Internet Protocol Next Generation:

Saving the Internet in the New Millennium

A Paper Presented By:

Robert A. Kondilas
Security Engineer
Technical Services
Systems Integrity
networkMCI Services
MCI

Preface

It is important to understand the motivation behind the need for change with the current Internet Protocol. The need for change is driven by a number of factors that each contributes to the building momentum of radically restructuring the protocol that delivered the Internet to the world and made it the success it is today.

This paper will discuss some of the pitfalls of the current Internet Protocol (IPv4) as well as what is in the proposed design for its successor...Internet Protocol Next Generation (IPng or IPv6). The paper is based on the information available through the Request for Comments (RFC) papers submitted to the Internet Engineering Task Force (IETF) which is tasked with governing the solution to the current Internet Protocol.

In the best estimate, the Internet Protocol Next Generation is still three (3) to seven (7) years away from wide scale implementation.

Fuel for the Fire of Change

Today new technologies are emerging on the Internet. The deployment of such online applications that allow banking, researching, as well as inquiring on one's phone bill are bringing more and more users to the Internet. However, the driving force does not lie solely with individuals. Companies, seeing this trend, are moving to the Internet as well. The explosion of technical support content, software patches, and customer surveys are also heavy contributing factors to the growth in the popularity in the Internet. Note that this environment is composed of two elements that feed off of each other's hunger. If companies did not initially see the potential of a presence on the World Wide Web (WWW), then their customers that had WWW access did and showed them. By the same token, for customers of companies that had a WWW presence, the demand by the customers drove the company to build out their WWW site with more functionality. It is the technology seen today that is driving the Internet of tomorrow.

Architecture Downfalls

Because of its inception over twenty (20) years ago, one of the basic downfalls of the Internet Protocol version 4 (IPv4) architecture is the scalability to more modern networks that exist today. It can be noted that there are three (3) basic types of networks that undermine IPv4's two-level addressing hierarchy:

- People who work within an organization that is connected to the Internet, and are considered members of their organization's network. There can be millions of such organizations worldwide, each with potentially millions of members. As Internet members, each intranet must have its own unique address. Within an intranet, each interface address (host number) must also be unique.

- People who work within an organization that is not connected to the Internet, and can be considered members of their organization's network. There can likewise be millions of people in this situation, scattered among millions of organizations, with an added wrinkle: by simply signing a service contract and installing some network hardware, they can all be made part of the Internet community instantly. How many intranet addresses will they need to change to make them unique within the Internet community?

- Individual network users, working from home or while traveling, connected through telephone lines or packet radio links.[1]

Each of the three types of users listed encompasses the majority of users that will be accessing services via the Internet.

Current and Future Growth

Tomorrow, a much different and varied Internet is approaching. Many trends are pushing towards this new Internet. Some of them include networked entertainment, desktop video conferencing, and device control. Networked entertainment is where a customer at home can request a movie from a source to be delivered at any time via the Internet. Video conferencing on the desktop not only aids corporate customers but can also bring families closer together. Lastly, device control could encompass anything from turning off the oven while you are on vacation to enabling the house alarm from work.

Not only are new technologies bringing about change but also direction of the work industry. Today, companies are looking for ways to reduce capital expenditures, increase income, as well

[1] Jim Bound and Al Cini. "IPv6 - Powering Up for Cyberspace." September 3, 1996. <http://www.digital.com/info/ipv6/ipng.html>. (1 November 1996).

as cover the world more effectively. They are equipping their sales forces with laptop computers with high-speed modems. Members of the sales force are meeting with clients and receiving orders that need to be processed immediately. Through the Internet, they can log in remotely to their company Intranet and request the product on their client's behalf. This means that every laptop must have a static IP (Internet Protocol) address or that the company must allocate the number of IP addresses to accommodate the maximum number of laptops that could possibly access the company Intranet at one time. Companies are also looking towards telecommuting as a way to reduce expenses. By moving their employees out of the office and into their individual homes, they will need to supply an IP address to any employee needing Internet access to get to the company intranet. It should be noted that while companies are moving to the mobile environment, they most likely are not dismantling their Local Area Networks (LANs) or Wide Area Networks (WANs) in the process.

With the advancement of technology and the movement of corporations towards the Internet, the size of the address space (the total number of IP addresses) is still limited to four billion addresses (2^{32}) . This is where the problem lies.

