

IAB Concerns and Recommendations
Regarding Internet Research and Evolution

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 (2004).

Abstract

This document discusses IAB concerns that ongoing research is needed to further the evolution of the Internet infrastructure, and that consistent, sufficient non-commercial funding is needed to enable such research.

Table of Contents

1.	Introduction	2
1.1.	Document Organization.	2
1.2.	IAB Concerns	3
1.3.	Contributions to this Document	4
2.	History of Internet Research and Research Funding.	4
2.1.	Prior to 1980.	4
2.2.	1980s and early 1990s.	5
2.3.	Mid-1990s to 2003.	6
2.4.	Current Status	6
3.	Open Internet Research Topics.	7
3.1.	Scope and Limitations.	7
3.2.	Naming	8
3.2.1.	Domain Name System (DNS).	8
3.2.2.	New Namespaces.	9
3.3.	Routing.	9
3.3.1.	Inter-domain Routing.	10
3.3.2.	Routing Integrity	11
3.3.3.	Routing Algorithms.	12
3.3.4.	Mobile and Ad-Hoc Routing	13
3.4.	Security	13

3.4.1.	Formal Methods.	14
3.4.2.	Key Management.	14
3.4.3.	Cryptography.	15
3.4.4.	Security for Distributed Computing.	15
3.4.5.	Deployment Considerations in Security	15
3.4.6.	Denial of Service Protection.	16
3.5.	Network Management	16
3.5.1.	Managing Networks, Not Devices.	16
3.5.2.	Enhanced Monitoring Capabilities.	17
3.5.3.	Customer Network Management	17
3.5.4.	Autonomous Network Management	17
3.6.	Quality of Service	17
3.6.1.	Inter-Domain QoS Architecture	18
3.6.2.	New Queuing Disciplines	19
3.7.	Congestion Control	19
3.8.	Studying the Evolution of the Internet Infrastructure.	20
3.9.	Middleboxes.	21
3.10.	Internet Measurement	21
3.11.	Applications	22
3.12.	Meeting the Needs of the Future.	22
3.13.	Freely Distributable Prototypes.	23
4.	Conclusions.	23
5.	Acknowledgements	23
6.	Security Considerations.	24
7.	Informative References	24
8.	Authors' Addresses	29
9.	Full Copyright Statement	30

1. Introduction

This document discusses the history of funding for Internet research, expresses concern about the current state of such funding, and outlines several specific areas that the IAB believes merit additional research. Current funding levels for Internet research are not generally adequate, and several important research areas are significantly underfunded. This situation needs to be rectified for the Internet to continue its evolution and development.

1.1. Document Organization

The first part of the document is a high-level discussion of the history of funding for Internet research to provide some historical context to this document. The early funding of Internet research was largely from the U.S. government, followed by a period in the second half of the 1990s of commercial funding and of funding from several governments. However, the commercial funding for Internet research has been reduced due to the recent economic downturn.

The second part of the document provides an incomplete set of open Internet research topics. These are only examples, intended to illustrate the breadth of open research topics. This second section supports the general thesis that ongoing research is needed to further the evolution of the Internet infrastructure. This includes research on the medium-time-scale evolution of the Internet infrastructure as well as research on longer-time-scale grand challenges. This also includes many research issues that are already being actively investigated in the Internet research community.

Areas that are discussed in this section include the following: naming, routing, security, network management, and transport. Issues that require more research also include more general architectural issues such as layering and communication between layers. In addition, general topics discussed in this section include modeling, measurement, simulation, test-beds, etc. We are focusing on topics that are related to the IETF and IRTF (Internet Research Task Force) agendas. (For example, Grid issues are not discussed in this document because they are addressed through the Global Grid Forum and other Grid-specific organizations, not in the IETF.)

Where possible, the examples in this document point to separate documents on these issues, and only give a high-level summary of the issues raised in those documents.

1.2. IAB Concerns

In the aftermath of September 11 2001, there seems to be a renewed interest by governments in funding research for Internet-related security issues. From [Jackson02]: "It is generally agreed that the security and reliability of the basic protocols underlying the Internet have not received enough attention because no one has a proprietary interest in them".

That quote brings out a key issue in funding for Internet research, which is that because no single organization (e.g., no single government, software company, equipment vendor, or network operator) has a sense of ownership of the global Internet infrastructure, research on the general issues of the Internet infrastructure are often not adequately funded. In our current challenging economic climate, it is not surprising that commercial funding sources are more likely to fund that research that leads to a direct competitive advantage.

The principal thesis of this document is that if commercial funding is the main source of funding for future Internet research, the future of the Internet infrastructure could be in trouble. In addition to issues about which projects are funded, the funding

source can also affect the content of the research, for example, towards or against the development of open standards, or taking varying degrees of care about the effect of the developed protocols on the other traffic on the Internet.

At the same time, many significant research contributions in networking have come from commercial funding. However, for most of the topics in this document, relying solely on commercially-funded research would not be adequate. Much of today's commercial funding is focused on technology transition, taking results from non-commercial research and putting them into shipping commercial products. We have not tried to delve into each of the research issues below to discuss, for each issue, what are the potentials and limitations of commercial funding for research in that area.

On a more practical note, if there was no commercial funding for Internet research, then few research projects would be taken to completion with implementations, deployment, and follow-up evaluation.

While it is theoretically possible for there to be too much funding for Internet research, that is far from the current problem. There is also much that could be done within the network research community to make Internet research more focused and productive, but that would belong in a separate document.

1.3. Contributions to this Document

A number of people have directly contributed text for this document, even though, following current conventions, the official RFC author list includes only the key editors of the document. The Acknowledgements section at the end of the document thanks other people who contributed to this document in some form.

2. History of Internet Research and Research Funding

2.1. Prior to 1980

Most of the early research into packet-switched networks was sponsored by the U.S. Defense Advanced Research Projects Agency (DARPA) [CSTB99]. This includes the initial design, implementation, and deployment of the ARPAnet connecting several universities and other DARPA contractors. The ARPAnet originally came online in the late 1960s. It grew in size during the 1970s, still chiefly with DARPA funding, and demonstrated the utility of packet-switched networking.

DARPA funding for Internet design started in 1973, just four years after the initial ARPAnet deployment. The support for Internet design was one result of prior DARPA funding for packet radio and packet satellite research. The existence of multiple networks (ARPAnet, packet radio, and packet satellite) drove the need for internetworking research. The Internet arose in large measure as a consequence of DARPA research funding for these three networks -- and arise only incidentally from the commercially-funded work at Xerox PARC on Ethernet.

