

A Large Corporate User's View of IPng

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Abstract

This document was submitted to the IETF IPng area in response to RFC 1550. Publication of this document does not imply acceptance by the IPng area of any ideas expressed within. Comments should be submitted to the big-internet@munnari.oz.au mailing list.

Disclaimer and Acknowledgments

Much of this draft has been adapted from the article "A User's View of IPng" by Eric Fleischman which was published in the September 1993 edition of *ConneXions Magazine* (Volume 7, Number 9, pages 36 - 40). The original *ConneXions* article represented an official position of The Boeing Company on IPng issues. This memo is an expansion of that original treatment. This version also represents a Boeing corporate opinion which we hope will be helpful to the on-going IPng discussions. An assumption of this paper is that other Fortune 100 companies which have non-computing-related products and services will tend to have a viewpoint about IPng which is similar to the one presented by this paper.

Executive Summary

Key points:

- 1) Large corporate users generally view IPng with disfavor.
- 2) Industry and the IETF community have very different values and viewpoints which lead to orthogonal assessments concerning the desirability of deploying IPng.
- 3) This paper provides insight into the mindset of a large corporate user concerning the relevant issues surrounding an IPng deployment. The bottom line is that a new deployment of IPng runs counter to several business drivers. A key point to

highlight is that end users actually buy applications -- not networking technologies.

- 4) There are really only two compelling reasons for a large end user to deploy IPng:
 - A) The existence of must-have products which are tightly coupled with IPng.
 - B) Receipt of a command to deploy IPng from senior management. The former would probably be a function of significant technological advances. The latter probably would be a function of a convergence of IPng with International Standards (OSI).
- 5) Five end user requirements for IPng are presented:
 - A) The IPng approach must permit piecemeal transitions.
 - B) The IPng approach must not hinder technological advances.
 - C) The IPng approach is expected to foster synergy with International Standards (OSI).
 - D) The IPng approach should have "Plug and Play" networking capabilities.
 - E) The IPng approach must have network security characteristics which are better than existing IPv4 protocols.

Introduction

The goal of this paper is to examine the implications of IPng from the point of view of Fortune 100 corporations which have heavily invested in TCP/IP technology in order to achieve their (non-computer related) business goals.

It is our perspective that End Users currently view IPng with disfavor. This note seeks to explain some of the reasons why an end user's viewpoint may differ significantly from a "traditional IETF" perspective. It addresses some of the reasons which cause IPng to be viewed by end users as a "threat" rather than as an "opportunity". It enumerates some existing End User dissatisfactions with IPv4 (i.e., current TCP/IP network layer). These dissatisfactions may perhaps be eventually exploited to "sell" IPng to users. Finally, it identifies the most compelling reasons for end users to deploy IPng. In any case, the IETF community should be warned that their own enthusiasm for IPng is generally not shared by end users and that convincing end users to deploy IPng technologies may be very difficult -- assuming it can be done at all.

The Internet and TCP/IP Protocols are not Identical

The Internet Engineering Task Force (IETF) community closely associates TCP/IP protocols with the Internet. In many cases it is difficult to discern from the IETF perspective where the world-wide Internet infrastructure ends and the services of the TCP/IP Protocol Suite begin -- they are not always distinguishable from each other. Historically they both stem from the same roots: DARPA was the creator of TCP/IP and of the seminal "Internet". The services provided by the Internet have been generally realized by the "TCP/IP protocol family". The Internet has, in turn, become a primary vehicle for the definition, development, and transmission of the various TCP/IP protocols in their various stages of maturity. Thus, the IETF community has a mindset which assumes that there is a strong symbiotic relationship between the two.

End users do not share this assumption -- despite the fact that many end users have widely deployed TCP/IP protocols and extensively use the Internet. It is important for the IETF community to realize, however, that TCP/IP protocols and the Internet are generally viewed to be two quite dissimilar things by the large end user. That is, while the Internet may be a partial selling point for some TCP/IP purchases, it is rarely even a primary motivation for the majority of purchases. Many end users, in fact, have sizable TCP/IP deployments with no Internet connectivity at all. Thus, many end users view the relationship between the Internet and TCP/IP protocols to be tenuous at best.

