

Developing a Next-Generation Internet Architecture

Robert Braden, David Clark, Scott Shenker, and John Wroclawski

July 15, 2000

Abstract

This document discusses goals and directions for a research effort aimed at developing a next-generation Internet architecture.

1 Introduction

The term "network architecture" is commonly used to describe a set of abstract principles for the technical design of protocols and mechanisms for computer communication. A network architecture represents a set of deliberate choices out of many design alternatives, where these choices are informed by an understanding of the requirements [CerfKahn74, Clark88]. In turn, the architecture provides a guide for the many technical decisions required to standardize network protocols and algorithms. The purpose of the architecture is provide coherence and consistency to these decisions and to ensure that the requirements are met.

The design of today's Internet technology was guided by an Internet architecture that was developed in the 1970s, under the Internet research program of the Defense Advanced Research Projects Agency (DARPA) of the US Department of Defense. Current reality and changing requirements are eating away at the viability of the original Internet architecture. Much of the coherence of the original architecture is being lost in a patchwork of technical embellishments, each intended to satisfy a particular new requirements.

We believe that it is now time to revisit the Internet architecture, to determine whether it can be changed to align better with current and future requirements. Many of the new requirements, present and future, are beginning to be apparent. It is our hope that a combination of abstract reasoning and practical experimentation can bring a new clarity to the architectural issues of the Internet.

For example, a new architecture might be designed to create greater functionality, generality, adaptability, and/or robustness in the future Internet. On the other hand, without a new long-term technical road map the Internet is likely to become decreasingly effective, often failing to meet the demands placed on it by applications.

The goal of a new-arch research program should be to define and prototype a future architecture towards which the Internet technology can evolve. A new-arch research program would differ from most other network research in the breadth of its viewpoint and in its relatively long time scale. Many Internet R&D projects are building upon the original Internet architecture, but few are looking at the basic architecture and requirements for a future Internet. The vendors who are largely driving the current technical evolution of the Internet do not believe they “own” the problem of the high-level design. Already, some individual ideas for architectural changes are floating around, and the new-arch goal would be to at least find a coherent, consistent synergy of the most useful ideas.

Using the experimental paradigm that worked well for the original Internet research program, the new-arch project should combine top-down reasoning about architectural principles with prototyping and simulation to explore the technical consequences of the new architecture. This combination of abstract reasoning and practical experimentation is not new; it will recapture the approach of the original DARPA Internet research program, still embodied in the “running code” mantra of the IETF.

The next section defines the problem space, exploring the meaning of “network architecture” and the gap between the original Internet architecture and current reality. Section 3 begins to discuss issues to be considered in the development of a new architecture. In particular, Sections 3.1 and 3.2 contain initial thoughts about the changing requirements and the design of a new architecture, respectively. Section 4 briefly comments on a possible new-arch research project organization and process.

2 Defining the Problem Space

2.1 What is Network architecture?

It is important to understand what the term “architecture” term means. The notion of network architecture was introduced during the Internet research phase by the research community that had developed the ARPAnet protocols. This community brought to bear on the computer communication problem the kind of abstract thinking about resources and relationships that came naturally to computer scientists who had designed hardware architectures and/or computer software systems.

This resulted in the development of an overarching “design philosophy” to accompany the design of the algorithms and protocols for the Internet (see for example [CerfKahn74, CerfKirs78, E2Earg81]). This philosophy was elaborated over time to create the complete original architecture of the Internet protocol suite [Clark88].

Network architecture is a set of high-level design principles that guides the technical design of the network, especially the engineering of its protocols and algorithms. To flesh out this simple definition, we have examples of the constituents of the architecture and how it is applied. A network architecture must typically specify:

- Where and how state is maintained and how it is removed.
- What entities are named
- How naming, addressing, and routing functions inter-relate and how they are performed.
- How communication functions are modularized, e.g., into “layers” to form a “protocol stack”.
- How network resources are divided between flows and how end-systems react to this division, i.e., fairness and congestion control.
- Where security boundaries are drawn and how they are enforced.
- How management boundaries are drawn and selectively pierced.
- How differing QoS is requested and achieved.

