ELSEVIER

network SECURITY

ISSN 1353-4858 December 2016

www.networksecuritynewsletter.com

Featured in this issue:

Addressing the security challenges of using containers

Containers have taken the open source ecosystem by storm, providing a fast and effective way to deliver and deploy a compact, portable environment for running applications across shared server resources.

Container technology is another breakthrough in the constant drive to increase development agility and get products to market more quickly. But that can't be at the expense of security. Enterprises need to take realistic precautions before completely embracing this approach to operating system-level virtualisation, explains Mike Pittenger of Black Duck Software.

Full story on page 5…

Opportunity vs risk with the Internet of Things

The Internet of Things (IoT) is no longer a new concept: it is beginning to appear in simple forms in our everyday lives – from smartphones to connected buildings and vehicles.

The IoT is promising a whole new market of opportunities but there is growing awareness of its potential challenges. There must be an increasing focus on its security risks. Until we have a clearer understanding of these risks, we must rely on programmed testing to develop knowledge and to support overall security, says Sameer Dixit of Spirent.

Full story on page 8…

Data classification: keeping track of your most precious asset

Many companies are now in the data business, whether they know it or not. Popular clichés about data turn on words such as 'lifeblood' or 'crown jewels'.

However, as we examine in this interview with Rui Melo Biscaia of Watchful Software, too many organisations fail to understand the true nature of their data – indeed, many don't know what data they have or where it is. But with data protection regulations starting to bite, it's time for organisations to get to grips with classifying and protecting their information.

Full story on page 10…

Major ISPs targeted in Internet of Things botnet attacks

Routers belonging to the customers of major Internet service providers (ISPs) have been attacked in what appears to be an attempt to use them to create botnets. This has resulted in thousands of customers temporarily losing Internet access.

Most reports suggest that these attacks are connected to the Mirai malware. The Mirai source code was released in a cybercrime forum recently, since when it has been used to build large botnets created from compromised Internet of Things (IoT) devices such as digital video recorders, CCTV systems and so on. These botnets were used to mount massive distributed denial of service

Continued on page 2...

Contents

Come and visit us at
www.networksecuritynewsletter.com

ISSN 1353-4858/16 © 2016 Elsevier Ltd. All rights reserved

This journal and the individual contributions contained in it are protected under copyright by Elsevier Ltd, and the following terms and conditions apply to their use:

Photocopying

Single photocopies of single articles may be made for personal use as allowed by national copyright laws. Permission of the publisher and payment of a fee is required for all other photocopying, including multiple or systematic copying, copying for advertising or promotional purposes, resale, and all forms of document delivery. Special rates are available for educational institutions that wish to make photocopies for non-profit educational classroom use.

Editorial Office:
Elsevier Ltd
The Boulevard, Langford Lane, Kidlington,
Oxford, OX5 1GB, United Kingdom
Tel: +44 1865 843239
Web: www.networksecuritynewsletter.com

Publishing Director: Bethan Keall

Editor: Steve Mansfield-Devine
E-mail: smd@contrarisk.com

Senior Editor: Sarah Gordon

Columnists: Karen Renaud, Colin Tankard

International Editoral Advisory Board:
Dario Forte, Edward Amoroso, AT&T Bell Laboratories; Fred Cohen, Fred Cohen & Associates; Jon David, The Fortress; Bill Hancock, Exodus Communications; Ken Lindup, Consultant at Cylink; Dennis Longley, Queensland University of Technology; Tim Myers, Novell; Tom Mulhall; Padget Petterson, Martin Marietta; Eugene Schultz, Hightower; Eugene Spafford, Purdue University; Winn Schwartau, Inter.Pact

Production Support Manager: Lin Lucas
E-mail: l.lucas@elsevier.com

Subscription Information
An annual subscription to Network Security includes 12 issues and online access for up to 5 users.
Prices:
€1424 for all European countries & Iran
US$1594 for all countries except Europe and Japan
¥189 000 for Japan
Subscriptions run for 12 months, from the date payment is received.
More information:
http://store.elsevier.com/product.jsp?isbn=13534858

Permissions may be sought directly from Elsevier Global Rights Department, PO Box 800, Oxford OX5 1DX, UK; phone: +44 1865 843830, fax: +44 1865 853333, email: permissions@elsevier.com. You may also contact Global Rights directly through Elsevier's home page (www.elsevier.com), selecting first 'Support & contact', then 'Copyright & permission'. In the USA, users may clear permissions and make payments through the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA; phone: +1 978 750 8400, fax: +1 978 750 4744, and in the UK through the Copyright Licensing Agency Rapid Clearance Service (CLARCS), 90 Tottenham Court Road, London W1P 0LP, UK; tel: +44 (0)20 7631 5555; fax: +44 (0)20 7631 5500. Other countries may have a local reprographic rights agency for payments.

Derivative Works
Subscribers may reproduce tables of contents or prepare lists of articles including abstracts for internal circulation within their institutions. Permission of the Publisher is required for resale or distribution outside the institution. Permission of the Publisher is required for all other derivative works, including compilations and translations.

Electronic Storage or Usage
Permission of the Publisher is required to store or use electronically any material contained in this journal, including any article or part of an article. Except as outlined above, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the Publisher. Address permissions requests to: Elsevier Science Global Rights Department, at the mail, fax and email addresses noted above.

Notice
No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein. Because of rapid advances in the medical sciences, in particular, independent verification of diagnoses and drug dosages should be made. Although all advertising material is expected to conform to ethical (medical) standards, inclusion in this publication does not constitute a guarantee or endorsement of the quality or value of such product or of the claims made of it by its manufacturer.

12987

Pre-press/Printed by
Mayfield Press (Oxford) Limited

...Continued from front page

(DDoS) attacks on targets such as the blog of security journalist Brian Krebs and DNS service provider Dyn.

New versions of Mirai quickly appeared as attackers modified the code to target additional devices, which could include home routers. However, a few security specialists have pointed out that the security of home routers has been a known problem for some time. So whether Mirai is directly involved is uncertain.

Deutsche Telekom in Germany said that the routers of around 900,000 of its 20 million customers had come under attack. The German Federal Office for Information Security (BSI) also issued a statement saying that government systems had been targeted too. The attacks, against the management interfaces on port 7547 of Speedport routers, were mostly unsuccessful. However, it left around 4% of Deutsche Telekom customers – those using specific Speedport models – unable to connect to the Internet (other users could reconnect by simply rebooting the router).

The company has since issued patches for affected routers and is giving inconvenienced customers a free 'day pass'. It is also now filtering traffic to prevent a similar attack in the future. In a statement, Deutsche Telekom said that it believed the same attack was being mounted against routers worldwide. The statement is here (in English): http://bit.ly/2hggiFU.

The Speedport routers are made by Arcadyan Technology in Taiwan. Deutsche Telekom said it is now reviewing its relationship with that company.

A few days after the Deutsche Telekom attack, routers belonging to customers of KCOM, TalkTalk and the Post Office in the UK came under a similar assault. Zyxel AMG1302 routers supplied to Post Office customers and D-Link DSL-3780 routers used by TalkTalk and KCOM were among the devices that were affected. Again, the attack targeted the maintenance interfaces on port 7547. The Post Office confirmed that around 100,000 of its customers had their Internet service disrupted.

Finally, a researcher going by the name 'kenzo' has published a proof-of-concept exploit that is capable of taking over the ZyXEL-built Eir D1000 routers distributed by Irish ISP Eir. As with the attacks on the other ISPs, this uses TR-064 commands on port 7547. Tens of thousands of customers are said to be vulnerable to the attack. There are more details here: http://bit.ly/2ga8Ufg.

Ransomware claims more victims

San Francisco's Municipal Railway (Muni) has become the latest high-profile victim of ransomware. It was a very public attack as the usual notification screen that tells the victim that the machine has been infected and demands the ransom ended up being shown on computer displays at Muni stations.

The attackers, who employed the HDDCryptor malware, demanded 100 bitcoins (around $73,000) for the decryption key. The San Francisco Municipal Transport Agency (SFMTA) confirmed that a number of payment systems were taken offline. Around 900 computers seem to have been infected, including office desktops, CAD workstations, email and print servers, laptops, payroll systems, SQL databases, lost and found property terminals and station kiosk PCs.

Muni was forced to open a number of fare gates on its system, effectively allowing many people to ride for free. The company estimated that it will have lost $50,000 in fares during the weekend in which its systems were down.

Someone who claimed responsibility for the attack also said that he had stolen 30GB of data, including databases and documents, belonging to SFMTA and would dump it on the Internet if the ransom wasn't paid. An SFMTA spokesperson denied the claim, adding that the organisation had no intention of paying.

Security journalist Brian Krebs reported that the email address used by the attacker had itself been hacked. An anonymous researcher simply guessed the security question for the cryptom27@yandex.com account and took it over. There's more information here: http://bit.ly/2gC8bj1.

Ransomware has also started spreading via Facebook and LinkedIn. According to Check Point, the Locky strain is being distributed through the use of SVG image files. When victims try to view them, they are prompted to download an image 'codec', which is actually the malware.

In brief

Uber tracks users

It has been revealed that Uber is routinely tracking the movements of its customers even when they aren't actively using the firm's app. Location information is recorded and uploaded to the firm's servers even when the app is put into background mode and the user is not ordering or undertaking rides using the service. In 2015, the Electronic Privacy Information Centre (Epic), filed a complaint against Uber with the US Federal Trade Commission (FTC) after Uber announced plans to carry out such tracking. However, the FTC failed to take any action and so now Uber has switched on the tracking system.

Dailymotion breached…

According to the data breach notification website LeakedSource, the French video sharing site Dailymotion has been breached, with the exposure of 85 million usernames and email addresses. The data haul also includes 18 million hashed passwords. The hack appears to have taken place on 20 October, but it's not sure when it was discovered.

…and adult sites exposed

The FriendFinder Network, which includes pornography sites and the Adult FriendFinder 'dating' site has had account details for over 412 million users stolen. The information includes email addresses and passwords, with nearly all of the latter being stored in plaintext or an easily crackable form, according to LeakedSource. The sites affected were AdultFriendFinder (339 million user records), Cams.com (62 million), Penthouse.com (7 million), Stripshow.com (1 million) and iCams.com (1 million). As with Dailymotion (above), the breach occurred in October, although it's not known if the attacks are linked. It appears that, as with the infamous Ashley Madison breach, the FriendFinder Network continued to store details for users who had requested that their accounts be closed. Around 15 million of the breached accounts fall into this category. ZDNet reported that the attack was performed by exploiting a local file inclusion flaw. Users of the sites were informed of the breach only when they logged in. Meanwhile, cyber-criminals are trading at least 380,000 user accounts belonging to the xHamster porn site, including usernames, email addresses and poorly hashed passwords (using only MD5 with no salt).

