INFORMATION SECURITY IN HEALTH CARE

S.C. HIMSS

10/30/2015

Mark Lachniet
mark.lachniet@cdw.com
(517-242-4874)

About The Speaker

- Information Security Solutions Manager, CDW (previously Security Engineer)
  - Penetration testing
  - Incident response & forensics
  - Regulatory compliance (HIPAA, PCI, NIST 800-53)
- Past employment:
  - K-12 Technology Director (Holt Schools)
  - Instructor, Masters in Information Assurance, Walsh College
  - Consulting at Analysts International, Promethean Security
- Industry certifications:
  - Certified Information Systems Security Professional (CISSP)
  - Certified Information Systems Auditor (CISA)
  - Licensed Private Investigator #3701-205679 (Michigan)

About The Speaker

- Assessment History
  - Hundreds of assessments over the last 15 years
  - Approximately 20% - 30% in Health Care or other industries that handle PHI
  - Penetration tests – i.e. “White Hat Hacking”
  - Policy and Procedure Gap Analysis
  - Incident Response
- Useful to compare Health Care to other industries
- What are the similarities and differences?
  - Perspective / viewpoint
  - Staffing
  - Technologies
  - Regulatory

Agenda

- Compare some HIMSS survey results with Verizon Breach Report results
- Discuss specific issues that I have seen while doing penetration tests and IT HIPAA assessments
  - Things that tend to be done well
  - Specific problem areas
  - Future concerns / looking forward

Comparing and Contrasting

- Results of 2015 HIMSS Cybersecurity Survey, dated October 15, 2015(1)
  - Versus….
- Results of the Verizon Data Breach Investigation Report for 2015 ← (The best resource I know of for good statistics)(2)
  - Versus….
- Personal experience!
- HIMSS: 279 Respondents
- Verizon: 70 Organizations (most of which are providing services for a large number of clients, so this probably represents hundreds or thousands of individual organizations)

(1) http://www.himss.org/2015-cybersecurity-survey

(2) http://www.verizonenterprise.com/DBIR/

Health Care – Disclosure Drive Results

- One trend that is immediately apparent in both reports is that events that would require disclosure are heavily represented:
  - Physical Security
  - Errors
  - Insiders
- HIMSS:

Insider Threat (54%)

- Negligent Insider (46%)
- Malicious Insider (12%)

External Threat (64%)

- Online Scam Artist (36%)
- Social Engineering (16%)
- Hacker (16%)

Health Care – Disclosure Drives Results

- Verizon Data Breach:

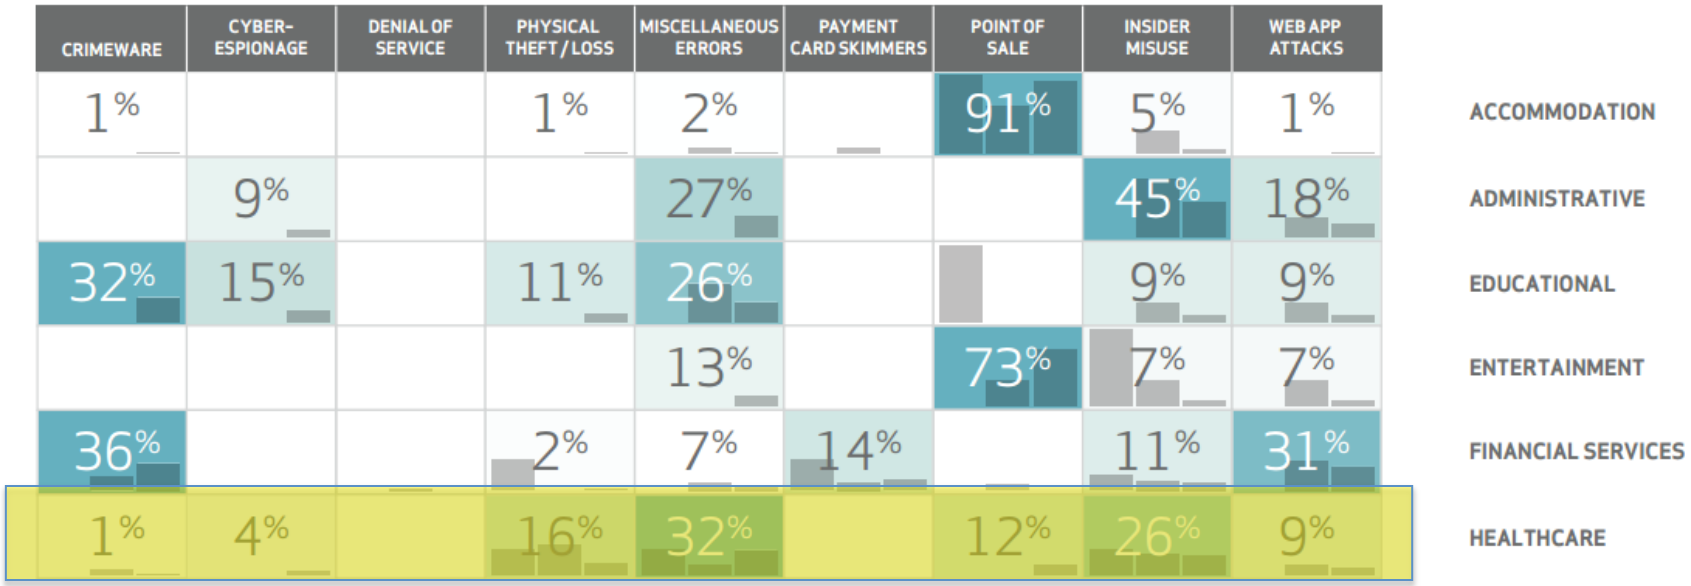

CRIMEWARE	CYBER-ESPIONAGE	DENIAL OF SERVICE	PHYSICAL THEFT / LOSS	MISCELLANEOUS ERRORS	PAYMENT CARD SKIMMERS	POINT OF SALE	INSIDER MISUSE	WEB APP ATTACKS	
1%			1%	2%		91%	5%	1%	ACCOMMODATION
	9%			27%			45%	18%	ADMINISTRATIVE
32%	15%		11%	26%			9%	9%	EDUCATIONAL
				13%		73%	7%	7%	ENTERTAINMENT
36%			2%	7%	14%		11%	31%	FINANCIAL SERVICES
1%	4%		16%	32%		12%	26%	9%	HEALTHCARE

- Physical theft: 16%
- Miscellaneous errors: 32%
- Insider misuse: 26%
- Heavily weighted towards personnel errors and omissions?

Insider Misuse – end user heavy

- This is about intentional acts
- Heavily weighted towards "end users" of technology
- Health Care is probably no different, despite doing a pretty good job of educating its employees about the rules and consequences