Ideally one would like to imagine using the architecture to “generate” the technical design, but making such a mapping in a mechanical fashion is clearly impossible. The architecture can only provide a set of abstract principles against which we can check each decision about the technical design. The role of the architecture is to ensure that the resulting technical design will be consistent and coherent – the pieces will fit together smoothly – and that the design will satisfy the requirements on network function associated with the architecture.

An architecture is more general than a particular conformant technical design. The technical design derived from a particular architecture is far from unique, and it may evolve over time in response to detailed changes in requirements; however, the same architecture may be maintained. Crucially, an architecture is expected to be relatively long-lived, applicable to more than one generation of the technology. A trivial example of this is IPv6, which represents a later generation of technology than IPv4 although both conform to the same Internet architecture.

An example of architectural thinking (which was, however, a speculative alternative to aspects of the current Internet architecture) was contained in the Application Level Framing (ALF) proposal of Clark and Tennenhouse [ALF90]. ALF was not a complete architecture, only an architectural component with specific goals: lower cost and more flexible implementation, more efficient operation over diverse infrastructure (packet and ATM), and effective support for a wider range of application requirements, and so on. ALF illustrates the practical short-term benefit of long-term architectural thinking. While the ALF idea as proposed was not cast as an incremental modification to the Internet but rather as a new and different approach, other researchers used the ALF idea as the basis for the implementation of new applications over the existing network.

2.2 Example: the Internet Architecture

To further clarify the meaning of “architecture”, we can consider the Internet architecture. The important features of the original Internet architecture [BCarp96] included:

- A connectionless packet-forwarding infrastructure (“dumb network”) that positioned higher-level functionality at the edge of the network for robustness (“fate-sharing”)

- A single least-common-denominator data delivery service at the inter-network layer, with different end-to-end services implemented in the transport (or application) layer above. This design supports both reliable stream and (unreliable) datagram service across the same connectionless infrastructure.
- Addresses that are fixed-size numerical quantities, with a simple (net, host) hierarchy
- Addresses that are applied to physical network interfaces, which can therefore be overloaded for both naming a node and for routing to it.

This Internet architecture evolved during the research phase, and it has continued to evolve.

- The most important change during the 1974 - 1980 research period was the separation of TCP into an inter-network (IP) layer and a transport layer (TCP), in version 4 of the protocols.
- Several important features were added to the architecture during the early 1980s, especially subnetting, autonomous systems, and the domain name system (DNS). All of these changes reflecting an increasing understanding and respect for the issues of scale, imposed hierarchical design on the architecture. Subnetting was the first step towards making addresses hierarchical, and this was extended later with the addition of classless addressing (CIDR). Autonomous systems made the organization of routing hierarchical, by distinguishing EGPs from IGPs. Finally, the DNS introduced hierarchical naming.
- Later, IP multicasting added logical addressing and multi-destination delivery fundamental parts of the architecture.
- Congestion control using packet loss as a congestion signal and an additive increase/multiplicative decrease algorithm at the end-systems was added in the late 1980s in response to congestion collapse events.

2.3 Why is Architecture Important?

Architecture sets a sense of direction. It provides a basis for individual technical decisions, as opposed to a random walk starting from today's design. An architec-

ture represents not just an abstract goal, but a bias that is applied to every decision made day-to-day. While architecture is long-term, it will often invariably have short-term implications. In general, long-range research usually has short-term spinoffs, but short-range research with no long view of where it is going cannot have a long-term impact.

Because design done in the presence of architectural guidelines tends to produce more general solutions, it can result in “platform” solutions – service layers that are able to support a range of applications. This generality has the potential to reduce overall the level of investment, since investment in a platform solution can be reused. However, architecture has to be sustained and grown. Left to itself, it erodes and atrophies. “Platform” solutions tend over time to be supplanted by “overlays” that supplant the nominal interface for emerging classes of applications. This point is clear to commercial entities that support interfaces that they view as “platform” solutions. The Microsoft Windows API is a good example of a “platform” interface that tries to support a range of applications over a common interface. Microsoft commits a lot of resources to sustaining their interface, to attempt to get new applications to run over it directly, as opposed to running over some application-specific overlay. The Internet, if it is to survive as a platform, has to make the same kind of investment; for the Internet, this investment is in design of the architecture.

