

Internet Engineering Task Force (IETF)
Request for Comments: 5969
Category: Standards Track
ISSN: 2070-1721

W. Townsley
O. Troan
Cisco
August 2010

IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) --
Protocol Specification

Abstract

This document specifies an automatic tunneling mechanism tailored to advance deployment of IPv6 to end users via a service provider's IPv4 network infrastructure. Key aspects include automatic IPv6 prefix delegation to sites, stateless operation, simple provisioning, and service, which is equivalent to native IPv6 at the sites that are served by the mechanism.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5969>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 2
- 2. Requirements Language 4
- 3. Terminology 4
- 4. 6rd Prefix Delegation 5
- 5. Troubleshooting and Traceability 7
- 6. Address Selection 7
- 7. 6rd Configuration 7
 - 7.1. Customer Edge Configuration 8
 - 7.1.1. 6rd DHCPv4 Option 9
 - 7.2. Border Relay Configuration 10
- 8. Neighbor Unreachability Detection 11
- 9. IPv6 in IPv4 Encapsulation 12
 - 9.1. Maximum Transmission Unit 13
 - 9.2. Receiving Rules 13
- 10. Transition Considerations 14
- 11. IPv6 Address Space Usage 14
- 12. Security Considerations 15
- 13. IANA Considerations 16
- 14. Acknowledgements 16
- 15. References 16
 - 15.1. Normative References 16
 - 15.2. Informative References 17

1. Introduction

The original idea and the name of the mechanism (6rd) described in [RFC5569] details a successful commercial "rapid deployment" of the 6rd mechanism by a residential service provider and is recommended reading. This document describes the 6rd mechanism, which has been extended for use in more general environments. Throughout this document, the term 6to4 is used to refer to the mechanism described in [RFC3056] and 6rd is the mechanism defined herein.

6rd specifies a protocol mechanism to deploy IPv6 to sites via a service provider's (SP's) IPv4 network. It builds on 6to4 [RFC3056], with the key differentiator that it utilizes an SP's own IPv6 address prefix rather than a well-known prefix (2002::/16). By using the SP's IPv6 prefix, the operational domain of 6rd is limited to the SP network and is under its direct control. From the perspective of customer sites and the IPv6 Internet at large, the IPv6 service provided is equivalent to native IPv6.

The 6rd mechanism relies upon an algorithmic mapping between the IPv6 and IPv4 addresses that are assigned for use within the SP network. This mapping allows for automatic determination of IPv4 tunnel endpoints from IPv6 prefixes, allowing stateless operation of 6rd.

6rd views the IPv4 network as a link layer for IPv6 and supports an automatic tunneling abstraction similar to the Non-Broadcast Multiple Access (NBMA) [RFC2491] model.

A 6rd domain consists of 6rd Customer Edge (CE) routers and one or more 6rd Border Relays (BRs). IPv6 packets encapsulated by 6rd follow the IPv4 routing topology within the SP network among CEs and BRs. 6rd BRs are traversed only for IPv6 packets that are destined to or are arriving from outside the SP's 6rd domain. As 6rd is stateless, BRs may be reached using anycast for failover and resiliency (in a similar fashion to [RFC3068]).

On the "customer-facing" (i.e., "LAN") side of a CE, IPv6 is implemented as it would be for any native IP service delivered by the SP, and further considerations for IPv6 operation on the LAN side of the CE is out of scope for this document. On the "SP-facing" (i.e., "WAN") side of the 6rd CE, the WAN interface itself, encapsulation over Ethernet, ATM or PPP, as well as control protocols such as PPPoE, IPCP, DHCP, etc. all remain unchanged from current IPv4 operation. Although 6rd was designed primarily to support IPv6 deployment to a customer site (such as a residential home network) by an SP, it can equally be applied to an individual IPv6 host acting as a CE.

6rd relies on IPv4 and is designed to deliver production-quality IPv6 alongside IPv4 with as little change to IPv4 networking and operations as possible. Native IPv6 deployment within the SP network itself may continue for the SP's own purposes while delivering IPv6 service to sites supported by 6rd. Once the SP network and operations can support fully native IPv6 access and transport, 6rd may be discontinued.

