

Internet Transparency

Status of this Memo

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

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This document describes the current state of the Internet from the architectural viewpoint, concentrating on issues of end-to-end connectivity and transparency. It concludes with a summary of some major architectural alternatives facing the Internet network layer.

This document was used as input to the IAB workshop on the future of the network layer held in July 1999. For this reason, it does not claim to be complete and definitive, and it refrains from making recommendations.

Table of Contents

1. Introduction.....	2
2. Aspects of end-to-end connectivity.....	3
2.1 The end-to-end argument.....	3
2.2 End-to-end performance.....	4
2.3 End-to-end address transparency.....	4
3. Multiple causes of loss of transparency.....	5
3.1 The Intranet model.....	6
3.2 Dynamic address allocation.....	6
3.2.1 SLIP and PPP.....	6
3.2.2 DHCP.....	6
3.3 Firewalls.....	6
3.3.1 Basic firewalls.....	6
3.3.2 SOCKS.....	7
3.4 Private addresses.....	7
3.5 Network address translators.....	7
3.6 Application level gateways, relays, proxies, and caches.....	8
3.7 Voluntary isolation and peer networks.....	8

3.8 Split DNS.....	9
3.9 Various load-sharing tricks.....	9
4. Summary of current status and impact.....	9
5. Possible future directions.....	11
5.1 Successful migration to IPv6.....	11
5.2 Complete failure of IPv6.....	12
5.2.1 Re-allocating the IPv4 address space.....	12
5.2.2 Exhaustion.....	13
5.3 Partial deployment of IPv6.....	13
6. Conclusion.....	13
7. Security Considerations.....	13
Acknowledgements.....	14
References.....	14
Author's Address.....	17
Full Copyright Statement.....	18

1. Introduction

"There's a freedom about the Internet: As long as we accept the rules of sending packets around, we can send packets containing anything to anywhere." [Berners-Lee]

The Internet is experiencing growing pains which are often referred to as "the end-to-end problem". This document attempts to analyse those growing pains by reviewing the current state of the network layer, especially its progressive loss of transparency. For the purposes of this document, "transparency" refers to the original Internet concept of a single universal logical addressing scheme, and the mechanisms by which packets may flow from source to destination essentially unaltered.

The causes of this loss of transparency are partly artefacts of parsimonious allocation of the limited address space available to IPv4, and partly the result of broader issues resulting from the widespread use of TCP/IP technology by businesses and consumers. For example, network address translation is an artefact, but Intranets are not.

Thus the way forward must recognise the fundamental changes in the usage of TCP/IP that are driving current Internet growth. In one scenario, a complete migration to IPv6 potentially allows the restoration of global address transparency, but without removing firewalls and proxies from the picture. At the other extreme, a total failure of IPv6 leads to complete fragmentation of the network layer, with global connectivity depending on endless patchwork.

This document does not discuss the routing implications of address space, nor the implications of quality of service management on router state, although both these matters interact with transparency to some extent. It also does not substantively discuss namespace issues.

2. Aspects of end-to-end connectivity

The phrase "end to end", often abbreviated as "e2e", is widely used in architectural discussions of the Internet. For the purposes of this paper, we first present three distinct aspects of end-to-endness.

2.1 The end-to-end argument

This is an argument first described in [Saltzer] and reviewed in [RFC 1958], from which an extended quotation follows:

"The basic argument is that, as a first principle, certain required end-to-end functions can only be performed correctly by the end-systems themselves. A specific case is that any network, however carefully designed, will be subject to failures of transmission at some statistically determined rate. The best way to cope with this is to accept it, and give responsibility for the integrity of communication to the end systems. Another specific case is end-to-end security.

"To quote from [Saltzer], 'The function in question can completely and correctly be implemented only with the knowledge and help of the application standing at the endpoints of the communication system. Therefore, providing that questioned function as a feature of the communication system itself is not possible. (Sometimes an incomplete version of the function provided by the communication system may be useful as a performance enhancement.)'