The Internet architecture is of particular importance to the military. In order that the DoD be able to use COTS technology for military needs, there must be a degree of generality in the devices, so that the differences between military and commercial objectives can be accommodated. To the extent that an overall architecture facilitates more general solutions, it directly enhances the COTS procurement objective.

Architecture guides technical development such as protocol design in a consistent direction; without adequate architectural principles, designers with short-term goals often implement point solutions. In general, an increasing collection of point solutions leads over time leads to a design that is complex, tangled, and inflexible.

The Internet is exhibiting this technical arteriosclerosis today. Extensions to the technical design of the Internet have been developed are being developed in the IETF at increasing rate. Examples include integrated service (IntServ) and differentiated service (DiffServ), IP Security (IPSEC), firewalls, mobile IP, Network Address Translator (NAT) devices, label switching, VPNs, and Web caches. We believe that certain of these extensions were created with some architectural sensitiv-

ity; IntServ [IntServ94], and DiffServ [DiffServ98XX], and IPSEC [IPSEC98XX] are examples. However, many of the recent extensions have been developed as point solutions for specific requirements of particular subsets of the Internet community – vendors, users, or ISPs – and largely in the absence of architectural thinking. They do not represent enhancements of the architecture, even though they were developed to meet legitimate short-term needs and requirements. Viewed as a whole, this situation demonstrates two types of negative effects: those that arise when technical development does not follow an architectural framework, and those that arise when an architecture becomes “tired” and requires rethinking in response to new requirements..

We can pick some illustrative examples from this list of extra-architectural extensions.

- IPSEC

IP Security (IPSEC) seems to be fully consistent with the original architectural model with respect to information hiding: it encrypts the transport-layer (UDP or TCP) headers that, which contain only end-to-end information, in principle of concern only to end hosts. However, network providers are very unhappy about losing the ability to “snoop” at port numbers and higher-layer protocol headers. This suggests that the original layering model was not in fact the right one, or at least it is no longer fully adequate for today’s requirements. Nobody owns the job of revisiting the general principle here, to provide a clean solution combining end-to-end security with the management information needed by the providers.

- NAT Devices

An architectural principle of the Internet was that addresses in packets are carried unchanged from source to destination. This principle is directly violated by the Network Address Translation (NAT) devices that have become popular. These devices were conceived to deal with exhaustion of the IP address space (an issue never considered in the original design), as a means to reduce the effort of address reassignment, and as a possible security enhancement.

However, NAT devices have significant high-level implications. Among the strongest of these is to move the Internet fundamentally away from a peer-to-peer model (in which any host is theoretically addressable) towards a client-server model (in which only “selected” hosts are addressable from the global network). This change appears relatively innocuous with today’s emphasis

on web applications, but has tremendous impact on the Internet's ability to support future applications and devices. This is an example of an "effective" solution to a point problem greatly restricting generality and future usability.

The example also demonstrates the difficulty of adapting an architecture "on the fly" to new circumstances. There is no accepted resolution of the architectural implications of NAT devices. At the same time that purists who defend the "old architecture" decry these devices as tasteless and destructive, advocates call them the wave of the future. Even after an ad hoc committee tried to catalog all the implications of address rewriting, there was still no community consensus about the correct set of design principles.

- Firewalls

Security firewalls have been introduced into the Internet design to protect regions of the network from outside attack. Firewalls violate the original architecture, and firewall design occurs in an undefined space of ad hoc evolution. There has been no overall agreement on the assumptions an end-node can or must make about what a firewall may do – what sorts of actions are acceptable. As a result, firewalls often result in a loss of end-to-end (i.e., user to server) functionality.

More important from the network perspective, firewalls and similar non-architectural changes often lead to unfortunate unintended consequences in precisely the functional areas they are attempting to address. A recent example arose when certain host vendors, acting in response to sites deploying overly aggressive firewalls that broke the standard TCP "MTU Discovery" protocol, developed and deployed a new version of MTU discovery that unwittingly enabled a serious network wide denial of service attack.

- Label Switching

Label switching represents an attempt to extend the architectural thinking to some extent, by asking what sort of layer two technology is actually needed to support the IP interface. But we should ask whether, given MPLS and IPSEC, the function of the IP layer needs rethinking.

