

Mapping the future of the Internet

David Clark
MIT CSAIL
August, 2010

Questions to consider

Why is the future of the Internet an important research question?

- Isn't the present one good enough?

What are the forces shaping the future of the Internet?

- The research community is not the only player here.

What are some interesting approaches?

- How would we evaluate them?

Background—U.S. research

FIND (Future Internet Design) is a U.S. NSF program to look at what our global network of 15 years from now should be.

- Project now about 5 years old.

NSF just completed a set of major awards (Future Internet Architecture) to teams that have proposed to develop and trial integrated concepts for a future Internet.

- Very sorry—awards are not public yet...

3

Motivation for FIND

Challenges community to think about why we built what we built.

- A lot we got right (perhaps surprising...)
- A lot is almost an accident.

Challenges us to envision a future.

- Not just improve the present.
- Free our minds from the constraints of what is, to imagine what might be.

The Internet is a success

So why would we want to rethink its design?

- It's *not* the data plane.
- Packets have proven their generality, and we have polished the data forwarding function for years.

It is *not* that some broad class of application is unsupported.

- Application designers have shown the broad utility of the Internet.

The issues are centered in the broader context within which the Internet is positioned.

- Need to consider a broad range of requirements.

5

Issues to consider

Security
Availability and resilience
Better management
Economic viability and industry structure
Meet society's needs
Support for tomorrow's computing
Exploit tomorrow's networking
Support tomorrow's applications
Fit for purpose (it works...)

6

Security

Use as a first example of a requirement.

- Hard and important.

Why is the problem so hard?

- We don't agree on the definition of good security
 - A balance among stake-holders.
- We want different outcomes in different contexts.
- We cannot correct the insecurity of end-nodes.

Old ideas: (good ideas, but not why we thought.)

- Disclosure, integrity, availability
- How does this relate to firewalls, VPNs?
 - After the fact—not a part of the network

7

A different modularity

Attacks on communication

- Confidentiality and integrity addressed with encryption.
- Availability?? The central objective of networks.
- What else? (Traffic analysis.)

Attacks on the host

- Infiltration (can lead to most anything)
- So either prevent infiltration or limit its consequences.

Denial of service

- A special case of availability.

Information assurance.

- Sign the information, not the connection.

8

Examples of attacks

Byzantine packet handling.

- Re-routing, adding and dropping.
- Only end-node can detect, so end-node must be able to request re-routing.
 - Explicit
 - Implicit
- Multi-homed end-nodes

DNS corruption (pharming)

- No architectural support today to mitigate this.
 - Design is redundant, but not in face of malice.

9

Availability

First, as much as possible, make attacks on communication into failures of availability.

- Limit the range of attacks and responses.
 - Think: what is excluded...?
- Mechanism: wrap an end-to-end confirmation of identity around a connection.
 - Cleanly makes many attacks on/by the network into an availability problem.

Second, develop a theory of availability.

- At a high level:
 - All critical resources must be supported in a rich, heterogeneous, diverse form.
 - It must be possible to detect and distinguish (to some degree) failures.
 - The point of detection must be able to invoke different resources.
 - In general, only the end-points can detect failures.

10

Network management

Even less structured than security.

- No real consideration in original design.
- Mostly remote management of boxes.

Possible decomposition:

- Fault isolation and resolution.
- Network planning and configuration.

Does this framing actually decompose the problem?

- Do we know the modules of management?

11

New ideas:

Critical interfaces:

- Between regions to allow cross-domain capabilities.
 - This interface is fundamental. It reflects reality.
- Between layers to integrate application, network and technology.

Expression of end-user *intent*.

- Critical in solving availability problem.

Better tools for abstracting the manager's job.

- Critical in solving availability problem.

Default management automatic, just like dynamic host configuration.

Instrumenting the data plane to detect problems.

12

Interfaces define the industry

ISPs exist because of IP, and the protocols that connect regions together.

- There is no fundamental reason why ISPs look the way they do.

Protocols define the services that can be created across multiple regions.

So by creating protocols, we create opportunities for service (e.g. revenue) creation.

- Which are possible, which are dangerous?

13

Region interconnection

Old idea: BGP.

New ideas:

- Interconnection of advanced services
- Direct expression of business constraints
- Routing overlays
- Fault localization and correction
- Interconnection of traffic aggregates
- Short-term markets for service
- Security issues
 - Control of DDoS
 - Detection of corrupted or untrustworthy regions

14

QoS as a case study

QoS (e.g. DiffServ) is a technical success.

- Used today in lots of enterprise networks.

It was an economic failure.

- Not used in public Internet.

We did not design the interface protocols that hook ISPs together.

- Business issues (who pays whom, what must be revealed, etc.)

Observations

Management has a lot to do with security, availability and economics.

- These areas are not “modules”.
- Cannot have a “security” or a “management” design sub-group.

For all these areas, we have lots of great ideas, but must sharpen the architectural framework.

16

Information--moving up-layer

Old idea: an application issue (ignore it.)