More importantly, many corporations have made substantial investments in (non-Internet) external communications infrastructures. A variety of reasons account for this including the fact that until recently the Internet was excluded from the bilateral agreements and international tariffs necessary for international commerce. In any case, end users today are not (in the general case) dependent upon the Internet to support their business processes. [Note: the previous sentence does not deny that many Fortune 100 employees (such as the author) are directly dependent upon the Internet to fulfill their job responsibilities: The Internet has become an invaluable tool for many corporations' "research and education" activities. However, it is rarely used today for activities which directly affect the corporations' financial "bottom line": commerce.] By contrast, large End Users with extensive internal TCP/IP deployments may perhaps view TCP/IP technology to be critically important to their corporation's core business processes.

Security Islands

Another core philosophical difference between large end users and the IETF is concerning the importance of Security Islands (i.e., firewalls). The prevalent IETF perspective is that Security Islands are "A Bad Thing". The basic IETF assumption is that the applications they are designing are universally needed and that Security Islands provide undesirable filters for that usage. That is, the IETF generally has a world view which presupposes that data access should be unrestricted and widely available.

By contrast, corporations generally regard data as being a "sensitive" corporate asset: If compromised the very viability of the corporation itself may in some cases be at risk. Corporations therefore presuppose that data exchange should be restricted.

Large end users also tend to believe that their employees have differing data access needs: Factory workers have different computing needs than accountants who have different needs than aeronautical engineers who have different needs than research scientists. A corporation's networking department(s) seeks to ensure that each class of employee actually receives the type of services they require. A security island is one of the mechanisms by which the appropriate service levels may be provided to the appropriate class of employee, particularly in regards to external access capabilities.

More importantly, there are differing classes of computer resources within a corporation. A certain percentage of these resources are absolutely critical to the continuing viability of that corporation. These systems should never (ever) be accessible from outside of the company. These "corporate jewels" must be protected by viable security mechanisms. Security islands are one very important component within a much larger total security solution.

For these reasons we concur with the observation made by Yakov Rekhter (of IBM) and Bob Moskowitz (of Chrysler) in their joint electronic mail message of January 28, 1994. They wrote:

"Hosts within sites that use IP can be partitioned into three categories:

- hosts that do not require Internet access.
- hosts that need access to a limited set of Internet services (e.g., Email, FTP, netnews, remote login) which can be handled by application layer relays.
- hosts that need unlimited access (provided via IP connectivity) to the Internet."

The exact mechanism by which a corporation will satisfy the differing needs of these three classes of devices must be independently determined by that corporation based upon a number of internal factors. Each independent solution will determine how that corporation defines their own version of "security island".

Thus, if end users use the Internet at all, they will generally do so through a "security island" of their own devising. The existence of the security island is yet another element to (physically and emotionally) decouple the End User from the Internet. That is, while the end user may use the Internet, their networks (in the general case) are neither directly attached to it nor are their core business processes today critically dependent upon it.

Networking from a Large End User's Perspective

The following five key characteristics describe Boeing's environment and are probably generally representative of other large TCP/IP deployments. The author believes that an understanding of these characteristics is very important for obtaining insight into how the large end user is likely to view IPng.

1) Host Ratio

Many corporations explicitly try to limit the number of their TCP/IP hosts that are directly accessible from the Internet. This is done for a variety of reasons (e.g., security). While the ratio of those hosts that have direct Internet access capabilities to those hosts without such capabilities will vary from company to company, ratios ranging from 1:1000 to 1:10,000 (or more) are not uncommon. The implication of this point is that the state of the world-wide (IPv4) Internet address space only directly impacts a tiny percentage of the currently deployed TCP/IP hosts within a large corporation. This is true even if the entire population is currently using Internet-assigned addresses.

2) Router-to-Host Ratio

Most corporations have significantly more TCP/IP hosts than they have IP routers. Ratios ranging between 100:1 to 600:1 (or more) are common. The implication of this point is that a transition approach which solely demands changes to routers is generally much less disruptive to a corporation than an approach which demands changes to both routers and hosts.

3) Business Factor

Large corporations exist to fulfill some business purpose such as the construction of airplanes, baseball bats, cars, or some other product or service offering. Computing is an essential tool to help automate business processes in order to more efficiently accomplish the business goals of the corporation. Automation is accomplished via applications. Data communications, operating systems, and computer hardware are the tools used by applications to accomplish their goals. Thus, users actually buy applications and not networking technologies. The central lesson of this point is that IPng will be deployed according to the applications which use it and not because it is a better technology.

