

IPv6 Rapid Deployment (6RD)

Abstract:

This document describes IPv6 Rapid Deployment (6RD) tunneling technique in detail along with configuration example to understand the 6RD deployment steps.

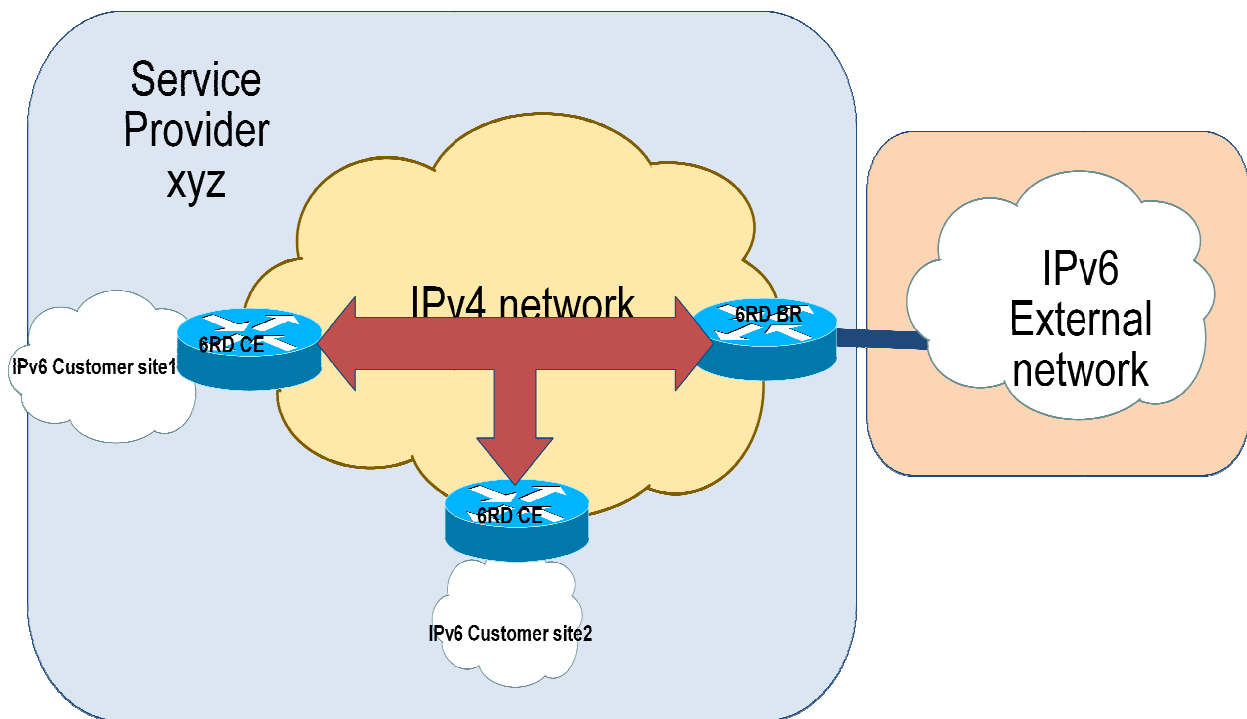
Introduction:

As part of migrating from IPv4 to IPv6, we may face up with situation where we need to connect IPv6 islands over IPv4 cloud. While the connectivity can be achieved by point to point tunnel like GRE, it is not scalable and so automatic 6to4 tunneling concept was introduced. Automatic 6to4 tunnel is a point to multipoint tunnel which treats IPv4 cloud as NBMA link. This technique requires that IPv6 cloud should be designed with special format address also known as 6to4 address that starts with prefix 2002::/16 followed by 32 bit IPv4 address. It derives the tunnel destination IPv4 address from the IPv6 address.

6RD, IPv6 Rapid Deployment is the latest tunneling technique which like automatic 6to4 tunnel utilizes stateless IPv6 in IPv4 encapsulation to transit IPv4 cloud. This is like superset of Automatic 6to4 technique with key differentiator that it utilizes service provider's own prefix instead of reserved prefix 2002::/16.

