Single Area OSPF Routing in AOS

Overview

This configuration guide explains the concepts behind configuring an ADTRAN Operating System (AOS) product for Open Shortest Path First (OSPF) routing protocol operation in a single area. For detailed information regarding specific command syntax, refer to the AOS Command Reference Guide on your ADTRAN OS Documentation CD.

Information in this document is relevant to the following sections of the AOS:

- Router (OSPF) Configuration Command Set
- Configuring Your Router
- Verifying Your Configuration Using Show Commands
- Managing Event Messages

This configuration guide contains the following information:

- Summary of OSPF
- Step-by-step instructions to configure OSPF in a single area
- Basic verification and troubleshooting tools for OSPF
- Full sample configuration for OSPF in a single area

Introduction

OSPF is a classless link-state routing protocol designed specifically as an Interior Gateway Protocol (IGP) for Internet Protocol (IP) networks. Key benefits of OSPF include the following:

- Fast convergence in the event of a network failure
- Scalability to include hundreds of networks
- Full support for Variable Length Subnet Masks (VLSM)
- Route summarization at Area Border Routers (ABR)
- Open standard supported by many network equipment vendors
- Rich metric for optimal path selection
- Open standard with multiple vendor support

OSPF Overview

The OSPF IP is defined in the Request for Standards (RFC) 2328 dated April, 1998. Routing protocols are designed to facilitate path-sharing to destination networks among peer gateways or routers. OSPF operates within one autonomous system (AS). A single AS includes all routers normally under the administrative control of a single networking team, as compared with the Internet, which consists of many AS's.

Historically, OSPF followed a class of routing protocols known as distance vector or Bellman-Ford protocols. Perhaps the best known of these is the Routing Information Protocol (RIP). RIP shared routes by advertising all the routes through one or more periodically sent (about every 30 seconds) RIP update packets. With RIP, very large networks could take a long time to propagate new information concerning links that had failed.

OSPF was developed by the OSPF working group of the Internet Engineering Task Force (IETF) to utilize the Shortest Path First (also know as the Dijkstra) algorithm. This link-state algorithm normally only forwards updates about changes in member links in the form of Link-State Advertisements (LSAs). All routers learn about the AS they are a part of by sharing LSAs, and the routers then develop a database of their neighbors.

The OSPF database can be compared to a map that has information on all nearby towns and the conditions of the roads leading to them. Based on this map, a traveler (the IP packet) can be sent quickly on a route that results in the shortest amount of time traveled. RIP does not allow consideration for link bandwidth when choosing a route; instead it bases its selection on the lowest hop count. OSPF chooses the route with highest bandwidth when populating the routing table. For this reason, OSPF makes better routing decisions than RIP.

OSPF Architecture

OSPF is a hierarchical routing protocol. As such, it takes advantage of pre-planning for the large network in a top-down fashion.

As shown in See OSPF Topology, a large internetwork can be subdivided into smaller "areas." These areas branch from a central or "core" area 0. If all routers reside in a single area, that area is normally referred to as
area 0.

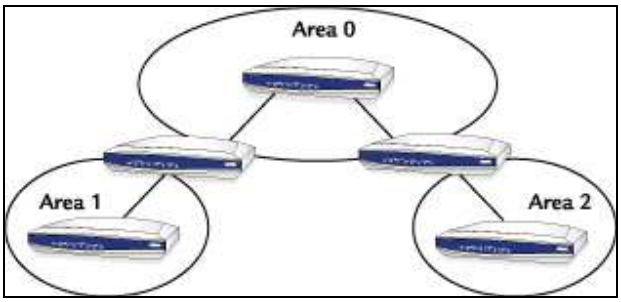

OSPF Topology

The benefits of hierarchical topologies include the following:

- Summarization - With a hierarchical topology, a single network address can represent an entire area.
- Control - Sensitive areas can be restricted from general access.
- Smaller routing tables - Only the summary routes are included.