"This principle has important consequences if we require applications to survive partial network failures. An end-to-end protocol design should not rely on the maintenance of state (i.e. information about the state of the end-to-end communication) inside the network. Such state should be maintained only in the endpoints, in such a way that the state can only be destroyed when the endpoint itself breaks (known as fate-sharing). An immediate consequence of this is that datagrams are better than classical virtual circuits. The network's job is to transmit datagrams as efficiently and flexibly as possible. Everything else should be done at the fringes."

Thus this first aspect of end-to-endness limits what the network is expected to do, and clarifies what the end-system is expected to do. The end-to-end argument underlies the rest of this document.

2.2 End-to-end performance

Another aspect, in which the behaviour of the network and that of the end-systems interact in a complex way, is performance, in a generalised sense. This is not a primary focus of the present document, but it is mentioned briefly since it is often referred to when discussing end-to-end issues.

Much work has been done over many years to improve and optimise the performance of TCP. Interestingly, this has led to comparatively minor changes to TCP itself; [STD 7] is still valid apart from minor additions [RFC 1323, RFC 2581, RFC 2018]. However a great deal of knowledge about good practice in TCP implementations has built up, and the queuing and discard mechanisms in routers have been fine-tuned to improve system performance in congested conditions.

Unfortunately all this experience in TCP performance does not help with transport protocols that do not exhibit TCP-like response to congestion [RFC 2309]. Also, the requirement for specified quality of service for different applications and/or customers has led to much new development, especially the Integrated Services [RFC 1633, RFC 2210] and Differentiated Services [RFC 2475] models. At the same time new transport-related protocols have appeared [RFC 1889, RFC 2326] or are in discussion in the IETF. It should also be noted that since the speed of light is not set by an IETF standard, our current notions of end-to-end performance will be largely irrelevant to interplanetary networking.

Thus, despite the fact that performance and congestion issues for TCP are now quite well understood, the arrival of QOS mechanisms and of new transport protocols raise new questions about end-to-end performance, but these are not further discussed here.

2.3 End-to-end address transparency

When the catenet concept (a network of networks) was first described by Cerf in 1978 [IEN 48] following an earlier suggestion by Pouzin in 1974 [CATENET], a clear assumption was that a single logical address space would cover the whole catenet (or Internet as we now know it). This applied not only to the early TCP/IP Internet, but also to the Xerox PUP design, the OSI connectionless network design, XNS, and numerous other proprietary network architectures.

This concept had two clear consequences - packets could flow essentially unaltered throughout the network, and their source and destination addresses could be used as unique labels for the end systems.

The first of these consequences is not absolute. In practice changes can be made to packets in transit. Some of these are reversible at the destination (such as fragmentation and compression). Others may be irreversible (such as changing type of service bits or decrementing a hop limit), but do not seriously obstruct the end-to-end principle of Section 2.1. However, any change made to a packet in transit that requires per-flow state information to be kept at an intermediate point would violate the fate-sharing aspect of the end-to-end principle.

The second consequence, using addresses as unique labels, was in a sense a side-effect of the catenet concept. However, it was a side-effect that came to be highly significant. The uniqueness and durability of addresses have been exploited in many ways, in particular by incorporating them in transport identifiers. Thus they have been built into transport checksums, cryptographic signatures, Web documents, and proprietary software licence servers. [RFC 2101] explores this topic in some detail. Its main conclusion is that IPv4 addresses can no longer be assumed to be either globally unique or invariant, and any protocol or applications design that assumes these properties will fail unpredictably. Work in the IAB and the NAT working group [NAT-ARCH] has analysed the impact of one specific cause of non-uniqueness and non-invariance, i.e., network address translators. Again the conclusion is that many applications will fail, unless they are specifically adapted to avoid the assumption of address transparency. One form of adaptation is the insertion of some form of application level gateway, and another form is for the NAT to modify payloads on the fly, but in either case the adaptation is application-specific.

Non-transparency of addresses is part of a more general phenomenon. We have to recognise that the Internet has lost end-to-end transparency, and this requires further analysis.