6RD Components

6RD domain consists of 2 components as 6RD Customer Edge (CE) routers and 6RD Border Relay (BR). IPv6 packets destined within same 6RD domain will traverse between 6RD CE routers while 6RD BR will connect 6RD domain to IPv6 Internet. Any IPv6 packet destined to outside world will traverse over IPv4 cloud to BR which will be sent to external IPv6 network after decapsulation.

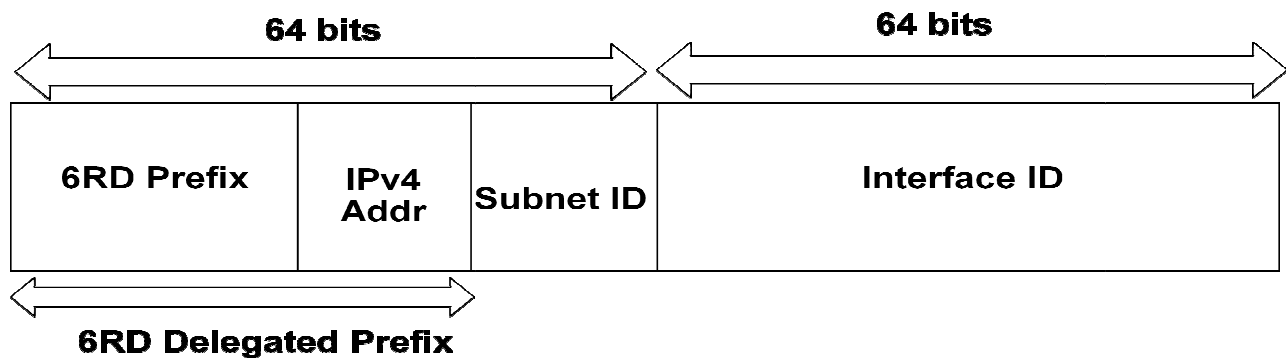


Each customer site connected to 6RD CE router will be configured with 6RD Delegated Prefix which will be created by combining 6RD prefix and all or part of CE router's IPv4 address. While Automatic 6to4 tunnel technique requires the entire IPv4 address to be embedded after 2002::/16 prefix, 6RD doesn't always require the entire 32 bits to be embedded. Assume that the entire IPv4 network (or at least all CE routers and BR routers) are reachable by 10.x.y.z, we can skip first 8 bit which is common and embed the remaining 24 bits which will be used to identify remote 6RD CE or BR device.

6RD Address Format

Below is the 6RD address format,

Figure 2.6



6RD prefix is a variable length prefix that can be of Service Providers choice from the prefix they received from registry. Since the prefix belongs to SP, it will be always unique in Internet as well.

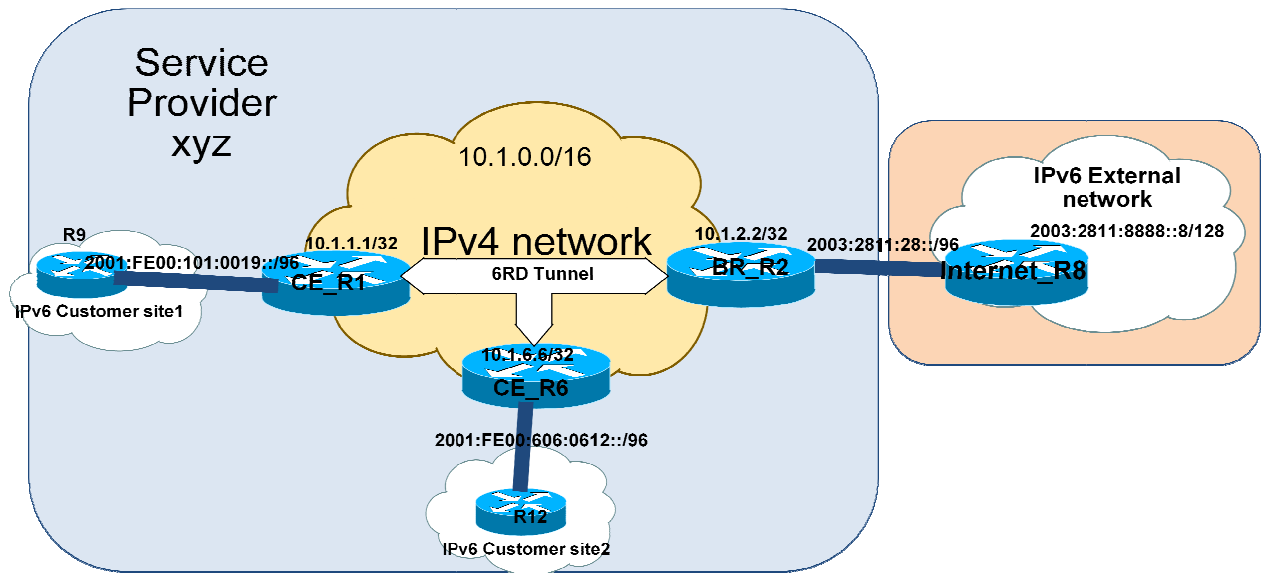
Embedded IPv4 address is a variable length field which will be used to identify remote CE or BR device. As mentioned earlier, 6RD supports prefix compression that if the entire service provider network address commonly starts with 10.x.y.z, we can omit 10 and only use 24 bits in Embedded IPv4 field to fit x.y.z.

