Clean-Slate Design for the Internet

A Research Program at Stanford University

Whitepaper Version 2.0, 18 April 2006

Edited by Nick McKeown and Bernd Girod

With contributions by Dan Boneh, Abbas El Gamal, Bernd Girod, Ashish Goel, Andrea Goldsmith, Mark Horowitz, Ramesh Johari, Joseph Kahn, David Mazieres, Nick McKeown, David Miller, Balaji Prabhakar, Tim Roughgarden

Overview

This whitepaper describes a new collaborative inter-disciplinary research program at Stanford. The document was prepared by a group of faculty from the Departments of Electrical Engineering, Computer Science, and Management Science and Engineering.

We believe that the current Internet has significant deficiencies that need to be solved before it can become a unified global communication infrastructure. Further, we believe the Internet's shortcomings will not be resolved by the conventional incremental and "backward-compatible" style of academic and industrial networking research. The proposed program will focus on unconventional, bold, and long-term research that tries to break the network's ossification. To this end, the research program can be characterized by two research questions: "*With what we know today, if we were to start again with a clean slate, how would we design a global communications infrastructure*?", and "*How should the Internet look in 15 years*?" We will measure our success in the long-term: We intend to look back in 15 years time and see significant impact from our program.

In the spirit of past successful inter-disciplinary research programs at Stanford, our program will be driven by research projects "from the ground up". Rather than build a grand infrastructure and tightly coordinated research agenda, we will create a loosely-coupled breeding ground for new ideas. Some projects will be very small, while others will involve multiple researchers; our goal is to be flexible, creating the structure and identifying and focusing funds to support the best research in clean-slate design.

The program will collaborate with, and be funded by, approximately seven industrial partners with interests in networking services, equipment, semiconductors and applications.

In this whitepaper, describe how the program is structured, and identify five key areas for research: 1. Network architecture; 2. Heterogeneous applications; 3. Heterogeneous physical layer technologies; 4. Security; and 5. Economics & policy. We expect these areas will evolve and perhaps change completely as the program progresses.

Table of Contents

Introduction

Shortcomings of the Internet

Designed over 30 years ago, the success of the Internet is a testament to the foresight of a handful of visionary researchers. Hundreds of millions of users rely on it for business and pleasure; and it is now hard to imagine a world without it.

But our reliance on the Internet makes us victims of its success, and vulnerable to its shortcomings. Some of the shortcomings are self-evident, such as the plague of security breaches, spread of worms, and denial of service attacks. Even without attacks, service is often not available due to failures in equipment or fragile routing protocols. And its behavior is unpredictable making it unsuitable for time-critical applications. Other shortcomings are less obvious: The Internet was designed for computers in fixed locations, and is ill-suited to support mobile end-hosts; it uses packet-switching making it hard to take advantage of improvements in optical switching technology; it neither ensures anonymity, nor facilitates accountability; and the demise and restructuring of most network service providers suggests that providing network service is not profitable.

In summary, we don't believe that we can or should continue to rely on a network that is often broken, frequently disconnected, unpredictable in its behavior, rampant with (and unprotected from) malicious users, and probably not economically sustainable.

Characteristics of a New Internet

It is not difficult to create a list of desired characteristics for a new Internet. Deciding how to design and deploy a network that achieves these goals is much harder. Over time, our list will evolve. It should be:

1. **Robust and available.** The network should be as robust, fault-tolerant and available as the wire-line telephone network is today.

2. **Inherently secure.** The network should be built on the premise that security is a must, and it should be protected from denial of service attacks. While it might be impractical and unwise to build a network that is completely impervious to attack – after all, end-hosts are complex and under human control – the network should be secure enough for critical applications such as finance and banking, air-traffic control, and military communications.

