RED HAT PRODUCT SECURITY RISK REPORT: 2016

This report explores the state of security risk for Red Hat® Products for calendar year 2016. We look at key metrics, specific vulnerabilities, and the most common ways users of Red Hat Products were affected by security issues. When we refer to a "Product" in this report, we mean a Red Hat offering listed at https://access.redhat.com/products.

Our methodology is to look at the vulnerabilities we addressed and their severity, then look at which issues were of meaningful risk, and which were exploited. All of the data used to create this report is available from public data maintained by Red Hat Product Security.

Red Hat Product Security assigns a Common Vulnerabilities and Exposures (CVE) name to every security issue we fix. If we fix a bug that later turns out to have had a security implication, we go back and assign a CVE name to that issue. Every CVE fixed has an entry in our public database in the Red Hat Customer Portal, as well as a public bug report with more technical detail. In this report, we will use "vulnerabilities" and "CVEs" interchangeably.

Note: With vulnerability counts, you can compare issues by Red Hat Products or dates because we use a consistent methodology to allocate names and score severity. You should not use vulnerability count data (such as the number of CVEs addressed) to compare our Products with any other company, because assigning and reporting practices vary considerably. Even among products from different Linux® vendors, the same CVE can have different effects, depending on how the product is built or integrated.

VULNERABILITIES

Across all Red Hat Products, and for all issue severities, we fixed more than 1,300 vulnerabilities by releasing more than 600 security advisories in 2016. That may seem like a lot of vulnerabilities, but for a given organization only a subset of those issues apply to Products and versions in use. Even then, within a product such as Red Hat Enterprise Linux, not every package is installed in a default or even likely installation.

We rate vulnerabilities using a four-point scale intended to quickly show how concerned we are about each security issue. We designed the scale to align as closely as possible with similar scales from other open source groups and enterprise vendors. The severity levels are intended to help users determine which issues matter the most. This prioritized risk assessment helps customers understand their exposure and better schedule upgrades to their systems, in an effort to improve their ability to assess how much risk each issue creates in their unique environment.

To aid customers who use the Common Vulnerability Scoring System (CVSS), we also publish CVSS scores for every vulnerability addressed. We started using CVSS v3 during 2016 and switched to that version by default at the start of 2017. However, CVSS scores have some limitations, and we do not use CVSS to prioritize vulnerabilities.

Our four-point scale rates vulnerabilities as Low, Moderate, Important, or Critical.

Critical vulnerabilities pose the most risk to an organization. By definition, a Critical vulnerability could be exploited remotely and automatically by a worm. However Red Hat, like other vendors, also stretches the definition to include flaws that affect web browser or plug-in users who visit malicious (or compromised) websites. These flaws actually account for the majority of the Critical issues fixed. If you're using a Red Hat product that does not include desktop, for example, you'll likely be affected by significantly fewer Critical issues.

Figure 1 gives advisory and vulnerability counts for a selected subset of Products and Product families. A single Red Hat advisory may fix multiple vulnerabilities across multiple versions of a product. As a general rule, the vulnerability count is indicative of the amount of effort a customer will expend to gain an understanding of the issues and fixes. The number of advisories reflects the effort customers can expect to expend deploying updates.

Red Hat security advisories and vulnerabilities for 2016

PRODUCT	CRITICAL ADVISORIES	IMPORTANT ADVISORIES	CRITICAL VULNERABILITIES	IMPORTANT VULNERABILITIES
All products	110	270	318	255
Red Hat Enterprise Linux 5,6,7	38	90	50	89
› Red Hat Enterprise Linux 6 (all)	30	43	50	66
› Red Hat Enterprise Linux 6 Server (Default)	7	17	9	22
› Red Hat Enterprise Linux Supplementary (5,6,7)	41	29	270	122
Red Hat JBoss® Middleware (all JBoss products)	21	34	2	15
Red Hat Storage (all storage products)	1	0	1	1
Red Hat OpenStack® Platform	0	31	0	9

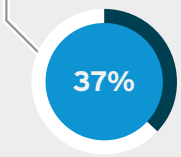

37% of all Critical advisories were for Red Hat Enterprise Linux supplementary channels (Java, Flash, and others)

100% of Red Hat Enterprise Linux Critical issues had updates the same or next day after public knowledge

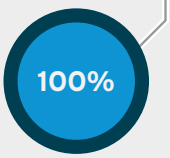

Figure 1

Red Hat Enterprise Linux 6 appears several times in Figure 1. During a Red Hat Enterprise Linux 6 installation, the user gets a choice of installing the default selection of packages or a custom selection. In 2016, we issued 7 Critical and 17 Important security advisories applicable to the default packages (those where the user installs a default server and does not add any additional packages or layered Products). We issued an additional 27 advisories that addressed Moderate or Low issues. Those numbers are similar to 2015, when we issued 6 Critical and 17 Important advisories.

