

OSPF虛擬鏈路：傳輸功能

目錄

[簡介](#)

[必要條件](#)

[背景資訊](#)

[案例 1](#)

[網路圖表:](#)

[RFC 2328第16.2節摘要](#)

[案例 2](#)

[網路圖表:](#)

[RFC 2328第6節摘要](#)

[RFC 2328第16.1節摘要](#)

[RFC 2328第16.1節摘要](#)

[RFC 2328第16.3節摘要](#)

簡介

本文的目的在於說明當V位元（虛擬連結位元）位於非主幹區域時，開放最短路徑優先(OSPF)的行為。僅當路由器是一個或多個完全相鄰虛擬鏈路的端點時，才会在第1類LSA中傳送V位訊號。設定V位後，可能會更改區域內和區域間路由之間的路徑計算優先順序。

必要條件

使用本文檔時，請參閱圖1中的網路圖：

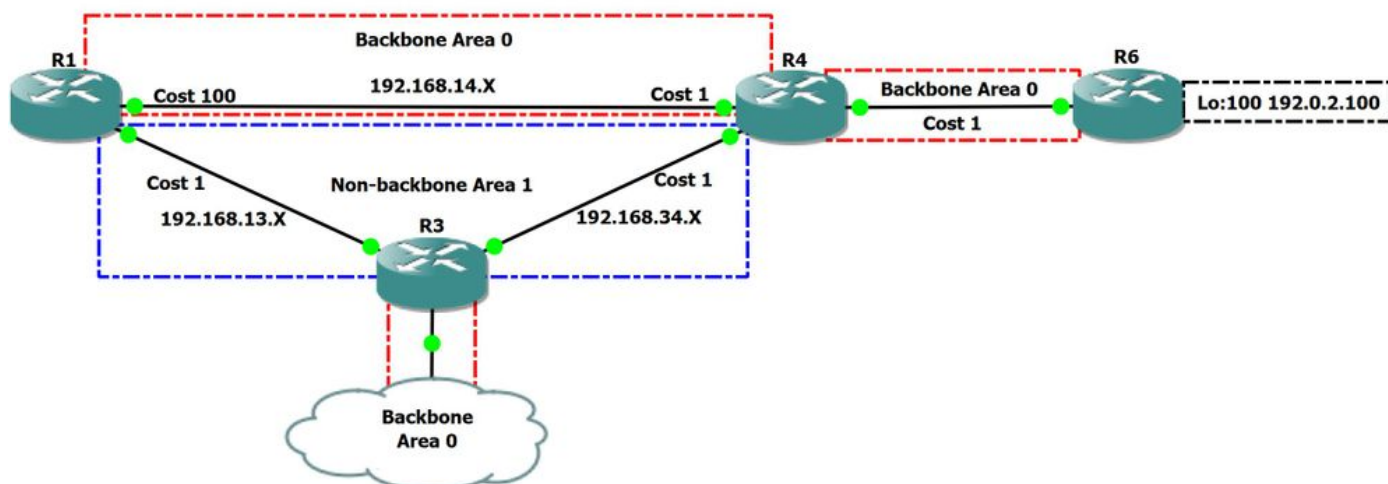


圖1

在上面的網路圖中，主幹區域0和非主幹區域1都存在。R1是連線區域0和區域1的區域邊界路由器(ABR),R4和R3在此網路中扮演著相似的角色。在此拓撲中，區域0不連續，因為R3和R4未通過區域0連線。

背景資訊

OSPF自治系統中的所有區域必須連線到主幹區域（區域0）。在某些情況下，在骨幹區域之間有一個非骨幹區域，這可能會導致自治系統的某些區域變得不可達，從而導致網路不連續。當不可能有連續的主幹區域時，您可以使用虛擬鏈路通過非主幹區域連線主幹。配置虛擬鏈路所經過的區域稱為中轉區域。

案例 1

網路圖表:

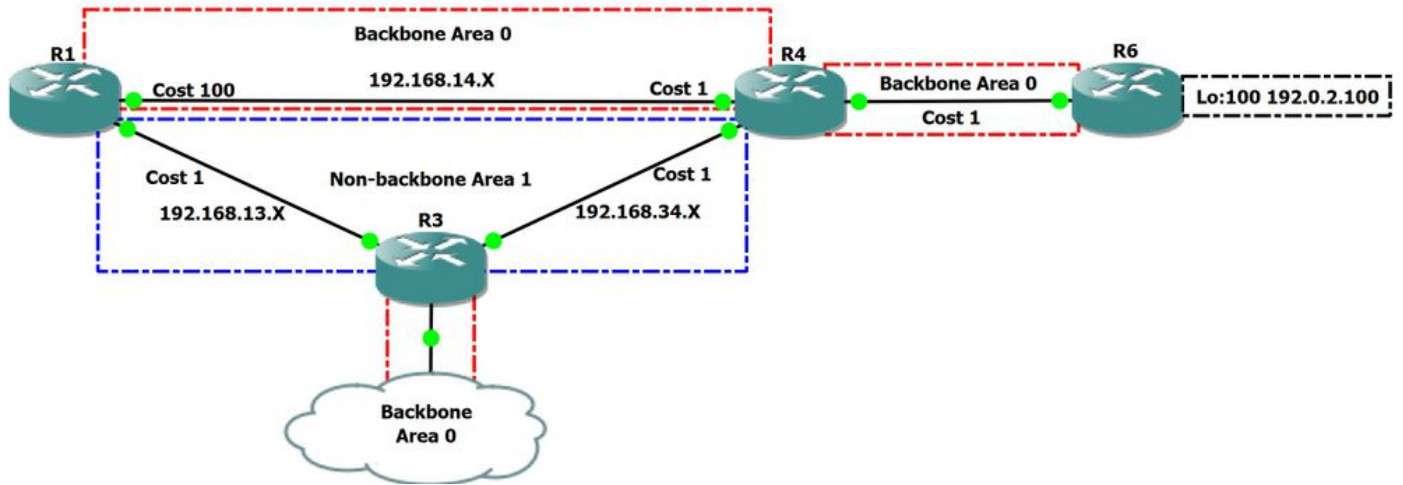


圖2

在此案例中，我們將檢查上述網路拓撲中的預期路徑計算。我們將調查從R1路由到R6環回100時，首選哪條路徑(IP地址為192.0.2.100/32)

下面我們來看一下R1上的OSPF資料庫，進一步瞭解拓撲：

```
R1#show ip ospf database
```

```
OSPF Router with ID (1.1.1.1) (Process ID 1)
```

```
Router Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	22	0x8000000C	0x00CD7A	2
4.4.4.4	4.4.4.4	289	0x8000000F	0x00434E	4
6.6.6.6	6.6.6.6	374	0x80000009	0x00630A	3

```
Summary Net Link States (Area 0)
```

Link ID	ADV Router	Age	Seq#	Checksum
192.168.13.0	1.1.1.1	18	0x80000001	0x00348D
192.168.13.0	4.4.4.4	207	0x80000001	0x00E3D0
192.168.34.0	1.1.1.1	8	0x80000001	0x005655
192.168.34.0	4.4.4.4	683	0x80000001	0x00F1AE