3. **Support mobile end-hosts.** Laptops, Wi-Fi and cellular telephones make mobility commonplace, rather than an exception. A new Internet should support mobility (and all the associated security, naming, routing and privileges) as seamlessly as it supports wired end-hosts today. It should also support future small sensor and ad-hoc networks, as well as RFID.

4. **Economically viable and profitable.** The network should be profitable for those who provide service and build equipment, and – if necessary – be coupled with suitable regulation to create competition and incentives for improvement.

5. **Evolvable.** The network architecture should pre-suppose that it will change and evolve over time – perhaps at the very lowest level. Its architecture and service model should not ossify and stifle improvement.

6. **Predictable.** The user should know what to expect from the network, and it should provide predictable and repeatable service. This might include guarantees on timely delivery of time-critical data, or guarantees that enough capacity is available when needed.

7. **Support anonymity where prudent, and accountability where necessary**.

Current Internet Research and Development Won't Get Us There

During the rapid growth of the Internet, academic research has understandably focused on short-term, immediately deployable mechanisms and services. There has been a tendency towards research that is incremental and – to be accepted for publication – research has been expected to be backwardly-compatible with the existing Internet, and not interfere with the basic IP service model.

In some ways the industrial environment has been more innovative than academia. Entrepreneurs and industry have driven most of the enormous growth in capacity through the rapid rollout and deployment of several generations of Ethernet, wireless Ethernet, middleware boxes for security, caching and content delivery, Internet routers and long-haul optical links. Increased connectivity has been quickly and successfully exploited by online services and commercial portals; and the introduction of broadband has spawned new services such as VOIP, peer-to-peer file sharing, and downloadable movies.

But the commercial success of equipment vendors and services built on top of the Internet (e.g. Google, Yahoo, Akamai, eBay, etc.) creates a strong vested interest and resistance to change. Those who profit most from the status quo are (understandably) least likely to rock the boat. They will tend to create barriers to the introduction of radically new technologies.

Resistance to change is compounded by the end-to-end design philosophy that makes the Internet "smart" at the edges and "dumb" in the middle. While a dumb infrastructure led to rapid growth, it doesn't have the flexibility or intelligence to allow new ideas to be tested and deployed. There are many examples of how the dumbness of the network has led to ossification, such as the long time it took to deploy IPv6, multicast, and the very limited deployment of differentiated qualities of service. Deploying these well-known ideas has been hard enough; deploying radically new architectures is unthinkable today.

Other Research Initiatives

This program complements a nationwide research effort to reconsider the design and architecture of the Internet. In 2003, the National Science Foundation (NSF) funded the "100x100 Clean Slate program" (http://**100x100**network.org) at CMU, Stanford, Rice and FraserResearch. Now in its third year, this program has developed novel "clean-slate" approaches to network design, access networks (optical and wireless), network control and congestion control. So as to involve the whole research community in clean-slate design, in 2006 NSF funded the FIND (Future Internet Network Design http://find.isi.edu/). FIND is a precursor to the proposed and much larger NSF GENI program (http://www.nsf.gov/cise/geni). GENI plans to build a nationwide programmable network – a research platform upon which experimenters can create, deploy and use whole new network architectures across a nationwide network. GENI is ambitious: the programmable network would support hundreds of simultaneous experiments, all running on the same virtualizable platform. If deployed (it will require about $250M from NSF), GENI will not be operational before about 2010.

The Stanford Clean Slate Program, while complementary with GENI (we expect the outcome of our work to lead to several network experiments on GENI), is unique in several ways. First, the Stanford program will be funded mostly (perhaps entirely) from about seven industrial sponsors. Second, all of the researchers are at the same location, allowing close collaboration across disciplines and among graduate students. Third, as described below, it allows us to bring together world-class researchers from several disciplines outside the traditional field of networking. We believe this will bring a broader and fresher perspective.

