

NSF Future Internet Summit

Washington, DC, October 12-15, 2009

Meeting summary

Version 7.0 of January 5, 2010.

Introduction

This is a summary of the NSF Future Internet Summit, held in Washington, DC on October 12-15, 2009¹. The meeting is part of the larger NSF Future Internet Design (FIND) initiative.

The FIND initiative is distinctive from many other NSF programs, both in its technical emphasis and its approach to research. The FIND program seeks to involve the research community in a collective exploration of what the Internet of 15 years from now should be. This project requires a long-range perspective, the invention of new concepts not fettered by the constraints of today's networks, an understanding of the larger social, economic and legal issues that arise from the interplay between the Internet and society, and a willingness to collaborate on larger visions that build on the ideas developed in individual research grants. After four years of initial research on concepts and mechanisms, the community has urged NSF to move the FIND initiative to a next phase², where approaches and mechanisms are woven together into coherent, overarching candidate designs for a future Internet. This meeting was organized to allow those interested in this next phase to meet, develop a basis for collaboration, share technical insights and develop new ideas, and begin to develop approaches to research. Because of the value that the FIND program places on integration and collaboration, NSF felt that it was important that there be an opportunity for members of the community interested in phase two of the program to meet and discuss objectives, approaches and opportunities.

The high-level goal of the meeting was to bring the research community together in order to develop new technical approaches and to build research collaborations that will move the current research on a future internet to the next stage, where integrated design proposals are developed and prototyped. The approach, format and scope of the meeting was intended to facilitate an intense and substantial interaction over four days centered on creativity, multi-disciplinary collaboration, the building of research teams and the development and evaluation of promising design approaches.

¹ This meeting summary was prepared by David Clark, as part of my role as PI for the meeting. The summary draws on presentations and other contributions from many meeting participants, and I acknowledge the range of material on which this report is based. The integration of that material into a single report is my undertaking, and any misunderstandings, misattributions or misplaced emphasis should be attributed to me.

² Earlier community discussion on the future of FIND can be found in the summary of the Spring 2008 FIND meeting, at http://www.nets-find.net/Meetings/FifthPIMeeting/Agenda_FIND_FifthPI_files/April2008summary.pdf, and the report of the observer team at the spring 2009 meeting, at http://www.nets-find.net/FIND_report_final.pdf.

A starting point

This workshop was designed as a starting point for a process that will help the research community to prepare and submit competitive proposals for a future Internet. It was a beginning, not a conclusion. While many ideas are captured in this report, they should not be taken as “the answer”, or constraining in any way what might finally be proposed.

The charge from NSF

The meeting began with comments from Ty Znati, the Division Director for Computing and Network Systems within CISE. This section summarizes his comments; his slides are available³. This summary also includes additional comments made at the end of the summit; see additional slides⁴.

The design of a future Internet is a critical research objective for the community. Society needs **YOU** to create the trustworthy Future Internets that meet the needs and challenges of the 21st Century. Society’s needs for an IT infrastructure may no longer be met by the current trajectory of incremental changes to the current Internet; the research community needs to explore both incremental and more revolutionary paths to anticipate the future.

The goal of this special funding opportunity is to (i) stimulate innovative and creative research from different perspectives that explore, design, and evaluate trustworthy overarching Future Internet architectures in the context of a range of scientific, technical, economic and social challenges, and (ii) grow inter-disciplinary communities of individuals from different domains, that currently do not exist, and foster collaborative free thinking among members of these communities, that otherwise would not have been possible.

Designing Future Internets that meet the emerging and future needs for our Society calls for a multi-disciplinary research agenda that **may** build on what is now known about requirements and mechanisms, considering lessons from the past, incorporating what is good, proposing new approaches where they are needed, and fitting these ideas into fresh overarching architectures that reach beyond core networking. Some areas of new Future Internet architectures are likely to raise economic, societal and legal issues, and to this end, participants were encouraged to conceive their research agenda in such a context. In this respect, it is expected that the research agenda include:

- Components and requirements, e.g.,
 - Trustworthiness, Economic viability and social values
- Underlying architectural principles
- Identification of architectural invariants
- Exposure of component interactions
- Prototyping of proposed architectures:
 - May require the construction of new artifacts
 - May use the prototype GENI, National Cyber Range (NCR) or others

³ See http://api.ning.com/files/RWAr-puboTdmKYkfuqXJFEfH3-LPrbz5eQGavn3gmGw_/TyFutureInternetSummitNSFOpeningRemarks.pptx

⁴ See http://api.ning.com/files/R1ad5qiBKnlrUFjnY5yfhX3cVeTKV4v3jNsGwCKurLI_/FutureInternetSummitNSFConcludingTyRemarks.pptx

- Metrics for architectural evaluation and a path to the comparison of alternative architectures

The FIND project, to this point, has put architectural research “back on the map”, with a refreshing and liberating impact on future network architectures. We are gradually moving toward new theories of network architecture, and new ground is being broken in a wide range of core networking areas such as naming, addressing, network management, access and transport technologies, sensing, content and media delivery, and network applications. Previous work has led to valuable outcomes and contributions, with solutions to point problems and Future Internet components and requirements. However, we are still far from the understanding needed to pick coherent architectural alternatives from among emerging ideas.

The goal of this meeting is to make progress towards the articulation and demonstration of candidate future networks, and to start the work of assembling and integrating coherent ideas and building blocks into possible architectures and architectural components. We have invited to the meeting professional facilitators, mentors, and participants drawn from a number of domains: networking researchers, experts in security and privacy, experts in economics, and individuals from other domains and interests, such as law, application designers, and social concerns. We also have participants from the government.

We expect you, the participants, to engage in brainstorming about candidate overarching Future Internet architectures, explore network architecture that includes necessary components and requirements, and identify underlying architectural principles and architectural invariants. As part of formulating new architectural proposals, you should suggest what future research is needed to expose and explore the multiple components of the architecture and their interactions, and ensure that candidate architectures consider important ethical, legal and societal dimensions. You should propose metrics for architecture validation, evaluation and assessment, and recommend how the overarching architectures should be tested on experimental infrastructure such as GENI, National Cyber Range (NCR) or other research platforms. You should find creative and effective ways to nurture the outcomes of this event and expand the scope of Future Internet Architecture research as the effort goes forward, and contribute to documenting the lessons learned and outcomes from participation in the Summit.

To help make this meeting a success, we have invited a group of experts to act as mentors to the process. The mentors are:

- Bruce Davie, Cisco
- Lance Hoffman, GWU
- Bruce Maggs, Duke University and Akamai
- Scott Midkiff, Virginia Tech
- Sue Moon, KAIST
- Craig Partridge, BBN
- Cetin Seren, Cisco
- David Tennenhouse, New Venture Partners

Here is the charge we have given to the mentors.

