IBM X-Force Threat Intelligence Quarterly, 3Q 2014

Get a closer look at Heartbleed—from the latest attack activity to mitigation strategies—using 2014 mid-year data and ongoing research

IBM

Contents

Executive overview

Welcome to the latest quarterly report from the IBM® X-Force® research and development team. In this report, we'll look at how the Heartbleed vulnerability—CVE-2014-0160, disclosed in April 2014—impacted organizations around the world. We'll focus on how attackers continue to take advantage of this pervasive vulnerability, review potential mitigation strategies and assess how the disclosure compares to the rest of our data from the first half of 2014.

So far, the disclosure of the Heartbleed vulnerability in the OpenSSL library has been the biggest event to hit the security industry in 2014. The bug permitted unauthenticated acccess from servers and clients alike. While the initial impact of Heartbleed is waning, a second wave of new vulnerabilities found within open-source and reusable software merits further discussion.

Servers worldwide continue to be affected by this serious vulnerability, so we wanted to investigate what has happened since the Heartbleed disclosure took so many organizations by surprise. Not only did the flaw focus the attention of researchers looking for new areas of vulnerabilities within open-source and reusable code, it also gave attackers another great opportunity to use one-day attack methods.

With the help of IBM Managed Security Services (MSS), we'll first look at how organizations dealt with the immediate aftermath of the Heartbleed announcement, while also adopting practical, large-scale mitigation strategies for ongoing protection. Then, from an attack perspective, our X-Force researchers will explain what the attackers might have been looking for and attempting to achieve with this type of vulnerability.

Throughout this report, you'll learn how an unexpected, widespread and difficult-to-patch vulnerability such as Heartbleed forces organizations to look deeper into their risk management and critical communication processes. This was especially true for major software vendors who had integrated OpenSSL into their commercial products and offerings. In addition to protecting against imminent attacks to their own potentially vulnerable systems, vendors also had to perform a thorough investigation of their own products—that is, they had to determine if any of their products were using this open-source library in order to provide patches.

For many of these reasons, Heartbleed had a far greater risk impact than other types of vulnerabilities; however, consequences were not as disastrous as they could have been. There were only a handful of breaches attributed to Heartbleed even though it was a vulnerability in the core technology that protects e-commerce and helps ensure privacy.

Finally, we'll close the report with a look at how Heartbleed compares to other publicly disclosed vulnerabilities, and how this midway point of 2014 compares to previous years. The good news is that the overall disclosure trend is decreasing. But when comparing the real-world impact of Heartbleed to its Common Vulnerability Scoring System (CVSS) ranking—which is only a 5.0, or "medium" risk—our researchers noted a significant disparity. We'll discuss some of the shortcomings in the current CVSS standard, the inconsistencies in scoring across different organizations, and the loss of confidence in the CVSS score as an accurate and reliable measure of risk. Then, we'll explain how the upcoming release of CVSS version 3 is expected to address many concerns of the security industry.

What is Heartbleed?

The Heartbleed vulnerability is a bug in OpenSSL, a popular open-source protocol used extensively on the Internet, which allows anyone who knows how to exploit the vulnerability to access and read the memory of systems thought to be protected.

Vulnerable versions of OpenSSL allow compromise of secret keys, user names, passwords and even actual content. Many security experts believe that this vulnerability has actually existed for at least two years and might have been exploited for just as long. Although many companies have issued statements claiming that they have now remedied the vulnerability in their environment, there is truly no way of knowing how much data has fallen into the wrong hands through the exploitation of this vulnerability.

For more information about Heartbleed, refer to the IBM Security Intelligence blog post from April 2014[1] or the Heartbleed website.[2]

Heartbleed attack activity: Then and now

What was the real-world impact of Heartbleed? Learn how the waves of attacks have affected IBM Managed Security Services customers.

On 7 April 2014, one of the most significant security events of the past few years occurred—that's when the Heartbleed vulnerability in OpenSSL (CVE-2014-0160) was publicly disclosed. The bug was introduced approximately two years ago and left more than half a million servers vulnerable to unencrypted data leaks of system memory with minimal trace of exploitation. The disclosure caused panic across a wide range of industries, government agencies and consumer groups that had used OpenSSL to keep their transactions private—and had instead found themselves vulnerable to an attack and without any proof of when leaks occurred (that is, no log-file evidence).

Almost as soon as the vulnerability was released as an OpenSSL advisory, IBM Managed Security Services (MSS) witnessed attackers immediately re-tooling and exploiting the bug on a global scale. Once the major vendors of intrusion detection and prevention systems created protection signatures, MSS was able to see just how bad the situation had become. On 15 April 2014, MSS witnessed the largest spike in activity across the customer base with more than 300,000 attacks in a single 24-hour period. That's an average of 3.47 attacks per second for more than hundreds of customers.

Figure 1. Attack activity related to the Heartbleed vulnerability, as noted for IBM Managed Security Services customers, in April 2014

Execution of the attack activity

Let's take a closer look at the attack activity after the Heartbleed disclosure. Rather than a single IP address executing the attack repeatedly, many of the attacks used a distributed method. A wide range of IP addresses across multiple autonomous system numbers (ASNs) attacked the networks monitored by MSS. In fact, entire ranges of IP addresses attacked several servers at once. This enabled attackers to have a large, diversified attack surface and the flexibility to overcome rudimentary blocking strategies.

What are autonomous system numbers?

On the Internet, an autonomous system refers to a connected group of one or more Internet Protocol routing prefixes run by one or more network operators to support a single, clearly defined routing policy. Each autonomous system is assigned a globally unique routing number, known as an autonomous system number (ASN).

Originally, autonomous systems were controlled on behalf of a single entity, such as an Internet service provider (ISP) or a very large organization with independent connections to multiple networks. Now, multiple organizations can run the Border Gateway Protocol (BGP) using private ASNs that sit behind an ISP. Although multiple autonomous systems may be supported by the ISP, the Internet only sees the routing policy of the ISP. Therefore, only the ISP must have an officially registered ASN.