Three reveals breach details

A data breach of UK mobile operator Three has resulted in the exposure of 130,000 customer records. The National Crime Agency (NCA) said it had arrested three men – two under the Computer Misuse Act and one for perverting the course of justice. It's alleged that they used unauthorised access to Three's systems in order to find customers who were eligible for phone upgrades. They then ordered the devices and intercepted the deliveries in order to sell the phones. A spokesperson for Three said that eight handsets had been illegally obtained this way. Although customer records were accessed, Three said that no bank information, passwords or other sensitive information was compromised. This comes at a time when the company has also suffered a spate of burglaries from retail stores resulting in the loss of 400 high-value phones.

Google malware breaches 1 million accounts

More than a million Android users have suffered Google account breaches as a result of malware. The 'Gooligan' malware – a variant of Ghost Push which was first seen in 2014 – has been found in at least 86 apps available in a number of third-party marketplaces. Most victims have been infected by being enticed to download what appear to be free versions of popular apps. The malware then steals account credentials for a number of Google services, including Gmail, Docs, Drive, Play and Photos. It's also capable of downloading modules that inject ads, generating revenue for the malware authors. Security firm Check Point believes that around three-quarters of Android handsets on the market today are vulnerable to the malware. More than half (57%) of the affected handsets are in Asia, with 20% in the Americas, 15% in Africa and 9% in Europe.

World gets a C- for security

Tenable Network Security has released its annual 'Global Cyber-security Assurance Report Card' which takes input from 700 information security practitioners to rate the cyber-security readiness of countries. On average, the world scored a C-, with an overall score of 70% – a drop of 6% compared with last year. This was due, says Tenable, to a 12-point drop in the 2017 Risk Assessment Index, which measured the ability of respondents to assess cyber-risk across 11 key components of the enterprise IT landscape. Among the key findings were a fairly dim view of cloud security, with software as a service (SaaS) and infrastructure as a service (IaaS) being two of the lowest-scoring risk assessment areas. Mobile security also showed up as one of the biggest enterprise security weaknesses. Nor were respondents convinced of firms' ability to assess security during the DevOps process. In general, India got the best score (84%), followed by the US (78%), Canada (75%) and France (74%), with the UK coming sixth (66%). The best-scoring industry was retail (76%) followed by financial services (72%). There's more information here: http://bit.ly/2hgWVwq.

Turks build attack platform

A Turkish group has created a distributed denial of service (DDoS) platform in which participants can earn rewards by taking part in attacks. The 'Surface Defense' platform is being promoted via dark web forums such as Turkhackteam and Root Developer, according to Forcepoint Security Labs, which discovered the scheme. Forcepoint says that the Surface Defense operators are attempting to 'gamify' DDoS operations by offering rewards such as hacking tools and click-fraud software. To earn points, participants need to download a stresser tool known as Sledgehammer that performs HTTP-based slow-loris-type denial of service attacks. The tool comes preconfigured with 24 targets chosen by the Surface Defense team. Forcepoint says it believes the hackers are mostly focusing on Turkish recruits sympathetic to nationalist ideologies. DDoS attacks launched from the platform have targeted groups such as the Kurdistan Workers Party, the German Christian Democratic Party and the Armenian National Institute website in Washington DC. There's more information here: http://bit.ly/2gaCGAH.

ICO criticises charities

The RSPCA and the British Heart Foundation (BHF), both leading UK charities, have been accused of serious privacy breaches because of their alleged misuse of donors' data. According to an investigation by the Information Commissioner's Office (ICO), both organisations conducted 'wealth screening', employing the services of wealth management companies to assess the financial status of donors to estimate how much more money they might be pressured to give. The charities also pieced together personal information about new and lapsed donors from other sources so that they could be better targeted. And they traded personal information with other charities. Donors were not aware of these activities and so could not give consent. The RSPCA was fined £25,000 and the BHF £18,000, with Elizabeth Denham, the Information Commissioner, saying that the amount of the fines was reduced to avoid causing distress to donors.

Hungry hackers hijack Deliveroo

Hackers have been hijacking accounts at online take-away firm Deliveroo and ordering food deliveries, the BBC's Watchdog programme has revealed. The scam appears to have affected people who have reused passwords, with their Deliveroo credentials matching those in user databases leaked from other sites. In several cases, the hackers changed the phone numbers and addresses of Deliveroo customers before running up large bills. Deliveroo said it has reimbursed affected customers.

Reviews

BOOK REVIEW

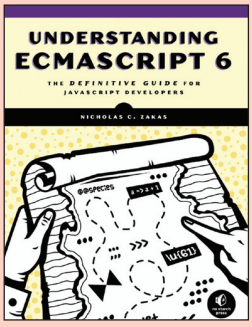

Understanding ECMAScript 6
Nicholas C Zakas.
Published by No Starch Press.
ISBN: 978-1-59327-757-4.
Price: $34.95, 352pgs, paperback and e-book editions available.

JavaScript is everywhere now. Few websites would function without it and with developments such as Node.js it has even become a widely used and powerful server-oriented programming language. And at the heart of JavaScript is ECMAScript, which has recently undergone a major upgrade.

The history of JavaScript is not without its issues and controversies. Almost from the beginning, competing standards and libraries emerged, some of which caused severe headaches for developers wanting to create websites that worked in every browser. Many of those issues have been resolved, yet different flavours of JavaScript and its derivatives remain.

However, at the core of all of them is ECMAScript, the standardised core of the language. In 2015, the latest iteration – formally known as ECMAScript 2015 but more commonly dubbed ECMAScript 6 – became feature-complete. It included new objects and patterns and some syntax changes. Not all of these features have been implemented in current versions of JavaScript and Node.js, but any developers looking to ensure the maximum portability of their code would do well to work to ECMAScript 6 standards.

Familiarity with JavaScript in all its guises is an important skill for information security practitioners who get involved at the sharp end of malware, web exploits and the like. Many web application vulnerabilities are to be found either in the poor use of JavaScript or through its employment as a hacking tool.

If you're completely new to JavaScript you'll need to look elsewhere for an introductory text. This book is intended to bring JavaScript users up to date with the latest features and capabilities enshrined in ECMAScript 6. Key topics covered include: iterators and generators; the specific features of arrow functions; storing data with sets, maps and other types; the inheritance capabilities of the language's objects; asynchronous programming; and how to best use modules.

This book, then, is for existing JavaScript coders who want to keep their skills sharp, their code standards-compliant and who feel they can make use of the latest functionality of ECMAScript 6.

There's more information available here: www.nostarch.com/ecmascript6.

– SM-D

BOOK REVIEW

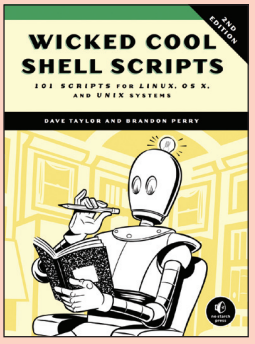

Wicked Cool Shell Scripts
Dave Taylor and Brandon Perry.
Second edition published by No Starch Press. ISBN: 978-1-59327-602-7.
Price: $34.95, 392pgs, paperback and e-book editions available.

For Unix admins, hacking out a quick shell script to automate a repetitive task or solve a problem is as natural as breathing. Indeed, sysadmins are quite open about their readiness to judge each other based on their scripting skills.

Shell scripts aren't just for the day-to-day task of running a Unix system, though. A script is essentially just a way of saving yourself a lot of typing. Whenever you have a combination of commands that you use frequently, it makes sense to gather them in a script. And this is a situation that will be familiar to any kind of hacker – although for the sake of propriety we'll assume we're talking about penetration testers and other white-hat hackers here.

A script might be something as simple as a single line that feeds some piece of input (perhaps the output of a command such as 'ifconfig') through filters (awk, sed and the like) to extract a single item of information (an IP or MAC address would be typical in the case of ifconfig). Or it might be many lines of code, performing multiple OS commands and complex manipulation of data. Some of us would argue, though, that if your shell script is longer than 20 or 30 lines you really should invest in a book on Python or Perl.

This book focuses on Bash as its shell of choice, which alone is probably enough to start a flame war. But it makes a lot of sense because Bash is arguably the most flexible and powerful shell and also the most commonly found. It's even available on Windows now.

"Bash is arguably the most flexible and powerful shell and also the most commonly found. It's even available on Windows now"

This is not a Bash primer (you'll find many excellent introductions freely available on the web). Nor is it an exhaustive guide to all of Bash's commands and capabilities. It's meant as a practical guide to performing specific tasks with scripts. In many cases you're left to work out for yourself what a particular command is doing and how it's doing it. You're likely, then, to want to resort to man pages and online guides if you want to fully understand the functioning of the scripts this book describes, or if you want to adapt or expand them to your own needs, as you almost certainly will.

The 101 scripts detailed in the book include the expected admin tasks, such as analysing disk usage, adding users, killing processes and so on. There are also net- and web-specific tasks, such as pulling URLs from a web page, reporting broken links or tracking changes to a page. There's a bunch of server admin stuff, some image manipulation techniques and some fun scripts – even games. To make all this work, the authors provide a library of essential functions that can be reused in your own scripts too.

The scripts are intended to be POSIX-compliant and so will run on pretty much all modern *nix systems. But as a bonus, there are some written specially for OS X. None of the scripts is aimed at security applications, but many could be adapted. And if you're just getting to grips with Bash scripting, following these examples will give you a great insight into how it all works.

There's more information available here: www.nostarch.com/wcss2.

– SM-D

Addressing the security challenges of using containers

Mike Pittenger

Mike Pittenger, Black Duck Software

Containers have taken the open source ecosystem by storm, providing a fast and effective way to deliver and deploy a compact, portable environment for running applications across shared server resources. Container technology is another breakthrough in the constant drive to increase development agility and get products to market more quickly. Speed and agility are key drivers for container adoption in the enterprise, but those drivers can't be at the expense of security. Enterprises need to take realistic precautions before completely embracing this approach to operating-system-level virtualisation.

The security risks that are unique to container technology can be exploited in a number of ways. Breaches can quickly spread from the shared server kernel to a collection of containers residing on a server. Untrusted code can take advantage of vulnerabilities in a system and – in the worst-case scenario – expose private, sensitive information to hackers. Enterprises that are subject to regulatory mandates and strict accountability cannot afford security liabilities and the potential penalties, including financial losses and loss of customer trust, that arise from breaches, unleashed malware and attacks on data integrity.

This article explores the security concerns around container use and possible avenues for exploitation and also examines how security professionals can enact measures to mitigate the dangers and protect code integrity when containers are in use.

A brief history

Containers decouple applications from operating systems, giving a clean and minimal Linux operating system with everything else run in one or more isolated containers. Because the operating system is separate from containers, you can move a container across any Linux server that supports the container runtime environment.

The technology behind containers has a long history, beginning with SELinux in 2000 and Solaris zones in 2005. The original container was LXC (also known as Linux Containers) a Linux operating system-level virtualisation method for running multiple isolated Linux systems on a single host.

The widespread adoption of containers is largely due to the development of standards aimed at making them easier to use, such as the Docker image format and distribution model.[1] Docker, which started as a project to build single-application LXC containers, makes containers even more portable and flexible for any infrastructure capable of running containers.

