[IN]SECURE

Issue 48 - December 2015

INSIDE THE LARGELY UNEXPLORED WORLD OF
MAINFRAME SECURITY

DEVELOPING AND IMPLEMENTING AN INFORMATION SECURITY PROGRAM

APPLYING MACHINE LEARNING TECHNIQUES ON CONTEXTUAL DATA FOR THREAT DETECTION

IVAN RISTIC AND SSL LABS: HOW ONE MAN CHANGED THE WAY WE UNDERSTAND SSL

People spend over 700 billion minutes per month on Facebook.

Research by Facebook

Quality web filter

Comprehensive web security

Highly competitive pricing

Thousands of customers

The Internet is full of temptations. Can your users resist them?

The Internet is one of the most useful resources in the office – but only if you can manage the potential issues:

- Productivity losses due to employees spending time on sites with little work-related content
- Security risks: from unsecure sites and from legitimate sites that have been compromised
- Bandwidth losses from people downloading large files or watching streaming media.

Run the 30-day trial of GFI WebMonitor to find out exactly how your Internet connection and remote machines are being used and what security risks you are exposed to.

Download your free trial from http://www.gfi.com/webmon

Microsoft Partner

GFI WebMonitor™

Web security, monitoring and Internet access control

TABLE OF CONTENTS

- **Andrew Ginter**, VP of Industrial Security at Waterfall Security Solutions
- **Carl Herberger**, VP Security Solutions at Radware
- **Brian Honan**, CEO at BH Consulting
- **Matt Jones**, Partner at Elttam
- **Wolfgang Kandek**, CTO at Qualys
- **Ganesh Kirti**, CTO at Palerra
- **Zoran Lalic**, Senior Security Engineer at a large corporation
- **James J. Treinen**, VP, Security Research at ProtectWise
- **Geoff Webb**, VP, Solutions Strategy for NetIQ, the security practice of Micro Focus.

Visit the magazine website at www.insecuremag.com

(IN)SECURE Magazine contacts

Feedback and contributions: **Mirko Zorz**, Editor in Chief - mzorz@net-security.org
News: **Zeljka Zorz**, Managing Editor - zzorz@net-security.org
Marketing: **Berislav Kucan**, Director of Operations - bkucan@net-security.org

Distribution

(IN)SECURE Magazine can be freely distributed in the form of the original, non-modified PDF document. Distribution of modified versions of (IN)SECURE Magazine content is prohibited without the explicit permission from the editor.

Copyright (IN)SECURE Magazine 2015.

The privacy risks of school technology tools

The Electronic Frontier Foundation (EFF) filed a complaint with the Federal Trade Commission (FTC) against Google for collecting and data mining school children's personal information, including their Internet searches—a practice EFF uncovered while researching its "Spying on Students" campaign.

The campaign was created to raise awareness about the privacy risks of school-supplied electronic devices and software. EFF examined Google's Chromebook and Google Apps for Education (GAFE), a suite of educational cloud-based software programs used in many schools across the country by students as young as seven years old.

While Google does not use student data for targeted advertising within a subset of Google sites, EFF found that Google's "Sync" feature for the Chrome browser is enabled by default on Chromebooks sold to schools. This allows Google to track, store on its servers, and data mine for non-advertising purposes, records of every Internet site students visit, every search term they use, the results they click on, videos they look for and watch on YouTube, and their saved passwords. Google doesn't first obtain permission from students or their parents and since some schools require students to use Chromebooks, many parents are unable to prevent Google's data collection.

Google's practices fly in the face of commitments made when it signed the Student Privacy Pledge, a legally enforceable document whereby companies promise to refrain from collecting, using, or sharing students' personal information except when needed for legitimate educational purposes or if parents provide permission.

"Despite publicly promising not to, Google mines students' browsing data and other information, and uses it for the company's own purposes. Making such promises and failing to live up to them is a violation of FTC rules against unfair and deceptive business practices," said EFF Staff Attorney Nate Cardozo.

"Minors shouldn't be tracked or used as guinea pigs, with their data treated as a profit center. If Google wants to use students' data to 'improve Google products,' then it needs to get express consent from parents."

Revealed: What info the FBI can collect with a National Security Letter

After winning an eleven-year-long legal battle, Nicholas Merrill can finally tell the public how the FBI has secretly construed its authority to issue National Security Letters (NSLs) to permit collection of vast amounts of private information on US citizens without a search warrant or any showing of probable cause.

The PATRIOT Act vastly expanded the domestic reach of the NSL program, which allows the FBI to compel disclosure of information from online companies and forbid recipients from disclosing they have received an NSL. The FBI has refused to detail publicly the kinds of private data it believes it can obtain with an NSL.

Merrill has been privy to this information since 2004, when the FBI served him with an NSL demanding that he turn over records about a customer of the Internet company he then owned, Calyx Internet Access. Until November 30, 2015, Merrill was subject to a gag order forbidding him from sharing this information with the public.

Merrill is now able to reveal that the FBI believes it can force online companies to turn over the following information simply by sending an NSL demanding it: an individual's complete web browsing history; the IP addresses of everyone a person has corresponded with; and records of all online purchases.

The FBI also claims authority to obtain cell-site location information with an NSL, which effectively turns a cell phone into a location tracking device. In court filings, the FBI said that at some point it stopped gathering location data as a matter of policy, but that it could secretly choose to resume the practice under existing authority.

"For more than a decade, the FBI has been demanding extremely sensitive personal information about private citizens just by issuing letters to online companies like mine," said Merrill. "The FBI has interpreted its NSL authority to encompass the websites we read, the web searches we conduct, the people we contact, and the places we go. This kind of data reveals the most intimate details of our lives, including our political activities, religious affiliations, private relationships, and even our private thoughts and beliefs," he explained.

The law authorizing NSLs allows the FBI to demand "electronic communications transactional records" from online companies, but the FBI has long refused to clarify just how broadly it construes this vaguely worded and undefined phrase.

The NSL that Merrill received in 2004 included an attachment listing the specific categories of highly sensitive personal information that the FBI was demanding he disclose under this authority. Merrill has repeatedly challenged the gag order that forbade him from disclosing this information. The Media Freedom & Information Access Clinic at Yale Law School represented Merrill in his current, successful First Amendment challenge.

Three months ago, in a partially redacted opinion, Judge Victor Marrero of the federal district court in Manhattan found that the gag order was no longer justified. Judge Marrero's decision described the FBI's position as "extreme and overly broad," affirming that "Courts cannot, consistent with the First Amendment, simply accept the Government's assertions that disclosure would implicate and create a risk." He also found that the FBI's overbroad gag order on Merrill "implicates serious issues, both with respect to the First Amendment and accountability of the government to the people." Judge Marrero's ruling goes into effect today and has just been published in full, without redaction, after the government declined to appeal.

More than ten thousand NSLs are issued to online companies by FBI officers every year, and almost all of those NSLs are accompanied by a complete gag order barring any public disclosure of what the FBI has requested and from whom. Merrill is the first person who has succeeded in completely lifting an NSL gag.

"The broad scope of the FBI's claimed NSL authority is deeply problematic because the government can issue NSLs without any judicial oversight," stated Lulu Pantin, a law student intern who represented Merrill in his successful lawsuit.

VPN protocol flaw allows attackers to discover users' true IP address

The team running the Perfect Privacy VPN service has discovered a serious vulnerability that affects all VPN providers that offer port forwarding, and which can be exploited to reveal the real IP address of users.

Dubbed Port Fail, the flaw affects all VPN protocols (IPSec, OpenVPN, PPTP, etc.) and all operating systems.

"The attacker needs to meet the following requirements: 1. Has an active account at the same VPN provider as the victim, 2. Knows victim's VPN exit IP address (can be obtained by various means, e.g. IRC or torrent client or by making the victim visit a website under the attackers control), and 3. The attacker sets up port forwarding. It makes no difference whether the victim has port forwarding activated or not," they shared in a blog post, along with a step-by-step explanation of how the bug can be exploited.

The company has offered advice for VPN providers on what to do to plug this hole, but also did something that they should definitely be praised for: they tested nine prominent VPN providers that offer port forwarding for the flaw, and notified the five that were vulnerable of the fact before they went public with the information.

Thank-you messages on Twitter revealed that among the affected providers were Private Internet Access (PIA) and nVPN.

"However, other VPN providers may be vulnerable to this attack as we could not possibly test all existing VPN providers," the team pointed out. Hopefully, these providers are working mitigating the issue.

Security researcher Darren Martyn noted: "I believe this kind of attack is probably going to be used heavily by copyright-litigation firms trying to prosecute Torrent users in the future, so it is probably best to double check that the VPN provider you are using does not suffer this vulnerability. If they do, notify them, and make sure they fix it."

More than 900 embedded devices share hard-coded certs, SSH host keys

Embedded devices of some 50 manufacturers have been found sharing the same hard-coded X.509 certificates (for HTTPS) and SSH host keys, a fact that can be exploited by a remote, unauthenticated attacker to carry out impersonation, man-in-the-middle, or passive decryption attacks, Carnegie Mellon University's CERT/CC warns.

Stefan Viehböck, Senior Security Consultant at SEC Consult, has analyzed firmware images of more than 4000 embedded devices of over 70 vendors - firmware of routers, IP cameras, VoIP phones, modems, etc. - and found that, in some cases, there are nearly half a million devices on the web using the same certificate.

Another aspect to the whole story is the large number of devices directly accessible on the web," Viehböck also noted. Just by looking at the numbers one can deduce that it is highly unlikely that each device is intentionally exposed on the web (remote management via HTTPS/SSH from WAN IP). Enabling remote management exposes an additional attack surface and enables attackers to exploit vulnerabilities in the device firmware as well as weak credentials set by the user."

According to the researcher, affected vendors are: ADB, AMX, Actiontec, Adtran, Alcatel-Lucent, Alpha Networks, Aruba Networks, Aztech, Bewan, Busch-Jaeger, CTC Union, Cisco, Clear, Comtrend, D-Link, Deutsche Telekom, DrayTek, Edimax, General Electric (GE), Green Packet, Huawei, Infomark, Innatech, Linksys, Motorola, Moxa, NETGEAR, NetComm Wireless, ONT, Observa Telecom, Opengear, Pace, Philips, Pirelli , Robustel, Sagemcom, Seagate, Seowon Intech, Sierra Wireless, Smart RG, TP-LINK, TRENDnet, Technicolor, Tenda, Totolink, unify, UPVEL, Ubee Interactive, Ubiquiti Networks, Vodafone, Western Digital, ZTE, Zhone and ZyXEL.

The top 7 improvements in Nmap 7

Nmap 7 is the product of three and a half years of work, nearly 3200 code commits, and more than a dozen point releases since the big Nmap 6 release in May 2012.

The top 7 improvements in Nmap 7:

1. Major Nmap Scripting Engine (NSE) expansion

As the Nmap core has matured, more and more new functionality is developed as part of the NSE subsystem instead. In fact, 171 new scripts and 20 libraries have been added since Nmap 6. Examples include firewall-bypass, supermicro-ipmi-conf, oracle-brute-stealth, and ssl-heartbleed.

NSE is now powerful enough that scripts can take on core functions such as host discovery (dns-ip6-arpa-scan), version scanning (ike-version, snmp-info, etc.), and RPC grinding (rpc-grind). There's even a proposal to implement port scanning in NSE.

2. Mature IPv6 support

IPv6 scanning improvements were a big item in the Nmap 6 release, but Nmap 7 outdoes them all with full IPv6 support for CIDR-style address ranges, Idle Scan, parallel reverse-DNS, and more NSE script coverage.

3. Infrastructure upgrades

The Nmap Project continues to adopt the latest technologies to enhance the development process and serve a growing user base. For example, the developers converted all of Nmap.Org to SSL to reduce the risk of Trojan binaries and reduce snooping in general.

They've also been using the Git version control system as a larger part of their workflow and have an official Github mirror of the Nmap Subversion source repository. They also created an official bug tracker which is also hosted on Github.

4. Faster scans

Nmap has continually pushed the speed boundaries of synchronous network scanning for 18 years, and this release is no exception. New Nsock engines give a performance boost to Windows and BSD systems, target reordering prevents a nasty edge case on multi-homed systems, and NSE tweaks lead to much faster -sV scans.

5. SSL/TLS scanning solution of choice

Transport Layer Security (TLS) and its predecessor, SSL, are the security underpinning of the web, so when big vulnerabilities like Heartbleed, POODLE, and FREAK come calling, Nmap answers with vulnerability detection NSE scripts.

The ssl-enum-ciphers script has been entirely revamped to perform fast analysis of TLS deployment problems, and version scanning probes have been tweaked to quickly detect the newest TLS handshake versions.

6. Ncat enhanced

Ncat has been adopted by the Red Hat/Fedora family of distributions as the default package to provide the "netcat" and "nc" commands. This cooperation has resulted in a lot of squashed bugs and enhanced compatibility with Netcat's options.

Also very exciting is the addition of an embedded Lua interpreter for creating simple, cross-platform daemons and traffic filters.

7. Extreme portability

Nmap is proudly cross-platform and runs on all sorts of esoteric and archaic systems. But their binary distributions have to be kept up-to-date with the latest popular operating systems.

Nmap 7 runs cleanly on Windows 10 all the way back to Windows Vista. By popular request, the developers even built it to run on Windows XP, though they suggest those users upgrade their systems.

OS X is supported from 10.8 Mountain Lion through 10.11 El Capitan. Plus, support for Solaris and AIX was updated.

Analytics services are tracking users via Chrome extensions

It's quite possible that, despite your belief that the Google Chrome is the safest browser there is and your use of extensions that prevent tracking, your online movements are still being tracked. The culprits? Popular Chrome extensions like HooverZoom, Free Smileys & Emoticons, Flash Player+, SuperBlock Adblocker and many more.

The fact was brought to the wider public's attention by Detectify Labs researchers, who have signed up for one of the analytics services that provides user information gathered by Chrome extensions.

This information includes URLs that users visited (browsing history), cookies, OAuth access-tokens, and shared links from sites such as Dropbox and Google Drive (which, when shared by employees, often lead to confidential company data).

After signing up for the service, the researchers were able to see common URLs used by employees on targeted companies, internal network URLs and separated websites for internal use only, internal PDFs, and pages which only one person had visited.

"The tracked browsing history data is made available through analytics services, where anyone can sign up to pay for a monthly subscription to analyze and dig through this traffic," the researchers explained.

"It is still unknown what happens with some of the data, such as your personal cookies, but there's a possibility that it is being used to enhance the profile of the user to make the analytics even more accurate in terms of location, gender, age and interests. Through these services, we've been able to confirm that even browsing patterns from only one user ended up in the search results, making it possible to fingerprint a specific user's browser history."

If you're wondering how you didn't notice this data collection before, the explanation is simple: the offending extensions use different tactics to hide their tracking scripts' activities - from running in a separate background instance of the extension (so that network traffic is hidden from tracking prevention tools) and packing data to make it difficult to identify, to using different subdomains for each extension and enabling tracking by default.

What's more, some third-party tracking services use a tracking script SDK inside the extensions, which allows them to download new scripts.

"Our guess is that this is a way to bypass any filters used by Chrome Web Store to identify malicious extensions and abuse of privacy. It's also a great way for the tracking scripts to be auto updated, without forcing the user or the owner of the extension to update the extension," the researchers posited.

And, if you believe that the developers of these extensions tricked you into allowing this, you haven't read carefully the information on each extension provided in the Chrome Web Store. Because the explanation IS there, but is difficult to notice due to the Chrome Web Store's GUI, and due to the fact that descriptions of why tracking scripts are included and the scripts' privacy policies being (quoting the researchers) "a complete joke."

But why would extension developers include these scripts in their offerings, you ask? The answer is money.

"Many of these extensions are being paid per user by the third party to install the tracking code in their extensions. We've seen some indications on Chrome Extension-forums that it's around $0.04 per user/month. For plugins with over tens and hundreds of thousands of users that equals a substantial amount of monthly income," the researchers noted.

Now that you know this, you might want to check whether the extensions you use are doing this, and uninstall them if they do. The researchers also advised that, if you need some of these extensions, you should use Incognito mode for your regular browsing and make sure no extension is enabled in Incognito mode.

Finally, they also urged users to send business documents via email instead of through a shared link on a file sharing service like Google Drive or Dropbox.

Microsoft's new security posture leads to baked-in security

More than ever, Microsoft wants its products to be the first choice for enterprises, organizations, and governments. And to do that, they embedded security in the core. At the Microsoft Government Cloud Forum 2015, Microsoft CEO Satya Nadella pointed out how the company has been working hard to respond to the demands put forth by the rapidly changing threat environment.

The breaking of the perimeter, corporate networks being extended to customers, an increased pace of connectivity, the incorporation of employees' devices into the enterprise environment, the Internet of Things (sensors in every room) - it all leads to one realization: we live in a world where attacks can come from anywhere.

Information security, Nadella said, is one of the most pressing issues of our times, as digital technology today is at the core of every industry. But users won't use technology if they don't trust it, and Microsoft is doing everything it can to build that trust. He pointed out that Microsoft's unique perspective on what's happening both when it comes to attack and response, as they have insight offered by more that 1 billion Windows devices, 300 billion users authentications each month, and 200 billion emails analyzed for spam and malware.

This perspective allowed them to create a specific operational security posture, characterized by consistency (it's like going to the gym - you have to constantly "exercise" security, he says), constant improvement in threat detection (moving from signatures to spotting unusual behavior), complete protection (endpoints, sensors, data centers, etc.).

This security posture spurred them to come up with new solutions that incorporate three points: platform, intelligence, and partners.

Julia White, general manager of Microsoft Office division, then entered the stage to demonstrate some of the embedded security technologies incorporated in Windows 10, Azure, and Office 365 (the platform component of the aforementioned triad): from Windows 10 Hello and Password features that allow user authentication without passwords and Azure Active Directory (user and device management for Windows domain networks), to the Office 365 suite (with its malware detection improvements such as the link "detonation chamber") and Windows 10's Device Guard, a system which allows administrators to block the execution of software that is not digitally signed by a trusted vendor or Microsoft.

She mentioned Credential Guard, which uses virtualization-based security to protect against credential theft attacks, and new protections to prevent enterprise data loss, which she demonstrated by trying - and failing - to exfiltrate potentially sensitive enterprise data via email or by uploading a file to a cloud storage account.

These features protect data across all devices, and IT admins can define what can and cannot be performed on each of them. When it comes to breach detection, spotting abnormal behavior is the key. Microsoft's Advanced Threat Analytics (ATA) helps IT pros quickly detect suspicious user and device activity within the enterprise network.

On the infrastructure level, detection and prevention is aided by the Azure security center, which offers a central view of security settings, constant monitoring, policy based recommendations, and partner solutions.

Nadella then took over again and explained the "intelligence" aspect of the operational security posture. With a quick nod to Microsoft's past efforts and collaborations with the industry partners and law enforcement, he explained that security managers from Microsoft’s various divisions and product groups are finally being placed physically together in a Cybersecurity Defense Operations Center, so that they can create an accurate intelligence graph - a representation of the curent threat situation - which is then shared with customers and partners. Finally, he noted that they the partnership part of the security equation is equally important as the first two. After all, as Nadella pointed out, they want to take advantage of tools each partner brings to the table, and partners to do the same with Microsoft's tools.

