

Internet2 Security Update: Some Excerpts From the 2nd Data Driven Collaborative Security Workshop and Some Timely Strategic Security Area You Should Be Thinking About

Joe St Sauver, Ph.D.

Internet2 Nationwide Security Programs Manager
(joe@uoregon.edu or joe@internet2.edu)

Internet2 Fall Members Meeting, Atlanta GA
Thursday, November 4th, 2010 10:30-11:45AM
Grand Ballroom III/IV

<http://pages.uoregon.edu/joe/sec-update-fall2010-mm/>

Introduction: We're All Busy, But...

- Many of us may all be preoccupied with major broadband stimulus-related infrastructure projects, but security issues continue to demand the community's attention:
 - Unpatched or incompletely patched systems and applications continue to get cracked, potentially resulting in breaches of personally identifiable information (PII),
 - Malware continues to outpace signature-based antivirus software, resulting in a steady supply of botnet hosts
 - Satisfying increasingly demanding compliance-related security requirements can also be daunting and time consuming.
- Given those pressures, it is pretty easy to fall into “reactive mode,” spending all our security related cycles just “fighting fires” and “trying to satisfy the auditors.”

We Need To Look For “Leverage Opportunities”

- The only way we can scale up to those day-to-day challenges is by looking for “leverage opportunities.”
- Think of “leverage opportunities” as times when we might be able to use technology to simultaneously fight the fires that break out (because we must continue to do that), while ALSO making substantive progress against vulnerabilities that are being actively targeted for exploitation.
- Doing this requires *Data, Analysis, Collaboration* and *Action*, the touchstones of the “Data Driven Collaborative Security” approach that we’ve been highlighting in the last couple Internet2 Data Driven Collaborative Security Workshops for High Performance Networks (DDCSW and DDCSW2).

The 2nd Internet2 Data Driven Collaborative Security Workshop

- Speaking of DDCSW2, we held the 2nd invitational Internet2 Data Driven Collaborative Security Workshop (“DDCSW2”) this summer from August 17th-18th, 2010 at the Knight Executive Education and Conference Center on the Washington University in St Louis campus. Thank you for sharing that facility with us!
- As was the case for the first DDCSW held at the University of Maryland Baltimore County, DDCSW2 included a mix of academic, corporate, non-profit and law-enforcement / government folks.
- Even if you did attend DDCSW2, unlike many closed cyber security meetings, you can check out some excellent presentations from that meeting online at security.internet2.edu/ddcsw2/

Three Topics From DDCSW2

- As a bit of a “teaser” to get you interested in learning more about DDCSW2, I wanted to highlight three immediately relevant tactical cyber security issues which were raised during that meeting, before covering some strategic cyber security issues.
- Three tactical cyber security issues from DDCSW2 included:
 - 1) Updates for PC Software OTHER THAN MS Windows, MS Office, Internet Explorer, etc.
 - 2) RPZ: DNS “Response Policy Zones,” and
 - 3) Dragon Research Group and DRG “Pods” (including the DRG ssh project)

1. Updates for PC Software **OTHER THAN Windows Itself, Office, Internet Explorer, etc.**

- Microsoft has done a great job of improving their software's code quality and helping users to keep Microsoft's own software (MS Windows, MS Office, Internet Explorer, etc) up-to-date.
- However, that's not the **only** software you've got on your PC.
- Most people also have third party applications installed such as:
 - Acrobat or Acrobat Reader
 - Flash Player
 - Third party browsers such as Firefox or Opera
 - Media helper applications such as QuickTime
 - Music players such as iTunes
 - Java
 - etc.
- Unfortunately you and your users may not be keeping up when it comes to keeping all those other applications patched up-to-date.

“The Proof Of The Pudding Is In The Eating”


- If you have a personally-owned Windows PC, try an experiment.
- Download Secunia PSI (free for personal use) from <http://secunia.com/products/> and run it on your personally owned system. (Secunia CSI is the institutional analogue of Secunia PSI)
- When you run PSI I would be extremely surprised if that tool doesn't find *at least* one third party application that is either end-of-life or less than fully patched on any given system you may happen to check.
- The problem of unpatched third party applications is endemic, and it IS getting noticed (and targeted!) by cyber attackers.

Consider PDF Attacks Last Year...

Rogue PDFs account for 80% of all exploits, says researcher

http://www.computerworld.com/s/article/print/9157438/Rogue_PDFs_account_for_80_of_all_exploits_says_researcher?

COMPUTERWORLD

 Print Article  Close Window

Rogue PDFs account for 80% of all exploits, says researcher

Adobe's Reader wins 2009 hacker honors by a landslide, says ScanSafe

Gregg Keizer

February 16, 2010 ([Computerworld](#))

Just hours before Adobe is slated to deliver the latest patches for its popular PDF viewer, a security firm announced that by its counting, malicious Reader documents made up 80% of all exploits at the end of 2009.

According to ScanSafe of San Bruno, Calif., vulnerabilities in Adobe's Reader and Acrobat applications were the most frequently targeted of any software during 2009, with hackers' PDF exploits growing throughout the year.

In the first quarter of 2009, malicious PDF files made up 56% of all exploits tracked by ScanSafe. That figure climbed above 60% in the second quarter, over 70% in the third and finished at 80% in the fourth quarter.

"PDF exploits are usually the first ones attempted by attackers," said Mary Landesman, a ScanSafe senior security researcher, referring to the multi-exploit hammering that hackers typically give visitors to malicious Web sites. "Attackers are choosing PDFs for a reason. It's not random. They're establishing a preference for Reader exploits."

Or Consider Java Today...

'Unprecedented wave' of Java exploits hits users, says Microsoft

://www.computerworld.com/s/article/print/9191640/_Unprecedented_wave_of_Java_exploits_hits_users_says_Microsoft

COMPUTERWORLD  Print Article  Close Window

'Unprecedented wave' of Java exploits hits users, says Microsoft

Java makes a tempting target, adds Symantec

Gregg Keizer

October 18, 2010 ([Computerworld](#))

Microsoft said Monday that an "unprecedented wave" of attacks are exploiting vulnerabilities in Oracle's Java software.

According to a manager at Microsoft's Malware Protection Center (MMPC), attempts to exploit Java bugs have skyrocketed in the past nine months, climbing from less than half a million in the first quarter of 2010 to more than 6 million in the third quarter.

