White Paper

IPv6 Rapid Deployment: Provide IPv6 Access to Customers over an IPv4-Only Network

What You Will Learn

IPv6 Rapid Deployment (6rd) (RFC 5969) 6rd is a stateless tunneling mechanism which allows an Service Provider to rapidly deploy IPv6 in a lightweight and secure manner without requiring upgrades to existing IPv4 access network infrastructure. While there are a number of methods for carrying IPv6 over IPv4, 6rd has been particularly successful due to its stateless mode of operation which is lightweight and naturally scalable, resilient, and simple to provision. The service provided by 6rd is production quality, it “Looks smells and feels like native IPv6” to the customer and the Internet at large.

This document presents:

- 6rd technology description
- 6rd deployment scenarios
- Cisco IOS® Software configuration for 6rd

Overview of 6rd Operation

Figure 1 provides an overview of 6rd operation.

Figure 1. Overview of 6rd Operation

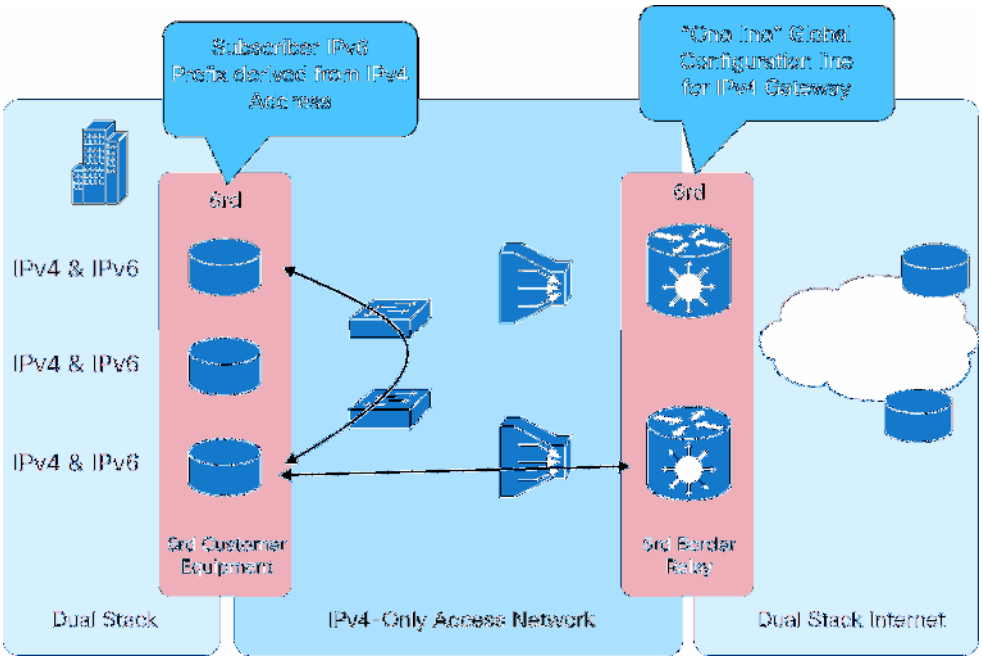

6rd consists of two main hardware components, the CE (Customer Equipment) router and the BR (Border Relay) router.

Customer Edge Router

The CE router sits at the edge of the service provider IPv4 access infrastructure and provides IPv6 connectivity to this end user's network. The native IPv6 traffic coming from the end user hosts is encapsulated in IPv4 by the CE router and tunneled to the BR router or directly to other CE routers in the same 6rd domain. Conversely, encapsulated 6rd traffic received from the Internet through the BR router and 6rd traffic from other CE routers will be de-capsulated and forwarded to the end-user nodes.

Border Relay Router

The BR router provides connectivity between the CE routers and the IPv6 network (public or private Internet). Both the CE and BR routers are dual-stack devices, and the devices between the BR and CE routers can be IPv4 only.

At the CE router, if the packet IPv6 destination address matches the locally configured 6rd prefix, the packet is considered to be part of the local 6rd domain and needs to be forwarded to another CE router. In such a case, the IPv4 address embedded in the IPv6 destination address is used as the IPv4 destination address of the 6rd tunnel, and the local WAN interface IPv4 address is used as the source address for the 6rd tunnel, which is an IPv6 packet directly encapsulated in IPv4. If the IPv6 destination address does not match the locally configured 6rd prefix—in other words, if the packet does not belong to the local 6rd domain—the packet will be tunneled to the BR router by a 6rd tunnel. In this case, the locally configured BR IPv4 address on the CE router is used as the destination address for the encapsulated packet.

6rd Operation Details

Figure 2 shows a 6rd network.

Figure 2.