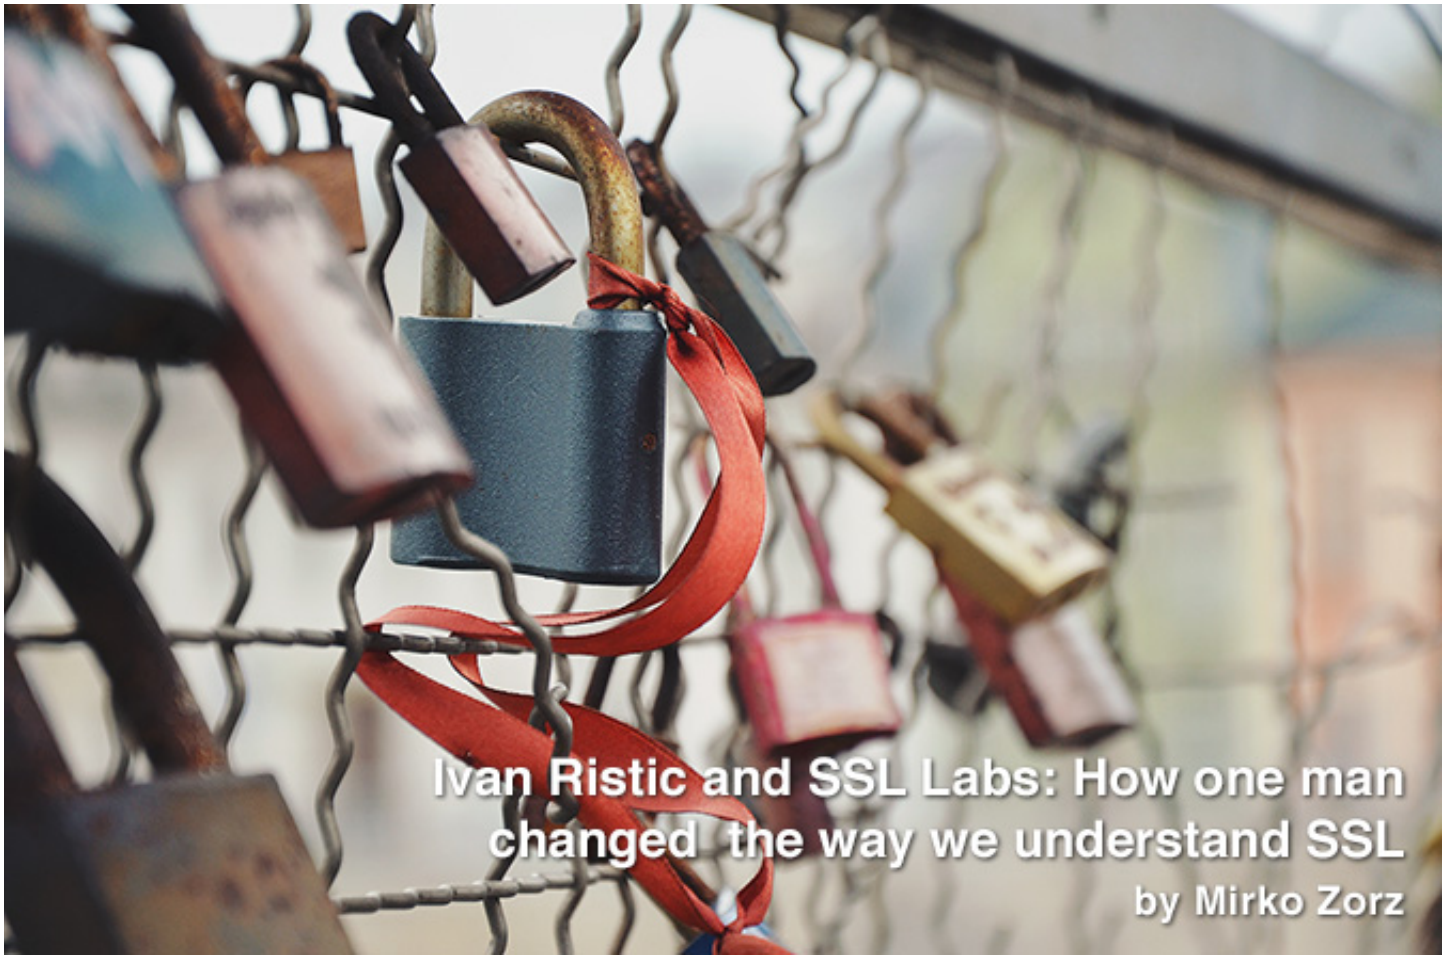

Ivan Ristic is well-known in the information security world, and his name has become almost a synonym for SSL Labs, a project he started in early 2009. Before that, he was mostly known for his work with OWASP and the development of the wildly popular open source web application firewall ModSecurity.

"When I originally came with the idea of SSL Labs, my primary audience were people like me, those who had to deploy encryption but were faced with poor documentation and behaviours. There were so many opportunities for mistakes and misconfiguration that the only way then (and today) was to inspect a running service to be absolutely sure," he explained to us his motivation for starting the project.

"I was well aware of the complexities of SSL deployments, because I had been using it for years. I was frustrated with the lack of tools and good documentation and I was sure that others were too. So I decided to create a tool to help myself as well as others."

SSL Labs was a pleasure project for Ristic, something he worked on in his spare time, so it evolved slowly at the beginning. But after he joined Qualys in May 2010 and became the company's director of engineering, he showed the project to Qualys CEO Philippe Courtot, who fell in love with it.

It took a couple more years for it to move from the status of "side project" to that of one of the main ones, but since 2013, it became Ristic's main focus at the company, and he gives Qualys much of the credit for the project's success.

"It's doubtful that I would have been able to spend adequate time on it were it not for the Qualys funding, and it was that which allowed me to respond to the challenges," he noted.

"Over the years, SSL Labs incorporated a great number of checks that are impossible to perform manually. With SSL Labs, you can do them in a minute. It's a game changer because, to assess your TLS configuration, you don't need to be an expert (which is extremely

difficult because of how quickly things change). In other words, you can focus on your job instead," he explained.

As time passed, there were other improvements. For example, organizations can perform automated assessments via the projects APIs - they can feed all their hostnames to the tool, automate the scanning, and know exactly when something changes (either because they broke something or because a new issue had been discovered).

"The usefulness of SSL Labs increased significantly when we started simulated capabilities of widely used clients (over 40 of them at this time), which helps with availability. Now you no longer have to be afraid if a change you're making is going to break something. Instead, you can see exactly how a particular client would behave," he added.

For years, and even after joining Qualys, SSL Labs' setup was one server hosted in the cloud and Ristic as the manager. But when Heartbleed hit in April 2014, they were inundated with a million sessions in only a couple of days, and they had to scramble to pad the backend.

"Luckily, it was easy to clone that server into six to handle the load," says Ristic. "The bigger problem was the fact that I was on vacation that week and with an unreliable Internet connection."

After that incident, SSL Labs was moved into the Qualys's data centre, where it remains today. Ristic remains the only developer, but the production servers are now maintained by the company's Ops team.

SSL Labs is not only helpful to organizations, but to end users as well. An increasing number of them started to care about security, and the project allowed them to gain some visibility into the security posture of a particular web site and, consequently, this gives them an idea of whether or not a particular organization is serious about security.

Ivan Ristic in his London office.

"Finally, SSL Labs also works as a great tool for raising awareness about various issues. It's now helping us transition from using weak ciphers and protocols to stronger configurations," he pointed out.

Future plans

The future of the project looks bright. Ristic plans to revamp the grading criteria to make it easier to understand, to remove some baggage (the current version is from 2009, when SSL/TLS security was vastly different), and to add support for the assessment of protocols other than HTTP. Many other improvements are planned, but we'll have to wait to hear about them until they are closer to becoming reality.

In the long term, the plan is to make SSL Labs better, either by adding new features or by making it more user-friendly.

"It's difficult to have a good plan when you are forced to react to external events," he says. "For example, progress on new features has been slow in the last two years because I had to instead spend my development time to handle various vulnerabilities: Heartbleed, POODLE, POODLE TLS, Freak, Logjam, and others. For a while it felt like I had to run just to stay in the same place."

Lessons learned

The project taught Ristic a great many things.

"As a user of TLS, you don't realise how many moving parts there are behind the scenes," he noted. "If I had to pick one thing, I'd say that I learned a lot about cryptography engineering. This comes from learning why certain features work in a certain way and, especially, why certain designs cause security issues. Apart from that, it was quite interesting to understand how much diversity there is in TLS deployments; so many different products with different capabilities and quirks. Although that doesn't seem to be very useful at first, it actually teaches you a lot about how to design a protocol that is used by billions of devices over several decades."

SSL Labs never stopped being a pleasure project for him. Part of the pleasure is that it is making a difference in a small way. Initially, he didn't think about where the effort would ultimately lead and he didn't think that SSL Labs would become so important.

"That's the beauty of building something because _you_ need it, you don't have to care about popularity," he says. "Of course, I'd lie if I said I didn't care. When you're sharing something with others and you're not asking them to give you money for it, the popularity of the product becomes the currency; what you're paid with. The more popular the product gets, the more motivated you feel to work on it. So it's like fuel for your development, in a sense."

"My philosophy has always been to pick one thing, then persistently work on it until you understand it fully and generally do the best job you can. SSL Labs was just the right size for this approach, a good project for one person to handle," he concluded.

Mirko Zorz is the Editor in Chief of (IN)SECURE Magazine and Help Net Security (www.net-security.org).

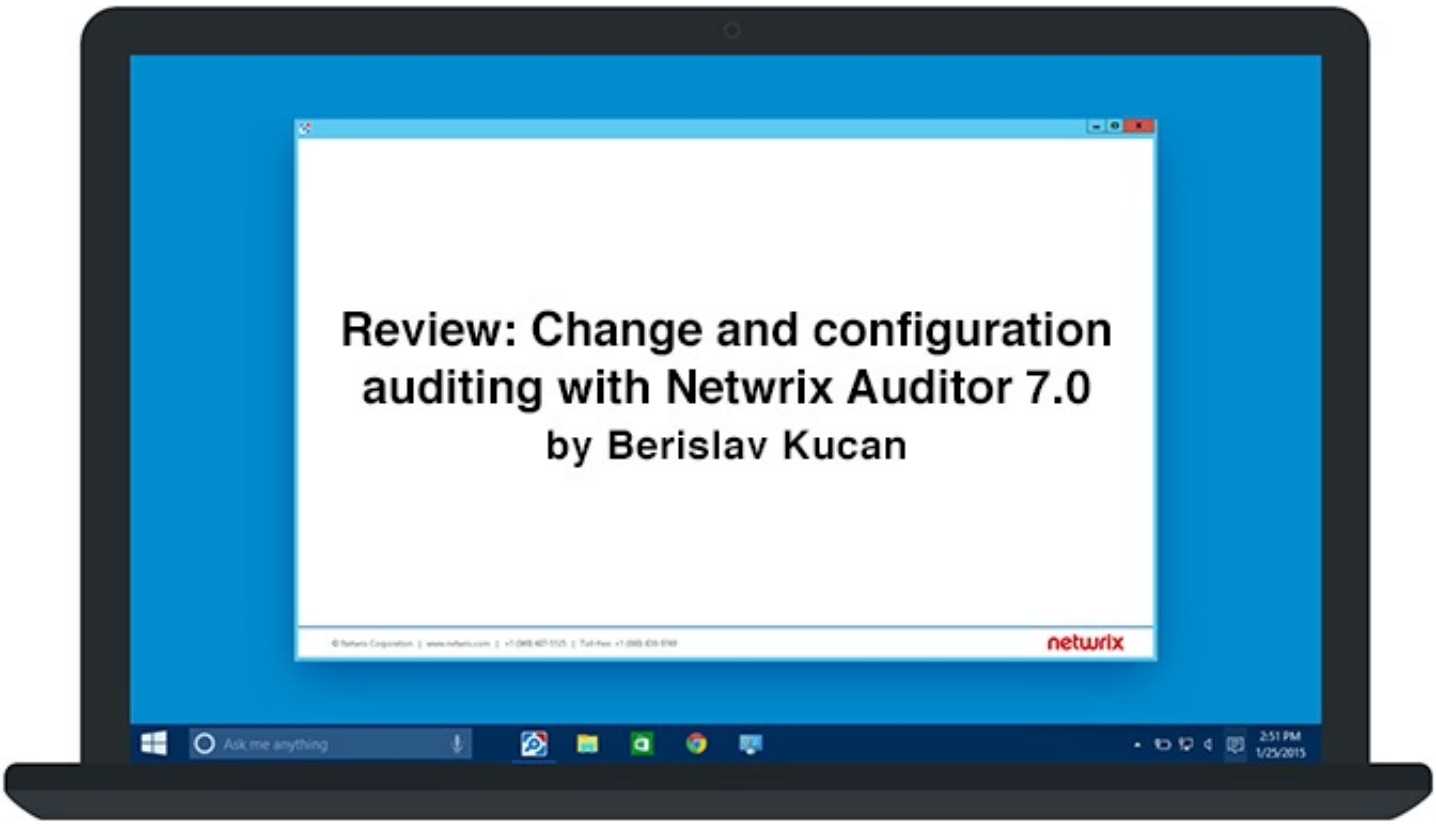

Netwrix Auditor is a powerful change and configuration auditing platform that leverages the data collected from all parts of the company network to provide detailed information on everything that is going on inside.

Installation

Installing Netwrix Auditor is a pretty straightforward procedure. You should install it on a workstation (more info on this below) and you will need administrator privileges to do it.

The solution stores the collected data in a two–tiered audit archive that includes a file–based local long–term archive and a short-term SQL-based audit database. For the latter, you'll obviously need a database server, so you can either type in the credentials of an already running server, or Netwrix Auditor will install and set up Microsoft SQL Server Express on your behalf. The whole installation, without the potential SQL server install, will be over in a minute or two.

The documentation suggests that Netwrix Auditor should be installed on a workstation rather than on a domain controller. This is probably because the software requires a SQL server and the Express version cannot be deployed from inside the Netwrix Auditor installation. Also, according to Microsoft, it is not recommended to install SQL Server on a domain controller because of specific security restrictions when running it in this configuration. Another thing is that, given the resource demands of a domain controller, SQL Server performance may also be degraded.

A successful installation on your system will generate two executables for you to use - Netwrix Auditor Client and Netwrix Auditor Administration Console.

For most of the systems you will audit, Netwrix provides both agent-based and agentless data collection methods. Installing agents is recommended when auditing SharePoint farms (*SharePoint_Shell_Access* role needs to be assigned and the agent needs to be manually installed) and when tracking user activity (done automatically without any intervention from the Administrator Console).

The Netwrix Auditor client can also be deployed on multiple computers through Group Policy.

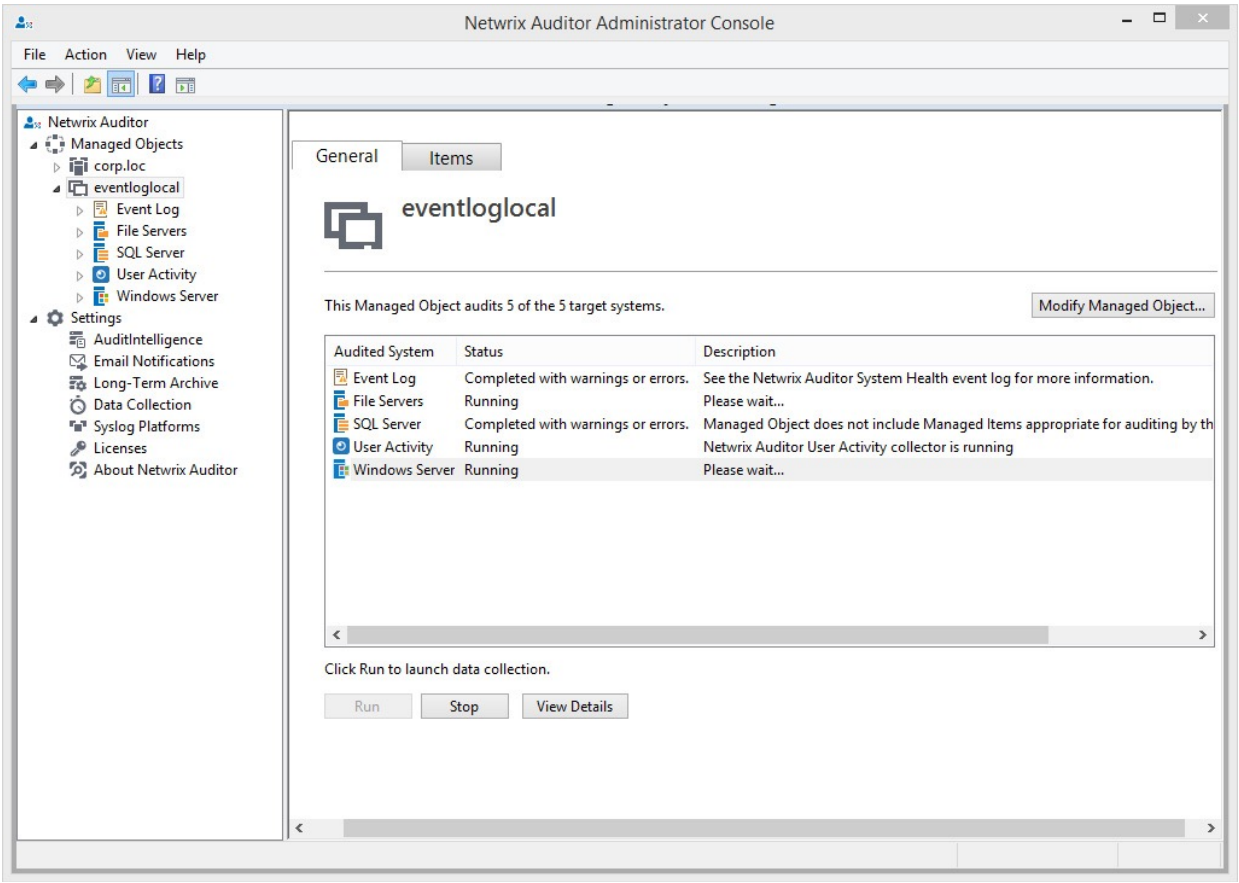

(Pre)configuration of audited systems

Before getting into the details on how to initiate audit procedures, it is very important to preconfigure all the systems that will be monitored by the Netwrix solution. The documentation that comes with the product specifies all the aspects you need to think about to prepare your environment for Netwrix Auditor workflows. Seasoned administrators will be on top of these things, but reading through the "Configure IT Infrastructure for Audit" part of the installation manual is nevertheless recommended.

I had some issues with (relatively) newer workstations and servers where Microsoft .Net Framework 4+ was installed. Getting some of the data from these machines didn't work and the root cause was detected when I read through the Netwrix Auditor System Health log inside the local Event viewer. The error shown was that Microsoft .NET framework 2.0 is required. As it turns out, Microsoft .NET 4.5 is not totally backward compatible with 2.0, meaning that some libraries are missing in action. The solution to this, which was the only minor bump in the road I had with Netwrix Auditor, was to enable Microsoft .NET 3.5 on the Microsoft Windows Server 2012 systems in question, as it contains all the needed dependencies.

Note: Netwrix Auditor System Health is a good tool for checking whether there are any errors in the connection between the Netwrix Auditor workstation and the audited environment.

Administrator Console

This is the part where the actual configuration starts. Netwrix Auditor uses a set of configurable Managed Objects to specify where the data will be collected. By default, the administrator can choose one of the following applications to create a Managed Object:

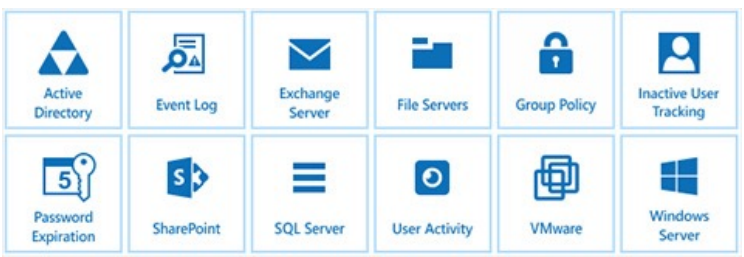

The actual objects that can be used by these specific built-in applications can be: *Domain, Computer Collection, Organizational Unit, SharePoint Farm* and *VMware Virtual Center*. Of course, specific audited system apps use just the appropriate object. It is usually just one, for instance for Active Directory you have the *Domain* as the managed object and for *Inactive Users Tracking* you can choose either a *Domain* or *Organizational Unit*. Every application presented here has its own context and therefore different levels of configuration options.