In addition to their direct consequences, these recent extensions have problems with "feature interactions". They have complex interactions with each other and with the rest of the architecture, leading to increasingly complex and baroque protocol engineering. Neither the technical nor the architectural implications of these design extensions are fully understood today. Here are some examples of such undesirable feature interactions.

1. NAT devices are incompatible with IPSEC encryption. This makes deployment of good network security at an appropriate level almost impossible.
2. Transparent web caches are incompatible with Secure Sockets Layer (SSL) or IPSEC authentication.
3. Diagnostic tools such as traceroute often give misleading results because some traffic is being intercepted by transparent network entities such as web caches.
4. NAT devices need to look inside control protocols such as the FTP control channel to modify addresses and ports that are specified.
5. Deployment of new application protocols is made difficult by firewalls, with the result that new protocols are sometimes tunneled over inappropriate existing protocols to allow firewall traversal.

As a result, current IETF protocol engineering is in a conceptual and technical muddle that will continue to lead to increasingly complex and costly engineering as well as loss of functionality. The only way to avoid this degeneration will be to restore coherence to the architecture, and that will happen only as the result of a deliberate effort to create a new architecture for the future.

The cause of the IETF's muddle is not that its members do not believe in the architecture abstraction. Basic Internet architecture concepts like the end-to-end argument, connectionless protocols, and universal connectivity are often cited in IETF discussions. IETF members accept and support the idea of architecture. Indeed, there are IETF efforts to define complete sub-architectures – the “security architecture”, the “routing architecture”, the “addressing architecture”, etc. The problem with these well-intentioned and often high-quality efforts is that they are piecemeal. The union of these sub-architectures does not form an Internet architecture, which is greater than the union of its parts. A decomposition may be useful, but it is insufficient; the requirements and principles need to be combined into a coherent whole. Finally, the IETF is unsuited by membership, organization, or mission to attempt a major overhaul of the overall architecture.

2.4 The Original Requirements

Network architecture could be considered to lie in the middle between specific technical design and the overall set of technical and non-technical requirements

for the network. Therefore, a fundamental requirement of a new-arch research effort must be the choice of high-level requirements and goals. The choice of requirements may be the most critical issue determining the ultimate usefulness of the new architecture that is developed.

As an example, the following list is a brief summary of the requirements underlying the original Internet architecture (quoted from [Clark88]). This list was ordered with the most important requirements first.

1. Internetworking: existing networks must be interconnected.
2. Robustness: Internet communication must continue despite loss of networks or [routers].
3. Heterogeneity: The Internet architecture must accommodate a variety of networks
4. Distributed management: The Internet architecture must permit distributed management of its resources
5. Cost: The Internet architecture must be cost effective.
6. Ease of Attachment: The Internet architecture must permit host attachment with a low level of effort.
7. Accountability: The resources used in the internet architecture must be accountable.

3 Developing A New Architecture

The previous section explained the importance of architecture and introduced evidence for a serious “architecture gap” in the current Internet. A new architecture should lead to greater functionality, lower costs, and increased adaptability for all types of communication.

3.1 New Architecture Requirements

The development of an architecture must be guided in part by an understanding of the requirements to be met. It is therefore vital to articulate a set of goals and re-

quirements. The technical requirements for the Internet have changed considerably since 1975, and they will continue to change. A new requirements list will form an important part of the results from the proposed research. This will be based on (1) the changing requirements for the Internet, and (2) the ways in which the Internet technology has tended to stray from the original architecture, reflecting tensions between design and reality.

The relationship between requirements and architecture is not simple. While major requirements arise from non-technical issues in the real world – e.g., business models, regulatory models, and politics – other requirements are themselves the product of earlier technical decisions, i.e., depend upon the architecture. As a result, a new-architecture design effort cannot be completely top-down. There is not likely to be a unique answer for the list of requirements, and every requirement has some cost. The cost of a particular requirement may become apparent only after exploration of the architectural consequences of meeting that objective, in conjunction with the other objectives. It therefore requires an iterative process, in which requirements can be reexamined and perhaps promoted or demoted during the effort.

A second crucial point is that with the transition of the Internet from research project to mainstream infrastructure the range of applicability of the requirements must be much broader. This implies that fewer and fewer of the requirements will be truly global - applying with the same importance everywhere. Many of the requirements that the architecture must meet will apply with different force, or not at all, in some situations and portions of the network.