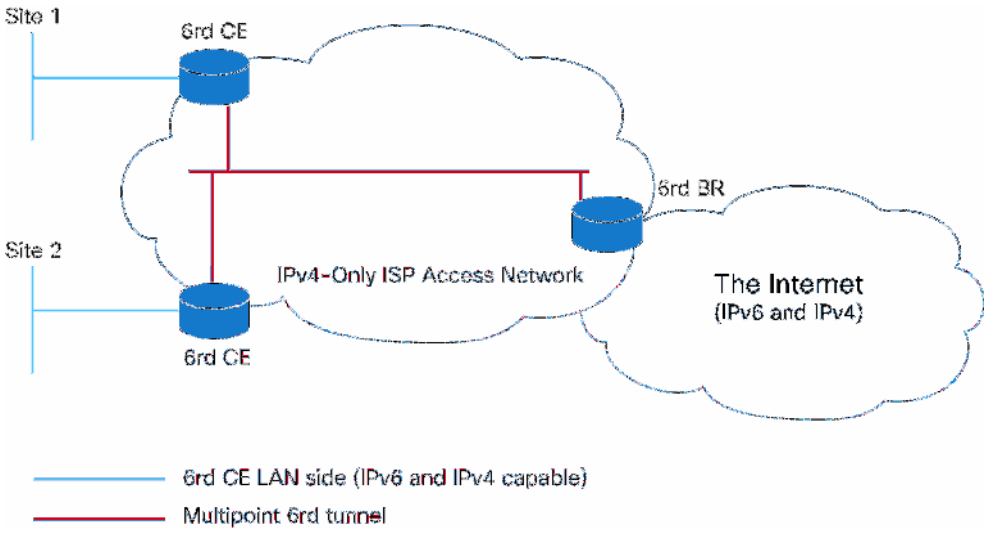

- The 6rd CE LAN-side interface carries traffic to and from IPv6 hosts.
- The multipoint tunnel interface carries tunnel encapsulated traffic to and from IPv6 hosts.

- The encapsulation used for the 6rd tunnel is a direct IPv6-in-IPv4 encapsulation. The IPv4 protocol field is set to Protocol 41.

6rd Delegated Prefix

In the network shown in Figure 2, the CE routers provide a range of prefixes to their sites. These prefixes are called 6rd delegated prefixes and are similar to IPv6 Domain Host Configuration Protocol Version 6 (DHCPv6) PD prefixes. A 6rd delegated prefix consists of the following elements:

- An IPv6 prefix selected by the service provider to be the common 6rd service provider prefix for the given 6rd deployment
- An assigned IPv4 address for the CE router; this address can be global or private, and 6rd does not have to use all 32 bits of the IPv4 address (as explained later in this document)

Consider the example in Figure 3.

Figure 3. Delegated Prefix Example

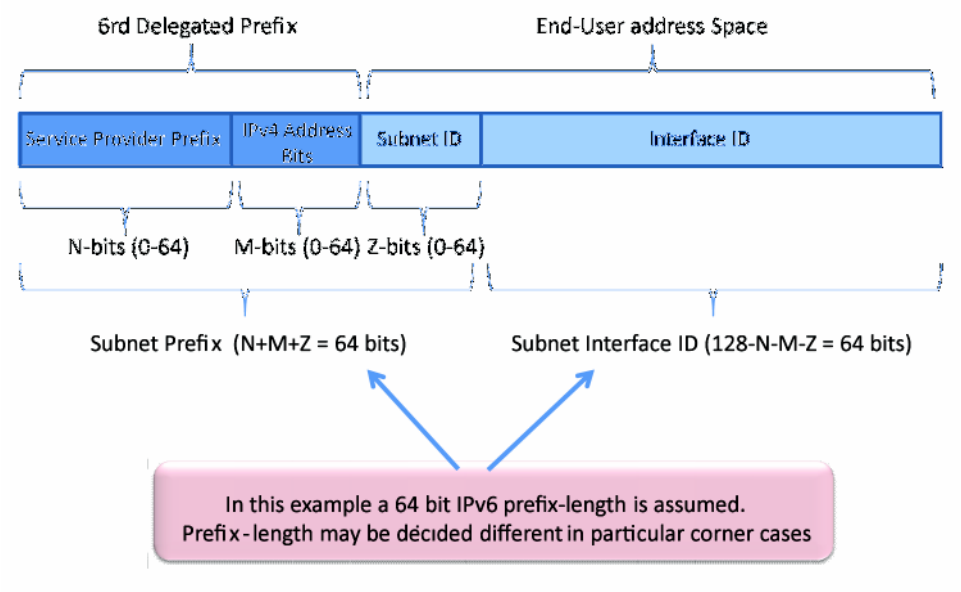

The example in Figure 3 shows the following:

- The service provider–selected prefix is 2001:DB8::/32.
- Each 6rd CE router uses an IPv4 address from the 10.0.0.0/8 block. In 6rd you can embed fewer than 32 bits of the IPv4 address in the 6rd delegated prefix. This action is made possible by the introduction of the following two concepts:
  - IPv4 common prefix: All the 6rd CE routers and the BR router in the 6rd domain can share a common IPv4 prefix for their IPv4 address blocks. This common prefix is provisioned to all nodes in the 6rd domain and therefore need not be carried in the IPv6 destination to identify a tunnel endpoint. In the example in Figure 3, the IPv4 common prefix is 10.0.0.0/8.
  - IPv4 common suffix: All the 6rd CE routers and the BR router can agree on a common tail portion of the IPv4 address to identify a tunnel endpoint. For instance, in this example, suppose that the IPv4 common

suffix is 0.0.0.1/8. This suffix implies that the following tunnel transport endpoints will be used (on the CE router, this address is typically the IPv4 address on the 6rd CE WAN-side interface, which is used as the tunnel source):

  - CE1: 10.1.1.1
  - CE2: 10.1.2.1
  - BR: 10.1.3.1

- The number of bits in the IPv4 address blocks that are unique—that is, not common to the 6rd CE routers and BR router in a domain can be calculated as follows:
  (32 bits)–(IPv4 common prefix length)–(IPv4 common suffix length)
  In the preceding example, the value would be:
  32 – 16 – 8 = 8.
  These 8 bits need to be embedded in the 6rd delegated prefix.
- The 6rd delegated prefix length therefore, is the sum of the service provider–selected prefix length and the number of bits in the IPv4 address blocks that are not common to the 6rd CE routers and BR router in a domain: In the preceding example, the length is:
  32 + 8 = 40
- The 6rd parameters are shown in Table 1.

Table 1.

Parameter	Value
Service provider prefix	2001: DB8::/32
IPv4 common prefix	10.1.0.0/16
IPv4 common suffix	0.0.0.1/8
CE1: Delegated 6rd prefix	2001: DB8:0100::/40
CE2: Delegated 6rd prefix	2001: DB8:0200::/40
BR: Delegated 6rd prefix	2001: DB8:0300::/40
CE1 (IPv4) tunnel transport source	10.1.1.1
CE2 (IPv4) tunnel transport source	10.1.2.1
BR (IPv4) tunnel transport source	10.1.3.1

Example 2

It is quite common that a challenger Service Provider is restricted by the incumbent Service Provider resulting in a fact that compression of the IPv4 address within the 6rd address space is not feasible. Due to that, the full 32-bits of the IPv4 address needs to be inserted into the 6rd IPv6 address as shown in picture 4 below:

Figure 4.

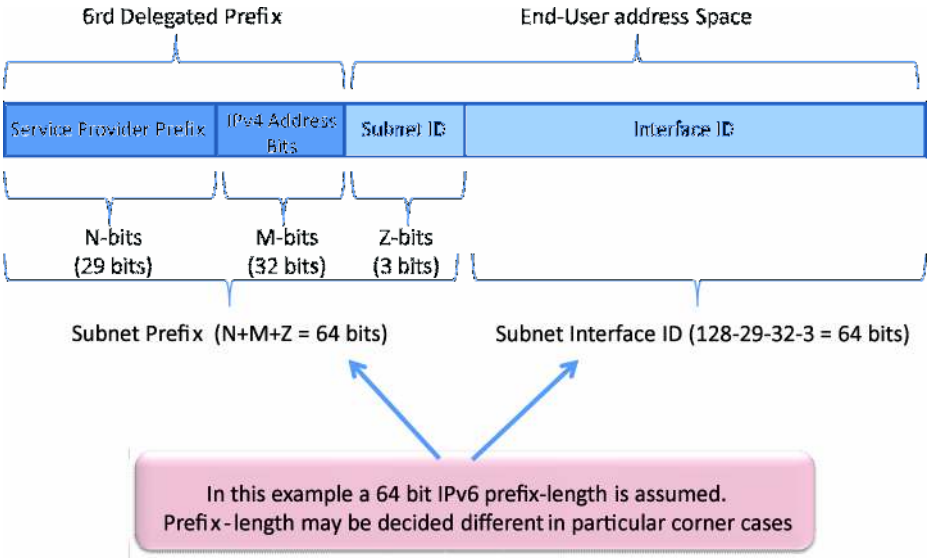

Regional registries (RIPE, ARIN, etc…) are working to a model where a Service Provider is providing IPv6 services through 6rd can get a /29 IPv6 address allocation instead of the minimal /32 address allocation. This would give 3 bits per end-user site to segment the local network, allowing 8 subnets per end-user site.

6rd Address Tunnel Endpoint Determination

When a native IPv6 packet destined to a 6rd domain address arrives at a 6rd CE router, it needs to be sent to the appropriate destination CE router. The destination IPv4 address for the 6rd tunnel is obtained using the following rules:

- Determine the number of bits of the IPv4 address carried in the IPv6 header, as follows:

  (32 bits)–(IPv4 common prefix length)–(IPv4 common suffix length)
- Determine the position of those bits in the IPv6 header. Figure 4 shows the 6rd domain prefix length.