- Guide groups of participants in their collaborative quest to generate overarching architectures

- Encourage and foster pursuit of novel, adventurous and innovative ideas
- Catalyze the participants to challenge their thinking and their assumptions
- Encourage and foster connections between participants and between groups of participants
- Provide feedback to participants and groups of participants
- Capture the key lessons learnt

Next steps

NSF is planning on producing a Dear Colleague Letter (DCL) informed by the summit, which would announce a one-time multi-million dollar funding opportunity for research on overarching Future Internet architectures. The project budgets would be larger than typical “large proposal” budgets. The target to release the DCL would be early November, and the proposal deadline would be 90 days from the issuance of the DCL, in accordance with normal NSF policy. A proposal, to be responsive to the anticipated DCL, will:

- Be composed of a cross-cutting team of experts to address multiple architectural components,
- Start with a well-articulated “vision” of the proposed overarching architecture and its high-level objectives,
- Address the requirements the architecture is expected to meet,
 - Include at a minimum **trustworthiness** broadly defined (security, privacy and usability) as fundamental design principles.
 - Management, Economic Viability and Social Values are of great interest
- Reason carefully about the architectural responses to the requirements, including trust, and the interactions between the building blocks, and define a coherent research agenda focused on the technical approaches to achieve these requirements,
- Identify underlying architectural principles and architectural invariants,
- Include prototyping and evaluation of the proposed overarching architecture.
- Provide a coherent management plan.

In detail, a strong proposal should strive to answer the following:

- What are the ideas that will drive the new architecture?
- How will these ideas be synthesized into an overarching architecture?
- What are the pillars/tenets and building blocks of the proposed architecture?

The standard proposal evaluation process will be used for selecting architectures to be pursued. A panel of experts will review proposals and make recommendations for awards. The charge to the Panel will be focused on “integrative” ideas, consistent with the framing of this summit. NSF will make final award decisions. It is expected that NSF will fund 2 to 4 proposals, based on availability of funds and the quality of research.

The following specific issues will be addressed in the DCL:

- Proposal submission is not limited to Summit attendees
 - Proposal teams may include any mix of those who did and did not attend workshop
 - Typical NSF eligibility rules apply
 - No impact on eligibility to core and cross-cutting programs
- A PI can be a lead in at most one proposal
 - Leads must be identified in the management plan

- No limitation on how many projects a PI can participate in
- Well justified support for technical staff and full-time researchers is allowed, commensurate with effort and duration of projects

NSF specifically reminds the research community that this special funding opportunity is not the only avenue by which to propose research on future Internet designs. Multiple venues are in place at NSF to support networking research with different scope and interest:

- NeTS, NetSE or TwC Large Competitions – Late November
- “Focused” ideas from the Summit can be turned into projects for NeTS, NetSE or TwC Small Competitions – Mid December deadline
- Expeditions, with an expected deadline – late 2010

Step one: defining the objective of a future Internet

The organization of the meeting was structured to postpone the discussion of mechanism—possible approaches to solving a problem—until we first discussed what the problem was. Different participants had different views about the objective of a future network, and what that implies for the requirements the network must meet.

To start the discussion, David Clark posed a list of possible objectives, which were called “use cases”, and a list of requirements that might be implied by a candidate use case. The three starting objectives were:

- The network “we all use”: the network of people and the applications that serve them—information, social context, interaction, etc.
- A network to support critical infrastructure.
- A network to connect sensors, actuators and the rich coming array of embedded processors.

The first step in the meeting was to get the group to propose their own candidate objectives, and to break up into groups to discuss each of these. On the first day, the following objectives were proposed and discussed in breakout groups. (In reading the following, it should be understood that these are summaries of discussions that were only about an hour in length. The ideas may seem fragmentary, and inconsistently developed. This summary is intended to give the reader a sense of what happened, as it was actually reported.)

1) An Internet for the other 3 billion. This objective requires us to change the rules of operation, technology, and economics. There will be a broad range of devices, many mobile. The discussion group identified five key problems: 1) how to make it easy for naïve users to connect, 2) how to make such a network economically viable, 3) how to manage a network without experts, 4) how to function in infrastructure-poor environments, and 5) how to deal with a range of social values. Some of the limitations of the infrastructure might suggest that the user should play a larger role in the operation of the network, but this network must deal with a population that is not technical, and in many cases that does not read.

2) A network of people with devices. This objective focused on a network based on human identity and relationships. The essence of the idea is that communication is with the person, not the device. When you change devices, (e.g. mobility), you are still you. So both mobility

and identity are first-order issues, as are issues of privacy, transparency and variation in social norms. Within the discussion group, there was no consensus on the range of issues. Perhaps a network of this sort will exploit social networks for routing, for example. It was not clear if this discussion was just about “the application layer”, or whether this objective required specific support in the network.

3) A network in which mobility is the norm. This breakout group identified a number of issues: 1) intermittent connectivity, which raises issues of binding to identity as location changes. The mobility of networks, as opposed to single devices, makes this a richer problem. 2) effects of wireless, which implies variation in the underlying performance. A research question is how much of this should be exposed to the application. Dynamic spectrum allocation makes this a richer issue. 3) location awareness, which will both be part of a service platform, a new component of privacy-related concerns and a tool to influence low-level functions like routing. 4) multihop. One speculation about multihop is that it can support emerging modes of social networking, such as flash crowds or clusters of vehicles. (Perhaps flash crowds and traffic jams are the same thing.) Again, the question is whether, and to what extent, this function should be exposed to the application layer.

4) Cyber-physical, power and energy networks. These networks are not just for information, but for control, for actuation, and for sensing. In all cases, the network is connected with the physical world. In terms of technology, this may imply wireless networking. In some cases it implies a desire for minimum power consumption. For example, in some designs, nodes are engineered with no receiver, but just a transmitter. How does that change the network design? At a fundamental level, these systems will have a fundamental spatio-temporal energy dimension. And the range of potential devices is huge—imagine a billion sensors each producing one bit, or one sensor producing a billion bits. Within this space is a high degree of variation in the criticality of the objective. Control of power grids and interaction among gaming sensors seems very different in this respect. This raises a question of which layer deals with issues such as availability and resilience.

A separate idea is a viewer for the physical world? Imagine looking at the physical world and seeing an overlay with cyber-annotations.

5) An application-based network. The starting point for this discussion was the vision of information everywhere, and the modes of access to it. Looking at the current modes of “application transport”, we see SMS, email, Web etc. Generalizing, we end up talking about everything. The issues that were identified were first, what kind of mechanism should we use to send information. Applications may require multiple paths, for good performance, reliability or security, or availability. Some information may be valued more than others. There may be a need for service creation—placement of caches, setting up peering, etc.. On the network side, there will be a need for signaling, SLA setup, enforcement, policy, etc.

The second and related issue is the design of the application interface. It seems that we will need a richer API, e.g. to specify how many paths, properties of the dissemination model, service components, willingness to pay, etc.

A number of questions were noted by the group but not resolved: are users willing to take find-grained decisions? Will better and more flexible contract systems sufficient? Should the network play any role at all in supporting these services? (The answer to this latter question seems to be yes.)

6) Bottom up design group, This group described itself as a Hegelian antithesis group. They felt that this “use-centric” or “objective-centric” approach was not the place to start. In this view, design is tangled with constraint. In the context of infrastructure-heavy systems such as networks, design is an agreement on constraint. We should design for the “use case” of the things we have not thought of yet. For example, we could take the idea of trustworthy behavior. We should provide tools to allow us to build trust as needed. We should try to design a network that allows us to establish whatever trust we need. So design basic infrastructure from the bottom up, based on least constraint so that it still works.

