Network Working Group
Request for Comments: 4192
Updates: 2072
Category: Informational

F. Baker
Cisco Systems
E. Lear
Cisco Systems GmbH
R. Droms
Cisco Systems

September 2005

Procedures for Renumbering an IPv6 Network without a Flag Day

Status of This Memo

Copyright Notice

Abstract

This document describes a procedure that can be used to renumber a network from one prefix to another. It uses IPv6's intrinsic ability to assign multiple addresses to a network interface to provide continuity of network service through a "make-before-break" transition, as well as addresses naming and configuration management issues. It also uses other IPv6 features to minimize the effort and time required to complete the transition from the old prefix to the new prefix.

Table of Contents

1. Introduction

The Prussian military theorist Carl von Clausewitz [Clausewitz] wrote, "Everything is very simple in war, but the simplest thing is difficult. These difficulties accumulate and produce a friction, which no man can imagine exactly who has not seen war.... So in war, through the influence of an 'infinity of petty circumstances' which cannot properly be described on paper, things disappoint us and we fall short of the mark". Operating a network is aptly compared to conducting a war. The difference is that the opponent has the futile expectation that homo ignoramus will behave intelligently.

A "flag day" is a procedure in which the network, or a part of it, is changed during a planned outage, or suddenly, causing an outage while the network recovers. Avoiding outages requires the network to be modified using what in mobility might be called a "make before break" procedure: the network is enabled to use a new prefix while the old one is still operational, operation is switched to that prefix, and then the old one is taken down.

This document addresses the key procedural issues in renumbering an IPv6 [RFC2460] network without a "flag day". The procedure is straightforward to describe, but operationally can be difficult to automate or execute due to issues of statically configured network state, which one might aptly describe as "an infinity of petty circumstances". As a result, in certain areas, this procedure is necessarily incomplete, as network environments vary widely and no one solution fits all. It points out a few of many areas where there are multiple approaches. This document updates [RFC2072]. This document also contains recommendations for application design and network management, which, if taken seriously, may avoid or minimize the impact of the issues.

1.1. Summary of the Renumbering Procedure

By "renumbering a network", we mean replacing the use of an existing (or "old") prefix throughout a network with a new prefix. Usually, both prefixes will be the same length. The procedures described in this document are, for the most part, equally applicable if the two prefixes are not the same length. During renumbering, sub-prefixes (or "link prefixes") from the old prefix, which have been assigned to links throughout the network, will be replaced by link prefixes from the new prefix. Interfaces on systems throughout the network will be configured with IPv6 addresses from the link prefixes of the new prefix, and any addresses from the old prefix in services like DNS [RFC1034][RFC1035] or configured into switches and routers and applications will be replaced by the appropriate addresses from the new prefix.

The renumbering procedure described in this document can be applied to part of a network as well as to an organization’s entire network. In the case of a large organization, it may be advantageous to treat the network as a collection of smaller networks. Renumbering each of the smaller networks separately will make the process more manageable. The process described in this document is generally applicable to any network, whether it is an entire organization network or part of a larger network.

1.2. Terminology

DDNS: Dynamic DNS [RFC2136][RFC3007] updates can be secured through the use of SIG(0) [RFC4033][RFC4034][RFC4035][RFC2931] and TSIG [RFC2845].

DHCP prefix delegation: An extension to DHCP [RFC3315] to automate the assignment of a prefix, for example, from an ISP to a customer [RFC3633].

flag day: A transition that involves a planned service outage.

ingress/egress filters: Filters applied to a router interface connected to an external organization, such as an ISP, to exclude traffic with inappropriate IPv6 addresses.

link prefix: A prefix, usually a /64 [RFC3177], assigned to a link.

SLAC: StateLess Address AutoConfiguration [RFC2462].

1.3. Summary of What Must Be Changed

Addresses from the old prefix that are affected by renumbering will appear in a wide variety of places in the components in the renumbered network. The following list gives some of the places that may include prefixes or addresses that are affected by renumbering, and gives some guidance about how the work required during renumbering might be minimized:

- Link prefixes assigned to links. Each link in the network must be assigned a link prefix from the new prefix.
- IPv6 addresses assigned to interfaces on switches and routers. These addresses are typically assigned manually, as part of configuring switches and routers.
- Routing information propagated by switches and routers.
- Link prefixes advertised by switches and routers [RFC2461].
- Ingress/egress filters.
- ACLs and other embedded addresses on switches and routers.
- IPv6 addresses assigned to interfaces on hosts. Use of StateLess Address Autoconfiguration (SLAC) [RFC2462] or DHCP [RFC3315] can mitigate the impact of renumbering the interfaces on hosts.

