Report From the Workshop On Research Challenges for 2015 Global Networks Security Breakout Session

Joe St Sauver, Internet2 Security Programs Manager
joe@oregon.uoregon.edu or joe@internet2.edu

November 20th, 2008

http://www.uoregon.edu/~joe/nitrd/november20th.pdf

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

1. Introduction

Networking Research Challenges Workshop

- This invitation-only workshop was held at the Edgewater Hotel in Seattle from September 28th-30th, 2008.

- It was sponsored by the National Science Foundation and the Department of Energy Office of Science, and was organized in cooperation with the US Federal Large Scale Networking Coordination Group and its Joint Engineering Team.

- We'll talk a little about that workshop and the results thereof today, and we'll try to leave a few minutes for questions at the end of this presentation.

Workshop Purpose and Goals

- "The **Federal Plan for Advanced Networking Research and Development** of March 2008 [for the final version, see http://www.nitrd.gov/pubs/ITFAN-FINAL.pdf] presents a **vision for advanced networking based on a design and architecture for security and reliability** that provides for heterogeneous, anytime-anywhere networking with capabilities such as federation of networks across domains and widely differing technologies; dynamic mobile networking with autonomous management; effective quality of service (QoS) management; support for large-scale data transport and sensornets; near-real-time autonomous discovery, configuration, and management of resources; and **end-to-end security tailored to the application and user.**

  "Goals for this research include:
  -- Next-Generation heterogeneous networking (convergence of optical, wireless, packets-switched, and dynamic circuit-switched network)

Workshop Purpose and Goals (2)

"-- Network security **<===**

-- Federated heterogeneous networking and internetworking: resource reservation, security, management, performance monitoring, fault diagnosis, etc.

-- Networking challenges: 100Gb+ data transport, end-to-end optical transport, cross-layer communications in dynamically reconfigurable optical networks, optical packet switching to eliminate router bottlenecks, sensor nets, secure dynamic mobile networking

"The Large Scale Networking Coordination Group proposes this workshop on Networking Research Challenges to identify the future challenges of networking associated with these four goals."

Format of the Workshop

- The meeting included:
  -- A variety of plenary presentations
  -- Several panels
  -- Four parallel breakout groups
  -- Interim reports from the breakout groups
  -- Dinner (with presentations)
  -- Discussions during breaks, etc.
- The four parallel breakout groups covered:
  -- Federated Optical Networking
  -- Heterogeneous Networking
  -- Network Science and Engineering
  -- Networking Security
- But before we talk about the work of the security breakout group, let's briefly touch on the security plenary presentation.

2. Security Plenary Presentation

Karl Levitt's Security Plenary Presentation

- Each breakout area had an introductor plenary talk.
- Karl Levitt (NSF/CISE and UC Davis) delivered the introductory plenary presentation for the security area, entitled "Network Security: Rethinking The Network To Support Security, Mobility, Management, Experimental Evaluation" (35 slides).
- A copy of that presentation is linked from the Security Breakout group page at http://www.uoregon.edu/~joe/nitrd/karl.ppt
- That talk was very thought provoking and somewhat controversial (exactly as a plenary should be!)
- Just to provide three quick examples, completely out of context:
  - What are the consequences of a simple routing core?
  - How should we be thinking about the botnet problem?
  - Is there a science of security?
- If I may, three slides stolen directly from Karl's talk…

Karl's "Consequences of a Simple Routing Core"

Benefits

- Universal connectivity
- Data forwarding permits packets to be sent from anywhere to anywhere
- Routers perform a very simple function and can be realized at any scale: central office to consumer devices
- Internet is *open*: supports creation of many applications and link technology
- Many faults are handled easily by the core

Problems

- Little support for management
- Diagnosis can be a nightmare
- Bad guys can launch attacks across Internet to any vulnerable node
- Impossible to trace attackers to their source
- Quality of service (especially RT) not easily achieved

Karl's "Thinking About the Botnet Problem "

Botnets will continue to be an issue

- Any vulnerable host can become a bot
- There will always be vulnerable hosts

The source of a Botnet will be difficult to determine

- Without accountability it is impossible to identify the commander of a Botnet

So, it is essential to stop or delay the growth or damage associated with Botnets; *only the network can do this*

- An ISP or an enterprise router can detect Bot-like traffic
- And, perhaps block or delay such traffic

But, there are consequences to blocking

- Blocking consumes precious human and device resources
- False positives will lead to many calls to a help desk

Karl's "Is There a Science of Security?"