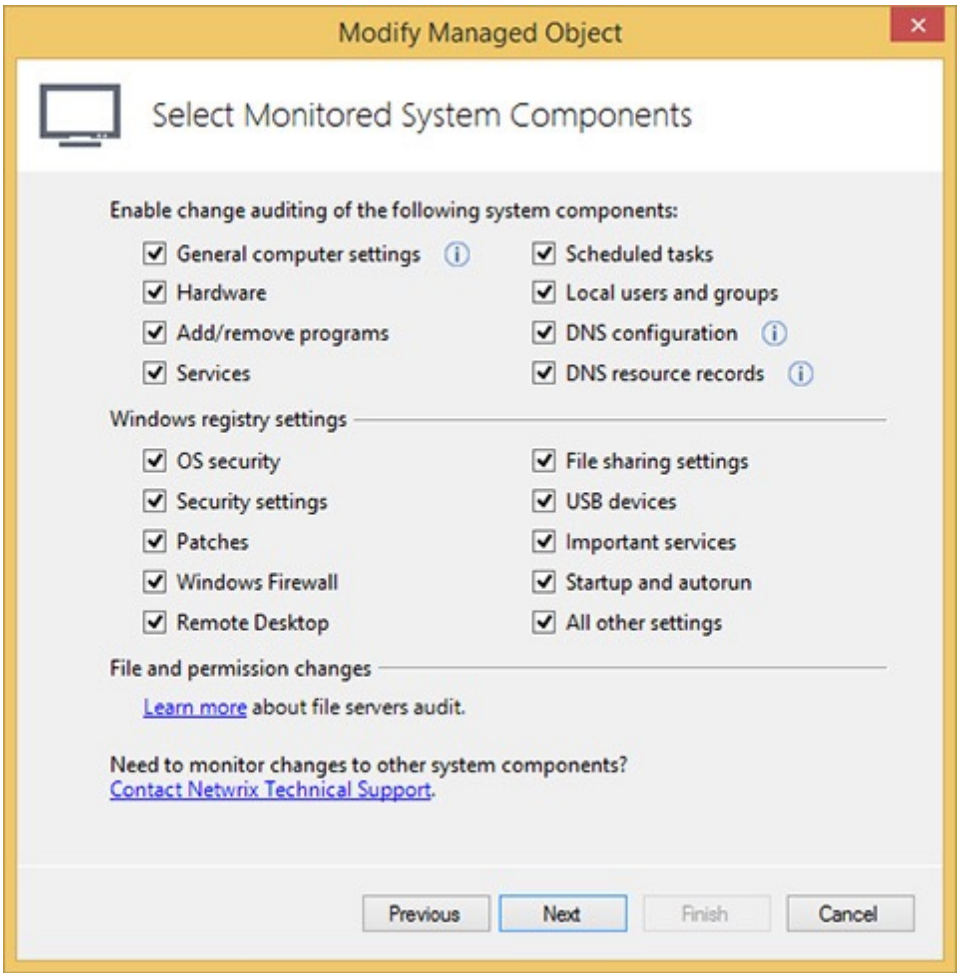

The interface of the Netwrix Auditor Administrator Console is your typical two-panel window where the list of objects and settings is on the left and the data or further details open on the right. The layout is very clean, with the focus on efficiency. All of the newly created managed objects are available under the Managed Objects listing and you can create folders and group them using the structure that suits you the best.

Besides building the audit objectives, the admin console provides a number of settings related to the general usage of the software. Here you can set up email notifications, location and retention settings, schedule data collection cycles and manage existing ones. There is also support for a couple of Syslog based platforms. When it comes to scheduling, the default option is to set up a task once per day, but you can also set up multiple ones.

Searching through the data

The more systems and targets you define in the Administrator Console, the more data you will have in your retainers. Collecting massive amounts of data from your networks can prove to be very valuable. Netwrix Auditor's search capabilities are immense, the system is very fast and provides a long list of search parameters. Besides the most basic ones like searching for a word or a set of characters, you can always choose to use one of the predefined search modules, which include:

1. **Who** - targeting a specific account
2. **Action** - choose one of the actions including what was added, removed, modified or read
3. **What** - searching for a specific object
4. **When** - aiming for a date range
5. **Where** - match just selected host, domain, etc.

While it seems you cannot use regular expressions in these predefined modules, there is an advanced section where the administrator can set up specific filters with six different operators (“contain”, “equal”, “starts with” + their opposites). These filters are a must when going deeper into mangling the data.

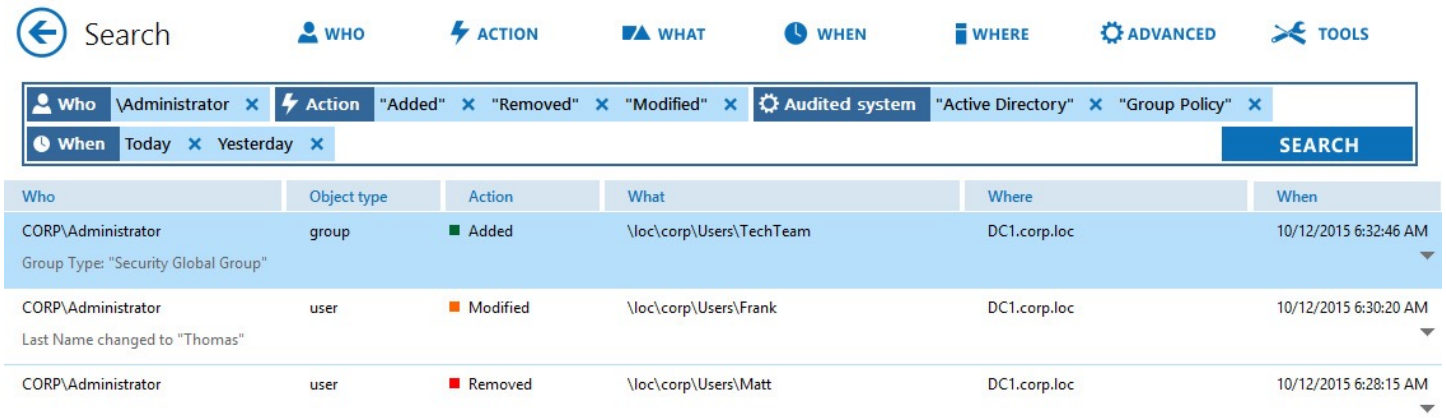

The results of a successful search query are laid out in rows and columns and are very easy to read through. Every result can be expanded to present a more detailed look. What I really like is that the values of specific objects from the search results, such as *Action, When* or let's say *Where*, can be reused to create new search queries. This comes quite handy when you come across something that sounds suspicious or interesting enough to make you dig deeper.

Custom search queries can be saved and if you do that, they will appear on the main screen of the application. If you save a large number of searches, the screen starts to look a bit cluttered, so it would be nice if one of the next versions enables some way of organizing the saved searches.

If you need detailed monitoring of user activity, you will be happy to know that aside from the textual logging of every move your users make, Netwrix Auditor also deploys a video recorder as well. Files can be accessed directly from the Reports section, or you can find them in the right folders inside the Program Data directory. By the default settings, the video quality is really good. FYI, a 23 minute long file of a Windows 8 user’s actions on his workstation was compressed into a 1.35 MB avi video file. The resolution in this case was 1024x768 and it was recorded in black and white. I am not into video editing, but taking into consideration the length and resolution of the video, the size of the file seems reasonably small.

Reporting and compliance

Big data and a powerful search engine should always be accompanied by detailed reporting capabilities. While browsing through the templates, I stopped counting, but I presume there were a couple of hundred of different reports you can run on your data. Templates are categorized into topics, where you can expand the icons to access more subsets (example: *Active Directory > Active Directory State-in-time > User accounts group memberships*).

In case your organization is audited and you need to prove that specific processes and controls are/were in place, just select one of the compliance reporting datasets which include PCI DSS 3.0, HIPAA, SOX, FISMA/NIST800-53 and ISO/IEC 27001.

Access to the audit data can be given to specific teams as the Netwrix Auditor client can be installed on multiple computers. It just needs to be configured to connect to the main workstation with the appropriate credentials.

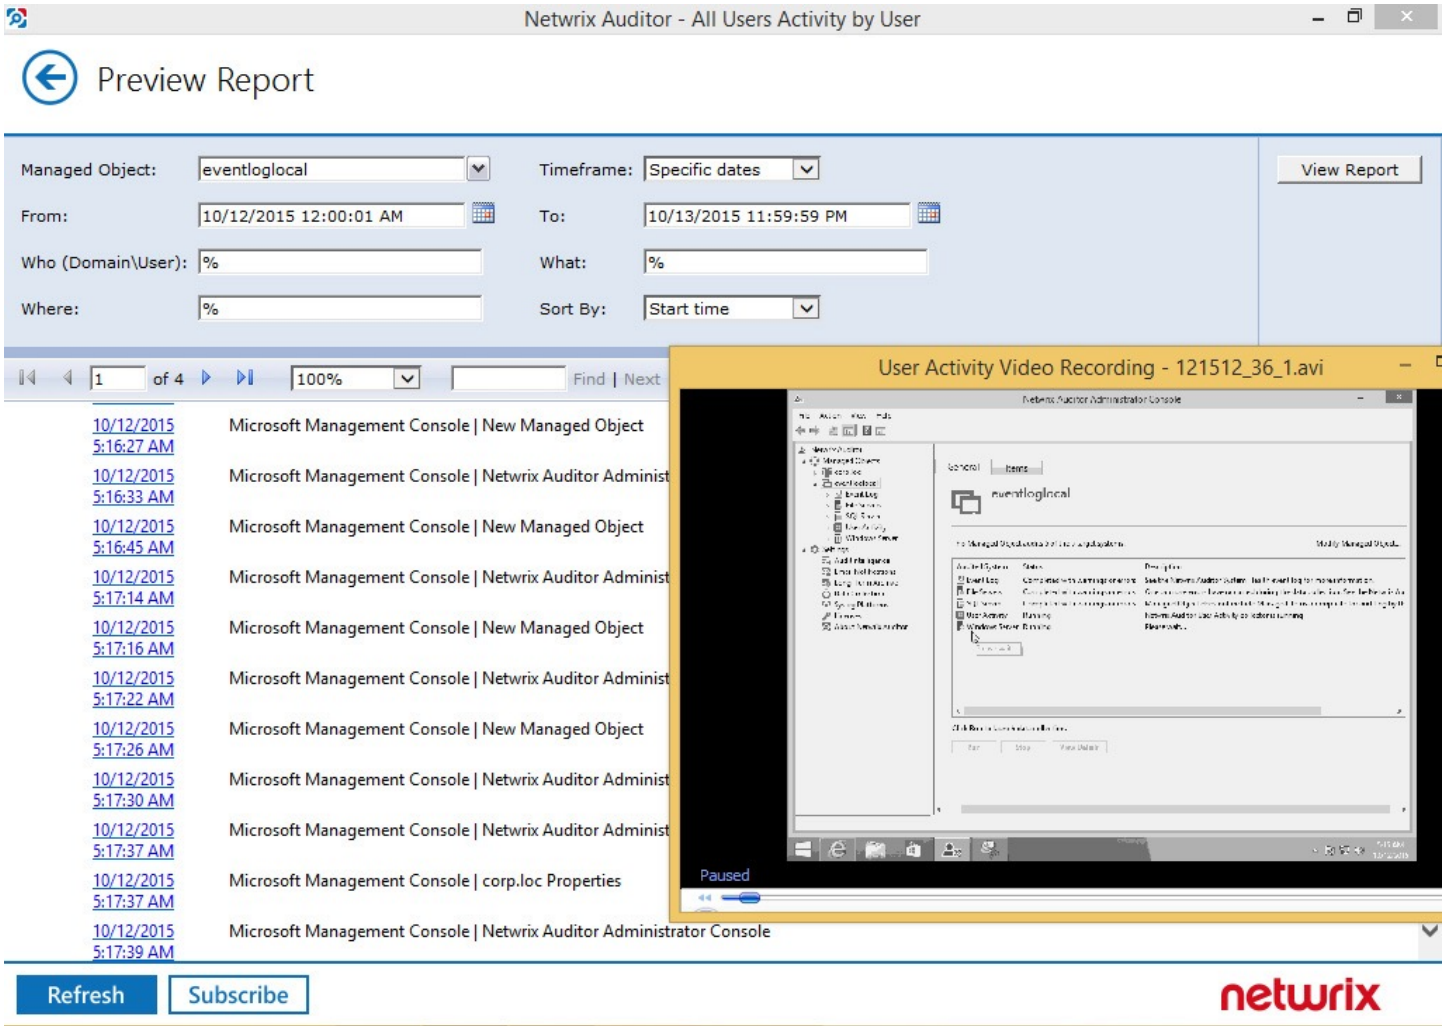

Besides providing the users with different methods of exporting the reports (PDF, Excel, Word and Atom data feed), there is a possibility of creating subscriptions for specific reports. Every generated report features a "Subscribe" button, where the information can be customized and associated with the selected recipients.

Delivery over email can be in PDF, Excel, Word and CSV data files and the minimum timeframe between sent reports is 24 hours.

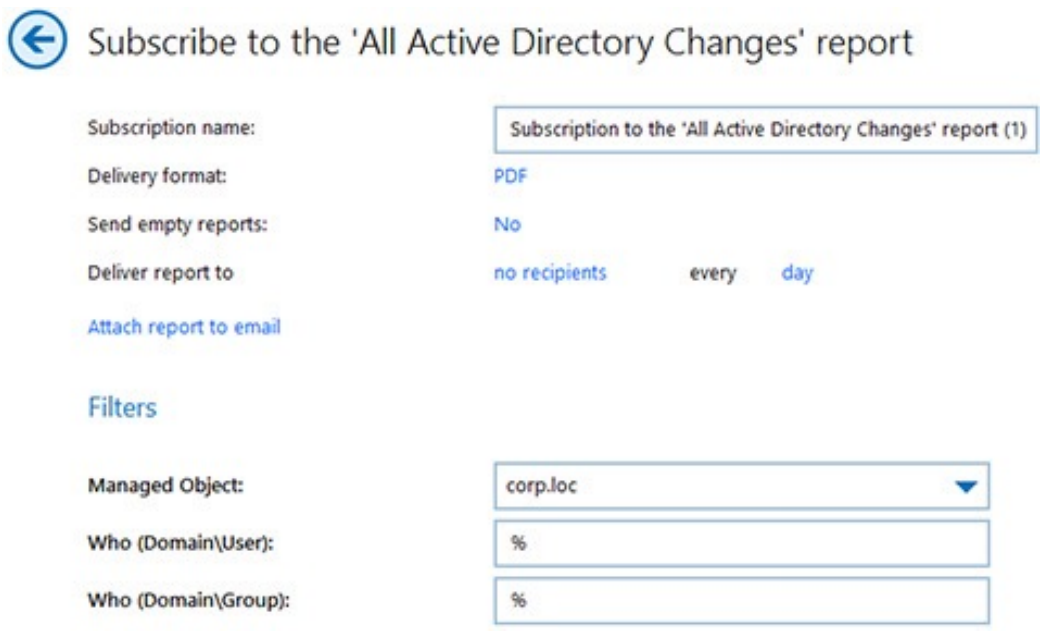

Documentation

As you can see so far, Netwrix Auditor is a robust solution that has a very broad coverage of audited systems. To function properly, each audited target needs to be configured to support data collection. The documentation provided by Netwrix is impressive and spreads over 430 pages in four different PDF documents. These include the installation and configuration guide, administrator guide, user guide and release notes.

The advice presented in the installation and configuration guide will be pivotal for successfully setting up every possible detail related to particular conditions of tracking changes inside your corporate network. As with the other documents, it is formatted very well and provides both a thorough step-by-step manual, as well as a quick reference guide with specific details such as how to configure object-level auditing for the domain partition.

The user guide is intended for Netwrix Auditor users (not admins) that are tasked for searching and filtering of audit data, generating various reports, etc. While the release notes seem to focus just on version history, it is actually much more than that. In a transparent and very helpful manner, Netwrix provides a 20 pages long list of known issues that you could potentially come across while running Netwrix Auditor. Every known issue is assigned its own ID and contains a detailed description. The majority of issues have workaround ideas or solutions, while others point the user to specify the ID in question to the Netwrix Technical support.

Pricing

Each Netwrix Auditor application (audit system type) is sold separately. Pricing starts from $3 and is calculated per enabled Active Directory user. First year of support and maintenance is included.

Final thoughts

By combining in-depth collection methods, powerful search engine and extensive reporting functionalities, Netwrix Auditor proves to be an impressive solution for maximizing visibility of every aspect of what's going on inside your IT infrastructure.

Berislav Kucan is the Director of Operations for (IN)SECURE Magazine and Help Net Security.

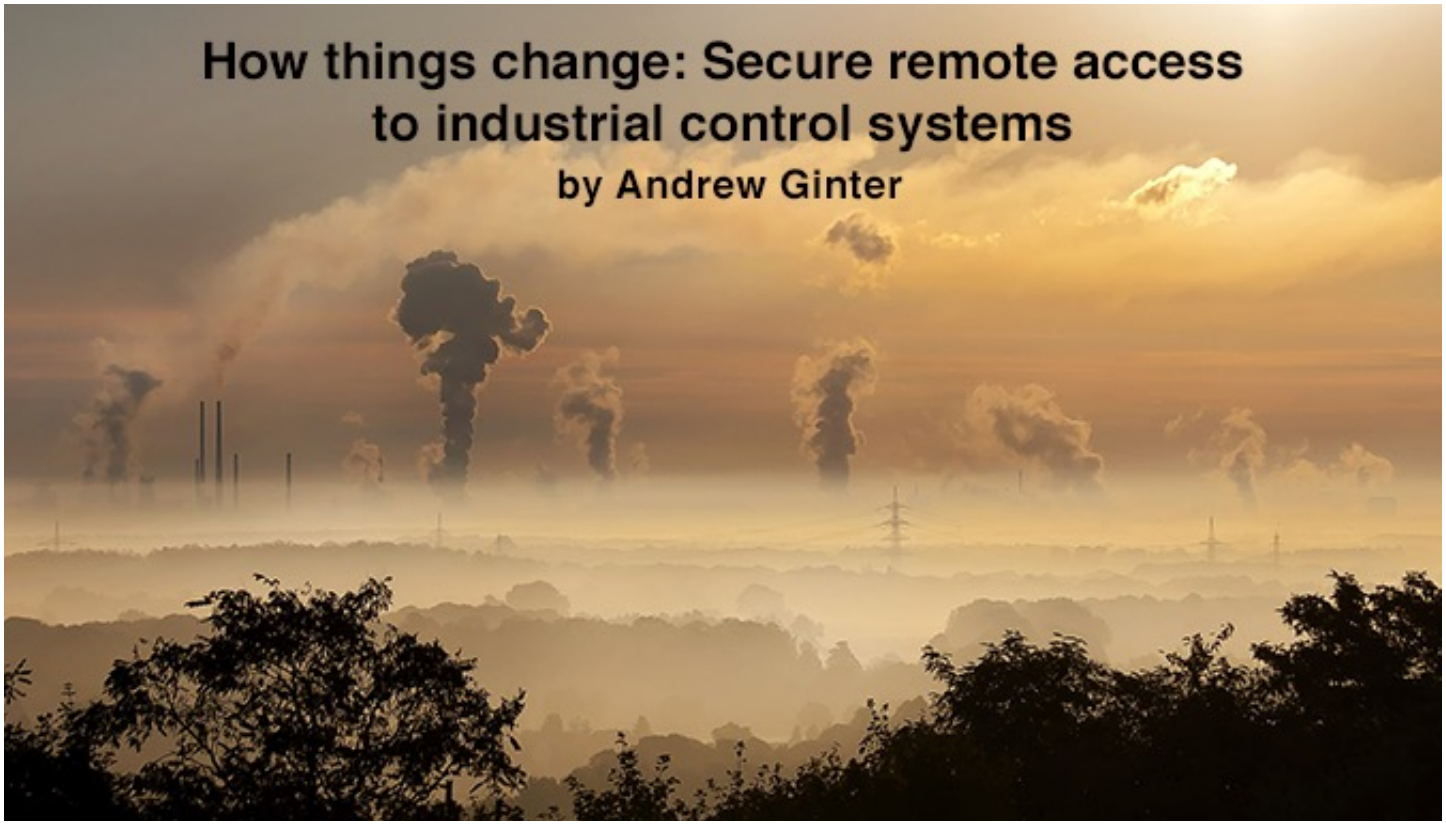

How things change: Secure remote access to industrial control systems

by Andrew Ginter

2016 is shaping up to be a year of curtailment for remote access to industrial control systems. Remote access is taken for granted in many organizations, enabling telecommuting and "anywhere, anytime availability" for a mobile workforce. Remote access carries risks of course, but these risks are seen as manageable and acceptable to most businesses. Most businesses, though, do not use industrial control system software to operate billion-dollar-plus physical infrastructures. Most businesses never experience the realization that they can't restore a physically damaged turbine or cracking tower from backups.

IT-style protections are failing control systems

IT firewalls are porous by design, and these porous firewalls are, in fact, what we need on corporate networks. Corporate firewalls must permit thousands, or hundreds of thousands, of email messages and Web pages into IT networks every day. Firewall vendors do what they can to filter out attacks, but no filter is perfect. All firewalls permit attacks into IT networks.

Given this reality, it seems only natural to accept the risks of remote access – what's another few attacks making their way into an IT network that deals with attacks more or less constantly? We watch our IT networks. We identify compromised computers. We isolate them, and restore them from backup.

On control-system networks, though, we need different protections. Control-system perimeters do not need to be porous; while there are many reasons to monitor a control-system network, there is rarely any reason to control a power plant or refinery from the Internet. Recognizing this, the Federal Energy Regulatory Commission (FERC) asked for public comments recently about the risks of remote access to industrial control systems. Taking over a control-system remote-access session is straightforward after all: seed a little malware on a technician's laptop with spear-phishing, evade anti-virus by infecting only a small number of laptops, gain administrative privileges through social engineering, and wait.