- DNS entries. New AAAA and PTR records are added and old ones removed in several phases to reflect the change of prefix. Caching times are adjusted accordingly during these phases.

- IPv6 addresses and other configuration information provided by DHCP.

- IPv6 addresses embedded in configuration files, applications, and elsewhere. Finding everything that must be updated and automating the process may require significant effort, which is discussed in more detail in Section 3. This process must be tailored to the needs of each network.

1.4. Multihoming Issues

In addition to the considerations presented, the operational matters of multihoming may need to be addressed. Networks are generally renumbered for one of three reasons: the network itself is changing its addressing policy and must renumber to implement the new policy (for example, a company has been acquired and is changing addresses to those used by its new owner), an upstream provider has changed its prefixes and its customers are forced to do so at the same time, or a company is changing providers and must perforce use addresses assigned by the new provider. The third case is common.

When a company changes providers, it is common to institute an overlap period, during which it is served by both providers. By definition, the company is multihomed during such a period. Although this document is not about multihoming per se, problems can arise as a result of ingress filtering policies applied by the upstream provider or one of its upstream providers, so the user of this document also needs to be cognizant of these issues. This is discussed in detail, and approaches to dealing with it are described, in [RFC2827] and [RFC3704].

2. Detailed Review of Procedure

During the renumbering process, the network transitions through eight states. In the initial state, the network uses just the prefix that is to be replaced during the renumbering process. At the end of the process, the old prefix has been entirely replaced by the new prefix, and the network is using just the new prefix. To avoid a flag day transition, the new prefix is deployed first and the network reaches an intermediate state in which either prefix can be used. In this state, individual hosts can make the transition to using the new prefix as appropriate to avoid disruption of applications. Once all

of the hosts have made the transition to the new prefix, the network is reconfigured so that the old prefix is no longer used in the network.

In this discussion, we assume that an entire prefix is being replaced with another entire prefix. It may be that only part of a prefix is being changed, or that more than one prefix is being changed to a single joined prefix. In such cases, the basic principles apply, but will need to be modified to address the exact situation. This procedure should be seen as a skeleton of a more detailed procedure that has been tailored to a specific environment. Put simply, season to taste.

2.1. Initial Condition: Stable Using the Old Prefix

Initially, the network is using an old prefix in routing, device interface addresses, filtering, firewalls, and other systems. This is a stable configuration.

2.2. Preparation for the Renumbering Process

The first step is to obtain the new prefix and new reverse zone from the delegating authority. These delegations are performed using established procedures, from either an internal or external delegating authority.

Before any devices are reconfigured as a result of the renumbering event, each link in the network must be assigned a sub-prefix from the new prefix. While this assigned link prefix does not explicitly appear in the configuration of any specific switch, router, or host, the network administrator performing the renumbering procedure must make these link prefix assignments prior to beginning the procedure to guide the configuration of switches and routers, assignment of addresses to interfaces, and modifications to network services such as DNS and DHCP.

Prior to renumbering, various processes will need to be reconfigured to confirm bindings between names and addresses more frequently. In normal operation, DNS name translations and DHCP bindings are often given relatively long lifetimes to limit server load. In order to reduce transition time from old to new prefix, it may be necessary to reduce the time to live (TTL) associated with DNS records and increase the frequency with which DHCP clients contact the DHCP server. At the same time, a procedure must be developed through which other configuration parameters will be updated during the transition period when both prefixes are available.

2.2.1. Domain Name Service

During the renumbering process, the DNS database must be updated to add information about addresses assigned to interfaces from the new prefix and to remove addresses assigned to interfaces from the old prefix. The changes to the DNS must be coordinated with the changes to the addresses assigned to interfaces.

Changes to the information in the DNS have to propagate from the server at which the change was made to the resolvers where the information is used. The speed of this propagation is controlled by the TTL for DNS records and the frequency of updates from primary to secondary servers.

The latency in propagating changes in the DNS can be managed through the TTL assigned to individual DNS records and through the timing of updates from primary to secondary servers. Appendix A gives an analysis of the factors controlling the propagation delays in the DNS.

