Redistribute OSPF Among Different OSPF Processes

Contents

Introduction

Background Information

Why Redistribute Between Two OSPF Processes

- Filter the OSPF Route
  - Intra-Area Routes
  - Inter-Area Routes
  - External Route Filtering
- Keep Different OSPF Domains Separate

RedistributeBetween Different OSPF Processes

- OSPF Route Preference Rule
- A Single Redistribution Point
- Two Redistribution Points
  - Administrative Distance
  - Network Operation Without Network Failure
  - Network Operation With Network Failure

Proposed Solution

- Usethe Distance 255 Command
- Filter Routes Based on Tags
- Use the Match Internal Keyword While Redistributing
- Prefix-Based Filtering
- Prefix-Based Filtering and Prefix-Based Administrative Distance

Summary

Related Information

Introduction

This document describes the guidelines for the redistribution of Open Shortest Path First (OSPF) between different processes.

Background Information

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:

- Filter an OSPF route from part of the domain

- Separate different OSPF domains
- Migrate between separate domains

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

Filter 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 in order 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, 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.

Keep 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 can 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 are 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 though BGP; and, because BGP has dampening capability, the instability in one domain can be less visible in the other domain.

Redistribute Between Different OSPF Processes

As mentioned previously, there can be an alternative solution to redistribution between multiple processes. The section shows how redistribution between different processes can be planned carefully, dependent upon 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 can 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) can use the administrative distance rule. However, because different OSPF processes have the same administrative distance by default, OSPF distance can be configured explicitly for different OSPF processes in order to achieve the desired behavior.

Note: Before Cisco bug ID CSCdi7001 - 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:

Image 1



Router A Configuration

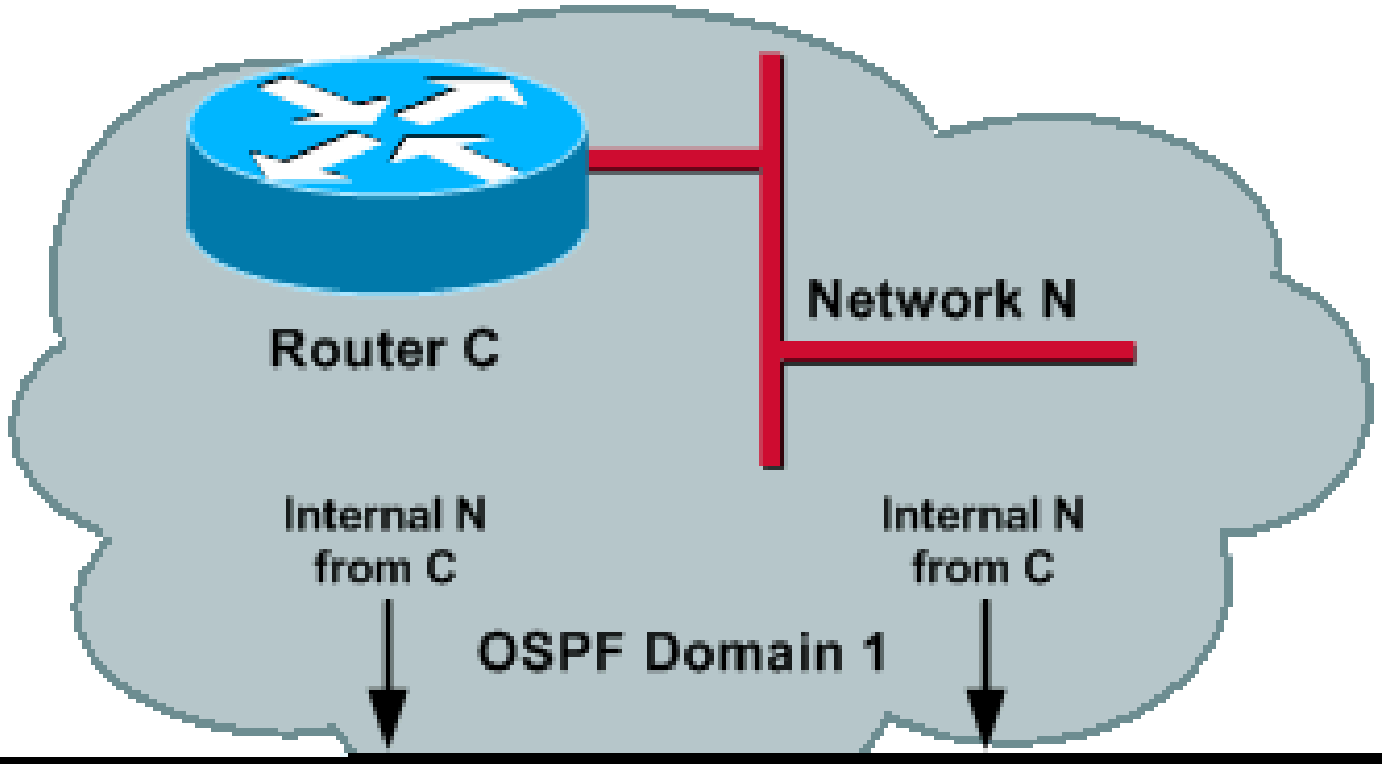
`router ospf 1`
`redistribute ospf 2 subnet`
 	
`router ospf 2`
`redistribute ospf 1 subnet`

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 can be unexpected results.

Consider the next topology, where Router A and Router B redistribute mutually between both domains. This configuration does not work, and is demonstrated later in this section.

Image 2



Routers A and B Configuration
router ospf 1 redistribute ospf 2 subnet router ospf 2 redistribute ospf 1 subnet

Given a Network N in Domain 1, Router A and B learn Network N as an internal route in Domain 1. Because they redistribute 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 competes 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 processes have the same administrative distance.

