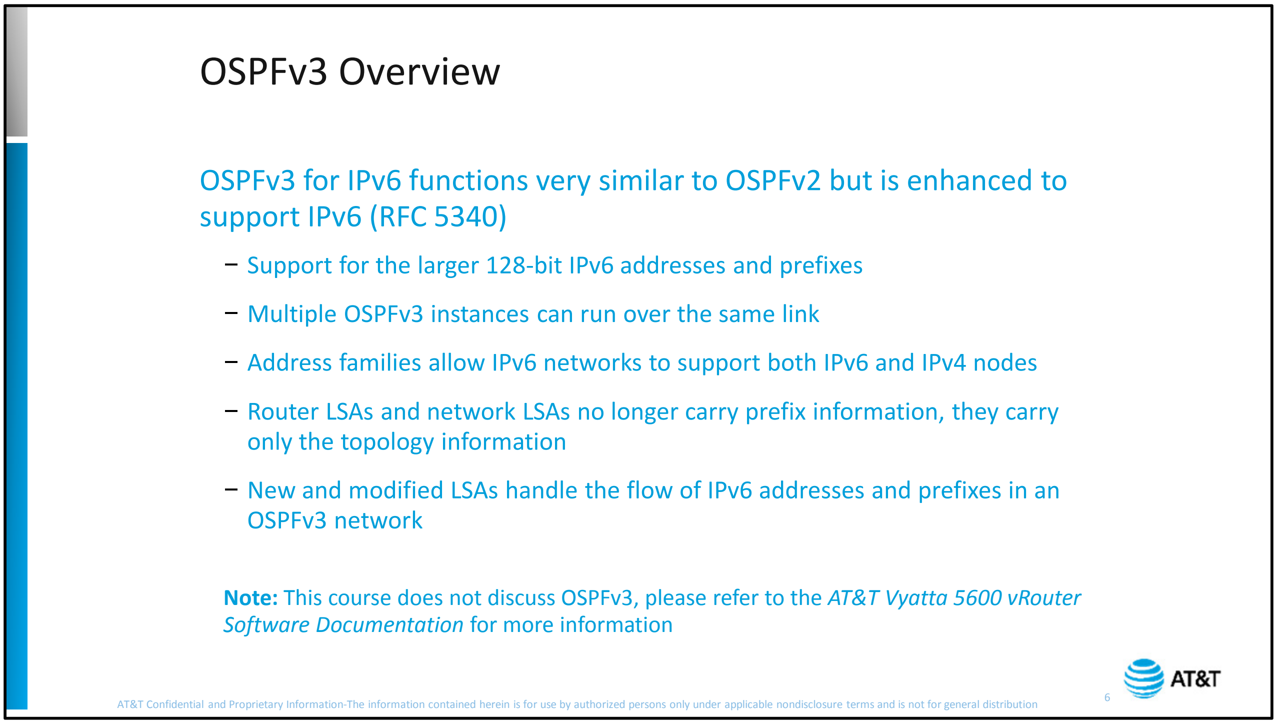

OSPFv3 adds support for IPv6 in the OSPF routing protocol and is described in RFC 5340. OSPFv3 works similar to OSPFv2, however, some significant changes exist in OSPFv3 for IPv6, we list some of these enhancements here:

- OSPFv3 supports the larger 128-bit IPv6 addresses and prefixes.
- Multiple OSPFv3 instances can run over the same link.
- OSPFv3 adds support for address families which helps OSPFv3 IPv6 networks to support both IPv6 and IPv4 nodes.
- Router LSAs and network LSAs no longer carry prefix information. In OSPFv3, these LSAs carry only the topology information.
- New and modified LSAs handle the flow of IPv6 addresses and prefixes in an OSPFv3 network

We do not cover the details and configuration of OSPFv3 in this course. For an in depth look at OSPFv3 on the 5600 vRouter, 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

What is a Link-State Protocol?

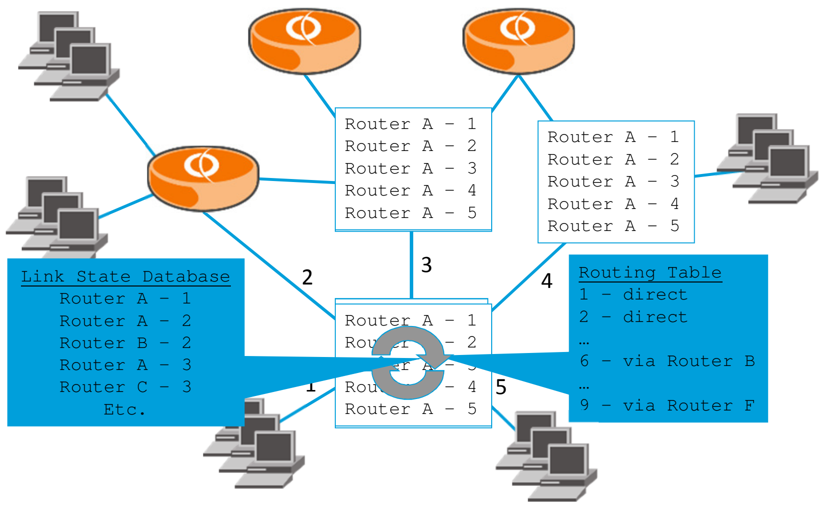

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution

Let’s look how a link state protocol works.
Instead of forwarding lists of all networks and subnets in the environment, an OSPF router only advertises directly-connected links. These announcements, called Link State Advertisements, or LSAs, are sent to an OSPF multicast IP address. All OSPF routers store the information in their local Link State Database.
Let’s look at an example.
Router A is directly connected to links 1, 2, 3, 4, and 5.
Therefore, router A’s LSA will contain information about itself and those 5 links.
The LSA will then be propagated to all other OSPF routers, which will store the information.
As all OSPF routers send their LSAs, each router builds a database. The database contains a complete list of routers and links in the network. All routers in the OSPF network must have the same link state information in their databases. The OSPF protocol includes acknowledgements and periodic checks to ensure the databases stay synchronized.
Each router then runs their local database through the Shortest Path First algorithm to generate a local 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

When are LSAs Sent?

When do OSPF routers send Link State Advertisements (LSA)

- When there is a change in the network
  - Affected devices send link change information immediately
  - Devices run SPF to calculate best routes
- When OSPF is configured or modified on a router it goes through a four step process
  - Locate neighbors
  - Establish adjacencies
  - Exchange LSDB
  - Run SPF to calculate best 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

Link state advertisements are sent when there is a change in the network, such as a link going down or coming back up.
Devices directly connected to the affected link send out a new LSA immediately, which is propagated via multicast to all other OSPF routers.
All devices then run the SPF algorithm to calculate a new routing table.
The other time LSAs are sent is when OSPF is configured or modified on an individual device. That device must first join the OSPF network. This is a four step process.

1. A device locates its OSPF neighbors
2. The device establishes adjacencies with its neighbors
3. The devices exchange link state database information
4. The devices run SPF to calculate the best routes.

