(IN)SECURE

Issue 41 - February 2014.

CLOUD INSECURITY? TIME TO BUST THE MYTH

LEVERAGING BIG DATA FOR SECURITY OPERATIONS

PRIVACY IN THE CLOUD: THE POWER OF ENCRYPTION

THE PAST, PRESENT & FUTURE OF BIG DATA SECURITY

TABLE OF CONTENTS

/// CLOUD SECURITY

/// BIG DATA

Big Data and cloud security have been some of the most discussed topics during the past year, and we decided it was time to dedicate an entire issue to them. We have a great lineup of industry leaders sharing their knowledge and exploring various areas.

As you're reading this, I'm about to start my yearly trip to San Francisco for the madness known as RSA Conference, without a doubt the most significant information security event of the year.

I'm looking forward to meeting many of you, discovering new companies and seeing innovative technologies in action. A special issue of (IN)SECURE Magazine after the event will showcase the best of what the show had to offer and put a spotlight on several companies. See you next week in San Francisco!

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.

Most organizations are unable to resolve a cyber attack

The lack of incident detection and investigation puts companies and their CISOs' jobs at significant risk, according to a new Ponemon Institute study.

In fact, when a CEO and Board of Directors ask a security team for a briefing immediately following an incident, 65% of respondents believe that the briefing would be purposefully modified, filtered or watered down.

Additionally, 78% of respondents believe most CISOs would make a “best effort guess” based on limited information, and they would also take action prematurely and report that the problem had been resolved without this actually being the case.

This disconnect results from several critical shortcomings in the current point solution approach to cybersecurity and incident response (IR), namely:

- Lack of timely compromise detection: 86% of respondents say detection of a cyber-attack takes too long

- Inability of point solutions to prioritize alerts as they come in: 85% say they suffer from a lack of prioritization of incidents

- Lack of integration between point solutions: 74% say poor or no integration between security products negatively affects response capabilities

- An overwhelming number of alerts paralyzing IR efforts: 61% say too many alerts from too many point solutions also hinders investigations.

“When a cyber attack happens, immediate reaction is needed in the minutes that follow, not hours or days,” said Dr. Larry Ponemon, chairman and founder of the Ponemon Institute. “It’s readily clear from the survey that IR processes need to incorporate powerful, intuitive technology that helps teams act quickly, effectively and with key evidence so their companies’ and clients’ time, resources and money are not lost in the immediate aftermath of the event.”

The sad state of cybersecurity readiness

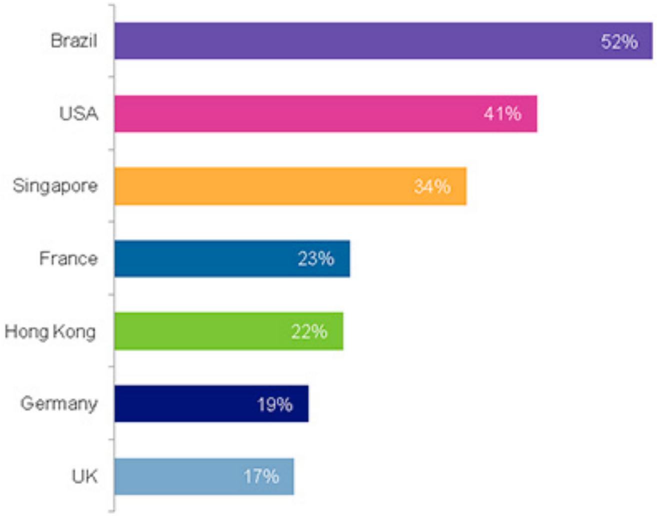

Just 17 percent of UK business leaders see cybersecurity as a major priority, compared to 41 percent in the US, research from BT has revealed.

The research, which assessed attitudes to cybersecurity and levels of preparedness among IT decision makers, highlights that UK businesses are lagging behind their US counterparts in crucial areas. Just one in five (21 percent) respondents in the UK are able to measure the ROI of their cybersecurity measures compared to nine in ten (90 percent) US companies. Similarly, 86 percent of US directors and senior decision makers are given IT security training, compared to just 37 percent in the UK.

More than half (58 percent) of IT decision-makers globally stated that their boards underestimate the importance of cybersecurity. This figure increases to 74 percent in the US but drops to 55 percent in the UK. The difference in levels of preparedness correlates with attitudes to threats. Non-malicious insider threats (e.g. accidental loss of data) are currently the most commonly cited security concern globally, being reported as a serious threat by 65 percent of IT decision makers. In the UK this falls to 60 percent and is followed by malicious insider threats (51 percent), hacktivism (37 percent) organized crime (32 percent), nation states (15 percent) and terrorism (12 percent).

In the US the proportion of IT decision makers who see non-malicious insider threats as a severe threat increases to 85 percent and is followed by malicious insider threats (79 percent), hacktivism (77 percent), organized crime (75 percent), terrorism (72 percent) and nation states (70 percent).

400Gbps NTP-based DDoS attack hits Cloudflare

Matthew Prince, CEO of content delivery network Cloudflare, has confirmed that one of its customers has been targeted with a very big Network Time Protocol (NTP) reflection attack - "bigger that the Spamhaus attack from last year."

He didn't name the customer, but he has shared that the attack reached the level of over 400 gigabits per second, that it probably caused congestion on some peering exchanges (mostly in Europe), that (based on sampled data) it misused just over 4,500 misconfigured NTP servers, and that the customer initially wanted to pay with a stolen credit card.

Despite the recommendation issued by US-CERT about updating public-facing NTP servers to a ntpd version that doesn't allow attackers to use them for NTP amplification attacks, there are still many vulnerable ones out there.

"The attack relies on the exploitation of the 'monlist' feature of NTP, as described in CVE-2013-5211, which is enabled by default on older NTP-capable devices. This command causes a list of the last 600 IP addresses which connected to the NTP server to be sent to the victim," explains US-CERT.

Intrinsic-ID enhances its Saturnus cloud security solution

Intrinsic-ID has released a new version of Saturnus (www.intrinsic-id.com/saturnus), its device-unique cloud security solution, that gives users total control over the protection of their data. The new version includes enhanced usability features that make using the cloud safer without compromising flexibility, performance or ease of use. As part of the launch of this new version, the company is offering a money back guarantee.

Saturnus is a hardware/software solution based on proprietary and patented Hardware Intrinsic Security (HIS) technology. The USB token contains a smart-card chip embedded with HIS, the strongest technology to generate and store security keys. On top of this, the Saturnus software application runs on the client device that will process the data. While working in the cloud, the Saturnus token is connected to the USB port of a device. The hardware-based security is augmented by a second step, a username and password-based login system. This combination is called two-factor authentication and is based on two unrelated factors: something known to the user (username and password) and something the user has (the hardware token). The result of two-factor authentication based on HIS is incredibly strong protection for all data in the cloud.

The USB token provides security based at the client site. When the USB token is connected, files are encrypted before they leave the device on the way to the cloud. This encryption is performed by using symmetric key cryptography, which means files are encrypted and decrypted using the same key. Since encryption and decryption are only performed on the client side and within the hardware of the token, the key never leaves the user and is therefore completely secure from malicious use.

The GUI enables an intuitive use of the application. All functions are visible on the application screen and files can be dragged into the secure Saturnus environment.

Infosecurity Europe to feature over 350 exhibitors

infosecurity
EUROPE

Infosecurity Europe (bit.ly/infosec-2014) is Europe's number one Information Security event. Featuring over 350 exhibitors, the most diverse range of new products and services, an education program and over 12,000 visitors from every segment of the industry, it's the most important date in the calendar for infosec professionals across Europe.

Take the chance to hear about new and existing products, services and solutions as exhibiting companies take to the stage to demonstrate the capabilities of their information security solutions. Pose your questions directly to the solution providers and find the answers you've been looking for.

The Business Strategy Theatre seminars feature case studies and best practices for addressing the challenges and issues facing management, CEOs and other board level directors. Benefit from the opportunity to:

- Discover how to tackle the key business challenges and issues impacting how an organization protects itself against the latest threats.
- Gain first-hand experience from vendors and end-users, sharing practice experience and real life learning.
- Access learning that can be applied directly to your business.

IE 0-day used in watering hole attack tied to previous campaigns

An Internet Explorer zero-day vulnerability (CVE-2014-0322) is actively exploited in the wild in a watering-hole attack targeting visitors to the official website of the U.S. Veterans of Foreign Wars, FireEye researchers warned.

"It's a brand new zero-day that targets IE 10 users visiting the compromised website – a classic drive-by download attack. Upon successful exploitation, this zero-day attack will download a XOR encoded payload from a remote server, decode and execute it," they explained.

"We believe the attack is a strategic Web compromise targeting American military personnel amid a paralyzing snowstorm at the U.S. Capitol in the days leading up to the Presidents Day holiday weekend. Based on infrastructure overlaps and tradecraft similarities, we believe the actors behind this campaign are associated with two previously identified campaigns (Operation DeputyDog and Operation Ephemeral Hydra)," they added.

This new campaign has been dubbed "Operation SnowMan," and the similarities with the aforementioned earlier campaigns are many: exploitation of an IE zero-day, delivery of remote access Trojan (Gh0st RAT), "watering hole" exploit delivery method, related C&C infrastructure, the use of a simple single-byte XOR encoded (0×95) payload obfuscated with a .jpg extension.

"The exploit targets IE 10 with Adobe Flash. It aborts exploitation if the user is browsing with a different version of IE or has installed Microsoft's Experience Mitigation Toolkit (EMET)," they shared, and pointed out that installing EMET or updating to IE 11 are perfect mitigation measures.

It is believed that the same actors have likely orchestrated all these campaigns. So far, the targets were US government agencies, defense companies, IT and law firms, NGOs, mining companies, so it's safe to say they were cyber espionage campaigns aimed at stealing confidential information.

Websense researchers say they have discovered the use of this same vulnerability as early as January 20, 2014 (FireEye detected the exploit on February 11), and that the targets were the visitors to a fake site mimicking that of the French aerospace association GIFAS, which includes contractors and firms in both the military and civilian aircraft industry.

Google offers five grants to women in security to attend HITB2014AMS

Google is offering five grants to women in security to attend the Hack In The Box Amsterdam conference in May.

These grants include a VIP ticket to the conference on the 29th and 30th of May, an exclusive invite to the HITBSecConf Speakers Reception on the 28th, an invite to the Girl Geek Dinner Amsterdam on the 29th and an invite to the HITB Post Conference Reception sponsored by Microsoft on the evening of the 30th.

Winners of the grant will also receive up to 1000 EUR towards travel costs (to be paid after the conference).

To find out more about the Google Women in Tech Travel and Conference Grant program, see here - www.google.com/edu/students/women-in-tech-conference-and-travel-grants/

As crimeware evolves, phishing attacks increase

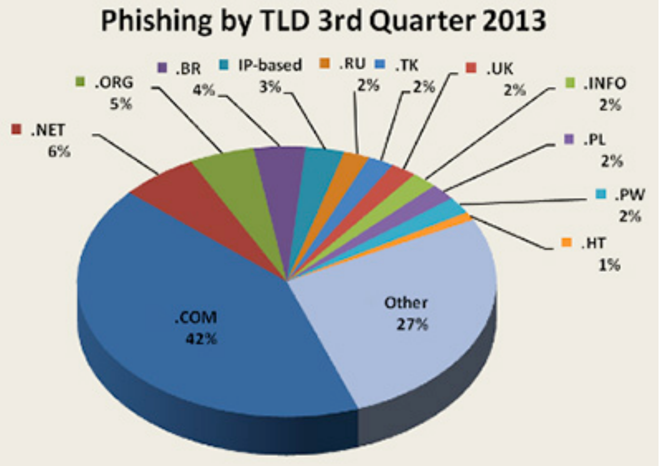

The number of phishing campaigns increased by more than 20 percent in the third quarter of 2013, with crimeware attacks evolving and proliferating, according to the APWG.

The total number of unique phishing websites observed rose to 143,353 in Q3, up from Q2's 119,101. The increase is generally attributable to rising numbers of attacks against money-transfer and retail/e-commerce websites.

During the same period, there was an 8 percent decline in the number of brands targeted by phishers, as the number of brands targeted fell from an all-time high of 441 in April 2013 to 379 in September 2013.

Attack vectors continued to evolve, placing social media at forefront of crimeware's vanguard in the quarter. "In the 3rd quarter of 2013, we also saw a change in the phishing themes used by malware authors. An emphasis on social media-themed subjects, such as 'Invitation to connect on LinkedIn', was used to entice users who would be used to seeing such subjects," said APWG contributor Carl Leonard of Websense Security Labs.

Encryption use continues to grow

Use of encryption continues to grow in response to consumer concerns, privacy compliance regulations and on-going cyber attacks and yet there are still major challenges in executing data encryption policy, according to a Ponemon Institute study.

Key findings:

- Steady increase in the deployment of encryption with 35% of organizations having an enterprise wide encryption strategy
- Most organizations deploy encryption to lessen the impact of data breaches
- The number one perceived threat to sensitive data is employee mistakes rather than external attack
- Two biggest challenges faced by organizations executing a data encryption policy are knowing where sensitive data resides and managing the actual technology
- Key management identified as a major issue by more than half of organizations
- Organizations with the highest security posture are now three times more likely to have a formal encryption strategy than those with the lowest security posture.

The results of the study show there has been a steady increase in the deployment of encryption solutions used by organizations over the past nine years, with 35% of organizations now having an encryption strategy applied consistently across the entire enterprise compared with 29% last year.

The survey also indicated that only 14% of organizations surveyed do not have any encryption strategy compared with 22% last year.

Mobility is the weakest security link

Surveying more than 750 security decision makers and practitioners, a CyberEdge Group report found that more than 60 percent had been breached in 2013 with a quarter of all participants citing a lack of employer investment in adequate defenses.

Key findings include:

Concern for mobile devices. Participants were asked to rate — on a scale of 1 to 5, with 5 being highest — their organization's ability to defend cyber threats across nine IT domains. Mobile devices (2.77) received the lowest marks, followed by laptops (2.92) and social media applications (2.93). Virtual servers (3.64) and physical servers (3.63) were deemed most secure.

The BYOD invasion. By 2016, 77 percent of responding organizations indicate they'll have BYOD policies in place. 31 percent have already implemented BYOD policies, 26 percent will follow within 12 months, and another 20 percent will follow within two years.

Inadequate security investments. Although 89 percent of respondents' IT security budgets are rising (48 percent) or holding steady (41 percent), one in four doubts whether their employer has invested adequately in cyber threat defenses.

Malware and phishing causing headaches. Of eight designated categories of cyber threats, malware and phishing/spear-phishing are top of mind and pose the greatest threat to responding organizations. Denial-of-service (DoS) attacks are of least concern.

Ignorance is bliss. Less than half (48 percent) of responding organizations conduct full-network active vulnerability scans more frequently than once per quarter, while 21 percent only conduct them annually.

Careless employees are to blame. When asked which factors inhibit IT security organizations from adequately defending cyber threats, "low security awareness among employees" was most commonly cited, just ahead of "lack of budget."

What do government security pros think?

Tripwire and the Government Technology Research Alliance (GTRA) announced the results of a U.S. government cybersecurity survey that evaluated the attitudes and responses of 111 security and compliance professionals from U.S. government agencies and contractors.

"Cybersecurity continues to be one of the top priorities of senior executives in the federal government," said Ron Ross, fellow at National Institute of Standards and Technology (NIST). "Studies, such as this one, bring together important data points that help decision makers assess trends and take part in an ongoing dialog that will help us craft effective solutions to our difficult and challenging cybersecurity problems."

Key findings include:

- 60 percent believe the new NIST framework will improve security.
- 55 percent believe government IT security has improved due to the administration's policies.
- 46 percent say they have seen reductions in risk due to continuous monitoring efforts.
- 43 percent of IT security and compliance employees consider poor governance and the dysfunctional Congress "the biggest security threat we face."
- 45 percent of respondents believe funding is the greatest challenge their agency faces in successfully implementing cybersecurity programs; only 37 percent believe they have adequate resources to properly implement policy; and when asked what federal security leaders should do to connect security to the agency mission, the second-most popular response was "more funding."

How Edward Snowden's actions impacted defense contractors

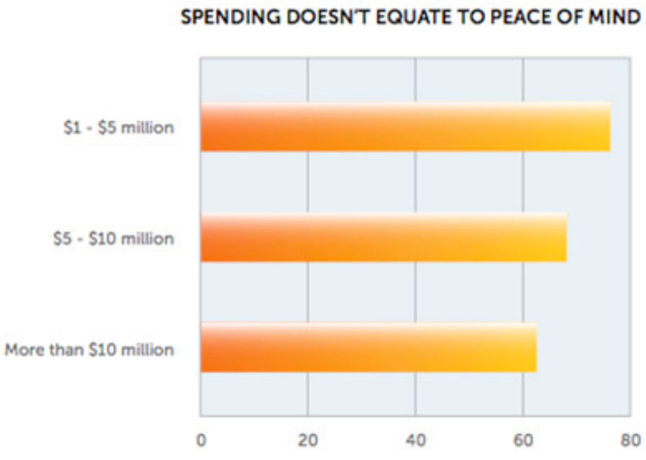

Despite big security budgets, defense contractors still express concern over their vulnerability to APTs.

A new ThreatTrack Security study sheds light on the attitudes of a very exclusive group of IT and security managers - those employed by U.S. defense contractors - at a time when national cybersecurity is under scrutiny.

75% of the respondents indicated that the Edward Snowden incident has changed their companies' cybersecurity practices in one of the following ways:

- 55% say their employees now receive more cybersecurity awareness training
- 52% have reviewed or re-evaluated employee data access privileges
- 47% are on higher alert for anomalous network activity by employees
- 41% have implemented stricter hiring practices
- 39% say their own IT administrative rights have been restricted.

63% of the survey respondents hold either secret, top secret or confidential clearances. However, of those who have access to or store confidential information, 27% do not hold such clearances. This represents a potential privileged access problem wherein contractor employees without such clearances may have easy access to sensitive government data.

Lack of skills hindering appsec programs

An ongoing shortage of skills in application security is severely hampering the implementation of effective Appsec programs, according to SANS. The 2014 Application Security Programs and Practices survey queried 488 IT and security professionals about the current and future state of application security in their organizations.

"One thing that stands out this year is the increase in number of organizations with a formal application security program in place. Approximately 83% of respondents (up from 66%) have an Appsec program in place, and more than 37% have a program that has been operating for more than five years," says SANS Analyst Frank Kim. "This indicates that a lot of progress is being made, but it also highlights that there is much more to do."

In the survey, more than 35% of respondents test the security of their business-critical applications on an ongoing basis, up from 23% in last year's survey. And, encouragingly, only a small percentage (fewer than 3%) of respondents left application security to chance and did not test at all. The survey found that a lack of qualified staff and lack of skills are seen as the major inhibitors to instituting Appsec programs.

"This year's survey provides valuable and surprising insights into the challenges that organizations face today in implementing a successful Appsec program," says SANS Analyst Jim Bird. "It's not only funding and getting management buy-in—there are other, more fundamental problems, including a shortage of skills, that are preventing people from taking care of security where it makes the most difference, upfront in design and development."

Windows, IE, Java are most vulnerable

When compared with the numbers from the previous year, 2013 has seen an increase in reported security vulnerabilities and, what's more, the number of critical vulnerabilities has also risen - although it's considerably smaller than in 2009.

GFI researchers have combed through the details provided by the US National Vulnerability Database (NVD), and have discovered that in 2013, an average of 13 new vulnerabilities were reported each day, bringing the total to 4794 - 447 more that in 2013.

50 percent of the flaws were found in products of only 10 vendors out of 760. The numbers are both a testament to the number of different offerings these big firms have and to their popularity, which naturally points to the conclusion that they are more often targeted by hackers and analysed by security researchers for security flaws.

Oracle has topped the list not only because of Java vulnerabilities, but also because of hardware flaws found in the company devices. Still, Microsoft can't sigh a sigh of relief, as the company has had a huge rise in "high severity" vulnerabilities when compared to 2012 numbers.

Critical vulnerabilities found in its various operating systems made Microsoft occupy 8 of the first 9 spots on the list of most targeted OSes in 2013. Finally, Microsoft's Internet Explorer, Oracle's Java and Google's Chrome have ended up occupying the first three spots (respectively) on the list of most targeted applications.

USA still the global spam king

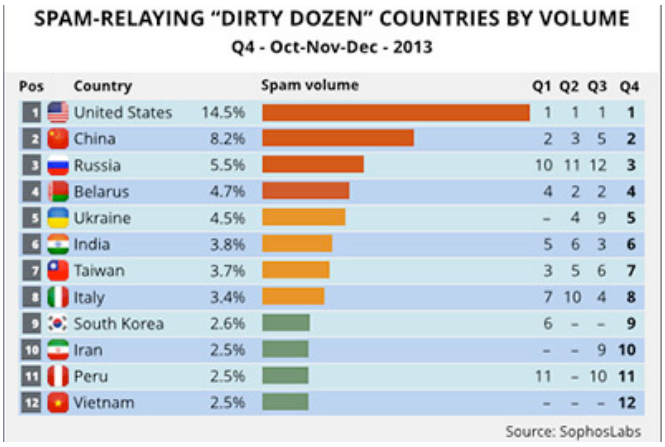

SPAM-RELAYING "DIRTY DOZEN" COUNTRIES BY VOLUME

Q4 - Oct-Nov-Dec - 2013

Pos	Country	Spam volume	Q1	Q2	Q3	Q4
1	United States	14.5%	1	1	1	1
2	China	8.2%	2	3	5	2
3	Russia	5.5%	10	11	12	3
4	Belarus	4.7%	4	2	2	4
5	Ukraine	4.5%	-	4	9	5
6	India	3.8%	5	6	3	6
7	Taiwan	3.7%	3	5	6	7
8	Italy	3.4%	7	10	4	8
9	South Korea	2.6%	6	-	-	9
10	Iran	2.5%	-	-	9	10
11	Peru	2.5%	11	-	10	11
12	Vietnam	2.5%	-	-	-	12

Source: SophosLabs

SophosLabs revealed the Dirty Dozen top spam-relaying nations, as it published the final “Spampionship” league table of 2013. Once again, it was the USA which earned the league’s top spot, generating 14.5 percent of the total spam volume sent during the last quarter of the year, giving it a clean sweep of top finishes for 2013. However, the gap to second place narrowed, with China re-emerging as a major player in the spam sending Dirty Dozen, leaping from 4.6 percent to 8.2 percent, while Russia’s spam contribution edged up from 3.0 percent in Q3 to 5.5 percent in Q4.

"The most obvious message from the Dirty Dozen charts is that the problem of zombified computers spewing spam is a truly global one,” says Sophos Senior Security Analyst, Paul Ducklin. “Every region of the world is strongly represented, with the exception of Africa.”

Spammers don't send spam themselves: they use botnets, or "zombie armies", of malware-infected computers to distribute their spam for them, almost always without the owners of the infected computers being aware.

Turning to the Spampionship table of spam-relaying countries by population, the numbers indicate the average “spamminess” per person compared to the USA. Results show things have stayed pretty stable, as Belarus retained its top spot, with the average computer there over 10 times more likely to send spam than if it were in the USA.

Cloud insecurity? Time to bust the myth

by Jeff Jones

Amidst the rapid growth of cloud computing, in multiple studies over the past several years security and privacy are commonly cited as top concerns. But a look at the actual experiences of cloud customers finds that often, those concerns are misplaced.

The existence of misconceptions about cloud computing is not necessarily surprising. The industry is still evolving and the range of cloud services continues to grow. Even the definition remains unclear to many people. Try asking several folks to explain “the cloud” and see what you get.

Among consumers in particular, the cloud is still a bit of a puzzle. Web hosting company Webfusion surveyed over 1,000 people in the UK in 2013 and found that only 34 percent of them claimed to understand what cloud computing means. Even smaller percentages correctly identified services like Dropbox, iTunes, and Gmail as cloud services. A companion survey conducted in the US turned up similar results.