Sizing Up the Task

One might think that an address space comprised for four billion addresses could adequately handle the needs on the Internet community for some time to come. When put into perspective, roughly two (2) out of every three (3) people on earth could have an IP address assigned to them. However, the problem does not necessarily pertain to the number of addresses but rather how those addresses are allocated. Twenty years ago it was believed that a large network capable of accommodating several million systems would be a permanent solution. However, the conception of a network any larger than the specifications that had been set was ludicrous. The possibility that hundreds of millions of systems could be attached to the Internet is no longer a dream but a potential reality. As the Internet becomes the medium by which to exchange data, the current addressing scheme is grinding to a halt.

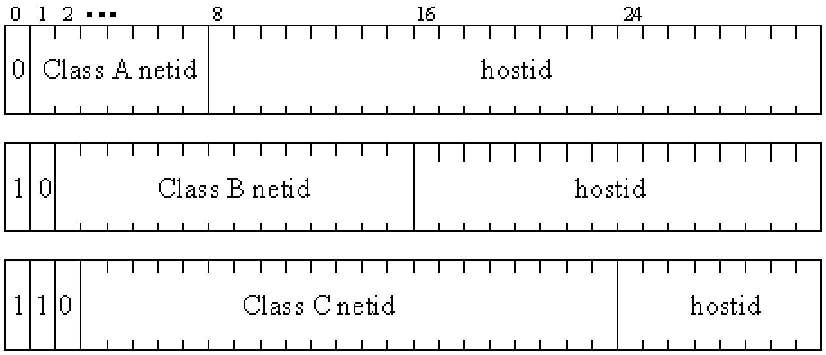

Figure 1. IPv4's 32-bit address field, showing its three unicast address formats.[2]

IPv4 was designed to create a hierarchy consisting of two levels: network and host numbers. In order to use the thirty-two (32) bit address size effectively as well as fulfill the varied needs of organizations requesting addresses, it was divided into three types of unicast addresses as pictured in Figure 1:

- **Class A** addresses consist of a 7-bit network number, followed by a 24-bit host number. Class A addresses are intended for use by the world's 128 largest organizations, each of which can then assign as many as 2^{24} unique addresses to its networked devices. Except for a few reserved and some perhaps reassignable values, Class A addresses are all taken.
- **Class B** addresses consist of a 14-bit network number, followed by a 16-bit host number. As many as 16,384 organizations can claim Class B addresses, supporting communities with up to 65,536 members. Class B addresses are, likewise, practically all used up.
- **Class C** addresses, with their 21-bit network numbers, are intended for very small network communities with 256 connections or less. From a pool of a little more than two million

[2] Ibid.

> possible values, growth in today's Internet is being sustained by assigning blocks of unused Class C addresses to new subscribers, as well as by creatively re-using certain previously assigned addresses.[3]

What makes this all become more lucid is the fact that it is not the number of addresses that is propelling us towards a change but rather how those addresses are dispersed. The two factors that contribute to this imminent problem are the excessive routing overhead and wasteful address assignment.

With any network, as the number of systems attached to it increases, the overhead of the system increases as well. In the 1970s, when relatively few systems were attached to the Internet and routers only had to maintain table sizes of one thousand (1000) entries, the overhead of routing data from one system to another was minimal. If the same routing scheme were still implemented, the routers would not be able to maintain tables sized in the millions of entries.

Assignment of addresses is the other concern. Until recently, it was, at best, grossly mismanaged. For example, a company with one thousand (1000) computers would not request four (4) Class C addresses (in order to handle 1024 computers) but rather a Class B address. This means the company would essentially be wasting 64,536 usable addresses. The utilization of the Class B address space that the example company registered for would be 0.015%. In comparison, if the same company had requested five (5) Class C addresses (in order to make way for future expansion totaling at the most 1280 computers), the utilization would be 78.125%.

Security

Currently, IPv4 has no security features. It was the decision that in order to have the highest level or service over the Internet, any security implementation would hinder traffic to some degree and thus slow service over the network. Without security the Internet, as well as any of its users, is a potential target of attacks.

[3] Ibid.

The New Challenger: IPng

The Internet Protocol Next Generation (IPng) is the successor designed to replace the current version of the Internet Protocol (IPv4). It has radically new features including addressing and security. The goal of IPng was simple: overhaul the current version of IPv4. More specifically, the outline of the areas that IPng would address were:

- Support billions of hosts, even with inefficient address space allocation.
- Reduce the size of routing tables.
- Simplify the protocol, to allow routers to process packets faster.
- Provide better security (authentication and privacy) than the current IP.
- Pay more attention to type of service, particularly for real-time data.
- Aid multicasting by allowing scopes to be specified.
- Make it possible for a host to roam without changing its address.
- Allow the protocol to evolve in the future.
- Permit the old and new protocols to coexist for years.[4]

In the following pages, some of the most significant features of IPng will be discussed.

