Deploying IPv6: Problems and Solutions

13 November 2011

NTT Information Sharing Platform Laboratories

Notes

- Articles based on the chapter 5 of this documents appeared in '*Nikkei Network*' issued by Nikkei Business Publications Inc. from April to June 2011 (in Japanese).
- The issue list in chapter 6 of this document is based on the document issued by IPv6 fix working group in IPv6 promotion council Japan (in Japanese).
  - http://www.v6pc.jp/jp/upload/pdf/2011093001_v6fix.pdf
- Network Security Project in NTT Information Sharing Platform Laboratories publishes this document. If there are any comments, questions and suggestions, please contact following address.
  - *E-mail: ipv6-deployment-issues@nttv6.com*
- We, NTT Information sharing platform laboratories carefully review information provided on this document for its accuracy, value, completeness, violation of right, and so on, however, we cannot guarantee the validity of the information. We do not take responsibility for any damage caused by utilization or by unavailability of this document. Contents of this document may be revised or deleted without notice.

1. Introduction

The central pool of IPv4 addresses was exhausted on February 3, 2011, and the address pool for the Asia Pacific Region was depleted in April 2011. We are thus faced with having to deal with the depletion of IPv4 addresses while at the same time aggressively promoting the rollout of IPv6 addresses. This document will catalog the issues and problems raised by the deployment of IPv6 networks, will outline the causes of these problems, and will offer solutions to overcome these problems. The *fallback problem, the rogue router advertisement (RA) problem and the path MTU black hole problem* are especially pervasive, so we will discuss these issues in some detail.

2. Status of Internet Resources and IPv4 Address Depletion

All the number resources used by the Internet—IP addresses, AS numbers, protocol numbers, and so on—must be uniformly managed throughout the world, and this job of centralized management is the responsibility of the Internet Assigned Numbers Authority (IANA), a division of the international not-for-profit organization Internet Corporation for Assigned Names and Numbers (ICANN). Figure 1 shows a schematic overview of Internet resource management.

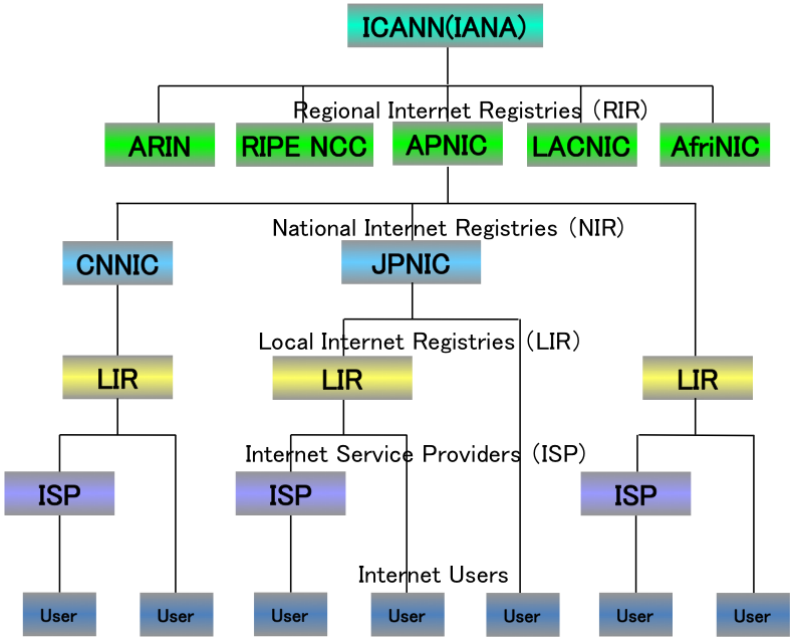

Figure 1 : Internet resource management structure

In order to effectively coordinate global Internet Protocol addressing systems, the IANA delegates responsibility to five subordinate Regional Internet Registries (RIRs), including the Asia Pacific Network Information Centre (APNIC), the RIR in charge of the Asia-Pacific region. The RIRs allocate IP addresses to National Internet Registries (NIRs) of countries in their respective regions, which are then allocated to ISPs requesting IP resources. The IANA and RIRs thus hold IP addresses and AS numbers as pools of Internet resources.

As the Internet has expanded to encompass the world, the volume of IP addresses and other Internet resources needed to connect to the Internet has also grown exponentially. Now there is great concern over the rapid depletion of IPv4 addresses, and indeed the IANA's pool of unallocated addresses was exhausted on February 3, 2011. Stocks of unallocated IPv4 address blocks held by the RIRs are rapidly being depleted, most notably addresses in the APNIC region are quickly being consumed and new IPv4 addresses are expected to be exhausted by April or May of 2011 (IPv4 Address Report).

2.1. Coping with IPv4 Address Pool Depletion

As unallocated IPv4 address space has continued to shrink, concerned scientists around the globe have been looking for ways to deal with the exhaustion of IPv4 addresses. Japan's address registry, the Japan Network Information Center (JPNIC), overseen by the Ministry of Internal Affairs and Communications (MIC), the Task Force on IPv4 Address Exhaustion, and other research groups have focused on this issue and come up with three strategies:

1. Make the most effective use of finite address resources using Network Address Translation (NAT) and other IPv4 address sharing schemes,
2. Recover and reuse IPv4 addresses that have never been used or needed, and
3. Deploy IPv6.

The first strategy, address sharing through NAT, can only be implemented on the user side and does not address the rapid proliferation of network servers. Regarding the second strategy, it is not at all clear to what extent IPv4 addresses can be recovered and circulated or whether IPv4 addresses could be provided in sufficient numbers when required. Moreover, recovery of unused addresses would involve fragmentation of IPv4 addresses, resulting in problems associating with increasing the IPv4 routing table. By process of elimination, we come to the third strategy—deployment of IPv6—as the only long-term solution to the problem of exhaustion of IPv4 addresses.

3. Path to IPv6 Deployment

The Internet Engineering Task Force (IETF), the premier Internet standardization body, has always been centrally involved not only in IPv6 standardization, but in developing scenarios and studying the migration from IPv4 to IPv6 and the deployment of IPv6. However, the deployment of IPv6 is not going as planned, so there has been an ongoing effort to reassess the rollout scenarios that were painstakingly developed. Figure 2 shows a schematic of the IETF working groups that have been involved in the IPv6 deployment.

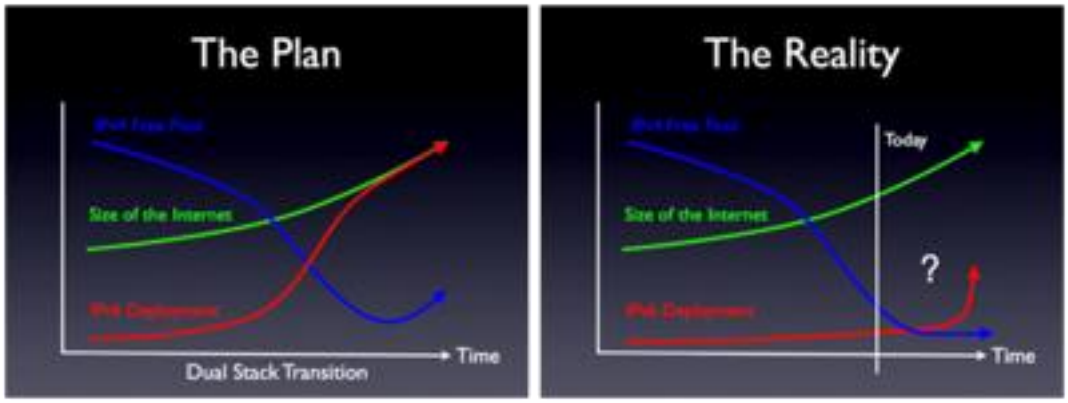

Extract from http://www.cs.jhu.edu/~bkhabs/v4v6/townsley-64-coexist-00.pdf

Figure 2 : IPv6 deployment: vision and reality

When it first began working on IPv6 standardization issues, the IETF operated on the assumption that IPv6 would be gradually rolled out as the pool of IPv4 addresses diminished, so by the time IPv4 addresses were exhausted, IPv6 would be up and running. Unfortunately the deployment of IPv6 has been held up, so we face the daunting prospect of dealing with the exhaustion of IPv4 addresses and the deployment of IPv6 addresses at the same time.

As circumstances have changed, standardization work is focusing on both support for IPv6 and dealing with the exhaustion of the pool of IPv4 addresses: Standardization is proceeding on CGN and DS-Lite (Durand) IPv4 address sharing schemes (permitting multiple parties to share use of a single IPv4 address) (Miyagawa, 2010) while at the same time addressing an IPv4/IPv6 interconversion scheme necessary for early deployment of IPv6 (Fred Baker). Schemes for rolling out IPv6 on already extensively deployed IPv4 networks have also been

standardized, including the 6to4 approach [RFC 3056] (Carpenter, 2001) that enables IPv6 Internet connectivity over IPv4, and the 6rd scheme [RFC 5969] (Mark Townsley, 2010) giving ISPs access to IPv4 private address space. Japan and other countries are already providing IPv6 Internet connectivity over 6rd links. Many technologies are moving through standardization that opens the way to IPv6, and the barriers to ISP adoption of IPv6 are clearly coming down. Even Application Service Providers, usually the last to accept changes, are gradually moving to embrace IPv6. Google has stepped up its efforts to make their services IPv6 compliant, and most Google services can now be accessed over IPv6. Facebook and Yahoo have also followed suit. Services like Akamai that exploit many other services has announced a roadmap to IPv6-compliance.

In addition, the ICANN, various RIRs, and other organizations charged with managing Internet resources are actively engaged in educational initiatives and promoting IPv6 technology to the ISPs and other interested parties in their respective areas. The UN, government agencies, and many other organizations around the world are also very actively involved in promoting the deployment and spread of IPv6. The U.S. has published *Technical Infrastructure of USGv6 Adoption* that outlines the requirements and the conditions to support IPv6 in order to provide many government agencies with supplies and equipment. The OECD has also kept abreast of the current state of IPv4 address depletion and IPv6 deployment, as evidenced by a recent study *OECD Resources on Internet Addressing: IPv4 and IPv6*.

It is clear that IPv6 promotion activities are making headway across different sectors around the world, and it is only a matter of time before IPv6 becomes the reigning Internet protocol.

4. Issues Arising from Deployment of IPv6

Even as the rollout of IPv6 continues, we can anticipate unforeseen problems will occur along the way. To focus attention on these problems, the IPv4/IPv6 Coexistence Working Group of the IPv6 Promotion Council set up the IPv6 Deployment Issues Sub Working Group (SWG) to compile a list of potential problems associated with the deployment of IPv6, and to offer recommended solutions to these problems. Principle problems that have been identified so far include the following.

1. IPv4/IPv6 fallback
2. Rogue RA
3. PMTUD black holes
4. Problems relating to the Domain Name System (DNS) when IPv6 is deployed
5. Captive portal and DSN problems (IPv6 uninstall at hotels).
6. Poor quality tunnels, transition technology related issues
7. Different QoS protocols at dual-stack sites, different QoS of IPv4 and IPv6
8. Address selection related problems (multi-prefix problems)
9. Problems with false recognition and IPv6-ready routers that only support IPv6 bridge functions (IPv6 pass-through functions)
10. Problems with bridge filters in IPv6-ready routers
11. DNS registration issues ("DNS registration, reverse lookup, forward lookup, DDNS")
12. Security and filtering issues (ICMP filtering problems, etc.)
13. IPv6-ready mail system issues (sending and receiving mail)
14. IPv6-ready mail system issues (anti-spam techniques)
15. Blacklist Database Service (DNSBL) Issues
16. Presence of unsupported L2 multicast equipment
17. Adverse effects of IPv6 multicast on home communications
18. IPv6 address notation
19. Implementations that do not meet minimum specifications
20. IPv6 privacy address (RFC 4941) issues
21. IPv6 address traceability (privacy) issues

22. CGN, translation issues
23. Expressions subject to misunderstanding, problems from sharing obsolete information
24. IPv6 impact on multiple IPv4 subnets
25. IPv6 impact on large-scale L2 networks
26. Problems that cannot be resolved within CPEs own domain
27. IRR registration issues
28. Number of DNS records and OS operation
29. Problems regarding how sites are viewed

In the next section, we will take a detailed look at the first three of these issues which have proven especially problematic. We will then provide a more cursory treatment of the remaining 26 issues.

5. IPv6 Deployment Issues: Fallback, Rouge RA, and Path MTU

5.1. IPv6/IPv4 Fallback

5.1.1. What is IPv6/IPv4 Fallback?

5.1.1.1. Network Topology after IPv6 Installed

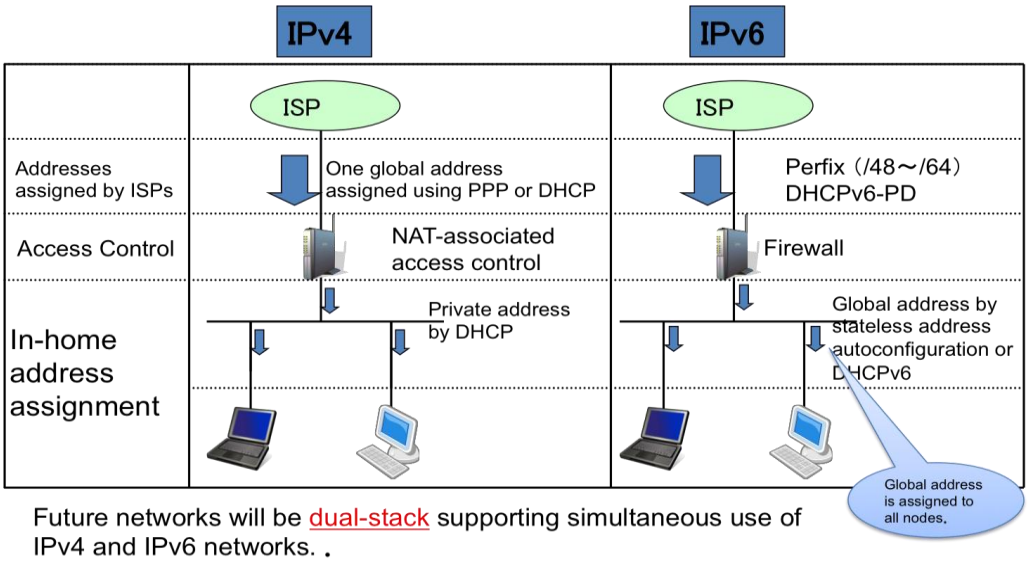

Figure 3 : Difference between IPv4 and IPv6

Internet Service Providers deliver IPv6 services to users, but the Internet is little changed from the IPv4 network. IPv6 connectivity to users' homes (the last mile) is also virtually the same as IPv4, so aside from 6rd and other IPv4 tunnel connectivity schemes that ISPs are already offering, IPv6 will be provided by PPPoE [RFC 2516] (L. Mamakos, 1999), network connectivity by Ethernet (non-tunnel interface), and the like. Essentially, the transition from IPv4 to IPv6 involves the protocol used to allocate addresses (IPv4 uses PPP/DHCP, whereas IPv6 uses DHCPv6-PD/stateless address auto configuration, etc.), the type of address allocated to homes (IPv4 allocates private addresses, IPv6 allocates global addresses), and whether or not Network Address Translation (NAT) is present. Figure 3 shows a summary overview of the main differences between IPv4 and IPv6. There are differences in detail, but the big environmental change is the transitions to a dual-stack environment that supports provisioning of IPv4 and IPv6 at the same time.

5.1.1.2. Dual-Stack Host Operation and Fallback

A great deal of home electronic equipment already supports IPv6 including PCs with Windows, Mac, and Linux OSes, smart phones, and some Internet-ready TVs. In the dual-stack environment, PCs already have both IPv4 and IPv6 addresses for accessing these IPv6-enabled nodes. Figure 4 illustrates how communications work in the dual-stack node.

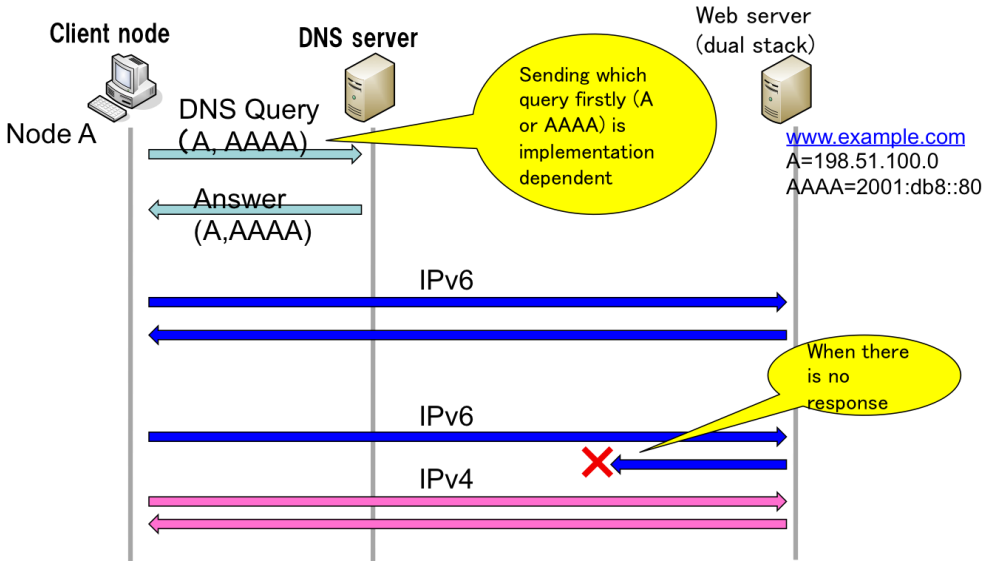

Figure 4 : Dual-stack node communications operations

Here we assume that Node A communicates with www.example.com on a server that supports both IPv4 and IPv6 addresses. Node A first queries the DNS server for www.example.com's IPv4 address (an A resource record) and IPv6 address (an AAAA resource record [RFC 3596] (S. Thomson, 2003)). As to which resource record will be queried first—the A record or the AAAA record—and the timing of the query this will depend on the implementation. If both addresses are obtained from the DNS server, many implementations are set up to prefer IPv6 communication. In this example, Node A tries IPv6 communication using the IPv6 address 2001:db8::80 that was supplied by the server. If the initial communication succeeds then information is exchanged over IPv6; if a problem is encountered that causes the IPv6 communication to fail, the link simply gives up on IPv6 and tries to communicate over IPv4. This switching from IPv6 to IPv4 communication is called *fallback*.