Where advisories outnumber vulnerabilities for a given Product (as was the case for OpenStack), the same vulnerability may have affected multiple supported versions of the Product. In those cases, each version got its own security advisory.

In 2016, we issued 110 Critical security advisories addressing 318 Critical vulnerabilities. We released updates to address 76% of Critical issues on the day they became public or one day later. We addressed 98% of them within a week.

Given that Products are released all the time, a year-to-year comparison isn't particularly useful. That said, our 2016 figures were similar to those of 2015, when we addressed 82% of issues the same or next day and 99% of them within a week.

Looking only at issues affecting base Red Hat Enterprise Linux releases, we released 38 Critical security advisories addressing 50 Critical vulnerabilities in 2016. Of those issues, 100% had updates available the day they became public or one day later. In 2015, there were 46 Critical advisories fixing 61 Critical issues, with 96% fixed by the next day.

For Red Hat Enterprise Linux, far fewer Critical vulnerabilities affect server installations. That's because most Critical vulnerabilities occur in browsers or browser components. One way to reduce risk when using our modular Products is to make sure you install the right variant and review the package set to remove packages you don't need.

VULNERABILITY TRENDS

Red Hat continually releases new Products, so the number of vulnerabilities we address generally increases each year. The numbers were similar, however, between 2015 and 2016. When looking at a specific Product, we find that we fix fewer vulnerabilities over time because of our security fix backporting practices.

We use the term "backporting" to describe the action of taking a fix for a security flaw out of the most recent version of an upstream software package and applying that fix to an older version of the package we distribute. Our backporting efforts permit us to deploy automated updates to customers with less risk.

Red Hat security timeline

Vulnerabilities increase as we deliver more products and versions.