It might seem understandable for consumers to have challenges describing a somewhat intangible, technical concept like cloud computing. But what about businesses?

Typically, business leaders seem aware of the benefits touted by cloud providers -- reduced capital costs, economies of scale, time savings, and flexibility. However, organizations that are considering cloud computing also voice a number of apprehensions.

To better understand the concerns that are acting as barriers to cloud adoption, and to see whether those barriers matched the experiences of actual cloud customers, in June of 2013 Microsoft commissioned an independent study by comScore (bit.ly/1lorQzR), which focused on SMBs.

Perceived Risks and Rewards of Cloud Adoption

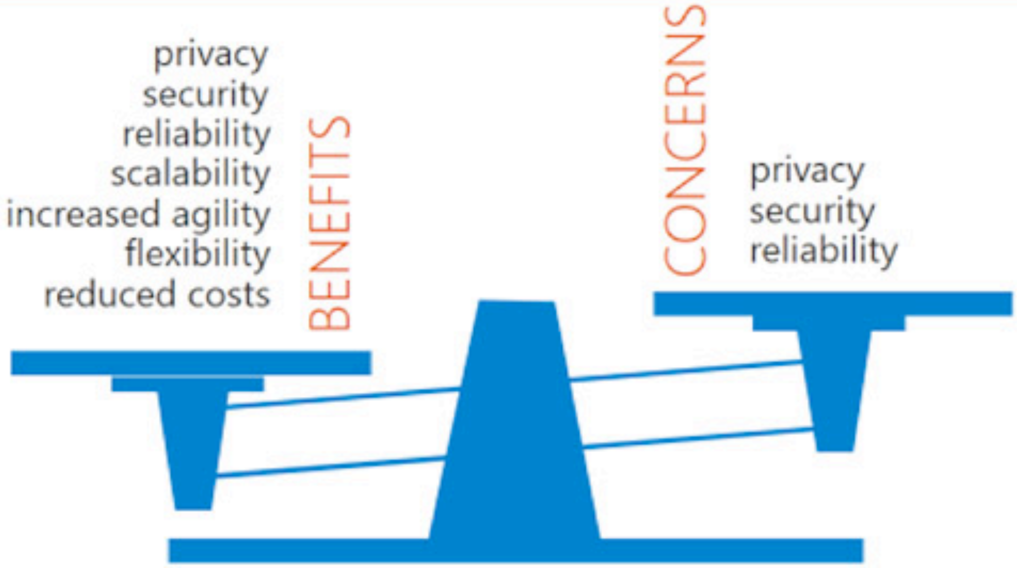

Respondents were not aware of the connection to Microsoft. Nor did comScore ask the respondents about specific offerings. The objective was to learn about their experiences with cloud services, regardless of the vendor. Looking first at SMBs that had not yet adopted the cloud, the study found:

• 60 percent cited concerns around data security as a barrier to adoption
• 45 percent said they were concerned that using the cloud would result in a lack of control over their data privacy
• 42 percent expressed concerns about the reliability of the cloud.

Ensuing surveys have produced similar results. In fact, the headline “Security concerns are still holding back cloud adoption” recently appeared on Help Net Security (bit.ly/1drdyfp).

Perhaps we shouldn’t be surprised. Plenty has been written about these “barriers” to adoption. However, these perceptions are actually strongly refuted by the realities (and benefits) reported by companies that use cloud services. From the comScore survey, among SMBs actually using the cloud:

• 94 percent said they had experienced security benefits in the cloud that they didn’t have with their former on-premises technology approach. Benefits included more consistent system updates, better spam email management and up-to-date antivirus protection
• 62 percent said that their levels of privacy protection increased as a result of moving to the cloud
• 75 percent said they experienced improved service availability since moving to the cloud.

Clearly, the benefits of the cloud outweigh the concerns. Improved reliability, security and privacy protections help give back both time and money. For example, the study showed that, as a result of moving to the cloud:

• 70 percent of SMBs were able to invest more in product development and innovation, demand creation and expansion into new markets
• 50 percent of SMBs have pursued new opportunities because of the time they saved managing security.

But even while cloud services are taking on much of their customers’ security work, it’s important for businesses to remember they still have some responsibilities. For example, cloud customers will still need to manage the security of their client devices – ensuring up-to-date antivirus software, and educating employees on the importance of using strong passwords.

The chart on the following page illustrates the mix of security responsibilities between customer and provider, depending on the service model deployed.

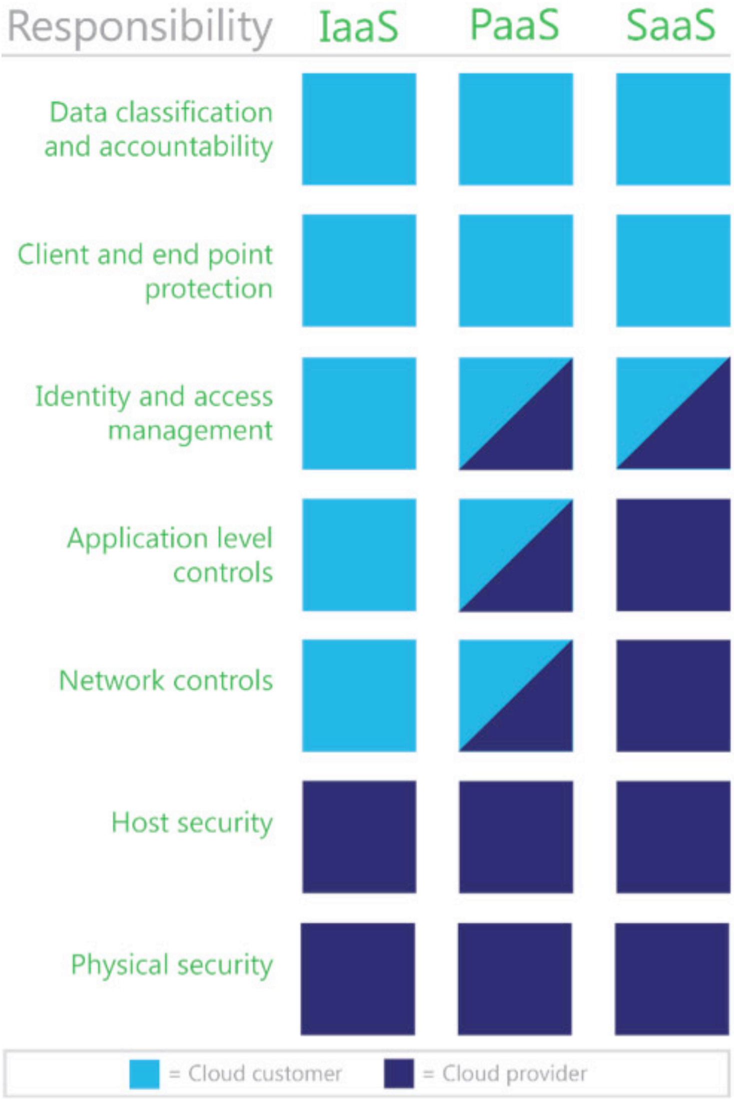

Responsibility	IaaS	PaaS	SaaS
Data classification and accountability	🟦	🟦	🟦
Client and end point protection	🟦	🟦	🟦
Identity and access management	🟦	🟦/⬛	🟦/⬛
Application level controls	🟦	🟦/⬛	⬛
Network controls	🟦	🟦/⬛	⬛
Host security	⬛	⬛	⬛
Physical security	⬛	⬛	⬛

🟦 = Cloud customer ⬛ = Cloud provider

Employees in particular play an important role in protecting an organization's data and other assets. Knowing how to spot phishing scams, and other types of social engineering is imperative. Employees should be trained to be alert for and to avoid bogus links in email and on suspicious web sites.

More and more people are also using their smartphones and other personal devices to access company data and systems remotely. To help organizations and their employees learn to defend against online fraud and other cybercrimes, Microsoft has published an "Internet Security at Work Toolkit" (bit.ly/1eNaXwL), with tips, fact sheets, videos and other information. IT Pros should consider downloading and sharing those resources across their organization. Specific guidance (bit.ly/1ehbdj3) on recognizing and avoiding scams that come through email or web sites is also available.

For businesses that haven't yet adopted cloud services, a good way to begin is by assessing the organization's current level of preparedness with the Cloud Security Readiness Tool (bit.ly/M4O7WV), released by Microsoft's Trustworthy Computing Group in 2012.

Knowing that the vast majority of customers have experienced security improvements after moving to the cloud should help ease concerns among potential adopters. It's time we busted the myth of the insecure cloud, and remove that perceived barrier, once and for all.

Jeff Jones is a 25-year security industry professional that has spent the last several years at Microsoft helping drive security progress as part of the Trustworthy Computing initiative. In the role of Director, Jeff draws upon his security experience to work with enterprise CSOs and Microsoft's internal security teams to drive practical and measurable security improvements into Microsoft process and products.

Executive hot seat: Cloud Security Alliance CEO

Interview by Mirko Zorz

For many years, Jim Reavis has worked in the information security industry as an entrepreneur, writer, speaker, technologist and business strategist. Jim is helping shape the future of information security and related technology industries as co-founder, CEO and driving force of the Cloud Security Alliance.

How has the cloud security landscape changed in the last five years? How mature is the cloud today?

Cloud computing has matured dramatically over the past five years, and the use cases are quite broad. Five years ago, many IaaS offerings were very simple, with perhaps five or so options.

Today many of those same providers have literally hundreds of offerings, management tools and product add-ons, and third party providers have added many new security solutions. Many of these offerings today are purpose-built for securing cloud providers, where five years ago they were often legacy security products “tweaked” for the cloud.

Some of the most interesting innovations are identity as a service, cloud aware encryption, cloud application control and log management.

Cloud adoption is at an all-time high, yet those that are not using it are still saying the biggest obstacle is security. What can service providers do to earn customer trust? What should customers be on the lookout for?

Certainly the Snowden issue, which I will discuss later, is a factor. However, the majority of customers I talk to say the main issue is compliance and governance. It is about showing proof of security, rather than technical security itself, which is a nuance to the issue that many surveys miss. Solving this is mostly an educational issue with the various players, I think.

Policy makers and auditors need to understand how the cloud really operates. Security professionals need to understand that the risks are not black and white - you may actually reduce risks by using a public cloud provider that has better firewalls and backup systems. The big thing providers need to do to increase trust is to be transparent with their security and governance practices. We think the CSA Security, Trust and Assurance Registry Program (cloudsecurityalliance.org/star) provides the global model for trust in the cloud. It is a control framework that is mapped against key security standards and requirements, it allows providers to publish their security practices for all to see, includes 3rd party certification where needed and will in the future provide continuous monitoring.

Last year you launched the Software Defined Perimeter Initiative, a project to develop an architecture for creating highly secure and trusted end-to-end networks between any IP addressable entities, allowing for systems that are highly resilient to network attacks. Are you satisfied with the response? What enterprises are working with you on the development?

We are quite pleased with the response so far, and a few very large enterprises have implemented prototypes and pilots with positive results. We have much more work to do in order to publish mature specifications and simplify the adoption. SDP is a big idea that says we are going to need to change our view of how we implement IP networks to decrease the global discovery and visibility of computers. The IP-controlled thermostat in my home is my business alone.

What trends do you expect to see in the next 12 months? Do you expect to see a notable increase in cloud security automation?

We see the growth in consumer-owned mobile devices and new generations of the Internet of Things creating a “force multiplier” that will lead to even more aggressive cloud adoption.

As organizations lose control of the endpoint device, they will have fewer options to prevent cloud usage, although solutions are coming to market to help them manage this usage. We do see a lot more automation in the entire lifecycle of acquisition, implementation and management of cloud services. Whether it is called cloud brokering or by another name, we see a lot of this automation coming from intermediaries.

AS ORGANIZATIONS LOSE CONTROL OF THE ENDPOINT DEVICE, THEY WILL HAVE FEWER OPTIONS TO PREVENT CLOUD USAGE

After Edward Snowden's revelations, there's been a growing trend of organizations moving or considering moving their data to cloud providers outside the US. What's your take on this situation?

We continue to monitor this highly dynamic situation. According to the survey we conducted in the immediate aftermath of the revelations, 10% of customers outside of the United States had stopped a cloud project using a US-based service, while another 56% said the news would negatively impact their future adoption of US-based cloud services.

I would say that in considering all of the data we have analyzed over the past several months, this trend is not growing and has somewhat stabilized for the following reasons:

- Some of the revelations have uncovered surveillance programs that have been conducted without the knowledge of the provider, such as via tapping fiber outside of a data center. This has created an opinion that there is perhaps less complicity than was originally imagined and that in some cases the provider could even be a considered a victim of surveillance.

- Generally speaking, customers seem to feel that the larger US cloud providers are taking positive actions to provide assurances that their information is being safeguarded from unwarranted government access. Many have announced greater efforts at end to end encryption. Some are even taking legal action against the US federal government. To be clear, the customers are telling us there is work yet to be done, and are forcefully pushing for greater transparency from the providers in their relationships with governments and in their own management of customer information.

• Still other revelations have shown that countries besides the US are conducting similar options to acquire large consumer data sets. Some of these countries collaborate with the NSA. Very recently, EU Justice Commissioner Viviane Redin was highly critical of European hypocrisy related to the Snowden revelations, given their own spying activities.

The mature conclusion that is being arrived at is that an enterprise's information must be safeguarded against a whole host of threats: domestic, international, private sector, public sector, criminals, hacktivists, etc.

The general consensus is that organizations need to take a holistic approach to increase the baseline of their security, which includes governance issues related to provider selection; technical architecture improvements, such as increased encryption and logfile monitoring and several other practices. More and more, the CSA best practices are being cited and used as the foundation of this increased security baseline.

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

Want to reach a large audience of
security pros by writing for (IN)SECURE?
STOP
Send your idea to mzorz@net-security.org

Saturnus
Saturnus

Saturnus
secure your digital life

Security uncertainty in the cloud: Problems and potential solutions

by Sergey Ignatchenko and Dmytro Ivanchykhin

This article is an attempt to perform an analysis of cloud security by taking into consideration two aspects. The first one is related to the different cloud security models (IaaS, PaaS, and SaaS), and we will attempt to assess some of the security risks for each. The second is related to the uncertainty of such assessments and the ways to deal with these uncertainties, in particular based on Service Level Agreements (SLAs), insurance, and certifications.

The basic premise for our analysis is that you, as a representative (CSO/CIO/…) of a company, are considering migrating to the cloud, but are concerned about the security implications of such a move.

It should be noted that our analysis is not intended for life-critical applications, or any industry where a security breach is considered completely unacceptable - our analysis is for a typical business application in a typical business enterprise, where there is a need to find the right balance between solution costs and solution risks.

Table 1 on the following page illustrates, more or less, a typical split of security responsibilities between a cloud service provider (CSP) and the customer for different cloud service models.

IaaS and security

In the “Infrastructure as a Service” (IaaS) world, a CSP usually provides you with a number of virtual servers and other virtual appliances that are essentially managed by yourself. One can reasonably argue that with IaaS, most of the security is still handled by the customer. If an organization is switching a part of its infrastructure to IaaS, it still needs to configure firewalls, IPSs, IDSs, etc. In this set up, the customer is usually also responsible for OS patching and application security. In theory, the CSP may provide help with setting the system up, but in practice that help is

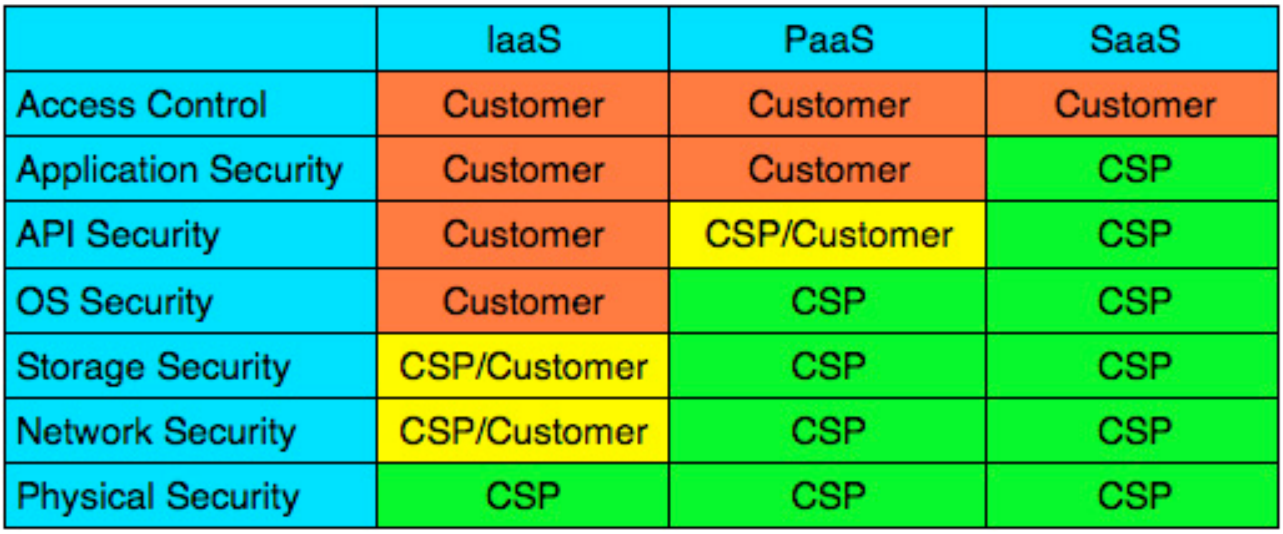

	IaaS	PaaS	SaaS
Access Control	Customer	Customer	Customer
Application Security	Customer	Customer	CSP
API Security	Customer	CSP/Customer	CSP
OS Security	Customer	CSP	CSP
Storage Security	CSP/Customer	CSP	CSP
Network Security	CSP/Customer	CSP	CSP
Physical Security	CSP	CSP	CSP

Table 1.

rather limited, so the balance of responsibilities usually looks like this: the CSP provides the (usually virtualized) servers, the customer does the rest.

This approach can be either a blessing or a curse, depending on your point of view. If you're looking to delegate to the cloud as much as possible, it is not exactly good news: you still need to run your own IT/security department (in a cloud/IaaS setup, IT/security department can be more easily outsourced to the cloud, but this is beyond the scope of this article). On the other hand, if you're more concerned about migration to the cloud being smooth, out of all cloud migration scenarios IaaS requires the least possible change, with the logical part of your infrastructure kept more or less intact - or at least not that affected as it would be when migrating to the other cloud models.

In a sense, IaaS is the closest you can get to managing your own infrastructure. From many points of view - including the security one - it is very close to how traditional hosting ISPs function, with the main differences being the following:

1. You get “on-demand” scaling and the associated reduction in costs, which is usually the main reason why enterprises are switching to the cloud in the first place.

2. In exchange for reduced costs, you need to deal with a set of cloud-specific security risks which include at least two rather broad categories:

- Category 1: Attacks coming from some other customer of the same CSP. Anybody, including a hacker from some distant country, can (attempt to) crack the security that separates virtual machines - yours and the attacker's - in the cloud.
- Category 2: Attacks via cloud administrator.

How big are these risks? In practice, we feel that the risk for attacks cracking the barrier between two virtual machines is not that big.

We've taken a look at CVE vulnerabilities for the last 3 months, and have found that only about 20 of 1300 vulnerabilities (about 1.5 %) are virtualization-related. It is also interesting to note that out of these 20 vulnerabilities, at least half are directly related to UI, so they're sitting somewhere in between the two risk categories mentioned above. Of course, the number of known vulnerabilities shouldn't be confused with real risks, since a single vulnerability can be enough to defeat security completely, but it still provides some information about the attack surface. Let’s say that this risk is not too high, and can be assessed with some level of certainty.

Unfortunately, the risk of being hit with an attack via cloud administrator is much more difficult to assess. Assuming that cloud administrators are running desktops/laptops (and they usually are), we need to recall that any system is only as strong as its weakest link. So, to break into your cloud (and to access your sensitive data) it is not necessary for the attacker to break into a neighboring VM via a weakness in the hypervisor.

Instead, it is sufficient to break into the desktop/laptop that the cloud administrator uses to manage the cloud, including your VM. Normally, administrators have the option to make an image of your VM while it is running. Such an image is rarely encrypted, and is even more rarely encrypted with a key that is not available to the same cloud administrator.

Also, such an image can be made without you noticing it, but even worse that that: even if you're using encryption, such an image will contain all your keys in unencrypted form. All this combined together, it means that whoever controls the desktop/laptop of the cloud administrator (the one who can create an image of your VM) should be considered as a person who can access your data.

To make things worse, it is not only the cloud administrator who has such powers, but also any hacker that is sitting half-a-world away but has managed to convince the cloud administrator to open a malicious e-mail attachment several months ago. And as the March 2011 RSA breach (tinyurl.com/o6s3o3k) has demonstrated, such spear-phishing attacks are extremely difficult to prevent even in the most security-conscious environments.

Compare it to the situation with more traditional ISP hosting: while the ISP administrator can physically take your HDD and compromise your data, taking out your HDD cannot be done remotely, and this makes a world of difference. With traditional ISP hosting, a remote attacker would need to bribe the ISP admin to obtain access to your sensitive data. This is theoretically possible, but not that likely. With a CSP, all the remote attacker needs is to spear-phish the laptop of the CSP administrator, and that is a much more realistic scenario.

The likelihood of such spear-phishing attacks being successful depends heavily on the security practices of a specific CSP, and may be at any point of the spectrum between “security as solid as rock” and “has holes large enough to let USS Enterprise through.”

In practice, it means that while risks in the first category can be assessed with at least some level of confidence and without a deep analysis of the security practices of a specific cloud provider, risks related to attacks via cloud administrator depend on the practices of a specific CSP, and this leads to uncertainty in security assessments.

PaaS and security

In the “Platform as a Service” (PaaS) world, you normally get an application platform that your developers use for developing new applications – the migration of existing applications without rewriting them is usually not an option. PaaS usually provides APIs to access platform services, such as network and storage. Unfortunately, for PaaS-based systems, security is even more difficult to assess than for IaaS-based ones. When using PaaS, you inherit all the security risks from IaaS, and face additional ones specific to PaaS.

With PaaS, separation between customers is handled by the PaaS provider. Database (or any other storage) separation is normally a responsibility of the CSP, too.

It leads us to the following categories of additional PaaS-specific risks:

- Category 3: Risks related to holes in PaaS API-based separation - unless the PaaS provider uses VMs to separate customers. APIs (especially higher-level APIs normally used in PaaS) are notoriously buggy and insecure, so there is considerable uncertainty when it comes to security assessments of these risks.

- Category 4: Risks related to database access separation. Data access separation is tricky – for each way to do it right there are at least several dozen ways to do it wrong. While we're shouldn’t automatically assume that all PaaS providers are doing it wrong, we still cannot rule it out; once again, it means that risk assessment is not possible without a detailed security analysis of the specific CSP.

SaaS and security

If you're dealing with the SaaS cloud model, things can be a bit simpler for you. Usually, SaaS CSPs are providing a very specialized environment for a very specific task, and there is little pressure to go down the “slippery slope” of moving more and more data into the cloud without understanding risks associated with such a move. If the data that you have moved into the cloud is not sensitive, you're perfectly fine.

If, on the other hand, that data is sensitive, from a security perspective most of the risk categories listed above still apply.

Summary of risk categories

Now let's summarize, in one table, all the CSP-related risks mentioned above. While the list of the risks represented by those four categories is by no means exhaustive, it is sufficient to illustrate the magnitude of security uncertainties faced by enterprises planning a migration to the cloud.

