Signals in the Noise:

A Cybersecurity Data Odyssey

Orange
Cyberdefense

orange™

About us

“Orange Cyberdefense is a specialist security research unit within the group that helps us fulfil our mission of being a trusted partner to our customers by ensuring that we identify, track, analyse, communicate and act upon significant developments in the security landscape that may impact them. Our team of dedicated researchers is globally recognised and frequently showcased at international security events and in leading publications. Their exceptional skills and unrivaled experience impact directly on our operations and are made accessible to our clients in various forms across our range of products and services”.

Charl van der Walt
Chief Security Strategy Officer
Orange Cyberdefense

Introduction:

Dispelling cybersecurity myths

Most of us are familiar with the story of Achilles, the heroic Greek warrior that survived many great battles during the Trojan War. Legend has it that Achilles gained the power of invincibility as a baby after his mother dipped him in the River Styx held only by his heel. This, ultimately, proved his undoing as it was in that same heel that he was fatally shot with a poisoned arrow. This popular Greek myth provides a useful analogy in the context of enterprise security, as even the greatest security systems and procedures can have their Achilles moment in light of seemingly remote vulnerabilities.

When evaluating and planning their organisation's security approach, IT leaders can easily fall into the trap of accepting common assumptions without questioning their validity. After all, cybersecurity is such a multi-faceted and fluid topic, even the most capable practitioners seek comfort in having some settled 'truths' on which to rest. In the hard reality of information security, however, assumptions can prove misconceived and can easily mislead us into doing the wrong things, or doing things wrong.

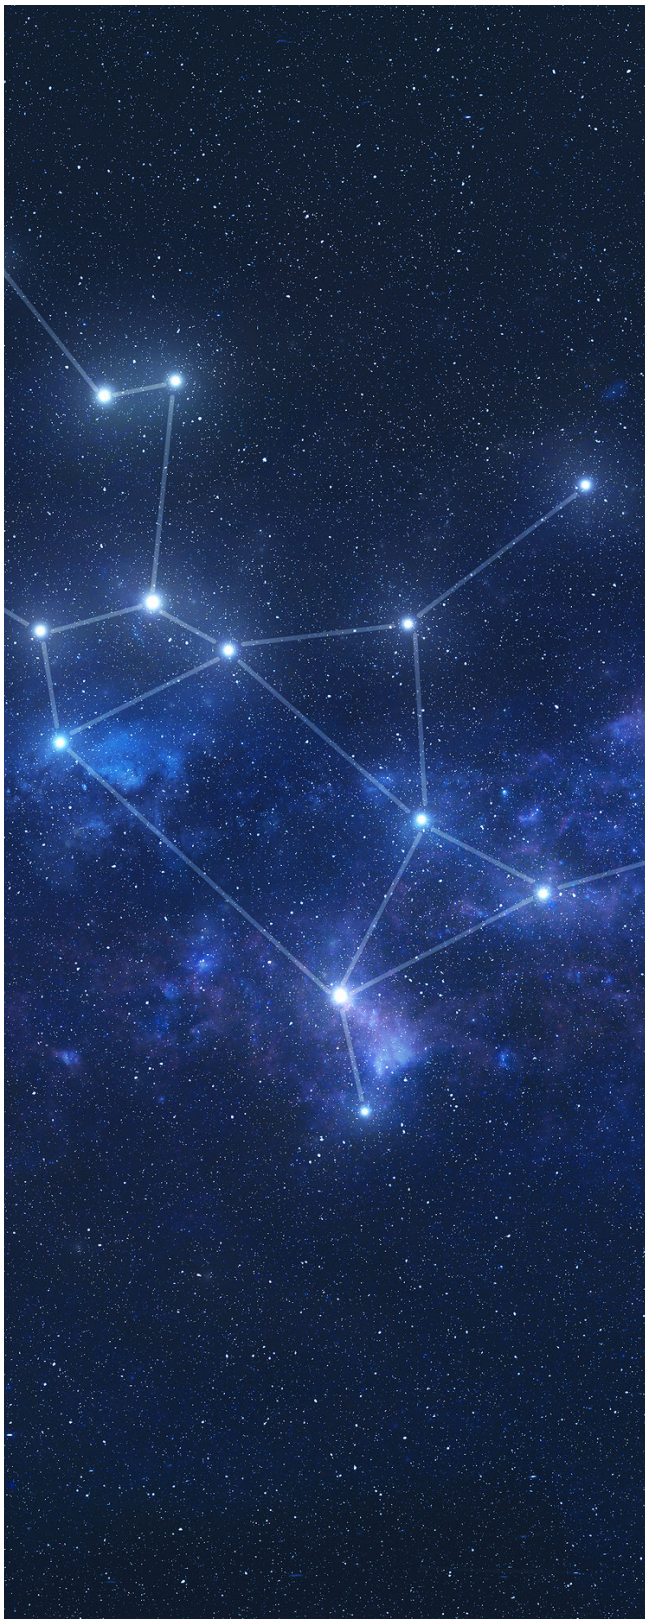

> **Assumptions can prove misconceived and can easily mislead us into doing the wrong things… or doing things wrong.**

It is this Achilles' heel that today's hackers employ to their advantage, but filtering facts from fallacies is no easy task for the average corporate defender. We set out to test widely-held assumptions that may be leading us to deploy our scarce resources in inefficient or even counter-productive ways. Using the volumes of data at our disposal as a managed security services provider, we started with a set of hypotheses and, after many long days and nights of testing, arrived at some surprising conclusions.

