Cyberinfrastructure Architectures, Security and Advanced Applications

Internet2 Member Meeting, Arlington VA
3:00-4:00 PM, April 22nd, 2008

Joe St Sauver, Ph.D.
joe@uoregon.edu or joe@internet2.edu
Manager, Internet2 Security Programs
Internet2 and the University of Oregon
http://www.uoregon.edu/~joe/architectures/

Disclaimer: All opinions expressed in this talk are strictly my own, and do not necessarily represent the opinions of any other entity. This talk is provided in a detailed written form to insure accessibility, and for ease of web indexing.

I. Introduction

The Abstract for Today's Talk

When talking with users about cyberinfrastructure and advanced applications, security is a topic which often comes up -- but not for the right reasons.

More often than should be the case, some security practices and some security-oriented network architectures hinder rather than help users to do their work. What can be done to avoid this?

How can we have both secure cyberinfrastructure and an application-friendly online environment at the same time?

My Inspiration for This Talk

- I was inspired to do this talk because researchers at some Internet2 connected sites were running into various local network “security measures” which were keeping them from getting their work done. I was specifically motivated by discussions I had with a number of researchers during the lunchtime cyberinfrastructure meeting Jill Arnold held during the San Diego Internet2 Member Meeting.

- Also during this time, I became concerned about two unrelated issues:

  -- architectural steps the federal government was undertaking in conjunction with its new TIC ("Trusted Internet Connection") program, and
  -- a lack of attention to some old attacks which might prove devastating at some point in the future.

SALSA and NetGuru

- Some of the topics I'll be covering today have previously been the subject of extensive discussions by participants in **Salsa** (an acronym which formerly stood for "Security At Line Speed"), Internet2's security advisory group (see http://security.internet2.edu/salsa/), and **NetGuru**, another Internet2 security activity (see for example http://security.internet2.edu/netguru/docs/internet2-salsa-netguru-200702.html) .I'd be remiss if I did not acknowledge the valuable insights I've received from listening to and participating in discussions with those groups.

- I, however, am solely responsible for the content of this talk including any errors expressed or implied by it. If you follow the recommendations in this talk and have unfavorable results, that is not the fault of anyone involved with SALSA or NetGuru.

Is the Security of Cyberinfrastructure Appropriately A Part of Internet2's Work?

- Yes, and for multiple reasons. Just to mention a few of them…
- The Internet2 network is a very high capacity network, and that capacity, while provisioned and intended solely for legitimate uses, could be a potent weapon if it were to be abused to attack other sites. We want to make sure that doesn't happen.
- We also need to make sure that security measures do not keep users from doing the very work Internet2 was meant to enable.
- Securing cyberinfrastructure requires us to go beyond thinking just about the network backbone. We need to think "end-to-end" rather than just "network node-to-network node." That means we need to care about the security of regional optical networks, the security of campus networks, and even the security of individual end hosts. This end-to-end focus is not new to Internet2; as evidence of this consider the earlier Internet2 Campus Expectations Task Force.

II. Cyberinfrastructure? What Is That?

There Are Two Types of "Cyberinfrastructure"

- "Cyberinfrastructure" is a term often used in two rather different contexts.

  For one group of cyber security practitioners, "cyberinfrastructure" is a term dating approximately from PDD (Presidential Decision Directive) 63, *Critical Infrastructure Protection*, signed by President Bill Clinton on May 22, 1998:

  > "Critical infrastructures are those physical and cyber-based systems essential to the minimum operations of the economy and government. […]"

  http://www.fas.org/irp/offdocs/pdd/pdd-63.htm

The Other Sort of Cyberinfrastructure

- For the rest of us, "cyberinfrastructure" is a term more commonly associated with the NSF (National Science Foundation) and its Office of Cyberinfrastructure, and with documents such as *Revolutionizing Science and Engineering Through Cyberinfrastructure,* aka the "Atkins's Report."

- That report declared five years ago that "[…] cyberinfrastructure refers to infrastructure based upon distributed computer, information and communication technology. If infrastructure is required for an industrial economy, then we could say that cyberinfrastructure is required for a knowledge economy."

  http://www.nsf.gov/od/oci/reports/atkins.pdf at page 5

An Example of "Our Kind" of Scientific Cyberinfrastructure

- A classic example of high end cyberinfrastructure would be the U.S. High End Computing (HEC) infrastructure, as described in the NITRD (National Coordination Office for Networking and Information Technology Research and Development) Supplement to the President's FY 2009 Budget (see www.nitrd.gov/pubs/2009supplement/NITRD-09Supp_FINAL-hec-ia.pdf or the screen shot of part of that document which is shown on the next slide)

President's 2009 Request

Strategic Priorities Underlying This Request

Ongoing investment in Federal HEC facilities and advanced applications supports Federal agencies' science, engineering, and national security missions and helps sustain U.S. scientific leadership. Priorities include:

Leadership-class systems: Continue acquisition of highest-capability systems for cutting-edge scientific research and national security applications

Production-quality HEC resources: Invest in capacity platforms to expand Federal computing resources for critical agency needs and for the science and engineering communities

Advanced applications: Develop data- and compute-intensive applications for current and new HEC platforms

Highlights of Request

Acquisition of prototype leadership-class and production R&D systems

NSF: Continue multiyear acquisitions of the Track 1 petascale system and other midrange Track 2 systems to capitalize on the growing importance of cyberinfrastructure for advanced scientific discovery and education; Track-2 system (504 TF) at the Texas Advanced Computing Center becomes operational

OSD (HPCMP): Upgrade HEC platforms at multiple supercomputing centers

NIH: Selected acquisition of cluster and midrange compute-intensive systems

DOE/SC: Upgrade LCF system at ORNL to 1 PF (early FY 2009); expand ANL's LCF resources by upgrading BlueGene/P to 250-500 TF (late FY 2008); NERSC 104-TF XT4 in full production and integrated into a common high-performance file system

NASA: Continue annual investments in supercomputing systems to track Moore's Law (4X capacity every 3 years) and meet NASA's rapidly growing requirements for large-scale numerical modeling and simulation

DOE/NNSA: Acquire new production system to replace ASC Purple; continue operation of RoadRunner base system; deploy and operate TLCC07 capacity clusters; initiate operation of Sequoia Initial Delivery (ID) system

Applications

NSF: Multidisciplinary Cyber-enabled Discovery & Innovation (CDI) program, including petascale applications that focus on understanding complexity in natural, built, and social systems and increasingly data-intensive applications; software for applications that need to integrate computation and data acquisition in heterogeneous, dynamic computing environments

OSD (HPCMP): CREATE program to develop highly scalable application codes (aircraft, ships, antennae)

But Cyberinfrastructure Isn't Just About High End Computing

- Cyberinfrastructure also encompasses large scale data storage; visualization; middleware, operating system and application software; collaboration tools; and yes, even networks. :-)

- Examples of relevant cyberinfrastructure-related networks include the Internet2 Network, federal mission networks and other national scale R&E networks at home and abroad, and experimental next generation network test beds.

Global Environment for Network Innovations

- GENI is…

  […] d*esigned to allow experiments on a wide variety of problems in communications, networking, distributed systems, cyber-security, and networked services and applications. The emphasis is on enabling researchers to experiment with radical network designs in a way that is far more realistic than they can today. Researchers will be able to build their own new versions of the "net" or to study the "net" in ways that are not possible today. Compatibility, with the Internet is NOT required. The purpose of GENI is to give researchers the opportunity to experiment unfettered by assumptions or requirements and to support those experiments at a large scale with real user populations.*

See http://www.geni.net/faq.html

The Most Compelling Reason to Redesign the Net (Per The GENI Research Plan)

- The GENI Research Plan (www.geni.net/GDD/GDD-06-28.pdf) has a section on pdf page 18 describing important requirements and opportunities associated with any design for a Future Internet. So **what's the first/"most compelling reason"** to redesign the net?

2.1.1 Security and Robustness

Perhaps the most compelling reason to redesign the Internet is to get a network with greatly improved security and robustness. The Internet of today has no overarching approach to dealing with security—it has lots of mechanisms but no "security architecture"—no set of rules for how these mechanisms should be combined to achieve overall good security. Security on the net today more resembles a growing mass of band-aids than a plan.

The GENI Research Plan Goes On to Say…

We take a broad definition of security and robustness. A traditional focus of the security research community has been on protection from unwanted disclosure and corruption of data. We propose to extend this to availability and resilience to attack and failure. Any Future Internet should attain the highest possible level of availability, so that it can be used for "mission-critical" activities, and it can serve the nation in times of crisis. We should do at least as well as the telephone system, and in fact better.

Many of the actual security problems that plague users today are not in the Internet itself, but in the personal computers that attach to the Internet. We cannot say we are going to address security and not deal with issues in the end-nodes as well as the network. This is a serious challenge, but it offers an opportunity for CISE to reach beyond the traditional network research community and engage groups that look at operating systems and distributed systems design.

Cybersecurity, Cyberinfrastructure and Network Architectures

- So as you think about things like GENI and other "clean slate" efforts to rework our struggling Internet, remember: cybersecurity is front and center when it comes to driving new cyberinfrastructure architectures.

- That is, while we may not know exactly what the next rendition of the Internet will look like, but, without question, cybersecurity considerations will be a fundamental consideration.

- So how do clean slate efforts relate to efforts by today's application programmers?