```
Router Link States (Area 1)
```

Link ID	ADV Router	Age	Seq#	Checksum	Link count
1.1.1.1	1.1.1.1	17	0x80000009	0x00EC2B	2
3.3.3.3	3.3.3.3	18	0x8000000E	0x005A64	4
4.4.4.4	4.4.4.4	544	0x80000005	0x0007CF	2

Summary Net Link States (Area 1)

Link ID	ADV Router	Age	Seq#	Checksum	
155.1.37.0	3.3.3.3	1558	0x80000004	0x00A7C3	
192.0.2.100	1.1.1.1	23	0x80000001	0x009F0C	<- R6 Loopback
192.0.2.100	4.4.4.4	370	0x80000001	0x0059AA	<- R6 Loopback
192.168.14.0	1.1.1.1	23	0x80000001	0x000B52	
192.168.14.0	4.4.4.4	331	0x80000001	0x00CEE5	
192.168.34.0	1.1.1.1	3608	0x80000002	0x00406C	
192.168.46.0	1.1.1.1	23	0x80000001	0x00B388	
192.168.46.0	4.4.4.4	484	0x80000001	0x006D27	

從上面的輸出我們可以看到，R1通過R4獲知R6 Lo100:192.0.2.100作為第3類彙總LSA，而R1也通過區域內得知了R6 Lo100:192.0.2.100，因此它自己也生成了第3類彙總LSA 骨幹。在下面的輸出中，我們可以看到R6直接連線了192.0.2.100。

```
R1#show ip ospf da router 6.6.6.6
```

```
OSPF Router with ID (1.1.1.1) (Process ID 1)
```

```
Router Link States (Area 0)
```

```
LS age: 614
```

```
Options: (No TOS-capability, DC)
```

```
LS Type: Router Links
```

```
Link State ID: 6.6.6.6
```

```
Advertising Router: 6.6.6.6
```

```
LS Seq Number: 8000000D
```

```
Checksum: 0x5B0E
```

```
Length: 60
```

```
Number of Links: 3
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 192.0.2.100 <-- Loopback 100 directly connected
```

```
(Link Data) Network Mask: 255.255.255.255
```

```
Number of MTID metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: another Router (point-to-point)
```

```
(Link ID) Neighboring Router ID: 4.4.4.4
```

```
(Link Data) Router Interface address: 192.168.46.6
```

```
Number of MTID metrics: 0
```

```
TOS 0 Metrics: 1
```

```
Link connected to: a Stub Network
```

```
(Link ID) Network/subnet number: 192.168.46.0
```

```
(Link Data) Network Mask: 255.255.255.0
```

```
Number of MTID metrics: 0
```

```
TOS 0 Metrics: 1
```

RFC 2328第16.2節摘要

16.2. Calculating the inter-area routes

- (5) Next, look up the routing table entry for the destination N.
(If N is an AS boundary router, look up the "router" routing table entry associated with Area A). If no entry exists for

N or if the entry's path type is "type 1 external" or "type 2 external", then install the inter-area path to N, with associated area Area A, cost IAC, next hop equal to the list of next hops to router BR, and Advertising router equal to BR.

- (6) Else, if the paths present in the table are intra-area paths, do nothing with the LSA (**intra-area paths are always preferred**).
- (7) **Else, the paths present in the routing table are also inter-area paths. Install the new path through BR if it is cheaper**, overriding the paths in the routing table. Otherwise, if the new path is the same cost, add it to the list of paths that appear in the routing table entry.

在上面的輸出中，我們可以看到區域內路由優先於區域間路由。因此，在我們的場景中，根據RFC 2328,R1應首選通過區域內主幹路由。

檢查在拓撲中是否觀察到此行為：

```
R1#show ip ospf rib 192.0.2.100
```

```
OSPF Router with ID (1.1.1.1) (Process ID 1)
```

```
Base Topology (MTID 0)
```

```
OSPF local RIB
```

```
Codes: * - Best, > - Installed in global RIB
```

```
LSA: type/LSID/originator
```

```
*> 192.0.2.100/32, Intra, cost 102, area 0
```

```
SPF Instance 9, age 02:19:34
```

```
Flags: RIB, HiPrio
```

```
via 192.168.14.4, GigabitEthernet3 label 1048578
```

```
Flags: RIB
```

```
LSA: 1/6.6.6.6/6.6.6.6
```

```
R1#show ip route 192.0.2.100
```

```
Routing entry for 192.0.2.100/32
```

```
Known via "ospf 1", distance 110, metric 102, type intra area
```

```
Last update from 192.168.14.4 on GigabitEthernet3, 02:26:29 ago
```

```
Routing Descriptor Blocks:
```

```
* 192.168.14.4, from 6.6.6.6, 02:26:29 ago, via GigabitEthernet3
```

```
Route metric is 102, traffic share count is 1
```

從上面的輸出中您可以看到，我們更傾向於通過主幹區域0向R6 loopback100遷移。在我們的鏈路狀態資料庫中，我們還知道通過R3（然後是R4）的區域間路徑。通過R4以2的開銷獲取的彙總LSA如下所示：