5.1.1.3. Problems related Fallback

There is no problem associated with falling back from IPv6 to IPv4 if the transition occurs quickly. In communication between dual-stack nodes, the communication protocol is redundant and there is a greater chance of network problems occurring on the communications route. But if fallback takes a long time, user convenience would be markedly degraded by deploying IPv6. In browsing the web, for example, users would have to wait much longer for pages to be displayed. For fallback to work smoothly, nodes must be capable of detecting the network state. As one can see in Figure 5, both IPv4 and IPv6 include schemes for sending error messages back to the node originating a communication in the event a problem occurs on the network.

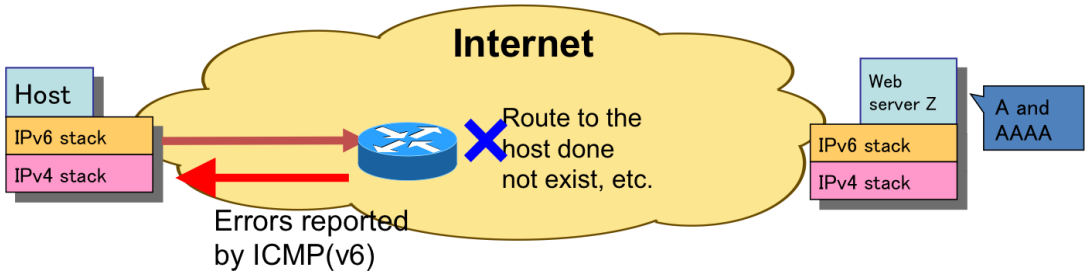

Figure 5 : Error notification from the network

Say a problem occurs on a communication route so there is no path to the destination, then an intermediate router along the path sends an ICMP message (Internet Control Message Protocol; called ICMPv6 in IPv6 [RFC 4443] (A. Conta, 2006)) back to the source address with notification that the communication has failed. Upon receiving this message, the source node then has the option of falling back to IPv4 (the specific action taken in response to an ICMP message is different depending on the type of ICMP which varies for different upper-level protocols). A communication failure is one thing, but a problem caused by equipment failure that is hard to detect or perhaps the ICMP message doesn't reach the source node at all due to a packet filter problem. In these latter cases, the source node remains unaware that a network problem has occurred. When TCP or some other protocol guaranteeing reliability is used, communication is suspended by a time out, and this provides time for the fallback transition to be made. Figure 6 illustrates how fallback occurs when accessing two IPv6 addresses on a web server. Many servers

today register multiple addresses for the same domain name to provide redundancy and load balancing, so fallback may occur a number of times corresponding to the number of IP addresses that are registered.

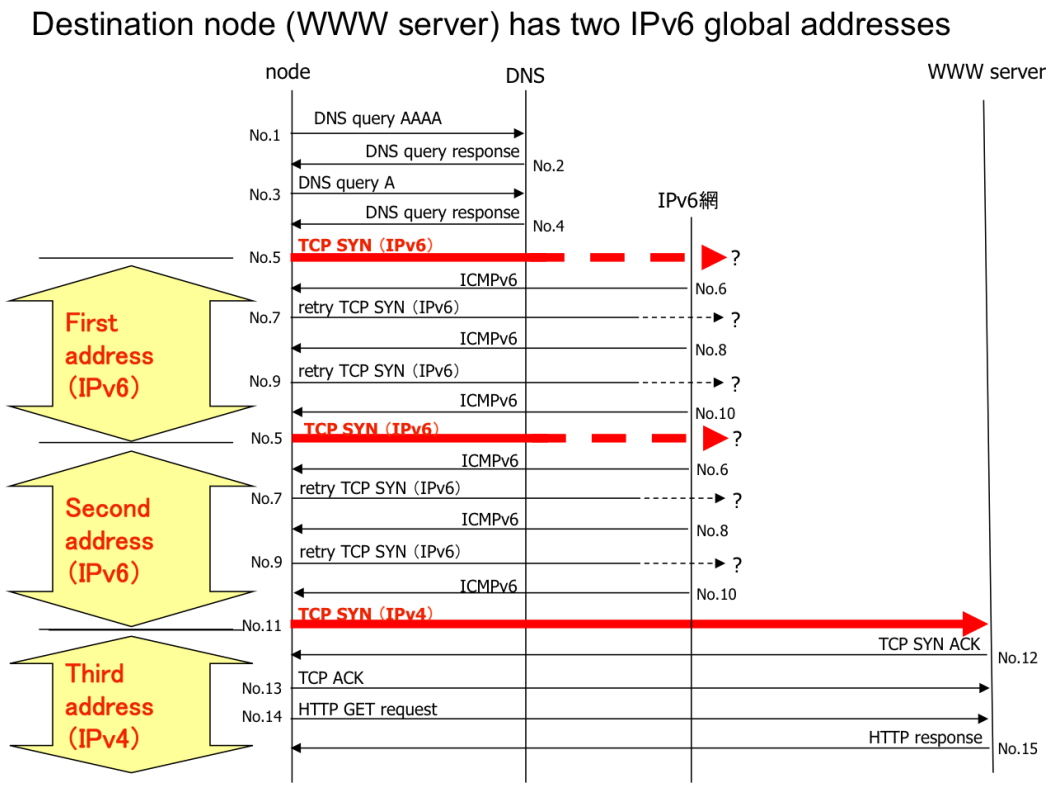

Figure 6 : Examples of fallback

Note that messages 6, 8, and 10 in Figure 6 are error messages sent from the network side, but if there is no notification that an error has occurred, this section goes into time out. According to the provisions of the TCP protocol, the network tries to reconnect up to three times for a single address if a session cannot be established (this procedure too is prescribed based on type of error). One might assume the procedure is implemented by the protocol, for the actual equipment may not operate in this way. Table 1 shows measurement results for fallback times from IPv6 to IPv4 for a number of different OSes in a actual forced fallback environments. One can see from the table that the time to respond to ICMPv6 error notifications and the time required for fallback varies considerably among the OSes. Depending on the OS, some do not perform fallback even though they respond to the ICMPv6 message, others suspend communication until a TCP timeout, while for others

OS	Browser	No response	Type=1 (Destination Unreachable)							TCP reset
			Code=0	Code=1	Code=2	Code=3	Code=4	Code=5	Code=6	
Windows Vista Home Premium	IE	21.00	21.01	21.02	20.99	21.02	21.00	21.00	21.01	1.00
	FireFox	21.00	21.01	21.00	20.99	21.01	21.00	20.99	21.00	1.00
	Google Chrome	21.00	21.00	21.00	21.00	20.99	20.99	21.00	21.00	1.00
Windows 7 Ultimate	IE	21.01	21.01	21.01	21.01	21.00	21.00	21.01	21.01	1.01
	FireFox	21.01	20.99	21.00	21.00	21.00	21.00	21.00	21.00	1.01
	Google Chrome	21.00	21.01	21.01	21.01	21.01	21.00	21.00	21.00	1.01
MAC OS X Ver 10.6.4	Safari	74.70	3.95	3.95	3.93	3.94	3.97	74.99	74.71	0.01
	FireFox	74.75	3.94	3.95	No Fallback	No fallback	3.95	74.78	74.75	0.01
	Google Chrome	74.76	3.95	3.94	No Fallback	No fallback	3.95	74.73	74.78	0.01
FreeBSD7.2	FireFox	75.01	12.64	12.64	No Fallback	No fallback	12.64	75.01	75.03	0.01
Fedora 13	FireFox	21.01	0.01	0.01	0.01	0.01	0.01	No fallback	No Fallback	0.01

Firefox version: 3.6.3
Safari version: 5.0
Chrome: 5.0.375

Code=0: no route to destination [RFC2463]
Code=1: communication with destination administratively prohibited [RFC2463]
Code=2: beyond scope of source address [RFC4443]
Code=3: address unreachable [RFC2463]
Code=4: port unreachable [RFC2463]
Code=5: source address failed ingress/egress policy [RFC4443]
Code=6: reject route to destination [RFC4443]

Table 1 : Time required by fallback

performance of a full fallback depends on the type of ICMPv6 message received. When a problem occurs with the OS or library, there are cases where fallback is instigated by HTTP communication [RFC 2068] (Fielding, 1997)], but not by HTTPS communication.

5.1.1.4. Events Causing Fallback

Fallback can be triggered by problems with IPv6 connectivity. For example, consider the following cases.

- **Use of unmanaged transition technologies**

Currently, 6to4 and Teredo are standard settings on the Windows OS. IPv6 connectivity is provided over the IPv4 Internet using tunnel technology. The communications quality of 6to4 in particular is generally poor, so stable IPv6 cannot be expected.

- **Destination host has AAAA record, but IPv6 is unavailable**

There are cases where an IPv6 address is registered as the server address, but IPv6 connectivity to the server is unavailable.

- **IPv6 Internet connectivity is unavailable**

Fallback is triggered on networks using IPv6 ULA address [RFC 4193] (R. Hinden, 2005) for VPN and IPv4 private addresses, but IPv6 connectivity is unavailable.

- **IPv6 connectivity environment fails**

Problems commonly occur in connected IPv6 network environments. For example, problems due to rogue RAs generated.

5.1.1.5. Ways to Detect Fallback

It is somewhat difficult for end-users to determine whether fallback has occurred. There are three main approaches:

1. determine that the access destination node has both IPv4 and IPv6 addresses registered in the DNS,
2. determine that the communication is running over an IPv4 or IPv6 network, and
3. determine that problems have arisen with IPv6 communication.

However, usually this information is generally concealed from the user. In order for IPv4 and IPv6 to coexist and transition smoothly from IPv4 to IPv6 without confusion, users should not have to concern themselves with communication details, yet this makes it difficult to isolate problems when they do occur. Let us consider these three approaches in greater detail.

In the first approach, many computer OSes have command interfaces that can be used to query the DNS. For example, a domain name can be converted to an IP address using the *nslookup command* on Windows 7 machines and using the *host command* on Mac OS X machines. Figure 7 illustrates the conversion commands to convert domain names to IP addresses for various operating systems. In all of these examples, requests are made to translate (or resolve) the domain name www.apnic.net, an external DNS server made available by Google (IP address

Figure 7 : Verifying IP Addresses

8.8.8.8), to the numeric IP address(es) assigned to the domain name. One can see that www.apnic.net has both IPv4 and IPv6 addresses: 202.12-29.211 and 2001:dc0:2001:11::211, respectively (normally one obtains the same result without specifying the DNS server). The second approach is to determine whether the actual communication is run over an IPv6 or an IPv4 network. Here there are a variety of techniques that might be used: accessing a website that displays differently depending on whether IPv6 or IPv4 is used, using a plug-in that displays whether the communication is run over IPv4 or IPv6, monitoring the network communication state of the OS, performing an actual packet dump, and so on.

- **Using a website**

There are sites that display differently depending on whether they were accessed over IPv6 or over IPv4. One such site is the Kame Project (http://www.kame.net) that is well-known among IPv6 users. Visitors to the site will see a non-mosaic version of a dancing turtle if they access the site over IPv6. The APNIC site mentioned earlier (http://www.apnic.net) displays the IP address of your own node

currently being used to access the site, so this can also be used to verify address information. In addition, there is a site at http://test-ipv6.com for testing IPv6 connectivity (the equivalent Japanese site is at http://test-ipv6.jp).

- **Using a plug-in**

Using the Firefox plugin ShowIP, one can display the IP address of a destination host. Figure 8 shows a screenshot of IPv6 connectivity environment when accessing the http://www.kame.net.

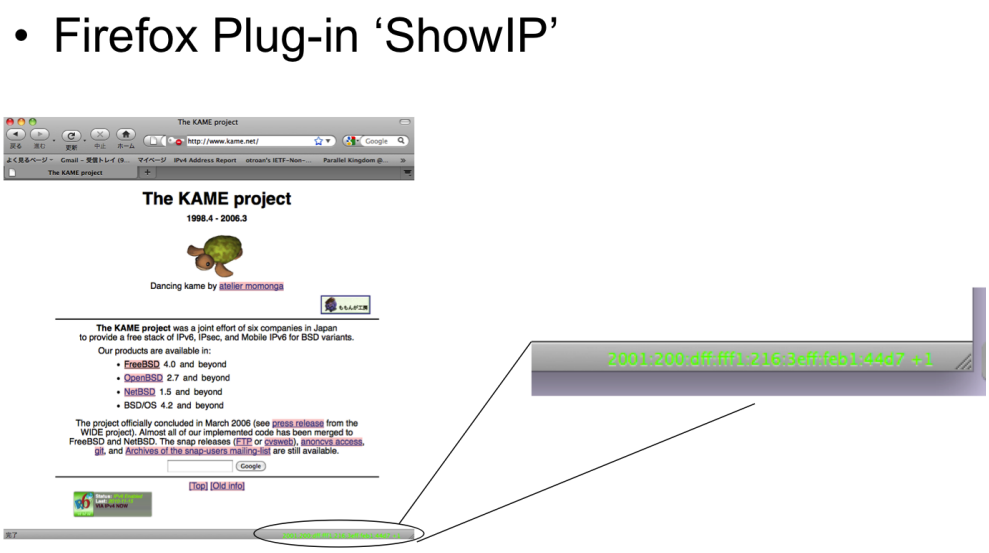

Figure 8 : Displaying address with a browser Plug-in

Clicking on the displayed information brings up the entire address of the server, which can be used to verify the address information of the destination node. This is a convenient indicator showing communication is running over IPv6, but at the time of writing Version 0.8.19 does not display the correct address when fallback has occurred. Apparently, when making the transition from an IPv6 environment to an internal environment, the information does not follow.

- **Monitoring OS network usage**

If communication is in progress, you can verify what protocol the OS is using with the 'netstat' command. This command is widely available across different operating systems including Windows OS, Mac OS, and UNIX-like OS. Figure 9 is a

screenshot showing the state when the OSes are running over IPv6 communication. When the state (stat) shows an entry of *Established* this means communication is currently taking place, and you can verify your own address, the destination node address, the protocol, and the port.

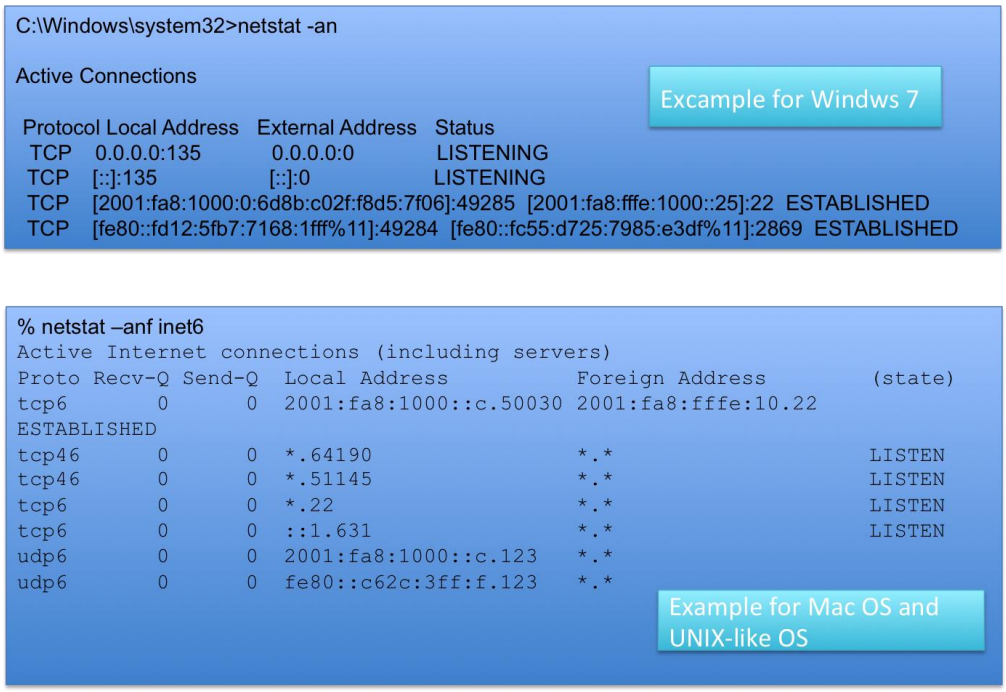

Figure 9 : Verifying communication status with a netstat command

- **Packet dump**

You can distinguish between different protocols by performing a packet dump while communication is in progress. In fact, you can detect whether fallback has occurred by observing the packet interaction. WireShark, tcpdump, and a host of other packet dumping tools are available, but with the enormous surge in communication packets for the many programs that exploit the network, it is difficult determine the purpose of a group of packets for most modern OSes. Figure 10 shows fallback data that was actually captured by Wireshark.

