Summary of the Workshop On Research Challenges for 2015 Global Networks Security Breakout Session

Interagency Working Group on

Cyber Security and Information Assurance (CSIA IWG)

7AM Pacific, December 10th, 2008

Joe St Sauver, Internet2 Security Programs Manager
joe@oregon.uoregon.edu or joe@internet2.edu
http://www.uoregon.edu/~joe/nitrd/december10th.pdf

Disclaimer: The opinions expressed are not necessarily those of any other entity

1. Introduction

What We'll Be Talking About Today

- Tomas Vagoun asked me to talk with you a little today about the invitational Networking Research Challenges Workshop which was held at the Edgewater in Seattle from September 28th-30th, 2008.
- I was also asked to "provide more context on the Internet2 initiative," and to discuss "what class/types of security issues are being investigated by the Internet2 security group, and what classes are not, and why"
- And finally, I was asked to address "where do we go from here? What are some opportunities for the Internet2 Security Group and the CSIA IWG interactions and linkages? How can Internet2 and CSIA support each other?"
- I only have twenty minutes to talk plus ten minutes for discussion), so if you want a more detailed report about the security breakout section, please see www.uoregon.edu/~joe/nitrd/november20th3.pdf

2. The Security Breakout Session

Participants Self-Selected A Breakout Section

- Each invited workshop participant self-selected a breakout section
- Not surprisingly, since most of the invited participants were network researchers (and not network security types), most of them selected areas other than the network security breakout section.
- We were a small group of just eight folks, and not all participants were able to be present for the entire time reserved for the breakout sessions either due to needing to participate in multiple breakout sessions, or due to Rosh Hashanah occurring during the time of the workshop.
- A larger group, comprised primarily of network security-focused researchers, seasoned with operational network security folks and technical participants from the commercial network security community (e.g., the sort folks who tend to gather at things like the annual San Francisco RSA Conference) would probably result in a broader/different set of perspectives if a follow-on workshop is subsequently convened.

Participants in the Security Breakout Section

- Security breakout participants were:
  - Matt Crawford, FNAL
  - Phil Dykstra, DREN
  - Chris Greer, NCO
  - Karl Levitt, NSF
  - Paul Love, NCO
  - Grant Miller, NCO
  - Thomas Ndousse, DOE
  - Joe St Sauver, Internet2 and U. Oregon
- Because the breakout session took place over multiple days and represents the opinions and work of many people, no opinion mentioned in this summary should be attributed as being the opinion of the facilitator or any particular participant unless they choose to express agreement with it.

The Charge to the Security Breakout Section

Participants in the security breakout section were asked by the conference organizers to think about eight questions. They were:

1. Visions for network security across multi-domain, multi-layer heterogeneous networks and what it will enable in 5-15 years. What applications will be enabled, based on advances in the capabilities of this breakout area.

2. Visions for a new trust model that will allow extending secure communications across federated, virtualized, multi-domain networks.

3. What basic research is needed in network security and/or trust models to enable end-to-end secure dynamic, seamless, transparent, heterogeneous network environments including foundational theory for risk modeling and analysis, vulnerabilities trends network protocols and services, cyber security simulations and testbeds?

4. How do we provide end-to-end security in virtualized networks, heterogeneous networks, dynamic optical networks, embedded networks, federated networks, sensor nets, hybrid packet/switched networks, etc?

5. How do we accomplish coordinated network security in a distributed autonomous network environment?

6. What are the research challenges of distributed intrusion protection/detection, performance measurement, management and incident response in a secure dynamic heterogeneous networking environment?

7. What are the security vulnerabilities of the emerging control plane and signaling technologies for dynamically switched optical networks?

8. Is there a need for a network security test bed?

That's a LOT of material to cover! We don’t claim to have exhaustively addressed any of the questions, this is just a first pass.

Time Horizon and Scope

- The workshop's time horizon, 2015, was both “very far in the future” and “almost upon us,” particularly w.r.t. security.
- Unlike some other research areas, security is prone to being very operationally focused (and reactive) due to the urgency of fighting today’s security “fires,” and that sometimes makes folks reluctant to think strategically/over a longer term time horizon.
- On the other hand, as we know from things like trying to deploy DNSSEC, developing and deploying new security technologies can easily take a decade or more. If we were to identify a new security technology today, it might easily by 2018 or 2019 (not just 2015!) before it was in production deployment.
- We also recognize that security issues can span both unclassified and classified networks, but our breakout group intentionally limited our consideration to unclassified topics only since this meeting included foreign nationals and others without government security clearances.
- Let’s dive in and look at some of the breakout group’s topics.

Topic #1: Vision of Network Security in 2015

"Visions for network security across multi-domain, multi-layer heterogeneous networks and what it will enable in 5-15 years. What applications will be enabled, based on advances in the capabilities of this breakout area."

- **Put another way, what do we (think) we know about the network and computing environment of 2015?**
- The network will be "way too fast"
- Everything will likely be encrypted
- The network will be truly multilayer: it won't be just a layer three world any more
- Security will enable applications ***largely in so far as it doesn't "get in the way" or interfere with applications working.***

Topic #1: Vision of Network Security in 2015 (2)

- Other factors impacting future network security developments
  -- Huge installed legacy/production base means new security technology introduction and diffusion may essentially follow equipment replacement lifetimes --> S L O W rollout…
  -- Costs and benefits are often asymmetric (my expenditure on network security may help <u>your</u> security, but paradoxically may not necessarily do much for my own security)
  -- We need the commercial sector to build the gear we need, but commercial differentiation favors new features and increased complexity over simplicity, performance and economy.
  -- Deployed complexity (example: firewalls) currently exceeds the administrative ability of amateurs; the supply of trained network engineers and security people remains insufficient
  -- Compliance related activities (paperwork) may drain additional resources away from actually fighting the cyber "wars"
  -- One size doesn't and cannot fit all; flexibility is important
  -- We will continue to overlook obvious solutions

Topic #2: Trust Models

"Visions for a new trust model that will allow extending secure communications across federated, virtualized, multi-domain networks."

