INCIDENT READINESS

Preparing a proactive response to attacks

An F-Secure Consulting whitepaper

F-SECURE
CONSULTING

F-Secure Consulting is a research-led cyber security consultancy, partnering with enterprises and early adopters worldwide.

We help protect organizations and individuals in a fast-changing world by providing evidence-based recommendations and services. Our evidence is the product of continuous research and development, conducted by consultants whose work has been recognized and utilized by the security community for almost two decades.

We are intellectually curious, obsessed with solving the most complex security challenges, and driven to build our clients' cyber resilience through collaboration and technological innovation.

www.f-secure.com/consulting
@FSecure_Consult
linkedin.com/company/f-secure-consulting

CONTENTS

BACKGROUND

Recent events have shown that it is possible to effectively respond to a crisis. Those that succeed call on past experience to proactively develop their readiness for inevitable future risks. The measures they take are pre-emptive and continuous. For organizations considering the impacts of a cyber attack, there is much to be learnt from this “ready” mindset.

Cyber attacks are a business reality, with 75% of organizations in the UK having experienced one between March 2019 and March 2020[1]. However, the impact of the attack and the damage caused, can be mitigated. Depending on an organization’s incident response (IR) and crisis management efforts, an attack with the potential to disrupt business continuity may have little impact. This can be achieved by anticipating incidents and developing a response around specific threats, their likelihood, and the motives of the attacker. It is never the result of a defensive arsenal of technology on its own.

> *“There are organizations that are spending a ton on cyber security and they have very bad risk postures. There are others that are not spending very much but they have very good risk postures. The bottom line is it’s about their level of readiness.”* – Paul Proctor, VP, Gartner

[1] https://www.gov.uk/government/publications/cyber-security-breaches-survey-2020/cyber-security-breaches-survey-2020

PROACTIVE VS REACTIVE RESPONSE

Readiness—in the context of cyber incidents—can be understood as a state of being and working to continually be prepared for compromise, such that business continuity can be maintained throughout an attack. It is both the measurement of an organization's current security posture and a set of actions. The latter, when improved, can increase overall organizational resilience.

Readiness can be aligned with the "preparation" phase in the IR lifecycle (see fig.1.), influencing all successive phases. It is broad in scope and designed to begin as far in advance of an incident as possible; this preparation phase may take place months or even years before an incident. Although this preparation phase is fundamental to IR as a practice, it is commonly overlooked.

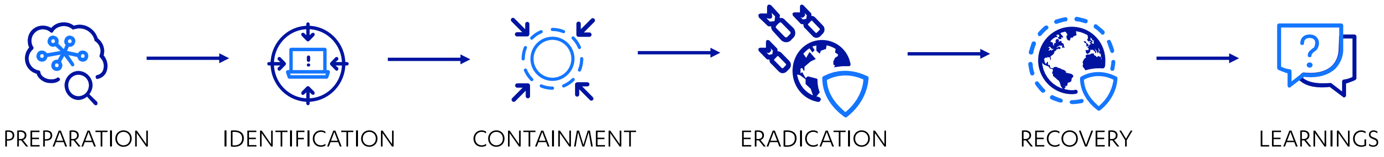

Fig. 1. Incident response lifecycle

INCIDENT READINESS IN ACTION

The following walkthrough is based on real IR engagements and contrasts the same attack being deployed against two near identical organizations in the financial sector. It illustrates the success of response with and without prior readiness activity and welcomes the reader to make their own judgement about the value of incident readiness and how they may apply such a mindset in their own organization.

SETTING THE SCENE

ENTERPRISE-A

Enterprise-A has a workforce of 500,000 employees, and operations in several geographic regions. One central group security function exists in addition to local security representatives in each regional office. Regions are responsible for their own security function and detection technology; no monitoring of this technology takes place at the central security group level.

Security initiatives include penetration testing for applications, infrastructure, and networks. Some policies and processes for IR exist, but are not widely understood due to a lack of training.