"Some of our exploit 'malware' families were telling a scary story ... an unprecedented wave of Java exploitation," said Holly Stewart, a senior program manager at the MMPC, in a post to the [team's blog](#) Monday.

Stefan Frei from Secunia at DDCSW2

- Given the timeliness of this issue, we were delighted when Stefan Frei, Research Analyst Director at Secunia, was able to come to DDCSW2 to talk about their experience with Secunia PSI on **2.6 million PCs**. See security.internet2.edu/ddcsw2/docs/sfrei.pdf
- Some highlights:
 - half of all users have >66 programs from >22 vendors (dang!)
 - The top-50 most common programs include 26 from Microsoft, plus 24 3rd party programs from 14 different vendors (with 14 different update mechanisms!)
 - Eight programs from three vendors all have a > 80% user share
 - All programs in the top-50 portfolio have a $\geq 24\%$ user share
 - In the 1st half of 2010, 3rd party programs in the top-50 portfolio had 275 vulnerabilities, 4.4X more than MS programs
 - One exploitable vulnerability is all you need to own a PC...¹⁰

Sample Secunia PSI Run Output

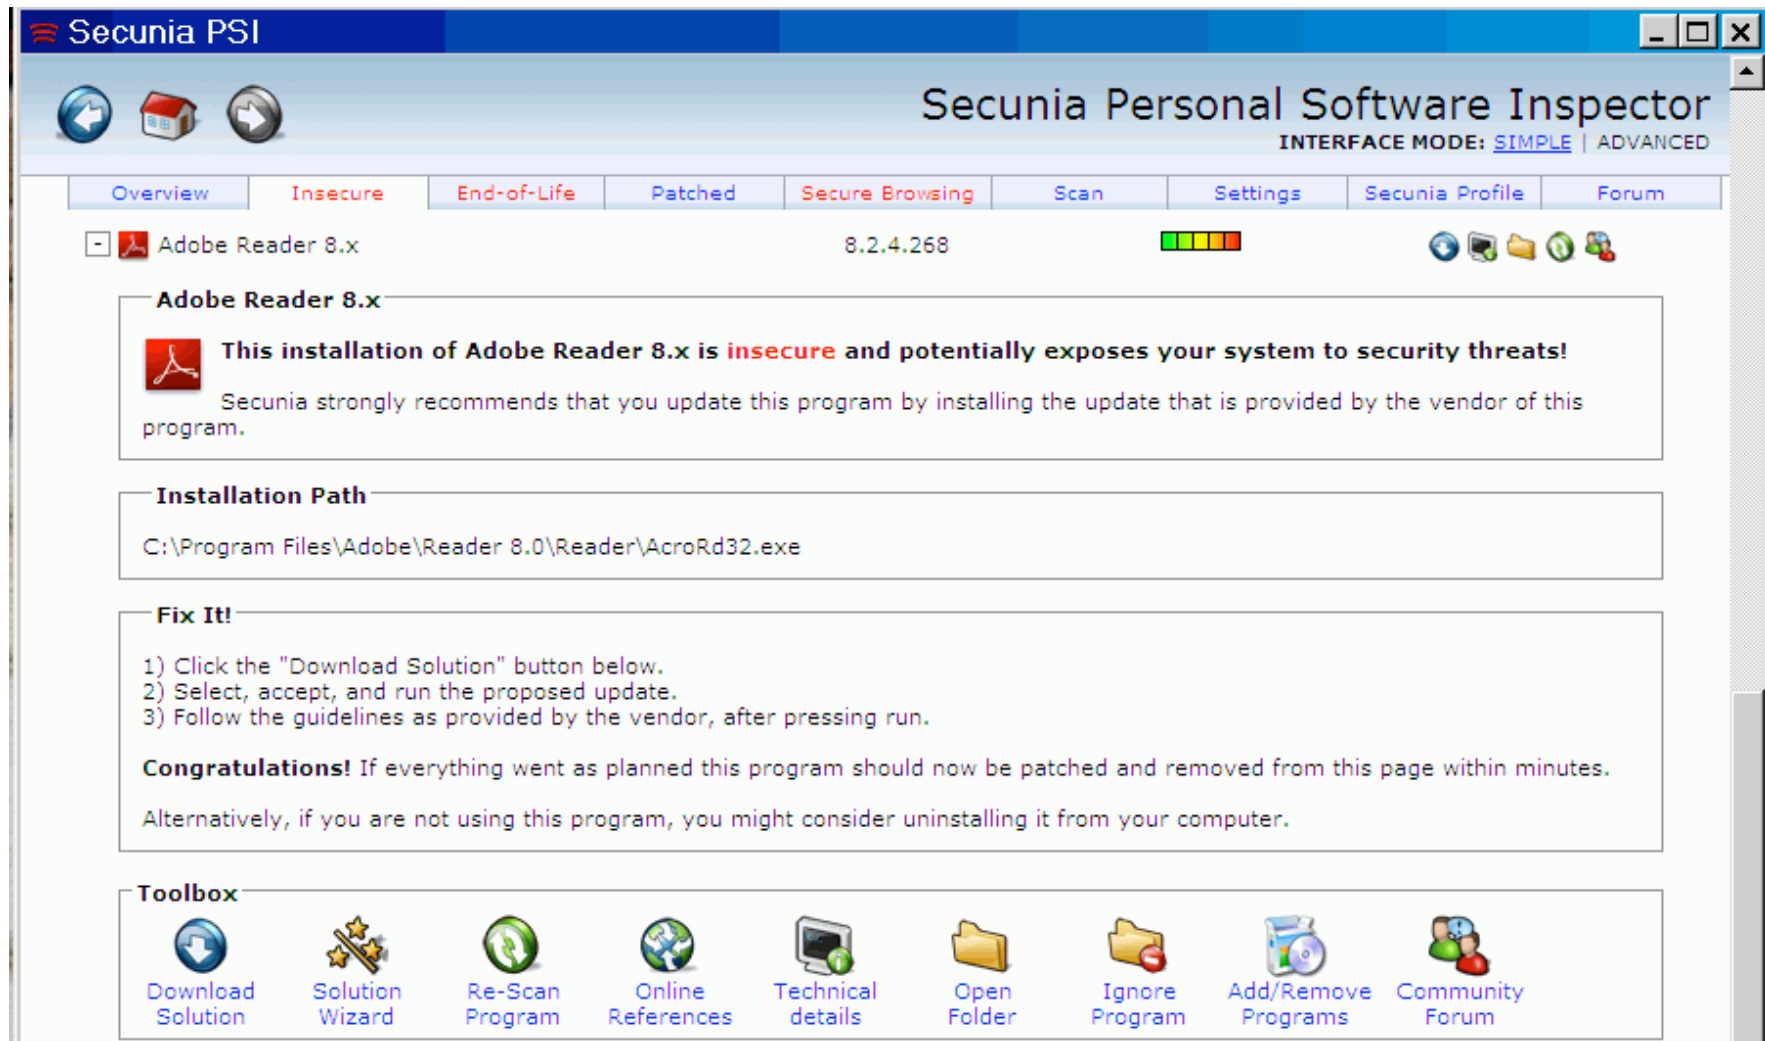
The screenshot displays the Secunia PSI application window. The title bar reads "Secunia PSI". Below the title bar is a navigation bar with icons for back, home, and forward, and the text "Secunia Personal Software Inspector" and "INTERFACE MODE: SIMPLE | ADVANCED". The main content area has a tabbed interface with "Overview", "Insecure", "End-of-Life", "Patched", "Secure Browsing", "Scan", "Settings", "Secunia Profile", and "Forum". The "Insecure" tab is selected, showing a list of insecure programs. A mouse cursor is hovering over the "Apple QuickTime 7.x" entry. Below the list is a link to "Help us improve our service to you: Program missing? Suggest it here!". The status bar at the bottom shows "Secunia's Privacy Statement", "Secunia PSI Status: Ready to scan.", and "Secunia PSI v1.5.0.2".

