Transition Scenarios and Solutions

draft-ietf-ngtrans-trans-scenes-00.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet- Drafts as reference material or to cite them other than as "work in progress."

To view the entire list of current Internet-Drafts, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Northern Europe), ftp.nis.garr.it (Southern Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this memo is unlimited.

Abstract

The document details scenarios and proposed solutions using the current transition mechanisms.
The first section categorises these mechanisms into 3 groups.
The second section describes appropriate scenarios and demonstrates how the mechanisms can be used and combined to form solutions.

Transition Mechanisms

Currently there are many different transition mechanisms with which you can use to implement IPv6.

The current mechanisms are:

Dual Stack
AIIH
Automatic Tunneling
Configured Tunneling
6over4
DTI
6to4
NAT-PT
SIIT
SOCKSv5
ALG
Bump in the Stack

These mechanisms can be categorised into 3 main areas:

- Those mechanisms that allow nodes to use either IPv4 or IPv6

  -- Dual Stack
  -- Assignment of IPv4 Global Addresses to IPv6 Hosts (AIIH)

- Those that are Tunneling and Encapsulation Mechanisms

  -- Automatic Tunneling
  -- Configured Tunneling
  -- Dynamic Tunneling Interface (DTI)
  -- 6over4 without explicit tunnels
  -- 6to4

- Those that are Translators

  -- Stateless IP/ICMP Translator (SIIT)
  -- Network Address Translation _ Protocol Translation (NAT-PT)
  -- SOCKSv5 Translator
  -- Application Layer Gateway (ALG)
  -- Bump in the Stack

Transition Scenarios

Scenario 1 Isolated IPv6 host in an IPv4 domain needing to communicate with an IPv6 network not directly connected.

Scenario 2 A large organisation with direct external connections.

Scenario 3 A small/medium organisation using a NAT.

Scenario 4 IPv4 dependent applications.

Justifying the choice of scenarios

SCENARIO 1
A typical situation that may occur in the early stages of migration where a few IPv6 nodes (1 in this scenario) in an IPv4 network require connection to an IPv6 network.

SCENARIO 2
Chosen to show the many possible ideas and migration techniques that could be used in a large organisation.

SCENARIO 3
This will offer ideas on how to migrate a small number of users and could be used in larger organisations, which have similar needs.

SCENARIO 4
This may occur frequently in the later stages or during migration when it is found that some IPv4 applications cannot be changed to support IPv6.

IMPORTANT ISSUES CONCERNING MY SCENARIOS AND SOLUTIONS
The document doesn't discuss how the implementation will be installed on each system, this will be vendor specific and would complicate matters if it listed all vendor-specific solutions.

Most scenarios and diagrams do not show any redundant equipment, this was done to simplify and to make solutions easier to understand and not to complicate the actual issues. Redundant equipment and networks will also be required to implement IPv6 in the ways discussed in each of the solutions and at the same time as updating the _live_ network.

Assumptions are specific, otherwise each scenario would become less well defined.

Scenario 1 - Isolated IPv6 host in an IPv4 Domain wishing to communicate with an IPv6 network not directly connected.

Introduction

A Corporate Customer requires connection to a new bank system. An IPv6 host in the customer's site is needed to connect to the bank's IPv6 only site. The bank decided to implement IPv6 only, to benefit from its enhanced functionality e.g. standardised security. The bank system is a specialised system and therefore only requires communication with customer sites. The bank's border router (R2) is a dual router but nodes within the Bank's network are IPv6 only.

`A diagram showing the current situation goes here'

Migration Requirements

- IPv6 host needs to communicate with all other nodes within the customer network.
- Continually maintain IPv6 functionality, therefore no use of translators should be permitted, need to maintain security and authentication procedures.
- No changes should be made to the Bank's network.

Suitability of the Transition Categories for this Scenario

IPv4 AND IPv6 MECHANISMS
One of these mechanisms will be required to allow the node within the customer site to communicate with IPv6 and IPv4 nodes.

TUNNELING AND ENCAPSULATION MECHANISMS
Tunneling will be required to allow IPv6 packets to traverse the IPv4 network.