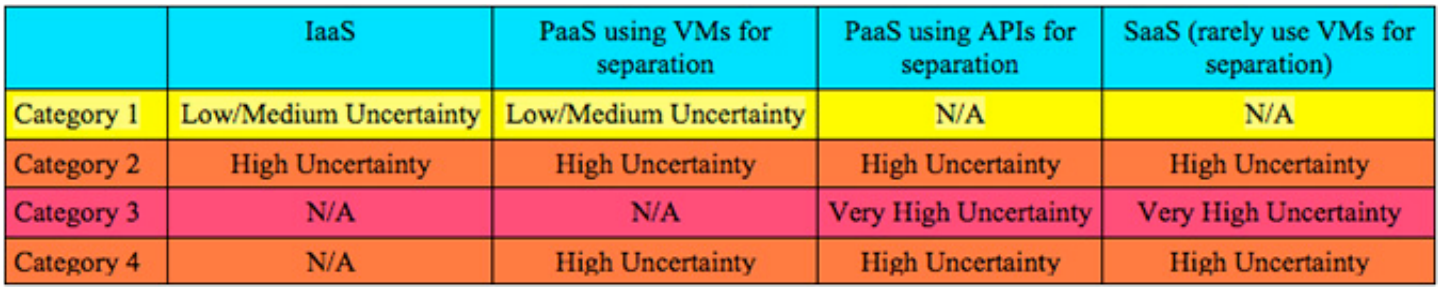

	IaaS	PaaS using VMs for separation	PaaS using APIs for separation	SaaS (rarely use VMs for separation)
Category 1	Low/Medium Uncertainty	Low/Medium Uncertainty	N/A	N/A
Category 2	High Uncertainty	High Uncertainty	High Uncertainty	High Uncertainty
Category 3	N/A	N/A	Very High Uncertainty	Very High Uncertainty
Category 4	N/A	High Uncertainty	High Uncertainty	High Uncertainty

Table 2.

While uncertainty estimates in Table 2 are subjective, we've discussed them with a few people from the security community and have found that there is not much disagreement on them. It should also be noted that for the following discussion we won't rely on specific uncertainties, but only on the assumption that for all cloud service models, there is substantial security uncertainty related to the specific security practices of a specific CSP.

When we take a look at Table 2, we can see that even for IaaS, which has the lowest overall CSP-related uncertainty compared to the other cloud service models, security uncertainty is high. In other words, without going deep into the details of a specific CSP, we cannot be sure how good their security is.

One (especially someone representing a CSP) may say: “Hey, you're speaking about security breaches affecting a CSP, but the same types of breaches can happen in your own environment. So, what's the difference?” Granted, most of the risk categories mentioned above also apply to your own environment, but there is one fundamental difference: in your own environment you can find out how risky it is (it will be very expensive and labor-intensive, but still possible). In a CSP environment it is much more difficult, up to the point of sometimes being next to impossible. In a nutshell, this “how difficult is to find out how secure your system is” is exactly what security uncertainty is about.

The importance of security certainty

The importance of being certain about the security of your system cannot be overestimated. Security is a very special entity compared to other technological issues businesses need to deal with. If the system does not do what is expected, or is not performing well, it is usually highly visible so the problem can be addressed right away. When it comes to security, you never know that the system is broken until it is too late. And given the scale of potential consequences arising from a security breach - which in worst case can easily cost hundreds of millions or even billions - taking uninformed security risks is the last thing a CTO/CIO/CSO should do.

Dealing with uncertainty: Audits, SLAs, insurance, and certifications

In business, there are four well-known ways to deal with similar uncertainties:

1. Perform a security audit of the CSP yourself. Fortunately or not, in most cases this won't fly (unless, of course, you're that big that you have your own cloud).
2. Rely on the guarantees that a CSP provides in its Service Level Agreement (SLA).

3. Rely on insurance (provided by a third-party insurance company)
4. Rely on a third-party security audit (ideally by a reputable authority).

Items 1, 3 and 4 on the list are rather obvious, but item 2 one may require a bit of explanation.

Service Level Agreements (SLAs) are a well-known way for service providers to provide certain guarantees for difficult-to-estimate aspects of their service. In a sense, it is the proverbial “putting your money where your mouth is”. If an ISP says “I guarantee that your network will be up 99.99% of time, or it is money back for this month”, it increases the customer's confidence that the ISP is able to achieve this target. Also, it creates a very specific incentive for the ISP to make sure that the network is indeed up and running at least 99.99% of time. These two sides of a SLA combined explain the high popularity of SLAs and SLA guarantees in the world of ISPs, mitigating to a great extent the inherent uncertainty of ISP service quality.

Security guarantees in CSP SLAs

You still want to migrate to the cloud, and there is no reliable third-party security audit. What kind of security guarantees you need from the CSP's SLA?

Let's assume that we are considering whether a migration to cloud makes business sense, and let's assume that the decision is purely money-based (which implies, among other things, that the potential of a single breach is not prohibitively expensive for the company).

A bit of math

Let's define the following:

S: Amount that could be saved due to transfer to the cloud during a certain time period **T** (**S** can be expressed, for example, in dollars/year)
L: Amount of loss if a security breach happens (per security breach; **L** can be expressed, for example, in dollars)
λ: Rate of security breaches - a frequency of security breaches expressed in terms of a number of security breaches per time period **T**. Somewhat similar to the failure rate from engineering. Under our conditions, $1/\lambda$ can be seen as “average time between security breaches” (somewhat similar to MTBF)
ξ: Losses due to security breaches happened during time period **T; $\xi = L * \lambda$**

The mathematical expectation of security breach losses **ξ** can be calculated as:

$$E[\xi] = E[L * \lambda] = L * \lambda$$

For the benefit arising from moving to the cloud to be positive, it is required that **$S > E[\xi]$**, that is, **$S > E[\xi] = L * \lambda$** or, respectively, **L** should be less than **S / λ**.

If compensation **C** is provided (for example, in SLA, or by a cyber-insurance company) in case of a security breach, then losses including compensation **$L1$**, if security breach happens, become **$L1 = L – C$** and **$\xi1$**, security breach losses including compensation, will have a mathematical expectation of
$E[\xi1] = E[(L-C)*\lambda] = (L-C) * \lambda$

Respectively, for the benefit arising from moving to the cloud to be positive, it is required that **$S > E[\xi1] = (L-C) * \lambda$** or **$L-C < S / \lambda$ [**]**

The value of **λ** depends heavily on the security procedures adopted by the CSP. To account for the worst-case scenario, one may assume that value of **λ** may happen to be arbitrary high; then, taking a limit as **λ** approaches infinity, **[**]** effectively becomes **$L – C < 0$, [***]** which, in simple words, means that the compensation per breach (from the CSP provider or from the insurance company) should be greater than the losses from the breach.

Back to business - SLAs

The formula above **[***]** can be summarized as follows: if considering the worst-case scenario in presence of an arbitrarily high security uncertainty, the move to the cloud will be cost-efficient only if the compensation in case of security breach (guaranteed by the SLA or by an insurance company) will be more than cost of security breach to the company. Or in even simpler form: with the conservative assumptions above, *In the absence of reliable third-party certifications, putting data in the cloud only makes financial sense if the cost of the data transferred into the cloud is less than the compensation provided by the SLA or by cyber-insurance in case of a security breach.*

It is interesting to note that the clause in the SLA you should be looking for doesn't depend on the type of the CSP (being IaaS, PaaS or SaaS); it is more a question of if the CSP will be willing to provide a compensation guarantee good enough for your migration.

In practice, we wouldn't expect CSPs providing SLA guarantees themselves in case of a security breach on the order of millions; the reason being that a CSP breach is likely to affect many (if not all) of its customers, and paying more than few months worth of their fees would be way too risky for the CSP. Still, having clauses in the SLA that guarantee you a few months worth of your CSP fees can allow you to move the least sensitive part of your expensive infrastructure to the cloud.

On the other hand, there is a possibility that the CSPs may obtain cyber-insurance covering security breaches, from independent insurance companies; this may allow them to substantially increase the guarantees in the SLAs.

Back to business: Cyber-insurance

As discussed above, CSPs themselves are not likely to provide compensation on the order of millions in case of a security breach because of risks being too high for a CSP. However, there are companies out there whose whole business revolves around taking that risk for others: insurance companies (in our case – cyber-insurance companies).

In theory, we could expect two different insurance models for our case. In the first model, a cyber-insurance company could insure customers against security breaches directly. In the second one, a cyber-insurance company may insure a CSP against security breaches, so that the CSP is able to put higher guarantees into their SLA. As it is usual with group insurance, we could expect that the second model would be more cost-efficient for the customer.

In practice, there are indeed companies that provide cyber-insurance against security breaches. One such company is CloudInsure, which has been reported to provide insurance with compensations up to $1 million (and insurance costs in the range from $5K to $10K per year). It has also been reported that CloudInsure is ready to insure CSPs (as described in second insurance model above), though as of now it is not clear if any CSP is including their (or anybody else's) insurance against security breaches in their SLAs. In any case, even with current pricing, cloud insurance may easily allow migration of the data worth on the order of $1M to the cloud (though you need to account for insurance cost when calculating cost efficiency of such a move).

The maze of certifications: Compliance and risks as two sides of the same coin

When you ask any sizable CSP about their security certifications, they will show you a very long list: HIPAA, SAS70, SSAE16 SOC1/SOC2, FIPS 140, ISO 27001, PCI DSS, and so on. The sheer number of items in this alphabet soup served by the CSP can be really impressive. Now, to understand the real value that these certifications have for you as a CSP customer, we need to define what your goal with certification is. There are two possible reasons why you may want a certification from your CSP – compliance and risk assessment.

It should be understood that, strictly speaking, being secure does not imply compliance with standards, and vice versa. While usually measures aimed at reaching compliance do help security (and therefore reduce risks), and often measures aimed to improve security help to reach compliance, the final states of “being secure” and “being compliant” are not 100% the same, and separate analysis of compliance and security may be necessary to ensure that the system is both compliant and secure.

Certification and compliance

The first reason for you to make sure that your cloud provider does have certification is that you need to be compliant with a certain regulation. For example, if you need to store medical data and you're located in the US, chances are that you need to be compliant with HIPAA data protection requirements, with no way around it. In a similar way, if you want to process credit cards, you likely need to be compliant with PCI DSS. It means that if you're moving respective portion of your data and/or processing to the cloud, you will very likely need your CSP to be compliant with an appropriate regulation and have appropriate certification. We should note that this aspect of certification is not exactly the focus of the article, so we won't concentrate on it too much. One word of caution though – even if your CSP is compliant with a certain regulation, it doesn't mean that by using their services you will be automatically compliant, too.

Just one example – your CSP can be an IaaS cloud provider, and can be PCI DSS compliant in a sense that they handle physical security according to PCI DSS, and that their own systems are secure according to PCI DSS requirements; however, as they're IaaS, they can't make sure that your web server sits in the DMZ, or that your firewall (provided by IaaS) is configured according to PCI DSS; it means that even if a CSP is perfectly PCI DSS compliant, you as a merchant can be easily non-compliant with all the potentially unpleasant consequences. Or if we express it in mathematical terms – as a rule of thumb, CSP compliance is usually a necessary condition for you to be compliant, but not a sufficient one.

Certification as a way to assess security risks

The second reason to choose a CSP with a certification is to reduce security uncertainty and therefore to allow you to move more sensitive applications and/or data to the cloud.

Here, however, some understanding of the real meaning of various certifications is necessary. There are several problems with understanding and comparing the certifications presented by most cloud providers. The very first word of caution is that you should never judge a CSP for risk accessing purposes based on the number of the certifications they have. Certifications should be taken one by one, and the merits for risk assessment should be evaluated for each certification.

Security-unrelated standards and certifications

There are many standards and certifications out there which are popular, but that have little to do with security. One such example is SAS 70 (in fact, the misuse of SAS 70 as a security certification was/is so frequent that Gartner has even issued a press release titled “Gartner Says SAS 70 Is Not Proof of Security, Continuity or Privacy Compliance”).

Paper-only compliance and under-specification

A big problem with many certifications is that a big part of them can be complied with without making any changes to the system, but merely by producing paperwork. While paperwork can have its own merits (for example, very few people will argue about the usefulness of security policy), certifications that assure very little beyond pure paperwork don't really provide much in terms of risk assessment. In other words, if the only thing a CSP needs to do to comply is to produce a pile of papers, the certification is not worth much for practical purposes (it may still be needed for compliance though).

A close cousin of paper-only compliance is under-specification. It happens when the standard to be complied with does provide some value beyond paperwork, but the set of security controls within is not sufficient to provide comprehensive coverage. For example, if a certain standard specifies that passwords need to be changed on regular basis, it is a good thing, but if the standard doesn't require using firewalls to protect certain parts of the CSP infrastructure, it may allow fully compliant organizations to not use firewalls at all. Taking into account the fact that a system is only as secure as its weakest link, this is clearly not enough to provide meaningful improvement in assessed risks.

Lack of cloud specifics

The cloud environment has its own unique challenges. For example, while Category 2 risks described earlier are not exactly unique to the cloud, in practice there is a substantial increase in Category 2 risks in cloud environments.

In particular, consolidation of control on administrative PCs used to control hundreds of CSP clients makes these PCs much more attractive attack targets, which in turn makes attacks on them much more likely and much more damaging. We strongly feel that to provide a meaningful capping of risks, such issues need to be addressed in the standards and certifications.

Navigating the alphabet soup of existing certifications

Let's take a look at existing security-related certifications from the point of view of risk assessment in general, and caveats listed above in particular.

SAS 70

SAS 70 (also known as "Statement on Auditing Standards No. 70 report") is one example of a statement that is widely (and wrongly) declared as a security certification. French Caldwell, research vice president at Gartner, has said about SAS 70: "SAS 70 is basically an expensive auditing process to support compliance with financial reporting rules like the Sarbanes-Oxley Act (SOX). [...] Chief information security officers (CISOs), compliance and risk managers, vendor managers, procurement professionals, and others involved in the purchase or sale of IT services and software need to recognize that SAS 70 is not a security, continuity or privacy compliance standard."

Another Gartner research vice president, Jay Heiser, has added that "Given that SAS 70 cannot be considered as proof that an offered IT service is secure, it should be a matter of suspicion when a vendor insists that it is".
In 2011, SAS 70 has been retired in favor of SSAE16.

SSAE 16 (SOC1/SOC2)

SSAE16 ("Statement on Standards for Attestation Engagements") is a standard that effectively replaced SAS 70. SSAE 16 audit reports come in three different flavors – Service Organization Control 1 (SOC 1), SOC 2, and SOC 3. Out of these, SOC 1 is a more or less direct replacement of SAS 70, and has been criticized by the security community because its controls are (like those in SAS 70) self-defined, and therefore merely having a SOC 1 report does not provide enough information about security properties of the system. SOC 2 reports, however, were received more favorably, but with a reservation that even SOC 2 is still underspecified with respect to security details.

It is worth noting that both SAS 70 and SSAE 16 are standards prepared by AICPA, which stands for "American Institute of Certified Public Accountants" and, as we understand, the reports are meant to be prepared by CPAs (Certified Public Accountants). We have no doubts that it has severely limited technical details in these statements (one cannot reasonably expect a CPA to know security at the level of a (ISC)2 CISSP).

HIPAA

While HIPAA has never been intended to certify the security of a cloud provider, it is still used by many CSPs as a part of alphabet soup of certifications they have. When looking at HIPAA regulations, it's easy to notice that there are two different sets of rules – Privacy Rule and Security Rule. HIPAA Privacy Rule has very little to do with security at all; HIPAA Security Rule has certain rules regarding security, but they're very few and too far between to provide substantial coverage for a cloud provider environment. It doesn't mean that you won't need a HIPAA-compliant CSP: if you need to store PHI data in the US, you certainly need a HIPAA-compliant provider. Still, having HIPAA compliance helps risk assessment of the CSP in a rather limited way.

ISO 27001

ISO 27001 (or more precisely, ISO/IEC 27001:2005 – or newer version 27001:2013) is a standard that aims to provide a framework for information security management systems (ISMS). In the industry, ISO 27001 has a longstanding reputation of being vague and underspecified (its cousin 27002 is better defined, but as far as we know, there is no ISO 27002 certification). With the release of 27001:2013 there was an improvement in this regard (now normative Annex A contains specific controls from 27002:2013 to be implemented), however, it still looks pretty much underspecified and lacking cloud specifics for the purposes of the risk assessment in the cloud. While covering pretty well such areas as organizational security, human resources security, and physical security, the important fields of network security and encryption are grossly underspecified (just as one example, we weren't able to find in ISO 27001 or ISO 27002 any reference to a minimal strength of encryption algorithms which are allowed to be used).

Based on this, our subjective opinion is that the value of ISO 27001/27002 for practical purposes such as risk assessment is rather limited, unless it is complemented by other means (such as the Cloud Control Matrix).

FIPS 140

FIPS 140 (Federal Information Protection Standard publication 140) is an interesting example of a standard that is fully specified, though, taken alone, it still cannot be used as a reliable metric of CSP security.

Most of the standards mentioned before were listing too few technical details, while staying at a level which is “too high” to provide necessary coverage. In contrast, FIPS publications tend to be very down-to-earth, and have all the necessary technical details.

The problem with the application of FIPS standards (including, but not limited to, FIPS 140) to CSPs is that the scope of FIPS is limited to cryptography. As it is well-known in the security industry, having good cryptography is only one of many prerequisites to having a secure system, so using only FIPS 140-compliant algorithms is clearly not enough to achieve meaningful system-wide security.

Moreover, in many cases even combining FIPS 140 with higher-level standards like ISO 27001 may be insufficient to guarantee security: the problem here is that between the high-level view of ISO 27001 and the low-level view of FIPS 140 there is still a substantial uncovered gap related in particular to security protocols used within the system.

As a big portion of security attacks comes exactly from breaches in the protocols (see, for example, recently developed BREACH attack on TLS), there is still a need for a more comprehensive security standard covering the whole stack of technology used, from high level to low level, including everything in between.

PCI DSS

PCI DSS is, in our opinion, a very nicely written standard, that doesn't suffer from the under-specification that most of previously discussed ones do. While it is reasonably detailed, it’s still possible to comply with it in practice.

The main problem we see with it for the purpose of CSP risk assessment is that PCI DSS has a very narrow scope. Everything that is not related to credit card numbers is out of PCI DSS scope, so if you're absolutely insecure but do not process credit cards at all, technically you're still compliant with PCI DSS.

On the other hand, if PCI DSS certification for a CSP can be interpreted as “it is safe to process credit card anywhere on this CSP” (and our point of view is that this is the only reasonable interpretation of PCI DSS claims for the CSP, but it is better to ask your specific CSP about their interpretation) – it will say quite a lot about security processes of a specific CSP. While still lacking cloud-specific details, it is probably one of the most comprehensive existing security standards.

Our subjective opinion is that when we see that certain CSP is PCI DSS certified (assuming that interpretation of being certified above stands), it makes us a bit more assured about this CSP’s security practices; still, due to the limited scope of PCI DSS, for the purposes of a risk assessment of a CSP we'd clearly prefer to rely on some certification with much wider scope.

New kids on the block: CCM and STAR

Probably understanding the limitations of the existing certifications, several years ago the Cloud Security Alliance (CSA) has started the development of a new industry guidance framework, the Cloud Control Matrix (CCM). In September 2013, CCM 3.0 has been released.

The CCM specifically aims to provide security guidelines for CSPs to follow, and to provide a framework for assessment of security controls. It covers many of the gaps of the previous standards; one of the most important advantages of the CCM is that it is cloud-specific, and as such, it addresses cloud-specific issues.

Just as one example: based on the cloud service model and the cloud being public or private, CCM suggests which controls might fall under CSP responsibility versus the consuming organization; this delineation of responsibility in the cloud is one major item desperately missing from previous standards. While we feel that there is still a long way ahead for the CCM (in particular, covering cloud-specific attack vectors, such as those described in the described four categories), we strongly feel that the CCM represents a desperately needed move in the right direction (moreover, as far as we know, CSA's efforts are the only efforts in this area).

The STAR (CSA Security, Trust and Assurance Registry) is an initiative parallel to the CCM, and is partially based on it. Up until recently, only STAR layer 1 (self-assessment) was available; as with any self-assessment, its value for risk assessment was rather limited. However, in September 2013 it has been announced that STAR certification layer 2, based on independent third-party assessments and audits, is available from BSI (British Standards Institute). The STAR certification is based on ISO 27001, but is aided with certain controls from the CCM (currently from CCM 1.4, though in the future support for later versions of CCM is planned), which allows it both to deal with under-specification and to provide cloud-specific controls.

As of now, STAR layer 2 certification is probably the best single thing a CSP can show you to demonstrate that they really care about security. At this time, STAR layer 2 certification is still so very new that no CSP has it yet; however, we could expect first STAR layer 2 certifications by mid-2014.

Conclusion

We have performed an analysis of security uncertainties inherent to cloud service providers (CSPs) and some possible ways to overcome these uncertainties. We have found that while compensations for security breaches (in SLAs and/or by cyber-insurance) can allow enterprises to push some of the less sensitive data to the cloud, further reduction of security uncertainty is still necessary. Such an uncertainty reduction can be achieved via cloud-specific efforts like the CCM and the STAR.

We feel that such efforts are absolutely necessary to deal with security uncertainty, which severely limits information being moved into the cloud.

This article has been prepared based on cloud security research project conducted by authors at OLogN Technologies AG, Liechtenstein.

Sergey Ignatchenko is a Security Researcher with security experience in various fields going back to 1996. Among other things, he has designed a comprehensive and regulations-compliant security solution for one of G8 stock exchanges, and a security architecture for a very large system with 100+ millions of users, 10+ thousands of RSA SecurID tokens, and capable of processing billions transactions per day. His systems have passed numerous security audits with flying colours. He can be reached at si@o-log-n.com

Dmytro Ivanchykhin is an Information Security Consultant with interests spanning areas from mathematical aspects of cryptography (with an emphasis on abstract algebra and Galois fields) to applied cryptography and secure protocols. He can be reached at di@o-log-n.com

Share with the world: Who reads my data in the cloud?

by Ken Ivanov

The advent of cloud services can be referred to as one of the most serious historical changes in the data storage and access model. Ease of deployment, scalability and economic efficiency spurred a mass migration of businesses from the standard CAPEX model to the cloud. In the rush of implementing these new business practices, the crucial point of observing the privacy of and controlling access to in-house data has somewhat fallen by the wayside.

Ironically, recent spy scandals revealed through the efforts of Edward Snowden and the Wikileaks project made it clear that the question of privacy and data security has never been as important as it is today.

Who is spying on me?

Revelations in information security lead us to realize that the world of information is much more dangerous than it appeared before. The NSA is developing wide-scale global surveillance projects (tinyurl.com/okdes6f), and is already capable of silently capturing Internet communication. At the same time, intelligence bureaus are working on establishing backstage contacts with first-tier online service providers such as Microsoft, Google and Amazon. Despite the providers' efforts to refute the collaboration claims by stating that no mass uncontrolled disclosures of private data had taken place, the public continues to rail against the connection and to call for the companies to respect their customers' right to privacy.

On the other hand, it is important to remember that while international cloud legislation is still quite immature, data centers and the data they store are normally subject to the laws of the country in which they physically reside. Data privacy and protection laws in that country may differ significantly from that of your homeland, and there's always the possibility of natural disasters, revolutions, and local wars creating a chaotic situation that can result in your data being accessed by people that should not have access to it.

The underworld

Aside from global surveillance threats, we should keep in mind that the data centers of major cloud providers are attractive targets for hacker gangs, since a single successful break-in attempt can result in access to an impressive amount of potentially valuable information. We are, therefore, expecting a sharp rise in the number of professional attacks targeting cloud services in the near future, and believe that a certain percentage of them will prove to be successful. The underlying technology is not going to be the only target; a bribed or blackmailed service provider's employee might be as helpful in gaining access to the data as a rootkited server, and at a much lower cost for the attacker.

What is worse is that it's really tempting for cloud service providers to hush up the hacks. Making this information public affects the provider's reputation, and privacy regulators may apply tangible financial sanctions for privacy legislation violations. Finally, there is usually little or no direct evidence of the theft, and there is always a chance that the attackers won't brag about it publicly. All this means that there is no real incentive for the service providers to announce the theft until it becomes evident from information provided by other, indirect sources.