2.2. 1980s and early 1990s

The ARPAnet converted to the Internet Protocol (IP) on January 1, 1983, approximately 20 years before this document was written. Throughout the 1980s, the U.S. Government continued strong research and development funding for Internet technology. DARPA continued to be the key funding source, but was supplemented by other DoD (U.S. Department of Defense) funding (e.g., via the Defense Data Network (DDN) program of the Defense Communication Agency (DCA)) and other U.S. Government funding (e.g., U.S. Department of Energy (DoE) funding for research networks at DoE national laboratories, (U.S.) National Science Foundation (NSF) funding for academic institutions). This funding included basic research, applied research (including freely distributable prototypes), the purchase of IP-capable products, and operating support for the IP-based government networks such as ARPAnet, ESnet, MILnet, the NASA Science Internet, and NSFnet.

During the 1980s, the U.S. DoD desired to leave the business of providing operational network services to academic institutions, so funding for most academic activities moved over to the NSF during the decade. NSF's initial work included sponsorship of CSnet in 1981. By 1986, NSF was also sponsoring various research projects into networking (e.g., Mills' work on Fuzzballs). In the late 1980s, NSF created the NSFnet backbone and sponsored the creation of several NSF regional networks (e.g., SURAnet) and interconnections with several international research networks. NSF also funded gigabit networking research, through the Corporation for National Research Initiatives (CNRI), starting in the late 1980s. It is important to note that the NSF sponsorship was focused on achieving core NSF goals, such as connecting scientists at leading universities to NSF supercomputing centers. The needs of high-performance remote access to supercomputers drove the overall NSFnet performance. As a side effect, this meant that students and faculty at those universities enjoyed a relatively high-performance Internet environment. As those students graduated, they drove both commercial use of the Internet and the nascent residential market. It is no accident that this was the environment from which the world wide web emerged.

Most research funding outside the U.S. during the 1980s and early 1990s was focused on the ISO OSI networking project or on then-new forms of network media (e.g., wireless, broadband access). The European Union was a significant source of research funding for the networking community in Europe during this period. Some of the best early work in gigabit networking was undertaken in the UK and Sweden.

2.3. Mid-1990s to 2003

Starting in the middle 1990s, U.S. Government funding for Internet research and development was significantly reduced. The premise for this was that the growing Internet industry would pay for whatever research and development that was needed. Some funding for Internet research and development has continued in this period from European and Asian organizations (e.g., the WIDE Project in Japan [WIDE]). Reseaux IP Europeens [RIPE] is an example of market-funded networking research in Europe during this period.

Experience during this period has been that commercial firms have often focused on donating equipment to academic institutions and promoting somewhat vocationally-focused educational projects. Many of the commercially-funded research and development projects appear to have been selected because they appeared likely to give the funding source a specific short-term economic advantage over its competitors. Higher risk, more innovative research proposals generally have not been funded by industry. A common view in Silicon Valley has been that established commercial firms are not very good at transitioning cutting edge research into products, but were instead good at buying small startup firms who had successfully transitioned such cutting edge research into products. Unfortunately, small startup companies are generally unable financially to fund any research themselves.

2.4. Current Status

The result of reduced U.S. Government funding and profit-focused, low-risk, short-term industry funding has been a decline in higher-risk but more innovative research activities. Industry has also been less interested in research to evolve the overall Internet architecture, because such work does not translate into a competitive advantage for the firm funding such work.

The IAB believes that it would be helpful for governments and other non-commercial sponsors to increase their funding of both basic research and applied research relating to the Internet, and to sustain these funding levels going forward.

3. Open Internet Research Topics

This section primarily discusses some specific topics that the IAB believes merit additional research. Research, of course, includes not just devising a theory, algorithm, or mechanism to accomplish a goal, but also evaluating the general efficacy of the approach and then the benefits vs. the costs of deploying that algorithm or mechanism. Important cautionary notes about this discussion are given in the next sub-section. This particular set of topics is not intended to be comprehensive, but instead is intended to demonstrate the breadth of open Internet research questions.

Other discussions of problems of the Internet that merit further research include the following:

[CIPB02, Claffy03a, Floyd, NSF03a, NSF03b].

3.1. Scope and Limitations

This document is NOT intended as a guide for public funding agencies as to exactly which projects or proposals should or should not be funded.

In particular, this document is NOT intended to be a comprehensive list of **all** of the research questions that are important to further the evolution of the Internet; that would be a daunting task, and would presuppose a wider and more intensive effort than we have undertaken in this document.

Similarly, this document is not intended to list the research questions that are judged to be only of peripheral importance, or to survey the current (global; governmental, commercial, and academic) avenues for funding for Internet research, or to make specific recommendations about which areas need additional funding. The purpose of the document is to persuade the reader that ongoing research is needed towards the continued evolution of the Internet infrastructure; the purpose is not to make binding pronouncements about which specific areas are and are not worthy of future funding.

For some research clearly relevant to the future evolution of the Internet, there are grand controversies between competing proposals or competing schools of thought; it is not the purpose of this document to take positions in these controversies, or to take positions on the nature of the solutions for areas needing further research.

That all carefully noted, the remainder of this section discusses a broad set of research areas, noting a subset of particular topics of interest in each of those research areas. Again, this list is NOT comprehensive, but rather is intended to suggest that a broad range of ongoing research is needed, and to propose some candidate topics.

3.1.1. Terminology

Several places in this document refer to 'network operators'. By that term, we intend to include anyone or any organization that operates an IP-based network; we are not using that term in the narrow meaning of commercial network service providers.

3.2. Naming

The Internet currently has several different namespaces, including IP addresses, sockets (specified by the IP address, upper-layer protocol, and upper-layer port number), Autonomous System (AS) number, and the Fully-Qualified Domain Name (FQDN). Many of the Internet's namespaces are supported by the widely deployed Domain Name System [RFC-3467] or by various Internet applications [RFC-2407, Section 4.6.2.1]

3.2.1. Domain Name System (DNS)

The DNS system, while it works well given its current constraints, has several stress points.

The current DNS system relies on UDP for transport, rather than SCTP or TCP. Given the very large number of clients using a typical DNS server, it is desirable to minimize the state on the DNS server side of the connection. UDP does this well, so it is a reasonable choice, though this has other implications, for example a reliance on UDP fragmentation. With IPv6, intermediate fragmentation is not allowed and Path MTU Discovery is mandated. However, the amount of state required to deploy Path MTU Discovery for IPv6 on a DNS server might be a significant practical problem.

One implication of this is that research into alternative transport protocols, designed more for DNS-like applications where there are very many clients using each server, might be useful. Of particular interest would be transport protocols with little burden for the DNS server, even if that increased the burden somewhat for the DNS client.