ENTERPRISE-B

Enterprise-B is identical in profile to Enterprise-A, except that it has recently undertaken an incident readiness project including the following elements:

- IR plans and playbooks developed at the regional level
- First responder training in place for both group and regional staff
- Tabletop incident simulation exercises carried out on a regular basis for each region, with group security playing the role of incident management
- Review of the central logging solution that collects and aggregates key logs, with an emphasis on logging strategy and policy for priority assets within group and regional offices
- A review of high privilege accounts and the management thereof

INCIDENT PHASE 1: INITIAL BREACH

ENTERPRISE-A

DAY 1

The third-party developer responsible for developing and maintaining the organization’s money transfer application contacts a representative from head office—suspicious files have appeared on the application servers of two regional offices. These files were not uploaded by the development team, which is the only team with access to the servers.

DAY 10

Group security at head office contacts the region where the suspicious files have been identified. It transpires that the office is receiving numerous client complaints regarding small, unrecognized transactions from their accounts. The application owner is currently investigating the matter but is unsure of the cause at present.

DAY 21

The developers are instructed by group security to delete the suspicious files and report any further unusual file activity on other regional servers.

Group security investigates the unrecognized transactions in the transaction logs and discovers a possible authentication flaw, which allows authentication keys to be reused. The team issues a patch and distributes it to all regions where the application is hosted.

ENTERPRISE-B

DAY 1

Upon receiving notification from the development team about a suspicious file on the application servers for its money transfer application, a ticket is raised and assigned as a low-risk priority. An analyst from the group security team calls the developer who raised the notification. They ask them to share their screen to enable a two-way discussion about the suspicious nature of the file. The servers are Linux, so the developer logs on via Secure Shell (SSH) and prints the file contents to the shared screen.

The analyst takes note of the file contents:

```
#<?php
/***********************************************
********************************
 * Copyright 2017 WhiteWinterWolf
 * https://www.whitewinterwolf.com/tags/php-
webshell/
 *
 * This file is part of wwolf-php-webshell.
 *
 * wwwolf-php-webshell is free software: you can
redistribute it and/or modify
 * it under the terms of the GNU General Public
License as published by
 * the Free Software Foundation, either version
3 of the License, or
 * (at your option) any later version.
```

A quick Google search of the string “WhiteWinterWolf” confirms that the files are web shells. The application servers are now considered compromised, so the analyst raises the ticket priority and directs the infrastructure team to snapshot the virtual server’s hard drives. These snapshots are sent for forensic analysis.

Operational obligations and a lack of functional backup servers prohibit the affected servers from being taken offline. Instead, the analyst locates the security team’s incident playbook for host compromise and follows the instructions to generate a hash for the suspicious file in question. The hash is sent as part of a notification to all regions, advising them to search for instances of the file on their servers.

DAY 10

Further investigation leads the analyst to a string of client complaints about unauthorized transactions on their accounts. After requesting application transaction logs for review, they quickly identify that an authentication key has been reused for two different accounts. This indicates an authentication flaw in the application. As the web shell and the unauthorized transactions are linked, the risk of the incident is reassessed by the incident manager; the decision is made to engage an external IR provider.

Virtual hard drive snapshots and the accompanying logs are sent to the provider for analysis.

Whilst the hard drive snapshots and logs are processed, Enterprise-B receives guidance from the IR provider to increase monitoring on the affected assets. Over the next few days, a Linux endpoint detection and response (EDR) agent and updated AV signatures are deployed to the application servers and the logs are actively monitored by the security team.

DAY 21

An update on the investigation thus far is received from the IR provider. It details the following:

- External reconnaissance activity is visible in the local server logs
- A web shell was uploaded via the file upload functionality of a web application used by the developers to upload and maintain code on the server
- Additional malware was then installed, allowing the attacker to interact directly with the server
- The attacker was able to view the transaction logs stored in clear text on the server, where they were able to access the authentication keys
- There is some indication that the attacker is using the application API directly from the server via the GNU Wget functionality

Group security establishes an internal cyber security incident response team (CSIRT)[2] and convenes to discuss the incident and its remediation. Representatives are involved from every region.

It becomes clear that the attacker has a foothold on all the regional application servers. Their objective is also clear: to extract funds by transferring money between accounts directly from the application servers. The initial malicious transactions of nominal amounts were designed to assess the effectiveness of this approach.

[2] https://www.f-secure.com/content/dam/f-secure/en/consulting/our-thinking/collaterals/digital/F-Secure-Consulting-Blue-Team-paper-2020.pdf

The CSIRT decides to issue a temporary patch, stopping the authentication key from being reused. Monitoring on the servers is increased while the application team builds a new application server for distribution to the affected regions.