- Reduced Link-State exchanges - Updates only reflect an area.
- Fewer SPF calculations - A single interface failure in the domain need not force a recalculation; this reduces the load on area routers.
- Larger internetworks - OSPF is designed to support large (greater than 15-hop) networks.

With Variable Length Subnet Masks, the internetwork can be designed such that summarization becomes a natural result. Suppose that the private network scheme 10.0.0.0 is used for the entire AS or domain (collection of all areas). Each area could then be provided with a single 16-bit subnet such as 10.1.0.0, 10.2.0.0, etc. Within a single area, the 16-bit subnet could then be further divided into 24-bit subnets (e.g., 10.1.1.0/24, 10.1.2.0/24, etc.). The routers between areas are known as Area Border Routers (ABRs); they perform summarization and control between areas. For connections to the Internet or to another AS, the Autonomous System Border Router (ASBR) performs Network Address Translation (NAT) on addresses or enforces policies between domains.

OSPF Network Types

OSPF defines three major network types, based on topology: broadcast, point-to-point, and Non-Broadcast Multi-Access (NBMA) (see See OSPF Network Types).

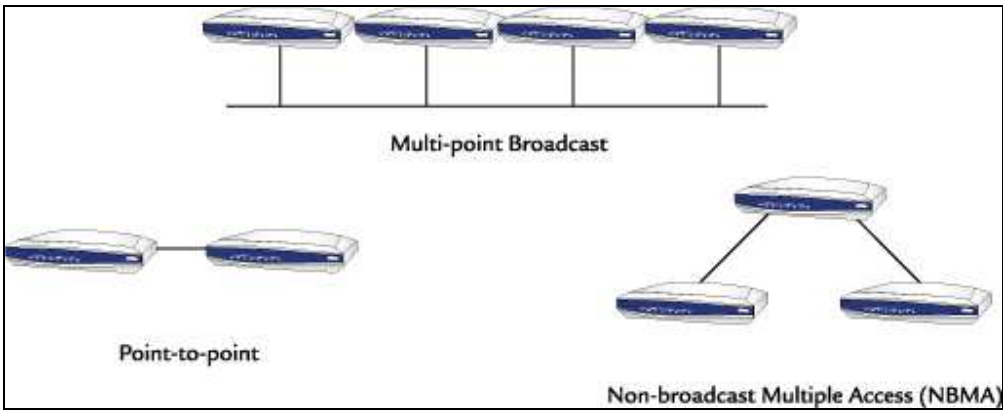

OSPF Network Types

In general, broadcast refers to networks in which routers can broadcast requests to each other, such as networks joined on a single Ethernet segment. To reduce updates on the network, the concepts of a Designated Router (DR) and Backup Designated Router (BDR) were developed within the OSPF protocol.

A DR is used to reduce routing traffic. Imagine that ten backbone routers are joined on a single 100-megabit Ethernet segment at the core or area 0. If a single router had to send every update to all other routers, it would open connections to nine other routers. If the other routers needed to do the same thing (at startup, for instance), a large amount of routing traffic would result. Now consider that all routers can elect a single "spokesrouter" (the DR) responsible for notifying all other routers about changes to the network. That single router could broadcast or flood the update to members of the group (multicast group) as it is told about any change. As a side benefit, all routers have the same routing table, reducing the opportunity for errors.

A special "Hello" packet is used to learn about each router. The Hello protocol allows each router to establish an adjacency with each neighbor. The Hello packet is used to share information about neighbors in the DR election process. The router with the highest priority becomes the DR; the next highest becomes the BDR.

The point-to-point network uses the Hello packet to establish an adjacency, but no DR or BDR is necessary. In general, Hello packets are sent by default at 10-second intervals for broadcast networks, and at 30-second intervals for point-to-point networks.