We will discuss this adjacency process in detail on the next few slides. Be aware that the majority of OSPF implementation problems occur due to issues with the adjacency process, so understanding how this works will help you in troubleshooting your OSPF 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

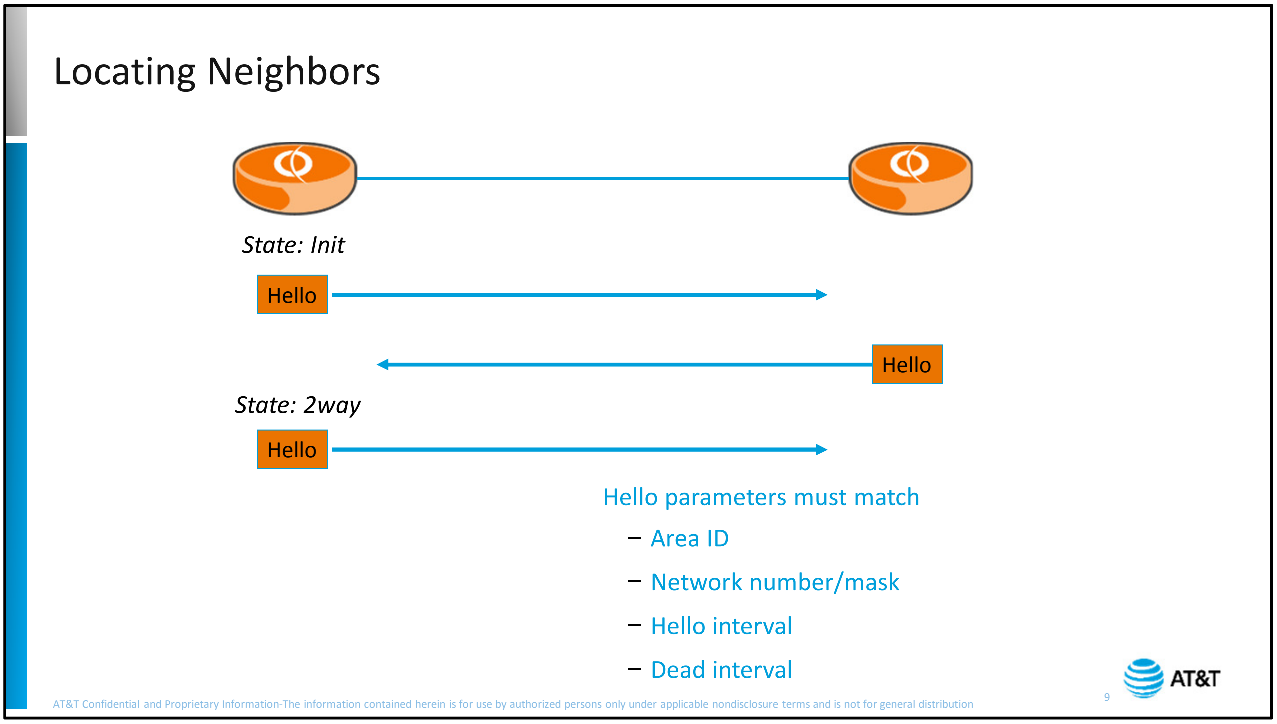

When a router comes on-line, it starts OSPF in the *Init* state.
It then begins sending hello packets, every 10 seconds by default, out every active OSPF interface.
Other OSPF routers are also sending hello packets. When a router receives a hello, it compares the values in the hello packet with its local settings.
Four parameters must match: area ID, network number and mask, hello interval, and dead interval.
When they match, the device state changes from *Init* to *Two Way (2way)*.
The OSPF router continues to send and receive hellos every 10 seconds. At this point the hellos are used to monitor the state of the link. As long as a device continues to send and receive hellos on the link, the link is active and is advertised in OSPF.

Forming Adjacencies

Only adjacent routers can synchronize LSDBs

- "Neighbor" does not mean "adjacent"
- Forming adjacencies is different for different types of networks
  - **Point-to-point:** if neighbor, automatically adjacent
  - **Broadcast (LAN):** adjacent with Designated Router (DR) and Backup Designated Router (BDR) only
  - **Point-to-multipoint:** if branch, adjacent with hub only; if hub, adjacent with all branches

AT&T Confidential 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

Just because two devices are OSPF neighbors, does not mean that they will exchange link state advertisements.
The devices must first form an adjacency. Only adjacent routers exchange link state information.
The process of forming an adjacency differs depending on the type of network interface.
A point-to-point link, such as a DSL or serial interface, only has one other device on the link. These two devices automatically form an adjacency.
On broadcast interfaces such as Ethernet, a new device will only form an adjacency with the designated router or DR and backup designated router or BDR on the link. We will look at the DR election process on the next slide. Limiting adjacencies to only two devices is designed to prevent OSPF link state advertisements from flooding a broadcast network.
A point-to-multipoint network is a network where multiple devices are on the same subnet but do not share a common physical link. This can occur in frame relay networks, or virtual private networks in hub-and-spoke or partial mesh configurations. In this case, one device on the network serves as a relay point for link state advertisements, resending it to all other devices on the point to multipoint network as necessary.

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

Designated Router Election

DR and BDR selected during hello exchange based on priority

- New device with higher priority does not take over DR role until next election
- Priority 0 never becomes DR
- If all priorities equal, highest router ID becomes DR

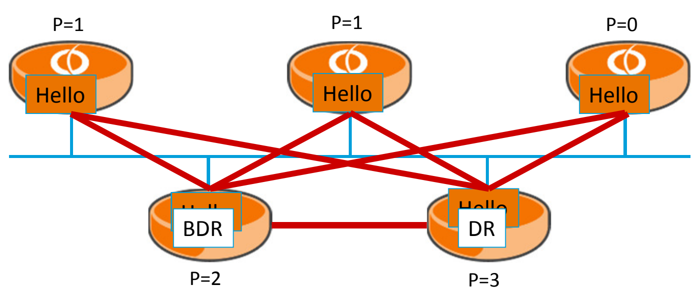

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution

AT&T

Let's look at the designated router election process.
The OSPF protocol automatically selects the designated router and backup designated router based on device priority, which is included in the hello messages. The higher the priority value, the more likely a router will become the DR.
Once a LAN has a DR and a BDR, new devices do not assume the DR role, even if they have higher priority. In other words, the DR selection process is not pre-emptive.
If you want to ensure a device never becomes the designated router, set the priority to zero.
If all devices have equal priority, then OSPF uses device router IDs to select the designated router and backup designated router.
Let's look at an example of the DR selection process.
Assuming all the devices in the example come on line at approximately the same time, they will all begin sending hello packets on the link. As the devices learn about each other, they use the advertised priority to designate the device with the highest priority as the DR, and the next-highest as the BDR.
Now the DR and the BDR form an adjacency with each other. All other routers will form adjacencies to the DR and the BDR only. This cuts down on the total amount of OSPF traffic sent on the broadcast link, since devices only send LSAs to their adjacent 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