Additional study of DNS caching, both currently available caching techniques and also of potential new caching techniques, might be helpful in finding ways to reduce the offered load for a typical DNS

server. In particular, examination of DNS caching through typical commercial firewalls might be interesting if it lead to alternative firewall implementations that were less of an obstacle to DNS caching.

The community lacks a widely-agreed-upon set of metrics for measuring DNS server performance. It would be helpful if people would seriously consider what characteristics of the DNS system should be measured.

Some in the community would advocate replacing the current DNS system with something better. Past attempts to devise a better approach have not yielded results that persuaded the community to change. Proposed work in this area could be very useful, but might require careful scrutiny to avoid falling into historic design pitfalls.

With regards to DNS security, major technical concerns include finding practical methods for signing very large DNS zones (e.g., and tools to make it easier to manage secure DNS infrastructure.

Most users are unable to distinguish a DNS-related failure from a more general network failure. Hence, maintaining the integrity and availability of the Domain Name System is very important for the future health of the Internet.

3.2.2. New Namespaces

Additionally, the Namespace Research Group (NSRG) of the Internet Research Task Force (IRTF) studied adding one or more additional namespaces to the Internet Architecture [LD2002]. Many members of the IRTF NSRG believe that there would be significant architectural benefit to adding one or more additional namespaces to the Internet Architecture. Because smooth consensus on that question or on the properties of a new namespace was not obtained, the IRTF NSRG did not make a formal recommendation to the IETF community regarding namespaces. The IAB believes that this is an open research question worth examining further.

Finally, we believe that future research into the evolution of Internet-based distributed computing might well benefit from studying adding additional namespaces as part of a new approach to distributed computing.

3.3. Routing

The currently deployed unicast routing system works reasonably well for most users. However, the current unicast routing architecture is suboptimal in several areas, including the following: end-to-end

convergence times in global-scale catenets (a system of networks interconnected via gateways); the ability of the existing inter-domain path-vector algorithm to scale well beyond 200K prefixes; the ability of both intra-domain and inter-domain routing to use multiple metrics and multiple kinds of metrics concurrently; and the ability of IPv4 and IPv6 to support widespread site multi-homing without undue adverse impact on the inter-domain routing system. Integrating policy into routing is also a general concern, both for intra-domain and inter-domain routing. In many cases, routing policy is directly tied to economic issues for the network operators, so applied research into routing ideally would consider economic considerations as well as technical considerations.

This is an issue for which the commercial interest is clear, but that seems unlikely to be solved through commercial funding for research, in the absence of a consortium of some type.

3.3.1. Inter-domain Routing

The current operational inter-domain routing system has between 150,000 and 200,000 routing prefixes in the default-free zone (DFZ) [RFC-3221]. ASIC technology obviates concerns about the ability to forward packets at very high speeds. ASIC technology also obviates concerns about the time required to perform longest-prefix-match computations. However, some senior members of the Internet routing community have concerns that the end-to-end convergence properties of the global Internet might hit fundamental algorithmic limitations (i.e., not hardware limitations) when the DFZ is somewhere between 200,000 and 300,000 prefixes. Research into whether this concern is well-founded in scientific terms seems very timely.

Separately from the above concern, recent work has shown that there can be significant BGP convergence issues today. At present, it appears that the currently observed convergence issues relate to how BGP has been configured by network operators, rather than being any sort of fundamental algorithmic limitation [MGVK02]. This convergence time issue makes the duration of the apparent network outage much longer than it should be. Additional applied research into which aspects of a BGP configuration have the strongest impact on convergence times would help mitigate the currently observed operational issues.

Also, inter-domain routing currently requires significant human engineering of specific inter-AS paths to ensure that reasonably optimal paths are used by actual traffic. Ideally, the inter-domain routing system would automatically cause reasonably optimal paths to be chosen. Recent work indicates that improved BGP policy mechanisms

might help ensure that reasonably optimal paths are normally used for inter-domain IP traffic. [SMA03] Continued applied research in this area might lead to substantially better technical approaches.

The current approach to site multi-homing has the highly undesirable side-effect of significantly increasing the growth rate of prefix entries in the DFZ (by impairing the deployment of prefix aggregation). Research is needed into new routing architectures that can support large-scale site multi-homing without the undesirable impacts on inter-domain routing of the current multi-homing technique.

The original application for BGP was in inter-domain routing, primarily within service provider networks but also with some use by multi-homed sites. However, some are now trying to use BGP in other contexts, for example highly mobile environments, where it is less obviously well suited. Research into inter-domain routing and/or intra-domain policy routing might lead to other approaches for any emerging environments where the current BGP approach is not the optimal one.

3.3.2. Routing Integrity

Recently there has been increased awareness of the longstanding issue of deploying strong authentication into the Internet inter-domain routing system. Currently deployed mechanisms (e.g., BGP TCP MD5 [RFC-2385], OSPF MD5, RIP MD5 [RFC-2082]) provide cryptographic authentication of routing protocol messages, but no authentication of the actual routing data. Recent proposals (e.g., S-BGP [KLMS2000]) for improving this in inter-domain routing appear difficult to deploy across the Internet, in part because of their reliance on a single trust hierarchy (e.g., a single PKI). Similar proposals (e.g., OSPF with Digital Signatures, [RFC-2154]) for intra-domain routing are argued to be computationally infeasible to deploy in a large network.

A recurring challenge with any form of inter-domain routing authentication is that there is no single completely accurate source of truth about which organizations have the authority to advertise which address blocks. Alternative approaches to authentication of data in the routing system need to be developed. In particular, the ability to perform partial authentication of routing data would facilitate incremental deployment of routing authentication mechanisms. Also, the ability to use non-hierarchical trust models (e.g., the web of trust used in the PGP application) might facilitate incremental deployment and might resolve existing concerns about centralized administration of the routing system, hence it merits additional study and consideration.

3.3.3. Routing Algorithms

The current Internet routing system relies primarily on two algorithms. Link-state routing uses the Dijkstra algorithm [Dijkstra59]. Distance-Vector routing (e.g., RIP) and Path-Vector routing (e.g., BGP) use the Bellman-Ford algorithm [Bellman1957, FF1962]. Additional ongoing basic research into graph theory as applied to routing is worthwhile and might yield algorithms that would enable a new routing architecture or otherwise provide improvements to the routing system.

Currently deployed multicast routing relies on the Deering RPF algorithm [Deering1988]. Ongoing research into alternative multicast routing algorithms and protocols might help alleviate current concerns with the scalability of multicast routing.