4) Integration Factor

Large corporations currently support many diverse computing environments. This diversity limits the effectiveness of a corporation's computing assets by hindering data sharing, application interoperability, "application portability", and software re-usability. The net effect is stunted application life cycles and increased support costs. Data communications is but one of the domains which contribute towards this diversity. For example, The Boeing Company currently has deployed at least sixteen different protocol families within its networks (e.g., TCP/IP, SNA, DECnet, OSI, IPX/SPX, AppleTalk, XNS, etc.). Each distinct Protocol Family population potentially implies unique training, administrative, support, and infrastructure requirements. Consequently, corporate goals often exist to eliminate or merge diverse Data Communications Protocol Family deployments in order to reduce network support costs and to increase the number of devices which can communicate together (i.e., foster interoperability). This results in a basic abhorrence to the possibility of introducing "Yet Another Protocol" (YAP). Consequently, an IPng solution which introduces an entirely new set of protocols will be negatively viewed simply because its by-products are more roadblocks to interoperability coupled with more work, expense, and risk to support the end users' computing resources and business goals. Having said this, it should be observed that this abhorrence may be partially overcome by "extenuating circumstances" such as applications using IPng which meet critical end-user requirements or by broad (international) commercial support.

5) Inertia Factor

There is a natural tendency to continue to use the current IP protocol (IPv4) regardless of the state of the Internet's IPv4 address space. Motivations supporting inertia include the following: existing application dependencies (including Application Programming Interface (API) dependencies); opposition to additional protocol complexity; budgetary constraints limiting additional hardware/software expenses; additional address management and naming service costs; transition costs; support costs; training costs; etc. As the number of Boeing's deployed TCP/IP hosts continues to grow towards the 100,000 mark, the inertial power of this population becomes increasingly strong. However, inertia even exists with smaller populations simply because the cost to convert or upgrade the systems are not warranted. Consequently, pockets of older "legacy system" technologies often exist in specific environments (e.g., we still have pockets of the archaic BSC protocol). The significance of this point is that unless there are significant business benefits to justify an IPng deployment, economics will oppose such a deployment. Thus, even if the forthcoming IPng protocol proves to be "the ultimate and perfect protocol", it is unrealistic to imagine that the entire IPv4 population will ever transition to IPng. This means that should we deploy IPng within our network, there will be an ongoing requirement for our internal IPng deployment to be able to communicate with our internal IPv4 community. This requirement is unlikely to go away with time.

Address Depletion Doesn't Resonate With Users

Thus, the central, bottom-line question concerning IPng from the large corporate user perspective is: What are the benefits which will justify the expense of deploying IPng?

At this time we can conceive of only four possible causes which may motivate us to consider deploying IPng:

Possible Cause:	Possible Corporate Response:
1) Many Remote (external) Peers solely use IPng.	Gateway external systems only.
2) Internet requires IPng usage.	Gateway external systems only.
3) "Must have" products are tightly coupled with IPng (e.g., "flows" for real-time applications).	Upgrade internal corporate network to support IPng where that functionality is needed.

- 4) Senior management directs IPng usage. Respond appropriately.

It should explicitly be noted that the reasons which are compelling the Internet Community to create IPng (i.e., the scalability of IPv4 over the Internet) are not themselves adequate motivations for users to deploy IPng within their own private networks. That is, should IPng usage become mandated as a prerequisite for Internet usage, a probable response to this mandate would be to convert our few hosts with external access capabilities to become IPng-to-IPv4 application-layer gateways. This would leave the remainder of our vast internal TCP/IP deployment unchanged. Consequently, given gateways for external access, there may be little motivation for a company's internal network to support IPng.

User's IPv4 "Itches" Needing Scratching

The end user's "loyalty" to IPv4 should not be interpreted to mean that everything is necessarily "perfect" with existing TCP/IP deployments and that there are therefore no "itches" which an improved IPv4 network layer -- or an IPng -- can't "scratch". The purpose of this section is to address some of the issues which are very troubling to many end users:

- A) Security. TCP/IP protocols are commonly deployed upon broadcast media (e.g., Ethernet Version 2). However, TCP/IP mechanisms to encrypt passwords or data which traverse this media are inadequate. This is a very serious matter which needs to be expeditiously resolved. An integrated and effective TCP/IP security architecture needs to be defined and become widely implemented across all vendors' TCP/IP products.
- B) User Address Space privacy. Current IPv4 network addressing policies require that end users go to external entities to obtain IP network numbers for use in their own internal networks. These external entities have the hubris to determine whether these network requests are "valid" or not. It is our belief that a corporation's internal addressing policies are their own private affair -- except in the specific instances in which they may affect others. Consequently, a real need exists for two classes of IPv4 network numbers: those which are (theoretically) visible to the Internet today (and thus are subject to external requirements) and those which will never be connected to the Internet (and thus are strictly private). We believe that the concept of "local addresses" is a viable compromise between the justifiable need of the Internet to steward scarce global resources and the corporate need for privacy. "Local addresses" by definition are non-globally-unique addresses which should

never be routed (or seen) by the Internet infrastructure.