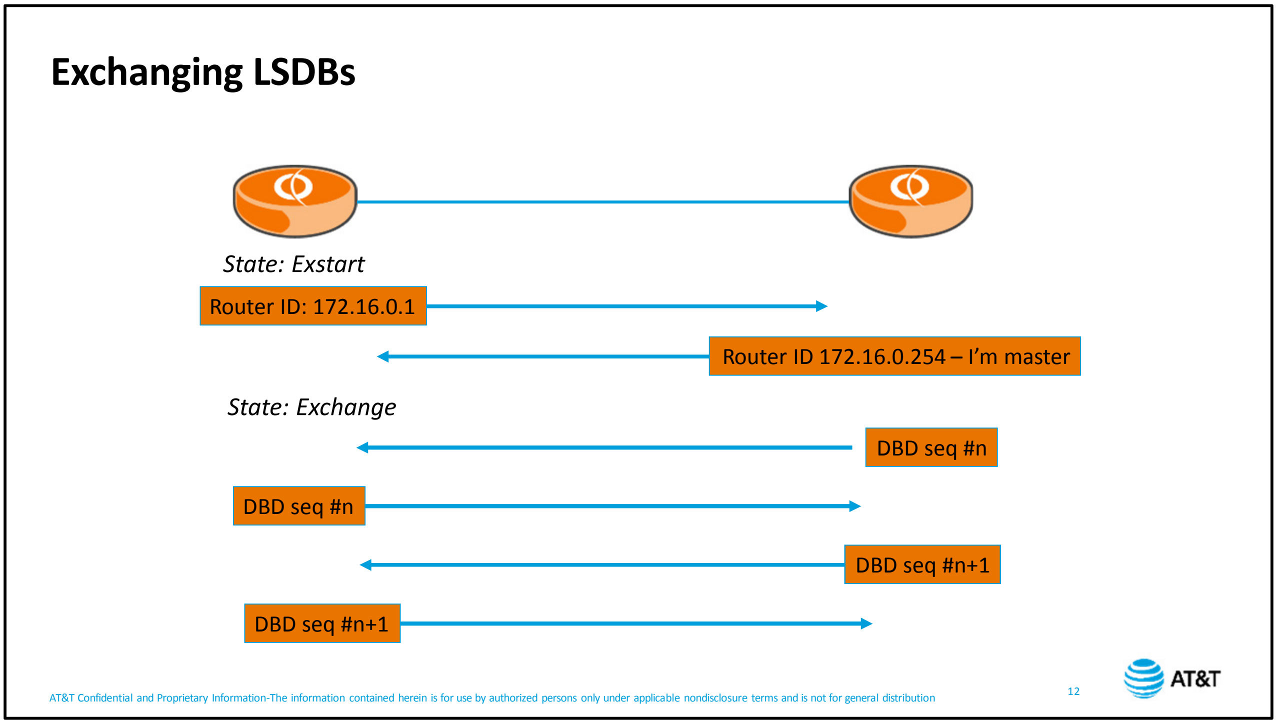

Once a device forms its adjacencies, the state changes to *Exstart* and prepares for the database exchange.

The database exchange is a master-slave exchange. The devices compare router IDs to determine which device is master. The device with the highest router ID becomes the master. The master device sets the initial sequence number for the exchange.

The device state now changes to *Exchange*.

The master initiates the exchange by sending the first Database Descriptor (DBD) packet. This packet is a summary of the database rather than the actual link information – we will discuss this process later in the module.

The slave device replies with its own DBDs using the same sequence number. This also serves as an acknowledgement of receipt of the master's DBD.

The master increments the sequence number and sends the next DBD. Incrementing the sequence number serves as an ACK to receipt of the slave's DBD.

The exchange continues until both devices have sent the entire database summary to the adjacent neighbor.

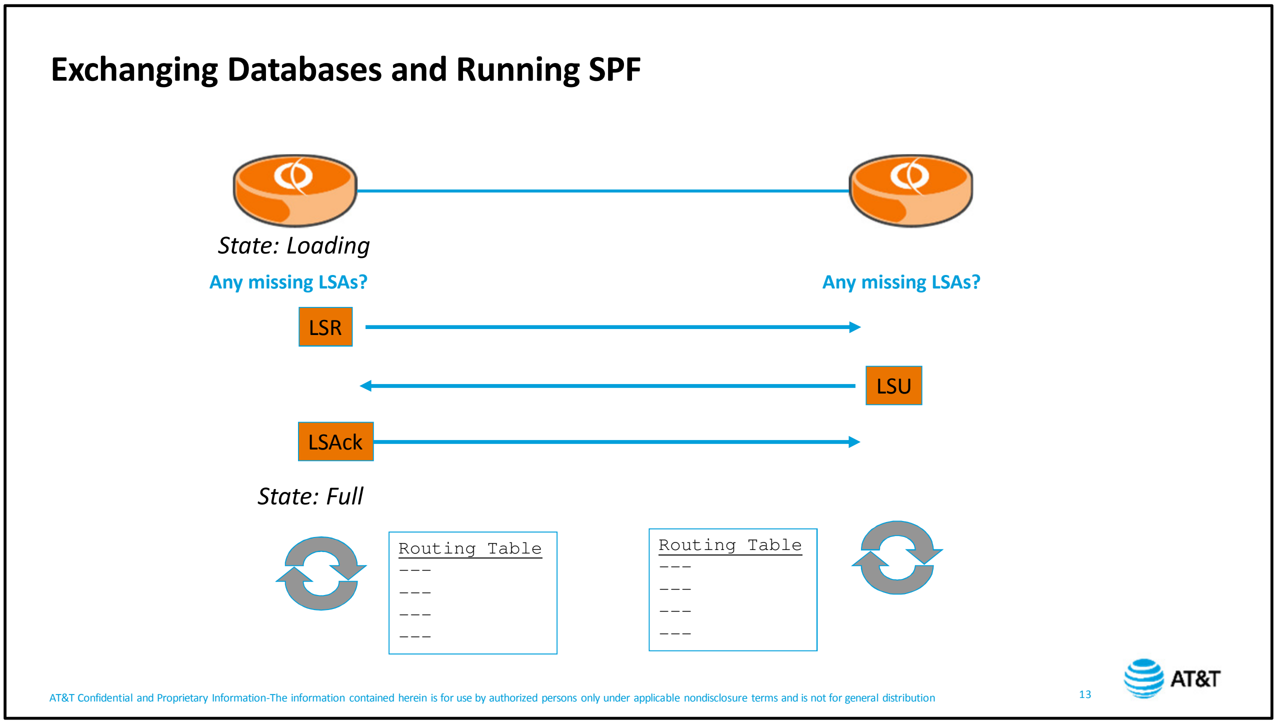

After receiving the complete database summary, the device state changes to *Loading*. Both devices compare the received summary with its own local database to determine if it is missing any LSAs. Remember, in order for OSPF to work properly, all devices must have the same information in the database.

If a device is missing a link state advertisement, it sends a link state request (LSR) to the adjacent neighbor.

The neighbor responds with the actual link information in a link state update (LSU).

The data exchange is connection-oriented, so the requester sends an acknowledgement (LSAck) when it receives the requested update. This process continues until both devices have synchronized their databases.

Each device now changes state to *Full*, indicating a complete adjacency.

Each device now runs the SPF algorithm to generate its own routing table.

