Nexo and the Cyber War

November 2016

According to Verizon's annual *Data Breach Investigations* reports, there were more than 3,100 reported data theft events last year, compared to 700 just four years ago. And this doesn't count undetected events. To paraphrase John Chambers of Cisco, "There are two kinds of companies - those that know they've been hacked, and those that don't."

The core of this pervasive, global problem is the overwhelming advantage enjoyed by attackers. They can throw tens of thousands of attacks at the same target over a very short timeframe, and *only one* needs to succeed to produce a full breach or to commandeer a key system. Put another way, if a corporation or cloud service is attacked 1,000,000 times, they have to stop all 1,000,000 of them to remain secure.

This paper describes how malware and unauthorized users have been getting into to our devices, applications and networks – and how a new generation of secure file sharing can flip the odds back in our favor.

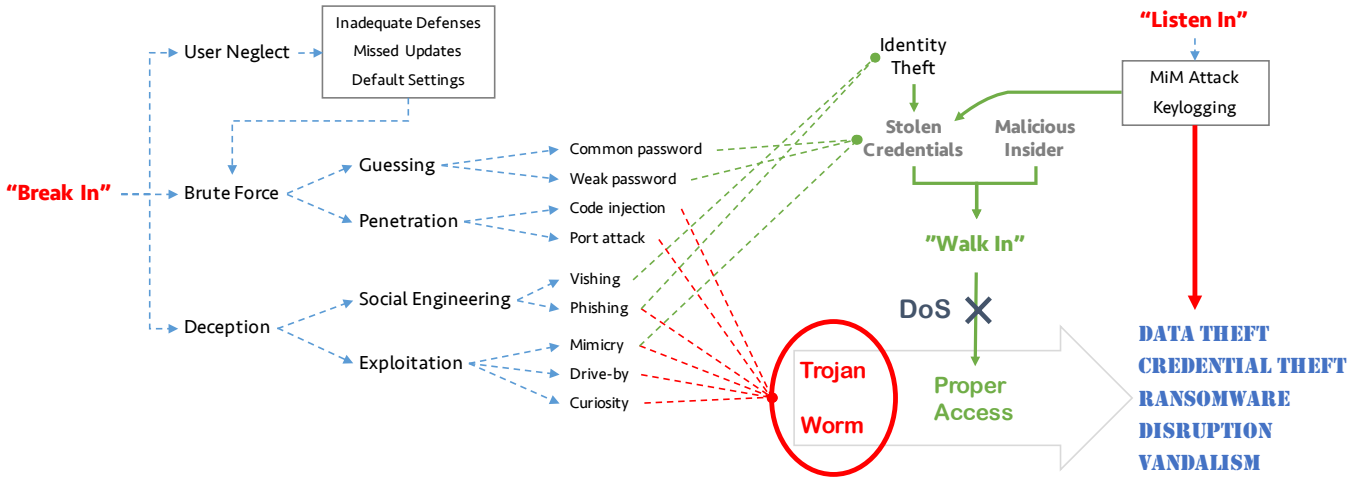

What's Happening Today

The graph above categorizes most of the popular tactics used by hackers to accomplish their missions (the most common missions are in blue at right). A simple way to explain these tactics is to compare them to how a burglar might rob a house. At the highest level, he has three options:

- Walking In. If a house is left open or he has a copy of the keys, getting inside is easy. (Even if there is an alarm, he might know its code or how to bypass it.) This is clearly the most desirable method, as it doesn't look any different than a routine home entry. And that's just the problem: from the system's perspective, a hacker logging in with the proper credentials looks legitimate, and can perform permitted tasks without raising an eyebrow.

- Breaking In. Though "walking in" is more common than one would expect (see "User Neglect" in the chart), most businesses have learned to take measures to protect themselves. Now, our burglar finds all doors and windows locked, the alarm system armed and surveillance cameras everywhere – he has to break in. First, he will probe every inch of its perimeter looking for a weakness. Failing that, he can use any number of tricks to fool the homeowner into granting him access. Once inside, he grabs a spare set of keys and can "walk in" at his pleasure.[1]

[1] In 2014, hackers stole 60 million ID/password combinations from Dropbox. This massive credential theft was not detected, and was only discovered two years later after they were published on the Web.

- Taking Hostages. A third option for our burglar is to extort the homeowners – coercing them into meeting his demands by threatening their home or their livelihoods. One way is to capture the house itself – changing the locks or physically barring entry altogether. On the Internet, this is done through a "denial of service" (DoS) attack.[2] Another is to take someone hostage or make the home unlivable until a ransom is paid. Called "ransomware," this increasingly popular type of malware seizes control of a computer, application, or critical data and does not give it back until a hefty ransom is paid.

