

OSPF Redistribution Among Different OSPF Processes

Document ID: 4170

Introduction

Why Redistribute Between Two OSPF Processes?

- Filtering the OSPF Route
- Keeping Different OSPF Domains Separate

Redistributing Between Different OSPF Processes

- OSPF Route Preference Rule
- A Single Redistribution Point
- Two Redistribution Points

Proposed Solution

- Using the distance 255 Command
- Filtering Routes Based on Tags
- Using the match internal Keyword While Redistributing
- Prefix-Based Filtering
- Prefix-Based Filtering and Prefix-Based Administrative Distance

Summary

[NetPro Discussion Forums – Featured Conversations](#)

Related Information

Introduction

This document provides guidelines for the redistribution of Open Shortest Path First (OSPF) between different processes. Redistribution between different processes is difficult, and special measures are necessary for the proper operation of the network. This document also highlights some changes introduced in Cisco IOS® software.

Why Redistribute Between Two OSPF Processes?

There can be several reasons for redistribution between multiple processes. These are a few examples:

- To filter an OSPF route from part of the domain
- To separate different OSPF domains
- To migrate between separate domains

Although redistribution between different processes might be necessary in certain cases, an alternate design solution (if possible) is a more appropriate choice, as will be discussed in the subsections of this section.

Filtering the OSPF Route

Intra–Area Routes

In OSPF, IP prefixes within an area are not exchanged directly between routers. They are part of the Link State Advertisement (LSA) that also announces the topology of the network; therefore, there is no way to filter routes within an area.

Note: Local filtering on a router (which can be done to prevent some routes from being installed in a given

router) is not considered to be actual route filtering. This is normally accomplished with the **distribute-list** command under router OSPF.

One solution would be to use a different process and filter the desired routes on the redistributing routers; however, doing this actually separates the area into two domains. A better design would be to separate the area into different areas and use the Cisco IOS type 3 filtering feature, which is explained later.

Inter–Area Routes

In OSPF, all routers within an area have the exact same topology. An area has no knowledge of the topology of another area; therefore, it relies on the information announced by the attached area border routers (ABRs).

The information advertised inside an area by an ABR (as a type 3 LSA) is actually the IP prefixes that are learned from remote areas or that are calculated for other attached areas.

An ABR originates these routes:

- Non–backbone intra–area routes into the backbone
- Backbone intra–area and inter–area routes into non–backbone area

Therefore, between areas there is a distance vector behavior that can be leveraged to filter routes between areas.

Cisco IOS software implemented an inter–area filtering feature. For more information on this feature, refer to OSPF ABR Type 3 LSA Filtering.

External Route Filtering

Because external routes are advertised as type 5 LSAs and are flooded domain–wide, except into stub areas and not–so–stubby areas (NSSAs), there is currently no way to filter a type 5 LSA. One solution is to have a different process and to filter between processes while redistributing. In the near future, Cisco IOS software will have a means to achieve type 5 filtering on area borders.

Keeping Different OSPF Domains Separate

It is seen as a common practice to use different OSPF processes in order to separate different IP routing domains, either for administrative purposes or to segment the routing domain and to control routing information in the redistribution point.

It should be noted, however, that instability in one domain could affect the other domain. For example, if there is a change in the OSPF network (type 1 and 2) where an autonomous system boundary router (ASBR) resides between the two domains, all the type 5 LSAs will be re–originated and flooded throughout the remote domain. Thus, if there is a constant instability in one network, this can lead to a constant injection and withdrawal of type 5 LSAs in the other domain.

A better design choice is to use Border Gateway Protocol (BGP) between different domains. In this case, OSPF exchange between different domains goes through BGP; and, because BGP has dampening capability, the instability in one domain will be less visible in the other domain.

Redistributing Between Different OSPF Processes

As mentioned previously, there can be an alternative solution to redistributing between multiple processes. The section shows how redistribution between different processes should be planned carefully, depending on

the number of redistribution points.