Simply wait for the user to log into the control system, over a thoroughly encrypted VPN connection, with a two-factor dongle, and a fingerprint swipe to boot.

Once the remote-access session is well and truly established, make the remote-access window disappear and give control of the window to a remote attacker to reprogram the control system.

FERC has asked if risks like this one are acceptable. This is a continuation of a discussion FERC initiated some time ago as to whether "cloud" control-system services pose acceptable risks. These services constitute more or less continuous remote access from some site on the Internet, straight into the brains of hundreds or even thousands of industrial control systems. A compromised cloud system is able to carry out widespread, simultaneous attacks on entire industries and nations.

The risk for control systems is, of course, cyber sabotage, rather than cyber espionage. While wholesale theft of data is detectable using intrusion detection systems, control-system sabotage can happen without much data at all flowing into the Internet to trigger intrusion detection systems. Worse, the average intrusion takes weeks or months to detect. For all that time, a remote intruder has remote control of equipment on our reliability-critical or even safety-critical networks. Should this really be a risk we accept?

Unidirectional security gateways needed for safe remote access

Industrial control system owners and operators are increasingly replacing firewalls on their control system networks with unidirectional security gateways. The gateways replicate database servers and other servers from industrial networks to corporate networks, where corporate applications and users can freely query those replicas. Since the gateways are physically unable to send any message back into control-system networks, a unidirectionally protected industrial network enjoys true freedom from Internet-based remote-control attacks.

Occasionally, legitimate remote access may still be needed though, even into unidirectionally protected sites. For example, while control-system software, equipment and even cloud vendors are able to use replicated servers for monitoring, those vendors' experts may still need to occasionally reach back into control-system networks to adjust systems or equipment. On the surface, hardware-enforced unidirectional server replication would seem to prohibit any such remote access. In practice though, a number of remote-access mechanisms are possible, even with unidirectional gateways, mechanisms that are often counter-intuitive to firewall practitioners.

For example, the most secure remote-access mechanism is known as "remote screen view" (RSV). RSV works vaguely like remote desktop, in the sense that RSV agents capture screen images and send those images through unidirectional gateways to external viewers. But no mouse-click or keystroke can penetrate the unidirectional security gateway hardware. Vendors wanting to use RSV to adjust control-system components must make an appointment with someone at the control-system site. An authorized individual at a site can take advice over the phone from any vendors "looking over their shoulder" with RSV. The individual at the site is the one moving the mouse, and is the person ultimately responsible for any changes made to equipment at the site. RSV provides a way for vendors to provide specialized expertise to an industrial site, without putting that site at risk of compromise.

French regulations for critical industrial infrastructures already demand that unidirectional gateways be the only connection between critical networks and less-trusted networks. FERC's expressed interest in the risks of remote access and especially the risks of cloud/vendor connections suggests that remote access to industrial control system networks is a topic that will see a lot of discussion in 2016. Whether FERC mandates additional protections for industrial networks or not, owners and operators are well advised to consider their own risk profiles in light of these widespread discussions, and consider protecting industrial networks unidirectionally, rather than relying on porous firewalls and "restoring from backups" their damaged equipment.

Andrew Ginter is the VP of Industrial Security at Waterfall Security Solutions (www.waterfall-security.com).

Developing and implementing an information security program

by Zoran Lalic

In today's digital age, every organization, regardless of its size, must have an information security program in place to adequately protect data – both their own and that of their customers. Whether you create your own program or adopt a framework that is several hundred pages long, it must be designed in a way to detect, prevent and significantly reduce the risks. Developing a comprehensive information security program that recognizes these risks is one of the major issues that organizations face today. The program should reflect a security strategy that goes beyond the traditional security tools, such as firewalls and anti-virus tools.

The most recent data breaches show that processes within the security programs have failed and programs were not comprehensive enough to detect attacks in time. Typically, a series of different security incidents and events happen over a certain time - days, months, or years - that organizations choose to ignore until it is too late.

It's no secret that in most data breaches, cybercriminals understood the target's defensive capabilities better than the targeted organization itself. Organizations must have a home court advantage; the best approach to accomplish it is to develop and follow an information security program. So let's piece together the puzzle. Here are seven key elements for a successful security program.

Executive management involvement

A successful security program has several pre-requisites, the most critical of which is executive management commitment and support. Without executive management leadership, buy-in and active approval and support, any effort to develop and implement a successful security program will most likely fail.

Let's compare two organizations. Organization X has an information security program implemented with executive management support. Organization Y also has an information security program implemented, but without active executive management support. Both organizations were alerted about a zero-day vulnerability affecting their Internet-facing web

server. After conducting their risk analyses, the security teams from both organizations made the decision to disable the vulnerable service on their web servers.

Communications have been sent out informing each organization about their security teams' decisions. At Organization X, the Director of Marketing goes to the CEO explaining that the disabled service is really necessary for the marketing team to perform their daily tasks. The CEO, who participates in a security risk committee, explains that, if exploited, the newly discovered vulnerability could potentially expose customer data and that it was necessary to disable it until patched. At organization Y, the Director of Marketing goes to her CEO with the same request. The CEO, who does not participate in a security risk committee, sends the request to the security team to enable the service to avoid an impact on company revenue.

Executive management must understand that information security is not only extremely critical to an organization's success, but that it is also their responsibility. Executive management needs to participate in information security committee meetings.

Governance

A successful security program must be governed. Information security governance is an essential component of an information security program because not only does it ensure strategic alignment between the business and security initiative, it also provides direction for an effective security program.

Let's apply the information security governance strategy of Organization Y. Organization Y has a security initiative to select and implement a DLP solution company-wide. The tool has been selected based on the defined requirements; the solution has been implemented. Soon after the implementation, Organization Y made the decision to move into the public cloud. The security team has realized that the DLP tool that was recently purchased is not supported in the public cloud. In other words, the security team has to understand the business and its mission so that the security program can be aligned with it. Information security is usually viewed as a business roadblock, but in reality, it is a key business enabler.

A critical piece of information security governance is the information security steering committee. This is a team that consists of security personnel and business unit leaders whose goal is to integrate information security into the business flow.

Some of the main responsibilities of this team are to:

- Define roles, responsibilities, accountability and communication of the security program and security initiatives
- Review and approve information security policies and processes
- Promote information security education and training
- Ensure security policies are followed
- Promote information security within their business units.

Executive management support and well-defined security governance are the foundation of a successful information security program. It is now time for the organization to start building and executing the program itself.

Risk management

Variations of each security program will depend on the organization and the industry in which the organization operates. Each organization has its own risks. Some organizations must adhere to compliance regulations; however, the main components of every security program are usually similar.

The organization needs to understand what risks exist in the environment before formulating proper security controls to address those risks. The organization needs to conduct a risk assessment to identify both external and internal threats and vulnerabilities that may be a danger to the organization. One element often overlooked is the asset and software inventory. Without a proper inventory, it would be nearly impossible to know what to protect within the organization.

Consider an identified risk where no risk analysis has been conducted for zero-day vulnerabilities that might be introduced.

A remediation for this risk might be creation of a vulnerability flow chart and vulnerability analysis team. This team would meet and analyze zero-day vulnerabilities and make a decision on a remediation process based on the analysis.

There is no silver bullet to make an organization risk-free, there will always be some risk. Risk assessments should prioritize risk and the security program should ensure that risks are properly managed and minimized, but risk will never be completely eliminated. Consider the organization's risk appetite in order to ensure the acceptable risk. In case of an exception, ensure that a proper exception process is in place and is approved by executive management.

Standards

Armed with knowledge from the risk assessment, the organization is now ready to address risks. The first step for most organizations is to select a security standard or framework such as NIST or ISO 27001/02. This standard will act as a baseline to implement proper security controls that address the risks identified during the risk assessment. When selecting the standard, the organization should take into consideration any regulatory requirements they might have to comply with, such as PCI DSS and HIPAA.

Upon the successful selection of the standard, the organization should conduct a gap analysis to identify which controls are required to address the risk factors. Consider risk assessment findings as a current state and the required security controls as a desired state. The transport from the current state to the desired state is known as a security roadmap.

The security roadmap should consist of projects and initiatives that will take the organization from the current state to the desired state. This road trip is known as an Information Security Program.

FOR A POLICY TO BE EFFECTIVE IT SHOULD BE DRIVEN BY EXECUTIVE MANAGEMENT AND IT SHOULD MANDATE WHAT NEEDS TO BE ACCOMPLISHED

Policies and procedures

It is now time to start shaping the security program by writing policies, standards, guidelines and procedures.

Policies are high-level statements relating to confidentiality, integrity and availability of information across the organization. Some organizations develop one large security policy, while others develop smaller, more specific policies.

For a policy to be effective it should be driven by executive management and it should mandate what needs to be accomplished. The policy needs to be enforced and accepted by the entire organization.

Security policies are usually classified as roadblocks, and they should, therefore, be written in a fashion that is not too restrictive and is easy for a user to understand. Some specific policies might include data classification, incident handling, and disaster recovery.

Security standards are more specific and directly support security policies. Standards usually define security controls required to accomplish a goal dictated by the policy.

Guidelines should be viewed as best practices that support standards, and not as mandatory requirements.

Procedures are step-by-step instructions on how to implement policies, standards and guidelines.

To illustrate the relationship between policies, standards, guidelines and procedures let's use encryption as an example.

- A policy might state that data-in-transit must be encrypted
- A standard will support this policy by defining which cryptographic algorithms and protocols are approved for use by the organization. It might state that SSLv3 is not approved for use
- A guideline will go a step further and define the best practices for data-in-transit. It might state that TLS v1.2 should be used if at all possible
- A procedure might provide instructions on how to disable/enable certain protocols. For example, it would provide instructions on how disable SSLv3 on an Internet-facing web server.

Employee engagement

A critical element of a security program is security education and awareness. The organization can spend millions of dollars on the security tools to protect its perimeter, but what would happen if a user within the organization opened a malicious attachment? Users need to be educated so that they become proud of protecting the organization's assets and not paranoid. Humans are usually the weakest link in a security chain and cybercriminals are well aware of that. Users need to understand their important role in this chain and, for example, how to recognize social engineering and malware attacks.

Organizations should make this part of security program fun so that users are engaged and more willing to participate. Consider the example:

• Create a small security contest. Ask your users to provide the top 10 places that they have heard users usually hide their passwords.
• Create another small security contest. Challenge your users to decrypt an information security related word.
• Offer a gift card for contest winners.

Some organizations stop at this point of the security program. However, the security program is a living thing and it doesn't remain fixed such as Java Array – it is never done. The program must be very flexible to adapt to the security threat landscape that changes on a daily basis. The best approach to address this variable landscape is by defining a security program lifecycle.

THE SECURITY PROGRAM LIFECYCLE IS NOT AN OPTION, BUT A REQUIREMENT FOR A SUCCESSFUL AND COMPREHENSIVE SECURITY PROGRAM

Program lifecycle

The security program lifecycle is not an option, but a requirement for a successful and comprehensive security program. It should bring continuous improvement to an organization's security posture.

Consider the following examples:

Organization Y made the decision to move some infrastructure into the public cloud. Now their external network perimeter extends even further. A good security program would now dictate a new risk assessment, implementation of new security controls, and an update to the security policies, standards, guidelines and procedures. The bottom line is that the program becomes a repeatable process.

In this example let's consider a vulnerability that was introduced on Organization's Y Internet-facing web server. Cybercriminals unleashed their vulnerability scanners and scanned Internet-facing servers and applications for vulnerable services.

This is one of the main reasons to ensure that the security program dictates how often the organization needs to re-assess their network externally as well as internally. It should include periodic risk assessment, quarterly vulnerability scanning and annual penetration testing.

Now it is time to start reporting back to management and other business units about the security program effectiveness. The regular updates should include the most significant risks, security goals and objectives. Every organization has its own approach to reporting, but what seems to be very effective with the executive management is a risk score.

Consider an example where executive management was provided with the risk score of 2.2 out of 5 for four months in a row, and the following month they are provided with the risk score 3.4. They will immediately ask why the score jumped, and what they can do to lower the risk score to a more favorable position.

An analogy that I often hear to simplify an information security program is a jigsaw puzzle approach. Look at the pieces of puzzle as different security components, such as firewalls, IDS and policies. Most organizations already have those pieces in place, but that approach does not address security in an organized and comprehensive way.

You can have all the puzzle pieces in front of you, but until you put it together you will not have a recognizable picture. Look at an information security program from the same perspective. In most cases you will have most of the pieces in front of you - you just have to assemble the puzzle the proper way.

In the digital world universe, it is not enough just to have a security plan in place; it needs to be communicated and updated constantly.

What good would a security policy do if users are not aware of its existence? The organization has to ensure that their security program is not caught off guard because there's a gap. Gaps are elements that cybercriminals seek out to exploit vulnerabilities, enabling cybercriminals to build bridges within the network to easily navigate around.

There are always tribulations during security program development and implementation, but you have to get in front of that first domino in order to prevent a domino effect.

Zoran Lalic is a Senior Security Engineer with extensive industry experience in information security program development, penetration testing, forensics analysis, vulnerability management, security architecture design and incident response. His experience spans environments of all sizes – small offices to global networks. Additionally, he helped companies become PCI DSS compliant. Zoran has been an active researcher of new techniques used to compromise networks.

Spyware/adware combo masquerading as AnonyPlayer hits Android users

If you suddenly start seeing random advertisements popping up on your Android device, you have likely been infected with adware. But if you're terribly unlucky, you might have also been hit with information-stealing malware.

Dr. Web researchers have recently uncovered and analyzed a Trojanized version of the legitimate AnonyPlayer media application (they dubbed it Android.Spy.510).

It works as expected, but in the background it collects the following information: the model of the mobile device, the SDK version of the OS, availability of root access, and login credentials for the user's Google Play account.

All this was obviously not enough for the criminals behind this piece of malware, as they also make it ask the user to install an app called AnonyService, which supposedly secures the user's privacy from third parties.

But users who choose to do so will actually install an advertising module, which will then ask the user to allow the use of the Accessibility Service. Those who don't find the warning problematic will be saddled with the advertising module.

The interesting thing about this module is that shows several behaviors aimed at making the compromise of the device less obvious.

For one thing, it waits a few days before springing into action. Secondly, it shows ads only when apps that are not on a hard-coded whitelist are run. For example, it will not show ads if you open the device's Settings, clock app, camera, contacts, and so on.

Every launch of an app that's not on the list will trigger the showing of an ad. "As a result, the owner of a compromised device may think that it is the launched application that is responsible for annoying notifications," the researchers noted.

Naturally, to divert suspicion from the Trojanized version of AnonyPlayer, no ads will be shown when that particular app is launched.

The legitimate AnonyPlayerand the Trojanized version can't be found on Google Play, but can be downloaded from third-party app stores.

Top malware families targeting business networks

Check Point has revealed the most common malware families being used to attack organisations' networks during October 2015. The top 10 malware families detected globally were:

1. Conficker – Worm that allows remote operations and malware download. The infected machine is controlled by a botnet, which contacts its Command & Control server to receive instructions.
2. Sality – Virus that allows remote operations and downloads of additional malware to infected systems by its operator. Its main goal is to persist in a system and provide means for remote control and installing further malware.
3. Cutwail – Botnet mostly involved in sending spam e-mails, as well as some DDoS attacks. Once installed, the bots connect directly to the command and control server, and receive instructions about the emails they should send.
4. Neutrino EK – Exploit Kit that can be used to attack computers using versions of the Java Runtime Environment. Attacks involving the Neutrino Exploit Kit have been associated with ransomware scams.

5. Gamarue – Used to download and install new versions of malicious programs, including Trojans and AdWare, on victim computers.
6. Agent – Trojan which downloads and installs adware or malware to the victim's machine. Agent variants may also change the configuration settings for Windows Explorer and/or for the Windows interface.
7. Pushdo – Trojan used to infect a system and then download the Cutwail spam module and can also be used to install additional third party malware.
8. Alman – Virus which infects all executable files in the system. The virus propagates over the network and also has rootkit capabilities.
9. ZeroAccess – Worm that targets Windows platforms allowing remote operations and malware download. Utilizes a peer-to-peer (P2P) protocol to download or update additional malware components from remote peers.
10. Fareit – Trojan used to steal sensative information such as user names and passwords stored in web browsers, as well as email and FTP credentials.

Vonteera adware blocks AVs, can install uninstallable Chrome extensions

The Vonteera adware family has been around for quite some time, but it is now slowly starting to cross the line between unwanted, potentially malicious software to malware.

According to Malwarebytes researchers, the adware has a new trick in its sleeve: it adds 13 certificates to the targeted systems' "Untrusted Certificates" list, and they all belong to companies developing popular AV and security software such as Avast, AVG, Baidu, Bitdefender, Malwarebytes, Trend Micro, and others. The list is used by Windows' User Account Control (UAC) to keep out untrusted software.

"The effect of this is potentially devastating since your system will refuse to run any applications signed with these certificates," Malwarebytes researcher Pieter Arntz explained.

This means that an affected user will have trouble cleaning their systems from malware - the fact that this happens only if Vonteera has managed to infect a system without triggering AV software means that the user either doesn't use protection or that it's not that good.

So what can users do to get rid of it? One option is to go into Certificate Manager and delete the certificates in question, then download an AV solution - preferably one developed by the aforementioned manufacturers, as they obviously detect the adware - and run it to find and remove the offending software.

"Make sure to check back on the certificates after you have removed the adware, since they might have been re-instated in the meantime," Arntz advised.

Another option is to temporarily disable UAC so that the needed AV can be downloaded, installed and run, or to use Task Scheduler to bypass UAC.

ModPOS: The most sophisticated POS malware to date

Elements of ModPOS date back as far as early 2012. It targeted US retailers in late 2013 and throughout 2014, and is expected to continue to do so in the future. According to iSight Partners, the malware is responsible for the theft of information tied to millions of payment cards so far.

How did it remain (mostly) hidden for so long, you ask? Well, the truth is that it's extremely stealthy, as well as extremely sophisticated, and malware analysts have been having a hell of a time reverse-engineering its modules.

The malware's individual modules are typically packed kernel drivers, which makes them difficult to detect. So far, researchers have managed to discover three of them: a downloader/uploader, a keylogger, and a POS scraper module. And only the downloader/uploader is detected (as Straxbot) by a single AV solution.

ModPOS also sports several plugins that are meant to collect information about the target system, about the domains, computers and network resources available to the infected system, and username and password information for local and domain accounts. All this information is sent to the attackers.

"From a coding perspective, these samples are much more complex than average malware; there is professional-level coding, and the size, implemented operational security and overall characteristics of the code likely required a significant amount of time and resources to create and debug, and an advanced understanding of how to undermine security identification and mitigation tools and tactics," the researchers found.

The drivers inject malicious code into a variety of processes, including system, winlogon.exe, firefox.exe, and credit.exe.

"The credit.exe process is notable and related to stealing credit card track data from the POS system's running memory. This is unique to POS vendors that use this executable as a part of their software," the researchers explained.

"[We are] confident that the actors customize the malware based on the targeted environment. This malware can also log keystrokes, upload stolen information and download other malware payloads. It uses AES-256-CBC encryption for data storage and transmission, and the encryption key is uniquely generated per victim system."

The researchers believe the authors have ties with Eastern Europe.

"ModPOS, and most POS malwares, have increased in sophistication. In September and October of 2015, there were several discussions within hacker forums to share information about current POS code and requests for assistance to add more functionality and test the results. The hacker community has been very active sharing information, conducting test, tweaking code and re-testing since the summer months…all preparing for the Holiday shopping season," says Paul Fletcher, cyber security evangelist at Alert Logic.

"In my opinion, the main points of interest about the increased sophistication of POS malware are the use of encryption and the "anti-forensics" (aka obfuscation or anti-analysis) concepts."

"The use of encryption by the attacker has been a long time coming, and it's interesting to me because one of the best practices for security professionals is to use encryption where possible. While some organisations have been slow to adopt the use of encryption, the hacker community embraces this concept and it gives them an edge. This point shows that tools and technology are generally the same being used by attackers and security professionals, giving more proof that security technology solutions alone aren't enough, people and process built around those security technology solutions are essential," he pointed out.

"The anti-forensics component of sophisticated malware is an indication that the hacking community has done extensive reconnaissance on multiple POS systems, as well as the support systems (back-end) within the retailers infrastructure."

Exploit kit activity up 75 percent

The creation of DNS infrastructure by cybercriminals to unleash exploit kits increased 75 percent in the third quarter of 2015 from the same period in 2014, according to Infoblox.