In this context, the most basic issue is sharing and isolation (or non-interference). When we consider a constraint, we should ask whether it affects the basic forwarding process, or whether it requires us to design something equally global, such as the DNS. If we try to drive design in this way, what resources define the constraint space? For example, what made the Internet possible was switching speeds that were fast enough to switch packets. So what do we have now? We have even faster switching speeds—speeds fast enough to build software-defined radios. Perhaps the next paradigm is “signal switching”. Another fundamental is lots of end points. At the low end, every batch of concrete may come with embedded processors. And there will also be end-points that move.

Discussion

In the discussion that followed the presentation of these six objectives, it was noted that there was considerable overlap among the topics. While the groups had begun thinking that they were talking about distinct topics, at the end the distinction was a matter of emphasis. Several of the discussion could be characterized as taking different cuts at the general question of what “the network we all use” will look like in fifteen years. The objectives of critical infrastructure, and of “the network of things” were also present here in various forms.

The following more detailed points came up in the discussion.

- There were a few specific use cases that were noted, because they might have interesting requirements: business to business applications, financial transactions that need to be real-time and secure, and machine to machine.
- The idea of *constraint*, from the “bottom up” group, was important. It was suggested that we should ask what was constraining the other use cases. We should challenge ourselves to question how many of the perceived constraints are in our heads.
- The point about “willingness to pay” should get more consideration. How can we build this in? is our future driven by economics? If so, does it change country to country or region to region? How can we harmonize this variation.
- The question was asked as to how can we automate things to get humans out of the loop and whether this a good idea?

In general, a high-level challenge is to go through our list of problems and ask why we are not solving them today. Will our efforts just carry this set of limitations forward?

In general, the term “use case” made many folks uncomfortable, but the discussion of objectives and capabilities was exciting.

Step two: requirements for a future network

The discussion of objectives or use-cases was very high-level. To try to dig deeper, the next exercise was to ask what each of these objectives might imply in terms of the requirements or attributes that the network would need to incorporate in order to meet the objective. The assumption underlying this exercise was that to tell whether two objectives were actually the same, or were different, the way to begin was to ask if they implied the same set of requirements on the network. To start this phase of the discussion, Clark again offered a straw-man framework—a list of candidate requirements that could be used to flesh out each of these high-level objectives. He noted that he had derived this list by thinking about one of his candidate objectives—a successor to the “Internet that we all use”. The hypothesis is that other objectives, such as a network for critical infrastructure, might imply a different mix of requirements, or a different “unpacking” of these requirements as we sort out what they really mean.

Security: The Internet of today is marked by a number of serious security issues, including weak defenses against attacks on hosts, attacks that attempt to disrupt communications, attacks on availability (Denial of Service or DoS attacks), and attacks on the proper operation of applications. A future Internet must have a coherent security architecture, which makes clear what role the network, the application, the end node, etc. each has in improving security. The goal of security will have a primary role in shaping the architecture of a successful future Internet.

Availability and resilience: These goals are sometimes lumped into security, but have been identified separately because of their importance, and because availability issues arise in the Internet of today independent of security attacks. Improving availability requires attention to security, to good network management and preventing errors by operators, and to good fault detection and recovery. Again, what is needed is an *architecture* for availability. While the Internet of today deals with specific sorts of faults and component failures (lost packets, links and routers that fail), it does not have an overall architecture.

Better manageability: Management has been a weak aspect of the current Internet from the beginning, to a considerable extent because the shape and nature of the management problem was not clear in the early days of the design. Among other things, it was not clear what aspects of network operation would (or should) involve human operators, and which would preferably be automated if possible.

Economic viability: A fundamental fact of the current Internet is that the physical assets out of which it built, the links, routers, wireless towers, etc., are expensive. These assets, often collectively called *facilities*, come into existence only if some actor chooses to invest in them. There is a tension between a core value of the current Internet—its open platform quality, and the desire of investors to capture the benefits of their investment. Clark called the working out of this tension the *fundamental tussle*. Any proposal for a future Internet must of necessity take a stance in this space. One tilts the fundamental tussle toward vertical integration and a more closed architecture if additional functions are bundled with (or to any extent replace) the basic forwarding function.

Suitable for the needs of society: The Internet is not just a technical artifact connecting computers, but a social artifact connecting people, deeply embedded in society. The success of a future Internet will depend on how it deals with important social issues (for example, *identity*), some of which will come to the front as we address issues above such as security.

Longevity: The proposed design must remain useful over time. One view is that a long-lived network must be *evolvable*; it must have the adaptability and flexibility to deal with changing requirements, while remaining architecturally coherent. The goal of evolution over time is closely linked to the goal of operating in different ways in different regions, in response to regional requirements such as security. On the other hand, a factor that can contribute to longevity is the *stability* of the system: the ability of the system to provide a platform that does not change in disruptive ways.

Support for tomorrow's computing: The Internet arose as a technology to hook computers together, so as the shape of computing evolves, so should the Internet. In 10 years, the dominant form of computing will not be the PC, nor even the PDA, but most probably the small, embedded processor acting as a sensor or actuator. At the same time, high-end processing will continue to grow, with huge server farms, cloud computing and the like. Any future Internet must somehow take this wide spectrum of computation into account.

Utilize tomorrow's networking: At least two technologies will be basic to tomorrow's networks, wireless and optical. Wireless (and mobility) implies new sorts of routing, the need for intermittent connectivity, and dealing with losses. Advanced optical networks can offer rapid reconfiguration of the network connectivity graph, which again has large implications for routing and traffic engineering.

Support tomorrow's applications: Today's Internet has proved versatile and flexible in supporting a range of applications. There is not some killer application that is blocked from emerging because of the current Internet. None the less, applications of today and tomorrow present requirements that a future Internet should take into account. These include a range of security requirements, support for highly available applications, new sorts of naming, and the like.

Further discussion

The breakout groups re-assembled for a second set of discussions to consider whether, and to what extent, this straw-man list of requirements applied to each of their different network objectives. They were then invited to modify or expand this list. In fact, at the high level at which this list was posed, there were few amendments proposed. The issue of energy efficiency and "being green" was an important addition to the list. The issues seem to emerge when one goes deeper into what is really implied by terms such as "security", which has to be unpacked in the context of any specific requirement.

Availability and resilience

With respect to availability and resilience, the group noted that the issue arises at several levels, from hardware connectivity up through applications to human users. Humans connect sporadically, so it may be useful to have proxy agents that are always present to represent the interests of the user.

It was also noted that different circumstances call for different degrees of resilience, so the architecture should not presume a single answer, but should allow different objectives to be "dialed in" to the running system.

Security

With respect to security, the group noted that the term had to be refined before it could be acted on. Some dimensions of security, such as provenance of information, deserve more attention. Unless we have some sort of metrics to describe security, it is hard to determine if we are getting better. The network can only deal with certain dimensions of security, and one point of view is that the network should only do the minimum necessary to allow applications to achieve their own needed level of security. For example, an open question is whether devices like NATs will be needed in a future net, and if so, how can they be non-disruptive.

Several people noted that we should try to learn from history, and be clear as to why the situation is so bad today. Learning from history, as well as designing effective systems going forward, requires open reporting of information.

The relationship between security and privacy received a lot of attention, and was the topic of a later presentation –see below.

