Spring | 14

Cloud Transparency Initiative

Jordan Atwood, Doug Wynkoop

[Type the abstract of the document here. The abstract is typically a short summary of the contents of the document.]

Information School, University of Washington
2014 Capstone

Table of Contents

Setting the Stage

It seems a week doesn't go by without another headline news story about a data breech or computer security snafu. Usually we think of large companies like financial institutions or national chains being the targets of cyber attack, but a recent Verizon study revealed that wasn't always the case. Small food and retail companies are also often the victims of cyber breaches as well.

Smaller companies may not have an IT department or a security specialist on hand to counter this threat. They're also among the companies that may to have the most to gain from switching to a cloud solution in other ways, too. For larger companies who have spent a lot of money building up their IT infrastructure and human resources, dropping all of that to move to a cloud service might not make sense. For a small company with almost no investment in computer and network technology, allowing another company to take care of their IT needs might well be a much cheaper and more secure option.

So, we wondered, were local companies doing this? The answer, we found, was no.

During out interview process, we encountered companies who had successful in-house IT solutions, and ones using ancient legacy systems developed in the nineties. We talked to companies who were doing their numbers by hand, and companies who simply sent all of their receipts to an accountant. We also talked to a company who had outsourced their IT, but to a local company who provided an in-house product instead of a cloud-based one.

Of the companies we interview in-depth, none of them were entirely happy with their IT solutions, and all would be willing and able to pay more for something that worked better. But none of them had moved into the cloud in any kind of meaningful way.

Although we did encounter a variety of reasons for this, such as cost effectiveness or concern about special business requirements, one item that was a deal breaker for a all of them was the security of the cloud itself.

Small businesses, then, are stuck. They want better network and computing solutions, but are being blocked by their problems accurately assessing security. To try and counter that, we've done research to help explain some of the biggest problems with cloud and network security and how to fix them. After reading this paper or reviewing the other resources on our site, git.io/cti, we hope small business owners and IT professionals will have a better idea about what kind of risks are involved in cloud adoption, and how to mitigate those risks to levels acceptable to your business.

What is the Cloud?

Before we really get started, let's take a moment to determine exactly what people are talking about when they refer to "the cloud." According to the definition given by the National Institute of Standards and Technology, a cloud service has five essential key factors that when taken all together make it "cloud:"

1. The service should be on demand – a customer should simply be able to ask for cloud functionality and receive it without having to interact with another person to receive it.
2. The service should be accessible over the internet and from many different platforms – everything from cell phones to laptops to desktop computers.
3. The cloud provider makes the service available by pooling multiple resources together. For instance, a cloud storage provider might store data in several different data centers, but to the customer it all looks exactly the same.
4. The fourth characteristic of a cloud service is the ability to adapt to changing demand – the more of a cloud service a customer uses, the more is made available. If they start using less, then the service is automatically scaled down along with the diminished demand.
5. As a result of the elasticity of the service, cloud services are billed according to how much the client actually uses.

In the end, the idea is to have a service that is simple and easy to buy, access, and use. Because the provider can pool many resources together to realize economies of scale, and because the client pays for only what they use, cloud services have the potential to be cheap, efficient, and effective. However, it also introduces security concerns.

At the core of these concerns is the fact that adopting a cloud service splits responsibility for keeping a company's IT infrastructure safe between two different parties. To complicate things even further, the cloud service's interests and the client's interests may not always be aligned. That isn't to say it's not possible to manage the risks of a cloud deployment, but it does mean that you have to go in with your eyes open.

Authentication and Accountability

One question is at the very heart of cyber security – how can we keep people we don't want away from private information while allowing access to those who need it? In order to begin to answer this question, you have to determine first who the person is, which is generally referred to as authentication, and what they should have access to, which is known as authorization. Without taking time to define who should have access to what, no other security measures will be effective.

Authentication

There are three factors we can use in order to identify someone attempting to access a computer system:

1. Something they know. The most common examples of this include passwords and PINs.
2. Something they have. This can either by a physical item, like a common access card, or a "virtual" item, like a digital certificate.
3. Something they are. Often referred to as biometric authentication, this measures something unique about a user, like a thumbprint or retina pattern.

Generally speaking, the more things a system checks – often referred to as "authentication factors" – the more certain and secure the authentication is. However, this also has to be balanced against ease of use.

The password is the most common kind of authentication currently in use today. Although passwords can be tricky to remember and awkward to use, alternatives may not be available. Because of this, it's necessary to take a moment to talk about a few vulnerabilities of passwords and how to counter them.

The simplest attack on passwords is for an attacker to simply keep guessing things until they get it right. This may sound like it would take forever, but if passwords are done incorrectly it may become viable.

Most of "brute force" attacks are done using an automated program that will do the guessing for an attacker. Using passwords comprised of dictionary words, birth dates, names, or "keyboard walks" like "12345qwert" are very easy for a password-cracking program to guess. If the program is allowed to make guesses as quickly as the authentication server will respond to it, an attacker may be able to try thousands, if not millions of attacks in a very short amount of time.

One way to counter this is by requiring strong passwords that require the use of numbers, lower and upper case letters, and special characters that are at least eight characters in length. Some experts suggest having a minimum length even greater than this. Pass phrases, which are strings of random words, may be easier to remember and also meet complexity requirements. Changing passwords on a recurring basis is also recommended.

Another method that can be used to make these attacks less practical is to only accept authentication requests that come straight from the login page, and to pause briefly before checking to see if the user name and password are correct. Even a one second delay can slow brute force attacks down significantly. Simply locking the account after a certain number of failed tries may also be a way to attack the problem, although a mechanism for unlocking the account again should be in place. Also, messages in response to a failed login attempt should not specify if it was the username or password that was incorrect. This will keep attackers from trying to determine what valid user names are in the system by guessing.

If your business has a large number of services that require different user names and passwords, it may be tempting for your employees to start re-using passwords, writing passwords down, or making weak passwords to manage the increasing complexity. This can be a problem if your company adopts a variety of cloud services, each with their own authentication systems. A way around this problem is federated identity.

In this system, a user will log in only once, and be given a token that will allow them to freely access any system they are authorized to use. This can greatly simplify the problem of generating and remembering dozens of passwords. However, it may be necessary to exclude administrator accounts from this system. Administrators often have broad abilities to issue permissions, add or remove programs, and other high-value, high-risk tools that are attractive targets for hacking attacks. Making access to these tools easier could be counter-productive to security.

Extra care should be taken if you plan to use biometric information in your authentication plan. While you can change a password, everyone comes with only one set of fingerprints. Once stolen, a retina scan or thumbprint is current forever. If it's possible to avoid storing actual copies of this data, it is highly advisable to do so.

Storing authentication credentials is also important. In order for a system to determine if a user has given it the right authentication information, it must store a copy of this information somewhere. That means that copies of passwords, biometric data, or security certificates will be kept somewhere on the network. Because it has to be able to be accessible even to people who are not authorized users of the system, this information may be at risk.

It's wise to assume that a determined attacker can eventually access stores of user names and passwords. Although it may be acceptable to store user names without protecting them, passwords should be "hashed and salted." Hashing refers to changing a piece of information to a long, scrambled, and unique pieces of data.

Hashing makes a password impossible for an attacker to read, but if two users happen to have the same password, the hashing algorithm will generate identical strings of gibberish for each of them. So, if a hacker has access to a large number of passwords, it's easy to look for patterns like this and decode all of them at once. "Salting" prevents this. Salting refers to the practice of adding a small, random piece of data to the password before it is hashed, so that identical passwords will appear completely differently when stored.

Authorization

Once the person has been identified as a valid user, the next step is to determine what they should have access to. It might be tempting to simply give all users the ability to do anything – and this might work in a very small environment – but the key to keeping information safe is the principle of least privilege. This principle simply states that any user of a system should have access to exactly what they need to perform their job, but nothing more. By maintaining the principle of least privilege, we're able to minimize the damage that can be done by an angry employee, or by an attacker who has stolen someone's authentication credentials to get access to the system.

Actually determining who should have access to what information can become quite complex. Although entire books have been written on the subject, unless your company handles classified or extremely sensitive information, there are some straightforward ways of managing this problem. The most commonly used one is called role-based authorization.

In role-based authorization, a company determines tools and information someone in a particular role needs to do their job. As a result, instead of having to grant permissions to each person individually, an administrator can simply add a person to a "role," and they'll automatically get all the permissions they need to accomplish their tasks.