Exploit kits are a particularly alarming category of malware because they represent automation. Highly skilled attackers can create exploit kits, which are packages for delivering a malware payload, and then sell or rent these toolkits to those with little technical experience - vastly increasing the ranks of malicious attackers capable of going after individuals, businesses, schools, and government agencies.

Four exploit kits - Angler, Magnitude, Neutrino, and Nuclear - accounted for 96 percent of total activity in the category for the third quarter.

Most exploit kit attacks are distributed through spam emails or compromised web sites, or are embedded in online ads. When users click a link in the emails or ads, the exploit kit takes advantage of vulnerabilities in popular software to deliver a malware payload that can perform actions such as planting ransomware, capturing passwords for bank accounts, or stealing an organization's data.

Cybercriminals need the DNS to register domains for building the "drive-by" locations where exploit kits lie in wait for users, and for communicating with command-and-control servers that send instructions to infected devices and extract information.

"Exploit kits are behind some of the highest-profile attacks in recent months," said Craig Sanderson, senior director of security products at Infoblox. "For example, a recent Angler attack on Mail Online implanted malicious ads on the site for five days, potentially exposing millions of online visitors to infection."

Rooted, Trojan-infected Android tablets sold on Amazon

If you want to buy a cheap Android-powered tablet, and you're searching for it on Amazon, the best thing you can do is carefully read all the negative reviews you can find. If you are lucky, you'll see some that will warn you about the device being rooted and coming pre-installed with malware.

Security researchers from Cheetah Mobile have recently discovered a slew of these devices - over 30 tablet brands in total - being sold on Amazon and other reputable online stores.

The malware in question is the Cloudsota Trojan, which allows remote control of the infected devices and conducts malicious activities without user consent.

It can install additional adware or malware, uninstall anti-virus and other security apps. It has root permissions, so it can automatically open all the additional apps it has installed. It also replaces the boot animation and wallpapers on the devices with advertisements, and changes the browser's homepage and redirects search results to strange ad pages.

Worst of all, even if the user manages to remove it, it will reappear after each reboot of the device.

The researchers posit that the attackers who did this are from China, as much of the Trojan's code is written in Chinese, its C&C server is registered in Shenzhen, and the manufacturers of tablets are all from China.

"According to our rough estimation, at least 17,233 infected tablets have been delivered to customers hands," they noted, but added that since many tablets are not protected by AV apps, the number could be much greater. "These tablets share some similarities that all of them are low-priced and manufactured by nameless small-scale workshops."

The devices have been shipped around the world, but Mexican, USA and Turkish buyers were most hit. The researchers have notified Amazon and other online retailers of the problem, and have advised manufacturers to investigate their system firmware.

UK parliament's secure network hit with crypto-ransomware

The UK parliament's secure network has been breached and several computers on it have been compromised by hackers, The Times has reported.

The newspaper makes is sound like the attack was targeted, and the target was Chi Onwurah, MP for Newcastle upon Tyne Central and the Labour Party's shadow minister for culture and the digital economy, and her employees. But they also say that they hackers compromised the systems by delivering and spreading crypto-ransomware, which ended up encrypting sensitive files stored in a shared drive used by some 8,500 government employees (MPs, lords, staff, etc.).

After MP Onwurah was faced with the ransom note, she reported the problem and the Parliamentary Digital Service (PDS) first cut her off the shared drive, than took her computers and cleaned them up by replacing the hard drives.

The MP said that no files containing information about constituents were compromised. The incident happened in May, and spurred the MP to start an investigation in order to find out the extent of cyberattacks on MPs - especially targeted ones - and how well the defenses put in place by the PDS are working to fend them off.

Exposing Rocket Kitten cyber-espionage group operations and targets

Check Point identified specific details and analyzed cyber-espionage activity conducted by the group Rocket Kitten, with possible ties to the Iranian Revolutionary Guard Corps.

Led by researchers in Check Point's Threat Intelligence and Research Area, the data paints a picture of strategic malware attacks supported by persistent spear phishing campaigns. The details show the Rocket Kitten group actively targeted individuals and organizations in the Middle East, as well as across Europe and in the United States, documenting specifics such as:

- Business and government sectors across Saudi Arabia, including news agencies and journalists; academic institutions and scholars; human rights activists; military generals; and members of the Saudi royal family
- Embassies, diplomats, military attachés and 'persons of interest' across Afghanistan, Turkey, Qatar, UAE, Iraq, Kuwait and Yemen, as well as NATO commands in the region
- Dozens of Iran researchers, as well as European Union Iran research groups, specifically in the fields of foreign policy, national security and nuclear energy.
- Venezuelan trade and finance targets
- Former Iranian citizens of various influential positions
- Islamic and anti-Islamic preachers and groups; famous columnists and cartoonists; TV show hosts; political parties; and government officials.

Researchers were also able to trace and unmask the true identity of an aliased attacker, identified as "Wool3n.H4T," as one of the prominent figures behind the campaign. Further, based on the nature of the attacks and their repercussions, researchers suggest Rocket Kitten's motives were aligned with nation-state intelligence interests, aimed at extracting sensitive information from targets.

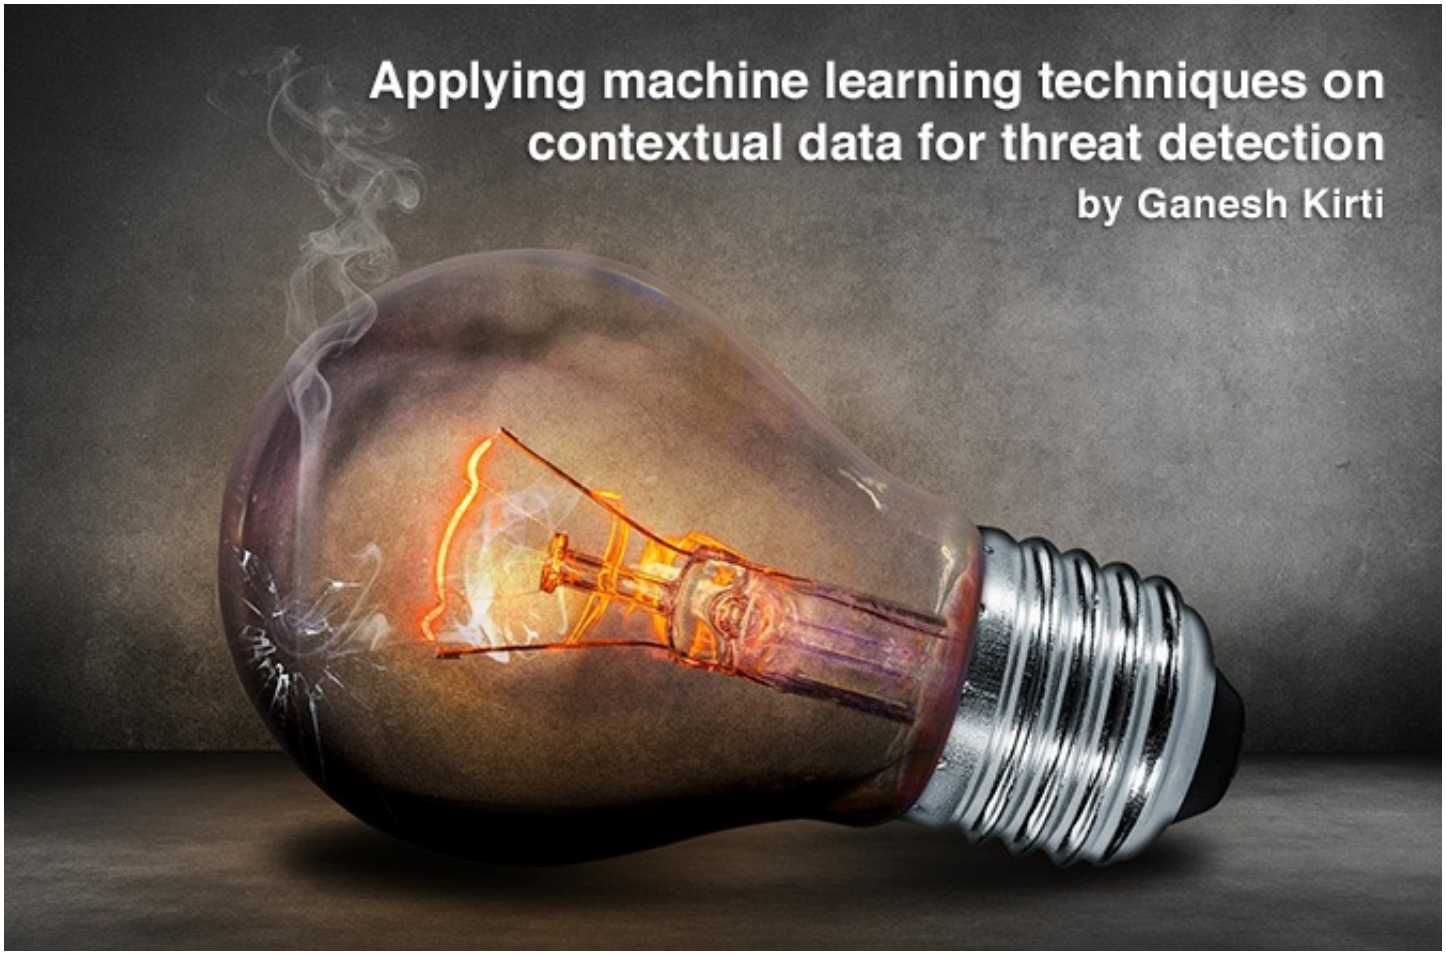

Applying machine learning techniques on contextual data for threat detection

by Ganesh Kirti

The momentum behind cloud computing couldn't be stronger as companies, governments and other organizations move to the cloud to lower costs and improve agility. However, you need look no further than the headlines about the latest data breach to know how extremely important security architectures are amid this rapid cloud adoption.

The question is on every CIO's and security officer's mind: What are the most efficient techniques to detect threats to cloud services?

Security technology is advancing to answer the challenge. Machine learning, threat intelligence and predictive analytics are among the combination of techniques being used to catch bad actors. Enterprises also can efficiently detect threats by using application and situational context in conjunction with machine learning techniques to reduce one of the biggest pitfalls of threat detection – false positives – and ultimately heighten security across the board.

The first thing to do to start wrapping your head around cloud security architecture is deciding what to monitor. Remember, threat landscape is dynamic in cloud workloads. Every source of activity should be monitored. This includes configurations, APIs, end users, administrators/privileged users, external federated users, service accounts and type of transactions made by users. Everything.

Second, it's essential to understand why context is important. It's the only way to understand the threat severity and decide whether a specific event or particular user behavior is anomalous. Examples of context are: a business user performing mass delete of objects after hours, a part-time contractor performing administrative operations in multiple cloud applications, an engineer cloning a source code repository from an unknown location.

By implementing a comprehensive approach - activity monitoring and user behavior analysis, and considering the context in which

those events happen - organizations can be confident that their clouds are secure. The strategy should follow these six tactics:

1. Threat analytics and detection architecture. It starts with an architecture that can analyze data from various sources to derive early indicators of threats. This architecture should accept data feeds from all of the sources mentioned above. Analytics architecture should leverage machine-learning techniques to efficiently consume data to identify anomalies. A combination of supervised and unsupervised techniques should be used.

2. Security configurations. The security posture of a service depends on how stringent security configurations are. A weak security configuration provides an entry point for malicious users. Examples of risks due to weak configurations are: administrator users with weak passwords, over-permissive access to servers, and anonymous users accessing sensitive content. It is important to configure stringent values and continuously monitor those values for drifts.

3. Contextual data feeds. A particular risk event should be analyzed in the context of occurrence. If the context is not used, then one will end up with high false positive rate. For example, alerting about a user with anomalous behavior by just looking at her login data in AD is insufficient. For improved accuracy, the user login behavior in the login session should correlate user attributes such as: transaction type, how sensitive the transaction is, is the user travelling, is the user a part-time employee, and what user roles are. Contextual data helps improve threat detection accuracy.

4. User behavior analytics. User behavior analytics model and analyze user-centric behavior. Users in the analysis include both end users and privileged users. A highly privileged user or an end user with access to lot of cloud services is in general a high risk. It is important that high-risk users are monitored continuously by adding them to a watch list. Their behavior, the strength of their passwords, the authentication policy and all sensitive privileges should be monitored and adjusted to avoid risks created by their activity.

5. Supervised and unsupervised machine learning techniques. Machine-learning techniques should be used to define a baseline and detect outliers. A practical approach is to use both unsupervised and supervised models to improve accuracy and reduce false positives. Many implementations use one or the other, causing a high false positive rate or an issue when their solution does not scale by demanding high volume of labeled data.

To improve accuracy and scalability of threat detection, use unsupervised learning to model clusters of users with normal behavior. Statistical and probabilistic mixture models are practically proven for this purpose and subsequently detect outliers that represent users with abnormal behavior - i.e. risky users. Also, use supervised models to get hints from security experts for determining risky users' patterns and actions. Based on these hints, build benchmark datasets for training, validation and testing of supervised models.

Though supervised models require more security expertise and manual efforts, they tend to present a lower false positive rate than the unsupervised. The best practice is to increase the effectiveness of supervised modeling by an unsupervised data pre-processing step that usually identifies highly risk users with a fair false positive rate that is minimized with the subsequent supervised learning models.

6. Threat intelligence feeds. Real-time collaboration with security communities and commercial intelligence feeds help detect threats at an early stage. For example, a hacker accessing an application using compromised user credentials from a blacklisted IP address can be detected if external intelligence feeds provide blacklisted network information.

As an organization's cloud footprint grows, it's vital to take a comprehensive approach to security that encompasses machine learning, threat intelligence, predictive analytics and context.

Ganesh Kirti is the CTO of Palerra (www.palerra.com).

Why governments need to take the lead in cybersecurity

by Brian Honan

Time and time again we hear people lament about the impact cybercrime has on our businesses, our individual lives, the economy, and on society. Report after report show the impact cybercrime is having on our economies, with some estimating the global cost of cybercrime is approaching $3 trillion per year. As each of these reports is published, there is the usual handwringing over why the state of cybersecurity is so bad.

We blame companies for not protecting our personal data properly, we blame the vendors for producing ineffective solutions that do not address our problems properly, we blame standards bodies for developing standards and frameworks that address only the basic elements of security, we blame users for falling victim to phishing emails and other scams, we blame law enforcement for lack of action and/or capability in dealing with cybercriminals, we blame academia for not training students in the proper skills or not conducting research in the proper areas, and finally, we blame criminals for conducting these attacks.

There is one group that I often see missing from all of the above finger pointing and arguably this group has the most influence in how we improve cybersecurity and how we tackle cybercrime: the governments of each of our countries. For the past number of decades, governments have failed to recognize or even acknowledge that cybersecurity is an important issue. The collective attitude has been that cybercrime or cyberattacks were not an issue that governments should be concerned with and that individuals and companies should protect themselves.

It is this short-sightedness that has led us to the poor state of cybersecurity we now face.

Lack of leadership and investment into cybersecurity by governments has resulted in many law enforcement agencies lacking the appropriate capabilities and resources available to tackle cybercrime. This lack of leadership has also resulted in many government systems being less secure than they should be.

It is said “nature abhors a vacuum” and so, too, does leadership. Without leadership from our governments, the private sector has stepped into the role of defining what good security practice is and we now have countless standards all competing for our attention. Due to the lack of resources and skilled staff, law enforcement agencies have had to look to private sector companies to bolster their capabilities. We regularly see security vendors working with law enforcement to take down botnets and disrupt online criminal activity. These services are offered to augment the technical capabilities of law enforcement and are often provided at no cost.

The value for the security companies is the media attention they get for doing this work. Law enforcement agencies welcome the help, but this practice highlights the severe lack of funding by governments in this area. When the marketing budget a security vendor can spend on its involvement in botnet takedowns exceeds the annual budget that the law enforcement cybercrime units receive, there is something seriously wrong with our priorities.

In effect, private sector companies are the ones who are driving the cybersecurity agenda and not governments. The danger is that the cybersecurity agenda will be driven by the goals of the private sector companies involved, which in many cases do not align with the greater requirements of society. We have seen companies create a niche in the market for their services and then campaign that their services should be government policy. The push by a number of companies promoting hacking back as a valid approach to deal with a cyberattack is a good example of this.

Due to the lack of resources and skilled staff, law enforcement agencies have had to look to private sector companies to bolster their capabilities.

But the biggest concern is the practice by security vendors of quickly attributing attacks to certain nation states based only on the information those private companies hold. As a result, we see press release after press release saying that certain countries are the source of major attacks, often with only the flimsiest pieces of evidence to support those claims. Time after time we have seen so-called facts and evidence from vendor reports being used to support political arguments, and then later witnessed that evidence being refuted.

This constant flow of “news” stories, no doubt supported by political lobbying on behalf of those cybersecurity companies, runs the risk of shaping public and political opinion on how government foreign and domestic policy should be formed in relation to cyberattacks. When government policy in relation to cyber security is based on marketing reports and press releases from private sector cybersecurity firms, we are opening ourselves to major problems in the future.

As security professionals, let’s make sure that when we see companies making their marketing propaganda part of the political agenda we call them out on their hype with fact-based arguments.

As private citizens, let’s make sure we lobby our politicians to take cybersecurity seriously and highlight to them where the real issues lie.

It’s time our governments focused their priorities on developing better policies regarding cybersecurity, so let’s make sure they develop those policies based on the greater needs of society and not the marketing requirements of private companies.

Brian Honan (www.bhconsulting.ie) is an independent security consultant based in Dublin, Ireland, and is the founder and head of IRISSCERT, Ireland's first CERT. He is a Special Advisor to the Europol Cybercrime Centre, an adjunct lecturer on Information Security in University College Dublin, and he sits on the Technical Advisory Board for several information security companies. He has addressed a number of major conferences, wrote ISO 27001 in a Windows Environment, and co-authored The Cloud Security Rules.

Ernest Hemmingway once said, “I have learned a great deal from listening carefully. Most people never listen.” Perhaps, like most of the things we do, technology will absolve us of that requirement too – it will listen for us. In fact, it seems that soon, technology will be listening to us all the time, everywhere.

Whether we’re talking to Siri, Cortana, Jibo, Google, Alexa, or one of many other “recognition” technologies designed to understand and respond to our speech, there are an increasing number of things listening.

As a species, voice has always been the primary way we communicate (with apologies to body language experts everywhere.) The use of voice to ask questions, exchange ideas, and issue instructions is profoundly, although not uniquely, human. It appears that as far back as 300,000 years ago ancient human ancestors where beginning to communicate through the use of a language of some kind. And we haven’t stopped since.

So commanding the world around us through the use of language is natural. Once devices are smart enough, there are a lot of advantages to telling them what to do rather than having to show them, by clicking, dragging, poking, pinching and swiping. For one, it leaves our hands free to do something more important, like grip the wheel of the car, or at least open a fresh bag of chips.

Most importantly, in the home and office environment, it cuts the tether that currently exists between us and the technology we are interacting with. No longer will we have to walk over to a device to interact with it. We’ll be able to tell it what to do, and let it get on with it.

Home IoT technology, whether that’s a Google Nest, an Amazon Echo, or the cute-and-chipper Jibo, offers the promise of managing your home life (or at least your shopping habit) without needing to manually interact with many of the actual devices.

Yelling “Hey Siri” or “OK, Google” is a first step in this direction, but of course we can (and therefore inevitably will) go much further.

Technology that listens to us will permeate our homes, cars, and offices. There's already smartcards for office workers that listen to the conversation around them and figure out how to make you more productive, what the general mood is, and what people are working on.

Which leads us to the inevitable question – who exactly are we talking to? Because of course, we're not really talking to a device at all. Rather, we'll be communicating with an entire eco-system of services, all tuned to parse our words, our gestures, and therefore our desires. And all that data will aggregate over time to let the masters of those same services understand us more profoundly than we could possibly imagine today.

Take that next logical step and your home automation device (or devices) of choice will quickly want to follow suit. Why bother having to tell the TV (another device who increasingly will be eavesdropping on your fireside chat) what channel to watch, when your home automation center can decide, based on what its hearing, what kind of mood you're in, who's in the room with you, and therefore what it should recommend for entertainment?

Listening devices offer the opportunity to far more deeply understand us as consumers, customers and people. And the ability to respond through speech and activity, Jibo-like, will make us more and more comfortable with sharing more and more with these machines.

The machines will know us better than anyone because, unlike everyone else, they will be with us 24/7. Our online footprint will be immense, because every facial twitch, every roll of the eye, every suppressed yawn will be stored, reviewed and quietly filed.

And we'll be happy to share it.