One should not underestimate the risks of data theft from the cloud service providers' facilities due to technical or human error. History shows us that even the largest and most respected companies make mistakes in implementing information security measures. The need to provide data reliability and accessibility makes the service providers' job difficult. The customers' data is usually stored in highly redundant form, across a variety of sites with synchronization links between them. For an attack to be successful, it's enough for just one of those sites or links to contain a flaw. And sometimes there is no need to attack the working infrastructure at all. Hard drives are known to eventually fail - can you be entirely sure that none of the failed hard drives, scrapped and sent to a landfill by the service provider, won't end up in the wrong hands at some stage of the disposal process?

Cloud: A safe or a cloakroom?

What I was trying to point out above is that storing sensitive data in the cloud poses many risks. “So what”, you ask, “show me a business with no risks at all”. That makes sense. As a matter of fact, the presence of risks and their potential impact are of little importance for a business. The only thing that matters is the business' ability to adequately mitigate those risks.

Imagine that you went to a bank in the morning and deposited $1000 in cash in your bank account. Later that day, two men in balaclavas forced the cashier to hand them all the cash – including your $1000. Does it mean that the bank will withdraw that $1000 from your account and leave you with nothing? No, because the bank is responsible for the money you deposit with them. They take the responsibility of storing your money safely, and take on all the risks associated with this commitment. The exact methods they use to deal with the risks – building underground bunkers, hiring security staff, or transferring risks to insurance companies - are not something you need to care about.

Cloud services have that in common with banks - the only difference is that it's not your money you hand to them, but your megabytes. They take on the responsibility of providing the maximal levels of data availability and protection. But what happens if your data is stolen or destroyed? The scope of the cloud providers' responsibility is written down in their service-level agreement (SLA), and I'll bet you won't be happy when you finally open and read it:

You are responsible for … taking your own steps to maintain appropriate security, protection and backup of Your Content, which may include the use of encryption technology to protect Your Content from unauthorized access...

NEITHER WE … WILL BE RESPONSIBLE FOR ANY COMPENSATION, REIMBURSEMENT, OR DAMAGES ARISING IN

CONNECTION WITH: ... ANY UNAUTHORIZED ACCESS TO, ALTERATION OF, OR THE DELETION, DESTRUCTION, DAMAGE, LOSS OR FAILURE TO STORE ANY OF YOUR CONTENT OR OTHER DATA.

Basically, cloud service providers disclaim any responsibility for disclosing your data to any other party, as well as for losing or damaging it. They also suggest that you "take your own steps to maintain appropriate security and protection of your data". Therefore, mitigating the corresponding risks becomes your own responsibility. A good solution would be to transfer the risks to specialized insurance companies. Unfortunately, due to the novelty of the cloud concept and apparent difficulties in assessment of financial equivalents of such risks as data loss or revelation, insurance companies struggle to offer any product suitable for the cloud realities. I have high expectations about insurance companies introducing proper security quality standards for cloud services, as it is difficult to imagine a better incentive for a service provider to improve the quality of their product than the cost of insurance premiums paid by their customers.

And while insurance companies are working hard on their new cloud-specific products (we hope), users of cloud services must take care of the security of their data themselves.

What's in the box?

The majority of cloud service providers have their own security measures in place. Let's summarize what is on offer and what kind of protection those measures actually provide. The majority of providers offer the following security instruments: secure TLS-driven data transfer, user authorization and, sometimes, server-side encryption (SSE). The scheme below shows a typical route of data between the customer's environment and the cloud.

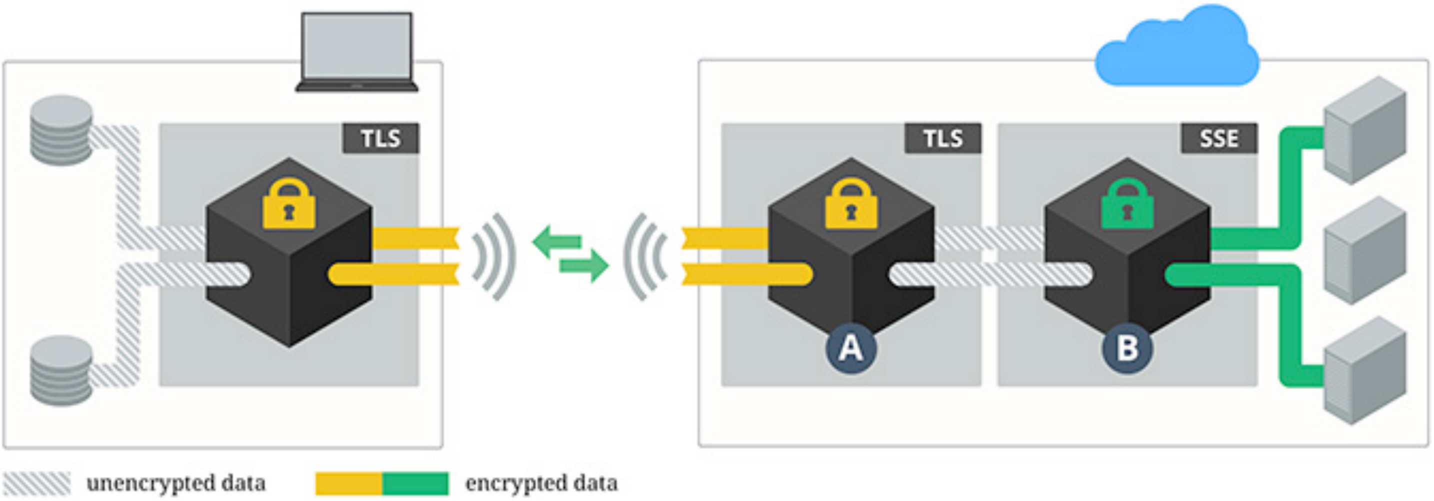

Picture 1. Data flow between the customer's premises and the cloud storage.

The TLS protocol secures data transfers between two peers, or, in our case, between the customer's computer and the cloud front end. The (correct) use of TLS gives the customer the ability to establish the genuineness of the service endpoint and protect the data exchanged from passive or active interception.

Each side of the protocol works like a black box, which takes plain data from the application layer and sends encrypted output to the network. When receiving data, the black box applies a reverse permutation to the encrypted data and passes the decrypted data up the stack. With the help of user authorization the service identifies a particular account holder and confirms his identity. The authorization is typically based on establishing the fact of ownership of an authorization token. The most widely used types of authorization tokens are access keys, digital certificates and passwords. It is important to keep in mind that any entity in possession of an authorization token can send authorized requests to the service on behalf of the legitimate account holder. Therefore it is really important to keep authorization tokens safe.

Server-side encryption (SSE) is an additional layer of security adopted by certain cloud service providers. Before sending the customer's data to the data centre, the provider encrypts it with a strong symmetric cipher. If an attacker gets access to the data centre, they will be unable to decrypt the data without having the encryption key.

The most critical point of this scheme is that the key and the data encrypted with it reside in the same environment, the cloud provider's network. Even though it may be stored in a completely isolated subnet whilst the data remains in the data centre, the key still "meets" the data at point B where the encryption or decryption is performed. Besides, the data in unencrypted form is fed to the input of the TLS server endpoint A, which normally is a front-end web server, subject to the relevant attack risks. In case an adversary gets access to one of those points, they will be able to read the customer's data in the clear without any need for the encryption key.

It should be noted that if your authorization token is leaked, server side encryption won't prevent data theft. With the token in his possession, an adversary can easily impersonate a legitimate user and get access to the data through the service's public REST or SOAP interface.

Finally, server-side encryption won't protect the data from the eyes of intelligence agencies.

Ready. Steady. Go.

I hope I persuaded you that data of any positive significance must always be covered with an extra layer of protection in addition to the security measures provided and adopted by the service providers. One of the simplest yet effective methods of building such an extra layer is to adopt client-side encryption (CSE) (in contrast to server-side encryption offered by the service providers).

The main idea behind CSE is that the customer encrypts their data before sending it over to the service provider, and decrypts the data after downloading it back. With CSE, the customer keeps the encryption key securely in his own environment rather than entrusting it to the service provider. This way, the customer's data goes through the A and B points in customer-encrypted form and can't be recovered by an adversary who controls those points.

Even if a lucky attacker or intelligence personnel gain full control over the provider's computational infrastructure, they will be unable to recover the data.

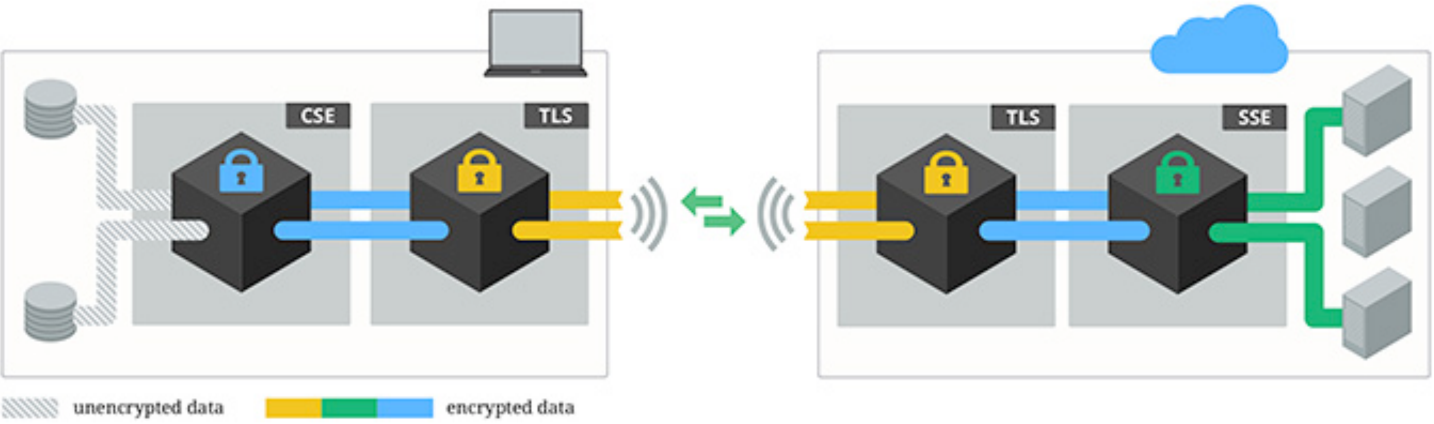

Picture 2. A modified data flow with CSE in force.

Even if an adversary manages to steal the victim customer's authorization token and impersonates them to the service, the data returned by the service will still be in encrypted form, and they will be unable to recover it without obtaining the CSE encryption key.

A variety of encryption keys can be used with the CSE scheme, from generic symmetric encryption keys and asymmetric key pairs up to passwords and customer's biometric information. This flexibility is achieved by separating customers' encryption keys ("user keys") from object encryption keys ("session keys").

The data is first encrypted with a randomly chosen session key (SK). The session key is then encrypted with a customer's user key (for example, a public RSA key). The encrypted data together with the encrypted session key is then sent to the cloud service provider. When reading the data back, the customer first decrypts the session key with their user key, and then uses the session key to decrypt the data.

The session keys-based scheme has a number of advantages, the most substantial of which is its ability to adopt user keys of virtually any nature and the uniqueness of per-object encryption keys. The latter makes it impossible for an attacker who recovers a single session key to decrypt the rest of encrypted objects, as they are encrypted with different session keys.

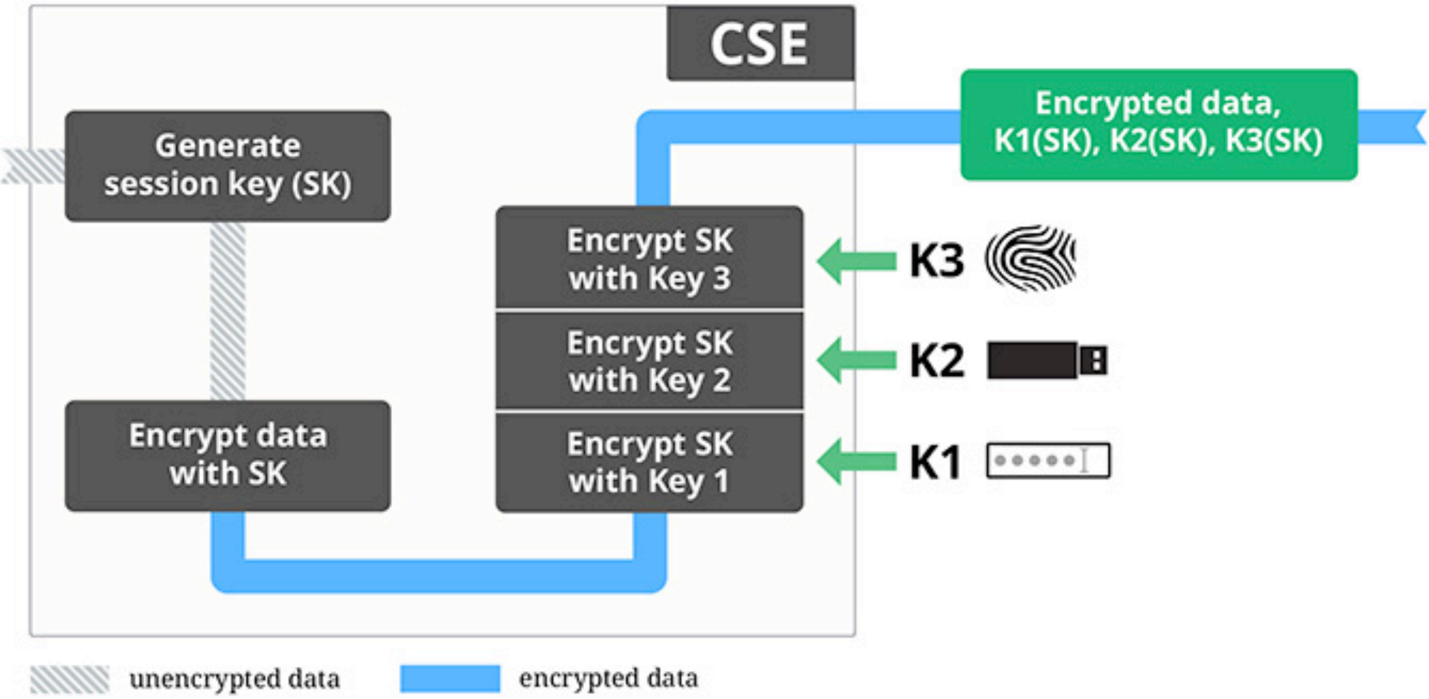

Picture 3. The CSE from the inside.

Aside from encryption, the customer may accompany protected objects with MDP (modification detection and prevention) records. An MDP record is basically a message digest computed by the writer over the unencrypted data and encrypted with it before sending it over to the cloud service. When reading the data back, the customer computes a message digest over the decrypted data and compares it to the message digest attached to the encrypted object. If an MDP-ed object was altered while residing in the cloud or in transit, the message digests won't match.

In certain scenarios it makes sense to use strong digital signatures instead of basic MDP records to provide for a higher level of modification detection in multi-user cloud environments. In particular, strong signatures might be useful for addressing proof of authorship and non-repudiation tasks.

Another attractive side of the CSE is that one can actually encrypt data with several different user keys at the same time. Objects encrypted in such way can be decrypted with any of the keys used for encrypting them. This feature can effectively be used to set up flexible access-rights schemes or establish secure cloud-driven document flow within the organization.

"Wait a minute," an attentive reader will say, "does it mean that the protection of my client-side encrypted data wholly depends on the encryption key?" Correct. The encryption key is a single point of access of the CSE scheme, just like it is in any other properly designed encryption scheme. That's why adequate protection of encryption keys is a task of the highest importance.

Per-object session keys are not an easy target for attackers.

In fact, the only way for an attacker to recover the session key is to guess it by trying different keys one by one. This task can be really tough; provided that the session key is of sufficient length and a cryptographically strong RNG is used for generating it, guessing the session key is an infeasible task. That's why the attacker is most likely to concentrate on gaining access to the user key, which might be much easier to achieve.

Operating systems and third-party vendors provide a variety of mechanisms for securing user keys, from protected operating system areas to dedicated cryptographic hardware. Some only give the key away after authenticating the user who requests it; the others do not export keys at all, instead performing the decryption on board on behalf of the requesting user. In either case it is the responsibility of the customer environment administrator to define and adopt proper key usage and storage policies in order to minimize the risk of revealing the keys to unauthorized parties. Now that you understand the basics of client-side encryption, we can consider the task of integrating the mechanism into existing cloud-driven infrastructures. As I described above, this is mainly a matter of adopting appropriate key management policies, techniques and procedures, and those should be chosen based on the role of the cloud storage service in your infrastructure and the types of logical links between the local environment and the service.

The one-to-one links arise where access to the storage is restricted to one or a fixed small number of the customer's computers. These links can often be employed in scenarios where the cloud is used for storing large amounts of data in a centralized way, mainly as a replacement for classic relational databases. Typical examples of one-to-one links are those used in personal backup and synchronization tools and links between the "worker" and "storage" roles within computational cloud environments.

Where the number of one-to-one links is small, or where there is no need for an extra database abstraction layer, the encryption keys can be stored directly on the computers that access the storage. In larger organizations, access to the storage can be organized through a small set of "cloud gateways" that will be responsible for storing the keys and will optionally provide a standard interface (e.g. ODBC) for outer data storage.

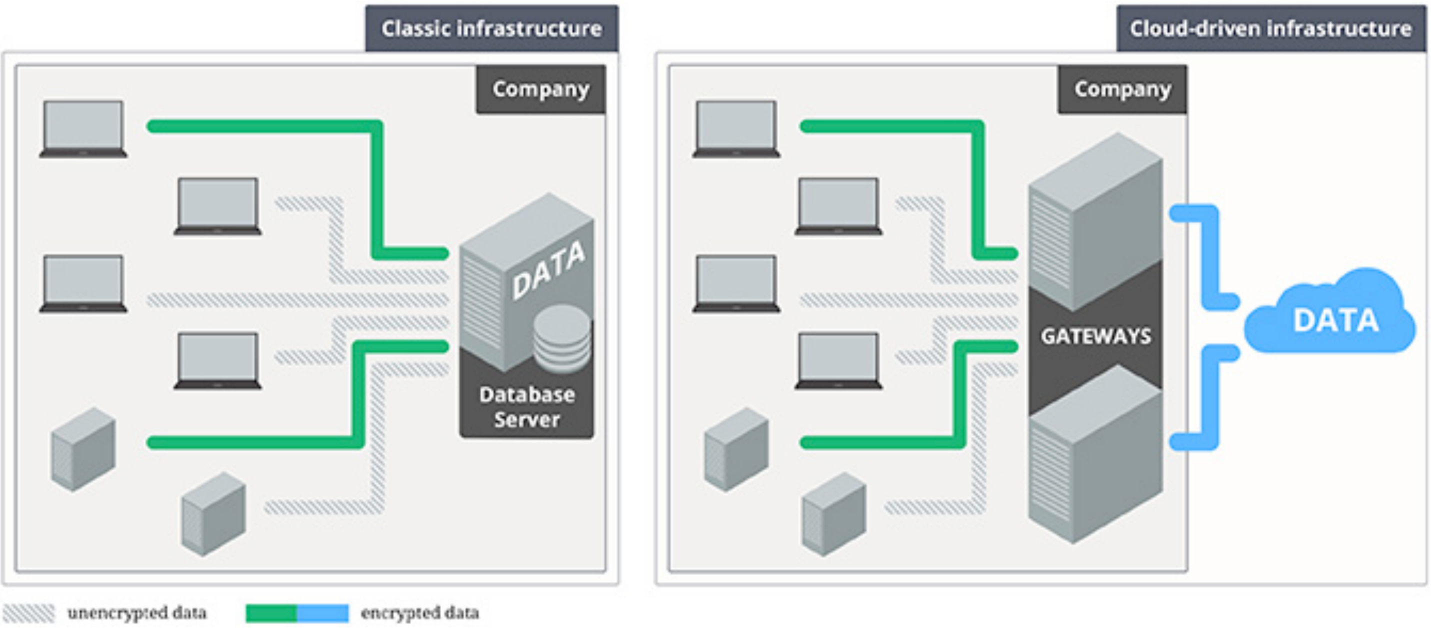

Picture 4. Classic and cloud-driven infrastructures.

The many-to-one links arise in environments characterized by widely decentralized access to the cloud storage and volatile structure, and are typical for scenarios where the cloud is used for sharing or synchronizing data between different company users or departments. As the set of users approved to access the storage is constantly changing with time, encryption keys require a higher level of protection and manageability. This is where enterprise-level hardware security modules may be of great help. HSMs can be effectively used for storing encryption keys securely and enforcing individual access rights to them for every user who needs to access the cloud. Every change in a particular user's position can be easily put in practice by altering the corresponding key access record on the HSM.

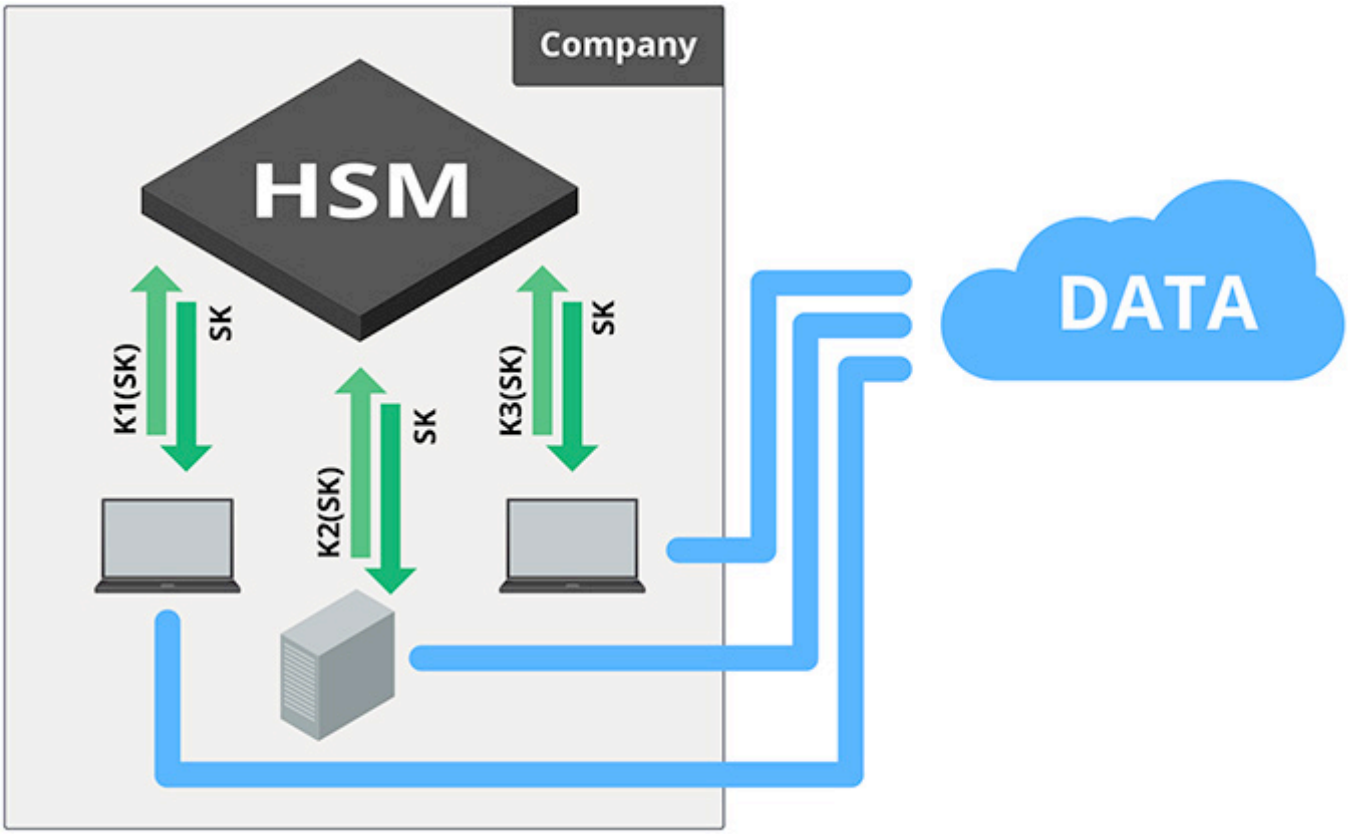

Picture 5. Cloud-driven infrastructure with advanced key management rules in force.