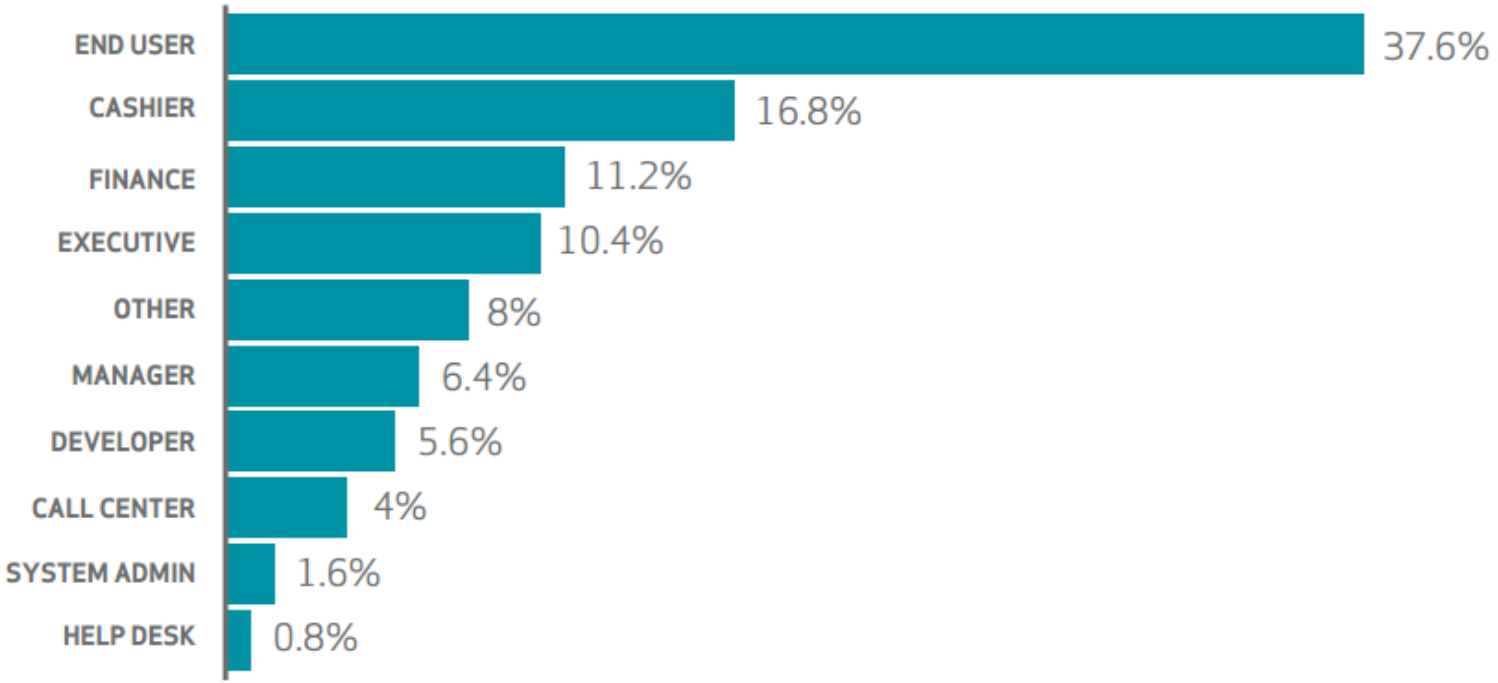

MISCELLANEOUS ERRORS

- Even across the entire IT industry, accidental disclosure of data is a majority of incident types
- Interestingly, disposal of data is still well-represented

"D'oh!"	Sensitive information reaching incorrect recipients	30% of incidents
"My bad!"	Publishing nonpublic data to public web servers	17% of incidents
"Oops!"	Insecure disposal of personal and medical data	12% of incidents

