

WORKSHOP DESCRIPTION:
Reconnections: Managing Academic Networks With New Requirements
7-September-2005

Manageability Issues In Future Academic Networking

I. PROJECT SUMMARY

The original design principles for the Internet are not fully aligned with the requirements that are being shaped by the new uses, users, and technologies of the network. Some early clean-slate thinking efforts have started to consider a new set of design goals. At the same time, in the academic networking community, there is growing concern over the increasing unmanageability of its environment, and the need to realign fundamental architectural principles for campus and interrealm networks in order to effectively manage and utilize what is being deployed.

This document describes a workshop, and subsequent activities, that will convene leading practitioners of academic networking (complex, advanced campus and enterprise networks with multiple, technically diverse external connections), along with a few related interests (e.g. network researchers, vendors, corporate network architects) to focus on the architectural and management requirements for future R&E networking. By providing an academic enterprise perspective, this work is intended to inform and be informed by other efforts on beginning a redesign and redeployment of internets.

II. PROJECT DESCRIPTION

A. Background

A number of network professionals have observed that there is an increasing disconnect between traditional Internet design assumptions and evolving user and operator requirements. Some of these conclusions represent the articulate and deep understandings of the design and development community (e.g. see Ref1), identifying grand challenges that illustrate how fragile some of our most cherished architectural foundations may be. First efforts at "clean-slate" thinking has already been commissioned to begin a process of redesign of the general Internet.

At the same time, there is increasing doubt within the "commingled universe" of R&E networking about the long-term viability of our current approaches. This academic environment is characterized by complex and advanced campus networks connected to a growing diversity of external connections, including both commodity Internet and new services such as Abilene and dedicated (aka "personal") lambdas. The introduction of additional, potentially disruptive approaches from clean-slate and other initiatives could present even greater difficulties for academic enterprises. First thoughts are now emerging from the deployment community, highlighting the inadequacy of marketplace responses to ever-increasing system complexity (see Ref2), and the operational consequences of ad hoc and incremental approaches to security and regulatory requirements (see Ref3).

The enterprise R&E network is evolving its focus from technologies to include management of technologies. New metrics such as mean time to diagnose have entered the vocabulary, representing that the stress points have moved from meeting the initial demands of the research community to maintaining the resulting mix of deployed technologies. The introduction of packet disruption devices, the mesh of firewalls and exceptions, the use of ubiquitous VLANs and other network trends all have greatly escalated the complexity, and correspondingly decrease the manageability of campus networks.

While it may be claimed that the new designs and technologies being explored will run to isolated machines or subnets in the academic setting, history has shown that not to be the general case. Connectivity to special purpose networks is a useful tool and connections of special-purpose hosts to the campus network has many benefits. As a result, hosts on campus that might be connected to non-IP off-campus services can pose significant security and management challenges to the IP regions of the enterprise, or vice versa.

As traditional backbones are joined by multiple dedicated links, there may be value in viewing the new academic networking environment as being managed from the outside in, from the enterprise out into the external network clouds. It is at the enterprise that the increasing demands and complexity from the campus and external links need to be integrated and maintained. The enterprise defines the security environment of its systems; it is the context in which cost-benefit analysis is done; it will be the locus of management in the old and new world orders.

The goals of the workshop are not to develop new networking technologies, but to bring forward the management requirements created by the new technologies, and new policy requirements, particularly as they are placed into the existing broth at the enterprise. Conversely, new approaches must seek to improve manageability, and we intend to identify those requirements for such planners. We seek to both advise other clean-slate thinking and to be advised.

B. Workshop and Outputs

The focus of the workshop, and subsequent work that might emerge, is on the future of architecture and management of campus and national lab networks that connect to a broad variety of external networks. Manageability is achieved through art and science applied to a deployed infrastructure built on strong architectural principles. This workshop intends to identify the architectural principles that serve to meet the R&E needs while providing the basis for a manageable infrastructure. It will also identify much of the science, and some of the art, that must be established to convert that infrastructure into robust, flexible, and sustainable networks.