Over the course of this report, we'll share with you our methods, findings and learnings. We hope this will help you to look at your IT estate in a different way, anticipating the threats that others ignore.

Part 1:

The price of a patchwork approach to mobile security

When securing a corporate IT network it's all too easy to focus on the obvious items like laptops and servers. The smaller supercomputers that every employee uses just as frequently, if not more, during the course of their working day can often be overlooked. But with mobile devices being used for all manner of working practices and accumulating masses of corporate data, it's vital that they're kept secure.

Figure 1 (right): Example - Exploitable vulnerability on iOS. Patching has an important role to play in this. Failing to install important updates to mobile operating systems, be it iOS or Android, leaves devices open to security threats. The prevailing sentiment seems to be that iOS devices 'patch themselves' and are therefore a better choice in terms of security. Our research suggests that this assumption does hold…but not entirely.

Figure 2 (right): Example -
76.3% of Apple devices we tracked are impacted by a new vulnerability.

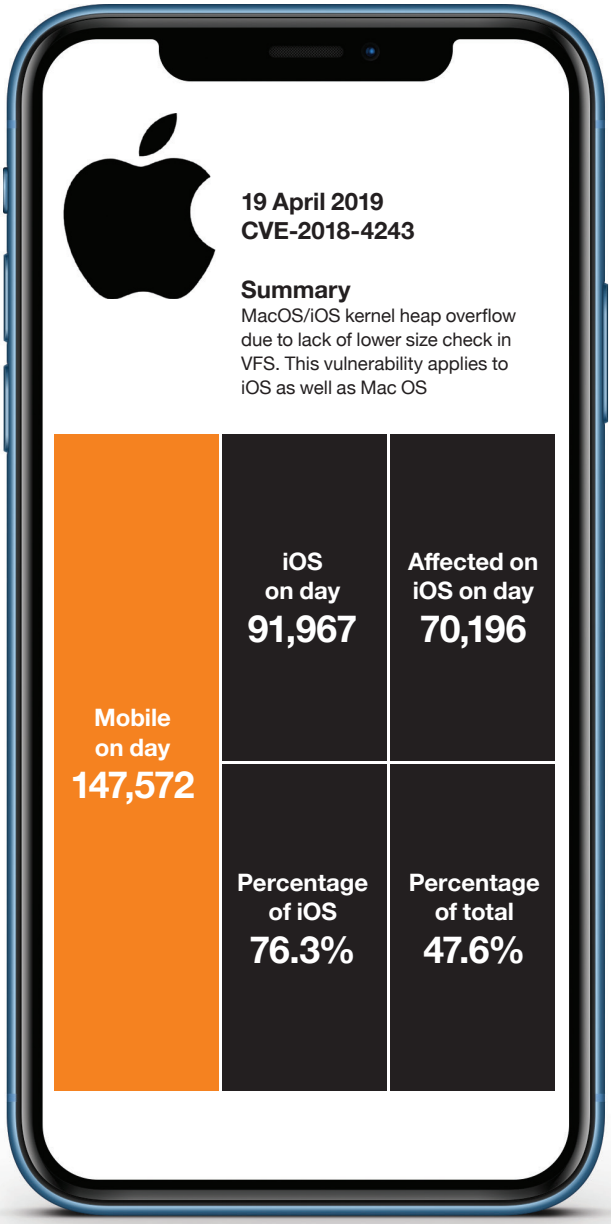

Using mobile device information from millions of unique web server logs, we analysed the rate and frequency at which new updates are applied to mobile devices.

We examined 85,000 iOS and 60,000 Android devices per day over 6 months. As we anticipated, the vast majority of iOS updates are applied promptly, limiting the potential for vulnerabilities to be exploited.

However, we also found that across both major OS families up to 20% of users avoid installing updates for some reason– almost in perpertuity - and therefore use older versions of the software, which increasingly become open to bugs and exploits.

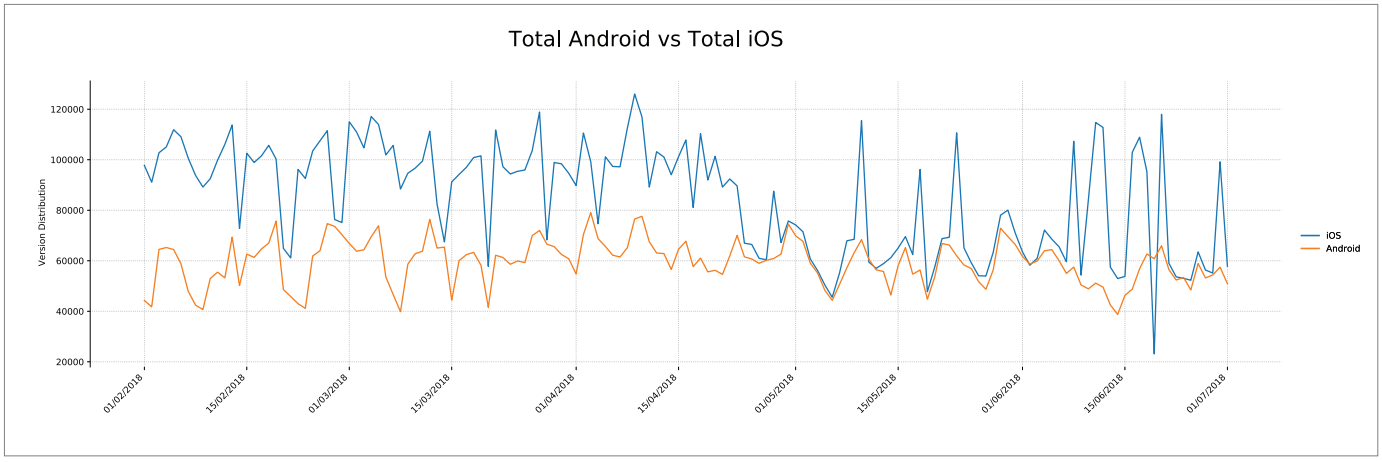

Figure 3 (above): Example -
Shows the distribution of iOS and android devices examined for this study.