We believe that 16 contiguous Class B "local addresses" need to immediately be made available for internal corporate usage. Such an availability may also reduce the long-term demand for new IPv4 network numbers (see RFC 1597).

- C) Self-Defining Networks. Large End Users have a pressing need for plug-and-play TCP/IP networks which auto-configure, auto-address, and auto-register. End users have repeatedly demonstrated our inability to make the current manual methods work (i.e., heavy penalties for human error). We believe that the existing DHCP technology is a good beginning in this direction.
- D) APIs and network integration. End users have deployed many differing complex protocol families. We need tools by which these diverse deployments may become integrated together along with viable transition tools to migrate proprietary alternatives to TCP/IP-based solutions. We also desire products to use "open" multi-vendor, multi-platform, exposed Application Programming Interfaces (APIs) which are supported across several data communications protocol "families" to aid in this integration effort.
- E) International Commerce. End users are generally unsure as to what extent TCP/IP can be universally used for international commerce today and whether this is a cost-effective and "safe" option to satisfy our business requirements.
- F) Technological Advances. We have ongoing application needs which demand a continual "pushing" of the existing technology. Among these needs are viable (e.g., integratable into our current infrastructures) solutions to the following: mobile hosts, multimedia applications, real-time applications, very high-bandwidth applications, improved very low-bandwidth (e.g., radio based) applications, standard-TCP/IP-based transaction processing applications (e.g., multi-vendor distributed databases).

Only Two Motivations For Users To Deploy IPng

Despite this list of IPv4 problem areas, we suspect that there are only two causes which may motivate users to widely deploy IPng:

- (1) If IPng products add critical functionality which IPv4 can't provide (e.g., real time applications, multimedia applications, genuine (scalable) plug-and-play networking, etc.), users would be motivated to deploy IPng where that functionality is needed.

However, these deployments must combat the "Integration Factor" and the "Inertia Factor" forces which have previously been described. This implies that there must be a significant business gain to justify such a deployment. While it is impossible to predict exactly how this conflict would "play out", it is reasonable to assume that IPng would probably be deployed according to an "as needed only" policy. Optimally, specific steps would be taken to protect the remainder of the network from the impact of these localized changes. Of course, should IPng become bundled with "killer applications" (i.e., applications which are extremely important to significantly many key business processes) then all bets are off: IPng will become widely deployed. However, it also should be recognized that virtually all (initial) IPng applications, unless they happen to be "killer applications", will have to overcome significant hurdles to be deployed simply because they represent risk and substantially increased deployment and support costs for the end user.

(2) Should IPng foster a convergence between Internet Standards and International Standards (i.e., OSI), this convergence could change IPng's destiny. That is, the networks of many large corporations are currently being driven by sets of strong, but contradictory, requirements: one set demanding compliance with Internet Standards (i.e., TCP/IP) and another set demanding compliance with International Standards. This paper assumes that the reader is already familiar with the many reasons why end users seek to deploy and use Internet Standards. The following is a partial list as to why End Users may be motivated to use International Standards (i.e., OSI) as well:

- A) World-wide commerce is regulated by governments in accordance with their treaties and legal agreements. World-wide telecommunications are regulated by the ITU (a United Nations chartered/authorized organization). International Standards (i.e., OSI) are the only government-sanctioned method for commercial data communications. Aspects of this picture are currently in the process of changing.

- B) The currently proprietary aeronautical world-wide air-to-ground and ground-to-ground communications are being replaced by an OSI-based (CLNP) Aeronautical Telecommunications Network (ATN) internet which is being built in a number of different national and international forums including:
 - * International Civil Aviation Organization (ICAO)
 - * International Air Transport Association (IATA)
 - * Airlines Electronic Engineering Committee (AEEC)

"Civil Aviation Authorities, airlines, and private aircraft will use the ATN to convey two major categories of data traffic among their computers: Air Traffic Services Communications (ATSC) and Aeronautical Industry Services Communication (AISC)." [Note: The data communications of airline passengers are not addressed by the directive.]

