Cybersecurity for businesses of all sizes

A blueprint for protection

PREPARED BY GIANNIS TZIAKOURIS & JERZY 'YURI' KRAMARZ

Talos INCIDENT RESPONSE

Updated May 30, 2023

TABLE OF CONTENTS

Executive Summary 3

Security baseline for endpoints 3

Multi-factor authentication 3

Third-party defense solutions 4

Data backups 4

Device hardening 5

Email security 6

User security awareness 7

Data loss prevention 8

Security via virtualization 9

Security baseline for networks 9

Wireless security 10

Virtual private networks 11

Network vulnerability assessment and management 12

Threat intelligence 13

Security baseline for cloud 13

Azure 14

Amazon Web Services 14

Google Cloud 15

Physical security baseline 16

Security staffing baseline 17

The importance of logs in the environment 17

Conclusion 18

EXECUTIVE SUMMARY

One of the primary reasons why cybersecurity remains a complex undertaking is the increased sophistication of modern cyber threats. As the internet and digital technologies continue to advance, so do the methods and tools cybercriminals use. This means that even the most secure systems are vulnerable to attacks. Detecting and preventing these attacks require constant vigilance and adaptation. The human element of security only makes this more complex, rendering well-secured systems vulnerable to compromise due to human error or negligence, such as weak passwords or falling victim to social engineering attacks.

Developing a robust cybersecurity practice involves implementing multiple layers of security measures that are interconnected and continually monitored, including training and awareness programs to ensure that employees follow best practices. Even with best practices in place, the potential for threats from cybercriminals remains a constant concern. To assist companies in their journey for cybersecurity, this paper outlines what businesses with a burgeoning security program need for a robust foundation and logging architecture for businesses, even those with limited resources. This paper focuses on key areas of cybersecurity aiming at the holistic protection of IT environments, including their comprising endpoints, networks, cloud services and physical security. The paper also poses a set of essential questions that, when answered at the strategic level, can increase the resilience and security of a business.

The advice we outline in this paper may not be for every organization — much of this depends on budgeting, human resources and the maturity of your security operations. However, we hope businesses of all sizes can find advice in any of these sections that could help them address gaps in organizational security or potential areas of investment.

SECURITY BASELINE FOR ENDPOINTS

The multi-faceted nature of cybersecurity demands the constant monitoring and protection of entire IT environments, including their evolving assets and users. To add to the complexity, threat actors are constantly evolving their TTPs (tactics, tools, and procedures) to devise more effective attacks and avoid detection.

Despite the presented complexities, several activities can make it exponentially harder for adversaries to successfully compromise an environment, while increasing visibility by leveraging best logging practices. The presented endpoint security topics in this section are listed in order of impact while considering the ease of deployment (if two security topics have the same impact but one is easier to implement, that will be presented first) to ensure that companies with limited resources approach security in the most effective way possible.

MULTI-FACTOR AUTHENTICATION

Multi-factor authentication (MFA) is one of the tools that, when added to a cybersecurity arsenal, can go a long way, and play a significant role in thwarting several types of attacks. MFA refers to the authentication method that requires users to provide two or more verification factors to gain access to a resource. Usually, it comes in the form of a text or a pushed code from an application. MFA has proven highly effective in stopping attacks such as brute forcing, credential theft, unauthorized access, etc. The use of MFA can be used across an entire organization or to critical or sensitive resources that require enhanced authentication and authorization.

MFA configuration can be as trivial as creating MFA challenges for all log-in activity to company systems and resources or something more advanced considering the location and browser or host requesting access. If there is a need for a managed, enhanced service, solutions such as Cisco Duo can be a great option.

When looking to improve multi-factor authentication protection, along with associated processes, there are few questions that should be considered to determine next steps:

- Are all critical users and assets covered by MFA? Is MFA enabled on VPN and other remote access services and applications?
- Have all users been sufficiently trained on the MFA system and proper use to avoid phishing attempts targeting MFA information?
- Are sufficient logs recorded for each MFA authentication, covering geo-location, IP source and other relevant information? Are these logs stored in central repository or locally on the MFA device?
- Is MFA covering critical internal applications and servers such as Active Directory?
- If MFA exceptions are granted, are these exceptions recorded and the associated account activity regularly audited?

THIRD-PARTY DEFENSE SOLUTIONS

Although the use of native Windows and Unix security defense mechanisms can serve as a strong deterrent for adversaries, a more elevated security posture can be achieved with third-party security solutions in conjunction with native security mechanisms. The simplest form of a third-party defense solution that can be deployed on clients and servers is anti-virus software, such as ClamAV. Despite the wide adoption of anti-virus tools, their capabilities are limited in detecting malicious file signatures which in some cases can be rendered ineffective.

An alternative is the use of Endpoint Detection and Response (EDR) solutions that provide a layered approach to endpoint protection leveraging tools for detecting, investigating and analyzing security threats in and around endpoints with rule-based automated response capabilities. An important consideration when deploying an EDR is its proper setup and configuration based on the operating environment to ensure a high-rate detection, while not interfering with legitimate user activities. A plethora of EDR solutions exist, such as Cisco Secure Endpoint. Such tools cover log analysis, integrity checking, Windows registry monitoring, rootkit detection, time-based alerting, and active response. It supports different operating systems, including Windows, Linux, FreeBSD and OpenBSD.

When looking to improve enterprise response and detection protection, along with associated processes, there are few questions that should be considered to determine next steps:

- Is there a need for an offline or cloud based EDR? Are there any data privacy considerations?
- Are there any systems that require passive monitoring? This should especially be considered in cases of operational technology (OT) environments where any changes to systems are prohibited.
- If the deployment of EDR is offline, where is the log data stored and replicated?
- Is the configuration of the EDR well evaluated to warrant the effective detection and response of threats?
- Who are the stakeholders responsible for maintaining and regularly analyze the information and alerts produced by the EDR?

DATA BACKUPS

Data backups are a crucial aspect of system recovery. In the case where sufficient operating system images and data backups exist, re-imaging a system without the loss of data can be a trivial task. Where backups are not present, the complexity of re-imaging a system exponentially increases due to data loss considerations, which in turn significantly increases the recovery time.

