

# Contemplating a new Internet: An upcoming NSF initiative

---

David D. Clark  
MIT CSAIL  
September 2005

# A 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.
-

# The NSF and its role

---

This talk is motivated by the upcoming initiative from the NSF.

- The changing face of research leadership.

This talk: Why?

Why needed? Why now? Why succeed?

---

# Isn't today's net good enough?

---

## Security and robustness.

- As available as the phone system
- Been trying for 15 years--try differently?

## Easier to manage.

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

## Recognize the importance of non-technical considerations

- Consider the economic landscape.
  - Consider the social context.
-

# What will be happening in 10 years

## New network technology.

- Wireless
  - Mobility
  - Dynamic capacity allocation
  - Dynamic impairments
- Advanced optics
  - Dynamic capacity allocation (again!)

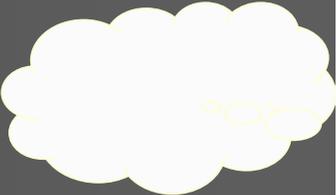
## New computing paradigms

- Embedded processor, sensors, everywhere

Whatever computing is, that is what the Internet should support.

- The Internet grew up in a stable “PC” time.

# The scope of the challenge



Is it “Internet classic”? A cloud of routers with general purpose computers at the edges?

No! The scope of the question is much bigger than that.

Ask: what will “the edge” look like. That is where the action is.

- Sensors. Embedded computers.

Ask: what is it that users do? Try to conceptualize a network that supports that.

- Information access and dissemination.
  - Location management and location-aware systems.
- 

# What should we reconsider?

---

For the moment, everything.

- Packets, datagrams, circuits--everything.
- Our religious beliefs
  - End to end, transparency
  - Our model for layering.

To conceive of a future, we have to let go of the present.

- This does *not* mean that we cannot get there incrementally.
-

# Timing

---

This is a long term effort.

- IPv6 started in 1990.

It is less important when we start, more important that we do so.

- We can and will do mid-course correction.
- Adjust the objective as we get closer.

Long term research has short-term fallout.

Short term research never achieves a long-term objective.

---

# Defining success

---

We throw away the current Internet.

- The most dramatic form of success.

We set a goal, and then we realize we can get there incrementally.

- Impose a bias or direction on change.

Lots of fresh ideas leak into the present Internet.

---

# The benefit

---

Today, we see erosion of clean design principles--architecture.

Clean architecture means clean interfaces, as well as better behavior.

Creates more opportunities for innovation.

- The NAT story.
-

# If we don't do this?

---

If we don't step up to conceive of what networking will be in 10 years:

- A narrowing of the utility of the Internet to specific purposes. E-commerce?
  - Limit our ability to exploit new technology.
  - A loss of funding (inside NSF) to sectors that seem more relevant and vigorous.
    - A gentle glide into irrelevance for research.
-