Security and the Applications Programmer Today

- From an application programmers point of view today, the network doesn't "exist to be secure or robust," as mentioned in the GENI report -- it's existence and satisficing/sufficient levels of security are taken as givens.
- Why? Well, the network exists to facilitate the researcher's substantive scientific, engineering or other work. They've got an application they want to run, or a dataset they need to move from one site to another, and security is a secondary consideration at best.
- This is gradually changing over time, as the miscreants become more focused on application-specific vulnerabilities (such as SQL injection attacks, or XSS (cross site scripting) attacks), and the applications guys HAVE to pay attention, but most users would still prefer the network to simply be a clean pipe that just moved bits from one place to another -- they **want** a network that "just works," and which they don't have to worry about.

The Empirical Reality Can Be a Bit Different

- Rather than having a transparent end-to-end pipe, today's application programmer knows that they must potentially navigate a network encrusted with layers of firewalls, antivirus gateways, traffic shapers, proxies, and other active network security devices. Instead of being **a content agnostic "dumb pipe," the network has become s very content-aware and very nosy participant in the delivery (or NON-delivery!) of network traffic.**
- In other cases, the network is neither a dumb transparent pipe nor an intelligent active network participant, it may simplyintentionally not work at all. Some traffic intended for external hosts may be completely blocked, or that traffic may be involuntarily redirected without any notice to a local server. This is increasingly true when it comes to email traffic which may be blocked for anti-spam reasons if it isn't sent through the institution's email server, and more recently, DNS traffic has also been the subject of blocking or redirection in an effort to cope with DNS-changing malware.
- **The foundation of most sites' network security is the firewall.**

But Firewalls Can Interfere With Advanced Apps

- For a nice rendition of many of the problems that advanced grid applications can encounter due to firewalls, you may want to see

  "Firewall Issues Overview," August 16, 2006, http://www.ogf.org/documents/GFD.83.pdf

- That document does an excellent job of explaining the challenges that grid applications face in a firewalled environment, which is why grid-related systems often end up positioned outside an institutional firewall, in the "DMZ," or connected via

  > "dedicated high-performance physical or logical links as fiber, wavelength, sub wavelength, VPN, VLAN, etc. Assuming that external sources cannot gain access and misuse these links they are rarely secured by firewalls."

- Nonetheless, let's look at firewalls a little.

IV. Firewalls and Security: The Conventional Wisdom

Firewalls

- The “conventional wisdom” when it comes to system and network security often begins with (and in some cases, unfortunately, **ends with**) “firewalls.”
- As an example, if you happen to visit another campus and mention an interest in network and system security, the first thing you’re likely to hear about is that site’s firewall (assuming they have one).
- For some organizations -- and for some combinations of network architectures, application loads and security requirements -- firewalls **can** be an important part of a site’s security program.
- In other circumstances, however, firewalls can introduce single points of failure to an otherwise robust network design while also interfering with the operation of mission critical applications and effectively hindering (rather than helping) the identification and isolation of security incidents if they occur.

But Firewalls Are Ubiquitous

- Especially if we're talking about corporate environments, firewalls, like antivirus software, are ubiquitous.

- The 2007 CSI Computer Crime and Security Survey indicates that at least among 494 computer security practitioners in US corporations, government agencies, financial institutions, medical institutions and universities, **97% of respondents used firewalls.** For comparison, 98% used antivirus software (those were the top two security technologies used). See the 2007 CSI Survey at PDF page 19, http://i.cmpnet.com/v2.gocsi.com/pdf/CSISurvey2007.pdf

- Firewalls really have become **that** ubiquitous.

But What Is A Firewall?

- People can have many different ideas when it comes to thinking about firewalls, so let's make sure we're all on the same page.
- For me, a firewall is a hardware appliance (or a software application) that establishes a security perimeter, dividing an "inner" "more trusted" region of the network from an "outer" "less trusted" region.
- Thus, for example, a company might have a hardware firewall appliance sitting between their connection to the Internet and their corporate LAN, with the primary purpose of that device being to hinder remote network scans and exploits which might be targeting the corporation's internal systems.
- In other cases, a firewall may protect only a single subnet, or even just a single system (as in the case of software host firewalls).
- I'll even squint and call home cable modem or dsl "routers" (such as those from Linksys and other consumer vendors), "firewalls"

Why Do We Even "Need" Firewalls?

- Let me run through some of the argument you'll traditionally hear advanced in favor of perimeter firewalls…

  -- The Internet can be a "rough neighborhood," and there are some people out there with bad intentions. Thus, an unshielded Internet-exposed system will be subjected to a constant barrage of scans and attempted exploits, and it is prudent to block those attacks as far upstream as we can, whenever we can.

  -- There are some networked resources intended solely for private use and which were never meant to be accessed by the public (e.g., an employee-only "intranet"); we should also shelter those sort of resources from unauthorized access attempts.

  -- Firewalls can also help us tolog problematic traffic, and give us insight into attacks which we may be seeing.

Why Do We "Need" Firewalls? (cont. 1)

-- Firewalls also serve as a policy enforcement point for local users so that if the local policy is "no user-administered servers," or "no peer-to-peer applications," or "all email must go through our official company mail servers," or "all web access must go through the content filter," those policies can be technically enforced.

-- Firewalls may be required as a defacto matter of popular "common sense." For example, if you don't have a perimeter firewall and your system does happen to get compromised for whatever reason, journalists and other "monday morning quarterbacks" may immediately raise their eyebrows and drop their jaws in disbelief:

> "What? You didn't even HAVE a **firewall**? Well of COURSE you'll get hacked then! Sheesh! Don't you academic guys know ANYTHING?"

Why Do We "Need" Firewalls? (cont. 2)

-- Firewalls may also be required for the institution to be in compliance with Payment Card Industry (PCI) requirements, or as a matter of meeting auditor findings and recommendations.

-- Firewalls may be important for their contribution to reducing overall noise levels in your logs. If you're administering hosts and are constantly subject to wave after wave of probes from the script kiddies, it can be easy to miss more sophisticated attacks simply because of all the other background noise you're also experiencing.

-- Firewalls may also enable network address translation, so that rather than giving every workstation a globally routable address, hosts within the firewall may be given RFC1918 private addresses instead. Doing NAT can reduce requirements for globally routable IPv4 address space, and may reduce the ability of external systems to reconnoiter internal systems, although systems are by no means immune from attack just because they're using private addresses.

An ‘Unspoken’ Reason Why People Buy Firewalls

- **An unspoken reason why many people may buy firewalls is simply the desire to feel “safe” online.**
- I’d originally considered making an analogy between a firewall and a child's "special blanket," but there are some icons which are sacrosanct I finally decided that firewalls are actually like a bed with a thick and fluffy eiderdown comforter on a cold winter’s day.
- Once you’ve got a perimeter firewall deployed, many people may mentally feel almost as if they’re in a mountain cabin in the middle of a blizzard, laying beside a roaring fire in a warm feather bed, safe and snug in a place where one can cozily wait out whatever network craziness may be raging in the wilderness “outside.”
- That psychological sense of protection may be the biggest (albeit unacknowledged!) reason why many people like to have a firewall. A firewall conveys a sense of safety, just like a parent's comforting arms.

Ascription of Pop Psychological Motives Aside…

- There are some times when firewalls **are** technically necessary (or at least quite helpful).

1st Example: Reinstallation of MS Windows

- A classic example of a time when having a hardware firewall is a necessary is during reinstallation of Microsoft Windows.
- Because reinstallation of Microsoft Windows normally begins with installation of an unpatched version of the operating software from CD, followed by downloading and installation of numerous patches obtained from Microsoft over the network, there will normally be a period of time when an incompletely patched (and thus easily compromised) system will be connected to the network.
- If that incompletely patched system isn't sheltered by a firewall (at least for the online patching process), it is virtually certain that that newly installed system will be compromised before all the required patches can be downloaded and installed, even if the installer works as quickly and as diligently as they possibly can (the SANS Survival Time graph generally shows unpatched systems getting owned in under ten minutes, see isc.sans.org/survivaltime.html)

2nd Example: The Zero Day "Patch Window"

- Another example of a time when firewalls can be very handy is when new host vulnerabilities get discovered.

- Firewalls can provide protection during the period of time between (a) the discovery and disclosure of a new vulnerability, (b) the active exploitation of that vulnerability by miscreants, and (c) the release and installation of a vendor patch/workaround.

- This is particularly important as miscreants accelerate their efforts to reverse-engineer new vulnerabilities revealed by vendor patches. If it takes multiple days to roll out new patches to all applicable systems, but miscreants can reverse engineer patches and generate attacks in mere hours, we have a real problem if vulnerable systems aren't protected by some other mechanism (such as a firewall).