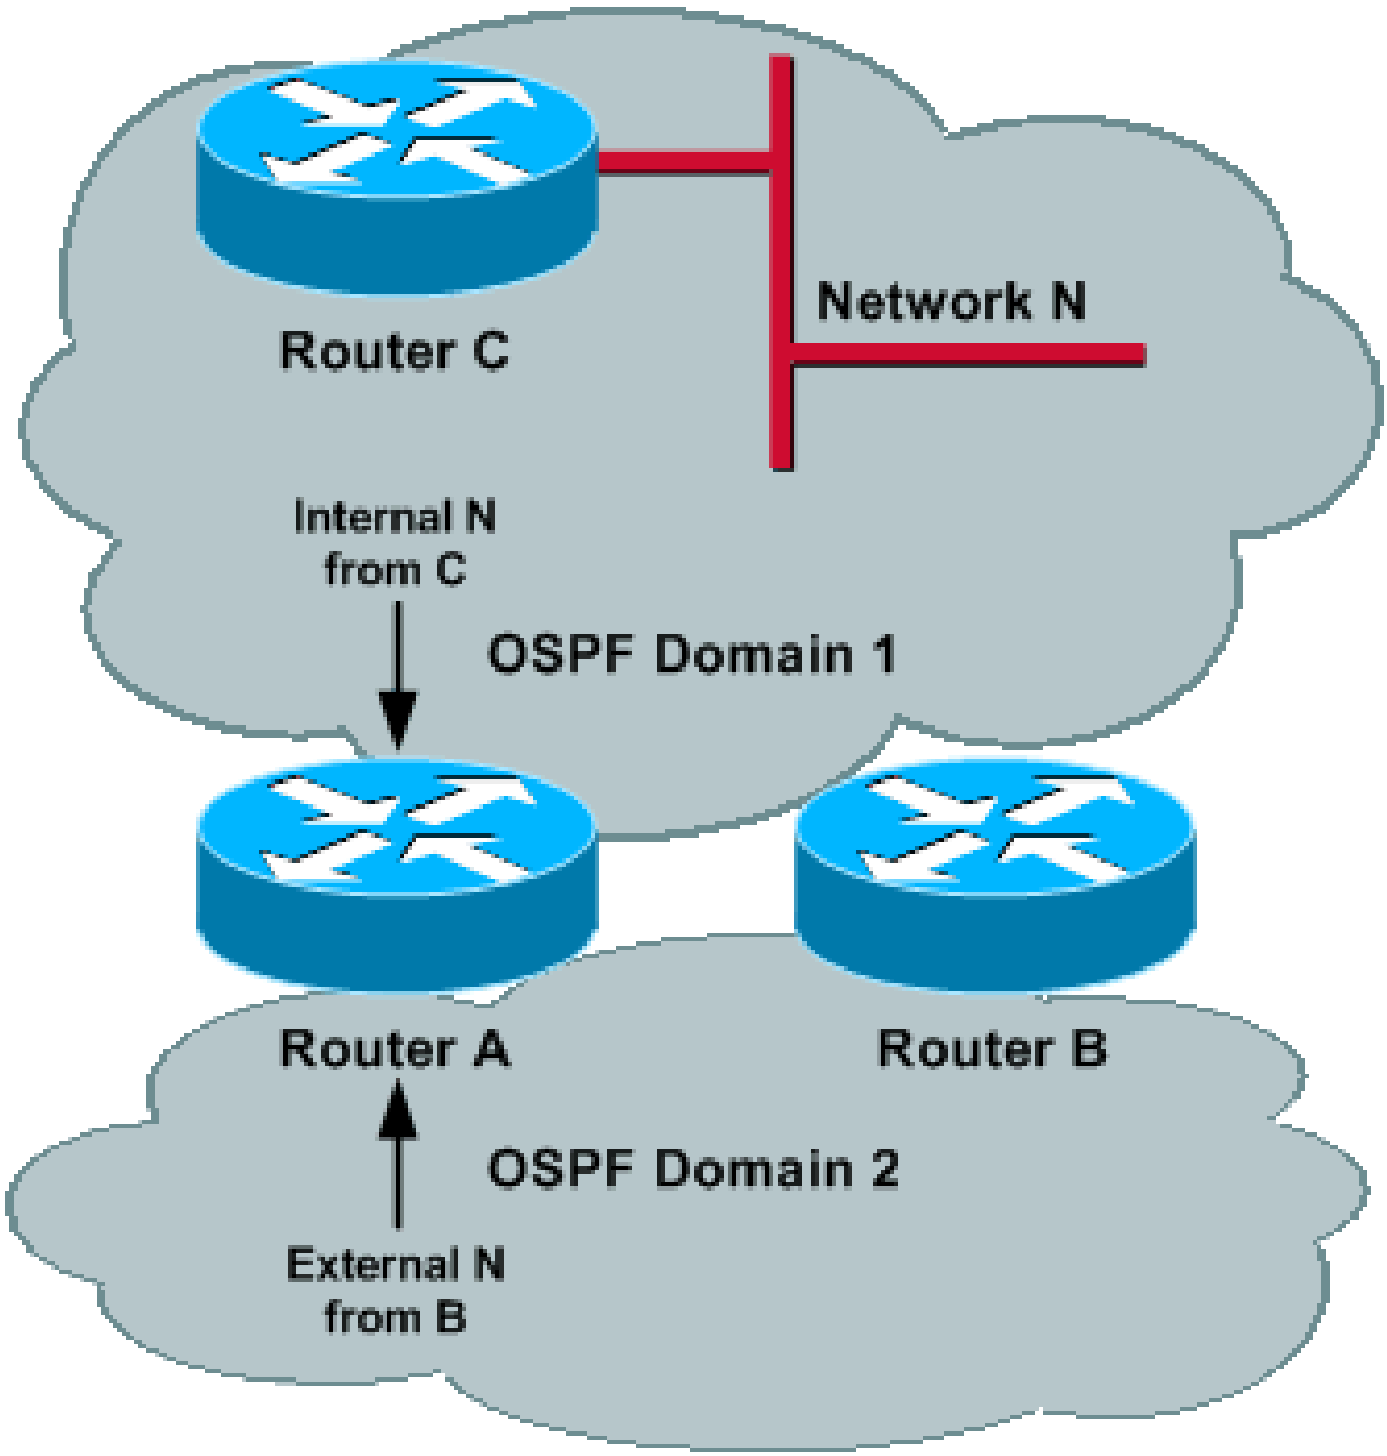
Note: This could lead to a constant injection and withdrawal of a type 5 from one process to the other.

Before Cisco bug ID CSCdw10987 (registered customers only and integrated in Cisco IOS® Software Releases 12.2(07.04)S, 12.2(07.04)T, and later), the last process to make a 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 use redistribution between multiple processes, you can use administrative distance in order 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.

Image 3



Routers A and B Configuration
<#root> router ospf 1 redistribute ospf 2 subnet **distance ospf external 200**   router ospf 2 redistribute ospf 1 subnet **distance ospf external 200**

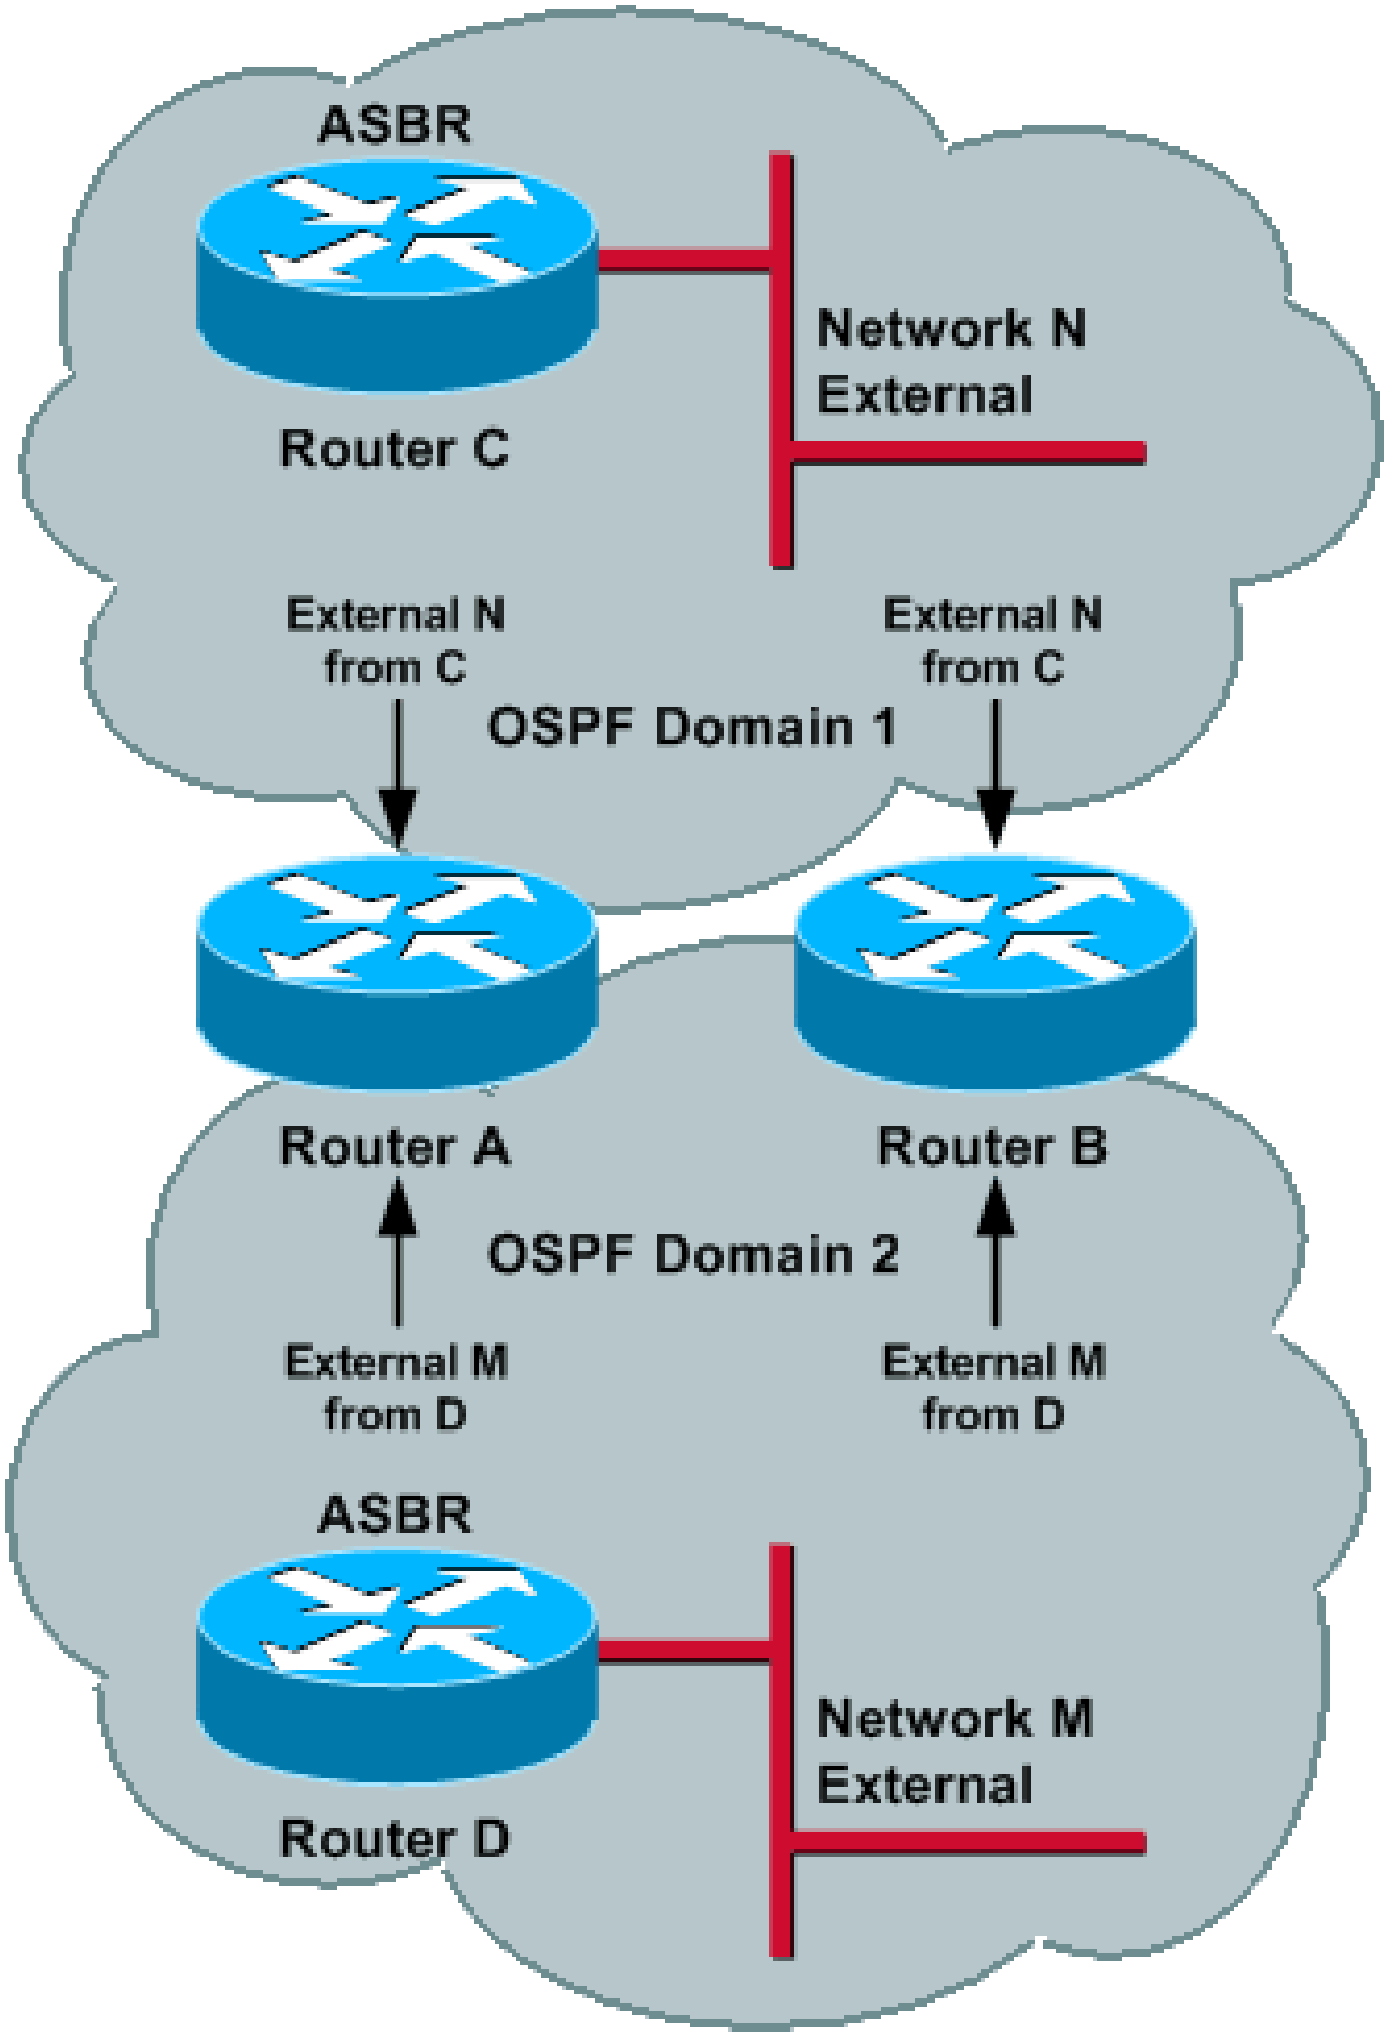
Network Operation Without Network Failure

