MANDIANT®
NOW PART OF Google Cloud

Cyber Snapshot Report
Issue 4

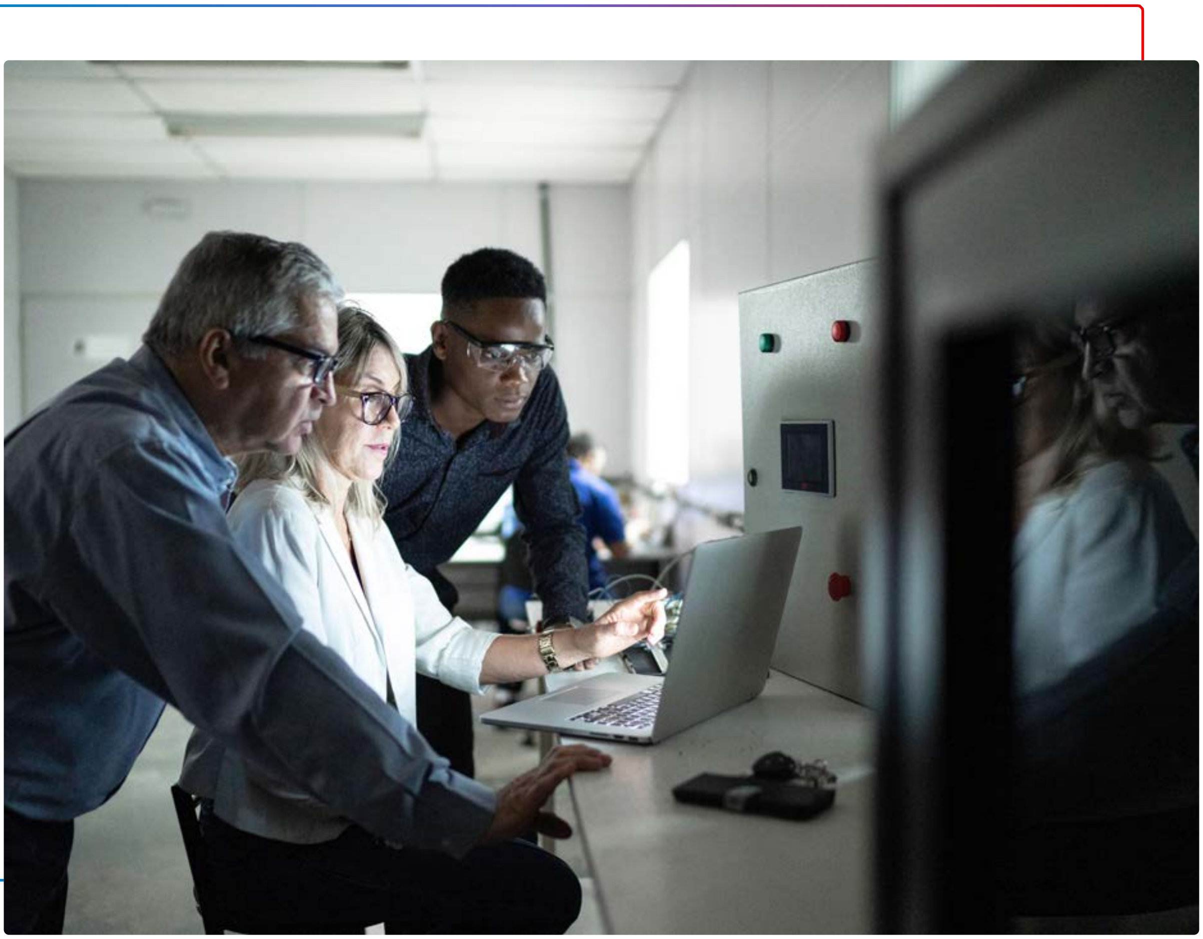

The Defender's Advantage Cyber Snapshot provides insights into cyber defense topics of growing importance based on Mandiant frontline observations and real-world experiences. This issue covers a wide range of topics, including building security into AI systems, best practices for effective crisis communications during an incident, and mitigating the latest risks to IoT and edge network infrastructure.

Securing Artificial Intelligence Systems

Artificial intelligence (AI) is advancing rapidly, and it's important that effective risk management strategies evolve along with it. Mandiant believes that it is important to build security into AI systems from the start to avoid the bolt-on security solutions we have seen plague networks and DevOps. To help achieve this, Google recently introduced the Secure AI Framework (SAIF), a conceptual framework for secure AI systems.

SAIF offers a practical approach to address top of mind concerns including security (e.g, access management, network and endpoint security, supply chain attacks, etc.), AI/ML model risk management (e.g., model transparency and accountability, data poisoning, data lineage, etc.), privacy and compliance (e.g., data privacy and usage of sensitive data), and people and organizations (e.g., talent gap, governance and Board reporting).

There are six core elements of SAIF that collectively guide organizations to build and deploy AI systems in a secure and responsible manner.

- Expand strong security foundations to the AI ecosystem
- Extend detection and response to bring AI into an organization's threat universe
- Automate defenses to keep pace with existing and new threats
- Harmonize platform level controls to ensure consistent security across the organization
- Adapt controls to adjust mitigations and create faster feedback loops for AI deployment
- Contextualize AI system risks in surrounding business processes

Expand strong security foundations to the AI ecosystem

- **Review what existing security controls across the security domains apply to AI systems**

  Existing security controls across the security domains apply to AI systems in a number of ways. For example, data security controls can be used to protect the data that AI systems use to train and operate; application security controls can be used to protect the software that AI systems are implemented in; infrastructure security controls can be used to protect the underlying infrastructure that AI systems rely on; and operational security controls can be used to ensure that AI systems are operated in a secure manner.

  The specific controls that are needed will vary depending on the use of AI, as well as the specific AI systems and environments.

- **Evaluate the relevance of traditional controls to AI threats and risks using available frameworks**

  Traditional security controls can be relevant to AI threats and risks, but they may need to be adapted to be effective, or additional layers added to the defense posture to help cover the AI specific risks. For example, data encryption can help to protect AI systems from unauthorized access by limiting the access of the keys to certain roles, but it may also need to be used to protect AI models and their underpinning data from being stolen or tampered with.

- **Perform an analysis to determine what security controls need to be added due to AI specific threats, regulations, etc.**

  Using the assembled team, review how your current controls map to your AI use case, do a fit for purpose evaluation of these controls and then create a plan to address the gap areas. Once all of that is done, also measure the effectiveness of these controls based on whether they lower the risk and how well they address your intended AI usage.

- **Prepare to store and track supply chain assets, code and training data**

  Organizations that use AI systems must prepare to store and track supply chain assets, code, and training data. This includes identifying, categorizing, and securing all assets, as well as monitoring for unauthorized access or use. By taking these steps, organizations can help protect their AI systems from attack.

- **Ensure your data governance and lifecycle management are scalable and adapted to AI**

  Depending on the definition of data governance you follow, there are up to six decision domains for data governance:

  - Data quality
  - Data security
  - Data architecture
  - Metadata
  - Data lifecycle
  - Data storage

  AI data governance will become more important than ever. For example, a key underpinning of the effectiveness of AI models are the training sets of data. Ensure that you have a proper lifecycle management system when it comes to data sets, with a strong emphasis on security as part of the lifecycle (i.e. have security measures from creation of data to the ultimate destruction of data embedded throughout the lifecycle). Data lineage will also play a key part and help to answer questions with regards to privacy and intellectual property. If you know who created the data, where it came from, and what makes up the dataset, it is much easier to answer questions on the aforementioned topics.

  As AI adoption grows, your organization's success will likely hinge on scaling these decision domains in an agile manner. To help support this effort, it is critical to review your data governance strategy with a cross functional team and potentially adjust it to ensure it reflects advances in AI.

- **Retain and retrain**

  We are not talking about AI, but rather people. For many organizations, finding the right talent in security, privacy and compliance can be a multi-year journey. Taking steps to retain this talent can add to your success, as they can be retrained with skills relevant to AI quicker than hiring talent externally that may have the specific AI knowledge, but lack the institutional knowledge that can take longer to acquire.

Extend detection and response to bring AI into an organization's threat universe

- **Develop understanding of threats that matter for AI usage scenarios, the types of AI used, etc.**

  Organizations that use AI systems must understand the threats relevant to their specific AI usage scenarios. This includes understanding the types of AI they use, the data they use to train AI systems, and the potential consequences of a security breach. By taking these steps, organizations can help protect their AI systems from attack.

- **Prepare to respond to attacks against AI and also to issues raised by AI output**

  Organizations that use AI systems must have a plan for detecting and responding to security incidents, and mitigate the risks of AI systems making harmful or biased decisions. By taking these steps, organizations can help protect their AI systems and users from harm.

- **Specifically, for Gen AI, focus on AI output - prepare to enforce content safety policies**

  Gen AI is a powerful tool for creating a variety of content, from text to images to videos. However, this power also comes with the potential for abuse. For example, Gen AI could be used to create harmful content, such as hate speech or violent images. To mitigate these risks, it is important to prepare to enforce content safety policies.

- **Adjust your abuse policy and incident response processes to AI-specific incident types, such as malicious content creation or AI privacy violations**

  As AI systems become more complex and pervasive, it is important to adjust your abuse policy to deal with use cases of abuse and then also adjust your incident response processes to account for AI-specific incident types. These types of incidents can include malicious content creation, AI privacy violations, AI bias and general abuse of the system.

Automate defenses to keep pace with existing and new threats

- **Identify the list of AI security capabilities focused on securing AI systems, training data pipelines, etc.**

  AI security technologies can protect AI systems from a variety of threats, including data breaches, malicious content creation, and AI bias. Some of these technologies include traditional data encryption, access control, auditing which can be augmented with AI and newer technologies that can perform training data protection, and model protection.

