



**BULLETPROOF ANNUAL
CYBER SECURITY
REPORT 2019**

www.bulletproof.co.uk

CONTENTS

3 A NOTE FROM THE MD

5 EXECUTIVE SUMMARY

Top targeted industries

Looking back at 2018

7 THE THREAT LANDSCAPE

Everything changes

Ransomware moves over for cryptomining

Cryptomining profitability

Summer to winter

Card skimming and Magecart

Nothing stands still

9 PENETRATION TESTING

Our results

Analysis

Frequently seen faults of 2018

Out-of-date software

Annual patterns

More vulnerabilities than ever

Compliance to drive improvement?

XSS

15 SOCIAL ENGINEERING

Payloads

Our phishing results

Summary

17 MANAGED SIEM AND THREAT HUNTING

False positives, open events and action taken

What were they?

SIEM has come a long way

21 CONCLUSION

Skeletonscare

AI and the scammer

IoT devices

The far future

A NOTE FROM THE MD

I'm sure many of you are currently suffering from report fatigue. You could be forgiven if you've started to feel as though most yearly round-up reports take a long time to say the same thing as everyone else. However, I can't help but think that that is quite interesting. Everyone in the security industry is largely saying the same thing and have been for a few years now.

At Bulletproof we are always pushing for innovation within the security industry. We don't want to be just another security provider. We want to make a difference, and I feel we really can. This report suggests that there are numerous areas we can really improve upon and are ripe for a little innovation. My concern when reading these security reports is that it can often seem like we are just covering old ground.

Take the OWASP top 10 for example. This hasn't really changed much over the years. Year in, year out we find ourselves talking about the same old threats. And yet, it hasn't got any easier to educate staff, and businesses still struggle to develop secure applications. This is because, although the threats haven't changed much, the environments around us have. They are getting more complex and this is leading to weaknesses.

As I speak to more businesses it becomes more apparent that people understand the importance of security. However, this usually leads businesses to focus on becoming compliant in various schemes such as ISO 27001 or PCI for example. This is perfectly understandable as not only does it feel like businesses should be more secure, but it also means they have a certificate they can proudly show to their customers to prove they are. Whilst compliance is to be encouraged and does provide numerous benefits if it's done right, some schemes are failing to keep up with the times. Also, different schemes have different criteria and, frustratingly, businesses often become quite lax in their processes and update schedules once they've got their



Oli Pinson-Roxburgh
Managing director



“Our security monitoring services are heavily weighted towards threat hunting activities, as relying on the tools alone is no longer enough.”

certificate. A combination of issues means that despite being compliant with various schemes, businesses are not really addressing the real issues with their security. This is all demonstrable in the fact that we are still seeing big companies getting hacked, often through well-known weaknesses.

How do we solve these security issues when it's like the floor is constantly moving beneath us? We have seen a huge uptake in cloud technologies and serverless is starting to make real traction. Whilst on the face of it, it seems like these will help simplify your operations, it often makes things more complex and confusing from a security perspective. How do you keep track of what you have? Who is responsible for what? As our data has shown, businesses regularly fail to keep traditional infrastructures secure, why would we be any better at securing the cloud?

From our own data, we can see that there is some synergy between common flaws identified through our penetration tests and the methods of attack Bulletproof analysts have blocked via our SIEM. This proves that automated tools can help detect attacks, but ultimately people need to be watching in order to head them off before they become a problem.

Spotting suspicious and stealthy attacks is key. Unlike a lot of providers, our security monitoring services are heavily weighted towards threat hunting activities, as relying on the tools alone is no longer enough. For example, modern malware is designed to avoid detection for as long as possible in order to mine or steal as much data as possible. Whilst we are making great strides with machine learning algorithms, you will never be able to beat skilled security experts when it comes to keeping ahead of the hackers.

At Bulletproof, we aren't short of knowledgeable security pros and we put them in control of some truly innovative tech and, as this report shows, we are making a difference to our customers.

EXECUTIVE SUMMARY

Cyber security is an industry that never stands still, so it can be hard to keep track of what's been going on. Events occur at such a rapid pace that it seems barely a week goes by without a mega data breach being reported. Bulletproof has sifted through this year's events and delved into our own data to provide a high-level analysis of 2018. What trends did we see? What changes were seen in the hacking community? What changes need to be made in the cyber security industry?

Top targeted industries

The following industries were among the most targeted



Retail
16.7%



Finance & Insurance
13.1%



Hospitality
11.9%



Service Providers (IT)
9.5%



Payment Services
4.8%

There was also a worrying rise in the number of attacks targeting the healthcare industry and government services.

Looking back at 2018

It seems to be “quite a year for cyber security” every year, and 2018 was no different. There were some high-profile data breaches scattered throughout the year from Dixons to British Airways. Marriott took the biggest blow as it emerged that the personal data of over 500 million customers had been compromised. This was particularly interesting, not only due to the sheer number of records, but also due to the fact that the hackers had unauthorised access to the network since 2014.

Bulletproof has been busy protecting businesses across the country with our penetration tests and managed SIEM with active threat hunting. We performed hundreds of pen tests on web apps, infrastructures and mobile applications with a nice mixture of authenticated and unauthenticated tests. Our SIEM filtered through millions of log files and raised thousands of events for our SOC analysts to sift through.

Looking back through our SIEM data, we were pleased to see a grand total of zero security incidents, which shows our SOC analysts are working hard to keep our customers secure. We also discovered that insider threats seem to be the most pressing issue, with 17% of events raised relating to suspicious user activity, and a further 11% from suspicious admin activity. Both of these represent a real danger to businesses as compromised users, be it through stolen credentials or malicious insiders, can cause severe damage.