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

In a more general way, all networks internal to Domain 1 are reached through Domain 1 and all networks internal to Domain 2 are 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).

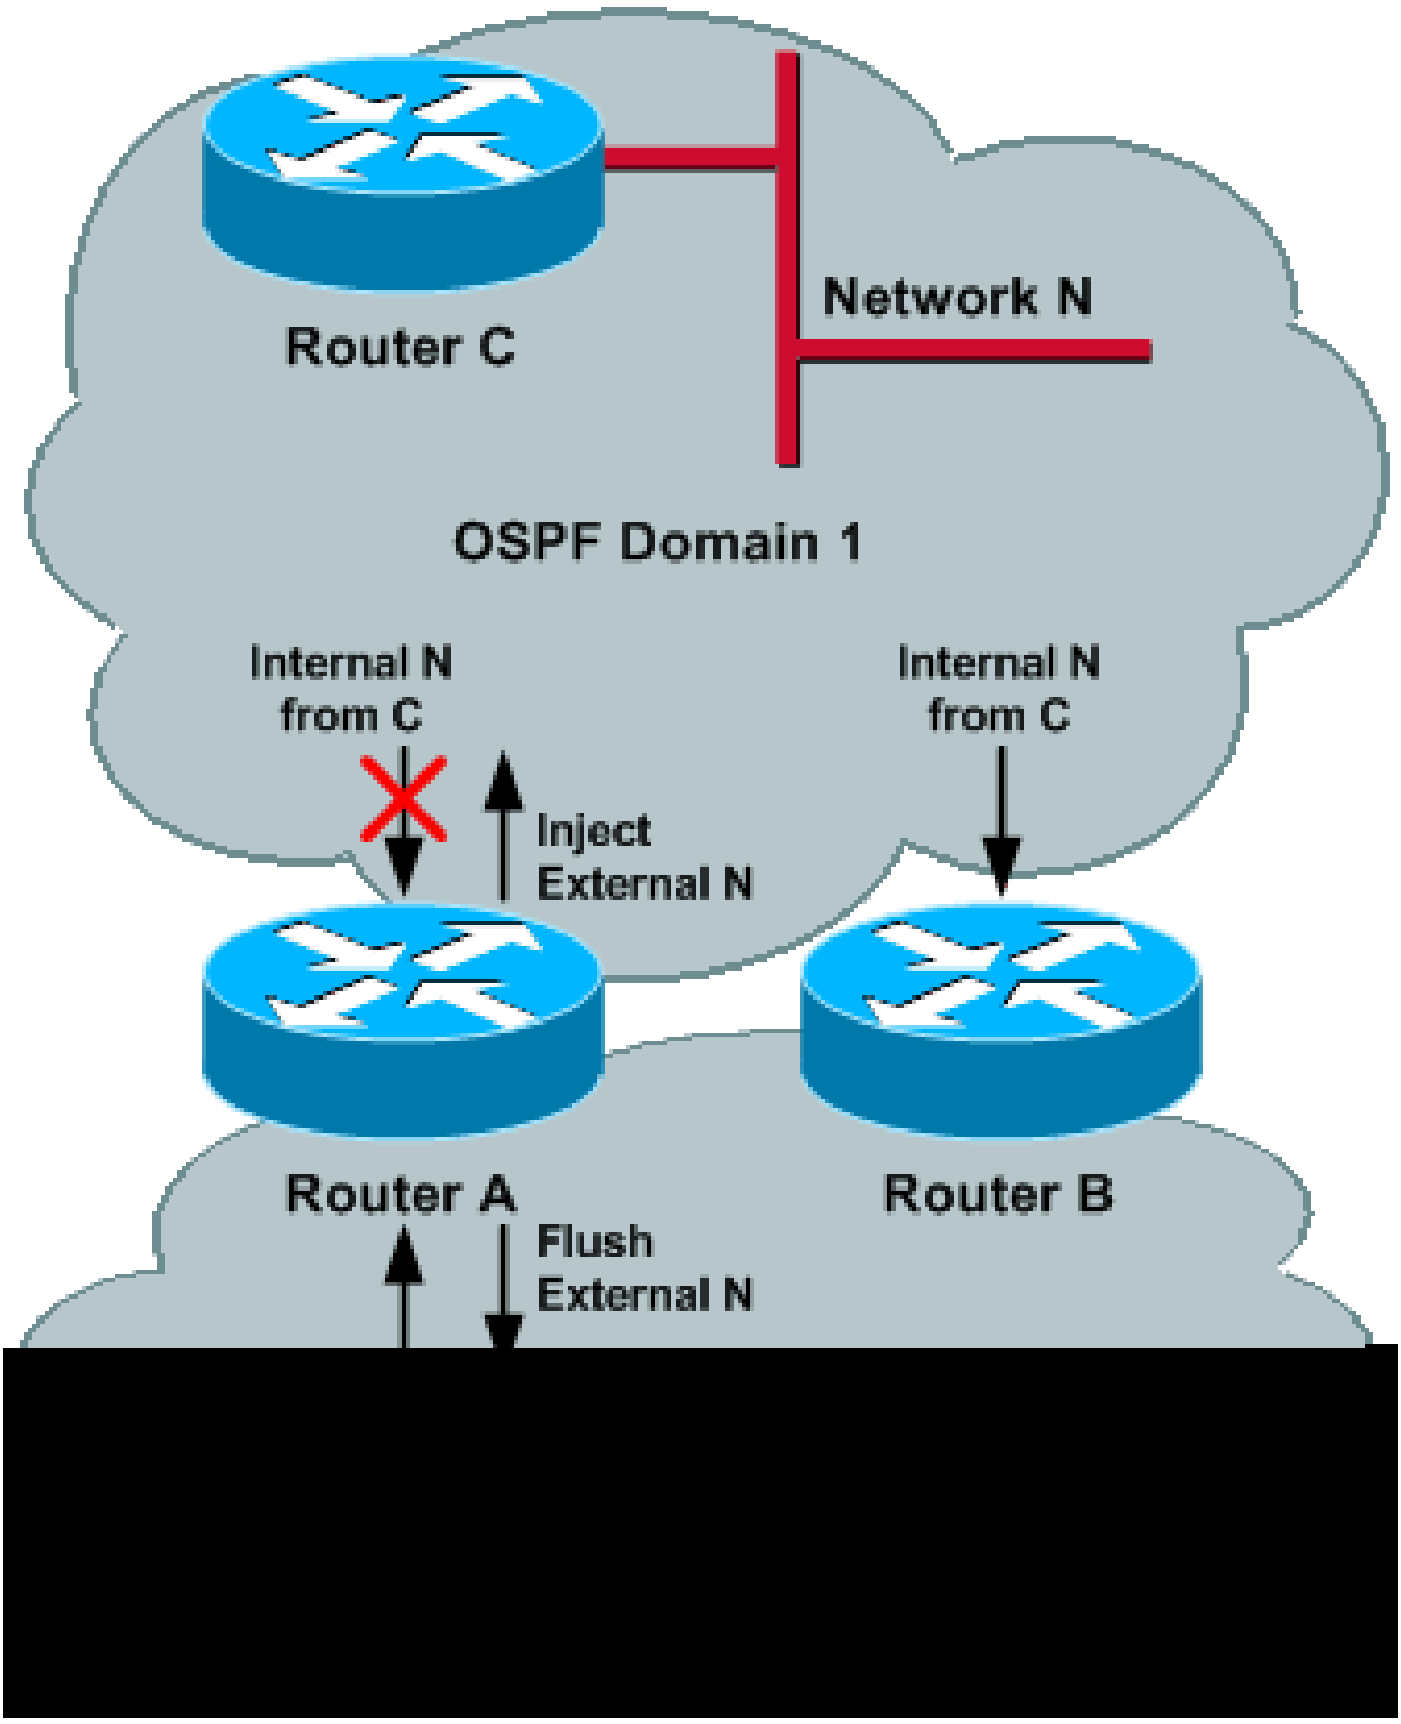
If there are prefixes external to both domains (that come from some other redistribution points), the same problem still occurs because the administrative distance for those external routes is the same in both processes. If you make the administrative distance for external processes different, it does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