Why Stanford is a Good Place for a Clean-Slate Program

Stanford is well-positioned to make a large impact on the future global communications infrastructure. We have an unusual and perhaps unique tradition of interdisciplinary research, with few structural boundaries to collaboration or co-advising of graduate students. Networking – and the Internet in particular – does not naturally belong in any particular academic department. Optical, wireless, an\][d wire-line might belong in an electrical engineering department, while protocol design and security might belong more naturally in a computer science department. Modeling and theory of networks find their home in EE, CS, MS&E and Statistics; and application design might belong in almost any department of the university and medical school. With a large, naturally collaborative, diverse, and world-class faculty, we believe Stanford brings a unique set of strengths to Internet research.

A Breeding Ground for Clean Slate Research

The Clean Slate Program is an open and inclusive research program in-keeping with the style of Stanford's successful research centers. Rather than putting in place a static "one size fits all" organization, we have created something more lightweight and flexible.

First and foremost, the program will be about research. The majority of effort and funds will be directed towards a mix of small, medium, and large collaborative research projects, so long as they fit the broad goals of the program. When evaluating and deciding upon research to include in the program, we will ask the following guiding questions:

1. *Is the research "clean-slate" enough? In other words, is it sufficiently bold, long-term and a departure from existing research.*

2. *Is it likely to be successful, and if so, is it likely to have a large impact on the future Internet?*

Organization of the Program

Executive Director: Nick McKeown, with help from Bernd Girod. Term is nominally for two years.

Advisory Board: Dan Boneh, Andrea Goldsmith, Mark Horowitz, Ramesh Johari, Balaji Prabhakar. Term is nominally for two years, with half of the board reappointed each year.

Faculty: All Stanford faculty and students are welcome to join the program and attend seminars. Participants so far have included more than 20 Stanford Professors from four departments: Computer Science (CS), Electrical Engineering (EE), Management Sciences and Engineering (MS&E), and the Graduate School of Business (GSB) and over 50 PhD students.

Funding Process

The Program will maintain an informal proposal process, designed to be lightweight and able to take risks. A Request for Proposals will be announced every 6months (Deadlines: beginning of May, and beginning of November). A proposal will consist of a two-page description of the research work, including a means to judge its success after one year (for review), and after five years (for impact). Proposals will be reviewed by the Advisory Board, and decisions made in consultation with the Director.

Our goal is to be flexible. By default, a project will be funded for two years (to give continuity of funding for graduate students). It will be reviewed by the Advisory Board after 12 months to decide if the research should be funded beyond two years. We expect to fund approximately 3-4 new projects per year (1-2 per RFP process); after the startup phase, we expect approximately eight projects to be running at any one time.

While different projects will be of different size and duration, we think of the "canonical" project as consisting of two professors, two PhD students, and lasting two years. In practice, we expect there to be variation from project to project, and we will welcome proposals and ideas for projects that fall outside this model.

Resources and Funding

We plan for the program to be funded by both industry and government. Assuming the program grows to support sixteen PhD students and postdocs, with associated travel and equipment, it will need approximately $1M per year. We don't anticipate that it will be difficult to find funds from industry or government for a high profile research program, on such a relevant topic, given the world-class reputation of our researchers.

The School of Engineering has recently committed substantial funds to kick off this new program. The goals of the program and the impact it could achieve are very much in-keeping with the original aspirations of the funders and founders of the Stanford Networking Center (SNRC), which is currently concluding 6 years of successful operation. We expect that several members of SNRC will transfer their involvement to the new research program.

So Far

The program started at an offsite workshop held in August 2005. Twelve faculty members attended and shared ideas and problems. Details, abstracts and slides can be found at: http://klamath.stanford.edu/csdi.

We have launched a weekly seminar to bring together faculty and students. The seminar is in its third quarter: Details can be found at: http://cleanslate.stanford.edu.

