

# 2016 Data Breach Investigations Report

89% of breaches had a  
financial or espionage motive.



# 2016 DBIR Contributors

(See Appendix B for a detailed list.)



Mishcon de Reya



CHAMPLAIN COLLEGE | LCDI Leahy Center for Digital Investigation





# Table of Contents

2016 DBIR – Introduction .....	1
Victim demographics.....	3
Breach trends .....	6
Points of focus .....	12
Vulnerabilities .....	13
Phishing .....	17
Credentials .....	20
Incident classification patterns.....	22
Web App Attacks.....	27
Point-of-Sale Intrusions .....	31
Insider and Privilege Misuse .....	35
Miscellaneous Errors .....	40
Physical Theft and Loss .....	43
Crimeware .....	45
Payment Card Skimmers.....	49
Cyber-espionage.....	52
Denial-of-Service Attacks .....	56
Everything Else .....	60
Wrap up.....	62
Varieties of impact .....	64
Appendix A: Post-compromise fraud .....	66
Appendix B: Contributing organizations .....	71
Appendix C: The Taupe Book .....	72
Appendix D: Attack graphs.....	74
Appendix E: Methodology and VERIS resources.....	76
Appendix F: Year in review .....	78

# 2016 DBIR—Introduction

**“It’s like déjà vu, all over again.”**  
–Yogi Berra

Well here we are again, and it is time to take the annual journey into our collection of real-world data breaches and information security incidents from the prior year. We have published this report nine times<sup>1</sup> and we truly appreciate you spending your valuable time with us, whether you have been with us since our humble, pie-chart-centric beginnings or if this is your first read.

We would be remiss if we did not begin by acknowledging the organizations that contributed data (and time) to this publication. Simply stated, we thank you for helping to make this possible. For a full list of contributors, mosey over to Appendix B.

The incident data is the workhorse of this report and is used to build out all the information within the Breach Trends and Incident Classification Patterns sections. We use non-incident security data to paint a fuller picture in the patterns as well as in stand-alone research. Any opportunity to take several organizations’ data and combine them for a research topic was pursued. The Gestalt principles in action!

The nine incident classification patterns we identified back in the 2014 report still reign supreme. And while there are no drastic shifts that have established a show-stopping talking point when looking at the patterns as a whole, we have searched for interesting tidbits in the actions that comprise them.

This year’s dataset is made up of over 100,000 incidents, of which 3,141 were confirmed data breaches. Of these, 64,199 incidents and 2,260 breaches comprise the finalized dataset that was used in the analysis and figures throughout the report. We address the reasons for culling the dataset in Victim Demographics and provide additional details when we discuss motives in Breach Trends. Of course, we would never suggest that every last security event of 2015 is in this report. We acknowledge sample bias, and provide information about our methodology as well as links to resources that we encourage you to look into to help collect and analyze incident data within your own organization, in Appendix E.

We will also acknowledge what isn’t in this report. For those looking for proclamations about this being the year that mobile attacks bring us to our knees or that the Internet of Things (IoT) is coming to kill us all, you will be disappointed. We still do not have significant real-world data on these

**The nine incident classification patterns we identified in 2014 still reign supreme.**

<sup>1</sup> [Nine times? Nine times.](#)

technologies as the vector of attack on organizations.<sup>2</sup> If you feel we are in error, put down the torches and pitchforks and share any breach data that you have. We are always looking for avenues to shine lights into areas in which we may not have sufficient illumination. Also, their absence is not a suggestion to ignore these areas in your risk management decision-making.

The report is designed so you can enjoy it like a prog-rock concept album, from beginning to end, or feel free to bounce around (the room). Enjoy the Breach Trends section for all your figure and chart needs. Get some knowledge on a few of the concepts that stretch across several patterns in our Points of Focus section and for those who want more factoids, pop over to the appendices and give our Taupe Book section a look.

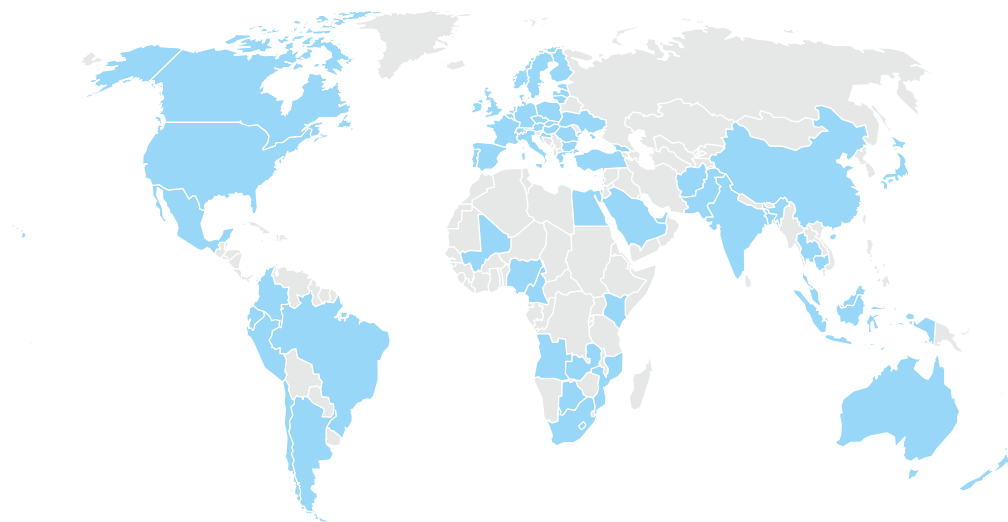
---

<sup>2</sup> Yes, we are aware of the xCode hack, but without confirmed organizations that suffered an attribute loss it will not be an influencer of this report.

## Victim demographics

Before we get into the adversaries behind the incidents and breaches that both underpin this report and keep information security professionals busy, let's acknowledge who is on the receiving end of these attacks. The 2016 report features incidents affecting organizations in 82 countries and across a myriad of industries.

**No locale, industry or organization is bulletproof when it comes to the compromise of data.**



**Figure 1.**

Countries represented in combined caseload.

No locale, industry or organization is bulletproof when it comes to the compromise of data. Some are notably more represented than others and this is not an indictment that the public sector is any less secure than any other industry. As with prior years, the numbers that follow are heavily influenced by US agency reporting requirements, which open up the fire hose of minor security incidents. Tables 1 and 2 show the number of incidents and breaches by victim industry and size. You may have noticed that the totals in Tables 1 and 2 feature fewer incidents and breaches than the previously advertised 100,000 and 3,141. None are typos—there are a couple of filters applied to the original total. We excluded incidents involving devices repurposed as infrastructure to be used against another target (more on this in the Secondary Motive sidebar in Breach Trends). We also had numerous incidents that failed the “You must be this detailed to enjoy this ride” test.<sup>3</sup>

<sup>3</sup> Complexity and completeness scoring is discussed in [Appendix E: Methodology and VERIS resources](#).

When we zoom in on just confirmed breaches, the numbers are less astronomical and we see industries such as Accommodation and Retail accounting for a more significant percentage of breaches (as opposed to incidents). This is unsurprising as they process information which is highly desirable to financially motivated criminals.

Industry	Total	Small	Large	Unknown
Accommodation (72)	362	140	79	143
Administrative (56)	44	6	3	35
Agriculture (11)	4	1	0	3
Construction (23)	9	0	4	5
Educational (61)	254	16	29	209
Entertainment (71)	2,707	18	1	2,688
Finance (52)	1,368	29	131	1,208
Healthcare (62)	166	21	25	120
Information (51)	1,028	18	38	972
Management (55)	1	0	1	0
Manufacturing (31-33)	171	7	61	103
Mining (21)	11	1	7	3
Other Services (81)	17	5	3	9
Professional (54)	916	24	9	883
Public (92)	47,237	6	46,973	258
Real Estate (53)	11	3	4	4
Retail (44-45)	370	109	23	238
Trade (42)	15	3	7	5
Transportation (48-49)	31	1	6	24
Utilities (22)	24	0	3	21
Unknown	9,453	113	1	9,339
Total	64,199	521	47,408	16,270

**Table 1.**

Number of security incidents by victim industry and organization size, 2015 dataset.



Industry	Total	Small	Large	Unknown
Accommodation (72)	282	136	10	136
Administrative (56)	18	6	2	10
Agriculture (11)	1	0	0	1
Construction (23)	4	0	1	3
Educational (61)	29	3	8	18
Entertainment (71)	38	18	1	19
Finance (52)	795	14	94	687
Healthcare (62)	115	18	20	77
Information (51)	194	12	12	170
Management (55)	0	0	0	0
Manufacturing (31-33)	37	5	11	21
Mining (21)	7	0	6	1
Other Services (81)	11	5	2	4
Professional (54)	53	10	4	39
Public (92)	193	4	122	67
Real Estate (53)	5	3	0	2
Retail (44-45)	182	101	14	67
Trade (42)	4	2	2	0
Transportation (48-49)	15	1	3	11
Utilities (22)	7	0	0	7
Unknown	270	109	0	161
Total	2,260	447	312	1501

Small = organizations with fewer than 1,000 employees, Large = organizations with 1,001+ employees.

**Table 2.**

Number of security incidents with confirmed data loss by victim industry and organization size, 2015 dataset.

### Breaches vs. Incidents

This report uses the following definitions:

**Incident:** A security event that compromises the integrity, confidentiality or availability of an information asset.

**Breach:** An incident that results in the confirmed disclosure (not just potential exposure) of data to an unauthorized party.

## Breach trends

Playing a part on the blue team in information security can, to a very small degree, be compared to the lot of a hapless soldier. The soldier is told to guard a certain hill and to keep it at all costs. However, he is not told who his enemy may be, what they look like, where they are coming from, or when (or how) they are likely to strike. To ride this analogous horse a bit further, the soldier is given a hand-me-down rifle with only a few rounds of ammunition to fulfill his task. It seems a bit unfair really—even the American Revolution got Paul Revere.

With that in mind, we hope that this section and the facts and figures contained in it will go some way toward making you better prepared than our friend mentioned above. After all, “forewarned is forearmed.”