3rd Example: There Are Some Protocols Which May "Require" Our Use of A Firewall

- Sometimes you may be told you need a firewall because of inherent vulnerabilities in specific network protocols…

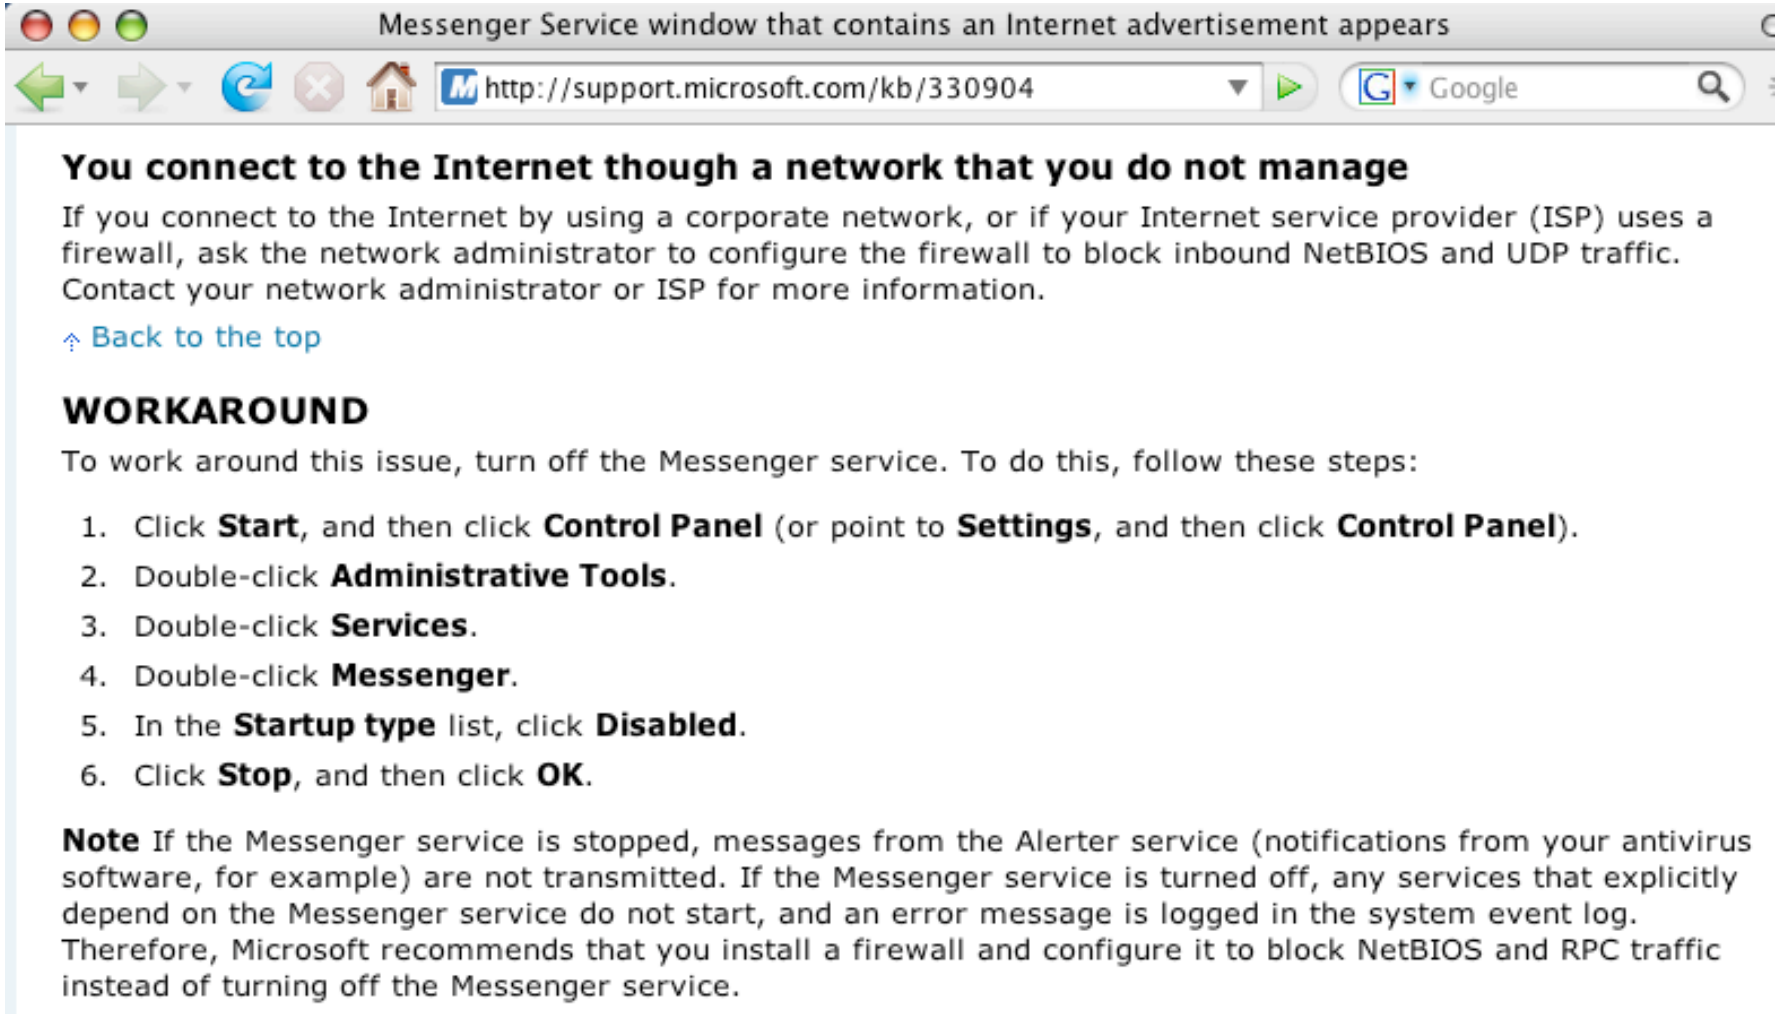

4th Example: Insecure VoIP Phones

- "One reason for having the VoIP phones on a separate VLAN is we firewall it. It turns out all these phones have Web servers — not browsers — in them and one way to configure them is to talk directly to the phone. All you need is the phone admin password, which is the same one in every phone and it's in the manual, so we don't let Web connections get to the VoIP phones, so security is at that level."

  "Behind The Scenes of MIT's Network: Network Manager/ Security Architect Jeff Schiller on Buying Into VoIP and Fiber In a Big Way," http://www.networkworld.com/news/2007/ 011907-mit-your-take.html

Are There Any General Security Principles Facilitated By Use of Firewalls?

- Sure… just to mention a few, how about:

  -- Least Privilege

  -- Defense in Depth

  -- Separation of Duties

Principle: Least Privilege

- One such security principle is that of "least privilege," or giving a person, program or computer only the access needed to allow the person/program/computer to do its required work.

- A classic example of the "least privilege" principle is embodied in the limited use of administrative ("root") accounts on systems. Since you generally don't need administrative privileges for most routine tasks, you use (or should use!) an unprivileged account for most of your day-to-day work, becoming a privileged user (e.g., via sudo) only when/if/for as long as may be necessary.

- Firewalls are a sort of network version of that same principle. If workstations aren't supposed to be running web servers, for example, why allow unsolicited incoming traffic from random Internet sources to go to port 80 or port 443 on those systems?

Principle: Defense In Depth

- Another key principle which firewalls facilitate is the notion of "defense in depth," a principle which some colloquially call the "belt and suspenders" principle. That's an image which nicely captures the rationale for this principle -- even if your belt breaks, if you're also wearing suspenders you'll avoid the embarrassment of inadvertently ending up with your pants around your knees.

- In computing and networking, defense in depth means that rather than relying solely on careful administration of networked computers to keep those computers safe (that's the "belt"), one can obtain additional insurance by also shielding those computers behind a firewall (that's the "suspenders"). That way, even if someone were to have an exploit which might work against some systems, they wouldn't be able to use that exploit if they can't get to those systems in the first place.

Principle: Separation of Duties

- Firewalls also provide a degree of what some might call "separation of duties." In accounting parlance, separation of duties means that whenever possible, you want two or more people to be involved in the execution of any potentially abuse-able processes.
- For example, you don't want the person who's requesting the purchase of equipment to also be the person who's verifying receipt of that equipment, and the person who's actually paying the vendor's invoices, and the person who's doing the annual inventory.It doesn't take a keen grasp of business to recognize that such a setup would make it very easy for a dishonest employee to rampantly abuse the purchasing process for personal financial gain.
- Firewalls and their administrators can play a similar role on the network. Without a firewall, or some means of monitoring/auditing network traffic, it may be hard for management to understand and control how their network's being used.

So, Then Firewalls Are Always Great, Right?

- After going through the preceding slides with me, you might think that firewalls are wonderful and that I'm a big fan of perimeter firewalls, just like "every other security guy," <cough>, <cough>!

  Well, you may be surprised to learn that **I'm actually NOT a big fan of traditional border firewalls**, and in fact, I think firewalls are fading as a mater of cultural interest… for example, if we check Google trends, the ultimate arbiter of all things "trendy," :-), we see…

Google Trends: firewalls

http://www.google.com/trends?q=firewalls&ctab=0&geo=all&date=all&sort=0

Google Trends

firewalls

Search Trends

Tip: You can compare searches by separating with commas.

Trend history

firewalls

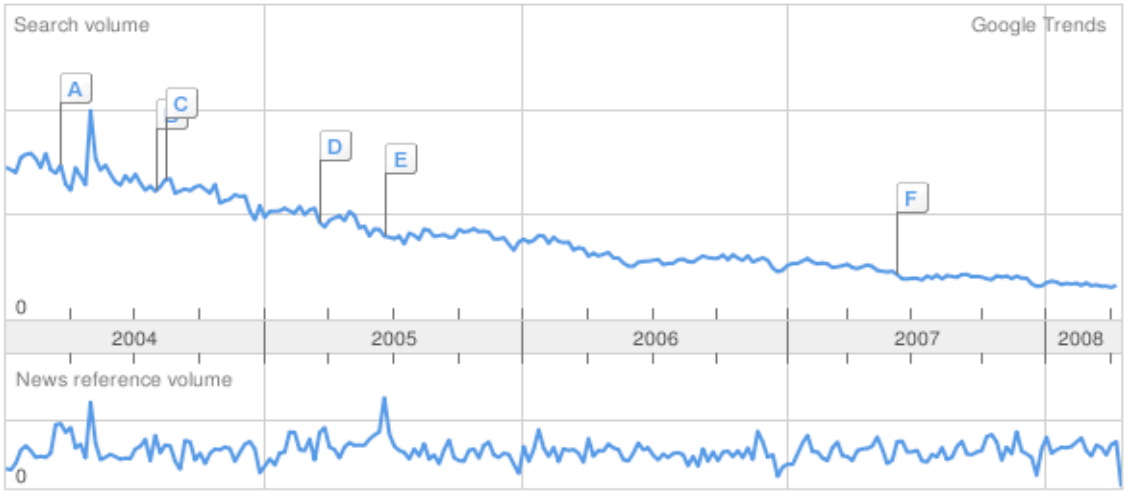

A Witty Worm Sneaks Through ISS Firewalls
InformationWeek - Mar 22 2004

B Honda Automobile, Thailand Drives Security Enhancements with Fortinet's FortiGate™ Antivirus Firewalls
Computerworld Australia - Aug 5 2004