There was some discussion about the need for the requirements for longevity, and whether the challenge of “better” manageability, were actually well-posed. A requirement for usability was noted. But it seemed as if this list worked for many participants as a high-level starting point to test their own ideas.

Wrapping up day one

At the end of day one, these various objectives were now set aside. They had served their purpose as a mechanism to get the meeting to try to frame the problem before exploring alternative approaches. But it was understood that with only an hour or two of discussion, this list was not a well-developed set of alternatives. Any proposal that is developed to respond to the Dear Colleague Letter will have to describe the objective that its candidate network will meet, but the objective will derive from a combination of top-down (e.g. objective-driven) and bottom up (e.g. mechanism-constrained) considerations, and will thus be a distinctive part of each proposal. With this in mind, day two was dedicated to discussion of design approaches—how a network might be designed. What might a solution look like, now that we have a rough idea of the problem space?

Step three: candidate approaches to designing a future Internet

Volunteers were solicited to give ten-minute talks on possible design approaches for a future Internet, and ten presentations were selected. Each of these talks, most of which represented work that had proceeded prior to this meeting, demonstrated a distinctive mix of objective, requirements, and approaches.

Post-modern architecture (speaker—Ken Calvert)

The point of view behind this proposal is that a network architecture for tomorrow needs to take more explicit account of the various stake-holders that make up the system. Customers, service providers, application developers, governments, and so on have concerns and objectives. The current “narrow waist” of the Internet works fine as a way to create a useful data plane, but does not allow these various actors to establish policy—it masks what is going on inside the network, and does not provide the right set of management tools. For example, the customer should have more control of route selection, but be more explicitly responsible for any packets they send. This requirement requires that we “re-factor” the basic functions of the network.

The RISE Internet Architecture (speaker—Nick McKeown)

RISE is an acronym for “Reliability, Innovation, Security and Economic viability”. The “mission statement” of this project is given as follows:

The Internet architecture should enable users to receive the data-delivery services they need, over time and in many contexts. These services should be delivered efficiently, reliably and securely over an economically viable infrastructure, in a manner that is compatible with relevant societal norms.

Some basic design principles are:

- Path diversity, and user choice over paths.
- Virtualization of facilities, and the enabling of “software-defined networks”.
- Extensive measurement capabilities.
- Decoupling of identity and provenance from trust.
- Decoupling of information security from path security.
- Design of a new generation of application APIs.

Choices, Clouds and Contracts (speakers—Nirmala Shenoy and Murat Yuksel)

This proposal has at its core a structured and regular approach to the design of addressing and forwarding. In this scheme a “cloud” is an autonomous entity of definable granularity, like an ISP, a pop or an AS. Clouds exist in tiers, and in this tiered structure, clouds can “float” and attach to other clouds at any tier, if policy permits. The goal is to exploit the topological hierarchy of inter- and intra-ISP structure in protocol design in an efficient way based on tier-based address aggregation.

Combined with this idea is the idea of “contract switching”, which is a mechanism by which providers of connectivity “in the clouds” can describe their capabilities (e.g. path fragments) and offer them in a market-based routing scheme. The objective is to embed flexible economic tools into inter-domain routing.

The SILO project (speaker—Rudra Dutta)

The vision of the SILO project is a set of design principles that provide increased service flexibility, manageability and evolvability by replacing the current Internet principles (layering, peering and a narrow waist) with a design based on fine-grained service elements that can be composed, based on an ontology of functions and interfaces. Applications should be able to specify high-level functional requirements, and request that elements be composed to meet these requirements. This can ideally be done automatically, and for common cases can be “pre-compiled”, thus capturing prior wisdom about effective service composition. In this way, common cases can be met efficiently, but other needs are not precluded by a rigid design.

Content-Centric Networks (speaker—Van Jacobson)

This proposal is based on the observation that what people care about is not devices, but content—information. People exchange content, and create it for others to use. The device (e.g. the computer) is a low-level aid to this task. The content-centric architecture responds to this observation by making the addresses in packets correspond to information elements, rather than machines. One sends a packet to a piece of information, asking that it be returned, and gets back a packet from that address containing the requested information.

This reformulation of what addresses means leads directly to new modes of dissemination and in-network storage, such as any-to-any dissemination, demand-driven buffering and directed diffusion distribution. Information is signed, so validity and provenance are explicit properties of the scheme. And since one does not receive information except when asking for it, one only receives information that is relevant. So certain security attributes are intrinsic to this scheme. The traditional service of the Internet today, conversations between explicit endpoints, can be derived as a tiny part of the total space.

Security and industry structure (speaker—David Clark)

This proposal describes an architecture suited to the requirements of “the Internet we all use”—a successor to today’s Internet that better addresses issues of security, economics, and longevity, as well as tomorrow’s networking technology, end-points and applications.

In common with other emerging designs, this proposal argues for a service layer “between” a packet forwarding layer and an application layer. The packet layer should be thought of as an “underlay” for the service layer. However, the two layers should be distinct, based on arguments deriving from the desire to induce an industry structure where the operator of the packet forwarding service does not control the service layer. The service layer, subject to certain constraints, can evolve over time driven by competitive innovation. The packet layer addresses (some) issues of security, usage management, and management; its central feature is a regionalized set of address spaces that define security and management domains. To assist application designers in exploiting the features of these two layers, this proposal includes a set of “application design patterns” that help guide application builders toward designs that complete the picture for security, management, usability and the like.

Virtualization (speaker—Aaron Falk)

Virtualization is a strategy for sharing physical infrastructure. Links and computing elements (e.g. routers) are virtualized, or partitioned into “slices”, and a management interface allows a developer of a higher-level service (e.g. a packet-level architecture) to requisition a slice to support that service.

The idea of virtual links and processors is not new. This proposal builds on that understanding to tie all those resources into a common management regime, and to link the various “virtual parts” together.

The benefits of this idea include ease of deployment for a new higher-level architecture, isolation of one architecture from another (with resulting benefits of security and manageability), and migration from the current to a future internet. Virtualization reduces the need to get a new design correct from the start, and allows both vendors and service providers to innovate without disrupting the core business.

Recursive Network Architecture (speaker—Joe Touch)

This proposal is motivated by an insight about regularity in design. Many protocols (e.g. protocols at different “layers”) seem very similar in abstract structure. What if the ideal abstract architecture had only one protocol? What if there was one abstract mechanism that could be recursively realized (in concrete terms) at different layers and across different scopes?

The Recursive Network Architecture (RNA) is based on recursive composition of a single, configurable protocol structure, derived from basic principles of multiparty interaction. RNA unifies forwarding, layering, recursion, and resolution in a single framework replicated

throughout a stack, and is compatible with the Internet as a degenerate case. RNA provides an opportunity to support dynamic late-binding of stack components to support increased flux in network conditions and capabilities.

An initial RNA prototype supported by the NSF FIND program was implemented to explore a mechanism based on this framework and its implications. Extending RNA as a complete architecture includes incorporating protocol negotiation, developing concepts of recursive security and recursive network management, and exploring the interaction between routing and discovery. Optimization can be supported through caching, precomputation, and dampening. RNA is natively compatible with compartmentalized security, and encourages delegation and enforcement through boundary verification.

