Internet Engineering Task Force
INTERNET DRAFT
draft-iab-research-funding-01.txt

Ran Atkinson, Editor
Sally Floyd, Editor
Internet Architecture Board

29 June 2003

IAB Concerns & Recommendations Regarding Internet Research & Evolution

Status of this Memo

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.

Feedback can be sent to the IAB mailing list at iab@ietf.org, or to the editors at rja@extremenetworks.com and floyd@icir.org. Feedback can also be sent to the mailing list set up for feedback at research-funding@ietf.org. Requests to join can be sent to research-funding-request@ietf.org, with "subscribe research-funding" in the body of the request.

Table of Contents

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. [Add a citation.] 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. (E.g., issues related to the global grid are not discussed in this document because these issues are addressed through the Global Grid Forum and other grid-specific organizations, not in the IETF.)

Where at all possible, the examples in the paper 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

Recently, 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 [J02]: "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 were 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 & 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.

2.2. 1980s and early 1990s

The ARPAnet converted to the Internet Protocol 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 (US Department of Defense) funding (e.g. via DCA's Defense Data Network (DDN) program) and other U.S. Government funding (e.g. US Department of Energy (DoE) funding for research networks at DoE national laboratories, (US) 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, and NSFnet.

In the late 1980s, the U.S. DoD desired to leave the business of providing operational network services to academic institutions, so funding for many academic activities moved over to the NSF. NSF funding included research projects into networking, as well as creating the NSFnet backbone and sponsoring the creation of several NSF regional networks (e.g. SURAnet) and interconnections with several international research networks.

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. Furthermore, those increased funding levels should be sustained and protected against inflation 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.

3.1. Scope & Limitations

This document is NOT intended as a guide for funding organizations 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.

3.2. 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), and the Fully-Qualified Domain Name (FQDN). Many of the Internet's namespaces are supported by the widely deployed Domain Name System [RFC refs] 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 minimise the state on the DNS server side of the connection. UDP does this well, so 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 are 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. .COM), practical methods for incremental deployment of DNS security, and tools to make it easier to manage secure DNS infrastructure.

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 participants in the IRTF NSRG membership 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..

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 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.

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.

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 [RFC2385], OSPF MD5, RIP MD5 [RFC2082]) provide cryptographic authentication of routing protocol messages, but no authentication of the actual routing data. Current proposals (e.g. S-BGP [KLMS2000]) for improving this in inter-domain routing are unduly challenging to deploy across the Internet because of their reliance on a single trust hierarchy (e.g., a single PKI). Similar proposals (e.g. OSPF with Digital Signatures, [RFC2154]) for intra-domain routing are argued to be computationally infeasible to deploy in a large network.

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 merits additional study and consideration.

3.3.3. Routing Algorithms

The current Internet routing system relies primarily on only three algorithms. Link-state routing uses the Dijkstra algorithm [Dijkstra59]. The Distance-Vector and Path-Vector algorithms 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/commercial considerations of the routing system and in part because of inadequate consideration of security implications.

3.3.4. Mobile & Ad-Hoc Routing

Mobile routing [IM1993] and mobile ad-hoc routing [RFC2501] 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 AAA and 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'" [S03].

3.4.1. Freely Distributable Prototypes

U.S.'s DARPA has historically funded development of freely distributable implementations of various security technologies, such as IP security, in a variety of operating systems. Experience has shown that a good way to speed deployment of a new technology is to provide an unencumbered, freely-distributable prototype. We believe that applied research projects in Internet security 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] and DARPA/USN funding of ESP/AH design and integration into 4.4 BSD [Atk96].

3.4.2. 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, while there has been significant work into hierarchical security models (e.g. Bell-Lapadula) [BL1976], there has not been adequate formal study of alternative security models (e.g. PGP's Web-of-Trust model). Use of a hierarchical trust model creates significant limitations in how one might approach securing components of the Internet, for example the DNS, or the inter-domain routing system. So research to develop new trust models or on the applicability of existing non-hierarchical trust models to existing 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.

3.4.3. Key Management

A recurring challenge to the Internet community is how to design, implement, and deploy key management appropriate to the myriad 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 absence 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.

3.4.4 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.

3.4.5 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.

3.4.6. Deployment Considerations in Security

Lots of work has been done on theoretically perfect security that is impossible to deploy. Unfortunately, Kent's 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 centralising route advertisement policy enforcement in the Routing Information Registries or some similiar body. Historically, public-key infrastructures have been either very difficult or impossible to deploy at large scale. Some have recently suggested that the PGP web-of-trust authentication model should be applied to inter-domain advertisement of routing prefixes [Schiller03]. 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.

3.4.7. 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 malware (viruses, worms, etc.) 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. This makes the design of anti-DoS measures of high importance to the Internet. Some work has been done on this front [Sav00], [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 hierarchical monitoring of individual devices.

Unfortunately, network management research has historically been very underfunded, because it is difficult to get funding bodies to recognize this as legitimate networking research.