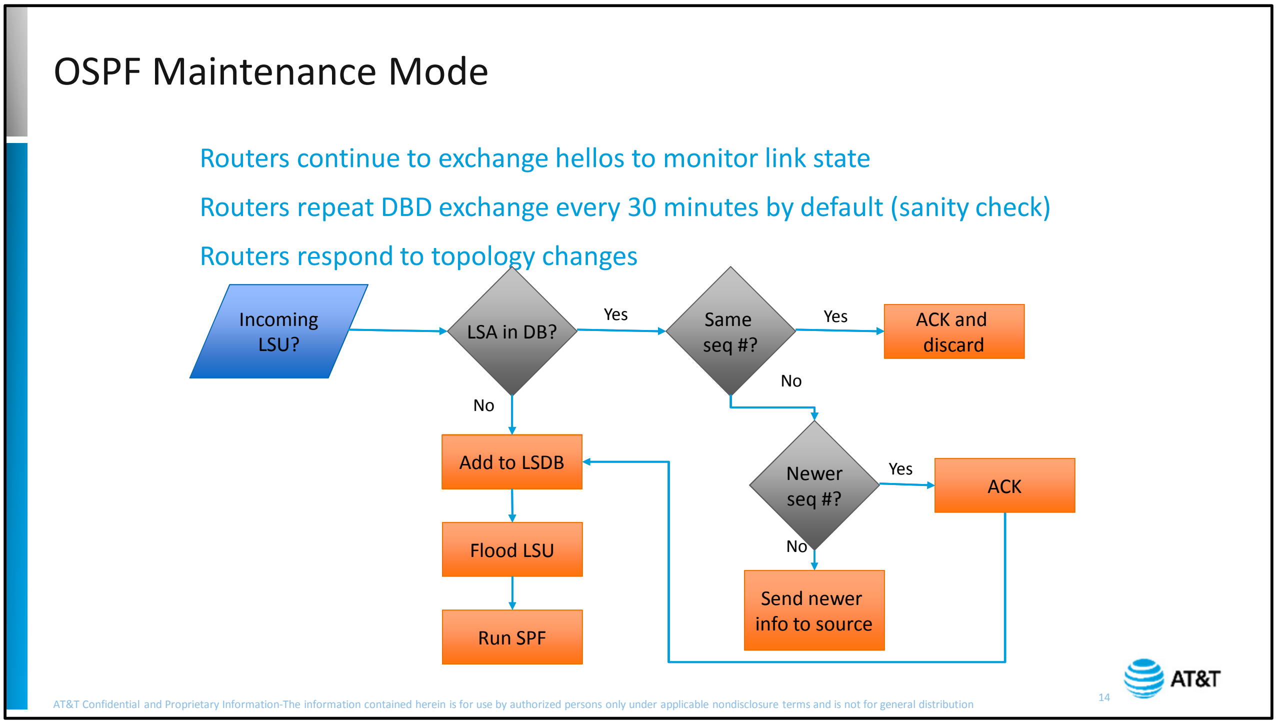

Now that each device has a routing table, OSPF goes into maintenance mode.
Routers continue to exchange hellos with their neighbors to monitor link state.
Routers also repeat the DBD exchange with their adjacent neighbors every 30 minutes by default. This is a “sanity check” built into OSPF to ensure that devices do not somehow miss link state information.
If the database is out of sync, or if a device receives a new link state advertisement (LSA), it goes through the following process.
The incoming link state update (LSU) is compared with the link state advertisement (LSA) already in the database.
If the data is not already present in the database, the device adds it to the database, floods the link state update out all interfaces, then runs the SPF algorithm to generate a new routing table.
If the link state advertisement is already in the database, then the device compares the sequence number of the new data with the one already in the database. If the sequence numbers match, then the device sends a simple acknowledgement and discards the link state update.
If the sequence number is different, then the device determines whether the link state update is newer than the data already in the database. If it is new, then the device acknowledges the update, then adds the new information, floods the link state update, and runs SPF.

If the sequence number received is older than the one in the database, the device responds with the more current link state update.

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

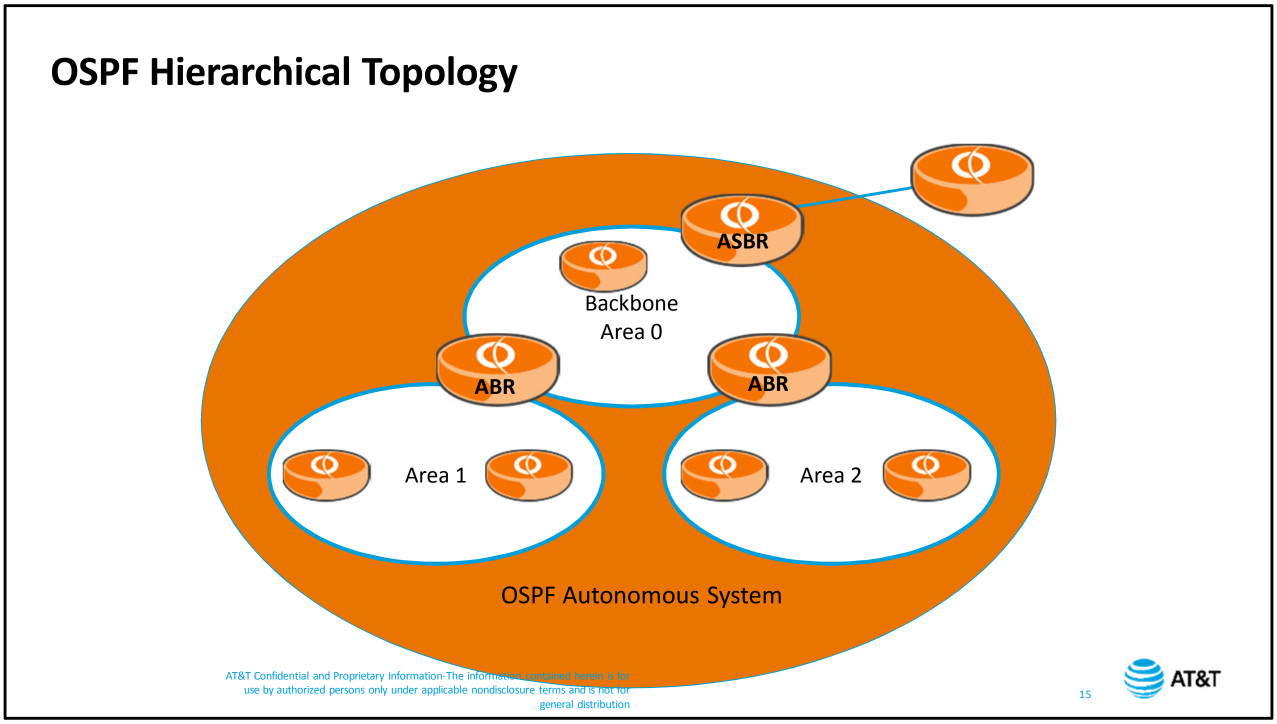

As we stated earlier, OSPF is designed as a hierarchical routing protocol. The complete network of routers comprises an OSPF autonomous system.
You then organize your network into areas.
A single backbone area provides connectivity between leaf areas. The backbone area is always area 0.
Area border routers, or ABRs, are routers connected to the backbone area and to one or more leaf areas. OSPF is designed so that routing updates cannot pass from one leaf area to another without going through a router connected to area 0.
A router entirely within an area is called an area router, or just a router.
Routers in area 0 are sometimes called backbone routers.
A router advertising routes from outside the OSPF autonomous system is called an Autonomous System Boundary Router, or ASBR. These outside routes could come from other routing protocols, such as BGP, or from local static routes that have been redistributed into OSPF. For more information on redistribution, please refer to the *AT&T Vyatta 5600 vRouter Software Documentation* on http://businesscenter.att.com.