6rd utilizes the same encapsulation and base mechanism as 6to4 and could be viewed as a superset of 6to4 (6to4 could be achieved by setting the 6rd prefix to 2002::/16). Unlike 6to4, 6rd is for use only in an environment where a service provider closely manages the delivery of IPv6 service. 6to4 routes with the 2002::/16 prefix may exist alongside 6rd in the 6rd CE router, and doing so may offer some efficiencies when communicating directly with 6to4 routers.

The 6rd link model can be extended to support IPv6 multicast. IPv6 multicast support is left for future consideration.

How this mechanism should be used and other deployment and operational considerations are considered out of scope for this document.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

3. Terminology

6rd prefix An IPv6 prefix selected by the service provider for use by a 6rd domain. There is exactly one 6rd prefix for a given 6rd domain. An SP may deploy 6rd with a single 6rd domain or multiple 6rd domains.

6rd Customer Edge (6rd CE) A device functioning as a Customer Edge router in a 6rd deployment. In a residential broadband deployment, this type of device is sometimes referred to as a "Residential Gateway" (RG) or "Customer Premises Equipment" (CPE). A typical 6rd CE serving a residential site has one WAN side interface, one or more LAN side interfaces, and a 6rd virtual interface. A 6rd CE may also be referred to simply as a "CE" within the context of 6rd.

6rd delegated prefix The IPv6 prefix calculated by the CE for use within the customer site by combining the 6rd prefix and the CE IPv4 address obtained via IPv4 configuration methods. This prefix can be considered logically equivalent to a DHCPv6 IPv6 delegated prefix [RFC3633].

6rd domain A set of 6rd CEs and BRs connected to the same virtual 6rd link. A service provider may deploy 6rd with a single 6rd domain, or may utilize multiple 6rd domains. Each domain requires a separate 6rd prefix.

CE LAN side The functionality of a 6rd CE that serves the "Local Area Network (LAN)" or "customer-facing" side of the CE. The CE LAN side interface is fully IPv6 enabled.

CE WAN side	The functionality of a 6rd CE that serves the "Wide Area Network (WAN)" or "Service Provider-facing" side of the CE. The CE WAN side is IPv4-only.
6rd Border Relay (BR)	A 6rd-enabled router managed by the service provider at the edge of a 6rd domain. A Border Relay router has at least one of each of the following: an IPv4-enabled interface, a 6rd virtual interface acting as an endpoint for the 6rd IPv6 in IPv4 tunnel, and an IPv6 interface connected to the native IPv6 network. A 6rd BR may also be referred to simply as a "BR" within the context of 6rd.
BR IPv4 address	The IPv4 address of the 6rd Border Relay for a given 6rd domain. This IPv4 address is used by the CE to send packets to a BR in order to reach IPv6 destinations outside of the 6rd domain.
6rd virtual interface	Internal multi-point tunnel interface where 6rd encapsulation and decapsulation of IPv6 packets inside IPv4 occurs. A typical CE or BR implementation requires only one 6rd virtual interface. A BR operating in multiple 6rd domains may require more than one 6rd virtual interface, but no more than one per 6rd domain.
CE IPv4 address	The IPv4 address given to the CE as part of normal IPv4 Internet access (i.e., configured via DHCP, PPP, or otherwise). This address may be global or private [RFC1918] within the 6rd domain. This address is used by a 6rd CE to create the 6rd delegated prefix as well as to send and receive IPv4-encapsulated IPv6 packets.

4. 6rd Prefix Delegation

The 6rd delegated prefix for use at a customer site is created by combining the 6rd prefix and all or part of the CE IPv4 address. From these elements, the 6rd delegated prefix is automatically created by the CE for the customer site when IPv4 service is obtained. This 6rd delegated prefix is used in the same manner as a prefix obtained via DHCPv6 prefix delegation [RFC3633].

In 6to4, a similar operation is performed by incorporating an entire IPv4 address at a fixed location following a well-known /16 IPv6 prefix. In 6rd, the IPv6 prefix as well as the position and number of bits of the IPv4 address incorporated varies from one 6rd domain to the next. 6rd allows the SP to adjust the size of the 6rd prefix, how many bits are used by the 6rd mechanism, and how many bits are left to be delegated to customer sites. To allow for stateless address auto-configuration on the CE LAN side, a 6rd delegated prefix SHOULD be /64 or shorter.