Any system is only as strong as its weakest part, and all to often management of authentication and authorization systems is where things go wrong. For starters, don't use any system that someone's just invented for authorization – there are often industry standards for this kind of thing, which are well tested and secure.

In addition, issuing, revoking, and reviewing the status of authentication credentials and authorization permissions is extremely important. Just like landlords often retrieve the keys or change the locks after a tenant leaves, IT managers should make sure that once someone is no longer in a specific job or employed by the company that their ability to access programs and information is changed accordingly.

In order to facilitate management, centralizing and automating portions of these tasks might be helpful.

When it comes to security when using a cloud product, authorization and authentication become extra important. This is because there's no way to physically block someone from accessing a cloud system, and older, network-based security tools designed to keep out harmful traffic may not work either. This means authentication may be one of the only things keeping attackers out, and authorization one of the only things limiting their access once they're in.

For this reason, it's important to determine what kind of system a cloud provider is using before you choose to adopt them. How many factors are they using for their authentication? How are authentication credentials stored, and where? Do they have a role-based authorization system, and if so, does it give you enough control? Finally, it's important to determine if these systems will work with the authentication and authorization systems in place on your own systems.

Network Security

In order to understand network security, we should take a moment to define a few key terms that a non-technical person may not be aware of. These definitions are a simplified, and designed to help someone with little network background knowledge.

Terminology

Servers and clients

A server is simply a computer that stays on line and provides a "service" to any computer that contacts it. The computer requesting the service is a client. The service servers provide can be anything from running printers to video games.

Ports and Services:

A port is a piece of software designed to listen for certain kinds of incoming network traffic. For instance, Internet traffic has its own port, certain kinds of email have their own ports, etc. A service is simply a grouping of ports. There are literally thousands of ports, most of which are rarely used.

Router

Routers a device found positioned between networks, and help network traffic find its way from one network to another. Switches and hubs perform similar functions, but in different contexts.

Firewall

A firewall is a software or hardware device that is designed to block malicious or unauthorized network traffic. It does this by keeping requests from reaching certain ports, or by filtering out incoming traffic based on its IP address.

IP Address

Every device hooked up to the Internet has a unique Internet Protocol Address. This allows traffic to go exactly where it needs to go without getting mixed up. Computers on a single network may share an IP address; IP addresses can also be faked.

Virtual Private Network (VPN)

VPNs allow a computer to remotely connect to a network as if it was physically on that network. It does this by creating a very secure encrypted connection across the Internet that would be extremely difficult to break into. VPNs can either go from an individual computer to a network, or connect two networks together.

Encryption

This involves scrambling data, either while it's traveling from one place to another, or is in storage. The trick to encryption is that it's relatively easy to open if you have an encryption key, but very difficult to decode if you do not. This allows people who have keys to communicate safely.

Supply Chain

Good network security starts at the most basic level – computer and network hardware. Picking hardware isn't just about getting a good price or choosing a reliable brand, however. Many computers and network devices have components built from many places around the world. It's possible that one of these sources may introduce a security flaw into one of these devices to allow an attack at a later date.

It's advisable to have as good an idea as you can of how and where your devices were assembled, but often this simply can't be done. Modern supply chains are just too complex. The safest alternative is to be vigilant in monitoring your network for possible rogue behavior by one of your network components. At the very least, hardware and supply chain vulnerabilities should not be taken for granted.

Another good way to help reduce the chance of leaving accidental vulnerabilities in your network architecture is to try and use the same kinds of device when possible, instead of mixing and matching. Although a network that has many different kinds of hardware on it may be less vulnerable to falling prey to a single attack, this kind of setup makes the network much harder to administer. If, for instance, all of your routers are exactly the same, keeping them safe, up to date, and running smoothly will be much easier.

Physical Security

A simple way to boost security is to simply restrict physical access to key network components. No one but administrators should be allowed to physically access your company's servers, and contractors should be vetted and accompanied while making repairs. Putting key devices behind locked doors, and encasing cables in metal ducting is a good way to ensure the key parts of your infrastructure remain safely tamper-free. More advanced set ups might require card-based access and other, more rigorous controls.

Removable Media

Another way a network can be compromised at a hardware level is through the use of removable media, like CDs and USB drives. Using these devices is a big security risk for a number of reasons. One risk is that it's very easy to steal information with removable media. They can fit large amounts of files into a device that can hang from a keychain, which makes data breaches of enormous size a real possibility, as shown in the recent Snowden NSA leak.

Removable media also allows data to enter your network without passing through any of the security devices that separate your systems from the Internet. An unwary user could easily infect one of your computers by mistake. Disabling removable media where possible or restricting its use can go a long way to prevent breaches and to keep your network safe.

Patching

There are certain threats that are almost impossible to defend against. These brand-new attacks, called "zero day exploits" are vulnerabilities no one has ever seen before. No one, that is, but the attacker using them. Against this kind of attack there's little defense, but once the attack is analyzed and countered, computer security firms and the producers of the affected software or hardware will issue a patch to fix it. However, these patches work only if they are installed!

For this reason, all network devices, servers, and computers, as well as all the programs on them, must be kept up to date and secure. Often times, this requires restarting the device in question. If the device is a personal computer, the user may keep putting it off. If the device is a server, turning it off may need to be done at a special time or a special way to avoid disrupting your business. Any safety-savvy company should have a plant to make sure these patches are installed soon after they are issued, and that the network and its devices are always up to date.

Ports and Services

Every open port or running service on your network is a potential vulnerability, especially if that port "faces," or is accessible from, the Internet. However, we can't simply close all network ports, because then there's no point in having a network or Internet connection at all. The key is to close all ports and services that are not in use and block them with a firewall. For some ports this is especially vital, as they are inherently unsecure. For instance, the Telnet service allows a computer to remotely execute commands on a machine that has Telnet running, while remote administration allows a foreign computer to entirely take over a machine with the service enabled. Unless there is a pressing need for one of these high-risk services to be running, turn them off!

Network DMZ

Once a network is connected to the Internet, it's faced with a problem. Allowing just anyone to try to connect to internal company infrastructure is deeply risky, but there are some systems that must face the Internet to work, such as email servers. To solve this problem, most network admins create a demilitarized zone, or DMZ, between the company and the internet.

There are a number of ways to set this up, but among the safest ways is to have one layer of firewalls between the DMZ and the internet, and another, more restrictive layer of firewalls separating the DMZ from the company's networks. This has the dual benefit of allowing services in the DMZ to run with the most security they can, while providing multiple layers of security between sensitive company information and potential attackers.

A number of services and security devices can (and often should) be located in the DMZ. Some of these include services that absolutely must face the Internet, such as email servers, authentication servers, and VPN servers. Security systems that might be present in the DMZ include intrusion detection systems, which inspect incoming and outgoing traffic for suspicious activity, and proxy servers, which can enforce a company's browsing policy and can hide the architecture of your network from anyone observing your outgoing traffic.

Defense in Depth and Segregation

Putting up multiple barriers between an attacker and vital infrastructure, like the multiple firewalls in a DMZ, is a strategy known as "defense in depth." Since compromise of at least one layer of security is likely, using multiple layers can help give security professionals time to identify the attack and stop it before it reaches vital data and processes.

To further help a defense in depth goal, certain portions of the company's internal network may be segregated from each other depending on the sensitivity of the data each part contains. This can be done through the use of virtual networks, which allows for the multiple networks to all share the same hardware, internal firewalls, or even complete physical segregation.

Wireless Technologies

DMZ architecture is good for separating the physical infrastructure of a network from the potential harm – but what about the infrastructure that isn't physical? There are a number of wireless technologies that can greatly benefit companies to use, but all of them come with their own risks as well as their rewards.

Wireless networks are among the most common kinds of wireless technology, and see a lot of use in both home and business contexts. However, they have two serious drawbacks – users to not have to be physically connected to the network to access it, and once a user is on a wireless network, it can be difficult to tell who they are and where they are located.

Keeping users off of a network involves using strong authentication paired with a good wireless encryption system. Today, the best kind of wireless encryption is the WPA2 protocol. Without both of these in place, it may be possible for a nearby attacker to read traffic being sent over your wireless network. WEP, an older protocol, is now out of date and no longer provides adequate security.

For larger companies, centralized administration of wireless access points can not only allow for advanced functions such as wireless roaming – that is moving seamlessly from one wireless access point to another as you move around the building or campus – but can also allow administrators to use new, advanced tools that can physically locate and identify devices on the wireless network. For smaller companies, or companies with lower security requirements, these tools may be overkill.