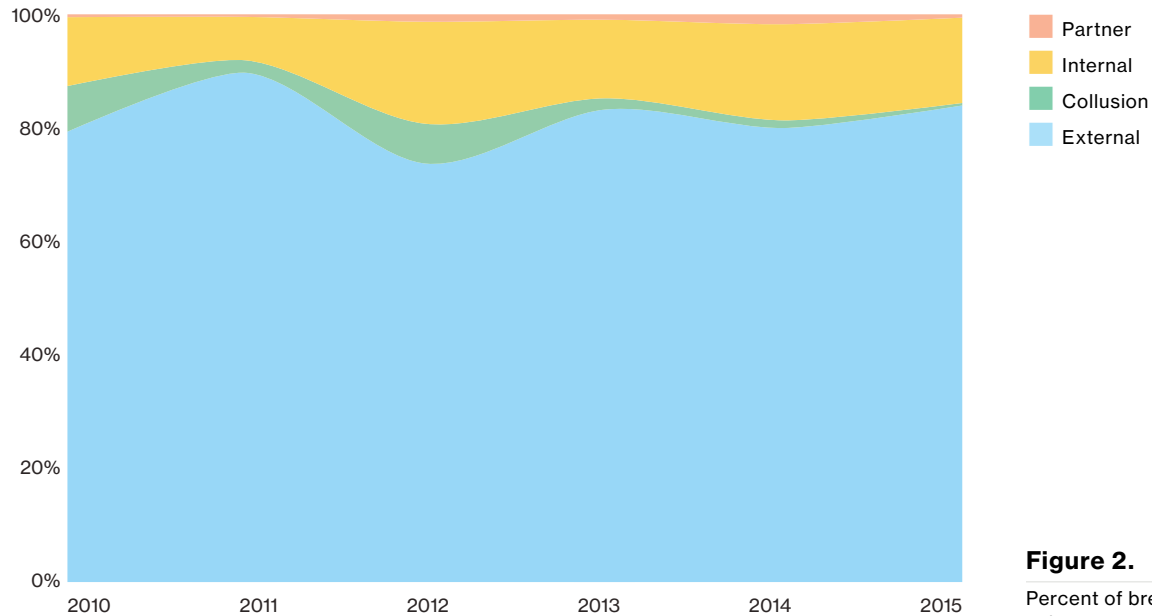
**Be prepared:  
forewarned is  
forearmed.**

### **A brief primer on VERIS**

This section, and many that follow, are based on the Vocabulary for Event Recording and Incident Sharing, or VERIS for short. VERIS is a framework to record and share your security events, incidents and breaches in a repeatable manner. It asks the question, what threat Actor took what Action on what Asset compromising what Attribute? We commonly refer to those as the 4As. In addition to the 4As, it captures timeline, victim demographics, discovery method, impact data and much more.

There are a lot of tools available for VERIS. Methods for creating, importing and analyzing the data are all freely available. More on that in Appendix E: Methodology and VERIS resources.

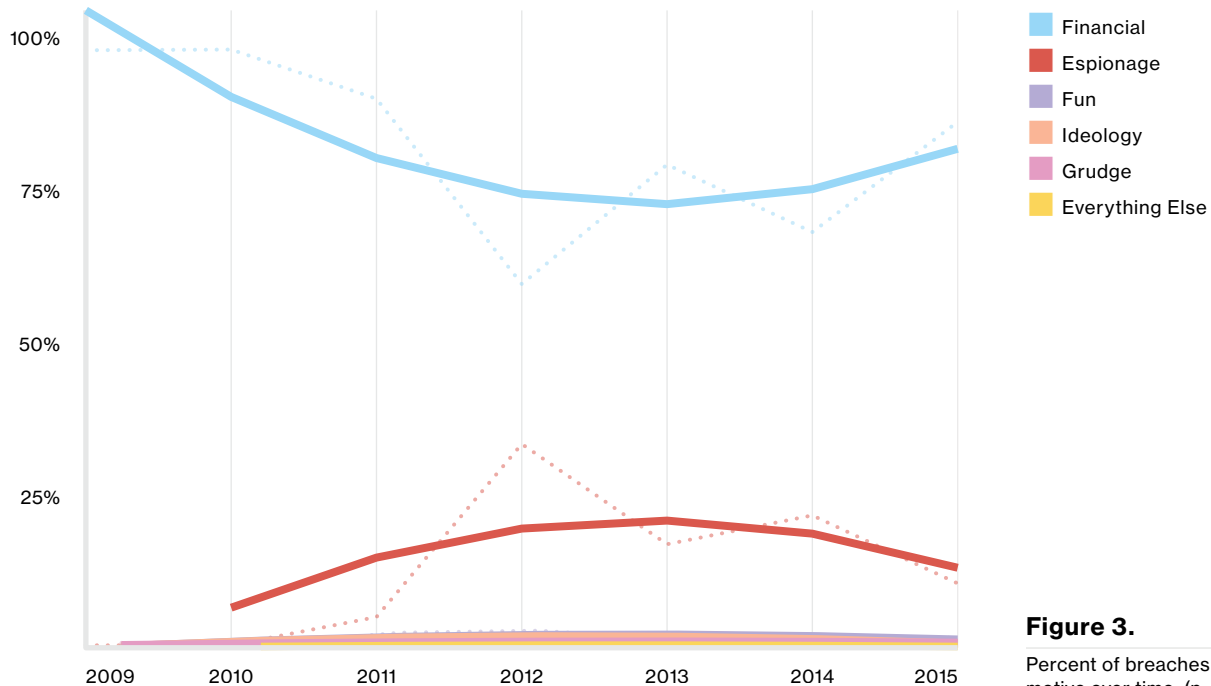
For those who have read the DBIR before, Figure 2 will come as no surprise. Again, the Actors in breaches are predominantly external. While this goes against InfoSec folklore, the story the data consistently tells is that, when it comes to data disclosure, the attacker is not coming from inside the house. And let's face it, no matter how big your house may be there are more folks outside it than there are inside it.



**Figure 2.**  
Percent of breaches per threat actor category over time, (n=8,158)

### Why are these people attacking me?

So why do the Actors do what they do? Money, loot, cash, filthy lucre, greed ... get the idea? In fact, it can be money even when it's not money (see Secondary Motive sidebar for more). In the 2013 DBIR it appeared that perhaps the reigning lothario of "financial gain" was in danger of being cast aside in favor of "espionage." Could such a thing come to pass? No, not really.



**Figure 3.**  
Percent of breaches per threat actor motive over time, (n=6,762)

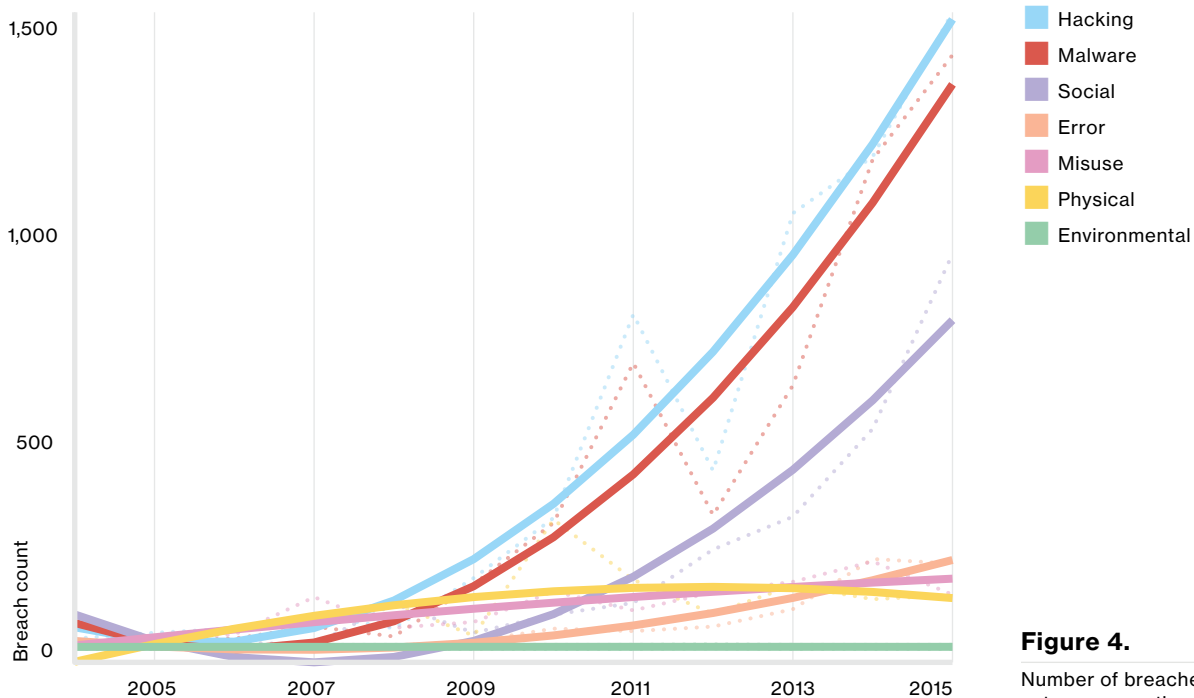
There was never any real danger of the financial motive losing its prominence, as even at its peak, espionage remained a far distant second. As illustrated by Figure 3, breaches with a financial motive dominate everything else, including espionage and fun.

**Secondary motive**

Many of the attacks discussed in this report have what we call a ‘secondary motive’, which we define as when the motive of the incident is to ‘aid in a different attack’. We filter these out of the report because it would overshadow everything else if we didn’t. One example is where the bad guy compromises a web server to repurpose it to his own uses (e.g., hosting malicious files or using it in a spam or DoS botnet). Even criminals need infrastructure. “It is a far, far better thing” that someone else manages it for free, rather than having to pay for it yourself. We had thousands of these incidents, as well as poorly configured NTP and DNS servers, leveraged to launch reflective DoS attacks.

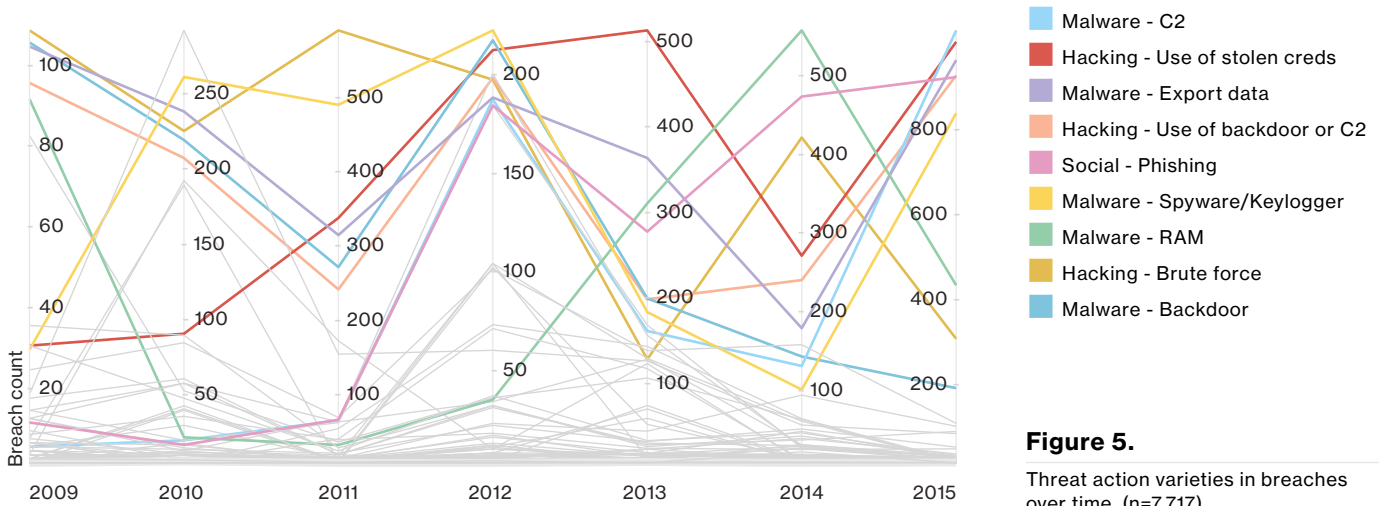
**Pistols at dawn, or knives at noon?**

Now that we know at least a very little bit more about who’s coming after us, the next logical question is: how are they armed? As a glance at Figures 4 and 5 can show you, it is often with phishing, which leads to other events that are not going to make your day. We also see the calling card of Point-of-Sale (POS) attacks. No need to go get in the weeds on this here, as these topics will reappear quite a bit in the pages to follow.