C How Do Network Firewalls Work?
WHIOtv.com - Aug 16 2004

D Hackers Target Browsers To Dodge Firewalls
E-Commerce Times - Mar 21 2005

E Firewalls for WiFi Handhelds
Handheld Learning - Jun 24 2005

F Quality of Service and Firewalls with Linux and Enterprise Open Source Tools at SOA World
SYS-CON Media - Jun 14 2007

More news results »

All regions | All years

Regions

1. India
2. New Zealand
3. Australia
4. United Kingdom
5. United States
6. Canada
7. Mexico

Cities

1. **Chennai**, India
2. **Delhi**, India
3. **Brisbane**, Australia
4. **Melbourne**, Australia
5. **Sydney**, Australia
6. **London**, United Kingdom
7. **Atlanta**, GA, USA

Languages

1. **English**
2. **German**
3. **Spanish**
4. **Dutch**
5. **Portuguese**
6. **French**

But Even Before Google Was Tracking A Loss of Interest in Firewalls, Others Were Speaking Up…

- Deke Kassabbian, University of Pennsylvania, **March 2003**
  http://dolphin.upenn.edu/~deke/writing/fwatpenn.html
  "I believe that there is sometimes a tendency to try to solve too many problems through the use of firewalls without acknowledging their downsides. [* * *] Most systems can be made network-safe without firewalls, though some of the most common operating systems are far from secure 'out of the box'. I believe that this option should be explored before deciding to use a firewall."

- Abe Singer, San Diego Supercomputing Center, "Life Without Firewalls," **login;** the Usenix Magazine, **December 2003**, p. 34-41; http://www.usenix.org/publications/login/2003-12/pdfs/singer.pdf

Or What Did Bill Cheswick, Co-Author of the "Firewall Bible" Say Earlier This Month?

- The firewall world's bible is Cheswick, Bellovin and Rubin's *Firewalls and Internet Security: Repelling the Wily Hacker*, originally published in 1994, now in a revised 2nd edition.

- Imagine my surprise, then, when, while at RSA 2008 in San Francisco, during one of the keynotes, Bill Cheswick was interviewed by Herbert Thompson, and as part of that interview Cheswick stated that "**I haven't used firewalls in, uh, well, mostly, for ten years or more." and "They still have their use, but I really want my hosts to be secure enough they don't need a firewall."** (media.omediaweb.com/rsa2008/webcast.htm?id=4_1 at around the 34 minute mark)

- What a cool and timely comment, eh?

V. Firewalls and Security: A Reconsideration

So Let's Think About Your Risk Model

- It is a good idea to have a clear idea about the risks you're trying to mitigate when planning your security strategy.

- Ideally, your security strategies should align well with the risks you face.

- So what are some of the risks your university may face?

The Insider Threat

- When it comes to damage caused by cybercrime, we know from things like the *2007 E-Crime Watch Survey*, conducted by the United States Secret Service, the Carnegie Mellon University SEI CERT Program, Microsoft, and CSO Magazine (see http://www.sei.cmu.edu/about/press/releases/2007ecrime.html) that "when asked who caused more damage (in terms of cost or operations), results were fairly close (insiders 34%, outsiders 37%, unknown 29%)." Clearly, the insider threat is a non-trivial animal creeping around your network jungle.

- But if you're worried about the damage associated with the "insider threat," traditional border firewalls won't help you, because the trusted insider is, well, already inside that perimeter…

Loss of Personally Identifiable Information

- The cybersecurity incident we seem to see reported in the news most often is the loss of personally identifiable information, perhaps in the form of thousands (or millions!) of records with social security numbers or credit cards or other sensitive information extracted from an institutional database and exfiltrated elsewhere.

- Once a bad guy is inside your perimeter, most firewalls don’t prevent, and aren't meant to prevent, the exfiltration of personally identifiable information (PII) -- most firewalls routinely permit arbitrary outbound traffic to all external destinations. In fact, when it comes to PII, we might say that your firewall is focused on the wrong sort of traffic, blocking traffic that's trying to come in, rather than paying attention to traffic that's trying to go out!

Distributed Denial of Service Attacks

- But there are other threats that firewalls also fail to handle. What if you're worried about distributed denial of service attacks?
- If you're not worried about DDoS, maybe you should be -- Arbor networks reports that on average they see 1,300 DDoS attacks per day, and DDoS attack traffic now consistently accounts for 1-3% of all inter-domain Internet traffic, see asert.arbornetworks.com/2008/03/2-of-internet-traffic-raw-sewage
- Surely the firewall must do a yeoman's job of stopping great gobs of packet dung from being hitting your systems from random destinations around the Internet? Well, yes, the traditional border firewall can be used to stop that sort of unwanted traffic, but at least in the case of raw traffic floods, doing so won't help you because by the time that DDoS traffic hits and gets blocked by your firewall, it's too late -- your upstream link(s) will already have been saturated. Firewalls can't protect your site against DDoS.

What About Malware?

- How about malware, things like viruses and trojan horses and all the rest of the animals in the evil software menagerie?
- Again, nope, firewalls won't help, at least not unless you proxy **all** traffic (unencrypted/in plain text!) through a firewall that actually includes an antivirus gateway or which acts as a unified threat management device, and even then we know that the bad guys can generate new malware, or tweak/repack old malware, more rapidly than the antivirus guys can release newly updated signatures.
- And of course, being a security conscious sort of entity, I'm sure much of your traffic WILL be encrypted end-to-end, which means that if won't be visible to a virus-scanning firewall or unified threat management device at all.
- Finally, if you're using a proxy-based firewall solution, that proxy will only handle the protocols it has been trained to understand, so I'm really sorry, but your allowed applications just contracted a bit.

So What About The *Traditional* Threats?

- Firewalls were traditionally intended to mitigate classic computer intrusions, e.g., "cracking/hacking," and all the attendant precursors thereto (such as automated scans and brute force login attempts) -- but is that really what you're worried about anymore?
- Important systems are no longer secured by just traditional passwords flowing in plain text over the wire. Now access control may include two factor methods such as hardware tokens (SecureID, CryptoCards, etc.), and traffic will more than likely be encrypted, reducing eavesdropping exposure.
- And, given improvements in passive traffic monitoring systems (such as Snort, Bro, etc.), do you really need a firewall to just document the ineffective attacks that you're seeing? Can you get any satisfaction when it comes to getting that abuse stopped?
- Let's assume you do stay with a traditional firewall. You can end up paying a large (albeit largely non-financial) price…

Firewalls Can Throttle Throughput

- If your application needs the ability to deliver traffic at a high rate of speed, firewalls can act as a choke point, throttling that throughput.

- For many years, the fastest border firewalls topped out at gigabit speeds; these days faster firewalls are available, but those high end systems aren't cheap and if you load them down with complex rules, they may begin to lose their ability to keep up.

- Edge firewalls may be particularly prone to throttling traffic. Those devices built to meet the needs of a price sensitive consumer broadband market dominated by DSL or cable modem connections, will be ill suited to protecting fast ethernet or even gigabit connections.

Sample Firewall Throughput Discussion

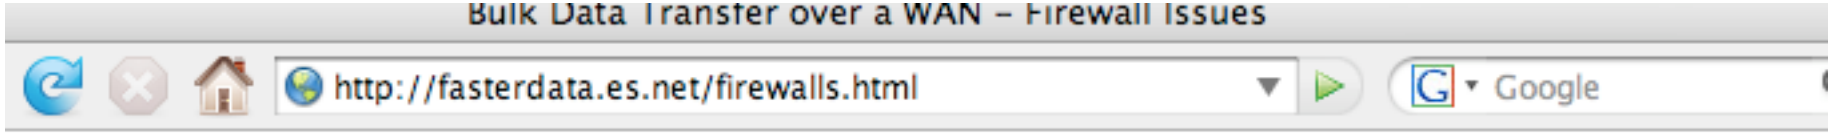

Firewall Performance Issues

Firewalls are often slower than the link speed of their network interfaces (e.g. many firewalls with Gigabit Ethernet interfaces have a maximum throughput rate of 800 Mbps). This causes a problem when a host with a network interface that is faster than the firewall's internal processor attempts to send data through the firewall (TCP bursts typically occur at or near the maximum data rate of the sending host's interface). Since the firewall must buffer the traffic bursts sent to it by the data transfer host until it can process all the packets in the burst, input buffer size is critical. Unfortunately firewalls often have small input buffers, since they are typically designed to scale to large numbers of low-speed flows, rather than a few high-speed data flows. If the firewall's input buffers are too small to hold the bursts from the data transfer host, packet loss will result -- often causing severe performance problems.

Firewalls Interfere With Active Scanning

- Interior firewalls, and "personal firewalls" such as Linksys broadband "routers" can interfere with active scanning of your networks with tools such as Nessus.

- This is sort of a good news/bad news story. On the one hand, if you can't scan that host, hopefully the bad guys can't either. On the other hand, while you can't scan that host maybe some remote party can (you can never tell what holes may have been punched in a firewall), and maybe they'll find that that host has a remotely exploitable vulnerability.