Minimal Internet Architecture (speaker—Henning Schulzrinne)

This proposal reflects both a view about tomorrow's requirements and a design philosophy. The philosophy is that the minimal nature of the Internet design has been a powerful contribution to generality and innovation, and any future Internet should specify as little as possible. At the same time, the current Internet does not address some key needs with respect to service creation, security and management.

With respect to the low level "packet" service, for example, the MIA specification just encompasses the delivery of packets from point A to point B, where A and B are globally unique identifiers, which could refer to machines, content or humans. The current Internet experimented with adding new services at this network layer, such as QoS, multicast, mobility and security. All have been, at best, partial successes. Instead, we should ask what set of minimal interfaces to a service platform (including processing and storage) will allow us to compose higher-level services with the same degree of minimality. In this respect, this proposal is a virtualization proposal at a higher level.

Internet 3.0 (speaker—Raj Jain)

Our goal is to develop an architecture that explicitly recognizes *users* and *contents* and their economic relationships. This will allow new class of requirements to be expressed and will shape the services that the network can provide enabling new business models and applications, e.g., user-to-content connections enable users to continue using data (e.g., watching a movie) as they move, change devices, or handle disruptions.

The key to innovation in the Future Internet will be a strengthening of rights and obligations of the diverse owners of network components that include devices and content. The economic reality of the Internet requires networks to provide means of enforcement and negotiation of policies including security, privacy, quality of service, energy efficiency in presence of mobility and disruptions. Internet 3.0's goal is to allow *policy-based communication* that is aware of different policies at the granularity of users, content, hosts, or infrastructure.

"Mobility First" (speaker—V. Arun)

The Mobility First proposal is based on the recognition that Internet usage is changing very rapidly from fixed hosts to mobile devices such as cell phones, portable computers, machines and sensors: it is anticipated that by 2015, mobile/wireless devices will vastly outnumber fixed hosts (~10B mobiles vs. ~1B hosts). Our vision is that of a future Internet architecture which supports mobile and embedded devices as "first-class" users (no gateways!), thus enabling a variety of new applications efficiently, securely, and at scale.

Such a network should lead to new types of economic models resulting from cellular/Internet convergence, while also serving as a key enabler for emerging cyber-physical or M2M applications involving networked observation and control of the physical world.

The central tenets of the proposed “mobility first” architecture are: dynamic end-point and network mobility as the norm; robustness in the presence of disconnection; support for heterogeneity of devices and technologies; opportunistic communications for energy constrained wireless devices; addressability and routing of content; multicast, anycast and multi-homing modes as basic services; and location as a fundamental network attribute.

The key building blocks of the proposed mobility first architecture are:

- Separation of naming from addressing.
- Routing protocols that deal with changing points of attachment, disconnections, and varying link quality.
- Robust hop-by-hop transport protocol vs. TCP-like end-to-end connections.
- Content naming, routing and storage within the network.
- Global location service and optional location-aware modes for routing and security.
- Security architecture based on strong authentication and encryption as the norm.
- Privacy architecture that supports different quantitative levels of privacy assurance.
- Network management with participation of end-user nodes.
- Network virtualization and programmability for creating distinct customized protocols.

Step 4: key research challenges

Following the presentations of these candidate approaches, the attendees were invited to break up into groups to further explore each of the ideas. The groups were specifically asked to identify perhaps three key research challenges raised by their approach. There were two motivations for this exercise: to see whether there were common research challenges, and to catalog research questions for those who might not want to participate in this funding opportunity but to submit a more traditional research proposal to NSF. Here is a list of some of the more general issues that were raised:

Services

Several proposals advocated that there should be a service layer between what we traditionally describe today as the packet forwarding layer and the application layer. There were a wide range of conceptions of this service layer, and what it might do, including delay tolerant information delivery, information dissemination more generally, security services such as protection from attack, and so on. The following questions emerged:

- How are new services created? How should service elements be composed, with respect to operating models, management, interfaces, and so on. Is automated composition driven by a service ontology a realistic idea? What are the minimal interfaces required?
- What are the economic implications of a service layer? What would such a layer imply with respect to industry structure? Should this be a layer of competing services, or a layer integrated with a packet forwarding layer?

- What is the right granularity for service elements, and how does this influence utility and generality?

Storage “in the network”

Several proposals advocated that storage (and also computation) should be “first-class” elements of a future network. However, this high-level proposal raises a number of questions.

- In common with the discussion of services in general, should such storage be integrated into a packet forwarding layer, or a separate layer above it, perhaps at a layer where service providers compete to offer services?
- How does the requirement for mobility affect the architecture for storage in the network? The particular cost tradeoffs of transmission and storage suggest that mobility may raise special requirements.
- Should storage be assumed to be reliable? What is the correct “service model” for storage?

Scale

Several proposals raise issues of scale, and the range of scales across which the proposal should work.

- If resources are to be “virtualized”, how many virtual systems should be anticipated? Ten, or a million? What aspects of the approach would have to change to accommodate this sort of variation?
- The future computing world will include both high-end powerful computing elements (e.g. large data centers and supercomputers) and low-performance sensors. To what extent can a single architecture serve both? (And do the issues of mobility change the answer?)

Trustworthy management interfaces

Proposals for virtualization require an interface to set up virtual systems. How can we address the requirement that this interface be trustworthy?

Proposals for a service layer require an interface to configure services out of service elements. How can this process be made trustworthy?

Layering and economics

A number of proposals described a relayering or remodularization of the architecture. Each of these raises questions about the incentives of investment and industry structure.

Virtualization schemes must resolve the tension between owner of resources and people who provide on top. Some schemes take a specific point of view, others (e.g. the proposal by Clark) presume that the facilities owners can provide service building blocks, on top of which competition will drive innovation.

User choice

Several schemes give the user more control over the service they invoke, whether this is choice in routing, cost-performance tradeoffs, or recovery from errors and service failures. These proposals raise important issues with respect to usability, and preventing the user from attacking the network via these tools.

Application design

Several proposals directly address the design of applications. We need to study how to build applications that can take advantage of additional services. Do we need tool kits, design patterns, libraries or something else? How can we help application designers achieve goals such as security, robust operation and manageability?

Other research questions

During the Tuesday presentations, we asked participants to jot down any other notes or questions that came up, and some of these are listed below.

- DDoS – will the future Internet offer better tools to handle it? If so, how? If not, why isn't it necessary?
- In addition to competition between providers (firms), we also need to facilitate competition between protocols and architectures, etc.
- What level of accountability should exist in the network (e.g. self-certifying vs. "legal name" ID)?
- What is "physical"?
- How can we have efficient support for accountability, provenance, trace-back and any other buzzword capability to counter forgery?
- What is the thin waist of security? Should we look for one?
- How much progress have we made in 30 years, or in 10. If we met 10 years from now, could we have exactly the same meeting?
- We have been talking about the future Internets based on new requirements. But do we have new techniques/tools/knowledge that we can use that the original designers didn't?
- Which "future architectures" will best map to virtualized physical infrastructure?
- How do we address privacy and security concerns in the cloud
- How do we understand "overhead" in these new architectures?

Step 5: elaboration

The material above summarizes (imperfectly and briefly) two days of interleaved breakout discussions and plenary reports. The next two days were set aside for a variety of breakout group activities. Some groups met to elaborate on the material presented above. Other groups continued to discuss requirements (e.g. how to build a network for the next three billion people), with the goal of discovering what this objective might mean for architecture.