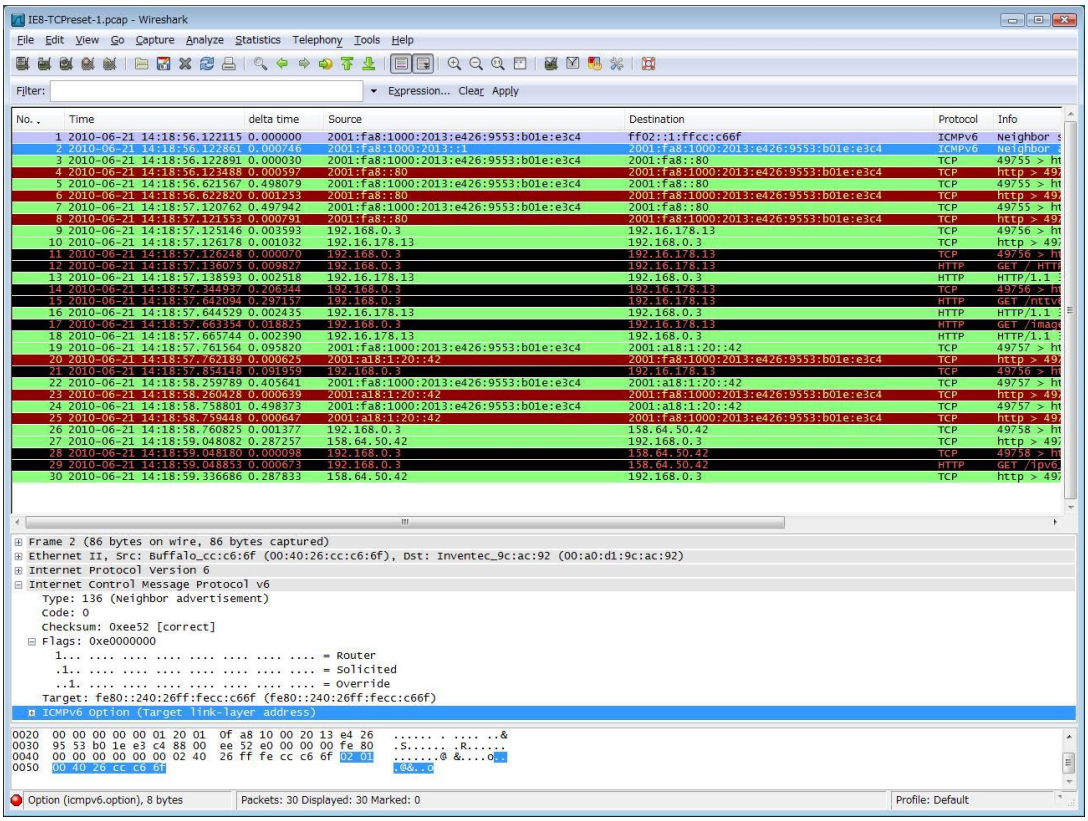

Figure 10 : Examples of Wireshark packet dump

Finally, the third way to determine if fallback has occurred is to employ a command to verify whether a destination host responds (ping/ping6), or a command to trace the route an IP packet follows to some Internet host (tracert/traceroute6). Figure 11 illustrates how the route is traced to server www.apnicnet using the traceroute6 command. Note that these commands use ICMPv6, but that ICMPv6 is sometime filtered by the server or an intermediate network for security reasons, so this technique may not work.

5.1.1.6. Dealing with the Fallback Problem

It is difficult to completely prevent fallback from IPv6 to IPv4, but there may be solutions available depending on the environment. When a technique such as 6to4 that cannot guarantee QoS is used to support IPv6 connectivity, the situation is markedly improved by switching over to IPv6 connectivity that does guarantee QoS. In cases where the destination host can be specified and you know the IPv6 connectivity is problematic, the situation can be avoided by implementing a mechanism for selecting the IPv6 address.

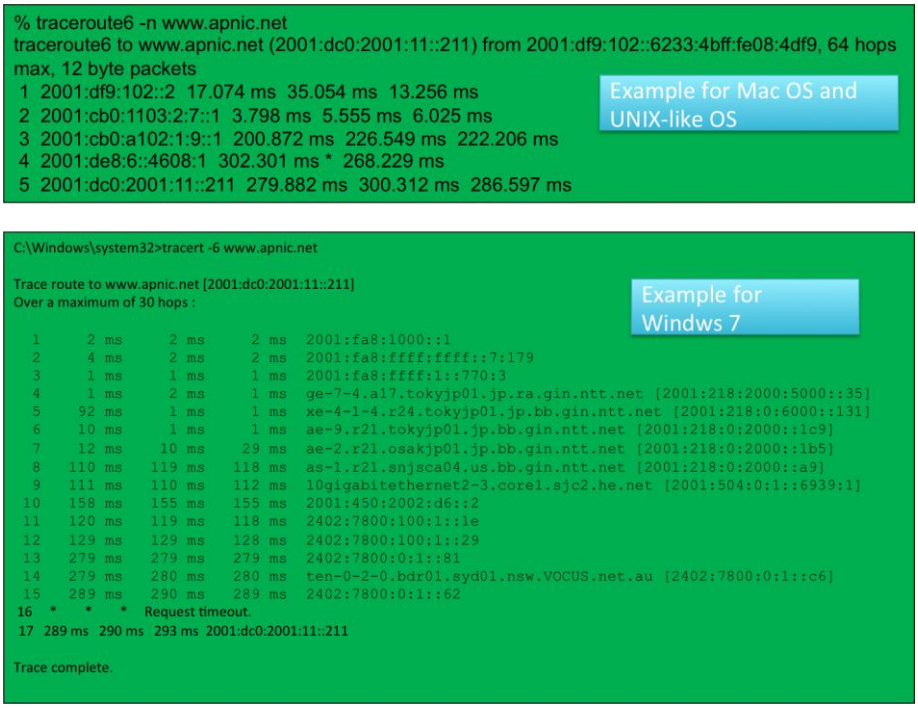

Figure 11 : Route verification with the traceroute command

Dealing with Fallback: IPv6 Address Selection Configuration Mechanism

It is assumed that IPv6 specifications will support multiple addresses, and the mechanism for selecting the address to be used is standardized in RFC 3484 (Draves, 2003). This mechanism has been implemented in Windows, FreeBSD, and most recently Linux, so by configuring an "address selection policy table,"

1. you can specify the connection address order in cases where the destination node has multiple addresses, and
2. you can select the source address of your own node that you want the destination node to use when your own node has multiple addresses

for both IPv4 and IPv6. Figure 12 illustrates how defaults are defined in the policy table in RFC 3484 and how the policy table works. In the default setting, the IPv6 address (::/0) has higher priority than the IPv4 address (::ffff:0:0/96), so IPv6 would be preferred. By switching the priority to IPv4, one can give priority to IPv4 for a specific address area to connect a VPN or closed network, or make other detailed setting changes.

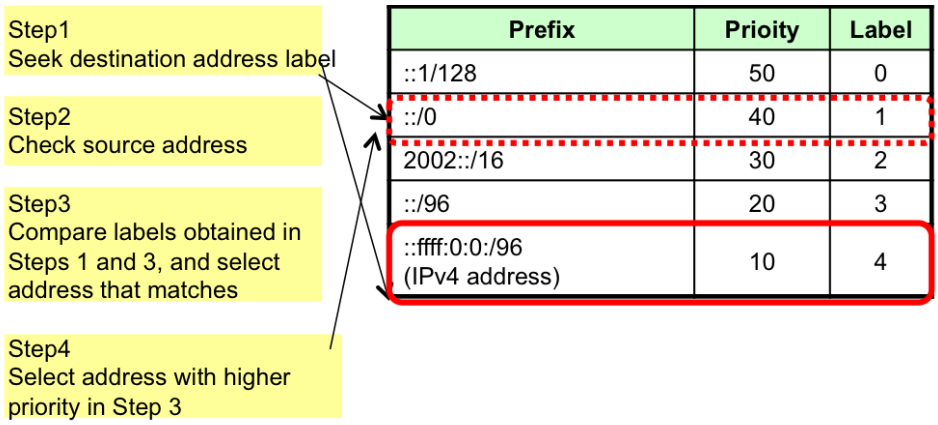

Figure 12 : Setting the RFC3484 address policy table

Dealing with Fallback: DNS Cache Server

In a recent implementation of BIND, the free DNS system developed by the Internet Systems Consortium (ISC), it can be set to not return an IPv6 address in response to AAAA resource record queries from clients. Using this function, one can prevent fallback by not responding to hosts that do not support IPv6 Internet connectivity with AAAA resource records.

Dealing with Fallback: Other

In cases where you know IPv6 connectivity is problematic, or in cases where a unique local address is used for an internal site network, a router deployed at the exit to the site can be set up to return ICMPv6 error messages that would enable fallback or the TCP session can be reset. Internet standards body IETF is also discussing ways of dealing with the fallback problem, and has proposed sending both IPv4 and IPv6 communications to destination nodes having both IPv4 and IPv6 addresses, and simply using whichever session responds first (D. Wing).

5.2. Rogue Router Advertisement (RA)

5.2.1. What is the Rogue RA Issue?

5.2.1.1. IPv6 Plug and Play Function and RA

When a host connects the network under IPv6, a standard *plug and play* mechanism is used to autoconfigure the information needed to communicate [RFC 4862] (S. Thomson, 2007). This mechanism is implemented using the Neighbor Discovery Protocol (NDP) [RFC 4861] (T. Narten E. N., 2007)]. NDP integrates IPv4 Application Service Providers [RFC 826] (Plummer, 1982), ICMP Router Discovery [RFC 1256] (Deering, 1991), and ICMP redirect. The protocol also adds neighbor unreachability detection, a number of other functions, and is the heart of IPv6. NDP implements these functions using the five types of messages shown in Table 2.

Message		Main functions
Router Solicitation	RS	Locates routers on the same link
Router Advertisement	RA	Notifies of router existence
Neighbor Solicitation	NS	Quires destination node layer 2 addresses on the same link
Neighbor Advertisement	NA	Notifies own layer addresses
Redirect		Notifies preferred next hop to hosts

Table 2 : Types of Neighbor Discovery messages

When IPv6 nodes connect to the network, they send out RS messages to find routers on the link. If routers are present on the link, they respond with RA messages. Figure 13 shows the packet format of an RA message.

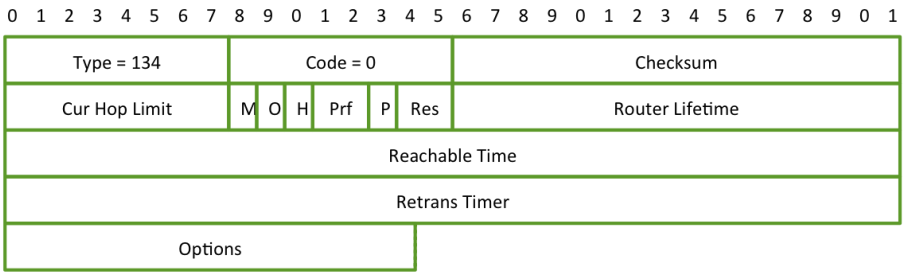

Field Name	Functions
Type	ICMPv6 message Type： 134
Cur Hop Limit	Indicates the default value that should be placed in the Hop Count field of the IP header
M Flag	When set, it indicates that addresses are available via Dynamic Host Configuration Protocol
O Flag	When set, it indicates that other configuration information (e.g. DNS server address) is available via DHCPv6
H Flag（RFC3775）	When set, it indicates that the router sending this Router Advertisement is also functioning as a Mobile IPv6 home agent on this link.
Prf Flag（RFC4191）	Indicates whether to prefer this router over other default routers. 00：Medium 01：High 11：Low
Router Lifetime	The lifetime associated with the default router in units of seconds. (A Lifetime of 0 indicates that the router is not a default router)
Reachable Time	The time, in milliseconds, that a node assumes a neighbor is reachable after having received a reachability confirmation.
Retrans Timer	The interval time, in milliseconds, between retransmitted Neighbor Solicitation messages.

Figure 13：RA Packet format

The RA contains a variety of information: whether the host can use the router that sent the RA as the default router, the time to live (TTL) and priority if it is used as the default router, timer for Neighbor Unreachability Detection (NUD), and information specified for address autoconfiguration.

Option information is also included in the RA. Table 3 shows typical parameters specified by the RA. Note that routers do not just send RAs in response to RS messages. As illustrated in Figure 14, routers periodically (default interval is 600 seconds) send out RAs to all nodes that are on the same link as the router.

Message	Function	Specification
Link layer address of source node	Indicate link layer address of a node sending the packet	RFC4861
Prefix Information	Determines if the destination host is on the same link, and the host uses it as its own address prefix .	RFC4861
MTU	Default MTU value used on the link .	RFC4861
Route information	Notifies host of received route for the next hop .	RFC4191
DNS Server	Notifies DNS server address	RFC6106

Table 3 : Parameters specified by RAs

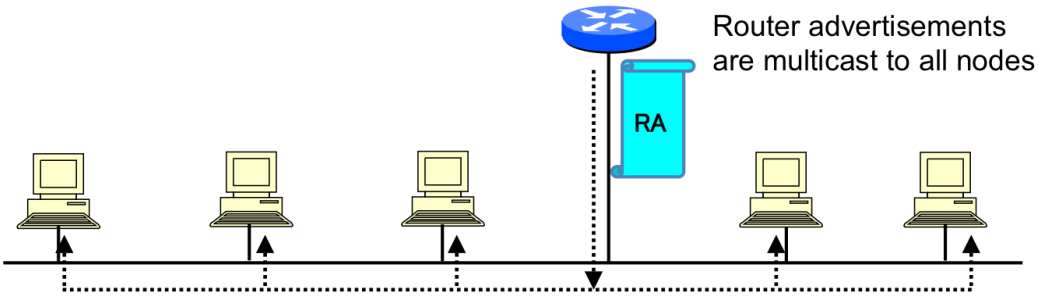

Figure 14 : RA delivery

5.2.1.2. Address Autoconfiguration Mechanism

RAs are also used to autoconfigure IPv6 addresses of hosts on the same link. Two types of address autoconfiguration are standardized in IPv6: stateless address autoconfiguration (SLAAC) [RFC 4862] (S. Thomson, 2007), and DHCPv6 [RFC 3315] (R. Droms, 2003). The M flag field (see Figure 13) in the RA message from the router tells the node which type of address autoconfiguration is used. When M is set to "1", the address is obtained from the DHCPv6 server, much the same as the address is obtained from the DHCP server in IPv4. When M is set to "0" and the prefix option flag is set for address configuration in the RA, the address is configured by stateless autoconfiguration. Stateless autoconfiguration is illustrated in Figure 15 (note that when a link local address is autoconfigured by the host, multiple addresses—i.e., the link local address and global address— are detected). The big advantage of IPv6 stateless address autoconfiguration is that it only

involves the RA message from the router and does not require any other server, and this user-friendliness accounts for its widespread use today in IPv6 networks.

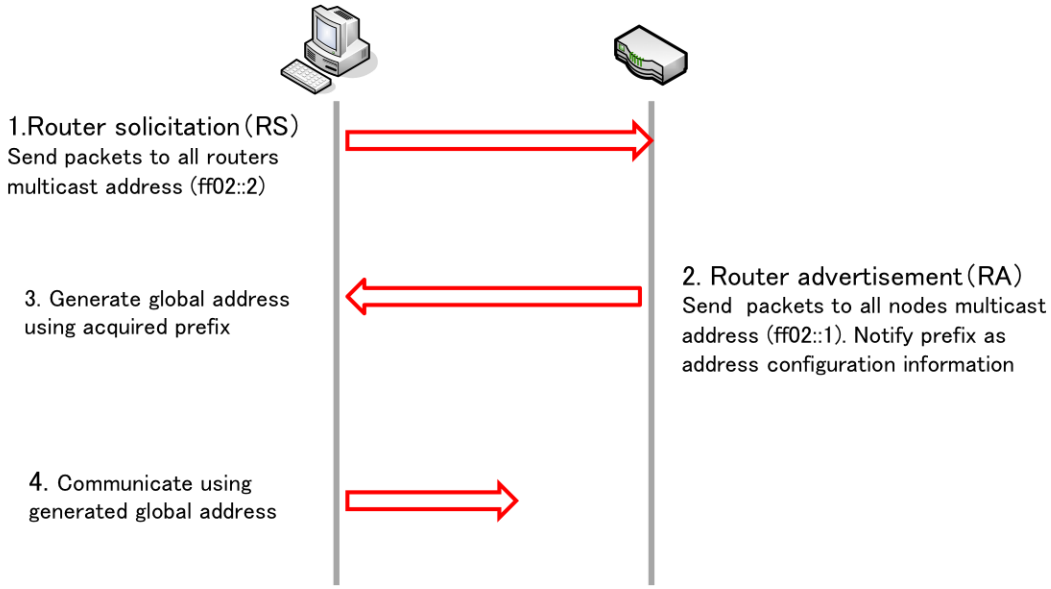

Figure 15 : Operations of stateless address autoconfiguration

5.2.1.3. Problems with Rogue Router Advertisements

It will be apparent from the previous section that the RA is critically important for implementing IPv6 plug & play and network configuration, but the RA is most commonly used to set the default router information and address information. This means that when unintended RA messages are delivered, this adversely affects all hosts on the same link and could cause the communication to fail (see Figure 16).

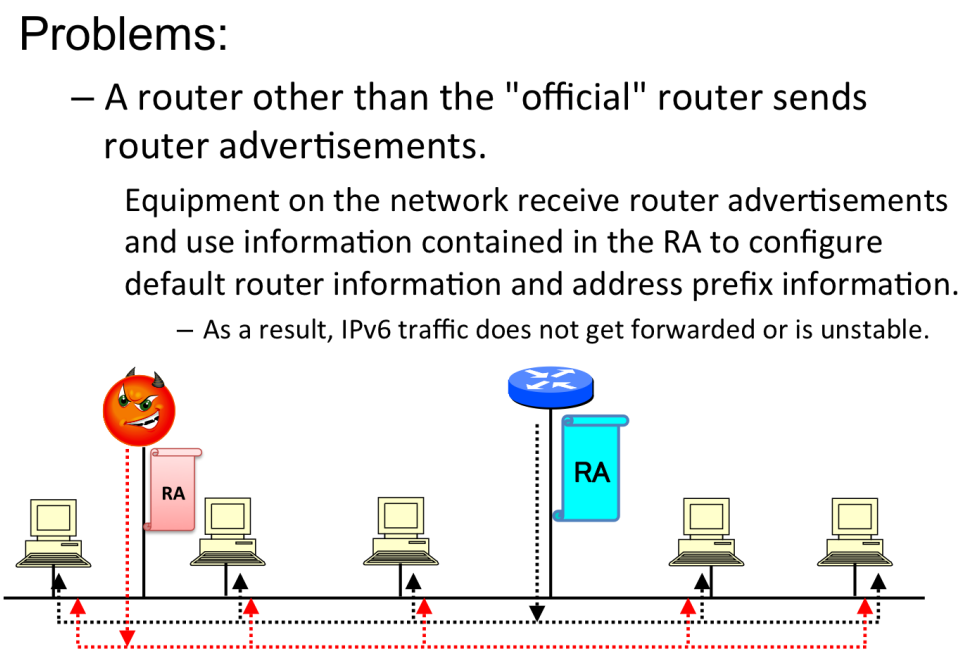

Figure 16 : Impact of rogue RAs

Here we refer to unintended RA messages as *rogue RAs*. There were rogue DHCP servers that could cause an incorrect default router or address to be configured in IPv4 as well, but in IPv6 the impact is much more pronounced since routers are capable sending messages to all hosts on the network.