Insecure Programs [?]	Version Detected [?]	Threat Rating [?]	Direct [?]
+ Adobe Flash Player 10.x	10.0.45.2 (NPAPI)		
+ Adobe Reader 8.x	8.2.4.268		
+ Adobe Shockwave Player 11.x	11.5.2.602 (NPAPI)		
+ Adobe Shockwave Player 11.x	11.5.2.602 (ActiveX)		
+ Apple QuickTime 7.x	7.55.90.70		
+ GnuPG / gpg 1.4.x	1.4.7.6410		
+ ImageMagick 6.x	6.3.1		
+ Mozilla Firefox 3.6.x	3.6.10		
+ SSH Secure Shell for Workstations 3.x	3.2.5		
+ Sun Java JRE 1.6.x / 6.x	6.0.200.2		

Help us improve our service to you:
[Program missing? Suggest it here!](#)

Secunia's Privacy Statement | Secunia PSI Status: Ready to scan. | Secunia PSI v1.5.0.2

Drilling Down on One of Those Programs



IMPORTANT: Don't forget to check and fix ALL **RED TABS**! ¹²

Interested in Using PSI/CSI At Your Site?



Collaboration

- Secunia PSI is free for private use
 - PSI 2.0 will include auto update for 3rd party programs (released this autumn)
- Special PSI partnership for US universities
 - free PSI redistribution for campus and student machines
 - free PSI/CSI integration (university can monitor their PSI users population)
 - Example: <http://informationsecurity.iu.edu/tools/>
- Research
 - on you own PSI data
 - in collaboration with Secunia on all data

BTW, Change Is Coming for Some 3rd Party Apps

- I think we're at something of a cusp when it comes to some third party software, at least when it comes to some vendors.
- For example, as of Mac OS X 10.6 update 3, the version of Java that is ported by Apple and ships with OS X will be "deprecated." (see <http://tinyurl.com/java-deprecated>). While it may be possible for a fully open source version of Java to be developed for OS X, it may be tricky to get the same seamless integration that the vendor supported version of Java currently provides. Apps using Java are also reportedly going to be rejected by the Apple iPhone App Store.
- Finally, it also appears that Apple will no longer be pre-installing Adobe Flash Player on Macs (although users can still download and install it themselves).
- Quoting Bob Dylan, "You better start swimmin' / Or you'll sink like a stone / For the times, they are a-changin'."

2. Another Excerpt From DDCSW2: RPZ

- RPZ stands for “DNS Response Policy Zones” and Eric Ziegast of ISC was good enough to come to DDCSW2 and do two talk for us, with one of them covering RPZ. See <http://security.internet2.edu/ddcsw2/docs/Ziegast-rpz.pdf>
- RPZ stems from a seminal July 30th, 2010 article by Paul Vixie of ISC in CircleID entitled, “Taking Back the DNS,” see http://www.circleid.com/posts/20100728_taking_back_the_dns/
- In a nutshell, Vixie’s insight was that it’s crazy for sites like ours to help the bad guys to commit their cyber crimes by providing trustworthy and reliable DNS service for evil purposes.
- For example, our name servers should NOT be docilely and dutifully resolving domain names known to lead to malware, thereby helping the bad guys to efficiently infect our systems.
- Think of RPZ as new real time “block listing” for DNS.

Some RPZ Pragmatic Details

- RPZ is currently available as a patch for BIND (see the links from Vixie's CircleID article).
- ISC is NOT providing a data feed for RPZ, just the protocol spec and a reference implementation (patch) for BIND.
- You could build your own RPZ zone, or select one supplied by a third party.
- If you do implement RPZ for typical users, you may want to also make sure you offer an unfiltered recursive resolver for any campus malware researchers or security researchers (or at least do not block their ability to run their own unfiltered recursive resolver, or their ability to reach Google's intentionally open recursive resolvers at 8.8.8.8 and 8.8.4.4).
- Because of the dismal status of malware protection right now, I think that we'll be hearing a lot about RPZ in the future.

3. The Dragon Research Group and DRG Pods (Including the DRG ssh Project)

- Many of you will already be familiar with Team Cymru (see <http://www.team-cymru.org/>) and the excellent work that Rob Thomas and his team do in furthering Internet security.
- You may *not* be as familiar with Dragon Research Group, the international all-volunteer research group offshoot of Team Cymru (even though they *are* available as a link from the top bar on the primary Team Cymru web site).
- We were fortunate to have Paul Tatarsky, Seth Hall and John Kristoff provide a briefing on the DRG for DDCSW2, see <http://security.internet2.edu/ddcsw2/docs/tatarsky.pdf>
- For today's update, we'll just highlight two things related to the that talk: volunteering to run a DRG "pod," and an example of one project enabled by DRG pod data, the DRG ssh project.