TRANSLATORS
Translators have been ruled out due to breaking end to end connectivity.

Suitability of these Transition Mechanisms for this Scenario

DUAL STACK
Dual stack will need to be deployed in the node installed in the customer's premises. To allow for some sort of routing through the IPv4 network, the border router (R1) may also need to be installed with both IPv4 and IPv6.

DUAL STACK WITH CONFIGURED TUNNELS
The IPv4/IPv6 host could be configured to tunnel all IPv6 packets to the default IPv4/IPv6 router. The packets would be encapsulated within IPv4 so that they could be routed to the router through IPv4 infrastructure. The problem is how does the host configure its IPv6 address can neighbour solicitations be transferred to the router through configured tunneling. If you are configuring a tunnel to the router R1 you might as well just tunnel all the way to the Bank's Router R2 and leave R1 as a normal IPv4 router.

DUAL STACK WITH AIIH
No need to use AIIH mechanism as only one host is connecting.

DUAL STACK WITH DTI
Not relevant as packets will mostly be traversing IPv4 networks.

DUAL STACK WITH 6OVER4
6over4 could be used if the network supports multicast routing. As 6over4 only works within a domain the IPv4 router (R1) would need to support IPv6 and also to have a 6over4 interface configured. The host could then communicate to router R1 using IPv4 multicast packets. The rest of the communication path would have to be carried out by another mechanism such as configured tunneling or 6to4.

6TO4
This could be implemented at the routers R1 and R2 using their unique IPv4 addresses as an NLA ID to create a unique IPv6 address.

Solution 1

For this solution the IPv4 Network does Support Multicasting.

Mechanisms Suggested in Solution

- Dual Stack
- 6over4
- 6to4 or configured tunneling.

ROUTER
The 6over4 mechanism will require the IPv4 network border router (R1) to be installed with IPv6 and a 6over4 interface. Note this router and the host does not need to be on the same segment, if in fact they were then there would be no requirement for 6over4. The router and the host are expected to have some IPv4 infrastructure between them. A Configured tunnel will need to be set up between R1 and R2 so the IPv6 packets can traverse the IPv4 Internet. The routers R1 and R2 could use the 6to4 method but this would mean that the router R2 would also have to implement 6to4 which in this case is not permitted as one of the requirements is that `No changes should be made to the Bank's network'.

HOST
The IPv4 host needs to firstly implement dual stack, keeping its original IPv4 address and then implementing 6over4 to allow the host to use encapsulation of IPv6 packets within IPv4 multicast packets. As this can only be used within an organisation, uses IPv4 Organisation-Local Scope (239.192.0.0) the router (R1) needs to be configured to support IPv6 routing. This host will find out its prefix by sending a router solicitation encapsulated within an IPv4 multicast packet to router R1, the router will then return with a router advertisement using the same method of encapsulation within an IPv4 multicast packet.

`A diagram showing the solution goes here'

Solution 2

For this solution the IPv4 Network does not support Multicasting.

Mechanisms Suggested in Solution

- Dual Stack
- Configured Tunneling

HOST
The host firstly needs to be implemented with the dual Stack mechanism. As the host is not going to be able to automatically be allocated a globally unique IPv6 address this will need to be input manually using the prefix of the router (R2). _Not sure if this can be done or is acceptable it could also find out its address by sending a router advertisement encapsulated within an IPv4 packet to the router R2_. Secondly the host has to be manually configured with a tunnel from host to the Bank's Router (R2). Once this is complete the Host will be able to communicate with the end node retaining the original functionality of the IPv6 packet.

SECURITY IMPLEMENTATION FOR BOTH SOLUTIONS
Both solutions will require some sort of Security implementation whether the use of MD5 as an authentication algorithm or DES-CBC as an encryption algorithm depends entirely on what the bank system uses.