Findings:

- There are basically two traditional trust models:
  - Hierarchical trust models, rooted at a trusted origin, such as PKI and other certificate-based models, and
  - Less structured ad hoc "web-of-trust" models, as used by PGP/Gnu Privacy Guard, where the trustworthiness of a credential is a function of attestation by multiple trusted peers
- Trust can sometimes be tightly coupled to notions of identity and reputation, although those are not ubiquitously present in all cases. For example, a trusted party's ultimate "real life" identity may not always be known.

Topic #2: Trust Models (2)

Findings (continued)

- Federated trust models, such as those based on Shibboleth & InCommon or Kerberos also are seeing active development and widespread deployment in some communities.

- There are many practical problems which remain unsolved: revocation lists are still problematic, for example, and the ad hoc nature of PGP/Gnu Privacy Guard's can deter adoption in some business application.

- **Secure communication is already possible across federated, virtualized, multi-domain networks.**

Topic #2: Trust Models (3)

- **Recommendations:**

- What is urgently needed is further exploration is work on making existing trust models more practically **usable.** (For example, what proportion of your current mail stream is digitally signed with either PGP/Gnu Privacy Guard or S/MIME? If the signed fraction is low, why?)*

- The linkages between concepts of trust, identity (or anonymity) and reputation also require additional research.

---

* For one consideration of that question, see "Why Johnny Can't Encrypt," http://gaudior.net/alma/johnny.pdf and "Johnny 2" at http://www.simson.net/clips/academic/2005.SOUPS.johnny2.pdf

Topic #3: End-to-End Security

What basic research is needed in network security and/or trust models to enable end-to-end secure dynamic, seamless, transparent, heterogeneous network environments including foundational theory for risk modeling and analysis, vulnerabilities trends, network protocols and services, cyber security simulations and testbeds?

- We believe that end-to-end secure, dynamic, seamless, transparent and heterogeneous network environments are already possible today via applications such as SSH.

- The real gaps may lay up or down the network stack.

Security Vulnerabilities: Up and Down the OSI Stack

- **Findings:**

  Just to review, the OSI stack model has seven layers. They are:
  - Layer 7: Application Layer
  - Layer 6: Presentation Layer
  - Layer 5: Session Layer
  - Layer 4: Transport Layer
  - Layer 3: Network Layer
  - Layer 2: Data Link Layer
  - Layer 1: Physical Layer
- By default, when thinking about network security, there is something of a tendency to focus on issues at Layer 3.
- However, in reality, we need to look both up and down the stack to address the security risks we face today.
- Let's begin by looking down the stack.

Down the OSI Stack

- It is a fundamental rule that higher layers cannot be secured without the lower layers also being secured, yet in recent years there has been limited attention to insecurities at the physical layer or data link layer, despite changes in network operational practice that include things like nation-wide layer two networks, and national and regional optical networks.
- Currently known/familiar threats at lower levels of the OSI stack include ARP spoofing MITM (man-in-the-middle) attacks at layer two, and physical layer attacks such as passive optical taps or the interception of wireless network signals by attackers. While these attacks are well known, little research is currently focused on detecting and addressing those concerns in scalable ways. That needs to be corrected.

Down the OSI Stack (2)

- Less familiar attacks which may be relevant to the lower levels of the OSI stack (such as the physical layer) over the next five to fifteen years include:

  - intentional attempts at **kinetic (physical) destruction of key national network infrastructure** by terrorists or hostile nation state actors

  - intentional attempts at **electromagnetic destruction of network assets via high power microwave weapons, or high altitude electromagnetic pulse effects**, a threat which was recently publicly reaffirmed by the 2008 blue ribbon Congressional Commission to Assess the Threat to the United States from Electromagnetic Pulse (EMP) Attack (see www.empcommission.org)

Up the OSI Stack

- Simultaneously, at the same time there is a need to look "down the stack" and insure that all higher layers are built upon a sound foundation, we note that there is also increased miscreant interest "up the OSI stack," particularly at the application layer.
- As noted by SANS Institute in their Top 20 Security Risks report, nearly half of the 4,396 total vulnerabilities reported in SANS @RISK data from November 2006 to October 2007 relate to web application vulnerabilities such as SQL Injection attacks, cross-site scripting, cross-site request forgeries, and PHP remote file inclusions (see www.sans.org/top20/#s1). While these are application layer vulnerabilities, they critically need our attention.
- This change of emphasis reflects miscreant efforts to obtain sensitive financial information such as credit card numbers or other personally identifiable information; in the government and commercial sector, an information-centric focus is presumed in counterintelligence and the protection of proprietary competitive information.

Topic #4: End to End Security in Diverse Network Environments

"*How do we provide end-to-end security in virtualized networks, heterogeneous networks, dynamic optical networks, embedded networks, federated networks, sensor nets, hybrid packet/switched networks, etc?"*

Findings:

- Network security shouldn't (and can't!) require knowledge of, or assumptions about underlying transport technologies. A user might be using ethernet, wireless, optical lambdas, packet over sonet, ATM, FDDI, etc., and in fact, in some cases, they might be using several of those technologies on a single connection.
- If the network is a passive transport media (rather than an active participant in the security process), security is a **host** or **application layer problem** rather than a network layer problem

Closer Coordination Between Security and Networking Folks, Security and Apps Folks, amd Security and System Administrators

- Currently these are largely silo'd communities; those silos need to be attacked and broken down so that communication and cooperation can occur.

- Similarly, there needs to be closer coordination between federal IT security entities and:
  - Higher education operational security folks
  - Higher education security researchers
  - Commercial system and networking security entities
  - Civil and criminal cyber law enforcement agencies

- **Recommendation:** develop mechanisms and opportunities to foster sharing and interaction.

We Need to Move The Perimeter Into The Host

- Chokepoints can't keep up at 10Gbps, 100Gbps is here, 1000GB will be here during the duration of the timeline for this workshop
- To scale border protection, we need to move the perimeter "two inches into the host" -- put network security policy onto a trusted network interface card/chip.
- Deals with the issue of the firewalls not being able to keep up at increasingly high rates

Recommendations:

- The NIC would be site-configurable, not host-configurable, and auditable. It would report events as required by configured security policy.
- Verifying host & OS integrity is probably out of scope.