The CSE scheme may effectively be used for managing access rights in environments with complex information access rules by adopting multiple encryption keys procedures. Each logical group of users (e.g. "developers", "colonels", "board members") is assigned a dedicated encryption key. Every object is then encrypted with keys of all groups whose members are allowed to access it. A multiple keys approach can be flexibly integrated into existing corporate group policy rules.

Dollars and cents

How much will it cost to deploy the CSE scheme in a live environment? The final bill will depend on the size of your infrastructure and the chosen key management strategy.

Extra traffic and cloud storage capacity costs are floating somewhere around zero. A protected object is only up to a kilobyte larger than a matching unprotected object, so from a point of view of traffic and space economy it virtually does not matter whether you work with protected or unprotected objects.

Compensation of productivity decrease caused by extra computational burden on the customer's resources. In practice, the burden is imperceptible. Modern processors provide encryption speeds of up to 8-10 MB/sec, much faster than an average speed the cloud storage providers are capable of accepting or giving away the data with. This means that a computer that encrypts outgoing data on-the-fly will be idling for a certain amount of time waiting for the preceding portion of encrypted data to be delivered to the cloud front-end.

Expenses related to key management can make a significant share of the migration budget, especially if many-to-one links are present and HSMs are adopted to manage the keys. The costs in this case directly depend on the size and complexity of the infrastructure, the type of encryption keys involved, and the expected frequency and number of read/write operations. The exact HSM to use should be chosen based on the key management requirements, and should be capable of handling the estimated operation load.

The cost of implementing the CSE protocol may pose another unpleasant surprise. The current labour rate of qualified software developers, especially those with good knowledge of cryptography, is not among the cheapest on the software market, and a proper implementation of the CSE takes a significant amount of time. Happily, the market offers some out-of-the-box software products that can help adopt the CSE in faster, cheaper and user-friendlier way.

Popular SecureBlackbox middleware from EldoS offers cloud storage access components that come with built-in support for CSE. The product supports the majority of popular cloud services, including S3, Azure, SkyDrive, Dropbox and Google Drive. Besides the encryption components themselves, SecureBlackbox offers out-of-the-box support for a variety of encryption key types, from easy-to-use password-based keys to cryptographically strong symmetric and asymmetric keys residing on hardware security devices.

The one-off asking price for the product is fairly competitive and equals to 3-4 hours of work of a qualified security expert - a good bargain, taking into account that integrating SecureBlackbox-driven CSE into an existing infrastructure is an easy-to-do task that only requires common programming skills.

Client-side encryption is an inexpensive and easy-to-integrate method that addresses the risk of private data revelation, topped up with useful access control features. Even if adopted in its simplest form, CSE provides that extremely durable extra layer of protection that your remotely deposited information needs in today's world.

Ken Ivanov, PhD, is a Director and a Chief Security Expert at EldoS Corporation (UK), a provider of security and file system solutions for software vendors. For more than 12 years Ken helped businesses, individuals and governmental agencies all over the world to adopt information security measures. Being really passionate about security and privacy, as well as liberal values, Ken is concerned about the increasing influence of international intelligence services. That's why the question of privacy for data residing in the cloud is of particular importance to him. He can be contacted at Ivanov@eldos.com.

HITB2014AMS

May 27th & 28th 2014 - Hands on Technical Training
May 29th & 30th 2014 - Triple Track Conference

Celebrating 5 years of the HITB Security Conference in The Netherlands

MISCHEL KWON (Former Director, US Computer Emergency Response Team)

KATIE MOUSSOURIS (Lead, Microsoft Security Response Center)

KRISTIN LOVEJOY (General Manager, IBM Security Services Division)

PAMELA FUSCO (Chief Information Security Officer, Apollo Group)

JENNIFER STEFFENS (Chief Executive Officer, IOActive)

JAYA BALOO (Chief Information Security Officer, KPN Telecom)

+ HITB Haxpo

May 28th, 29th & 30th 2014

A 3-day IT security exhibition for hackers // makers // breakers // builders

LOCK PICKING VILLAGE

HAXPO // IT SECURITY EXHIBITION // TECHNOLOGY SHOWCASE AREA

HACKWEEKDAY

CAPTURE THE FLAG (CTF)

HANDS-ON TECHNICAL TRAININGS

SOCIAL ENGINEERING CHALLENGE

MULTI-TRACK SECURITY CONFERENCE

HITB LABS

Registration Opens December 2013

Venue: De Beurs van Berlage
Website: http://haxpo.nl
Follow us: @HITBHaxpo / @HITBSecConf

Supported & Endorsed By

I amsterdam.

Executive hot seat: Intrinsic-ID CEO

Interview by Mirko Zorz

Pim Tuyls is the Founder and CEO of Intrinsic-ID. Tuyls initiated work on Hardware Intrinsic Security within Philips Research in 2002. As a principal scientist, he managed the cryptography cluster in Philips Research in which the initial research was carried out. Later he transferred this work to Intrinsic-ID and headed the technology development.

Based on your experience, what strategic errors do SMBs make when using/moving to the cloud? Does security play a big role in their decision-making?

SMBs mainly think about the benefits of the cloud: scaling and elastic IT-enabled capabilities, higher productivity, leveraging a mobile workforce, to mention a few. But they don't make a proper risk analysis since they are not aware of the possible consequences.
Breached customer data might lead to a lawsuit. By storing data unprotected in the cloud specific laws and regulation might be violated. Recently, we encountered a few cases where SMBs were not aware that their back-up system stores their data unprotected in the cloud. This can lead to huge fines that could have been avoided if an appropriate pre-study was done and security measures had been taken.

Currently, security does not play a big role in their decision-making process, but awareness is growing. The revelations made by Edward Snowden last year have clearly awakened many people and organizations. In-house security expertise is usually missing in small organizations. It is important that they partner with a security expert in time to make sure that their transition to the cloud goes smoothly and leads to success instead of a big headache.

What advice would you give to organizations that have mission-critical data in the cloud? What critical steps should they take in order to ensure ongoing security?

My recommendation to organizations with mission-critical data stored in the could is to conduct a risk and vulnerability profile to understand what kind of data they store in the cloud; what type of cloud the data is stored in (private versus public), and what the consequences are if the data in the cloud gets exposed. In all cases, we recommend implementing the best possible security solution that the budget will allow.

Once organizations understand what is at risk they need to take measures to find and implement a security solution. We recommend a solution that uses a two-factor authentication system based on an independent hardware root of trust, which will ensure that the data stored in the cloud is properly encrypted before it leaves the device. By independent root of trust we mean that the security keys should not be influenced nor known by any other party, and should be generated within the device itself.

Early in the process of moving to the cloud, companies should evaluate cloud security solutions that are best suited to their needs and in line with their budget. In most cases, software-only solutions are not adequate because they have several weaknesses and the highest protection is usually achieved with hardware-based systems that offer client-side protection. Client-side protection means that the keys are in the hands of the customer and cannot be learned by a “trusted party”.

> **SINCE IT’S A CLOSED SYSTEM, A PRIVATE CLOUD OFFERS TEMPTING SECURITY ADVANTAGES**

While security is still the most significant obstacle to cloud adoption, companies also wonder if they should use a private or a public cloud. What criteria should they use to decide?

There are various ways to look at this, but currently we see the following main aspects. What we see more and more in the market is that companies’ tools have to be compatible with those that their employees use in their private life. Employees are accustomed to using their own mobile devices and even their own (most often) public clouds. In that sense, using a public cloud is attractive and companies often decide to use those.

However, this opens the door for all kind of attacks on the companies’ network since there is no control of the software installed on the devices of employees. Such a policy has to be handled with care and security solutions need to be installed on the user’s devices.

Since it’s a closed system, a private cloud offers tempting security advantages. For an outsider, it’s much more difficult to attack such a cloud. But even in that case, a company still depends on the trustworthiness of the owner or administrator of the private cloud. Nevertheless, I think that the main consideration should be that they use a system that guarantees that they have full control of their data and absolute privacy.

Even if a company is inclined to opt for a private cloud, they should choose one that has a security system in place (maybe via a partner) – one that ensures full protection of the company data so that even if/when the cloud provider is hacked, the privacy of the company data is guaranteed.

Edward Snowden's revelations have impacted on, among other things, the way organizations think about their data in the cloud. As a European company, have you seen an uptick in business or concern? What have been the reactions among your clients?

Absolutely! We see that many organizations are beginning to take data privacy very serious, especially in Europe and the US. They have started projects to implement data protection systems or are setting up projects to analyze their vulnerabilities. Some projects aimed at further integrating the cloud have been delayed till a clear picture of the vulnerabilities and risks are in place.

OUR VISION IS TO ELIMINATE DATA SECURITY AND USER AUTHENTICATION ISSUES IN THE CONNECTED WORLD

This year almost all organizations, both large and small, that we are in contact with are considering how to improve their security against attacks and accidents.

Clearly, small companies are moving faster with implementations of systems to protect their data in the cloud, while large organizations need more time to build a vulnerability and risk profile and depend more on consensus building.

How does Intrinsic-ID enable organizations to protect their cloud data? What makes you stand out in the marketplace?

What Intrinsic offers customers is peace of mind knowing that they are getting the highest-level of security solution for the price and in a way that is easy to implement. We offer our customers a full solution to authenticate to the cloud and protect data in the cloud based on an independent root of trust that is unique. It makes sure that nobody outside the company can access the company's data.

Our solutions consist of two components: software that runs on PCs or mobile platforms and a physical token that implements the independent root of trust. This root of trust is based on HIS technology or the electronic fingerprint of a chip. Since all chips are unique due the deep-sub micron process variations, this guarantees unique, very high quality keys to start with: the keys have full entropy or randomness and are therefore almost impossible to guess even for governmental organizations.

On top of this, our approach makes sure that our clients have full control over their security and don't have to be nervous about the fact that a backdoor might be in place that allows organizations to listen in or break in to their systems. Our secure cloud product not only delivers security but also allows our customers to work securely in the cloud. Therefore we have implemented a secure sharing mechanism: it allows users to share a document with one or more colleagues without disclosing any information to outsiders.

Finally, our solutions have been developed by a team with in-depth security expertise in all the aspects of a security system. The team consists of hardware security specialists, software security experts and system security architects. This combination of expertise allows us to build a strong, integrated security product.

What are your plans for the rest of the year? What areas are you focusing on?

Our vision is to eliminate data security and user authentication issues in the connected world, and we will continue to work to simplify and improve methods for achieving that. We intent to further extend our secure cloud offering to other platforms and enhance it with more features that allow it to work together with various mobile and cloud services, thus providing a complete solution to our customers.

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

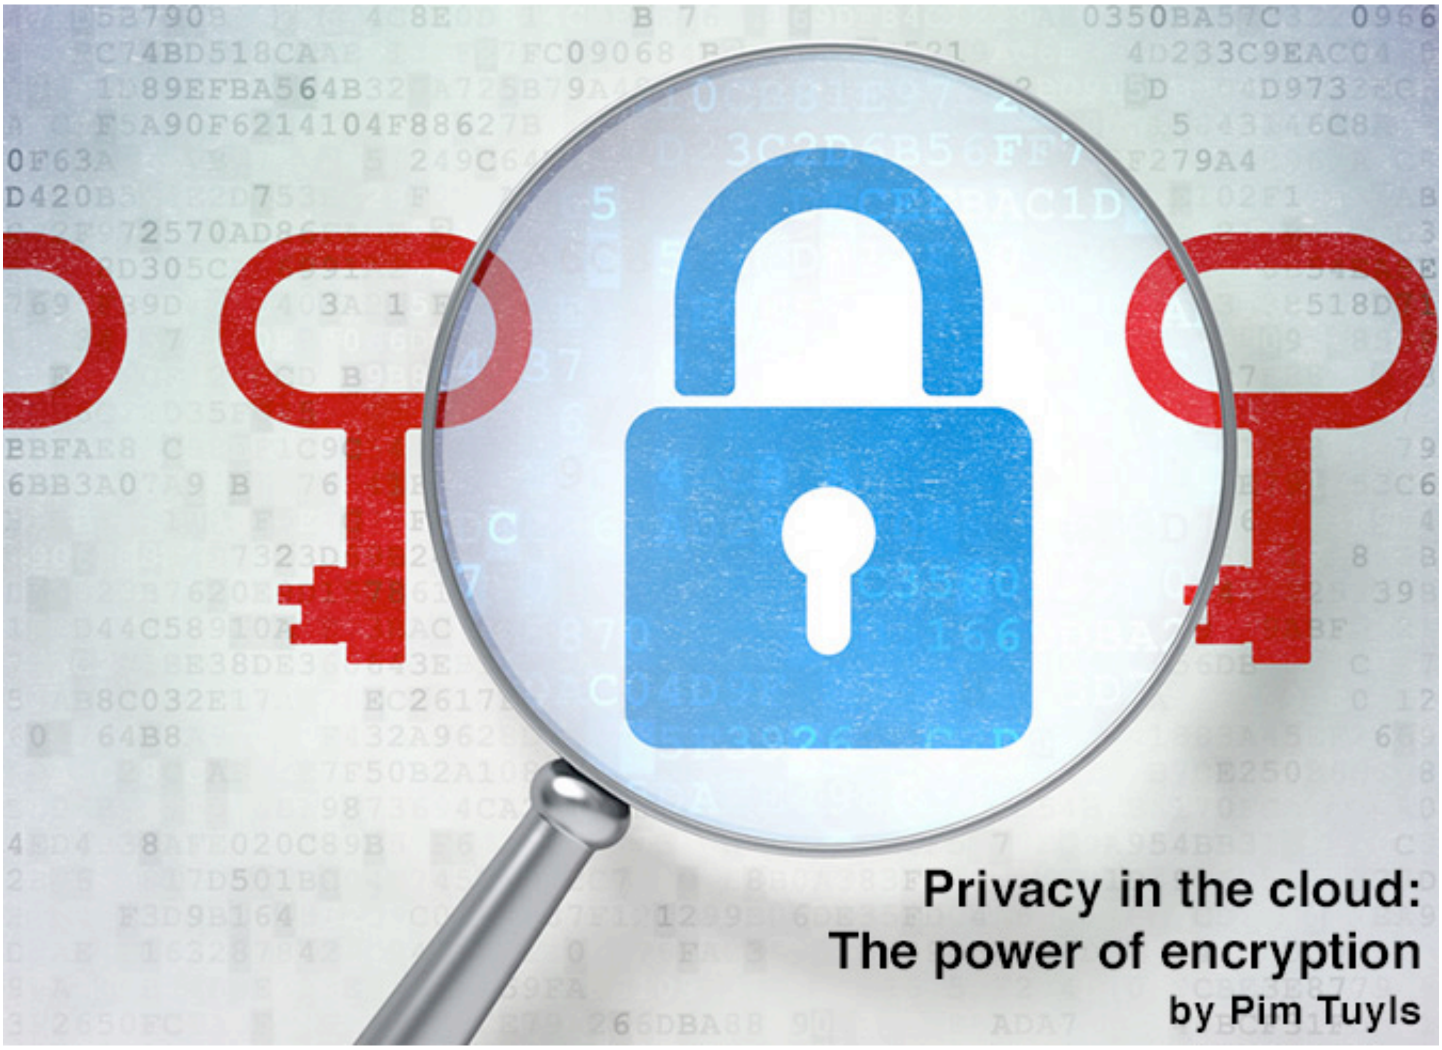

Privacy in the cloud: The power of encryption

by Pim Tuyls

Adequate protection of your digital data can be very straightforward. Just like your own house, the best way to secure it is with a unique physical key that you keep in your own hands.

Our virtual world keeps expanding and the “anytime, anywhere” cloud is growing. More and more, employees integrate their own devices and favorite apps in their daily work.

Traditional security measures are no longer enough, and coming up with good security solutions for the cloud has been challenging.

CIOs realize that there is no single solution to curb the increasing risks, and that a layered approach is needed.

Encryption systems in the cloud

An age-old method for preventing sensitive information falling into the hands of unauthorized people is encryption. Also, concerning the cloud, protection by encryption has become an increasingly important activity, especially when it comes to the mitigation of the risk of breaches.

More and more, cloud providers, but also specialized companies, offer encryption/authentication tools for data in or on the way to the cloud. In Figure 1 I illustrate three categories of solutions: server-based, gateway-based and client-based encryption systems.

a. Server-based encryption

In a server-based encryption system, the server encrypts all data stored on the server. This implies that a (master) key is hidden on the server, too. Such a system offers a basic level of protection, but it does not solve the problem of secure cloud storage.

On the one hand, this system assumes implicit trust in the owner of the server encryption key. On the other hand, when the encryption key is compromised in an attack, all the data of all the users of the system are exposed simultaneously.

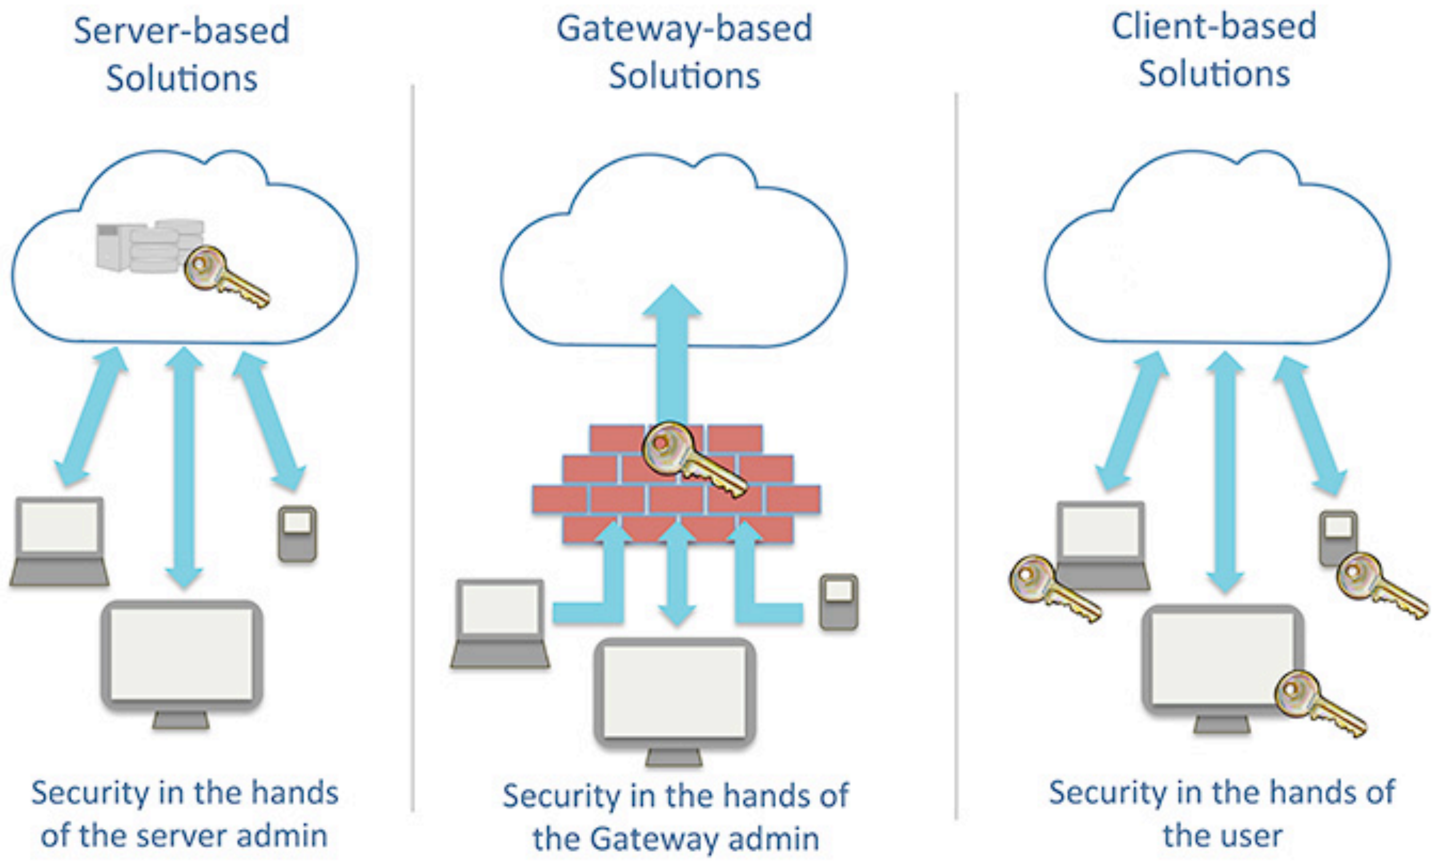
Server-based Solutions
Gateway-based Solutions
Client-based Solutions
Security in the hands of the server admin
Security in the hands of the Gateway admin
Security in the hands of the user

Figure 1. Encryption/authentication solutions for the cloud can be divided in three categories, according to who controls the security keys.

b. Gateway-based encryption

Gateway-based systems install a gateway within the company. All company traffic that goes to the outside world will be encrypted with a key that is stored in the gateway.

The advantage of this system stems from the fact that enterprise unique keys are used to protect the data stored in the cloud. A successful attack on the system of one enterprise will not automatically translate to a successful attack on the data of a neighboring enterprise. But, it will reveal all the data of the enterprise victim. Furthermore, there is still an implicit assumption of trust in the owner of the keys that are stored on the gateway.

c. Client-based encryption

Client-based encryption systems encrypt the data at the client side. This means that the keys are generated at the client side and all files are encrypted on the client device before being sent to the cloud. No trust has to be given to another party for the generation and the management of the keys.

Within this category there are software- and hardware-based systems. In software-based systems, the cryptographic keys are derived from passwords that have to be provided and remembered by the user.

Due to this human component (creating and remembering a password with high randomness is difficult) the cryptographic keys will not contain a lot of entropy. When cryptographic keys are not sufficiently random, there is a risk that these keys will be compromised using brute-force (guessing) attacks.

In hardware-based systems, the keys are generated from physical randomness of the hardware that the user owns: a smartcard, USB stick or Micro SD card. This guarantees the highest entropy level of the keys, i.e. they cannot be brute-forced. A hardware-based client side encryption system encrypts the data at the client side with the strongest keys without having to trust another party.

This solution gives control back to the user and has a clear parallel with securing our home in the physical world.

Strongest (unclonable) key in the hand of the user

Cryptographic keys used in electronic devices are traditionally stored in non-volatile memory (typically secure EEPROM or E-fuses). However, this approach is sensitive to specific security issues, like the eavesdropping and tampering of the keys.

Instead of storing keys in non-volatile memory, it is nowadays possible to generate and store secure keys based on unique physical properties of the underlying hardware.

This approach is called Hardware Intrinsic Security (HIS) and makes use of the concept of Physical Unclonable Functions (PUFs). The principle of a PUF can best be described as "biometrics for electronic devices" and is illustrated in Figure 2.

The different threshold voltages for the transistors in SRAM result in a cell-unique start-up behavior. Start-up values establish a unique and robust fingerprint that is turned into a secure secret key. Keys are not present in the hardware when the device is not powered. The fact that the keys are only generated when they are required minimizes the "window of opportunity" for an attacker to compromise them.

The strength of this kind of hardware-based key generation and storage can be combined with a username and password-based login system. This combination is called two-factor authentication. This means that security is based on two unrelated factors - "something you know" and "something you have".

Given that both of these factors are required to access cloud data, it will become even more difficult for an attacker to gain unauthorized access to or manipulate the data. Two-factor authentication provides incredibly strong protection for all data in the cloud.

d. Secure data storage in the cloud

Secure data storage in the cloud can be achieved by using symmetric key cryptography based on, for example, AES. In symmetric key cryptography files are encrypted and decrypted using the same key. I strongly recommend a key that is securely generated from and stored in the hardware. Since encryption and decryption are only performed on the client side, this key never leaves the user's hardware and is therefore completely secure from malicious use.