The 6rd delegated prefix is created by concatenating the 6rd prefix and a consecutive set of bits from the CE IPv4 address in order. The length of the 6rd delegated prefix is equal to length of the 6rd prefix (n) plus the number of bits from the CE IPv4 address (o).

The figure shows the format of an IPv6 address (Section 2.5.4 of [RFC4291]) with a 6rd prefix and an embedded CE IPv4 address:




Figure 1

For example, if the 6rd prefix is /32 and 24 bits of the CE IPv4 address is used (e.g., all CE IPv4 addresses can be aggregated by a 10.0.0.0/8), then the size of the 6rd delegated prefix for each CE is automatically calculated to be /56 (32 + 24 = 56).

Embedding less than the full 32 bits of a CE IPv4 address is possible only when an aggregated block of IPv4 addresses is available for a given 6rd domain. This may not be practical with global IPv4 addresses, but is quite likely in a deployment where private addresses are being assigned to CEs. If private addresses overlap within a given 6rd deployment, the deployment may be divided into separate 6rd domains, likely along the same topology lines the NAT-based IPv4 deployment itself would require. In this case, each domain is addressed with a different 6rd prefix.

Each 6rd domain may use a different encoding of the embedded IPv4 address, even within the same service provider. For example, if multiple IPv4 address blocks with different levels of aggregation are used at the same service provider, the number of IPv4 bits needed to encode the 6rd delegated prefix may vary between each block. In this case, different 6rd prefixes, and hence separate 6rd domains, may be used to support the different encodings.

Since 6rd delegated prefixes are selected algorithmically from an IPv4 address, changing the IPv4 address will cause a change in the IPv6 delegated prefix which would ripple through the site's network and could be disruptive. As such, it is recommended that the service provider assign CE IPv4 addresses with relatively long lifetimes.

6rd IPv6 address assignment, and hence the IPv6 service itself, is tied to the IPv4 address lease; thus, the 6rd service is also tied to this in terms of authorization, accounting, etc. For example, the 6rd delegated prefix has the same lifetime as its associated IPv4 address. The prefix lifetimes advertised in Router Advertisements or used by DHCP on the CE LAN side MUST be equal to or shorter than the IPv4 address lease time. If the IPv4 lease time is not known, the lifetime of the 6rd delegated prefix SHOULD follow the defaults specified in [RFC4861].

5. Troubleshooting and Traceability

A 6rd IPv6 address and associated IPv4 address for a given customer can always be determined algorithmically by the service provider that operates the given 6rd domain. This may be useful for referencing logs and other data at a service provider that may have more robust operational tools for IPv4 than IPv6. This also allows IPv4 data path, node, and endpoint monitoring to be applicable to IPv6.

The 6rd CE and BR SHOULD support the IPv6 Subnet-Router anycast address [RFC4291] for its own 6rd delegated prefix. This allows, for example, IPv6 ICMP echo messages to be sent to the 6rd virtual interface itself for additional troubleshooting of the internal operation of 6rd at a given CE or BR. In the case of the BR, the IPv4 address used to calculate the 6rd delegated prefix is the configured BR IPv4 address.

6. Address Selection

All addresses assigned from 6rd delegated prefixes should be treated as native IPv6. No changes to the source address selection or destination address selection policy table [RFC3484] are necessary.

7. 6rd Configuration

For a given 6rd domain, the BR and CE MUST be configured with the following four 6rd elements. The configured values for these four 6rd elements are identical for all CEs and BRs within a given 6rd domain.

IPv4MaskLen	The number of high-order bits that are identical across all CE IPv4 addresses within a given 6rd domain. For example, if there are no identical bits, IPv4MaskLen is 0 and the entire CE IPv4 address is used to create the 6rd delegated prefix. If there are 8 identical bits (e.g., the Private IPv4 address range 10.0.0.0/8 is being used), IPv4MaskLen is equal to 8 and IPv4MaskLen high-order bits are stripped from the IPv4 address before constructing the corresponding 6rd delegated prefix.
6rdPrefix	The 6rd IPv6 prefix for the given 6rd domain.
6rdPrefixLen	The length of the 6rd IPv6 prefix for the given 6rd domain.
6rdBRIPv4Address	The IPv4 address of the 6rd Border Relay for a given 6rd domain.