Single-Area OSPF 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

Now we'll look at configuring OSPF for single-area 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

Single-Area OSPF

- Hierarchical design not required
- Suitable for smaller networks
  - Under 100 networks/subnets
- All routers must be in area 0
- No summarization available
- Can use ASBRs to connect to non-OSPF networks

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution 17

AT&T

We talked about OSPF being a hierarchical routing protocol with route summarization occurring at area borders. However, you do not have to deploy OSPF in a hierarchical fashion.
Deploying OSPF in a single-area design is appropriate for smaller networks containing 100 or fewer subnets in the OSPF autonomous system.
If you do use single-area OSPF, you need to use area number 0 on all your routers.
Of course, you will not be able to summarize routes within your OSPF network, but you can still use autonomous system boundary routers to connect to non-OSPF networks. This includes your Internet 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

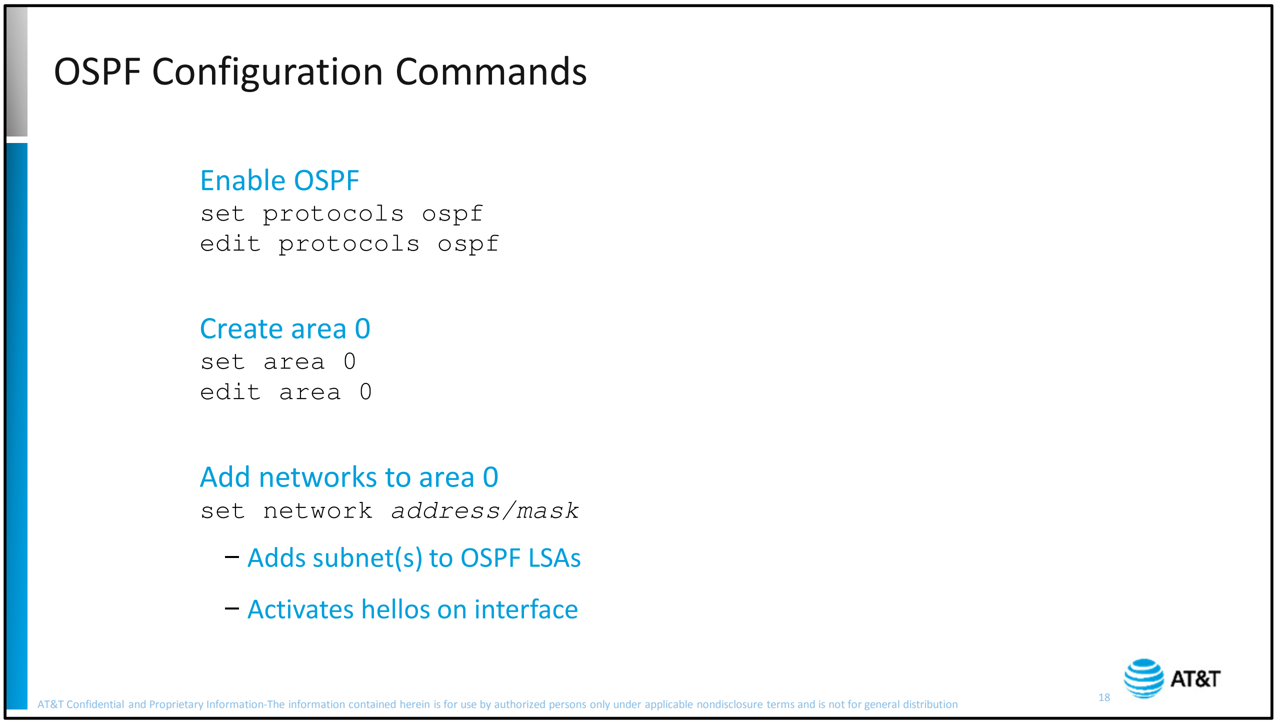

To enable OSPF on a vRouter, enter the command `set protocols ospf` in Configuration mode.

Because you will be setting multiple parameters within OSPF, we recommend the use of the `edit` command. This not only enables OSPF but places you within the OSPF level of the configuration hierarchy. This saves you on typing `set protocols ospf` at the beginning of every command.

The next step is to create area 0. Again, we recommend using the `edit` command so that all subsequent commands are associated with the same area number.

Next, add the networks you want active in OSPF to the area.

The `network` statement adds the subnets you specify to OSPF LSAs, and activates hellos on the interfaces connected to those subnets.

OSPF Scenario

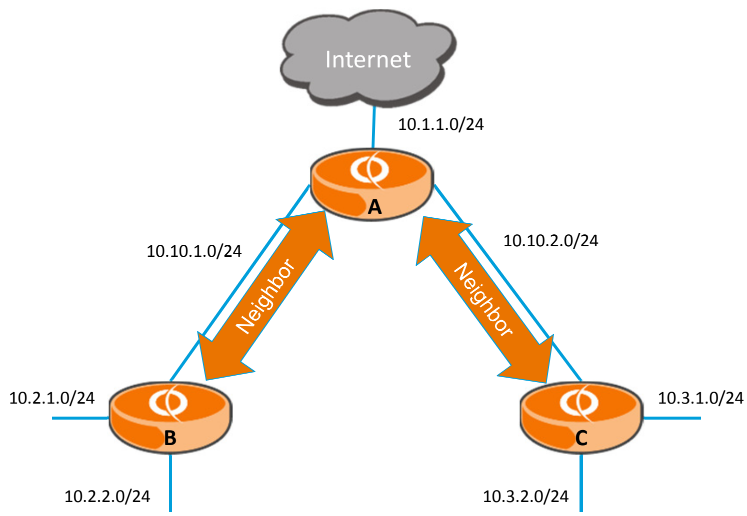

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution

19

Now let's apply these commands to a network of routers.
Router A is connected to the Internet and to 2 other routers.
Routers B and C are connected to subnets with end users and servers.
When we finish configuring all the devices, we expect to see neighbor relationships form on the links between the routers.

AT&T Confidential 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 Configuration

```
[edit]
vyatta@VY-A# edit protocols ospf area 0
[edit protocols ospf area 0]
vyatta@VY-A# set network 10.1.1.0/24
[edit protocols ospf area 0]
vyatta@VY-A# set network 10.10.1.0/24
[edit protocols ospf area 0]
vyatta@VY-A# set network 10.10.2.0/24
[edit protocols ospf area 0]
vyatta@VY-A# commit
[edit protocols ospf area 0]
vyatta@VY-A# save
[edit]
vyatta@VY-A#
```

AT&T Confidential 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 configure router A first.
We begin with the `edit` command. We combine the `protocols ospf` and `area 0` commands into a single `edit` command. Note the change in the prompt.
We can now activate OSPF for each of the subnets directly connected to our device. First the link to the Internet, then the link to router B, and then the link to router C.
We commit our changes to make them active, then save them to make them permanent.

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

Configurations Side-by-Side

```
[edit]
vyatta@VY-A# show protocols ospf
    area 0 {
        network 10.1.1.0/24
        network 10.10.1.0/24
        network 10.10.2.0/24
    }
[edit]
vyatta@VY-A#
```

```
[edit]
vyatta@VY-B# show protocols ospf
    area 0 {
        network 10.10.1.0/24
        network 10.2.1.0/24
        network 10.2.2.0/24
    }
[edit]
vyatta@VY-B#
```

```
[edit]
vyatta@VY-C# show protocols ospf
    area 0 {
        network 10.10.2.0/24
        network 10.3.1.0/24
        network 10.3.2.0/24
    }
[edit]
vyatta@VY-C#
```

AT&T Confidential 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 complete our configuration of all three routers, we can see that each device has placed all of its networks into the same area, area 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