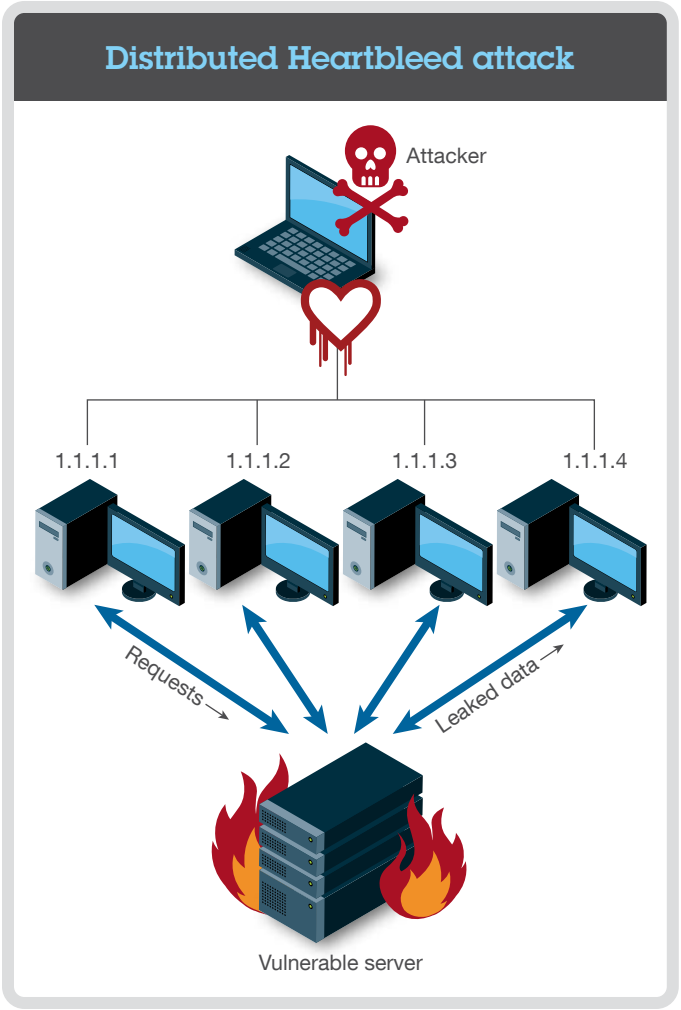

Graphic 1. Distributed Heartbleed attack

The attacks slowed down after 22 April 2014. However, Figure 2 shows that despite the leveling off of attack activity, the number of attacked customers has remained relatively consistent over time. Why? As large organizations vulnerable to Heartbleed were able to apply the issued patch to their infrastructure, they rendered the attacks less fruitful. As a result, attackers focused on other exploits. MSS witnessed a significant drop in both the number of source IPs generating attacks and the total number of attacks globally against the MSS customer base.

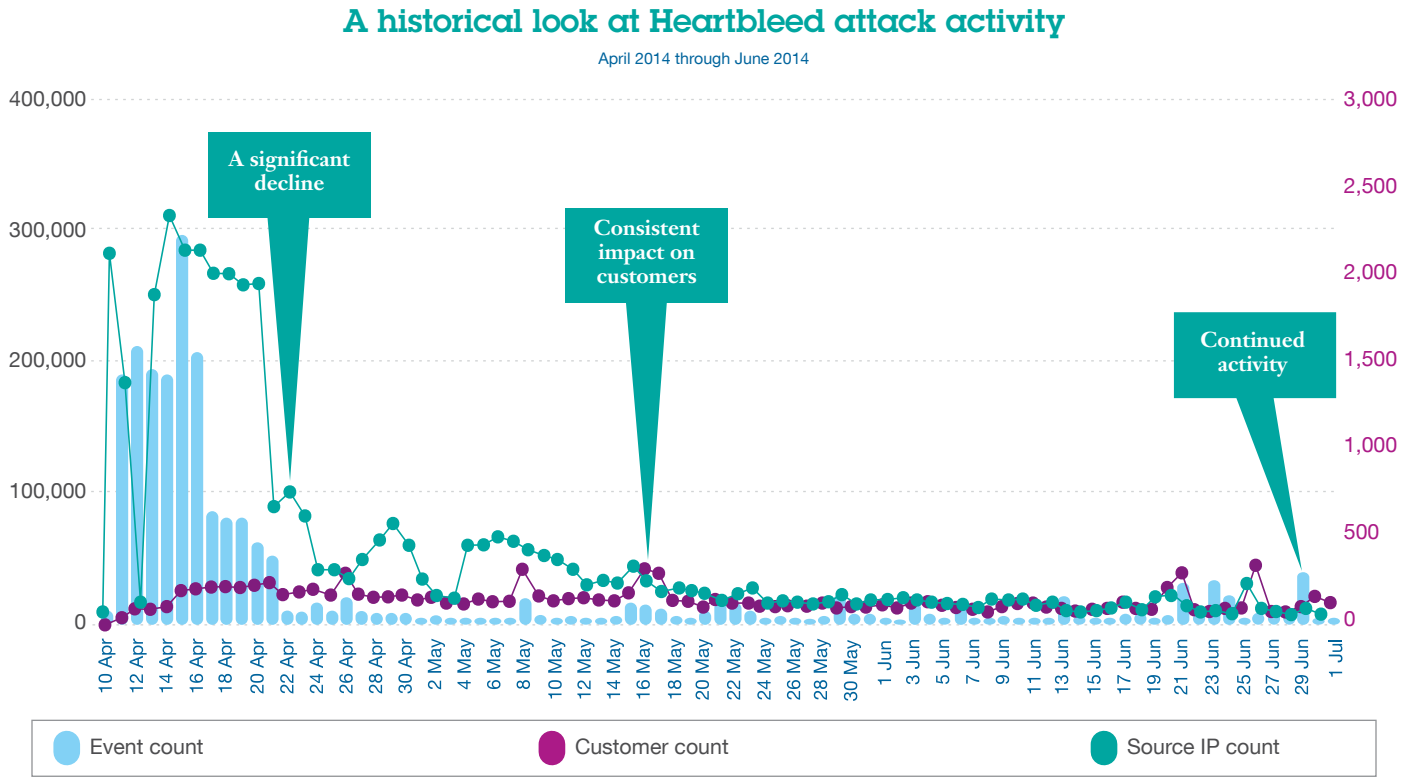

Figure 2. A historical look at Heartbleed attack activity for IBM Managed Security Services customers, April 2014 through June 2014

But the current status of attacks is still significant. MSS sees an average of 7,000 attacks per day across a large attack surface. Organizations that have applied the OpenSSL patch to their infrastructure and deployed blocking mechanisms, such as intrusion detection or intrusion prevention systems, can breathe a little easier. According to recent MSS reports, we see that despite the initial rush to patch systems, approximately 50 percent of potentially vulnerable servers have been left unpatched—making Heartbleed an ongoing, critical threat.