The deployed Internet routing system assumes that the shortest path is always the best path. This is provably false, however it is a reasonable compromise given the routing protocols currently available. The Internet lacks deployable approaches for policy-based routing or routing with alternative metrics (i.e., some metric other than the number of hops to the destination). Examples of alternative policies include: the path with lowest monetary cost; the path with the lowest probability of packet loss; the path with minimized jitter; and the path with minimized latency. Policy metrics also need to take business relationships into account. Historic work on QoS-based routing has tended to be unsuccessful in part because it did not adequately consider economic and commercial considerations of the routing system and in part because of inadequate consideration of security implications.

Transitioning from the current inter-domain routing system to any new inter-domain routing system is unlikely to be a trivial exercise. So any proposal for a new routing system needs to carefully consider and document deployment strategies, transition mechanisms, and other operational considerations. Because of the cross-domain interoperability aspect of inter-domain routing, smooth transitions from one inter-domain routing system are likely to be difficult to accomplish. Separately, the inter-domain routing system lacks strong market forces that would encourage migration to better technical approaches. Hence, it appears unlikely that the commercial sector will be the source of a significantly improved inter-domain routing system.

3.3.4. Mobile and Ad-Hoc Routing

While some of the earliest DARPA-sponsored networking research involved packet radio networks, mobile routing [IM1993] and mobile ad-hoc routing [RFC-2501] are relatively recent arrivals in the Internet, and are not yet widely deployed. The current approaches are not the last word in either of those arenas. We believe that additional research into routing support for mobile hosts and mobile networks is needed. Additional research for ad-hoc mobile hosts and mobile networks is also worthwhile. Ideally, mobile routing and mobile ad-hoc routing capabilities should be native inherent capabilities of the Internet routing architecture. This probably will require a significant evolution from the existing Internet routing architecture. (NB: The term "mobility" as used here is not limited to mobile telephones, but instead is very broadly defined, including laptops that people carry, cars/trains/aircraft, and so forth.)

Included in this topic are a wide variety of issues. The more distributed and dynamic nature of partially or completely self-organizing routing systems (including the associated end nodes) creates unique security challenges (especially relating to Authorization, Authentication, and Accounting, and relating to key management). Scalability of wireless networks can be difficult to measure or to achieve. Enforced hierarchy is one approach, but can be very limiting. Alternative, less constraining approaches to wireless scalability are desired. Because wireless link-layer protocols usually have some knowledge of current link characteristics such as link quality, sublayer congestion conditions, or transient channel behavior, it is desirable to find ways to let network-layer routing use such data. This raises architectural questions of what the proper layering should be, which functions should be in which layer, and also practical considerations of how and when such information sharing should occur in real implementations.

3.4. Security

The Internet has a reputation for not having sufficient security. In fact, the Internet has a number of security mechanisms standardized, some of which are widely deployed. However, there are a number of open research questions relating to Internet security. In particular, security mechanisms need to be incrementally deployable and easy to use. "[Security] technology must be easy to use, or it will not be configured correctly. If mis-configured, security will be lost, but things will 'work'" [Schiller03].

3.4.1. Formal Methods

There is an ongoing need for funding of basic research relating to Internet security, including funding of formal methods research that relates to security algorithms, protocols, and systems.

For example, it would be beneficial to have more formal study of non-hierarchical trust models (e.g., PGP's Web-of-Trust model). Use of a hierarchical trust model can create significant limitations in how one might approach securing components of the Internet, for example the inter-domain routing system. So research to develop new trust models suited for the Internet or on the applicability of existing non-hierarchical trust models to existing Internet problems would be worthwhile.

While there has been some work on the application of formal methods to cryptographic algorithms and cryptographic protocols, existing techniques for formal evaluation of algorithms and protocols lack sufficient automation. This lack of automation means that many protocols aren't formally evaluated in a timely manner. This is problematic for the Internet because formal evaluation has often uncovered serious anomalies in cryptographic protocols. The creation of automated tools for applying formal methods to cryptographic algorithms and/or protocols would be very helpful.

3.4.2. Key Management

A recurring challenge to the Internet community is how to design, implement, and deploy key management appropriate to the myriad of security contexts existing in the global Internet. Most current work in unicast key management has focused on hierarchical trust models, because much of the existing work has been driven by corporate or military "top-down" operating models.

The paucity of key management methods applicable to non-hierarchical trust models (see above) is a significant constraint on the approaches that might be taken to secure components of the Internet.

Research focused on removing those constraints by developing practical key management methods applicable to non-hierarchical trust models would be very helpful.

Topics worthy of additional research include key management techniques, such as non-hierarchical key management architectures (e.g., to support non-hierarchical trust models; see above), that are useful by ad-hoc groups in mobile networks and/or distributed computing.

Although some progress has been made in recent years, scalable multicast key management is far from being a solved problem. Existing approaches to scalable multicast key management add significant constraints on the problem scope in order to come up with a deployable technical solution. Having a more general approach to scalable multicast key management (i.e., one having broader applicability and fewer constraints) would enhance the Internet's capabilities.

In many cases, attribute negotiation is an important capability of a key management protocol. Experience with the Internet Key Exchange (IKE) to date has been that it is unduly complex. Much of IKE's complexity derives from its very general attribute negotiation capabilities. A new key management approach that supported significant attribute negotiation without creating challenging levels of deployment and operations complexity would be helpful.

3.4.3. Cryptography

There is an ongoing need to continue the open-world research funding into both cryptography and cryptanalysis. Most governments focus their cryptographic research in the military-sector. While this is understandable, those efforts often have limited (or no) publications in the open literature. Since the Internet engineering community must work from the open literature, it is important that open-world research continues in the future.

3.4.4. Security for Distributed Computing

MIT's Project Athena was an important and broadly successful research project into distributed computing. Project Athena developed the Kerberos [RFC-1510] security system, which has significant deployment today in campus environments. However, inter-realm Kerberos is neither as widely deployed nor perceived as widely successful as single-realm Kerberos. The need for scalable inter-domain user authentication is increasingly acute as ad-hoc computing and mobile computing become more widely deployed. Thus, work on scalable mechanisms for mobile, ad-hoc, and non-hierarchical inter-domain authentication would be very helpful.

3.4.5. Deployment Considerations in Security

Lots of work has been done on theoretically perfect security that is impossible to deploy. Unfortunately, the S-BGP proposal is an example of a good research product that has significant unresolved deployment challenges. It is far from obvious how one could widely deploy S-BGP without previously deploying a large-scale inter-domain public-key infrastructure and also centralizing route advertisement

policy enforcement in the Routing Information Registries or some similar body. Historically, public-key infrastructures have been either very difficult or impossible to deploy at large scale. Security mechanisms that need additional infrastructure have not been deployed well. We desperately need security that is general, easy to install, and easy to manage.