Maintaining good backup practices is required to ensure effective recovery. As such, golden images for endpoints and servers should be readily available for re-imaging any malfunctioning or compromised system in a timely manner. Such snapshots can be taken every 6 to 12 months or whenever substantial changes are made to system configuration (e.g., new software or security stack used, etc.).

Data backups should be taken every 15 days, performing a complete backup is required on at least business-critical hosts (where possible, snapshots should be taken from every system) and incremental backups daily in between. Old backups should be retained for at least 90 days. All backups should be scanned for infections to avoid the re-introduction of malicious software in the IT or OT environment upon re-imaging systems.

An additional consideration for good backups is warranting the integrity of the data to thwart adversaries from tampering with them (e.g., change security configuration) and protecting any sensitive data in backups from being stolen. To achieve this, strong encryption should be enforced, such as the Advance Encryption Standard (AES). Even when backups are stored on cloud, a provider with strong encryption should be selected. If there is a need for a personal cloud solution for backups, tools such as Seafile and OwnCloud can be explored. Careful consideration should be taken for the protection of cryptographic key(s) from adversaries, ensuring that access to the encrypted backups is not lost. Administrative access to the backup storage system should be carefully controlled, audited, and protected by MFA.

Several tools can be used for backups, storing them locally or online. Software like Code42 and OneDrive can be used for data backups on both Windows and Unix. For Unix (Mac) users, the native Time Machine software can be used for local snapshots on endpoints and MacOS servers, including automated hourly, daily or monthly backups. In the case where iCloud is used, it is possible to enable iCloud synchronization to store all backup files online for a cost (Apple monthly subscription). Alternatively, for Linux file systems it is possible to use ufsdump and ufsrestore commands in a pipeline to copy a file system by writing to standard output with the ufsdump command and reading from standard input with the ufsrestore command. The resulting files will need to be stored in a backup server or big backups clusters in the cloud such as S3.

Alternatively, for more comprehensive data backup open-source tools the use of Borg and Duplicati is recommended to replicate and encrypt data for protection, as well as upload backups to several cloud systems from both Windows and Unix systems.

When looking to improve device backups there are few questions that should be considered:

- Are there tools and sufficient storage for the effective capture of data and OS backups?
- How frequent should backups be performed (both incremental and full backups)?
- Are backups tested and scanned for infections? And if yes, how often are they scanned?
- Is there a redundancy plan for backups in the case of damaged or compromised backups?
- Is there a quick process in place to access backups and re-image compromised systems?
- Are there sufficient measures taken to protect the integrity of backups (e.g., encryption)?

DEVICE HARDENING

An important aspect of cybersecurity is building secure systems from the ground up and then adding additional layers of security to control and protect other applications. A well-configured system that is properly hardened can significantly elevate an organization's security posture, forcing adversaries to search for exploitable paths that are often detected by commercial security solutions such as anti-virus or EDR. At its simplest form, hardening relies on the configuration of already available features that are often overlooked during system deployment. Several configuration guides exist for the security configuration of various systems and applications such as Apache or Windows by CIS:

- Windows
- Linux
- IBM AIX
- Oracle Solaris
- Apple macOS
- Red Hat Enterprise Linux
- Ubuntu Linux
- Web Servers (Microsoft IIS, Apache, NGIX)
- Databases (MySQL, Microsoft SQL, Oracle)
- Virtualization (VMware)

- Cloud systems
- Applications (Google Chrome, Mozilla Firefox, Microsoft Web Browser)

A major focus of hardening is ensuring that service and user accounts utilize strong passwords, administrator access is managed and secured, and that systems and residue software are patched to ensure that weaknesses are not present in the default application and operating system stack. Once basic security tasks are complete, each of the benchmarks explores specific options that can be deployed for additional security on the operating system or application layer via configuration files or manual setup. Security practitioners or IT admins should enable Credential Guard and Windows Firewall and remove unnecessary services or network. These should be part of the "golden image" that can be used at a later stage to set up additional systems. Properly setup "Golden images" can save a significant amount of time when setting up a new environment, while minimizing the attack surface on the applications and an operating system. Logs present in these hardened images can also be a valuable source of information for detecting any malicious activity. These logs should be collected, stored and hunted via manual or automated solutions that can identify common sigma rule matches indicating malicious activity or other activities that could indicate adversarial activity on a given device.

When looking to improve device hardening in businesses of all size the following questions should be considered:

- Are well-known, established security standards for hardening devices and applications followed? Is the hardening standard applied consistently across all devices and systems?
- When was the last time core system golden images were used across the enterprise, was the golden image hardened and reviewed for weaknesses? Are golden images part of security reviews? Do the available hardened images contain appropriate logs that could potentially detect a security breach?
- Are all applications, operating systems and infrastructure sufficiently patched to minimize the attack surface? How is it warranted that devices remain hardened over time, and what is the process for updating or changing the enforced hardening policies? How can device hardening processes be integrated into the existing security processes?

EMAIL SECURITY

Phishing is one of the top initial infection vectors, according to Talos IR. Adversaries regularly to bypass security controls and gain initial entry into an internal system. While there are many types of phishing – such as stealing money from targets or tricking users into opening malicious attachments – all types of phishing are generally successful. Therefore, it is important to enforce strong security practices on email servers, making it exponentially difficult for adversaries to compromise an organization.

Email security starts with good configuration, including setting up the email relay to restrict the domains or IP addresses that a mail server can relay emails from. If relay configuration is not applied correctly, adversaries can leverage the server to forward malicious emails and use the company's own reputation to distribute spam. By adding strong SMPT authentication, preferably with MFA, only known accounts can use the SMPT server to send email.

An organization can also scan links and attachments to make sure the email gateway automatically reviews content sent to users. Furthermore, the use of a domain name system blocklist or real-time blackhole list (RBL) can stop malicious activities by accessing known domains and IP addresses that have a reputation of being spam or malicious. The URIBL service is a list of domains detected as sending spam email. Where a DNSBL is a software mechanism rather than a specific list, leveraging a wide array of criteria to get an address listed or unlisted, such as email activity or their internet service providers (ISPs) that are more known to host spammers. Spamcop is a Cisco-owned service that blacklists domains found in spam messages and listed as having a poor reputation. An important aspect to note is that links or attachments from trusted senders can still be malicious as it is still possible that attackers compromise the environment of a trusted sender and use their compromised email servers to forward malicious links and binaries to trusted peers hence attachment and link scanning is an important modern feature of any email gateway.