- **Use AI defenses to counter AI threats, but keep humans in the loop for decisions when necessary**

  AI can be used to detect and respond to AI threats, such as data breaches, malicious content creation, and AI bias. However, humans must remain in the loop for important decisions, such as determining what constitutes a threat and how to respond to it. This is because AI systems can be biased or make mistakes, and human oversight is necessary to ensure that AI systems are used ethically and responsibly.

- **Use AI to automate time consuming tasks, reduce toil, and speed up defensive mechanisms**

  Although it seems like a more simplistic point in light of the uses for AI, using AI to speed up time consuming tasks will ultimately lead to faster outcomes. For example, it can be time consuming to reverse engineer a malware binary. However, AI can quickly review the relevant code and provide an analyst with actionable information. Using this information, the analyst could then ask the system to generate a YARA rule looking for these actions. In this example, there is an immediate reduction of toil and faster output for the defensive posture.

Harmonize platform level controls to ensure consistent security across the organization

- **Review usage of AI and life cycle of AI based apps**

  As mentioned in Step 1, understanding the use of AI is a key component. Once AI becomes more widely used in your organization, you should implement a process for periodic review of usage to identify and mitigate security risks. This includes reviewing the types of AI models and applications being used, the data used to train and run AI models, the security measures in place to protect AI models and applications, the procedures for monitoring and responding to AI security incidents, and AI security risk awareness and training for all employees.

- **Prevent fragmentation of controls by trying to standardize on tooling and frameworks**

  With the aforementioned process in place, you can better understand the existing tooling, security controls, and frameworks currently in place. At the same time, it is important to examine whether your organization has different or overlapping frameworks for security and compliance controls to help reduce fragmentation. Fragmentation will increase complexity and create significant overlap, increasing costs and inefficiencies. By harmonizing your frameworks and controls, and understanding their applicability to your AI usage context, you will limit fragmentation and provide a 'right fit' approach to controls to mitigate risk. This guidance primarily refers to existing control frameworks and standards, but the same principle (e.g. try to keep the overall number as small as possible) would apply to new and emerging frameworks and standards for AI.

Adapt controls to adjust mitigations and create faster feedback loops for AI deployment

- **Conduct Red Team exercises to improve safety and security for AI-powered products and capabilities**

  Red Team exercises are a security testing method where a team of ethical hackers attempts to exploit vulnerabilities in an organization's systems and applications. This can help organizations identify and mitigate security risks in their AI systems before they can be exploited by malicious actors.

- **Stay on top of novel attacks including prompt injection, data poisoning and evasion attacks**

  These attacks can exploit vulnerabilities in AI systems to cause harm, such as leaking sensitive data, making incorrect predictions, or disrupting operations. By staying up-to-date on the latest attack methods, organizations can take steps to mitigate these risks.

- **Apply machine learning techniques to improve detection accuracy and speed**

  Although it is critical to focus on securing the use of AI, AI can also help organizations achieve better security outcomes at scale. AI-assisted detection and response capabilities, for example, can be an important asset for any organization. At the same time, it is essential to keep humans in the loop to oversee relevant AI systems, processes, and decisions.
  Over time, this effort can drive continuous learning to improve AI base protections, update training and fine-tuning of data sets for foundation models, and the ML models used for building protections. In turn, this will enable organizations to strategically respond to attacks as the threat environment evolves. Continuous learning is also critical for improving accuracy, reducing latency and increasing efficiency of protections.

- **Create a feedback loop**

  To maximize the impact of the previous three elements, it is critical to create a feedback loop. For example, if your Red Team discovers a way to misuse your AI system, that information should be fed back into your organization to help improve defenses, rather than focusing solely on remediation. Similarly, if your organization discovers a new attack vector, it should be fed back into your training data set as part of continuous learning. To ensure that feedback is put to good use, it is important to consider various ingestion avenues and have a good understanding of how quickly feedback can be incorporated into your protections.

Contextualize AI system risks in surrounding business processes

- **Establish a model risk management framework and build a team that understands AI-related risks**

  Organizations should develop a process for identifying, assessing, and mitigating the risks associated with AI models. The team should be composed of experts in AI, security, and risk management.

- **Build an inventory of AI models and their risk profile based on the specific use cases and shared responsibility when leveraging third-party solutions and services**

  Organizations should build a comprehensive inventory of AI models and assess their risk profile based on the specific use cases, data sensitivity, and shared responsibility when leveraging third-party solutions and services. This means identifying all AI models in use, understanding the specific risks associated with each model, and implementing security controls to mitigate those risks along with having clear roles and responsibilities.

- **Implement data privacy, cyber risk, and third-party risk policies, protocols and controls throughout the ML model lifecycle to guide the model development, implementation, monitoring, and validation**

  Organizations should implement data privacy, cyber risk, and third-party risk policies, protocols and controls throughout the ML model lifecycle to guide the model development, implementation, monitoring, and validation. This means developing and implementing policies, protocols, and controls that address the specific risks associated with each stage of the ML model lifecycle. Keep the fourth element of the framework in mind to ensure you do not create undue fragmentation.

- **Perform a risk assessment that considers organizational use of AI**

  Organizations should identify and assess the risks associated with the use of AI, and implement security controls to mitigate those risks. Organizations should also cover security practices to monitor and validate control effectiveness, including model output explainability and monitoring for drift. As referenced in the first two elements, it is important to create a cross functional team and build a deeper understanding of the relevant use cases to support this effort. Organizations can use existing frameworks for risk assessment to help guide their work, but will likely need to augment or adapt their approach to address new emerging AI risk management frameworks.

- **Incorporate the shared responsibility for securing AI depending on who develops AI systems, deploys models developed by model provider, tunes the models or uses off-the-shelf solutions**

  The security of AI systems is a shared responsibility between the developers, deployers, and users of those systems. The specific responsibilities of each party will vary depending on their role in the development and deployment of the AI system. For example, the AI system developers are responsible for developing AI systems that are secure by design. This includes using secure coding practices, training AI models on clean data, and implementing security controls to protect AI systems from attack.

- **Match the AI use cases to risk tolerances**

  This means understanding the specific risks associated with each AI use case and implementing security measures to mitigate those risks. For example, AI systems that are used to help make decisions that could significantly impact people's lives, such as healthcare or finance, will likely need to be more heavily secured than AI systems that are used for less urgent tasks, such as marketing or customer service.

In Conclusion

AI has captured the world's imagination and many organizations are seeing opportunities to boost creativity and improve productivity by leveraging this emerging technology. SAIF is designed to help raise the security bar and reduce overall risk when developing and deploying AI systems.

4 Phases of Cybersecurity Crisis Communications

Communicating during a crisis is tough for even the savviest and most well-prepared organizations. The unique attributes of a cyber attack, and the increasing use of the public domain by threat actors, mean how a victim organization communicates with their stakeholders during an incident can impact their brand long after the technical remediation is wrapped up.

To further complicate the response, the process of managing communications can compete for the time and attention of the crisis responders and executive leaders as the organization works to quickly restore business operations and remediate networks during a cyber incident. Moreover, there is often confusion on the scope of Cybersecurity Crisis Communications. While media relations is a very visible part of the communication strategic response, it is only one audience. In practice, organizations should develop a comprehensive communications strategy that informs all of its internal and external stakeholders.

To avoid communication missteps, particularly at a stressful time when every second counts, it helps to have seasoned cybersecurity crisis communications experts providing advice and expertise to help organizations and their governing boards respond appropriately. As the threat landscape evolves and threat actors incorporate new techniques, Mandiant now offers cybersecurity crisis communications specialists alongside its incident response team to help customers navigate incidents, evaluate stakeholder engagement, and strategize for the associated cascading communications.

What is Cybersecurity Crisis Communications?

Cybersecurity-specific Crisis Communications is a combination of incident response and crisis management operations, where tailored messaging is developed for a variety of stakeholders and channels, with intricately timed delivery. During a crisis, the communications strategy should consider several factors beyond impact to business operations, risk appetite, and the potential for brand or reputation damage. For example, threat actor behavior and intelligence trends are important considerations in deciding how, what, and when to communicate. Sometimes a “strategic non-response” is the best response as certain messaging may tip off the threat actor, causing them to change their tactics, techniques, and procedures.

Trust and brand resilience are notably tested during a cyber incident - and the middle of an incident is not the time to start building trust. Rather, an organization's approach to crisis communications and failure to provide information and be transparent can further erode trust. Communication missteps compound the overall loss and impact to business operations. Therefore, it is imperative to have a strong understanding of what crisis communications is, best practices for response during times of crisis, and how to prepare and plan for game day.

What leads to the best Cybersecurity Crisis Communications response?

From Mandiant's experience, success starts well before day one of a breach or incident. Rather, it is a continual cycle of review, analysis, and refinement of the organization's incident response and business continuity plans. With specific attention to crisis communications, we will share lessons learned from our specialists' first-hand experience addressing cybersecurity crisis communications planning. The cycle of these activities is grouped into four phases — strategic readiness, assurance, response, and post-incident review.

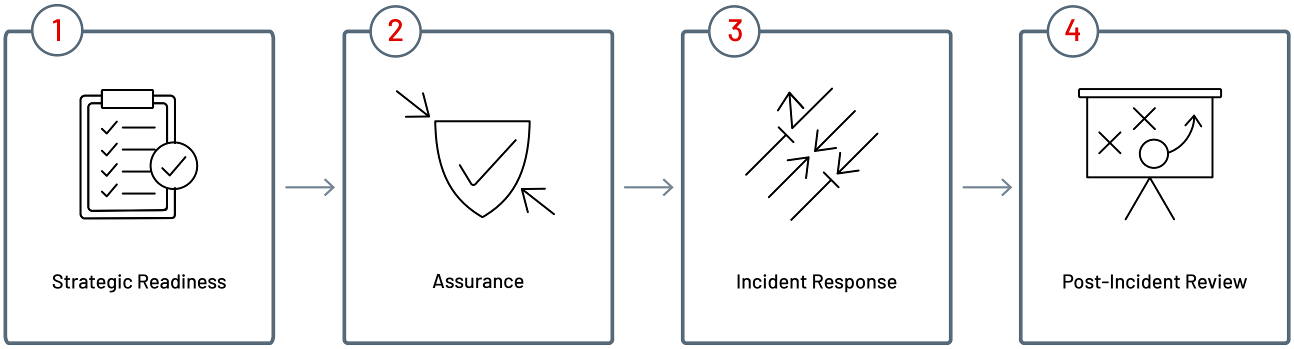

FIGURE 1: Cybersecurity Crisis Communications Phases

Phase 1: Strategic Readiness

First in the cycle is the pre-breach "Strategic Readiness" phase or simply stated, the planning phase. This phase is a foundational and essential activity for all organizations, regardless of size, sector, or location. The approach should be customized to the organization, providing a written and repeatable plan with clearly defined roles and responsibilities, a governance structure with formal decision authority levels, and a framework for response. Like many athletic coaches, for responders it is our playbook and is based on potential activities. This should also be thought of and serve as a living breathing document that is regularly reviewed and shared with those individuals that will be part of the response team during an incident.