OSPF Route Preference Rule

The OSPF route selection rule is that intra-area routes are preferred over inter-area routes, which are preferred over external routes. However, this rule should apply to routes learned via the same process. In other words, there is *no* preference between external routes from one process compared to internal routes of other process.

The preference rule between a given OSPF process and any other process (be it OSPF or another routing protocol) should follow the administrative distance rule. However, because different OSPF processes will have the same administrative distance by default, OSPF distance should be configured explicitly for different OSPF processes, to achieve the desired behavior.

Note: Before Cisco bug ID CSCdi7001 (registered customers only) fixed in Cisco IOS Software Release 11.1 and later administrative distance between processes did not work properly, and the internal routes of one process were preferred over the external routes of another process.

A Single Redistribution Point

When there is a single redistribution point, all exchange between the domains occurs at a single point, and there is no way that a redistribution loop can form. This is a configuration example:

Figure 1




Router A Configuration
<pre>router ospf 1 redistribute ospf 2 subnet router ospf 2 redistribute ospf 1 subnet</pre>

Two Redistribution Points

It is more complicated when there are two points of redistribution. If redistribution is done in both points of a network without any special care, there may be unexpected results.

Consider the next topology, where Router A and Router B are redistributing mutually between both domains. This configuration is *not* working, as will be demonstrated later in this section.

Figure 2



Routers A and B Configuration
<pre> router ospf 1 redistribute ospf 2 subnet router ospf 2 redistribute ospf 1 subnet </pre>

Given a Network N in Domain 1, Router A and B learn Network N as an internal route in Domain 1. Because they are redistributing process 1 into process 2, the same Network N is learned in Domain 2 as an external route.

Now, in each router, the internal network learned via one process is competing against the external network of another process. As previously mentioned, there is no preference rule between different processes; therefore, the result would be indeterministic, as both process have the same administrative distance.


Note: This could lead to a constant injection and withdrawal of a type 5 from one process to the other.

Before Cisco bug ID CSCdw10987 (registered customers only) (integrated in Cisco IOS Software Releases 12.2(07.04)S, 12.2(07.04)T, and later), the last process to make an shortest path first algorithm (SPF) would have won, and the two processes overwrite other routes in the routing table. Now, if a route is installed via one process, it is not overwritten by another OSPF process with the same administrative domain (AD), unless the route is first deleted from the routing table by the process that initially installed the route in the routing table.

Administrative Distance

When you are using redistribution between multiple processes, you can use administrative distance to prefer one process over another process, because OSPF route preferences apply only within the same process. However, this is not enough for proper operation in the network, as explained later in this section.

Figure 3



Routers A and B Configuration
<pre> router ospf 1 redistribute ospf 2 subnet distance ospf external 200 router ospf 2 redistribute ospf 1 subnet distance ospf external 200 </pre>


Network Operation Without Network Failure

Consider a Network N in Domain 1, where N will be known as an internal route in Domain 1 and will be redistributed by both Router A and by Router B. Because the administrative distance of external routes has been increased, Router A and Router B will chose OSPF process 1 to reach Network N.

In a more general way, all networks internal to Domain 1 will be reached through Domain 1 and all networks internal to Domain 2 will be reached through Domain 2, by both Router A and Router B. Other routers in each domain pick up the closest ASBR (if metric type 2 is used) or the shortest path through one of the ASBRs (if metric type 1 is used).

If there are prefixes external to both domains (that come from some other redistribution points), the same problem will still occur because the administrative distance for those external routes is the same in both processes. If you make the administrative distance for external processes different, it will not resolve the problem. This is an example:

Figure 4



Router C (ASBR) advertises external N into Domain 1. This prefix gets redistributed by Router A and Router B into Domain 2 and will reach each of the routers; therefore, N will be external in both domains. To have proper operations, the administrative distance of external routes needs to be different for the two processes, so that one domain is preferred over the other. Assume that the administrative distance for Domain 1 is set lower than Domain 2.