Implementers should be aware that, in addition to possible attacks against IPv6, security attacks against IPv4 must also be considered. Use of IP security at both IPv4 and IPv6 levels should nevertheless be avoided, for efficiency reasons. For example, if IPv6 is running encrypted, encryption of IPv4 would be redundant except if traffic analysis is felt to be a threat. If IPv6 is running authenticated, the authentication of IPv4 will add little. Conversely, IPv4 security will not protect IPv6 traffic once it leaves the IPv6-over-IPv4 domain. Therefore, implementing IPv6 security is required even if IPv4 security is available [6over4].

Although the above was written for 6over4, it is also particularly relevant to all tunneling mechanisms used.

Scenario 2 - Large Organisation with direct external connections

Introduction
A Large network with direct Internet connectivity (without the use of a NAT) using globally unique IP addresses. In this scenario Sites 2 and 3 (small sites of about 100 users) have already been upgraded to support both IPv4 and IPv6 and it is a requirement that the head office (Site 1) be migrated so that IPv6 communication can be used between sites. Site 1 also wishes to expand its IP network but cannot be allocated anymore IPv4 addresses for its needs. The site will be converting in the earlier stages when most external sites are predominantly still IPv4. Some applications can only operate using IPv4, these are contained in one subnet (Domain 5). For IPv4 dependent applications distributed over many subnets look at Scenario 4 (IPv4 dependent applications).

`A diagram showing the distribution of Sites and the current IP addressing goes here'

Migration Requirements

- Need to communicate with IPv6 only hosts (maybe on the Internet)
- Need to communicate with IPv4 only hosts (a small number in the organisation and the Internet)
- Minimise IPv4 traffic

HEAD OFFICE (SITE 1)
The site uses the internal routing protocol OSPF and BGP at the border router. The network doesn't support multicasting. As shown in the diagram below, the hosts in Domain 5 will be running IPv4 dependent applications. Domain 6 is the new domain required.

`Diagram of the Head Office (Site 1) goes here'

DOMAINS
Domains 1 to 6 all contain the same number of devices. Each has a router connecting the subnets in each domain with each other and the backbone. DHCP, DNS, Mail and File Servers are contained within each domain.

Suitability of the Transition Categories for this Scenario

IPV4 OR IPV6 MECHANISMS
These mechanisms will be required to allow nodes to communicate with both IPv4 only nodes and IPv6 only nodes.

TUNNELING AND ENCAPSULATION MECHANISMS
The Internet is predominantly IPv4 so communication from one site to another will require the use of tunneling and encapsulation.

TRANSLATORS
This will be difficult to implement as the situation where IPv4 hosts will need to communicate with IPv6 nodes through a translator, almost impossible to set up. A translator within nodes such as BiS would be most suitable for this situation.

Suitability of these Transition Mechanisms for this Scenario

DUAL STACK
The dual stack mechanism will be required in each device to allow all nodes to communicate freely with each other. There are problems with scalability using only this method as there will not be enough unique IPv4 addresses for each IPv4 interface in each node.

Manageability of two IP addresses could be come very complex in such a large organisation.

DUAL STACK WITH CONFIGURED TUNNELS
Configured tunnels may need to be used between border network routers at each site for communication over the IPv4 network. In another case where it is not possible to convert the border router (R1) to a dual router then configured tunnels could be set up between the default domain routers (R4 to R9) or between the backbone routers (R2 and R3) to the border router at other sites. Could end up with configured tunnels set up between not only the other 2 sites but also other IPv6 sites on the Internet. This will lead to complex maintenance and administration of all these tunnels.

DUAL STACK WITH AUTOMATIC TUNNELS
Automatic tunneling could be implemented at each node for end to end tunneling. Each of the endpoints will need to support IPv4 compatible addressing. We also want tunneling encapsulation to occur for the shortest distance possible. How will `flat' IPv4 compatible addresses be routed through an IPv6 router to the correct segment.

DUAL STACK WITH 6OVER4
Multicasting is not currently implemented in this network. If it was supported this method would still be inappropriate as there will be vast numbers of IPv6 implementations and not just a few.

6to4
The 6to4 mechanism could be used between border routers e.g. R1 and border routers at other sites as long as they also implement 6to4 as an alternative to configured tunneling.