The suggestions for reducing the delay in the transition to new IPv6 addresses applies when the DNS service can be given prior notice about a renumbering event. However, the DNS service for a host may be in a different administrative domain than the network to which the host is attached. For example, a device from organization A that roams to a network belonging to organization B, but the device's DNS A record is still managed by organization A, where the DNS service won't be given advance notice of a renumbering event in organization B.

One strategy for updating the DNS is to allow each system to manage its own DNS information through Dynamic DNS (DDNS) [RFC2136][RFC3007]. Authentication of these DDNS updates is strongly recommended and can be accomplished through TSIG and SIG(0). Both TSIG and SIG(0) require configuration and distribution of keys to hosts and name servers in advance of the renumbering event.

2.2.2. Mechanisms for Address Assignment to Interfaces

IPv6 addresses may be assigned through SLAC, DHCP, and manual processes. If DHCP is used for IPv6 address assignment, there may be some delay in the assignment of IPv6 addresses from the new prefix because hosts using DHCP only contact the server periodically to extend the lifetimes on assigned addresses. This delay can be reduced in two ways:

- Prior to the renumbering event, the T1 parameter (which controls the time at which a host using DHCP contacts the server) may be reduced.

- The DHCP Reconfigure message may also be sent from the server to the hosts to trigger the hosts to contact the server immediately.

2.3. Configuring Network Elements for the New Prefix

In this step, switches and routers and services are prepared for the new prefix but the new prefix is not used for any datagram forwarding. Throughout this step, the new prefix is added to the network infrastructure in parallel with (and without interfering with) the old prefix. For example, addresses assigned from the new prefix are configured in addition to any addresses from the old prefix assigned to interfaces on the switches and routers. Changes to the routing infrastructure for the new prefix are added in parallel with the old prefix so that forwarding for both prefixes operates in parallel. At the end of this step, the network is still running on the old prefix but is ready to begin using the new prefix.

The new prefix is added to the routing infrastructure, firewall filters, ingress/egress filters, and other forwarding and filtering functions. Routes for the new link prefixes may be injected by routing protocols into the routing subsystem, but the router advertisements should not cause hosts to perform SLAC on the new link prefixes; in particular the "autonomous address-configuration" flag [RFC2461] should not be set in the advertisements for the new link prefixes. The reason hosts should not be forming addresses at this point is that routing to the new addresses may not yet be stable.

The details of this step will depend on the specific architecture of the network being renumbered and the capabilities of the components that make up the network infrastructure. The effort required to complete this step may be mitigated by the use of DNS, DHCP prefix delegation [RFC3633], and other automated configuration tools.

While the new prefix is being added, it will of necessity not be working everywhere in the network, and unless properly protected by some means such as ingress and egress access lists, the network may be attacked through the new prefix in those places where it is operational.

Once the new prefix has been added to the network infrastructure, access-lists, route-maps, and other network configuration options that use IP addresses should be checked to ensure that hosts and services that use the new prefix will behave as they did with the old one. Name services other than DNS and other services that provide

information that will be affected by renumbering must be updated in such a way as to avoid responding with stale information. There are several useful approaches to identify and augment configurations:

- Develop a mapping from each network and address derived from the old prefix to each network and address derived from the new prefix. Tools such as the UNIX "sed" or "perl" utilities are useful to then find and augment access-lists, route-maps, and the like.

- A similar approach involves the use of such mechanisms as DHCP prefix delegation to abstract networks and addresses.

Switches and routers or manually configured hosts that have IPv6 addresses assigned from the new prefix may be used at this point to test the network infrastructure.

Advertisement of the prefix outside its network is the last thing to be configured during this phase. One wants to have all of one's defenses in place before advertising the prefix, if only because the prefix may come under immediate attack.

At the end of this phase, routing, access control, and other network services should work interchangeably for both old and new prefixes.

2.4. Adding New Host Addresses

Once the network infrastructure for the new prefix is in place and tested, IPv6 addresses from the new prefix may be assigned to host interfaces while the addresses from the old prefix are retained on those interfaces. The new IPv6 addresses may be assigned through SLAC, DHCP, and manual processes. If SLAC is used in the network, the switches and routers are configured to indicate that hosts should use SLAC to assign IPv6 addresses from the new prefix. If DHCP is used for IPv6 address assignment, the DHCP service is configured to assign addresses from both prefixes to hosts. The addresses from the new prefixes will not be used until they are inserted into the DNS.