Figure 5. IPv4 Address Bits in 6rd IPv6 Address Using CE1 as the Example

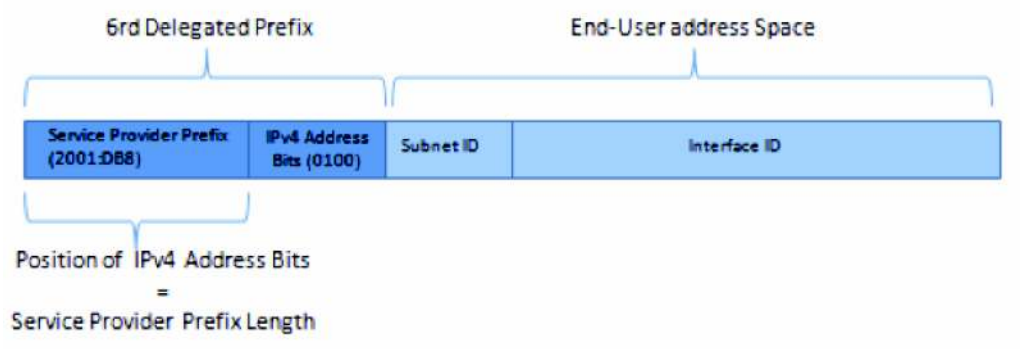

- Extract the bits of the IPv4 address carried in the IPv6 destination address header. This extraction can be performed now that the 6rd domain address and the length of the common prefix is known.
- Start with the IPv4 common prefix, then append the bits extracted from the IPv6 header, and then append the IPv4 common suffix (Figure 6).

Figure 6. IPv4 Destination Address

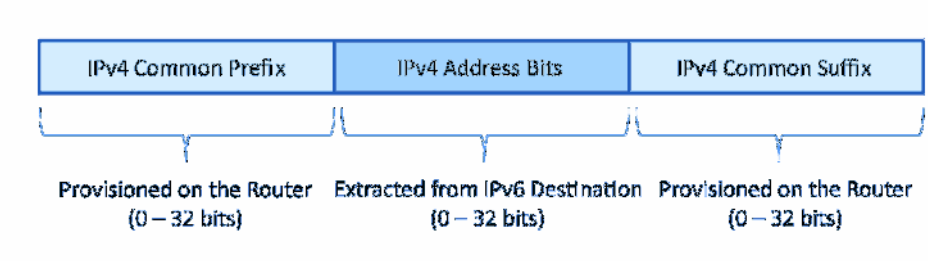

Now apply the preceding algorithm to the 6rd network presented earlier. Consider a packet destined for host 2001:DB8:0100::11, which is a host in CE site 1.

- The number of bits carried in the IPv6 header is:
  (32 bits) – (IPv4 common prefix length) – (IPv4 common suffix length)
  Here, the value is:
  32 – 16 – 8 = 8
- Determine the location of those bits in the IPv6 header. This is the 6rd domain prefix length:
  32
- Extract the bits of the IPv4 address carried in the IPv6 header. This extraction can be performed now that you know the location and the length. The result is a binary value between the bits 33 and 40 for 8 bits, the example shown here:
  Extracting 8 bits from the bit 32 in 2001:DB8:0100::11 yields hexadecimal 01 (IPv6 addresses are hexadecimal).
  The binary value is 00000001, which is 1 in the decimal numbering system.
- To reconstruct the address of the remote 6rd CE router, start with the IPv4 common prefix, then append the bits extracted from the IPv6 header, and then append the IPv4 common suffix.
  The result is 10.1.1.1, which is the address for CE1

The preceding 6rd tunnel endpoint determination is run each time the CE router receives an IPv6 packet destined for a destination within the 6rd local domain.

If the received native IPv6 packet is destined for the IPv6 Internet beyond the 6rd local domain, then the IPv4 address of the 6rd BR router will be provided by manual input.

Routing Considerations

Native IPv4 routing is used between the CE and BR routers in the 6rd domain. For high availability, more then one BR router can be configured. To achieve this goal, the BR routers must use an IPv4 anycast address advertised in the IPv4 Interior Gateway Protocol (IGP), resulting in multiple 6rd BR routers in the 6rd domain. The CE router will then use the closest BR router based on the IGP selection rules.

The service provider must announce the registered IPv6 address range (6rd delegated prefix) to the IPv6 Internet for global reachability.

Lifecycle of a 6rd Packet

This section describes step by step how a packet goes from the CE router to another CE router or to the IPv6 Internet and back (Figure 7).