DUAL STACK WITH AIIH
AIIH is a solution that will be required in this scenario as it allows dynamic allocation of IPv4 address when communication with another IPv4 node is needed. DNS servers will be configured to only send an A record in response to a query when the end node is IPv4 only. If the end node is using dual stack it will respond with an AAAA record. This is to save the scarce IPv4 addresses.

DUAL STACK WITH DTI
Requires that all routers would need to support IPv6. This could be used in the later stages of this particular transition where IPv4 only nodes in Domain 5 need to communicate with other IPv4 domains/sites. During the early stages DTI will not be much use as the network will still be predominantly IPv4.

DUAL STACK AND A TRANSLATOR
Translators would be a simple solution but they will simply not be scaleable in this scenario and could not operate at the network boundary due to the excessive traffic flow.

TRANSLATOR
Same as above.

Solution

I have included all my thoughts into this one solution for this particularly large scenario. Included are different options depending on the requirements.

Mechanisms Suggested in Solution

- Dual Stack in each domain
- AIIH
- Configured Tunneling at domain routers
- Bump in the stack for IPv4 hosts in domain 5
- 6to4 at border routers

Stage 1 - Backbone Routers

The ideal solution would be to install IPv6 firstly onto R1 the IPv4 border router and then onto routers R2 and R3 (OPTION 1). This may not be possible in all situations and R1 may need to be upgraded at a later more convenient stage (OPTION 2).

OPTION 1
The router R1 will need to be upgraded to a dual router and configured to support IPv6 OSPF. Configured tunnels will have to be set-up to each of the external sites. This will allow internal IPv6 packets destined for one of the 2 sites to be encapsulated in IPv4 and sent over the v4 network. Installation and configuration of BGP4 will also be required at this router to become an IPv6 border router, tunneling will need to be configured between neighbouring IPv6 border routers. Once this router has been upgraded the 2 backbone routers R2 and R3 need to be upgraded and configured to support IPv6 OSPF.

OPTION 2
If R1 cannot be upgraded immediately the next solution would be to upgrade the routers R2 and R3 connecting the 2 FDDI backbone rings and assigning configured tunnels to each of these to the 2 external sites. Each could be individually taken off-line and IPv6 could be installed and IPv6 OSPF could be configured while traffic is routed through the other. These routers would also need to be configured as IPv6 border routers using BGP4 and configured tunnels would need to exist between these and neighbouring IPv6 border routers to exchange routing information. At a later more convenient stage R1 could be upgraded to support IPv6 and could take over as the IPv6 border router.

Stage 2 - Domain 1

The order of implementation should be as follows: -

ROUTER
Domain router R4 needs to be upgraded to a dual router and configured to support IPv6 OSPF.

DNS SERVER
Upgraded to dual stack and configured to allow AAAA records. DNS should be configured to return IPv6 addresses, if the node is dual stack, to reserve IPv4 addresses for nodes that actually need them and also to make use of the functionality offered by IPv6 wherever possible.

DHCP SERVER
Upgraded to allow for allocation of stateful IPv6 addresses if required and also distribution of IPv4 addresses to all dual stack nodes within the domain.

MAIL/FILE SERVERS
Upgrade to dual stack. Each server should be given a permanent an IPv4 addresses (keep there original IPv4 addresses). This is to make sure that servers can always communicate.

Stage 3 - Domain 2, 3 and 4

All these domains have the same characteristics as Domain 1 so the ideas and implementation order can be simply used for these domains.

Stage 4 - Domain 5

As all nodes in this domain need to support an application that cannot be converted to use IPv6 we are stuck with using IPv4 `forever'.  `Forever' meaning the lifetime of the application. Presumably the most suitable solution would be to use a translator at the border of the domain so that all IPv4 only nodes would be able to communicate with IPv6 nodes outside. This is not the case, this solution would only allow IPv6 nodes to communicate with the IPv4 nodes through the translator. To allow IPv4 nodes to communicate with IPv6 through a translator would be virtually impossible (apart from using a translator that is incorporated into the node is i.e. Bump in the Stack mechanism).
There are 2 main options available and these are:

OPTION 1.
The Router and Servers should be upgraded as shown in Domain 1 to 4. Dual stack should be implemented on all hosts. This would allow IPv6 application running on these nodes to use the IPv6 stack but the IPv4 dependent application will still only be able to use IPv4. This means that all nodes on the network will need to keep using IPv4 `forever' to communicate with these nodes. This may be the case

anyway, as we don't know how long if ever nodes on the Internet will take to be converted to IPv6.