3. Multiple causes of loss of transparency

This section describes various recent inventions that have led to the loss of end-to-end transparency in the Internet.

3.1 The Intranet model

Underlying a number of the specific developments mentioned below is the concept of an "Intranet", loosely defined as a private corporate network using TCP/IP technology, and connected to the Internet at large in a carefully controlled manner. The Intranet is presumed to be used by corporate employees for business purposes, and to interconnect hosts that carry sensitive or confidential information. It is also held to a higher standard of operational availability than the Internet at large. Its usage can be monitored and controlled, and its resources can be better planned and tuned than those of the public network. These arguments of security and resource management have ensured the dominance of the Intranet model in most corporations and campuses.

The emergence of the Intranet model has had a profound effect on the notion of application transparency. Many corporate network managers feel it is for them alone to determine which applications can traverse the Internet/Intranet boundary. In this world view, address transparency may seem to be an unimportant consideration.

3.2 Dynamic address allocation

3.2.1 SLIP and PPP

It is to be noted that with the advent of vast numbers of dial-up Internet users, whose addresses are allocated at dial-up time, and whose traffic may be tunneled back to their home ISP, the actual IP addresses of such users are purely transient. During their period of validity they can be relied on end-to-end, but they must be forgotten at the end of every session. In particular they can have no permanent association with the domain name of the host borrowing them.

3.2.2 DHCP

Similarly, LAN-based users of the Internet today frequently use DHCP to acquire a new address at system restart, so here again the actual value of the address is potentially transient and must not be stored between sessions.

3.3 Firewalls

3.3.1 Basic firewalls

Intranet managers have a major concern about security: unauthorised traffic must be kept out of the Intranet at all costs. This concern led directly to the firewall concept (a system that intercepts all traffic between the Internet and the Intranet, and only lets through

selected traffic, usually belonging to a very limited set of applications). Firewalls, by their nature, fundamentally limit transparency.

3.3.2 SOCKS

A footnote to the effect of firewalls is the SOCKS mechanism [RFC 1928] by which untrusted applications such as telnet and ftp can punch through a firewall. SOCKS requires a shim library in the Intranet client, and a server in the firewall which is essentially an application level relay. As a result, the remote server does not see the real client; it believes that the firewall is the client.

3.4 Private addresses

When the threat of IPv4 address exhaustion first arose, and in some cases user sites were known to be "pirating" addresses for private use, a set of official private addresses were hurriedly allocated [RFC 1597] and later more carefully defined [BCP 5]. The legitimate existence of such an address allocation proved to very appealing, so Intranets with large numbers of non-global addresses came into existence. Unfortunately, such addresses by their nature cannot be used for communication across the public Internet; without special measures, hosts using private addresses are cut off from the world.

Note that private address space is sometimes asserted to be a security feature, based on the notion that outside knowledge of internal addresses might help intruders. This is a false argument, since it is trivial to hide addresses by suitable access control lists, even if they are globally unique - indeed that is a basic feature of a filtering router, the simplest form of firewall. A system with a hidden address is just as private as a system with a private address. There is of course no possible point in hiding the addresses of servers to which outside access is required.

It is also worth noting that the IPv6 equivalent of private addresses, i.e. site-local addresses, have similar characteristics to BCP 5 addresses, but their use will not be forced by a lack of globally unique IPv6 addresses.

3.5 Network address translators

Network address translators (NATs) are an almost inevitable consequence of the existence of Intranets using private addresses yet needing to communicate with the Internet at large. Their architectural implications are discussed at length in [NAT-ARCH], the fundamental point being that address translation on the fly destroys end-to-end address transparency and breaks any middleware or

applications that depend on it. Numerous protocols, for example H.323, carry IP addresses at application level and fail to traverse a simple NAT box correctly. If the full range of Internet applications is to be used, NATs have to be coupled with application level gateways (ALGs) or proxies. Furthermore, the ALG or proxy must be updated whenever a new address-dependent application comes along. In practice, NAT functionality is built into many firewall products, and all useful NATs have associated ALGs, so it is difficult to disentangle their various impacts.