Due to the client-side encryption, files intercepted during transmission or illegally accessed in storage are unusable.

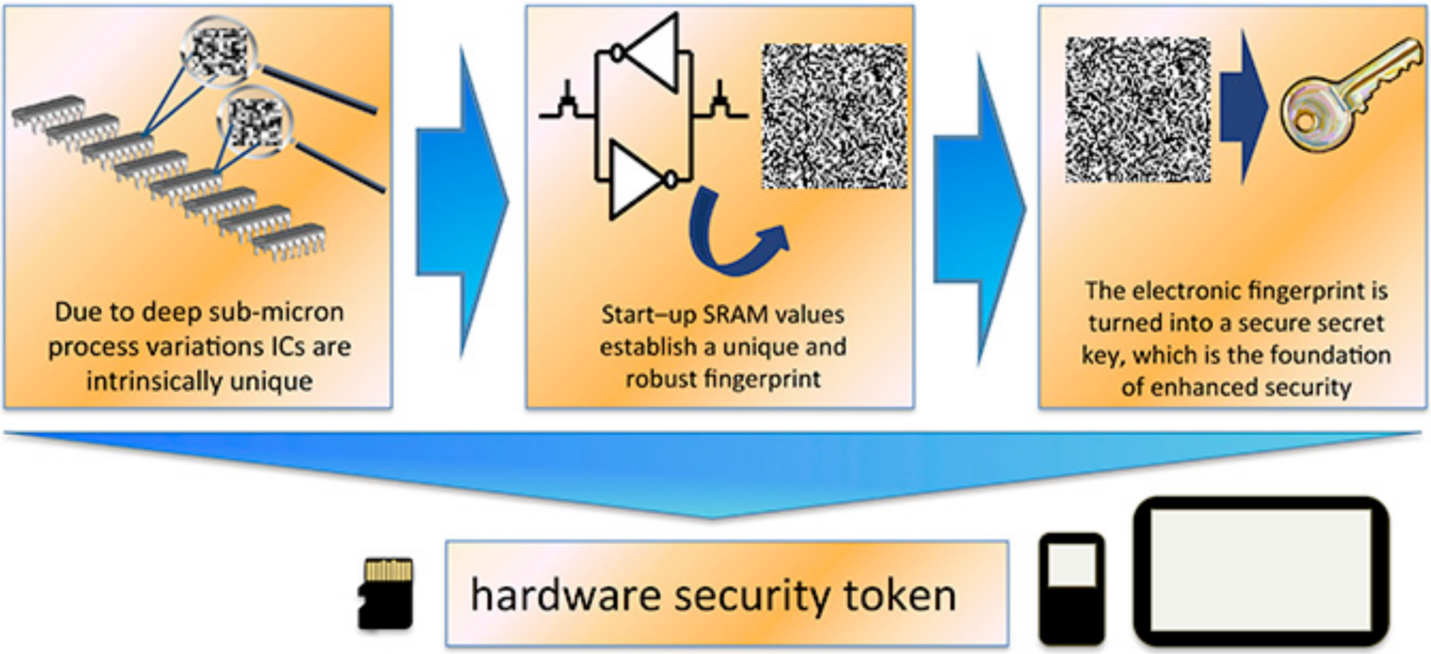
Due to deep sub-micron process variations ICs are intrinsically unique
Start-up SRAM values establish a unique and robust fingerprint
The electronic fingerprint is turned into a secure secret key, which is the foundation of enhanced security
hardware security token

Figure 2. The principle of Physical Unclonable Functions (PUFs): secret keys are extracted from the "biometrics of an electronic device".

Only the legitimate user (password) that is in possession of the key (hardware) can download the files stored in the cloud onto his device, after which they can be automatically decrypted in the background.

e. Secure sharing of data in the cloud

Secure file sharing can be done using asymmetric cryptography, also known as public-key cryptography. In that case, a public/private key-pair has to be generated during the installation of the hardware for each user. Although different, the two parts of such a key pair are mathematically linked. The public key is used to encrypt files, while the private key is used to decrypt files. While the private key never leaves the hardware token of the user, the public key should be stored in a public key database (see Figure 3).

Since this public key cannot be used to decrypt data stored in the cloud, there is no security risk involved for the user when publicly storing this key. Only with the securely-stored private key can data from the cloud be decrypted.

Before sharing a file with another user in the cloud, the public key of the receiving user must be retrieved from the public key database. The file is encrypted with the public key before it is stored in the cloud in a location that the recipient can access. After the recipient has been notified that a shared file is available, the file can be downloaded from the cloud onto the recipient's device. On this device, where the recipient's private key is available, the downloaded file can be decrypted and used. Since the private key of the recipient is the only key that mathematically matches the public key used for encryption, the recipient is the only person that can decrypt this file. In case attackers are able to intercept the file, they will not be able to decrypt it because they lack the private key.

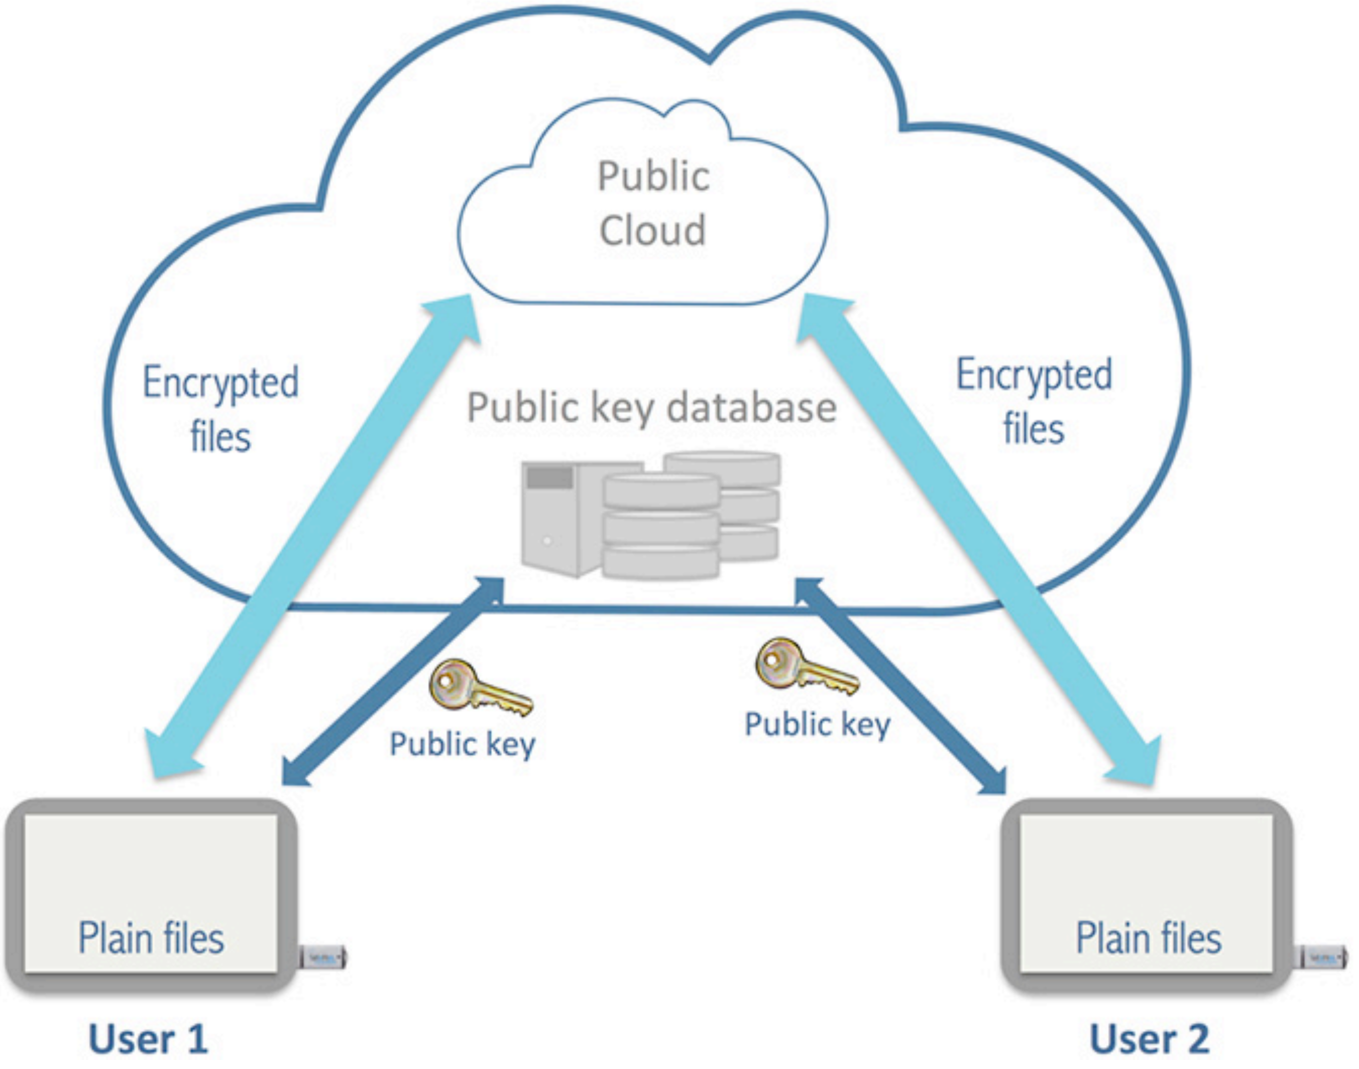

Figure 3. The flow of information for sharing encrypted data in the cloud.

The future

The cloud is a blessing for the modern economy, but recent events have also made it very clear that it can easily turn into a nightmare if security issues are not properly managed. In order to avoid disasters and enable the cloud's true potential, security has to be an integral aspect of any cloud-based system.

A strong two-factor authentication implementation as described in this article provides the strongest security of data in the cloud without compromising flexibility, performance or ease of use. Data can be safely available anytime, anywhere from multiple devices.

HIS-enabled chips and tokens have the potential to offer top-level security and key management flexibility to protect a whole new class of applications. Apart from securing payment transactions and provisioning media content, these applications may include machine-to-machine, smart grid, track-and-trace and many others as they emerge from the rapidly developing Internet of Things.

Dr. Pim Tuyls is the CEO of Intrinsic-ID, a provider of security IP cores and applications based on patented Hardware Intrinsic Security (HIS) technology.

Prior to Intrinsic-ID, Pim was at Philips Electronics where he initiated work on Hardware Intrinsic Security as a principal scientist within Philips Research. Several of Pim's papers relating to secure implementations of Physical Unclonable Functions (PUF) technology have been published at security conferences.

Thousands of FTP sites compromised to serve malware and scams

Some 7,000 FTP sites and servers have been compromised to serve malware, and its administrators are usually none the wiser, claim Hold Security researchers.

FTP sites function as online file caches and are accessible remotely. Users who have the required login credentials can upload and download files from them, but other users can also retrieve certain files hosted on such a server if given a specific link that leads to the file (and without needing to provide login credentials). It is this latter capacity that makes login credentials to FTP servers a prized haul for cyber scammers, as they upload malware and malicious links to the server, then embed direct links to them in spam emails delivered to potential victims. Access to an FTP server can also be occasionally leveraged by the attackers to compromise connected web services.
"The victim companies hosting exploited FTP sites are spread across the spectrum – from small companies and individual accounts with ISPs to major multi-national corporations," noted the researchers.

"Hackers planted PHP scripts armed with backdoors (shells) and viruses in multiple directories hoping that these directories map to web servers of the victim companies to gain control of the web services. They also uploaded HTML files with seamless re-directs to malicious sites."

Alex Holden, the company's CISO, has shared that among the compromised file transfer servers are also some that belong to The New York Times and UNICEF. Affected organizations have been notified of the problem and some have already moved to fix it. It is unknown who stole the FTP credentials, and who is using them, but judging by the complexity of some of the passwords, it's natural to assume that they haven't been guessed, but stolen via information-stealing malware.

PoC mobile malware records swipes on touch screen smartphones

A security researcher has developed proof-of-concept malware capable of capturing screenshots and finger swipes on mobile devices, and is set to demonstrate his creation at the RSA Conference in San Francisco.

As users inexorably switch from computers to mobile phones and tablets for their online shopping and banking needs, malware developers will surely set their sights on creating financial malware adapted to the new platforms.

Neal Hindocha, senior security consultant for Trustwave, has proved that the feat can be done, but his proof-of-concept malware does not lend itself to industrial scale data collection.

You see, this "touchlogger" is capable of logging the coordinates of any swipe or touch, and of taking screen captures, but it can't (yet) distinguish between general use and use for online payment and banking purposes.
Going through the swipe logs and screenshots to search for relevant information is a very time-consuming and labor-intensive task that's difficult to automate, so using it against specific targets seems more logical and more efficient than infecting a huge number of devices.

Hindocha has managed to make the PoC code work on jailbroken iOS and rooted Android devices, but shared with Forbes that it's possible get it to work on non-rooted Android devices as well - possibly by infecting the device when it's plugged into a PC.

The malware might be aimed at capturing financial information, but since it records every touch, an attacker would also discover additional relevant information such as the user's security code for unlocking the device, image and alpha-numeric passwords, etc.

Java-based malware hits Windows, Mac and Linux

Kaspersky Lab researchers have recently analyzed a piece of malware that works well on all three of the most popular computer operating systems - the only thing that it needs to compromise targeted computers is for them to run a flawed version of Java.

The Trojan is written wholly in Java, and exploits an unspecified vulnerability (CVE-2013-2465) in the JRE component in Oracle Java SE 7 Update 21 and earlier, 6 Update 45 and earlier, and 5.0 Update 45 and earlier.

Once the malware is launched, it copies itself into the user’s home directory and sets itself to run every time the system is booted. It then contacts the botmasters' IRC server via the IRC protocol, and identifies itself via a unique identifier it generated.

The malware's main reason of existence is to make the infected machine flood specified IP addresses with requests when ordered to via a predefined IRC channel. The botmasters simply have to define the address of the computer to be attacked, the port number, the duration of the attack, and the number of threads to be used in it.

At the time of analysis, the botnet formed by machines "zombified" by this particular Trojan was targeting a bulk email service.

Sophisticated cyber-espionage tool uncovered

Kaspersky Lab discovered “The Mask” (aka Careto), an advanced Spanish-language speaking threat actor that has been involved in global cyber-espionage operations since at least 2007. The primary targets are government institutions, diplomatic offices and embassies, energy, oil and gas companies, research organizations and activists. Victims of this targeted attack have been found in 31 countries around the world – from the Middle East and Europe to Africa and the Americas.

The main objective of the attackers is to gather sensitive data from the infected systems. These include office documents, but also various encryption keys, VPN configurations, SSH keys (serving as a means of identifying a user to an SSH server) and RDP files (used by the Remote Desktop Client to automatically open a connection to the reserved computer).

The campaign was active for at least five years until January 2014. During the course of Kaspersky Lab’s investigations, the C&C servers were shut down. 380 unique victims between 1000+ IPs have been counted.

The complexity and universality of the toolset used by the attackers makes this cyber-espionage operation very special. This includes leveraging high-end exploits, an extremely sophisticated piece of malware, a rootkit, a bootkit, Mac OS X and Linux versions and possibly versions for Android and iOS. The Mask also used a customized attack against Kaspersky Lab’s products.

Among the attack’s vectors, at least one Adobe Flash Player exploit (CVE-2012-0773) was used. It was designed for Flash Player versions prior to 10.3 and 11.2.

The malware intercepts all the communication channels and collects the most vital information from the infected system. Detection is extremely difficult because of stealth rootkit capabilities.

Fully functional trojanized FileZilla client steals FTP logins

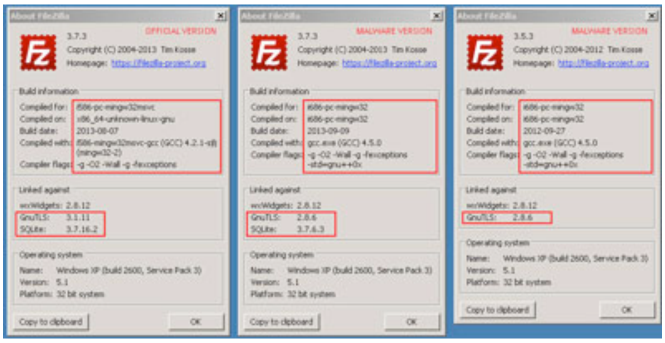

```
About FileZilla                                   ×
3.7.3                              OFFICIAL VERSION
Copyright (C) 2004-2013  Tim Kosse
Homepage: https://filezilla-project.org

Build information
Compiled for:   i586-pc-mingw32msvc
Compiled on:    x86_64-unknown-linux-gnu
Build date:     2013-08-07
Compiled with:  i586-mingw32msvc-gcc (GCC) 4.2.1-sjlj
                (mingw32-2)
Compiler flags: -g -O2 -Wall -g -fexceptions

Linked against
wxWidgets: 2.8.12
GnuTLS:    3.1.11
SQLite:    3.7.16.2

Operating system
Name:     Windows XP (build 2600, Service Pack 3)
Version:  5.1
Platform: 32 bit system

Copy to clipboard                               OK
```

```
About FileZilla                                   ×
3.7.3                               MALWARE VERSION
Copyright (C) 2004-2013  Tim Kosse
Homepage: https://filezilla-project.org

Build information
Compiled for:   i686-pc-mingw32
Compiled on:    i686-pc-mingw32
Build date:     2013-09-09
Compiled with:  gcc.exe (GCC) 4.5.0
Compiler flags: -g -O2 -Wall -g -fexceptions
                -std=gnu++0x

Linked against
wxWidgets: 2.8.12
GnuTLS:    2.8.6
SQLite:    3.7.6.3

Operating system
Name:     Windows XP (build 2600, Service Pack 3)
Version:  5.1
Platform: 32 bit system

Copy to clipboard                               OK
```

```
About FileZilla                                   ×
3.5.3                               MALWARE VERSION
Copyright (C) 2004-2012  Tim Kosse
Homepage: http://filezilla-project.org

Build information
Compiled for:   i686-pc-mingw32
Compiled on:    i686-pc-mingw32
Build date:     2012-09-27
Compiled with:  gcc.exe (GCC) 4.5.0
Compiler flags: -g -O2 -Wall -g -fexceptions
                -std=gnu++0x

Linked against
wxWidgets: 2.8.12
GnuTLS:    2.8.6

Operating system
Name:     Windows XP (build 2600, Service Pack 3)
Version:  5.1
Platform: 32 bit system

Copy to clipboard                               OK
```

Trojanized versions of the popular FileZilla FTP client are being offered to unsuspecting users via hacked websites with fake content.

"Malware installer GUI is almost identical to the official version. The only slight difference is version of NullSoft installer where malware uses 2.46.3-Unicode and the official installer uses v2.45-Unicode. All other elements like texts, buttons, icons and images are the same," Avast researchers warn.
"The installed malware FTP client looks like the official version and it is fully functional! You can’t find any suspicious behavior, entries in the system registry, communication or changes in application GUI."

It's interesting to note that one of the malicious versions has been compiled way back in September 2012, and is still detected by just a couple of commercial AV solutions. Another one dates back to September 2013, and is also poorly detected.

"We assume that the stolen FTP accounts are further abused for upload and spread of malware. Attackers also can download whole webpage source code containing database log in, payment system, customer private information etc," the researchers pointed out.

Researchers uncover months-old POS malware botnet

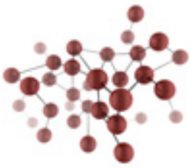

With the Target and Neiman Marcus breaches, the topic of malware that collects card data directly from Point-of-Sale devices has received renewed interest. The PoS malware used in the former has been identified as a modified version of the BlackPOS malware, but there are other similar ones currently in use out there.

RSA researchers have discovered the entire server infrastructure used in a global PoS malware operation that targets retailers in the US, Russia, Canada and Australia, and have managed to access part of it. The malware in question is the ChewBacca Trojan, which is capable of logging keystrokes and scraping the memory of PoS systems and the card magnetic stripe data they contain.

"RSA observed that communication is handled through the TOR network, concealing the real IP address of the C&C server(s), encrypting traffic, and avoiding network-level detection," they noted. "The server address uses the pseudo-TLD '.onion' that is not resolvable outside of a TOR network and requires a TOR proxy app which is installed by the bot on the infected machine."

"The ChewBacca Trojan appears to be a simple piece of malware that, despite its lack of sophistication and defense mechanisms, succeeded in stealing payment card information from several dozen retailers around the world in a little more than two months," the researchers noted.

"Retailers have a few choices against these attackers. They can increase staffing levels and develop leading-edge capabilities to detect and stop attackers (comprehensive monitoring and incident response), or they can encrypt or tokenize data at the point of capture and ensure that it is not in plaintext view on their networks, thereby shifting the risk and burden of protection to the card issuers and their payment processors."

Android bootkit infects 350,000 devices

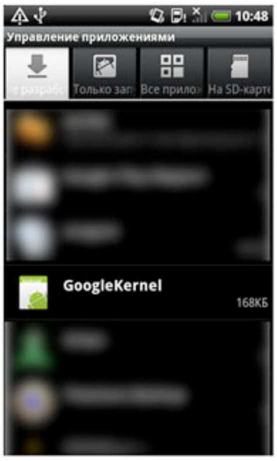

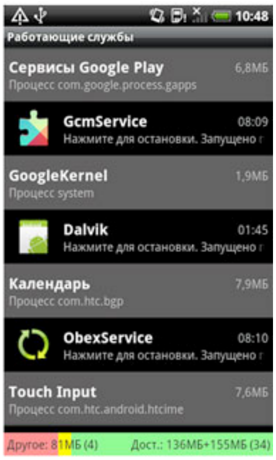

The first ever Android Trojan with bootkit capabilities has been discovered and analyzed by Dr.Web researchers, who warn that the malware is already operating on some 350,000 mobile devices around the world. The malware - dubbed Oldboot - resides in the memory of infected devices and launches itself early on in the OS loading stage, they say, and believe that the Trojan is being distributed via modified firmware. To ensure persistence, the attackers have inserted one of the Trojan's components into the boot partition of the file system, and have altered the script that is tasked with initializing the OS components.

"When the mobile phone is turned on, this script loads the code of the Trojan Linux-library imei_chk, which extracts the files libgooglekernel.so and GoogleKernel.apk and places them in /system/lib and /system/app, respectively," the researchers explained.

"Thus, part of the Trojan Android.Oldboot is installed as a typical application which further functions as a system service and uses the libgooglekernel.so library to connect to a remote server and receive various commands, most notably, to download, install or remove certain applications."

Even if other elements of the Trojan are removed successfully, the modified script will restart the installation process by triggering the imei_chk each time the device is rebooted.

Beware of malicious specialized software keygens

Masquerading malware as key generators for popular games is a well-known malware-delivery tactic, but it's not often that you see malicious keygens for other types of software.

Nevertheless, it happens occasionally. Trend Micro researchers warn that malware peddlers have lately been targeting with this approach professionals working in a variety of industries. They have spotted fake generators for specialized (and expensive) engineering (Aveva) and automotive repair (AllData) software, multimedia tools (Bigasoft), benchmarking software (Geekbench), software for chemists and biologists (CambridgeSoft), computing software (Wolfram Mathematica) and, yes, some games.

Unfortunately, the offered executables are not what they seem. Once installed, they pave the way for other malicious software to be installed on the compromised computer, and lately that software is often a fake AV variant.

"Fake antivirus software has declined significantly from its heyday several years ago (in part due to crackdowns on their payment systems)," the researchers pointed out. "Since then, it has been overshadowed by first police ransomware and then in more recent months by CryptoLocker."

Mac Bitcoin-stealing Trojan lurks on download sites and GitHub

CoinThief, the recently discovered Bitcoin-stealing Trojan that targets Mac users, has been spotted being offered on several download websites such as CNET's Download.com and MacUpdate.com, as well as masquerading as precompiled binaries in several GitHub projects.