Bluetooth is another commonly used wireless technology. However ubiquitous it may be, Bluetooth was not designed with security in mind. Many devices have Bluetooth discovery left on by default, which may allow an attacker access to the device. Further, Bluetooth encryption and authentication are viewed as lacking by some security professionals. Alternative protocols designed with security specifically in mind, such as ZigBee, are more secure, but also more costly.

More and more companies are adopting RFID tags. These are small tags that can be attached to just about anything and used to store information. This can be extremely helpful in inventory management, shipping, maintenance logging, and more. Passive tags can only be read from about three feet away, but active tags, which have a small battery, can be read from much further away from that.

As convenient as these little devices may be, most to not support authentication, encryption, or activity logging. As a result, anyone can read or change the information on these tags (if the tags are writable) without leaving any evidence that they have done so. Although solutions for this problem are up and coming, RFID tags should be considered largely insecure for the time being.

Work Laptops

Another way data can bypass a network's security is via laptops and other mobile devices. If the company owns the laptop, then it can use all of the same security controls that it would ordinarily use on a given company machine to keep it safe – enforcing updates, for instance, or disabling boot from USB or CD, a measure that's very important to take with these devices. However, there are two new major vulnerabilities introduced by a laptop that can be removed from the building – loss, and use over unsecured networks.

Loss can be addressed in a number of different ways. For starters, there are various different kinds of anti-loss and anti-theft software and hardware available. This includes GPS tracking software and a variety of locks, safes, and security devices that can be physically attached to the laptop.

Encryption can make the device useless even if it is completely lost or stolen. There are two ways this can be done. One way is to encrypt the entire hard drive. This makes the laptop nothing more than an expensive paperweight to a prospective thief, but once the laptop is powered on and unlocked it is relatively unprotected.

The other method is to use file-level encryption, where sensitive files are locked down individually. This means that even if the computer is stolen while powered on and unlocked, a thief would still not be able to access sensitive data. On the other hand, they would have access to anything else on the device. Choosing how to encrypt work laptops is largely down to how they are to be used and the needs of each individual company.

BYOD

Bringing your own device to work, or BYOD, is becoming increasingly popular, especially when work can benefit from the use of a smart phone. However, adoption does open a can of worms. Questions arise about who is responsible for the device, how they should be supported, who repairs them if they break, how they should be integrated into the company – and all of this has to be made easy to really see the benefit of this kind of program! Worse, these devices have all the risks of a work laptop without any of the controls.

Authentication is one of the main issues – typing in a strong password on a small smartphone screen is a real problem, and if repeated login attempts lock down or wipe the phone, it's likely to happen a lot. It's possible to take measures like installing security software, or trying to protect the phone from malicious apps by using a "blacklist" that bans known troublemakers, but the former tends to aggravate users, and the latter is hard to keep up to date. In addition, users have a great deal of power to bypass security protocols on phones they own, and catching them can be difficult.

To counter this, it's often easier to try and secure the data instead of trying to secure the phone platform. Encrypt any data sent to the phone. Use role-based authorization to limit what a user can access or store, or consider allowing access only to pre-approved information and content. Try "sandboxing" work applications on the phone – that is, running them in their own separate environment apart from the rest of the phone's operating system. This has the added advantage of requiring tough authentication only when trying to access valuable data. Also, if the phone is lost and a remote wipe becomes necessary, it may be possible to remove only proprietary work information instead of deleting all of the users personal data.

Using a centralized management system has advantages as well. Keeping track of what data has been sent where and accessed by who boosts security and accountability. Finally, educating employees on responsibilities and mobile phone risks complements these other methods and addresses what is often the weakest point in any IT system – the end user.

Application Security

Any application that runs on your network also should be vetted. Whether malicious or not, a bad application can destabilize your machines, cause crashes or data losses, or even compromise your security. To ensure that the programs you run are safe, it's possible to get them scanned for vulnerabilities. This is not an inexpensive process, and is not one that can be done in house, so working with the vendor and a third-party scanner may be the best way to go. Scanning applications you have developed in house may also be advisable.

These scans can be done if you have either the source code of the program, or just the executable file. They can reveal a number of weaknesses, including "dead" or redundant code, as well as branches, a term that refers to possible security problems built into the code, which may be accidental or maliciously installed "back doors" that can allow an attacker unauthorized access.

Finally, although this may go without saying, it is absolutely vital to change all vendor-default user names and passwords on all computers, servers, and network devices. You may laugh, but this is a shockingly common vulnerability around the world.

At this point you may be wondering what all of this network security has to do with a safe cloud deployment. For one, a cloud deployment is only as secure as your own network. Even if your provider had the best security in the entire world, if your end of the equation is unsafe it's all for naught.

The other reason is that cloud providers have their own corporate networks, and cloud solutions use network technology to run. This means that the cloud company should be addressing the potential vulnerabilities and using some or all of the controls mentioned here to keep themselves – and their clients – safe.

Since this kind of information is often proprietary or secret, however, it may be that the company is reluctant to give you details on the effectiveness and comprehensiveness of their network countermeasures. As I will reiterate over the course of the paper, it will be up to you as a customer to demand that they provide enough information that you can make an informed risk decision.

Logging, Auditing, and Monitoring

Logs

Logs are simply records of what your network has been doing. Every device on your network will generate some sort of log. Although often dry and tiring to read, logs are absolutely vital security tools. Logs are used to determine what normal behavior in a system looks like so that anomalies can be detected, and are one of the only ways to catch intruders and fix gaps in security measures. Without logs, there is little an administrator can actually know about the status of their network at any given time.

Ideally, logs should be recorded in a centralized database or file system. Some devices may send out only summaries, or use a proprietary logging format, so be sure your logging method can account for these potential problems.

Logs must contain certain key pieces of information, including who, what, where, and when. The "when" is particularly sensitive, and relies on properly implemented Network Time Protocol, or NTP. Without precisely accurate and synchronized time, it may prove impossible definitively identify prosecute attackers. The log should also contain information on the severity of the event, which is score from 0 (highest) to 7 (lowest). Each event should use the lowest appropriate severity.

Which events should be logged vary from company to company, but keep in mind that they are your main record of any potential problem or attack. Certain high-risk events, such as failed logins, or the granting of new user privileges should almost always be logged. Information that points to possible attacks or malfeasance should also be logged.

There are also things logs should *not* contain. Sensitive information, such as passwords or encryption keys should never be logged, nor should any private information. It may be possible to log semi-sensitive information in some circumstances, provided that it is scrubbed out when the log is displayed.

Log information must be protected both in storage an in transit, as it is an important target for a hacker looking to cover up his footprints. Access and changes to logs should be recorded, and logs should be transferred to read-only systems quickly. Logs in transit should be encrypted and their origin verified to be sure they're not being spoofed or tampered with. Requests for log information from third parties should also be carefully vetted.

In order for logs to be useful, they must be available to the right people and systems, stored for the proper length of time, and reviewed regularly in order to be useful. In order to make sure your log system is meeting all of your needs, it should be periodically audited and reviewed.

Scanning and Testing

Other vital tools for understanding how your network is running and how safe it is are vulnerability scans and penetration testing. Vulnerability scans check your network against a long list of known potential intrusion points, including recorded exploits, open ports, and unsecured services.

Although patching, antivirus, and careful firewall maintenance measures may prevent much of this, changes over time in your network may result in the accidental introduction of a problem that this kind of scan can detect. Checking each one of these potential problems by hand is far too laborious for any one person or team, making these scans an extremely helpful auditing tool.

Penetration testing is also a valuable way to assess your actually level of security. This kind of testing involves contracting a third party to try and overcome your defenses and tell you of any potential weak spots they encounter on the way. In addition to simply being a good idea to perform periodically, penetration testing is actually required by certain compliance regimes.

Logging, testing, and auditing are vital to ensure that your network is behaving as it should. In absence of these measures, the only way to detect a breach is when something goes horribly wrong. Lack of these tools is a serious potential stumbling block in any cloud adoption. Because of the nature multi-tenant – that is, many users on the same service – nature of cloud services, vulnerability scans and penetration testing may not be allowed. In absence of these abilities, you must determine what kind of assurances the cloud provider can give that their infrastructure and yours are free from the kind of security weaknesses best detected with these auditing tools.

Some services also may not offer logging support, provide insufficient logging, or logs that do not integrate properly with your log storage system. It may be that they offer other kinds of visibility data that would make up for this, but adopting a cloud service that offers no logs is probably unwise.