The centrality of networks that are trustworthy

NSF has stated that while networks with different objectives may have to meet different requirements, no proposal should ignore the requirement for trustworthy operation and suitable security. A number of security experts were present at this meeting, and they had a side-meeting to try to distill what advice they might offer, independent of the specifics of a particular architecture. They presented this advice to the larger group.

- First, any network of global scope will be composed of regions operated by different actors with different interests. Some regions may be trustworthy, some not. It does not make sense to require that "the network" do something, since not all parts can be equally trusted to do it as expected. However, a particular end-point may be

attached to a region of the network that is trustworthy, and the user should be able to count on that region doing things to enhance the security of the end-node.

As a means to make the term “security” more specific and actionable, the group worked with this list of sub-objectives.

- Self-defense for the network
- Protecting communication among end-points.
- Mitigating attacks on machines.
- Mitigating DDoS attacks.

Whatever a network does to improve the security of attached end-nodes, it must first be able to protect *itself* from attack, and this objective includes attacks by one region of the network on another. Examples of issues might include malformed or false routing messages, flooding attacks on links, misuse of new service-creation protocols and the like.

With respect to protecting communication among willing end-points, the traditional taxonomy of confidentiality, integrity and availability is helpful. Confidentiality and integrity can be addressed on an end to end basis with encryption. This leaves us with the challenge of dealing with availability in the presence of attacks by parts of the network. We must also decide whether (and to what extent) we should deal with the issue of traffic analysis (see below). An architecture that declares that it will not deal with these issues must justify this position.

With respect to protecting the host from attacks, different proposals may take different points of view about whether and to what extent the network (and network-based services) can be part of a overall mechanism to protect hosts from attack. But any proposal must take a position on this point. Additionally, the network should be prepared to make network security and diagnosis information available to participants.

With respect to mitigating DDoS attacks, again different proposals may take different approaches to dealing with DDoS attacks, but some approach is required—this is a problem that a network must (at least in part) deal with, since it involves network resources.

There were a number of additional issues that the group discussed but did not reach consensus:

- Privacy and confidentiality of users
- Security metrics built into the architecture
- Network recovery from security failures
- Architectural audit trails
- Granularity of identities
- Network support of data objects.

This last point raised definitional questions about what the word “architecture” means. Several people (and several of the proposals presented to the plenary group) proposed that information objects should be signed, so that information assurance (e.g. authenticity) can be separated from network security. However, another point of view is that since such a feature cannot be enforced by an architecture, it should not be included in the description of the architecture. The issue is whether “architecture” can include guidance to the users of that architecture.

Traffic analysis

Transparency is a feature of the current network architecture. Internet communications by default are akin to postcards--intermediaries can view information about the sender, recipient, and the content of the communication. Some users employ encryption to make communications less transparent. However, encryption does not obscure the identities of the parties to the communication, and other facts can be inferred from the communication even when it is encrypted. Thus, traffic analysis has become a staple technology of both private-sector and public-sector actors that wish to monitor users of the internet.

To promote a trustworthy environment, a next generation architecture should address the privacy and security risks associated with analysis of routing and other information contained within packets. In one point of view, a trustworthy communications infrastructure would enable obfuscation of both traffic data about a communication, and the content of the communication itself.

Other activities and discussions

The Mini-Sandpit Forum

As groups started to form on Tuesday afternoon, there were a handful of participants who weren't yet inspired by the topics from the morning presentations. These people were invited to attend a Sandpit Forum, a short meeting modeled on the framework of the Sandpit meeting format that the facilitation team often employs. Eight people attended, with stated reasons like they were looking for synergy with different people, they wanted to have more influence in the development of a group, or to have a unique interest that the other groups hadn't yet embraced or incorporated.

The objective of the Forum was to attempt to define some other problems that perhaps weren't addressed in the other proposals, and give participants a chance to verbalize these and discuss them. With this, participants could either start a new group their own, or get clarity about the contribution they might want to try to insert into the other existing groups.

Most of the participants were somewhat (varying in degree) dissatisfied with the presentations made so far, feeling that they hadn't sufficiently addressed or incorporated key issues. This is by no means the complete list of issues, but these element were strongly emphasized or shared by several of the participants:

- Energy efficiency and environmental impact, intelligent and efficient power supplies
- Continuation of smart infrastructure from core to edge
- Use of cognitive tools to organize data and instant access to information in a cognitive sense.
- Anticipating the acceleration of technology
- Preserving competition; protecting against monopolization and single- company, organization or government dominance
- Intellectual property rights and ownership of data

Participants were encouraged to list the problems and challenges that currently keep us from achieving some of these goals now. About 80 problem statements (ranging in scope) were generated. The full list is included as an appendix to this report. The next step would

have been to cluster these into topic areas, and use a ladder-of-abstraction technique to mine them for more probing and provocative questions. However, the priorities of the participants did not permit these next steps.

A network for the next three billion (Group lead—Ellen Zegura)

On the first day of the summit, a brainstorming exercise about use cases identified “the 3 Billion people in the less developed world” as a use case with some distinct characteristics. On the second morning of the summit, a set of candidate future architectures were presented, none of which appeared to address the 3B challenges head on. In response, a group was formed to think more deeply about the issues and their implications for architecture design.

By Thursday, the group variously known as 3 Billion (3B), Other 3 Billion (O3B), and Internet for Everyone (Really) (IFER), came to the following observations:

- The developing world has constraints that are unique and/or more acute than in the developed world; at the same time, other resources, such as human capital and spectrum are often more abundant.
- Some of those constraints seem to have architectural implications, e.g., power availability and reliability
- Any architecture for *everyone* must accommodate these constraints
- Starting from a clean slate with these constraints may lead to a good architecture for everyone (robust, green, usable, secure...) and to interesting components for other architectures. Such an architecture should allow for implementations that make varying tradeoffs between capital and operating cost, system performance, system features, etc. depending on the environment in which it is deployed.
- We might avoid “second system syndrome” by designing starting with developing world constraints rather than starting with developed world desiderata.
-

The group further identified a list of primary constraints that any architecture must satisfy to be suitable in the current developing world. These include: robust to significant, frequent disruptions, both planned and unplanned; low cost; low power; support for mobility as the common case and of many forms (including device sharing); low barrier to innovation by local people for local needs; appropriate and usable security; appropriate and usable manageability.

These constraints seem to imply a set of useful architectural components including developer services, management services, storage as a first class object, simple and usable security mechanisms, visibility mechanisms, robust-yet-low-cost mechanisms, a mobility architecture, DTN, and pervasive considerations for power and other scarce resource.

As final remarks, those interested in this area are willing and eager to synthesize best ideas from other architectures, perhaps more substantially than in other efforts. Those involved share the conviction that an architecture that starts with these considerations can be robust, green, secure and manageable, hence good for everyone.

Green networks (speaker, Ken Christensen)

There are strong arguments for energy efficiency in the ICT sector, and networking in particular. Sustainability may be the biggest challenge of the 21st century. ICT contributes 2% of CO₂, the same as aviation (according to the Gartner group), and PCs consume 2% of electricity in US, data centers about 1.5% (according to the EPA).

Since many ITC assets are often lightly loaded, one path to energy efficiency is to develop devices where the instantaneous power consumption is proportional to actual instantaneous load, as opposed to total potential capacity. One step in this direction is Energy Efficient Ethernet.

