TECHNOLOGY WHITE PAPER

IPv6 in Fixed and Mobile Networks

IPv4, designed roughly two decades ago, has shown its durability in recent years as the Internet has grown at an exceptional rate. While it has adapted to an increasing number of users and applications, it has been showing signs of weakness. Concerns about available address space, poor adoption of security, imperfect mobility support, and apprehension about scalability prompted the development of IPv6. While IPv6 addresses these issues, it also introduces new auto-configuration mechanisms, which should significantly reduce network operating costs. These changes will allow for the current rate of growth of the Internet, not only in terms of the number of connected devices, but also in the number of applications, to be maintained well into the future.

ALCATEL

D. J. Wilson, R. Dragnea

IPV6 IN FIXED AND MOBILE NETWORKS

While IPv4 has proved resilient to the enormous growth of the Internet, IPv6 offers better support for mobility allied to superior scalability for the future

Introduction

In recent years, the Internet has seen an explosion in the number of users as well as in the volume of traffic as more and more media-rich applications are deployed. These changes have strained not only the deployed Internet infrastructure, but also the protocols themselves.

Introduced over 20 years ago, the Internet Protocol (IPv4) [1] was designed at a time when the Internet consisted of a handful of nodes connected by unreliable low speed links. At that time, the Internet was fueled by relatively lightweight applications, such as TELNET, File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), Unix to Unix Copy Program (UUCP), Gopher and the Network News Transfer Protocol (NNTP). IPv4 was tailored for these applications and network conditions.

Over time, and primarily as a result of the World Wide Web, the popularity of the Internet increased, providing momentum for the development of new media- and bandwidth-rich applications, such as peer-to-peer file sharing, online gaming, streaming applications, Voice over Internet Protocol (VoIP), and voice- and videoconferencing. Many of these applications have stringent real-time requirements; although they can handle loss, they rely on very little traffic being lost, as well as on low latency and jitter. This is a far cry from the early days of the Internet when IPv4 was designed.

With recent developments in wireless networking technology, mobile computing and Internet-enabled handheld devices, such as mobile phones, the number of mobile devices connected to the Internet has increased dramatically, and has been projected to exceed the number of fixed nodes in the next few years[1]. While supported by IPv4, mobility involves new network requirements for which IPv4 was not designed.

In recent years, IPv4 has struggled to keep up with new applications and concern over diminishing unallocated address space has worsened. IPv6 [2] was developed to address concerns over network scaling and the problem of address space. It will allow for continued growth in the Internet and enable new applications, such as mobility, to be widely adopted.

The Internet plays a unifying role in many activities, but its evolution to a new protocol varies in different regions of the world. IPv6 drivers in Asia were originally political, but are becoming increasingly commercial. In Europe, the focus is on academic and research networks while North America has taken more of a "wait and see" attitude. However, recently the US government has mandated the use of IPv6 in government networks by 2008. More information on the IPv6 market can be found in [3].

Evolution to a New Internet Protocol: IPv6

IP address depletion

When IPv4 was introduced, it made use of a class-based addressing scheme, and its address space seemed inexhaustible. Nevertheless, the limited granularity of address allocations resulted in significant waste. About a decade after its introduction, it was becoming apparent that the IPv4 class-based addressing scheme would in time

Fig. 1 IPv4 and IPv6 header formats

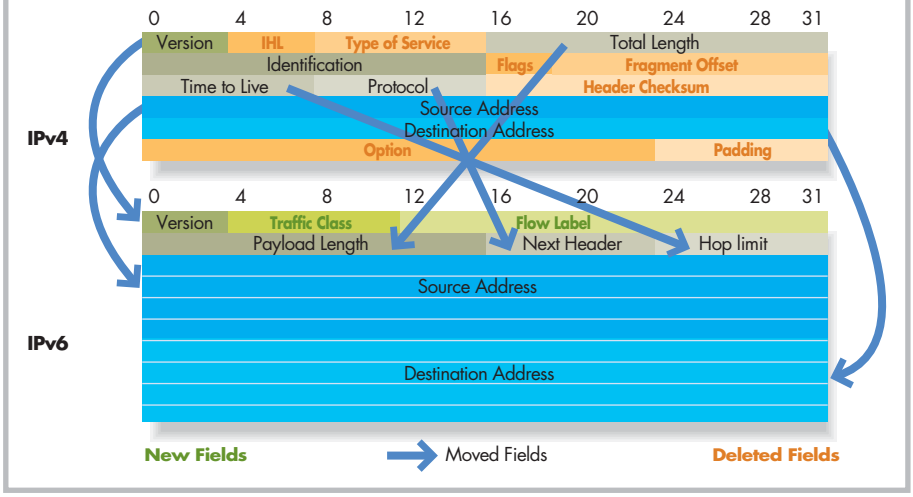

[1] Pyramid Research 2004

result in a shortage of addresses. This prompted the development of Classless Inter-Domain Routing (CIDR) [4], which allows the size of a network to be specified with significantly more granularity, reducing the number of unused, yet allocated, addresses.