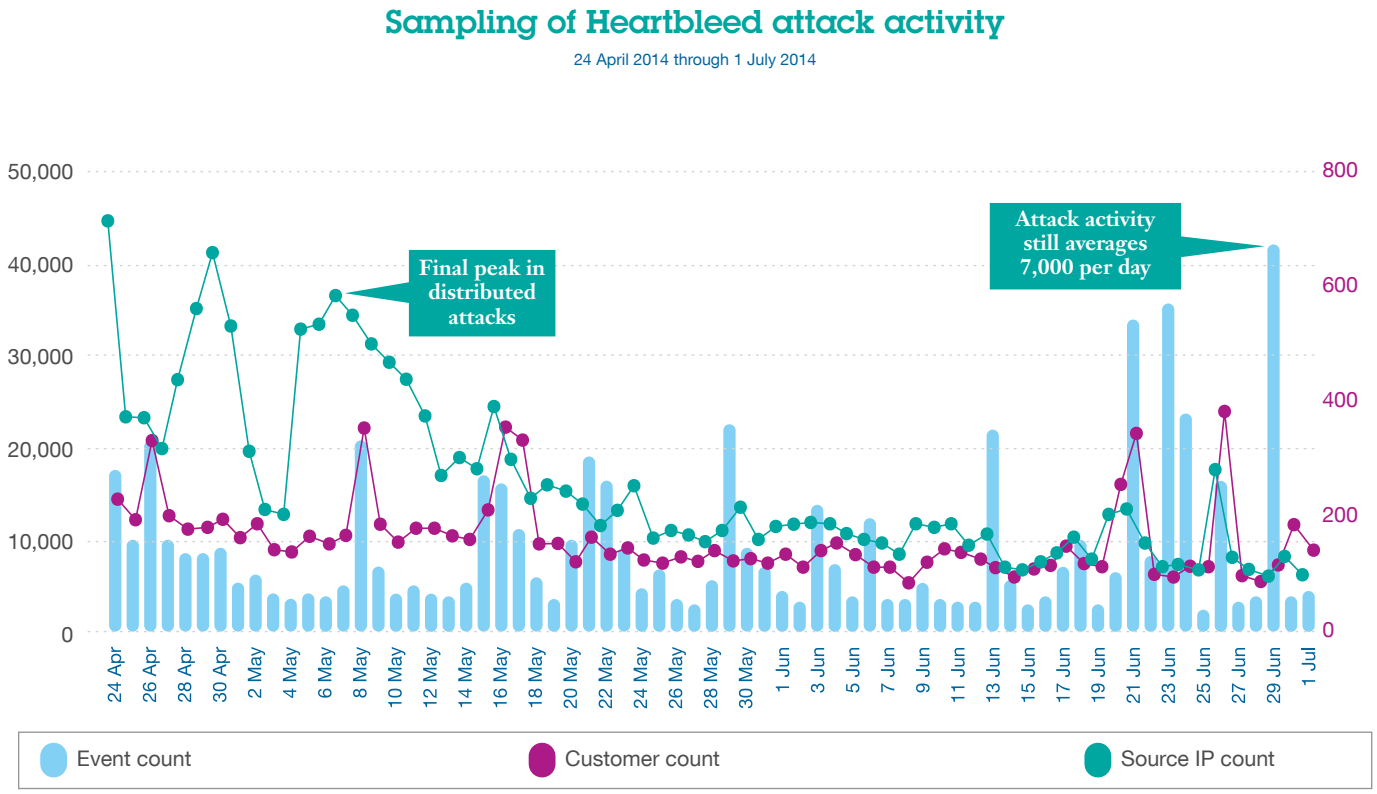

Figure 3. Sampling of Heartbleed attack activity for IBM Managed Security Services customers, 24 April 2014 through 1 July 2014

Lessons learned and recommendations

There were many lessons learned from the Heartbleed attacks. For example, MSS found that having an incident response plan—and maintaining an asset database—were both absolutely critical to reducing exposure to the attacks. Organizations that had struggled to maintain a current asset database were left blind to which systems were vulnerable and which systems were critical. Even if they had an incident response plan, they needed an up-to-date asset database in order to deploy it.

On the other hand, companies that had maintained their asset database and incident response plan were able to rapidly deploy patches on critical systems vulnerable to attack, thereby reducing their exposure to Heartbleed. They also face significantly less risk for threats in the future.

It's also important to understand the detection and defense strategies for attacks such as Heartbleed. In certain scenarios, organizations can utilize firewalls to block out the bulk of the attacks toward their networks. The MSS security operations center (SOC) applies this methodology when large global attacks happen and the majority of the attacks stem from a small subset of hosts. This blocking technique can provide short, temporary reprieve from attack activity, providing valuable time for critical systems to be patched.

Firewalls are an excellent defense when a small subset of hosts are generating the attacks. In addition, intrusion detection and prevention devices can provide an even greater protection by blocking attacks at the offending packet level. This alleviates the need for maintaining an active list of attackers and reduces the risk involved while systems are patched.

The race to prevent one-day attacks

Discover how quickly attackers rush to exploit a vulnerability such as Heartbleed—and how you can mitigate the threat.

"How fast can we get a patch deployed?" That is the question most organizations asked themselves when the Heartbleed security advisory (CVE-2014-0160) was published on 7 April 2014. Immediately after the announcement, organizations rushed to patch their systems. Meanwhile, just one day after the disclosure, a proof-of-concept tool[3] capable of exploiting the Heartbleed bug began circulating, exposing unpatched systems to skilled and unskilled attackers alike. But more troubling is the fact that also a day after the disclosure, attacks leveraging the vulnerability began to occur,[4] and the actions of some of the affected organizations in mitigating attacks or patching the vulnerability were already too late.[5, 6]