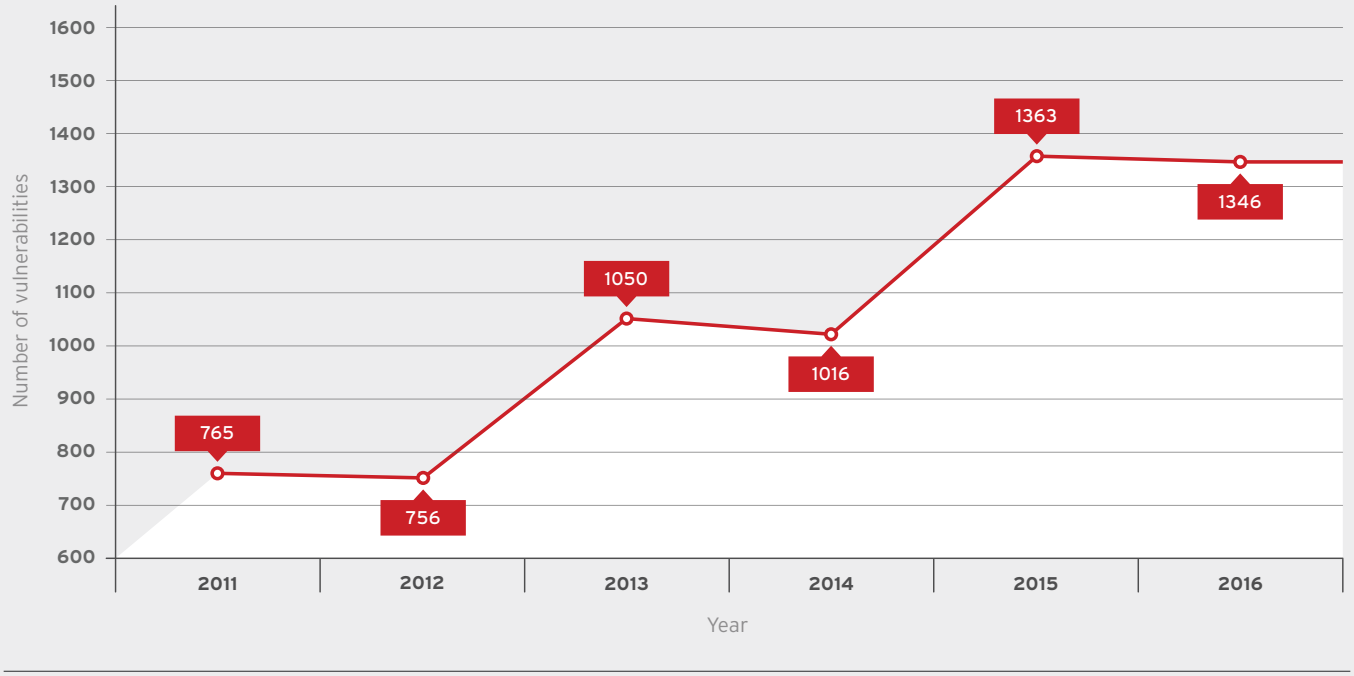

Figure 2

The trends can be investigated using our public data, and we occasionally do risk reports that delve into a given Product and version. For an example, see our Red Hat Enterprise Linux 6.5 to 6.6 risk report.

WHICH ISSUES WERE BRANDED, AND WHICH REALLY MATTERED IN 2016

In 2014, the OpenSSL Heartbleed vulnerability started a trend of branding vulnerabilities, changing the way security vulnerabilities that affected open source software were being reported and perceived. Vulnerabilities are found and fixed all the time. A vulnerability may get a catchy name, fancy logo, or media attention. That doesn't mean it poses a material risk to users.

So, let's take a chronological tour through 2016 to see which issues were branded or generated attention and, more importantly, which issues we think actually mattered for Red Hat customers because they were high risk.

kernel keychain overflow (19 Jan 2016) CVE-2016-0728

Severity: Important

A bug was found in the Linux kernel keyring facility that could have led to local privilege escalation. Privilege escalation flaws in the kernel like this require an untrusted local user (or other exploit that gains an attacker an untrusted account), and they can lead to gaining root privileges.

This issue only affected the Linux kernels as shipped with Red Hat Enterprise MRG 2 and Red Hat Enterprise Linux 7. We supplied mitigations for the issue through a systemtap script. Final updates were available within a week of the issue becoming public. This was one of the five most-viewed issues in the Red Hat Customer Portal. Working public exploits are available for this issue.

glibc getaddrinfo overflow (16 February 2016) CVE-2015-7547

Severity: Critical

A bug was found in the libresolv function of glibc, used to perform Domain Name Server (DNS) lookups. A remote attacker could have created specially crafted DNS responses, causing libresolv (when called in a certain way) to crash or potentially execute code with the permissions of the user running the library.

This issue affected the versions of glibc shipped with Red Hat Enterprise Linux 6 and 7. We released updates the day this issue became public. It received the second-highest number of views in the Red Hat Customer Portal in 2016. A proof-of-concept exploit is available for this issue, but due to various factors it is unlikely to see widespread exploitation.

"DROWN" (1 March 2016) CVE-2016-0800

Severity: Important

DROWN was a branded bug affecting the SSLv2 protocol. SSLv2 was found to be vulnerable to the Bleichenbacher RSA padding oracle attack, which can be used to decrypt RSA cipher text.

This issue affected the OpenSSL library as shipped across many Red Hat Products. We issued Red Hat Enterprise Linux updates the same day, with some Products following later. We gave this issue enhanced coverage in the Red Hat Customer Portal, including a banner on all pages and a customer outreach email campaign. It became one of the five most-viewed issues of 2016. The researchers who found the issue indicated they were able to perform attacks, but this requires observation of a large number of encrypted handshakes. We are not aware of any meaningful exploitation of this issue.

"Badlock" (12 April 2016) CVE-2016-2118

Severity: Important

Badlock was a branded bug affecting Samba remote protocols. A man-in-the-middle attacker could have gained access to passwords or other sensitive information from the Security Account Manager database.

Red Hat Enterprise Linux (all versions) and Red Hat Storage were affected. We released updates the day this issue became public. We gave this issue enhanced coverage in the Red Hat Customer Portal, including a banner on all pages and a customer outreach email campaign. We are not aware of any public exploits for this issue.

Samba DCE/RPC critical (12 April 2016) CVE-2015-5370

Severity: Critical

Multiple bugs were found in Samba's DCE/RPC protocol implementation. A remote, authenticated attacker could have used these flaws to cause a denial of service against the Samba server or possibly execute arbitrary code as root.

Red Hat Enterprise Linux (all versions) and Red Hat Storage were affected. Updates were available the day the issue became public. We addressed the issue simultaneously with "Badlock," but this one is of higher severity and risk. We are not aware of any public exploits for this issue.

"ImageTragick" (3 May 2016) CVE-2016-3714

Severity: Important

ImageTragick was a branded bug affecting ImageMagick, and it related to the software not properly sanitizing certain input. An attacker could create an image that, when processed by an application using ImageMagick or the command line utilities, would lead to arbitrary execution of shell commands with the privileges of the user running the application.