Security Implications of Circuit-Oriented Architectures

- **Findings:**
- Circuit-oriented architectures may be an exception to the comment that "transport doesn't matter"
- Circuit-oriented point-to-point wide-area optical architectures are a major focus of the government and academic advanced networking computing community, particularly for high bandwidth science applications.
- Ironically, however, security concerns may have limited the deployment of these facilities, with the kernel of those concerns typically relating to circuit oriented architectures bypassing traditional perimeter security appliances such as firewalls or intrusion detection systems.
- These concerns generalize beyond optical networks to a variety of other point-to-point environments, including tunnels and VPNs.
- Further work is needed in this area.

Topic #5: Network security meets secure network traffic

"How do we accomplish coordinated network security in a distributed autonomous network environment?"

- **Findings:**
- We find it likely that traffic in future networks will be encrypted end-to-end.
- Traffic monitoring and filtering may have no more inputs than source and destination addresses, plus traffic history.
- Traffic analysis will become an important part of network-based security systems.
- Even when traffic is sent in the clear, as is the norm for open science data, the sheer volume of data flows guarantees that pattern-based detection will misfire often, again shifting the burden to traffic analysis.

Topic #5: Network security meets secure network traffic (2)

- **Recommendations:**

- Network-based intrusion detection and prevention systems must incorporate content-blind rules or heuristics. The nature of these methods is an area for study. The inputs to such rules can include source and destination addresses, security association ID, times of observation, and possibly some key negotiation traffic.

Topic #6: Challenges of Distributed Security

"What are the research challenges of distributed intrusion protection/detection, performance measurement, management and incident response in a secure dynamic heterogeneous networkings environment?"

- **Findings**
- Security attacks are increasingly distributed, therefore their detection and defense often requires a distributed solution.
- Optical circuit switched paths may cross several administrative domains, adding to the complexity of solutions.
- Traffic flows often take asymmetric paths, making monitoring and control from a single location impossible.
- No current intrusion prevention systems work in the face of distributed asymmetric flows.
- Coordination between incident response groups tends to flow up and down a pyramid with little lateral interaction.

Topic #6: Challenges of Distributed Security (2)

- **Recommendations**

- Optical switched paths provide an opportunity to perform authentication prior to establishing connections.

- Research in distributed intrusion prevention systems.

- Methods should be defined for more direct sharing of performance and incident detection data across domains.

Topic #7: Control Plane Security

- "What are the security vulnerabilities of the emerging control plane and signaling technologies for dynamically switched optical networks?"

- **Findings:**

- Systems at the endpoints of dynamically switched optical paths may make assumptions about the origin of traffic arriving on those paths. Compromise of the control plane - or accidental flaws in its design or operation - can invalidate those assumptions, with effects that cannot be predicted.

- Control plane traffic is commonly carried in-band. Even when it is isolated, the possibility of it appearing in-band by error may exist.

Topic #7: Control Plane Security (2)

- **Recommendations:**

- Elements of the control plane are end systems with respect to control plane functionality. End-to-end security mechanisms for the control plane should be developed, possibly in parallel with methods for the isolation of control plane traffic. These security mechanisms must be particularly robust against partial network failures and against active attacks through the physical media.

Topic #8: Is there a need for a network security test bed?

- **Findings:**
- In the immediate term, a test bed is needed to test and deploy capabilities and to see how the community of users and network engineers respond to them. The potential value of test beds in verifying the usability of security designs should not be overlooked.
- These concerns generalize beyond optical networks to a variety of other point-to-point environments, including tunnels and VPNs.

Topic #8: Is there a need for a network security test bed? (2)

- **Recommendations:**

- A network security test bed should be built on the GENI infrastructure.

- Attack traffic datasets would be a useful component of a testbed environment.

- Additionally, applications that run on GENI (and other experimental test bed networks) should include security metrics and a discussion of security considerations.

- All new network architectures must include a security model.

Some Potential Next Steps

- Solicit, analyze and synthesize network security R&D roadmaps and plans which have been generated at the agency level
- Determine whom those agencies have consulted for expert level advice on network security research and development directions, and insure that people specializing in areas we weren't equipped to consider have a chance to provide input, particularly in areas such as sensor networks, wireless and embedded system networks
- Confirm we're asking the right research and development questions, questions which will improve the security of our networks and systems, while preserving the performance and usability of those environments

3. If We Have Time, A Quick Overview of Internet2 And A Few Words About Internet2 Security-Related Activities

For more information about Internet2, please see the full 43 slide October 2008 "Internet2 Overview" presentation at http://www.internet2.edu/pubs/Internet2-Overview.ppt

The Internet2 Network Is Bigger Than You Think

- **Internet2 is NOT just a high bandwidth national network linking 200 or so leading American R&E universities**
  -- we also connect over 40,000 K12 schools; over 4,200 public libraries; over 1,000 four year colleges and universities; nearly 700 community colleges and nearly 200 museums, zoos, aquariums, and science centers via Internet2's Sponsored Educational Group Participant (SEGP) program
  -- we also peer with federal mission networks and international research and education networks
  -- Most recently we've begun working on health-related networking as a result of the FCC's Rural Health Care pilot program (see http://www.internet2.edu/health/ for more information), so expect to also see Internet2 also connecting hospitals and clinics, etc.
  -- Internet2 also has a commercial peering service with over 60,000 commercial routes available to members who want to leverage their Internet2 connectivity to help meet their commercial Internet connectivity needs

Advanced Networking

- While Internet2 provides production-quality transport for bulk unicast IPv4 traffic, it also supports advanced protocols, including:

  -- IPv6
  -- IP multicast
  -- jumbo frames (9K MTU)

- Most recently, the Internet2 community has also been hard at work on the Dynamic Circuit Network (DCN).