Subnet ID is another variable length field. All 6RD CE and 6RD BR will be given with same 6RD prefix to which it will append/embed its own IPv4 address (after compression if required). Now Subnet ID will be used to differentiate between subnets.

Interface ID is 64 bit field which will be derived from host MAC address.

The combination of 6RD prefix and all or part of CE IPv4 address is known as **6RD delegated prefix**. Each customer site will be allotted with one 6RD delegated prefix and subnet ID will be used to differentiate subnets within customer site.

6RD Configuration Example:



In the above topology, assume the SP core is designed with 10.1.0.0/16 network. So any 6RD CE or BR will have address in range of 10.1.0.0/16 network. This SP selects 2001:FE00::/32 as the 6RD prefix. CE_R1 and CE_R6 will act as 6RD CE nodes while BR_R2 will act as 6RD BR node connecting to external IPv6 domain.

On CE_R1, loopback address is 10.1.1.1/32 which is reachable from other devices in IPv4 cloud. Since the first 16 bits are common among all devices, we omit the same and use only last 16 bit to identify other 6RD CE and BR devices. Now concatenating 6RD prefix which is 2001:FE00::/32 and last 16 bits of CE device will make 6RD delegated prefix for Customer site1 as 2001:FE00:0101::/48 which further will be divided as different subnet and assigned to different network in customer site1 connected to CE_R1. In our topology, 2001:FE00:101:0019::/96 is assigned between R1 and R9. Below is the summary of 6RD configuration required on CE device,

- Create tunnel interface and enable the mode as 6RD using “tunnel mode ipv6ip 6rd”
- Specify the 6RD delegated prefix under tunnel interface using “tunnel 6rd prefix <6rd-delegated-prefix>”
- Specify the IPv4 field length under tunnel interface using “tunnel 6rd ipv6 prefix-len <>”
- Specify the 6RD BR address under tunnel interface using “tunnel 6rd br <br-ipv4-addr>”
- Configure a static route for 6RD prefix (not 6RD delegated prefix) pointing tunnel interface as outgoing interface using “ipv6 route <6rd-prefix>::/<len> tunnel<>”
- Configure a default static route and point towards 6RD BR using “ipv6 route ::/0 tunnel<> <6rd-br-address>”

On CE_R6, loopback address is 10.1.6.6/32 which is reachable from other devices in IPv4 cloud. 6RD delegated prefix to customer site2 connected to CE_R6 will be 2001:FE00:606::/48. This further will be divided as different subnet and assigned to network. In our topology, 2001:FE00:606:0612::/96 is assigned between R12 and R6. Above same 6RD CE related configuration to be done.