The malware's initial variant installs browser extensions for Safari and Google Chrome that monitor all web browsing traffic, looking specifically for login credentials for many popular Bitcoin websites as well as Bitcoin wallet sites and login credentials. These newer variants have already been made to include also a browser extension for Firefox ("Pop-Up Blocker 1.0.0").

"The malware is being distributed disguised as price tickers for Bitcoin ("Bitcoin Ticker TTM for Mac") and Litecoin ("Litecoin Ticker"), which have been available on download.com since early December. According to the download stats, the malware has been downloaded 57 times," SecureMac researchers noted. Fortunately, the two websites have already reacted and removed the malware.

In a Reddit thread initiated by Nicholas Ptacek, lead developer at SecureMac, the developer of Bitcoin Ticker TTM has noted that his original app was never open source, so it seems like his app was never trojanized, and that only its and his name was used to trick users into downloading the malware. Ptacek also shared that the malware is being distributed on GitHub in the BitVanity and StealthBit projects and wrote in details about how to remove the malware from the system if you have been infected.

Still, it would be probably wrong to assume that the malware is not still being distributed on other download sites and under different names, so be careful when downloading anything, and check for the malicious extension.

How to recover deleted or corrupted digital currency

by Zeljka Zorz

The popularity of Bitcoin and other digital/cryptographic currency cannot be denied. Different users like using it for different reasons, but many agree that the question of keeping their stash safe is something that occasionally keeps them up at night.

With the currency's rising popularity, different services started popping up to help solve that dilemma. Still, some chose to keep their digital money on their own devices. But what happens if these devices break down, get corrupted, or get accidentally erased?

Chris Bross, Senior Enterprise Recovery Engineer at well-known data recovery firm DriveSavers, has recently been involved in a few cases where customers needed their digital currency recovered.

"In the past, we have typically been asked to recover user-created data files like photos, QuickBooks, videos, etc. Now we are being asked to recovery digital currency," he shared with (IN)SECURE Magazine. "We did not particularly start offering recovery services to Bitcoin users, customers just came to us begging for help."

Usually, the value of the data on any device is gauged by the user to whom it belongs. The recovery process can be relatively expensive, but when Bitcoins are at stake, the value of the recovery makes it even more worthwhile.

So far, the company has had 8 requests for digital currency recovery, 7 for Bitcoin and 1 for Litecoin.

"In general, hardware failures of the storage devices are more common, but in many of these cases the users were completely at fault because they deleted the data, and in two cases, even overwrote it," he explained.

These two cases ended with the experts recovering the wallet.dat files, but the encrypted data was unfortunately corrupted beyond repair.

"Having said that, no situation is entirely hopeless and we always want to attempt the data recovery, as early as possible after the data loss event has occurred, to prevent the user from making a poor choice that leads to additional loss," he notes.

While the company usually deals more with companies needing to recover data, all of their requests for Bitcoin recovery so far have been from individuals.

"In almost all cases, the user was keeping only a single copy of their wallet.dat file on a single storage device," says Bross. "Three of the eight total cases were stored on a SSD (sold state drive), while the others were on HDDs. In all cases, the users did not trust the available cloud backup or other options to replicate their critical Bitcoin data."

In one particular case where the customer stored his digital wallet on his Microsoft Surface Pro tablet, performed a number of steps that led to data loss AND disabled the computer's ability to boot up, they had to develop a new method to be able to create an image of the SSD without physically removing it from the tablet.

"Since it was an SSD, there were concerns about BGC (background garbage collection) processes that would erase and sanitize the missing data as a normal function of system maintenance. We needed to intercept and halt those processes to mitigate any additional loss of data," he shared.

"Do you also offer the service of tracing and recovering stolen Bitcoin?" I asked, and received a negative answer. "Although, we do offer forensic data recovery services for any case that may end up in litigation or prosecution," Bross noted.

Finally, I asked him to share some tips on how to keep one's digital currency safe.

"Print it on paper. Ironically, in these modern times of digital everything, paper and ink are more reliable and last longer than magnetic or solid state storage," he offered.

"In addition, keep multiple copies of the files on other secured and encrypted media that you personally control yourself. If you choose to store this data in the cloud, be well aware of potential cloud security vulnerabilities and risks."

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

> Visit www.insecuremag.com
> SUBSCRIBE TO (IN)SECURE MAGAZINE

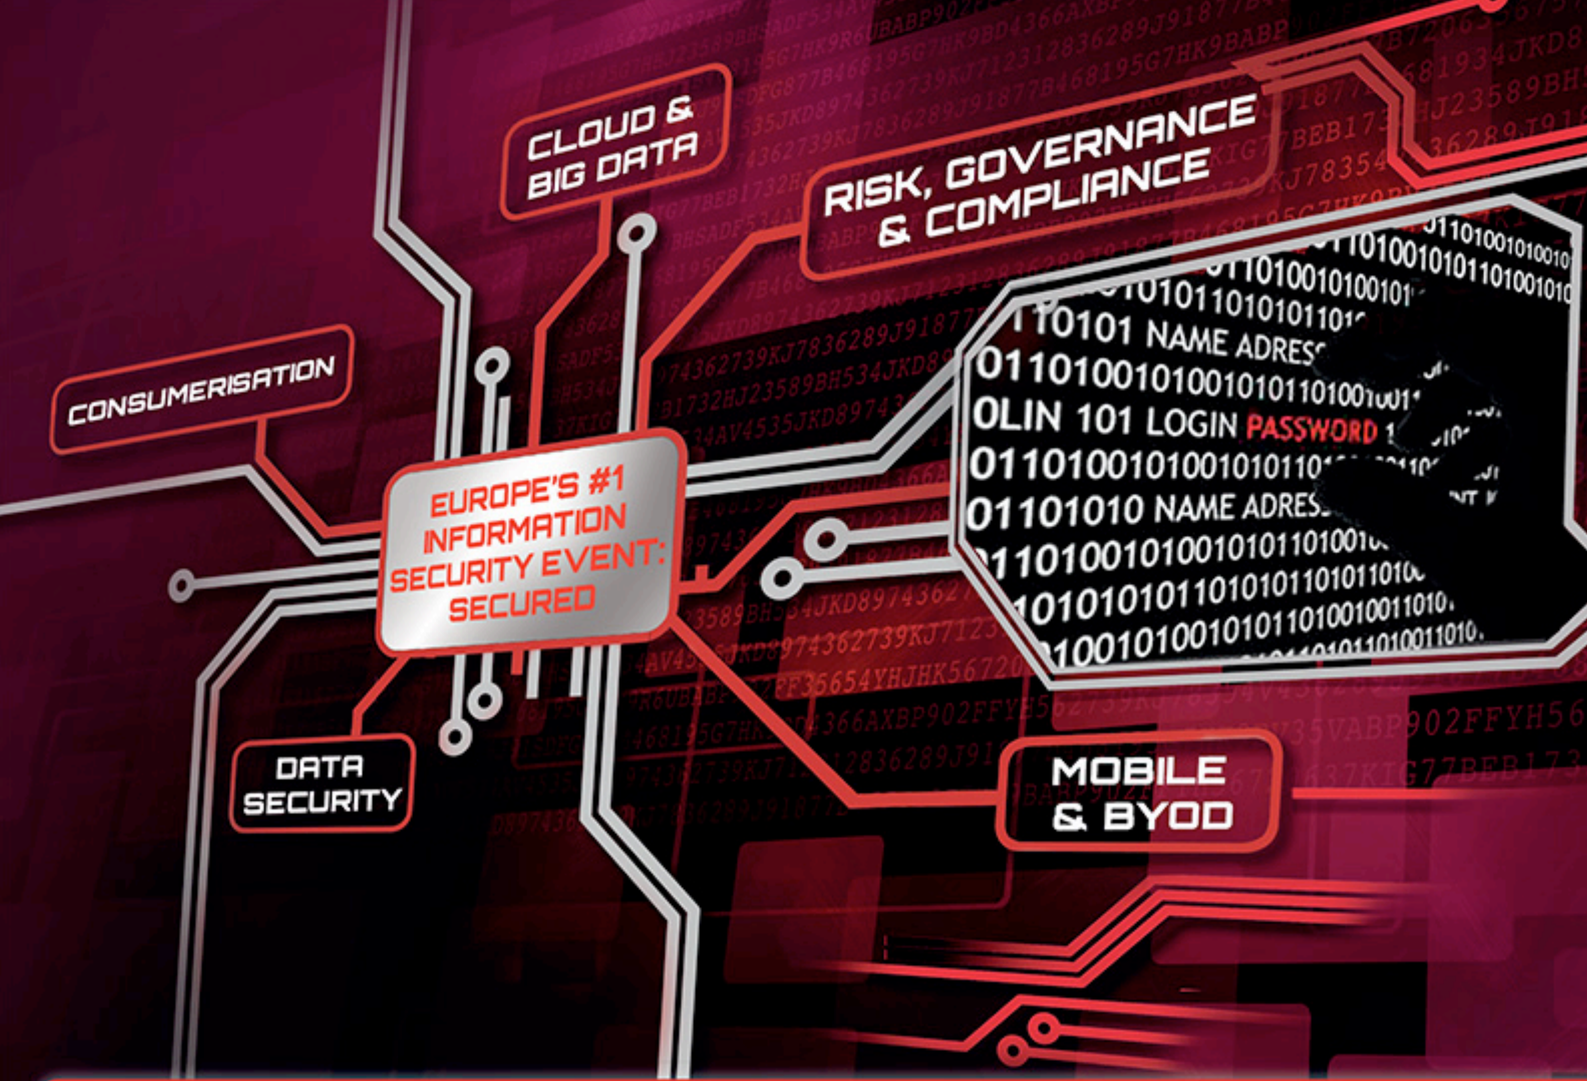

CLOUD & BIG DATA
RISK, GOVERNANCE & COMPLIANCE
CONSUMERISATION
EUROPE'S #1 INFORMATION SECURITY EVENT: SECURED
DATA SECURITY
MOBILE & BYOD
NAME ADRESS
LOGIN PASSWORD

COLLECT CPE • CREDITS • COLLECT CPD

CYBERCRIME, THREATS & ATTACK VECTORS

Is your business secure?

Meet industry leading vendors & suppliers and learn how to optimise your organisation's information security posture!

Infosecurity secures business – come and secure yours!

3-Day programme highlights at a glance!

Tuesday 29th April	
10.00 - 10.30	OPENING KEYNOTE INTERVIEW: Establishing a framework for collaboration and intelligence sharing between government and industry to address new cyber threats
10.45 - 11.50	Security as an enabler: Supporting enterprise innovation and transformation
12.05 - 13.05	Actionable intelligence: Building an holistic security threat intelligence capability
13.20 - 14.20	The 2014 Cyber Security Breaches Survey official launch
14.35 - 15.35	'Applification' of business and implications for security: Securing agile development
15.50 - 16.40	Balancing security versus usability: Optimising the user experience to maximise security
16.55 - 17.30	RESEARCH KEYNOTE: Mapping the threat horizon
Wednesday 30th April	
10.00 - 10.35	INFOSEC PERSPECTIVES: Whistle Blowing: Threat or Opportunity?
10.50 - 11.45	Building transparency and trust in the supply chain: How to secure data in the extended, connected enterprise
12.00 - 13.00	What's new in cybercrime?: Keeping up with the cybercriminal
13.15 - 14.15	Future-proofing information security and protecting connected legacy systems
14.30 - 15.30	Why can't we do that, why can't we have that….? Rethinking information security education strategies to engage generation Y
15.45 - 16.35	Big Data security intelligence: Using Big Data to detect threats, fraud and breaches
16.50 - 17.30	EU General Data Protection Legislation: Status update and key action points for organisations
Thursday 1st May	
10.00-10.35	HALL OF FAME KEYNOTE INTERVIEW 'In the News'
10.50-11.35	Risk and control: Effective risk assessment methodologies to drive security strategy and investment
11.50-12.20	Case study: Protect, detect, respond: Anatomy of an effective incident response strategy
12.35-13.30	Privacy & Compliance Think Tank: Utilising compliance as an information security asset
13.45-14.45	Redefining cyber security for critical national infrastructure in a hyper-connected world: Securing SCADA, ISC and smart grids
15.00-16.00	Data centric security: Protecting data in the interconnected, mobile, social, cloud-based enterprise

Download the Infosecurity Europe App

Engage with Infosecurity Europe on Twitter: @infosecurity #infosec14

Leveraging Big Data for security operations

by Josh Goldfarb

Talk of Big Data seems to be everywhere these days. A simple web search for the phrase returns countless vendors offering solutions to the challenge, myriad events discussing the topic, and blogs, articles, and forums filled with talking points.

Amidst this sea of information, organizations are starved for practical knowledge they can leverage internally to make some sense of the Big Data puzzle in order to meet their challenging operational requirements.

This article aims to provide some practical guidance that organizations can leverage to begin to navigate their way through the Big Data challenge.

Data value vs. data volume

Distilled to its essence, Big Data is about two symbiotic, but somewhat diametrically opposed components: collection and analysis. Both are equally important, but the more data one collects, the more difficult analysis becomes due to the volume and variety of data. Much of this article will discuss ideas relating to the analysis component of this challenge, but a brief discussion regarding collection is helpful to frame the discussion. As the readers of this article are aware, the proper instrumentation of an enterprise network for network security monitoring soon results in both an overwhelming variety of data sources and a nearly unmanageable volume of total data.

As one might expect, different data sources carry different value for security operations. For example, proxy logs, with their rich metadata and relatively compact size provide high value to security operations.

Conversely, logs from routers or switches, though helpful for network operators, are significantly more voluminous but provide less value for security operations purposes.

When architecting a log collection strategy, most organizations do not assess the value of each data source to security operations and instead resort to collecting every data source they can obtain from the enterprise. This approach produces a few cascading effects:

- More data requires more storage, which adds cost.
- In place of additional storage, organizations typically cut the retention period of logs (e.g., from six months to three months), which hinders the ability to perform historical forensics
- The additional volume of data creates a larger volume of noise, which in turn makes identifying anomalous traffic more difficult
- Performance slows down (e.g., queries return less quickly), which in turn reduces productivity and efficiency.

Additionally, regardless of the actual length of the retention period, most organizations apply a uniform retention period. In other words, the policy is that after N days (whatever N is), all data rolls off to archive and ceases to be accessible as part of day to day operations.

Another benefit of assessing data value is that organizations can better understand which data sources should be retained longer than others.

When attempting to manage enormous volumes of data effectively, it's important to streamline and optimize workflow at every opportunity.

Streamlining workflow

When attempting to manage enormous volumes of data effectively, it's important to streamline and optimize workflow at every opportunity. It is most effective to facilitate the organization's human resources (most often the scarcest) working smarter, not harder. How to achieve this in practice requires more details than this article allows for, but some general guidelines are provided here:

- Develop efficient processes for all functions (e.g., incident response for a malicious code incident)
- Automate within the process where possible (not for automation's sake)
- Enrich data where possible to further optimize the process (not for enrichment's sake)
- Focus analysts on one centralized work queue/alerting stream
- Enable/facilitate rapid and efficient investigation/analysis with the end goal of resolution (no analysis for analysis' sake).

Integrating actionable intelligence

Intelligence can greatly add to the both the detection and response capabilities of the enterprise. There are a lot of sources (whether free or by subscription) referring to themselves as security intelligence sources. In order for information to be intelligence, it must be timely, reliable, high fidelity, and actionable.

In order for information to meet these requirements, it must have context and a use case. Organizations use the context and use case to best leverage the intelligence. For example, integrating intelligence regarding malicious email attachment MD5 hashes into the alerting stream requires a much different approach than integrating intelligence regarding command and control (C2) URL patterns/substrings.

If the intelligence comes at a relevant time (timely), does not produce high volumes of false positives (reliable), is of high quality (high fidelity), and produces relevant alerting (actionable), it can be leveraged according to its specific context and use case.

Communal presence

The security community is a relatively small and tight-knit community built almost entirely on trust. Many of the relationships and the most trusted information-sharing exchanges within the industry are built on a strong foundation of trust. These exchanges require a give and take –sharing information and methods, as well as receiving information and methods. The information and methods exchanged this way are most often of extremely high quality. A strong presence in the community, particularly by the organization's leadership can greatly aid in obtaining practical, timely, and high quality information and methods to assist with the Big Data challenge. Conversely, leadership that maintains a strong presence in the community is more likely to motivate the organization to be an active participant in trusted information sharing exchanges.

Remembering the user

While systems get infected, users create, edit, use, and share data. In many sophisticated attacks, the adversary is after the organization's sensitive, proprietary, and/or confidential data. With the volume of data (legitimately) entering and leaving the network on a daily basis, it is quite difficult to identify anomalous / malicious activity within this data. Most contemporary analysis techniques are IP (or domain) centric. If the vantage point (for some analytical techniques) is shifted to a user-based perspective, additional possibilities emerge. When examining human beings, as opposed to machines, it becomes much easier to analyze activity (both historically and in near real-time) to identify trends and departures from normal behavior. This enables the enterprise to build richer analytical techniques and alerting.

Summary

The challenge of Big Data is maturing from a marketing buzzword to a reality confronting large enterprises. The volume of data brings new challenges, but it also brings new technologies, opportunities, and capabilities. Big Data reminds us to consider both aspects of collection and analysis, where one necessitates the other. As new technologies emerge, it is important to remember that they will help with the challenge, but they alone cannot solve the problem. Smooth and seamless integration of people, process, and technology is just as important as always.

Joshua Goldfarb (www.yourcyberanalyst.com) is a cyber security analyst with over a decade of experience building, operating, and running SOCs. Before joining nPulse Technologies (www.npulsetech.com) as its CSO, Goldfarb worked as an independent consultant, applying his analytical methodology to help enterprises build and enhance their network traffic analysis, security operations, and incident response capabilities to improve their information security postures. Earlier in his career Goldfarb served as the Chief of Analysis for US-CERT.

The past, present, and future of Big Data security

by Ulf Mattsson

While Apache Hadoop and the craze around Big Data seem to have exploded out into the market, there are still a lot more questions than answers about this new environment. One of the biggest concerns, second perhaps only to ROI, is security.

This is primarily due to the fact that many have yet to grasp the paradigm shift from traditional database platforms to Hadoop. Traditional security tools address a separation of duties, access control, encryption options, and more, but they are designed for a structured, limited environment, where data is carefully collected and cultivated.

Hadoop, on the other hand, is an environment with limited structure, high ingestion volume, massive scalability and redundancy, designed for access to a vast pool of multi-structured data. What's missing is new security tools to match.

Another challenge with securing Hadoop comes from the rapid expansion of the environment itself. Since its initial development, new tools and modules have been coming out not only from Apache, but nearly every other third-party vendor as well. While security is tested and implemented for one module, three more have come out and are waiting for the same treatment. This makes it very difficult to create an overall security architecture for the entire Hadoop ecosystem as it continues to grow. However, some security tools have been released over the last few years, including Kerberos, which provides strong authentication. But Kerberos does little to protect data flowing in and out of Hadoop, or to prevent privileged users such as DBA's or SA's from abusing the data. While authentication remains an important part of the data security structure in Hadoop, on its own it falls short of adequate data protection.

Another development was the addition of coarse-grained volume or disk encryption, usually provided by data security vendors. This solved one problem (protecting data at rest) but considering one of the primary goals behind Hadoop is *using* the data, one might suggest that it provided little in the grand scheme of Big Data security. Sensitive data in use for analytics, traveling between nodes, sent to other systems, or even just being viewed is subject to full exposure.

Up until recently, Big Data technology vendors have often left it to customers to protect their environments, and they, too, feel the burden of limited options.

Today, vendors such as Teradata, Hortonworks, and Cloudera, have partnered with data security vendors to help fill the security gap. What they're seeking is advanced functionality equal to the task of balancing security and regulatory compliance with data insights and "big answers".

The key to this balance lies not in protecting the growing ecosystem, or blanketing entire nodes with volume encryption, but targeting the sensitive data itself at a very fine-grained level, with flexible, transparent security. Applying this security through a comprehensive policy-based system can provide further control and additional options to protect sensitive data, including multiple levels of access to various users or processes. Once secured, the data can travel throughout the Hadoop ecosystem and even to outside systems and remain protected.

The options for fine-grained data security in Hadoop now include encryption (AES or format-preserving), masking, and Vaultless Tokenization.

Typically, encryption is the least desirable option, as standard strong encryption produces values that are unreadable to the tools and modules in Hadoop, format-preserving encryption is typically much slower than masking or Vaultless Tokenization, and both require complicated cryptographic key management across tens or even hundreds of nodes.

Masking was developed for non-production systems and testing, and has found a home in Hadoop's early, experimental phase. Individual data elements are either replaced with random values or generalized so that they are no longer identifiable. It is fast, produces values that are readable to systems and processes, and requires no key management. However, because masking was designed for non-production, it is usually not reversible, and is therefore not ideal for any situations where the original data may be needed sometime after the data is masked.

Vaultless Tokenization, similar to masking, also replaces data elements with random values of the same data type and length. It is also much faster than format-preserving encryption, virtually eliminates key management, and is transparent to processes. The added benefit comes from the ability to perform both one-way protection and reversible security.

This provides ideal protection for test/dev environments and can also allow retrieval of the original data when required by authorized users or processes.

THE OPTIONS FOR FINE-GRAINED DATA SECURITY IN HADOOP NOW INCLUDE ENCRYPTION, MASKING, AND VAULTLESS TOKENIZATION.

Due to the read-only nature of the Hadoop environment (files cannot be updated, you can only create a file, read it and delete it), application of these fine-grained protection methods requires a unique approach.

This is typically performed in one of two ways. The first is a secured gateway, situated in front of Hadoop, which parses incoming data to identify sensitive data elements, and applies the selected protection method before passing the data on to Hadoop. The second is a secured landing zone, which may be a node or partition within Hadoop that is protected with coarse-grained encryption. Files arrive in the landing zone, and are then parsed by one of the processing applications in Hadoop (MapReduce, Hive, Pig, etc.), identifying and protecting sensitive data elements before ingesting the data into the main Hadoop cluster. This method utilizes the massively parallel processing of Hadoop to efficiently protect data.

In the next five years, the creation of data by more and more people and devices will continue to drive the companies towards Hadoop and other Big Data platforms. The requirements for handling extreme levels of volume, velocity, variety, and veracity will only increase, and Big Data will assume more and more critical business functions.

As the environment becomes more established, usability and enterprise integration will improve, new data exchange protocols will be used, and a set of security tools will be standardized and made native to platforms.

Laws and regulations relating to privacy and security will also continue to increase, and security will become an even more vital component in Big Data. Companies will be unable to harness the massive amounts of machine-generated data from the Internet of Things without implementing comprehensive data security - first in the area of industrial environment (power grids, etc.) and later on consumer use (healthcare, etc.). Security will be viewed not only in terms of loss-prevention, but value creation, enabling compliant data collection, use, analysis, and monetization.

Big Data security will evolve, becoming increasingly intelligent and data-driven in its own right. We will see more tools that can translate security event statistics into actionable information. Data security policies will be intricately designed, and likely multi-layered, utilizing a combination of coarse- and fine-grained security methods, access control, authentication, and monitoring.

In the exciting near future, the data is only getting bigger, but we must not allow it to outgrow security.

Ulf T. Mattsson is the CTO of Protegrity. Ulf created the initial architecture of Protegrity’s database security technology, for which the company owns several key patents. His extensive IT and security industry experience includes 20 years with IBM as a manager of software development and a consulting resource to IBM's Research and Development organization, in the areas of IT Architecture and IT Security. Ulf holds a degree in electrical engineering from Polhem University, a degree in Finance from University of Stockholm and a master's degree in physics from Chalmers University of Technology.