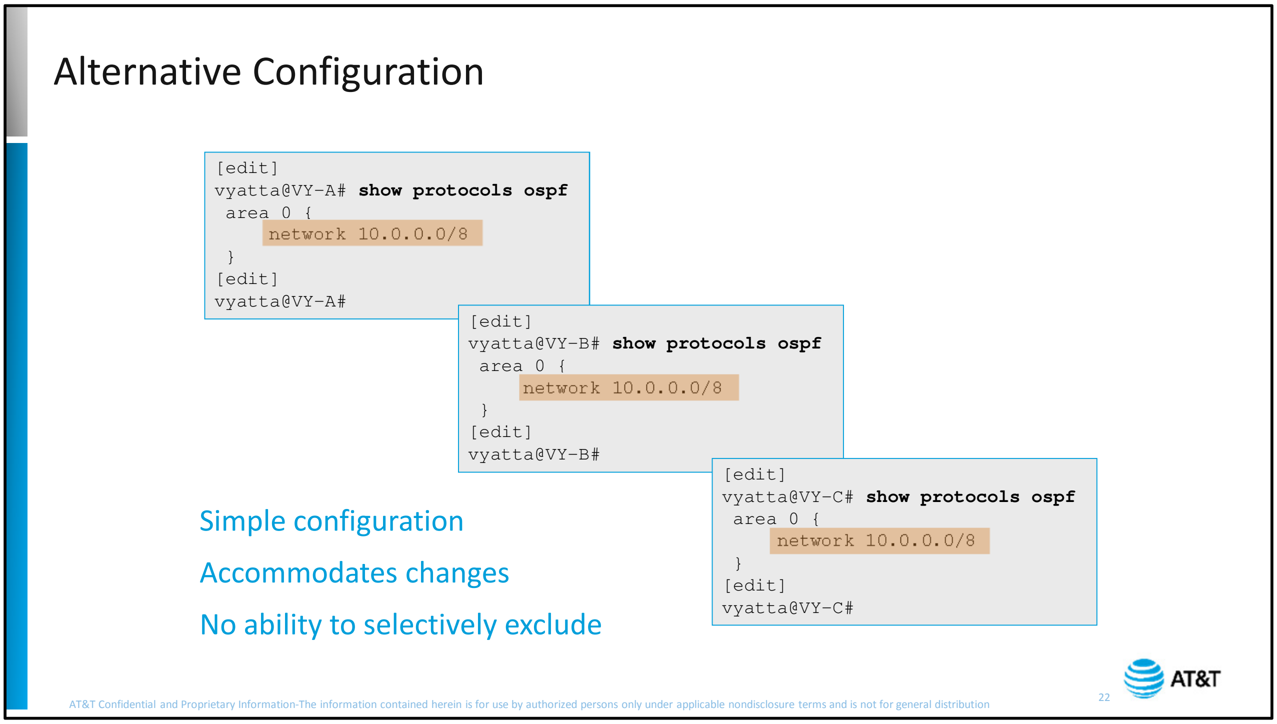

You can also use summary network and mask combinations in your network statements. In this case, all of the directly-connected subnets can be summarized into network 10.0.0.0 with an 8-bit mask. If we configure OSPF using the summary, we achieve the same effect of activating OSPF on all interfaces.

Using a summary range may simplify your OSPF configuration.

If you add or remove networks in the future, you do not need to worry about modifying your OSPF configuration; the summary will ensure that OSPF stays active on all interfaces within the range.

However, if you use a summary network statement, you cannot easily exclude individual subnets or addresses from routing updates. If you only want certain networks advertised within OSPF, you should use specific network statements rather than summaries.

AT&T Confidential 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 OSPF Routing

Use `show ip route` in Operational mode to display OSPF routes

```
vyatta@VY-A:~$ 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

O       10.1.1.0/24 [110/1] is directly connected, dp0p1p1, 00:02:57
C    *> 10.1.1.0/24 is directly connected, dp0p1p1
O    *> 10.2.1.0/24 [110/2] via 10.10.1.2, dp0p1p2, 00:01:30
O    *> 10.2.2.0/24 [110/2] via 10.10.1.2, dp0p1p2, 00:01:30
O    *> 10.3.1.0/24 [110/2] via 10.10.2.2, dp0p1p3, 00:00:41
O    *> 10.3.2.0/24 [110/2] via 10.10.2.2, dp0p1p3, 00:00:41
O       10.10.1.0/24 [110/1] is directly connected, dp0p1p2, 00:02:57
<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

23

To verify that OSPF is working, use the command `show ip route` and look for active OSPF routes.

In this case, we are looking at router A.

We see that we have active routes learned via OSPF in our routing table. We know the routes are active because of the asterisk (*), and we know they are from OSPF because of the *O* in the first column.

OSPF has calculated that the best path to this particular network is via the peer at 10.10.1.2, reachable via data plane interface 2.

This route was added to the table by the SPF algorithm 1 minute and 30 seconds ago.

Note that we also have inactive OSPF routes. These are for the directly connected networks. They are included in the routing table as OSPF routes because we have LSAs for these networks in our database. However, because they are directly connected, the OSPF path to the network is not the “best path” according to precedence. For more information on precedence, please refer to the *AT&T Vyatta 5600 vRouter Software Documentation* on http://businesscenter.att.com.

Verifying OSPF Neighbors

Use `show ip ospf neighbor` in Operational mode to display neighbor information

```
vyatta@VY-A:~$ show ip ospf neighbor

OSPF Process 0:
Neighbor ID     Pri State           Dead Time Address         Interface
172.24.42.52      1 Full/DR           31.721s 10.10.1.2       dp0p1p1
172.24.42.53      1 Full/DR           35.533s 10.10.2.2       dp0p1p2
vyatta@VY-A:~$
```

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution
24

AT&T

We can view the status of our OSPF neighbors with the command `show ip ospf neighbor`.
The output shows the Neighbor ID which is the router ID of the neighbor, the State of the connection (in this case, *Full*), and it also shows that the neighbor device is the DR. Note that the *State* reported in the output is the neighbor's state and not our own.
The Dead Time counts from 40 seconds down to 0, and is reset by hello messages.
The Address is the physical address of the neighbor, and the Interface is the interface used to reach 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

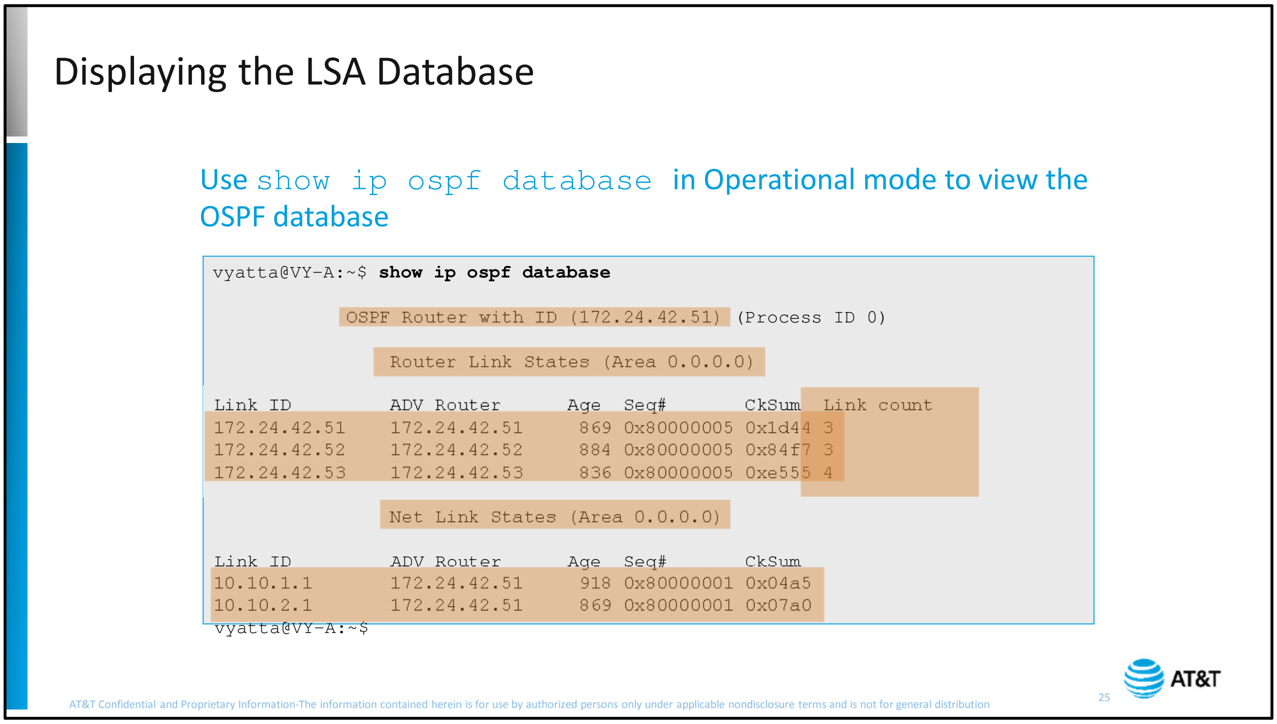

You can display the contents of the OSPF LSA database with the command `show ip ospf database`. In single-area OSPF configurations, you will see two or three categories of link states listed.

Router link states are a list of all the OSPF routers in the area.

Network link states are a list of networks that are connected to more than one OSPF router. We will discuss network link states later in this module.

Not shown is a list of external links – network advertisements that originated outside of the OSPF network. We will see this when we add a default route to our OSPF network.

Things to look for in this output: How many links each router in the network is advertising. You can then look at the network links to see which router is advertising which particular network.

For example, in our database, we see that this local router – identified at the top of the output – is advertising 3 links. Two of the links display in the Net Link States list. But where is the third?

OSPF Database Details

Display OSPF database details

`show ip ospf database router` *`address`*