DAY 34

The new server image and patched application undergo a penetration test before being deployed as the replacement in all regions. Approval of the single server image and application version happens quickly to avoid any further operational issues. It is imperative that all vulnerabilities are addressed before deployment to the regions.

INCIDENT PHASE 2: DOMAIN COMPROMISE

ENTERPRISE-A

DAY 182

An anti-virus alert is triggered on three servers in the regional office where activity on the application servers was first observed. The alerts indicate use of the credential harvesting tool Mimikatz[3].

An administrator is assigned to log on to the affected hosts using the local administrator account. Their attempts to identify the alerted files are ineffectual, as they have been successfully quarantined and are not present on disk.

DAY 190

After speaking to the administrator, the security manager concludes that the threat has been contained by the anti-virus solution. The administrator had attempted to extract the Windows event logs but advised that the logs had since rolled over and were not preserved. Every user password within the regional office network is reset as a precaution and the investigation is closed off.

[3] https://github.com/gentilkiwi/mimikatz

ENTERPRISE-B

Day 182

An anti-virus alert triggers an investigation into three servers in a regional office. The regional security manager creates a ticket and assigns it as medium risk due to the nature of the alert—that a known hacking tool has triggered alerts on multiple servers. The security manager instructs the administrator to take snapshots of the virtual servers and capture the event logs as per the first responder process.

Meanwhile, the security manager reviews the access logs in the security information and event management (SIEM) system to identify any accounts that have recently logged on to the servers. One account stands out as anomalous, as it belongs to a manager in the money transfer application team. This account has never previously accessed the servers. Despite no business case, it also has local administrator permissions for them.

Day 183

The security manager consults the application manager, and it is concluded that the activity was not legitimately performed by the organization. After arranging a replacement laptop for the application manager, the security manager takes possession of the infected laptop. A forensic imaging tool is used to capture a disk image of the laptop onto an encrypted external hard drive, as per the instructions in the first responder guide.

The external hard drive and copies of the server snapshots are shipped via courier with next-day delivery to group security for forensic investigation. The team has already been briefed on the incident and the incident ticket is updated with relevant details.

The password for the administrator account is reset and the affected servers are restored from a previous backup.

Day 185

Group security forms a CSIRT with key representatives from the security, infrastructure, and executive teams. The decision is made to engage Enterprise-B's external IR provider for reinforcement. The organization's incident manager invokes its incident response retainer and briefs the provider's consultants on what is known thus far. All available artifacts are shared with the consultants for analysis.

DAY 191

The investigation conducted by the IR consultants highlights the following events:

- The application manager was sent a phishing email containing a link to a fake Linkedin profile three weeks prior to the first anti-virus alerts. A two-stage payload was then downloaded onto their host. As the user had local administrator permissions, a registry key was used for persistence of the payload
- Several randomly-named text files were found in the C:\Windows\Temp\ directory of the laptop, which appear to be the output of a keylogger. These include the contents of emails, as well as passwords for internal platforms. The latter includes the SSH password to a jumpbox that provides access to the Linux environment where the application servers are hosted
- The application manager's credentials were then used to log on to the three servers flagged by the anti-virus alerts
- The investigation streams for the servers are ongoing. The IR provider has so far identified a folder containing several executable files on each server. These executables appear to have been used by the attacker for enumeration
- One such executable file named "<COMPANY_NAME>.exe" has been submitted to the malware analysis team for reverse engineering to scrutinize its structure and purpose. Preliminary analysis indicates that some extracted strings have keylogging capability and that the executable is a custom piece of malware

DAY 193

As a result of these findings, group security meets to assess the threat and plan remediation. The consequent actions are informed by both the findings and are in line with its domain compromise playbook:

- Exchange administrators from each region are instructed to search for instances of the phishing email in their regional user base. Those regions with a proxy are instructed to search the proxy logs for the phishing URL
- A list of hashes for the files identified in the attacker's tool folder is sent to each regional office. It comes with instruction for the regional security teams to use their anti-virus solution to search for instances of these on their estates
- The regional administrator for the Linux environment hosting the money transfer application is instructed to search for instances of logons using the compromised application manager's credentials.