DRG “Pods”

- Dragon Research Group makes available a customized Linux Live CD distribution that securely converts a system (or virtual machine) into a DRG data collection endpoint (or “pod”).
- A full description of the distribution and how you can sign up to participate is at www.dragonresearchgroup.org/drg-distro.html
- Because network activity policies vary from site to site, the DRG distribution intentionally provides substantial flexibility. Thus, for example, if your site will only permit passive measurement activities, the pod can be configured to carefully support that policy, while if your site allows active measurements, that more liberal framework can also be accommodated.
- All DRG pod locations are confidential.
- A nice example of the sort of work that the DRG pods can enable is the DRG ssh project, which we’ll describe next.

The DRG ssh Project

- ssh (secure shell, e.g., an “encrypted version of telnet”) is the preferred way that most security-conscious individuals remotely login to Unix boxes and other systems. On many hardened systems, sshd may be the only network service that’s accessible.
- Because sshd may be the only service that’s open, it gets a lot of attention from cyber criminals who scan the Internet looking for vulnerable hosts. Anyone running sshd is all too familiar with failed ssh login attempts from random sources in their syslogs.
- Wouldn’t it be nice if you could see a list of all the IP addresses that have recently been seen ssh scanning? Wouldn’t it be particularly nice to know if one of those actively scanning hosts is actually a (likely compromised) system on your campus?
- You can read more about the DRG ssh project at <http://www.dragonresearchgroup.org/insight/>


Part of a Recent DRG ssh Password Auth Report

Mozilla Firefox					
http://www.dragonresearchgroup.org/insight/sshpwauth.txt					
137	ASGARR GARR Italian academic a	131.114.27.122	2010-10-24 02:32:27	sshpwauth	
766	REDIRIS RedIRIS Autonomous Sys	161.111.114.58	2010-10-24 11:55:45	sshpwauth	
1239	SPRINTLINK - Sprint	204.215.65.203	2010-10-21 15:55:06	sshpwauth	
1241	FORTHNET-GR FORTHnet	193.92.98.180	2010-10-23 23:44:13	sshpwauth	
1267	ASN-INFOSTRADA Infostrada S.p.	151.3.164.146	2010-10-22 07:27:10	sshpwauth	
1736	MU-AS - Marquette University	134.48.6.45	2010-10-22 02:25:01	sshpwauth	
1916	Rede Nacional de Ensino e Pesq	200.143.196.24	2010-10-24 20:40:21	sshpwauth	
2514	INFOSPHERE NTT PC Communicatio	219.117.232.106	2010-10-21 17:18:28	sshpwauth	
2527	SO-NET So-net Entertainment Co	121.2.67.42	2010-10-19 16:28:16	sshpwauth	
2764	AAPT AAPT Limited	203.63.104.225	2010-10-22 14:50:48	sshpwauth	
3202	St. Andrews University (SuperJ	193.194.64.243	2010-10-20 09:15:26	sshpwauth	
3208	ARN	193.194.64.243	2010-10-20 09:15:26	sshpwauth	
3215	AS3215 France Telecom - Orange	80.14.189.178	2010-10-23 13:02:43	sshpwauth	
3261	DIPT-AS Donbass Network	92.242.121.254	2010-10-20 17:26:02	sshpwauth	
3269	ASN-IBSNAZ Telecom Italia S.p.	94.95.129.131	2010-10-23 20:45:13	sshpwauth	
3301	TELIA-AS SWEDEN TeliaNet Swede	90.224.37.166	2010-10-21 01:57:16	sshpwauth	
3352	TELEFONICA-DATA-ESPANA TELEFON	213.99.240.243	2010-10-24 10:52:28	sshpwauth	
3352	TELEFONICA-DATA-ESPANA TELEFON	88.26.231.136	2010-10-18 14:13:08	sshpwauth	
3462	HINET Data Communication Busin	220.133.136.182	2010-10-21 11:23:00	sshpwauth	
3462	HINET Data Communication Busin	60.250.203.250	2010-10-19 22:34:33	sshpwauth	
3549	GBLX Global Crossing Ltd.	201.234.213.82	2010-10-23 20:44:41	sshpwauth	
3549	GBLX Global Crossing Ltd.	208.51.7.247	2010-10-23 17:04:16	sshpwauth	
3758	ERX-SINGNET SingNet	203.126.53.110	2010-10-23 13:27:31	sshpwauth	
3786	LGDACOM LG DACOM Corporation	123.142.80.122	2010-10-20 15:00:16	sshpwauth	
3816	COLOMBIA TELECOMUNICACIONES S.	200.21.232.166	2010-10-24 01:57:31	sshpwauth	
3816	COLOMBIA TELECOMUNICACIONES S.	200.21.228.182	2010-10-24 10:45:09	sshpwauth	
3816	COLOMBIA TELECOMUNICACIONES S.	190.254.22.94	2010-10-23 21:44:29	sshpwauth	
4134	CHINANET-BACKBONE No.31,Jin-ro	121.14.195.176	2010-10-18 09:23:26	sshpwauth	
4134	CHINANET-BACKBONE No.31,Jin-ro	118.122.32.46	2010-10-22 04:27:47	sshpwauth	
4134	CHINANET-BACKBONE No.31,Jin-ro	121.14.195.83	2010-10-24 08:54:28	sshpwauth	
4134	CHINANET-BACKBONE No.31,Jin-ro	221.232.155.6	2010-10-22 20:34:58	sshpwauth	