Now, if Router D (ASBR) is advertising external M into Domain 2, then this prefix is redistributed by Router A and Router B into Domain 1 and it will reach each of the routers. Thus, M will be external in both domains and, because the administrative distance is lower for Domain 1, M will be reachable via Domain 1. This sequence of events could occur:

1. Router A (Router B) redistributes M into Domain 1, and external M will reach Router B (Router A).
2. Because the administrative distance of Domain 1 is lower than Domain 2, Router A (Router B) will install M through Domain 1 and will set to maxage its previous originated LSA (event 1) into Domain 1.
3. Because M has been set to maxage in Domain 2, Router A (Router B) will install M though Domain 2 and, therefore, will redistribute M into Domain 2.
4. Same as event 1.

This cycle continues, and the way to fix it is to have the Domain 2 prefix reachable via Domain 2. However, if the administrative distance is set lower for Domain 2, then the same problem will occur for Domain 1 and for prefix N.

The solution is to set administrative distance based on prefix. See the [Prefix-Based Filtering](#) and [Prefix-Based Filtering and Prefix-Based Administrative Distance](#) sections for more information.

- 1.
5. Router A (Router B) receives the flushed LSA (event 1) from Router B (Router A). It will remove the Network N through process 2 and, therefore, flush the external N into Domain 1. Network N was learned through Domain 2 and redistributed into Domain 1.
6. Before Router A (Router B) receives the flushed LSA from Router B (Router A), it will install the external Network N through Domain 1, because N has been flushed through Domain 2.
7. Since Router A (Router B) installed Network N through process 1, it will generate an external N into Domain 2.

You can see that there is a race condition that could appear from one domain to the other domain. In events 1, 4, and 7, Router A generates an external Network N into Domain 2; and in events 2 and 5, Router A withdraws the prefix. The problem occurs because the routes learned through one domain are redistributed back to the same domain.

Proposed Solution

This section shows how to prevent a route that belongs to a domain from being redistributed back to the same domain, to avoid routing loops.

Using the distance 255 Command

The previous section shows how a routing loop is created if the prefixes learned from one domain are redistributed back to the same domain. Because redistribution is occurring from a routing table, you can prevent a route that belongs to Domain 1 and that is learned from the remote router over Domain 2 from being installed in the routing table. Therefore, the router will not redistribute those routes back to Domain 1.

To do this, issue the **distance 255 router_ID inverse_mask access-list** command. This command tells the router to deny all the prefixes that are received by a remote router with the specified router ID and that match the access control list (ACL) from the routing table.


Note: The **distance 255** command gives a distance of 255 to those routes and, therefore, prevents their installation into the routing table.

In Figure 6, Router A uses the **access-list 1** command to match all the routes in Domain 1 and uses the **distance 255** command under process 1 to deny routes received from Router B that match prefixes that belong to Domain 1.

When you use the **distance 255** command, it denies any route received from Router B that belongs to Domain 1. Because Router B redistributes all routes in Domain 1 into Domain 2, Router A will not install those routes and, therefore, will not redistribute back into Domain 1 again.

Note: The connected interface of Router B in Domain 1 should be excluded from the ACL.

Figure 6



Router A Configuration	Router B Configuration
<pre> router ospf 1 redistribute ospf 2 subnet distance 255 <Router B> 0.0.0.0 2 ! access-list 1 !--- Matches the router in Domain 2. router ospf 2 redistribute ospf 1 subnet distance 255 <Router B> 0.0.0.0 1 ! access-list 2 !--- Matches the route in Domain 1. </pre>	<pre> router ospf 1 redistribute ospf 2 subnet distance 255 <Router A> 0.0.0.0 2 ! access-list 1 !--- Matches the router in Domain 2. </pre>

The previous **distance ospf external 200** command is no longer needed because routes learned from the remote router via one of the processes are not installed.

