


Responding to the growing threat of human-operated ransomware attacks

PwC Cyber Security





Contents

Human-operated ransomware attacks	1
Why are organisations vulnerable?	4
How should organisations respond?	6
How can PwC help?	12


Human-operated ransomware attacks

Human-operated ransomware attacks are now one of the top priority cyber threats faced by most organisations. In this type of attack, cyber criminals gain access to internal corporate networks and deploy ransomware to encrypt data – often to devastating effect – before attempting to extort organisations into paying seven or eight figure ransoms to recover access and restore systems. Attackers also steal and threaten to leak sensitive data, to provide additional leverage when extorting their victims.

These attacks represent a more challenging threat than previous well-known ransomware attacks, such as NotPetya and WannaCry. This results from skilled and adaptable financially-motivated people behind the attacks, who can identify and overcome defences, as well as evolve their tactics to maximise their chances of getting organisations to successfully pay out. This is unlike previous high profile attacks which relied on wormlike functionality to spread ransomware.

Given these attackers have now started stealing and threatening to leak sensitive data, the majority of improvement efforts should be focused on preventing these attacks. Focusing solely on backup and recovery strategies is no longer a viable option, as these do not prevent the attacker from stealing data in the first place, or help with the resulting regulatory implications. Even when backup and recovery strategies are in place, for large organisations an enterprise-wide recovery from backups can take weeks and in some cases be practically unfeasible.

Organisations who have not already taken steps to understand and reduce their vulnerability to these attacks should act now. This is especially important as organisations across a wide-range of sectors have recently been affected, and the frequency of these attacks is highly likely to continue to rise over the coming months. The improvements required to reduce the risk of these attacks are not anything surprising to cyber security teams, likely already forming part of organisations' existing improvement plans. However, the escalation in the threat should cause organisations to re-prioritise ongoing and planned activities, as well as consider what actions they can take to immediately reduce their vulnerability to these attacks.





How do human-operated ransomware attacks work?


Ransomware attackers often work with other criminal groups who use **automated and mass-scale techniques** to gain access to company networks, for example by distributing banking trojans via phishing emails. In many cases investigated by PwC's incident response practice, these relatively simple techniques are effective at compromising organisations as they have yet to solve challenging business problems like how to restrict Microsoft Office macros.

Once initial access has been gained into a target organisation, attackers compromise privileged accounts and further systems, using a combination of **legitimate administration tools and security testing tools** (e.g. Cobalt Strike, PowerShell Empire, and BloodHound). We have seen these tools as sufficient to gain privileged access in internal corporate networks due to widespread IT and Active Directory hygiene issues, and detection capabilities that fail to detect the techniques used in these attacks.

When attackers have gained privileged and widespread access to the target's network, sensitive data is extracted and ransomware deployed as widely as possible, in most cases causing significant and long-term disruption to business operations.

The number of ransomware actors has grown steadily throughout 2020, encouraged by the profits derived from high-profile attacks

Typical path of a human-operated ransomware attack



Key

Automated and mass scale

'Human-operated'



What's driving the growth in human-operated ransomware attacks?

The rapid growth in human-operated ransomware attacks has been driven by several factors, including the:

Growing number of actors attracted by perceived easy revenue.

The number of ransomware actors has grown steadily throughout 2020, encouraged by the profits derived from high-profile attacks. For example, the NetWalker affiliate programme was launched in March 2020 and claimed to have amassed \$4.5 million in its first six weeks. After four months of operation, this revenue had reportedly increased to \$25 million.

Popularity of affiliate programmes, lowering the barrier to entry for newcomers.

Twelve well-known schemes are run as affiliate programmes where ransomware developers lease access to their malware in exchange for a share of profits. This trend significantly lowers the barrier to entry, as malware development is no longer a requirement; allowing them to specialise in spreading through organisations' IT environments and deploying ransomware at scale.

Emergence of leak sites, placing additional pressure on victims to pay ransom demands.