DAY 198

While no other instances of the phishing email are found, the malware hash search shows hits on six additional servers in the region where the malicious files were first identified. The additional server list includes a domain controller.

No anomalous authentication events are identified in the Linux environment. However, it is recognized that limitations in the logging prohibit evidence being collected across the entire period of interest.

The administrator increases the log retention policy to 60 days, resets all user credentials, and hardens the firewall rules to ensure that only listed IP addresses can access the Linux servers. A collection script is run on all Linux hosts in the environment, the output of which is shared with the IR provider for analysis.

DAY 201

Following the processes and playbook guidance defined for regional domain compromise, group security obtains executive approval to sever the domain trust between the affected region and head office.

The following instructions are sent to the regional security team to eradicate the threat:

1. Restore backups of the affected servers from one month ago[4]
2. Reset all user and administrator passwords via Active Directory (AD)
3. Reset the KRBTGT password as per Microsoft's best practice guide
4. Identify additional accounts which may have excessive privileges
5. Increase monitoring for all affected assets
6. Monitor for operational impact and report back to group security
7. Schedule a domain hardening assessment with the organization's cyber security partner

By following these processes, Enterprise-B prevents the potential business impact of the attack and successfully contains the incident. It may now evaluate the incident with its IR provider and report back to the business. As part of the lessons learned process defined in each of the organization's playbooks, the CSIRT conducts a meeting to discuss the information ascertained during the incident. An action plan to address the identified shortcomings is developed during the meeting, and the resulting actions list is assigned to key individuals for completion.

[4] The IR provider confirms these are clean, as they predate the servers' compromise. The executive team know that the backup restoration may cause some business impact and loss of production data. Stakeholders are on standby to deal with any fallout.

INCIDENT PHASE 3: ATTACKER CASH-OUT

ENTERPRISE-A

DAY 362 (ATTACKER OBJECTIVE ACHIEVED)

Enterprise-A receives an influx of reports from clients who have lost significant sums of money through its money transfer application. These come from multiple regions. The organization's executive team convenes a crisis management meeting with group security and regional security representatives to investigate and control the situation.

An external IR provider is engaged and briefed on the findings so far. The provider requests the following artifacts to support the investigation:

- Disk images from the Linux application servers
- Transaction logs from the money transfer application
- Firewall logs for the Linux environment

DAY 368

As most of the application server disk images need to be couriered from the respective regions, limited artifacts are available for analysis. One disk image of a regional application server and some network logs are provided initially.

Preliminary findings reveal that malware has been running on the host for over a year. The transaction logs again show that authentication keys have been reused in money transfers, for different accounts, via the API on the application server. The patch issued by group security to prevent the reuse of authentication keys was reversed by the attacker shortly before the bulk of the transactions took place.

A timeline of key events from the attack confirms that the provided disk image was not the first in the attack chain. The region where the domain compromise occurred a few months prior is instead determined as the attack's most likely origin. This conclusion is drawn from the timestamps of the fraudulent transactions and connections on the firewall. Due to the time that has passed since this point, the available historic evidence is limited.

DAY 376

Additional disk images are received for analysis from the regional offices, including the region suspected as the starting point for the attack.

DAY 386

The attacker's foothold on the Linux servers is confirmed to originate from the corporate Windows environment, with the source of the SSH connection being the application manager's host. The IR provider requests a disk image of the host, and the laptop is shipped to head office for analysis; the absence of equipment or experience to acquire a disk image prohibits the regional security team to do so independently. This significantly hinders the investigation and causes a substantial delay in the analysis process.

No connection exists between the regional application servers. The servers in each region were individually compromised, rather than via lateral movement. Each was compromised via web shells uploaded via the developer's code maintenance application, followed by the installation of persistent malware.

The development team concedes that it could see the web shells being created and wrote a script to delete any files matching the names of said shells. They did not notify the organization, as they believed appropriate mitigation steps were taken by deleting the files.

The IR provider identifies variations of the web shells, which were not deleted by the developer's script. This is the mechanism that allowed the attacker to maintain persistence on the application servers.

DAY 399

The laptop belonging to the application manager is received. It takes the regional team a few hours to find and provide its BitLocker recovery key, after which time the IR provider acquires a disk image and begins analysis.