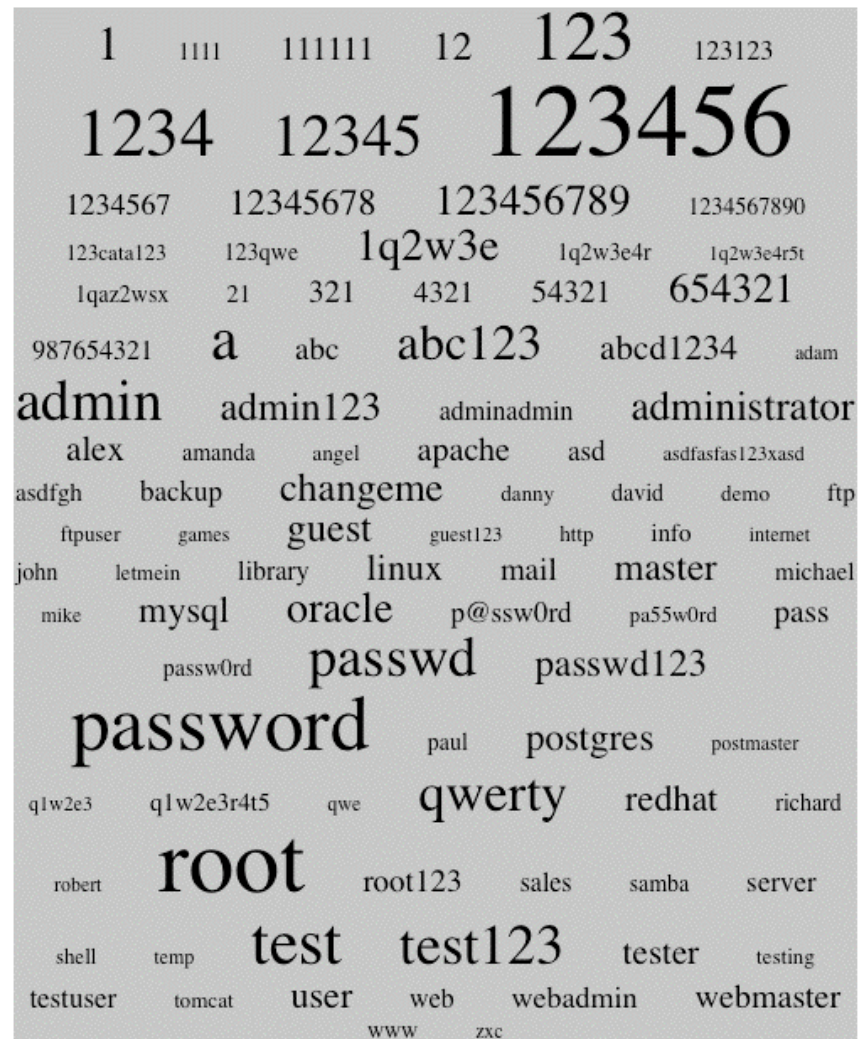
DRG ssh Username/Password Tag Clouds



most popular usernames



most popular passwords



An Aside on Ssh Scanning Tools

- At least some of the hosts that are engaged in ssh scanning/brute forcing are likely infested with the dd_ssh brute forcing script. For more information about this attack tool, see:
<http://isc.sans.edu/diary.html?storyid=9370>
- Metasploit Framework 3.4.0 (released May 18th, 2010) also now includes “strong support” for brute forcing network protocols, including support for brute forcing ssh, see
http://blog.metasploit.com/2010_05_01_archive.html
- These sort of brute forcing tools mean that brute forcing attacks are likely here to stay...
- Many sites may want to consider deploying anti-brute forcing scripts as part of their system configuration. One such tool is fail2ban, see http://www.fail2ban.org/wiki/index.php/Main_Page however there are many others you might also try.

DRG Will Be Doing More Cool Projects

- So as cool as the preceding ssh analyses are, they are really just an example, the tip of the proverbial iceberg, if you will.
- With your help, many further interesting projects may become possible.
- We'd encourage you to consider participating in the DRG's activities by hosting a pod at your site.
- If nothing else, you'll at least want to keep an eye on their ssh scanner/ssh brute forcer report to make sure your ASN or ASNs used by your colleagues, don't show up as a source of abusive ssh brute force traffic!

Should We Continue Having DDCSW Meetings?

- So now you know a little about three of the great security-related presentations that were shared at the last DDCSW.
- An open question to those of you in the Internet2 community:
“Should we have further DDCSW events in the future?”
- We think the quality of the material presented at both DDCSWs was outstanding, but we recognize that everyone in the security community is very busy. Some might go so far as to say that the biggest “gift” we could give the security community would be to REFRAIN from offering yet another security meeting competing for limited time and travel resources.
- **So should we consider merging DDCSW with another meeting? Which one? Should we drop DDCSW entirely?** We’d appreciate your feedback! (please send it to joe@internet2.edu)
- If we do decide to hold another DDCSW, would you be interested in attending and presenting at it? Or maybe hosting it?

That's It For Our Brief “Taste of DDCSW” and Overview of A Few “Tactical” Security Topics

- Now let's move on and talk a little about some timely “big picture” or “strategic” security topics.

Three Strategic Security Topics

- While there are many important strategic security topics we could talk about today, there are three strategic security challenges which have largely received short shrift at most of our sites:

4) IPv4 Exhaustion and IPv6 Deployment

5) Security of the Domain Name System and DNSSEC, and

6) The Security of Mobile Internet Devices

- Let's briefly talk about each of those topics.

4. IPv4 Runout and IPv6: IPv4 Runout Is Nigh

- **Only 5% of global IPv4 address space remains unallocated.**
- The last large unallocated IPv4 netblocks (“/8’s”, each 1/256 of the total IPv4 address space) will be allocated by IANA on or about 4 June 2011.
- The regional Internet registries (such as ARIN) will begin to exhaust their last IPv4 allocations on or about 27 January 2012.
- Neither of those dates are very far from now:

4 Nov 2010 --> 4 Jun 2011: **212 days**

4 Nov 2010 --> 27 Jan 2012: **1 year, 2 months, 23 days**

Preparing for Imminent IPv4 Runout

- Between now and then, **you should be doing three things:**
 - 1) If you have legacy IPv4 address space, review your records documenting that allocation (if you have any and if you can find them) and decide if you're going to sign the **ARIN Legacy Registration Services Agreement**. (See <https://www.arin.net/resources/legacy/>)
 - 2) **If you have a legitimate need for additional IPv4 address space for any pending projects, request that space NOW.**
If you wait six months to make that request, it may be too late.
(Note: I am NOT suggesting that you request space you don't legitimately need – PLEASE be reasonable and responsible)
 - 3) **Everyone should be proceeding with deployment of IPv6 on the networks and systems they operate.**

Most Universities Have NOT Deployed IPv6

- **Only a few universities have deployed IPv6 both throughout their infrastructure AND on all their public-facing servers.**
- See “IPv6 Status Survey,” http://www.mrp.net/IPv6_Survey.html

If your site isn't listed, you can check it using the form that's at:
<http://www.mrp.net/cgi-bin/ipv6-status.cgi>

- Note: this test only checks public services for IPv6-accessibility.