There are a number of issues with respect to the power consumption of hosts. At present, most hosts need to be connected the network and active at all times. Some reasons are perhaps easy to mitigate, such as keeping DHCP information active, but some are more complex, such as participation in peer-to-peer systems. There is standards addressing idea of a *network connectivity proxy*, such as ECMA TC32-TG21 – Proxying support for sleep modes.

A future architecture should include a new view of network connectivity, with attention to issues such as:

- Assistants (proxies?)
- Exposing selective connectivity
- Evolving soft state
- Host-based control
- Application primitives
- Security

The economics of competition in networks (speaker—John Chuang)

Conventional (economics) wisdom is that competition is a good thing. Competition and the discipline of consumer selection among competing alternatives drives innovation and increased quality and consumer satisfaction. This would seem to imply that we should design our networks to allow for competition. This simple statement, however, does not reveal a useful level of detail: there are many dimensions: where, what, who, why, when, how?

Here is a simple, two-dimensional matrix that can help sort out the landscape of competition in networks.

	Edge (E)	Core (C)
Logical (L)	Number of access service providers	Number & coverage of backbone networks
Physical (P)	Number of wires (or wireless channels) into each home	Number & coverage of wide-area physical networks

Using this matrix, we can catalog various sorts of competition, or (in other words) the different sorts of choice options that might exist.

- End users choose: EL (multihoming), EP (facilities-based competition), CL (user directed routing)
- Access network providers choose: EP (open access), CL (transit providers)
- Backbone networks choose: CL (interconnection, inter-domain routing), CP (bandwidth markets)

Given this list of options, we can now catalog a number of interactions that might occur among them, many of which can be found in one or another proposal for a network design.

- EP: facilities-based competition relieves pressure for open access (or vice-versa)
 - e.g., local loop open access before/after proliferation of mobile telephony service
- EL and CL: multihoming (1st hop choice) and user-directed routing (ith hop choice) as *imperfect substitutes*
 - What is optimal form of “end user empowerment”?
- CP and CL: interaction between physical topology and interconnection
- Vertical integration and expansion:
 - What if one firm operates in more than one cell?
 - E.g., CP and CL; EP and EL; EL and CL; EP and CP; all four

This talk did not express a preference among these options, but instead just pointed out that there are many options if one decides to design a network that exploits the discipline of competition, and that these different choices will in some cases imperfectly substitute for each other and will impose different pressures on innovation and investment.

Who needs applications anyway (speaker—Jeff Burke)

It has been said (perhaps cynically) that asking the application designers what the network should do is a clear road to failure. Perhaps this saying arises from asking them the wrong question.

The problem faced by application designers is *designability*, which might be seen as a superset of manageability, one of our candidate requirements. Once we as architects and protocol designers have done our work, others come who build physical networks out of our designs, and application designers build complete systems out of our designs. An important way to learn from creators is to understand their *process*, not just their products. Perhaps we, as network architects, should study and *facilitate design processes*, as an alternative to use cases: the design processes of which the network is now both *component* and *facilitator*.

Unlike other infrastructure, networks have the potential to embed knowledge *that may only be gathered in one ‘place’ at the time of design* that impacts everything from physical topology, to addressing, NAT and many other things. Other materials and infrastructures *cannot* by their nature intrinsically incorporate information about how their subcomponents were designed and connected. They can’t remember their mapping between principles (or business logic, or policy, if you like) and their physical reality. They can’t adopt aspects of the ontologies of their creators, while retaining generality at an underlying level.

As an application designer, I see this group struggling mightily, across a variety of disciplines, to keep track of the mapping between what’s important in their design process and network elements. Can we *facilitate the persistence and evolution of design knowledge*

throughout the life of a network? Could architectural components help network owners, designers, and users discover and manage their own specific knowledge about specific networks, with no loss in generality? As our architecture(s) improve, could they better embody and communicate design decisions and principles that are increasingly hidden by complexity?

Privacy in cyberspace (speaker—Chris Hoofnagle)

This talk quickly covered the following topics:

- Fair information practices
- Sources of privacy law
- Electronic Communications Privacy Act of 1986
- Third party doctrine
- Free speech & anonymity

Privacy is not secrecy, nor is it security. Privacy, in cyberspace, relates in part to how data is used, and the rights and constraints on its use. In this respect, it is important to understand *fair information practices*, which include collection limitations, data quality, purpose specification, use limitation, security safeguards, openness, individual participation, and accountability.

There are many sources of privacy policy in the United States, including contract, property law, torts, norms, inefficiency, self-regulatory codes, administrative guidelines, regulations and decisions (e.g. FCC and FTC), statutory decisions (both state and federal), constitutions (both state and federal), and international constraints.

The US has no comprehensive privacy law to regulate data collection, use, or dissemination. Many privacy protections are accidents and outrages of history.

The rules are totally incoherent: your cable viewing records have stronger protections than your medical records or your book buying habits. However, all business practices are subject to state and federal “unfair and deceptive trade practices” statutes.

Currently, different laws cover data while it is being transmitted and while it is being stored, with different limitations on the ability of police to obtain access to the information (depending, for example, on the age of the stored information), and different penalties for abuse.

The Third Party Doctrine states “...What a person knowingly exposes to the public, even in his own home or office, is not a subject of Fourth Amendment protection. But what he seeks to preserve as private, even in an area accessible to the public, may be constitutionally protected.” System designs that places data within the possession of the individual enjoys more protection. This current reality (which might be changed) has implications for the storage of personal information “in the cloud”.

Finally, the future of anonymity and free speech is unclear. Tools that are intended to facilitate anonymity and speech may serve only to single out and spotlight those who choose to use them.

Mentor advice to the group

As part of the meeting, eight mentors were asked to attend, with the charter of encouraging groups to think creatively and boldly, to help identify groups with common approaches and ideas, and to offer advice about how to improve and shape emerging ideas.

During the meeting, the mentors were asked to meet with the breakout groups, and to give advice following plenary presentations. Their specific comments have been incorporated into this report, but they presented some general advice about what they would expect to see in a competitive proposal in this context. They offered the following general outline:

Why this idea? A proposal must lead off with a mission statement, or a topic sentence that describes the vision; or an “elevator pitch” The mentors felt that most of the proposals had such a vision, but it was sometimes buried. Sometimes team got so deep into mechanisms it seemed to have forgotten why it went looking for those mechanisms.

The path to progress. Why should a listener/reader believe that your team would make progress on the hard problems? There is a nice balance that any competitive proposal must strike here. It is not necessary that the work be done in order to be competitive. On the other hand, just identifying the problems is not sufficient, and saying “We know how to solve this” does not inspire confidence. A proposal should show that you’ve got an initial path (a first few steps) towards a possible solution, and even better, describe what you would try next if the initial approach fails.

Identify the gaps. It is not likely that a proposal will be able to address fully all of the relevant requirements. However, a failure to clearly identify gaps suggests a lack of a comprehensive perspective on your own work. It is important to characterize the gap clearly enough that an expert in that area understands where to start helping. Simply saying we need “security” or “better economic models” isn’t enough.