Areas of Research

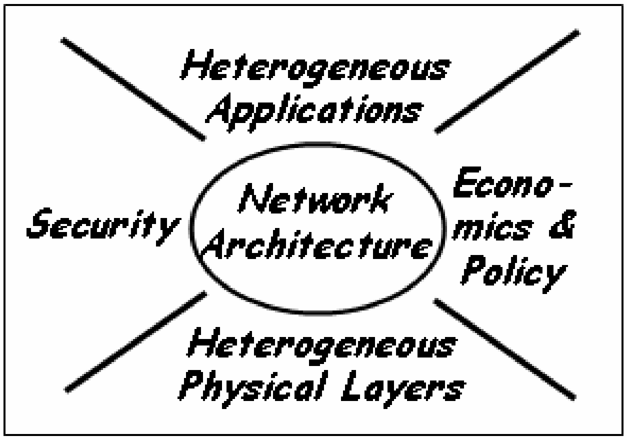

Figure 1: Research Areas

At our retreat in August 2005, we identified key areas and research questions that must be addressed for a clean slate internet design. In what follows, we introduce five research areas (Figure 1). Security and economics/policy are important areas that tie together high level application heterogeneity and low level physical layer heterogeneity. They are all at their core driven by network architecture, and inform network architecture.

There is no natural order in which to discuss our five research areas, as they are all interact in numerous ways. For example, while it is tempting to think of our work motivated by applications from the top-down, it is worth noting that predicting future applications is rarely successful (just think back fifteen years and ask if you would have predicted the web, search engines, peer-to-peer file transfer, and the growth of IP telephony). Our goal is to put on the table the set of topics identified at this time, recognizing that our list will change as we learn more.

Within each area we list a number of research topics we plan to explore. Our goal is to create an open program that encourages new ideas, even if they are not mutually compatible. We feel that a tightly integrated program at this stage would constrain the research and unduly favor incremental ideas.

Area 1: Network Architecture

The original Internet architecture provided a method to transport packets, and has changed little since it was first proposed. It provides a "dumb" connectionless packet-forwarding packet-switched infrastructure, with high functionality at the edge (the so-called "end-to-end principle")**.** The Internet provides a single, simple lowest-common denominator best-effort packet-switched datagram delivery service (IP), with fixed-size numerical addresses (one per physical network interface). If an application requires a reliable stream service, it can optionally be provided on top of the underlying unreliable service.

Adherence to the end-to-end principle has come with two main costs: loss of functionality within the network, and a lack of innovation. Although the Internet has evolved, it has done so surprisingly little. New functionality has only been added in response to pressing problems that threatened the Internet's operation. For example, when congestion caused the network to collapse in the 1980s, congestion control and avoidance were quickly deployed. Similarly, when the network almost ran out of IP addresses in 1993, CIDR (Classless Interdomain Routing) was deployed in just a few months. And "subnetting" (more efficient use of addresses in enterprises), DHCP (to dynamically allocate addresses to temporary hosts, such as laptops), NAT (network address translation for firewalls), and access control were all added out of necessity – to solve an urgent problem. While the ability to deploy new functionality is good, there have been two main problems. First, by deploying techniques in a hurry, some half-baked mechanisms have infiltrated the network and have been hard to improve or remove (e.g. policies in the routing protocols that lead to unstable routing protocols). Second, desirable but non-urgent mechanisms take more than a decade to deploy (e.g. multicast and IPv6). The network has essentially ossified. As a result, most innovation in the infrastructure has taken place in "gated communities": inside enterprises that have isolated themselves from the public network.

Researchers and network operators have tried to enable innovation by introducing overlay networks and MPLS (multi-protocol label switching). More recently, there have been proposals to introduce virtualization into the network, in which links and routers are "sliced" to support multiple experimental networks simultaneously.

In our work, we plan to revisit many of the basic assumptions of network architecture. Here are some topics we will explore:

1. **Flows as first-class citizens.** One innovation that we believe to be important is the recognition of *flows* in the network. We believe flows should be treated as

first-class citizens,[1] perhaps replacing the packet as the predominant unit for manipulation inside switches and routers. To some extent routers already identify flows for bandwidth partitioning, preferential service, signaling congestion, security and flow switching. Flows can also be used to determine the economic value of a flow and to offer differential service on this basis.