**Figure 4.** Number of breaches per threat action category over time, (n=9,009)

Now, to be fair to the other hardworking threat action types in our list, phishing (and the higher level threat action category of Social) was given a leg up this year by the ‘Dridex’ campaign. We had several contributors who combined to provide a great amount of insight into that naughtiness and this skewed the results somewhat.

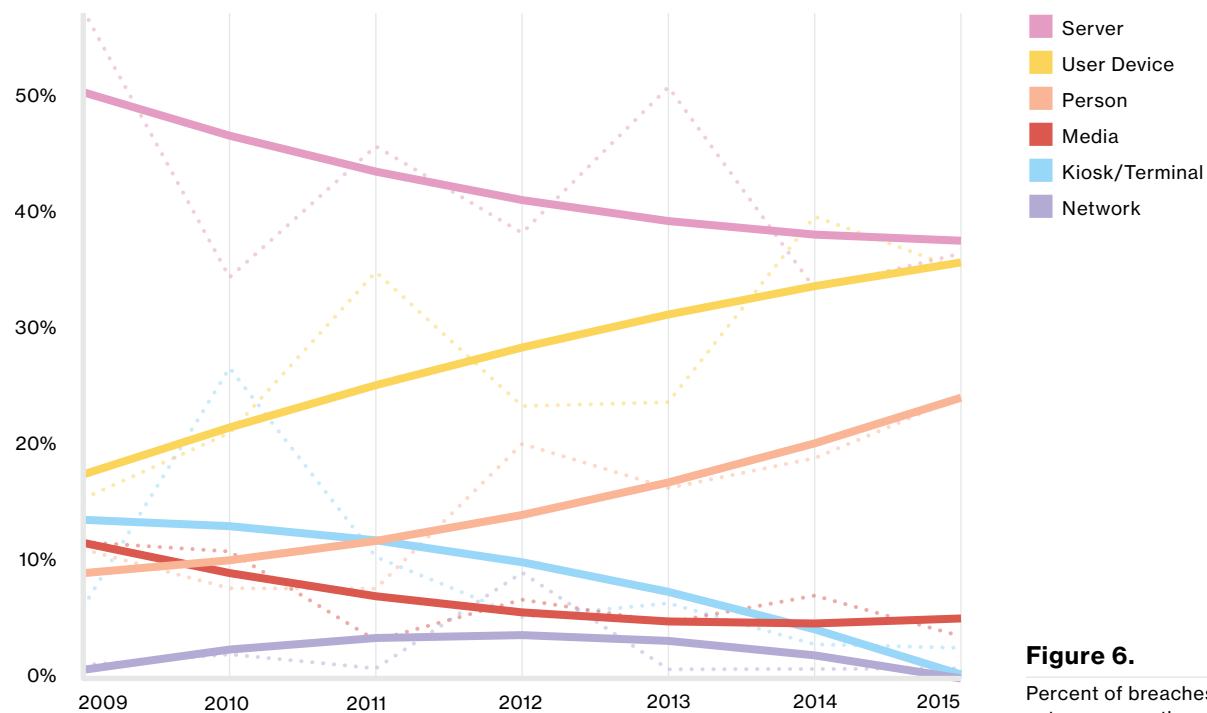


**Figure 5.**  
Threat action varieties in breaches over time, (n=7,717)

Nevertheless, at this point, we think both Phishing and Point-of-Sale could safely say, in their best Ron Burgundy voice, “You might have heard of me, I’m kind of a big deal.” Due to this rock-star status, we’re going to dig a little deeper into POS attacks later in the Patterns section and also in the Post-Compromise Fraud write-up. Likewise, we discuss phishing in greater detail in the Phishing section and Cyber-espionage pattern. We even have a section on credentials this year. Credentials have made numerous cameo appearances in this report for years, but never before have they had a speaking part. (Always a bridesmaid, never a bride.)

**The many facets of assets**

Guess what? When the bad guys’ actions are centered around phishing and POS devices, the asset varieties displayed in Figure 6 reflect this. That lovely “Person” line trending up is due to the human asset falling victim to phishing attacks<sup>4</sup>. The “User device” line upward trend is based on desktops being infected with malware, as well as POS terminals getting popped.

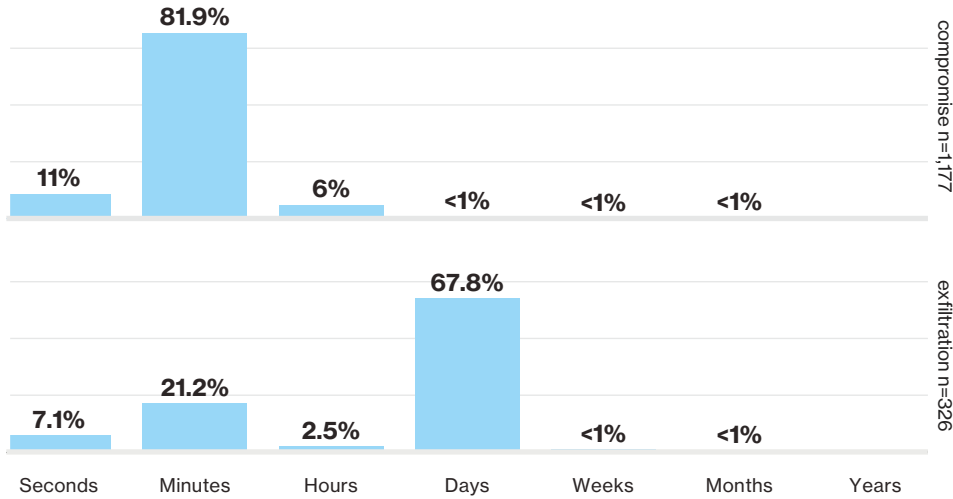


**Figure 6.**  
Percent of breaches per asset category over time, (n=7,736)

4 In VERIS we model this stage of the attack as a loss of Integrity based on the influencing of human behavior.

**Mick was wrong – time is not on our side.**

Rome wasn't built in a day, but data breaches frequently were. Figure 7 illustrates how quickly the threat Actor gets in and out of your network. The large spikes, however, are driven by very specific threats. The compromise time of minutes, while depressing to look at, is actually another reflection of the ubiquitous 'Dridex' breaches in this year's dataset. As previously alluded to, these cases begin with a phishing, featuring an attachment whose mission in its malware life is to steal credentials. If you have legit creds, it doesn't take a very long time to unlock the door, walk in and help yourself to what's in the fridge. Conversely, the exfiltration time being so weighted in the 'days' category is heavily representative of attacks against POS devices where malware is dropped to capture, package and execute scheduled exports.

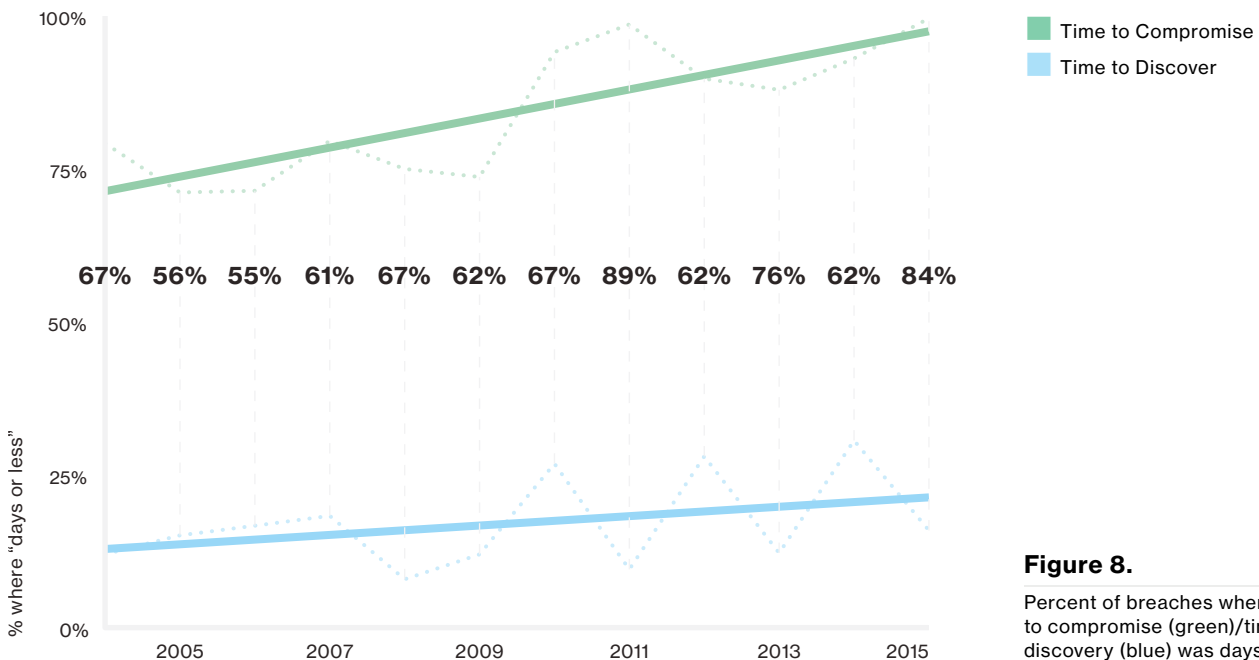


**Figure 7.**

Time to compromise and exfiltration.

**Bad news travels fast, with one exception.**

We like this next graph – one line goes one way and the other line goes the other way. Actually we would like it even more if the lines took different paths. The bad news is, the detection deficit in Figure 8 is getting worse.



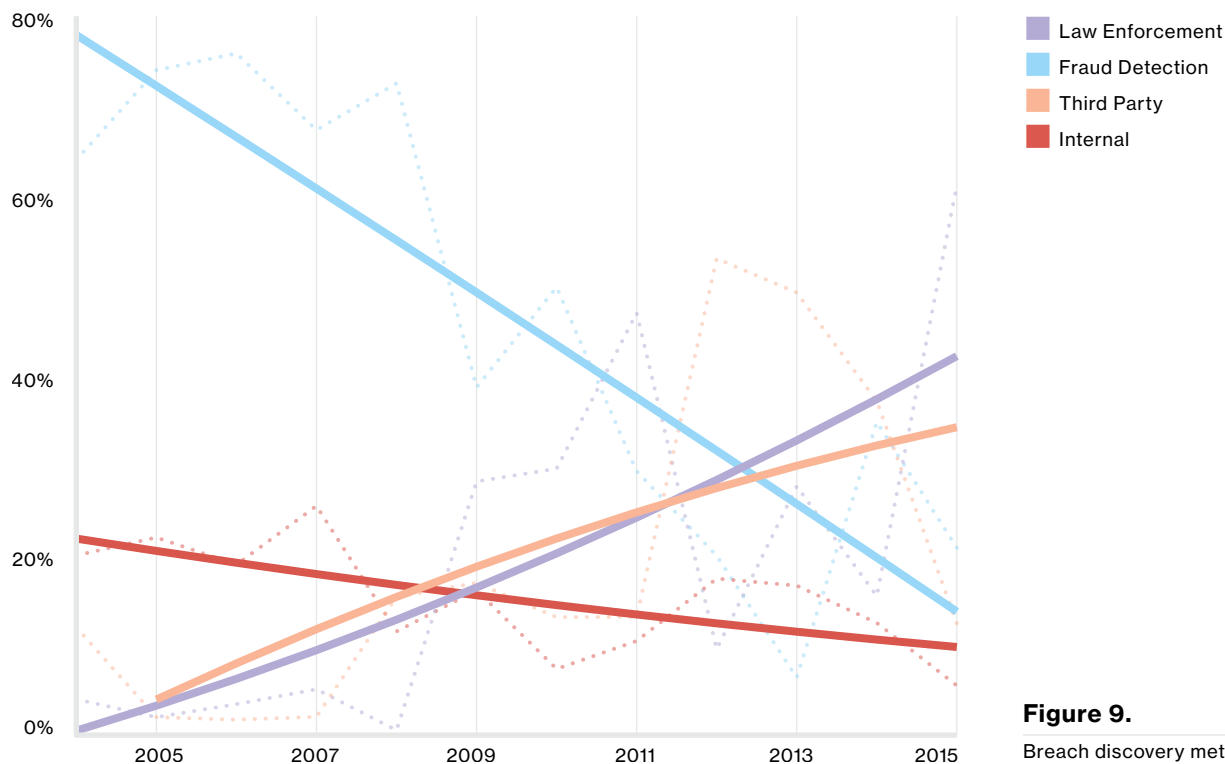
**Figure 8.**

Percent of breaches where time to compromise (green)/time to discovery (blue) was days or less