The NBMA network is used in partially meshed network types such as Frame Relay point-to-multipoint topologies. In general, Frame Relay networks employ a hub and spoke design. In some rare instances, a partially meshed or fully meshed topology could be used, but costs rise as the number of links increases. Fully meshed networks links increase exponentially for every new node added! A fully meshed 20-site network requires 190 links.

The NetVanta access routers are specially designed to function in the point-to-point mode in their role as remote routers in a point-to-multipoint network. In this topology, all routers are in the same subnet in a manner similar to broadcast networks. The same multipoint architecture could also use pure point-to-point configurations since many commercial routers can utilize sub-interface or virtual interface technologies in Frame Relay routing networks. In the case where point-to-point links are used, all subinterface pairs have their own subnets. Both topologies are supported and allow the network designer flexibility for his design.

Single Area OSPF Configuration

As discussed previously, an AS is a group of network devices under the administrative control of a single networking team. Usually your organization's network is considered one AS, although some very large companies may have several (for example, each division has its own AS).

While OSPF can scale to very large networks by breaking the AS into smaller areas, it may also make good sense to use OSPF in smaller networks with as few as five to ten routers. In these smaller networks (typically no more than 80 subnets), it is not always necessary to break the AS into areas.

See Enable the OSPF Process and See Specify OSPF networks delineate a sample configuration for OSPF in a single area. See Single Area OSPF Configuration illustrates the sample configuration.

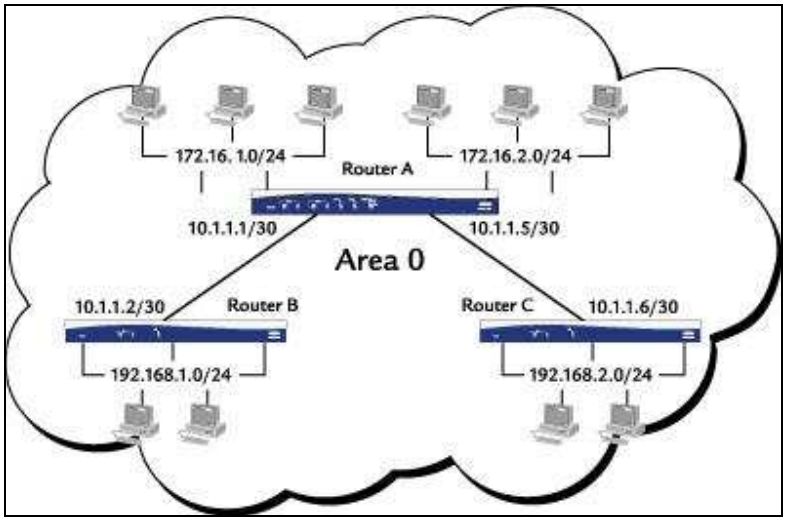

Single Area OSPF Configuration

Enable the OSPF Process

Prior to configuring OSPF on your router, you should execute the show ip protocols command (see See show ip protocols) to see all routing protocols that are currently configured. This may eliminate conflicting configurations or duplicate configuration effort.

Configuration for Router A

From the RouterA> prompt, enter into privileged mode by typing enable as shown in the following figure. If an enable password has been set, you will be prompted to enter it. You should now be at the enable prompt: RouterA#. Next, enter into global configuration mode by typing config t, and start the OSPF process by entering router ospf. Notice that it is not necessary to specify an OSPF process ID, since only one instance of OSPF is allowed in the ADTRAN OS.

```
RouterA> enable
RouterA# config t
RouterA(config)# router ospf
```