2. **Network addressing.** Internet addresses identify a physical network interface. Instead, can addressing be made more intuitive, referring to services and people, rather than physical interfaces?

3. **Routing protocols.** How can we simplify routing protocols and make them more reliable and stable?

4. **Exploiting structure.** Can we exploit the way in which networks are used, and their inherent structure? For example, it seems likely that tree-like access networks (optical, electrical, wireless and hybrids) will persist, interconnected by a richly connected core. In North America, the core consists of approximately 100 switching centers (based on population centers) interconnected by high capacity long haul optical links.

5. **Dynamic circuit switching.** If the core of the network is to benefit from high capacity all-optical switching, then should we deploy dynamic circuit switching? If so, how?

6. **Backbone design.** Today, backbone networks are hugely over-provisioned. Can they be designed more efficiently, to be tolerant to failure, and predictable throughout their lifetime?

7. **Models of the end-to-end principle**. Can we capture the costs and utility of the end-to-end principle in a mathematical model (analogous, say, to the successful "price of anarchy" program, which considered the cost of source routing based on delay information)?

8. **Cross-layer design.** While we recognize the power of layering in network design, it has inevitable inefficiencies. We will explore where interfaces belong, and what services each layer should provide.

9. **Network virtualization.** Is it possible to create a network infrastructure that is continuously evolvable? This is currently under consideration as part of a large NSF initiative, called Geni, in which links and routers are virtualized to build a nationwide research infrastructure. We will be involved in the Geni research program, and plan to explore network virtualization as an extension to the NetFPGA work already underway at Stanford.

Area 2: Accommodating Heterogeneous Applications

Due to its open architecture, today’s Internet enables a broad range of applications which can be implemented without regard of the underlying infrastructure. Support of *heterogeneous applications* has arguably been the most important driver for the Internet's

[1] A phrase introduced by Sandy Fraser of Fraser Research.

rapid growth, yet the protocols which comprise the Internet are woefully inadequate to support this heterogeneity. Some applications (such as voice telephony or tele-surgery) may require very strict delay guarantees. Others, such as email delivery, primarily require a sufficient aggregate data volume over a period of time. Some applications (e.g. video streaming) might need a certain minimum bit-rate, but can tolerate packet losses, while others (e.g. file transfer) are highly elastic in terms of bit-rate, but need reliable delivery of each byte.

Emerging applications cannot always be easily supported by today's Internet. This is particularly true of *sensor networks*. It is expected that most of the network-attached devices in the future will have sensing capabilities. Both wireless sensor networks and networked imaging sensors will be widely deployed. Low data-rate, cheap, energy-limited sensor nodes distributed over a relatively small geographical area have required the research community to develop specialized networks that are distinct from the Internet. Existing sensor deployments are rarely connected to the Internet due to the challenges of secure and reliable transport.

Resource allocation, performed in today's Internet by the end hosts through the use of the *TCP protocol*, does not recognize heterogeneous application requirements. The combination of error control, flow control, and rate allocation in one protocol is a major limitation of current Internet technology. Even if TCP is not used, e.g. for real-time streams, the "fair sharing" of TCP is adopted in the form of TFRC (TCP-friendly rate control). TCP or TFRC do not recognize that different applications have vastly different elasticity to the rate allocated to them. Moreover, there are, in general, no mechanisms to prevent malicious applications from usurping more than their fair share. Malicious traffic can easily shut down other flows. Denial-of-service attacks exploit this weakness.