This configuration works properly in cases where both routers have loose connectivity to the network (as described in [Network Operation Without Network Failure](#) and [Network Operation With Network Failure](#)). However, because the prefixes are denied from the routing table, the domains can not back up each other.


Note: You must explicitly list all of the prefixes of each domain in an ACL. Maintenance of such an ACL can be very difficult.

Filtering Routes Based on Tags

There is a new feature in Cisco IOS software (from Cisco bug ID CSCdt43016 (registered customers only)) that allows you to filter routes based on the tag. To prevent the redistribution of routes from one domain back into the same domain, a router can tag a route that belongs to a domain while it is redistributing, and you can filter those routes on the remote router based on the same tag. Because the routes will not be installed into the

routing table, they will not be redistributed back into the same domain.

Figure 7



Routers A and B Configuration
<pre>router ospf 1 redistribute ospf 2 subnet tag 1 distribute-list 1 route-map filter_domain2 in ! route-map filter_domain2 deny 10 match tag 2 route-map filter_domain2 permit 20 router ospf 2 redistribute ospf 1 subnet tag 2 distribute-list 1 route-map filter_domain1 in ! route-map filter_domain1 deny 10 match tag 1 route-map filter_domain1 permit 20</pre>

When you are redistributing from Domain 1, the routes are tagged with tag 1 and are filtered on the remote router based on the same tag. When you are redistributing from Domain 2, the routes are tagged with tag 2 and are filtered on the remote router based on the same tag.


Note: The previous **distance ospf external 200** command is no longer needed because the route learned from the remote router via one of the processes is not installed.

This configuration works properly in cases where both routers lose connectivity to the network (as described in Network Operation Without Network Failure and Network Operation With Network Failure). However, because the prefixes are denied from the routing table, the domains can not back up each other.

Using the match internal Keyword While Redistributing

When you are redistributing from a domain, you can use the **match internal** keyword to redistribute only the internal routes that belong to one domain into another domain. This prevents the redistribution of prefixes that are already external back into the same domain.

Figure 8



Routers A and B Configuration
<pre>router ospf 1 redistribute ospf 2 subnet match internal distance ospf external 200 ! router ospf 2 redistribute ospf 1 subnet match internal distance ospf external 200 !</pre>


This configuration works properly in cases where both routers lose connectivity to the network (as described in Network Operation Without Network Failure and Network Operation With Network Failure). One domain could back up the other domain.

If there are already external prefixes in either of the domains (such as external prefixes that were redistributed via another protocol), then those prefixes will not be redistributed to other domains, because only internal prefixes are redistributed. Also, there is no control over external prefixes, and all external prefixes will be blocked.

Prefix-Based Filtering

When you are redistributing from a domain, prefixes can be matched against an ACL to avoid the redistribution of prefixes that belong to one domain back into the same domain.

Figure 9



Router A and B Configuration

```

router ospf 1
redistribute ospf 2 subnet route-map filter_domain2
distance ospf external 200
!
route-map filter_domain2 permit 10
match ip address 1
!
access-list 1

!--- Matches the prefix in Domain 1.

router ospf 2
redistribute ospf 1 subnet route-map filter_domain1
distance ospf external 200
!
route-map filter_domain1 permit 20
match ip address 2
!
access-list 2


!--- Matches the prefix in Domain 2.

```

This configuration works properly in cases where both routers lose connectivity to the network (as described in Network Operation Without Network Failure and Network Operation With Network Failure). One domain could back up the other domain.

Note: You must explicitly list all the prefixes of each domain in an ACL. Maintenance of such an ACL can be very difficult. Another solution is to tag prefixes during distribution and then filter the corresponding tags.