You'll have seen in other security reports that companies seem reluctant to talk about false positives. To some extent, this is understandable as false positives could be seen as mistakes. However, at Bulletproof, we think a certain number of false positives is a good thing for a number of reasons. For starters, it never hurts to be too careful, and false positives show that investigations into potentially suspicious activity are taking place.

They're also good for making continual improvements to a service. They can help our SOC analysts fine tune the alerting system and provide a more efficient product to clients in the future. That's not to say businesses should be flooded with false positives at all times. They should be kept to a minimum, but they should still be there. It's for this reason we have included our false-positive results in the data shown in this report.

The most frequent flaws found by our penetration testers mirrored our findings in 2017, in that 22% of high and critical-risk issues consisted of missing patches and out-of-date software or software that is no longer supported. Whilst there are numerous reasons for this, not keeping up with patches or replacing unsupported software is just asking for trouble.

5% of critical and high-risk issues involved poor or default passwords, which is the worst of all the cyber sins. Other frequent finds included Cross Site Scripting (XSS) and SQL injection vulnerabilities that have the potential to leak sensitive information, both of which are avoidable.

At Bulletproof, we truly live and breathe cyber security. In looking through the data and findings documented within this report, it was encouraging to see that we were making a difference to our customers. We were able to detect the prevailing trends of 2018 which will help us when looking to the near future.

Cyber security is an ever-changing environment that's hard to predict. However, if pressed, we'd say that we might see an improvement when it comes to patching, what with the risk of hefty GDPR fines if a serious breach is found out to be due to poor patch management. We expect to see more instances of card skimming in the early days as companies start to uncover more compromised payment pages. It wouldn't be overly surprising to see a decline in cryptojacking. The fluctuating value of digital currencies means these operations are inherently uncertain. If it ceases to be profitable, hackers will stop doing it. One thing is for certain though: 2019 will be quite the year for cyber security.

“In sifting through the data and findings documented within this report, it was encouraging to see that we were making a difference to our customers.”

THE THREAT LANDSCAPE

Everything changes

The cyber threat landscape doesn't stay the same for very long. Popular attack methods come and go and sometimes even come back again, often bigger, better and scarier. 2017 was very much the year of ransomware. At one point, 6/10 payloads contained ransomware.¹ There are numerous ransomware families out in the wild, all with multiple variants, but the big players last year were WannaCry and Petya/NotPetya which received a considerable amount of press coverage. This led to many UK businesses stockpiling Bitcoin to pay off ransomware attacks², no doubt thinking that would be cheaper than investing in cyber security.

Ransomware moves over for cryptomining

As we moved into 2018, cases of ransomware plummeted. This was a trend that we were already seeing towards the end of 2017. By December, ransomware infections were scoring a mere 10% infection rate³, whereas cryptomining malware was rising rapidly. In fact, there was a staggering 629% increase in cryptomining malware in Q1 of 2018 compared to Q4 of 2017⁴. Furthermore, research has revealed that 59% of UK companies have been hit by cryptojacking of some form.

Cryptojacking is the act of hijacking a user's CPU to mine for cryptocurrencies, and it's not just the work of a few dedicated organisations. In the first half of 2018, 47 new families of cryptomining malware were discovered⁵, showing that there are plenty of people giving it a go.

This is not to say that ransomware was not a threat this year, just that it was not as prevalent.

There is some case to be made that the remaining examples are using increasingly sophisticated software⁶ and techniques when conducting these types of attack. For example, some ransomware creators are coding their software to slow down the encryption rate, keeping below the threshold of any detection tools. Some are deploying strains of ransomware that make slight changes to its code as it spreads to another victim so that it's harder for anti-virus systems to detect them.

This trend can be at least partially attributed to the mass media exposure ransomware received last year. A raised awareness would have led to greater diligence. Businesses would have felt more pressure to patch their systems against certain threats. There's also the fact that very few businesses who paid the ransom (19.1%)⁷ actually got their files back. Many discovered that the 'ransomware' infecting their network was in fact just posing as such and simply deleted their files. Others discovered that groups illegally infecting businesses with malware simply weren't true to their word, which I'm sure came as a shock to all. This will have had an impact and discouraged businesses from paying ransoms, making ransomware no longer as lucrative.

With the goal of ransomware being to extort a payment in Bitcoin or other cryptocurrencies out of businesses, hackers no doubt realised it was easier to exploit a company's systems to mine directly. The cryptomining craze started with Bitcoin. Originally, this could be done with home mining rigs, but due to being subject to artificial scarcity, more CPU power is now needed to mine efficiently. The average home mining rig is no longer sufficient, but a huge rack of servers is ideal.

“629% increase in cryptomining malware in Q1 of 2018”

Cryptomining profitability

Monero seems to be the digital currency of choice⁸. This may be partially due to the algorithm (Cryptonight) used to mine Monero (XMR) being well suited for browser-based mining⁹. 2018 saw hackers take to injecting a line of malicious JavaScript into websites to siphon off visitors' CPU. The benefit of this method is it's difficult to detect, though thousands of websites are suspected to contain such scripts¹⁰.

Whilst this is a concern for the time being, the fluctuating value of cryptocurrencies means cryptojacking's future is uncertain. At the time of writing, the XMR/GBP exchange rate sits at £47.61. One of the largest Coinhive campaigns to date, consisting of over 11,000 parked websites earned just \$7.69 over a period of three months¹¹. The less profitable cryptojacking becomes, the fewer people will be drawn to it.



At Bulletproof, we aren't short of knowledgeable security pros and we put them in control of some truly innovative tech and, as this report shows, we are making a difference to our customers.

Summer to winter

We saw a sharp rise in fileless malware in the first half of 2018¹². Whilst not new, this kind of malware is becoming increasingly more advanced. One of the main concerns is that fileless malware makes threat hunting difficult, and any sort of forensic investigation close to impossible as next to no evidence is left behind.