The "sawtooth" nature of throughput and delay make TCP famously unsuitable for *real-time applications*. It is nevertheless used when firewalls prevent the use of UDP. The overhead associated with reliable transport on top of stateless routing unduly penalizes small packet payloads, required for low-latency applications such as voice telephony or remote control. TCP interprets packet loss as an indicator of congestion and thus behaves erratically when confronted with wireless losses. All these short-comings are well-understood and have led to a flood of incremental "stop-gap" research over the last decade. However, the success of the Internet has prevented these problems from being tackled in a comprehensive and fundamental way.

We believe that the future Internet should provide much better support for a broad range of applications and enable new applications, such as distributed control, camera networks, or anycast. In our work, we plan to explore the following topics:

1. **Maximum utility resource allocation.** What are the right notions of fairness for heterogeneous applications? Utility captures the value an application derives from a certain allocation of network resources (bandwidth, delay, etc.). Is a distributed resource allocation scheme feasible which maximizes the total utility across all users? How should utility be exposed? Can such a scheme prevent malicious applications from grabbing more than their fair share? How can it be extended to distributed applications?

2. **Application-aware congestion control.** Can we design a parameterized congestion control mechanism which has varying behavior depending on the utility function of the application? Can we develop simple routing/queue-management mechanisms that support this parameterized protocol?

3. **Multi-path source routing.** What benefits can multi-path routing provide to alleviate congestion and increase robustness? Should it be under the control of the application or the network?

4. **Flexible transport service.** What should a future reliable transport protocol look like? Should it include forward error correction in combination with automatic retransmission? Should it provide bandwidth, delay, jitter and loss guarantees? Which limitations can and which cannot efficiently be overcome by over-provisioning?

5. **Multicast and anycast transport.** How can point-to-point transport extend to one-to-many, many-to-one, and many-to-many transport?

6. **Benefits of flow-based network design.** We conjecture that flows must be first-class citizens to best support heterogeneous applications. Can we rigorously quantify the benefits of such a flow-based network design?

7. **Location-based services.** Future network nodes will often be aware of their fixed or mobile geographic location. How can nearby peers find each other? How can applications query nodes in a certain geographic area? What services should be universally provided by the network?

Area 3: Accommodating Heterogeneous Physical Layers

While the Internet was created originally for a fixed mesh of relatively slow wired links, we have seen an explosion of technologies that support the network---from optical fibers to wireless mobile access. This *physical layer heterogeneity* poses tremendous challenges for network architecture, resource allocation, reliable transport and security.

Optical fiber is the backbone of the Internet because it is, and is very likely to remain, the best and highest-capacity means of transporting large amounts of information over long distances. Scaling fiber transmission capacity will rely on increasing the number of parallel channels through wavelength-division multiplexing, as per-channel bit rates approach their practical limits (about 40 Gb/s). The potential bit-rate of each optical fiber (tens of terabits per second) is so high that in practice, it is the cost to light and operate the fibers that limits transmission bandwidth. Nevertheless, for most users, end-to-end bandwidth is still constrained by limitations of currently deployed access technology.

Wireless access to the Internet has grown dramatically, and wireless technology poses very different design challenges than wired systems. Wireless channels have much lower capacity than, say, optical fiber, and wireless users wireless users experience time-varying channel quality and also interfere with each other. Thus, bandwidth is expected to remain scarce, with an increasing mismatch between wireless access and wired backbone. Mobility leads to rapidly changing channel conditions. While losses in the wired network primarily occur only due to buffer overflow in routers, losses in a wireless channel are frequently due to interference, signal fading, or noise.. TCP interprets packet loss as an indicator of congestion and thus behaves erratically when confronted with wireless losses.

Mobility was not anticipated in the Internet’s addressing scheme; IP addressing is tied to a user’s physical location. Later patches, such as Mobile IP lack support for local or hierarchical mobility management and do not provide the desired level of support for seamless and efficient network operation.