The growing use of leak sites, where victim data is released if a ransom is not paid, further increases the pressure on organisations to meet demands, and in turn the profitability of these attacks. After the Maze Group hosted a leak site in January 2020, a further 18 actors have followed suit. By the end of September, over 750 organisations have had data exposed, with 80% of leaks occurring since late April 2020.

Arrival of new attackers, some of which are highly aggressive.

Despite their recent arrival on the ransomware scene, operations like Suncrypt have released data more frequently and in larger numbers than some of their more established rivals. Conti, which we assess to be the replacement for the notorious Ryuk ransomware, launched its leak site in August and released data on over 100 victims in its first eight weeks of activity.

Running total of ransomware data leaks

November 2019 – September 2020



Why are organisations vulnerable?

Our experience helping organisations respond to human-operated ransomware attacks has shown that attackers exploit several commonly occurring IT and security weaknesses. Outlined below are the themes we observe most responding to these incidents and helping clients to deliver targeted security improvements.

Legacy IT creates security weaknesses attackers can exploit.

Most organisations have legacy IT of some form; many still rely on this heavily, however this has significant security implications. The risk of out-of-support operating systems increases over time as vulnerabilities are identified and remain unpatched. In addition, legacy operating systems are nearly always incompatible with modern security tooling and lack the security features required to defend against these types of attacks. In many instances legacy systems host some of an organisation's most critical applications; those typically targeted in human-operated ransomware attacks.

Technology is not securely configured to prevent common cyber attack techniques.

In most cases, initial access in human-operated ransomware attacks stems from the compromise of workstations with phishing emails or servers by exploiting unpatched vulnerabilities in internet-facing services. Attackers exploit the poor configuration of these systems, with security controls either absent or not effectively configured. These issues often remain unfixed due to a lack of awareness around security good practice, a lack of prioritisation by IT teams when delivering security fixes, and the unknown or perceived business impact of delivering the fixes.

Poor protection of privileged accounts allows attackers to compromise credentials.

We have seen skilled ransomware operators obtain Domain Administrator privileges within 72 hours, as organisations have not adequately protected privileged accounts. The most common problems are large numbers of overly privileged accounts, risky operating practices by domain administrators, and the use of insecure passwords and authentication mechanisms. The root cause of this is often that Active Directory (and other identity and access management systems) is seen as an IT tool, rather than a security tool that is essential to preventing attackers from gaining privileged access; and is therefore not configured with security in mind.

Ineffective detection and response capabilities give attackers freedom to operate.

Attackers often remain in internal corporate networks for days or weeks, escalating privileges and expanding their footholds, before deploying destructive ransomware. During this time, there are usually many detection and response opportunities that organisations fail to identify and take advantage of. This is most often because, traditional signature-based security tooling is ineffective at detecting the techniques used in these attacks, alerts are often lost in the noise as security operations teams are overwhelmed by false positives, and containment processes are not effective. Also, the use of commodity trojans such as Emotet, Dridex or Qakbot as an initial infection vector is often highly effective, as detecting and remediating these 'commodity malware' infections is not prioritised by security teams.

Organisations have retained on-premise infrastructure and struggled to adopt the SaaS cloud.

Cloud "software-as-a-service" business applications (such as email, file storage, and CRM) can significantly reduce the impact of a human-operated ransomware attack, yet many organisations continue to rely on and invest in on-premise infrastructure and applications. This type of environment was rarely architected with security in mind (or indeed to prevent modern security threats), meaning attackers can easily use now widely available offensive security tooling to exploit the resulting weaknesses. There are no quick-fixes, as effectively retrofitting modern cyber security controls on IT infrastructure can be costly and complex, requiring IT to be modernised before it can become securable.



What questions should you ask to assess your vulnerability?

Management teams should question their security teams to determine how susceptible they are to ransomware attacks, including:

How are we using technical security controls to limit the risk of a workstation being compromised by phishing attacks?