OPTION 2.
Each IPv4 node requires the installation of the dual stack and to implement the Bump in the stack mechanism. This would allow these IPv4 dependent applications to communicate with IPv6 only nodes without any necessary changes to the application code.

Stage 5 - Domain 6

Obviously we need to obtain some IPv4 addresses for this domain, as we cannot just implement IPv6 only. It is not possible to just take a few IPv4 addresses from the DHCP scope of each domain. So instead there are 2 possible options: -

OPTION 1.
Once all domains, apart from domain 6, have been upgraded for IPv6 support then a traffic analysis will need to be carried out on each of these domains. The amount of IPv4 traffic carried on each domain and how many nodes using IPv4 will need to be measured. The domain with the lowest number of nodes communicating using IPv4 will have some of it's IPv4 address space taken away from it and given to Domain 6 and this proportion will need to be calculated depending on the no of nodes in each of the domains. Changing the subnet mask for this scope of addresses. Both of these domains will have to implement AIIH components on their DHCP and DNS servers, as there will obviously not be enough IPv4 addresses for each dual stack node. Routers will need to be configured for these subnet changes.

OPTION 2.
A slightly more drastic option would be to reallocate all IPv4 addresses within the network and reconfigure allocation for each domain and its subnet mask. The IPv4 addresses allocated to the whole site could then be spread evenly across all 6 domains. AIIH components would need to be installed in each of the domain's DNS and DHCP servers, as there would not be enough IPv4 addresses for each dual stack node. All servers should be given a permanent IPv4 address, as it is essential these always have IPv4 connectivity.

Scenario 3 - Small/Medium Organisation using a NAT

Introduction

There are 9 offices, each office is linked using a point-to-point connection as shown in the diagram. Each site contains a DHCP, DNS and Mail server and routers are used between offices (as opposed to half bridges) to minimise traffic. Each router uses the RIP routing protocol.

The network currently uses a private address space of 192.168/16 prefix with each site using a /24 prefix as shown below:

`Diagram No of users and addressing in each office goes here'

HEAD OFFICE

The Head Offices in London has a DNS server and a NAT at the border to convert the non-globally unique IP addresses to globally unique IP addresses, this is used for security purposes as well as allowing its own internal addressing structure. All external traffic and traffic that is destined for external sources is sent through the NAT. This external traffic is minimal with a large percentage of this being SMTP traffic.

Additionally the Head Office has a firewall which is configured to route all incoming SMTP traffic from the ISP's servers IP address to the internal mail router which is a Linux machine running SENDMAIL. The internal mail router then looks at the domain name in the message header and directly sends it to the relevant Mail server in each of the offices. On this same machine runs the Proxy Daemon Squid and NAT. An Intranet Server runs on a separate Linux machine using Apache.

The NAT in the Billingsgate Office (Head Office) is used for external communication and is assigned one IPv4 address (194.14.1.1). All external communication will pass through this device.

`Diagram showing the layout and communication links between the offices goes here'

Assumptions

Assume all testing and coding has been carried out on applications to allow them to run on IPv6 only hosts. If any applications cannot be run on IPv6 hosts this is detailed in Scenario 4 (IPv4 dependent applications) and therefore is not covered in this scenario.

Migration Requirements

- To maintain a translator at the network border for ease of maintenance.
- To allow for communication with IPv4 only hosts at all times during the transition.
- Eventually eliminate all IPv4 traffic within the network.

Suitability of the Transition Categories for this Scenario

IPV4 OR IPV6 MECHANISMS
One of these mechanisms will be required to allow nodes to communicate with both IPv4 only nodes and IPv6 only nodes.