New idea: need a framework

- Naming and identity of information.
 - Independent of how you get it.
 - But: think about privacy.
 - If you shout for information, the whole world hears.
- Dissemination
 - Swarms, P2P: (heterogeneous).
 - Should this be the basic service, or on top of a transport service?
 - Improves availability of information if it is pushed into the network.
- Economics: one service or many competing?

17

Issues to consider

Security
Availability and resilience
Better management
Economic viability and industry structure
Meet society's needs
Support for tomorrow's computing
Exploit tomorrow's networking
Support tomorrow's applications
Fit for purpose (it works...)

18

The role of identity

A requirement for identity comes up often:

- Detect misdirection attacks on communication.
- Detect invalid (unauthentic) pieces of information.
- Validate identity/authority of incoming connections to prevent infiltration attacks.
- Allow application/network to pick desired communication pattern, to insert the desired degree of checking into the path between communicating parties, depending on the degree of trust between the parties.

19

Designing identity schemes

There is more than one way we could approach identity.

- A private matter among end-nodes.
 - E.g. encrypted or meaningless except at end points.
- Signal of identity that is visible in the network.
 - Surveillance cameras in cyberspace.
 - Facilitate both policing (perhaps) and repression.
- Third-party credentials vs. continuity-based familiarity.
- Revocable anonymity.
 - Anonymity can only be revoked by its creators.
- Probably need all in different circumstances, so architecture should not constrain.

These are not choices to be made by technologists alone.

- Need a multi-disciplinary conversation.
 - I am very fearful of getting this wrong.

20

Application design

If the Internet is a platform (as it is), what is the shape of the systems/services/applications that run on that platform?

Applications run “on” the Internet. They are not the Internet.

Many approaches to construction

- Patterns of communication.
- Use of end node software and server software.

They are full of servers and services run by unaffiliated parties.

- Not just by ISPs—the Internet is an open platform for innovation.

The changing structure...

In the old days, there were two sorts of devices:

- Routers
- End-node computers.

Now:

- Server farms
- Cloud computing (latest buzzword...)

So where should computation be placed?

- And why?

Placement of computation

“The Internet” is not changed by where computation is placed.

- Except that we need some really high-capacity circuits...

But the user and the industry structure are strongly influenced.

Issues with application design

Ease of use

Ease of deployment

Performance

Robustness

Security

Economic (industry) structure

Who is in control?

Application design patterns

Should we care about the design of applications?

- We cannot *control* how applications are designed. So in the past, we left the design to others.

But application design influences those important factors I listed.

We should offer guidance as to proper design approaches.

- *Application design patterns.*

Forces that shape the future

Better security (or not).

Issues of government control.

Economics and investment.

- “The best way to predict the future is to invest in it.”

The shape of future computing.

The health of research and innovation.

The future of information.

The future of the user experience.

The trajectory of the developing world.

Security

We may be at an inflection point.

- Some users and usages may be retreating from the Internet due to security concerns.

Sudden (and very belated) attention to the problem.

- Will we actually be serious this time?
- Some of us have been trying to get attention since 1990.

I believe we require a serious, long-term assault on the problem.

A driver of change?

We *have* to address issues of security.

But the drive for better security may change the Internet.

- Perhaps in ways some of us do not want.
- Cyber-attack does not change the Internet, cyber-defense does.
- Better identity and accountability, more logging of activity, etc.

And who will do all this?

The concerns of government

The Internet was built by the private sector.

- Emerged initially (especially in the U.S.) without government intervention or regulation.

But governments have objectives, and will want control.

- Security.
- Control of objectionable content and behavior, etc.

The ISP as agent of the state

Deep packet inspection (e.g. UK proposal).

Logging of online behavior.

Regulation of infected machines.

- Consider theory of liability.

Regulation of content. (also private sector concern)

- Conflict with neutrality mandates.

Address many other security problems.

Jurisdictional boundaries.

Different countries will act differently. How should network designers think about this?

Economic viability

Fundamentals:

- Different parts of the network are built by different actors.
- Physical facilities (fibers, towers, etc.) require capital investment.
- Investors must be motivated to invest.

Our preferences:

- Facilities owners must not *control* the future of the network. Just invest in it.

31

What happens today?

How do facilities owners operate and interact?

- One answer is that they become ISPs.
 - Measure/model usage
 - Track customers and markets
 - Control routing.

ISPs serve a critical business function today.

- They don't just move packets, but manage capital and risk. Important economic role.

But ISPs were defined by the interface protocols.

- Is their current structure fundamental?

32

Don't run aground

Must steer the future between two unwelcome outcomes.

- Open access mandates leads to insufficient return on investment and deployment/upgrade stagnates.
 - Consider Australia, stimulus, OECD data.
- Return to more control (more vertical integration) leads to degradation of platform for innovation, which stagnates.
 - Consider legislation of neutrality.
 - A technical debate with larger consequences.
- Also avoid instability of the overall system.