3.6 Application level gateways, relays, proxies, and caches

It is reasonable to position application level gateways, relays, proxies, and caches at certain critical topological points, especially the Intranet/Internet boundary. For example, if an Intranet does not use SMTP as its mail protocol, an SMTP gateway is needed. Even in the normal case, an SMTP relay is common, and can perform useful mail routing functions, spam filtering, etc. (It may be observed that spam filtering is in some ways a firewall function, but the store-and-forward nature of electronic mail and the availability of MX records allow it to be a distinct and separate function.)

Similarly, for a protocol such as HTTP with a well-defined voluntary proxy mechanism, application proxies also serving as caches are very useful. Although these devices interfere with transparency, they do so in a precise way, correctly terminating network, transport and application protocols on both sides. They can however exhibit some shortfalls in ease of configuration and failover.

However, there appear to be cases of "involuntary" applications level devices such as proxies that grab and modify HTTP traffic without using the appropriate mechanisms, sometimes known as "transparent caches", or mail relays that purport to remove undesirable words. These devices are by definition not transparent, and may have totally unforeseeable side effects. (A possible conclusion is that even for non-store-and-forward protocols, a generic diversion mechanism analogous to the MX record would be of benefit. The SRV record [RFC 2052] is a step in this direction.)

3.7 Voluntary isolation and peer networks

There are communities that think of themselves as being so different that they require isolation via an explicit proxy, and even by using proprietary protocols and addressing schemes within their network. An example is the WAP Forum which targets very small phone-like devices with some capabilities for Internet connectivity. However, it's not

the Internet they're connecting directly to. They have to go through a proxy. This could potentially mean that millions of devices will never know the benefits of end-to-end connectivity to the Internet.

A similar effect arises when applications such as telephony span both an IP network and a peer network layer using a different technology. Although the application may work end-to-end, there is no possibility of end-to-end packet transmission.

3.8 Split DNS

Another consequence of the Intranet/Internet split is "split DNS" or "two faced DNS", where a corporate network serves up partly or completely different DNS inside and outside its firewall. There are many possible variants on this; the basic point is that the correspondence between a given FQDN (fully qualified domain name) and a given IPv4 address is no longer universal and stable over long periods.

3.9 Various load-sharing tricks

IPv4 was not designed to support anycast [RFC 1546], so there is no natural approach to load-sharing when one server cannot do the job. Various tricks have been used to resolve this (multicast to find a free server, the DNS returns different addresses for the same FQDN in a round-robin, a router actually routes packets sent to the same address automatically to different servers, etc.). While these tricks are not particularly harmful in the overall picture, they can be implemented in such a way as to interfere with name or address transparency.

4. Summary of current status and impact

It is impossible to estimate with any numerical reliability how widely the above inventions have been deployed. Since many of them preserve the illusion of transparency while actually interfering with it, they are extremely difficult to measure.

However it is certain that all the mechanisms just described are in very widespread use; they are not a marginal phenomenon. In corporate networks, many of them are the norm. Some of them (firewalls, relays, proxies and caches) clearly have intrinsic value given the Intranet concept. The others are largely artefacts and pragmatic responses to various pressures in the operational Internet, and they must be costing the industry very dearly in constant administration and complex fault diagnosis.

In particular, the decline of transparency is having a severe effect on deployment of end-to-end IP security. The Internet security model [SECMECH] calls for security at several levels (roughly, network, applications, and object levels). The current network level security model [RFC 2401] was constructed prior to the recognition that end-to-end address transparency was under severe threat. Although alternative proposals have begun to emerge [HIP] the current reality is that IPSEC cannot be deployed end-to-end in the general case. Tunnel-mode IPSEC can be deployed between corporate gateways or firewalls. Transport-mode IPSEC can be deployed within a corporate network in some cases, but it cannot span from Intranet to Internet and back to another Intranet if there is any chance of a NAT along the way.

Indeed, NAT breaks other security mechanisms as well, such as DNSSEC and Kerberos, since they rely on address values.