Firewalls Can Add Significant Complexity

- Suddenly, instead of having a network that simply passes packets, and which is either ALWAYS passing packets (if it is up) or NEVER passing packets (if it is down), we have a network which SOMETIMES passes packets, and which other times intentionally (and/or inadvertently) does not.
- That can add tremendous complexity when it comes to trying to debug why an application does or doesn't work…only works intermittently… or sort of works (but only poorly).
- Suddenly, anytime there's a problem with that application, in addition to everything else, the possibility that a firewall may be interfering needs to be investigated and ruled out, both locally and potentially at remote sites as well.

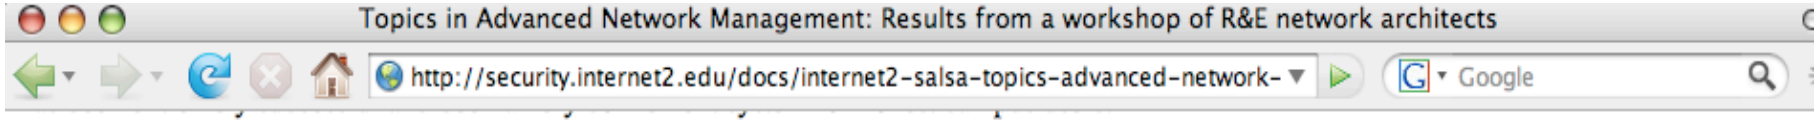

4. 0 Packet Disruption Devices

The afternoon session on the second day went deep into the ecosystem of packet disruption and shaping devices, including firewalls, load balancers, and intrusion prevention systems (IPSs). Deployments and philosophy amongst campuses showed more variability than was seen with other parts of the discussion.

There are a lot of deployment concerns that are common to all packet disruption devices regardless of the purpose of the device. Mitigation techniques exist to address some of the problems for some devices, but in general, most of these challenges will have to be the subject of future research to enhance the devices and protocols used to perform these functions or reduce the need for them entirely. Three concerns were dominant.

First, these devices may limit network availability through their own failure, due to device failure or simple volume of traffic that these devices can handle relative to what the underlying network would otherwise be able to carry. This ratio becomes worse for the systems designed to provide protection at higher levels of the protocol stack, such as intrusion prevention systems (IPSs), due to the increased processing implied. Deployments must take these limitations into account.

There are also limitations imposed on the set of architectures that can be used in conjunction with the devices, since these boxes generally operate as single points on the network. Virtual and physical networks in particular may be used to route traffic selectively around or through these devices. If the system is deployed without full awareness of the infrastructure into which it's placed, security vulnerabilities may arise. On the other hand, it is also necessary to structure the deployments such that a set of hosts on the internal network may be excluded from these devices. Only two schools present even had complete control of the network from backbone uplink down to the wall jack throughout their entire network.

Lastly, and perhaps the hardest of the problems is the loss of end-to-end transparency and diagnostic ability. Firewalls are the most notorious example, but other packet distruption devices can give misleading or imperfect information about the state of the network itself. (E.g. one copyright music detection appliance scans network traffic for music signatures; when finding such a stream, the appliance emits a TCP reset command to sever the connection and confuse the user and diagnostician alike.) If some form of packet disruption is responsible for service failure, very little beyond thorough knowledge of the fingerprints left by these devices, the entire network structure, and intuition can be any guide to diagnosis.

Nevertheless, these systems are important tools and the group spent the afternoon of the second day discussing in detail the wide variety in packet disruption devices and their deployment in higher ed.

You May Not Be Facing Just One Firewall

- In addition to your institutional perimeter firewall, you may also need to get traffic through an interior departmental firewall, and perhaps also a software firewall running directly on a given system, and the same may be true on the opposite end of the connection as well. Think about that -- your application traffic might need to transit no less than half a dozen nested or "daisy chained" firewalls!

- If your chances of getting traffic through one firewall are slim, imagine how things will go when you try to transit six of them!

- At some point the whole firewall business becomes rather absurd.

Your Firewall Administrator Is Unknown

- Many times the only person who can definitively tell you if a/the firewall is in fact the culprit is the firewall's administrator.

- Survey your users:
  -- If you had trouble with your desktop PC, such as a problem with Microsoft Word, or Excel, who would you ask for help?
  -- What if you had a problem with your email, such as perhaps you'd forgot your password? Who/how would you reset it?
  -- What if you had network problems, such as your network jack suddenly stopped working -- who would you call then?
- **Then** ask, "What if you installed a new application and it didn't work, and someone suggested that your campus firewall might be keeping it from working. Who would you contact for that sort of firewall-related problems?" I suspect that far fewer people know who runs their school's firewall or their department's firewall.

Crunchy on the Outside, Soft on the Inside

- Another problem with firewalls can be the phenomenon known as being "crunchy on the outside, but soft on the inside," or network services which are superficially secure (at the firewall), but which are exceedingly vulnerable when you begin to look at hosts living within the firewall.

- Why's that? Well, some may ascribe this phenomena to unwarranted overconfidence or a false sense of security:

  "Heh, isn't it great to have a firewall? Now we no longer need to worry about keeping all our systems patched up to date! When the bad guys attack, they'll just get bounced off our firewall!"

Twenty Thousand of Your Closest (And Presumably Most Trustworthy) Friends

- There are other issues with perimeter firewalls, particularly when they're applied in a higher education context.
- If a perimeter firewall is being deployed at a large research university, it might be separating the billions of users on the Internet from the local university community (let's call that "just" twenty thousand faculty, students and staff, although obviously some schools will be larger and some smaller).
- How secure do you feel having 20,000 users within your "circle of trust?" If even 1/10th of 1% of those "insiders" are untrustworthy (or even just careless or easily tricked), that means you still have 20 points of vulnerability within your security perimeter!
- If you've got a security mindset, you'd probably be much happier if you were trusting a much smaller number of people.

Firewalls Can Discourage Experimentation and Innovation

- Academia is a funny place -- you never can predict who's going to develop some cool new innovative application. It might be a person employed as a member of the institutional IT staff, but it might also be a faculty member in the humanities or social sciences, or an undergraduate student -- you just can't tell.

- But now let's put that person behind a firewall, and see what happens. Does the volume and quality of her experimental code and the quality of her innovations increase, or decrease?

- Oh, I'll certainly grant that you're likely to see many innovative approaches to overcoming the challenges that firewalls impose, but I don't think overcoming firewalls is the one and only topic meriting developer attention.

Sometimes Firewalls Will Definitely Break Apps

- Consider H.323 video conferencing as an example. Because of the complexity of the H.323 protocol, it has historically been quite difficult to work with H.323 video in firewalled environments, particularly if the firewall is doing NAT.
- There has been effort over the last few years by the ITU and some vendors to make H.323 firewall traversal work (see the Radvision white paper on H.460.17, H.460.18, and H.460.19 at http://www.h323forum.org/papers/301005_Firewall_NAT_Traversal_White_Paper.pdf), but how widely deployed and how interoperable are different vendors implementations of those protocols?
- And what of legacy equipment which may not support H.323 firewall traversal enhancements?
- What do we actually here from those “on the street” who are doing H.323 video conferencing support?