In addition to CIDR, Network Address Translation (NAT) [5] has been used to extend the life of the IPv4 address space. NAT enables several nodes to concurrently share a single globally routable IPv4 address. However, NAT has its downsides as it breaks end-to-end connectivity, preventing some applications from functioning without application proxies.

Although CIDR and NAT have extended the life of the IPv4 address space for another decade, an address shortage is again looming. Extending the current address allocation trend shows that the IPv4 address space could be exhausted as soon as 2015 [6][2].

As *Figure 1* shows, the most visible change between the IPv4 and IPv6 headers is the dramatic increase in the address size from 32 bits to 128 bits. For IPv6, this means there are 340,282,366,920,938,463,463,374,607,431,768,211,456 unique addresses compared with the comparatively miniscule 4,294,967,296 for IPv4. To put this number in perspective, IPv6 provides approximately 10^{22} addresses per square meter of the Earth's surface, so even with an inefficient allocation scheme, it should provide an adequate number of addresses for quite some time.

Simplified header

IPv6 reduces the workload required to process the IP header by reducing the number of fields compared to the IPv4 header, and by giving the header a fixed length. *Figure 1* shows the evolution from the IPv4 header to the IPv6 header.

The Header Checksum (HC) field was removed on the premise that higher layer protocols, such as the Transmission Control Protocol (TCP), would detect and recover from transmission errors. This simplified IP header processing as the HC needed to be recalculated at every router as a result of modifying the Time To Live (TTL) field.

Fixing the size of the IPv6 header was achieved by replacing the IPv4 Options field with Extension Headers, which reside between the IPv6 header and any upper layer headers. Extension headers also increase the flexibility of how and what optional Internet-layer information is included, and make it possible to specify the intended recipients of the information.

Fragmentation

Another burden imposed on routers by IPv4 was packet fragmentation, which is the result of a node transmitting a packet larger than the maximum size packet allowed on a link on the way to the destination. Fragmentation is expensive as it requires a packet to be split into segments and an IP header to be added to each segment. IPv6 specifies that fragmentation is to be performed by the source node and not routers. Consequently, the "Identification", "Flags", and "Fragment Offset" fields of the IPv4 header, which were used for fragmentation, have been moved into a fragmentation extension header in IPv6, further reducing the number of header fields.

Integrated security

Several authentication and encryption security mechanisms were created in parallel with the development of IPv6. Outlined in RFC 2401 [7], these security mechanisms can be used with both IPv4 and IPv6, but only IPv6 mandates that all nodes implement these security features. In the case of IPv6, this means that applications can now reliably use standard security mechanisms at the IP level, instead of implementing security at the application level on a per-application basis.

Auto-configuration

While most of the stateful configuration mechanisms available in IPv4, such as the Dynamic Host Configuration Protocol (DHCP) [8], have been adapted to IPv6, stateless auto-configuration mechanisms have been added to IPv6, allowing devices to use discovery protocols to learn the information needed to configure themselves. These mechanisms should dramatically reduce the cost of building and maintaining an IPv6 network.

Mobile IP

The goal of Mobile IP is to allow nodes to remain reachable at a single address while moving from network to network. While dynamic address assignment schemes, such as DHCP, allow nodes to obtain an address belonging to the network to which they are attached, any movement between networks using this method will result in any established transport layer connections being lost.

To enable mobile nodes to move freely while maintaining any active connections, a global unicast address is required in addition to an address that identifies the current location of the mobile node. How these addresses are used and who owns them is the major difference between Mobile IPv4 [9] and Mobile IPv6 [10].

Figures 2 and *3* show two common configurations for Mobile IPv4. While in its home network, the Mobile Node (MN) obtains a local address using an address assignment mechanism, such as DHCP or static assignment. This home address is the MN's point of presence on the Internet both while residing in its home network and when visiting foreign networks. Regular routing mechanisms will result in any packet sent to an MN's home address being forwarded to its home network.

When the MN detects that it has left its home network, it obtains a Care-of Address (CoA) from either a Foreign

[2] The referenced web site recalculates estimates for IPv4 address exhaustion daily based on historical data. Consequently, the estimated exhaustion date predicted on the site may vary.