Governance, Organization & Risk Management

Cloud Governance

To be able to really trust a cloud provider, it's important that you know as much about how they do business as possible. To begin with, how is their business doing? How old are they? How competitive are they in their space? Are other customers pleased with the service they're getting? Are they making money, or are they in the red? And finally, what will happen should the company fail or be bought out?

Have they oversold their service? This is a common practice among cloud companies – they may not have enough backup capacity in case they experience some kind of failure, or be in the habit of compressing client's files. Should a file be corrupted and clients begin to panic, this can lead to a "run on the cloud" if something goes wrong. Network connection speed could also just normally be very low, sometimes for reasons the cloud provider themselves control. Be sure the connection is fast enough to actually move your data back and forth as needed.

Cloud Organization

Proper security requires that certain organizational conventions be met. Most importantly, there must be a segregation of duties, especially between the people who design a company's security architecture and those who implement it. Comingling these responsibilities can lead to conflicts of interest, too little oversight, or too much power in the hands of too few. In addition, it's good to know the size and experience of the company's security staff. How many problems can they handle and once? What's their record on dealing with security problems?

Cloud Risk Management

How a company deals with risk is also extremely important. The priorities of a cloud provider and a cloud customer may not always align. How to they vet new security employees, onboard new applications, or pick out hardware?

Documentation on these processes can tell you a lot about a company and how it handles risk. This is important, because once you use their product, their risks become yours as well.

Moreover, cloud adoption imposes additional problems for a client company in terms of governance, organization, and risk management as well. In order to deal with these changes, a prospective client must do their due diligence.

Local Governance

To begin with, how easy or difficult would it be to switch services? If the cloud is using an open-source storage standard, moving back from the cloud or to another provider should be relatively smooth – if not, it might be difficult or even impossible to get all your data back out.

How do you get billed for the cloud service? Are their terms clear and easy to understand? Do they provide you with real-time billing information, or simply hand you the tab at the end of the month? How easy is the bill to understand when you do get it? It may be necessary to have help to fully break down the statement; some large customers of Amazon Web Services have such detailed and lengthy bills they can't be contained in an Excel spreadsheet without crashing it.

Before moving any data into the cloud, consider what compliance issues you may have to deal with. While the easiest way to determine if a cloud service allows for compliance with a given regime, it's vital to remember that just because they are compliant does not mean that you automatically inherit their rating. Nor does the fact that one of their clients has managed to become compliant mean that you will be able to replicate the achievement.

Oftentimes what a cloud provider being "compliant" means is that the specific security benchmarks that are their responsibility have been met. This leaves it to you to take care of everything else. It is also the client's responsibility to get definitive documentation and proof of what parts of the provider exactly have been certified, and when. Compliance is an ongoing process, and their status may have lapsed.

Local Organization

Contracting with a cloud company can decrease your IT overhead, but this should be approached with care. Besides the moral impact of potentially firing IT staff, outsourcing responsibilities to the cloud could lead to a decrease in their overall level of skill. Most companies of sufficient size need in-house IT skills, so it's important to be aware of this potential organizational risk. This can be worsened by the fact that the more cloud services are adopted, the more security regimes the client's IT staff must learn. There may come a point when their ability to coordinate between multiple services and environments breaks down.

Local Risk Management

Lastly, risk management becomes difficult when you don't fully understand how a cloud provider is set up or does business. Some cloud providers may provide excellent evidence for the effectiveness and readiness of their security measures; others may provide little to none. How much your company needs is something it must determine for itself. Before it does, though, it must find out the answer to one last question: does the cloud provider subcontract with other cloud companies or third party firms? If so, the evaluation process may need to start all over again.

Network Control

Sacrificing control over part of your IT infrastructure is a necessary, even desirable part of adopting a cloud service. However, it can pay dividends to be sure of exactly what you're giving up and what you're getting back in return. Any function that you turn over requires some assurance that they're handle it properly, but some particularly vital ones are highlighted here.

Authentication and authorization mention a special mention again here. If the provider is handling authentication, how are they storing the authentication credentials? And how do they handle password recovery?

Does the provider have a good authorization plan that allows you to maintain least privilege? And how to do you know that authentication and authorization is being granted only to your employees?

Encryption and encryption key management are also extremely important. As we will cover later, encryption is absolutely vital in the cloud. How are they implementing it? Are they using a sufficiently strong and safe standard? Who controls the keys? If they do, how can you be sure that your data is safe and the keys properly stored? If you control the keys, what happens if a key is lost? Can the encrypted data be recovered?

Lastly, as we covered earlier with ports and services, having the minimum number of programs and tools needed for the job – referred to as the principle of least function – is often vital for maintaining security. If the cloud provider makes services available that the client doesn't need, is it possible to turn them off or restrict their use?

Accidents, Disasters, Recovery, and Continuity

Sooner or later, accidents or disasters are bound to occur, and can potentially cause far more damage than any hacker. The impact of these mishaps should be minimized, recovery planned in advance, and measures put in place to maximize capacity during and after the event.

Types of Accidents

Disasters can be anything as simple as a power outage to as cataclysmic as a major earthquake. Anything that causes a power loss or cuts a connection can impact a customer's ability to connect to a cloud provider. Accidents can happen on either side of the connection as well, and can less predictable and even more damaging than a disaster. This can include something like inadvertently reformatting the wrong hard drive to spilling a drink on a server rack. Lastly, simple wear and tear will eventually cause equipment to start failing.

Equipment failure is perhaps the easiest of the three to address, and can be addressed the same way for both cloud providers and cloud clients – have a good equipment replacement plan. It may be painful to replace working equipment, but nowhere near as painful as losing an entire hard drive full of data, or losing a connection to the Internet (and your cloud service) for half the day.

Some accidents are simply unavoidable and unpredictable, but any company can adopt some best practices to minimize them. Keep food and drink outside of sensitive areas, and restrict access to only those who need to be there. Properly train all of your employees, and restrict permissions to functions that can do real damage to only those who need them most.

Recovery and Continuity

Disasters can strike rapidly and endanger both life and property. The first priority of any disaster response should be to preserve life and limb. The rest is just stuff. However, once the disaster has passed, or if the disaster is ongoing, like a snowstorm, a company should have a plan for who will be in charge of making repairs and recovering, and how they will go about it, as well as a team dedicated to keeping things running as best they can in the mean time.

In any event, redundant systems and hardware may be extremely helpful for avoiding outages on both sides. For a client, this might involve a backup generator and uninterruptable power sources, as well as multiple copies of devices that connect the network to the internet, such as routers and firewalls, set up to take the full load should one fail. For a cloud provider, this might involve contracts with multiple power companies for their data centers, and massive redundancy in power, Internet and cooling connections

Backups

Backups are a must for any company with electronic data. There are three different kinds of backups – cold, warm, and hot. Cold backups involve writing old data onto formats like tapes, which are secure, but take a long time to recover from. Hot backups often involve mirroring – that is, copying an entire database or production environment, and automatic failover, so that should one database fail, the other will immediately pick up with little or no downtime. A warm backup is somewhere between the two extremes.

Likewise, backup locations come in similar flavors. A cold backup may simply be a reserved space with little or nothing in the way of hardware or backup data on hand, whereas a hot backup is a complete reproduction of the original facility, with a warm backup, again, being somewhere in between.

Another thing to consider is where the backup is physically located. If the backup is in the same geographic area, or worse, in the same building, it will likely be affected by a disaster in the same way the original data or location would be.

In addition to planning for these contingencies themselves, a cloud adopter must also learn what plans the service provider has made in order to accurately assess the risks of using the service and preparing to work with them should a mishap or disaster occur.

Attacks

Wherever a service exists, there is an attack that can exploit it. Many of these attacks can be extremely subtle, and there even a small mistake can leave a web service or web site vulnerable. This is especially important for cloud services, as many of them run as "thin clients" – that is, the heavy lifting is done remotely, and the controls and navigation simply displayed in a web browser window.

Injection Attacks

Injection attacks, cross-site scripting, and buffer overflows are attacks on websites that all exploit a single type of weakness – user input. Sometimes this input can come through comment boxes or text fields, or even through radio buttons and drop down menu, which some consider relatively safe. Even the site's own URL can be used against it, especially if PHP, a popular server-side scripting language, is running on the site's host server.

