

# OSPF虚拟链路：中转能力

## 目录

[简介](#)

[先决条件](#)

[背景信息](#)

[场景 1](#)

[网络图：](#)

[摘自RFC 2328第16.2节](#)

[场景 2](#)

[网络图：](#)

[RFC 2328第6部分的摘要](#)

[摘自RFC 2328第16.1节](#)

[摘自RFC 2328第16.1节](#)

[摘自RFC 2328第16.3节](#)

## 简介

本文档旨在演示V位（虚拟链路位）存在于非主干区域时的开放最短路径优先(OSPF)行为。仅当路由器是一个或多个完全邻接虚拟链路的终端时，才会在第1类LSA中发出V位信号。当设置V位时，这可能会更改区域内路由和区域间路由之间的路径计算首选项。

## 先决条件

使用本文档时，请参阅图1中的网络图：

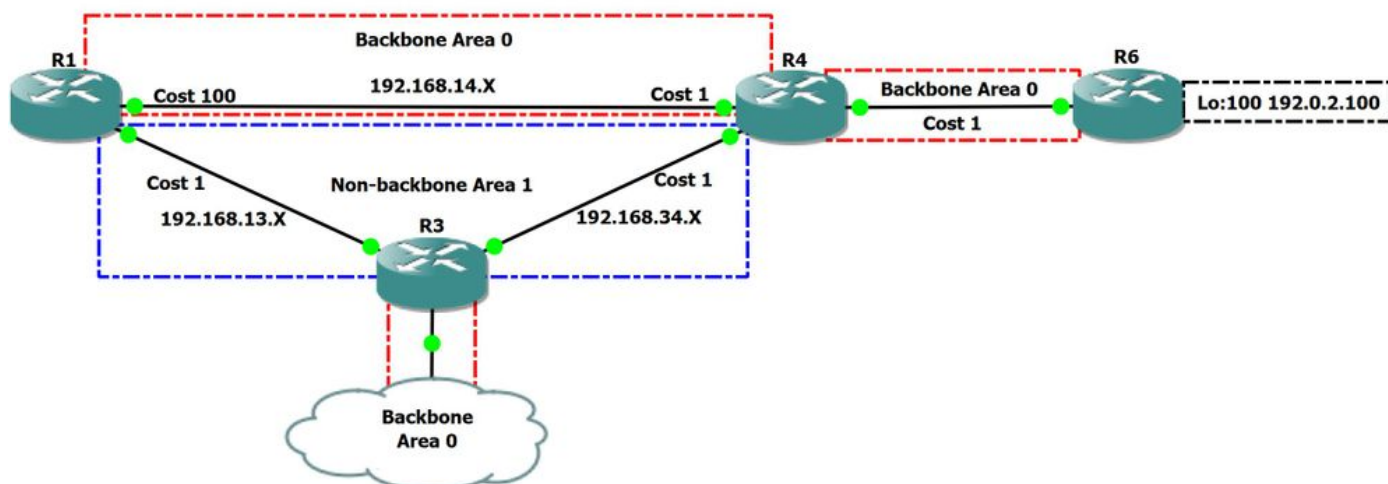


图 1

在上面的网络图中，我们同时拥有主干区域0和非主干区域1。R1是连接区域0和区域1的区域边界路由器(ABR),R4和R3在此网络中具有类似的作用。在此拓扑区域0中，由于R3和R4未通过区域0连接，因此区域0不连续。

## 背景信息

OSPF自治系统中的所有区域都必须连接到主干区域（区域0）。在主干区域之间有非主干区域的情况下，这可能会导致自治系统的某些区域无法访问，并导致网络不连续。如果无法有连续的主干区域，您可以使用虚拟链路通过非主干区域连接您的主干。配置虚拟链路所经过的区域称为中转区域。