In October, Bloomberg released an article¹³ stating that motherboards imported from Supermicro in China were coming complete with tiny spy chips. The widely discredited article stated that these chips would send information from leading western brands, including Amazon and Apple. Whilst all companies involved outright denied this was the case, and no supporting evidence has been forthcoming, this article managed to knock a considerable amount off of Supermicro's share price.

Despite this particular case being seemingly untrue, it does bring the supply chain into focus. Towards the beginning of the year, there was still much discussion being had concerning new Meltdown and Spectre vulnerabilities. These showed that serious flaws can exist at the hardware level and there will almost certainly be others found in the future. Therefore, it's not unreasonable to suggest that the future may well hold incidents that focus on this area, particularly with the rising concern of state-backed attacks.

Card skimming and Magecart

Throughout the year we saw numerous data breaches. The largest by far was the data theft of 500 million Marriott customers. Others that stick out are British Airways (in which 380,000 payment cards were compromised), Ticketmaster (40,000 card details) and Dixons (105,000).

¹ <https://blog.barkly.com/ransomware-statistics-2017>

² <https://www.citrix.com/blogs/2017/06/06/ransomware-in-the-uk-one-year-on/>

³ <https://blog.malwarebytes.com/cybercrime/2018/02/ransomwares-difficult-second-album/>

⁴ <https://www.mcafee.com/enterprise/en-us/assets/reports/rp-quarterly-threats-jun-2018.pdf>

⁵ <https://documents.trendmicro.com/assets/rpt/rpt-2018-Midyear-Security-Roundup-unseen-threats-imminent-losses.pdf>

⁶ <https://blog.malwarebytes.com/cybercrime/2018/02/ransomwares-difficult-second-album/>

⁷ <https://cyber-edge.com/wp-content/uploads/2018/03/CyberEdge-2018-CDR.pdf>

⁸ <https://www.newsbtc.com/2018/08/15/report-cryptojacking-coincides-with-crypto-popularity-monero-the-choice-currency/>

⁹ <https://coinhive.com/#hash-rate>

¹⁰ <https://scotthelme.co.uk/protect-site-from-cryptojacking-csp-sri/>

¹¹ <https://arxiv.org/pdf/1803.02887.pdf>

¹² <https://www.sentinelone.com/blog/fileless-malware-changes-way-treat-cyber-threats/>

¹³ <https://www.bloomberg.com/news/features/2018-10-04/the-big-hack-how-china-used-a-tiny-chip-to-infiltrate-america-s-top-companies>

¹⁴ <https://www.riskiq.com/blog/labs/magecart-ticketmaster-breach/>

For a number of these hacks, card skimming was to blame. Many point the finger towards Magecart, which became one of the most common methods of stealing card data towards the end of the year.

Card skimming, wherein hackers steal data as it's inputted by compromising forms with, yes you guessed it, malicious JavaScript code, is not a new phenomenon. However, the Mage.js script was identified on over 800 e-commerce sites in 2018¹⁴ and, chances are, there are a lot more that haven't yet been discovered. Frustratingly, even those who are PCI compliant can still be targeted.

We are likely to see more reports of this in 2019. Not necessarily because attacks will continue or increase (though it's likely), but because it can take a long time for a business to become


aware that they have been breached. This is largely because hackers are not usually directly attacking the vendor. In worst case scenarios it can take months or even a year to detect this kind of attack, or any kind of breach for that matter, as is apparent with the case of Marriott. Hackers had unauthorised access to their database for four years. It's also worth noting that there are multiple groups using Magecart and the more successful stories we hear, the more other groups will try their hand at it.

Nothing stands still

So, as can be seen, the threat landscape can change drastically in the space of a year. As soon as industries become aware of an attack vector and adopt adequate defences against them, hackers look for ways to overcome these defences.

“The largest by far was the data theft of 500 million Marriott customers.”

PENETRATION TESTING



Our results

Our penetration testers have been working flat out this year probing infrastructures and testing apps up and down the UK. We've seen well-managed environments and carefully designed apps, as well as infrastructures and applications riddled with flaws. There have been instances where we've been able to waltz into a company's systems, escalate privileges and view all sorts of sensitive data. Whilst this makes for more interesting reading, it ultimately means that businesses are in for some heavy losses if they fail to make the right changes.

We found issues in every single one of our penetration tests conducted this year. Whilst some are far better than others, no app or infrastructure is perfect. We looked over all the results from each individual test and broke them down by severity.

¹⁴ <https://www.riskiq.com/blog/labs/magecart-ticketmaster-breach/>

Analysis

Whilst critical risks occupy the smallest wedge of the pie as we'd expect, 6% is still worryingly high. The presence of such risks can often indicate that malicious parties can access sensitive data from an unauthenticated perspective.

We at Bulletproof define a critical vulnerability as, "a serious and immediate risk of compromise to both systems and data." Following the Common Vulnerability scoring system, a critical risk gets a score of 9-10, so it's the worst one. If there's a critical flaw in your infrastructure or application, you're in trouble. We uncover high-risk issues more often, defining these as indicating a "serious weakness or exposure that should be addressed immediately." Whilst hackers may have to work a little harder to exploit these issues, a high risk will lead to compromise.


Medium risks should certainly be remedied as soon as possible. In isolation, they may not allow a hacker direct access to the network, but can lead to unwanted functionality from an app. Also, if exploited in conjunction with other vulnerabilities or combined with an element of social engineering, or if enough is known about the environment beforehand, medium vulnerabilities can pose just as much risk a high or critical vulnerability. If more information comes to light in the future, a medium risk can become much more severe. Put simply, why build a business on issues you know to be there?

Frequently seen faults of 2018

With critical and high risks posing an immediate risk to environments and the data they contain, we had a look to see if there were any patterns amongst them. The most frequent issues we saw involved missing patches, out-of-date software or software that is no longer supported. Being similar in nature and, in most cases, leading to the same outcome, we have grouped all of these together. Following this, making up very similar sized chunks, were Cross Site Scripting (XSS) vulnerabilities, default/poor passwords and SQL injection vulnerabilities.