```
Link connected to: Stub Network
 (Link ID) Network/subnet number: 10.1.1.0
 (Link Data) Network Mask: 255.255.255.0
  Number of TOS metrics: 0
   TOS 0 Metric: 1

Link connected to: a Transit Network
 (Link ID) Designated Router address: 10.10.1.1
 (Link Data) Router Interface address: 10.10.1.1
  Number of TOS metrics: 0
   TOS 0 Metric: 1

Link connected to: a Transit Network
 (Link ID) Designated Router address: 10.10.2.1
 (Link Data) Router Interface address: 10.10.2.1
  Number of TOS metrics: 0
   TOS 0 Metric: 1

vyatta@VY-A:~$
```

AT&T Confidential 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 see everything advertised by a router by adding the router address to the command. The first part of the output expands on the database summary, including the number of links.

The second part of the output shows all the links being advertised.

Transit links are connected to other OSPF routers. These links also show up in the database summary.

Stub links are networks that do not have any other OSPF routers advertising them. These links do not show up in the database summary, but are still included in the SPF calculation and show up in the resulting 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

Originating a Default Route

Used at edge of OSPF autonomous system

Eliminates need to redistribute a static default route

`set protocols ospf default-information originate [always]`

- `always` parameter sends default LSAs even if device has no local default route

AT&T Confidential 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

An addition we need to make to our scenario is the distribution of a default route.
In our scenario, router A is at the edge of the network, connected to the Internet. We can assume that router A has some kind of route to the Internet, either via BGP or a local default route.
By configuring this edge router to originate a default route into the OSPF network, we will add an LSA for network 0.0.0.0/0, showing the edge router is directly connected to the network exit point. This eliminates the need to redistribute a static default route, or manually configure a static default route on each device.
The command is `set protocols ospf default-information originate`. Using the keyword `always`, instructs the device to generate the default route even if no default route is currently present in the local routing table. This might occur with a BGP edge router.

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

Default Route in Our Scenario

```
vyatta@VY-A# show protocols
 ospf {
     area 0 {
         network 10.0.0.0/8
     }
     default-information {
         originate {
         }
     }
 }
 static {
     route 0.0.0.0/0 {
         next-hop 10.1.1.2 {
         }
     }
 }
```

```
vyatta@VY-C:~$ show ip route
<Headers omitted>

Gateway of last resort is 10.10.2.1 to network 0.0.0.0

O>* 0.0.0.0/0 [110/10] via 10.10.2.1, 00:00:01
O>* 10.1.1.0/24 [110/20] via 10.10.2.1, dp0p3p1, 00:10:52
<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

28

AT&T

In our scenario, we have added the `default-information originate` command to our edge router.

We have also added a static default route. As long as this route is in the routing table – that is, as long as the interface connected to subnet 10.1.1.0 is active – the device will advertise a default route in OSPF.

When we look at the routing table on router B, we now see the default route learned via OSPF.

Default Route in OSPF Database

```
vyatta@VY-A:~$ show ip ospf database
            OSPF Router with ID (192.168.200.1) (Process ID 0)

                Router Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum  Link count
192.168.200.1   192.168.200.1     11 0x80000247 0x9882 3
192.168.200.2   192.168.200.2     44 0x80000247 0xf93e 3
192.168.200.3   192.168.200.3     44 0x80000247 0x5b9b 4

                Net Link States (Area 0.0.0.0)

Link ID         ADV Router      Age  Seq#       CkSum
10.10.1.1       192.168.200.1    163 0x80000242 0x7bea
10.10.2.1       192.168.200.1    223 0x80000242 0x7ee5

                AS External Link States

Link ID         ADV Router      Age  Seq#       CkSum  Route             Tag
0.0.0.0         192.168.200.1     10 0x80000002 0xa3f1 E2 0.0.0.0/0       254
vyatta@VY-A:~$
```

AT&T Confidential 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

When we add the default route to OSPF, it is distributed as an external link. This is because the actual path does not originate within the OSPF 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

Disable Hellos on Stub Interfaces

- Configure an interface as passive to disable hellos
  `set protocols ospf passive-interface` *`dpxpypz`*
  - Subnet still advertised in OSPF
  - Eliminates unnecessary traffic

```
[edit]
vyatta@VY-B# show protocols ospf
   area 0 {
       network 10.10.1.0/24
       network 10.2.1.0/24
       network 10.2.2.0/24
   }
   passive-interface dp0p2p1
   passive-interface dp0p2p2
```

```
[edit]
vyatta@VY-C# show protocols ospf
   area 0 {
       network 10.10.2.0/24
       network 10.3.1.0/24
       network 10.3.2.0/24
   }
   passive-interface dp0p3p1
   passive-interface dp0p3p2
```

AT&T Confidential 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 last thing we should do in our scenario is disable OSPF hellos on the interfaces where there are no other OSPF routers.

We still want the subnets advertised in OSPF, so we do not want to delete the network statement, but we do not want those hellos sent every 10 seconds when there is no one there to listen to them.

To disable hellos, we need to make the interface passive for OSPF. The command is `set protocols ospf passive-interface`.

In our scenario, routers B and C have stub networks, so we make those interfaces passive for OSPF.