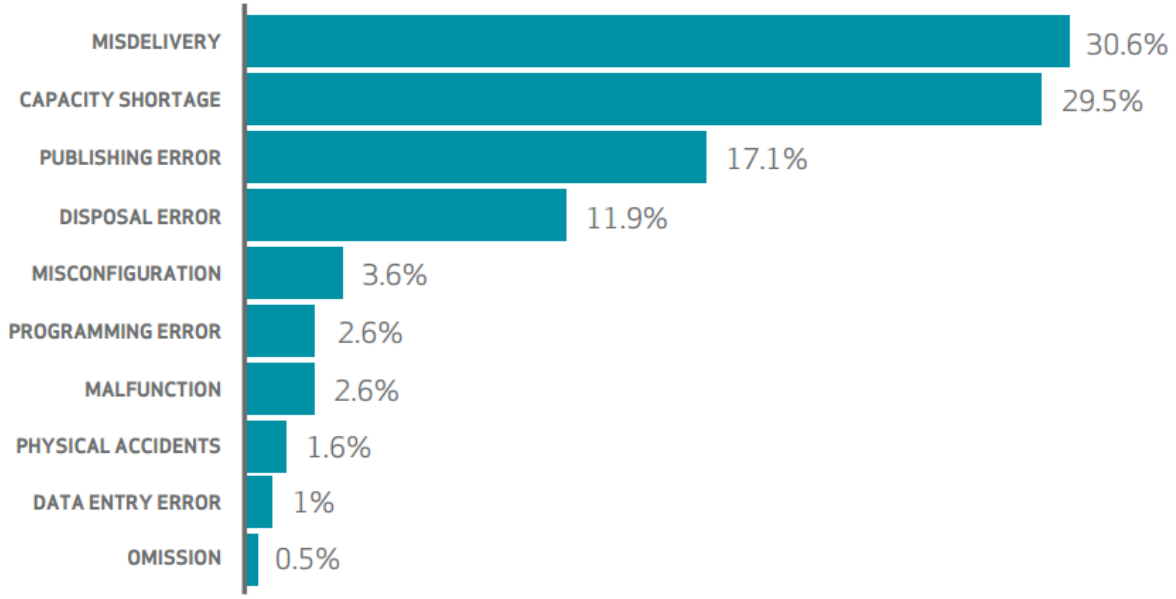

Miscellaneous Issues and I.T.

- Mis-delivery: 30.6%
  - Disclosure to inappropriate parties (are they checking the authorized parties list?)
  - The problem with fax machines – shared faxes, incorrectly given demographic information
  - Sending information to the wrong e-mail address
- Publishing Error: 17.1% / Misconfiguration: 3.6% / Programming Error: 2.6% / Data Entry Error: 1%
    - Development / test / QA servers that are publicly accessible
    - Mistakes on permissions and accounts
    - Inadequately vetted data being made public
- Errors that can be laid directly at the feet of I.T. are minimal – is this because they are not discovered, or not reported?

Incident Response in Health Care

- Based on my experience, technical IR in Health Care is one of the areas that needs the most attention
- Many IR plans are based on IR plans that are used in other parts of the organization (medication errors, slip and fall, "malpractice", etc.)
- Health Care IR is generally good in procedural areas:
  - Who to call
  - What form to fill out
  - How to contact compliance or legal
- But lacking in technical areas:
  - How to identify an incident on an IT system
  - Evidence preservation and investigation
  - Identifying the scope of a breach

Who is reporting these incidents?

- HIMSS (2015):

- Verizon (2014 report):

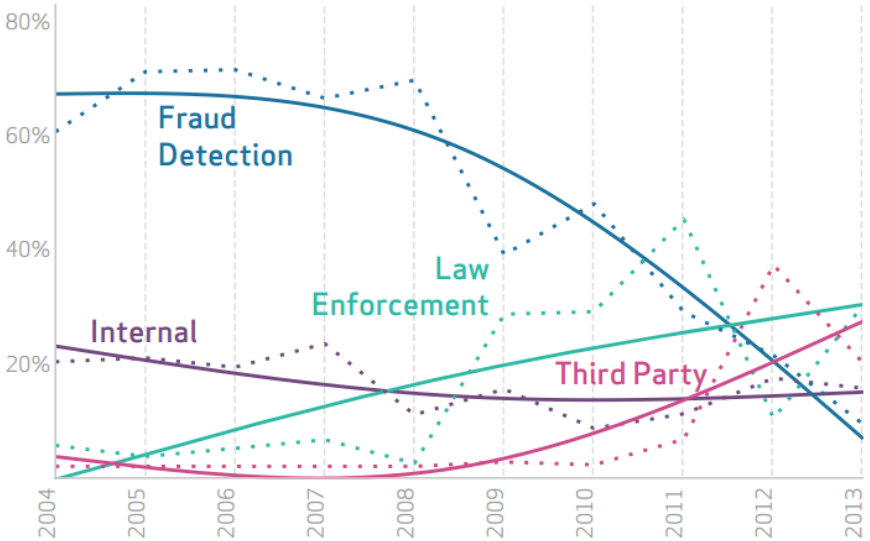

Who is reporting these incidents?