In the 2015 report, we mentioned that there was some improvement in discovery in the ‘days or less’ category, however, that improvement was short-lived. We also pointed out that we would need more than one year’s data to verify that as a definite trend and sadly we did not get that verification. Moreover, readers with excellent memories will notice that the deficit in 2014 grew from last year’s report. Data for the year-to-year graphs is filtered by incident year (i.e., when the compromise occurred). We continue to add incidents and breaches to prior calendar years post-report to enrich our data. Also, some breaches will occur late in the year and are discovered the next year.

To add another ray to this sunbeam, attackers are getting even quicker at compromising their victims. When you review the leading threat actions again, this really won’t come as a surprise. The phishing scenario is going to work quickly, with the dropping of malware via malicious attachments occurring within seconds. Physical compromises of ATMs and gas pumps also happen in seconds. In the majority of confirmed data breaches, the modus operandi of nation-states as well as financially motivated attackers is to establish control via malware and, when successful, it is lightning fast. As this figure is for confirmed breaches only, it makes sense that the time to compromise is almost always days or less (if not minutes or less). If—and some have called “if” the biggest word in the language—there’s any good news, it’s that the number of breaches staying open months or more continues to decline slightly.

**The time to compromise is almost always days or less, if not minutes or less.**



**Figure 9.** Breach discovery methods over time, (n=6,133).

When it comes to external<sup>5</sup> breach discovery, fraud detection and law enforcement notification are battling it out like the Celtics and Lakers in the ‘80s. Figure 9 shows that law enforcement will raise the banner for 2015, due (again) to a botnet takedown and the subsequent notifications to members of the botnet. All in all, external notification is up. And when you have to wait on external detection to tell you you’re popped, it’s probably too late to keep the horses in the barn.

<sup>5</sup> External is everything but internal detection and when a partner supplies a monitoring or AV service.

## Points of focus

One last thing before we get to the patterns. There are a couple of topics that are omnipresent in many of the patterns that we use to classify incidents. While they will receive credit where credit is due, in the pattern sections, we feel that we also need to put the spotlight on them here.

We have numerous breaches where we can infer that some Common Vulnerabilities and Exposures (CVE) were used in order for the attack to advance. Hey, we're looking at you, drive-by downloads! Unfortunately, we don't have a tremendous amount of CVE data in our corpus, either because it was not measured or was unable to be identified. This lack of detail makes us an embarrassment of sad pandas. (Yes, we wanted to say "sleuth", but apparently we can't. Look it up.) Luckily we have contributors in the vulnerability space that can lighten our mood.

Phishing has continued to trend upward (like spawning salmon?) and is found in the most opportunistic attacks as well as the sophisticated nation state tomfoolery. We feature a section where we dive into the human element a bit deeper, with some data on our innate need to click stuff.

Lastly, we strike a deceased equine a bit more with a section on credentials (of the static variety). Don't get us wrong—passwords are great, kind of like salt. Wonderful as an addition to something else, but you wouldn't consume it on its own.

**We don't have a tremendous amount of CVE data because it wasn't measured or was unable to be identified.**



# Vulnerabilities



## At a glance

<b>Description</b>	A look into software vulnerabilities, whether we are making any progress in addressing them and ways to improve.
<b>Contributors</b>	Kenna Security (formerly Risk I/O) collaborated with us again to leverage their vulnerability and exploitation data. We also utilized vulnerability scan data provided by Beyond Trust, Qualys and Tripwire in support of this section.
<b>Key findings</b>	Older vulnerabilities are still heavily targeted; a methodical patch approach that emphasizes consistency and coverage is more important than expedient patching.

**New vulnerabilities come out every day.**

### Methodology

The visualizations and statements regarding rates of exploitation in this section are underpinned by vulnerability exploitation data provided by Kenna Security. This dataset spans millions of successful real-world exploitations, and is derived from hunting down exploitation signatures in security information and event management (SIEM) logs and correlating those with vulnerability scan data to find pairings that would be indicative of a successful exploitation.

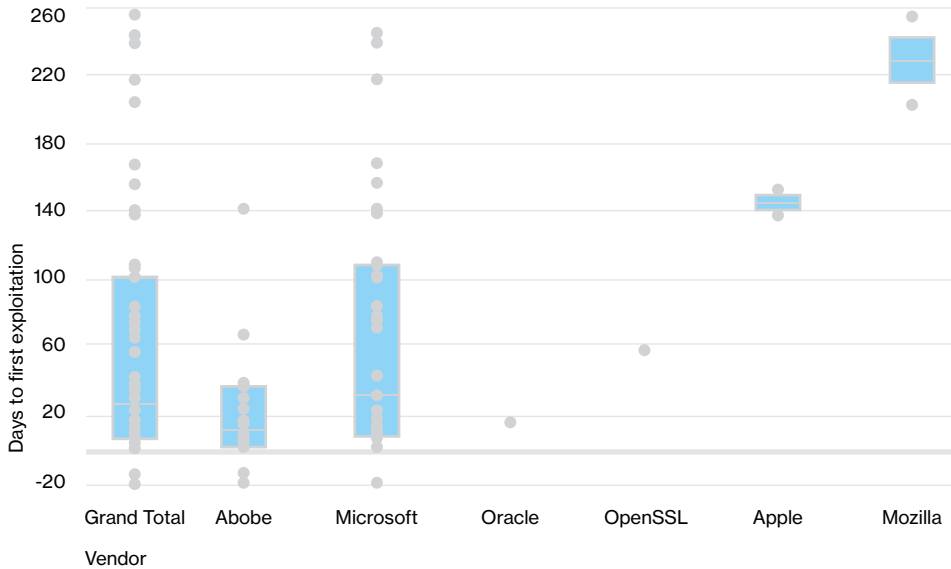
### The tortoise and the hare

Vulnerability management has been a Sisyphean endeavor for decades. Attacks come in millions, exploits are automated and every enterprise is subject to the wrath of the quick-to-catch-on hacker. What's worse, new vulnerabilities come out every day. Since the first DBIR, we've been advocating the turtle's approach to vulnerability management (slow and steady wins the race).

This year we revisit this data to see whether the trends hold, but in typical DBIR fashion, we dig a little deeper, to look at not just how attackers are interacting with vulnerabilities (exploitation), but also how well and how fast enterprises are executing remediation. If we can measure both of these routinely, then we can provide much-needed answers about how the tortoise won the race—and so learn how to close the gap between attackers and enterprises.

### Slow and steady—but how slow?

This year we take a different approach to measuring the time from publication to exploitation. Figure 10 is a box plot, which plots the time between publication and the first observed successful exploit by vendors.<sup>6</sup> We can see that Adobe vulnerabilities are exploited quickly, while Mozilla vulnerabilities take much longer to exploit after disclosure. Half of all exploitations happen between 10 and 100 days after the vulnerability is published, with the median around 30 days. This provides us with some general guidelines on which software vulnerabilities to prioritize along with some guidance on time-to-patch targets.

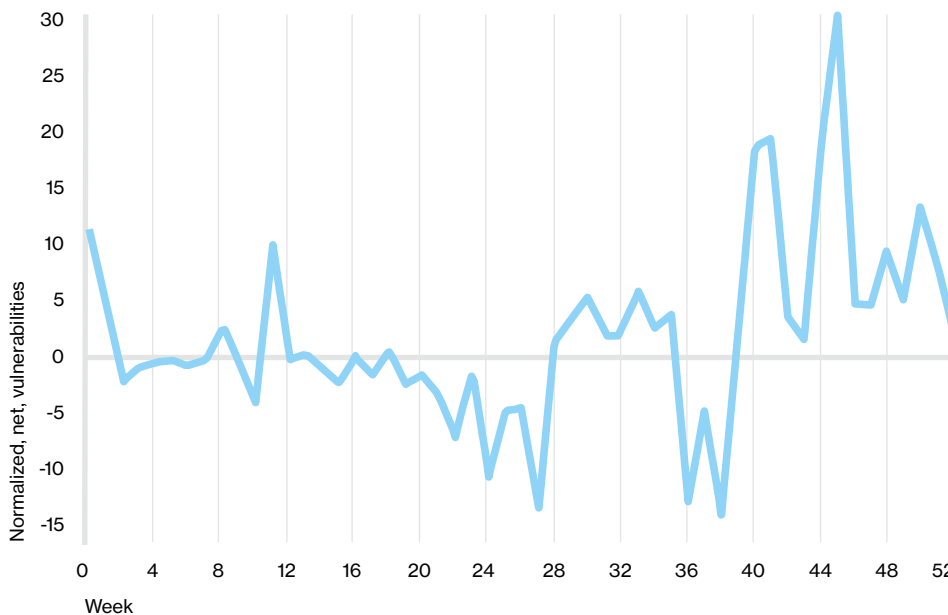


**Figure 10.**

Time to first-known exploitation by vulnerability category.

### Treading water

Figure 11 shows the number of vulnerabilities opened each week minus the number of vulnerabilities (aka “vulns”) closed, scaled by the number of assets in the dataset during each week of 2015. When the line is above zero, it means that more vulns are being opened than closed (new vulns disclosed, more



**Figure 11.**

Delta of number of vulnerabilities opened each week and number closed.

<sup>6</sup> The blue boxes in Figure 10 represent 50% of the values for a given category and the gray line within the box is the median value. The dots represent individual values.

machines entering the environment, new software installed). When it's below zero, remediation efforts are driving down vulnerability counts faster than new vulns are entering the enterprise.