Whilst there were a number of one-off vulnerabilities unique to particular apps or environments, these are the faults that occurred most frequently:

2018 most frequent critical/high faults



It should be immediately obvious why two of these issues are alarming. Default credentials on server components or anywhere in a network should never be used. Practically all default credentials are publicly available in some form. There are even sources that collate them together in one place (<http://www.defaultpassword.com/> for example). If malicious actors can learn more about a network's configuration and set up, these credentials will be tried. It's surprising just how many organisations have integral pieces of kit still allowing the default credentials, practically giving hackers complete control. Providing hackers with free reign over your environment is bad enough, but the use of default credentials also makes things more complicated from a monitoring perspective. If a login looks legitimate, it'll take longer to work out something is amiss.

Out-of-date software

Perhaps more concerning is the number of businesses that had unpatched, out-of-date or unsupported software in use somewhere within their infrastructure or app. Of course, software that is simply 'no longer supported', is not necessarily as dangerous as a missing patch. Patches are often released specifically to address security issues, whereas software that is no longer supported may not have any known exploitable issues. However, it does mean that, should any issues come to light in the future, the manufacturer will not be releasing any fixes.

With so many past examples of organisations getting compromised due to unpatched flaws, one could reasonably question the mindset of those who keep putting patching on the back burner. The infamous WannaCry outbreak springs to mind. We know that this was caused by malicious actors making use of the EternalBlue exploit, a patch for which already existed. It's estimated that 57% of breach victims are breached due to an unpatched vulnerability.¹⁵

"It's estimated that 57% of breach victims are breached due to an unpatched vulnerability."



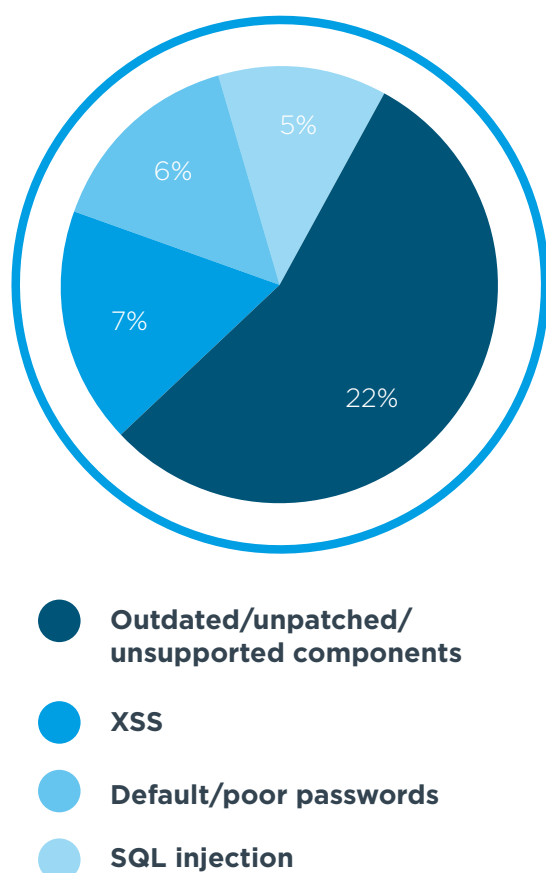

```
#selection at the end -add back the deselected mirror modifier object
mirror_ob.select= 1
modifier_ob.select=1
bpy.context.scene.objects.active = modifier_ob
print("Selected" + str(modifier_ob)) # modifier ob is the active ob
mirror_ob.select = 0
```

“Systems and applications will remain just as complex and patching will always carry an element of risk.”

Annual patterns

This appears to be a recurring trend. Below are the top recurring high and critical risks for 2017:

2017 most frequent critical/high faults



Negligence is one of them, both on a business and individual level. Businesses overlook them in favour of other things and users tend to be reluctant to shut down their machine to spend ten minutes waiting for updates to download and install. Patching or updating environments is not necessarily a quick win either. It's definitely a win but could well be a time-consuming one. This is particularly the case with software that is no longer supported. Upgrading to something that is supported will require a substantial amount of research and testing to ensure there are no knock-on effects to other applications.

Whilst it's an issue for all companies, it's the smaller ones who are more likely to find getting on top of patching difficult. The larger corporations are likely to have a dedicated infrastructure team who arrange, test and roll-out all relevant patches to a schedule. Smaller companies might not have the staff resource for this and instead have a small team in charge of all things IT, meaning the day-to-day dealings get in the way of the important background events.

It's also a case of business infrastructures becoming more complex, making it harder to keep track of all those patches. It's not just your OS you need to keep up to date. Third-party software also requires diligent patching, as does firmware, application libraries and just about anything else you have running that keeps your business afloat, not to mention cloud and emerging technologies. Each year the list of reported vulnerabilities grows. 2017 smashed previous records with 14,714 reported vulnerabilities. This was beaten in 2018 which hit over 15,000 before December.¹⁶

The exact same percentage was down to a lack of patching or up-to-date software. Bar the odd 1%, the other sections are the same too.


It is Bulletproof's assumption that we'll have just as many, if not more penetration tests littered with outdated or unpatched services throughout 2019. There are a number of contributing factors as to why this problem has become what it is.

¹⁵ <https://www.servicenow.com/content/dam/servicenow-assets/public/en-us/doc-type/resource-center/analyst-report/ponemon-state-of-vulnerability-response.pdf>

¹⁶ <https://www.cvedetails.com/browse-by-date.php>

More vulnerabilities than ever

Reported CVEs



The recent boom in popularity of IoT devices will have contributed to the number of these reported vulnerabilities. Often IoT devices find their way into office environments and no one gives it much thought. It could be a cool coffee machine that can be interacted with via a Wi-Fi app or an air-conditioning unit.