The loss of transparency brought about by private addresses and NATs affects many applications protocols to a greater or lesser extent. This is explored in detail in [NAT-PROT]. A more subtle effect is that the prevalence of dynamic addresses (from DHCP, SLIP and PPP) has fed upon the trend towards client/server computing. Today it is largely true that servers have fixed addresses, clients have dynamic addresses, and servers can in no way assume that a client's IP address identifies the client. On the other hand, clients rely on servers having stable addresses since a DNS lookup is the only generally deployed mechanism by which a client can find a server and solicit service. In this environment, there is little scope for true peer-to-peer applications protocols, and no easy solution for mobile servers. Indeed, the very limited demand for Mobile IP might be partly attributed to the market dominance of client/server applications in which the client's address is of transient significance. We also see a trend towards single points of failure such as media gateways, again resulting from the difficulty of implementing peer-to-peer solutions directly between clients with no fixed address.

The notion that servers can use precious globally unique addresses from a small pool, because there will always be fewer servers than clients, may become anachronistic when most electrical devices become network-manageable and thus become servers for a management protocol. Similarly, if every PC becomes a telephone (or the converse), capable of receiving unsolicited incoming calls, the lack of stable IP addresses for PCs will be an issue. Another impending paradigm shift is when domestic and small-office subscribers move from dial-up to always-on Internet connectivity, at which point transient address assignment from a pool becomes much less appropriate.

Many of the inventions described in the previous section lead to the datagram traffic between two hosts being directly or indirectly mediated by at least one other host. For example a client may depend on a DHCP server, a server may depend on a NAT, and any host may depend on a firewall. This violates the fate-sharing principle of [Saltzer] and introduces single points of failure. Worse, most of these points of failure require configuration data, yet another source of operational risk. The original notion that datagrams would find their way around failures, especially around failed routers, has been lost; indeed the overloading of border routers with additional functions has turned them into critical rather than redundant components, even for multihomed sites.

The loss of address transparency has other negative effects. For example, large scale servers may use heuristics or even formal policies that assign different priorities to service for different clients, based on their addresses. As addresses lose their global meaning, this mechanism will fail. Similarly, any anti-spam or anti-spoofing techniques that rely on reverse DNS lookup of address values can be confused by translated addresses. (Uncoordinated renumbering can have similar disadvantages.)

The above issues are not academic. They add up to complexity in applications design, complexity in network configuration, complexity in security mechanisms, and complexity in network management. Specifically, they make fault diagnosis much harder, and by introducing more single points of failure, they make faults more likely to occur.

5. Possible future directions

5.1 Successful migration to IPv6

In this scenario, IPv6 becomes fully implemented on all hosts and routers, including the adaptation of middleware, applications, and management systems. Since the address space then becomes big enough for all conceivable needs, address transparency can be restored. Transport-mode IPSEC can in principle deploy, given adequate security policy tools and a key infrastructure. However, it is widely believed that the Intranet/firewall model will certainly persist.

Note that it is a basic assumption of IPv6 that no artificial constraints will be placed on the supply of addresses, given that there are so many of them. Current practices by which some ISPs strongly limit the number of IPv4 addresses per client will have no reason to exist for IPv6. (However, addresses will still be assigned prudently, according to guidelines designed to favour hierarchical routing.)

Clearly this is in any case a very long term scenario, since it assumes that IPv4 has declined to the point where IPv6 is required for universal connectivity. Thus, a viable version of Scenario 5.3 is a prerequisite for Scenario 5.1.

5.2 Complete failure of IPv6

In this scenario, IPv6 fails to reach any significant level of operational deployment, IPv4 addressing is the only available mechanism, and address transparency cannot be restored. IPSEC cannot be deployed globally in its current form. In the very long term, the pool of globally unique IPv4 addresses will be nearly totally allocated, and new addresses will generally not be available for any purpose.

