Insight Space

cyber insights
programme

Issue 01

Welcome to the NCC Group Insight Space

Cyber security and business resilience go hand-in-hand – and in today's world, they are both more important than ever. A solid understanding of the common techniques employed by threat actors, and what you need to do to counter them, can enable your organisation to operate with confidence.

Whether you're in an executive or technical role, we'll seek to understand your unique challenges and provide you with the tools you need to build organisational resilience, in a language that you can understand. Here, you'll find insights from our experts to help you stay secure, including tips from our global CTO, Ollie Whitehouse, technical director Paul Vlissidis, and many more.

You'll also be able to access these tools in a range of formats, so that you can read through or listen on the go, on any device. Most importantly, we're always here to work with you to help increase your resilience – if you'd like to find out more about any of the content on offer here, don't hesitate to get in touch.

"Organisations in every sector are currently taking a step back and examining their own resilience. From a cyber security perspective, this could mean putting measures in place to reduce risk, or implementing a robust business continuity strategy.

"There's a lot to explore, but we're here to advise organisations and join them on their cyber security journey, no matter where they are currently. We'll work with our clients every step of the way to help them through this period of uncertainty, providing them with reassurance that their business is ultimately safer and more secure."

- Ian Thomas,
Managing Director at NCC Group

Contents

Technical Insights

Ollie Whitehouse

Chief Technical Officer

Liam Stevenson

Associate Director

Business Insights

Paul Vlissidis

Technical Director & Senior Adviser

Ade Clewlow

Senior Adviser

Tim Rawlins

Senior Adviser

Insight Space

cyber insights programme

nccgroup

Threat monitor

Threat Intelligence Report Q1 2020

A sample report from the Fox-IT Research Fusion and Intelligence Team

FOX IT
part of nccgroup

Report context

This report is a sample version of Fox-IT's quarterly insights into cyber intelligence and client-oriented threat analysis services, designed specifically to fit the needs of our managed detection and response (MDR) customer community.

This report outlines cyber threats to financial services and critical infrastructure, and supports decision-making processes and preparedness against the offensive capabilities of sophisticated threat actors and financially-motivated fraud operators.

Fox-IT, as part of the global NCC Group team, believes in contributing to a more secure society, and, by offering the expertise of their Threat Intelligence team, will help to tip the scales in favour of your security teams.

COVID-19

Working on the security and availability of IT systems is perhaps more important than ever. Much of the remaining day-to-day order is highly reliant on information technology, and the presence of well-functioning IT systems in health care is therefore of vital importance.

To this end, Fox-IT has prepared and delivered customized and actionable insights for health care providers. We encourage everyone to collaborate and contribute to this cause by volunteering and sharing suspicious activity for further investigation.

Content outline

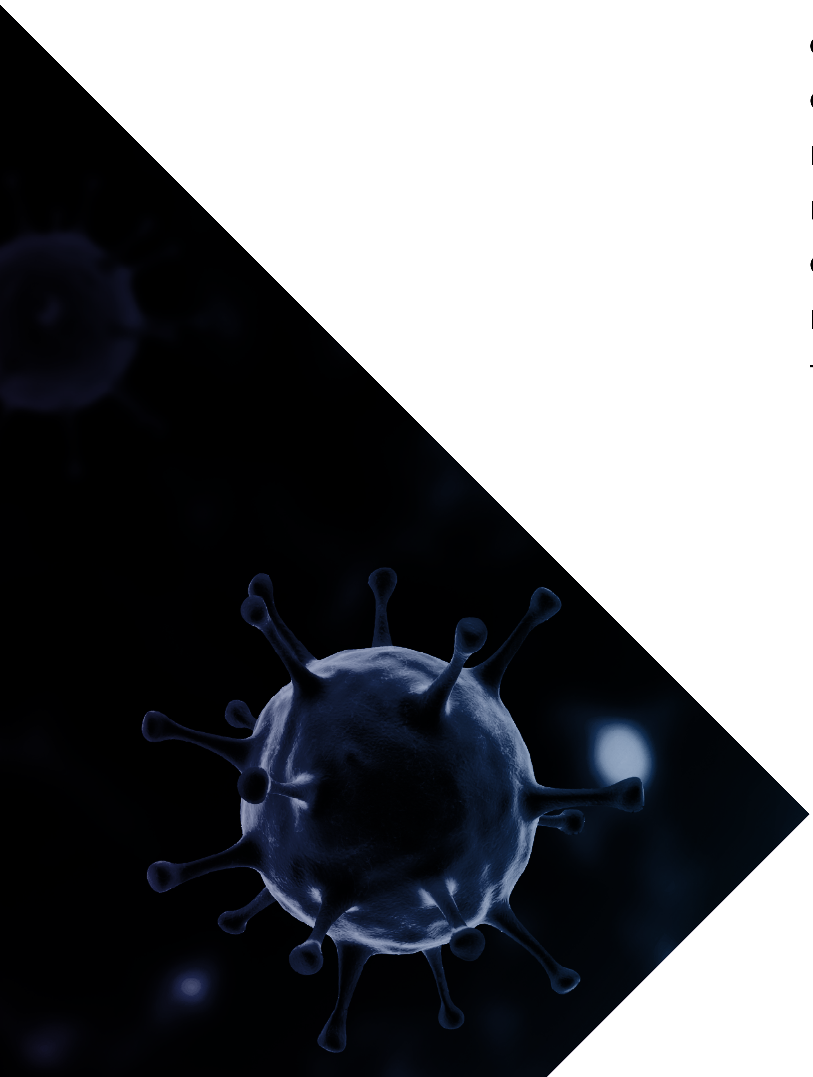

Critical events monitor

The timeline highlights the most strategically impactful incidents on the global threat landscape during this period.

All content from the below:

JANUARY

01/
Travelex forced to take digital services offline after Sodinokibi ransomware breach

14/
Microsoft patches crypto vulnerability reported by NSA

19/
Citrix releases fix for widely exploited CVE-2019-19781 in Citrix ADC and Gateway

24/
Ryuk info stealer targets government and military secrets

FEBRUARY

03/
Doppelpaymer leaks victim data on darknet if ransom unpaid

10/
Chinese Military Personnel charged with hacking Equifax in 2017 by US Justice Department

16/
Iranian Fox Kitten group exploits VPN flaws worldwide

20/
UK blames Russian GRU over 2019 Georgia defacement attacks

MARCH

05/
New PwndLocker ransomware fixes flaw allowing file recovery

09/
Compromise European power grid organization ENTSO-E restricted to office network

13/
Czech Republic hospital with Covid-19 testing laboratory hit forcing shutdown

26/
Navigator group delivers USB by post to install Griffon backdoor

COMMENTARY

The mitigation steps to reduce risk depends on the version of Citrix. Security teams at any organisations running Citrix ADC and Citrix Gateway Release 12.1 build 50.28, or those that saw a time-lapse between implementation of the mitigation steps after the release of the public exploit code on the 10th of January, should check for forensics artefacts to find indications that the vulnerability has been successfully exploited.

TAKEAWAYS

The US government has again chosen to bring criminal charges to counter actions of foreign governments in cyberspace. The unsealed indictment alleges that four officials hacked into credit reporting agency Equifax in 2017, stealing personal records such as driver's license numbers and Social Security numbers on about 145 Americans. They are members of China's armed forces, specifically PLA's 54th Research Institute.
The intrusion was initiated by exploiting a vulnerability in Apache Struts software. Equifax has been faulted for lax security measures and is set to pay $380.5 million to victims of the breach as part of a class-action lawsuit.

COMMENTARY

The European Network of Transmission System Operators for Electricity (ENTSO-E), a coordinating body for utilities delivering electricity, encountered malicious evidence of an intrusion into its office network. ENTSO-E stated the compromise did not spread to any operation electric transmission system impacting critical control systems. Reports indicate there was repeated, high-volume communication between a compromised mail server and threat actor infrastructure. Allegedly, the open-source Pupy RAT was used, further obfuscating the origin of the intrusion.

This section outlines developments on financially-motivated attackers. The inside story on the most relevant organizational, architectural and geographical factors are detailed to enable a timely response.

Figure 4: Total target list distribution by malware family 2019 Q4 – 2020 Q1

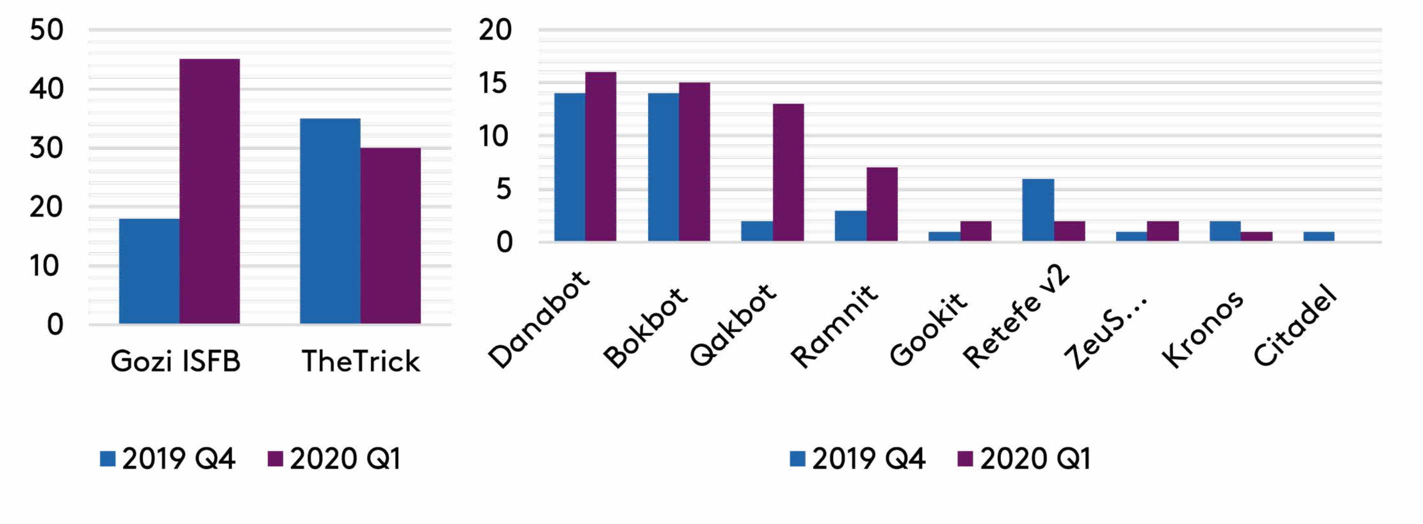

GOZI ISFB FLYING UNDER THE RADAR

Gozi ISFB is an established name within the banking malware scene and, in terms of config distribution, is the most active banking malware family during 2020 Q1. Four different Gozi ISFB offshoots are currently widely distributed by a number of different threat actors and causing financial losses by either performing transactional fraud themselves or by facilitating other types of malicious activity.

TRW PLUS DATA BREACH

Organized crime groups carry out targeted, multi-stage ransomware attacks to lock organizations out of their critical data and information systems. Recently, invented by the Maze ransomware gang, an increasing number of TRW actors are leaking sensitive data of non-paying victims to damage organizational reputations and inflict regulatory fines as an added pressure point to force a victim to pay the ransom. Dedicated leak websites are put into production and the stolen data is subsequently released in multiple parts, leaving a line of communication open to the victim to allow them to retract the leaks upon payment.

NAVIGATOR BADUSB ATTACK

FIN7 – tracked by Fox-IT as Navigator – is a threat group involved in high volume credit card theft and notorious for its use of crafty social engineering techniques to obtain a foothold within target networks. Recently, and fitting the group's modus operandi, Navigator distributed a malicious USB thumb drive, including a gift card and letter to victims in order to deliver the Navigator GRIFFON backdoor. The USB exploit is known as 'BadUSB' and acts as an USB keyboard to inject malicious commands, more specifically to retrieve two pieces of PowerShell code.