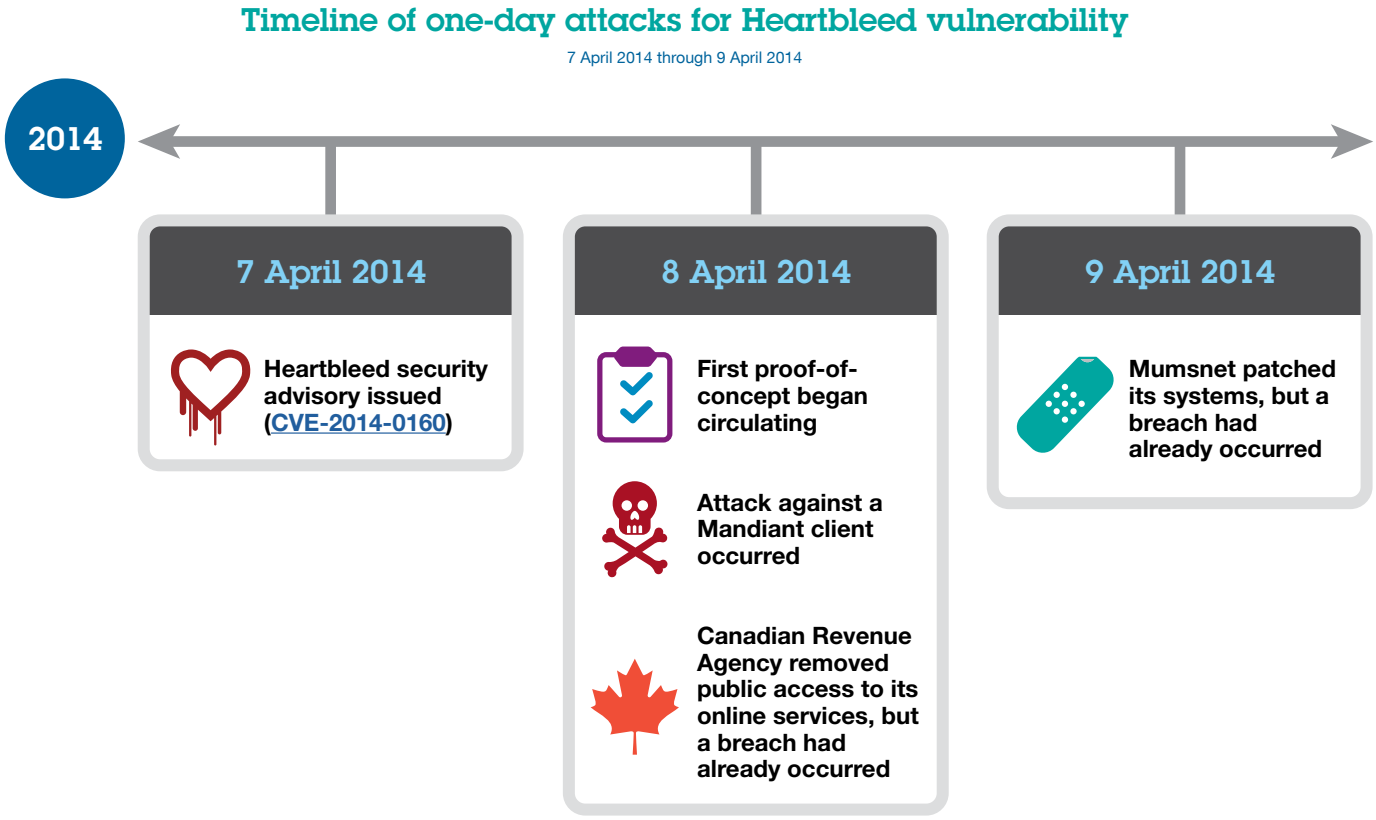

Figure 4. Timeline of one-day attacks for Heartbleed vulnerability (CVE-2014-0160), 7 April 2014 through 9 April 2014

Reflections of Java attacks in 2012

Heartbleed isn't the first time one-day attacks have occurred—that is, attacks leveraging an already-patched vulnerability. In fact, X-Force analysts noted this trend after the disclosure of the 2012 Java vulnerability (CVE-2012-1723), as discussed in our IBM X-Force 2012 Trend and Risk Report.

In the case of this Java vulnerability, a security researcher claimed that he was able to create a proof-of-concept exploit just a day after a fix was issued for the vulnerability, and a week later, published the details of the vulnerability.[7] Fortunately, in this particular case, the proof-of-concept code was not released. However, the integration of working exploit code in a popular mass exploit kit[8]—posted just a month after the patch was made available for the vulnerability—meant that many still-unpatched systems became potential targets via drive-by attacks. Later, the exploit for the vulnerability was integrated into other exploit kits, further increasing the risk of unpatched systems becoming compromised.

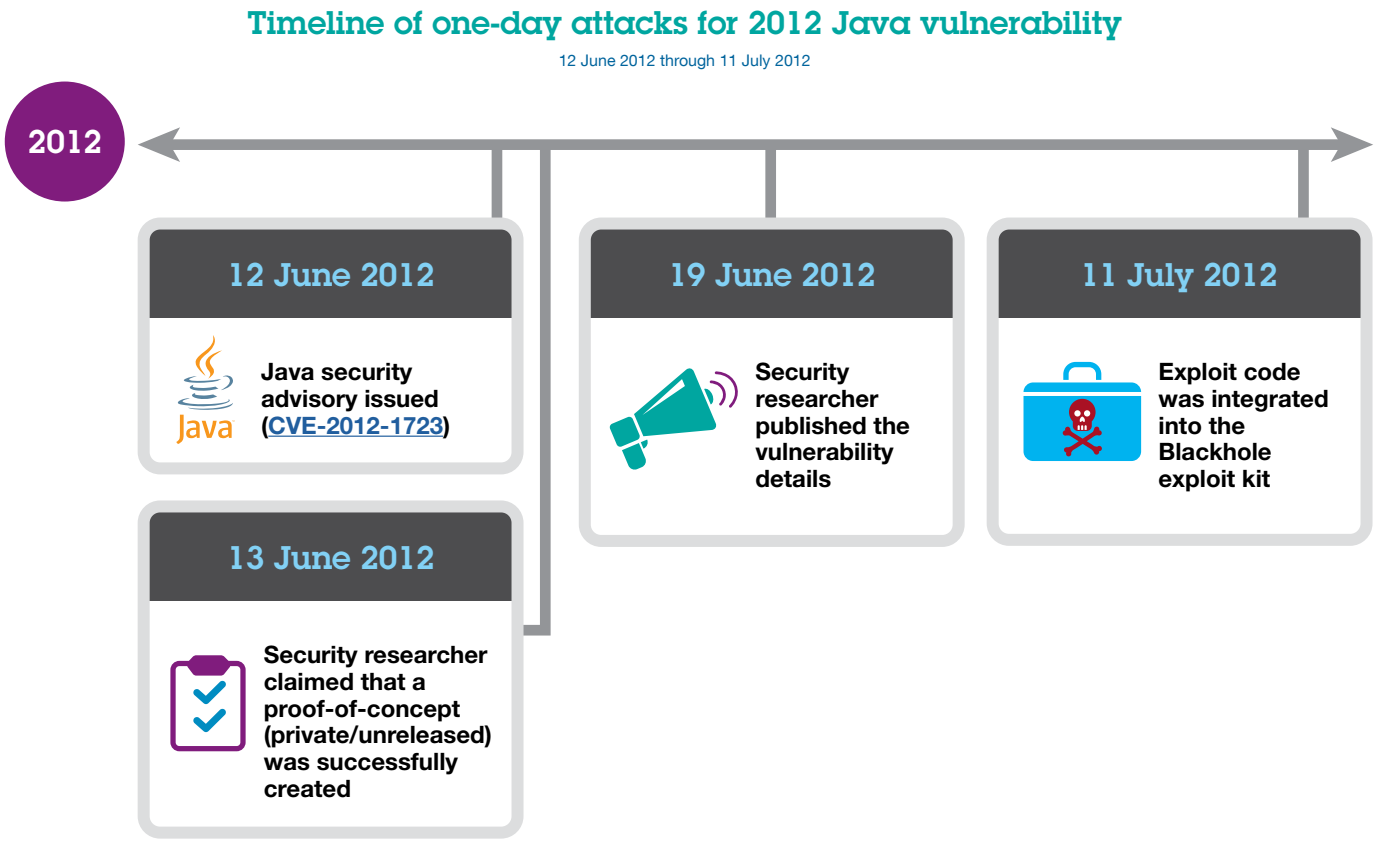

Figure 5. Timeline of one-day attacks for 2012 Java vulnerability (CVE-2012-1723), 12 June 2012 through 11 July 2012

A race to patch

Both the Heartbleed and Java vulnerability examples demonstrate that one-day attacks can be just as dangerous as zero-day attacks. Unlike zero-day attacks—where the vulnerability is unknown and a patch is not available—the issue with one-day attacks is in how long it takes a patch to be deployed. Patch management is a complex challenge, especially for large deployments and for mission-critical systems that require extensive pre-deployment testing.