The machines will know us better than anyone because, unlike everyone else, they will be with us 24/7

We're trained to share with tech – to have radio stations build playlists for us, to have websites suggest things we might be interested in, or would like to buy, all based on our behavior. So the idea that the IoT-enabled devices in our home would learn about us, would listen and respond, is really only a short, highly convenient step away. And as low-power, convenient technology for voice recognition designed for IoT devices becomes available, so the ease with which everything around us can listen in becomes even greater.

Of course, the obvious questions about who can also listen in, and who gets to keep a copy of our inner-most desires, is really at the root of concerns over this trend – this move to ever smarter, ever more engaged technology.

Can we trust the machines to only listen when we want them to? If we trust the manufacturers, can we trust that the devices themselves are sufficiently secure to fend of malicious hacking attacks? And even if all the above is true, can we also trust governments not to extend the long arm of national security interest into our private conversations? Does it even matter? If we are happy to share, should we really worry about the privacy implications?

As the IoT starts to become part of our daily lives, we will need to become comfortable with, not only the devices themselves, but also the constant trade-off we must make between utilizing the promise of the IoT, and managing the loss of control and privacy it will entail.

Will trade privacy for services!

The full value of the IoT isn't that it's just a bunch of smart gadgets sitting around talking to each other – the real impact is the way the IoT will change the fundamentals of how we think about and interact with technology.

Because no longer will we go log in, or switch on when we want to do something. Instead the IoT will surround us with a smart, communicating and response fabric that manages and monitors almost everything we do, that build products better, controls our environment, and deals with our wants and needs, often before we are aware of them fully.

The IoT, in other words, delivers the full promise of its value when it is completely embedded in our world. And that process requires us to open the door to a very different way of thinking about information, and especially privacy.

Ultimately, we're doing what we, especially in the West, have been doing for a very long time: we're trading. I have something to sell – access to my thoughts and needs, and you have something to give me for it – better services, more targeted offers, cut rate delivery, more interesting TV viewing, maybe safer homes and smarter kids.

There's actually nothing wrong with choosing the hand over some of our privacy to the next generation of smart tech (or, to be optimistic, probably the generation after that) but we need to do so with open eyes. And when they tell us what these devices do, and how they do it – when they explain what kinds of information they gather and who the device shares it with – we better listen very closely.

Geoff Webb is the Vice President, Solutions Strategy, for NetIQ, the security practice of Micro Focus.

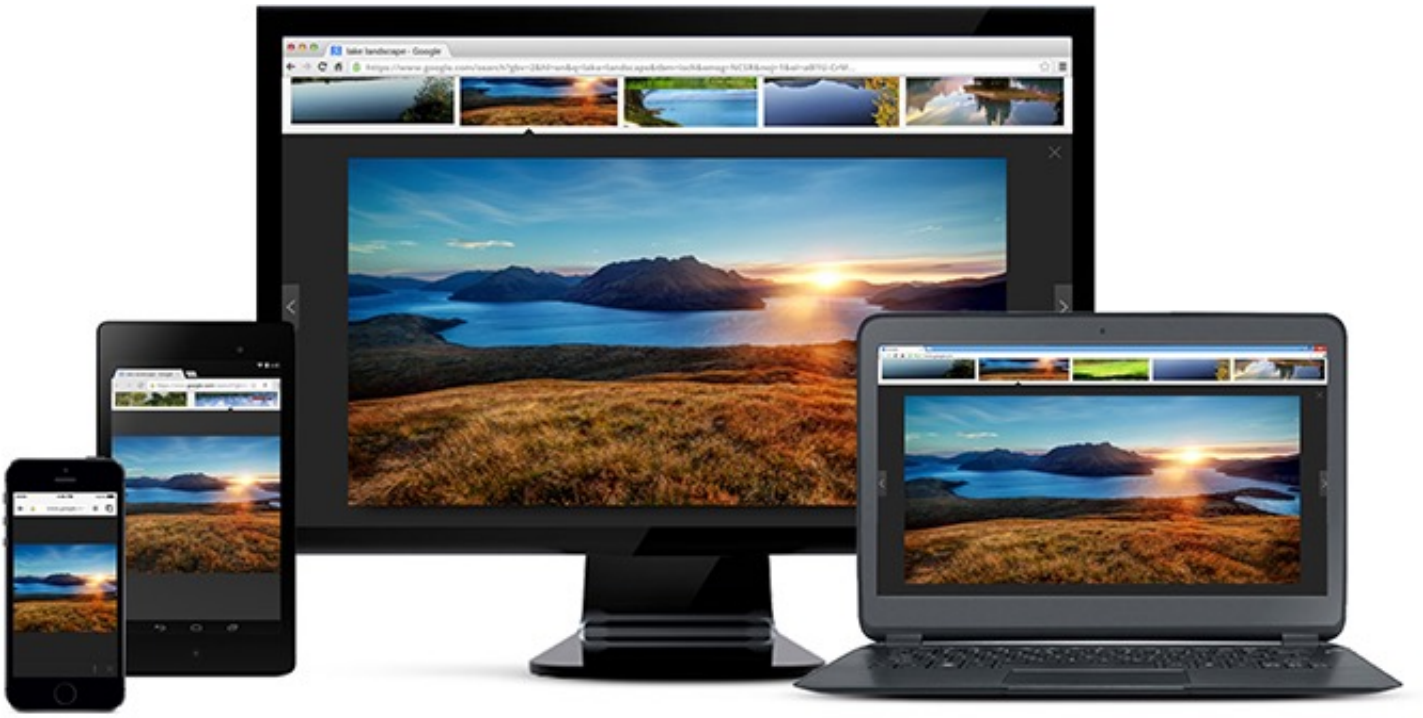

Why I recommend Chrome to family

by Matt Jones

I recommend Chrome often for browsing to friends and family who are concerned about malware and the like. Here's my rationale: I believe the Google security team has done an excellent job maturing it over the years.

In this article I first look at the importance of secure software architectures and hardened platforms, and then look at a few examples of important and strategic security research that's helped strengthen its security posture when it comes to protecting system integrity.

The second part of the article goes over a few relevant considerations for those managing security threats, and cautions about the "drain the swamp" approach, which I believe to be a pitfall when preparing threat management security activities.

Secure architectures

Defining logical boundaries in software and then implementing containment and platform hardening has slowly migrated, for well over a decade, from critical server side daemons to client-side software.

One of the more commonly known daemon implementations is in OpenSSH with privsep by Niels Provos, however there are numerous other examples around this time. One of the pioneers of privilege separation was Chris Evans. Design notes for his vsftpd (very secure FTP daemon) include this great excerpt:

"Unfortunately, this is not an ideal world, and coders make plenty of mistakes. Even the careful coders make mistakes. Code auditing is important, and goes some way towards eliminating coding mistakes after the fact. However, we have no guarantee that an audit will catch all the flaws."

As a teenager I remember looking at Mark Dowd's pre-authentication OpenSSH challenge-response bug back in 2002, and how if privilege separation was enabled it would land an attacker in a chroot jailed and unprivileged child process. This type of hurdle is an immediate mitigating factor requiring a pivot to break out of the unprivileged jailed child process and continue, with the kernel being a popular target.

From my perspective, strengthening kernels has been a critical pursuit for many years, and I've personally seen the grsecurity/PaX

project as a key project to help set a standard for major operating systems to harden user-land, tighten low-level OS controls and help raise the cost of attacking bug classes or boundaries. As the cost of attack rises, some bugs become less valuable, and many attackers will likely move to another vector.

Over the past several years, client-side software and mobile platforms have adopted similar containment approaches, and the same types of principles often apply. Watching Pwn2own unfold over the years confirmed this to me - targets that don't adopt a security conscious architecture in their foundation fall quickly and cheaply as bugs can have an immediate critical impact.

When Vupen demonstrated exploiting Flash and first took out Chrome at Pwn2own in 2012 via the non-sandboxed plugin, it showed the importance of strong software architectures and imitated an approach typically taken by attackers (taking the cheapest path).

Chrome maturity

Hardened systems and secure architectural decisions in software indicate a mature approach to building secure systems, and appears to be often coupled with a good understanding both of the immediate attack surface but also of attack chains.

Chrome is an excellent case study for this topic, and enough time has passed that we can look back at a few things and see how it all turned out.

Google has an internal Chrome security team and, undoubtedly, a lot of work happens behind the scenes. In this article I'd like to highlight a few interesting pieces of public technical research, which I think has been important to help evolve their security posture in a strategic way.

Abstract vectors

The Internet Printing Protocol (IPP) code in the Common Unix Printing Services Daemon (cupsd) has been for quite some time an interesting potential vector for attacking Chrome on Linux, but the whole cross-origin thing was a fairly effective excuse for me to procrastinate. Then one morning I woke up and read an amazing guest Google Project Zero blog post by Neel Mehta, who pulled off that attack in style.

This work is great for understanding abstract vectors - via a Cross Site Scripting (XSS) in the CUPS templating engine, a malicious webpage could interface with cupsd on the localhost, bypass the browser security constraint, then open a new vector to exhaust and attack the huge C-based IPP parsing codebase that sat quietly and peacefully for years thanks to the old binding-to-localhost obscurity technique.

This attack is platform specific (multiple Linux flavors and OS X run CUPS by default) and also not an immediate attack surface, but the attack requires a level of understanding of the full stack and the skill to uncover exploitable bugs in this attack path. At the end of Mehta's post is a perfect example of follow-up hardening specific to this vector and I highly recommend reading it.

This raises a point about mitigating bypasses, and also non-default configurations. The latter case may not be the most appealing type of vulnerability to pursue, but this lack of appeal can also make it an easier attack path in many code bases.

Library attack surface

There is also the problem of libraries and external dependencies - many libraries are there for heavy lifting, and often come with a significant attack surface (e.g. WebKit). Attacks against them have proved fruitful for countless of attackers who got in early, ran a web-rendering fuzzer, and found many parsing related vulnerabilities. Another example is libstagefright in Android - a ripe attack surface that spans many projects, and where the impact is generally critical.

However, for targets such as Chrome, on modern up-to-date platforms, the first hurdle for exploitation is usually user-land mitigation controls. And then, if there is sandboxing of the library component, an attacker will likely require a pivot to affect the integrity of the system.

Back in 2010, at CanSecWest, Tavis Ormandy and Julien Tinnes gave their "There's a party at Ring0, and you're invited" talk, and presented their (at the time) recent work on researching kernel security in both Linux and Windows. One key point mentioned was the importance (or even dependency) of kernel security for user-land sandboxing, which is relevant for Chrome's sandbox to help protect system integrity. The topic of attacking the kernel attack surface (e.g. GDI) via Chrome to bypass the sandbox and user-land mitigating controls to me highlights a lot of consideration in both attack surface across privilege boundaries and exploit economics.

The Chrome team has responded to their complicated threat landscape quite well, by building strong foundations and by performing research-driven offense and defense activities

Target specific attacks

One of the most valuable aspects of running events like Pwn2own is, in my opinion, that of getting in-depth attack trees, which demonstrate intricate vectors that run through the components of complicated software.

The two part blog series by Chrome software engineers dubbed "A Tale of Two Pwnies", published in May and June 2012, shows off the work of two individuals (Pinkie Pie and Sergey Glazunov) who have dedicated time and technical skill to understand software beyond a shallow layer, and then crafted effective attacks (by exploiting somewhat flamboyant attack paths).

Receiving such specific technical details and sitting down with an attacker to talk about his approach and execution is incredibly valuable - gaining insights into an attacker's intuition can extend a threat model beyond what's been initially anticipated. In my opinion, this approach should be used whenever conducting red-team style penetration tests, as documenting intuition and approaches from the testers can end up being more valuable than the findings of the tests.

When reaching this level of maturity there is an idealistic state where both vulnerabilities and exploits become target specific (and not generic bug-classes/exploit techniques). This results in the bar rising substantially and, consequently, the threat actors shifting.

Drain the swamp

Now that we've looked at a few offensive examples that have helped mature Chrome, I'd like to take a look at the "drain the swamp" analogy used by Haroon Meer in his 2011 44-Con presentation titled "Penetration Testing considered harmful":

"When you're up to your neck in alligators, it's easy to forget that the initial objective was to drain the swamp."

The analogy can be summarized simply as "losing sight of your initial objective", and I consider it perfect when applied to scoping defensive activities in security.

While all activities can serve a valuable purpose, it's important to be a realist and carefully consider your objective and your current level of maturity before jumping in.

I believe the Chrome team has responded to their complicated threat landscape quite well, by building strong foundations and by performing research-driven offense and defense activities. However, prioritizing security activities requires careful thought on a case-by-case basis.

Penetration testing and public bug bounties

Penetration testing can take many forms and is a fundamental activity for identifying weaknesses and guiding improvements. I see bug bounties as a fairly progressive activity that can show maturity and promote an open channel to security researchers.

Having third parties perform penetration tests and running public bug bounties comes with a set of pros and cons. Before deciding to do either of these things, careful evaluation is needed, and once the decision is made to move forward, adequate preparation has to be performed in order for these activities to bring real ROI.

In his talk, Meer pointed out that penetration testers also introduce risk. This brings up an interesting point: When public bug bounties are aimed at securing applications and infrastructure, money rewards for bugs can turn bug bounty hunters into the new primary threat actor. In case they get access to systems or internal data, how can we trust their data safeguarding methods and maturity? What new risks does this introduce?

Web-based attacks generally have a lower bar for entry, and there's a lot of resources to help get people started and grow skills quickly. Yet, sometimes, we see companies paying high amounts for common bugs that are relatively cheap and easy to identify and exploit. Perhaps, in such instances, when the ROI for bugs is so askew while also introducing risk to the organization, focusing on other activities first may have been wiser.

Another point in Meer's presentation is how a few penetration testers could each successfully compromise a target but all take different paths. Testers and bug hunters are often objective-focused rather than assurance-focused - their objective is to successfully exploit a bug in something important and get paid.

Focusing on assurance, which would be preferred, would mean to focus on breadth, and then prioritize depth and coverage for each layer, based on experience of what's most likely for that target.

All security activities will have varying levels of ROI and understanding the threat landscape and current maturity helps to prioritize a budget.

Systemic patterns

From both my own experience and observing publicized vulnerabilities over a number of years, systemic patterns are repeated throughout branches, and this is why it's important to leverage both whitebox (i.e. code review) and blackbox security assessments.

Fabian Yamaguchi's excellent paper titled "Vulnerability Extrapolation: Assisted Vulnerability Discovery Using Machine Learning" was to me a refreshing piece of work that tries to dig layers beneath the surface, which can have a wider tangible effect than squashing an individual bug. I mentioned the OpenSSH challenge-response bug earlier - this vulnerability had an updated advisory because that same int overflow construct had to be patched in the PAM module too (but separately). While blackbox testing can successfully unearth bugs, a whitebox or greybox test may be able to find the systemic patterns where these flaws originate and they can then be addressed in a more broad-sweeping way than isolated patching.

And this brings up an interesting case about fuzzing for vulnerability discovery. In his "Babysitting an Army of Monkeys" presentation, Charlie Miller talked about how he made a dumb-fuzzer with several lines of Python, and how he successfully used it to crash multiple high-target client-side applications. In his slide-deck however, he has some interesting fuzzing statistics on the number of crashes that were unique and the manifestations of single vulnerabilities, which were actually quite interesting.

Motivation and beyond bugs

Michal Zalewski recently wrote on his blog about "Understanding the process of finding serious vulns" and announced that he started an informal survey: he's reaching out to the discoverers of high-impact flaws in commonly used software and asks them about their methodology, vendor communication channels, bug disclosure decisions, and their

motivation for the research. These are important questions and it will be interesting to read the results.

In my opinion it's important to factor in motivation and even experience when managing submissions. An experienced, assurance-focused bug hunter may (needlessly) spend days or more preparing a PoC for an individual bug so that it is treated seriously, instead of maybe just being allowed to (due to his or her experience) to explain at code-level the potential of the flaw.

If it's known that a bug hunter focuses on assurance, it may be easier and better to keep them focused on creating a more intimate understanding of the target – an understanding that is needed to come up with results beyond typical vulnerabilities with accompanying exploits.

There's a related point where for some targets it should be beyond "bugs", and also encompasses attack surface reduction and defense in depth approaches for potential risks that are very costly to pursue or even theoretical in nature.

Sometimes such issues can exist for some time and due to a small change in circumstance (a code change or a new technique surfaces) a potential technical risk can become a vulnerability that's practical to exploit.

The ecosystem for security bugs in Chrome is tightly managed and controlled because they understand the economics of attacks affecting system integrity

Conclusion

I believe the ecosystem for security bugs in Chrome is tightly managed and controlled because they understand the economics of attacks affecting system integrity. And that's why I recommend Chrome (either on a Chromebook, grsec hardened Linux, or an up-to-date Windows) to family and friends when they bring up a recent tale about "viruses".

The short answer to "why?" is usually "because Java" out of laziness, but there's obviously been a lot of strategic, well-thought out research and activities beyond the obvious things such as automatic updates, safe-browsing, smart UX, bug bounties, etc.

I previously mentioned that the Flash exploit path used by Vupen was preempted at a Pwn2own - while it was clearly successfully exploited and correctly proved a point at the time, this path was also known to be weak (or even seen as the weakest/most likely) and I'm sure hardening was on the roadmap even if it wasn't ready at the time.

To me that's a perfect example of understanding your threat model and being a realist - some things take time to improve, but with each small iteration of hardening it's important to have a tangible effect and, in the meantime, it's good to consider defensive countermeasures.

At the core of it, it's about raising the cost of attacks after building your threat model, then continuously maturing and tuning over time. For something as complicated as a browser, this requires a lot of consideration and skilled execution over years, and I'm thankful it's been in good hands. Not everyone has Google-like resources however, but with this mindset and attitude it's possible to make incremental improvements efficiently by first understanding the problem you're up against.

Matt Jones (@volvent) is Partner at elttam (elttam.com.au). He specializes in threat modeling, code review and data analysis. He has over a decade of industry experience in both offensive and defensive roles.

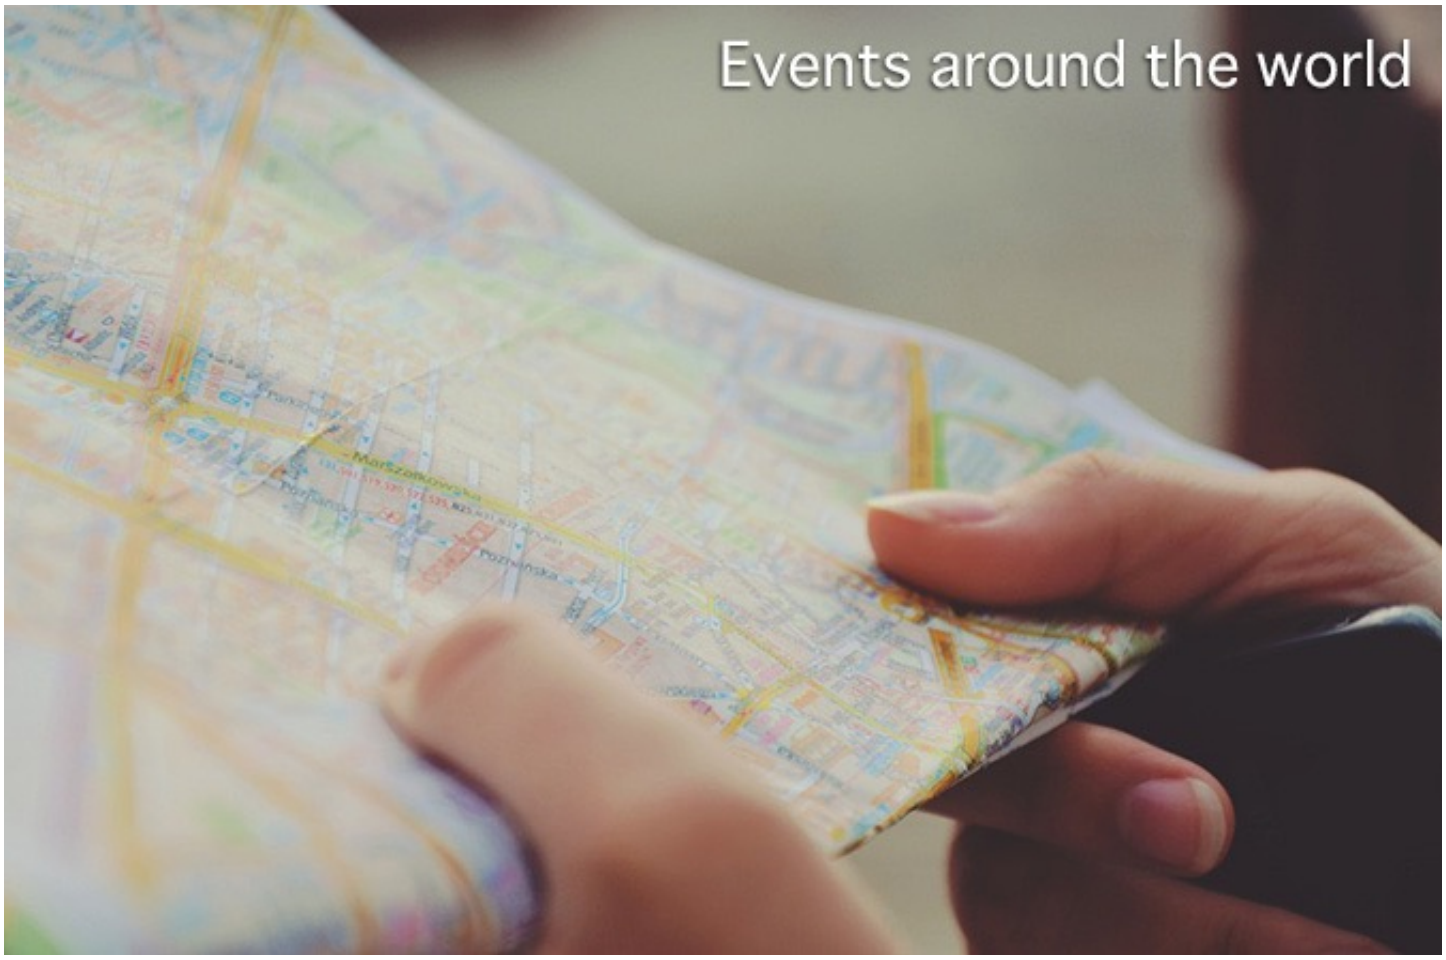

RSA Conference 2016

www.rsaconference.com - San Francisco, USA / 29 February - 4 March 2016.