How are we using secure administration practices and restricting the use of domain administrator accounts to limit the risk of credential-theft attacks?

How could we still access our email and messaging if a ransomware attack impacted our on-premise IT infrastructure?

How quickly and comprehensively are we detecting and remediating 'commodity malware' infections on workstations?

How confident are we that we can detect the compromise and abuse of privileged accounts by an attacker?

How do we restrict the connectivity and Active Directory trust relationships between different areas of the business to slow down and limit the spread of ransomware attacks?

How confident are we that we can detect common attacker tools being used within our network, for example Cobalt Strike or BloodHound?

How are we mitigating the risk of legacy IT to ensure this does not provide a point of weakness that allows an attacker to compromise our entire environment?

What have we done to validate that any network security controls are still effective with an increasingly remote workforce?

How should organisations respond?

Given the growing cyber threat presented by human-operated ransomware attacks, we recommend cyber security teams take a three step approach to both immediately and sustainably reduce their organisation's vulnerability to these attacks:

1 Understand and report on their organisation's vulnerability to the threat

Use security testing to assess whether the techniques used in human-operated ransomware attacks can be prevented and detected by defences in place, and to identify the vulnerabilities and weaknesses that could be exploited by ransomware attackers. Security teams should provide ongoing reporting to management on risk, using the results of this threat-focused testing approach.

2 Deliver targeted improvements to immediately reduce risk, and validate their implementation

Targeted improvements should be delivered to prevent and detect the techniques used in human-operated ransomware attacks, and address the identified vulnerabilities and weaknesses. We have seen customers successfully achieve this by bringing together IT and security teams to collaborate on developing actionable fixes appropriate for the environment. Security testing should then be used to validate that improvements have been correctly implemented to address risks identified.

3 Build capabilities to deliver sustainable cyber risk reduction

Where the root cause of security vulnerabilities and weaknesses has not been remediated in the previous steps, strategic initiatives should be designed to deliver sustainable cyber risk reduction. For example moving away from legacy IT or reworking how privileged access is managed within the organisation.

Using this approach we recommend six priority areas of focus to reduce the risk of human-operated ransomware attacks.