Docker containers comprise a file system, network stack, process space and anything else needed to run an application, such as system tools and system libraries. Each Docker container includes the designated application and its dependencies, which vary from application to application, but are identical across different copies of the same container. Docker containers are not tied to any specific infrastructure. They run on any computer, on any infrastructure and in any cloud.

"Docker and application container technologies like it are the next step on the journey from physical, single-tenanted computing resources to more efficient, virtual, multi-tenanted infrastructure"

At the core of the Docker platform is Docker Engine, a software layer that builds and runs Docker containers to create the operating environment for

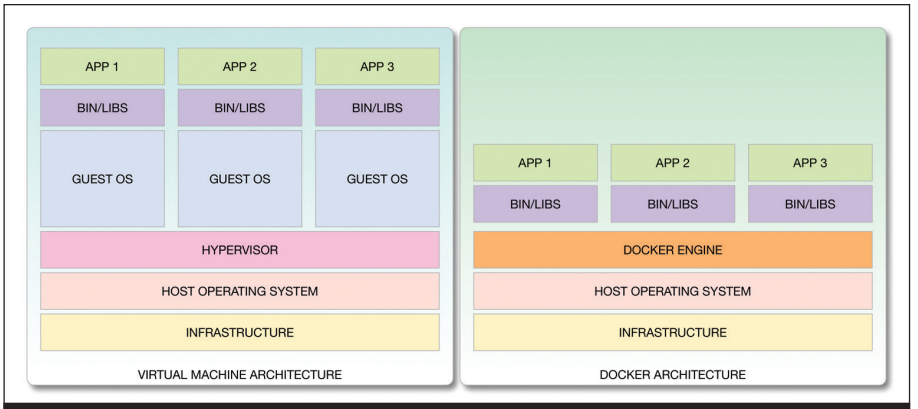

The Docker container architecture compared to conventional virtual machines.

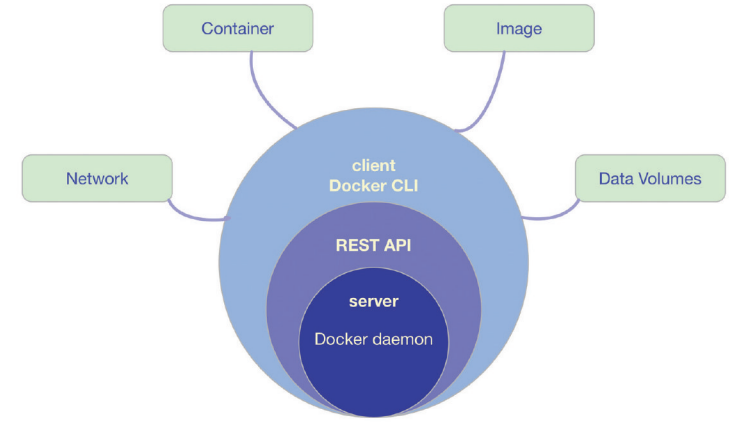

The Docker engine.

distributed applications. The in-host daemon communicates with the Docker client to execute commands to build, ship and run containers.

Docker and application container technologies similar to it are the next step on the journey from physical, single-tenanted computing resources to more efficient, virtual, multi-tenanted infrastructure that can run in traditional IT environments and in the cloud. Among its other benefits, Docker is also ideally suited to so-called CI/CD – or Continuous Integration/Continuous Delivery – environments, which seek to accelerate development practices and streamline the path between development and production environments.

Docker allows software publishers to realise substantial efficiencies over traditional IT and even other virtualisation technologies. A typical server can run a thousand or more Docker containers at native speeds. Application processes within Docker environments run directly on the host, but are kept isolated from other processes. CPU and memory, network and disk I/O performance within a Docker container are virtually identical to what a developer would see running the application in a native environment.

The security challenge

As an editorial in *Container Journal* noted, security always seems an afterthought when it comes to any new technology, with container use being no exception.[2] Enterprises jumping on the container bandwagon require up-front planning, vigilance and an effective solution designed to reduce risk.

One of the advantages of container technology is the compact nature of the deployment model, which is achieved by mutually sharing key resources on the server, including the operating system and – in many instances – bins and libraries. This results in a substantially smaller footprint for deploying tested applications, but also makes it extremely important that none of the shared resources be contaminated by rogue code. Both virtual machines and containers are isolated constructs, but containers share key resources.

Many of the key container technologies, by design, require root access to operate. For example, Docker states in its security documentation that when applications and containers are used with Docker, it implies running the Docker daemon, which itself requires root privileges. Even restricting Docker use to 'trusted users' only, there are potential avenues that can be exploited just using the Docker daemon or taking advantage of applications built on the Representational State Transfer (REST) model that can open potential vulnerabilities to hackers. Ultimately, the security risks come down to gaining assurance that the software in the container is clean and does not feature malicious code.

"There are potential avenues that can be exploited just using the Docker daemon or taking advantage of applications built on the REST model that can open potential vulnerabilities to hackers"

The presumed security flaws of containers are often a matter of how you set and administer the policies that affect the interactions between container and host. When you containerise any application, it still has to rely on operating system services. The OS is often running in a host virtualisation environment, with the app in the guest's container. That boundary, which allows the guest container to call into the host OS environment, is where you get security problems if the policies are not sourced out correctly.

Reducing risk

For all their benefits, containers represent a new layer in the application stack that can serve as a source of exploitable vulnerabilities and risk for the application owner and hosting firm alike. The ease and speed with which containers

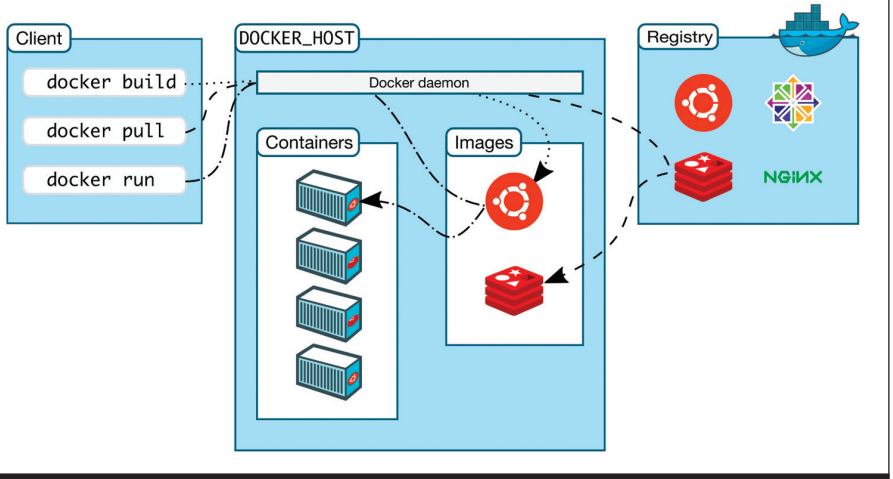

The Docker architecture. Source: Docker.com.

can be configured and launched can amplify mistakes, making it difficult to track and manage deployed applications within a dynamic IT environment.

Visibility into open source code is vital to security and you should monitor and thoroughly scan every container under development before deployment. Container technology can deliver a great deal of value, but you can only capitalise on that value if security demands are adequately met. To leverage the benefits that application container technologies offer, your organisation must understand the possible risks. Specifically, your application container deployment cannot proceed at the cost of security and visibility.

Certification and provenance

With the increased use of open source code in enterprise-scale applications, the repositories that serve to provide code to development teams building container images are only secure if they can be trusted. Files and images from repositories need to be treated with the same caution given to random downloads from the Internet: you should treat each instance and every download as though someone could have embedded malicious ingredients into the code.

"Verifying the publisher of a container doesn't guarantee that the software application and supporting files within the downloaded container don't have flaws or exploitable vulnerabilities"

Without supporting systems in place vetting the contents of container images, compromised or malicious container images that are offered via a repository like Docker Hub might be distributed to unsuspecting organisations. Registry servers such as Docker Hub may offer administrators features that help to mitigate the risk of dodgy containers. For example, administrators may be able to leverage features within the registry to limit the types of container images they will allow into their network.

The key risks associated with open source software.

Docker has taken steps to provide additional layers of accountability, such as Docker Content Trust which uses public key cryptography to allow publishers to 'sign' Docker containers and vouch for the integrity of the code they contain. Aligned projects like Notary and The Update Framework (TUF) promise the same capabilities for non-trusted actors who wish to publish Docker images.

Vulnerabilities within images

Vouching for the provenance of containers is necessary but not sufficient to secure their deployments. Verifying the publisher of a container doesn't guarantee that the software application and supporting files within the downloaded container don't have flaws or exploitable vulnerabilities. Containers may well bundle outdated and insecure components, especially when the underlying operating system is not the most current version.

Privately funded research suggests that security flaws in Docker images are not uncommon. A survey of images hosted on Docker Hub, a central repository for Docker developers to pull and push container images, found that more than 30% of official repositories contained images that were 'highly susceptible' to attacks targeting known vulnerabilities such as Shellshock, HeartBleed and Poodle. Some 40% of general images on Docker (images not explicitly verified by any authority) were found to contain known and exploitable security flaws.

In other words: organisations that wish to leverage Docker or any container must be able to both 'trust' and 'verify'. That is, they need to establish the bona fides of the publisher of the container they wish to use and verify that the contents of that container won't introduce serious and exploitable security vulnerabilities into their environment that could put the enterprise at risk.

Container management

Even when adequate precautions are put into verifying the provenance and security of containers at the time they are deployed, organisations must maintain vigilance of their deployed, containerised applications.

Like any other applications, applications deployed within containers age. In the process, they may become vulnerable to newly discovered security vulnerabilities or other risks. Applications deployed within containers may also contain data or configuration settings that are insecure or out of sync with your current applications or network environment. As one example, a prominent crowdfunding site pushed a container for its web-based funding platform into production with a development debugger enabled. That gave attackers a ready avenue to run malicious code on the vulnerable system.

Finally, the ease with which application containers can be assembled from different technology layers and deployed means that vulnerabilities, misconfigurations or flaws in any one layer can easily be reproduced across multiple applications.

End-to-end solution

Developing effective strategies for securing a system running containers requires that you rethink traditional virtualisation methods. The isolation layer that separates the kernel from the container is less robust than the isolation provided by a virtual machine, but very manageable if security provisions are implemented consistently. The term 'container' suggests that the contents of the image are confined and separated from other shared resources, but the primary concern is the integrity of the contents of the runtime environment residing within the container. The level of security depends largely upon whether the container content has been scanned or vetted or if it consists of unidentified code from unspecified sources. Hidden vulnerabilities can reside within container images created from repositories that have not been fully checked and analysed.

To minimise this type of security flaw, an end-to-end solution that can vet components of applications and container images from the earliest stages of deployment through deployment and ongoing future maintenance is required. Container contents should be screened and analysed component by component, including all third-party applications as well as software developed in-house from open source components. Ongoing monitoring will ensure that policies set up to control open source use within the enterprise can be effectively enforced. In scenarios where older open source software needs to be kept in use, this solution can offer insights into any existing vulnerabilities.