It is important to have the right team in the room with clearly defined roles and responsibilities. This team should include representation from across the organization (including HR, Procurement, Communications, Legal, Logistics, and Operations to name a few). You can't anticipate what you'll need, especially when it comes to provisioning hardware, getting out cascading communications, and conducting insightful data impact assessments. The team should also implement a governance and management model, with specific working groups aligned to functional responsibilities. One of the deliverables developed during the planning phase is a Crisis Communications annex to the Incident Response Playbook. This playbook should be specific to the organization and include sections on incident and crisis response, key messaging based on hypothetical scenarios, and stakeholder identification and channel mapping. One additional consideration is the importance of having alternative communication mechanisms, commonly referred to as "out of band" communications, in the event your primary way of communication is compromised. In data breach and cybersecurity incidents, a threat actor may have persistence in the network, requiring leaders and responders to use these alternative communication methods.

Phase 2: Assurance

The second phase, also part of the proactive and pre-breach response, is the "Assurance" or exercise phase. During this phase, companies should exercise their team's response based on real-world attacks and scenarios. Some states are even moving to mandate this as part of the board response[1]. It certainly helps to bring in well-trained specialists to develop and facilitate exercises based on tailored, realistic scenarios. During these exercises, the team can practice their plan, test their playbook, and identify gaps for remediation. The team members also develop muscle memory in a safe and less stressful environment. Come game time, the consequence of a mistake is more significant and more likely under the higher-pressure situation. It is much easier to stay calm and respond clear-headed when you can anticipate what is next in your expected delivery and execution.

It is also imperative as part of the assurance phase for teams to be receptive to advice and feedback. This phase should also be a recurring activity, and not a "check the box" exercise, with individuals from across the organization, in various job roles and levels, well beyond the executive leadership team and the board. Lastly, the exercise should include the "reinforcement" or surge team, and this should be a deep bench of talent. Response team planning should account for sustained efforts covering at least the first 30 days, with shifts of personnel. The initial response will likely require 24/7 coverage — and to prevent burnout and exhaustion—it helps to have a ready relief roster trained and set for response. It is important to ensure your organization has a communications annex or section in the organization's Business Continuity Plan, the Disaster Recovery Plan, and the Incident Response Plan.

[1] New York State Department of Financial Services, DFS SUPERINTENDENT ADRIENNE A. HARRIS ANNOUNCES UPDATED CYBERSECURITY REGULATION, November 2022

Phase 3: Incident Response

The third phase is the reactive phase of “Incident Response.” Response execution will be defined by the priority and attention you put into the first two phases. The adage that is you should spend 80% of your time planning, and 20% on execution is certainly true. When the day comes, it is imperative that companies are able to quickly spin up their teams for response. They will know their roles and responsibilities and have a working governance structure to respond. They will be able to organize the requisite information exchange sessions and track the action items and tasks. They will have already mapped their stakeholders and communication channels and be able to quickly assess channel readiness.

The smoothest and most-effective responders are usually those who are well-trained, well-equipped, and have pre-staged the requisite tools ahead of time. They respond with dignity, respect for the team, and consideration for pace — recognizing that it is a marathon not a sprint. They closely collaborate and share information prior to making decisions, but they also don’t get into analysis paralysis. Organizations that fail or stumble are typically those not open to advice or feedback, don’t recognize their performance failures, or are poorly organized and coordinated in their response and communications.

Phase 4: Post-Incident Review

Managing a breach is hard, both from an emotional and an operational standpoint and many people never want to talk about the incident again. However, as difficult as it may be, it's important to move to the final phase, the Post-Mortem Assessment. This phase starts just as the dust settles — the investigation is complete, the remediation activities restored business operations, and notifications have been made to regulators or victims. Some may also call this the "After Action" or "Lessons Learned" phase and second to planning, it is one of the most important phases to be thorough. Specialists can work alongside clients to identify gaps and solutions to mitigate the impact of future incidents.

Some of the best practices garnered from Mandiant's client cases surfaced during Post-Mortem Assessments. Each incident and each response is different – some have false starts and recover well; others are shining examples of industry best practices. What is important is to share lessons learned for the benefit of others. As Winston Churchill famously said, "those that fail to learn from history are doomed to repeat it."

Business IoT Targeted by Espionage Groups

The number of active Internet of Things (IoT) connected devices is expected to reach nearly 42 billion in 2023[2], helping to accelerate innovation and automation across sectors from smart manufacturing, retail inventory management, digital payments, and physical security and surveillance. As with nearly every technology advancement, cyber risk is a side effect every business must expect.

In the past, Mandiant has observed IoT devices, smart devices, and routers compromised and used to create botnets to perpetrate large scale financially motivated cyber crime operations. A botnet is a network of compromised devices that a threat actor can use to conduct a variety of threat activity, such as distributed-denial-of-service (DDoS) attacks and malware distribution. However, Mandiant assesses with moderate confidence that state-linked espionage groups have also leveraged botnets for multiple purposes[3]. This attacker behavior underscores the opportunity large scale adoption of IoT and smart devices presents for state-linked threat actors looking to acquire strategic intelligence and intellectual property from global businesses.

Organizations looking to continue their digital transformation, accelerate automation, recover lost value chains after the economic impacts of the COVID-19 pandemic, or leverage the rollout of 5G connectivity networks[4] are encouraged to work closely with their cyber security teams to ensure a comprehensive cyber defense plan is in place to help protect the organization.

IoT Device, Smart Device, and Router Botnets Useful for Obfuscating Activity

Mandiant assesses that state-linked espionage groups use botnets consisting of IoT, smart devices, and routers to obfuscate malicious activity, based on multiple campaign observations from Mandiant and other private and public sector security researchers. Reported instances of compromised device botnet use by espionage groups include the following.

- In April 2022 Mandiant reported[5] on a campaign by APT29 using a botnet of IoT cameras as part of command and control (C2) activities using the QUIETEXIT malware (Figure 2). The domains used in this C2 activity appeared designed to blend in with legitimate traffic from the infected IoT devices, apparently to hide the activity from anyone reviewing logs.

[2] Frost and Sullivan, Internet of Things (IoT) Predictions Outlook, November 2022
[3] Mandiant, Espionage Actors Lurk in Compromised Device Botnets, April 2023
[4] Frost and Sullivan, The Top Growth Opportunities for IoT in 2023, March 2023
[5] Mandiant, https://www.mandiant.com/resources/blog/unc3524-eye-spy-email

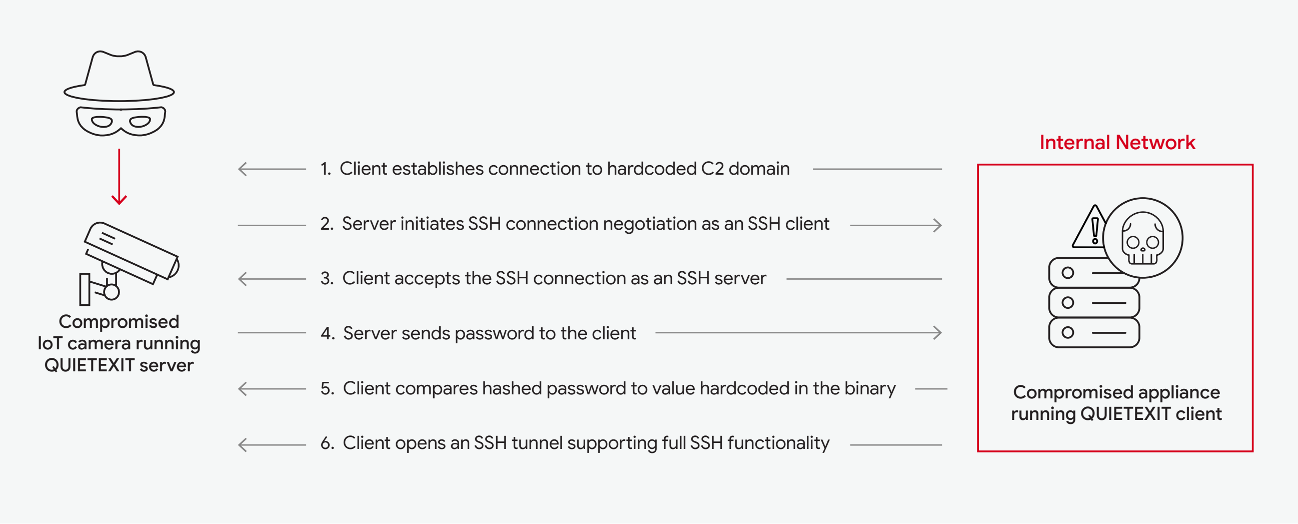
Compromised IoT camera running QUIETEXIT server
1. Client establishes connection to hardcoded C2 domain
2. Server initiates SSH connection negotiation as an SSH client
3. Client accepts the SSH connection as an SSH server
4. Server sends password to the client
5. Client compares hashed password to value hardcoded in the binary
6. Client opens an SSH tunnel supporting full SSH functionality
Internal Network
Compromised appliance running QUIETEXIT client

FIGURE 2: How QUIETEXIT works with IoT devices

- A 2021 report[6] from France's Agence nationale de la sécurité des systèmes d'information (ANSSI, French National Agency for the Security of Information Systems) detailed a campaign linked to the Chinese group APT31 that reportedly used a botnet of routers and possibly other small office and home office devices to obfuscate activities within targeted networks.
- In 2022 PricewaterhouseCoopers reported[7] on malware observed during an engagement that they named "BPFDoor," which Mandiant has linked to APT41. In the reported campaign, the malware allegedly received commands from virtual private servers (VPS) that were controlled by a network of Taiwan-based compromised routers.
- Chinese security firm Antiy reported[8] in 2022 that it had observed a large network of compromised IoT devices and Linux devices routing traffic between C2 servers and Torii malware. According to the firm, they were able to attribute the activity to OceanLotus, referred to by Mandiant as APT32; however, Mandiant has not confirmed this attribution.
- In 2018 researchers publicly reported[9] use of VPNFILTER malware in campaigns targeting networking devices and network-attached storage (NAS) devices globally, with a heavy concentration of devices in Ukraine. Some samples reportedly integrated adversary-in-the-middle (aitm) and destructive capabilities, but it is possible that these modules were intended for other purposes. Mandiant believes this use of VPNFILTER is consistent with Russian-sponsored cyber espionage activity.