Across both major OS models up to 10-20% of users avoid update patches almost in perpertuity.

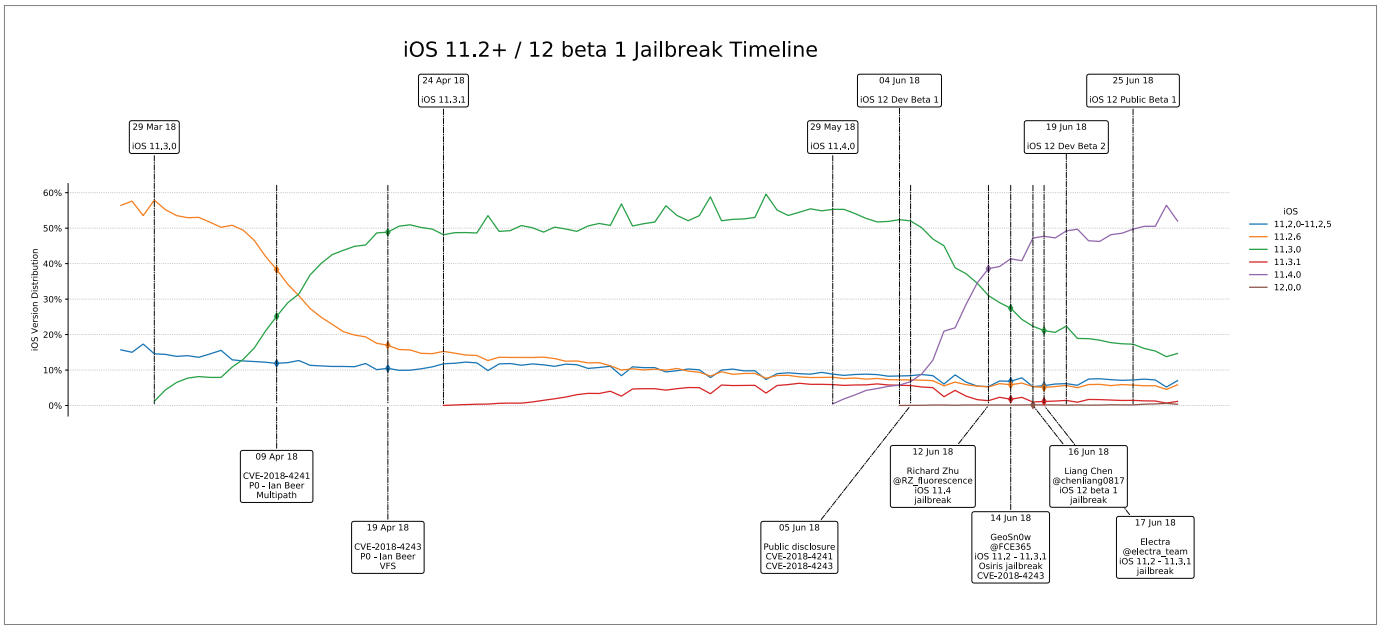

Figure 4 (above): Patching behaviour for iOS, relative to important vulnerability disclosures