Rogue RA problem is described in RFC6014[RFC6104] (Tim, et al., 2011) in detail.

5.2.2. Causes of Rogue RAs

Rogue RAs are attributed to the following causes.

1. Administrator Makes a Configuration Error

Because RAs can be configured so easily, RAs with the wrong address configuration data can be sent or RAs can be unintentionally sent out when a router is connected to the network. There are also some routers available (e.g., Cisco routers) that automatically send out RAs (including the address prefix attached to the address in the interface) when they are recognized as the default router. Moreover, when using a VLAN on a corporate network, cases have been reported of RAs from another segment getting mixed up with VLAN traffic due to an error in configuring the VLAN, or multiple segments on different networks inadvertently being connected

on Layer 2 via a switch. Cases have been reported of networks that appear to operate normally in the IPv4 environment, but failing when multiple RAs are mixed up when IPv6 is introduced.

2. User Configuration Errors

We are seeing increasing numbers of home routers that support IPv6, including some that offer 6to4 and other IPv6 Internet access functions for the transition period. If IPv6 is deployed on a network and operates at the same time as 6to4, this can result in multiple RAs and create problems. Cases have also been reported of RAs being unintentionally delivered by the Internet Connection Sharing (ICS) function that comes standard on the Windows OS. The Windows ICS function essentially turns a Windows machine into a router in IPv4 environments that uses IPv4 DHCP to distribute addresses to subordinate devices. When done intentionally, it operates much like a NAT router enabling multiple equipments to share a single Internet access line, and is a very convenient function (Figure 17). But when this function is turned on in a wireless LAN environment and you try to operate a 6to4, if you try to provide IPv6 Internet connectivity using 6to4, RA messages are sent out over the wireless LAN.

This cause many IPv6-ready devices on the wireless LAN to fail. This sort of confusion in wireless LAN environments due to ICS-induced rogue RAs frequently occurs in event networks, and other networks.

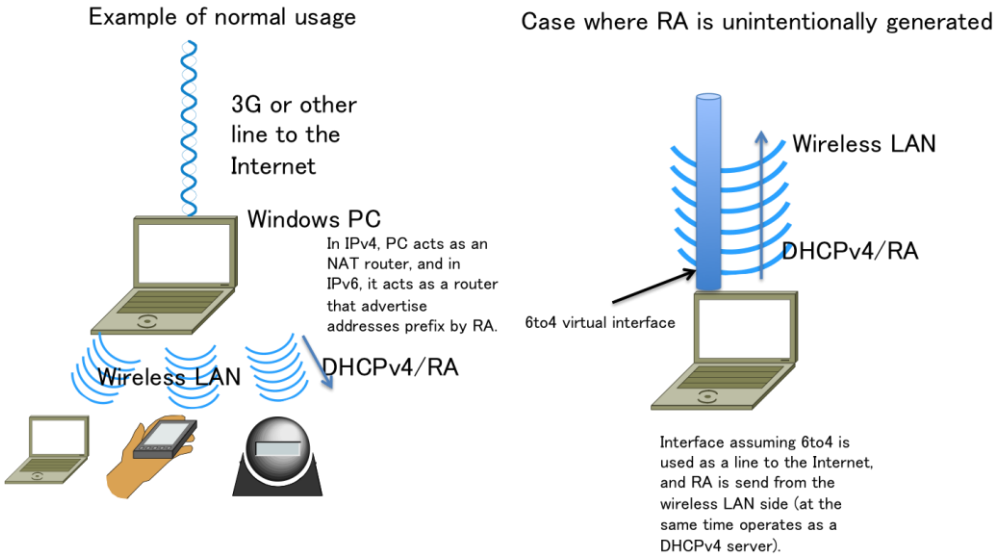

Figure 17 : "Internet Connection Sharing" function

There are other scenarios in which RAs originally from different segments can become intermingled: multiple segments can be linked by a PC with switching or bridging capabilities, a network cable could be hooked up incorrectly, and so on.

3. Intentional Advertisements

Deliberate distribution of rogue RA messages can subject networks to denial-of-service attacks. And since default router information can be disclosed, networks are also vulnerable to man-in-the-middle (MITM) attacks in the form of active eavesdropping and tampering with packets in transit. This kind of malicious rogue RA demands special vigilance, particularly in the case of networks that are accessed by unspecified large numbers of users.

5.2.3. Finding Rogue RAs

IPv6 communication generally becomes unstable when rogue RA messages are sent, and this is often the key to identifying the problem. When the problem is caused by 6to4 such as ICS-induced rogue RAs, the 6to4 address of the local host starts with 2002::/32 which is added onto the regular address, so in this case it is relatively easy to detect the rogue RA. Figure 18 shows an example of actual address information after an ICS-induced rogue RA has been generated. On Windows machines, you can verify an appended address using the netsh command.

```
% ifconfig -a
lnc0:
flags=108843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST,NEEDSGIAN
T> mtu 1500
        inet6 fe80::20c:29ff:fea1:xxxx%lnc0 prefixlen 64 scopeid 0x1
        inet6 2001:xx8:1000:0:20c:29ff:fea1:xxxx prefixlen 64 autoconf
        inet6 2001:xx0:640:f152:20c:29ff:fea1:xxxx prefixlen 64 detached autoconf
        inet6 2002:72xx:27xx:1:20c:29ff:fea1:xxxx prefixlen 64 autoconf
        ether 00:0c:29:xx:xx:xx
```

Figure 18 : Address state after receiving a rogue RA

Not only addresses but multiple default routers can also be registered, and this can be readily verified from the routing information. And as illustrated in Figure 19, several operating systems can display RA sender information that is attached to the incoming address prefix, so rogue RAs can also be detected in this way (J. Arkko, 2005).

Example for MacOS and BSD-like OS

```
speed(2)  ndp –p
2001:db8:1000:1000::/64 if=bge0
flags=LO vltime=infinity, pltime=infinity, expire=Never, ref=2
  No advertising router
fe80::%bge0/64 if=bge0
flags=LAO vltime=infinity, pltime=infinity, expire=Never, ref=0
  No advertising router
```

Example for Windows 7

```
C:\Users\foo>netsh interface ipv6 show siteprefixes
Prefix                  Lifetime    Interface
----------------------  ----------  ---------------------------
2001:0:4137:9e76::/64       infinite  Local Area Connection*
```

Figure 19 : Verifying a rouge RA has been sent

5.2.4. Dealing with Rogue RAs

RA message based stateless address autoconfiguration is extremely convenient and simple to use. Configuring the router involves practically no special settings, and we have heard reports of reduced setup costs to deploy IPv6 over large networks compared to IPv4 because so much time and effort are saved. Yet precisely because it is so easy and convenient, configuration errors have resulted in serious disruptions and the security risk has increased. In order to exploit RA's plug and play capabilities, a wide range measures have been implemented at standards and operations levels. Here we will highlight a few of these techniques.

- *Terminate autoconfiguration*

Many OSes allow you to turn off receiving RAs and configure address and default router information manually. In the server environment, autoconfiguration is not used and it is preferable to set parameters manually. Administrators of enterprise networks might also be tempted to suspend receipt of RA's and configure their networks manually, but configuring IPv6 addresses manually opens the door to human error, and extreme caution should be exercised.

- *Exploit RA snooping*

Special treatment of RAs can be implemented using L2 switching devices. This capability can be used to only accept RAs from designated ports in much the same way that DHCP snooping works in the IPv4 environment. Procedures are documented in RFC 6105 [RFC6105] (Eric, et al., 2011). L2 switches can also be set to filter out packets coming from suspect sources.

Currently these capabilities are only available in high-end switches, but we expect to see these features introduced in lower cost switches in the near term.

- *Secure Neighbor Discovery*

Secure Neighbor Discovery (SEND), which adds security to the Neighbor Discovery Protocol, has now been standardized [RFC 3791] (J. Arkko, 2005). While SEND provides some measure of protection against rogue RAs, it is complicated to configure (both router and host must be configured), and therefore is not widely implemented.

- *Router priority*

As we saw earlier in Figure 13, the default router priority can be set in the RA. If ICS (discussed earlier) is implemented, it is configured to send RA messages where the router priority is low, so by having the correct routers send high-priority RAs, the host can be prevented from changing the default router. Note however that this approach cannot address rogue RAs that are generated maliciously and of course only works on platforms supporting this option (e.g., Windows), so caution is called for in attempting to apply this method.

- *Rogue RA guard software*

When rogue RAs are detected, tools have become available enabling you to send an RA that blocks the rogue RA or the prefix information attached to the RA (rafixd developed by the KAME Project, RAMOND, and others), and in fact this software has already been applied to event and other networks. It is fairly effective for dealing with ICS vulnerabilities and improperly configured routers, but does not help much in the case of deliberate malicious attacks.

Restricting access to L2 networks and/or only allowing network resources to be used by authorized personnel can curb the generation of rogue RAs, but again will not help in the case of ICS when routers have been improperly configured.

Rogue RAs are not just problematic in terms of IPv6 communication failure due to improper configuration, in most cases it presents a serious security challenge to IPv6. Elevated security risk that comes with the deployment of simple plug and play merely highlights the tradeoff between convenience and security. Given the presence of RA messages for autoconfiguration in the current IPv6 specification, it is mandatory that users choose the best combination of measures that provide the level of security they need in constructing networks.

5.3. Path MTU Black Hole Issue

5.3.1. What is the Path MTU Issue?

5.3.1.1. IP Communication and Path MTU

The Internet Protocol operates over various links. MTU stands for Maximum Transmission Unit: the maximum data size (maximum IP packet length) that can be sent over a link at one time. For perspective, the MTU of the Ethernet that is extensively deployed in the LAN environment is 1,500 octets. As one can see in Table 4, the Ethernet MUT varies with the type of link (S. Deering, 1998). If the data to be sent exceeds the size of the MTU, then the source node breaks up or fragments the data.

Link	MTU (octets)
FDDI	4352
Ethernet	1500
IEEE 802.3	1492
SLIP	1006
IP over ATM	9180

Table 4 : Example of link MTU

When traffic is routed over the Internet, it may traverse various types of paths, and the link having the smallest MTU defines the largest packet size that can traverse this path without fragmentation. This is called the *path MTU*. Moreover, IPv6 standardizes the minimum MTU size to be supported by links as 1,280 octets [RFC 2460] (S. Deering, 1998).

If the IPv6 MTU size proves to be too large for some intermediate IPv4 subnet, packet fragmentation will ensue. This fragmentation is generally carried out by link ingress routers. As one can see in Figure 20, packets that are fragmented in transit are reassembled at the terminal node. To enable fragmentation, IPv4 provides a fragmentation field in the IPv4 header.

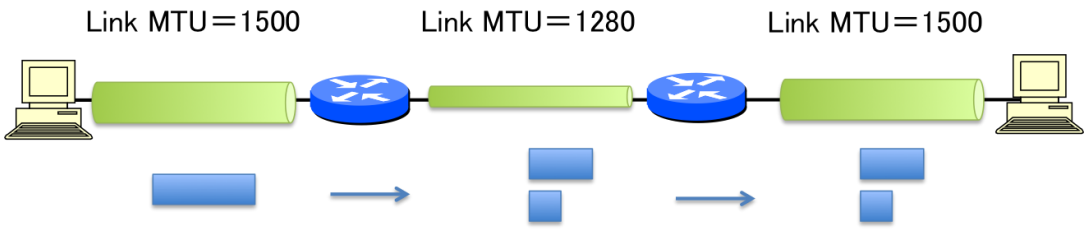

Figure 20 : IPv4 packet fragmentation

The IPv6 specification reduces the packet processing load on intermediate nodes by disallowing fragmentation at intermediate nodes and implementing packet fragmentation only at the source node (J.C. Mogul, 1990). Optimal efficiency is achieved by sending packets that match the actual path MTU of the link (the

maximum size of IP packets that can be transmitted over IPv4 without fragmentation). However, IP network paths vary for different destinations and packet transmission paths are subject to dynamic change, and this makes it difficult to determine the path MTU in advance. This led to the development of the Path MTU Discovery protocol [RFC 1191] (J.C. Mogul, 1990) [RFC 1981] (J. McCann, 1996), a technique for determining the path MTU between two IP hosts.

5.3.1.2. Path MTU Discovery

The Path MTU Discovery scheme in IPv6 is shown schematically in Figure 21. The client PC accesses Server A. The client sends packets to Server A at the maximum MTU supported by the link (1,500 octets). Router A intends to forward the packets on to the server, but because the link MTU only supports smaller size packets, an ICMPv6 unreachable "Packet Too Big" message is sent to notify the client that the router cannot forward the packet. Note that the ICMPv6 packet error message also tells the client the smaller link MTU that could be forwarded (e.g., 1,454 octets).

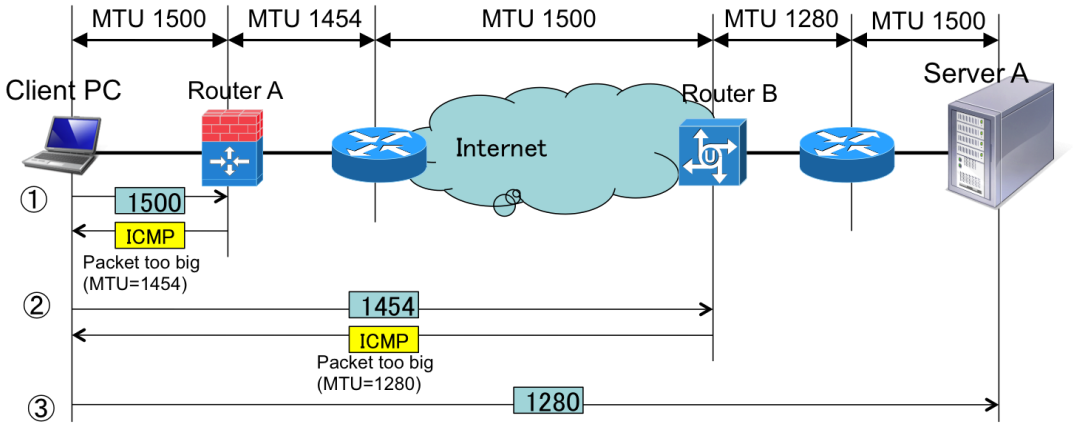

Figure 21 : Path MTU Discovery operation

The client PC receives the message, remembers the MTU size that can be sent, then if the client wants to proceed, instructs Server A to send the packet on at the smaller MTU size. Then, as one can see in Figure 21, since the size of the link MTU beyond Router B is smaller still, the "Packet Too Big" message is again sent to the client PC who now knows that the path MTU to Server A is 1,280 octets. Note that Path MUT Discovery is also defined in IPv4. A Don't Fragment (DF) flag is defined in the IPv4 header of outgoing packets, so if the flag is set packets will not be

fragmented, thus achieving the same result as when packets are sent by router as IPv6.

5.3.2. Causes of MTU Problems

Path MTU Discovery employs ICMPv6, as we saw in the previous section. For whatever reason, if the ICMPv6 "Packet Too Big" message does not get back to the client, then there is no way for the client PC to know what the path MTU is. For security reasons, currently all ICMP messages coming from outside IPv4 networks are blocked by firewalls, so we are now in the process of implementing a filtering solution. If we were to extend this same approach to IPv6, then we would lose Path MTU Discovery capability as shown in Figure 22 (E. Davies J.M., 2007). Essentially, this means that smaller packets from a client PC reach their destination, but packets exceeding a certain size do not. The cause is all the more difficult to detect since the initial ping or traceroute goes through successfully, since the default packet size for these commands is small.

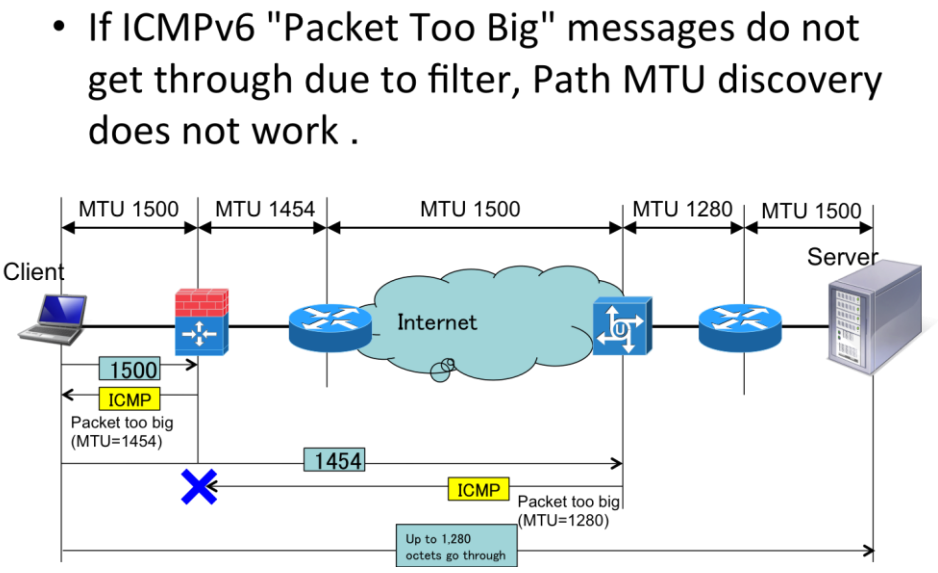

Figure 22 : Path MTU issued caused by filtering

Indeed, some of the problems recognized include (1) smaller email messages reach their destination but longer messages do not, (2) I can see simple webpages, but cannot access complex pages, (3) and when logging in from a remote terminal the login is successful, but the terminal hangs when displaying large files or a directory listing of many files. When propagating packets that are too large for the link MTU

seem to disappear (are swallowed up), the packets are said to have fallen into an *MTU black hole*.

One would expect this phenomenon to occur on IPv6 networks where packets are not fragmented in transit, but the same thing occurs on IPv4 networks when packets are sent with the *Do Not Fragment* flag set. Particularly when PPPoE is used to access the network, MTU black holes can occur because point-to-point connections have an MTU lower than that of standard Ethernet (when the PPP tunneling protocol is used over Ethernet).