Image 4


ASBR
Network N
External
Router C
External N
from C
External N
from C
OSPF Domain 1
Router A
Router B
OSPF Domain 2
External M
from D
External M
from D
ASBR
Network M
External
Router D

When Router A had connectivity to Network N, it used process 1 because of the better administrative distance, and process 2 was kept for backup information. Once the path through process 1 becomes unreachable, process 2 is used for connectivity.

Image 5



Now, all routers in Domain 2 use Router B to reach Network N; and Router A (or the part of Domain 1 which has lost connectivity to Network N through Domain 1) uses Domain 2 for connectivity to Network N. This scenario remains valid if Router B had lost connectivity to Network N, instead of Router A.

If Router A and Router B both lose connectivity to Network N (for example, if Router C goes down), then

this sequence of events could occur:

1. Before Network N becomes unreachable, Router A and Router B learned Network N though process 1 and redistributed it into process 2 as external.

2. Router A and Router B detect (almost at the same time) that Network N is unreachable through Domain 1; therefore, they flush their previously external N in Domain 2.

3. Before Router A (Router B) receives the flushed LSA from Router B (Router A), it installs the external N through Domain 2 (a higher administrative distance) as the backup route.

4. Since Router A (Router B) installed N through process 2, it generates an external N into Domain 1.

5. Router A (Router B) receives the flushed LSA (event 1) from Router B (Router A). It removes 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 installs the external Network N through Domain 1, because N has been flushed though Domain 2.

7. Since Router A (Router B) installed Network N through process 1, it generates 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 redistribution back to the same domain in order to avoid routing loops.

Use 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 occurs 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 does not redistribute those routes back to Domain 1.

In order to do this, enter 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.

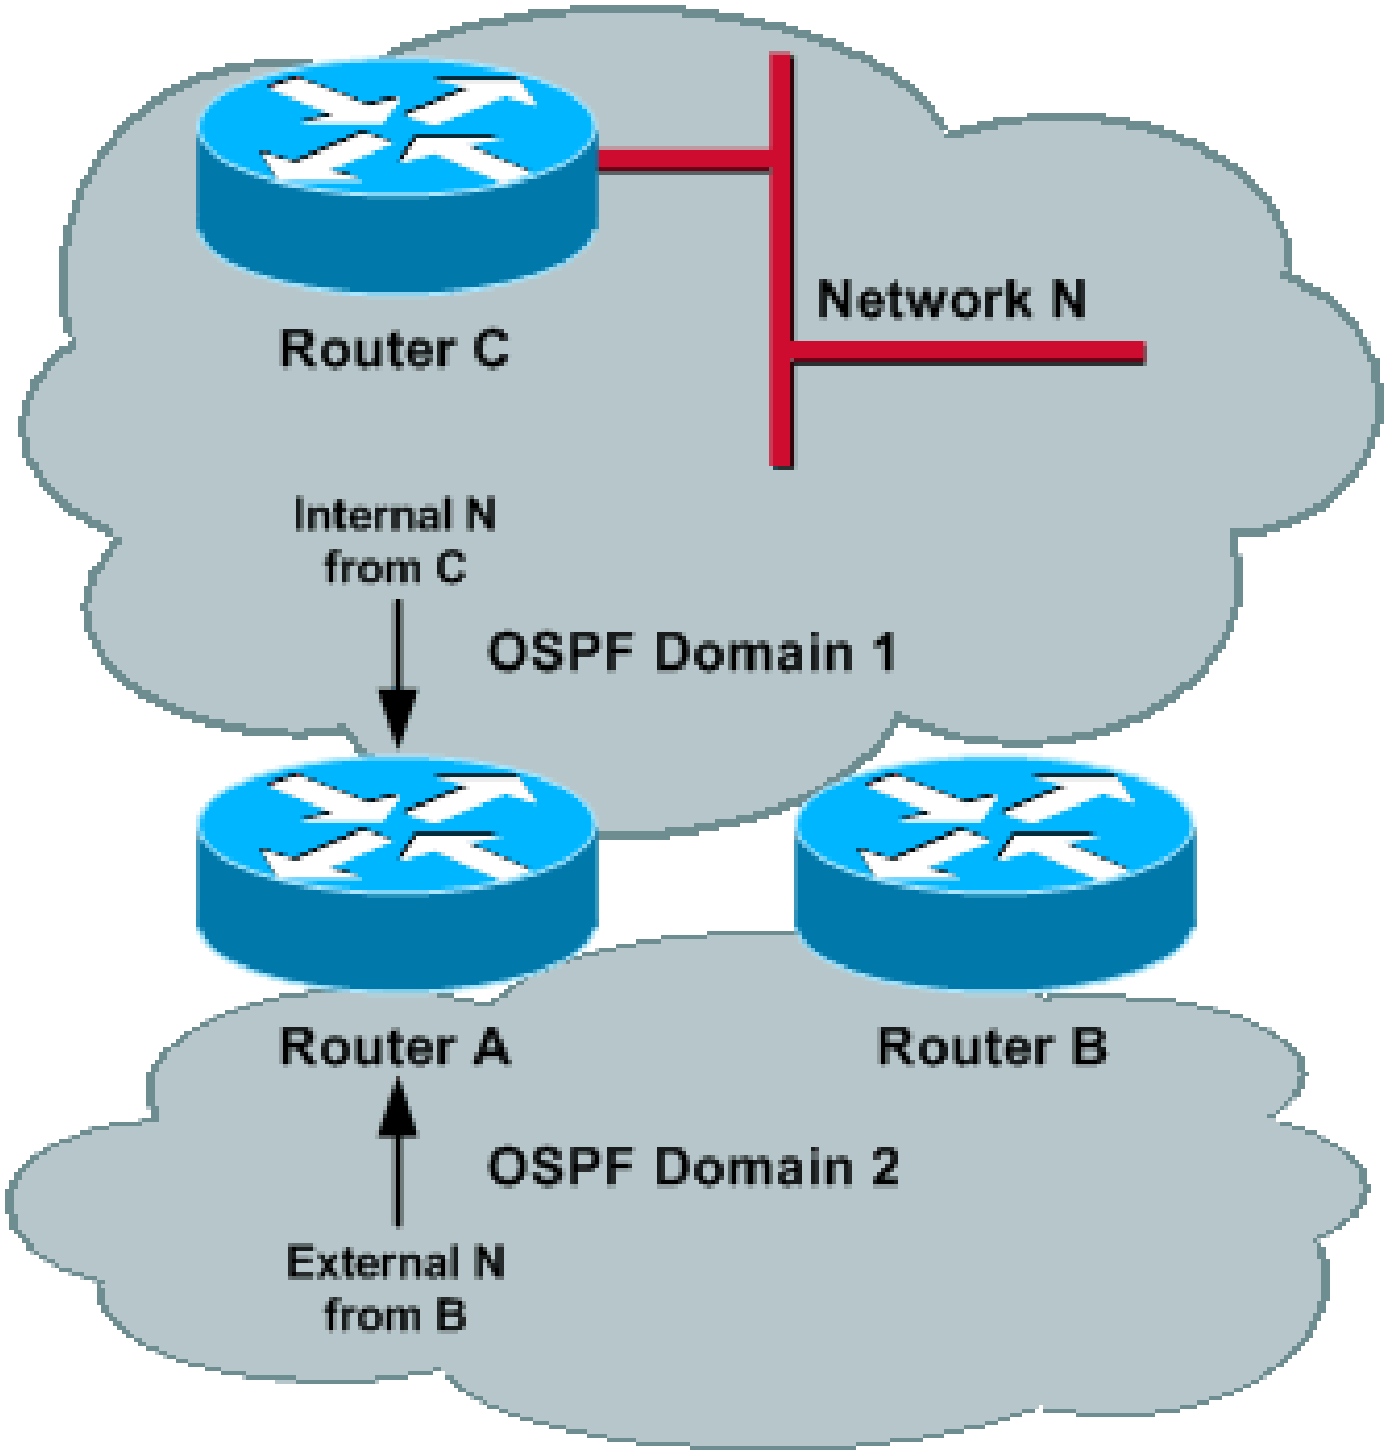
Note: The **distance 255** command gives a distance of 255 to those routes and, therefore, prevents their installation into the routing table.