Additionally, if the vulnerability is in a widely-used library, such as the OpenSSL library as in the case of Heartbleed, the issue is further complicated because organizations are dependent on the additional time it takes for software vendors to integrate and test the patches in their own products before the patched product is handed off to them. This additional waiting time means an increased exposure window that attackers can turn to their own advantage.

A race to exploit

From an attacker's standpoint, recently-patched vulnerabilities represent an opportunity, as they are potential (albeit temporary) weak points in their targets' infrastructure. For one-day attacks, the goal of the attacker is to take advantage of the exposure window of organizations between when the patches are announced and when the patches are actually deployed.

The following sections explain how attackers create one-day exploits, so you can understand how quickly a weaponized attack can take advantage of a patched vulnerability.

Locating the patched bug

Typically, attackers first need to identify the code that has been patched in response to a vulnerability. For open-source projects, this is a straightforward task because they can simply review publicly accessible source-code repositories and source-code check-ins relating to the vulnerability. For closed-source applications, they can use a process called "binary diffing" to find which parts of the binary code have changed, narrowing it down to changed functions and, eventually, the vulnerable code. For experienced attackers, binary diffing can be as straightforward as finding differences in source code.

Our research has shown that an experienced attacker can find vulnerable code in just a few minutes (in the case of open-source applications with commented check-ins) to a maximum of a few hours (in the case of binary applications with multiple unrelated changes).

Exploiting the bug

The major factor that affects when an attacker might leverage a vulnerability is the difficulty involved with developing an exploit, or *weaponizing* it. On the one hand, there are vulnerabilities that can only be exploited if the application is in a particular state and/or exploit mitigations are bypassed. For these cases, more research time is required in order to develop a working exploit. But on the other hand, there are vulnerabilities that can easily be exploited and the associated exploit mitigations can be bypassed using previously-used or published techniques.

For an experienced attacker, an exploit code for easy-to-exploit vulnerabilities can usually be developed within a few hours. In contrast, an exploit code for difficult-to-exploit vulnerabilities can take up to a few days, weeks or even months to develop, depending on the level of difficulty.

Unfortunately, in the case of Heartbleed, development of exploit code was straightforward, as evidenced by the release of the exploit code just a day after the disclosure.

Mitigations

The race to deploy patches and develop exploits will not always be won by organizations, and it is prudent to assume that most of the time, attackers may win. However, there are ways for organizations to improve their posture against one-day attacks while patches are being tested and deployed:

- **Apply workarounds.** Check if the vendor provides guidance for a temporary workaround that can help prevent exploitation of the vulnerability. This may involve changing a particular configuration or temporarily disabling a module or a feature where the vulnerability exists or is being used as a vector for exploiting the vulnerability.

- **Block attacks.** Security products—such as intrusion detection or intrusion prevention systems and anti-virus software—can serve as a first line of defense against exploitation of vulnerabilities while patches are being tested and deployed. Vendor programs such as the Microsoft Active Protections Program (MAPP)[9] provide security software providers with early access to vulnerability information so that when Microsoft patches are released, security vendors can immediately release security content that can help detect and block exploits against the recently patched vulnerabilities.

- **Shut down systems temporarily.** Although business leaders may object, another solution is to temporarily shut down or disconnect the affected system while a patch is being tested. This option may be the best way to help prevent the loss of customers' personal or financial information. If a temporary shutdown of a vulnerable system will help stop their information from being stolen, customers will likely understand why a service is temporarily unavailable.

Attackers are opportunistic; they will grab every opportunity to attack when a target is in a weak state. An organization's best defense against one-day attacks is to be ready—to have action plans prepared and mitigations in place when a critical vulnerability is reported.

Vulnerability disclosures in the first half of 2014

What's the state of security beyond Heartbleed? Find out about this year's disclosure trends and proposed changes for vulnerability scoring.

Since 1997, X-Force has been tracking public disclosures of vulnerabilities in software products, such as the Heartbleed disclosure (CVE-2014-0160) from earlier this year. Our X-Force researchers collect software advisories from vendors, read security-related mailing lists, and analyze hundreds of vulnerability web pages where remedy data, exploits and vulnerabilities are disclosed.

In the first half of 2014, we reported just over 3,900 new security vulnerabilities affecting 926 unique vendors. If this trend continues through the end of the year, the total projected vulnerabilities would fall below 8,000 total vulnerabilities for the first time since 2011.

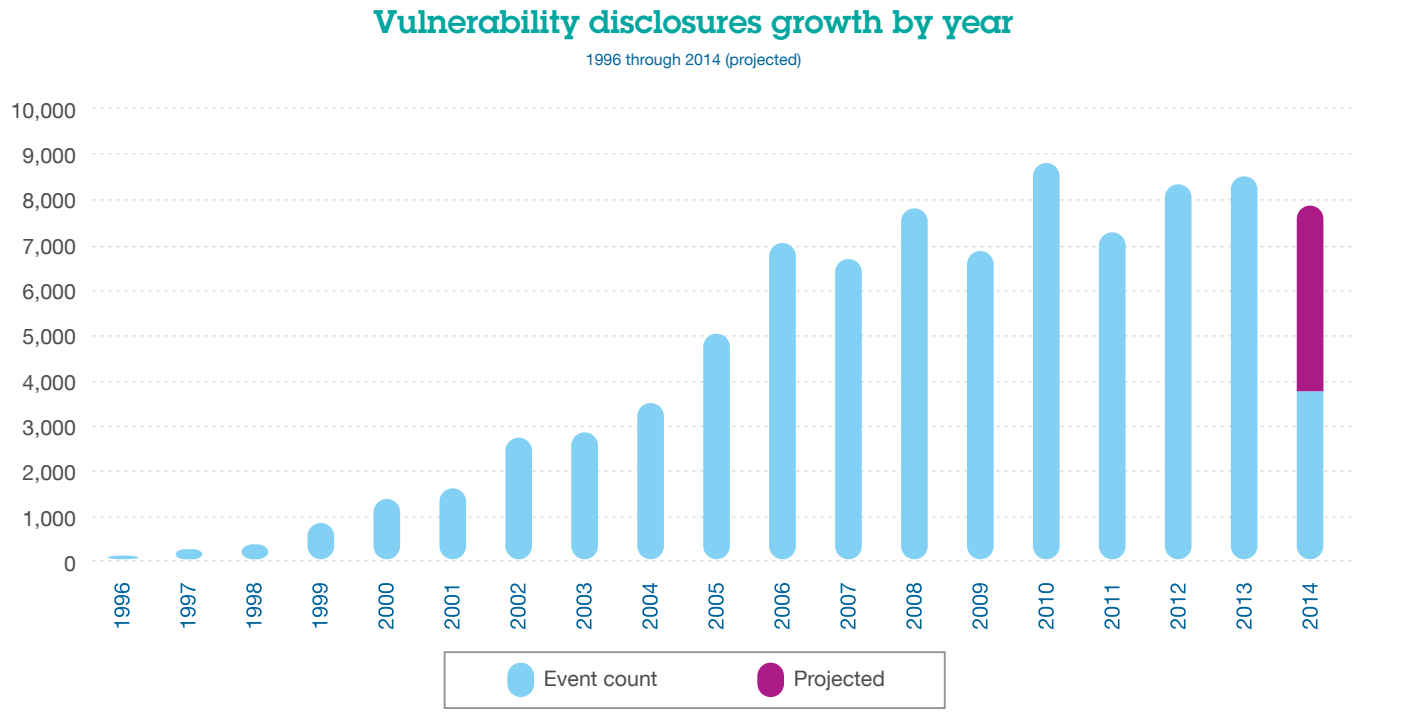

Figure 6. Vulnerability disclosures growth by year, 1996 through 2014 (projected)

It is difficult to point to any one factor that has contributed to the decline in the number of vulnerability disclosures in 2014. However, it is interesting to note that the total number of vendors disclosing vulnerabilities has decreased year over year (1,602 vendors in 2013, compared to 926 vendors in 2014).