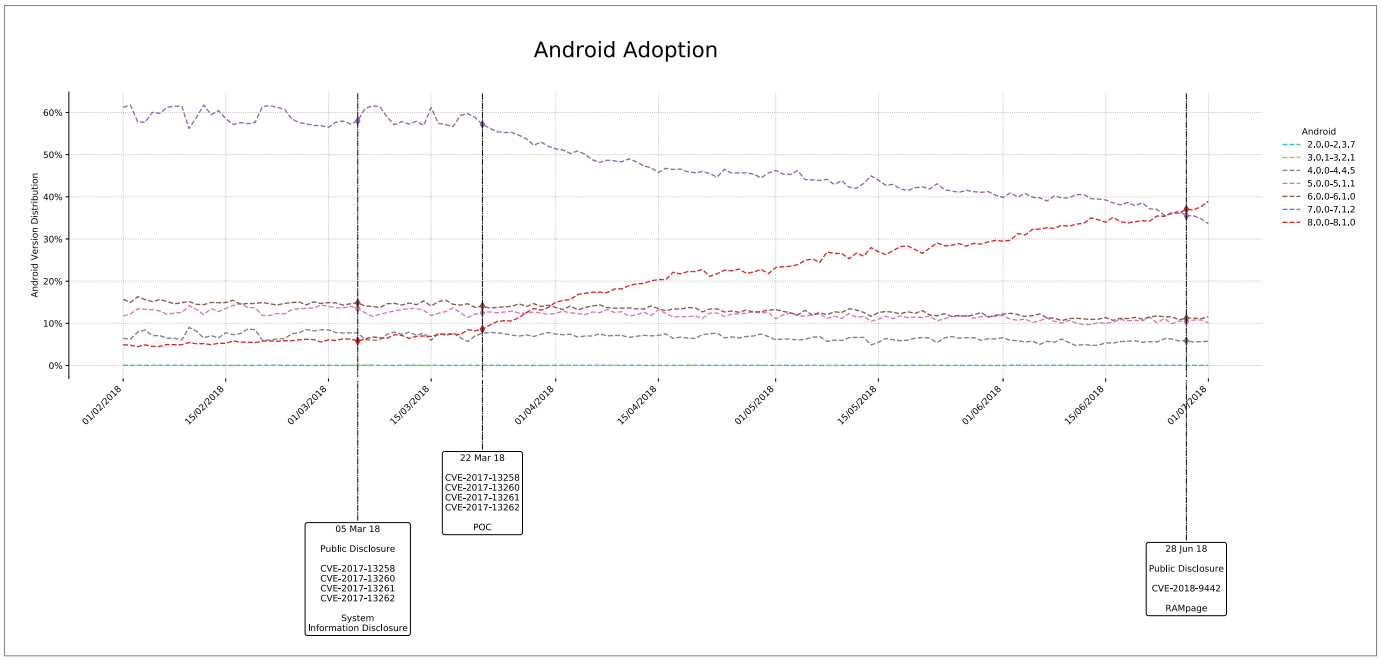

Figure 5 (above): Patching behaviour for Android devices, relating to important vulnerability discosures

The longer these devices remain unpatched, the greater the risk of exploitation. And, when we consider the vast number of mobile devices that are used within corporate networks, we can better appreciate the scale of the risk.

The data in the two charts above suggests that an iOS vulnerability has a much shorter life expectancy for the bulk of the iOS population than a similar vulnerability would have in the Android population: Apple's philosophy of strict central control and automated patches has a marked impact on the efficiency with which vulnerabilities are addressed.

Although the bulk of devices are patched more quickly in the Apple ecosystem, both ecosystems manifest a significant portion of device owners (between 10% and 20%) who it seems will never patch, regardless of how much the manufacturer tries to push them. We can see this as an example of the 'human factor' in which security conflicts with usability regardless of how hard the vendor tries to encourage the right security decision. Both operating system families demonstrate this characteristic.

For IT, this situation presents a diplomatic challenge: employees are typically more territorial when it comes to mobile devices than they are with other hardware. They feel it's up to them whether they install updates and may resist pressure – either from IT or from the manufacturer – to implement updates. IT however has a duty to the organisation to make sure that all software connected to the network is up-to-date and therefore needs to push on this policy, despite the resistance they receive.

> **IT has a duty to the organisation to make sure that all software connected to the network is up-to-date, despite any resistance they receive.**

Apart from the obvious benefit in investing in the more tightly-controlled Apple ecosystem for enterprise mobile IT, businesses need to consider the use of Mobile Device Management (MDM). MDM helps to ensure that anyone who's connecting their device to the business' network is using the latest software update, and has the appropriate policies applied, reinforced with regular checks to make sure this is being adhered to.

Key takeaways:
For security professionals

- **Mobile platforms are arguably less of a target for attacks than desktops, but are still a security risk and are a growing area of concern.**
- **Mobile devices have to be patched in the same way as desktops do, but often fall outside the traditional sphere of influence and control for corporate IT.**
- **The Apple iOS ecosystem is tightly controlled and centrally managed, which has a demonstrable impact on the speed and scale with which new patches are deployed and vulnerabilities mitigated.**
- **Android on the other hand is generally decentralised and poor at distributing patches to devices.**
- **We therefore recommend Apple iOS over Android as a choice of corporate mobile platform.**
- **Nevertheless both ecosystems contain a portion of devices that are never patched and therefore remain vulnerable.**
- **We recommend the use of MDM as a means of exerting control over devices from both platforms and thus mitigating this risk.**

Part 2:

The limits of threat intelligence

Threat intelligence technology is quickly gaining traction in cybersecurity. IT networks are generating so much data that the idea of being able to easily harness this data to predict, detect and address security threats is every IT security manager's dream. Businesses using threat intelligence systems are eager to have as much data at their disposal as possible.

This has bred a whole commoditised industry of shared data feeds – threat identifiers sold by vendors and service providers to enterprises as indicators of potential threats their network may face.

The logic behind this burgeoning market is quite simple: if an IP address has demonstrated suspicious behaviour on one network, that same address may well appear on another network attempting similar techniques. Shared feeds have therefore become a common source of data for many threat intelligence systems for their supposed predictive value. However, taking action against a suspected 'indicator' consumes valuable time and resources, no matter how slight. We therefore need to be sure that the IP address indicators contained in these thread feeds have the predictive value required to justify the time and energy required to respond to them.

We investigated whether using data from these feeds amounts to an effective and valuable way to identify threats. Over the course of one month, we observed and recorded 118,000 unique 'suspicious' events, involving 12,750 unique IP address indicators, with a view to determining the likelihood that any IP marked as suspicious would reappear as an IP indicator.

Figure 6 (below): Our Threat Intelligence research process illustrated that only 0,01% of IP indicators were eventually confirmed suspicious, and 49% of those were already identified by our honeypots.