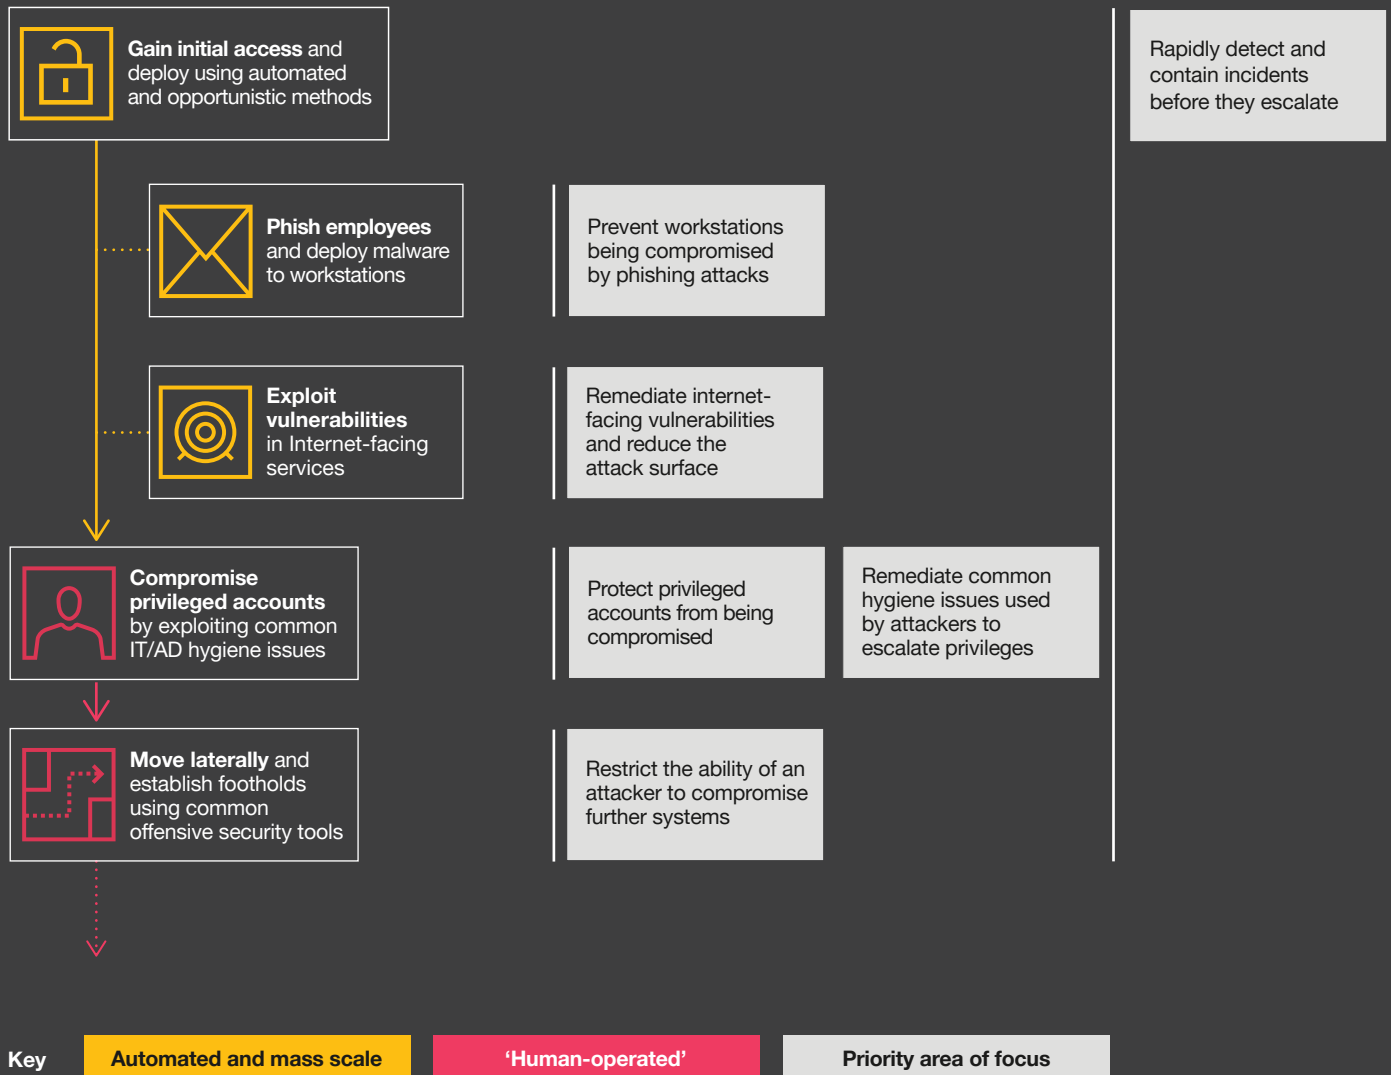
- 1. Prevent workstations being compromised by phishing attacks**
- 2. Remediate internet-facing vulnerabilities and reduce the attack surface**
- 3. Protect privileged accounts from being compromised**
- 4. Remediate common hygiene issues used by attackers to escalate privileges**
- 5. Restrict the ability of an attacker to compromise further systems**
- 6. Rapidly detect and contain incidents before they escalate**

For each, we have listed several targeted improvements that, with the right delivery approaches and prioritisation, can be delivered within three months. While many of these improvements may seem obvious at first glance, their value (or the challenge in their delivery) should not be underestimated – it is the “hard basics” of security that still represent the greatest challenge, and corresponding benefit.