Since BR_R2 is connected to external domain, it doesn't require address assigned from 6RD delegated prefix except on tunnel interface. Below is the summary of 6RD configuration required on BR device,

- Create tunnel interface and enable the mode as 6RD using "tunnel mode ipv6ip 6rd"
- Specify the 6RD delegated prefix under tunnel interface using "tunnel 6rd prefix <6rd-delegated-prefix>"
- Specify the IPv4 field length under tunnel interface using "tunnel 6rd ipv4 prefix-len <>"
- Configure a static route for 6RD prefix (not 6RD delegated prefix) pointing tunnel interface as outgoing interface using "ipv6 route <6rd-prefix>::/<len> tunnel<>"

CE_R1 Configuration:

```
hostname R1
!
ip cef
ipv6 general-prefix PREFIX 6rd Tunnel0
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 ip address 10.1.1.1 255.255.255.0
!
interface Tunnel0
 no ip address
 no ip redirects
 ipv6 enable
 tunnel source Loopback0
 tunnel mode ipv6ip 6rd
 tunnel 6rd ipv4 prefix-len 16
 tunnel 6rd prefix 2001:FE00::/32
 tunnel 6rd br 10.1.2.2
!
interface Ethernet0/0
 no ip address
!
interface Ethernet0/0.13
 encapsulation dot1Q 13
 description core facing interface
 ip address 10.1.13.1 255.255.255.0
!
interface Ethernet0/0.19
 encapsulation dot1Q 19
 description CE facing interface
 ipv6 address 2001:FE00:101:19::1/96
!
router ospf 1
 network 10.1.0.0 0.0.255.255 area 0
```

```
!  
ipv6 route 2001:FE00::/32 Tunnel0  
ipv6 route ::/0 Tunnel0 2001:FE00:202::  
ipv6 router ospf 1910  
  default-information originate  
  redistribute static  
!  
end  
  
R1#  
R1#show tunnel 6rd  
Interface Tunnel0:  
  Tunnel Source: 10.1.1.1  
  6RD: Operational, V6 Prefix: 2001:FE00::/32  
    V4 Prefix, Length: 16, Value: 10.1.0.0  
    V4 Suffix, Length: 0, Value: 0.0.0.0  
    Border Relay address: 10.1.2.2  
  General Prefix: 2001:FE00:101::/48  
R1#show ipv6 interface tunnel0  
Tunnel0 is up, line protocol is up  
IPv6 is enabled, link-local address is FE80::A01:101  
No Virtual link-local address(es):  
No global unicast address is configured  
Joined group address(es):  
  FF02::1  
  FF02::2  
  FF02::1:FF01:101  
MTU is 1480 bytes  
ICMP error messages limited to one every 100 milliseconds  
ICMP redirects are enabled  
ICMP unreachable are sent  
Post_Encap features: Tunnel 6RD  
ND DAD is not supported  
ND reachable time is 30000 milliseconds (using 30000)  
ND RAs are suppressed (periodic)  
Hosts use stateless autoconfig for addresses.  
R1#
```

CE_R6 Configuration:

```
hostname R6  
!  
ip cef  
ipv6 general-prefix PREFIX 6rd Tunnel0  
ipv6 unicast-routing  
ipv6 cef  
!  
interface Loopback0
```

```
ip address 10.1.6.6 255.255.255.255
!
interface Tunnel0
no ip address
no ip redirects
ipv6 enable
tunnel source Loopback0
tunnel mode ipv6ip 6rd
tunnel 6rd ipv4 prefix-len 16
tunnel 6rd prefix 2001:FE00::/32
tunnel 6rd br 10.1.2.2
!
interface Ethernet0/0
no ip address
!
interface Ethernet0/0.56
encapsulation dot1Q 56
description core facing interface
ip address 10.1.56.6 255.255.255.0
!
interface Ethernet0/0.612
encapsulation dot1Q 612
ipv6 address 2001:FE00:606:612::6/96
description Customer site facing interface
ipv6 ospf 6712 area 0
!
router ospf 1
network 10.1.0.0 0.0.255.255 area 0
!
ipv6 route 2001:FE00::/32 Tunnel0
ipv6 router ospf 6712
redistribute static
!
end

R6#
R6#show tunnel 6rd
Interface Tunnel0:
  Tunnel Source: 10.1.6.6
  6RD: Operational, V6 Prefix: 2001:FE00::/32
    V4 Prefix, Length: 16, Value: 10.1.0.0
    V4 Suffix, Length: 0, Value: 0.0.0.0
    Border Relay address: 10.1.2.2
  General Prefix: 2001:FE00:606::/48
R6#show ipv6 interface tunnel0
Tunnel0 is up, line protocol is up
IPv6 is enabled, link-local address is FE80::A01:606
```

```
No Virtual link-local address(es):
No global unicast address is configured
Joined group address(es):
  FF02::1
  FF02::2
  FF02::1:FF01:606
MTU is 1480 bytes
ICMP error messages limited to one every 100 milliseconds
ICMP redirects are enabled
ICMP unreachable are sent
Post_Encap features: Tunnel 6RD
ND DAD is not supported
ND reachable time is 30000 milliseconds (using 30000)
ND RAs are suppressed (periodic)
Hosts use stateless autoconfig for addresses.
R6#
```