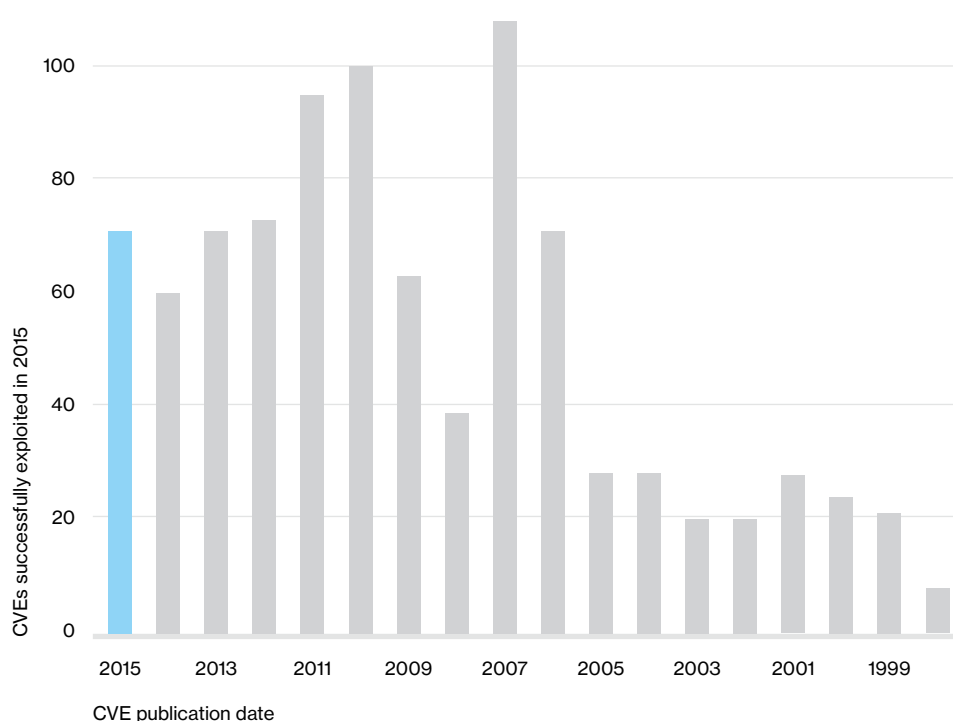
Basically, we confirmed across multiple datasets that we are treading water—we aren't sinking in new vulnerabilities, but we're also not swimming to the land of instantaneous remediation and vuln-free assets. However, all that patching is for naught if we're not patching the right things. If we're going to tread, let's tread wisely.

**All that patching is for naught if we're not patching the right things.**

**What should we mitigate? Hacker economics.**

So what are the right things? The 2015 DBIR gave us an idea and since then, not much has changed.

Revisiting last year's trends, we find that the two golden rules of vulnerabilities still hold.



**Figure 12.**  
Count of CVEs exploited in 2015 by CVE publication date.

First, Figure 12 arranges CVEs according to publication year and gives a count of CVEs for each year. While 2015 was no chump when it came to successfully exploited CVEs, the tally of really old CVEs which still get exploited in 2015 suggests that the oldies are still goodies. Hackers use what works and what works doesn't seem to change all that often.<sup>7</sup> Secondly, attackers automate certain weaponized vulnerabilities and spray and pray them across the internet, sometimes yielding incredible success. The distribution is very similar to last year, with the top 10 vulnerabilities accounting for 85% of successful exploit traffic.<sup>8</sup> While being aware of and fixing these mega-vulns is a solid first step, don't forget that the other 15% consists of over 900 CVEs, which are also being actively exploited in the wild.

<sup>7</sup> Astute and frequent readers of the DBIR will notice one more gem in this chart—last year, the numbers of published CVEs exploited were lower across the board—and this year, we have more and better data. Those newly exploited CVEs however, are mostly—and consistently—older than one year.

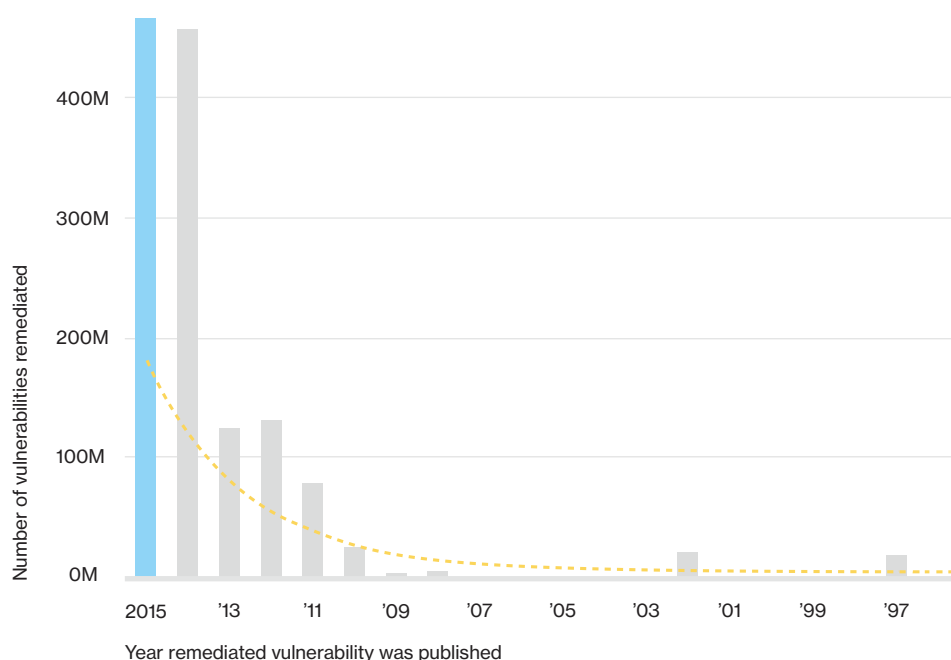
<sup>8</sup> Here are the results from the initial analysis on the dataset. We would expect that in your environment the top CVEs would be different, but a similar distribution would still be realized. CVE-2001-0876, CVE-2001-0877, CVE-2002-0953, CVE-2001-0680, CVE-2002-1054, CVE-2015-0204, CVE-2015-1637, CVE-2003-0818, CVE-2002-0126, CVE-1999-1058.

We have received several criticisms on the data behind the exploitation rates, and how it has been presented in this section. It is not our intent for anyone to rely on the individual CVEs listed in Footnote 8 to base their vulnerability management strategy on (or any external list of CVE's for that matter). Our full response to the criticisms as well as links to Kenna's responses can be found here: <https://securityblog.verizonenterprise.com/?p=7544>

### Can't solve everything

In Figure 13, we see that during 2015, vulnerabilities published in 2015 and 2014 were being patched. After that though, the vulnerabilities begin to drop off and really hit a steady state. This gets at a core and often ignored vulnerability management constraint—sometimes you just can't fix a vulnerability—be it because of a business process, a lack of a patch, or incompatibilities. At that point, for whatever reason, you may have to live with those residual vulnerabilities. It's important to realize that mitigation is often just as useful as remediation—and sometimes it's your only option.

**Mitigation is often just as useful as remediation—and sometimes your only option.**



**Figure 13.**  
Closure rate of CVEs by CVE publication date.

## Recommended controls

### Knowledge is power.

Establish a process for vulnerability remediation that targets vulnerabilities which attackers are exploiting in the wild, followed by vulnerabilities with known exploits or proof-of-concept code.

### Have a Plan B.

If you have a system that cannot be patched or receive the latest-and-greatest software update, identify it, and apply other risk mitigations in the form of configuration changes or isolation. Discuss a plan on how the device(s) could be replaced without causing severe business disruption.

### At your service

Vulnerability scanning is also useful in identifying new devices and new services. Review scan-to-scan changes as another control to identify unknown devices and deviations from standard configurations.

# Phishing



## At a glance

<b>Description</b>	A form of social engineering in which a message, typically an email, with a malicious attachment or link is sent to a victim with the intent of tricking the recipient to open an attachment.
<b>Contributors</b>	Anti-Phishing Working Group, Lares Consulting, SANS Securing the Human and Wombat Security provided the non-incident data for this section.
<b>Top patterns</b>	Everything Else, Web App Attacks, Cyber-espionage
<b>Frequency</b>	9,576 total incidents, 916 with confirmed data disclosure.
<b>Key findings</b>	13% of people tested click on a phishing attachment; median time to click is very short.

**The majority of phishing cases feature phishing as a means to install persistent malware.**

### **You can't fool all the people all the time. Or can you?**

Social engineering in its basic form is simply to dupe or trick someone into doing something they would not otherwise do (not unlike some online dating). Social tactics can take many forms such as pretexting,<sup>9</sup> elicitation (the subtle art of extracting information from a subject via conversation), baiting (planting infected media in victim areas), and a myriad of other lowdown and dirty tricks. However, by far its most successful variety is phishing, which as the name implies is malicious correspondence trying to get the recipient to take the bait in the form of an attachment or embedded link. It is important to note that 'pretexting' via email (a back-and-forth dialogue leveraging an invented scenario to gain a certain end) and a phishing email are similar, but not the same. In the case of a pretexting email, the criminal is primarily purporting to be someone they are not, usually within the victim organization (e.g., the CFO who instructs the victim to approve a fraudulent Automated Clearing House (ACH) transfer).

### **Bummed is what you are...**

...when you click on that attachment and get owned. The basic structure of phishing attacks remains the same—user clicks, malware drops, foothold is

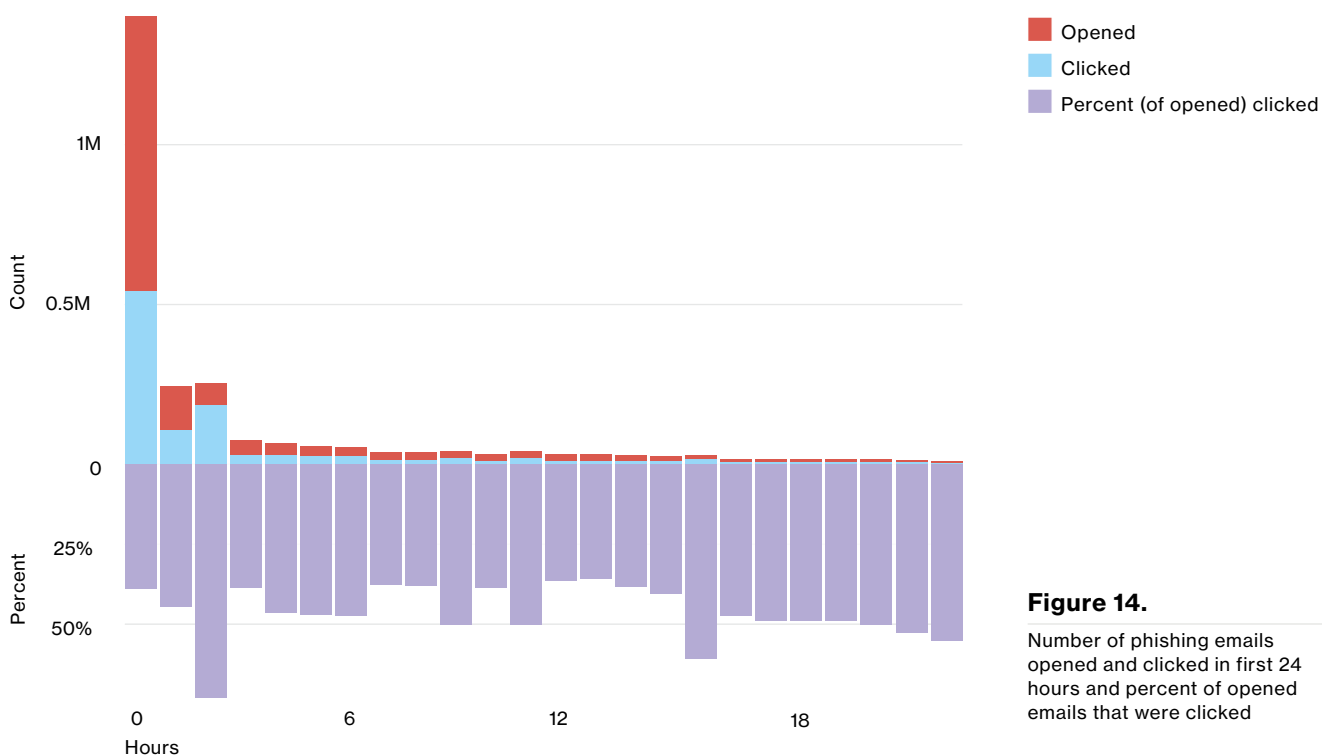
<sup>9</sup> I'm Frieda's boss.

gained. There are still cases where the phishing email leads users to phony sites, which are used to capture user input, but the majority of phishing cases in our data feature phishing as a means to install persistent malware. The victim opens the email, sees the attachment that contains the malware du jour and says “That file looks good, I’ll have that”. What happens next is dictated by the end goal of the phisher.