This makes the development of a single ordered list of requirements, as was done to motivate the original Internet research program, deeply problematic. Instead, a new Internet architecture must deal with a multi-ordered requirements set; with many requirements taking on different importance at different times, and in different regions of the network. It seems likely that such a “meta-requirement” will have a significant impact on the technical architecture. We touch briefly on one possible strategy for addressing it below; in any case we believe that meeting this need represents one of the most challenging aspects of designing a new architecture.

The commercialization of the Internet has led to many of the new requirements. An architecture of tomorrow must take into account the needs and concerns of commercial providers if it is to be accepted and thus to be able to influence overall direction. Examples of these concerns include (1) a framework for policy controls on inter-provider routing, (2) recognition that service providers need some ability

to see parts of the header for purposes of traffic planning, regulation of usage, etc., and (3) support for a variety of payment models for network usage.

For example, since today there is no way to assign a “value assertion” to traffic flows, there is no way to determine “settlements” by observing traffic patterns. One can count packets, but this does not indicate which end paid to have them sent. One of the motivations for some of the overlay delivery mechanisms that have recently been built over the Internet today, including the Akamai and RealAudio delivery infrastructure, is that they implement a specific payment model (sender pays), so that a class of users who match that value equation can associate themselves with this service.

Internet requirements continue to change. Some important new requirements that may influence the new architecture are as follows.

- Mobility

The Internet architecture should support flexible, efficient, highly-dynamic mobility.

- Policy-driven Auto-Configuration

The Internet architecture should provide auto-configuration of end systems and routers, subject to policy and administrative constraints.

- Highly time-variable resources

The Internet architecture should support resources that are highly variable over short time-scales. This may for example be due to switched backbone links, or due to mobile devices that can switch physical transmission medium as the node moves.

- Allocation of Capacity

An architecture of tomorrow must give users and network administrators the ability to allocate capacity among users and applications. In today’s Internet, allocation occurs implicitly as a result of congestion control. The goal has generally been some approximation of “fairness”; all slow down together, but this is not always the right model. For commercial activities, there is a desire to allocate capacity based on willingness to pay. For operational government activities, e.g., disaster response, there is a need to allocate capacity based on priority of task. It is not (always) the role of the network to tell the user how fast to go. The administrator should be able to ask the network

for resources, and the network should be able to inform the user if it cannot meet the requests due to resource limitations.

- Extremely long propagation delays

This requirement arises particularly in the proposed Interplanetary Internet, using the Internet technology for NASA's planetary exploration program. It is an extension of the more traditional "high bandwidth-delay product" requirement; reflecting the fact that both delay itself and delay-bandwidth interactions complicate the architecture of a network.

This discussion has dealt with technical requirements, but it is important to note that there are significant non-technical drivers on Internet design. There are obvious commercial drivers, as network providers learn how to make a profit from the Internet. Increasingly, there are also legal and public policy drivers, including intellectual property law, encryption export law, police surveillance, privacy and free speech, telecommunications laws, charging, and taxation. These are all subject to national variation, since the Internet is worldwide. We must be aware of these issues, but our job is to concentrate on the technical requirements within this broader context. We note that the proposing team brings both Internet industry structure and economics research credentials and pragmatic experience within the changing Internet environment to the table, and has a track record of effectiveness working within this space.

3.2 New Architecture Design

The proposed new-arch research team would develop a fresh design for the Internet architecture. This effort would not start from scratch, discarding all Internet principles and technology. As in any good science, this research should work from established principles as much as possible.

While it should retain what is useful from the present Internet technology, this research should consider backwards compatibility issues to be of relatively low priority; there is no other way to reach the goal of greater architectural coherence, or to define and validate a future architecture towards which the Internet technology can evolve. For example, it may propose an architectural change that would ultimately result in every router and host in the world being changed. To make progress, the project must take an optimistic viewpoint about future development of bridging technology to allow the existing Internet to move in the new technical

directions indicated by the new architecture. Development of such bridging technology is an engineering task that would come later. Furthermore, if it is possible to pick a future point in the architectural design space, it should be possible to use that sighting point to provide an improved framework for nearer-term technical decisions.