Figure 7. 6rd Packet Lifecycle

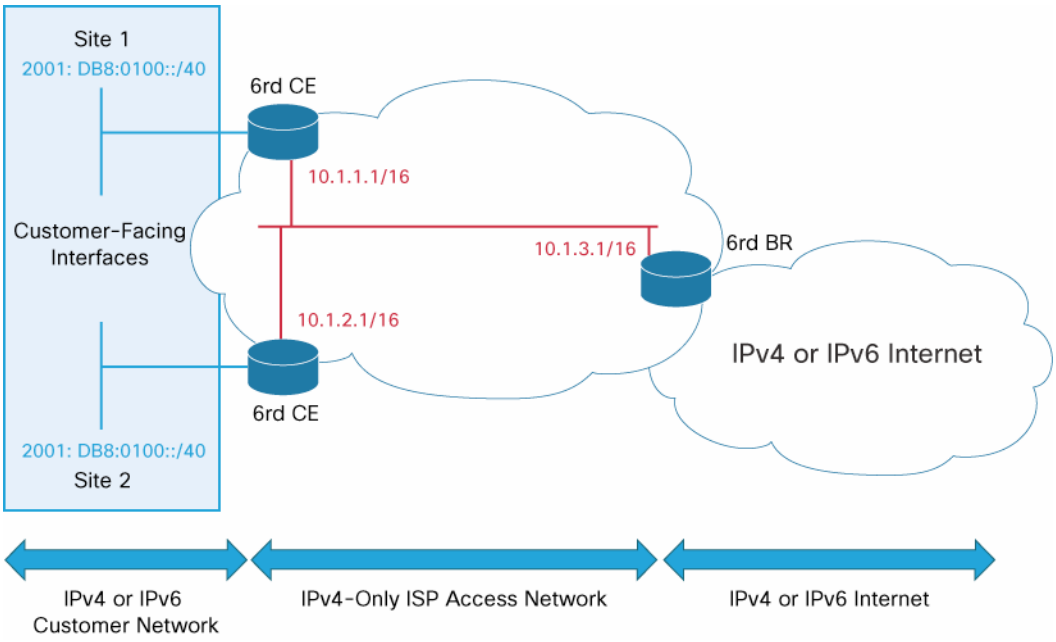

CE to CE

IPv6 traffic is received natively on a customer-facing interface if the IPv6 destination address falls into the range of the locally configured 6rd domain prefix and then needs to be forwarded to another 6rd CE router (Figure 8).

Figure 8. CE to CE

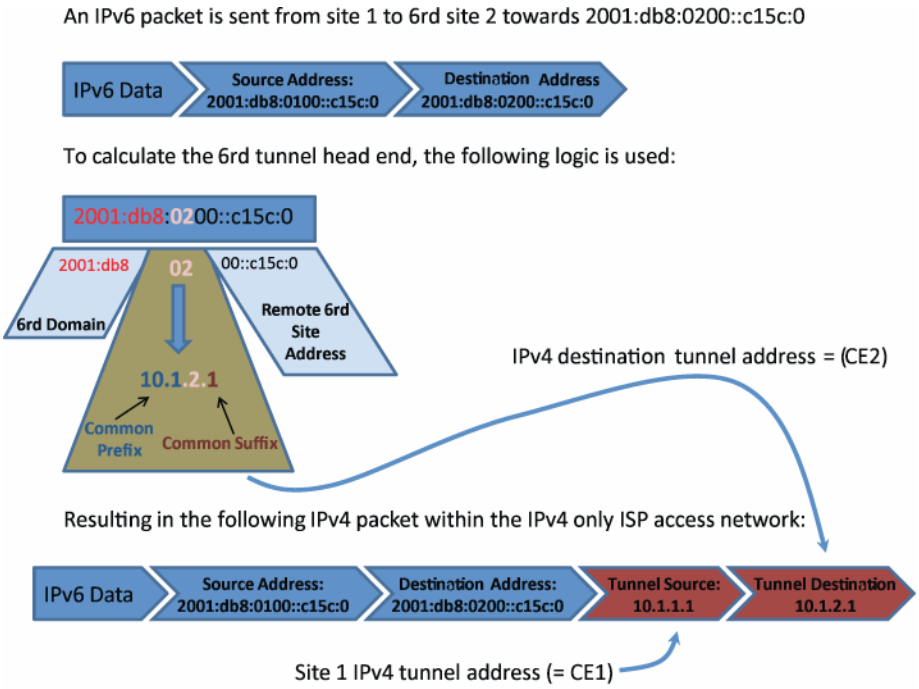

The IPv6 packet is encapsulated in an IPv4 header. The embedded IPv4 address is copied to the IPv4 destination address. The locally configured tunnel source is copied to the IPv4 source address. The IPv4 tunnel header field protocol type is set to 41 (IPv6 in IPv4).

The IPv4 packet, tunneling the IPv6 packet, is forwarded to the destination CE router through the IPv4 domain following the IPv4 routing table.

The destination CE router receives the tunneled IPv6 packet, and the IPv4 header is removed. As a security measure, the IPv4 header source address is compared to the embedded IPv4 address in the IPv6 source address. The packet is discarded if there is no match. If there is a match, the IPv6 packet is forwarded as a native IPv6 packet to the IPv6 destination address on the CE LAN.

CE to IPv6 Internet

In a CE to IPv6 Internet scenario, IPv6 traffic is received natively on a customer-facing interface. The IPv6 destination address does not fall into the range of the locally configured 6rd prefix, which means that it is not targeted at a destination inside the local 6rd domain. In this case, the packet needs to be forwarded to a 6rd BR router.