5.3.3. Identifying MTU Problems

While it may seem that the MTU problems sketched above—sometimes packets get through and sometimes they do not—are inexplicable, there are ways to determine if the problem can be attributed to an MTU black hole. For example, a black hole can be detected by changing the packet size of the ping command using the –s option available in FreeBSD, Linux, Mac, and other OSes. Figure 4 illustrates the procedure used by the author when he ran into an actual path MTU problem (segment, address, and host names have been changed). Here the Traceroute6 command was used to check the route to the destination www.example.com. Ping6 signals are sent to various routers along the path, and if a response does not come back, MTU problems are pinpointed. If you are using Linux, the check is even easier using the tracepath6 command. In Figure 5 shows the output of the tracepath6 command.

Figure 23 : Detecting path MTU black holes

An alternative approach, when traceroute6 or ping6 are stymied by firewall filters, is to send an ICMPv6 "Packet Too Big" message (e.g., a packet dump) back to the source node and check to see if it gets through. Or you can access the web server using terminal software and see if you get a response.

5.3.4. Solving MTU Problems

Once you have localized where the path MTU problem is occurring, you can request that the ICMP filter be adjusted or the link MTU be adjusted (switching from tunnel connectivity to native connectivity). MTU problems frequently occur somewhere along the transmission path, and solving them often involves cooperation from others outside your own organization. Regarding adjustment of the filter, as we observed earlier in discussing how to deal with IPv6/IPv4 fallback, care must be taken since many features of IPv6 networks depend on ICMPv6 capabilities. When IPv6 is deployed, firewalls should be set up in such a way to enable essential ICMPv6 messages (see Table 5) to pass through. For a detailed discussion of ICMPv6 messages, refer to RFC 4890: "Recommendations for Filtering ICMPv6 Messages in Firewalls" (E. Davies, 2007).

Message	Type	Code numbers should be allowed through firewalls
Destination Unreachable	1	All
Packet Too Big	2	All
Time Exceed	3	0
Parameter Problem	4	1,2

Table 5 : ICMPv6 messages that should not be blocked

If setting up the firewall proves difficult or communication efficiency deteriorates, MTU problems can be avoided by setting the interface MTU to the minimum size for IPv6 (1,280 octets). This approach works well for large servers in circumventing MTU problems. If setting up the firewall proves problematic, MTU difficulties can be avoided by setting the MTU to the minimum size (1,280 octets) for the client interface. Configuration examples for Windows 7 and Free BSD are shown in Figures 24 and 25, respectively. In both examples, the MTU of the interface has been reduced from 1,500 octets standard for Ethernet to 1,280 octets (only users with super user or administrative privilege can make these changes on both OSes). Note that scaling back the interface MTU reduces the amount of data that can be sent at the same time, so performance may suffer. This approach has been used by large servers to circumvent MTU problems.

```
C:\Windows\system32>netsh interface ipv6 show interfaces

Idx  Met    MTU     State      Name
---  ----------  ----------  ------------  ---------------------------
 1       50  4294967295  connected     Loopback Pseudo-Interface 1
11       50        1280  connected     Local area connection*
10       10        1500  connected     Local area connection
14       50        1280  disconnected  isatap.nttv6.com
20       10        1280  connected     6TO4 Adapter

C:\Windows\system32>netsh interface ipv6 set interface "10" mtu=1280
OK

C:\Windows\system32>netsh interface ipv6 show interfaces

Idx  Met    MTU     State      Name
---  ----------  ----------  ------------  ---------------------------
 1       50  4294967295  connected     Loopback Pseudo-Interface 1
11       50        1280  connected     Local area connection*
10       10        1280  connected     Local area connection
14       50        1280  disconnected  isatap.nttv6.com
20       10        1280  connected     6TO4 Adapter

C:\Windows\system32>
```

Figure 24 : Setting the interface MTU on Windows 7

```
# ifconfig -a
em0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu 1500
        options=9b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING,VLAN_HWCS
        ether 00:0c:29:xx:xx:xx
        inet6 fe80::20c:29ff:xxxx:xxxx%em0 prefixlen 64 scopeid 0x1
        inet6 2001:db8:1000:0:20c:29ff:xxxx:xxxx prefixlen 64 autoconf
        inet 192.16.178.100 netmask 0xffffff00 broadcast 192.16.178.100
        nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
        media: Ethernet autoselect (1000baseT <full-duplex>)
        status: active
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> metric 0 mtu 16384
        options=3<RXCSUM,TXCSUM>
        inet 127.0.0.1 netmask 0xff000000
        inet6 ::1 prefixlen 128
        inet6 fe80::1%lo0 prefixlen 64 scopeid 0x4
        nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
# ifconfig em0 mtu 1280
# ifconfig -a
em0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu 1280
        options=9b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING,VLAN_HWCS
        ether 00:0c:29:xx:xx:xx
        inet6 fe80::20c:29ff:xxxx:xxxx%em0 prefixlen 64 scopeid 0x1
        inet6 2001:db8:1000:0:20c:29ff:xxxx:xxxx prefixlen 64 autoconf
        inet 192.16.178.100 netmask 0xffffff00 broadcast 192.16.178.100
        nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
        media: Ethernet autoselect (1000baseT <full-duplex>)
        status: active
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> metric 0 mtu 16384
        options=3<RXCSUM,TXCSUM>
        inet 127.0.0.1 netmask 0xff000000
        inet6 ::1 prefixlen 128
        inet6 fe80::1%lo0 prefixlen 64 scopeid 0x4
        nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
```

Figure 25 : Setting the interface MTU on FreeBSD

5.3.5. Dealing with MTU Problems

In TCP based communication, one can rely on Path MTU Discovery as little as possible by adjusting the Maximum Segment Size (MSS). This is how many IPv4 NAT routers are implemented, and particularly when PPPoE is used to access the Internet, the NAT router at the exit adjusts the TCP MCC in line with the link MTU in cases where the link MTU has been reduced. It is thus possible to reduce the Path MTU Discovery overhead by reducing the size of the access link MTU.

Some recent OSes have been developed that implement Path MTU Discovery at the TCP layer, and adjust the MSS accordingly (RFC 4821 (M. Mathis, 2007)). Be aware that whether a problem occurs or not depends on the operating system. These various MTU-related problems are summarized in RFC 2923 (Lahey, 2000).

6. Other Issues Associated with Deployment of IPv6

Now that we have discussed the first three big issues on the IPv6 Promotion Council's IPv6 Deployment Issues SWG's list of potential problems associated with the deployment of IPv6—IPv4/IPv6 fallback, rogue RAs, and PMTUD black holes—in this section we sill provide a more cursory examination of the remaining 26 issues on the SWG's list.

6.1. Problems Relating to the Domain Name System (DNS) when IPv6 is Deployed

One must specify an IP address destination when communicating over the Internet, but IP addresses are long numerical strings not suitable for easy memorization. This led to the development of the Domain Name System (DNS) that translates machine-readable Internet Protocol addresses into human-friendly alphabetical domain names such as www.example.com. Today the DNS is extensively used as an integral part of the Internet. It has a particularly important role in maintaining the address space which has been expanded in IPv6 (32-bit IPv4 address numbers have been extended to 128-bit IPv6 address numbers).

6.1.1. Nature of the Problem

1. In communication by IPv6, DNS name resolution may not be able to retrieve the correct AAAA resource record (data representing the IPv6 address) (S. Thomson, 2003). This makes it impossible to communicate over IPv6. The "desire to communicate by IPv6" is motivated by one of two considerations:

    - Desire to communication with a host that only has an IPv6 address, or
    - Communication over IPv6 offers significantly better quality than communication over IPv4 (e.g., using carrier-grade NAT).

2. Whether IPv4 or IPv6 is used to query the DNS server depends on current circumstances and implementation. Although users can anticipate the same result whether IPv4 or IPv6 is used, sometimes problems occur: you are taken to some other site than the one you are looking for or you are unable to communicate at all.

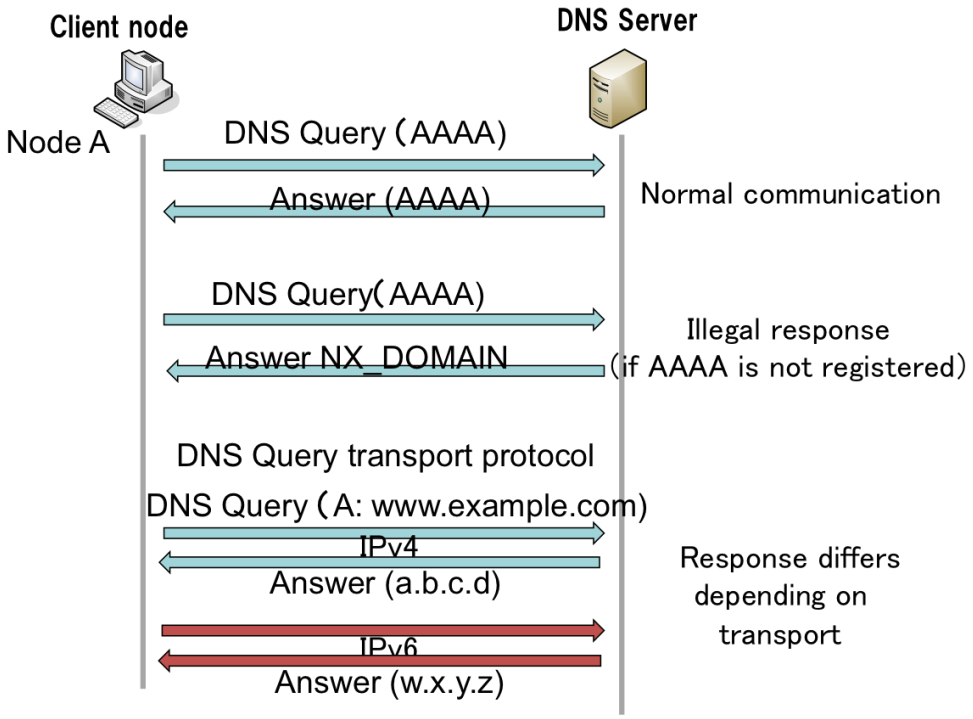

Figure 26 : DNS-related issues

6.1.2. Causes

- When just A and AAAA records are registered through implementation of a DNS server, a Non-Existent Domain Error (NXDOMAIN) flag is returned when a resource record is queried. This means that in some cases either IPv6 or IPv4 communication is not possible (Issue 1).

  - Implementation differs depending on the OS and application. The priority and resolution order of AAAA and A records are different.

  - In DNS implementations of load balancing, some DNS servers do not reply to AAAA queries (see References).

  - Not only do different OSes behave differently, there are cases where the resolver acts differently for different versions of the same OS (e.g., Windows).

3. AAAA resource records may not be returned depending on how the DNS is configured (e.g., BIND configuration). In these cases, users cannot use IPv6 even if that is their preference (Issue 1).

   - DNS server is implemented in line with query transport (AAAA record is returned when queried by IPv6). Communication is supported by dual-stack services, but if the DNS server can only notify IPv4 addresses, it will only support IPv4.

   - There are client OSes that are only capable of handling IPv4 for DNS transport (e.g., WindowsXP).

4. When a DNS server is not accessible, access is delayed or impossible as a result of DNS name resolution (Issues 1 and 2).

   - If an OS has the ability to look up both IPv4 and IPv6 addresses on the DNS server, IPv6 has priority. If communication with a IPv6 DNS server is unavailable, name resolution takes more time since it must switch over to a IPv4 DNS server.

5. DNS proxies and cache servers are unable to transmit AAAA resource records or resolve IPv6 addresses. This is particularly relevant in the case of home routers, most of which are implemented with a simple DNS proxy function. Inability to resolve address can be attributed to the following (Issue 1):

   - Implementation is unable to correctly handle old AAAA resource records.

   - Cannot resolve addresses because unable to properly handle DNS packets longer than 512 bytes. AAAA resource records are easier to generate because they are 12 bytes longer than A resource records.

   - Where IPv6 is used to implement a closed network, AAAA resource records are not returned intentionally in order to not disclose the IPv6 address.[1]

[1] IPv6 Home Router Guideline, Version 2.0 (2010) states that routers "must be capable of transparently processing resource records (RRs) regardless of type.

6.1.3. Security Considerations

Data and services supporting both IPv4 and IPv6 must reconcile the DNS with applications.

- Lack of compatibility introduces the same security vulnerability as cross-site scripting (XSS).
- Exposure to phishing attacks.

6.1.4. Solutions

- Addressing the first issue, providers must avoid using defective DNS implementations, verify DNS-related configurations, and so on. There are cases where AAAA resource records are deliberately not returned. Users may be able to deal with this situation by resolving the address using Google public DNS server 8.8.8.8 or some other public server, but this runs the risk of causing the fallback-related issues discussed earlier, or some other problem.
- Turning to the second issue, this is addressed by maintaining a cache DNS server that is capable of resolving both IPv4 and IPv6 transport, and configuring user terminals so they notify both IPv4 and IPv6 addresses.
  - ➢ It is assumed that User Datagram Protocol (UDP) fragments can be transferred.
    - ✧ They are permitted to pass through end node firewalls.
    - ✧ Note that some home broadband routers do not deal with fragments.
  - ➢ Configure so TCP queries can be sent and received.
    - ✧ Allow TCP level DNS communication to pass through firewalls.
  - ➢ Addresses may be resolvable using 8.8.8.8 or other public DNS server.

6.1.5. References

RFC 3596, "DNS Extensions to Support IP Version 6" (S. Thomson, 2003)

RFC 3901, "DNS IPv6 Transport Operational Guidelines" (A. Durand, 2004)

RFC 4294, "IPv6 Node Requirements" (Loughney, 2006)

RFC 4472, "Operational Considerations and Issues with IPv6 DNS" (A. Durand, 2006)

RFC 4942, "IPv6 Transition/Co-existence Security Considerations" (E. Davies, 2007)

draft-ietf-6man-node-req-bis: "IPv6 Node Requirements RFC 4294-bis" (Work in Progress)

"JPRS IPv4 address exhaustion and DNS: IPv6-aware DNS," http://www.kokatsu.jp/blog/ipv4/data/interop2009/11_JPRS_TAKASHIMA.pdf

IPv6 Promotional Council, "IPv6 Home Router Guideline (ver. 2)," http://www.v6pc.jp/jp/upload/pdf/v6hgw_Guideline_2.0.pdf

6.2. Captive Portal and DNS Problems (IPv6 Uninstall at Hotels)

6.2.1. Nature of the Problem

When a user opens a browser and tries to access the Internet from a hotel room or other public accommodation, he is likely to be redirected to a web page which may require authentication and/or payment. It has been observed that, for some implementation of this scheme, some OSes are unable to communicate over IPv4 if IPv6 is turned on. To cope with this problem, some hotels have adopted a standing policy to "uninstall IPv6," and this is thwarting or at least slowing the deployment of IPv6.

6.2.2. Causes

- The DNS cache of hotel captive portals returns an A record to AAAA queries.

- Windows XP returns an A response to AAAA queries, so fallback based on the A response of the URL of the redirect display cannot occur and communication is stymied.
- Windows since Vista, Linux, MacOSX, and FreeBSD do not use an A response to AAAA queries; instead, these OSes only use A responses to A queries, so they operate normally without a problem.
- For a more detailed explanation of this issue and insight into the operating environment, point your browser here: http://v6fix.net/docs/hotel.html.ja.

6.2.3. Solutions

- Dealing with the problems associated with network equipment installed in hotels and other similar facilities is fundamentally difficult.
- As a stop-gap measure there is no alternative but to temporarily uninstall client IPv6 communication functions.

6.3. Poor Quality Tunnels, Transition Technology Related Issues (Teredo, 6to4)

6.3.1. Nature of the Problem

The transition to IPv6 is being promoted by the use to tunnel technology over the IPv4 Internet, by defining transition technologies (e.g., 6to4 and Teredo) providing access to the IPv6 Internet, and by providing special connectivity equipment (Internet relay routers) at no charge to users. But relay routers have no QoS guarantee which means that communications may be carried over poor-quality paths, so we are likely to see problems with quality and inability to get through. There are also relay routers in service that are poorly or inadequately managed, and this can lead to breakdowns in communication, eavesdropping, and other problems.

6.3.2. Causes

- Some client OSes (e.g., Windows XP, Vista, and 7) are configured to automatically exploit these capabilities.
    - Users perceive communication based on Teredo or 6to4 addresses as slow.
- According to the provision of RFC 3484 [Draves, 2003], IPv4 is preferred over tunneling when accessing dual-stack servers on Windows XP, Vista, and 7 machines. Regarding 6to4, there is no particular problem except when only IPv6 communication is used, since a tunnel address is not provided if a IPv6 global unicast address is given.
    - Earlier versions of MacOSX (up to Version 10.6.4) were susceptible to this problem since 6to4 is included but IPv6 is preferred.

Regarding operation of the policy table, refer to the recommended solutions to the "IPv6/IPv4 Fallback Problem."

6.3.3. Security Considerations

- Caution: the IPv4 global address of the NAT router employed by the user as well as the L4 port number are used by the Teredo address.
    - Note that NAT routers allow packets from outside to the address/port used, thus opening access to the outside.
- Caution: 6to4 addresses are encapsulated in IPv4 address, so beware of attacks coming via IPv4.
- If a relay router cannot be trusted, this should raise concern that packets may be compromised.
    - There is currently no way of guaranteeing that a router is trustworthy.

6.3.4. Identifying the Problem

Check to see if a tunnel interface is being used. Every OS provides a way to check the status of the network interface.

- For example, in Windows this is done using the ipconfig command + packet capture, or ipconfig + netstat, as shown in Figure 27.

```
C:¥Users¥foo> ipconfig /all
…
Tunnel adapter Local area connection:

   Connection-specific DNS Suffix . :
   Description . . . . . . . . . . . . . . . : Microsoft 6to4 Adapter
   Physical Address. . . . . . . . . . . . . : 00-00-00-00-00-00-00-E0
   DHCP Enabled . . . . . . . . . . . . . . . : No
   Autoconfiguration Enabled. . . . . . . . . . . . : Yes
   IPv6 Address . . . . . . . . . . . . . . . : 2002:xxxx:e30c::813c:e30c(Preferred)
    Default Gateway . . . . . . . . . . . . . : 2002:c058:6301::c058:6301
    DNS Server . . . . . . . . . . . . . . . : xxx.xx.5.12
                                               xxx.xx.5.13
   NetBIOS over TCP/IP . . . . . . . . . . . : Disabled
```

Figure 27 : Checking if 6to4 isused on a Windows machine

- MacOS and other UNIX-like systems: ifconfig, netstat command, packet capture

6.3.5. Solutions

If it is known that a tunnel is being used as a transition technology, one solution would be to suspend use of the tunnel and go with high-quality IPv6 communication provided by a commercial service. In its most recent upgrade, MacOS X gives users the option of disabling 6to4.