In Image 6, Router A uses the **access-list 1** command in order to match all the routes in Domain 1 and uses the **distance 255** command under process 1 in order 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 does not install those routes and, therefore, does not redistribute back into Domain 1 again.

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

Image 6



Router A Configuration	Router B Configuration

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

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.

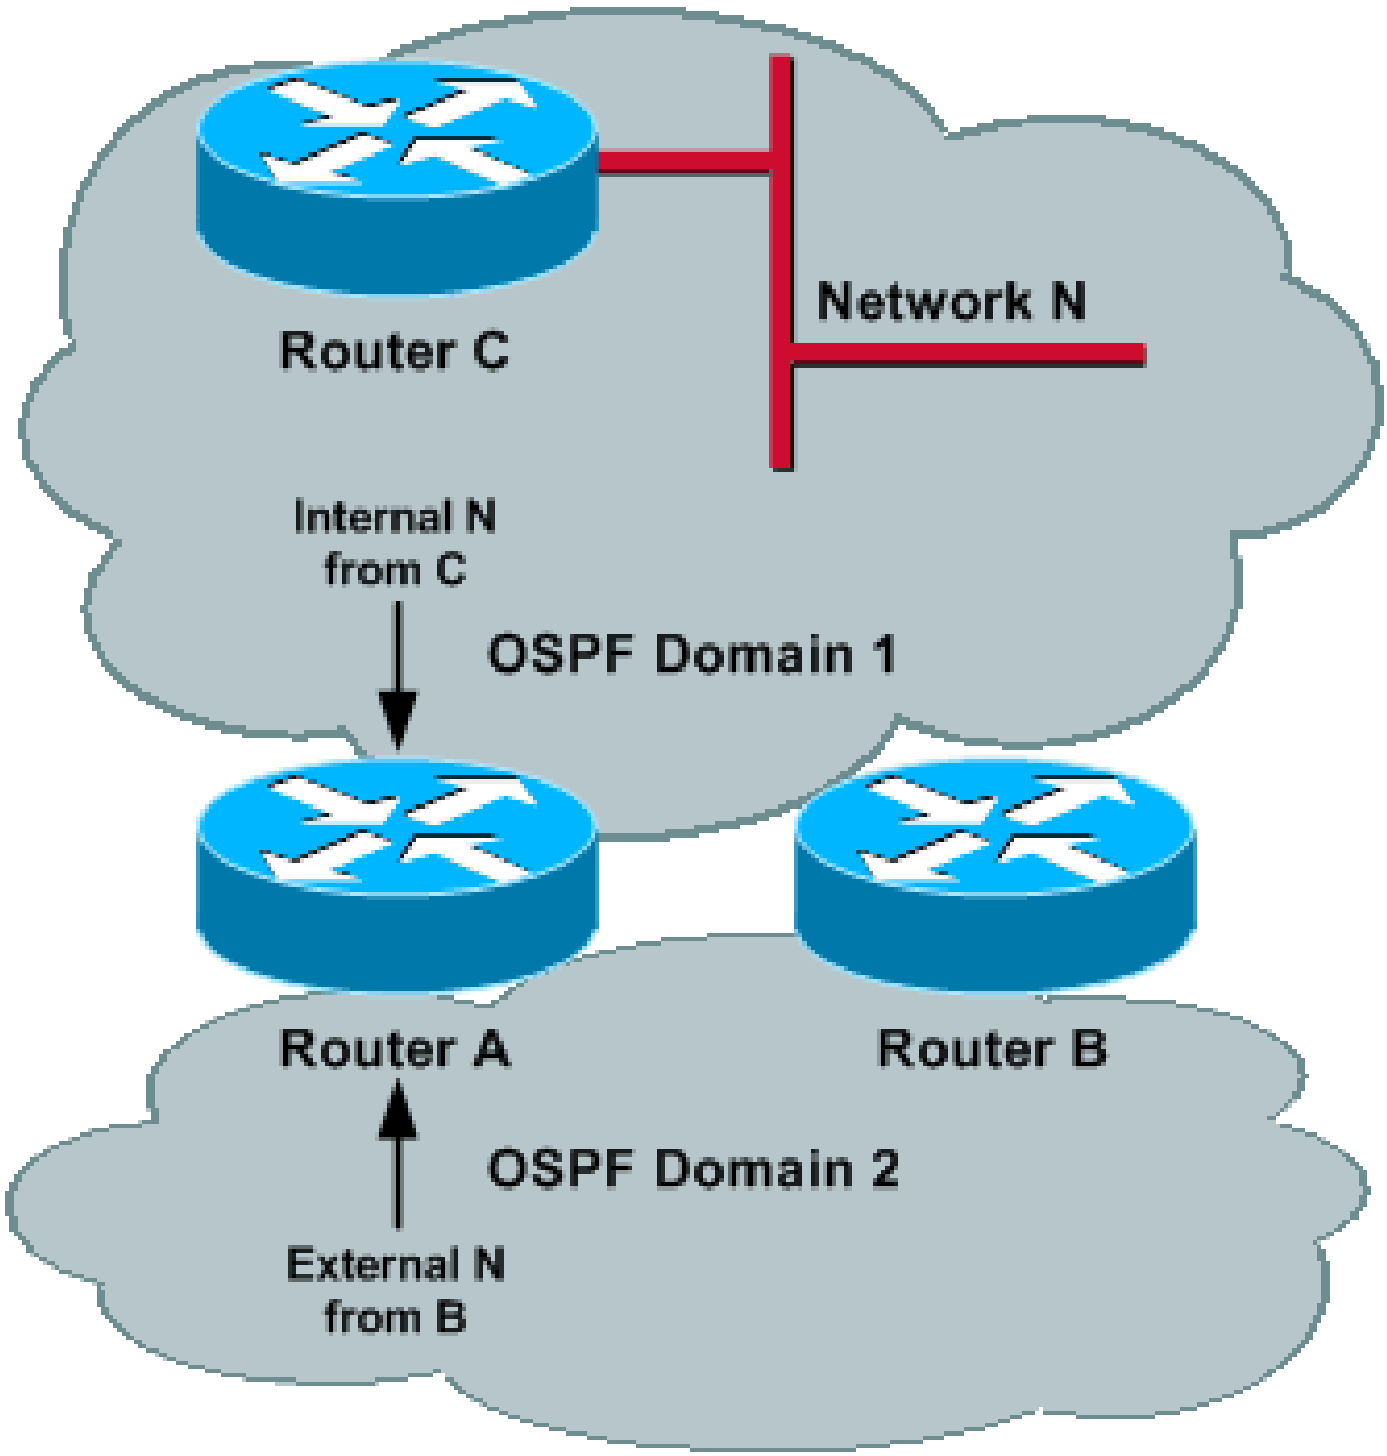
This configuration works properly in cases where both routers lose connectivity to the network (as described in Network Operation Without Network Failure and Network Operation With Network Failure). However, because the prefixes are denied from the routing table, the domains cannot back up each other.