3.5.1. Configuration Management

Operators at the IAB Network Management Workshop [RFC-3535] held in 2002 reported that scalable distributed configuration management for sets of network devices is a significant challenge today. An enhanced network management architecture that more fully supports real operational needs is desirable. Even individual improvements in configuration management for sets of networked devices would be very welcome. Such improvements would need to include an integrated approach to security for the configuration data.

3.5.1. Enhanced Monitoring Capabilities

SNMP does not scale very 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 networks-as-a-whole.

3.5.2. Managing Networks, Not Devices

In particular, at present there are few or no good tools for managing a whole network of devices, though SNMP (Simple Network Management Protocol) and CMIP (Common Management Information Protocol) are fine for reading status of well-defined objects from individual boxes. Applied research into methods of managing sets of networked devices seems worthwhile. Ideally this configuration management approach would support distributed management, rather than being strictly hierarchical.

As an example, the current set of network management tools for managing multimedia (voice and video) IP networks is inadequate, and research would be useful in this area. The lack of appropriate network management tools has also been cited as one of the major barriers to the deployment of IP multicast [D00, SP03].

3.5.3. Improving the Scalability of 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 something else. Current approaches to network management do not scale sufficiently, so network service providers and enterprises often have difficulty operating their network(s) as successfully and economically as desired. Hence, more work is needed to improve the scalability of network management systems. This might involve application of artificial intelligence, expert systems technology, or other mechanisms, for example.

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]. 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.

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

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, 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 reasonably well, it is possible that improved 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 control mechanisms, from 1988 [J88], 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 stresses, 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.

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 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 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 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 [K00], 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.

3.9. Middleboxes

Research is needed to address the challenges posed by middleboxes. This includes issues of security, control, and 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 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 behaviour, 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 behavioural characteristics.

3.11. 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.

3.12. Additional topics

We have not included in this document discussions about the need for additional research in providing tools for researchers (e.g., modeling, simulations, test-beds).

Any new sections in this document should be focused on the problems that need to be addressed, rather than focused on the new approaches or technologies that might be promising answers to those problems.

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, we are not in this document 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, Rob Austein, Jon Crowcroft, Sally Floyd, James Kempf, Craig Partridge, Vern Paxson, and Mike St. Johns. We are also grateful to Kim Claffy, Andrei Gurtov, Hilarie Orman for feedback.

We have also drawn widely on the following sources: [CIPB02], [IST02], [NV02], [NSF02], [NSF03], [NSF03a].

Upcoming workshops include the following: [COST-NSF03].

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. IANA Considerations

There are no IANA considerations regarding this document.

Normative References

There are no Normative References because this is an Informational document.

Informative References

[Atk96] R. Atkinson et alia, "Implementation of IPv6 in 4.4 BSD", Proceedings of USENIX 1996 Annual Technical Conference, USENIX Association, Berkeley, CA, January 1996.

[Bellman1957] R.E. Bellman, "Dynamic Programming", Princeton University Press, Princeton, NJ, 1957.

[BL1976] D. E. Bell & L. J. LaPadula, "Secure Computer Systems: Unified Exposition and Multics Interpretation", MITRE Technical Report NMTR-1997 (ESD-TR-75-306), The Mitre Corporation, March 1976.

[Claffy03] K. Claffy, "Priorities and Challenges in Internet Measurement, Simulation, and Analysis", NSF PI meeting, January 2003. URL "http://www.caida.org/outreach/presentations/2003/nsfpi0301/".

[Clark02] D. D. Clark, "Deploying the Internet - why does it take so long and, can research help ?", Large-Scale Networking Distinguished Lecture Series, (US) National Science Foundation, Arlington, VA, 8 January 2002. URL: http://www.ngi-supernet.org/conferences.html

[CSTB99] Computer Science & Telecommunications Board, (US) 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".

[CIPB02] Critical Infrastructure Protection Board, "National Strategy to Secure Cyberspace", The White House, Washington, DC, September 2002, URL "http://www.whitehouse.gov/pcipb".

[COST-NSF03] COST-IST(EU)--NSF(USA) Workshop on Networking, June, 2003. URL "http://cgi.di.uoa.gr/~istavrak/costnsf/".

[CWWS92] J. Crowcroft, I Wakeman, Z. Wang, & D. Sirovica, "Is Layering Harmful ?", IEEE Networks, January 1992.

[Diot00] C. Diot, et alia, "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", Numerishe Mathematik, 1, 1959, pp.269-271.

[FF1962] L.R. Ford Jr. & D.R. Fulkerson, "Flows in Networks", Princeton University Press, Princeton, NJ, 1962.

[FK02] S. Floyd and E. Kohler, Internet Research Needs Better Models, Hotnets-I. October 2002. URL "http://www.icir.org/models/bettermodels.html".

[Handley02] Mark Handley's viewgraphs to an NSF meeting, 2002.