  Please see the illustration on the following slide showing one DCN usage scenario (this example involves offloading a large science data set transfer from the University of Nebraska/Fermilab shared IP infrastructure onto a DCN connection “on the fly”)

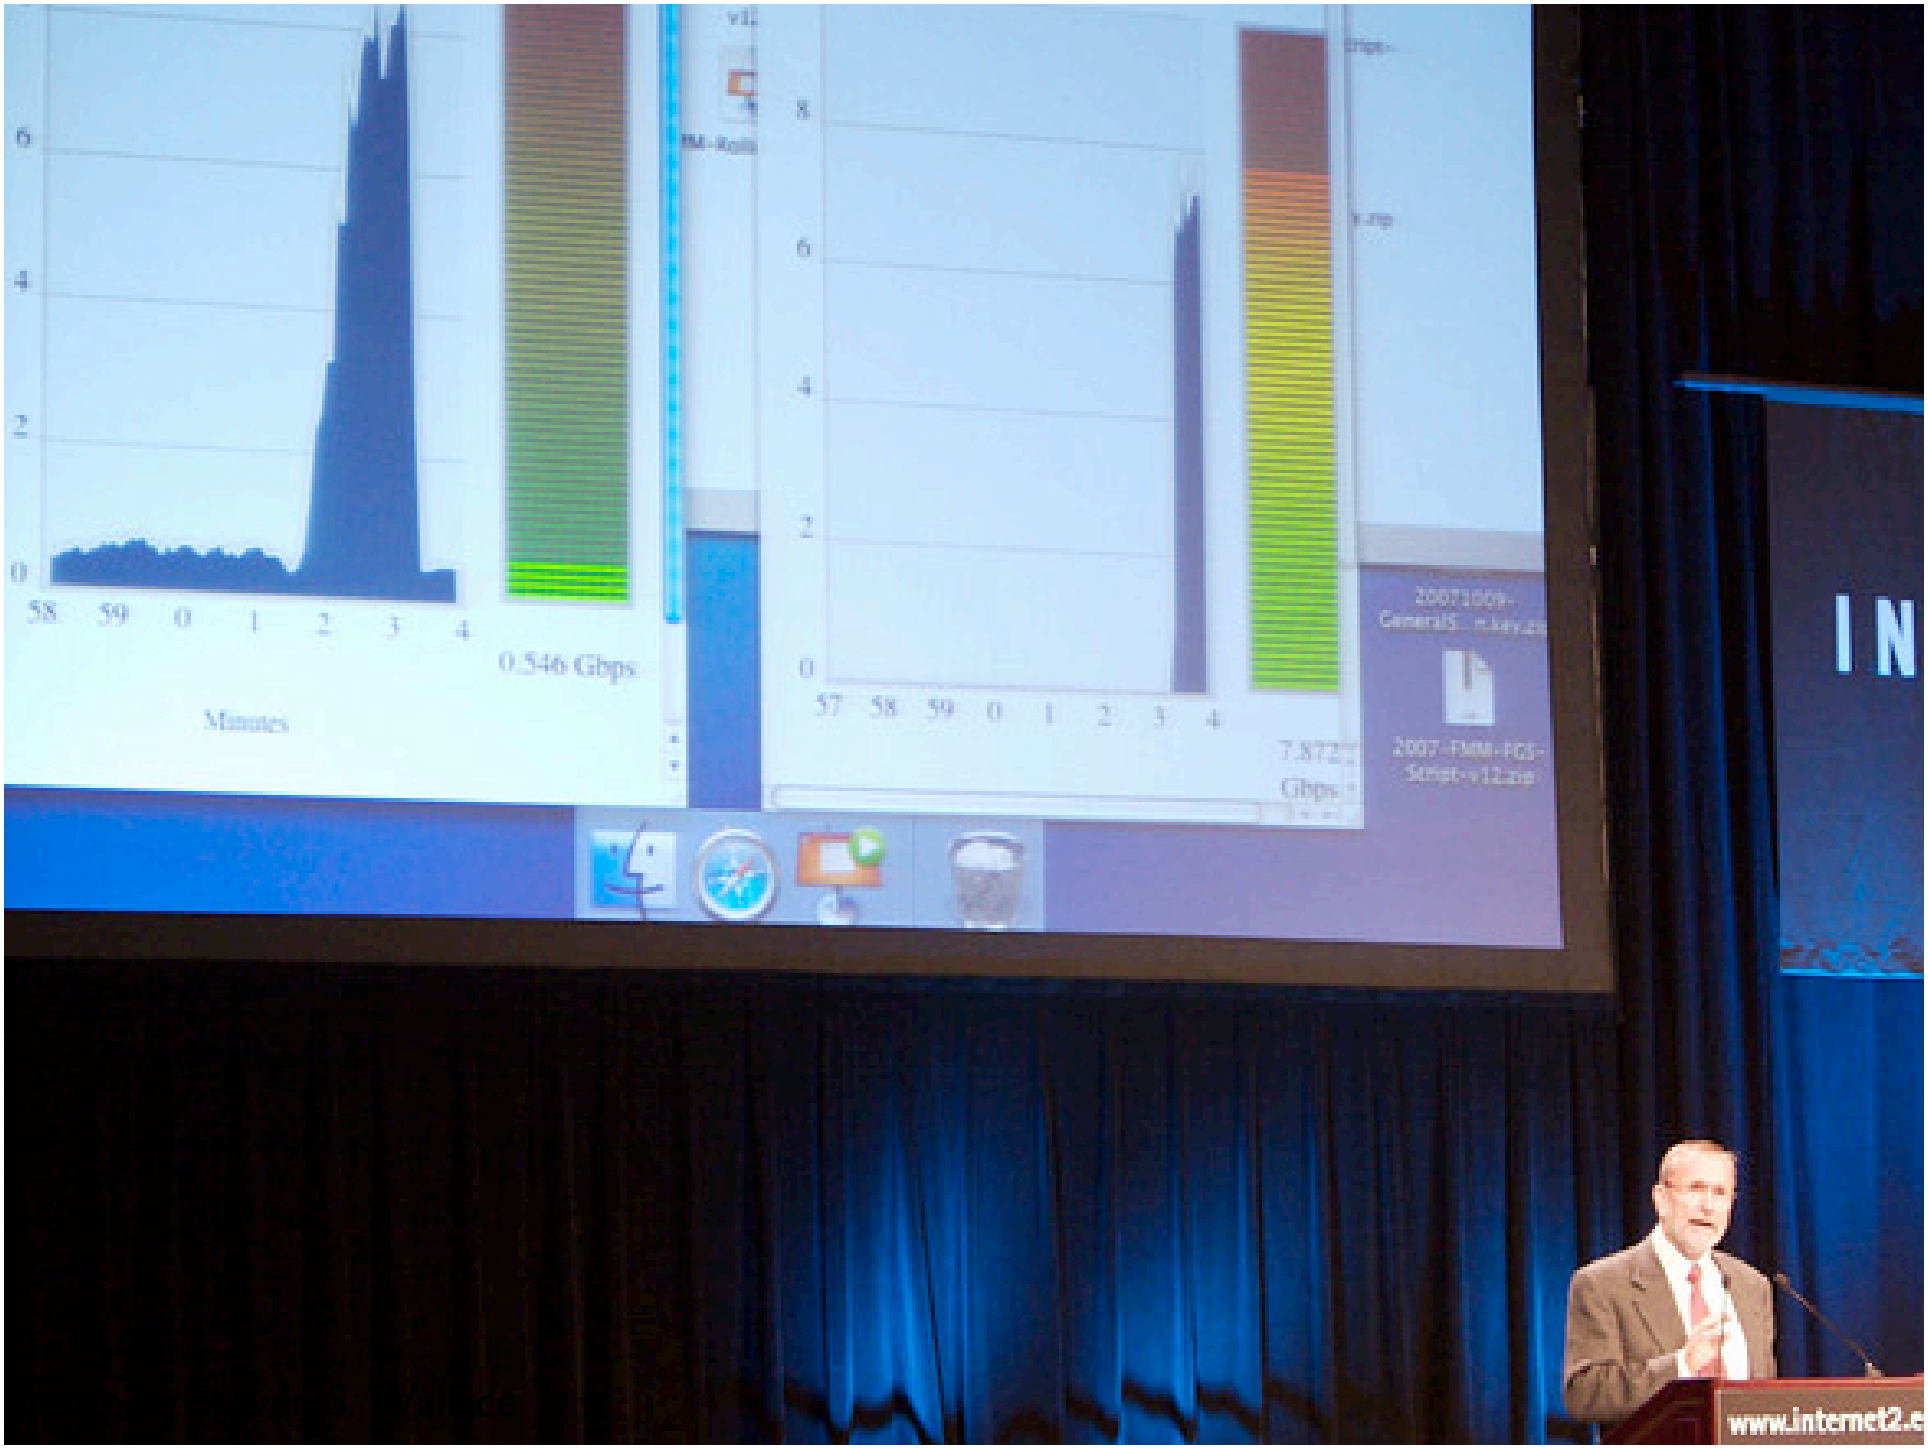
6
4
2
0
58 59 0 1 2 3 4
Minutes
0.546 Gbps
8
6
4
2
0
57 58 59 0 1 2 3 4
Gbps
2007-FMM-RDS-Smat-v12.zip
IN
www.internet2.e