7.1. Customer Edge Configuration

The four 6rd elements are set to values that are the same across all CEs within a 6rd domain. The values may be configured in a variety of manners, including provisioning methods such as the Broadband Forum's "TR-69" [TR069] Residential Gateway management interface, an XML-based object retrieved after IPv4 connectivity is established, a DNS record, an SMIPv2 MIB [RFC2578], PPP IPCP, or manual configuration by an administrator. This document describes how to configure the necessary parameters via a single DHCP option. A CE that allows IPv4 configuration by DHCP SHOULD implement this option. Other configuration and management methods may use the format described by this option for consistency and convenience of implementation on CEs that support multiple configuration methods.

The only remaining provisioning information the CE requires in order to calculate the 6rd delegated prefix and enable IPv6 connectivity is an IPv4 address for the CE. This CE IPv4 address is configured as part of obtaining IPv4 Internet access (i.e., configured via DHCP, PPP, or otherwise). This address may be global or private [RFC1918] within the 6rd domain.

A single 6rd CE MAY be connected to more than one 6rd domain, just as any router may have more than one IPv6-enabled service provider facing interface and more than one set of associated delegated prefixes assigned by DHCPv6 prefix delegation or other means. Each

domain a given CE operates within would require its own set of 6rd configuration elements and would generate its own 6rd delegated prefix.

7.1.1.1. 6rd DHCPv4 Option




Figure 2

option-code	OPTION_6RD (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 SP'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 Border Relay(s) for a given 6rd domain.

6rdPrefix The service provider's 6rd IPv6 prefix represented as a 16-octet IPv6 address. The bits in the prefix after the 6rdPrefixLen number of bits are reserved and MUST be initialized to zero by the sender and ignored by the receiver.

The CE MUST include a Parameter Request List Option [RFC2132] for the OPTION_6RD. Because the OPTION_6RD contains one IPv4MaskLen/6rdPrefixLen/6rdPrefix block, and because DHCP cannot convey more than one instance of an option, OPTION_6RD is limited to provision at most a single 6rd domain. Provisioning of a CE router connected to multiple 6rd domains is outside the scope of this protocol specification.

The presence of the OPTION_6RD DHCP option is an indication of the availability of the 6rd service. By default, receipt of a valid 6rd DHCP option by a 6rd-capable CE results in configuration of the 6rd virtual interface and associated delegated prefix for use on the CE's LAN side. The CE MUST be able to configure the 6rd mechanism to be disabled, in which case the 6rd DHCP option, if received, is silently ignored.

A detailed description of CE behavior using multiple BR IPv4 addresses is left for future consideration. In such a case, a CE MUST support at least one BR IPv4 address and MAY support more than one.

When 6rd is enabled, a typical CE router will install a default route to the BR, a black hole route for the 6rd delegated prefix, and routes for any LAN side assigned and advertised prefixes. For example, using a CE IPv4 address of 10.100.100.1, a BR IPv4 address of 10.0.0.1, an IPv4MaskLen of 8, 2001:db8::/32 as the 6rdPrefix, and one /64 prefix assigned to a LAN side interface, a typical CE routing table will look like:

```
::/0 -> 6rd-virtual-int0 via 2001:db8:0:100:: (default route)
2001:db8::/32 -> 6rd-virtual-int0 (direct connect to 6rd)
2001:db8:6464:100::/56 -> Null0 (delegated prefix null route)
2001:db8:6464:100::/64 -> Ethernet0 (LAN interface)
```

7.2. Border Relay Configuration

The 6rd BR MUST be configured with the same 6rd elements as the 6rd CEs operating within the same domain.

For increased reliability and load balancing, the BR IPv4 address may be an anycast address shared across a given 6rd domain. As 6rd is stateless, any BR may be used at any time. If the BR IPv4 address is

anycast the relay MUST use this anycast IPv4 address as the source address in packets relayed to CEs.

Since 6rd uses provider address space, no specific routes need to be advertised externally for 6rd to operate, neither in IPv6 nor IPv4 BGP. However, if anycast is used for the 6rd IPv4 relays, the anycast addresses must be advertised in the service provider's IGP.

8. Neighbor Unreachability Detection

Neighbor Unreachability Detection (NUD) for tunnels is described in Section 3.8 of [RFC4213]. In 6rd, all CEs and BRs can be considered as connected to the same virtual link and therefore neighbors to each other. This section describes how to utilize neighbor unreachability detection without negatively impacting the scalability of a 6rd deployment.