3.4.6. Denial of Service Protection

Historically, the Internet community has mostly ignored pure Denial of Service (DoS) attacks. This was appropriate at one time since such attacks were rare and are hard to defend against. However, one of the recent trends in adversarial software (e.g., viruses, worms) has been the incorporation of features that turn the infected host into a "zombie". Such zombies can be remotely controlled to mount a distributed denial of service attack on some victim machine. In many cases, the authorized operators of systems are not aware that some or all of their systems have become zombies. It appears that the presence of non-trivial numbers of zombies in the global Internet is now endemic, which makes distributed denial of service attacks a much larger concern. So Internet threat models need to assume the presence of such zombies in significant numbers. This makes the design of protocols resilient in the presence of distributed denial of service attacks very important to the health of the Internet. Some work has been done on this front [Savage00], [MBFIPS01], but more is needed.

3.5. Network Management

The Internet had early success in network device monitoring with the Simple Network Management Protocol (SNMP) and its associated Management Information Base (MIB). There has been comparatively less success in managing networks, in contrast to the monitoring of individual devices. Furthermore, there are a number of operator requirements not well supported by the current Internet management framework. It is desirable to enhance the current Internet network management architecture to more fully support operational needs.

Unfortunately, network management research has historically been very underfunded. Operators have complained that existing solutions are inadequate. Research is needed to find better solutions.

3.5.1. Managing Networks, Not Devices

At present there are few or no good tools for managing a whole network instead of isolated devices. For example, the lack of appropriate network management tools has been cited as one of the major barriers to the widespread deployment of IP multicast [Diot00,

SM03]. Current network management protocols, such as the Simple Network Management Protocol (SNMP), are fine for reading status of well-defined objects from individual boxes. Managing networks instead of isolated devices requires the ability to view the network as a large distributed system. Research is needed on scalable distributed data aggregation mechanisms, scalable distributed event correlation mechanisms, and distributed and dependable control mechanisms.

Applied research into methods of managing sets of networked devices seems worthwhile. Ideally, such a management approach would support distributed management, rather than being strictly centralized.

3.5.2. Enhanced Monitoring Capabilities

SNMP does not always scale well to monitoring large numbers of objects in many devices in different parts of the network. An alternative approach worth exploring is how to provide scalable and distributed monitoring, not on individual devices, but instead on groups of devices and the network-as-a-whole. This requires scalable techniques for data aggregation and event correlation of network status data originating from numerous locations in the network.

3.5.3. Customer Network Management

An open issue related to network management is helping users and others to identify and resolve problems in the network. If a user can't access a web page, it would be useful if the user could find out, easily, without having to run ping and traceroute, whether the problem was that the web server was down, that the network was partitioned due to a link failure, that there was heavy congestion along the path, that the DNS name couldn't be resolved, that the firewall prohibited the access, or that some other specific event occurred.

3.5.4. Autonomous Network Management

More research is needed to improve the degree of automation achieved by network management systems and to localize management. Autonomous network management might involve the application of control theory, artificial intelligence or expert system technologies to network management problems.

3.6. Quality of Service

There has been an intensive body of research and development work on adding QoS to the Internet architecture for more than ten years now [RFC-1633, RFC-2474, RFC-3260, RFC-2205, RFC-2210], yet we still

don't have end-to-end QoS in the Internet [RFC-2990, RFC-3387]. The IETF is good at defining individual QoS mechanisms, but poor at work on deployable QoS architectures. Thus, while Differentiated Services (DiffServ) mechanisms have been standardized as per-hop behaviors, there is still much to be learned about the deployment of that or other QoS mechanisms for end-to-end QoS. In addition to work on purely technical issues, this includes close attention to the economic models and deployment strategies that would enable an increased deployment of QoS in the network.

In many cases, deployment of QoS mechanisms would significantly increase operational security risks [RFC-2990], so any new research on QoS mechanisms or architectures ought to specifically discuss the potential security issues associated with the new proposal(s) and how to mitigate those security issues.

In some cases, the demand for QoS mechanisms has been diminished by the development of more resilient voice/video coding techniques that are better suited for the best-effort Internet than the older coding techniques that were originally designed for circuit-switched networks.

One of the factors that has blunted the demand for QoS has been the transition of the Internet infrastructure from heavy congestion in the early 1990s, to overprovisioning in backbones and in many international links now. Thus, research in QoS mechanisms also has to include some careful attention to the relative costs and benefits of QoS in different places in the network. Applied research into QoS should include explicit consideration of economic issues of deploying and operating a QoS-enabled IP network [Clark02].

3.6.1. Inter-Domain QoS Architecture

Typically, a router in the deployed inter-domain Internet provides best-effort forwarding of IP packets, without regard for whether the source or destination of the packet is a direct customer of the operator of the router. This property is a significant contributor to the current scalability of the global Internet and contributes to the difficulty of deploying inter-domain Quality of Service (QoS) mechanisms.

Deploying existing Quality-of-Service (QoS) mechanisms, for example Differentiated Services or Integrated Services, across an inter-domain boundary creates a significant and easily exploited denial-of-service vulnerability for any network that provides inter-domain QoS support. This has caused network operators to refrain from supporting inter-domain QoS. The Internet would benefit from

additional research into alternative approaches to QoS, particularly into approaches that do not create such vulnerabilities and can be deployed end-to-end [RFC-2990].

Also, current business models are not consistent with inter-domain QoS, in large part because it is impractical or impossible to authenticate the identity of the sender of would-be preferred traffic while still forwarding traffic at line-rate. Absent such an ability, it is unclear how a network operator could bill or otherwise recover costs associated with providing that preferred service. So any new work on inter-domain QoS mechanisms and architectures needs to carefully consider the economic and security implications of such proposals.

3.6.2. New Queuing Disciplines

The overall Quality-of-Service for traffic is in part determined by the scheduling and queue management mechanisms at the routers. While there are a number of existing mechanisms (e.g., RED) that work well, it is possible that improved active queuing strategies might be devised. Mechanisms that lowered the implementation cost in IP routers might help increase deployment of active queue management, for example.

3.7. Congestion Control.

TCP's congestion avoidance and control mechanisms, from 1988 [Jacobson88], have been a key factor in maintaining the stability of the Internet, and are used by the bulk of the Internet's traffic. However, the congestion control mechanisms of the Internet need to be expanded and modified to meet a wide range of new requirements, from new applications such as streaming media and multicast to new environments such as wireless networks or very high bandwidth paths, and new requirements for minimizing queueing delay. While there are significant bodies of work in several of these issues, considerably more needs to be done.