Businesses of all sizes should enforce TLS to ensure all information sent via email is encrypted and only recipients with issued certificates can access the data to avoid data leakage. PGP/MIME and S/MIME (Secure/Multipurpose Internet Mail Extensions) are two options for encrypting emails end-to-end. S/MIME uses asymmetric cryptography and digital certificates for emails to help authenticate the email sender. These certificates have a similar function to TLS certificates. However, they can sign each email digitally and encrypt them at rest and in transit. This also prevents anyone from impersonating users since each email has a specific digital signature. Specialist programs and configurations can be used for this purpose.

Another measure for consideration when dealing with email security is the use of the Domain Name Services (DNS) protocol to appending additional records to the domain records to improve email security. DNS records such as Sender Policy Framework (SPF) can prevent spoofed sender address while DomainKeys Identified Mail (DKIM) record leverage a fingerprint hash, which validates the email so that the receiving mail server identifies the sender. Finally, Domain-based Message Authentication Reporting & Conformance (DMARC) record uses SPF and DKIM protocols for message authentication. Email service providers often check emails against all three of these records to see if they are from the place they claim to be from and have not been tampered with in transit. Most messaging systems use DNS lookups to verify the existence of the sender's email domain before accepting a message. A reverse lookup is also a viable option for identifying bogus mail senders. Once Reverse DNS Lookup is activated, the SMTP server verifies that the sender's IP address matches both the host and domain names that were submitted by the SMTP client in the EHLO/HELO command.

When looking to improve email security a few questions should be considered to determine next steps:

- Is the email server default security configuration evaluated? Are all incoming emails scanned to identify malicious attachments and links?
- Has two factor authentication been enabled for user email to ensure that identity is verified? Is there a well-defined and tested process on the provision of email accounts to users and administrators? Is the list of users and administrators frequently checked?
- Are emails monitored by a DLP solution to avoid data leakage?
- Are there any systems or checks in place to automatically check the integrity and malicious score of incoming emails by checking parameters such as DKIM and SPF? Are these checks managed by the email gateway? Are the existing domain records using DKIM, SPF and DMARC to ensure optimized security? When was the last time these records were evaluated?

USER SECURITY AWARENESS

Users are often seen as the weakest link of defense, serving as the foothold to enterprise networks and infrastructure through vishing, phishing, and credential theft. Such attacks are extremely complex to identify and stop, making them extremely effective.

To mitigate the effects of such attacks, frequent awareness campaigns are necessary to ensure that employees are well-informed on the tactics, methodologies and tools leveraged by attackers so they are better prepared and protected. Several methods exist to spread awareness, including holding cybersecurity training, disseminating hacking case studies, and performing phishing or vishing campaigns.

Security training should be short, not exceeding two hours, to hold employees' attention. This training should be simple and interactive so all employees can grasp the principles of healthy cybersecurity practices with ease and apply them in their daily use. Several security training providers exist but it is strongly recommended that all training material is created from scratch based on the needs, environment, and security goals of a corporation to warrant direct applicability. Inspiration can be drawn from FTC, National Cybersecurity Alliance and Cisco Networking Academy. At a minimum, the following topics should be visited during security trainings:

- Best practices on password creation and management such as ensuring that passwords are not reset frequently as users will simply write them down. Instead, focus on the importance of password length or secure password

storage, such as the use of password vault storage with MFA.

- The malicious effects of unauthorized software on corporate devices. A simple demonstration of the installation of malicious software on a company system and its effects should provide an effective lesson.
- Best security practices on organization networks, including the risks entailed in using unauthorized VPN solutions and cloud services that can lead to data exfiltration and compromised networks.
- An in-depth training on data classification and sharing, including appropriate cases where the sharing of sensitive data is applicable and the stakeholders eligible of initiating such requests. This should educate employees on how to easily differentiate between fraud and legitimate requests.
- Which stakeholders should be informed in an emergency or suspicious behavior? This information should be readily available to all employees and should be encouraged and rewarded for reporting incidents.

Another tool for user awareness is the dissemination of case studies of recent adversarial activity to employees, providing them with a better understanding of how adversaries operate and their potential effects. The case studies should include an overview of the malicious activity, the root cause of the attack and the impact of the attack on the environment and business. By aligning case studies with the company vertical, it is easier for employees to visualize the potential impact on the business.

Lastly, phishing and voice phishing via phone calls (vishing) campaigns should be carried out frequently (at least twice a year) to assess whether employees can identify suspicious calls or messages. IT teams or administrators should keep statistics on each of these tests to track the team's security awareness progress. This could include information on the number of employees that successfully got phished or vished, the number of employees that identified the malicious messages or calls and the number of employees that reported the incident to security personnel. All employees that fail to identify the phishing or vishing correspondence should be alerted, providing to them the campaign details, the information that was able to be "stolen" from them and the potential impact of their activity in the case of a real attack. Employees that repeatedly fall victim to phishing or vishing campaigns should be registered on a best practices cybersecurity course to help increase awareness. Phishing quizzes created by OpenDNS and PhishingQuiz can be leveraged to further spread awareness.

When looking to improve user awareness a few questions should be considered to determine next steps:

- When was the last time employees were trained in cybersecurity?
- What are the best delivery methods for security awareness for this organization?
- Is the general employee awareness level sufficient and frequently tested?
- Is sufficient information recorded for the performance of employees on cybersecurity awareness training? What do the results indicate? How can it improve security awareness where the information presented is not satisfactory?
- Are latest threats and trends included in the cybersecurity awareness training? How often is the awareness training material revised?

DATA LOSS PREVENTION

An added security should focus on detecting attempts of data exfiltration via Data Loss Prevention tools (DLP). DLP is a set of tools and processes that ensure sensitive data is not lost, misused, or accessed by unauthorized users. DLP software classifies data based on criticality and identifies violations of policies defined by organizations or within a predefined policy pack. Once such violations are identified, DLP enforces remediation with alerts, encryption and other actions preventing end users and adversaries from leaking data. DLP solutions can also provide reporting for compliance and auditing purposes.