Celebrating its 25th anniversary, RSA Conference continues to drive the information security agenda forward. Connect with industry leaders at RSA Conference 2016.

HITBSecConf2016 Amsterdam

conference.hitb.org - Amsterdam, The Netherlands / 23-27 May 2016.

HITB2016AMS features 2 and 3 days of technical trainings followed by a 2-day conference with a Capture the Flag competition, a technology exhibition and mini Haxpo hacker-spaces village for hackers, makers, builders and breakers.

Infosecurity Europe 2016

www.infosecurityeurope.com - London, UK / 7-9 June 2016.

Infosecurity Europe is Europe's number one information security event featuring over 315 exhibitors showcasing the most diverse range of products and services to 12,000 visitors.

Inside the largely unexplored world of mainframe security

by Zeljka Zorz

The security of mainframe computers - the so-called "big iron", which is mainly used by large organizations for critical applications, bulk data and transaction processing - is not a topic that has garnered much interest from the public. And, according to Phil "Soldier of Fortran" Young, the security community has not shown much interest so far, either.

"The two biggest misconceptions about mainframe computers is that they are unhackable, and that they are legacy and therefore don't deserve our attention or focus," says Young, who helps financial institutions protect their mainframes.

"The belief of them being 'unhackable' usually stems from a misunderstanding of how hacking works. They think someone has to find 0-days to exploit a mainframe, when in reality all they need to do is find a misconfigured webserver, a user account that has an easy-to-guess password, and so on."

Even though US-CERT rates traditional mainframes as one of the most secure computer systems due to a small number of vulnerabilities (when compared to the thousands affecting Windows, Linux, and other similar systems), Young believes that another reason for the belief is that fewer people know about mainframes and even fewer target them. If they are less likely to be attacked then, yes, technically, mainframes are "more secure" - but not "unhackable".

"As for the idea that these machines are 'legacy' and on the way out, that is totally false," Young notes. "They are modern operating system with their own nomenclature. They offer the same, and sometimes better, controls that other operating systems offer. Just because the operating system originated in the 70s (and was re-written in the 90s) doesn't make it legacy. Like I said in my BlackHat talk - parts of the NT likely still exist in Windows 10. But does that make it a 'legacy' operating system? No."

The fact is, mainframes are extremely useful computers, and are at the basis of almost every big and important service and business - retailers, banks, insurers, governments. Mainframes are backward compatible, and have high hardware and computational utilization rates and extensive input-output facilities. And they are highly reliable, which makes them a much better alternative to a cloud infrastructure.

"These systems are nowhere near leaving the enterprise," says Young.

"Sure, we hear from time to time that a company is planning on switching to an alternative. But usually after looking at the costs they change their mind."

Still, many organizations that use mainframes never test them - mostly because they are afraid that a penetration test could bring down one of their core systems.

"If a network security expert, with no knowledge about the mainframe, is able to bring it down with a simple Nmap scan, then that should be fixed, not ignored," he opined. "However, the likelyhood of that happening today is almost zero and is fueled mostly by old wives' tales from the late 90s when Nmap could bring the mainframe down under specific circumstances."

Young became interested in mainframe security in 2011. He scoured the Internet for tools, guides, anything to help him out do an audit of a mainframe and, when he found nothing, it became a problem that he set out to fix.

"When I say there was nothing online, I mean there was nothing - there was a link to a password cracker from 2000 and a post to the Nessus mailing list. That was it," he explained.

"When I started my blog and talks I figured there would be no interest. But year after year I get more interest as people start to do research in this space. So much so, in fact, that we had a little 'mainframe hacker meetup' at DEF CON this year."

Young and his colleague Chad "Big Endian Smalls" Rikansrud have been working on spreading the word about the issue of mainframe security.

They have been doing presentations about their work on security conferences, writing blogs, developing tools, and listing online resources in an effort to get the conversation and research started.

"I started out small with the tools. Mostly simple scripts as PoC. A perfect example of this is a shell script I wrote called Enumerate TSO, which would check user IDs of a mainframe (and works due to the way the TSO panel divulges information)," says Young. "It was slow but it worked. It has since been replaced with an Nmap script which does the same but is much faster:

```
Initiating NSE at 09:57
NSE: [tso-enum 10.10.0.200:23] Trying User ID: DB2
NSE: [tso-enum 10.10.0.200:23] Trying User ID: CICS
NSE: [tso-enum 10.10.0.200:23] Trying User ID: DADE
NSE: [tso-enum 10.10.0.200:23] Trying User ID: CHAD
NSE: [tso-enum 10.10.0.200:23] Trying User ID: NOPE
NSE: [tso-enum 10.10.0.200:23] Trying User ID: TRYME
NSE: [tso-enum 10.10.0.200:23] Trying User ID: ADCDMST
NSE: [tso-enum 10.10.0.200:23] Valid TSO User ID: ADCDMST
NSE: [tso-enum 10.10.0.200:23] Trying User ID: FAKED
NSE: [tso-enum 10.10.0.200:23] Trying User ID: DCDB
NSE: [tso-enum 10.10.0.200:23] Trying User ID: ADCDG
NSE: [tso-enum 10.10.0.200:23] Valid TSO User ID: ADCDG
NSE: [tso-enum 10.10.0.200:23] Trying User ID: TRIED
NSE: [tso-enum 10.10.0.200:23] Trying User ID: forced
NSE: [tso-enum 10.10.0.200:23] Trying User ID: adcdj
NSE: [tso-enum 10.10.0.200:23] Valid TSO User ID: ADCDJ
NSE: [tso-enum 10.10.0.200:23] Trying User ID: wrong
NSE: [tso-enum 10.10.0.200:23] Trying User ID: wronged
NSE: [tso-enum 10.10.0.200:23] Trying User ID: bl5ke
NSE: [tso-enum 10.10.0.200:23] Trying User ID: da5id
NSE: [tso-enum 10.10.0.200:23] Trying User ID: hahahah
Completed NSE at 09:57, 23.86s elapsed
Nmap scan report for 10.10.0.200
Host is up (0.0096s latency).
PORT   STATE SERVICE VERSION
23/tcp open  tn3270  IBM Telnet TN3270
| tso-enum:
|   TSO User ID:
|     TSO User:ADCDMST -  Valid User ID
|     TSO User:ADCDG -  Valid User ID
|     TSO User:ADCDJ -  Valid User ID
|_  Statistics: Performed 18 guesses in 22 seconds, average tps: 0
```

"Generally, I'm making tools to make a penetration testing easier. Either I wrote it because I needed it to do a pentest or because I wondered if it could be done or not. Other tools for Nmap include a Network Job Entry (NJE) named node brute forcer."

Another tool he wrote (and released at DEF CON) allows for submitting messages using NJE commands. "Once you have the node names (using the Nmap script) you can use iNJEctor to submit commands," he explained.

"I think the most interesting thing I've learned in all the time I've been exploring mainframes is, by far, how unexplored this area really is. It almost feels like I'm exploring a vast unexplored island with an old castle - who knows if I'm going to find anything, but the journey is fun and interesting," he shared.

"For example, I gave a talk at Skytalks about finding mainframes on the Internet. It started with a though 'I bet there's internet facing mainframes', and it eventually led to me having the only database that I know of of internet facing mainframes.

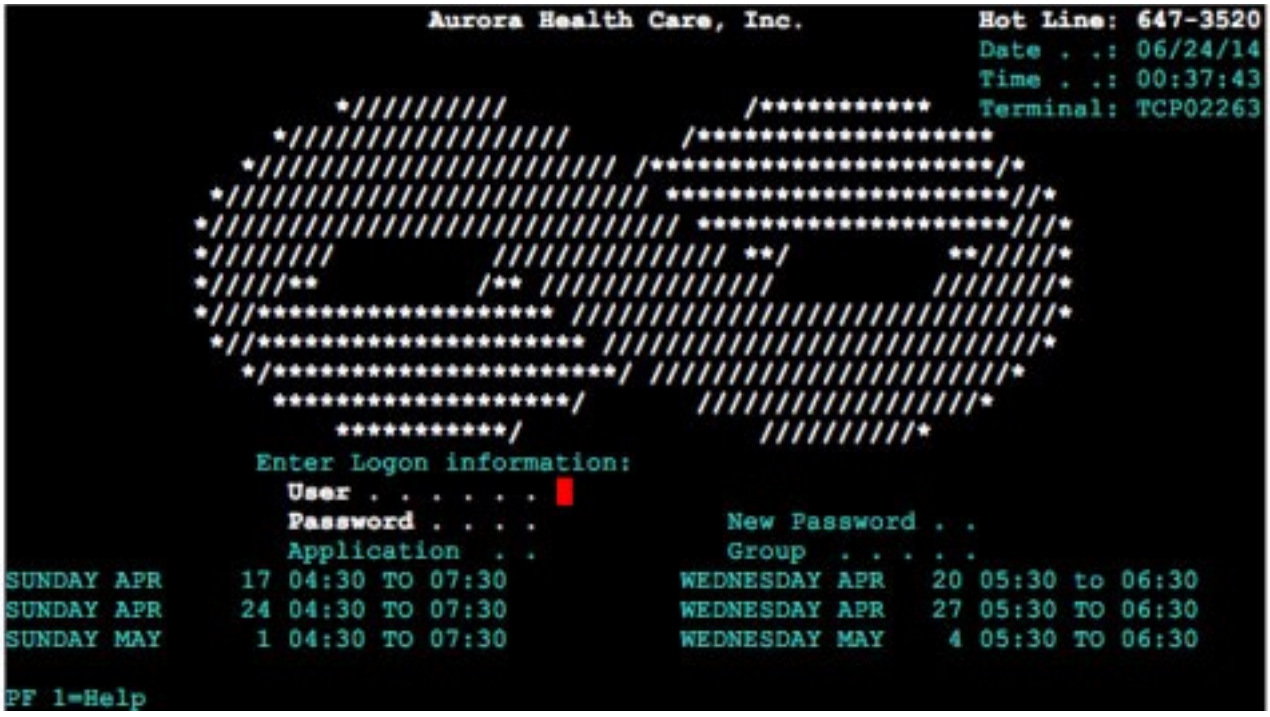

"On top of that I've learned more about the inner workings of various tools and scipts than I ever would have if I just used those tools," he noted. "At one time I wrote a Metasploit meterpreter in REXX (a mainframe scripting language). This required me to learn the inner workings of both Metasploit, meterpreter, REXX and z/OS. Another great example is writing scripts for the Nmap scripting engine - I am intimately more familiar with the inner workings of Nmap now than I ever was."

Another plus of his research is the community. "I know I've ragged on them in the past in my talks (and if you ever dare go to mainframe forums, turn back, do not go there) but when I had the opportunity to meet people in real life at conferences it was truly far more welcoming and open than I had ever expected."

In the end, I wanted to know what practical advice could he offer to organizations running mainframes to improve their security.

"Get a pentest done. Just do it. It won't be great because the skillset isn't there today but you have to start somewhere," he advised. "After that, use standards like the DoD DISA STIG for z/OS to lock it down and make sure your auditors are actually auditing your mainframe properly. Finally, keep your mainframe up to date. If you have a mainframe you need to sign up for IBMs z/OS security portal where they will give you patches to security vulnerabilities for the system."

Zeljka Zorz is the Managing Editor of (IN)SECURE Magazine and Help Net Security (www.net-security.org).

The Lord of the Hacktivist Rings

by Carl Herberger

Cyber attacks against websites have been happening for about a decade. But what puts one company at high risk of attack and into what we call 'The Ring of Fire?' The Cyber Attack Ring of Fire maps out vertical markets based on the likelihood that organizations in these sectors will experience attacks. It reflects five risk levels, with organizations closer to the red center more likely to experience DoS/DDoS and other forms of cyber attacks with greater frequency.

The Ashley Madison attack is a prime example of a company that would be listed in the high-risk section within the Ring of Fire. The company dealt with an attack so severe, it (reportedly) ultimately led to suicides after the release of the stolen information. Even so, the Ashley Madison attack was NOT different than your run-of-the-mill hacktivist cyber attacks (with a motive other than money).

Think of these attacks as digital versions of protests or ideological fights.

If recent cyber attacks have any uniqueness to them, it is that they have built upon the lessons of previous attacks and have become somewhat more calculating and laser-guided with their cyber-ordnance. Along these lines, heinous hacks are here to stay and there are strong, immediate calls-to-action for those operating in the Ring Of Fire to prepare and defend themselves. Those behind large-scale attacks include people who have been scorned, ideologues who now have a new form of communication, and protestors of all sorts. These individuals and groups have a new, modern avenue to pursue, and they are likely to use it to express their displeasure about other issues in the future.

Having said that, why are companies in these rings suddenly at increased risk of attack? Why do their profiles differ from other companies? Should there really be such a difference in risk?

Let's take a look at why these companies are affected by cyber attacks now more than ever.

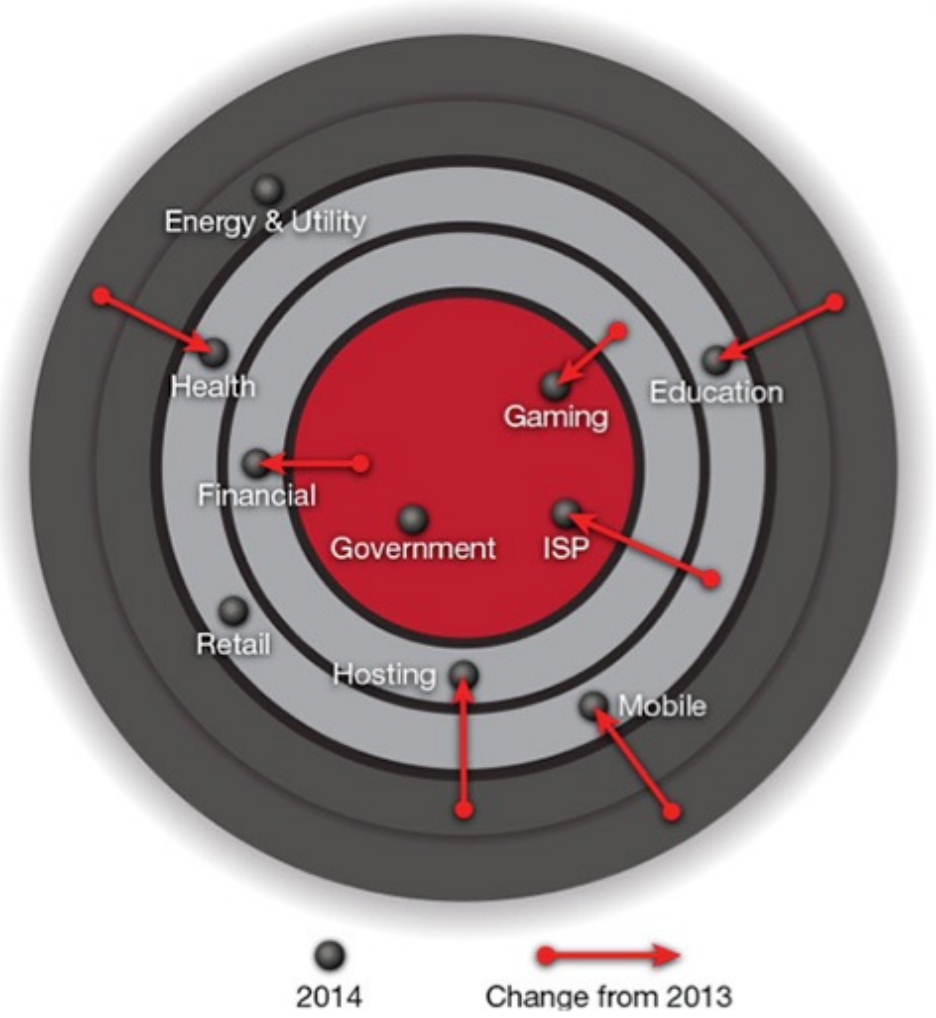

What attributes increase the chances that a company will be cyber attacked?

Availability. Does a second of outage mean something to your business? Online businesses that require high-availability are increasingly attractive targets.

The more companies conduct business online, the more disruptive a cyber attack, such as a DDoS, can be. This is particularly relevant to e-commerce sites, but certainly not limited to them. Similar to Ferguson protestors, hacktivists behind recent, major attacks were mainly gunning for attention, and attention they got. One such example is Lizard Squad's DDoS attack on XBOX and PlayStation during the week between Christmas and New Year's Day. Now, as in the Ashley Madison attack, the intent may to be shut down or shut up a business' message or operations altogether.

As more and more companies increasingly put their operations fully online, the Internet becomes an even more attractive place to conduct a protest (or a believed "anonymous" attack). Holding ill-written signs in front of business doors does not grab the attention or win the will of predisposed audiences as it once did. However, removing messaging and taking down important websites or businesses yields highly desirable results.

Aggressive or ideological business models. Does your business generally produce a percentage of dissatisfied or distraught customers? Do you run a business in an area that is morally objectionable to some? Are you affiliated with political movements or ideological pursuits? Do you compete intensely for customers or on slim margins? Is your business model incredibly disruptive to large populations of employed people? If so, you are at a much higher risk for cyber attack than more mundane businesses or ones with virtuous pursuits.

Even though this category is self-explanatory, it should be pointed out that nearly every major Western election, as well as many others worldwide, have experienced a cyber attack within the last three years.

Some were heinous and prevented proper tabulations and victor timing. To plan an election without cyber defenses is to be remiss these days. Moreover, it is now obvious from influential companies like Uber to Planned Parenthood and everything in-between, how easy it is for a company with a distinct message to invite a cyber attack.

What trends are increasing the chances that your company will be cyber attacked?

Imbalanced applications or business models

Generally it's great for an attacker to have an imbalance between a technical request and a technical reply.

In other words, if you search a website for all PDF files it contains, the request for this information is low, but the reply is potentially huge. This idea pervades the DNS service provider space and other techniques such as NTP and brute force attacks.

Use of cloud technologies

Denial of service attacks are not particularly complicated to pull off, technologically, but you often need a number of things to come together to make them work properly. Some examples of this include:

- The ability to anonymize yourself
- The ability to make an attack hard to mitigate
- Complicating the detection algorithm
- Complicating the effectiveness of mitigation techniques.

The use and expansion of cloud technologies dramatically complicates the protection against cyber attacks and makes it easier for hackers to go on the offensive and improve their chances of being effective.

Internet of Things

Traditional hackers use computers they've infected without the owners' consent. Future attacks will involve "things"- Internet connected and often ill-secured microcomputers that will be conscripted into a "bot army" and participate in magnifying volume-based attacks. In addition to being easy to recruit into a bot army, these "things" will have several nefarious benefits, including usurping most modern day HTTP security protocols.