We plan to address the following research challenges:

1. **Optical Internet.** Given that there will be a continued need to scale bandwidth, how can we build an infrastructure with scalable bandwidth and affordable design costs. What will be the principal technology-imposed constraint on the system? To what extent should all-optical switching continue to evolve from circuit switching toward flow and packet switching? Will continued scaling of the network require substantial changes in physical technology? Does building optical buffers to enable all-optical packet switching really matter?

2. **Impact of CMOS scaling.** While CMOS scaling will continue, the energy efficiency of CMOS systems will scale slowly in the future. How should large switching systems be built, when cost of power and cooling dominate and scaling no longer hides the cost of increasing electronic complexity? Will power limitations constrain switch complexity? Do power constraints in future systems drive buffering (and thus routing) out of the core of the network?

3. **Wireless Internet.** How should wireless communication with its intrinsic mobility, interference between users, broadcast capabilities, and dynamically changing link performance and network topology impact network design? How can we ensure end-to-end performance that is commensurate with the application requirements (e.g. rate and delay) in the face of rapidly changing mobile radio channels. Such performance “guarantees” are hard even in static networks, but become even more challenging when the end-to-end routes change based on node movement.

4. **Mobility support.** A network design that supports mobility must ensure that applications perform seamlessly despite movement of network nodes. Data forwarding cannot be based on a static node location or network topology and data must follow a mobile node as it moves through a network or across networks. How can seamless hand-over be achieved between different cells of the same network or between different networks? How can trust between mobile nodes and the network be established and maintained.

5. **Ad hoc networks.** How should resource constraints in low-power mobile devices affect the network architecture and protocols? How can trust schemes for mobility support be extended to mobile ad hoc networks, where nodes themselves become relays? How should future transport and routing protocol accommodate wireless sensor nodes, which “sleep” most of the time? How can reliable and predictable performance be achieve in mobile ad hoc networks, e.g., cars on a highway?

6. **Resource allocation in heterogeneous networks.** When multiple heterogeneous networks are available, when should flows be routed over which network? Where should the packets be buffered? Is there a role for multiple description codes and/or multipath routing? How can application requirements be incorporated in physical resource allocation?

Area 4: Security

Any future network design must be built with security in mind from the start. We envision four high level requirements. First, a future network design should make it harder to mount the kinds of attack that are currently prevalent. The network should provide tools to quarantine fast-spreading infections, mitigate Denial of Service attacks, and provide better source authentication. Nevertheless, we fully expect next-generation unanticipated attacks to emerge. Thus, our second high level requirement is to ensure that the network is designed so that detection and recovery from attacks is much easier than it is today. A third component of future networks may include tools to help law enforcement identify the origin of an attack. This is a controversial topic and the extent of such tools must be limited by an appropriate privacy policy; most likely by drawing on analogies from the physical world. Clearly, all these enhanced security features should not diminish the usefulness of the Internet. Thus, our fourth high level requirement is that the Internet remain a general-purpose communication medium as it is today.

A network designed to meet these requirements should provide the following capabilities: (i) block malware from spreading, (ii) identify compromised hosts in case malware does spread, and (iii) quarantine compromised hosts until they are fixed. A better network design can improve all three tasks, as discussed below.

One approach to preventing malware from spreading is to restrict the full connectivity model available in current networks. For example, although machines in a LAN typically communicate with few other machines (the file server, mail server, printer, and web proxy), current networks provide full connectivity: any machine in a LAN can communicate with any other --- a fact frequently exploited by network worms. A future LAN architecture could make it harder for worms to spread by limiting the full connectivity available today.

Future networks must include capabilities to mitigate network-layer Denial of Service attacks. To prevent a compromised host from participating in a Distributed Denial of Service attack (DDoS) a future network should be able to limit the amount of traffic generated by end-hosts. In particular, a site being attacked by DDoS should be able to submit evidence of the attack to a quarantine service. The service will then restrict traffic to the site from the sources participating in the attack. Such rapid response to attacks could eliminate network-layer DDoS altogether.