These attacks can be prevented, however, if extreme care is taken to validate, sanitize, and convert output based on user content into a harmless format. This can be done in a number of ways. First, it's important to check that any input from a web field meets business rules – i.e., that someone has not entered a birth year of 2791, or "Potato." Then the content must be sanitized for malicious content. This generally involves steps like removing all punctuation. Then any output to the site is checked to make sure no harm can be done by sending back the requested information. It may also be possible to set up the programs running on the back end, such as databases, to handle user input in very strictly proscribed ways, so as to further prevent this style of attack.

Web Applications

Web-based applications, in addition to having the potential problems discussed earlier in the paper, can have other vulnerabilities. That's because these applications don't exist on their own, floating in a void, the rest on top of several other programs that run below it, unseen to most outsiders, but attackable by a clever hacker. This means all the programs must work by themselves as well as working together as a whole. A strong patching program and rigorous testing can establish the safety of this kind of application. Like web pages, all applications must also be able to resist malicious or invalid input.

Admin Pages

Many cloud services have an administrator control web page. Access to this page comes with a wide variety of powers that can do a great deal of damage in the wrong hands. In addition to strong authorization, this page and its login screen must be completely protected against any of the attacks listed above. Further restricting access to the admin page by IP address or physical location may also be desirable, if the cloud provider offers that kind of service.

Forced Browsing

To try and find admin pages, attackers may make use of a simple kind of attack called a forced browsing attack. Some website or web service administrators believe that simply not linking to a page and giving it an obscure name means no-one will be able to get to it. This is not actually the case, especially since administrator pages have predictable names, and hacking tools allow very thorough explorations of website structure. This becomes an extremely serious problem in sites with multiple levels of functionality – like customer, employee, manager, admin, etc. In order to keep these attacks from working, each page must be tied to an authentication/authorization level, not simply hidden from view.

Cloud APIs

Many cloud services provide an API, or application programming interface. This allows cloud clients to write programs that can interact directly with the service. APIs are great tools and allow for powerful customization, but they can also be potential weak points in the security of the cloud system. A cloud provider should be able to certify that this interface is free from vulnerabilities, and that using responsibly it will not weaken the security of your cloud deployment.

Malware

Malicious software, or malware, is possibly the greatest security threat to any IT infrastructure today. It can be used to delete, steal, alter, or even hold data ransom. Traditional network security tools, such as firewalls, are of little help against this data-centric threat. No security solution is complete without addressing this menace.

Malware often comes into a system via browsing the Internet or via email. Browsing is particularly germane to a cloud user, considering the same program they're using to access their cloud service could simultaneously be downloading a key logger onto their computer.

Browsing

Disabling browser plug-ins and client side scripts and languages such as ActiveX and JavaScript go a long way towards making browsing a safer experience. However, this may not always be practical. At the very least, these programs should have to ask permission before running.

A good browsing policy – one that restricts traffic from going to high-risk sites from work computers – might also help lower risk. Other measures to consider are having one browser for work and one for play, or restricting leisure browsing to a virtual machine – a simulated computer inside a computer that is largely separate from the environment that hosts it.

Sometimes malware may arrive even through relatively safe sites via malicious advertizing. This attack vector remains a problem because pinning down who exactly is responsible for tackling it has proven tricky. Keeping browsers and Adobe Flash patched and up to date and blocking pop-ups may help. Blocking ads from the ad firms most notorious for providing malicious ads has been shown to dramatically cut down on the number of “malvertisements” companies encounter.

Email

Email is another very common vector for malware. Blocking links in email (or at the very least, masked links), scanning attachments, or banning certain high-risk extensions, such as .exe from being sent as attachments, reduce the risk of email significantly. Banning attachments altogether and using a secure drop box might also be an option.

Content filtering, or filtering out suspect email origins might help weed out spam or email-based attacks before they arrive. Employees should also be trained to recognize deceptive email techniques. Adopting the use of digital signatures in emails may help make telling real messages from fake ones easier.

Despite all these countermeasures, however, it’s still possible that malware may find its way into a network. Luckily, there’s another line of defense. To begin with, simply having anti-malware programs installed on all company devices is an excellent countermeasure. Locking down end users is also helpful – malware has to install itself, and if the user does not have the permission to install things on the computer they are on, the malware cannot install. Keeping track of all programs and changes on company computers, and reviewing this information can also show infections; “whitelisting,” or only giving permission for certain programs to run can go hand-in-hand with this measure. Finally, deep-packet inspection using a device such as an intrusion detection system, or IDS, can identify malware signatures and alert web admins that an infection has occurred.

Social Engineering

The last kind of attack is social engineering, a novel and difficult attack vector. Using this method, an attacker learns about a person or organization, then poses as a friend or colleague to try and extract information. This can be as simple as using an employee’s Facebook page to find the answers to their password recovery question, or to physically show up on the premises, pose as a contractor or bigwig, and start asking for access to restricted systems. This can only really be countered by everyone at a company being aware and cautious at all times.

Cloud providers have to deal with many of these same problems and should be willing to inform you of how they deal with them and how they will work with you to resolve them should they occur.

Shared Environment and Trust

Co-Tenancy

The shared, or co-tenant environment of the cloud is one of its largest security risks. It's difficult to be sure that data sent to the cloud is accessible to you and only to you. There are three possible points of exposure for this data: in transit, at rest, and in use.

Transit

Data in transit has the potential to be intercepted on the way to and from a cloud service, and the potential to be misdirected after it arrives. The solution in both cases is encryption, either through the "https" SSL connection, or, preferably, through a stronger VPN connection. This encryption means that even if someone is listening in, or the data goes to the wrong place, the only thing that will spill out of a secured area is almost pure gibberish to anyone without the encryption key.

Rest

Similarly, data at rest also has the potential to leak over or even to be accessed by other people in the cloud environment. This is especially true with deleted data, which will be covered in more detail later. In addition, it's impossible to know who exactly at the cloud company might be looking at your information without your knowledge or permission. The solution, again, is encryption.

For these data encryptions to work, you must manage the encryption yourself, or select a so-called "zero visibility" cloud service, that encrypts your data for you, but hands over all of the keys. However, potential flaws have been found even in these services when moving data around in the cloud.

Processing

The last place data is at risk in the shared cloud environment is in processing. This is perhaps the most difficult place for an attacker to compromise or for a mistake to be made, which is fortunate because of the dearth of good countermeasures currently on the market. Most operating systems and processors were not designed to keep data separate in processing; those used in most cloud systems are no exception. Some companies, such as Green Hills have designed chips and operating systems specifically for this task, but they have not been widely adopted. Depending on your needs, this is potentially a vulnerability that many cloud clients will simply have to accept.

Metering

The cloud introduces other problems with trust as well. These are often related to how the cloud provider handles your information. For instance, a two-way risk in this regard is the possibility of metering manipulation. Both sides must work to make sure that no attacker or malicious insider is tinkering with the amount owed to the cloud service provider. In addition, procedures should be in place to deal with a scenario in which the cloud client's account is hijacked by an attacker who runs up an enormous bill on the client's account.

Ancillary Data

It's also important for a cloud client to know exactly how the ancillary data they provide to a cloud service is being stored, used and protected. This includes things like payment information, but can also include less obviously sensitive but still revealing information derived from usage habits, or even data mined from the files you have uploaded to their system. Many cloud companies may keep this information private, but asking about how this kind of data is used is still part of a client's due diligence.

Insider Threat

Unavoidably, there's the potential of an insider threat. These can be either malicious, or simply accidental. Using good authorization practices within your own network can help reduce the damage done by an insider, but in the cloud it may not be apparent who a potential inside attack is, what powers they might have, or what damage they could wreak.

Encrypting data sent to and stored in the cloud will largely mitigate the theft of data by a malicious insider, but can a cloud employee accidentally delete your files? In order to understand this threat, it's important to know exactly who will have access to your data and what they can do.

Incident Response

The biggest of all shared responsibilities and trust issues, however, is the matter of incident response. This matter is complex as it is sensitive, as it may well involve multiple incident handling (IH) teams, legal departments, criminal investigations, media entanglements, and legal obligations that span multiple jurisdictions. Besides which, the attack could have any number of different profiles, and be against just the client, the provider, or both at once, further muddying who is responsible for what.

A plan must be put into place before hand, contact channels clearly established, documentation and obligations to share information and inform the other of developing situations clearly delineated. Ideally, both the provider and the client should practice incident response to keep both teams sharp.

The only way these requirements can be defined in a binding way is formally within the Service Level Agreement, or SLA. This document is key for a variety of cloud contractual obligations, but perhaps none more vital than delineating duties before, during, and after a security incident.