The adoption and success of container-based application delivery requires reliable mechanisms to efficiently develop, deploy, operate and manage these applications in a secure and consistent manner, making it possible to have a strongly supported development and deployment platform with built-in management features that follow each application's lifecycle.

About the author

Mike Pittenger is vice-president of security strategy at Black Duck Software (www.blackducksoftware.com). At Black Duck, he is responsible for strategic leadership of security solutions, including product direction and strategic alliances. Black Duck helps organisations remediate open source security risks in the applications they build by finding the open source components in use, mapping those components to current known security vulnerabilities and alerting them when new vulnerabilities are identified.

References

1. 'Docker overview'. Docker. Accessed Nov 2016. https://docs.docker.com/engine/understanding-docker/.
2. Tozzi, Christopher. 'The state of container security today'. Container Journal, 19 Sep 2016. Accessed Nov 2016. http://containerjournal.com/2016/09/19/state-container-security-today/.

Opportunity vs risk with the Internet of Things


Sameer Dixit

Sameer Dixit, Spirent

The Internet of Things (IoT) is promising a whole new market of opportunities and there is growing awareness of its potential challenges. Until we have a clearer understanding of these risks, we must rely on programmed testing to develop knowledge and to support overall security.

Burning issue

The first burning issue with any new technology is always, will it work? Then comes a wave of excitement about the new opportunities it offers. Only then do we seriously address the possibilities for misuse or exploitation of the innovation.

The IoT is no longer a new concept: it is beginning to appear in simple forms in our everyday lives – from smartphones to connected buildings and vehicles. As IoT becomes mainstream, there must be an increasing focus on its security risks and how to reduce them.

Connecting any simple, low-cost device into the Internet is a risk in itself. Whereas a PC is a relatively sophisticated endpoint with plenty of resources to host its own internal firewall and anti-malware systems, you cannot build expensive protection into a $5 monitor. There is a lot of interest in a new generation of tiny low-cost computers such as the Raspberry Pi to bridge that gap, but much of the security burden falls on the network itself, which must be secured if we are going to get the full benefits of an IoT connecting millions of tiny sensors.

Complexity problem

Then comes the mounting challenge of complexity, when many different protocols share the same complex network. Choosing which type of wireless connection for an IoT device means balancing a number of factors such as data throughput, range of dispersion, endpoint battery life and latency, as well as the size and cost of the receiver. Each system has its own benefits and drawbacks. For example:

- Cellular data is available almost everywhere, but uses a lot of battery power, relatively expensive hardware and ongoing network charges.

- Wifi is nearly everywhere in office buildings, uses relatively inexpensive hardware, but its password security provisioning would be unsuitable for, say, linking smart lightbulbs.
- Bluetooth is even cheaper and offers long battery life, but operates over a very short range and would need a central access point to link back to the IoT network.
- Mesh networks are much more complex to set up, but do offer a highly resilient network solution.
- Inexpensive, low-power wide area network (LPWAN) technologies such as Sigfox, LoRaWAN and Ingenu are ideal for low-data traffic such as smart meters, but have little value for more sophisticated communications.

Learning through testing

The companies developing IoT technologies and devices do have core competencies in these markets, but they have less exposure or knowledge of the ever-changing threat landscape. This is where companies with in-depth protocol knowledge and experience of testing complex environments end-to-end can play a role. Testing the performance and security of mobile phones, for example, requires a deep understanding not only of the various mobile telephony technologies, but also of Bluetooth, wifi, GPS and the additional vulnerabilities of a handset operating simultaneously on all those different networks.

"Criminal attacks come in two main flavours: theft and blackmail. The criminal can either steal money or data from a system, or threaten to damage it unless paid a ransom"

Testing such networks end-to-end requires close collaboration with a very diverse customer base, spanning many vertical markets. It is clear from contacts in these markets that there are currently five key areas of IoT innovation, namely:

- Smart cities.
- Heavy equipment.
- Automotive.
- Healthcare.
- Monitoring.

It is also clear that criminal attacks come in two main flavours: theft and blackmail. The criminal can either steal money or data from a system, or threaten to damage it unless paid a ransom. If the driver is terrorism or sabotage, instead of criminal gain, the perpetrator can simply inflict that damage without any warning.

"Criminals could simply threaten to make a public announcement that the trusted hospital system is far from secure. At least one large US hospital has already suffered such attacks"

Taking the automotive industry as an example, connected cars embrace many IoT facets, from simply delivering Internet connectivity and GPS location to the vehicle, through monitoring the vehicle for optimum safe performance and anti-theft measures, to sophisticated accident avoidance or even driverless vehicle functions.

Criminals have a choice of either hacking the system to steal the vehicle or its contents, or else blackmailing the manufacturer. GPS jammers are already being used to hide the location of stolen trucks, and remote locking systems have been hacked to enable car theft. Such attacks are obvious, but represent just the tip of an iceberg. A connected car might allow the owner to restrict its speed to 100kph when his son is driving it, but the son could hack it and go racing, or his rival could set the speed threshold embarrassingly low. Then what about busy executives using the car's Internet connection for a bank transfer – can they be sure that their transactions are every bit as secure as they would be on a home or office connection?

Healthcare is especially vulnerable to the blackmail threat: one's patient medical record might not have much value to the criminal, but the threat to upload all the medical records from a private hospital to the public Internet could be worth a colossal ransom. Or the criminals could simply threaten to make a public announcement that the trusted hospital system is far from secure. At least one large US hospital has already suffered such attacks.

Large industrial control systems are another very tempting target for blackmail. Tampering with such processes can not only lose money during the shutdown but also leave the whole system badly damaged – as when the Stuxnet worm destroyed a thousand centrifuges in Iran. And the sheer scale of utility networks makes them ideal for terrorist attacks that could disrupt electricity, water or communications across a vast area.

We have not included finance in this hit list, although criminals have long been targeting financial systems as the most obvious way to steal money. Although mobile banking apps and ATMs can be seen as forming an IoT, these endpoint apps and devices are expected to be regularly tested for security issues. So this does not present the same problem as the more general IoT of millions of devices that were not historically designed to be connected and have no inbuilt security against hacking.

Security through testing

Testing an Internet or cloud is necessary because the system is too complex for anyone to analyse and predict every possible form of vulnerability or failure. Apart from the physical complexity of a large network, there are many different levels to analyse, beginning with the firmware and different operating systems on the network. Then there is a constantly evolving population of applications on the network, plus all the different protocol and security systems. In a cellular network, access security can be in the SIM card, where a passcode is needed for wifi access, where Bluetooth can either use a password or pre-shared application key.

Next there are specific security functions such as authentication of human or machine users and determining what level of interaction they are authorised to have when accessing the network. Finally there are questions of physical security: are the endpoints and the actual network sufficiently rugged? Are they vulnerable to being tampered with?

A lot of thought will go into each of these aspects but when it comes to determining the overall end-to-end security, there is only one way to know and that is to test it. The network can be modelled accurately under laboratory conditions and tested exhaustively for functionality. Then it can be tested for performance to see if it works reliably under all normal operating conditions. It is also possible to monitor an operating network continuously for signs of trouble.

Reliability through testing

Then there is the question of how reliably the IoT will perform under extreme or stress conditions. There are many obvious and less obvious ways these extremes can arise.

Taking for example an IoT that connects intrusion alarm systems across a metropolitan district: typically these alarms use the local wifi to maintain connection with headquarters and only turn up a 3G connection if the wifi fails. So what impact will it have on the mobile network if there is a power cut across a whole district and thousands of alarm systems are simultaneously connecting to the mobile network? How might that briefly affect other systems on the network, let alone the actual alarm response?

Being able to anticipate such problems and recommend suitable tests demands experience and in-depth understanding of the many factors involved. The good news is that test solutions, readily programmable to emulate every such situation, are available. The tester does not have to build up realistic traffic scenarios parameter by parameter – though that is certainly possible. Instead, the system can record real-life traffic data and then multiply it many times to emulate traffic storms – as when an emergency causes a surge in activity.

Independently, or at the same time as the performance is being tested, the network can be subjected to any number of known malware attacks – and if it is a cloud-based test solution it will be kept up to date with the very latest malware. It's also possible to use 'fuzz testing', where you are not just testing for known attacks but also for borderline errors that can arise when wrong data is accidentally or deliberately injected into the system.

Balancing opportunities and risks

There are endless opportunities offered by connecting machines to machines via an IoT, but we have a lot to learn about the risks that might arise from such a complex system – especially the risks posed by deliberate criminal intentions. Key areas that will need very careful and comprehensive security measures include the automotive industry, healthcare and industrial and utility control systems.

Although the IoT adds many new factors and combinations of risks, the basic techniques for testing complex, multi-protocol networks are already well established. Sophisticated test solutions are already available and have long proven their ability to anticipate problems and ensure security under every normal and extreme operating condition.

About the author

Sameer Dixit is senior director for security consulting at Spirent Communications and has more than 15 years experience in penetration testing and security research. Dixit leads Spirent's ethical hacking and security research team, Spirent Security Labs. Prior to Spirent, he worked for leading security companies such as Trustwave-SpiderLabs and Cenzic, where he led the penetration testing, vulnerability scanning and managed security testing services team. Dixit has contributed research for OWASP, been quoted in various industry-leading security and business publications, and regularly speaks at cyber-security events. He blogs on emerging web and mobile security trends.

Data classification: keeping track of your most precious asset

Steve Mansfield-Devine

Steve Mansfield-Devine, editor, *Network Security*

Many companies are now in the data business, whether they know it or not. Popular clichés about data turn on words such as 'lifeblood' or 'crown jewels', underscoring the inescapable fact that information is an organisation's most critical asset. But, as we examine in this interview with Rui Melo Biscaia, director of product management at data security specialist Watchful Software, too many organisations fail to understand how their data should be categorised and handled. And a lot of them don't even know what data they've got.

Classified information

Data classification is a critical first step in the implementation of many security solutions. So how do you get started?

"You need to understand what data you have," explains Biscaia. "And then you need to understand how sensitive that data is for you and how important is it for you to protect it from falling into the wrong hands. So let's take the step back into data discovery. Data discovery can be thought of as two processes, or two ways that are not mutually exclusive. One has to do with data at rest – the terabytes and terabytes and gazillions of files that you already have in place."

It sounds easy when you put it like that. But many organisations fall at this first hurdle because they simply don't know where their files are – not all of them, anyway.

"When you're talking about data discovery for data at rest – data that has already been produced by someone and is just sitting out there somewhere – you need to have multiple ways to search for it"

"Some of them you know about because they are sitting in a file server somewhere, or any other repository that you control," says Biscaia. "But let's face it, they're everywhere. They're on private clouds, they're in public Dropbox accounts or the like, they're in a laptop that someone has at home, or even on the mobile phones that we all operate. So when you're talking about data discovery for data at rest – data that has already been produced by someone and is just sitting out there somewhere – you need to have multiple ways to search for it."

There are plenty of tools out there that will perform data discovery on endpoints and servers. And some will even perform some degree of analysis of the files, examining metadata and contents to find keywords and specific types of data, which can be gathered in a centralised database. But that only gets you so far.