Red Hat Enterprise Linux versions 6 and 7 were affected. We released updates within a week of the issue becoming public. We gave ImageTragick enhanced coverage in the Red Hat Customer Portal, including a banner on all pages. A working public exploit with a Metasploit module exists for this issue.

Exploiting this issue remotely depends on a given application parsing an untrusted image in a vulnerable way, and we're therefore not aware of widespread exploitation of this issue. However, since ImageMagick is very commonly used to convert images, both through libraries and the command line "convert" utility, we've kept this as a meaningful issue.

Overcloud images default password (13 June 2016) CVE-2016-4474

Severity: Important

An issue was discovered in the build process for the overcloud images, as used by OpenStack Director, resulting in images having a default root password.

Red Hat OpenStack 7.0 and 8.0 Director were affected. Updates were available the same day the issue was disclosed. However, remote root access (such as via SSH) is disabled by default, reducing the impact of this issue.

Authorization bypass in JGroups (23 June 2016) CVE-2016-2141

Severity: Critical

JGroups, a middleware messaging library, did not require necessary headers from new nodes joining a cluster. An attacker could have used this bug to bypass security restrictions to send and receive messages within the cluster, leading to information disclosure, message spoofing, or further possible attacks.

Various Red Hat JBoss Middleware Products were affected. We released updates for some Products the day the issue became public. Other product updates followed in the coming weeks. We're not aware of widespread exploitation of this issue.

"Challenge Ack" (12 July 2016) CVE-2016-5696

Severity: Important

A bug in the Linux kernel networking implementation could have allowed attackers to terminate or inject payloads into unsecured Transmission Control Protocol (TCP) connections.

This issue affected Red Hat Enterprise Linux 6, 7, and MRG 2 Products. We released a configuration option mitigation the day it became public. Final updates to correct this were available at the next update cycle, six weeks after the issue became public. Challenge Ack was one of the five most-viewed issues in the Red Hat Customer Portal. Public exploits are available for it.

"httpoxy" (18 July 2016) CVE-2016-5387

Severity: Important

httpoxy was a branded bug caused by an unexpected interaction between web servers and applications. Some web servers and frameworks pass incoming headers to scripts in environment variables, and some applications would use a proxy password as an environment variable. A remote attacker could have sent a carefully crafted header that set this proxy variable, potentially intercepting any requests being made by the applications.

This issue affected various Red Hat Products. Updates for Red Hat Enterprise Linux releases were available the same day the issue became public, with middleware updates several weeks later. We gave this issue enhanced coverage in the Red Hat Customer Portal, including a banner on all pages. No special exploit is needed to trigger this issue, but the outcome depends entirely on the applications targeted. We're not aware of widespread exploitation of this issue.

"SWEET32" (24 August 2016) CVE-2016-2183

Severity: Moderate

SWEET32 was a branded bug affecting OpenSSL and Network Security Services (NSS) when the Data Encryption Standard/Triple DES (DES/3DES) cipher was used as part of the Transport Layer Security/Secure Sockets Layer (TLS/SSL) protocol. A man-in-the-middle attacker could use this flaw to recover plaintext data by capturing large amounts of encrypted traffic if the communication used a DES/3DES-based ciphersuite.

This issue affected various Red Hat Products. An update for OpenSSL, released a month after the issue became public, mitigated this issue by lowering the priority of DES cipher suites. As with many TLS issues, this issue is hard to exploit as it requires a man-in-the-middle attack. We're therefore not aware of active exploitation of it.

BIND DoS (27 September 2016) CVE-2016-2776

Severity: Important

A bug in the Berkeley Internet Name Domain (BIND) allowed a remote attacker to cause named to exit, causing a denial of service.

This issue affected all versions of Red Hat Enterprise Linux. Updates were available the day after the issue became public. A public exploit and a Metasploit module exist for this issue.

Through the year, there were other similar BIND issues that could result in BIND crashing, leading to a denial of service: CVE-2015-8704 (January), CVE-2016-1286 (March), and CVE-2016-8864 (November). Although these other issues do not have public exploits, they did gain customer attention, being amongst the 20 most-viewed issues in the Red Hat Customer Portal.

"Dirty COW" (19 October 2016) CVE-2016-5195

Severity: Important

A bug in the Linux kernel led to privilege escalation. A local, unprivileged user could use this flaw to run arbitrary code as root.