Required objectives. NSF has made clear that attention to certain requirements, such as trustworthy operation, will be required in any competitive proposal. Lack of a good security model will imperil your proposal, and homegrown security that hasn’t been vetted by an expert is unwise.

Meeting summary (David Clark—one man’s view)

While a number of diverse objectives and approaches were put forward at this meeting, in some respects the starting points showed a high degree of alignment.

With respect to objectives for a future network, many of the ideas centered on what we had initially called “the network we all use”—the successor to the Internet that links us together today. A network that makes information and people a part of the design is centered on that objective, as is the objective of serving the next three billion people as well as the billion we serve today. Networks of devices also received attention, but the meeting did not identify other high-level objectives that were strongly divergent from these.

With respect to requirements, the list that formed our starting point remained largely unmodified. There was discussion of whether longevity was a requirement in its own right, and whether we had properly captured the goal that related to manageability. We agreed that terms such as “security” were at such a high level that they needed to be made concrete

in the context of a specific objective. But the initial list continued to serve throughout the meeting.

With respect to approaches—different points of view about how to meet these objectives and satisfy the resulting requirements—there was both some commonality and some significant divergence.

One way to classify the various proposals is to ask what their common point of agreement is—what the core interfaces of the architecture are. Compared to the point of agreement in the Internet, the “narrow waist” at the IP layer, some proposals push the point of agreement down, and some push it up.

Virtualization schemes push that point of agreement down, and into the management system. In low-level schemes for virtualization, each of the elements—links, processors, etc.—appears to the next layer as a virtualized slice of its native appearance. There is no abstraction or overlay of a common interface. The common interface is the management system by which a slice is requested, allocated and controlled.

In contrast, *information-centric* schemes push the point of agreement higher—to the point where information objects are named and retrieved. In some proposals, such as Content-Centric Networks, the information layer becomes the *only* point of common agreement—machines and services no longer have names, and are not directly invoked in the design.

Several schemes have a two-layer character, with some sort of packet delivery layer as an underlay for some sort of service layer that in turn supports applications. The diversity arises in how these two layers are designed. The packet layer must deal with issues such as resource allocation (congestion), security, the economics of investment in facilities, and the like—traditional concerns that take on a new character if this layer is an underlay for a service layer.

In some conceptions of the service layer, the service architecture is essentially one of virtualization at a higher layer—processing and storage nodes distributed throughout the network can be sliced and made available to high-level service providers. In these schemes, the architecture does not speak to what those services might be, but just permits them to be instantiated. The minimalist point of view suggests that the architecture should do nothing more than what is necessary to support service instantiation at this level, and leaves the details of how services are designed and installed as an unconstrained exercise for the service designer.

In another view, the network is seeded with more complex service building blocks, which can be composed as needed to build integrated services. The specification of a desired service might be provided by an application when it is needed, and would be composed automatically out of the building blocks, perhaps guided by some sort of service element ontology.

Several proposals recognize that the user—the actual person—needs to have some representation in the architecture of tomorrow. This representation is not just about *identity*—a means to associate rights and responsibilities with a user. It is about building systems in which the person and the context in which that person functions, the social network, the location, and the like, is available so that applications can tailor their behavior accordingly.

Finally, there was attention to the design of applications themselves. While an architecture cannot *constrain* what an application does, this should not mean the architecture should be silent as to what an application *should* do in order to achieve overall goals such as secure, resilient, and robust operation.

Some proposals are driven by new technology, such as wireless. The designs implied by that focus seem consistent with those that arise from other motivations, such as flexible service offerings or healthy economic incentives. Many details may vary, but the focus on information, services and people emerges as a common theme.

Of course, the devil is in the detail, to use an old saying, and as different proposals deal with issues such as security and management, many different concepts emerge. Exploring that diversity was an exciting aspect of this summit.

Appendix: Raw “how to” questions from the Sandpit Forum

- How to solve scalability problems?
- How to deploy public cryptographic keys?
- How to have incremental infrastructure – like cable TV?
- How to create incentives to nudge the network?
- How to create an “accordion” effect to deal with accelerated technology?
- How will real-time control of physical devices over the Internet be accomplished?
- How to define classes of service?
- What about pay-per-view-per-quality?
- How to make every contributor make a buck? (ISPs, information sellers, consumers)
- How to create/maintain competition in the last mile? (DSL/EMTS/Wireless/Sat/Cell/Multi hop/Wifi)
- What about remote manufacturing automated factors with human supervision and exception handling?
- What about autonomous vehicle transportation?
- What about remote medicine?
- How to nudge the network?
- What about holographic interactive communications?
- How to build an Internet that will enable future applications, which are socially as much as (or more than) technically shapes?
- How to make it interactive voice and video P2P?
- How to deliver video content?
- How to look beyond throwing bandwidth to the problem?
- What about end to end delay?
- How to handle massive information showers in a room?
- How to define a data-centric network?
- How to increase cognitive capacity and information?
- How to enable users to add data to online knowledge base?
- How to identify data?
- How to get researchers to understand carrier revenue models?
- How to have small information explosions?
- How to increase the carry of wireless services?
- How to keep a single organization or government from having too much control?
- How to define incremental progress?
- How to keep working opportunistically and not get stuck on problems?
- How to know what to call ugly?
- How to keep the Internet free?
- How to make the Internet worldwide?
- How to develop habits to review the past completely?
- How to identify what we did wrong?
- How to reduce power usage?
- How to solve the name revolution problem?
- How to minimize or eradicate spamming?
- If spam is here to stay, what to do about it?
- How to create smart, cognitive, changeable components that organize data?
- How to accelerate technology at a moderate pace?
- How to give the power to the end user?
- How to give more power to the end user?
- How to induce competition within the ISP community?
- Wireless is a different channel than wired/optics and nobody is talking about this, yet this will be the entire edge? How to we bring this into the Internet community?

- How to create a smart infrastructure from core to edge?
- Wireless network management must be integrated to the very edge and into the office. How do we put this intelligence in to every end service?
- How do we solve the PHY, MAC, Network contention in the last 10 feet, that will evolve from massively broadband wireless and the edge?
- How to create incentives for green nets?
- How to create incentives for competition in the last mile, applied to nets and users?
- How to create incentives?
- How to simplify service and data offerings?
- How to know if end users are experienced enough to design and purchase their own service? (Good defaults/templates are needed)
- How to understand ATT's revenue model?
- How to create a basic minimal structure that allows users and community to create from it what they need?
- How to put infrastructure knowledge in every entity on the edge?
- How to look at technology advances as a way to predict the future
- How to reduce the energy use of the infrastructure systems? How to change thinking so that channel off/on is not the mindset?
- How to protect revenue for royalties?
- How about intellectual property and copyright?
- How to solve the intellectual property problem?
- What about exploits? (Identity safety, freeloaders/abusers, denial of service/extortion, PCI needs)
- How to recognize information? How to tag or code information, files and bits?
- How to develop new layer that can 'wormhole' through the ISO?
- How to make energy (for energy use and economics) a "first class" citizen in the network architectures?
- How to deploy intelligent, energy-efficient power supply?
- How to ensure that "trust" is proxy-enabled?
- How to keep any power from controlling a substantial part of network operation?
- How to build good defaults for end users?
- How to not change the network fundamentally?
- How to recognize and manage overhead in new architectures?
- How to re-define intellectual property for and on the Internet?