We would note that research on TCP congestion control is also not yet "done", with much still to be accomplished in high-speed TCP, or in adding robust performance over paths with significant reordering, intermittent connectivity, non-congestive packet loss, and the like.

Several of these issues bring up difficult fundamental questions about the potential costs and benefits of increased communication between layers. Would it help transport to receive hints or other information from routing, from link layers, or from other transport-level connections? If so, what would be the cost to robust operation across diverse environments?

For congestion control mechanisms in routers, active queue management and Explicit Congestion Notification are generally not yet deployed, and there are a range of proposals, in various states of maturity, in this area. At the same time, there is a great deal that we still do not understand about the interactions of queue management mechanisms with other factors in the network. Router-based congestion control mechanisms are also needed for detecting and responding to aggregate congestion such as in Distributed Denial of Service attacks and flash crowds.

As more applications have the need to transfer very large files over high delay-bandwidth-product paths, the stresses on current congestion control mechanisms raise the question of whether we need more fine-grained feedback from routers. This includes the challenge of allowing connections to avoid the delays of slow-start, and to rapidly make use of newly-available bandwidth. On a more general level, we don't understand the potential and limitations for best-effort traffic over high delay-bandwidth-product paths, given the current feedback from routers, or the range of possibilities for more explicit feedback from routers.

There is also a need for long-term research in congestion control that is separate from specific functional requirements like the ones listed above. We know very little about congestion control dynamics or traffic dynamics of a large, complex network like the global Internet, with its heterogeneous and changing traffic mixes, link-level technologies, network protocols and router mechanisms, patterns of congestion, pricing models, and the like. Expanding our knowledge in this area seems likely to require a rich mix of measurement, analysis, simulations, and experimentation.

3.8. Studying the Evolution of the Internet Infrastructure

The evolution of the Internet infrastructure has been frustratingly slow and difficult, with long stories about the difficulties in adding IPv6, QoS, multicast, and other functionality to the Internet. We need a more scientific understanding of the evolutionary potentials and evolutionary difficulties of the Internet infrastructure.

This evolutionary potential is affected not only by the technical issues of the layered IP architecture, but by other factors as well. These factors include the changes in the environment over time (e.g., the recent overprovisioning of backbones, the deployment of firewalls), and the role of the standardization process. Economic and public policy factors are also critical, including the central fact of the Internet as a decentralized system, with key players being not only individuals, but also ISPs, companies, and entire

industries. Deployment issues are also key factors in the evolution of the Internet, including the continual chicken-and-egg problem of having enough customers to merit rolling out a service whose utility depends on the size of the customer base in the first place.

Overlay networks might serve as a transition technology for some new functionality, with an initial deployment in overlay networks, and with the new functionality moving later into the core if it seems warranted.

There are also increased obstacles to the evolution of the Internet in the form of increased complexity [WD02], unanticipated feature interactions [Kruse00], interactions between layers [CWWS92], interventions by middleboxes [RFC-3424], and the like. Because increasing complexity appears inevitable, research is needed to understand architectural mechanisms that can accommodate increased complexity without decreasing robustness of performance in unknown environments, and without closing off future possibilities for evolution. More concretely, research is needed on how to evolve the Internet while still maintaining its core strengths, such as the current degree of global addressability of hosts, end-to-end transparency of packet forwarding, and good performance for best-effort traffic.

3.9. Middleboxes

Research is needed to address the challenges posed by the wide range of middleboxes [RFC-3234]. This includes issues of security, control, data integrity, and on the general impact of middleboxes on the architecture.

In many ways middleboxes are a direct outgrowth of commercial interests, but there is a need to look beyond the near-term needs for the technology, to research its broader implications and to explore ways to improve how middleboxes are integrated into the architecture.

3.10. Internet Measurement

A recurring challenge is measuring the Internet; there have been many discussions about the need for measurement studies as an integral part of Internet research [Claffy03]. In this discussion, we define measurement quite broadly. For example, there are numerous challenges in measuring performance along any substantial Internet path, particularly when the path crosses administrative domain boundaries. There are also challenges in measuring protocol/application usage on any high-speed Internet link. Many of

the problems discussed above would benefit from increased frequency of measurement as well as improved quality of measurement on the deployed Internet.

A key issue in network measurement is that most commercial Internet Service Providers consider the particular characteristics of their production IP network(s) to be trade secrets. Ways need to be found for cooperative measurement studies, e.g., to allow legitimate non-commercial researchers to be able to measure relevant network parameters while also protecting the privacy rights of the measured ISPs.

Absent measured data, there is possibly an over-reliance on network simulations in some parts of the Internet research community and probably insufficient validation that existing network simulation models are reasonably good representations of the deployed Internet (or of some plausible future Internet) [FK02].

Without solid measurement of the current Internet behavior, it is very difficult to know what otherwise unknown operational problems exist that require attention, and it is equally difficult to fully understand the impact of changes (past or future) upon the Internet's actual behavioral characteristics.

3.11. Applications

Research is needed on a wide range of issues related to Internet applications.

Taking email as one example application, research is needed on understanding the spam problem, and on investigating tools and techniques to mitigate the effects of spam, including tools and techniques that aid the implementation of legal and other non-technical anti-spam measures [ASRG]. "Spam" is a generic term for a range of significantly different types of unwanted bulk email, with many types of senders, content and traffic-generating techniques. As one part of controlling spam, we need to develop a much better understanding of its many, different characteristics and their interactions with each other.

3.12. Meeting the Needs of the Future

As network size, link bandwidth, CPU capacity, and the number of users all increase, research will be needed to ensure that the Internet of the future scales to meet these increasing demands. We have discussed some of these scaling issues in specific sections above.

However, for all of the research questions discussed in this document, the goal of the research must be not only to meet the challenges already experienced today, but also to meet the challenges that can be expected to emerge in the future.

3.13. Freely Distributable Prototypes

U.S.'s DARPA has historically funded development of freely distributable implementations of various Internet technologies (e.g., TCP/IPv4, RSVP, IPv6, and IP security) in a variety of operating systems (e.g., 4.2 BSD, 4.3 BSD, 4.4 BSD, Tenex). Experience has shown that a good way to speed deployment of a new technology is to provide an unencumbered, freely-distributable prototype that can be incorporated into commercial products as well as non-commercial prototypes. Japan's WIDE Project has also funded some such work, primarily focused on IPv6 implementation for 4.4 BSD and Linux. [WIDE] We believe that applied research projects in networking will have an increased probability of success if the research project teams make their resulting software implementations freely available for both commercial and non-commercial uses. Examples of successes here include the DARPA funding of TCP/IPv4 integration into the 4.x BSD operating system [MBKQ96], DARPA/USN funding of ESP/AH design and integration into 4.4 BSD [Atk96], as well as separate DARPA/USN and WIDE funding of freely distributable IPv6 prototypes [Atk96, WIDE].