This branded issue was unique in that the name, domain, and logo were created by a third party after the issue was made public. These elements were not related to the upstream project or finder of the issue. This was the most-viewed security issue on the Red Hat Customer Portal for 2016.

Red Hat Enterprise Linux, MRG, and Virtualization Products were affected. Mitigations were available the same day the issue became public. Kpatch updates were also available for Red Hat Enterprise Linux 7 customers prior to the final updates, which were released one week after the issue became public. Working public exploits and a Metasploit module are available for this issue.

mAlert OpenSSL DoS (24 October 2016) CVE-2016-8610

Severity: NA

An issue was found in the TLS protocol relating to handling of Alert packets. An attacker could send a large number of packets during a handshake, potentially causing a denial of service against servers that are not able to gracefully handle them.

This was not treated as a security issue by upstream cryptographic libraries such as OpenSSL, and servers such as Apache httpd were not affected. We released no updates for this issue.

Various bugs in Firefox and supplementary channels (various dates)

Severity: Critical

By definition, a Critical vulnerability is one that could be exploited remotely and automatically by a worm. However we, like other vendors, also stretch that definition to include web browser or plug-in flaws that affect users visiting malicious or compromised websites. Red Hat Enterprise Linux releases include the Firefox browser. We rate these memory flaws in Firefox and the crypto library NSS Critical because, while they are unproven to lead to code execution, there is a possibility that one may be exploited. In 2016 this led to 10 advisories addressing 37 Critical vulnerabilities.

Red Hat provides some packages that are not open source software in supplementary channels for users of Red Hat Enterprise Linux. This channel contains software such as Adobe Flash Player, IBM Java, Oracle Java, and Chromium browser.

A large number of Critical flaws affected these packages. For example, for Adobe Flash Player in 2016, we issued 11 Critical advisories to address more than 250 Critical vulnerabilities. As these projects release security updates, we ship appropriate updated packages to customers.

A tour of vulnerabilities in 2016

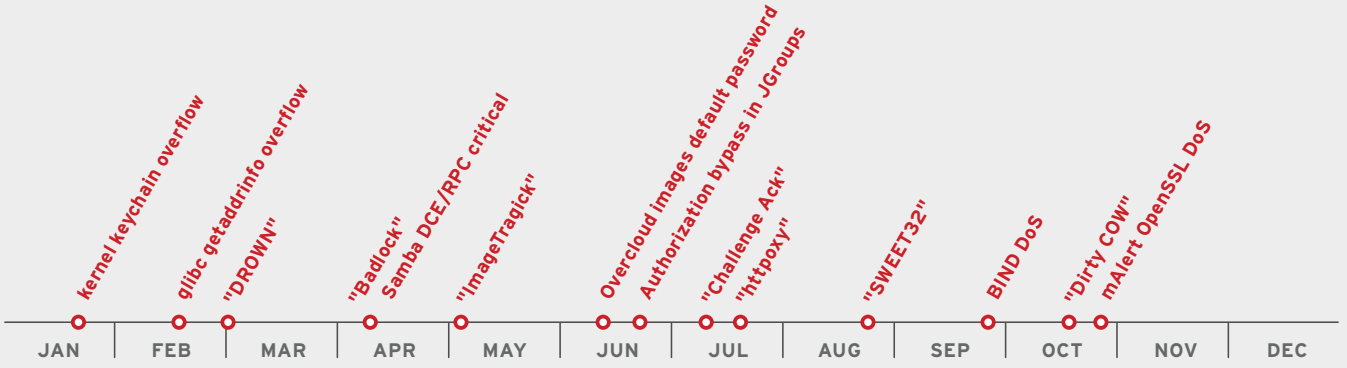

Figure 3

We examined the issues in this section because they were meaningful. Some are issues of high severity and likely to be exploited (or already have a public working exploit). Some were highly visible because of branding (with a name, logo, enhanced media attention), regardless of their severity. Some issues had both high severity and high visibility.

Branded issues / issues of high risk

A branded issue isn't always one of high risk.