Everything needs to be thoroughly hardened before it's incorporated into a network as they can offer a substantial foothold on the whole environment, as proved when a casino was hacked via a fish tank¹⁷.

Whilst not all of the 15,000+ vulnerabilities yet have a patch or will indeed receive a patch, chances are a great deal of them will have been released. If the company is large, keeping up with all of these and organising an appropriate schedule to get them installed requires skills, knowledge and time to adequately manage. This is something many businesses feel they simply don't have the resources for.

Then there is the fear of patching. In a perfect world, each and every update would install quickly, solve any underlying issue and that would be that. However, the very real fear of downtime or technical issues as a result of the patching leaves some businesses reluctant to push updates out. With so many apps and services designed for interoperability, a change to one piece of software can have a drastic impact on another.

“With penetration testing, we tend to press the importance of testing the internal environment as well as the external facing systems. This simulates the damage a hacker can do if they get past your perimeter defences, or if they acquire a user's credentials.”



Compliance to drive improvement?

Ordinarily, we would say that there's unlikely to be a dramatic improvement here as we move into 2019. This is because systems and applications will remain just as complex and patching will always carry an element of risk. Generally speaking, businesses start by making do before they have the resources to make do properly.

However, May 25th saw the much-talked about GDPR become fully enforceable. This means that any data breaches that put personal information of European and UK citizens at risk could lead to hefty fines. If the ICO find out that the breach was due to out-of-date software, then they are less likely to be lenient. The cost of updating won't seem as bad when faced with potential multi-million-pound fines. With GDPR and other compliance standards raising the need for regular vulnerability scans and penetration tests, it's likely that these issues will be picked up upon and remedied.

XSS

Whilst a bit more complicated than unpatched software or default passwords, Cross Site Scripting (XSS) is not difficult to defend against. It's also worth noting that an XSS vulnerability can in fact vary in severity. In some cases, they can be little more than a nuisance and in other cases, they can be quite severe and lead to users accessing sensitive data.

If the correct validation and sanitisation is occurring, then XSS should not be a threat. It is odd to see these vulnerabilities occurring so often. We put it down to simple oversight more often than not. In a rush to get a web application live, this simple flaw is often forgotten or simply not even thought of.

With penetration testing, we tend to press the importance of testing the internal environment as well as the external facing systems. This simulates the damage a hacker can do if they get past your perimeter defences, or if they acquire a user's credentials. We've seen some well segmented and protected internal infrastructures in 2018, however, the prevailing trend was to see weakened internal infrastructures behind moderate external-facing services. We tended to find more high and critical issues when conducting an internal test. Anyone can gain access if they're dedicated enough. In fact, it's not necessarily that hard to gain access to an environment, as we discovered.

"Anyone can gain access if they're dedicated enough. In fact, it's not necessarily that hard to gain access to an environment, as we discovered."

¹⁷ https://www.washingtonpost.com/news/innovations/wp/2017/07/21/how-a-fish-tank-helped-hack-a-casino/?noredirect=on&utm_term=.9358ced39815

SOCIAL ENGINEERING

When scoping for our penetration tests, we always encourage including an element of social engineering. Leveraging the human element is rapidly becoming one of the most efficient ways of compromising a network. The most successful social engineering tactic is a good old-fashioned phishing campaign. This approach is still one of the most popular methods for circumventing security. It's a pervasive threat that won't be going away any time soon. In Q1 of 2018, rates of spam or phishing emails seemed to be going down ever so slightly. However, going into Q2 this saw cases of web-based social engineering rise substantially and this increase continued into Q3, growing by a noteworthy 233%¹⁸.

There are a number of factors that come into play here. Isolating one reason to explain a sudden spike is next to impossible. For starters, it could be down to the natural ebb and flow of malicious emails and messages. It could be that there were a number of particularly severe data breaches in 2018, which would have led to a large quantity of active emails becoming known to spammers and scammers. There is some weight to this idea as roughly 4.5 billion data records were compromised in the first half of the year alone¹⁹. Whilst that doesn't mean 4.5 billion individual email addresses became widely known, it does equate to a frighteningly high number. New malware can trigger groups to start sending out batches of malicious emails too, not to mention that it just takes one hacking group to start a particularly wide-spread campaign to see figures start to shoot up.

Payloads


The preferred delivery method of payloads seemed to switch between malicious attachments and URLs periodically throughout the year. The figures from Q3 shows that the use of malicious URLs outnumbers attachments²⁰. This is partly due to the fact that such links could be sent via messaging apps. With businesses using apps such as Slack or Teams, it's not unreasonable to assume that these could easily become attack vectors.

Back in Q1, the most common malicious email attachment was a Trojan. Specifically it was Trojan-PSW.Win32.Fareit, which is primarily used to harvest user data²¹, such as account data for cloud storage services, browser cookies and account data for mail clients. Q2 saw this replaced with Exploit.Win32.CVE-2017-11882.

This can lead to a command and control RAT being executed and continued to reign supreme in Q3²².

Our phishing results

Through 2018, our penetration testers sent thousands of targeted phishing emails. Claiming to be from the target company's IT team, spoofing their email where we could, we would advise the users that, due to a security event, they needed to change their passwords. We would provide them with a link to a mock portal designed to look like their own, or Outlook's Web App.



Only one set of credentials is required to get authorised access to a network. Once inside, users can drop malware, insert a backdoor or make off with any data available to them. If combined with issues existing within the internal infrastructure, then privilege escalation and data theft will be a simple task, not to mention a great deal of sensitive information will be available via a compromised email account, which could then be used for spear phishing.