"There's no way that you can control every single piece of the cloud out there, every single piece of those cloud repositories, or even your mobile phones," says Biscaia. In fact, there are several areas of technology in which you'll need to call on the services of specific solutions, he says. "For cloud, you need something that Gartner calls 'cloud access security brokers', which work like a proxy that filters everything that is going upwards to a cloud somewhere, or downwards from a cloud. For mobile phones, you need a mobile device management [MDM] system.

Rui Melo Biscaia, Watchful: "You need to understand what data you have."

"A combination of different technologies, or rather different approaches, will allow you to do data discovery at rest. And when you find it, then you need to apply some sort of policy-driven automation that will allow you to understand and reason better with regard to how sensitive it is."

Unstructured data

It's natural to think about an organisation's data in terms of databases. However, the highly organised and well-managed information sitting in databases is often just the tip of the iceberg. The vast majority of data is unstructured, scattered around the place in emails, spreadsheets and other files with no readily discernible or rational organisation to it.

"That's the tricky thing," says Biscaia. "While that data sits on a controlled repository that is protected by all sorts of boundaries and security measures and procedures, that's all fine. But as soon as someone extracts that information – copies and pastes it, or generates a report – it becomes an unstructured data file format that is just sitting everywhere."

So while you might have some form of proxy approach carrying out data discovery, the fact is that data doesn't always stay in the same place. It's constantly being copied, moved and incorporated into new forms – for example, when an employee copies database records into a spreadsheet or presentation. You need to go out and discover that data all over again. And if that wasn't enough of a Sisyphean task, there's an added burden.

"We've talked about data that already exists," says Biscaia. "What about data that we are all producing at any given moment in time? That also requires discovery – discovering who is producing data in your company – and how. Data discovery needs to focus not only on data at rest – data that already exists, including unstructured data file formats – but also data that is produced as we speak, data that is just placed into an email, or placed into a file that is put in a Dropbox folder, or shared via Bluetooth, or instant messaging or whatever."

In context

Data is not an island unto itself. A seemingly innocent number or piece of text gains significance from its context and that's a factor that Biscaia believes isn't given enough weight.

"Data discovery needs to focus, if you will, not only on data at rest but also data that is produced as we speak, data that is just placed into an email, or placed into a file that is put in a Dropbox folder, or shared via Bluetooth, or shared via instant messaging"

"This is one of the interesting aspects of data discovery that not a lot of people have thought about," he says. "It's not just about PII [personally identifiable information], PHI [protected health information] or whatever; it's also, in which context is that data being handled? And should I even care about that data?"

A combination of factors – such as who is accessing the information, by what means and within what scope of reference, along with the content of the data itself – all need to be considered together. All of this conspires to make

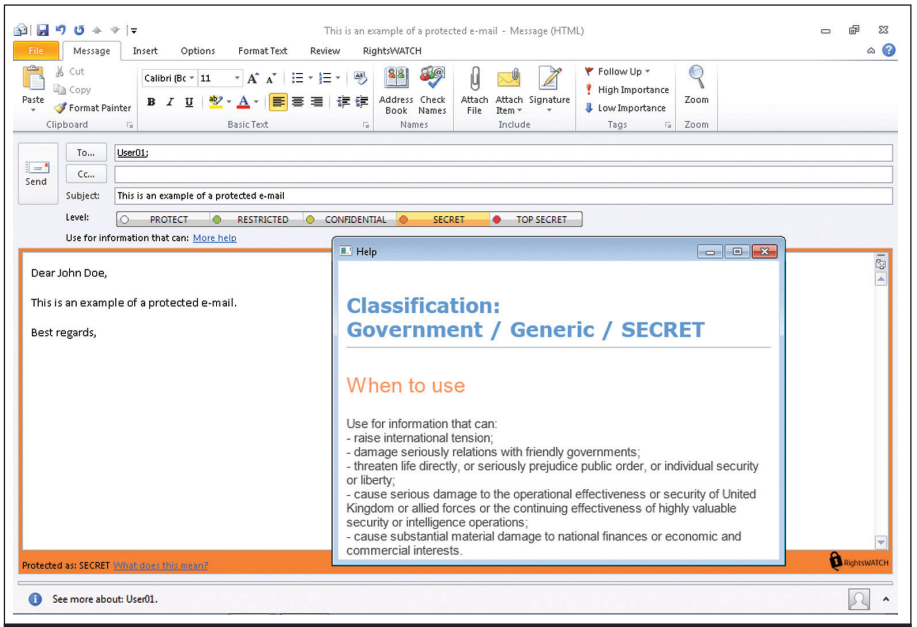

By integrating data classification into users' normal workflow, and presenting them with explanatory information, they can better understand the importance of classification.

data discovery a continuous process.

"It's not a one-time project," says Biscaia. "Basically, you need to keep on doing it and analysing the data. You need to deploy tools and processes so that you can rest assured that every single file – every single unstructured file format, if you will – that is placed somewhere in your corporate network, or even outside of that network, gets classification applied to it. It will receive some sort of tag so that you can analyse later on where that file is, what it has on it and how sensitive it is for you. It's impossible, in my opinion, to distinguish between where data discovery ends and where it starts, because those two things are interconnected and need to play along very well together."

"Of course you'll not be able to categorise and classify everything, everywhere at all times. But at least every little step you make in that direction is a step towards being more protected and mitigating more of a risk that will always exist"

Given the complex nature of finding these files and the continuous flood of new data, can you be confident that you've accounted for every byte?

"You can't, it's impossible," says Biscaia. "From a technological standpoint, nothing is impossible, but it's highly improbable that you can be sure you've discovered everything. Whenever someone says: 'I need to control everything, I need to know everything,' what I usually ask them is: 'What do you have right now?' Any little step that you can make in order to be closer to a more secure infrastructure or approach is better than what you have right now. Of course you won't be able to discover everything. Of course you won't be able to categorise and classify everything, everywhere, at all times. But at least every little step you make in that direction is a step towards being more protected and mitigating more of a risk that will always exist."

A matter of class

One you (mostly) know what data you've got, and where, the next step is to work out its importance and the degree of protection it needs. If you're starting from scratch on a data classification programme, what approach should you take?

"There are basically two ways that you can look at it from a data classification standpoint and we're talking about unstructured data file formats," says Biscaia. "You can trust users to make those decisions, or you can trust corporate policy to make that decision." At least, that's how it's often done, he adds. "But it shouldn't be one or the other - it should be one *and* the other."

Many solutions – and, for that matter, vendors and customers – assume that data classification must be user-driven because the individuals creating and working with the data understand best how important and sensitive it is. According to Biscaia, though, this overlooks one important factor.

"Users don't have 'security' in their job descriptions," he says. Attending to the security of the organisation's data and systems isn't what most people were hired to do, nor is it part of their skillset. This doesn't mean they're bad people, he points out. "It just means that they want to do their job to the best of their abilities. They need to share that data. They need to consume and produce data. But to be aware of what the security should be, or how sensitive that data is – sometimes they don't even know about that."

Most employees' encounters with corporate security policies happen when they are first hired, when they agree to abide by policies that they've probably never seen. If they're lucky they get a brief course on security on induction and then maybe a one-day refresher once or twice a year. Even if the employer is more diligent and provides frequent training and constant daily reminders, this doesn't turn employees into security professionals because that's not their function.

The consequence of this, Biscaia says, is that it's unreasonable to assume that employees are in any kind of position to judge the importance of a piece of data or file, how it should be classified and the consequences of that action. This isn't to say that employees have no value in data classification: being close to the data they create and manage, they clearly do. It's just that it's unreasonable to load the entire burden for classifying all the organisation's information on to people who need to spend their time on their real job.

The other side of the coin, then, is policy-driven classification. Here, the organisation uses centrally managed policies to decide how data should be classified depending on the content, the format (eg, spreadsheet files) and the context (eg, specific users, job functions,

departments and so on). That context could also include the type of device – for example, whether it's a corporate managed device. By combining all these factors, this approach can be quite fine-grained and flexible. But it's still not a perfect solution, according to Biscaia.

"In order to apply a policy beforehand, you need to foresee every single possible combination of keywords, environments, scopes of reference and context that just might happen," he explains.

The solution lies in a combination of these approaches, where the basic classification is implemented through policies but users get to see the classification at work. Ideally, says Biscaia, this process should be as automatic as possible, not deflecting people from their normal workflows or interrupting their tasks and also not asking them to come to decisions about data that they're not qualified to make. However, the system should allow users to override the automatic decisions for specific items where they have been given the authority to do so. In many cases this is likely to be a matter of imposing more restrictive controls on the individual file than is decided by corporate policies. But whatever the change made, it's important that the system is capable of logging the decision in a way that leaves users in no doubt about what just happened and their personal responsibility for it.

"If you don't have those policies, if you don't know even which keywords, which strings of text, which content to search for, trust the users to do that. Grow your policy via learning and observing the normal workflow of the users"

"They cannot say that they didn't see it happen and they were unable to understand that data classification," he says.

Wrong policies

There will be some security practitioners who, as soon as they hear the word 'policy' will give an involuntary shudder. Who hasn't come to grief over an incorrectly set firewall rule or intrusion detection system filter? A policy is something that's easy to get wrong. And once in place, it's not always obvious that something is amiss, even though the potential for damage is great. On the other hand, maybe you've spent years honing your security policies and want, somehow, to connect a new data classification and data loss prevention (DLP) programme to them. How do you set out to ensure you get this right?

Every organisation is different because of their varying levels of maturity when it comes to information security and their specific information technology architectures. But let's say you're an organisation with no data classification or security policies and you want to fix both of those issues.

"Start in a binary world," says Biscaia. "Is it sensitive, is it non-sensitive? If you don't have those policies, if you don't know even which keywords, which strings of text, which content to search for, trust the users to do that. Grow your policy via learning and observing the normal

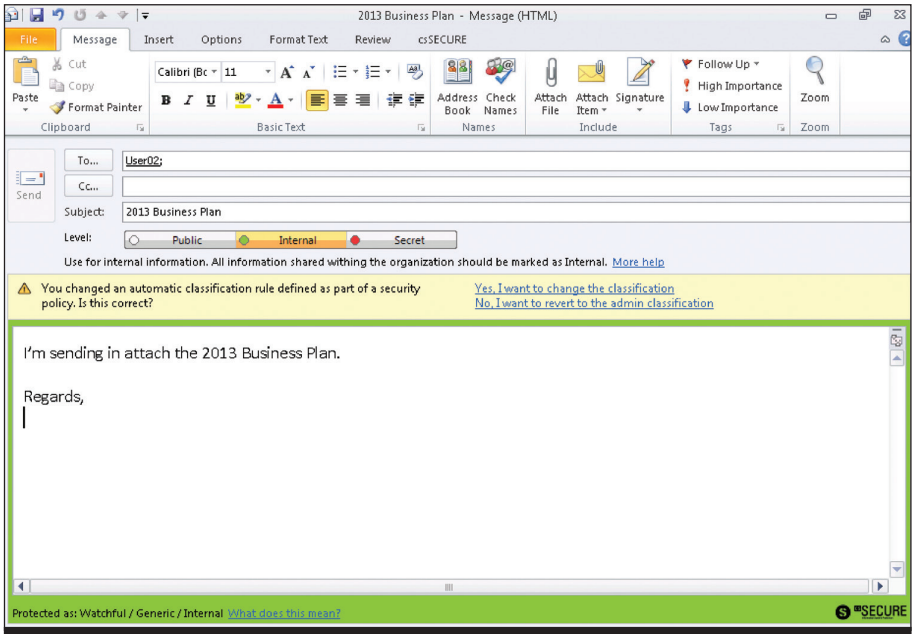

It's possible to include capabilities for users to override default classification decisions while also making them understand that these decisions will be logged.