The development of a new Internet architecture is likely to include the following components.

- Examination of the areas in which the original architecture is known to have failed.
- Examination of the changed and changing requirements.
- Exploration and development of some proposed new architectural changes that have already been suggested to meet these requirements (a particular example is discussed below.)
- Exploration of possible new meta-principles for an architecture (a particular example is presented below).
- Outline of several candidate architectures.
- Consultation with experts in relevant technical areas, such as mobility, economic models, and embedded computing.
- Agreement on a single candidate for further exploration.
- Implementation of a proof-of-concept environment sufficient to evaluate the candidate architecture through experiments and simulation.
- Iteration based on feedback from proof-of-concept experiments.

3.2.1 New Architectural Principles

To give the flavor of the work that a new-arch effort might perform, we can look briefly at two past examples of architectural “thought experiments”. The first was Application Layer Framing or ALF [ALF90ref]. ALF was proposed as a new architectural component with specific goals – lower cost and more flexible implementation, more efficient operation over diverse infrastructure (packet and ATM),

effective support for a wider range of application requirements, and so on. While the ALF idea as proposed was not cast as in incremental modification to the Internet architecture but rather as a new and different approach, other researchers used the ALF idea as the basis for the implementation of new applications over the existing network. It could be reconsidered as a part of a new architecture.

One of the most basic and famous principles of the Internet architecture is the end-to-end argument [E2EArg90], which justifies the Internet design of “smart” end systems and a “dumb” network. However, many modern developments – e.g., firewalls, NAT devices, Web caches, traffic shapers and profile meters – violate the end-to-end argument. Some thought is being given to new architectural approaches that essentially modify the end2end end-to-end argument to fit the new reality. For example, Clark’s talk “A New Layer in the Reference Model” [Clark98] suggested as a thought experiment that the Internet be composed of regions that could be diverse along one or more dimensions – e.g., addressing, trust assumptions, or performance assumptions. Working through the consequences of this new principle led to the definition of a new building block – an “isolator” – to connect regions, and an additional layer in the protocol stack – a “Trust Enforcement Layer”. This sort of reconsideration of the end-to-end argument is very likely to play a significant role in our work.

3.2.2 New Meta-Architectural Principle

The architectural principle of the preceding section is a special case of a possible new approach to the Internet architecture. This new approach would recognize the fundamental importance of heterogeneity in the architecture. This would be analogous to the extensions already made in the current architecture as a result of the recognition of the importance of scale. One objective of this potential new architectural approach is to specifically acknowledge the variability of requirements within different regions of the network. Another objective is to more effectively support technological evolution, and to allow different regions of the Internet to be specialized towards potentially conflicting objectives.

The goal of the meta-principle suggested here is to explicitly minimize the degree of required global architectural consistency within the Internet. The principle recognizes that a complete architecture includes sub-architectures, i.e., design principles, for a number of conceptually-independent issues such as congestion control, security, or QoS. Suppose that the Internet were divided into regions (we avoid the

term “domain” to avoid overloading), in which different sub-architectures could be used. Thus, one region might take one (sub-)architectural approach to congestion control, while another might take another approach. There would still be a global architecture that established rules for end-to-end communication across a concatenation of regions. This over-arching global architectural principle would define the minimal set of globally agreed mechanisms, leaving the rest of the architecture to be defined by sub-architectures components that could vary from region to region within a well-defined set of functional rules.. For example, the addressing model might be a global property, consistent everywhere, while different regions might use widely different data transport models, and congestion control mechanisms might be different depending on whether the transport layer in a region was packet-based or circuit-switched.

The major technical challenge of such an approach would be to allow as much regional flexibility as possible while still supporting the definition of end-to-end services. It is a research issue to determine the extent to which this approach is feasible. It is attractive because heterogeneity is clearly second only to scale as an important meta principle for the Internet.

4 Project Organization

There is no commercial provider who believes that they hold the responsibility for the Internet architecture. Because the Internet is open, this is a classic role for a neutral party such as a government-funded research effort. Such an effort could easily be justified because the government is itself a sophisticated user that directly benefits from the generality of a coherent architecture and the platform that results. Besides the broad commercial Internet, an improved architecture will also benefit users with specialized needs, such as military or civilian disaster response, by facilitating the availability of suitable COTS products.