Out of thousands of emails, 7% gave us their credentials. On average, 14.5 billion spam messages are sent a day, of which 73% are phishing emails²³. That's a lot of credentials at risk. If we take our success rate as a model, that's roughly 740,950,000 credentials that would have been compromised. Of course, not all phishing campaigns are made equal. Some are more convincing than others, and some will be caught up in spam filters.

The only real defence against phishing is education. Unfortunately, we don't foresee a situation where no one falls for a phishing email, at least not in the near future. As we go into 2019, we expect phishing to remain a dominant threat in the cyber security industry. It's a pattern that seems to remain consistent year in year out. The weakest link in the security chain is people. Awareness is growing, as we have seen with an increased interest in our cyber security training packages. However, we don't expect to see any immediate changes. There is a reason phishing hasn't slowed down over the years: it works.

Certain payloads and attachments are likely to fall out of style and perhaps the preference for malicious URLs will endure. It's reasonable to suggest we may start to see a decline in the use of attachments, though the widespread use of PDFs and Microsoft products in business environments means they're not likely to ever go away completely.

Summary

Throughout the year we saw a number of more interesting, standalone high or critical-rated flaws, but they are scattered here and there. Due to human nature and the intricacies of typical infrastructure setups, it's perhaps not all that surprising to see that the majority of critical or high-risk flaws putting data at risk are easily avoided or down to simple mistakes or misconfigurations. These flaws shouldn't really be there, but they are, and will keep coming up for the time being.

“There is a reason phishing hasn't slowed down over the years: it works.”

¹⁸ <https://www.cbronline.com/news/global-data-breaches-2018>

¹⁹ <https://blog.barkly.com/phishing-statistics-2018>

²⁰ <https://threats.kaspersky.com/en/threat/Trojan-PSW.Win32.Fareit/>

²¹ <https://securelist.com/spam-and-phishing-in-q3-2018/88686/>

²² <https://securelist.com/spam-and-phishing-in-q3-2018/88686/>

²³ <https://www.spamlaws.com/spam-stats.html>



MANAGED SIEM AND THREAT HUNTING

The interesting thing about providing 24/7 monitoring and threat hunting for a variety of businesses is we not only get to see activity across different networks, but different industries too. Over the course of 2018, Bulletproof’s SIEM filtered through millions of alerts and our dedicated SOC analysts raised thousands of events. Whilst there are some differences between clients in terms of thresholding and what they want raising as events, most businesses want similar information. Is this traffic malicious? Are restricted files being accessed? Is malware getting through the firewall? And so on.

Our raised events, once investigated by analysts, fall into the following categories:

Area	Description
Security Incident	Events that are escalated via email and/or phone and require immediate attention, as there is a confirmed security incident/violation.
Open Events	Events that are escalated via email and are continuously investigated by the SOC to verify and give updates on the activity spotted.
Action Taken	Events that are escalated via email and an action is taken by COMPANY to mitigate the alert.
False Positives	Events that have been investigated and have been found not to pose any security related concerns. These are raised on the portal and archived in case of future reference.
Warnings	Events that have been investigated and have been found not to pose any security related concerns. These events usually suggest that changes on the alert threshold must be made.

Security incidents are the most serious ones. The whole purpose of a SIEM and active threat hunting is to reduce the risk of these occurring and help contain and remove any threats as soon as possible, keeping damage

and downtime to a minimum. We’re pleased to say that out of our thousands of raised events, none of them were security incidents. Our SOC analysts have certainly been doing their jobs.

False positives, open events and action taken

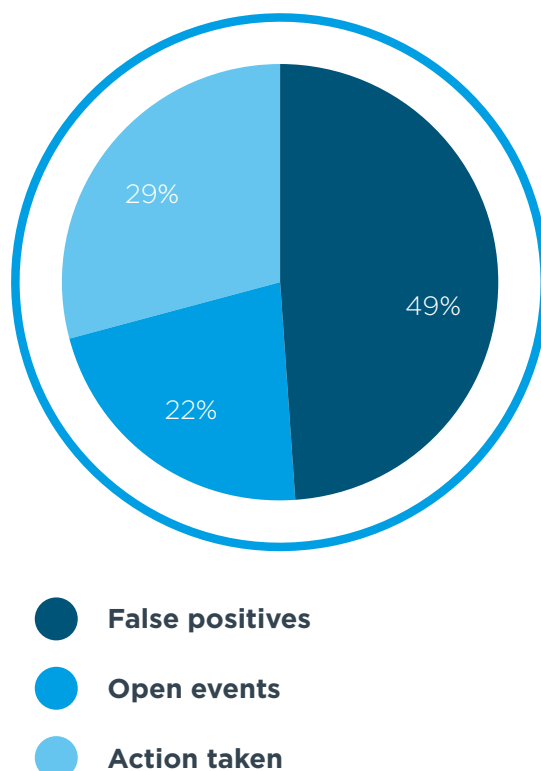
Unsurprisingly, false positives make up the largest portion of our results. These can be anything from suspicious IPs making contact with the environment, which turn out to be completely legitimate, to out-of-hours logins that are revealed to be down to emergency or planned maintenance. False positives are inevitable with threat monitoring. Unfortunately, this is something that has often put people off investing in a quality SIEM system, as companies fear getting swamped by alerts that turn out to be nothing.

We have fine-tuned our platform with the help of machine learning algorithms, to keep false positives to a minimum. However, they'll never go away completely. If a company isn't seeing it's fair share of false positives, then they should be questioning their level of monitoring. It's better to investigate something that turns out to be nothing than ignore something that turns out to be disastrous.

'Open events' are the most important category in the chart, as they could be indications of compromise. This is why regular communication and constant monitoring of these is crucial.

On a month-to-month basis, the number of open events stays relatively consistent very rarely rising above 5% of our monthly events.