Financial malware trends

The financial sector and online payment systems specifically are a continuous target of fraud attacks by financially motivated hacking. The malware tailored to exploit the online banking processes known as banking trojans are detailed over time to capture trends and patterns.

TROJAN ACTIVITY

Bokbot is responsible for the majority of malicious activity, and is continuing its proven approach of targeting North American financials while updating its functionality at the same time. Although TheTrick is mostly moving away from wire fraud, the group is still responsible for a fair amount of global activity. Growth is again observed for both Danabot and Gozi ISFB, with the latter initiating new worldwide campaigns.

Figure 1: Threat size of financially motivated threat actors 2020 Q1

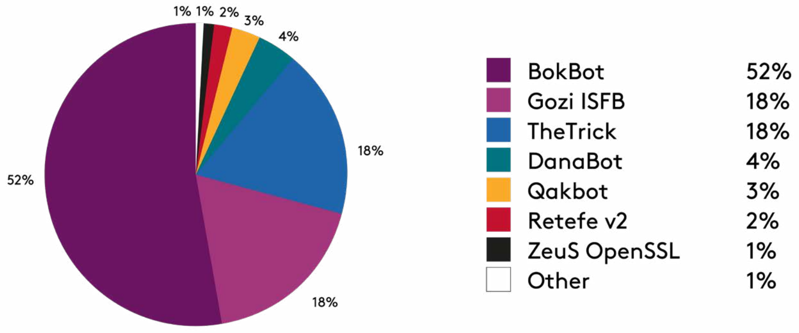

ATTACK TYPES DEPLOYED

When it comes to banking trojans, webfake redirect attacks are often deployed as the weapon of choice in order to commit fraud, in which victims are redirected to a phishing page representing the online banking page. For the client-side webinject, bank interfaces are adapted with an overlay to illicitly access credentials and mislead users during transfers. Overall, server-side inject attacks are decreasing and are solely deployed by TheTrick group to selective targets.

Figure 2: Attack types leveraged by financial malware 2020 Q1

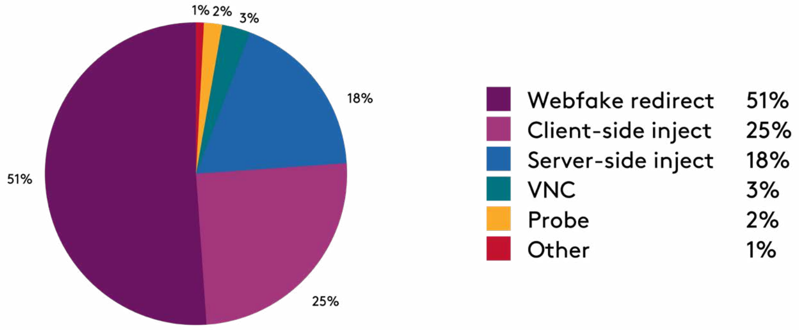

TARGETING BY COUNTRY – TOP 10

Attacks are primarily concentrated in Europe and North America, again representing a continuation of the previous quarter. Both Poland and Italy are increasingly scoped for exploitation, the US however remains the overall top targeted country. Note that the Top 10 countries amount to more than half of all targeting: for 2020 Q1 up to 63% and for 2019 Q4 as much as 69%.

Figure 3: Top 10 targeted countries by financial malware 2020 Q1

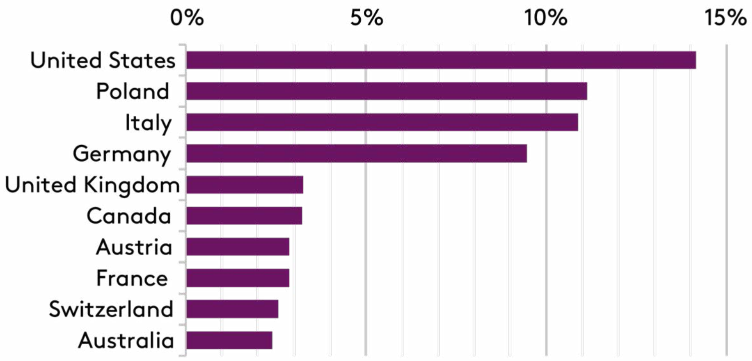

Credit card recovery

During the first quarter of 2020, Fox-IT recovered over 700,000 compromised cards originating from the source codenamed “Octopus”.

Octopus was however subject of a Russian law enforcement operation in March this year, which effectively cut off our visibility. New sources for credit card data are however uncovered and subsequently prepped for distribution.

700,000

Compromised cards recovered

On the 20th of March the Russian FSB initiated a takedown operation on a credit card theft ring Fox-IT refered to as “Octopus”. Octopus is essentially a platform that hosts a large number of carding shops selling credit card track data. We have written extensively about Octopus in our Threat Monitor Annual Report 2019. According to our insights the entire infrastructure has been taken down and thus we are no longer able to provide cards from this source. Further information is detailed in the section below.

Figure 5: Count of recovered cards by source

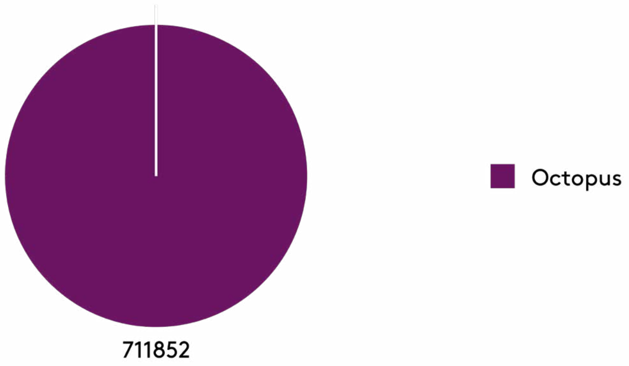

Law & enforcement

Here, primary law enforcement and legislative actions impacting cyberspace are summed up. Most significant are operations that have an immediate impact on active threat actors and the infrastructure, malware and monetization channels they leverage.

RUSSIA BLOCKS SECURE EMAIL

The Russian government is increasingly blocking access to secure email services such as ProtonMail, Tutanota Mail and Startmail. Shifting more and more towards China's model of internet censorship, Russia's state communications watchdog "Roskomnadzor" is through legislation now allowed to control information and block (messaging) applications. This signals new important steps towards an indepen-dent Russian internet not reliant on routing through other countries.

OCTOPUS TAKEDOWN OPERATION

On March 20th, 2020 the Russian Federal Security Service (FSB) arrested 30 individuals that were involved in a high volume credit card data theft and distribution network Fox-IT referred to as "Octopus". Russian law enforcement action on cybercriminals residing within its own borders is exceptionally rare, however, according to our insights (confirmed by the FSB statement) the group sold card data originating from both Russian and foreign financials thus revealing a potential motive.

CORONA CRIMINALS PRIORITIZED

The current COVID-19 pandemic provides cybercriminals with ample opportunity to exploit these unprecedented times for various nefarious means. Several governments are considering their options to combat this malicious activity, e.g. with the U.S. prioritizing the prosecution of cybercriminals exploiting fears about the corona virus. The Australian Signals Directorate stated that it had "mobilized its offensive cyber capabilities to disrupt foreign cyber criminals responsible for a spate of malicious activities during COVID-19", advertising its hacking potential.

Understanding the objectives and methods of threat actors is key to building a stronger security strategy.

Our experts are always here to help your organisation become more resilient - if you have any questions about any of the content in this paper, or would like to access the full version of this report, please contact your account manager.

Terms of use

Fox-IT tracks global cybercrime activity. We base our intelligence on tracking threat actors, darkweb research, forensic investigations, internationally deployed sensors and fraud monitoring services.

Going beyond botnet & malware information, we provide a global picture of trends, geographical activity, actors, their motivations and their evolving business models. We provide links to campaigns, tactics, procedures and individual IoCs to feed network security components. Customers become part of a global community, with live threat tracking, collaboration, and the largest criminal threat database, with over a decade of experience.

The data and charts contained within this report represents Fox-IT its own dataset collected within its malware lab. The data from this lab should be considered a sample including factors potentially skewing the analysis: our lab does not analyze every malware sample on the threat landscape, merely those assessed to represent a cross-section from a variety of sources. Our sources may be skewed towards certain types, families or regions which can introduce further bias. The report documents the dataset over a fixed period of time allowing for comparative analysis, whereas when referring to previous datasets a discrepancy with previous reports may seemingly occur due to inclusion of the updated dataset that may contain recent data impacting the statistical outcome.

Furthermore, the lists of data we use to identify targets for attacks can also be biased because they will naturally contain more data pertaining to Fox-IT customers than organizations not part of the MDR community. Although we augment customer supplied data (such as URLs for online banking and BINs) with autonomously collected data, the customer supplied data will always be more detailed and extensive. In short, these charts provide indications, and should be incorporated by interested parties as such. Customers are advised to incorporate and correlate multiple feeds with internal network telemetry.

Insight Space

cyber insights programme

Technical Deep-Dive

Thematic for Success in Real-World Offensive Cyber Operations

How to make threat actors work harder and fail more often

Executive Summary

The purpose of this paper is to assist organisations in prioritising their security activities, to thwart attack techniques successfully utilised during Red Team engagements and other offensive operations by real-world threat actors.

Organisations that effectively address the issues discussed in this paper will dramatically limit the options available and inhibit effectiveness of easy to execute yet highly impactful techniques, forcing an attacker into taking alternate riskier approaches and increase the potential for detection.

Offensive cyber operations, or Red Teaming as it is known when delivered commercially, provides value to organisations by assessing the defensive operational efficacy against real-world end-to-end attack scenarios and threat actors. The value is however generally only obtained where a mature cyber resilience program is in place.

For most organisations, Red Teaming is not the correct mechanism to employ in the first instance. Instead, we would recommend less combative and more collaborative approaches between cyber offense and defence functions. These approaches, often referred to as Purple Teaming (the combination of Red – Offensive and Blue – Defence), lead to materially better resilience outcomes both technically and culturally.

This paper is based on the techniques NCC Group employs most commonly in successfully breaching the largest and more sophisticated organisations across the globe in the wide range of Red Team engagements we deliver annually. Each stage of an attack is described with a reference to the respective Mitre ATT&CK technique for further reference.

In reality, the attacker wins most of the time because of poor operational hygiene inside and outside of the organisation in relation to digital assets. This poor hygiene provides the window for initial compromise coupled regularly with an inability to detect, contain or effectively respond to a breach.

Reconnaissance

The reconnaissance phase is where a threat actor surveys the target organisation attempting to gain information about its people, processes, systems and partners. Information is everywhere (ATT&CK TA0015, TA0016, T1526)

All organisations will leak information and once it has gone attempts to remove it will have limited success due to the nature of the Internet. Awareness of the information an attacker has is therefore crucial meaning organisations should regularly perform reconnaissance against themselves.

Developer repositories such as GitHub, BitBucket, GitLab, SourceForge, etc. are rich sources of information perform searches for organisation name, internal domains and words specific to your organisation or critical systems. Where sensitive information is obtained ensure actions are taken to limit its effectiveness, remove accounts, change hostnames or ultimately use newly written code that has not been exposed.

Staff members within organisations commonly use Social Media which is outside of the control of the organisation. Information such as key technologies used for critical systems and the specific staff members associated with it make them targets for spear phishing attacks. Organisations that perform their own reconnaissance in these areas can more effectively educate the specific staff members to alert them that they and their families are at increased risk of being targeted. Organisations should consider additional defensive monitoring for digital assets of key staff.

Organisations should understand where they themselves reveal information, job advertisements detailing key technologies for critical systems, published documents containing meta data of internal resources and technical presentations given to third parties or customers can all be harvested to better target organisations. Limiting job and presentation information to be generic and having a process to sanitise all documents regardless of whether they are intended for the Internet can stop the usefulness of them.

The adoption of cloud systems (Office 365, G Suite etc) can often be discovered by reviewing a organisations publicly facing DNS, providing an attacker with both an idea of the defences they will need to overcome and a list of portals/endpoints with which to attempt user/credential enumeration. Organisations should firstly understand the number and location of all authentication portals including staff, customer and third parties then ensure they know what is normal in regards to when they are used, where they are used from and the numbers of success and failure events.