As the “devil is in the detail” when reducing the risk posed by ransomware, we recommend that organisations that have already implemented these improvements still take steps to validate that they are effectively preventing and detecting the techniques used in these attacks. In cases investigated by PwC's incident response practice, we regularly see a disconnect between the believed and actual levels of risk reduction brought about by improvements that have been made (for example, the deployment of tooling without having appropriately configured it).

Typical path of a human-operated ransomware attack

Six priority areas of focus to reduce the risk of human-operated ransomware attacks



Prevent workstations being compromised by phishing attacks

As well as raising employee awareness of phishing emails, organisations should implement technical security controls to prevent phishing emails with malicious payloads compromising workstations. These are increasingly important in light of multiple high profile cybercrime actors hijacking legitimate email threads to distribute malicious payloads, combating many common employee awareness training initiatives. Key actions include:

Ensure email filtering tooling is appropriately configured to block phishing emails.

Email filtering tooling should be configured to scan attachments for malicious files and links, block file-types commonly used by attackers (e.g. script files such as HTA and PS1), and detect techniques used by attackers to fake legitimacy (e.g. the spoofing of internal email addresses, or sending from recently registered domains). This tooling should also be integrated with up-to-date threat intelligence (including both atomic and behavioural indicators) to block suspected likely phishing emails.

Deploy and configure web filtering tooling to prevent users from downloading malicious files.

Web filtering tooling should be configured to block the download of file-types commonly used to deliver malware, and scan all others for malicious content. Sites should be assigned a risk rating and blocked accordingly; high risk sites might include those categorised as malicious, those using newly registered domains, or commonly abused top-level domains.

Restrict Microsoft Office macros to prevent attackers using these to deliver malicious payloads.

Microsoft Office macros should be disabled where possible for employees, teams and departments without a suitable business case. Where this is not possible, their use should be restricted by: blocking execution in documents downloaded from the internet, only allowing execution for documents in trusted locations, or only allowing signed macros to be run. For organisations with Microsoft 365 E5 licenses, Application Guard for Office should be enabled where possible.

Restrict execution of scripts on workstations to prevent attackers using these to bypass defences.

The use of scripts on workstations should be restricted, for example by using PowerShell constrained language mode to limit the full use of PowerShell to authorised users. Common scripting file extensions should be set to open in text editors by default, rather than being executed, to reduce the risk of users executing malicious files delivered via phishing emails.

Remediate internet-facing vulnerabilities and reduce the attack surface

Many of the groups recently involved in ransomware attacks exploit vulnerabilities in internet-facing services to gain initial access. Yet in many cases we see these vulnerabilities remaining unfixed, as there are issues in the coverage and configuration of vulnerability scanning tools, and remediation efforts are not effectively prioritised, tracked or escalated. Key actions include:

Perform vulnerability scanning and monitoring to ensure vulnerabilities are rapidly remediated.

Vulnerability scanning tooling should be configured to perform regular scans, and alert for any new internet-facing vulnerabilities or exposed services. Security teams should also ensure all internet-facing IP address ranges are covered by scanning tools and review any existing exceptions to ensure they have not been granted for vulnerabilities that could be exploited by attackers.

Disable or restrict access to internet-facing services to reduce the attack surface.

Externally accessible systems should be regularly audited and organisations should look to disable or restrict access to any internet-facing services (e.g. RDP, SSH, SMB) that are not strictly required; this mitigates the risk of future vulnerabilities by reducing the overall attack surface. Where services are required then compensating controls should be implemented to mitigate risk (for example, IP address whitelisting and strong authentication).

Enforce multi-factor authentication to reduce the impact of compromised credentials.

Multi-factor risk-based authentication should be configured on all remote access systems and internet-facing services. Authentication logs should also be collected from these services, and monitored for use cases that could represent compromised accounts, for example sign-in events indicating 'impossible travel' by the user, or logins from unexpected countries and devices.