Breakdown of events raised 2018



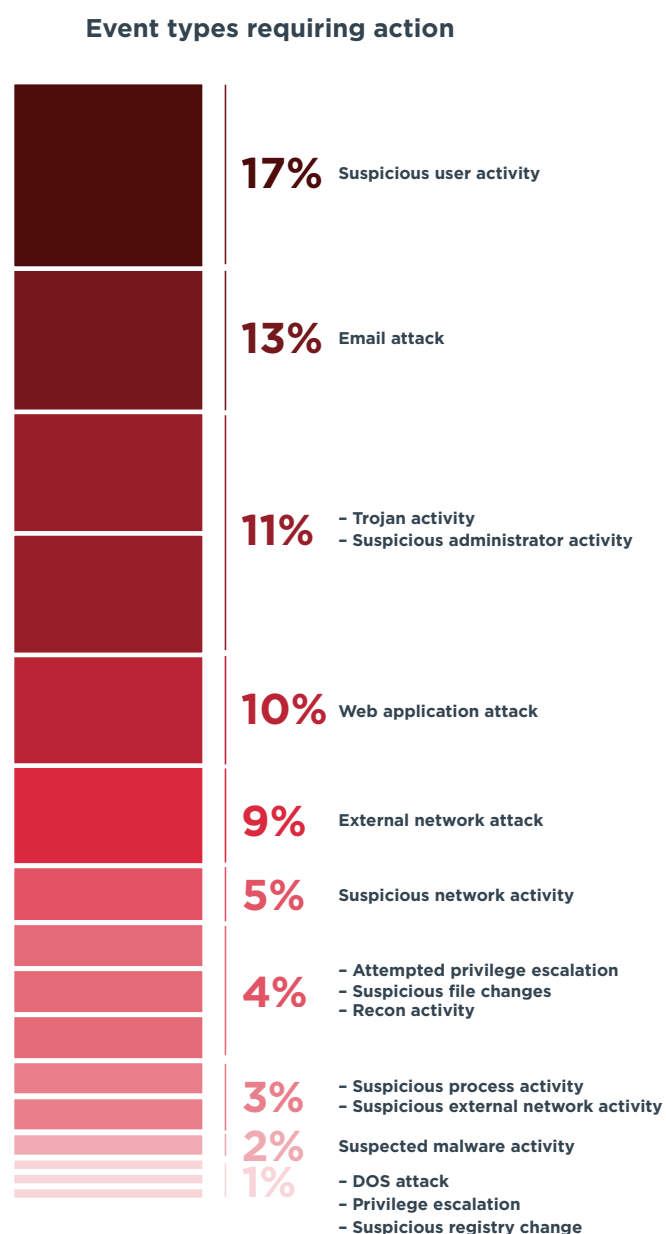
“We’re pleased to say that out of our thousands of raised events, none of them were security incidents.”





What were they?

Throughout the year, hundreds of events required action to prevent potentially malicious incidents on our clients' networks. That's hundreds of potential data breaches or attacks blocked. The type of events seen over the year can be broken down as below:



The largest portion is for suspicious user activity. There's a variety of actions that could fall into this, from logging into environments not usually associated with the user in question to accessing files unrelated to the user's job role or logging in at bizarre hours. It could even indicate that an account has been compromised and is being used by a malicious entity. It could indicate foul play on the employee's part or it could be innocent but unpermitted activity. Either way, it highlights that the most persistent threat in any organisation lies within. It is unfortunately the case that a compromised account or an employee willing to abuse their access can end up doing the most damage to an environment. A further 11% was suspicious administrator activity, which is more worrying. Occasionally, these alerts can be triggered by a user using their admin account instead of their regular one. This is not sticking to best practices and these events should always be investigated.

There were times throughout the year where this wasn't the case and required intervention from our clients. Obviously, the severity of a compromised administrator account, be it through hacking or malicious insiders, cannot be overstated. Administrator accounts will have near unrestricted access to data and will be able to implement changes on a network. Any suspicious administrator activity needs to be treated as a high-priority incident and investigated immediately.

Other commonly occurring events were 'email attacks' and 'privilege escalation attempts'. Email attacks could consist of many things from spam and identity theft to the dropping of malicious payloads. These represent a persistent threat worldwide to just about every industry.

Privilege escalation may well indicate compromised accounts attempting to gain admin access and also suggests a poorly managed AD environment. The more paths there are to admin access, the easier it is for a hacker to move up the ladder. Detecting and stopping these attacks is crucial. Fortunately for our clients, with the right monitoring in place it's easy to spot and prevent these attempts.

Although only occupying a tiny percentage of events, preventing DoS attacks is key. Denial of Service attacks could cost a business a lot of money. It's estimated that a small company could lose up to \$120,000 if they fell victim to a DoS or DDoS attack, and for an enterprise it can be more than \$2 million²⁵. It's interesting to see the number of DoS and DDoS attacks are among the lowest of recorded events. With the rise of botnets such as Mirai using insecure IoT devices to conduct widespread and frequent DDoS attacks in 2016 and 2017, fear of these attacks was justifiably high. However, the rate of attacks actually fell in 2018. We should hold on the celebrations however, as whilst the frequency rate was down, the size of attack was up, with a record 1.7 Tbps attack recorded²⁶.

It seems hackers are capable of devastating attacks and have some quite powerful tools at their disposal. With that in mind, we need to be vigilant moving into 2019 as these hard-hitting attacks might start to come in greater numbers.

SIEM has come a long way

SIEM with active threat hunting has come a long way in recent years and, as we have seen, can be of a huge benefit to businesses. Over half of our events raised over the year, of which there were thousands, needed action or may go on to need action. Without competent alerting, monitoring and investigation these could have all gone on to cause damage to businesses. We've managed to keep false positives down to a minimum, showing that people don't need to fear being bombarded with alerts that ultimately go nowhere. Our results show that 24/7 security monitoring and log analysis could easily become an essential part of anyone's security practices.