[6] https://www.cert.ssi.gouv.fr/cti/CERTFR-2021-CTI-013/
[7] https://www.pwc.com/gx/en/issues/cybersecurity/cyber-threat-intelligence/cyber-year-in-retrospect/yir-cyber-threats-report-download.pdf
[8] https://mp.weixin.qq.com/s/2RIuW4O56UWiNSQB2hQtGA
[9] https://blog.talosintelligence.com/vpnfilter/
[10] https://thehackernews.com/2018/06/vpnfilter-router-malware.html

Public reporting and Mandiant observations indicate that some actors have compromised or used existing botnets created by other threat actors. Mandiant suspects that this tactic is useful for espionage actors in very limited circumstances and will therefore not significantly increase in usage in the future.

- In September 2022 Mandiant identified[11] a campaign by UNC4210, which is suspected to be linked to Turla Team, in which the actors hijacked at least three C2 domains associated with an ANDROMEDA malware botnet. The version of ANDROMEDA associated with the botnet was first uploaded to VirusTotal in 2013 and spread from infected USB keys. After re-registering the expired C2 domains, Turla was seemingly able to use the remaining infections that contacted the servers to profile and select victims (Figure 3).

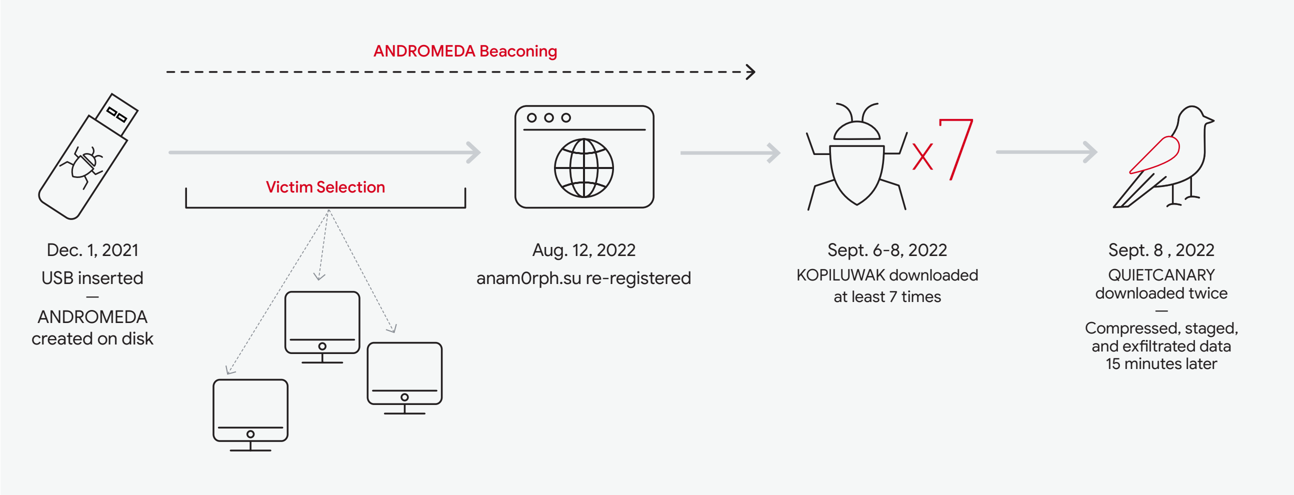

FIGURE 3: Timeline of ANDROMEDA to Turla Intrusion

[11] https://www.mandiant.com/resources/blog/turla-galaxy-opportunity

Securing IoT Devices

IoT and smart devices are often not designed to be secure and at times have hard-coded credentials and/or are difficult or impossible to patch when software vulnerabilities are discovered. Organizations actively deploying these devices, or including IoT in digital transformation plans, should ensure that they are able to be properly secured and regularly checked for suspicious activity. Figure 4 outlines security risks related to IoT device manufacturing and operation that asset owners should consider alongside plans to deploy these devices.

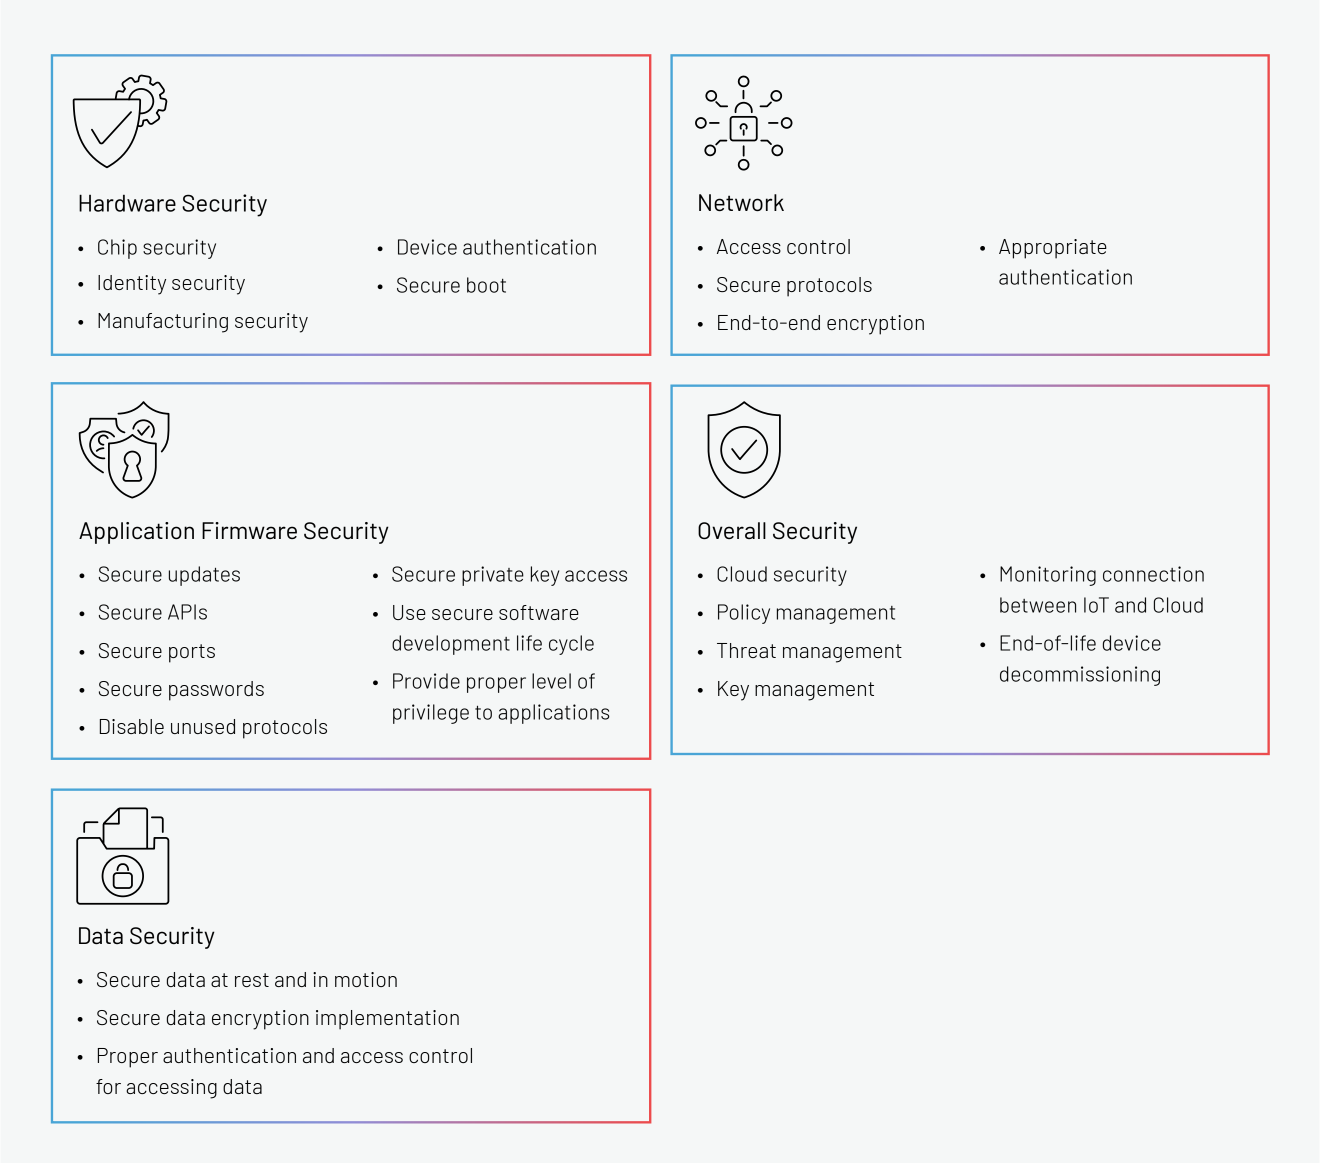

Hardware Security

- Chip security
- Identity security
- Manufacturing security
- Device authentication
- Secure boot