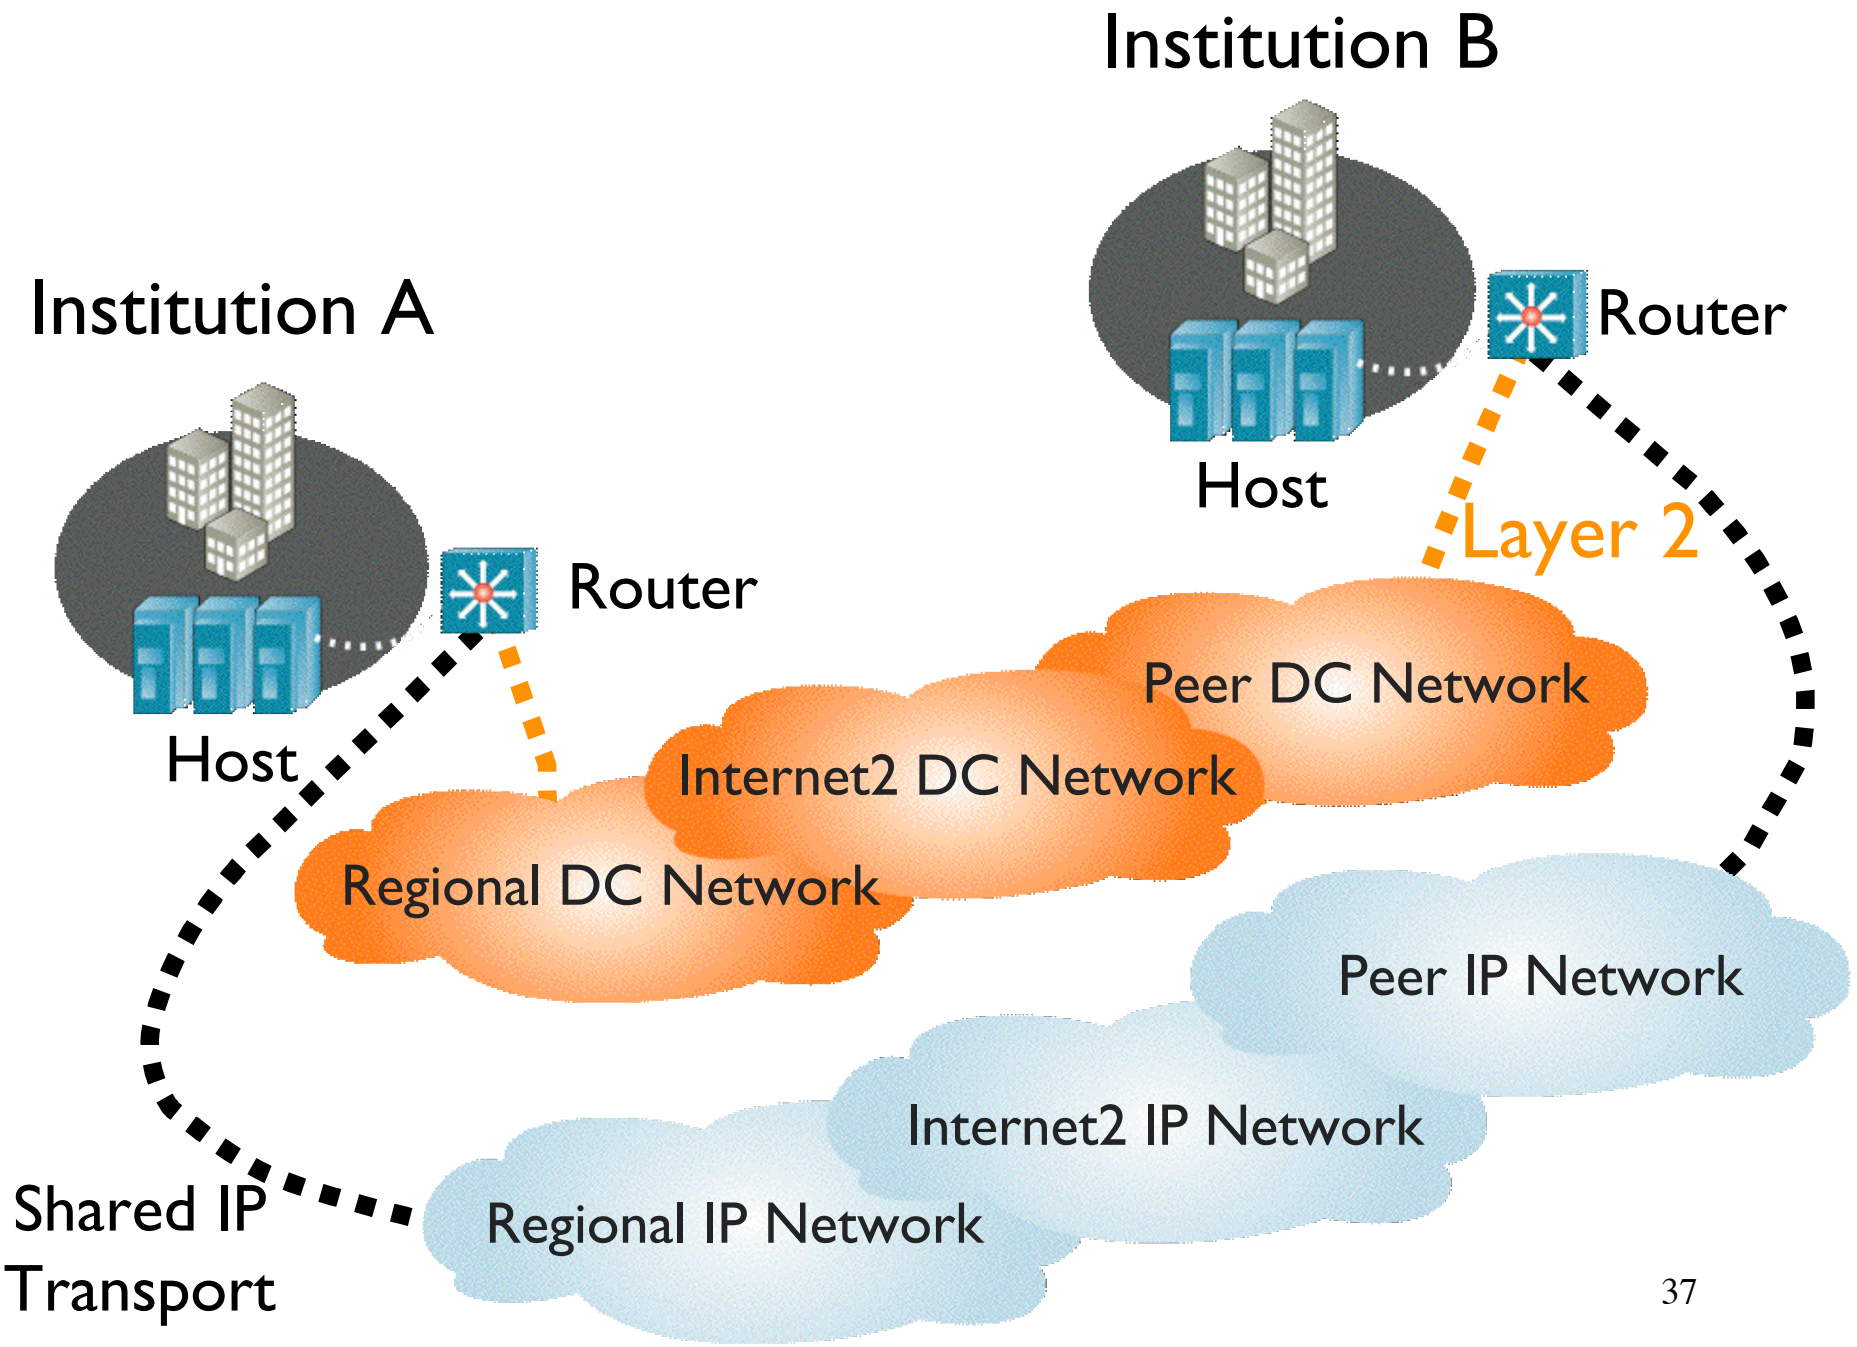
Institution B
Institution A
Router
Host
Layer 2
Router
Host
Peer DC Network
Internet2 DC Network
Regional DC Network
Peer IP Network
Internet2 IP Network
Regional IP Network
Shared IP
Transport