A variety of common Internet attacks exploit the weak source authentication in current network protocols. For example, spammers and phishers are able to fool users by sending mail that appears to come from reputable sources. A future network design should include stronger source authentication so that end-host filters can block such forgeries.

Area 5: Economics and Policy

The past hundred years have shown that market forces have tremendous impact on the structure and operation of communications networks. In the current Internet serious deficiencies in the market structure have been exposed in recent years. The current Internet has not converged on a balance between regulation and competition; observe, for example, the fact that six of the seven largest national ISPs in 2002 have since undergone corporate restructuring. They are simply not profitable.

At a regional level, operators lack incentives to deploy next generation network technologies. This is compounded by the fact that regulators do not yet understand the proper role they should play – note, for example, the confusion surrounding the bundling and unbundling of local DSL services. In the last decade, the FCC has been more reactive rather proactive in addressing the economics of network management and investment, due to a lack of understanding of the competitive tradeoffs that affect network evolution.

Some of the economic travails of the current Internet can be traced to a failure of engineering. The Internet lacks explicit economic primitives, hindering its functionality in several ways. In particular, the Internet provides no support for determining the value of a packet (to the sender, the receiver, or the service provider). Such information could be used, for example, to better allocate the resources of the network, providing high-value traffic with higher bandwidth, more reliability, or lower latency paths. A related issue is that the current Internet does not provide support for differentiating between different packets on economic grounds. For example, two packets with the same origin and destination will typically be routed on the same path through the network, even if the packets have very different values. Even if these values were known to the network, the current routing protocols would not permit the packets to travel on different paths. Finally, the lack of economic primitives in the current Internet makes charging for traffic, and micropayments in particular, a challenge to implement. Such payments could contribute to both the prevention of near-valueless uses of the network (spam) and to defraying the network maintenance costs.

We plan to investigate two broad areas: high level market structure and low level engineering of economic primitives.

High level market structure. This set of questions considers broad economic questions regarding sustainability of the network infrastructure. However, answering them requires a clear understanding of engineering details, particularly cost structure.

1. **Investment costs.** How do network costs and structure affect incentives for investment and operation, and what are the consequences for regulation and competition? Specifically, what is the influence of investment and operation cost structure on industry architecture?

2. **Regulation.** Where can competition sustain the network, and where is regulation needed—at what regional levels of aggregation, and for what types of network technology?

3. **Demand.** What is the effect of demand for new content and services on market structure? Should network service provision and content provision be bundled or decoupled? Can content-generated revenues sustain investment in network infrastructure?

Low level economic primitives. At a more fundamental level, we must understand which economic primitives should be integrated into the network infrastructure itself, and how this integration should be accomplished. This set of questions is rich from an engineering perspective, because we must understand the pieces of information needed at a protocol level to communicate *value*.

1. **Packet-based and flow-based value identification**. Should packets and flows be assigned an explicit value? Can such identification be used to achieve price differentiation?

2. **Contractual granularity.** What is the proper "granularity" for the formation of contracts between network agents (including providers and users)? Here "granularity" includes both the timescale of contracts, as well as their quality of service requirements (bandwidth, delay, jitter, and loss).

3. **Incentive compatibility.** As the information provided by a protocol increases, gaming behavior may become more widespread; e.g., a provider may use knowledge about other providers' networks to its own competitive advantage. What are the consequences of incentives on the performance and robustness of network resource allocation protocols? To what extent can feedback such as

loss, delay, or even reputation be used to manage incentives in the absence of currency-based transactions?

4. **Wireless spectrum allocation.** A major specific design problem is emerging in the wireless industry: how should spectrum be shared among competing uses? The FCC is considering radical new ways of allocating spectrum (opportunistic radios that find and use "white spaces" in space and time in the frequency spectrum). Technical, economic, and policy issues will inform the FCC's decision, and the result will have a major impact on the engineering of future wireless systems.