Specify OSPF networks

Use the network area command to specify which networks will participate in OSPF. The network area command serves two functions: (1) inform the OSPF process which interfaces will be participating in OSPF routing, and (2) specify which networks should be advertised by this router. Notice that the network area command uses a wildcard mask (sometimes referred to as an inverted mask) to indicate the subnet to be advertised. When using a wildcard mask, simply remember that the zeros portion of the mask defines the network address that must match and the ones portion of the mask represents the range of hosts in that subnet that do not have to match. Also note that area 0 has been specified. OSPF rules dictate that all areas must connect to area 0 (the transit or backbone area). Even though you are only configuring OSPF in a single area, always use area 0. This will be very helpful in the event that your network grows to include multiple areas.

```
RouterA(config-router)# network 10.1.1.0 0.0.0.3 area 0
RouterA(config-router)# network 10.1.1.4.0 0.0.0.3 area 0
RouterA(config-router)# network 172.16.1.0 0.0.0.255 area 0
RouterA(config-router)# network 172.16.2.0 0.0.0.255 area 0
```

Verification of OSPF Operation by using Show Commands

Show commands are useful in verifying OSPF configuration and operation. For example, you can display existing configuration information about the OSPF database, timing, neighbor adjacencies, protocols, and link states in your network.

show ip route

Once OSPF is configured, you can examine the routing table to determine whether routes are being learned via OSPF. From the RouterA# prompt, enter show ip route as shown in the following figure. Notice that routes learned via OSPF are marked with an "O." In a single area configuration intra-area there will be no inter-area (IA) entries, although you may have external entries (N1, N2, E1, or E2) in the route table if you are redistributing routes from another routing protocol such as RIP.

```
RouterA#show ip route
Codes: C - connected, S - static, R - RIP, O - OSPF, B - BGP
       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

Gateway of last resort is not set

C    10.1.1.0/30 is directly connected, ppp 1
C    10.1.1.2/32 is directly connected, ppp 1
C    10.1.1.4/30 is directly connected, ppp 2
C    10.1.1.6/32 is directly connected, ppp 2
C    172.16.1.0/24 is directly connected, eth 0/1
C    172.16.2.0/24 is directly connected, eth 0/2
O    192.168.1.0/24 [110/66] via 10.1.1.2, ppp 1
O    192.168.2.0/24 [110/66] via 10.1.1.6, ppp 2

RouterA#
```

show ip ospf

Verify global OSPF configuration parameters by using the show ip ospf command (see figure below). Notice that the router ID (172.16.2.1) and the global OSPF timers are listed with this command as well as a count of interfaces and areas participating in OSPF. If you suspect that the SPF algorithm may be running too often, check the number of times the SPF has been executed with this command. This could be an indication of instability within your network. An interface that is cycling between an "up" state and a "down" state frequently ("flapping") can be the source of frequent SPF calculations. Depending on the size of your network, the SPF calculation can be very processor intensive and thus negatively impact network performance.

```
RouterA#show ip ospf
Summary of OSPF Process with ID: 172.16.2.1
  Supports only single Type Of Service routes (TOS 0)
  SPF delay timer: 5 seconds, Hold time between SPFs: 10 seconds
  LSA interval: 240 seconds
  Number of external LSAs: 0, Checksum Sum: 0x0
  Number of areas: 1, normal: 1, stub: 0, NSSA: 0
    Area (0) 172.16.2.1
      Number of interfaces in this area: 4
      Authentication type: 0
      SPF algorithm execution count: 25
      Number of LSAs: 2, Checksum Sum: 0x157e9
RouterA#
```

show ip ospf neighbor

Display a summary of OSPF neighbor adjacencies by using the show ip ospf neighbor command (see figure below). The output of this command includes the following headers:

- Neighbor ID - The router ID of this neighbor. Note that this may differ from the IP address of the interface or link over which we have formed an adjacency.
- Pri - OSPF Router priority. This is used for the election of DR and BDR on a broadcast network such as an Ethernet LAN. The default OSPF Router priority is 1, but the priority can be manipulated on a per interface basis with the ip ospf priority interface configuration command. A higher number increases the preference of this router in the DR/BDR election. Setting the OSPF router priority to 0 would prevent this router from becoming either the DR or BDR on a particular broadcast link.
- State - Current state of the neighbor relationship. A state of "FULL" indicates proper operation for point-to-point links and broadcast links for the DR or BDR routers. All other routers on a broadcast network should reach the "Two-Way" state.
- Address - The address of the neighboring router on the interface specified.
- Interface - The interface on which this neighbor adjacency is formed.

```
RouterA#show ip ospf neighbor
Neighbor ID     Pri State           Dead Time   Address         Interface
192.168.1.1       1 FULL/PTPT       00:00:38    10.1.1.2        ppp 1
192.168.2.1       1 FULL/PTPT       00:00:35    10.1.1.6        ppp 2
192.168.103.193   1 FULL/BDR        00:00:38    172.16.1.2      eth 0/1
RouterA#
```

show ip ospf neighbor detail

Verify complete detail of OSPF neighbor adjacencies by using the show ip ospf neighbor detail command (see figure below). The detail keyword displays the number of state changes and the DR/BDR information. Numerous state changes occurring over a short interval may indicate an interface or link problem with this neighbor.

```
RouterA#showip ospf neighbor detail
Neighbor ID: 192.168.1.1, Interface IP Address: 10.1.1.2
   In area: 0, Via interface ppp 1
   Priority: 1, State: FULL, State changes: 5
   DR: 0.0.0.0, BDR: 0.0.0.0
   Options: 1
   Dead timer due: 00:00:31
Neighbor ID: 192.168.2.1, Interface IP Address: 10.1.1.6
   In area: 0, Via interface ppp 2
   Priority: 1, State: FULL, State changes: 5
   DR: 0.0.0.0, BDR: 0.0.0.0
   Options: 1
   Dead timer due: 00:00:38
Neighbor ID: 192.168.103.193, Interface IP Address: 172.16.1.2
   In area: 0, Via interface eth 0/1
   Priority: 1, State: FULL, State changes: 5
   DR: 172.16 1.1, BDR: 172.16.1.2
   Options: 1
   Dead timer due: 00:00:39
RouterA#
```

show ip protocols

Display a summary of all routing protocols enabled on a router by using the show ip protocols command as shown in the following figure. It is recommended that the show ip protocols command be used prior to configuring any routing protocol to avoid duplicate effort, for example if the routing protocol is already configured or to ensure consistent routing protocol configuration. This command is also useful for confirming which networks have been configured for a particular routing protocol.

```
RouterA#show ip protocols

Routing protocol is "ospf"
  Redistributing: ospf
  Routing for networks:
    10.1.1.0 0.0.0.3
    10.1.1.4 0.0.0.3
    172.16.1.0 0.0.0.255
    172.16.2.0 0.0.0.255
  Routing Information Sources:
    Gateway         Distance      Last Update
    172.16.2.1      110           02:30:05
    192.168.1.1     110           02:36:54
    192.168.2.1     110           02:53:33
RouterA#
```

show ip ospf database

Display a summary of the OSPF link state database by using the show ip ospf database command as shown in the following figure. All routers in an OSPF area must have an identical link state database. Included in the output of this command are the following headers:

- Link ID - Identification for each LSA.
- Adv Router - Router ID of the router that sourced this LSA.
- Age - The age of this LSA in seconds. The maximum value is one hour (3600 seconds).
- Seq# - Sequence number of the LSA. Begins with 0x80000001 and counts up to 0x8fffffff.

- Checksum - LSA checksum. This value is used to ensure that the LSA is uncorrupted as it is flooded throughout the area.

```
RouterA#showip ospf database

Router Link States, Area 0
  Link ID         Adv Router      Age    Seq #       Checksum
  172.16.2.1      172.16.2.1      899    0x800000CE  0x8EDC
  192.168.1.1     192.168.1.1     900    0x8000009D  0x1DDF
  192.168.2.1     192.168.2.1     1521   0x80000071  0x1509

  Net Link States, Area 0
  Link ID         Adv Router      Age    Seq #       Checksum
  172.16.1.1      172.16.2.1      376    0x8000005F  0xD257
RouterA#
```