Defending against information leakage

Defending against disclosure is difficult and often impractical. In order to recruit the right people into the organisation, you need to specify the skills and experiences they are expected to have. Similarly, staff will have a desire to demonstrate their skills by gaining professional certifications and advertising these to their peers and potential employers. The best defence is educating staff on the importance of being generic in social media and being mindful of receiving very specific requests for technical help from unsolicited sources. These requests should be reported to the security team where appropriate.

The adoption of cloud systems is similarly hard to disguise. No real practical advice exists to effectively mitigate this disclosure, only the correct hardening and monitoring of the services themselves can provide assurance against abuse.

With regards to source code and sensitive information contained within, several automated tools can be used as part of a CI/CD pipeline exist to flag hardcoded values. One opensource tool is the excellent "Whispers" by Skyscanner https://github.com/Skyscanner/whispers

Overall, whilst "security through obscurity" is not true security, it is undeniable that limiting the amount of information readily available to an attacker forces them to enumerate and perform discovery activities more, increasing the chances for detection and limiting their ability to provide specific spear-phishing payloads.

In Phase

The in phase is where a threat actor attempts to gain their initial foothold within an organisations system, be they in the cloud or within private networks.

EXPLOITATION OF VULNERABILITIES (ATT&CK T1190)

Most organisations mature enough to commission Red Team engagements have a timely patching process for Internet facing services. However, there are occasions where due to a lack of vendor patch, exploitation is possible, though this is uncommon. Internally (once access has been obtained) there are typically far more opportunities as the patching cycle is often slower and outdated unsupported operating systems are more commonly found. This is an area where organisations are taking appropriate actions externally but internally removal of any operating system that is unsupported should be prioritised along with any reduction in time to patch known vulnerabilities.

Defending against known vulnerabilities

Whilst the obvious answer is "patch your systems in a timely fashion", this is a simplistic view that does not reflect the complexity of most enterprises. The reliance on managed service providers (MSPs), competing business priorities (uptime vs anything else) and a lack of internal asset management can all contribute to difficultly in patching systems both regularly and out of band.

Ensuring the organisation has an accurate database of assets, their locations, who is responsible, and the software versions installed is the first step in being able to identify systems affected by new vulnerabilities as they are released. Ensuring that the relevant people in the organisation are subscribed to vulnerability/vendor security mailing lists and have a process for raising an urgent change is also key to ensuring the timely patching of new vulnerabilities.

MSPs should be asked to clearly define their patching process for assets they are responsible for (both regularly and out of band) as well as the service level agreement (SLA) they will do this in.

Where a patch or mitigation does not exist, additional protective monitoring and firewalling should be placed around assets at risk, ensuring that only the services and network routes required for operation remain open. This is particularly true for unsupported operating systems that may still be in use for business continuity reasons, but it is important to stress that these mitigations only serve to lower the risk and should be applied alongside a comprehensive plan to migrate to a supported platform as soon as is reasonably possible.

Adopting an "assume breach" mentality to internal patching, assuming that an attacker will be able to breach the perimeter and have opportunity to exploit these vulnerabilities inside your internal network should allow the organisation to prioritise internal patching and not fall foul of the common reliance on the perimeter.

Detecting exploitation of known vulnerabilities

Ensure that all hosts are protected with updated Anti-Virus products and Endpoint Detection & Response (EDR) tooling that sends its logs to a SIEM for correlation and resolution. Ensure that any network-based IDS/IPS has coverage of all network traffic and is running regularly updated signature definitions.

EXTERNAL AUTHENTICATION EXPLOITATION (ATT&CK T1078)

Organisations should prioritise removing any single factor authentication exposed to the Internet.

Even where there is two factor authentication in use, organisations should ensure they have effective monitoring of endpoints as in many cases a valid set of credentials can be identified before the second factor is required. Understand what is normal such as number of failed authentication attempts in certain periods. Where increased activity is identified investigate the source from which it came, is it likely that multiple users are at the same source? Is it normal for those users to log in at the attempted times? Are failed account names seen repeatedly (could be relatively slow such as once per day/ week) which would indicate repeated attempts using a single password across a list of gathered usernames. Where suspicious activity is identified then organisations need to identify which (if any) of the successful attempts were the attacker and not the valid user.

Defending against external authentication exploitation:

Where possible, ensure that all external methods of authentication require a second factor and do not give a positive indication until the whole process is completed.

If using Office 365, ensure that legacy authentication mechanisms are disabled https://docs.microsoft.com/en-us/azure/active-directory/conditional-access/block-legacy-authentication

Ensure all authentication attempts are centrally logged to allow for event correlation and anomaly detection.

PHISHING\VISHING (ATT&CK T1192, T1193, T1194, TA0003)

A very common initial access technique is conducting phishing attacks against users identified during the reconnaissance phase. These attacks generally consist of attempts to compromise enterprise credentials or sending a malicious file that executes arbitrary code when the user opens it. Starting a conversation with a simple email without malicious content to avoid detection and allay recipient suspicion is also useful at times.

Most organisations do avail themselves of some form of filtering or content checking so attacks utilising domain fronting and keyed payloads are necessary. Organisations should ensure they understand what Internet resources their users communicate with and wherever possible limit it to known good ones. Where attachments are assessed for malicious content then any sandbox would need to simulate a domain joined host with standard Internet connectivity to try to activate the keying to reveal the actual payload.

Vishing in addition to, or to supplement spear phishing attacks are difficult to detect or prevent. These can be used to troubleshoot phishing campaigns or where no attachments were used on purpose to elicit a response from users. A process to verify calls are from genuine technical or other colleagues such as HR should be implemented for staff members to identify malicious calls.

Depending on the method of entry, it may be desirable for the attacker to install a persistence mechanism to ensure continued access to the target environment. This is particularly pressing when initial access was via an email attachment or some other "one-shot" mechanism where a system restart or log-off would cause the payload to exit. Common methods of persistence include scheduled tasks, start-up tasks, DLL Hijacking, COM Hijacking and service installation. The risks associated with writing payloads to disk are often outweighed by the necessity of maintaining access to achieve objectives. The use of Endpoint Detection and Response is more prevalent within organisations and they do identify these types of activities with varying levels of success depending on the product and the technique used. Organisations should use them if they do not already and regularly review the products themselves for the protection it affords against changing techniques.

Defending against Phishing:

NCC Group has previously written about Phishing Mitigations in 2016 https://www.nccgroup.com/uk/about-us/newsroom-and-events/blogs/2016/january/phishing-mitigations-configuring-microsoft-exchange-to-clearly-identify-external-emails/

The advice in the blog post is still true today, the importance of helping users identify emails from external sources, regardless of how convincing the domain, sender information or content is, is still of the utmost importance, alongside a range of technical controls to prevent spoofing.

Defending against code execution

Attackers looking to gain access to the environment via sending malicious files are relying on the ability of their victim to execute arbitrary code on their endpoint. This can be via variants of office documents with Macros or other active content, which may call out to other executables on the host to facilitate code execution. Several technologies and strategies exist to inhibit the ability for attackers to execute arbitrary code:

AppLocker (ATT&CK M1038)

AppLocker is a technology that allows system administrators to define allow/block lists for which applications and files can be executed on a host. These rules can be applied dynamically based on user/group membership, the host being executed on, the path of the file, its cryptographic signature or its file hash. This provides a flexible way of defining what user can execute what code, on what host, from where.

Reference AppLocker policies exist that may be of use to any systems administrators looking to implement it. Two such examples are https://github.com/microsoft/AaronLocker and https://github.com/api0cradle/UltimateAppLockerByPassList/tree/master/AppLocker-BlockPolicies. NCC Group has not tested these. Be sure to test in your environment before enforcing any AppLocker policies.

Modifying default handlers

By default, .hta files are executed by the mshta application. mshta can access several scripting engines, allowing code execution (ATT&CK T1170). These files can be heavily obfuscated to prevent anti-virus products from analysing them (the NCC Group Demiguise tool is one such example: https://github.com/nccgroup/demiguise).

By changing the default file hander for .hta files to something innocuous such as notepad, system administrators can prevent .hta files from automatically being executed by the mshta engine.

Disabling Office Macros

Whilst many multi-billion pound business still run on the back of Office Macros, there is no denying the security risk they pose. With access to several scripting engines, Office Macros can be used to write complex payloads in order to further the attacker's objectives (ATT&CK T1204).

Several approaches to limiting the ability for macros to run exist:

Disable all

Whilst initially this sounds extreme, it should be remembered that not all users have a need to run macros. Using Group Policy, it is possible to disable macros for all users except those with a business need for their execution. With additional education for users who are allowed to execute macros, their widespread deactivation would require an attacker to find a user who had the appropriate rights to execute the file.

Only run signed macros

It is possible to cryptographically sign Office macros to ensure they are written and trusted by an organisation. Whilst there are undeniable maintenance overheads to this, it would an organisation to ensure the only macros being executed have been audited and signed by them.

Trusted Locations

An alternative to the maintenance heavy approach of digitally signing all macros may be to only allow their execution from trusted locations. A read-only share containing the documents containing the macros can be added as a "Trusted Location" via GPO. This would allow the execution of these macros only from this location.

Do not run macros from the internet (ATT&CK M1021)

Using the Attachment Execution Service (AES), Office can recognise that files have been downloaded from the internet, as opposed to being created on the local PC or copied from an internal source. Disabling execution of Macros from the internet would prevent attachments that have been emailed from executing, preventing attackers from simply emailing in documents, but still allowing ease of use for end users.

Attack Surface Reduction

The Attack Surface Reduction (ASR) rules https://docs.microsoft.com/en-us/windows/security/threat-protection/microsoft-defender-atp/attack-surface-reduction#attack-surface-reduction-rules are a list of additional settings within Windows Defender that can reduce the avenues available to attackers to gain code execution from common applications. Options available include:

- Blocking a variety of scripting engines from downloaded files
- Preventing office processes from injecting into other processes
- Preventing Office and Adobe Reader from spawning child processes
- Preventing the execution of potentially obfuscated scripts.

By implementing these additional protections, most methods attackers use to gain execution from emailed files will be blocked and logged

The through phase is where a threat actor moves laterally through an environment gaining access to additional hosts or functionality moving towards their ultimate objective.

USING INTERNAL INFORMATION REPOSITORIES (ATT&CK T1213, T1039, T1081)

Once an initial foothold has been achieved corporate communication platforms, information repositories and network shares are all searched for information. These platforms can be searched as the compromised user looking for information on the objective, credentials in scripts/user guides and password reset processes and systems. Where credentials are identified attempts to use them to move laterally to other hosts is relatively simple and can appear as normal behaviour.

Defending and detecting against the abuse of internal information usage

Organisations should firstly perform regular internal audits of information repositories to identify who has access to them and remove those that do not require it. Following that an assessment of the information they contain and whether it is appropriate, where credentials are identified then they should be considered compromised and changed. Logging of repository search strings should be captured, and alerts created for specific keywords such as password. Consider creating canary content that creates an alert when accessed or used which could include a specific domain account.

MAINTAINING AND ELEVATING ACCESS THROUGH MOVEMENT (ATT&CK T1075, T1076, T1028)

Once an attacker can execute code on an end user device, they will seek to establish command and control (C2) from the implanted device to an external C2 server. This allows them to carry out actions in the background as the user performs their normal activities. C2 implants communicate back to the C2 server using already known protocols and channels such as HTTP(s), DNS and TCP. More advanced C2 implants abuse existing services (Slack, DropBox, Google Drive etc) to avoid any anomalous connections from being spotted.

