



OSPF -- Part III

Peter J. Welcher

Introduction

Three months back we looked at basic configuration of OSPF routing on Cisco routers. In the last two months we looked at summarizable ranges of addresses, and how to configure summarization for OSPF. If you missed some of these articles, please look for them at <http://www.netcraftsmen.net/welcher/papers/index.htm> , along with other previous articles.

This month we'll round out our coverage of OSPF with some of the other features that can be configured in Cisco routers.

OSPF Network Type

OSPF has three network types:

- Broadcast networks (Ethernet, Token Ring, FDDI)
- Nonbroadcast multiaccess networks (SMDS, Frame Relay, X.25, Classic IP Over ATM)
- Point-to-point networks (HDLC, PPP)

You can configure an interface as any type you wish.

Broadcast is typically used on networks where broadcast and multicast are supported. The idea is to send out one OSPF multicast that reaches multiple receivers. This can be appropriate with X.25 and Frame Relay if you have a full mesh of PVC's. Loss of a PVC can cause problems with OSPF broadcast networks, however, since some routers with a common subnet will no longer be able to transmit directly to each other.

Non-broadcast is used where broadcast/multicast is not feasible. The drawback to this interface type is that neighbors must be explicitly configured, which can get to be a nuisance.

Point-to-multipoint is for use in partial mesh situations, such as typical Frame Relay or ATM networks. It causes the generation of multiple host routers. It saves you from having to configure neighbor statements. It allows you to use the NBMA addressing model, where the FR or ATM cloud is one IP subnet. And it tolerates loss of virtual circuits.

An example of how to configure OSPF network type:

```
interface serial 0
encapsulation frame-relay
ip address 10.1.1.1 255.255.0.0
ip ospf network point-to-multipoint
```

To do this as non-broadcast

```
interface serial 0
encapsulation frame-relay
ip address 10.1.1.1 255.255.0.0
ip ospf network non-broadcast

router ospf 1
network 10.0.0.0 255.0.0.0 area 10
neighbor 10.1.2.3
neighbor 10.2.3.4
(etc.)
```

Redistribution into OSPF

Let's leave this topic a bit quickly (preserving it for a future article?) by noting that yes, you can redistribute other IP routing protocols into OSPF in the usual way.

If you want redistribution of full tables of subnets, and not just summaries, you may wish to add the `subnets` keyword to your redistribution command. If you wish to control summarization, you can do so with the router mode `summary-address` command. For example,

```
router ospf 1
summary-address 10.1.0.0 255.255.0.0
```

causes subnets 10.1.1.0, 10.1.2.0, etc. to be summarized as 10.1.0.0 /16 instead.

Sometimes when you are at an Autonomous System boundary, i.e. where a company connects to the Internet, it suffices to advertise a default route. You might consider the following OSPF command for this:

```
router ospf 1
default-information originate
```

Note that you'll either want to use a route map to specify what routes should trigger generation and advertisement of the default, or you'll already have a source for a default route, or you'll use the keyword `"always"`. You'll probably also want to use a route map or distribute list to filter out unnecessary routes, that is, those which the default route renders superfluous.

Routes redistributed into OSPF are known as **external routes**.

Stub Areas

Stub areas are a powerful OSPF design technique. Stub areas substitute a default route for external routes. Suppose you were silly enough to redistribute 20,000 or 50,000 routes into your OSPF area 0. Making area 1 stubby means that only 0.0.0.0 /0 is advertised into area 1, greatly reducing the load on the routers there (also the traffic announcing link state changes). Even with just your average corporate

network where redistribution is present, stubby areas can reduce traffic and make OSPF more stable.

To configure an area as stubby:

```
router ospf 1
area 1 stub
```

Note that all routers in an area must be configured to think the area is stubby. The OSPF Hellos contain information about not only the area but the stubbiness of the area. If two neighboring routers don't agree, they won't form an OSPF adjacency, and won't exchange OSPF routing information.

There is a further Cisco refinement of this. Configuring an stub area border router (ABR) with the keyword `no-summary` causes it to withhold information from the area, advertising solely a default into the area, vice information about inter-area and external routes. This is not part of the OSPF standard, but as long as the ABR's to an area are Cisco routers and configured similarly, no harm done. This greatly reduces the size of the routing tables and the OSPF traffic in the **totally stubby** area. Example:

```
router ospf 1
area 1 stub no-summary
area 1 default-cost 10000
```

The default cost command specifies the cost of the default route originated into the stub area.

Another way of thinking about these two is to think about what you lose. With a stub area, you lose the ability to pick between multiple ABR's for routes to external networks. I generally think of external routes, which were redistributed into OSPF from another protocol, as being "far away". So unless you require that traffic to those destinations be able to wisely choose the best ABR to exit the area via, by all means, make the area stubby. With a totally stubby area, you lose the ability to choose between multiple ABR's for any traffic leaving the area (and going to area 0, of course). This will probably not matter to you unless the area is geographically rather spread out.

Rules: Area 0 cannot be stubby. You cannot redistribute (originate external routes) within stubby areas.

To work around the latter, OSPF has an option, the Not-So-Stubby-Area (NSSA) option. This is supported in IOS 11.2.

To configure it:

```
router ospf 1
area 1 nssa
```

You should probably read more about NSSA's before trying to use this feature, so we won't go into the configuration options.

Routers in the area must all agree the area is NSSA. You should avoid redistributing on a NSSA ABR because confusion may result.

Authentication

OSPF has authentication, both to control formation of neighbor relationships and to validate OSPF update authenticity. Configure:

```
interface ethernet 0
```

```
ip ospf authentication-key topsecretkey

router ospf 1
area 1 authentication
```

This causes simple password authentication to be used. This is not particularly secure but makes sure you hear about router wannabe's. (For example, someone introduces an NT or UNIX box that speaks OSPF, which is readily available for either OS. If you accept their routing information, your network can break if they misconfigure the address or subnet mask on any interface.)