Once the new IPv6 addresses have been assigned to the host interfaces, both the forward and reverse maps within DNS should be updated for the new addresses, either through automated or manual means. In particular, some clients may be able to update their forward maps through DDNS, but automating the update of the reverse zone may be more difficult as discussed in Section 4.2.

2.5. Stable Use of Either Prefix

Once the network has been configured with the new prefix and has had sufficient time to stabilize, it becomes a stable platform with two addresses configured on each and every infrastructure component interface (apart from interfaces that use only the link-local address), and two non-link-local addresses are available for the use of any host, one in the old prefix and one in the new. This is a stable configuration.

2.6. Transition from Use of the Old Prefix to the New Prefix

When the new prefix has been fully integrated into the network infrastructure and has been tested for stable operation, hosts, switches, and routers can begin using the new prefix. Once the transition has completed, the old prefix will not be in use in the network.

2.6.1. Transition of DNS Service to the New Prefix

The DNS service is configured to use the new prefix by removing any IPv6 addresses from the old prefix from the DNS server configuration. External references to the DNS servers, such as in the DNS service from which this DNS domain was delegated, are updated to use the IPv6 addresses from the new prefix.

2.6.2. Transition to Use of New Addresses

When both prefixes are usable in the network, each host can make the transition from using the old prefix to the new prefix at a time that is appropriate for the applications on the host. If the host transitions are randomized, DNS dynamic update mechanisms can better scale to accommodate the changes to the DNS.

As services become available through addresses from the new prefix, references to the hosts providing those services are updated to use the new prefix. Addresses obtained through DNS will be automatically updated when the DNS names are resolved. Addresses may also be obtained through DHCP and will be updated as hosts contact DHCP servers. Addresses that are otherwise configured must be updated appropriately.

It may be necessary to provide users with tools or other explicit procedures to complete the transition from the use of the old prefix to the new prefix, because some applications and operating system functions may be configured in ways that do not use DNS at all or will not use DNS to resolve a domain name to a new address once the new prefix is available. For example, a device that only uses DNS to

resolve the name of an NTP server when the device is initialized will not obtain the address from the new prefix for that server at this point in the renumbering process.

This last point warrants repeating (in a slightly different form). Applications may cache addressing information in different ways, for varying lengths of time. They may cache this information in memory, on a file system, or in a database. Only after careful observation and consideration of one's environment should one conclude that a prefix is no longer in use. For more information on this issue, see [DNSOP].

The transition of critical services such as DNS, DHCP, NTP [RFC1305], and important business services should be managed and tested carefully to avoid service outages. Each host should take reasonable precautions prior to changing to the use of the new prefix to minimize the chance of broken connections. For example, utilities such as netstat and network analyzers can be used to determine if any existing connections to the host are still using the address from the old prefix for that host.

Link prefixes from the old prefix in router advertisements and addresses from the old prefix provided through DHCP should have their preferred lifetimes set to zero at this point, so that hosts will not use the old prefixes for new communications.

2.7. Removing the Old Prefix

Once all sessions are deemed to have completed, there will be no dependence on the old prefix. It may be removed from the configuration of the routing system and from any static configurations that depend on it. If any configuration has been created based on DNS information, the configuration should be refreshed after the old prefixes have been removed from the DNS.

During this phase, the old prefix may be reclaimed by the provider or Regional Internet Registry that granted it, and addresses within that prefix are removed from the DNS.

In addition, DNS reverse maps for the old prefix may be removed from the primary name server and the zone delegation may be removed from the parent zone. Any DNS, DHCP, or SLAC timers that were changed should be reset to their original values (most notably the DNS forward map TTL).

2.8. Final Condition: Stable Using the New Prefix

This is equivalent to the first state, but using the new prefix.

3. How to Avoid Shooting Yourself in the Foot

The difficult operational issues in Section 2.3, Section 2.6, and Section 2.7 are in dealing with the configurations of routers and hosts that are not under the control of the network administrator or are manually configured. Examples of such devices include Voice over IP (VoIP) telephones with static configuration of boot or name servers, dedicated devices used in manufacturing that are configured with the IP addresses for specific services, the boot servers of routers and switches, etc.

3.1. Applications Affected by Renumbering

Applications may inadvertently ignore DNS caching semantics associated with IP addresses obtained through DNS resolution. The result is that a long-lived application may continue to use a stale IP address beyond the time at which the TTL for that address has expired, even if the DNS is updated with new addresses during a renumbering event.