BR_R2 Configuration:

```
hostname R2
!
ip cef
ipv6 general-prefix PREFIX 6rd Tunnel0
ipv6 unicast-routing
ipv6 cef
!
interface Loopback0
 ip address 10.1.2.2 255.255.255.0
!
interface Tunnel0
 no ip address
 no ip redirects
 ipv6 enable
 tunnel source Loopback0
 tunnel mode ipv6ip 6rd
 tunnel 6rd ipv4 prefix-len 16
 tunnel 6rd prefix 2001:FE00::/32
!
interface Ethernet0/0
 no ip address
!
interface Ethernet0/0.24
 encapsulation dot1Q 24
 description core facing interface
 ip address 10.1.24.2 255.255.255.0
!
interface Ethernet0/0.28
 encapsulation dot1Q 28
```

```
description External connected interface
ipv6 address 2003:2811:28::2/96
!
router ospf 1
network 10.1.0.0 0.0.255.255 area 0
!
router bgp 65500
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 2003:2811:28::8 remote-as 65535
!
address-family ipv4
exit-address-family
!
address-family ipv6
redistribute static
network 2001:FE00::/32
neighbor 2003:2811:28::8 activate
exit-address-family
!
ipv6 route 2001:FE00::/32 Tunnel0
!
end
```

R2#

```
R2#show tunnel 6rd
```

```
Interface Tunnel0:
```

```
Tunnel Source: 10.1.2.2
```

```
6RD: Operational, V6 Prefix: 2001:FE00::/32
```

```
  V4 Prefix, Length: 16, Value: 10.1.0.0
```

```
  V4 Suffix, Length: 0, Value: 0.0.0.0
```

```
General Prefix: 2001:FE00:202::/48
```

```
R2#show ipv6 interface tunnel0
```

```
Tunnel0 is up, line protocol is up
```

```
IPv6 is enabled, link-local address is FE80::A01:202
```

```
No Virtual link-local address(es):
```

```
No global unicast address is configured
```

```
Joined group address(es):
```

```
  FE02::1
```

```
  FE02::2
```

```
  FE02::1:FF01:202
```

```
MTU is 1480 bytes
```

```
ICMP error messages limited to one every 100 milliseconds
```

```
ICMP redirects are enabled
```

```
ICMP unreachable are sent
```

```
Post_Encap features: Tunnel 6RD
```

```
ND DAD is not supported
```



```
ND reachable time is 30000 milliseconds (using 30000)
ND RAs are suppressed (periodic)
Hosts use stateless autoconfig for addresses.
R2#
```