```
R1#show ip ospf database summary 192.0.2.100
```

```
OSPF Router with ID (1.1.1.1) (Process ID 1)
```

```
Summary Net Link States (Area 1)
```

```
LS age: 523
```

```
Options: (No TOS-capability, DC, Upward)
```

```
LS Type: Summary Links(Network)
```

```
Link State ID: 192.0.2.100 (summary Network Number)
```

```
Advertising Router: 1.1.1.1
```

```
LS Seq Number: 80000005
```

```
Checksum: 0x9710
Length: 28
Network Mask: /32
      MTID: 0      Metric: 102
```

```
LS age: 973
Options: (No TOS-capability, DC, Upward)
LS Type: Summary Links(Network)
Link State ID: 192.0.2.100 (summary Network Number)
Advertising Router: 4.4.4.4
LS Seq Number: 80000005
Checksum: 0x51AE
Length: 28
Network Mask: /32
      MTID: 0      Metric: 2
```

<- This is Type-3 LSA injected by ABR R4

請注意，2的開銷反映了ABR對目的地字首的開銷。第3類LSA從區域0泛洪到非主幹區域，反之亦然，它描述了ABR對其他區域鏈路的可達性。它包括從注射第3類LSA的ABR角度計算的成本，但對接收第3類LSA的路由器隱藏全部成本。

從上面的輸出中，我們現在知道我們有兩條路徑可用來從R1到達R6環回：

1. 區域內費用102美元
2. 通過Type-3 LSA + R1到R4的開銷為2的區域間，也就是2。這讓我們總開銷為4

在此案例中，我們已觀察到我們偏好成本更高的區域內路徑，因為RFC 2328中定義區域內路徑優先於區域間路徑。

進行場景2之前，以下是OSPF如何解釋第3類LSA的示例：

- ABR R4可以到達A區域內鏈路，開銷為X
- R1可到達ABR R4，開銷為Y
- 意味著R1可以通過SPT到達鏈路A，開銷為X + Y




圖3

這就是區域間路由通常與距離向量協定相比的原因，因為區域間資訊是隱藏的。由於區域間OSPF是距離向量，因此容易受到路由環路的影響。它通過強制無環的區域間拓撲來避免環路，在該拓撲中，一個區域的流量只能通過區域0到達另一個區域。

案例 2

網路圖表:

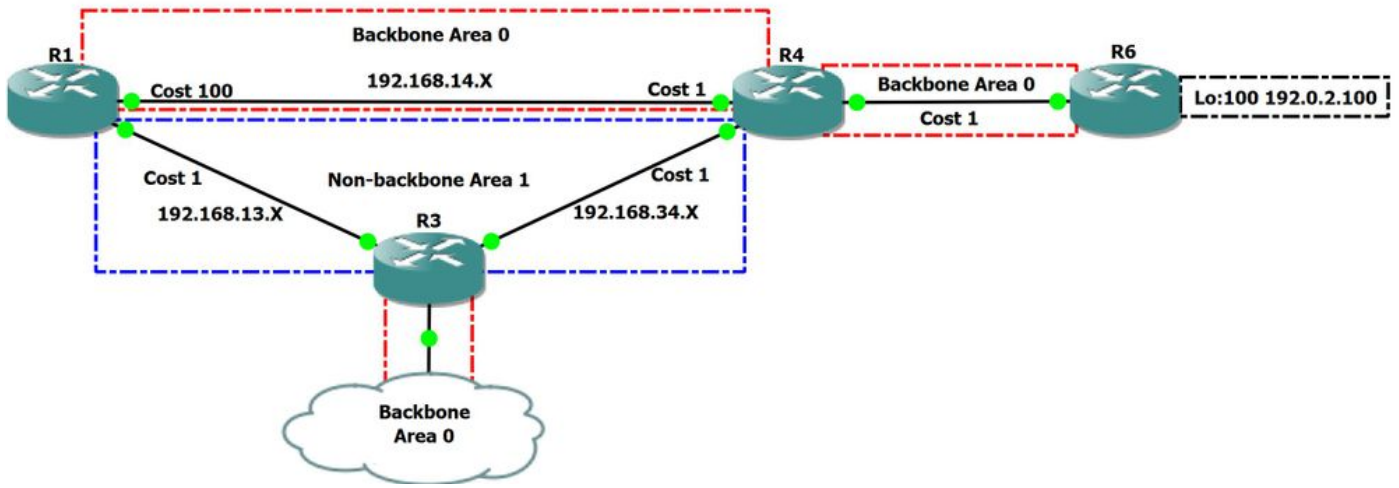


圖4

在此案例中，我們在R3和R4上設定V位元，如此一來，就可以在非主幹區域1的型別1 LSA中存在此位元時檢查路徑優先順序。

RFC 2328第6節摘要

6. The Area Data Structure

TransitCapability

This parameter indicates whether the area can carry data traffic that neither originates nor terminates in the area itself. This parameter is calculated when the area's shortest-path tree is built (see Section 16.1, where TransitCapability is set to TRUE if and only if there are one or more fully adjacent virtual links using the area as Transit area), and is used as an input to a subsequent step of the routing table build process (see Section 16.3). When an area's TransitCapability is set to TRUE, the area is said to be a "transit area".

RFC 2328第16.1節摘要

16.1 Calculating the shortest-path tree for an area

- (2) Call the vertex just added to the tree vertex V. Examine the LSA associated with vertex V. This is a lookup in the Area A's link state database based on the Vertex ID. **If this is a router-LSA, and bit V of the router-LSA (see Section A.4.2) is set, set Area A's TransitCapability to TRUE.** In any case, each link described by the LSA gives the cost to an adjacent vertex. For each described link, (say it joins vertex V to vertex W):

從RFC中的上述語句可以看出，當在router-LSA中設定V-bit時，我們知道該位元被設定為能夠傳輸的區域，或者換句話說，當運行Dijkstra演算法時，該區域的TransitCapability為true。

一旦我們知道，如果設定了V位，則可以考慮某個區域進行功能傳輸，我們必須檢查是否已配置此功能：預設情況下，OSPF區域傳輸功能已啟用。