Our home invasion analogy may be disturbing or sound extreme. Unfortunately, it is a 24/7 reality for businesses – especially for large corporations, where the reward is commensurate with their size. And now, it is becoming a reality for individuals like you and me. Yet we continue to use 40-year-old technologies like email, which is widely known to be unsecure. We also share files containing sensitive or personal information using vulnerable, decade-old public clouds like Dropbox or Google Drive.

Why? For one, we are creatures of habit and are reluctant to change. But the real barriers to acceptance have been, a) the trade-off between security and ease of use, and b) a lack of motivation, because we perceive that we're safe due to the lack of evidence to the contrary. With the number of personal horror stories about email breaches, ransomware and identity theft on the rise, the proper motivation is finally emerging. But ease of use is a tougher challenge. For a new generation of information sharing software to take hold, it must deliver a step-function increase in security without burdening the user experience. (This is precisely one of our core goals at Nexo.)

What Happens During a Breach

The term "breach" is evocative of medieval battles, where after a fashion, the attacker manages to gain entry into a castle by bashing through its gate or walls. In information security, it most often takes the form of an implant (malicious code that executes unnoticed inside a secure perimeter) or an impostor (appearing to the system as a legitimate user). We'll review examples of impersonation via stolen credentials later in this paper. You can imagine the damage that can be done by a bad guy with, say, system-level administrator access. They can query databases, copy files, download messages, and cripple equipment and software before countermeasures can be taken.

Implant-style breaches typically involve the introduction of malware in the form of trojans (aptly-named per our castle analogy) or worms.[3] While trojans can hide in almost any kind of computing device, they most often take up residence in servers. They are nefarious in that they can replace system-level code and delete their original form, making the compromise not only hard to detect but difficult to diagnose as well.

At their most basic, servers are souped-up computers designed to handle lots of concurrent users. They are typically specialized for a particular task like processing, storage, databases or communications. They mostly run Windows or Linux – the former favored more by corporations, and the latter by Internet companies – and are extensively interconnected by an internal network (the "Intranet").

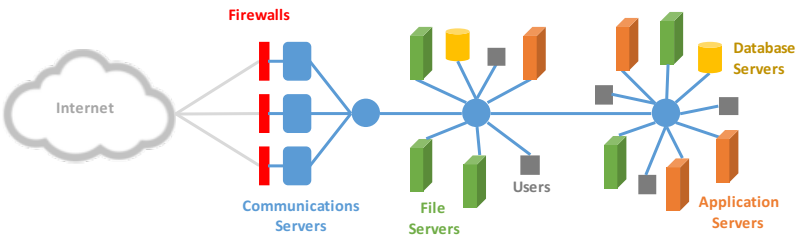

A small company might just have one multitasking server. Larger businesses have banks of hundreds of special-purpose servers in one or more data centers. And a Cloud is essentially a souped-up data center, serving multiple customers. In general, all three of these scenarios employ the same operating systems, applications and similar equipment, and are thus susceptible to the same types of malware.

[2] In October 2016, cyber-criminals launched a distributed denial of service attack (DDoS) on Dyn, a key provider of data that translates URLs into numeric addresses for routing users to websites. By infecting millions of devices, it caused them to simultaneously flood Dyn's servers with traffic, tying them up and preventing millions of legitimate users from accessing major sites, in some cases, for days.

[3] The terms trojan and worm are often used interchangeably. "Virus" is becoming less frequently used.

IT has traditionally invested in a perimeter defense system, where firewalls "wall off" the internal network from the outside world at every point they interconnect. More recently, security software has turned its attention inward to Intranet-connected devices like servers, user devices and communications gear, in an attempt to detect malware that has gotten past the perimeter. Unfortunately, as evidenced by the exponential rise in breaches, hackers have become so sophisticated that they regularly defeat one or more of these measures.

An undetected trojan, then, can grant itself access to any number of systems, servers and applications. It can monitor systems that manage user IDs, email, chat, transactions and databases. It can then open up an outbound connection and transmit the data back to the perpetrators.