A new-arch research project could proceed in two phases. During the first phase, a team of individuals knowledgeable about network architecture would develop a new prioritized requirement list and an architecture to support those requirements. The second phase would be concerned with proof of concept; it would elaborate, explore, and demonstrate the designs using an appropriate combination of simulation, prototyping, and testing. These two phases would partially overlap and form a loop, so that the abstract reasoning required for architecture design will be grounded in experimentation.

4.1 Design Team

We believe that the basic architecture design can only be accomplished by a small team of people, ideally in the range of 6 to 10. If the results of this phase are promising, the team will develop an organizational as well as intellectual framework within which a larger group of people can become involved with the design and development of specific protocols and algorithms. In the most optimistic scenario, the effort will eventually take on a life of its own, and DARPA, NSF, and industrial funding will be used to follow up on promising new research directions that are suggested by the draft architecture.

No individual researcher possesses divine knowledge, and a small group of must be a very small subset of today's large and competent Internet R&D community. However, we believe that this project could make a contribution by working in a small group that has explicitly taken a longer range view. The quality and breadth of relevant experience of the core team are also important.

Once a draft architecture is completed, it will be useful to involve a larger subset of the research community. The team may catalyze this process by organizing workshops devoted to discussion of the Internet architecture. This task might be undertaken by a research group of the IRTF (Internet Research Task Force).

4.2 Proof-of-Concept for New Architecture

The major end result of this project is expected to be sets of abstractions: design principles, requirements, and objectives. However, past experience has shown the limitations of protocol suite designs that are essentially top-down. To ensure that the conceptual work on a new architecture does not become an idle exercise, the project includes a major proof-of-concept component. This activity will build and test prototype protocols that are consistent with the architecture, using an appropriate combination of experimental code and simulation. This reduction of theory to practice will provide feedback to the architectural design, revealing many secondary technical issues that arise from the architectural abstractions.

Clearly, the limited scope of the proposed project will not allow it to produce a complete prototype of a new protocol suite, should that be called for by the new architecture. Indeed, the purpose of a new architecture is to guide a large volume of further protocol engineering by researchers, vendors, and other members

of the Internet community; this will certainly not be accomplished directly under the proposed project. Rather, the research project would produce protocol code to serve as a proof of concept for particular important elements of the new design, using existing Internet protocol stacks wherever the resulting deviation from the new architecture is unimportant.

The objective of this proof-of-concept effort is somewhat analogous to the working models of inventions that used to be required by the US Patent Office. To patent a new mechanism, it was not necessary to produce a complete instance of a machine using that mechanism, it was only necessary to present enough mechanical context to demonstrate the mechanism being patented.

The new architecture is likely to imply some changes from the Internet protocol suite below the application layer, i.e., in parts of the stack that have traditionally been implemented within the kernel of a conventional operating system. To prototype and test such a different protocol stack will require programmable nodes. This can be achieved (1) in a local laboratory environment, (2) in a research testbed, or (3) tunneled across the Internet. It is likely that all of these approaches will be used.

However, requirements and objectives are abstractions that may not be directly quantifiable, and some important principles cannot be tested in a testbed of feasible cost. In particular, requirements in the following four areas must be validated using a combination of simulation and plausible argumentation: (1) scaling issues, (2) heterogeneity, (3) high performance, and (4) interaction with economic and business models. In this context, it is particularly important to note recent advances in multi-level and parallel simulation algorithms. The team expects to make use of such next-generation simulation techniques to the extent supportable by the technology at the time they are required.

4.3 Technology Transfer

The primary purpose of the proposed effort will be to create an architecture – a set of conceptual rules that meet specified requirements. This result may be considered to be a meta-technology, since it is intended to guide the future development of a manifold of network protocol technology. The problem, then, is meta-technology transfer.

The primary transfer task will be to convince the larger network technical commu-

nity of the viability and value of the new architecture, and then to enlist them in pursuing the detailed technical consequences that flow from it. The initial route for meta-technology transfer must be the traditional academic channels of technical papers and presentations. In addition, the proof-of-concept code will provide a platform for other researchers individuals and organizations to get started with the new concepts. After this, we can hope that the broader Internet community of vendors, providers, and users will buy into the new architecture.