Recently, attackers have been abusing a technique known as “domain fronting”. This is where an attacker registers a service on a well-known public content delivery network (CDN) provider that other legitimate businesses use. The attacker is then able to point their service (attackerservice[.]provider[.]com) to their C2 server. They can then configure their implant to abuse the way in which CDNs route traffic, by ensuring the content of their HTTPS connection is made to legit[.]com, but the “Host Header” in the HTTP request is attackerservice[.]provider[.]com. To a network observer, the Server Name Indicator (SNI) field of the HTTPS Request/TLS Handshake is going to legit[.]com, but once that traffic is received by the CDN, it is decrypted and forwarded onto the server the attacker specified for attackerservice[.]provider[.]com. This can make detecting C2 additionally difficult for defenders.

An alternative to persisting on end user devices is to move laterally to a server or other system when processes such as command and control can run without the risk of regular log-off/shutdown. Virtual Desktop Infrastructure is a common location where a beacon can execute for a long period of time. This can negate the risk of writing to disk but may result in a loss of access if the compromised user ends their remote session or the VDI has a maximum session length. Organisations should consider the length of time a VDI or other virtual hosts can be active for before a restart, reducing this to the minimum required for the business.

Where user and server environments are not segregated effectively, and passwords discovered or obtained then a move to the server estate can prove fruitful especially where they have Internet access. Obtaining command and control from a server environment generally reduces the requirement for persistence though servers can also be overlooked in terms of EDR provision making things easier.

Lateral movement itself is often achieved by simply reusing the local administrator hash of a compromised host. It is extremely common for "gold builds" to have a static local administrator password that is never changed once the image is deployed. This can be by design (troubleshooting, break glass, ease of remote administration, familiarity) or simply because the risks are not immediately obvious.

This means that a suitably privileged compromise of one of these hosts allows the attacker to extract the password hash of the administrator account and perform a "pass the hash" attack with it. The password itself does not need to be recovered or "cracked", it can be used in its hashed form to authenticate with other hosts built from the same image. This can be even more successful for attackers when multiple different types of host (laptop, workstation, server) all have the same local administrator password.

Another method of enabling lateral movement from privileged access to a compromised host is harvesting credentials from memory.

Credentials obtained will be used in the stealthiest manner available to appear as normal as possible. Remote Desktop Protocol from the initial victim or from a jump host if that is what staff usually do makes detection difficult.

Given that it is often the case that privileged users have an elevated level of privilege across the entirety of the estate, this can mean that the compromise of a single end-user device with those credentials can lead to the compromise of all hosts within the estate.

Defending against persistence

Ensure that where practicable, VDI sessions have a limited run time and regularly close all processes opened by a user.

Defending against lateral movement

Practice good network segregation to limit the ability for attackers to move between end user devices and servers. Business roles and processes should be separated to avoid one being compromised and affecting all others. Servers should not have Internet access unless there is a business reason for it and then it should be limited to only specific protocols to designated endpoints with anything else causing an alert.

The "Local Administrator Password Solution" (LAPS) is a free package from Microsoft that allows system administrators to ensure that all hosts in their Active Directory estate have a suitably unique and random local administrator password.

This password is stored as a protected attribute against the computer object in Active Directory, allowing delegated access to it by users with a requirement to view the password. Microsoft provide a comprehensive implementation tutorial (https://gallery.technet.microsoft.com/step-by-step-deploy-local-7c9ef772/file/150657/1/step%20by%20step%20guide%20to%20deploy%20microsoft%20laps.pdf) for deploying LAPS.

It is vital during the deployment of LAPS that an organisation ensures the permissions to view and modify the LAPs password attributes are appropriate. It is our experience that a poorly deployed LAPS implementation can be just as useful to an attacker as not having it at all. Common mistakes include allowing all users or computers to read the attribute and not enforcing the policy across all hosts..

It should be noted that LAPS is a Windows only solution. Organisations should ensure they don't create the same pitfall of a shared administrator password with other operating systems (although, some LAPS like solutions do now exist for MacOS and flavours of Linux)

Defending against password theft from memory (ATT&CK M1043)

Credential Guard

Windows Defender Credential Guard is a feature of modern versions of Windows that virtualises the Local Security Authority Subsystem Service (LSASS) process outside of the running operating system. By preventing any user (including the NT AUTHORITY\SYSTEM) from accessing LSASS directly, instead forcing them through a proxy service, it becomes almost impossible for attackers to access these credentials in LSASS directly. Credential Guard can be enabled on appropriate hardware using Group Policy, InTune or via registry edits (https://docs.microsoft.com/en-us/windows/security/identity-protection/credential-guard/credential-guard-manage)

Protecting LSASS

If Credential Guard is not an option, then running LSASS as a protected process may be an option in many environments.

In Windows 8.1 and Server 2012 R2 and above, the LSASS.EXE process can be configured to run as a protected process. This provides additional protection for the LSA to prevent reading memory and code injection by non-protected processes and can help prevent Mimikatz and other tools from retrieving passwords and hashes from memory. It also prevents dumping the LSASS.EXE process memory, which can be used offline for the same purpose.

Although this setting can be bypassed with a driver designed for this purpose, the driver needs to be signed, may be susceptible to detection by anti-virus or EDR, making additional noise in the event logs. As such it raises the bar for an attacker considerably.

Microsoft have documented enabling LSASS protection here: https://technet.microsoft.com/en-us/library/dn408187.aspx

Limiting the impact of credential theft

Whilst focus should rightly be given to preventing the theft of credentials, it is important that any compromise of credentials has the least impact possible. For several years now, Microsoft have published guides on "Securing Privileged Access" (https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/securing-privileged-access). These documents detail a number of architectural and operational changes that can be made to limit the impact of credential compromise of a privileged user. We will not cover everything in these documents (they are a must-read for all Blue Teams/ System Administrators), but we would like to highlight a few key points and concepts:

Tiering

The concept of tiering is that every asset on your network has a worth/criticality/sensitivity/different impact from loss of confidentiality, integrity, availability.

In many organisations, this will range from end-user devices (least critical) to Domain Controllers (most critical). The idea of classifying devices into different tiers is to segregate administrator access to different tiers, so that an administrator of a lower tier device does not have administrative access to a higher tiered device on the same account.

This results in the compromise of a lower tier asset not automatically resulting in the compromise of a more critical asset.

Enhanced Security Administrative Environment (ESAE aka Red Forest)

An Enhanced Security Administrative Environment is dedicated administrative environment to host administrative accounts, workstations, and groups in a hardened and highly network. This allows for additional controls and protections to be applied that would not be practical or desirable in the user/ production environment. Typically, this will have a one way trust down to the environment it controls to facilitate authentication.

This can be combined with the concept of "Just enough Administration" (JEA) to ensure that even if an account was compromised, it only has administrative privileges for a specific task for a short amount of time (https://docs.microsoft.com/en-us/powershell/scripting/learn/remoting/jea/overview?view=powershell-7)

Privileged Access Workstations (PAWs)

A privileged access workstation is a dedicated host that is hardened against and insulated from internet-based attack vectors. By removing the day-to-day work from the PAW and focussing on administrative tasks, the attack surface of internet browsing, email and logging into less secure systems is removed.
A robust AppLocker policy can be applied, as the software running on the PAW should be vetted and explicitly allowed.

In order to save cost and prevent administrators from requiring two devices, a common deployment pattern is to host a day-to-day user Virtual Machine on the PAW. Given the relative rarity of virtual machine breakouts, this is a good compromise between security and usability. (https://docs.microsoft.com/en-us/windows-server/identity/securing-privileged-access/privileged-access-workstations)

Detecting and defending against Command and Control

Defeating Domain Fronting

Many TLS Intercepting Proxies are now able to detect the mismatch between the SNI of the TLS Handshake and the content of the host header. This would prevent (or at least, flag) any communication using this technique. It should be noted that public services that are routinely certificate pinned (therefore, unlikely to be subjected to TLS Inspection) are also vulnerable to domain fronting, so may allow this to bypass the proxy.

USING THE ACCESS ALREADY SECURED (ATT&CK T1078)

If the initial staff targeting and spear phishing were successful, then the compromised staff member may have access to the objective system or use other functionality such as Citrix to access it. With direct access to the objective the attacker can identify if their privileges are sufficient to achieve the goal. Regular monitoring by the attacker of the staff member to identify how they access the objective system may be required. Staff members accessing critical systems is necessary for businesses to operate, identifying malicious from normal behaviour can be extremely difficult.

Defending against abuse of legitimate access

Ensure that all staff with administrative access to systems are provided with a separate account from their day to day account. This account should not be used for activities such as email or web browsing and should have suitable restrictions in place to ensure it can only access systems required by the staff to carry out their work.

Attackers may need to perform actions outside of normal working hours for the staff member or take actions within the system that are unusual. Organisations that understand normal behaviour and can identify anomalies for investigation will have a much better chance of identifying malicious behaviour.

EXPLOIT CENTRALIZED IDENTITY AND ACCESS MANAGEMENT (ATT&CK T1078)

Organisation's environments often use a variety of IAM systems, from Active Directory to cloud IAM services to custom-built solutions. Where initial attacker access is obtained but the staff member does not have access to the objective then NCC Group's red teams often targets IAM systems. Often these systems have been created with lower security requirements as a level of trust is assumed for the internal network or the information presented by the staff member is deemed secure. When these systems are compromised they regularly provide privileged access to critical systems.

Defending against Password reset misuse

Organisations should ensure that the process to reset a password should require more than simple information obtainable within the corporate domain. Actions such as requiring a manager to approve and them having to contact the member of staff making the request by phone adds complexity that an attacker may find difficult to circumvent.

Out Phase

The out phase is where the threat actor concludes the operation and achieves their goal.

SECURING THE REQUIRED ACCESS (ATT&CK T1078)

With access to a range of user accounts from previous activities then obtaining the right account access is usually a matter of time. Where separate accounts are used between the office domain and the objective system then identification of the correct jump host or access application is required. Accessing specific user drives with elevated privileges can obtain this type of information and can also reveal passwords being stored by the user in their home folder. Staff members remembering multiple passwords and changing them regularly remains an issue and therefore reuse occurs. Where password creation is left to users then meeting the minimum complexity requirements with a word they can remember is regularly the answer.

Defending against password reuse and minimal complexity

Organisations should provide appropriate password creation and management solutions to their staff. Auditing password hashes looking for duplicates can identify those that are using the same minimum complexity password (ATT&CK M1027).

Azure Active Directory

For organisations using AzureAD, Password Protection features exist that enforce custom password filters both on-premises and in the cloud. This allows defining a list of banned words/phrases that may commonly be used a root words for passwords in an enterprise (Company name, office names, department names, product names etc).

Microsoft document the deployment of this technology here: https://aka.ms/deploypasswordprotection

Password Filtering

It could be said that the default Windows password complexity setting lacks any finesse. By default, "Password1" would be sufficient to meet the complexity requirements for an 8-character password. The default rules (https://docs.microsoft.com/en-us/windows/security/threat-protection/security-policy-settings/password-must-meet-complexity-requirements#reference) do not allow the system administrators to define a list of disallowed words or phrases.

Microsoft do support allowing system administrators to shim in their own password filter (https://docs.microsoft.com/en-gb/windows/win32/secmgmt/password-filters?redirectedfrom=MSDN) to implement a list of banned words/phrases/write custom rules. This solution is not without a significant overhead, as it does require installation and maintenance of the password filtering software on each Domain Controller. A number of open-source implementations of password filtering exist that could be evaluated if AzureAD is not an option.

OBJECTIVE ACTIONS

Confidentiality; Data exfiltration via the C2 channel which if we are at this stage has not been detected. Screenshots are the simplest method for the attacker though make the data less useable. There is a possibility that any bulk data needs to be staged on hosts before exfiltration. Identifying odd network traffic between hosts that would not normally communicate may catch the internal transfer. Depending on the size of the data then large transfers outside of the network may be identified but given modern communications this is unlikely. (ATT&CK T1041)

Integrity; this can be difficult for an attacker to avoid being detected though it is dependent on the objective system. Changing small pieces of information within applications or databases that staff members have access to is possible. Complex transactional systems with subsequent business processes where integrity checks are performed can identify changes and are not known by the attacker until too late. (ATT&CK TA0040)

Availability; privileged access to underlying operating systems is the simplest method and by the time this has been achieved there is little an organisation can do. (ATT&CK TA0040)

Conclusion

Offence is comparatively far easier than defence due to the ability to leverage all sources of information and techniques within legal and ethical boundaries. An attacker often has to only succeed just once where most enterprises need to defend against everything. This imbalance stacks the deck significantly to the advantage of attackers.