```
R1#show run all | sec ospf
router ospf 1
capability opaque
capability lls
capability transit
```

要在區域1中設定V位，我們將建立一個從R3到R4的虛擬鏈路。當虛擬鏈路啟動時，我們應看到在第1類LSA中設定V位。

```
R3(config)#router ospf 1
R3(config-router)#area 1 virtual-link 4.4.4.4
```

```
R3#show ip ospf interface brief
```

Interface	PID	Area	IP Address/Mask	Cost	State	Nbrs	F/C
VL0	1	0	192.168.34.3/24	1	P2P	1/1	<-- Here we have
Virtual-link present and 1 neighborhood over VLO							
Gi3	1	0	192.168.80.3/24	1	DR	0/0	
Gi2	1	1	192.168.13.3/24	1	P2P	1/1	
Gi1	1	1	192.168.34.3/24	1	P2P	1/1	

```
R3#
```

現在讓我們檢查R3區域1的型別1 LSA。

```
R3#show ip ospf 1 1 database router 3.3.3.3 OSPF Router with ID (3.3.3.3) (Process ID 1) Router Link States (Area 1)
LS age: 189 Options: (No TOS-capability, DC) LS Type: Router Links Link State ID: 3.3.3.3 Advertising Router:
3.3.3.3 LS Seq Number: 80000018 Checksum: 0x525E Length: 72 Area Border Router Virtual Link Endpoint <- V-bit
set
```