Protect privileged accounts from being compromised

Once attackers have compromised workstations with phishing attacks, one of their first goals is to escalate privileges. The primary way organisations can make it harder for attackers to do this is by protecting the credentials of privileged accounts from being exposed on those systems most at risk of being compromised. Key actions include:

Restrict the use of domain administrator accounts to prevent their credentials being exposed.

Security teams should work with IT teams to develop secure administration practices that reduce the risk of credential-theft attacks. For example, restricting domain administrator accounts from logging into workstations and servers, and only using accounts with domain administrator privileges where strictly required. Domain administrator accounts should be controlled by a privileged access management tool that securely manages credentials and isolates administrator sessions.

Identify and remediate any attack paths to privileged accounts in Active Directory.

BloodHound (an offensive security tool) should be used to identify, investigate and eliminate attack paths to privileged accounts. These are often present due to hidden and unintended trust relationships within Active Directory environments, for example misconfigured or complex inheritance-based permissions. This should also be an opportunity for security teams to gain confidence BloodHound's use is reliably detected.

Restrict accounts in local administrator groups to reduce the identity attack surface of endpoints.

Users and groups in the local administrator group on workstations and servers should be reviewed to ensure accounts are only added where strictly necessary. This reduces the number of accounts that, if compromised, could allow an attacker to gain widespread administrative access to systems. Local administrator groups should also be monitored to ensure any modifications are detected.

Monitor Active Directory to detect the insecure use, compromise and abuse of privileged accounts.

Security teams should deploy tooling, such as Microsoft Defender for Identity, to monitor and detect anomalies involving privileged accounts, e.g. logging in from a new location. Accounts suspected of compromise should be blocked and investigated thoroughly before reinstating access.

Remediate common hygiene issues used by attackers to escalate privileges

The most common way we see attackers escalate privileges within enterprise IT environments is by exploiting vulnerabilities and weaknesses resulting from IT and AD hygiene issues. These issues are prevalent in sprawling and poorly understood internal corporate networks that have evolved over time without adequate security governance and investment. Key actions include:

Enforce strong passwords on service accounts to make these more difficult to crack.

Strong passwords should be set on all service accounts, prioritising any associated with Service Principal Names (SPNs). A precursor to this should be auditing and maintaining an inventory of service accounts to identify account owners, confirming they are still required, and ensuring they are appropriately labelled and that a strong password policy is enforced. Longer-term service accounts should be onboarded onto a privileged account management tool.

Remove credentials stored in network shares to prevent attackers using these to compromise accounts.

In many cases investigated by PwC's incident response practice, organisations unknowingly have plaintext privileged credentials stored in network file shares (or other easily accessible locations) which are accessed and exploited by attackers. Offensive security tools, e.g. PowerShell scripts should be used to scan file shares for credentials, so these can be removed and the exposed passwords reset. Access to network shares should also be restricted as much as possible, as these provide a common way for attackers to gain access to sensitive data.

Use security testing and Microsoft Secure Score to identify IT and Active Directory hygiene issues.

Security testing should be used to identify weaknesses and vulnerabilities an attacker could exploit by simulating cyber attack techniques. For organisations using Microsoft 365, Secure Score should also be used to identify fixes to improve security posture and reduce attack surface, for example by enabling Credential Guard to improve workstation secure configuration.

Restrict the ability of an attacker to compromise further systems

Attackers attempt to compromise further workstations and servers to escalate privileges and gain the access they need to deploy ransomware widely across the environment. Making it more difficult for an attacker to be able to move laterally increases the chance attackers will be detected before deploying ransomware. Key actions include:

Remediate exploitable vulnerabilities on internal systems to remove trivial routes to compromise further systems.

Vulnerability scanning should be used to identify vulnerabilities on the internal network. The scope and coverage of vulnerability scanning tools should be reviewed and any exceptions challenged. Where system stability is a concern, bespoke or manual penetration testing should be carried out to check for vulnerabilities (rather than simply not performing testing). Systems should be segmented from the network where they cannot be patched.