- It was not until 2013 that Verizon received survey results showing that more incidents were discovered by internal staff than by external referral
- Still seems to be a significant discrepancy between Health Care and overall IT
- Possibly due to breach reporting standards?
- Many reported incidents are PHI related rather than strictly technical such as hacking
- Requires us to put a strong emphasis on helping end users with training and awareness – especially anti-phishing training
- HIMSS on phishing:

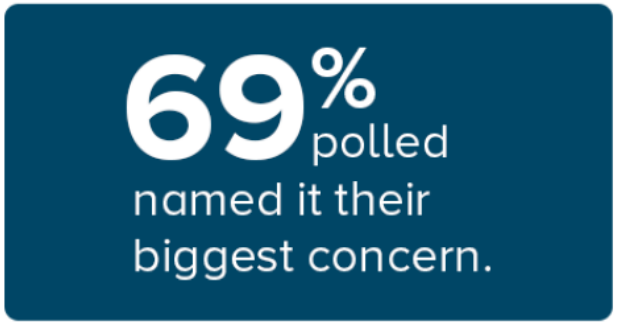

Phishing

- 22% of HIMSS respondents listed Mock Phishing exercises as a defense mechanism – this should be 100%
- My own company did a phishing exercise and a frightening number of our employees clicked on the link. NO employees reported it to information security
- This resulted in a significant corporate effort – mandatory training and a second phishing exercise
- As a penetration tester, I can say that attacking an organization through phishing is FAR easier than attacking it through technical means
- Humans have an in-built desire to be helpful, and attackers take advantage of this (and will continue to do so at an increasing rate)

Phishing – How I do it

- The first step is to do research using public records:
  - Social media (LinkedIn, Facebook, etc.,)
  - Scripts and software to enumerate names and e-mail addresses
  - Look for directories on official web sites
  - Identify generic inboxes such as marketing, accounts payable, IT helpdesk, etc.
  - Metadata from word and PDF documents – shows actual usernames and software packages used
- Free Tool: FOCA
  - https://www.elevenpaths.com/labstools/foca/index.html
- Free Tool: Maletego
  - https://www.paterva.com/web6/products/maltego.php

Phishing – How I do it

- TIP: Using a different format for user ID's and e-mail addresses makes it harder for attackers, as its much harder to find the login ID than the e-mail address
- Focus on: Management, billing, HR
- Avoid: IT, Risk Management, legal
- Create customized phishing emails:
  - "bypass your organization's firewall and content filter"
  - General messages from I.T. – new requirements, testing
  - Amazon gift card for participating in a survey
  - Infected PDF documents – tracking from UPS or a vendor invoice that looks just legit enough to open
  - Free iPad! (who falls for this any more!?!)

Phishing – The Citrix Server

- Create a fake Citrix web site registered under a name such as http://www.organization-beta.com that looks exactly like the official Citrix server (costs about $15)
- Send a phishing e-mail saying that IT is responding to user demand and rolling out a new, much faster, Citrix server and that they have been selected to test it. Fake the IT director as the source with a perfectly copied signature at the end
- The e-mail is from the lookalike domain, so any responses go to the attacker and not the IT director
- The fake web site will take their login information (user ID and password) and log it to a text file. After submitting their login, they get redirected to the real Citrix server
- User believes that they must have made a mistake typing in their password and often doesn't notice the change
- Sometimes take 3-4 logins before redirecting – the users will type in every password they know which is useful to the attacker

Incident Response – PHI Inventory

- The weakest part of IR that I have observed is in the detection phase. Out of all of the organizations I have assessed, I have never found an organization that had a good inventory of PHI locations
- TIP: PHI inventories are hard and constantly changing. IT cannot do it alone. Consider regular end-user self-assessments!
- If you don't know all of the places where PHI lives, how can you monitor it for abuse?
- Major applications tend to have logging, but these are only a small part of places where PHI lives. What about:
  - File shares and workstations – Word and Excel files
  - SQL databases, data warehouses
  - Single-purpose applications (pharmacy, radiology)
- Well run places have someone responsible for looking for public references to their organization and leaks through web searches, social media, complaints, etc.