A critical aspect of DLP solutions is the set configuration on the classification and management of data, including the tasks to be executed according to different monitored activities against different classifications of data. The accuracy of anomaly detection and the associated containment activities are completely reliant on the classification of data and the criticality assigned to different

assets. If data is misclassified, a DLP solution can be rendered ineffective.

Consider the following two examples: If a file containing sensitive data is not tagged as classified but is left with the standard tag of "public information," a DLP solution will allow for the export of the file as it is not "considered important" for the organization as per the classification provided. Contrary if an empty document with no value to the company is tagged as "Classified," in the case where a user tries to send that file to another machine or user, a DLP solution can identify this as a potential violation due to the provisioned classification.

When looking to improve data loss prevention a few questions should be considered to determine next steps:

- Are warnings from the DLP solution reviewed in real time or at least regularly? Do logs generated by the DLP stored in centralized platform for ease of access and review?
- Is there a data classification policy that is applied across all assets under protection? How is it enforced? Are email and file sharing platforms, including cloud covered by the deployed DLP solution?
- How often is our data classification policy reviewed?

SECURITY VIA VIRTUALIZATION

The use of virtualization to run services and operations is a simple defense measure that can provide a higher degree of availability and security throughout an organization. In contrast to hardware-based security, virtualized security is flexible and dynamic.

Virtualization adds an additional layer of security, including additional authentication, authorization and encryption ensuring that, even if there is an underlying physical machine that is compromised, the data enclosed in the virtualized environment is not easily accessible. From an incident response perspective, the use of virtualization allows for fast recovery time and lower recovery costs. In the case of physical compromised machines with limited backups, it can take days or even weeks to return an environment back to its original state. A security team can use virtualized environments to import a copy of a file of a virtual machine in the virtualization software to recover a compromised machine in minutes. A quick recovery process can also reduce the cost associated with recovery, as it keeps the overhead cost low in terms of work hours spent by employees restoring the environment.

However, virtual machines can also be attacked at scale, as seen in the recent ESXiArgs campaign that targeted VMWare hypervisors, encrypting dozens or hundreds of virtual machines at a time. Virtualization can be set up locally or be cloud-based. Virtualization technologies can incur a significant computational overhead, both in terms of bandwidth and system resources due to their complex nature. We recommend using high-performing systems for hosting virtualized environments, such as Docker, Virtual Box and VMware.

When looking to improve virtualization, a few questions that should be considered to determine the next steps:

- Is a cloud-based or local deployment better for the organization? Are there any privacy or data concerns for cloud-based virtualization?
- Are images of virtualized machines readily available for recovery? Are they hardened and built with specific security controls in mind? How are we managing VM backup?
- Is there a process for recovering compromised virtual machines? How often is it evaluated? Are all VM backups evaluated for malware and other malicious configurations prior to their re-introduction in the environment?

SECURITY BASELINE FOR NETWORKS

A collective response is required on the network and endpoint layer to hinder cyber-attacks, while maintaining better visibility through extensive data logging. Therefore, it is also important for businesses to also consider network security. This section presents the key areas of network security in the order of importance while considering the ease of deployment.

WIRELESS SECURITY

Wi-Fi access points are often exploited and used as entry points to organizations as they are susceptible to several types of attacks, including wireless sniffing, man-in-the-middle attacks, and brute forcing. Many actions can be taken to protect wireless infrastructure.

First, it is critical to enforce segregation between core business networks and networks intended for guests due to the higher exposure of threats that guest networks naturally bring. There are often many non-validated devices connecting to guest networks. The initial focus should be on changing the default credentials used for the administrator portal and Wi-Fi access with strong, long passwords. This will ensure that attackers cannot easily access routers by leveraging default credentials, which can be easily found online.

An added measure is to disable non-essential Wi-Fi features that can be a catalyst for an attacker. All features can be configured accordingly in a router's administrator page. Organizations may want to consider disabling remote management, Wi-Fi Protected Setup (WPS) and Universal Plug and Play (UPnP). Although such features were introduced for an elevated user experience, such features can also pose a serious security threat. For example, remote management was introduced to allow IT employees to remotely access routers and manage them, however, adversaries can also leverage this to remotely attack and compromise such devices. Another example is WPS, a physical button on a router that allows for the seamless connectivity of other devices to the network without the need for password. Such features allow potential attackers to physically bypass authentication and connect to such devices if they have physical access. Lastly, UPnP enables devices to actively scan and detect other devices in the network allowing easy connectivity, however, it also makes such systems more visible and susceptible to attackers.

An added consideration when securing Wi-Fi networks would be to hide the SSID (Service Set Identifier) of an access point/router, where applicable especially in business related networks, to protect it against attackers that would try to impersonate the Wi-fi access point by setting up another access point with a same SSID, an Evil Twin Attack. Hiding an SSID is not an advanced security measure and attackers can still bypass it by using scanning techniques.

Another significant aspect of Wi-Fi security is the protocol used for authentication and encryption. Protocols such as Wired Equivalent Privacy (WEP) and Wi-Fi Protected Access (WPA) should be actively avoided as they are obsolete. By default, the use of WPA2 should be configured where in newer routers WPA3 should be used instead. WPA3 was launched in 2018 as an improvement to WPA2 after they witnessed an attack to WPA2 called KRACK attack, where it was possible to perform key reinstallation and compromise the Wi-fi network. WPA3 introduced a new exchange protocol called Simultaneous Authentication of Equals (SAE) providing a strong encryption based on 192-bit AES encryption key for the enterprise edition. An added consideration for when selecting a Wi-Fi encryption protocol is the use of personal or enterprise configuration. WPA3 enterprise adds another layer of security compared to the personal configuration by leveraging the Protected Management Frames to improve the privacy of data packages. In addition, it allows for more ways of authentication based on the Extensible authentication protocol (EAP), including the use of digital certificates issued only to trusted entities which is highly recommended for critical environments comprising of high value assets such as servers or databases.

Wi-Fi routers and access-points should also be protected by a firewall or an intrusion detection system (IDS) such as Snort to ensure the early detection of attacks. Extensive logging should be maintained for the identification of anomalies and attacks.

