



# Contemplating a future Internet

---

David D. Clark  
MIT CSAIL  
July 2007

# What I want to talk about

---

- Process and program structure
    - Just a little
  - Some research ideas
    - The major topic.
  
  - First, the “why”...
-

# The research challenge

---

- The Internet is a tremendous success, but...
  - Can we meet tomorrow's needs by incremental improvement of today's design?
    - Hypothesis: No!
  - US National Science Foundation 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.
-

# 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.
-

# 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.
-

# 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
-

# 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.
-

# Status

---

## Three phases:

- Phase 1 (current phase): exploratory grants, meetings to facilitate interaction and collaboration. Three annual award cycles.
- Phase 2: awards for integrated proposals.
- Phase 3: demonstration of ideas on experimental infrastructure. (GENI)

First year awards made in summer 2006.

Second year proposals now being evaluated.

Starting to develop process of collaboration and consensus.

---

# Structuring the research

---

FIND embodies an “unusual” approach (in the NSF context) to collaboration and cooperation in achieving a large vision.

- Traditional: give a single large grant, and hope.
- Now: use traditional “small grant” merit review process and then create means to encourage working together post-grant.

Now, we must make this collaboration happen internationally.

---

# International activities

---

EU--Eiffel proposal; FIRE

Country-specific activities in Europe

Korea

Japan

Canada (perhaps)

---

# FIND and GENI

---

## FIND is a research agenda

- There are others inside NSF:
  - Cyber-trust
  - SING (theory of networks)
- And there are others in the U.S outside NSF

## GENI is infrastructure to demonstrate research.

- A big idea going after big funding.
  - Support multiple experiments.
    - Network architecture to distributed systems (think PlanetLab).
  - Shape and schedule dictated by the funding strategy.
    - At least two years to funding, so have to launch in parallel.
-

# Some research ideas

---

## Start at the “traditional” layers

- People have trouble conceiving a “not like the Internet” Internet.

But the real action will be at higher layers.

The ideas here based (to some extent) on current funded work in FIND

---

# Start with the basics

---

## ■ Packets?

- Most folks think packets are the right way to go “at the edge”.
  - Lots of bursty traffic, high variance.
- But not in the middle.
  - Deal with aggregates of packets
  - E.g. “circuits”.
  - This needs to be part of the architecture.
    - Management issues.

## ■ Two questions

- Are the packets the same everywhere.
  - Are they a “universal”?
- Should we assume universal interactive connectivity?



# Universal packet: two options

---

- Today's answer: yes.
    - The devil you know.
  - Or: no.
    - Motivation: better exploit the diverse features of wireless (and other?) networks.
      - Assertion: packet processing cost is not the issue
    - Conclusion: conversion must either be “very limited” (not worth the trouble?), or involves knowledge of application semantics.
      - Prior work on ALF.
-

# Application-level converters

---

- Do we want application-level converters in the network?
    - A barrier to the deployment of new applications?
      - Implies: must be optional.
      - Universal packet as a baseline function.
    - A point of excessive control?
      - Implies that third parties must be able to deploy them.
    - Implies they may not be at the physical point of connection. Hmm...
-

# Application services

---

- There *are* going to be application-level servers/services “in the application”, whether or not we have a universal packet.
    - Lots of reasons: performance, resilience, reformatting, staging, filtering and protection (of and by whom?), etc.
    - Design the network to support this.
      - But what does this imply?
-

# Tussle argument

---

- I (the user) want to be able to connect to the servers and services of my choice.
    - Implies that my choice should *not* be based on physical topology.
  - I (the user) want to be able to establish a protected path (a VPN) to the point of my choosing.
    - Implies either universal packet carriage or that VPNS are an “application”.
      - Who can control it under these two models?
  - The future of E2E is defined by trust.
-

# DTNs

---

- For lots of reasons, should not assume that “source” and “destination” are always on the net.
    - Mobility, developing world, times of crisis..
    - Begs the question of what “source” and “destination” mean.
  - The idea of DTNs should be a fundamental part of architecture.
    - Management analysis.
  - How does the DTN model relate to application-level services?
    - Can applications switch from interactive to staged mode “seamlessly”?
-

# Next topic: addressing

---

- Yesterday: global addresses.
  - Today, NAT and address rewriting.
    - We see a hint of the problems conversion can cause to new applications.
  - Tomorrow:
    - Idea 1: Indirection
    - Idea 2: Capabilities
    - Idea 3: Overlays
-