As in the CE to CE scenario, the IPv6 packet is encapsulated in an IPv4 header; however, the difference is that the locally configured BR IPv4 address is then copied to the IPv4 destination address. In addition, the locally configured tunnel source is copied to the IPv4 source address. The protocol field is set to 41 (IPv6 in IPv4), and the encapsulated packet is then forwarded to the BR router through the IPv4 domain following the normal IPv4 routing table.

The BR router receives the IPv4 packet and removes the IPv4 header encapsulation. The IPv4 header source address is compared to the IPv4 address embedded in the IPv6 source address. The packet is discarded if the addresses do not match. Otherwise, the IPv6 packet is forwarded natively to the IPv6 destination address.

IPv6 Internet to CE

In the IPv6 Internet to CE scenario, the BR router receives a native IPv6 packet from one of its IPv6 network-facing interfaces. The IPv6 destination address falls into the range of the locally configured 6rd prefix, which means that it is targeted at a destination inside the local 6rd domain. In this case, the packet needs to be forwarded to the appropriate CE router (Figure 9).

Figure 9. IPv6 Internet to CE

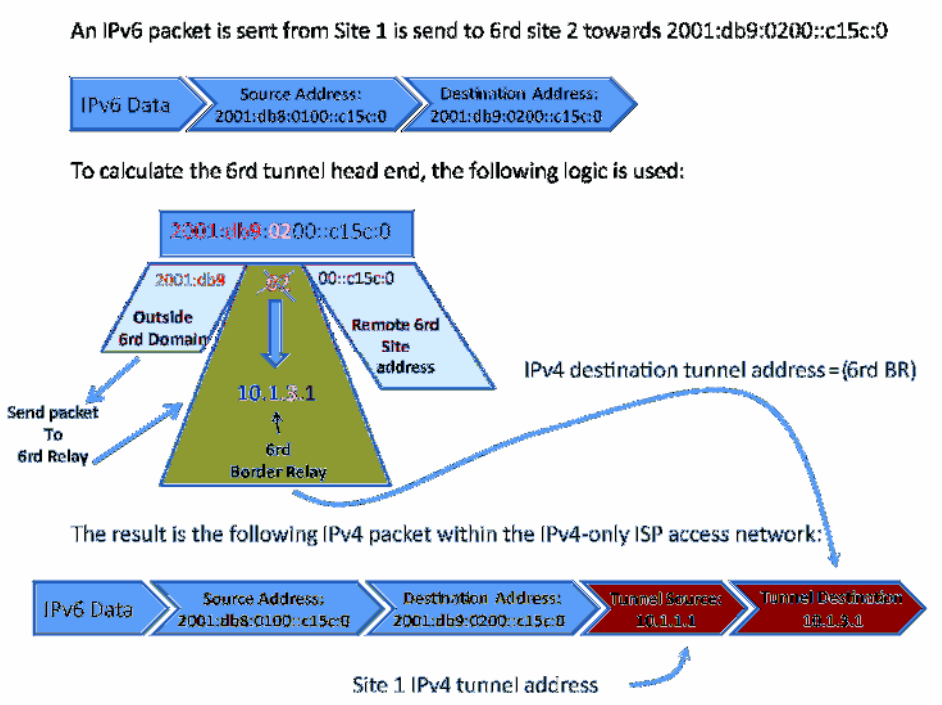

The 6rd BR router will encapsulate the IPv6 packet in an IPv4 header. The embedded IPv4 address within the IPv6 header is used as the IPv4 destination address. The source address of the tunnel will be the IPv4 address configured on the 6rd tunnel interface on the 6rd BR router. The protocol field of the IPv4 packet within the header is set to 41 (IPv6 in IPv4). The packet is then forwarded by the 6rd BR router to the 6rd CE router using the traditional IPv4 forwarding path.

When the CE router receives the tunneled 6rd IPv4 packet, the IPv4 header is removed to expose the encapsulated IPv6 packet. First, the IPv4 header source address is compared to the locally configured 6rd BR IPv4 address. The packet is discarded if the addresses do not match. If there is a match, then the IPv6 packet is forwarded natively over IPv6 to the IPv6 destination address by the 6rd CE device.

Provisioning

The 6rd CE router–delegated IPv6 address is derived from the local WAN interface IPv4 address in combination with the 6rd prefix. It is possible to configure the address manually; however, often the IPv4 address is received through DHCPv4 from the connected service provider.

In addition to this IPv4 address, additional information must be configured.

- **6rd prefix:** This is the common prefix that is used throughout the 6rd domain and that determines whether a packet is targeted at a destination inside or outside the 6rd domain.
- **6rd prefix length:** This parameter provides information about the value bits in the 6rd prefix.
- **IPv4 mask length:** This parameter specifies how many bits are common in all 6rd CE and BR router IPv4 addresses, and therefore can be compressed within the 6rd prefix for the site. Within Cisco IOS Software, it is possible to compress bits from both the IPv4 prefix and suffix.
- **BR IPv4 address:** This parameter specifies the IPv4 address to use as a tunnel destination or source address for 6rd tunnels.