It is unclear exactly what is likely to happen if the Internet continues to rely exclusively on IPv4, because in that eventuality a variety of schemes are likely to promulgated, with a view toward providing an acceptable evolutionary path for the network. However, we can examine two of the more simplistic sub-scenarios which are possible, while realising that the future would be unlikely to match either one exactly:

5.2.1 Re-allocating the IPv4 address space

Suppose that a mechanism is created to continuously recover and re-allocate IPv4 addresses. A single global address space is maintained, with all sites progressively moving to an Intranet private address model, with global addresses being assigned temporarily from a pool of several billion.

5.2.1.1 A sub-sub-scenario of this is generalised use of NAT and NAPT (NAT with port number translation). This has the disadvantage that the host is unaware of the unique address being used for its traffic, being aware only of its ambiguous private address, with all the problems that this generates. See [NAT-ARCH].

5.2.1.2 Another sub-sub-scenario is the use of realm-specific IP addressing implemented at the host rather than by a NAT box. See [RSIP]. In this case the host is aware of its unique address, allowing for substantial restoration of the end-to-end usefulness of addresses, e.g. for TCP or cryptographic checksums.

5.2.1.3 A final sub-sub-scenario is the "map and encapsulate" model in which address translation is replaced by systematic encapsulation of all packets for wide-area transport. This model has never been fully developed, although it is completely compatible with end-to-end addressing.

5.2.2 Exhaustion

Suppose that no mechanism is created to recover addresses (except perhaps black or open market trading). Sites with large address blocks keep them. All the scenarios of 5.2.1 appear but are insufficient. Eventually the global address space is no longer adequate. This is a nightmare scenario - NATs appear within the "global" address space, for example at ISP-ISP boundaries. It is unclear how a global routing system and a global DNS system can be maintained; the Internet is permanently fragmented.

5.3 Partial deployment of IPv6

In this scenario, IPv6 is completely implemented but only deploys in certain segments of the network (e.g. in certain countries, or in certain areas of application such as mobile hand-held devices). Instead of being transitional in nature, some of the IPv6 transition techniques become permanent features of the landscape. Sometimes addresses are 32 bits, sometimes they are 128 bits. DNS lookups may return either or both. In the 32 bit world, the scenarios 5.2.1 and 5.2.2 may occur. IPSEC can only deploy partially.

6. Conclusion

None of the above scenarios is clean, simple and straightforward. Although the pure IPv6 scenario is the cleanest and simplest, it is not straightforward to reach it. The various scenarios without use of IPv6 are all messy and ultimately seem to lead to dead ends of one kind or another. Partial deployment of IPv6, which is a required step on the road to full deployment, is also messy but avoids the dead ends.

7. Security Considerations

The loss of transparency is both a bug and a feature from the security viewpoint. To the extent that it prevents the end-to-end deployment of IPSEC, it damages security and creates vulnerabilities. For example, if a standard NAT box is in the path, then the best that can be done is to decrypt and re-encrypt IP traffic in the NAT. The traffic will therefore be momentarily in clear text and potentially vulnerable. Furthermore, the NAT will possess many keys and will be a prime target of attack. In environments where this is unacceptable,

encryption must be applied above the network layer instead, of course with no usage whatever of IP addresses in data that is cryptographically protected. See section 4 for further discussion.

In certain scenarios, i.e. 5.1 (full IPv6) and 5.2.1.2 (RSIP), end-to-end IPSEC would become possible, especially using the "distributed firewalls" model advocated in [Bellovin].

The loss of transparency at the Intranet/Internet boundary may be considered a security feature, since it provides a well defined point at which to apply restrictions. This form of security is subject to the "crunchy outside, soft inside" risk, whereby any successful penetration of the boundary exposes the entire Intranet to trivial attack. The lack of end-to-end security applied within the Intranet also ignores insider threats.

It should be noted that security applied above the transport level, such as SSL, SSH, PGP or S/MIME, is not affected by network layer transparency issues.