Day 416

Significant time has passed since the domain compromise, and many laptop artifacts have rolled over or been overwritten. The IR provider highlights that the host was compromised when the user clicked a phishing URL. Custom malware is also running on the host as well as the suspected output from a keylogger. Another investigation is launched to determine if the domain controller is still compromised.

POST-INCIDENT LEARNINGS

Following extensive forensic investigation, remediation of the incident takes several iterations. These cause long periods of operational downtime for Enterprise-A, impacting the business. The investigation finds that the attacker targeted the organization one year prior to their objectives being reached. They maintained persistence on several key assets during this time and systematically gathered data that would later inform and enable their final attack.

With the benefit of hindsight, three key detection events are identified. These represent the point at which incident response processes could have been invoked and may ultimately have resulted in a different outcome for Enterprise-A.

Enterprise-A's response approach is illustrated via the following key detection events, the response action taken for each, and the subsequent impact.

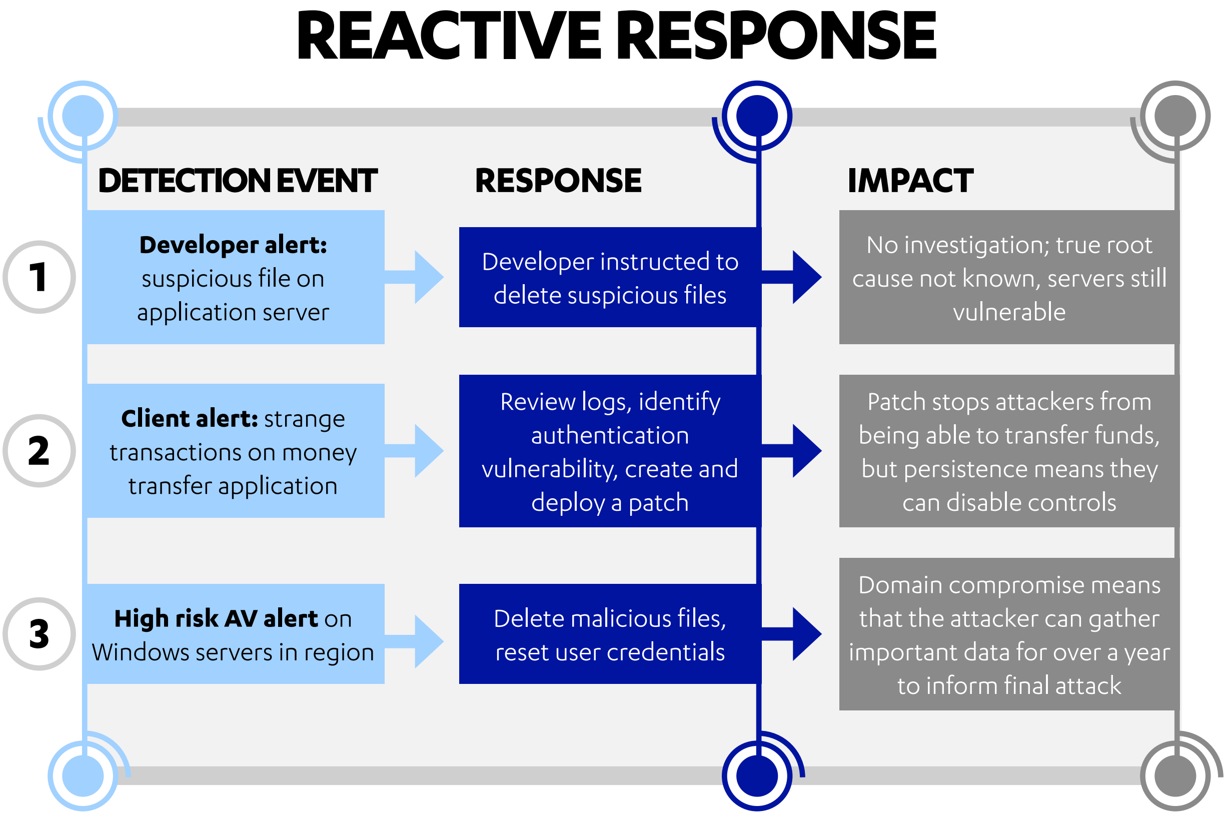

REACTIVE RESPONSE

	DETECTION EVENT	RESPONSE	IMPACT
1	**Developer alert:** suspicious file on application server	Developer instructed to delete suspicious files	No investigation; true root cause not known, servers still vulnerable
2	**Client alert:** strange transactions on money transfer application	Review logs, identify authentication vulnerability, create and deploy a patch	Patch stops attackers from being able to transfer funds, but persistence means they can disable controls
3	**High risk AV alert** on Windows servers in region	Delete malicious files, reset user credentials	Domain compromise means that the attacker can gather important data for over a year to inform final attack