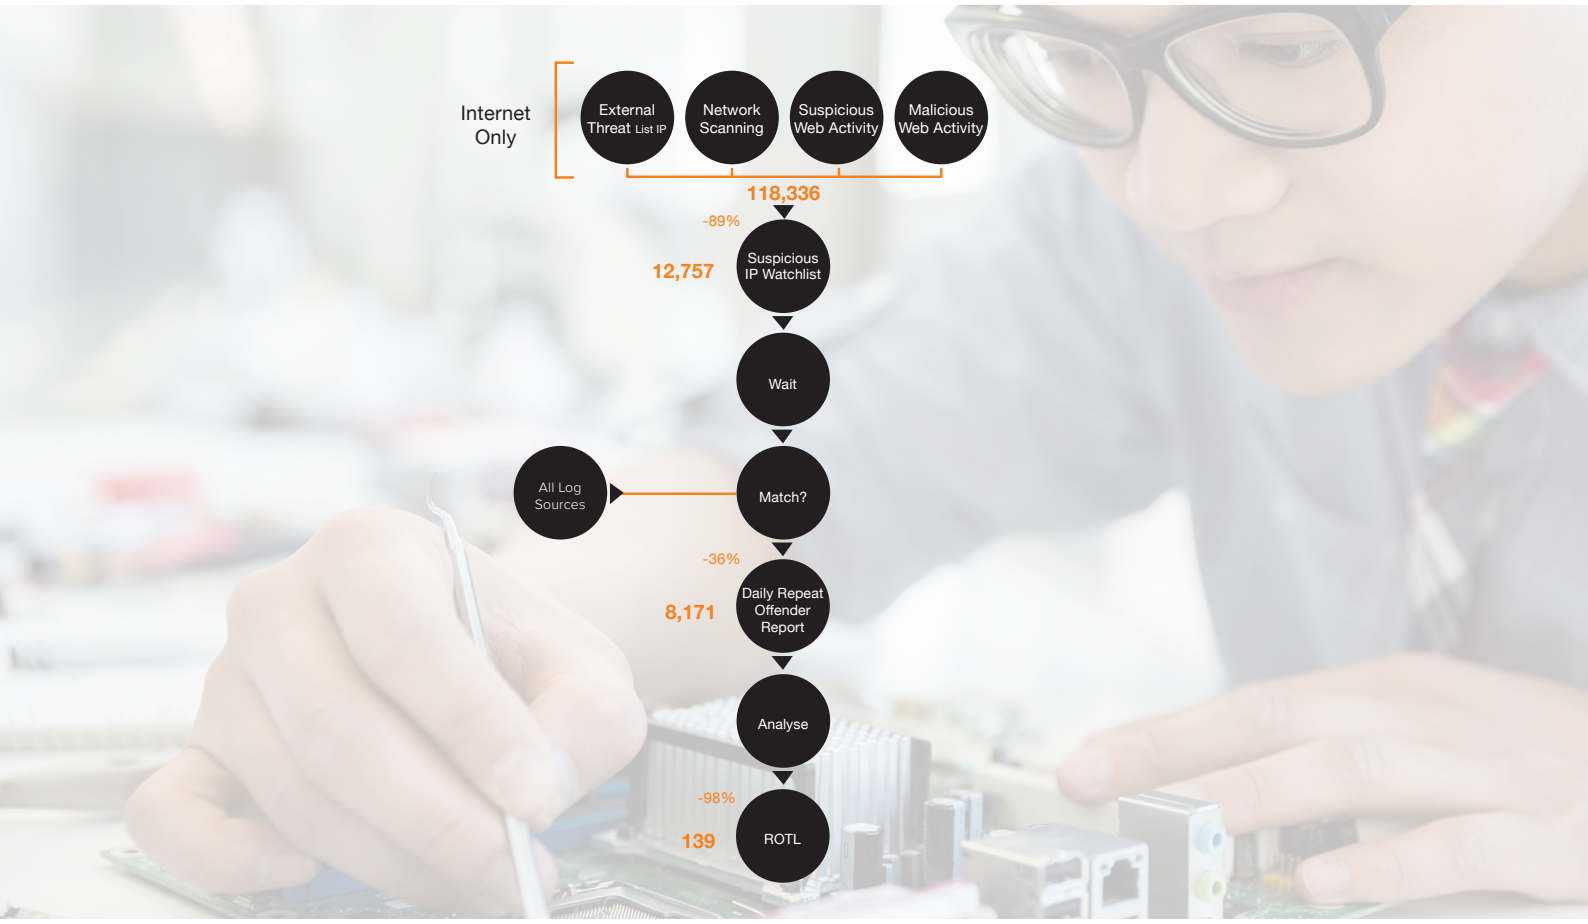

We found that the utility is – at best – limited. Of all the IP addresses tracked, only 0.01% of suspicious IP addresses were confirmed as malicious through manual investigation. The level of attrition required to find legitimate threats among the data assessed was therefore staggeringly large, indicating a very low probability of identifying a threat through this method.

Instead, we found more value in less fashionable approaches, specifically 'honeypotting'. This strategy involves luring potential attackers towards a network removed from any production infrastructure, using multiple sensors to track their activity and details. We found that honeypotting offers a better 'signal-to-noise' ratio, as any suspicious activity on the honeypot network was far less likely to have benign intent, and therefore more likely to present a genuine threat. Indeed, around one in seven malicious IP addresses detected by a honeypot in one country was also noted by a honeypot in a different country (nearly 15%).

Of all the IP addresses tracked, only 0.01% of suspicious IP addresses were confirmed as malicious through manual investigation.