You should also check to see if your institution has enabled IPv6 throughout your local area network for use by end user workstations.

IPv6 Status Survey					
http://www.mrp.net/IPv6_Survey.html					
Internet2 Members					
Organisation (domain)	Web	Mail	DNS	NTP	XMPP
American University (american.edu)	FAIL	FAIL	0/2 0/2	FAIL	
Arizona State University (asu.edu)	FAIL	FAIL	0/4 0/4		FAIL
Arkansas State University (astate.edu)	FAIL	FAIL	0/1 0/2		FAIL
Auburn University (auburn.edu)	FAIL	FAIL	0/3 0/3	FAIL	
Baylor College of Medicine (bcm.edu)	FAIL	FAIL	0/4 0/4		
Baylor University (baylor.edu)	FAIL	FAIL	0/2 0/4	FAIL	
Binghamton University (binghamton.edu)	FAIL	FAIL (G)	0/3 0/3		
Boston College (bc.edu)	FAIL	FAIL	0/2 0/2	FAIL	
Boston University (bu.edu)	FAIL	FAIL	0/3 0/3	FAIL	
Bowling Green State University (bgsu.edu)	FAIL	FAIL	0/2 0/4	FAIL	
Bradley University (bradley.edu)	FAIL	FAIL	0/2 2/4		
Brandeis University (brandeis.edu)	FAIL	FAIL	0/2 0/9	FAIL	FAIL
Brigham Young University (byu.edu)	FAIL	FAIL	0/8 0/8		
Brown University (brown.edu)	FAIL	FAIL	0/2 4/8	FAIL	
California Institute of Technology (caltech.edu)	FAIL	FAIL	0/3 0/6	FAIL	

“We’ve *Intentionally* Decided to NOT Do IPv6”

- Some universities may be aware of IPv4 runout, AND may have made an intentional decision to NOT deploy native IPv6 for their users. You may even be from one of those universities.
- If so, I would urge you to reconsider that decision.
- If you do NOT deploy native IPv6, your users will (intentionally or inadvertently) end up transparently accessing IPv6 content via a variety of non-native transition mechanisms such as Teredo, 6to4, ad hoc manually configured tunnels, etc., whether you support native IPv6 or not. This will ultimately be a mess, and far less secure than just “biting the bullet” and doing native IPv6. See “IPv6 and the Security of Your Network and Systems,” pages.uoregon.edu/joe/i2mm-spring2009/i2mm-spring2009.pdf

Large Scale Network Address Translation

- If you do find yourself talking to those who aren't planning to add IPv6, and you ask them “How will you scale IPv4 addressing post-IPv4 runout?” the most common answer you'll hear is that they plan to do large scale NAT (you may also hear this called “carrier grade NAT,” although most large scale NAT solutions are not really “carrier grade”).
- Sites that try large scale NAT will be sharing a single public IPv4 address across dozens or sometimes even hundreds of users.
- Large scale NAT will pose many challenges, and after you think about them a little, we hope that you will reconsider your decision to go down that road. For example...

Incident Handling in a Large Scale NAT World

- Incident handlers and security staff know that abuse complaints involving dynamic addresses need both the address of the problematic host, AND the timestamp/time zone when the incident was observed in order to be actionable.
- As large scale NAT becomes more widely deployed, actionable abuse reports will now need to have THREE items: the address of the problematic host, the timestamp/time zone when the incident was observed, AND the source port number.
- Unfortunately, many abuse records do not currently include source port info. For example, if you look at Received: headers in mail messages, you will NOT see source port information listed. Many other sources of backtracking information are similarly bereft.

Loss of Transparency (and Loss of Innovation, and Loss of Throughput, and...)

- Large scale NAT may work adequately well for users with simple mainstream needs (such as browsing the web, or sending email via a third party web email service), but those sort of applications should NOT be the epitome of “advanced applications” or “high performance applications” in our community!
- Innovative advanced applications and high performance data transfers almost always work better when Internet connected hosts have globally routed unique IP addresses.
- For that matter, even some pretty basic applications, such as video conferencing, often ONLY work if you have a public address.

“There Are A Million Different Really Good Reasons Why We Just *Can’t* Deploy IPv6!”

- There may be. Unfortunately, you really don’t have any good alternative (as Iljitsch van Beijnum wrote in Ars Technica a month or so ago, “There is No Plan B: Why The IPv4-to-IPv6 Transition Will Be Ugly,” see arstechnica.com/business/news/2010/09/there-is-no-plan-b-why-the-ipv4-to-ipv6-transition-will-be-ugly.ars)
- **The time has come to get IPv6 deployed on your campus, and on your servers, and on your regional networks.**


5. Security of the Domain Name System (DNS) and DNSSEC

- Pretty much everything on the Internet relies on the ability of users to safely refer to sites by symbolic names (such as `www.internet2.edu`) rather than IP addresses (such as `207.75.165.151`), trusting DNS to do that translation for them.
- If that translation process is untrustworthy, instead of going where you wanted to go, you might end up being taken to a site that will drop malware on your system, or you might be diverted from your bank or brokerage to a fake financial site run by some offshore cracker/hacker.
- It is absolutely critical that DNS be trustworthy.
- DNSSEC, a system of cryptographic signatures that can help insure that DNS results haven't been tampered with, can help secure DNS results -- IF it gets used.

Two DNSSEC Tasks: Signing and Checking

- For DNSSEC to work, two things need to happen:
 - sites need to cryptographically sign their own DNS records
 - other sites need to check, or verify, that the DNSSEC-signed results they receive are valid
- Many sites have held off signing their site's DNS records because for a long time the DNS root ("dot") and the EDU top level domain weren't signed. That's no longer a problem: both have now been signed.
- At the same time, many recursive resolvers haven't bothered to check DNSSEC signatures because "no one" has bothered to sign their zones.
- DNSSEC thus formerly epitomized the classic Internet "chicken and egg" deployment problem.

Nonetheless, Deployment IS Beginning To Happen!