However, papers and presentations in existing forums may be insufficient to accomplish this effectively. The new-arch research team should identify opportunities to organize and hold workshops and retreats specifically targeted at advancing the transfer of results from this work. To do this effectively, the team should seek out workshop sponsors and partners that contribute both resources (intellectual, organizational, and financial) and visibility.

The time-scale for the effect of the new architecture on the Internet may be very long, of the order of 10 years or more. However, there may be much shorter-term effects on technical decisions, once there is a future towards which we believe the technology should evolve.

5 Conclusions

The Internet architecture continues to a central role in the technical development of the Internet. The importance of this conceptual abstraction is broadly acknowledged by the Internet technical community and in the IETF, even though vendors, ISPs, and carriers with practical technical problems have largely replaced network researchers as IETF members. For example, the index of RFCs lists 42 documents with the word “architecture” in their titles in the last 10 years, and 14 in the last 3 years; the rate is essentially constant. However, the rapid proliferation of un-architected extensions with complex and conflicting feature interactions has created a serious technical muddle in the IETF, and without an updated architecture the result will be a loss of functionality, adaptability, and robustness.

The importance of the Internet architecture to technical development is enshrined in the continuing existence of the Internet Architecture Board (IAB) within the Internet technical infrastructure. The IAB and the IETF are chartered by the Internet Society. While it has few real powers in the Internet standards process, the IAB plays the role of a watchdog on the architecture; it acts as an independent

investigator of architectural [BCarp96] and other technical issues relevant to the IETF. Membership on the IAB is determined by a quasi-democratic but meritocratic selection process. Although the IAB serves a highly valuable function, its organization and membership make it highly unlikely that the IAB will design or promulgate more than minor evolutionary changes to the architecture.

We believe strongly that the long-term viability of the Internet requires a more revolutionary approach. This can be accomplished best (and perhaps only) by a small group of network researchers who are given freedom to work on updating the architecture. These people must be able to build conceptual models, work through their consequences, and then iterate, in a manner that is independent of short-range issues and does not require backwards compatibility. In other words, an effort like that proposed here will be needed.

Specific results of such a research effort should include a requirements definition and an architectural design as well as the design and validation of protocols to realize the architecture. The long-term result will be a design basis for tomorrow's Internet technology that will better position it to meet military and commercial needs.

Evaluation of the concrete results of this project would use conventional simulation and network experiments on prototype implementations. However, evaluation of the abstract results, the architecture and requirements, must be qualitative. Ultimately, the ability of the architecture to influence the future direction of Internet technical evolution will be the true measure of success.

- [ALF90] D. Clark and D. Tennenhouse, "Architectural Considerations for a New Generation of Protocols". Proc ACM SIGCOMM, Sept 1990.
- [BCarp96] B. Carpenter, Editor, "Architectural Principles of the Internet". Internet Architecture Board, RFC-1958, June 1996.
- [Clark88] D. Clark, "The Design Philosophy of the DARPA Internet Protocols". Proc SIGCOMM 1988, Sept 1988.
- [Clark98] D. Clark; Talk at DARPA NGI PI Meeting, Tucson, AZ, March 1998.; XXXX?
- [CerfKahn74] V. Cerf and R. Kahn, "A Protocol for Packet Network Intercommunication". IEEE Trans on Comm, COM-22, No. 5, May 1974, pp. 637-648.
- [CerfKirs78] V. Cerf and P. Kirstein, "Issues in Packet Network Interconnection", Proc. IEEE, v.66, 11, November 1978.
- [Diffserv98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services", Network Working Group RFC-2475, December 1998.
- [E2Earg81] J. Saltzer, D. Reed, and D. Clark, "End-To-End Arguments in System Design". 2nd International Conf on Dist Systems, Paris France, April 1981.
- [FutArch901] D. Clark, L. Chapin, V. Cerf, R. Braden, and R. Hobby, "Towards the Future Internet Architecture". Network Working Group RFC-1287, December 1991.
- [IntServ94] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture". Network Working Group RFC-1633, June 1994.
- [IPSEC98] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol". Network Working Group RFC-2401, November 1998.