[IM1993] J. Ioannidis & G. Maguire Jr., "The Design and Implementation of a Mobile Internetworking Architecture", Proceedings of the Winter USENIX Technical Conference, pages 489-500, 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".

[J88] Van Jacobson, Congestion Avoidance and Control, SIGCOMM, 1988. URL "http://citeseer.nj.nec.com/jacobson88congestion.html".

[J02] William Jackson, "U.S. should fund R&D for secure Internet protocols, Clarke says", 10/31/02, URL "http://www.gcn.com/vol1_no1/security/20382-1.html".

[K00] Hans Kruse, The Pitfalls of Distributed Protocol Development: Unintentional Interactions between Network Operations and Applications Protocols, 8th International Conference on Telecommunication Systems Design, Nashville, March 2000. URL "http://www.csm.ohiou.edu/kruse/publications/TSYS2000.pdf".

[KLMS2000] S. Kent, C. Lynn, J. Mikkelson, & K. Seo, "Secure Border Gateway Protocol (S-BGP)", Proceedings of ISoc Network & Distributed Systems Security Symposium, Internet Society, Reston, VA, February 2000.

[LD2002] E. Lear & R. Droms, "What's in a Name: Thoughts from the NSRG", 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 (Extended Version), July, 2001. URL "http://www.icir.org/pushback/".

[MBKQ96] M. McKusick, K. Bostic, M. Karels, & J. Quarterman, "Design and Implementation of the 4.4 BSD Operating System", Addison-Wesley, Reading, MA, 1996.

[S03] J. I. Schiller, "Interception Technology: The Good, The Bad, and The Ugly!", NANOG, June 2003. URL "http://www.nanog.org/mtg-0306/schiller.html".

[Schiller03] J. I. Schiller, Private Communication, MIT, Cambridge, MA. 2003.

[NV02] NetVision 2012 Committee,"DARPA's Ten-Year Strategic Plan for Networking Research", (US) Defence Advanced Research Projects Agency, October 2002. Citation for acknowledgement purposes only.

[NSF02] NSF Workshop on Network Research Testbeds, October 2002. URL "http://www-net.cs.umass.edu/testbed_workshop/".

[NSF03] NSF ANIR Principal Investigator meeting, January 9-10, 2003, URL "http://www.ncne.org/training/nsf-pi/2003/nsfpimain.html".

[NSF03a] Revolutionizing Science and Engineering Through Cyberinfrastructure, NSF Report, January 2003. URL "http://www.cise.nsf.gov/evnt/reports/atkins_annc_020303.htm".

[ResearchQuestions] Web Page on "Papers about Research Questions for the Internet", URL "http://www.icir.org/floyd/research_questions.html".

[RFC-1510] J. Kohl & C. Neuman, "The Kerberos Network Authentication Service (V5)", RFC 1510, September 1993.

[RFC-2082] F. Baker & R. Atkinson, "RIPv2 MD5 Authentication", RFC-2082, January 1997.

[RFC-2154] S. Murphy, M. Badger, & B. Wellington, "OSPF with Digital Signatures", RFC-2154, June 1997.

[RFC-2385] A. Heffernan, "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC-2385, August 1998.

[RFC-2407] D. Piper, "The Internet IP Security Domain of Interpretation for ISAKMP", RFC-2407, November 1998.

[RFC-2501] S. Corson & J. Macker, "Mobile Ad Hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations", RFC-2501, January 1999.

[RFC-2990] G. Huston, "Next Steps for the IP QoS Architecture", RFC-1990, November 2000.

[RFC-3221] G. Huston, "Commentary on Inter-Domain Routing in the Internet", RFC-3221, December 2001.

[RFC-3424] L. Daigle (Ed.), "IAB Considerations for Unilateral Self-Address Fixing (UNSAF) Across Network Address Translation", RFC-3424, Internet Architecture Board, November 2002.

[RFC-3535] J. Schoenwalder, Editor, "Overfiew of the 2002 IAB Network

Management Workshop", RFC-3535, May 2003.

[RIPE] RIPE (Reseaux IP Europeens), URL "http://www.ripe.net/ripe/".

[Sav00] Savage, S., Wetherall, D., Karlink, A. R., and Anderson, T., Practical Network Support for IP Traceback, SIGCOMM 2000.

[SP03] P. Sharma & R. Malpani, "IP Multicast Operational Network Management: Design, Challenges, and Experiences", IEEE Network, March 2003.

[WD02] Walter Willinger and John Doyle, "Robustness and the Internet: Design and Evolution", 2002, URL "http://netlab.caltech.edu/internet/".

[WIDE] WIDE Project, URL "http://www.wide.ad.jp/".

9. 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
Jun-ichiro Itojun Hagino
Mark Handley
Geoff Huston (IAB Executive Director)
Charlie Kaufman
James Kempf
Eric Rescorla
Mike St. Johns

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

This draft was created in November 2002 and revised January 2003, February 2003, and June 2003.