Recommendations for H.323 From One Network

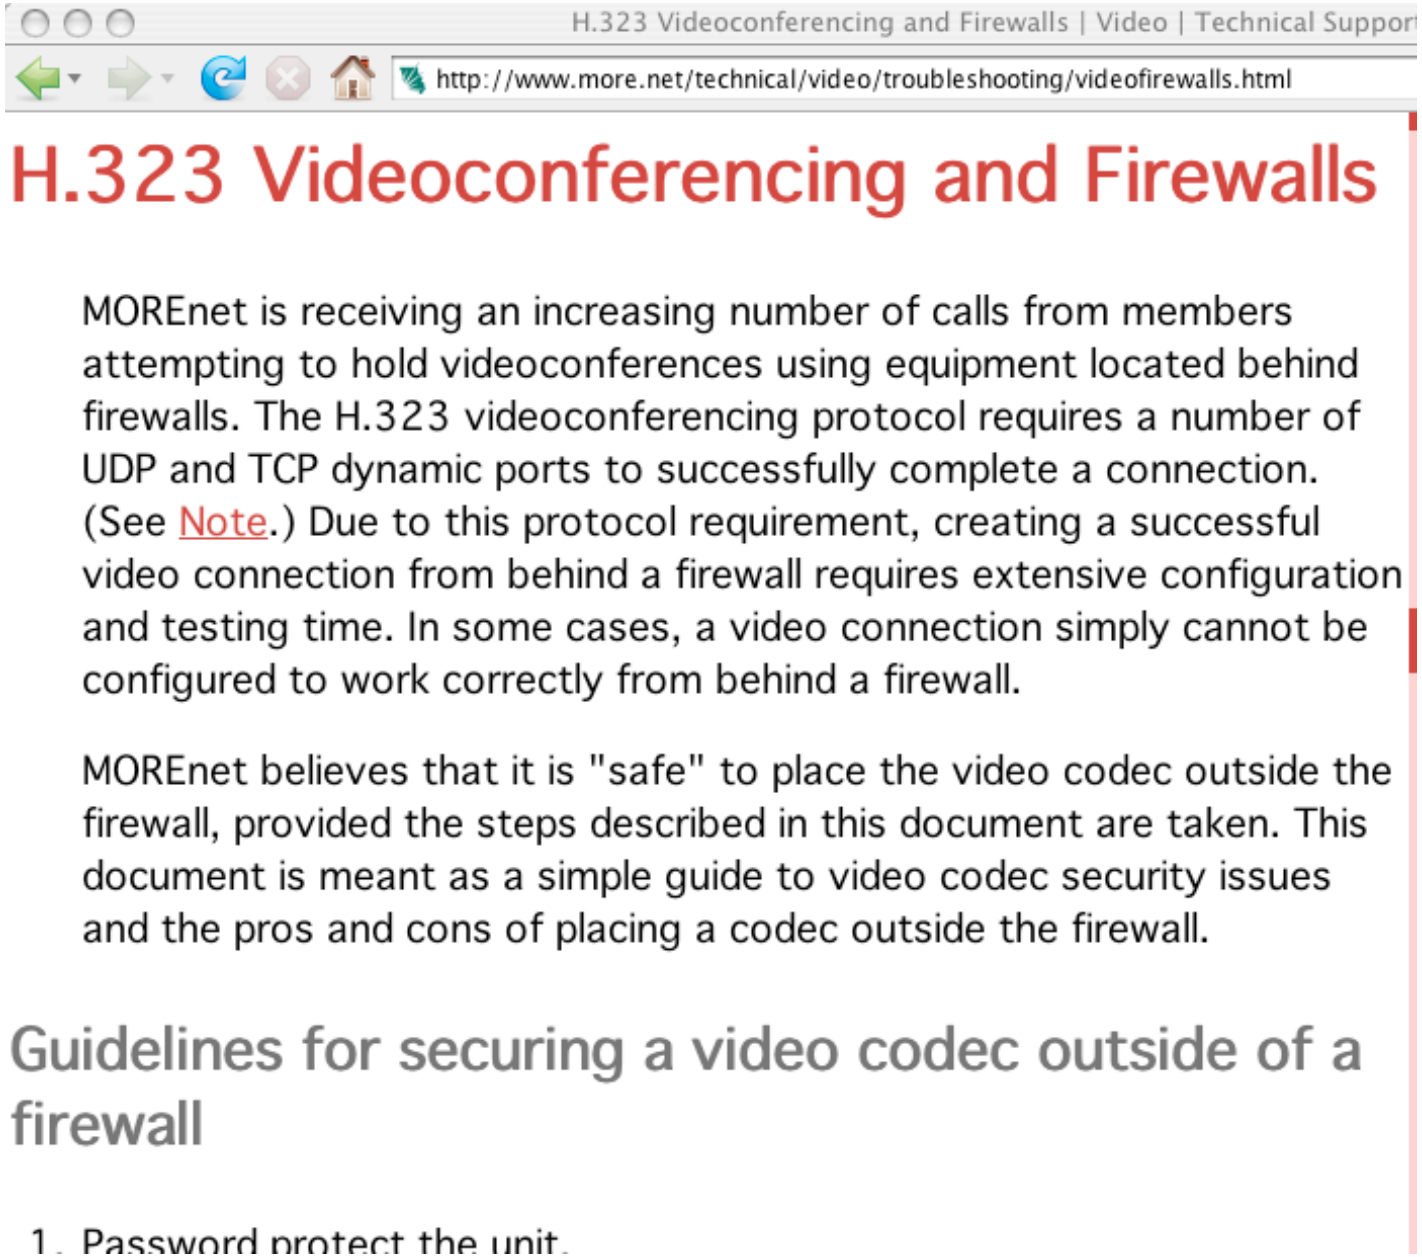

H.323 Videoconferencing and Firewalls | Video | Technical Support

http://www.more.net/technical/video/troubleshooting/videofirewalls.html

H.323 Videoconferencing and Firewalls

MOREnet is receiving an increasing number of calls from members attempting to hold videoconferences using equipment located behind firewalls. The H.323 videoconferencing protocol requires a number of UDP and TCP dynamic ports to successfully complete a connection. (See Note.) Due to this protocol requirement, creating a successful video connection from behind a firewall requires extensive configuration and testing time. In some cases, a video connection simply cannot be configured to work correctly from behind a firewall.

MOREnet believes that it is "safe" to place the video codec outside the firewall, provided the steps described in this document are taken. This document is meant as a simple guide to video codec security issues and the pros and cons of placing a codec outside the firewall.

Guidelines for securing a video codec outside of a firewall

1. Password protect the unit.

Some Tools Have Given Up/Changed Protocol Architectures; Consider EVO/Koala

APENDIX C. – Firewall settings for Koala

Mandatory:
While EVO works fine in a Network Address Translation environment (NAT), the local or institute firewall (if any) should permit communication on the following port.

- ***IN/OUT - UDP/TCP***: **46015**

Optional:
Allowing these TCP outgoing ports will offer the possibility to your Koala client to estimate in real-time what are the best Panda servers to connect to in function of your location, network parameters (bandwidth, packet loss...) and load.

- ***OUT - TCP*** on specific IP
  1) LUSs Services: **4042, 4043, 4044**
     evo01.cern.ch (192.91.244.138)
     evo01.caltech.edu (131.215.116.151)

  2) Proxy Services: **60001, 60002, 60003**
     evo01.cern.ch (192.91.244.138)
     evo01.caltech.edu (131.215.116.151)

  3) Topology Services: **10090**
     evo01.cern.ch (192.91.244.138)

Firewalls and IPv6: Oops!

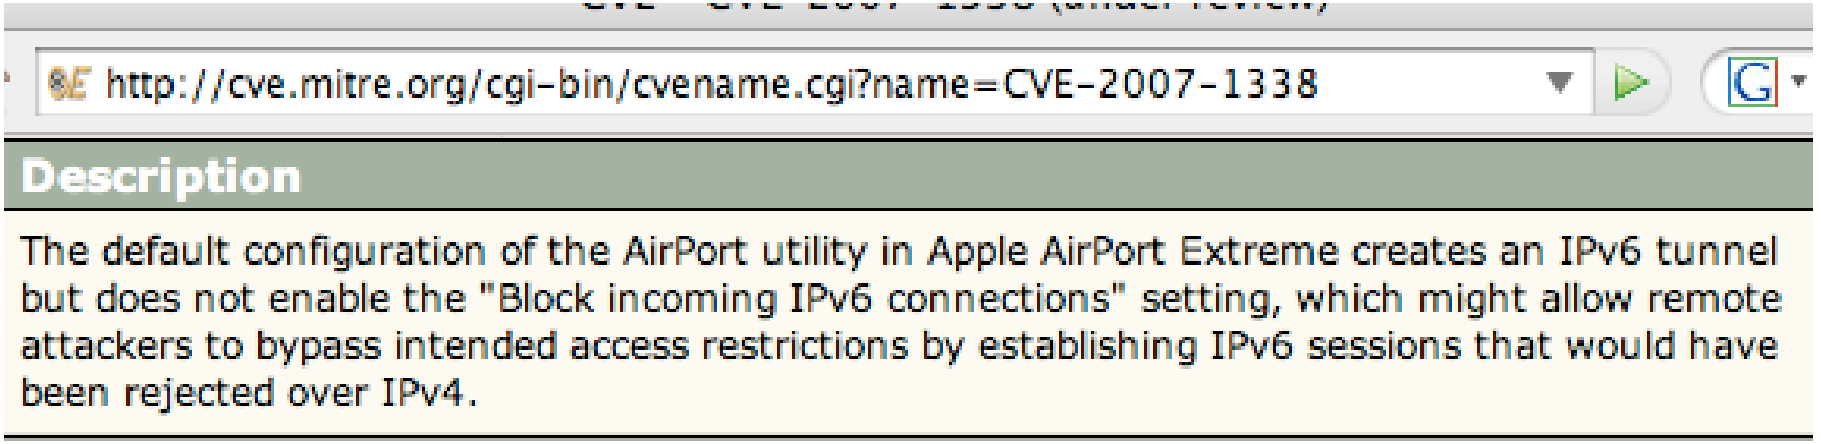

- The point? There's sometimes a tendency to sometimes forget that IPv6 exists when planning perimeter security. (This "vulnerability" associated with this default condition has been "corrected")

Another Take on IPv6 and Firewalls…

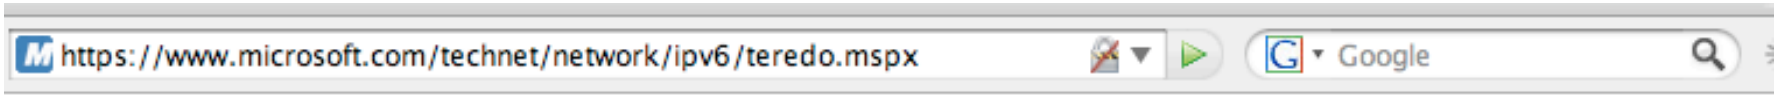

Teredo and Protection from Unsolicited Incoming IPv6 Traffic

As described in Using IPv6 and Teredo, IPv6 traffic that is tunneled with Teredo is not subject to the IPv4 packet filtering function of typical NATs. Although this might sound like Teredo is bypassing the NAT and allowing potentially malicious IPv6 traffic on private networks, consider the following:

- Teredo does not change the behavior of NATs. Teredo clients create dynamic NAT translation table entries for their own Teredo traffic. The NAT forwards incoming Teredo traffic to the host that created the matching NAT translation table entry. The NAT will not forward Teredo traffic to computers on the private network that are not Teredo clients.

- Teredo clients that use a host-based, stateful firewall that supports IPv6 traffic, such as Windows Firewall, are protected from unsolicited, unwanted, incoming IPv6 traffic. Windows Firewall is enabled by default for Windows XP with SP2, Windows Vista, and Windows Server 2008.

The combination of IPv6, Teredo, and a host-based, stateful, IPv6 firewall does not affect the packet filtering function of the NAT for IPv4-based traffic and does not make your Windows-based computer more susceptible to attacks by malicious users and programs that use IPv6 traffic, rather than IPv4 traffic.