Whilst SIEM as we know it will continue to evolve, we at Bulletproof foresee that number growing as current cyber pressures push more companies towards it. Proactive monitoring and investigation can and will keep many businesses secure against a wide variety of attacks.


²⁵ <https://www.netscout.com/threatreport>

²⁶ <https://securityintelligence.com/how-will-you-face-the-high-price-of-ddos-attacks/>

“Technology is developing at a tremendous pace and if cyber security doesn't develop with it, then there could be trouble.”

CONCLUSION

As long as there are ways for criminals to make money out of their activities, they'll continue to develop new strains of malware, discover new flaws or conduct more sophisticated attacks. Historical data shows that cyber security threats have been continually growing since we all became connected and there's no real reason to expect that this will change any time soon. Vulnerabilities, bugs and flaws can be seen as a by-product of innovation and progress. We all do much more online than ever before. We willingly submit data into the ether, relinquishing control over it and placing our trust into others to keep it safe. Whilst it makes our lives more convenient, it means there's more data out there to be had by malicious entities and it's becoming harder and harder to keep track of it all.



“Technology is developing at a tremendous pace and if cyber security doesn't develop with it, then there could be trouble.”

What do we have to look forward to? Well, agile environments, ephemeral systems, more flaws found in hardware, hackers making use of AI and machine learning, the possibilities are almost endless. What we can say for certain however, is that cyber threats are here to stay and there will always be a need for skilled cyber security professionals. We cannot rely on tech alone, which is why Bulletproof will continue to put the right people behind innovative tech to continue to provide the best service we can.

Our penetration testers are picking up on all the things that real world hackers are targeting and our SOC analysts are blocking serious threats before they do any damage to our customers' infrastructure. We are prepared for a busy 2019 and innovating for the future, but what does the future look like? I tend to believe that 90% of security predictions never happen, but here are mine regardless. Let's see next year how well I've done.

Skeletonscare

This year saw a number of instances of what I like to call 'Skeletonscare', in which a scam email makes use of old personal information to add a layer of authenticity (and fear) to a threat. The majority of Skeletonscare emails we saw tended to feature a recurring theme. The email would state that malware had been dropped onto the victim's machine and had given a hacker access to their webcam. The message went on to state that the hacker had compromising videos of the victim accessing adult content and was ready to pass them on to their contacts unless they paid a significant sum in Bitcoin. What sets this apart from other scams is the message contained a password once used on the victim's account (or an account at least). This password would often be out of date, but would have been used once. The logic here being, 'if they know that information, maybe the account is compromised.' We know that

there was no compromise and that malicious parties were just using information leaked in previous data breaches in order to convince their victims that they had something, praying on the skeletons they may have in their closet (hence Skeletonscare). There have been some claims that these campaigns have earned over \$50,000, showing that some people have enough of a guilty conscience to pay up.

We found a public database which contained over 1 billion usernames and passwords, and we know there are a number of Wikis on the darkweb with similar volumes. The issue is, some data breaches contain more information than others. The Marriott leak, for example, contained all sorts of info, from check-in and departure dates to passport numbers. That's a lot of info that will lend some believability to a scam. I believe we'll start to see a rise in these sorts of scams, some growing in sophistication.

AI and the scammer

Conducting a good phishing campaign which makes use of data pulled from previous breaches is going to be time consuming, particularly when you factor in that not all emails found via data breaches are going to be active. In order to have any amount of success, they'll need to be thousands of emails produced. Unfortunately, AI and machine learning are developing at such a rate that it's not unreasonable to assume that, eventually, intelligent and affordable tech will help automate this process.

AI could pull all the required data found in breach databases (usernames, passwords, addresses, passport information or anything else that could be useful) along with any useful information found on social media sites into an email template. These can then be sent en masse. Really, the technology required to do this already exists, it's just a question of whether it's affordable or worthwhile for hackers to adopt this process over cryptomining or ransomware.

IoT devices

Smart home hubs are becoming increasingly popular allowing people to stream music, use search engines and control the lighting with voice commands alone. The likes of the Google Home Hub and Amazon Echo are amazing pieces

of technology that can integrate seamlessly with the rest of the home's Wi-Fi enabled products. They are also vulnerable. If your smart home hub gets hacked, then the hacker could listen to every conversation you have nearby. Think of the data they could harvest from that.

The far future

This year saw Google develop AI that successfully booked a haircut through a conversation (<https://www.youtube.com/watch?v=tYJ-stQfD4A>). It responded quickly to any questions asked and, presumably, the person on the other end of the phone was unaware that they were speaking to a machine. Whilst this is good news for anyone looking to book a haircut with minimal human interaction, it's a hypothetical nightmare for cyber security.

Theoretically, the same technology could be applied to the phishing campaign or the extortion scam. People are far more likely to fall for something if they think they're convinced they're speaking to an actual person. Imagine a piece of AI with access to information found through the Marriott data breach (usernames, passwords, passport numbers, arrival & departure dates etc). Whether it's a phone call, an online chat service or a string of emails, a smart piece of AI fed the right data could easily convince people they are speaking to legitimate persons and subsequently trick them into making a payment or giving away their more sensitive data.

To take it to the extreme, we've recently seen actors brought back from the dead on the big screen by clever CGI. This sort of technology can only get better. It's not beyond the realms of plausibility to think one day, we could see an extreme form of identity theft involving video conferences with AI fed CGI CEOs. This is obviously veering into the realms of fiction and we don't foresee these so-called 'deepfake' videos happening any time soon. The frightening thing, however, is it's not altogether unbelievable. Technology is developing at a tremendous pace and if cyber security doesn't develop with it, then there could be trouble. Also, if the next series of Black Mirror has an episode featuring a CGI CEO, then they got the idea from here.



 01438 532 900

 contact@bulletproof.co.uk

 www.bulletproof.co.uk