Threat Intelligence is a sound proposition that has its place in a mature security operation. But, based on our research, we recommend that organisations think carefully about whether it's right for their security approach. No two IT environments are alike, so before you invest heavily in intelligence data you must realise the inherent limitations of that data and the level of manual work required to make it actionable.

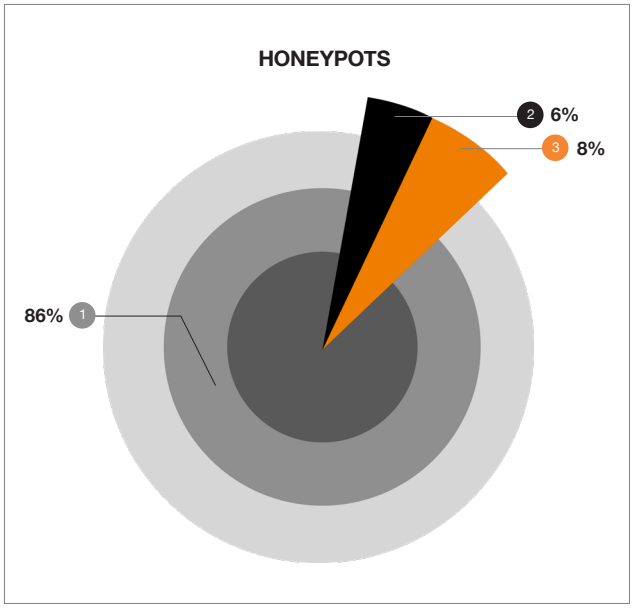

Figure 7 (above): 14% of all IP indicators detected by one honeypot were also detected by at least one other honeypot Ultimately, we recommend that before considering purchase of any "predictive" Threat Intelligence lists, you should think carefully about your threat intelligence approach, and whether they deliver value for money in terms of the cost of identifying genuine threats.

Before you invest heavily in intelligence data you must realise the inherent limitations of that data and the level of manual work required to make it actionable.

Key takeaways:
For security professionals

Threat intelligence promises to allow us to focus our limited resources on responding to specific, identified threats.

Generally Internet threat intelligence takes the form of a list of 'indicators' (IP addresses) that have been observed acting maliciously or suspiciously elsewhere. These observations are thought of as 'predictors' of future misbehaviour. Our research suggests, however, that previous misbehaviour on live environments is actually a very poor indicator for future misbehaviour.

We believe that independent honeypots are more efficient at identifying IP addresses that are likely to misbehave in the future and are worth taking action against.

Part 3:

Cracking the password code

In today's business environment, employees commonly access sensitive information from an array of devices and locations. In most cases, there's an inherent battle between convenience and security – the more security steps and provisions in place, the more frictions an employee may face when trying to access their data. It's little surprise therefore that employees tend to bend more towards convenience, consequently leaving security holes for attackers to exploit.

One such security stress-point is corporate passwords, often a source of jest in office environments due to the lack of care they receive. We analysed UK organisation DNS domains for vulnerability to exploits via weak passwords and Outlook Web Access (OWA) with a view to using our proprietary tool - Ruler - in order to gain full network access where possible.

The goal of this activity was to use public breach data dumps to assess how severe the threat of compromise is that businesses face from insufficient password protection.

Based on the extent of the weak password controls we've seen businesses employ over the years, we expected that the odds of success would be in Ruler's favour.

Figure 8 (below): Top 10 passwords from cracked corporate password hashes.

We were able to combine breached user passwords with OWA instances to project that 2,800 businesses were at risk of compromise using this one, simple technique.

But what we found surprised even us. Through our analysis, we were able to combine breached user passwords with OWA instances to project that 2,800 businesses were at risk of compromise from this one simple vector. But leaked data breaches are not the only way we can derive user passwords. Upon deeper analysis of the enterprise passwords themselves, we found that even enterprise passwords that comply with generic password complexity rules, could be easily predicted by analysing the patterns that users typically follow.

To determine this, we attempted to crack almost 600,000 enterprise password hashes to understand how those passwords are generally constructed. 85% of the password hashes we collected could be cracked, revealing a wealth of information about the passwords themselves and the difficulty level they present in deterring attackers.