A typical 6rd deployment may consist of a very large number of CEs within the same domain. Reachability between CEs is based on IPv4 routing, and sending NUD or any periodic packets between 6rd CE devices beyond isolated troubleshooting of the 6rd mechanism is NOT RECOMMENDED.

While reachability detection between a given 6rd CE and BR is not necessary for the proper operation of 6rd, in cases where a CE has alternate paths for BR reachability to choose from, it could be useful. Sending NUD messages to a BR, in particular periodic messages from a very large number of CEs, could result in overloading of the BR control message processing path, negatively affecting scalability of the 6rd deployment. Instead, a CE that needs to determine BR reachability MUST utilize a method that allows reachability detection packets to follow a typical data forwarding path without special processing by the BR. One such method is described below.

1. The CE constructs a payload of any size and content to be sent to the BR (e.g., a zero-length null payload, a padded payload designed to test a certain MTU, a NUD message, etc.). The exact format of the message payload is not important as the BR will not be processing it directly.
2. The desired payload is encapsulated with the inner IPv6 and outer IPv4 headers as follows:
 - * The IPv6 destination address is set to an address from the CE's 6rd delegated prefix that is assigned to a virtual interface on the CE.

- * The IPv6 source address is set to an address from the CE's 6rd delegated prefix as well, including the same as used for the IPv6 destination address.
 - * The IPv4 header is then added as it normally would be for any packet destined for the BR. That is, the IPv4 destination address is that of the BR, and the source address is the CE IPv4 address.
3. The CE sends the constructed packet out the interface on which BR reachability is being monitored. On successful receipt at the BR, the BR MUST decapsulate and forward the packet normally. That is, the IPv4 header is decapsulated normally, revealing the IPv6 destination as the CE, which in turn results in the packet being forwarded to that CE via the 6rd mechanism (i.e., the IPv4 destination is that of the CE that originated the packet, and the IPv4 source is that of the BR).
 4. Arrival of the constructed IPv6 packet at the CE's IPv6 address completes one round trip to and from the BR, without causing the BR to process the message outside of its normal data forwarding path. The CE then processes the IPv6 packet accordingly (updating keepalive timers, metrics, etc.).

The payload may be empty or could contain values that are meaningful to the CE. Sending a proper NUD message could be convenient for some implementations (note that the BR will decrement the IPv6 hop limit). Since the BR forwards the packet as any other data packet without any processing of the payload itself, the format of the payload is left as a choice to the implementer.

9. IPv6 in IPv4 Encapsulation

IPv6 in IPv4 encapsulation and forwarding manipulations (e.g., handling packet markings, checksumming, etc.) is performed as specified in Section 3.5 of "Basic Transition Mechanisms for IPv6 Hosts and Routers" [RFC4213], which is the same mechanism used by 6to4 [RFC3056]. ICMPv4 errors are handled as specified in Section 3.4 of [RFC4213]. By default, the IPv6 Traffic Class field MUST be copied to the IPv4 ToS (Type of Service) field. This default behavior MAY be overridden by configuration. See [RFC2983] and [RFC3168] for further information related to IP Differentiated Services and tunneling.

IPv6 packets from a CE are encapsulated in IPv4 packets when they leave the site via its CE WAN side interface. The CE IPv4 address MUST be configured to send and receive packets on this interface.

The 6rd link is modeled as an NBMA link similar to other automatic IPv6 in IPv4 tunneling mechanisms like [RFC5214], with all 6rd CEs and BRs defined as off-link neighbors from one other. The link-local address of a 6rd virtual interface performing the 6rd encapsulation would, if needed, be formed as described in Section 3.7 of [RFC4213]. However, no communication using link-local addresses will occur.

9.1. Maximum Transmission Unit

Maximum transmission unit (MTU) and fragmentation issues for IPv6 in IPv4 tunneling are discussed in detail in Section 3.2 of RFC 4213 [RFC4213]. 6rd's scope is limited to a service provider network. IPv4 Path MTU discovery MAY be used to adjust the MTU of the tunnel as described in Section 3.2.2 of RFC 4213 [RFC4213], or the 6rd Tunnel MTU might be explicitly configured.