When looking to improve wireless security a few questions should be considered to help organizations to determine next steps:

- Are the deployed wireless access points running the latest security firmware? Have the default configuration and credentials been changed? Did access points undergo security evaluation to ensure that vulnerabilities are identified?
- Are the latest security standards for Wi-Fi leveraged such as WPA3?
- Is client isolation enabled across all the Wireless Access Points (APs)?
- Are all access points covered by security solutions such as Intrusion Detection System (IDS) or DNS monitoring?

- Are access points used for critical business operations and public access sufficiently isolated from each other to ensure that public, guest traffic, cannot cross network boundary into corporate environment?
- Is there a process for new certificate migration if certification is used for Wi-Fi authentication? How are certificates managed, along with the public and private keys?

VIRTUAL PRIVATE NETWORKS

The use of Virtual Private Networks (VPN), such as Cisco Secure Access, has been widely embraced by organizations for the added security and enhanced management of their network they offer. The use of VPN allows the enforcement of different types of security such as filtering of incoming, outgoing network traffic, and blocking potentially malicious and anomalous connections. A complete description of VPN terminology and types of VPN setups can be found at rfc4026.

A VPN makes it more difficult for attackers to eavesdrop or modify network connections and access network resources. Creating a VPN that covers several systems can incur significant computational and bandwidth overhead. If this is a concern, organizations can restrict the use of VPN over high-value systems and sensitive communication to at least safeguard business-critical services and information.

Although VPNs provide an additional layer of security and control over a managed network, it is important to be aware that once information leaves the secured tunnel at egress points (gateway), the data is not protected anymore and is susceptible to various attacks depending on the destination of the data.

If the primary goal of the VPN is to gain anonymity and privacy, setting up a VPN environment is not sufficient. Anonymity is usually achieved using commercial VPN providers, in which case all network connections are forwarded through the servers of the VPN providers and meshed with the communications of other users and organizations. A careful selection of a VPN provider should be performed as some VPN providers keep detailed logs on their users and their activity and they share this information with third parties to assist in attribution.

When looking to improve remote access into the organization a few questions should be considered to help organizations to determine next steps:

- What systems should be accessible via a VPN? Are these systems hardened and necessary to allow employees to perform their business functions? Is the whole network open to VPN access or just specific applications and systems?
- How are VPN connections being logged? Do we capture geo-location, username, and source IP, along with host information of the systems connecting to our infrastructure?
- Can VPN connectors be used to enforce security policies, such as deploying EDR or Anti-Virus solutions?
- Is MFA used as an authentication mechanism for VPN access? Are there any accounts outside of MFA that have a single factor authentication enabled?
- What security, encryption, authorization protocols are used as part of a VPN solution? Are they strong, modern security protocols without legacy support? What type of VPN configuration is required in the existing environment?
- How many servers will be used as egress points (gateways) for VPN solution in the case where load balancing and different geographies are considered?
- How often are VPN logs reviewed for potential anomalies and successful authentication from unexpected countries?
- How is VPN performance monitored to ensure that its effectiveness over time?

NETWORK DEVICE DISCOVERABILITY

Maintaining a high degree of visibility in a network, with the ability to monitor all assets connected to a network at a given time, is necessary for good security. It allows for the detection of rogue or legacy or forgotten devices amid 'shadow IT' that could be vulnerable or unpatched and thus pose a serious threat to the environment.

A straightforward way to monitor devices connected to a network is by leveraging the "Network Map" recorded by routers, listing all the devices connected to them with both wired and wireless connections. To check the devices connected to a router, the administrator portal of an access point can be checked under the "Network Map" or "Connected Devices" to identify potential unauthorized devices in the network. Cisco Secure Network Analytics is an excellent alternative to manual device identification, allowing for large-scale passive asset discovery.

An alternative method to list the devices connected in a network is the use of an IP scanner software, such as Advanced IP Scanner. Although Active IP scanner is a good addition to the security arsenal, caution is needed due to the amount of traffic it generates to detect network devices. The deployment of such tools needs to be validated and used by authorized users to warrant an elevated level of control and security. Performing a network scan is a trivial task, as it is only required to specify the IP range for the scan and run the scanner. The provided result will return information including device hostname, assigned local IP, manufacturer, and MAC address for each of the identified devices in the network. The devices listed in a scan should be checked against an asset database to ensure they are legitimate. It is important to build a process around the network discovery activity to ensure that network and device records are updated frequently and are accurate as assets do change with time. This process can take the form of automated scans which would update the asset database to ensure that records are up to date. The asset link should also be made to specific asset owner, location, role and description of the asset so that network and security teams can easily identify the resources as they appear on the network.

When looking to improve asset discovery and management a few questions can help organizations to determine next steps:

- Is the asset inventory up to date? How frequently are asset discovery scans performed?
- What types of assets are being discovered (e.g., servers, workstations, mobile devices, IoT devices, etc.)?
- Are there any new devices appearing on the latest network scan? Are the asset owners of these devices known? Are there any blind spots or areas of the network that may not be covered by asset discovery scans?
- Is it possible to determine if any new devices are managed by an existing centralized management system such as Puppet or Windows Active Directory (if applicable)?
- Is there a well-defined incident response procedure for unknown devices appearing on the network?
- How is the effectiveness of the asset discovery process measured and evaluated over time? If no asset owner is identified, is it possible to disable the identified device(s)?

NETWORK VULNERABILITY ASSESSMENT AND MANAGEMENT

As technology is rapidly evolving, changes to devices can introduce new vulnerabilities such as Log4j. These vulnerabilities can be proactively identified with the help of network security scanners or penetration testing activities aiming to discover exploitable systems. Vulnerability assessment is a broad topic aiming to better understand what is exposed to potential adversaries. Such assessments often focus on the analysis of vulnerable applications, exposed network ports, internally exposed systems which can be present in vast and complex internal and external environments. Security teams should consider executing vulnerability assessment and penetration testing activities regularly to ensure that a network level attack surface is discovered and can be mitigated. Cisco Vulnerability Management (formerly Kenna Security) can assist with this. Depending on the findings from a penetration testing activity, such as those offered by Talos IR, corrective actions should be taken, including re-configuring firewalls, IPS solutions, hardening devices or servers, patching or simply remove vulnerable components to ensure that discovered vulnerabilities will not be exploited by attackers.