Prevent and detect common techniques for mass-deployment of ransomware.

Host-based firewalls should be configured on workstations to block inbound connections by default, in order to prevent ransomware deployment at scale to workstations using SMB-based lateral movement. Use of legitimate software deployment mechanisms (e.g. SCCM) and remote administration tools (e.g. PsExec, WMI and GPO) should be monitored in order to detect unauthorised use. Domain controllers should be monitored for the execution of script files, commonly used by attackers to deploy malware and disable security tooling.

Segment business units and high-risk networks to limit the blast radius of ransomware attacks.

Business units and high-risk networks should be segmented to limit the number of systems impacted by a ransomware attack. This should be done by blocking network connectivity, especially protocols commonly used for lateral movement, and breaking or hardening Active Directory trust relationships.

Rapidly detect and contain incidents before they escalate

As the deployment of ransomware is the final stage of an attack that may have lasted months, there are almost always opportunities to detect and contain these attacks before data is encrypted or stolen. By effectively detecting and containing “commodity malware” infections, organisations can also prevent opportunities for the ransomware attackers to gain access in the first place. Key actions include:

Deploy a capable endpoint security agent to detect and prevent attacker activity.

An endpoint agent should be chosen that detects and prevents suspicious activity on workstations and servers using behavioural analytics, as well as providing support for the Anti Malware Scan Interface (AMSI) to detect the malicious use of scripting languages, and which empowers security teams with rapid investigative and response capabilities.

Onboard a managed detection and response (MDR) service to automate the response to common threats and ensure that “commodity malware” is detected and remediated.

Many organisations do not have the capability to effectively monitor their estate and respond to the volume of alerts created; lacking either in expertise, tooling or staff. A quick win is to onboard a managed detection and response service, as this provides the necessary tooling with the required people, process and automation/orchestration wrappers, and comprehensive, demonstrable detection coverage of attacker techniques.

Ensure common attacker tools are detected and alerts are effectively remediated.

Security testing should be used to ensure an organisation can effectively detect common attacker tools, and that the necessary people and processes are in place to investigate and respond to alerts. Capability gaps should be identified and remediated through the deployment and configuration of detection and response tooling, development of incident response processes and training of suitable individuals.




How should organisations prepare to respond to a ransomware attack?

While the majority of efforts should be focused on preventing these attacks, it is also vital that organisations plan and exercise their response to a major ransomware incident. In many cases where PwC's incident response team have been brought in to assist, days have already been lost due to a lack of clear response and recovery plans, and leadership's failure to understand the scale of the challenge.

We have seen that the 'realities of recovery' in organisations are far more challenging than they anticipate, as IT environments are complex and information about critical systems is unclear. Even when organisations have gained access to the decryption keys, in many cases they are still unable to decrypt large amounts of data as it has been left corrupted by ransomware tools, or they lack the technical ability to do so.

We recommend organisations carry out the following actions to prepare for a major ransomware attack and mobilise an effective response:

- 1. Develop and exercise incident response and crisis plans** that clearly outline how a response should be managed and coordinated. These should be challenged to ensure they are effective in a catastrophic ransomware scenario, where common security and IT tools may be unavailable and recovery efforts may have to be sustained for weeks or months.
- 2. Understand where critical data is**, what the implications would be on that data if systems were unavailable, the regulatory requirements attached to this, and what would need to be recovered in order to create a 'minimum viable company'.
- 3. Ensure that offline backups have been created and validated** for all critical systems and data, including Active Directory, with a well-defined and tested restore procedure.
- 4. Build or retain the technical expertise to investigate and respond** to the attack (e.g. investigating the extent of compromise, identifying attacker access and eradicating the attacker from the environment).



For our latest insights and resources, please visit

www.pwc.co.uk/cybersecurity

@2020 PwC. All rights reserved. PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details. This content is for general information purposes only, and should not be used as a substitute for consultation with professional advisors.

Mitie Design RITM3624127 (10/20).