Worms are like trojans except that they rapidly self-propagate to other servers and devices, including end user computers. Left to their own devices, worms can so thoroughly compromise an operation that nearly every system can be accessed, monitored and reported out.[4]

Criminal Motivation

The ubiquity of the Internet shifted many applications from the desktop to the cloud (that's certainly the case with file sharing). Unfortunately, public cloud services are incredibly enticing to cyber-criminals because they've *gathered millions of users at one attack point*. Before the cloud, bad guys had to go after each target one at a time, with uncertain results. But getting access to millions of cloud users and their information – for a modest incremental effort – seriously upgraded their motivation.

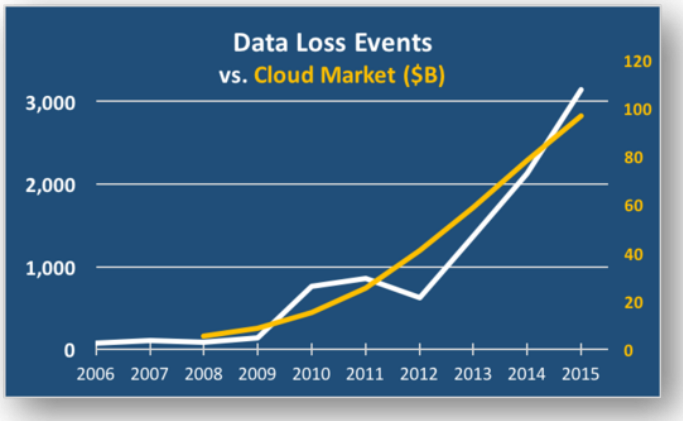

The white line in the chart (right) shows the alarming increase in data theft over the past 10 years.[5] "Data theft" involves cyber-criminals extracting information and/or login credentials from cloud or corporate users. Next to it, the orange line shows the growth of the cloud computing market. The conclusion is inescapable: *cyber-crime was begotten from – and is fueled by – the success of cloud computing*.

Nexo: Compared to the top file sharing services, Nexo is unique because it does not store anything in the cloud. When a Nexo user shares files or chats with other users, all of the data flows directly between the participants' devices. There is no intermediate storage and no vulnerable, central repository of user data and files. In fact, with Nexo, the largest target presented to a hacker is a single user.

Cloud Vulnerability

It stands to reason, then, that many of the tactics used by cyber-criminals are aimed at clouds and other large aggregations of user data. Clouds are basically oversized corporate data centers, except that 100% of their users externally access them via the Internet. While this section focuses on cloud vulnerability, corporate data centers fall prey to most of the same tactics.

[4] In June, 2016, it was reported that Democratic National Committee (DNC) systems had been breached by two groups of Russian hackers. DNC officials told the Washington Post, "the intruders so thoroughly compromised the DNC's system that they also were able to read all email and chat traffic" and also gained access to research files.

[5] *Verizon Data Breach Investigations*, years 2007-2016. (A given year's report covers the previous year's activity.)

What is a cloud, anyway?

A cloud is a concentration of computing resources like processing, storage and applications, that are provided as a service to one or more customers. Unlike in-house facilities, where capacity is determined by physical equipment (e.g., servers), clouds outsource the equivalent computing environment, but one that can more readily scale up or down as needs change. A public cloud is accessible by any Internet user with the proper credentials, whereas a private cloud is sectioned off for exclusive use by one customer.

If you were to visit a cloud facility, it would look similar to a very large corporate data center – you'd see servers, communications gear and storage. The big difference is that it's physically separated from the customer and users, and is accessed via the Internet. It is that *separation* between the user and the computer that created a new entry point to vast collections of information. And it provoked an arms race – and all-out war – with malevolent hackers.