The use of an anycast source address could lead to any ICMP error message generated on the path being sent to a different BR. Therefore, using dynamic tunnel MTU Section 3.2.2 of [RFC4213] is subject to IPv4 Path MTU blackholes.

Multiple BRs using the same anycast source address could send fragmented packets to the same IPv6 CE at the same time. If the fragmented packets from different BRs happen to use the same fragment ID, incorrect reassembly might occur. For this reason, a BR using an anycast source address MUST set the IPv4 Don't Fragment flag.

If the MTU is well-managed such that the IPv4 MTU on the CE WAN side interface is set so that no fragmentation occurs within the boundary of the SP, then the 6rd Tunnel MTU should be set to the known IPv4 MTU minus the size of the encapsulating IPv4 header (20 bytes). For example, if the IPv4 MTU is known to be 1500 bytes, the 6rd Tunnel MTU might be set to 1480 bytes. Absent more specific information, the 6rd Tunnel MTU SHOULD default to 1280 bytes.

9.2. Receiving Rules

In order to prevent spoofing of IPv6 addresses, the 6rd BR and CE MUST validate the embedded IPv4 source address of the encapsulated IPv6 packet with the IPv4 source address it is encapsulated by according to the configured parameters of the 6rd domain. If the two source addresses do not match, the packet MUST be dropped and a counter incremented to indicate that a potential spoofing attack may be underway. Additionally, a CE MUST allow forwarding of packets sourced by the configured BR IPv4 address.

By default, the CE router MUST drop packets received on the 6rd virtual interface (i.e., after decapsulation of IPv4) for IPv6 destinations not within its own 6rd delegated prefix.

10. Transition Considerations

An SP network can migrate to IPv6 at its own pace with little or no effect on customers being provided IPv6 via 6rd. When native IPv6 connectivity is available, an administrator can choose to disable 6rd.

The SP can choose to provision a separate IPv6 address block for native service, or reuse the 6rd prefix block itself. If the SP uses a separate address block, moving from 6rd to native IPv6 is seen as a normal IPv6 renumbering event for the customer. Renumbering may also be avoided by injecting the 6rd delegated prefix into the SP's IPv6 routing domain. Further considerations with regards to transitioning from 6rd to native IPv6 are not covered in this protocol specification.

11. IPv6 Address Space Usage

As 6rd uses service-provider address space, 6rd uses the normal address delegation a service provider gets from its Regional Internet Registry (RIR) and no global allocation of a single 6rd IANA-assigned address block like the 6to4 2002::/16 is needed.

The service provider's prefix must be short enough to encode the unique bits of all IPv4 addresses within a given 6rd domain and still provide enough IPv6 address space to the residential site. Assuming a worst case scenario using the full 32 bits for the IPv4 address, assigning a /56 for customer sites would mean that each service provider using 6rd would require a /24 for 6rd in addition to other IPv6 addressing needs. Assuming that 6rd would be stunningly successful and taken up by almost all Autonomous System (AS) number holders (32K today), then the total address usage of 6rd would be equivalent to a /9. If the SP instead delegated /60s to sites, the service provider would require a /28, and the total global address consumption by 6rd would be equivalent to a /13. Again, this assumes that 6rd is used by all AS number holders in the IPv4 Internet today at the same time, that none have used any of 6rd's address compression techniques, and that none have moved to native IPv6 and reclaimed the 6rd space that was being used for other purposes.

To alleviate concerns about address usage, 6rd allows for leaving out redundant IPv4 prefix bits in the encoding of the IPv4 address inside the 6rd IPv6 address. This is most useful where the IPv4 address space is very well aggregated. For example, to provide each customer

with a /60, if a service provider has all its IPv4 customers under a /12 then only 20 bits needs to be used to encode the IPv4 address and the service provider would only need a /40 IPv6 allocation for 6rd. If private address space is used, then a 10/8 would require a /36. If multiple 10/8 domains are used, then up to 16 could be supported within a /32.

If a service provider has a non-aggregatable IPv4 space and requiring the use of the full 32-bit IPv4 address in the encoding of the 6rd IPv6 address, the 6rd prefix MUST be no longer than /32 in order to offer a 6rd delegated prefix of at least /64.