6.3.6. References

RFC 3056 "Connection of IPv6 Domains via IPv4 Clouds" (Carpenter, 2001)

RFC 3068 "An Anycast Prefix for 6to4 Relay Routers" (Huitema, An Anycast Prefix for 6to4 Relay Routers, 2001)

RFC 3964 "Security Considerations for 6to4" (P. Savola, 2004)

RFC 4380 "Teredo: Tunneling IPv6 over UDP through Network Address Translations" (Huitema, Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs), 2006)

RFC 6081 "Teredo Extensions" (Thaler, 2011)

draft-ietf-v6ops-6to4-to-historic: "Request to move Connection of IPv6 Domains via IPv4 Clouds (6to4) to Historic status" (work in progress)

draft-ietf-v6ops-6to4-advisory: "Advisory Guidelines for 6to4 Deployment" (work in progress)

6.4. Different QoS at Dual-Stack Sites, Different QoS of IPv4 and IPv6

6.4.1. Nature of the Problem

The quality and service processing capability of IPv4 and IPv6 may be different. When accessing sites with both A and AAAA resource records, the response time for IPv4 and IPv6 will differ and you may not be able to access either IPv4 or IPv6.

- It is very apparent when IPv6 runs slow, but when IPv4 and IPv6 are supported by different servers, sometimes IPv4 runs slow when the IPv4 server is overloaded and the IPv6 server is empty.

6.4.2. Causes

These problems are attributed to the following.

- Transmission lines on the server side
  - Throughput of IPv4 and IPv6 lines is different
  - The round-trip time of IPv4 and IPv6 differs (IPv6 is more pervasive overseas).
- In some cases performance of firewalls, server OSes, applications, and so on suffers when IPv6 is used.
- AAAA resource records are registered in the DNS so they are returned even though IPv6 is not enabled.

- Routing or other problems occur while packets are carried over the network.
- A significant difference in communication quality on the last mile (including tunnel) when either IPv4 or IPv6 uses the tunnel.
- Protocol processing speed varies due to problems with the implementation of the client OS.
- When equipment is divided by protocol on the server side, a difference in quality emerges due to different processing capacity and number of accesses of the servers.

6.4.3. Identifying the Problem

Perform accesses over both IPv6 and IPv4, and compare communication speed. Note that the "poor quality tunnel" problem may also be present, so it is necessary to keep these problems separated.

6.4.4. Solutions

- End users
  - Use higher quality protocol (reduce IPv6 priority, access over IPv4, etc.).
  - Complain to service provider.
- Service providers
  - Make efforts to minimize differences in quality between IPv4 and IPv6 services offered (as dual-stack user terminals become more common and IPv6 communications are preferred on those terminals, particular attention must be given to the quality of IPv6 services provided).
  - Improve transmission quality of IPv6 services.
  - Improve robustness of firewalls and server equipment.

6.5. Address Selection Related Problems (Multi-Prefix Problems)

6.5.1. Nature of the Problem

Multiple IPv6 addresses can be allocated to equipment interfaces in IPv6. When users with multiple global IPv6 prefixes attempt to communicate, they may not be able to communicate due to the source address selected (see Figure 28).

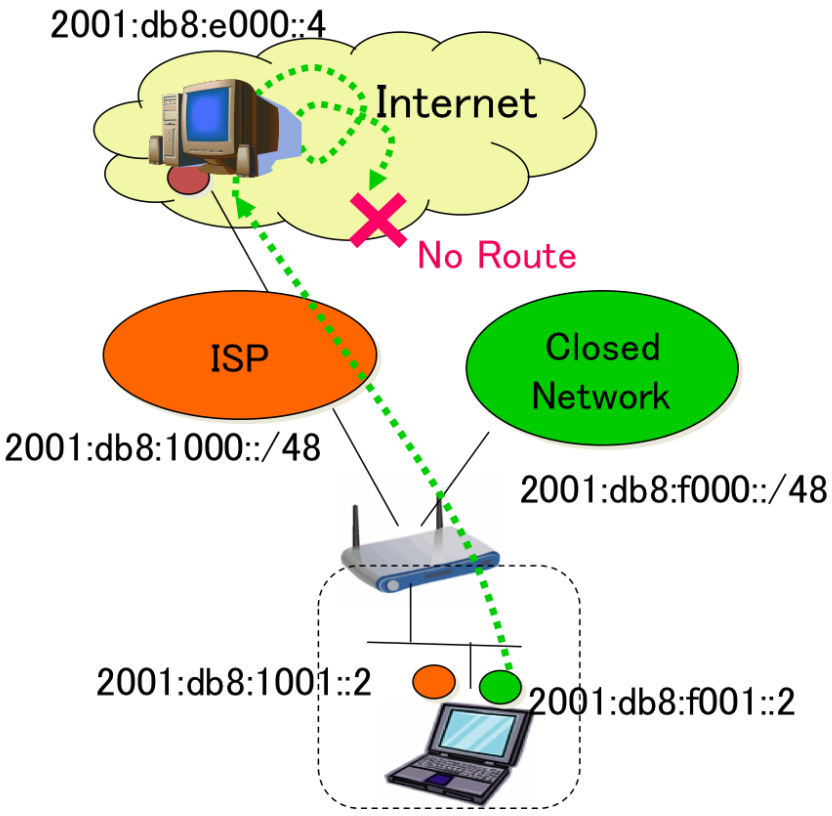

Figure 28 : Communication failure due to multi-prefixes

6.5.2. Causes

Source address selection errors for terminals that have multiple IPv6 prefixes.

- When packets originating from source addresses other than service providers allocating IPv6 addresses are sent to service providers allocating prefixes, addresses other than those allocated by Unicast Reverse Path Forwarding (uRPF) (Savola, 2004) are filtered, so communication sometimes fails. In the

case of a closed network service, communication packets will not get through if a source address selection error occurs even if it is not filtered.

Consider the following examples of users having multiple IPv6 prefixes.

- FLETS Service is used in combination with IPv6 service provided by an ISP.
- 6to4-enabled IPv6 router (AirMac, etc.) is used at the same time as IPv6 service from an ISP.

6.5.3. Identifying the Problem

- Check to see if multiple IPv4 addresses have been allocated to a single terminal interface. Use the ipconfig command on Windows machines, and the ifconfig command on UNIX-like and Mac OS machines. Figure 29 illustrates how the check is done on a FreeBSD machine.

```
%  ifconfig –a
lnc0: flags=108843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST,NEEDSGIANT>
mtu 1500
        inet6 fe80::20c:29ff:fea1:xxxx%lnc0 prefixlen 64 scopeid 0x1
        inet6 2001:xx8:1000:0:20c:29ff:fea1:xxxx prefixlen 64 detached autoconf
        inet6 2001:xx0:640:f152:20c:29ff:fea1:xxxx prefixlen 64 detached autoconf
        inet6 2002:72xx:27xx:1:20c:29ff:fea1:xxxx prefixlen 64 autoconf
        ether 00:0c:29:xx:xx:xx
```

Figure 29 : Multi-prefix state interface

- Perform a packet dump.
  - Verify the source address of the communications packet is correct.
  - Verify no communications error messages have been sent by ICMPv6.

6.5.4. Solutions

- Configure terminal to select the correct address. Use the policy table in RFC 3484 to determine the correct settings, which are already defined for Windows and UNIX-like OSes (and will be defined for the coming version (LION) of the MacOS). Refer to "IPv6 Address Configuration for Fallback Mapping" for configuration details.

- Use only one IPv6 address for selecting services.

6.5.5. References

RFC 3484: "Default Address Selection for Internet Protocol version 6 (IPv6)" (Draves, 2003)

RFC 5220: "Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules" (A. Matsumoto, 2008)

draft-ietf-v6ops-ipv6-multihoming-without-ipv6nat: IPv6 "Multihoming without Network Address Translation" (Work in Progress)

6.6. Problems with False Recognition and IPv6-Ready Routers that Only Support IPv6 Bridge Functions (IPv6 Pass-Through Functions)

6.6.1. Nature of the Problem

Users intent on accessing IPv6 Internet connectivity services are purchasing routers marketed as "IPv6-ready routers" but that are actually only equipped IPv6 bridge capability. There is a good chance that users will not be able to access IPv6 connectivity services using these routers.

6.6.2. Causes

There is no unified definition of "IPv6-ready," so there are varying perceptions of the meaning of the term.

6.6.3. Analysis

Users buy what they think is a IPv6-ready router, only to find that it's only capable of handling IPv6 bridge traffic, and does not support IPv6 connectivity services.

6.6.4. Solutions

Spread awareness that so-called "IPv6-ready routers" may not be capable of handling the full range of IPv6 connectivity services.

- We should be able to mitigate this problem by clearly defining exactly what "IPv6-ready router" means, and by encouraging vendors to disclose the specific IPv6 connectivity services their products can handle in a format that can be made publicly available.

6.6.5. References

- "IPv6 bridge functionality," http://bb.watch.impress.co.jp/cda/koko_osa/18406.html
- If you do a Google search on the term "IPv6-ready router," you will get many hits for routers featuring IPv6 bridge function.
- Table showing the status of FLET'S HIKARI network-compliant routers (non-NTT broadband routers), http://flets.com/next/list_router.html

6.7. Problems with Bridge Filters in IPv6-Ready Routers

6.7.1. Nature of the Problem

Home routers with IPv6 pass-through function do not support IPv6 filtering, and therefore present a security risk.

6.7.2. Analysis

Home routers with "IPv6 pass-through function" can only access IPv4 through NAT and only support IPv6 through simple bridging without filtering capability. The fact that IPv6 may not provide the same level of security as IPv4 could be problematic. While the same problem could occur on IPv4 as well, that fact that so many routers have been sold and continue to be sold featuring a "IPv6 pass-through" function" means the impact is far greater on the IPv6 environment.

6.7.3. Identifying the Problem

A determination can be made as to whether a problem actually exists after examining the home router "IPv6 pass-through" function and assessing the environment and services for which it is used.

6.7.4. Solutions

- Deploy home routers with both IPv6 filtering and the IPv6 pass-through function (currently there are no low-cost routers offering this combination of features).
- Be aware in using routers with IPv6 bridge function that they do not support filtering capability (other security measures must be implemented).

6.7.5. References

- "Bridge function security issues," http://121ware.com/product/atermstation/product/function/33.html

6.8. DNS Registration Issues ("DNS Registration, Reverse Lookup, Forward Lookup, DDNS")

6.8.1. Nature of the Problem

In IPv6, addresses are set in client equipment primarily by an automatic address configuration feature (address auto-configuration). This raises issues for forward DNS lookup and reverse DNS lookup.

- If DHCPv6 is used, it works the same way as DHCP in IPv4.
- If stateless address autoconfiguration (SLAAC) is used, this presents problems with the address registration.
- If a temporary address is used, the registration method and the pro and cons of registration become problematic.

There is no generalized method of forward and reverse DNS lookup registration in IPv6 and dual-stack environments. Each provider implements his own solution, and this lack of standardization presents a huge obstacle to providers and users employing the same technology.

6.8.2. Solutions

The following four strategies could offer a solution, and should be followed up and studied by the industry.

- Do not register forward and reverse lookups.
- Automatically generate resource records to be registered.
- Use a wildcard record.
- Use Dynamic DNS (DDNS).

The following factors would affect the implementation.

- Log analysis would become much more difficult without the ability to perform reserve lookups.
- Access control would become more difficult.
- It becomes more difficult to verify with paranoid checking.

6.8.3. References

"IPv6 Compliance and IPv6 Functional Use Guideline for IPv6 Terminal OSes,"
http://www.v6pc.jp/pdf/v6TermOs_2006Guideline-0.pdf

"Reverse DNS in IPv6 for Internet Service Providers,"
http://datatracker.ietf.org/doc/draft-howard-isp-ip6rdns/

"IPv6-compliant DNS,"
http://v6ops-f.jp/index.php?plugin=attach&refer=meeting%2F%C2%E81%B2%F3IPv6%A5%AA%A5%DA%A5%EC%A1%BC%A5%B7%A5%E7%A5%F3%A5%BA%A5%D5%A5%A9%A1%BC%A5%E9%A5%E0&openfile=v6ops-f-dns-ito.pdf

"Considering DNS Configuration in the IPv6 Age,"
http://v6ops-f.jp/index.php?plugin=attach&refer=meeting%2F%C2%E81

%B2%F3IPv6%A5%AA%A5%DA%A5%EC%A1%BC%A5%B7%A5%E7%A5%F3%A5%BA%A5%D5%A5%A9%A1%BC%A5%E9%A5%E0&openfile=v6ops-f-dns-shin.pdf

"IPv6 Reverse Lookup Auto-Generation DNS Server,"
http://v6ops-f.jp/index.php?plugin=attach&refer=meeting%2F%C2%E82%B2%F3IPv6%A5%AA%A5%DA%A5%EC%A1%BC%A5%B7%A5%E7%A5%F3%A5%BA%A5%D5%A5%A9%A1%BC%A5%E9%A5%E0&openfile=2_06_v6rev.pdf

"One Implementation of a IPv6 Reverse DNS Server,"
http://member.wide.ad.jp/~fujiwara/v6rev.html

"IPv6 Reverse Zone Maker,"
http://negi.ipv6labs.jp/shared/ipv6_reverse-zone-maker.html

6.9. Security and Filtering Issues (ICMP Filtering Problems, etc.)

6.9.1. Nature of the Problem

If the same filtering and other policies used in IPv4 were applied to IPv6, this would unintentionally create communication barriers. The net effect would be to prevent Internet communication in the IPv6 environment.

6.9.2. Causes

IPv6 makes frequent use of ICMPv6 for communication. Configuring filters without understanding this crucial aspect of IPv6 communication could cause IPv6 communication to fail.

6.9.3. Analysis

- ICMPv6 filter (Path MTU Discovery aware)
  - IPv4 connections are commonly implemented in recent years to filter incoming ICMP packets from outside the network to prevent DoS attacks. It appears that IPv6 connections are following suit by filtering incoming ICMPv6 packets. The only problem with this approach is that it screens out ICMP error packets ("Packet Too Big" messages), which causes MTU Discovery to fail on subordinate terminals (the Path MTU Black Hole problem).

6.9.4. Identifying the Problem

Check to see if ICMPv6 packets are being filtered out at ingress routers.

6.9.5. Solutions

Change filters in Internet connection routers.

6.9.6. References

RFC4890, "Recommendations for Filtering ICMPv6 Messages in Firewalls" (E. Davies, 2007)

6.10. IPv6-Ready Mail System Issues (Sending and Receiving Mail)

The Simple Mail Transfer Protocol (SMTP) for outgoing mail transport over the Internet and the Post Office Protocol (POP) and Internet Message Access Protocol (IMAP) used by mail e-mail clients (MUAs) for receiving messages do not depend on different types of protocols, yet special precaution is called for in implementing mail, filtering, and anti-spam techniques on IPv6 networks.

6.10.1. Issues Involved in Sending and Receiving Mail

In making the mail systems compatible with IPv6, you may lose the ability to send and receive mail.

6.10.2. Analysis

- Issue 1: Mail exchange (MX) records support both AAAA and A hosts, so the MTA may create problems with sending and/or receiving mail if only AAAA hosts are preferred.
- Issue 2: One can easily envision situations where mail could not be sent depending on the recipient.

6.10.3. Causes

- Issue 1: The Mail Transfer Agent (MTA) fails. Because older MTAs can not understand IPv6 addresses, they cannot initiate fallback when they encounter an IPv6 address (nor can IPv4 be used even if an IPv4 address is specified for the destination at the same time) and fail.

- Issue 2: If the receiving MTA does not set up reverse DNS lookup in the sending MTA as an anti-spam measure, incoming mail will be rejected. This is because the vast majority of mail coming from MTAs with no reverse DNS is spam, which typically is coming from a botnet. Mail on IPv4 exploits this fact by treating sending MTAs with no reverse DNS as spammers, so receiving MTAs are configured to not accept mail from such sources.

6.10.4. Security Considerations

This technique of rejecting mail from MTAs with no reverse DNS has not been widely adopted by service providers because some legitimate messages could be blocked (not all mail from MTAs with no reverse DNS is spam). But because this anti-spam measure is so easy to implement, it has been widely adopted by many companies. Indeed, we can safely assume that this is the default setting on many mail appliances. While some believe that reverse DNS is not needed in IPv6, on the contrary applying reverse DNS to MTAs is still a valuable anti-spam tool.

Considering that mail already cannot be sent from MTAs with no reverse DNS in IPv4, continuing the practice in IPv6 should not present a problem.

6.10.5. Identifying the Problem

Issues 1 and 2: After implementing IPv6, check to make sure mail is properly sent and received. Also, verify that delivery error status is working properly. If there are some hosts that can send mail and others that cannot, this may be due to MTA problems at the host or because there is a problem with the MTA reverse DNS settings on your own MTA. Generally, mail delivery errors are sent after the default delivery time times-out (the default is often several hours), so the user often becomes aware of the problem very late or not at all.

6.10.6. Solutions

Issue 1: On the receive side, verify that MX records exist for hosts, and avoid registering hosts with only AAAA records. Modify the software if necessary. On the send side, negotiate with the destination host, and configure the designated site to send IPv4 traffic.

Issue 2: Check to see if MTA reverse DNS is set up correctly, and if not take appropriate action.

6.11. IPv6-Ready Mail System Issues (Anti-Spam Techniques)

6.11.1. Greylisting Issues

The effectiveness of greylisting for IPv6-compliant mail systems is unclear.