```
Number of Links: 4
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 1.1.1.1
(Link Data) Router Interface address: 192.168.13.3
Number of MTID metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 192.168.13.0
(Link Data) Network Mask: 255.255.255.0
Number of MTID metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: another Router (point-to-point)
(Link ID) Neighboring Router ID: 4.4.4.4
(Link Data) Router Interface address: 192.168.34.3
Number of MTID metrics: 0
TOS 0 Metrics: 1
```

```
Link connected to: a Stub Network
(Link ID) Network/subnet number: 192.168.34.0
(Link Data) Network Mask: 255.255.255.0
Number of MTID metrics: 0
TOS 0 Metrics: 1
```

正如我們在上述輸出中所看到的，R3現在已在區域1的Type-1 LSA上設定V位，並在路由進程級別啟用功能傳輸。

我們還可以看到，R1在下面的輸出中為區域1啟用了電纜傳輸：

```
R1#show ip ospf
Routing Process "ospf 1" with ID 1.1.1.1
Start time: 00:02:48.412, Time elapsed: 01:27:00.690
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
Supports area transit capability
Supports NSSA (compatible with RFC 3101)
Supports Database Exchange Summary List Optimization (RFC 5243)
Event-log enabled, Maximum number of events: 1000, Mode: cyclic
It is an area border router
Router is not originating router-LSAs with maximum metric
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPF's 10000 msec
Maximum wait time between two consecutive SPF's 10000 msec
Incremental-SPF disabled
Minimum LSA interval 5 secs
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
EXCHANGE/LOADING adjacency limit: initial 300, process maximum 300
Number of external LSA 0. Checksum Sum 0x000000
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 2. 2 normal 0 stub 0 nssa
Number of areas transit capable is 1
External flood list length 0
IETF NSF helper support enabled
Cisco NSF helper support enabled
Reference bandwidth unit is 100 mbps
  Area BACKBONE(0)
    Number of interfaces in this area is 1
    Area has no authentication
    SPF algorithm last executed 00:00:33.554 ago
    SPF algorithm executed 11 times
    Area ranges are
    Number of LSA 10. Checksum Sum 0x05EB7B
    Number of opaque link LSA 0. Checksum Sum 0x000000
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 3
    Flood list length 0
  Area 1
    Number of interfaces in this area is 1
    This area has transit capability          <-- This area is transit capable
    Area has no authentication
    SPF algorithm last executed 00:00:04.259 ago
    SPF algorithm executed 8 times
    Area ranges are
    Number of LSA 10. Checksum Sum 0x0517AA
    Number of opaque link LSA 0. Checksum Sum 0x000000
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```


由於區域1現在通過所有標準成為中轉區域，因此我們現在應該觀察與我們在第一個場景中看到的不同的路徑計算/優先選擇。

RFC 2328規定，如果某個區域被視為中轉區域，則其檢查應與非中轉區域不同

RFC 2328第16.1節摘要

16.3. Examining transit areas' summary-LSAs

This step is only performed by area border routers attached to one or more non-backbone areas that are capable of carrying transit traffic (i.e., "transit areas", or those areas whose TransitCapability parameter has been set to TRUE in Step 2 of the Dijkstra algorithm (see Section 16.1).

The purpose of the calculation below is to examine the transit areas to see whether they provide any better (shorter) paths than the paths previously calculated in Sections 16.1 and 16.2. Any paths found that are better than or equal to previously discovered paths are installed in the routing table.