Figure 10




Routers A and B Configuration
<pre> router ospf 1 redistribute ospf 2 subnet tag 1 route-map filter_domain2 distance ospf 2 external 200 ! route-map filter_domain2 deny 10 match tag 2 route-map filter_domain2 permit 20 router ospf 2 redistribute ospf 1 subnet tag 2 route-map filter_domain1 distance ospf 1 external 200 ! route-map filter_domain1 deny 10 match tag 1 route-map filter_domain1 permit 20 </pre>

Prefix-Based Filtering and Prefix-Based Administrative Distance

As mentioned in the Administrative Distance section, there is a need for a prefix-based administrative distance where there are external prefixes originated by other ASBRs in each domain. In the next example topology, ASBR1 and ASBR2 are redistributing Networks X and Y into Domain 1 and Domain 2, respectively.

This example uses an ACL to match all of the prefixes (internal and external) that belong to a domain, and it uses the **distance** command to increase the administrative distance of prefixes that do not initially belong to the corresponding domain.

Figure 11



Routers A and B Configuration

```

router ospf 1
redistribute ospf 2 subnet route-map filter_domain2
distance 200 0.0.0.0 255.255.255.255 2
!
route-map filter_domain2 permit 10
match ip address 1
!
access-list 1

!--- Matches the prefixes in Domain 1.

access-list 2

!--- Matches the prefixes in Domain 2.

router ospf 2
redistribute ospf 1 subnet route-map filter_domain1
distance 200 0.0.0.0 255.255.255.255 1
!
route-map filter_domain1 permit 10
match ip address 2
!
access-list 1

!--- Matches the prefixes in Domain 1.

access-list 2

!--- Matches the prefixes in Domain 2.

```

The **distance 200 0.0.0.0 255.255.255.255 2** command under process 1 sets the administrative distance of all prefixes that belong to Domain 2 to 200; therefore, Routers A and B use Domain 1 to reach prefixes that belong to Domain 1.

Note: You must explicitly list all the external prefixes of each domain in an ACL. Maintenance of such an ACL can be very difficult.

Summary

When there is more than one redistribution point between OSPF domains, routing loops can easily occur. To prevent routing loops, prefixes that belong to a domain should not be redistributed back to the same domain. Also, the administrative distances of the OSPF processes should be set correctly. These five methods have been proposed in this document:

- Use the **distance 255** command.
- Filter based on tags.
- Use the **match internal** keyword during redistribution.
- Use prefix-based filtering during redistribution.
- Use prefix-based filtering and prefix-based administrative distance.

The first two solutions prevent the routes that belong to a domain from being installed in the routing table, which prevents their redistribution back to the same domain.

Note: Because the prefixes are denied from the routing table, the domains can not back up each other.

You can use the last three solutions to back up a domain with another domain, if needed. However, you should note these caveats:

- The **match internal** solution does not allow you to have control over prefixes, and all external prefixes will be blocked from redistribution. In other words, if there are external prefixes from other ASBRs, then those LSAs will not be redistributed from one domain to the other.
- The use prefix-based filtering during redistribution solution allows a domain to back up another domain. However, backup only works correctly when there are no external routes from the other ASBR.
- The use prefix-based filtering and prefix-based administrative distance solution is the only solution that allows a domain to back up another domain in the presence of external routes from other ASBRs.

This document repeatedly refers to the use of one domain to back up another domain. It should be noted that backup means that, should Router A lose its connection to part of the domain through a given domain (such as Domain 1), then it could use the other domain (Domain 2) to correctly route to those destinations that can not be reached through Domain 1.

However, if a domain is partitioned because prefixes are not redistributed back to the original domain, then the other domain could not back up the partitioned domain unless prefixes are redistributed back to the original domain. However, as noted in the Administrative Distance and the Network Operation With Network Failure sections, this will introduce other problems.

NetPro Discussion Forums – Featured Conversations

Networking Professionals Connection is a forum for networking professionals to share questions, suggestions, and information about networking solutions, products, and technologies. The featured links are some of the most recent conversations available in this technology.

NetPro Discussion Forums – Featured Conversations for RP
Service Providers: MPLS
Virtual Private Networks: Services
Virtual Private Networks: Security

Related Information

- [OSPF Support Page](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2007 – 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jan 23, 2006

Document ID: 4170
