President's Information Technology Advisory Committee

Scalable Information Infrastructure

Les Vadasz, Chair
Steve Dorfman
Dave Dorman
Danny Hillis
Dave Nagel

February 17, 1999

Since We Met Last Time...

- Made November 1998 suggested updates
- Resulting draft was network-centric, did not discuss software
- Further efforts on scalability, software, and budget
- Added scalability-specific software recommendations
- Added funding guidance
- Result: Three findings, six recommendations

SII Findings

1. Our Nation's dependence on the Internet as the basis for its information infrastructure continues to grow at a dramatic rate.
2. The Internet is growing well beyond the intent of its original designers, to such an extent that we no longer understand it and cannot confidently continue to extend it.
3. Learning how to build and use large, complex, highly-reliable and secure systems requires research.

Finding #1: Our Nation's dependence on the Internet as the basis for its information infrastructure continues to grow at a dramatic rate.

- The Internet is increasingly becoming an integral part of our daily lives.
- The information infrastructure is becoming critical to our Nation's well being.
- Like our electricity system, we must be able to manage and control it.
- America leads the world in these technologies. It's crucial to continue investing to maintain our leadership position.

Finding #2: The Internet is growing well beyond the intent of its original designers, to such an extent that we no longer understand it and cannot confidently continue to extend it.

- The Internet is growing to a size that we have never seen before (from 2K devices in 1985 to >50M today).
- The Internet is growing in capabilities:
  - Anytime, anywhere communications
  - Electronic commerce
  - Streaming audio and video
  - Building blocks for new applications, e.g. digital signatures, shared environments
- Extending the Internet is challenging.
  - A network of this size and complexity hasn't been built.
  - We need reliability equal to or better than the phone system.

Finding #3: Learning how to build and use large, complex, highly-reliable, and secure systems requires research.

- Need to analyze this large, complex system to better understand its behavior
- Need to develop revolutionary new networking research ideas to keep up with Internet growth > Moore's law
- Scaling advances needed in multiple dimensions:
  - High speed, quality of service, ubiquitous access, infrastructure services, security, servers accessing exabyte data stores
- Research should encompass the complete end-to-end system

SII Recommendations (1)

1. Fund research in global-scale network behavior, including collecting and analyzing performance data analysis/collection, and as well as modeling and simulating network behavior
2. Support research on the physics of the network, including optical, wireless technologies including satellite, and bandwidth issues
3. Support research to anticipate and plan for scaling the Internet

SII Recommendations (2)

4. Support research on middleware that enables large scale systems
5. Support research on large scale applications and the scalable services they require
6. Fund a balanced set of testbeds that serve the needs of networking research as well as research in enabling information technologies and advanced applications

Recommendation #1: Fund research in the behavior of the global-scale network. This should include collecting and analyzing performance data as well as modeling and simulating network behavior.

- Mathematical models for telephony took over 30 years to develop, and they don't apply to the Internet
- Telephony an effective Internet deployment vehicle today, but differences limit future deployment ability
  - IP vs. switched networks, wide-range of applications vs. voice, wireline/wireless/cable/broadcast vs. twisted pair, wider performance range, broader statistical distribution characteristics
- Must develop new models and simulation methods
- Will require data sharing among network service providers

Recommendation #2: Support research on the physics of the network, including optical, wireless technologies including satellite, and bandwidth issues

- Improving transmission technologies profoundly impacts infrastructure capability
- Requires research in physical and signal processing: waveforms, signal quality, and manufacturing techniques
- Pressing issues: optical, wireless, and local-loop transmission

Recommendation #3: Support research to anticipate and plan for scaling the Internet

- #1 problem on the Internet today is scalability
  - # nodes, # users
  - Geographic dispersion
  - Heterogeneity of the networking environment and end nodes
  - New traffic types such as voice and streamed content
  - Availability, bandwidth, efficiency, reliability, security
  - New features, uses, and applications
- How can we design, build, test, manage, and evolve a network that's 10,000 times larger than today's Internet?

Recommendation #4: Support research on middleware that enables large scale systems

- Middleware: shared software that improves software development, helps the infrastructure run efficiently, makes large software systems possible
- Example: network aware applications
- Focus middleware research on:
  - Information management
    - Managing and presenting data, monitoring resources, integrating data from diverse sources, tracking information lineage, controlling access,...
  - Information services and survivability
    - Authentication and security, detecting intrusion, preventing human errors, long-term information preservation

Recommendation #5: Support research on large scale applications and the scalable services they require

- Very large applications have unique scalability needs that today's technology can't meet, e.g., scalable application development and quality control.
- Example applications: DOE nuclear stockpile stewardship, NASA EOSDIS Earth sensing database, FAA air traffic control system
- Large applications will be common in 20 years
- Potentials applications that can drive this research:
  - National Digital Library
  - Next Generation World Wide Web

Recommendation #6: Fund a balanced set of testbeds that serve the needs of networking research as well as research in enabling information technologies and advanced applications

- Distinguish between two types of testbeds
- 100X: primarily for "revolutionary" applications
  - Connects several Federal research networks: NSF's vBNS, DOE's ESnet, NASA's NREN, DoD's DREN.
  - Users expect high availability, performance
- 1000X: primarily for networking research
  - Stresses new technologies
  - More fragile, may experience outages
- Recommend periodic review of testbed activities
  - Ensure funding/usage consistent with original intent
  - Review participation by all appropriate parties
  - Costs for 100X testbed to be shared with applications communities

Recommended Funding Levels

- Current SII funding estimated @ $200M annually, about equally divided between network research and testbeds (exclusive of applications)
- Recommend linear increase from $200M to $500M between FY2000 and FY2004
- Emphasis on networking research (3/2 ratio)
- Relative project ratios estimated from a bottom up approach: 1 large (>$4M): 4 medium ($1-2M): 20 small ($100K-200K).
- Also fund two networking expedition centers for $5M/each

Funding Observations

- $300M = 19 large, 75 medium, and 375 small projects
- ~100^{+} upper-tier research universities
- ~50^{+} universities doing networking research
    - 375/50 = ~7 projects/university

➨ 75% of $ goes to medium/large projects

➨ # of small projects limited by # researchers