Access to the Cloud

Even if your cloud deployment is as secure as can be, that doesn't do you any good if you can't access it. A cloud provider should easily be able to tell you information on their "up-time," or how reliably they're online, but there are other factors consider as well.

DNS

One potential area of attack for either your website or the cloud provider's portal is an attack on your domain name server, or DNS, called cache poising. DNS's take in text Internet addresses – like "www.google.com," and turn them into numeric IP addresses computers can understand. A cache poisoning attack will cause a DNS server to route traffic to the wrong site, which can be loaded with malware or attempt to steal a user's authentication credentials or other information.

There are a number of technical countermeasures a DNS can take, but most small companies do not run their own DNS servers, which means these measures may be out of your hands. Running the most recent version of the Berkeley Internet Name Domain (BIND) system, as well as using tools like DNSSec, can help a DNS resist this kind of attack. Clients can check a websites security certificate or use a VPN connection to ensure they're really at the right place.

DNS servers, individual websites, and even entire Internet service providers are vulnerable to Denial of Service and Distributed Denial of Service attacks. These attacks are designed to force a service offline, and can target everything from hardware to applications with the intent of consuming bandwidth, processing power, or service space until something breaks. Attacks can be recognized via signature similar to malware, or by unusual activity.

DDoS and ISPs

DDoS attacks, which are among the most famous and common of these attacks, simply try to overwhelm an internet site or service with meaningless web traffic from an army of malware-infected computers until the site crashes. The goal of the entity being attacked is to sort out the good traffic from the bad. This can be done by shutting down while the attack is happening, shifting to a new IP address to try and buy some time, or by trying to filter out the bad traffic from the good on the fly. Some cloud providers and ISPs offer this last option as a premium service.

Properly secured ISPs and DNS servers make this kind of attack much harder to accomplish. Your cloud provider should definitely have a plan to deal with this kind of attack, which may involve a special arrangement with their ISP.

From the client side, simply having two ISPs in case one fails is a relatively good way of mitigating this threat.

Data Deletion, Preservation, and Storage

One of the biggest problems with storing data in the cloud is that it's not always clear where the data is being stored. If the cloud company is contracting with another cloud company for this service, it's possible they might not even know themselves. However, this is a vital piece of risk-management and policy information.

Location

All cloud adopters should ask if:

1. The cloud company will allow them to chose where data is stored
2. If not, will they inform the client where the data ends up?

If they will not give you either of these pieces of information, be wary.

The reason you should be cautious is because privacy protections and other data laws may vary widely from jurisdiction to jurisdiction. If you do not know where your data is, it may be impossible to tell what your obligations are, or what your provider is legally allowed to do with your data. For instance it's illegal to send certain kinds personal data over international borders in the EU. Once the data ends up there, removing it again may be impossible.

Data Remanence

Deletion and preservation of data is very important for two main reasons. The first is pure security. Data remanence, when saved data is not entirely erased from memory, is a huge risk the shared environment of a cloud provider. How can they guarantee that they have truly destroyed the information you entrusted to them?

eDiscovery

The other is legal. eDiscovery is the obligation of a company to turn over electric documents during legal action. To reduce the burden of this activity, which can involve huge amounts of data, a company should have a robust data retention and deletion policy that reacts to business needs. For instance, instead of deleting all files after a set amount of time, the completion of a project should trigger the review of the documents it generated. The easiest way to do this is with a centralized system that can support these policies.

When a lawsuit is pending, the client must quickly inform all of their IT staff, including their cloud provider, who must be able to protect data from further alteration and deletion. Actually retrieving the data involves the ability to provide an audit trail for all documents, provide data in its original format, and preserve file metadata. It may even involve forensic analysis of the memory medium.

Can your cloud service provide this? You should hope so. Failure to comply with eDiscovery can result in your entire database being opened to the opposing party, huge fines, or even the reversal of the burden of proof. Since any company can be sued, everyone should take this very seriously.

Conclusion

Although cloud adoption is not simple, nor is it for everyone, it's important to keep in mind, that, done right, it can save your company money, allow an increased focus on your core product or service, and increase security all that the same time. The important thing is to plan ahead, keep security in mind, ask the important questions, and get everything in writing. Despite the fact that shared responsibilities can make things a bit complicated, overcoming this obstacle is both possible, and often times a profitable undertaking

References

A guide to network vulnerability management. (2012, August 9). *Dark Reading*. Retrieved from http://www.darkreading.com/attacks-breaches/a-guide-to-network-vulnerability-management/d/d-id/1138344?

Ashford, W. (2013, August 2). Cloud service providers often not set up for incident response. *Computer Weekly*. Retrieved from http://www.computerweekly.com/news/2240203007/Cloud-service-providers-often-not-set-up-for-incident-response

AT&T Distributed Denial of Service (DDoS) Defense. (2014). *AT&T Product Brief*. Retrieved from http://www.business.att.com/content/productbrochures/ddos_prodbrief.pdf

Authorization and permissions in SQL servers. (n.d.). *Microsoft developer network*. Retrieved from http://msdn.microsoft.com/en-us/library/bb669084(v=vs.110).aspx

Badger, L., Bohn, R, Chu, S., Hogan, M., Liu, F., et al. (2011, November 1) Useful information for cloud adopters. *US Government Cloud Computing Technology Roadmap, Volume II, Release 1.0 (Draft)*. National Institute of Standards and Technology. Retrieved from http://www.nist.gov/itl/cloud/upload/SP_500_293_volumeII.pdf

Balding, C. (n.d.). Cloud storage. Cloudsecurity.org. Retrieved from http://cloudsecurity.org/cloud-storage.html

Basic security practices for web application. (n.c.) . Microsoft. Retrieved from http://msdn.microsoft.com/en-us/library/vstudio/zdh19h94(v=vs.100).aspx

Barker, E., Barker, W., Burr, W., Polk, W., and Smid, M. (2012, July). Recommendation for key management—Part 1. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57_part1_rev3_general.pdf

Beckham, J. (2011, May 3). The top 5 security risks of cloud computing. Retrieved from http://blogs.cisco.com/smallbusiness/the-top-5-security-risks-of-cloud-computing/

Bird, J., & Manico, J. (2014, April 7). Attack surface analysis cheat sheet. *OWASP*. Retrieved from https://www.owasp.org/index.php/Attack_Surface_Analysis_Cheat_Sheet

Bresnick, J. (2013, July 25). Does cloud computing threaten patient privacy, data security? Health IT Security. Retrieved from http://healthitsecurity.com/2013/07/25/does-cloud-computing-threaten-patient-privacy-data-security/

Brute force attacks. (n.d.). Wordpress.org Retrieved from http://codex.wordpress.org/Brute_Force_Attacks

Butler, B. (2014, April 21) Even the most secure cloud storage may not be so secure, study finds. Network World. Retrieved from http://www.networkworld.com/article/2176237/cloud-computing/even-the-most-secure-cloud-storage-may-not-be-so-secure--study-finds--.html

BYOD-Bring your own device. (n.d.). Dell.com. Retrieved from http://www.dell.com/learn/us/en/555/mobility-byod

Cloud top ten security risks. (2014, January 23). *OWASP*. Retrieved from https://www.owasp.org/index.php/category:owasp_cloud_%e2%80%90_10_project

Chaturvedi, P., Gupta, K. (2013, April). Detection and prevention of various types of jamming attacks in wireless networks. *international journal of computer networks and wireless communications* (IJCNWC), ISSN: 2250-3501 Vol.3, No2. Retrieved from http://www.ijcnwc.org/papers/vol3no22013/3vol3no2.pdf

Cohen, R. (2009, March 1). Navigating the fog -- billing, metering & measuring the cloud. *Ruv.net*. Retreived from http://www.elasticvapor.com/2009/03/navigating-fog-billing-metering.html

Collins, H. (2010, September 20). Top 10 network security threats. *Top 10 Network Threats*. Retrieved June 12, 2014, from http://www.govtech.com/security/Top-10-Network-Security-Threats.html

Coty, S. (2014, February 8). Computer forensics and incident response in the cloud. *RSA Conference 2014*. Retrieved from http://www.rsaconference.com/writable/presentations/file_upload/anf-t07a-computer-forensics-and-incident-response-in-the-cloud.pdf

CPU tech secure processor solutions built with INTEGRITY. (2010, March 3). *Green Hills Software*. Retrieved from http://www.ghs.com/news/20100303_CPU_tech.html