Network

- Access control
- Secure protocols
- End-to-end encryption
- Appropriate authentication

Application Firmware Security

- Secure updates
- Secure APIs
- Secure ports
- Secure passwords
- Disable unused protocols
- Secure private key access
- Use secure software development life cycle
- Provide proper level of privilege to applications

Overall Security

- Cloud security
- Policy management
- Threat management
- Key management
- Monitoring connection between IoT and Cloud
- End-of-life device decommissioning

Data Security

- Secure data at rest and in motion
- Secure data encryption implementation
- Proper authentication and access control for accessing data

FIGURE 4: Considerations for securing IoT devices

What this means for organizations within a digital transformation

Mandiant anticipates that cyber espionage actors will continue to use this tactic because it provides the attackers with an effective tactical advantage for a relatively low investment of time and resources, as IoT and smart devices are often poorly secured and continue to proliferate. Mandiant also speculates that as IoT and smart devices continue to grow in popularity and tools specifically targeting these devices become more available in underground markets and freely online, espionage actors may show increased interest in use of botnets as a means of disguising intelligence gathering activity as benign or opportunistic, financially motivated cyber crime.

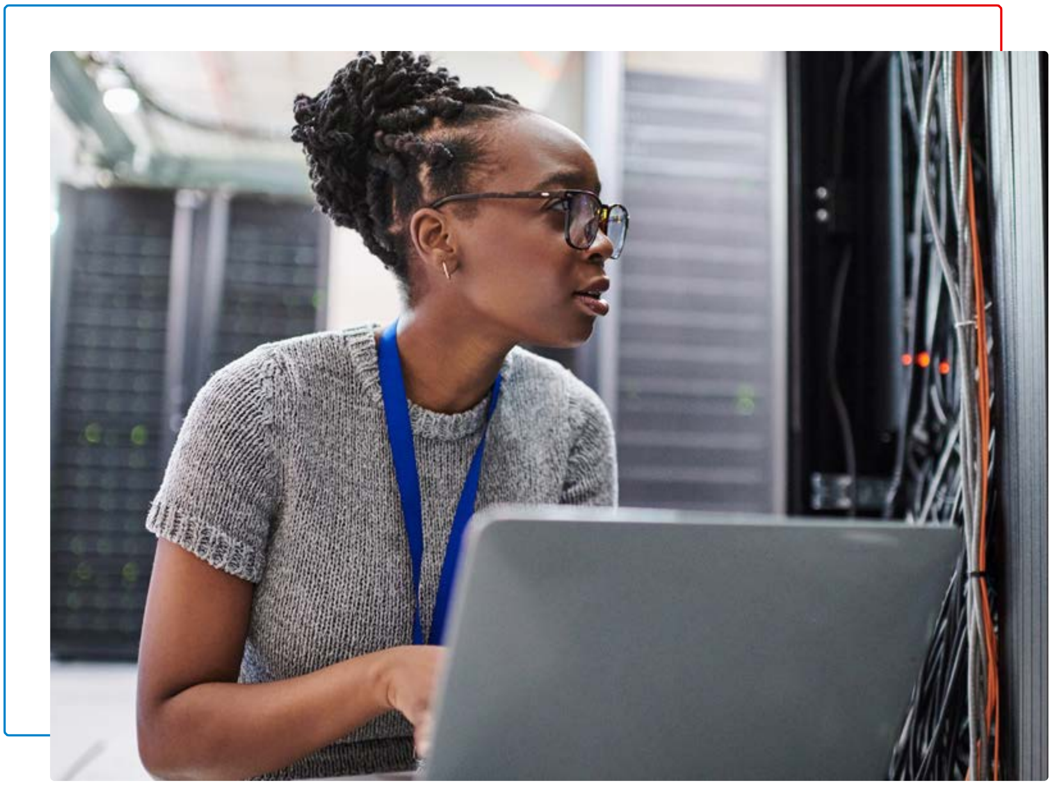

Establishing Resilience Against Edge Device Attacks

Over the past 10 years, organizations have increased visibility throughout their digital environments. As a result, they are detecting attackers faster[12] and have made significant progress in proactively securing their environments from threats like password reuse and brute force attacks as they continue to move towards defense in depth style architecture.

While this trend progresses in the right direction, most organizations center detection and response around the visibility provided by their endpoint detection and response (EDR) solutions. However, EDR solutions are deployed, as the name implies, on endpoints. In other words, firewalls, IOT devices, VPNs, hypervisors, and many other devices are not typically supported by EDR, and are therefore commonly referred to as "edge devices." What happens when malicious actors start targeting those devices?

Because edge devices by definition sit outside the typical detection range of most organizations, they provide attackers with enormous value during intrusions. Edge devices will always be targets to adversaries, just in different ways. These edge devices provide many valuable services to organizations such as monitoring internal security tools, but historically have not been supported by EDR solutions and are rarely monitored at the system level. This type of system-level monitoring is needed to identify if code changes or targeted malware is installed.

Edge devices are leveraged for security hunting and protection and are not inherently protected themselves. More to the point, vendors typically do not enable direct access to the operating system or filesystem for users. Because detections aren't extended to these edge devices and systems, defenders are limited in their capacity to perform analysis into underlying, potentially anomalous behavior.

[12] M-Trends 2023, Mandiant April 2023

Over the past five years, Mandiant has seen increasing evidence to suggest nation-backed adversaries are targeting edge devices. This focus on edge devices is as concerning for defenders as it is advantageous for attackers. Malicious intrusions are targeting edge devices likely to gain a foothold or maintain persistence in the target environment. Beyond a simple foothold, edge devices offer malicious actors a host of advantages. First among them being that edge devices have elevated visibility and privileges within the environment to provide network monitoring or a secure point of access. Access to these devices also allows the attacker to control the timing of the operation and can reduce the chances of detection. Edge devices, by definition, are not visible to EDR solutions, meaning that all these advantages are conferred on attacks as well as the ability to remain hidden from defenders.

Nation-backed adversaries often dedicate considerable time and effort for extensive research and development cycles to identify and create exploits for previously unknown vulnerabilities. Mandiant has investigated dozens of intrusions over the years where suspected China-nexus groups have exploited zero-day vulnerabilities and deployed custom malware to steal user credentials and maintain long-term access to the victim environments. For example in 2022, UNC3886 targeted edge devices such as firewalls and later in the attack life cycle, hypervisor technologies.

UNC3886 Case Study

Multiple components of the Fortinet[13] ecosystem were targeted by UNC3886 before they moved laterally to VMWare infrastructure. These components and their associated versions, at the time of compromise, are listed as follows:

- **FortiGate: 6.2.7** – FortiGate units are network firewall devices which allow for the control and monitoring of network traffic passing through the devices.
- **FortiManager 6.4.7** – The FortiManager acts as a centralized management platform for managing Fortinet devices.
- **FortiAnalyzer 6.4.7** – The FortiAnalyzer acts as a centralized log management solution for Fortinet devices as well as a reporting platform.