Architecture

While IPv4 used an address that was thirty-two (32) bits in length, IPng uses 128. This equivocates to the address of IPv4 squared twice. What makes this so unique is that this addressing scheme could give an address to possibly every atom in the composition of the earth. However, it is not merely the size of the address space that makes IPng so promising. It is the structure in which it is implemented.

In the Figure 2, the three types of unicast addressing that IPng recognizes are depicted. Defining the three (3) types of unicast addressing gives a better picture:

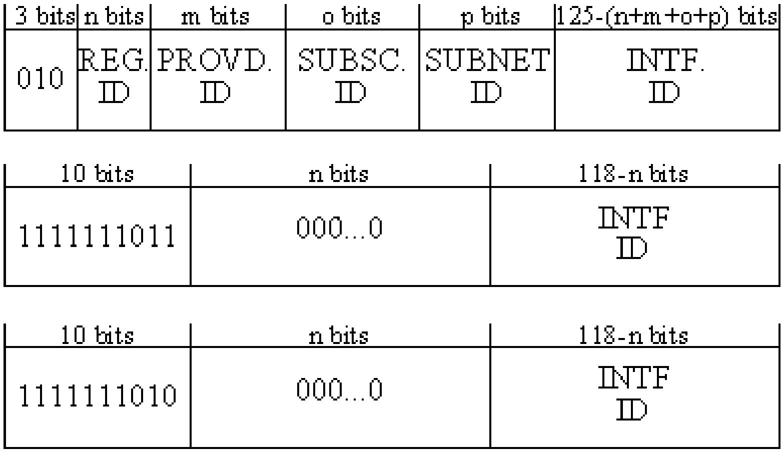

3 bits	n bits	m bits	o bits	p bits	125-(n+m+o+p) bits
010	REG. ID	PROVD. ID	SUBSC. ID	SUBNET ID	INTF. ID

10 bits	n bits	118-n bits
1111111011	000...0	INTF ID

10 bits	n bits	118-n bits
1111111010	000...0	INTF ID

<u>Figure 2.</u> IPng's 128-bit address field, showing its three unicast address formats: provider-based addresses (top), site-local-use addresses (middle), and link-local-use addresses (bottom)[5]

- **Provider-based unicast** addresses are assigned by an Internet Service Provider (ISP) to an organization, offering globally unique Internet addresses to all of the organization's members for easy integration within the worldwide Internet community. Devised as part of CIDR, the basic mechanism for assigning these addresses through ISPs is already in place.

- **Site-local-use addresses** can be assigned to the network devices within an isolated intranet. Later, should the organization decide to join the Internet community, all of its sitewide local addresses automatically become globally unique provider-based addresses with one

[4] Andrew S. Tanenbaum. *Computer Networks.* 3rd Edition. Upper Saddle River, New Jersey: Prentice Hall. 1996. p. 437.

[5] Jim Bound and Al Cini. Ibid.

administrative operation. This is much easier than the corresponding IPv4 procedure, which usually involves tediously changing every address on every computer within the intranet.

- **Link-local-use addresses** are designed for use by individuals on a single communications link – such as mobile laptop computer users connected through phone lines (voice or ISDN) or radio links.[6]

In addition to the unicast addressing, it also supports multicast and anycast addressing. Multicasting will probably have the widest use among the three. It will involve identifying two or more recipients of the data in transit. The scope of a particular multicast address can be confined to a single system, restricted within a specific site, associated with a particular network link, or distributed worldwide.[7] The newest member to the addressing family is the anycast address. With this scheme, two or more systems with the same value can be recipients of the data packets. The distance from the system originating the packet determines the specific system that gets the packets. The distance however is not a measure of length per se, but rather how the protocol measures the distance (i.e. one protocol could use the number of routers that the packet must go through before reaching its destination as the distance).

One last architecture consideration that needed to be addresses with IPng was its inherent compatibility with IPv4. Without a solid path of interoperability between the two protocols, the deployment of IPng would be stillborn. After all, with all of the capital invested by governments and companies all over the world into IPv4, it would not be prudent to require the purchase of all new equipment to support the new protocol. Therefore, the new architects of IPng implemented a design that would allow administrators of networks to upgrade their systems and routers as needed or as time permitted. The ingenuity of this design is that it allows applications running on the current protocol to be ported to IPng without any development downtime. In addition, hardware at the network and computer level would not need to be replaced but rather inexpensively upgraded through the manufacturer via a firmware revision.