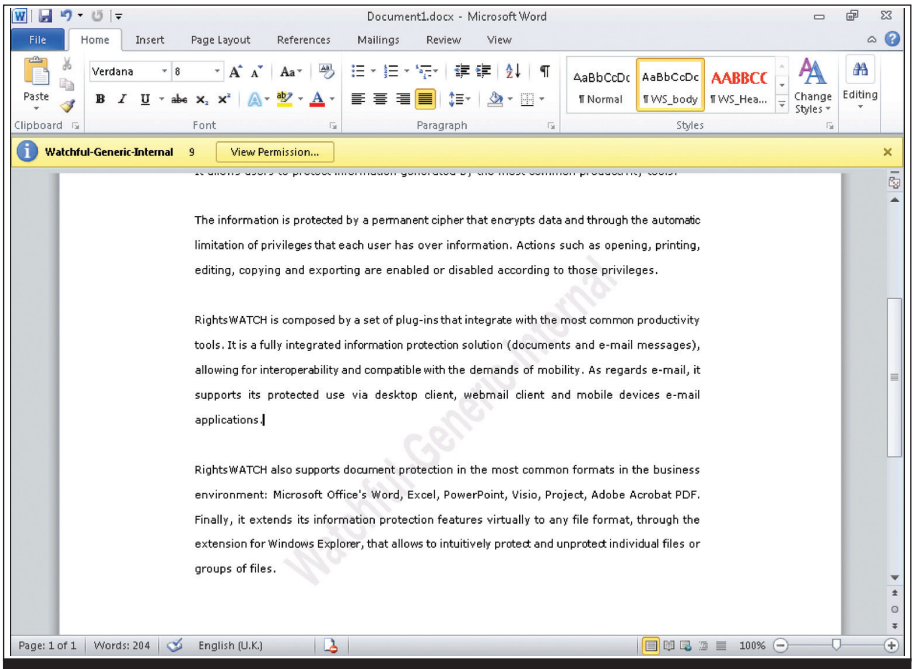

In addition to setting classifications in metadata, it's also possible to embed them in documents – in this case, a Word file.

workflow of the users. You will understand that the majority of your users, in a specific context, will classify specific files in a particular way. That will lead you to understand that you should have a policy to allow that to happen without asking the user. And, of course, you can fine tune it as you go along. Don't start running, don't start building your complex matrices and policies and start enforcing it. Start small, grow from there, and then adapt it as you go along."

Biscaia says he's seen this work well with organisations that, in spite of data discovery, still don't know where a lot of their data is and have no idea what classification categories they should apply. Starting with the users can be the best way to define those classification levels.

At the other end of the spectrum are large organisations that have put a lot of resources into data security policies.

"I've seen information classification documents 400 pages long, written by KPMG, Deloitte and the like, that are very, very comprehensive, very powerful, and they cover all the angles," says Biscaia, "but no-one even knows that the thing exists. They have all the policies written, but then you have no way to put them into practice. If you want to boil the ocean, you will fail. If you want to put that comprehensive classification and data discovery into the equation, you will fail. Again, you will have to grow them – start small and grow from there. But the difference is, in this case, you know where you want to go. You want to implement that document, but let's start small."

Joined-up approach

As for creating a joined-up environment with other security solutions, the obvious candidates here are DLP and other forms of data protection, including encryption.

"We are all thinking about encryption these days," says Biscaia. "We're all thinking about protection, about how you make the data file format accessible only to the person that should see it. It is about protecting that data wherever that data is. I don't care if it is in the corporate repository, in my laptop, in a USB key or in my mobile phone – that data is protected, and if someone tries to open that data file, only if they have a security clearance driven by authentication will they be able to open and leverage that data. To achieve that, you need encryption."

The trick is, what kind of encryption? Various forms of digital rights management (DRM) have become popular as a way of limiting the propagation of data to people who shouldn't have access to it. But there lies the rub – do you know who those people are for every piece of data? Encrypting data willy-nilly without a clear understanding of who needs access to it and in what contexts is a path to failure because it will severely interfere with the ability of your staff to do their jobs.

Conversely, if you have an effective data classification system upstream, with well-defined policies that relate to user roles, data types and contexts, then you'll stand a much better chance of being able to roll out an enterprise DRM system that actually works.

"The idea," says Biscaia, "is that you have data discovery and data classification from a user-driven, policy-driven approach that drives the encryption through the taxonomy that you've established for that data file format."

Automating the process

It's all very well have a user-driven approach to evolving your policies, but how do you capture those user decisions? And to what degree can you automate the process?

"What we need to have is that every single time those users are interacting with that data, they see the policies. And they need to be educated as to which policies should be applied to that data file"

"When you have these processes in place – data discovery, data classification, data protection, even DLPs and MDMs – you need to have all those logged," says Biscaia. "You need to log who did what, when and how. But bear in mind that every piece of this enterprise security puzzle plays an important role. A data classification tool is good at doing data classification, but it is not good at replacing a security incident and event management [SIEM] tool, for example."

In other words, when you've logged all the user decisions, all you really have are, well, logs. Making sense of them and turning them into policies takes more work. You can feed the data classification logs into a SIEM as part of your security infrastructure and you can use reporting tools to inform your policy-making. Properly designed reports will allow you to make sense of what's going on, but ultimately this is a task that calls for some good, old fashioned human intelligence.

Staff training

It's clear that, even with well-developed corporate policies, users play a key role in defining and running a successful data classification programme. But since we've already established that security isn't really part of this job, how do they know how to do that? This is where training enters the picture.

"The traditional way of training," says Biscaia, "is to gather a bunch of key users into a room every or six months. We'll do some training to make sure that they are at least aware of the existence of these policies." That's fine to an extent, he says. But there's a more effective approach. "What we need to have is that every single time those users are interacting with that data, they see the policies. And they need to be educated as to which policies should be applied to that data file, and which types of actions will they be able to perform. These tool tips, these awareness mechanisms, will help them to be trained in their daily jobs. It means that the training never ends, because every time they save a document, every time they send an email, every time they share that data with someone, that policy is shown to them."

This constant interaction helps users to understand not just that a policy has been applied, but why: why they can't send that email that has sensitive information to the wrong recipient, because that person is not an authorised recipient of such content; why they can't save

a specific document that holds sensitive information into a Dropbox account, but they can save the very same document into their desktops.

This method greatly lessens the learning curve, claims Biscaia. "It may also mean – and I emphasise the word 'may' – that the liability of the enterprise decreases, if and when the enterprise has to go to court and explain why a security breach occurred."

This interactive method puts data classification and security at the forefront of employees' use of the data, maintaining an awareness that information is something that must be treated with proper respect, but without interfering with work processes. At the same time, by logging every interaction and every decision (such as a user overriding a policy, where allowed), an organisation can demonstrate due diligence when it comes to protecting information.

"If I'm a C-level executive in China handling this data file, there are certain policies that need to be applied, but if I'm the same C-level exec handling the same data file when I'm in Germany, different policies may be applied. For that you need data classification"

This could be of great significance when the EU's General Data Protection Regulation (GDPR) comes into force in 2018. This regulation could potentially lead to very large fines for organisations that suffer data breaches. But unlike many other regulations – such as the Payment Card Industry Data Security Standard (PCI DSS), the GDPR offers no framework or guidelines for how you should ensure data protection. It simply demands that you do it. Following a breach, an organisation can't simply point to compliance with a number of technical standards. A system that logs all decisions made about the confidentiality of data could provide valuable evidence that the relevant efforts were made by a breached firm.

"Everything that we've been talking about thus far is really mapped into the

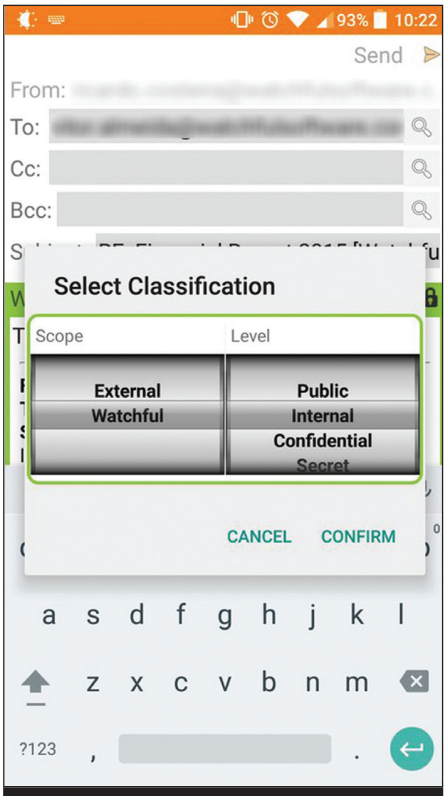

Data classification isn't just for workstations and laptops: with so much data being held on mobile devices these days, it's important that these platforms are also part of the data classification and management process.

GDPR," says Biscaia. "Enterprises are liable to protect their sensitive data from falling into the wrong hands. Where is that data? I don't know. As an enterprise CISO, I don't even know where that data is. So first I need to deploy data discovery capabilities to understand what that data is."

Into the cloud

Biscaia points out that the GDPR makes firms liable even for data that they own but which isn't stored on their own servers. "As I'm moving to the cloud and third-party suppliers that host datacentres for me, I am liable and responsible for their policies. Under the GDPR I will need to certify and be able to trust that the security policies of my third-party suppliers are in accordance with best practices from a security standpoint."

The GDPR is helping, finally, to push some firms towards encryption. "I need to have a policy-driven capability to apply different encryption types, different encryption strengths, different encryption keys into different types of data, otherwise I get a one-size-fits-all type of encryption, which basically means that it's all or nothing – either everyone can do everything, or no-one can do anything, and that's not really what I want to do," says Biscaia. "If I'm a C-level executive in China handling this data file, there are certain policies that need to be applied, but if I'm the same C-level exec handling the same data file when I'm in Germany, different policies may be applied. For that you need data classification."

The GDPR may also be driving firms that have previously been reluctant to finally embrace data classification, but there are a lot of doubts, says Biscaia. "Enterprises are saying that they know they need to be prepared and the thing that comes into their mind is encryption. Once they do that, without thinking about the upstream and downstream consequences, they fail, they create a monster. They have all sorts of encryption, they have all sorts of protection, but then no-one can use the data. And then you have to manage devices and you have to on-board devices and stuff like that, so it becomes very complicated for an enterprise to manage from a security standpoint. That applies to the users, too, which is the fundamental piece of the puzzle here. They rebel, they start calling the helpdesk and the helpdesk will uninstall any agents from the machines, so it gets messy."