Internet2 Network

www.internet2.edu/network

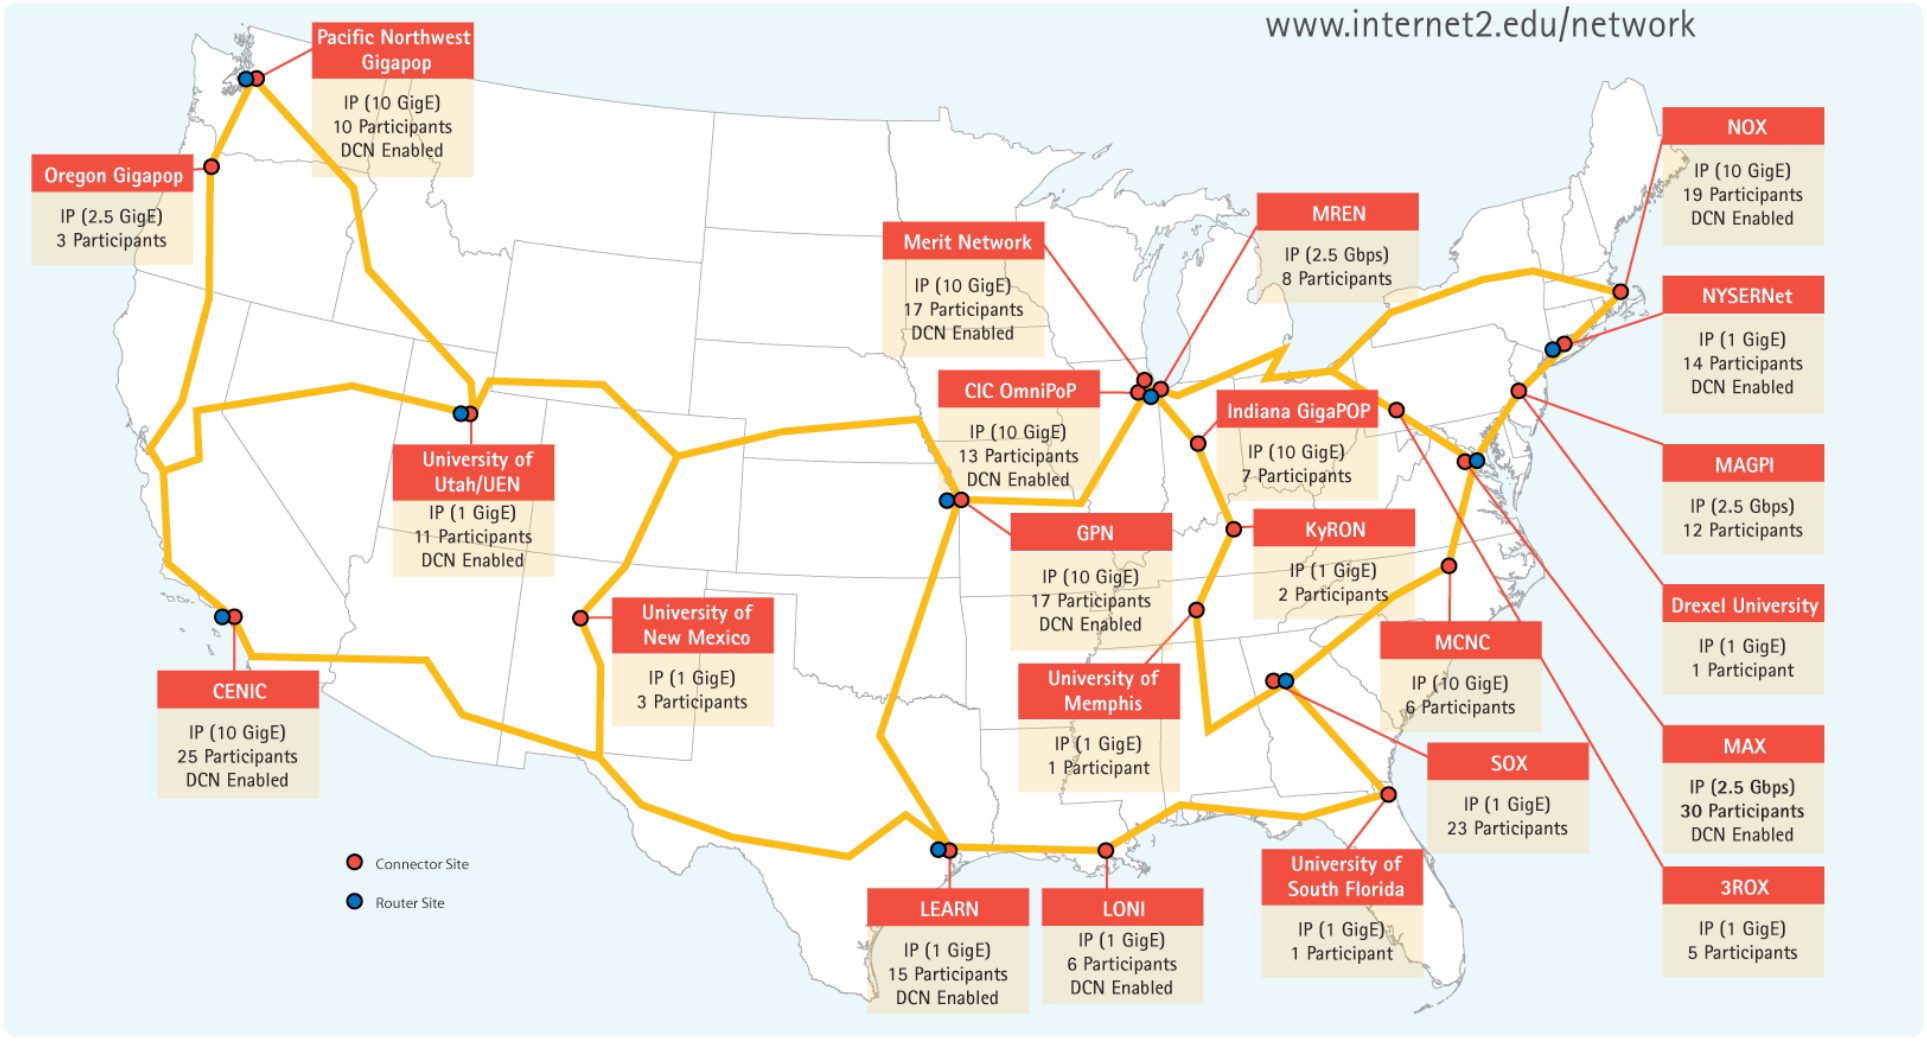

ciena | INDIANA UNIVERSITY | infinera | Juniper NETWORKS | Level(3) COMMUNICATIONS