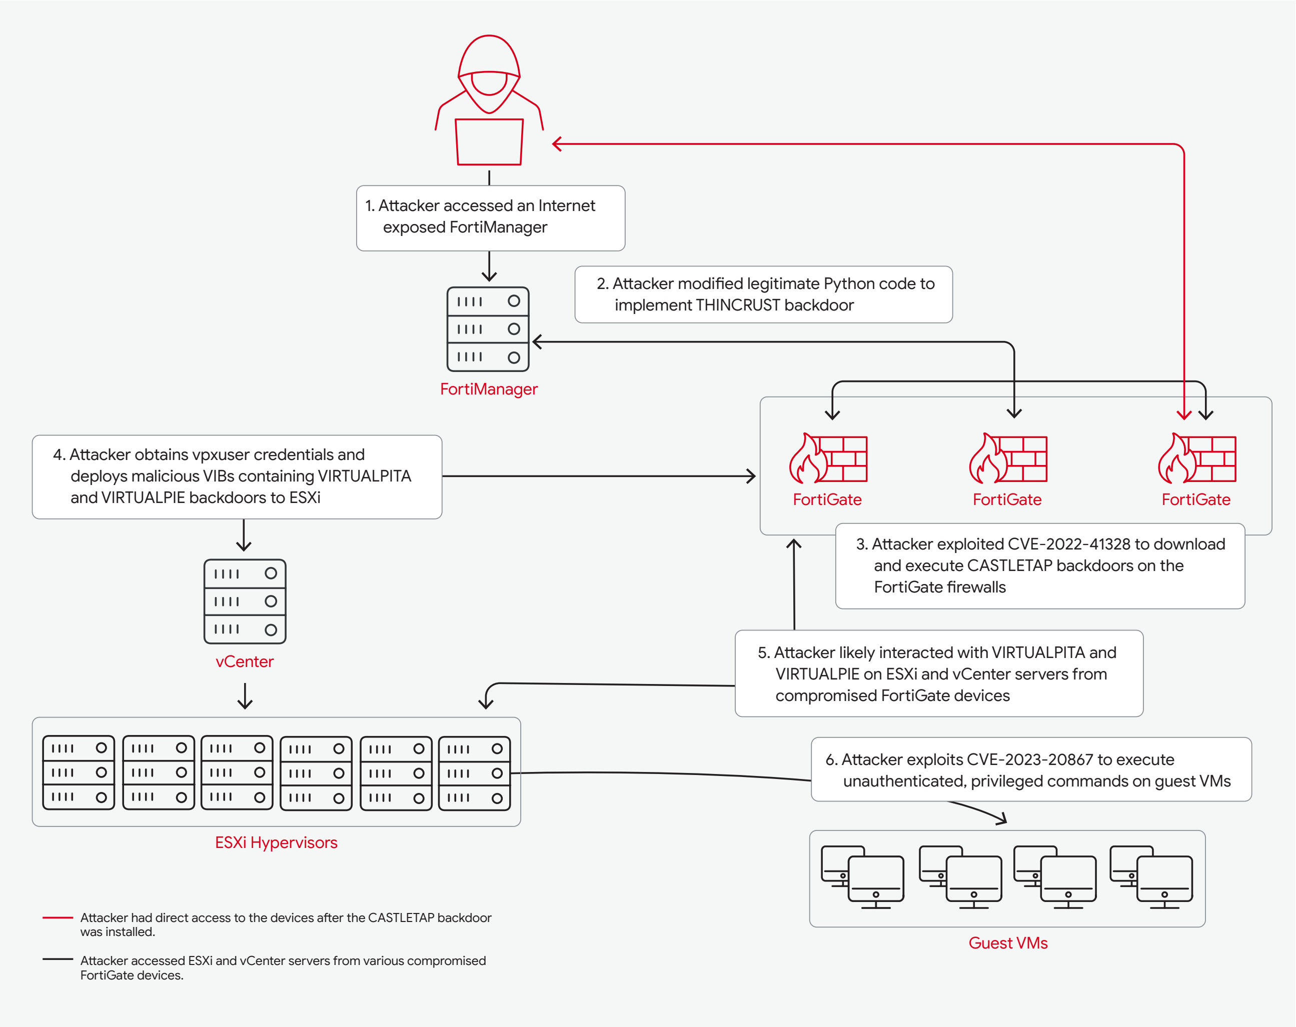

1. Attacker accessed an Internet exposed FortiManager

2. Attacker modified legitimate Python code to implement THINCRUST backdoor

3. Attacker exploited CVE-2022-41328 to download and execute CASTLETAP backdoors on the FortiGate firewalls

4. Attacker obtains vpxuser credentials and deploys malicious VIBs containing VIRTUALPITA and VIRTUALPIE backdoors to ESXi

5. Attacker likely interacted with VIRTUALPITA and VIRTUALPIE on ESXi and vCenter servers from compromised FortiGate devices

6. Attacker exploits CVE-2023-20867 to execute unauthenticated, privileged commands on guest VMs

— Attacker had direct access to the devices after the CASTLETAP backdoor was installed.

— Attacker accessed ESXi and vCenter servers from various compromised FortiGate devices.

FIGURE 5: Activity after internet access restrictions implemented to FortiManager

[13] https://www.mandiant.com/resources/blog/fortinet-malware-ecosystem

In 2022, Mandiant began tracking UNC3886, a group with a suspected China-nexus. This group specifically targeted the Fortinet ecosystem and eventually moved laterally to access VMWare Infrastructure within targeted environments. To gain this access, UNC3886 proved to have sufficient knowledge of multiple Fortinet solutions including FortiGate (firewall), FortiManager (centralized management solution), and FortiAnalyzer (log management, analytics and reporting platform). With this knowledge, UNC3886 deployed a backdoor, tracked by Mandiant as **THINCRUST**, across FortiManager and FortiAnalyzer devices to gain persistence. Then UNC3886 leveraged access to FortiManager native scripts to exploit CVE-2022-41328 to download and execute another backdoor, **CASTLETAP**, across FortiGate devices to further maintain access within the environment.

Mandiant observed SSH connections from the Fortinet devices to ESXi servers within the target environment followed by the installation of vSphere Installation Bundles[14] that contained **VIRTUALPITA** and **VIRTUALPIE** backdoors.

In another scenario, where the FortiManager was restricted from the internet, UNC3886 leveraged previously established access to install a network traffic redirection utility Mandiant tracks as **TABLEFLIP**, and a reverse shell backdoor variant of **REPTILE**, on the FortiManager. This combined use of malware allowed UNC3886 to circumvent network access control lists (ACLs) in place to restrict external access.

In both of these scenarios, malicious activity was detected following a full compromise of both the Fortinet ecosystem and the VMware hypervisor, once UNC3886 began performing reconnaissance commands and exfiltrating data using legitimate system processes.

For a detailed account of this case study, please refer to: the blog "Fortinet Zero-Day and Custom Malware Used by Suspected Chinese Actor in Espionage Operation".

CASTLETAP is a Linux binary that passively listens for packets and activates the backdoor functionality when it receives an ICMP Echo packet. Within these packets, the malware also searches for C2 server information that it can connect back to over SSL socket. Its capabilities include uploading and downloading files, spawning normal, and busybox-based shell.

THINCRUST is a Python backdoor embedded in a third-party library code that allows remote command execution, reading, and writing files via HTTP requests. The encrypted commands are stored in HTTP cookies.

VIRTUALPITA is a 64-bit passive backdoor for Linux and VMware ESXi that creates a listener on a hardcoded TCP or VMCI port numbers. It supports arbitrary command execution, file upload and download, and the ability to start and stop vmsyslogd.

VIRTUALPIE is a backdoor written in Python that spawns a demonized IPv6 listener on a hardcoded TCP port. It supports file transfer, arbitrary command execution, and reverse shell capabilities. It communicates using a custom protocol and the data is encrypted using RC4.

TABLEFLIP is a Linux utility that performs traffic redirection. It passively listens on all active interfaces for specialized command packets. These packets contain XOR encoded IP address and port number to redirect traffic to using iptable commands.

REPTILE is a publicly available Linux rootkit written in C. It supports backdoor functionality which can be activated through ICMP, UDP or TCP packets via port-knocking. Additional capabilities include reverse shell and file transfer.

[14] https://blogs.vmware.com/vsphere/2011/09/whats-in-a-vib.html

APT29 Case Study

Mandiant has also observed nation-backed actors, like APT29, targeting similar types of edge device appliances with a novel tunneler.

In early 2022, after gaining access to the target environment, APT29 deployed **QUIETEXIT** to endpoints throughout the environment. In one case, APT29 hijacked legitimate application specific startup scripts to enable **QUIETTEXT** to run at startup, as it does not have native persistent mechanisms. **QUIETEXIT** supports full SSH functionality and APT29 leveraged a SOCKS tunnel into the target environment. This allowed APT29 to execute tools to steal data with little to no evidence on the target computer. APT29 targeted network attached storage (NAS) masquerading the binary name to blend in with legitimate files on the file system. To maintain additional access, APT29 deployed a secondary backdoor, **REGEORG** web shell, on a DMZ web server. This, combined with a lack of supported anti-virus or EDR solutions, aided in a prolonged dwell time.

QUIETEXIT is a reverse SSH tunneler that connects out to a remote C2, but requires a password to authenticate. QUIETEXIT can execute commands or proxy traffic via SOCKS. QUIETEXIT is derived from the open source DROPBEAR SSL client-server software.

REGEORG is an open-source utility used to tunnel webshell traffic.

QUIETEXIT. Mandiant observed command and control (C2) systems were primarily legacy conference room camera systems, which were likely infected with the server component of QUIETTEXT. By targeting these trusted systems, APT29 remained undetected in target environments for at least 18 months.

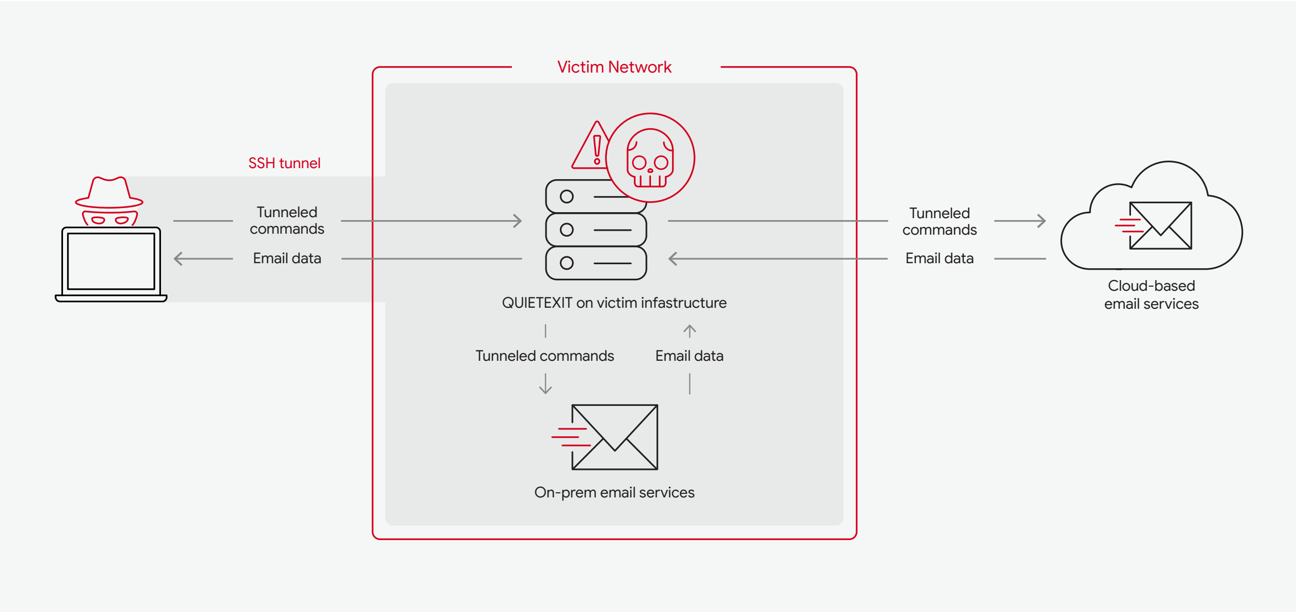

FIGURE 6: Tunneling though QUIETTEXIT

Completing the mission, APT29 successfully obtained privileged credentials to the target's email environment and focused efforts on executive teams and employees who work with corporate development, mergers and acquisitions, or the IT security staff. In some cases, APT29 leveraged the same eDiscovery and Graph API tools used to perform programmatic searching and access to email data that investigators use to conduct response efforts. These tools allowed APT29 to conduct bulk email exfiltration.