Acknowledgements

This document and the related issues have been discussed extensively by the IAB. Special thanks to Steve Deering for a detailed review and to Noel Chiappa. Useful comments or ideas were also received from Rob Austein, John Bartas, Jim Bound, Scott Bradner, Vint Cerf, Spencer Dawkins, Anoop Ghanwani, Erik Guttman, Eric A. Hall, Graham Klyne, Dan Kohn, Gabriel Montenegro, Thomas Narten, Erik Nordmark, Vern Paxson, Michael Quinlan, Eric Rosen, Daniel Senie, Henning Schulzrinne, Bill Sommerfeld, and George Tsirtsis.

References

- [Bellovin] Distributed Firewalls, S. Bellovin, available at <http://www.research.att.com/~smb/papers/distfw.pdf> or <http://www.research.att.com/~smb/papers/distfw.ps> (work in progress).
- [Berners-Lee] Weaving the Web, T. Berners-Lee, M. Fischetti, HarperCollins, 1999, p 208.
- [Saltzer] End-To-End Arguments in System Design, J.H. Saltzer, D.P.Reed, D.D.Clark, ACM TOCS, Vol 2, Number 4, November 1984, pp 277-288.
- [IEN 48] Cerf, V., "The Catenet Model for Internetworking," Information Processing Techniques Office, Defense Advanced Research Projects Agency, IEN 48, July 1978.

- [CATENET] Pouzin, L., "A Proposal for Interconnecting Packet Switching Networks," Proceedings of EUROCOMP, Brunel University, May 1974, pp. 1023-36.
- [STD 7] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC 1546] Partridge, C., Mendez, T. and W. Milliken, "Host Anycasting Service", RFC 1546, November 1993.
- [RFC 1597] Rekhter, Y., Moskowitz, B., Karrenberg, D. and G. de Groot, "Address Allocation for Private Internets", RFC 1597, March 1994.
- [RFC 1633] Braden, R., Clark, D. and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.
- [RFC 1889] Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", RFC 1889, January 1996.
- [BCP 5] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC 1928] Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D. and L. Jones, "SOCKS Protocol Version 5", RFC 1928, March 1996.
- [RFC 1958] Carpenter, B., "Architectural Principles of the Internet", RFC 1958, June 1996.
- [RFC 2018] Mathis, M., Mahdavi, J., Floyd, S. and A. Romanow, "TCP Selective Acknowledgement Options", RFC 2018, October 1996.
- [RFC 2052] Gulbrandsen, A. and P. Vixie, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2052, October 1996.
- [RFC 2101] Carpenter, B., Crowcroft, J. and Y. Rekhter, "IPv4 Address Behaviour Today", RFC 2101, February 1997.
- [RFC 2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997.

- [RFC 2309] Braden, B., Clark, D., Crowcroft, J., Davie, B., Deering, S., Estrin, D., Floyd, S., Jacobson, V., Minshall, G., Partridge, C., Peterson, L., Ramakrishnan, K., Shenker, S., Wroclawski, J. and L. Zhang, "Recommendations on Queue Management and Congestion Avoidance in the Internet", RFC 2309, April 1998.
- [RFC 2326] Schulzrinne, H., Rao, A. and R. Lanphier, "Real Time Streaming Protocol (RTSP)", RFC 2326, April 1998.
- [RFC 2401] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [RFC 2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, "An Architecture for Differentiated Service", RFC 2475, December 1998.
- [RFC 2581] Allman, M., Paxson, V. and W. Stevens, "TCP Congestion Control", RFC 2581, April 1999.
- [NAT-ARCH] Hain, T., "Architectural Implications of NAT", Work in Progress.
- [NAT-PROT] Holdrege, M. and P. Srisuresh, "Protocol Complications with the IP Network Address Translator (NAT)", Work in Progress.
- [SECMECH] Bellovin, S., "Security Mechanisms for the Internet", Work in Progress.
- [RSIP] Lo, J., Borella, M. and D. Grabelsky, "Realm Specific IP: A Framework", Work in Progress.
- [HIP] Moskowitz, R., "The Host Identity Payload", Work in Progress.

Author's Address

Brian E. Carpenter
IBM
c/o iCAIR
Suite 150
1890 Maple Avenue
Evanston, IL 60201
USA

EMail: brian@icair.org

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