A few questions should be considered when dealing with network vulnerability assessments:

- Are the results of cybersecurity assessment or penetration test stored securely and accessible only by authorized personnel?
- How often is penetration testing performed on external, internal and application infrastructure?
- Who decides which vulnerabilities to address and when? Does that same person accept the risk for vulnerability management actions that are delayed or deferred entirely, and how is that risk acceptance process formalized?
- Who is responsible for the follow up remediation actions? How are engagements tracked along with their outputs to ensure that corrective actions are taken?
- Is there a project manager associated with security assessments who ensures the timely execution of the engagements along with remediation efforts? How are vulnerabilities being remediated identified, and what is the process for verifying that remediation is successful?
- Is the attack surface reduced by conducting regular assessments? How is the effectiveness of the vulnerability management process being measured and evaluated over time?

THREAT INTELLIGENCE

Having visibility over the latest threats and indicators of compromise (e.g., malicious hashes, IP addresses and domains) and using such information to proactively identify attack patterns can play a critical role in the security of an organization. In the case where there is no intelligence it would render the detection of attacks exponentially difficult, requiring the time-consuming practice of in-depth digital forensic analysis. As such it is imperative to be aware of the threat intelligence landscape and available platforms to ingest data from. Leveraging threat intelligence data can be treated in two ways: on an ad-hoc basis, visiting threat intelligence platforms to check specific indicators of compromise to better understand the threat; or by establishing a relation with a threat intelligence vendor and receive a continuously threat intelligence data feed and ingest this information in available security tools such as Firewalls, IPS, EDRs, Anti-virus to ensure that once such IOCs are present in an environment they will be timely detected and eradicated.

Different standards exist and used by different providers when it comes to standardizing data intelligence format and exchange. Depending on the security stack available in an environment and the standard believed to be the most comprehensive for an organization, a respective threat intelligence provider should be identified and selected to ensure best compatibility. Well known examples of threat intelligence standards are MISP, IODEF, STIX, Sigma and Yara.

Several threat intelligence organizations exist, providing customers with the latest indicators of compromise. A few examples of threat intelligence sources can come from companies such as Cisco Talos or any Cisco Secure products.

When looking to improve threat intelligence there are few questions that should be considered:

- Is there an acceptable standard identified for the exchange and ingestion of threat intelligence? How is the threat intelligence stored and used across the business? Is there a centralized database to record all ingested threat intelligence when performing investigations? How long is threat intelligence stored for?
- Are there relevant threat intelligence sources and providers identified? Is there a capability to ingest available threat intelligence data in security tools to better defend against them and make available threat intelligence part of regular security activities such as threat hunts?
- Who is the responsible stakeholder(s) for collecting, preparing, and ingesting intelligence feeds into the various security solutions? How are security employees trained to use threat intelligence effectively?
- How is the effectiveness of the existing threat intelligence program measured over time? How is threat intelligence prioritized and analyzed, focusing on the most significant threats for an organization?

SECURITY BASELINE FOR CLOUD

An increasing number of companies adopt cloud-based solutions for their operations for efficiency and cost reduction. Despite the benefits of cloud solutions, their scalability, and the ease of control over data sovereignty, the security of the platform and its various backends is generally placed on cloud providers. Companies need to ensure that they have strong visibility into applicable controls and understand the auditing applied to the various environments. Although cloud providers take appropriate security measures to protect their infrastructure and technological stack from attacks, it also requires a significant effort from users to secure the rented cloud environments, such as SaaS, PaaS and IaaS, as some aspects of security are left up to them. Cloud providers are often responsible for securing the hardware and software used for their infrastructure, storage, databases, networking and services. But customer data, identity access management, installed applications, operating system and network configuration lies with the users.

One of the focus areas should be the protection of the connection between an organization and a cloud environment. All access requests to cloud services and infrastructure from a company's environment should be closely monitored and forwarded through an encrypted network channel. All API and application gateways that point to cloud resources should be monitored and protected. This can detect cases where a company's infrastructure is compromised and adversaries are attempting to access the cloud from the compromised organization or when passively monitoring the network for data.

Another essential aspect of cloud security is access management and resource authorization. All cloud vendors provide a suite of tools for administrators to specify the groups and roles of users accessing specific resources. Administrators should ensure that all access to cloud resources is provided on an ad-hoc basis according to the needs of the business, always following the least privilege principle. In addition, all cloud user accounts should be covered by a policy enforcing the creation of strong credentials and multi-factor authentication to warrant that only authorized users can access relevant cloud resources.

In an infrastructure-as-a-service (IaaS) setting, a more in-depth security approach can be followed due to the extensive features of IaaS that compile an entire computing environment. As such, special focus on network segregation, firewall configuration and device hardening should be revisited to ensure that deployed systems do have adequate protection.

AZURE

The Azure operational security suite is the go-to place for configuring access controls and security measures. Special consideration should be given to the security roles and access controls (Azure RBAC) for users, groups and applications to set access as per business scope. In terms of data and storage protection, HTTPs can be configured for Azure file shares and enforce general storage service encryption with Azure storage service. For enhanced monitoring for the detection of anomalous access in relation to data storage, the Azure storage analytics can be used. Microsoft Defender for Cloud can be deployed on endpoint solutions in conjunction with the Azure Firewall to secure connections and system behavior. The enablement of MFA and the use of premium features such as Azure AD identity protection can be of value when securing accounts and monitoring the health of accounts, respectively.

For enhanced monitoring and visibility, the use of Azure Monitor logs, can reveal malware activity; in conjunction with the Azure Active Directory portal which can provide a holistic view of the integrity and security of an organization. A more enhanced security approach is the use of Microsoft Sentinel, which can be used for security information and event management as well as for security orchestration, automation and response, serving as the single point of monitoring and analysis of threats across an Azure environment. A complete guide on Azure security can be found here.

A minimum of the Activity logs, Network security group flow logs and Process data/security alerts in Azure should be captured and recorded for at least 12 months. A complete list of all security logs and auditing options in Azure can be found here.