Troubleshooting OSPF by using Debug Commands

Debug commands help identify and eliminate adjacency problems within a network.

debug ip ospf events

Since the formation of OSPF neighbor relationships is one of the more common sources of OSPF problems, the debug ip ospf events command is a particularly useful troubleshooting tool. From the RouterA# prompt, enter debug ip ospf events as shown in the following figure. In the example, a new neighbor relationship is being formed. Notice that the output shows each transition in the formation of the neighbor relationship.

```
RouterA#debug ip ospf events
RouterA#22:53:00: OSPF: Neighbor 192.168.103.193 state changed from just created
 to DOWN
22:53:00: OSPF: Neighbor 192.168.103.193 state changed from DOWN to INIT
22:53:00: OSPF: Neighbor 192.168.103.193 state changed from INIT to EX_START
22:53:01: OSPF: Neighbor 192.168.103.193 state changed from EX_START to
EXCHANGE
22:53:01: OSPF: Neighbor 192.168.103.193 state changed from EXCHANGE to
LOADING
22:53:01: OSPF: Neighbor 192.168.103.193 state changed from LOADING to FULL
```

debug ip ospf adj

Another useful tool for troubleshooting OSPF adjacency problems is the debug ip ospf adj command. From the RouterA# prompt, enter debug ip ospf adj as shown in the following figure. Again, a new neighbor has been discovered and the neighborship is being formed. This command helps to debug issues involving the election of the DR and BDR on a broadcast network such as an Ethernet LAN.

```
outerA#debug ip ospf adj
00:33:52: OSPF: Newneighbor discovered: 192.168.2.1
00:33:52: OSPF: Adjacency should be established with neighbor 192.168.2.1 because
interface is point-to-point
00:33:52: OSPF: Establishing adjacency with 192.168.2.1
00:33:52: OSPF: Sending database description packet to 192.168.2.1 with sequence
0222a2da on ppp 2
00:33:52: OSPF: Neighbor 192.168.2.1 is slave
00:33:52: OSPF: Sending database description packet to 192.168.2.1 with sequence
0000559c on ppp 2
00:33:52: OSPF: Sending database description packet to 192.168.2.1 with sequence
0000559d on ppp 2

RouterA#
```

Sample Configuration

The entire NetVanta 3305 configuration used in the sample network is listed below.

```
!
hostname "RouterA"
no enable password
!
ip subnet-zero
ip classless
ip routing
!
event-history on
no logging forwarding
no logging email
logging email priority-level info
!
!
!
!
interface eth 0/1
  ip address 172.16.1.1 255.255.255.0
  no shutdown
!
interface eth 0/2
  ip address 172.16.2.1 255.255.255.0
  no shutdown
!
!
!
interface t1 1/1
  clock source internal
  tdm-group 1 timeslots 1-24 speed 64
  no shutdown
!
interface t1 1/2
  shutdown
!
interface t1 2/1
  clock source internal
  tdm-group 2 timeslots 1-24 speed 64
  no shutdown
!
```

```
interface t1 2/2
  shutdown
!
interface ppp 1
  ip address 10.1.1.1 255.255.255.252
  no shutdown
  cross-connect 1 t1 1/1 1 ppp 1
!
interface ppp 2
  ip address 10.1.1.5 255.255.255.252
  ip ospf network point-to-point
  mtu 1520
  no shutdown
  cross-connect 2 t1 2/1 2 ppp 2
!
!
router ospf
network 10.1.1.0 0.0.0.3 area 0
network 10.1.1.4 0.0.0.3 area 0
network 172.16.1.0 0.0.0.255 area 0
network 172.16.2.0 0.0.0.255 area 0
!
!
!
no ip n-form agent
no ip http server
no ip http secure-server
no ip snmp agent
no ip ftp agent
!
!
!
!
line con 0
no login
!
line telnet 0 4
login
!
end
```