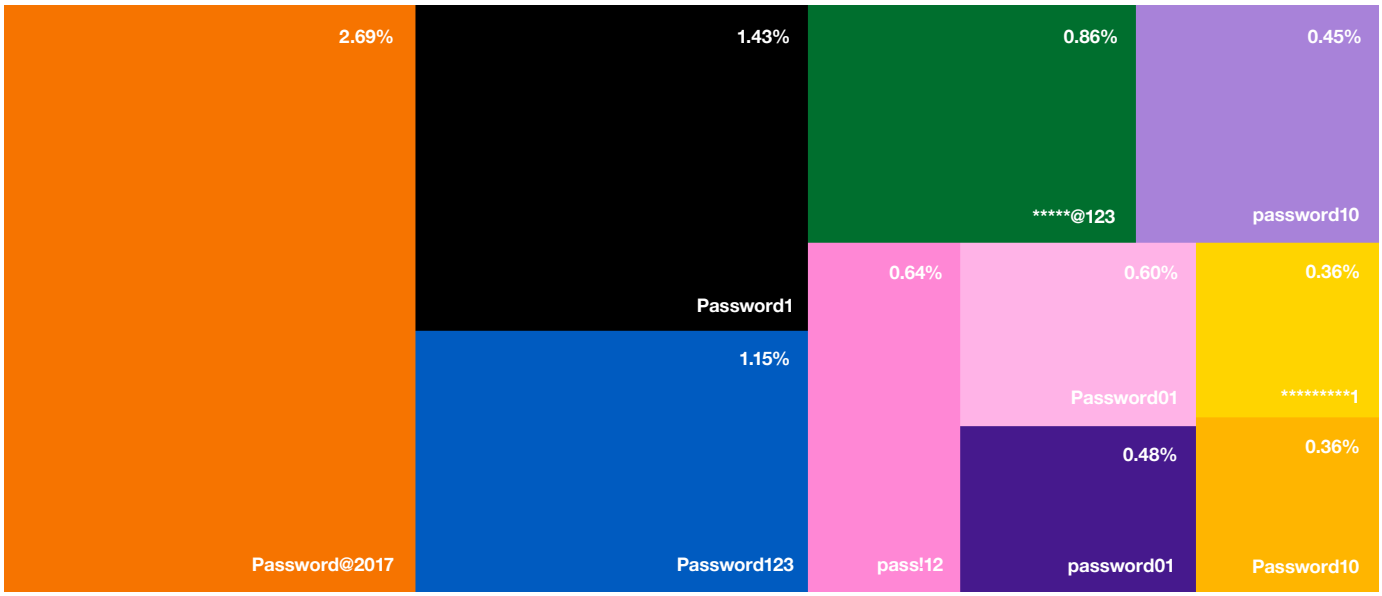

The time taken to crack the passwords was also noteworthy, ranging from two days for those of eight characters in length, to less than two minutes for those of six characters. These tests demonstrated that it's not just 'pass1234' that's prone to cracking, but any password that adheres to the common formats employed by most IT departments.

This has two major implications. Firstly, that once a hacker gets hold of an employee's password they will be able to try that same password over different systems in the expectation that the employee uses a single password for multiple sites. But also, once a hacker has worked out the password format your organisation uses, they are armed with a repeatable template they can use for follow-up attacks.

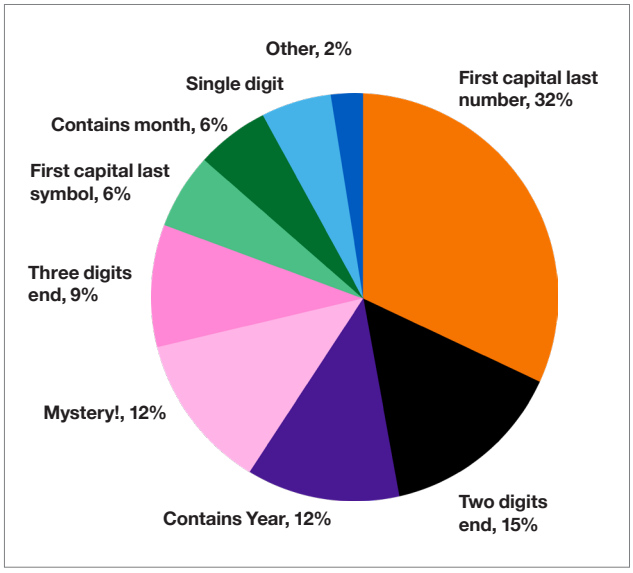

Figure 9 (above): Predictable password formats
The results of our experiments offer a wakeup call on the vulnerable nature of corporate password security. What's often treated as routine or even an inconvenience by many employees can be a matter of major organisational security when confronted with the sophisticated tools and tactics of today's cyber criminals.

To address this problem, we recommend avoiding standard word-and-numbers combinations for passwords and encourage employees to use pass phrases instead. These are more difficult for attackers to crack due to their unpredictable nature, and easier for employees to remember. Furthermore, passwords alone will not suffice to protect internet-facing environments, so consider the additional layer of mandatory two-factor authentication (2FA) tools such as Okta, Yubikey, Google Authenticator or SMS authentication for all employees.

What's often treated as an inconvenience can be a matter of major organisational security when confronted with the tools and tactics of cyber criminals.

Key takeaways:
For security professionals

Passwords are a still a key mechanism though which access to networks and other resources is controlled. Password complexity policies used by our customers appear to be having a positive impact in terms of the complexity and length of the passwords their users choose. However those same policies often encourage users to select passwords with predictable patterns.

Access to billions of passwords through leaked data breaches has given hackers the data they need to develop very sophisticated models for predicting user passwords leaving passwords even more vulnerable than before.

Using a single factor of authentication, like a password, is simply no longer an acceptable option when protecting systems and services on the Internet.

Part 4:

Putting hackers To the test

The core principles of operating a business are common across all sectors. The return on any product or service must be greater than the resource put in for that business to be sustainable. In service-based businesses, this is usually measured in terms of whether worker outputs (services rendered) produce a greater value per hour than inputs (wages and other expenses).

Your potential attackers are also resource-constrained; for them the return of breaching a company's security must be greater than the work that went into it.

This is all rudimentary stuff, yet it's useful to think of cyber threats in this way. Your potential attackers are running a business too; they have limited resources just like this, and for them the return on breaching a company's security must be greater than the work that went into it. This invites us to think about the goal of security as making it so economically unviable for a hacker to breach your security that it's not worth their time.

We had this logic in mind when we investigated the topic of vulnerability assessment. We undertook an experiment designed to show the time taken to breach a network pre- and post-vulnerability assessment. Using data from once-off penetration tests against specific targeted systems and continuous vulnerability scans against thousands of IPs both inside networks and on the Internet, we wanted to see what the findings could tell us about our customers' basic security hygiene.