- Are there *impossibility* results?
- Are there powerful *models* (like Shannon's binary symmetric channel) so that realistic security and privacy properties can be computed? Possibilies include:
  - Control Theory for security
  - Kirchoff-like laws to capture normal behavior for routers
- Is there a theory that enables:
  - Secure systems to be *composed* from insecure components, or even
  - Secure systems to be composed from secure components
- *Metrics:* Is there a theory such that systems can be ordered (or even partially ordered) with respect to their security or privacy?
- Can entire systems (hosts, networks) and their "defenses" be *formally verified* with respect to realistic security objectives and threats?
- Are there security-related hypotheses that can be validated *experimentally*?
- What kind of an instrument (*testbed*) is needed to validate such hypotheses?

And There's Lots More

- We don't have time to review Karl's entire presentation today, but truly, it is well worth a look.

- Having hopefully whetted your appetite for Karl's talk, let's go on to the security breakout section, the portion of the workshop which solicited active input from participants.

3. The Security Breakout Sessions

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 RSA Conference or at NANOG meetings) 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.

This Workshop vs. The Just Issued Report

- Because the Federal Plan for Advanced Networking Research and Development had just been issued in final form, there was some question from participants about how the workshop should incorporate that work in their own deliberations.

- We were asked to NOT spend our time rehashing or critiquing or elaborating on that report or its findings during our breakout sessions, but to focus on the breakout charge/discussion questions we'd received.

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?

Some participants also shared thoughts about other security-related topics, building on the specific charge items mentioned.

That's a Lot of Material To Cover

- Those eight topics represented a lot of potential ground to cover, although in some cases we felt as if several questions asked more or less the same question, albeit in slightly different ways.
- We don't claim to have exhaustively addressed any of the questions, this is just a first pass.
- If you were a participant and I grotesquely screwed up something you were trying to say, please let me know and I'll fix it.
- We welcome your contribution, or the contribution of those you may work with (feel free to send comments to me).
- Please don't wait too long, I need to finish up my own written report for inclusion in the workshop's final report.
- The material we'll go over today doesn't represent all the work that the group covered, consider it just some selected material. To see additional items, go to http://www.uoregon.edu/~joe/nitrd/

Were These The "Right" Questions?

- There was also some concern among some participants that these may not have been the key questions to focus on from a network security research and development planning point of view.

- If we don't ask the right or key questions, we may find ourselves answering the questions which were asked, **but** ultimately not thinking about all the things which may be critical.

- Thus, it may be helpful to step back a bit, and to ask some meta questions first.

Some Meta Questions to Ponder For the Future

- What exactly are the cyber threats we're worried about? Intrusions? Eavesdropping on network traffic? Distributed denial of service attacks? Malware? Physical attacks on community critical systems? Insider threats? All of the above and more?
- Do we have a reference or "model" security environment in mind? For example, securing an intentionally simple and transparent research network that supports multiple advanced protocols is a lot different than securing a unicast-only IPv4-only production network that is heavily encrusted with firewalls and active network security middle boxes.
- If there's a tension between security and usability, or security and performance, how are, and how <u>should</u>, those conflicts be resolved?
- Security involves the network, but it also involves systems, and applications, and users…but the focus of this workshop was primarily on the network. Does it make sense for us to just look at ONLY the network?
- Multiple participating agencies have been working on their own agency network security R&D roadmaps. How will they fit into this?

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 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.

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)

Down the OSI Stack (3)

- Addressing those known and other reasonably anticipated threats will require a substantial program of additional research, including:

  - **identification of key Internet assets** (such as transoceanic cable landing points, major network traffic exchange points, locations where multiple long haul networks are channeled into common corridors due to a lack of alternatives, etc.),
  - the development of methods to **harden** or otherwise improve the survivability of those high value assets in a physically hostile environment
  - **systematic testing to quantify the vulnerability of commercial network equipment to electromagnetic pulse effects** (it is difficult to believe, but unfortunately the recent EMP Commission largely failed to evaluate the vulnerability of packet switched equipment to EMP, focusing instead on conventional telephone switching equipment in the telecommunications chapter of their report)
  - approaches to countering potential electromagnetic pulse risks, whether through improved shielding at time of manufacturing, through post hoc shielding, through the use of purely optical (rather than OEO) interconnects, or other methods

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.