Note that these parameters can be manually configured or acquired through a variety of provisioning techniques or tools, such as DHCP.

6rd Basic Configuration Guidelines

Interface Command Line

[no] tunnel mode ipv6ip 6rd

- This command can be configured only on a tunnel interface.
- This command specifies that the tunnel is to be used for 6rd.

[no] tunnel 6rd prefix <ipv6 address>/<length>

- This command is mandatory for the operation of 6rd.
- This command specifies the common IPv6 prefix.
- The length here indicates the position of the IPv4 address in the 6rd delegated prefix and payload destination.
- The tunnel line state of a 6rd tunnel will stay down until this command is configured.
- Configuring a prefix length of 0 is equivalent to removing this command.

[no] tunnel 6rd ipv4 {prefix-len<length>} {suffix-len<length>}

- This command is optional for the operation of 6rd.
- This command specifies the most significant bits and least significant bits of the IPv4 transport address (that is, the tunnel source) that are common to all the 6rd routers in a domain.
- The valid range is from 0 to 31. The sum of the IPv4 prefix length and the IPv4 suffix length cannot exceed 31.
- If this command is not configured and the 6rd IPv6 prefix is configured, the system will use 0 as the default.

[no] tunnel 6rd br<ipv4-address>

By default, at a 6rd router all incoming packets from the 6rd domain will require their outer IPv4 source address to be embedded in the 6rd encoded IPv6 source address. Packets not meeting this criterion will be dropped. Configuring this command allows packets with the specified source to be exempt from this check. This command is needed on the CE router because packets arriving at the CE router from the BR router will typically be traffic from a native IPv6 host, which is not required to have a 6rd encoded source address.

Show Command Line

show tunnel 6rd { <interface> | <cr> }

This command can be used to display 6rd-related information about a tunnel. If an interface is not specified, information about all the 6rd tunnels on the router will be displayed.

show tunnel 6rd destination <ipv6-prefix><tunnel interface>

This command translates a 6rd IPv6 prefix to the corresponding IPv4 endpoint. It uses the tunnel 6rd parameters configured on the specified tunnel to compute the IPv4 endpoint.

show tunnel 6rd prefix <ipv4-destination><tunnel interface>

This command translates an IPv4 endpoint address to the corresponding 6rd IPv6 prefix. It uses the tunnel 6rd parameters configured on the specified tunnel to compute the IPv6 prefix.

IPv6 General Prefix

ipv6 general-prefix <name> 6rd <tunnel interface>

This command creates a general prefix that can be referred to by <name> based on the supplied tunnel interface.

Configuration Examples

Both the BR and CE routers have a shared common 6rd delegated prefix. You should configure the IPv6 subnet router anycast for this prefix to assist with troubleshooting.

The IPv6 general prefix is used to represent an abstraction of the 6rd delegated prefix. This configuration will help in provisioning through a variety of options or DHCP.

Basic Configuration Examples

BR Router

```
ipv6 general-prefix DELEGATED_PREFIX 6rd Tunnel0
interface Loopback0
ip address 10.0.0.1 255.255.255.0
!
interface Tunnel0
tunnel source Loopback0
tunnel mode ipv6ip 6rd
tunnel6rd ipv4 prefix-len 8
tunnel6rd prefix 2001:db80::/32
ipv6 address DELEGATED_PREFIX::/128 anycast
!
```

```
ipv6 route 2001:db80::/32 Tunnel0
ipv6 route ::/0 2001:babe::1
```

CE Router

```
ipv6 general-prefix DELEGATED_PREFIX6rd Tunnel0
interface Dialer0
ip address dhcp  ! (10.1.1.1)
!
interface Tunnel0
tunnel source Dialer0
tunnel mode ipv6ip 6rd
tunnel 6rd ipv4 prefix-len 8
tunnel 6rd prefix 2001:db80::/32
tunnel 6rd br 10.1.3.1
ipv6 address DELEGATED_PREFIX ::/128 anycast
!
interface Ethernet0
ipv6 address DELEGATED_PREFIX ::/64 eui-64
!
ipv6 route 2001:db80::/28 Tunnel0
ipv6 route ::/0 Tunnel0, 2001:db80:a000:0010::
ipv6 route 2001:db80:0:A00::/56 Null0
```

Provisioning with DHCPv4

Currently, the Cisco IOS Software 6rd CE implementation does not support DHCPv4 option 212 for 6rd provisioning. The Cisco IOS Software DHCP server, however, supports the option as a binary string and so does Cisco® Network Registrar.