For example, many existing applications make use of standard POSIX functions such as getaddrinfo(), which do not preserve DNS caching semantics. If the application caches the response or for whatever reason actually records the response on disk, the application will have no way to know when the TTL for the response has expired. Any application that requires repeated use of an IP address should either not cache the result or make use of an appropriate function that also conveys the TTL of the record (e.g., getrrsetbyname()).

Application designers, equipment vendors, and the Open Source community should take note. There is an opportunity to serve their customers well in this area, and network operators should either develop or purchase appropriate tools.

3.2. Renumbering Switch and Router Interfaces

The configuration and operation of switches and routers are often designed to use static configuration with IP addresses or to resolve domain names only once and use the resulting IP addresses until the element is restarted. These static configurations complicate the process of renumbering, requiring administration of all of the static information and manual configuration during a renumbering event.

Because switches and routers are usually single-purpose devices, the user interface and operating functions (software and hardware) are often better integrated than independent services running on a server platform. Thus, it is likely that switch vendors and router vendors

can design and implement consistent support for renumbering across all of the functions of switches and routers.

To better support renumbering, switches and routers should use domain names for configuration wherever appropriate, and they should resolve those names using the DNS when the lifetime on the name expires.

3.3. Ingress Filtering

An important consideration in Section 2.3, in the case where the network being renumbered is connected to an external provider, is the network's ingress filtering policy and its provider's ingress filtering policy. Both the network firewall's ingress filter and the provider's ingress filter on the access link to the network should be configured to prevent attacks that use source address spoofing. Ingress filtering is considered in detail in "Ingress Filtering for Multihomed Networks" [RFC3704].

3.4. Link Flaps in BGP Routing

A subtle case arises during step 2 in BGP routing when renumbering the address(es) used to name the BGP routers. Two practices are common: one is to identify a BGP router by a stable address such as a loopback address; another is to use the interface address facing the BGP peer. In each case, when adding a new prefix, a certain ambiguity is added: the systems must choose between the addresses, and depending on how they choose, different events can happen.

- If the existing address remains in use until removed, then this is minimized to a routing flap on that event.

- If both systems decide to use the address in the new prefix simultaneously, the link flap may occur earlier in the process, and if this is being done automatically (such as via the router renumbering protocol), it may result in route flaps throughout the network.

- If the two systems choose differently (one uses the old address and one uses the new address), a stable routing outage occurs.

This is not addressed by proposals such as [IDR-RESTART], as it changes the "name" of the system, making the matter not one of a flap in an existing relationship but (from BGP's perspective) the replacement of one routing neighbor with another. Ideally, one should bring up the new BGP connection for the new address while the old remains stable and in use, and only then take down the old. In this manner, while there is a TCP connection flap, routing remains stable.

4. Call to Action for the IETF

The more automated one can make the renumbering process, the better for everyone. Sadly, there are several mechanisms that either have not been automated or have not been automated consistently across platforms.

4.1. Dynamic Updates to DNS Across Administrative Domains

The configuration files for a DNS server (such as named.conf) will contain addresses that must be reconfigured manually during a renumbering event. There is currently no easy way to automate the update of these addresses, as the updates require both complex trust relationships and automation to verify them. For instance, a reverse zone is delegated by an upstream ISP, but there is currently no mechanism to note additional delegations.

4.2. Management of the Reverse Zone

In networks where hosts obtain IPv6 addresses through SLAC, updates of reverse zone are problematic because of lack of trust relationship between administrative domain owning the prefix and the host assigning the low 64 bits using SLAC. For example, suppose a host, H, from organization A is connected to a network owned by organization B. When H obtains a new address during a renumbering event through SLAC, H will need to update its reverse entry in the DNS through a DNS server from B that owns the reverse zone for the new address. For H to update its reverse entry, the DNS server from B must accept a DDNS request from H, requiring that an inter-administrative domain trust relationship exist between H and B. The IETF should develop a BCP recommendation for addressing this problem.

5. Security Considerations

The process of renumbering is straightforward in theory but can be difficult and dangerous in practice. The threats fall into two broad categories: those arising from misconfiguration and those that are actual attacks.