AT&T Confidential 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 OSPF Interface Configuration

```
vyatta@VY-B:~$ show ip ospf interface
dp0p2p1 is up
  ifindex 2, MTU 1500 bytes, BW 0 Kbit <UP,BROADCAST,RUNNING,MULTICAST>
  Internet Address 10.10.1.2/24, Broadcast 10.10.1.255, Area 0.0.0.0
  MTU mismatch detection:enabled
  Router ID 10.100.100.2, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State Backup, Priority 1
  Designated Router (ID) 10.100.100.1, Interface Address 10.10.1.1
  Backup Designated Router (ID) 10.100.100.2, Interface Address 10.10.1.2
  Multicast group memberships: OSPFAllRouters OSPFDesignatedRouters
  Timer intervals configured, Hello 10s, Dead 40s, Wait 40s, Retransmit 5
    Hello due in 6.251s
  Neighbor Count is 1, Adjacent neighbor count is 1
dp0p2p2 is up
  ifindex 3, MTU 1500 bytes, BW 0 Kbit <UP,BROADCAST,RUNNING,MULTICAST>
  Internet Address 10.2.1.1/24, Broadcast 10.2.1.255, Area 0.0.0.0
  MTU mismatch detection:enabled
  Router ID 10.100.100.2, Network Type BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 10.100.100.2, Interface Address 10.2.1.1
  No backup designated router on this network
  Multicast group memberships: <None>
  Timer intervals configured, Hello 10s, Dead 40s, Wait 40s, Retransmit 5
    No Hellos (Passive interface)
  Neighbor Count is 0, Adjacent neighbor count is 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

AT&T

To view the interface-level operations and configuration, use the command `show ip ospf interface`. This command can be very useful in troubleshooting adjacency problems.

The output displays the address and mask associated with the interface.

The area ID, expressed in dotted decimal format.

The designated router status of the device if this is a broadcast link – in this case, this device is the backup DR.

The priority setting for designated router selection,

The hello and dead interval timers, which must match if a device is to form an adjacency,

And the timer for when the next hello packet is due to be transmitted.

If an interface is passive, then the output will indicate that no hellos are being sent.

AT&T Confidential 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

AT&T Confidential 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 conclude with a brief discussion on troubleshooting OSPF.

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

Set the Router ID

Dedicate a range of addresses for loopback interfaces

Configure the loopback interface

Set router ID to address of loopback interface

```
set protocols ospf parameters router-id address
```

Advertise loopback addresses in OSPF

- OSPF command output displays router ID of neighbors, advertising routers, etc.
- If addresses are in routing table, devices are directly reachable

Note: changing router ID may require a process reset using the `reset ip ospf process` 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

Something you can do to facilitate troubleshooting in an OSPF network is set the router ID on all your OSPF devices.

First, dedicate a range of addresses for the loopback interfaces on your OSPF routers. Each device will have its own address with a 32-bit mask.

Next, configure each device's loopback interface.

If OSPF is already running, you'll need to set the router ID, then restart the OSPF process. If you have not configured OSPF yet, the device will automatically use the loopback address for the router ID.

Next, add the loopback address to the OSPF network statements.

Now, when you look at OSPF operational command output or at log messages, the router ID can easily be associated with a specific device.

Because the addresses are explicitly in the routing table, you will be able to access the device remotely as long as at least one interface is active and advertising OSPF.

Note that, depending on your version of vRouter software, you may have to reset the OSPF process in order for the router ID change to become active. If this is the case, you will see a message when you commit the 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

Adding Router IDs

RouterID range:
10.100.100.0-50

A: 10.100.100.1/32
B: 10.100.100.2/32
C: 10.100.100.3/32

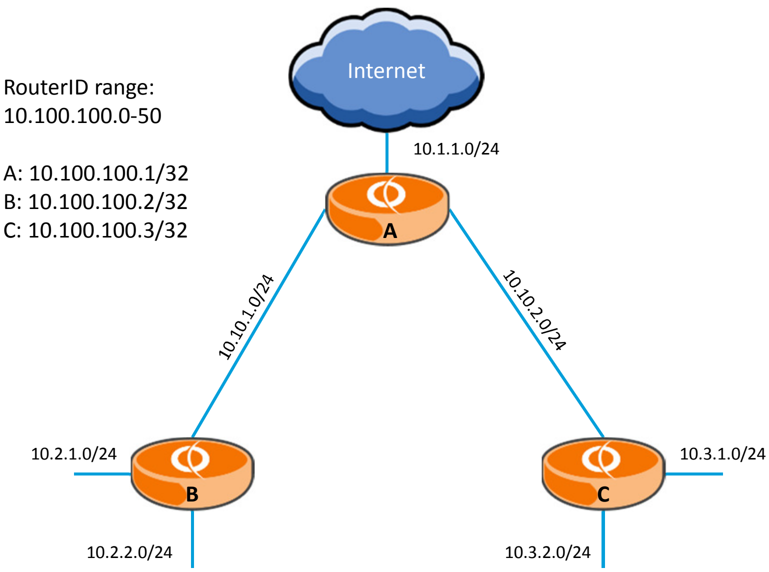

AT&T Confidential and Proprietary Information-The information contained herein is for use by authorized persons only under applicable nondisclosure terms and is not for general distribution 34

We will add router IDs to our scenario.
According to the table displayed on the left side of the example.

AT&T Confidential 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 Router ID Configuration

```
vyatta@VY-A:~$ show ip route
<Truncated Output>
O   10.100.100.1/32 [110/10] is directly connected, lo, 00:06:25
C>* 10.100.100.1/32 is directly connected, lo
O>* 10.100.100.2/32 [110/20] via 10.10.1.2, dp0p1p1, 00:01:30
O>* 10.100.100.3/32 [110/20] via 10.10.2.2, dp0p1p2, 00:00:21
C>* 127.0.0.0/8 is directly connected, lo
C>* 172.24.42.0/24 is directly connected, dp0p1p3
vyatta@VY-A:~$
```

```
vyatta@VY-A:~$ show ip ospf neighbor

Neighbor ID     Pri State           Dead Time Address         Interface
10.100.100.2      1 Full/Backup       39.709s 10.10.1.2       dp0p1p1
10.100.100.3      1 Full/Backup       34.051s 10.10.2.2       dp0p1p2
vyatta@VY-A:~$
```

AT&T Confidential 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 routing table, we can see entries for the individual addresses in the loopback range.
These entries match the router IDs we see in the `show ip ospf neighbor` command 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

Troubleshooting Tips

The most common OSPF configuration issues include

- Networks not included in OSPF configuration
  - Check OSPF database
  - Check network statement in configuration for local router(s)
- Neighbors cannot form adjacency

```
vyatta@VY-A:~$ show ip ospf neighbor

Neighbor ID     Pri State           Dead Time Address         Interface
10.100.100.2      0 2-Way/DROther     38.138s 10.10.1.2       dp0p1p1
```

  - Area ID
  - Network/mask
  - Hello interval
  - Dead interval
  - Check DR priority

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

36 AT&T

When troubleshooting an OSPF configuration, keep in mind that single-area OSPF issues generally come down to one of two problems.

If an expected route is missing from the routing table and you have verified the underlying physical networks, then chances are the network is not included in the OSPF configuration. You can check the OSPF database for the link. Remember that stub networks will not display in the database summary. You may have to look at individual routers.

You can also go to the routers directly connected to the missing link and check the configuration. Make sure the missing network is included in the OSPF network statement.

The second problem is that you are not receiving LSAs from a neighbor router when you should be. This generally means your devices have not been able to form an adjacency. The neighbor output shows that you are still in a 2-way state with your neighbor rather than fully adjacent.

First, check the configuration for the four required matching parameters: area ID, network and mask configured on the interface, hello interval, and dead interval. These must be the same on all devices sharing a link. If these check out, then check the designated router priority on both devices. If both devices have priority set to 0, neither can become the DR. If there is no DR, there is no adjacency, which means there is no exchange of routing information.

Summary

You should now be able to

- Explain basic OSPF operations
- Configure the vRouter for single-area OSPF networks
- Verify and troubleshoot OSPF 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

AT&T

This concludes the AT&T Vyatta 5600 vRouter OSPF Basics module.

You should now be able to:

- Explain basic OSPF operations
- Configure vRouters for a single-area OSPF network
- Verify and troubleshoot OSPF operations

We hope that this information has been useful, and that you will take additional courses in the future.

Thank you.

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