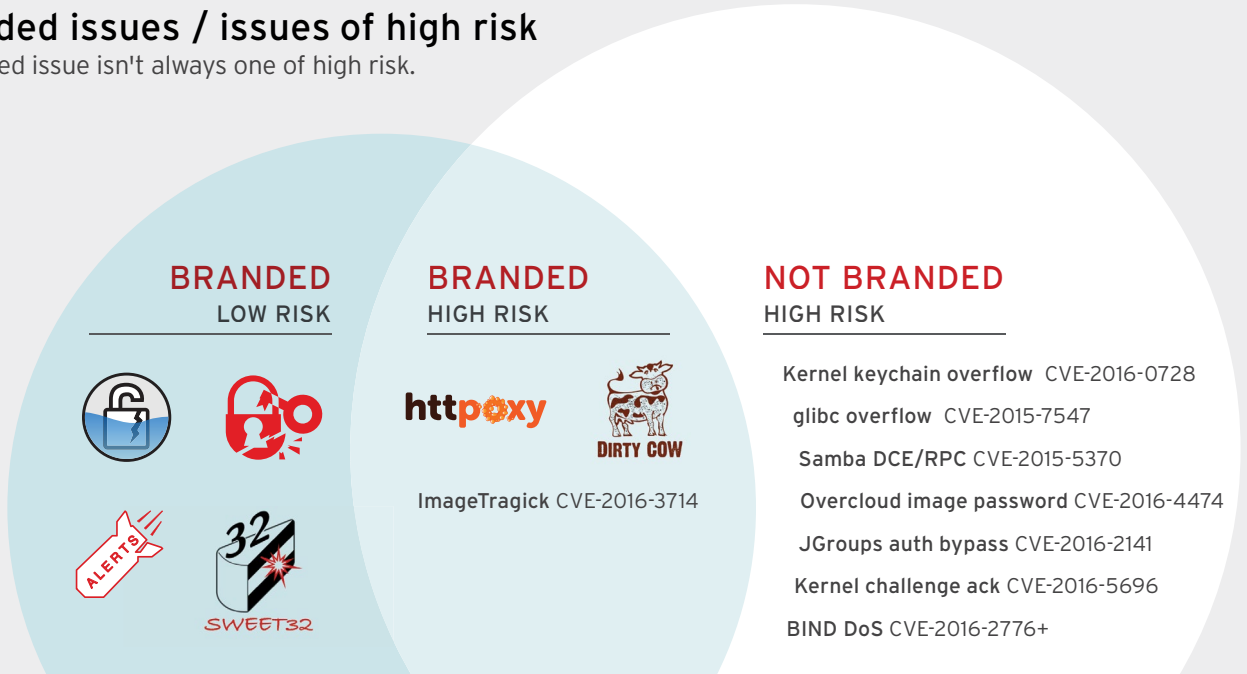

Figure 4

LOWER-RISK ISSUES WITH INCREASED CUSTOMER ATTENTION

Another way we gauge customer concern around an issue is to measure web traffic, specifically views for each CVE page in the Red Hat Customer Portal.