How companies can protect themselves

Whether you believe it or not, see it or not, understand it or not, each of these trends has the ability to change the information security landscape immeasurably going forward. If all three trends come to pass, the resulting changes will serve to have today's CISO's job and responsibilities look like working for a horse and buggy manufacturer. The best among us will know when to "fold'em and when to hold'em" and it is my highest recommendation to begin the following:

- A process of decaying endpoint protection investments. Instead, move to a collection of collectors, detectors, command and control applications, and strong mitigation technologies.
- New "entryway" security investments. Consider new "fingerprinting" ideas that are agnostic to IP, and technologies that enrich your visibility.
- Become obsessive about application security.
- Availability will be challenged, as access will come from disparate devices and technologies (IoT).
- Prepare for large volume attacks. Cyber-attacks will conscript consumer (not just phones) and industrial devices in attacks against you.
- Software-Defined Networking (SDN) security. Are you ready? Start a personal and professional project as SDN is here – are attacks far behind?
- Pick your security vendors wisely. Those with no SDN strategy will leave you high-and-dry.

I wish you luck with your thoughts and pursuits and look forward to the exciting times of change and challenge before us.

Carl Herberger is the VP Security Solutions at Radware (www.radware.com).

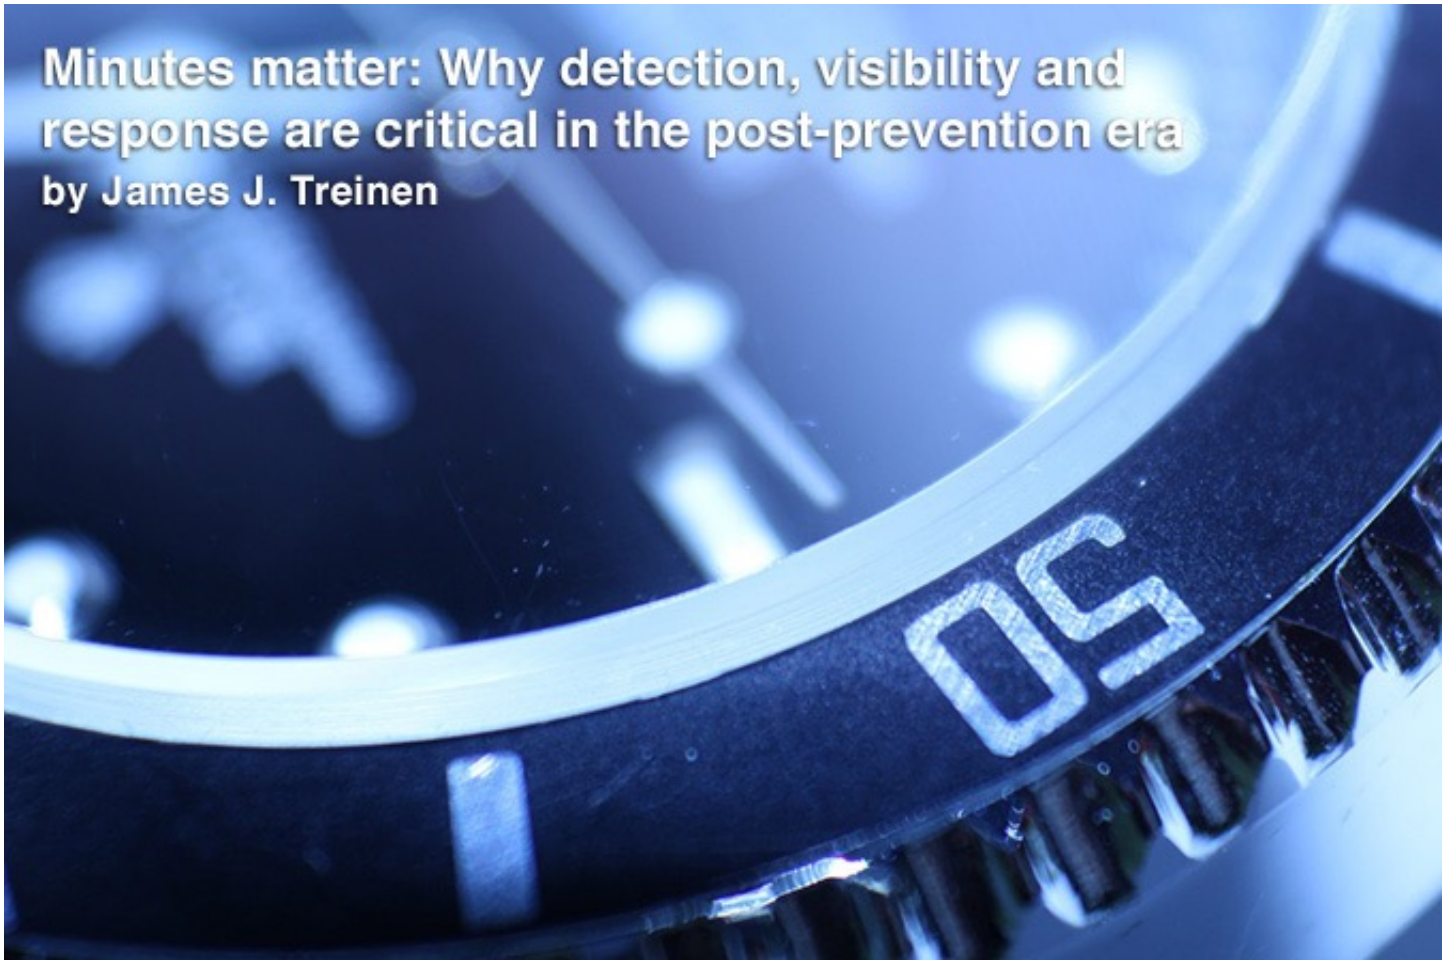

The network security landscape is flooded with threat detection solutions, yet the volume and complexity of today's attacks still appear to be getting the best of many organizations. Last year, reported US data breaches reached a record high of 783, with an estimated 43 percent of companies experiencing a breach.

This year, we've seen high profile data breaches hit nearly every sector including the healthcare, financial, federal, and even information security industries, proving that even the most sophisticated organizations in the world are not safe.

Password management company LastPass suffered a cyberattack in June, which resulted in the compromise of email addresses, password reminders and authentication hashes; healthcare giant Anthem revealed a data breach in February which exposed more than 80 million patient and employee records; and the data breach at the OPM affected 21.5 million federal workers and exposed the data of 4.2 million individuals. To make matters worse, research shows that attackers are present on victims' networks an average of more than six months before they are detected. Combine this with the increasing number of devices connecting to corporate networks driven in part by the Internet of Things – think of all those phones, tablets and smartwatches – and the job of a network security professional has become much more expansive in recent years. Cyber attackers identify new methods of attack faster than security professionals can preventatively defend against them. It is no longer a matter of if a company will experience a breach, but when, and to what extent.

In order to survive in today's constantly evolving security environment, organizations can no longer rely solely on preventative measures. Real-time visibility and detection are necessary, but organizations must also adopt faster and more effective incident response to enhance the resilience of their critical infrastructure.

Over the course of my career, I've learned how difficult it can be to detect threats within an organization's network. Much of my work as a research scientist has involved creating tools to help make a security analyst's life

easier. The goal of these tools is to help analysts answer the seemingly simple questions around the nature of a security incident, such as “Did any hosts communicate with a particular IP or Domain? If so, what was the nature of the connection? Did data theft occur, and if so, what data was lost?”, which become increasingly complicated as the size of the network and the number of alerts grow. Even if some of the answers are simple to find, addressing them over long periods of time is challenging, especially with existing solutions.

The trouble with a point-in-time approach

The trouble with current solutions, which take a point-in-time approach - where network analysis is performed by looking only at the present and in a very narrowly defined time window - is that they don’t have the ability to establish the specific nature of an incident. Network recording devices were introduced to help fill this gap, but the current generation of on-premise appliances offer windows of time that are much shorter than the six month statistic mentioned above.

The lack of long term, full fidelity network detail greatly hinders incident responders and leaves them guessing as to what actually happened. This leaves only logs and NetFlow, which are a good start, but leave many unanswered questions.

In addition to not being able to look back in time, point-in-time solutions don’t detect threats that may have occurred in the past and for which no known detection technique existed (e.g. a zero-day vulnerability). Many exploits are carried out without detection, and if full packet capture (PCAP) isn’t saved, then there is no way to go back and search for newly discovered exploits. More advanced attackers will use new variants of malware, or bespoke malware, which can enter an environment undetected, and will use it only a few times, often successfully. Once this malware has been discovered, it is typically too late for the early victims. By going back in time, organizations can re-examine the stored PCAP to find instances of these early exploits.

It’s interesting to note that the types of exploits found using these techniques tend to be much more targeted (e.g. data theft and credential theft), which makes sense. These variants are typically of crimeware families such as Dyre or Drydex, modified to evade detection. Once the new variants are discovered, signatures are written, and the obvious question becomes whether the new variant was used against an organization prior to the new signature’s arrival. Current point-in-time tech has no way of answering that very pressing question.

Combatting alarm fatigue

In addition to the problems created by this weakness, the job of a security analyst is made more difficult by the alarm fatigue created by low fidelity alarm streams. In a 2015 survey titled “The Cost of Malware Containment,” Ponemon Institute found that the average enterprise receives 16,937 malware alerts per week from their IT security products, of which only 19 percent are deemed reliable, and only 4 percent are investigated. The volume of alarms has resulted in security professionals becoming deaf to them.

At the beginning of my career, I worked at an MSSP where we would process more than 100 million alarms per day. The sources included network monitoring devices, host monitoring devices, proxy logs, firewall logs and more, and it was impossible to triage the volume of alarms effectively. SIEM was eventually introduced, with the idea to build rules that looked for predefined patterns in the alarm stream. This was effective at tamping down some of the alarm fatigue issues, but my team still encountered the same problems as with traditional signature-based intrusion detection systems - rules could only be fired for which predefined patterns were known.

Next was the era of anomaly detection - the idea that it’s possible to find new needles in the haystack. While many anomaly detection techniques enable security professionals to parse through larger volumes and variety of data to detect and prevent fraudulent activities, many tend to create false positives and wind up contributing to the alarm fatigue issue instead of helping solve it.

Neither of these techniques are silver bullets, which further emphasizes the need for analytical tooling that can alleviate the burden placed on security analysts.

Addressing security's talent shortage

The problem is further exacerbated by the industry's talent shortage. Even with the talent we have, many security professionals lack the skills required for deep network analysis. In fact, in a recent ISACA study titled "The State of Cybersecurity: Implications for 2015," less than half of those surveyed believe their current security teams have the ability to detect and respond to complex incidents.

The skills shortage could be attributed to the fact that it wasn't until a few years ago that universities started offering computer science degrees specializing in security, and even now, these programs are few and far between. Most security professionals have a background in computer science, but without the opportunity to study new threats and forms of attack while in school, many find it challenging to keep pace with today's security environment.

Not only that, but the industry still lacks the tools that can help security professionals scale themselves and their teams, ultimately making their jobs more manageable. Very few solutions on the market today provide deep network analytics that can look back in time, filling in the missing dots and increasing overall visibility into an organization's health. I firmly believe the security industry needs to help the pros work better, faster and more efficiently.

SECURITY OF THE PAST FOCUSED ON DETECTING AND BLOCKING ATTACKS, BUT IN TODAY'S POST-PREVENTION ERA, AN ORGANIZATION MUST BE ABLE TO SEE MORE, HUNT SMARTER AND RESPOND FASTER

Surviving in the post-prevention era

Security of the past focused on detecting and blocking attacks, but in today's post-prevention era, an organization must be able to see more, hunt smarter and respond faster. Alarm fatigue, talent shortages and point-in-time detection may plague security teams, but by incorporating better visibility and detection capabilities into existing solutions, security teams will spend less time inspecting security incidents and more time resolving them.

Ultimately, organizations need to accept the fact that they won't be able to prevent every security incident moving forward. If you still don't believe me, consider the adversaries - organized crime and nation states, to name a few. They are well-funded, well-trained and well-staffed. Even though the US recently signed a cybersecurity agreement with the Chinese government, it appears that it has already been breached. Highly trained operatives from nations and criminal syndicates are crafting new exploits every day, and they are likely already entrenched across your network.

The security perimeter is now either full of holes or doesn't exist at all. To this end, you can't prevent every attack, so you must be able to detect them. Tooling that is built to survive this reality will excel at distilling a wide number of data points down into highly focused, highly confident alerts and address the data breaches of the future in a much quicker and more intelligent way.

James J. Treinen, Ph.D. is the VP, Security Research at ProtectWise (www.protectwise.com).

Web application fingerprinting with Blind Elephant

by Wolfgang Kandek

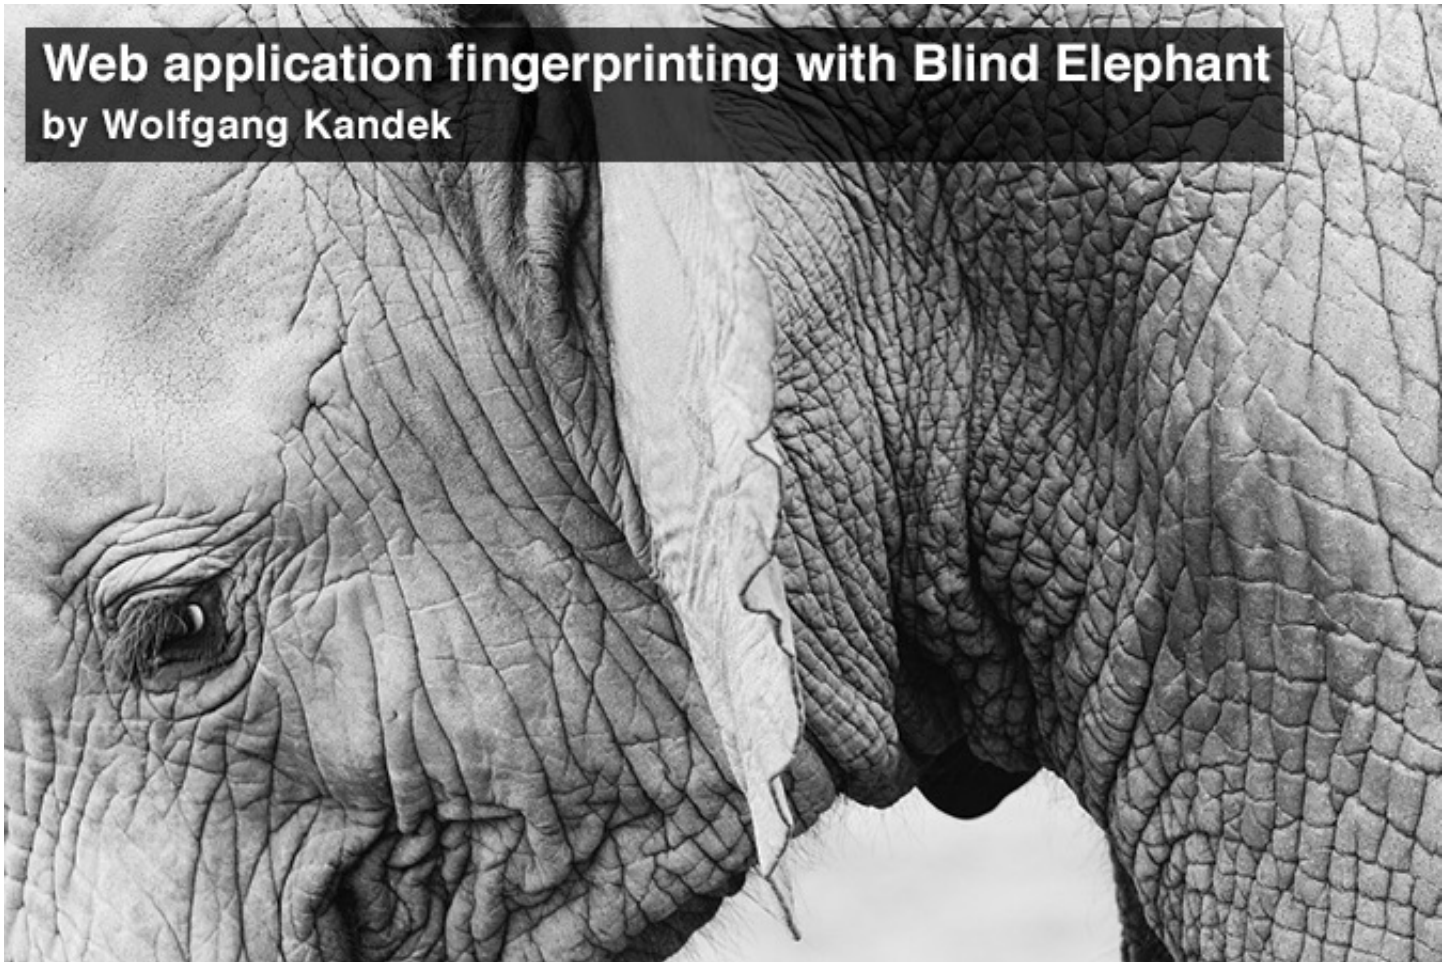

Web applications can present serious security risks to enterprise IT environments, and identifying the vulnerabilities they present can be a tricky business when traditional scanning methods fall short or disrupt the applications they are inspecting. The alternative is to use an indirect process called fingerprinting. The open source Blind Elephant project was developed five years ago to build a fingerprinting tool and to continually add support for different web apps and plugins.

Sometimes standard web application scanning techniques are too intrusive or require access that is not available or have unacceptably negative side effects. In these cases, it is necessary to use an indirect method like web application fingerprinting to determine the web application's version by inspecting static files it contains, and then report the known vulnerabilities for that version.

When you need web application fingerprinting

There are several classes of vulnerabilities for which standard web application scanning techniques don't work well and where web application fingerprinting makes sense:

- **Vulnerabilities with little or no information:** Sometimes Proof-of-Concept code or information pertaining to the vulnerability is unavailable, making it difficult or impossible to create active vulnerability detection. In the case of CVE-2015-7319 (WordPress Appointment Booking Calendar Plugin SQL Injection) and CVE-2015-7320 (Multiple Cross-Site Scripting Vulnerabilities), even though partial information is available about the vulnerabilities, it is not enough to build a reliable exploit.
- **Post-authentication vulnerabilities:** Persistent cross-site scripting vulnerabilities like Moodle Multiple Security Vulnerabilities need a user with certain rights in order to be successfully exploited. If authentication is not provided, a detection is not possible. CSRF vulnerabilities, ranked in the 8th position by OWASP, for example CVE-2015-6655 that affects Pligg CMS,

also require authentication, making detections difficult for the same reasons.

- **File upload vulnerabilities:** Vulnerabilities such as Vtiger CRM Remote Code Execution Vulnerabilities require the upload of arbitrary data on a customer's application. In this case, even though the PoC is available, authenticated access is needed to upload arbitrary files on the affected host and exploit malicious code, so it's easier and just as accurate to use web application fingerprinting.

- **RCE vulnerabilities:** Remote code execution vulnerabilities, like the one recently identified in concrete5, require execution of arbitrary code on a targeted system. While it's possible to build a detection from the available proof of concept, it would compromise the customer's instance. It's therefore better to use web application fingerprinting.

- **SQL injection:** For SQL injection vulnerabilities like CVE-2014-6293, Statistics (ke_stats) Extension in TYPO3, the number and names of tables may vary with the implementation of the application, so it's not possible to automate the table lookups required for a detection.

Blind Elephant is an open-source static-file web application fingerprinter. It attempts to discover the version of a (known) web application by comparing static files at known locations against pre-computed hashes for versions of those files in all available releases.

About Blind Elephant

Blind Elephant is a trustworthy open-source static-file web application fingerprinter. It attempts to discover the version of a (known) web application by comparing static files at known locations against pre-computed hashes for versions of those files in all available releases. This technique works well when the static files change with every release, allowing the fingerprinter to identify the application version based on the contents of the files. This technique is non-invasive and generic, and the use of pre-computed hashes means it is fast, low-bandwidth and highly automatable.

Current detections

Five years after integrating Blind Elephant (blindelephant.sourceforge.net) with Qualys Cloud Suite and adding detectable applications, the Qualys / Blind Elephant integration now detects over 200 web applications, plugins and extensions, and this number continues to grow every week. Organizations using the tool can look it up in their scan reports and see a listing of the web applications found in their environment.

In order to add a detection, the contributor needs access to the source code and a few versions of the web application. With too few versions or files that remain unchanged across versions, it is not possible to create detections.

Future detections

When support is added for different web-applications and their extensions/plugins, it gets posted online. Anyone who wants a detection added for a certain open source web application can post a request to Qualys' Blind Elephant community (community.qualys.com/community/blindelephant).

Wolfgang Kandek is the CTO at Qualys (www.qualys.com).