6.11.2. Analysis

Greylisting is a method of protecting e-mail users against spam by temporarily rejecting incoming mail from IP addresses not listed in the mail server's database. Processing of current filtering programs is based on IPv4, so greylisting programs must be adapted to IPv6 addresses. However, IPv6 address space is so vast that there will likely be operational difficulties.

6.11.3. Causes

Processing of existing programs is premised on IPv4 addresses. Moreover, the address space of IPv6 is so vast that spammers can continue to send spam while constantly changing IP addresses, so it is not at all clear if IPv6 will offer the same protection as IPv4.

6.11.4. Identifying the Problem

Ascertain whether greylisting is currently in use from the IPv4 operating time (delay time).

6.11.5. Solutions

- Evaluate the effectiveness of the greylisting programs listed in references below to assess their support for IPv6.
- Consider identifying spam not based on IP address but based on the reputation of send domain authentication results and domain name.

6.11.6. References

"Greylisting for qmail with IPv6 support," http://gurubert.de/greylisting

"Milter-greylist home page," http://hcpnet.free.fr/milter-greylist/

6.12. Blacklist Database Service (DNSBL) Issues

6.12.1. Nature of the Problem

If your mail system is IPv6 compliant, you may not be able to use blacklists or other DNS-based blackhole lists (DNSBLs) to defend e-mail users against spam.

6.12.2. Analysis

DNSBLs are lists of IP addresses of computers or networks linked to spamming that are used to reject messages sent from sites on such lists. DNSBLs are based on IPv4, so blacklists must be adapted to IPv6 addresses.

6.12.3. Causes

- Existing DNSBL spam blocking assumes IPv4 addresses.
- IPv6 addresses are longer, have a different format, and other differences with IPv4, so programs must be adapted and databases expanded if DNSBLs are to be used in the same way to prevent spam on IPv6 networks.

6.12.4. Security Considerations

- Currently there are virtually no implementations of DNSBL that support IPv6.
- It is currently not clear which of two competing methodologies for implementing IPv6-enabled DNSBL will prevail: separate registration of expanded IPv6 address and grouping suspect addresses by prefix.
    - If IPv6 addresses are registered separately, spammers could continue to send spam by merely changing addresses, so the effectiveness of the database would be questionable.
    - If the prefix approach is used, addresses of non-spammers might also be blocked, thus risking excessive filtering.

6.12.5. Identifying the Problem

Check to see if DNSBL is currently being used on IPv4. If so, it is likely that DNSBL is the default configuration on the email appliance.

6.12.6. Solutions

- Wait until consensus is reached about the best method for implementing DSNBL on IPv6.
- Consider putting some IP addresses on the DNSBL whitelisted. Verify this approach to see if it works.
- Consider identifying spam not based on IP address but based on the reputation of send domain authentication results and domain name.

6.13. Localizing Problems on Access Lines: Troubleshooting When Multiple Providers are Involved in Providing Service

6.13.1. Nature of the Problem

When a user runs into communication problems and there are a number of different service providers involved—access line provider, Internet Service Provider, Virtual Network Enabler, and so on—the user often does not know which service provider to contact. Even when the user contacts the primary contractor (ISP or access line provider), sometimes the problem cannot be identified. Typically

- solving the problem takes considerable time,
- the root cause cannot be specified, and
- the user gets passed from one call center to the next.

6.13.2. Causes

Besides the providers that the user knows—the access provider and ISP—there are often other providers in the background that the user is not aware of (VNE, roaming). In situations where a problem or failure is caused by one of these elusive background providers, the call centers that users has access to may not be able to solve the problem.

6.13.3. Analysis

- In the case of a major failure, information will probably be shared.

- In the rare case where a user has some influence, the problem may not be solved.

6.13.4. Solutions

- Cooperation among providers involved.
- Activities to raise awareness plays a role.

6.13.5. References

http://www.soumu.go.jp/main_content/000009743.pdf

6.14. Presence of Unsupported L2 Multicast Equipment

6.14.1. Nature of the Problem

If you try to use L2 communication equipment that does not support multicast or is improperly implemented, the Neighbor Discovery Protocol (NDP) used for multicast functions (T. Narten, 2007) will not work, making IPv6 communication impossible. You can expect

- IPv6 communication between nodes on LANs to fail,
- MAC address resolution by NDP to fail, and
- certain nodes will not support IPv6 communication. In addition, IPv6 packets for some node addresses will not arrive.

6.14.2. Causes

- The problem is caused by L2 communications equipment that does not support multicast capability or is improperly implemented present on the LAN. Typical L2 equipment includes
  - L2 switches (hardware, firmware)
  - Ethernet cards (hardware, firmware, driver)
  - PCs, etc. (OS driver)

- There are two potential problems that could be involved: sending multicast packets over the LAN fails (problem on the L2 switch), or receiving multicast packets at the node fails (problem on the Ethernet card or the OS driver).
    - If some nodes on the LAN are not capable of receiving multicast, then an asymmetrical failure occurs: the nodes can send IPv6 packets to other nodes, but they cannot receive IPv6 packets sent from other nodes.
- This is not an IPv6 problem, but rather a multicast capability problem in L2.
    - IPv4 uses the broadcast-based Address Resolution Protocol (ARP) to resolve addresses, but NDP is the multicast-based equivalent to ARP in IPv6.
    - Since there are relatively few opportunities to use L2 multicast in the IPv4 environment, there is little chance that this problem would come to light.

6.14.3. Security Considerations

If you apply the workaround of promiscuous mode settings, this requires administrative (superuser) privilege, involves dealing with a type of packet not required before, and is likely open up other security issues. It could also cause node overload problems.

6.14.4. Identifying the Problem

- Inability to carry IPv6 traffic could be caused by some of the equipment or all of the equipment, so it is difficult to isolate the cause.
- Between two nodes, packets can only move in one direction.
- Packet capture in promiscuous mode captures all packets (not just the packets addressed to it), and is an effective way to diagnose network connectivity issues from a PC.

6.14.5. Solutions

- Update L2 communications equipment firmware and drivers with the latest versions.

- Replace L2 communications equipment with equipment that has multicast capability.
- Set the network interface to promiscuous mode and capture all packet on the LAN whether multicast is available or not.
- If your switch lacks multicast capability and you are able to specify the destination node, sometimes you can solve the problem by setting a specific multicast address and see if it gets through.

6.14.6. References

RFC 4861: "Neighbor Discovery for IP version 6 (IPv6)" (T. Narten, 2007)

6.15. Adverse Effects of IPv6 Multicast on Home Communications

6.15.1. Nature of the Problem

In environments receiving services delivered by IPv6 multicast, the multicast send traffic even to equipment where it is not needed. This creates excess load on the network that has an adverse effect on normal communications.

6.15.2. Causes

The problem stems from the specifications for multicast, which specifies that packets will be multicast to all nodes on the same segment.

6.15.3. Analysis

When multicast traffic volume is heavy, say when delivering video, the switches and nodes that do not need to receive this data can have a major impact on the communication. Especially when wireless access points are installed with a bridge, because wireless bandwidth is significantly narrower than wireline bandwidth, the congestion that occurs on wireless networks can quickly become problematic. Note that this problem also occurs when using multicast on IPv4 networks.

6.15.4. Identifying the Problem

- You may be able to verify the communications state of the switch by just checking to see if the communication ON/OFF lamp is lit.
- Check the status of packet capture, communication state, and so on.

6.15.5. Solutions

Multicast service subscribers can remedy this problem by dividing the equipment into two sets—equipment that receives the multicast and equipment that does not—by using equipment with MLD snooping capability, and if multicast is not really needed, using wireless equipment that lacks multicast forwarding capability.

6.15.6. References

RFC 4541: "Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches" (M. Christensen, 2006)

6.16. IPv6 Address Notation

6.16.1. Nature of the Problem

IPv6 supports different address notation schemes and this can lead to a number of problems: it complicates searching addresses, results in inconsistent output from logs, and leads to inappropriate advice from customer support.

6.16.2. Causes

Acceptance of abbreviated notation and certain flexibility in the notation of IPv6 addresses has resulted in IPv6 addresses being represented in different ways.

6.16.3. Analysis

- Around the time RFC 3513 was issued, specifications for address notation were changed (when just a single 16bit 0 field is present).
- No uniform address notation scheme was recommended.

- IPv4 addresses do not use abbreviations or letters, and while there is some flexibility in the treatment of single leading zeros, this does not present a problem because everyone understands leading zeros.

6.16.4. Security Considerations

When X.509 certificates are used for access control, if a simple text-based comparison is done when certificates are authenticated, it could produce a valid/invalid false positive, and thus represent a security risk.

6.16.5. Identifying the Problem

IPv6 address search fails, log outputs are not uniform.

6.16.6. Solutions

- Use products and systems that are RFC 5952 compliant (permanent solution).
- For products and systems that do not comply with RFC 5952, encourage suppliers to bring their products into compliance while avoiding the issue by bringing address notation into compliance with RFC 5952 (interim solution).
- Spread awareness and inform engineers and customer support personnel about RFC 5952 compliant address notation.
- Use RFC 5952 compliant abbreviations when conveying IPv6 addresses verbally such as over the telephone. Phonetic codes—A as in Alfa, B as in Bravo, C as in Charlie…—work well for conveying letters verbally.

6.16.7. Referencess

RFC 5952: "A Recommendation for IPv6 Address Text Representation" (S. Kawamura, 2010)

6.17. Implementations That Do Not Meet Minimum Specifications

6.17.1. Nature of the Problem

Consumer electronics or sensor devices that do not meet or share minimum required specifications are connected to an IPv6 network.

6.17.2. Analysis

There are terminal implementations that do not meet the minimum specifications to acquire IPv6 addresses, DNS, and other options in the IPv6 environment. The same problem could also occur in the IPv4 environment, but is less likely considering the technical maturity of IPv4.

6.17.3. Solutions

It is necessary to clearly define the minimum specifications for terminals in the IPv6 environment.

6.17.4. Identifying the Problem

If terminal specifications vary, equipment may not support connectivity to the network, communication may not work normally, and various other problems could occur.

6.17.5. References

"IPv6 Compatibility and IPv6 Functional Use Guideline for IPv6 Terminal OSes," http://www.v6pc.jp/pdf/v6TermOs_2006Guideline-0.pdf

"IPv6 Home Router Guideline," http://www.v6pc.jp/jp/upload/pdf/v6hgw_Guideline_2.0.pdf

"RFC 4294: IPv6 Node Requirements" (Loughney, 2006)

"Requirements For IPv6 in ICT Equipment," ripe-501, http://www.ripe.net/ripe/docs/ripe-501

6.18. IPv6 Privacy Address (RFC 4941) Issues

6.18.1. Nature of the Problem

The lower 64 bits of IPv6 addresses are automatically generated from the interface's MAC address using the modified EUI64 format, but these lower 64 bits of the IPv6 address do not change so long as the device doesn't change. This raises concern this unchanging information could be used to track the activities of an individual or the

usage of a particular machine. RFC 3041 proposed a scheme for generating anonymous addresses to protect privacy that was further revised in RFC 4941, but the scheme has never been adopted as intended. Specifically, the host becomes difficult to manage

- when the address is used as a listening address by the server, and
- when used by the network (e.g., an enterprise network), assuming the "IP address remains static and unchanged."

6.18.2. Causes

RFC 4941 proposes random shuffling of the lower 64 bits of IPv6 addresses, but it is not clear in which cases this would be recommended. Nor is it understood what is resolved and what is not resolved. Moreover, there are cases where this approach is unintentionally used.

6.18.3. Analysis

- Is RFC 4941 recommended for actual services?
  Assuming the "IPv6 address is a fixed address," two approaches are being studied—an IPv6 address authentication key application and a push-type service—but caution is called for since anonymous addresses would become unavailable (a way of dealing with DDNS and MobileIPv6 would thus be required). When the constraints of a push-type service become apparent, this could diminish the advantages of IPv6. Stakeholders must be kept well informed of these developments.

- What is resolved by RFC 4941?
  Since the lower 64 bits are randomly varied, the MAC address of the host is concealed from outside, thus mitigating privacy concerns associated with communications records and tracking. But if the upper 64 bits do not change, concerns over privacy associated with traceability are not alleviated.

- The same concern exists in IPv4 for broadband services when a fixed IPv4 address is given if the address is somehow tied to the individual's personal information.

6.18.4. Security Considerations

The primary assumption of RFC 4941 is that the loss of the MAC address can be prevented and that security issues will not arise when using conventional EU164-based address generation (RFC 4941, Section 7).

6.18.5. Solutions

Educate others and spread awareness as to when RFC 4941 can be used effectively.

- Caution: Different operating system behave differently (e.g., the default on Windows machines is ON).
- Where appropriate, hosts under one management should be configured to not use RFC 4941.

6.18.6. References

draft-iesg-serno-privacy: "Privacy Considerations for the Use of Hardware Serial Numbers in End-to-End Network Protocols" (Expired)

RFC4941: "Privacy Extensions for Address Configuration in IPv6" (T. Narten, 2007)

Ministry of Internal Affairs and Communications: "Guideline to Enable IPv6 for E-Government Systems," March 30, 2007, p. 22.

6.19. IPv6 Address Traceability (Privacy) Issues

6.19.1. Nature of the Problem

Use of IPv6 introduces a different level of traceability than IPv4, and measures must be implemented to deal with this difference. Consider the implications.

- Traceability could be abused by malicious services. Such services could target IP-linked access of multiple sites that do not have a login mechanism, and exploit other vulnerabilities.
- If the IP of a router is changed, it will be difficult to come up with some kind of workaround.

6.19.2. Causes

- It is anticipated that highly static operations will become more common in the way ISPs allocate IPv6 prefixes compared to the way IPv4 global addresses have been allocated.

- This is because the terminal interface IDs (last 64 bits) is automatically generated as a static operation from its 48-bit MAC address (see the previous section).

6.19.3. Analysis

- IP address traceability that we refer to here is the ability to surmise that messages are coming from the same user if multiple connections to a destination are made over the Internet from the same IP address. The destination does not have to be the same host.

- If you continue to use the same IP address over a prolonged period, this opens the way to traceability over a period of time.

- In earlier IPv4 operations, the global address was generally dynamically allocated to home routers from the ISP. As a service option, the user could select static operation. Even if a slightly different address is not actually allocated when static addressing is not guaranteed, a different address is obtained when power is restored to a home router, and there are various other possibilities. Thus, if you include selection of the ISP by the user, one is able to exercise some control over the traceability of one's home router.

- Consequently, when using IPv4, the prospect of ISPs enabling long-term traceability of home routers is not a big issue.

- But in IPv6, the traceability of a home router can be obtained as in IPv4 by looking at the prefix allocated to the router. It is apparent even in this preliminary stage before IPv6 is widely deployed, that prefixes are likely to used dynamically to the same extent as IPv4 global address. This is because dynamic use of IPv4 global addressing was motivated in part by economy, but there is little need for this with the IPv6 prefix. While modifying the IPv4 global

address only affected the WAN address of the home router, modifying the IPv6 prefix changes the IPv6 addresses of all the equipment linked to the home router. The means that the users' options for controlling the traceability of home routers is severely restricted (only option is to change ISPs).

- If practically all households are opened up to long-term traceability, this will motive the development of new services. And as these services become generally available, it will become more difficult to change prefixes since this would interferes with the services.
- IPv4 has the same problem, but the impact is far more pervasive in IPv6.

6.19.4. Identifying the Problem

- Verify your ISP's management policy.
- Actually monitor the prefix over a prolonged period.

6.20. CGN, Translation Issues

6.20.1. Nature of the Problem

As IPv4 addresses are depleted and IPv6 is deployed, we will see increasing effects from the use of Carrier Grade NAT (CGN) and IPv6/IPv4 translators. Problems of applications and services not working as users expect have already occurred.

6.20.2. Causes

- IPv4/IPv6 translators are implemented to provide the following functions.
  - To ensure reachability of IPv4-only equipment on the IPv6 Internet.
  - To ensure reachability of IPv6-only equipment on the IPv4 Internet.
- CGN is implemented to provide the same capabilities (IPv4 ⇒ IPv4 communication is problematic).

6.20.3. Analysis

Problems caused by CGN and translators depend on the performance and operations of equipment and on the circumstances of individual applications employed by users. If applications, services, and user environments support both IPv4 and IPv6, no problems occurs since IPv6 is generally preferred. In other words, these problems should be limited to transition period.

- Problems arising from limiting the number of concurrent sessions
  According to Reference 1, 99% of users will not be affected if the upper limit on concurrent sessions is set at 1000, but at most this is no more than 1/64th the pace that global IPv4 addresses are being consumed (the upper limit for 90% of users to be satisfied is 100 sessions).

- Inability of users to specify IP addresses on the service side could lead to a number of problems.

  - From outside the CGN and translation function between IPv6 and IPv4 it appears to be the same global IPv4 address.

  - Poor interaction with services run by IP-based access management.

  - Greater difficulty analyzing logs for sites restricting bulletin board access, prohibiting parallel downloads, etc.

- When CGN is deployed, some NAT traversal technologies will not work. UPnP and other applications will not work under double NAT.

- Problem with IP addresses included in some application protocols

  IP address information will not be correctly translated if the CGN/translation function Application Layer Gateway (ALG) is not supported.

- ISP level and residential service uses the same subnet address in NAT 4-4-4

  In cases where communication to a user in the same CGN is transmitted directly without going through the CGN, there are home routers that have

trouble selecting a route in the residential service and CGN if the subnet address is covered.

6.20.4. Identifying the Problem

It is easy to tell if applications are not working as expected, but users will generally not be able to pinpoint the translation function as the source of the problem.

6.20.5. Solutions

- Steps available to users:
  - Reduce the number of concurrent CGN/translation sessions (reduce the number of applications and devices concurrently in use).
  - Try changing the CGN home network subnet address.
  - Give up on CGN/ translation, and try different applications and services.
- Steps available to developers:
  - CGN and translation function: Reduce the number of concurrent sessions to under 100 (user settings, dynamic number control, etc.)
  - CGN and translation function: Analyze access using non source IP information.
  - CGN: Use NAT traversal technology taking double NAT into account.
  - CGN and translation function: Do not encapsulate the IP address for your own protocol.
  - Translation functions: IPv6 works normally, and take a practical attitude toward IPv4 even if user experience is poor.