Even with the projected decline in the overall number of vulnerability disclosures in 2014, the number of vulnerabilities disclosed by the largest enterprise software vendors remains relatively unchanged year over year (34 percent in 2013, compared to 32 percent in 2014), as shown in Figure 7.

When looking at trends in enterprise software, the X-Force team looks at major software vendors who create the widest variety of enterprise software. We observed that out of thousands of vendors, these companies consistently disclose a significant number of security vulnerabilities. We categorize these vendors in a top 10 group, leaving out the content management system (CMS) vulnerabilities since the majority of those are in third-party plug-ins and add-ons and not widely used as enterprise-level software. These vendors typically have a more comprehensive approach to security that includes policies and practices for properly addressing and responding to security vulnerabilities, which likely leads to a larger number of public vulnerability disclosures.[10]

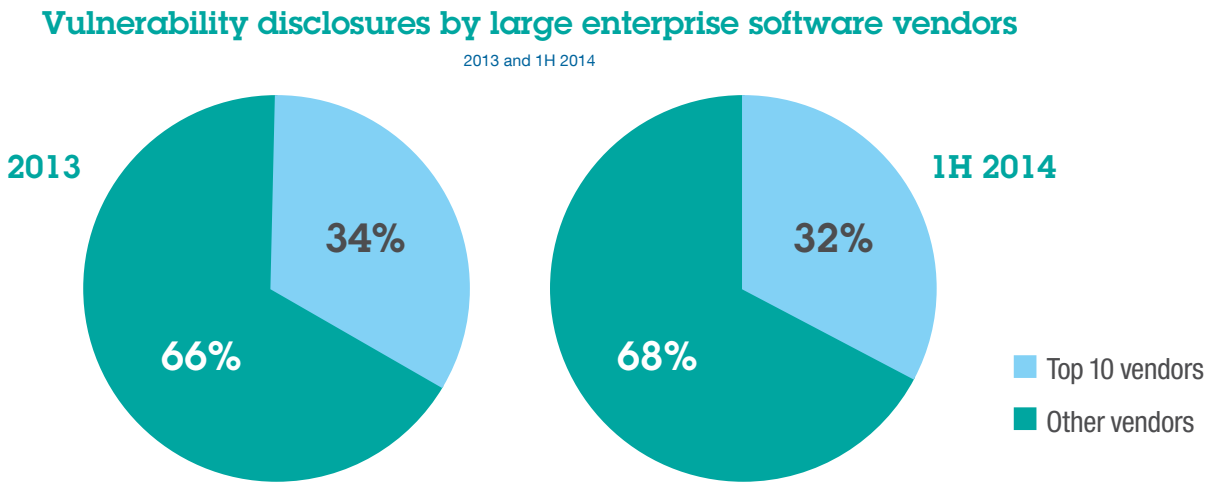

Figure 7. Vulnerability disclosures by large enterprise software vendors, 2013 and 1H 2014

Vulnerabilities in content management systems

Vulnerabilities in CMS continue to be some of the most reported in 2014, accounting for nearly 10 percent of the total vulnerabilities researched. What's more, the largest percentage of CMS vulnerabilities occur in plug-ins or modules written by third-party sources—not by the core CMS vendor.

Many CMS plug-ins are maintained by one person or a small group of people, and they may have infrequent updates (or none at all). Therefore, many plug-ins contain tempting, unpatched security vulnerabilities. Figure 8 shows the lag in patch rates for plug-ins, as compared to core CMS platforms, and the data is unchanged since our 2013 reporting.

While X-Force has previously cautioned against using CMS plug-ins, a new wave of attacks against these platforms was launched in recent months, showing the continued risk. Russian site Yandex reported on a malware dubbed the Mayhem Virus[11] that seeks to compromise web servers through CMS vulnerabilities and brute-force attacks of weak or default credentials.

After these web servers are compromised, they can be used to serve malware or carry out large-scale, high-bandwidth distributed-denial-of-service (DDoS) attacks against other sites and targets. For example, WordPress was used in an amplification DDoS attack in March 2014 that affected more than 162,000 sites.[12] In this case, attackers used the legitimate functionality of the XML-RPC pingback feature to link blog content from different authors to a third-party website.