# Patterns of communication

---

- Is two-party e2e communication the right paradigm?
    - What is happening at the service level?
      - Dissemination?
      - Diffusion?
    - What do addresses at the packet level have to do with this question?
      - Multicast.
      - Data-driven delivery.
    - Two contradictory ideas (?)
      - Pre-position my content near me. (Dissemination.)
      - Widespread mobility.
-

# Indirection

---

- A generalization of:
    - Multicast
    - Mobile IP
    - Anycast
  - And other things today done at a higher level.
    - Server selection.
  - And proposed as an aid to
    - Security and prevention of DoS attacks.
  - Where to start to evaluate this idea?
-



# Two ways to start

---

- Do a security analysis of indirection.
    - In general, if attacker can find your true address, seems they can still attack you.
      - Echoes of magic and “True Names”.
    - Capabilities try to sidestep this, but themselves seem to generate a complex security analysis.
  - Note that different uses of indirection may benefit from a different routing scheme.
    - Akamai makes their routing a differentiator.
      - Does this require the deployment of new routers, or can we use a common platform?
-

# Next topic: routing

---

- Today, routing and forwarding done by same hardware.
  - Emerging idea: compute routes more centrally, and download into forwarding engine.
    - Can there be competing route computation schemes (perhaps based on different address ranges?)
    - What are the forwarding primitives?
-

# So a possible idea

---

- Might call this “partial virtualization”.
  - One plane of forwarding engines
  - Multiple co-existing route computations.
  - Points where addresses get rewritten.
    - Very stateful. Can we do stateful anycast?
-

# Security of indirection

---

- Use anycast to diffuse an attack (or a flash crowd) across many points of entry.
    - Anycast so cannot gang up on specific indirection point.
  - But must control consequence of attacker forging a “converted” packet.
    - Does this necessarily imply encryption?
  - Only if forwarders are trusted can we assume that an attack will be deflected.
  - Routing itself must be secure and robust.
-

# Management

---

- Do multiple routing protocols imply multiple management of aggregates?
  - Increased need to integrate routing and route recovery with lower level tools for fault recovery.
    - Must bring this stuff inside a common management architecture.
-

# Economics

---

- What is the motivation/reward for deploying a forwarder?
  - How does the facilities provider make long-term provisioning decisions?
  - What is the structure of the “route computation” industry?
  - What is the basis to negotiate interconnection?
-

# How much should be built in?

---

- Today, the idea of “overlay” is to do something the “underlay” did not do.
    - But this is not fundamental.
    - What is?
  - What we “build in” is easier for applications to use
    - Easier to manage, easier to reason about.
    - Example, a common address format with different delivery modes “underneath”.
  - Having a baseline routing service is “helpful”.
-

# The future of routing

---

- The photonics folks predict a fiber core in which the connectivity can be re-arranged in a time-scale of seconds.
    - Today, routing, traffic engineering and connectivity occupy different time scales.
  - If they blur, then we have to rethink routing.
  - What would this mean if we have competing routing systems?
-



# User choice

---

- Should we let users pick routes?
  - Current motivation seems to be performance.
  - In future, access to enhanced services and other differentiators.
    - Economic implications:
      - Pro: driver of service innovation
      - Con: even more disconnect from routing and planning.
    - Management implications: many...
-

# Validating the connection

---

- How can the receiver decide if it wants to receive the connection?
  - Can it “outsource” the decision?
  - Idea: Instead of a “per-layer” open, devise a cross-layer, single packet session initiation request.
    - Design it to have minimal cost to the receiver
    - Design it so the state (if any) can be handed off.)
    - Use this to re-establish soft state in the network?
-

# Congestion and resource mgt

---

- Next time, design into the packet layer.
    - But: explicit, implicit, feedback/forward, etc.?
  - A techno/economics/mgt problem.
    - How interact with new routing?
  - Route diversity and other aspects of service assurance.
  - Relate to traffic engineering
  - What must be in packet to control access to QoS and enhanced network services?
-

# Identity vs. location.

---

- A well-known idea at this point.
    - I discussed location above.
  - But what is identity?
    - Distinguish between what the end nodes want and what is required to be visible in the network.
      - Control of DoS. But is it pushback, deterrence, or what?
      - Access to enhanced network services.
  - Do we know what the end-nodes really need?
-

# Higher level architecture

---

- Identity
    - Need many systems, so just leave “space” for it.
  - Location
    - Another technical/economic issue.
    - Many ways to capture and represent.
    - Security analysis?
  - Information authenticity
    - Not derived from where it came from.
-