For a detailed account of this case study, please refer to: the blog "Eye Spy on Your Email".

APT28 Case Study

In 2022, Mandiant observed APT28 deviate from historic activity. This group demonstrated a preference towards compromising edge infrastructure to conduct a variety of operations, a technique referred to as "Living on the Edge." Since the outset of the war in Ukraine, the Russian Military Intelligence, or known as the GRU, has attempted to conduct successive and almost constant campaigns of cyber espionage and disruption aimed against key services and organizations within Ukraine. This balance of access to and actions against target organizations relies on the compromise of edge infrastructure such as routers and other internet connected devices.

For a detailed account of this case study, please refer to:
M-Trends 2023, "The Invasion of Ukraine: Cyber Operations During Wartime".

Key takeaways

In these case studies, evidence of compromise was detected within the environment during post exploitation activity, as by design, actors target edge networks to remain undetected. During the investigations, Mandiant conducted thorough reviews of impacted systems to identify the initial entry vector. In these cases, evidence existed to trace access back to edge device IP addresses. This led investigators down the path of working with vendors to collect forensic images of these devices to perform further analysis. Cross organizational communication and collaboration is key to providing both manufacturers with early notice of new attack methods in the wild before they are made public and investigators with expertise to better shed light on these new attacks.

What you can do to protect against these attacks

Cyber espionage related actors have increased their investment in research and development of tooling and exploits against systems that do not generally support EDR. These types of tooling and exploits require a deep understanding of the targeted operating systems. While organizations continue to build out security operations centers (SOCs), organizations should also continue to expand visibility further than endpoint detection. Visibility gaps allow threat actors to evade detection with minimal effort. Determining those visibility gaps is the next step to build an efficient SOC to support the security of the organization. Organizations should inventory devices on the network and evaluate if monitoring tools are available for each. Each device that does not support monitoring tools likely has vendor-specific hardening actions to ensure proper logging is enabled. Organizations should also ensure that these vendor-specific logs are forwarded to a central repository. Utilizing network access controls to limit or completely restrict egress traffic from these devices should also be evaluated. Implementing additional network monitoring and hunting for anomalous traffic to and from edge devices and other non-EDR enabled technologies allows further detection capabilities if these network controls are not feasible.

For additional resources please refer to the following:

Mandiant's Microsoft 365 Hardening Guide

Detection and Hardening within ESXi Hypervisors

6 Tips for Implementing Privileged Asset Management

Increased adoption of cloud services and SaaS applications is exponentially growing the number of accounts organizations must operate and manage. For example, today the average employee can access 30 corporate accounts and applications. Further, machine identities, digital certificates and keys now outnumber human identities by a factor of 45x[15].

For organizations that are struggling to reduce unnecessary accounts and remove excessive privileges for humans and systems that do not require it, implementing Privileged Access Management (PAM) can help.

PAM is a practice of controlling and securing access to assets within a business, by

Creating authorization workflows

Securely storing and encrypting secrets

Auditing, monitoring, and logging privileged access events

Setting policies for secrets management (e.g. Password Changes)

Securing and isolating access to target systems via a session manager

Traditionally, organizations lean on multi-factor authentication (MFA) solutions as a primary technology in their approach to PAM. If not implemented properly and maintained, the MFA solution can present unintended risk to the organization.

[15] 5 Reasons to Prioritize Privileged Access Management, CyberArk, 2022

2017	2019	2021	2022
Equifax Attackers gain across to PII of approximately 147 million consumers.	**Australian National University** Attackers access 19 years worth of PII of staff and students.	**Verkada** A supply chain attack in which attackers access the Verkada security camera system used by hospitals, schools and prisons.	**U.S. Dept. of Veterans Affairs** Sensitive credentials to systems containing health records exposed on GitHub.

FIGURE 7: Timeline of Breaches caused when attackers exploit Privileged Access Management solutions.

Attackers have historically exploited vulnerabilities in access management solutions with a high degree of success. Notable data breaches in size and scope cite PAM vulnerabilities dating back to 2017 with Equifax where attackers gained access to personal and privileged information for 147 million consumers[16]. Followed by the Australian National University where sensitive credentials to systems containing health records were exposed on GitHub[17]. In 2021 a supply chain attack against physical security vendor Verkanda exposed access to security camera systems used by hospitals, schools, and prisons[18]. Finally, in 2022 the U.S. Department of Veterans Affairs was victim to a data exposure of privileged account credentials by a contractor[19].

Mandiant has observed threat actors successfully bypassing MFA controls on multiple instances. In one case, Russian-based Advanced Persistent Threat (APT) groups performed MFA Fatigue Attacks[20] by repeatedly pushing second-factor authentication requests to the target victim's email, phone, or registered devices to gain access to email accounts resulting in wire fraud incidents.

In another example, Mandiant observed APT29[21] taking advantage of the self-enrollment process for MFA in an organization, which allowed anyone with a username and password to enroll a device. APT29 performed password guessing attacks to attempt to find accounts without enrolled devices and added their own.

Regardless of the organization's size or the maturity of the PAM program, security leaders should take the time to review these 7 tips when implementing PAM to help secure their business.

[16] Wallix Cybersecurity, Equifax Breach: Preventing Data Breaches with Privileged Access Management
[17] Australian National University, Incident Report on the Breach of the Australian national Universities Administrative Systems, 2019
[18] Verkanda, Summary: March 9, 2021 Security Incident Report, 2021
[19] FedScoop, VA investigates breach after federal contractor publishes source code, September 2022
[20] Mandiant, Suspected Russian Activity Targeting Government and Business Entities Around the Globe, December 2021
[21] Mandiant, You Can't Audit Me: APT29 Continues Targeting Microsoft 365, August 2022

01

Understand Privileged Accounts

Security and IT leaders are often asked, "What is a privileged account?" While the generic answer is that all accounts may have some level of privilege, following are several categories of accounts that provide higher privileges:

- **Domain Administrators:** Users that have full control over a domain.
- **Personal Privileged Accounts:** User accounts with more privileges than a regular user. Users utilize this on a case by case basis.
- **Default Accounts:** Accounts created automatically by the system or application (e.g. SA, Root, mysql, ec2-user).
- **Service Accounts:** Accounts that are assigned to machines and provide access to corporate systems, services and applications.
- **Root, Super Administrator, or Global Admin (Cloud):** Additional administrator accounts for a system that grants user full control over the local device.
- **Break-glass Accounts:** Accounts used to gain access to systems in the event of a security incident.
- **Security Accounts:** Accounts used by security personnel to access systems to perform security audits and investigations.

It is important to understand the risk associated with the misuse of accounts that provide privileged access. Start with least privilege to ensure that each user is only able to perform the actions defined by their role.

Pay special attention to roles that access personally identifiable information (PII) or intellectual property (IP).

Actions to take

- Perform a risk assessment of privileged access within your organization. Identify accounts for both humans and systems that present risk to critical assets and information. Consider the types of permissions, including identification procedures such as interactive logon. Prioritize the high-risk PAM accounts.
- Get buy-in from senior and executive management to drive the implementation of tools that are going to reduce the overall risk within an organization.
- Ensure that security and information technology teams collaborate in the implementation to account for the needs of various user groups and the communications and change management required in PAM implementations.
- Perform recertification and validation of permissions assigned to accounts. When maintaining this account, it should not change, if there is a new use case then the right type of account should be created, or a time limited policy for access should be granted, following the right approvals.

02

Establish a Continuous Process for Account Creation, Discovery and Onboarding

As PAM is implemented across an environment, PAM teams should be proactively addressing security gaps within the organization. A critical aspect of this effort is to onboard all necessary accounts that have been identified by application teams and any implications of managing these accounts through the PAM solution are understood.

Failing to onboard all necessary accounts can result in the proliferation of unsecured privileged accounts, leaving an organization vulnerable. Attackers can exploit these accounts to gain access to your systems, elevate privileges, move laterally, and establish persistence.

The problem of unsecured privileged accounts is particularly challenging when new accounts and services are added to an environment. These additions further increase the attack surface and scope of discovery for privileged accounts, exacerbating the risk of a security breach.

By maintaining a comprehensive inventory of all old and new accounts within an environment, organizations can quickly identify which accounts are at risk during a security incident. This helps to secure those accounts, identify the systems they have access to, and create trusted routes for accessing critical assets. This can alleviate the pressure on your security team, incident responders, and incident managers when responding to security incidents.

Actions to take

- Onboard accounts when they are created and avoid increasing the 'Discovery Scope'
- Use discovery tools to identify accounts that have been missed, onboard them, and implement effective controls to manage the account's lifecycle.
- Understand the scope of privileged accounts. Consider where Intellectual property, PII or PHI is stored, and how it is accessed.
- Establish a continuous process for account discovery. Work with the teams to adopt automation for creation, discovery and onboarding.
- Leverage the MITRE ATT&CK® framework to review dozens of commonly abused adversary techniques used in privilege escalation attacks.

03

Ensure Proper Access Controls for the PAM Implementation

PAM is designed to protect the keys to the kingdom. Therefore access to PAM solutions should be managed as these user accounts and systems can become targets.

Authorizations for PAM administration should not be linked to the directories it is protecting. Use the PAM tool's built in Directory to manage this access.

Also consider the workflow for authorization. This control improves the ability for the organization to defend against insider threats, and helps the PAM team to understand how to control the system implemented for PAM.

Actions to take

- Onboard these types of accounts
  - PAM Tool administrators
  - Application accounts
  - Server accounts
  - Automation and Scripting
- Review the Access Model for using the PAM solution and identify toxic combinations of privilege.
- Make sure the administration accounts for the PAM solution are correctly permissioned, limiting access to the credentials the application secures.
- Set up authorization and access workflows to secure critical accounts.
- Enable your teams with essential security knowledge to continuously improve the organization's security posture through training, enablement and support from experts.

