

IPng Requirements of Large Corporate Networks

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. This draft summarizes some of the requirements of large corporate networks for the next generation of the Internet protocol suite.

Executive Overview

As more and more corporations are using TCP/IP for their mission-critical applications, they are bringing additional requirements, summarized below, the satisfaction of which would make TCP/IP even more appealing to businesses. Since these are requirements rather than solutions, we include capabilities that might be provided in protocol layers other than the one that IPv4 occupies; i.e., these items might lie outside the scope typically envisioned for IPng, but we'll refer to them as IPng requirements nonetheless. When we mention potential solutions, it is not to suggest that they are the best approach, but merely to clarify the requirement.

Among business users the major requirements we see for IPng are:

- smooth migration from, and coexistence with, IPv4;
- predictable levels of service for predictable costs;
- security; and
- accommodation of multiple protocols suites.

We also mention several more specific requirements.

IPng must have a viable strategy for migration from, and coexistence with, IPv4. IPv4 and IPng must coexist well, because they will need to do so for several years. To encourage IPv4 users to upgrade to

IPng, IPng must offer compelling advantages and an easy migration path.

Corporate networks must meet promised levels of service while controlling costs through efficient use of resources. The IETF should consider both technical solutions (such as service classes and priorities) and administrative ones (such as accounting) to promote economy.

Many businesses will not connect to a network until they are confident that it will not significantly threaten the confidentiality, integrity, or availability of their data.

Corporations tend to use multiple protocols. Numerous forces stymie the desire to settle on just one protocol for a large corporation: diverse installed bases, skills, technical factors, and the general trend toward corporate decentralization. The IETF needs a strategy for heterogeneity flexible enough to accommodate the principal multiprotocol techniques, including multiprotocol transport, tunneling, and link sharing.

Some of these requirements might be satisfied by more extensive deployment of existing Internet architectures (e.g., Generic Security Service and IPv4 type of service). The current Internet protocols could be enhanced to satisfy most of the remaining requirements of commercial users while retaining IPv4. Nevertheless, some corporations will be scared away from TCP/IP by the publicity about the address space until the IETF sets a direction for its expansion.

Migration and Coexistence

As the use of IPv4 continues to grow, the day may come when no more IPv4 network addresses will be left, and no additional networks will be able to connect to the Internet. Classless Inter-Domain Routing (CIDR, RFC 1519) and careful gleaning of the address space will postpone that cutoff for several years. The hundreds of millions of people on networks that do get IPv4 addresses won't be affected directly by the exhaustion of the address space, but they will miss the opportunity to communicate with those less lucky.

Because the Internet is too large for all its users to cutover to IPng quickly, IPng must coexist well with IPv4. Furthermore, IPv4 users won't upgrade to IPng without a compelling reason. Access to new services will not be a strong motivation, since new services will want to support both the IPng users and the IPv4 users. Only services that cannot exist on IPv4 will be willing to use IPng exclusively. Moreover, if IPng requires more resources (e.g., storage, memory, or administrative complexity) than IPv4, users will

not install IPng unless it has clear benefits over IPv4. Indeed, the millions of users of low-end systems (DOS, sub-notebooks) might not ever be able to use IPng if it takes more memory. Thus there will be a long period of coexistence between IPng and IPv4, so the coexistence needs to be quite painless, and not based on any assumption that IPv4 use will diminish quickly.

Service Level Agreements

If a corporation depends on its network for applications that are critical to its business (such as airlines do for reservations, and brokerages do for stock and bond trades), then the corporation insists that the network provide the needed service level for a predictable cost, so they can allow for it in their budget ahead of time. A service level agreement (SLA) is a contract between network's provider and users that defines the service level which a user will see and the cost associated with that level of service. Measurements in an SLA may include response times (average and maximum), availability percentages, number of active sessions, throughput rates, etc.. Businesses need to be able to predict and guarantee the service levels and costs (routing capacity, link bandwidth, etc.) for their traffic patterns on a TCP/IP network.