Misconfigurations can easily arise if any system in the network "knows" the old prefix, or an address in it, a priori and is not configured with the new prefix, or if the new prefix is configured in a manner that replaces the old instead of being co-equal to it for a period of time. Simplistic examples include the following:

Neglecting to reconfigure a system that is using the old prefix in some static configuration: in this case, when the old prefix is removed from the network, whatever feature was so configured becomes inoperative - it is not configured for the new prefix, and the old prefix is irrelevant.

Configuring a system via an IPv6 address, and replacing that old address with a new address: because the TCP connection is using the old and now invalid IPv6 address, the SSH session will be terminated and you will have to use SSH through the new address for additional configuration changes.

Removing the old configuration before supplying the new: in this case, it may be necessary to obtain on-site support or travel to the system and access it via its console.

Clearly, taking the extra time to add the new prefix to the configuration, allowing the network to settle, and then removing the old obviates this class of issue. A special consideration applies when some devices are only occasionally used; the administration must allow a sufficient length of time in Section 2.6 or apply other verification procedures to ensure that their likelihood of detection is sufficiently high.

A subtle case of this type can result when the DNS is used to populate access control lists and similar security or QoS configurations. DNS names used to translate between system or service names and corresponding addresses are treated in this procedure as providing the address in the preferred prefix, which is either the old or new prefix but not both. Such DNS names provide a means, as described in Section 2.6, to cause systems in the network to stop using the old prefix to access servers or peers and cause them to start using the new prefix. DNS names used for access control lists, however, need to go through the same three-step procedure used for other access control lists, having the new prefix added to them as discussed in Section 2.3 and the old prefix removed as discussed in Section 2.7.

It should be noted that the use of DNS names in this way is not universally accepted as a solution to this problem; [RFC3871] especially notes cases where static IP addresses are preferred over DNS names, in order to avoid a name lookup when the naming system is inaccessible or when the result of the lookup may be one of several interfaces or systems. In such cases, extra care must be taken to manage renumbering properly.

Attacks are also possible. Suppose, for example, that the new prefix has been presented by a service provider, and the service provider

starts advertising the prefix before the customer network is ready. The new prefix might be targeted in a distributed denial of service attack, or a system might be broken into using an application that would not cross the firewall using the old prefix, before the network's defenses have been configured. Clearly, one wants to configure the defenses first and only then accessibility and routing, as described in Section 2.3 and Section 3.3.

The SLAC procedure described in [RFC2462] renumbers hosts. Dynamic DNS provides a capability for updating DNS accordingly. Managing configuration items apart from those procedures is most obviously straightforward if all such configurations are generated from a central configuration repository or database, or if they can all be read into a temporary database, changed using appropriate scripts, and applied to the appropriate systems. Any place where scripted configuration management is not possible or is not used must be tracked and managed manually. Here, there be dragons.

In ingress filtering of a multihomed network, an easy solution to the issues raised in Section 3.3 might recommend that ingress filtering should not be done for multihomed customers or that ingress filtering should be special-cased. However, this has an impact on Internet security. A sufficient level of ingress filtering is needed to prevent attacks using spoofed source addresses. Another problem comes from the fact that if ingress filtering is made too difficult (e.g., by requiring special-casing in every ISP doing it), it might not be done at an ISP at all. Therefore, any mechanism depending on relaxing ingress filtering checks should be dealt with with extreme care.

6. Acknowledgements

This document grew out of a discussion on the IETF list. Commentary on the document came from Bill Fenner, Christian Huitema, Craig Huegen, Dan Wing, Fred Templin, Hans Kruse, Harald Tveit Alvestrand, Iljitsch van Beijnum, Jeff Wells, John Schnizlein, Laurent Nicolas, Michael Thomas, Michel Py, Ole Troan, Pekka Savola, Peter Elford, Roland Dobbins, Scott Bradner, Sean Convery, and Tony Hain.

Some took it on themselves to convince the authors that the concept of network renumbering as a normal or frequent procedure is daft. Their comments, if they result in improved address management practices in networks, may be the best contribution this note has to offer.

Christian Huitema, Pekka Savola, and Iljitsch van Beijnum described the ingress filtering issues. These made their way separately into [RFC3704], which should be read and understood by anyone who will

temporarily or permanently create a multihomed network by renumbering from one provider to another.

In addition, the 6NET consortium, notably Alan Ford, Bernard Tuy, Christian Schild, Graham Holmes, Gunter Van de Velde, Mark Thompson, Nick Lamb, Stig Venaas, Tim Chown, and Tina Strauf, took it upon themselves to test the procedure. Some outcomes of that testing have been documented here, as they seemed of immediate significance to the procedure; 6NET will also be documenting its own "lessons learned".