Header Composition

One of the obstacles encountered with IPv4 is the timely delivery of packets through the network to its destination. While the addressing, as mentioned earlier, can be attributed to the delay in service during transmission, another factor exists. As packets of data travel through the IP network, the data contained in the header of each packet plays a significant role in how the packet

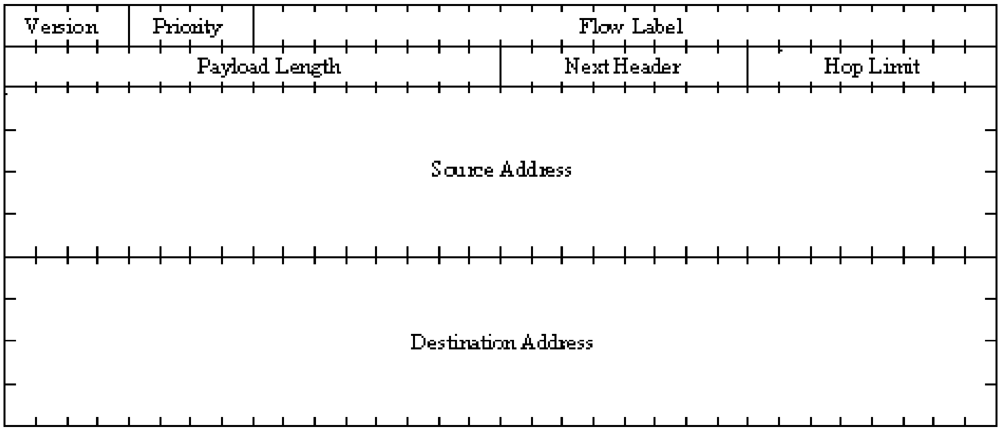

Figure 3. The IPng fixed header[8]

[6] Ibid.
[7] Ibid.
[8] Andrew S. Tanenbaum. Ibid. p. 439.

is handled. A less than optimal implementation of a header design is as effective as a detour on the highway: the data gets to it destination but it takes a little longer than necessary. The design of the IPv4 header, while effective with the Internet of ten (10) and twenty (20) years ago, is a source of burden today.

It is for this reason that the architects of the IPng protocol offered a solution that said:

- A simple, minimalist approach should be taken, avoiding lots of statically allocated but rarely used option fields.
- The architecture should be extensible, allowing the easy addition of optional features.[9]

This allowed for a design that offered flexibility as well as increased performance. Of note is that while the packet size of IPng increased four fold over its predecessor, the size of the header contained within a packet only doubled in size. Put simply, IPng is able to carry almost four times the amount of data that IPv4 could carry and yet carry it through an identical network at a faster rate.

Extension Header Composition

IPng also has the capability of transporting information in addition to the standard header. The information is carried in an Extension Header layer following the Header. It contains data considered not necessary to the delivery of the packet and therefore optional. Some of the fields that will be supported and their purpose:

• Routing	Full or partial route to follow
• Fragmentation	Management of datagram fragments
• Authentication	Verification of the sender's identity
• Encrypted Security Payload	Information about the encrypted contents
• Hop-by-Hop Option	Miscellaneous information for routers
• Destination Options	Additional information for the destination[10]

It is important to note that each of the options described above will not be processed at any time through the network but rather at the packet's destination. Any number of the options can be utilized but none of them are necessary for a transmission of packets.

Security

One of the most significant enhancements of IPng over IPv4 is the introduction of security. Essentially, there is no effective privacy or authentication mechanism. The implementation of security is achieved in two ways:

- Authentication Header
- Encapsulating Security Header

The Authentication Header is an Extension Header that provides authenticity and integrity to the data packet but provides no measure of privacy. MD5 has been proposed as a standard to implement globally. The fact remains that it is algorithm independent and can be customized within a network or among specified networks to utilize another algorithm with interoperability unaffected. From a network intrusion standpoint, this will virtually eliminate all attacks of masquerade (otherwise known as "IP Spoofing"). Because the Authentication Header provides

[9] Jim Bound and Al Cini. Ibid.
[10] Andrew S. Tanenbaum. Ibid. p. 444.

only authenticity and integrity of the data, with no capability for confidentiality, this feature could see widespread use throughout the Internet community.