根據RFC，如果區域具備傳輸能力，則需遵守RFC 2328第16.3節所述的路徑計算

附註：在本示例中，虛擬鏈路允許傳輸資料流量通過區域1轉發，但是傳輸資料流量採用的實際路徑不需要遵循虛擬鏈路。換句話說，虛擬鏈路允許傳輸流量通過區域轉發，但並不規定流量將採用的準確路徑。

假設在R1上禁用了功能傳輸。讓我們使用traceroute檢查通向目標R6環回的路徑：100 192.0.2.100。

```
R1#traceroute 192.0.2.100
Tracing the route to 192.0.2.100
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.14.4 2 msec 2 msec 2 msec    <--R4
 2 192.168.46.6 3 msec 3 msec *        <--R6
```

在區域1中設定V位後，一旦開啟此功能，我們會觀察以下日誌：

```
R1#debug ip ospf spf intra
OSPF SPF intra debugging is on
R1#debug ip ospf spf inter OSPF SPF inter debugging is on R1#conf Enter configuration commands,
one per line. End with CNTL/Z. R1(config)#router ospf 1 R1(config-router)#capability transit
R1(config-router)#
*Aug 14 15:28:07.934: OSPF-1 INTER: Running spf for summaries in transit area 1
*Aug 14 15:28:07.934: OSPF-1 INTER: Summary transit processing lsid 192.0.2.100 adv_rtr 4.4.4.4
type 3 seq 0x8000000B
*Aug 14 15:28:07.934: OSPF-1 INTER: Summary metric 2
*Aug 14 15:28:07.934: OSPF-1 INTER: found best path to adv_rtr:
i,ABR [2] via 192.168.13.3, GigabitEthernet1, Area 1 orp_txit_adv_rtr 0.0.0.0 pathflag 0x0
*Aug 14 15:28:07.934: OSPF-1 INTER: Add transit path via area 1
*Aug 14 15:28:07.934: OSPF-1 SPF : Exist path: next-hop 192.168.13.3, interface GigabitEthernet1
*Aug 14 15:28:07.934: OSPF-1 INTRA: Route update succeeded for 192.0.2.100/255.255.255.255,
metric 4, Next Hop: GigabitEthernet1/192.168.13.3 area 0
```

現在讓我們檢查R1如何路由到R6 loopback100

```
R1#show ip ospf rib 192.0.2.100
```

OSPF Router with ID (1.1.1.1) (Process ID 1)

Base Topology (MTID 0)

OSPF local RIB

Codes: * - Best, > - Installed in global RIB

LSA: type/LSID/originator

```
*> 192.0.2.100/32, Intra, cost 4, area 0
    SPF Instance 14, age 00:12:28
    Flags: RIB, HiPrio, Transit
    via 192.168.13.3, GigabitEthernet1 label 1048578
    Flags: RIB
    LSA: 1/6.6.6.6/6.6.6.6
```

```
R1#show ip route 192.0.2.100
Routing entry for 192.0.2.100/32
Known via "ospf 1", distance 110, metric 4, type intra area
Last update from 192.168.13.3 on GigabitEthernet1, 00:01:26 ago
Routing Descriptor Blocks:
```

```
 * 192.168.13.3, from 6.6.6.6, 00:01:26 ago, via GigabitEthernet1
Route metric is 4, traffic share count is 1
```

為什麼會看到區域內而不是區域間的？在RFC 2328第16.3節中，我們提到當執行路徑計算時，如果我們有在傳輸區域（型別3）上開銷較低的路由，我們就應該更新字首的下一跳。這的確是我們在上述產出中看到的行為。提到的下一跳是正確的，但型別具有誤導性。

RFC 2328第16.3節摘要

16.3. Examining transit areas' summary-LSAs

(4) Look up the routing table entry for the advertising router BR associated with the Area A. If it is unreachable, examine the next LSA. Otherwise, the cost to destination N is the sum of the cost in BR's Area A routing table entry and the cost advertised in the LSA. Call this cost IAC.

(5) **If this cost is less than the cost occurring in N's routing table entry, overwrite N's list of next hops with those used for BR, and set N's routing table cost to IAC.** Else, if IAC is the same as N's current cost, add BR's list of next hops to N's list of next hops. In any case, the area associated with N's routing table entry must remain the backbone area, and the path type (either intra-area or inter-area) must also remain the same.