However, a resilient organisation can make the offensive side work far harder and fail more often if:

- They know their threat model and the external information available to an attacker
- They know their external attack surface and actively manages and monitors it
- They know what is in their estate, manage it, have effectively segregated it by its function, criticality status and business owner and are able to monitor it in near-real-time
- They have the right visibility and controls across their estate with an understanding of what is normal
- They have an ability to; detect, contain and respond to the anomalous
- They have strong identity and access management and have provided staff with tooling to complement it and manage their own passwords effectively including mandates multi-factor authentication
- They ensure information repositories have appropriate access controls and do not contain credentials

Technical Analysis

How to build resilience against offensive cyber operations

By Liam Stevenson,
Associate Director
at NCC Group

IT leaders have accepted that it is a question of when, not if they will experience a cyber attack. However, the question of their organisation’s resilience against those attacks can be more difficult to answer.

At NCC Group, we conduct offensive cyber operations, or Red Teaming as it is known when delivered commercially, to help organisations assess their defensive security posture against the attack scenarios used by real-world threat actors.

With that in mind, here are some of the most commonly successful techniques that we use, and how you can strengthen your organisation’s resilience against them at every phase of an attack.

Reconnaissance

Before an attack, threat actors will survey target organisations to gain information about its people, processes, systems and partners. All organisations leak information that is difficult to remove, so conduct regular reconnaissance against yourself to establish the scale of leakage in your business.

For example, developer repositories can contain internal domains and words specific to your organisation and critical systems, so you should remove accounts, change hostnames or use newly written code to mitigate their effectiveness for attackers. In addition, you should check for credentials or private keys that are committed to repositories by accident, changing these immediately if found.

Similarly, threat actors can often view an organisation's cloud-based systems on its public DNS, so ensure that you understand the number, location and typical activity of your authentication portals to make it harder for them to breach with this information.

Finally, avoid leaking information about key technologies and internal resources through social media, job advertisements and external presentations. Conduct reconnaissance to identify the problem areas, sanitise your documents and educate staff to be generic in their social media activities.

"Security through obscurity" isn't perfect. However, by limiting the information available to an attacker, you can increase your chances of detecting them by forcing them to take more risks in the reconnaissance phase.

> "Security through obscurity" isn't perfect.

Our Red Team typically breaches an organisation in one of three ways: exploitation of known vulnerabilities, exploitation of external authentications and phishing attacks.

Many organisations have a timely patching process for their external services, but their internal patching is often slower. As a result, opportunities to breach through known vulnerabilities on outdated operating systems can often be found on the internal network.

To mitigate this, ensure that you have an accurate database of assets which includes their locations, assigned responsibilities for monitoring and raising new patches and current software versions installed. You should also ask your managed service providers to define the patching processes for their assets and push them to do better if it does not meet your standards.

If patching isn’t possible, build additional monitoring and controls around your vulnerable systems to ensure that only the services and network routes required for their operation remain open.

Externally, you can make it harder for threat actors to breach by removing single factor authentication for any services that are exposed to the internet. Even if two factor authentication is in place, you should log authentication attempts and monitor your endpoints for unusual activity including the number, time and location of attempts and repeated attempts with the same account names.

Finally, phishing attacks remain successful as an initial access technique, so deploy filtering systems, limit your communication resources to those with secure reputations and use endpoint detection and response to protect against persistent attacks.

You should have a process for staff members to report suspected phishing emails to their IT team, and a process for your colleagues to verify that calls are genuine to mitigate malicious ‘vishing’ calls that are commonly used to support phishing campaigns.

Through Phase

Once an initial foothold has been achieved, threat actors can search communication platforms, information repositories and network shares to gain access to additional hosts and functionalities. To slow their progress, remove repository access rights for those that do not require them, remove sensitive information and change exposed credentials.

To run command and control processes without the risk of being logged off or shut down, attackers can also move from end-user devices to locations including Virtual Desktop Infrastructures (VDI). To hinder this, reduce your VDIs' maximum session lengths to the minimum required for the business.

You should also segregate your network to limit any movement between end user devices and servers and restrict internet access to the servers that require it.

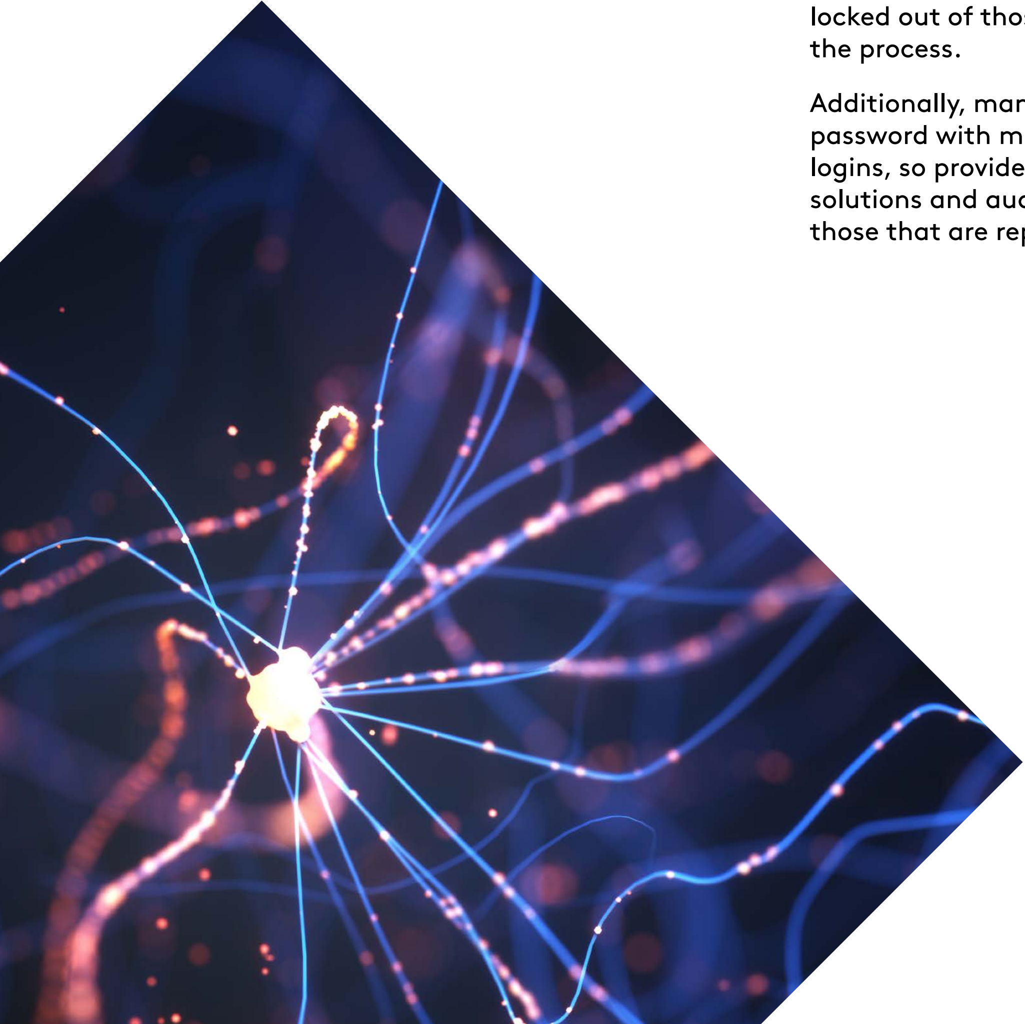

Attackers will assess the access of users they have compromised to determine how to reach or achieve their objective. As such, ensure that all staff with privileged access to systems have unique administrative accounts that are different to their everyday accounts.

If the staff member does not have the required access, our Red Teams often target Identity Access Management (IAM) systems, which are often created with lower security requirements due to an assumed level of trust on the internal network. Compromised IAM systems can provide access to critical systems, but you can strengthen your resilience and avoid being locked out of those systems by adding complexity to the process.

Additionally, many staff members still use one password with minimal complexity for separate logins, so provide password creation and management solutions and audit password hashes to identify those that are repeatedly using the same ones.

Out Phase

Even when a threat actor is ready to conclude their offensive cyber operations, there are still things you can do to make life difficult for them.

For example, attackers might need to stage large quantities of data on hosts before extracting it, so monitor for odd network traffic between hosts that would not normally communicate. Finally, establish standard ways of working for sharing information outside of the organisation, such as a secure file share. By doing this attackers will be forced to use 'non-standard' ways of exfiltrating information making it easier to spot and prevent.

Conclusion

Whether it’s a Red Team or real-world threat actor, the wealth of information and techniques available to offensive cyber operators means that they will always have the advantage over a defending organisation.

However, by following the steps outlined in this blog, you can significantly improve your resilience to cyber attacks by making the offensive side work harder and fail more often. When it comes to protecting your assets, those extra seconds, minutes and days can make all the difference.

If you want to take the next step on your cyber security journey, we can help - get in touch with our team today.

Insight Space

cyber insights
programme

Ollie Whitehouse

Technical Viewpoint

Building your cyber defence:

start with the basics

Cyber defence is not about magic amulets, it is about discipline

When we talk about cyber resilience and defence, organisations are often swayed by solutions that promise untold success with comparatively little effort. These solutions will often talk about how they will be able to shore up an organisation's cyber defences and increase its overall resilience quickly, cheaply and with very little effort.

> The reality, however, is that resilience takes hard work. In order to establish and successfully achieve a state of robust cyber defence, far more effort has to be put into transformation and operational improvement than in magical solutions.

Insights from a nation-grade offensive capability

NCC Group is one of the few organisations in the private sector that operates world-class, at-scale offensive and defensive capabilities under one roof.

On one hand, we help clients attack hardware, software, systems, processes and people to assess real-world impact and resilience. On the other, we respond to cyber incidents globally, monitor hundreds of organisations, millions of endpoints and track threat actors of all types. All against a backdrop of helping organisations design, build and operate resilient and compliant products, services and organisations.

Based on these insights, we recently had our Red Team distil the factors that make their lives easier, as well as those that force them to take more risks and invariably fail more if implemented.