**“What we have here is a failure to communicate.”**

Apparently, the communication between the criminal and the victim is much more effective than the communication between employees and security staff. We combined over eight million results of sanctioned phishing tests in 2015 from multiple security awareness vendors aiming to fix just that. Figure 14 is jam-packed with information. In this year’s dataset, 30% of phishing messages were opened by the target across all campaigns.<sup>10</sup> “But wait, there’s more!” (in our best infomercial voice) About 12% went on to click the malicious attachment or link and thus enabled the attack to succeed. That indicates a significant rise from last year’s report in the number of folks who opened the email (23% in the 2014 dataset) and a minimal increase in the number who clicked on the attachment (11% in the 2014 dataset). The median time for the first user of a phishing campaign to open the malicious email is 1 minute, 40 seconds. The median time to the first click on the attachment was 3 minutes, 45 seconds, thus proving that most people are clearly more on top of their email than I am.

**The main perpetrators for phishing attacks are organized crime syndicates and state-affiliated actors.**

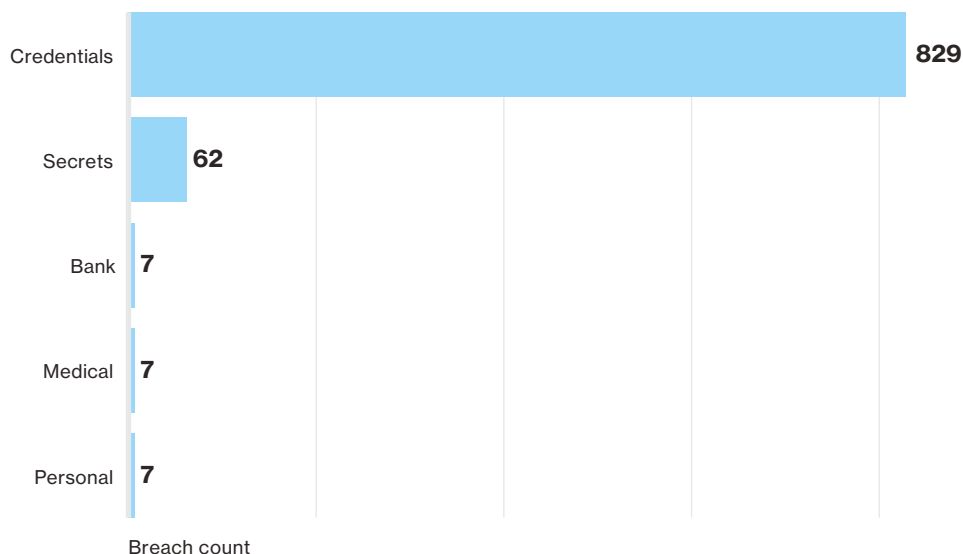


**Figure 14.** Number of phishing emails opened and clicked in first 24 hours and percent of opened emails that were clicked

However, before we drag these individuals outside and collectively stone them, keep in mind that the main perpetrators for these types of attacks are organized crime syndicates (89%) and state-affiliated Actors (9%) who can put some thought into the ruse they use (yeah, I know). In roughly 636,000 sanctioned phishing emails, we captured whether the email was reported. Approximately 3% of targeted individuals alerted management of a possible phishing email. We did not verify by what means the email was reported, or whether it was because they were savvy enough to avoid the trap or because they only realized it once they had fallen in themselves.

<sup>10</sup> Granted this could be affected by preview pane opening of emails or people not loading images in emails.

As an aside, the smaller proportion of nation-state Actors in this year's data is due to a large contribution from a particular contributor who saw a great deal of 'Dridex' campaigns which skewed the data toward organized crime. We should not conclude from this that certain groups from East Asia have had a crisis of conscience and mended their wicked ways.



**Figure 15.**

Top five data varieties breached by phishing attacks, (n=905)

What do the attackers ultimately steal? A heck of a lot of credentials (mostly due to the large amount of opportunistic banking Trojans—beware of Greeks bearing gifts), but also trade secrets.

## Recommended controls

### Filter it! Filter it real good!

“An ounce of prevention is worth a pound of cure.” It was good advice when Ben said it and so it remains. The first opportunity to defend against email-borne threats is (thankfully) before a human can interact with it. Email filtering is your buddy in this fight and you need to have an understanding of your current solution, and test its implementation.

### Talk amongst yourselves (I'm verklempt)!

Provide employees with awareness training and information so they can tell if there is something 'phishy' (couldn't resist) going on. Also, provide them with a means for reporting these events. We recommend a button on their taskbar, but whatever works for you.

### One click does not a catastrophe make.


So, it snuck past your email filters and someone went clicky-clicky. There is still ample opportunity to limit the impact. Assuming the organization's "seekrit stuff" isn't resident on the initial foothold, make it hard to pivot from the user device to other assets in the organization. Protect the rest of your network from compromised desktops and laptops by segmenting the network and implementing strong authentication between the user networks and anything of importance. Static passwords are adorable, but sophisticated attackers don't just bypass them, they utilize them to advance their attack.

### Keep your eye on the ball.

You increase your chances of catching signs that you have fallen victim to a phishing attack if you monitor outbound traffic for suspicious connections and potential exfiltration of data to remote hosts.

**Protect the rest of your network from compromised desktops and laptops by segmenting the network and implementing strong authentication.**

# Credentials

 At a glance	
<b>Description</b>	Use of stolen credentials and other hacking and malware actions targeting traditional username and password authentication are prevalent across numerous patterns.
<b>Top patterns</b>	Web App Attacks, POS Intrusions
<b>Frequency</b>	1,429 incidents with confirmed data disclosure.
<b>Key findings</b>	Static credentials continue to be targeted by several of the top hacking action varieties and malware functionalities.

**63% of confirmed data breaches involved weak, default or stolen passwords.**

## **We're not mad, just disappointed.**

The use of stolen, weak or default credentials in breaches is not new, is not bleeding edge, is not glamorous, but boy howdy it works. Static authentication mechanisms have been attacked for as long as we can remember. Password guessing from an InfoSec perspective has been around at least as long as the Morris worm, and has evolved to prominent malware families like Dyre and Zeus that are designed to (among other bad things) capture keystrokes from an infected device. All those efforts to get users to use special characters, upper/lower case numbers and minimum lengths are nullified by this ubiquitous malware functionality.

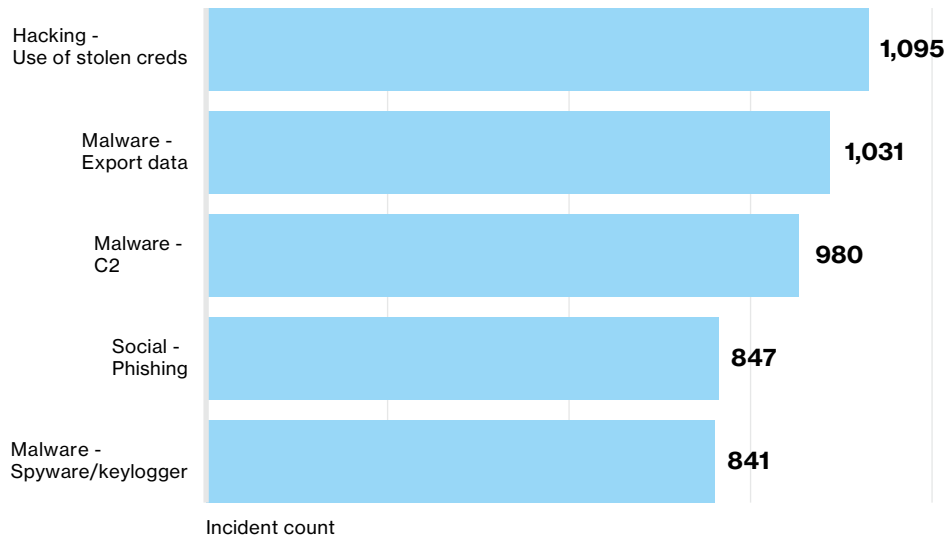
The capture and/or reuse of credentials is used in numerous incident classification patterns. It is used in highly targeted attacks as well as in opportunistic malware infections. It is in the standard toolkit of organized criminal groups and state-affiliated attackers alike. Even fraud committed with stolen payment card data often relies on the static Card Verification Value (CVV) information on the magnetic stripe.<sup>11</sup>

We are realists here, we know that implementation of multi-factor authentication is not easy. We know that a standard username and password combo may very well be enough to protect your fantasy football league. We also know that implementation of stronger authentication mechanisms is a bar

<sup>11</sup> More on this in the [Post-Compromise Fraud appendix](#).



raise, not a panacea. Even with all of that, 63%<sup>12</sup> of confirmed data breaches involved leveraging weak/default/stolen passwords. This statistic drives our recommendation that this is a bar worth raising. Figure 16 shows the most common threat action varieties associated with attacks involving legitimate credentials. The obvious action of the use of stolen credentials is numero uno, but we see some other common actions used in conjunction, including C2 malware, exporting of data, phishing and keyloggers.



**Figure 16.**

Top threat action varieties within incidents involving credentials, (n=1,462)

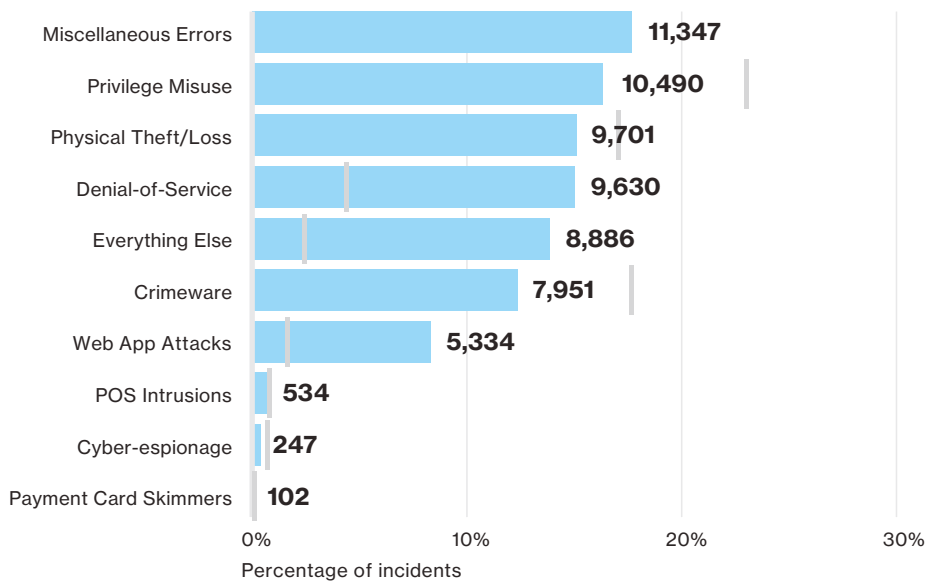
<sup>12</sup> We combined all incidents with confirmed data disclosure AND use of stolen creds OR brute force OR password dumpers OR a data variety of credentials.