Verification:

```
R9#ping 2003:2811:8888::8 source e0/0.19
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2003:2811:8888::8, timeout is 2 seconds:
Packet sent with a source address of 2001:FE00:101:19::9
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 4/4/5 ms
R9#

R1#
*Jun 20 16:19:39.206: Tunnel0: IPv6/IP adjacency fixup, 10.1.1.1->10.1.2.2, tos set to 0x0
*Jun 20 16:19:39.209: Tunnel0: IPv6/IP to classify 10.1.2.2->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:19:39.213: Tunnel0: IPv6/IP adjacency fixup, 10.1.1.1->10.1.2.2, tos set to 0x0
*Jun 20 16:19:39.216: Tunnel0: IPv6/IP to classify 10.1.2.2->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:19:39.218: Tunnel0: IPv6/IP adjacency fixup, 10.1.1.1->10.1.2.2, tos set to 0x0
*Jun 20 16:19:39.222: Tunnel0: IPv6/IP to classify 10.1.2.2->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
R1#
*Jun 20 16:19:39.225: Tunnel0: IPv6/IP adjacency fixup, 10.1.1.1->10.1.2.2, tos set to 0x0
*Jun 20 16:19:39.228: Tunnel0: IPv6/IP to classify 10.1.2.2->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:19:39.231: Tunnel0: IPv6/IP adjacency fixup, 10.1.1.1->10.1.2.2, tos set to 0x0
*Jun 20 16:19:39.235: Tunnel0: IPv6/IP to classify 10.1.2.2->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
R1#

R2#
R2#
*Jun 20 16:19:39.207: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.2.2 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:19:39.208: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.2.2->10.1.1.1, tos set to
0x0
*Jun 20 16:19:39.214: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.2.2 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:19:39.215: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.2.2->10.1.1.1, tos set to
0x0
*Jun 20 16:19:39.219: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.2.2 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
```

```
*Jun 20 16:19:39.221: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.2.2->10.1.1.1
R2#, tos set to 0x0
*Jun 20 16:19:39.226: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.2.2 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:19:39.227: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.2.2->10.1.1.1, tos set to
0x0
*Jun 20 16:19:39.232: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.2.2 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:19:39.234: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.2.2->10.1.1.1, tos set to
0x0
R2#
```

```
R9#ping 2001:fe00:0606:0612::12 sou e0/0.19
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:FE00:606:612::12, timeout is 2 seconds:
Packet sent with a source address of 2001:FE00:101:19::9
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 5/5/6 ms
R9#
```

```
R1#
*Jun 20 16:20:31.355: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.1.1->10.1.6.6, tos set to
0x0
*Jun 20 16:20:31.360: Tunnel0: IPv6/IP to classify 10.1.6.6->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:20:31.363: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.1.1->10.1.6.6, tos set to
0x0
*Jun 20 16:20:31.367: Tunnel0: IPv6/IP to classify 10.1.6.6->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:20:31.369: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.1.1->10.1.6.6, tos set to
0x0
R1#
```

```
*Jun 20 16:20:31.373: Tunnel0: IPv6/IP to classify 10.1.6.6->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:20:31.377: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.1.1->10.1.6.6, tos set to
0x0
*Jun 20 16:20:31.381: Tunnel0: IPv6/IP to classify 10.1.6.6->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
*Jun 20 16:20:31.383: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.1.1->10.1.6.6, tos set to
0x0
*Jun 20 16:20:31.388: Tunnel0: IPv6/IP to classify 10.1.6.6->10.1.1.1 (tbl=0,"default" len=120 ttl=252
tos=0x0) ok, oce_rc=0x0
R1#
```

```
R6#
```

*Jun 20 16:20:31.356: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.6.6 (tbl=0,"default" len=120 ttl=252 tos=0x0) ok, oce_rc=0x0

*Jun 20 16:20:31.359: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.6.6->10.1.1.1, tos set to 0x0

*Jun 20 16:20:31.364: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.6.6 (tbl=0,"default" len=120 ttl=252 tos=0x0) ok, oce_rc=0x0

*Jun 20 16:20:31.366: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.6.6->10.1.1.1, tos set to 0x0

*Jun 20 16:20:31.370: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.6.6 (tbl=0,"default" len=120 ttl=252 tos=0x0) ok, oce_rc=0x0

*Jun 20 16:20:31.372: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.6.6->10.1.1.1
R6#, tos set to 0x0

*Jun 20 16:20:31.378: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.6.6 (tbl=0,"default" len=120 ttl=252 tos=0x0) ok, oce_rc=0x0

*Jun 20 16:20:31.380: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.6.6->10.1.1.1, tos set to 0x0

*Jun 20 16:20:31.384: Tunnel0: IPv6/IP to classify 10.1.1.1->10.1.6.6 (tbl=0,"default" len=120 ttl=252 tos=0x0) ok, oce_rc=0x0

*Jun 20 16:20:31.386: Tunnel0: IPv6/IP 6RD post encapsulation feature, 10.1.6.6->10.1.1.1, tos set to 0x0
R6#