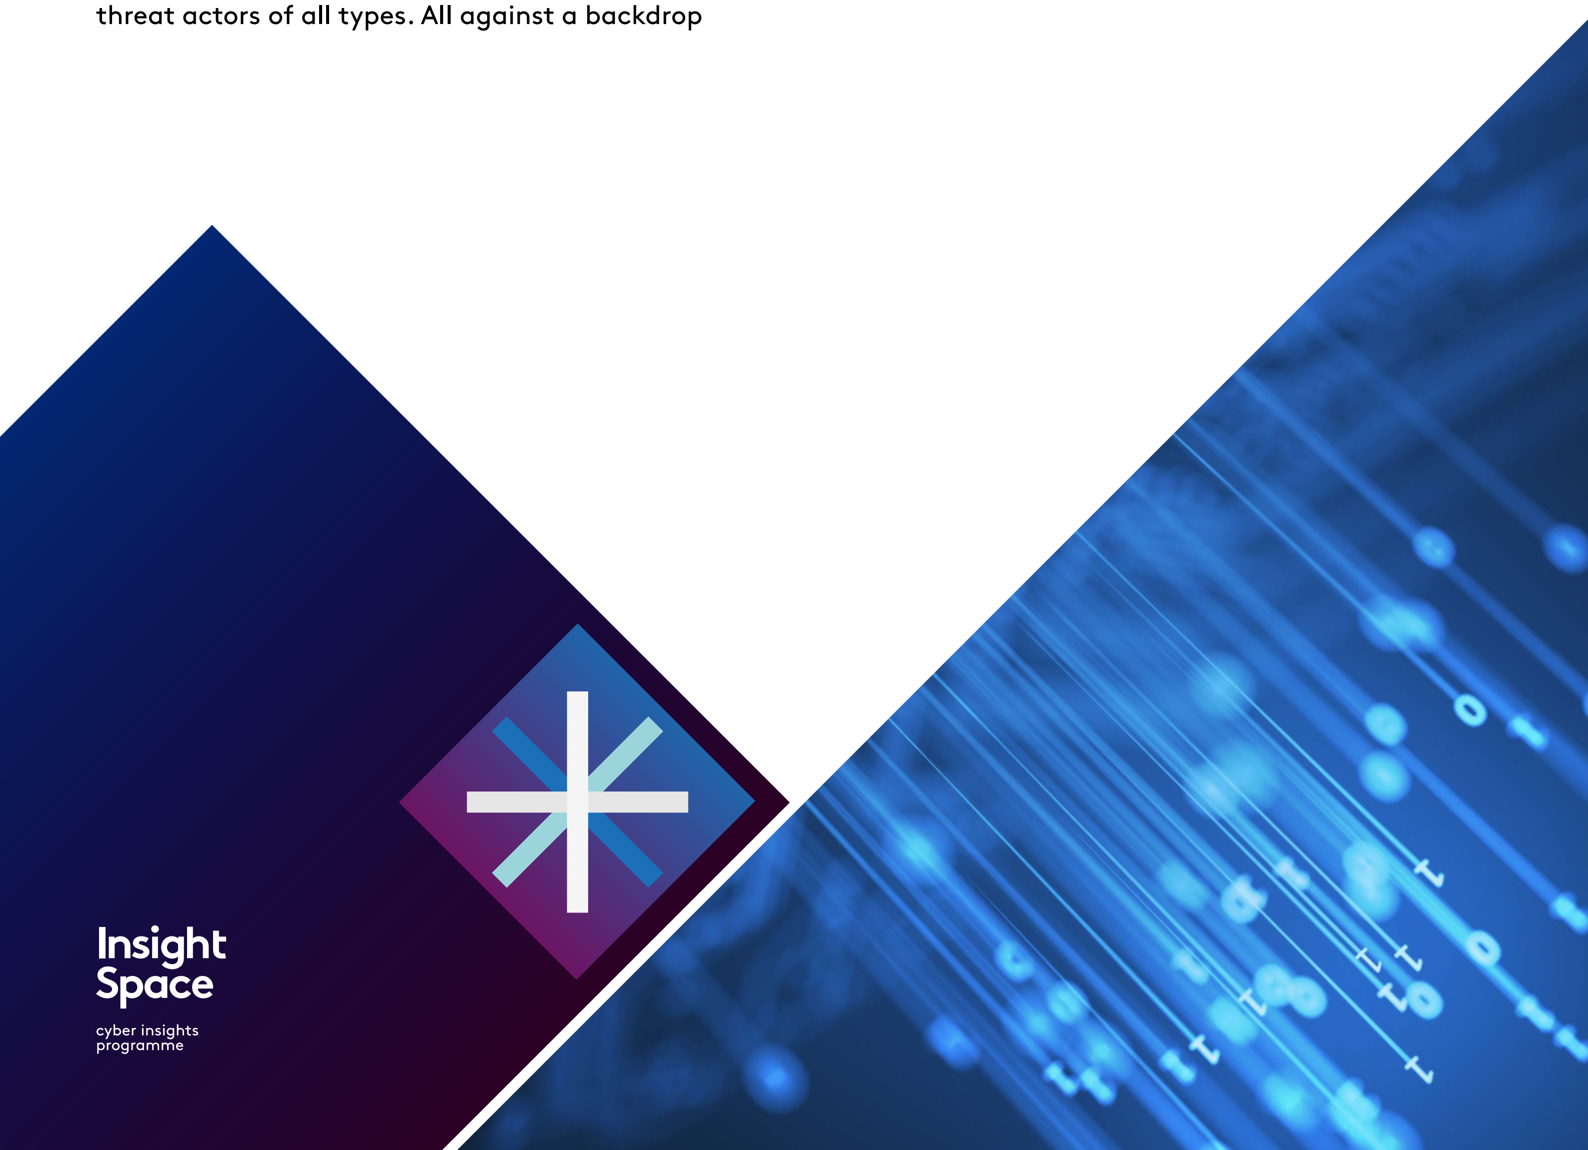

The basics are crucial

The key findings from our Red Team were:

- Organisations leak too much information.
- Organisations don't know what is in their computing estate.
- Organisations don't reliably maintain their computer estate.
- Organisations place trust in networks and not solely on identity of devices, its software and users.
- Organisations don't have visibility of what is happening in their estates.
- Organisations place too much faith in their security technologies and their efficacy.

The above makes pretty uncomfortable reading. However, the reality is that organisational complexity, apathy towards robust technology management and, often, an unwillingness to change due to fear of failure can lead to a toxic mix which undermines cyber defence.

Without paring back this problem, addressing the root causes, and driving discipline and rigour, there is little chance organisations will repel an attack, let alone detect and contain one in a reasonable period of time.

Addressing the things we can

Within this list, there are some issues which can be difficult to address. For example, trying to stop all information leaking is a path to likely failure. The goal should be that the leaks don't impact the security of the organisation materially.

Instead we want to focus on the areas that, if addressed, will have the greatest impact.

Threat actors will always be able to adjust their methods to counteract defences or particular approaches. They will be able to identify new and novel techniques to breach networks and systems, know where to persist (or not), and establish command and control channels.

Instead we should focus on making it awkward for them, and ensure that anything – and we mean anything – which happens within a material environment is logged. Through robust understanding and visibility we can make it difficult for threat actors to carry out attacks.

Know what is out there, who owns it, what it does, its lifetime and its purpose

Creating an asset inventory of both physical and virtual assets, including those which are ephemeral, is key. Build it, maintain it, and care for it – it is the single source of truth of what makes up an organisations technology estate and its systems.

Without a near real-time asset inventory, it's almost impossible to know when a part of your system is end-of-life, unmaintained, or previously compromised.

Your asset inventory in 2020 should go all the way up the stack including various "as a service" components, whether it's Infrastructure, Platform or Software, including serverless concepts.

Log everything that is critical and practicable for as long as possible

Detection and analysis rely on effective logging. With incomplete or inconsistent coverage, detection, containment and response missions will not be as effective. We've seen many massive breakthroughs in either detecting or understanding the activity of a sophisticated actor because of high quality logs stretching back years.

If we solely rely on the logs from security products to feed detection, then there is significant exposure. NCC Group's Red Team, like other sophisticated threat actors, spends time and money to acquire and subvert security products and logging in general.

However, the reality is that attackers can't spend time trying to avoid it all if we expect to be successful in our mission. As such, trying to access an environment with mature and persistent logging is like trying to break into a car in a flood lit car park with CCTV and a helicopter circling overhead, whilst trying to break into a car – in other words, awkward and difficult, with a high chance of being caught.

Enrichment in the understanding of assets and users

Not all assets within organisations are created with equal value. As such, understanding of a system, function or user and their role provides valuable context. This context is useful in understanding the risk they present, as well as the impact when security events occur.

By ingesting a wide range of supplementary information, we can build a level of understanding of an organisation's detection and response capabilities, as well as an understanding of what things do and why they matter.

Once we have the foundations we can build with confidence

Much like business operations, good management information and intelligence enables you to operate and make decisions with confidence.

Similarly, with cyber defence, if we have visibility over our systems, their state, their purpose and who is accessing them, we can have confidence that if something suspicious happens now or in the future, we'll have information on the event.

These individual threads on their own may not tell us anything, but when weaved together they provide a rich tapestry, enabling understanding. This understanding gives confidence, and confidence in turn enables business agility and quantified risk taking.

If you want to take the next step on your cyber security journey, we can help - get in touch with our team today.

Insight Space

cyber insights programme

Business Analysis

Four resilience questions that boards should ask before commissioning a red team

By Paul Vlissidis, Technical Director and Senior Advisor at NCC Group

Cyber is one of the biggest risks to businesses worldwide, but many boards still consider it an IT problem rather than a core business issue. It's vital that business leaders take responsibility for their organisation's resilience against cyber attacks, but knowing where to start can be difficult.

At NCC Group, we conduct offensive cyber operations, or Red Teaming as it is known when delivered commercially, to help organisations assess their defensive security posture against the attack scenarios used by real-world hackers.

Red Teaming can establish whether your cyber investment is paying off and identify where you should allocate resources to improve your resilience.

However, you should only commission a Red Team if you know the answers to some fundamental questions about your current level of resilience. Otherwise, you risk raising more questions than answers about your business's ability to deal with a cyber attack.

With that in mind, here are four questions to ask before you commission a Red Team.

1. WHAT ARE THE MOST CRITICAL ASSETS IN MY BUSINESS?

It’s impossible to achieve total security across your IT estate, so it’s important to prioritise the most critical assets in your business. If you cannot identify what those assets are or where they are located, a Red Team exercise might uncover gaps in your security posture that do not present a significant risk to your organisation.

On the other hand, testing your most important assets in a Red Team exercise can help you understand how they could be compromised and used against you. It can also reveal whether hackers could use other parts of your network to breach those assets.

Speak to your IT team to find out where you hold sensitive company information and which systems are business-critical.

2. WHAT ARE THE THREATS TO MY BUSINESS?

Any analysis of an organisation’s security posture should start with the realistic threats that they are exposed to, and the same is true of Red Teaming. For example, if your biggest threat is theft of company information by a disgruntled employee, it would not be helpful to engage a Red Team equipped with state-sponsored external hacking skills.

By replicating the attack scenarios that you are most likely to face, Red Teams can deliver more relevant insights that will better inform your investment to improve resilience.

Cyber attacks are often perceived as threats from the outside that must be defended internally. However, attacks can originate internally, externally or from third parties such as suppliers with access to your systems, so think carefully about who might target you and why, and how they would carry out an attack.

3. WHAT IS MY BUSINESS'S LEVEL OF CYBER MATURITY?

Before you commission a Red Team, it's crucial to understand where you are in terms of your cyber maturity. If you are not practicing basic security hygiene throughout your organisation, you probably aren't ready for a Red Team exercise. If this is the case, consider working towards a standard like Cyber Essentials to get the basics in place.

However, if you are already using processes and systems like Managed Detection and Response (MDR) to protect your assets, Red Teams can be a useful test of their effectiveness. Again, their insights can advise you on how to complement and improve your security estate to increase your overall resilience.

4. THIS IS JUST THE BEGINNING OF THE JOURNEY – ARE WE READY?

For many businesses, undertaking a Red Team represents a significant security milestone but the issues it identifies often require a long term security transformation to improve resilience.

It is very rare indeed for a Red Team exercise to deliver a clean bill of health and their reports often make uncomfortable reading for the board. You should not commit to a Red Team unless you are prepared to follow up with the security improvements it recommends.

Improving Resilience

No organisation will ever be completely secure against a cyber attack. However, by asking yourself these questions, you can ensure that your Red Team exercise is more effective and take responsibility for making your organisation as resilient as possible.

For more information on how to increase your resilience, get in touch or read the full report based on our own Red Team operations here.

If you want to take the next step on your cyber security journey, we can help - get in touch with our team today.

Insight Space

cyber insights programme

Business Viewpoint

Think like a threat actor:

Counter common techniques to secure your organisation

The techniques used by criminals to attack IT systems, whether they're criminals or nation-state actors, are often the same. It can mean that it's difficult to tell them apart.

However, it also means that if you harden your systems and educate your staff about commonly-used techniques, you will make your organisation more resilient.

This paper is designed to help you understand the techniques used by the criminals and the common methods to stop, delay, or detect them. This should allow any executive to challenge their IT team, security manager, chief information security officer or IT guru, at any level, on the practical security measures in place to defend their organisation. Finding the optimal level of security to reflect the risk and investment required to defend against it is difficult. But with this knowledge, at least a better-informed conversation can take place.