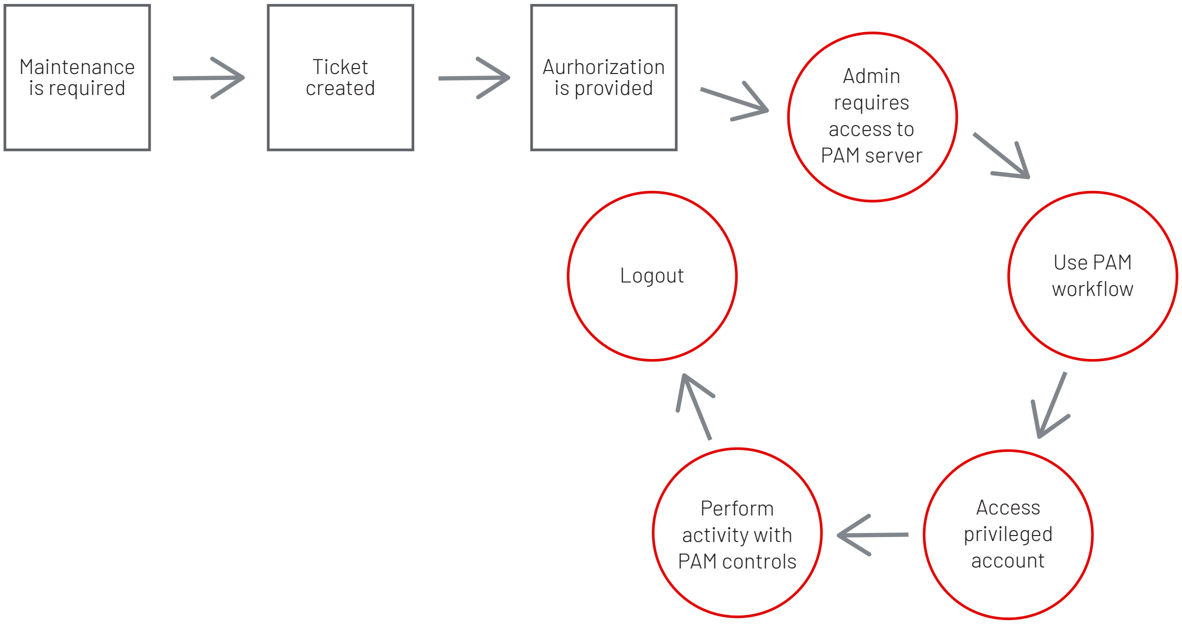
Maintenance is required
Ticket created
Aurhorization is provided
Admin requires access to PAM server
Use PAM workflow
Access privileged account
Perform activity with PAM controls
Logout

FIGURE 8: PAM authorization workflow

04

Map and Secure Access Routes

When accessing assets within the organization there are key activities to secure privileged access. One action is to avoid letting passwords become exposed on an endpoint system. There are also considerations of what happens between the user and the target system.

Is traffic being sent through or across a trusted path?
How can this be enforced?

Other questions to ask include:

- Does the connection involve a web resource?
- Is the web resource being connected to directly from a workstation?
- Is HTTPS being used for connections?

Understanding the flow and path to targets helps
to calculate the risk of a threat. How applications are accessing secrets and how those secrets are used between applications and servers should also be considered.

Actions to take

Create a map of access within the environment, and check that the architecture maps to what is provided with the PAM tools.

- Identify the secure protocols that should be used and identify the routes that must be taken to gain access to session managers and targets.
  - Connect over HTTPS instead of over RDP
  - Consider using a reverse listener to keep NACL's clear and only allow outbound access.
  - Establish a credential and access tiering model.
- Make sure third parties are secured with secure authentication mechanisms and authorization has been correctly provisioned.
- Confirm that third parties are connecting over a secure method.
- Identify access routes:
  - Path through internal and public networks
  - Clean source systems
  - Strong access tiers
  - Credential protection
  - Test access routes

05

Implement Logging with Adequate Retention

Logging and monitoring provide valuable information that is critical in the aftermath of a cyber attack. Logging and monitoring aim to assist in identifying the scope and impact of a cyber incident. Forensic investigators use log sources to answer many questions during a cyber incident. For example, in order to identify data exfiltration, forensic investigators rely heavily on firewall or netflow log data. These logs can answer questions around data exfiltration and how much data has left the network.
Another example involves tracking user activity. If a privileged account is compromised, logging can help track down actions performed by that user.

Actions to take

- Implement a logging and monitoring solution that captures privileged accounts activities across the organization.
- Ensure adequate log retention: Develop a logging and monitoring policy which outlines the types of activities that need to be logged and the retention for those logs.
- Confirm that log data is being sent to appropriate systems, but more importantly that the data is being used to enrich defenses within the organization. Security teams can leverage threat analytics to further improve the controls that have been put in place, and as indicators of a threat actor landing within the environment.
- Look for anomalies of user activity including system access.
- Understand the plan of action when an incident occurs. Create a plan to review the audit logs and data if not storing these in a SIEM.

06

Implement MFA

Multi-factor authentication (MFA) is important to prove digital identity and secure access to a system via a single actor or entity.

MFA requires users to provide multiple authentication factors to access an application. Two of the most common forms of MFA are one-time passcodes (OTP) and push notifications. OTP are codes that users receive on their mobile devices through MFA applications (e.g. Google Authenticator), which can be used to authenticate. Push notifications send a notification to a user's mobile device to approve or reject a login attempt.

Both of these methods are vulnerable to phishing attempts or man in the middle attacks. Recent attacks have shown that MFA push notifications, or SMS delivered codes are not enough to protect access. For example, an MFA fatigue[22] attack occurs when threat actors bombard a user with push notifications in hopes of the user getting frustrated and hitting accept.

Actions to take

- Implement strong and phishing resistant MFA
  - FIDO2 authentication by using biometrics or hardware keys (e.g. YubiKey)
  - Use challenge response mechanisms that do not simply allow accept (number matching)
- Conduct employee-wide training to ensure employees have the knowledge available to utilize MFA correctly
- Introduce alerting and risk-based tagging against accounts that seem to be under attack
  - Geo-location
  - Abnormal hours of access
  - Excessive requests for MFA challenge response
- Review the authentication assurance of the authenticators being used
- Perform threat hunting against these identities

[22] Mandiant, Suspected Russian Activity Targeting Government and Business Entities Around the Globe, December 2021

In Conclusion

As more enterprises move towards implementing PAM solutions to enhance their security posture, it is critical that these tools are appropriately configured and implemented. Misconfigurations, lack of appropriate access management and not fully utilizing built-in capabilities are some of the key factors that can result in creating a false sense of security and, in some unique cases, may result in increasing enterprise risk.

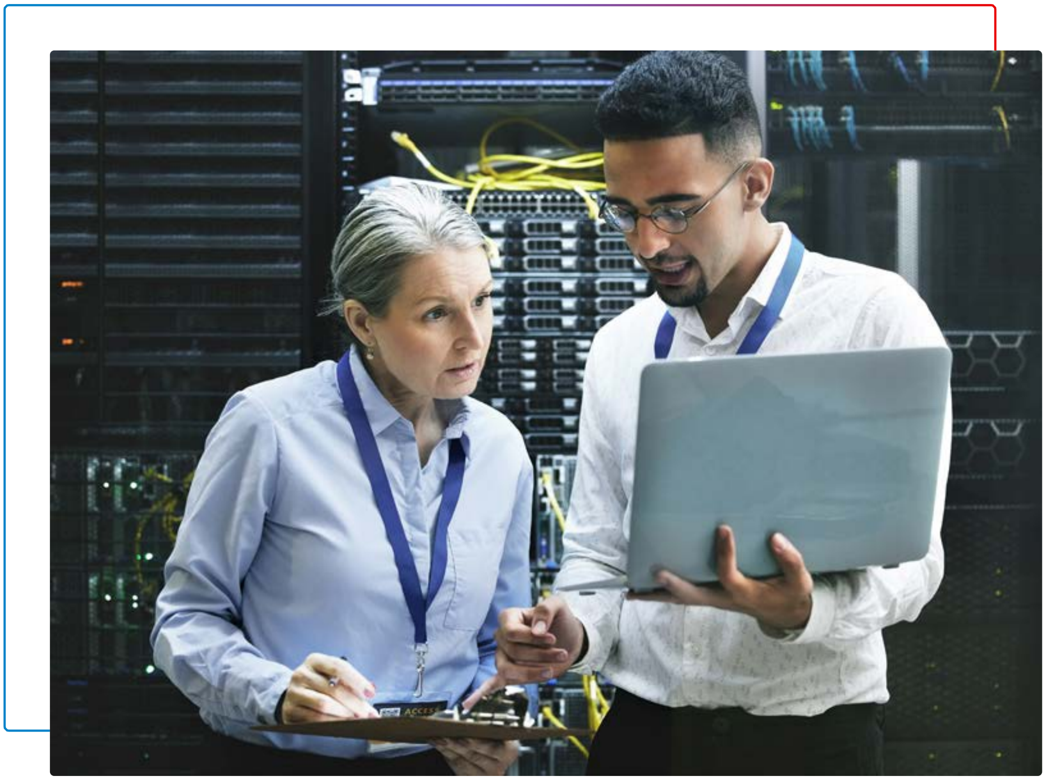

Learn more at www.mandiant.com

Mandiant

11951 Freedom Dr, 6th Fl, Reston, VA 20190
(703) 935-1700
833.3MANDIANT (362.6342)
info@mandiant.com

About Mandiant

Mandiant is a recognized leader in dynamic cyber defense, threat intelligence and incident response services. By scaling decades of frontline experience, Mandiant helps organizations to be confident in their readiness to defend against and respond to cyber threats. Mandiant is now part of Google Cloud.

MANDIANT
NOW PART OF Google Cloud