The Encapsulating Security header is a feature that provides a means of integrity and confidentiality of data transmitted in packets. It is flexible enough to handle a variety of algorithm-based encryption methods. Globally, however, the implementation of the Data Encryption Standard Cipher Block Chaining (DES-CBC) will be used as the standard algorithm. It is expected that export regulations will affect the use and implementation of this feature depending on the origin and destination countries and the laws that govern data traffic in those specific areas. In addition, this feature will protect networks from “sniffing” attacks, where clear text transmissions of data (such as usernames and passwords) are intercepted.

Is IPng the Answer?

In order to understand the movement towards the IPng, it must first be decided that there is neither one definitive feature that IPv4 lacks nor one feature that IPng offers that has fueled the fire of change. IPv4 is a victim to the Darwinian Theory of Evolution: Only the Strong Survive. It has lasted for over twenty-five (25) years and now a stronger generation is about to succeed it.

The simple fact is that the features that IPng has to offer are too promising to surpass. To state its case :

- It solves the Internet scaling problem, provides a flexible transition mechanism for the current Internet, and was designed to meet the needs of new markets such as nomadic personal computing devices, networked entertainment, and device control. It does this in a evolutionary way which reduces the risk of architectural problems.

- Ease of transition is a key point in the design of IPng. It is not something that was added in at the end. IPng is designed to interoperate with IPv4. Specific mechanisms (embedded IPv4 addresses, pseudo-checksum rules etc.) were built into IPng to support transition and compatibility with IPv4. It was designed to permit a gradual and piecemeal deployment with a minimum of dependencies.

- IPng supports large hierarchical addresses that will allow the Internet to continue to grow and provide new routing capabilities not built into IPv4. It has anycast addresses that can be used for policy route selection and has scoped multicast addresses that provide improved scalability over IPv4 multicast. It also has local use address mechanisms that provide the ability for "plug and play" installation.

- The address structure of IPng was also designed to support carrying the addresses of other Internet protocol suites. Space was allocated in the addressing plan for IPX and NSAP addresses. This was done to facilitate migration of these Internet protocols to IPng.

- IPng provides a platform for new Internet functionality. This includes support for real-time flows, provider selection, host mobility, end-to- end security, auto-configuration, and auto-reconfiguration.[11]

In conclusion, IPng is the future. It is no longer a question of “Why?” but rather a question of “When?” While it may still have some issues that need modification or resolution, its framework stands as the successor to the protocol that carried the world to the 21st Century only to make the necessary handoff.

[11] Robert M. Hinden. “IP Next Generation Overview.” May 14,1995. <http://sol.ibr.cs.tu-bs.de/~strauss/ipng/INET-IPng-Paper.html>. (4 November 1996).

Works Cited

Bound, Jim and Cini, Al. “IPv6 - Powering Up for Cyberspace.” September 3, 1996. <http://www.digital.com/info/ipv6/ipng.htm>. (1 November 1996).

Hinden, Robert M. “IP Next Generation Overview.” May 14,1995. <http://sol.ibr.cs.tu-bs.de/~strauss/ipng/INET-IPng-Paper.html>. (4 November 1996).

Tanenbaum, Andrew S. *Computer Networks*. 3rd Edition. Upper Saddle River, New Jersey: Prentice Hall. 1996.

Works Referenced

Atkinson, Ran. “IP Authentication Header.” August, 1995.<ftp://ds.internic.net/rfc/rfc1826.txt>. (4 November 1996).

Atkinson, Ran. “IP Encapsulating Security Payload (ESP).” August, 1995.<ftp://ds.internic.net/rfc/rfc1827.txt>. (4 November 1996).

Atkinson, Ran. “IP Authentication using Keyed MD5.” August, 1995. <ftp://ds.internic.net/rfc/rfc1828.txt>. (4 November 1996).

Atkinson, Ran. “Security Architecture for the Internet Protocol.” August, 1995. <ftp://ds.internic.net/rfc/rfc1825.txt>. (4 November 1996).

Bound, Jim and Cini, Al. “IPv6 - Powering Up for Cyberspace.” September 3, 1996. <http://www.digital.com/info/ipv6/ipng.htm>. (1 November 1996).

Hinden, Robert M. “IP Next Generation Overview.” May 14,1995. <http://sol.ibr.cs.tu-bs.de/~strauss/ipng/INET-IPng-Paper.html>. (4 November 1996).

Hinden, Robert M. and Deering, Steve. “Internet Protocol, Version 6 (IPv6) Specification.” December, 1995. <ftp://ds.internic.net/rfc/rfc1883.txt>. (4 November 1996).

Tanenbaum, Andrew S. *Computer Networks*. 3rd Edition. Upper Saddle River, New Jersey: Prentice Hall. 1996.