



Defining a future network: A new research agenda

David Clark
MIT CSAIL
January, 2007

The research challenge

- The Internet is a tremendous success, but...
 - Can we meet tomorrow's needs by incremental improvement of today's design?
 - NSF and its research community have concluded that they must take a leadership position with respect to revolutionary network research, and must provide suitable infrastructure for this research.
-

FIND: An NSF challenge question

1) What are the requirements for the global network of 10 or 15 years from now, and what should that network look like?

To conceive the future, it helps to let go of the present:

2) How would we re-conceive tomorrow's global network today, if we could design it from scratch?

- This is not change for the sake of change, but a chance to free our minds.
-

What I want to cover...

- Why should we undertake to re-conceive the Internet?
 - What are the compelling challenges that justify the research?
 - What features might define a global network of the future?
 - What approaches do the research community have to address these challenges?
-

Isn't today's net good enough?

Must start with serious discussion of requirements:

- It's not just about cool new apps.

Security and robustness.

- Been trying for 20 years--try differently?

Recognize the importance of considerations beyond the technical.

- The economic landscape.
- The social context.
- The international scope.

Easier to manage.

- Really hard intellectual problem.
 - No framework in original design.
-

Security and reliability

Define the objective broadly.

- “Classic” security, availability, resilience.

Hard because:

- Many problems are in the end-hosts.
 - Many problems involve a balance of interests.
 - Among actors, states and societies.
 - We don't have agreement about the objective.
 - Different contexts call for different answers.
 - We don't have a coherent approach.
-

Options for a more secure future

Building a framework around trust.

- Tools to enhance it.
- Tools to interact without it.

Rethink how we manage identity.

- Who controls, who chooses?

Use diffusion to disperse attacks.

Shift focus to security of information, not connections.

Use virtual machines to isolate activities.

Economic landscape

In 1975, it was not clear to the early designers that we were designing the landscape of investment and competition.

- Now it is.

Could we do a better job to shape:

- Regulation (or lack of)?
 - Continued investment and innovation?
 - Options for user choice?
 - Deployment of new services?
 - Health of the value chain?
 - Consider the role of facilities providers, for example.
 - Role of advertising?
-

Social context

Failure to understand and respond to larger social concerns will lead to the eventual rejection of new concepts, and doom the venture.

- The opposite can lead to success.

Examples of important issues.

- Loss of anonymity and privacy.
 - Data mining and profiling.
 - Correlation and linking across people.
 - Tomorrow: location and presence.
 - Issues around access to information.
 - Excessive controls, limits on speech, IPR, forgery.
 - Instability of personal information.
 - Access and ease of use.
 - Variation in local values.
-

Identify desirable social outcomes

Service in times of crisis.

Networks for the developing world.

Managing the personal experience

- Managing personal information
- Ubiquitous health care
- Telepresence
- Citizen as sensor

Focus on the user

Technology drivers

New network technology.

- Usual place to start, but I will get to it later.

New computing technology.

- Whatever computing is, that is what the Internet should support.
- The Internet grew up in a stable “PC” time.
- The cellular industry evolved independently.
- Tomorrow: many different views; sensors, cell phones, embedded processors, \$100 laptops, etc.

Rich space of services and servers.

- Design alternatives will have important influence on personal choice, control, innovation, etc.
-

Define a broad scope to research

A problem with the word “Internet”.

- It is too constraining, but otherwise nobody knows what you are talking about...

Future networking is not just about a new kind of packet.

- Robust content distribution
 - Naming, security, resilience
 - Management and sharing of personal information
 - Real time multi-media distribution
 - Multicast
 - Network-embedded storage and computation
 - Location mgt (human and object)
 - Identity mgt. (human and object)
 - Distributed name management
-

New network technology

Wireless

- Mobility and ubiquitous access
- Not well supported in current Internet
- Great diversity in approach and function
- Raises technical and social issues.
 - Location, identity, security.

Optical technology

- Not just cheap pipes,(but predictions of 10^{-4} cost/performance)
- Rapid reconfiguration of core
- No bottlenecks at edge

What are suitable technology choices for different parts of the world?

Summary--bold objectives

Achieve a materially improved level of security, availability and resilience.

Encourage a healthy industry and continued innovation and growth.

Address rapid migration away from the “PC paradigm”.

Increase usability and utility to a broader base of business, the citizen and the student.

Balance social needs such as privacy and accountability, in the context of diverse cultural and regional norms.

Summary--integrative visions

What are candidate designs of a future global network suited for the needs of 2020?

What are good design principles for large scale future distributed systems and services?

- Alternative approaches will lead to different outcomes in the socio-technical space.
-