TUNNELING AND ENCAPSULATION MECHANISMS
The Internet is predominantly IPv4 so communication from one site to another will require the use of tunneling and encapsulation.

TRANSLATORS
A translator could be used to replace the NAT at the border of the network. This would work as communication outside the private network is only to one particular end-point the ISP's server, so IPv4 to IPv6 translation could occur.

Suitability of these Transition Mechanisms for this Scenario

DUAL STACK AND NAT
The dual stack mechanism if implemented in the correct order could be used on its own for communication within the private network but this would not allow communication with external nodes due to non-globally unique addressing. IPv6 addressing could be used for communication with other nodes in the network and IPv4 addressing for communication with the NAT. This mechanism will not suffer from scalability issues in this scenario, there are enough IPv4 addresses to support dual hosts as the address space is private. Manageability of two different IP addresses for each node is an issue, which will complicate administration.

DUAL STACK WITH CONFIGURED TUNNELS
To infer that you need configured tunnels means that you are likely to have some IPv4 infrastructure between IPv6 router or between a host and a router. In this scenario, upgrading all routers before hosts will be easier than configuring tunnels between hosts and routers and routers to routers. This mechanism could be used in a situation where regional offices e.g. Birmingham and Edinburgh have upgraded to IPv6 and the Head Office still using IPv4. A tunnel would need to be configured from Birmingham to Edinburgh, encapsulating the IPv6 packet within IPv4 so that it can be routed at the Head Office. Configured tunneling is more likely to be used

in large establishments or communication over a WAN where there is a large IPv4 infrastructure.

DUAL STACK WITH AUTOMATIC TUNNELS
If used would allow hosts to be updated before routers. Each host would need to be configured to use IPv4-Compatible IPv6 addresses, tunneling would then occur between end-points. This network is only small with only a few routers, all routers and routing are internal to the organisation and as such would be easier to upgrade than have the added problems of routing `flat' addresses and performance degradation of encapsulating most IPv6 packets within IPv4 packet.

DUAL STACK WITH AIIH
One of the main reasons in using the AIIH mechanism is if there are not enough IPv4 addresses for each Dual Stack node on the network. In this scenario they are using a private address space and therefore are not limited (within reason) to a number of IPv4 addresses.

DUAL STACK WITH DTI
Requires that all routers would need to support IPv6. If this is the case then if you have the Dual stack and IPv6 routing through the private network why use DTI? This mechanism would be used as part of a complex solution for larger organisations with direct external connections to the Internet and especially in the later stages of transitioning. I don't think its benefits could be of use in this scenario.

DUAL STACK AND TRANSLATOR
A NAT is already used at the border of the network which suits their needs. All nodes in the organisation can be upgraded to dual stack and the NAT be upgraded to convert IPv6 to IPv4 addresses. This would allow all nodes in the network to be able to communicate using only IPv6 and the translator used for converting IPv6 headers to IPv4.

TRANSLATOR
Could be used on its own to replace the NAT already installed meaning that the internal structure of the network could remain the same without any alterations within the private network. Could be used in the later stages once migration has been completed and most sites are using IPv6.

6OVER4
The internal network does not use multicasting so this mechanism is not relevant for this scenario.

DUAL STACK, 6TO4 AND TRANSLATOR
Could be used to give the Translator a unique IPv6 address by using

the unique IPv4 address for the translator and using it as an NLA. This would allow each node internal to the organisation to have a unique IPv6 address.

Solution 1

Reasoning behind the solution

Currently uses private address space so there is no problem with limited IPv4 addresses. The easiest approach in transitioning would be to use dual stack on all hosts. This solution does not require the complex methods of encapsulation.

Mechanisms Suggested in Solution

- Dual Stack
- Translator
- 6to4 (Option)

Stage 1 - Head Office

Starting with the Head Office in London as most traffic will be routed through here, it is essential that this is the first to be upgraded to IPv6 to allow for communication to allow routing from regional sites.