Firewalls May Be Easily Circumvented

- Consider, for example, a reflexive firewall policy which:

  -- denies inbound traffic from the Internet by default,
  -- permits outbound traffic, and
  -- allows inbound traffic in response to output requests.

- If a miscreant can just convince one of your users to visit a malicious web page, at that point:

  -- the miscreant can deliver content which can be run by the user's browser inside your firewall, or

  -- the miscreant can circumvent your firewall entirely by creating outbound traffic from the user's system to arbitrary servers of the miscreant's choice, and then "responding" to those requests

If You've Gotten 0wn3d, It's Too Late For Firewalls

- Some technologies (such as backups, or file alteration checksum-based technologies) are of value even if (especially if!) you've been compromised.

- Firewalls, on the other hand, tend to be like contraception for a pregnant couple -- once you're pregnant (or once your system has been compromised), the relevance of "barrier methods" drops!

Firewalls Don't Even Make Networks Opaque

- Sometimes you'll hear security officers talk about how, if nothing else, firewalls at least can provide some level of network opacity, blocking things like traceroute from showing the precise path that traffic may take to a given host.

- That belief, like many other beliefs relating to firewalls, is rather ill-founded

Normal Traceroute, Encountering a Firewall

```
traceroute to www.cnn.com (64.236.91.22), 30 hops max, 40 byte packets

[* * *]
12  pop2-sun-p5-0.atdn.net (66.185.147.25)  24.425 ms  24.399 ms  24.308 ms
13  bb1-sun-p0-1.atdn.net (66.185.140.192)  73.108 ms  89.468 ms  199.94 ms
14  bb1-den-p7-0.atdn.net (66.185.152.252)  49.082 ms  49.261 ms  48.919 ms
15  bb2-den-p1-0.atdn.net (66.185.152.137)  49.822 ms  49.476 ms  49.17 ms
16  bb2-kcy-p7-0.atdn.net (66.185.152.189)  66.478 ms  66.55 ms  66.665 ms
17  bb1-kcy-p1-0.atdn.net (66.185.152.126)  66.201 ms  66.359 ms  65.715 ms
18  bb1-chi-p6-0.atdn.net (66.185.152.124)  86.64 ms  121.681 ms  200.262 ms
19  bb2-chi-p7-0.atdn.net (66.185.152.131)  65.529 ms  65.715 ms  66.461 ms
20  bb2-vie-p14-0.atdn.net (66.185.152.215)  85.269 ms  84.892 ms  85.044 ms
21  pop1-vie-p2-0.atdn.net (66.185.139.83)  85.145 ms  85.454 ms  84.76 ms
22  dar1-mtc-s3-0-0.atdn.net (66.185.139.134)  86.395 ms  86.802 ms  85.668
    ms
23  * * *
24  * * *
25  * * *
26  * * *
27  * * *
28  * * *
29  * * *
30  * * *
```

The point? Using **normal** traceroute, we don't know what's happening at hops 23 and beyond.

tcptraceroute? The Firewall's Irrelevant

```
Tracing the path to www.cnn.com (64.236.91.22) on TCP port 80 (http), 30 hops
max
[* * *]
12  pop2-sun-p5-0.atdn.net (66.185.147.25)  25.522 ms  24.905 ms  24.748 ms
13  bb1-sun-p0-1.atdn.net (66.185.140.192)  24.551 ms  24.120 ms  24.530 ms
14  bb1-den-p7-0.atdn.net (66.185.152.252)  49.380 ms  49.404 ms  49.384 ms
15  bb2-den-p1-0.atdn.net (66.185.152.137)  49.270 ms  49.323 ms  49.535 ms
16  bb2-kcy-p7-0.atdn.net (66.185.152.189)  66.628 ms  66.818 ms  67.929 ms
17  bb1-kcy-p1-0.atdn.net (66.185.152.126)  66.392 ms  66.711 ms  67.843 ms
18  bb1-chi-p6-0.atdn.net (66.185.152.124)  66.834 ms  66.283 ms  65.865 ms
19  bb2-chi-p7-0.atdn.net (66.185.152.131)  131.483 ms  204.054 ms  202.901
    ms
20  bb2-vie-p14-0.atdn.net (66.185.152.215)  85.445 ms  87.108 ms  85.174 ms
21  pop1-vie-p2-0.atdn.net (66.185.139.83)  226.130 ms  199.190 ms  200.41 ms
22  dar1-mtc-s3-0-0.atdn.net (66.185.139.134)  86.276 ms  86.551 ms  86.754
    ms
23  cnn.atdn.net (66.185.144.58)  86.504 ms  86.956 ms  86.521 ms
24  www.cnn.com (64.236.91.22) [open]  86.140 ms  86.576 ms *
```

Note that firewalls don't even do a very good job of helping you hide your internal infrastructure (reverse proxies may be a different issue). Using tcptraceroute, we can easily see the last couple hops. (To get tcptraceroute, see michael.toren.net/code/tcptraceroute/.)

So Let Us Be Bold Today

- I believe the time has come for us to move beyond the traditional firewall, or even beyond things like unified threat management systems.
- Let us decide to look critically at firewalls wherever they may be deployed on our networks, and let us reconsider if we **really** need to have those boxes everywhere on our networks.
- We know that a ship in the harbor is safe, but that's not what a ship is for. We have ships to go forth boldly and sail the oceans of the world. A ship represents an opportunity for the self-assured to do great things.
- Our networks and systems are just like ships. We can hunker down and continually attempt to stay safe in the harbor, hiding behind institutional firewalls, or we can boldly go forth and do great things, and return safely to tell the tales of our adventures.

You Won't Be Alone, Living Without Firewalls

- Plenty of people have all sorts of systems exposed to the Internet, **without** a traditional border firewall, and without getting compromised, and they do it safely day after day.

- Am I telling you that you can safely put a poorly administered, unpatched and unhardened system on the network without a firewall, and expect to avoid getting 0wn3d? No, I'm **NOT** saying that.

- I AM saying that if you do a careful job of patching and hardening, MANY systems CAN be safely exposed to the Internet without a traditional border firewall without dire things happening, and on balance, the benefits associated with having a clear channel to the Internet will likely outweigh the incremental risks and the effort associated with patching and hardening hosts to be ready to do so.

"But I HAVE To Firewall My Admin Systems!"

- Okay, I hear you. You may need to have a firewall to be in compliance with some regulatory requirements. Fine. So be it.
- If you **have** to have a firewall, or even if you just **WANT** to have a firewall, push that firewall as close to the systems that need to be protected as you can.
- Rather than putting that firewall at the border, with 20,000 potentially compromised systems or potentially untrustworthy users inside the perimeter, push that firewall back until it is just in front of the sensitive subnet, or even consider a host-based firewall running on each of your sensitive hosts.

- An excellent discussion of this can be found in Terry Gray's *Firewalls: Friends or Foes*, Educause Review, Jan-Feb 2003, http://staff.washington.edu/gray/papers/fff-final.htm

"You Sure Seem Anti-Firewall All Right! Do You Even Oppose Antispoofing Filters?"

- No, I have no problem with people filtering spoofed traffic, and in fact I encourage people to do so. (See my earlier talk, 'A Brief Practical Security "Punchlist,"' Internet2 Member Meeting, April 25th, 2006, http://www.uoregon.edu/~joe/punchlist/punchlist.pdf)

- All I ask is that you **balance** network security and network usability. Antispoofing filters are a nice example of a security enhancement which should have NO effect on network usability for legitimate applications.

- Helping universities to make those sort of decisions is something that a "Network Usability Officer" might be able to do.

VI. What's A Network Usability Officer?

It's Time For a “Network Usability Officer"

- Many of our institutions already have:

  -- An **Information Security Officer** or a **Chief Information Security Officer**

  -- Acting as a partial check/balance to the ISO/CISO, there may also be an institutional **Privacy Officer**, insuring that faculty/staff/student/patient privacy rights are respected, both from a philosophical and from a regulatory compliance perspective

- But has the time come to consider a 3rd role, what might be called a **Network Usability Officer (NUO)**, someone dedicated to insuring that as we act to preserve our security and protect our privacy, we don’t simultaneously and inadvertently destroy the very usability of our networks and computer systems?

The Case for a Network Usability Officer Role

- ISOs/CISOs and institutional privacy officers often have strong institutional mandates backed up by statutes, senior management demands, popular user demand and press oversight. It can be hard to say “no” to people in those roles given the ever-lingering specter of the institution being hit by security breaches or privacy-related lawsuits.
- And yet, without someone representing the other side of the coin, it is easy to get the balance wrong, and go TOO FAR in the direction of security and/or privacy, losing site of the need to preserve the usability of the network as an institutional asset.
- You may have invested millions in your network and associated systems -- wouldn’t it be good if you could actually USE those resources for their funded purpose?
- You **need** someone who's “at the table” to insure that usability is also given due consideration when decisions are being made.

Won't the Users/Applications People Just Argue Their Own Case for The Importance of Usability?

- They might, if the institution is fortunate enough to have users or applications people with unusual levels of self-confidence and leadership… but will we hear them and will we listen to them?
- Do their voices have "official" institutionally-acknowledged legitimacy, on par with those speaking for the security community and privacy interests on campus?
- Will they know who they need to speak to, and how to frame their concerns in light of the other issues that institutional policy makers are concerned with? Will they understand budget issues, and historical security/privacy issues which may have occurred?
- Will they speak with a coordinated, consistent and uniform voice, or will the institution face a blizzard of conflicting, inherently inconsistent, and constantly changing requests? I think it is in the institution's best interest to have **someone** "own this" problem.