Curtis, W., Mew C. (2008, February 28). Preparing for E-discovery. *NACUA Notes, 6 (2)*. Retrieved from http://www.ncsu.edu/general_counsel/legal_topics/documents/PreparingforE-Discovery.pdf

Data Validation. (2013, December). *OWASP* Retrieved from https://www.owasp.org/index.php/Data_Validation#Best_Method

Di Bello, A. (N.D). How cloud computing changes incident response. *Guidance Software*. Retrieved from http://endpoint-intelligence.blogspot.com/2012/07/how-cloud-computing-changes-incident.html

Dirking, B. Kodali, K. (2008, June). Strategies for preparing for E-discovery. *The Information Management Journal*. Retrieved from http://www.oracle.com/us/products/middleware/content-management/059431.pdf

DNS cahce poisoning: The next generation. (2007, August 13). Dell SecureWorks. Retrieved from http://www.secureworks.com/resources/articles/other_articles/dns-cache-poisoning/

Dou, W., Chen, Qi., Chen, J. (2013, September). A confidence-based filtering method for DDoS attack defense in cloud environment. *Future Generation Computer Systems, 29 (7)*. Retrieved from http://www.sciencedirect.com/science/article/pii/S0167739X12002312

Ducklin, P. (2013, August 16). Anatomy of a brute force attack - how important is password complexity? Sophos: Naked Security. Retrieved from http://nakedsecurity.sophos.com/2013/08/16/anatomy-of-a-brute-force-attack-how-important-is-password-complexity/

Eisner, D. (2014, April 14). Overcoming cloud storage security concerns: Seven key steps. CircleID. Retrieved from http://www.circleid.com/posts/20140425_overcoming_cloud_storage_security_concerns_7_key_steps/

Emigh, J. (2013, September 27) How to secure your laptop PC. Notebookreview.com Retrieved from http://www.notebookreview.com/howto/how-to-secure-your-laptop-pc/

Gabriel, C. (2013, January 21). No BYOD policy? Time to grasp the nettle. CXO Unplugged. Retrieved from http://cxounplugged.com/2013/01/byod-policy/

GadAllah, S. (2003, December 30). The importance of logging and traffic monitoring for information security. SANS Institute InfoSec Reading Room. Retrieved from http://www.sans.org/reading-room/whitepapers/logging/importance-logging-traffic-monitoring-information-security-1379

Grobauer B., Walloschek T., & Stocker E. (2011, August 15). Understanding cloud computing vulnerabilities. Retrieved from www.infoq.com/articles/ieee-cloud-computing-vulnerabilities

Halley, B. (2008, October 20). How DNS cache poisoning works. NetworkWorld.com. Retrieved from http://www.networkworld.com/article/2277316/tech-primers/how-dns-cache-poisoning-works.html

Harbert, T. (2012, April 23). E-discovery in the cloud? Not so easy. *Computerworld*. Retrieved from http://www.computerworld.com/s/article/9226375/E_discovery_in_the_Cloud?taxonomyId=19&pageNumber=1

Honorof, M. (2013, July 29) How to secure your cloud storage. *Tom's Guide*. Retrieved from http://www.tomsguide.com/us/how-to-secure-cloud-storage,review-1799.html

How to secure your TCP/IP ports. (2009, September 9). PC Plus. Retrieved from http://www.techradar.com/us/news/networking/how-to-secure-your-tcp-ip-ports-633089/3%22%20/l%20%22articleContent

Hubbard, D., & Sutton, M. (2010, March).Top threats to cloud computing. Retrieved from https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf

Hurst, D. (2009, October 29). Cloud security and it's (sic) effect on application security. *OWASP*. Retrieved from https://www.owasp.org/images/a/a6/understanding_the_implications_of_cloud_computing_on_application_security-dennis_hurst.pdf

Information Supplement: PCI DSS Cloud computing guidelines. (2014, February). PCI Security Standards Council. Retrieved from https://www.pcisecuritystandards.org/pdfs/PCI_DSS_v2_Cloud_Guidelines.pdf

Information supplement: PCI DSS Virtualization guidelines. (2011, June 1). PCI Security Standards Council. Retrieved from https://www.pcisecuritystandards.org/documents/Virtualization_InfoSupp_v2.pdf

Jansen, W., Grance, T. (2011, December). Guidelines on security and privacy in public cloud computing. National Institute of Standards and Technology. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf

Kent, K., Souppaya, M. (2006, September). Guide to computer security log management. National Institute of Standards and Technology. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf

Kesan, J., Hayes, C., Bashir. (2013). Information privacy and data control in cloud computing: consumers, privacy preferences, and market efficiency. *Washington and Lee Law Review 70 (1)*. http://scholarlycommons.law.wlu.edu/wlulr/vol70/iss1/6/

Kontzer, T. (2013, December 13). Tips for preventing data center outages. Networkcomputing.com. Retrieved from http://www.networkcomputing.com/data-centers/tips-for-preventing-data-center-outages/d/d-id/1234588?

Krig, D., Lee, R. (2001, October). Remote Denial of Service Attacks and Countermeasures. *Princeton University Department of Electrical Engineering Technical Report*. Retrieved from http://www.princeton.edu/~rblee/ELE572Papers/karig01DoS.pdf?q=tilde/rblee/ELE572Papers/karig01DoS.pdf

Laptop computer security policy (n.d.). Internal revenue service manual. Retrieved from http://www.irs.gov/irm/part10/irm_10-008-026.html#d0e245

Lemos, R. (2013, November 18). Enterprises should practice for cloud security breaches. *Darkreading.com*. http://www.darkreading.com/risk/enterprises-should-practice-for-cloud-security-breaches/d/d-id/1140912?

Los, R., Shackleford, D., & Sullivan, B. (2013, February). *The Notorious nine: cloud computing top threats in 2013*. Retrieved from https://downloads.cloudsecurityalliance.org/initiatives/top_threats/The_Notorious_Nine_Cloud_Computing_Top_Threats_in_2013.pdf

Mell, P., Bergeron, T., Henning, D. (2005, November). Creating a patch and vulnerability management program. *National Institute of Standards and Technology*. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-40-Ver2/SP800-40v2.pdf

Mell, P., Grance, T., (2011, September). The NIST definition of cloud computing. *National Institute of Standards and Technology*. Retrieved from http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Mogull, R. (2011, March). Incident response process in a cloud environment. *TechTarget SearchCloudSecurity*. Retrieved from http://searchcloudsecurity.techtarget.com/tip/Incident-response-process-in-a-cloud-environment

Morrisdale, P. A. (2005, September 1). The six dumbest ideas in computer security. Retrieved from http://www.ranum.com/security/computer_security/editorials/dumb/

Moyle, E. (2011, March). Meeting the PCI requirement for Web security in the cloud. TechTarget, SearchCloudSecurity. Retrived from http://searchcloudsecurity.techtarget.com/tip/Meeting-the-PCI-requirement-for-Web-security-in-the-cloud

Multi-factor authentication introduction. (n.d.) SafeNet, Inc. Retrieved from http://www.safenet-inc.com/multi-factor-authentication/

Munsch, D., Lerchey, J. (2010, April 27). Electronic discovery: Litigation holds, data preservation and production. Retrieved from http://www.cmu.edu/iso/aware/presentation/legal_hold.pdf

Network and Information Security Standards Report. (2007, May 11). European Committee for Standardization (CEN). Retrieved from https://web.archive.org/web/20130803074337/http://www.cen.eu/cen/Sectors/Sectors/ISSS/Activity/Pages/NISSG%20Report%20Table%20of%20Content.aspx

Olzack, T. (2012, January 30). Building the foundation: Architecture design, chapter 3*Enterprise security: A practitioner's guide*. InfoSec Institute. Retrieved from http://resources.infosecinstitute.com/architecture-design-chapter-3/

Olzack, T. (2012, February 17). Attack surface reduction : Architecture design, chapter 4. *Enterprise security: A practitioner's guide*. InfoSec Institute. Retrieved from http://resources.infosecinstitute.com/attack-surface-reduction/

Overview of cyber vulnerabilities. (n.d.). *ICS-CERT*. Retrieved June 12, 2014, from http://ics-cert.us-cert.gov/content/overview-cyber-vulnerabilities

Parkinson, A. (n.d.). Federal information laws: does your business cloud storage comply? Top Ten Reviews. Retrieved from http://business-cloud-storage-services.toptenreviews.com/federal-information-laws-does-your-business-cloud-storage-comply-.html

Patch Management. (2008, February). The government of the Hong Kong Special Administrative Region. Retrieved from http://www.infosec.gov.hk/english/technical/files/patch.pdf