# Incident classification patterns

What began with a muttered complaint of “ugh, another one of these” during data conversion a couple of years ago grew into a shift in how we present our core results and analysis. The nine incident classification patterns were born of recurring combinations of the who (Actors), what (assets), how (actions) and why (motive) among other incident characteristics.

In the 2014 report, we found that over 90% of breaches fell into one of the nine buckets and this year’s dataset is no different. We hope that by discussing security incidents, both for this year and historically, and using these clusters as the foundation, we can allow security folks to gain the most from the entire (huge) dataset. Understanding that you don’t have to necessarily worry about 2,260 different breach possibilities, but only a select number of nine patterns (depending on your industry) makes the life of a CISO less of a daily Kobayashi Maru.

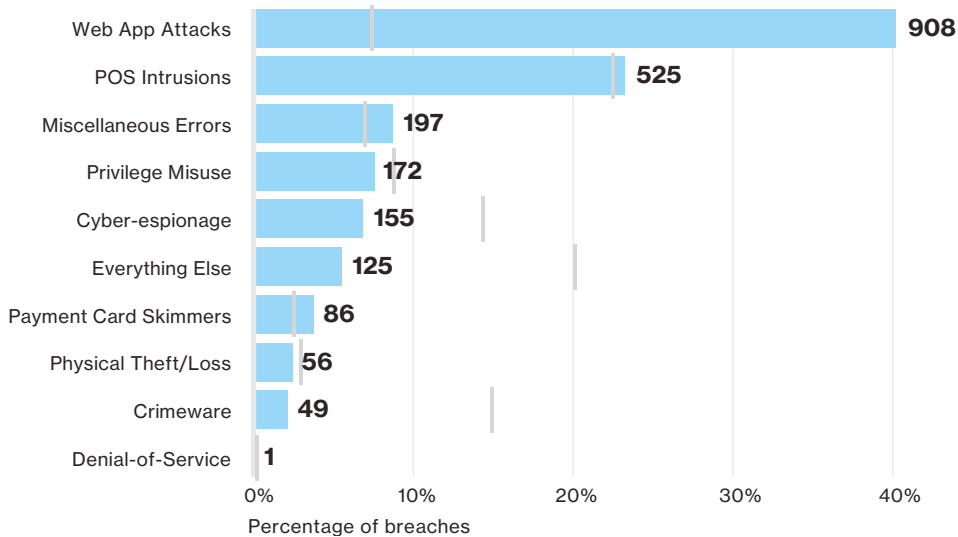
Before we dive deeper into changes over time and the individual patterns (and don’t fret, we will), let’s take a moment and look at the incident and breach breakouts for 2015 in Figures 17 and 18.



**The nine classification patterns were born of recurring combinations of the who, what, how and why.**

**Figure 17.**

Percentage (blue bar), and count of incidents per pattern. The gray line represents the percentage of incidents from the 2015 DBIR. (n=64,199)

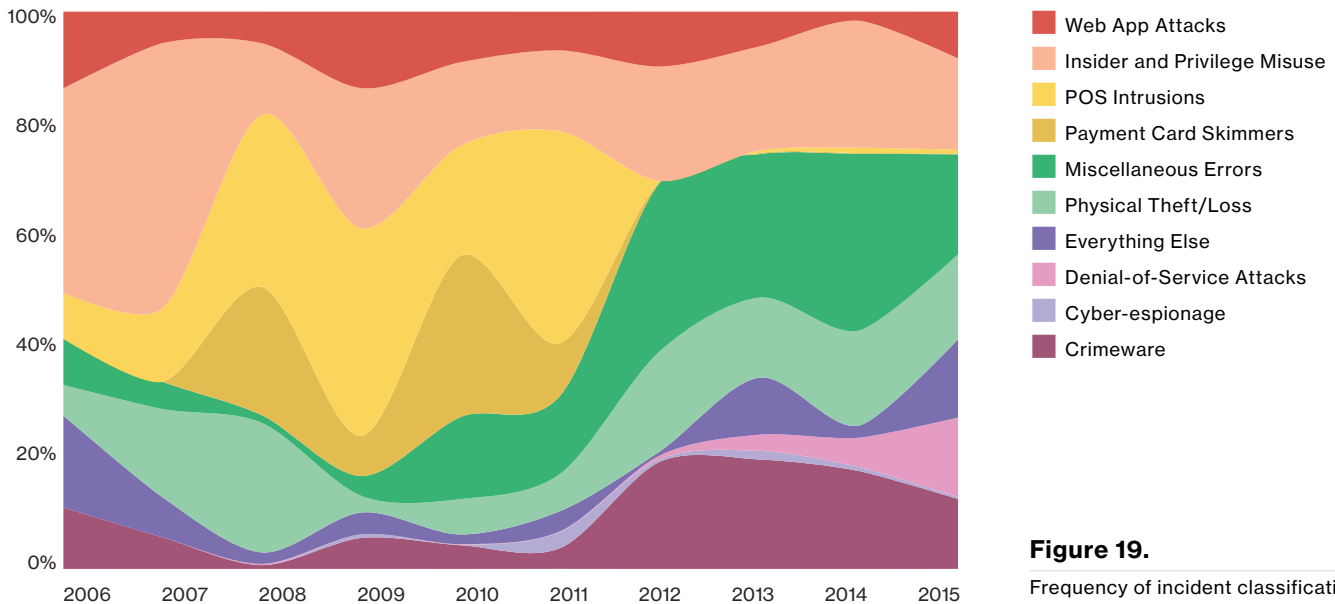


**Figure 18.**

Percentage (blue bar), and count of breaches per pattern. The gray line represents the percentage of breaches from the 2015 DBIR. (n=2,260)

Much to the chagrin of Jerry Lee Lewis, there was not a whole lot of moving and shaking going on in the pattern rankings compared to last year and looking at all incidents, only one pattern moved in the pecking order. Crimeware was the third most common pattern last year and has moved to sixth. The reason is the filter on the secondary motive we discussed in the Breach Trends section. Thousands of incidents where we know a device was participating in a denial-of-service (DoS) bot (but nothing else) were not sent to /dev/null per se, but you won't find them here.<sup>13</sup>

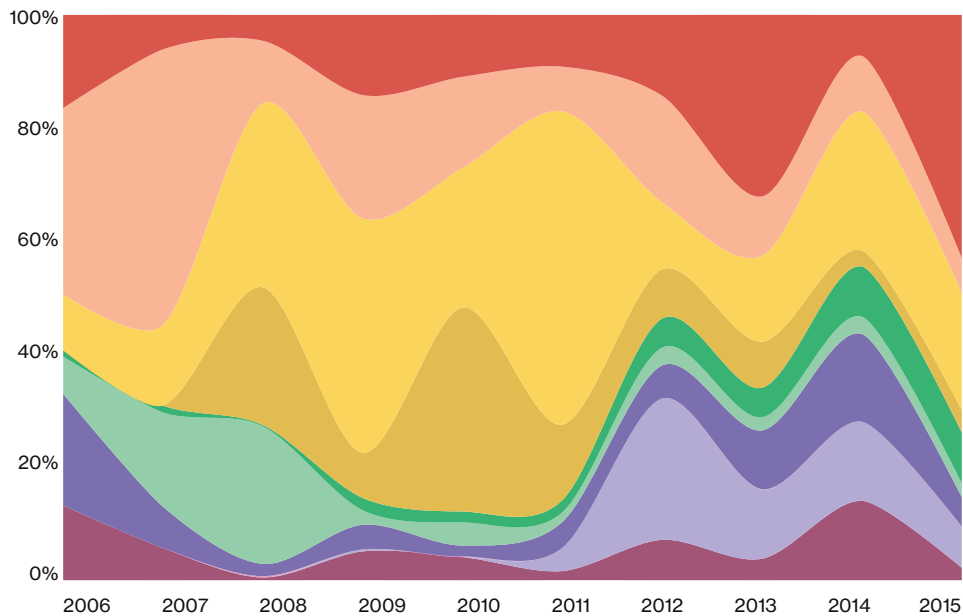
The fact is that our dataset is constantly evolving with contributors joining (yay) and others not able to participate for a year. Many of our contributors have a certain specialty or discipline that results in their data being associated with a certain victim industry, or threat Actor type, or country ... you get the picture. Because of this fact, the ebbs and flows in the patterns from year to year are attributed more to changes in our data than changes in the threat landscape. Bad guy trends would likely be best gleaned from the threat action variety level within a pattern and again, the deeper dives are coming. Having said all of that, Figures 19 and 20 represent the obligatory "trend" graphs.



**Figure 19.**

Frequency of incident classification patterns over time across security incidents.

<sup>13</sup> There are thousands of compromised web servers used as phishing sites that did not make the cut either. No information on how the server was compromised, or if it was owned or maintained by an organization, was available.



- Web App Attacks
- Insider and Privilege Misuse
- POS Intrusions
- Payment Card Skimmers
- Miscellaneous Errors
- Physical Theft/Loss
- Everything Else
- Denial-of-Service Attacks
- Cyber-espionage
- Crimeware

**Figure 20.**

Frequency of incident classification patterns over time across confirmed data breaches.

OK, in lieu of worrying about how patterns rank overall compared to each other, let's get to the good stuff. The best way to use the patterns is to understand the applicability of each of them to your organization. The following charts show the frequency of each of the patterns relative to each industry. In other words, it shows for all the incidents (Figure 21) and breaches (Figure 22) within your industry, those patterns which were common and those which didn't make an appearance. We have included the incident and breach totals again as some of the combinations are a small percentage, but still represent a significant number of events. We use the North American Industry Classification System (NAICS) to classify the victim industry—go to the NAICS website<sup>14</sup> if you're unsure where your organization fits. Of course if you are an E Corp-like conglomerate, you can have business units that fall into several industry categories.