IPng should allow control of the cost of networking, a major concern for corporations. Teleprocessing lines are a significant cost in corporate networks. Although the cost per bit-per-second tends to be lower on higher-bandwidth links, high-bandwidth links can be hard to get, particularly in emerging nations. In many places it is difficult to acquire a 64 kpbs line, and T1 service might not exist. Furthermore, lead times can be over six months. Even in the US the cost of transcontinental T1 service is high enough to encourage high utilization. Cost-conscious businesses want IPng to allow high utilization of teleprocessing links, but without requiring excessive processing power to achieve the high utilization. There has been considerable speculation concerning the goodput through congested routes when using the Internet's current congestion control algorithms; instead, it should be measured in a range of realistic cases. If peak-busy-hour goodput under congestion is near the theoretical maximum, publicize the data and move on to other requirements. If not, then the IETF should seek a better standard (e.g., they might explore XTP's adaptive rate-based approach and other proposals).

Functions, such as class of service and priority, that let an enterprise control use of bandwidth also may help meet service level agreements. On the one hand, it has been said that the absence of these inhibits TCP/IP usage in corporate networks, especially when predictable interactive response times are required. On the other

hand, few vendors have felt motivated to implement TCP's architected type-of-service, and priority tends to be handled in a non-standard way (e.g., to assure that interactive well-known ports, such as Telnet, get faster response times than non-interactive well-known ports, such as file transfer). The IETF should sort out these apparently conflicting perspectives. If the ad hoc techniques can be demonstrated to be adequate, then they should be standardized; otherwise, effective techniques should be developed and standardized.

Commercial users often require the options of a higher level of service for a higher cost, or a lower level of service for a lower cost; e.g., some businesses pay top dollar to assure fast response time during business hours, but choose less expensive satellite services for data backup during the night. Pervasive use of IPv4's type-of-service markings might satisfy this requirement.

To discourage waste of bandwidth and other expensive resources, corporations want to account for their use. Direct cost recovery would let an entity measure and benchmark its efficiency with minimal economic distortion. Alternatives, such as placing these costs into corporate overhead or charging per connection, make sense when the administrative cost of implementing usage-based accounting is high enough to introduce more economic distortion than the alternatives would. For example, connection-based costs alone may be adequate for a resource (such as LAN bandwidth) that is not scarce or expensive, but a combination of a connection cost and a usage cost may be more appropriate for a more scarce or expensive resource (such as WAN bandwidth). Balance must be maintained between the overhead of accounting and the granularity of cost allocation.

Security

Many corporations will stick with their private networks until public ones can guarantee equivalent confidentiality, integrity, and availability. It is not clear that additional architecture is needed to satisfy this requirement; perhaps more wide spread use of existing security technology would suffice. For example, the Internet could encourage wide deployment of Generic Security Service, and then solicit feedback on whether additional security requirements need to be satisfied. Note that businesses are so concerned about network cost control mechanisms that they want them secured against tampering. IPng should not interfere with firewalls, which many corporations consider essential.

Heterogeneity

Corporate users want the Internet to accommodate multiple protocol suites. Several different protocol suites are growing in use, and some older ones will be used for many more years. Although many people wish there were only one protocol in the world, there is little agreement on which one it should be.

Since the marketplace has not settled on one approach to handling multiple protocols, IPng should be flexible enough to accommodate a variety of technical approaches to achieving heterogeneity. For example, most networking protocols assume they will be the dominate protocol that transports all others; protocol designers should pay more attention to making their protocols easily transported by others. IPng needs to be flexible enough to accommodate the major multiprotocol trends, including multiprotocol transport networking (for an example, see X/OPEN document G306), tunneling (both IP being the tunnel and being tunneled), and link sharing (e.g., point-to-point protocol and frame relay). Fair sharing of bandwidth by protocols with different congestion control mechanisms is a particularly interesting subject.

Flow and Resource Reservation

Corporate users are becoming more interested in transmitting both non-isochronous and isochronous information together across the same link. IPng should coexist effectively with the isochronous protocols being developed for the Internet.