6.20.6. References

"Assessing Impact of Users when ISPs Introduce NAT,"
http://www.ieice.org/jpn/books/kaishikiji/2010/201006.pdf

"Use of NAT Goes According to Plan,"
http://www.janog.gr.jp/meeting/janog24/program/d2p5.html

"Difficult for Consumers to Understand: Network Appliance IPv6 Issues,"
http://itpro.nikkeibp.co.jp/article/Watcher/20091015/338865/

6.20.7. Expressions Subject to Misunderstanding, Problems from Sharing Obsolete Information

6.20.8. Nature of the Problem

Problems occur as a result of passing on obsolete information or information that is simply wrong about IPv6. Consider the following examples.

- IPsec is always implemented in IPv6 environments.
  Implementation is required, but using IPsec is not. Saying that IPsec is always used is mistaken.
- Use of global addresses diminishes security and causes one to become overly cautious.
- Meaning of response to multicast.
- Ambiguity of the term "IPv6-ready" and the misapprehension "IPv6-ready browser" is fully IPv6 compatible when it really only supports IPv6 bridge function (IPv6 pass-through function).
- The notion that if you uninstall or disable IPv6, equipment run faster (disable recommendation).

Constructing a network based on this kind of information would almost certainly result in problems due to the non-standard network configuration.

6.20.9. Causes

This problem can be attributed to lack of awareness of new information, and insufficient collection and archiving of information. Essentially, it is a lack of knowledge on the part of users.

6.20.10. Security Considerations

Misunderstanding of IPsec and the overcautious attitude toward global addresses could certainly affect security.

6.20.11. Identifying the Problem

The lack of knowledge is harder to deal with. The only solution is to ask question and find out how much people know.

6.20.12. Solutions

Circulate information and knowledge. Come up with question-format materials that convey the actual situation to a wider audience. Another approach would be for companies to put up a trusted site where employees can ask questions [Sire?].

6.21. IPv6 Impact on Multiple IPv4 Subnets

6.21.1. Nature of the Problem

The following items must be considered when deploying IPv6 in networks integrating several subnets over IPv4.

- Prefix length allocated by ISPs.
  In IPv6, prefixes allocated range from /64, /48. Addresses are allocated to each subnet either manually or by DHCP-PD.

- You have the option choosing the same topology for IPv6 if the IPv4 network is divided into multiple subnets or configuring a separate topology when using IPv6 pass-through function, but using IPv6 pass-though must be thought through carefully if there is any reason for dividing the IPv4 network into multiple subsets.

6.21.2. Security Considerations

Note that if the IPv4 and IPv6 topologies are different, it is likely that the IPv4 network access policy the IPv6 access policy also do not correspond.

6.22. IPv6 Impact on Large-Scale L2 Networks

6.22.1. Nature of the Problem

Operational problems occur when deploying IPv6 in an IPv4 operating environment, or segments of an IPv4 network appear to be linked by IPv6.

6.22.2. Cause

If an L2 design is implemented incorrectly on an existing IPv4 network, VLANs will cause problems due to IPv4-dependent functions of L2 equipment. And if a user attempts to interconnect two L2 networks with a hub or switch, this will also cause trouble.

6.22.3. Analysis

- On IPv4, even poorly designed and implemented L2 networks will work satisfactorily since the interface only has to deal with one kind of data, but the complex network information provided by IPv6 creates problems if the L2 network is poorly designed.

- While integrating authentication with a dynamic VLAN is not an issue on IPv4, it creates problems on IPv6 since multicast messages such as router advertisements are broadcast to all ports on the link.

6.22.4. Security Considerations

Sometimes unintended segments are configured.

6.22.5. Identifying the Problem

Check to see if IPv6 router advertisements reach unintended segments.

6.22.6. Solutions

- Implement L2 network designs carefully and correctly (multicast-aware design).

- Deploy VLANS and other L2 equipment that is capable of handling multicast.

6.23. Problems that Cannot be Resolved Within CPEs Own Domain

6.23.1. Nature of the Problem

When configuring with CPE, there are cases where one's own domain is used as an access destination URL (e.g., http://.setup). The problem is that domain cannot be resolved, and in some cases, cannot be accessed by the configuration screen.

6.23.2. Causes

- There are cases where broadband router unique domains such as ".setup" cannot be resolved.

  - If the IPv6 DNS server address is configured using a host under a router with IPv6 pass-through function, problems occur when a DNS query is sent and IPv6 has priority.

  - When CPE expects a query to come to the DNS via IPv4 and uses a unique domain, the CPE passes the query through if it comes via IPv6. A global DNS query returns a name error (NXDOMAIN), the name is not resolved, and knowledge that the name does not exist is stored in the host (negative caching).

6.23.3. Analysis

The problem occurs when a unique domain is used and a mixed IPv4/IPv6 environment is not expected.

6.23.4. Identifying the Problem

Check the DNS configuration of the host.

6.23.5. Solutions

- Input the CPE IP address directly.
- Try customizing the policy table to assign priority to IPv4.

6.24. IRR Registration Issues

6.24.1. Nature of the Problem

Routing advertisements are blocked because routing information is not being registered, so communication is stymied.

6.24.2. Causes

- Upstream: A database called the Internet Routing Registry (IRR) is required to share routing information—routing polices and routing advertisements—among ISP.
- Another common condition, in tandem with the IRR, is the WHOIS system, a registry database for managing IP address assignment administrators.
- If information is not registered in the IRR or if the registered data contains errors, the routing information probably will not work as intended.
- In these circumstance, communication between terminals fails.

6.24.3. Analysis

- The problem is caused by improper configuration, someone neglected to register the information, or some other disaster caused by human error.
- On rare occasions, traffic may be filtered or a peer-to-peer link disconnected because of payment problems or some other financial reason.
- These same problems occur on IPv4.

6.24.4. Identifying the Problem

Verify communication is not reaching its intended target using ping or traceroute, then check the following.

- Verify registration in a public IRR.
  - http://www.irr.net/docs/list.html}
- For commercial IRRs, the query must be sent to the ISP managing the IRR.
- Check to see if you are registered in the registry database.
- Verify with a BGP Looking Glass server.
  - http://www.bgp4.as/looking-glasses
  - http://neptune.dti.ad.jp/
  - http://lg.he.net/
  - http://www.ipv6tf.org/index.php?page=using/connectivity/looking_glass
  - http://www.switch.ch/network/tools/ipv6lookingglass/
  - http://bgp4.jp/
  - http://lg01.colo01.bbtower.ad.jp/
- Use mail or a dedicated tool to send routing information over an Internet exchange (IX). It may be necessary to turn off filters first.

6.24.5. Solutions

- Verify registration when IPv6 is deployed.
- Verify connectivity to a number of different widely dispersed sites when IPv6 is deployed.

6.24.6. References

RFC 1786/RIPE-181: "Representation of IP Routing Policies in a Routing Registry," (T. Bates, 1995)

RFC 2650: "Using RPSL in Practice" (D. Meyer, 1999)

RFC 2726: "PGP Authentication for RIPE Database Updates" (J. Zsako, 1999)

RFC 2769: "Routing Policy System Replication" (C. Villamizar, 2000)

RFC 4012: "Routing Policy Specification Language Next Generation (RPSLng)," (L. Blunk, 2005)

RFC 5943: "A Dedicated Routing Policy Specification Language Interface Identifier for Operational Testing" (Haberman, 2010)

"IPv6 BGP filter recommendations, RIPE31(Sept. 23, 1998),"
http://www.space.net/~gert/RIPE/ipv6-filters.html

"JANOG Comment 1006," Aug. 26, 2008,
http://www.janog.gr.jp/doc/janog-comment/jc1006.txt

"Reference for IPv6 Router settings," http://www.cymru.com/Bogons/v6top.html

JPIRR, http://www.nic.ad.jp/ja/ip/irr/

6.25. Number of DNS Records and OS Operation

6.25.1. Nature of the Problem

In Microsoft Internet Explorer, users can control the number of fallback attempts (refer to "IPv6/IPv4 Fallback") with a setting in the registry. The default is 5 attempts. If it takes more than 5 attempts to access an AAAA record, there will be problems establishing an IPv6 connection since fallback to IPv4 will fail.

6.25.2. Causes

The default number of fallback attempts is set too low.

6.25.3. Analysis

Currently there are many A records registered, so there is a good probability that there are just as many AAAA records registered.

6.25.4. Identifying the Problesm

If communication is held up in an IPv6/IPv4 dual-stack environment, check for congestion on the route and also check the number of A and AAAA records registered.

6.25.5. Solutions

- Reduce the number of registered DNS records (see Google comments on the lead-up to World IPv6 day).
- Adjust the default value in the registry to increase the number of fallback attempts.

6.25.6. References

http://support.microsoft.com/kb/2293762/ja

6.26. Problems Regarding How Sites are Viewed

6.26.1. Nature of the Problem

- Problems have been reported of IPv4 and IPv6 sites looking different, or different data being obtained when accessed via IPv4 and accessed via IPv6 (excluding cases where the appearance has been deliberately changed). There have also been reports of older content being accessed by IPv6, and data results obtained via IPv4 and via IPv6 being different.
- There is a problem of no IPv6 connectivity to the content referent through IPv6/IPv4 dual-stack access.

6.26.2. Causes

- When IPv6 and IPv4 content servers are separate, data synchronization may not work properly.
- Lack of IPv6 compliance of Content Management Systems (CMSs).
  - Checking for broken links cannot be implemented on IPv6.

6.26.3. Security Considerations

Security concerns over different management levels when there are multiple servers.

6.26.4. Identifying the Problem

The problem is revealed by differences in way sites look and in the data obtained.

6.26.5. Solutions

Depending on the configuration, maintain firm control over data management. Notify the site administer if problems are encountered.

Appendix A: Abbreviation and Acronyms

IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
RA	Router Advertisement
AS	Autonomous System
ICANN	Internet Corporation for Assigned Names and Numbers
IANA	Internet Assigned Numbers Authority
RIR	Regional Internet Registry
ARIN	American Registry for Internet Numbers
RIPE NCC	Réseaux IP Européens Network Coordination Centre
APNIC	Asia Pacific Network Information Centre
LACNIC	Latin American and Caribbean Internet Addresses Registry
AfriNIC	The Registry of Internet Number Resources for Africa
CNNIC	China Internet Network Information Center
JPNIC	Japan Network Information Center
ISP	Internet Service Provider
NIR	National Internet Registry
LIR	Local Internet Registry
NAT	Network Address Translation
IETF	Internet Engineering Task Force
CGN	Carrier Grade NAT
DS-Lite	Dual-Stack Lite
OECD	Organization for Economic Co-operation and Development
MTU	Maximum Transmission Unit
PMTUD	Path MTU Discovery
DNS	Domain Name System
DDNS	Dynamic Domain Name System
ICMP	Internet Control Message Protocol

ICMPv6	Internet Control Message Protocol version 6
L2	Layer 2
CPE	Customer Premises Equipment
IRR	Internet Routing Registry
OS	Operating System
PPPoE	Point to Point Protocol over Ethernet
DHCP	Dynamic Host Configuration Protocol
DHCPv6	Dynamic Host Configuration Protocol for IPv6
DHCPv6-PD	DHCPv6 Prefix Delegation
TCP	Transmission Control Protocol
HTTP	Hyper Text Transmission Protocol
HTTPS	HTTP over SSL/TLS
VPN	Virtual Private Network
ULA	Unique Local IPv6 unicast Address
SLAAC	StateLess Address AutoConfiguration
VLAN	Virtual Private Network
LAN	Local Area Network
SEND	Secure neighbor Discovery
MSS	Maximum Segment Size
uRPF	unicast Reverse Path Forwarding
SMTP	Simple Mail Transfer Protocol
MUA	Mail User Agent
IMAP	Internet Message Access Protocol
POP	Post Office Protocol
MTA	Mail Transfer Agent
VNE	Virtual Network Enabler
URL	Uniform Resource Locator
IX	Internet Exchange

Appendix B: Terminology

For IPv6 related terminology, refer to "IPv6 Terminology" compiled by the Internet Association of Japan (http://www.iajapan.org/ipv6/v6termwg.html#glossary_02).

Update History

Date	Version	Difference
26 July 2011	1.0	Initial version
12th September 2011	1.1	Published on the web page
13th November 2011	1.2	Correct some typo and add copy right

References

A. Conta, S. Deering, M. Gupta. 2006. Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification. 3 2006. RFC4443.

A. Durand, J. Ihren. 2004. DNS IPv6 Transport Operational Guidelines. 9 2004. RFC3901.

A. Durand, J. Ihren, P. Savola. 2006. Operational Considerations and Issues with IPv6 DNS. 4 2006. RFC4472.

A. Matsumoto, T. Fujisaki, R. Hiromi, K. Kanayama. 2008. Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules. 7 2008. RFC5220.

C. Villamizar, C. Alaettinoglu, R. Govindan, D. Meyer. 2000. Routing Policy System Replication. 2 2000. RFC2769.

B. Carpenter, K. Moore. 2001. Connection of IPv6 Domains via IPv4 Clouds. 2 2001. RFC3056.

D. Meyer, J. Schmitz, C. Orange, M.Prior, C. Alaettinoglu. 1999. Using RPSL in Practice. 8 1999. RFC2650.

D. Wing, A. Yourtchenko. *Happy Eyeballs: Trending Towards Success with Dual-Stack Hosts.* draft-ietf-v6ops-happy-eyeballs. Work in progress. Internet Draft.

S. Deering. 1991. ICMP Router Discovery Messages. 9 1991. RFC1256.

R. Draves. 2003. Default Address Selection for Internet Protocol version 6 (IPv6). 2 2003. RFC3484.

A. Durand. Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion. Work in progress.

E. Davies, J. Mohacsi. 2007. Recommendations for Filtering ICMPv6 Messages in Firewalls. 5 2007. RFC4890.

E. Davies, S. Krishnan, P. Savola. 2007. IPv6 Transition/Co-existence Security Considerations. 9 2007. RFC4942.

E. Levy-Abegnoli, et al. 2011. IPv6 Router Advertisement Guard. 2011. RFC6105.

R. Fielding, J. Gettys, J. Mogul, H. Frystyk, T. Berners-Lee. 1997. Hypertext Transfer Protocol -- HTTP/1.1. 1 1997. RFC2068.

F. Baker, X. Li, at el. *Framework for IPv4/IPv6 Translation.* Work in progress. draft-baker-behave-v4v6-framework.

B. Haberman. 2010. A Dedicated Routing Policy Specification Language Interface Identifier for Operational Testing. 8 2010. RFC5943.

C. Huitema. 2001. An Anycast Prefix for 6to4 Relay Routers. 6 2001. RFC3068.

—. 2006. Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs). 2 2006. RFC4380.

IPv4 address Report. *IPv4 address Report.* [Online] http://www.potaroo.net/tools/ipv4/.

J. Arkko, Ed., J. Kempf, B. Zill, P. Nikander. 2005. SEcure Neighbor Discovery (SEND). 3 2005. RFC3971.

J. McCann, S. Deering, J. Mogul. 1996. Path MTU Discovery for IP version 6. 8 1996. RFC1981.

J.C. Mogul, S.E. Deering. 1990. Path MTU discovery. 11 1990. RFC1191.

L. Blunk, J. Damas, F. Parent, A. Robachevsky. 2005. Routing Policy Specification Language next generation (RPSLng). 3 2005. RFC4012.

L. Mamakos, K. Lidl, J. Evarts, D. Carrel, D. Simone, R. 1999. A Method for Transmitting PPP Over Ethernet (PPPoE). 2 1999. RFC2516.

K. Lahey. 2000. TCP Problems with Path MTU Discovery. 9 2000. RFC2923.

Loughney, J. 2006. IPv6 Node Requirements. 4 2006. RFC4294.

M. Christensen, K. Kimball, F. Solensky. 2006. Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches. 5 2006. RFC4541.

M. Mathis, J. Heffner. 2007. Packetization Layer Path MTU Discovery. 3 2007. RFC4821.

M. Townsley, O. Troan. 2010. IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification. 8 2010. RFC5969.

OECD resources on Internet addressing: IPv4 and IPv6. [Online] http://www.oecd.org/document/14/0,3343,en_2649_34223_44954318_1_1_1_1,00.html.

P. Savola, C. Patel. 2004. Security Considerations for 6to4. 12 2004. RFC3964.

D. Plummer. 1982. Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on thernet Hardware. 12 1982. RFC826.

R. Droms, Ed., J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney. 2003. Dynamic Host Configuration Protocol for IPv6 (DHCPv6). 7 2003. RFC3315.

R. Hinden, B. Haberman. 2005. Unique Local IPv6 Unicast Addresses. 10 2005. RFC4193.

S. Thomson, T. Narten, T. Jinmei. 2007. IPv6 Stateless Address Autoconfiguration. 9 2007. RFC4862.

S. Deering, R. Hinden. 1998. Internet Protocol, Version 6 (IPv6) Specification. 12 1998. RFC2460.

S. Kawamura, M. Kawashima. 2010. A Recommendation for IPv6 Address Text Representation. 8 2010. RFC5952.

S. Thomson, C. Huitema, V. Ksinant, M. Souissi. 2003. DNS Extensions to Support IP Version 6. 10 2003. RFC3596.

P. Savola, F. Baker. 2004. Ingress Filtering for Multihomed Networks. 3 2004. RFC3704.

T. Bates, E. Gerich, L. Joncheray, J-M. Jouanigot, D. Karrenberg, M. Terpstra, J. Yu. 1995. Representation of IP Routing Policies in a Routing Registry. 3 1995. RFC1786.

T. Narten, E. Nordmark, W. Simpson, H. Soliman. 2007. Neighbor Discovery for IP version 6 (IPv6). 9 2007. RFC4861.

T. Narten, R. Draves, S. Krishnan. 2007. Privacy Extensions for Stateless Address Autoconfiguration in IPv6. 7 2007. RFC4941.

Technical Infrastructure for USGv6 Adoption. [Online] http://www.antd.nist.gov/usgv6/.

Thaler, D. 2011. Teredo Extensions. 1 2011. RFC6081.

T. Chown, S. Venaas. 2011. Rogue IPv6 Router Advertisement Problem Statement. 2011. RFC6104.

Zsako, J. 1999. PGP Authentication for RIPE Database Updates. 12 1999. RFC2726.