PHI – Finding It

- As an attacker one of the first things that I do after getting a user ID and password is to search available file shares for passwords and PHI
- TIP: When you get back to the office, try doing a simple Windows file search of your shares with terms such as: *password*, *social*, *account* and see what you can find. This is only going to show file name hits. Then try a better tool that will look inside of files such as WinGrep.
- Typically will find passwords that get me into other systems – especially in instructions for IT Test / QA departments and for dealing with temporary workers
- There are tools available that can do this in an advanced way, some vendors will offer free assessments

PHI – Finding It

- The biggest problem with automated detection is false positives – there is simply too much data that looks like a SSN or account number out there. Workers get desensitized to reviewing false positives data all the time and may miss the actually important information when it comes along
- The same problem exists for secure mail appliances, which can be configured to identify what it believes to be PHI and block or encrypt it automatically
- However, there are so many false positives that it tends to annoy end users, and really annoys the IT administrator that is reviewing the rejected e-mails
- Systems have difficulty with passworded ZIP files
- This tends to lead to systems that rely on end-user discretion to manually flag a message with a checkbox in their mail client or a subject header
- Leaving it to the end user is more error-prone and requires more training to use it correctly
- Beware of SQL database risks - once I get a basic level of user access I start enumerating the tables and data contained in SQL servers with manual poking and automated scripts.

PHI – SQL Server Tips

- TIP: One of the best ways to minimize the risk of disclosure of PHI through databases is to ensure that strong development and DBA practices are in place. Use the Open Web Application Security Project (http://www.owasp.org) as a guideline:
  - Make a 1:1 relationship between SQL accounts and tables (avoid sharing the same user ID on the back end SQL database between applications, and never use the 'sa' account!)
  - Use each SQL instance for a single purpose instead of having a clustered set of servers for all databases. This is expensive, but it makes it much easier to control access
  - Many SQL databases do a daily export of their contents to a flat file so that they can be backed up. How are these files protected?
  - Use SQL data-at-rest encryption so that SQL fields are not in plain text
  - Beware of plaintext SQL login credentials in .INI, .BAT and other configuration files
  - Don't hard-code passwords in applications, it makes changing passwords extremely difficult, and finding the credentials very easy

Logging Systems

- With PHI located in so many places, how can you possibly log it all, let alone monitor it?
- Logging in SQL databases adds significant overhead
- Logging in web applications is important, but rarely well implemented (especially for internally developed apps)
- How can you log access to files such as spreadsheets?
- Can you identify multiple logins to an application from different locations?
- Can you log Internet traffic at the packet level for multiple devices? Could you go back to work and print a list of all Internet IP addresses and ports that a specified internal IP address has talked to in the last 60 days?
- Even if you are able to log most of this activity, how can you monitor it?
- I have yet to see what I consider effective proactive monitoring of PHI access – it is largely reactive due to complaints, celebrity visits, etc.
- SEIM and Log Aggregation tools exist but are expensive and still require manual care and feeding

Issues and Countermeasures: Vendors

- Vendors! There are too many of them, and they do not keep their systems updated. Never trust a vendor.
- In some cases, local IT may not be allowed to update them due to regulation or support agreements
- One good thing is that access to a vendor-managed system does not usually give access to other systems
- TIP: Ensure that patching and maintenance of IT systems is part of the purchasing process and final contract. If its not, its unlikely to ever be done. If it is in there, demand KPIs
- TIP: Do regular vulnerability scans or pentests to identify insecure systems and track their status
- TIP: Place vendor systems on a different network and enforce access control between them and the rest of the network (segmentation)
- TIP: Avoid using the same passwords on multiple systems – especially putting common admin passwords on vendor systems