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

Filter 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 are not be installed into the routing table, they are not redistributed back into the same domain.

Image 7



Routers A and B Configuration
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
```

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.

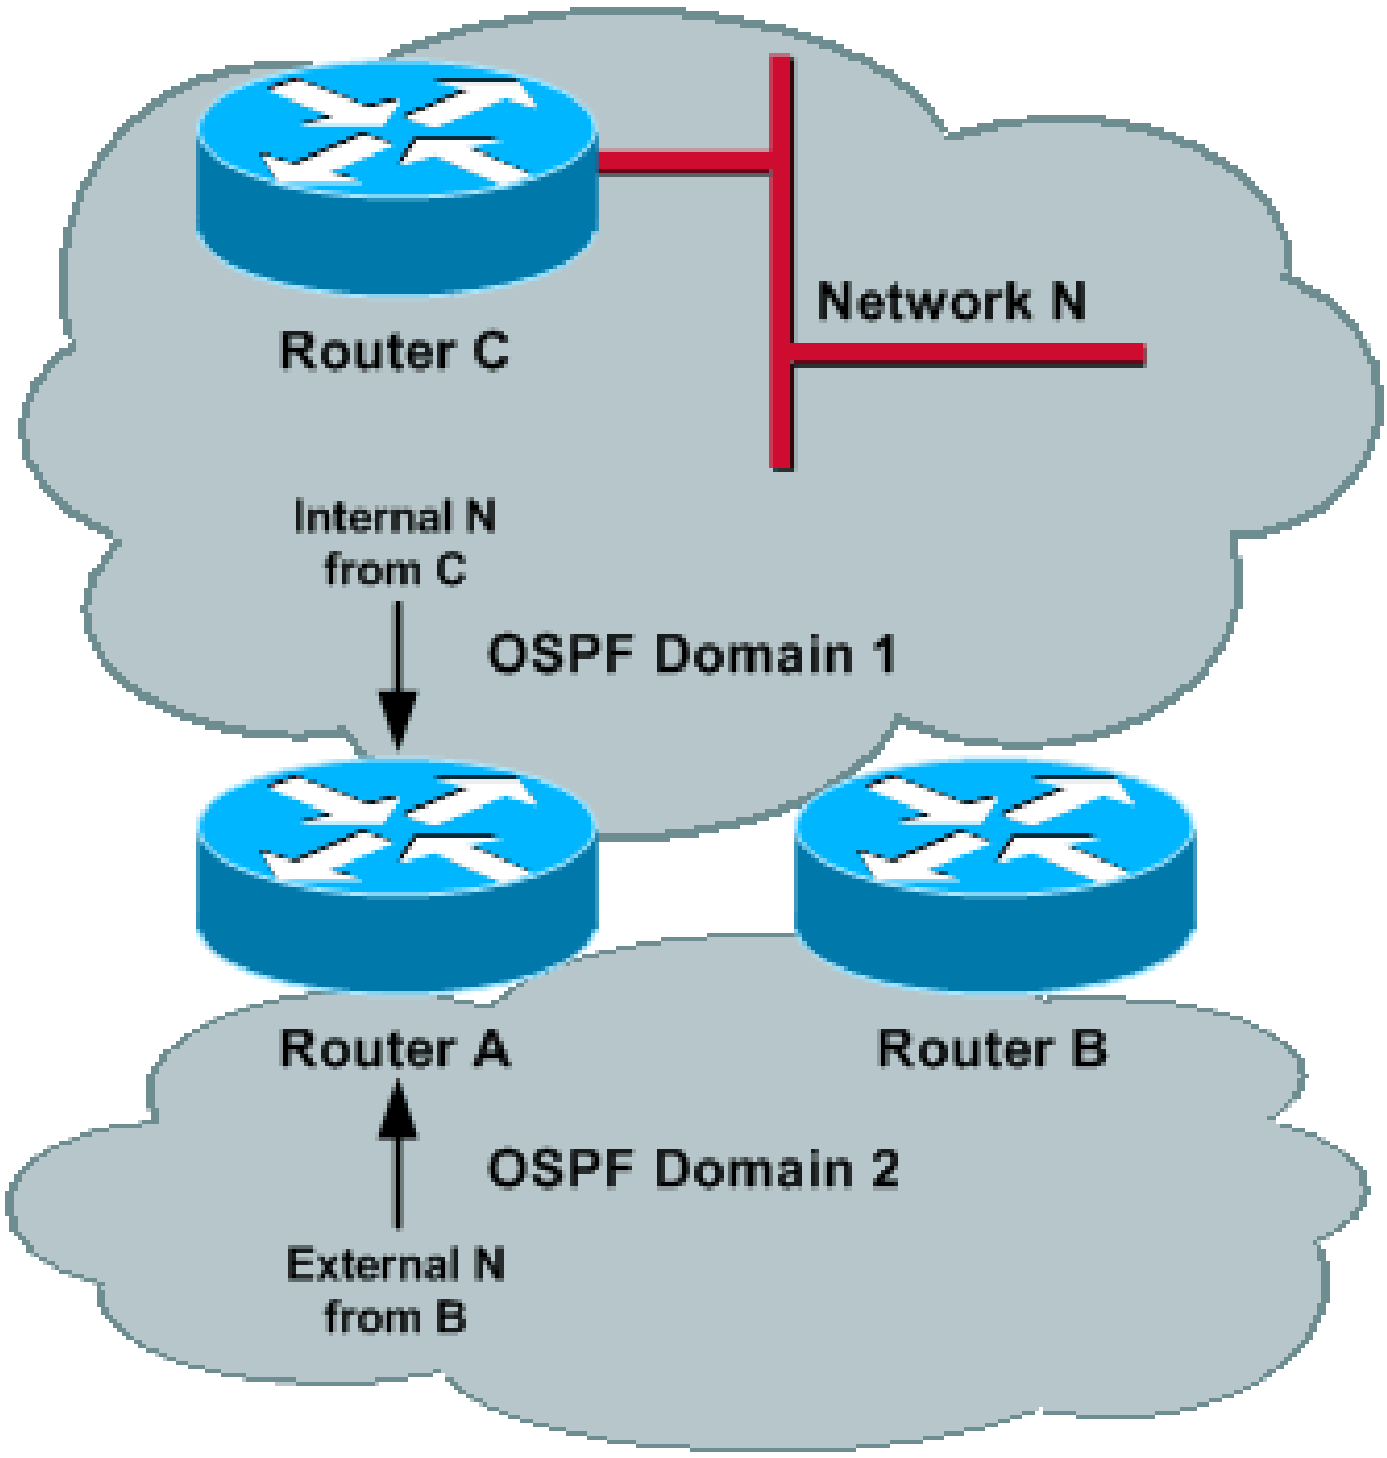
Note: The previous **distance ospf external 200** command is no longer needed because the route learned from the remote router via one of the processes is not installed.

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.

Use 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.