Phifer, L. (2013, Jan 28) BYOD security strategies: Balancing BYOD risks and rewards. Search Security. Retrieved from http://searchsecurity.techtarget.com/feature/BYOD-security-strategies-Balancing-BYOD-risks-and-rewards

Phneah, E. (2013, February 4). Five security risks of moving data in BYOD era. CBS Interactive. Retrieved from http://www.zdnet.com/five-security-risks-of-moving-data-in-byod-era-7000010665/

Pinzon, Scott. (2002) Foundations: What is a port? (and why should I block it?) WatchGuard Technologies, Inc. Retrieved from https://www.watchguard.com/infocenter/editorial/135090.asp

PCI PSS Quick Reference Guide. (2010, October). PCI Security Standards Council. Retrieved from https://www.pcisecuritystandards.org/documents/PCI%20SSC%20Quick%20Reference%20Guide.pdf

Presti, K. (2012, June 28). Seven security threats circling your network. *CRN*. Retrieved from http://www.crn.com/slide-shows/security/240002785/7-security-threats-circling-your-network.htm

Rains, T. (2014, February 4). Threats in the cloud, part 1: DNS attacks. *Microsoft Security Blogs*. Retrieved from http://blogs.technet.com/b/security/archive/2014/02/04/threats-in-the-cloud-part-1-dns-attacks.aspx

Rains, T. (2014, February 6). Threats in the cloud, part 2: Distributed denial of service attacks. *Microsoft Security Blog*. Retrieved from http://blogs.technet.com/b/security/archive/2014/02/06/threats-in-the-cloud-part-2-distributed-denial-of-service-attacks.aspx

Reed, J. (2010, September 10). Following Incidents into the Cloud *SANS Institute InfoSec Reading Room*. Retrieved from http://www.sans.org/reading-room/whitepapers/incident/incidents-cloud-33619\

Richmond, R. (2010, May 19). Five ways to keep online criminals at bay. *The New York Times*. Retrieved from http://www.nytimes.com/2010/05/20/technology/personaltech/20basics.html?_r=4&.

Risks. (2011, November 4) *Cloudcontrols.org*. Retrieved from http://www.cloudcontrols.org/cloudcontrols/risks/

Roman, J. (2013, May 21). Safeguarding ISPs from DDos attacks. Information Security Media Group, Corp. Retrieved from http://www.bankinfosecurity.com/isp-security-needs-to-be-improved-a-5773/op-1

Rothstein, B., Hedges, R., Wiggins, E. (2007). Managing discovery of electronic information: A pocket guide for judges. *Federal Judicial Center*. Retrieved from http://www.fjc.gov/public/pdf.nsf/lookup/eldscpkt.pdf/$file/eldscpkt.pdf

Rouse, M. (2005, September). Cache poisoning (domain name system poisoning or dns cache poisoning). Search Security. Retrieved from http://searchsecurity.techtarget.com/definition/cache-poisoning

Rubens, P. (2012, April 9) Four steps to securing Mobile devices and apps in the workplace. . *eSecurity Planet*. Retrieved from http://www.esecurityplanet.com/mobile-security/4-steps-to-securing-mobile-devices-and-apps-in-the-workplace-mdm-byod.html

Rubens, P. (2013, March 25). Six emerging security threats, and how to fight them. *eSecurity Planet*. Retrieved from http://www.esecurityplanet.com/network-security/6-emerging-security-threats-and-how-to-fight-them.html

Rubens, P. (2013, May 2). Ten steps you can take to secure a laptop. Techradar.com Retrieved from http://www.techradar.com/us/news/mobile-computing/laptops/10-ways-to-secure-a-laptop-1148348/1

Ryder, J. (2001, July 30). Laptop security, part one. Symantec. Retrieved from http://www.symantec.com/connect/articles/laptop-security-part-one-preventing-laptop-theft

Sambasivam, S. (2008, June) On the road to E-discovery compliance. *International Auditor*. Retrieved from http://www.theiia.org/intAuditor/itaudit/archives/2008/june/on-the-road-to-e-discovery-compliance/

Savage, M. (2011, March 27). PCI DSS compliant cloud providers: no PCI panacea. *TechTarget, SearchCloudSecurity*. Retreived from http://searchcloudsecurity.techtarget.com/news/2240033583/PCI-DSS-compliant-cloud-providers-No-PCI-panacea

Salted password hashing - doing it right. (2014, February 24). CrackStation.net. Retrieved from https://crackstation.net/hashing-security.htm

Savitz, E. (2012, January 19). Storing data in the cloud raises compliance challenges. Forbes.com. Retrieved from http://www.forbes.com/sites/ciocentral/2012/01/19/storing-data-in-the-cloud-raises-compliance-challenges/

Securing your network with firewalls and ports. (2008). Microsoft developer network. Retrieved from http://msdn.microsoft.com/en-us/library/ms864793.aspx

Security incident management. (2014). *Cloud.cio.gov*. Retrieved from http://cloud.cio.gov/topics/security-incident-management

Shapland, R. (2012, March 1). Forced browsing: Understanding and halting simple browser attacks. Computerweekly.com. Retrieved from http://www.computerweekly.com/answer/Forced-browsing-Understanding-and-halting-simple-browser-attacks

Souppaya, M., Scarfone, K. (2013, July). Guide to enterprise patch management technologies. National Institute of Standards and Technology. Retrieved from http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-40r3.pdf

Souppaya, M., Scarfone, K. (2013, July). Guide to malware incident prevention and handling for desktops and laptops. National Institute of Standards and Technology. Retrieved from http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-83r1.pdf

Strom, D. (2007, November 18). Ten tips to secure your laptop. Informationweek.com. Retrieved from http://www.informationweek.com/10-tips-to-secure-your-laptop/d/d-id/1061655?

Strong Authentication Best Practices. (n.d.) SafeNet, Inc. Retrieved from http://www.safenet-inc.com/multi-factor-authentication/strong-authentication-best-practices/

Stults, G. (2004, May 9). An overview of Sarbanes-Oxley for the information security professional. SANS Institute InfoSec Reading Room. Retrieved from https://www.sans.org/reading-room/whitepapers/legal/overview-sarbanes-oxley-information-security-professional-1426

Sullivan, D. (2006). Malware: The ever-evolving threat. *The shortcut guide to protecting business internet usage*. Retrieved from http://searchsecurity.techtarget.com/feature/Malware-The-ever-evolving-threat

Tellis, P. (2011, January 11) Keeping web users safe by sanitizing input data. Smashing Magazine. Retrieved from http://www.smashingmagazine.com/2011/01/11/keeping-web-users-safe-by-sanitizing-input-data/

Threat analysis of cloud services (initial thoughts for discussion). (n.d.). Collaborate.nist.gov. Retrieved from http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/CloudSecurity/Threat_Analysis_of_Cloud_Services.pdf

Threat sources by cloud architecture component. (n.d.). Collaborate.nist.gov. Retrieved from http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/CloudSecurity/Cloud_Threats_by_Sources.pdf

Top 20 security controls. (n.d.). Collaborate.nist.gov. Retrieved from http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/CloudSecurity/Top_20_Security_Controls-excerpt.pdf

The three methods an ISP uses to defend against DoS and DDoS. (2013, July 17). DDoSAttacks.biz. Retrieved from http://www.ddosattacks.biz/protection/the-three-methods-an-isp-uses-to-defend-against-dos-and-ddos/

Wang, C. (2010, October 29). *Q&A: Demystifying Cloud Security*. Retrieved from http://resources.idgenterprise.com/original/AST-0036145_G2A_demystifying_cloud_security.pdf

Wang, W. (2013, January 10). Five checks you must run to ensure your netweork is secure. The Hacker News. Retrieved from http://thehackernews.com/2013/01/5-checks-you-must-run-to-ensure-your.html

Watson, C., Keary, E., Fitzgerald, A. (2014, April 7). Logging cheat sheet. *OWASP*. Retrieved from https://www.owasp.org/index.php/Logging_Cheat_Sheet

Weiss, A. (2012, October 28). Prevent web attacks using input sanitation. eSecurity Planet. Retrieved from http://www.esecurityplanet.com/browser-security/prevent-web-attacks-using-input-sanitization.html

Zeltser, L. (2011, June 13). Malvertising: The use of malicious ads to install malware. Infosec Island. http://www.infosecisland.com/blogview/14371-Malvertising-The-Use-of-Malicious-Ads-to-Install-Malware.html