Issues and Countermeasures: Staffing

- It probably comes as no surprise to anyone here but IT in Health Care is usually vastly under-funded
- The ratio of support staff to systems is low
- Inadequate time to track and follow up on assessment and audit findings – often just enough time to get remediation from one year's assessment done in time for the next years assessment
- Staff tends to be compliance and business intelligence focused, fewer generalists with wide ranging skills
- Support staff often focused on just one area, often in their own "island" and not coordinated with other employees (often in a completely different reporting structure)
- Rarely employ trained security staff that know how to attack and perform triage on systems. If there is IT security staff, they tend to be focused on Anti-Virus, firewalls, and intrusion detection systems (and home-grown)
- TIP: Try to employ at least one IT staff member that has experience in doing penetration tests and incident response. This "real world" experience can be invaluable compared to folks with just a few training classes or a product focus

Issues and Countermeasures: Rogues

- Many departments that do not "play nice" with centralized I.T.
- Individual departments purchase (and sometimes support) their own equipment without involving IT
- Example: One hospital I assessed had a pharmacist that had decided to purchase his own pharmacy software:
  - UNIX "green screen" application
  - Configured so that anyone from the Internet could access it via Telnet
  - User ID and root password were the same as the hostname (i.e. pharmacy.hospital.org, with the login of "pharmacy" and password of "pharmacy" providing administrator access)
  - Software updates performed by pharmacist or his favorite local "mom-and-pop" IT vendor
- Applications like this are often not even known, let alone protected – how are you going to find these systems? (answer: penetration tests)
- Will the media and public care whose fault it is?

Issues and Countermeasures: Big and Flat

- Other than for IP Telephony and remote buildings, most organizations have all of their systems on one large, flat network (or possibly a large number of subnets with no access control between them)
- This aids attackers – if I can get into one system, I can attack other systems on the network at will (the whole reason we use Internet DMZs in the first place)
- Modern attacks have a large manual component, and pivoting between systems is one of the major goals of an attacker
- TIP: Place different categories of systems on different VLANs and enforce access control between them. Pay special attention to anything not maintained by IT, servers vs. client workstations, SQL databases etc.
- TIP: If you cannot separate clients from servers, at least configure workstation firewalls. If I cannot attack clients, it will be much harder for me to get into servers
- The obvious problem to this recommendation is that figuring out exactly which systems need to talk to which other systems is painful and time consuming

Issues and Countermeasures: Authentication

- Authentication for remote access users is often based simply on their domain user ID and password
- Most organizations know that two-factor authentication is ideal, but there are often a very large number of users, making this cost prohibitive
- Authentication for workstations is another perpetual problem – need to balance security versus usability. If you make it too painful and time consuming to log in, people simply won't do it
- Use of proximity cards and virtual desktops that follow the user have been used with great success
- Even these have their problems – notably having a workable form of single-sign-on to link applications to a single account or card
- TIP: Use VPN portals to restrict remote users to just those applications they truly need, rather than providing packet-level access to the entire internal network, and save two-factor for vendors and IT staff that truly do need full access

Issues and Countermeasures: Portals

- Information portals, especially physician portals are a problem
- Well-run organizations will have a formal process for a doctor to request access to a portal, and will require some minimum level of training and acceptance of acceptable use policies
- In practice, though, it is often the front desk people at the physician's office that do all of the work using their doctor's account
- Office staff have a high turn-over rate and will be replaced often without informing the organization
- Passwords for portals will rarely be changed – prior staff can probably access a portal for months or even years after they leave the employment of a doctor's office
- Very rarely any notification of termination of employees and even doctors by the offices
- The access dilemma: use granular access control "by patient" that requires a lot of maintenance or dangerous "break the glass"?