Fig. 2. Points of detection, response, and impact within a reactive response approach

Following the investigation's conclusion, critical gaps in Enterprise-A's readiness posture are identified. At a minimum, the following high-level readiness measures could reduce the impact of incidents of this nature:

- The creation of an IR plan detailing relevant regulations, thresholds, incident severity levels, escalation matrices, and roles and responsibilities to be assumed during a cyber security incident
- The development of IR playbooks detailing high-level process flows for different categories of incidents; these would include incident management at the global level and incident categories at the regional level
- First responder training for regional representatives to advocate their responsibility for preserving evidence in suitable circumstances. This would also provide them with the skills required to conduct initial analysis from triaged data
- Incident response simulations developed and delivered systematically to each region, encouraging buy-in from the regional offices for the incident response plan and playbook

We can now consider what a proactive response may have looked like for Enterprise-A had the above elements been in place at the time of the incident:

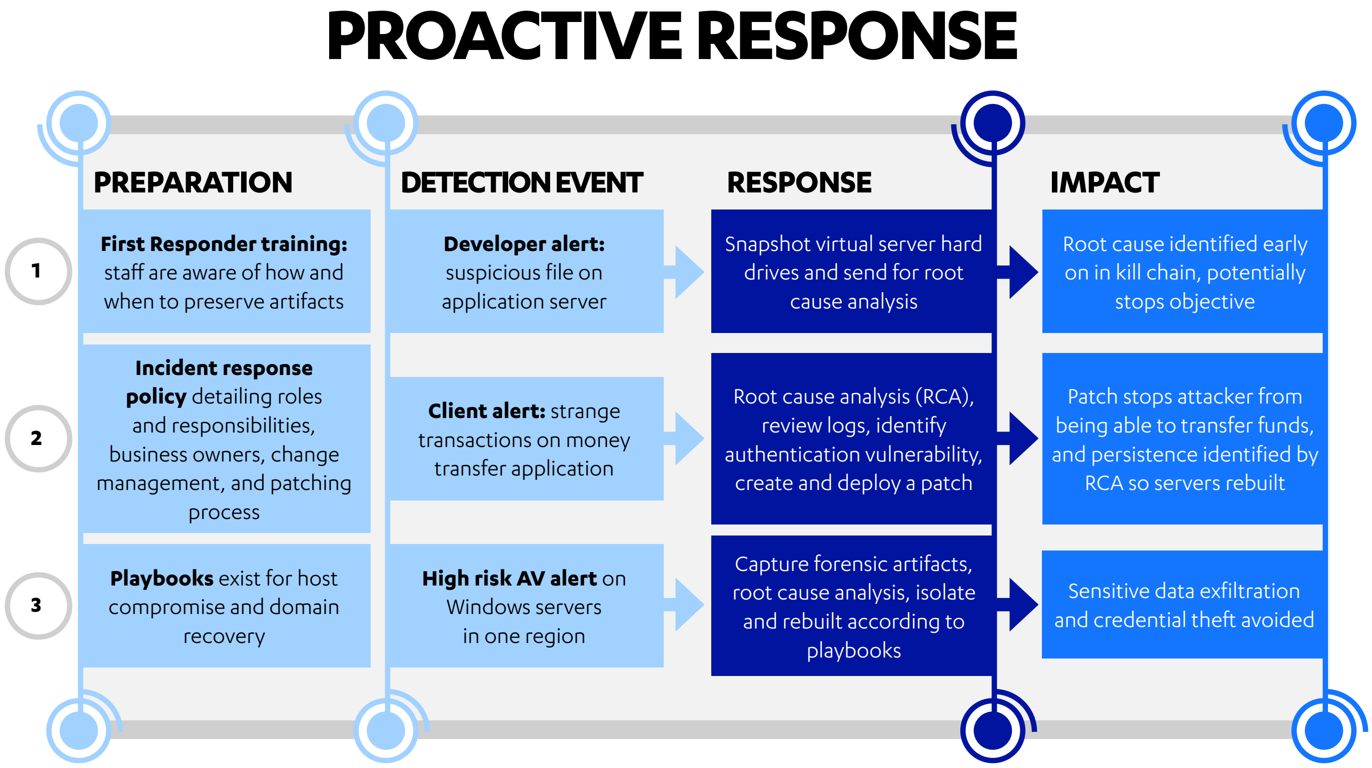

PROACTIVE RESPONSE

	PREPARATION	DETECTION EVENT	RESPONSE	IMPACT
1	**First Responder training:** staff are aware of how and when to preserve artifacts	**Developer alert:** suspicious file on application server	Snapshot virtual server hard drives and send for root cause analysis	Root cause identified early on in kill chain, potentially stops objective
2	**Incident response policy** detailing roles and responsibilities, business owners, change management, and patching process	**Client alert:** strange transactions on money transfer application	Root cause analysis (RCA), review logs, identify authentication vulnerability, create and deploy a patch	Patch stops attacker from being able to transfer funds, and persistence identified by RCA so servers rebuilt
3	**Playbooks** exist for host compromise and domain recovery	**High risk AV alert** on Windows servers in one region	Capture forensic artifacts, root cause analysis, isolate and rebuilt according to playbooks	Sensitive data exfiltration and credential theft avoided

Fig. 2. Points of preparation, detection, response, and impact within a proactive response approach

Enterprises that prioritize readiness are better equipped to raise the barriers between advanced persistent threats (APTs) and their goals, prior to a compromise. Such threats are systematic and require sufficient resource to succeed. By preserving key data and developing learnings early in the kill-chain, the resource required by the attacker diminishes the value of its persistence. Instead of relying on prevention, and responding reactively when it fails, this approach proactively deters threats by making an organization an uneconomical target.