Up the OSI Stack (2)

- Arguably, proper application of encryption to data in transit and data at rest, along with improved application development practices to eliminate things like SQL injection issues, should largely mitigate these risks, and yet we know that is not the case.

- Phishing, a social engineering attack on confidential data, continues to be a problem, for example. Because system integrity can be undercut by users volunteering their passwords, we need additional research into human factors so we can better understand how to keep human participants in complex security systems from serving as the "weakest link."

- Similarly, SSH and SSL/TLS encryption along with two factor authentication (the use of both something you know, such as a password, and something you have, such as a hardware cryptographic token), should largely make technical credential capture attempts a futile exercise, yet we know that end-to-end strong encryption and two factor authentication is still the exception rather than the rule, nominally because of economic and ease-of-use issues.

Up the OSI Stack (3)

- We urgently need research work into how we can eliminate continued reliance on simple passwords transmitted in plain text, an outdated and incredibly insecure foundation security technology that is still rife across the Internet.

- We don't know how to deploy two factor authentication at scale. Users don't want to tote a bandoleer of hardware tokens with them wherever they go, with perhaps one token for access to routers and other network devices, another for access to servers, and still others for commercial sector tasks such as personal bank access and stock brokers. Federated approaches based on Shibboleth have great potential in this area, but deployment/adoption has been slow to-date.

Up the OSI Stack (4)

- Or consider messaging security: while PGP/GnuPrivacyGuard has the potential to substantially improve the privacy and integrity of a ubiquitous application (email), we know that the uptake of that technology has been virtually non-existent beyond a small number of technical elites. We need to understand how to overcome those obstacles.
- We also know that spam is now rampant. In fact, spam now constitutes 90% of all email, and it would not be an exaggeration to say that within five to fifteen years, 99%, 99.9% or even a greater percentage of all email may be spam unless effective measures take place.
- When we get to the point where only one message in a thousand or one message in ten thousand is "real" (non-spam) will email continue to be viable as a foundation collaboration technology supporting scientific research?

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.

Security Implications of Circuit-Oriented Architectures (2)

- This is somewhat counterintuitive: if one system, or a small subnet of systems, connects via a switched optical network connection to another small subnet, thereby forming a small closed collaborative enclave, that would appear to provide a reduced attack surface, most notably limits on the potential population of attackers who may have access to those interconnected resources.

- The scenario that concerns some, however, is one which uses the circuit-based architecture to bridge sensitive networks to public networks. Imagine a scenario with two sites interconnected by a point-to-point optical network:

  Internet ==> host at site one ==> optical network
  ==> host at site two ==> sensitive internal network

Security Implications of Circuit-Oriented Architectures (3)

- The optical network element in that diagram might explicitly avoid institutional firewalls.
- That scenario would thus potentially enable synchronous or phased undesirable access: That is, end-to-end malicious access would not necessarily need to occur. Content from the Internet could be introduced at one time, and only subsequently obtain access to sensitive internal networks. (e.g., the path does not need to exist end-to-end in order for contamination to occur)
- We believe that cross-contamination can be prevented through use of a partitioned "red/black" network architecture, much as secure government networks are currently air-gapped from the Internet, but that strict partitioning comes with substantial real and intangible costs. That approach, applied to an unclassified environment, needs to be carefully studied.
- These concerns generalize beyond optical networks to a variety of other point-to-point environments, including tunnels and VPNs.

Security Implications of Circuit-Oriented Architectures (4)

- **Recommendation:**

- Fundamental research is needed to develop strategies to address these concerns as they are essential to enabling broad deployment of circuit based networking solutions.

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.

- Currently available ones are small in scale, have limits on the acceptable range of what can be tested or are classified and therefore unavailable for non-classified research.

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.

What We Didn't Talk About

- There were some areas where folks either didn't have strong opinions or lacked a background/basis for comments, including:

  -- Networks for embedded systems
  -- Sensor networks
  -- Wireless networks

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
- .

Thanks for the Chance to Talk

- Are there any questions?