Image 8



Routers A and B Configuration
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 !

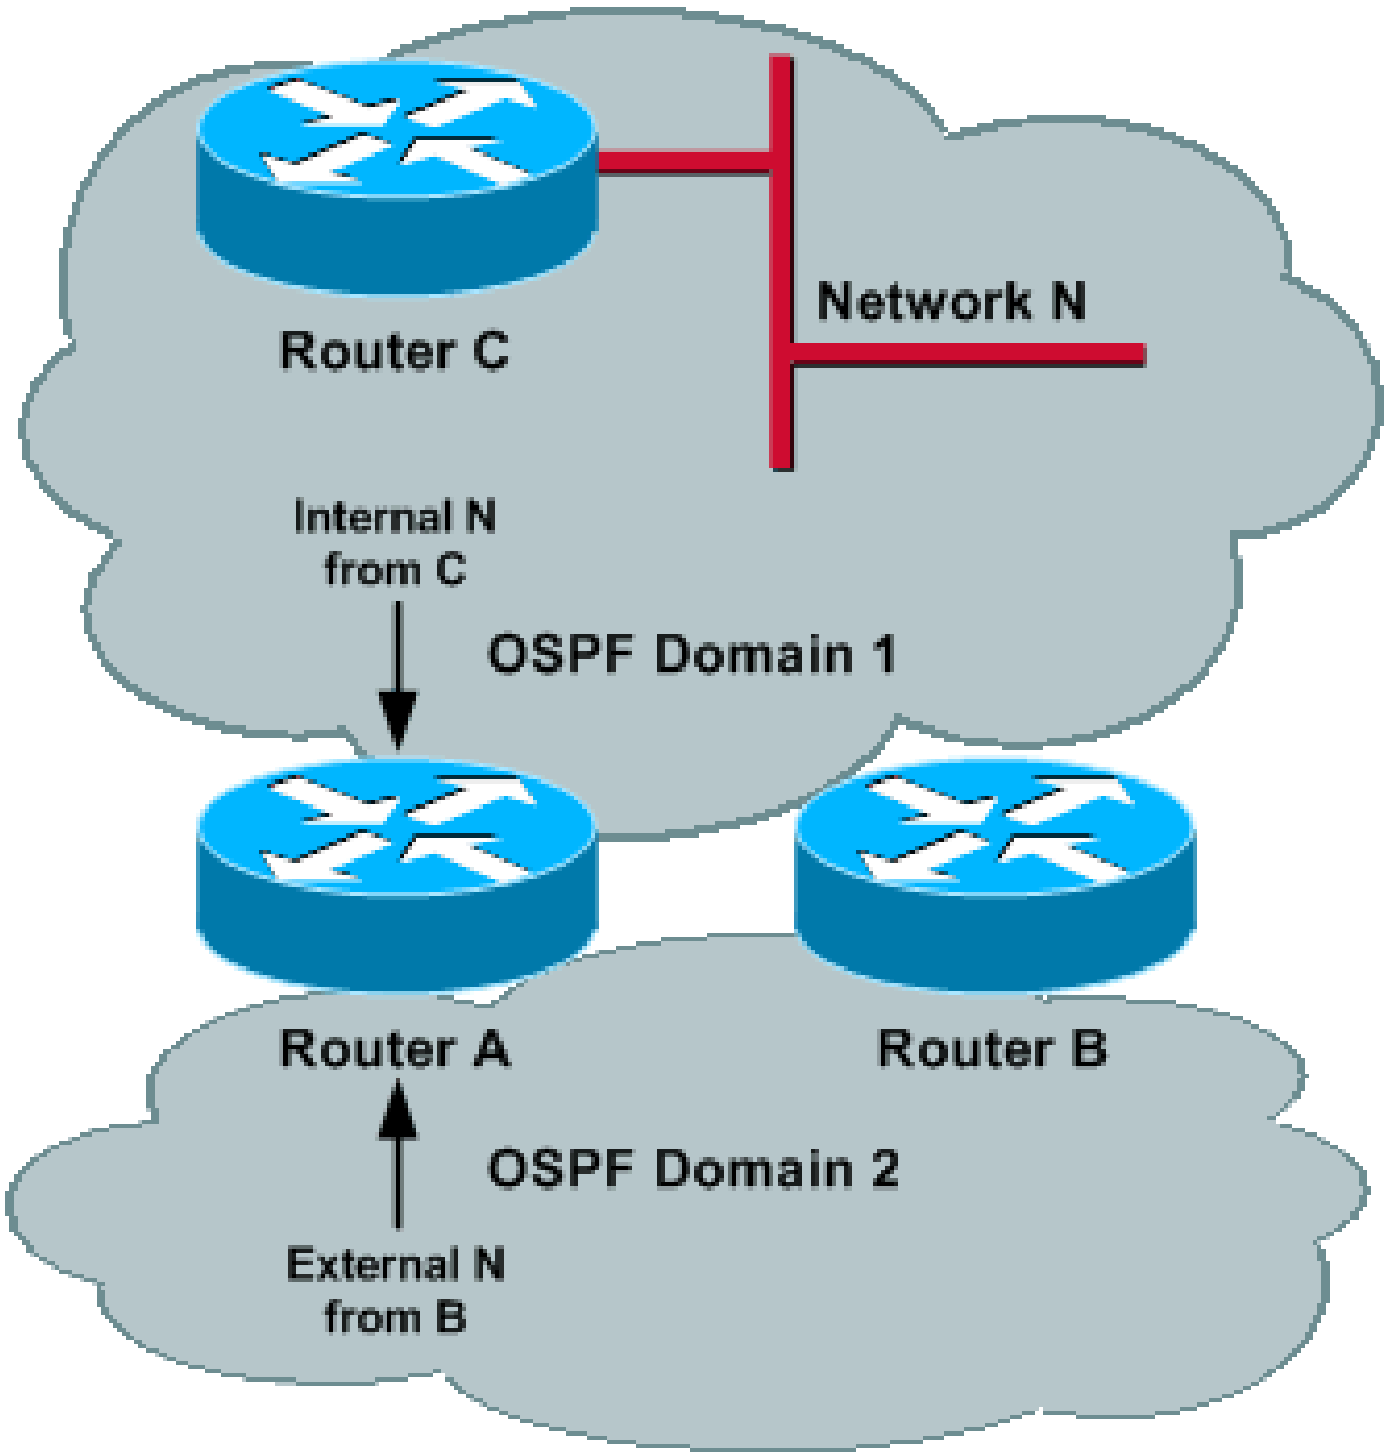
This configuration works properly in cases where both routers lose connectivity to the network (as described in Network Operation Without Network Failure and Network Operation With Network Failure). One domain could back up the other domain.

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

Image 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.

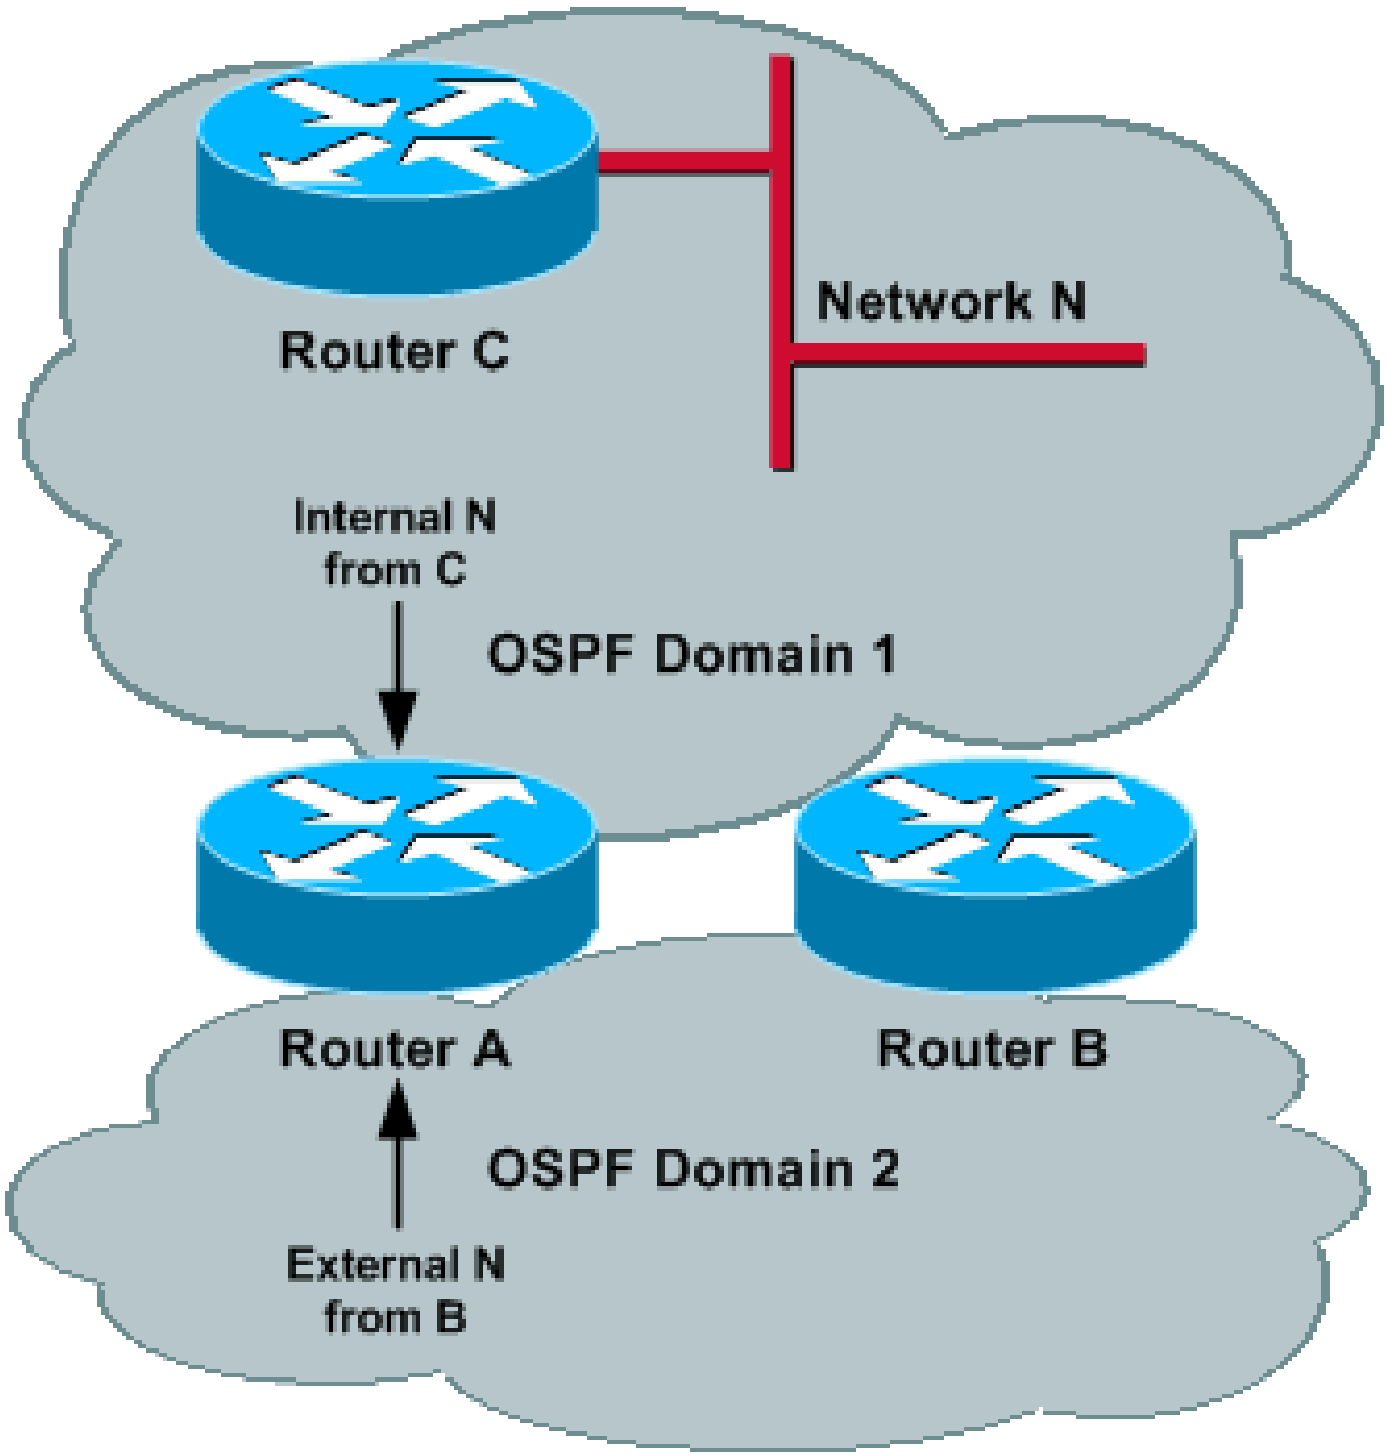
```
router ospf 2
redistribute ospf 1 subnet route-map filter_domain1
distance ospf external 200
!
route-map filter_domain1 permit 20
match ip address 2
!
access-list 2