So What If A School Doesn't Add a NUO?

- I have few illusions -- I fully understand that most campuses will not create Network Usability Officer positions, at least not right now. Budgets are tight, staff are needed for other projects, and there's not a large cadre of applicants perfect for this work. Heck, there's not even a NUO certification program yet! :-)
- But, if your institution does continue to ignore the need for someone who's focused on making sure that networks remain usable, over time, like an insidious illness, your network will inexorably become less useful, less flexible, and less productive.
- Some newly released applications will never work. Other older applications which once worked may cease working.
- Application programmers may devote tremendous effort to adapting their applications to whatever connectivity remains: "The only port that's still usable is 80/tcp -- the web -- so I guess that's just how we'll have to implement our protocol."

BCP 56

- The IETF can read those programmers' minds. They KNOW that application programmers, out of firewall-related frustrations, have turned to "everything over port 80", and fortunately they provide at least some guidance (although a lot of water has flowed down the http river since February 2002).

- If you haven't read Best Common Practice 56 (http://tools.ietf.org/html/bcp56), "On the use of HTTP as a Substrate," K. Moore, University of Tennessee, February 2002, I recommend it to you. It explains some of the factors which motivate application designers to use http as a foundation for their applications, and some of the architectural considerations which should be carefully weighed before using http as a substrate for other application protocols.

But Even Port 80 Isn't A "Clear Channel"

- Part of the problem with doing everything over port 80 is that even port 80 (especially port 80!) isn't a "clear channel."
- Port 80 traffic may be proxied/cached, inspected/filtered, accelerated or traffic shaped (or both simultaneously!), virus/spyware scanned, despammed, and generally pushed, pulled, poked, prodded, mangled and tinkered with…
- Heck, since this is basically the only path from behind the firewall to the world, you'd have to be crazy NOT to "throw everything you've got" at potentially malicious traffic on it…
- So go ahead, base your mission critical application on traffic flowing over this "theme park thrill ride" of a twisting, turning and constantly changing/always surprising network channel. If for some reason your application doesn't work as it should, I'm sure you'll have a very easy time debugging what went wrong, and where (and why) those problems occurred…. NOT!

‘True! Port 80 IS Miserable! But That’s Okay… I Can Also Use Port 443 (“https”) And Since That’s Encrypted, My Traffic Will Be Safe! And Private, Too!’

- At one point, I thought that, too.

- Unfortunately, Man In The Middle (MITM) attacks have reduced my confidence in that proposition, including traffic that's been TLS/SSL encrypted.

VII. Speaking of MITM, Remember Layer 2 Security?

The Problem With Ethernet Switches

- All of us have ethernet switches on our campuses, however in many cases the layer 2 security of those devices has been neglected relative to other potential cyberinfrastructure threats. Ethernet switches, perhaps more than any other type of networking gear, are wonderful in that they basically "just work" out of the box with minimum configuration required. Unfortunately, ethernet switches are simultaneously terrible for exactly the same reason.

- Depending on the amount of attention you've devoted to your ethernet switches, you may be vulnerable to ARP poisoning and other Layer 2 attacks that often receive comparatively little attention. If you or your staff have already addressed layer two security issues, feel free to skip this section, and sorry for bringing up topics which you already have well in hand.

Main L2 Risks: Redirection and MITM/Sniffing

- From time to time you may run into administrators who assume that ethernet switches provide protection against traffic being sniffed on the wire. That is an erroneous assumption. Ethernet switches can be forced to forward traffic to a local attacker who may then be able to do man-in-the-middle eavesdroping attacks against those packets.

- Three of the better known tools used for this purpose are Dug Song's dsniff (see http://monkey.org/~dugsong/dsniff/), Cain and Abel (see http://www.oxid.it/cain.html), and ettercap (see http://ettercap.sourceforge.net/) Particularly see:

  -- http://www.oxid.it/ca_um/topics/apr.htm and
  -- http://www.monkey.org/~dugsong/dsniff/faq.html

How Do I Fix These Vulnerability?

- I claim zero L2 expertise (but what a great talk from someone else at the next Joint Techs or next Member Meeting, eh?), but some of the L2-targeted solutions which are often recommended include:

  -- enabling port security, or using static ARP entries (ugh, painful/poor scaling properties)
  -- monitoring ARP traffic with something like ARPWatch
  -- improving segmentation (moving to a unique VLAN per host, although that may not be practical for large networks)
  -- DHCP snooping
  -- BPDU (bridge protocol data unit) filtering traffic from hosts
  -- DHCP Option 82
  -- disabling trunking except where needed

  although you also need user education about SSL certs, etc.

Some Sources for Further Information

- "Cisco IOS Switch Security Configuration Guide," 21 June 2004, www.nsa.gov/snac/os/switch-guide-version1_01.pdf

- "Catalyst Secure Template," 11/01/2002, www.cymru.com/gillsr/documents/catalyst-secure-template.htm

- LAN Security: What Hackers Know About Your Switches, Cisco Press, September 2007, ISBN 978-1587052569, 360 pps., www.amazon.com/LAN-Switch-Security-Networking-Technology/dp/1587052563

VIII. One Last Cybersecurity Architecture Example: The Federal TIC Program

The TIC Program and The Reduction In Agency Connections (Thousands Down to 50)

- The federal government recently announced that it plans to reduce the number of connections it supports between federal agencies and the Internet from thousands down to just a target of 50. See www.whitehouse.gov/omb/memoranda/fy2008/m08-05.pdf and www.whitehouse.gov/omb/egov/documents/TIC_ImplementationPlanningGuidance.pdf
- That reduction in connections is reportedly motivated by a desire to get the number of network connections down to a smaller, more easily managed and monitored number of connections, while also striving to contain costs.
- There are many details about the TIC program which have not been publicly disclosed, but we also know that those smaller number of connections will be monitored by US-CERT using "Einstein" boxes (additional monitoring may also take place)

A Sense of Where "Einstein"'s Currently At

- "Robert Jamison, undersecretary for national protection and programs at the Department of Homeland Security and Charbo's boss, told the hearing that Einstein, currently deployed at DHS and a handful of other government agencies, was being re-vamped for its rollout across all the federal networks.

  "Einstein currently collects information about traffic flows, and network managers analyze it daily, looking at where on the Internet so-called data packets that make up Web traffic are headed. But Jamison told the hearing that the new version, for which officials have requested an additional $115 million this year, will collect network traffic flow data in real time and also analyze the content of some communications, looking for malicious code, for example in e-mail attachments."

  (http://www.upi.com/International_Security/Emerging_Threats/Analysis/2008/03/03/analysis_einstein_and_us_cybersecurity/2343

But Coming Back to TIC, 50 Connections Aren't All That Many…

- When I first heard about the TIC program's target of just fifty connections, my first thought was, "Oh, that means they're going to have one connection per state, eh?"

- But as you think about things a bit more, though, I don't think that's how things will work out for a variety of reasons:

  -- Some states don't have established Internet exchange points
  -- Other states may not have substantial federal Internet traffic
  -- Some connections will be needed overseas as well as in the US, reducing the number of connections available for domestic use
  -- Some large states (such as CA, FL, NY or TX for example) may justify "expenditure" of multiple connections, also "using up" some of the limited 50 connection "budget"

A Need for Extensive Private Backhaul Networks

- So how will federal connectivity requirements be handled in areas which don't have a local TIC node? I believe there will need to be an extensive private backhaul network, much like the regional optical networks which are now a routine part of Internet2's own network architecture.

- When deploying that infrastructure, I believe that the feds will need to exercise some care, since long non-redundant backhaul circuits can potentially introduce single points of failure.

- Maintaining long dedicated backhaul circuits may also have cost impacts unless remote agencies tunnel traffic back to TIC nodes over the Internet via VPNs, or backhaul is provided free of charge.

A Small, Finite, Discoverable (and *DDoS-able!*) Number of Connections

- If we know that the US Government is connected to the Internet via only a small, finite, number of connections, miscreants (or hostile foreign entities for that matter) will probably make it a priority to identify the location and capacity of those connections.
- While we don't know the capacity of those connections, let's assume that they are probably going to be in the OC12 (622Mbps), gigabit, OC48 (2.4Gbps), or ten gigabit range (and oh, I suppose that there may be some faster connections but probably not many if any above 10Gbps, and those may be offset by some 100Mbps's)
- Assuming I'm correct about that connection speed distribution, an entity able to generate a suitably targeted and distributed DDoS in the range of 31.1 Gbps (50x622Mbps) to 500Gbps (50x10Gbps) might be able to take the entire federal government connection to the Internet down. That strikes me as a pretty bad thing event.

Speaking of DDoS, As The Number of Network Interfaces Goes Down, Localizing Spoofed Traffic Becomes Harder

- Another paradox of consolidating connections: because the government will be reducing the number of connections to the Internet down to just fifty, that action may have the paradoxical effect of making it substantially harder to localize the true source of spoofed attack traffic since all traffic in a single location may be coming in over a single interface.
- It may even obscure the true TARGET of an attack, since a given connection might be shared between a high value, high profile target and one or more obscure agencies. There will be a temptation to assume that cyber attackers will always be going after the high value, high profile target, but that may not necessarily always be the case.

Reservations About TIC

- After looking at the information that's available about the Trusted Internet Connection Program, I must admit, I remain concerned.

- I certainly understand and appreciate the desire to get the number of Internet connections down to a more manageable number, but I worry that doing so will make us more vulnerable to cyber attack, not less. Sometimes architectural complexity and "extra" connections may actually be your friend!

- I hope that before we go too far down this radical new direction, folks give the TIC concept a second, closer, look.

Thanks for the Chance to Talk Today

- Are there any questions?