Encryption is undoubtedly an important part of the solution to data protection, Biscaia says: "It's good IT security practice to have it. But you need that data classification upstream, to let the encryption work in your favour. Data classification, as small as it is from the perspective of the enterprise security puzzle, will help you drive that encryption and data protection better, which ultimately means you won't fail."

About the author

Steve Mansfield-Devine is a freelance journalist specialising in information security. He is the editor of Network Security and its sister publication Computer Fraud & Security. He also blogs and podcasts on infosecurity issues at Contrarisk.com.

The hard truth about hardware in cyber-security: it's more important

Mathias Wagner

Mathias Wagner, NXP Semiconductors

When it comes to cyber-security, there is currently an emphasis on the need for software to be secured, while hardware appears to have taken a back seat. However, a recent Kaspersky report showed that 92% of hosts connected to Industrial Control Systems (ICS) contained vulnerabilities that can be exploited to attack, take over or even harm devices and their normal mode of operation.[1]

Hardware attacks aren't prominent, as the criminal often needs access to the device and/or chip itself, making them much more difficult to perform. But this doesn't mean that they aren't happening.

There has been a dramatic increase in the number of professional cyber gangs using crypto-extortion, which is when criminals demand money for unlocking the victim's device. Norton estimates that 5,700 computers are infected every day – that's 171,000 per month.[2]

In addition, ransomware attacks against businesses are also on the rise. Nearly 40% of all businesses worldwide experienced this type of attack in the past year, according to a survey of over 500 companies in four countries by computer security firm Malwarebytes.[3] In a ransomware attack, attackers use malware to lock down clients and systems to put them out of business. These types of attacks can lead to the loss of highly sensitive customer and patient data and have been seen in a variety of different business types and sectors, including hospitals and public administrations.

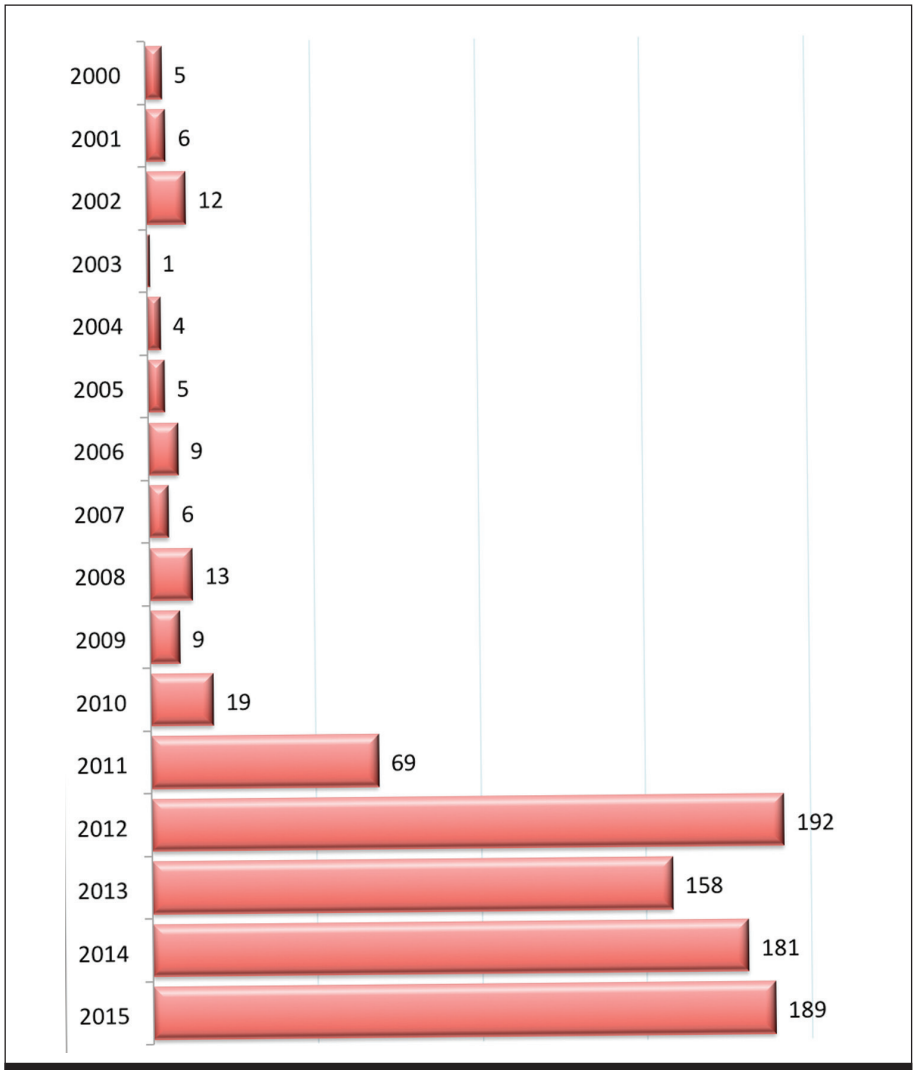

The number of vulnerabilities in industrial control systems, by year. Source: Kaspersky.

Security by design

Security providers are becoming increasingly concerned about these types of attacks. With the rise of the Internet of Things (IoT) and the proliferation of devices such as smartphones and tablets, smart appliances and meters as well as self-driving vehicles, the threat of hardware attacks is more omnipresent and potentially damaging than ever. Think of it like the 'hands and feet of the Internet' – with more devices, there is an ever-increasing number of physical assets being connected. In order for them to be secure, connected devices need to be designed in a way that even knowing the design does not give any substantial advantage to the attacker.

According to *Technology Review*, there will be an estimated 50 billion connected

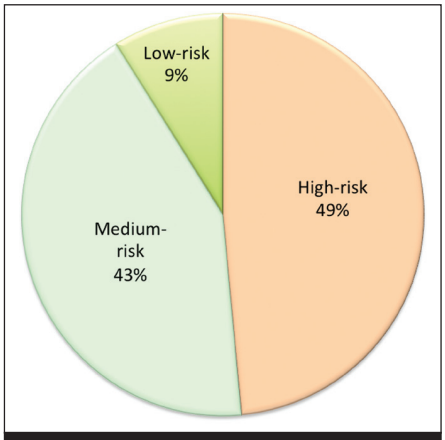

Types of vulnerabilities in industrial control systems in 2015. Source: Kaspersky.

devices by 2020. This will be followed by data manipulation, data theft and other cyber-attacks that will weaken peoples' trust in digital products and services.[4] That's why it's imperative for secure hardware to become a condition for maximum security in sensitive environments, such as securing IoT devices and self-driving cars.

After the US, Germany suffers the second highest amount of loss from cyber-attacks. According to *Wirtschaftswoche,* for a German company, data theft costs an average of $3.1m per case – more than any other European country.[5]

With software issues, solutions such as anti-virus or firewall software can help combat certain types of attacks. However, because of the nature of hardware, security has to be inherent in its design. This threat is something the industry and businesses can't afford to ignore.

Let's look at some examples of the more common types of hardware attacks and how businesses can protect themselves.

Physical or implementation attacks

In most cases with implementation attacks on microchips, the criminal needs to be in possession of or have access to the device to be able to tamper with it.

A powerful type of attack in this case is when a chip is exposed to a laser or an electric power glitch by sending a high voltage level that exceeds the limits of the chip's power supply to make the microchip trip. If a light is shone onto the chip, it generates additional charges that aren't intended to be there and can create havoc. The result is a chip that doesn't execute as expected.

Side-channel attacks

There are, however, other powerful attacks where the adversary doesn't need to actively manipulate the device's operating conditions – and sometimes doesn't even need physical access to the device/chip – to carry out a physical attack. This type of attack is referred to as a 'side-channel' attack where information leaks out of a device unintentionally.

In cryptography, a side-channel attack can occur when the attacker gains access to information through the physical implementation of a cryptosystem. All of the following methods provide attackers with extra information they can then use to break the system.

Information can be obtained through timing information (ie, how much time various computations take to perform), power consumption (ie, making use of varying power consumption by the hardware during computation and correlating the crypto computation to uncover the secret key), electromagnetic leaks (ie, obtaining plaintext and other information from leaked electromagnetic radiation), or even through sound (similar to a power-monitoring attack where the attacker obtains information by monitoring sound produced during a computation).

Hardware reverse engineering attacks

Probably the most difficult type of chip attack is what's known as 'hardware reverse engineering'. Again, this type of attack requires the person to be in physical possession of the chip.

"Hardware attacks are just not as visible. One reason for this is the perception that hardware's not as flexible as software and is harder to change and therefore inherently more secure"

However, in this case, the person would use a microscope, x-ray machine, MRI or similar to zoom into the chip so he can understand all of the transistors in order to reconstruct the architecture of the chip. That way, he can understand when the best time is to attack. This is the most invasive of hardware

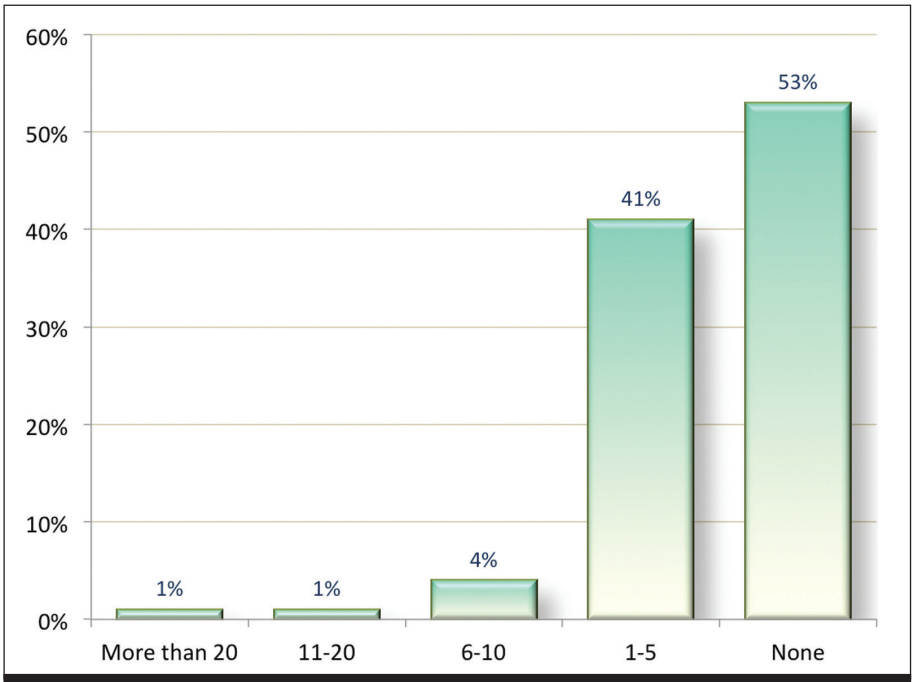

Number of ransomware attacks suffered by surveyed firms in the past 12 months. Source: Malwarebytes.

attacks but it takes heavy machinery to succeed.