The order of Implementation within the Head Office can be followed from that in Scenario 2, Solution 1 but has been detailed again below:

DEFAULT ROUTER (R1)
Connects with all other offices. A software upgrade will be required to allow this to operate as a dual router.
The router will treat IPv6 as an independent protocol so therefore RIPv2 will need to be activated and configured for IPv6.

DHCP SERVER
Depending on whether this server is necessary is dependent on whether stateful auto-configuration is required. If required will need to be upgraded to dual stack to allow allocation of IPv4 addresses out of the private address space and also stateful IPv6 addresses.

DNS SERVER
Upgraded to a dual stack, a requirement for hosts to look up a destination node using DNS to find out if v4 or v6 address is to be used.

PROXY SERVER/TRANSLATOR
Will need to be bilingual. This is the conversion point for the network and will be translating between IPv6 and IPv4 and vice versa. The firewall will need no new configuration as the data sent to and from will be IPv4 format, until ISP migrates to IPv6.

SERVERS
Once the above have been converted then the Mail Server and any File Servers will need IPv6 installed. Again the Mail server may require some extra configuration.

WORKSTATION
Now all the dependencies have been configured all workstations will need to be upgraded to support IPv6. This can be in any order and there is no time limit.

Stage 2 _ London Offices

Once the head office has been upgraded, the regional offices in London can be upgraded. These can be carried out in whichever order desired. The following implementation rule shown in stage 1 must apply to each site:

- Default Router
- DHCP Server
- DNS Server
- Other Servers e.g. Mail Server etc
- Workstations

Stage 3 - Regional Offices

Once the London sites have been upgraded the regional sites can be upgraded. The order is as follows:
Birmingham
Manchester
Glasgow
Edinburgh

The order doesn't have to be followed but it should be noted that if Glasgow is upgraded before Manchester and Birmingham then tunnels will have to be configured from Glasgow's Router to the Head Office Router. Implementation in each office should be followed in accordance with Stage 1 and Stage 2.

Stage 4 - Final Stage

Once all nodes within the organisation have been upgraded to IPv6, the IPv4 component in each node can be deactivated to allow for just IPv6 traffic on the network. Deactivation will obviously have to be done in reverse order to the implementation i.e. disable IPv4 on the workstations first and routers last. The translator at the border will convert all IPv6 headers into IPv4 headers and vice versa.

Normally it is difficult for translators to convert from IPv4 to IPv6 e.g. when external to internal communication is initiated. Only one source in this case will be externally initiating communication and this will be the ISP server when sending SMTP traffic. The translator (will have to be an application translator) at the border of the network can detect and forward SMTP traffic to the correct node internally.

6TO4 OPTION
The 6to4 mechanism could be used in conjunction with the translator. The one unique IP address associated with the translator could be assigned to the NLA field creating a globally unique IPv6 prefix. This would allow all nodes within the organisation to have a globally unique IPv6 address and allow the NAT to receive either IPv6 or IPv4 packets.

Solution 2

Mechanisms Suggested in Solution

- Translator

If there was absolutely no need for the implementation of IPv6 within the organisation or if all IPv4 applications required intensive configuration to convert for IPv6 support there is another solution. In this case the NAT could just be upgraded to a translator supporting external IPv6 traffic and leaving the current internal infrastructure the same.
This adds to the problem of how IPv4 nodes can work out how to send to an IPv6 address external to the organisation. As all external traffic would be sent to the ISP address, the translator could be configured to send all external traffic to this one IPv6 address. The nodes could be configured manually to send any external data to a certain IPv4 address, which could be configured by routers to send on to the translator. The translator would need to know that this IPv4 address should be converted to the IPv6 address of the ISP

server.
This would be quite complex and it would be far easier, for long term administration and maintenance, to migrate the network to IPv6.

Scenario 4 - IPv4 dependent Applications

Introduction

There may be occasions in later stages where IPv4 hosts haven't been converted because applications running on them are IPv4 dependent and the application code cannot be changed but these need to exist within an IPv6 network. The diagram below shows a private IPv6/IPv4 network with IPv4 hosts (using IPv4 dependent applications).