Agent (FA) located in the foreign network, or using an address assignment mechanism like DHCP. The MN registers its CoA with a router in the home network, known as a Home Agent (HA), which acts on behalf of the MN when it is visiting a foreign network. This results in a binding in the HA between the home address and the CoA.

When a Correspondent Node (CN) sends a packet to the MN's home address, the packet is routed to the home network. If the MN is away, the packet is intercepted by the HA, which tunnels it to the CoA bound to the MN's home address. Depending on the owner of the CoA, either the foreign agent or the MN receives the tunneled packet and performs the decapsulation. *Figure 2* shows the case in which the MN owns the CoA and terminates the tunnel itself. If the foreign agent owns the CoA, it forwards the packet directly to the MN, as shown in *Figure 3*.

To respond to the CN, the MN sends packets using its home address as the source IP address. Only normal routing mechanisms are required for this packet to reach the CN. However, there are a number of issues associated with the configurations shown in *Figures 2* and *3*, and Mobile IPv4 in general.

The first issue is ownership of the CoA, which is owned by either the FA or the MN. Both options have disadvantages. If the MN owns the CoA, every MN in the foreign network consumes at least two global addresses. Since address space in IPv4 is already scarce, it does not make sense to assign two or more addresses to every mobile node as this would further accelerate the allocation of IPv4 addresses. In the other case, an FA allows several MNs to share a single CoA. However, for the foreign agent to forward packets to the MN, the two nodes must reside on the same link, otherwise the packet will be routed back to the home network, creating a routing loop and eventually resulting in the packet being discarded. This scenario solves the address scarcity problem, but requires a foreign agent to reside on each link to which a mobile node could attach.

Another reason for using FAs relates to the use of ingress filtering on many routers. Security concerns surrounding IP address spoofing prompted changes in the way that packets are forwarded. RFC 2267 [11] and other security advisories recommend that IP packets be forwarded based not only on the destination address, but also on the source address. Consequently, packets that have a source address that does not belong to a network downstream from the link on which the packet arrived will be dropped. In order to bypass ingress filtering, reverse tunneling (see *Figure 4*) was proposed in RFC 3024 [12]. This specifies that the MN encapsulates all of its packets to the FA, at which point the FA will replace the encapsulation and tunnel the packets to the HA using the CoA belonging to the FA.

Fig. 2 Mobile IPv4 communication without a foreign agent

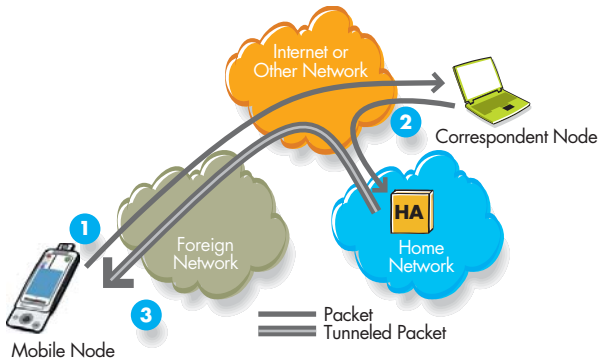

Fig. 3 Mobile IPv4 communication with a foreign agent

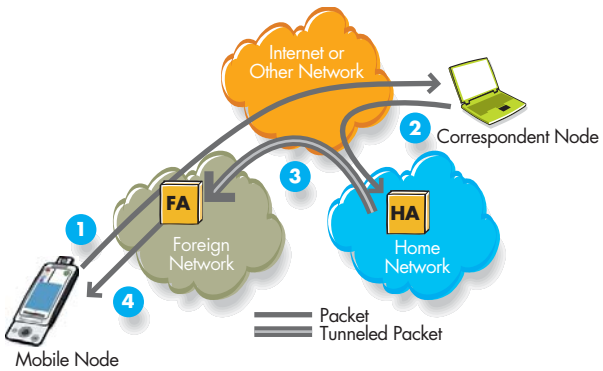

Fig. 4 Mobile IPv4 communication using reverse tunneling

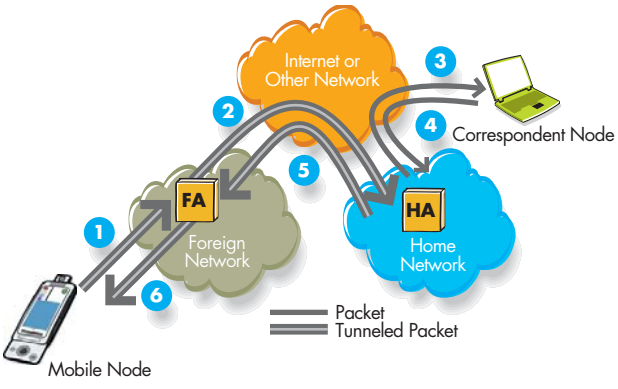

The HA removes the encapsulation and forwards the packet on behalf of the MN. While this circumvents ingress filters on routers, it has several downsides, including the additional overhead associated with tunneling packets between the MN, FA and HA.