Network neutrality: a case study

A technical problem:

- Congestion is real, so it has to be managed.
- Arises directly from the goal of cost-effective sharing.
- Any decision about allocation is policy.

An interconnection problem:

- Actions at the edge are motivated by interconnection.

A fight over the character of the platform.

- Defining the character of "open".
- Limiting the actions of service providers and the customers.

A hypothesis about the natural outcome under competition.

Many contending players

ISP, governments, criminals, application designers, content owners, users, etc.

- Some fight with money, some with laws.

What are they fighting over?

- Power and control
 - E.g. sovereignty, jurisdiction, civil stability, social norms.

Money

- For example, advertizing revenues (U.S.).
 - IDC predicts that in 2011 online advertizing spending will be \$45B (out of \$265B total).
 - By comparison, consumers pay about \$32B/year today for broadband access.

Control point analysis

When looking at current architecture, and alternative proposals for the future:

Find the points of control.

Ask who will control them and how the fight will play out.

- Strategically, decide if we care.

Ask what power arises from that control.

- Economic.
- Political.
- Cultural and social.

Control

Yesterday we fought over:

- Encryption

Today we fight over:

- The IP address space.
- The DNS (an obvious point of control).
- Access to and ownership of information.
- Jurisdictional authority.

Tomorrow we will fight over:

- Who owns my social network and knowledge of my location?
- Who owns my reputation? My online identity?
- The controls of future security mechanisms.

A big idea--virtual networks


In a virtual network, facilities (routers, links, etc.) are virtualized and then used by higher-level service providers to implement different networks, possibly using very different architectures.

- VPNs and overlays are a limited version of this idea today.
- A new form of competition.

38

A big idea because it "re-layers"

In a narrow sense:



Push the narrow waist "down"

In an architecture based on virtual networks, there is no single agreements on packets, addresses, formats etc.

All we agree on is how to get a "slice" of the virtual network and install software.

How should we analyze?

Look at requirements and drivers.

- My starting point: economic viability.

In a world of virtual networks, why would someone invest in expensive facilities?

- Owner does not control routing, so where should the links go?

Control point analysis: find the points of control.

- In an "unusual" place: the management interface.
- Will this be open or closed?

Now design the economic models.

Another new ideas: futures

If investment in facilities is a "up-front" or "sunk" cost, with a long period of depreciation and cost recovery;

And virtual networks anticipate flexible access to resources over a short term;

Then there must be some way to insulate facilities investors from risk so that they will invest.

Consider a futures market for bandwidth.

- Happens today with really expensive cables.

42

A new interface

Do we need to standardize the interface that defines this futures market?

- Has a lot in common with other commodity markets.

Not sure, but if we do, it is an odd sort of standard.

- Not moving packets, but money.

Not just bandwidth, but in a location.

- Compare to spectrum auctions.
- "Path futures"?

43

Another big idea

Information dissemination as basic service.

- Move the narrow waist "up".
- See, for example, Content Centric Networking, by Van Jacobson at PARC.

The basic (and very clever) idea:

- Don't send packets to a service, send a packet to find the information.
- The ACK is the information.

Now do control point analysis.

- Is this an overlay, or do the router owners control access to information?

Future computing

The Internet hooks computers together.

- (It hooks people and information together, but mediated by computers...)

More diversity is coming.

- The era of the PC is ending.
 - The Internet and the PC co-evolved.

Extremes:

- Million server "clouds".
- \$.10 sensors and controllers.

The future of information

Now: the Web.

The future:

- The Semantic Web
- Information about information (meta-data).
 - Context, location, rights,
- Information formatted on demand.
 - The personal experience, erosion of authoritative archives, erosion of search... (not just ads).
- Erosion of editorship: wiki vs. garbage
- Data from sensors.

The user experience

The movement of social context online.

- Massive revolution.
- Move from "communication" to much richer interaction.
- The new platform. (Who controls?)

Capture of physical location.

Linkage of real and cyberspace.

Massive multi-player games and living online.

- Making a living in a virtual world.

Can IT reduce energy consumption?

The future of the Internet?

What is important?

- The economic landscape and commercial viability.
- The tussle of interests and societal forces.

Technology plays second fiddle unless it can mitigate these issues.

- Not the way a techie thinks....
- The design is being debugged at run time.

Conflicting mega-trends

ICT (the Internet, cyberspace, etc.) seems to have empowered the individual and “organizations of the willing” at the expense of traditional institutions.

ICT has been the foundation for the construction of bureaucracy and control.

- The surveillance state, data mining, etc.

How do these reconcile/collide?

- Technology will influence the balance.

Feasible futures

Internet economics “succeeds”

- Open model persists.

Internet economics falters

- Downward spiral of consumer interest.

Investment in bandwidth falters.

- Stagnation at current speeds.

Wireless stagnates.

- Limited access and closed structure.

“Tussle lockup”.

- Conflict of interests halts progress.

For further examples of ideas

Go to the FIND web site:

www.nets-find.net.

Look for the link to the summary report of the “Future Internet Summit”.