R1優先使用區域間Type-3而不是Type-1區域內路由，儘管輸出中將其描述為區域內。我們清楚地看到下一跳與區域0無關

```
R1#show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
4.4.4.4	0	FULL/ -	00:00:39	192.168.14.4	GigabitEthernet3
3.3.3.3	0	FULL/ -	00:00:32	192.168.13.3	GigabitEthernet1

```
R1#show ip ospf neighbor detail
```

```
Neighbor 4.4.4.4, interface address 192.168.14.4
  In the area 0 via interface GigabitEthernet3
```

```
Neighbor priority is 0, State is FULL, 6 state changes
DR is 0.0.0.0 BDR is 0.0.0.0
Options is 0x12 in Hello (E-bit, L-bit)
Options is 0x52 in DBD (E-bit, L-bit, O-bit)
LLS Options is 0x1 (LR)
Dead timer due in 00:00:36
Neighbor is up for 00:30:20
Index 1/1/1, retransmission queue length 0, number of retransmission 3
First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
Last retransmission scan length is 1, maximum is 2
Last retransmission scan time is 135 msec, maximum is 135 msec
```

Neighbor 3.3.3.3, interface address 192.168.13.3

In the area 1 via interface GigabitEthernet1

```
Neighbor priority is 0, State is FULL, 6 state changes
DR is 0.0.0.0 BDR is 0.0.0.0
Options is 0x12 in Hello (E-bit, L-bit)
Options is 0x52 in DBD (E-bit, L-bit, O-bit)
LLS Options is 0x1 (LR)
Dead timer due in 00:00:39
Neighbor is up for 00:30:20
Index 1/1/2, retransmission queue length 0, number of retransmission 3
First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
Last retransmission scan length is 4, maximum is 4
Last retransmission scan time is 126 msec, maximum is 126 msec
```

我們還將跟蹤指向R6 loopback100的目標的路由：

```
R1#traceroute 192.0.2.100
Tracing the route to 192.0.2.100
VRF info: (vrf in name/id, vrf out name/id)
 1 192.168.13.3 2 msec 4 msec 3 msec          <-- R3
 2 192.168.34.4 5 msec 3 msec 3 msec          <-- R4
 3 192.168.46.6 5 msec 6 msec *              <-- R6
R1#
```

因此，在上面的輸出中，我們可以看到主幹區域1比主幹區域0更優先以到達R6環回100。

如果區域內和區域間路由之間的開銷相等，也可以使用兩者的ECMP（等價多重路徑）。在我們的拓撲中，這可以通過將指向R4的R1鏈路從100減少到2來實現。

完成後，RIB和OSPF RIB中都有以下輸出：

```
R1#show ip ospf rib 192.0.2.100 OSPF Router with ID (1.1.1.1) (Process ID 1) Base Topology (MTID 0) OSPF local RIB
Codes: * - Best, > - Installed in global RIB LSA: type/LSID/originator *> 192.0.2.100/32, Intra, cost 4, area 0 SPF
Instance 14, age 00:13:08 Flags: RIB, HiPrio, Transit, OldTrans via 192.168.13.3, GigabitEthernet1 label 1048578
Flags: RIB LSA: 1/6.6.6.6/6.6.6.6 via 192.168.14.4, GigabitEthernet3 label 1048578 Flags: RIB LSA: 1/6.6.6.6/6.6.6.6
```

```
R1#show ip route 192.0.2.100
Routing entry for 192.0.2.100/32
Known via "ospf 1", distance 110, metric 4, type intra area
Last update from 192.168.14.4 on GigabitEthernet3, 00:12:44 ago
Routing Descriptor Blocks:
 192.168.14.4, from 6.6.6.6, 00:12:44 ago, via GigabitEthernet3
Route metric is 4, traffic share count is 1
 * 192.168.13.3, from 6.6.6.6, 00:12:44 ago, via GigabitEthernet1
Route metric is 4, traffic share count is 1
```