7. References

7.1. Normative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC2072] Berkowitz, H., "Router Renumbering Guide", RFC 2072, January 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

[RFC2462] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

7.2. Informative References

[Clausewitz] von Clausewitz, C., Howard, M., Paret, P. and D. Brodie, "On War, Chapter VII, 'Friction in War'", June 1989.

[DNSOP] Durand, A., Ihren, J. and P. Savola, "Operational Considerations and Issues with IPv6 DNS", Work in Progress, October 2004.

[IDR-RESTART] Sangli, S., Rekhter, Y., Fernando, R., Scudder, J. and E. Chen, "Graceful Restart Mechanism for BGP", Work in Progress, June 2004.

[RFC1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation and Analysis", RFC 1305, March 1992.

[RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, August 1996.

[RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", RFC 1996, August 1996.

[RFC2136] Vixie, P., Thomson, S., Rekhter, Y., and J. Bound, "Dynamic Updates in the Domain Name System (DNS UPDATE)", RFC 2136, April 1997.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

[RFC2845] Vixie, P., Gudmundsson, O., Eastlake 3rd, D., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", RFC 2845, May 2000.

[RFC2931] Eastlake 3rd, D., "DNS Request and Transaction Signatures (SIG(0)s)", RFC 2931, September 2000.

[RFC3007] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.

[RFC3177] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.

[RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

[RFC3871] Jones, G., "Operational Security Requirements for Large Internet Service Provider (ISP) IP Network Infrastructure", RFC 3871, September 2004.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

Appendix A. Managing Latency in the DNS

The procedure in this section can be used to determine and manage the latency in updates to information a DNS resource record (RR).

There are several kinds of possible delays that are ignored in these calculations:

- the time it takes for the administrators to make the changes;
- the time it may take to wait for the DNS update, if the secondaries are only updated at regular intervals, and not immediately; and
- the time the updating to all the secondaries takes.

Assume the use of NOTIFY [RFC1996] and IXFR [RFC1995] to transfer updated information from the primary DNS server to any secondary servers; this is a very quick update process, and the actual time to update of information is not considered significant.

There is a target time, TC, at which we want to change the contents of a DNS RR. The RR is currently configured with TTL == TTLOLD. Any cached references to the RR will expire no more than TTLOLD in the future.

At time TC - (TTLOLD + TTLNEW), the RR in the primary is configured with TTLNEW (TTLNEW < TTLOLD). The update process is initiated to push the RR to the secondaries. After the update, responses to queries for the RR are returned with TTLNEW. There are still some cached references with TTLOLD.

At time TC - TTLNEW, the RR in the primary is configured with the new address. The update process is initiated to push the RR to the secondaries. After the update, responses to queries for the RR return the new address. All the cached references have TTLNEW. Between this time and TC, responses to queries for the RR may be returned with either the old address or the new address. This ambiguity is acceptable, assuming the host is configured to respond to both addresses.

At time TC, all the cached references with the old address have expired, and all subsequent queries will return the new address. After TC (corresponding to the final state described in Section 2.8), the TTL on the RR can be set to the initial value TTLOLD.

The network administrator can choose TTLOLD and TTLNEW to meet local requirements.

As a concrete example, consider a case where TTLOLD is a week (168 hours) and TTLNEW is an hour. The preparation for the change of addresses begins 169 hours before the address change. After 168 hours have passed and only one hour is left, the TTLNEW has propagated everywhere, and one can change the address record(s). These are propagated within the hour, after which one can restore TTL value to a larger value. This approach minimizes time where it is uncertain what kind of (address) information is returned from the DNS.

Authors' Addresses

Fred Baker
Cisco Systems
1121 Via Del Rey
Santa Barbara, CA 93117
US

Phone: 408-526-4257
Fax: 413-473-2403
EMail: fred@cisco.com

Eliot Lear
Cisco Systems GmbH
Glatt-com 2nd Floor
CH-8301 Glattzentrum
Switzerland

Phone: +41 1 878 9200
EMail: lear@cisco.com

Ralph Droms
Cisco Systems
200 Beaver Brook Road
Boxborough, MA 01719
US

Phone: +1 978 936-1674
EMail: rdroms@cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.