# 场景 1

网络图：

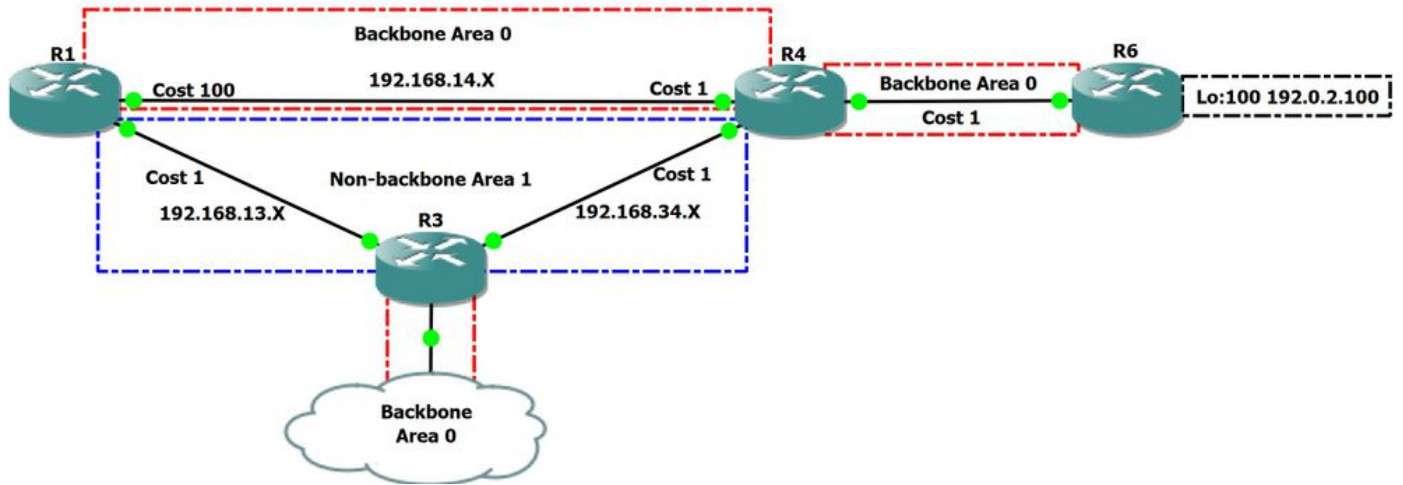


图 2

在此场景中，我们将回顾上述网络拓扑中的预期路径计算。我们将研究从R1路由到IP地址为192.0.2.100/32的R6环回100时首选哪条路径

让我们查看R1上的OSPF数据库，进一步了解拓扑：

```
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	

从上述输出中，我们可以看到R1通过R4获知R6 Lo100:192.0.2.100作为第3类汇总LSA，R1也自身发起第3类汇总LSA，因为它知道R6 Lo100:192.0.2.100通过区域内 主干。在以下输出中，我们可以看到R6直接连接了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

## 摘自RFC 2328第16.2节

### 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.

在上述输出中，我们可以看到它表示区域内路由优先于区域间路由。因此，在我们的场景中，根据RFC 2328,R1应更愿意通过区域内主干。

让我们检查拓扑中是否观察到此行为：

```
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
```

从上面的输出中您可以看到，我们更喜欢将主干区域0转向R6 loopback100。在链路状态数据库中，我们还知道通过R3和R4的区域间路径。通过R4获取的总结LSA的开销为2，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          <- This is Type-3 LSA injected by ABR R4
LS Seq Number: 80000005
Checksum: 0x51AE
Length: 28
Network Mask: /32
          MTID: 0          Metric: 2

```

请考虑，此开销为2表示ABR到目标前缀的开销。第3类LSA从区域0泛洪到非主干区域，反之亦然，它描述ABR到其他区域链路的可达性。它包括从ABR角度考虑的开销，这些开销会拒绝第3类LSA，但会隐藏收到第3类LSA的路由器的全部开销。

从上面的输出中，我们现在知道，我们可以采用两条路径从R1到达R6环回接口：

1. 区域内费用为102
2. 区域间，通过第3类LSA + R1到R4的开销为2，也就是2。这样，总开销为4

在此场景中，我们已经观察到，我们倾向于使用成本更高的区域内路径，因为RFC 2328中定义区域内路径优先于区域间路径。

继续场景2之前，OSPF会先解释第3类LSA:

- ABR R4可以到达区域内链路A，开销为X
- R1可以到达ABR R4，开销为Y
- 表示R1可以通过SPT以X + Y的开销到达链路A




图 3

这就是区域间路由通常与距离矢量协议进行比较的原因，因为区域间的信息是隐藏的。由于区域间OSPF是距离矢量，因此容易出现路由环路。它通过强制使用无环路区域间拓扑来避免环路，在该拓扑中，来自一个区域的流量只能通过区域0到达另一个区域。