A small community within the Internet2 campuses has had a set of planning calls that have identified a key set of issues that are drivers for workshop discussion. These issues often are new requirements, some specific to the R&E community, that must be framed as architectural needs that are mapped against design for manageability principles. Others are approaches failing under scale. The issues include:

- Given the increasing need for the academic enterprise to have policy affect both traffic and routing, what underlying network architecture principles need to be modified?
- With the diversity of systems (desktops, hi-performance workstations, instrumentation, telephones, elevator and parking control, soda machines, etc) now attaching to the network, what architectural principles can guide crafting of networks with multiple levels of guarantee of service (reliability, SLA definition and monitoring)?
- How should design reflect the need to provision multiple levels of quality of service (bandwidth, latency, etc) across a distributed campus network, with the particular geo-location characteristics of academic research partnerships?
- With the decentralized decision making structure in academia, how should network architecture evolve to permit manageability but reduce friction between the enterprise and the departmental research community in areas such as security, address space, etc?
- How can the emergence of new technologies, such as trust fabrics, be leveraged in the design of networks?
- Given the unusual funding models within higher-ed, (e.g. one-time versus continuing research funds, disparity in local funding bases, uneven legacy conditions, etc), how should design principles reflect these conditions?
- How must manageability itself, and network management tools, including our approach to diagnostics, change to meet the new demands for service and complexity of the environment?
- Has the approach of complexity in the network rather than at the end station reached its limit in terms of network and desktop management, and what can improve the situation?

The planning group will continue to refine the list, and then present a set of structured discussion items to the attendees prior to the meeting, in order to prime the pump.

Attendees will be primarily be drawn from a select group of network architect practitioners from leading edge campuses and national labs. They will be chosen to reflect the long-range view of the workshop as well as the need to manage a complex mesh of enterprise and external networks. In addition, a small group of pertinent vendors (desktops, network equipment, etc) will be asked to send their chief architectural designers. Finally, we will augment this group with other clean-slate thinkers to maximize cross-pollination.

During the workshop we will progress through three successive assessments:

1. Identify parts of our current approaches are failing, or are soon to fail, in managing enterprise R&E networks and underlying principles that need to be modified
2. Develop a sense of the requirements for academic networking in the future, acknowledging drivers in collaboration and science, technology, policies (including compliance and audit), funding, trust fabrics, external network trends, and others

3. Translate the requirements and manageability gap analysis into a set of architectural principles that can guide future design and deployment of manageable networks to the academic community. Identify specific science that has to be pursued for manageability needs and processes to develop good practices.

Deliverables will include the following:

- Two reports from the workshop – Within a month of the workshop, a report and web site will be developed that contains polished minutes of the workshop. Two months later, after review and refinement by participants, a final report will be submitted to NSF and circulated as appropriate. That report will include recommendations on subsequent steps that NSF and other agencies could pursue towards the design of better academic networking environments.
- Strengthening of ties within the future network community – Through the planning and the workshop itself, we intend to continue discussions that have already shaped this effort. Beyond the manageability expertise in the R&E community, there are other community strengths, such as emerging trust fabrics, that can be shared with other clean-slate efforts.
- Early dissemination to key constituencies – One of the distinctive aspects of the workshop participants is the ability to elevate the issues to key campus constituencies, including the campus CIO's, virtual organizations, and other national and international groups. In that these constituencies will have to be a significant part of implementing new approaches, it is important to get their attention and involvement early. The participants in this workshop have numerous and effective ways to communicate awareness and outcomes.

C. Workshop Steering Committee

The workshop's Steering Committee will include representatives from Internet2's Security Advisory and Oversight Group (SALSA); and other representatives from academia and corporate entities. Members will include:

- Deke Kassabian, Senior Technology Director, Information Systems and Computing, University of Pennsylvania (Steering Committee Chair)
- Larry Acciaioli - IP Network Technology Supervisor, Ford Motor Company
- Robert Aiken - Director of Engineering, Academic Research & Technology Initiatives, Cisco Systems
- Robert Brammer - Sector VP and Chief Technology Officer, Northrop Grumman Information Technology
- Vinton Cerf - Senior VP of Technology Strategy, MCI
- Gary Dobbins - Director, Information Security, University of Notre Dame
- Terry Gray - Associate Vice President, IT Infrastructure, University of Washington
- Farnam Jahanian - Founder and Chairman of the Board, Arbor Networks
- Ken Klingenstein - Director, Middleware and Security, Internet2
- Chris Misra - Network Analyst, University of Massachusetts
- James Pepin - Chief Technology Officer, Information Sciences Division, and Director of High Performance Computing and Communications, University of Southern California
- Mark Poepping - Director for Architecture, Discovery & Integration for Computing Services, Carnegie Mellon University
- Jack Suess - Chief Information Officer, University of Maryland, Baltimore County

D. References

1. Clark et al, report on the January 2005 meeting of the End-to-End Research Group, which is part of the Internet Research Task Force.
<http://www.ir.bbn.com/~craig/e2e-vision.pdf>
2. Gray, T. "'Finger Pointing Tools' for Isolating Distributed System Performance Problems"
<http://staff.washington.edu/gray/papers/fpt.html>
3. Gray, T. "Security in the Post-Internet Era: the Indeterminate Internet"
<http://staff.washington.edu/gray/talks/2004/netsec-09mar04.ppt>
4. Internet2 Security
<http://security.internet2.edu/>