- C) A corporation's customers may have data communications requirements which are levied upon them by the governments in which they operate which they, in turn, must support in their own products in order to fulfill their customers' needs. For example, Boeing is influenced by existing:
- * Computer Aided Logistics Support (CALs; i.e., these are GOSIP (OSI)-based) requirements for US Department of Defense contractors.
 - * Airline requirements emanating from A and B above.
- D) The end user perception that once we have deployed International Standards we will not subsequently be compelled to migrate by external factors to another technology. Thus, we would have a "safe" foundation to concentrate upon our real computing issues such as increased customer satisfaction, business process flow-time improvements, legacy system modernization, and cost avoidance.
- E) The proposals of entities desiring to obtain contracts with Governments are evaluated on many subjective and objective bases. One of the subjective issues may well be the "responsibility" and "dependability" of the bidder company including such intangibles as its corporate like-mindedness. For this reason, as long as the Government has OSI as their official standard, the bidder may have a subjective advantage if its corporate policy also includes a similar standard, particularly if data communications services are being negotiated.
- F) The perception that the need for IPng may imply that IPv4 is unfit to be a strategic end user alternative. Also, IPng is not a viable deployment option at this time.
- G) Doubts concerning IPv4 scalability (e.g., toasternet: an algorithmic change in which currently "dumb devices" become intelligent and suddenly require Internet connectivity).

It currently appears that many of these "OSI motivations" are undergoing change at this time. This possibility must be tracked with interest. However, a key point of this section is that a

corporation must base its data communications decisions upon business requirements. That is, corporations exist to sell products and services, not to play "networking games".

Thus, if a means could be found to achieve greater synergy (integration/ adoption) between Internet Standards and International Standards then corporate management may be inclined to mandate internal deployment of the merged standards and promote their external use. Optimally, such a synergy should offer the promise of reducing currently deployed protocol diversity (i.e., supports the "Integration Factor" force). Depending on the specific method by which this convergence is achieved, it may also partially offset the previously mentioned "Inertia Factor" force, especially if IPng proves to be a protocol which has already been deployed.

User-based IPng Requirements

From the above one can see that a mandate to use IPng to communicate over the Internet does not correspondingly imply the need for large corporate networks to generally support IPng within their networks. Thus, while the IPv4 scalability limitations are a compelling reason to identify a specific IPv4 replacement protocol for the Internet, other factors are at work within private corporate networks. These factors imply that large TCP/IP end users will have a continuing need to purchase IPv4 products even after IPng products have become generally available.

However, since the IETF community is actively engaged in identifying an IPng solution, it is desirable that the solution satisfy as many end user needs as possible. For this reason, we would like to suggest that the following are important "user requirements" for any IPng solution:

- 1) The IPng approach must permit users to slowly transition to IPng in a piecemeal fashion. Even if IPng becomes widely deployed, it is unrealistic to expect that users will ever transition all of the extensive IPv4 installed base to IPng. Consequently, the approach must indefinitely support corporate-internal communication between IPng hosts and IPv4 hosts regardless of the requirements of the world-wide Internet.
- 2) The IPng approach must not hinder technological advances from being implemented.
- 3) The IPng approach is expected to eventually foster greater synergy (integration/adoption) between Internet Standards and International Standards (i.e., OSI). [Note: This may be accomplished in a variety of ways including having the Internet

Standards adopted as International Standards or else having the International Standards adopted as Internet Standards.]

- 4) The IPng approach should have "self-defining network" (i.e., "plug & play") capabilities. That is, large installations require device portability in which one may readily move devices within one's corporate network and have them autoconfigure, autoaddress, autoregister, etc. without explicit human administrative overhead at the new location -- assuming that the security criteria of the new location have been met.
- 5) The approach must have network security characteristics which are better than existing IPv4 protocols.

Conclusion

In summary, the key factor which will determine whether -- and to what extent -- IPng will be deployed by large end users is whether IPng will become an essential element for the construction of applications which are critically needed by our businesses. If IPng is bundled with applications which satisfy critical business needs, it will be deployed. If it isn't, it is of little relevance to the large end user. Regardless of what happens to IPng, the large mass of IPv4 devices will ensure that IPv4 will remain an important protocol for the foreseeable future and that continued development of IPv4 products is advisable.

Security Considerations

Security issues discussed throughout this memo.

Author's Address

Eric Fleischman
Network Architect
Boeing Computer Services
P.O. Box 24346, MS 7M-HA
Seattle, WA 98124-0346 USA

E-Mail: ericf@atc.boeing.com