We use the same methods employed by criminals to test an organisation's ability to stop, or spot and react to, an attack. Our 'Red Team' simulates criminal activity to help identify any weaknesses, so that you can fix them before suffering a real attack.

We know that criminals and hostile states regularly follow the same routes during an attack and so the methods and techniques have been codified into the "Mitre ATT&CK" framework. This framework makes it easier for organisations to examine their defences to see what they have in place to stop or detect an attack at the various different stages.

We've identified issues throughout this paper that have an associated Mitre ATT&CK number, which describes a particular technique, as well as outlining the steps that attackers typically take when trying to break into business systems. To get reassurance that things are under control, ask to see the evidence that each issue and technique has a mitigation in place.

This is where the criminals, known in the cyber security world as "threat actors", look for information about the organisation they are targeting.

Our 'Red Team' simulates criminal activity to help identify any weaknesses, so that you can fix them before suffering a real attack

Reconnaissance

Most organisations and their staff leak information and it is very difficult to control such leaks. But identifying what information is available, most often on the internet, can help the organisation identify where accidental oversharing of internal information is taking place and help address it.

Common places to look for information leaking include:

- Staff job adverts - which may reveal your IT systems, either current or planned
- LinkedIn - where project names, IT systems, security clearances, roles and responsibilities are public, and email and phone numbers are often available
- Staff and organisational social media accounts - particularly where staff passes are included on photos, and people talk about internal issues and projects and identify colleagues
- Council, real estate and office building websites - these often include floorplans with too much detail, such as access control points and IT rooms, and policies for getting into a building
- IT developers' code stores - such as GitHub, BitBucket, GitLab, and SourceForge where internal information including individuals' names, IT details and even passwords are regularly found
- Breach databases - where email addresses and passwords can be found, giving a helping hand to the criminals

Once the threat actors have enough information to target an organisation, they move on to the next stage known as In Phase.

What to ask your security team:

- Do we actively look for publicly available information about our organisation? If it is sensitive, what steps can we take to remove or amend it?
- How can we avoid this information being published in the future?

In Phase

In this stage, attackers look to gain an initial foothold on the network and, using the Mitre ATT&CK framework, may seek to **exploit vulnerabilities**, known as ATT&CK T1190.

Trying to remove all the vulnerabilities from an IT system is extremely hard. The usual mantra of 'patch your systems' to ensure that the latest version of software is in use on every machine across an entire network of laptops, routers, switches, servers, printers and firewalls etc. is very difficult. This is particularly true where old, legacy systems and third-party equipment and processes are necessary to deliver business services.

Start by ensuring that systems that can be seen from the internet are as up-to-date as possible. This does mean that an accurate record of everything in the system, no matter how old or unused, will help ensure that nothing is missed. Too often compromises have come via old equipment or websites that have been forgotten about but are still connected internally.

> Regular scanning for known vulnerabilities should help identify when something needs to be looked at and prioritised for replacement.

If the system is no longer supported by the manufacturer or supplier, but is business-critical, then additional protection must be put in place to defend it.

You should assume that your IT systems have been attacked and the attacker is on the inside, known as an 'assume breach' mentality. Having this mindset means that you will understand that even internal systems, not visible from the internet, are vulnerable.

To maintain security, you should put in place a regular review to identify updates and new vulnerabilities. This means an internal scanning process should be in place, as well as an external scan.

Another essential way to improve resilience, described as the ability to resist and then adapt to change, is to ensure that any system that can be seen from the internet is hardened. This is to counter the **external authentication exploitation** known as ATT&CK T1078.

Regular scans can also help to identify misconfigured systems that are connected to the internet – often, systems that were thought to be 'internal only' have been a regular way in for both threat actors and our Red Team. Identifying these weak spots and removing any points of access that only require a username and a password from internet-facing websites and systems is essential.

Instead, these systems should require a form of multi-factor authentication. Ideally, this will reply on the users entering a username, a password and a randomly generated number from an app on their phone, although other methods using keys, text messages, and dongles are also available.

Just having multi-factor authentication is not enough. There needs to be some monitoring or logging of attempts to gain access. And then an investigation of unusual attempts such as logging in from unusual locations, at strange times, using odd equipment or from different locations with the same account details in a very short period (known as impossible travel). There are lots of unusual behaviours that the threat actors use to log in, but effective multi-factor authentication will resist many of them. It will also help to limit success of the phishing (email-based) and 'vishing' (voice or social engineering attack) identified as ATT&CK T1192, T1193, T1194, TA0003.

Phishing emails are one of the most common ways that organisations are attacked. The more sophisticated attackers will use the information gained during the reconnaissance phase to send emails targeted at specific individuals. This is known as 'spear phishing' and frequently targets IT administrators with high levels of access, personal assistants with access to their boss' emails, finance departments, and individuals with high levels of authorisation to internal systems – including the executives such as you.

Phishing emails usually seek to get a response, such as opening a document or clicking on a link, although it may be the precursor to a telephone call or vishing from an individual claiming to be from an IT department, a client changing account details or a head hunter wanting to get information on your availability for a new job.

Security awareness training is the only way of preparing staff to manage and resist vishing attempts. It needs to be bite-sized, frequent, targeted to particular internal audiences and focused on the organisation.

> The impact of phishing emails can be reduced with technical measures and should be the focus of a regular conversation with your security team to ensure that these measures are cost-effective and up to date.

What to ask your security team:

- Do we have visibility over our entire IT estate, and have processes in place to keep devices and systems up to date wherever possible?
- Do we regularly carry out internal and external vulnerability scans of our systems to uncover and resolve known issues?
- Do we have regular and targeted security awareness training in place?

Through Phase

If the measures to keep the threat actor out of your systems have failed, they will seek to move across the network. This is because the initial point of the compromise is rarely the place with the information or access that they want.

Often, they will seek to use **internal information repositories** ATT&CK T1213, T1039, T1081 to further their attack.

These internal repositories might be something as simple as a list of passwords stored in a Word document on a desktop, password reset processes and systems, or an internal site or HR database. This information could give the attacker essential access to other parts of the network, known as lateral movement, or to higher levels of authorisation, known as privilege escalation.

To defend against the harvesting of such information internally, a regular review of where sensitive information is stored, and who has access to it, should be carried out. The logging of wide-scale internal searches, and alerts for specific keyword searches, such as “password”, can also warn that an attack is taking place.

You could also ask about the use of ‘canaries’, not the birds but a small snippet of code that is disguised as something interesting such as a file, document, photo or even an entire fictional computer that sends an alert when looked at.

The threat actor might be able to **use the access already secured** ATT&CK T1078 to obtain the information they are looking for. One simple measure is to ensure that staff with high privileges use separate accounts, or even separate machines, to do technical administration and not the same one they use for email and surfing the internet. This restricts access and reduces vulnerability.

The threat actor will probably want to ensure that, having gained access to the system on one machine, they can continue to do so even if it is turned off or rebooted. This may require them to **maintain access through movement** ATT&CK T1075, T1076, T1028 which essentially involves moving to a better place on the network.

In many modern systems the servers are not regularly turned off so the attacker may look to sit there as it gives them access to different users and often a path out to the internet.

The separation of internal networks with additional control, limiting access to the internet, and ensuring that endpoint detection and response software is installed on servers as well as laptops or computers used by staff (endpoints, in IT terms), will all help limit the criminals' activity.

Regularly targeted servers or systems often control identity and access management, such as the Windows Active Directory. So, an attempt to **exploit centralized identity and access management** ATT&CK T1078 is common among criminals and hostile states.

Unfortunately, these are often seen as internal systems and are usually not as well protected as they should be. However, when they are compromised, they regularly provide high level or privileged access to critical systems. Although the resetting of passwords is a frequent occurrence in many organisations, a good process, such as having a manager approve it or requiring a telephone call to the member of staff requesting it, can be an additional hurdle that an attacker may find difficult to overcome.

Having gained access and the ability to move across the network the attacker will begin to look to move to the phase where they can see their objective – known as the Out Phase.

What to ask your security team:

- Do we know where sensitive information is stored, and who has access to this? If not, can we regulate this?
- Can we put stronger password reset processes in place?

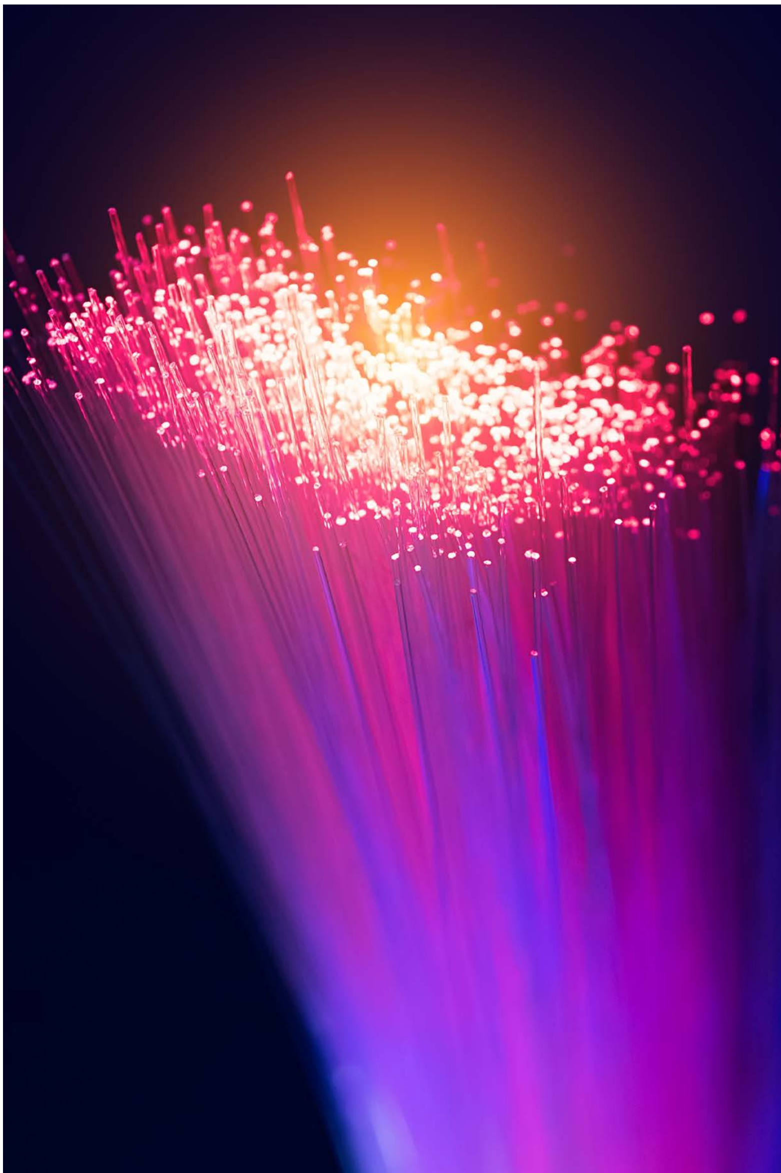

Out Phase

The attacker will need to be confident of **securing the required access** ATT&CK T1078, which will probably involve using the right password.

Our Red Team frequently see passwords based on a single word with a number and changed regularly with a predictable pattern. January2020, February2020, March2020 is very likely to be followed by April2020 if the system requires a new password every 30 days. Password managers are a very sensible way forward, particularly when combined with other secure policies.

Having secured their access, they will be looking to steal or exfiltrate the data. There are many different ways of doing this so ask if exfiltration over a command and control channel ATT&CK T1041 has been mitigated. Then ask whether all the techniques covered by ATT&CK TA0040, which sees the confidentiality and availability of data compromised, would be detected and stopped.

What to ask your security team:

- Could we implement password managers?
- Have we mitigated exfiltration over a command and control channel?
- Can we detect and stop attack techniques covered by ATT&CK TA0040?