Taking a back seat

It's not necessarily that hardware attacks have taken a back seat to software attacks but rather they're just not as visible. One reason for this is the perception that hardware's not as flexible as software and is harder to change and therefore inherently more secure. For example, modern, high-end smartphones use secure elements inside because the manufacturers have the budget to put the costly security measures in place. However, security implemented in hardware benefits from Moore's law: it gets cheaper over time.

Apart from hardware's inherent characteristics, software also isn't subject to the same level of certification schemes as hardware. Hardware can be secured with a pure software solution but it depends on where its intended use is. For example, if a white box model is used to hide cryptographic keys – eg, for secure payments on a smartphone – the current deployed solutions have been shown to be vulnerable to attack. On top of this, the manufacturer has to take a risk management-based approach similar to credit card companies, which add substantial running costs for monitoring, maintenance and software patching. For instance, if the token on the white box implementation is hacked but risk management is running on the cloud, the provider will be able to trace unusual activity such as your phone being used to buy something in Brazil when you were in Belgium.

Standing on guard

So what are the steps that companies can take to ensure their hardware is secure? In order to carry out the majority of these types of attacks, the person needs access to the actual device. There are relatively few attacks – around 5% or less – that the person can carry out remotely if the device is online.

For example, the strongest form of side-channel attack (in an information theoretic sense), known as a template attack, requires the person to have access to an identical experimental device that can be programmed to his liking. In this case, the person buys a chip somewhere else and gains access to the actual target of the device from which he reads out the characteristics and puts them into a shape sorting machine to see if they work or not. Think of it like a shape sorter that children play with – the peg either fits or it doesn't.

"In the era of the IoT, the old model of cryptography seen in the 1980s, where two endpoints (humans and/or computers) were sending messages to each other, is no longer relevant as there are multiple endpoints"

This is why chips have to be secure by design. For example, to shield its products against known and future hardware attacks, NXP has set up a Vulnerability Analysis Team to test products against 200-250 types of attack.

The role of cryptography

Cryptography plays a huge role in the security of hardware. By design, cryptography is a mathematical method to secure things – be it hardware or software. However, in the era of the IoT, the old model of cryptography seen in the 1980s, where two endpoints (humans and/or computers) were sending messages to each other, is no longer relevant as there are multiple endpoints. That's why it's even more imperative that hardware is tamper resistant.

To ensure this, the entire architecture and value chain of the connected world – infrastructure, hardware (including gateways, end nodes, actuators, sensors), interfaces and communication, software, network and cloud – need to be part of a comprehensive approach, incorporating security and privacy 'by design' as starting principles.

Research has resulted in the development of a number of fundamental proposals for businesses when it comes to guarding against hardware attacks, as stated below.

Integrity of information: the integrity of information from the connected world – of devices, objects and sensors – needs to be protected at all times. Protecting the users' rights in this context demands a decoupling of 'human' identities from the identities of the devices involved.

Privacy protection: in view of their data sovereignty, people and businesses need to have absolute transparency over who can access data from the devices they're using and what data should be accessible. At the same time, it is essential to ensure that a digital device protects the user's anonymity or identity.

Security-by-design rules: the preceding issues create a need for security-by-design rules for such devices. In addition, the already available privacy-by-design solutions should be used to protect people's private lives (eg, multiple keys, anonymous attestation and zero-knowledge protocols).

Common framework: also needed is a new common framework establishing basic minimum requirements for security and privacy in the highly connected world. These should be comprehensive and should apply to all interconnected objects – from sensor to cloud – whether a connected car or a heating system in a smart home.

System of security certification: implementing the legal standards demands a system of security certification that will enable us to guarantee compliance with those minimum requirements and thereby certify both hardware- and software-based security technology.

Proceeding in this way offers a unique chance to build secure, more robust systems right from the start, providing a more secure environment to industry and consumers.

Working together

Hardware solutions are faster respective to a simple software implementation. But are hardware solutions appropriate

to answer the security requirements of the future?

Hardware solutions must be appropriate in view of future security standards – simply because there are attack vectors that can't be dealt with by software alone. The optimal security can only be achieved by a combination of both secure hardware and secure software. This begins with the conceptual phase. Hardware must be designed so that it offers appropriate interfaces to the software. Of course, the software must also address these in the right way.

"The development of connected devices must incorporate security by design as a principle. The goal is to develop an architecture and/or a design that denies attackers promising and exploitable angles of approach and protects the user's privacy"

Besides ensuring that software and hardware are more secure, the IT industry and governments around the globe as a whole need to look at creating and establishing reliable security standards. This is one of the greatest challenges in the near and distant future for the secure design of the IoT.

One challenge for the whole system will be how to proceed with personal data, anonymity and privacy. In the future, the protection of users, clients and their data will become increasingly important. Among other issues, the question will be how to effectively guard and control one's data.

It will become even more important to define and evaluate individual levels of security. This will be key in order to develop hardware and software that are ideally attuned to each other. Therefore, the development of connected devices must incorporate security by design as a principle. The goal is to develop an architecture and/or a design that denies attackers promising and exploitable angles of approach and protects the user's privacy.

Perhaps the most decisive aspect of secure systems is the ability for people to trust them. This is why we need verifiable minimum standards for system security. In other words, an independent security authority and a transparent framework of rules that guarantees compliance with standards. This will increase the level of security in the IoT and consolidate the trust of businesses and consumers into connected technologies.

Security is more about the people and processes rather than the technology alone. And that's the hardest variable to predict.

About the author

Dr Mathias Wagner is CTO of the Innovation Centre Security in the Security & Connectivity business unit of NXP Semiconductors. He is also technology manager for hardware security and tamper resistance within the global NXP Technology Competence Framework. In this role, he is responsible for driving innovation in the area of security and cryptography as well as managing the development of all security IP, be it hardware or software. As senior fellow he advises the business unit and NXP management on all aspects of security matters and is closely engaged with the leading players in a wide variety of security ecosystems. Prior to his current position, Wagner served in various security technology functions in NXP, thereby gaining knowledge in all aspects of hardware attacks targeting smartcards and related embedded devices. Before joining NXP (Philips at the time), he held a variety of research positions with Hitachi, spending nine years in Cambridge, UK, and 1.5 years at the Hitachi Central Research Laboratory, Tokyo. Wagner graduated from the University of Hamburg, Germany, with a first class Diploma in Theoretical Physics (major) and Mathematics (minor), and subsequently a first class Doctor of Natural Science in Theoretical Semiconductor Physics. He is a member of the 'Studienstiftung des Deutschen Volkes', an organisation that sponsors the 0.5% best German students.

References

1. 'Industrial cyber-security threat landscape'. Kaspersky/Securelist, 11 Jul 2016. Accessed Dec 2016. https://securelist.com/analysis/publications/75343/industrial-cyber-security-threat-landscape/.
2. 'Ransomware on the rise: Norton tips on how to prevent getting infected'. Norton/Symantec. Accessed Dec 2016. https://us.norton.com/ransomware/article.
3. 'International study finds nearly 40% of enterprises hit by ransomware in the last year'. MalwareBytes, 3 Aug 2016. Accessed Dec 2016. https://press.malwarebytes.com/2016/08/03/international-study-finds-nearly-40-percent-of-enterprises-hit-by-ransomware-in-the-last-year-2/.
4. Buck, Christian. 'Bis 2020 gibt es 50 Milliarden vernetzte Geräte'. Technology Review via Heise Online (in German), 11 Nov 2013. Accessed Dec 2016. www.heise.de/tr/artikel/Bis-2020-gibt-es-50-Milliarden-vernetzte-Geraete-2041999.html.
5. 'Wenn Cyber-attacken in den Bankrott führen'. Wirtschafts Woche (in German), 25 Nov 2015. Accessed Dec 2016. www.wiwo.de/unternehmen/it/hackerangriffe-auf-unternehmen-wenn-cyber-attacken-in-den-bankrott-fuehren/12632916.html.

The Firewall

The danger within

Colin Tankard, Digital Pathways

While many security professionals and budgets are focused on threats from external actors, the insider threat looms large. According to Vormetric, 89% of organisations are at least somewhat vulnerable to insider attacks. It states that privileged users are considered to be the most dangerous, primarily owing to their access to systems and information considered to be particularly sensitive.

This is echoed by research from the Ponemon Institute, which found that almost half of respondents believe the insider threat to be growing.

The insider threat can come from those inside the organisation that have had their credentials compromised, or are negligent or malicious. Increasingly, external actors are looking to gain a foothold on the network through social engineering exploits, often stealing the credentials of a victim within the target organisation. They then look to move laterally through the network, seeking credentials with higher levels of privileged access to gain entry to more valuable information. According to Ponemon, in 2011, just 21% of respondents reported this kind of activity but this had increased to 46% by 2016.

Many organisations rely on technologies such as security information and event management (SIEM) systems that monitor and review log files from multiple sources to determine what is happening on the network and whether behaviour constitutes a threat. However, SIEM systems are often unwieldy, owing to the vast amount of data that flows through them and the high likelihood of false positives. As a result, user behaviour analytics is increasingly being used to counter the insider threat. Combined with risk scoring capabilities, such technology can assess the risk of each user, scoring all activity against expected baselines according to types of role and activity. These systems go further in that they can instantly start to record a session, allowing the organisation to see exactly what was being typed.

Such user monitoring systems can educate users in good behaviour. In the Verizon 2016 Security Report, 46% of all user alerts were simply down to the user making an error or forgetting the company policy on data handling. Using a monitoring system to send users on-screen messages advising of their error enables the user to be educated. Research has shown that with better education, the number of data handling errors can be greatly reduced, thus making detecting the deliberate exfiltration of data much easier.

SIEM combined with a dedicated insider threat technology solution will do much to rein in unwanted user behaviour. But best practice is to also deploy a comprehensive encryption solution combined with access control capabilities.

Tracking insider threats is essential for every organisation. Insiders have access to the most valuable information and can therefore pose the greatest risk to any organisation. But no one technology by itself is sufficient. Organisations should consider using a set of technologies that work together so that they can be sure that all bases are covered.

EVENTS CALENDAR

4–6 January 2017
Real World Cryptography Conference
New York City, NY, US
www.realworldcrypto.com/rwc2017

13–15 January 2017
Shmoocon 2016
Washington, DC, US
www.shmoocon.org

24–25 January 2017
FIC 2017
Lille, France
www.forum-fic.com

25 January–1 February 2017
SANS Cyber Threat Intelligence Summit
Arlington, VA, US
www.sans.org/event/cyber-threat-intelligence-summit-2017

25–26 January 2017
Cyber Defence and Network Security
London, UK
www.cdans.org

27–29 January 2017
REcon Brussels
Brussels, Belgium
https://recon.cx

31 January–1 February 2017
Cyber Tech International Conference & Exhibition
Tel Aviv, Israel
http://10times.com/cyber-tech

13–17 February 2017
RSA Conference 2016
San Francisco, US
www.rsaconference.com

19–21 February 2017
International Conference on Information Systems Security & Privacy 2017
Porto, Portugal
www.icissp.org/