AMAZON WEB SERVICES

AWS security can be treated in a comparable way to Azure. At a fundamental level, the creation of strong passwords for AWS resources and MFA authentication for users should be enforced. Furthermore, management of users, groups and roles should be carefully performed to warrant authorized, secure access to cloud resources, based on least privilege principle. It is essential that all access rights are frequently

assessed with the Access Analyzer to ensure that all access activity matches the assigned privileges and operates under the least privilege principle.

To control network traffic is possible to use network access control lists (ACLs) on the subnet level to allowlist and blocklist inbound and outbound traffic. For more extensive filtering and protection, it is possible to leverage the AWS network firewall to enforce strong blocking to malicious behavior. All network connections can be audited through the Network Access Analyzer to identify unintended network connectivity to resources in VPCs.

Another noteworthy feature of AWS is CloudTrail. It monitors and records API account activity across the AWS infrastructure, allowing for extensive control over storage, analysis, and remediation actions, making it an effective tool for incident response. Custom alerts for potential misuse of data or anomalous behavior can be set to notify administrators. In addition, VPC Flow Logs should be enforced to allow for the monitoring of IP traffic going to and from a VPC, subnet, or network interface. Lastly, the use of GuardDuty can be leveraged for the continuous monitoring of threats on AWS accounts and workloads, providing extensive visibility, detection, and remediation features.

Data can roll over and be lost or that can result in a significant monetary cost. It is recommended that logs from these features are captured and recorded for at least 12 months. It is possible to extract the logs and store them in an S3 bucket or locally to lower associated costs, where applicable.

GOOGLE CLOUD

A similar methodology, outlined above, can be followed in terms of security and logging in Google Cloud (GCP).

On the network level, the proper configuration of network security is essential. Such configuration can be performed through the Google Virtual Private Cloud (VPC) to create subnets and segregate information and user activity in a Cloud environment for added security. Cloud router is a tool capable of enforcing a dynamic exchange of routes between different VPCs. To further control network traffic is possible to leverage the GCP firewall to block malicious behavior. If DDoS and web application exploitation is a concern for internet facing applications, the use of Cloud Armor can be a good defense mechanism.

Another common issue with Google Cloud is the configuration of Cloud storage buckets that are sometimes made publicly available. IAM should be used to manage access to storage resources. The use of the data loss prevention API tool can provide added security by protecting data and blocking exfiltration attempts.

In terms of logs, VPC flow logs should be enabled to capture information about the incoming and outgoing IP traffic of network interfaces. Flow logs are updated every five seconds, providing immediate visibility. Audit logs, should also be recorded and analyzed to detect anomalous user activity, including which user did what activity in Google cloud resources. Firewall rules logging can also be a reliable source of information to detect changes to firewall rules. Lastly, special consideration should be given to data access audit logs to gain visibility in the API calls that read the configuration or metadata of resources, as well as user-driven API calls that create, modify, or read user-provided resource data. All logs should be captured and recorded for at least 12 months. Logging data in Google Cloud can incur additional costs for an organization.

When looking to improve various aspects of cloud security a few questions should be considered to determine next steps:

- Are strong credentials and MFA used throughout the cloud environment? Are the same security controls used across IaaS, SaaS, and PaaS offerings? Are all accounts created based on the least privilege principle? Are all accounts and their roles regularly audited?
- Is there a satisfactory level of visibility on the internal cloud traffic, including inbound and outbound connections from internal LANs or network ranges used by the applications? Is there a satisfactory level of visibility on user account activity and what resources users' access across cloud offerings and applications?
- Is there a possibility to set alerts for anomalous user behavior or connections? Is this an added offer in the SaaS or PaaS environment? Is the same level of detail available through IaaS offering?

- Are there sufficient data backups on the cloud data? How are backups managed, which availability zones are backups set for and how would recovery from an incident in one of the cloud environments that could impact production environment?
- Is data and activity segregation enforced (e.g., subnets, security groups, etc.)? When was the last time this security setup was reviewed?
- Is all access to cloud resources performed only from authorized, company owned systems?
- Are there technical incident response playbooks on how to access cloud backups and defend the environment? Do existing playbooks cover different regions and systems exposed throughout cloud infrastructure?

PHYSICAL SECURITY BASELINE

Although cybersecurity focuses on software-based hardening and the deployment of security tools, there is no complete security without considering physical security measures. Adversaries can leverage weak physical security to introduce malicious devices in networks, such as malicious Wi-Fi access points or deploy malicious USB drives and other monitoring devices. Circumventing physical security often originates from social engineering or insider activity. To reduce the effectiveness of such threats, the deployment of cameras, padlocks and smart gates in locations where high-value digital assets reside should be considered.

In terms of IP cameras, several options exist with advanced capabilities such as night vision, motion detection, footage recording (on cloud and offline) that can be acquired and deployed. An important aspect to remember is that digital security equipment such as cameras, gates and electronic padlocks can still be a prime target for attackers and as such they should also be properly configured and secured. This includes activities such as performing regular updates of firmware and assigning strong passwords for accessing the camera and changing those frequently (by default, several IP camera brands come with password protection off). Users should rename the default administrative account and set a new admin password. If an IoT device such as camera uses the wireless network, turn on WPA2 Encryption to ensure that adversaries cannot connect and access the video feed and ensure that they cannot place any IP cameras facing sensitive information. And if privacy is a concern, cameras should be isolated on a local network by assigning a non-routable internet IP address, such as 192.168.1[.]10.

Certain vendors set port forwarding as part of the camera software or have UPnP enabled that can potentially expose cameras to the internet. To check whether such configuration exists and whether it is possible to disable them, the official documentation of the device should be reviewed to identify the setup required for a local-only mode. If a camera cannot be configured for local operation, there are two options: Utilize the cameras in a separate VLAN with no access to the internet or block the respective IP and port on the firewall.

Padlocks and smart gates should be deployed in locations storing high-value digital assets. The access key for the padlocks should only be accessible by authorized personnel performing relevant work to the protected site. Smart gates can also be used to verify the identity of users attempting to enter a protected site. A more cost sensitive solution is a swipe card reader that will only give access to verified employees (consider the case where an employee's card is misused by an attacker or another employee) where a more advanced and expensive approach is the installment of biometrics enabled gates that leverages techniques such as fingerprint analysis or face recognition.