A further limitation of Mobile IPv4, which is visible in *Figure 2*, is known as triangle routing. The first major disadvantage of triangle routing is that the HA becomes a bottleneck. Every MN from a particular home network has its traffic forwarded through the HA. If reverse tunneling is used, this becomes even worse as traffic in both directions is relayed through the HA, and packets cross the Internet twice.

Mobile IPv6

Mobile IPv6 addresses many of the problems inherent to Mobile IPv4. One distinct advantage is that it was defined at roughly the same time as IPv6. In contrast, Mobile IPv4 was introduced nearly a decade and a half after IPv4. While similar to Mobile IPv4, Mobile IPv6 makes numerous small changes to avoid some of the problems associated with its predecessor. Several other benefits are also achieved as a result of it being based on IPv6.

Figure 5 shows the simplest configuration for Mobile IPv6, which mandates the use of reverse tunneling, known in Mobile IPv6 as bidirectional tunneling. It is used for the same reason as in Mobile IPv4, that is, to prevent packets from being discarded as a result of ingress filtering. The main difference from Mobile IPv4 is that the tunnel terminates at the MN, and as a result, no FA is required. This enables mobile nodes to travel to any foreign network and operate correctly without relying on foreign agents, as long as they can obtain a CoA.

The configuration shown in *Figure 5* suffers from the same triangle routing problems as with Mobile IPv4. For this reason, Mobile IPv6 specifies a second mode of operation, known as "route optimization", shown in *Figure 6*. In this mode, a mobile node can establish a mobility binding directly with a correspondent node, in the same way that it establishes a binding with the HA. This allows the MN and CN to communicate without relaying packets through a home agent, and to remain connected as the MN moves between networks. Additionally, in this mode of operation, tunneling is not used for communication between the MN and CN. Instead, a new IPv6 "Destination Option" and "Routing Header" are defined. The new destination option, or "Home Address Option", is used by a mobile node to notify the recipient of its home address. In the opposite direction, the Type 2 routing header is used by a CN to route a packet to a mobile network's CoA. Upon receiving the packet, the MN can fetch its home address, which is used as the final destination for the packet, from the routing header. The motivation for introducing a new routing header type, which allows for only a single IPv6 address, is that it enables firewalls to apply different filters to Mobile IPv6 packets than to regular source routed packets.

The lack of addresses in IPv4 resulted in a need for foreign agents. As this limitation no longer exists in IPv6,

Fig. 5 Mobile IPv6 communication using bidirectional tunneling

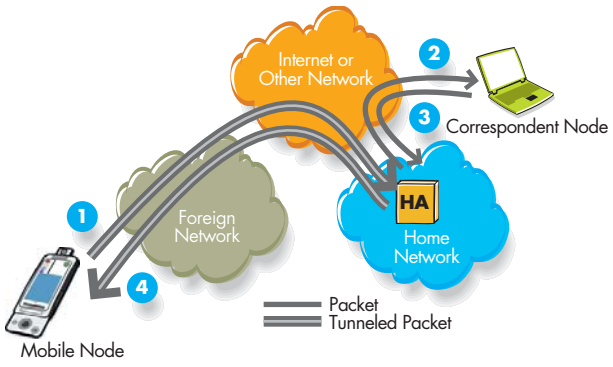

Fig. 6 Mobile IPv6 communication using route optimization

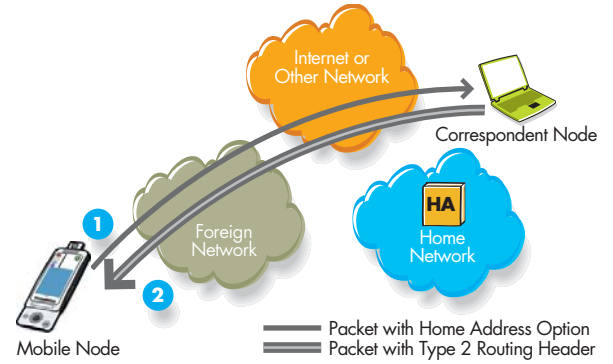

Mobile IPv6 has no use for foreign agents. Consequently, every mobile node can be assigned several addresses, one home address, and one CoA for each foreign network to which it is attached. Furthermore, configuration of the CoAs can rely on the IPv6 auto-configuration mechanisms rather than on DHCP.