2nd Level .edu's Which ARE Signed (10/12/10)

- berkeley.edu
- cmu.edu
- desales.edu
- example.edu
- fhsu.edu
- indiana.edu
- internet2.edu
- iu.edu
- iub.edu
- iupui.edu
- k-state.edu
- ksu.edu
- lsu.edu
- merit.edu
- monmouth.edu
- penn.edu
- psc.edu
- suu.edu
- ucaid.edu
- upenn.edu
- weber.edu
- What about YOUR school???

Data from:


<http://secspider.cs.ucla.edu/> ³⁹

Some Universities Are Now Validating DNSSEC Signatures, Too


- For example, the University of Oregon is now verifying DNSSEC signatures on its production recursive resolvers, and this has generally been going just fine.
- If you need a simple test to see whether your current recursive resolvers are verifying DNSSEC signatures, try the (somewhat irreverent but quite straightforward) “thumbs up”/“eyes down” DNSSEC validation tester that’s available at:

<http://test.dnssec-or-not.org/>


For Example...



Yes, you are using DNSSEC



No, you are not using DNSSEC



Labs

Verifying DNSSEC Signatures Is Not Completely Without Risk

- In many ways, the most serious risk you face when validating DNSSEC signatures is that DNSSEC will “work as advertised.”
- That is, a domain may accidentally end up with invalid DNSSEC signatures for a variety of reasons, and once they’ve done that, their site will then (correctly) become inaccessible to those of us who are verifying DNSSEC signatures.
- Paradoxically, when that happens, the site will continue to work just fine for everyone who is NOT doing DNSSEC, and the DNSSEC problem may thus go unnoticed by the site.
- This may be an irritating experience for your users if a critical site ends up being inaccessible.
- <http://dnsviz.net> is a great resource for visualizing and debugging these sort of issues when they arise. If you want an intentionally broker domain to try testing, try using dnssec-failed.org

Not Ready to Jump In? Try Taking Baby Steps

- Maybe you can at least either:
 - sign your own domain or at least
 - begin to validate the signatures that others have added?

You don't need to immediately do both simultaneously!

- Maybe you can sign just part of your domain (such as your cs or engineering subdomains), or you can just try signing a couple of less-important institutional “test” domains...
- Maybe you can create additional “opt-in” validating resolvers, even if you don’t enable DNSSEC by default on your production recursive resolvers?

6. Security of Mobile Internet Devices

- There's a huge temptation to just focus on traditional networks, servers, desktop workstations and laptops, but there's been a real revolution quietly going on: **we're entering an age where mobile Internet devices are becoming virtually ubiquitous.**
- For example, the 2009 ECAR Study of Undergraduate Students and Information Technology (<http://www.educause.edu/ers0906>) reported that 51.2% of respondents owned an Internet capable handheld device, and another 11.8% indicated that they planned to purchase one in the next 12 months...
- What about faculty/staff? While mobile Internet devices and cell phones have formerly been treated as “listed property” by the IRS, Section 2043 of H.R. 5297 (the “Small Business Jobs Act of 2010”) was signed into law Sept 27, 2010, fixing that. Because of that recent change, expect to see a lot more institutionally owned faculty/staff mobile Internet devices soon...

Mobile Internet Devices Raise LOTS of Questions

- I've got a full 110 slide presentation discussing the security of mobile Internet devices that I recently gave as the closing session for the Northwest Academic Computing Consortium (NWACC) 2010 Network Security Workshop in Portland (see <http://pages.uoregon.edu/joe/nwacc-mobile-security/> (PDF or PPT formats)).
- Given our limited time together today, I'm obviously not going to be able to cover all that material.
- Recognizing how common mobile Internet devices have become, however, I do want to at least alert you to some of the security issues that you face from mobile devices, leaving you to see the full presentation for details and additional issues.
- To keep this simple, we'll largely focus on the Apple iPhone for the rest of this quick discussion.

A Few Mobile Internet Device Security Questions

- What type(s) of mobile Internet devices should we support? Blackberries? iPhones? Android devices? Does it matter?
- Is cellular wireless connectivity secure enough to protect PCI-DSS or HIPAA or FERPA data that may be transmitted?
- Should we centrally manage our mobile devices? If so, how?
- Is there PII on our users' mobile Internet devices? Do those devices have hardware “whole device” encryption to protect that data?
- What if one of these mobile devices get lost or stolen? Can we send the device a remote “wipe” or “kill” code?
- Do we need antivirus protection for mobile devices?
- What if users want to “jailbreak” their device? Is that okay?
- **And there are many more security questions, but few people are talking about these issues in our community. Why?**

Are We Seeing a Recapitulation Of the Old “Managed” vs. “Unmanaged” PC Wars?

- For a long time, way back in the “bad old days,” traditional IT management simply pretended that PCs didn’t exist.
- While they were “in denial,” people bought whatever PCs they wanted and “administered” them themselves. Sometimes that worked well, other times chaos reigned.
- Today's more closely managed “enterprise” model was the result of that anarchy. At some sites, standardized PC configurations are purchased and tightly locked down and are then centrally administered. While I’m not a fan of this paradigm, I recognize that it is increasingly common.
- Are we re-experiencing that same evolutionary process for mobile Internet devices?
- What might we be able to do if we did use a managed model?

An Example of One Simple Mobile Internet Device Policy Question: Device Passwords

- If a mobile Internet device is lost or stolen, a primary technical control preventing access to/use of the device is the device's password.
- Users hate passwords, but left to their own devices (so to speak), if they use one at all, they might just use a short (and easily overcome) one such as 1234
- You and your school might prefer that users use a longer and more complex password, particularly if that mobile Internet device has sensitive PII on it.
- You might even require the device to wipe itself if it detects that it is the target of an in-person password brute force attack.
- If the device is managed, you **can** require these things – but are your mobile Internet devices managed? Many aren't.

Other Potential Local iPhone “Policies” Include

- Adding or removing root certs
 - Configuring WiFi including trusted SSIDs, passwords, etc.
 - Configuring VPN settings and usage
 - Blocking installation of additional apps from the AppStore
 - Blocking Safari (e.g., blocking general web browsing)
 - Blocking use of the iPhone’s camera
 - Blocking screen captures
 - Blocking use of the iTunes Music Store
 - Blocking use of YouTube
 - Blocking explicit content
-
- Some of these settings may be less applicable or less important to higher ed folks than to corp/gov users.