Internet2 Is Also About More Than Just A Physical Network…

- It includes software development efforts, such as the Internet2 Middleware work on federated identity management which has resulted in Shibboleth and related work, work that has seen practical application in government as well as higher education. An example of this is the Department of Justice's Global Federated Identity and Privilege Management (GFIPM) Initiative, see www.it.ojp.gov/default.aspx?area=nationalInitiatives&page=1179
- Internet2's work also includes security. Like middleware, security is handled via a collaborative and community-driven model that does not rely on a large central staff. Also, rather than attempting to compete in an adversarial way with our colleagues, we prefer to work collegially with them. One example of that sort of collaboration is the Internet2/Educause Security Task Force (see http://www.educause.edu/security/16030), and another fine example is the Research and Education Network Information Sharing and Analysis Center (REN-ISAC) hosted by Indiana University (see http://www.ren-isac.net/)

Top 10 Issues for IT in Higher Education, 2008

- Educause annually surveys higher ed technology leaders to identify the top ten issues they face. This year's report at http://net.educause.edu/ir/library/pdf/ERM0831.pdf says that

- ***The number ONE issue for 2008 was "Security"***

- **The number five issue was "Identity and Access Management," and the number six issue was "Disaster Recovery/Business Continuity."**

- Clearly a variety of IT security-related issues remain very salient and a top priority for higher education IT leadership.

The Tricky Bits

- Everyone's already really, really, really busy, and there are a tremendous number of potentially relevant IT security issues
- Attacking some issues is distinctly non-trivial and may involve significant pain (paid in cash or karma)
- We've got no direct authority to compel sites to do (or not do) things: we need to persuade or advise, not direct or command
- Meetings may (or may not) have the appropriate folks – security issues of concern may be **policy** level issues which need to be addressed by CIOs, **technical network** issues appropriate to network architects or senior network engineers, **technical system/server issues**, or **end user issues**, so you need to assume that the meeting attendees may only be conduits to the right people "back at the ranch."
- Many security issues go FAR beyond just higher ed
- You also have to avoid accidentally educating the bad guys.

General Criteria for Security Areas Meriting Priority Security Attention

- **Areas affecting or particularly relevant to backbone network operations, or campus systems and networks.**
- **Mass scale phenomena** involving millions of users (or more): spam, worms, bots/zombies, malware, etc.
- **High impact phenomena** which can really hurt: distributed denial of service attacks; attacks which employ cyber events to affect tangible facilities, such as SCADA systems which control pipelines, factories and other facilities; etc.
- **Highly publicized phenomena –** if the media broadly covers an area (such as system breaches involving personally identifiable information), it is hard for that area **not** to become a priority area.
- **Emerging threats** which **aren't** being adequately covered.

Some Topics We've Worked On

-- Botnets
-- Capacity Planning and System and Network Security
-- Cybercrime, Cyberwar, Cyber Terrorism and Cyber Espionage
-- Cyberinfrastructure Architectures, Security and Usability
-- DNS Security and DNSSEC
-- Domain Names and Registrars
-- Fast Flux
-- Lawful Intercept and CALEA
-- Loss of Network Control Incidents/Insider Threat Management
-- Malware
-- National Scale Disaster Planning (EMP and Pandemic Flu)
-- Network Traffic Analysis
-- Phishing
-- Real Time Emergency Notification and the Clery Act
-- Route Injection
-- Security and Performance
-- Security and Privacy
-- Spam

There Are Lots More Areas Still To Cover

- As long as the list on the preceding page is, there are still many more security areas still to cover. A short list might include:
  -- Anonymity, International Censorship Circumvention, and Tor/Onion Routing
  -- Campus Physical Plant Control Systems and Their Security
  -- Data Storage and Data Backup in a <$200/TB World
  -- IPv6 and the Security Implications of Deploying It On Campus
  -- Passwords and Authentication
  -- Physical Security of Critical Internet Infrastructure
  -- P2P Applications, DMCA and Managing Network Traffic
  -- S*BGP and Securing the Routing Infrastructure
  -- Security of Optical Network Elements
  -- Tunnels
  -- VoIP Security
  -- Wireless Security (Including Sensor Nets and Mobile Devices)
- We invite you to work with us on these, or on topics from the preceding slide, if you're interested in any of these areas!

What Security Areas *Aren't* We Doing?

- Production traffic monitoring and community incident reporting is being well handled by the **REN-ISAC**, so we don't have much interest in duplicating/competing with their efforts
- We also generally try to avoid focusing on security-related areas which **Educause** may already be addressing. Thus, for example, Internet2's not actively working on security awareness and training since there's already an Educause group working on that topic.
- We do our utmost not to compete with Internet2 member institutions for security **research funding**, although we're happy to partner with those institutions on projects if invited to do so.
- We've had limited involvement at the protocol development level (e.g., **IETF** work), however that may be changing.
- Not being lawyers, we tend to stay clear of IT legal/policy topics
- Not being EE's, we don't tend to work on chip leve security issues
- Since we don't have government clearances and don't work in SCIFs, we don't investigate **classified** security topics

Opportunities to Work Together and To Help Each Other On Security Issues

- We welcome involvement and participation on security-related topics from our Federal partners, particularly at Internet2 Member Meetings and Joint Tech meetings, and on working groups.
- We've been delighted to participate in a variety of federal agency IT security research roadmap workshops, and we'd welcome the opportunity to participate in future federal IT security activities.
- We applaud efforts such as your Cyber Leap Year program, and note that that's an area that we've personally encouraged the higher ed cyber security community to contribute to…
- We'd like to better understand the unclassified system and network security challenges you face, and how we can work together to address them so feel free to send me email at joe@internet2.edu or joe@oregon.uoregon.edu, or give me a call at 541-346-1720
- Thanks for the chance to talk today! Are there any questions?