!--- 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.

Image 10



Routers A and B Configuration
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

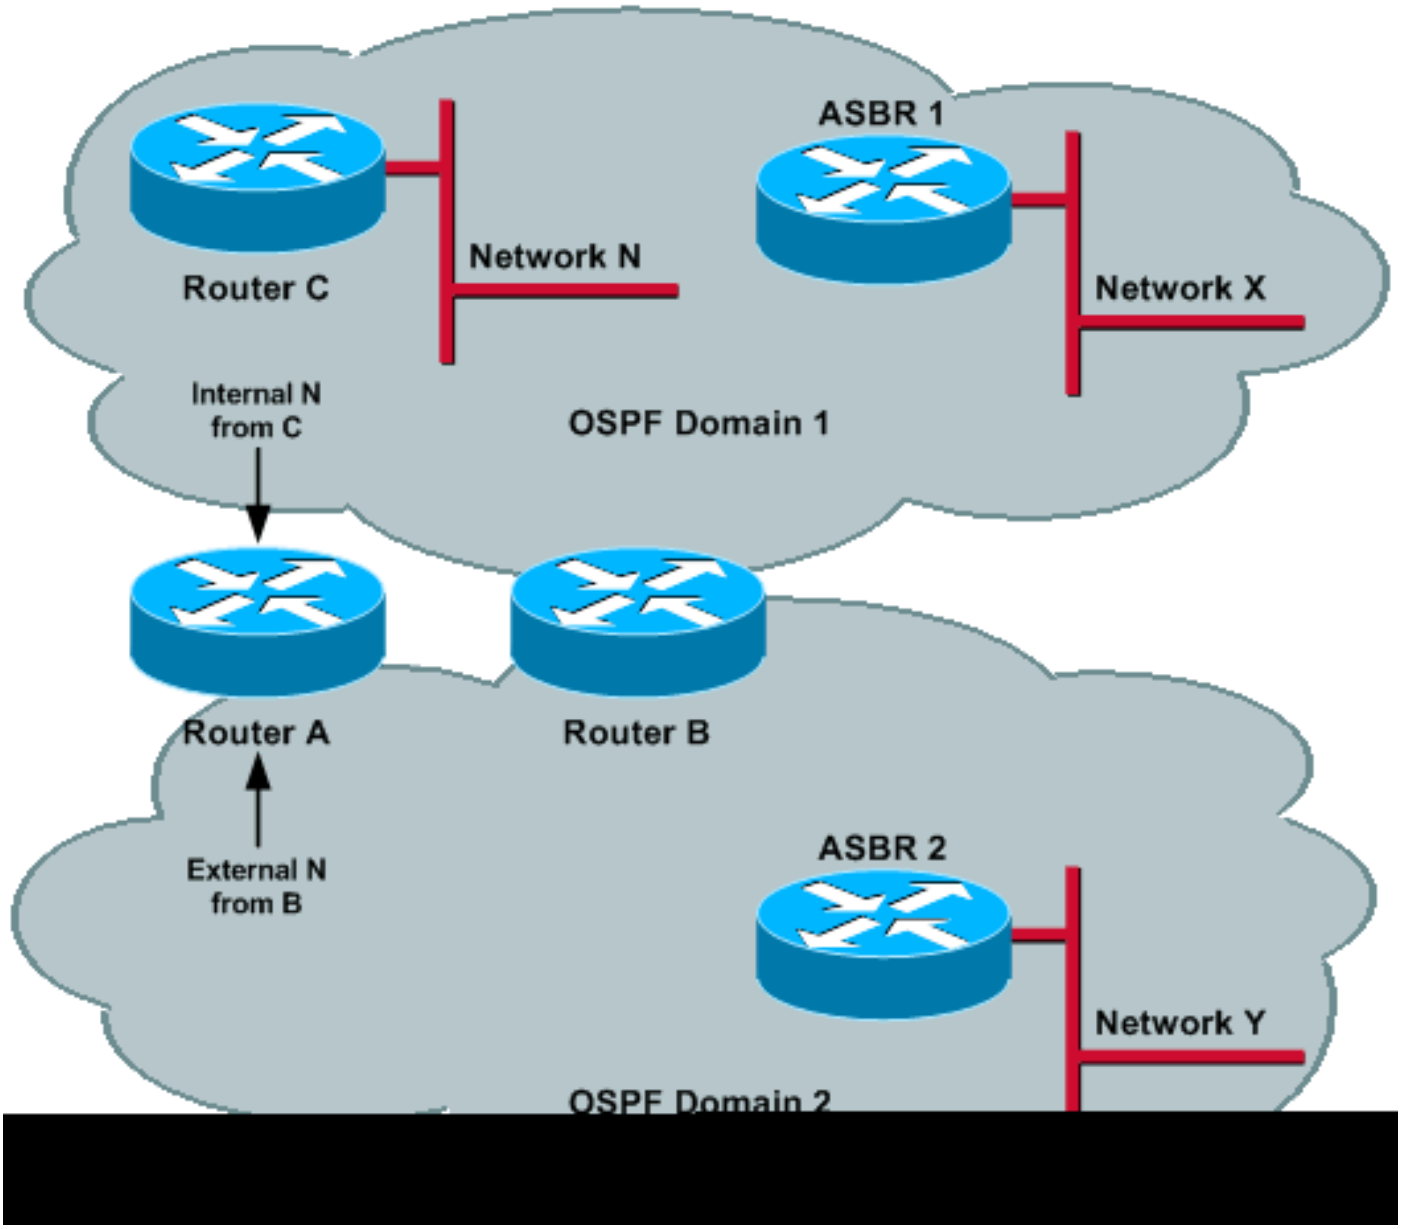
```
!
route-map filter_domain1 deny 10
match tag 1
route-map filter_domain1 permit 20
```

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 redistribute Networks X and Y into Domain 1 and Domain 2, respectively.

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

Image 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 2
!
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 1
!
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. In order to prevent routing loops, prefixes that belong to a domain cannot be redistributed back to the same domain. Also, the administrative distances of the OSPF processes can 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 cannot back up each other.

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

- The match internal solution does not allow you to have control over prefixes, and all external prefixes are blocked from redistribution. In other words, if there are external prefixes from other ASBRs, then those LSAs are not 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 can be noted that backup means that, can 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) in order 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 can introduce other problems.

Related Information

- OSPF Support Page
- Technical Support & Documentation - Cisco Systems