The Internet protocols should take advantage of services that may be offered by an underlying fast packet switching service. Constant-bit-rate and variable-bit-rate services typically require specification of, and conformance to, traffic descriptors and specification of quality-of-service objectives from applications or users. The Internet's isochronous protocols should provide mechanisms to take advantage of multimedia services that will be offered by fast packet switching networks, and must ensure that quality-of-service guarantees are preserved all the way up the protocol stacks to the applications. Protocols using available-bit-rate services may achieve better bandwidth utilization if they react to congestion messages from a fast packet switching network, and if they consider consequences of cell discard (e.g., if one cell of an IP datagram is discarded, it would be a waste to continue forwarding the rest of the cells in that datagram; also, selective retransmit should be revisited in this context).

When the Internet protocol suite allows mixing of non-isochronous and isochronous traffic on one medium, it should provide mechanisms to

discourage inappropriate reservation of resources; e.g., a Telnet connection probably doesn't need to reserve 45Mbps. Accounting, class-of-service, and well-known-port distinctions are possible ways to satisfy that requirement.

Mobile Hosts

Wireless technology opens up opportunities for new TCP/IP applications that are specific to mobile hosts. In addition to coordinating with organizations developing wireless standards, the IETF also should encourage the specification of new TCP/IP applications enabled by wireless, such as connectionless messaging.

IPng should deal well with the characteristics (delay, error rates, etc.) peculiar to wireless.

Topological flexibility

Today a TCP/IP host moved to a different subnet needs a new IP address. Such moves and changes can become a significant administrative cost. Moreover, mobile hosts require flexible topology. Note how the wireless world is trying to defeat the subnet model of addressing either by proxy or by IPaddress servers. Perhaps IPng needs an addressing model more flexible than subnetting, both to reduce the administrative burden and to facilitate roaming users.

The need to eliminate single points of failure drives the business requirement for multi-tail attachment of hosts to networks. Corporate users complain that TCP/IP can non-disruptively switch a connection from a broken route to a working one only if the new route leads to the same adapter on the end system.

Configuration, Administration and Operation

Businesses would like dynamic but secure updating of Domain Name Servers, both to ease moves and changes and to facilitate cutover to backup hosts. In this vein, secure and dynamic interaction between DNS and Dynamic Host Configuration Protocol (DHCP, RFC 1541) is required. The IETF should encourage wide deployment of DHCP, and then solicit feedback on whether additional configuration requirements need to be satisfied.

Policy-Based Routing

Policy-based routing is a more a solution than a requirement. Businesses rarely require a general purpose policy architecture, although they do state requirements that policy-based routing could satisfy. For example, corporations do not want to carry for free the

transit traffic of other enterprises, and they may not want their sensitive data to flow through links controlled by certain other enterprises. Policy-based routing is one possible way to satisfy those requirements, but there seems to be a concern that general purpose policy-based routing may have high administrative cost and low routing performance.

Scaling

If IPng satisfies the scaling requirement of the Internet, then it satisfies it for corporate networks a fortiori.

Conclusions

Enhancements to the Internet protocol suite, together with wider deployment of some of its existing architectures, could satisfy these requirement of commercial customers while retaining IPv4. Expansion of the address space eventually will be necessary to allow continued Internet growth, but in RFC 1518 Tony Li and Yakov Rehkter have shown that from a technical perspective the addressing issue of IPng is not an immediate concern.

Nevertheless, the TCP/IP community should establish a direction for enlargement of the address space, because unfounded publicity about the address space is scaring away potential TCP/IP users. If the IETF does not provide direction on how its address space will grow, then people may use non-standard, and probably incompatible, approaches.

Security Considerations

The IETF should encourage wide deployment of GSS API, and then solicit feedback on whether additional security requirements need to be satisfied. Businesses are so concerned about network cost control mechanisms that they want them secured against tampering. IPng should not interfere with firewalls, which many corporations consider essential. See other comments on Security throughout this memo.

Authors' Addresses

Edward Britton
IBM Corp.
E69/503
P.O.Box 12195
Research Triangle Park, NC 27709

Phone: (919) 254-6037
EMail: brittone@vnet.ibm.com

John Tavs
IBM Corp.
E69/503
P.O.Box 12195
Research Triangle Park, NC 27709

Phone: (919) 245-7610
EMail: tavs@vnet.ibm.com