`Diagram goes here'

Requirements

- IPv4 hosts with IPv4 only applications need to be able to interoperate with IPv6 only hosts and vice versa.
- Reduce IPv4 traffic on the network.

Suitability of the Transition Categories for this Scenario

IPV4 OR IPV6 MECHANISMS
Required to allow nodes to communicate with both IPv4 only nodes and IPv6 only nodes. Doesn't solve the problem of how IPv4 applications can communicate with IPv6 only nodes.

TUNNELING AND ENCAPSULATION MECHANISMS
A tunneling method could be used to allow IPv4 nodes to communicate with other IPv4 nodes on different segments.

TRANSLATORS
A translator placed at the border of a segment could be used for translation of IPv6 to IPv4 network but still there would be the same problem of converting IPv4 to IPv6. A translator that could be

placed in each node would be more appropriate.

Suitability of the Transition Mechanisms for this Scenario

DUAL STACK
Each node with IPv4 dependent applications installed will use the dual stack mechanism, every time the IPv4 application needs to communicate with another node it will only be able to do so using an IPv4 address. For nodes to communicate with these IPv4 only applications means they will need to each be assigned an IPv4 address. This would allow all hosts to communicate but the IPv4 application would not be able to communicate with IPv6 only nodes and vice versa.

DUAL STACK AND AIIH
May be used depending on the size of the network and size of IPv4 address pool and number of hosts requiring IPv4 addresses e.g. is there a big enough allocation of IPv4 addresses.

DUAL STACK, TRANSLATOR AND DTI
A translator could be used only if all the IPv4 Hosts were contained in one subnet. If this was not the case and the IPv4 hosts were distributed around the network a DTI mechanism could used to encapsulate the IPv4 packets inside IPv6 for routing. You would need DTI on every subnet that has IPv4 applications the DTI tunnel end-point could be a translator that could then convert the packets from IPv4 to IPv6 and vice-versa. The DTI solves the tunneling problem which would allow IPv6 hosts to communicate with an IPv4 only host but not the other way round.

DUAL STACK WITH BUMP IN THE STACK
This is by far the simplest and cleverest method. Each IPv4 dependant application's host once again implements the dual stack mechanism. The host is given both an IPv6 address and IPv4 address. The IPv4 applications will assume it is communicating using IPv4 when the host is actually using IPv6 at the network layer. The bump in the stack mechanism installed would then allow communications with other IPv4 only nodes and IPv6 only nodes. The whole network can use IPv6 addressing and only the few IPv4 dependant application's hosts need ever use IPv4 addresses.

Solution 1

Mechanisms Suggested in Solution

- Dual Stack
- AIIH

HOST IMPLEMENTATION
Installation of dual stack on each host will already have been done. Dual stack needs to be installed on each host and would then have to obtain an IPv4 address (using the principles of AIIH) or have a permanent IPv4 address for communication with the IPv4 dependent application's hosts. This would require a fair amount of work.

Solution 2

Mechanisms Suggested in Solution

- Dual Stack
- Bump in the Stack

HOST IMPLEMENTATION
Each host that is using the IPv4 dependent applications needs to be installed with the dual stack mechanism. This would allow these nodes to communicate with all other hosts but will only allow IPv4 dependent applications to communicate with other IPv4 hosts and not

IPv6 only hosts. To solve this problem the Bump In The Stack mechanism should then be installed on each of the IPv4 dependent application hosts. If it was a private network it would be possible to then eliminate all IPv4 traffic by disabling the IPv4 implementation on all nodes that are using the dual stack.

References

[6over4] B. Carpenter & C. Jung,'Transmission of IPv6 over IPv4 Domains without Explicit Tunnels', RFC 2529

Author's Address

Tim Larder
Barbrook Cottage,
Holmesdale Road,
South Nutfield,
Surrey, RH1 4JE, UK.

Phone: 0411 645514
Email: Tim.Larder@virgin.net

--=====================_878331==_
Content-Type: text/plain; charset="us-ascii"

--=====================_878331==_--