MIS TRAINING INSTITUTE'S

INFOSEC WORLD

CONFERENCE & EXPO 2014

April 7-9, 2014 | Disney's Contemporary Resort | Orlando, FL | Bonus Workshops: April 5, 6, 9, 10, 11

70+ Sessions Hand-Picked by Leading InfoSec Experts to Give You the Tools You Need to Protect Your Organization

Register Today!

www.misti.com/infosecworld

KEYNOTE SPEAKERS

Stewart A. Baker
Partner, Steptoe & Johnson LLP; Former First Assistant Secretary for Policy, United States Department of Homeland Security

Rob DuBois
Retired Navy SEAL; Security and Policy Advisor, Speaker and Author

Malcolm Harkins
Vice President, Chief Security and Privacy Officer, Intel Corporation

Chris Hazelton
Research Director, Mobile & Wireless, 451 Research

CO-LOCATED SUMMITS:
- CISO EXECUTIVE SUMMIT - APRIL 6
- IT AUDIT MANAGEMENT SUMMIT - APRIL 9, 10, 11
- CLOUD SECURITY SUMMIT - APRIL 10

www.misti.com/infosecworld

MIS TRAINING INSTITUTE The International Leader in Audit & Information Security Training

PLATINUM SPONSOR LIEBERMANSOFTWARE

GOLD SPONSOR IBM

SILVER SPONSOR LogRhythm

CISO SUMMIT SPONSOR hp

GLOBAL EDUCATION SPONSOR (ISC)²

Information stewardship: Avoiding data breaches and managing Big Data

by Mike Small

Information security continues to be an increasing problem because of the volume, velocity and variety associated with Big Data. This article summarizes the nature of information security risks and describes how better information stewardship based on information centric security is essential to manage these risks.

The term “Big Data” is as much a reflection of the limitations of the current technology as it is a statement on the quantity, speed and variety of data being generated. Big Data needs to be understood as data that has greater volume, variety or velocity than can be comfortably processed using the technology that we already have.

Big Data comes from a number of sources, both internal and external. Many organizations have accumulated large amounts of data that is not being exploited.

There is an even larger amount of data that is held in publicly available sources, like government databases, and social media. In addition to this, the inbuilt instrumentation of smart systems generates a massive amount of still untapped data.

To realize its potential value Big Data needs to be transformed into Smart Information, which can then be used to improve planning and increase efficiency as well as to create new kinds of products.

Information security challenges

The underlying information security challenges of malice, misuse and mistake apply equally to Big Data. Big Data techniques can also be used by criminals to improve their exploits, provide insight that facilitates security breaches, and aggregate data to assist with identity theft. Big Data can be misused through abuse of privilege by those with access to the data and analysis tools; curiosity may lead to unauthorized access and information may be deliberately leaked. Mistakes can also cause problems - corner cutting could lead to disclosure or incorrect analysis.

There are three major risk areas that need to be considered:

Information lifecycle: Big Data turns the classical information lifecycle on its head. The provenance of the data may be doubtful, the ownership of the data may be subject to dispute, the classification of the information discovered may not be feasible until after analysis. For all of these reasons the compliance requirements and needed controls cannot easily be predetermined.

Data provenance: Big Data involves absorbing and analyzing large amounts of data that may have originated outside of the organization that is using it. If you don't control the data creation and collection process, how can you be sure of the source and the integrity of the data? How do you know that you have the right to use the data in the way that is being planned? These points are defined very clearly in a UK report on the use of smart metering for power consumption by utility companies.

Technology unknowns: The technology that underlies the processing of Big Data was conceived to provide massively scalable processing rather than to enforce security controls. While this is not a new phenomenon in the IT industry, there has not been sufficient time for the inherent vulnerabilities and security weaknesses to become manifest.

Looking after property that is not your own is called stewardship. Information stewardship is not a new term; it has been in use since the 1990s and covers a wide range of challenges involved in managing information as a key organizational asset. These challenges include the management of the whole information lifecycle from ownership to deletion, as well as aspects like business value, data architecture, information quality, compliance and security.

The basic objectives of information security for Big Data are the same as for normal data: to ensure its confidentiality, availability, and integrity.

To achieve these objectives certain processes and security elements must be in place. There is a large overlap with the normal information security management processes, however specific attention is needed in the following areas:

Everyone is responsible: The unstructured nature of Big Data means that it is difficult to assign the responsibility to a single person. Everyone in an organization needs to understand their responsibility for the security of all of the data they create or handle. This means creating a culture of security.

Verification of data source: Technical mechanisms are needed to verify the source of the external data used (for example, digital signatures).

System integrity: Good oversight and control over the integrity of the systems used for analysis is needed, and that also goes for privilege management and change control. Be careful to validate conclusions – if you can't explain why the results make sense, they probably don't. Always build in a way to check – don't let Big Data lead you to stupid conclusions.

Secure processing: Measures to secure the data within the analysis infrastructure are needed to mitigate potential vulnerabilities and to secure against leakage. These could include disk level encryption and a high level of network isolation. Big Data should be secured in transit, preferably by using encryption, but at minimum by using SSL/TLS. If the cloud is being used to process Big Data, you must know how to verify that it is secured.

Access management: Access to the analysis infrastructure, the data being analyzed, and the results should be subject to proper IAM controls.

Audit: Logging and monitoring of activities on the analysis infrastructure is crucial to proper auditing.

Mike Small (CEng, FBCS, CITP) is a fellow of the BCS, a member of the London Chapter of ISACA Security Advisory Group, and a senior analyst at KuppingerCole. Until 2009, Mike worked for CA where he developed CA's identity and access management product strategy. He is a frequent speaker at IT security events around EMEA.

Generating value from Big Data analytics

by Norman Marks

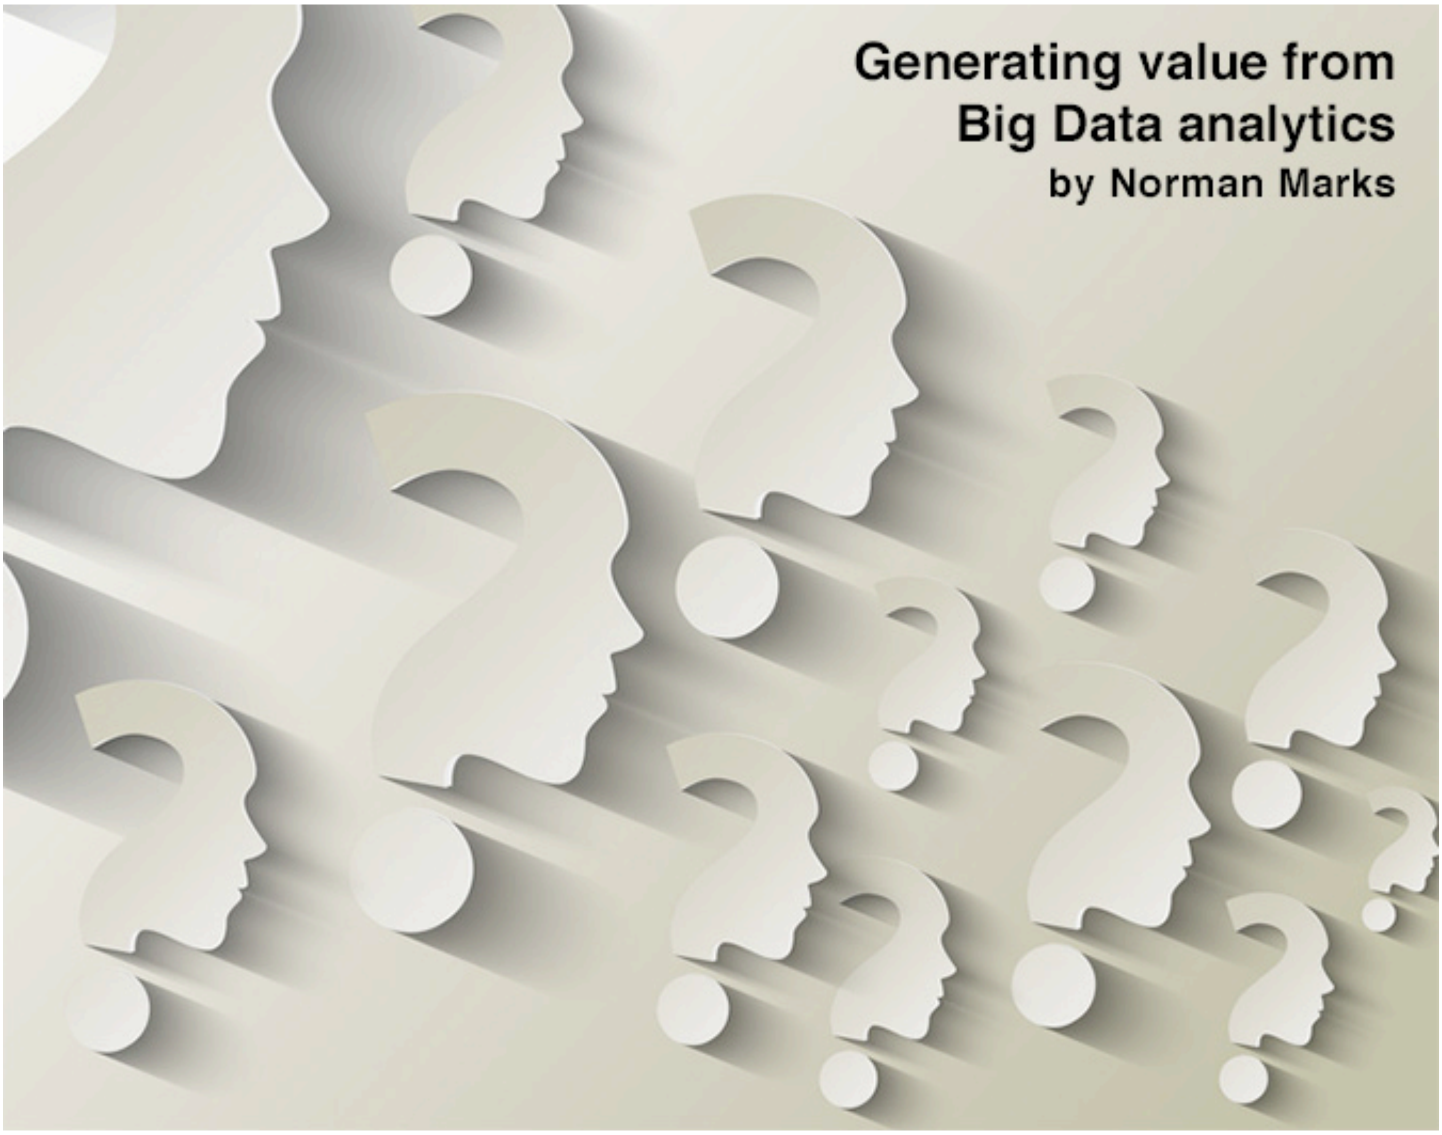

Traditional business intelligence has generally targeted structured data that can be easily parsed and analyzed, but advances in analytics methods now allow the examination of more varied data types. New analytics tools and methods are expanding the possibilities for how enterprises can derive value from existing data within their organizations and from freely available external information sources, such as Software as a Service (SaaS), social media and commercial data sources.

These advances allow enterprises to make better business decisions and increase competitive advantage. But this renaissance of analytics capability can also introduce additional technical and operational risk, so enterprises must weigh the technical and operational risk against the business risk that is associated with failure to adopt Big Data analytics.

As with any potential investment that is intended to bring about an improvement in efficacy or efficiency of business activities, several key elements must be well understood to enable systematic strategic planning:

- Anticipated returns and potential impacts to competitiveness through adoption
- Potential impact to the current operational ecosystem
- Opportunity cost for the investment (i.e., what else the enterprise might have invested in instead)
- Loss of value for investments already made.

Objective and systematic analysis of these factors becomes increasingly challenging as the industry hype that surrounds a new technology or business trend increases: hype can, in some cases, create unfounded pressure to adopt, or create barriers to adoption in others.

For technologists, the potential technical, operational and compliance risk that is associated with maintaining and operating on a large volume of potentially sensitive data is very apparent; however, the business-relevant factors that provided the initial impetus for adoption of these analytics tools and methods may be less apparent.

However, understanding the business case - the rationale for adoption, the anticipated return that the business hopes to achieve, and the competitiveness impact to the business if the enterprise chooses not to adopt while its competitors do adopt - is equally important.

For information security, audit and governance professionals, lack of clarity about the business case may stifle organizational success and lead to role and responsibility confusion.

Understanding the business case

Big Data refers to large, quickly growing or varied types of information ("high volume, high velocity, and/or high variety information assets that require new forms of processing to enable enhanced decision making, insight discovery and process optimization").

Big Data analytics is the application of emerging statistical, processing and analytics techniques for Big Data for the purpose of advancing the business— applying statistical models and techniques to business information to derive conclusions that are beneficial to that business.

Big Data analytics is particularly appealing to many enterprises because, in many cases, they have already made some investment in both business analytics and the collection of large data sets, on which analysis can be applied.

This means that the foundation from which to draw new conclusions, explore new ways of doing business and open up new avenues of competitive advantage may already be in place.

What is this competitive advantage specifically? Some data suggests a direct correlation between the use of Big Data analytics and profitability. For example, one study cites an increase in overall profitability of six percent as a direct result of using Big Data effectively. That metric, while appealing in the abstract, lacks sufficient underlying context and level of detail to be able to understand precisely how that correlation is made. Specific case studies, by contrast, provide a clearer picture of how increases in competitiveness are achieved and why these analytics techniques provide value. These studies show the imaginative uses to which pre-existing data are leveraged as a result of better analytics techniques, and the transformative impacts that are achieved.

What makes improvements for Big Data analytics particularly compelling from a business perspective is that the data already exists.

Data exists about customers, such as purchases they make and their receptiveness and responsiveness to marketing efforts. During the normal course of business, many enterprises collect large volumes of data about their customers—their habits, preferences, the specifics of individual transactions, fraud history, etc. When analyzed, this data allow enterprises to make changes and, subsequently, measure the performance of those changes. This allows those enterprises to dynamically shift inventory and / or pricing in response to consumer demand.

Furthermore, data analytics allows enterprises to create better-targeted marketing campaigns, to better measure the efficacy of those campaigns, and to launch new products and service offerings in response to customer demand. From a business standpoint, investment in Big Data analytics is compelling because it leverages the otherwise latent or unused resource of already-collected data.

Not every enterprise will be equipped to make use of Big Data analytic techniques. Some enterprises may be missing key skills in their existing personnel, or they may be missing critical portions of the technological ecosystem. Also, the technical ecosystem may not be laid out in a way that allows the techniques to operate; and enterprises may lack the processes to gain access to data and make use of the intelligence they collect as a result of the application of these methods.

It's important to find out the exact current situation, because investments made before readiness is fully achieved may be inefficient, suboptimal in terms of the results they produce or, as a worst case, may represent needless expense.

From a process standpoint, existing silos should be evaluated to determine whether individual business units, departments and personnel are willing and able to share information and act on information received. This consensus needs to happen so that analysis can be performed (disparate sources of data may need to be consolidated to operate on them) and so that the derived conclusions can be put to productive use. Enterprises need to consider that these areas may not share information currently and may have a history (depending on the culture) of competitiveness, antagonism, or resistance to outside influence.

These cultural barriers can impede open and collaborative exchange of important data elements and act as a barrier to adaptation in response to conclusions drawn. This consensus among silos can extend beyond the department level and down to the level of individual personnel. For example, key stakeholders may not know precisely where key data elements reside or how to access those data elements.

TECHNICAL AND OPERATIONAL RISK SHOULD CONSIDER THAT CERTAIN DATA ELEMENTS MAY BE GOVERNED BY REGULATORY OR CONTRACTUAL REQUIREMENTS

Likewise, "Shadow IT" (technology adopted without the direct oversight or, in many cases, awareness of the IT organization) can complicate information sharing because technology adopted without centralized oversight may represent a significant repository of critical information, and lack of central awareness of the information may limit the ability to include it in the scope of analysis.

Because of these factors, some degree of organizational self-awareness is required to think through supporting processes and identify potential problem areas before enterprises undertake significant investment in Big Data analytics.

Lastly, the technology implementation plays a role in determining organizational readiness. In many cases, new tools are required to support the analysis to be conducted, and capabilities for data storage and computation may need to be evaluated to ensure sufficiency. Moreover, sufficient data on which to operate need to exist and be accessible to analysts.

Data sources must be identified, which involves locating structured data (e.g., data organized in a relational database) and unstructured data (e.g., data stored ad hoc on a file system or in a loose collection.) Identifying data sources can likewise involve data in a variety of different formats, including video, audio, images and text. Computational resources may need to be expanded to enable operation and analysis of these data.

Information security and audit practitioners that are evaluating a Big Data analytics initiative in their enterprise need to weigh the management and mitigation of the technology risk of adoption against the business risk to the enterprise if they choose to not adopt Big Data analytics.

Technical and operational risk should consider that certain data elements may be governed by regulatory or contractual requirements and that data elements may need to be centralized in one place (or at least be accessible centrally) so that the data can be analyzed. In some cases, this centralization can compound technical risk. For example:

Amplified technical impact—If an unauthorized user were to gain access to centralized repositories, it puts the entirety of that data in jeopardy rather than just a subset of the data.

Privacy (data collection)—Analytics techniques can impact privacy; for example, individuals whose data is being analyzed may feel that the revealed information about them is overly intrusive.

Privacy (re-identification)—Likewise, when data is aggregated, semi-anonymous information or information that is not individually identifiable information might become non-anonymous or identifiable in the process.

These risk areas can cause some practitioners to be understandably wary. However, analytics efforts can be used to offset risk by applying the tools and techniques to security-event information, to transaction information for the purpose of detecting fraud, or to other technical information for risk-reduction purposes.

Stockpiles of security-relevant information, such as user and system activity, can be logged and examined the same way as more business-facing data can be logged and examined. The same analytics techniques and tools that streamline and increase the quality of business processes can likewise streamline and increase the quality of other risk-mitigation processes.

Tools purchased and analytics techniques that are acquired to help enable business-facing efforts can, with planning, be adopted by information security and risk management areas to help advance their goals as well.

The security and audit practitioners' consideration of risk can and should be holistic. If an enterprise elects not to employ these techniques, there is a risk to the business, because competitors will capitalize on the opportunity.

This result could have ramifications just as serious to the enterprise as a security or privacy breach or dreaded business continuity implications.

A holistic view of the risk in an enterprise should seek to account for all sides of the risk equation, including the following:

- Business value of adoption
- Business risk of non-adoption
- Technical/security/privacy risk that may increase depending on the implementation used to support the Big Data analytics approaches at the technical level
- Possible risk-offsetting benefits of the technology at the technical level

A number of business dynamics make the application of new and better analytics appealing to enterprises.

By looking at how these analytics techniques are transforming enterprises in real-world scenarios, the value becomes apparent as enterprises start to realize dramatic gains in the efficiency, efficacy and performance of mission-critical business processes.

The business case is made even more compelling by the fact that most enterprises already have in place the foundation for analysis in the form of more data than they can currently use productively.

Most enterprises already retain a large amount of data, such as information about their customers, metrics about the performance of internal business processes, data about information systems and their technology ecosystem, transactional information about sales and marketing and numerous other data items about how they do business.

While some new areas of technical risk may arise as a result of more voluminous and concentrated data, the business consequences of not adopting Big Data analytics may outweigh the technology risk.

Understanding this business case can help security, audit and governance practitioners in two ways: It helps them to understand the motivation and rationale driving their business partners who want to apply Big Data analytics techniques within their enterprises, and it helps balance the risk equation so that technical risk and business risk are addressed.

Norman Marks is a member of ISACA's Emerging Business and Technology Committee.

Too big to fail: The Big Data dilemma

by Prakash Panjwani

As data grows exponentially, and with it the ability to use it effectively, how can organizations ensure that the analysis of data sets containing sensitive information doesn't result in a costly data breach? If an organization can access this data and analyze it, then who else can access it?

Organizations are struggling to understand how the collection of extremely large and complex structured and unstructured data sets can be protected. And it is clear some traditional approaches to database and application security are not well suited to Big Data deployments.

Imagine a supercomputer available in seconds thanks to cloud computing where many data sources are merged for analysis from across an enterprise. It is easy to see how sensitive information could be exposed or created in such a scenario.

There is also a risk that human error could combine the data or allow access by unauthorized users. It's fair to say that by its very makeup – the aggregation of multiple information sources – Big Data is sensitive information. Stakeholders expect this information to be secure in the hands of organizations that use it. So, if Big Data is so large and complex, how can it be protected?

The answer, as Google executive chairman Eric Schmidt recently said, is to "encrypt everything." Admittedly, Schmidt was responding to the issue of government surveillance, but he is nonetheless right. Encryption does protect your data, big or small, when a breach occurs. Encryption ensures that your data remains secure, regardless of where it resides – in the data center, the cloud, a mobile device or even in the hands of hackers from enemy organizations, nation states or malicious insiders.

The caveat here is that the encryption algorithms need to be secured and the keys to decipher the encryption need to be protected. In addition, Big Data processes expect data in real world formats so a complete data protection strategy has to include format-preserving approaches such as tokenization.

Organizations must first understand what to protect (encrypt or tokenize) and then must learn how to manage the encryption keys used for this protection. Specifically, how are the keys generated? Where are they stored? What is the strength (or size) of keys? How often are keys changed? These are just some of the questions that will help you build a strong encryption strategy for protecting data.

A key management solution needs to allow a comprehensive data protection strategy. Segregating encryption and tokenization technologies and managing them individually is a recipe for failure.

A security strategy is only as strong as its weakest link. Solutions must also protect data efficiently. This means that the impact of the encryption on business and cost must be minimal. The solution should, preferably, be transparent, especially to the end users who are less likely to be concerned with security and compliance issues.

Most users simply want to get their job done quickly and without "technical" issues. So as part of the protection approach and best practices, authentication and access controls need to be seamlessly integrated to ensure that only the correct users are accessing the right information.

WHILE EVERYTHING IS BIGGER AND FASTER WITH BIG DATA, ULTIMATELY ENCRYPTION AND TOKENIZATION REMAIN CRITICAL ELEMENTS OF PROTECTING SENSITIVE DATA

At the forefront of Big Data protection is the Cloud Security Alliance. The group is specifically working to address the security and privacy issues magnified today by the velocity, volume, and variety of Big Data, such as large-scale cloud infrastructures, diversity of data sources and formats, streaming nature of data acquisition and high volume inter-cloud migration.

The CSA draws attention to this issue when it states: "Securing Big Data stores: this document focused on using Big Data for security, but the other side of the coin is the security of Big Data."

The CSA highlights the way Big Data differs from regular data: Big Data is differentiated from traditional technologies in three ways: the amount of data (volume), the rate of data generation and transmission (velocity), and the types of structured and unstructured data (variety) (Laney, 2001).

While everything is bigger and faster with Big Data, ultimately encryption and tokenization remain critical elements of protecting sensitive data, whether in transit across high speed networks, stored encrypted in large volumes of data at rest, or tokenized during processing.

Confidential data is everywhere in every organization. Therefore a comprehensive and holistic data protection strategy includes encryption, tokenization, authentication and access management, with the understanding that key management is critical.

With great power comes great responsibility, and that power must be used wisely to ensure that the analytics are put to good use, that the data will remain in the control of those people who will do that, and that it's protected from those who may have other intentions.

Prakash Panjwani is Senior Vice President & General Manager, Data Protection Solutions at SafeNet (www.safenet-inc.com).

Events around the world

RSA Conference 2014

www.rsaconference.com/helpnet

Moscone Center, San Francisco, CA, USA

24 February - 28 February 2014

InfoSec World Conference & Expo 2014

www.infosec-world.com

Disney's Contemporary Resort, FL, USA

7 April - 9 April 2014

Infosecurity Europe 2014

www.infosec.co.uk

Earls Court, London, UK

29 April - 1 May 2014

HITBSecConf2014 Amsterdam & HITB Haxpo

conference.hitb.org

De Beurs van Berlage, Amsterdam, The Netherlands

29 May - 30 May 2014

SECURITY AWARENESS

FOR THE 21st CENTURY

- Go beyond compliance and focus on changing behaviors.
- Create your own program by choosing from 30 different training modules.
- Meets requirements of the Data Protection Act and PCI DSS.
- Training is mapped against the 20 Critical Control framework.
- For more information visit us at www.securingthehuman.eu

SANS SECURING THE HUMAN

www.securingthehuman.eu