Gauging customer concern

Most-viewed vulnerabilities from the Red Hat CVE database

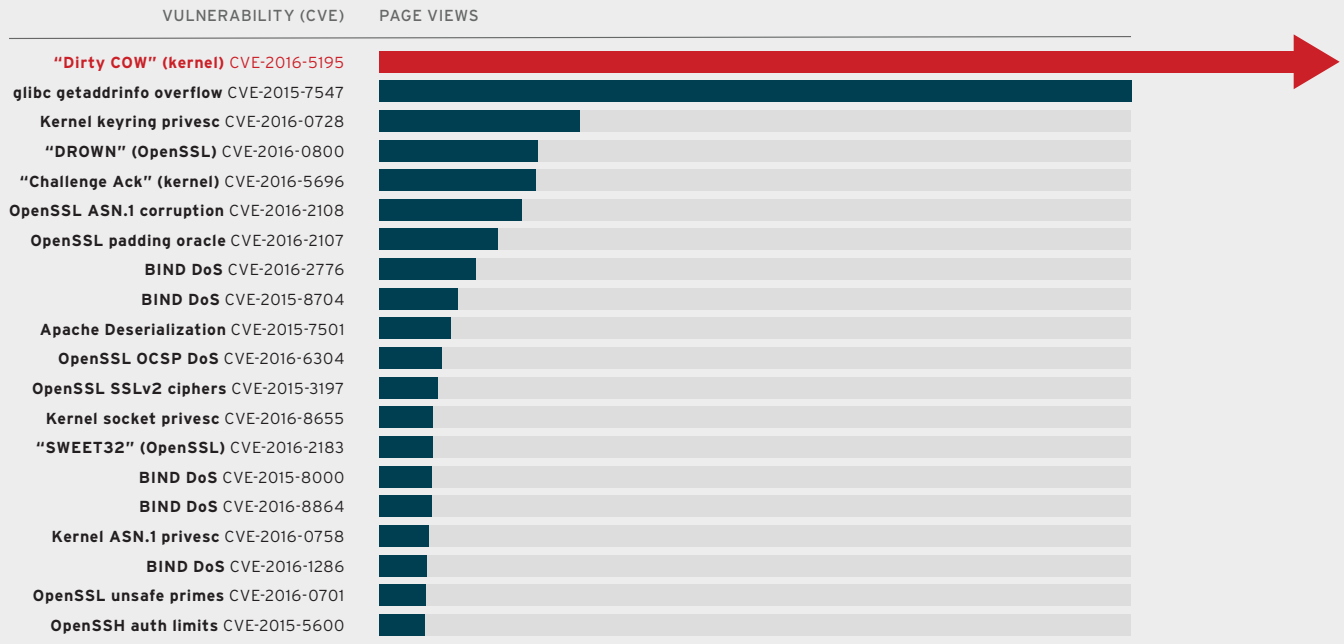

Figure 5

"Dirty COW" was by far the most-viewed issue in the Red Hat Customer Portal for 2016.

Figure 5 gives an indication of customer interest in given vulnerabilities. Many of the top issues appear elsewhere in this report. Of the rest, the most-viewed issues were:

- A November 2015 issue in Java Object Serialization (CVE-2015-7501) that affected many JBoss Middleware Products. We previously covered this issue in the 2015 report.
- A flaw in OpenSSH (CVE-2015-5600) that did not affect Red Hat Products in a default configuration and was rated Low impact. It was also a top hit in 2015.
- Two flaws in the Linux kernel: CVE-2016-8655, which affected Red Hat Enterprise Linux 7 and MRG 2 but did not allow privilege escalation in default or common use, and CVE-2016-0758, which affected Red Hat Enterprise Linux 7 and MRG 2 and could potentially lead to privilege escalation. There are no known exploits for these issues.

- Various flaws in OpenSSL. After high-profile issues such as Heartbleed and Poodle in previous years, OpenSSL issues tend to generate increased customer interest, independent of the actual severity or risk:
  - Two issues we rated Important: CVE-2016-2108, relating to ASN.1 corruption, where a malicious certificate could cause a buffer overflow, and CVE-2016-6304, relating to Online Certificate Status Protocol (OCSP) stapling support, which is unlikely to affect many applications.
  - Two issues we rated Moderate or Low: CVE-2016-2107, relating to padding oracles; and CVE-2015-3197, regarding SSLv2 ciphers, which requires a malicious client.
  - One issue that did not affect Red Hat Products, CVE-2016-0701, regarding unsafe primes.

THE OPEN SOURCE SUPPLY CHAIN

Red Hat Products are based on open source software. Some Red Hat Products contain several thousand individual packages, each of which may be based on separate, third-party software from upstream. While Red Hat engineers play a part in many upstream components, handling and managing vulnerabilities across thousands of third-party components is non-trivial.

Red Hat has a dedicated Product Security team that monitors issues affecting Red Hat Products and works closely in relationships with upstream projects. In 2016, we investigated more than 2,600 vulnerabilities that potentially affected parts of our Products, leading to fixes for 1,346 vulnerabilities. This is up by 30% from 2015, when the team investigated 2,000 vulnerabilities.

Every one of those 2,600+ vulnerabilities is tracked in the Red Hat bugzilla tool and is publicly accessible. Each vulnerability has a master bug including the CVE name as an alias and metadata including the dates we found out about the issue, its severity, and its source. Issues that are not yet public still get an entry in bugzilla, but they are initially private to Red Hat. Once an issue becomes public, the associated bugzilla is updated and made public.

This data is also available via the Red Hat Security Data API. We use this data to create metrics and spot trends. In addition to tracking vulnerabilities themselves, we quantify the ways we learn about potential issues. We do this by looking at the "source" metadata to see how we first heard about each of the issues we fixed in 2016.

How Red Hat finds vulnerabilities

Sources of vulnerabilities fixed by Red Hat in 2016

ADVANCE SOURCE	NUMBER OF ISSUES	PERCENT
Internet	882	65.5%
Relationship	172	12.8%
Red Hat	143	10.6%
Individual	75	5.6%
Peer vendors	49	3.6%
CVE	25	1.9%

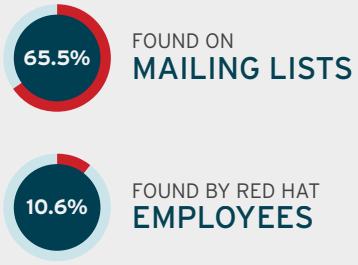

Figure 6

1,346
VULNERABILITIES
FIXED IN 2016

KEY:

- **Internet**: for issues not disclosed in advance, we monitor a number of mailing lists and security webpages of upstream projects
- **Relationship**: issues reported to us via upstream projects, generally in advance of public disclosure
- **Red Hat**: issues found by Red Hat employees
- **Individual**: issues reported to Red Hat Product Security directly by a customer or researcher
- **Peer vendors**: issues reported to us by other open source distributions, through relationships, or via a shared private forum
- **CVE**: if we haven't found out about an issue any other way, we can catch it from the list of public assigned CVE names from Mitre
- **CERT**: issues reported to us from a national Computer Emergency Response Team (CERT) like CERT/CC or CPNI

Red Hat employees find many of the vulnerabilities we fix. We don't wait for others to find flaws for us to fix. Our engineering, quality assurance, and security teams actively look for and find issues. Red Hat employees found 11% of the issues we fixed in 2016, a slight decrease from 12% in 2015. We share these issues back upstream and with other peer vendors (generally via the "distros" shared private forum) if they are risky. In addition to those 143 issues, Red Hat also finds and reports flaws in software that isn't part of a current shipped product or issues that affect other vendors' software.

Relationships matter. When you are fixing vulnerabilities in third-party software, having a relationship with the upstream community makes a big difference. If an upstream community is willing to give information about flaws in advance, we feel a responsibility to give value back for that notification. At Red Hat we do this by reviewing draft advisories, checking patches, and feeding back the results from our quality testing when there is enough time.

Finding, fixing, and sharing issues in such a large software community demands a non-trivial amount of time and effort. If your organization uses open source software that you manage yourself, you need to be able to find out about vulnerabilities that affect those components so you can analyze and remediate them. Vendors without a sizable, dedicated security team often must rely on watching what other vendors do or monitoring on other vulnerability feeds, such as the list of assigned CVE names from Mitre. Red Hat chooses to invest in a dedicated team handling vulnerability notifications to help us find out about issues that affect our Products and build upstream relationships.

EMBARGO AND RELEASE TIMINGS

Vulnerabilities known to Red Hat in advance of being made public are known as "under embargo." This mirrors the way journalists use the term for stories which are not to be published until a specific date and time.

Red Hat Products contain open source components, and therefore Red Hat is often not the only vendor shipping those components. This means that Red Hat is not in sole control of the date each flaw is made public. We find that this tends to lead to much shorter times between when a flaw is first reported and when it becomes public, shortening an attacker's window of opportunity to exploit those flaws.

When we find security issues, we choose to embargo only the ones that we believe pose a material risk. Even then, we use embargos sparingly. If we do choose to embargo an issue due to severity, we share details with the relevant upstreams as well as other peer vendors, working together to address the issues. We talk about this more in our blog post, "The hidden costs of embargos." In 2016, we chose not to embargo 30% of the 143 issues found by Red Hat employees.

For 2016, we knew about 394 (29%) of the vulnerabilities we addressed in advance of them being public, down slightly from 32% in 2015. We expect this figure to vary from year to year. Across all Products and vulnerabilities of all severities known to us in advance, the median embargo was seven days. This is much lower than 2015, when the median embargo was 13 days.

There are many positives to quickly releasing fixes for issues of higher risk, but the drawback to not having a regular patch day is that you need to respond to more issues as they happen. We do suggest embargo dates that avoid weekends and major holidays.

CONCLUSION

This report looked at the security risk to users of Red Hat Products in 2016 by providing metrics around vulnerabilities, highlighting those that were the most severe and those that were exploited, and showing which were branded or gained media attention.

We found roughly as many vulnerabilities in 2016 as in 2015, although the number of issues we found out about in advance of the vulnerability being public did drop slightly. The median embargo length for those was reduced to just 7 days, down from 13 in 2015.

There are other types of security risks, such as malware or ransomware, that we haven't covered in this report. They rely on an attacker having access to a system through an intrusion or by exploiting a vulnerability.

For the last year of vulnerabilities affecting Red Hat Products, the issues of material risk and the issues that got branded do have an overlap, but they certainly don't closely match. Just because an issue gets a name, a logo, or press attention does not mean it's of increased risk. We've also shown some vulnerabilities of increased risk that did not get branded or draw media attention at all.

At Red Hat, our dedicated Product Security team analyzes threats and vulnerabilities against all our Products. We provide advice and updates through the Red Hat Customer Portal. Customers can call on this expertise to help them respond quickly to the issues of material risk, while avoiding the media whirlwind around those that are not.

facebook.com/redhatinc
@redhatnews
linkedin.com/company/red-hat

redhat.com
F6404_0317