SUMMARY AND CONCLUSION

An organization's ability to maintain business continuity through a cyber incident, and recover any losses in the aftermath, is dependent on its governance of the IR lifecycle. This includes the preparation phase where controlled and strategic actions can be taken, based on learnings from tabletop exercises, capability-specific assessments, and real attacks.

Examples like the one in this paper show that preparation reaches beyond the deployment, configuration, and testing of tooling. Whilst suitable tooling is essential, it cannot provide incident readiness on its own. Instead, its inclusion alongside strategic management of people and processes—via continually-improved IR plans, playbooks, and low-level procedures—make for an effective response Operational resilience during a real incident is the result of skilled risk management and the technical ability of the CSIRT, in addition to the performance of tooling.

For the internal CSIRT managing the frontline during a compromise, proactive response can help circumvent burnout and improve channels of communication with the executive team. Such cross-departmental collaboration is a step towards security becoming a business priority.

Incidents are indeed a business consideration not a pure technical one. The scrutiny of customers, partners, regulators, and the media that results from an incident forms the case to develop security teams equipped to defend against modern threat actors—both constantly evolving and persistent in nature. The value of readiness investment can be represented to senior stakeholders in this context as a route to increased credibility, continuity, and growth.

REFERENCES

1. **Cyber security breaches survey 2020**
   https://www.gov.uk/government/publications/cyber-security-breaches-survey-2020/cyber-security-breaches-survey-2020

2. **The F-Secure guide to rainbow teaming: blue team | building resilience through response process development and simulation**
   https://www.f-secure.com/content/dam/f-secure/en/consulting/our-thinking/collaterals/digital/F-Secure-Consulting-Blue-Team-paper-2020.pdf

3. **Mimikatz**
   https://github.com/gentilkiwi/mimikatz

4. **Enterprise-B is able to restore backups of the affected servers from one month prior to day 201 of the attack**
   The IR provider confirms these are clean, as they predate the servers' compromise. The executive team know that the backup restoration may cause some business impact and loss of production data. Stakeholders are on standby to deal with any fallout.

F-SECURE **INCIDENT RESPONSE TEAM**

US
+1 (917) 341-2116

UK
+44 (0) 333 311 0014

Finland
+358 9 4245 0223

Denmark
+45 89 88 21 10

South Africa
+27 (10) 500-1921

Singapore
+65 3159 1795

F-Secure.

www.f-secure.com/consulting