## 场景 2

网络图：

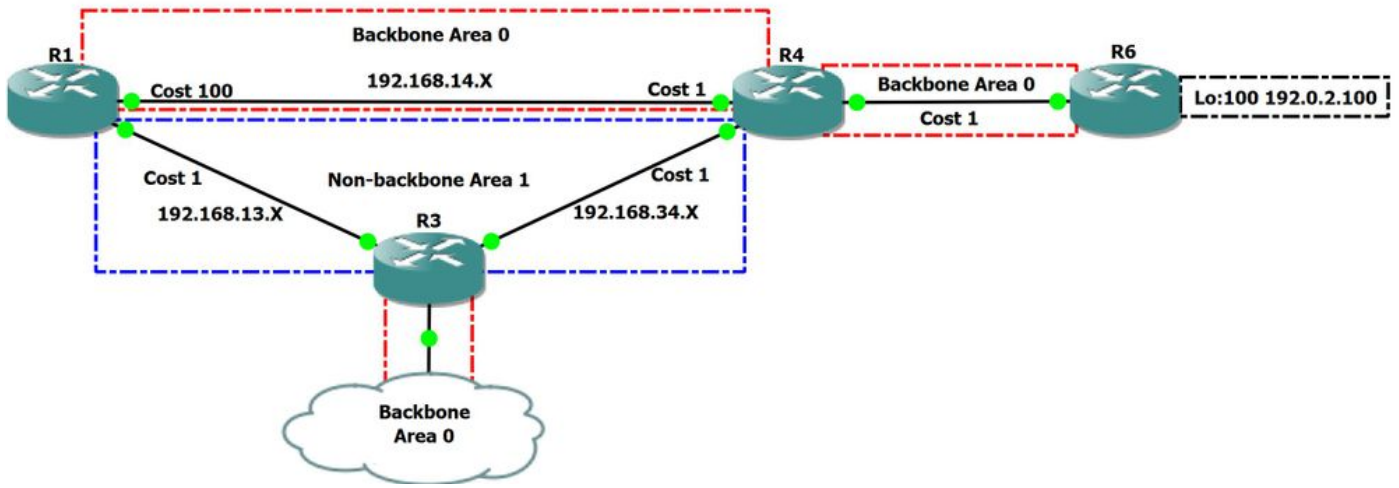


图 4

在此场景中，我们在R3和R4上设置V位，以便当此位存在于非主干区域1的第1类LSA中时，我们可以检查路径优先级。

## RFC 2328第6部分的摘要

### 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".

## 摘自RFC 2328第16.1节

### 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):

从RFC中的上述语句中，我们可以看到，当路由器LSA中设置了V位时，我们知道在其中设置了该位的传输能力区域，或者，在运行Dijkstra算法时，该区域的TransitCapability是正确的。

一旦我们知道如果设置了V位，可以将某个区域视为功能传输区域，则必须检查是否配置了此功能：默认情况下，OSPF区域传输功能已启用。

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

要在区域1中设置V位，我们将创建从R3到R4的虚拟链路。启动虚拟链路时，我们应看到第1类LSA中设置的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#
```

现在，我们检查R3区域1的第1类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
```

如上输出所示，R3现在在其区域1的第1类LSA上设置了V位，并在路由进程级别启用了功能传输。

我们还可以看到，R1在以下输出中为区域1启用了电缆传输：

```
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
```



由于区域1现在通过所有条件成为中转区域，因此我们现在应该观察到不同的路径计算/首选项，在我们的第一个场景中，我们以前看到过。

如果某区域被视为中转区域，则其检查方式应与非中转区域不同，RFC 2328中将说明

## 摘自RFC 2328第16.1节

### 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.

根据RFC，如果区域支持中转，则需按照RFC 2328第16.3节中所述的路径计算

**注意：在本例中，虚拟链路允许通过区域1转发中转数据流量，但中转数据流量采用的实际路径不需要跟随虚拟链路。换句话说，虚拟链路允许通过区域转发中转流量，但不规定流量将采用的精确路径。**

假设R1上禁用了功能传输。让我们使用traceroute检查通往目的R6环回地址：100 192.0.2.100的路径。

```
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
```

在区域1中通过V位设置启用此功能后，我们会观察以下日志：

```
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
```

现在，让我们检查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

为什么我们看到区域内而不是区域间？在RFC 2328第16.3节中，我们提到，在进行路径计算时，如果路由在中转区域（类型3）上具有较低的开销，我们应更新前缀的下一跳。这正是我们在上述输出中看到的行。所提到的下一跳正确，但是类型具有误导性。

## 摘自RFC 2328第16.3节

### 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优先选择区域间第3类，而不是第1类区域内路由，尽管它在输出中表示为区域内路由。我们清楚地看到下一跳与区域0无关

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
<b>3.3.3.3</b>	<b>0</b>	<b>FULL/ -</b>	<b>00:00:32</b>	<b>192.168.13.3</b>	<b>GigabitEthernet1</b>

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
```

我们还要跟踪通往R6环回接口100的目的地址：

```
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，在上述输出中，我们看到非主干区域1优先于主干区域0以到达R6环回100。

如果区域内路由和区域间路由的开销相等，则也可以同时使用ECMP（等价多路径）。这可以通过将R1到R4的链路从100减少到2来实现。

完成此操作后，RIB和OSPF RIB中都有以下输出：

```
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
```