Scalably Pushing Policies to the iPhone

- To configure policies such as those just mentioned on the iPhone, you can use configuration profiles created via the iPhone Configuration Utility (downloadable from <http://www.apple.com/support/iphone/enterprise/>)
- Those configuration files can be downloaded directly to an iPhone which is physically connected to a PC or Mac running iTunes -- but that's not a particularly scalable approach. The configuration files can also be emailed to your user's iPhones, or downloaded from the web per chapter two of the Apple Enterprise Deployment Guide.
- **While those configuration files need to be signed (and can be encrypted), there have been reports of flaws with the security of this process; see “iPhone PKI handling flaws” at cryptopath.wordpress.com/2010/01/**

What's The 'Big Deal' About Bad Config Files?

- If I can feed an iPhone user a bad config file and convince that user to actually install it, I can:
 - change their name servers (and if I can change their name servers, I can totally control where they go)
 - add my own root certs (allowing me to MITM their supposedly “secure” connections)
 - change email, WiFi or VPN settings, thereby allowing me to sniff their connections and credentials
 - conduct denial of service attacks against the user, including blocking their access to email or the web
- **These config files also can be made non-removable (except through wiping and restoring the device).**

We Need to Encourage “Healthy Paranoia”

- Because of the risks associated with bad config files, and because the config files be set up with attributes which increase the likelihood that users may accept and load a malicious configuration file, **iPhone users should be told to NEVER, EVER under any circumstances install a config file received by email or from a web site.**
- Of course, this sort of absolute prohibition potentially reduces your ability to scalably and securely push mobile Internet device security configurations to iPhones, but...
- This issue also underscores the importance of users routinely sync'ing/backing up their mobile devices so that if they have to wipe their device and restore it from scratch, they can do so without losing critical content.

What About Hardware Encryption?

- Another example of a common security control designed to protect PII from unauthorized access is hardware encryption.
- Many sites require “whole disk” encryption on all institutional devices containing PII.
- Some mobile Internet devices (such as earlier versions of the iPhone) didn’t offer hardware encryption; 3GS and 4G iPhones now do. **However, folks have demonstrated that at least for the 3Gs (and at least for some versions of iOS) was less-than-completely bullet proof; see for example Mr NerveGas (aka Jonathan Zdziarski’s) demo “Removing iPhone 3G[s] Passcode and Encryption,” www.youtube.com/watch?v=5wS3AMbXRLs**
- This may be a consideration if you are planning to use certain types of iPhones for PII or other sensitive data.

Remotely Zapping Compromised Mobile Devices

- Strong device passwords and hardware encryption are primary protections against PII getting compromised, but another potentially important option is being able to remotely wipe the hardware with a magic “kill code.” Both iPhones and BlackBerry devices support this option.
- Important notes:
 - If a device is taken off the air (e.g., the SIM card has been removed, or the device has been put into a electromagnetic isolation bag), a device kill code may not be able to be received and processed.
 - Some devices (including BlackBerries) acknowledge receipt and execution of the kill code, others may not.
 - Pre-3GS versions of the iPhone may take an hour per 8GB of storage to wipe (3GS’s wipe instantaneously).

Terminating Mobile Device-Equipped Workers

- A reviewer who looked at an earlier draft of some of these slides pointed out an interesting corner case for remote zapping:
 - Zap codes are usually transmitted via Exchange Active Sync when the mobile device connects to the site's Exchange Server, and the user's device authenticates
 - HR departments in many high tech companies will routinely kill network access and email accounts when an employee is being discharged to prevent “incidents”
 - If HR gets network access and email access killed before the zap code gets collected, the device may not be able to login (and get zapped), leaving the now ex-employee with the complete contents of the device See: <http://tinyurl.com/zap-then-fire>
- Of course, complete user level device backups may *also* exist₅₅.

Malware and A/V on the Non-Jailbroken iPhone

- Because earlier versions of the iPhone disallowed applications running in the background, it was difficult for traditional antivirus products to be successfully ported to the iPhone.
- To the best of my knowledge, your options for antivirus software on the iPhone are still “quite limited,” with no offering from traditional market leaders such as Symantec and McAfee at that time.
- On the other hand, since the iPhone used/uses a sandbox-and-cryptographically "signed app" model, it was hard for the iPhone to get infected.
- Will you allow users to jail break that security model?

And If There's NOT A/V For Mobile Devices...

- Some sites may “accidentally” adopt an “overly broad” policy when it comes to deploying antivirus, perhaps decreeing that **“If it can't run antivirus, it can't run.”**

As you might expect, I believe this is a mistake when there are compensating controls (such as use of a signed-app model in the case of the iPhone), or cases where the demand for A/V on a platform is so minimal there's not even a commercial A/V product available.

There are ways to avoid malware besides just running antivirus software!

- Remember “compensating controls!”

What About Jailbroken iPhones?

- Normally only Apple-approved applications run on the iPhone. However, some users have developed hacks (NOT blessed by Apple!) that will allow users to “break out of that jail” and run whatever applications they want.
- Jailbreaking your iPhone violates the license agreement and voids its warranty, but it is estimated that 5-10% of all iPhone users have done so.
- Q: “Is jailbreaking my iPhone legal?”
A: I am not a lawyer and this is not legal advice, but see:

**“EFF Wins New Legal Protections for Video Artists, Cell Phone Jailbreakers, and Unlockers,” July 26, 2010,
<http://www.eff.org/press/archives/2010/07/26>**

Jailbroken iPhones and Upgrades

- When a jail broken iPhones gets an OS upgrade, the jailbreak gets reversed and would typically need to be redone.
- This may cause some users of jail broken iPhones to be reluctant to apply upgrades (even upgrades with critical security patches!), until the newly released version of iOS also gets jailbroken.
- That's obviously a security issue and cause for concern.
- If you do successfully jailbreak your iPhone, your exposure to malware *will* increase.

Your Should Be Talking About These Issues

- If your user support and security staff aren't talking about these sort of issues at your site, you're likely not ready to address the security issues that will arise in conjunction with mobile devices.
- I'd urge you to review the full talk about mobile Internet device security that I mentioned on slide 45, and to initiate local conversations about mobile Internet device security as soon as you can reasonably do so.
- That's all I've got for you for you today for my part of the security update session. I assume we'll hold questions till the end of the session.
- Thanks!