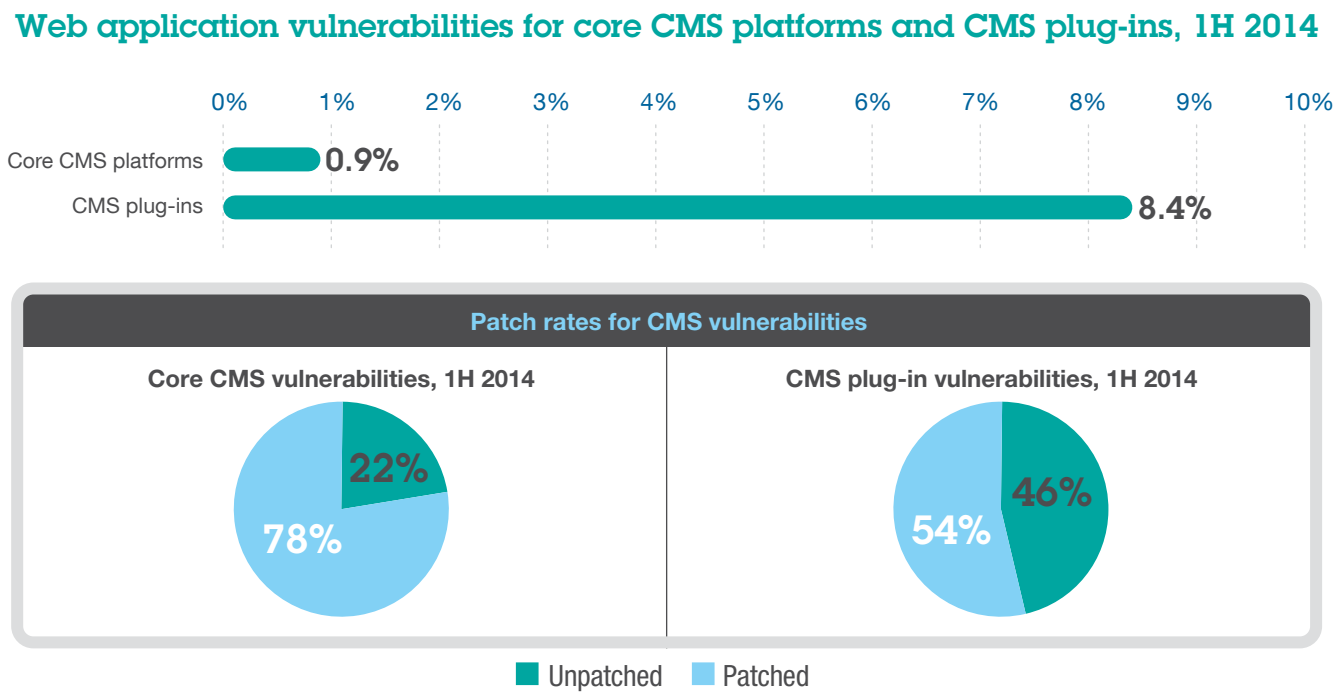

Figure 8. Web application vulnerabilities for core CMS platforms and CMS plug-ins, as a percentage of all disclosures and corresponding patch rates, 1H 2014

CVSS scoring

X-Force uses version 2 of the Common Vulnerability Scoring System (CVSS) to communicate the severity of vulnerabilities. We score vulnerabilities from three different perspectives: as a vulnerability database that tracks third-party vulnerability disclosures, as a security research organization that discovers new vulnerabilities, and as a large software vendor that needs to help customers accurately assess the severity of vulnerabilities within its products. X-Force is currently working alongside other organizations on developing the new CVSS, version 3.[13]

In the scoring of vulnerabilities for the first half of 2014, we found that the majority of issues fall into the CVSS medium-severity range (67 percent), with 24 percent of all vulnerabilities rated critical or high. As shown in Figure 9, these results are unchanged from 2013, and this is the third consecutive year where the majority of vulnerabilities have been rated as medium-level risks.

CVSS base scores, 2012 through 1H 2014

CVSS score	Severity level
10	**Critical** A successful exploit is likely to have catastrophic adverse effects
7.0 – 9.9	**High** A successful exploit is likely to have significant adverse effects
4.0 – 6.9	**Medium** A successful exploit is likely to have moderate adverse effects
0.0 – 3.9	**Low** A successful exploit is likely to have limited adverse effects

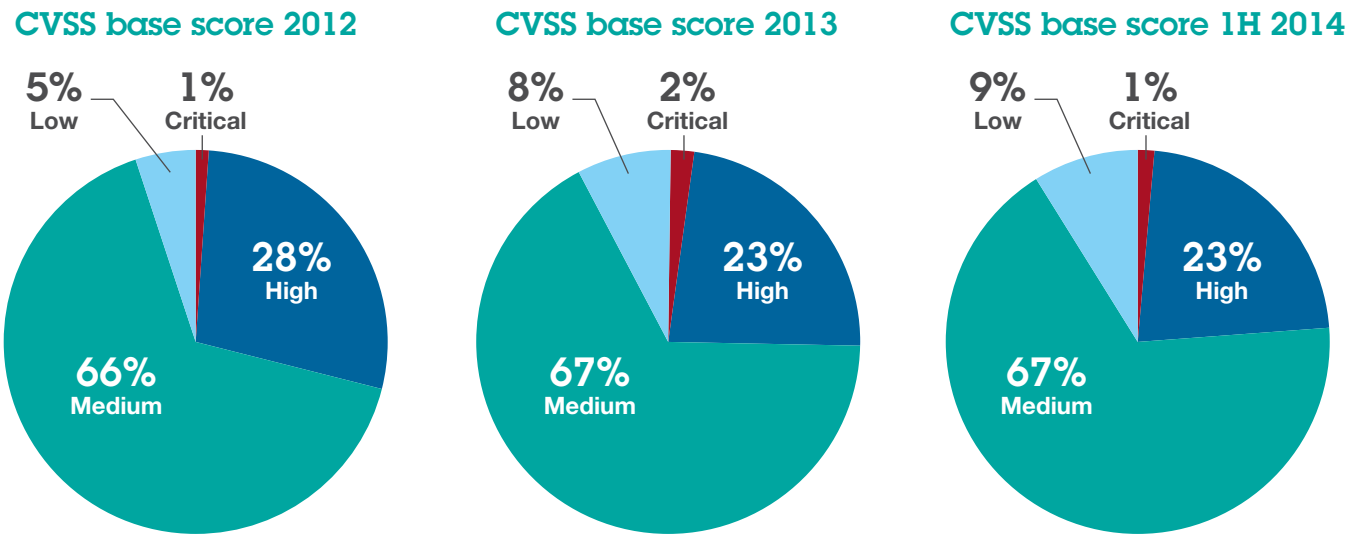

Figure 9. CVSS base scores, 2012 through 1H 2014

Many in the industry, including security analysts, corporate incident response teams and enterprise software consumers, have become dissatisfied with scoring inconsistencies that often occur across different organizations. Sometimes the inconsistencies are the result of the subjectivity that can go into how an individual or organization scores vulnerabilities, but they can also result from some of the inherent flaws in the current CVSS standard and a lack of clear guidelines on how to objectively assess certain types of vulnerabilities. As a result, the CVSS score often fails to reflect the true risk a vulnerability may pose to an organization, causing an overall loss of confidence in the CVSS score as an accurate and reliable measure of risk.

To learn more about the concerns surrounding CVSS and the development of the new CVSS, version 3, refer to the CVSS v3 Development website, hosted by the Forum of Incident Response and Security Teams (FIRST).[13]

The most obvious example of how some CVSS scores do not always represent true risk and impact to an organization is the Heartbleed vulnerability. As mentioned earlier in this report, Heartbleed was disclosed in April 2014, but it had actually existed for two years. This vulnerability received a CVSS base score of 5.0, which falls into the medium-risk level—along with 67 percent of all other vulnerabilities reported during the first half of 2014.

However, with the number of products impacted, the time and attention IT teams spent patching systems and responding to customer inquiries, as well as the potential sensitivity of data exposed, the true impact of the Heartbleed vulnerability was greater than the CVSS base score would indicate. This also brings to question what other vulnerabilities fell into the medium-risk category (CVSS base score 4.0 to 6.9) that may have been disregarded by organizations, but that also had potential large-scale impacts similar to Heartbleed.

Final thoughts on the first half of 2014

Although overall vulnerability numbers are down for the first half of 2014, the impact to the top 10 enterprise software vendors remains consistent. It is uncertain at this point whether this trend will continue through the end of the year as attackers continue to seek higher impact/higher potential reward targets or whether we will see an increase in the second half of the year in the number of disclosed vulnerabilities against smaller vendors and components, such as CMS plug-ins.