Conclusion

The Internet has evolved dramatically since IPv4 was developed, with the result that the dynamic nature of today's Internet has stretched IPv4 in ways for which it was not designed. Concerns over the ability of IPv4 to continue to scale safely prompted the development of IPv6.

IPv6 was built on two decades of experience with IPv4. To meet the current strong growth in the number of connected devices, the address size has been increased significantly, enabling an enormous number of devices to be supported. In addition, simplified headers and hierarchical addresses have been implemented to reduce the burden on routers, to allow for traffic growth and to support an increasing number of applications. Integrated security

mechanisms will allow applications to seamlessly use authentication and encryption functionality, without the difficulties of application level security which were typical with IPv4. Auto-configuration mechanisms allow devices to statelessly configure themselves simply by connecting to a link. Finally, while mobile IP is not something new for IPv6, the new protocol offers mobility improvements that simplify deployment and offer better performance than IPv4.

Although the switch from IPv4 to IPv6 will probably take more than a decade, it is clear that it is needed to allow the Internet to grow further. This transition is likely to spawn new services and revenue-generating applications that are not possible in today's Internet.

References

[1] J. Postel (Editor): "Internet Protocol", *RFC 791*, ftp://ftp.rfc-editor.org/in-notes/rfc791.txt, September 1981.

[2] S. Deering, R. Hinden: "Internet Protocol, Version 6 (IPv6) Specification", *RFC 2460*, ftp://ftp.rfc-editor.org/in-notes/rfc2460.txt, December 1998.

[3] A. Chaudhry, B. Crosby, A. Zinin: "Making the Move to IPv6", *Alcatel Telecommunications Review*, 4th Quarter 2004, Web supplement of this issue.

[4] V. Fuller et al: "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", *RFC 1519*, ftp://ftp.rfc-editor.org/in-notes/rfc1519.txt, September 1993.

[5] P. Srisuresh, K. Egevang: "Traditional IP Network Address Translator (Traditional NAT)", *RFC 3022*, ftp://ftp.rfc-editor.org/in-notes/rfc3022.txt, January 2001.

[6] G. Huston: "IPv4 Address Space Report", http://bgp.potaroo.net/ipv4/, July 2004.

[7] S. Kent, R. Atkinson: "Security Architecture for the Internet Protocol", *RFC 2401*, ftp://ftp.rfc-editor.org/in-notes/rfc2401.txt, November 1998.

[8] R. Droms: "Dynamic Host Configuration Protocol", *RFC 2131*, ftp://ftp.rfc-editor.org/in-notes/rfc2131.txt, March 1997.

[9] C. Perkins (Editor): "Mobility Support for IPv4", *RFC 3344*, ftp://ftp.rfc-editor.org/in-notes/rfc3344.txt, August 2002.

[10] D. Johnson, C. Perkins, J. Arkko: "Mobility Support in IPv6", *RFC 3775*, ftp://ftp.rfc-editor.org/in-notes/rfc3775.txt, June 2004.

[11] P. Ferguson, D. Senie: "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", *RFC 2267*, ftp://ftp.rfc-editor.org/in-notes/rfc2267.txt, January 1998.

[12] G. Montenegro (Editor): "Reverse Tunneling for Mobile IP, revised", *RFC 3024*, ftp://ftp.rfc-editor.org/in-notes/rfc3024.txt, January 2001.

David J. Wilson is a Research Engineer working on the Intelligent Switching and Routing project in the Alcatel Research & Innovation Center in Ottawa, Canada. (david.j.wilson@alcatel.com)

Raluca Dragnea is Project Leader and Local Unit Director of the Intelligent Switching and Routing project in the Alcatel Research & Innovation Center in Ottawa, Canada. She is a distinguished member of the Alcatel Technical Academy. She is a Distinguished Member of the Alcatel Technical Academy. (raluca.dragnea@alcatel.com)

Abbreviations

CIDR Classless Inter-Domain Routing
CN Correspondent Node
CoA Care-of Address
DHCP Dynamic Host Configuration Protocol
FA Foreign Agent
HA Home Agent
HC Header Checksum
HN Home Network
IP Internet Protocol
MN Mobile Node
NAT Network Address Translation
SMTP Simple Mail Transfer Protocol
TCP Transmission Control Protocol
TTL Time to Live
VoIP Voice over Internet Protocol