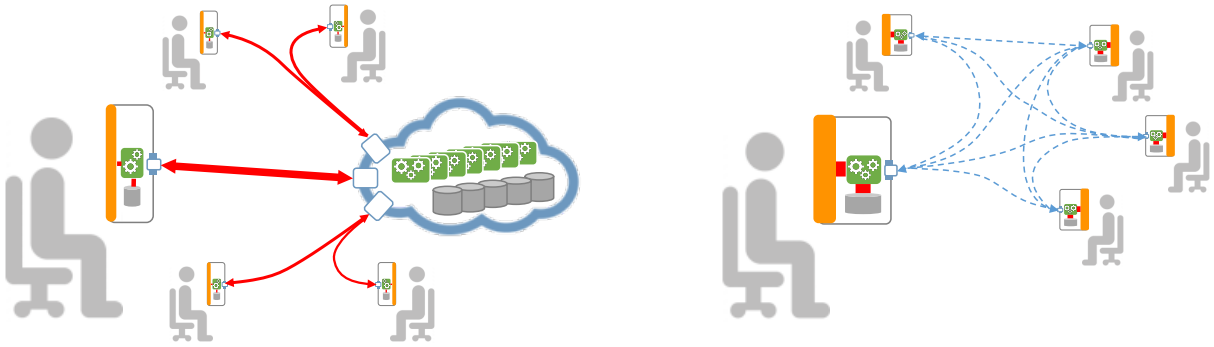

Nexo: The diagram above shows a cloud scenario (left) and peer-to-peer (right). The cloud user runs a browser (thin client) that connects to the cloud through the Internet for processing and storage (red lines). The cloud connects to the public Internet via multiple ports. The P2P user runs an application (thick client) that interacts with its own processor and storage. These are internal to the device and not publicly accessible. Files, messages and chat are then shared directly between the users via a mesh network that disappears after each session is complete. There is so much processing power and storage on our own devices nowadays that there is no need to introduce expensive, vulnerable infrastructure on top of what we already have.

Information Persistence

One of the problems with cheap, plentiful storage is the tendency to save information far beyond its useful life. For example, many email users have a habit of retaining their deleted messages (with attachments) in a trash folder. It isn't limited to email, either, as file directories behave the same way. The lifespan of these dead files is then further extended – practically indefinitely – when they're picked up by a backup cycle.

Because there are so many files, it's impossible for users to distinguish those with sensitive information from those without. As a result, an extreme quantity of data is unnecessarily sitting around either in clouds or corporate servers, and is vulnerable to seizure by hackers.

Nexo: When you share a file using today's leading services, you relinquish full control of the file and its contents to every recipient. They are free to save, copy, print or forward each file. Except in the minority of cases where groups are jointly creating a document, collaborators should only need to *view* each other's files. Nexo defaults to this sharing mode, so that you have to explicitly provide full access to a shared file. With our "screen-only" feature, control over a file's contents remains fully in the hands of its owner. You can also limit the number of times a file is opened, restrict it to geographic areas and timeframes, and set conditions for when the file is permanently deleted from a given recipient's device.

Penetration Attacks

With the Internet as the primary source of access for users and devices outside of the firewall, access points (called "ports") are required to provide entry. Hackers carefully analyze the code and behavior of commercially available routers to find weaknesses and security work-arounds within their ports. When surveilling a target, they detect the brand and model used and launch the appropriate attack. The objective is to deposit malware inside the target's perimeter defenses so that it can autonomously carry out its mission.

These attacks are not limited to access ports. Other approaches include "injecting" malware into browser code or database commands. As vendors learn about successful hacker exploits, they hurriedly modify their code and distribute it to customers via updates or "patches." But this takes time. Meanwhile, hackers share the successful new strategy with each other, and launch thousands of new attacks before the window closes.

The onus of keeping these patches up-to-date is on the customer, and invariably, some devices are missed and leave the door wide open.

Nexo: Unlike remotely accessible networks or clouds, your computer and mobile devices are not set up for remote access. Windows, Mac OS and others ship with strong firewalls that don't permit inbound traffic unless an application, like a browser or Nexo, controls the connection. And Nexo has robust node authentication so that you are only sharing with verified users. It's highly unlikely that a single user is targeted, but if they are, your OS (with automatic updates enabled) will prevent any attack.

Password Guessing

In some cases, devices are installed without changing the generic factory-set password (e.g., "password") or with weak passwords that can be guessed using brute force. Now, the hacker can log into a device with administrative privileges, and wreak havoc on the inside. Many cloud services limit the number of incorrect password entries – but not all. And some devices (like Wi-Fi access points) have no limit by default.

The typical Internet user has to manage dozens of ID/password combinations for the sites and apps they use. It's simply too difficult to "follow the rules" and manage 20 different, cryptic passwords. As a result, the same password is used across many sites and hackers need only discover one of them to get into the rest. By the same token, using a Google or Facebook login implies that the hacker need only solve for that one. The point is, the extensive use of cloud services exposes users to password vulnerabilities.

Nexo: First, no part of Nexo is publicly accessible. All information exchanges are established on-demand and do not persist beyond their completion. Second, each user's "account" resides on their own device, so there are no shared access scenarios. Also, most OSs provide password protection for the entire device, and Nexo also requires an ID and password.

A real-world example of this was a security crisis at Yahoo, where user accounts were being "hijacked" by hackers who would properly log in, reset the password, and have exclusive access to someone else's email. It turns out that this began with a major ID/password theft from Amazon. Because millions of Amazon users use their Yahoo Mail account as their login ID, it wasn't difficult to gain access to the corresponding Yahoo accounts because they shared the same password. The resolution was to have all users change their password.