The cost to an attacker becomes nearly seven timesmore expensive when penetration testing has been conducted and findings acted upon.

We found that when conducting penetration testing on internet systems, a 'serious' vulnerability is identified at least once in every 3.3 days worked, yet this increases to 21 days on retests, once a business has put our recommendations into action. In dollar terms, the cost to an attacker becomes nearly seven times more expensive when penetration testing has been conducted on a network and the findings put into effect. In other words, if you're able to find vulnerabilities before your attackers do and take corrective action, it'll take them far longer to uncover another one.

Figure 10 (below): Pentester effort in days required to identify a serious vulnerability

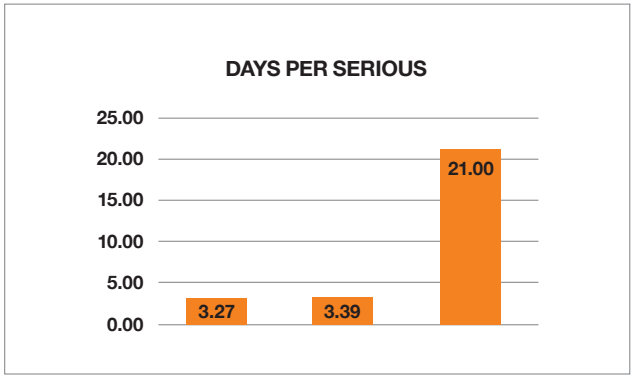

This shows the value of penetration testing in rendering cyber-attacks economically less feasible for hackers. When the returns don't add up, they're less likely to put time into attempting to breach your systems.

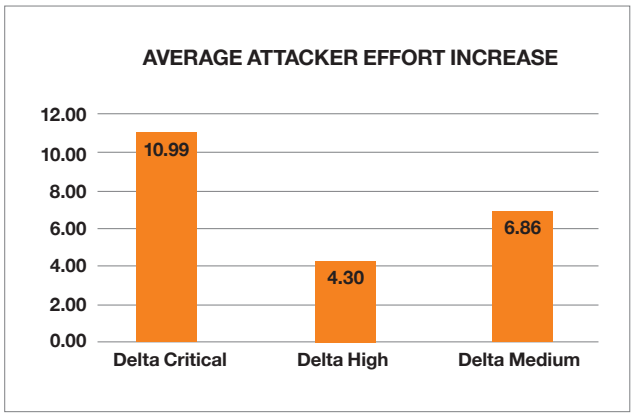

Figure 11 (above): The increase in attacker effort in days required after testing and remediation

Putting these findings into practice requires a change of mindset. The reason that many organisations choose not to pursue penetration testing is rarely down to resource and more often due to fear of discovering what vulnerabilities exist in their estate.

Hackers benefit from organisations that underestimate the risks they face.

Yet when it comes to cybersecurity, ignoring threats is the worst thing you can do. Hackers benefit from organisations that underestimate the risks they face, typically thinking either that they're not a plausible target or that the hackers lack the ingenuity to breach their defences.

This is where penetration testing becomes incredibly valuable. By identifying weak points you can take the appropriate steps to make them less appealing to hackers – more specifically, not worth their time to pursue.

Key takeaways:
For security professionals

It takes attackers time and energy to find and exploit vulnerabilities in your environment. The goal of security controls is to increase that effort for attackers, raising the cost and reducing the value of an attack for the hacker

By finding and fixing vulnerabilities before hackers do, we increase the time and effort they require to find exploitable problems.

Our research suggests that the incremental gains from testing and fixing are significant and demonstrable.

Conclusion:

An analytical approach To network security

Security is an industry dominated by Fear, Uncertainty and Doubt, perpetuated by strong personalities and vendors with specific, limited agendas. It may be that our real Achilles' heel is the willingness with which we accept and act upon generic and biased guidance at the expense of deliberate, thoughtful and analytical strategy. We need the courage to explore our own holy ground.

It may be that our real Achilles' heel is the willingness with which we accept and act upon generic and biased guidance.

As the four data experiments we've performed to date have all shown, small changes to how these areas of organisationalsecurity are treated can have a massive impact on securing your IT environment from cyber attack. Some practices fall firmly within IT's responsibilities – conducting penetration testing and creating honeypots for attackers, for example. Others, such as patching mobile devices and reinforcing password security, cannot be achieved by IT alone and need behavioural change at an organisational level.

Taking a comprehensive and analytical view of your IT estate helps to highlight the areas on which your organisation needs to focus, and what systems or behaviours need to be corrected in order to put the odds of remaining secure firmly in your favour. We hope that our hours of testing and analysis of data have helped you to spot some potential weak points in your own IT, and will encourage you to take preventative action before it's too late.

Why Orange Cyberdefense

Orange Cyberdefense is the expert cybersecurity business unit of the Orange Group. As Europe's go-to security provider, we strive to build a safer digital society.

We are a threat research and intelligence-driven security provider offering unparalleled access to current and emerging threats.

Orange Cyberdefense retains a 25+ year track record in information security, 250+ researchers and analysts 16 SOCs, 10 CyberSOCs and 4 CERTs distributed across the world and sales and services support in 160 countries. We are proud to say we can offer global protection with local expertise and support our customers throughout the entire threat lifecycle.

Orange Cyberdefense

orange™