The alternative is MD5 authentication:

```
interface ethernet 0
ip ospf message-digest-key 100 md5 topsecretkey

router ospf 1
area 1 authentication message-digest
```

This is intended to provide more stringent security. With that go increased management challenges. See the reference guide for the process to change MD5 authentication strings on routers in an area.

Router ID

The router ID in OSPF is the highest IP address on any **active** interface. Configuring an address on a loopback interface over-rides this. It also means you don't lose your OSPF if you save your configuration with no active interfaces, for example as part of performing password recovery.

Virtual Links

Virtual links allow you to reconnect a broken area 0 across another area. They also allow you to take a router that connects say areas 2 and 3, and make it part of area 0, allowing it to route packets between areas 2 and 3. You can also use it to "connect two ABR's on both sides of the border". That is, if the LAN connection between two ABR's is in area 1, you can make it also usable as part of area 0 with a virtual link. This latter usage can get out of hand if carried to extremes.

If you read the OSPF RFC, you'll see that the behavior of routers with virtual links is a bit more complex than anything else in there. I generally prefer to avoid virtual links where possible in designs. There are usually quite palatable alternatives. If nothing else, virtual links increase the conceptual complexity of the OSPF design. Some would reason that they also increase its fragility.

To configure a virtual link, you point two routers at each other. One should be connected to area 0 already.

```
router ospf 1
area 2 virtual link 10.1.1.1
```

(assuming 10.1.1.1 is the router id for the other end of the virtual link).

OSPF Interface Parameters

There are a number of interface tweaks for OSPF. They allow you to configure the cost on an

interface, if you don't like the default or need to match another brand of equipment. You can also alter the OSPF timers: be sure to do this consistently throughout an area! And you can influence Designated Router (DR) election on a LAN segment with a priority, although practically speaking there appears to be no reason to do so (the first router to boot completely may become the DR on a segment, and there is no way to displace the acting DR).

```
interface ethernet 0
ip ospf cost cost
ip ospf retransmit-interval seconds
ip ospf transmit-delay seconds
ip ospf hello-interval seconds
ip ospf dead-interval seconds
ip ospf priority number
```

Troubleshooting OSPF

Configure

```
ip ospf name-lookup
```

to use DNS to display names instead of router ID or neighbor ID in IOS show command output.

From there, I'd start with

```
show ip ospf interface
```

to make sure all the appropriate interface addresses have matched the base address/wildcard mask pairs in network commands (and have been placed into the correct areas, with correct stubbiness and timers).

After checking that, you might try

```
show ip ospf neighbor
```

since if OSPF routing isn't taking place, the usual reason is that some routers have not formed the necessary adjacencies. The output from this can show if communication is one way, which neighbor relationships if any have formed, and so on.

There are two OSPF-specific debug commands:

```
debug ip ospf events
debug ip ospf packet
```

The former of these can be helpful in troubleshooting problems involving failure to form an adjacency (as revealed by the `show ip ospf neighbor` command). It shows Hello transmission, reception, and any problems with the Hellos. If you don't see Hellos going out, check `show ip ospf interface`, the interface apparently didn't match any of the network commands.

Other OSPF show commands can be somewhat overwhelming at first, but can provide a good way to learn more about OSPF and what the router is doing in your OSPF network:

```
show ip ospf [process-id]
show ip ospf [process-id area-id] database
show ip ospf [process-id area-id] database [router][link-state-id]
```

```
show ip ospf [process-id area-id] database [network] [link-state-id]
show ip ospf [process-id area-id] database [summary] [link-state-id]
show ip ospf [process-id area-id] database [asb-summary] [link-state-id]
show ip ospf [process-id] database [external][link-state-id]
show ip ospf [process-id area-id] database [database-summary]
show ip ospf border-routers
show ip ospf neighbor [interface-name] [neighbor-id] detail
show ip ospf virtual-links
```

OSPF Metrics

In IOS 10.3 and later, by default, OSPF calculates the OSPF metric for an interface by dividing ten to the eighth by the configured bandwidth. For example, a 64K link gets a metric of 1562, while a T1 link gets a metric of 64. This calculation gives FDDI a metric of 1. If you have multiple links with high bandwidth, you might want to specify a larger numerator to be able to differentiate the cost on those links. (Of course, you need to do this consistently on all your OSPF routers). To make the numerator 10 to the 10-th, configure:

```
router ospf 1
ospf auto-cost reference-bandwidth 10000000000
```

Summary

For more information about OSPF, there are various design guides and informational documents about OSPF on the Cisco documentation CD and on CCO. The Christian Huitema routing book is well thought of. Terry Slattery and Bill Burton of Mentor Technologies have authored a McGraw-Hill book on Advanced IP Routing which includes coverage of OSPF. Finally, the OSPF RFC's are actually pretty readable. Start with RFC 2178.

Dr. Peter J. Welcher (CCIE #1773, CCSI #94014) is a Senior Consultant with Chesapeake NetCraftsmen. NetCraftsmen is a high-end consulting firm and Cisco Premier Partner dedicated to quality consulting and knowledge transfer. NetCraftsmen has nine CCIE's, with expertise including large network high-availability routing/switching and design, VoIP, QoS, MPLS, network management, security, IP multicast, and other areas. See <http://www.netcraftsmen.net> for more information about NetCraftsmen. Pete's links start at <http://www.netcraftsmen.net/welcher> . New articles will be posted under the Articles link. Questions, suggestions for articles, etc. can be sent to pjw@netcraftsmen.net .

10/98
Copyright 1998, Peter J. Welcher