X-Force also anticipates the release and adoption of CVSS version 3 to help foster more consistency in risk assessment across organizations and more confidence in the use of CVSS as one of the primary components within an organization's overall incident response plan. This way, when disclosures such as Heartbleed occur in the future, the industry as a whole will be better prepared for potential threats.

About X-Force

Advanced threats are everywhere. Help minimize your risk with insights from the experts at IBM.

The IBM X-Force research and development team studies and monitors the latest threat trends including vulnerabilities, exploits, malware, spam, phishing and malicious web content. The research team includes a variety of skill sets and backgrounds, leveraging IBM acquisitions in the Internet security market—including Internet Security Systems, IBM Trusteer™[14] and IBM Security AppScan®—and combining them with intelligence from active network monitoring from the IBM Managed Security Services group. In addition to advising customers and the general public about emerging and critical threats, X-Force also develops security content and protection techniques to help protect IBM customers from these threats.

IBM Security collaboration

IBM Security represents several brands that provide a broad spectrum of security competency:

- The IBM X-Force vulnerability research and development team discovers, analyzes, monitors and records a broad range of computer security threats, vulnerabilities, and the latest trends and methods used by attackers. Other groups within IBM use this rich data to develop protection techniques for our customers.
- The IBM Trusteer product family delivers a holistic endpoint cybercrime prevention platform that helps protect organizations against financial fraud and data breaches. Hundreds of organizations and tens of millions of end users rely on these products from IBM Security to protect their web applications, computers and mobile devices from online threats (such as advanced malware and phishing attacks).
- The IBM X-Force content security team independently scours and categorizes the web by crawling, independent discoveries, and through the feeds provided by IBM Managed Security Services.
- IBM Managed Security Services is responsible for monitoring exploits related to endpoints, servers (including web servers) and general network infrastructure. This team tracks exploits delivered over the web as well as via other vectors such as email and instant messaging.
- IBM Professional Security Services delivers enterprise-wide security assessment, design and deployment services to help build effective information security solutions.
- IBM QRadar® Security Intelligence Platform offers an integrated solution for security intelligence and event management (SIEM), log management, configuration management, vulnerability assessment and anomaly detection. It provides a unified dashboard and real-time insight into security and compliance risks across people, data, applications and infrastructure.
- IBM Security AppScan enables organizations to assess the security of web and mobile applications, strengthen application security program management and achieve regulatory compliance by identifying vulnerabilities and generating reports with intelligent fix recommendations to ease remediation. IBM Hosted Application Security Management service is a cloud-based solution for dynamic testing of web applications using AppScan in both pre-production and production environments.

Contributors

Producing the IBM X-Force Threat Intelligence Quarterly is a dedicated collaboration across all of IBM. We would like to thank the following individuals for their attention and contribution to the publication of this report.

Contributor	Title
Brad Sherrill	Manager, IBM X-Force Threat Intelligence Database
John Kuhn	Senior Threat Researcher, IBM Managed Security Services
Leslie Horacek	Manager, IBM X-Force Threat Response
Lyndon Sutherland	Threat Intelligence Analyst, IBM Managed Security Services
Mark Vincent Yason	Senior Threat Researcher, IBM X-Force Advanced Research
Nick Bradley	Practice Lead, Threat Research Group
Pamela Cobb	Worldwide Market Segment Manager, IBM X-Force and IBM Security Threat Portfolio
Robert Freeman	Manager, IBM X-Force Advanced Research
Scott Moore	Software Developer, Team Lead, IBM X-Force Threat Intelligence Database
Thomas Van Tongerloo	Senior Cyber Threat Analyst, IBM Managed Security Services
Troy Bollinger	Threat Intelligence Analyst, IBM Managed Security Services

For more information

To learn more about IBM X-Force, please visit:
ibm.com/security/xforce/

[1] Chris Poulin, "What to Do to Protect against Heartbleed OpenSSL Vulnerability," *IBM Security Intelligence Blog*, 10 April 2014. http://securityintelligence.com/heartbleed-openssl-vulnerability-what-to-do-protect/

[2] "The Heartbleed Bug," *codenomicon*, 29 April 2014. http://heartbleed.com

[3] "Untitled," *pastebin*, Accessed 25 July 2014. http://pastebin.com/qyXE7myF

[4] Christopher Glyer and Chris DiGiamo, "Attackers Exploit the Heartbleed OpenSSL Vulnerability to Circumvent Multi-factor Authentication on VPNs," *Mandiant M-unition Blog*, 18 April 2014. https://www.mandiant.com/blog/attackers-exploit-heartbleed-openssl-vulnerability-circumvent-multifactor-authentication-vpns/

[5] "Statement by the Commissioner of the Canada Revenue Agency on the Heartbleed Bug," *Reuters*, 14 April 2014. http://uk.reuters.com/article/2014/04/14/idUKnMKWbNWG8a+1dc+MKW20140414

[6] "The Heartbleed security breed - and what to do," *Mumsnet*, 16 April 2014. http://www.mumsnet.com/info/the-heartbleed-security-breach-to-do

[7] Michael Schierl, "CVE-2012-1723 – Oracle Java Applet Field Bytecode Verifier Cache Remote Code Execution," Accessed 25 July 2014. http://schierlm.users.sourceforge.net/CVE-2012-1723.html

[8] Kafeine, "Inside Blackhole Exploits Kit v1.2.4 - Exploit Kit Control Panel," *Malware don't need Coffee*, 22 July 2012. http://malware.dontneedcoffee.com/2012/07/inside-blackhole-exploits-kit-v124.html

[9] "Microsoft Active Protections Program," *Security TechCenter*, Accessed 25 July 2014. http://technet.microsoft.com/en-US/security/dn467918

[10] John Lucassen, "Are Vendors Doing What Is Needed to Mitigate Security Vulnerabilities?" *IBM Security Intelligence Blog*, 30 June 2014. http://securityintelligence.com/are-vendors-doing-what-is-needed-to-mitigate-security-vulnerabilities/#.U9FWSrG7IUi

[11] Swati Khandelwal, "Mayhem — A New Malware Targets Linux and FreeBSD Web Servers," *The Hacker News*, 24 July 2014. http://thehackernews.com/2014/07/mayhem-new-malware-targets-linux-and_24.html

[12] Ryan Barnett, "More than 162,000 WordPress sites used in DDoS attack," *Trustwave SpiderLabs*, 12 March 2014. http://blog.spiderlabs.com/2014/03/wordpress-xml-rpc-pingback-vulnerability-analysis.html

[13] Seth Hanford, "Common Vulnerability Scoring System, V3 Development Update," *FIRST*, June 2014. http://www.first.org/cvss/v3/development

[14] Trusteer, Ltd. was acquired by IBM in September of 2013.

Please Recycle

IBM Corporation
Software Group
Route 100
Somers, NY 10589

Produced in the United States of America
August 2014

WGL03057-USEN-00