4. Conclusions

This document has summarized the history of research funding for the Internet and highlighted examples of open research questions. The IAB believes that more research is required to further the evolution of the Internet infrastructure, and that consistent, sufficient non-commercial funding is needed to enable such research.

In case there is any confusion, in this document we are not suggesting any direct or indirect role for the IAB, the IETF, or the IRTF in handling any funding for Internet research.

5. Acknowledgements

The people who directly contributed to this document in some form include the following: Ran Atkinson, Guy Almes, Rob Austein, Vint Cerf, Jon Crowcroft, Sally Floyd, James Kempf, Joe Macker, Craig Partridge, Vern Paxson, Juergen Schoenwaelder, and Mike St. Johns.

We are also grateful to Kim Claffy, Dave Crocker, Michael Eder, Eric Fleischman, Andrei Gurtov, Stephen Kent, J.P. Martin-Flatin, and Hilarie Orman for feedback on earlier drafts of this document.

We have also drawn from the following reports:
[CIPB02,IST02,NV02,NSF02,NSF03,NSF03a].

6. Security Considerations

This document does not itself create any new security issues for the Internet community. Security issues within the Internet Architecture primarily are discussed in Section 3.4 above.

7. Informative References

- [ASRG] Anti-Spam Research Group (ASRG) of the IRTF. URL ["http://asrg.sp.am/"](http://asrg.sp.am/).
- [Atk96] R. Atkinson et al., "Implementation of IPv6 in 4.4 BSD", Proceedings of USENIX 1996 Annual Technical Conference, USENIX Association, Berkeley, CA, USA. January 1996. URL <http://www.chacs.itd.nrl.navy.mil/publications/CHACS/1996/1996atkinson-USENIX.pdf>
- [Bellman1957] R.E. Bellman, "Dynamic Programming", Princeton University Press, Princeton, NJ, 1957.
- [Claffy03] K. Claffy, "Priorities and Challenges in Internet Measurement, Simulation, and Analysis", Large Scale Network meeting, (US) National Science Foundation, Arlington, VA, USA. 10 June 2003. URL ["http://www.caida.org/outreach/presentations/2003/lsn20030610/"](http://www.caida.org/outreach/presentations/2003/lsn20030610/).
- [Claffy03a] K. Claffy, "Top Problems of the Internet and What Sysadmins and Researchers Can Do To Help", plenary talk at LISA'03, October 2003. URL ["http://www.caida.org/outreach/presentations/2003/netproblems_lisa03/"](http://www.caida.org/outreach/presentations/2003/netproblems_lisa03/).
- [Clark02] D. D. Clark, "Deploying the Internet - why does it take so long and, can research help?", Large-Scale Networking Distinguished Lecture Series, (U.S.) National Science Foundation, Arlington, VA, 8 January 2002. URL: <http://www.ngi-supernet.org/conferences.html>

- [CSTB99] Computer Science and Telecommunications Board, (U.S.) National Research Council, "Funding a Revolution: Government Support for Computing Research", National Academy Press, Washington, DC, 1999. URL ["http://www7.nationalacademies.org/cstb/pub_revolution.html"](http://www7.nationalacademies.org/cstb/pub_revolution.html).
- [CIPB02] Critical Infrastructure Protection Board, "National Strategy to Secure Cyberspace", The White House, Washington, DC, USA. September 2002, URL ["http://www.whitehouse.gov/pcipb"](http://www.whitehouse.gov/pcipb).
- [CWWS92] J. Crowcroft, I. Wakeman, Z. Wang, and D. Sirovica, "Is Layering Harmful?", IEEE Networks, Vol. 6, Issue 1, pp 20-24, January 1992.
- [Diot00] C. Diot, et al., "Deployment Issues for the IP Multicast Service and Architecture", IEEE Network, January/February 2000.
- [Deering1988] S. Deering, "Multicast Routing in Internetworks and LANs", ACM Computer Communications Review, Volume 18, Issue 4, August 1988.
- [Dijkstra59] E. Dijkstra, "A Note on Two Problems in Connexion with Graphs", Numerische Mathematik, 1, 1959, pp.269-271.
- [FF1962] L. R. Ford Jr. and D.R. Fulkerson, "Flows in Networks", Princeton University Press, Princeton, NJ, 1962.
- [FK02] S. Floyd and E. Kohler, "Internet Research Needs Better Models", Proceedings of 1st Workshop on Hot Topics in Networks (Hotnets-I), Princeton, NJ, USA. October 2002. URL ["http://www.icir.org/models/bettermodels.html"](http://www.icir.org/models/bettermodels.html).
- [IM1993] J. Ioannidis and G. Maguire Jr., "The Design and Implementation of a Mobile Internetworking Architecture", Proceedings of the Winter USENIX Technical Conference, pages 489-500, Berkeley, CA, USA, January 1993.
- [IST02] Research Networking in Europe - Striving for Global Leadership, Information Society Technologies, 2002. URL ["http://www.cordis.lu/ist/rn/rn-brochure.htm"](http://www.cordis.lu/ist/rn/rn-brochure.htm).