Crimeware	Cyber-espionage	Denial of Service	Everything Else	Stolen Assets	Misc. Errors	Card Skimmers	Point of Sale	Privilege Misuse	Web Apps
<1%	<1%	20%	1%	1%	1%	<1%	74%	2%	1%
		56%	4%		2%		4%	22%	11%
2%	2%	81%	2%	3%	4%			1%	5%
		99%		<1%			1%		1%
2%	<1%	34%	5%	<1%	1%	6%	<1%	3%	48%
4%	2%		11%	32%	18%		5%	23%	4%
4%	3%	46%	21%	<1%	11%		<1%	2%	12%
5%	16%	33%	33%		1%		1%	6%	6%
1%	2%	90%	2%	1%	1%			2%	1%
16%	<1%	1%	17%	20%	24%		<1%	22%	<1%
1%	<1%	45%	2%		1%	3%	32%	1%	13%
10%	16%	26%			6%			6%	35%

**Figure 21.**

Incident patterns by industry minimum 25 incidents

- Accommodation (72), n=362
- Administrative (56), n=44
- Educational (61), n=254
- Entertainment (71), n=2,707
- Finance (52), n=1,368
- Healthcare (62), n=166
- Information (51)n, 1,028
- Manufacturing (31-33), n=171
- Professional (54), n=916
- Public (92), n=47,237
- Retail (44-45), n=370
- Transportation (48-49), n=31

<sup>14</sup> [Census.gov/cgi-bin/sssd/naics/naicsrch?chart=2012](http://Census.gov/cgi-bin/sssd/naics/naicsrch?chart=2012)

From an incident standpoint, Denial-of-Service stands out like “a zoot suit at a Quaker funeral”. This is partly due to the fact that DoS attacks are in fact, happening all the time – remember all those popped boxes in the DoS botnets we filtered out? Another reality is that the other patterns that are more commonly classified as incidents as opposed to confirmed data breaches (Crimeware, Insider and Privilege Misuse, and Physical Theft and Loss) are mostly provided by the public sector and healthcare. Those are the top three incident patterns and we are confident that in the real world they are taking some of that market share from DoS in other industries.

Crimeware	Cyber-espionage	Denial of Service	Everything Else	Stolen Assets	Misc. Errors	Card Skimmers	Point of Sale	Privilege Misuse	Web Apps
			1%	<1%	1%	<1%	95%	1%	1%
	7%		17%	17%	27%			3%	30%
				3%			47%		50%
1%	<1%	<1%	2%	<1%	2%	9%		4%	82%
3%	3%		11%	19%	22%		7%	32%	3%
1%	3%		4%		25%		1%	11%	57%
3%	47%		3%				3%	24%	21%
4%	19%		25%	4%	15%			21%	13%
12%	16%		4%	9%	37%			13%	9%
1%	1%		4%		1%	3%	64%	2%	26%

**Figure 22.**

Incident patterns by industry minimum 25 incidents (only confirmed data breaches)

Accommodation (72), n=282

Educational (61), n=29

Entertainment (71), n=38

Finance (52), n=795

Healthcare (62), n=115

Information (51), n=194

Manufacturing (31-33), n=37

Professional (54), n=53

Public (92), n=193

Retail (44-45), n=182

The most interesting discovery in the breach patterns to industry matrix was the rise of Web App Attacks across the board, but especially for financial services organizations (up from 31% in the 2015 DBIR). The next item that raised an eyebrow or two (or perhaps a unibrow) was the decline (down from 36% last year) in Crimeware, also in Finance. Is there anything to this? Actually, yes. This year, again thanks to the organizations involved in the Dridex takedown, we have even more data involving the reuse of stolen credentials. This caused the spike in the Web App Attack pattern and if we removed these breaches, the numbers would be more in line with 2014. On the flip side, in 2014 we received more data on malware infections within organizations, leading to breaches that landed in our Crimeware bucket. Is Crimeware not playing as big a role in breaches? The perspective of the reporting contributor has a lot to do with the pattern breakdowns as well. Using the banking Trojan example:

**Event 1:** Organization A is infected with a Zeus variant via a drive-by download

**Event 2:** Malware has a keylogging functionality that captures banking credentials

**Event 3:** Malware exports captured data to command and control (C2) server

## Intermission music

**Event 4:** Credentials are used to log into Organization B web server

**Event 5:** Fraudulent transaction is initiated

Organization B may be quick to say “We didn’t have a malware incident” and if events 4–5 are provided to us, the incident would find a good home in the Web App Attacks section. But if we received data from Organization A and only events 1–3 are documented, it now becomes a newly minted Crimeware breach.

It is important to realize that there are interrelations between the incident patterns that aren’t always evident. Crimeware in one organization leads to DoS against another; or to fraudulent transactions on another’s application. Remember we’re all in this together: the security ecosystem, Kumbaya and trust falls folks...

## Appendix F: Year in review

The year began with the Verizon Cyber Intelligence Center (VCIC) tracking incidents that would emerge as 2015's major risk trends. We were seeking actionable intelligence from the mega-data breach at Sony Pictures Entertainment (SPE) in November 2014. Online wire-transfer provider Xoom was probably the year's first victim of a Business Email Compromise (BEC) to the tune of \$31 million. Palo Alto Networks reported Dridex banking Trojans "began 2015 with a bang." Chick-fil-A and OneStopParking were the victims of payment card breaches which hit the headlines. Sadly, headlines on sites like AOL and Huffington Post also led to the year's first major malvertisement campaign with an exploit kit (EK) attacking browsers with unpatched Adobe Flash Player. Later in **January**, Adobe released a new version of Flash Player to mitigate a zero-day vulnerability being exploited in three advertising networks.

On **February 4**, Blue Cross health insurance member-company Anthem announced they were the victims of a data breach along with almost 80 million people. And on February 27, ThreatConnect reported Chinese threat Actor "Deep Panda" was probably Anthem's attacker. Invincea and iSight partners each released intelligence on a Chinese cyber-espionage campaign that occurred in November 2014. Dyre, Vawtrak and Carbanak joined the list of active banking Trojans. Symantec and Microsoft announced the first major malware takedown of 2015 after the seizure of the infrastructure for the Ramnit botnet. With no arrests reported in the takedown, it came as no surprise Dr. Web reported signs of a Ramnit comeback about a month later.

In **March**, Premera, another Blue Cross member, announced a data breach affecting 11 million people. ThreatConnect's intelligence attributed the Premera breach to Deep Panda. The Mandarin Hotel Group reported a payment card data breach. POS vendor NEXTEP also reported a breach. March's takedown of the "Evolution" deep web marketplace included arrests and it stayed down. A day after the Canadian Security Intelligence Service (CSIS) reported Vawtrak was targeting Canadian banks, AVG reported a Vawtrak campaign collecting banking credentials globally.

Early **April** brought reports that threat Actors in China had launched "Great Cannon" DDoS attacks on GitHub, probably targeting censorship-evasion projects, and Great Cannon also attacked anti-censorship organization GreatFire. The Drudge Report was one of the sites serving up malvertisements leading to an EK and the click-fraud Trojan Bedep. Interpol, Microsoft and several security companies collaborated on two takedown operations seizing the infrastructure hosting the Simda and Beebone botnets. Pawn Storm and CozyDuke cyber-espionage campaigns aligned with Russian national security were the focus of several intelligence reports we collected in April. InterContinental Hotel Group, Sally Beauty and FireKeeper's Hotel and Casino joined the list of payment card data breaches in May. Healthcare sector data breaches proliferated with reports from Partners HealthCare, CareFirst Blue

**JAN**  
**Xoom**  
\$31 million  
business email compromise

**FEB**  
**Deep Panda**  
Likely cause of breach  
with 80 million victims

**MAR**  
**Premera**  
Data breach affecting  
11 million people

**APR**  
**Great Cannon**  
DDoS attacks on  
GitHub, GreatFire

Cross and Blue Shield, MetroHealth and Bellvue Hospital. We collected reports of cyber-espionage attacks on the German Parliament, the Bundestag and Penn State University but details were scarce and actionable intelligence was absent altogether. The banking Trojans leading reports in **May** were Vawtrak, Dyre and Tinba.

Health insurance breaches were bumped off the top of the headlines for mega-breaches in **June** when the US Office of Personnel Management (OPM) reported another breach. OPM had been breached in March 2014 according to a New York Times report. The initial tally for the 2015 OPM breach was 4 million persons, but eventually grew to 21 million. ThreatConnect was able to connect the OPM breach to Anthem. Fortune magazine published a four-part investigative report on the SPE breach. Wired and Der Spiegel published reports on the cyber-espionage attacks on the Bundestag initially reported in May. Cisco reported three security products had a common default Secure Socket Shell (SSH) key for remote support.

**July** ushered in a bonanza of data breach reports including Harvard University, a second breach at Penn State University, Trump Hotels and UCLA. Two other breaches would echo for several weeks. Social network/online dating site Ashley Madison suffered a data breach and almost 100 GB of stolen data was exposed. Italian security and surveillance company Hacking Team was also breached and 400 GB of data was exposed. Events would unfold and reveal several previously unknown vulnerabilities in Hacking Team's stolen data.

The breach bonanza continued in **August** with reports from American Airlines, the US Department of Defense, the US Department of Health and Human Services and the US Internal Revenue Service. The data breach at Carphone Warehouse was the first report the VCIC collected of a compound attack when the victim is targeted with a DDoS attack to occupy and distract defenders while a data breach attack is launched. Wireless networking company Ubiquity reported it was the victim of a \$47 million BEC. AOL and the Huffington Post were serving up malvertising again. Another malvertising campaign struck MSN, Telstra and dating site PlentyofFish.com.

New intelligence on the Chinese cyber-espionage Actor Blue Termite emerged in **September** in multiple reports of attacks on Japanese companies. Proofpoint contributed a report on a different Chinese cyber-espionage operation targeting Russian military and telecoms. Yet another Blue Cross and Blue Shield member reported a data breach when Excellus announced a breach that began in December 2013 compromising the PII and personal financial information (PFI) of 10 million people.

Data breach reports resumed in **October** when Experion reported their system with personal information for 15 million T-Mobile customers had been breached. UK wireless provider TalkTalk and four million of its customers made up another breach reported in October. The Daily Mail exposed as many as 15 million visitors to malvertisements. Trend Micro connected Pawn Storm to multiple attacks using Adobe Flash and Java vulnerabilities first discovered in the Hacking Team data cache. Another major botnet takedown took place with seizure of the Dridex banking Trojan's infrastructure and arrests of Andrey Ghinkul, Dridex's author.

In early **November** the VCIC began collecting intelligence that Dridex was recovering and resuming operations. Extortion DDoS threat Actor "The Armada" appeared on the scene attacking several email service providers. Indictments for the criminals responsible for 2014's breach of JP Morgan Chase were made public revealing the bank attacks were part of a stock fraud scheme. Australian grocery retailer Farmer's Direct reported the breach of the account registration information of more than 5,000 customers, but their payment information was not compromised.

● **MAY**  
● **Healthcare**  
Data breaches cause problems for insurance providers

● **JUN**  
● **OPM breach**  
21 million victims

● **JUL**  
● **Ashley Madison**  
100 GB of stolen data in high-profile compromise

● **AUG**  
● **Ubiquity**  
\$47 million business email compromise

● **SEP**  
● **Blue Termite**  
Chinese cyber-espionage attack on Japanese companies

● **OCT**  
● **Experion**  
Breach affects 15 million customers

● **NOV**  
● **Dridex**  
Banking malware shows up again



It seems every year ends with the InfoSec community fixated on the most-recent mega-breach. In **December**, it seemed that it would be the breach at the Australian Bureau of Meteorology (BOM). Leaks from the investigation attributed it to Chinese threat Actors. Virtually no details accompanied any reports or leaks from the BOM breach. Malvertisements struck The Independent, The Guardian and The Daily Motion. Juniper reported the discovery of backdoor vulnerabilities in ScreenOS. As the month and year were winding up, news broke of power outages that occurred on December 23 in Ukraine. BlackEnergy malware was found on systems in Ukrainian power companies. It was this breach that the VCIC and many of our colleagues in InfoSec were focused on at the end of the year.



# DEC

● **BlackEnergy**

Malware causes power outages in Ukraine