The 6rd option in DHCP is coded as follows, following RFC 5969:

Figure 10. DHCP option for 6rd

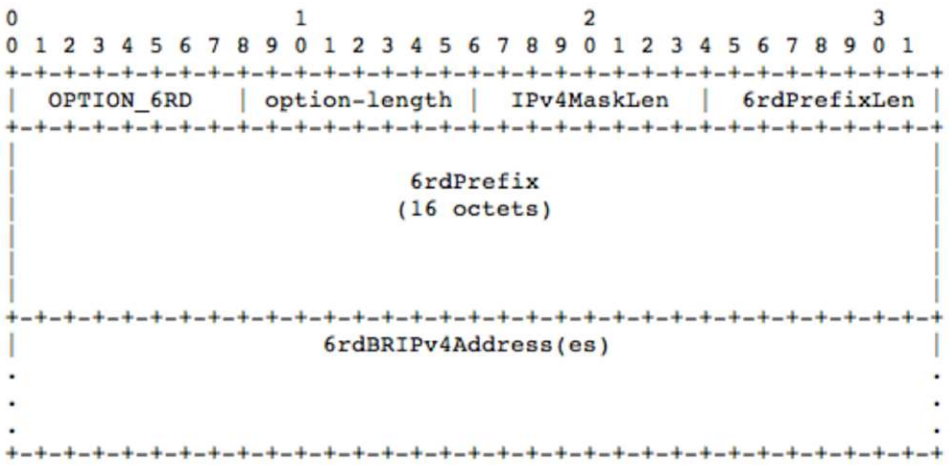

Here, the following fields are valid:

option-code

The OPTION_6RD is set to value 212

option-length:

The length of the DHCP option in octets (22 octets with one BR IPv4 address)

IPv4MaskLen:

The number of high-order bits that are identical across all CE IPv4 addresses within a given 6rd domain; this may be any value between 0 and 32 (any value greater than 32 is invalid)

6rdPrefixLen:

The IPv6 prefix length of the service provider's 6rd IPv6 prefix in number of bits; for the purpose of bounds checking by DHCP option processing, the sum of (32 – IPv4MaskLen) + 6rdPrefixLen must be less than or equal to 128

6rdBRIPv4Address:

One or more IPv4 addresses of the 6rd BR routers for a given 6rd domain

These settings lead to the following configuration on the Cisco IOS Software DHCPv4 server:

```
!
ipdhcp pool DHCP_UT610
host 192.168.0.61 255.255.255.0
client-identifier 0100.226b.68a8.5b
dns-server 192.168.0.4
default-router 192.168.0.1
domain-name foo.org
option212 hex 0009.1812.1614.3420.0106.f814.68f0.0000.0000.0000.0000.00c0.a806.01
!
```

Operational Considerations

Maximum Transmission Unit and Fragmentation

The maximum transmission unit (MTU) in a 6rd domain must be well managed. Ideally, the MTU should be the same on all links; however, ensuring a common MTU across the 6rd domain will not prevent other MTU issues along the end-to-end path. The combination of an anycast IPv4 source address and fragmentation is not recommended and must be avoided. If 6rd BR router reliability is ensured with IPv4 anycast, then the IPv4 do-not-fragment bit (DF-bit) bit must be set. Otherwise, incorrect reassembly on the CE routers may occur.

Both the CE and BR routers should support IPv6 path MTU discovery (PMTU), but if the IPv4 MTU is the same across the 6rd domain, then support for IPv4 PMTU discovery for the tunnel is not required. The typical 6rd IPv6 MTU is 1480 bytes (1500 bytes—20 bytes, due to the tunnel encapsulation overhead).

6rd and 6to4

6rd is a generalization of the automatic 6to4 tunneling mechanism (RFC 3056). It overcomes what is seen as the biggest shortcoming of the automatic 6to4 tunneling mechanism: the use of the well-known fixed prefix 2002::/16, for all 6to4 sites. This IPv6 prefix is injected by many routers (also known as 6to4 relay routers) on the IPv6 Internet. As direct consequences, traffic flows may be asymmetric and service providers have no control over the 6to4 relay of the return path, and sites have to renumber when native IPv6 becomes available. The 6rd

mechanism removes these disadvantages by allowing each service provider to use a unique IPv6 prefix for each customer and hence help ensure that no additional untrusted third-party relays are required.

Platform Support

6rd is supported on a variety of Cisco products. Use Cisco Feature Navigator to identify 6rd support for your platform. Cisco Feature Navigator can be found at http://www.cisco.com/go/fn/.

Security

Unlike 6to4, through which the CE router can virtually receive encapsulated IPv6 traffic from hundreds of relay routers located anywhere on the IPv6 Internet, 6rd exchanges encapsulated traffic only from BR routers located within the 6rd domain infrastructure. This restriction reduces the potential for spoofing attacks from external sources on the Internet.

Conclusion

When a service provider has an access network with full IPv4 support but for which the quick addition of IPv6 technology is impractical or too resource intensive, the service provider can use 6rd. 6rd is tailored for this scenario and simple to implement and has a proven field deployment track record.

For More Information

- http://www.cisco.com/go/ipv6
- http://www.cisco.com/en/US/prod/collateral/iosswrel/ps6537/ps6553/white_paper_c11-558744-00_ns1017_Networking_Solutions_White_Paper.html

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-665758-00 10/11