Security measures can serve as a deterioration for physical attacks but can still be circumvented with careful planning. Therefore, recording of pictures of users entering critical areas where motion detection enabled cameras are present should be performed. If smart gates are deployed, information such as the time of access, unique user ID, duration of access (tap in and out system) should be recorded.

When looking to improve physical security a few questions should be considered to determine next steps:

- Are there any privacy or security concerns that require a fully local deployment of cameras? If yes, what are the measures taken to ensure a local-only mode operation?

- How often are physical controls patched and checked for vulnerabilities, weak passwords, or other weaknesses? Are physical devices part of penetration testing activities?
- Are cameras and other sensors deployed in a manner where they preserve critical information, such as users typing their login passwords?
- Are sufficient logs recorded for the easy identification of users based on the times entered and exited critical sites? How are these logs stored and for how long?

SECURITY STAFFING BASELINE

Technology, system hardening and building secure networks are all a significant part of cybersecurity resilience. Businesses cannot realize these goals without the support of subject matter experts. Deploying state-of-the-art security monitoring technologies in an environment with too few security analysts can render such solutions ineffective. Any strategic cybersecurity plans without appropriate support from individuals and cross-functional teams is typically destined to fail. While Talos IR recommends that security solutions, configurations and hardening are deployed across the enterprise, equally, we recommend that long-term strategic security is built with employees as the core pillar.

Optimally, the cybersecurity team should grow at the same pace as the company's growth. The priority for companies should be on hiring security architects and security analysts in the first instance to ensure that suitable network and end-host security measures are taken as well as maintaining the capability to monitor threats and perform analysis on the available telemetry. Longer term, risk-based approach to security would require hiring incident responders, risk management officers and building a security operations center (SOC) to ensure that business operates in a secure fashion.

A significant aspect of staffing, especially in an industry such as cybersecurity, is having duplicated roles to ensure that staff are not overwhelmed due to the complex, demanding and time-consuming nature of the work, especially in cases of emergencies. Some cybersecurity emergencies can last weeks, and staff are expected to work around the clock to protect the environment. Another concern is that, due to the critical nature of cybersecurity, if a key employee takes leave, an entire environment can be left unprotected if no other employee can cover the workload. In the case where a more flexible approach to staffing is required, outsourcing cybersecurity and incident detection and response to vendors should be considered. If outsourcing is preferred, it is still necessary to ensure that the onsite IT staff are still well educated and trained to work closely with security, IR and Managed Detection and Response (MDR) service providers for performing any required activities locally when requested.

When looking to improve internal staffing, a few questions should be considered to determine next steps:

- What cybersecurity staffing model is required? Is there a need for 24/7 coverage? Is there sufficient redundancy in terms of staff ensuring smooth operations in the case of potential downtime (vacation, illness, etc.) of key staff?
- Is there a need for outsourcing some functions? Is this done for redundancy (an additional layer of security) or for ease of operation?
- How often is the staff trained? Are the staff sent to conferences and training courses to ensure that they understand the latest threats? Does the existing staff hold sufficient certifications?

THE IMPORTANCE OF LOGS IN THE ENVIRONMENT

Logs can be of critical importance when dealing with security emergencies where there is a need for enhanced visibility for the identification of a root cause or effects of an attack. Logs can also help to discover a variety of network and host issues that might disclose early signs of malicious activity in the environment. Good logging, which does not over- or under-log notable events, is an exceedingly challenging task and requires a detailed understanding of the environment that businesses are trying to protect. The list below contains

the most important **endpoint artefacts** that should be recorded as part of a layered security approach, allowing for data correlation between the network and endpoint layer. Talos IR recommends that all artefacts are recorded for at least three months, but 12 months is preferable.

- All authentication requests to computing resources (of at least high-value systems) should be recorded and at a minimum, including timestamp, username, hostnames, result of authentication request (success, failure), IP and location of the requester.
- All EDR and anti-virus notifications, including associated activity and other monitored telemetry in EDR solutions.
- Windows security, application and system event logs covering at least critical systems.
- Scheduled tasks, services, and processes, including timestamps and associated users.
- PowerShell and Terminal executions, including timestamps and associated users.
- For Linux systems, a copy of entries from folders such as /var/log directory should be recorded.
- File hashes of the binaries present in the environment.
- Email headers and a hash of their attached files should be recorded to easily identify phishing attempts.
- Warnings and messages from DLP solutions where applicable, including the timestamp of the event, originating host, user and IP address, files, and their associated hashes in question.

Similarly, to the above, a list of the most important **networking artefacts** is provided below. It is once more recommended that all artefacts are recorded for at least three, but preferably 12 months.

- Wi-Fi active sessions, including MAC address of connected device, username (if applicable), hostname, date and time of connection and duration of connection.
- NetFlow activity for various boundary and internal devices to be able to easily identify network traffic between devices.
- Web proxy logs, including source and destination IP addresses, requested URLs, timestamp and user requesting the data.
- Network Firewalls, IDS and IPS events including allowed and blocked attempts.
- Application firewalls events such as filters, guards, XML gateways, database firewalls, web application firewalls (WAFs).
- VPN access logs, including timestamps, usernames, hostnames, source client IP, authentication type.
- Changes in network topology, including a network map depicting the changes to ensure the easy attribution of activity to respective systems.

CONCLUSION

This paper serves as an outline of the key areas and methodologies used for establishing a strong cybersecurity baseline and high visibility through hardening, logging, and best security practices. The paper visited key security areas in endpoint, network, cloud and physical security to allow companies to select appropriate methods to increase their cybersecurity resilience in a holistic manner. Where applicable, specific guiding questions were presented which, when answered, would provide readers with a potential path to improving the people, process, and technology aspects of their defense. There might be a combination of key factors that eliminate a part of the advice provided in this paper, as some businesses will operate in vastly different environments. For example, some businesses might have almost zero cloud footprint, in which case the endpoint, network and physical aspect of security is only relevant. Whereas other companies might have a substantial cloud investment to support their operation through various SaaS or PaaS applications, and their focus should be on cloud, endpoint and staffing sections of this paper. Whatever the case may be, we hope that the advice and questions posed in this paper will allow our readers to reflect on their current cybersecurity countermeasures and provide them with guidance on potential future improvements.