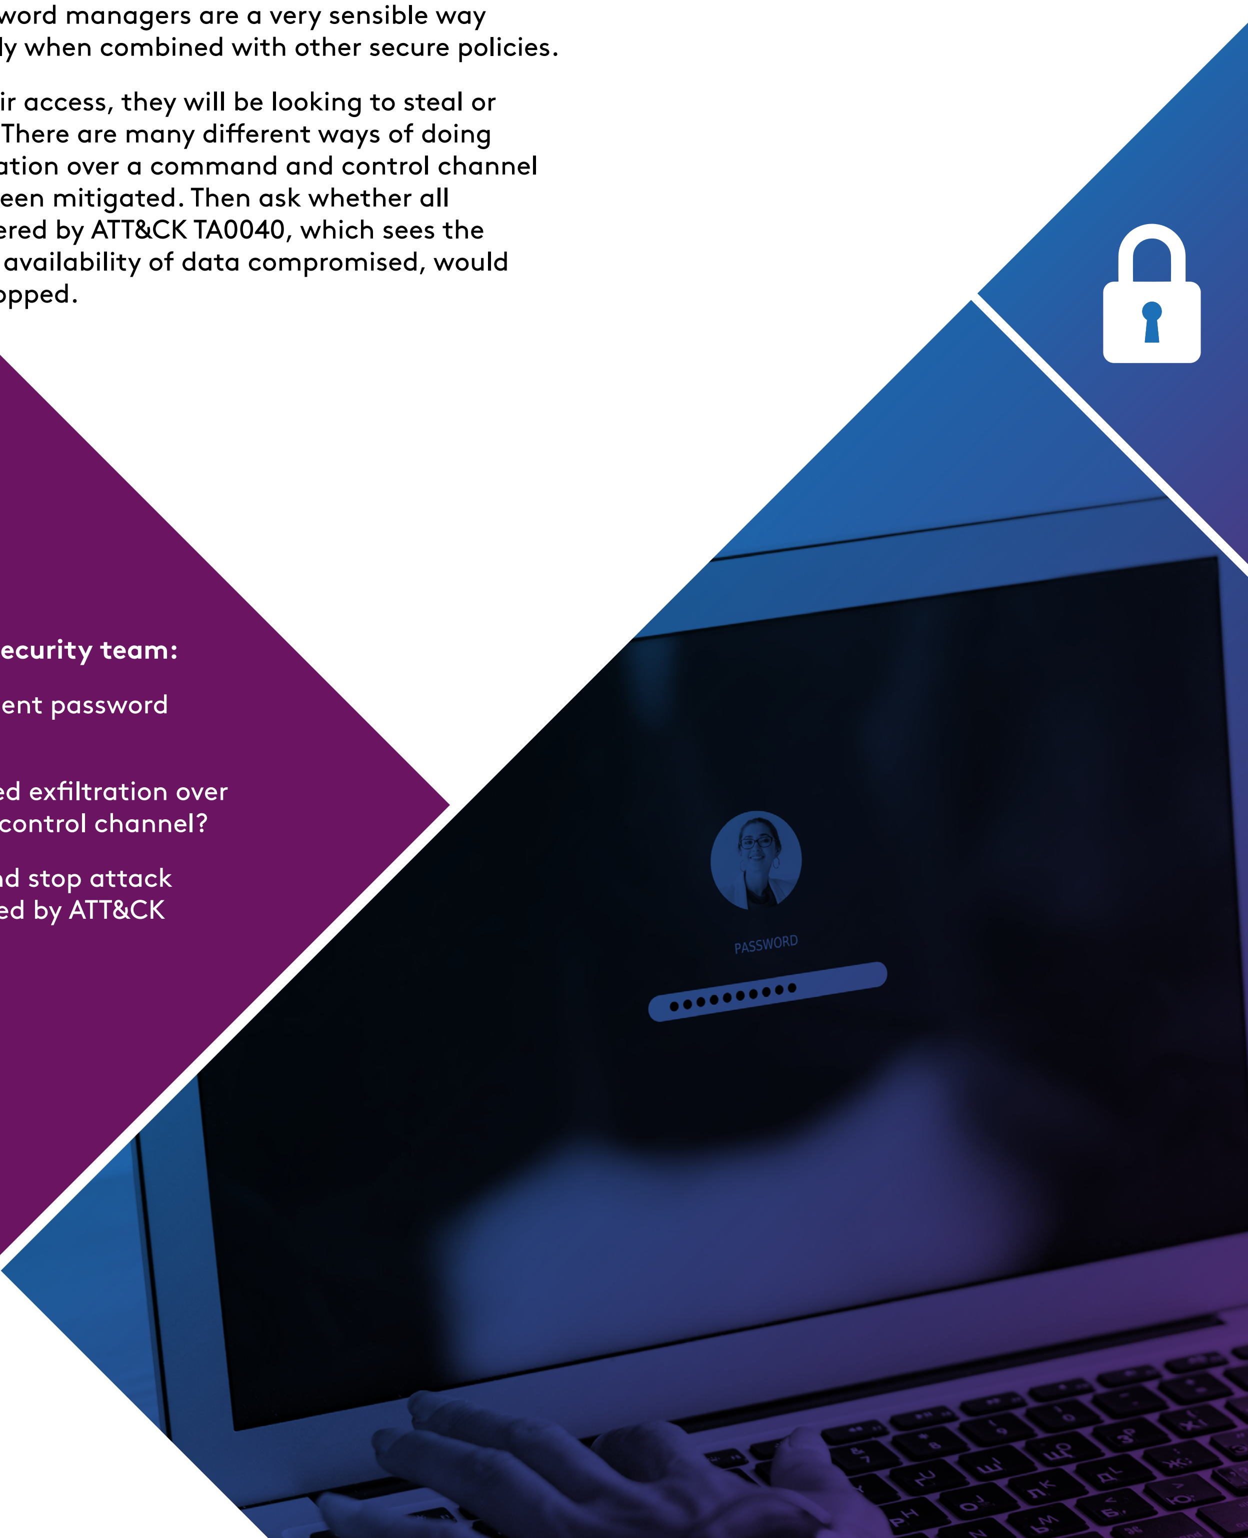

Conclusion

Securing your data from all the different methods and techniques the opposition use is undoubtedly difficult.

However, if you can challenge your security team to use the Mitre ATT&CK framework to identify the way that they operate, you stand a better chance than many.

Red Team attacks are a great way of checking that these mitigations are in place, but if you are starting out and want to give your security team a bit less of a challenge you might look at Purple Teaming. This is where our Red Team sit with your security, known as the Blue Team.

The two teams then work together to test the methods and techniques described above to test the security of your organisation and make your Blue Team even stronger.

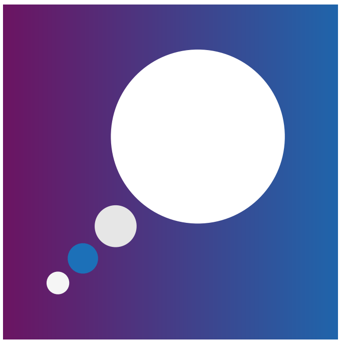

Whatever you do, make the opposition work harder by ensuring the Blue Team:

- Know their threat model and the external information available to an attacker
- Know their external attack surface and actively manage and monitor it
- Know what is in their estate, manage it, have effectively segregated it by its function, criticality status and business owner and are able to monitor it in near-real-time
- Have the right visibility and controls across their estate with an understanding of what is normal
- Have an ability to detect, contain and respond to the anomalous
- Have strong identity and access management system, and have provided staff with tooling to complement it and manage their own passwords effectively, including mandatory multi-factor authentication
- Ensure information repositories and internal file stores have appropriate access controls and do not contain credentials

To make your organisation more resilient, challenge your security team today and start more regular conversations with them. And remember, you can call your NCC Group Senior Adviser to support and guide you as you make your organisation more secure.

Insight Space

cyber insights programme

Business Panel

Making your cyber resilience budget work smarter

By Ade Clewlow, senior advisor at NCC Group

Around the world, businesses in all sectors are in the midst of huge change. With a global recession, changing ways of working and reduced investment, business leaders are now having to approach cyber security in new ways.

But with security budgets under increasing pressure, how can CISOs and senior leaders prioritise spending and continue to build resilience within their organisation?

I asked NCC Group experts, Paul Vlissidis, NCC Group Technical Director, and Lawrence Munro, our Director of Innovation, three of the most pressing questions facing CISOs today.

1. HOW CAN CISOS AND IT SECURITY PROFESSIONALS PRIORITISE INVESTMENT IN A CHALLENGING ECONOMIC ENVIRONMENT, WITHOUT PUTTING THEIR NETWORKS AT RISK?

Paul Vlissidis: "When budgets are tight, CISOs have to get used to doing the same, or more, with less. Unfortunately, it's a less than ideal time to be cutting security spending. Research from Portsmouth University has revealed that fraud rises during recessions, often linked to an increase in ransomware and other potentially lucrative attack methods.

"So, what can businesses do? Over the last few months, organisations have had to accelerate their digital transformation journey – and now is a perfect time to revisit and re-evaluate the decisions that have been made.

"As always, basic security hygiene – credential management, multi-factor authentication, and patching – has to be a priority, closely followed by the ability to detect and respond to attacks."

Lawrence Munro: "Right now, the opportunity is to focus on the fundamentals of security. When it comes to red team exercises or technical assessments, we often find that organisations fail at the basics, so re-examine your attack surface with this in mind.

"There's also an opportunity to demand more value from your suppliers and vendors – lean on the partners and advisors that you trust for support as you re-evaluate your approach to security."

2. WHAT APPROACH SHOULD CISOS AND SENIOR LEADERS PRIORITISE WHEN CONSIDERING THEIR NEXT STEPS TO SECURE ASSETS?

PV: "Processes have changed, with many employees now working remotely, and security policies may have been relaxed over the last few months to allow this to happen as quickly as possible. However, these decisions need to be kept front of mind so that you have a strategy in place to allow you to deal with this moving forward.

LM: "It's time to take stock and look back at any hastily made decisions or policy changes that have been made in the last few months. As people start to adopt a hybrid model of going into the office and working from home, it raises the risk of shadow IT, which has been more of an issue during COVID-19 as people battle with systems that were not designed for this new way of working.

"Assessing your current processes and systems can be useful in mapping out any weaknesses.

This can provide you with a snapshot of what your security looks like now, as it may well be different compared to what it was before."

3. HOW CAN CISOS AND IT SECURITY PROFESSIONALS ARTICULATE RISK IN MEASURABLE BUSINESS TERMS? HOW CAN THEY MAKE THEMSELVES HEARD?

PV: "Firstly, you need a solid evidence base. You might have data from networks and systems that you can point to, and red teaming is also a great way to map out potential routes into your organisation's systems.

"Secondly, when it comes to convincing the board, your biggest allies are non-exec directors, so it's worth making sure that they're informed about pressing threats. The UK's National Cyber Security Centre (NCSC) has a great toolkit to encourage discussions about cyber security between board members."

LM: "If you don't have evidence to back up your business case, look to partners who can help you provide this. You can also use open source projects such as MISP to help you gather threat intelligence in one place.

"If you can talk about the threat landscape in terms of how it specifically applies to your industry, this will also resonate within your board and help you to get the investment you need."

Missed our webinar? To get more in-depth advice from our experts, access the full webinar, 'How to make your cyber resilience budget work smarter', on-demand here.

Insight Space

cyber insights
programme

About NCC Group

NCC Group exists to make the world safer and more secure.

As global experts in cyber security and risk mitigation, NCC Group is trusted by over 15,000 clients worldwide to protect their most critical assets from the ever-changing threat landscape.

With the company's knowledge, experience and global footprint, it is best placed to help businesses identify, assess, mitigate and respond to the evolving cyber risks they face.

To support its mission, NCC Group continually invests in research and innovation, and is passionate about developing the next generation of cyber scientists.

With over 1,800 colleagues in 12 countries, NCC Group has a significant market presence in North America, continental Europe and the UK, and a rapidly growing footprint in Asia Pacific with offices in Australia and Singapore.

www.nccgroup.com