Issues and Countermeasures: Passwords

- It seems obvious but passwords are still a problem, even when you require complexity
- Given the default Windows password complexity requirements, many users will still select guessable passwords
- The more accounts there are, the greater the chance of a bad password cropping up
- Examples:
  - Season-based passwords (Fall2015, Winter15, Winter15!)
  - Passwords based on a dictionary word with a number at end
  - Passwords that have the same base and increment the number
- Even when new users must change their password immediately, there are usually a few accounts that still have the default starting password
- Non-user and service account passwords (test, vendor, backup)

Issues and Countermeasures: Trusts

- One of the biggest things that we attack are trust relationships
- For example, shared local passwords and admin accounts
- Once we get an end-user password, we will dump all of the passwords that have been cached on that user's workstation
- Obtain the local administrator password, if an administrator has signed onto that machine, we will get their domain administrator credentials as well
- Using the user, domain admin, or local admin password we will connect to every other Windows machine and repeat the process, gathering yet more passwords, essentially mapping Windows
- Eventually we will find a Domain Administrator password and own the domain
- It often does not matter if the password is complex, as we do not need to "crack" it, we can use it in its encrypted form using a pass-the-hash attack

Issues and Countermeasures: Trusts

- TIP: Use different local administrator passwords for each machine. Yes, this is extremely inconvenient, but there are password safe tools that can do this automatically for you.
- TIP: Do not let administrators stay logged into systems with their admin account. Use split accounts so that day-to-day use is under a regular user account. Use "run-as" for administration purposes with a different account. This has the added benefit of limiting the impact of attacks like Crypto-Locker
- TIP: Use firewalls to stop attackers from connecting to and enumerating passwords on systems that end users would not normally need access to across the network (like workstations). If we cannot connect, we cannot as easily enumerate credentials!
- TIP: Make sure to well harden Citrix or Remote Desktop servers. If we can get on there, we will find a large number of credentials! It is often easy to "escape" from the constrained user interface
- TIP: Do not grant administrator privileges to local machines. Without this, we cannot dump the credentials as easily

Issues and Countermeasures: Patching

- Most organizations are fairly good at patching servers that they directly manage
- However, this requires maintenance windows! Dealing with true 24x7 shops can be very difficult
- Server applications tend to be poorly managed, due to vendor requirements and the large number of applications in use
- TIP: Have an inventory of critical applications (at least all of the ones that host or process PHI) and ensure that at least one IT person is responsible for monitoring vendor bulletins so they know when critical updates are available and required. Automated inventory tools can help to create and maintain the list. Consider using a Business Impact Analysis (ala BCP/DR).
- Third-party applications, particularly on workstations are dangerous! If you cannot keep Adobe Acrobat patched, you are going to have malware. I rarely seen an organization that did this well

Issues and Countermeasures: Patching

- TIP: Use automated patching systems for 3rd party apps, preferably with a service provider creating and pushing the updates so you don't have to
- TIP: Minimize the amount of patching that you have to do by using virtual desktop systems and "dumb" clients. It is much easier to patch a smaller number of RDP servers than to patch clients. It also allows you to update software one server at a time in a clustered environment.
- TIP: Use vulnerability scanning software such as Nessus to perform credentialed scans on workstations. You can use a service account to log into workstations and identify outdated software on a machine-by-machine basis. It can also find unapproved applications such as VNC or PCAnywhere running
- TIP: Don't forget about printers and appliances, particularly if they interact with other systems. We are often able to compromise a network by attacking an outdated printer that has "scan to desktop" type functionality

Good Example: Cincinnati Children's Hospital

- One organization that I think has a nice model. I was given permission to mention them, but obviously cannot give specifics
- They have a large research function – many powerful servers in clusters, a lot of higher education collaboration, grants, etc.
- Separate, segmented network for all of the research systems
- Separate authentication systems, logging, patching, etc.
- Regular vulnerability assessments (performed by internal staff)
- Dedicated security and management staff just for the research network
- IT administrators trained on their technologies and on security
- Security is formally built into the System Development Life Cycle (SDLC) and reinforced in database admin, programming, and system support functions
- Supported by a dedicated security function in the parent organization

Q&A / Discussion

????

Thank You!

Mark Lachniet
mark.lachniet@cdw.com
Manager, Information Security Solutions
CDW
1000 Town Center Suite 1800 Southfield, MI 48075
Mobile: 616-304-3526