- [Jacobson88] Van Jacobson, "Congestion Avoidance and Control", Proceedings of ACM SIGCOMM 1988 Symposium, ACM SIGCOMM, Stanford, CA, August 1988. URL ["http://citeseer.nj.nec.com/jacobson88congestion.html"](http://citeseer.nj.nec.com/jacobson88congestion.html).
- [Jackson02] William Jackson, "U.S. should fund R&D for secure Internet protocols, Clarke says", Government Computer News, 31 October 2002. URL ["http://www.gcn.com/voll_no1/security/20382-1.html"](http://www.gcn.com/voll_no1/security/20382-1.html).
- [Kruse00] Hans Kruse, "The Pitfalls of Distributed Protocol Development: Unintentional Interactions between Network Operations and Applications Protocols", Proceedings of the 8th International Conference on Telecommunication Systems Design, Nashville, TN, USA, March 2000. URL ["http://www.csm.ohiou.edu/kruse/publications/TSYS2000.pdf"](http://www.csm.ohiou.edu/kruse/publications/TSYS2000.pdf).
- [KLMS2000] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo, "Secure Border Gateway Protocol (S-BGP)", Proceedings of ISOC Network and Distributed Systems Security Symposium, Internet Society, Reston, VA, February 2000.
- [LD2002] E. Lear and R. Droms, "What's in a Name: Thoughts from the NSRG", expired Internet-Draft, December 2002.
- [MBFIPS01] Ratul Mahajan, Steven M. Bellovin, Sally Floyd, John Ioannidis, Vern Paxson, and Scott Shenker, "Controlling High Bandwidth Aggregates in the Network", ACM Computer Communications Review, Vol. 32, No. 3, July 2002. URL ["http://www.icir.org/pushback/"](http://www.icir.org/pushback/).
- [MBKQ96] M. McKusick, K. Bostic, M. Karels, and J. Quarterman, "Design and Implementation of the 4.4 BSD Operating System", Addison-Wesley, Reading, MA, 1996.
- [MGVK02] Z. Mao, R. Govindan, G. Varghese, & R. Katz, "Route Flap Dampening Exacerbates Internet Routing Convergence", Proceedings of ACM SIGCOMM 2002, ACM, Pittsburgh, PA, USA, August 2002.
- [NV02] NetVision 2012 Committee, "DARPA's Ten-Year Strategic Plan for Networking Research", (U.S.) Defense Advanced Research Projects Agency, October 2002. Citation for acknowledgement purposes only.

- [NSF02] NSF Workshop on Network Research Testbeds, National Science Foundation, Directorate for Computer and Information Science & Engineering, Advanced Networking Infrastructure & Research Division, Arlington, VA, USA, October 2002. URL "http://www-net.cs.umass.edu/testbed_workshop/".
- [NSF03] NSF ANIR Principal Investigator meeting, National Science Foundation, Arlington, VA, USA. January 9-10, 2003, URL "<http://www.ncne.org/training/nsf-pi/2003/nsfpimain.html>".
- [NSF03a] D. E. Atkins, et al., "Revolutionizing Science and Engineering Through Cyberinfrastructure", Report of NSF Advisory Panel on Cyberinfrastructure, January 2003. URL "http://www.cise.nsf.gov/evnt/reports/atkins_annc_020303.htm".
- [NSF03b] Report of the National Science Foundation Workshop on Fundamental Research in Networking. April 24-25, 2003. URL "<http://www.cs.virginia.edu/~jorg/workshop1/NSF-NetWorkshop-2003.pdf>".
- [Floyd] S. Floyd, "Papers about Research Questions for the Internet", web page, ICSI Center for Internet Research (ICIR), Berkeley, CA, 2003 URL "http://www.icir.org/floyd/research_questions.html".
- [RFC-1510] Kohl, J. and C. Neuman, "The Kerberos Network Authentication Service (V5)", RFC 1510, September 1993.
- [RFC-1633] Braden, R., Clark, D., and S. Shenker, "Integrated Services in the Internet Architecture: an Overview", RFC 1633, June 1994.
- [RFC-2082] Baker, F. and R. Atkinson, "RIP-2 MD5 Authentication", RFC 2082, January 1997.
- [RFC-2210] Wroclawski, J., "The Use of RSVP with IETF Integrated Services", RFC 2210, September 1997.
- [RFC-2154] Murphy, S., Badger, M., and B. Wellington, "OSPF with Digital Signatures", RFC 2154, June 1997.
- [RFC-2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.

- [RFC-2407] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998.
- [RFC-2501] Corson, S. and J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", RFC 2501, January 1999.
- [RFC-2990] Huston, G., "Next Steps for the IP QoS Architecture", RFC 2990, November 2000.
- [RFC-3221] Huston, G., "Commentary on Inter-Domain Routing in the Internet", RFC 3221, December 2001.
- [RFC-3234] Carpenter, B. and S. Brim, "Middleboxes: Taxonomy and Issues", RFC 3234, February 2002.
- [RFC-3424] Daigle, L. and IAB, "IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC 3424, November 2002.
- [RFC-3467] Klensin, J., "Role of the Domain Name System (DNS)", RFC 3467, February 2003.
- [RFC-3535] Schoenwaelder, J., "Overview of the 2002 IAB Network Management Workshop", RFC 3535, May 2003.
- [RFC-3387] Eder, M., Chaskar, H., and S. Nag, "Considerations from the Service Management Research Group (SMRG) on Quality of Service (QoS) in the IP Network", RFC 3387, September 2002.
- [RIPE] RIPE (Reseaux IP Europeens), Amsterdam, NL. URL "<http://www.ripe.net/ripe/>".
- [Savage00] Savage, S., Wetherall, D., Karlink, A. R., and Anderson, T., "Practical Network Support for IP Traceback", Proceedings of 2000 ACM SIGCOMM Conference, ACM SIGCOMM, Stockholm, SE, pp. 295-306. August 2000.
- [Schiller03] J. I. Schiller, "Interception Technology: The Good, The Bad, and The Ugly!", Presentation at 28th NANOG Meeting, North American Network Operators Group (NANOG), Ann Arbor, MI, USA, June 2003. URL "<http://www.nanog.org/mtg-0306/schiller.html>".

- [SM03] P. Sharma and R. Malpani, "IP Multicast Operational Network Management: Design, Challenges, and Experiences", IEEE Network, Vol. 17, No. 2, March 2003.
- [SMA03] N. Spring, R. Mahajan, & T. Anderson, "Quantifying the Causes of Path Inflation", Proceedings of ACM SIGCOMM 2003, ACM, Karlsruhe, Germany, August 2003.
- [WD02] Walter Willinger and John Doyle, "Robustness and the Internet: Design and Evolution", Unpublished/Preprint, 1 March 2002, URL
"http://netlab.caltech.edu/internet/".
- [WIDE] WIDE Project, Japan. URL "http://www.wide.ad.jp/".

8. Authors' Addresses

Internet Architecture Board
EMail: iab@iab.org

Internet Architecture Board Members
at the time this document was published were:

Bernard Aboba
Harald Alvestrand (IETF chair)
Rob Austein
Leslie Daigle (IAB chair)
Patrik Faltstrom
Sally Floyd
Mark Handley
Bob Hinden
Geoff Huston (IAB Executive Director)
Jun-ichiro Itojun Hagino
Eric Rescorla
Pete Resnick
Jonathan Rosenberg

We note that Ran Atkinson, one of the editors of the document, was an IAB member at the time that this document was first created, in November 2002, and that Vern Paxson, the IRTF chair, is an ex-officio member of the IAB.

Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/S HE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the IETF's procedures with respect to rights in IETF Documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