The 6rd address block can be reclaimed when all users of it have transitioned to native IPv6 service. This may require renumbering of customer sites and use of additional address space during the transition period.

12. Security Considerations

A 6to4 relay router as specified in [RFC3056] can be used as an open relay. It can be used to relay IPv6 traffic and as a traffic anonymizer. By restricting the 6rd domain to within a provider network, a CE only needs to accept packets from a single or small set of known 6rd BR IPv4 addresses. As such, many of the threats against 6to4 as described in [RFC3964] do not apply.

When applying the receiving rules in Section 9.2, IPv6 packets are as well protected against spoofing as IPv4 packets are within an SP network.

A malicious user that is aware of a 6rd domain and the BR IPv4 address could use this information to construct a packet that would cause a Border Relay to reflect tunneled packets outside of the domain that it is serving. If the attacker constructs the packet accordingly and can inject a packet with an IPv6 source address that looks as if it originates from within another 6rd domain, forwarding loops between 6rd domains may be created, allowing the malicious user to launch a packet amplification attack between 6rd domains [RoutingLoop].

One possible mitigation for this is to simply not allow the BR IPv4 address to be reachable from outside the SP's 6rd domain. In this case, carefully constructed IPv6 packets still could be reflected off a single BR, but the looping condition will not occur. Tunneled packets with the BR IPv4 address as the source address might also be filtered to prohibit 6rd tunnels from exiting the 6rd domain.

To avoid forwarding loops via other internal relays, the BR should employ outgoing and incoming IPv4 packets filters, filtering out all known relay addresses for internal 6rd BRs, ISATAP routers, or 6to4 relays, including the well-known anycast address space for 6to4.

Another possible mitigation to the routing loop issue is described in [V6OPS-LOOPS].

The BR MUST install a null route [RFC4632] for its 6rd delegated prefix created based on its BR IPv4 address, with the exception of the IPv6 Subnet-Router anycast address.

13. IANA Considerations

IANA assigned a new DHCP Option code point for OPTION_6RD (212) with a data length of 18 + N (OPTION_6RD with N/4 6rd BR addresses).

14. Acknowledgements

This RFC is based on Remi Despres' original idea described in [RFC5569] and the work done by Rani Assaf, Alexandre Cassen, and Maxime Bizon at Free Telecom. Brian Carpenter and Keith Moore documented 6to4, which all of this work is based upon. We thank Fred Templin for his review and contributions, and for sharing his experience with ISATAP. Review and encouragement have been provided by many others and in particular Chris Chase, Thomas Clausen, Wouter Cloetens, Wojciech Dec, Bruno Decraene, Remi Despres, Alain Durand, Washam Fan, Martin Gysi, David Harrington, Jerry Huang, Peter McCann, Alexey Melnikov, Dave Thaler, Eric Voit, and David Ward.

15. References

15.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.
- [RFC2491] Armitage, G., Schulter, P., Jork, M., and G. Harter, "IPv6 over Non-Broadcast Multiple Access (NBMA) networks", RFC 2491, January 1999.

- [RFC3056] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.
- [RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

15.2. Informative References

- [RFC2578] McCloghrie, K., Ed., Perkins, D., Ed., and J. Schoenwaelder, Ed., "Structure of Management Information Version 2 (SMIv2)", STD 58, RFC 2578, April 1999.
- [RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, October 2000.
- [RFC3068] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.
- [RFC3964] Savola, P. and C. Patel, "Security Considerations for 6to4", RFC 3964, December 2004.
- [RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, August 2006.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

- [RFC5569] Despres, R., "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)", RFC 5569, January 2010.
- [RoutingLoop] Nakibly and Arov, "Routing Loop Attacks using IPv6 Tunnels", August 2009, <http://www.usenix.org/event/woot09/tech/full_papers/nakibly.pdf>.
- [TR069] "TR-069, CPE WAN Management Protocol v1.1, Version: Issue 1 Amendment 2", December 2007.
- [V6OPS-LOOPS] Nakibly, G. and F. Templin, "Routing Loop Attack using IPv6 Automatic Tunnels: Problem Statement and Proposed Mitigations", Work in Progress, May 2010.

Authors' Addresses

Mark Townsley
Cisco
Paris,
France

EMail: mark@townsley.net

Ole Troan
Cisco
Bergen,
Norway

EMail: ot@cisco.com