Nexo: We offer a mobile feature called the Personal Key. If membership in a Channel or access to a file is configured to require a Personal Key, then the recipient must have a verified mobile device that is within Bluetooth range of their computer. Because your mobile phone is the strongest proxy for your identity, this provides a powerful – yet passive – method for ensuring that the recipient is the person you intended.

Other Hacker Tactics

Phishing and Social Engineering

It's hard to believe that email's been around for half a century. It hasn't changed much since the 1990s and is a relatively open system in that a) the message body and attachments are sent in *plain text*, and b) the envelope data – "From" and "To" – are easily forged. These fields can be further concealed with the use of a "display name" so that the originating email address is not displayed by default. Even educated users are typically so busy handling dozens of emails at a time and filtering spam that they inadvertently fall for the deception.

Phishing is the *#1 method for infiltrating secure networks* and involves sending users messages made to appear as if they're from a normally trusted source. These can be elaborate and replicate graphical/HTML messages to the pixel. Links in the message can take users to infected sites where the mere visit transfers malware (called a "drive-by") or navigate to realistic-looking fake sites that ask to update personal information or reset a password.[6] If an attachment is opened, it immediately deposits malware on the user's computer.

Social engineering is the act of exploiting human tendencies to entice users into activating phishing messages or otherwise introducing malware to a device. Spearphishing is when specific users are targeted with content that is informed by combining known information with observed social network and other activity. For example, say an executive posts a photo of a corporate event on Facebook, which identifies other attendees. The spearphisher, posing as a fellow event-goer, sends a message to that exec and others with attached photos thereof. This is highly contextual and therefore less suspicious, and surprisingly effective.

Nexo: Unlike email, Nexo will use a proprietary ID scoring algorithm that uses multiple factors to verify that a user is who they say they are. For example, accounts with a verified credit card, mobile number and corporate email address receive the highest ID score, which is displayed to the recipient of a message from them. Another difference between Nexo and email is that intended recipients can screen inbound sharing requests before accepting them – emails just arrive in your inbox whether you want them or not. Nexo will also compare a sender's profile data (e.g., mobile number or email address) against your contacts and will flag exceptions.

Man in the Middle (MiM) Attacks

A MiM attack is when a hacker inserts himself between two connected and authenticated parties (this is a "Listen In" type of cyber crime). Typically, it's between a user and a computer, or between two users. The hacker has already figured out the protocol and the expected identifying data for each end of the communication. He then inserts a node in the middle of the session, as in the diagram above, causing A to believe it's still talking directly to B, and vice-versa. But in fact, the MiM is directly in the path of all the data flows, and can view or capture any part of them. The MiM can even pose as one of the parties.

[6] In October 2016, the DNC was unfortunately back in the news, this time with the release of John Podesta's email messages. He'd been phished by an email that looked convincingly like Google urging him to change his password due to its apparent theft. The URL in the email had been shortened using bit.ly and didn't arouse suspicion. After clicking it, he was sent to a fake Gmail page and was asked for his current password and a new password. The hackers simply passed the information to Gmail but now had his new password.

A variant of this attack is commonly done at unsecured, public Wi-Fi access points (e.g., a Starbuck's). The hacker can put a similarly-named access point in front of the legitimate one, and capture all of the network traffic associated with each user. We strongly advise not using such unprotected access points for this reason.

Nexo: We use our own proprietary, secure protocols when users establish a mesh network. When the sender requests the coordinates and keys for a given session, an encrypted list of valid nodes is created in real-time and distributed to all the participants. A MiM attempt will be immediately detected because it will not possess the right codes. The mesh algorithm will blacklist the fake node and ignore any packets coming from it.

Conclusion

While a good number of hacker tactics were described above, unfortunately, there are completely new types of attack discovered all the time. It's safe to say that any time you interact with a public cloud-based service, you and your information are exposed. The reason hacking is so rampant is because the Internet is based on standards-based protocols and systems that simply don't evolve as quickly as the attackers do. There's a logistical challenge as well – Internet-wide patches have to be applied to tens of thousands of nodes, and their propagation is slowed by the sheer magnitude of the effort, the need to apply patches to machines that are in full production mode without causing an outage, and certain administrators that might not wish to comply.

The only way to defeat threats that exploit innate vulnerabilities is to create a new protocol layer *on top of* what's there today. This makes it possible to employ inherently resilient architectures like peer-to-peer and to protect them with proprietary security schemes that go far beyond generic measures.